

# Semi-Automated Reasoning About Non-Determinism in C Expressions

Léon Gondelman

joint work with Dan Frumin and Robbert Krebbers

Radboud University Nijmegen

23 February, 2019 @ EUTypes, Krakow, Poland  
(in relation with Lambda Days 2019)

```
int main() {  
    int x;  
    int y = (x = 3) + (x = 4);  
    printf("%d, %d\n", x, y);  
}
```

What is the expected outcome of this program ?

```
int main() {  
    int x;  
    int y = (x = 3) + (x = 4);  
    printf("%d, %d\n", x, y);  
}
```

a small experiment with [existing compilers](#) gives

compiler	outcome	warnings
compcert	4, 7	no
clang	4, 7	yes
gcc-4.9	4, 8	no

```
int main() {  
    int x;  
    int y = (x = 3) + (x = 4);  
    printf("%d, %d\n", x, y);  
}
```

according to C standard, the program is allowed do... **anything**,  
it is even allowed to **crash**

this program violates the sequence point restriction:

- the order of evaluation in C expressions is **unspecified**
- concurrent memory access is allowed
- multiple unsequenced modifications result in **undefined behavior**

**the problem:** sequence point violations may cause a C program to crash or to have arbitrary results

**the goal:** guarantee the absence of undefined behavior in a given C program for any evaluation order

**in this talk:**

- (1) use concurrent separation logic to reason about C  
(previous work, Krebbers POPL'14)
- (2) turn it into a semi-automated reasoning procedure  
(our contributions)

**observation:** view non-determinism through **concurrency**

**idea:** use the **concurrent separation logic**

$$\frac{\{P_1\} e_1 \{\Psi_1\} \quad \{P_2\} e_2 \{\Psi_2\} \quad \forall v_1 v_2. \Psi_1 v_1 * \Psi_2 v_2 \vdash \Phi(w_1 \llbracket \odot \rrbracket w_2)}{\{P_1 * P_2\} e_1 \odot e_2 \{\Phi\}}$$

using the rules of this logic we can

- split the memory resources into two disjoint parts
- independently prove that each subexpression executes safely

**observation:** view non-determinism through **concurrency**

**idea:** use the **concurrent separation logic**

$$\frac{\{P_1\} e_1 \{\Psi_1\} \quad \{P_2\} e_2 \{\Psi_2\} \quad \forall v_1 v_2. \Psi_1 v_1 * \Psi_2 v_2 \vdash \Phi(w_1 \llbracket \odot \rrbracket w_2)}{\{P_1 * P_2\} e_1 \odot e_2 \{\Phi\}}$$

**limitations:**

- no support for automation
- difficult to conduct even a manual proof in Coq

instead of Hoare triples, we model our program logic using **weakest precondition calculus**

$$\text{wp } e \{ \Phi \}$$

- $e$  is safe (has defined behavior),
- if  $e$  terminates with a value  $v$ , then  $v$  satisfies the predicate  $\Phi$

the non-determinism is reflected in a similar but more concise way:

$$\frac{\text{wp } e_1 \{ \Psi_1 \} \quad \text{wp } e_2 \{ \Psi_2 \} \quad (\forall w_1 w_2. \Psi_1 w_1 * \Psi_2 w_2 \rightarrow \Phi(w_1 \llbracket \odot \rrbracket w_2))}{\text{wp } (e_1 \odot e_2) \{ \Phi \}}$$



a possible candidate for the load operation:

$$\frac{\text{wp } e \{l. \exists w. l \mapsto w * (l \mapsto w -* \Phi w)\}}{\text{wp } (*e) \{\Phi\}}$$

a possible candidate for the load operation:

$$\frac{\text{wp } e \{l. \exists w. l \mapsto w * (l \mapsto w -* \Phi w)\}}{\text{wp } (*e) \{\Phi\}}$$

**too weak:** does not allow sharing e.g.,  $*1 + *1$

**fractional permissions** enable sharing of resources:

$$1 \xrightarrow{q_1+q_2} v \dashv\vdash 1 \xrightarrow{q_1} v * 1 \xrightarrow{q_2} v$$

so multiple subexpressions **can safely read** from the same location

the rule for load becomes

$$\frac{\text{wp } e \left\{ 1. \exists w q. 1 \xrightarrow{q} w * (1 \xrightarrow{q} w -* \Phi w) \right\}}{\text{wp } (*e) \{ \Phi \}}$$

so we can now prove programs like  $*1 + *1$

a possible candidate for the assignment operation:

$$\frac{\text{wp } e_1 \{ \Psi_1 \} \quad \text{wp } e_2 \{ \Psi_2 \} \quad (\forall l \ w. \Psi_1 \ l * \Psi_2 \ w * \exists v. l \xrightarrow{1} v * (l \xrightarrow{1} w * \Phi \ w))}{\text{wp } (e_1 = e_2) \{ \Phi \}}$$

a possible candidate for the assignment operation:

$$\frac{\text{wp } e_1 \{ \Psi_1 \} \quad \text{wp } e_2 \{ \Psi_2 \} \quad (\forall l \ w. \Psi_1 \ l * \Psi_2 \ w \rightarrow \exists v. l \xrightarrow{1} v * (l \xrightarrow{1} w \rightarrow \Phi \ w))}{\text{wp } (e_1 = e_2) \{ \Phi \}}$$

**unsound:** does not account for sequence point violations

for example, we could verify programs like  $1 = (1 = 3)$

to account for sequence point violations, we decorate fractional permissions with two **access levels** :

$$l \xrightarrow{q}_{\xi} v, \quad \xi \in \{L, U\}$$

- permission  $l \xrightarrow{q}_U v$  states that the location is **unlocked**, so one can read from/write to the location  $l$
- permission  $l \xrightarrow{q}_L v$  states that the location has been **locked**, someone is **already** writing to it, so reads/writes are **forbidden**

the rule for assignment becomes

$$\frac{\text{wp } e_1 \{ \Psi_1 \} \quad \text{wp } e_2 \{ \Psi_2 \} \quad (\forall l \ w. \Psi_1 \ l * \Psi_2 \ w \rightarrow \exists v. l \xrightarrow{1} v * (l \xrightarrow{1} w \rightarrow \Phi \ w))}{\text{wp } (e_1 = e_2) \{ \Phi \}}$$

programs like  $1 = (1 = 3)$  cannot be verified any more

**remark:** we want to access locked pointers later again

$$l = 4; *l$$

we use the **unlocking modality**  $\mathbb{U}$  that unlocks all locked locations at the sequence point :

$$\frac{\text{wp } e_1 \{ \dots \mathbb{U}(\text{wp } e_2 \{ \Phi \}) \}}{\text{wp } (e_1 ; e_2) \{ \Phi \}}$$

$$\frac{l \xrightarrow{q}_L v}{\mathbb{U}(l \xrightarrow{q}_U v)}$$

$$\frac{P * Q}{\mathbb{U}P * \mathbb{U}Q}$$



## reasoning about programs

usually we prove programs **assuming** some logical context:

$$P \vdash \text{wp } e \{ \Phi \}$$

we intertwine the **application of wp rules** with other **logical steps** (*splitting resources, discharging side-conditions, ...*)

manual proof quickly becomes tedious even for small programs

e.g., to reason about binary operators we have to

- infer manually the intermediate postconditions
- subdivide resources all the time

turn program logic into an algorithm procedure  
using a novel **symbolic execution** algorithm:

**input**

precondition

program

--&gt;

**output**

postcondition

value

**frame** = resources not used

$l \mapsto v1 * k \mapsto v2 * r \mapsto v3$

-----

$l = *k + 10$

**postcondition:**  $\top$

**frame:**  $\top$

$l \mapsto v1 * k \mapsto v2 * r \mapsto v3$

-----  
 $l = v2 + 10$

**postcondition:**  $k \xrightarrow{0.5} v2$

**frame:**  $k \xrightarrow{0.5} v2$

$$\cancel{l \mapsto v1} * \cancel{k \mapsto v2} * r \mapsto v3$$

-----

$$v2 + 10$$

**postcondition:**  $k \xrightarrow{0.5} v2 * 1 \mapsto_L (v2 + 10)$

**frame:**  $k \xrightarrow{0.5} v2$

~~$l \mapsto v1 * k \mapsto v2 * r \mapsto v3$~~ 

-----

 $v2 + 10$ 

**postcondition:**  $k \xrightarrow{0.5} v2 * l \mapsto_L (v2 + 10)$

**frame:**  $k \xrightarrow{0.5} v2 * r \mapsto v3$

## example (continued)

$l \mapsto v1 * k \mapsto v2 * r \mapsto v3$

-----

$(l = *k + 10) + (r = *k + 10)$

**postcondition:**  $\top$

**frame:**  $\top$

after executing the LHS

~~$l \mapsto v1 * k \mapsto v2 * r \mapsto v3$~~

-----  
 $(v2 + 10) + (r = *k + 10)$

**postcondition:**  $k \xrightarrow{0.5} v2 * l \mapsto_L (v2 + 10)$

**frame:**  $k \xrightarrow{0.5} v2 * r \mapsto v3$



before executing the RHS

~~$l \mapsto v1$~~  \*  $k \xrightarrow{0.5} v2$  \*  ~~$r \mapsto v3$~~

-----  
 $(v2 + 10) + (r = *k + 10)$

**postcondition:**  $k \xrightarrow{0.5} v2$  \*  $l \mapsto_L (v2 + 10)$

**frame:**  ~~$k \xrightarrow{0.5} v2$~~  \*  ~~$r \mapsto v3$~~

## executing the RHS

~~$l \mapsto v1 * k \mapsto v2 * r \mapsto v3$~~ <sup>0.5</sup>

-----  
 $(v2 + 10) + (r = *k + 10)$

**postcondition:**  $k \xrightarrow{3/4} v2 * 1 \mapsto_L (v2 + 10)$

**frame:**  $k \xrightarrow{1/4} v2 * \del{r \mapsto v3}$

final result

$$\cancel{l \mapsto v1} * \cancel{k \xrightarrow{0.5} v2} * \cancel{r \mapsto v3}$$

-----

$$(v2 + 10) + (v2 + 10)$$

**postcondition:**  $k \xrightarrow{3/4} v2 * l \mapsto_L (v2 + 10) * r \mapsto_L (v2 + 10)$

**frame:**  $k \xrightarrow{1/4} v2 * \cancel{r \mapsto (v2 + 10)}$

our symbolic execution algorithm is a partial function  
restricted to *symbolic heaps* ( $m \in \text{sheap}$ ):

$$\text{forward} : (\text{sheap} \times \text{expr}) \rightarrow (\text{val} \times \text{sheap} \times \text{sheap})$$

satisfying the following specification:


$$\frac{\text{forward}(m, e) = (w, m_1^o, m_1)}{\llbracket m \rrbracket \vdash \text{wp } e \{v. v = w * \llbracket m_1^o \rrbracket\} * \llbracket m_1 \rrbracket}$$

symbolic execution helps to make the wp rules algorithmic  
but the algorithm itself may fail for several reasons:

- the program is not of the right shape
- the precondition is not a symbolic heap
- needed resource is missing in the precondition

to turn the program logic into an **automated procedure**  
we integrate the symbolic executor algorithm into a  
**verification condition generator** (vcgen)

design an **interactive** verification condition generator



vcgen automates the proof **as long as** forward does not fail,  
and when forward fails,

- vcgen **returns** to the user a partially solved goal
- from which it can be **called back** after the user helped out

### main message:

*symbolic execution with frames is a key to enable  
a semi-automated about non-determinism in C  
in an interactive theorem prover*

### other contributions:

- a definitional semantics to a fragment of C in Coq
- soundness proof for symbolic executor and vcgen
- development built on top of the Iris framework

**thank you !**