


Will It Fit?



Alexandre Moine

Arthur Charguéraud and François Pottier

Utrecht, April 25, 2024

Key Idea

Space as a resource.

Following Hofmann [1999], let $\diamond 1$ represent one space credit.

Key Idea

Space as a resource.

Following Hofmann [1999], let $\diamond 1$ represent one **space credit**.

- A space credit represents a **permission to allocate** one memory word.
- Space credits are **non-negative** numbers.
- Space credits can be **split and joined**: $\diamond(n_1 + n_2) \equiv \diamond n_1 * \diamond n_2$

Without garbage collection, one can use the following two reasoning rules:

- “alloc” consumes space:

$$\{ \diamond n \} \text{ alloc } n \{ \lambda l. l \mapsto [0, \dots, 0] \}$$

- “free” produces space:

$$\{ l \mapsto [v_1, \dots, v_n] \} \text{ free } l \{ \lambda(). \diamond n \}$$

Without garbage collection, one can use the following two reasoning rules:

- “alloc” consumes space:

$$\{ \diamond n \} \text{ alloc } n \{ \lambda l. l \mapsto [0, \dots, 0] \}$$

- “free” produces space:

$$\{ l \mapsto [v_1, \dots, v_n] \} \text{ free } l \{ \lambda(). \diamond n \}$$

and one can easily prove the following:

Adequacy theorem

If $\{ \diamond S \} t \{ \lambda_. \text{True} \}$ holds, then under heap limit S , the program t will not crash.

The GC Enters the Scene

- OCaml, Haskell, Scala, ... have **garbage collection**.
- There is no “free” instruction.
- At arbitrary times, the garbage collector can reclaim **unreachable** blocks.

The GC Enters the Scene

- OCaml, Haskell, Scala, ... have **garbage collection**.
- There is no “free” instruction.
- At arbitrary times, the garbage collector can reclaim **unreachable** blocks.
- This makes programming easier...

The GC Enters the Scene

- OCaml, Haskell, Scala, ... have **garbage collection**.
- There is no “free” instruction.
- At arbitrary times, the garbage collector can reclaim **unreachable** blocks.
- This makes programming easier...
- ...but makes **reasoning** more difficult. Where can space credits be recovered?

The GC Enters the Scene

- OCaml, Haskell, Scala, ... have **garbage collection**.
- There is no “free” instruction.
- At arbitrary times, the garbage collector can reclaim **unreachable** blocks.
- This makes programming easier...
- ...but makes **reasoning** more difficult. Where can space credits be recovered?

In this work:

- we answer this question,
- while supporting fine-grained **shared-memory concurrency**.

Free is a Ghost Update

Following [Madiot and Pottier \[2022\]](#), we make deallocation a **logical** operation:

$$\text{"}\ell \text{ is unreachable" } * \ell \mapsto [v_1, \dots, v_n] \Rightarrow \diamond n$$

If a block is unreachable then by giving up its ownership one obtains space credits.

$\diamond n$ means that n words are free **or can be freed** by the GC, once it runs.

Free is a Ghost Update

Following [Madiot and Pottier \[2022\]](#), we make deallocation a **logical** operation:

$$\text{"}l \text{ is unreachable" } * l \mapsto [v_1, \dots, v_n] \Rightarrow \diamond n$$

If a block is unreachable then by giving up its ownership one obtains space credits.

$\diamond n$ means that n words are free **or can be freed** by the GC, once it runs.

How does one prove that a location is unreachable?

A location ℓ is reachable if:

1. from some **root**,
 - a location that is currently stored in a **live** local variable x of some thread π
2. there is a **path**, through the heap, to ℓ .

A location ℓ is reachable if:

1. from some **root**,
 - a location that is currently stored in a **live** local variable x of some thread π
2. there is a **path**, through the heap, to ℓ .

Therefore, a location ℓ is unreachable if:

1. ℓ is not a root of any thread; and
2. ℓ has no (reachable) predecessor in the heap

A location ℓ is reachable if:

1. from some **root**,
 - a location that is currently stored in a **live** local variable x of some thread π
2. there is a **path**, through the heap, to ℓ .

Therefore, a location ℓ is unreachable if:

1. ℓ is not a root of any thread; and \rightsquigarrow **the pointed-by-thread assertion**
2. ℓ has no (reachable) predecessor in the heap \rightsquigarrow **the pointed-by-heap assertion**

From [Kassios and Kritikos \[2013\]](#), [Madiot and Pottier \[2022\]](#), [Moine et al. \[2023\]](#).

- $\ell \leftarrow_1 A$ asserts that A is an over-approximation of the (reachable) predecessors of ℓ .

From [Kassios and Kritikos \[2013\]](#), [Madiot and Pottier \[2022\]](#), [Moine et al. \[2023\]](#).

- $\ell \leftarrow_1 A$ asserts that A is an over-approximation of the (reachable) predecessors of ℓ .
- $\ell \leftarrow_1 \emptyset$ means that ℓ has no (reachable) predecessors.
- Pointed-by-heap assertions can be split and joined:

$$\ell \leftarrow_1 \{+l_1; +l_2\} \quad \equiv \quad \ell \leftarrow_{\frac{1}{2}} \{+l_1\} * \ell \leftarrow_{\frac{1}{2}} \{+l_2\}$$

- For greater comfort, we use signed multisets (not shown here).

What Are the Roots Considered by Real-Life GCs?

What Are the Roots Considered by Real-Life GCs?

The Free Variable Rule (FVR) [[Felleisen and Hieb, 1992](#)]

In a substitution-based semantics,
the roots are the locations that occur in the terms that remain to be evaluated.

What Are the Roots Considered by Real-Life GCs?

The Free Variable Rule (FVR) [Felleisen and Hieb, 1992]

In a substitution-based semantics,
the roots are the locations that occur in the terms that remain to be evaluated.

Term	Reachable Heap
$\text{let } a = (\text{ref } 4) \text{ in let } b = (\text{ref } 2) \text{ in } !a + !b$	\emptyset

What Are the Roots Considered by Real-Life GCs?

The Free Variable Rule (FVR) [Felleisen and Hieb, 1992]

In a substitution-based semantics,
the roots are the locations that occur in the terms that remain to be evaluated.

Term	Reachable Heap
$\text{let } a = (\text{ref } 4) \text{ in let } b = (\text{ref } 2) \text{ in } !a + !b$	\emptyset
$\longrightarrow \text{let } a = \ell_a \quad \text{in let } b = (\text{ref } 2) \text{ in } !a + !b$	$\{\ell_a \mapsto 4\}$

What Are the Roots Considered by Real-Life GCs?

The Free Variable Rule (FVR) [Felleisen and Hieb, 1992]

In a substitution-based semantics,
the roots are the locations that occur in the terms that remain to be evaluated.

Term	Reachable Heap
$\text{let } a = (\text{ref } 4) \text{ in let } b = (\text{ref } 2) \text{ in } !a + !b$	\emptyset
$\longrightarrow \text{let } a = \ell_a \quad \text{in let } b = (\text{ref } 2) \text{ in } !a + !b$	$\{\ell_a \mapsto 4\}$
$\longrightarrow \quad \quad \quad \text{let } b = (\text{ref } 2) \text{ in } !\ell_a + !b$	$\{\ell_a \mapsto 4\}$

What Are the Roots Considered by Real-Life GCs?

The Free Variable Rule (FVR) [Felleisen and Hieb, 1992]

In a substitution-based semantics,
the roots are the locations that occur in the terms that remain to be evaluated.

Term	Reachable Heap
$\text{let } a = (\text{ref } 4) \text{ in let } b = (\text{ref } 2) \text{ in } !a + !b$	\emptyset
$\longrightarrow \text{let } a = \ell_a \quad \text{in let } b = (\text{ref } 2) \text{ in } !a + !b$	$\{\ell_a \mapsto 4\}$
$\longrightarrow \quad \quad \quad \text{let } b = (\text{ref } 2) \text{ in } !\ell_a + !b$	$\{\ell_a \mapsto 4\}$
$\longrightarrow \quad \quad \quad \text{let } b = \ell_b \quad \text{in } !\ell_a + !b$	$\{\ell_a \mapsto 4, \ell_b \mapsto 2\}$

What Are the Roots Considered by Real-Life GCs?

The Free Variable Rule (FVR) [Felleisen and Hieb, 1992]

In a substitution-based semantics,
the roots are the locations that occur in the terms that remain to be evaluated.

Term	Reachable Heap
$\text{let } a = (\text{ref } 4) \text{ in let } b = (\text{ref } 2) \text{ in } !a + !b$	\emptyset
$\longrightarrow \text{let } a = \ell_a \quad \text{in let } b = (\text{ref } 2) \text{ in } !a + !b$	$\{\ell_a \mapsto 4\}$
$\longrightarrow \quad \quad \quad \text{let } b = (\text{ref } 2) \text{ in } !\ell_a + !b$	$\{\ell_a \mapsto 4\}$
$\longrightarrow \quad \quad \quad \text{let } b = \ell_b \quad \text{in } !\ell_a + !b$	$\{\ell_a \mapsto 4, \ell_b \mapsto 2\}$
$\longrightarrow \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad !\ell_a + !\ell_b$	$\{\ell_a \mapsto 4, \ell_b \mapsto 2\}$

What Are the Roots Considered by Real-Life GCs?

The Free Variable Rule (FVR) [Felleisen and Hieb, 1992]

In a substitution-based semantics,
the roots are the locations that occur in the terms that remain to be evaluated.

Term	Reachable Heap
$\text{let } a = (\text{ref } 4) \text{ in let } b = (\text{ref } 2) \text{ in } !a + !b$	\emptyset
$\longrightarrow \text{let } a = l_a \quad \text{in let } b = (\text{ref } 2) \text{ in } !a + !b$	$\{l_a \mapsto 4\}$
$\longrightarrow \quad \quad \quad \text{let } b = (\text{ref } 2) \text{ in } !l_a + !b$	$\{l_a \mapsto 4\}$
$\longrightarrow \quad \quad \quad \text{let } b = l_b \quad \text{in } !l_a + !b$	$\{l_a \mapsto 4, l_b \mapsto 2\}$
$\longrightarrow \quad \quad \quad \quad \quad \quad \quad \quad !l_a + !l_b$	$\{l_a \mapsto 4, l_b \mapsto 2\}$
$\longrightarrow \quad \quad \quad \quad \quad \quad \quad \quad 4 + !l_b$	$\{\cancel{l_a \mapsto 4}, l_b \mapsto 2\}$
$\longrightarrow^* \quad \quad \quad \quad \quad \quad \quad \quad 6$	$\{\cancel{l_a \mapsto 4}, \cancel{l_b \mapsto 2}\}$

Pointed-By-Thread Assertions

We use a **ghost thread identifier** π to identify a thread. $\{ \Phi \} \pi : t \{ \Psi \}$

Pointed-By-Thread Assertions

We use a **ghost thread identifier** π to identify a thread. $\{ \Phi \} \pi : t \{ \Psi \}$

We introduce **pointed-by-thread** assertions:

- $\ell \Leftarrow_1 \Pi$ asserts that Π is an over-approximation of the threads in which ℓ is a root.
- $\ell \Leftarrow_1 \emptyset$ asserts that ℓ is not a root in any thread.
- Pointed-by-thread assertions can be split and joined:

$$\ell \Leftarrow_{(p_1+p_2)} (\Pi_1 \cup \Pi_2) \quad \equiv \quad \ell \Leftarrow_{p_1} \Pi_1 * \ell \Leftarrow_{p_2} \Pi_2$$

Some Reasoning Rules: Alloc, Load and Fork

Allocation produces points-to, **pointed-by-heap**, and **pointed-by-thread** assertions:

$$\{ \diamond n \} \pi : \text{alloc } n \{ \lambda l. l \mapsto [0, \dots, 0] * l \leftarrow_1 \emptyset * l \Leftarrow_1 \{ \pi \} \}$$

Some Reasoning Rules: Alloc, Load and Fork

Allocation produces points-to, **pointed-by-heap**, and **pointed-by-thread** assertions:

$$\{ \diamond n \} \pi : \text{alloc } n \{ \lambda \ell. \ell \mapsto [0, \dots, 0] * \ell \leftarrow_1 \emptyset * \ell \Leftarrow_1 \{ \pi \} \}$$

Reading a location ℓ' from the heap makes it a **root** in this thread:

$$\frac{0 \leq i < |\vec{v}| \quad \vec{v}(i) = \ell'}{\{ \ell \mapsto \vec{v} * \ell' \Leftarrow_p \emptyset \} \pi : \ell[i] \{ \lambda w. \ulcorner w = \ell' \urcorner * \ell \mapsto \vec{v} * \ell' \Leftarrow_p \{ \pi \} \}}$$

Some Reasoning Rules: Alloc, Load and Fork

Allocation produces points-to, **pointed-by-heap**, and **pointed-by-thread** assertions:

$$\{ \Diamond n \} \pi : \text{alloc } n \{ \lambda \ell. \ell \mapsto [0, \dots, 0] * \ell \leftarrow_1 \emptyset * \ell \Leftarrow_1 \{ \pi \} \}$$

Reading a location ℓ' from the heap makes it a **root** in this thread:

$$\frac{0 \leq i < |\vec{v}| \quad \vec{v}(i) = \ell'}{\{ \ell \mapsto \vec{v} * \ell' \Leftarrow_p \emptyset \} \pi : \ell[i] \{ \lambda w. \ulcorner w = \ell' \urcorner * \ell \mapsto \vec{v} * \ell' \Leftarrow_p \{ \pi \} \}}$$

Spawning a new thread transfers some **roots** to it. Here is a (simplified) rule:

$$\frac{\forall \pi'. \{ \Phi \} \pi' : t \{ \lambda(). \text{True} \}}{\{ \Phi \} \pi : \text{fork } t \{ \lambda(). \text{True} \}}$$

Some Reasoning Rules: Alloc, Load and Fork

Allocation produces points-to, **pointed-by-heap**, and **pointed-by-thread** assertions:

$$\{ \Diamond n \} \pi: \text{alloc } n \{ \lambda \ell. \ell \mapsto [0, \dots, 0] * \ell \leftarrow_1 \emptyset * \ell \Leftarrow_1 \{ \pi \} \}$$

Reading a location ℓ' from the heap makes it a **root** in this thread:

$$\frac{0 \leq i < |\vec{v}| \quad \vec{v}(i) = \ell'}{\{ \ell \mapsto \vec{v} * \ell' \Leftarrow_p \emptyset \} \pi: \ell[i] \{ \lambda w. \ulcorner w = \ell' \urcorner * \ell \mapsto \vec{v} * \ell' \Leftarrow_p \{ \pi \} \}}$$

Spawning a new thread transfers some **roots** to it. Here is a (simplified) rule:

$$\frac{\text{locs}(t) = \{ \ell \} \quad \forall \pi'. \{ \Phi \} \pi': t \{ \lambda(). \text{True} \}}{\{ \Phi \} \pi: \text{fork } t \{ \lambda(). \text{True} \}}$$

Some Reasoning Rules: Alloc, Load and Fork

Allocation produces points-to, **pointed-by-heap**, and **pointed-by-thread** assertions:

$$\{ \Diamond n \} \pi: \text{alloc } n \{ \lambda l. l \mapsto [0, \dots, 0] * l \leftarrow_1 \emptyset * l \Leftarrow_1 \{ \pi \} \}$$

Reading a location ℓ' from the heap makes it a **root** in this thread:

$$\frac{0 \leq i < |\vec{v}| \quad \vec{v}(i) = \ell'}{\{ l \mapsto \vec{v} * l' \Leftarrow_p \emptyset \} \pi: l[i] \{ \lambda w. \ulcorner w = \ell' \urcorner * l \mapsto \vec{v} * l' \Leftarrow_p \{ \pi \} \}}$$

Spawning a new thread transfers some **roots** to it. Here is a (simplified) rule:

$$\frac{\text{locs}(t) = \{ \ell \} \quad \forall \pi'. \{ l \Leftarrow_p \{ \pi' \} * \Phi \} \pi': t \{ \lambda(). \text{True} \}}{\{ l \Leftarrow_p \{ \pi \} * \Phi \} \pi: \text{fork } t \{ \lambda(). \text{True} \}}$$

Some Reasoning Rules: Trimming and Logical Deallocation

Once l is no longer a root in thread π , pointed-by-thread assertions can be trimmed:

$$\frac{l \notin \text{locs}(t) \quad \{ l \Leftarrow_p \emptyset * \Phi \} \pi : t \{ \Psi \}}{\{ l \Leftarrow_p \{ \pi \} * \Phi \} \pi : t \{ \Psi \}} \text{TRIM}$$

Some Reasoning Rules: Trimming and Logical Deallocation

Once l is no longer a root in thread π , pointed-by-thread assertions can be **trimmed**:

$$\frac{l \notin \text{locs}(t) \quad \{ l \Leftarrow_p \emptyset * \Phi \} \pi : t \{ \Psi \}}{\{ l \Leftarrow_p \{ \pi \} * \Phi \} \pi : t \{ \Psi \}} \text{TRIM}$$

The logical deallocation rule can now be shown:

$$\begin{array}{c}
 \begin{array}{c}
 \text{\textit{l} is not a root} \\
 \downarrow \\
 \boxed{l \Leftarrow_1 \emptyset}
 \end{array} \\
 l \mapsto [v_1, \dots, v_n] * \boxed{l \Leftarrow_1 \emptyset} * \boxed{l \Leftarrow_1 \emptyset} \Rightarrow l \mapsto [v_1, \dots, v_n] * \diamond n * \boxed{\dagger l} \\
 \begin{array}{c}
 \text{\textit{l} has no predecessors in the heap} \uparrow \\
 \text{\textit{deallocation witness}} \uparrow
 \end{array}
 \end{array}$$

The points-to assertion is **not consumed**.

The Adequacy Theorem

A configuration is **stuck** if

- even after garbage collection has taken place,
- some thread has not terminated and cannot make progress.

The Adequacy Theorem

A configuration is **stuck** if

- even after garbage collection has taken place,
- some thread has not terminated and cannot make progress.

In particular, a **memory allocation request that cannot be satisfied** is stuck.

The Adequacy Theorem

A configuration is **stuck** if

- even after garbage collection has taken place,
- some thread has not terminated and cannot make progress.



In particular, a **memory allocation request that cannot be satisfied** is stuck.

Adequacy Theorem

If $\{ \diamond S \} \pi : t \{ \lambda_{-}. \text{True} \}$ holds, then under heap limit S , t cannot become stuck.



The Case of Lock-Free Data Structures: Treiber's Stack

A linearizable lock-free stack, implemented as a reference on an immutable list.




The Case of Lock-Free Data Structures: Treiber's Stack

A linearizable lock-free stack, implemented as a reference on an immutable list.




```
let rec pop s =  
  let xs = !s in  
  match xs with  
  | nil -> assert false  
  | y::ys -> if CAS s xs ys then y else pop s
```

A Problematic Interleaving for pop



A Problematic Interleaving for pop




pop s

```
pop s;  
pop s;  
pop s;  
...
```



A Problematic Interleaving for pop




```
let xs = !s in  
match xs with  
...
```

```
pop s;  
pop s;  
pop s;  
...
```



A Problematic Interleaving for pop




```
let xs = c in  
match xs with  
...
```

```
pop s;  
pop s;  
pop s;  
...
```



A Problematic Interleaving for pop




```
let xs = c in  
match xs with  
...
```



```
pop s;  
pop s;  
pop s;  
...
```

A Problematic Interleaving for pop




```
let xs = c in  
match xs with  
...
```



```
pop s;  
pop s;  
pop s;  
...
```

A Problematic Interleaving for pop




```
let xs = c in  
match xs with  
...
```



```
pop s;  
pop s;  
pop s;  
...
```

A Problematic Interleaving for pop



```
let xs = c in  
match xs with  
...
```



```
pop s;  
pop s;  
pop s;  
...
```

- The sleeping thread holds the root `c`, therefore **keeps the whole list reachable**.
- We **cannot prove** that `pop` produces space credits!
- **Key idea:** `c` is a **temporary root**. Do not allow the GC to run while it is held.

Our calculus, LambdaFit, has several novel features:

- **protected sections**, where the GC is disabled.
 - the GC can run only if **all** threads are outside protected sections
 - inside protected sections, allocation and divergence are forbidden
- **polling points**, which cannot be passed if any thread is waiting for the GC to run;
- there is a heap size limit S ,
- and a memory allocation request **blocks** until it can be satisfied.

Temporary Roots

Inside a critical section, a root can be **temporary** – unobservable by the GC.

Temporary Roots

Inside a critical section, a root can be **temporary** – unobservable by the GC.

Our reasoning rules take advantage of protected sections:

Key Idea

In the logical deallocation rule, ignore temporary roots.

Treiber's Stack, Fixed: pop

```
let rec pop s =  
  enter ();  
  let xs = !s in  
  match xs with  
  | nil -> assert false  
  | y::ys ->  
    if CAS s xs ys then (exit (); y) else (exit (); pop s)
```

- The variable `xs` is a root **inside a protected section** only.
- The GC **can never observe** this root.
- The problematic scenario disappears.

Inside and Outside Assertions

- Two new assertions: outside π and inside πT .
- The parameter T is the set of **temporary roots**.

$\{ \text{outside } \pi \} \pi: \text{enter} () \{ \lambda(). \text{inside } \pi \emptyset \}$ $\{ \text{inside } \pi \emptyset \} \pi: \text{exit} () \{ \lambda(). \text{outside } \pi \}$

Inside and Outside Assertions

- Two new assertions: outside π and inside πT .
- The parameter T is the set of **temporary roots**.

$$\{ \text{outside } \pi \} \pi: \text{enter} () \{ \lambda(). \text{inside } \pi \emptyset \} \quad \{ \text{inside } \pi \emptyset \} \pi: \text{exit} () \{ \lambda(). \text{outside } \pi \}$$

The assertion inside πT forms an **escape hatch** to the pointed-by-thread discipline.

$$\begin{array}{lll} \text{inside } \pi T * l \Leftarrow_p \{ \pi \} & \Rightarrow & \text{inside } \pi (T \cup \{ l \}) * l \Leftarrow_p \emptyset & \text{ADDTEMPORARY} \\ \text{inside } \pi T * l \Leftarrow_p \emptyset & \Rightarrow & \text{inside } \pi (T \setminus \{ l \}) * l \Leftarrow_p \{ \pi \} & \text{REMTEMPORARY} \end{array}$$

$$\frac{\{ \text{inside } \pi (T \cap \text{locs}(t)) * \Phi \} \pi : t \{ \Psi \}}{\{ \text{inside } \pi T * \Phi \} \pi : t \{ \Psi \}} \text{TRIMINSIDE}$$

$$\frac{\{ \text{inside } \pi (T \cap \text{locs}(t)) * \Phi \} \pi : t \{ \Psi \}}{\{ \text{inside } \pi T * \Phi \} \pi : t \{ \Psi \}} \text{TRIMINSIDE}$$

- All of the previous rules remain sound, including **logical deallocation**.
- Therefore **logical deallocation** is not impeded by temporary roots.

Treiber's Stack, Fixed: push

```
1  let rec push s v =  
2    let xs = new_cell () in  
3    set_data xs v;  
4    enter ();  
5    let ys = !s in  
6    set_tail xs ys;  
7    if compare_and_swap s ys xs  
8    then exit ()  
9    else (exit (); push s v)
```


Treiber's Stack, Fixed: push

```
1  let rec push s v =  
2    let xs = new_cell () in  
3    set_data xs v;  
4    enter ();  
5    let ys = !s in  
6    set_tail xs ys;  
7    if compare_and_swap s ys xs  
8    then exit ()  
9    else (exit (); push s v)
```

- What if this thread falls asleep at the start of line 6 or 7?

Treiber's Stack, Fixed: push

```
1  let rec push s v =  
2    let xs = new_cell () in  
3    set_data xs v;  
4    enter ();  
5    let ys = !s in  
6    set_tail xs ys;  
7    if compare_and_swap s ys xs  
8    then exit ()  
9    else (exit (); push s v)
```

- What if this thread falls asleep at the start of line 6 or 7?
- Some other thread may successfully pop ys.

Treiber's Stack, Fixed: push

```
1  let rec push s v =  
2    let xs = new_cell () in  
3    set_data xs v;  
4    enter ();  
5    let ys = !s in  
6    set_tail xs ys;  
7    if compare_and_swap s ys xs  
8    then exit ()  
9    else (exit (); push s v)
```

- What if this thread falls asleep at the start of line 6 or 7?
- Some other thread may successfully pop `ys`.
- That thread **must** logically deallocate `ys`,

Treiber's Stack, Fixed: push

```
1  let rec push s v =  
2    let xs = new_cell () in  
3    set_data xs v;  
4    enter ();  
5    let ys = !s in  
6    set_tail xs ys;  
7    if compare_and_swap s ys xs  
8    then exit ()  
9    else (exit (); push s v)
```

- What if this thread falls asleep at the start of line 6 or 7?
- Some other thread may successfully pop `ys`.
- That thread **must** logically deallocate `ys`,
- and to do so, **must also** logically deallocate `xs`.

Treiber's Stack, Fixed: push

```
1  let rec push s v =  
2    let xs = new_cell () in  
3    set_data xs v;  
4    enter ();  
5    let ys = !s in  
6    set_tail xs ys;  
7    if compare_and_swap s ys xs  
8    then exit ()  
9    else (exit (); push s v)
```

- What if this thread falls asleep at the start of line 6 or 7?
- Some other thread may successfully pop *ys*.
- That thread **must** logically deallocate *ys*,
- and to do so, **must also** logically deallocate *xs*.
- This is possible, as *ys* is a **temporary** root.

Treiber's Stack, Fixed: push

```
1 let rec push s v =  
2   let xs = new_cell () in  
3   set_data xs v;  
4   enter ();  
5   let ys = !s in  
6   set_tail xs ys;  
7   if compare_and_swap s ys xs  
8   then exit ()  
9   else (exit (); push s v)
```

- What if this thread falls asleep at the start of line 6 or 7?
- Some other thread may successfully pop *ys*.
- That thread **must** logically deallocate *ys*,
- and to do so, **must also** logically deallocate *xs*.
- This is possible, as *ys* is a **temporary** root.
- Once this thread wakes up, lines 6 or 7 **access a logically deallocated block!**
- This is possible, as logical deallocation does not consume the points-to assertion.

Specification of Treiber's Stack

$$\frac{\{ \diamond 2 \}}{\langle \forall L. \text{stack } s \ L \rangle}$$
$$\pi: \text{push } s \ v$$
$$\langle \text{stack } s \ (v :: L) \rangle$$
$$\frac{}{\{ \lambda(). \text{True} \}}$$

Specification of Treiber's Stack

$$\begin{array}{c} \text{Private precondition} \longrightarrow \{ \diamond 2 \} \\ \hline \langle \forall L. \text{stack } s \ L \rangle \\ \pi: \text{push } s \ v \\ \langle \text{stack } s \ (v :: L) \rangle \\ \hline \{ \lambda(). \text{True} \} \end{array}$$


Specification of Treiber's Stack

$$\begin{array}{c} \text{Private precondition} \longrightarrow \{ \diamond 2 \} \\ \hline \langle \forall L. \text{stack } s \ L \rangle \longleftarrow \text{Public precondition} \\ \pi: \text{push } s \ v \\ \langle \text{stack } s \ (v :: L) \rangle \\ \hline \{ \lambda(). \text{True} \} \end{array}$$

Specification of Treiber's Stack

$$\begin{array}{c} \text{Private precondition} \longrightarrow \{ \diamond 2 \} \\ \hline \langle \forall L. \text{stack } s \ L \rangle \longleftarrow \text{Public precondition} \\ \pi: \text{push } s \ v \\ \langle \text{stack } s \ (v :: L) \rangle \longleftarrow \text{Public postcondition} \\ \hline \{ \lambda(). \text{True} \} \end{array}$$

Specification of Treiber's Stack



Specification of Treiber's Stack

$$\begin{array}{l} \text{Private precondition} \longrightarrow \{ \diamond 2 \} \\ \hline \langle \forall L. \text{stack } s \ L \rangle \longleftarrow \text{Public precondition} \\ \pi: \text{push } s \ v \\ \langle \text{stack } s \ (v :: L) \rangle \longleftarrow \text{Public postcondition} \\ \hline \text{Private postcondition} \longrightarrow \{ \lambda(). \text{True} \} \end{array}$$

$$\frac{\{ \text{True} \}}{\langle \forall v \ L. \text{stack } s \ (v :: L) \rangle} \frac{\pi: \text{pop } s \quad \langle \text{stack } s \ L \rangle}{\{ \lambda v'. \ulcorner v' = v \urcorner * \diamond 2 \}}$$

Polling Point Insertion

Polling points have no effect on reasoning.

$$\{ \text{outside } \pi \} \pi : \text{poll } () \{ \lambda(). \text{outside } \pi \}$$

They can be **automatically inserted** by the compiler outside of protected sections.

If enough polling points are inserted,

then at all times,

every thread must reach a polling point in bounded time.

Safety and Liveness Theorems

Under the assumption that the program is verified:

$$\{ \diamond S * \text{outside } \pi \} \pi : t \{ \lambda _ . \text{outside } \pi \}$$

Safety and Liveness Theorems

Under the assumption that the program is verified:

$$\{ \diamond S * \text{outside } \pi \} \pi : t \{ \lambda _ . \text{outside } \pi \}$$

Then, at all times,

Safety

Every thread has terminated, is blocked on allocation, or can make progress.

Safety and Liveness Theorems

Under the assumption that the program is verified:

$$\{ \diamond S * \text{outside } \pi \} \pi : t \{ \lambda _ . \text{outside } \pi \}$$

Then, at all times,

Safety

Every thread has terminated, is blocked on allocation, or can make progress.

and assuming that enough polling points have been inserted, at all times,

Liveness

Every blocked thread is eventually unblocked.



Under review (journal paper):

Will it Fit? Verifying Heap Space Bounds of Concurrent Programs under Garbage Collection with Separation Logic

See also Alexandre's dissertation (upcoming).



Thank you for your attention!

alexandre.moine [at] inria.fr

<https://cambium.inria.fr/~amoine/>

- Matthias Felleisen and Robert Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103(2):235–271, 1992. URL <https://www2.ccs.neu.edu/racket/pubs/tcs92-fh.pdf>.
- Martin Hofmann. Linear types and non-size-increasing polynomial time computation. In *Logic in Computer Science (LICS)*, pages 464–473, July 1999. URL <https://doi.org/10.1109/LICS.1999.782641>.
- Ioannis T. Kassios and Eleftherios Kritikos. A discipline for program verification based on backpointers and its use in observational disjointness. In *European Symposium on Programming (ESOP)*, volume 7792 of *Lecture Notes in Computer Science*, pages 149–168. Springer, March 2013. URL https://doi.org/10.1007/978-3-642-37036-6_10.

- Jean-Marie Madiot and François Pottier. A separation logic for heap space under garbage collection. *Proceedings of the ACM on Programming Languages*, 6(POPL), January 2022. URL <http://cambium.inria.fr/~fpottier/publis/madiot-pottier-diamonds-2022.pdf>.
- Alexandre Moine, Arthur Charguéraud, and François Pottier. A high-level separation logic for heap space under garbage collection. *Proc. ACM Program. Lang.*, 7(POPL), jan 2023. doi: 10.1145/3571218. URL <https://doi.org/10.1145/3571218>.