

Temporary Read-Only Permissions for Separation Logic

Making Separation Logic's
Small Axioms
Smaller

Arthur Charguéraud François Pottier



LTP meeting
Saclay, November 28, 2016

Motivation

More Motivation

Separation Logic with Read-Only Permissions

Separation and ownership

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

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

Therefore, it is about **unique ownership**.

- ▶ if we don't own a memory cell, we cannot write it, or even read it.
- ▶ if we own it, we can read **and** write it.

We have either **no permission** or **read-write permission**.

The read and write axioms

The reasoning rule for writing requires and returns a unique permission :

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

So does the reasoning rule for reading :

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

They are known as “small axioms”.

But are they as small as they could be ? ...

Consequences

From memory cells and arrays,
the dichotomy extends to **user-defined data structures**.

For every data structure, we have either **no permission** or **read-write permission**.

Consequences

Here a specification of an array concatenation function :

$$\{a_1 \rightsquigarrow \text{Array } L_1 \star a_2 \rightsquigarrow \text{Array } L_2\}$$
$$(\text{Array.append } a_1 \ a_2)$$
$$\{\lambda a_3. a_3 \rightsquigarrow \text{Array } (L_1 \ ++ \ L_2) \star a_1 \rightsquigarrow \text{Array } L_1 \star a_2 \rightsquigarrow \text{Array } L_2\}$$

It is a bit **noisy**.

It also has several deeper drawbacks (see next slide).

Our goal

We would like the specification to look like this instead :

$$\begin{aligned} & \{RO(a_1 \rightsquigarrow \text{Array } L_1) \star RO(a_2 \rightsquigarrow \text{Array } L_2)\} \\ & (\text{Array.append } a_1 \ a_2) \\ & \{\lambda a_3. a_3 \rightsquigarrow \text{Array } (L_1 \ ++ \ L_2)\} \end{aligned}$$

This would be **more concise**,
require **less bookkeeping**,
make it clear that **the arrays are unmodified**,
and in fact **would not require the arrays to be distinct**.

Our means

For this purpose, we introduce
temporary read-only permissions.

Thank you for your attention.



Remboursez !

What!?

Remboursez !

What!?

Couldn't one view $RO(\cdot)$ as syntactic sugar?

- ▶ No.

Remboursez !

What!?

Couldn't one view $RO(\cdot)$ as **syntactic sugar** ?

- ▶ No.

Couldn't one express this using **fractional permissions** ?

- ▶ Yes. More heavily.

Remboursez !

What!?

Couldn't one view $RO(\cdot)$ as **syntactic sugar** ?

- ▶ No.

Couldn't one express this using **fractional permissions** ?

- ▶ Yes. More heavily.

Isn't the metatheory of $RO(\cdot)$ **very simple** ?

- ▶ Yes, it is. If and once you get it right. That's the point !

Motivation

More Motivation

Separation Logic with Read-Only Permissions

The sugar hypothesis



The sugar hypothesis

Could the Hoare triple :

$$\{RO(H_1) \star H_2\} t \{Q\}$$

be syntactic sugar for :

$$\{H_1 \star H_2\} t \{H_1 \star Q\}$$

?

Sugar does not reduce work

Sugar reduces **apparent** redundancy in specifications,
but has no effect on the proof obligations,
so **does not reduce** redundancy and bookkeeping in proofs.

If we must prove this :

$$\{H_1 \star H_2\} t \{H_1 \star Q\}$$

then we must work to ensure and argue that the permission H_1 is returned.

If “RO” was native, proving $\{RO(H_1) \star H_2\} t \{Q\}$ would require no such work.

Sugar does not allow aliasing

If “RO” is sugar, then this specification requires a_1 and a_2 to be disjoint arrays :

$$\begin{aligned} & \{ \text{RO}(a_1 \rightsquigarrow \text{Array } L_1) \star \text{RO}(a_2 \rightsquigarrow \text{Array } L_2) \} \\ & (\text{Array}.\text{append } a_1 \ a_2) \\ & \{ \lambda a_3. a_3 \rightsquigarrow \text{Array } (L_1 \text{ ++ } L_2) \} \end{aligned}$$

As a result, we must prove another specification to allow aliasing :

$$\begin{aligned} & \{ a \rightsquigarrow \text{Array } L \} \\ & (\text{Array}.\text{append } a \ a) \\ & \{ \lambda a_3. a_3 \rightsquigarrow \text{Array } (L \text{ ++ } L) \star a \rightsquigarrow \text{Array } L \} \end{aligned}$$

Duplicate work for us ; increased complication for the user.

If “RO” was native and duplicable, the first spec above would allow aliasing.

Sugar is deceptive

A read-only function admits an “RO” specification.

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

If “RO” is sugar, a function that **can** have an effect **also** admits an “RO” spec.

$$\{\text{RO}(h \rightsquigarrow \text{HashTable } M)\} (\textit{resize } h) \{\lambda(). []\}$$

An “RO” specification, interpreted as sugar, does **not** mean “read-only”.

Such sugar, if adopted, should use another keyword, e.g., **preserves**.

If “RO” was native, *resize* would not admit the second spec above.

Sugar causes amnesia and weakness

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, then the proof of *foo* runs into a problem...

Sugar causes amnesia and weakness

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.

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

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

Motivation

More Motivation

Separation Logic with Read-Only Permissions

Permissions

Permissions are as follows :

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

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

Properties of RO

RO is well-behaved :

$$\begin{aligned} \text{RO}([P]) &= [P] \\ \text{RO}(H_1 \star H_2) &\triangleright \text{RO}(H_1) \star \text{RO}(H_2) && \text{(the reverse is false)} \\ \text{RO}(H_1 \vee H_2) &= \text{RO}(H_1) \vee \text{RO}(H_2) \\ \text{RO}(\exists x. H) &= \exists x. \text{RO}(H) \\ \text{RO}(\text{RO}(H)) &= \text{RO}(H) \\ \text{RO}(H) &\triangleright \text{RO}(H') && \text{if } H \triangleright H' \\ \text{RO}(H) &= \text{RO}(H) \star \text{RO}(H) \end{aligned}$$

A new read axiom

The traditional read axiom :

TRADITIONAL READ AXIOM

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

is replaced with a “smaller” axiom :

NEW READ AXIOM

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

A new frame rule

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

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

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 “ H' has no RO components”, so $\text{RO}(H')$ cannot escape through Q .

That's all, folks !

That's all there is to it !

That's all, folks !

The paper gives a simple **model** that explains why the logic is sound.

The proof is machine-checked.

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.