

# Diamonds Are Forever: Reasoning about Heap Space in a Concurrent and Garbage Collected Language

---

Alexandre Moine

Arthur Charguéraud

François Pottier

Iris'23



*Inria*

# Formal Verification of Heap Space Bounds

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

$\rightsquigarrow$  the right to allocate one memory word.

Without GC:

- alloc consumes space

$$\{ \diamond \text{size}(b) \} \ell := \text{alloc}(b) \{ \ell \mapsto b \}$$

- free produces space

$$\{ \ell \mapsto b \} \text{free}(\ell) \{ \diamond \text{size}(b) \}$$

With GC: **no syntactical point to recover space credits.**


# Formal Verification of Heap Space Bounds Under Garbage Collection

Space credits can be recovered as soon as a location becomes **unreachable**:  
~> from **the roots** following **heap paths**.

Separation Logic for Heap Space with GC	High-Level	Concurrency
Madiot and Pottier [2022]	X	✓ (no examples)
Moine, Charguéraud and Pottier [2023]	✓	X
This work	✓	✓


## A Motivating Example: The Treiber's Stack

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




What is a space-aware specification for push and pop?

## A Space-Leaking Interleaving for pop



# A Space-Leaking Interleaving for pop




pop s

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



## A Space-Leaking Interleaving for pop




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

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



## A Space-Leaking Interleaving for pop




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

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





## A Space-Leaking Interleaving for pop




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



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

## A Space-Leaking Interleaving for pop




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



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

## A Space-Leaking Interleaving for pop




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



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

## A Space-Leaking Interleaving for pop



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



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

- The sleeping thread maintains reachable a morally dead structure.
- pop cannot produce space credits!

# Contributions

We present the first program logic to reason about

heap space usage for a high-level and concurrent language with GC.

Key contributions:

- A new pointed-by-thread assertion, tracking in which thread a location is a root.
- Examples:
  - Lock-free data structures: Treiber's stack, Michael and Scott's queue
  - Closures: Concurrent counter
- Theory and examples are fully mechanized in Iris.

## Prior Work: Pointed-by-Heap Assertions to Track Heap Predecessors

From Kassios and Kritikos [2013], Moine et al. [2023]:

$$l \leftarrow_q L$$

(signed) multiset

$\in [0, 1]$

- $l \leftarrow_1 L$  asserts that  $L$  is an over-approximation of the reachable predecessors of  $l$ .
- $l \leftarrow_1 \emptyset$  asserts that  $l$  is unreachable from the heap.

$$l \leftarrow_1 \{+l_1; +l_2\} \quad -* \quad l \leftarrow_{\frac{1}{2}} \{+l_1\} * l \leftarrow_{\frac{1}{2}} \{+l_2\}$$
$$l \leftarrow_{\frac{1}{2}} \{+l_1\} * l \leftarrow_0 \{-l_1\} \quad -* \quad l \leftarrow_{\frac{1}{2}} \emptyset$$

Main invariant: if  $l \leftarrow_0 L$  then  $L$  must only contain **negative** elements.

## What Are the Roots?

- We consider an interleaving semantics.
- The GC is interleaved with the per-thread reduction.

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

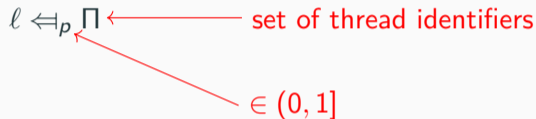
In a substitution-based semantics, the roots of a threadpool consist of the union of the locations occurring in each thread.

## Separation Logic Triples & Pointed-By-Thread Assertions

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

The **pointed-by-thread** assertion:

$$\ell \Leftarrow_p \Pi$$



- $\ell \Leftarrow_1 \Pi$  asserts that  $\Pi$  is an over-approximation of the threads in which  $\ell$  is a root.
- $\ell \Leftarrow_1 \emptyset$  asserts that  $\ell$  is not a root.

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



## Our Logical Deallocation Rule

$$l \mapsto_p \vec{w} * l \leftarrow_1 \emptyset * l \Leftarrow_1 \emptyset \quad \Rightarrow \quad l \mapsto_p \vec{w} * \diamond \text{size}(\vec{w}) * \dagger l$$

- Free does **not** consume the points-to. Useful for:
  - Persistent objects
  - Closing invariants

## Our Logical Deallocation Rule

$$\frac{\text{\underline{\mathit{l} is a valid block}}}{\mathit{l} \mapsto_p \vec{w} * \mathit{l} \leftarrow_1 \emptyset * \mathit{l} \leftarrow_1 \emptyset} \quad \frac{\text{\underline{\mathit{l} is not a root}}}{\mathit{l} \leftarrow_1 \emptyset} \quad \Rightarrow \quad \mathit{l} \mapsto_p \vec{w} * \diamond \text{size}(\vec{w}) * \dagger \mathit{l}$$

$\frac{\text{\underline{\mathit{l} is unreachable from the heap}}}{\phantom{\mathit{l} \leftarrow_1 \emptyset}}$

- Free does **not** consume the points-to. Useful for:
  - Persistent objects
  - Closing invariants

## Our Logical Deallocation Rule

only applicable in a precondition

$$l \mapsto_p \vec{w} * l \leftarrow_1 \emptyset * l \Leftarrow_1 \emptyset \quad \Rightarrow \quad l \mapsto_p \vec{w} * \Diamond \text{size}(\vec{w}) * \dagger l$$

- Free does **not** consume the points-to. Useful for:
  - Persistent objects
  - Closing invariants

# Our Logical Deallocation Rule

$$l \mapsto_p \vec{w} * l \leftarrow_1 \emptyset * l \Leftarrow_1 \emptyset \quad \Rightarrow \quad l \mapsto_p \vec{w} * \underbrace{\diamond \text{size}(\vec{w})}_{\text{space credits}} * \dagger l$$

$\underbrace{\hspace{15em}}_{l \text{ is still a valid block}} \quad \quad \quad \underbrace{\hspace{15em}}_{l \text{ is unreachable}}$

- Free does **not** consume the points-to. Useful for:
  - Persistent objects
  - Closing invariants

# Life of the Pointed-by-Thread Assertion, Part 1

The rule for allocation.

$$\{ \Diamond n \} \pi : (\text{alloc } n) \{ \lambda \ell. \ell \mapsto [\overbrace{(), \dots, ()}^{n \text{ times}}] * \ell \leftarrow \emptyset * \ell \Leftarrow \{ \pi \} \}$$

After a load, a thread points-to the loaded value.

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

## Life of the Pointed-by-Thread Assertion, Part 2

The  $l \Leftarrow_p \{\pi\}$  assertion can be **cleaned** when  $l$  does not appear in the focused part of  $\pi$ .

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

$l \Leftarrow_p \{\pi\}$  is **force-framed** when  $l$  becomes a root of the evaluation context in  $\pi$ .

$$\frac{\text{locs}(t_2) = \{l\} \quad \{ \Phi \} \pi: t_1 \{ \Psi' \} \quad \forall v. \{ l \Leftarrow_p \{\pi\} * \Psi' \ v \} \pi: [v/x]t_2 \{ \Psi \}}{\{ l \Leftarrow_p \{\pi\} * \Phi \} \pi: (\text{let } x = t_1 \text{ in } t_2) \{ \Psi \}}$$

A fork updates pointed-by-thread assertions.

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

# The Soundness Theorem

Our semantics

- is parameterized by a **maximal** heap size  $S$
- **interleaves** reduction steps and GC steps

An allocation is stuck if, after a full GC, there is not enough free space.

## Soundness Theorem

If  $\{ \diamond S \} \pi : t \{ \Psi \}$  holds, then  $t$  cannot reach a stuck configuration.

Reformulation: the live heap space of any execution of  $t$  cannot exceed  $S$ .



# Treiber's Push – The Standard Specification Using Atomic Triples

$$\begin{array}{l} \text{Private precondition} \longrightarrow \{ \text{stack-inv } s \} \\ \hline \langle \forall L. \text{content } s \ L \rangle \longleftarrow \text{Public precondition} \\ \pi: \text{push } s \ v \\ \langle \text{content } s \ (v::L) \rangle \longleftarrow \text{Public postcondition} \\ \hline \text{Private postcondition} \longrightarrow \{ \lambda \_ . \ulcorner \text{True} \urcorner \} \end{array}$$

- The public precondition is **atomically updated** into the public postcondition.
- The user can **open invariants** around a public precondition and postcondition.

## Space-Aware Treiber's Push

- Stack of **unboxed values**  $\rightsquigarrow$  focus on the space consumption of the structure.
- A fractional access token for the stack ( $r \in (0, 1]$ ):

stack  $s$   $r$

If  $r = 1$  then the stack is not being concurrently used.

$$\frac{\frac{\langle \forall L. \text{content } s L \rangle}{\pi: \text{push } s v}}{\langle \text{content } s (v::L) \rangle}}{\{ \lambda \_ . \text{stack } s r \}}$$

## Space-Aware Treiber's Pop

- pop cannot give back space.
- It returns **virtual credits**  $\blacklozenge_s 2$ .

$$\frac{\frac{\{ \text{stack } s \ r \}}{\langle \forall v \ L. \text{content } s \ (v :: L) \rangle} \quad \pi : \text{pop } s}{\langle \text{content } s \ L \rangle} \quad \frac{\langle \text{content } s \ L \rangle}{\{ \lambda v'. \ulcorner v' = v^\top * \text{stack } s \ r * \blacklozenge_s 2 \}}$$

If the stack is not being concurrently used,  
virtual credits can be converted to physical credits.

$$\text{stack } s \ 1 * \blacklozenge_s 2 \quad \Rightarrow \quad \text{stack } s \ 1 * \blacklozenge 2$$

## Allocation and Deallocation of a Stack

Allocation of a stack:

$$\{ \diamond 1 \} \pi : \text{create} () \{ \lambda s. \text{stack } s \ 1 * \text{content } s \ [] * s \leftarrow \emptyset * s \Leftarrow \{ \pi \} \}$$

Logical deallocation of a stack:

$$\begin{aligned} \text{stack } s \ 1 * \text{content } s \ L * s \leftarrow \emptyset * s \Leftarrow \emptyset &\Rightarrow \\ \text{stack } s \ 1 * \diamond(1 + 2 \times |L|) & \end{aligned}$$


## Stack With Boxed Values

$$\frac{\{ \text{stack } s \ r * v \leftarrow_q \emptyset * v \Leftarrow_p \{ \pi \} * \diamond 2 \}}{\langle \forall L. \text{content } s \ L \rangle}$$
$$\pi: \text{push } s \ v$$
$$\frac{\langle \text{content } s \ ((v, q, p) :: L) \rangle}{\{ \lambda \_ . \text{stack } s \ r \}}$$


The specification of `pop` is also richer:

- The pointed-by-heap assertion cannot be returned without an additional write.
- Can the pointed-by-thread assertion be returned?  
 $\rightsquigarrow$  **yes**, but we need to change the semantics.

## Stack With Boxed Values – Another Dangerous Interleaving



# Stack With Boxed Values – Another Dangerous Interleaving




```
push s v1
```

```
pop s;
```

```
...
```



## Stack With Boxed Values – Another Dangerous Interleaving




```
if CAS s c c'  
then ()  
else push s v1 -- retry
```

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





## Stack With Boxed Values – Another Dangerous Interleaving




```
if true
then ()
else push s v1 -- retry
```

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




## Stack With Boxed Values – Another Dangerous Interleaving



```
if true
then ()
else push s v1 -- retry
```

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

## Stack With Boxed Values – Another Dangerous Interleaving



```
if true
then ()
else push s v1 -- retry
```

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

The sleeping thread apparently **maintains reachable** the pushed value.

## Safe Points to the Rescue

The scenario of the previous slide **does not occur** in OCaml 5.

- The GC is stop-the-world.
- Threads run independently and, at **safe points**, lookup if a GC is pending.
- The GC runs only when **every** thread reached a safe point.
- In particular, no safe point before **if true then \_ else \_**.
- This expression **always** releases the roots of the **else** branch.
- Interleaving of small steps is **too fine**. We need a **coarser** interleaving.

For now: an atomic ifCAS primitive.

More meta-theory:

- Encoding of the logic for sequential programs
- Simplified mode where no deallocation is required
- Closures

More examples (with the indirect help of some of you!):

- Michael & Scott's queue, based on the proof of [Vindum and Birkedal \[2021\]](#)
- Async-Finish library (proof involving later credits [\[Spies et al., 2022\]](#))

↪ Some automation thanks to Diaframe [\[Mulder et al., 2022\]](#)

## Conclusion & Future Work

We present the first program logic to reason about

heap space usage for a high-level and concurrent language with GC.

Please talk to me if you know data structures with a cool space usage under GC!

Independently of heap space, our logic allows reasoning about **unreachability**.

↪ Can we apply our ideas to other areas?

**Thank you for your attention!**

alexandre.moine [at] inria.fr  
arthur.chargueraud [at] inria.fr  
francois.pottier [at] inria.fr

We handle cycles following the approach of [Madiot and Pottier \[2022\]](#).

$$\lceil \text{True} \rceil \multimap \emptyset \text{cloud}^0 P$$

$$l \mapsto_p \vec{v} \multimap l \leftarrow L \multimap l \Leftarrow \emptyset \quad D \text{cloud}^n P \quad \multimap (\{l\} \cup D) \text{cloud}^{n+\text{size}(\vec{v})} P \quad \text{if } L \subseteq P$$

$$D \text{cloud}^n D \Rightarrow \diamond n \multimap \left( \bigstar_{\ell \in D} \dagger \ell \right) \quad \text{if } D \cap \text{locs}(t) = \emptyset$$



## Proof Insight: Liveness-Based Cancellable Invariants

Recall standard cancellable invariants:

$$\text{CInv}^\gamma \Phi \triangleq \boxed{\Phi \vee \boxed{\mathbf{1}}^\gamma}$$

We can define liveness-based invariants:

$$\text{LInv}_\ell \Phi \triangleq \boxed{\Phi \vee \dagger \ell}$$

- The user can access the invariant as long as the location is logically allocated.
- Deallocation cancels the invariant.
- Useful to avoid another fractional token.

The semantics ensures no dangling pointers.

$$\frac{\ell \in \text{locs}(t)}{\{\dagger \ell\} \pi : t \{\Psi\}}$$

## Triples with Souvenir

Pointed-by-thread assertions are easy to manage in practice.

Introducing **triples with souvenir**  $[R]\{\Phi\} \pi: t \{ \Psi \}$

*“Give a pointed-by-thread assertion once and that’s it”*

$$\frac{\begin{array}{l} \text{locs}(t_2) = \{l\} \\ [R \cup \{l\}]\{\Phi\} \pi: t_1 \{ \Psi' \} \quad \forall v. [R]\{l \Leftarrow_p \{\pi\} * \Psi' \ v\} \pi: [v/x]t_2 \{ \Psi \} \end{array}}{[R]\{l \Leftarrow_p \{\pi\} * \Phi\} \pi: \text{let } x = t_1 \text{ in } t_2 \{ \Psi \}}$$

$$\frac{\begin{array}{l} \text{locs}(t_2) = \{l\} \quad l \in R \\ [R]\{\Phi\} \pi: t_1 \{ \Psi' \} \quad \forall v. [R]\{\Psi' \ v\} \pi: [v/x]t_2 \{ \Psi \} \end{array}}{[R]\{\Phi\} \pi: \text{let } x = t_1 \text{ in } t_2 \{ \Psi \}}$$

## The NoFree Mode

If the user pledges not to clean, framing pointed-by-thread assertions is not needed.

$$\frac{\text{LETNOCLEAN} \quad [\text{NoClean}] \{ \Phi \} \pi : t_1 \{ \Psi' \} \quad \forall v. [R] \{ \Psi' v \} \pi : [v/x] t_2 \{ \Psi \}}{[R] \{ \Phi \} \pi : \text{let } x = t_1 \text{ in } t_2 \{ \Psi \}}$$

- 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](https://doi.org/10.1007/978-3-642-37036-6_10).

## References] References

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>.

[] [

Ike Mulder, Robbert Krebbers, and Herman Geuvers. Diaframe: Automated verification of fine-grained concurrent programs in iris. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2022*, page 809–824, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392655. doi: 10.1145/3519939.3523432. URL <https://doi.org/10.1145/3519939.3523432>.

Simon Spies, Lennard Gäher, Joseph Tassarotti, Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. Later credits: Resourceful reasoning for the later modality. *Proc. ACM Program. Lang.*, 6(ICFP), aug 2022. doi: 10.1145/3547631. URL <https://doi.org/10.1145/3547631>.

[] [

Simon Friis Vindum and Lars Birkedal. Contextual refinement of the Michael-Scott queue. In *Certified Programs and Proofs (CPP)*, pages 76–90, January 2021. URL <https://cs.au.dk/~birke/papers/2021-ms-queue-final.pdf>.