

François Pottier
Jean-Marie Madiot

SLO

A Separation Logic for Heap Space under GC

We wish to verify a program's *heap space* usage,

- using *separation logic*,
- by viewing heap space as a *resource*...

Reasoning about Heap Space, without GC

Idea 1. Following Hofmann (1999), let $\diamond 1$ represent one *space credit*.

Allocation consumes credits; deallocation produces credits.

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

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

A function's space requirement is visible in its specification.

In the presence of GC, what Happens?

In the presence of GC,

- deallocation becomes *implicit*,
- so we lose the ability to recover space credits while reasoning.

A Ghost Deallocation Operation?

Idea 2. Switch to a *logical deallocation* operation:

$$x \mapsto b \quad \equiv \quad \diamond \text{size}(b)$$

A ghost update \Rightarrow *consumes* an assertion and *produces* an assertion.

This marries

- *manual reasoning* about memory at verification time
- *automatic management* of memory at runtime.

Is logical deallocation sound?

$$x \mapsto b \quad \Rightarrow \quad \diamond \text{size}(b)$$

It does have a few good properties: *no double-free*, *no use-after-free*.

Because $x \mapsto b$ is consumed,

- a block cannot be logically deallocated twice;
- a block cannot be accessed after it has been logically deallocated.

Unfortunately, logical deallocation in this form is *not sound*.

$$x \mapsto \ell \quad \not\equiv \quad \diamond \text{size}(\ell)$$

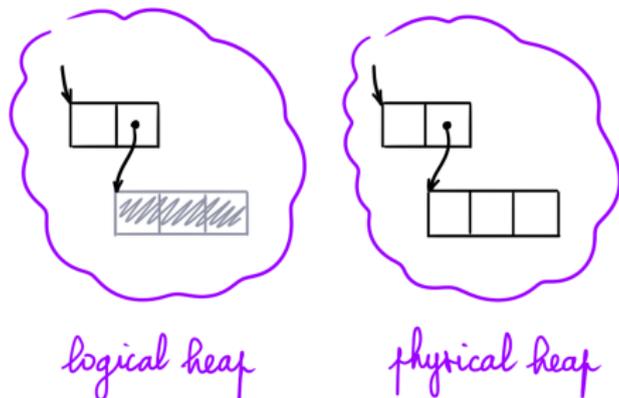
Introducing logical deallocation creates a distinction between

- the *logical heap* that the programmer keeps in mind,
- the *physical heap* that exists at runtime.

Logical versus Physical Heaps

The following situation is problematic.

The programmer has logically deallocated a block and obtained $\diamond 3$,



but this block is *reachable* and cannot be reclaimed by the GC.

We have 3 space credits but *no free space* in the physical heap!

To avoid this problem, we must *restrict logical deallocation*:

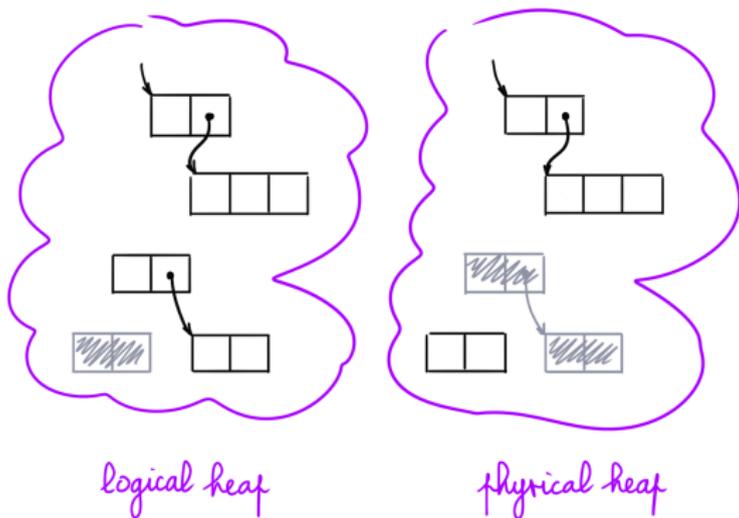
- A reachable block must not be deallocated.

In the contrapositive,

- A block should be logically deallocatable only if it is *unreachable*,
- so the GC *can* reclaim this block,
- so the logical and physical heaps remain *synchronized*.

The Desired Global Invariant

The logical and physical heaps *coincide on their reachable fragments*.



Restricting Logical Deallocation: How?

How do we restrict logical deallocation?

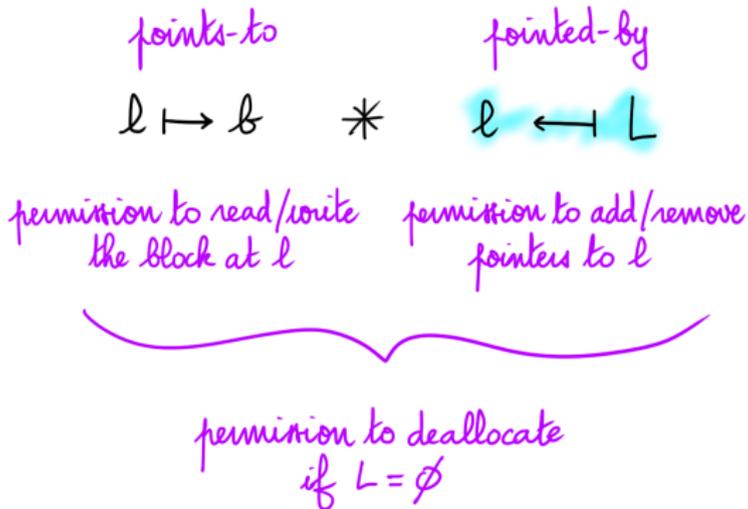
- We want to disallow deallocating a *reachable* block,
- but Separation Logic lets us reason about *ownership*.
- Reachability is a *nonlocal* property.

Idea 3. Following Kassios and Kritikos (2013),

- we *keep track of the predecessors* of every block.
- If a block has no predecessor, *then* it is unreachable,
- therefore it can be logically deallocated.

Points-To and Pointed-By Assertions

In addition to *points-to*, we use *pointed-by* assertions:



We get a sound logical deallocation axiom:

$$x \mapsto b * x \leftarrow \emptyset \Rightarrow \diamond \text{size}(b)$$

This axiom deallocates one block.

There is also a *bulk logical deallocation* axiom.

We want the pointers *from the stack(s) to the heap* to be explicit,

- so the operational semantics views them as GC *roots*,
- so our predecessor-tracking logic keeps track of them.

Idea 4. Use a low-level calculus where *stack cells* are explicit.

- 1 A Glimpse of SpaceLang
- 2 A Glimpse of the Reasoning Rules
- 3 Specification of a Stack
- 4 Conclusion



SpaceLang is imperative. An *instruction* i does not return a value.

skip	<i>no-op</i>	$*\rho = \text{alloc } n$	<i>heap allocation</i>
$i; i$	<i>sequencing</i>	$*\rho = [*\rho + o]$	<i>heap load</i>
if $*\rho$ then i else i	<i>conditional</i>	$[*\rho + o] = *\rho$	<i>heap store</i>
$*\rho(\vec{\rho})$	<i>procedure call</i>	$*\rho = (*\rho == *\rho)$	<i>address comparison</i>
$*\rho = v$	<i>constant load</i>	alloca x in i	<i>stack allocation</i>
$*\rho = *\rho$	<i>move</i>	fork $*\rho$ as x in i	<i>thread creation</i>

The operands of every instruction are stack cells ρ .

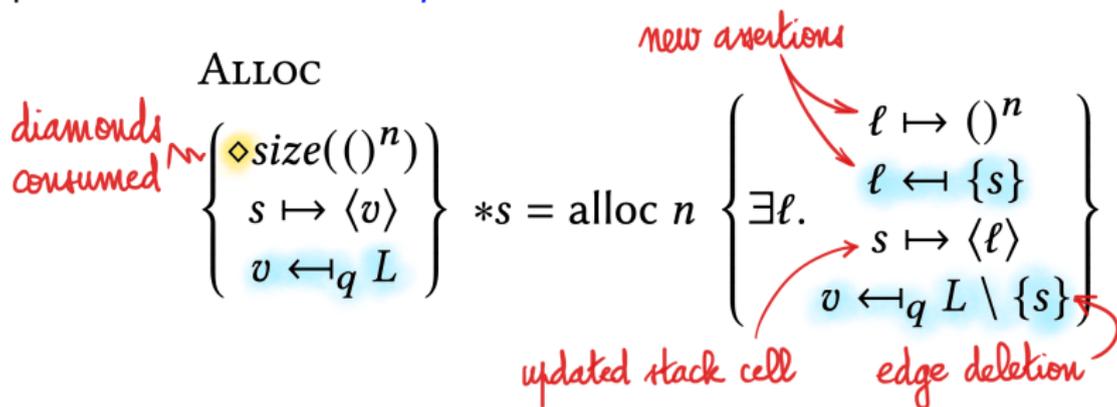
There is *no heap deallocation* instruction.

A small-step operational semantics, with a few unique features:

- *Garbage collection* takes place before every reduction step.
- The GC *roots* are the stack cells.
- Heap allocation *fails* if the heap size exceeds a fixed limit S .

- ① A Glimpse of SpaceLang
- ② A Glimpse of the Reasoning Rules**
- ③ Specification of a Stack
- ④ Conclusion

Heap allocation *consumes space credits*.



Points-to and pointed-by assertions for the new location appear.

One pointer to the value v is *deleted*.

Reasoning about a heap store involves some administration...

STORE

$$\vec{v}(o) = v$$

$$\frac{\left\{ \begin{array}{l} s \mapsto \langle \ell \rangle \\ r \mapsto \langle v' \rangle \\ \ell \mapsto_1 \vec{v} \\ v \leftarrow_q L \\ v' \leftarrow_{q'} L' \end{array} \right\} \quad [*s + o] = *r \quad \left\{ \begin{array}{l} s \mapsto \langle \ell \rangle \\ r \mapsto \langle v' \rangle \\ \ell \mapsto_1 [o := v'] \vec{v} \\ v \leftarrow_q L \setminus \{\ell\} \\ v' \leftarrow_{q'} L' \uplus \{\ell\} \end{array} \right\}}{\text{operands } \swarrow \quad \swarrow \text{ updated heap cell} \quad \swarrow \text{ edge deletion} \quad \swarrow \text{ edge addition}}$$

One pointer to v is *deleted*; one pointer to v' is *created*.

Logical deallocation of a block is a *ghost operation*:

$$\begin{array}{c}
 \text{Knowledge of} \\
 \text{all antecedents} \\
 \xrightarrow{\quad} \\
 l \mapsto_1 \vec{v} * l \xleftarrow{1} L * \text{dom}(L) \subseteq \{l\} \quad \Rightarrow_1 \quad \ddagger\{l\} * \diamond \text{size}(\vec{v}) \\
 \underbrace{\quad}_{\text{ownership}} \quad \underbrace{\quad}_{\text{no antecedent}} \quad \underbrace{\quad}_{\text{(but self)}} \quad \underbrace{\quad}_{\text{location now dead}} \quad \underbrace{\quad}_{\text{credit!}} \\
 \text{of the block}
 \end{array}$$

Theorem (Soundness)

If $\{\diamond S\} i \{True\}$ holds, then, executing i in an empty store cannot lead to a situation where a thread is stuck.

*If, under a precondition of S space credits, the code can be verified, then its *live* heap space cannot exceed S .*

This holds *regardless of the value of S* (the heap size limit).
Furthermore, the reasoning rules are *independent* of S .

The rules allow *compositional reasoning* about space.

- ① A Glimpse of SpaceLang
- ② A Glimpse of the Reasoning Rules
- ③ Specification of a Stack**
- ④ Conclusion

The user may define *custom* (simplified) predecessor tracking disciplines.
For example, sometimes, *counting* predecessors is enough.

$$v \leftarrow_1 n \triangleq \exists L. (v \leftarrow_1 L \star |L| = n)$$

Edge addition and deletion increment and decrement n .

Creating a stack *consumes 4 space credits*.

$$\left\{ \begin{array}{l} f \mapsto \langle \text{create} \rangle \\ \text{stack} \mapsto \langle () \rangle \\ \diamond 4 \end{array} \right\} *f(\text{stack}) \left\{ \begin{array}{l} f \mapsto \langle \text{create} \rangle \\ \exists \ell. \text{stack} \mapsto \langle \ell \rangle \\ \text{isStack } \ell [] \star \ell \leftarrow 1 \end{array} \right\}$$

We get unique ownership of the stack and *we have the sole pointer* to it.

Pushing *consumes 4 space credits*.

$$\left\{ \begin{array}{l} f \mapsto \langle \text{push} \rangle \\ \text{stack} \mapsto \langle \ell \rangle \\ \text{elem} \mapsto \langle v \rangle \\ \diamond 4 \star \text{isStack } \ell \text{ vs} \\ v \leftarrow n \end{array} \right\} *f(\text{stack}, \text{elem}) \left\{ \begin{array}{l} f \mapsto \langle \text{push} \rangle \\ \text{stack} \mapsto \langle \ell \rangle \\ \text{elem} \mapsto \langle v \rangle \\ \text{isStack } \ell (v :: \text{vs}) \\ v \leftarrow n + 1 \end{array} \right\}$$

The value v receives *one more antecedent*.

Popping *freed up 4 space credits*.

$$\left\{ \begin{array}{l} f \mapsto \langle \text{pop} \rangle \\ \text{stack} \mapsto \langle \ell \rangle \\ \text{elem} \mapsto \langle () \rangle \\ \text{isStack } \ell (v :: vs) \\ v \leftarrow n \end{array} \right\} *f(\text{stack}, \text{elem}) \left\{ \begin{array}{l} f \mapsto \langle \text{pop} \rangle \\ \text{stack} \mapsto \langle \ell \rangle \\ \text{elem} \mapsto \langle v \rangle \\ \diamond 4 * \text{isStack } \ell \text{ vs} \\ v \leftarrow n \end{array} \right\}$$

The number of predecessors of v is unchanged,
because the out-parameter $elem$ receives a pointer to it.

Logically deallocation of the stack, a *ghost operation*, is part of the API. It requires proving that the stack has *zero predecessors*.

$$\left\{ \begin{array}{l} isStack \ell \text{ vs} \star \ell \leftarrow 0 \\ \star \quad v \leftarrow n \\ (v,n) \in vns \end{array} \right\} \Rightarrow_I \left\{ \begin{array}{l} \diamond(4 + 4 \times |vs|) \\ \star \quad v \leftarrow n - (v \$ vs) \\ (v,n) \in vns \end{array} \right\}$$

It frees up *a linear number of space credits*.

- ① A Glimpse of SpaceLang
- ② A Glimpse of the Reasoning Rules
- ③ Specification of a Stack
- ④ Conclusion



A sound logic to reason about heap space usage in the presence of GC.

Our main insights:

- Allocation consumes *space credits* $\diamond n$.
- *Logical deallocation*, a ghost operation, produces space credits.
- Logical deallocation requires *predecessor tracking*, which we perform via *pointed-by assertions* $v \leftarrow L$.

Currently, predecessor tracking requires *heavy bookkeeping*.

We are investigating

- a more flexible *deferred* logical edge deletion mechanism;
- coarse-grained predecessor tracking based on *islands*;
- *simpler / more automated* tracking of *roots*;
- reasoning directly about call-by-value *λ -calculus*.

- 5 Syntax, Semantics of SpaceLang
- 6 Reasoning Rules of SL \diamond
- 7 Specification of List Copy

Memory locations: $l, c, r, s \in \mathcal{L}$.

Values include constants, memory locations, and *closed procedures*:

$$v ::= () \mid k \mid \ell \mid \lambda \vec{x}.i$$

Memory blocks include *heap tuples*, *stack cells*, and deallocated blocks:

$$b ::= \vec{v} \mid \langle v \rangle \mid \blacklightning$$

A *store* maps locations to blocks, encompassing the heap and stack(s).

The *size* of a block:

$$\text{size}(\vec{v}) = 1 + |\vec{v}| \quad \text{size}(\langle v \rangle) = \text{size}(\blacklightning) = 0$$

The size of the store is the sum of the sizes of all blocks.

A *reference* is a variable or a (stack) location and denotes a *stack cell*.

$$\rho ::= x \mid c$$

SpaceLang uses *call-by-reference*.

A variable denotes a closed reference, *not* a closed value as is usual.

The operational semantics involves substitutions $[c/x]$.

This preserves the property that *the code never points to the heap*.

The *roots* of the garbage collection process are *the stack cells*.

SpaceLang is imperative. An *instruction* i does not return a value.

skip	<i>no-op</i>	$*\rho = \text{alloc } n$	<i>heap allocation</i>
$i; i$	<i>sequencing</i>	$*\rho = [* \rho + o]$	<i>heap load</i>
if $*\rho$ then i else i	<i>conditional</i>	$[* \rho + o] = * \rho$	<i>heap store</i>
$*\rho(\vec{\rho})$	<i>procedure call</i>	$*\rho = (* \rho == * \rho)$	<i>address comparison</i>
$*\rho = v$	<i>constant load</i>	alloca x in i	<i>stack allocation</i>
$*\rho = * \rho$	<i>move</i>	alloca c in i	<i>active stack cell</i>
		fork $*\rho$ as x in i	<i>thread creation</i>

The operands of every instruction are stack cells (ρ).

There is no deallocation instruction for heap blocks.

We fix a *maximum heap size* S .

Heap allocation *fails* if the heap size exceeds S .

$$\text{STEPALLOC} \frac{\sigma' = [\ell += ()^n]\sigma \quad \text{size}(\sigma') \leq S \quad \sigma'' = \langle s := \ell \rangle \sigma'}{*s = \text{alloc } n / \sigma \longrightarrow \text{skip} / \sigma''}$$

S is a parameter of the operational semantics,

but the reasoning rules of $\text{SL}\diamond$ are independent of S .

The dynamic semantics of stack allocation is in *three steps*:

$$\frac{\text{STEPALLOCAENTRY} \quad \sigma' = [c += \langle () \rangle] \sigma}{\text{alloca } x \text{ in } i / \sigma \longrightarrow \text{alloca } c \text{ in } [c/x]i / \sigma'}$$

$$\frac{\text{STEPALLOCAEXIT} \quad \sigma(c) = \langle v \rangle \quad \sigma' = [c := \text{!}] \sigma}{\text{alloca } c \text{ in skip} / \sigma \longrightarrow \text{skip} / \sigma'}$$

Evaluation contexts: $K ::= [] \mid K; i \mid \text{alloca } c \text{ in } K.$

To complete the definition of the operational semantics,

- allow *garbage collection* before every reduction step.

$\sigma \text{ } \text{⊞} \text{ } \sigma'$ holds if

- the stores σ and σ' have the same domain;
 - for every ℓ in this domain,
either $\sigma'(\ell) = \sigma(\ell)$, or ℓ is unreachable in σ and $\sigma'(\ell) = \text{⚡}$.
- allow *thread interleavings* (comes for free with Iris).

Complete Operational Semantics

STEPSEQSKIP
 $\text{skip}; i / \sigma \longrightarrow i / \sigma$

STEPIF

$$\frac{\sigma(r) = \langle k \rangle}{\text{if } *r \text{ then } i_1 \text{ else } i_2 / \sigma \longrightarrow k \neq 0 ? i_1 : i_2 / \sigma}$$

STEPCALL

$$\frac{\sigma(r) = \langle \lambda \bar{x}. i \rangle \quad |\bar{x}| = |\bar{s}|}{*r(\bar{s}) / \sigma \longrightarrow [\bar{s}/\bar{x}]i / \sigma}$$

STEPCONST

$$\frac{\sigma' = \langle s := v \rangle \sigma \quad \text{pointers}(v) = \emptyset}{*s = v / \sigma \longrightarrow \text{skip} / \sigma'}$$

STEPMOVE

$$\frac{\sigma(r) = \langle v \rangle \quad \sigma' = \langle s := v \rangle \sigma}{*s = *r / \sigma \longrightarrow \text{skip} / \sigma'}$$

STEPALLOC

$$\frac{\sigma' = [\ell += ()^n] \sigma \quad \text{size}(\sigma') \leq S \quad \sigma'' = \langle s := \ell \rangle \sigma'}{*s = \text{alloc } n / \sigma \longrightarrow \text{skip} / \sigma''}$$

STEPLOAD

$$\frac{\sigma(r) = \langle \ell \rangle \quad \sigma(\ell) = \vec{v} \quad 0 \leq o < |\vec{v}| \quad \vec{v}(o) = v \quad \sigma' = \langle s := v \rangle \sigma}{*s = [*r + o] / \sigma \longrightarrow \text{skip} / \sigma'}$$

STEPSTORE

$$\frac{\sigma(r) = \langle v \rangle \quad \sigma(s) = \langle \ell \rangle \quad \sigma(\ell) = \vec{v} \quad 0 \leq o < |\vec{v}| \quad \sigma' = [\ell := [o := v]\vec{v}] \sigma}{[*s + o] = *r / \sigma \longrightarrow \text{skip} / \sigma'}$$

STEPLOCSEQ

$$\frac{\sigma(r_1) = \langle \ell_1 \rangle \quad \sigma(r_2) = \langle \ell_2 \rangle \quad \sigma' = \langle s := (\ell_1 = \ell_2 ? 1 : 0) \rangle \sigma}{*s = (*r_1 == *r_2) / \sigma \longrightarrow \text{skip} / \sigma'}$$

STEPALLOCAENTRY

$$\frac{\sigma' = [c += ()] \sigma}{\text{alloca } x \text{ in } i / \sigma \longrightarrow \text{alloca } c \text{ in } [c/x]i / \sigma'}$$

STEPALLOCAEXIT

$$\frac{\sigma(c) = \langle v \rangle \quad \sigma' = [c := \text{!}] \sigma}{\text{alloca } c \text{ in skip} / \sigma \longrightarrow \text{skip} / \sigma'}$$

STEPFORK

$$\frac{\sigma(r) = \langle v \rangle \quad \sigma' = [r := ()][c += \langle v \rangle] \sigma}{\text{fork } *r \text{ as } x \text{ in } i / \sigma \longrightarrow \text{skip} / \sigma' \quad \text{spawning } \text{alloca } c \text{ in } [c/x]i}$$

STEPCONTEXT

$$\frac{i / \sigma \longrightarrow i' / \sigma' \quad \text{spawning } \vec{i}}{K[i] / \sigma \longrightarrow K[i'] / \sigma' \quad \text{spawning } \vec{i}}$$

- 5 Syntax, Semantics of SpaceLang
- 6 Reasoning Rules of SL \diamond**
- 7 Specification of List Copy

Heap allocation *consumes space credits*.

$$\begin{array}{l}
 \text{ALLOC} \\
 \left\{ \begin{array}{l}
 \diamond \text{size}(()^n) \\
 s \mapsto \langle v \rangle \\
 v \leftarrow_q L
 \end{array} \right\} *s = \text{alloc } n \left\{ \begin{array}{l}
 \exists \ell. \\
 \ell \mapsto ()^n \\
 \ell \leftarrow \{s\} \\
 s \mapsto \langle \ell \rangle \\
 v \leftarrow_q L \setminus \{s\}
 \end{array} \right\}
 \end{array}$$

diamonds consumed (pointing to $\diamond \text{size}(()^n)$)
new assertions (pointing to $\ell \mapsto ()^n$ and $\ell \leftarrow \{s\}$)
updated stack cell (pointing to $v \leftarrow_q L \setminus \{s\}$)
edge deletion (pointing to $v \leftarrow_q L \setminus \{s\}$)

Points-to and pointed-by assertions for the new location appear.

One pointer to the value v is *deleted*. (This aspect is optional.)

Writing a heap cell is simple... but involves some administration.

STORE

$$\vec{v}(o) = v$$

$$\frac{\left\{ \begin{array}{l} s \mapsto \langle \ell \rangle \\ r \mapsto \langle v' \rangle \\ \ell \mapsto_1 \vec{v} \\ v \leftarrow_q L \\ v' \leftarrow_{q'} L' \end{array} \right\} \quad [*s + o] = *r}{\left\{ \begin{array}{l} s \mapsto \langle \ell \rangle \\ r \mapsto \langle v' \rangle \\ \ell \mapsto_1 [o := v'] \vec{v} \\ v \leftarrow_q L \setminus \{\ell\} \\ v' \leftarrow_{q'} L' \uplus \{\ell\} \end{array} \right\}}$$

operands \swarrow *edge addition* \rightarrow *updated heap cell* \swarrow
edge deletion \leftarrow

One pointer to v is deleted; one pointer to v' is *created*.

A points-to assertion for the new stack cell exists throughout its lifetime.

$$\begin{array}{c}
 \text{ALLOCA} \\
 \frac{\forall c. \{\Phi \star c \mapsto \langle () \rangle\} [c/x] i \{c \mapsto \langle () \rangle \star \Psi\}}{\{\Phi\} \text{alloca } x \text{ in } i \{\Psi\}}
 \end{array}$$

stack cell appears *stack cell disappears*

No pointed-by assertion is provided. (A design choice.)

- No pointers (from the heap or stack) to the stack.

Logical deallocation of a block is a *ghost operation*:

$$\begin{array}{c}
 \text{Knowledge of} \\
 \text{all antecedents} \\
 \xrightarrow{\quad} \\
 l \mapsto_1 \vec{v} * l \xleftarrow{1} L * \text{dom}(L) \subseteq \{l\} \quad \Rightarrow_1 \quad \dagger\{l\} * \diamond \text{size}(\vec{v}) \\
 \underbrace{\quad}_{\text{ownership}} \quad \underbrace{\quad}_{\text{no antecedent}} \quad \underbrace{\quad}_{\text{(but self)}} \quad \underbrace{\quad}_{\text{location now dead}} \quad \underbrace{\quad}_{\text{credit!}} \\
 \text{of the block} \quad \text{(but self)} \quad \text{credit!}
 \end{array}$$

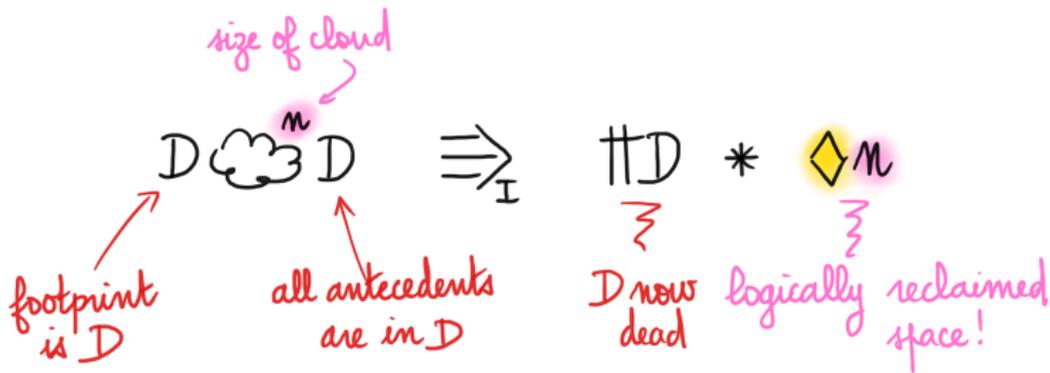
Deletion of deallocated predecessors can be *deferred*:

$$\nu \leftarrow_q L * \prod_{\substack{\mathcal{D} \\ \text{dead} \\ \text{locations}}} \Rightarrow_{\mathcal{I}} \nu \leftarrow_q L' \quad \text{removed from antecedents}$$

if $\text{dom}(L \setminus L') \subseteq \mathcal{D}$

A key rule: if L' is empty, then ν becomes eligible for deallocation.

A group that is *closed under predecessors* can be deallocated at once:



The rules for constructing a "cloud" (omitted) are straightforward.

Points-to and pointed-by assertions can be *split* and *joined*.

$$l \mapsto_{q_1+q_2} b \equiv l \mapsto_{q_1} b * l \mapsto_{q_2} b$$

$$v \leftarrow_{q_1+q_2} L_1 \uplus L_2 \equiv v \leftarrow_{q_1} L_1 * v \leftarrow_{q_2} L_2$$

$$v \leftarrow_q L \longrightarrow * v \leftarrow_q L' \quad \text{if } L \subseteq L'$$

$$l \mapsto_q b * l' \leftarrow_{q_1} L \stackrel{\exists \mathbb{I}}{\implies} l \mapsto_q b * l' \leftarrow_{q_1} L * \text{'}l' \$ pointers(b) \leq l \$ L \text{'}$$

Pointed-by assertions are *covariant*.

Points-to and pointed-by assertions can be *confronted*.

Space credits can be *split* and *joined*.

$$\begin{aligned} \text{True} &\equiv_{\text{I}} \diamond 0 \\ \diamond(m_1 + m_2) &\equiv_{\text{I}} \diamond m_1 * \diamond m_2 \end{aligned}$$

- 5 Syntax, Semantics of SpaceLang
- 6 Reasoning Rules of SL \diamond
- 7 Specification of List Copy**

Each cell owns the next cell and possesses *the sole pointer* to it.

$$\begin{aligned}
 isList\ l\ [] &\triangleq l \mapsto [0] \\
 isList\ l\ (v :: vs) &\triangleq \exists l'. l \mapsto [1; v; l'] \star l' \leftarrow 1 \star isList\ l'\ vs
 \end{aligned}$$

Let's now have a look at *list copy* and its spec. (Fasten seatbelts!)

$copy \triangleq \lambda(self, dst, src).$

 alloca *tag* in $*tag = [*src + 0];$

 if $*tag$ then

 alloca *head* in $*head = [*src + 1];$

 alloca *tail* in $*tail = [*src + 2];$

$*src = ();$

 alloca *dst'* in $*self(self, dst', tail);$

$*dst = \text{alloc } 3;$

$[*dst + 0] = *tag;$

$[*dst + 1] = *head;$

$[*dst + 2] = *dst'$

 else

$*src = ();$

$*dst = \text{alloc } 1;$

$[*dst + 0] = *tag$

- read the list's tag
- if this is a cons cell, then
- read the list's head
- read the list's tail
- clobber this root
- copy the list's tail
- allocate a new cons cell
- and initialize it

- this must be a nil cell
- clobber this root
- allocate a new nil cell
- and initialize it

Specification of List Copy

The case $m = 1$, where we have *the sole pointer* to the list, is special.

$$\left\{ \begin{array}{l}
 f \mapsto \langle \text{copy} \rangle * dst \mapsto \langle () \rangle * src \mapsto \langle \ell \rangle \\
 \text{isList } \ell \text{ vs } * \ell \leftarrow m \\
 m = 1 ? 0 : (2 + 4 \times |vs|) \\
 \forall v \in vs. \exists n. (v, n) \in vns \\
 *_{(v,n) \in vns} v \leftarrow n
 \end{array} \right\} \begin{array}{l}
 \text{need no space} \\
 \text{or} \\
 \underline{\text{linear space}}
 \end{array}$$

$$*f(f, dst, src)$$

$$\exists \ell'. \left\{ \begin{array}{l}
 f \mapsto \langle \text{copy} \rangle * dst \mapsto \langle \ell' \rangle * src \mapsto \langle () \rangle \\
 m = 1 ? \text{True} : (\text{isList } \ell \text{ vs } * \ell \leftarrow m - 1) \\
 \text{isList } \ell' \text{ vs } * \ell' \leftarrow 1 \\
 *_{(v,n) \in vns} v \leftarrow n + (m = 1 ? 0 : v \$ vs)
 \end{array} \right\} \begin{array}{l}
 \text{orig. list is} \\
 \text{deallocated} \\
 \text{or} \\
 \underline{\text{preserved}}
 \end{array}$$

each element receives
 zero new antecedent
 or
 a number of new antecedents