

Relational Views Framework for Proving Linearizability (of concurrent libraries)

Artem Khyzha Alexey Gotsman Matthew Parkinson

Concurrent libraries

- Encapsulate efficient concurrent algorithms
 - Java: `java.util.concurrent`
 - C++: Intel Threading Building Blocks
 - C#: `System.Collections.Concurrent`
- Implement stacks, queues, skip lists, hash tables, etc
- It is not easy to understand why they are correct

Concurrent library specification

- A standard correctness criterion — linearizability

Implementation

concrete
library

\sqsubseteq

Specification

abstract
library

Concurrent library specification

Non-blocking stack

```
struct Node {
  Node *next; int val;
} *Top;

void push(int v) {
  Node *t, *x;
  x = new Node;
  x->val = v;
  do {
    t = Top;
    x->next = t;
  } while(!CAS(&Top,t,x)); }
```



Atomic stack ADT

```
Sequence S;

void push(int v) {
  atomic { S = v :: S; }
}
```

Contextual refinement

- If $L \sqsubseteq L'$, then $C(L')$ can recreate executions of $C(L)$.
- It is sound to replace a library with its specification in reasoning about its client.
- Proofs about $C(L')$ work for $C(L)$.

Linearizability

H

t1: $\underbrace{t1, \text{ call push}(42) \quad t1, \text{ ret push}(42)}$

$\underbrace{t1, \text{ call isEmpty}() \quad t1, \text{ ret isEmpty}(\text{yes})}$

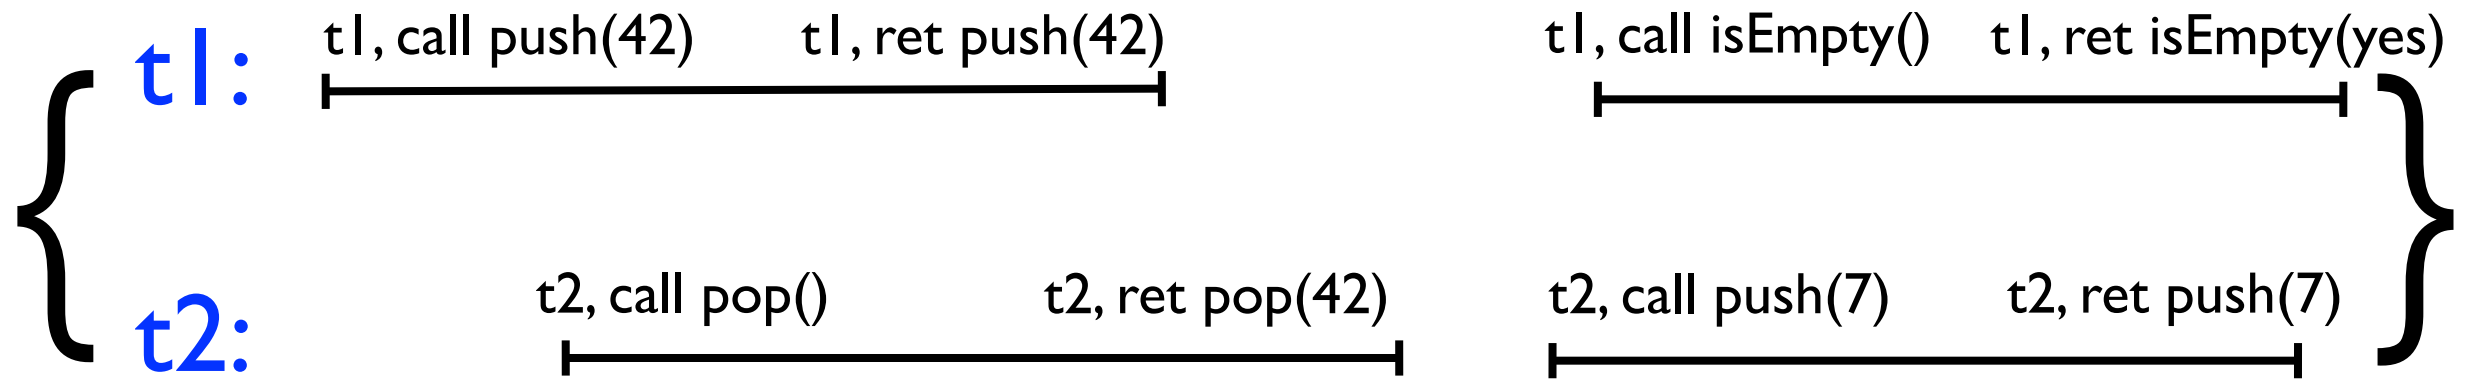
t2: $\underbrace{t2, \text{ call pop}() \quad t2, \text{ ret pop}(42)}$

$\underbrace{t2, \text{ call push}(7) \quad t2, \text{ ret push}(7)}$

- Client observes events **t, call m(a)**
and **t, ret m(r)**
- A history — a trace of events

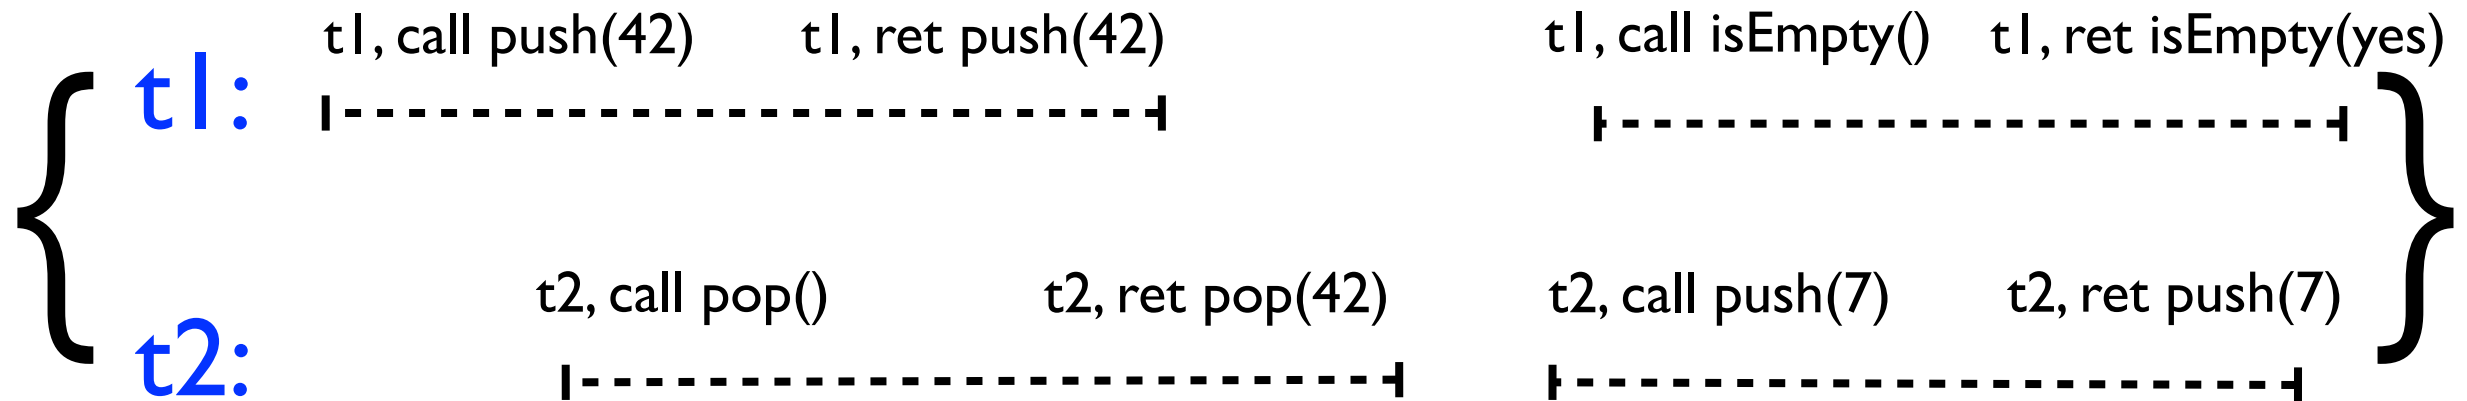
Linearizability

H



\subseteq

H'



Relational Views Framework

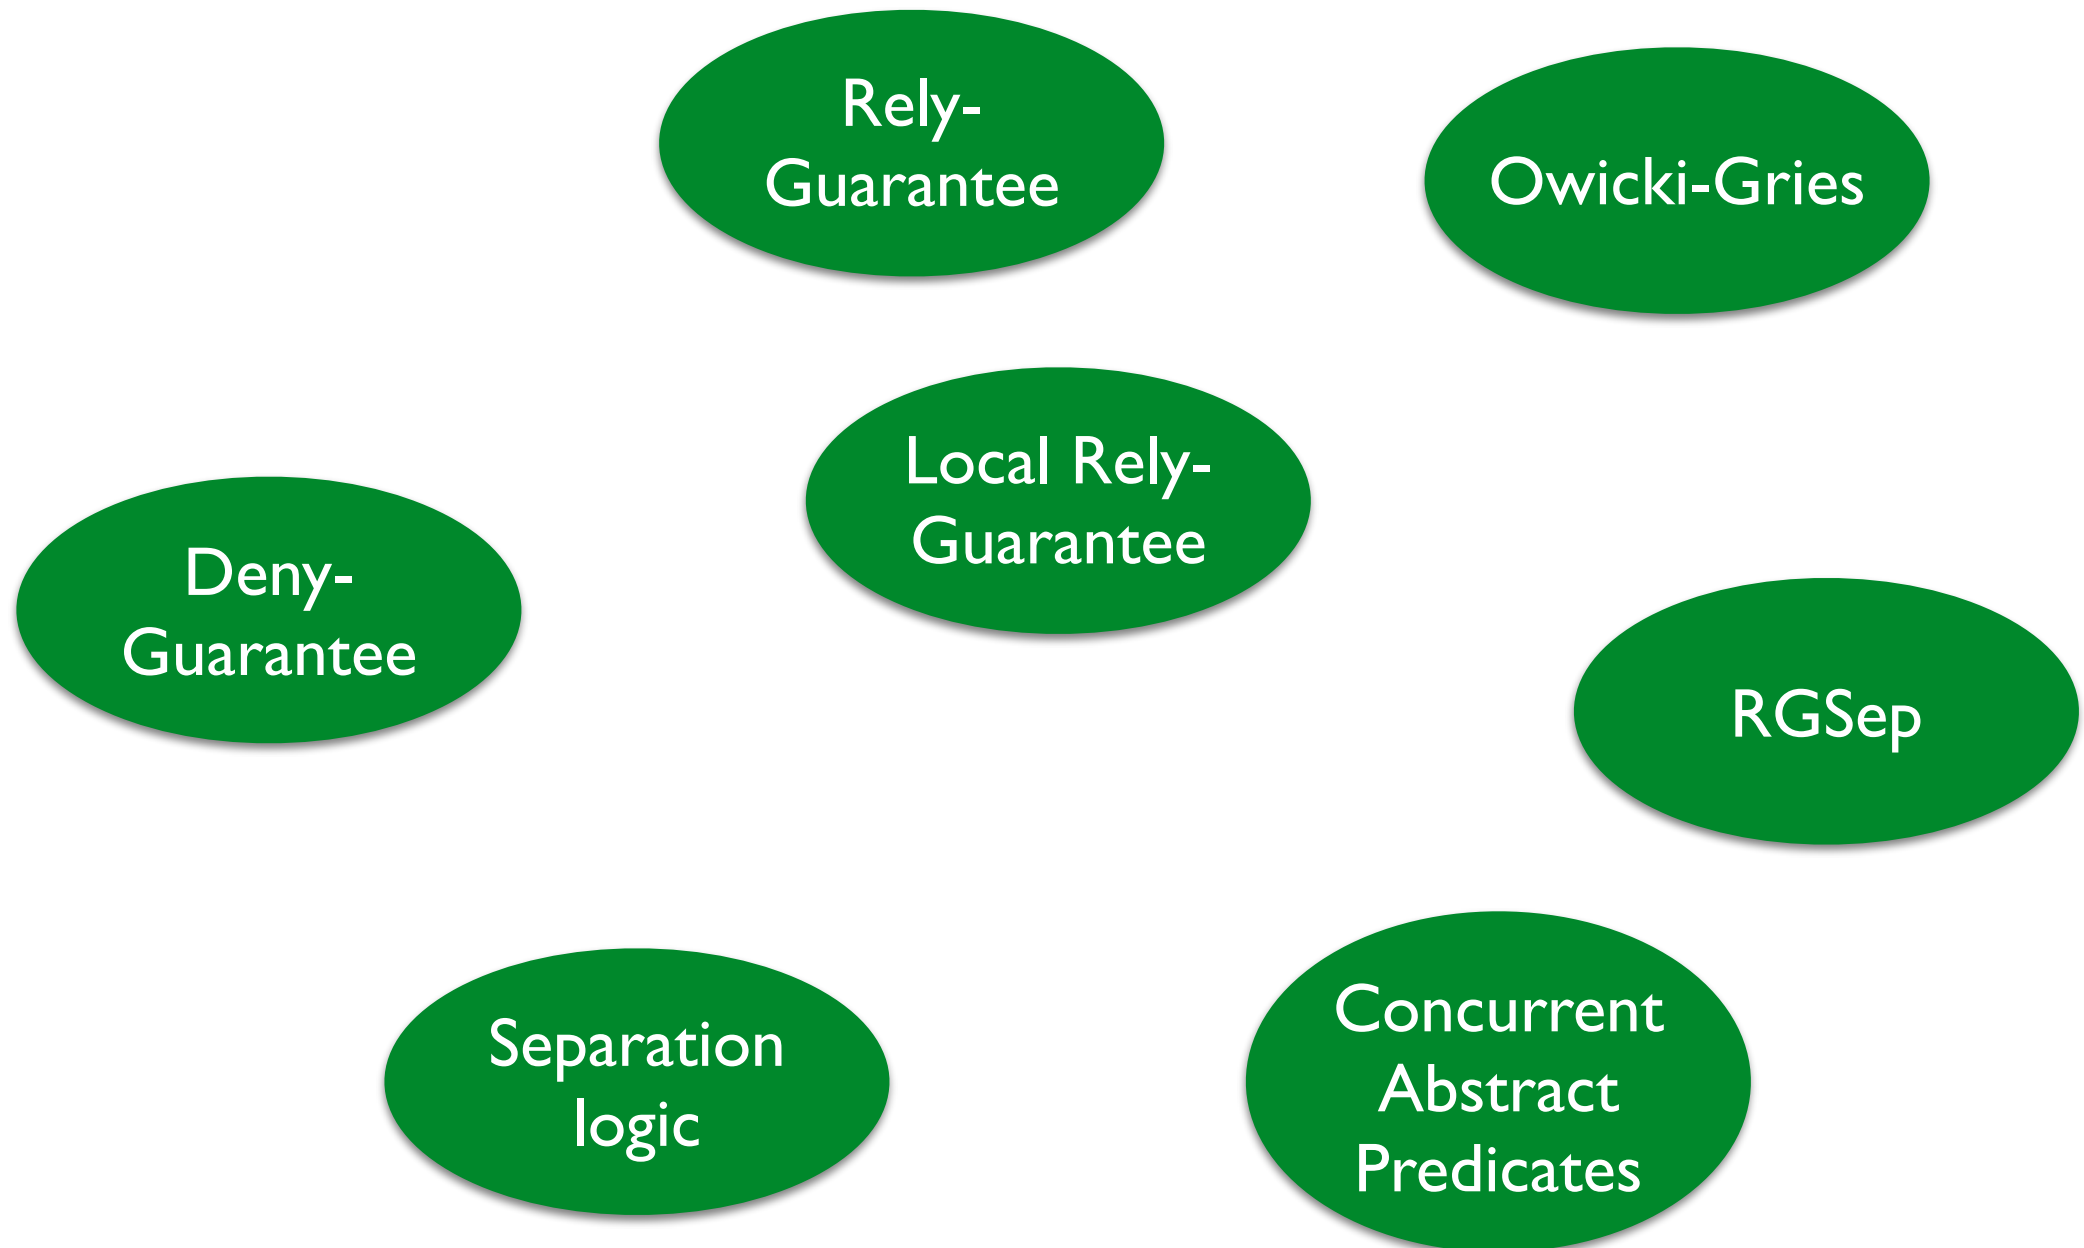
- Recipe:
 - linearization points-based proof method
 - Views Framework for thread-modular reasoning about concurrency

A simple and generic logic for proving linearizability

Views Framework

- The Views Framework is a generalisation of Hoare-style program logics for concurrency.
- Framework can be instantiated into existing logics by adjusting parameters.

Views Framework




Views Framework


- The Views Framework is a generalisation of Hoare-style program logics for concurrency.
- Framework can be instantiated into existing logics by adjusting parameters.

-
- We extend the original framework to support proving properties of pairs of programs
-

Linearization points



- implemented methods take effect instantly



```
struct Node {  
    Node *next; int val;  
} *Top;
```

```
void push(int v) {  
    Node *t, *x;  
    x = new Node;  
    x->val = v;  
    do {  
        t = Top;  
        x->next = t;  
    } while(!CAS(&Top,t,x)); }
```


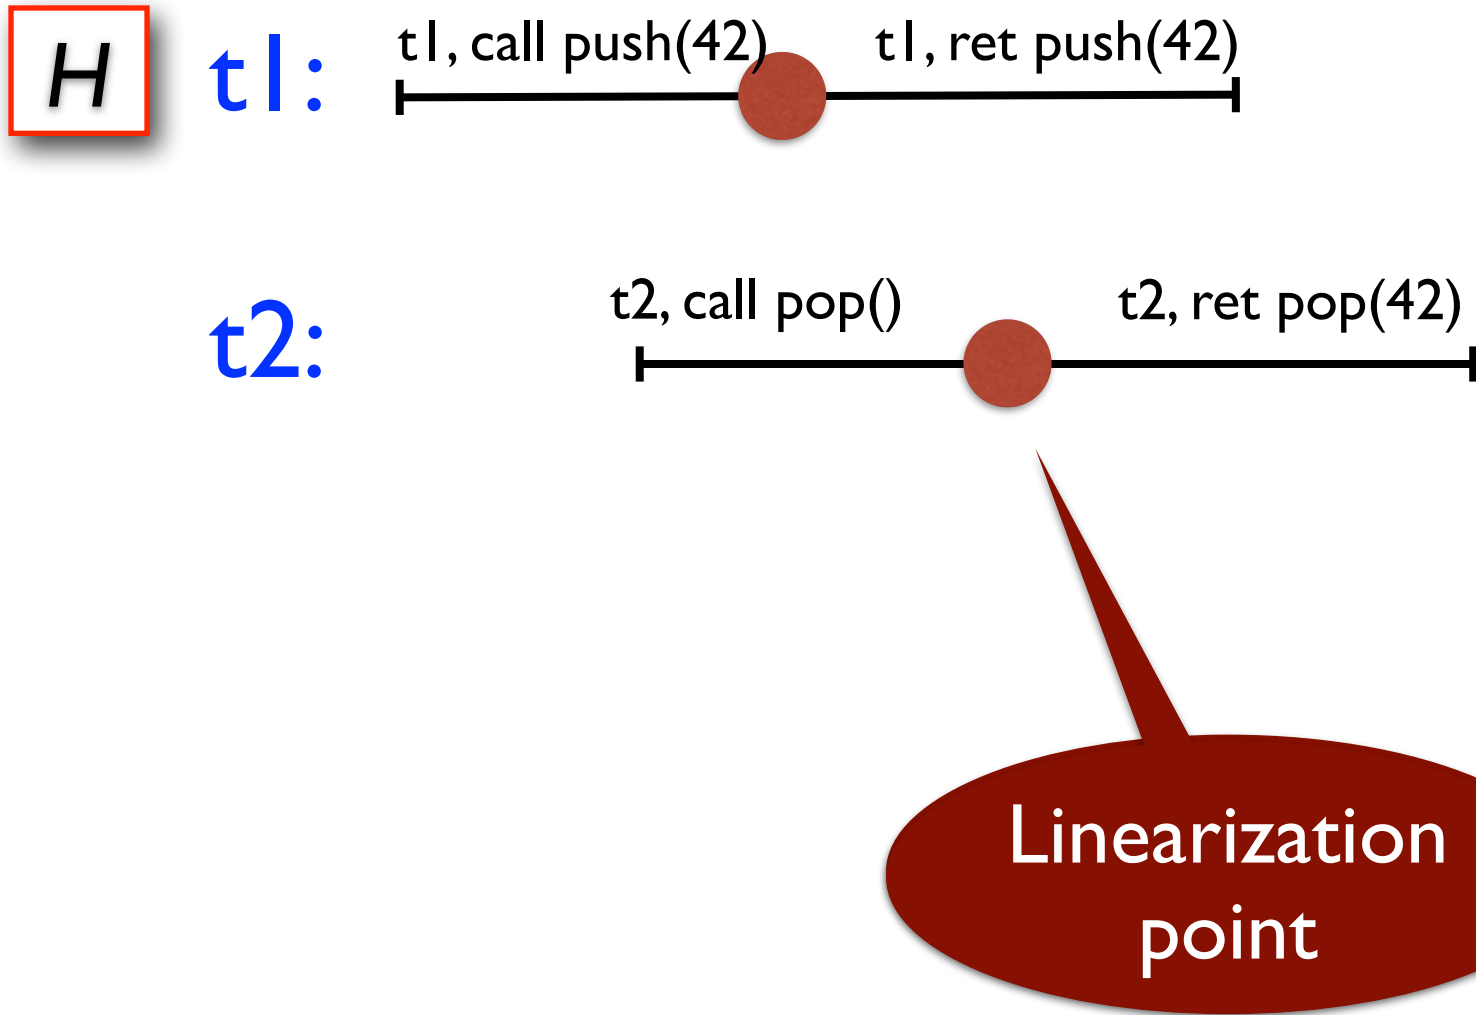
Sequence S;

```
void push(int v) {  
    atomic { S = v · S; }  
}
```

Linearizability

```
struct Node {  
    Node *next; int val;  
} *Top;
```



```
void push(int v) {  
    Node *t, *x;  
    x = new Node;  
    x->val = v;  
    do {  
        t = Top;  
        x->next = t;  
    } while(!CAS(&Top,t,x)); }
```



Sequence S;

```
void push(int v) {  
    atomic { S = v · S; }  
}
```



Proving linearizability



Facts to prove:

- there is an LP in m
- it is passed once
- return values of m and M are the same

Proving linearizability



Facts to prove:

- there is an LP in m
- it is passed once
- return values of m and M are the same

Two forms of auxiliary state:

- a state of the spec
- one-time permission to pass a LP — a **token**

Specifying methods

$$\vdash_t \{ I([\text{todo}(M)]_t) \} m \{ I([\text{done}(M)]_t) \}$$

- m starts in $(s, \sigma, \text{tokens})$ sat. $I([\text{todo}(M)]_t)$,
- uses $[\text{todo}(M)]_t$ as a permission to pass an LP
- finishes in $(s', \sigma', \text{tokens}')$ sat. $I([\text{done}(M)]_t)$

Relational views

- A set of views — assertions in a logic

$I([\text{todo}(M)]_t), I([\text{done}(M)]_t)$

- Reification — interpretation of views

$(s, \sigma, \text{tokens}) \text{ sat. } I([\text{todo}(M)]_t)$

- A composition operation *

- encodes thread-modular reasoning method

Relational views

(inspired by separation logic)

- Views are sets of partial states w\ tokens
- Views = sets of $(s, \sigma, \text{tokens})$
- States = heap configurations $(\text{Loc} \longrightarrow \text{Val})$
- Tokens = sets of $[\text{todo inc}(3)]_t$ and $[\text{done inc}(7)]_t$

state of the impl.

$c \mapsto 42$

abstract state

$C \Rightarrow 42$

tokens

$[\text{todo inc}(7)]_t$

Relational views

(inspired by separation logic)

- Views are sets of partial states w\ tokens
 - represent an ownership of a part of a heap
- Composition forces views to describe disjoint parts of a heap

$$c \mapsto 42 * C \Leftrightarrow 42 * [\text{todo}(\text{inc}(3))]t$$

$$p * p' = \{ (s \uplus s', \sigma \uplus \sigma', \tau \cup \tau') \mid \\ (s, \sigma, \tau) \in p \text{ and } (s', \sigma', \tau') \in p' \}$$

Relational views

(inspired by separation logic)

- Reification function:
 - complements partial heap configurations arbitrarily

$$\lfloor p \rfloor = \{ (s \uplus s', \sigma \uplus \sigma', \tau) \mid (s, \sigma, \tau) \in p \}$$

Relational views: rely/guarantee

- Views are triples of (**assertion**, **rely**, **guarantee**)
 - **assertion** is a set of (state, state, tokens)
 - **guarantee** — thread's own transitions
 - **rely** — transitions by other threads
- $\vdash_t \{(P, R, G)\} \subset \{(Q, R, G)\}$

Relational views: rely/guarantee

- Views are triples of (**assertion**, **rely**, **guarantee**)
 - **assertion** is a set of (state, state, tokens)
- Reification erases all extra information:
 - $\lfloor (\text{assertion}, \text{rely}, \text{guarantee}) \rfloor = \text{assertion}$

Relational views: rely/guarantee

- Views are triples of (**assertion**, **rely**, **guarantee**)
 - **assertion** is a set of (state, state, tokens)

$$\bullet (P, R, G) * (P', R', G') = (P \cup P', R \cap R', G \cup G'),$$

when $G' \subseteq R, G \subseteq R'$

- Encodes a parallel composition rule

Proof rules

$$\frac{\Vdash_t \{p\} \alpha \{q\}: _}{\vdash_t \{P\} \alpha \{Q\}}$$

$$\frac{\vdash_t \{P\} C_1 \{P'\} \quad \vdash_t \{P'\} C_2 \{Q\}}{\vdash_t \{P\} C_1 ; C_2 \{Q\}}$$

$$\frac{\vdash_t \{P\} C \{Q\}}{\vdash_t \{P * R\} C \{Q * R\}}$$

$$\frac{\vdash_t \{P\} C_1 \{Q\} \quad \vdash_t \{P\} C_2 \{Q\}}{\vdash_t \{P\} C_1 + C_2 \{Q\}}$$

$$\frac{\vdash_t \{P\} C \{Q\}}{\vdash_t \{\exists X. P\} C \{\exists X. Q\}}$$

$$\frac{P' \Rightarrow P \quad \vdash_t \{P\} C \{Q\} \quad Q \Rightarrow Q'}{\vdash_t \{P'\} C \{Q'\}}$$

$$\frac{\vdash_t \{P_1\} C \{Q_1\} \quad \vdash_t \{P_2\} C \{Q_2\}}{\vdash_t \{P_1 \vee P_2\} C \{Q_1 \vee Q_2\}}$$

$$\frac{\vdash_t \{P\} C \{P\}}{\vdash_t \{P\} C^* \{P\}}$$

Proof rules

$$\frac{\Vdash_t \{p\} \alpha \{q\}: _}{\vdash_t \{P\} \alpha \{Q\}}$$

$$\frac{\vdash_t \{P\} C_1 \{P'\} \quad \vdash_t \{P'\} C_2 \{Q\}}{\vdash_t \{P\} C_1 ; C_2 \{Q\}}$$

**Rule for atomic
commands
AND
linearization points**

$$\frac{\vdash_t \{P\} C \{Q\}}{\vdash_t \{P * R\} C \{Q * R\}}$$

$$\frac{\vdash_t \{P\} C \{Q\}}{\vdash_t \{P\} C \{Q\}}$$

$$\frac{\vdash_t \{P\} C \{Q\}}{\vdash_t \{\exists X. P\} C \{\exists X. Q\}}$$

$$\frac{P' \Rightarrow P \quad \vdash_t \{P\} C \{Q\} \quad Q \Rightarrow Q'}{\vdash_t \{P'\} C \{Q'\}}$$

$$\frac{\vdash_t \{P_1\} C \{Q_1\} \quad \vdash_t \{P_2\} C \{Q_2\}}{\vdash_t \{P_1 \vee P_2\} C \{Q_1 \vee Q_2\}}$$

$$\frac{\vdash_t \{P\} C \{P\}}{\vdash_t \{P\} C^* \{P\}}$$

Rule for atomic commands

$\Vdash_{\tau} \{p\} \alpha \{q\}:$

Rely-
Guarantee

Local Rely-
Guarantee

Deny-
Guarantee

requirements
to reasoning
techniques

RGSep

Separation
logic

Concurrent
Abstract
Predicates

Rule for atomic commands

$\Vdash_t \{p\} \alpha \{q\}$:

- follows a concrete semantics of α ;



$\{ x \rightarrow 42 \} ++x \{ x \rightarrow 43 \}$



$\{ x \rightarrow 42 \} --x \{ x \rightarrow 43 \}$

Rule for atomic commands

$\Vdash_t \{p\} \alpha \{q\}$:

- follows a concrete semantics of α ;
- requires simulation by an (optional) update to the abstract state;




$\{ x \rightarrow 42 \} ++x \{ x \rightarrow 43 \}$




$\{ x \rightarrow 42 \} --x \{ x \rightarrow 43 \}$

Rule for atomic commands



Rule for atomic commands



Rule for atomic commands

for all $(s, \sigma, \tau) \in \lfloor p * r \rfloor$

for all state updates $s \xrightarrow{\alpha} s'$

there exists $\sigma \xrightarrow{A} \sigma'$ such that

in τ , $[\text{todo}(A)]$ changes to $[\text{done}(A)]$

or $A = \text{nop}$, $\sigma = \sigma'$

and $(s', \sigma', \tau') \in \lfloor q * r \rfloor$

$\vdash \{p\} \alpha \{q\}$

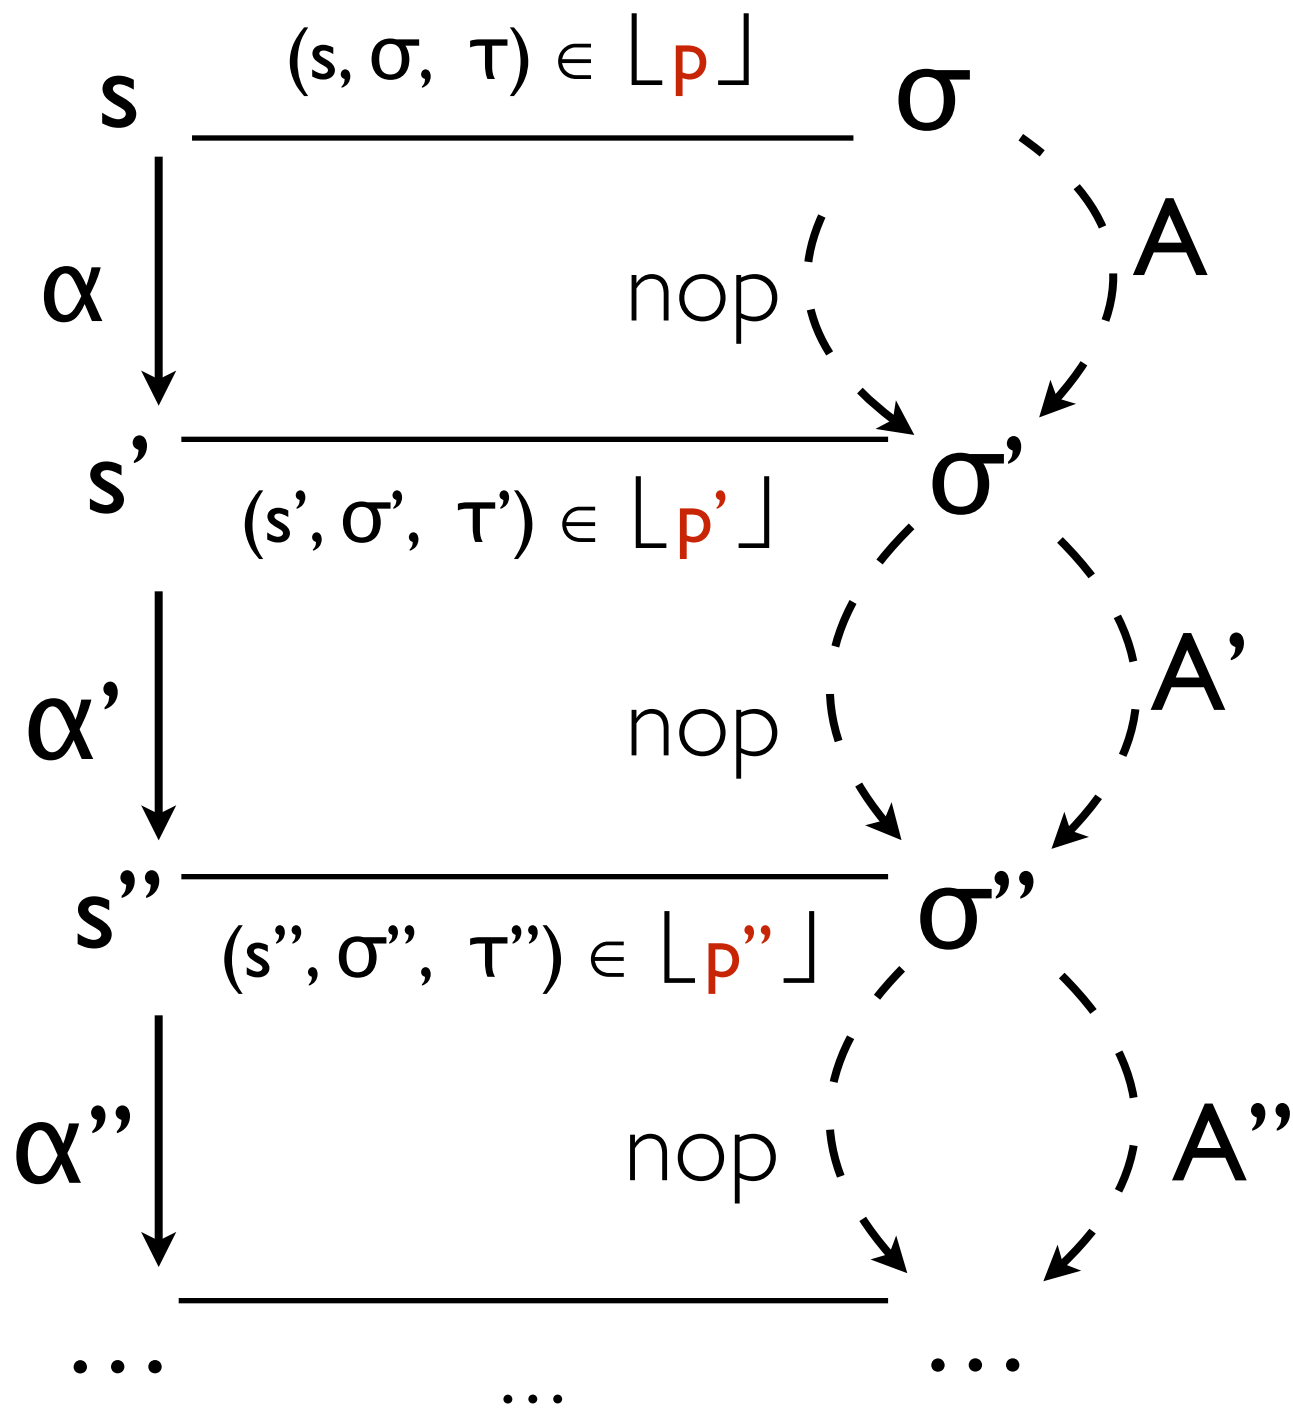
Specifying methods

$$\vdash_t \{ I([\text{todo}(M)]_t) \} m \{ I([\text{done}(M)]_t) \}$$
$$\Vdash_t \{ p \} C \{ q \} =$$

for all $C \xrightarrow{\alpha} C'$ exists p' .

$$\Vdash_t \{ p \} \alpha \{ q \}, \text{ and}$$
$$\Vdash_t \{ p' \} C' \{ q \}$$

Specifying methods



$$\Vdash_t \{ p \} C \{ q \} =$$

for all $C \xrightarrow{\alpha} C'$ exists p' .

$\Vdash_t \{ p \} \alpha \{ q \}$, and

$\Vdash_t \{ p' \} C' \{ q \}$

Soundness

- L — library/impl. L' — library/spec.
- When every method is specified like this:
 - $\vdash_t \{ I([\text{todo}(M)]_t) \} m \{ I([\text{done}(M)]_t) \}$
- $\text{histories}(L) \subseteq \text{histories}(L')$

Work in progress

- ✔ Simple logic and semantics
- ✔ Reasoning about LPs and helping
- ✘ Speculation

Instantiations:

✔ RGsep, RGsim

✘ CaReSL