

François Pottier  
Jean-Marie Madiot

SLO

A Separation Logic for Heap Space under GC

How can we establish *formal (verified)* bounds  
on a program's *heap space* usage?

We wish to

- work in the setting of a *program logic*,
- view heap space as a *resource*.

# Reasoning about Heap Space, without GC

Following Hofmann (1999, 2000), 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.

End of talk...?

# In the presence of GC, what Happens?

Garbage collection can offer superior *simplicity, safety, performance*.

In the presence of GC,

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

# A Ghost Deallocation Operation?

It is tempting to switch to a *logical deallocation* operation:

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

This would marry

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

At least two questions spring to mind:

Is this approach *practical*? Is it *sound*?

A pitfall would be to get *the worst of both worlds*:

- mental burden of manual reasoning about memory deallocation,
- performance issues sometimes caused by GC.

Yet we can strive to get the *best* of each:

- simplicity and possibly superior performance afforded by GC,
- reasoning at a suitable level of abstraction:  
e.g., via *bulk logical deallocation*.

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*.

- 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  $\diamond 3$ ,



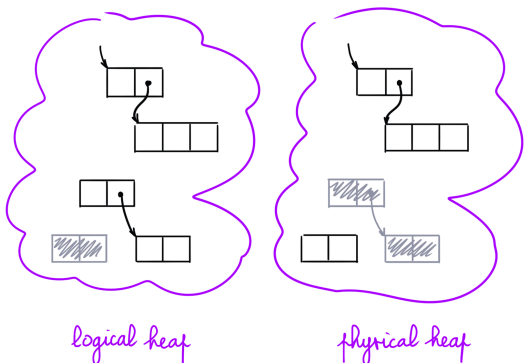
but this block is *reachable* and cannot be reclaimed by the GC.

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

To avoid this problem, we want to *restrict logical deallocation*.

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

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



So,  $\diamond k$  implies  $k$  free words *exist* in the logical heap  
 implies  $k$  free words *can be created* in the physical heap.

The outstanding problem is, *how* do we restrict logical deallocation?

- We want to disallow deallocating a *reachable* block,
- but Separation Logic lets us reason about *ownership*.
- Ownership and reachability are *unrelated!*
- Furthermore, reachability is a *nonlocal* property.

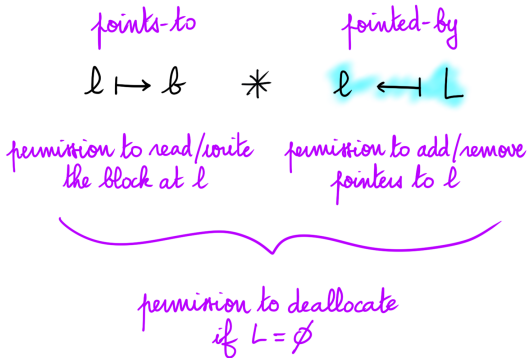
*Not* requiring reachability reasoning is a strength of traditional SL.

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, for a single block:

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



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.

This leads to a calculus where *stack cells* are explicit  
and *a variable denotes an address* on the stack.

- 1 Syntax, Semantics of SpaceLang
- 2 Reasoning Rules of SL $\diamond$
- 3 Ghost Reference Counting
- 4 Examples of Specifications
- 5 Conclusion



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 ( $\rho$ ).

There is no deallocation instruction for heap blocks.

We fix a *maximum heap size*  $S$ .

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

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

$S$  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.

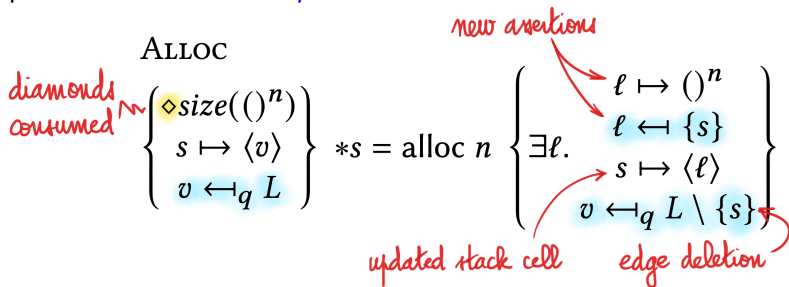
$\sigma \text{ } \text{☒} \text{ } \sigma'$  holds if

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



- 1 Syntax, Semantics of SpaceLang
- 2 Reasoning Rules of SL $\diamond$**
- 3 Ghost Reference Counting
- 4 Examples of Specifications
- 5 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*. (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*:

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


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  $\nu$  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*.

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

## 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 the code is verified under  $S$  space credits, then its heap space usage cannot exceed  $S$ .

This guarantee holds *for every*  $S$ .

The reasoning rules are *independent* of  $S$ .

The rules allow *compositional reasoning* about space.

- 1 Syntax, Semantics of SpaceLang
- 2 Reasoning Rules of SL $\diamond$
- 3 Ghost Reference Counting**
- 4 Examples of Specifications
- 5 Conclusion

# Choice of a Predecessor Tracking Discipline

Keeping track of a *multiset* of predecessors can be heavy.

Sometimes

- *counting* predecessors is enough,
- or recording what *regions* the predecessors inhabit is enough.

Can *high-level predecessor tracking disciplines* be defined on top of  $SL_{\diamond}$ ?

# Example: Ghost Reference Counting

The simplified pointed-by assertion  $v \leftarrow n$  counts predecessors:

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

Edge addition / deletion increment / decrement  $n$ .

- 1 Syntax, Semantics of SpaceLang
- 2 Reasoning Rules of SL $\diamond$
- 3 Ghost Reference Counting
- 4 Examples of Specifications**
  - A Stack
  - List Copy
- 5 Conclusion

## 4 Examples of Specifications

A Stack

List Copy

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*.

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

The number of antecedents of  $v$  is unchanged, as  $\text{elem}$  points to it.

Logically deallocating the entire stack is a *ghost operation*.

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

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

The ghost reference counters of the stack elements are decremented.

## 4 Examples of Specifications

A Stack

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.

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


---


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


---


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

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

- 1 Syntax, Semantics of SpaceLang
- 2 Reasoning Rules of SL $\diamond$
- 3 Ghost Reference Counting
- 4 Examples of Specifications
- 5 Conclusion





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

- Allocation consumes *space credits*  $\diamond n$ .
- *Logical deallocation* is a ghost operation.
- Logical deallocation requires *predecessor tracking*  $v \leftrightarrow L$ .

Predecessor tracking still requires *too much administration*.

We are investigating

- *deferred* edge deletion;
- *automated or simplified* tracking of *roots*;
- predecessor tracking based on *regions*;
- notions of *single-entry-point* regions.

We would also like to adapt  $SL_{\diamond}$  directly to call-by-value  $\lambda$ -calculus.

# A Bit of Controversy about OCaml

During this traversal, *which part of the tree is live?*

```
type tree = Leaf | Node of tree * tree
```

```
let rec walk t =  
  match t with  
  | Leaf          -> ()  
  | Node (t1, t2) -> walk t1; walk t2
```

# A Bit of Controversy about OCaml

During this traversal, *which part of the tree is live?*

```
type tree = Leaf | Node of tree * tree
```

```
let rec walk t =  
  match t with  
  | Leaf          -> ()  
  | Node (t1, t2) -> walk t1; walk t2
```

It could (should?) be *the subtrees that have not yet been traversed*, because t2 remains live while walk t1 is executed...

# A Bit of Controversy about OCaml

But the OCaml compiler transforms the code roughly as follows:

```
type tree = Leaf | Node of tree * tree

let rec walk t =
  match t with
  | Leaf          -> ()
  | Node ( _, _ ) -> walk t.1 ; walk t.2
```

Thus, `t` remains live while `walk t.1` is executed.

Every *left subtree remains live* until it has been entirely traversed.

Reasoning about space at this level

requires a *precise definition* of where each variable is a root.

# 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  

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