

Mechanised Owicki-Gries Proofs for C11

Brijesh Dongol
University of Surrey

Joint work with

Sadegh Dalvandi (University of Surrey)

Simon Doherty (University of Sheffield)

Heike Wehrheim (University of Paderborn)

John Derrick (University of Sheffield)

A weak memory talk

```
{talk = weak_memory}
  reaction := listen(talk)
{reaction = 😊 ∨ reaction = 🙄 }
```

A weak memory talk

```
{talk = weak_memory}
  reaction := listen(talk)
{reaction = 😊 ∨ reaction = 🙄} || this(talk)
{ reaction = 😊 }
```

A weak memory talk

$$\begin{array}{l} \{ \text{talk} = \text{weak_memory} \} \\ \text{reaction} := \text{listen}(\text{talk}) \\ \{ \text{reaction} = \text{😊} \vee \text{reaction} = \text{🙄} \} \end{array} \quad \parallel \quad \begin{array}{l} \text{this}(\text{talk}) \\ \\ \{ \text{reaction} = \text{😊} \} \end{array}$$

Turning 🙄 into 😊 — relate weak memory semantics to Hoare logic and Owicki-Gries style proof rules

Outline

C11 Axiomatic Semantics

C11 Operational Semantics

C11 Owicki-Gries Proofs in Isabelle

C11 Axiomatic Semantics

Axiomatic C11 semantics

Example (Message Passing).


```
Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ← f
  f := 1;   until r1 = 1;
            r2 ← d;
```

Axiomatic C11 semantics

Example (Message Passing).

```
Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ← f
  f := 1;   until r1 = 1;
            r2 ← d;
```

In C11, r2 can have a final value 0
— the execution below is **allowed**




Axiomatic C11 semantics

Example (Message Passing).

```

Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ← f
  f := 1;   until r1 = 1;
            r2 ← d;
  
```

In C11, r2 can have a final value 0 — the execution below is **allowed**



Axiomatic C11 semantics

Example (Message Passing).

```


Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ← f
  f := 1;   until r1 = 1;
            r2 ← d;
  
```

Corrected Message Passing.

```

Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ←A f
  f :=R 1;  until r1 = 1;
            r2 ← d;
  
```

In C11, r2 can have a final value 0
 — the execution below is **allowed**




Axiomatic C11 semantics

Example (Message Passing).

```

Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ← f
  f := 1;   until r1 = 1;
            r2 ← d;
  
```

In C11, r2 can have a final value 0 — the execution below is **allowed**




Corrected Message Passing.

```

Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ←A f
  f :=R 1;  until r1 = 1;
            r2 ← d;
  
```

The following execution is now **disallowed**




Axiomatic C11 semantics

Example (Message Passing).

```

Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ← f
  f := 1;   until r1 = 1;
            r2 ← d;
  
```

In C11, r2 can have a final value 0 — the execution below is **allowed**




Corrected Message Passing.

```

Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ←A f
  f :=R 1;  until r1 = 1;
            r2 ← d;
  
```

The following execution is now **disallowed**




Axiomatic C11 semantics

Example (Message Passing).

```

Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ← f
  f := 1;   until r1 = 1;
            r2 ← d;
  
```

In C11, r2 can have a final value 0 — the execution below is **allowed**




Corrected Message Passing.

```

Init:  f := 0; d := 0;
thread 1  ||  thread 2
  d := 5;   do r1 ←A f
  f :=R 1;  until r1 = 1;
            r2 ← d;
  
```

The following execution is now **disallowed**



What about verification?

- ▶ Axiomatic semantics useful for certain forms of verification, e.g., SMT, BMC, ...
- ▶ But how can we link with existing works
 - Hoare Logic, Owicki/Gries, Rely/Guarantee ?

We need an operational semantics for C11

C11 Operational Semantics

Point of departure

- ▶ Start with operational semantics by Doherty et al (2019)
 - proved sound and complete with respect to RC11
- ▶ For the experts: restrict attention to a fragment of C11
 - ▶ All operations are either *relaxed*, *write-releasing*, or *read-acquiring*
 - ▶ Do **not** model *fences* or *release-sequences*
 - ▶ Assume *no-thin-air*, i.e., $sb \cup rf$ acyclic
- ▶ Strategy: construct valid C11 graphs by stepping through program in thread order (without consulting axioms)
- ▶ Brings us back to well understood (programmer friendly) notion
Concurrency = Interleaving of threads
- ▶ What's different?
 - ▶ More non-determinism in choosing the next C11 state
 - ▶ Both reads **and** writes may change state configuration

Observing a C11 state

Key point.


- ▶ Each thread has its own **observable set** of writes
- ▶ Observable writes can be determined from the current C11 state

Observing a C11 state

Key point.

- ▶ Each thread has its own **observable set** of writes
- ▶ Observable writes can be determined from the current C11 state

Example. Restricting $mo \cup rf \cup fr$ to a single variable, we have:



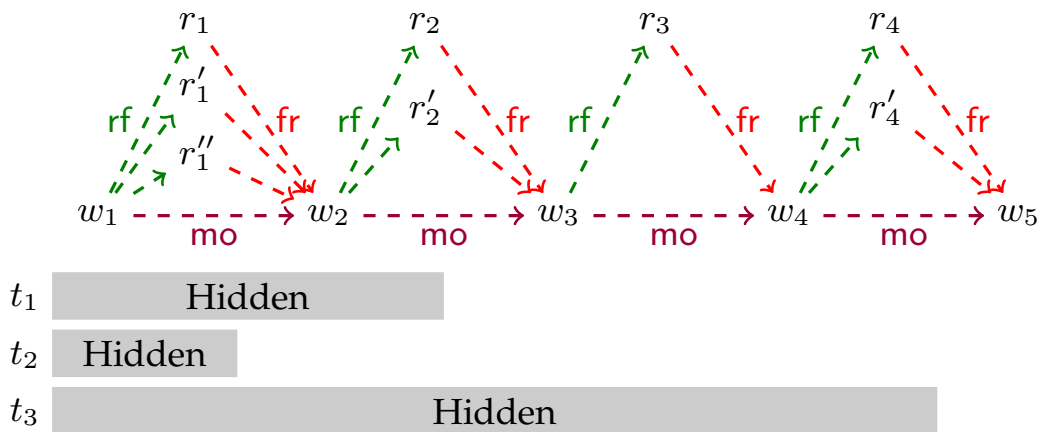
- ▶ Thread t_1 can observe w_3, w_4, w_5
- ▶ Thread t_2 can observe w_2, w_3, w_4, w_5
- ▶ Thread t_3 can observe w_5

Observing a C11 state

Key point.

- ▶ Each thread has its own **observable set** of writes
- ▶ Observable writes can be determined from the current C11 state

Example. Restricting $mo \cup rf \cup fr$ to a single variable, we have:




- ▶ Thread t_1 can observe w_3, w_4, w_5
- ▶ Thread t_2 can observe w_2, w_3, w_4, w_5
- ▶ Thread t_3 can observe w_5

Observable set changes as threads interact with the C11 state

Message passing with “bad” transition

```
Init:  f := 0; d := 0  
thread 1  ||  thread 2  
  d := 5;    do r1 ← f until r1 = 1;  
  f := 1;    r2 ← d;
```


Pre-state



Message passing with “bad” transition

```
Init:  f := 0; d := 0
thread 1  ||  thread 2
  d := 5;   ||   do r1 ← f until r1 = 1;
  f := 1;   ||   r2 ← d;
```

Pre-state




Thread 2 can observe
both writes to *d*

Message passing with “bad” transition

```


Init:  f := 0; d := 0
thread 1  ||  thread 2
  d := 5;   ||   do r1 ← f until r1 = 1;
  f := 1;   ||   r2 ← d;
  
```

Pre-state



Thread 2 can observe both writes to d

Possible post-state




“Bad” transition with read from $wr(d, 0)$ is possible

Message passing with release/acquire annotations

```
Init:  f := 0; d := 0  
thread 1  ||  thread 2  
  d := 5;    do r1 ←A 1 until r1 = 1;  
  f :=R 1;  r2 := d;
```


Pre-state



Message passing with release/acquire annotations

```
Init:  f := 0; d := 0
thread 1  ||  thread 2
  d := 5;   do r1 ←A 1 until r1 = 1;
  f :=R 1;  r2 := d;
```

Pre-state




Thread 2 can only observe $wr_1(d, 5)$

Message passing with release/acquire annotations

```


Init:  f := 0; d := 0
thread 1  ||  thread 2
  d := 5;   do r1 ←A 1 until r1 = 1;
  f :=R 1;   r2 := d;
  
```

Pre-state



Thread 2 can only observe $wr_1(d, 5)$

Only possible post-state



Only the "good" transition is available

C11 Owicki-Gries Proofs in Isabelle

Proof outline for message passing

Init: $d := 0; f := 0;$

$d := 5;$

$f :=^R 1;$



do $r1 \leftarrow^A f$ **until** $r1 = 1;$

$r2 \leftarrow d;$

Proof outline for message passing

Init: $d := 0; f := 0;$

$\{d =_1 0 \wedge d =_2 0 \wedge f =_1 0 \wedge f =_2 0\}$

$\{\neg(f \approx_2 1) \wedge d =_1 0\}$

$d := 5;$

$\{\neg(f \approx_2 1) \wedge d =_1 5\}$

$f :=^R 1;$

$\{true\}$

|||

$\{[f = 1]_2(d = 5)\}$

do $r1 \leftarrow^A f$ **until** $r1 = 1;$

$\{d =_2 5\}$

$r2 \leftarrow d;$

$\{r2 = 5\}$

$\{r2 = 5\}$

Proof outline for message passing

$$\begin{array}{c}
 \text{Init: } d := 0; f := 0; \\
 \{d =_1 0 \wedge d =_2 0 \wedge f =_1 0 \wedge f =_2 0\} \\
 \{ \neg(f \approx_2 1) \wedge d =_1 0 \} \quad \parallel \quad \{ [f = 1]_2 (d = 5) \} \\
 d := 5; \quad \text{do } r1 \leftarrow^A f \text{ until } r1 = 1; \\
 \{ \neg(f \approx_2 1) \wedge d =_1 5 \} \quad \parallel \quad \{ d =_2 5 \} \\
 f :=^R 1; \quad r2 \leftarrow d; \\
 \{ true \} \quad \parallel \quad \{ r2 = 5 \} \\
 \{ r2 = 5 \}
 \end{array}$$

Recall the Owicki-Gries technique:

$$\frac{\vdash \{P_1\}C_1\{Q_1\} \parallel \{P_2\}C_2\{Q_2\} \quad P \Rightarrow P_1 \wedge P_2 \quad Q_1 \wedge Q_2 \Rightarrow Q}{\vdash \{P\} (\{P_1\}C_1\{Q_1\} \parallel \{P_2\}C_2\{Q_2\}) \{Q\}}$$

$$\frac{\begin{array}{l} \vdash \{P_1\}C_1\{Q_1\} \\ \vdash \{P_2\}C_2\{Q_2\} \end{array} \quad \begin{array}{l} \{P_1\}C_1\{Q_1\} \text{ is interference free wrt } C_2 \\ \{P_2\}C_2\{Q_2\} \text{ is interference free wrt } C_1 \end{array}}{\vdash \{P_1\}C_1\{Q_1\} \parallel \{P_2\}C_2\{Q_2\}}$$

Proof outline for message passing

Init: $d := 0; f := 0;$

$\{d =_1 0 \wedge d =_2 0 \wedge f =_1 0 \wedge f =_2 0\}$

$\{\neg(f \approx_2 1) \wedge d =_1 0\}$		$\{[f = 1]_2(d = 5)\}$
$d := 5;$		do $r1 \leftarrow^A f$ until $r1 = 1;$
$\{\neg(f \approx_2 1) \wedge d =_1 5\}$		$\{d =_2 5\}$
$f :=^R 1;$		$r2 \leftarrow d;$
$\{true\}$		$\{r2 = 5\}$
		$\{r2 = 5\}$

- ▶ The **C11 state** is a special implicit variable in the program
- ▶ Assertions are **predicates over program states** (including the C11 states)

Proof outline for message passing

Init: $d := 0; f := 0;$
 $\{d =_1 0 \wedge d =_2 0 \wedge f =_1 0 \wedge f =_2 0\}$
 $\{ \neg(f \approx_2 1) \wedge d =_1 0 \}$ $\{ [f = 1]_2(d = 5) \}$
 $d := 5;$ **do** $r1 \leftarrow^A f$ **until** $r1 = 1;$
 $\{ \neg(f \approx_2 1) \wedge d =_1 5 \}$ $\{ d =_2 5 \}$
 $f :=^R 1;$ $r2 \leftarrow d;$
 $\{ true \}$ $\{ r2 = 5 \}$
 $\{ r2 = 5 \}$

- ▶ The **C11 state** is a special implicit variable in the program
- ▶ Assertions are **predicates over program states** (including the C11 states)
- ▶ We define **special assertions** on C11 state:

$x \approx_t v \iff$ Thread t **possibly** observes value v for x
 $x =_t v \iff$ Thread t **definitely** observes value v for x
 $[x = u]_t(y = v) \iff$ **If** thread t observes $x = u$
then it will definitely observe $y = v$

Hoare-style axioms

- ▶ Rules for compound statements are exactly as in Hoare logic
- ▶ But have a new set of **basic axioms** for (atomic) reads and writes (76 at last count), e.g.,

$$\text{d_obs_WrX_set} \frac{}{\{x =_t u\} [x := v]_t \{x =_t v\}}$$

$$\text{not_pobs_RdA_pres} \frac{}{\{\neg(x \approx_t u)\} [v \leftarrow^A y]_{t'} \{\neg(x \approx_t u)\}}$$

$$\text{c_obs_WrR_pres} \frac{z \neq y \quad z \neq x \quad x \neq y}{\{[x = u]_t(y = v)\} [z :=^R w]_t \{[x = u]_t(y = v)\}}$$

Hoare-style axioms

- ▶ Rules for compound statements are exactly as in Hoare logic
- ▶ But have a new set of **basic axioms** for (atomic) reads and writes (76 at last count), e.g.,

$$\text{d_obs_WrX_set} \frac{}{\{x =_t u\} [x := v]_t \{x =_t v\}}$$

$$\text{not_pobs_RdA_pres} \frac{}{\{\neg(x \approx_t u)\} [v \leftarrow^A y]_{t'} \{\neg(x \approx_t u)\}}$$

$$\text{c_obs_WrR_pres} \frac{z \neq y \quad z \neq x \quad x \neq y}{\{[x = u]_t(y = v)\} [z :=^R w]_t \{[x = u]_t(y = v)\}}$$

- ▶ All basic axioms verified in Isabelle, e.g.,

corollary d_obs_RdX_other:

"wfs $\sigma \implies x \neq y \implies$

$[x =_t u] \sigma \implies \sigma [v \leftarrow y]_t \sigma' \implies [x =_t u] \sigma'$ "

by (metis RdX_def avar.simps(1) d_obs_other)

C11 Owicki-Gries in Isabelle

- ▶ Owicki-Gries theory is included in standard Isabelle distribution (Nieto and Nipkow, 2002)

C11 Owicki-Gries in Isabelle

- ▶ Owicki-Gries theory is included in standard Isabelle distribution (Nieto and Nipkow, 2002)
- ▶ We have extended Nieto-Nipkow's WHILE language with relaxed / release-acquire statements
- ▶ C11 state is embedded in the standard state, e.g., for message passing

```
record MP =  
  d :: V  
  f :: V  
  r1 :: V  
  r2 :: V  
   $\sigma$  :: C11_state
```
- ▶ C11 states updated w.r.t. our operational semantics

Proof of message passing in Isabelle

lemma MessagePassing:

```
"||- { (wfs `σ `f `d) ∧ [ `d =1 0 ] `σ ∧ [ `d =2 0 ] `σ
      ∧ [ `f =1 0 ] `σ ∧ [ `f =2 0 ] `σ }
COBEGIN { (wfs `σ `f `d) ∧ ¬[ `f ≈2 1 ] `σ ∧ [ `d =1 0 ] `σ }
  < `d [ `σ ] :=1 5 >;
  { (wfs `σ `f `d) ∧ ¬[ `f ≈2 1 ] `σ ∧ [ `d =1 5 ] `σ }
  < `f [ `σ ]R :=1 1 >
  { [ `d =1 5 ] `σ }
||
  { (wfs `σ `f `d) ∧ [ `f = 1 ]2( `d = 5 ) `σ }
DO { (wfs `σ `f `d) ∧ [ `f = 1 ]2( `d = 5 ) `σ }
  < `r1 [ `σ ]A ←2 `f >
UNTIL `r1 = 1
INV { (wfs `σ `f `d) ∧ [ `f = 1 ]2( `d = 5 ) `σ
      ∧ ( `r1 = 1 → [ `d =2 5 ] `σ ) }
OD;; { (wfs `σ `f `d) ∧ [ `d =2 5 ] `σ }
  < `r2 [ `σ ] ←2 `d >
  { `r2 = 5 }
COEND
{ `r2 = 5 }"
apply oghoare
apply auto
using d_obs_diff_false zero_neq_numeral
by blast+
```

Case study 2: Peterson's mutual exclusion

```
Init:  flag1 := false; flag2 := false; turn = 1
thread 1
  flag1 := true;
  swapRA(turn, 2);
  do
    r1 ←A flag2;
    r2 ← turn;
  until ¬r1 ∨ r2 = 1;
  //CS1;
  flag1 :=R false;
thread 2
  flag2 := true;
  swapRA(turn, 1);
  do
    r3 ←A flag1;
    r4 ← turn;
  until ¬r3 ∨ r4 = 2;
  //CS2;
  flag2 :=R false;
```

- ▶ Encoded and verified in Isabelle
- ▶ Requires new types of assertions describing the C11 state
- ▶ Same auxiliary variable as proof in sequentially consistent setting (Apt and Olderog, 2009)
- ▶ However, proof requires more work beyond oghoare and auto
 - currently investigating ways to speed this up

Conclusions

- ▶ Operational semantics by Doherty et al (2019) makes deductive verification possible for (a realistic fragment of) C11
- ▶ Verification based on well-understood Owicki-Gries theory
- ▶ Straightforward extension of Nieto and Nipkow's mechanisations of Owicki-Gries in Isabelle
- ▶ Paper describing these works is forthcoming
- ▶ Currently investigating links with distributed correctness (with Philippa Gardner)
- ▶ Any questions, please e-mail: b.dongol@surrey.ac.uk