

# Verified Compilation of Linearisable Data Structures

**Yannick Zakowski**

David Cachera


Delphine Demange

David Pichardie




# Introduction: a motivating example

# Verifying an on-the-fly garbage collector




**With a *sequential* GC, the main program pauses during collection**

# Verifying an on-the-fly garbage collector



*An on-the-fly GC is hosted in a different thread, and collects the memory without ever pausing the main program*

# Verifying an on-the-fly garbage collector

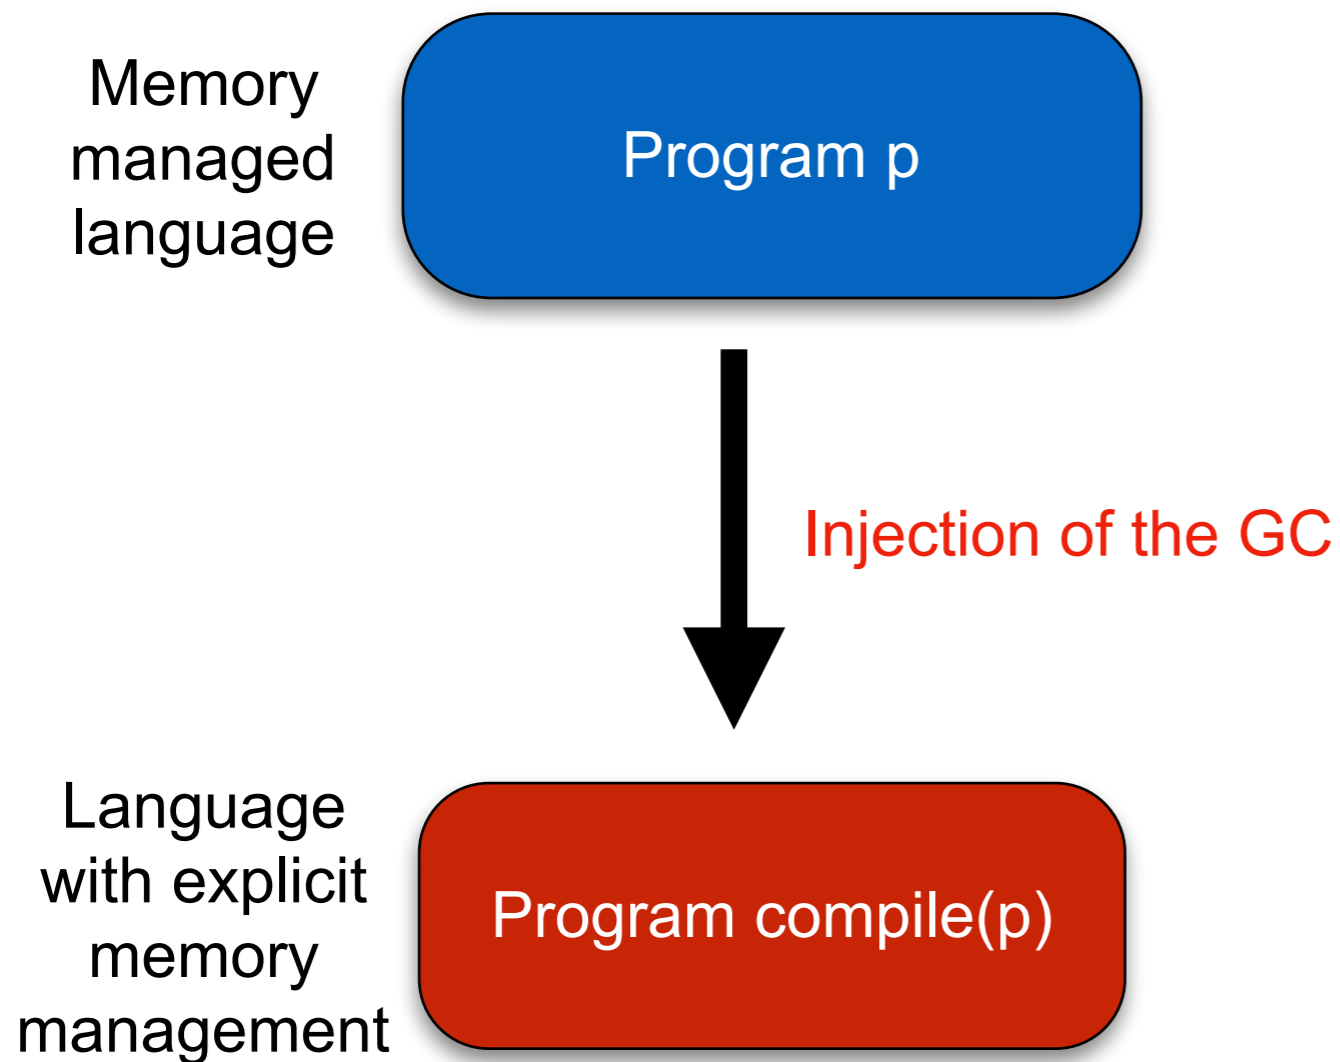


*An on-the-fly GC is hosted in a different thread, and collects the memory without ever pausing the main program*

## **Theorem (informal)**


The collector never reclaims a part of the memory that can still be accessed by the program

# Verifying an on-the-fly garbage collector in the context of verified compilation



# Verifying an on-the-fly garbage collector

## in the context of verified compilation



### Observational refinement

$$\forall P, P', \text{obs},$$
$$\text{compiler } P = \text{OK } P' \wedge$$
$$\text{low\_exec } P' \text{ obs} \Rightarrow$$
$$\text{high\_exec } P \text{ obs}$$

# A verified on-the-fly garbage collector

Scan:

**repeat**

no\_gray = **true**;

**foreach**  $x \in \text{OBJECTS}$

**if**  $x.\text{color} == \text{GRAY}$

no\_gray = **false**;

**foreach**  $f \in \text{fields}(x)$  **do**

MarkGray( $x.f$ );

$x.\text{color} = \text{BLACK}$

**until** no\_gray

Sweep:

**foreach**  $x \in \text{OBJECTS}$

**if**  $x.\text{color} == \text{WHITE}$

**then** FREE( $x$ )

Clear:

**foreach**  $x \in \text{OBJECTS}$

$x.\text{color} = \text{WHITE}$



**if**  $x.\text{color} = \text{WHITE}$  **then**  
**push(buffer[m], x)**



# A verified on-the-fly garbage collector

Scan:

**repeat**

no\_gray = **true**;

**foreach**  $x \in \text{OBJECTS}$

**if**  $x.\text{color} == \text{GRAY}$

no\_gray = **false**;

**foreach**  $f \in \text{fields}(x)$  **do**

MarkGray( $x.f$ );

$x.\text{color} = \text{BLACK}$

**until** no\_gray

Sweep:

**foreach**  $x \in \text{OBJECTS}$

**if**  $x.\text{color} == \text{WHITE}$

**then** FREE( $x$ )

Clear:

**foreach**  $x \in \text{OBJECTS}$

$x.\text{color} = \text{WHITE}$



if  $x.\text{color} = \text{WHITE}$  then  
**push(buffer[m], x)**

$nw = m.\text{next\_write}$   
 $nr = m.\text{next\_read}$   
 $d = m.\text{data}$   
 $d[nw] = x$   
 $nw = (nw+1) \bmod \text{SIZE}$   
assume ( $nr == nw$ )  
 $m.\text{next\_write} = nw$

# A verified on-the-fly garbage collector ?

?

Scan:

**repeat**

no\_gray = **true**;

**foreach** x ∈ OBJECTS

**if** x.color == GRAY

no\_gray = **false**;

**foreach** f ∈ fields(x) **do**

MarkGray(x.f);

x.color = BLACK

**until** no\_gray

Sweep:

**foreach** x ∈ OBJECTS

**if** x.color == WHITE

**then** FREE(x)

Clear:

**foreach** x ∈ OBJECTS

x.color = WHITE



if x.color = WHITE then  
**push(buffer[m], x)**

nw = m.next\_write  
nr = m.next\_read  
d = m.data  
d[nw] = x  
nw = (nw+1) mod SIZE  
assume (nr == nw)  
m.next\_write = nw

**1. Linearisability**

**2. Using our theorem: proving linearisability through Rely-Guarantee**

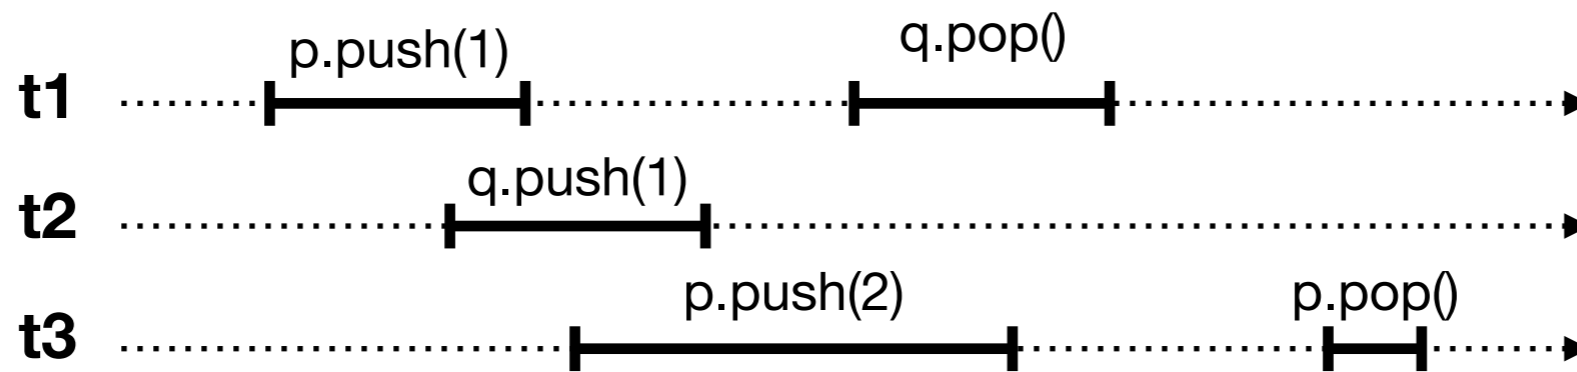
**3. Under the hood: systematic derivation of a simulation**

# Linearisability

# Linearisability

[Herlihy and Wing 90]


A notion of coherence for concurrent data structures



# Linearisability

[Herlihy and Wing 90]

A notion of coherence for concurrent data structures



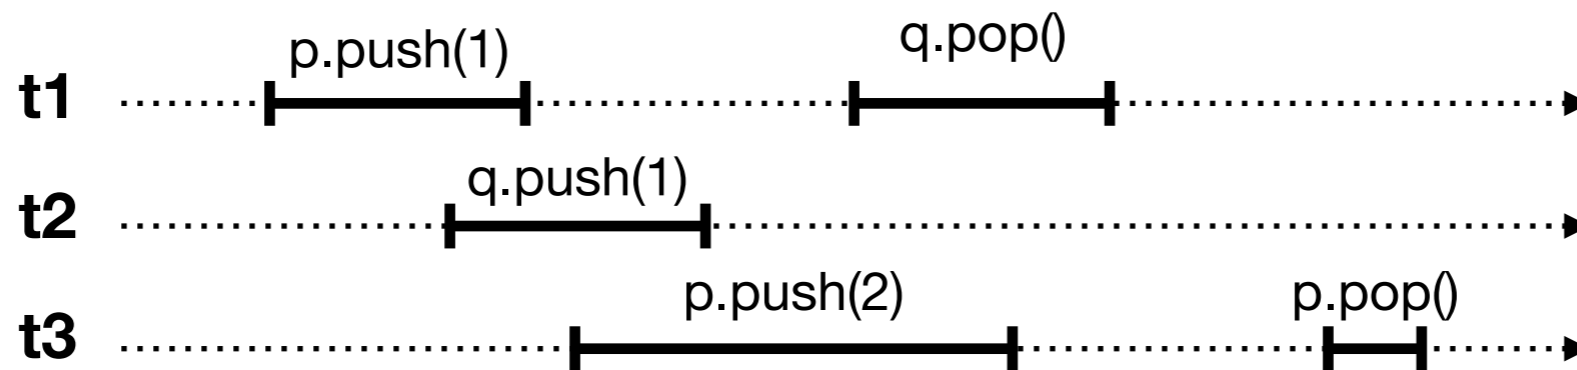
## Principle 1.

Any method should appear to happen in a one-at-a-time order

# Linearisability

[Herlihy and Wing 90]

A notion of coherence for concurrent data structures



## Principle 1.

Any method should appear to happen in a one-at-a-time order

## Principle 2. (Linearisability)

Any method should appear to take effect instantaneously  
*at some moment between its call and return*

# Linearisability

## Original formal definition

- Expressed in terms of traces of events (histories)
- For all possible history, there exists an “equivalent” well-behaved history
- Great, but does not fit our story

## Two main caveats

- The property is not explicitly usable for verified compilation purpose  
→ Change definition!
- Histories are global objects, difficult to reason about  
→ Derive it from RG proof obligations!



# Linearisability as an observational refinement

We see refinement as a compilation pass

- Source language:
  - abstract data structure
  - atomic operations over it
- Target language:
  - only concrete operations
- Compilation pass:
  - provides a concrete implementation

# Linearisability as an observational refinement

We see refinement as a compilation pass

- Source language:
  - abstract data structure
  - atomic operations over it
- Target language:  
only concrete operations
- Compilation pass:  
provides a concrete implementation

```
if x.color = WHITE then  
  push(buffer[m], x)
```



```
if x.color = WHITE then  
  nw = m.next_write  
  nr = m.next_read  
  d = m.data  
  d[nw] = x  
  nw = (nw+1) mod SIZE  
  assume (nr == nw)  
  m.next_write = nw
```

# Linearisability as an observational refinement

We see refinement as a compilation pass

- Source language:
  - abstract data structure
  - atomic operations over it
- Target language:  
only concrete operations
- Compilation pass:  
provides a concrete implementation

$$\text{Obs}(\mathcal{T}(p)) \subseteq \text{Obs}(p)$$

```
if x.color = WHITE then  
  push(buffer[m], x)
```



```
if x.color = WHITE then  
  nw = m.next_write  
  nr = m.next_read  
  d = m.data  
  d[nw] = x  
  nw = (nw+1) mod SIZE  
  assume (nr == nw)  
  m.next_write = nw
```

# Using our result: proving linearisability via Rely-Guarantee

# Rely Guarantee reasoning

[Jones81]

$R, G, I \vdash \{P\} \quad c \quad \{Q\}$

Environment  
R: Rely  
G: Guarantee

Global Correctness  
Invariant

Annotations

# Rely Guarantee reasoning

[Jones81]

$$\text{Environment } R, G, I \vdash \{P\} \quad c \quad \{Q\}$$


R: Rely  
G: Guarantee

Global Correctness  
Invariant

Annotations


# Rely Guarantee reasoning

[Jones81]



# Rely Guarantee reasoning

[Jones81]




$R$  : Rely, approximates the effect of the environment

$G$  : Guarantee, approximates the effect of the thread



# Rely Guarantee reasoning

[Jones81]



$R$  : Rely, approximates the effect of the environment

$G$  : Guarantee, approximates the effect of the thread

**A thread is proved against a contract.  
The notion of interference is checked against this contract.**

# Reasoning about linearisation using Rely-Guarantee

Introduction of an intermediate, *instrumented*, language.

- Explicit annotation of *linearisation points*
- Hybrid states, both concrete and abstract
- Linearisation points trigger the abstract semantics

# Reasoning about linearisation using Rely-Guarantee

Introduction of an intermediate, *instrumented*, language.

- Explicit annotation of *linearisation points*
- Hybrid states, both concrete and abstract
- Linearisation points trigger the abstract semantics


```
nw = m.next_write  
nr = m.next_read  
d = m.data  
d[nw] = x  
nw = (nw+1) mod SIZE  
assume (nr == nw)  
<m.next_write = nw; LIN>
```

# Reasoning about linearisation using Rely-Guarantee

Introduction of an intermediate, *instrumented*, language.

- Explicit annotation of *linearisation points*
- Hybrid states, both concrete and abstract
- Linearisation points trigger the abstract semantics

local map	$\rho_1$
shared heap	$\sigma_1$
abstract data-structure	$p_1$



```
nw = m.next_write
nr = m.next_read
d = m.data
d[nw] = x
nw = (nw+1) mod SIZE
assume (nr == nw)
<m.next_write = nw; LIN>
```

# Reasoning about linearisation using Rely-Guarantee

## Introduction of an intermediate, *instrumented*, language.

- Explicit annotation of *linearisation points*
- Hybrid states, both concrete and abstract
- Linearisation points trigger the abstract semantics

local map	$\rho_1$	$\rho_2$	$\rho_3$	$\rho_4$	$\rho_4$	$\rho_5$	$\rho_5$
shared heap	$\sigma_1$	$\sigma_1$	$\sigma_1$	$\sigma_1$	$\sigma_2$	$\sigma_2$	$\sigma_2$
abstract data-structure	$p_1$	$p_1$	$p_1$	$p_1$	$p_1$	$p_1$	$p_1$

→  $\langle$   
nw = m.next\_write  
nr = m.next\_read  
d = m.data  
d[nw] = x  
nw = (nw+1) mod SIZE  
assume (nr == nw)  
 $\langle$ m.next\_write = nw; LIN $\rangle$

# Reasoning about linearisation using Rely-Guarantee

## Introduction of an intermediate, *instrumented*, language.

- Explicit annotation of *linearisation points*
- Hybrid states, both concrete and abstract
- Linearisation points trigger the abstract semantics

local map	$\rho_1$	$\rho_2$	$\rho_3$	$\rho_4$	$\rho_4$	$\rho_5$	$\rho_5$	$\rho_5$
shared heap	$\sigma_1$	$\sigma_1$	$\sigma_1$	$\sigma_1$	$\sigma_2$	$\sigma_2$	$\sigma_2$	$\sigma_3$
abstract data-structure	$p_1$	$p_1$	$p_1$	$p_1$	$p_1$	$p_1$	$p_1$	$p_2$

```
nw = m.next_write
nr = m.next_read
d = m.data
d[nw] = x
nw = (nw+1) mod SIZE
assume (nr == nw)
<m.next_write = nw; LIN>
```



# Reasoning about linearisation using Rely-Guarantee

## Introduction of an intermediate, *instrumented*, language.

- Explicit annotation of *linearisation points*
- Hybrid states, both concrete and abstract
- Linearisation points trigger the abstract semantics

local map	$\rho_1$	$\rho_2$	$\rho_3$	$\rho_4$	$\rho_4$	$\rho_5$	$\rho_5$	$\rho_5$
shared heap	$\sigma_1$	$\sigma_1$	$\sigma_1$	$\sigma_1$	$\sigma_2$	$\sigma_2$	$\sigma_2$	$\sigma_3$
abstract data-structure	$p_1$	$p_1$	$p_1$	$p_1$	$p_1$	$p_1$	$p_1$	$p_2$
linearisation state	$B$	$B$	$B$	$B$	$B$	$B$	$B$	$A(v) \rightarrow$

```

nw = m.next_write
nr = m.next_read
d = m.data
d[nw] = x
nw = (nw+1) mod SIZE
assume (nr == nw)
<m.next_write = nw; LIN>

```

# Proving linearisability: the perspective of a user



# Proving linearisability: the perspective of a user

**Abstract data structure**

Buf := Empty | Cons x b

b.Push(x) = Cons x b

# Proving linearisability: the perspective of a user

**Abstract data structure**

Buf := Empty | Cons x b

b.Push(x) = Cons x b

**Concrete implementation of methods**

```
nw = m.next_write  
nr = m.next_read  
d = m.data  
d[nw] = x  
nw = (nw+1) mod SIZE  
assume (nr == nw)  
<m.next_write = nw; LIN>
```

# Proving linearisability: the perspective of a user

Abstract data structure


Concrete implementation of methods

Coherence invariant  $I_c$

$\text{Buf} := \text{Empty} \mid \text{Cons } x \text{ b}$

$\text{b.Push}(x) = \text{Cons } x \text{ b}$

```
nw = m.next_write  
nr = m.next_read  
d = m.data  
d[nw] = x  
nw = (nw+1) mod SIZE  
assume (nr == nw)  
<m.next_write = nw; LIN>
```



# Proving linearisability: the perspective of a user

Abstract data structure

Concrete implementation of methods


Coherence invariant  $I_c$

Relies and guarantees  $R_m$   $G_m$

Buf := Empty | Cons x b

b.Push(x) = Cons x b

```
nw = m.next_write  
nr = m.next_read  
d = m.data  
d[nw] = x  
nw = (nw+1) mod SIZE  
assume (nr == nw)  
<m.next_write = nw; LIN>
```



# Proving linearisability: the perspective of a user

**Abstract data structure**

**Concrete implementation of methods**


**Coherence invariant**  $I_c$

**Relies and guarantees**  $R_m$   $G_m$

$\text{Buf} := \text{Empty} \mid \text{Cons } x \text{ b}$

$\text{b.Push}(x) = \text{Cons } x \text{ b}$

```
nw = m.next_write
nr = m.next_read
d = m.data
d[nw] = x
nw = (nw+1) mod SIZE
assume (nr == nw)
<m.next_write = nw; LIN>
```



# Proving linearisability: the perspective of a user

**Abstract data structure**

**Concrete implementation of methods**

**Coherence invariant**  $I_c$

**Relies and guarantees**  $R_m$   $G_m$

# Proving linearisability: the perspective of a user

**Abstract data structure**

**Concrete implementation of methods**

**Coherence invariant**  $I_c$

**Relies and guarantees**  $R_m$   $G_m$

$R_{push}, G_{push}, I_c \vdash$

$\{\mathbf{ln} = B\}$

$p.push(v)$

$\{\mathbf{ln} = A(v_1) \wedge \mathbf{ret} = v_1\}$

**RG method specification**

# Proving linearisability: the perspective of a user

**Abstract data structure**

**Concrete implementation of methods**

**Coherence invariant**  $I_c$

**Relies and guarantees**  $R_m$   $G_m$

$R_{push}, G_{push}, I_c \vdash$

$\{\mathbf{ln} = B\}$

$p.push(v)$

$\{\mathbf{ln} = A(v_1) \wedge \mathbf{ret} = v_1\}$

$I_c$  stable under  $R_{push}$

**RG method specification**

**Stability obligations**



# Proving linearisability: the perspective of a user

**Abstract data structure**

**Concrete implementation of methods**

**Coherence invariant**  $I_c$

**Relies and guarantees**  $R_m$   $G_m$

$R_{push}, G_{push}, I_c \vdash$

$\{\mathbf{ln} = B\}$

$p.push(v)$

$\{\mathbf{ln} = A(v_1) \wedge \mathbf{ret} = v_1\}$

$I_c$  stable under  $R_{push}$

$G_{push} \subseteq R_{pop}$

$G_{pop} \subseteq R_{push}$

**RG method specification**

**Stability obligations**

**RG consistency**

# Proving linearisability: the perspective of a user

**Abstract data structure**

**Concrete implementation of methods**

**Coherence invariant**  $I_c$

**Relies and guarantees**  $R_m$   $G_m$

**RG method specification**

**Stability obligations**

**RG consistency**

Reasoning locally  
exclusively on  
methods



**Automatically  
obtain**

Observational refinement of  
the compilation pass  
implementing the methods  
*for any client*

# Refining linearisable data-structures



Scan:

**repeat**

no\_gray = **true**;

**foreach**  $x \in \text{OBJECTS}$

**if**  $x.\text{color} == \text{GRAY}$

no\_gray = **false**;

**foreach**  $f \in \text{fields}(x)$  **do**

MarkGray( $x.f$ );

$x.\text{color} = \text{BLACK}$

**until** no\_gray

Sweep:

**foreach**  $x \in \text{OBJECTS}$

**if**  $x.\text{color} == \text{WHITE}$

**then** FREE( $x$ )

Clear:

**foreach**  $x \in \text{OBJECTS}$

$x.\text{color} = \text{WHITE}$

**if**  $x.\text{color} = \text{WHITE}$  **then**  
push(buffer[m],  $x$ )

# Refining linearisable data-structures



Scan:

**repeat**

no\_gray = **true**;

**foreach** x ∈ OBJECTS

**if** x.color == GRAY

no\_gray = **false**;

**foreach** f ∈ fields(x) **do**

MarkGray(x.f);

x.color = BLACK

**until** no\_gray

Sweep:

**foreach** x ∈ OBJECTS

**if** x.color == WHITE

**then** FREE(x)

Clear:

**foreach** x ∈ OBJECTS

x.color = WHITE

**if** x.color = WHITE **then**  
push(buffer[m], x)


nw = m.next\_write  
nr = m.next\_read  
d = m.data  
d[nw] = x  
nw = (nw+1) mod SIZE  
assume (nr == nw)  
m.next\_write = nw

# A quick peak under the hood

# Backward simulations

Inductive step used to prove observational refinement

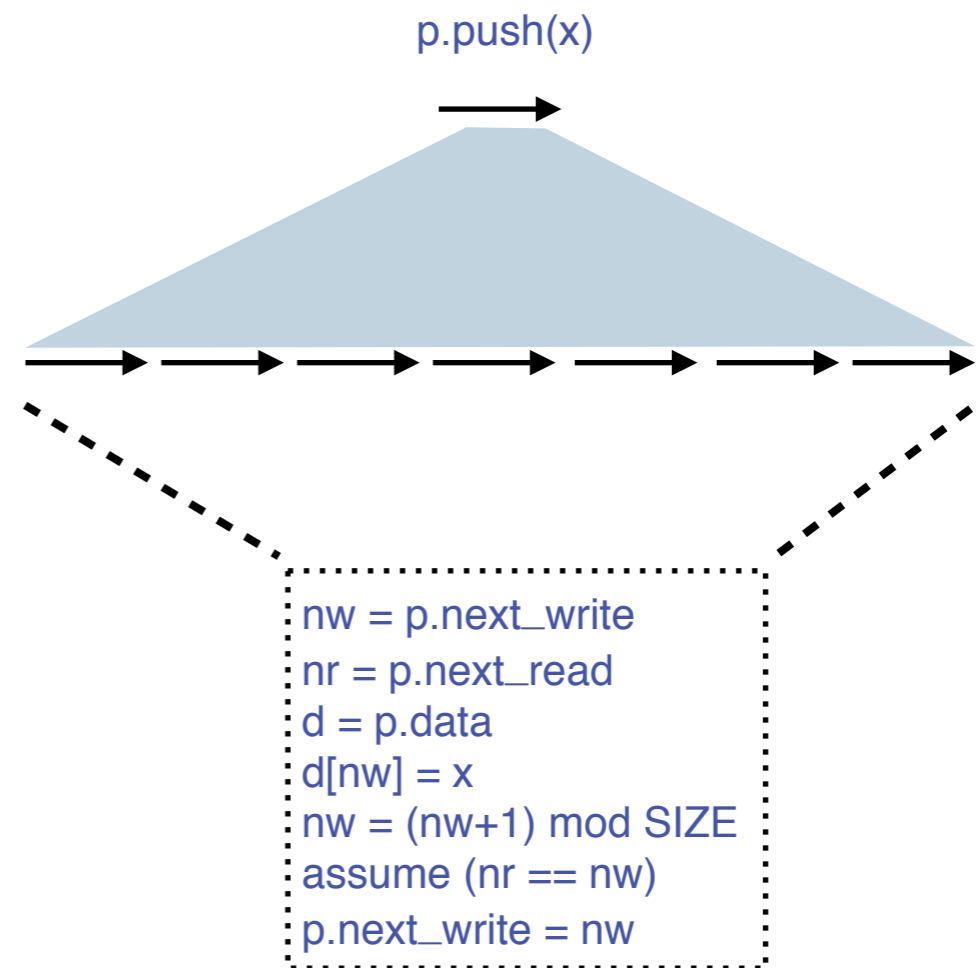

~ Relation between states of the source and target language



# Backward simulations

## Inductive step used to prove observational refinement

~ Relation between states of the source and target language




# Two simulations composed

The compilation pass is split in two phases

- Implementation of the data structure
- Cleaning of the instrumentation

We therefore build two simulations, and compose them





# Structure of the proof: an intuition

## Design and prove a rich invariant at the instrumented level

Objective: carry enough information to leverage the RG specification

- Maintains the coherence invariant
- Builds partial executions of encountered methods

## Prove thread local simulations

- For each thread, build a simulation parameterised by its rely
- Use the partial execution of methods to invoke the RG specification when needed


## Combine the simulations using the stability assumptions

# Conclusion



- Linearisability expressed in term of observational refinement
- A *local*, sufficient condition expressed in terms of Rely-Guarantee
- A *generic* meta-theorem: can be instantiated with any data structure (provided you manage to discharge the proof obligations)
- Provide strong semantic foundations:
  - all theorem expressed wrt an operational semantics
  - everything formalised in Coq
- Instantiated on a realistic example used in another project
- ~13.5 kloc

# Thank you



# Appendix

# Linearisability: limits of our result

## Future-dependent linearisation points

- Example: pair snapshot
- Linearisation is confirmed at a later point of execution
- Need: Maintain two speculative simulations in parallel

## Helping-based linearisation

- Example: HSY elimination-based stack
- Linearisation of thread A is performed by a step from thread B
- Need: Global view of the situation of each thread inside their method

# Separation logic

- **Rely-Guarantee: reasoning about races**
- **Separation logic: proving concisely the absence of races**

Assertions describe more precisely the memory.  
They can be interpreted as ownership of resources.

$$\llbracket r \mapsto v \rrbracket = \{h \mid h(r) = v \wedge \text{dom}(h) = \{r\}\}$$

Achieves great modularity through the *frame rule*

$$\frac{\vdash \{P\} c \{Q\}}{\vdash \{P * R\} c \{Q * R\}}$$

Several works combine RG and SL: RGSep, SAGL, Iris, ...