

Exploring limits of Rely-Guarantee approach

Cliff Jones
(joint work with with Nisansala Yatapanage)

Newcastle University

WG 2.3 Providence, RI
2018-05-07

Example: *GC*

NB: it is the underlying points, not the example, that matter

- concurrent Garbage Collection: Ben Ari [BA84]
- an 'Owicki/Gries' proof by Jan van de Snepscheut [vdS87]
- illustrates:
 - asymmetric processes (and therefore R/G):
Collector || *Mutator*
 - three *Mutator* options: *Redirect*/*Malloc*/*Zap*
- highlights a limitation of R/G
 - lets us explore options

A (second) data reification introduces marking

$$\Sigma_2 \text{ :: } \begin{array}{l} \text{roots: Addr-}\mathbf{set} \\ \text{hp: Heap} \\ \text{free: Addr-}\mathbf{set} \\ \text{marked: Addr-}\mathbf{set} \end{array}$$

where

$$\text{inv-}\Sigma_2(\text{mk-}\Sigma_2(\text{roots}, \text{hp}, \text{free}, \text{marked})) \triangleq \dots$$
$$\text{Heap} = \text{Addr} \xrightarrow{m} \text{Node}$$
$$\text{Node} = [\text{Addr}]^*$$

Mark trivial version (non-interfering)

Collector \triangleq (Unmark; *Mark*; Sweep)

Mark

wr *marked*

rd *hp, roots, free*

pre **true**

rely $\textit{marked}' = \textit{marked} \wedge \textit{free}' = \textit{free} \wedge \textit{hp}' = \textit{hp}$

guar **true**

post $\textit{marked}' = (\textit{free} \cup \textit{reach}(\textit{roots}, \textit{hp}))$

transition from sequential to interfering shows an R/G 'pattern'

Code for *Mark* (all three versions)

```
Mark  $\triangleq$   
  repeat  
     $mc \leftarrow \mathbf{card\ marked}$ ;  
    Propagate  
  until  $\mathbf{card\ marked} = mc$ 
```

```
Propagate:  
   $consid \leftarrow \{\}$ ;  
  do while  $consid \neq Addr$   
    let  $x \in (Addr - consid)$  in  
    if  $x \in \mathbf{marked}$  then  $\mathbf{Mark-kids}(x)$  else skip;  
     $consid \leftarrow consid \cup \{x\}$   
  od
```

proofs of (sequential) loops are straightforward — but differ

Collector has to enlist *Mutator* to help with marking process!

now:

Unmark

wr *marked*

rd *roots, free*

pre *true*

rely $free' \subseteq free$

guar $marked' \subseteq marked$

post $\forall a \in (Addr - (roots \cup free)) \cdot \exists m \in \overbrace{marked} \cdot a \notin m$

simplification: upper bound of marking (lower of GC) ignored for now

Possible values

- in the presence of interference

$$\{P\}x \leftarrow y\{x' = y \vee x' = y'\}$$

is not enough — instead we now write:

$$\{P\}x \leftarrow y\{x' \in \widehat{y}\}$$

- concept identified in [JP11]
 - development of Simpson's '4-slot' algorithm for ACMs
 - (asymmetric *WRITE* || *READ*)
- possible values ('posvals') extends expressive power
 - avoids the 'need' for some ghost variables
 - [JH16] has a rather clear specification of ACMs

$$r \leftarrow \text{READ}(): r \in \widehat{\text{buf}}$$

$$\text{WRITE}(x): \text{buf}' \neq \text{buf} \Rightarrow \text{buf}' = x$$

Mark simplified version (atomic)

Pretend *Mutator(Redirect)* can make change/mark atomically

$$\langle hp(a), \text{marked} \leftarrow hp(a) \dagger \{i \mapsto b\}, \text{marked} \cup \{b\} \rangle$$

code shown for brevity – paper has R/G specification

Mark

wr *marked*

rd *roots, hp, free*

pre true

rely $\text{marked} \subseteq \text{marked}' \wedge$

$$\forall (a, i) \in \text{dom } hp \cdot hp'(a, i) \neq hp(a, i) \Rightarrow hp'(a, i) \in \text{marked}'$$

guar $\text{marked} \subseteq \text{marked}'$

post $\text{reach}(\text{roots}, hp') \subseteq \text{marked}'$

slightly more complicated but previous proof *strategy* works

with exactly the same code — different R/G

Concurrent version (finally)

- the *Mutator* cannot redirect/mark atomically
 $hp(a) \leftarrow hp(a) \dagger \{i \mapsto b\}; \text{marked} \leftarrow \text{marked} \cup \{b\}$
- I have worked on ‘fiction of atomicity/atomicity refinement’ but I don’t think that’s the right approach here
- the essence of the correctness is a ‘three state’ argument:
Mutator can change $hp(a, i)$ to point to b
... then go to sleep (before marking b)
if the *Collector* has passed a , *post-Propagate* could fail
however there must be another pointer (say $hp(c, j)$) to b
if b not marked by *Collector*, c is still to be handled
BUT what if the *Mutator* destroys the $hp(c, j)$ link first?
... to do so, the *Mutator* must have marked b — phew!
- such a three state argument is not expressible as rely

Various ways to undermine compositionality

- is R/G expressively weak? intentionally so!
- minimal information about the (concurrent) environment
 - R/G facilitates this
 - interesting examples: asymmetric concurrent processes
- in Owicki/Gries [Owi75] one reasons about the final code!
- 'ghost' variables dent/destroy compositionality
 - in the extreme, they can dictate the entire environment code
e.g. $(PC = 1 \Rightarrow x = 5) \wedge (PC = 2 \Rightarrow x = 6) \wedge \dots$
 - in many cases, such 'auxiliary' variables can be avoided
often with an apposite abstraction
 - if I have to use them, I want a test
(cf. homomorphic rule vs. 'biased' specifications)

Concurrent version: alternative (i)

concede a ghost variable! (as in the SETTA paper)

$tbm: [Addr]$

$\langle hp(a), tbm \leftarrow hp(a) \dagger \{i \mapsto b\}, b \rangle;$
 $\langle marked, tbm \leftarrow marked \cup \{tbm\}, \mathbf{nil} \rangle$

but we do have a test!

$rely\text{-}Collector_c : \Sigma_2 \times \Sigma_2 \rightarrow \mathbb{B}$

$rely\text{-}Collector_c(\sigma, \sigma') \triangleq$
 $free' \subseteq free \wedge marked \subseteq marked' \wedge$
 $(\forall (a, i) \in hp \cdot$
 $hp'(a, i) \neq hp(a, i) \wedge hp'(a, i) \in Addr \Rightarrow$
 $hp'(a, i) \in marked' \vee tbm' = hp'(a, i)) \wedge$
 $(tbm \neq \mathbf{nil} \wedge tbm' \neq tbm \Rightarrow tbm \in marked' \wedge tbm' = \mathbf{nil})$

Concurrent version: more alternatives

- reject the restrictions of R/G and use Temporal Logic?
 - I worry that RGITL [STER11] is too expressive!
 - slippery slope to arguing about the combined code
- could specify that the *Mutator* preserves *can-be-completed*
- also damages compositionality! the designer of *Mutator* ...

(My) current preferred alternative

again, code as abbreviation for R/G specifications

$$mr-1: \langle hp(a) \leftarrow hp(a) \dagger \{i \mapsto b\} \rangle$$
$$mr-2: \langle marked \leftarrow marked \cup \{b\} \rangle$$

two roles for *tbm* in *rely-Collector_c*: remember value + mark point

Useful Lemma:

$$tbm \neq \mathbf{nil} \Rightarrow$$
$$\exists \{(a, i), (b, j)\} \subseteq \mathbf{dom} \, hp \cdot (a, i) \neq (b, j) \wedge hp(a, i) = hp(b, j) = tbm$$

So, at the point between *mr-1*/*mr-2*: rely on second path
but this is a 'local' rely — cf. [JH16]

Clearly this argument also limits compositionality — less so?

- compositionality is difficult to achieve for concurrency
 - see [dR01]
 - interference is the essence of concurrency
 - (even with process algebras!)
- recording interference (e.g. Rely/Guarantee)
 - with symmetric processes (e.g. $SIEVE = ||_i REM(i)$)
 - more interesting with asymmetric processes: different R/Gs
 - e.g. *4-slot*, *QREL* (aka 'union/find'), *GC*

Conclusions

- compositionality = good engineering
 - 'compositionality' wrt a notation?
- R/G can't specify GC fully compositionally
 - but the (related) versions of R/Gs are nice!
 - '*Collector/Mutator* were designed together' Simon Doherty
- adding 'ghost variables'
 - the 3-state argument looks like a clear test
- the jury is still out on the least damaging solution
 - actually, I think all candidate workarounds tell us something
- (for me) examples are essential!
 - N.B. two ways of using \hat{y}
- comments?



Mordechai Ben-Ari.

Algorithms for on-the-fly garbage collection.

ACM Transactions on Programming Languages and Systems, 6(3):333–344, 1984.



W.-P. de Roever.

Concurrency Verification: Introduction to Compositional and Noncompositional Methods.

Cambridge University Press, 2001.



Cliff B. Jones and Ian J. Hayes.

Possible values: Exploring a concept for concurrency.

Journal of Logical and Algebraic Methods in Programming, 2016.



Cliff B. Jones and Ken G. Pierce.

Elucidating concurrent algorithms via layers of abstraction and reification.

Formal Aspects of Computing, 23(3):289–306, 2011.



S. Owicki.

Axiomatic Proof Techniques for Parallel Programs.

PhD thesis, Department of Computer Science, Cornell University, 1975.



G. Schellhorn, B. Tofan, G. Ernst, and W. Reif.

Interleaved programs and rely-guarantee reasoning with ITL.

In *Temporal Representation and Reasoning (TIME)*, 2011 Eighteenth International Symposium on, pages 99–106, 2011.



Jan LA van de Snepscheut.

“Algorithms for on-the-fly garbage collection” revisited.

Information Processing Letters, 24(4):211–216, 1987.

(Concurrent) garbage collection: abstract specification

$\Sigma_0 :: \text{busy: Addr-set}$
 free: Addr-set

where

$\text{inv-}\Sigma_0(\text{mk-}\Sigma_0(\text{busy}, \text{free})) \triangleq \text{busy} \cap \text{free} = \{\}$

the invariant sets an upper bound for GC

*GC: Collector**

Collector

wr *free*

rd *busy*

pre true

...

post $(\text{Addr} - \text{busy}) \subseteq \widehat{\cup} \text{free}$

lower bound for GC

First data reification — gets us to:

$\Sigma_1 ::$ *roots*: **Addr-set**
hp: *Heap*
free: **Addr-set**

where

$inv\text{-}\Sigma_1(mk\text{-}\Sigma_1(\text{roots}, hp, free)) \triangleq$
dom *hp* = *Addr* \wedge
free \cap *reach*(*roots*, *hp*) = { }

upper bound for GC

Heap = *Addr* \xrightarrow{m} *Node*

Node = *Addr**

Simplification: here **nil** is ignored