

Temporary Read-Only Permissions for Separation Logic

Making Separation Logic's
Small Axioms
Smaller

Arthur Charguéraud François Pottier



Gallium seminar
Paris, April 10, 2017

Separation Logic: to own, or not to own

Separation Logic (Reynolds, 2002) is about **disjointness** of heap fragments.

- ▶ what “we” own, versus what “others” own.

Therefore, it is about **unique ownership**. A dichotomy arises:

- ▶ Of every memory cell, **either** we have ownership, **or** we don't.
- ▶ If we do, then we can **read and write** this cell.
- ▶ If we don't, then we can **neither write nor even read** this cell.

From memory cells (and arrays), this dichotomy extends to data structures:

*To a (user-defined) data structure,
either we have **no access at all**, or we have **full read-write access**.*

Separation Logic's read and write axioms

The reasoning rule for writing a cell requires (and returns) a full permission:

$$\text{SET} \\ \{l \hookrightarrow v'\} (\text{set } l \ v) \{\lambda y. l \hookrightarrow v\}$$

So does the reasoning rule for reading a cell:

$$\text{TRADITIONAL READ AXIOM} \\ \{l \hookrightarrow v\} (\text{get } l) \{\lambda y. [y = v] \star l \hookrightarrow v\}$$

They are known as “**small axioms**”, because they require minimum permission.

Separation Logic's read and write axioms

The reasoning rule for writing a cell requires (and returns) a full permission:

$$\text{SET} \\ \{l \hookrightarrow v'\} (\text{set } l \ v) \{\lambda y. l \hookrightarrow v\}$$

So does the reasoning rule for reading a cell:

$$\text{TRADITIONAL READ AXIOM} \\ \{l \hookrightarrow v\} (\text{get } l) \{\lambda y. [y = v] \star l \hookrightarrow v\}$$

terminology:
“permission” = “assertion”

They are known as “**small axioms**”, because they require minimum permission.

Separation Logic's read and write axioms

The reasoning rule for writing a cell requires (and returns) a full permission:

$$\text{SET} \\ \{l \hookrightarrow v'\} (\text{set } l \ v) \{\lambda y. l \hookrightarrow v\}$$

So does the reasoning rule for reading a cell:

$$\text{TRADITIONAL READ AXIOM} \\ \{l \hookrightarrow v\} (\text{get } l) \{\lambda y. [y = v] \star l \hookrightarrow v\}$$

They are known as “**small axioms**”, because they require minimum permission.

But are they as small as they could be?

Isn't it excessive for reading to require a full permission?

Are there adverse consequences of working with such coarse permissions?

The problem

Suppose we are implementing an abstract data type of (mutable) sequences.

Here is a typical specification of sequence concatenation:

$$\{s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2\}$$

(*append* s_1 s_2)

$$\{\lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 ++ L_2) \star s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2\}$$

The problem

Suppose we are implementing an abstract data type of (mutable) sequences.

Here is a typical specification of sequence concatenation:

$$\{s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2\}$$

(*append* s_1 s_2)

$$\{\lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 ++ L_2) \star s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2\}$$

Although correct, this style of specification can be criticized on several grounds:

- ▶ It is a bit **noisy**.
- ▶ It requires the permissions $s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2$ to be **threaded throughout the proof of *append***.
- ▶ It actually **does not guarantee that s_1 and s_2 are unmodified**. — (next slide)
- ▶ It requires s_1 and s_2 to be **distinct** data structures. — (slide after next)

The problem, facet 3: temporary modifications are not forbidden

Repeating “ $s \rightsquigarrow \text{Seq } L$ ” in the pre- and postcondition can be deceiving.

This does **not** forbid changes to the **concrete** data structure in memory.

Here is a function that really just reads the data structure:

$$\{s \rightsquigarrow \text{Seq } L\} (\text{length } s) \{\lambda y. s \rightsquigarrow \text{Seq } L \star [y = |L|]\}$$

The problem, facet 3: temporary modifications are not forbidden

Repeating “ $s \rightsquigarrow \text{Seq } L$ ” in the pre- and postcondition can be deceiving.

This does **not** forbid changes to the **concrete** data structure in memory.

Here is a function that really just reads the data structure:

$$\{s \rightsquigarrow \text{Seq } L\} (\text{length } s) \{\lambda y. s \rightsquigarrow \text{Seq } L \star [y = |L|]\}$$

And a function that actually modifies the data structure:

$$\{s \rightsquigarrow \text{Seq } L \star [|L| \leq n]\} (\text{resize } s \ n) \{\lambda(). s \rightsquigarrow \text{Seq } L\}$$

The problem, facet 4: sharing is not permitted

The specification of *append* requires s_1 and s_2 to be *distinct* data structures:

$$\{s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2\}$$

(*append* s_1 s_2)

$$\{\lambda s_3. s_3 \rightsquigarrow \text{Seq } (L_1 ++ L_2) \star s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2\}$$

Indeed, $s \rightsquigarrow \text{Seq } L \star s \rightsquigarrow \text{Seq } L$ is equivalent to *false*.

The problem, facet 4: sharing is not permitted

The specification of *append* requires s_1 and s_2 to be **distinct** data structures:

$$\begin{aligned} & \{s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2\} \\ & (\text{append } s_1 \ s_2) \\ & \{\lambda s_3. s_3 \rightsquigarrow \text{Seq } (L_1 \ ++ \ L_2) \star s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2\} \end{aligned}$$

Indeed, $s \rightsquigarrow \text{Seq } L \star s \rightsquigarrow \text{Seq } L$ is equivalent to *false*.

As a result, to allow sharing, we must establish **another specification**:

$$\begin{aligned} & \{s \rightsquigarrow \text{Seq } L\} \\ & (\text{append } s \ s) \\ & \{\lambda s_3. s_3 \rightsquigarrow \text{Seq } (L \ ++ \ L) \star s \rightsquigarrow \text{Seq } L\} \end{aligned}$$

Duplicate work for us. **Increased complication** and/or duplicate work for the user.

Fractional permissions to the rescue...?

Could sequence concatenation be specified as follows in Concurrent SL?

$$\begin{aligned} & \forall \pi_1, \pi_2. \{ \pi_1 \cdot (s_1 \rightsquigarrow \text{Seq } L_1) \star \pi_2 \cdot (s_2 \rightsquigarrow \text{Seq } L_2) \} \\ & \quad (\text{append } s_1 \ s_2) \\ & \quad \{ \lambda s_3. s_3 \rightsquigarrow \text{Seq } (L_1 \text{ ++ } L_2) \star \pi_1 \cdot (s_1 \rightsquigarrow \text{Seq } L_1) \star \pi_2 \cdot (s_2 \rightsquigarrow \text{Seq } L_2) \} \end{aligned}$$

Yes, if the logic allows **scaling**, $\pi \cdot H$. This requires existential quantification to be restricted so as to be precise (Boyland, 2010).

Without scaling, one must define $s \rightsquigarrow \text{Seq } \pi \ L$.

Fractional permissions to the rescue...?

Could sequence concatenation be specified as follows in Concurrent SL?

$$\begin{aligned} & \forall \pi_1, \pi_2. \{ \pi_1 \cdot (s_1 \rightsquigarrow \text{Seq } L_1) \star \pi_2 \cdot (s_2 \rightsquigarrow \text{Seq } L_2) \} \\ & \quad (\text{append } s_1 \ s_2) \\ & \quad \{ \lambda s_3. s_3 \rightsquigarrow \text{Seq } (L_1 \dashv\vdash L_2) \star \pi_1 \cdot (s_1 \rightsquigarrow \text{Seq } L_1) \star \pi_2 \cdot (s_2 \rightsquigarrow \text{Seq } L_2) \} \end{aligned}$$

Yes, if the logic allows [scaling](#), $\pi \cdot H$. This requires existential quantification to be restricted so as to be precise ([Boyland, 2010](#)).

Without scaling, one must define $s \rightsquigarrow \text{Seq } \pi L$.

This addresses problem facets 3 and 4,

- ▶ but is still [noisy](#),
- ▶ and still requires careful [splitting](#), [threading](#), and [joining](#) of permissions.

“Hiding the fractions” ([Heule et al, 2013](#)) is cool but requires yet more machinery.

In this paper

We propose a solution that is:

- ▶ not as powerful as fractional permissions (or other share algebras),
- ▶ but significantly simpler.

Our contributions:

- ▶ introducing a read-only modality, RO.
RO(H) represents temporary read-only access to the memory governed by H .
- ▶ finding simple and sound reasoning rules for RO.
- ▶ proposing a model that justifies these rules.

Some Intuition

Reasoning Rules

Model

Conclusion

Our solution

We would like the specification of *append* to look like this:

$$\begin{aligned} & \{RO(s_1 \rightsquigarrow \text{Seq } L_1) \star RO(s_2 \rightsquigarrow \text{Seq } L_2)\} \\ & (\text{append } s_1 \ s_2) \\ & \{\lambda s_3. \ s_3 \rightsquigarrow \text{Seq } (L_1 \ ++ \ L_2)\} \end{aligned}$$

Our solution

We would like the specification of *append* to look like this:

$$\begin{aligned} & \{RO(s_1 \rightsquigarrow \text{Seq } L_1) \star RO(s_2 \rightsquigarrow \text{Seq } L_2)\} \\ & (\text{append } s_1 \ s_2) \\ & \{\lambda s_3. s_3 \rightsquigarrow \text{Seq } (L_1 \ ++ \ L_2)\} \end{aligned}$$

Compared with the earlier specification based on unique read-write permissions,

- ▶ this specification is **more concise**,
- ▶ imposes **fewer proof obligations**,
- ▶ makes it clear that **the data structures cannot be modified** by *append*,
- ▶ and **does not require** s_1 and s_2 to be distinct. — (next slide)
- ▶ Furthermore, this spec **implies** the earlier spec. — (slide after next)

Our solution, facet 4: sharing is permitted

The top Hoare triple is the new spec of *append*, where s_1 and s_2 are instantiated with s .

$$\frac{\begin{array}{l} \{ \text{RO}(s \rightsquigarrow \text{Seq } L) \star \text{RO}(s \rightsquigarrow \text{Seq } L) \} \\ (\text{append } s \ s) \\ \{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L \ ++ \ L) \} \end{array}}{\begin{array}{l} \{ \text{RO}(s \rightsquigarrow \text{Seq } L) \} \\ (\text{append } s \ s) \\ \{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L \ ++ \ L) \} \end{array}} \text{CONSEQUENCE}$$

The bottom triple states that, with read-only access to s , *append s s* is **permitted**.

Our solution, facet 4: sharing is permitted

The top Hoare triple is the new spec of *append*, where s_1 and s_2 are instantiated with s .

$$\frac{\begin{array}{l} \{ \text{RO}(s \rightsquigarrow \text{Seq } L) \star \text{RO}(s \rightsquigarrow \text{Seq } L) \} \\ (\text{append } s \ s) \\ \{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L \ ++ \ L) \} \end{array}}{\begin{array}{l} \{ \text{RO}(s \rightsquigarrow \text{Seq } L) \} \\ (\text{append } s \ s) \\ \{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L \ ++ \ L) \} \end{array}} \text{CONSEQUENCE}$$

The bottom triple states that, with read-only access to s , *append* $s \ s$ is **permitted**.

Our solution, facet 5: the earlier specification can be derived

The Hoare triple at the top is the new spec of *append*.

$$\frac{\begin{array}{l} \{ \text{RO}(s_1 \rightsquigarrow \text{Seq } L_1) \star \text{RO}(s_2 \rightsquigarrow \text{Seq } L_2) \} \\ (\text{append } s_1 \ s_2) \\ \{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 \ ++ \ L_2) \} \end{array}}{\text{CONSEQUENCE}} \frac{\begin{array}{l} \{ \text{RO}(s_1 \rightsquigarrow \text{Seq } L_1 \ \star \ s_2 \rightsquigarrow \text{Seq } L_2) \} \\ (\text{append } s_1 \ s_2) \\ \{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 \ ++ \ L_2) \} \end{array}}{\text{READ-ONLY FRAME}} \begin{array}{l} \{ s_1 \rightsquigarrow \text{Seq } L_1 \ \star \ s_2 \rightsquigarrow \text{Seq } L_2 \} \\ (\text{append } s_1 \ s_2) \\ \{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 \ ++ \ L_2) \ \star \ s_1 \rightsquigarrow \text{Seq } L_1 \ \star \ s_2 \rightsquigarrow \text{Seq } L_2 \} \end{array}$$

The triple at the bottom is the earlier spec of *append*.

Our solution, facet 5: the earlier specification can be derived

The Hoare triple at the top is the new spec of *append*.

permission
becomes
read-only

$$\{ \text{RO}(s_1 \rightsquigarrow \text{Seq } L_1) \star \text{RO}(s_2 \rightsquigarrow \text{Seq } L_2) \}$$
$$(\text{append } s_1 \ s_2)$$
$$\{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 \ ++ \ L_2) \}$$

CONSEQUENCE

$$\{ \text{RO}(s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2) \}$$
$$(\text{append } s_1 \ s_2)$$
$$\{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 \ ++ \ L_2) \}$$

READ-ONLY FRAME

$$\{ s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2 \}$$
$$(\text{append } s_1 \ s_2)$$
$$\{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 \ ++ \ L_2) \star s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2 \}$$

The triple at the bottom is the earlier spec of *append*.

Our solution, facet 5: the earlier specification can be derived

The Hoare triple at the top is the new spec of *append*.

$\{ \text{RO}(s_1 \rightsquigarrow \text{Seq } L_1) \star \text{RO}(s_2 \rightsquigarrow \text{Seq } L_2) \}$

$(\text{append } s_1 \ s_2)$

$\{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 \ ++ \ L_2) \}$

CONSEQUENCE

$\{ \text{RO}(s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2) \}$

$(\text{append } s_1 \ s_2)$

$\{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 \ ++ \ L_2) \}$

READ-ONLY FRAME

$\{ s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2 \}$

$(\text{append } s_1 \ s_2)$

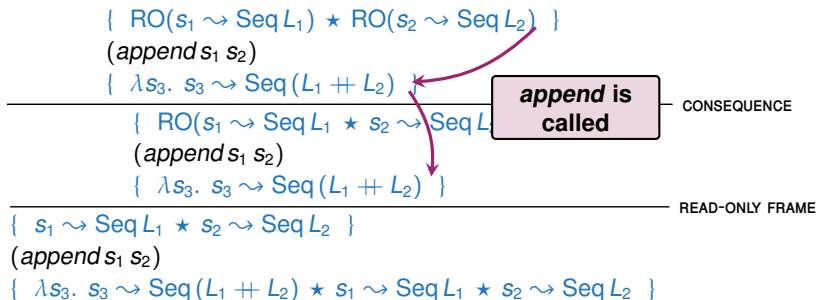
$\{ \lambda s_3. s_3 \rightsquigarrow \text{Seq}(L_1 \ ++ \ L_2) \star s_1 \rightsquigarrow \text{Seq } L_1 \star s_2 \rightsquigarrow \text{Seq } L_2 \}$

The triple at the bottom is the earlier spec of *append*.

**read-only
permission
is weakened**

Our solution, facet 5: the earlier specification can be derived

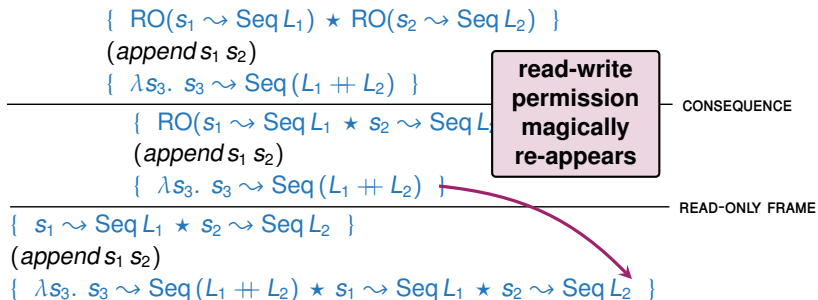
The Hoare triple at the top is the new spec of *append*.



The triple at the bottom is the earlier spec of *append*.

Our solution, facet 5: the earlier specification can be derived

The Hoare triple at the top is the new spec of *append*.



The triple at the bottom is the earlier spec of *append*.

Some Intuition

Reasoning Rules

Model

Conclusion

Permissions

The syntax of permissions is as follows:

$$H := [P] \mid l \hookrightarrow v \mid H_1 \star H_2 \mid H_1 \wp H_2 \mid \exists x. H \mid \text{RO}(H)$$

Every permission H has a read-only form $\text{RO}(H)$.

Properties of RO

Read-only access to a data structure entails read-only access to its parts:

$$\text{RO}(H_1 \star H_2) \triangleright \text{RO}(H_1) \star \text{RO}(H_2) \quad (\text{the reverse is false})$$

Read-only permissions are duplicable (therefore, no need to count them!):

$$\text{RO}(H) = \text{RO}(H) \star \text{RO}(H)$$

Read-only permissions are generally well-behaved:

$$\begin{aligned} \text{RO}([P]) &= [P] \\ \text{RO}(H_1 \wp H_2) &= \text{RO}(H_1) \wp \text{RO}(H_2) \\ \text{RO}(\exists x. H) &= \exists x. \text{RO}(H) \\ \text{RO}(\text{RO}(H)) &= \text{RO}(H) \\ \text{RO}(H) &\triangleright \text{RO}(H') \quad \text{if } H \triangleright H' \end{aligned}$$

A new read axiom

The traditional read axiom:

TRADITIONAL READ AXIOM

$$\{l \hookrightarrow v\} (\text{get } l) \{\lambda y. [y = v] \star l \hookrightarrow v\}$$

is replaced with a “smaller” axiom:

NEW READ AXIOM

$$\{\text{RO}(l \hookrightarrow v)\} (\text{get } l) \{\lambda y. [y = v]\}$$

The traditional axiom can be derived from the new axiom.

A new frame rule

The traditional frame rule is subsumed by a new “read-only frame rule”:

$$\frac{\text{FRAME RULE} \quad \{H\} t \{Q\} \quad \text{normal } H'}{\{H \star H'\} t \{Q \star H'\}} \qquad \frac{\text{READ-ONLY FRAME RULE} \quad \{H \star \text{RO}(H')\} t \{Q\} \quad \text{normal } H'}{\{H \star H'\} t \{Q \star H'\}}$$

This says: upon entry into a block, H' is temporarily replaced with $\text{RO}(H')$, and upon exit, magically re-appears.

The side condition normal H' means roughly that H' has no RO components.

This means that read-only permissions cannot be framed out.

If they could, the read-only frame rule would clearly be unsound. (Exercise!)

How read-only permissions are used

No reasoning rule involves a triple whose postcondition contains RO.

Read-only permissions always appear in **preconditions**, never in postconditions.

They are always **passed down**, never returned.

In fact, in the model, we will see that, in a triple $\{H\} t \{Q\}$:

- ▶ the postcondition applies only to **some read-write fragment** of the final heap;
- ▶ the read-only part of the heap **must be preserved** anyway, so there is no need for the postcondition to describe it.

Some Intuition

Reasoning Rules

Model

Conclusion

Memories & heaps – a simple model of access rights

A **memory** is a finite map of locations to values.

A **heap** h is a pair of two **disjoint** memories $h.f$ and $h.r$.


- ▶ $h.f$ represents the locations to which we have **full access**;
- ▶ $h.r$ represents the locations to which we have **read-only access**.

An **assertion**, or **permission**, is a predicate over heaps (or: a set of heaps).

Heap composition & separating conjunction

The combination of two heaps:

$$h_1 + h_2 = (h_1.f \uplus h_2.f, h_1.r \cup h_2.r)$$




is defined only if:

Heap composition & separating conjunction

The combination of two heaps:

$$h_1 + h_2 = (h_1.f \uplus h_2.f, h_1.r \cup h_2.r)$$




is defined only if:

- ▶ the read-write components $h_1.f$ and $h_2.f$ are **disjoint**;

Heap composition & separating conjunction

The combination of two heaps:

$$h_1 + h_2 = (h_1.f \uplus h_2.f, h_1.r \cup h_2.r)$$




is defined only if:

- ▶ the read-write components $h_1.f$ and $h_2.f$ are **disjoint**;
- ▶ the read-only components $h_1.r$ and $h_2.r$ **agree** where they overlap;

Heap composition & separating conjunction

The combination of two heaps:

$$h_1 + h_2 = (h_1.f \uplus h_2.f, h_1.r \cup h_2.r)$$




is defined only if:

- ▶ the read-write components $h_1.f$ and $h_2.f$ are **disjoint**;
- ▶ the read-only components $h_1.r$ and $h_2.r$ **agree** where they overlap;
- ▶ the read-write component $h_1.f$ is **disjoint** with the read-only component $h_2.r$, and vice-versa.

Heap composition & separating conjunction

The combination of two heaps:

$$h_1 + h_2 = (h_1.f \uplus h_2.f, h_1.r \cup h_2.r)$$




is defined only if:

- ▶ the read-write components $h_1.f$ and $h_2.f$ are **disjoint**;
- ▶ the read-only components $h_1.r$ and $h_2.r$ **agree** where they overlap;
- ▶ the read-write component $h_1.f$ is **disjoint** with the read-only component $h_2.r$, and vice-versa.

With this in mind, **separating conjunction** is interpreted as usual:

$$H_1 \star H_2 = \lambda h. \exists h_1 h_2. (h_1 + h_2 \text{ is defined}) \wedge h = h_1 + h_2 \wedge H_1 h_1 \wedge H_2 h_2$$


The read-only modality



$\text{RO}(H)$ is interpreted as follows:

$$\text{RO}(H) = \lambda h. (h.f = \emptyset) \wedge \exists h'. (h.r = h'.f \uplus h'.r) \wedge H h'$$

The read-only modality




$RO(H)$ is interpreted as follows:

$$RO(H) = \lambda h. (h.f = \emptyset) \wedge \exists h'. (h.r = h'.f \uplus h'.r) \wedge H h'$$

This means:

- ▶ we have write access to **nothing**.

The read-only modality



$RO(H)$ is interpreted as follows:

$$RO(H) = \lambda h. (h.f = \emptyset) \wedge \exists h'. (h.r = h'.f \uplus h'.r) \wedge H h'$$

This means:

- ▶ we have write access to **nothing**.
- ▶ if we had write access to certain locations for which we have read access, then H would hold.

The rest of the connectives

$$[P] = \lambda h. (h.f = \emptyset) \wedge (h.r = \emptyset) \wedge P$$

$$l \leftrightarrow v = \lambda h. (h.f = (l \mapsto v)) \wedge (h.r = \emptyset)$$

$$H_1 \wp H_2 = \lambda h. H_1 h \vee H_2 h$$

$$\exists x. H = \lambda h. \exists x. H h$$

$$\text{normal}(H) = \forall h. H h \Rightarrow h.r = \emptyset$$

Interpretation of triples


The meaning of the Hoare triple $\{H\} t \{Q\}$ is as follows:

$$\forall h_1 h_2. \left\{ \begin{array}{l} h_1 + h_2 \text{ is defined} \\ H h_1 \end{array} \right\} \Rightarrow \exists v h'_1. \left\{ \begin{array}{l} h'_1 + h_2 \text{ is defined} \\ t / [h_1 + h_2] \Downarrow v / [h'_1 + h_2] \\ h'_1.r = h_1.r \\ \text{on-some-rw-frag}(Q v) h'_1 \end{array} \right\}$$

What's nonstandard?

Interpretation of triples

The meaning of the Hoare triple $\{H\} t \{Q\}$ is as follows:

$$\forall h_1 h_2. \left\{ \begin{array}{l} h_1 + h_2 \text{ is defined} \\ H h_1 \end{array} \right\} \Rightarrow \exists v h'_1. \left\{ \begin{array}{l} h'_1 + h_2 \text{ is defined} \\ t / \lfloor h_1 + h_2 \rfloor \Downarrow v / \lfloor h'_1 + h_2 \rfloor \\ h'_1.r = h_1.r \\ \text{on-some-rw-frag}(Q v) h'_1 \end{array} \right\}$$


What's nonstandard?

- ▶ The read-only part of the heap must be **preserved**.

Interpretation of triples

The meaning of the Hoare triple $\{H\} t \{Q\}$ is as follows:

$$\forall h_1 h_2. \left\{ \begin{array}{l} h_1 + h_2 \text{ is defined} \\ H h_1 \end{array} \right\} \Rightarrow \exists v h'_1. \left\{ \begin{array}{l} h'_1 + h_2 \text{ is defined} \\ t / \lfloor h_1 + h_2 \rfloor \Downarrow v / \lfloor h'_1 + h_2 \rfloor \\ h'_1.r = h_1.r \\ \text{on-some-rw-frag}(Q v) h'_1 \end{array} \right\}$$

What's nonstandard?

- ▶ The read-only part of the heap must be **preserved**.
- ▶ The postcondition describes only **a read-write fragment of the final heap**.

$\text{on-some-rw-frag}(H) =$

$$\lambda h. \exists h_1 h_2. (h_1 + h_2 \text{ is defined}) \wedge h = h_1 + h_2 \wedge h_1.r = \emptyset \wedge H h_1$$

Soundness

Theorem

With respect to this interpretation of triples, every reasoning rule is sound.

Proof.

“Straightforward”. Machine-checked.



Some Intuition

Reasoning Rules

Model

Conclusion

We propose:

- ▶ a **simple extension** of Separation Logic with a read-only modality;
- ▶ a **simple model** that explains why this is sound.

We believe that temporary read-only permissions sometimes help state more **concise**, **accurate**, **useful** specifications, and lead to simpler proofs.

Possible future work: an implementation in CFML (Charguéraud).

Amnesia (1/2)

Suppose *population* has this “RO” specification:

$$\{\text{RO}(h \rightsquigarrow \text{HashTable } M)\} (\text{population } h) \{\lambda y. [y = \text{card } M]\}$$

Suppose a hash table is a mutable record whose *data* field points to an array:

$$\begin{aligned} h \rightsquigarrow \text{HashTable } M &:= \\ &\exists! a. \exists! L. (h \rightsquigarrow \{\text{data} = a; \dots\} \star a \rightsquigarrow \text{Array } L \star \dots) \end{aligned}$$

Suppose there is an operation *foo* on hash tables:

```
let foo h =  
  let d = h.data in           – read the address of the array  
  let p = population h in    – call population  
  ...
```

If “RO” is sugar for repeating $h \rightsquigarrow \text{HashTable } M$ in the pre and post, then the proof of *foo* runs into a problem...

Amnesia (2/2)

Reasoning about *foo* might go like this:

```
1 let foo h =
2   {h ~ HashTable M}                                – foo's precondition
3   {h ~ {data = a; ...} ★ a ~ Array L ★ ...}        – by unfolding
4   let d = h.data in
5   {h ~ {data = a; ...} ★ a ~ Array L ★ ... ★ [d = a]} – by reading
6   {h ~ HashTable M ★ [d = a]}                     – by folding
7   let p = population h in                          – we have to fold
8   {h ~ HashTable M ★ [d = a] ★ [p = #M]}
9   ...
```

At line 8, the equation $d = a$ is useless.

We have **forgotten** what d represents, and **lost the benefit** of the read at line 4.

If “RO” is sugar, the specification of *population* is **weaker** than it seems.

If “RO” is native, there is a way around this problem. (Details omitted.)