

# Verifying a Concurrent Garbage Collector using a Rely-Guarantee Methodology

**Yannick Zakowski**

David Cachera

Delphine Demange

Gustavo Petri

David Pichardie

Suresh Jagannathan

Jan Vitek




# Automatic memory management

## Dynamic allocation of memory (in the heap)

In C: `int * array = malloc(10 * sizeof(int))`

In Java: `Point originOne = new Point(23, 94)`


In OCaml: `let u = 'c'::'a'::'m'::'l'::[]`



# Automatic memory management

## Dynamic allocation of memory (in the heap)

In C: `int * array = malloc(10 * sizeof(int))`  
In Java: `Point originOne = new Point(23, 94)`  
In OCaml: `let u = 'c'::'a'::'m'::'l'::[]`



## Manual memory management

Programmer responsible for deallocation (C, C++...)

Risks: premature/double free, memory leak


# Automatic memory management

## Dynamic allocation of memory (in the heap)

In C: `int * array = malloc(10 * sizeof(int))`

In Java: `Point originOne = new Point(23, 94)`

In OCaml: `let u = 'c'::'a'::'m'::'l'::[]`



## Manual memory management

Programmer responsible for deallocation (C, C++...)

Risks: premature/double free, memory leak

## Automatic memory management

Memory reclaimed automatically: Garbage Collector (Lisp, OCaml, Java...)



# Garbage Collection

## Sequential Mark-Sweep Collectors (McCarthy, 1960)

On allocation, if few memory available:

1. Stop the user program
2. Perform full cycle: unreachable memory is reclaimed

Stop-the-world  
Single thread





# Garbage Collection

## Sequential Mark-Sweep Collectors (McCarthy, 1960)

On allocation, if few memory available:

1. Stop the user program
2. Perform full cycle: unreachable memory is reclaimed

Stop-the-world  
Single thread




# Garbage Collection

## Sequential Mark-Sweep Collectors (McCarthy, 1960)

On allocation, if few memory available:

1. Stop the user program
2. Perform full cycle: unreachable memory is reclaimed



Stop-the-world  
Single thread



## With concurrency

Stop-the-world?

Collector  
Multithread  
user code




# Garbage Collection

## Sequential Mark-Sweep Collectors (McCarthy, 1960)

On allocation, if few memory available:

1. Stop the user program
2. Perform full cycle: unreachable memory is reclaimed

Stop-the-world  
Single thread





## With concurrency

Mostly-concurrent?

Collector

Multithread  
user code






# Garbage Collection

## Sequential Mark-Sweep Collectors (McCarthy, 1960)

On allocation, if few memory available:



1. Stop the user program
2. Perform full cycle: unreachable memory is reclaimed

Stop-the-world  
Single thread



With concurrency

On-the-fly!



# On-the-fly garbage collection

**CACM '78**

---

## On-the-Fly Garbage Collection: An Exercise in Cooperation

Edsger W. Dijkstra  
Burroughs Corporation

Leslie Lamport  
SRI International

A.J. Martin, C.S. Scholten, and  
E.F.M. Steffens  
Philips Research Laboratories

---

**As an example of cooperation between sequential  
processes with very little mutual interference despite**

# On-the-fly garbage collection

CACM '78

## On-the-Fly Garbage Collection: An Exercise in Cooperation

Edsger W. Dijkstra  
Burroughs Corporation

Leslie Lamport  
SRI International

technique is developed which allows nearly all of the activity needed for garbage detection and collection to be performed by an additional processor operating concurrently with the processor devoted to the computation proper. Exclusion and synchronization constraints have been kept as weak as could be achieved; the severe complexities engendered by doing so are illustrated.

# On-the-fly garbage collection

CACM '78

## On-the-Fly Garbage Collection: An Exercise in Cooperation

Edsger W. Dijkstra  
Burroughs Corporation

Leslie Lamport  
SRI International



POPL '93 '94

A concurrent, generational garbage collector for a multithreaded implementation of ML

Damien Doligez

Xavier Leroy

## Portable, Unobtrusive Garbage Collection for Multiprocessor Systems

Damien Doligez  
École Normale Supérieure  
INRIA Rocquencourt  
École Polytechnique  
Damien.Doligez@inria.fr

Georges Gonthier\*  
INRIA Rocquencourt  
78153 LE CHESNAY CEDEX  
FRANCE  
Georges.Gonthier@inria.fr

technique is developed which allows nearly all of the activity needed for garbage detection and collection to be performed by an additional processor operating concurrently with the processor devoted to the computation proper. Exclusion and synchronization constraints have been kept as weak as could be achieved; the severe complexities engendered by doing so are illustrated.

# On-the-fly garbage collection

CACM '78

## On-the-Fly Garbage Collection: An Exercise in Cooperation

Edsger W. Dijkstra  
Burroughs Corporation

Leslie Lamport  
SRI International



POPL '93 '94

A concurrent, generational garbage collector for a multithreaded implementation of ML

Damien Doligez

Xavier Leroy

Portable, Unobtrusive Garbage Collection for Multiprocessor Systems

Damien Doligez

Georges Gonthier\*

technique is developed which allows nearly all of the activity needed for garbage detection and collection to be performed by an additional processor operating concurrently with the processor devoted to the computation proper. Exclusion and synchronization constraints have been kept as weak as could be achieved; the severe complexities engendered by doing so are illustrated.

is the d  
the basic algorithm of [9], and expose a series of counterexamples to explain why a straightforward adaptation of this algorithm to multiple mutators would not work; we also address some efficiency issues. In section 4 we describe the

# On-the-fly garbage collection

CACM '78

## On-the-Fly Garbage Collection: An Exercise in Cooperation

Edsger W. Dijkstra  
Burroughs Corporation

Leslie Lamport  
SRI International



POPL '93 '94

A concurrent, generational garbage collector for a multithreaded implementation of ML

Damien Doligez

Xavier Leroy

Portable, Unobtrusive Garbage Collection for Multiprocessor Systems

Damien Doligez

Georges Gonthier\*

is the d  
the basic algorithm of [9], and expose a series of counterexamples to explain why a straightforward adaptation of this algorithm to **multiple mutators would not work**; we also address some efficiency issues. In section 4 we describe the

PLDI '00

## Implementing an On-the-fly Garbage Collector for Java

Tamar Domani

Elliot K. Kolodner\*

Ethan Lewis

Eliot E. Salant

Katherine Barabash

Itai Lahan

Erez Petrank

Igor Yanover

Yossi Levroni

Abstract

introduced for LISP in 1960. Since then, garbage collection has been adapted for many



# On-the-fly garbage collection

CACM '78

## On-the-Fly Garbage Collection: An Exercise in Cooperation

Edsger W. Dijkstra  
Burroughs Corporation

Leslie Lamport  
SRI International



POPL '93 '94

A concurrent, generational garbage collector for a multithreaded implementation of ML

Damien Doligez

Xavier Leroy

Portable, Unobtrusive Garbage Collection for Multiprocessor Systems

Damien Doligez

Georges Gonthier\*

is the d  
the basic algorithm of [9], and expose a series of counterexamples to explain why a straightforward adaptation of this algorithm to multiple mutators would not work; we also address some efficiency issues. In section 4 we describe the

PLDI '00

## Implementing an On-the-fly Garbage Collector for Java

Tamar Domani

Elliot K. Kolodner\*

Ethan Lewis

Eliot E. Salant

Katherine Barabash

Itai Lahan

Erez Petrank

Igor Yanover

Yossi Levroni

Abstract

introduce mark buffers, adapt the algorithm to Java



# An on-the-fly garbage collector




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked






# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked

Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked






# Sequential mark & sweep

Periodically stop user code and perform:

- Mark: graph traversal from the roots
- Sweep: free cells remaining unmarked


Color conventions:

- Blue: free cells
- White: not marked
- Gray: being visited (pending nodes)
- Black: marked




# Concurrent mark & sweep

## 1. A need for the *mutators*' collaboration




# Concurrent mark & sweep



While the collector runs, the user thread can do:


# Concurrent mark & sweep



While the collector runs, the user thread can do:

1. an update;

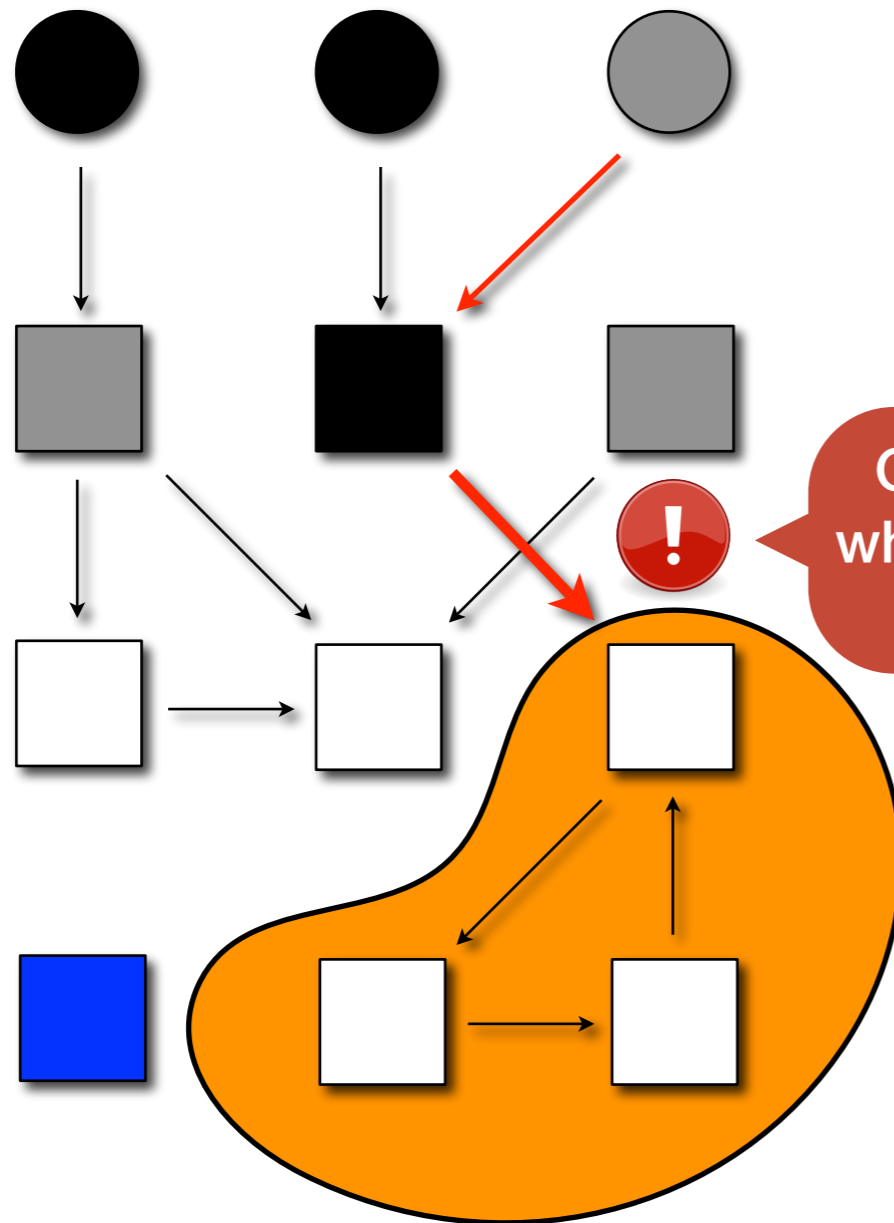
# Concurrent mark & sweep



While the collector runs, the user thread can do:

1. an update;
2. a store.

# Concurrent mark & sweep

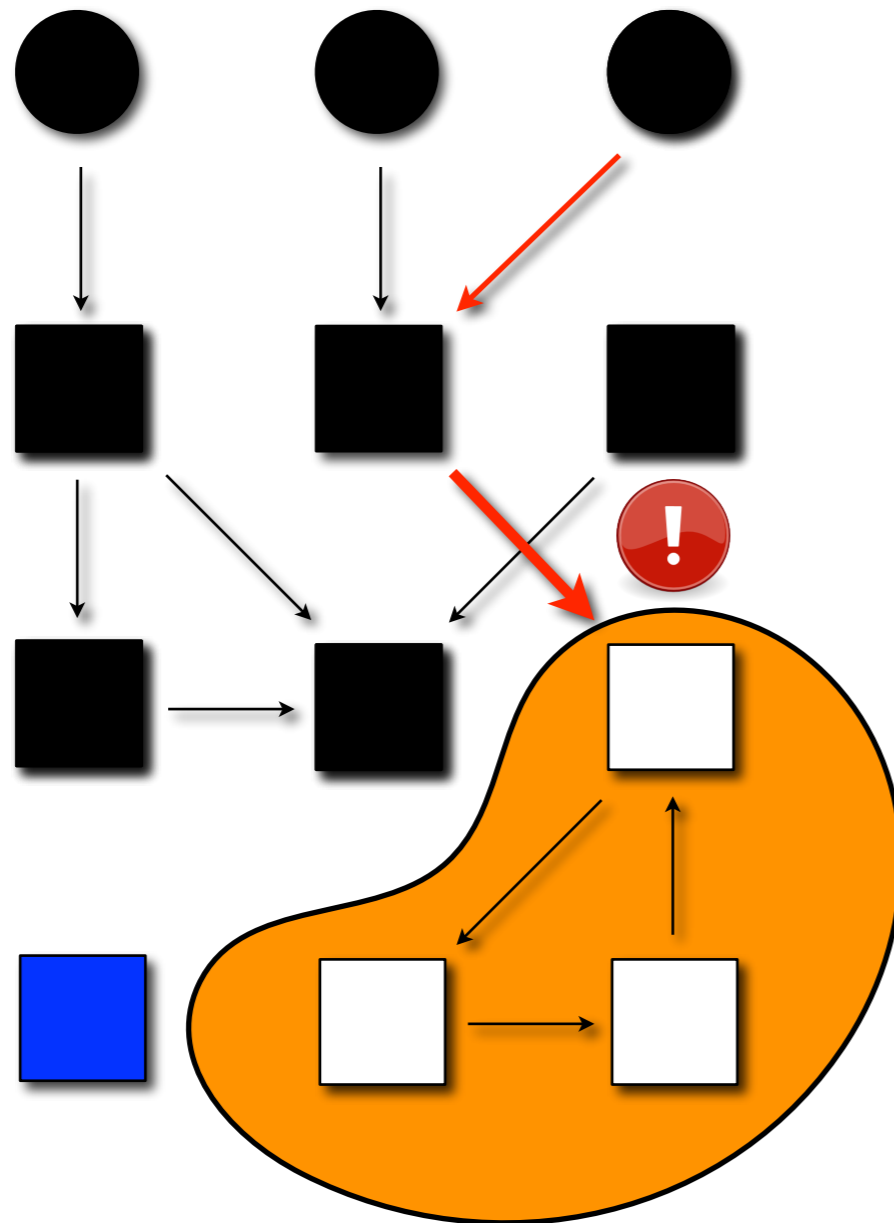


While the collector runs, the user thread can do:

1. an update;
2. a store.

Only access to white subgraph is black!


# Concurrent mark & sweep



While the collector runs, the user thread can do:

1. an update;
2. a store.

# Concurrent mark & sweep




While the collector runs, the user thread can do:

1. an update;
2. a store.



# Concurrent mark & sweep



While the collector runs, the user thread can do:


1. an update;
2. a store.

Updates go through write barriers

```
Update(x,f,y) ==  
  MarkGray(y);  
  x.f = y
```

```
MarkGray(x) ==  
  if x.color = WHITE  
  then x.color = GRAY
```

# Concurrent mark & sweep



While the collector runs, the user thread can do:


1. an update;
2. a store.

Updates go through write barriers

```
Update(x,f,y) ==  
  MarkGray(y);  
  x.f = y
```

```
MarkGray(x) ==  
  if x.color = WHITE  
  then x.color = GRAY
```

# Concurrent mark & sweep



While the collector runs, the user thread can do:


1. an update;
2. a store.

Updates go through write barriers

```
Update(x,f,y) ==  
  MarkGray(y);  
  x.f = y
```

```
MarkGray(x) ==  
  if x.color = WHITE  
  then x.color = GRAY
```

# Concurrent mark & sweep



While the collector runs, the user thread can do:


1. an update;
2. a store.

Updates go through write barriers

```
Update(x,f,y) ==  
  MarkGray(y);  
  x.f = y
```

```
MarkGray(x) ==  
  if x.color = WHITE  
  then x.color = GRAY
```

# Concurrent mark & sweep



While the collector runs, the user thread can do:

1. an update;
2. a store.

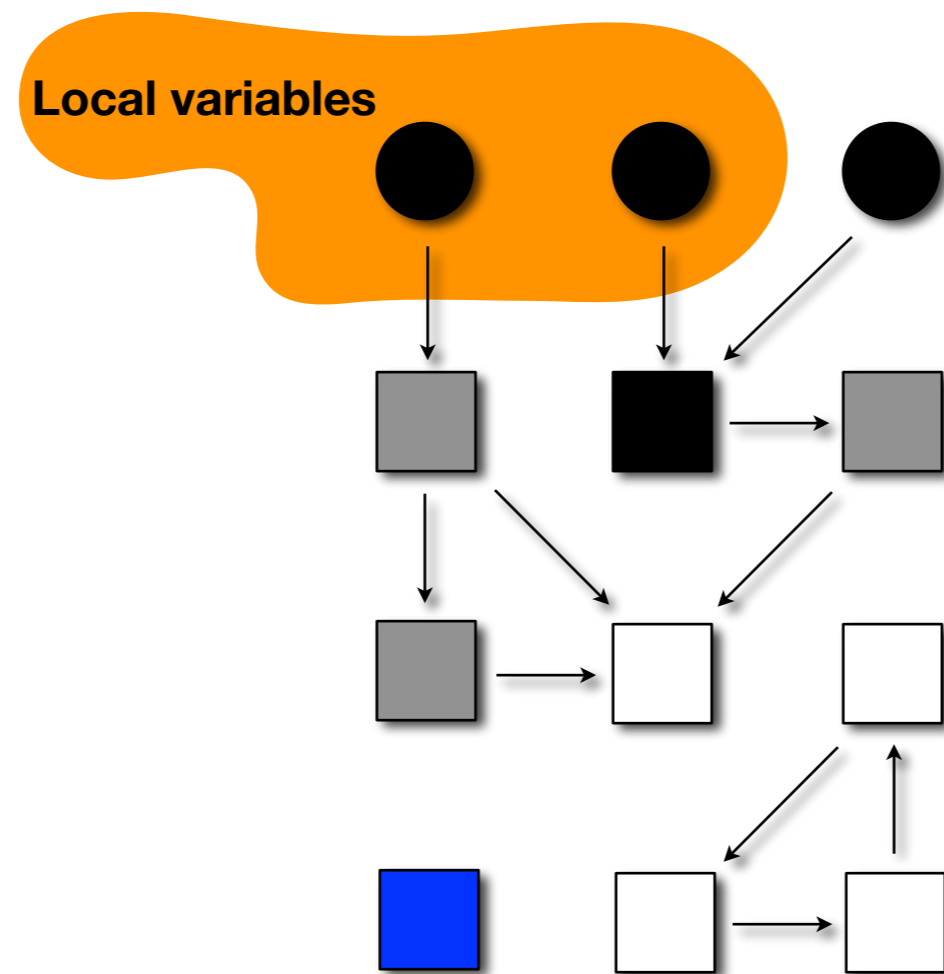
Updates go through write barriers

```
Update(x,f,y) ==  
  MarkGray(y);  
  x.f = y
```

```
MarkGray(x) ==  
  if x.color = WHITE  
  then x.color = GRAY
```

# Concurrent mark & sweep

## 2. A need for synchronisation



# Concurrent mark & sweep

## MUTATOR

```
[...]  
  Update(x,f,y);  
[...]  
  Alloc();  
[...]
```

## 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
```

# Concurrent mark & sweep

## MUTATOR

```
[...]  
  Update(x,f,y);  
[...]  
  Alloc();  
[...]
```

## 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
```

The collector has not access to all mutator roots...



# Concurrent mark & sweep

MUTATOR

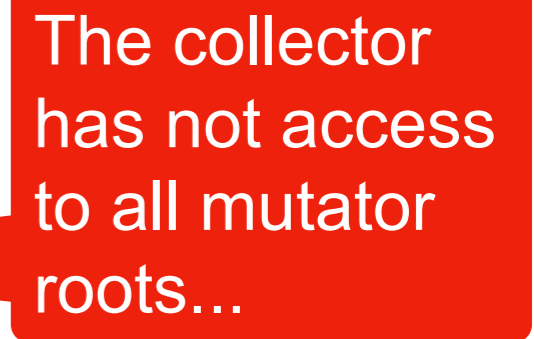
```
[...]  
  Update(x,f,y);  
[...]  
  Alloc();  
[...]
```

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
```



mark your roots please



The collector has not access to all mutator roots...

# Concurrent mark & sweep

MUTATOR

```
[...]  
  Udpate(x1,f1,y1);  
[...]  
  Cooperate();  
[...]  
  Alloc();  
[...]
```

Mark:

```
  Handshake();  
  foreach x ∈ GLOBALS  
  do MarkGray(x)
```

They need to  
synchronise!

```
  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
```

# Concurrent mark & sweep

MUTATOR

```
[...]  
  Udpate(x1,f1,y1);  
[...]  
  Cooperate();  
[...]  
  Alloc();  
[...]
```

Mark:

```
  Handshake();  
  foreach x ∈ GLOBALS  
  do MarkGray(x)
```

They need to  
synchronise!

```
    gray = true;  
    foreach x ∈ OBJECTS  
    x.color == GRAY  
    no_gray = false;  
    foreach f ∈ fields(x) do  
      MarkGray(x.f);  
    color = BLACK  
    no_gray
```

But user threads  
should not wait!

```
    foreach x ∈ OBJECTS  
    x.color == WHITE  
    then FREE(x)
```

Clear:

```
  foreach x ∈ OBJECTS  
  x.color = WHITE
```

# Concurrent mark & sweep

Mutators mark  
their roots

```
Cooperate() =  
  if statusm = statusc  
  then  
    skip  
  else  
    foreach r ∈ LOCAL_ROOTS do  
      MarkGray(r);  
    statusm = statusc;
```

Collector awaits for mutators  
to mark their roots

```
Handshake() =  
  statusc = Next(statusc);  
  while (statusm ≠ statusc) skip;
```

# Concurrent mark & sweep

Mutators mark  
their roots

```
Cooperate() =  
  if statusm = statusc  
  then  
    skip  
  else  
    foreach r ∈ LOCAL_ROOTS do  
      MarkGray(r);  
    statusm = statusc;
```

Collector awaits for mutators  
to mark their roots

```
Handshake() =  
  statusc = Next(statusc);  
  while (statusm ≠ statusc) skip;
```

**Updates  
its status  
Actively waits**

# Concurrent mark & sweep

Mutators mark  
their roots

```
Cooperate() =  
  if statusm = statusc  
  then  
    skip  
  else  
    foreach r ∈ LOCAL_ROOTS do  
      MarkGray(r);  
    statusm = statusc;
```

Checks if in sync  
Yes? Back to work  
No? Publishes roots } Never waits  
for anyone


Collector awaits for mutators  
to mark their roots

```
Handshake() =  
  statusc = Next(statusc);  
  while (statusm ≠ statusc) skip;
```

Updates  
its status  
Actively waits

# Timeline of a collection cycle (DLG 93/94)

# Timeline of a collection cycle (DLG 93/94)



- Graph algorithms
  - Subtle synchronisation
  - Lots of concurrent accesses
- } Sophisticated invariants



# Verifying the garbage collector

1. What do we prove?
2. How do we prove it?

# What do we prove?

**We rely on a *most general client* (mgc)**

`mgc  $\triangleq$`   
`collector || mutator || ... || mutator`

`mutator  $\triangleq$`   
`loop( update(x, f, v)`  
`$\oplus$  load(x, f)  $\oplus$  alloc()`  
`$\oplus$  cooperate()`  
`$\oplus$  changeRoots())`

# What do we prove?

**We rely on a *most general client* (mgc)**

`mgc`  $\triangleq$

`collector || mutator || ... || mutator`

`mutator`  $\triangleq$


`loop( update(x, f, v)  
⊕ load(x, f) ⊕ alloc()  
⊕ cooperate()  
⊕ changeRoots())`

Abstraction of  
any user thread

# What do we prove?

## Theorem:

In any execution state of the mgc, cells reachable by a mutator are never blue



# An IR to program the GC

<code>cmd</code>	<code>:=</code>	
	<code>skip</code>	<code>x = alloc(rn)</code>
	<code>assume e</code>	<code>free(x)</code>
	<code>x = [y].f</code>	<code>x = y.empty?()</code>
	<code>[x].f = e</code>	<code>x = y.top()</code>
	<code>atomic c</code>	<code>x.push(y)</code>
	<code>c<sub>1</sub> ; c<sub>2</sub></code>	<code>x.pop()</code>
	<code>c<sub>1</sub> ⊕ c<sub>2</sub></code>	<code>isFree?(x)</code>
	<code>loop(c)</code>	<code>foreach (x in l) do c od</code>

# An IR to program the GC

cmd :=

skip  
assume  $e$   
 $x = [y].f$   
 $[x].f = e$   
atomic  $c$   
 $c_1 ; c_2$   
 $c_1 \oplus c_2$   
loop( $c$ )

$x = \text{alloc}(rn)$   
 $\text{free}(x)$   
 $x = y.\text{empty}?$   
 $x = y.\text{top}()$   
 $x.\text{push}(y)$   
 $x.\text{pop}()$   
 $\text{isFree?}(x)$   
foreach ( $x$  in  $l$ ) do  $c$  od

# An IR to program the GC

cmd :=

skip  
assume  $e$   
 $x = [y].f$   
 $[x].f = e$   
atomic  $c$   
 $c_1 ; c_2$   
 $c_1 \oplus c_2$   
loop( $c$ )

$x = \text{alloc}(rn)$   
free( $x$ )  
 $x = y.\text{empty}?$   
 $x = y.\text{top}()$   
 $x.\text{push}(y)$   
 $x.\text{pop}()$   
isFree?( $x$ )  
foreach ( $x$  in  $l$ ) do  $c$  od

# An IR to program the GC

cmd :=

```
skip  
assume  $e$   
 $x = [y].f$   
 $[x].f = e$   
atomic  $c$   
 $c_1 ; c_2$   
 $c_1 \oplus c_2$   
loop( $c$ )
```

```
 $x = \text{alloc}(rn)$   
free( $x$ )  
 $x = y.\text{empty}?$   
 $x = y.\text{top}()$   
 $x.\text{push}(y)$   
 $x.\text{pop}()$   
isFree?( $x$ )  
foreach ( $x$  in  $l$ ) do  $c$  od
```

- Abstract buffers:
  - concrete implementations are linearizable

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



# An IR to program the GC

cmd :=

```
skip
assume  $e$ 
 $x = [y].f$ 
 $[x].f = e$ 
atomic  $c$ 
 $c_1 ; c_2$ 
 $c_1 \oplus c_2$ 
loop( $c$ )
```

```
 $x = \text{alloc}(rn)$ 
free( $x$ )
```

```
 $x = y.\text{empty}?$ ()
 $x = y.\text{top}()$ 
```

```
 $x.\text{push}(y)$ 
```

```
 $x.\text{pop}()$ 
```

```
isFree?( $x$ )
```

```
foreach ( $x$  in  $l$ ) do  $c$  od
```

- Intrinsic support for threads, roots, and objects
- Built-in iterator constructs : disciplined access

# An IR to program the GC

cmd :=

skip  
assume  $e$   
 $x = [y].f$   
 $[x].f = e$   
atomic  $c$   
 $c_1 ; c_2$   
 $c_1 \oplus c_2$   
loop( $c$ )

$x = \text{alloc}(rn)$   
free( $x$ )  
 $x = y.\text{empty}?$   
 $x = y.\text{top}()$   
 $x.\text{push}(y)$   
 $x.\text{pop}()$   
isFree?( $x$ )  
foreach ( $x$  in  $l$ ) do  $c$  od

Right level of abstraction: proofs are conducted with respect to the operational semantics of the IR, directly over the code

# Rely Guarantee reasoning

[Jones81]

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

Environment  
R: Rely  
G: Guarantee

Global Correctness  
Invariant

Annotations

# Rely Guarantee reasoning

[Jones81]

$R, G, I \vdash_t \{P\} c \{Q\}$


Environment  
R: Rely  
G: Guarantee

Global Correctness  
Invariant

Annotations


# Rely Guarantee reasoning

[Jones81]



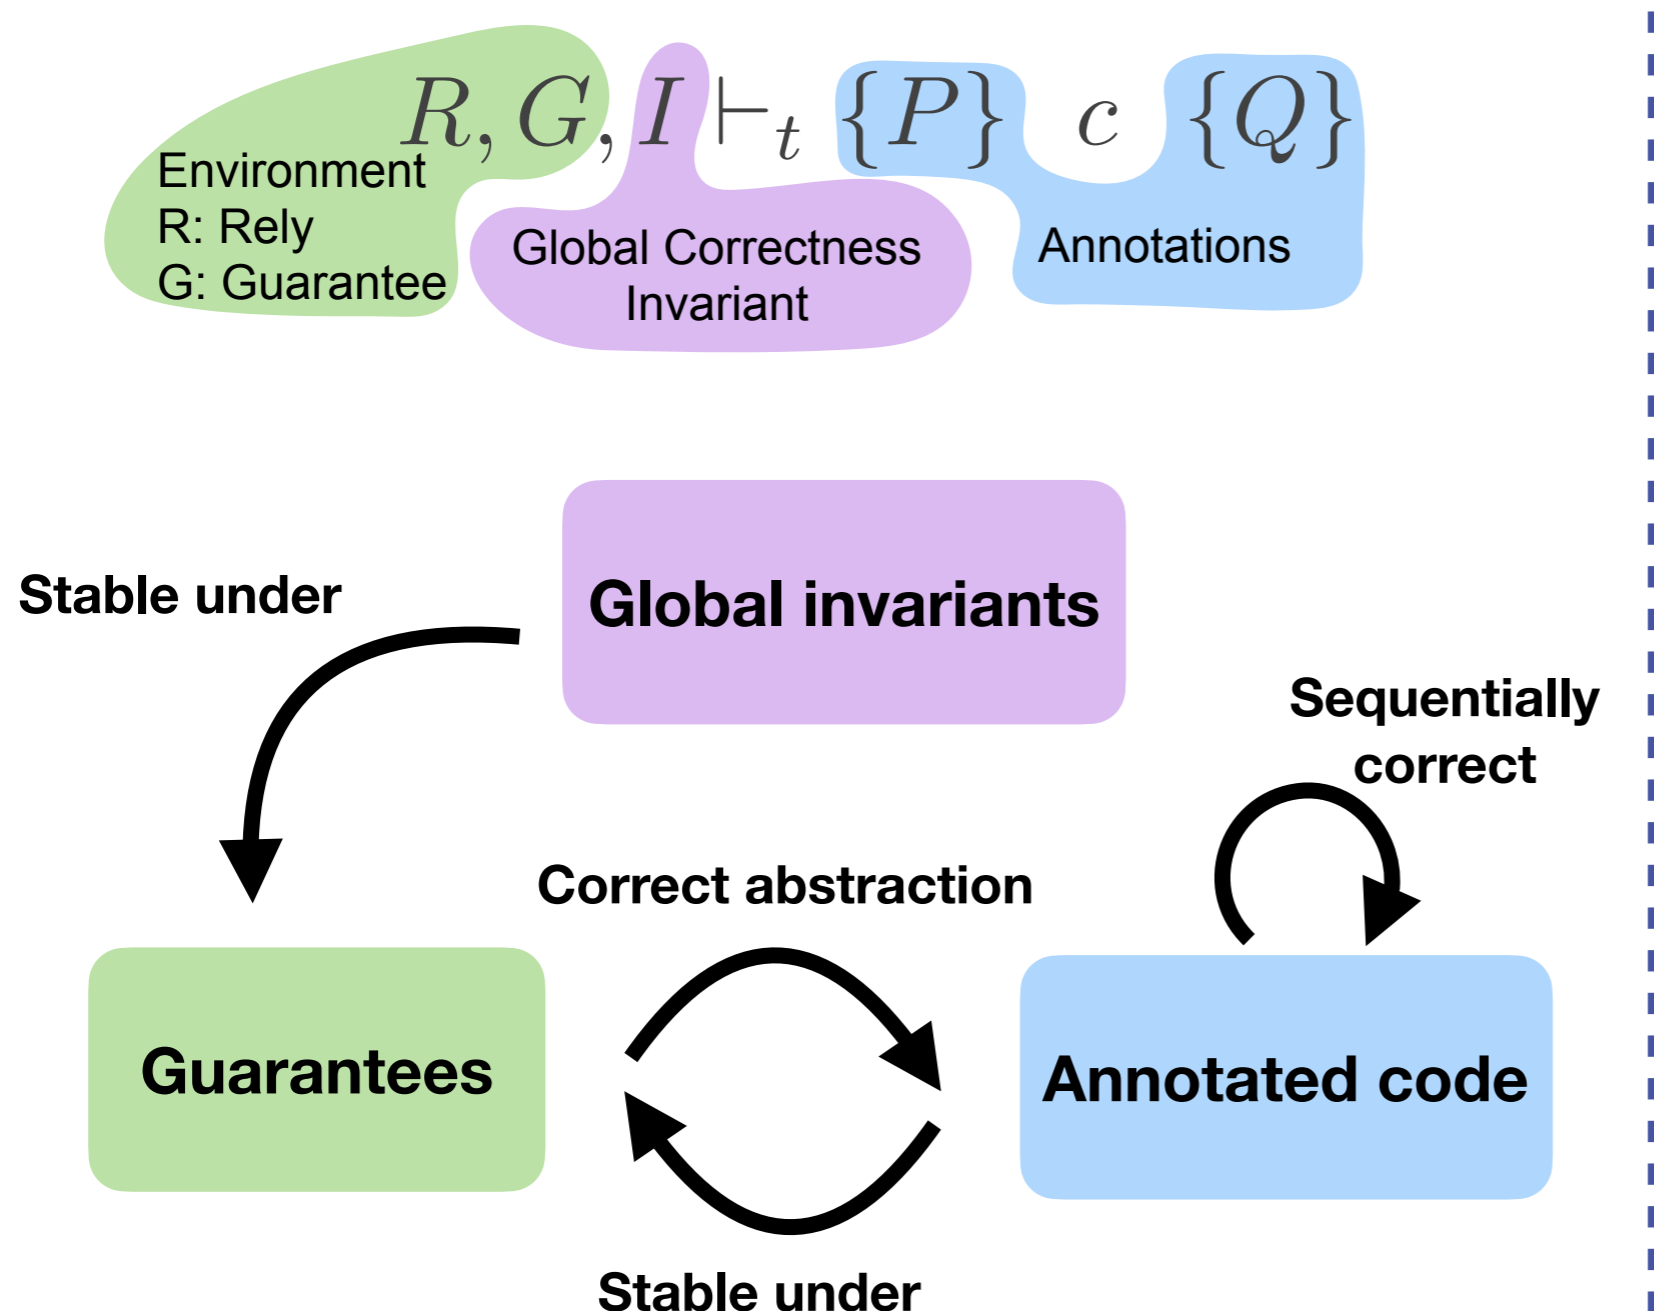
# Rely Guarantee reasoning

[Jones81]





# Rely Guarantee reasoning

[Jones81]



# Rely Guarantee reasoning


[Jones81]





# Rely Guarantee reasoning

[Jones81]

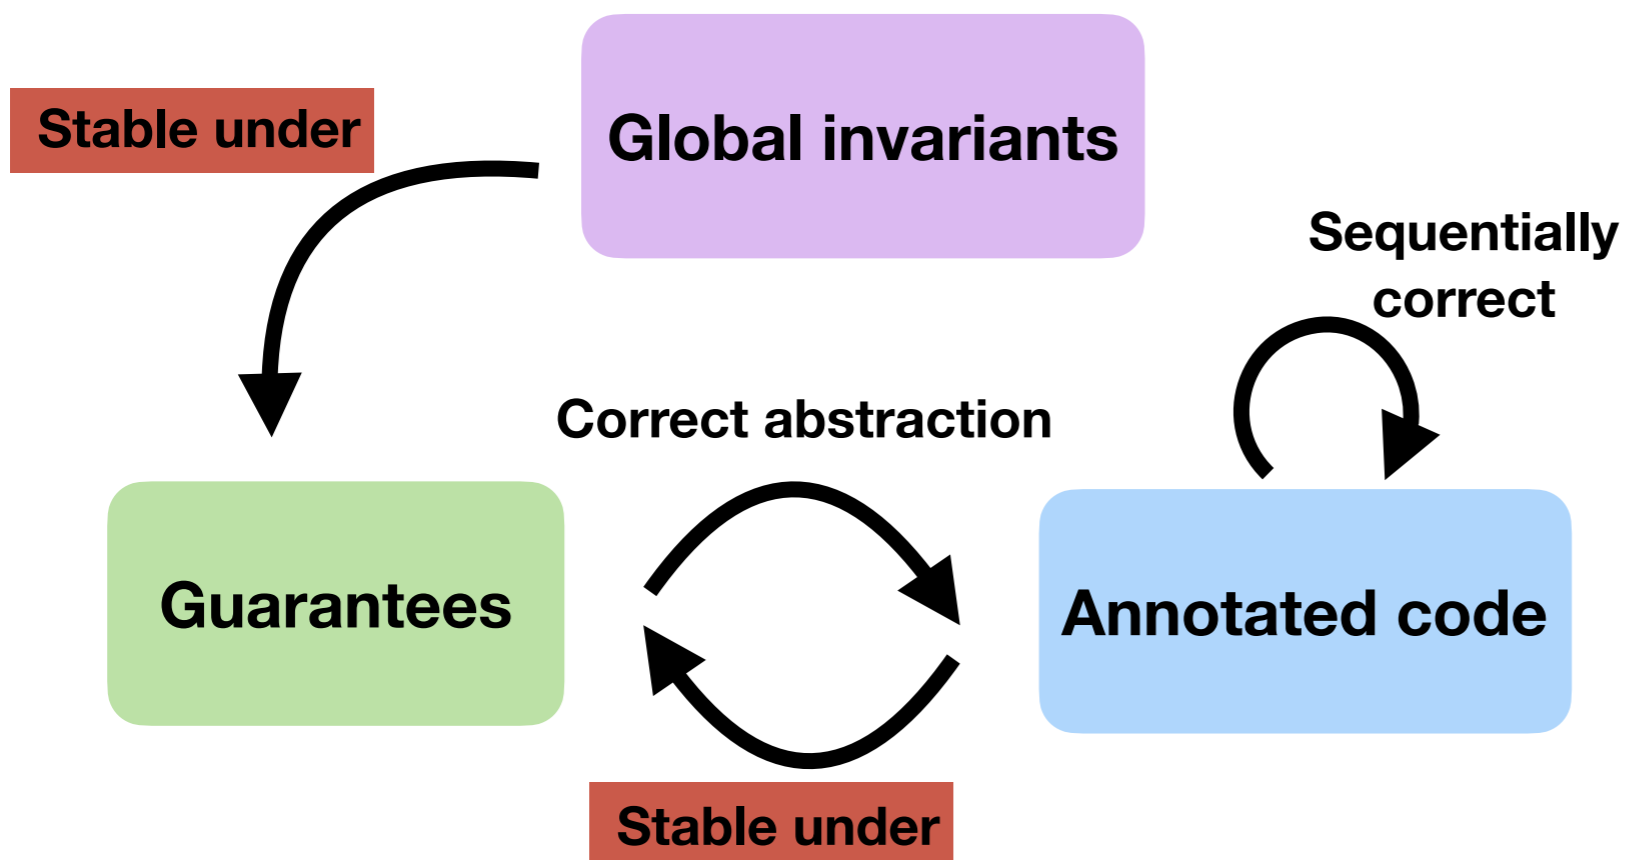


Predicate  $P$

*stable under*



interference  $R$

$$\forall \sigma \vdash P, \forall \sigma', \sigma \ R \ \sigma' \rightarrow \sigma' \vdash P$$



# Rely Guarantee reasoning

[Jones81]



Predicate  $P$

*stable under*

interference  $R$


$$\forall \sigma \vdash P, \forall \sigma', \sigma R \sigma' \rightarrow \sigma' \vdash P$$

**(#annotations \* #interferences)  
stability proofs required!**

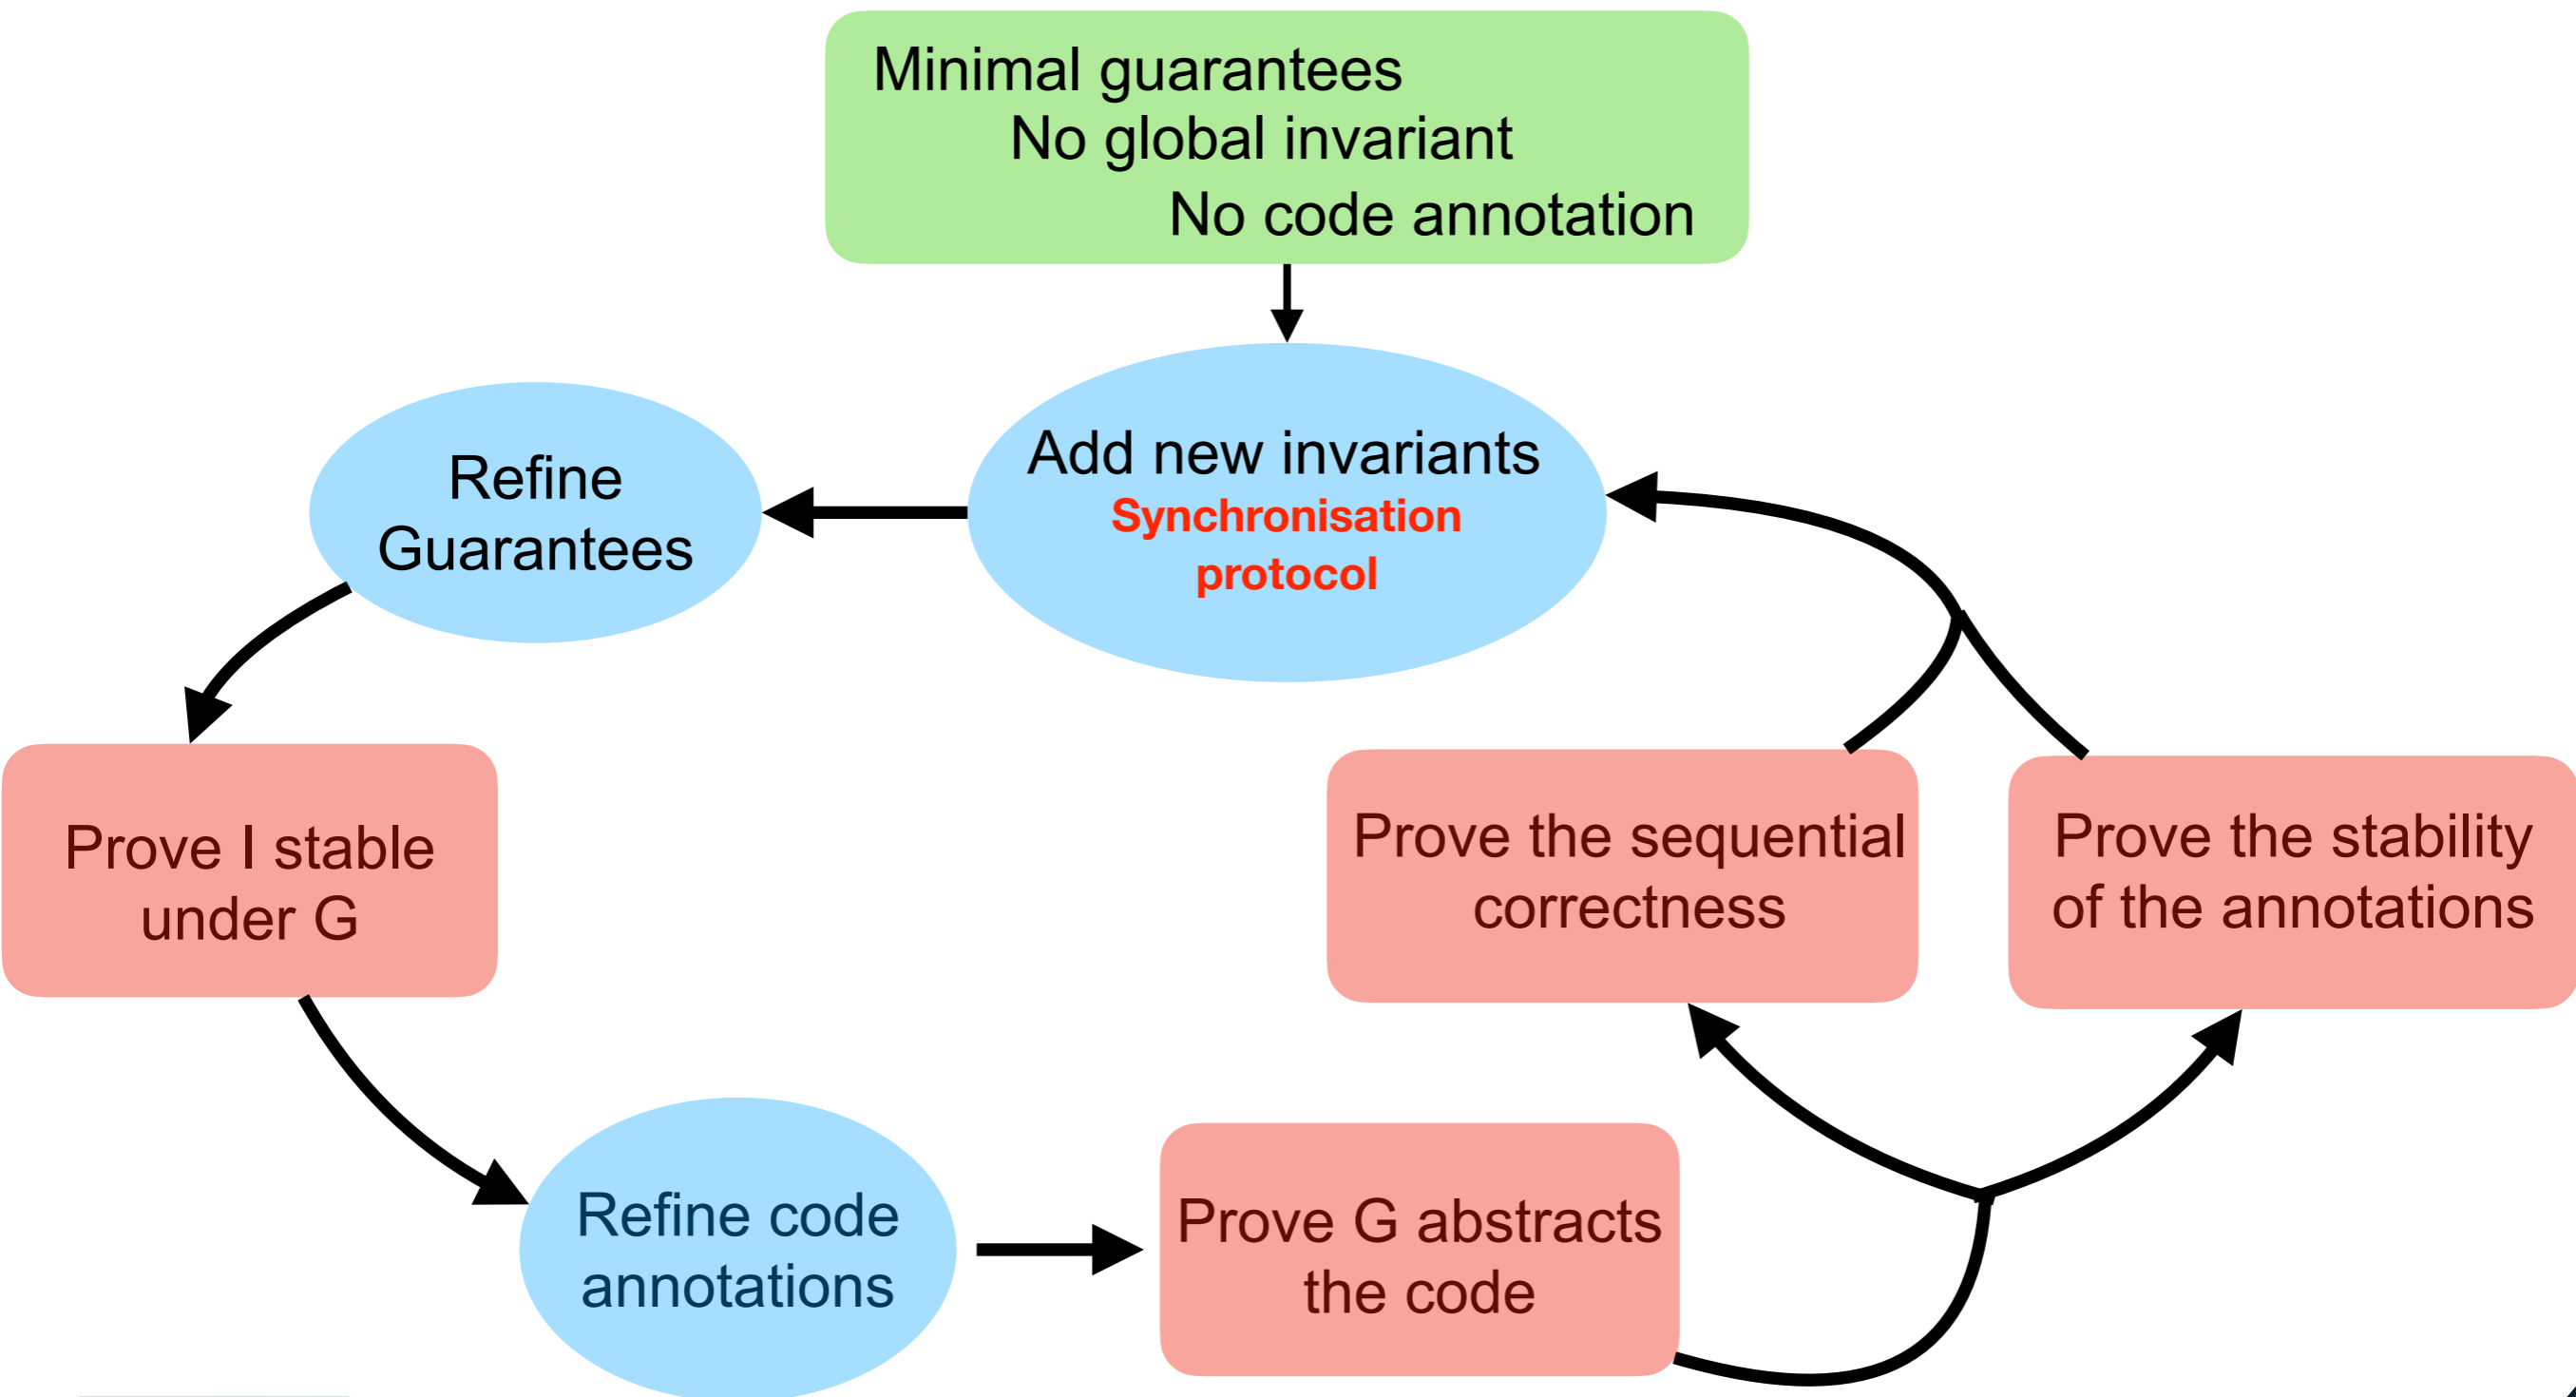
# Reusing proofs incrementally

Minimal guarantees  
No global invariant  
No code annotation

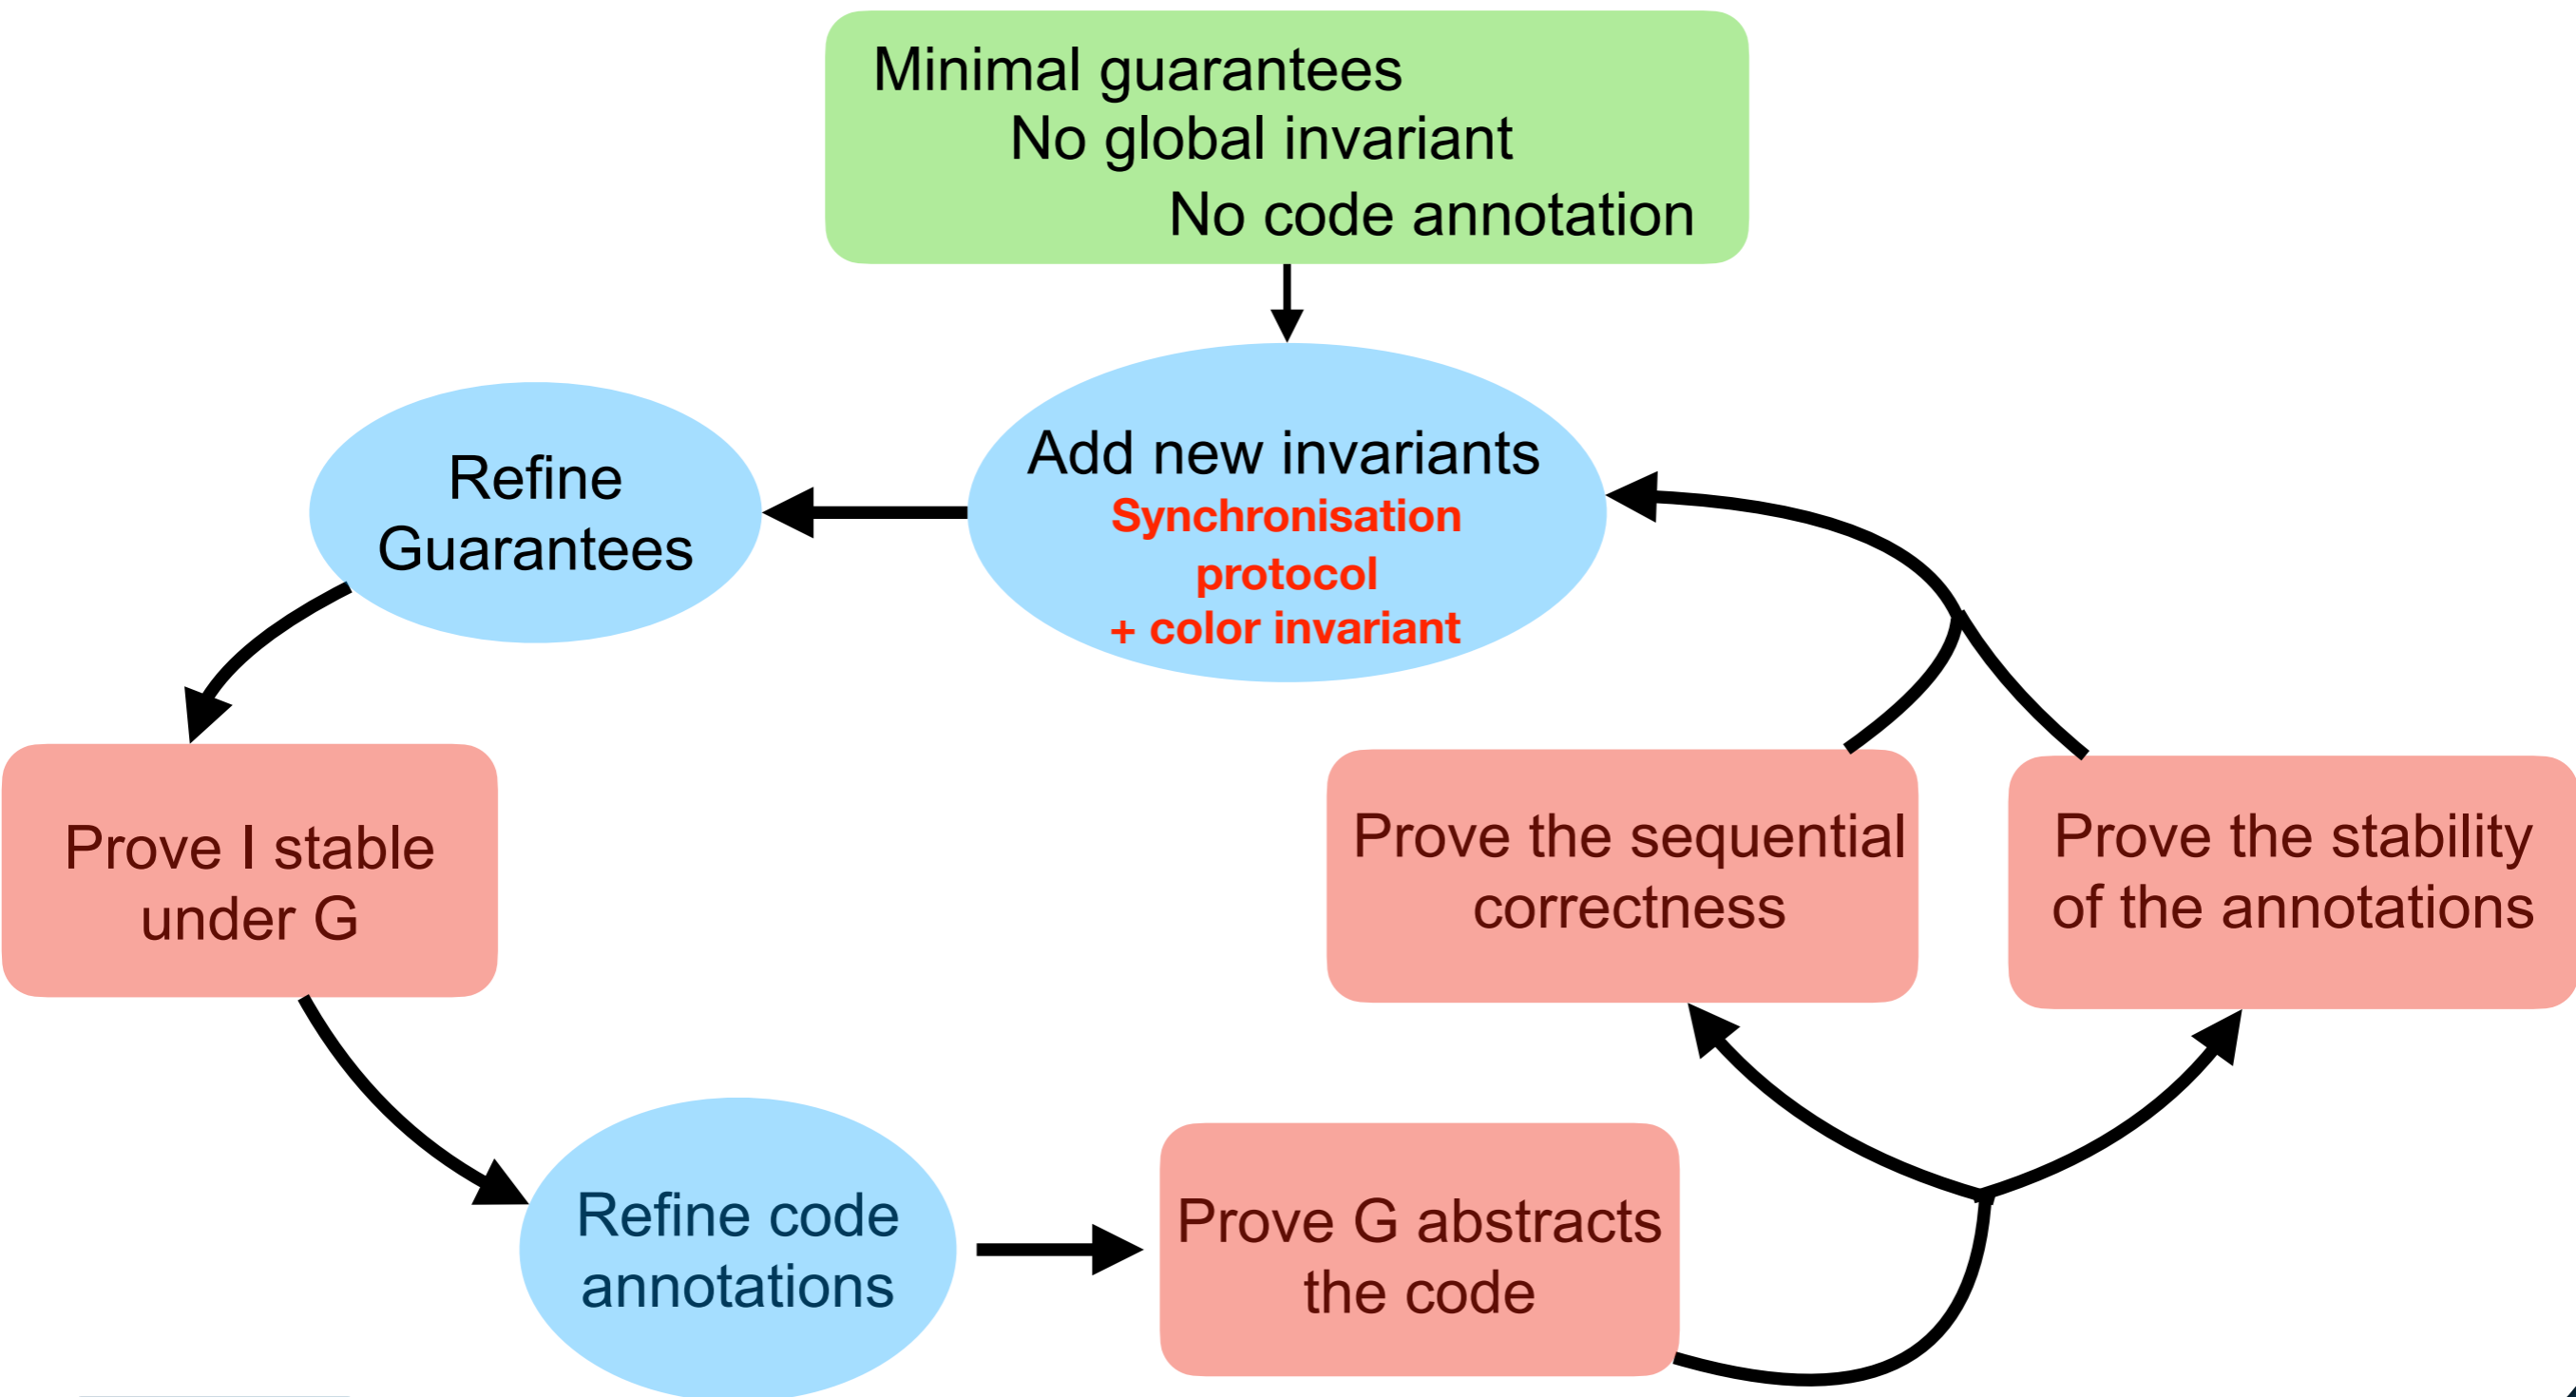
# Reusing proofs incrementally



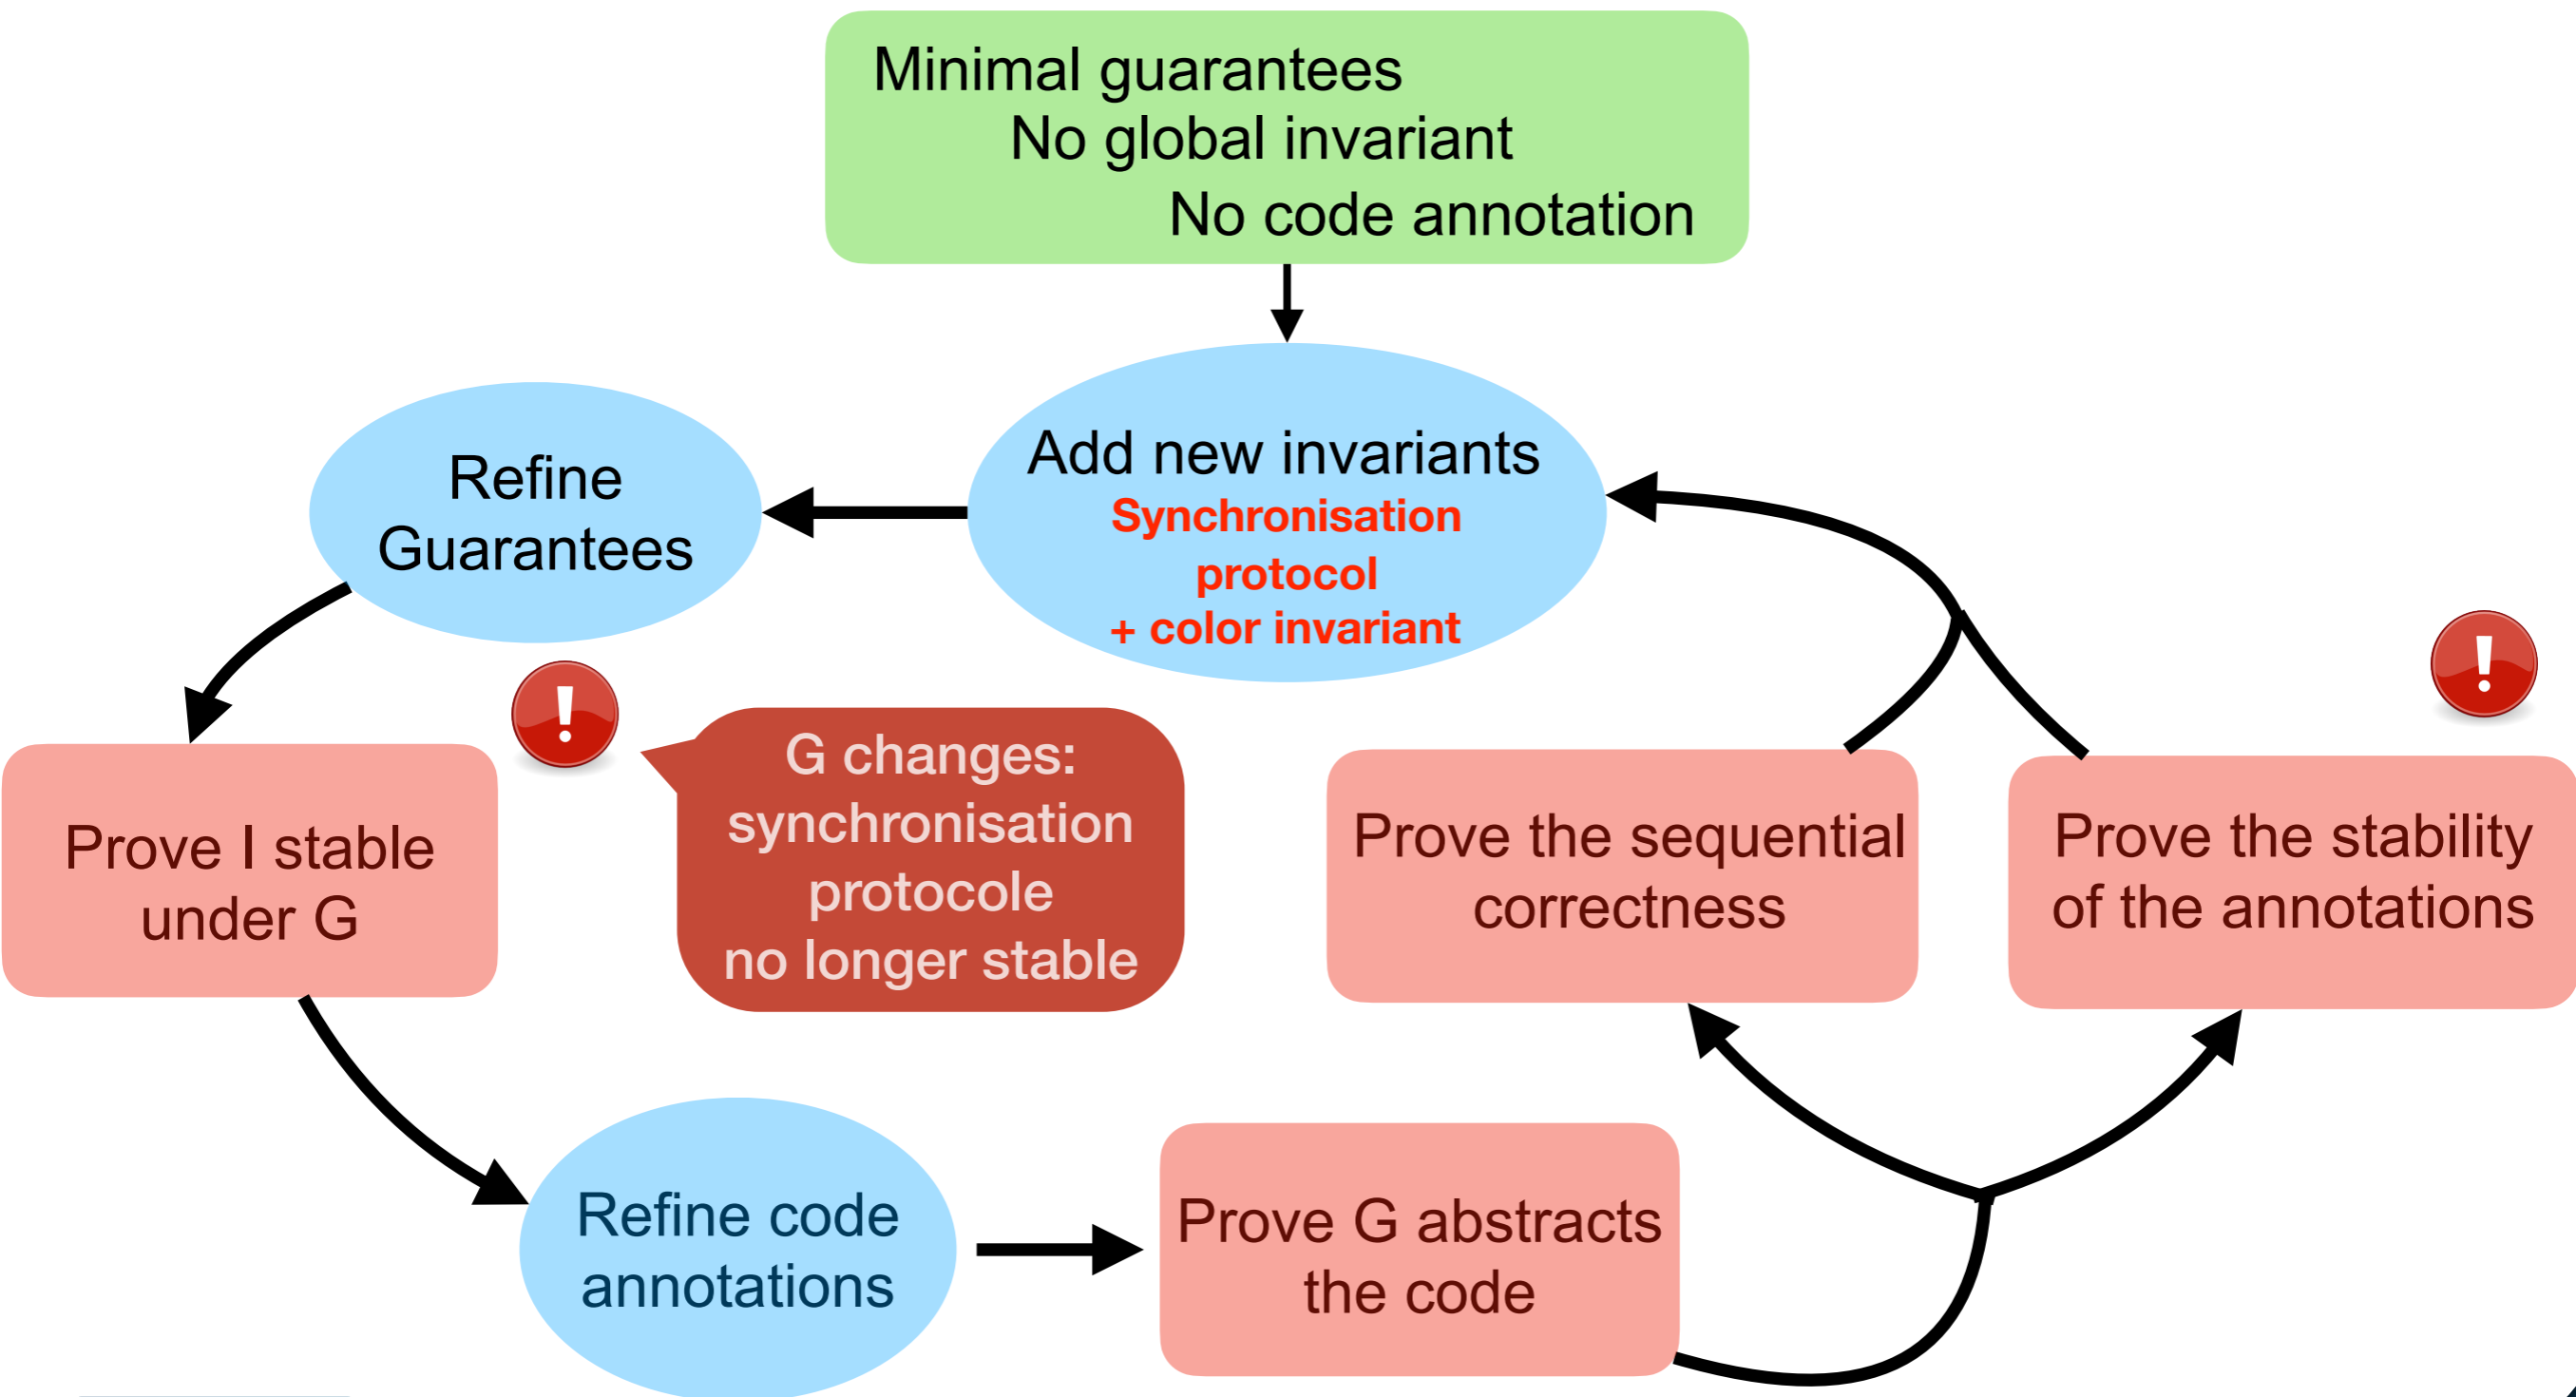
# Reusing proofs incrementally



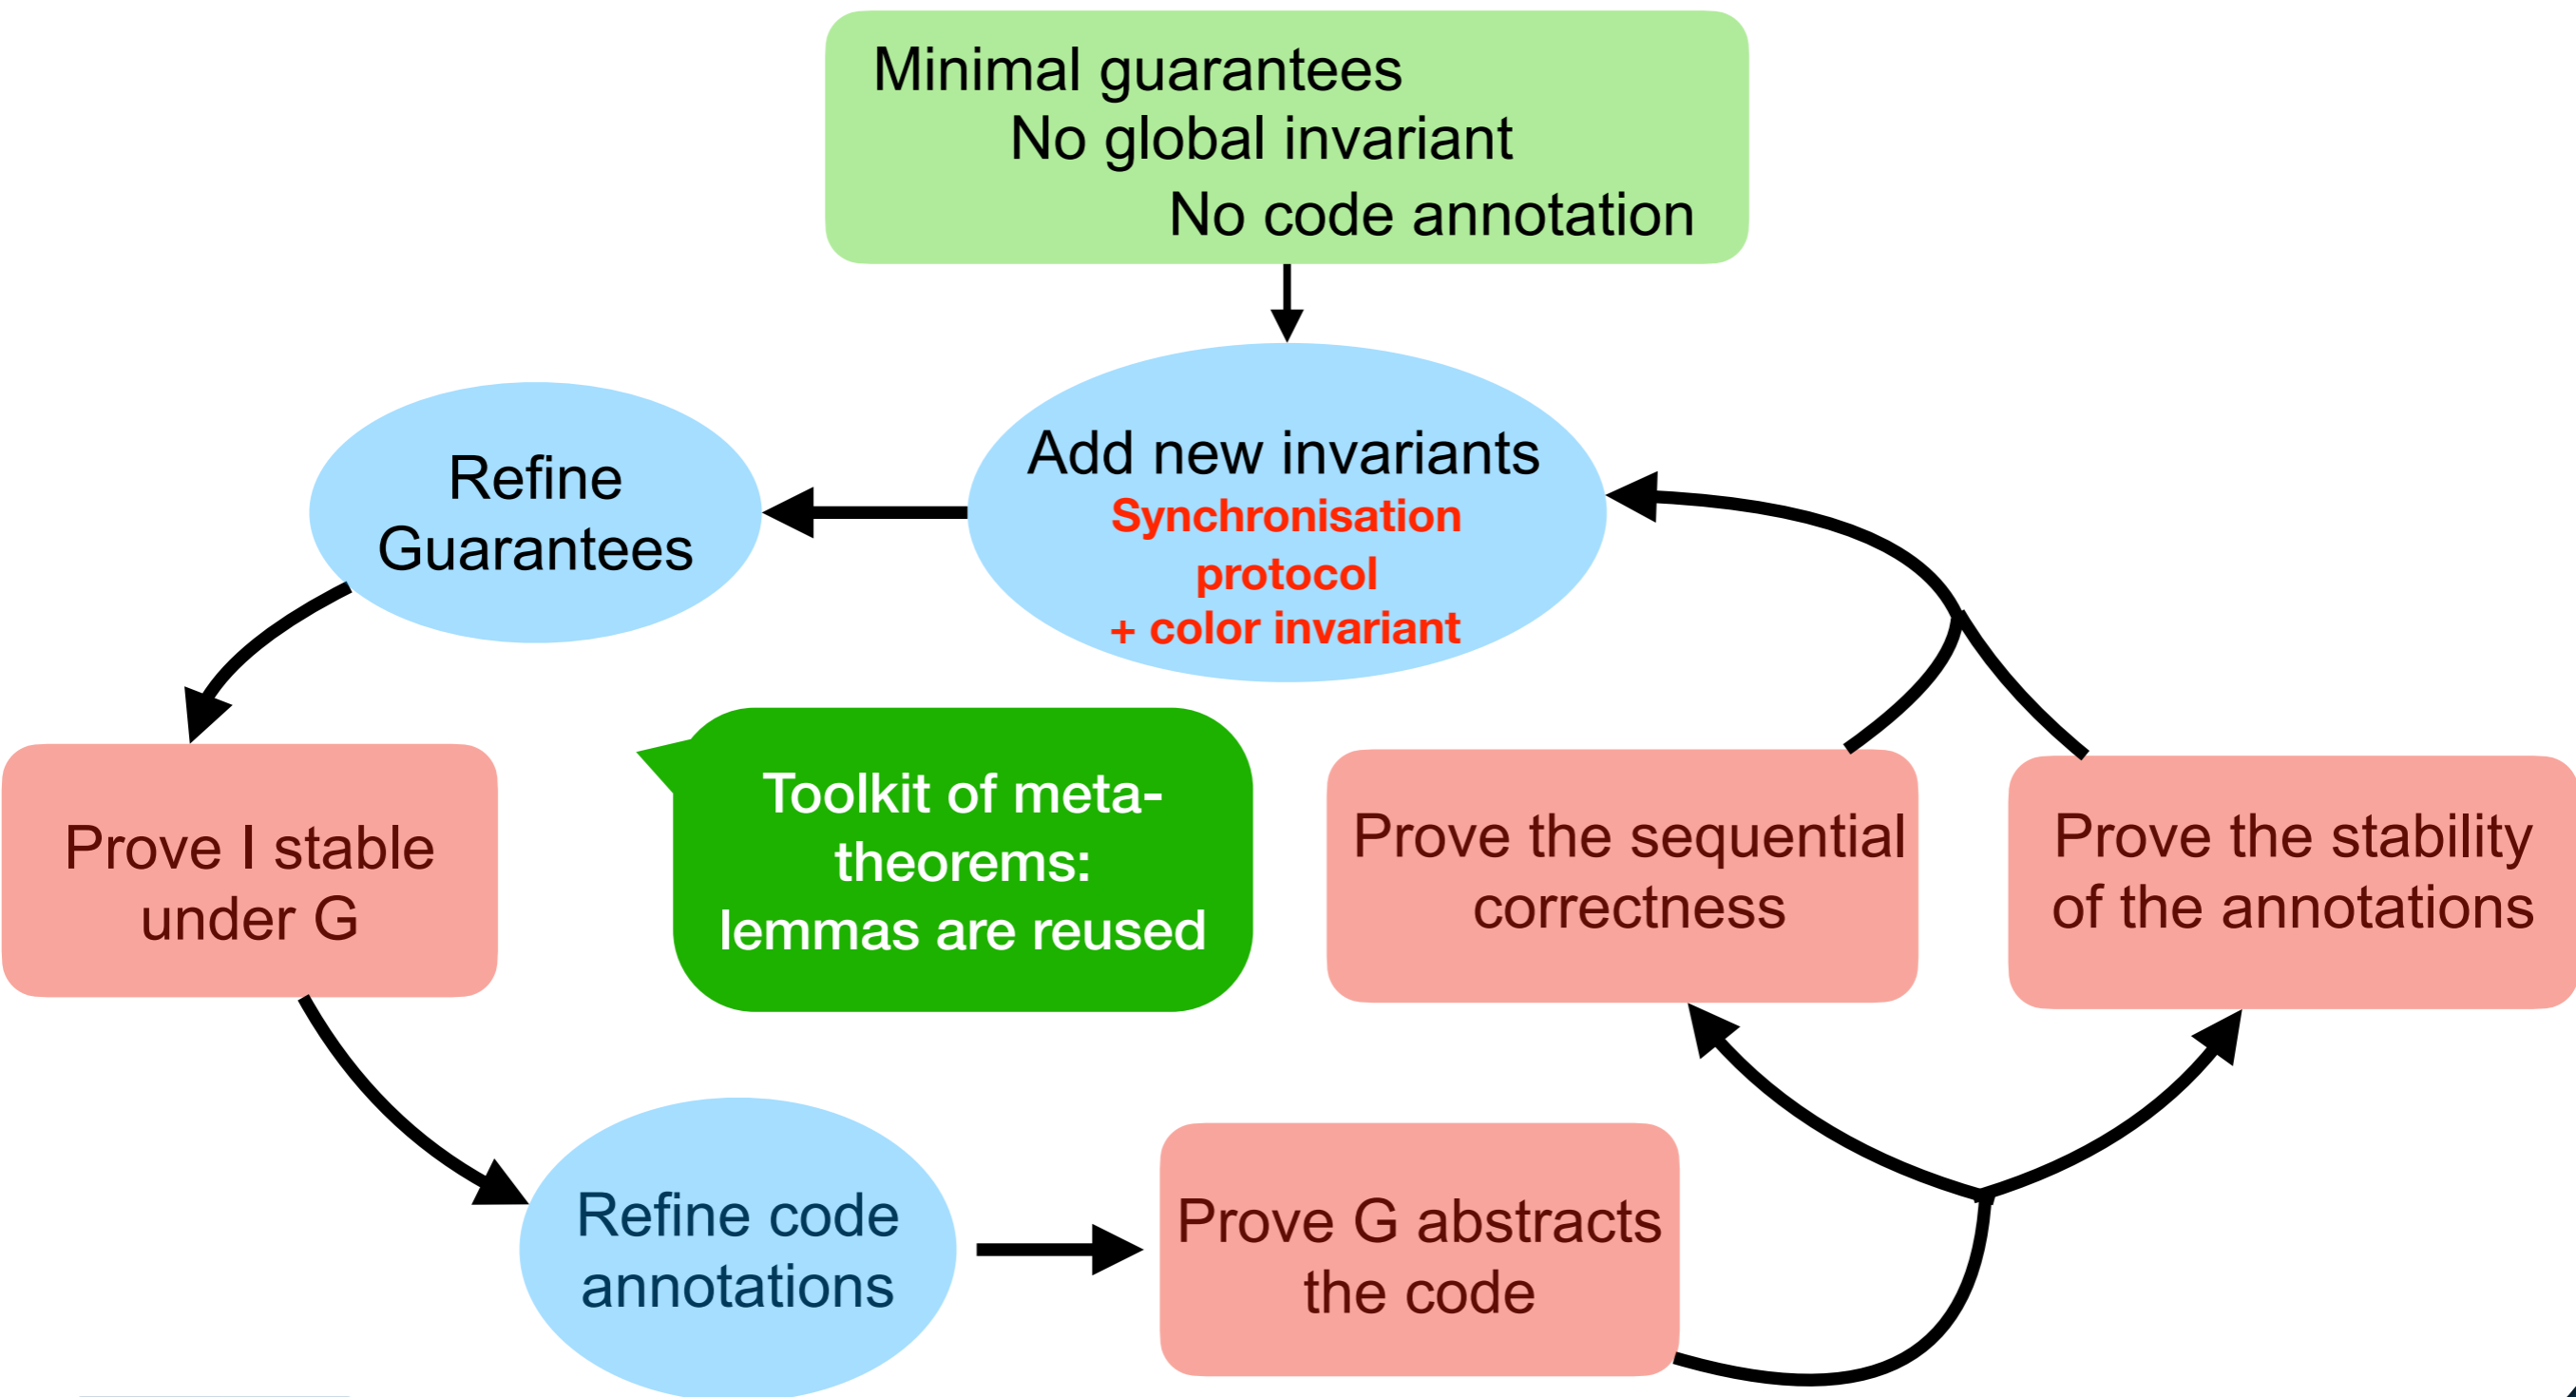
# Reusing proofs incrementally



# Reusing proofs incrementally



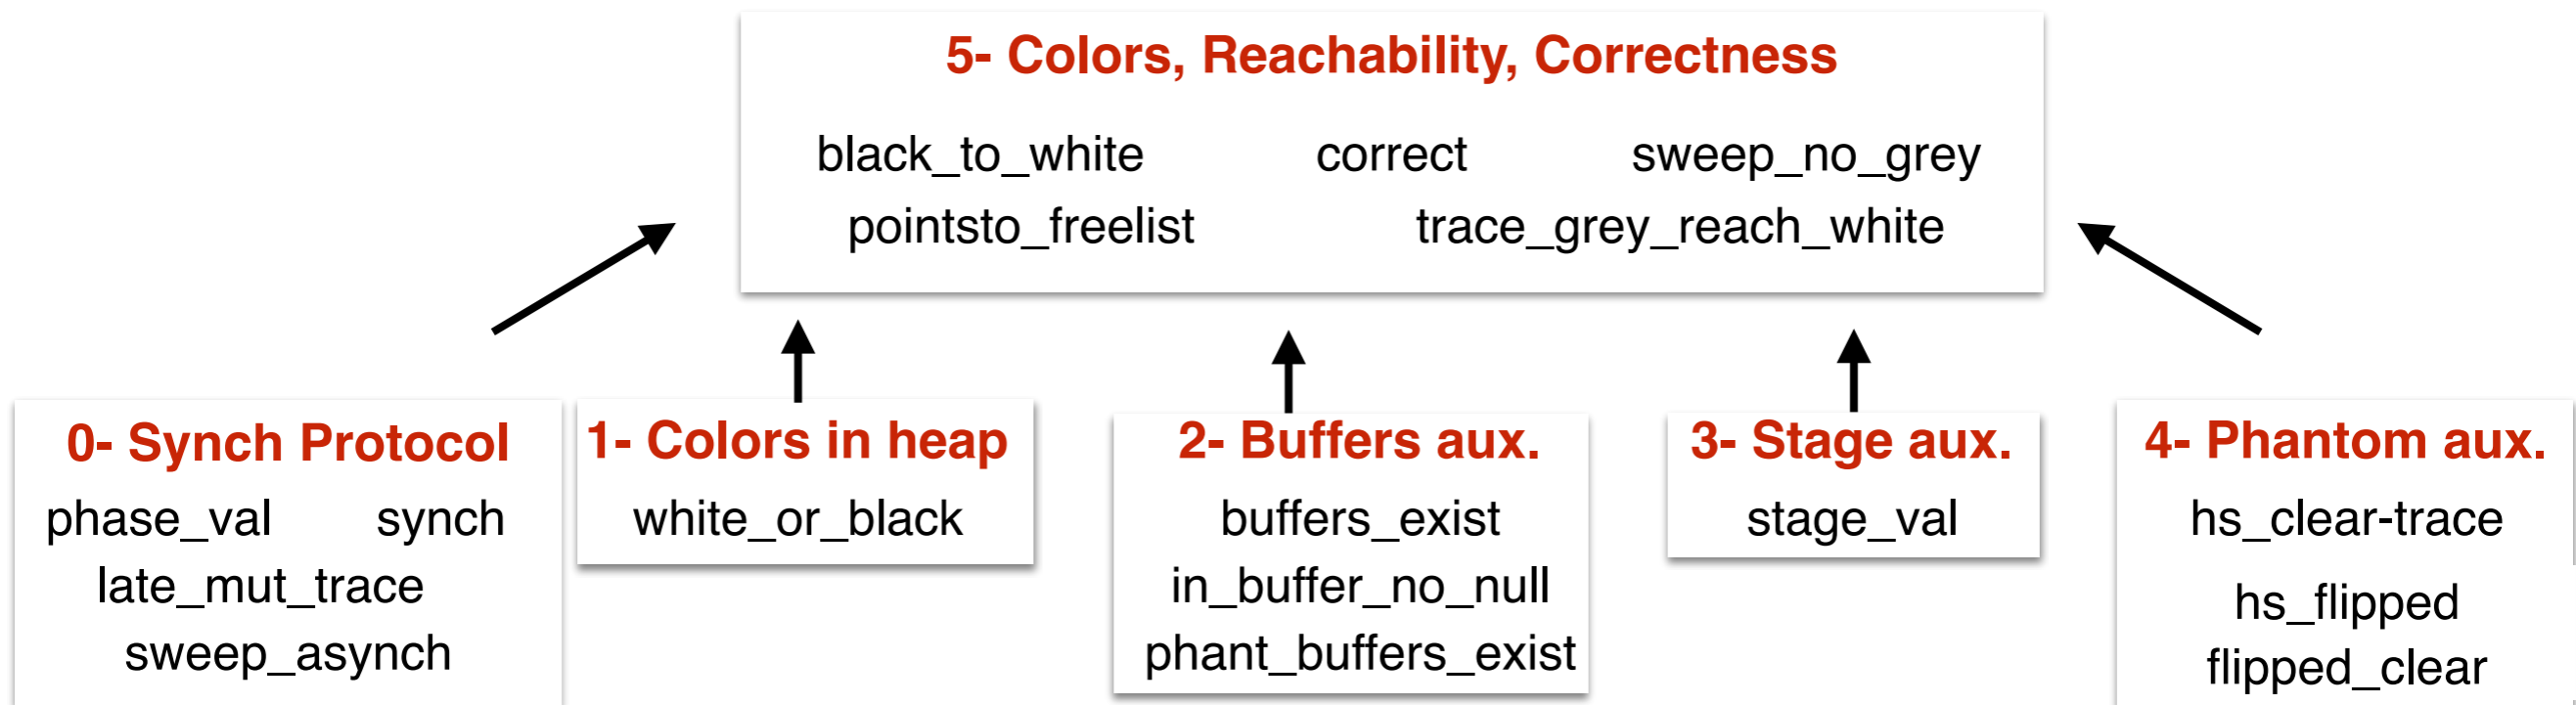
# Reusing proofs incrementally





# Incremental invariants of the GC

Proof = six layers of invariants (partially ordered)



# Conclusion

# Summary



- **Dedicated IR: right level of abstraction**
  - **Abstract concurrent queues**
  - **Native support of roots, objects, freelist**
  - **Iterators**
- **RG logics + soundness theorem + incremental workflow**
- **Realistic On-The-Fly garbage collector**
  - **Significant subset of Domani et al. GC**
  - **Proofs conducted w.r.t. code's semantics**
  - **Most-General-Client theorem**
  - **Proof: incremental invariants**

# Perspectives

- Dedicated IR: right level of abstraction
  - **Abstract concurrent queues**
  - Native support of roots, objects, freelist
  - Iterators
- RG logics + soundness theorem + incremental workflow
- Realistic On-The-Fly garbage collector
  - **Significant subset of Domani et al. GC**
  - Proofs conducted w.r.t. code's semantics
  - Most-General-Client theorem
  - Proof: incremental invariants

Proving  
atomic refinement  
of linearisable  
data-structures  
—> ongoing

We left out  
orthogonal optimisations  
ex : Generational GC