

Verifying Concurrent Data Structures

Roland Meyer

Joint work with Thomas Wies and Sebastian Wolff

Motivation

Programs = Algorithms + Data Structures

–Niklaus Wirth, 1978



Motivation

Programs = Algorithms + Data Structures

–Niklaus Wirth, 1978



Verification



Motivation

Programs = Algorithms + Data Structures

–Niklaus Wirth, 1978

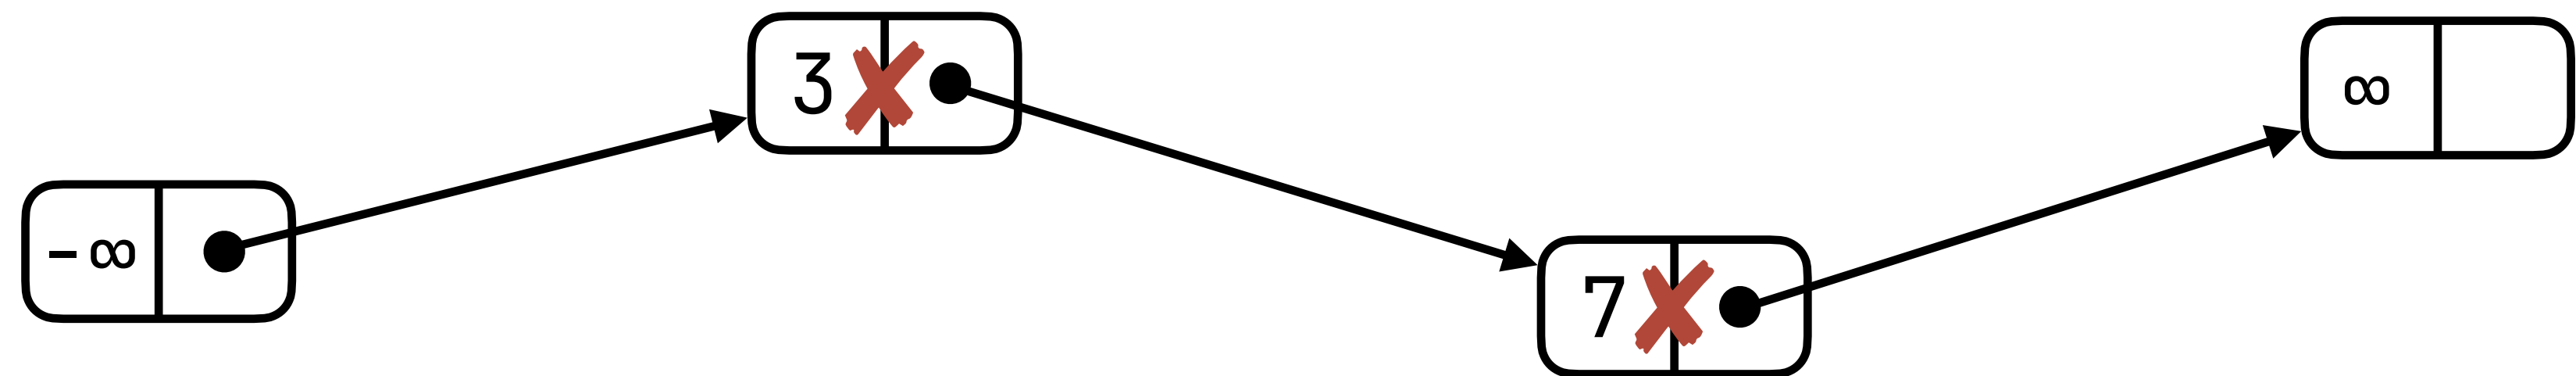
↓
Verification

↓
Automation



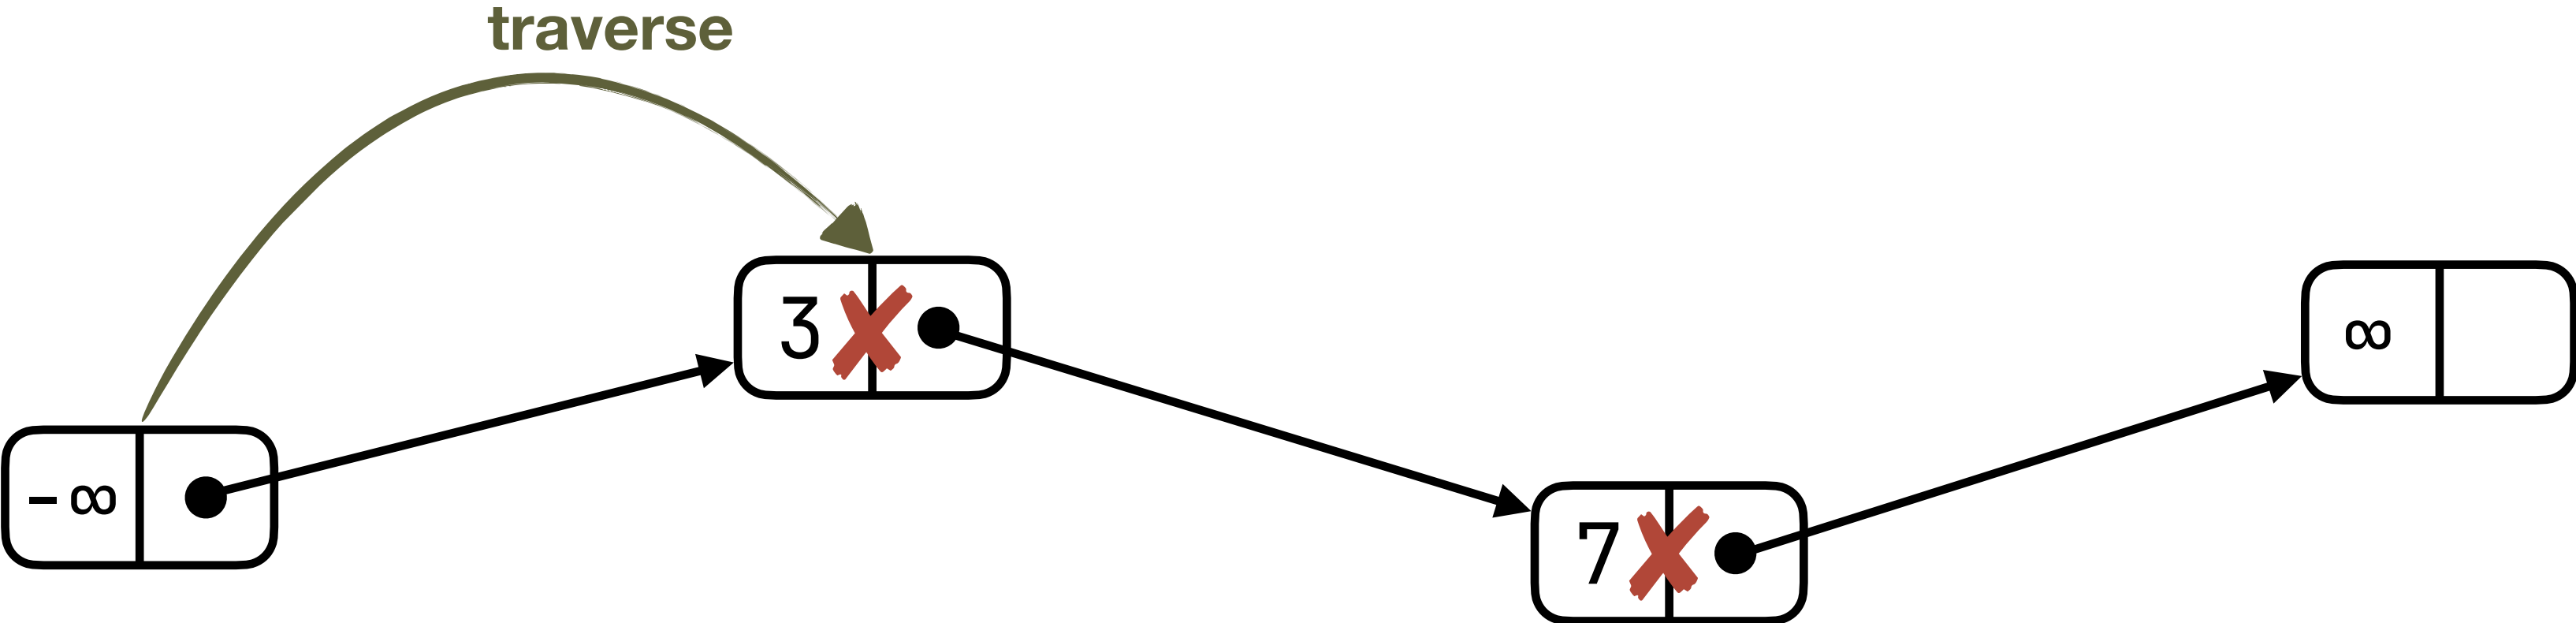
Example: contains(5)

set({}):



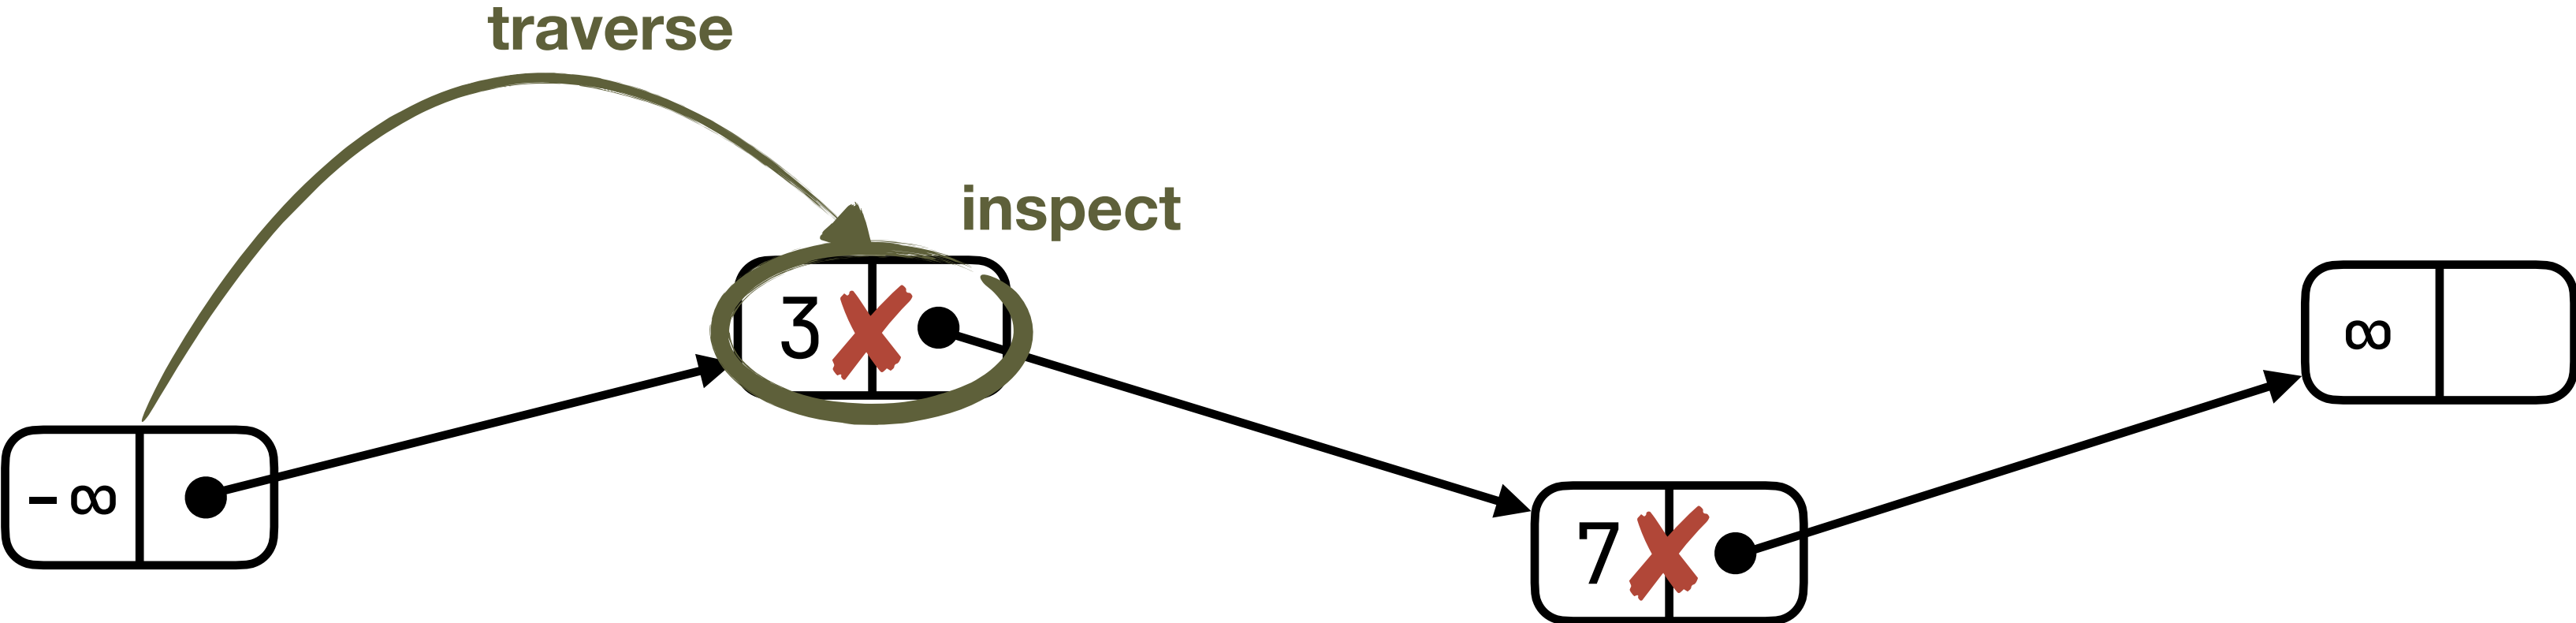
Example: contains(5)

set({}):



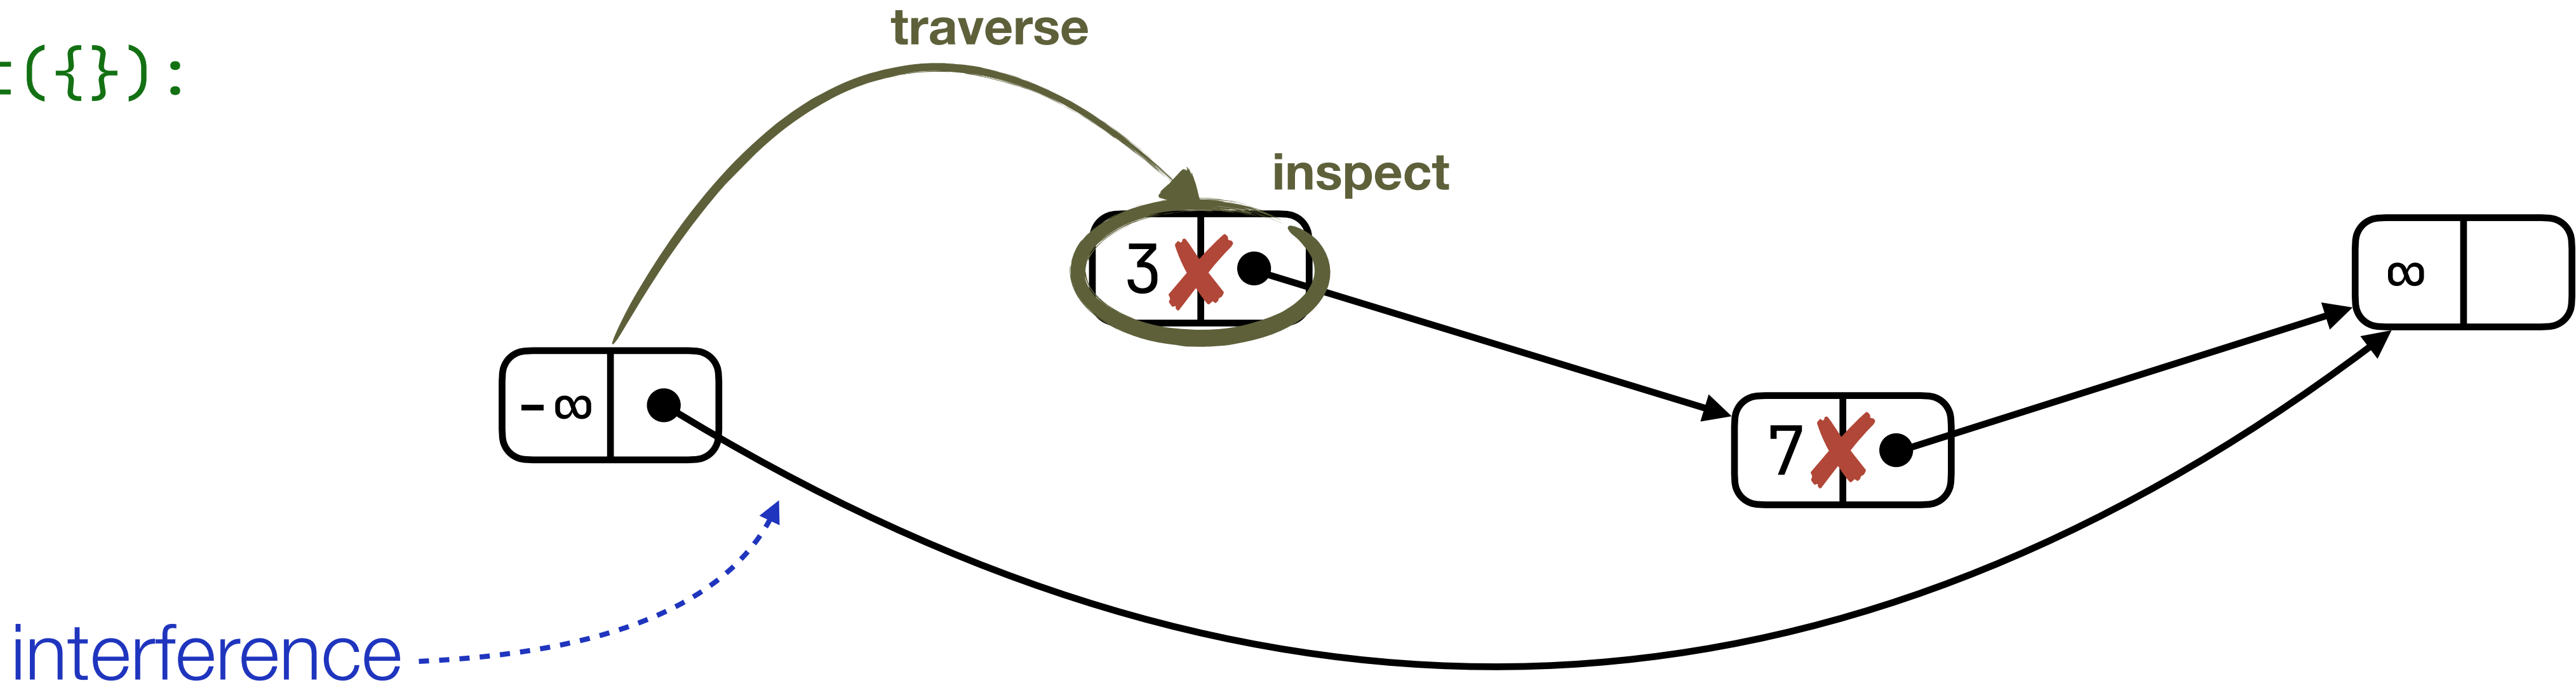
Example: contains(5)

set({}):



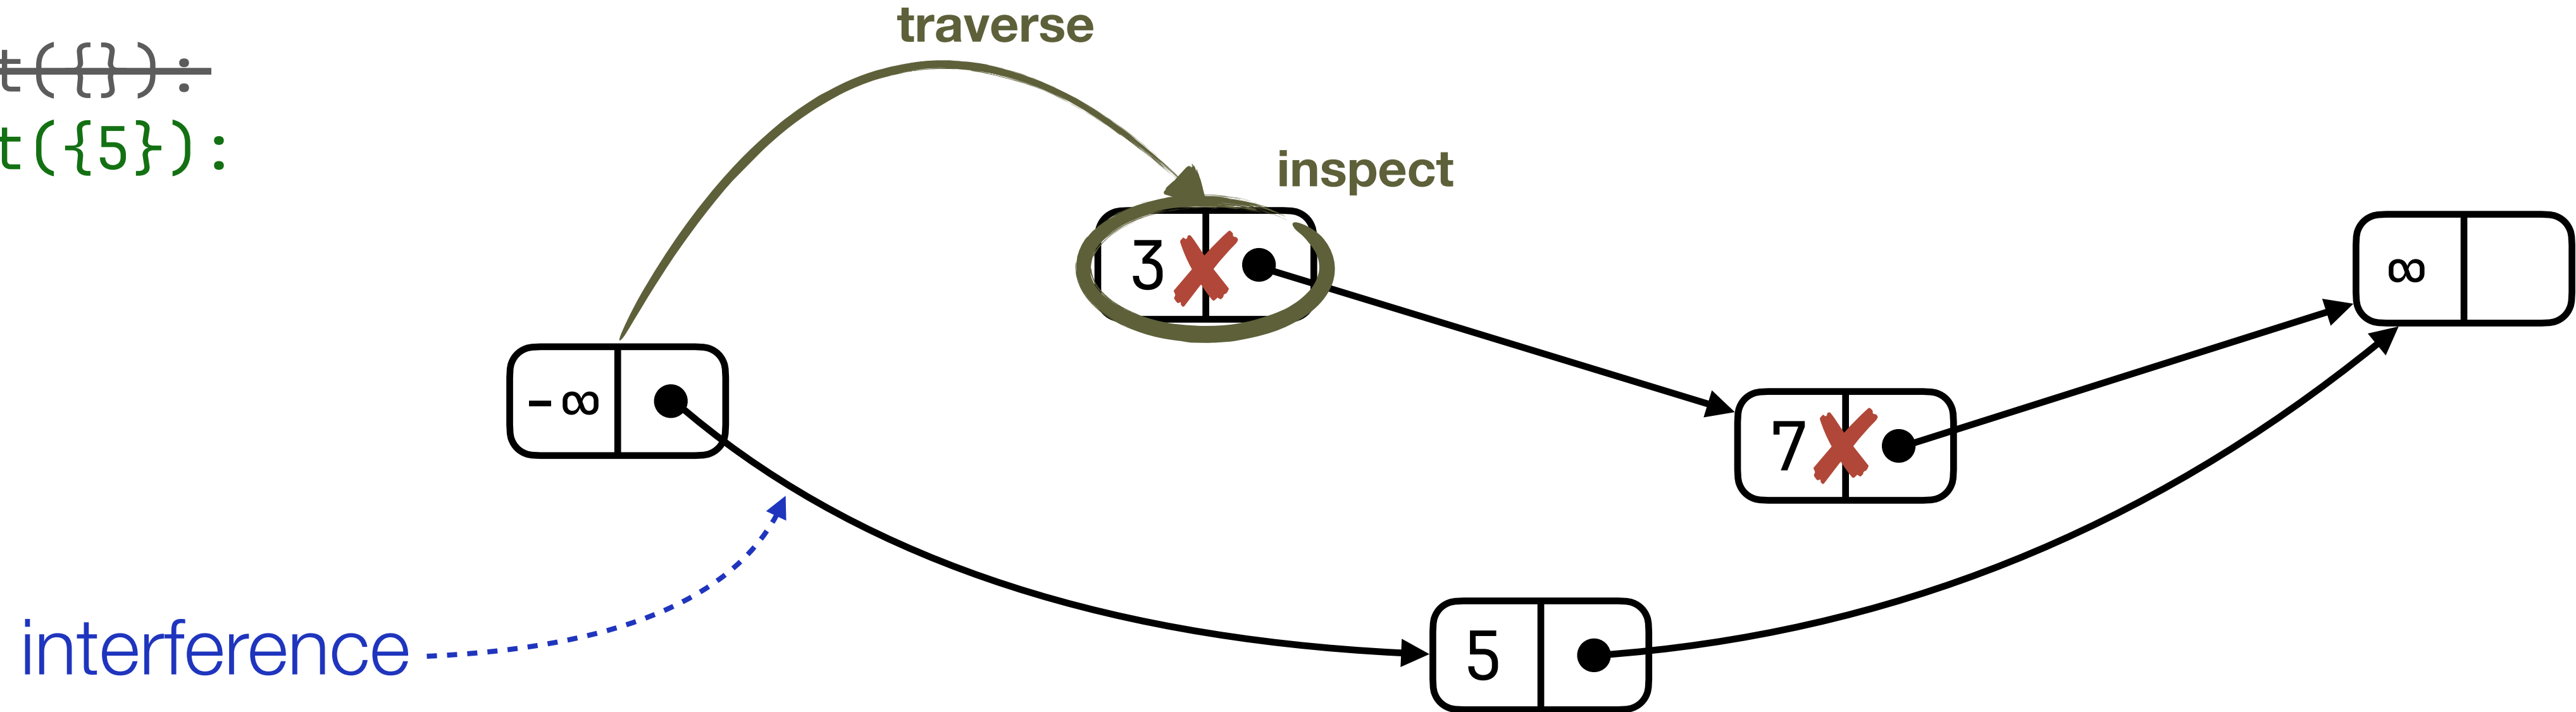
Example: contains(5)

set({}):



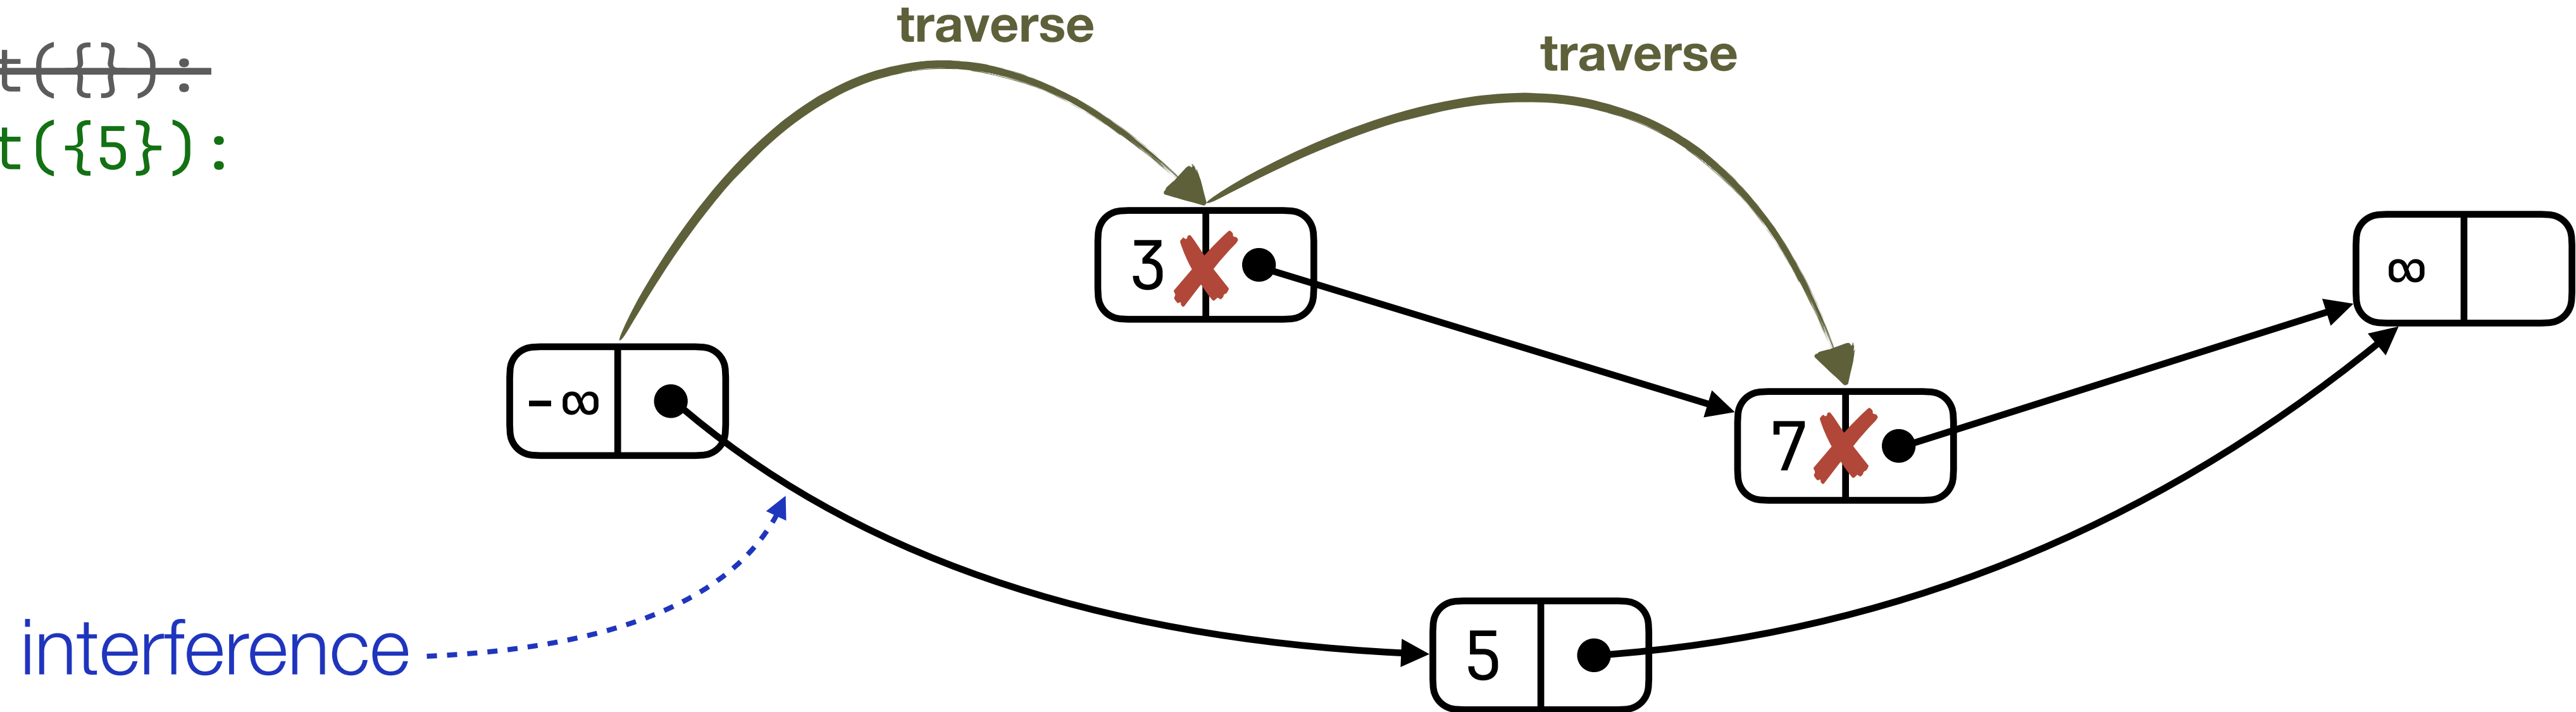
Example: contains(5)

~~set({}):~~
set({5}):



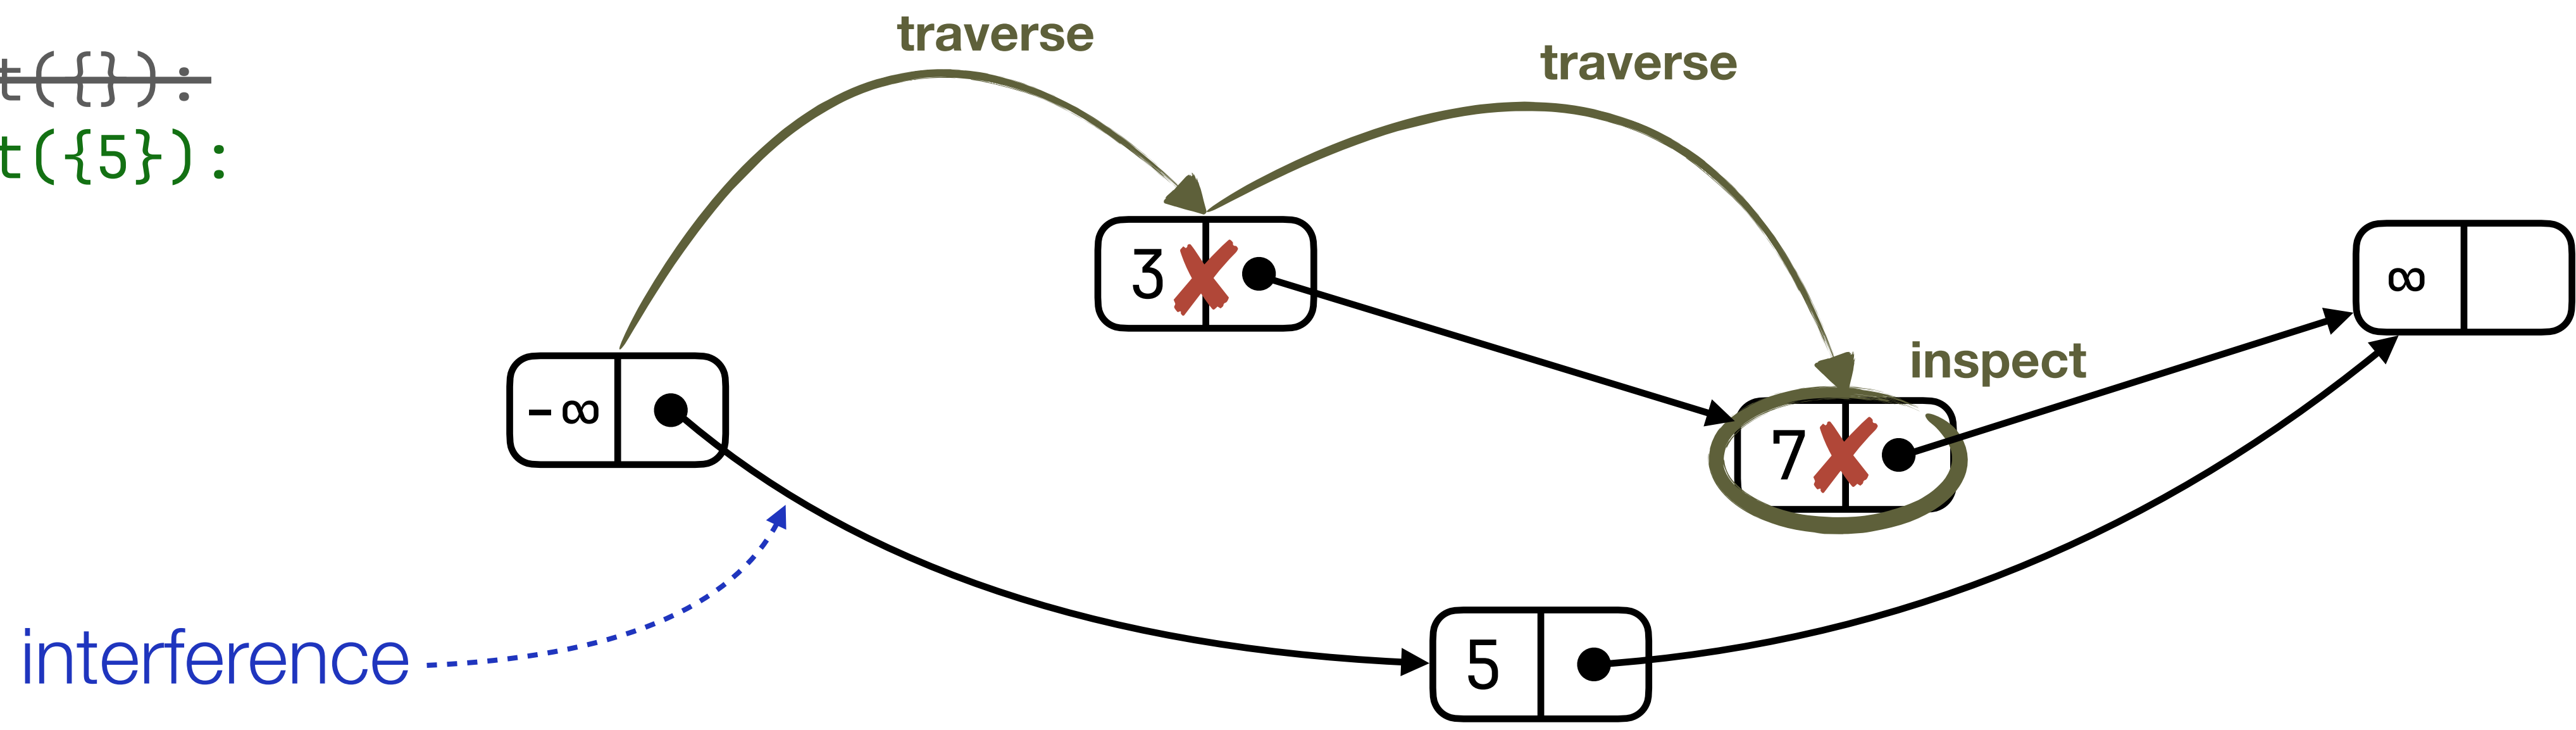
Example: contains(5)

~~set({}):~~
set({5}):

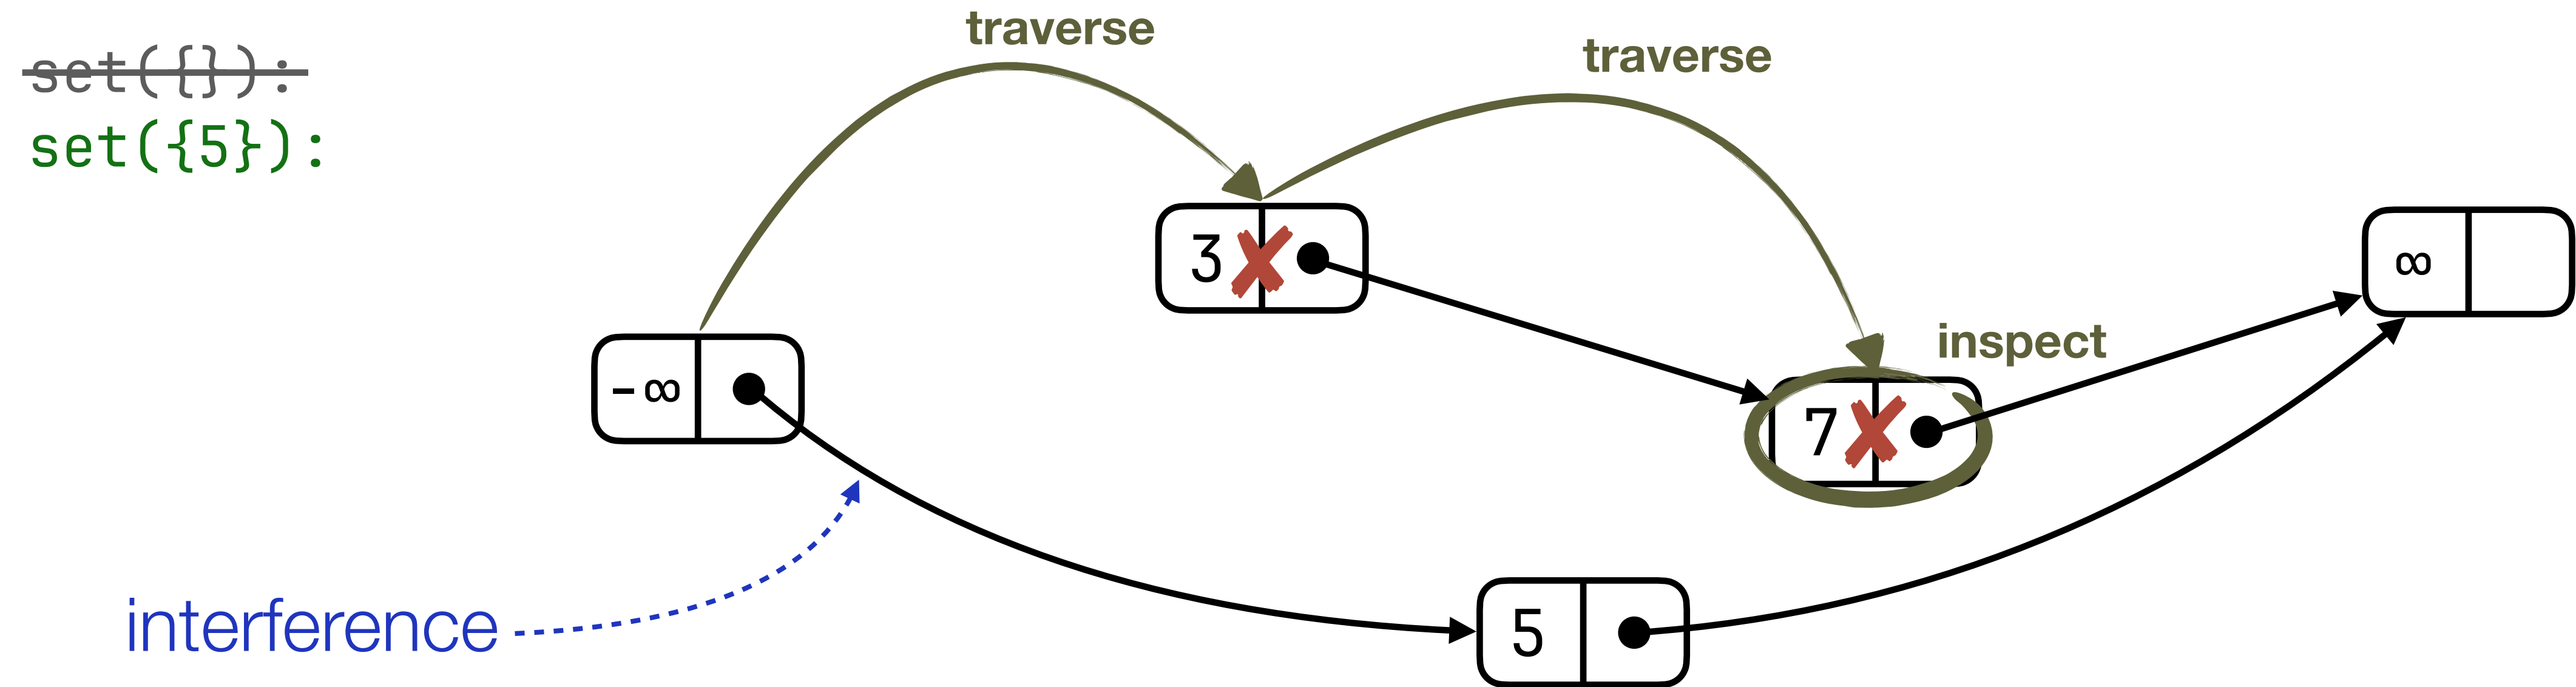


Example: contains(5)

~~set({}):~~
set({5}):



Example: contains(5)



- Result: `contains(5) = false`

➔ linearizable (i.e. correct); proof obligation:

`traversal` \models `set(M), 5 ∉ M`
at some point

Approach: Program Logic

**manual/mechanized/automated
proof construction**

- Owicki-Gries
 - ➔ Thread modularity
 - ➔ Implicit interference handling
- Separation logic
 - ➔ Parametric in the domain
 - ➔ Local reasoning
 - ➔ Arsenal of proof techniques

compact proofs

applicability

Contributions



Flow Framework

Reasoning about heap graphs without recursive predicates.

[POPL'18, ESOP'20, TACAS'23]



Hindsight

Non-fixed linearization points without prophecies.

[OOPSLA'22, PLDI'23]



Decomposing Updates

Unbounded footprints without induction.

[OOPSLA'22, *under submission*]

Contributions



Flow Framework

Reasoning about heap graphs without recursive predicates.

[POPL'18, ESOP'20, TACAS'23]



Hindsight

Non-fixed linearization points without prophecies.

[OOPSLA'22, PLDI'23]



Decomposing Updates

Unbounded footprints without induction.

[OOPSLA'22, *under submission*]

Goal

Reason about heap graphs of concurrent data structures.

Problem

No inductive definition, no recursive predicates.

Solution

Fixed points, data-flow analysis on heap graphs.

Flow Framework

Flow Framework

flow graphs = heap graphs + data flow values (ghost)

Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from commutative monoid $(\mathbb{M}, +, 0)$

$$m \in \mathbb{M} \quad \boxed{a}^x$$

Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from commutative monoid $(\mathbb{M}, +, 0)$



- Flow propagation via continuous **edge functions**



Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from commutative monoid $(\mathbb{M}, +, 0)$
- Flow propagation via continuous **edge functions**
- **The flow:**
fixed point, wrt. initial value



Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from commutative monoid $(M, +, 0)$
- Flow propagation via continuous edge functions
- **The flow:** some fixed point, wrt. initial value

$m \in M$

$m \in M$

$m \in M$

[POPL'18]

Go with the Flow: Compositional Abstractions for Concurrent Data Structures*

SIDDHARTH KRISHNA, New York University, USA
DENNIS SHASHA, New York University, USA
THOMAS WIES, New York University, USA

Concurrent s
structures. H
the sequent
overlays. To
the data stru
quantity asso
quantity as a
root form a t
which expres
flow invariari
semantic mod
describing ce
wide variety
logic RGSep
nontrivial co
dictionary al
proofs canno

CCS Concep
algorithms;

Additional Ke

ACM Refere
Siddharth Kr
for Concurr
<https://doi.org>

1 INTRO

With the ad
into solving
da Rocha Pi

*This work is f
1339362, MCB

Authors' addr
University, US

Permission to
provided that
the full citation
Abstracting wi
prior specific p
© 2018 Associa
2475-1421/2018
<https://doi.org>

Proceed

[ESOP'20]

Local Reasoning for Global Graph Properties

Siddharth Krishna¹, Alexander J. Summers², and Thomas Wies¹

¹ New York University, New York, NY, USA, {siddharth,wies}@cs.nyu.edu
² ETH Zürich, Zurich, Switzerland, alexander.summers@inf.ethz.ch

Abstract. Separation logics are widely used for verifying programs that manipulate complex heap-based data structures. These logics build on so-called *separation algebras*, which allow expressing properties of heap regions such that modifications to a region do not invalidate properties stated about the remainder of the heap. This concept is key to enabling modular reasoning and also extends to concurrency. While heaps are naturally related to mathematical graphs, many ubiquitous graph properties are non-local in character, such as reachability between nodes, path lengths, acyclicity and other structural invariants, as well as data invariants which combine with these notions. Reasoning modularly about such graph properties remains notoriously difficult, since a local modification can have side-effects on a global property that cannot be easily confined to a small region.

In this paper, we address the question: What separation algebra can be used to avoid proof arguments reverting back to tedious global reasoning in such cases? To this end, we consider a general class of global graph properties expressed as fixpoints of algebraic equations over graphs. We present mathematical foundations for reasoning about this class of properties, imposing minimal requirements on the underlying theory that allow us to define a suitable separation algebra. Building on this theory, we develop a general proof technique for modular reasoning about global graph properties expressed over program heaps, in a way which can be directly integrated with existing separation logics. To demonstrate our approach, we present local proofs for two challenging examples: a priority inheritance protocol and the non-blocking concurrent Harris list.

1 Introduction

Separation logic (SL) [31,37] provides the basis of many successful verification tools that can verify programs manipulating complex data structures [1, 4, 17, 29]. This success is due to the logic's support for reasoning modularly about modifications to heap-based data. For simple inductive data structures such as lists and trees, much of this reasoning can be automated [2, 11, 20, 33]. However, these techniques often fail when data structures are less regular (e.g. multiple overlaid data structures) or provide multiple traversal patterns (e.g. threaded trees). Such idioms are prevalent in real-world implementations

→ • • •
∈ M

Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from commutative monoid $(M, +, 0)$
- Flow propagation via continuous edge functions
- **The flow:** some fixed point, wrt. initial value

$m \in M$

$m \in M$

$m \in M$

[POPL'18]

Go with the Flow: Compositional Abstractions for Concurrent Data Structures*

SIDDHARTH KRISHNA, New York University, USA
DENNIS SHASHA, New York University, USA
THOMAS WIES, New York University, USA

Concurrent s
structures. H
the sequent
overlays. To
the data stru
quantity asso
quantity as a
root form a t
which expres
flow invarian
semantic mo
describing ce
wide variety
logic RGSep
nontrivial co
dictionary al
proofs canno

CCS Concep
algorithms;

Additional Ke

ACM Refere
Siddharth Kr
for Concurr
<https://doi.org/10.1145/3176951>

1 INTRO

With the ad
into solving
da Rocha Pi

*This work is f
1339362, MCB-

Authors' addr
University, US

Permission to
provided that
the full citation
Abstracting wi
prior specific p
© 2018 Associa
2475-1421/2018
<https://doi.org/10.1145/3176951>

Proceed

[ESOP'20]

Local Reasoning for Global Graph Properties

Siddharth Krishna¹, Alexander J. Summers², and Thomas Wies¹

¹ New York University, New York, NY, USA, {siddharth,wies}@cs.nyu.edu
² ETH Zürich, Zurich, Switzerland, alexander.summers@inf.ethz.ch

Abstract. Separation logics are widely used for verifying programs that manipulate complex heap-based data structures. These logics build on so-called *separation algebras*, which allow expressing properties of heap regions such that modifications to a region do not invalidate properties stated about the remainder of the heap. This concept is key to enabling modular reasoning and also extends to concurrency. While heaps are naturally related to mathematical graphs, many ubiquitous graph properties are non-local in character, such as reachability between nodes, path lengths, acyclicity and other structural invariants, as well as data invariants which combine with these notions. Reasoning modularly about such graph properties remains notoriously difficult, since a local modification can have side-effects on a global property that cannot be easily confined to a small region.

In this paper, we address the question: What separation algebra can be used to avoid proof arguments reverting back to tedious global reasoning in such cases? To this end, we consider a general class of global graph properties expressed as fixpoints of algebraic equations over graphs. We present mathematical foundations for reasoning about this class of properties, imposing minimal requirements on the underlying theory that allow us to define a suitable separation algebra. Building on this theory, we develop a general proof technique for modular reasoning about global graph properties expressed over program heaps, in a way which can be directly integrated with existing separation logics. To demonstrate our approach, we present local proofs for two challenging examples: a priority inheritance protocol and the non-blocking concurrent Harris list.

1 Introduction

Separation logic (SL) [31,37] provides the basis of many successful verification tools that can verify programs manipulating complex data structures [1, 4, 17, 29]. This success is due to the logic's support for reasoning modularly about modifications to heap-based data. For simple inductive data structures such as lists and trees, much of this reasoning can be automated [2, 11, 20, 33]. However, these techniques often fail when data structures are less regular (e.g. multiple overlaid data structures) or provide multiple traversal patterns (e.g. threaded trees). Such idioms are prevalent in real-world implementations

→ $\bullet \bullet \bullet$
 $\in M$

Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from commutative monoid $(M, +, 0)$
- Flow propagation via continuous **edge functions**
- **The flow:** **some** fixed point, wrt. initial value

$m \in M$

$m \in M$

$m \in M$

[POPL'18]

Go with the Flow: Compositional Abstractions for Concurrent Data Structures*

SIDDHARTH KRISHNA, New York University, USA
DENNIS SHASHA, New York University, USA
THOMAS WIES, New York University, USA

Concurrent structures. H the sequent overlays. To the data stru quantity asso quantity as a root form a t which expres flow invariari semantic mo describing ce wide variety logic RGSep nontrivial co dictionary al proofs canno

CCS Concep **algorithms;**

Additional Ke

ACM Refere Siddharth Kr for Concurr <https://doi.org/10.1145/3133936.3133962>

1 INTRO

With the ad into solving da Rocha Pi

*This work is f 1339362, MCB-

[ESOP'20]

Local Reasoning for Global Graph Properties

Siddharth Krishna¹, Alexander J. Summers², and Thomas Wies¹

¹ New York University, New York, NY, USA, {siddharth,wies}@cs.nyu.edu
² ETH Zürich, Zurich, Switzerland, alexander.summers@inf.ethz.ch

Abstract. Separation logics are widely used for verifying programs that manipulate complex heap-based data structures. These logics build on so-called *separation algebras*, which allow expressing properties of heap regions such that modifications to a region do not invalidate properties stated about the remainder of the heap. This concept is key to enabling modular reasoning and also extends to concurrency. While heaps are naturally related to mathematical graphs, many ubiquitous graph properties are non-local in character, such as reachability between nodes, path lengths, acyclicity and other structural invariants, as well as data invariants which combine with these notions. Reasoning modularly about such graph properties remains notoriously difficult, since a local modification can have side-effects on a global property that cannot be easily confined to a small region.

In this paper, we address the question: What separation algebra can be used to avoid proof arguments reverting back to tedious global reasoning in such cases? To this end, we consider a general class of global graph properties expressed as fixpoints of algebraic equations over graphs. We present mathematical foundations for reasoning about this class of properties, imposing minimal requirements on the underlying theory that allow us to define a suitable separation algebra. Building on this theory, we develop a general proof technique for modular reasoning about

→ $\in M$

Lower bounds meaningless (flow may vanish in compositions).
Computation has to restrict graph structure (e.g. acyclicity).

Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from commutative monoid $(\mathbb{M}, +, 0)$
- Flow propagation via continuous **edge functions**
- **The flow:**
some fixed point, wrt. initial value

$m \in \mathbb{M}$

$m \in \mathbb{M}$

$m \in \mathbb{M}$

[POPL'18]

Go with the Flow: Compositional Abstractions for Concurrent Data Structures*

SIDDHARTH KRISHNA, New York University, USA
DENNIS SHASHA, New York University, USA
THOMAS WIES, New York University, USA

Concurrent structures. H the sequent overlays. To the data stru quantity asso quantity as a root form a t which expres flow invariari semantic mo describing ce wide variety logic RGSep nontrivial co dictionary al proofs canno

CCS Concep **algorithms;**

Additional Ke

ACM Refere Siddharth Kr for Concurr <https://doi.org/>

1 INTRO

With the ad into solving da Rocha Pi

*This work is f 1339362, MCB-

[ESOP'20]

Local Reasoning for Global Graph Properties

Siddharth Krishna¹, Alexander J. Summers², and Thomas Wies¹

¹ New York University, New York, NY, USA, {siddharth,wies}@cs.nyu.edu
² ETH Zürich, Zurich, Switzerland, alexander.summers@inf.ethz.ch

Abstract. Separation logics are widely used for verifying programs that manipulate complex heap-based data structures. These logics build on so-called *separation algebras*, which allow expressing properties of heap regions such that modifications to a region do not invalidate properties stated about the remainder of the heap. This concept is key to enabling modular reasoning and also extends to concurrency. While heaps are naturally related to mathematical graphs, many ubiquitous graph properties are non-local in character, such as reachability between nodes, path lengths, acyclicity and other structural invariants, as well as data invariants which combine with these notions. Reasoning modularly about such graph properties remains notoriously difficult, since a local modification can have side-effects on a global property that cannot be easily confined to a small region.

In this paper, we address the question: What separation algebra can be used to avoid proof arguments reverting back to tedious global reasoning in such cases? To this end, we consider a general class of global graph properties expressed as fixpoints of algebraic equations over graphs. We present mathematical foundations for reasoning about this class of properties, imposing minimal requirements on the underlying theory that allow us to define a suitable separation algebra. Building on this theory, we develop a general proof technique for modular reasoning about

→ $\in \mathbb{M}$

Lower bounds meaningless (flow may vanish in compositions).
Computation has to restrict graph structure (e.g. acyclicity).

Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from commutative monoid $(M, +, 0)$
- Flow propagation via continuous **edge functions**
- **The flow:** fixed point, wrt. initial value

$m \in M$
 $m \in M$
 $m \in M$

NEW

Lower bounds!
Computation has mild assumptions about monoid + edge functions.

[POPL'18]
Go with the Flow: Compositional Abstractions for Concurrent Data Structures*
SIDDHARTH KRISHNA, New York University, USA
DENNIS SHASHA, New York University, USA
THOMAS WIES, New York University, USA

[ESOP'20]
Local Reasoning for Global Graph Properties
Siddharth Krishna¹, Alexander J. Summers², and Thomas Wies¹
¹ New York University, New York, NY, USA, {siddharth,wies}@cs.nyu.edu
² ETH Zürich, Zurich, Switzerland, alexander.summers@inf.ethz.ch

[TACAS'23]
Make flows small again: revisiting the flow framework
Roland Meyer¹, Thomas Wies², and Sebastian Wolff²
¹ TU Braunschweig, Braunschweig, Germany, meyer@tu-bs.de
² New York University, New York, USA, {wies,sebastian.wolff}@cs.nyu.edu

Abstract We present a new flow framework for separation logic reasoning about programs that manipulate general graphs. The framework overcomes problems in earlier developments: it is based on standard fixed point theory, guarantees least flows, rules out vanishing flows, and has an easy to understand notion of footprint as needed for soundness of the frame rule. In addition, we present algorithms for automating the framework, which are evaluated on benchmarks selected from onstrates ve previ-

ration logic [5, general graphs in algorithms that tance Protocol, [22, 24, 27, 34]. underlying meta

Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from commutative monoid $(M, +, 0)$
- Flow propagation via continuous **edge functions**
- **The flow:** **least** fixed point, wrt. initial value

$m \in M$

$m \in M$

$m \in M$



Lower bounds!
Computation has mild assumptions
about monoid + edge functions.

Go with the Flow: Compositional Abstractions for Concurrent Data Structures*

SIDDHARTH KRISHNA, New York University, USA
DENNIS SHASHA, New York University, USA
THOMAS WIES, New York University, USA

Concurrent structures. H the sequentia overlays. To the data stru quantity asso quantity as a root form a ti which expres flow invarian semantic mo describing ce wide variety logic RGSep nontrivial co dictionary al proofs canno

CCS Concep **algorithms;**

Additional Ke
ACM Refer
Siddharth Kr
for Concurr
<https://doi.org/>

1 INTRO
With the ad
into solving
da Rocha Pi

*This work is f

[POPL'18]

Local Reasoning for Global Graph Properties

Siddharth Krishna¹, Alexander J. Summers², and Thomas Wies¹

¹ New York University, New York, NY, USA, {siddharth,wies}@cs.nyu.edu
² ETH Zürich, Zurich, Switzerland, alexander.summers@inf.ethz.ch

[ESOP'20]

Make flows small again: revisiting the flow framework

Roland Meyer¹, Thomas Wies², and Sebastian Wolff^{2(✉)}

¹ TU Braunschweig, Braunschweig, Germany, meyer@tu-bs.de
² New York University, New York, USA, {wies,sebastian.wolff}@cs.nyu.edu

Abstract We present a new flow framework for separation logic reasoning about programs that manipulate general graphs. The framework overcomes problems in earlier developments: it is based on standard fixed point theory, guarantees least flows, rules out vanishing flows, and has an easy to understand notion of footprint as needed for soundness of the frame rule. In addition, we present algorithms for

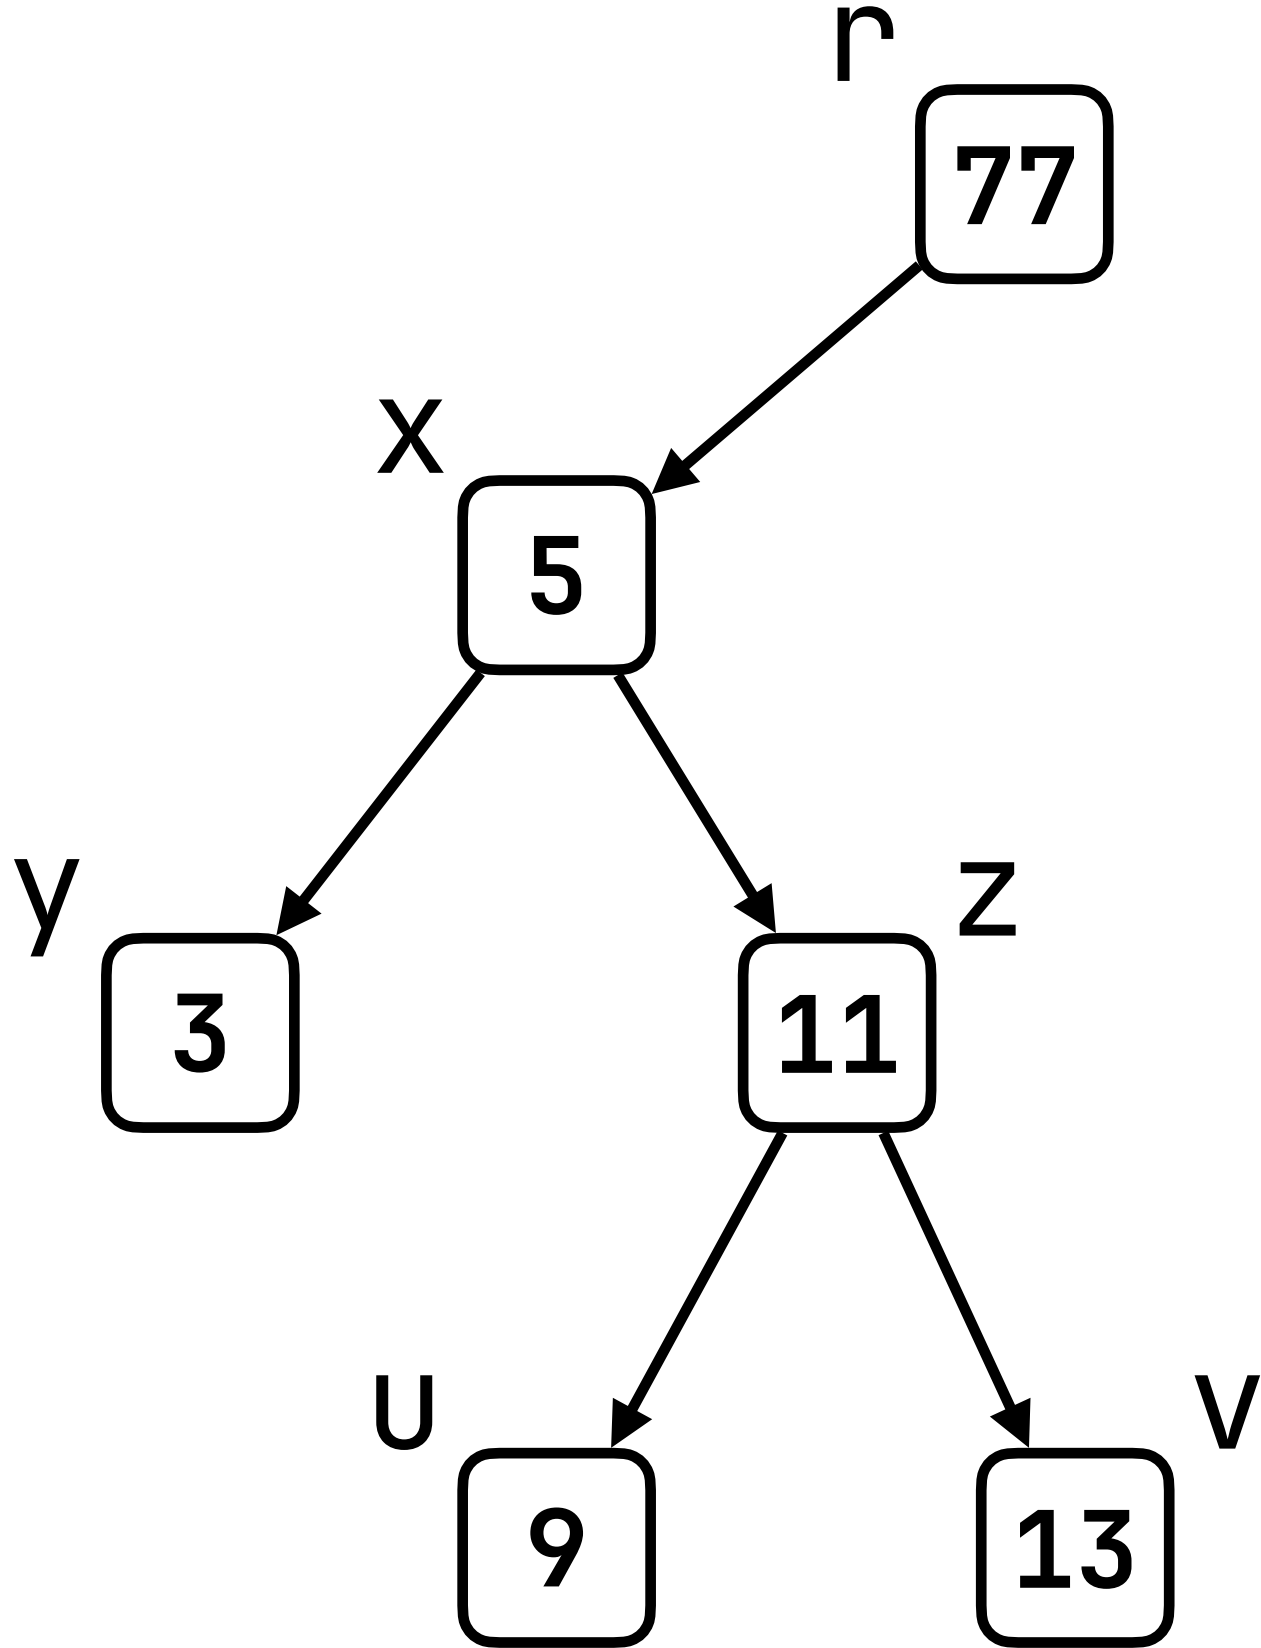
onstrates
ve previ-

eration logic [5,
neral graphs in
algorithms that
tance Protocol,
[22, 24, 27, 34].
nderlying meta

[TACAS'23]

Flow Framework

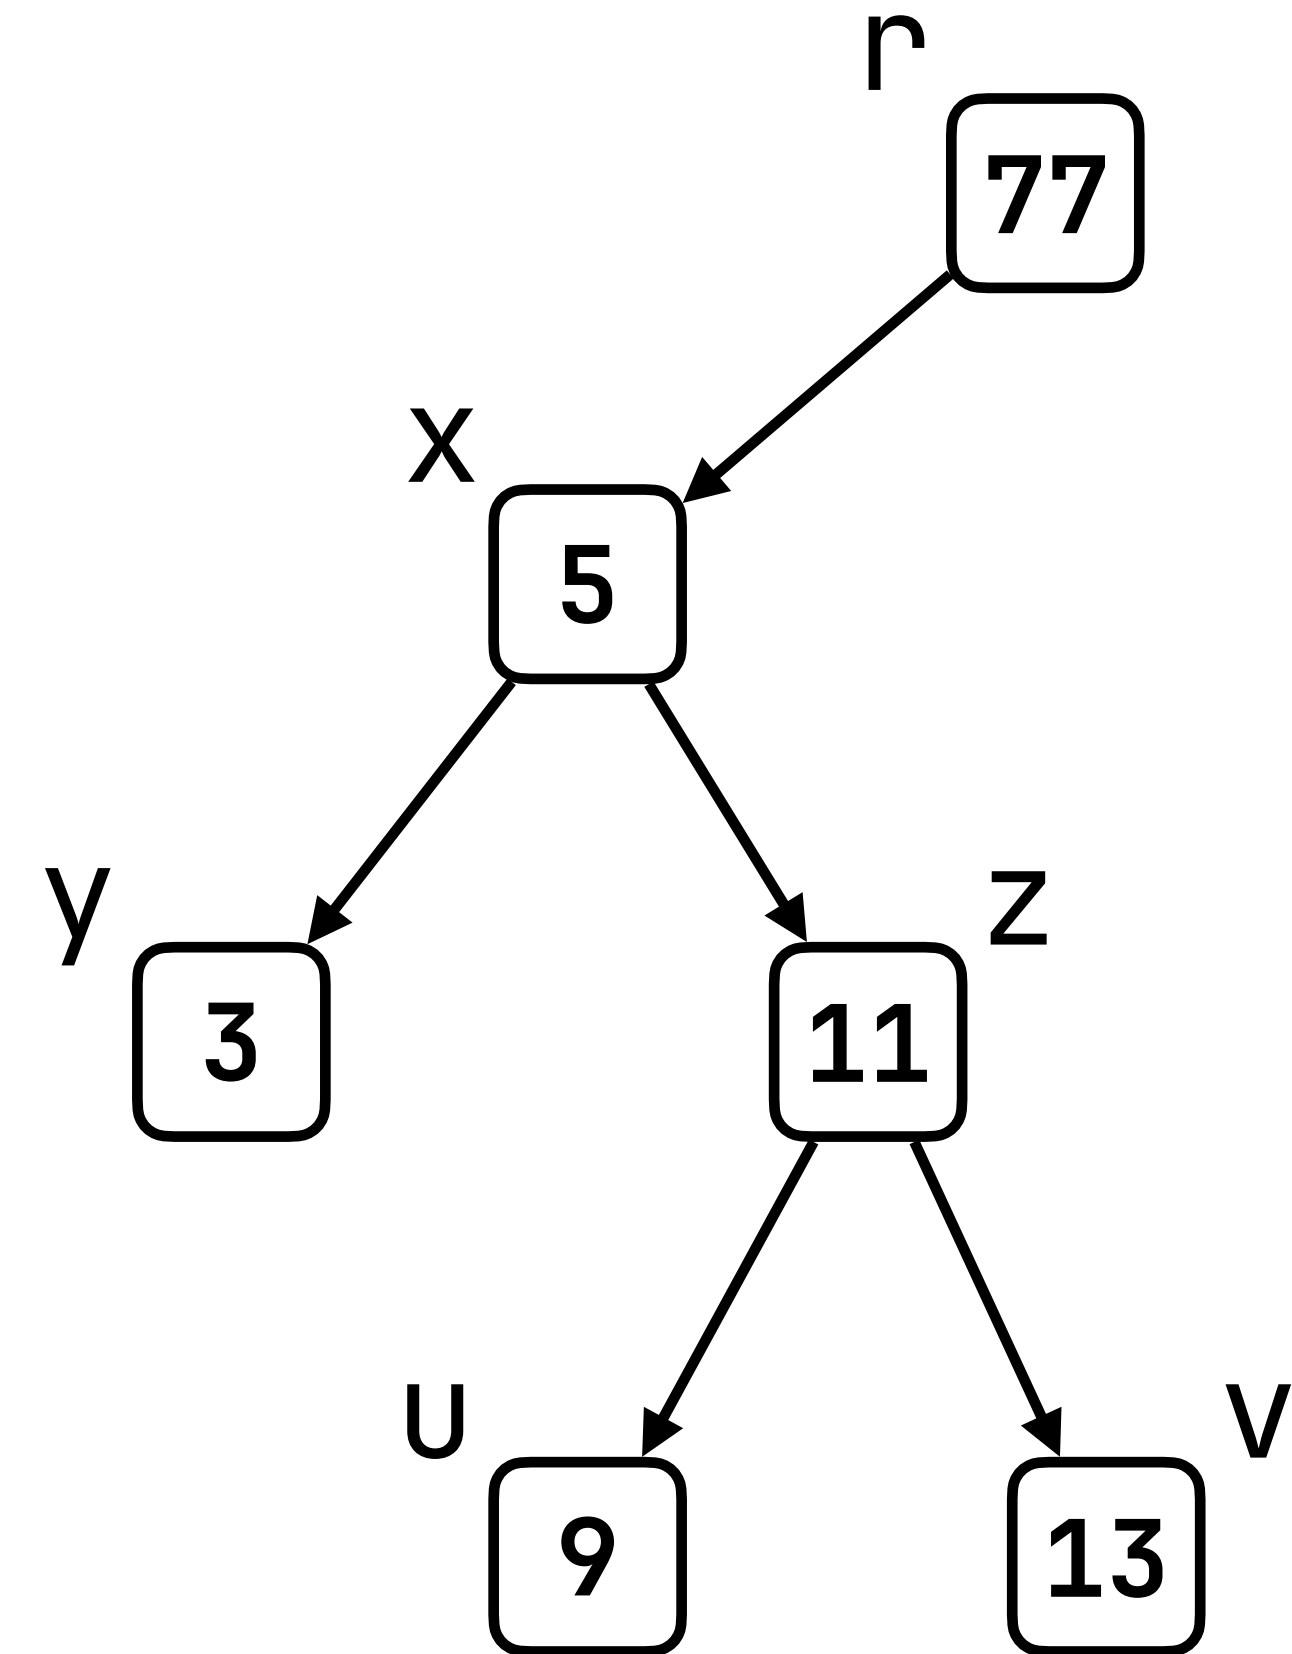
- Flow values from
- Flow propagation via
- ***The flow:***



Flow Framework

flow graphs = heap graphs + data flow values (ghost)

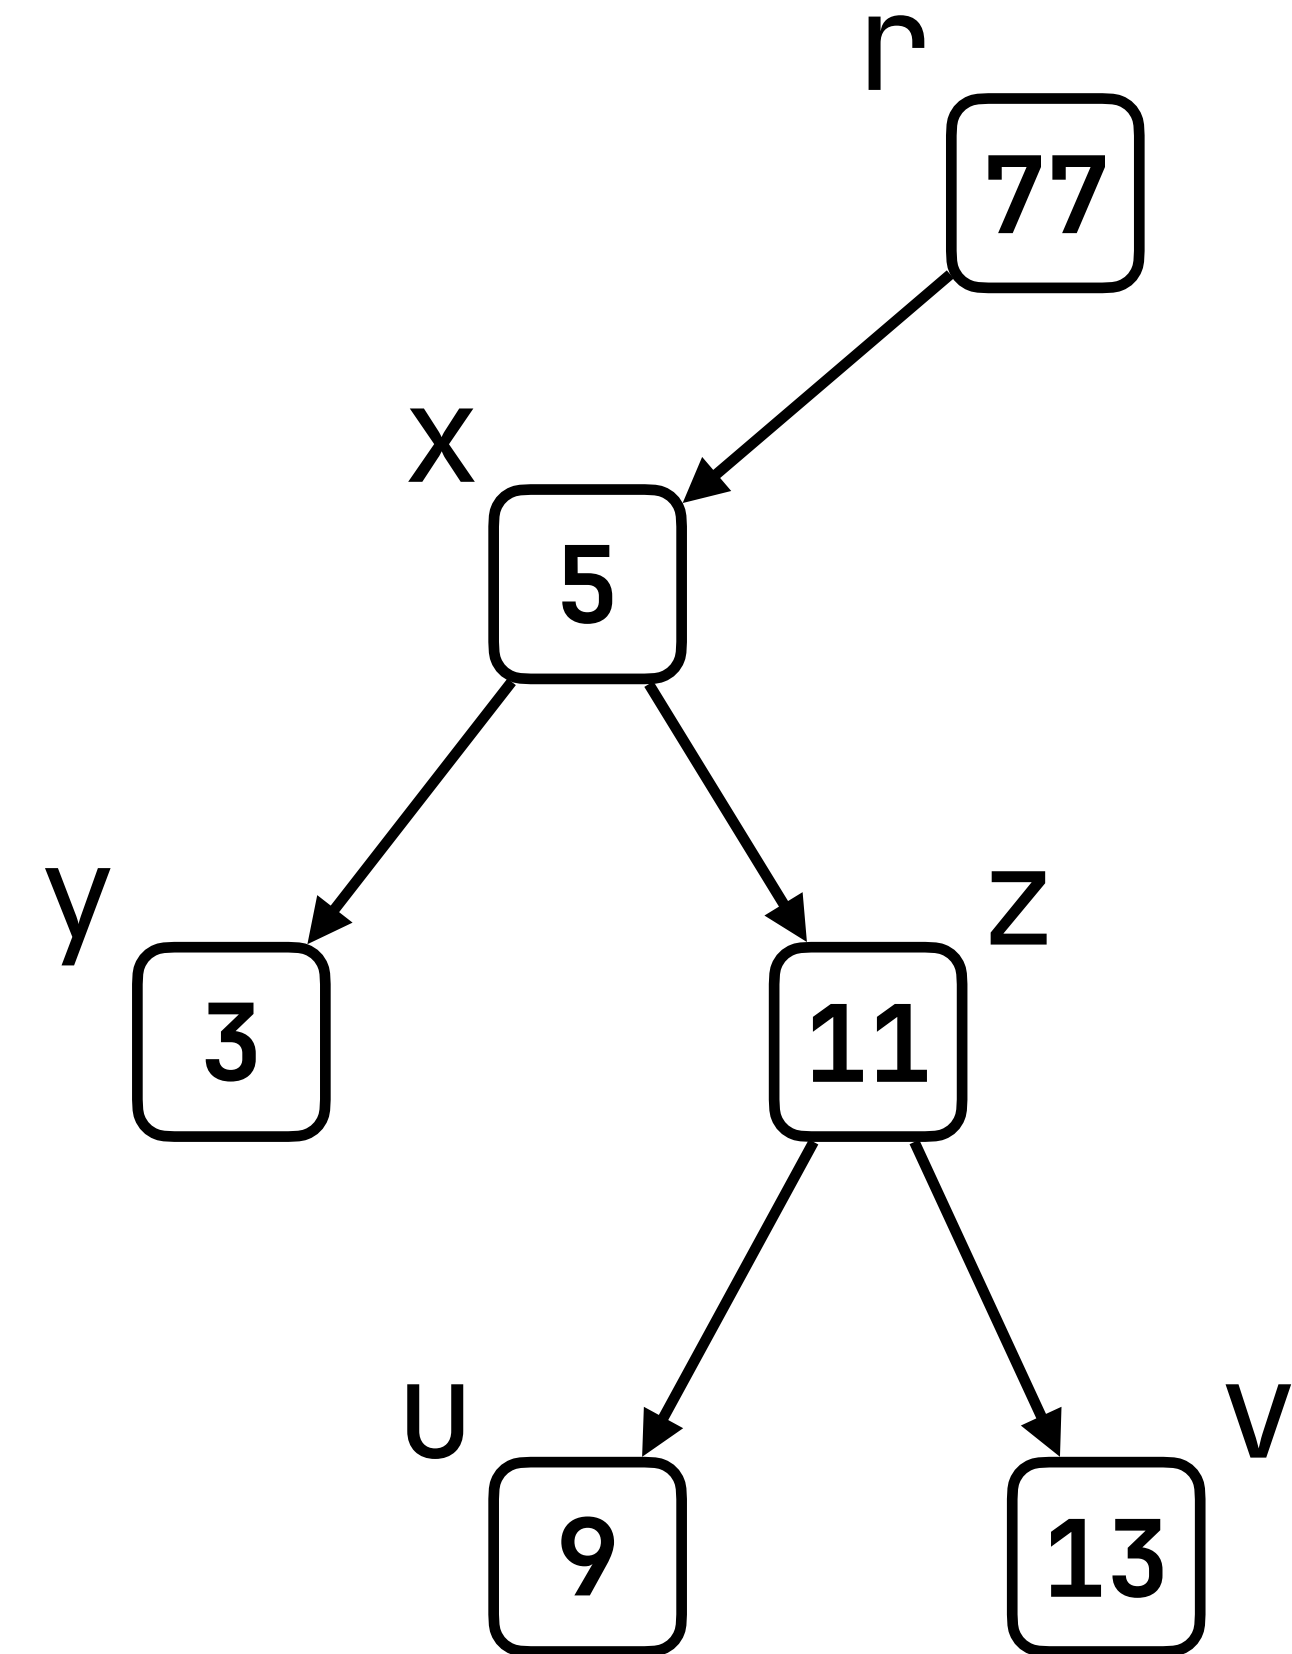
- Flow values from
- Flow propagation via
- ***The flow:***



Flow Framework

flow graphs = heap graphs + data flow values (ghost)

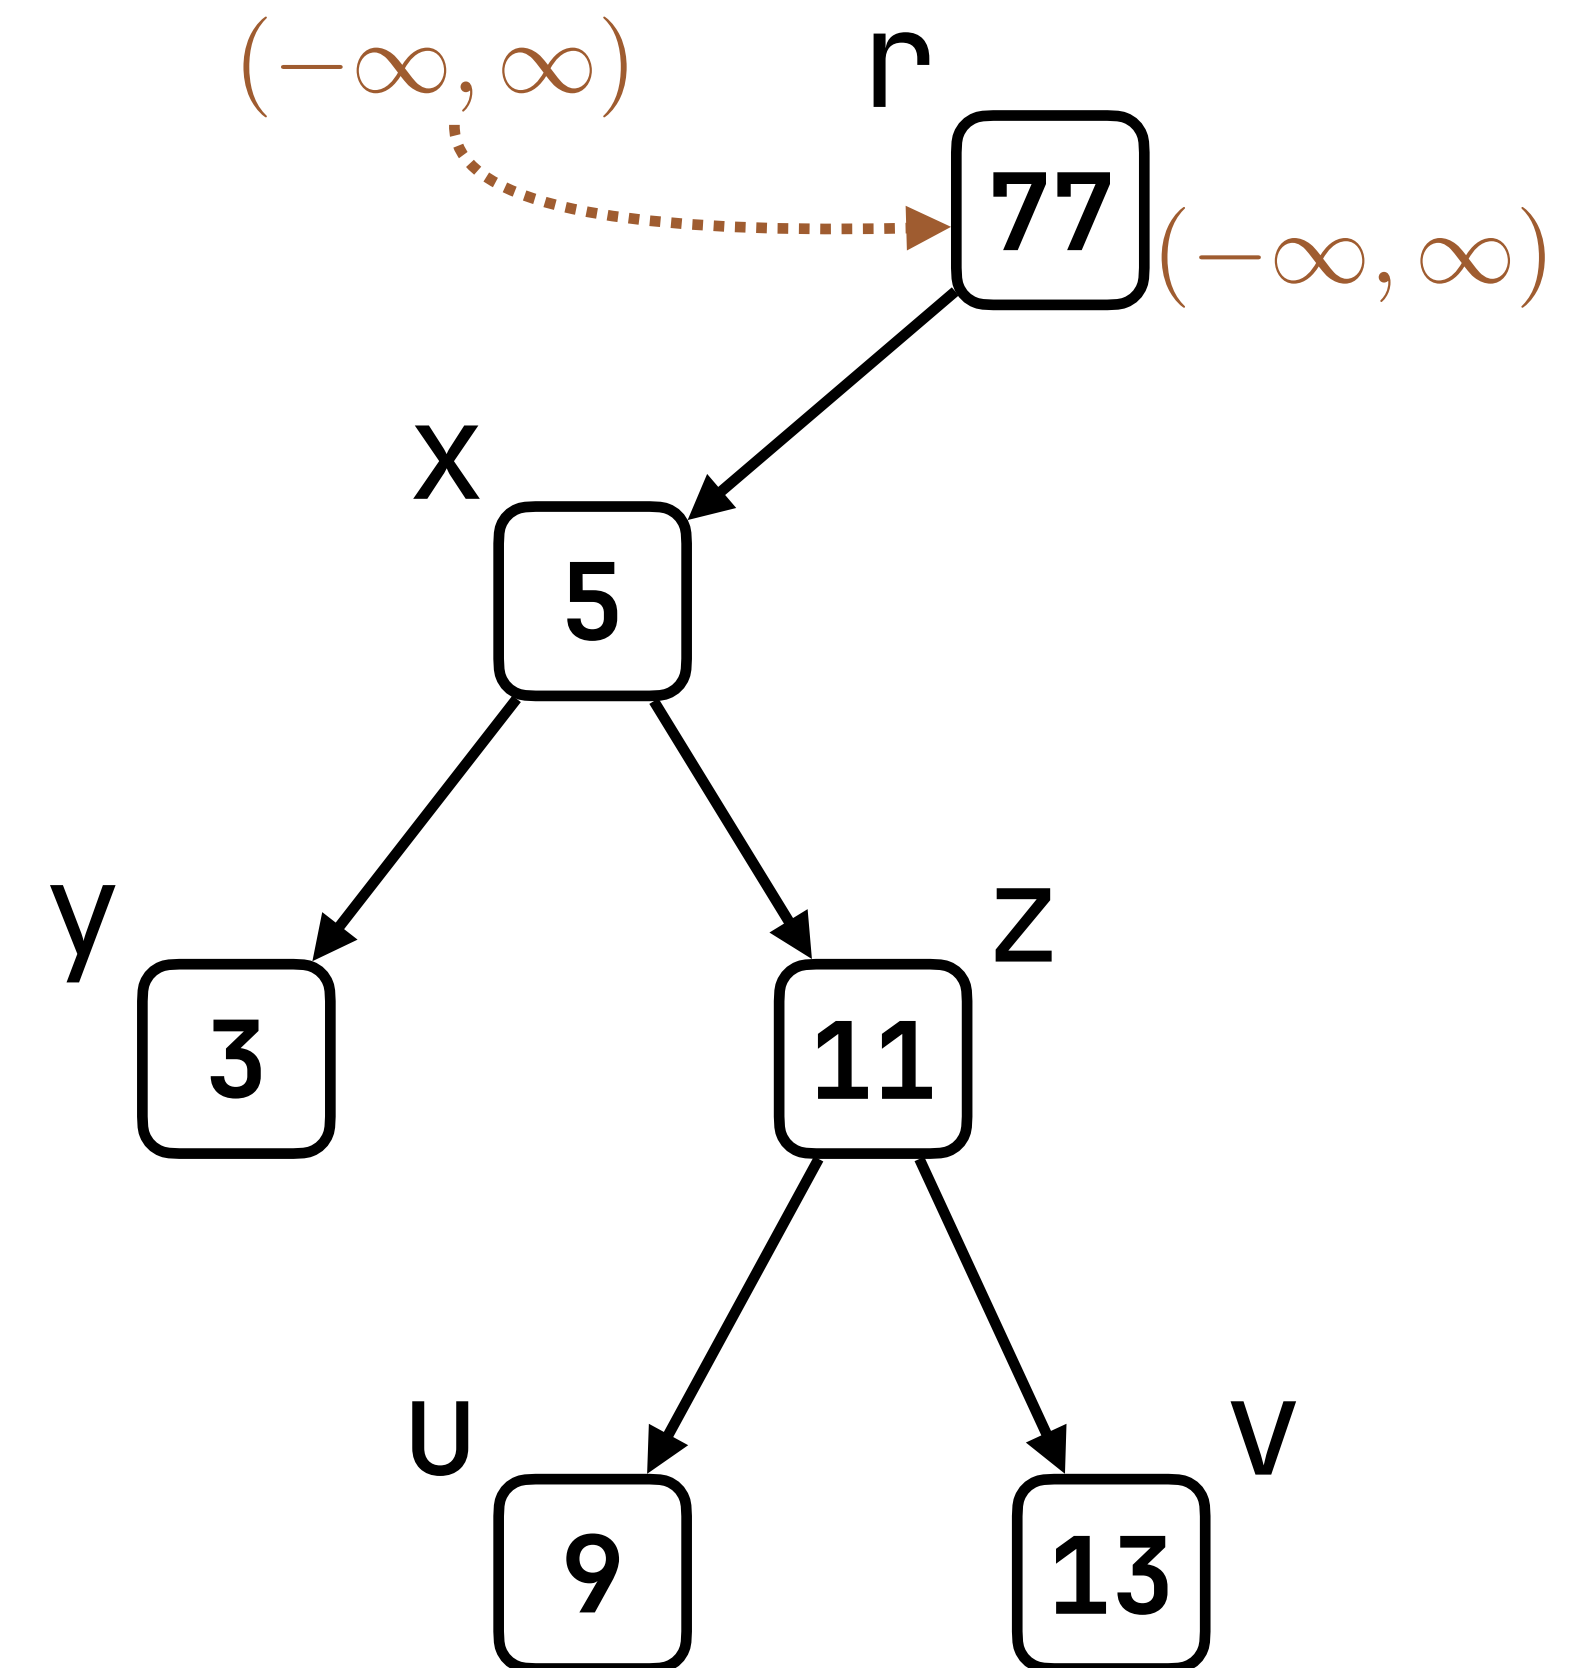
- Flow values from search path monoid $(2^{\mathbb{Z} \cup \{-\infty, \infty\}}, \cup, \emptyset)$
- Flow propagation via
- ***The flow:***



Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from search path monoid $(2^{\mathbb{Z} \cup \{-\infty, \infty\}}, \cup, \emptyset)$
- Flow propagation via
- ***The flow:***



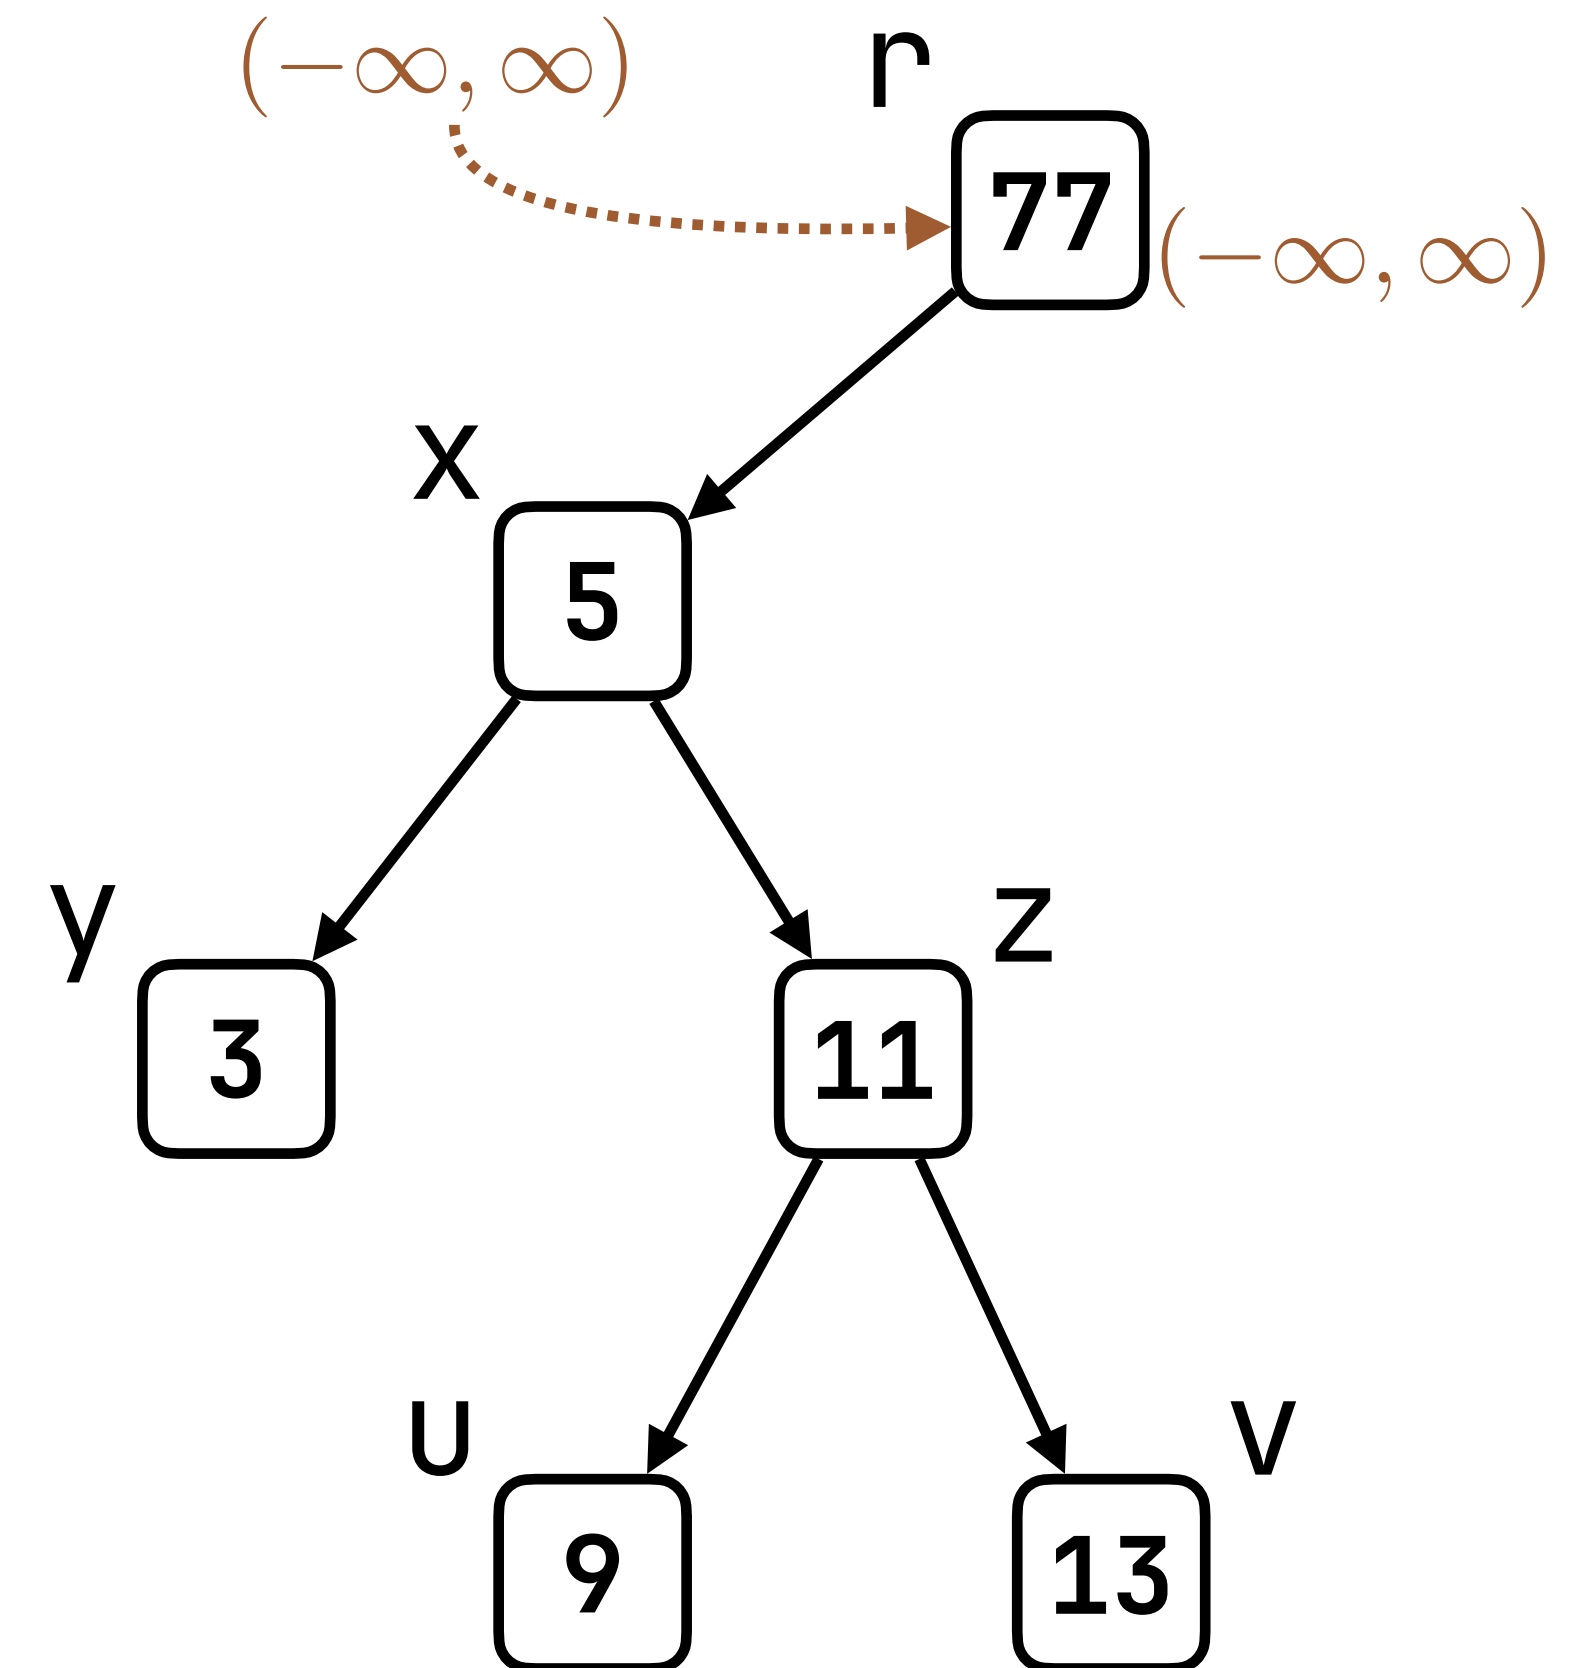
Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from search path monoid $(2^{\mathbb{Z} \cup \{-\infty, \infty\}}, \cup, \emptyset)$

- Flow propagation via
 - $\lambda_{<k} := \lambda m. m \cap (-\infty, k)$
 - $\lambda_{>k} := \lambda m. m \cap (k, \infty)$

- The flow:**



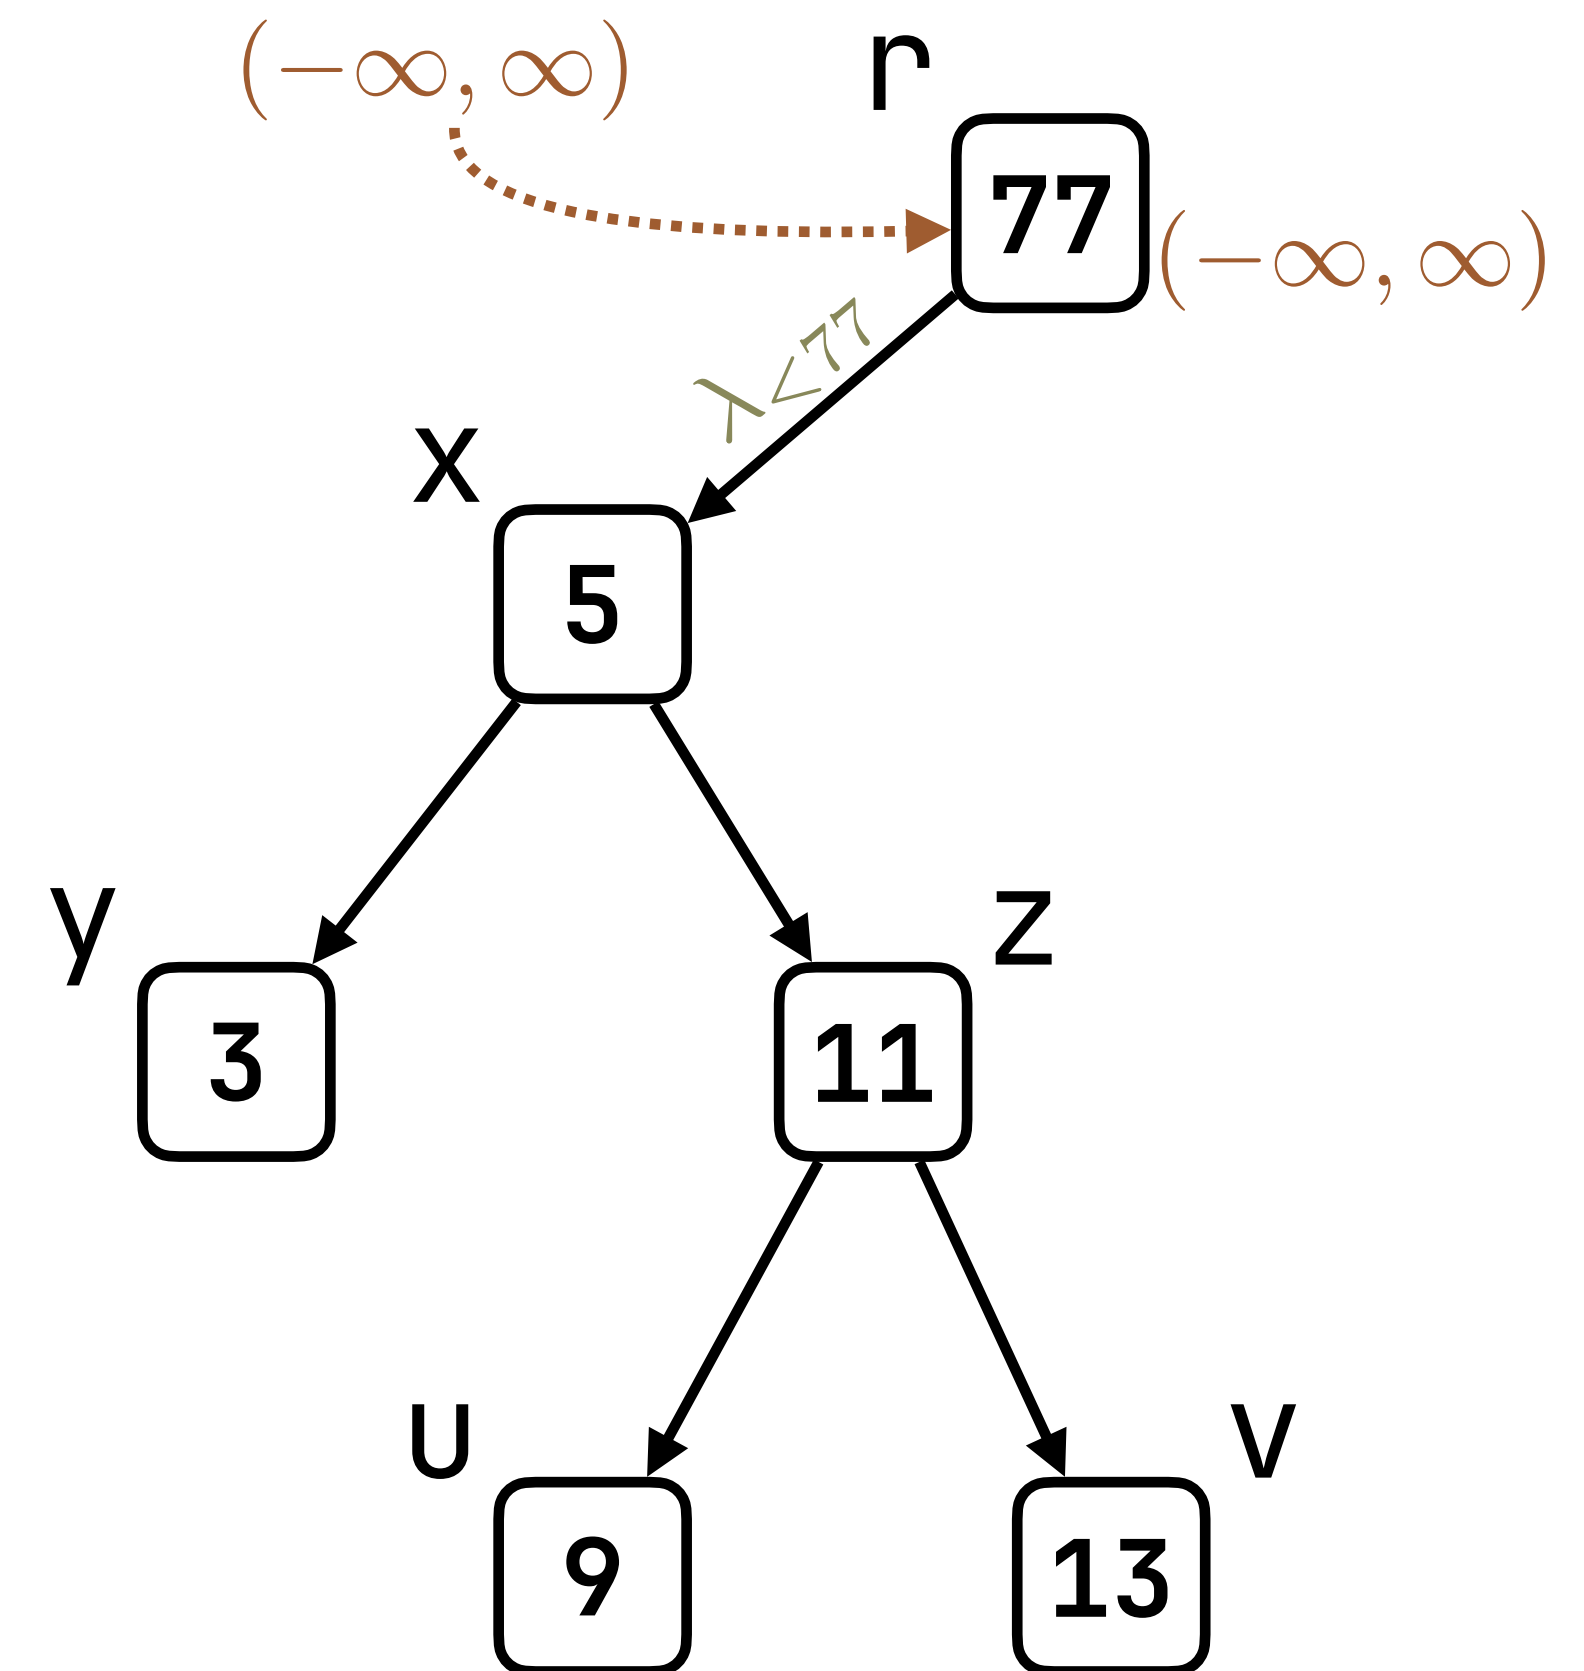
Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from search path monoid $(2^{\mathbb{Z} \cup \{-\infty, \infty\}}, \cup, \emptyset)$

- Flow propagation via
$$\lambda_{<k} := \lambda m. m \cap (-\infty, k)$$
$$\lambda_{>k} := \lambda m. m \cap (k, \infty)$$

- The flow:**



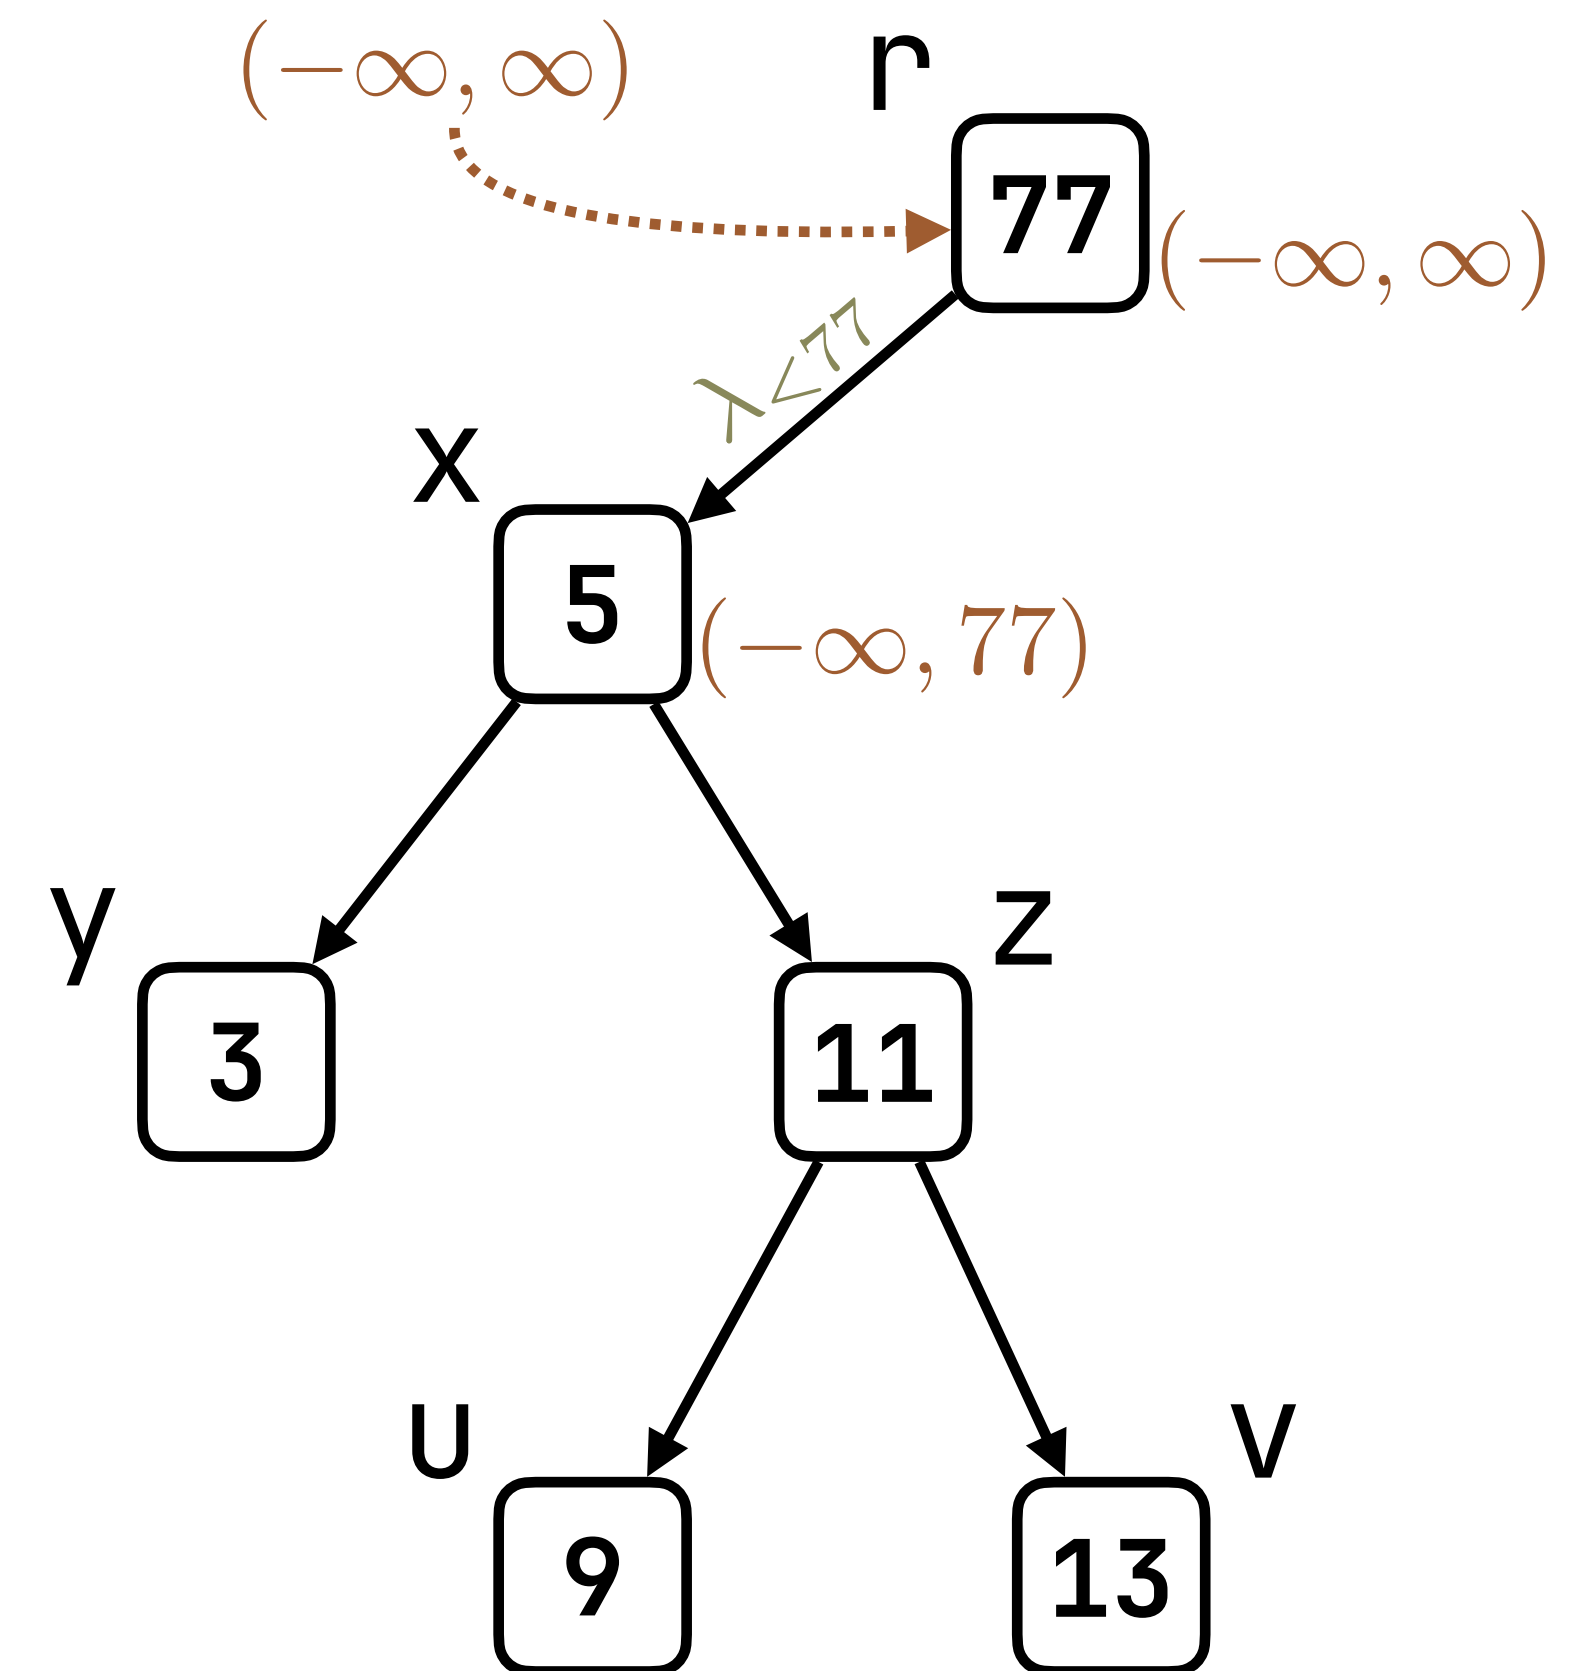
Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from search path monoid $(2^{\mathbb{Z} \cup \{-\infty, \infty\}}, \cup, \emptyset)$

- Flow propagation via
$$\lambda_{<k} := \lambda m. m \cap (-\infty, k)$$
$$\lambda_{>k} := \lambda m. m \cap (k, \infty)$$

- The flow:**



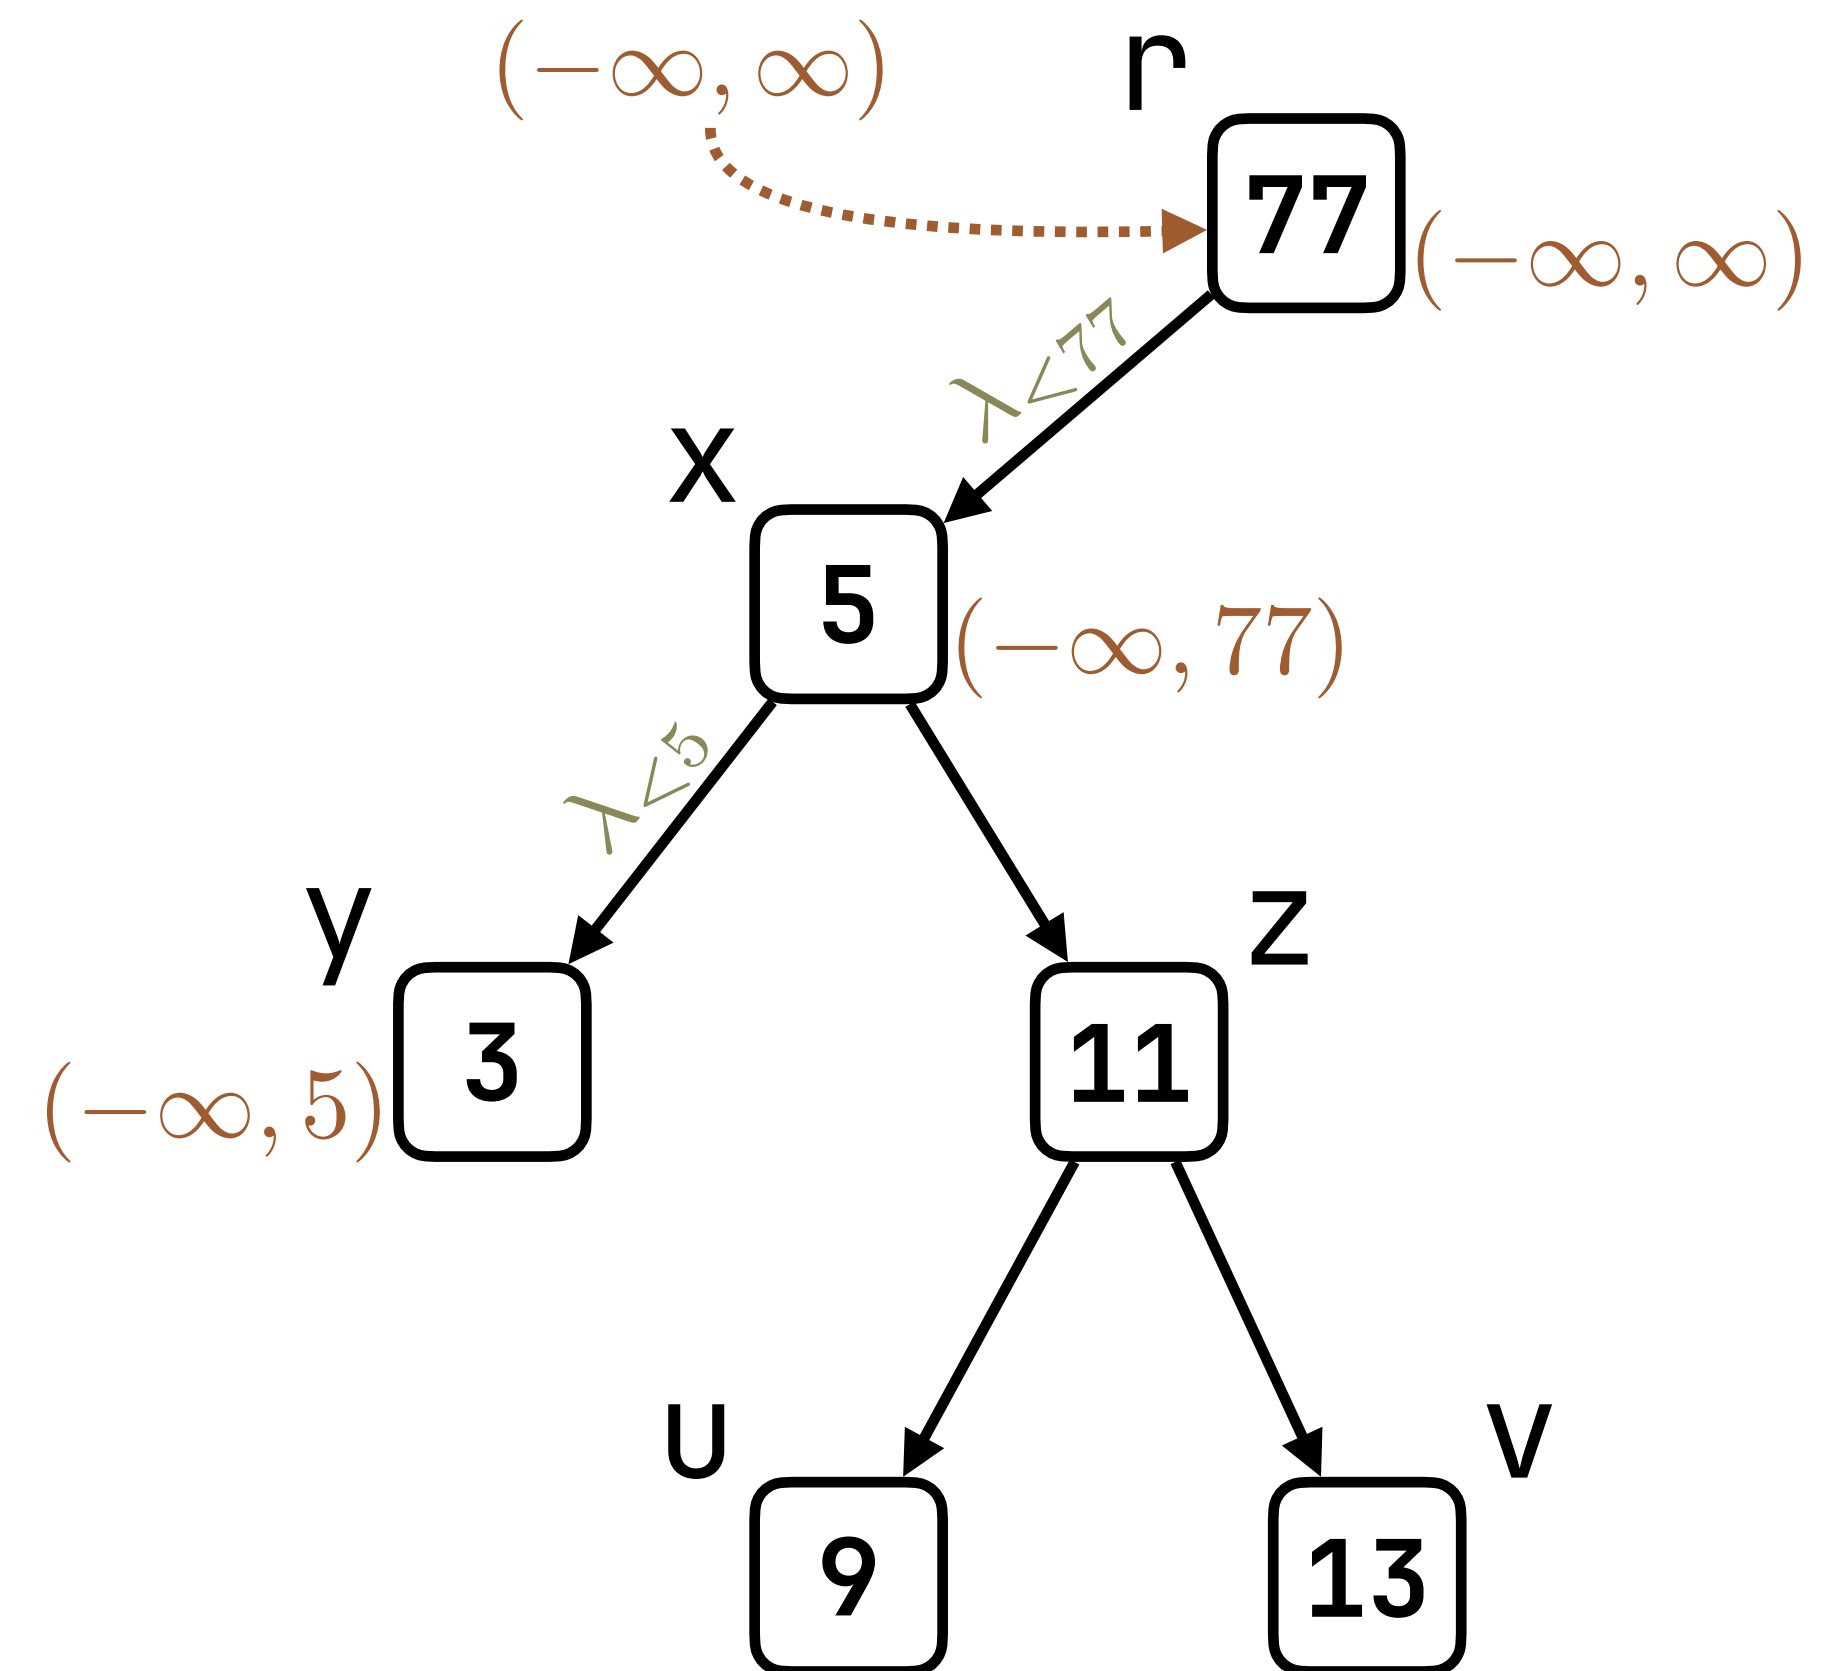
Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from search path monoid $(2^{\mathbb{Z} \cup \{-\infty, \infty\}}, \cup, \emptyset)$

- Flow propagation via
$$\lambda_{<k} := \lambda m. m \cap (-\infty, k)$$
$$\lambda_{>k} := \lambda m. m \cap (k, \infty)$$

- The flow:**



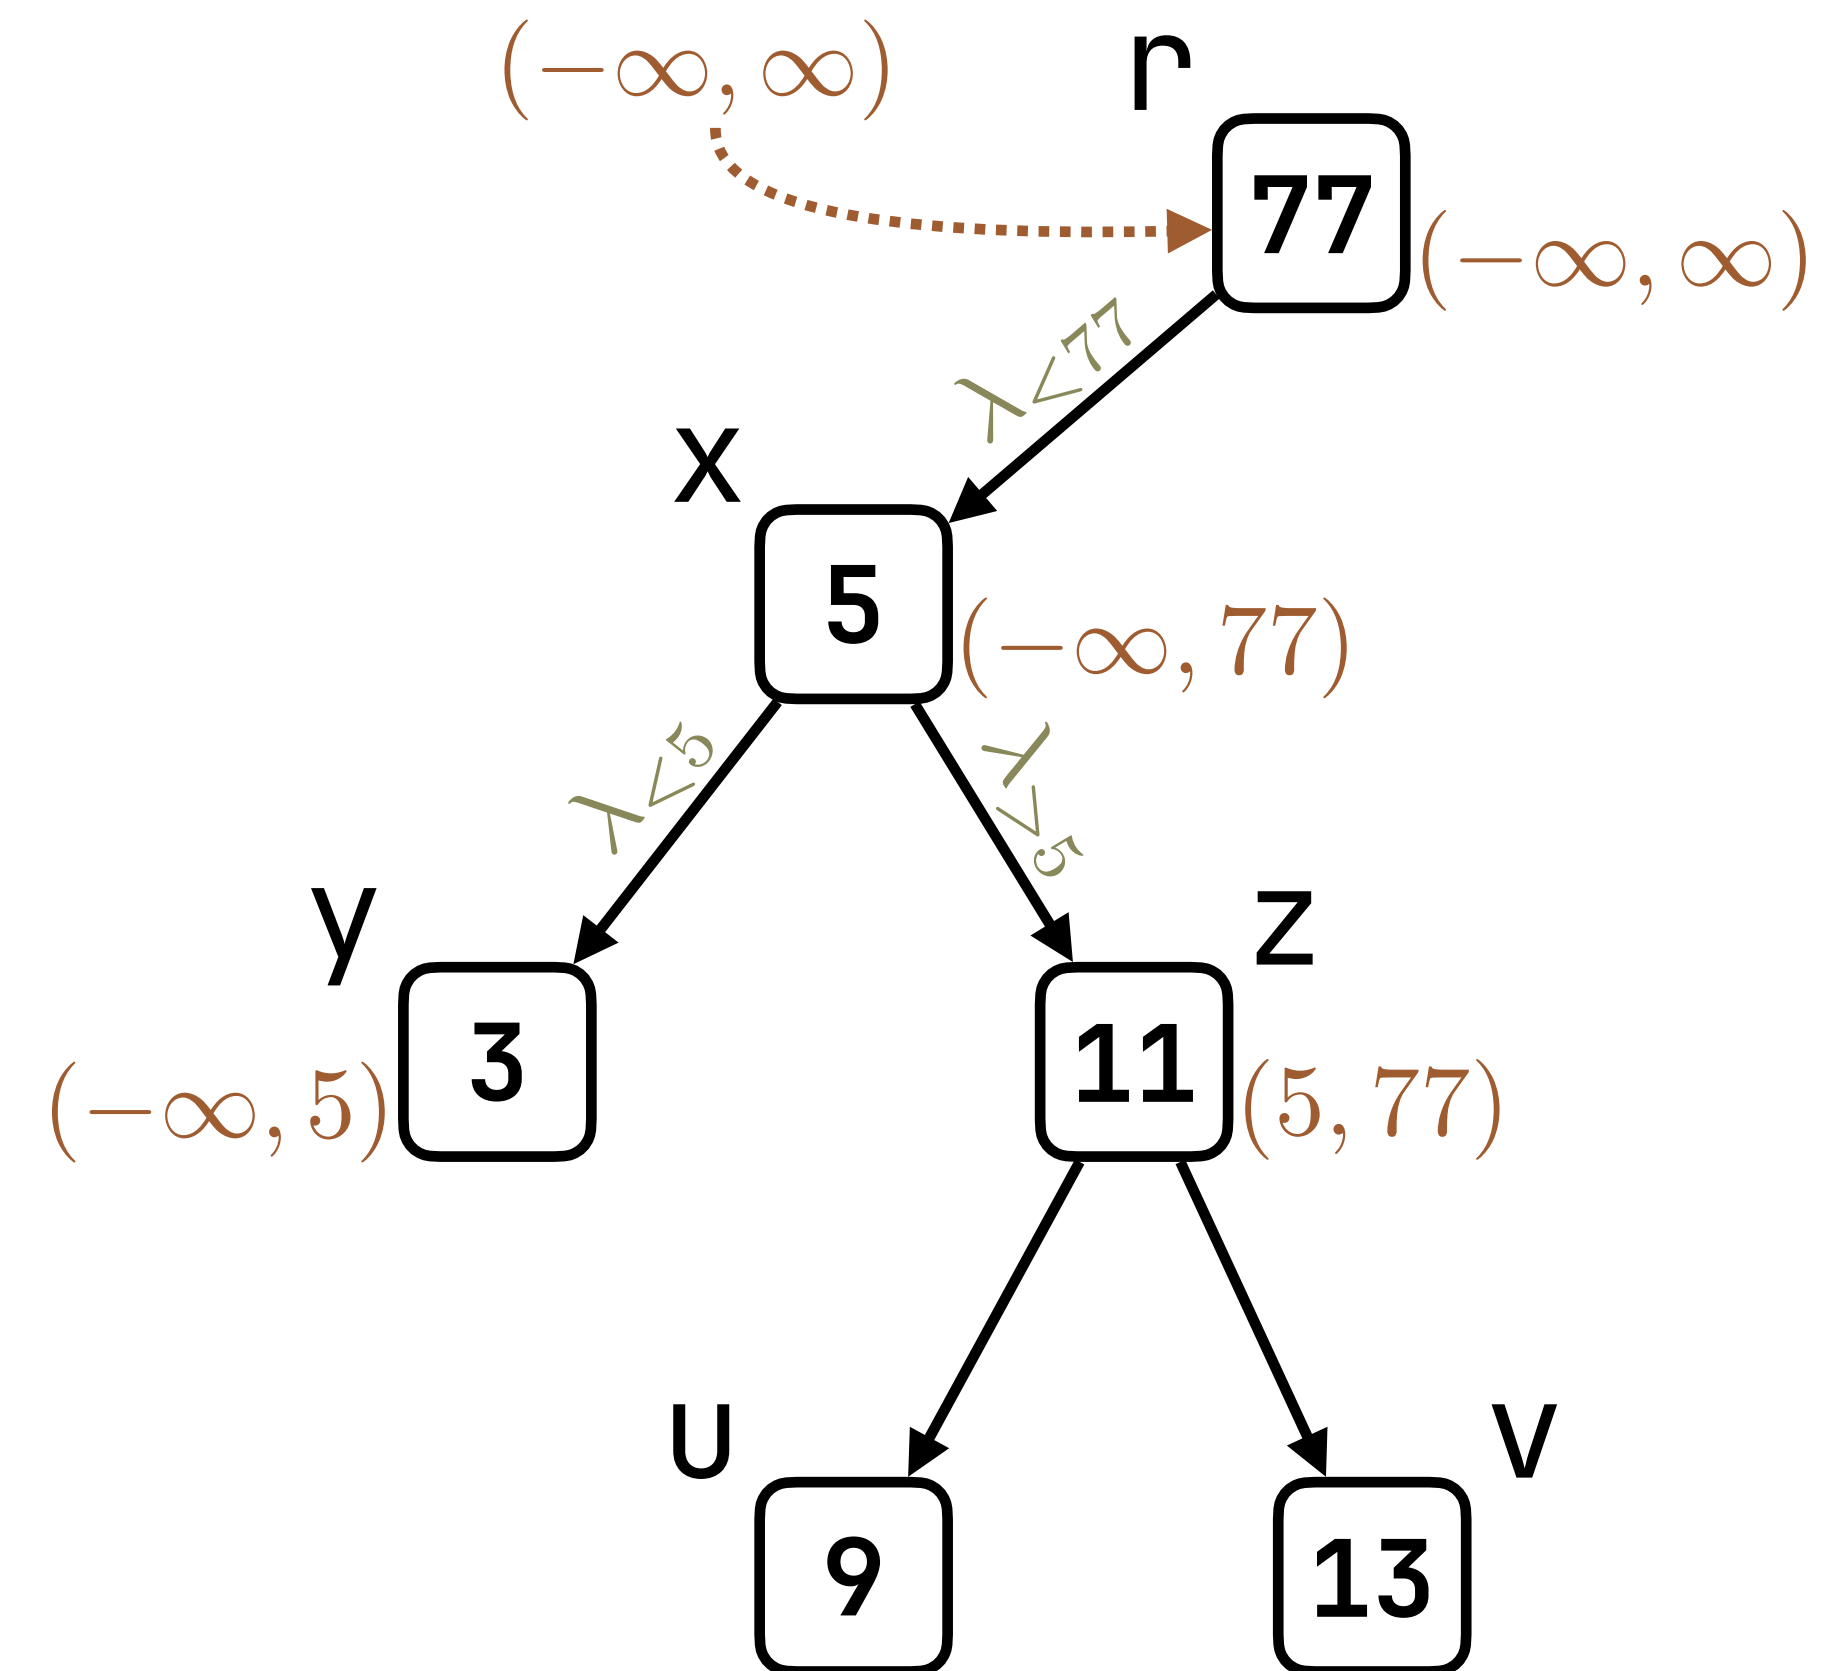
Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from search path monoid $(2^{\mathbb{Z} \cup \{-\infty, \infty\}}, \cup, \emptyset)$

- Flow propagation via
$$\lambda_{<k} := \lambda m. m \cap (-\infty, k)$$
$$\lambda_{>k} := \lambda m. m \cap (k, \infty)$$

- The flow:**



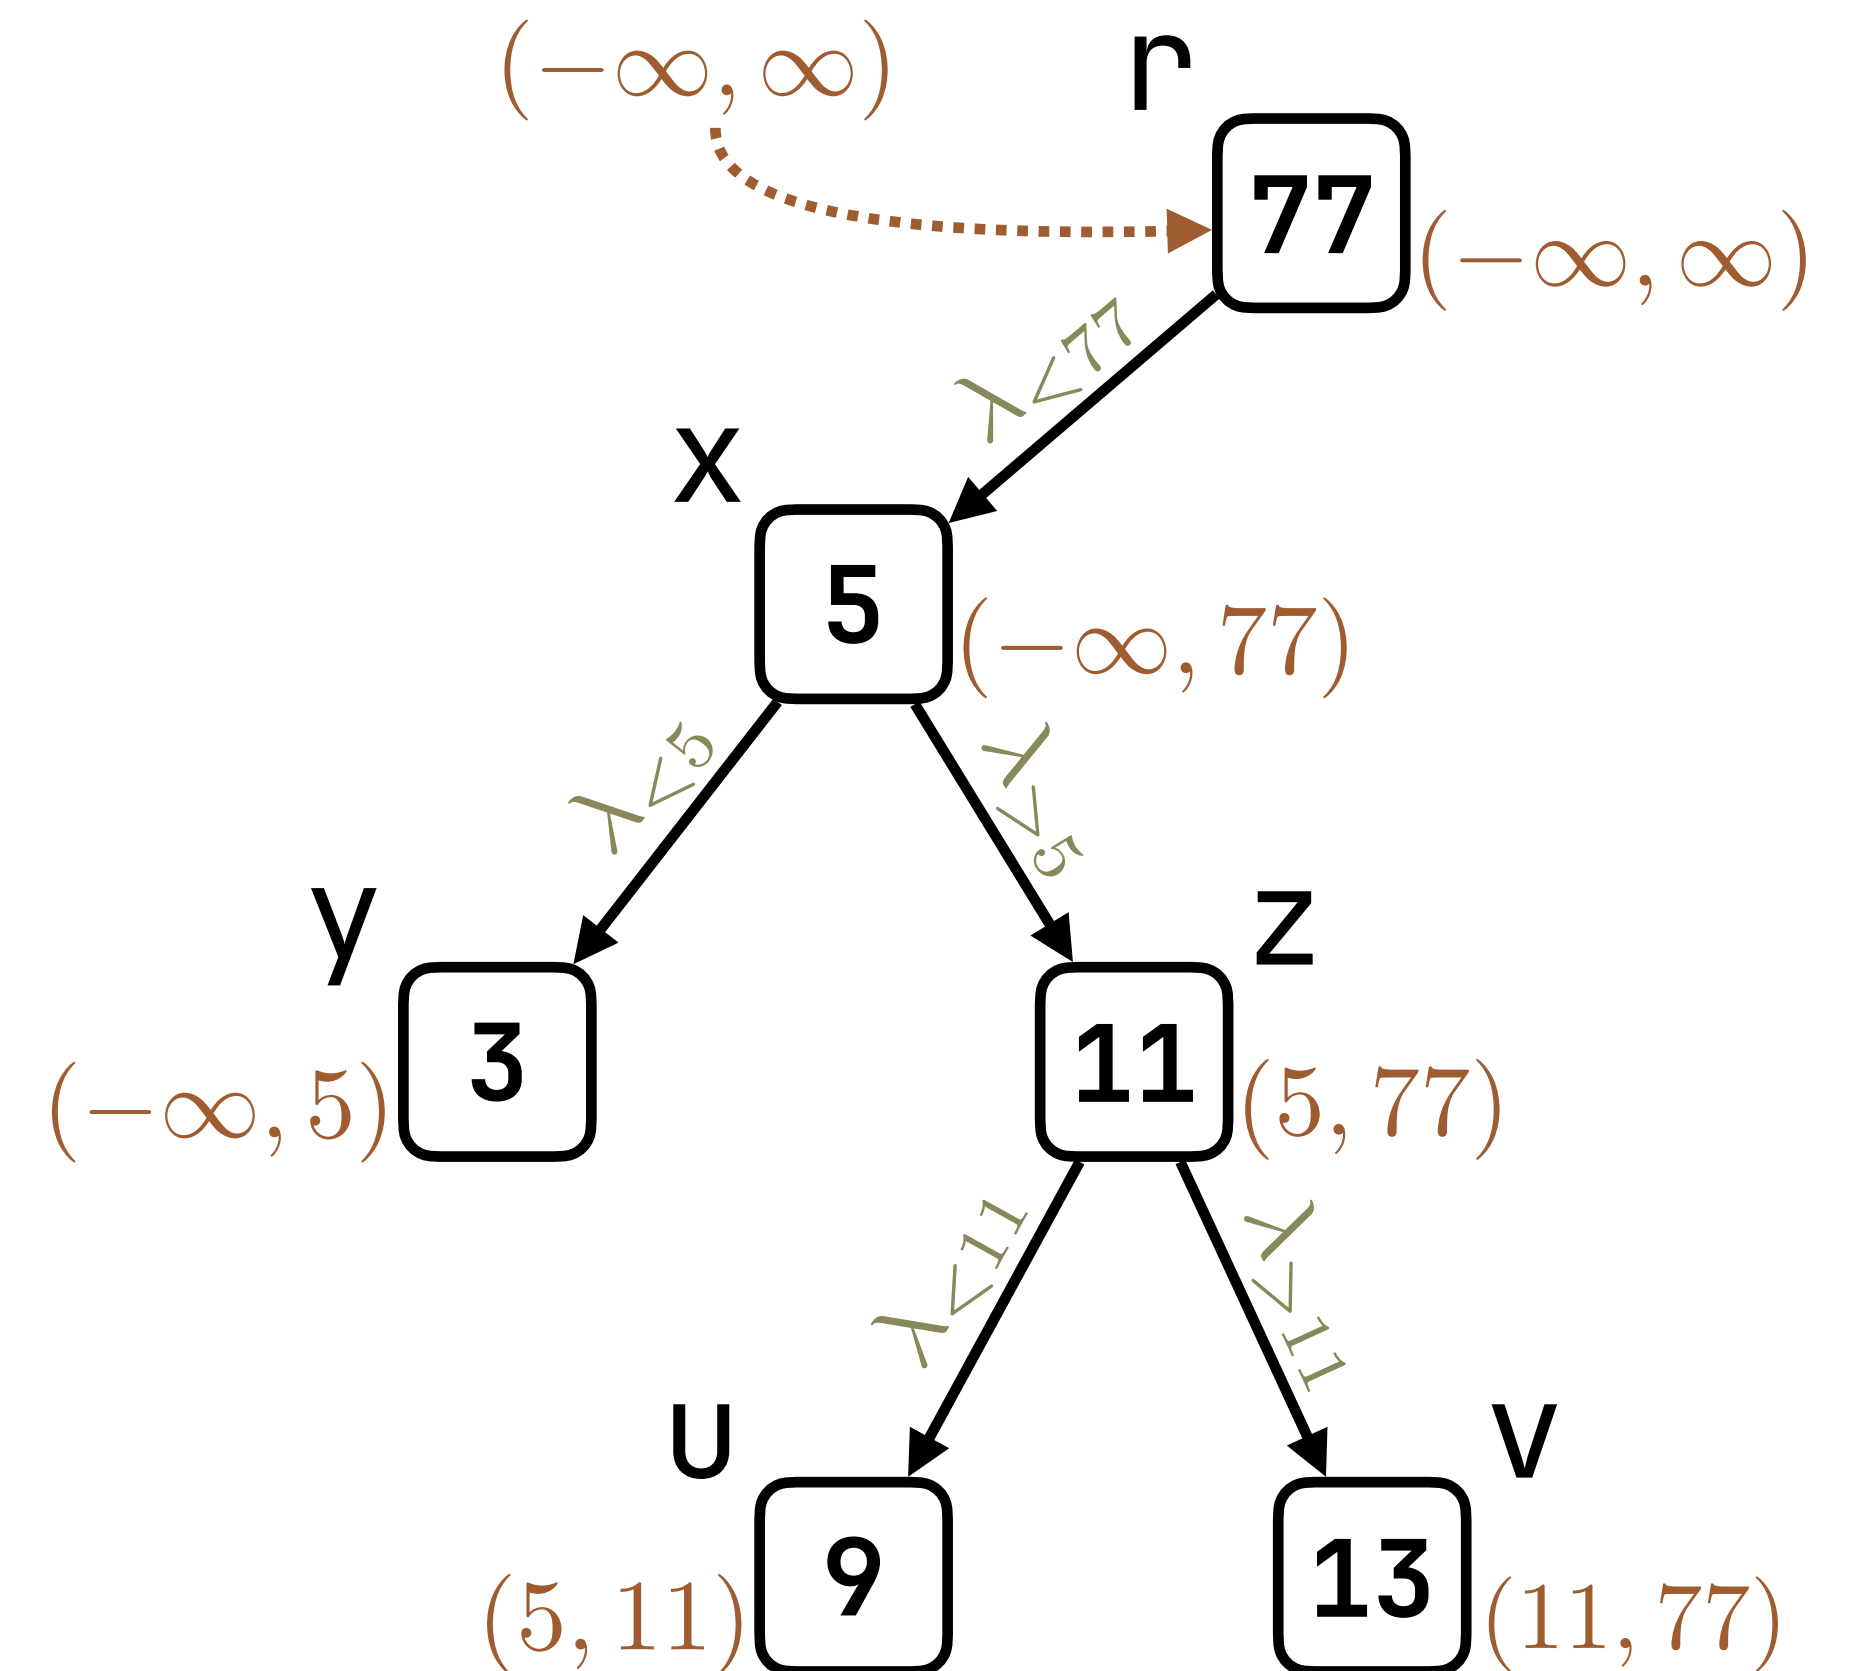
Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from search path monoid $(2^{\mathbb{Z} \cup \{-\infty, \infty\}}, \cup, \emptyset)$

- Flow propagation via
$$\lambda_{<k} := \lambda m. m \cap (-\infty, k)$$
$$\lambda_{>k} := \lambda m. m \cap (k, \infty)$$

- **The flow:**



Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from search path monoid $(2^{\mathbb{Z} \cup \{-\infty, \infty\}}, \cup, \emptyset)$

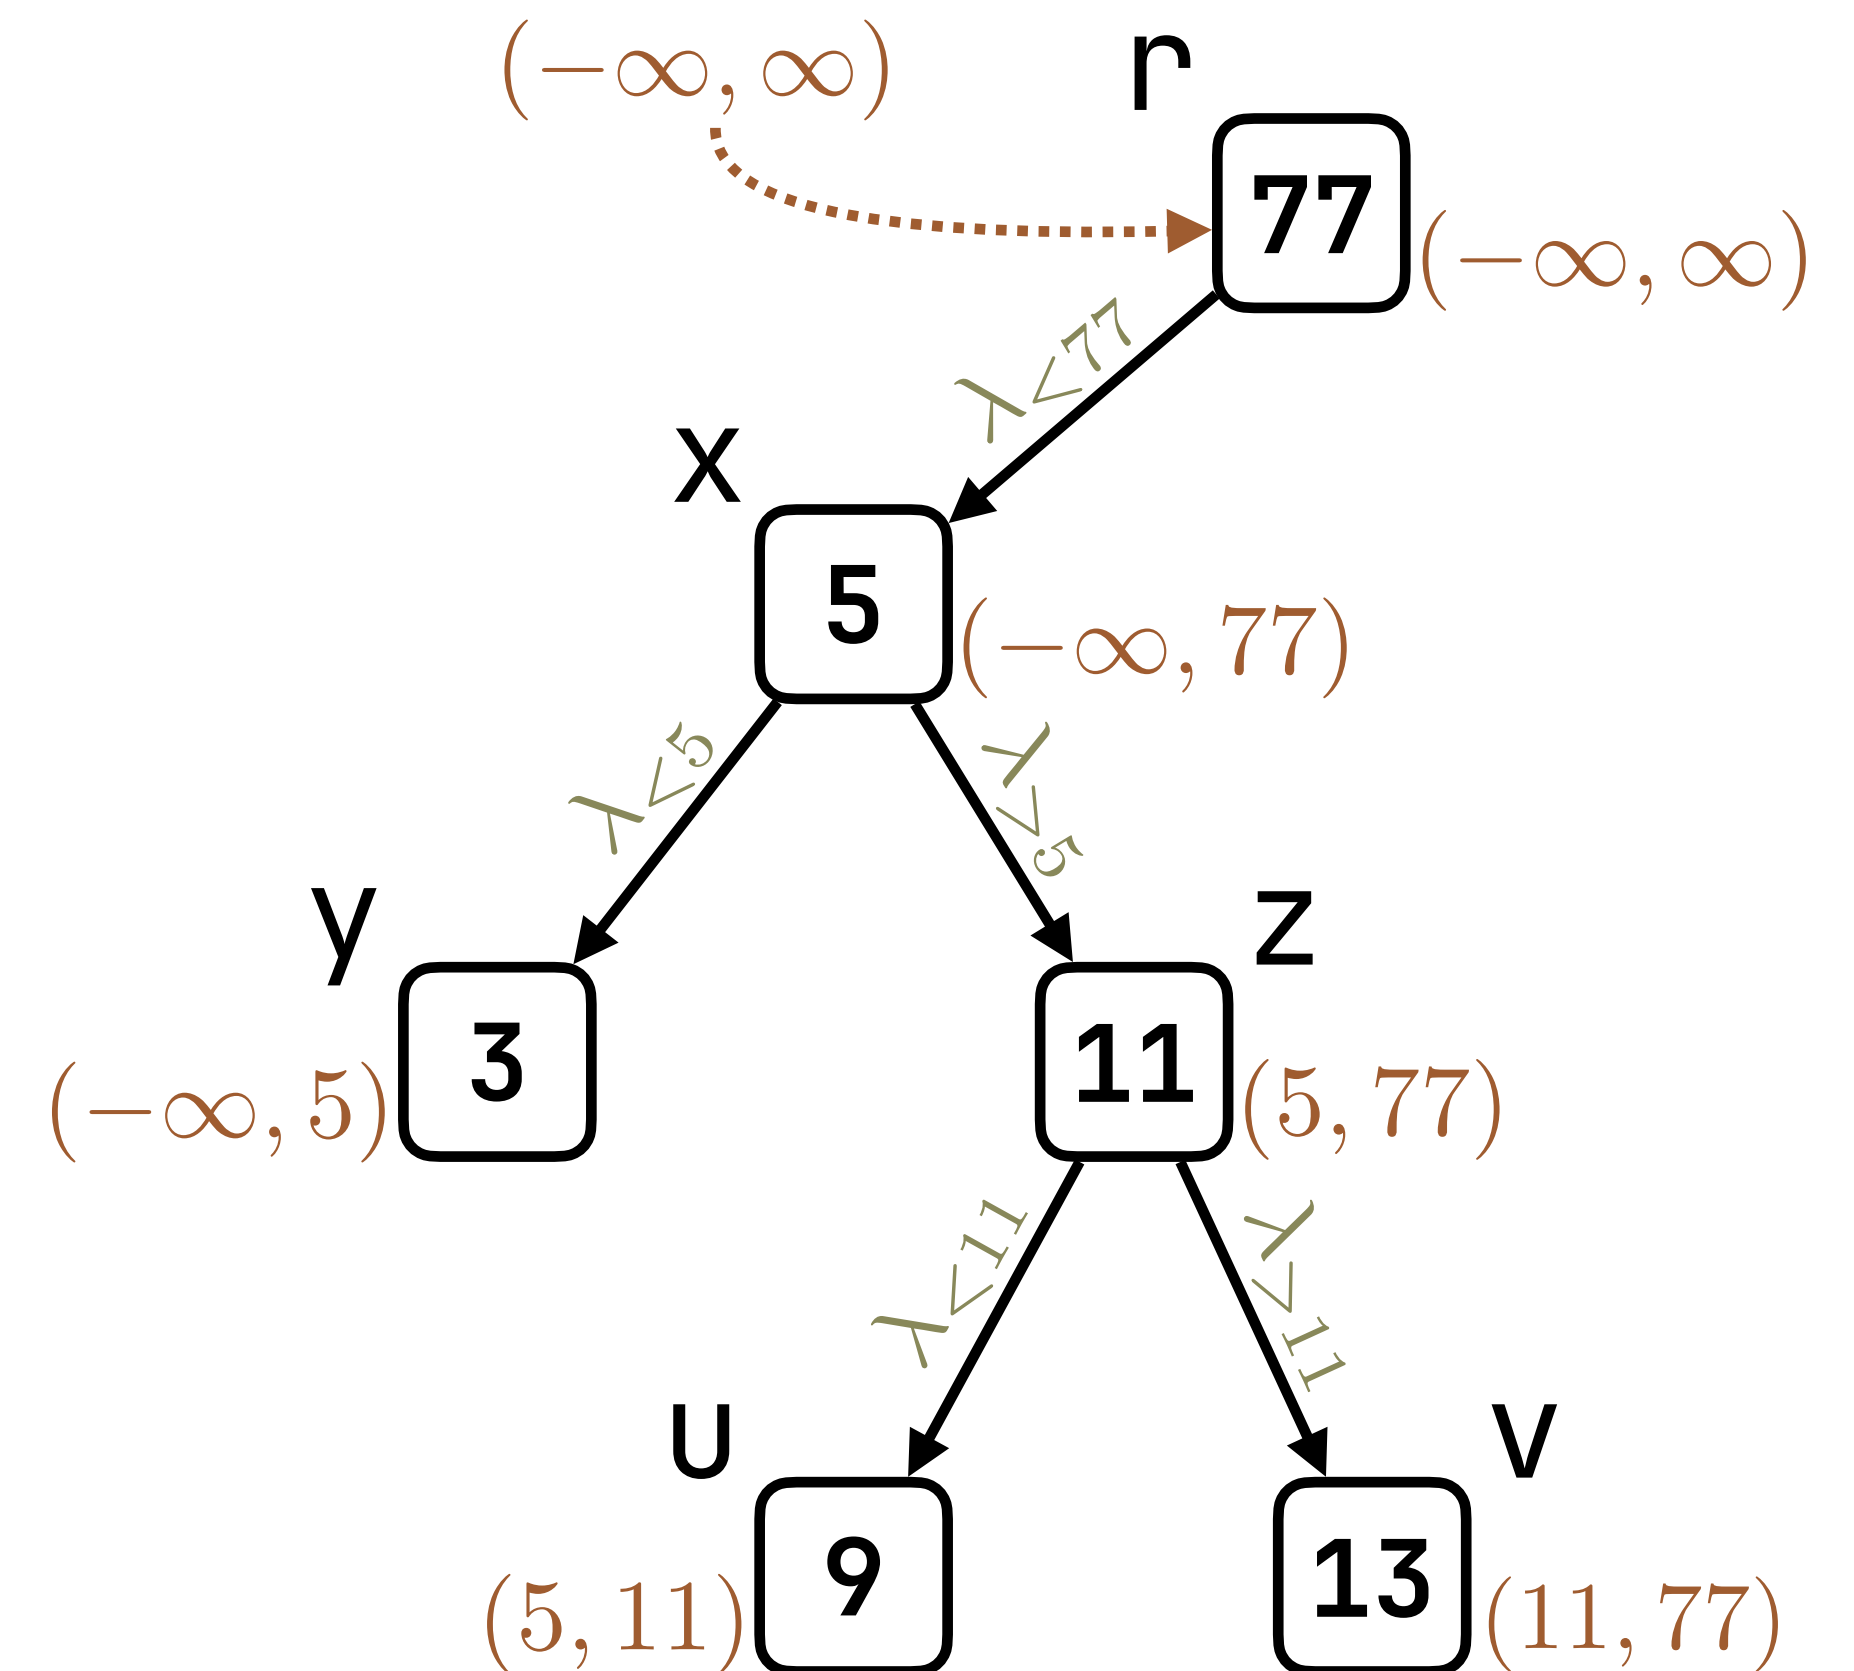
- Flow propagation via
$$\lambda_{<k} := \lambda m. m \cap (-\infty, k)$$
$$\lambda_{>k} := \lambda m. m \cap (k, \infty)$$

- The flow:**

$\text{Flow}(x)$ = "search paths that x lies on"

$\text{Outflow}(x)$ = "searches that continue from x "

$\text{Flow}(x) \setminus \text{Outflow}(x)$ = "responsibility of x "



Flow Framework

flow graphs = heap graphs + data flow values (ghost)

- Flow values from search path monoid $(\mathbb{Z} \cup \{-\infty, \infty\} \cup \emptyset)$

Sufficient for functional correctness.

- Flow propagation

$$\lambda_{<k} := \lambda_{m.r}$$

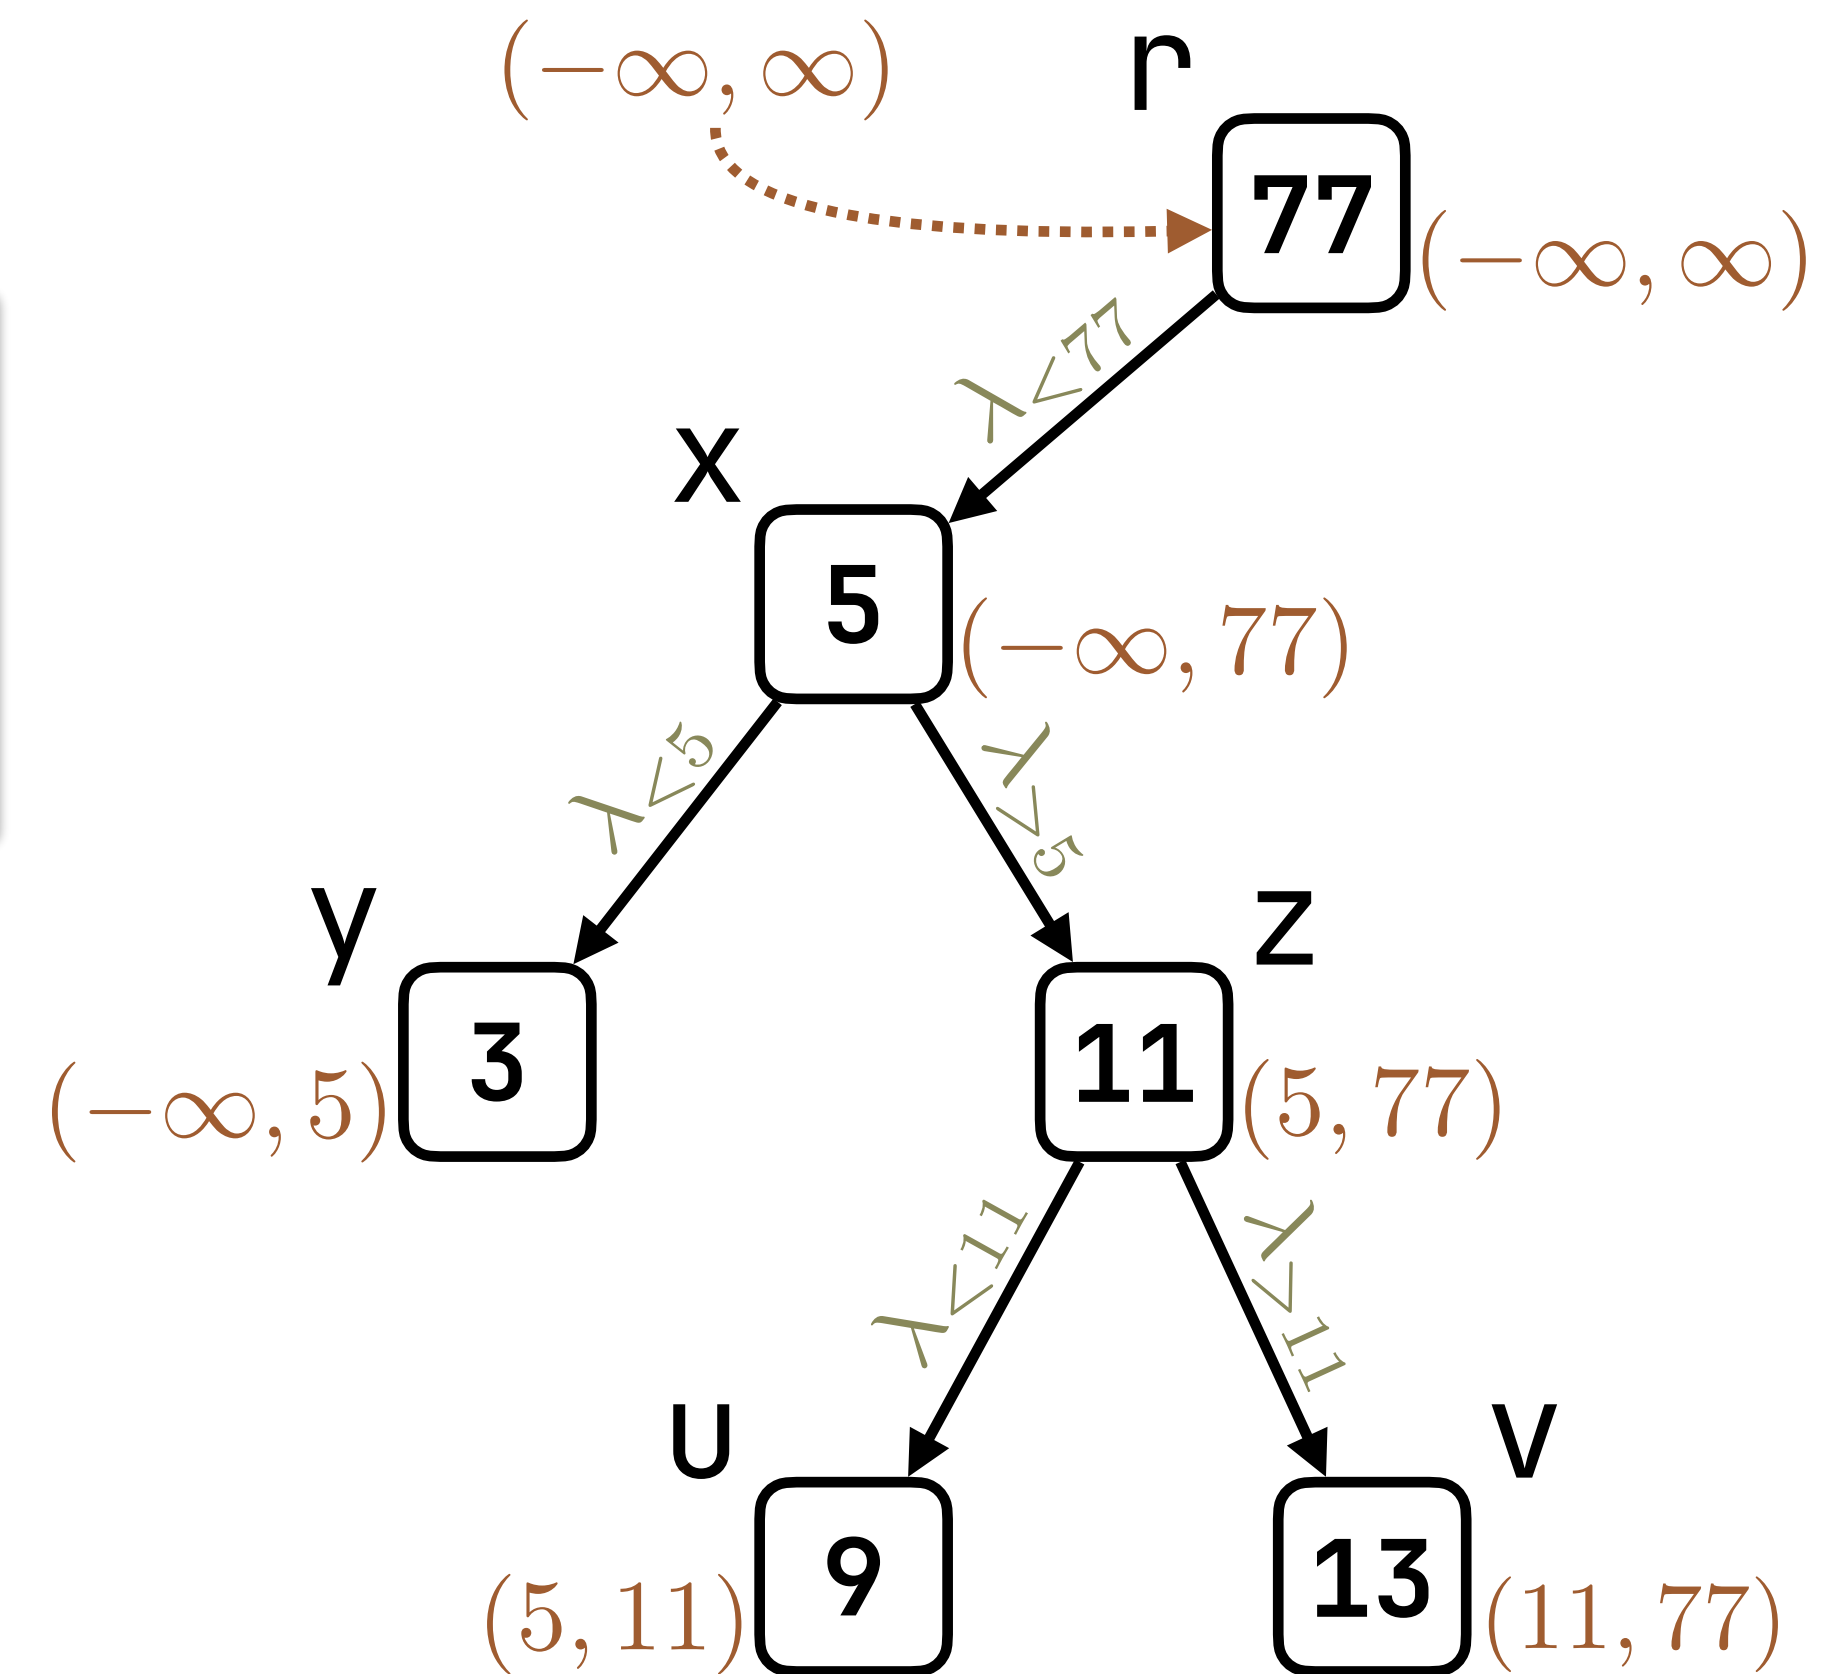
$$\lambda_{>k} := \lambda_{n.\infty}$$

- The flow:**

$\text{Flow}(x)$ = "search paths that x lies on"

$\text{Outflow}(x)$ = "searches that continue from x "

$\text{Flow}(x) \setminus \text{Outflow}(x)$ = "responsibility of x "



Separation

Separation

- Flow graphs form a **separation algebra**:

Separation

$(\Sigma, \star, \text{emp})$ with
 $\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
emp set of units

- Flow graphs form a **separation algebra**:

Separation

$(\Sigma, \star, \text{emp})$ with

$\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
emp set of units

- Flow graphs form a **separation algebra**:
 - ➔ framing: cut the graph, maintain inflow from **each node** outside (assoc. + cancel.).

Separation

$(\Sigma, \star, \text{emp})$ with

$\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
emp set of units

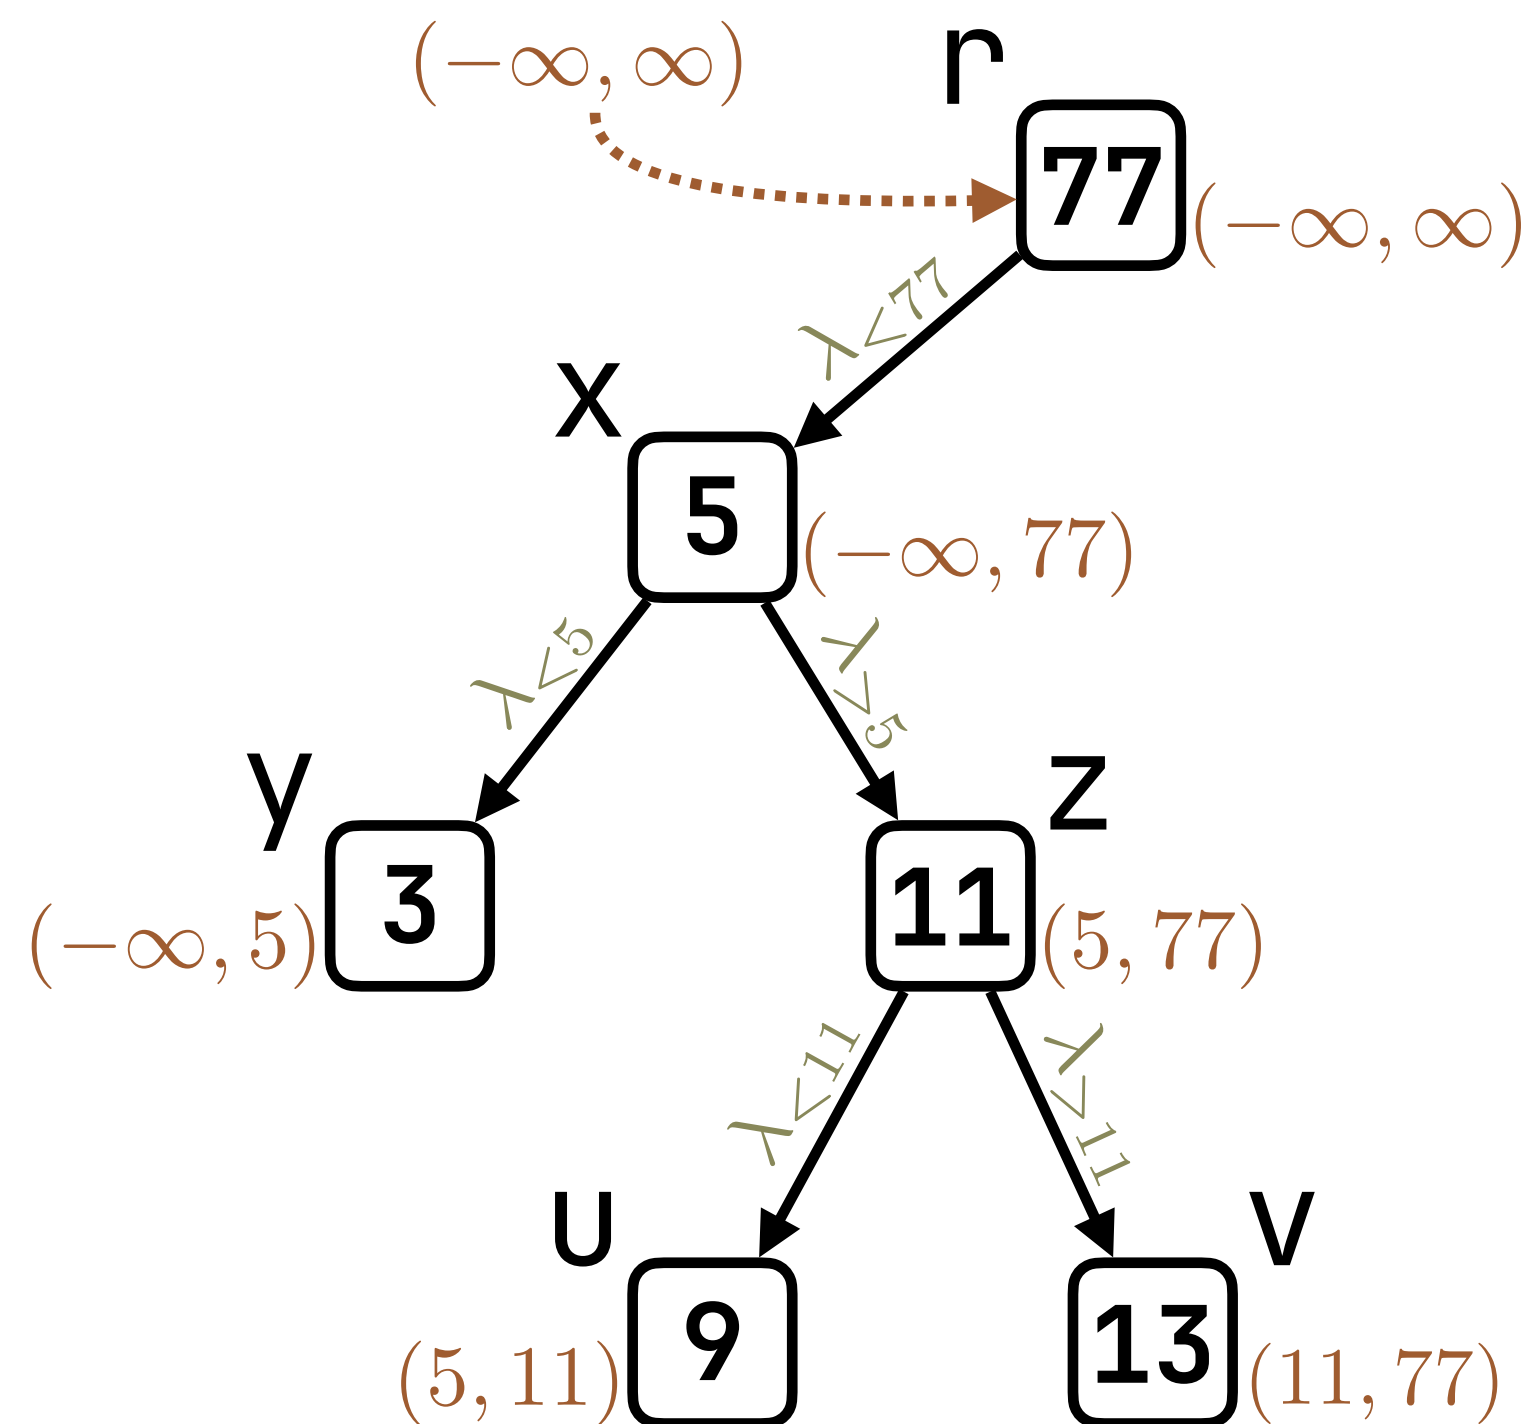
- Flow graphs form a **separation algebra**:
 - ➔ framing: cut the graph, maintain inflow from **each node** outside (assoc. + cancel.).
 - ➔ composition: defined if inflow & outflow match + **no vanishing flow**.

Separation

$(\Sigma, \star, \text{emp})$ with

$\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
emp set of units

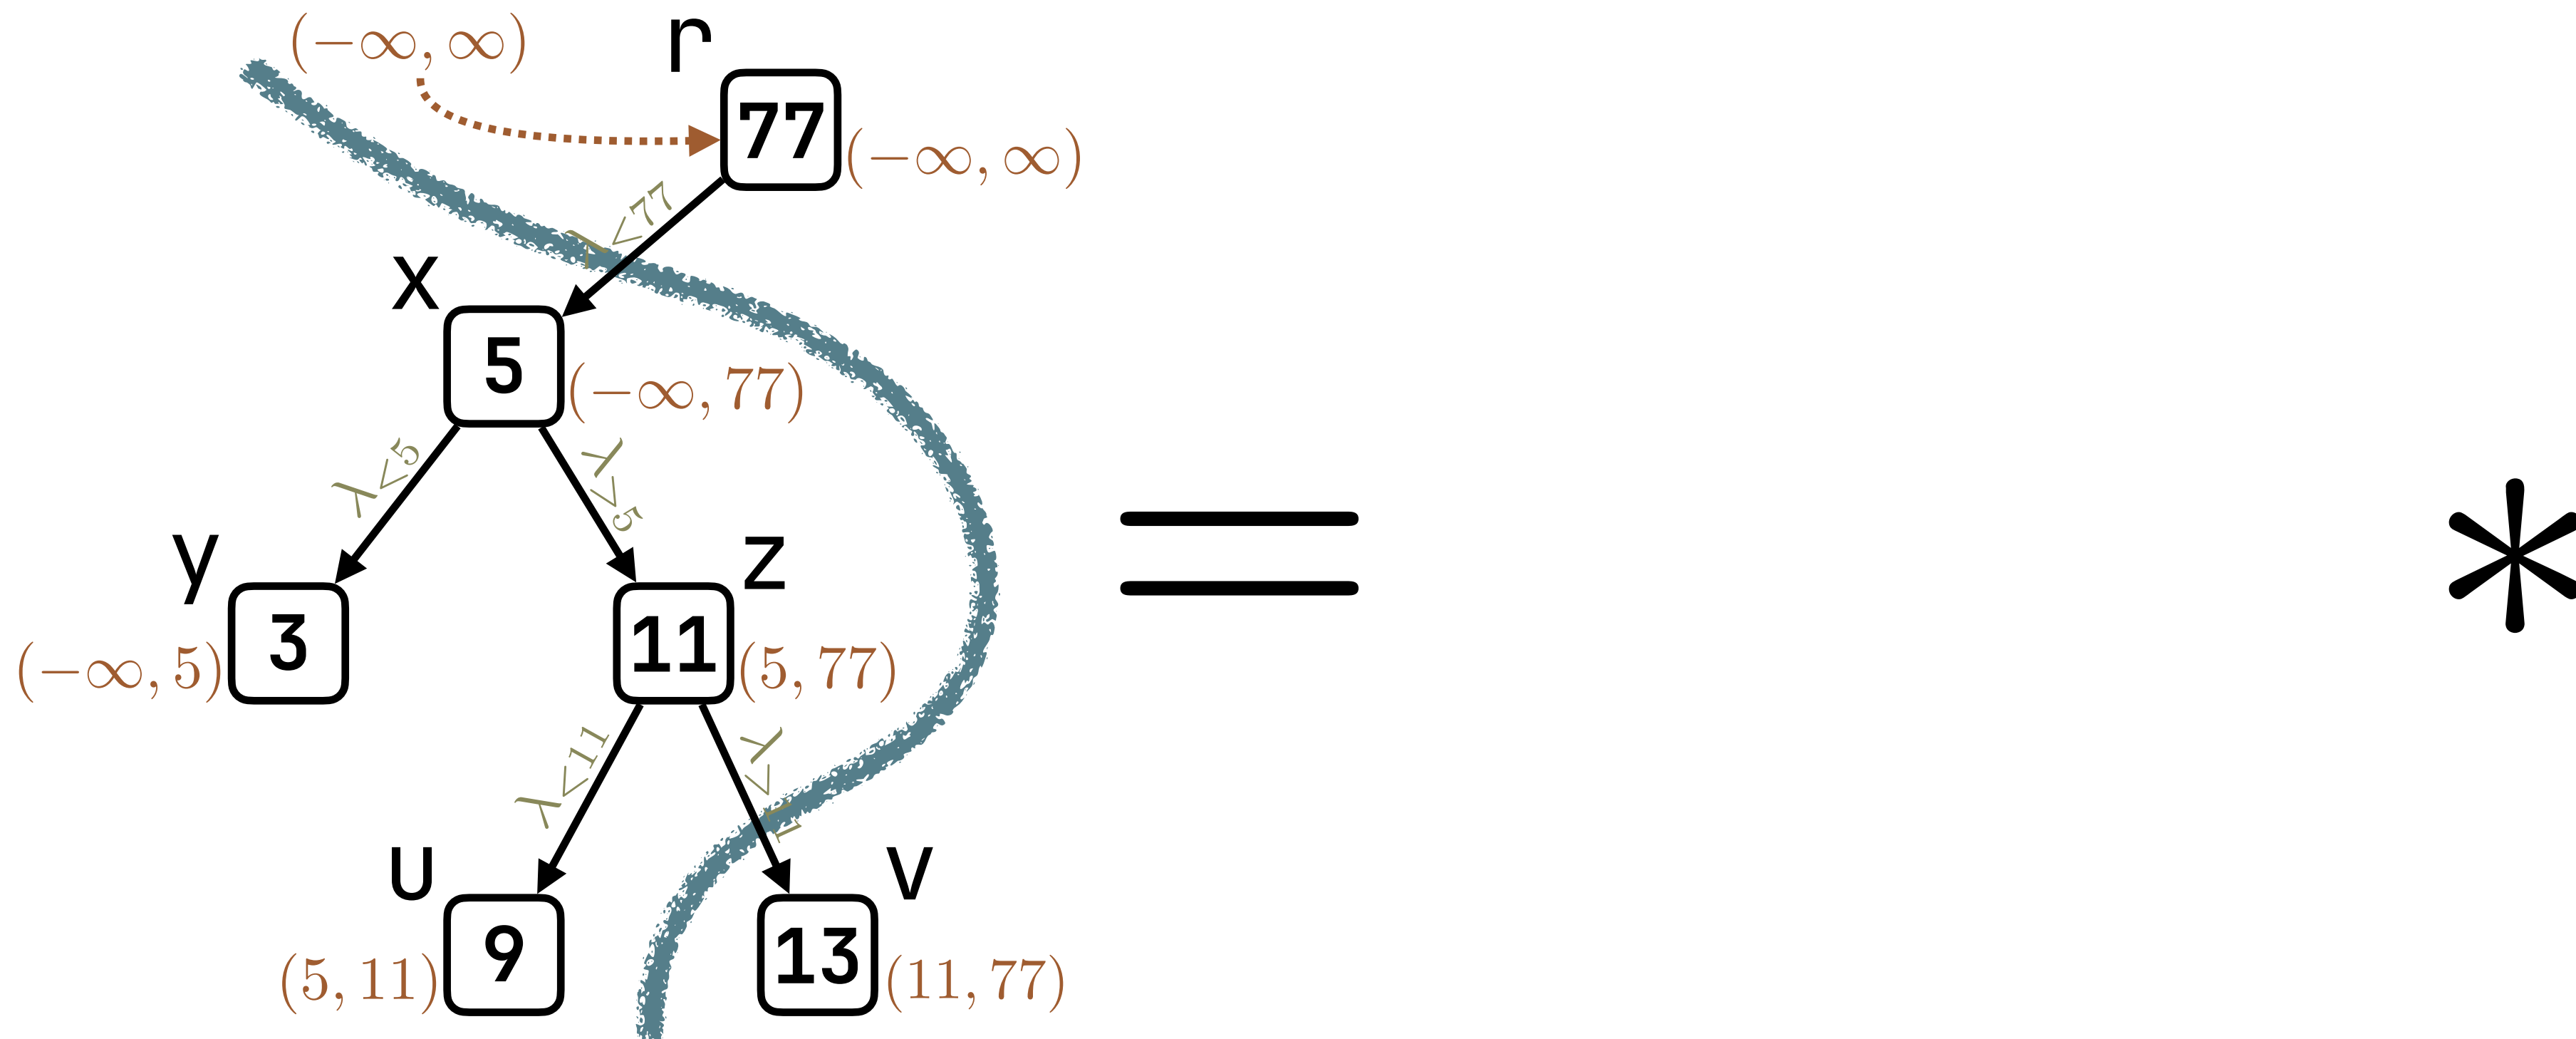
- Flow graphs form a **separation algebra**:
 - framing: cut the graph, maintain inflow from **each node** outside (assoc. + cancel.).
 - composition: defined if inflow & outflow match + **no vanishing flow**.



Separation

$(\Sigma, \star, \text{emp})$ with
 $\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
 emp set of units

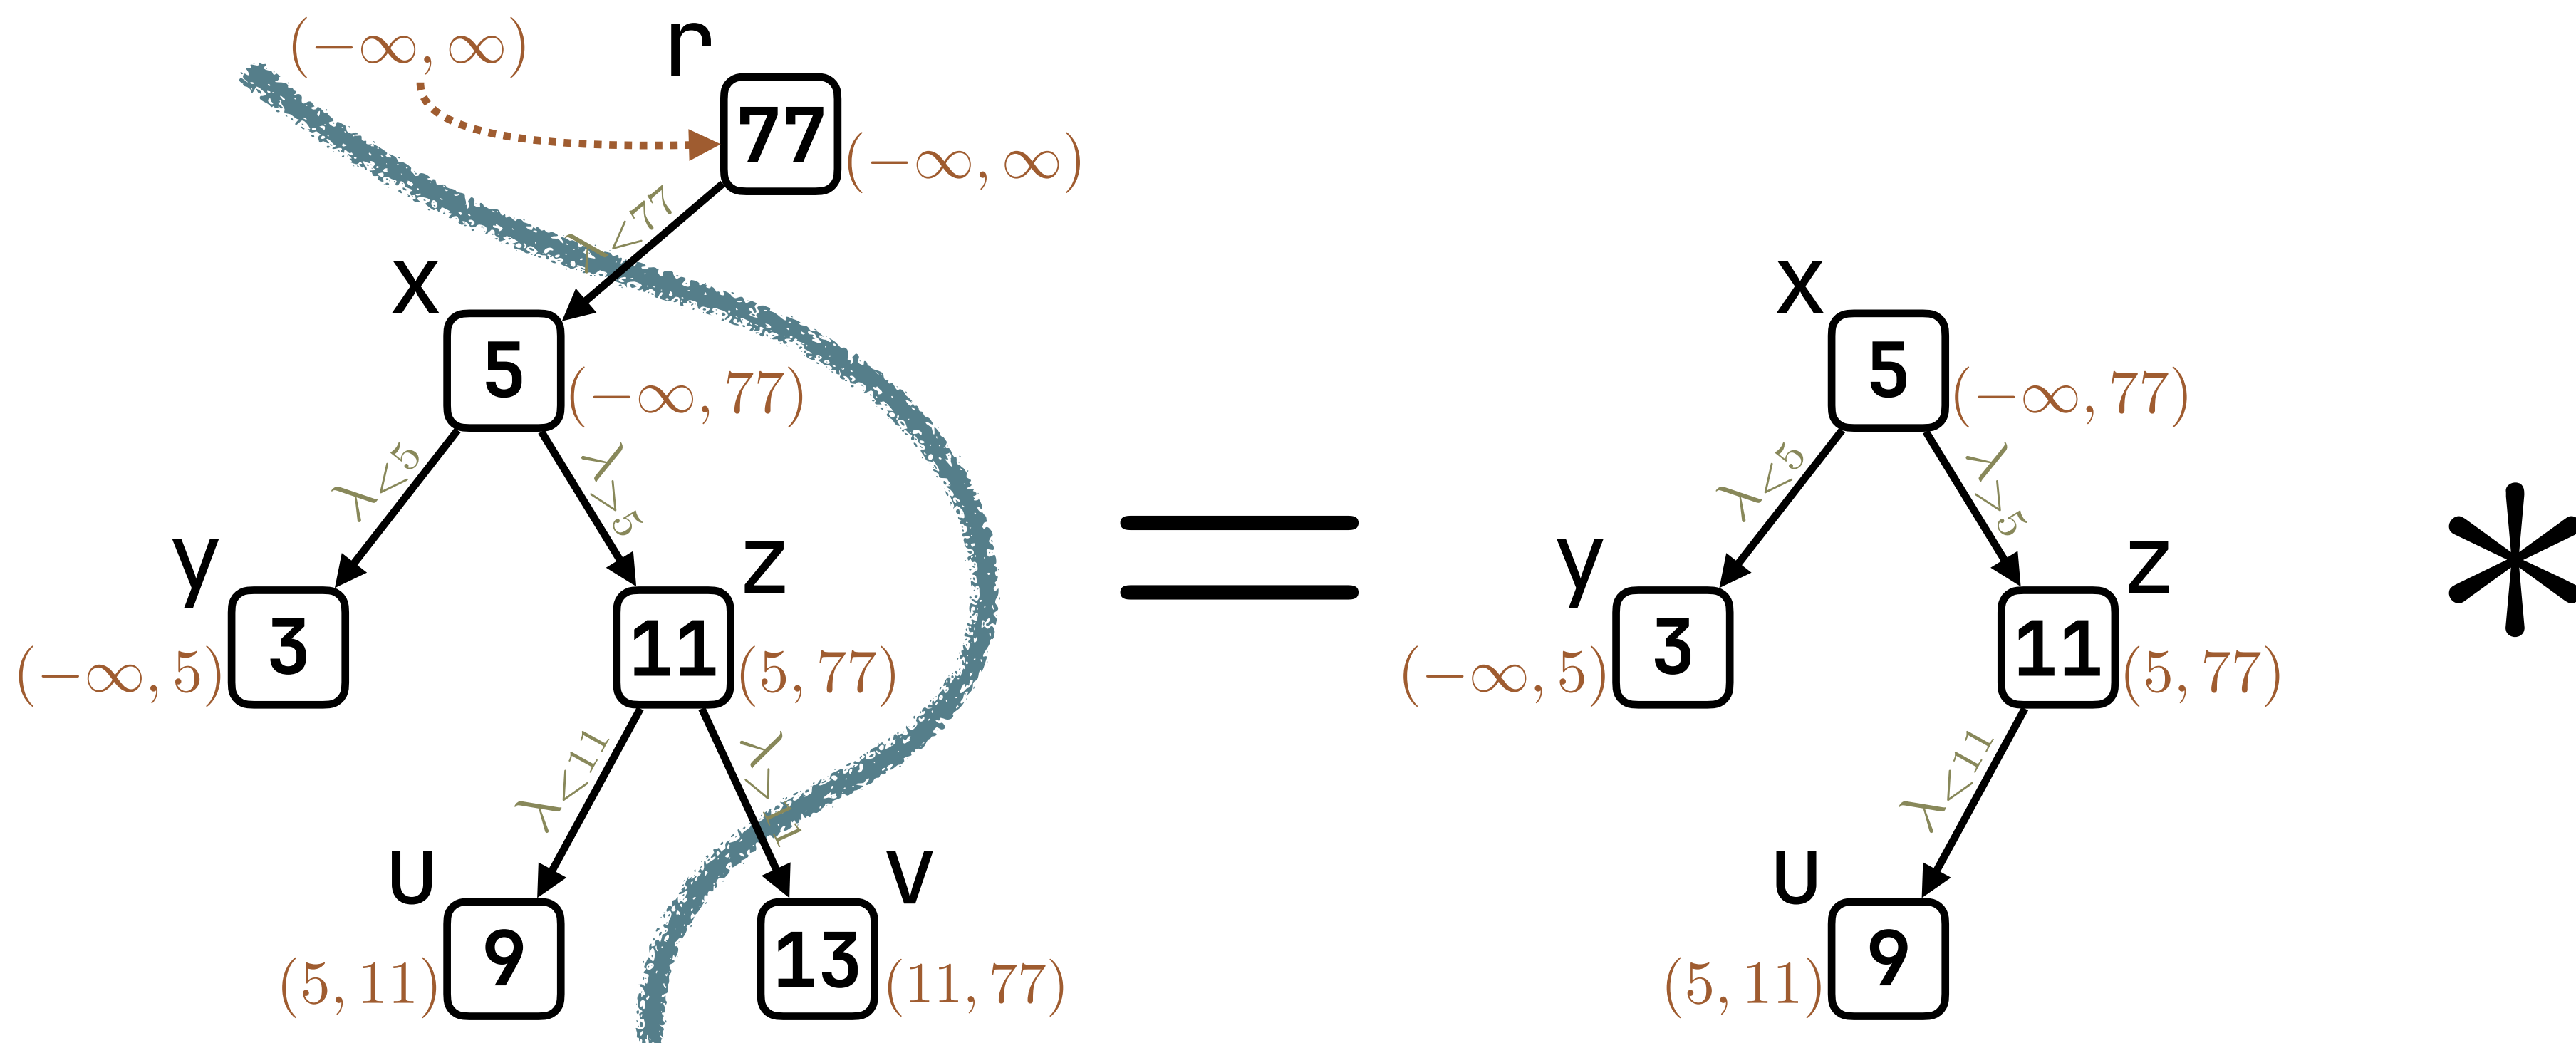
- Flow graphs form a **separation algebra**:
 - framing: cut the graph, maintain inflow from **each node** outside (assoc. + cancel.).
 - composition: defined if inflow & outflow match + **no vanishing flow**.



Separation

$(\Sigma, \star, \text{emp})$ with
 $\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
 emp set of units

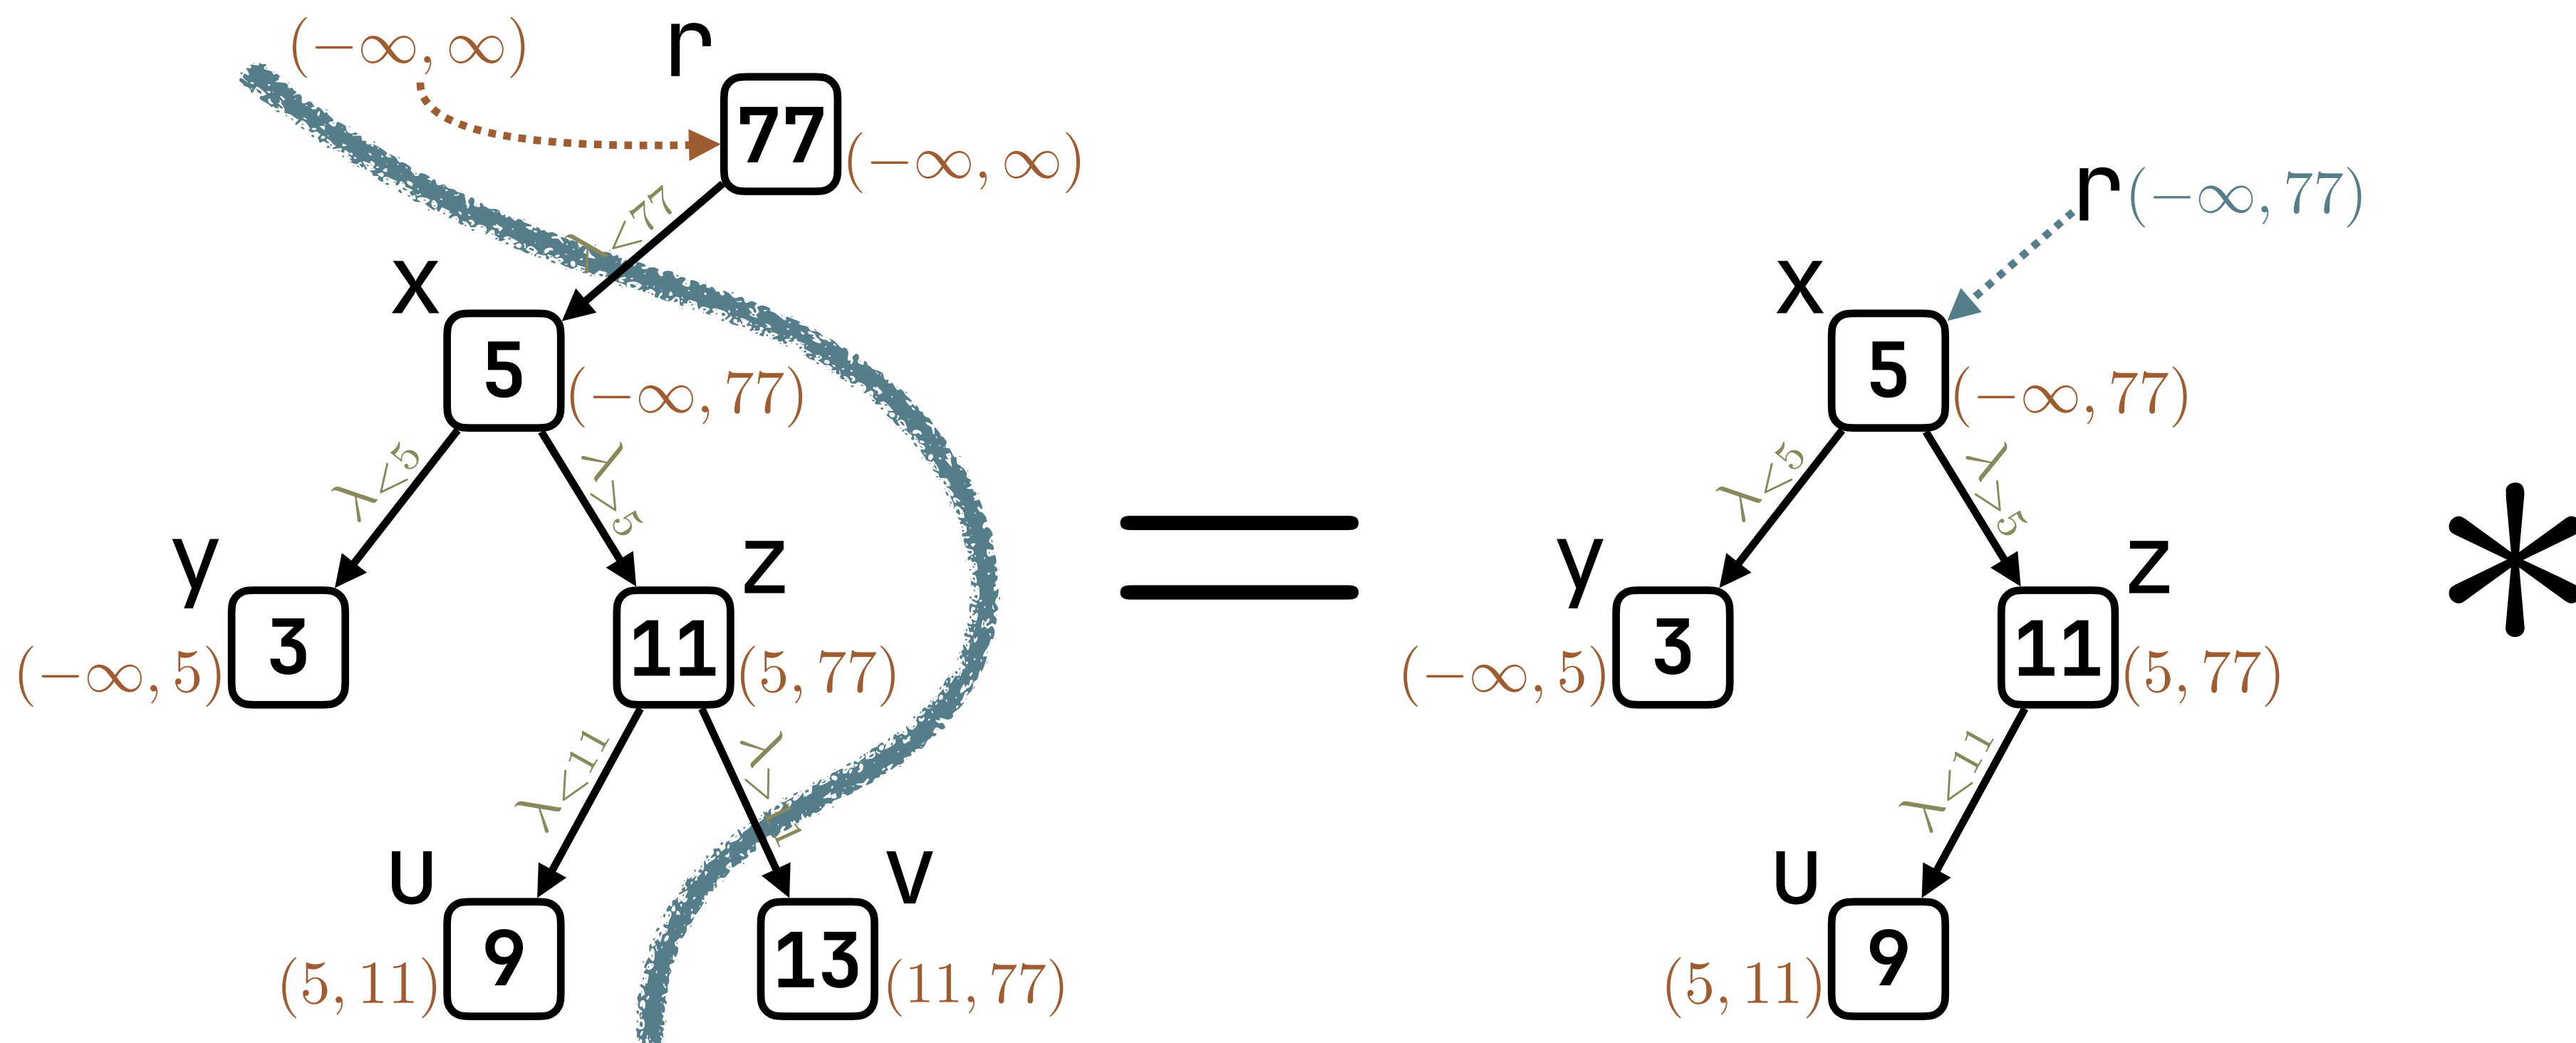
- Flow graphs form a **separation algebra**:
 - framing: cut the graph, maintain inflow from **each node** outside (assoc. + cancel.).
 - composition: defined if inflow & outflow match + **no vanishing flow**.



Separation

$(\Sigma, \star, \text{emp})$ with
 $\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
 emp set of units

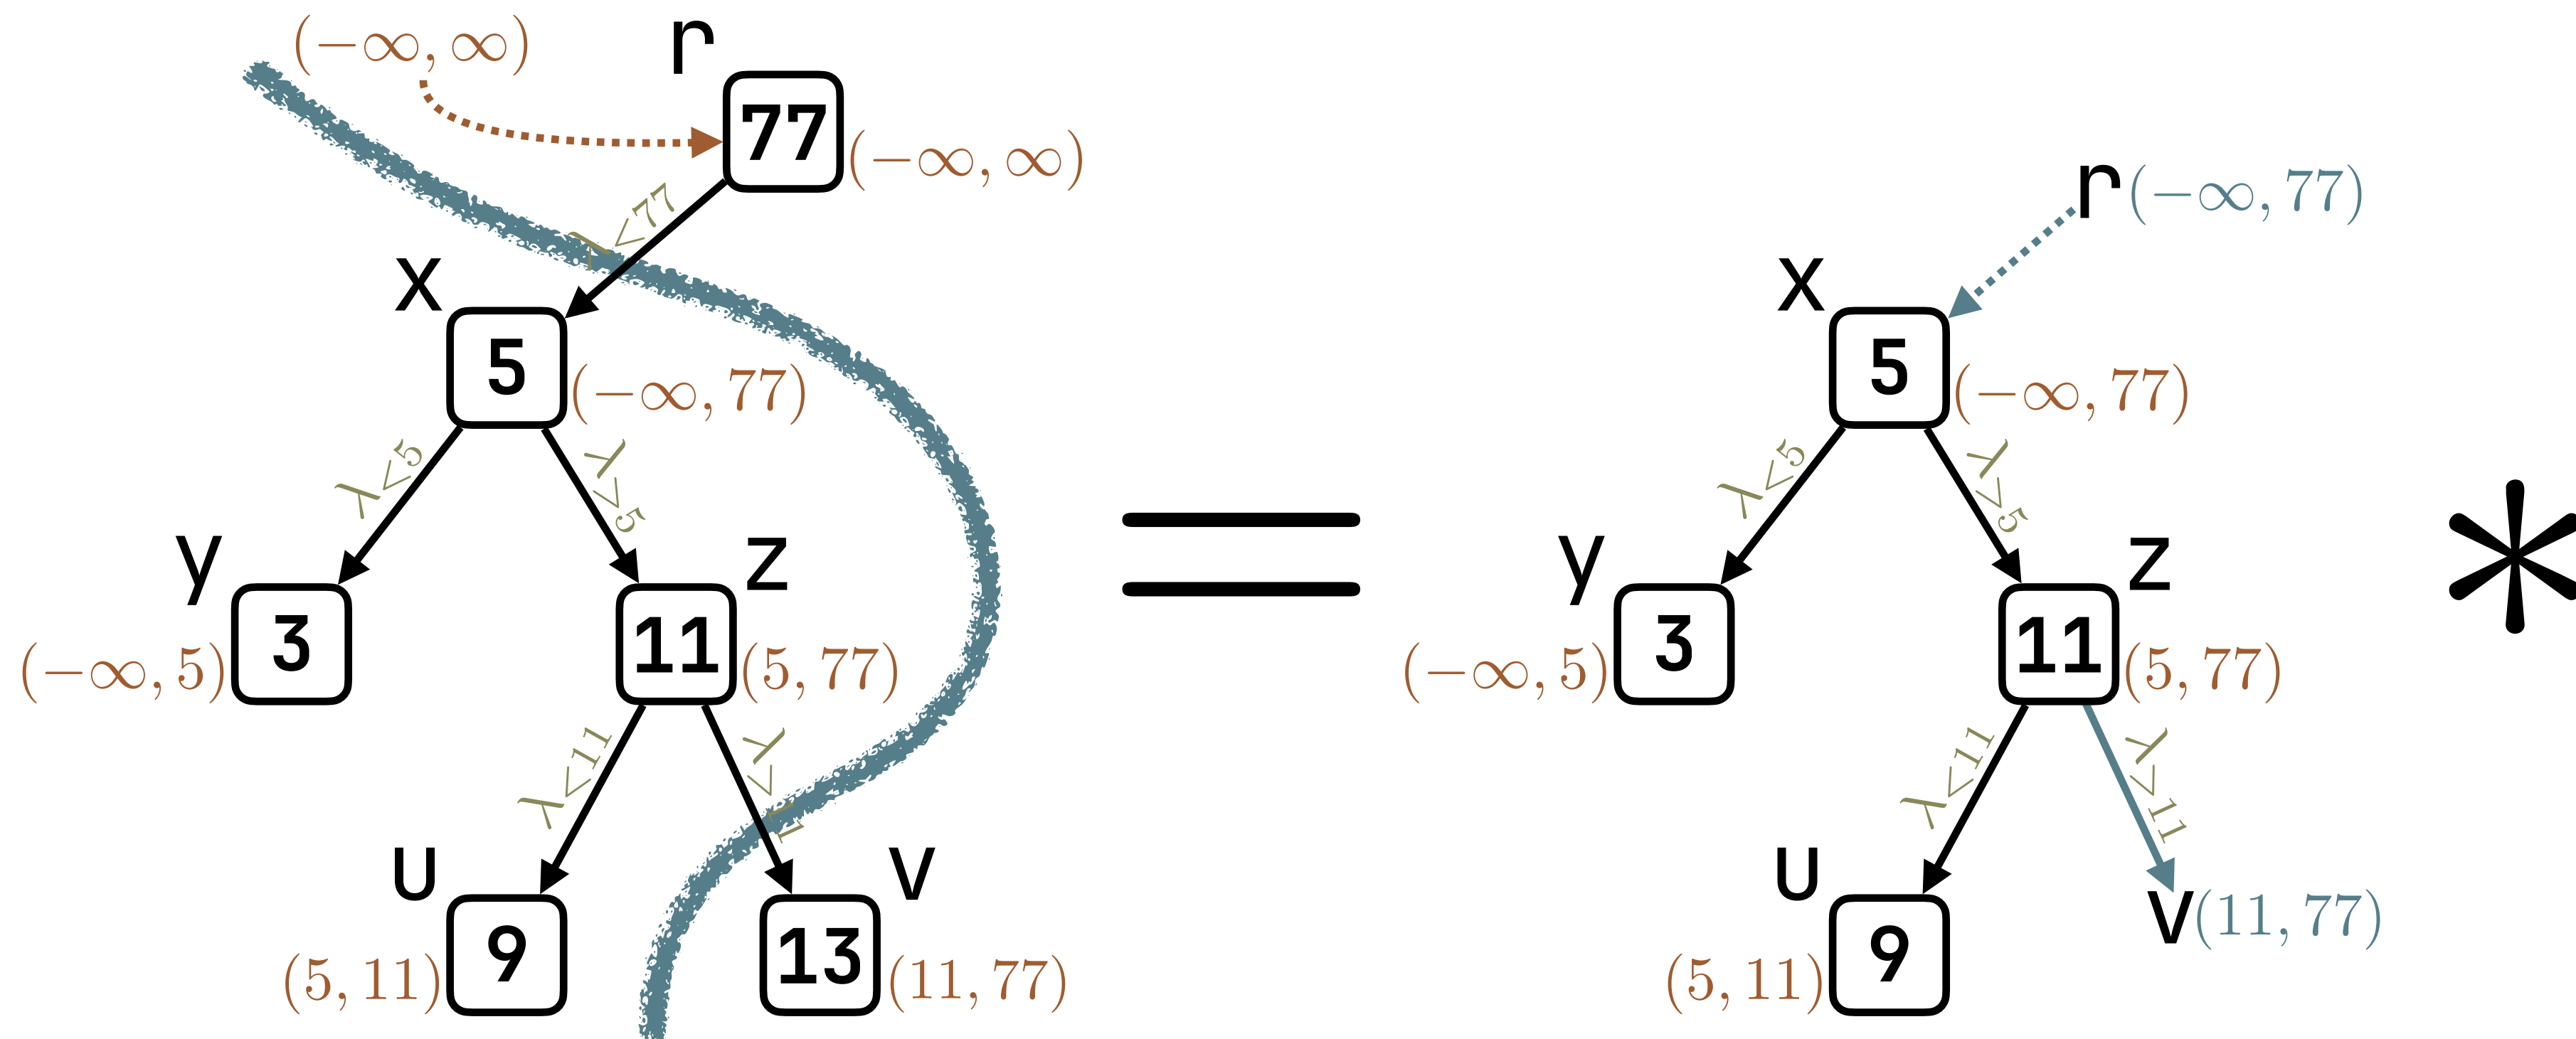
- Flow graphs form a **separation algebra**:
 - framing: cut the graph, maintain inflow from **each node** outside (assoc. + cancel.).
 - composition: defined if inflow & outflow match + **no vanishing flow**.



Separation

$(\Sigma, \star, \text{emp})$ with
 $\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
 emp set of units

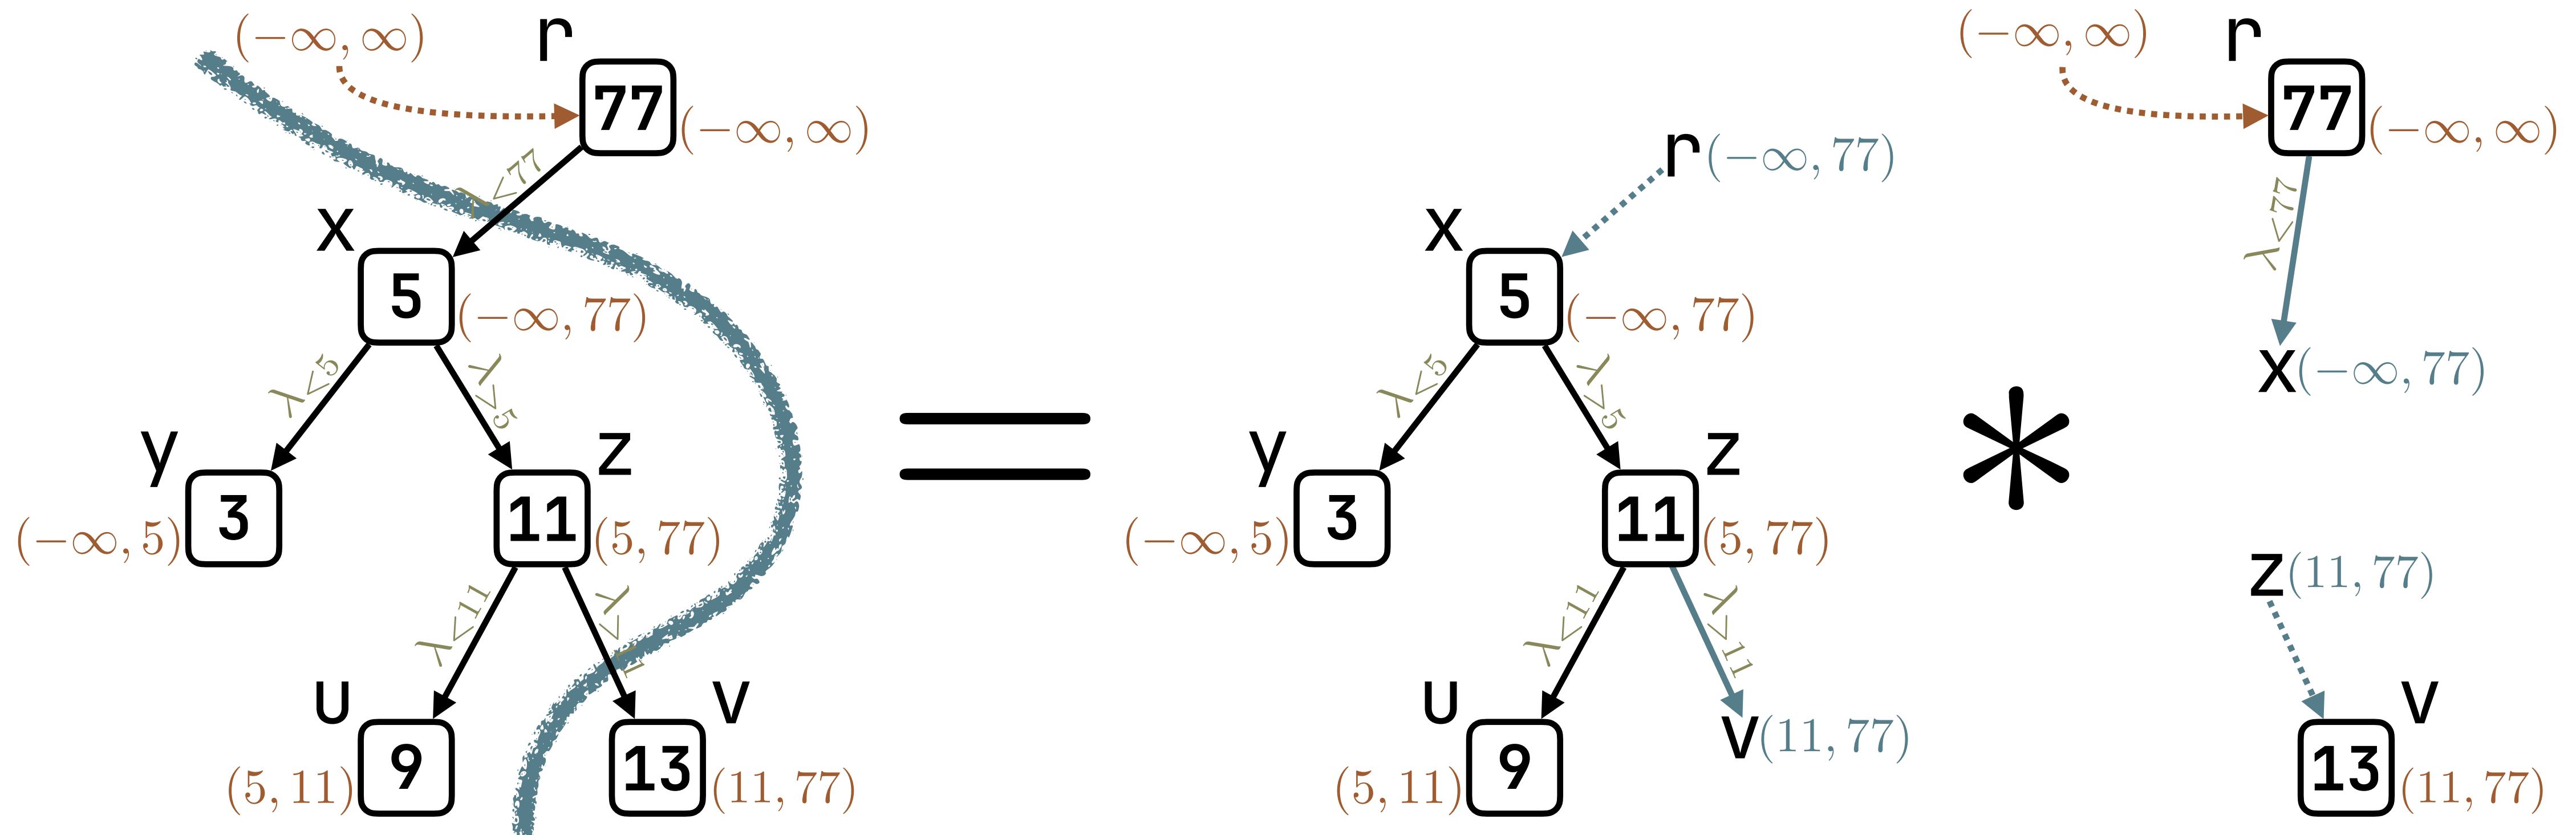
- Flow graphs form a **separation algebra**:
 - framing: cut the graph, maintain inflow from **each node** outside (assoc. + cancel.).
 - composition: defined if inflow & outflow match + **no vanishing flow**.



Separation

$(\Sigma, \star, \text{emp})$ with
 $\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
 emp set of units

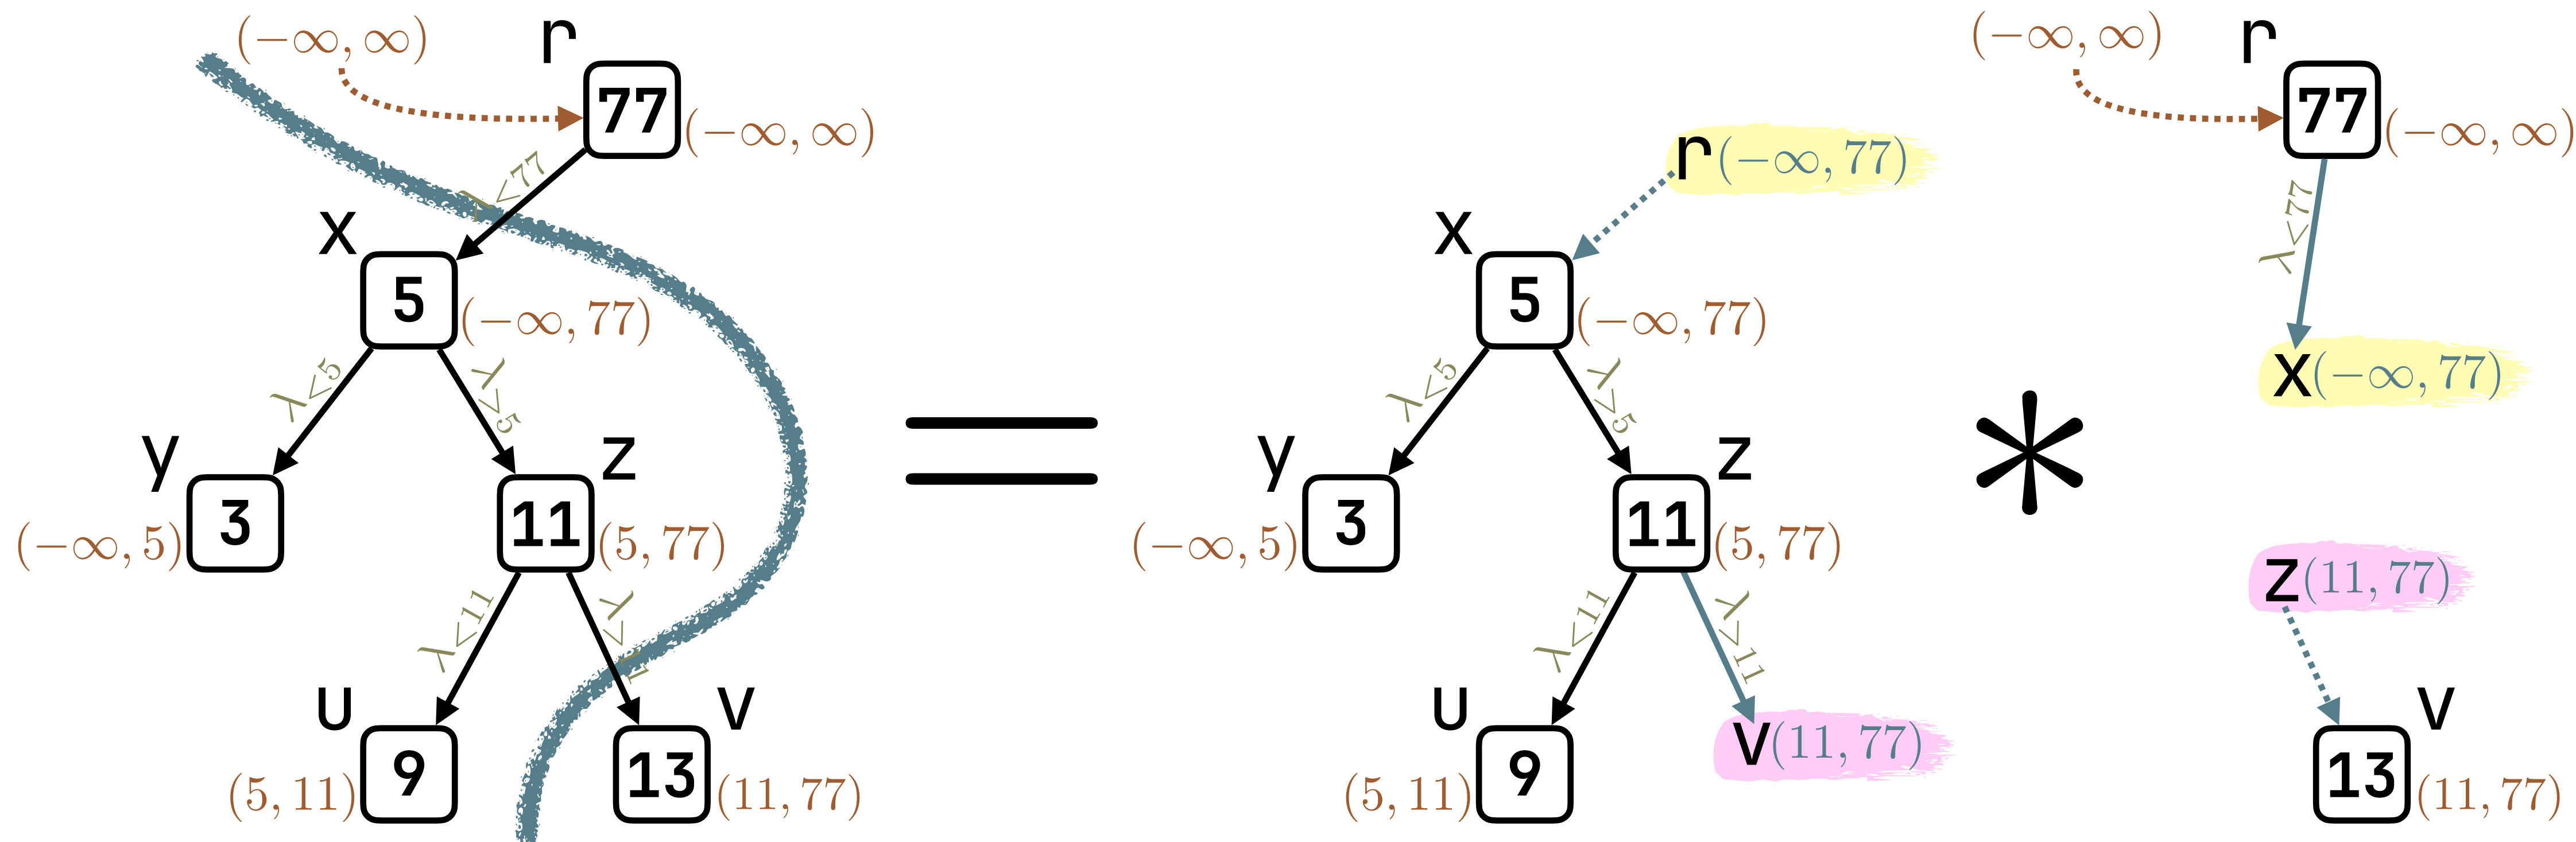
- Flow graphs form a **separation algebra**:
 - framing: cut the graph, maintain inflow from **each node** outside (assoc. + cancel.).
 - composition: defined if inflow & outflow match + **no vanishing flow**.



Separation

$(\Sigma, \star, \text{emp})$ with
 $\star : \Sigma \times \Sigma \rightarrow \Sigma$ associative, commutative, cancellative
 emp set of units

- Flow graphs form a **separation algebra**:
 - framing: cut the graph, maintain inflow from **each node** outside (assoc. + cancel.).
 - composition: defined if inflow & outflow match + **no vanishing flow**.



Frame-Preserving Updates

- Footprint
 - the region  affected by an update

Frame-Preserving Updates

- Footprint
 - the region  affected by an update
 - is **frame-preserving** if, for all F :

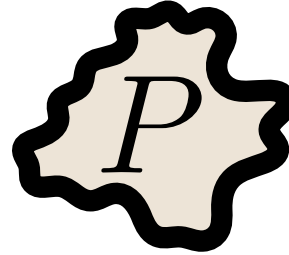
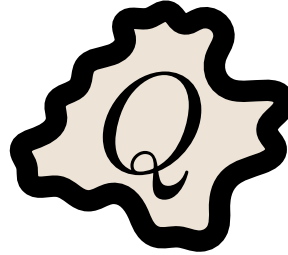
$$\left\{ \begin{array}{c} \text{irregular shape} \\ P \end{array} \right\} \text{ com } \left\{ \begin{array}{c} \text{irregular shape} \\ Q \end{array} \right\} \Rightarrow \left\{ \begin{array}{c} \text{irregular shape} \\ P * F \end{array} \right\} \text{ com } \left\{ \begin{array}{c} \text{irregular shape} \\ Q * F \end{array} \right\}$$

Frame-Preserving Updates

- Footprint
 - the region  affected by an update
 - is **frame-preserving** if, for all F :

$$\left\{ \text{blob } P \right\} \text{ com } \left\{ \text{blob } Q \right\} \Rightarrow \left\{ \text{blob } P * F \right\} \text{ com } \left\{ \text{blob } Q * F \right\}$$

- Theorem:

Update  \rightarrow  frame-preserving

if  and  have the same **inflow-outflow transformer** (up to lfp).

Frame-Preserving Updates

- Footprint
 - the region  affected by an update
 - is **frame-preserving** if, for all F :

$$\left\{ \begin{array}{c} \text{irregular blob} \\ P \end{array} \right\} \text{ com } \left\{ \begin{array}{c} \text{irregular blob} \\ Q \end{array} \right\} \Rightarrow \left\{ \begin{array}{c} \text{irregular blob} \\ P * F \end{array} \right\} \text{ com } \left\{ \begin{array}{c} \text{irregular blob} \\ Q * F \end{array} \right\}$$

Bekić's lemma

frame-preserving

- Theorem:
Update
if and have the same **inflow-outflow transformer** (up to lfp).

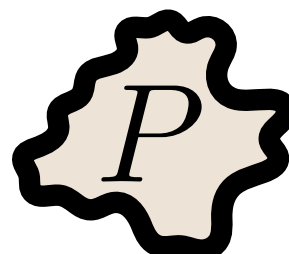
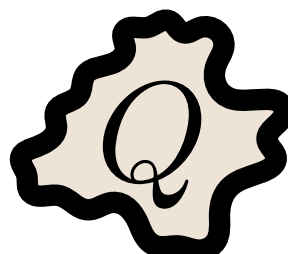
Frame-Preserving Updates

- Footprint
 - the region  affected by an update
 - is **frame-preserving** if, for all F :

$$\left\{ \begin{array}{c} \text{irregular shape} \\ P \end{array} \right\} \text{ com } \left\{ \begin{array}{c} \text{irregular shape} \\ Q \end{array} \right\} \Rightarrow \left\{ \begin{array}{c} \text{irregular shape} \\ P * F \end{array} \right\} \text{ com } \left\{ \begin{array}{c} \text{irregular shape} \\ Q * F \end{array} \right\}$$

Bekić's lemma

- Theorem:

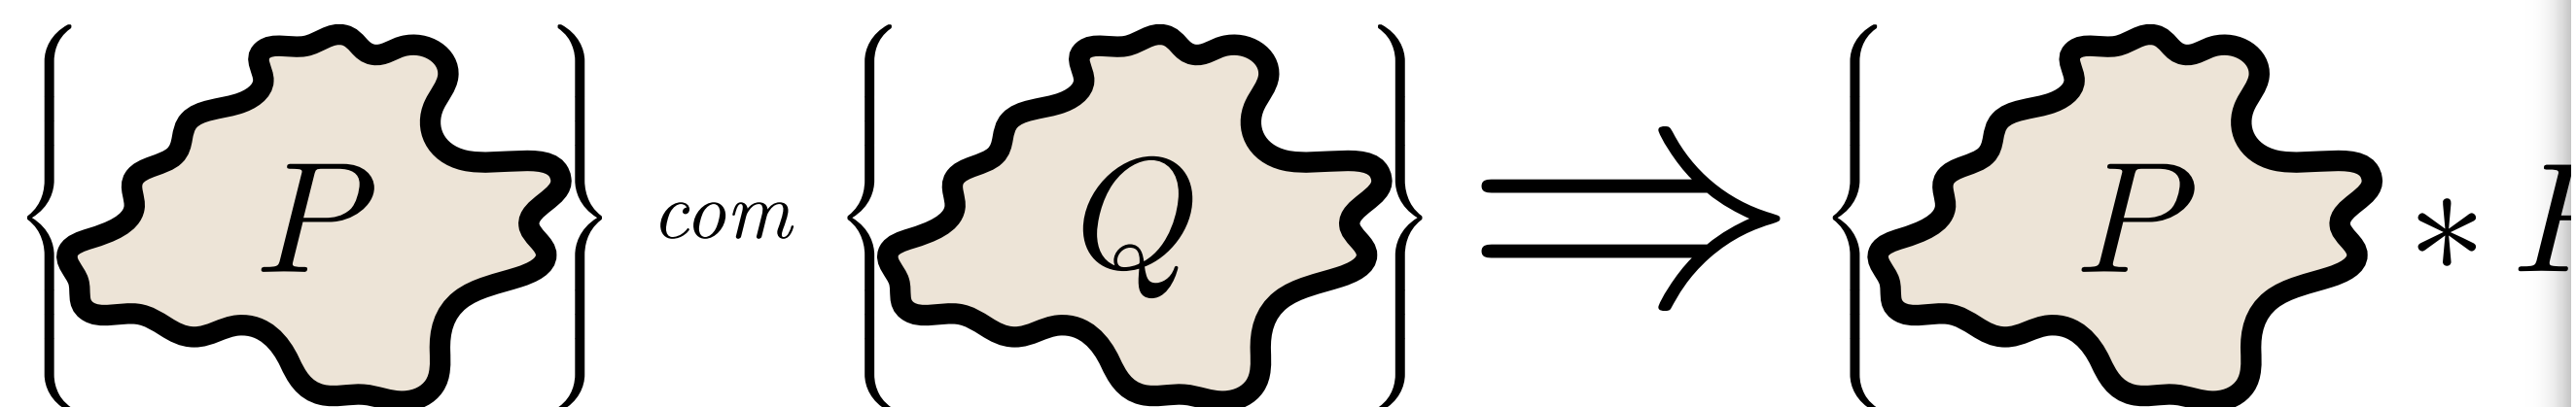
Update  \rightarrow  frame-preserving

if  and  have the same **inflow-outflow transformer** (up to lfp).

Frame-Preserving Updates

- Footprint

- the region  affected by an update
- is **frame-preserving** if, for all F :



Bekić's lemma

- Theorem:

Update

frame-preserving

if

and

have the same **inflow-outflow**

Make flows small again: revisiting the flow framework

Roland Meyer¹, Thomas Wies², and Sebastian Wolff²

¹ TU Braunschweig, Braunschweig, Germany, meyer@tu-bs.de

² New York University, New York, USA, {wies,sebastian.wolff}@cs.nyu.edu

Abstract We present a new flow framework for separation logic reasoning about programs that manipulate general graphs. The framework overcomes problems in earlier developments: it is based on standard fixed point theory, guarantees least flows, rules out vanishing flows, and has an easy to understand notion of footprint as needed for soundness of the frame rule. In addition, we present algorithms for automating the frame rule, which we evaluate on graph updates extracted from linearizability proofs for concurrent data structures. The evaluation demonstrates that our algorithms help to automate key aspects of these proofs that have previously relied on user guidance or heuristics.

Keywords: Separation Logic · Graph Algorithms · Frame Inference.

1 Introduction

The flow framework [23, 24] is an abstraction mechanism based on separation logic [5, 32, 40] that enables reasoning about global inductive invariants of general graphs in a local manner. The framework has proved useful to verify intricate algorithms that are difficult to handle by other techniques, such as the Priority Inheritance Protocol, object-oriented design patterns, and complex concurrent data structures [22, 24, 27, 34]. However, these efforts have also exposed some rough corners in the underlying meta theory that either limit expressivity or automation. In this paper, we propose a new meta theory for the flow framework that aims to strike a balance between these conflicting requirements. In addition, we present algorithms that aid proof automation.

Background. The central notion of the flow framework is that of a *flow*. Given a commutative monoid $(\mathbb{M}, +, 0)$ (e.g. natural numbers with addition), and a graph with nodes X and an *edge function* $E: X \times X \rightarrow \mathbb{M} \rightarrow \mathbb{M}$, a flow is a function $f_l: X \rightarrow \mathbb{M}$ that satisfies the *flow equation*:

$$\forall x \in X. f_l(x) = in_x + \sum_{y \in X} E_{(y,x)}(f_l(y)) .$$

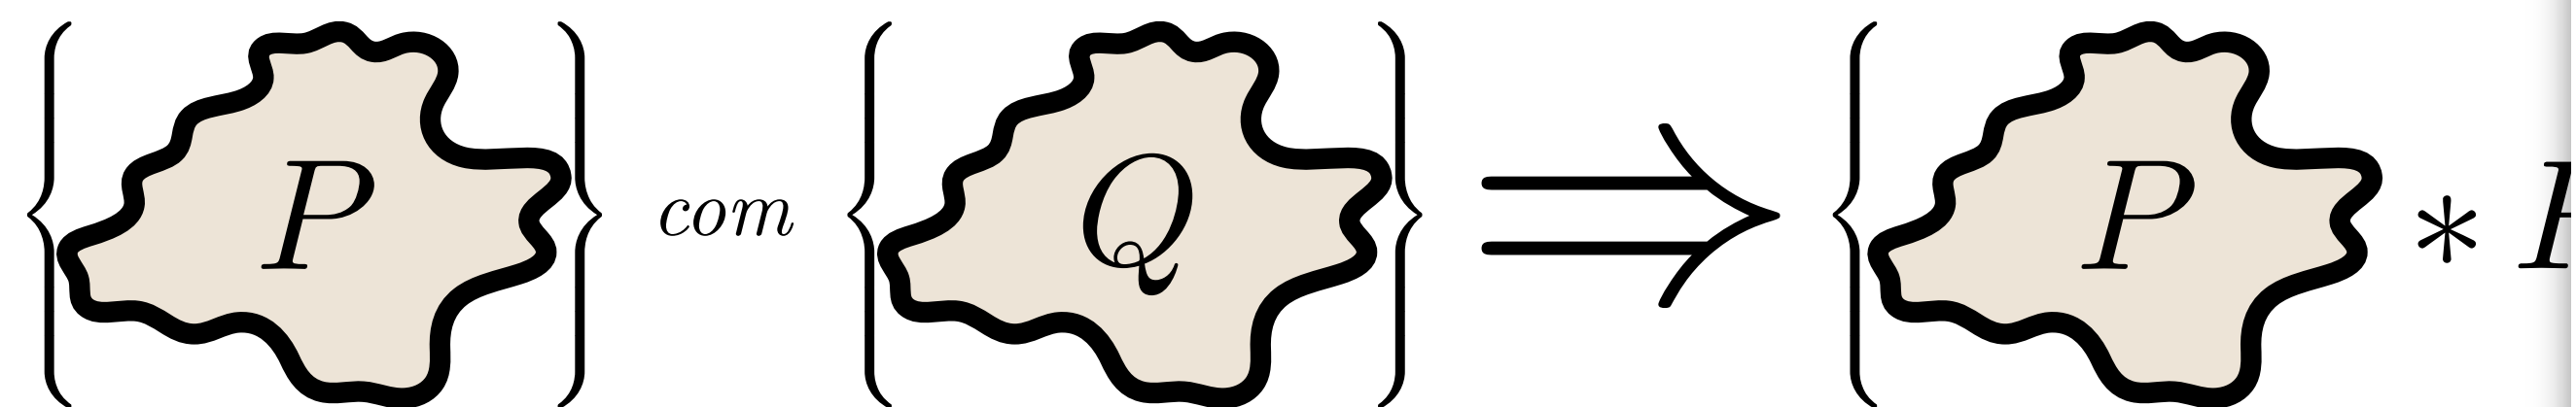
That is, f_l is a fixed point of the function that assigns every node x an initial value $in_x \in \mathbb{M}$, its *inflow*, and then propagates these values through the graph according to the edge function. This is akin to a forward data flow analysis where the monoid operation $+$ is used as the join. By choosing an appropriate flow monoid, inflow, and edge function, one can express inductive properties of graphs (reachability, sortedness, etc.) in terms of conditions that refer only to each node's flow value $f_l(x)$.

A graph endowed with an inflow and associated flow is a *flow graph*. An example flow graph h is shown on the right-hand side of Fig. 1a. Here, the flow value $f_l(w)$ for

Frame-Preserving Updates

- Footprint

- the region  affected by an update
- is **frame-preserving** if, for all F :



Bekić's lemma

- Theorem:



Make flows small again: revisiting the flow framework

Roland Meyer¹, Thomas Wies², and Sebastian Wolff²

¹ TU Braunschweig, Braunschweig, Germany, meyer@tu-bs.de

² New York University, New York, USA, {wies,sebastian.wolff}@cs.nyu.edu

Abstract We present a new flow framework for separation logic reasoning about programs that manipulate general graphs. The framework overcomes problems in earlier developments: it is based on standard fixed point theory, guarantees least flows, rules out vanishing flows, and has an easy to understand notion of footprint as needed for soundness of the frame rule. In addition, we present algorithms for automating the frame rule, which we evaluate on graph updates extracted from linearizability proofs for concurrent data structures. The evaluation demonstrates that our algorithms help to automate key aspects of these proofs that have previously relied on user guidance or heuristics.

Keywords: Separation Logic · Graph Algorithms · Frame Inference.

1 Introduction

The flow framework [23, 24] is an abstraction mechanism based on separation logic [5, 32, 40] that enables reasoning about global inductive invariants of general graphs in a local manner. The framework has proved useful to verify intricate algorithms that are difficult to handle by other techniques, such as the Priority Inheritance Protocol, object-oriented design patterns, and complex concurrent data structures [22, 24, 27, 34]. However, these efforts have also exposed some rough corners in the underlying meta theory that either limit expressivity or automation. In this paper, we propose a new meta theory for the flow framework that aims to strike a balance between these conflicting requirements. In addition, we present algorithms that aid proof automation.

Background. The central notion of the flow framework is that of a *flow*. Given a commutative monoid $(\mathbb{M}, +, 0)$ (e.g. natural numbers with addition), and a graph with nodes X and an *edge function* $E: X \times X \rightarrow \mathbb{M} \rightarrow \mathbb{M}$, a flow is a function $fl: X \rightarrow \mathbb{M}$ that satisfies the *flow equation*:

$$\forall x \in X. fl(x) = in_x + \sum_{y \in X} E_{(y,x)}(fl(y)) .$$

That is, fl is a fixed point of the function that assigns every node x an initial value $in_x \in \mathbb{M}$, its *inflow*, and then propagates these values through the graph according to the edge function. This is akin to a forward data flow analysis where the monoid operation $+$ is used as the join. By choosing an appropriate flow monoid, inflow, and edge function, one can express inductive properties of graphs (reachability, sortedness, etc.) in terms of conditions that refer only to each node's flow value $fl(x)$.

A graph endowed with an inflow and associated flow is a *flow graph*. An example flow graph h is shown on the right-hand side of Fig. 1a. Here, the flow value $fl(w)$ for

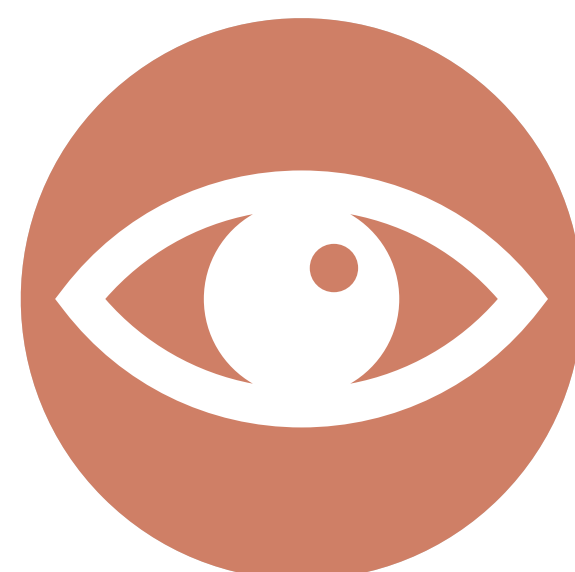
Contributions



Flow Framework

Reasoning about heap graphs without recursive predicates.

[POPL'18, ESOP'20, TACAS'23]



Hindsight

Non-fixed linearization points without prophecies.

[OOPSLA'22, PLDI'23]



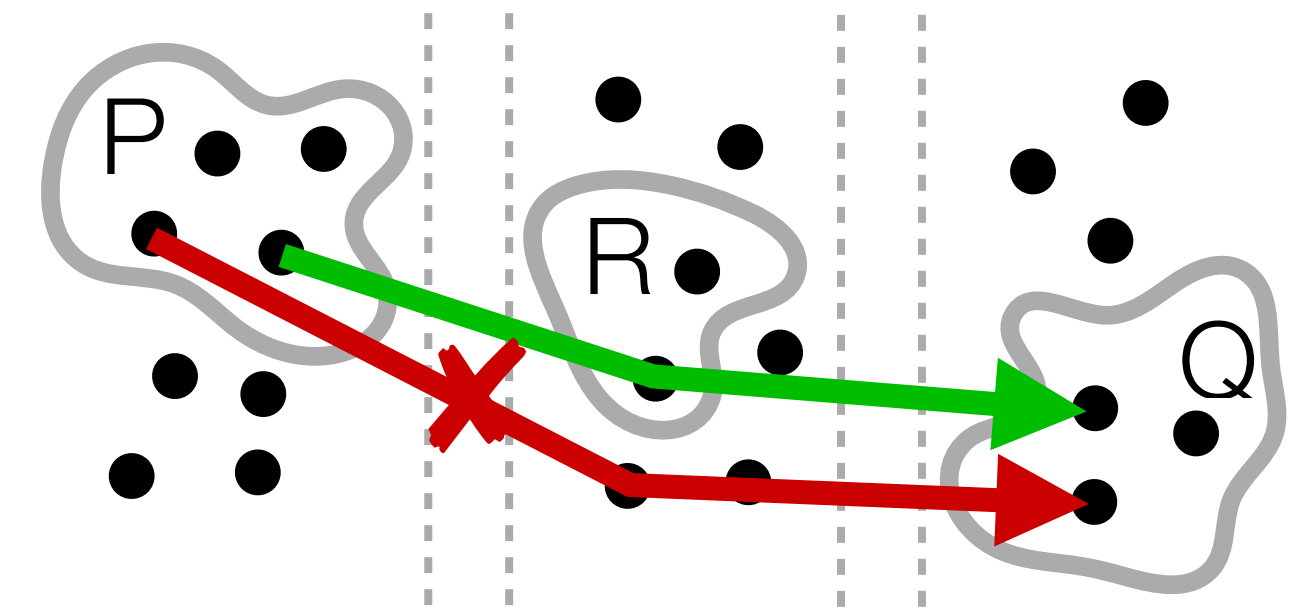
Decomposing Updates

Unbounded footprints without induction.

[OOPSLA'22, *under submission*]

Goal

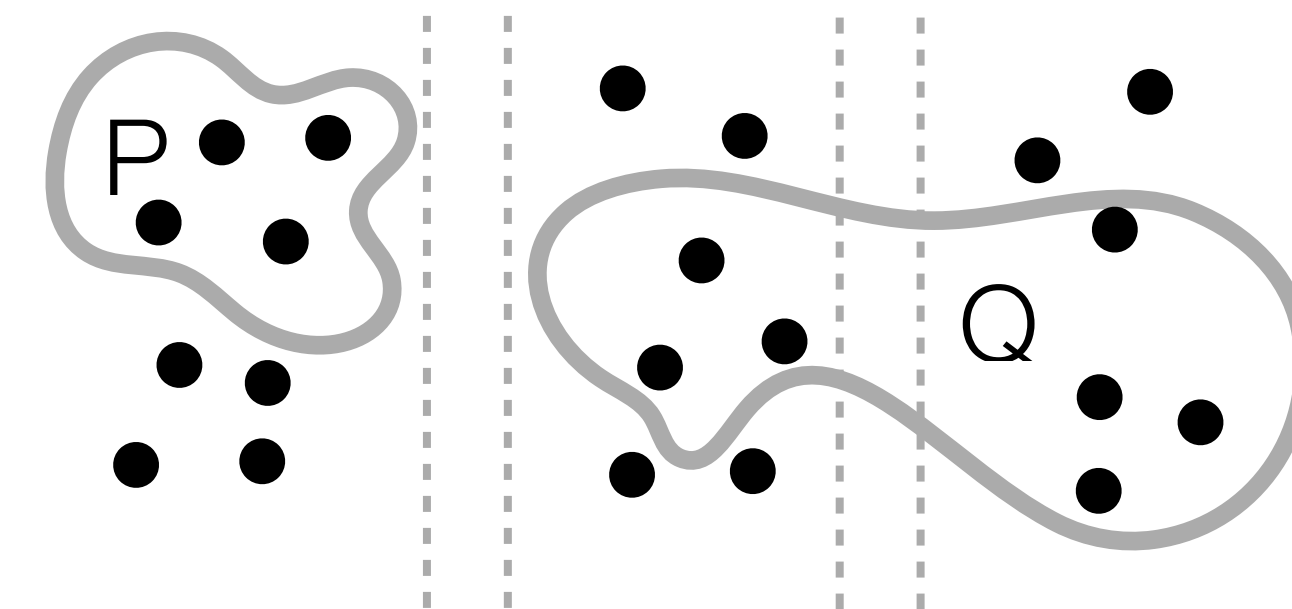
Show that a state (linearization point) has occurred.



Problem

The state may undergo interference.

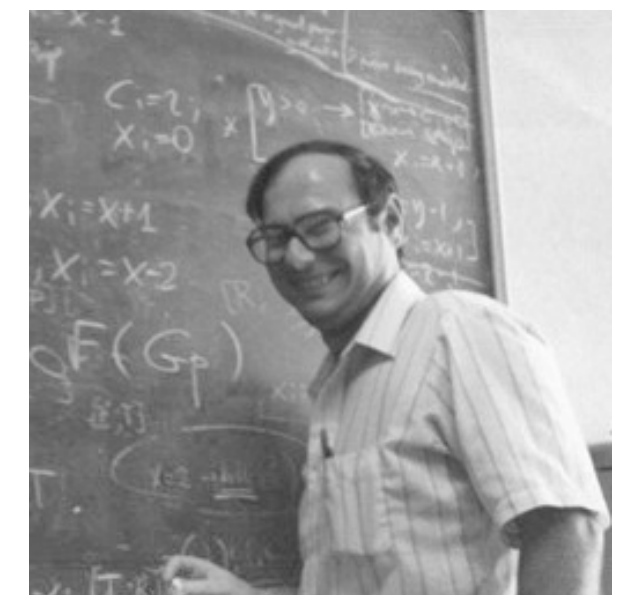
Its occurrence is not visible in the assertion.



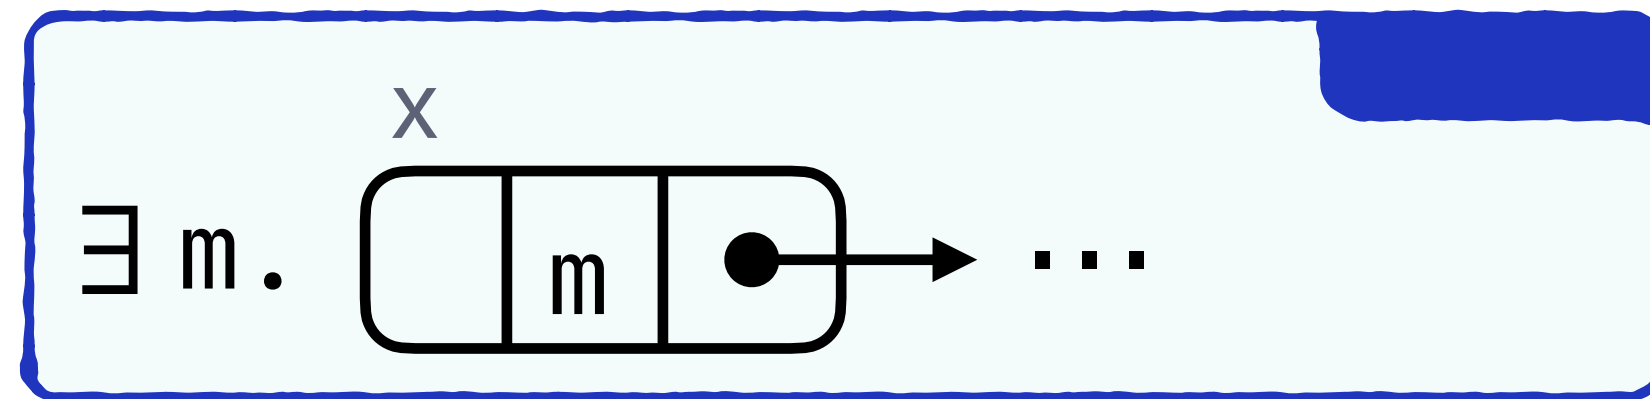
Solution

Introduce **past predicates** to keep track of the occurrence.

Past predicates are interference-free.



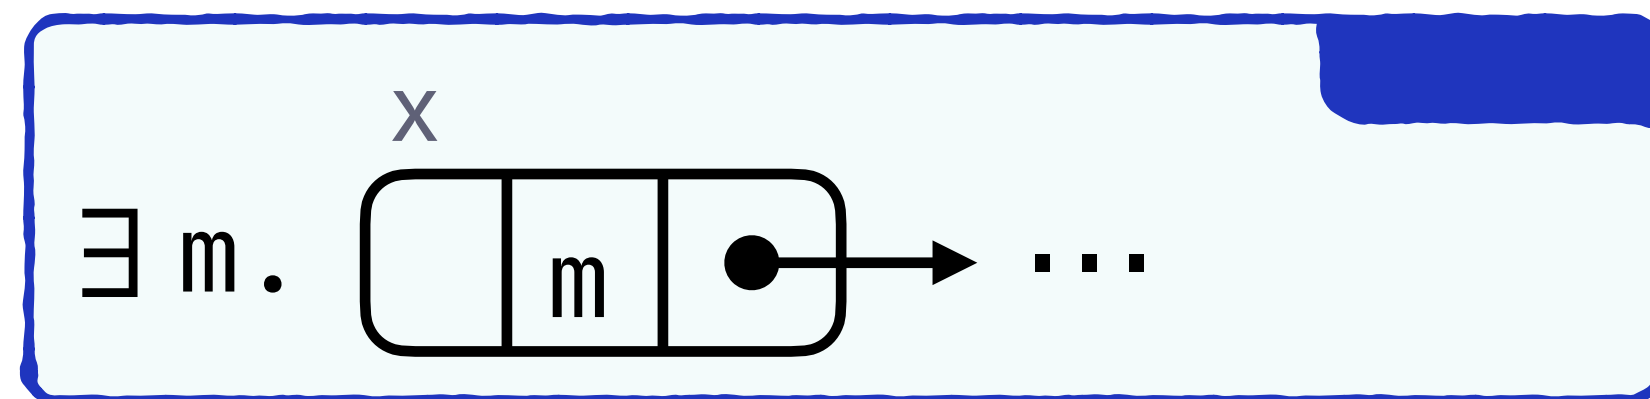
Example: find(k)



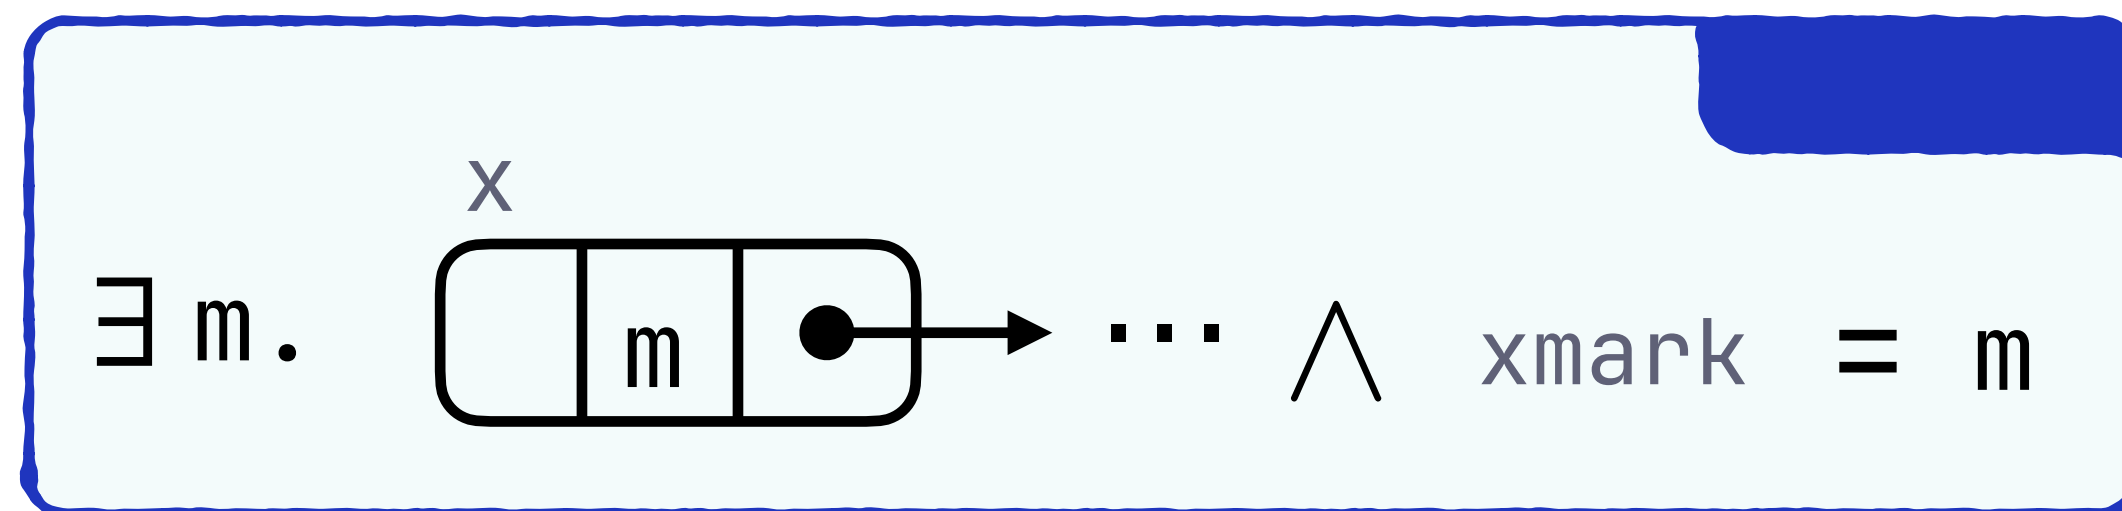
`xmark = x→mark;`

`if (!xmark){`

Example: find(k)

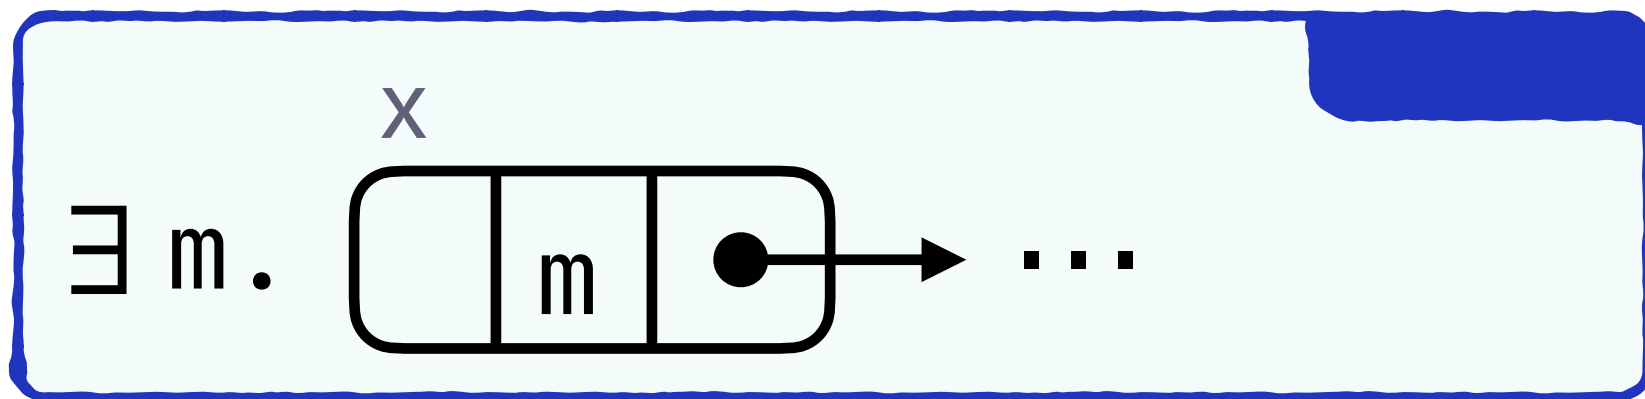


`xmark = x→mark;`

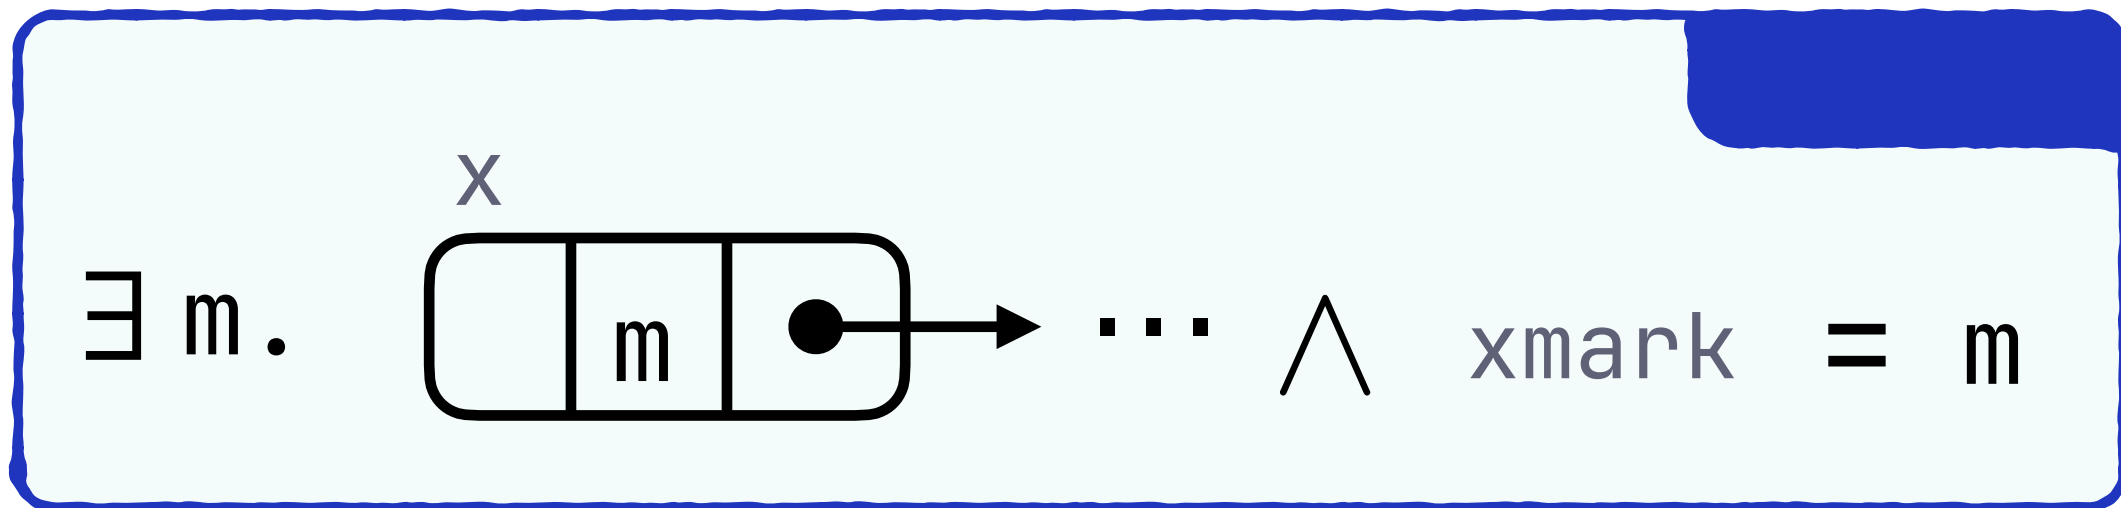


`if (!xmark){`

Example: find(k)



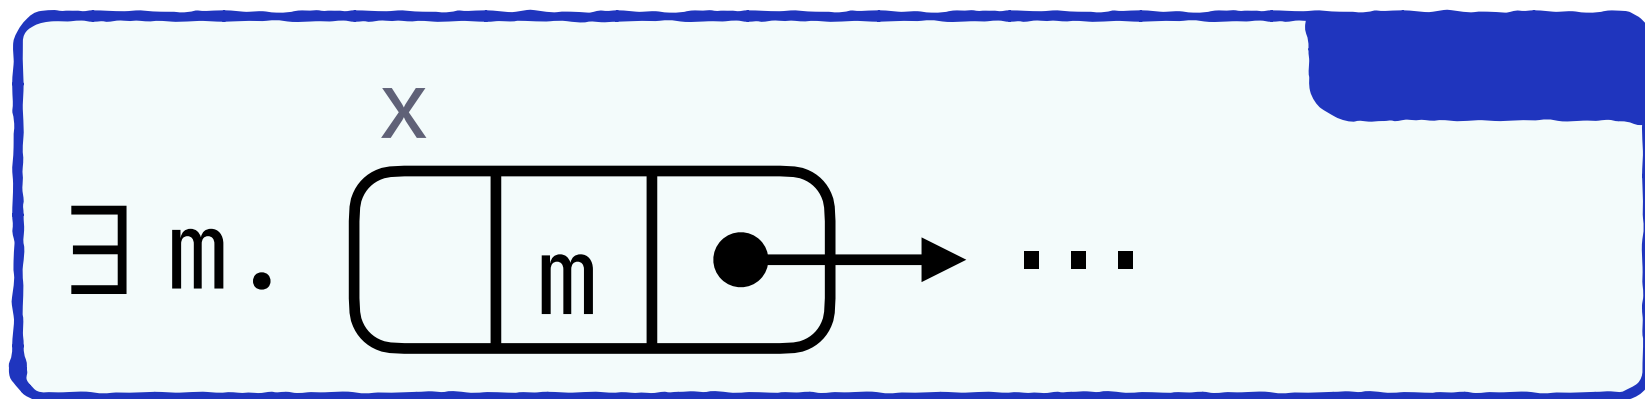
`xmark = x → mark;`



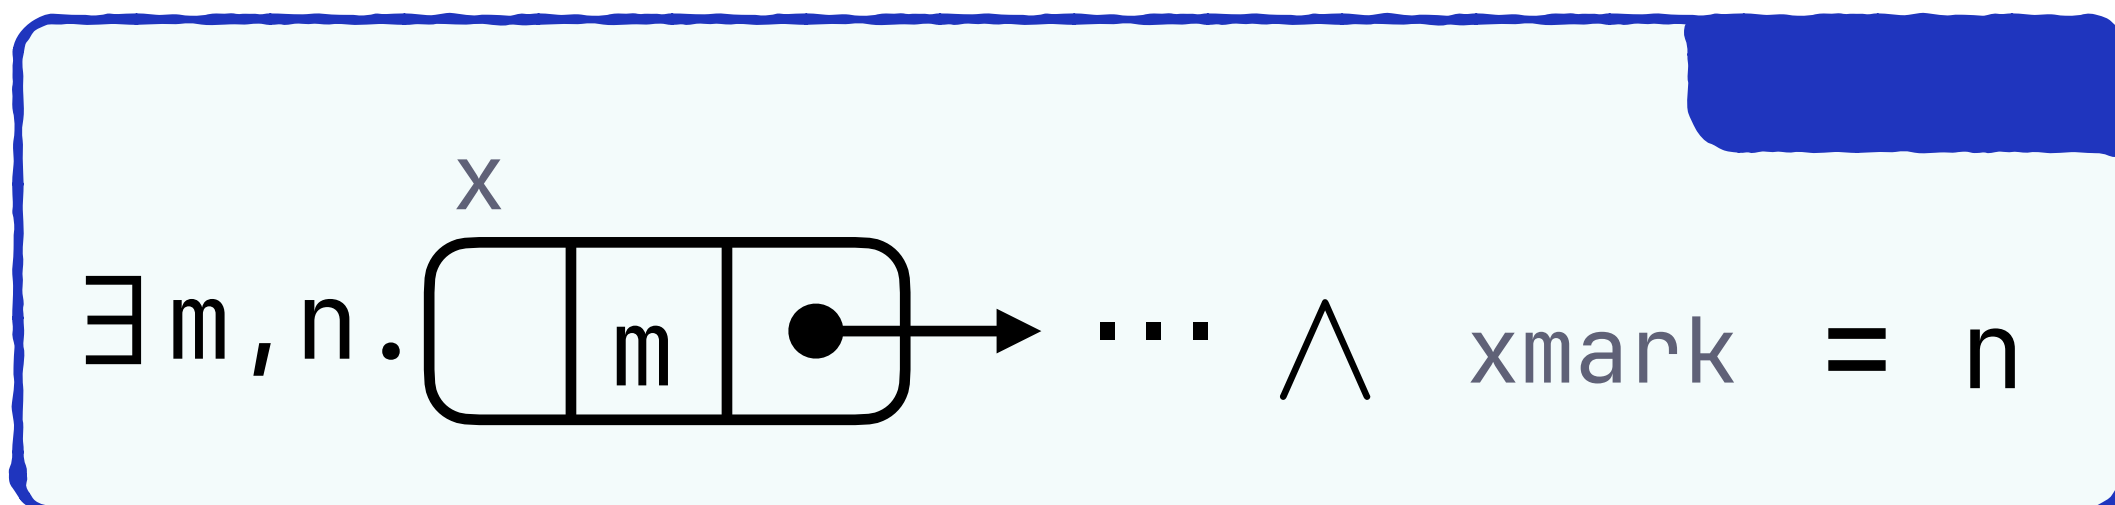
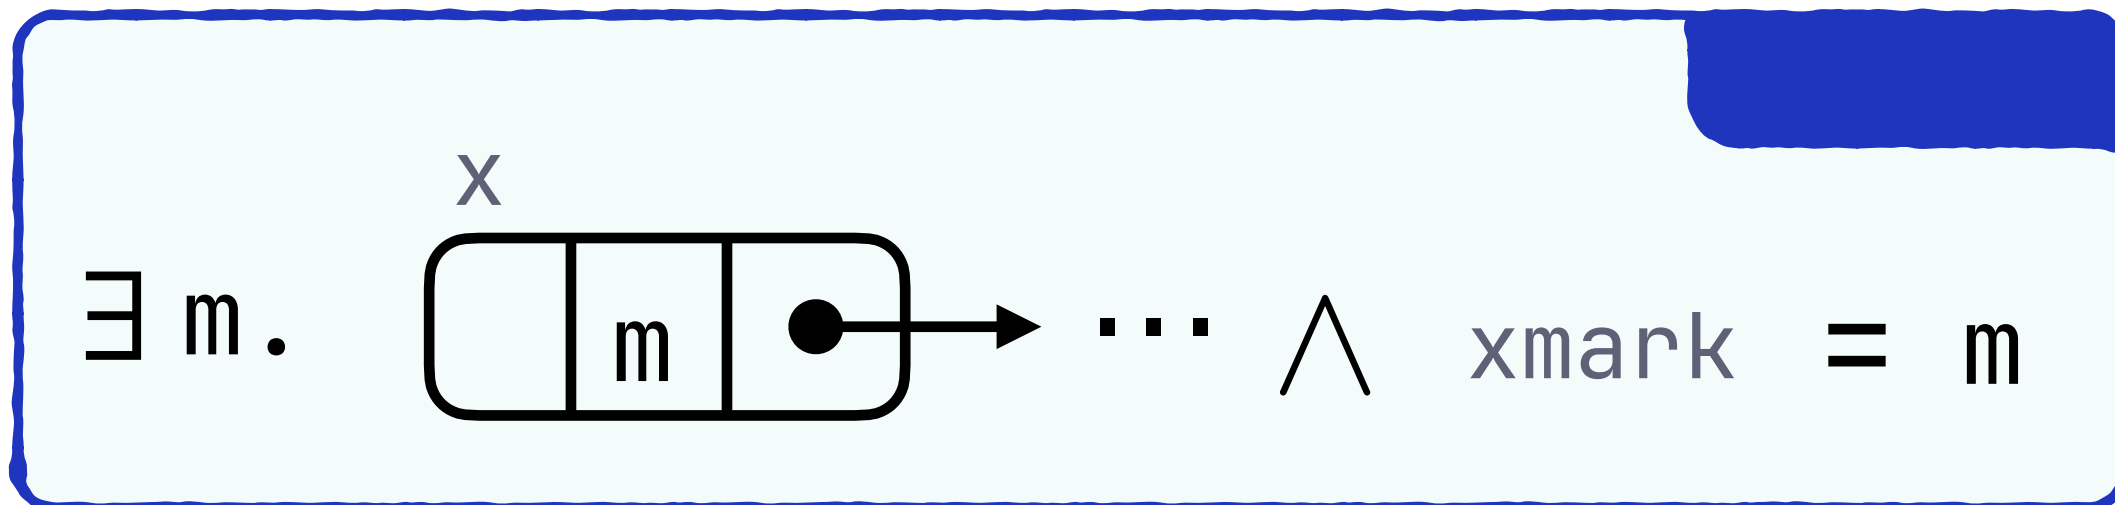
not interference-free

`if (!xmark){`

Example: find(k)

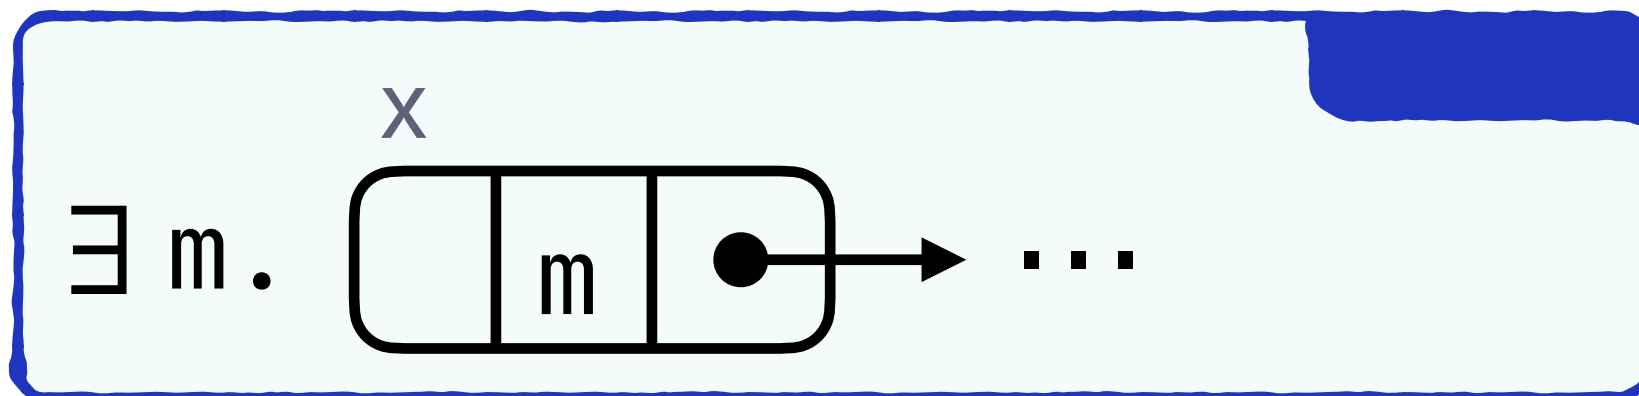


`xmark = x→mark;`

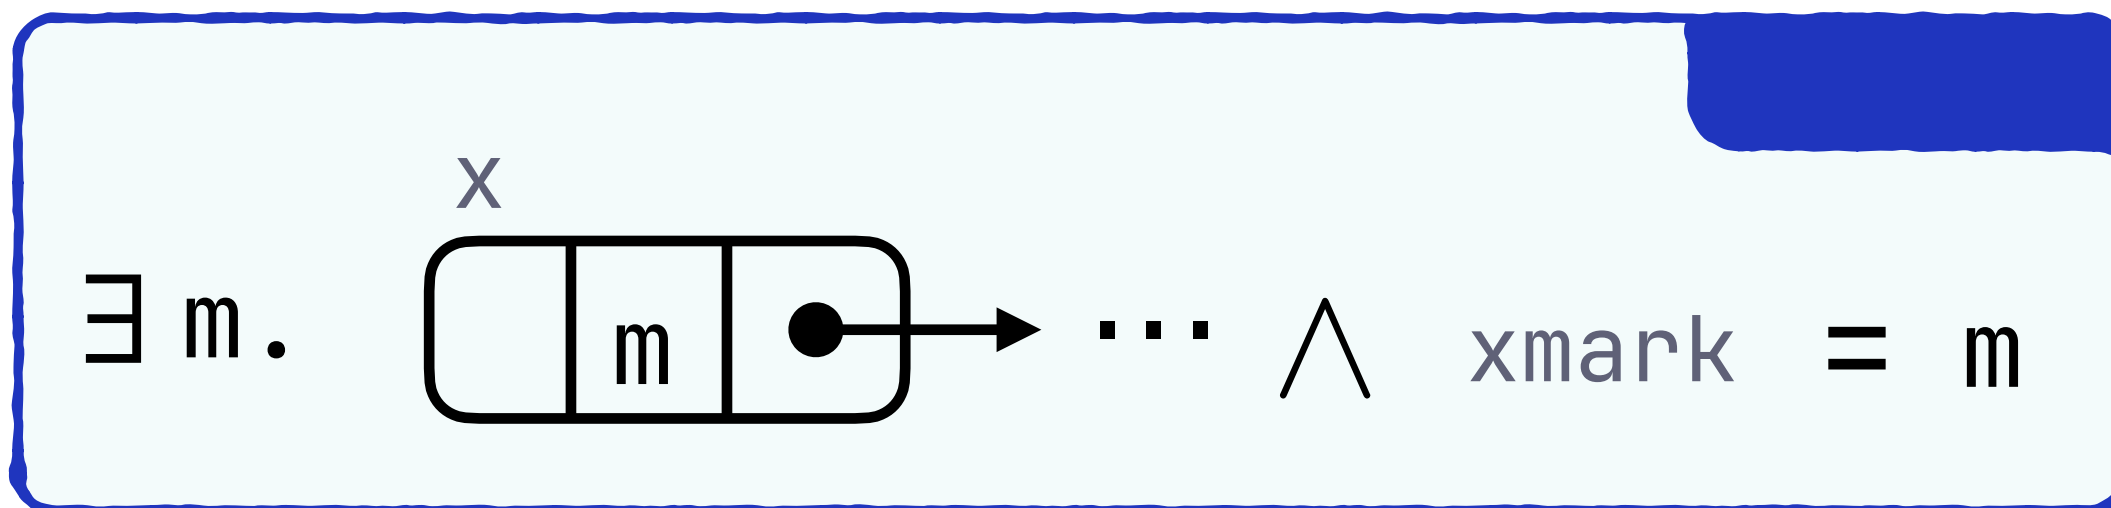


`if (!xmark){`

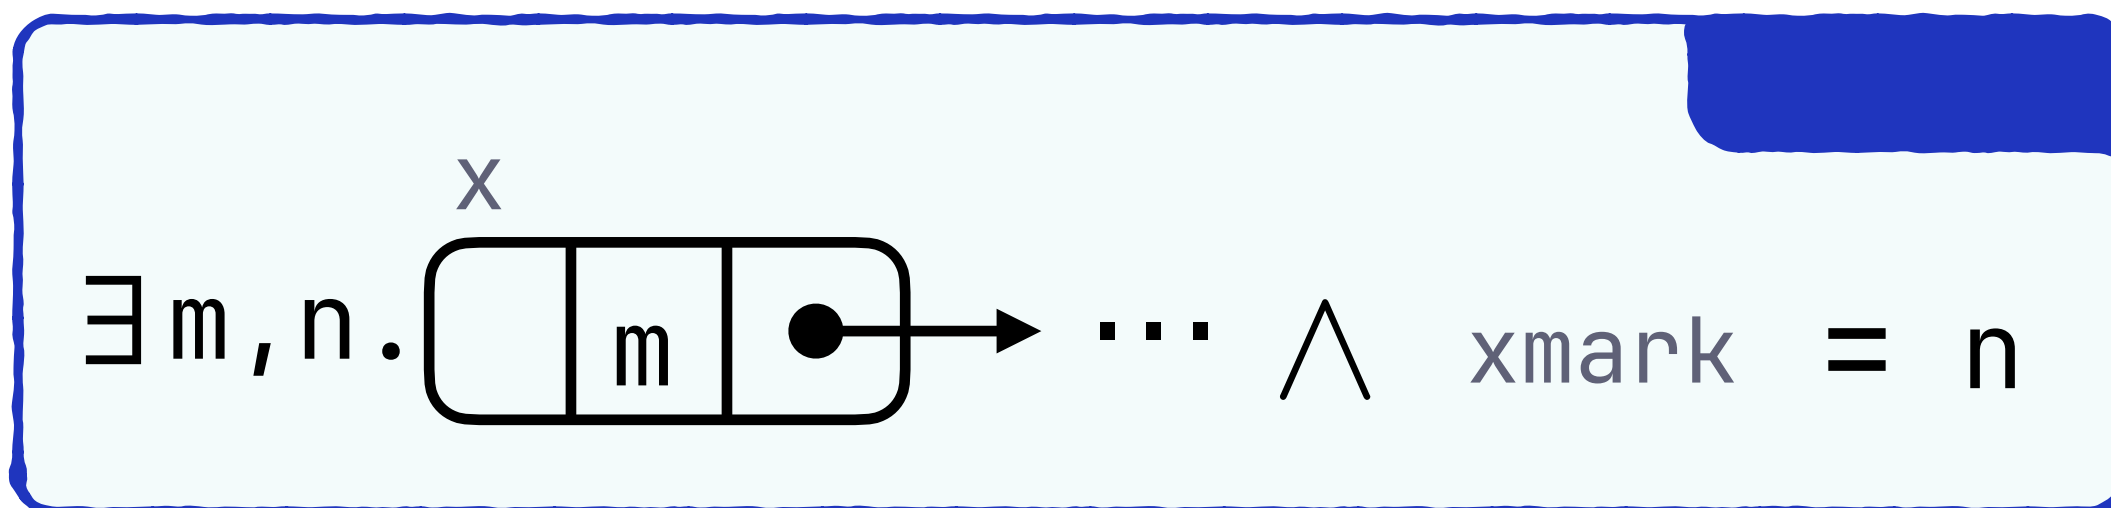
Example: find(k)



`xmark = x → mark;`



not interference-free



**interference-free
but not helpful**

`if (!xmark){`

State of the Art Solution

- **Prophecy variables**

[Abadi and Lamport 1991; Vafeiadis 2008; Liang and Feng 2013; Jung et al. 2020]

- ➔ case split along all possible computations
- ➔ predict each computation **eagerly** upfront
- ➔ fixes the linearization point

State of the Art Solution

- **Prophecy variables**

[Abadi and Lamport 1991; Vafeiadis 2008; Liang and Feng 2013; Jung et al. 2020]

- ➔ case split along all possible computations
- ➔ predict each computation **eagerly** upfront
- ➔ fixes the linearization point

**difficult to automate
scalability**


Separation Logic over Histories

- Let $(\Sigma, \star, \text{emp})$ be the separation algebra of states.

Separation Logic over Histories

- Let $(\Sigma, \star, \text{emp})$ be the separation algebra of states.
-  **separation algebra of computations** $(\Sigma^+, \star^+, \text{emp}^+)$.

Separation Logic over Histories

- Let $(\Sigma, \star, \text{emp})$ be the separation algebra of states.
-  **separation algebra of computations** $(\Sigma^+, \star^+, \text{emp}^+)$.
- State predicates are sets of states $Q \subseteq \Sigma$.

Separation Logic over Histories

- Let $(\Sigma, \star, \text{emp})$ be the separation algebra of states.
- **NEW** **separation algebra of computations** $(\Sigma^+, \star^+, \text{emp}^+)$.
- State predicates are sets of states $Q \subseteq \Sigma$.
- **NEW** **computation predicates:**

$$\boxed{Q \text{ Now}} = \Sigma^* \cdot Q$$

Separation Logic over Histories

- Let $(\Sigma, \star, \text{emp})$ be the separation algebra of states.
- **NEW** separation algebra of computations $(\Sigma^+, \star^+, \text{emp}^+)$.
- State predicates are sets of states $Q \subseteq \Sigma$.
- **NEW** computation predicates:

$$\boxed{Q \text{ Now}} = \Sigma^* . Q$$

$$\boxed{P \text{ Past}} = \Sigma^* . P . \Sigma^*$$

Separation Logic over Histories

- Let $(\Sigma, \star, \text{emp})$ be the separation algebra of states.
- **NEW** separation algebra of computations $(\Sigma^+, \star^+, \text{emp}^+)$.

- State predicates are sets of states $Q \subseteq \Sigma$.

- **NEW** computation predicates:

$$\boxed{Q \text{ Now}} = \Sigma^* . Q$$

$$\boxed{P \text{ Past}} = \Sigma^* . P . \Sigma^*$$

- Quantify on the level of computations (state-independently).



Separation Logic over Histories

- Let $(\Sigma, \star, \text{emp})$ be the separation algebra of states

- NEW** separation algebra of computations (Σ^+, \star)

- State predicates are sets of states $Q \subseteq \Sigma$.

- NEW** computation predicates:

$$\boxed{Q \text{ Now}} = \Sigma^* \cdot Q$$

$$\boxed{P \text{ Past}} = \Sigma^* \cdot P \cdot \Sigma^*$$

- Quantify on the level of computations (state-indepen



A Concurrent Program Logic with a Future and History

ROLAND MEYER, TU Braunschweig, Germany

THOMAS WIES, New York University, USA

SEBASTIAN WOLFF, New York University, USA

Verifying fine-grained optimistic concurrent programs remains an open problem. Modern program logics provide abstraction mechanisms and compositional reasoning principles to deal with the inherent complexity. However, their use is mostly confined to pencil-and-paper or mechanized proofs. We devise a new separation logic geared towards the lacking automation. While local reasoning is known to be crucial for automation, we are the first to show how to retain this locality for (i) reasoning about inductive properties without the need for ghost code, and (ii) reasoning about computation histories in hindsight. We implemented our new logic in a tool and used it to automatically verify challenging concurrent search structures that require inductive properties and hindsight reasoning, such as the Harris set.

CCS Concepts: • **Theory of computation** → **Separation logic; Hoare logic; Automated reasoning; Program verification; Programming logic.**

Additional Key Words and Phrases: Linearizability, Non-blocking Data Structures, Harris Set

ACM Reference Format:

Roland Meyer, Thomas Wies, and Sebastian Wolff. 2022. A Concurrent Program Logic with a Future and History. *Proc. ACM Program. Lang.* 6, OOPSLA2, Article 174 (October 2022), 30 pages. <https://doi.org/10.1145/3563337>

1 INTRODUCTION

Concurrency comes at a cost, at least, in terms of increased effort when verifying program correctness. There has been a proliferation of concurrent program logics that provide an arsenal of reasoning techniques to address this challenge [Bell et al. 2010; Delbianco et al. 2017; Elmas et al. 2010; Fu et al. 2010; Gotsman et al. 2013; Gu et al. 2018; Hemed et al. 2015; Jung et al. 2018; Liang and Feng 2013; Manna and Pnueli 1995; Parkinson et al. 2007; Sergey et al. 2015; Vafeiadis and Parkinson 2007]. In addition, a number of general approaches have been developed to help structure the high-level proof argument [Feldman et al. 2018, 2020; Kragl et al. 2020; O’Hearn et al. 2010; Shasha and Goodman 1988]. However, the use of these techniques has been mostly confined to manual proofs done on paper, or mechanized proofs constructed in interactive proof assistants. We distill from these works a concurrent separation logic suitable for automating the construction of local correctness proofs for highly concurrent data structures. We focus on concurrent search structures (sets and maps indexed by keys), but the developed techniques apply more broadly. Our guiding principle is to perform all inductive reasoning, both in time and space, in lock-step with the program execution. The reasoning about inductive properties of graph structures and computation histories is relegated to the meta-theory of the logic by choosing appropriate semantic models.

Running Example. We motivate our work using the Harris non-blocking set data structure [Harris 2001], which we will also use as a running example throughout the paper.

Authors’ addresses: Roland Meyer, TU Braunschweig, Germany, roland.meyer@tu-bs.de; Thomas Wies, New York University, USA, wies@cs.nyu.edu; Sebastian Wolff, New York University, USA, sebastian.wolff@cs.nyu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

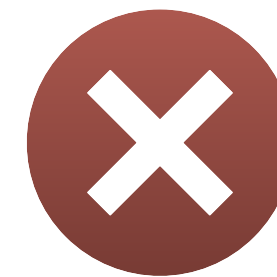
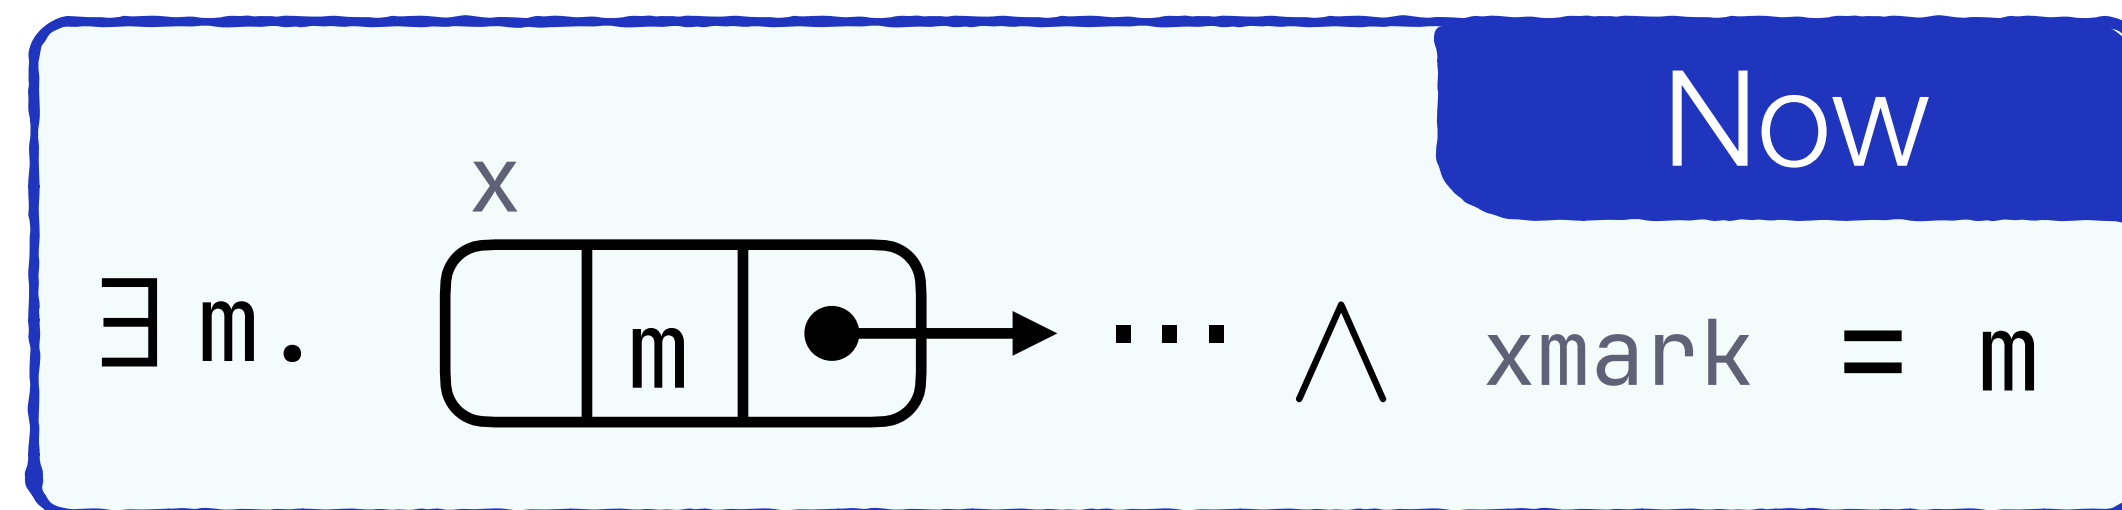
2475-1421/2022/10-ART174

<https://doi.org/10.1145/3563337>

Example: find(k)

...

```
xmark = x → mark;
```



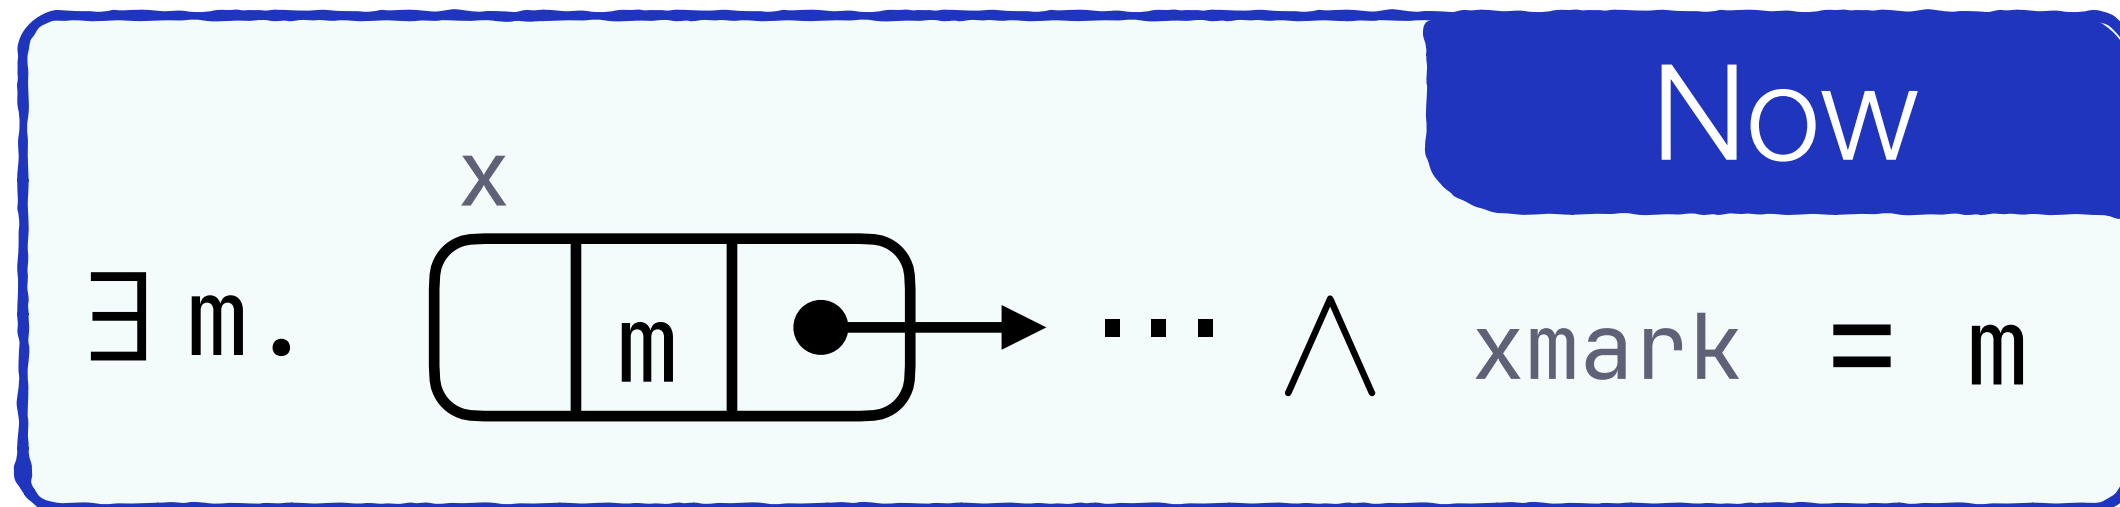
not interference-free

```
if (!xmark){
```

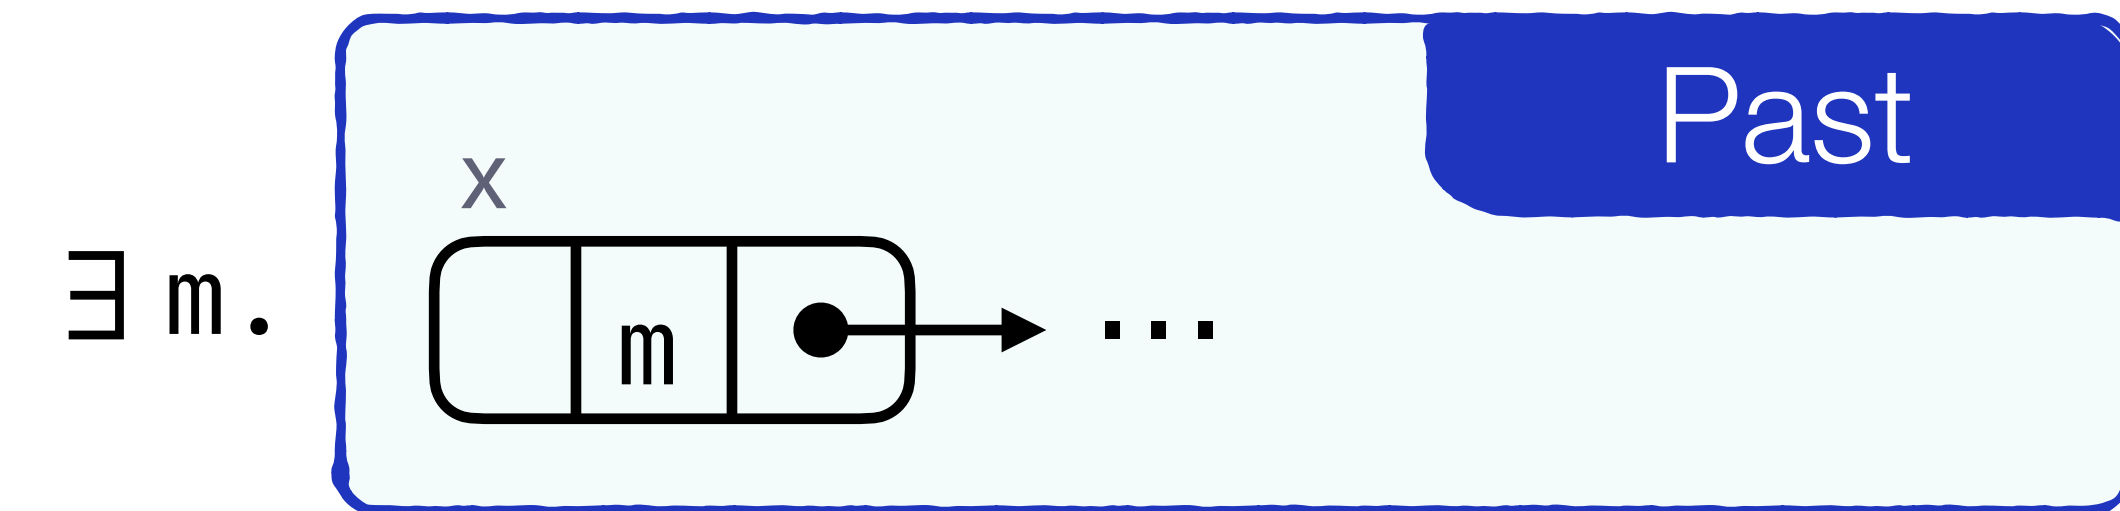
Example: find(k)

...

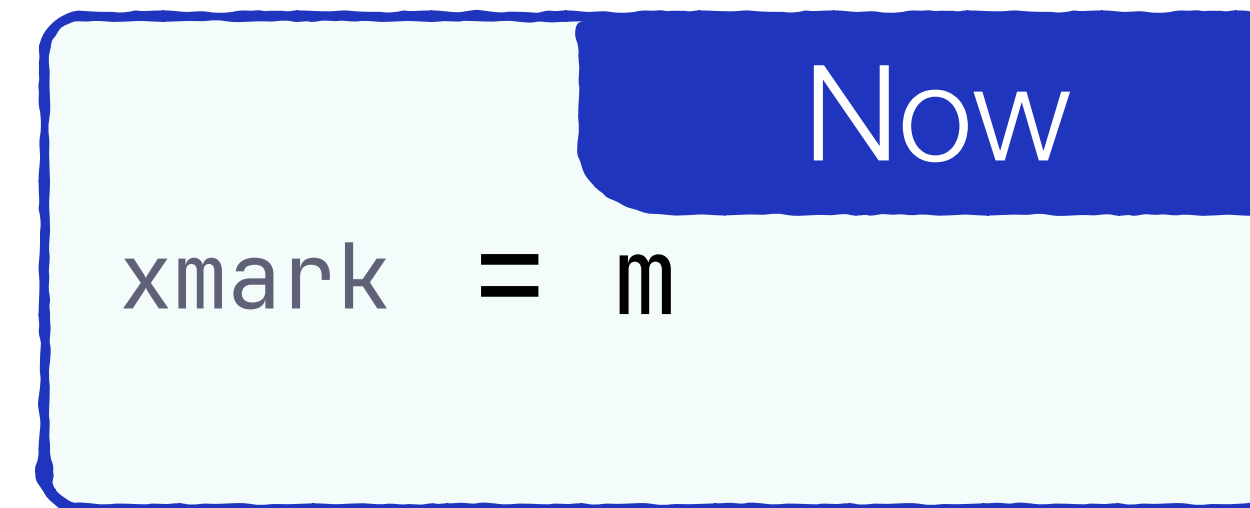
`xmark = x→mark;`



not interference-free



\wedge

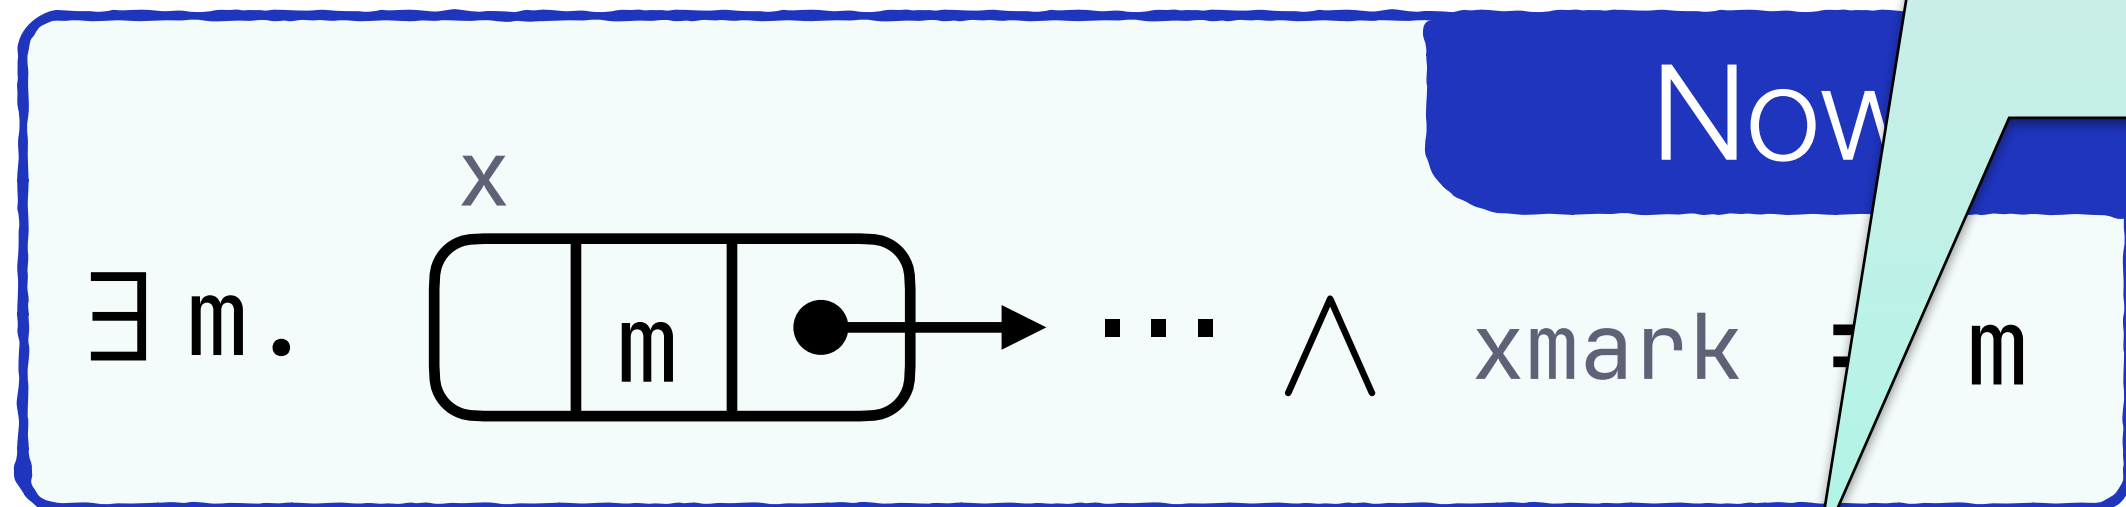


`if (!xmark){`

Example: find(k)

...

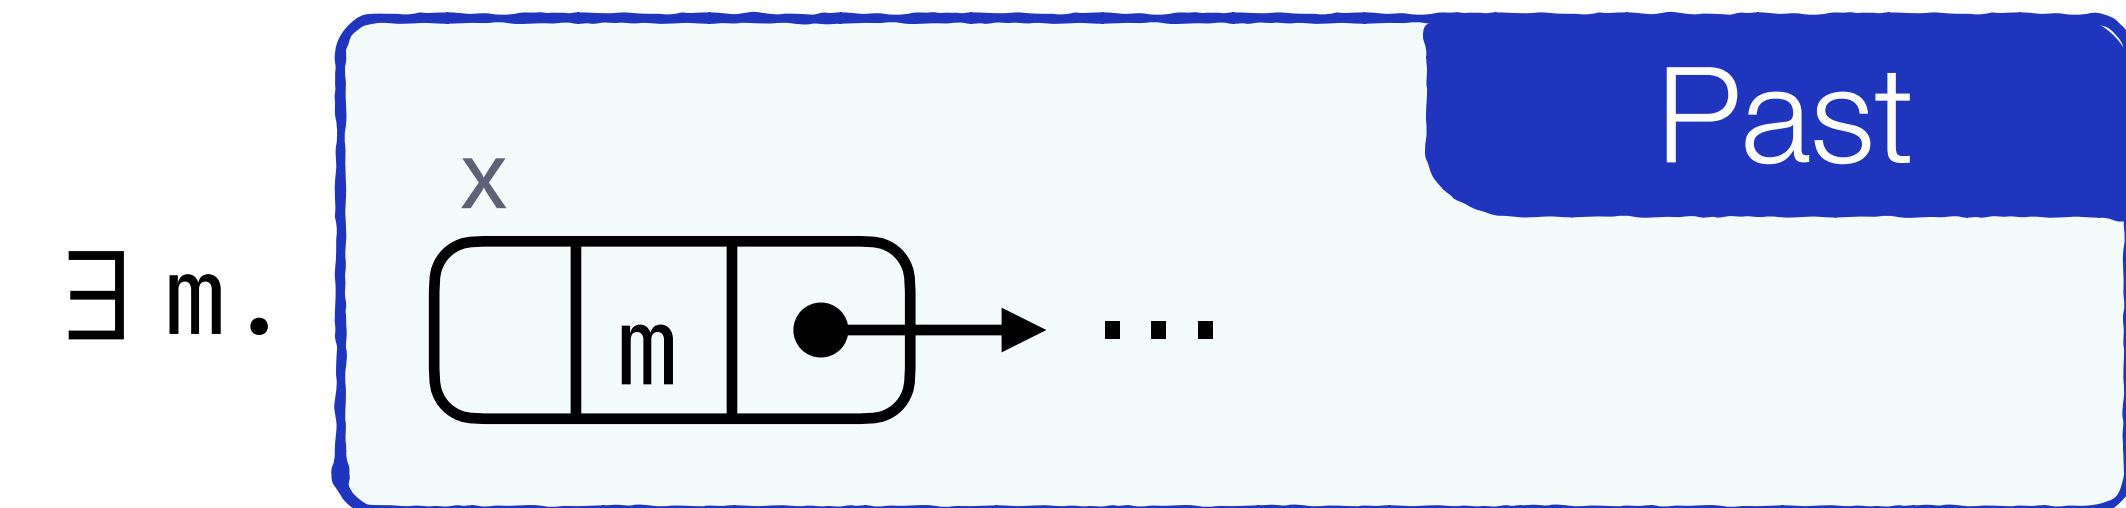
xmark = x → mark;



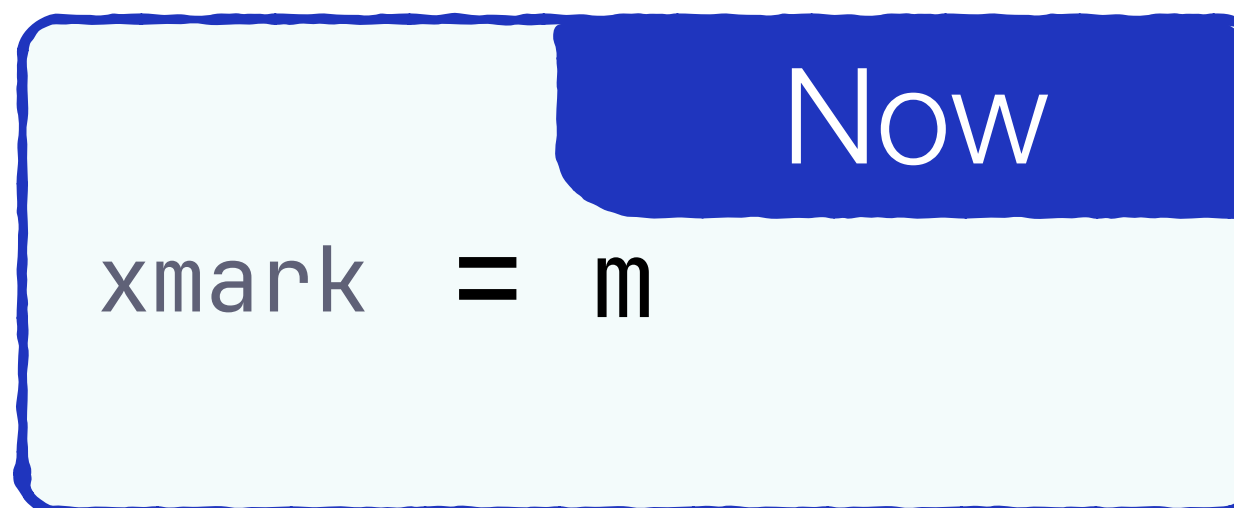
Past predicates are interference-free



not interference-free



\wedge

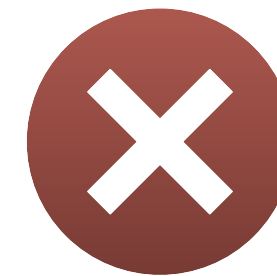
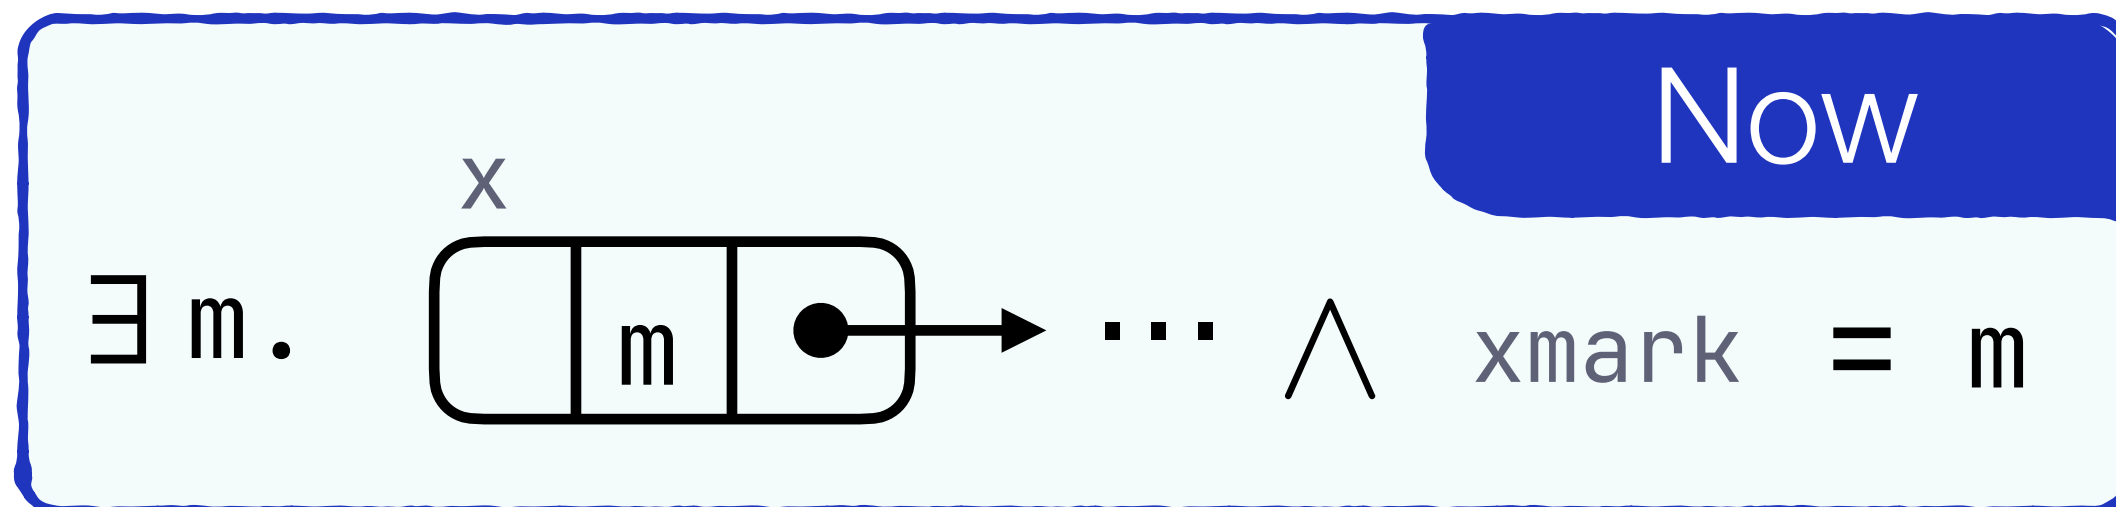


if (!xmark){

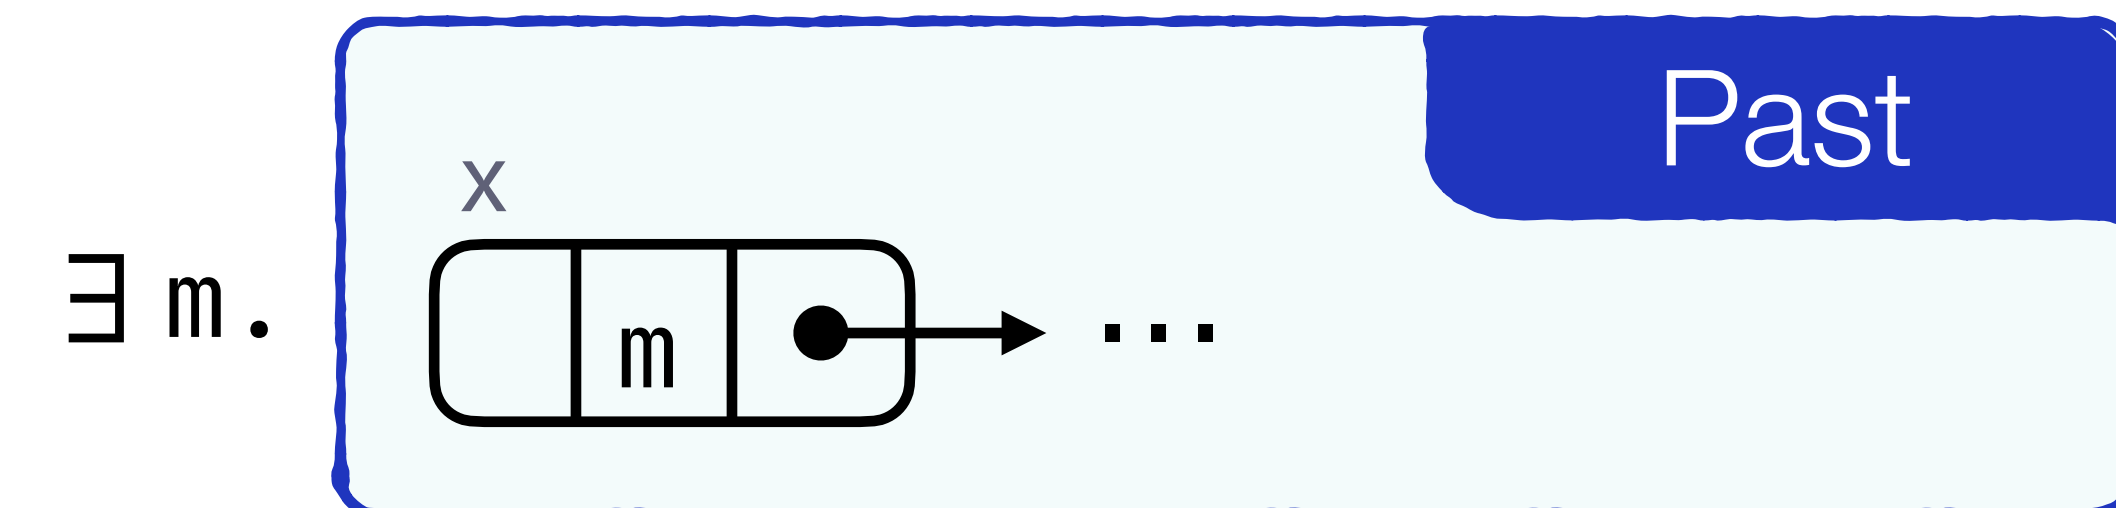
Example: find(k)

...

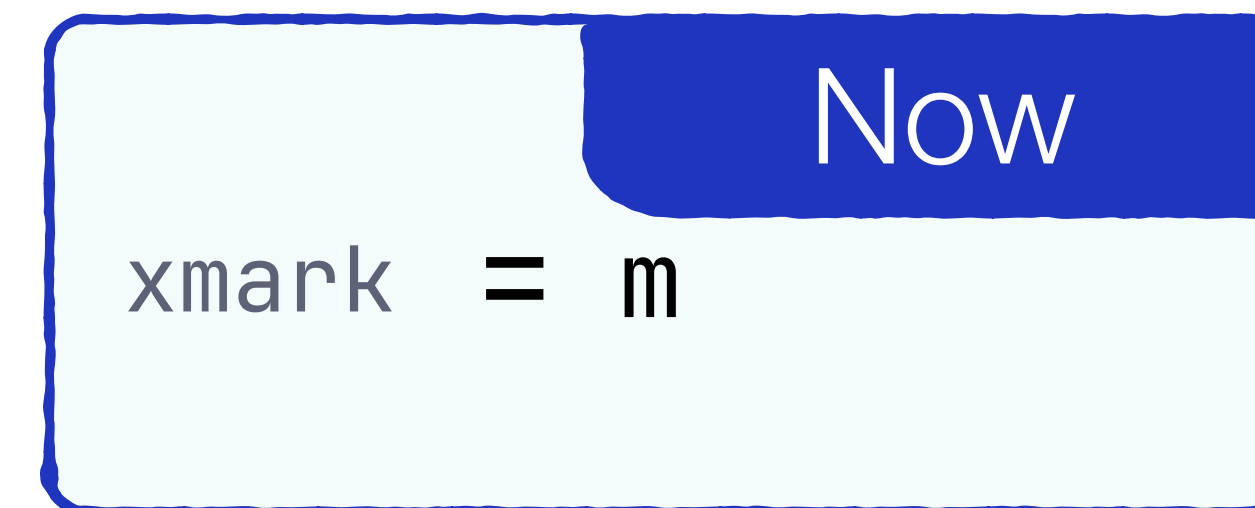
xmark = x → mark;



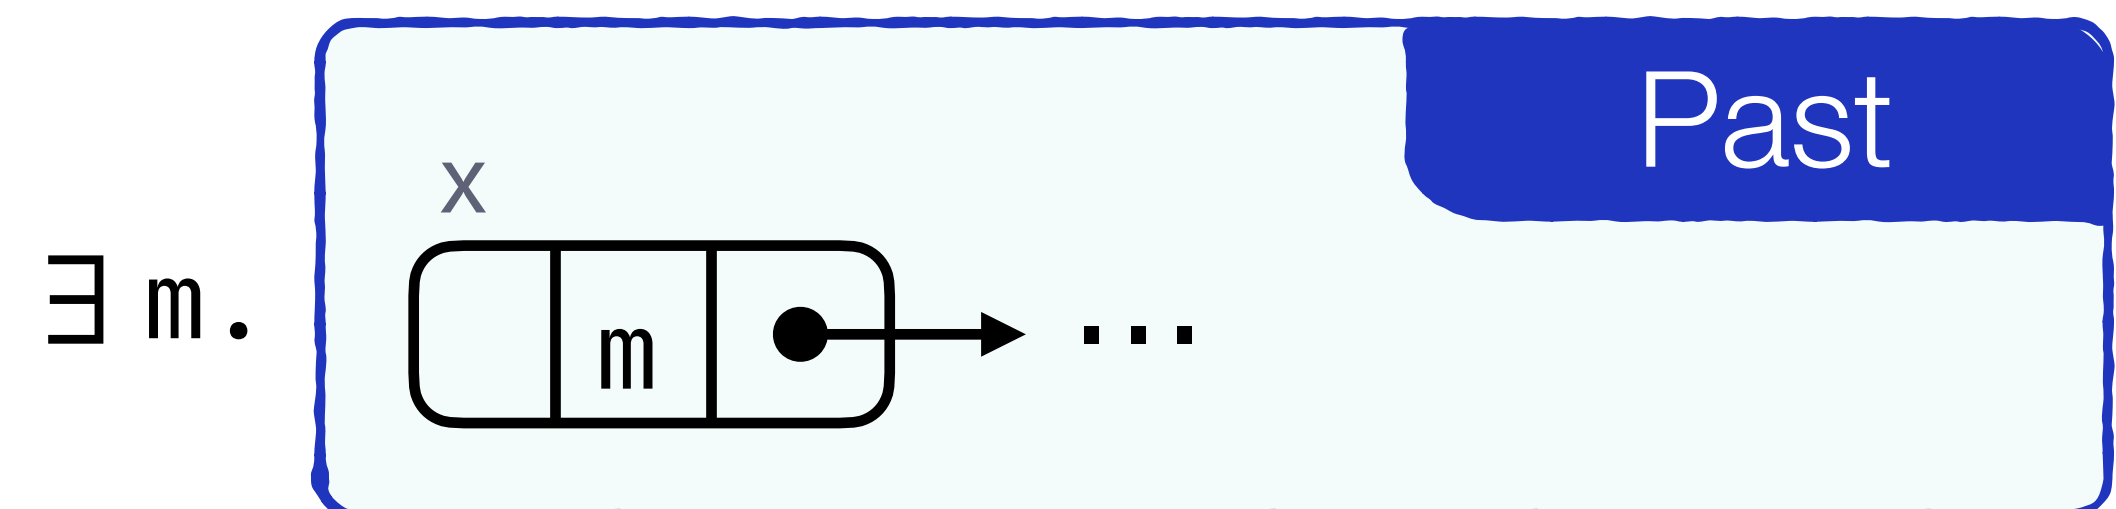
not interference-free



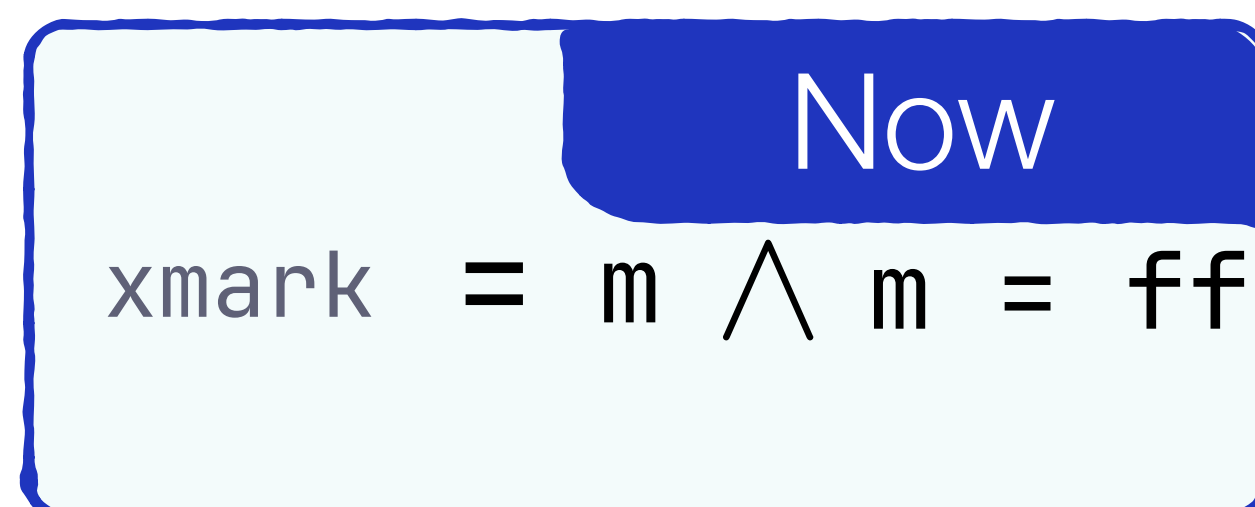
\wedge



if (!xmark){



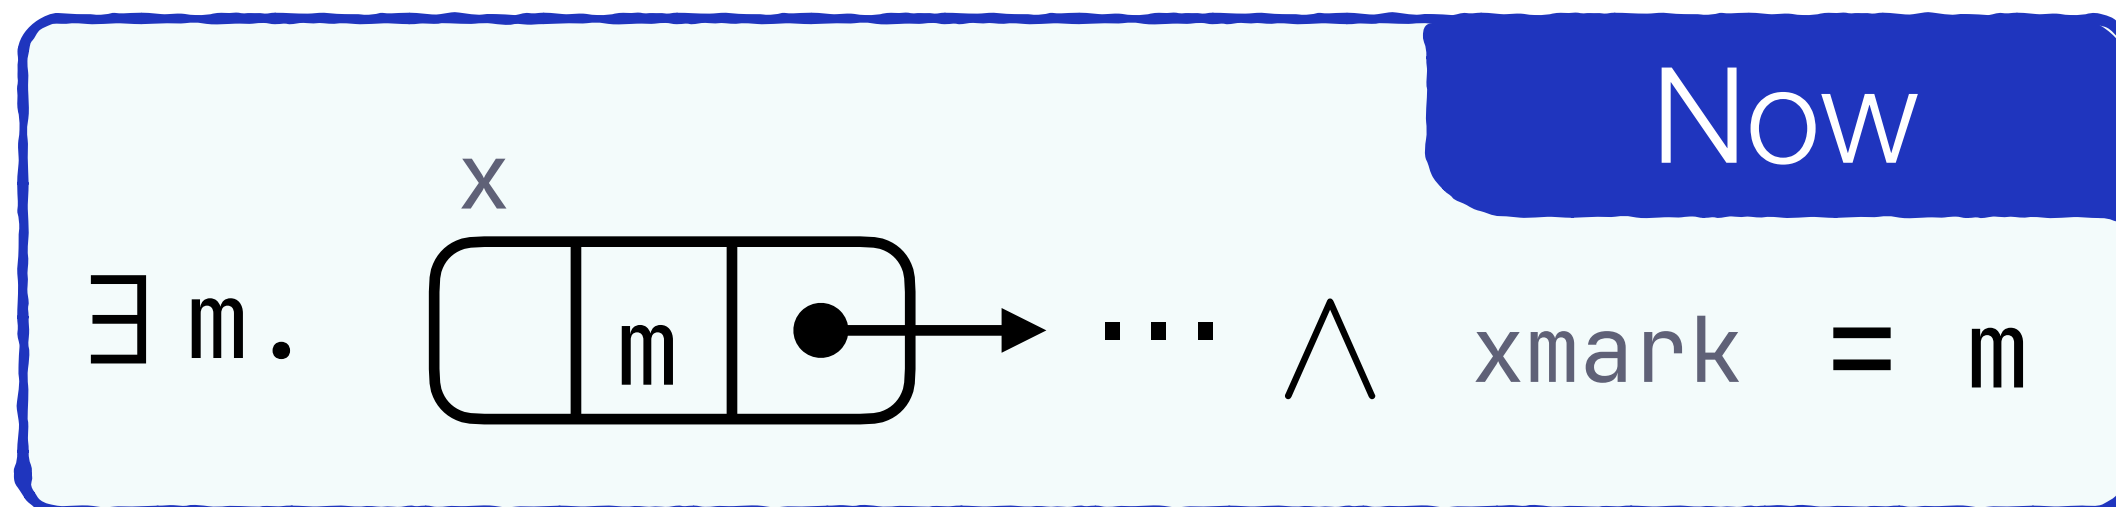
\wedge



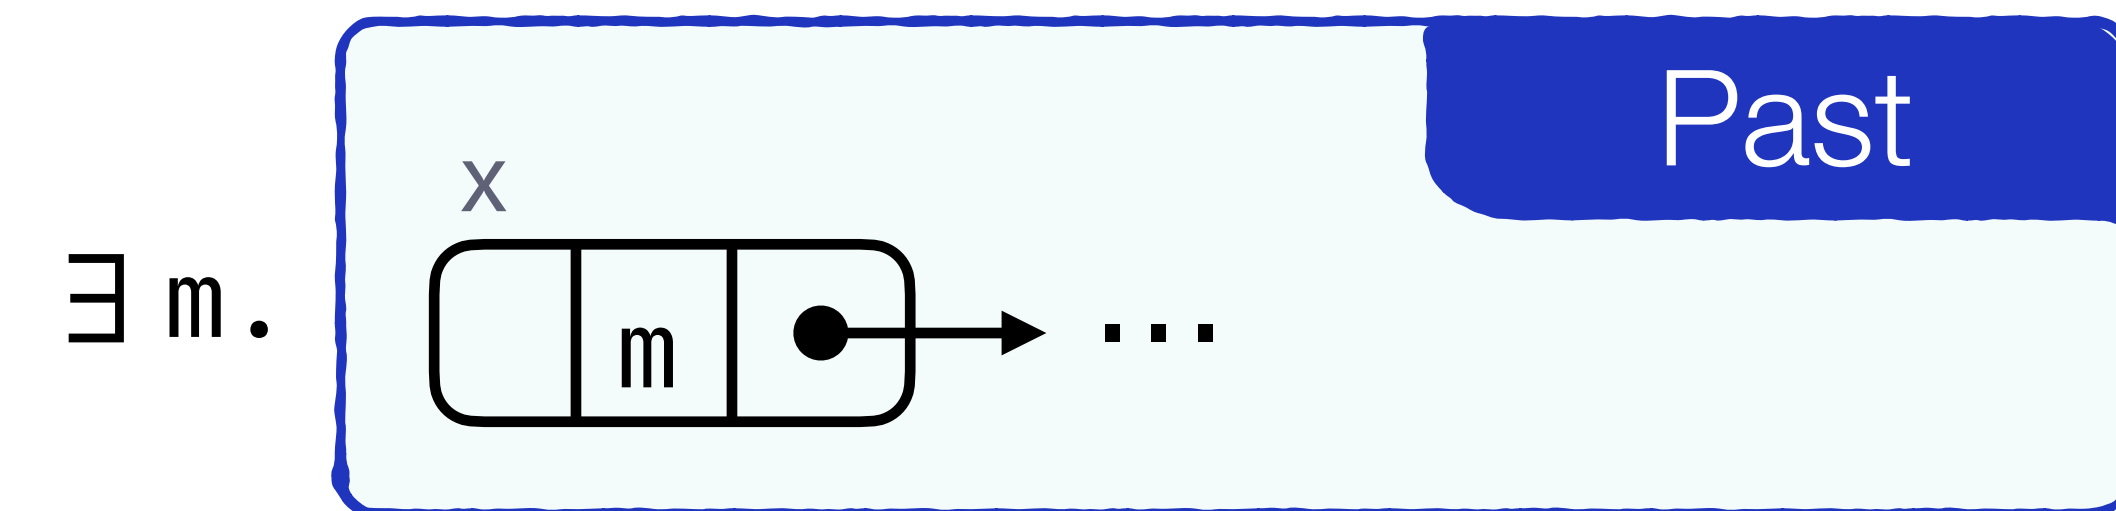
Example: find(k)

...

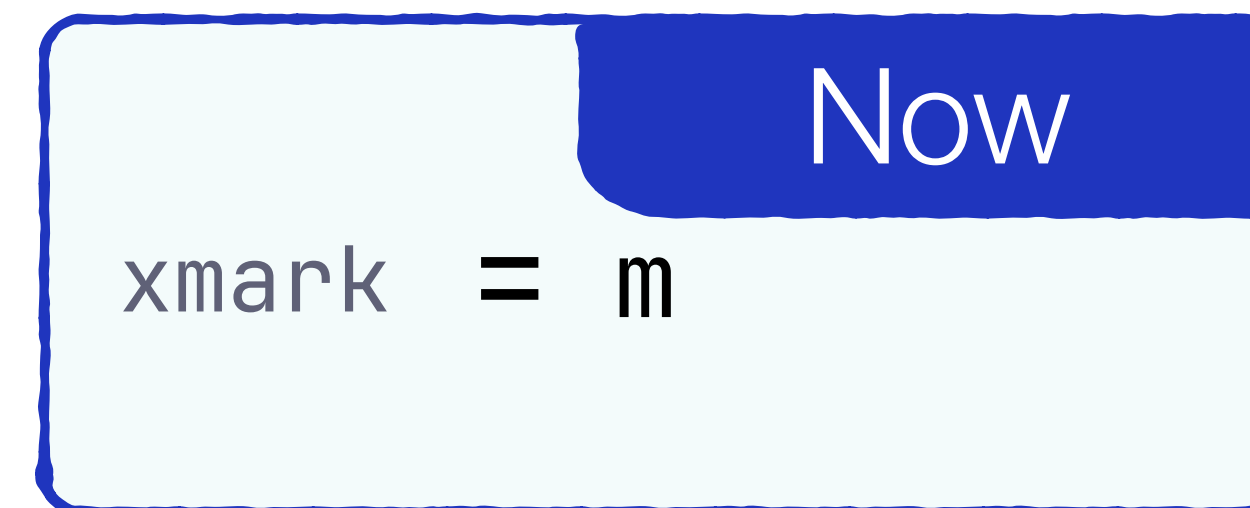
xmark = x → mark;



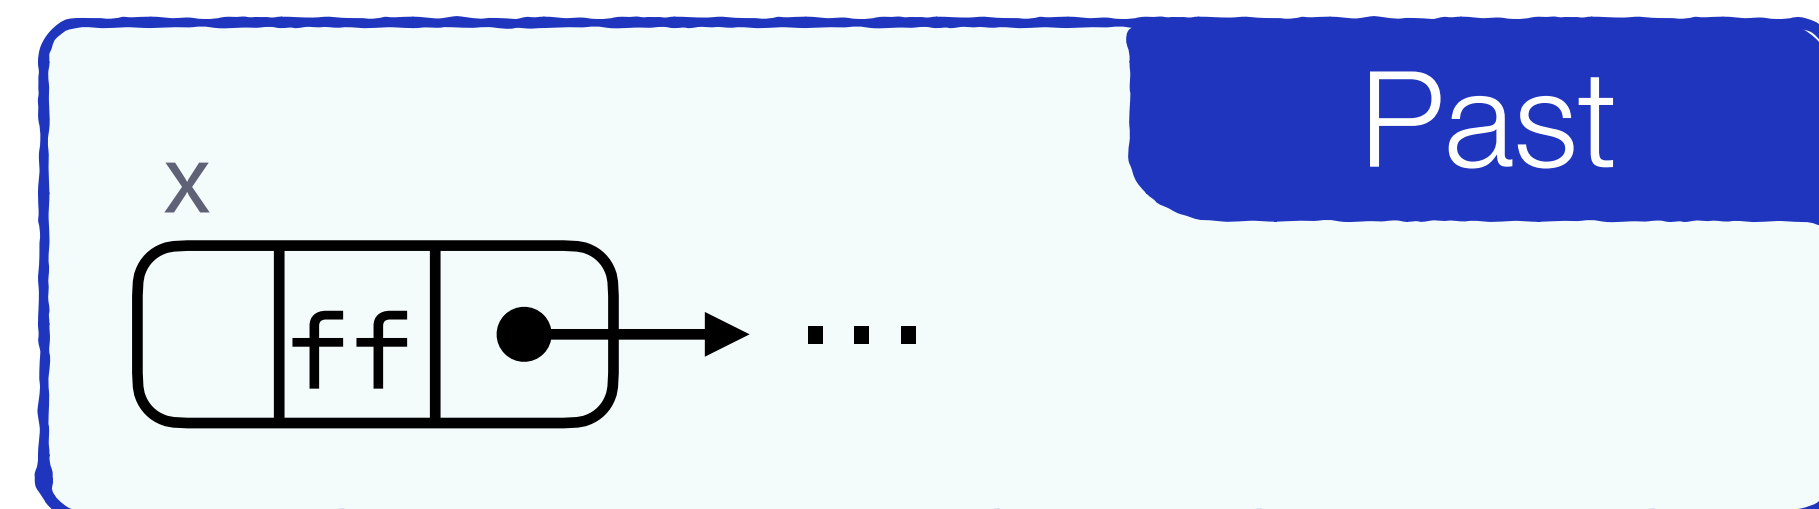
not interference-free



\wedge



if (!xmark){



Goal

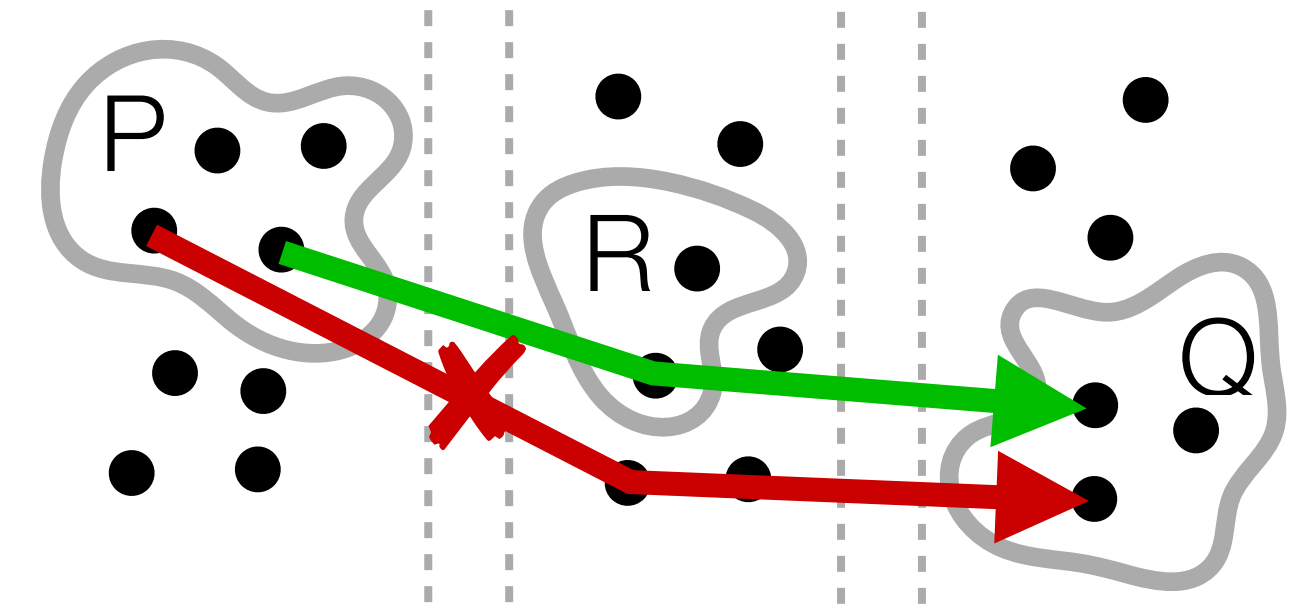
Show that a state (linearization point) has occurred.

Problem

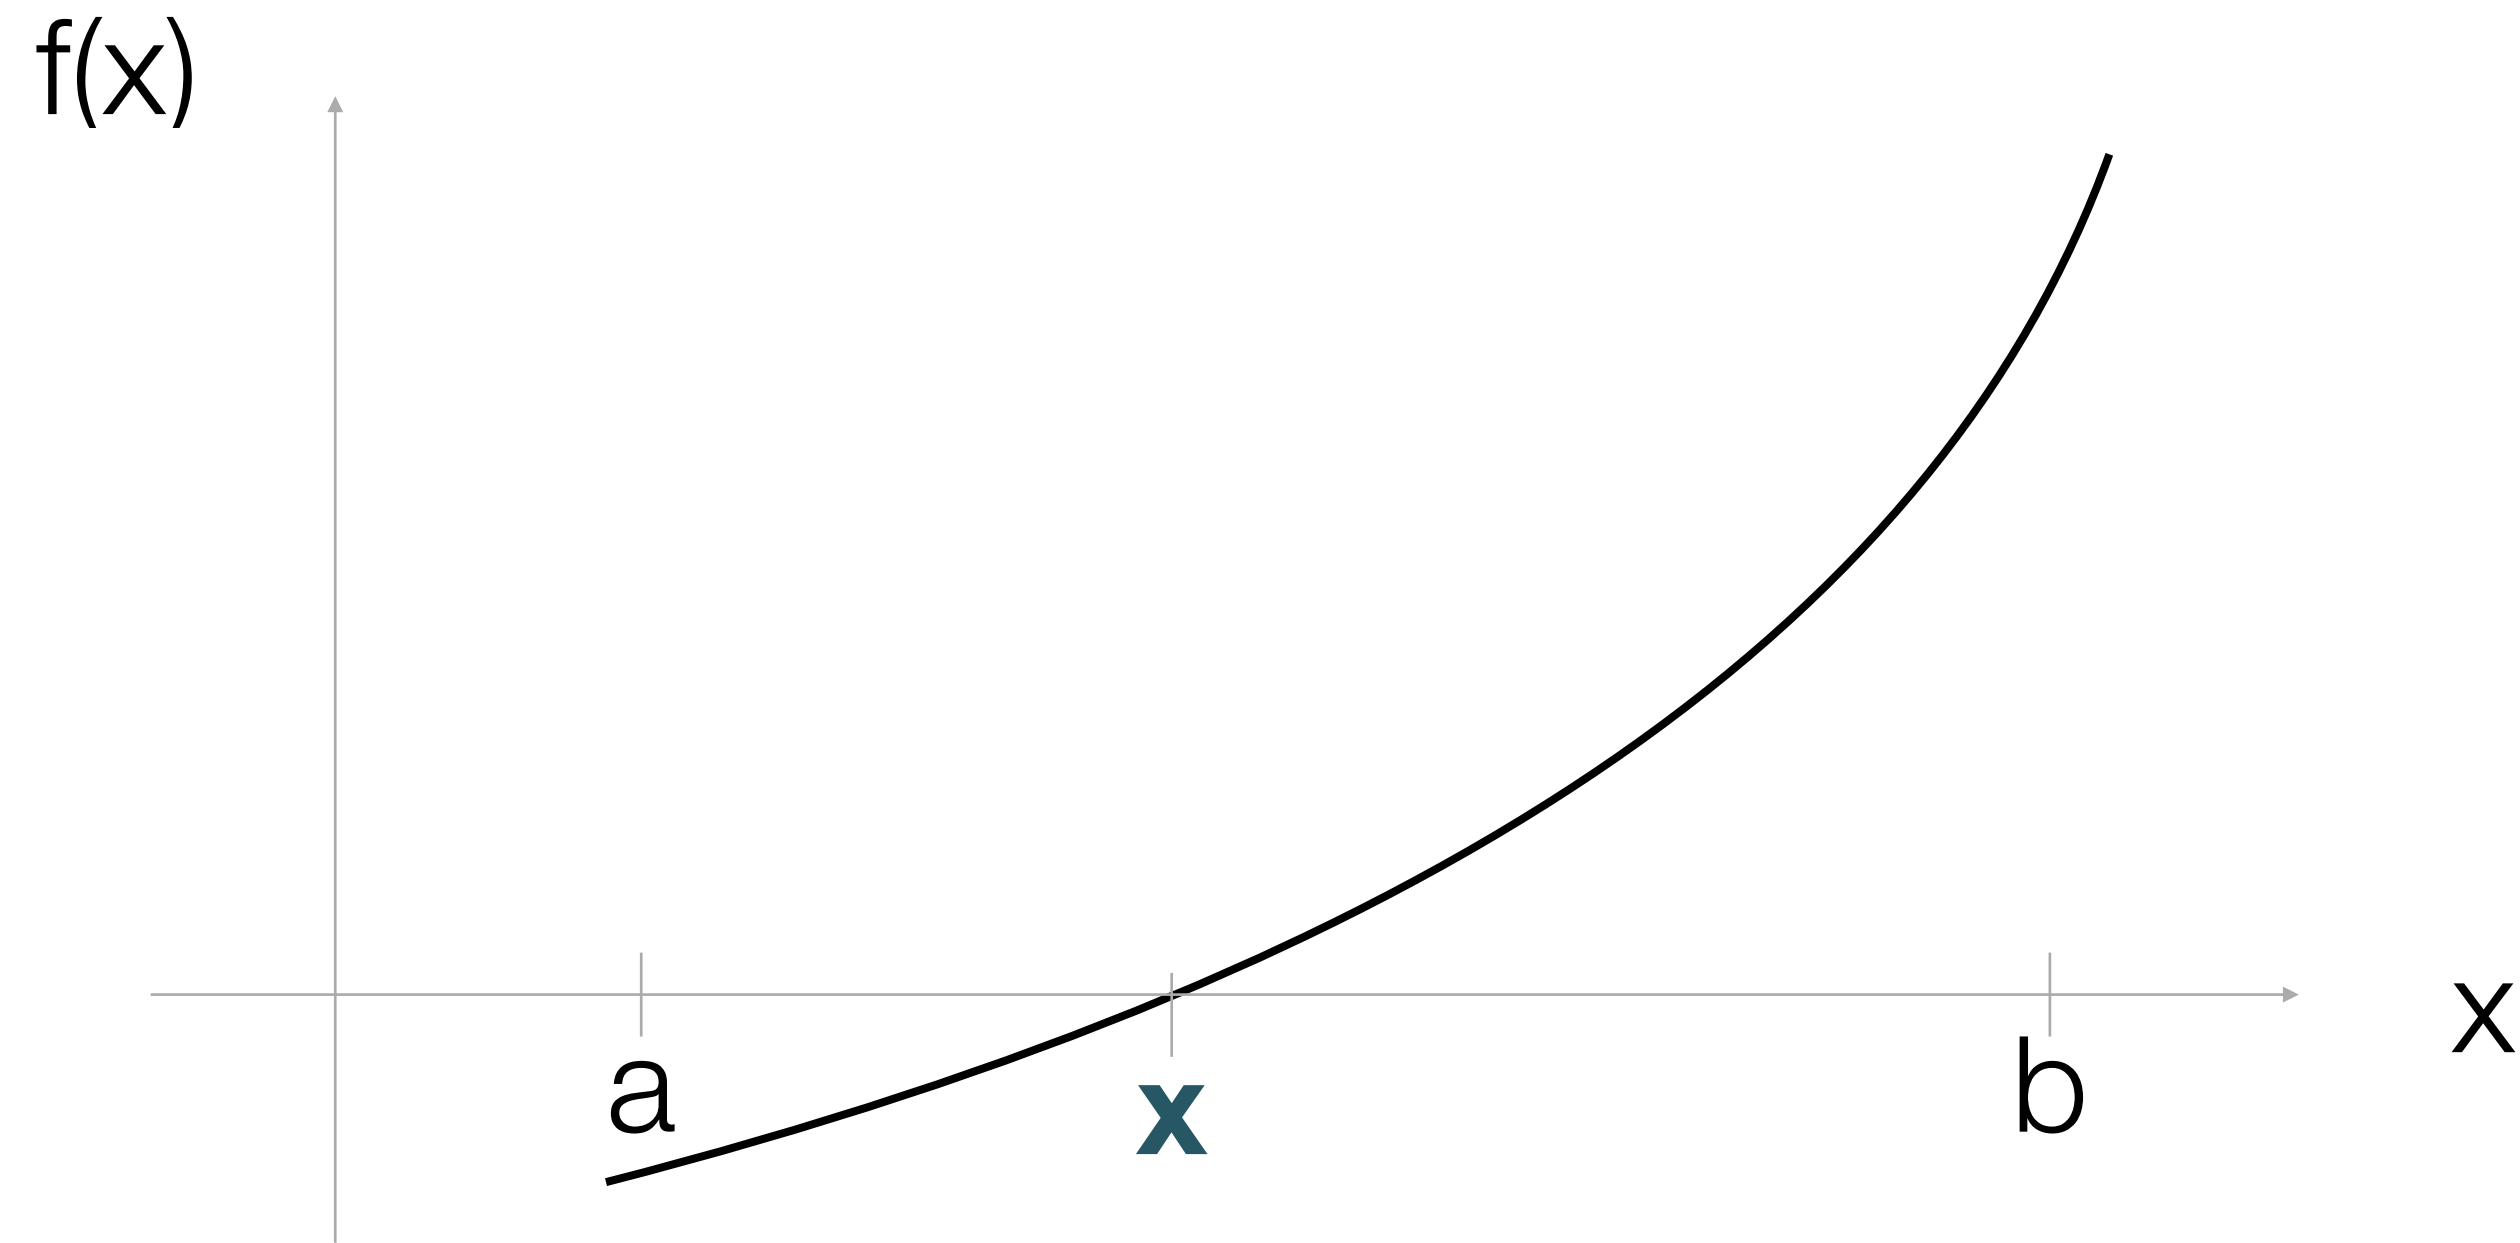
State-independent quantification not powerful enough.

Solution

Hindsight reasoning/temporal interpolation.



Hindsight Reasoning [O'Hearn et al. 2010; Lev-Ari et al. 2015; Feldman et al. 2018, 2020]

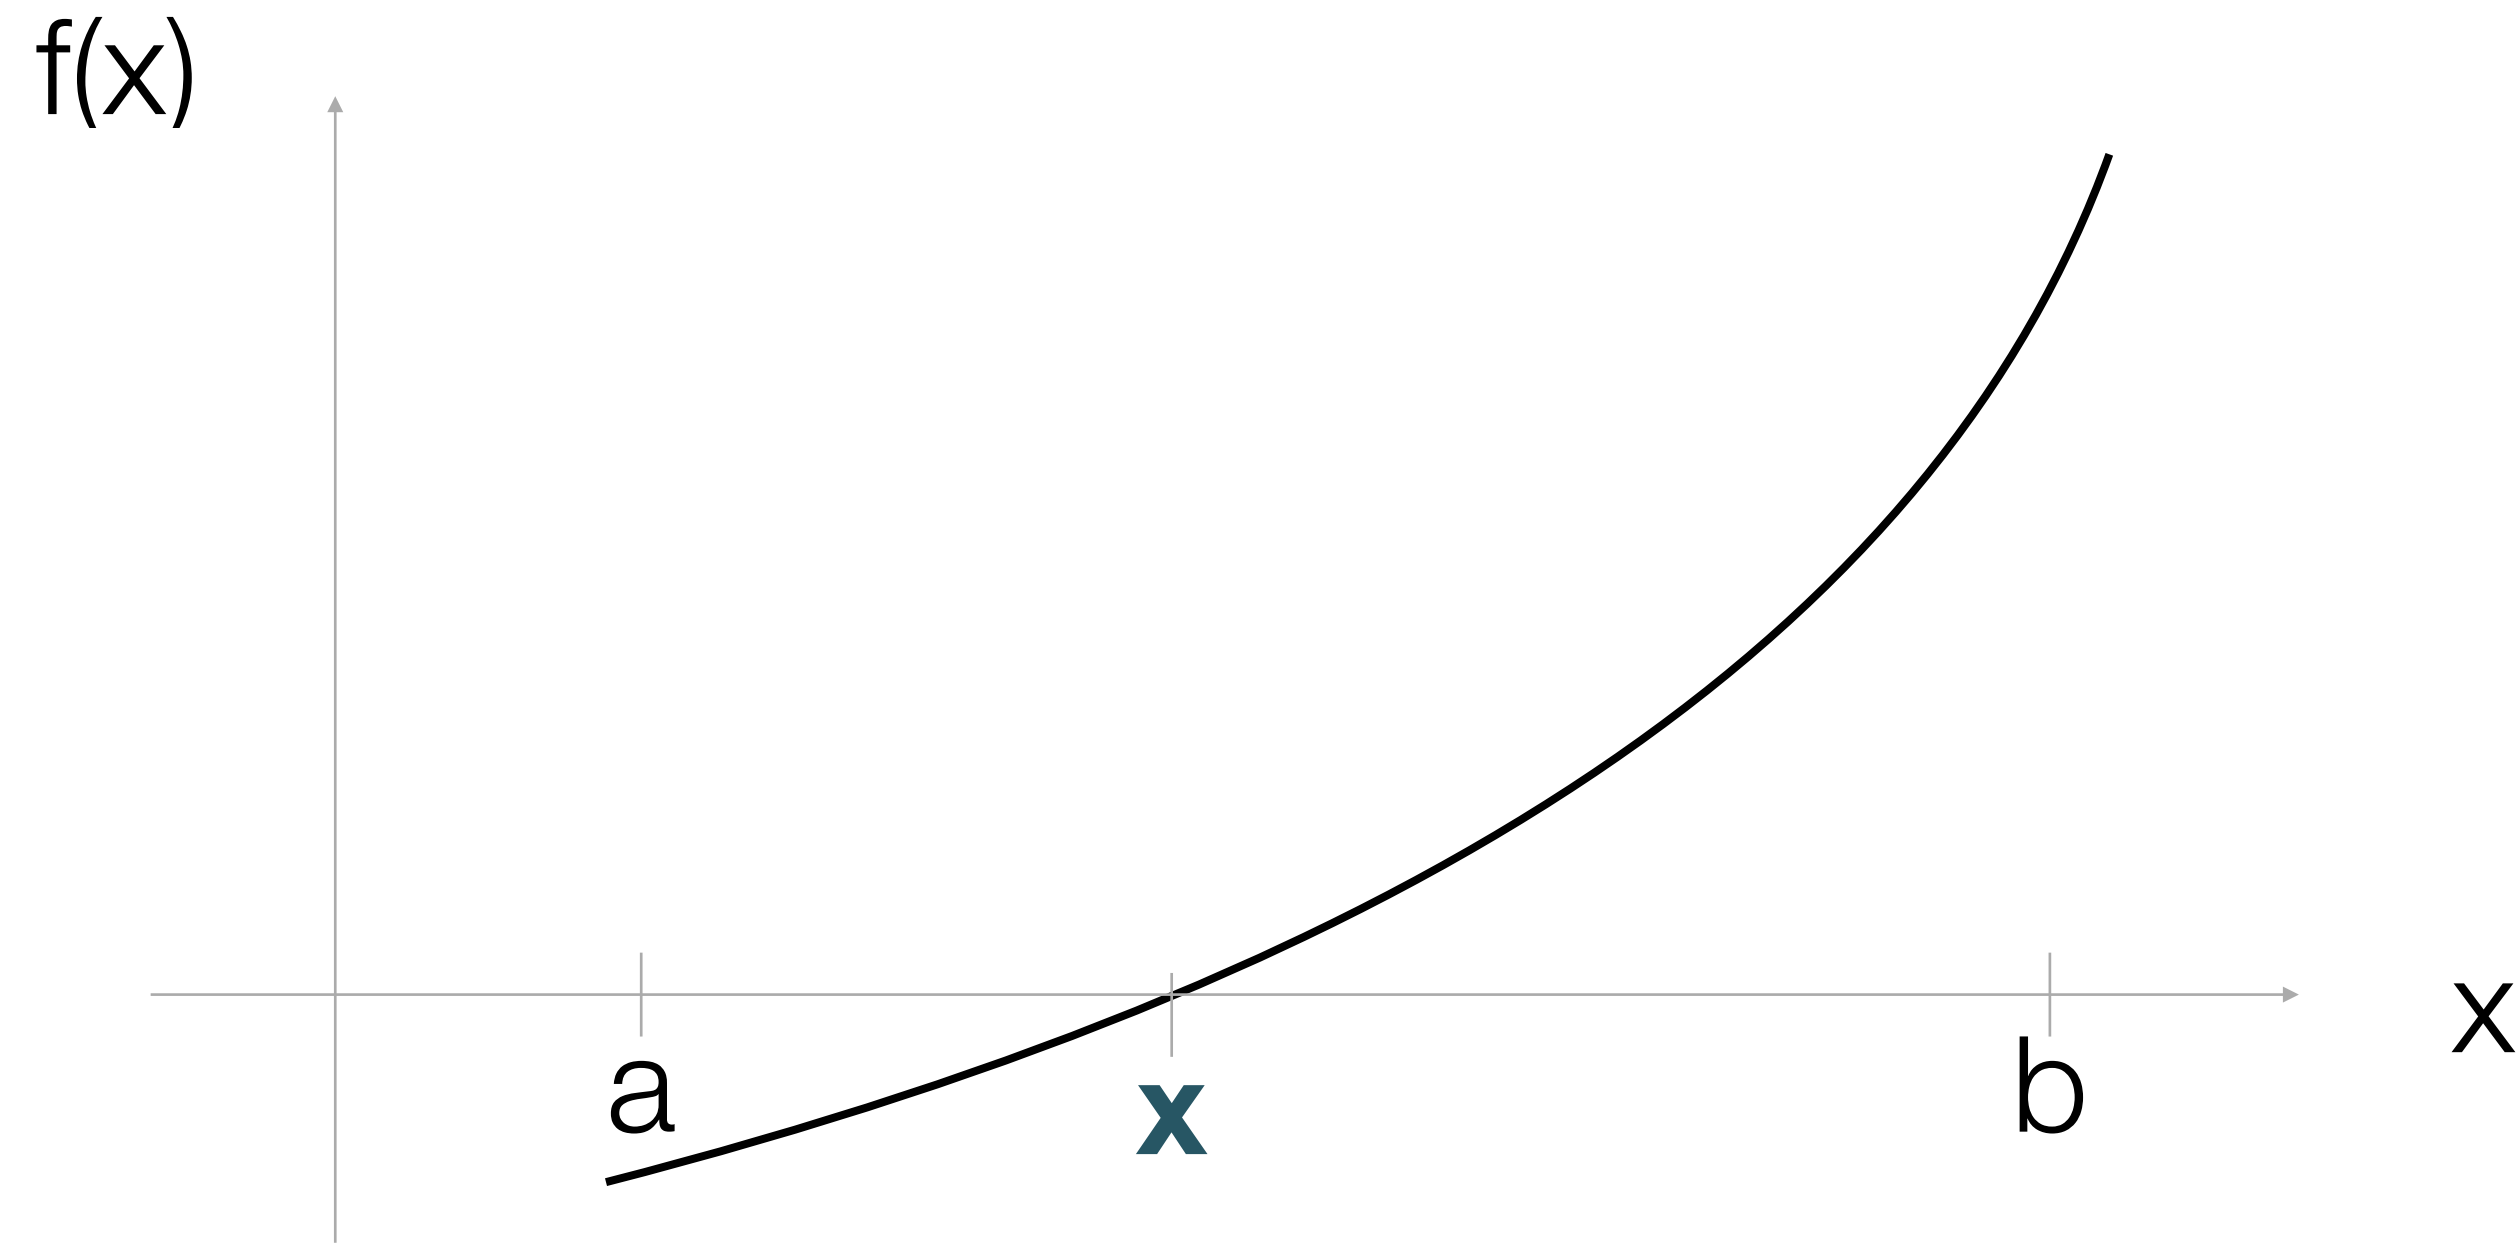


Nullstellensatz (Bolzano, 1817):

Let f be continuous on $[a, b]$, $f(a) < 0$, and $f(b) > 0$.

Then **there is x** in $[a, b]$ with $f(x) = 0$

Hindsight Reasoning [O'Hearn et al. 2010; Lev-Ari et al. 2015; Feldman et al. 2018, 2020]



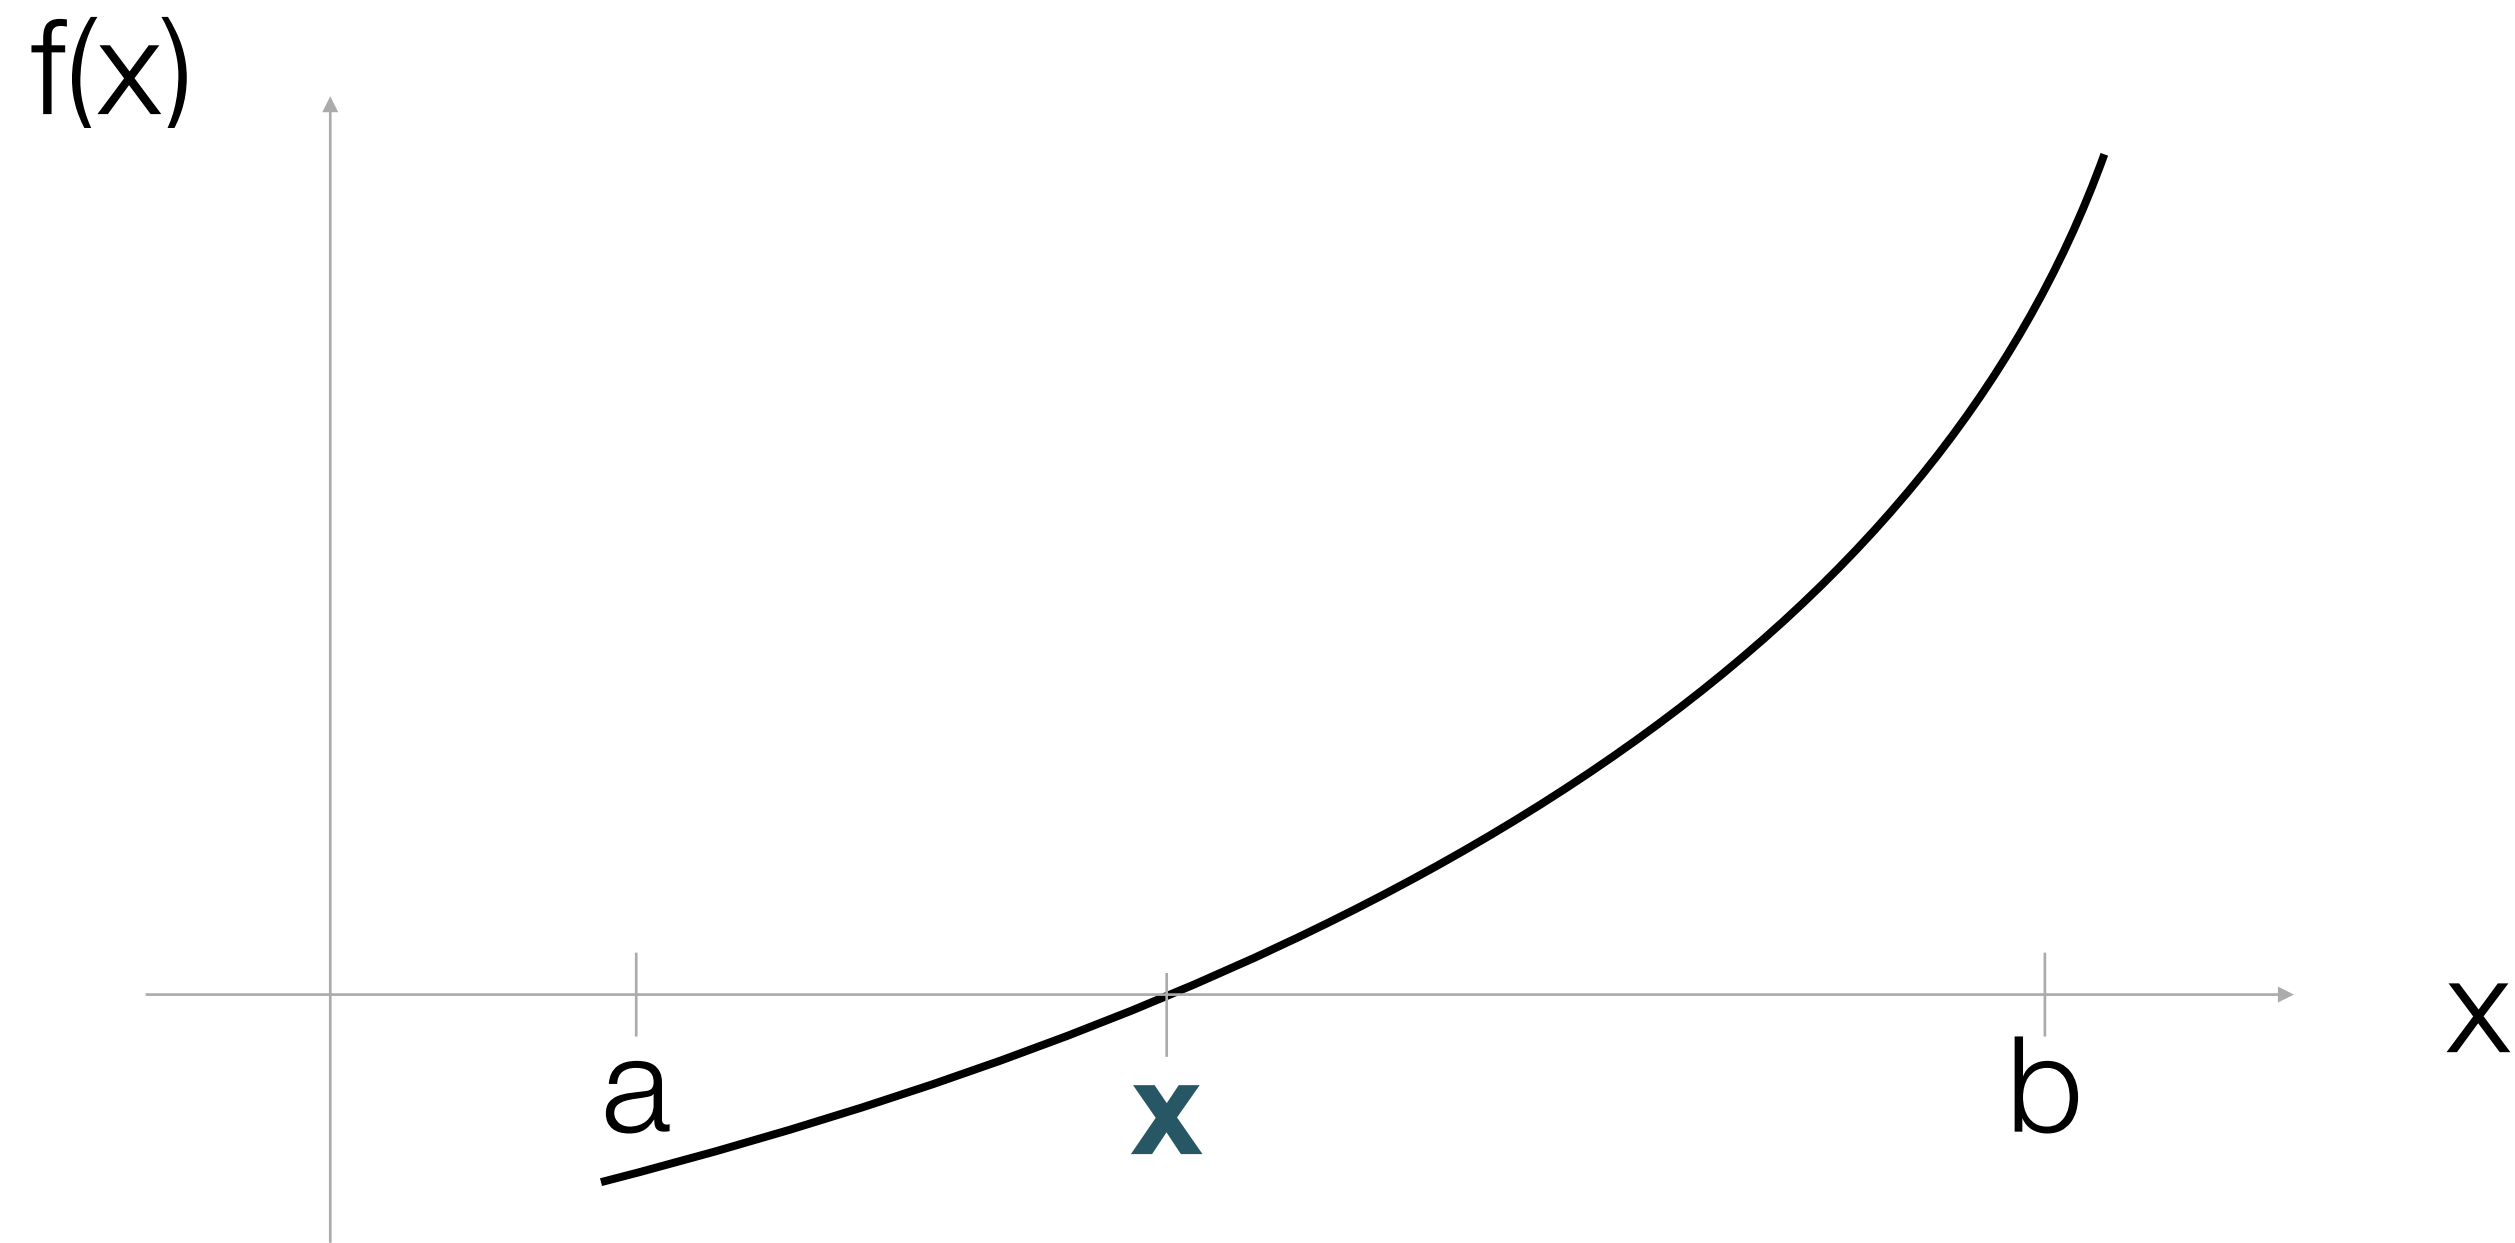
Nullstellensatz (Bolzano, 1817):

Let f be continuous on $[a, b]$, $f(a) < 0$, and $f(b) > 0$.

Then **there is x** in $[a, b]$ with $f(x) = 0$

— the **linearization point**.

Hindsight Reasoning [O'Hearn et al. 2010; Lev-Ari et al. 2015; Feldman et al. 2018, 2020]



- record how computation unfolds
- derive what **must have happened**
- reflects programming practice

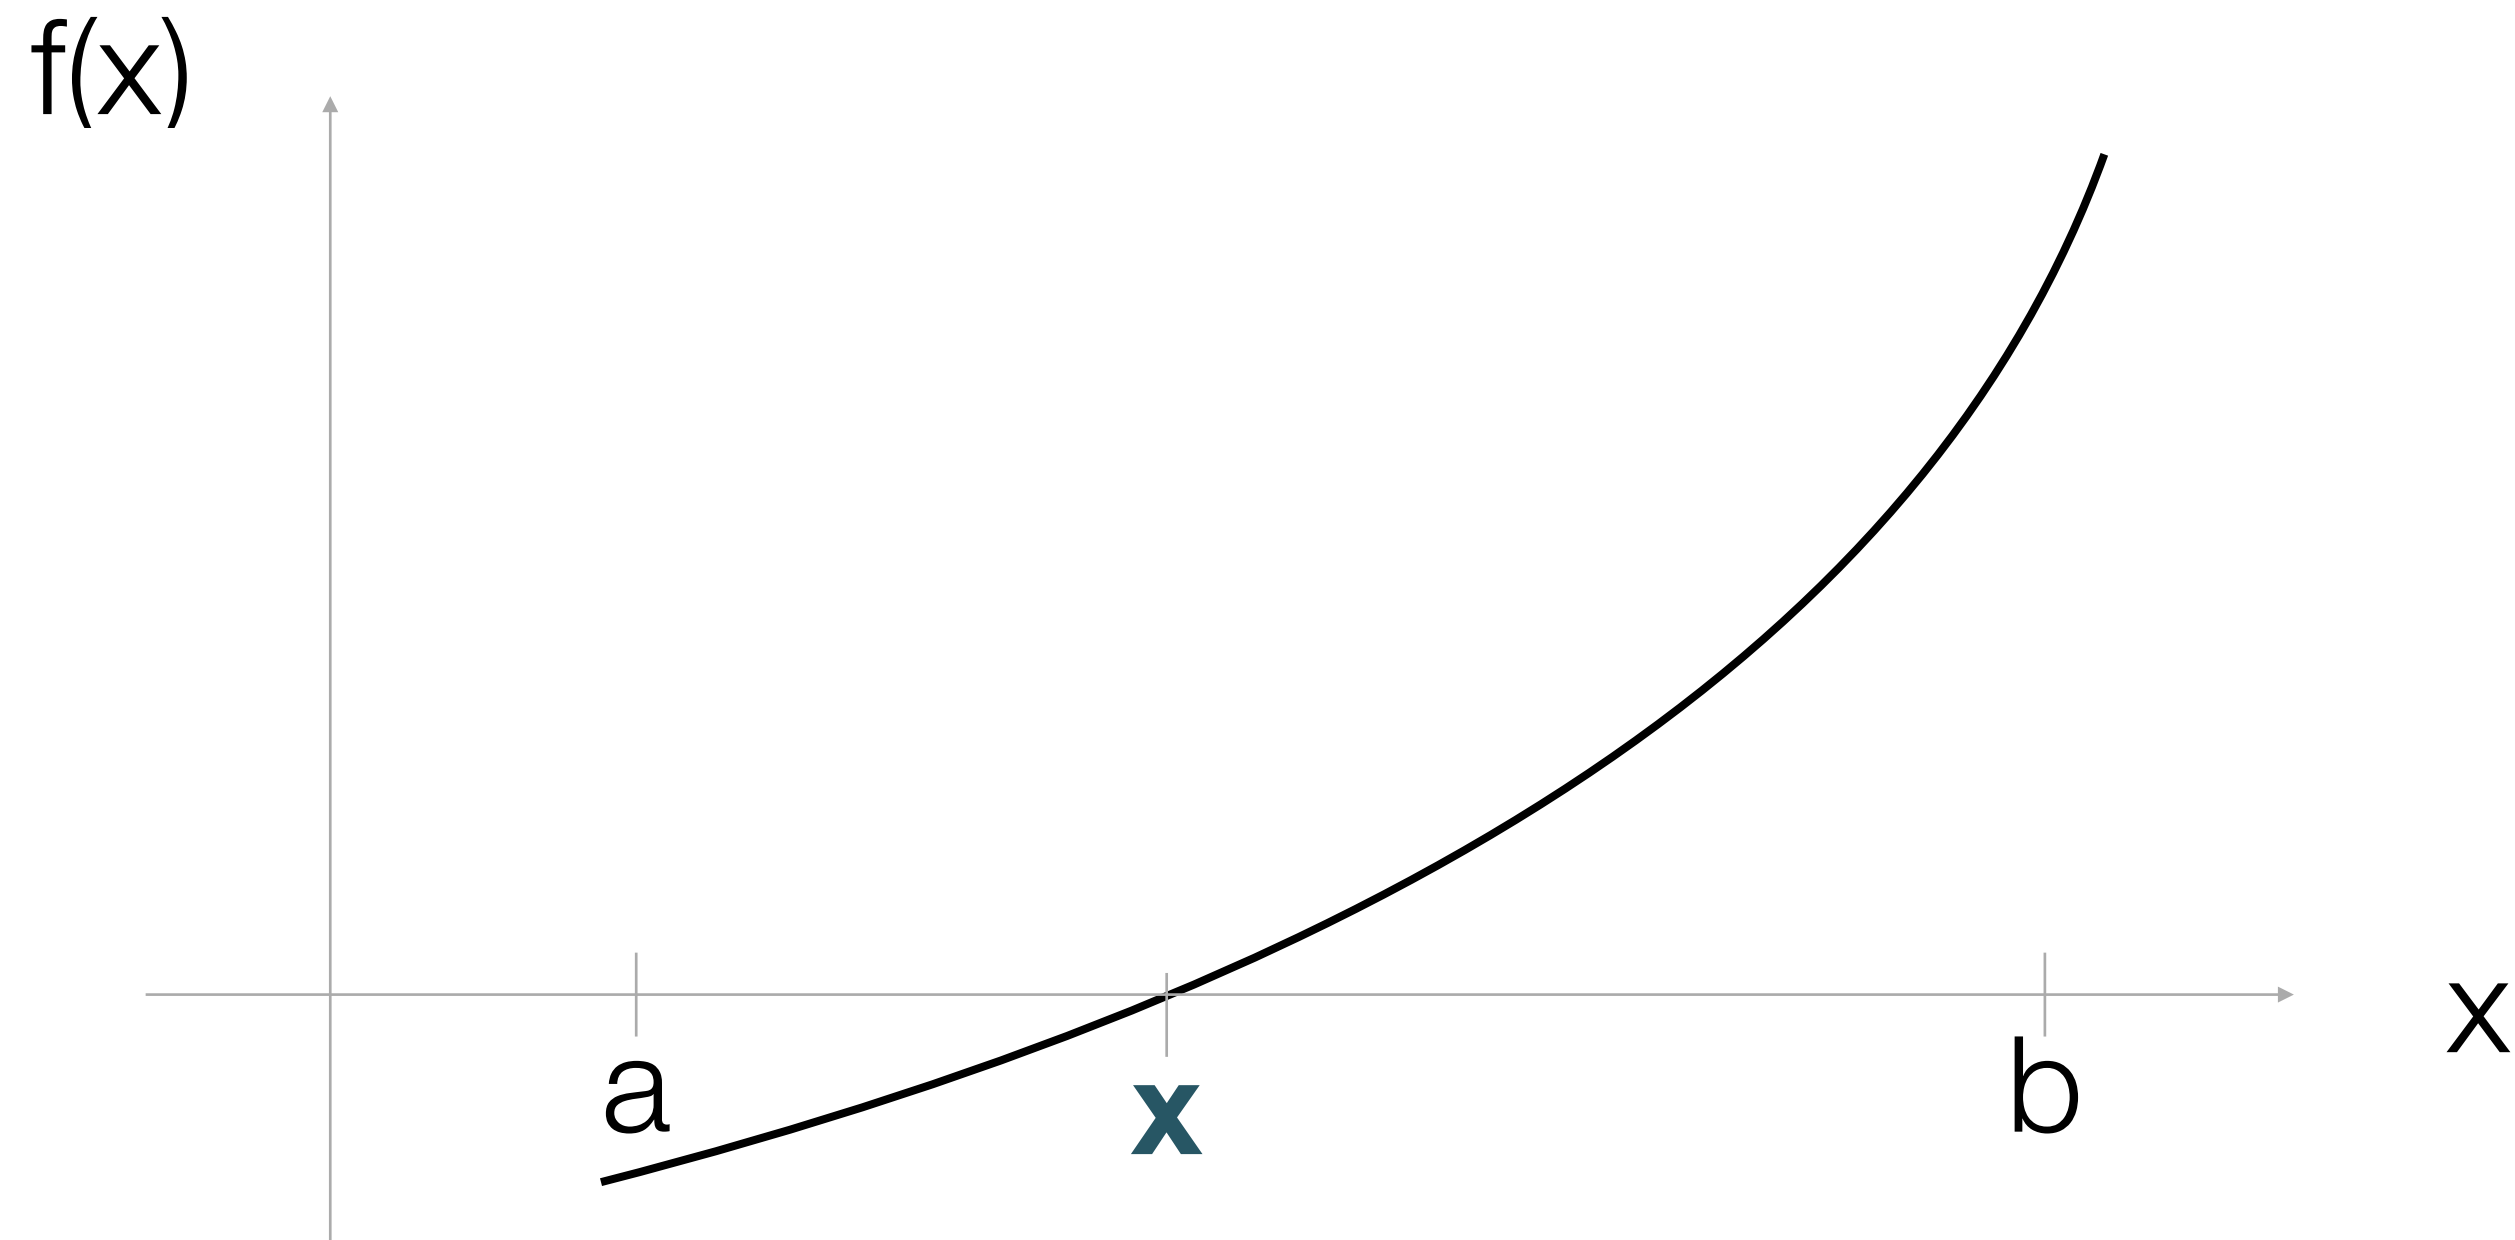
Nullstellensatz (Bolzano, 1817):

Let f be continuous on $[a, b]$, $f(a) < 0$, and $f(b) > 0$.

Then **there is x** in $[a, b]$ with $f(x) = 0$

— the **linearization point**.

Hindsight Reasoning [O'Hearn et al. 2010; Lev-Ari et al. 2015; Feldman et al. 2018, 2020]



- record how computation unfolds
- derive what **must have happened**
- reflects programming practice

**lazy approach,
easier to automate**

Nullstellensatz (Bolzano, 1817):

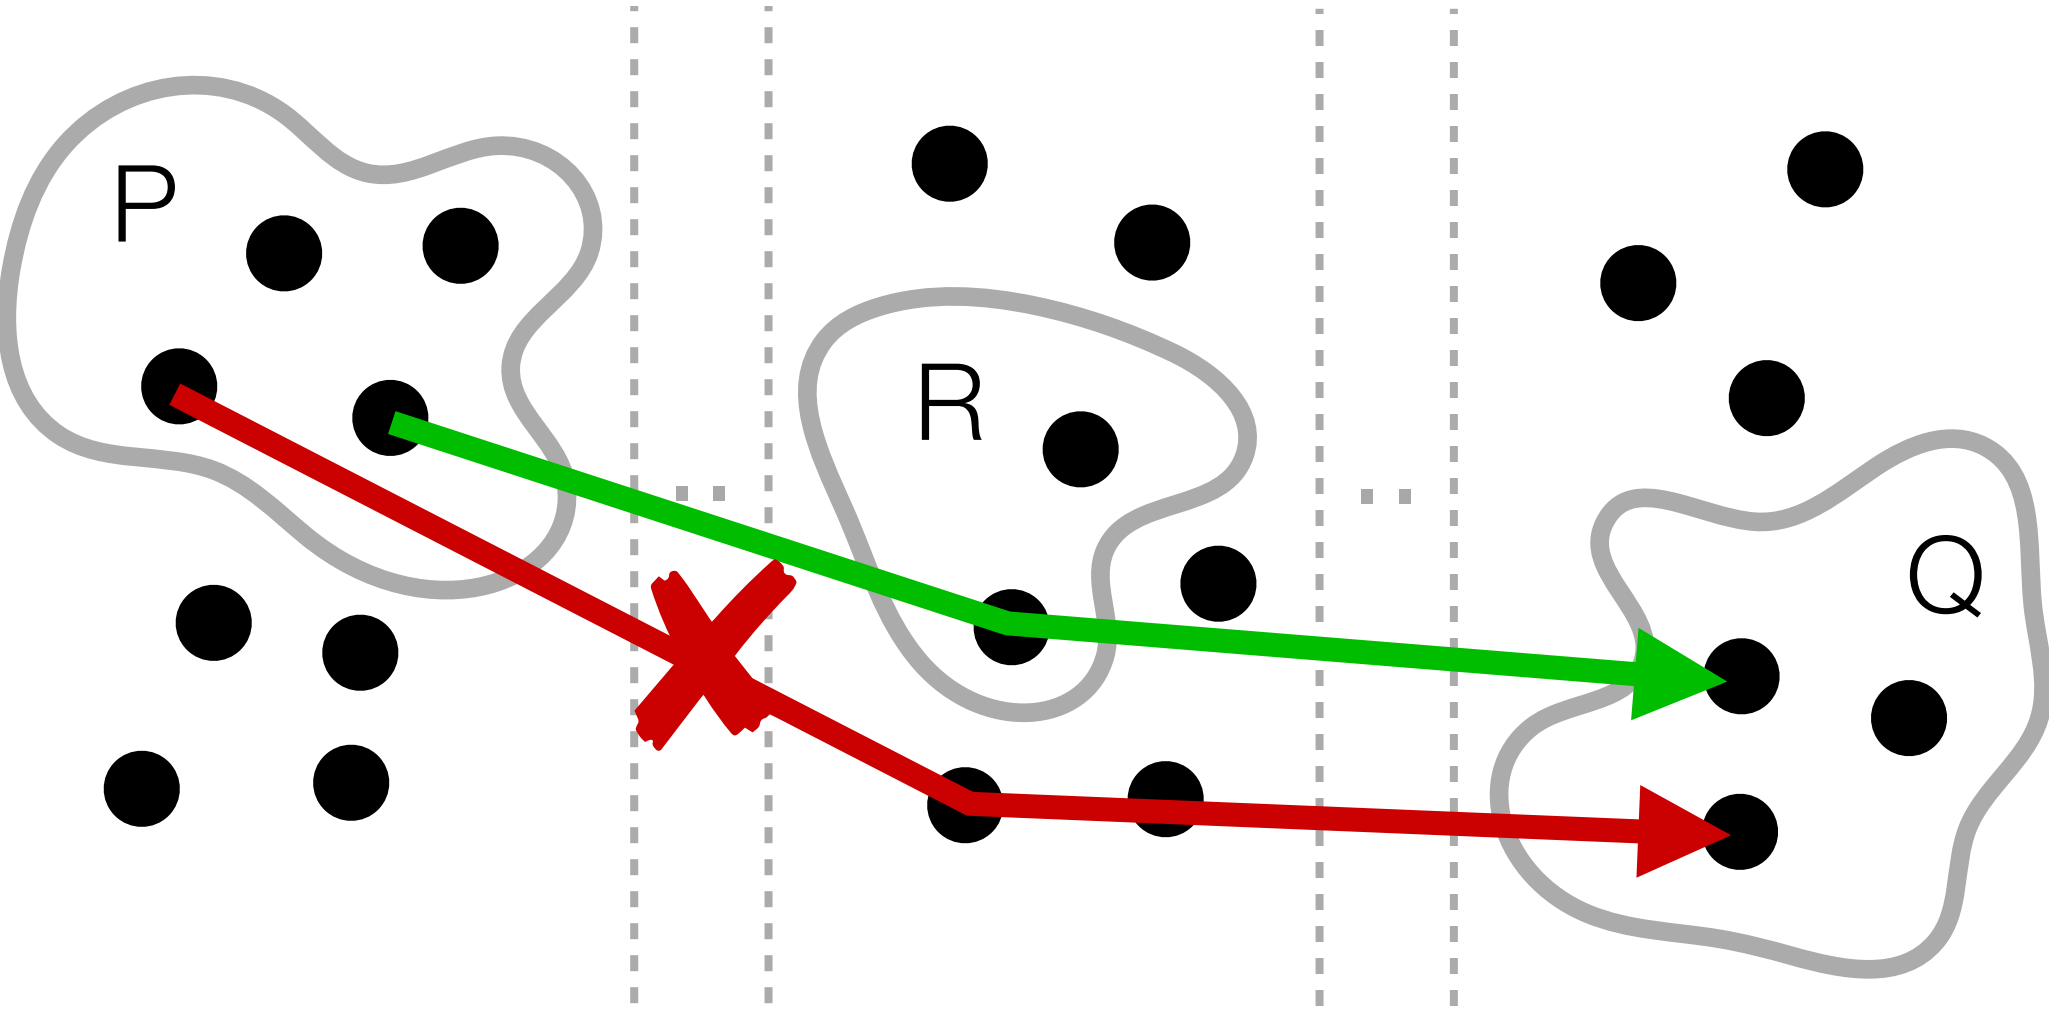
Let f be continuous on $[a, b]$, $f(a) < 0$, and $f(b) > 0$.

Then **there is x** in $[a, b]$ with $f(x) = 0$

— the **linearization point**.

Hindsight Reasoning

- Record execution history **P** Past
- not just current state **Q** Now



Hindsight Reasoning

- Record execution history P Past

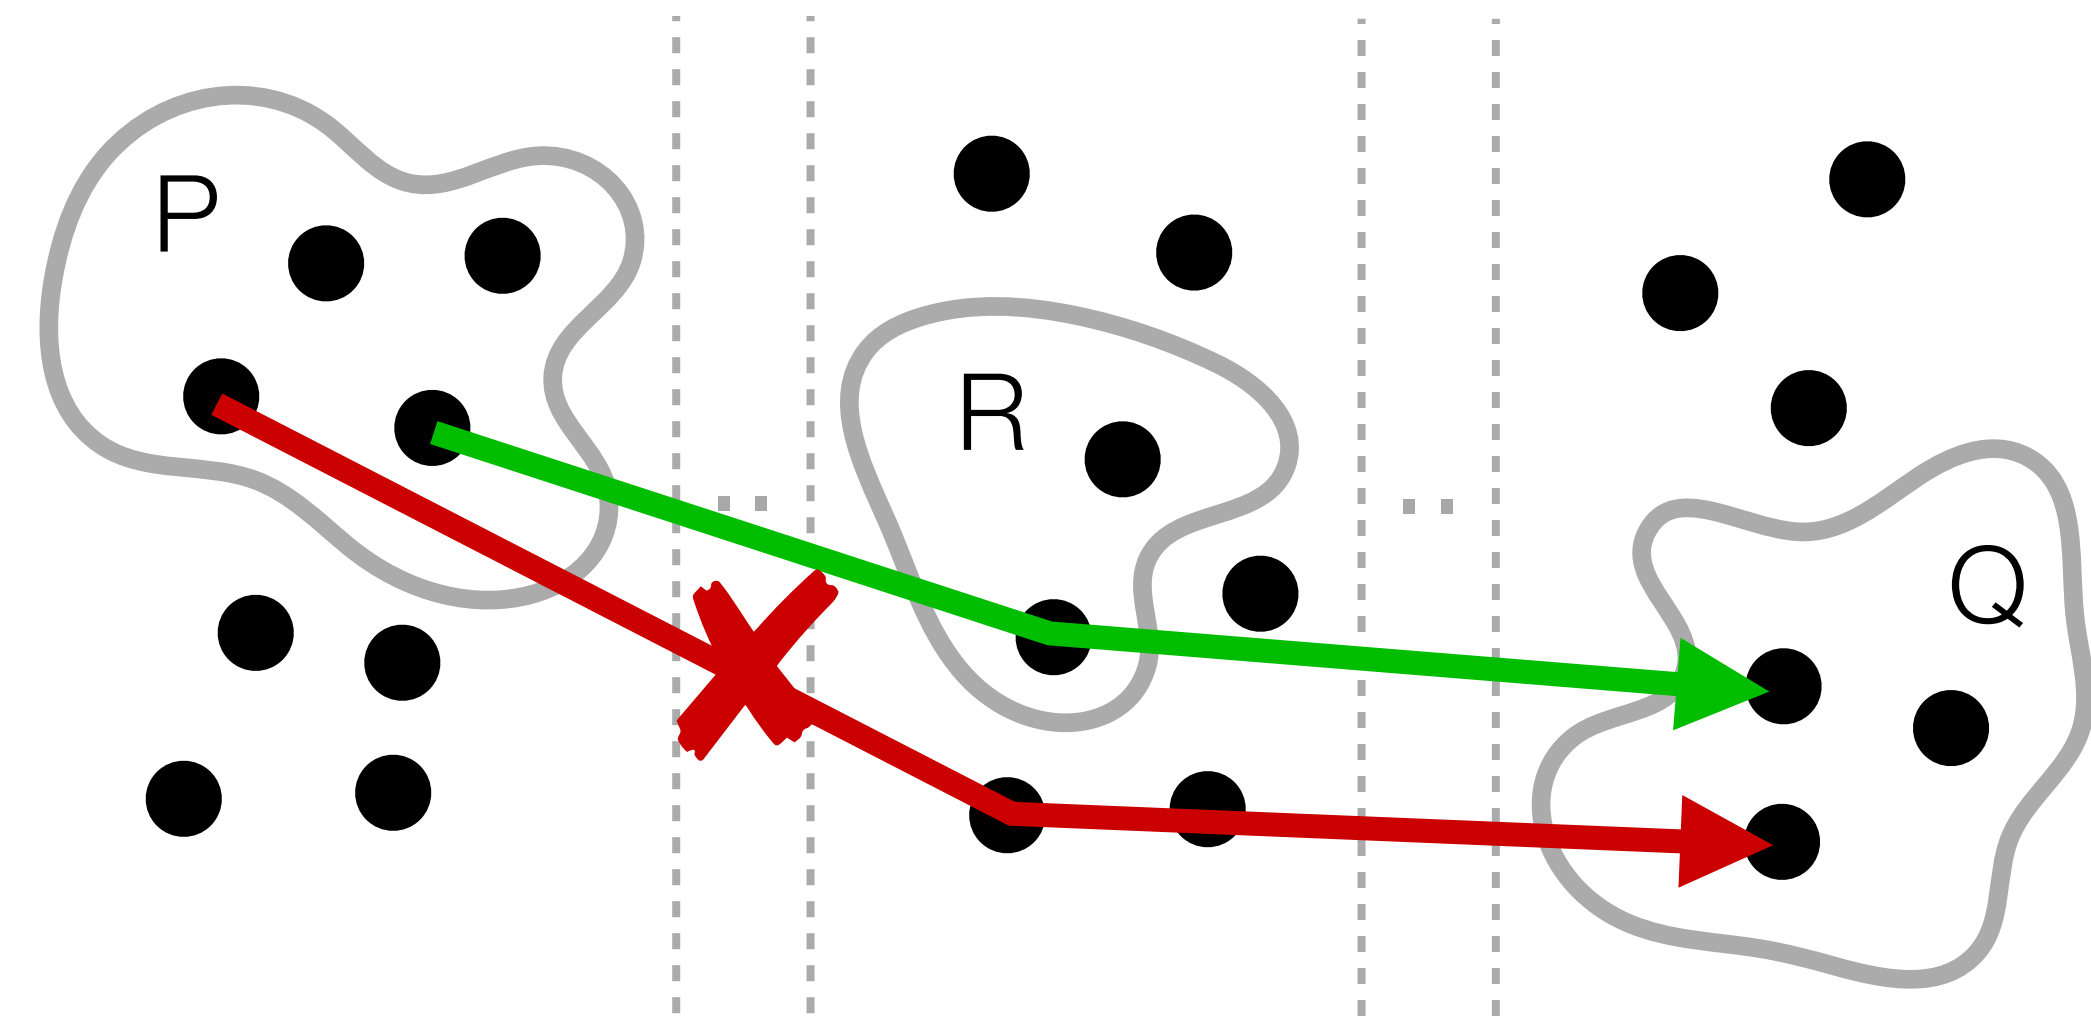
not just current state Q Now

- Derive **inevitable** facts:



if **all paths** from P to Q

pass through R .



Hindsight Reasoning

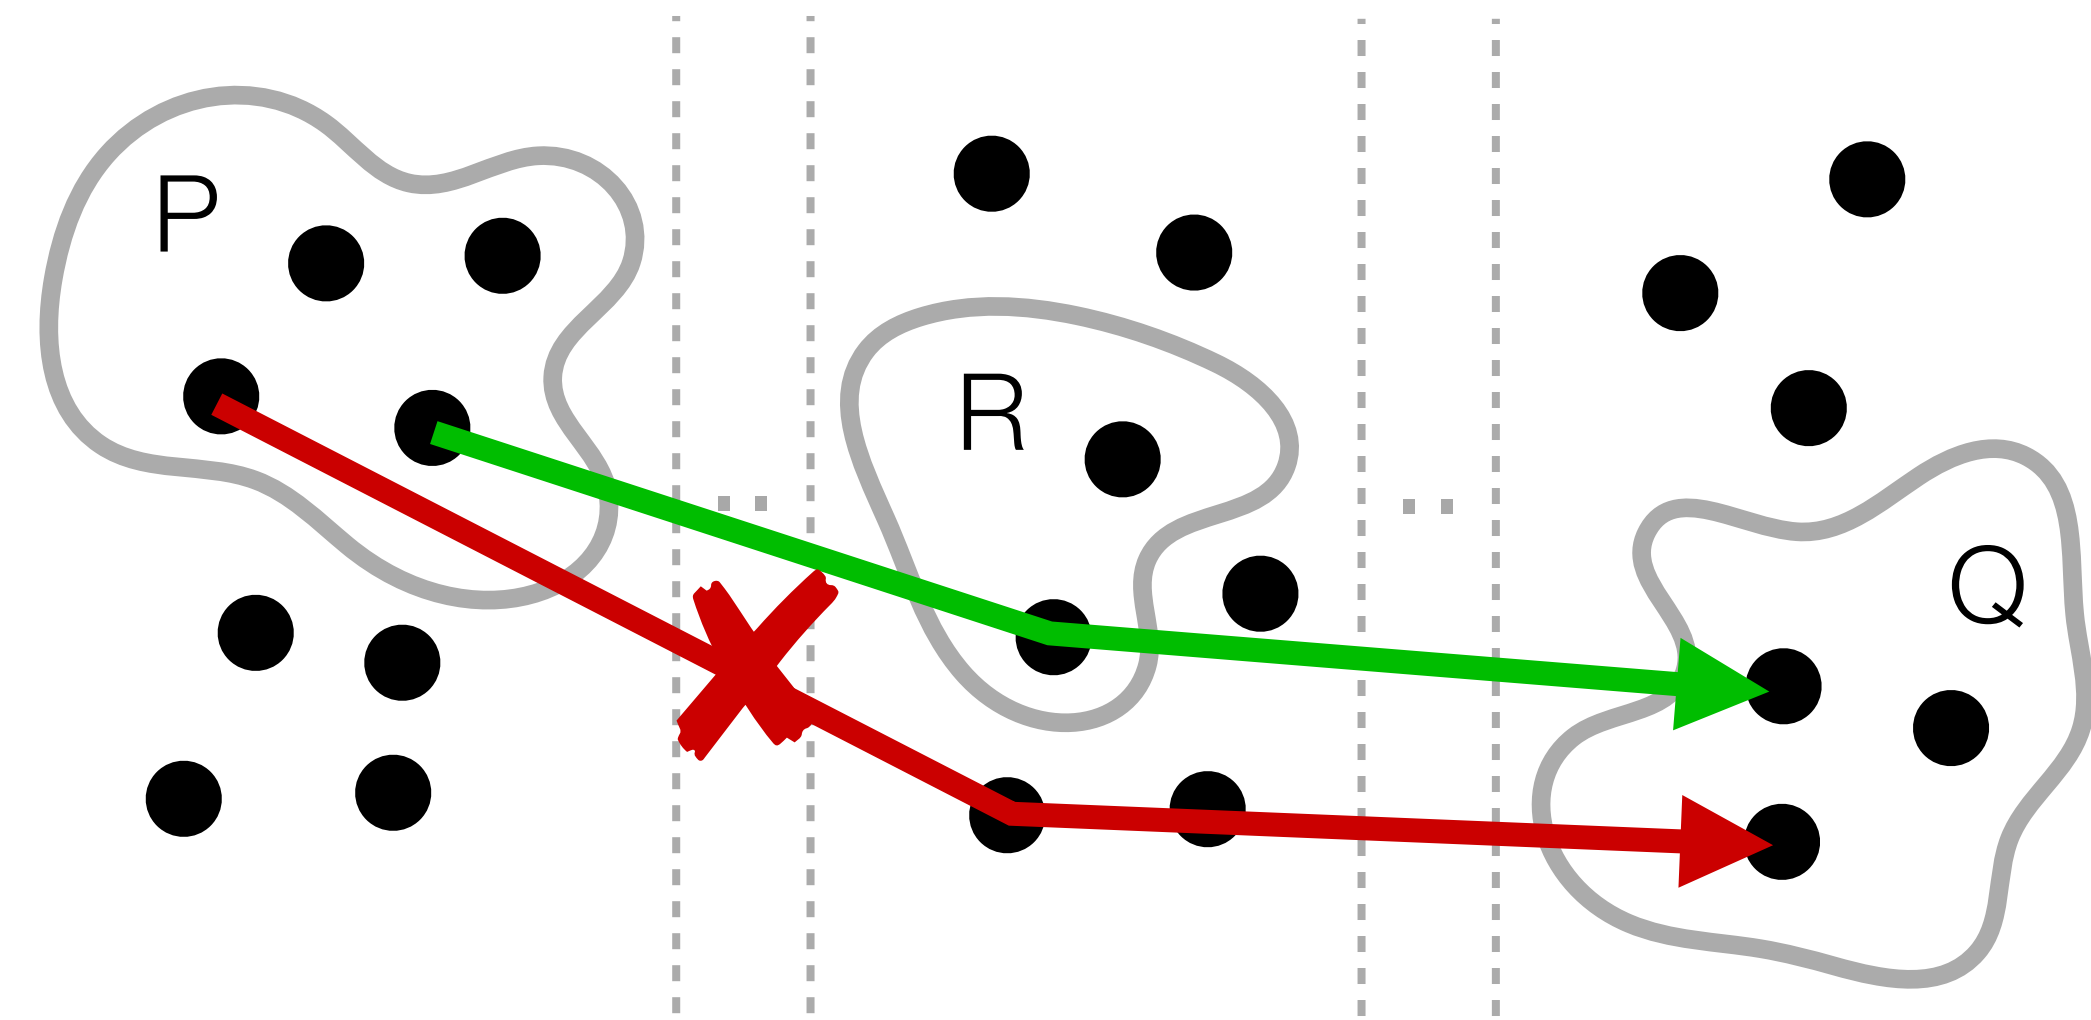
- Record execution history P Past

not just current state Q Now

- Derive **inevitable** facts:

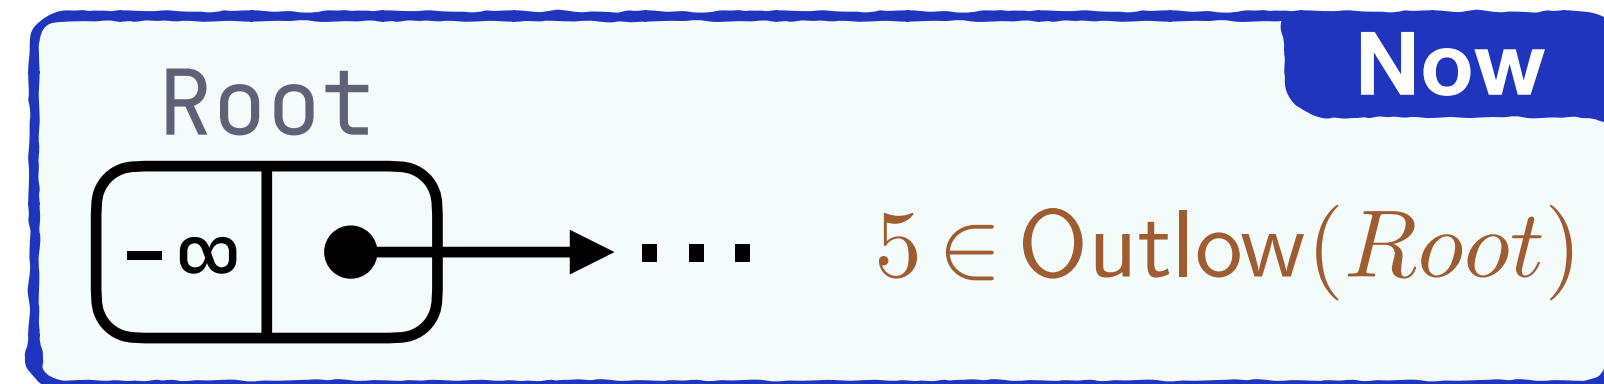
$$\boxed{P \text{ Past}} \wedge \boxed{Q \text{ Now}} \rightsquigarrow \boxed{R \text{ Past}}$$

if **all paths** from P to Q
pass through R .



formed from **interferences**
in the Owicki-Gries proof

Example: contains(5)



```
x = Root→next;
```

```
while (x→key < 5):
```

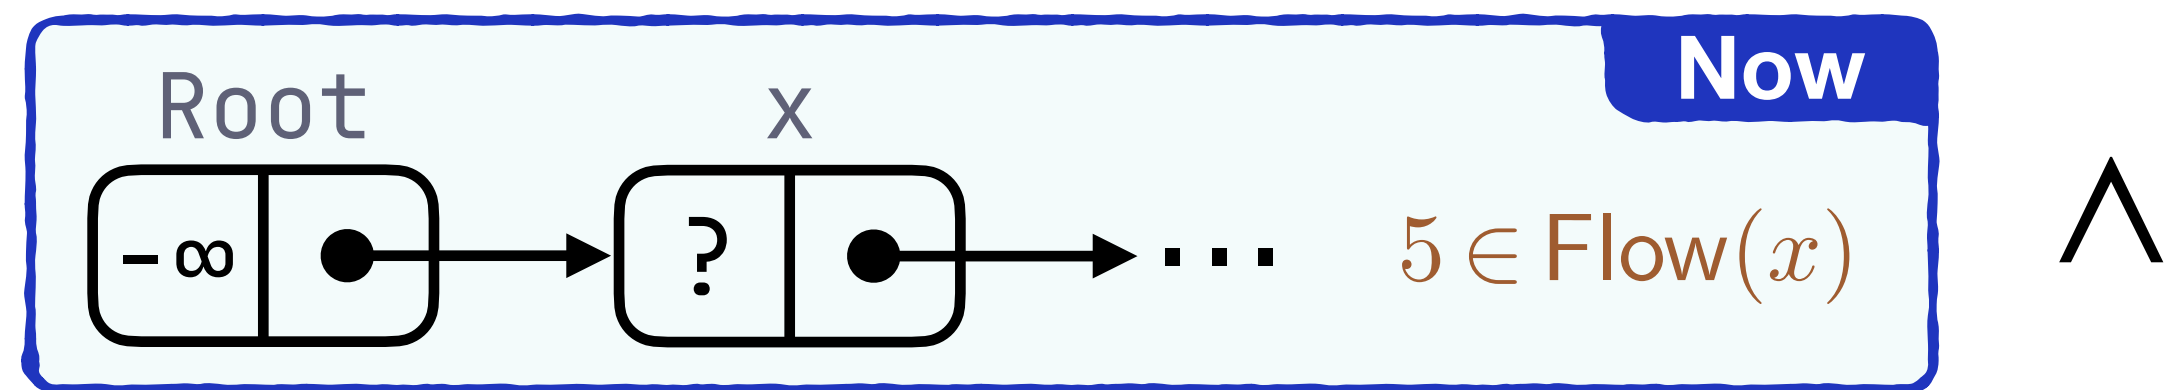
```
    y = x→next;
```

```
    x = y;
```

```
return x→key == 5;
```


Example: contains(5)

```
x = Root→next;
```



```
while (x→key < 5):
```

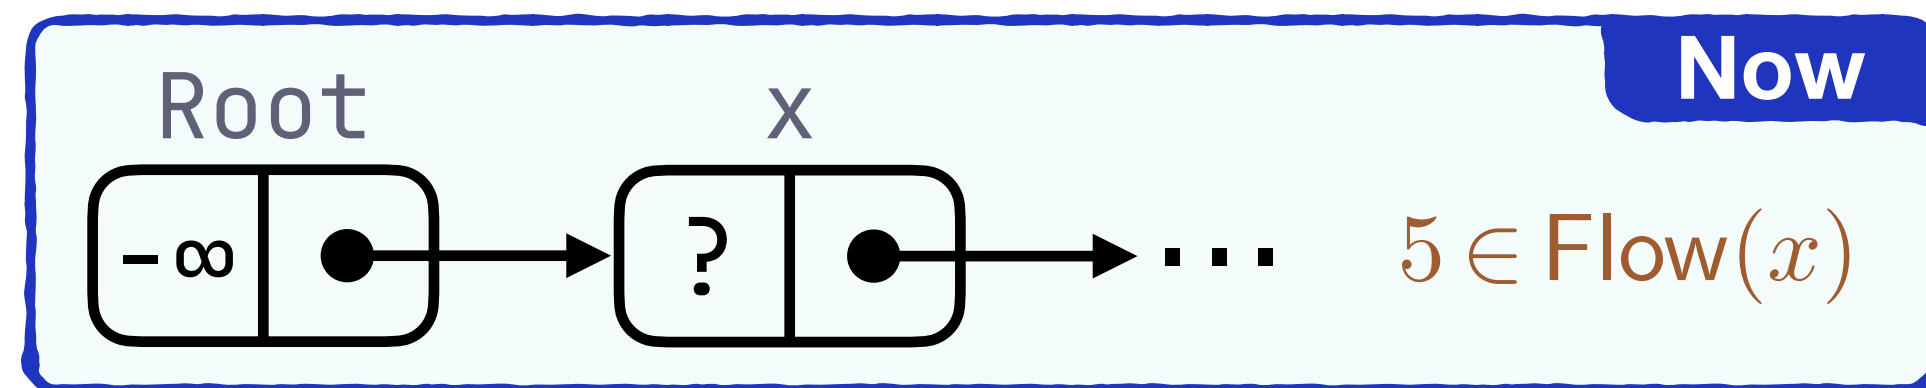
```
    y = x→next;
```

```
    x = y;
```

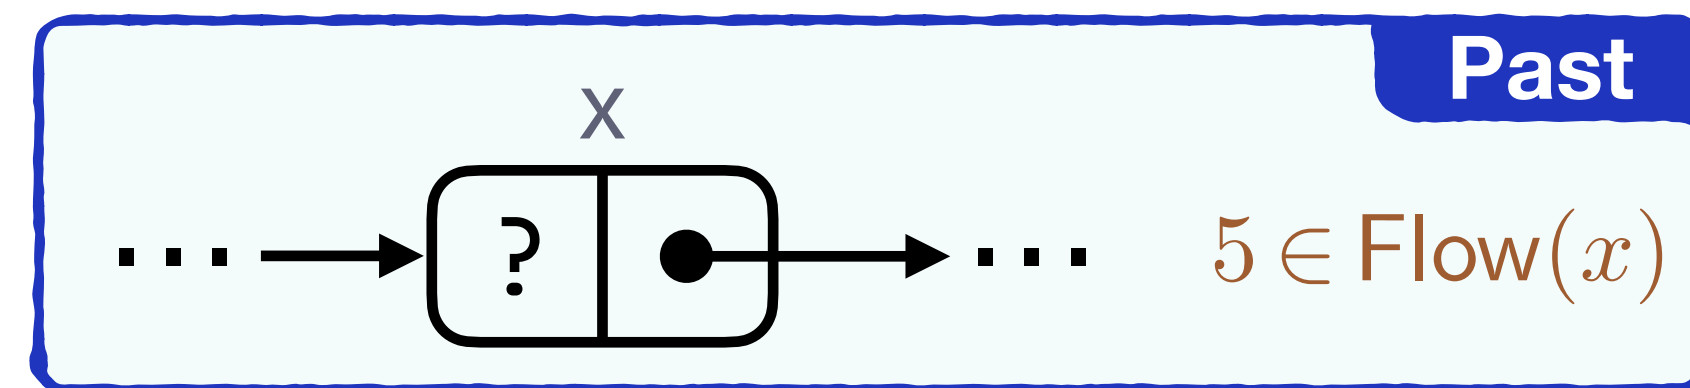
```
return x→key == 5;
```

Example: contains(5)

```
x = Root→next;
```



\wedge



```
while (x→key < 5):
```

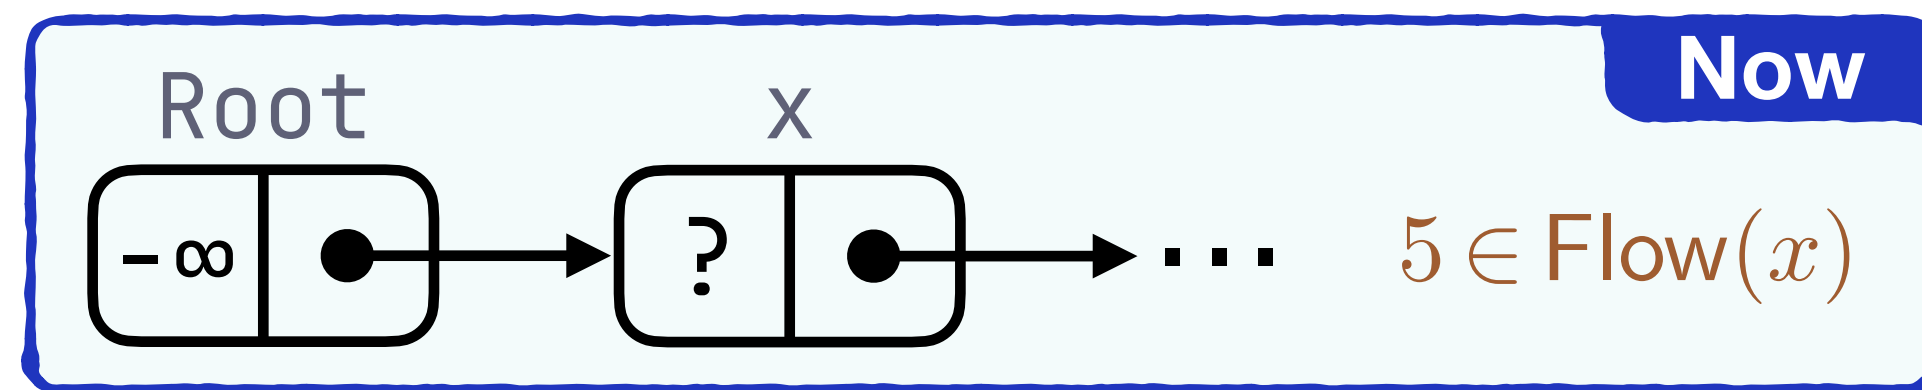
```
    y = x→next;
```

```
    x = y;
```

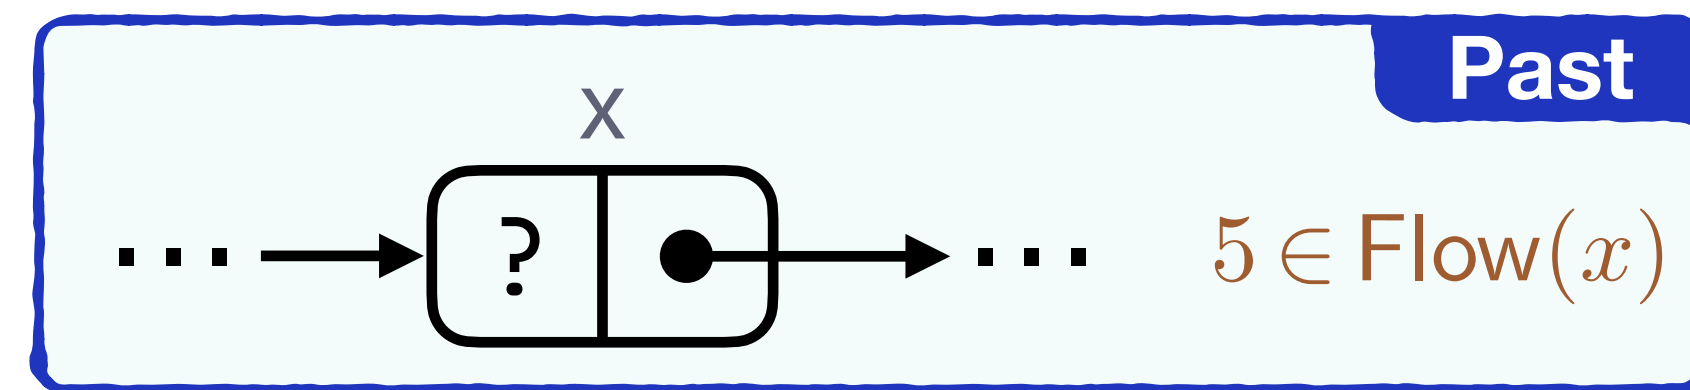
```
return x→key == 5;
```

Example: contains(5)

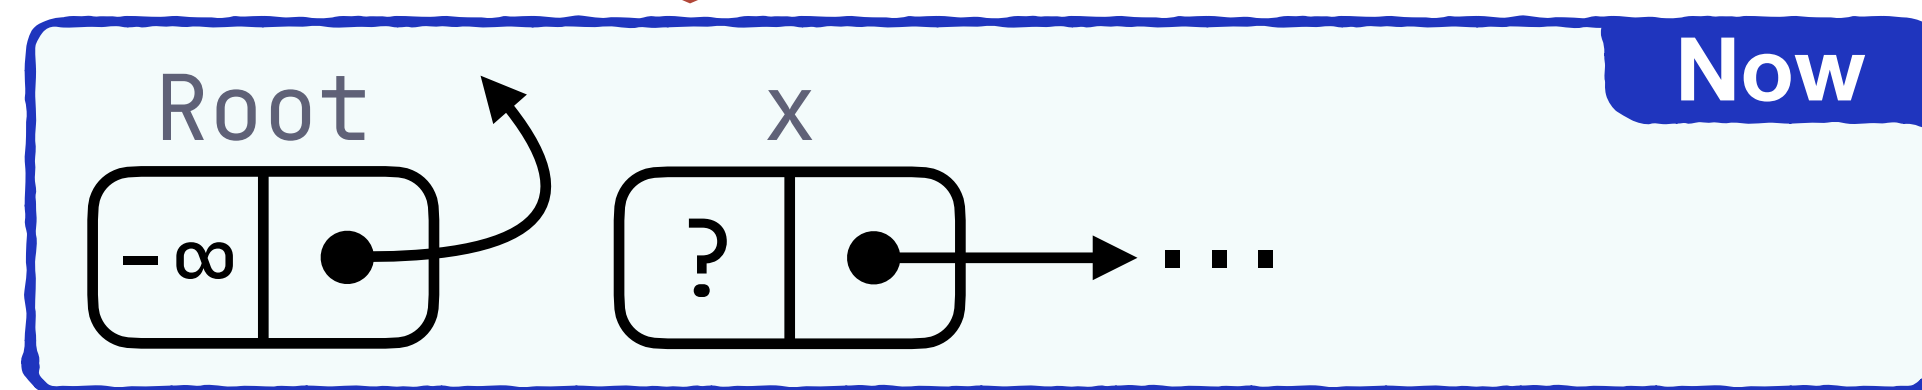
```
x = Root→next;
```



\wedge



interference



\wedge

```
while (x→key < 5):
```

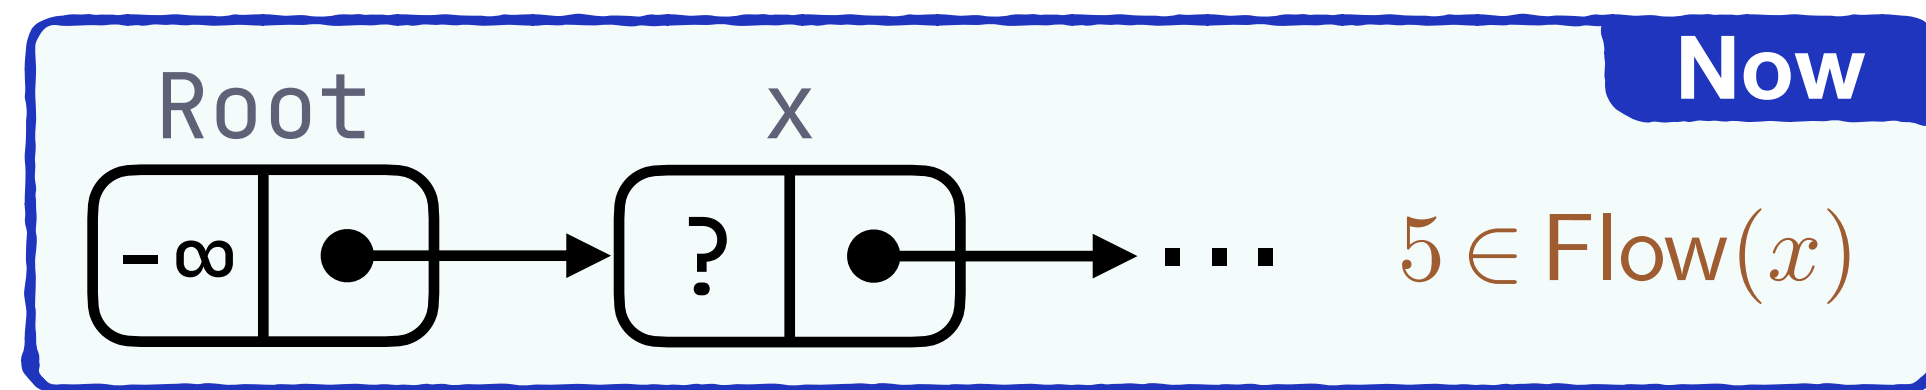
```
    y = x→next;
```

```
    x = y;
```

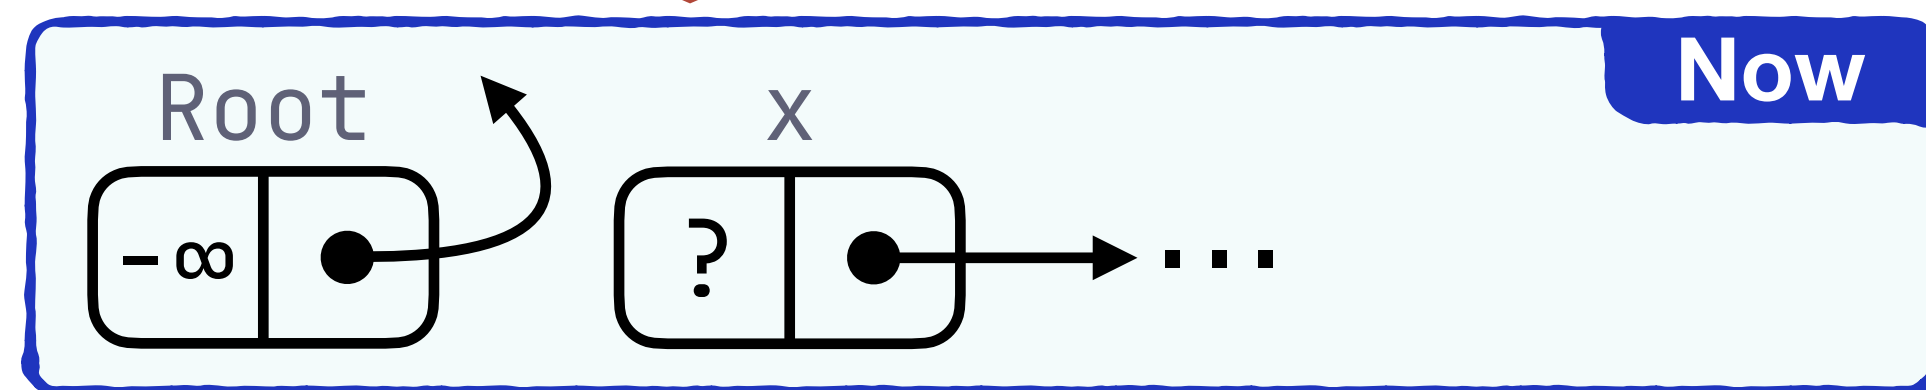
```
return x→key == 5;
```

Example: contains(5)

```
x = Root→next;
```



interference



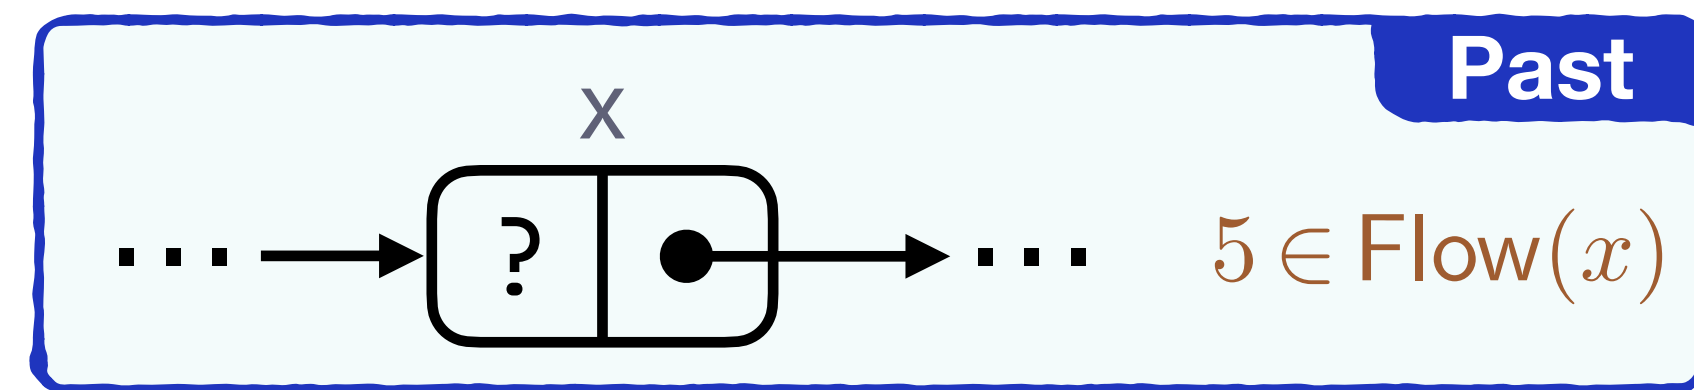
```
while (x→key < 5):
```

```
    y = x→next;
```

```
    x = y;
```

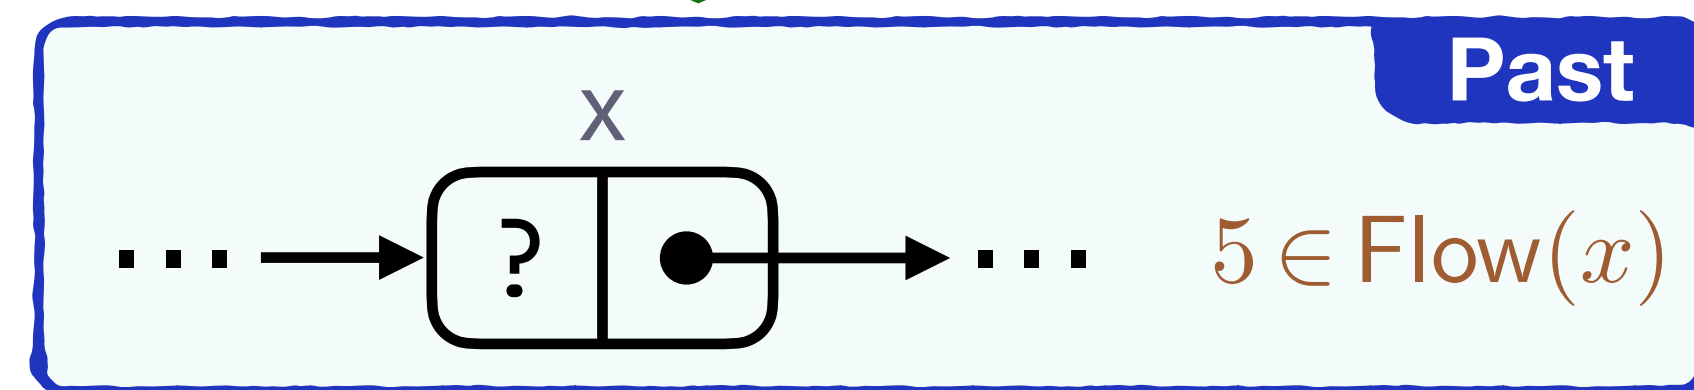
```
return x→key == 5;
```

\wedge



unchanged

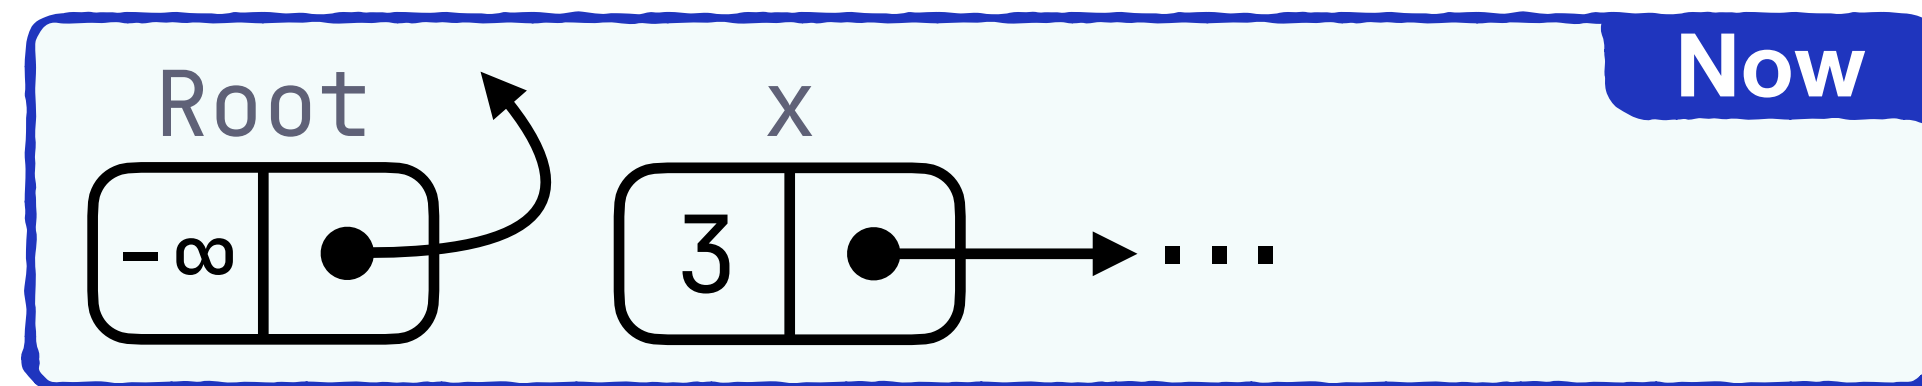
\wedge



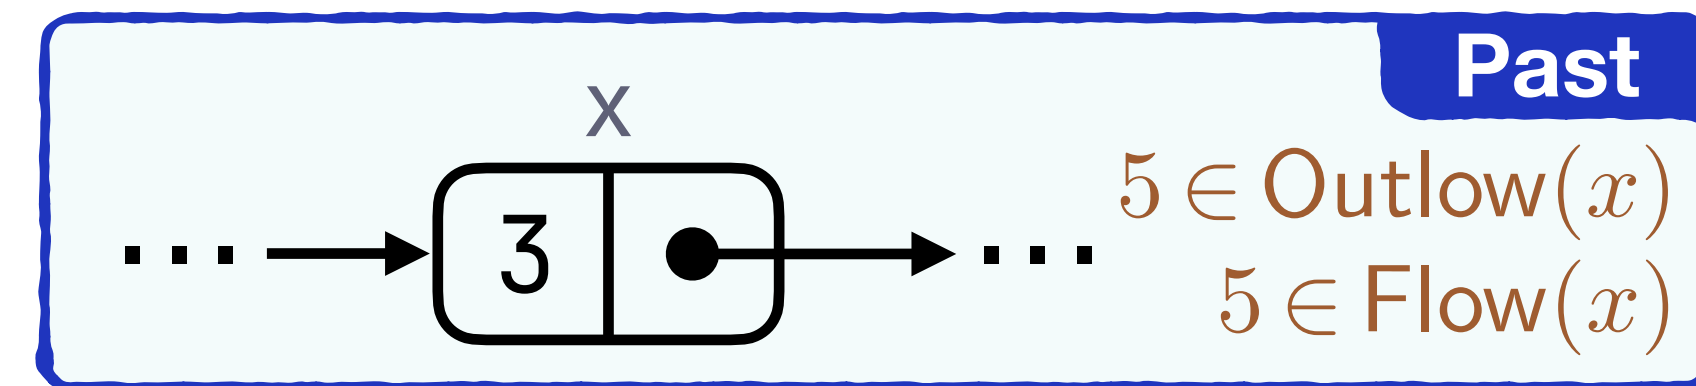
Example: contains(5)

```
x = Root→next;
```

```
while (x→key < 5):
```



^



```
    y = x→next;
```

```
    x = y;
```

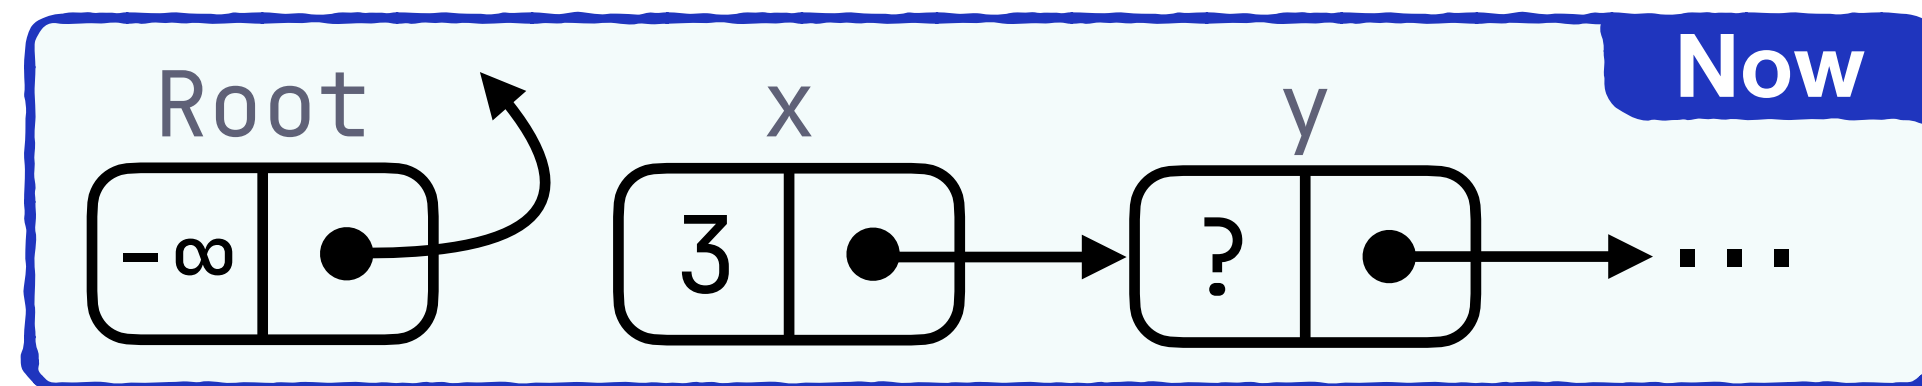
```
return x→key == 5;
```

Example: contains(5)

```
x = Root→next;
```

```
while (x→key < 5):
```

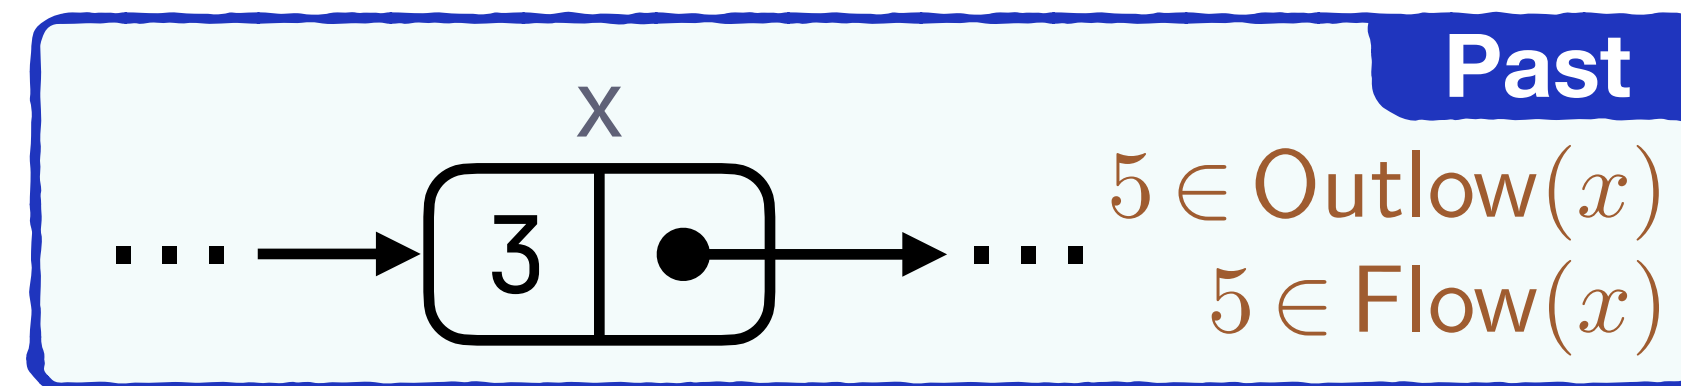
```
    y = x→next;
```



```
    x = y;
```

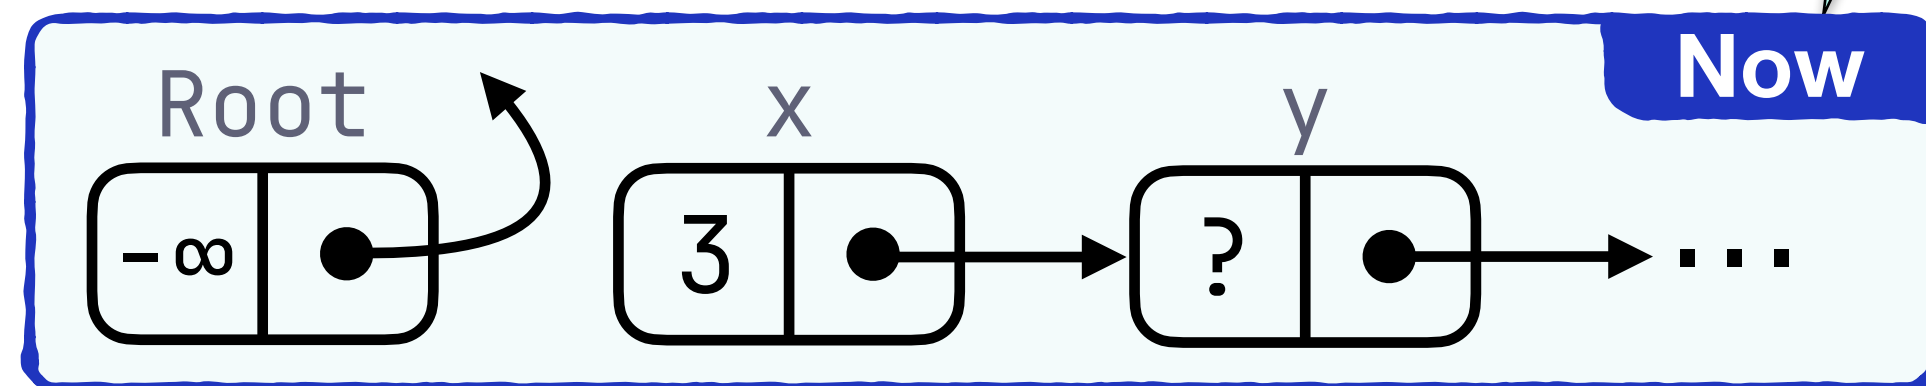
```
return x→key == 5;
```

^

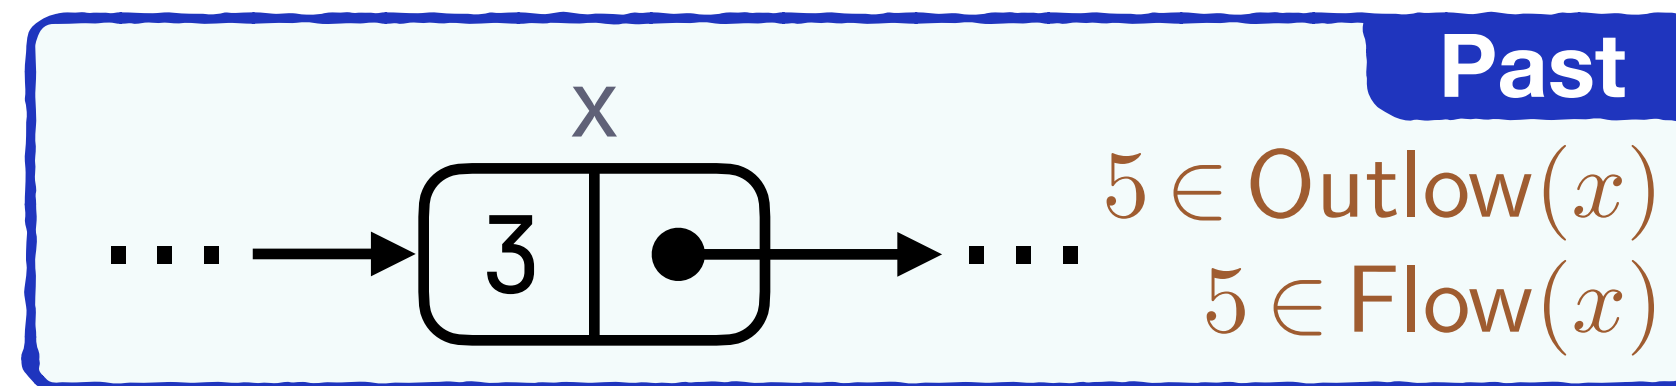


Example: contains(5)

```
x = Root→next;  
while (x→key < 5):  
    y = x→next;
```



^



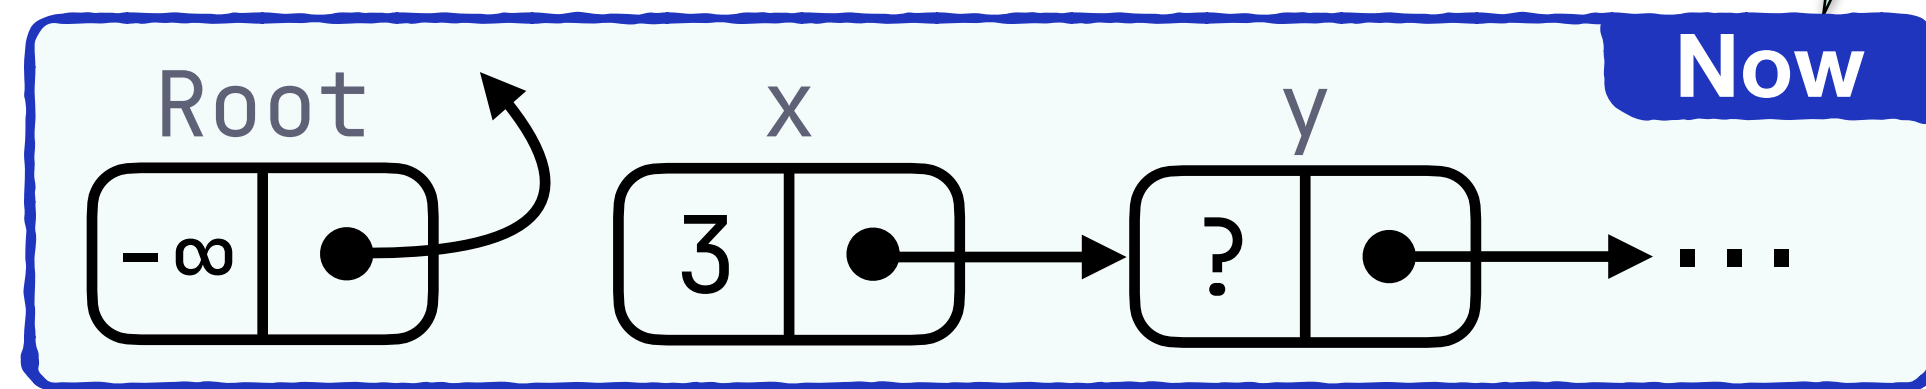
```
x = y;
```

```
return x→key == 5;
```

Traversal *learns* facts about Now, not the Past.

Example: contains(5)

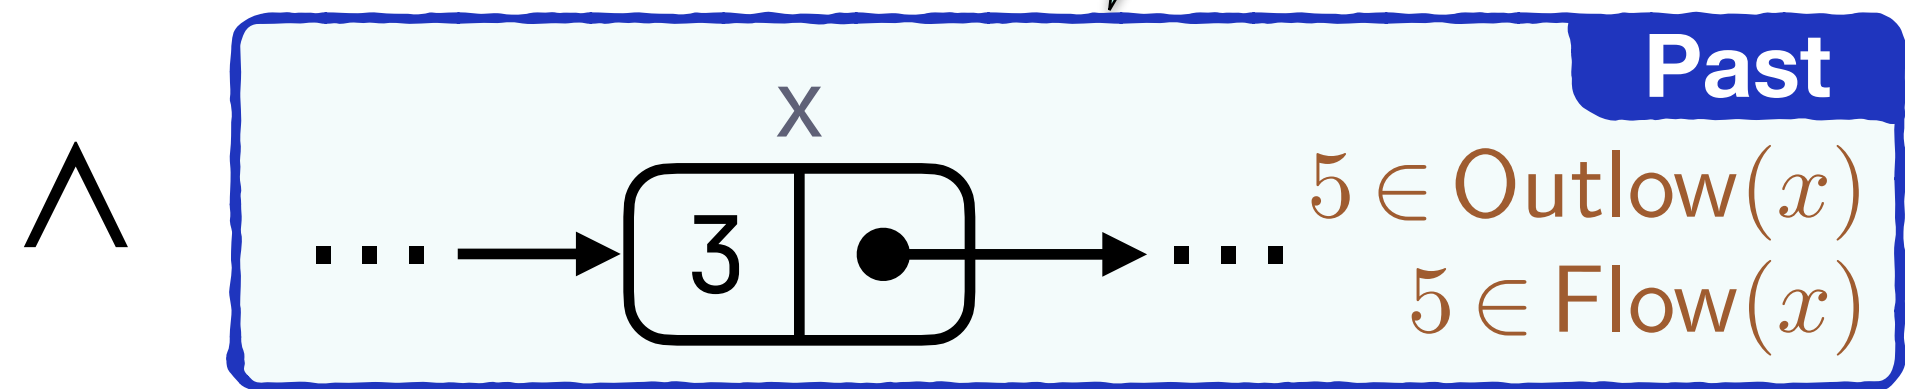
```
x = Root->next;  
while (x->key < 5):  
    y = x->next;
```



```
x = y;  
return x->key == 5;
```

Traversal *learns* facts about Now, not the Past.

Use **hindsight** to propagate Now facts into the Past.

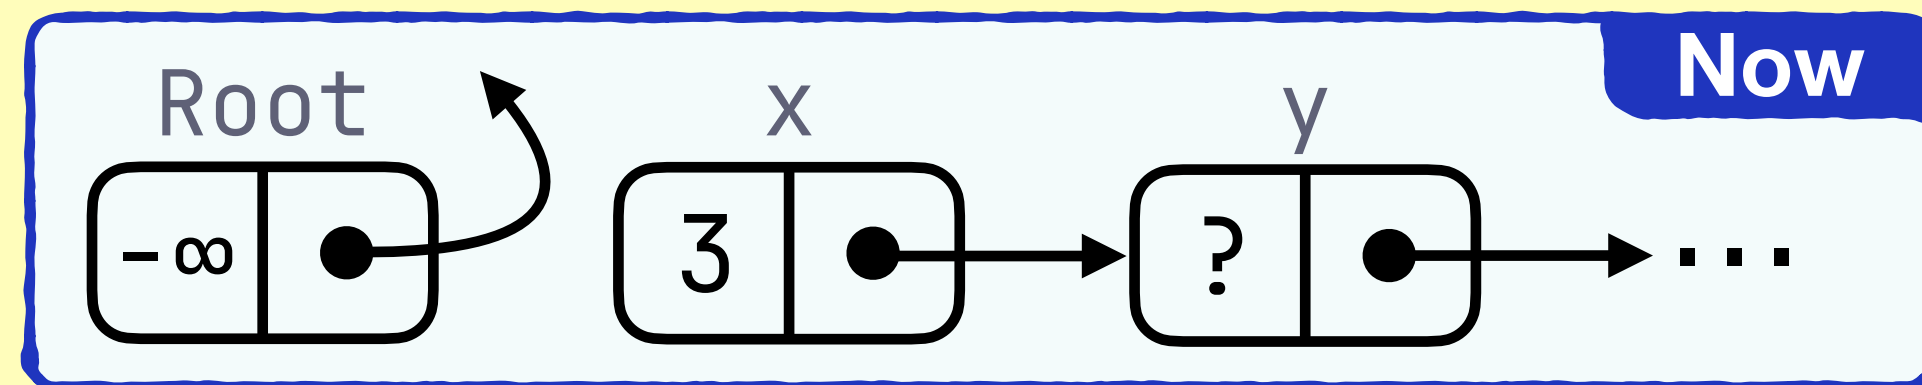


Example: contains(5)

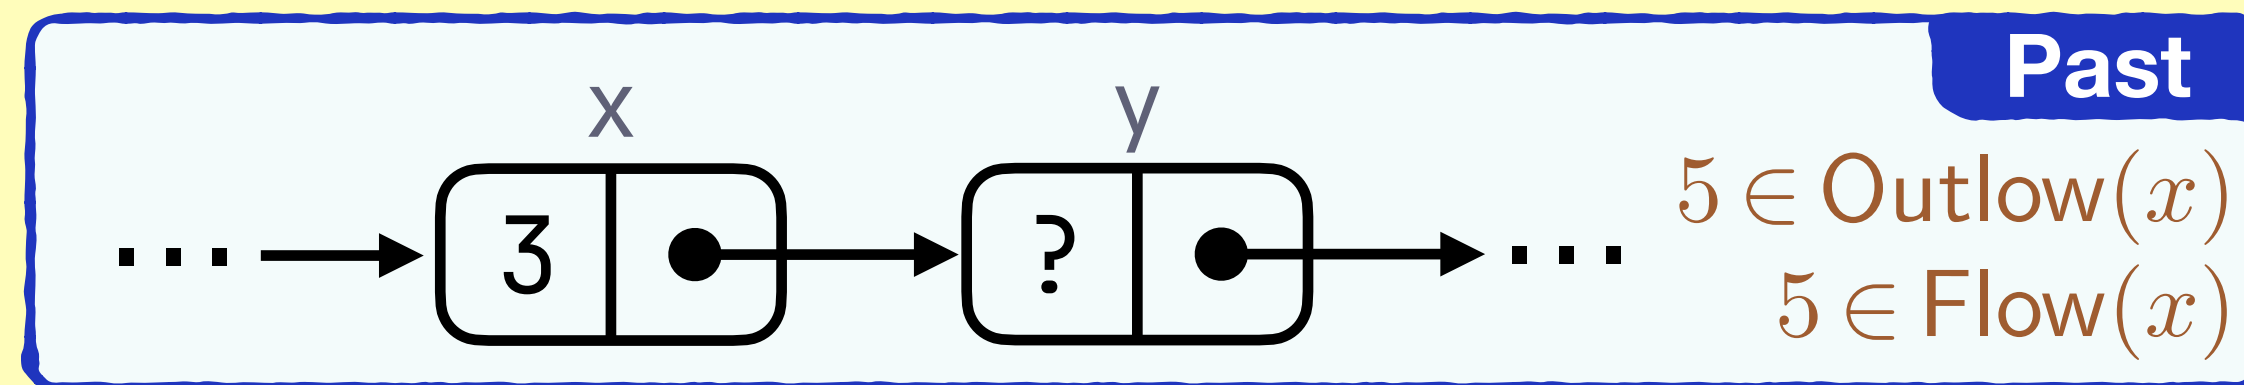
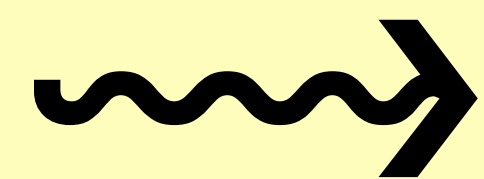
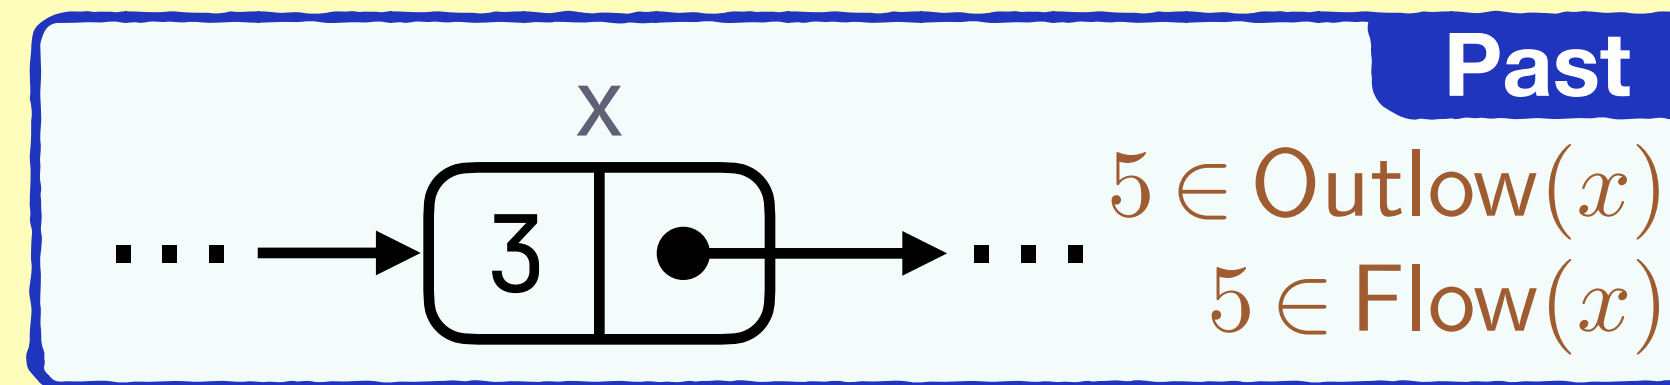
```
x = Root->next;  
while (x->key < 5):  
    y = x->next;
```

Traversal *learns* facts about Now, not the Past.

Use **hindsight** to propagate Now facts into the Past.



∧



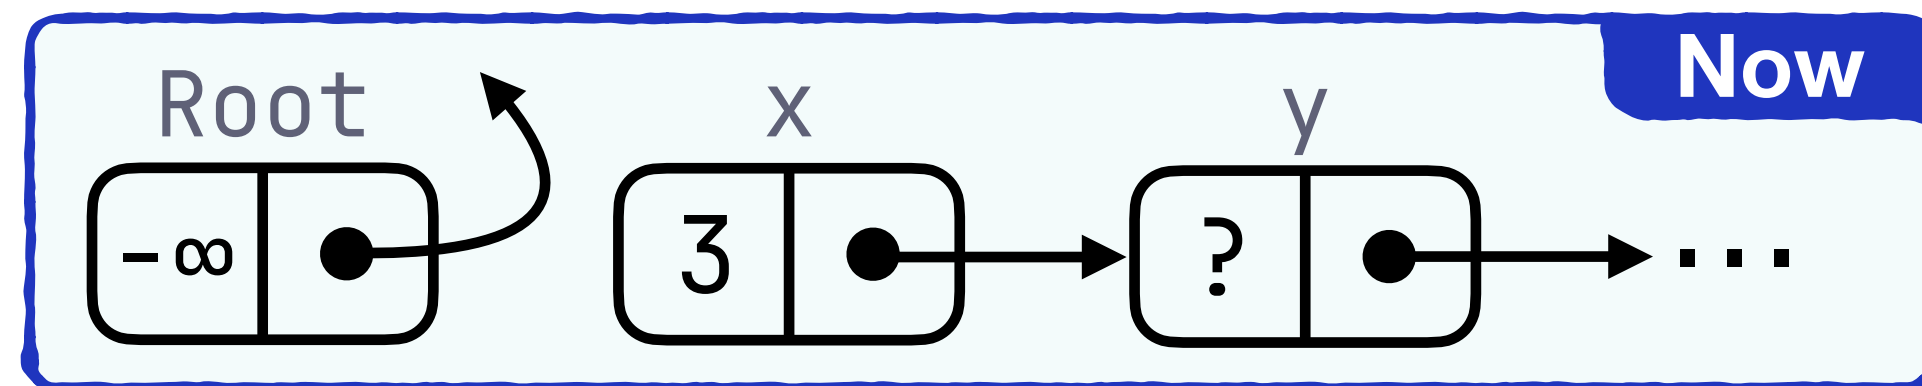
Temporal invariant:
unreachable nodes never change

Example: contains(5)

```
x = Root→next;
```

```
while (x→key < 5):
```

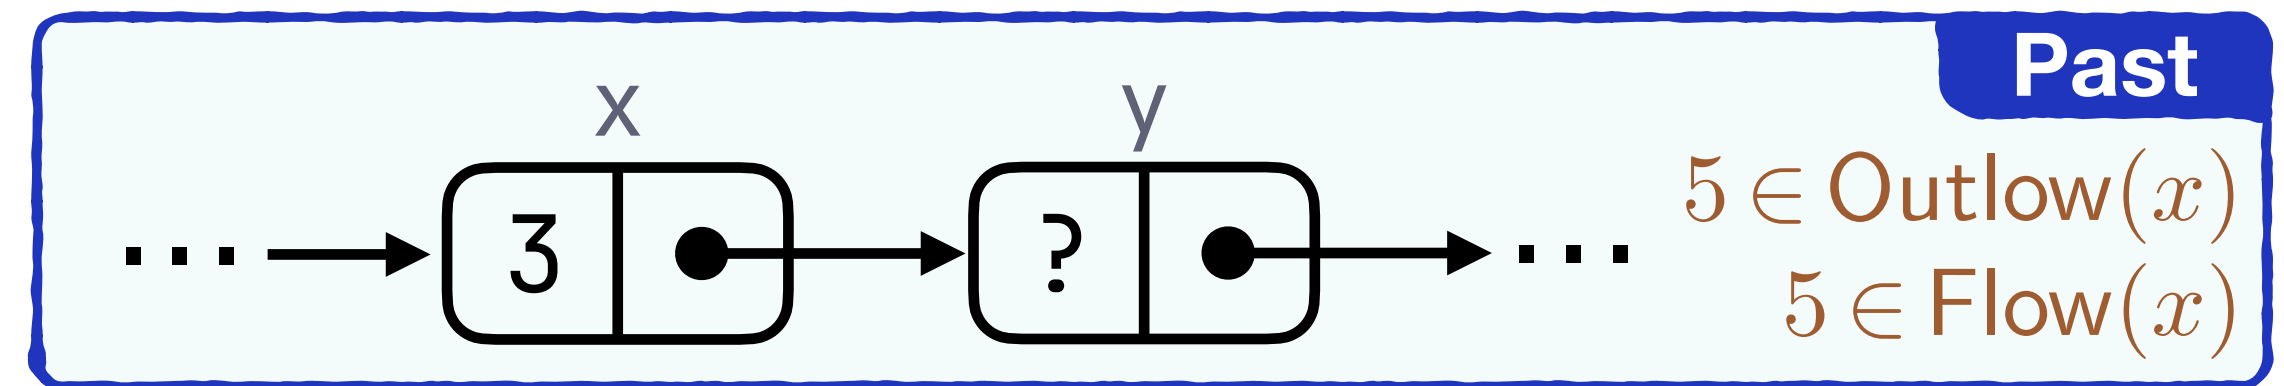
```
    y = x→next;
```



```
    x = y;
```

```
return x→key == 5;
```

∧

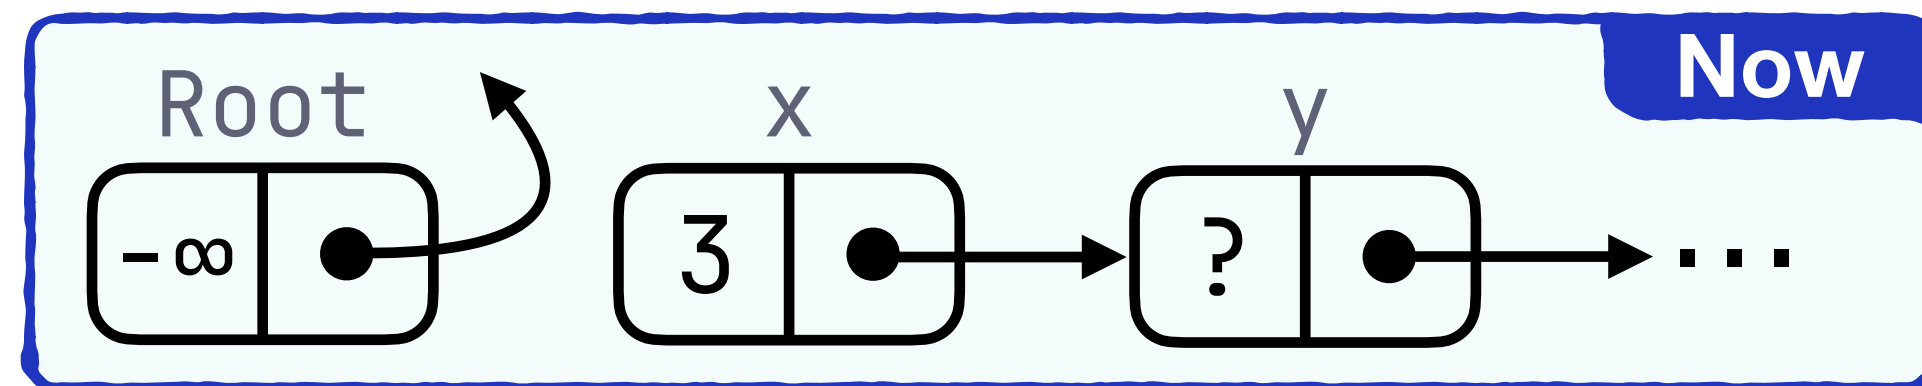


Example: contains(5)

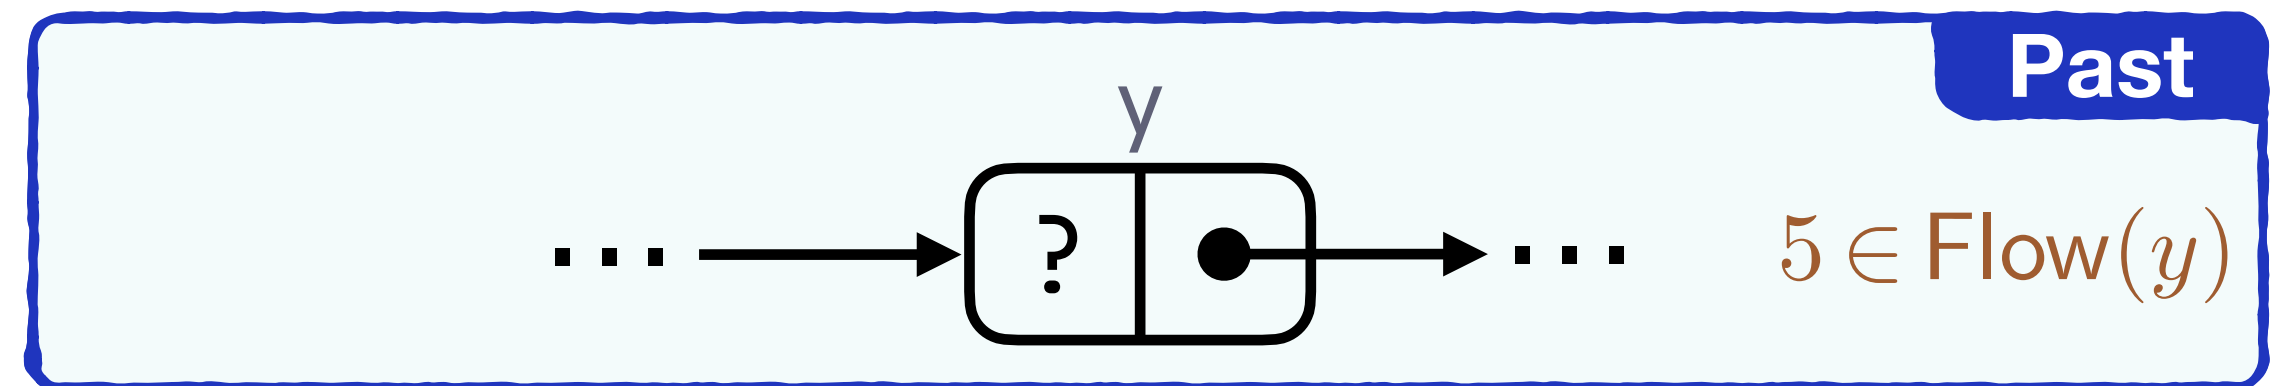
```
x = Root→next;
```

```
while (x→key < 5):
```

```
    y = x→next;
```



^



```
    x = y;
```

```
return x→key == 5;
```

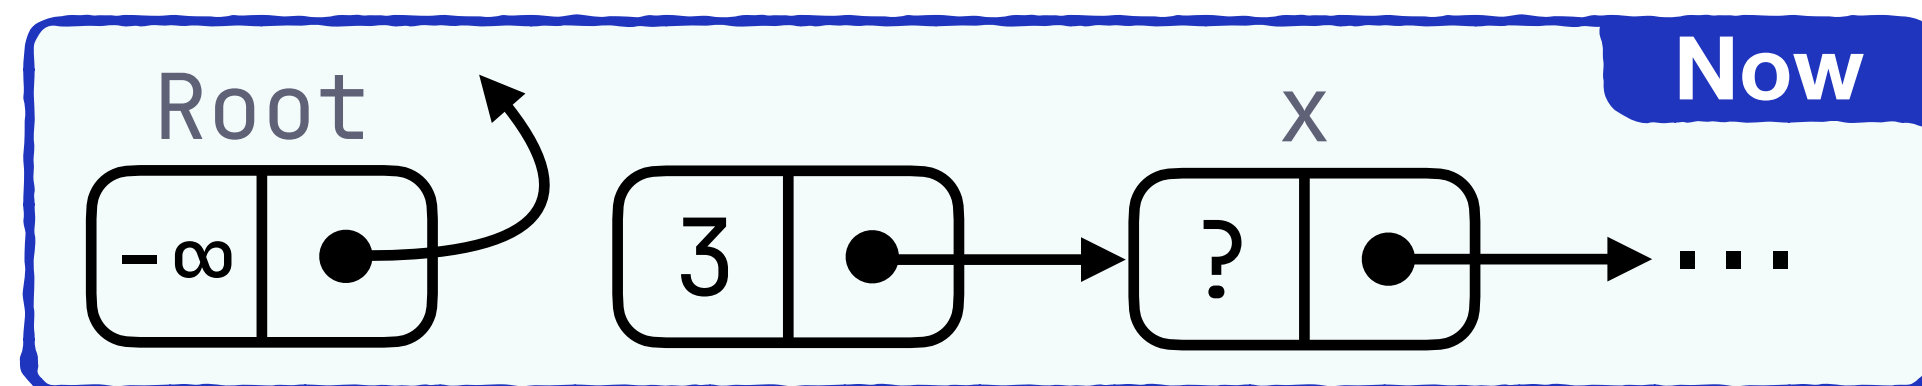
Example: contains(5)

```
x = Root→next;
```

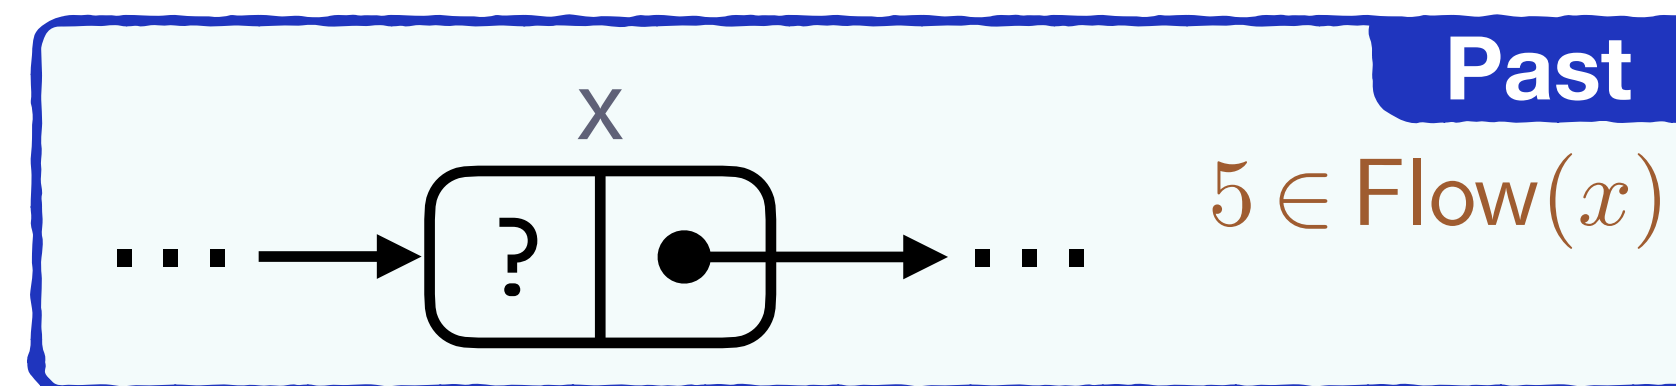
```
while (x→key < 5):
```

```
    y = x→next;
```

```
    x = y;
```



∧



```
return x→key == 5;
```

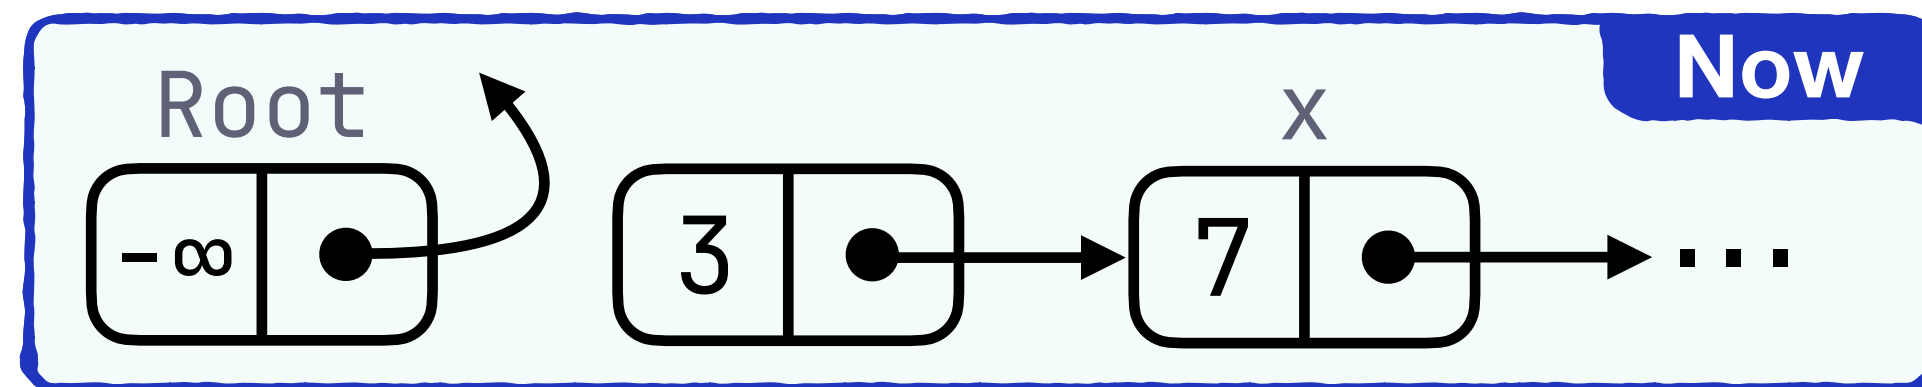
Example: contains(5)

```
x = Root→next;
```

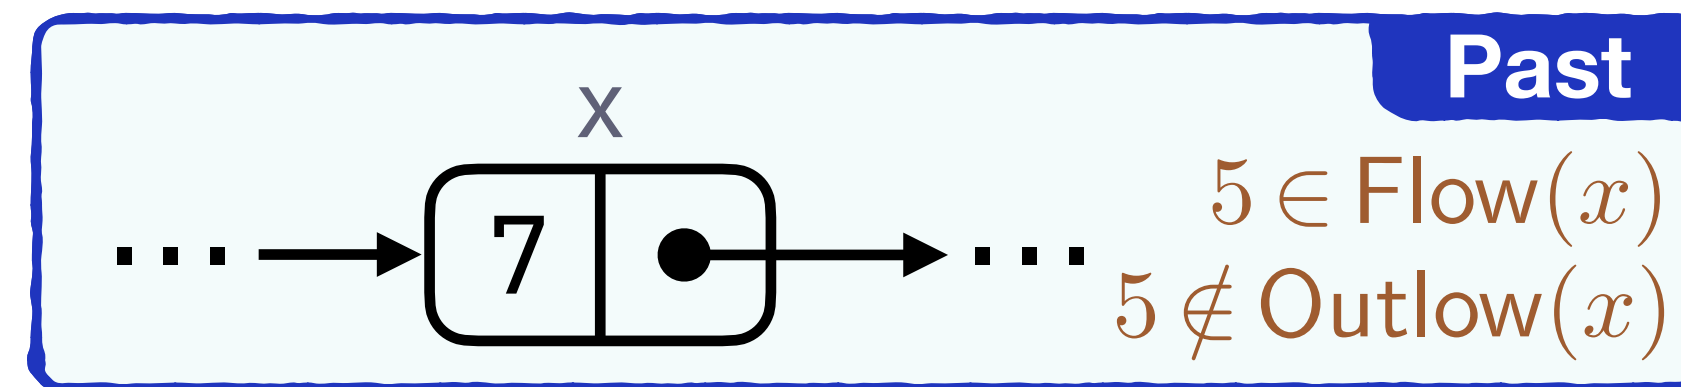
```
while (x→key < 5):
```

```
    y = x→next;
```

```
    x = y;
```



\wedge



```
return x→key == 5;
```

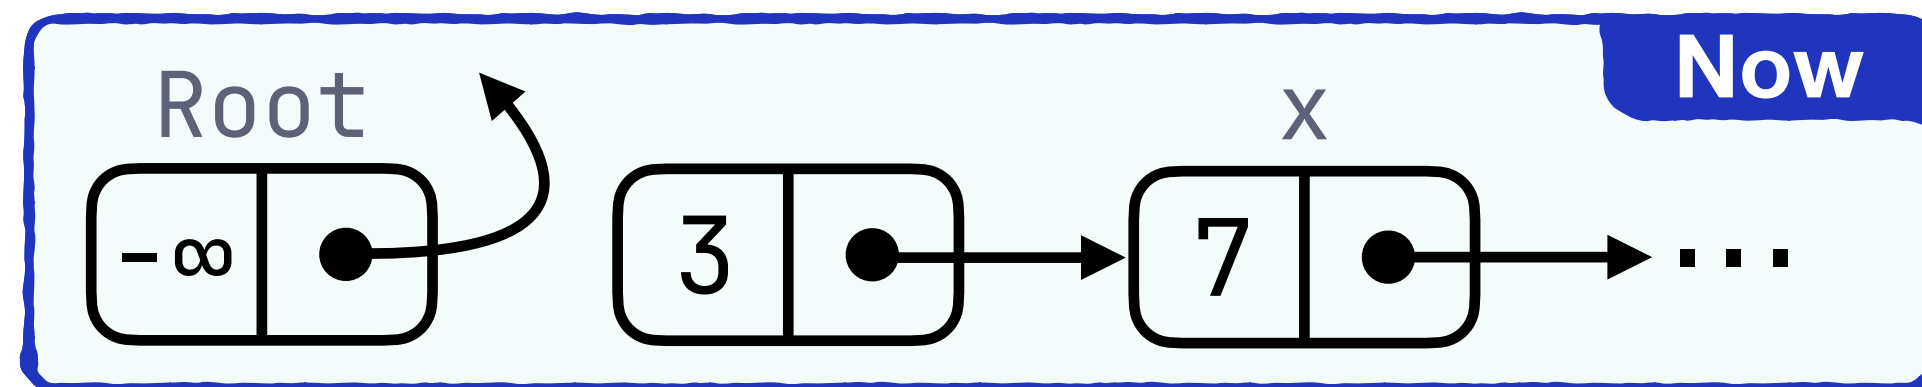
Example: contains(5)

```
x = Root→next;
```

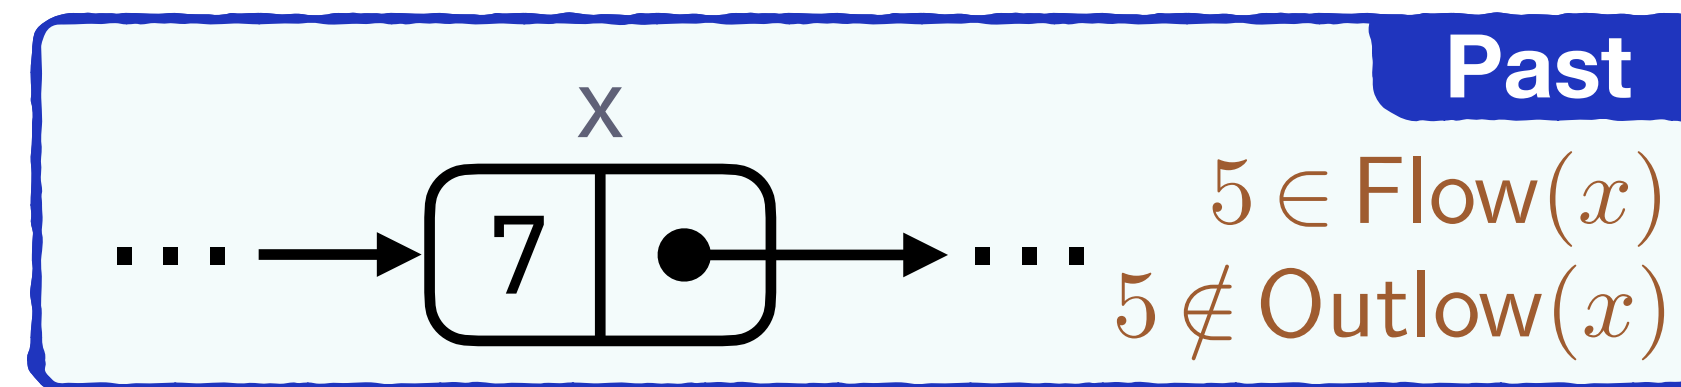
```
while (x→key < 5):
```

```
    y = x→next;
```

```
    x = y;
```



\wedge



```
return x→key == 5;
```

$\models \text{set}(M), 5 \notin M$ at some point

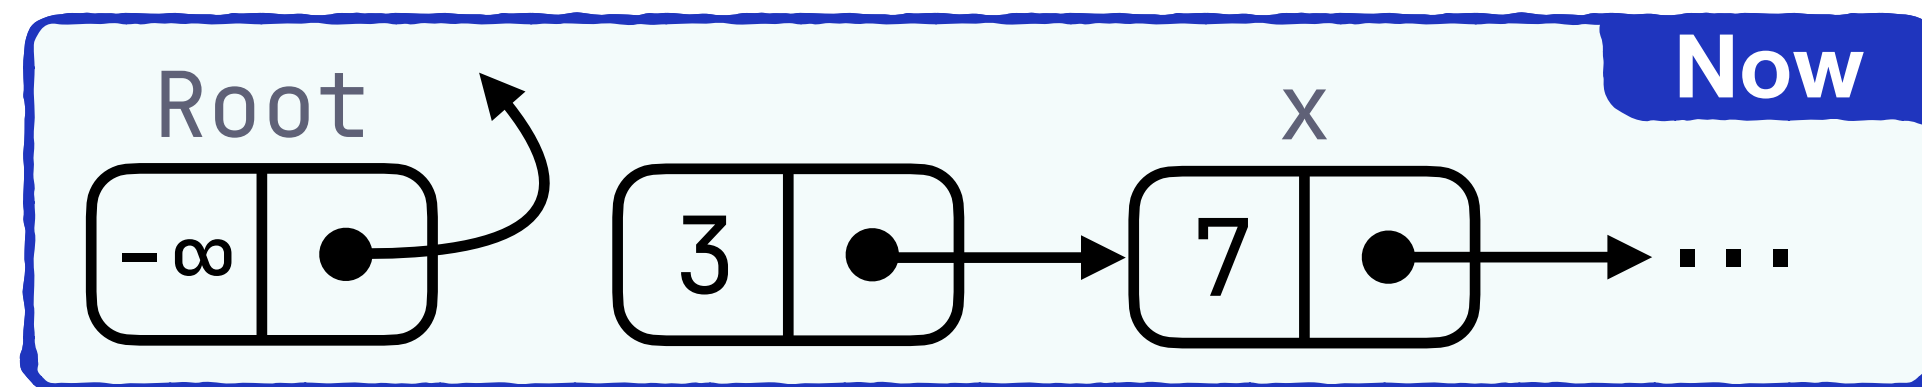
Example: contains(5)

```
x = Root→next;
```

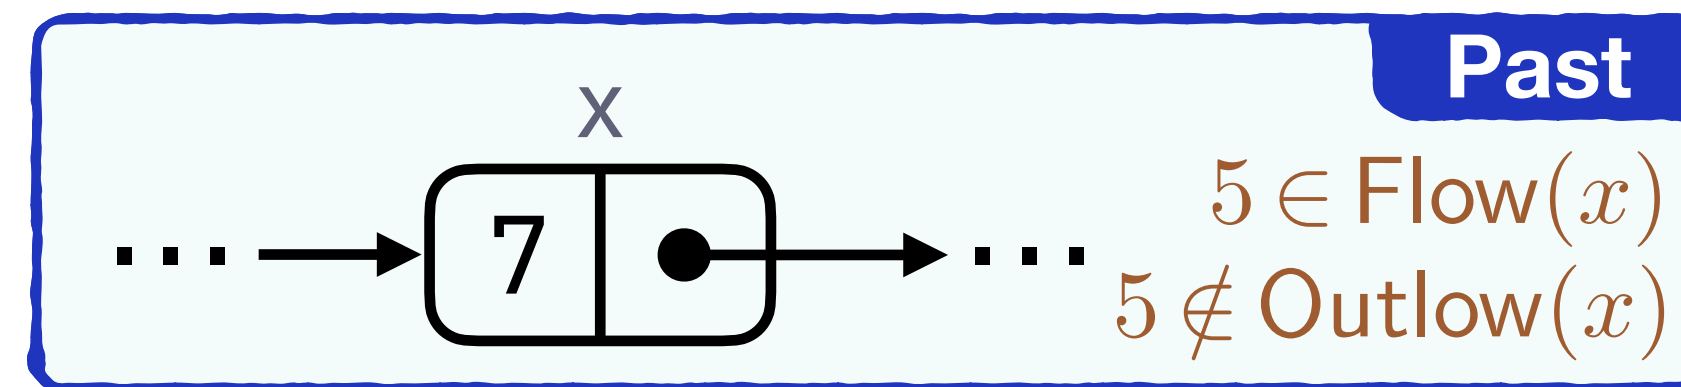
```
while (x→key < 5):
```

```
    y = x→next;
```

```
    x = y;
```



∧



```
return x→key == 5;
```

⊨ set(M), 5 ∉ M at some point

establishes linearizability proof obligation

Technical Contribution

- **Temporal interpolation** proof rule

Technical Contribution

- **Temporal interpolation** proof rule

$$\frac{\text{Interpolate}_{\mathbb{I}}(P, Q) \rightsquigarrow R}{\mathbb{I} \vdash \{S \wedge \text{Past}(P) \wedge \text{Now}(Q)\} \text{ skip } \{S \wedge \text{Past}(R)\}}$$

Technical Contribution

- **Temporal interpolation** proof rule
- Soundness
 - ➔ **eliminate** frame rule

Technical Contribution

- **Temporal interpolation** proof rule

$$\frac{\text{Interpolate}_{\mathbb{I}}(P, Q) \rightsquigarrow R}{\mathbb{I} \vdash \{S \wedge \text{Past}(P) \wedge \text{Now}(Q)\} \text{ skip } \{S \wedge \text{Past}(R)\}}$$

- Soundness
 - ➔ **eliminate** frame rule

Technical Contribution

- **Temporal interpolation** proof rule
- Soundness
 - ➔ **eliminate** frame rule
 - ➔ **eliminate** temporal interpolation rule

Technical Contribution

- **Temporal interpolation** proof rule

$$\frac{\text{Interpolate}_{\mathbb{I}}(P, Q) \rightsquigarrow R}{\mathbb{I} \vdash \{S \wedge \text{Past}(P) \wedge \text{Now}(Q)\} \text{ skip } \{S \wedge \text{Past}(R)\}}$$

- Soundness
 - ➔ **eliminate** frame rule
 - ➔ **eliminate** temporal interpolation rule

Technical Contribution

- **Temporal interpolation** proof rule
- Soundness
 - ➔ **eliminate** frame rule
 - ➔ **eliminate** temporal interpolation rule

Technical Contribution

- **Temporal interpolation** proof rule

$$\frac{\text{Interpolate}_{\mathbb{I}}(P, Q) \rightsquigarrow R}{\mathbb{I} \vdash \{S \wedge \text{Past}(P) \wedge \text{Now}(Q)\} \text{ skip } \{S \wedge \text{Past}(R)\}}$$

- Soundness
 - ➔ **eliminate** frame rule
 - ➔ **eliminate** temporal interpolation rule

Technical Contribution

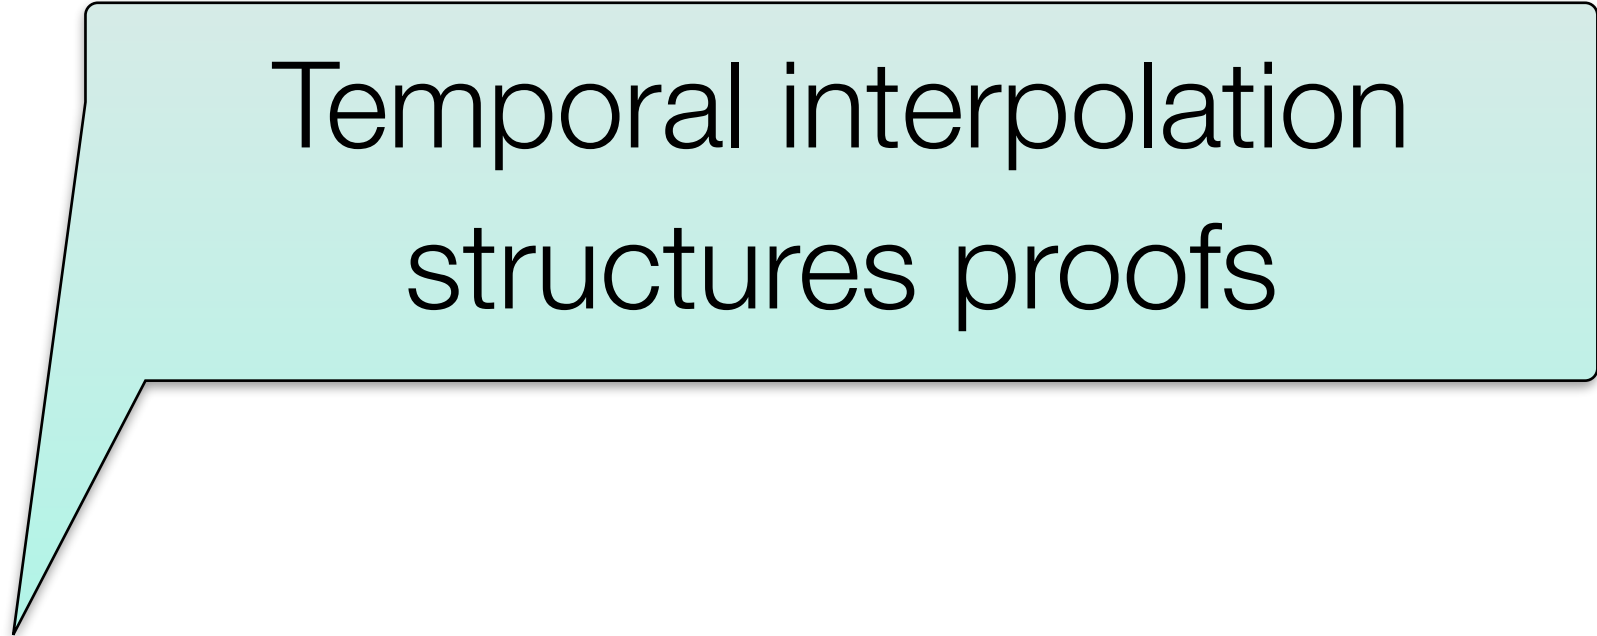
- **Temporal interpolation** proof rule

$$\frac{\text{Interpolate}_{\text{II}}(P, Q) \rightsquigarrow R}{\text{II} \vdash \{S \wedge \text{Past}(P) \wedge \text{Now}(Q)\} \text{ skip } \{S \wedge \text{Past}(R)\}}$$

- Soundness

- **eliminate** frame rule

- **eliminate** temporal interpolation rule



Temporal interpolation
structures proofs

Technical Contribution

- Interpolation requirements

$$\frac{\{ \text{Now}(P) \} \text{II}^* \{ \text{Now}(Q) \implies \text{Past}(R) \}}{\text{Interpolate}_{\text{II}}(P, Q) \rightsquigarrow R}$$

Technical Contribution

- Interpolation requirements

$$\frac{\{\text{Now}(P)\} \text{II}^* \{\text{Now}(Q) \implies \text{Past}(R)\}}{\text{Interpolate}_{\text{II}}(P, Q) \rightsquigarrow R}$$

- Interpolation strategy

$$\frac{\exists I. \quad \text{Now}(P) \implies I \quad I \text{ interference-free} \quad I \wedge \text{Now}(Q) \implies \text{Past}(R)}{\text{Interpolate}_{\text{II}}(P, Q) \rightsquigarrow R}$$

Technical Contribution

- Interpolation requirements

$$\frac{\{\text{Now}(P)\} \text{II}^* \{\text{Now}(Q) \implies \text{Past}(R)\}}{\text{Interpolate}_{\text{II}}(P, Q) \rightsquigarrow R}$$

- Interpolation strategy

$$\frac{\exists I. \quad \text{Now}(P) \implies I \quad I \text{ interference-free} \quad I \wedge \text{Now}(Q) \implies \text{Past}(R)}{\text{Interpolate}_{\text{II}}(P, Q) \rightsquigarrow R}$$

- Great choice: $R = P \wedge Q$

Technical Contribution

- Interpolation requirements

$$\frac{\{ \text{Now}(P) \} \text{II}^* \{ \text{Now}(Q) \} \implies P}{\text{Interpolate}_{\text{II}}(P, Q) \rightsquigarrow I}$$

- Interpolation strategy

$$\frac{\exists I. \text{Now}(P) \implies I \quad I \text{ interference-free} \quad I}{\text{Interpolate}_{\text{II}}(P, Q) \rightsquigarrow I}$$

- Great choice: $R = P \wedge Q$

Embedding Hindsight Reasoning in Separation Logic

ROLAND MEYER, TU Braunschweig, Germany

THOMAS WIES, New York University, USA

SEBASTIAN WOLFF, New York University, USA

Automatically proving linearizability of concurrent data structures remains a key challenge for verification. We present temporal interpolation as a new proof principle to guide automated proof search using hindsight arguments within concurrent separation logic. Temporal interpolation offers an easy-to-automate alternative to prophecy variables and has the advantage of structuring proofs into easy-to-discharge hypotheses. Additionally, we advance hindsight theory by integrating it into a program logic, bringing formal rigor and complementary proof machinery. We substantiate the usefulness of temporal interpolation by implementing it in a tool and using it to automatically verify the Logical Ordering tree. The proof is challenging due to future-dependent linearization points and complex structure overlays. It is the first formal proof of this data structure. Interestingly, our formalization revealed an unknown bug and an existing informal proof as erroneous.

Conference Version:

Roland Meyer, Thomas Wies, and Sebastian Wolff. 2023. Embedding Hindsight Reasoning in Separation Logic. *Proc. ACM Program. Lang.* 7, PLDI, Article 182 (June 2023), 24 pages. <https://doi.org/10.1145/3591296>

1 INTRODUCTION

We are concerned with automatically proving linearizability, the standard correctness criterion for concurrent data structures [Herlihy and Wing 1990]. A concurrent data structure is linearizable subject to a sequential specification of its methods, if each method takes effect in a single atomic step of its concurrent execution, the method's *linearization point*, and satisfies the sequential specification in this step.



Concurrent separation logics [Bell et al. 2010; Delbianco et al. 2017; Elmas et al. 2010; Fu et al. 2010; Gotsman et al. 2013; Gu et al. 2018; Hemed et al. 2015; Parkinson et al. 2007; Sergey et al. 2015; Vafeiadis and Parkinson 2007] provide a powerful toolbox of deductive reasoning techniques to verify complex concurrent data structures. However, the proof construction heavily relies on the proof author's creativity and expertise in wielding the available tools effectively. For instance, in order to construct the inductive invariant of the data structure, the proof author may have to devise proof-specific resource algebras to express ghost state that captures the key aspects of the computation history. This hinders proof automation due to the vast complexity of the proof space that needs to be explored. Similarly, the proofs may make use of prophecy variables [Abadi and Lamport 1991] to predict future-dependent linearization points [Jung et al. 2020; Liang and Feng 2013; Vafeiadis 2008]. Constructing such proofs involves backward reasoning, which is difficult to automate [Bouajjani et al. 2017]. It stands to reason that there is a need for guiding principles that help to structure the proof and that provide effective strategies for automated tools to prune the search space.

Hindsight theory [Feldman et al. 2018, 2020; Lev-Ari et al. 2015; O'Hearn et al. 2010] provides such a guiding principle, which we refer to as *temporal interpolation*. One proves lemmas of the form: if there existed a past state that satisfied property p and the current state satisfies q , then there must have existed an intermediate state that satisfied o . Such lemmas can then be applied, e.g., to prove the



© 2023 Copyright held by the owner/author(s).



This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the ACM on Programming Languages*, <https://doi.org/10.1145/3591296>.

Implementations

- pLankton  
 - linearizability checker for concurrent search structures
 - input: edge functions + node invariant
 - hard-coded: search path flow monoid
 - automatically constructs proofs

Implementations

- **pLankton**  
 - linearizability checker for concurrent search structures
 - input: edge functions + node invariant
 - hard-coded: search path flow monoid
 - automatically constructs proofs

- **nektion**  
 - additional input: custom flow monoid + proof outline
 - checks proof outline for validity

Implementations

- p_lankton



- linearizability checker for concurrent search

- input: edge functions + node invariant

- hard-coded: search path flow monoid

- automatically constructs proofs

- nektion



- additional input: custom flow monoid + proc

- checks proof outline for validity

[PLDI'23]

Embedding Hindsight Reasoning in Separation Logic

ROLAND MEYER, TU Braunschweig, Germany

THOMAS WIES, New York University, USA

SEBASTIAN WOLFF, New York University, USA

Automatically proving linearizability of concurrent data structures remains a key challenge for verification.

We present arguments relative to properties. Additionally, we complemented the tool with an independent interesting

Conference
Roland Meyer
Proc. ACM

1 INTRODUCTION

We are considering concurrent search structures. The tool's unique features are its parametric heap abstraction based on separation logic and the flow framework, and its support for hindsight arguments about future-dependent linearization points. We describe the tool, present a case study, and discuss implementation details.

© 2023 Copyright
This is the author's
of Record work

[CAV'23]

nektion: a linearizability proof checker

Roland Meyer¹, Anton Opaterny¹, Thomas Wies², and Sebastian Wolff²



¹ TU Braunschweig, Braunschweig, Germany
{anton.opaterny, roland.meyer}@tu-bs.de

² New York University, New York, USA
{wies, sebastian.wolff}@cs.nyu.edu



Abstract nektion is a new tool for checking linearizability proofs of highly complex concurrent search structures. The tool's unique features are its parametric heap abstraction based on separation logic and the flow framework, and its support for hindsight arguments about future-dependent linearization points. We describe the tool, present a case study, and discuss implementation details.

Keywords: separation logic · proof checker · linearizability · flow framework

1 Introduction

We present nektion, a mostly automated deductive program verifier based on separation logic (SL) [23,27]. The tool is designed to aid the construction of linearizability proofs for complex concurrent search structures. Similar to many other SL-based tools [2,8,14,22,33,33], nektion uses an SMT solver to automate basic SL reasoning. Similar to the original implementation of CIVL [7], it uses non-interference reasoning à la Owicki-Gries [25] to automate thread modularity. What makes nektion stand out among these relatives is its inbuilt support for expressing complex inductive heap invariants using the flow framework [12,13,20] and the ability to (partially) automate complex linearizability arguments that require hindsight reasoning [4,5,15,18,19,24]. Together, these features enable nektion to verify challenging concurrent data structures such as the FEMRS tree [4] with little user guidance.

nektion [17] is derived from the tool plankton [18,19], which shares the same overall goals and features as nektion but strives for full proof automation at the expense of generality. In terms of the trade-off between automation and expressivity, nektion aims to occupy a sweet spot between plankton and general purpose program verifiers. In the following, we discuss nektion's unique features in more detail and explain how it deviates from plankton's design.

The flow framework can be used to express global properties of graph structures in a node-local manner, aiding compositional verification of recursive data structures. The framework is parametric in a *flow domain* which determines what global information about the graph is provided at each node. Various flow domains have been proposed that have shown to be useful in concurrency proofs [11,26]. To simplify proof automation, plankton uses a fixed flow domain that is geared towards verifying functional correctness of search structures. In contrast, nektion is parametric in the flow domain. For instance, it supports custom domains for reasoning about overlaid structures and

pLankton



- Verifies future-dependent linearization points
- **Logical Ordering Tree**
 - found bug in implementation
 - found bug in previous proof
 - first proof of (fixed) version

Benchmark	Verification Time
Fine-Grained set	45s ✓
Lazy set	2m 13s ✓
FEMRS tree (no maintenance)	3m 50s ✓
Vechev&Yahav 2CAS set	1m 15s ✓
Vechev&Yahav CAS set	2m 20s ✓
ORVYY set	1m 36s ✓
Michael set	6m 53s ✓
Michael set (wait-free search)	6m 53s ✓
Harris set	57m 20s ✓
Harris set (wait-free search)	43m 00s ✓
LO-tree (generalized maintenance)	16m 43s ✓

Overview



Flow Framework

Data-flow-like ghost state for inductive heap invariants.

[POPL'18, ESOP'20, TACAS'23]



Hindsight

Non-fixed linearization points without prophecies.

[OOPSLA'22, PLDI'23]



Decomposing Updates

Unbounded footprints without induction.

[OOPSLA'22, *under submission*]

Goal

Localize the proof to a small number of nodes.

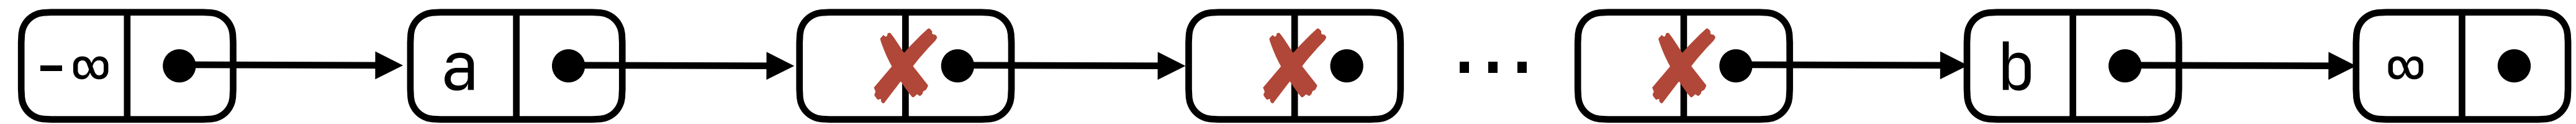
Problem

Complex updates affect an unbounded number of nodes.

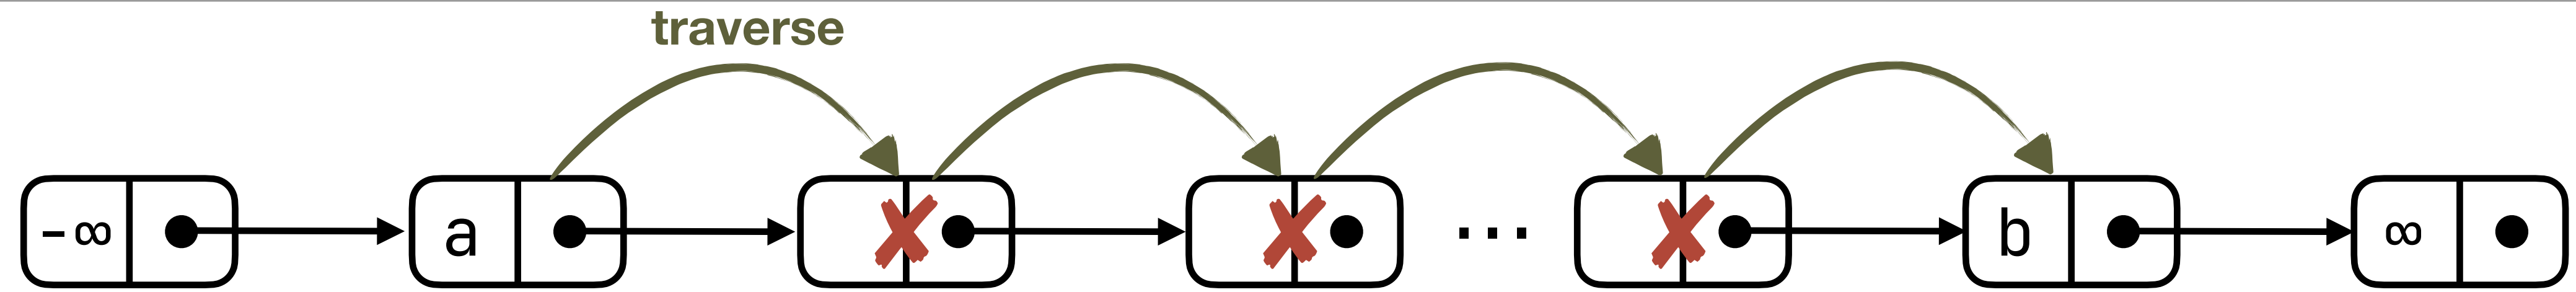
Solution

Keep track of ghost updates that are local.

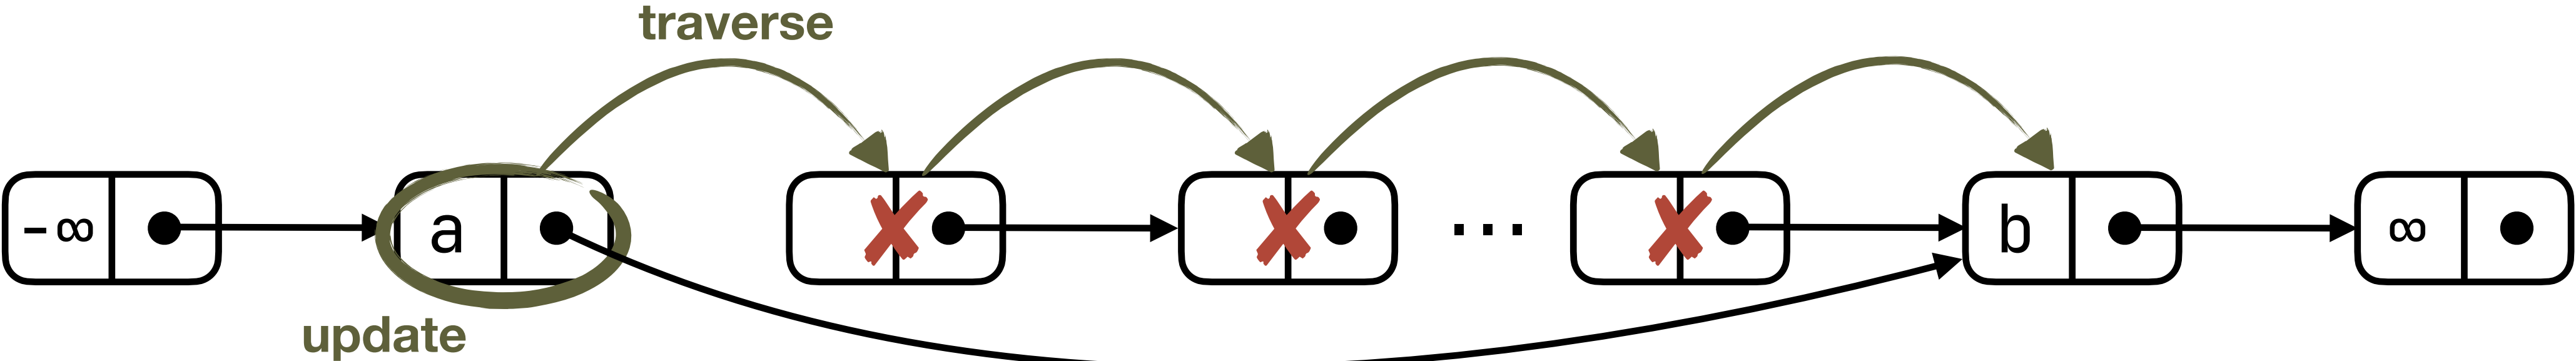
Complex Updates



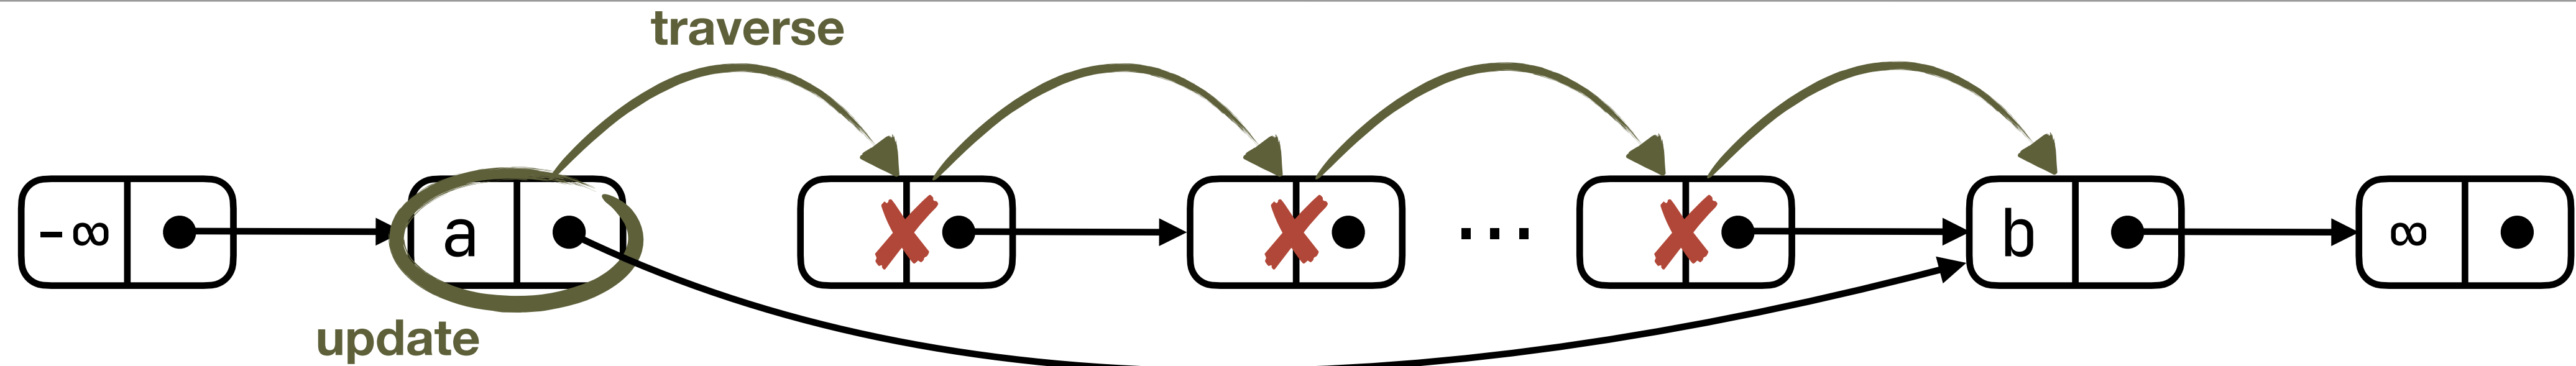
Complex Updates



Complex Updates



Complex Updates



State of the art: **recursive predicates + induction**

hard to automate

- custom predicate captures to-be-updated **region**
- **induction** over predicate upon actual update

Lesson in Life

Computation is **local**



and

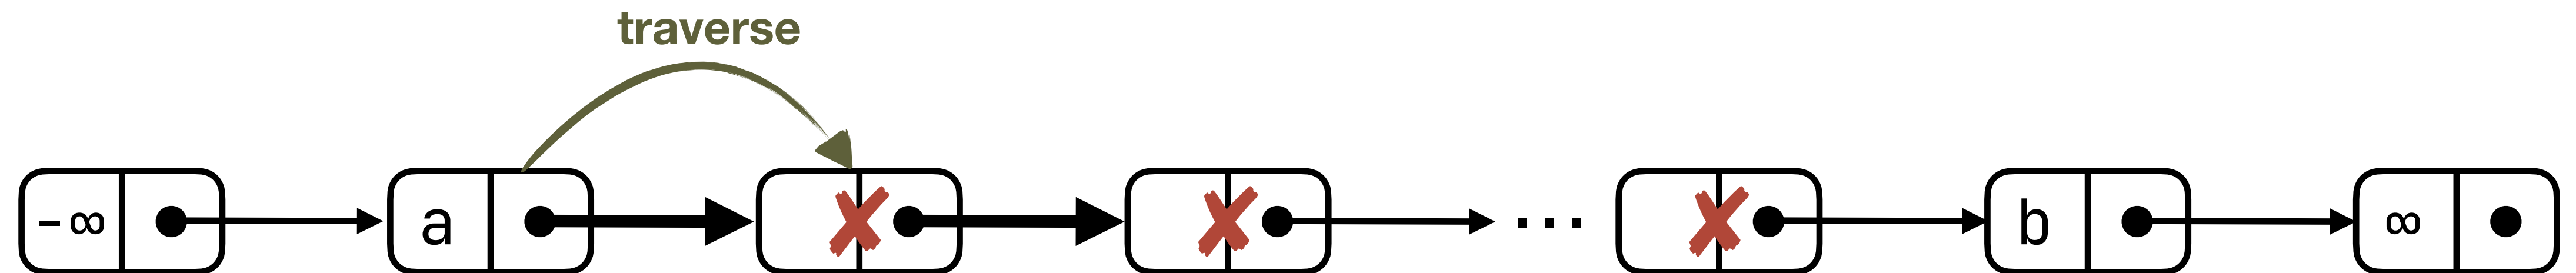
continuous



Complex Updates

Gost update chunks

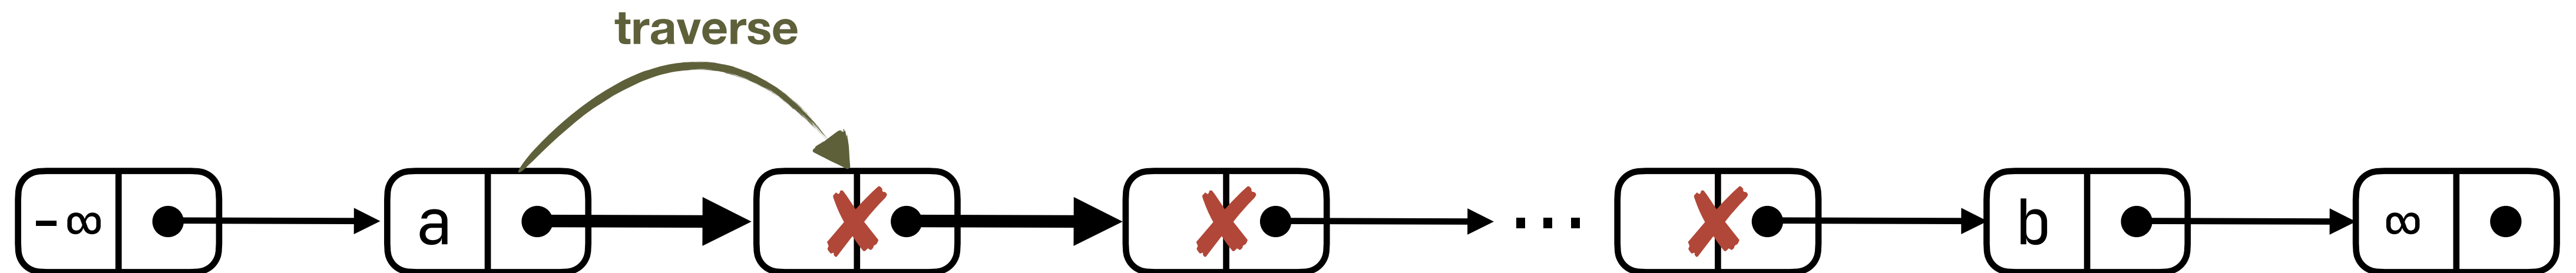
- ➔ compose complex update from smaller chunks
- ➔ the chunks are hypothetical/ghost information for the proof
- ➔ can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

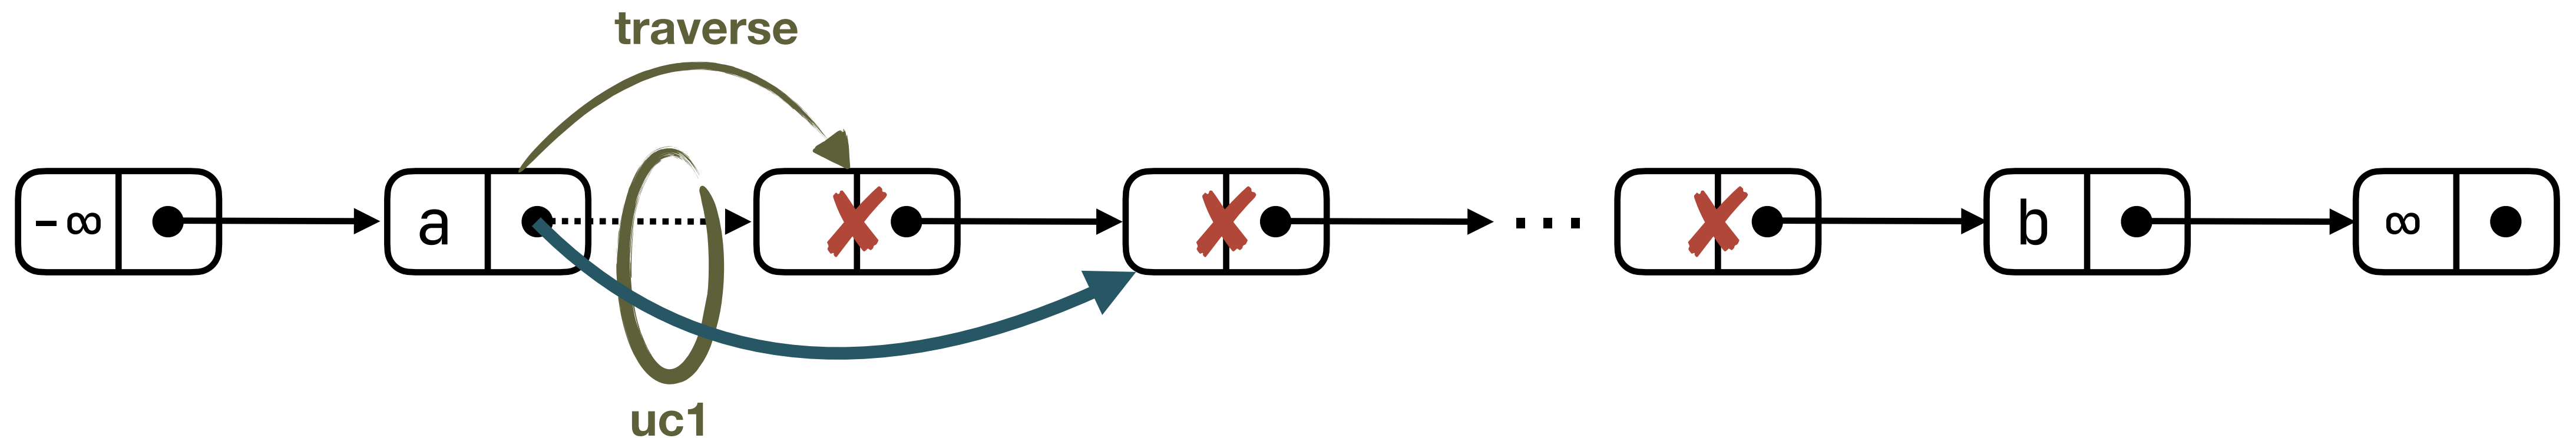
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

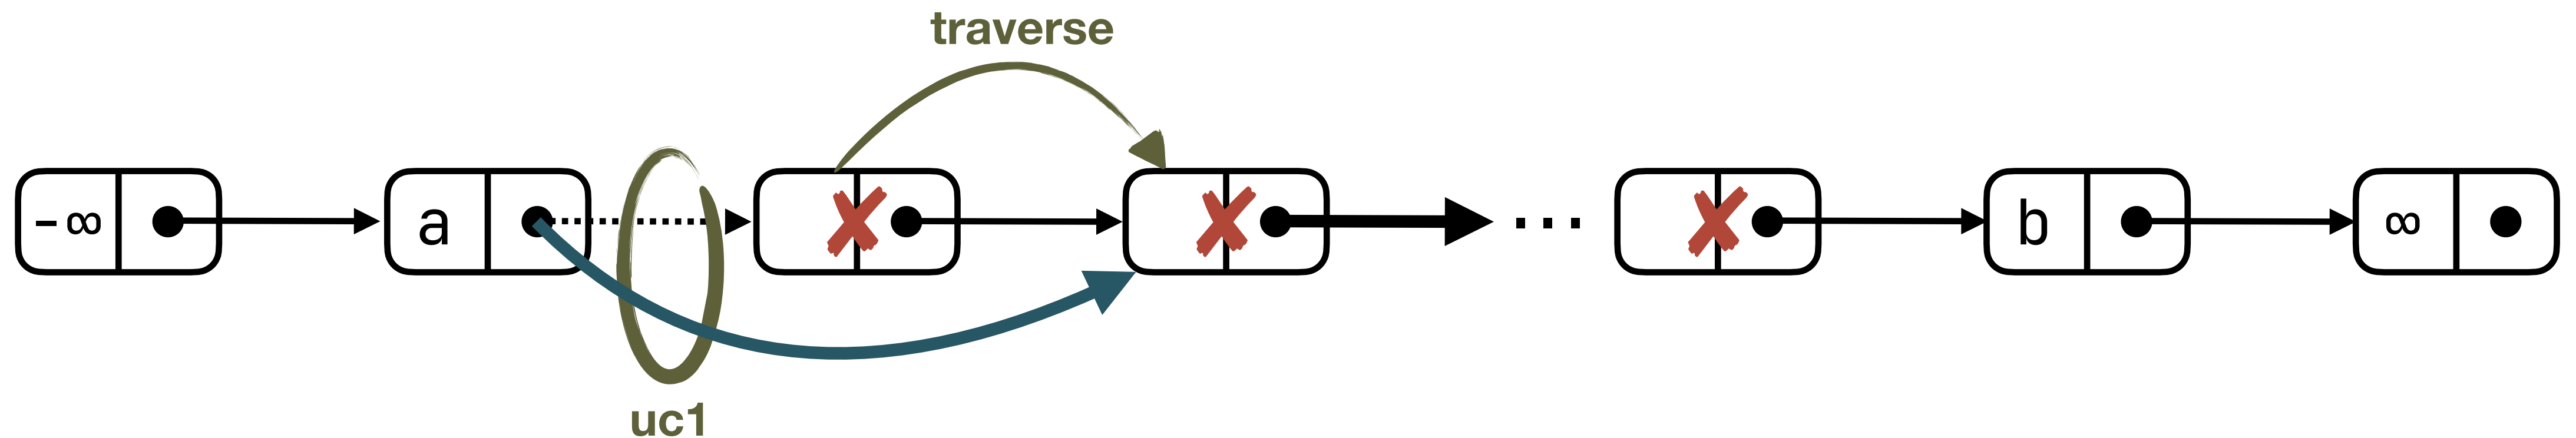
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

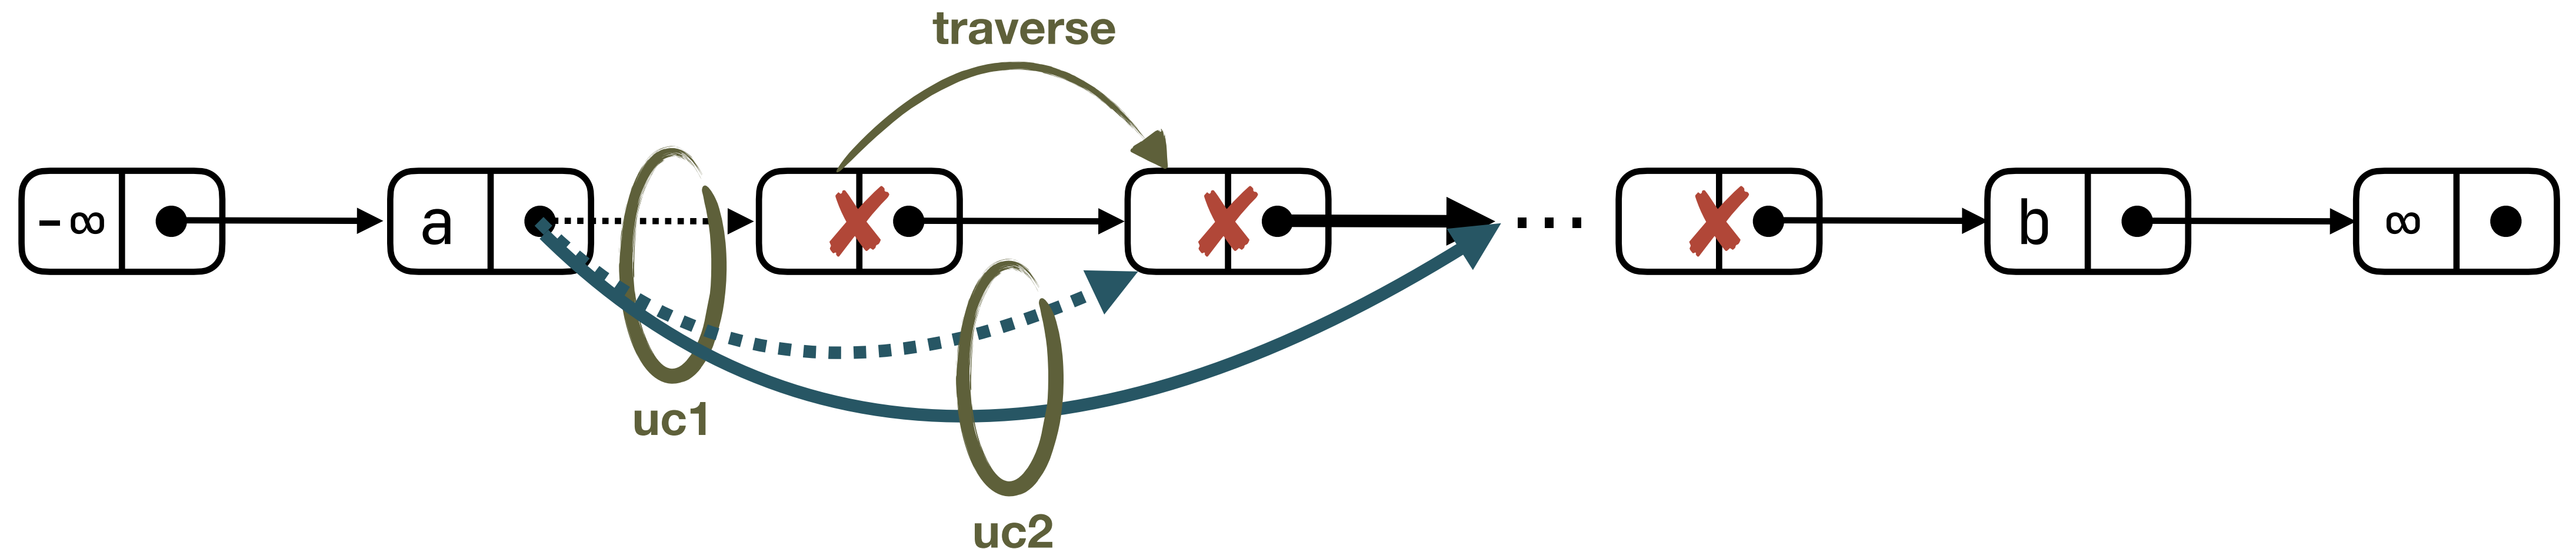
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

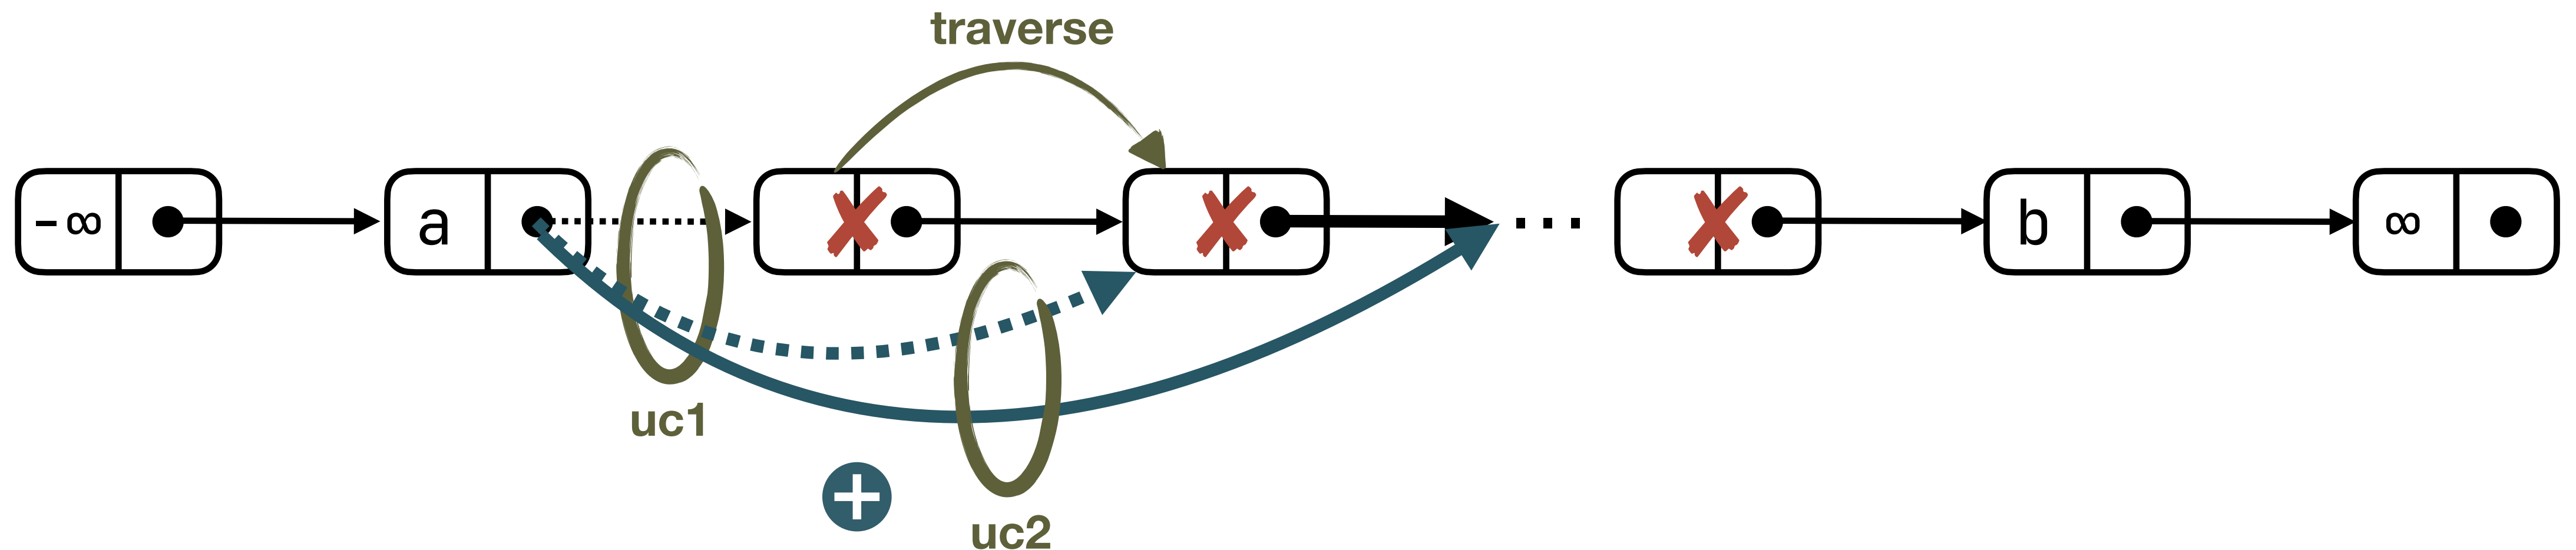
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

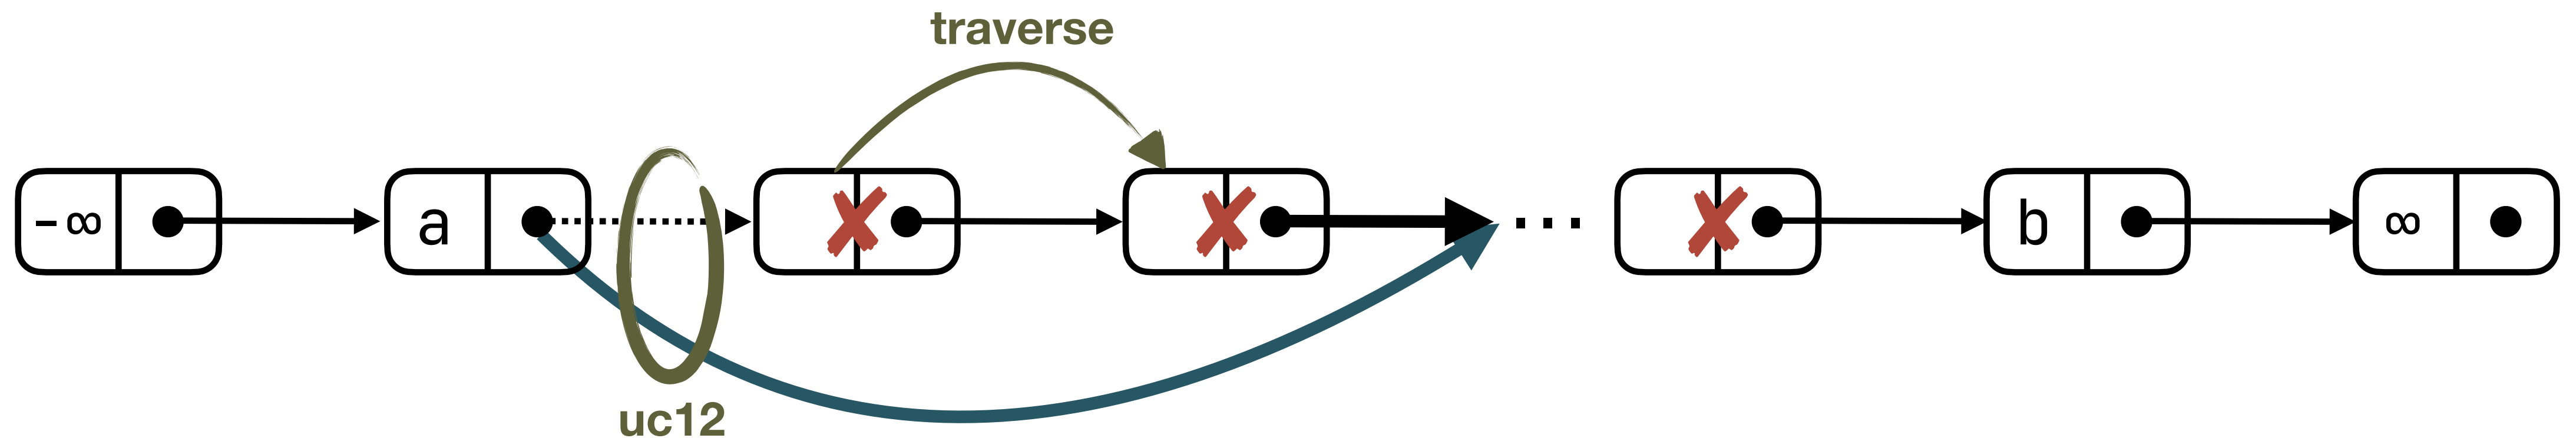
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

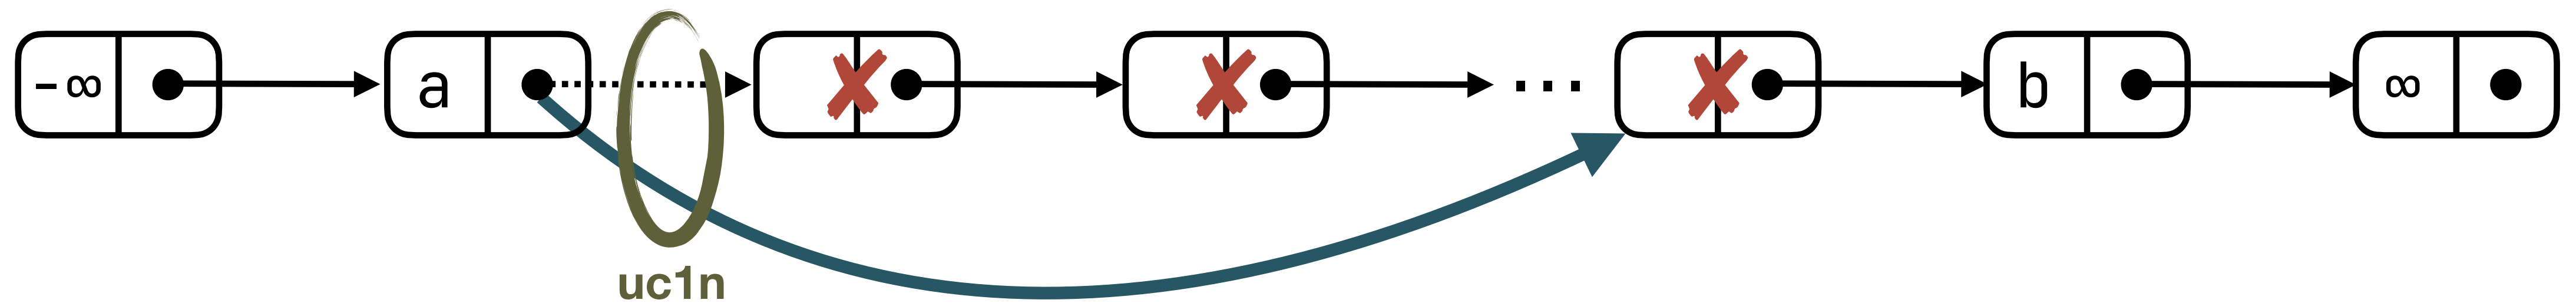
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

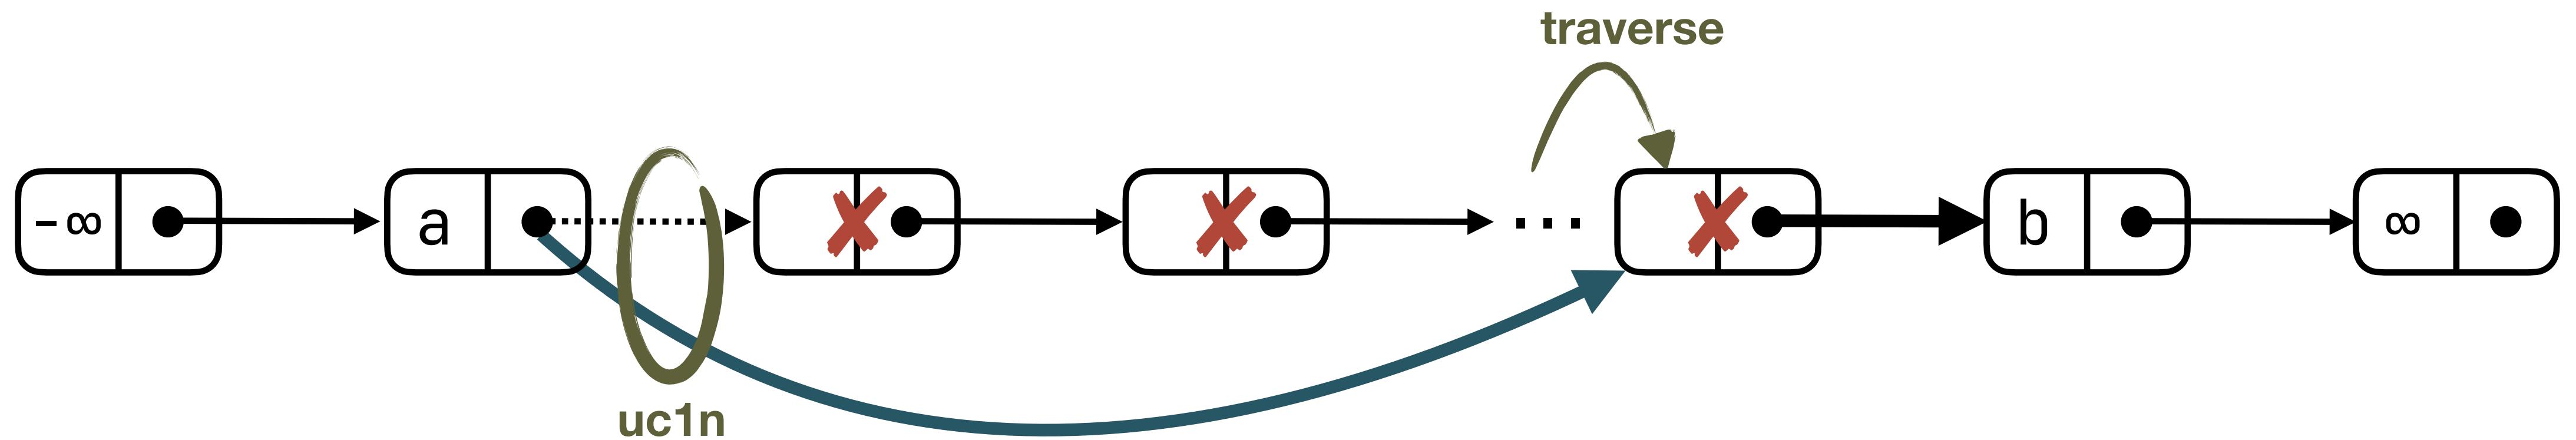
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

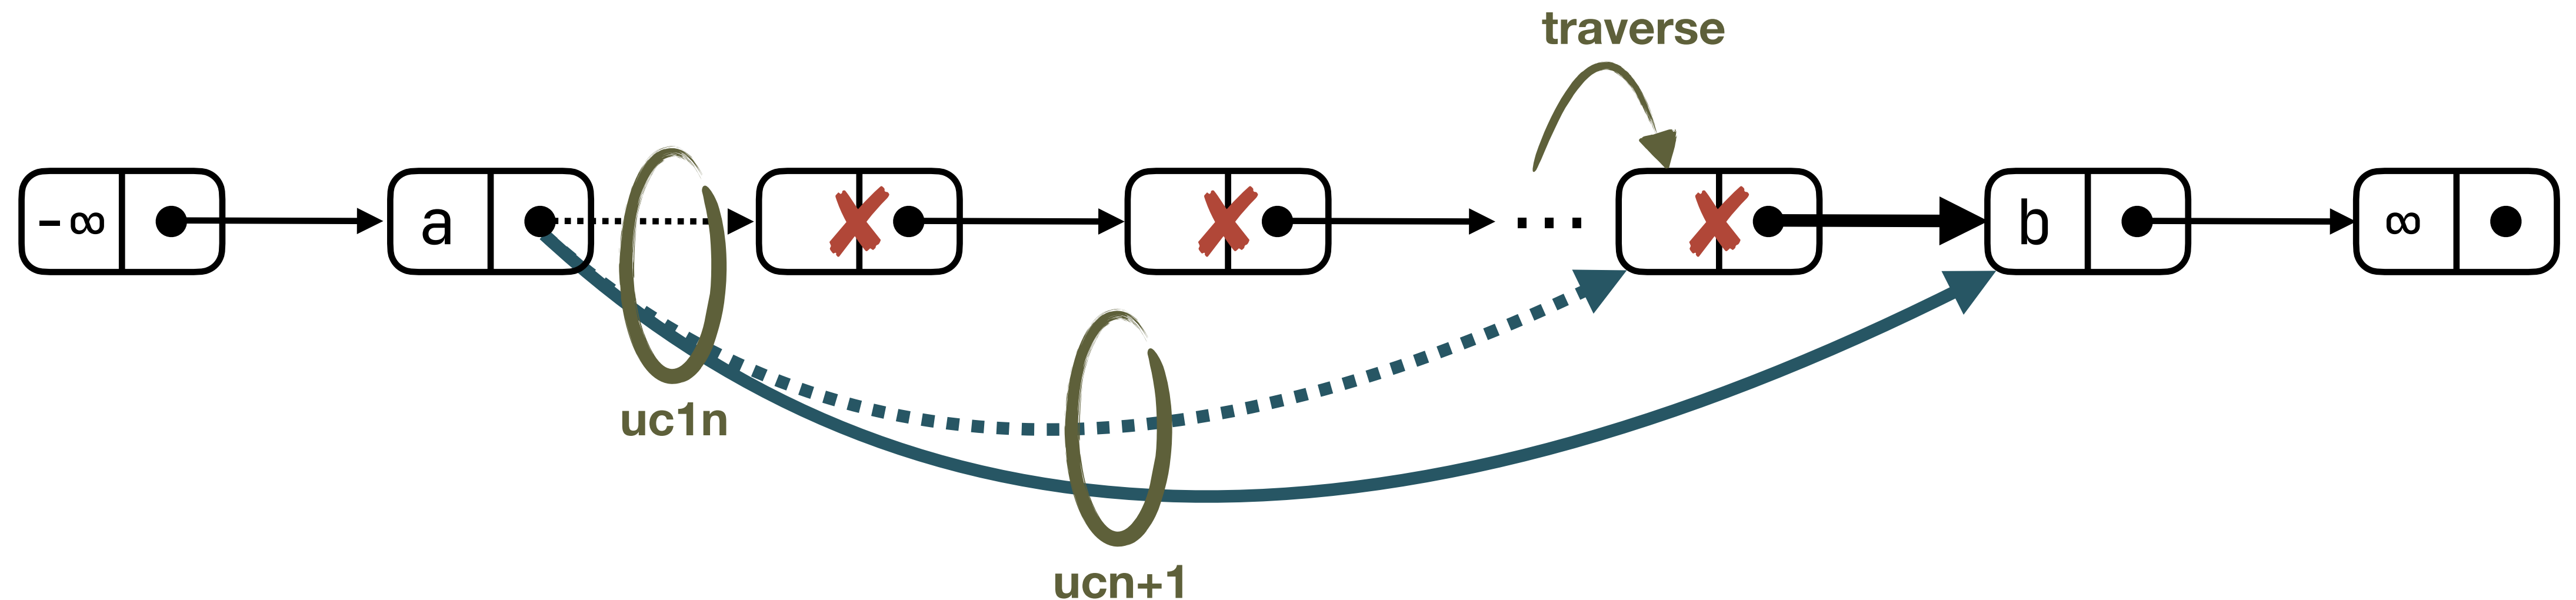
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

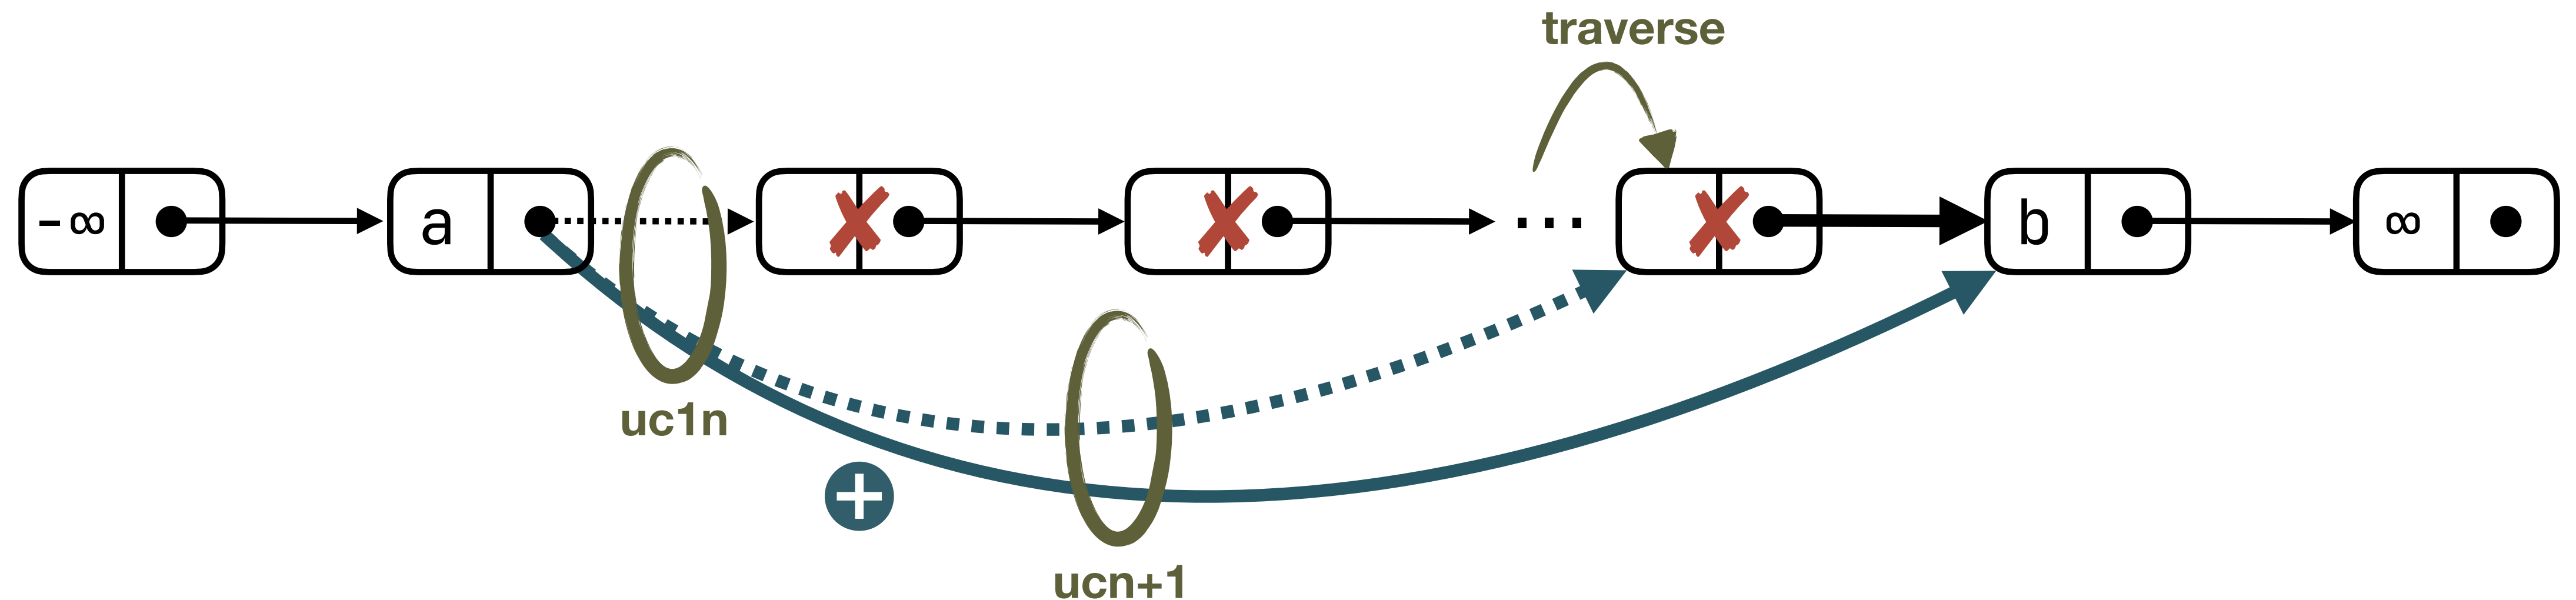
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

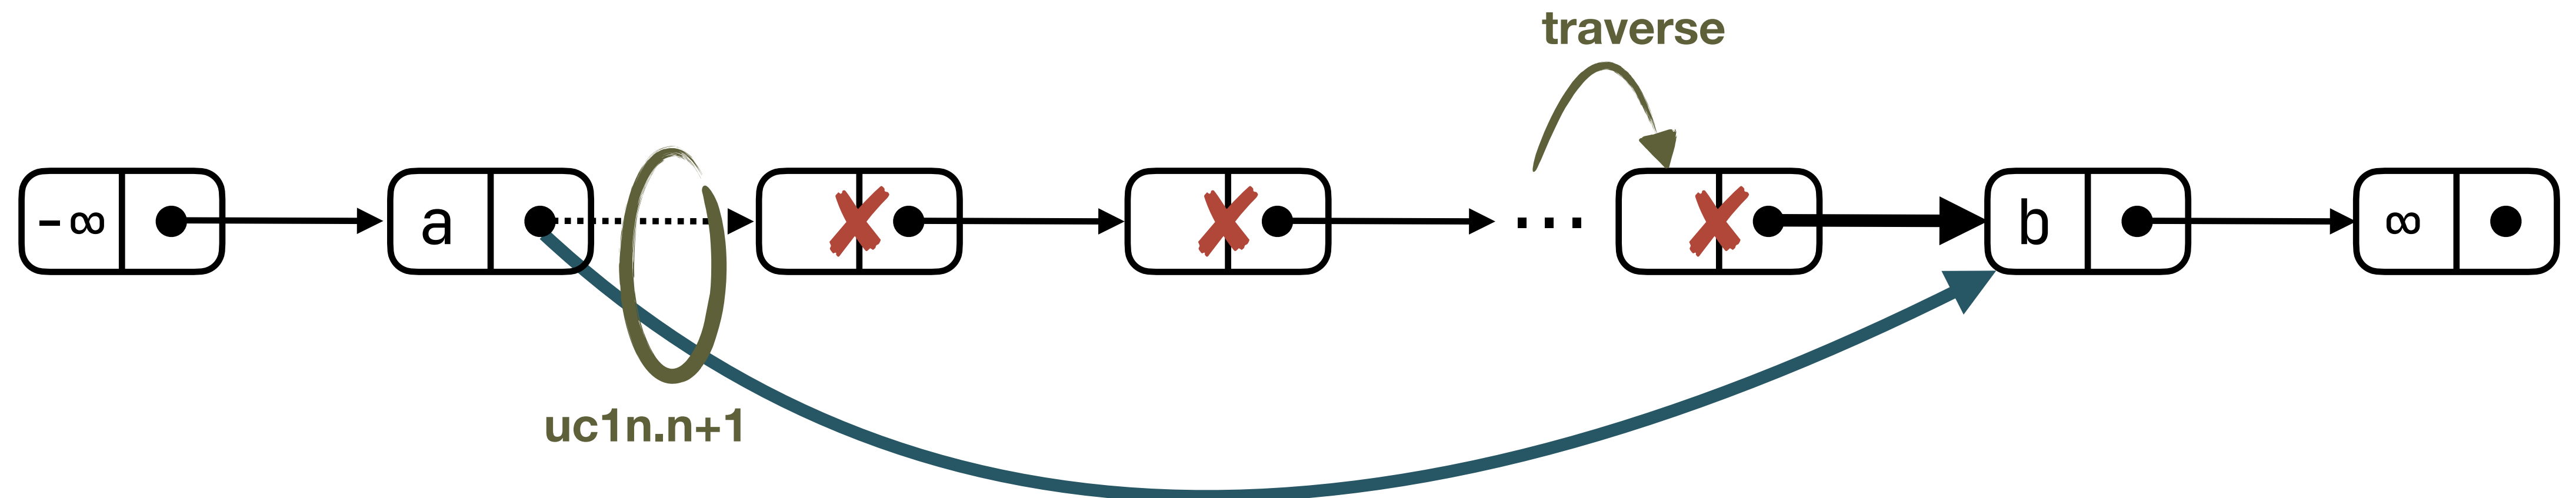
- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



Complex Updates

Gost update chunks

- compose complex update from smaller chunks **easier to automate**
- the chunks are hypothetical/ghost information for the proof
- can prove the composition **along the traversal**, no induction required



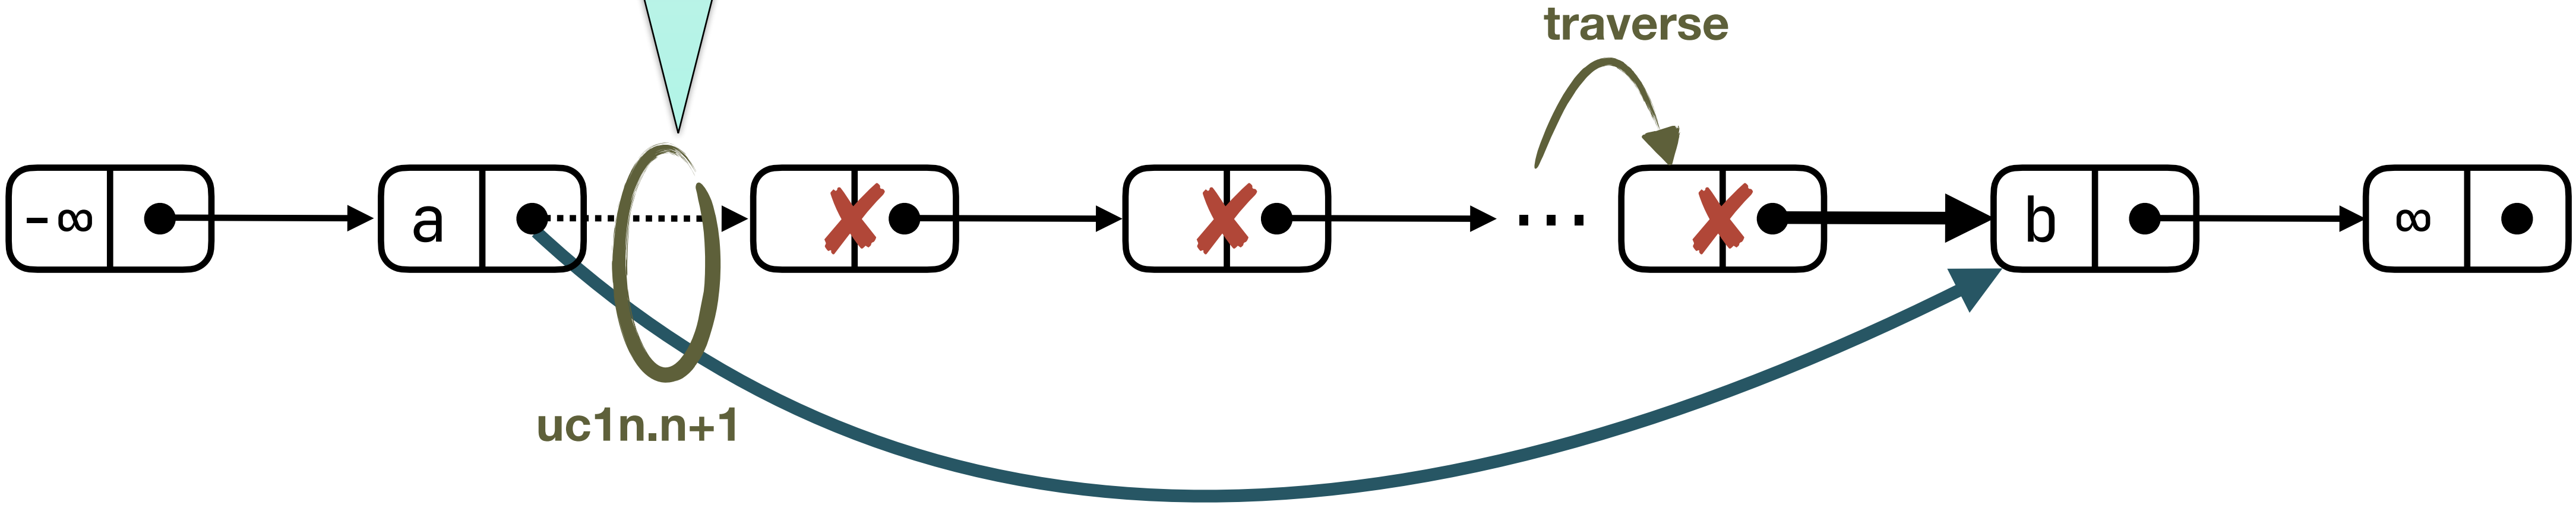
Complex Updates

Gost update chunks

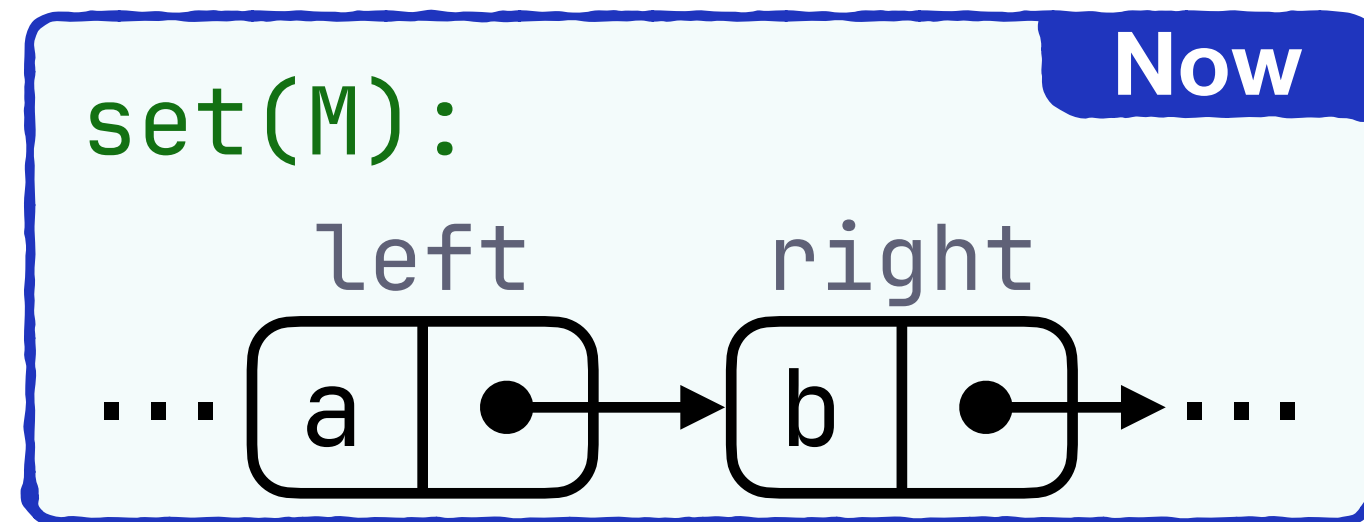
easier to automate

The update of interest, constructed from **local** (three nodes) ghost update chunks as the **limit of a computation**.

the proof
induction required



Example: removing marked nodes



```
while (right→mark):
```

```
    next = right→next;
```

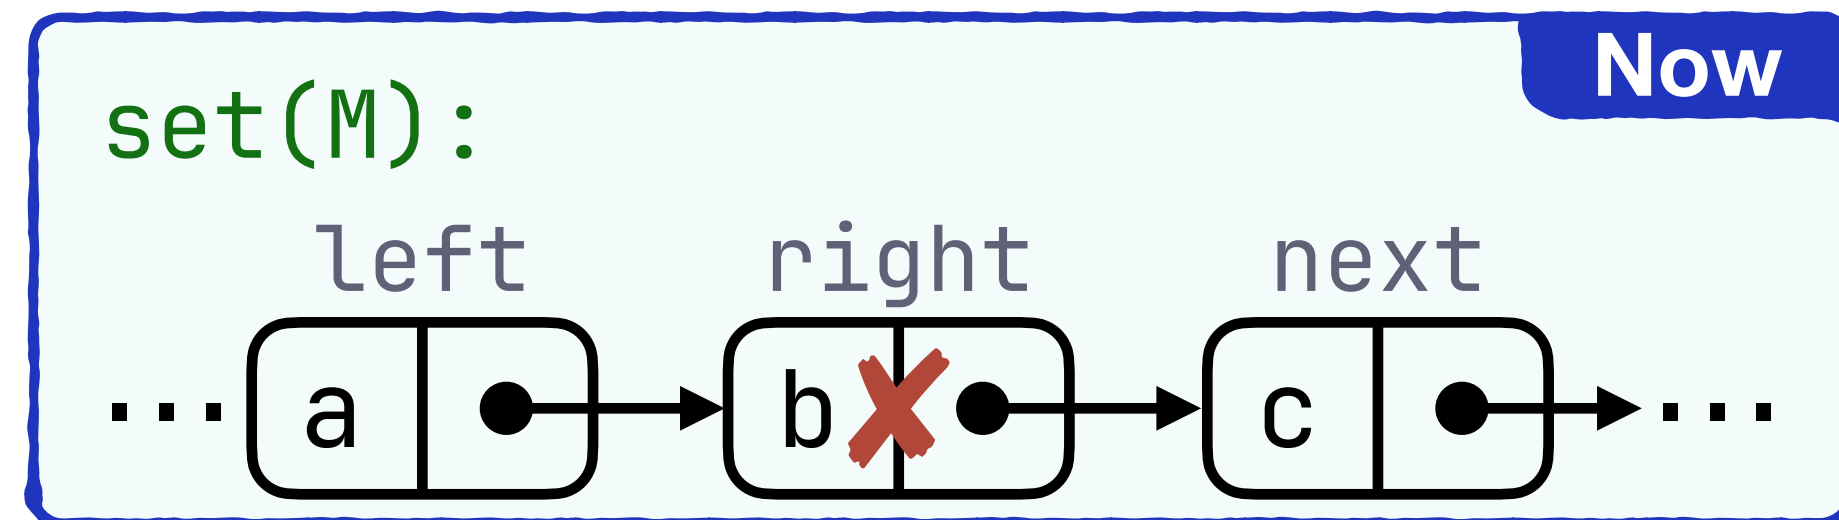
```
    right = next;
```

```
left→next = right;
```

Example: removing marked nodes

```
while (right → mark):
```

```
    next = right → next;
```



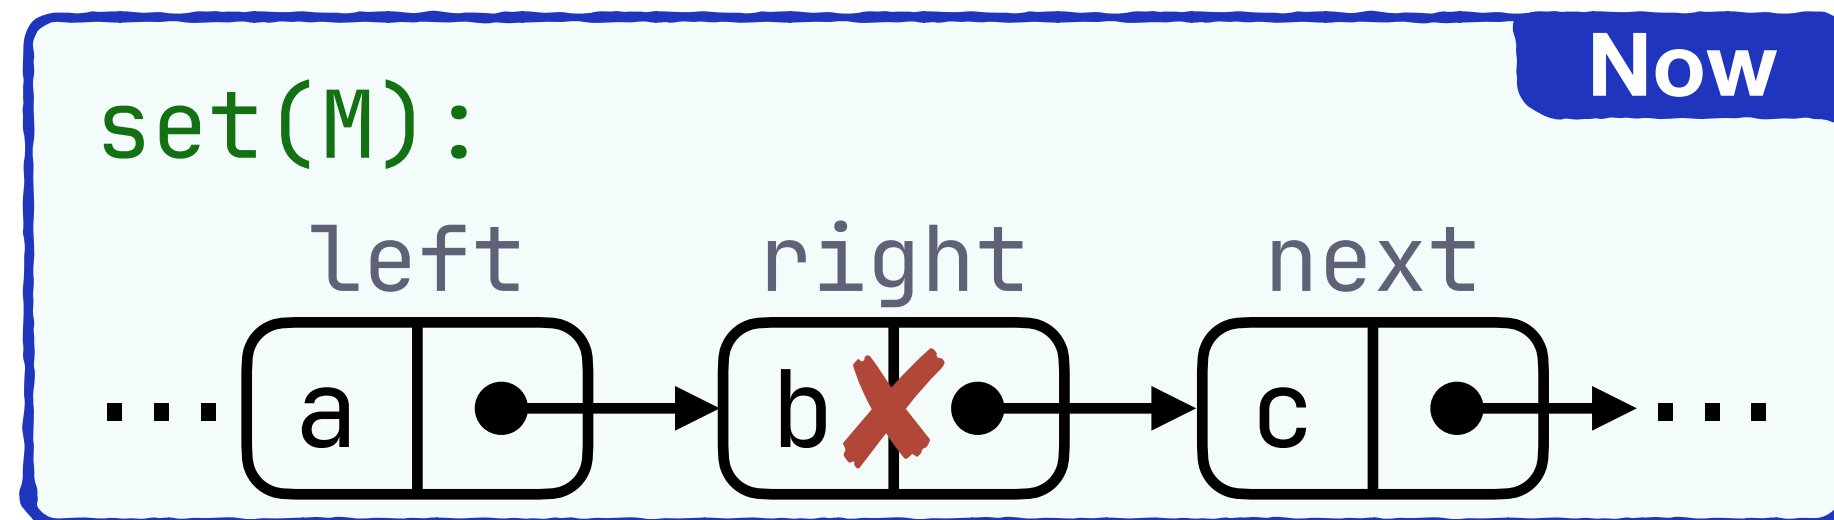
```
        right = next;
```

```
left → next = right;
```

Example: removing marked nodes

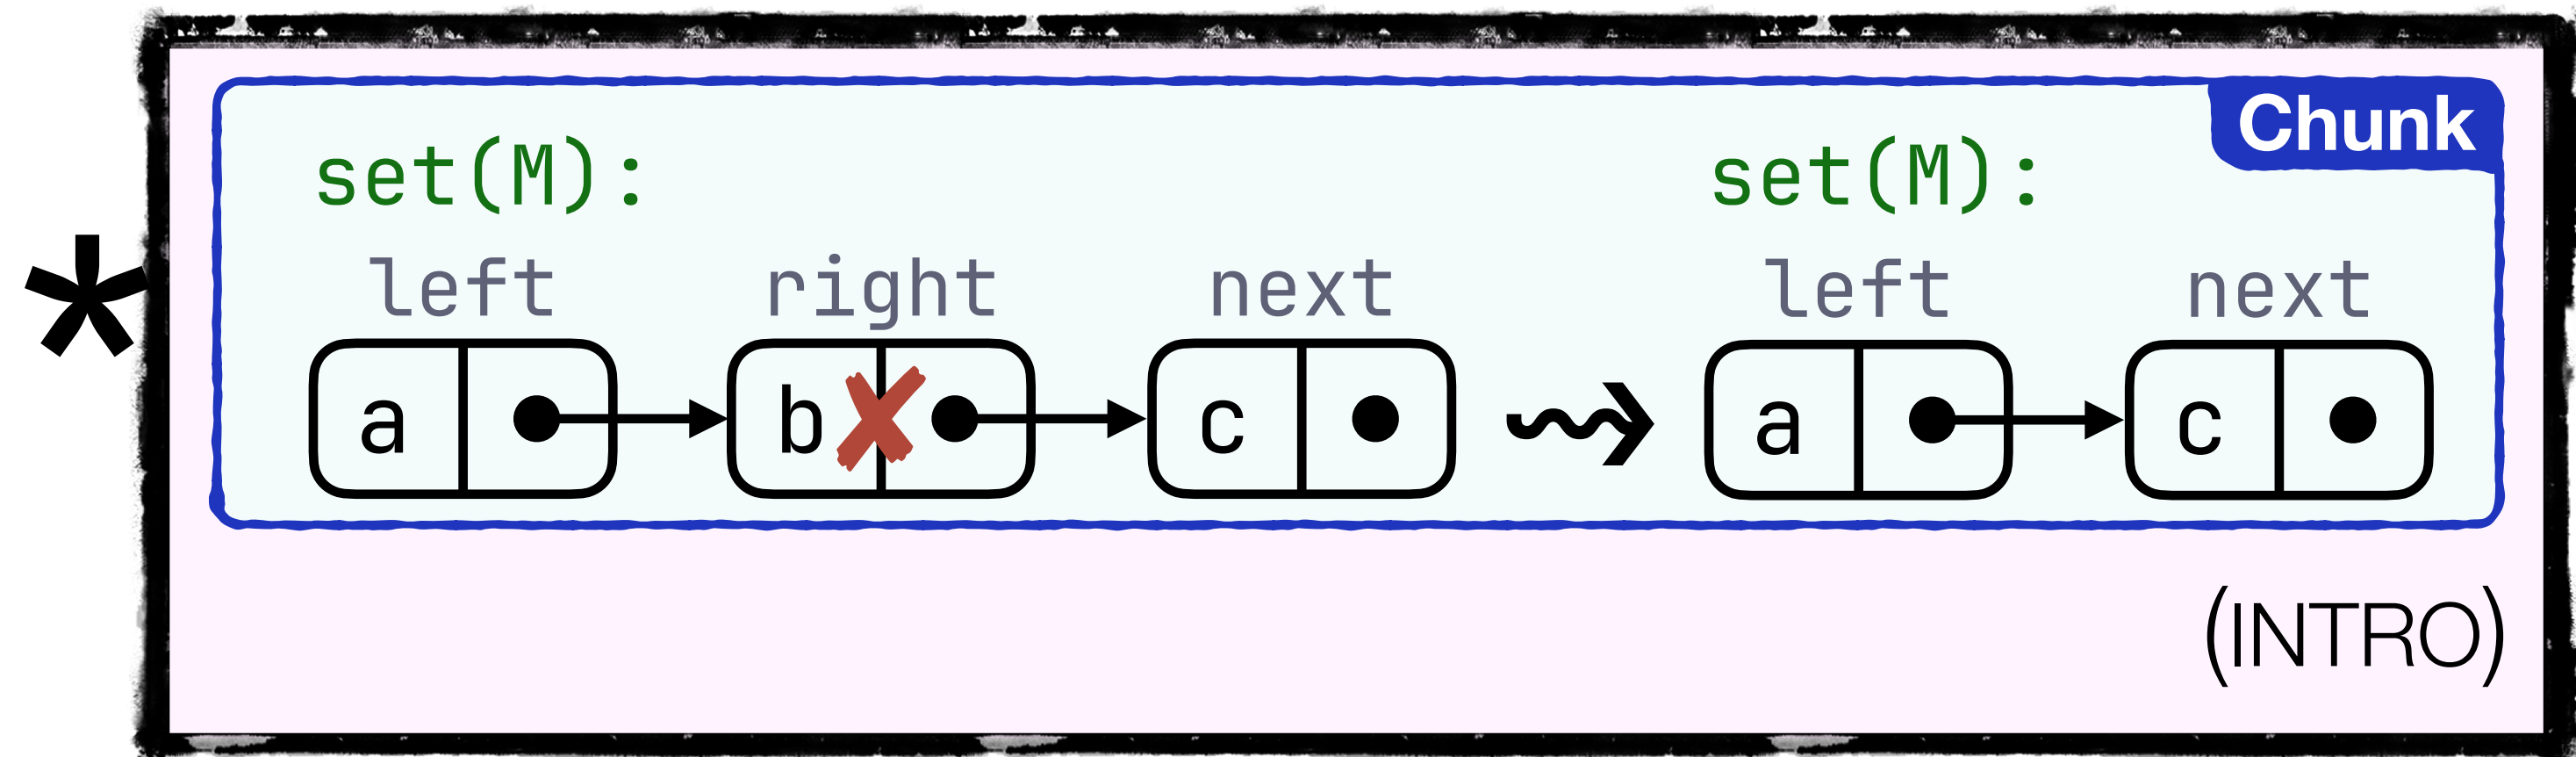
```
while (right → mark):
```

```
    next = right → next;
```



```
        right = next;
```

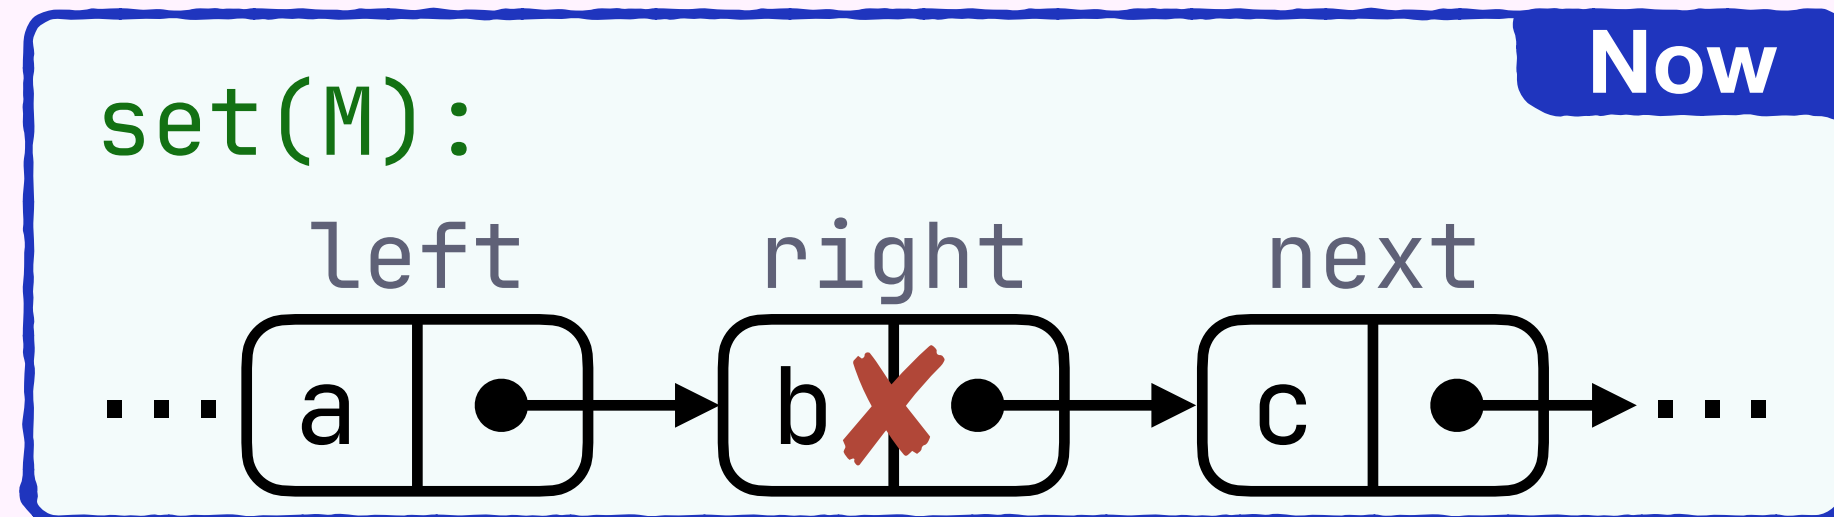
```
left → next = right;
```



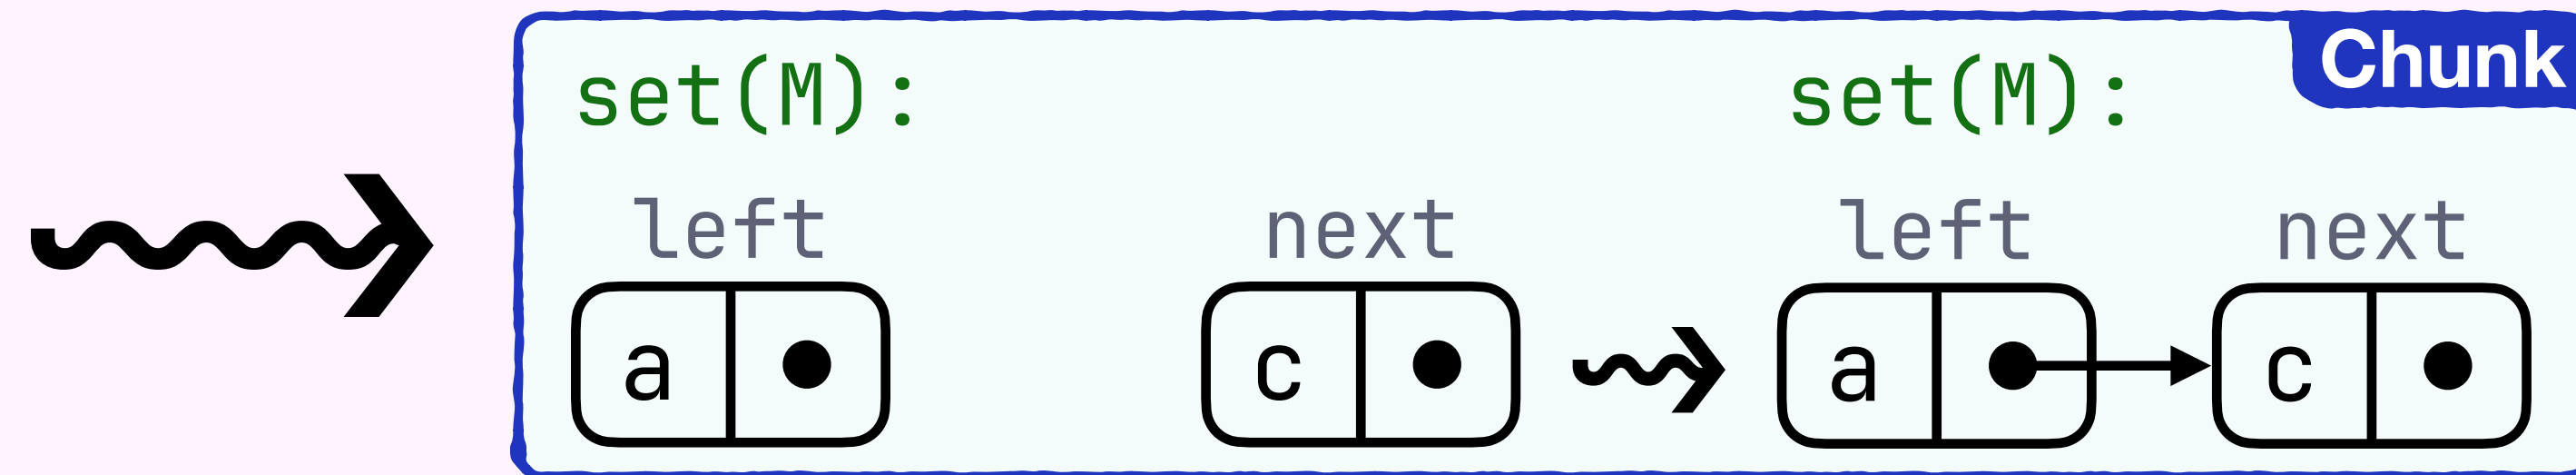
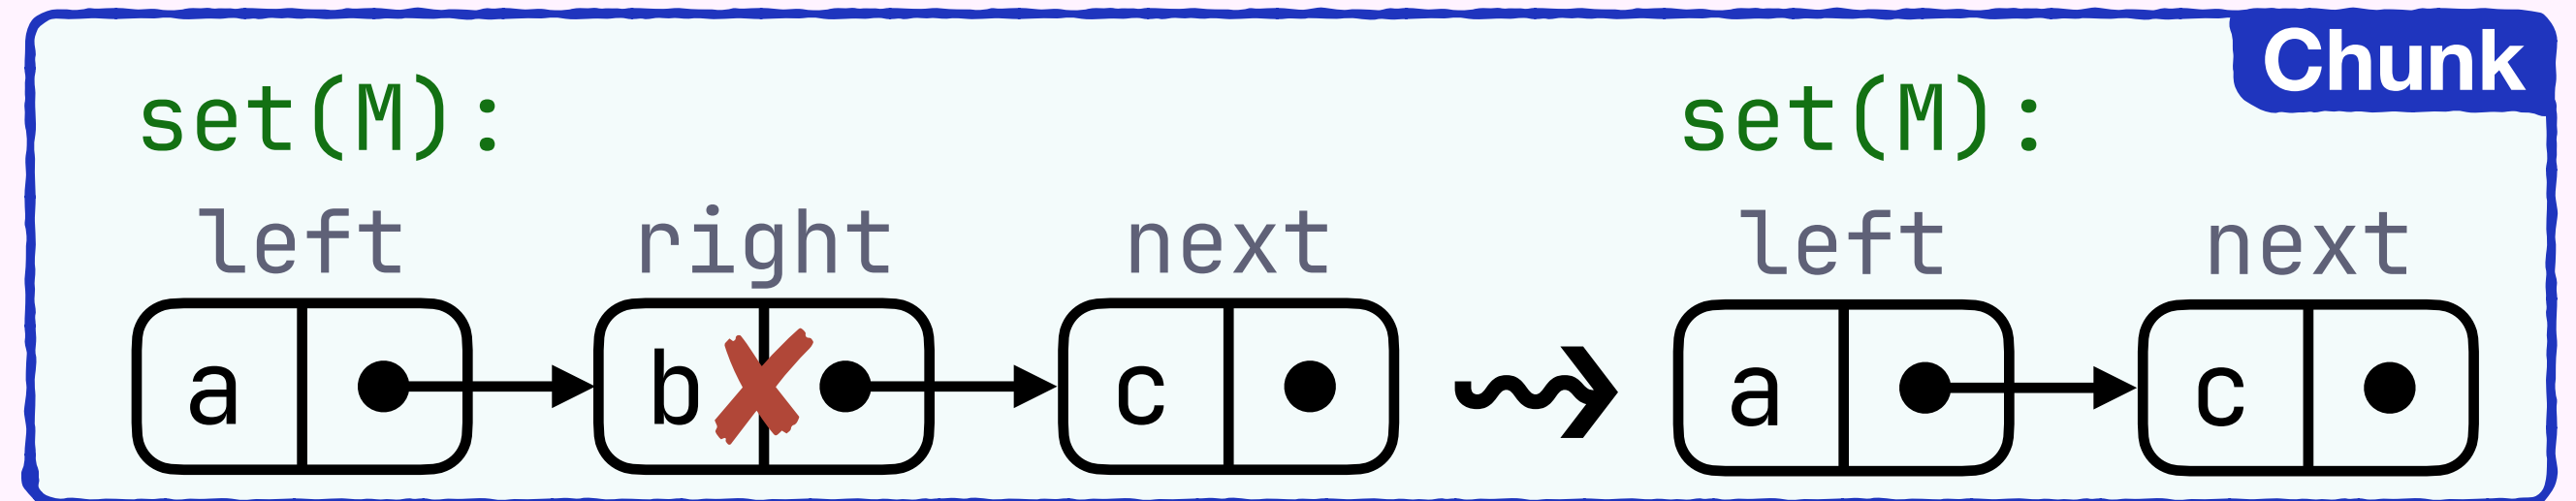
Example: removing marked nodes

```
while (right → mark):
```

```
    next = right → next;
```



*

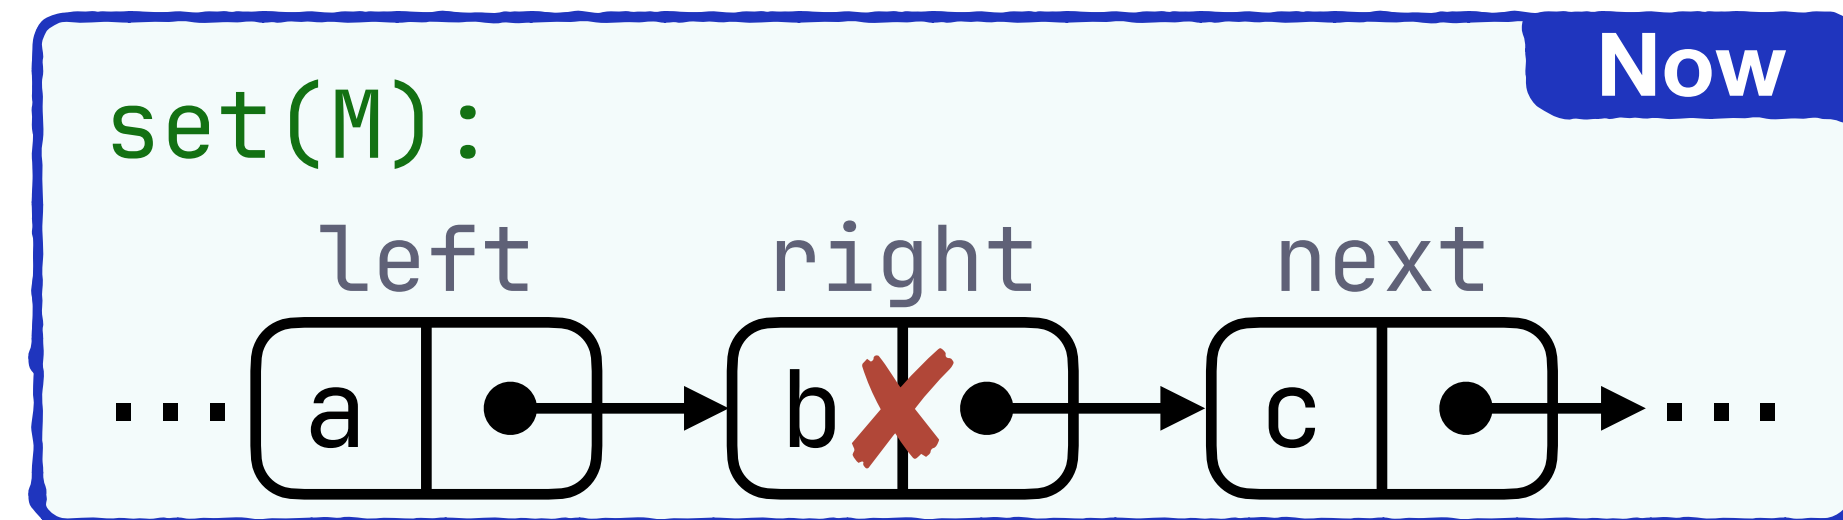


(DISCHARGE)

Example: removing marked nodes

```
while (right → mark):
```

```
    next = right → next;
```



*

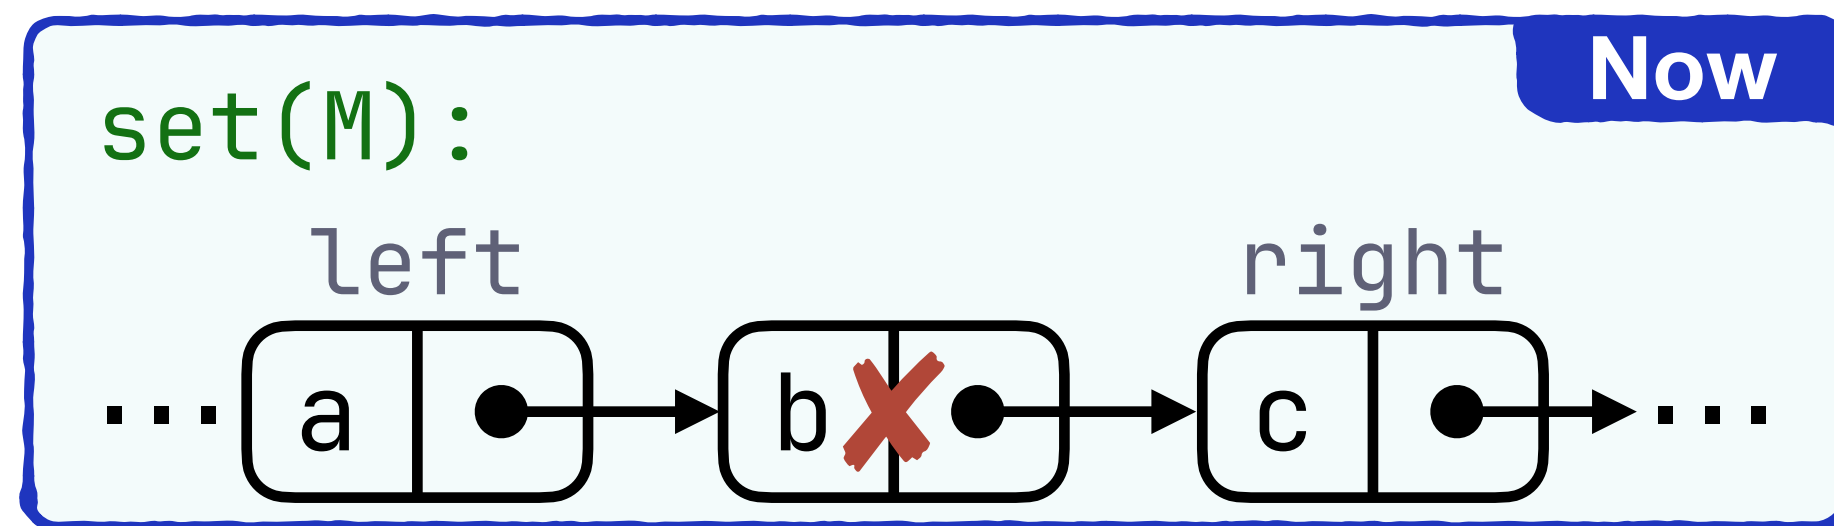


```
        right = next;
```

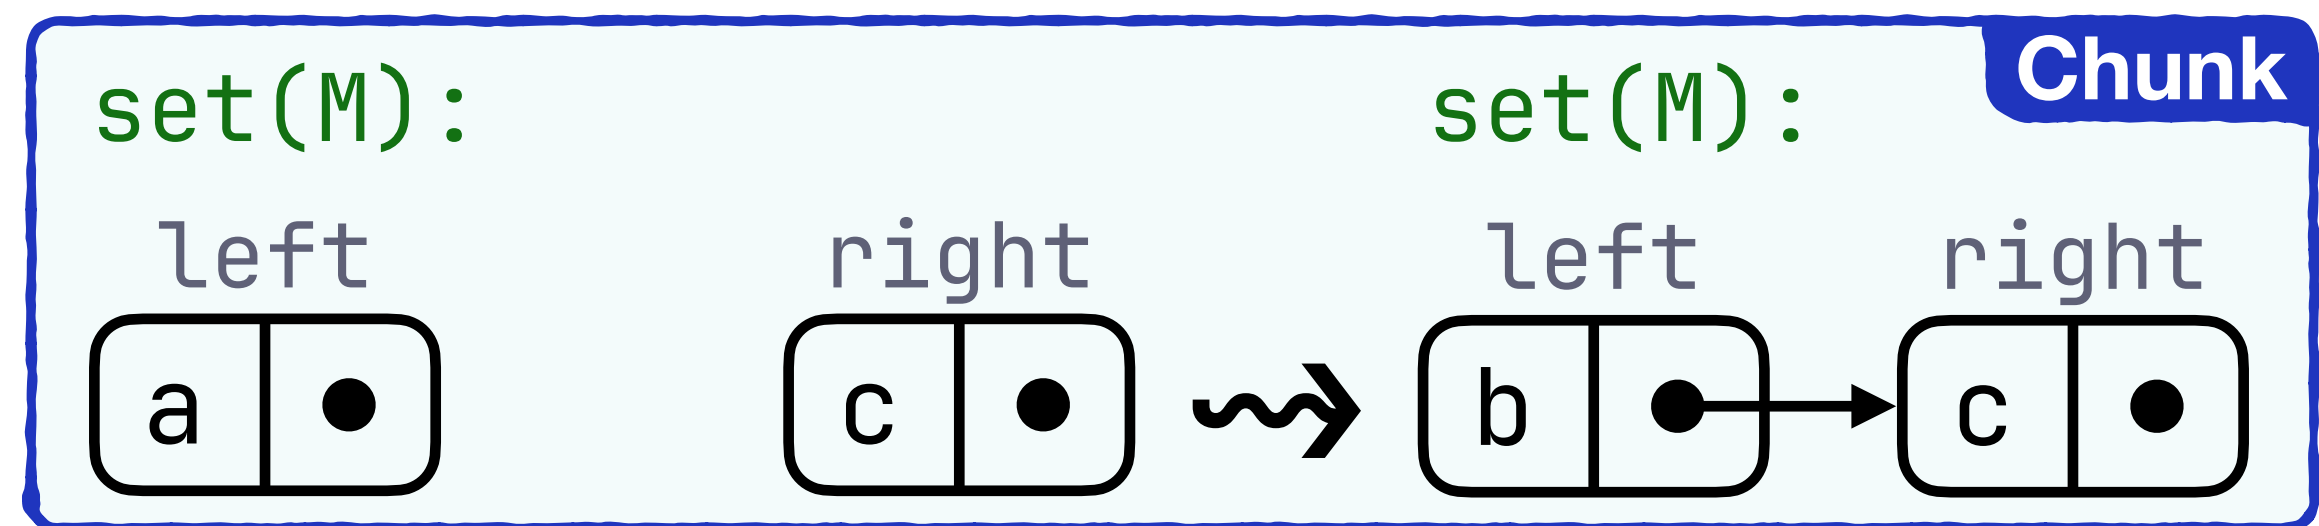
```
left → next = right;
```

Example: removing marked nodes

```
while (right → mark):  
    next = right → next;  
    right = next;
```



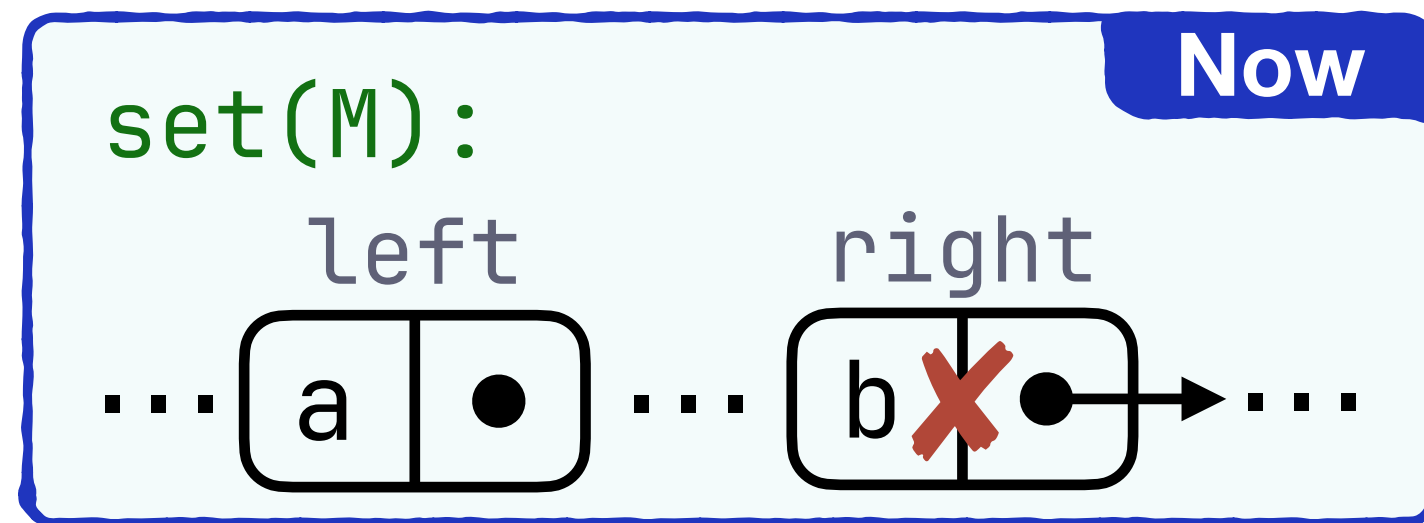
*



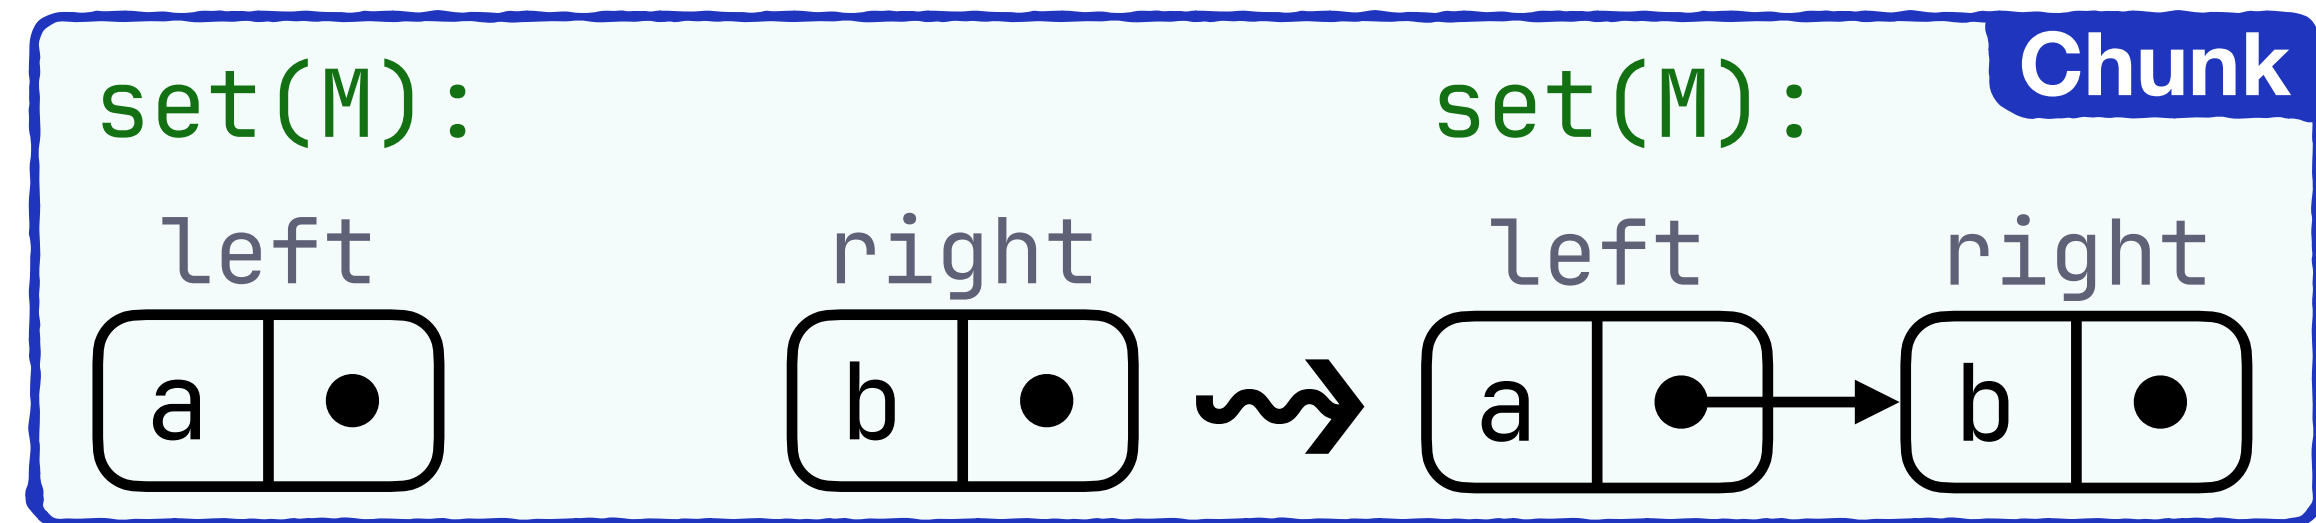
```
left → next = right;
```

Example: removing marked nodes

`while (right → mark):`



*



`next = right → next;`

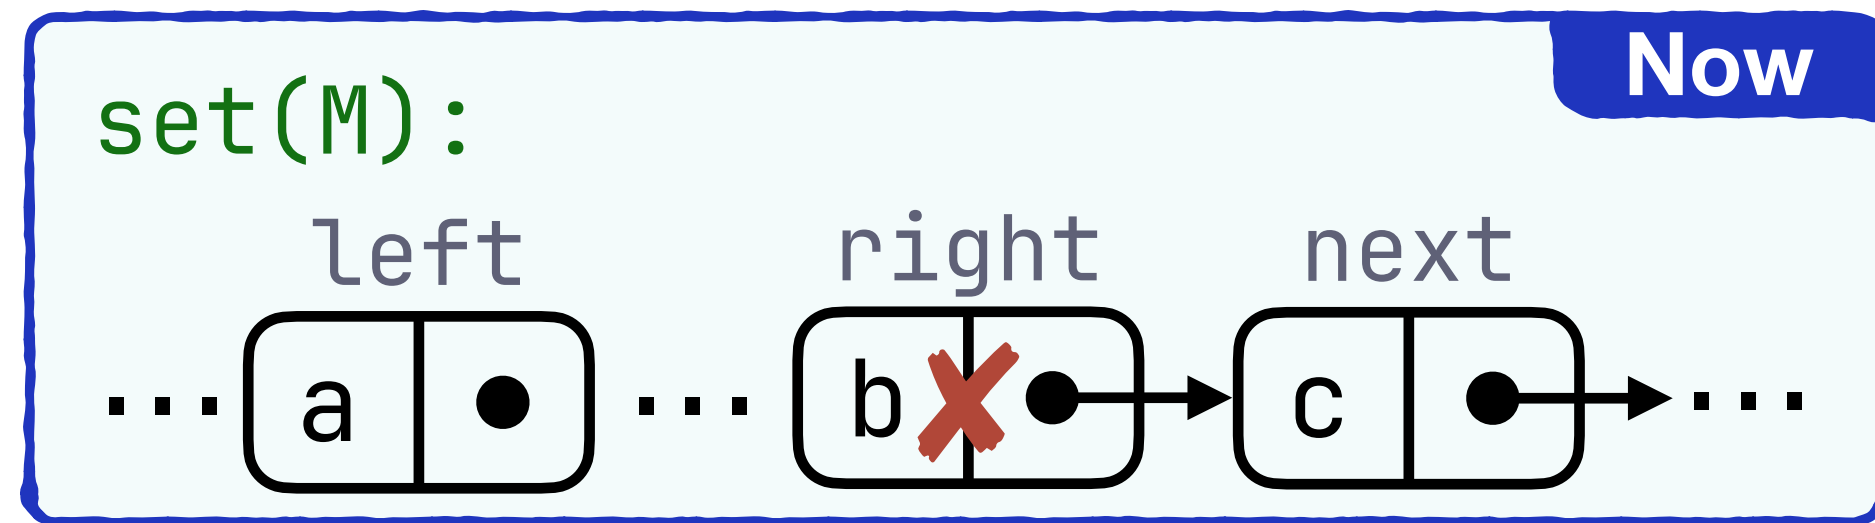
`right = next;`

`left → next = right;`

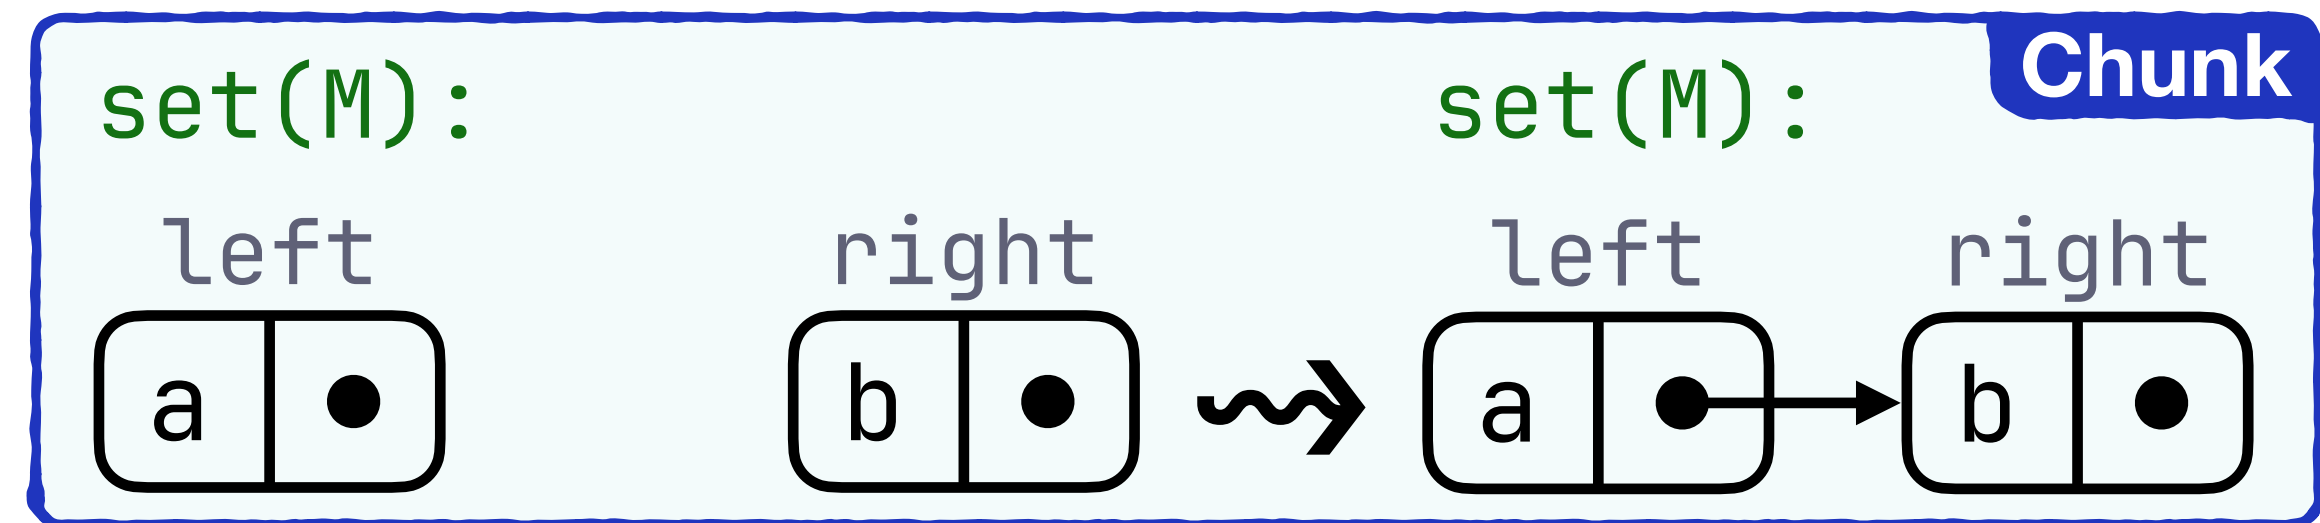
Example: removing marked nodes

```
while (right→mark):
```

```
    next = right→next;
```



*

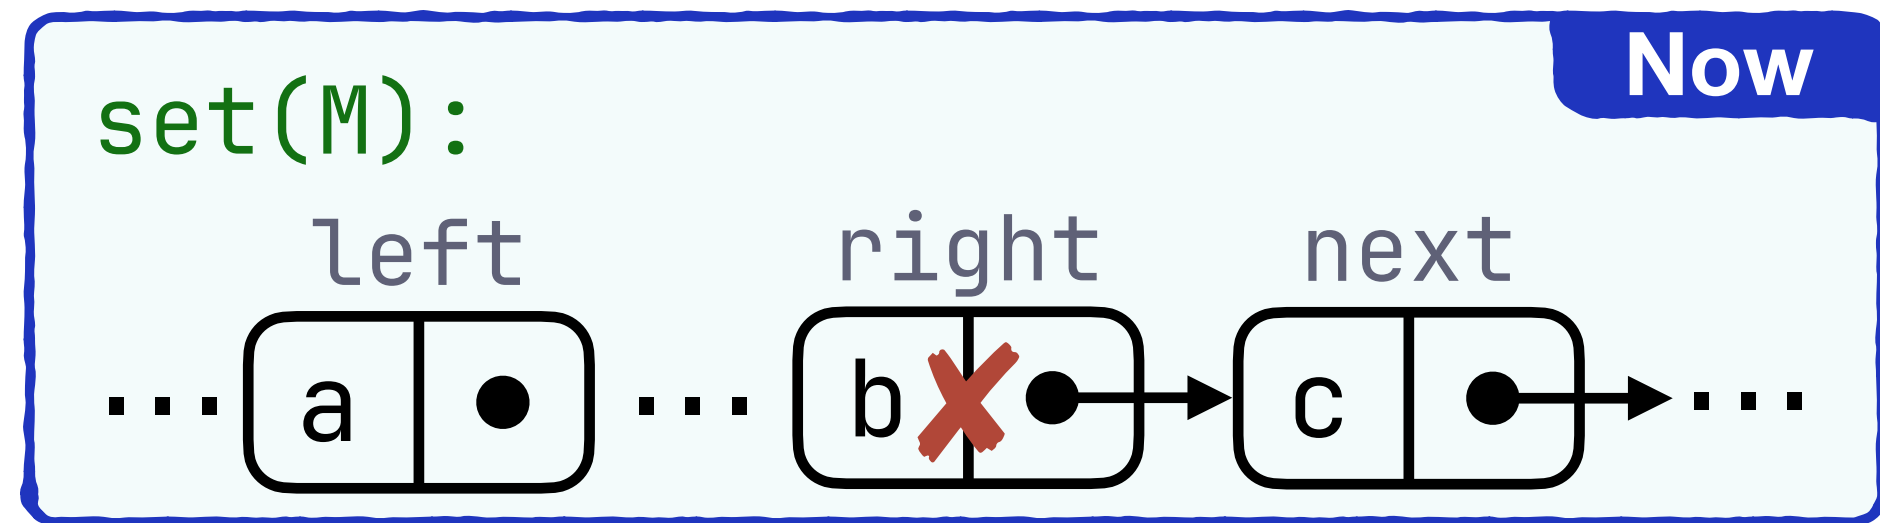
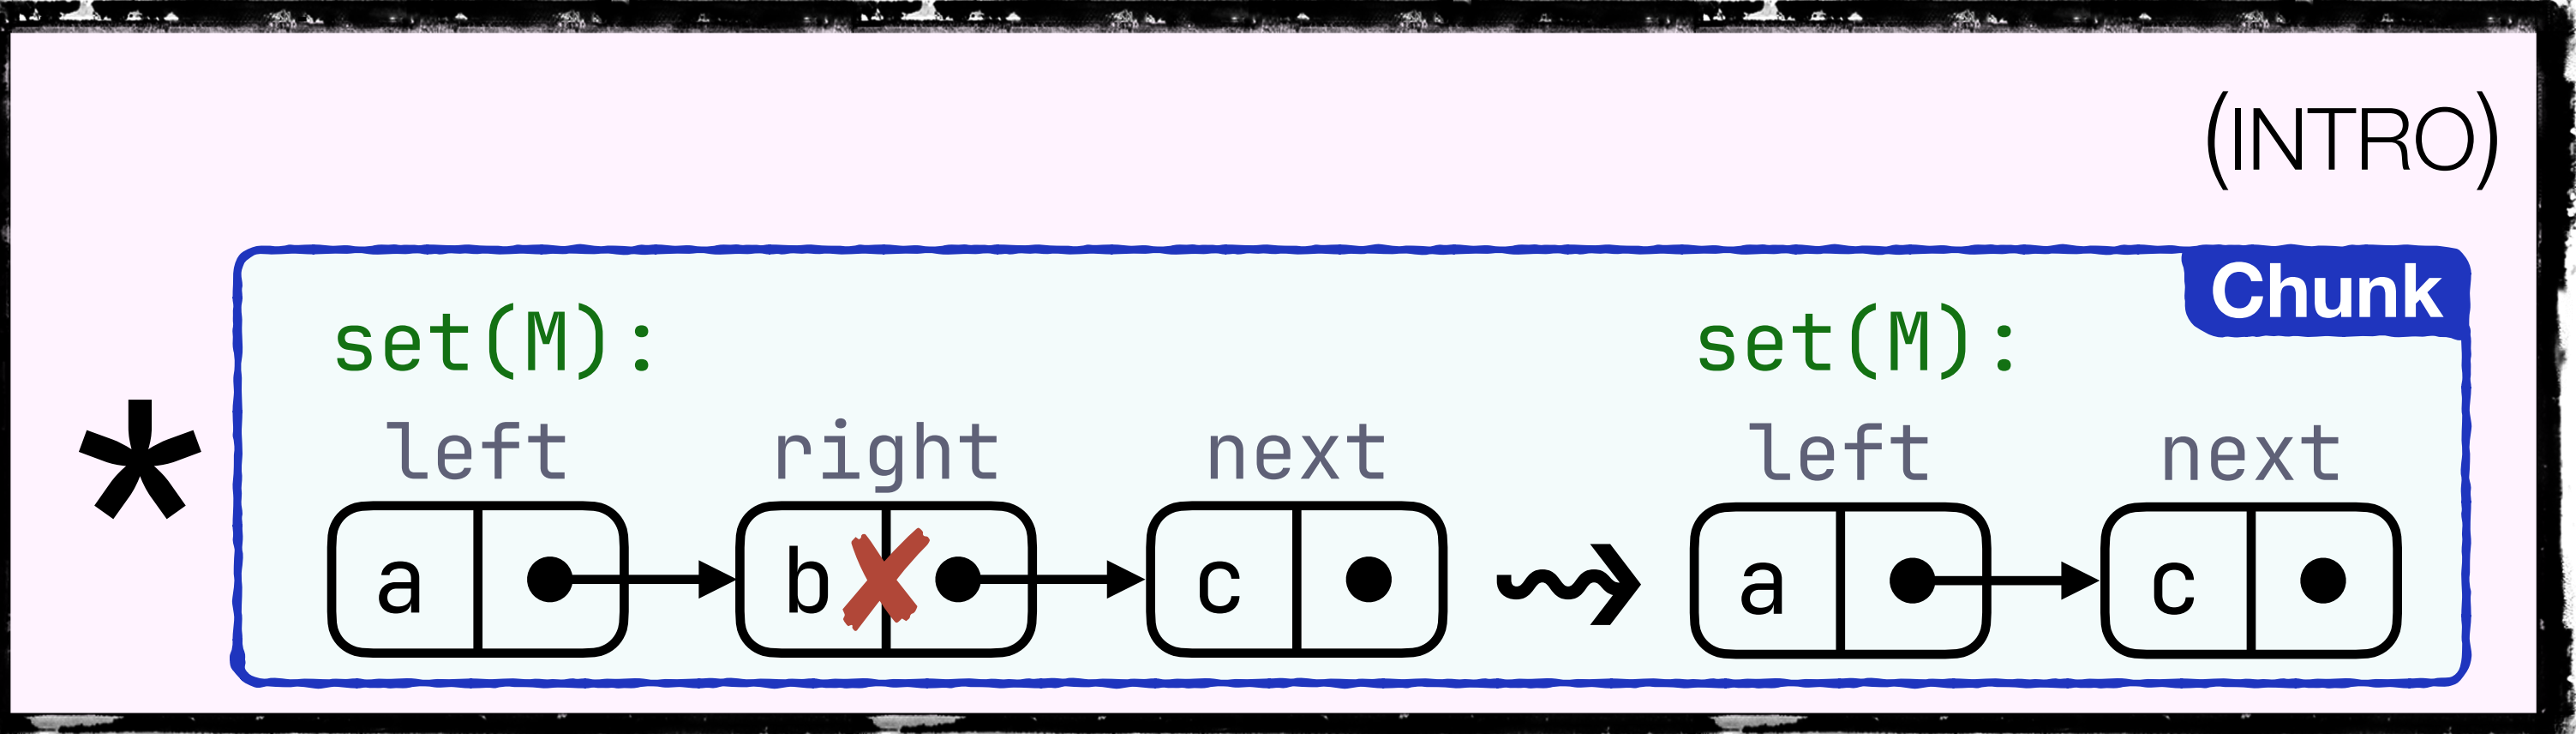


```
        right = next;
```

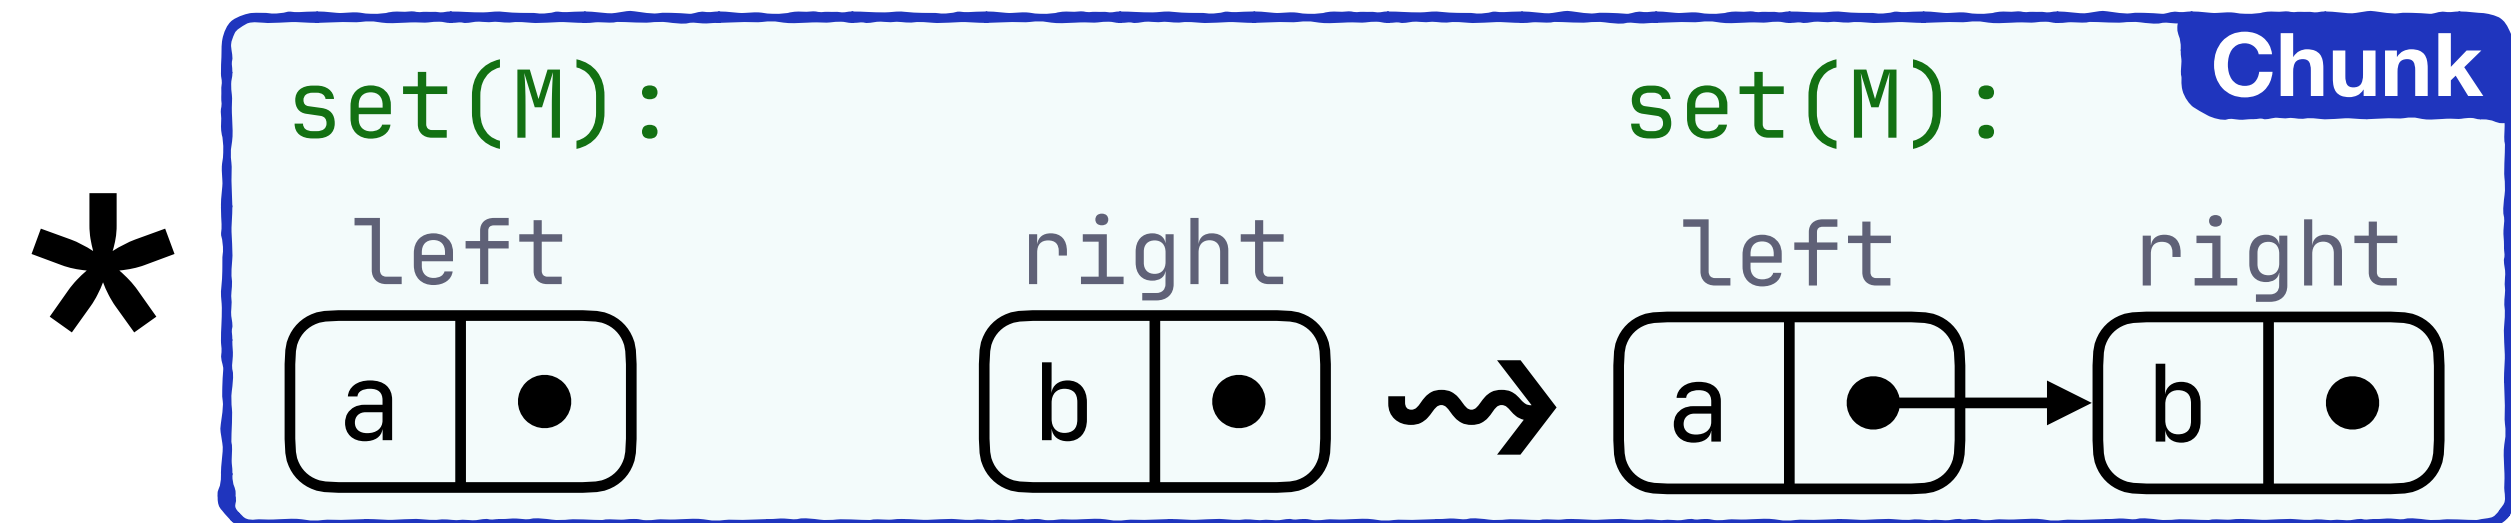
```
left→next = right;
```

Example: removing marke

```
while (right → mark):  
    next = right → next
```



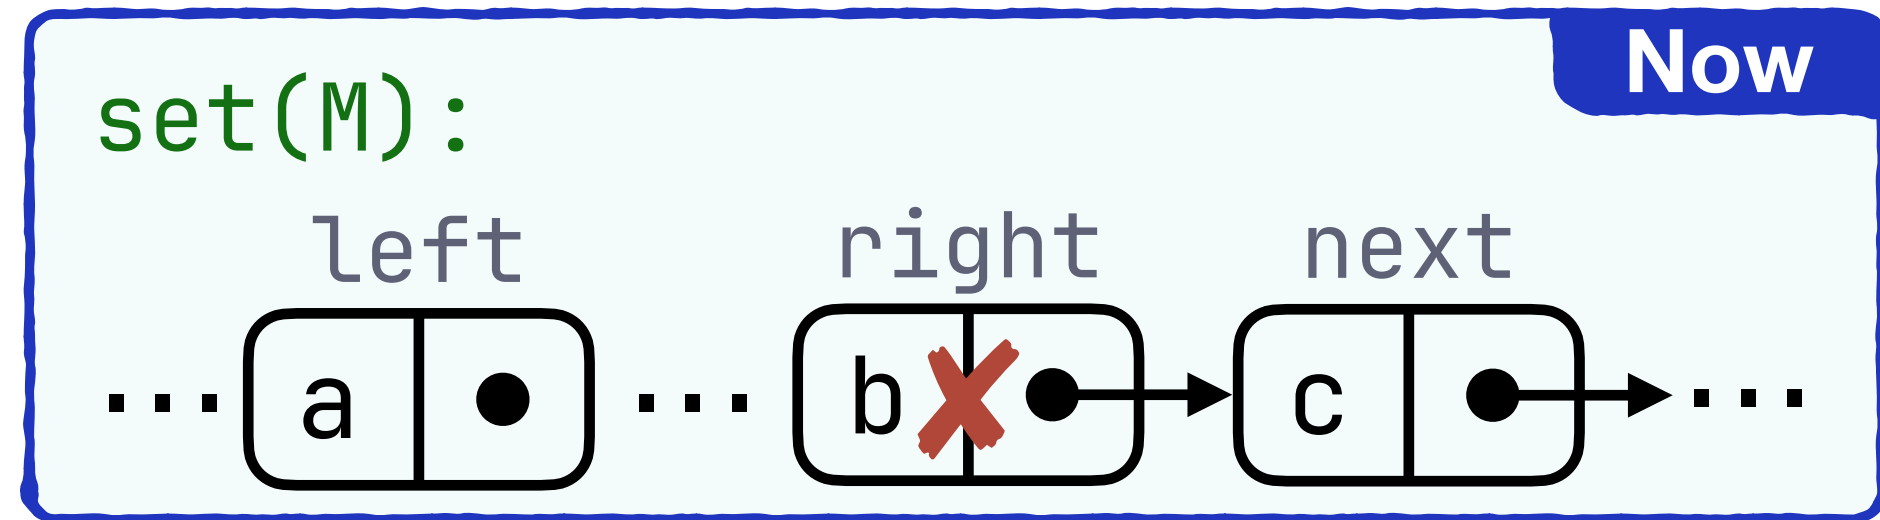
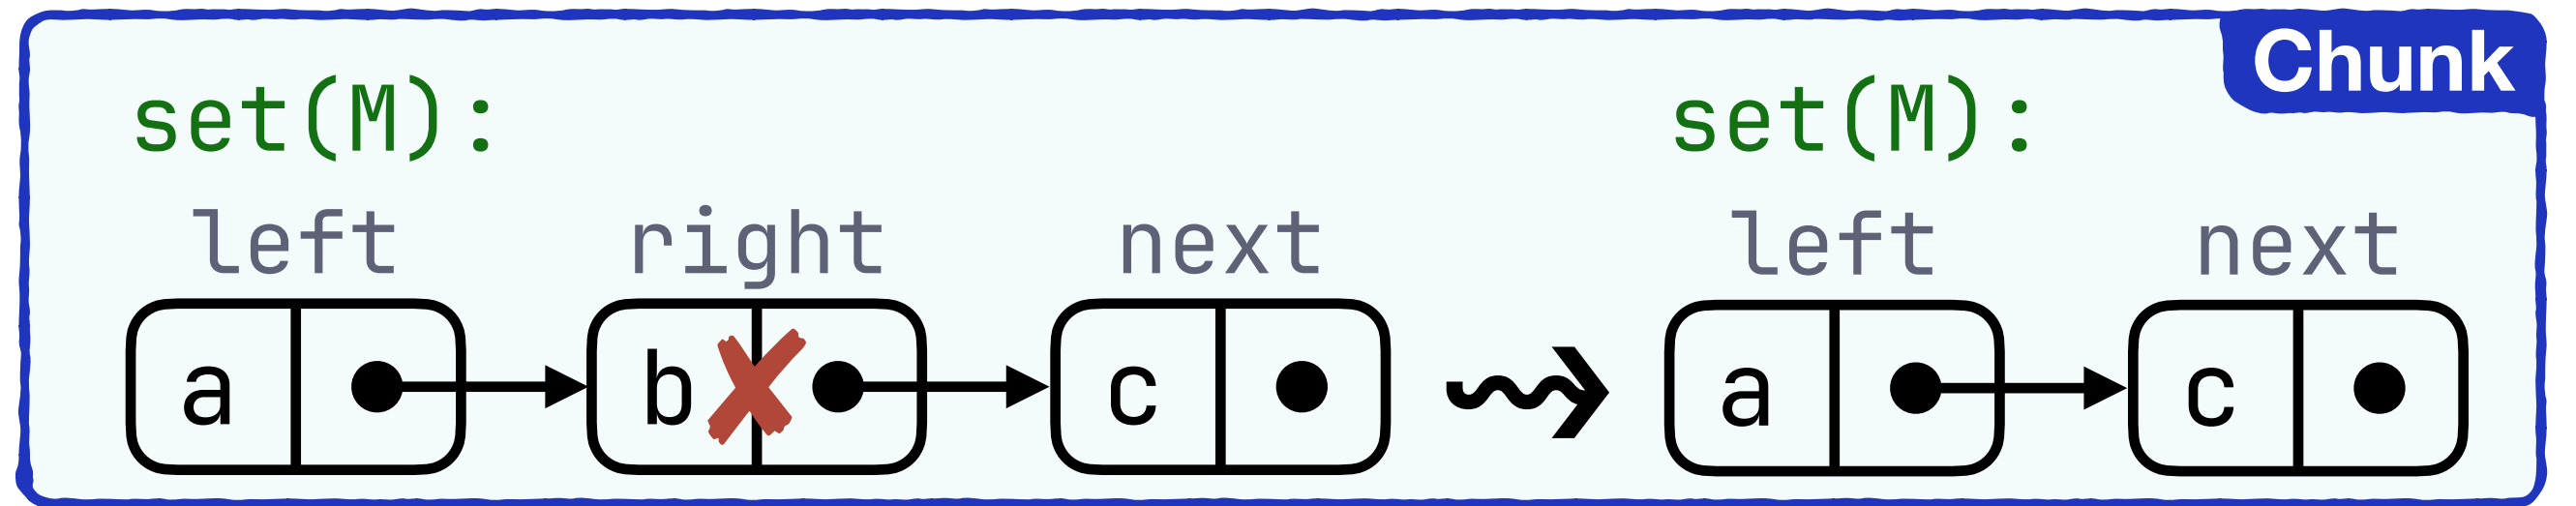
```
right = next;  
left → next = right;
```



Example: removing marked nodes

```
while (right → mark):  
    next = right → next;
```

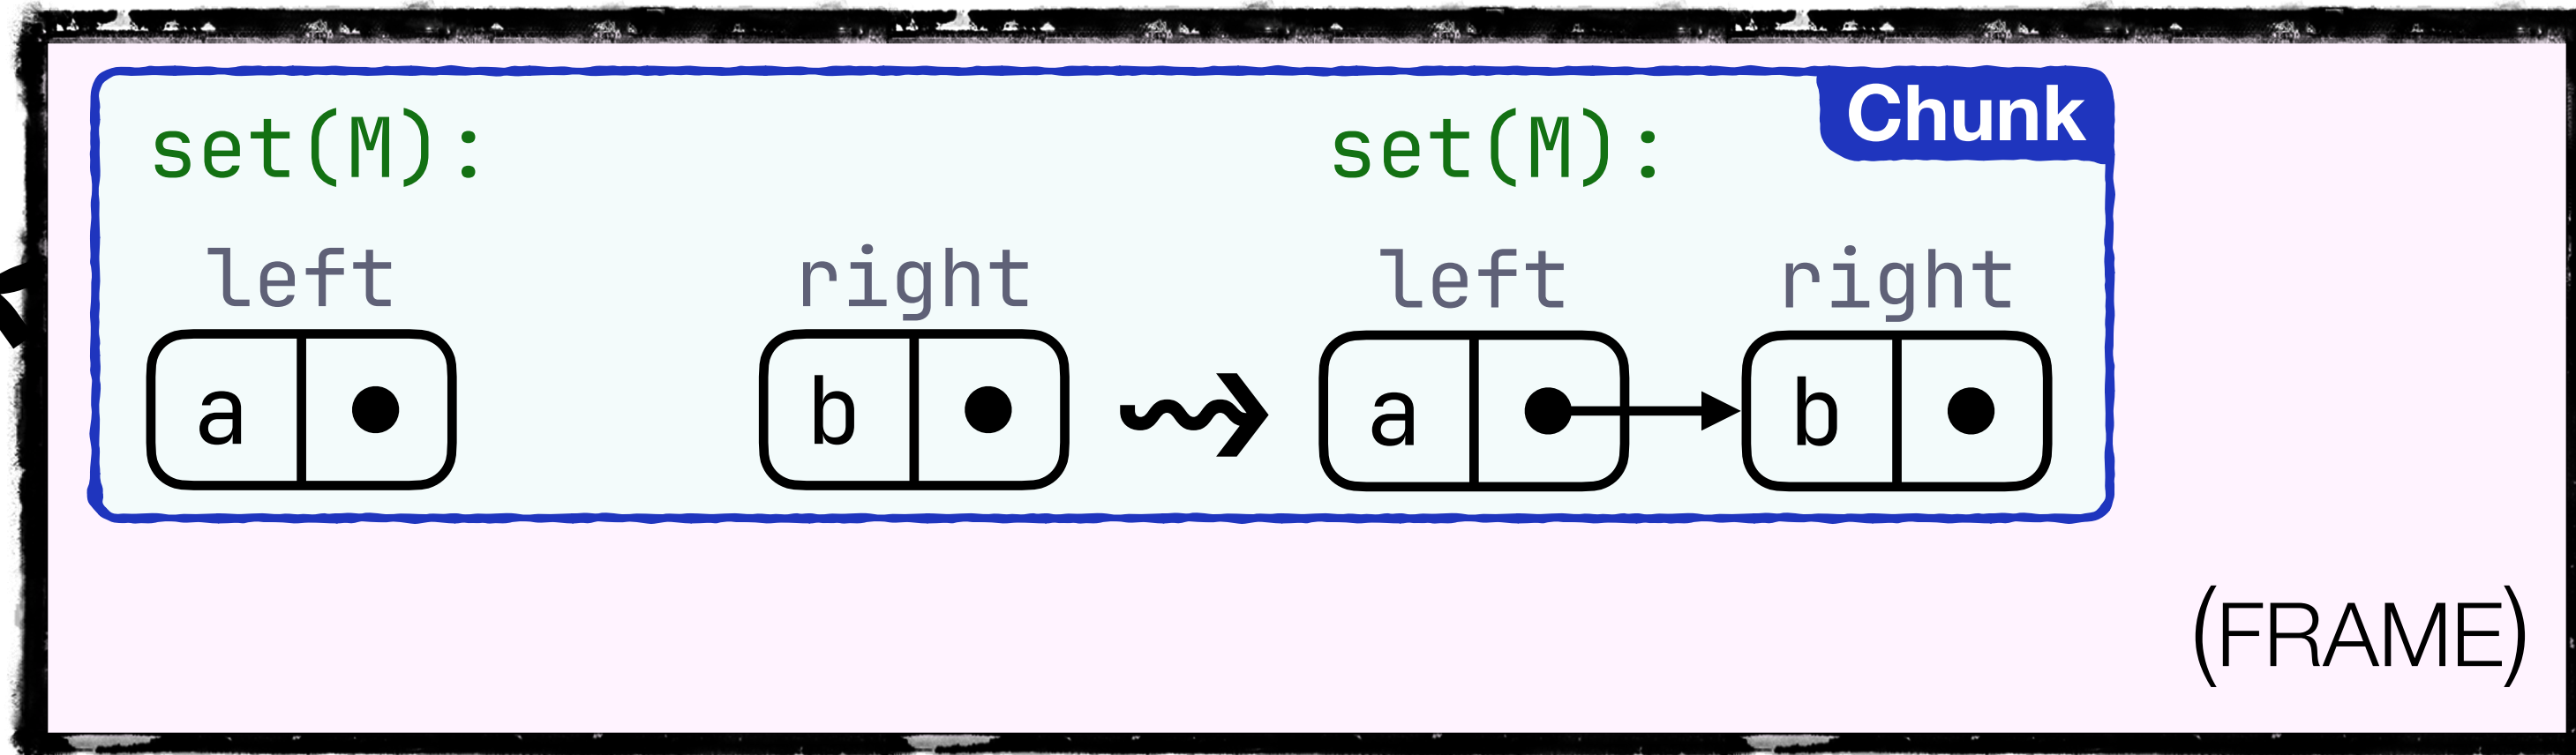
*



```
right = next;
```

```
left → next = right;
```

*

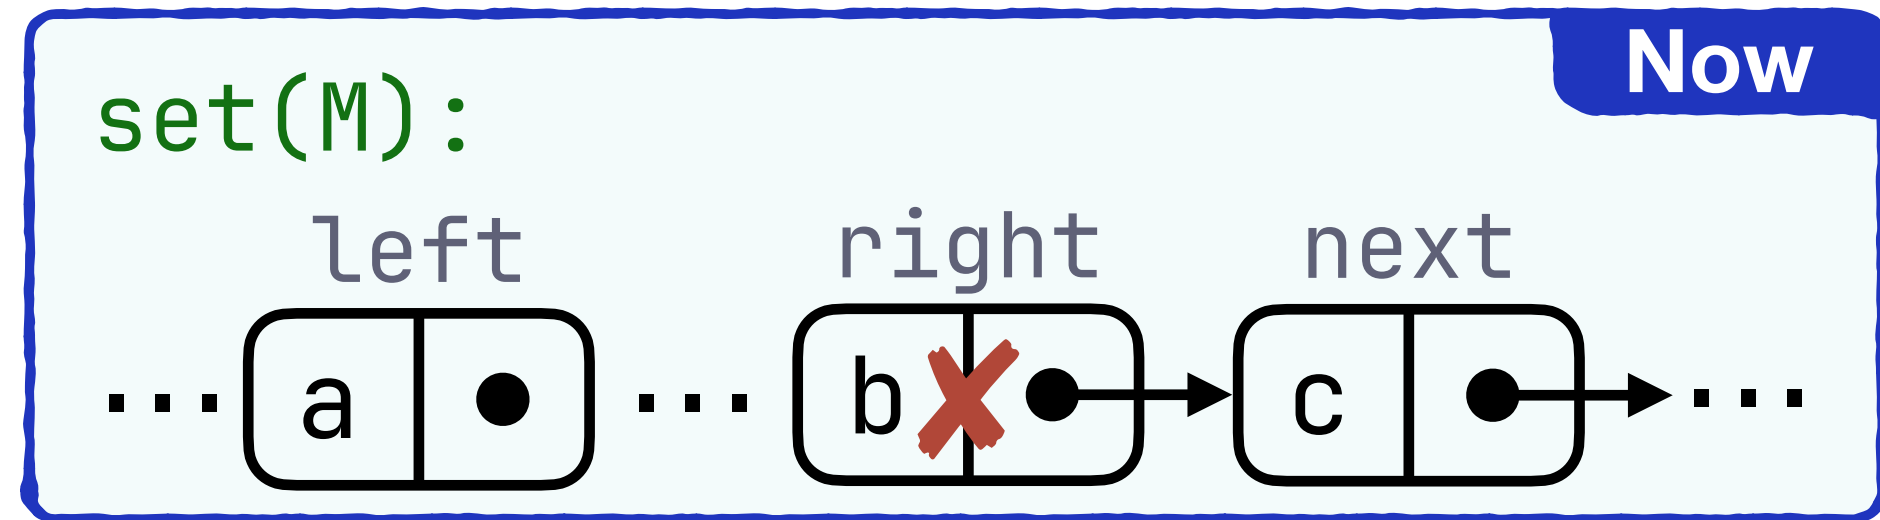
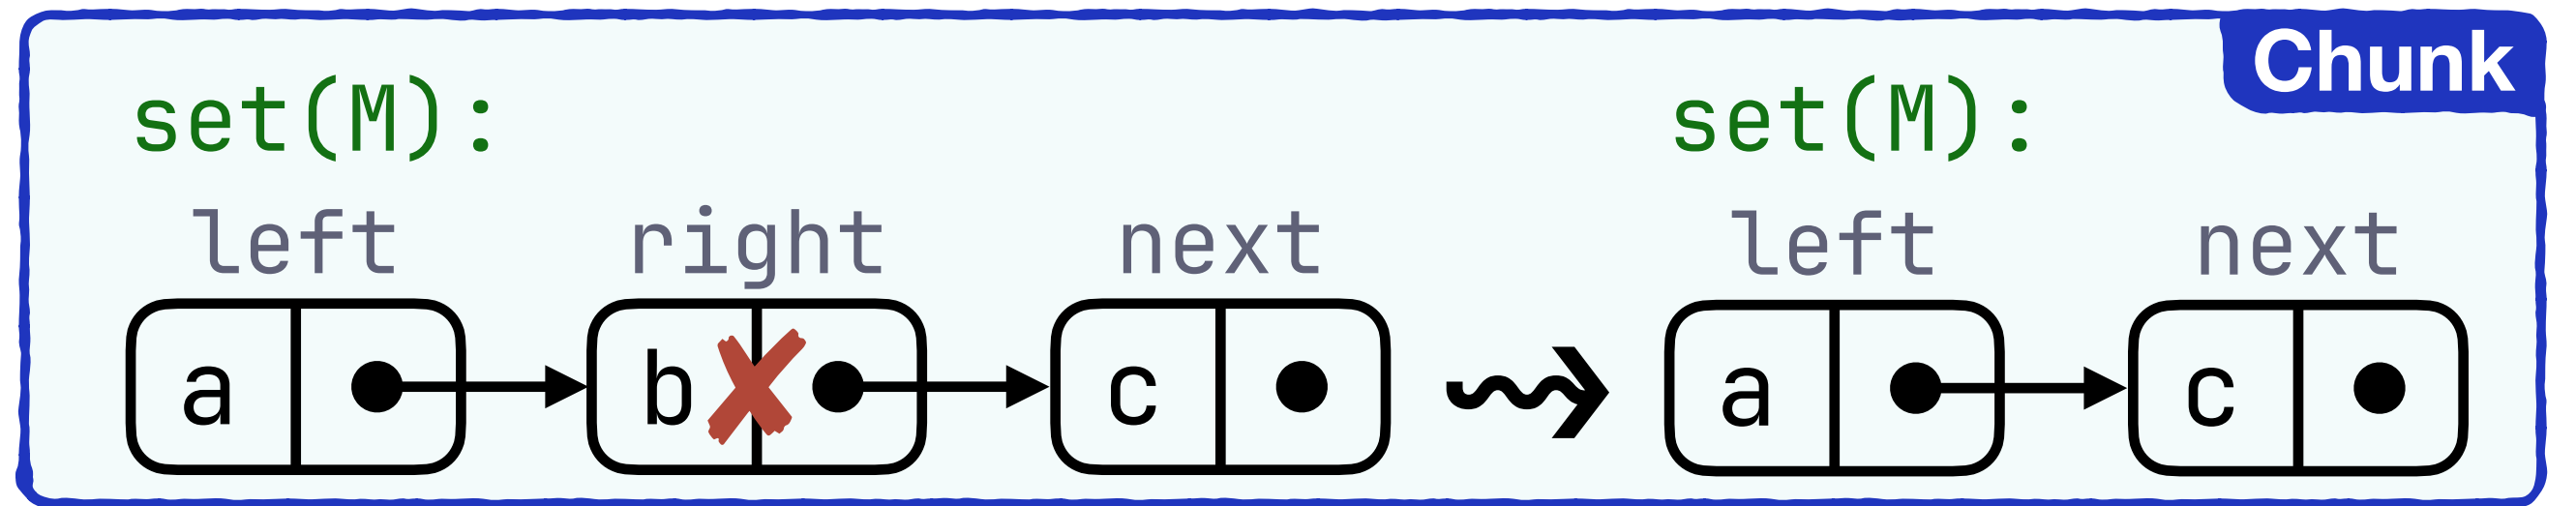


(FRAME)

Example: removing marked nodes

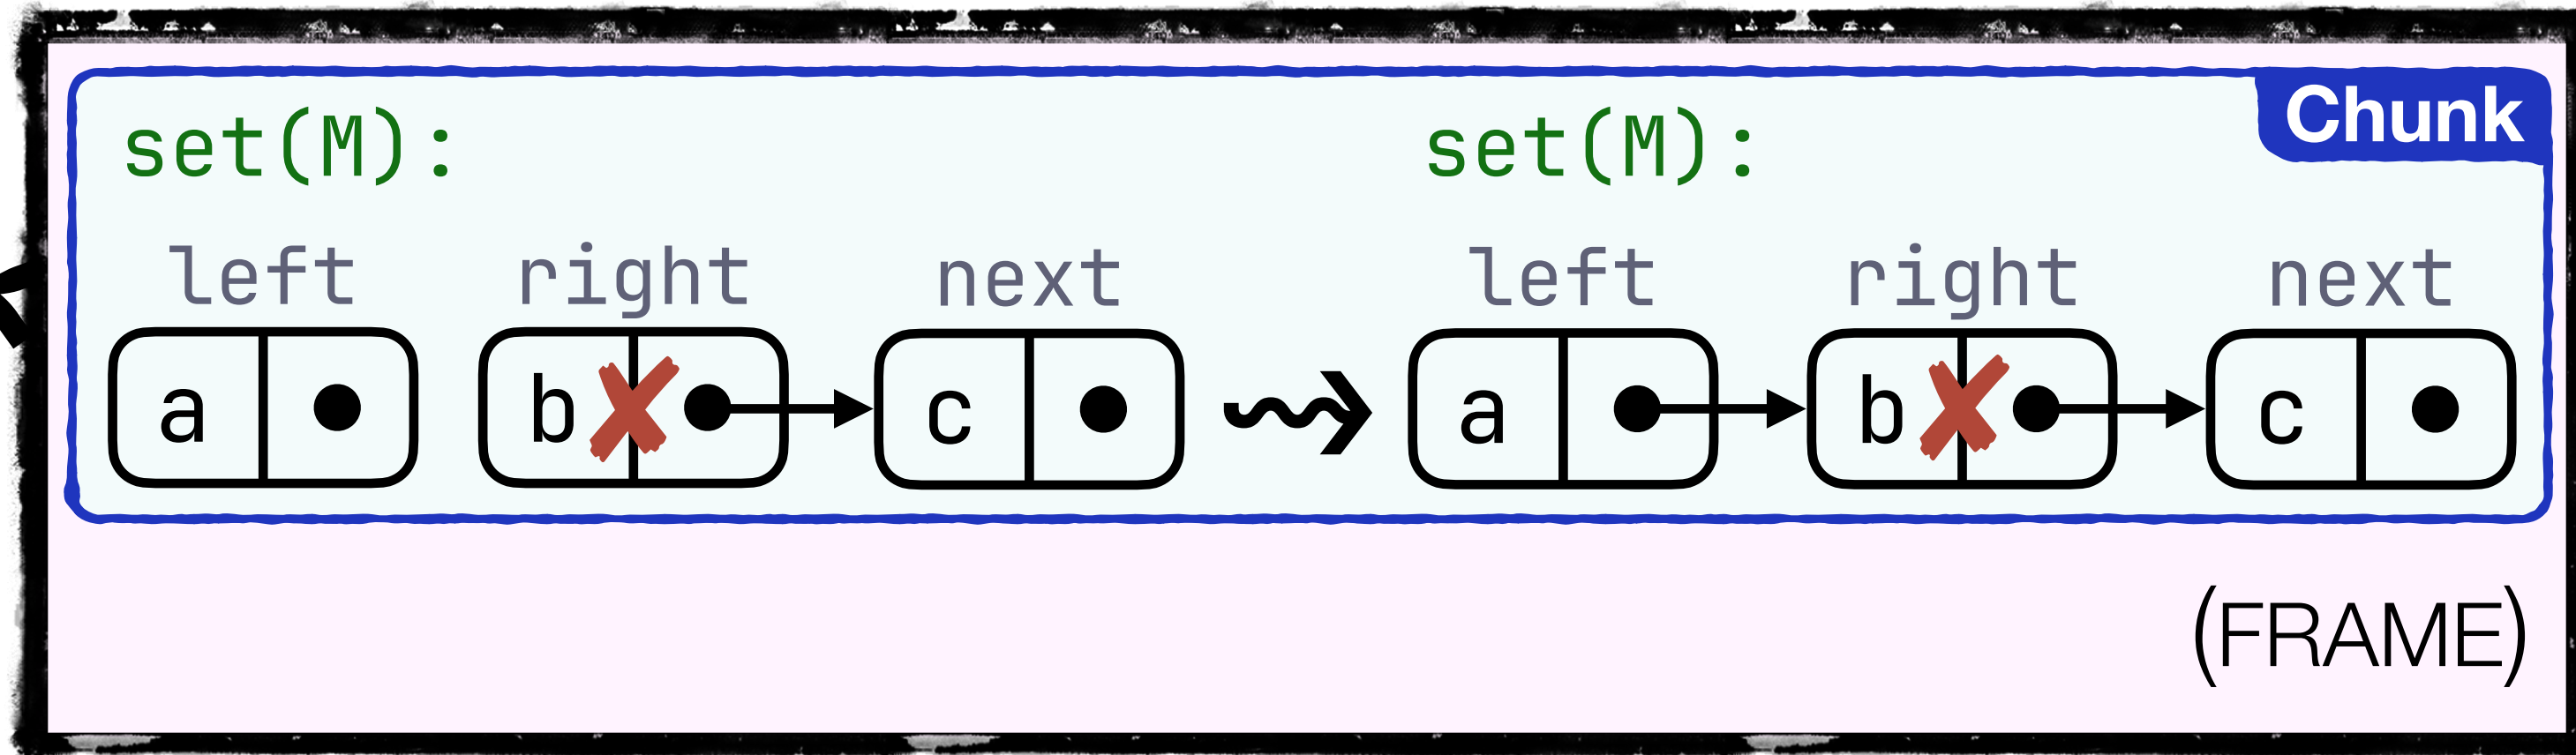
```
while (right → mark):  
    next = right → next;
```

*



```
right = next;
```

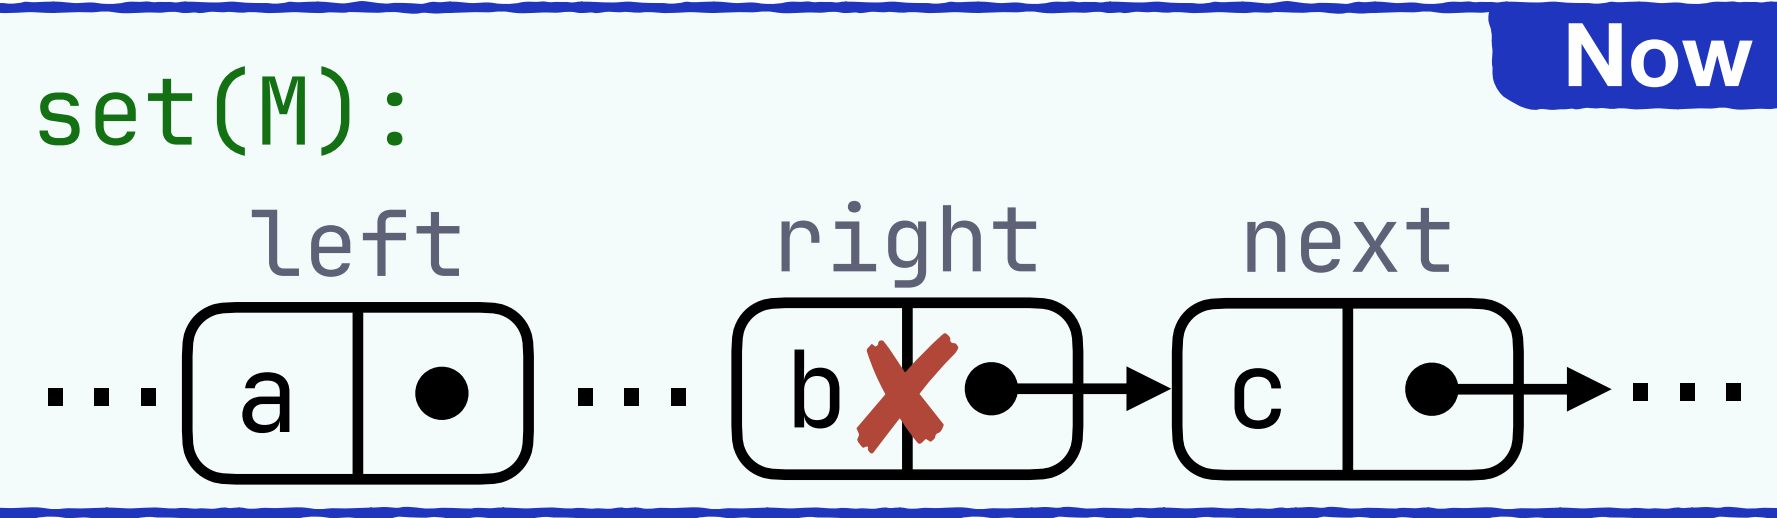
*



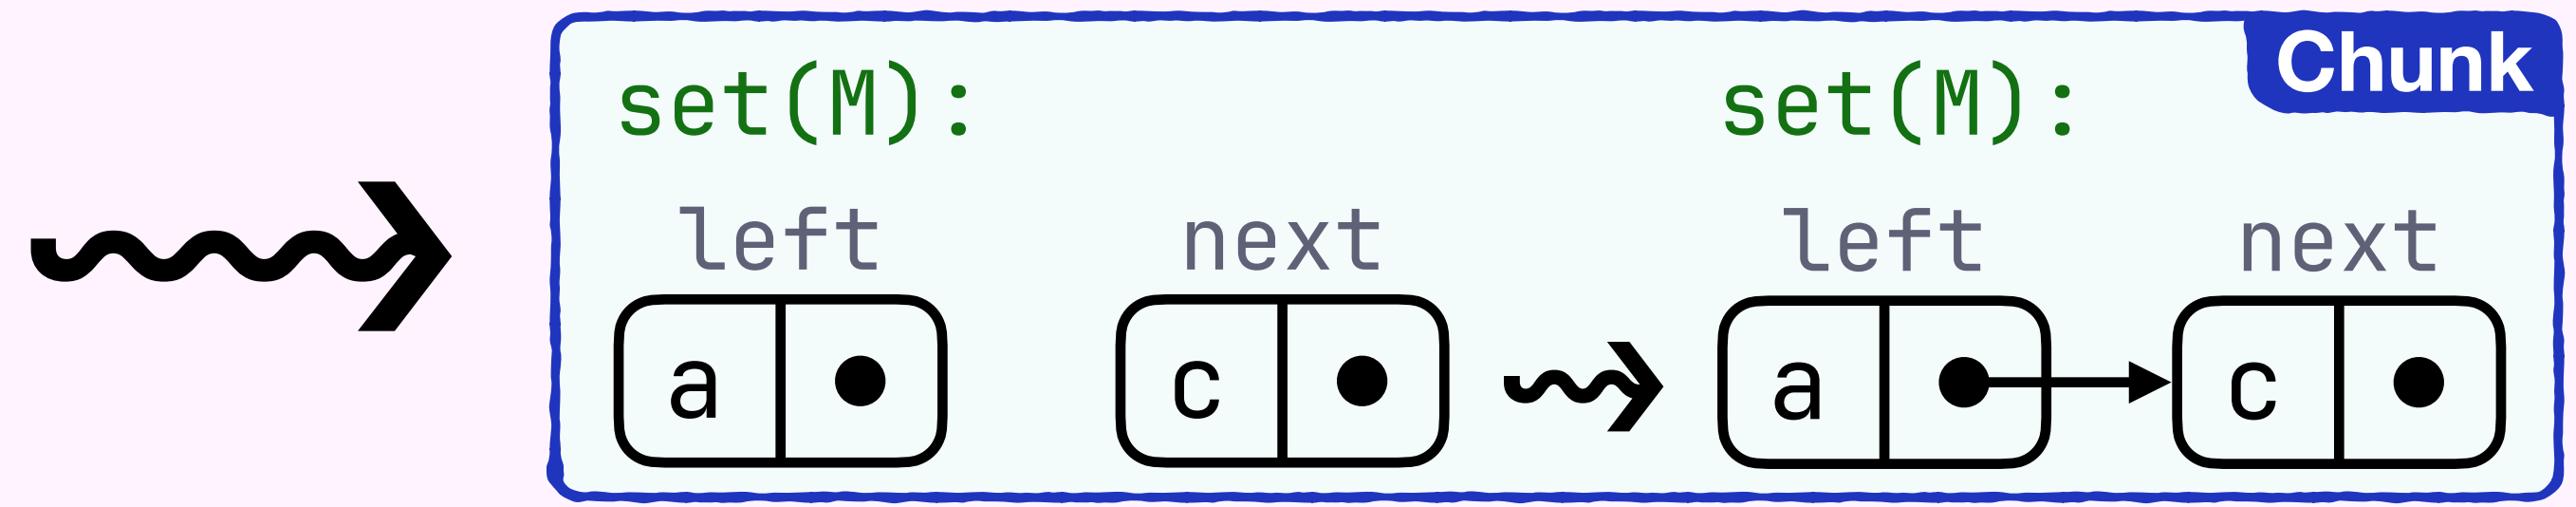
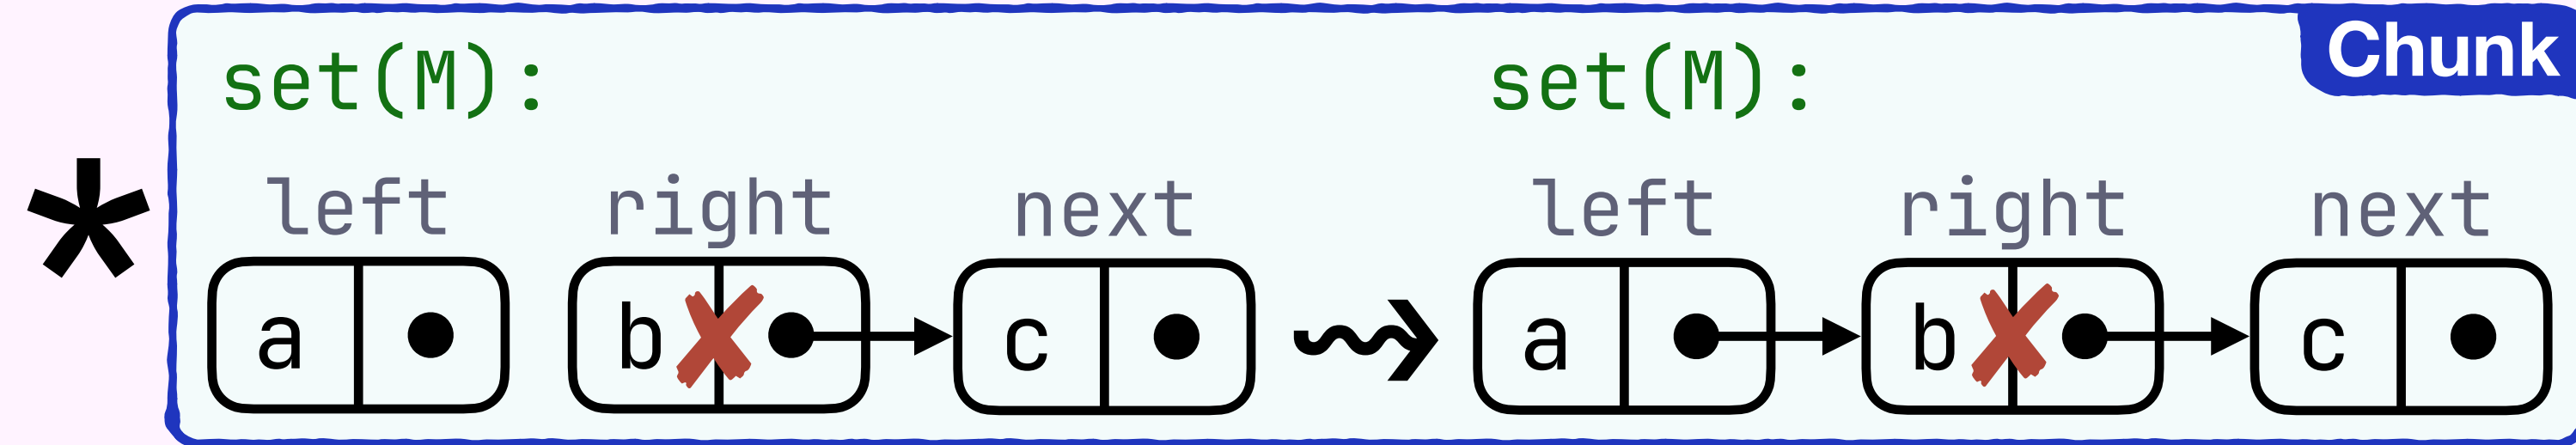
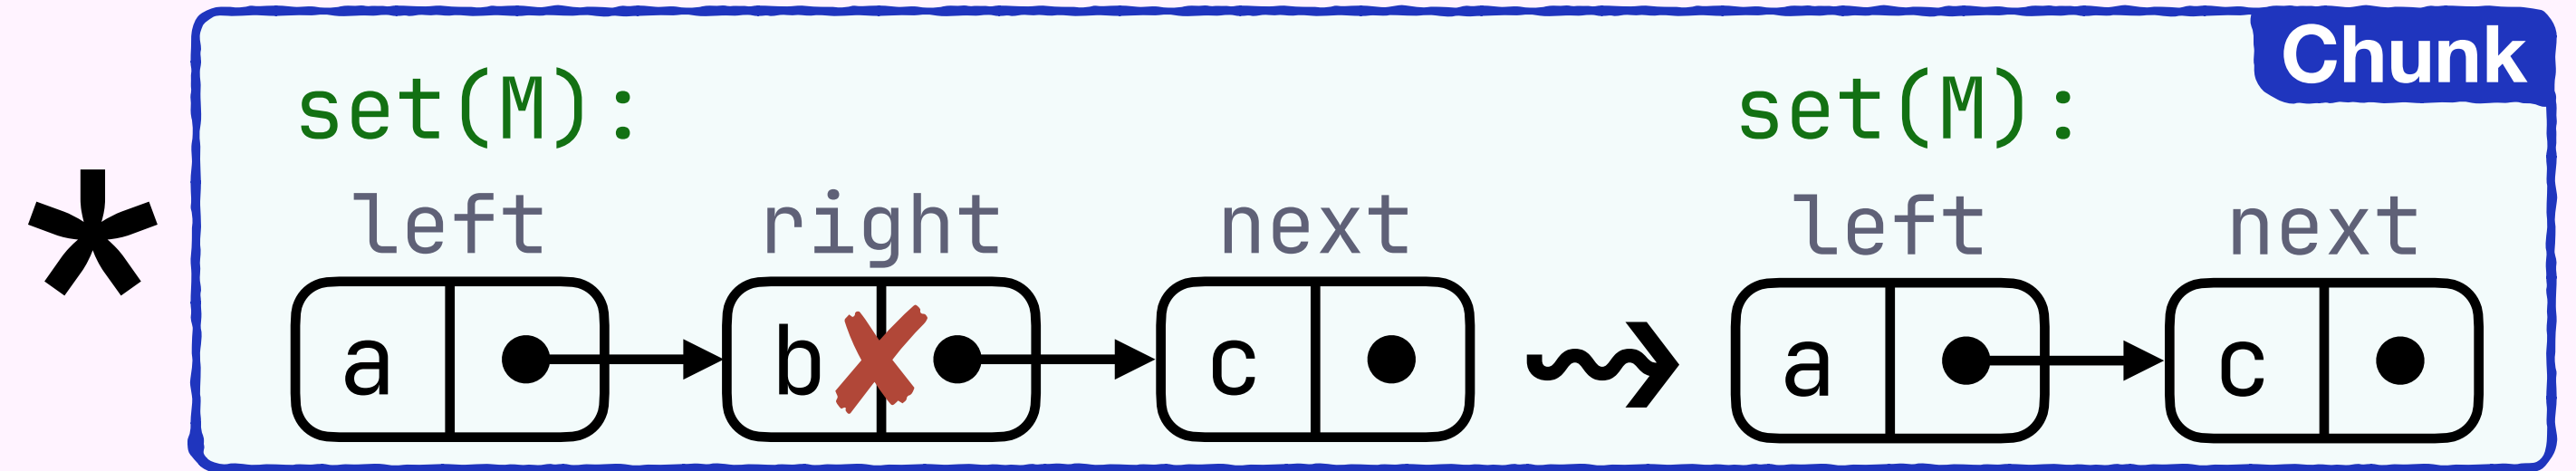
```
left → next = right;
```

Example: removing marked nodes

```
while (right → mark):  
    next = right → next
```



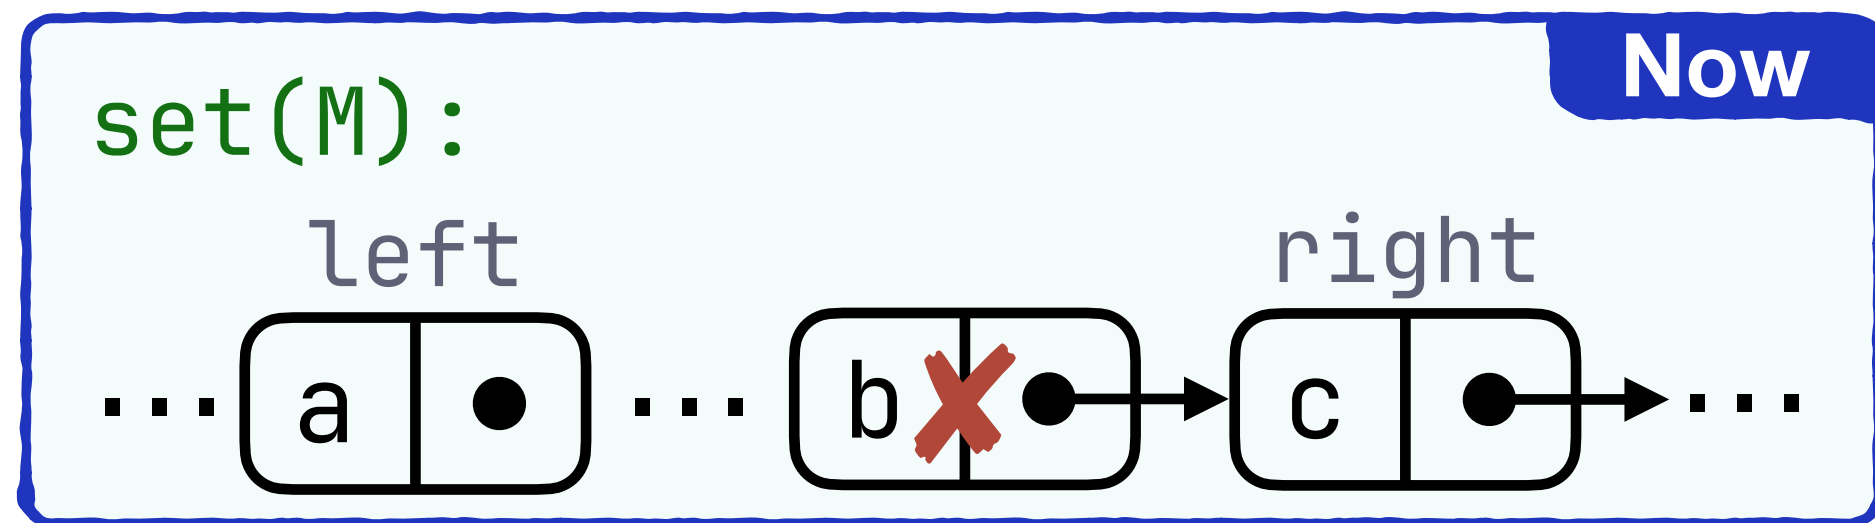
```
    right = next;  
    left → next = right;
```



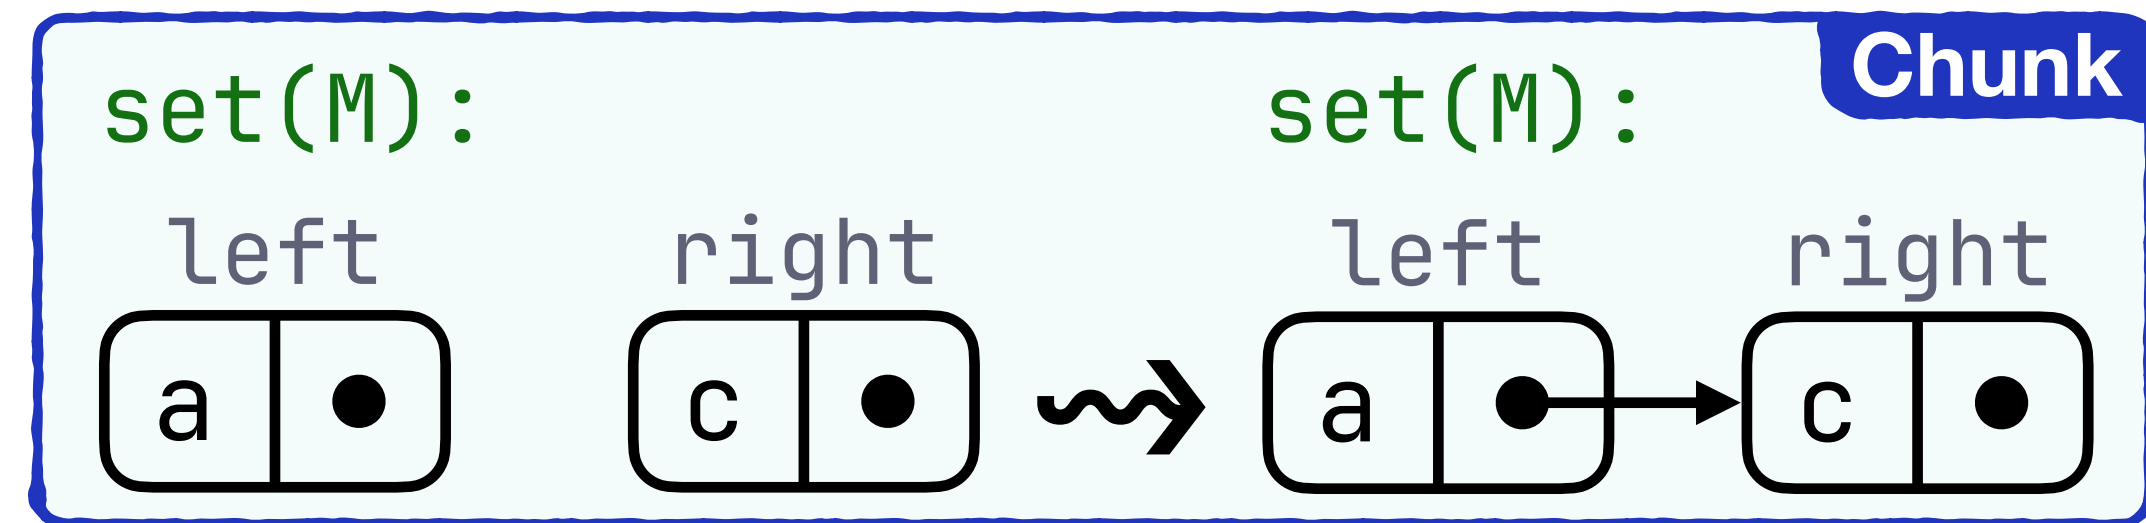
(MERGE) + (DISCHARGE)

Example: removing marked nodes

```
while (right → mark):  
    next = right → next;  
    right = next;
```



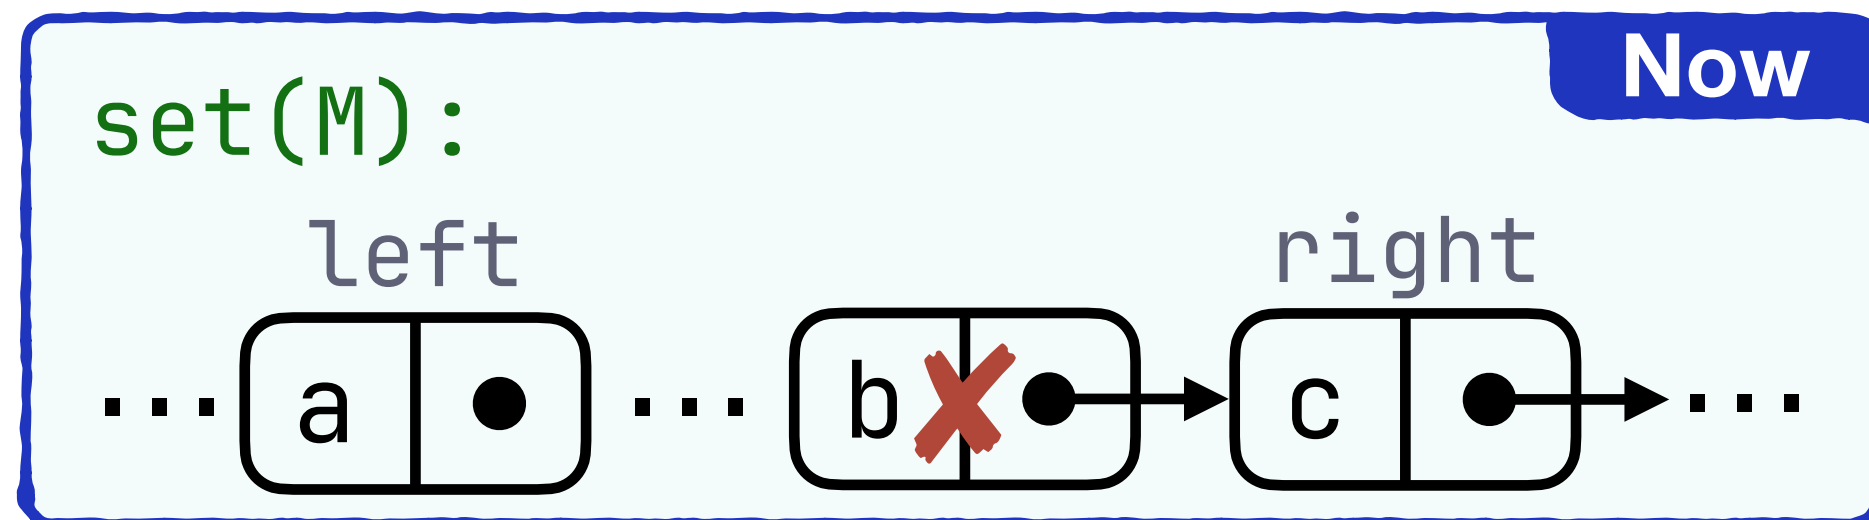
*



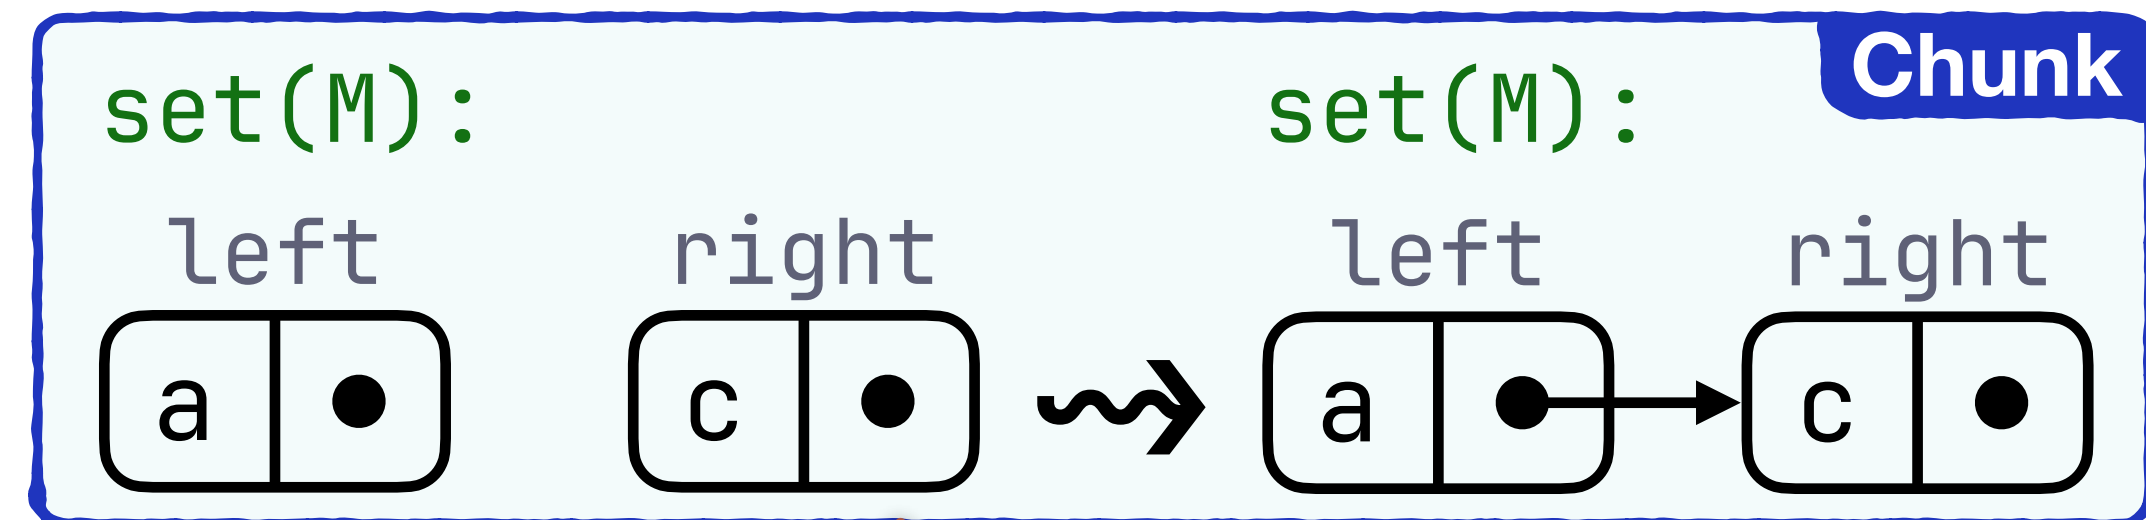
```
left → next = right;
```

Example: removing marked nodes

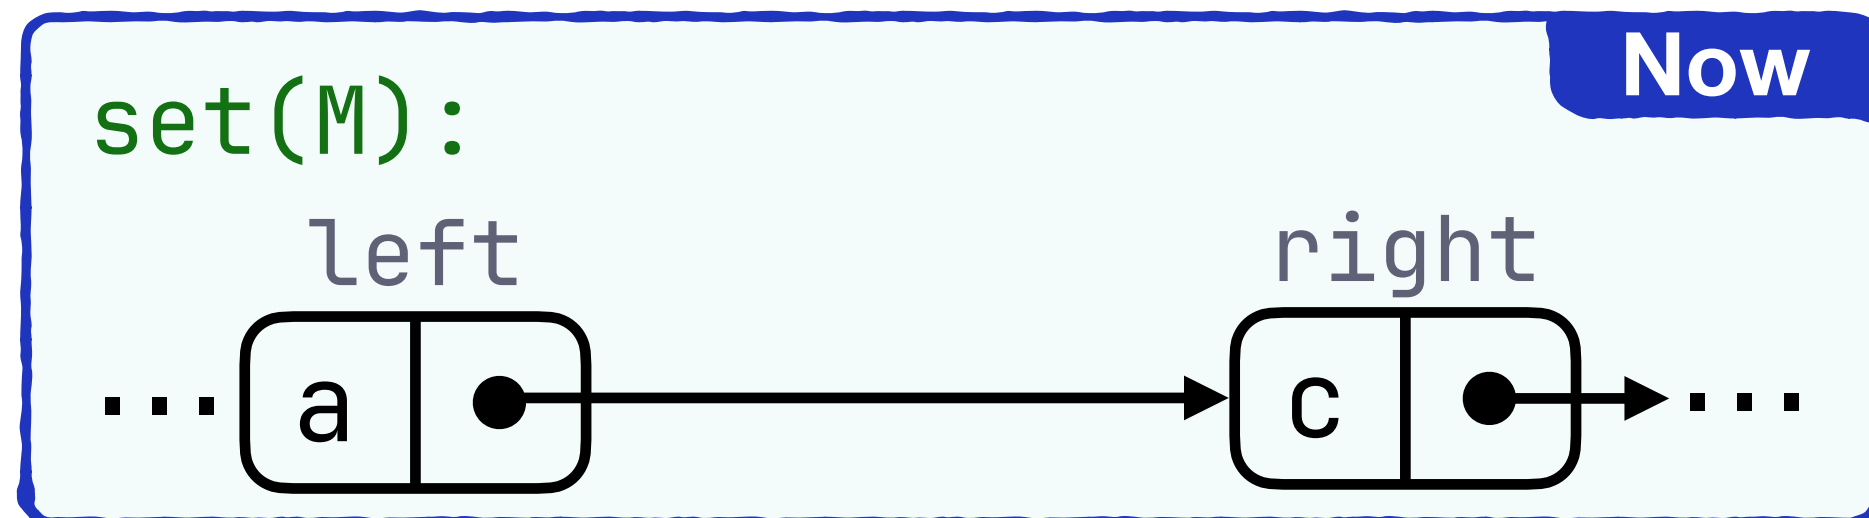
```
while (right → mark):  
    next = right → next;  
    right = next;
```



*



```
left → next = right;
```



PROVES

Technical Contribution

- Abstract Owicki-Gries separation logic, with
 - update chunks: $\langle P \rangle \text{ com } \langle Q \rangle$

Technical Contribution

- Abstract Owicki-Gries separation logic, with

- update chunks: $\langle P \rangle \text{ com } \langle Q \rangle$

- proof rules (excerpt):

$$\frac{}{R * \langle P * R \rangle \text{ com } \langle Q \rangle \implies \langle P \rangle \text{ com } \langle Q \rangle} \text{ (DISCHARGE)}$$

$$\frac{\text{com}_1 ; \text{com}_2 \approx \text{com}}{\langle P \rangle \text{ com } \langle Q \rangle * \langle Q \rangle \text{ com } \langle R \rangle \implies \langle P \rangle \text{ com } \langle R \rangle} \text{ (MERGE)}$$

Technical Contribution

- Abstract Owicki-Gries separation logic, with

- update chunks: $\langle P \rangle \text{ com } \langle Q \rangle$

- proof rules (excerpt):

$$\frac{}{R * \langle P * R \rangle \text{ com } \langle Q \rangle \implies \langle P \rangle \text{ com } \langle Q \rangle} \text{ (DISCHARGE)}$$

$$\frac{\text{com}_1 ; \text{com}_2 \approx \text{com}}{\langle P \rangle \text{ com } \langle Q \rangle * \langle Q \rangle \text{ com } \langle R \rangle \implies \langle P \rangle \text{ com } \langle R \rangle} \text{ (MERGE)}$$

- Soundness by elimination

Technical Contribution

- Abstract Owicki-Gries separation logic, with
 - update chunks: $\langle P \rangle \text{ com } \langle Q \rangle$
 - proof rules (excerpt):

$$\frac{R * \langle P * R \rangle \text{ com } \langle Q \rangle \implies \text{com}_1 ; \text{com}_2 \approx \text{com}_2 ; \text{com}_1}{\langle P \rangle \text{ com } \langle Q \rangle * \langle Q \rangle \text{ com } \langle R \rangle = \langle Q \rangle \text{ com } \langle P * R \rangle}$$

- Soundness by elimination

A Concurrent Program Logic with a Future and History

ROLAND MEYER, TU Braunschweig, Germany
THOMAS WIES, New York University, USA
SEBASTIAN WOLFF, New York University, USA

Verifying fine-grained optimistic concurrent programs remains an open problem. Modern program logics provide abstraction mechanisms and compositional reasoning principles to deal with the inherent complexity. However, their use is mostly confined to pencil-and-paper or mechanized proofs. We devise a new separation logic geared towards the lacking automation. While local reasoning is known to be crucial for automation, we are the first to show how to retain this locality for (i) reasoning about inductive properties without the need for ghost code, and (ii) reasoning about computation histories in hindsight. We implemented our new logic in a tool and used it to automatically verify challenging concurrent search structures that require inductive properties and hindsight reasoning, such as the Harris set.

CCS Concepts: • **Theory of computation** → **Separation logic; Hoare logic; Automated reasoning; Program verification; Programming logic.**

Additional Key Words and Phrases: Linearizability, Non-blocking Data Structures, Harris Set

ACM Reference Format:

Roland Meyer, Thomas Wies, and Sebastian Wolff. 2022. A Concurrent Program Logic with a Future and History. *Proc. ACM Program. Lang.* 6, OOPSLA2, Article 174 (October 2022), 30 pages. <https://doi.org/10.1145/3563337>

1 INTRODUCTION

Concurrency comes at a cost, at least, in terms of increased effort when verifying program correctness. There has been a proliferation of concurrent program logics that provide an arsenal of reasoning techniques to address this challenge [Bell et al. 2010; Delbianco et al. 2017; Elmas et al. 2010; Fu et al. 2010; Gotsman et al. 2013; Gu et al. 2018; Jung et al. 2018; Liang and Feng 2013; Manna and Pnueli 1995; Parkinson et al. 2007; Sergey et al. 2015; Vafeiadis and Parkinson 2007]. In addition, a number of general approaches have been developed to help structure the high-level proof argument [Feldman et al. 2018, 2020; Kragl et al. 2020; O'Hearn et al. 2010; Shasha and Goodman 1988]. However, the use of these techniques has been mostly confined to manual proofs done on paper, or mechanized proofs constructed in interactive proof assistants. We distill from these works a concurrent separation logic suitable for automating the construction of local correctness proofs for highly concurrent data structures. We focus on concurrent search structures (sets and maps indexed by keys), but the developed techniques apply more broadly. Our guiding principle is to perform all inductive reasoning, both in time and space, in lock-step with the program execution. The reasoning about inductive properties of graph structures and computation histories is relegated to the meta-theory of the logic by choosing appropriate semantic models.

Running Example. We motivate our work using the Harris non-blocking set data structure [Harris 2001], which we will also use as a running example throughout the paper.

Authors' addresses: Roland Meyer, TU Braunschweig, Germany, roland.meyer@tu-bs.de; Thomas Wies, New York University, USA, wies@cs.nyu.edu; Sebastian Wolff, New York University, USA, sebastian.wolff@cs.nyu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).
2475-1421/2022/10-ART174
<https://doi.org/10.1145/3563337>

pLankton



- Automatically decomposes complex updates
 - ➔ eager heuristic
- First linearizability proof for Harris' set
 - ➔ automated

Benchmark	Verification Time
Fine-Grained set	45s ✓
Lazy set	2m 13s ✓
FEMRS tree (no maintenance)	3m 50s ✓
Vechev&Yahav 2CAS set	1m 15s ✓
Vechev&Yahav CAS set	2m 20s ✓
ORVYY set	1m 36s ✓
Michael set	6m 53s ✓
Michael set (wait-free search)	6m 53s ✓
Harris set	57m 20s ✓
Harris set (wait-free search)	43m 00s ✓
LO-tree (generalized maintenance)	16m 43s ✓

Summary

Automated verification of concurrent search structures, using



the flow framework



hindsight reasoning



update chunks

Benchmark	Verification Time
Fine-Grained set	45s ✓
Lazy set	2m 13s ✓
FEMRS tree (no maintenance)	3m 50s ✓
Vechev&Yahav 2CAS set	1m 15s ✓
Vechev&Yahav CAS set	2m 20s ✓
ORVYY set	1m 36s ✓
Michael set	6m 53s ✓
Michael set (wait-free search)	6m 53s ✓
Harris set	57m 20s ✓
Harris set (wait-free search)	43m 00s ✓
LO-tree (generalized maintenance)	16m 43s ✓

Thanks

Summary

Automated verification of concurrent search structures,
using



the flow framework



hindsight reasoning



update chunks

Benchmark	Verification Time
Fine-Grained set	45s ✓
Lazy set	2m 13s ✓
FEMRS tree (no maintenance)	3m 50s ✓
Vechev&Yahav 2CAS set	1m 15s ✓
Vechev&Yahav CAS set	2m 20s ✓
ORVYY set	1m 36s ✓
Michael set	6m 53s ✓
Michael set (wait-free search)	6m 53s ✓
Harris set	57m 20s ✓
Harris set (wait-free search)	43m 00s ✓
LO-tree (generalized maintenance)	16m 43s ✓
