

Validating optimisations of concurrent C/C++ programs

Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS)

Joint work with Soham Chakraborty

Some questions:

- ▶ When is a compiler optimisation correct?
- ▶ Which optimisations are correct?
- ▶ How does one check that an optimisation is correct?
- ▶ Are compilers actually correct?
- ▶ What else can we do?

Compilation:

$$\text{source} \xrightarrow{\text{transform}} \text{target}$$

Correctness:

$$\text{Behaviours}(\text{target}) \subseteq \text{Behaviours}(\text{source})$$

Examples:

$x = 1; \text{print}(x);$	\rightsquigarrow	$x = 1; \text{print}(1);$	(✓)
$\text{print}(\text{unitialised_variable});$	\rightsquigarrow	$\text{print}(78);$	(✓)
$\text{print}(1); \text{print}(2)$	\rightsquigarrow	$\text{print}(1) \parallel \text{print}(2)$	(✗)
$\text{print}(1) \parallel \text{print}(2)$	\rightsquigarrow	$\text{print}(1); \text{print}(2)$	(✓)

Defining concurrent program behaviours

Sequential consistency (SC):

- ▶ Interleave each thread's atomic accesses.
- ▶ Most verification work assumes this model.

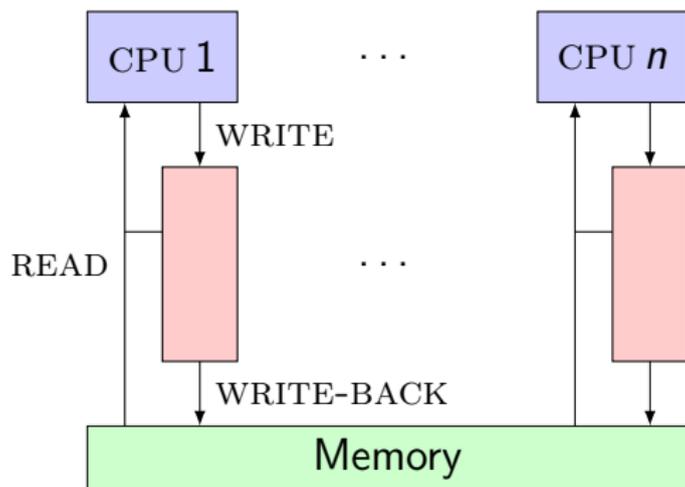
Initially, $X = Y = 0$.

$$\begin{array}{l} X = 1; \\ a = Y; \end{array} \parallel \begin{array}{l} Y = 1; \\ b = X; \end{array}$$

Under SC, this program cannot return $a = b = 0$.

SC is, however, obsolete. It is invalidated by modern hardware and compiler optimisations.

Store buffering in x86-TSO



Initially, $X = Y = 0$.

$$\begin{array}{l} X = 1; \\ a = Y; \end{array} \parallel \begin{array}{l} Y = 1; \\ b = X; \end{array}$$

Allowed outcome: $a = b = 0$.

Basic compiler optimisations break SC / TSO

Initially, $X = Y = 0$.

```
X = 1;  |||  print(X);  
Y = 1;  |||  print(Y);  
                |||  print(X);
```

Can the program print 010?

Justification:

The compiler may perform CSE:
Load X into a temporary t
and print t , Y , and t .

Observe:

- ▶ Programs with weak behaviours have racy accesses.

$$\begin{array}{l}
 X = 1; \\
 a = Y;
 \end{array}
 \parallel
 \begin{array}{l}
 Y = 1; \\
 b = X;
 \end{array}
 \qquad
 \begin{array}{l}
 X = 1; \\
 Y = 1;
 \end{array}
 \parallel
 \begin{array}{l}
 \text{print}(X); \\
 \text{print}(Y); \\
 \text{print}(X);
 \end{array}$$

- ▶ Declare such programs wrong/undefined.
- ▶ Provide some support for intensional races.
This is needed, e.g., to implement locks.

Consequences:

- ▶ Blame the programmer for writing racy programs!
- ▶ Allow the compiler to optimise aggressively!

- ▶ Even so, some optimisations that are correct for sequential programs are **not** correct for concurrent programs.

```
atomic_int L = 0; int X = 0;
lock() {...}
unlock() {L = 0; /* atomic access */}
```

lock();		lock();	↔	lock();		lock();
X = 42;		r = X;		unlock();		r = X;
unlock();		unlock();		X = 42;		unlock();

- ▶ The transformation above introduces a race.

Should these transformations be allowed?

1. CSE over a lock acquire:

$$\begin{array}{l} t_1 = X; \\ \text{lock}(); \\ t_2 = X; \end{array} \rightsquigarrow \begin{array}{l} t_1 = X; \\ \text{lock}(); \\ t_2 = t_1; \end{array}$$

If X changes in between, the program is racy.

2. Load hoisting:

$$\begin{array}{l} \text{if}(c) \\ r = X; \end{array} \rightsquigarrow \begin{array}{l} t = X; \\ r = c ? t : r; \end{array}$$

This may introduce a race, but the racy value is not used.

Allowing both is clearly wrong!

Consider the transformation sequence:

<code>if (c)</code>		<code>t = X;</code>		<code>t = X;</code>
<code> r₁ = X;</code>	\rightsquigarrow	<code>r₁ = c ? t : r₁;</code>	\rightsquigarrow	<code>r₁ = c ? t : r₁;</code>
<code>lock();</code>		<code>lock();</code>		<code>lock();</code>
<code>r₂ = X;</code>		<code>r₂ = X;</code>		<code>r₂ = t;</code>

When c is false, X is moved out of the critical region!

So we have to forbid one transformation.

- ▶ C11 forbids load hoisting, allows CSE over `lock()`.
- ▶ LLVM allows load hoisting, forbids CSE over `lock()`.

The C11 memory model

- ▶ Introduced by the C and C++ standards in 2011.
- ▶ First formalized by Batty et al., POPL 2011.
- ▶ Several subsequent fixes.

The C11 memory model

Two types of locations: ordinary and atomic

- ▶ Races on ordinary accesses \rightsquigarrow error

A spectrum of atomic accesses:

- ▶ Relaxed \rightsquigarrow no fence
- ▶ Consume reads \rightsquigarrow no fence, but preserve deps
- ▶ Release writes \rightsquigarrow no fence (x86); lwsync (PPC)
- ▶ Acquire reads \rightsquigarrow no fence (x86); isync (PPC)
- ▶ Seq. consistent \rightsquigarrow full memory fence

Primitives for explicit fences

Roach motel example in C11

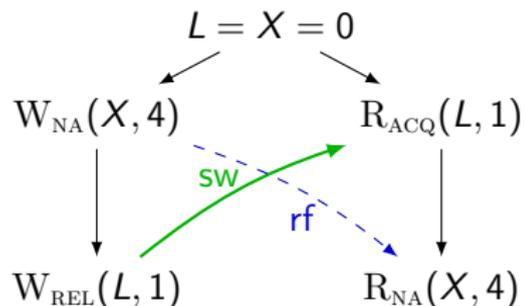
```
atomic_int L = 0; int X = 0;
```

```
 $X_{\text{NA}} = 4;$  ||  $\text{if}(L_{\text{ACQ}})$   
 $L_{\text{REL}} = 1;$  ||  $r = X_{\text{NA}};$ 
```

Roach motel example in C11

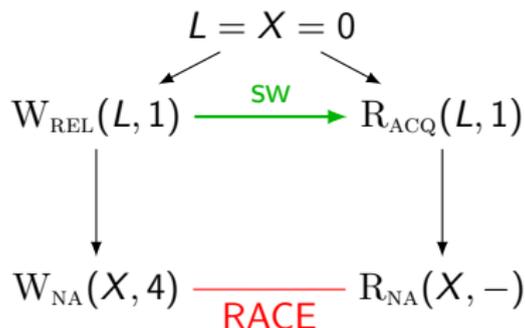
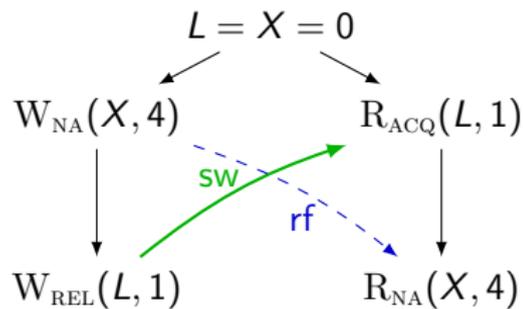
```
atomic_int L = 0; int X = 0;
```

```
 $X_{\text{NA}} = 4;$  ||  $\text{if}(L_{\text{ACQ}})$   
 $L_{\text{REL}} = 1;$  ||  $r = X_{\text{NA}};$ 
```



Roach motel example in C11

```
atomic_int L = 0; int X = 0;
```

$$\begin{array}{l} X_{\text{NA}} = 4; \\ L_{\text{REL}} = 1; \end{array} \parallel \begin{array}{l} \text{if}(L_{\text{ACQ}}) \\ r = X_{\text{NA}}; \end{array} \rightsquigarrow \begin{array}{l} L_{\text{REL}} = 1; \\ X_{\text{NA}} = 4; \end{array} \parallel \begin{array}{l} \text{if}(L_{\text{ACQ}}) \\ r = X_{\text{NA}}; \end{array}$$


Reorderings

✓ $unlock(); X = 1; \rightsquigarrow X = 1; unlock();$

✗ $X = 1; unlock(); \rightsquigarrow unlock(); X = 1;$

Eliminations

✓ $t = X; t' = X; \rightsquigarrow t = X; t' = t;$

✓ $t = X; lock(); t' = X; \rightsquigarrow t = X; lock(); t' = t;$

✓ $t = X; unlock(); t' = X; \rightsquigarrow t = X; unlock(); t' = t;$

✗ $t = X; unlock(); lock(); t' = X;$
 $\rightsquigarrow t = X; unlock(); lock(); t' = t;$

$\downarrow a \setminus b \rightarrow$	$R_{\neq sc}$	R_{sc}	W_{na}	W_{rlx}	$W_{\sqsupseteq rel}$	$C_{rlx acq}$	$C_{\sqsupseteq rel}$	F_{acq}	F_{rel}
R_{na}	✓	✓	✓	✓	✗	✓	✗	✓	✗
R_{rlx}	✓	✓	✓	(✓)	✗	(✓)	✗	✗	✗
$R_{\sqsupseteq acq}$	✗	✗	✗	✗	✗	✗	✗	✓	✗
$W_{\neq sc}$	✓	✓	✓	✓	✗	✓	✗	✓	✗
W_{sc}	✓	✗	✓	✓	✗	✓	✗	✓	✗
$C_{rlx rel}$	✓	✓	✓	(✓)	✗	(✓)	✗	✗	✗
$C_{\sqsupseteq acq}$	✗	✗	✗	✗	✗	✗	✗	✓	✗
F_{acq}	✗	✗	✗	✗	✗	✗	✗	=	✗
F_{rel}	✓	✓	✓	✗	✓	✗	✓	✓	=

Overwritten write:

$$\begin{array}{l} x.\text{store}(v, M); C; x.\text{store}(v', M) \\ \rightsquigarrow C; x.\text{store}(v', M) \end{array} \quad \begin{array}{l} C \text{ has no rel} \\ \& \text{ no } x \text{ accesses} \end{array}$$

Read after write:

$$\begin{array}{l} x.\text{store}(v, M); C; t = x.\text{load}(M') \\ \rightsquigarrow x.\text{store}(v, M); C; t = v \end{array} \quad \begin{array}{l} C \text{ has no acq} \\ \& \text{ no } x \text{ accesses} \end{array}$$

Read after read:

$$\begin{array}{l} t = x.\text{load}(M); C; t' = x.\text{load}(M) \\ \rightsquigarrow t = x.\text{load}(M); C; t' = t \end{array} \quad \begin{array}{l} C \text{ has no acq} \\ \& \text{ no } x \text{ accesses} \end{array}$$

The LLVM memory model

- ▶ Described in prose in the LLVM documentation.
- ▶ Formalization in progress.

Overview

- ▶ LLVM concurrency constructs match those of C11.
- ▶ But they give different semantics to racy programs.

C11 semantics

- ▶ In C11, racy programs have **undefined** behavior.

LLVM semantics

- ▶ *Write-write* races result in **undefined** behavior.
- ▶ *Write-read* races have **defined** behavior.
- ▶ The racy reads return **undef** (\approx any value).

The semantics of “undef” in LLVM

The semantics of “undef” in LLVM is very weak.

- ▶ Arithmetic on “undef” returns “undef”.
- ▶ Branching on “undef” could go either way.
- ▶ The following program, starting with $X = Y = Z = 0$,

$$Z_{\text{NA}} = 5; \left\| \begin{array}{l} t = Z_{\text{NA}}; \\ \text{if}(t == 1) X_{\text{NA}} = 1; \\ \text{if}(t == 7) Y_{\text{NA}} = 1; \end{array} \right.$$

is allowed to return $X = Y = 1$.

- ▶ This semantics allows some very aggressive optimisations. For example, $X = \text{undef}; \rightsquigarrow \text{skip}$;

Impact on optimization: C11 vs LLVM

Non-atomic read introduction:

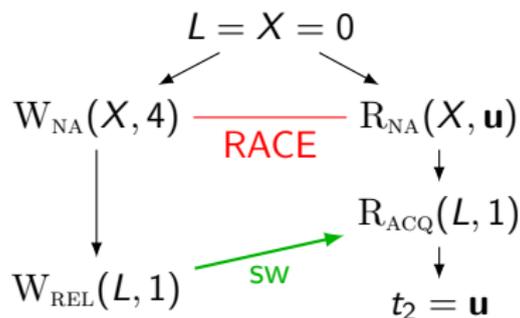
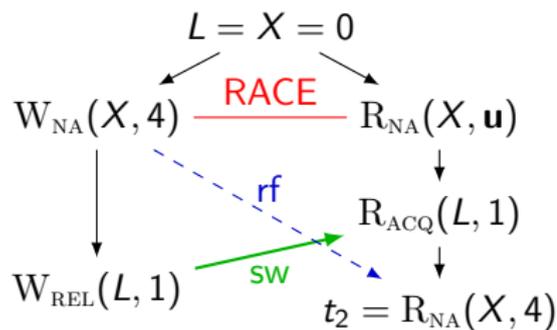
- ▶ Unsafe in C11
- ▶ Allowed in LLVM

$$X_{\text{NA}} = 1; \left\| \begin{array}{l} r = 0; \\ \text{if } (f) \\ r = X_{\text{NA}}; \end{array} \right. \rightsquigarrow X_{\text{NA}} = 1; \left\| \begin{array}{l} r = 0; \\ t = X_{\text{NA}}; \\ r = f ? t : 0; \end{array} \right.$$

Read elimination over an acquire action

Allowed in C11 but unsafe in LLVM.

$$\begin{array}{l} X_{\text{NA}} = 4; \\ L_{\text{REL}} = 1; \end{array} \parallel \begin{array}{l} t_1 = X_{\text{NA}}; \\ \text{if}(L_{\text{ACQ}}) \\ t_2 = X_{\text{NA}}; \end{array} \rightsquigarrow \begin{array}{l} X_{\text{NA}} = 4; \\ L_{\text{REL}} = 1; \end{array} \parallel \begin{array}{l} t_1 = X_{\text{NA}}; \\ \text{if}(L_{\text{ACQ}}) \\ t_2 = t_1; \end{array}$$



Compiler validation

- ▶ Translation validation is a generic approach.
- ▶ We apply it to the LLVM compiler.
- ▶ Specifically to check for concurrency errors.

$(P_{src} \xrightarrow{comp.} P_{tgt}) ?$ **Correct** : **Error**



1. Identify safe reorderings(R) & safe eliminations(E)

2. $(P_{src} \xrightarrow{(RUE)^*} P_{tgt}) ?$ **Correct** : **Error**

Metadata-based validation (MD)

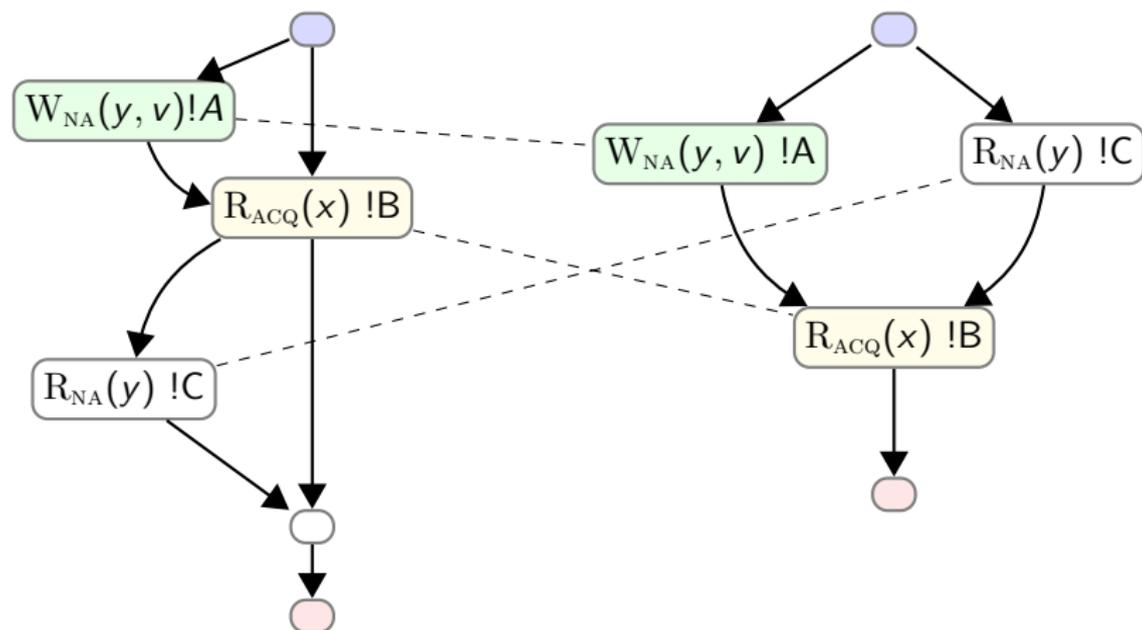
- ▶ Use LLVM metadata capability.
- ▶ Modify compiler to generate required metadata.

Compiler-independent matching (CIM)

- ▶ Independent of compiler internals.
- ▶ Applicable to any compiler.

Metadata-based validation (MD)

- ▶ Match accesses based on attached metadata.
- ▶ Analyze the matching on source and target CFGs.



Metadata-based validation (MD)

Can give false positives!



MD Report: **Error**

Compiler-independent matching (CIM)

- ▶ Mark action deletability/non-deletability
- ▶ Match access sequences and analyze



CIM reports: **Correct**

source

$$(C) \otimes a : W_{\text{RLX}}(X, v_x)$$

$$\checkmark b : W_{\text{REL}}(Y, v_y)$$

$$\checkmark c : R_{\text{NA}}(g)$$

$$\checkmark d : W_{\text{RLX}}(X, v'_x)$$

$$\checkmark e : R_{\text{ACQ}}(Z)$$

$$\times f : R_{\text{NA}}(g)$$

$$C = [a; W_{\text{RLX}}(X, -); b]$$

target

$$a' : W_{\text{RLX}}(X, v'_x) \checkmark$$

$$b' : R_{\text{NA}}(g) \checkmark$$

$$c' : R_{\text{NA}}(g) \times$$

$$d' : R_{\text{ACQ}}(Z) \checkmark$$

$$e' : W_{\text{REL}}(Y, v_y) \checkmark$$

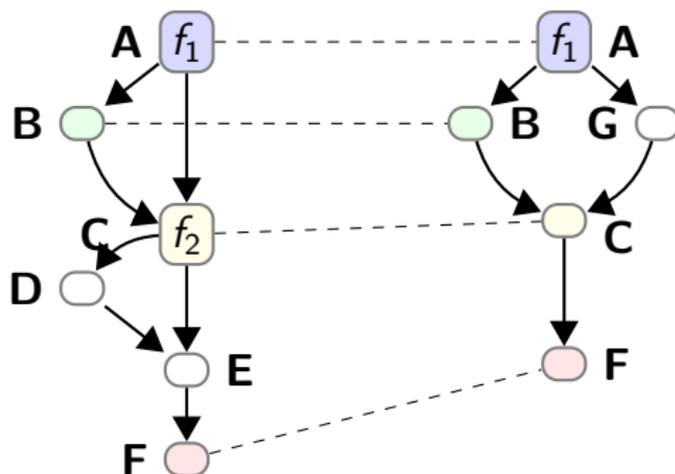
source	target
(C) $\otimes a : W_{\text{RLX}}(X, v_x)$	$a' : W_{\text{RLX}}(X, v'_x)$ ✓
✓ $b : W_{\text{REL}}(Y, v_y)$	$b' : R_{\text{NA}}(g)$ ✓
✓ $c : R_{\text{NA}}(g)$	$c' : R_{\text{NA}}(g)$ ✗
✓ $d : W_{\text{RLX}}(X, v'_x)$	$d' : R_{\text{ACQ}}(Z)$ ✓
✓ $e : R_{\text{ACQ}}(Z)$	$e' : W_{\text{REL}}(Y, v_y)$ ✓
✗ $f : R_{\text{NA}}(g)$	

$$C = [a; W_{\text{RLX}}(X, -); b]$$

Report: **Correct**

Control flow matching

- Match equivalent paths between source and target



$ABCF(f_1) \rightarrow \{ABCDEF(f_1 \wedge f_2), ABCEF(f_1 \wedge \neg f_2)\}$

$AGCF(\neg f_1) \rightarrow \{ACDEF(\neg f_1 \wedge f_2), ACEF(\neg f_1 \wedge \neg f_2)\}$

Experimental results

Tests	Model	End-to-end				Stepwise			
		llvm 3.6		llvm 3.7rc2		llvm 3.6		llvm 3.7rc2	
		CIM	MD	CIM	MD	CIM	MD	CIM	MD
straight	LLVM	95	95	0	0	95 †	95 †	0	0
	C11	0	0	0	0	0	0	0	0
& cond	LLVM	66	74	0	6	64 †	77 ††	0	0
	C11	12	39	0	27	15 ††	43 ††	1 †	31 ††
& dead	LLVM	58	74	0	3	57 †	79 ††	0	0
	C11	6	40	0	25	11 †	49 ††	0	26 ††
& loop	LLVM	48	72	0	27	48 †	90 ††	0	0
	C11	6	48	0	32	7 †	54 ††	0	44 ††
small	LLVM	32	38	0	8	32 †	40 ††	0	0
	C11	7	20	5	19	7 ††	24 ††	6 †	23 ††

- ▶ In total, 500 tests and about 7000 non-trivial steps.
- ▶ Report # failed validations.
- ▶ Erroneous passes: † GVN and †† SimplifyCFG

Experimental results

Tests	Model	End-to-end				Stepwise			
		llvm 3.6		llvm 3.7rc2		llvm 3.6		llvm 3.7rc2	
		CIM	MD	CIM	MD	CIM	MD	CIM	MD
straight	LLVM	95	95	0	0	95 †	95 †	0	0
	C11	0	0	0	0	0	0	0	0
& cond	LLVM	66	74	0	6	64 †	77 ††	0	0
	C11	12	39	0	27	15 ††	43 ††	1 †	31 ††
& dead	LLVM	58	74	0	3	57 †	79 ††	0	0
	C11	6	40	0	25	11 †	49 ††	0	26 ††
& loop	LLVM	48	72	0	27	48 †	90 ††	0	0
	C11	6	48	0	32	7 †	54 ††	0	44 ††
small	LLVM	32	38	0	8	32 †	40 ††	0	0
	C11	7	20	5	19	7 ††	24 ††	6 †	23 ††

Compiler independent matching (CIM) finds:

- ▶ plenty of errors on LLVM 3.6;
- ▶ no errors on LLVM 3.7rc2 (i.e., also no false positives);
- ▶ relatively few errors on LLVM 3.7rc2 w.r.t. the C11 model.

Experimental results

Tests	Model	End-to-end				Stepwise			
		llvm 3.6		llvm 3.7rc2		llvm 3.6		llvm 3.7rc2	
		CIM	MD	CIM	MD	CIM	MD	CIM	MD
straight	LLVM	95	95	0	0	95 †	95 †	0	0
	C11	0	0	0	0	0	0	0	0
& cond	LLVM	66	74	0	6	64 †	77 ††	0	0
	C11	12	39	0	27	15 ††	43 ††	1 †	31 ††
& dead	LLVM	58	74	0	3	57 †	79 ††	0	0
	C11	6	40	0	25	11 †	49 ††	0	26 ††
& loop	LLVM	48	72	0	27	48 †	90 ††	0	0
	C11	6	48	0	32	7 †	54 ††	0	44 ††
small	LLVM	32	38	0	8	32 †	40 ††	0	0
	C11	7	20	5	19	7 ††	24 ††	6 †	23 ††

Some stepwise errors are masked by other transformations:

$$\begin{array}{l}
 r = X_{ACQ}; \\
 \text{if}(a) \ r' = X_{NA};
 \end{array}
 \quad
 \begin{array}{l}
 \text{(1)} \\
 \rightsquigarrow
 \end{array}
 \begin{array}{l}
 r = X_{ACQ}; \\
 t = X_{NA}; \\
 r' = a? t : r';
 \end{array}
 \quad
 \begin{array}{l}
 \text{(2)} \\
 \rightsquigarrow
 \end{array}
 \begin{array}{l}
 r = X_{ACQ}; \\
 t = r; \\
 r' = a? t : r';
 \end{array}$$

Experimental results

Tests	Model	End-to-end				Stepwise			
		llvm 3.6		llvm 3.7rc2		llvm 3.6		llvm 3.7rc2	
		CIM	MD	CIM	MD	CIM	MD	CIM	MD
straight	LLVM	95	95	0	0	95 †	95 †	0	0
	C11	0	0	0	0	0	0	0	0
& cond	LLVM	66	74	0	6	64 †	77 ††	0	0
	C11	12	39	0	27	15 ††	43 ††	1 †	31 ††
& dead	LLVM	58	74	0	3	57 †	79 ††	0	0
	C11	6	40	0	25	11 †	49 ††	0	26 ††
& loop	LLVM	48	72	0	27	48 †	90 ††	0	0
	C11	6	48	0	32	7 †	54 ††	0	44 ††
small	LLVM	32	38	0	8	32 †	40 ††	0	0
	C11	7	20	5	19	7 ††	24 ††	6 †	23 ††

- ▶ Stepwise errors are also masked by adjacent accesses.
- ▶ Use MD or smaller programs.

Experimental results

Tests	Model	End-to-end				Stepwise			
		llvm 3.6		llvm 3.7rc2		llvm 3.6		llvm 3.7rc2	
		CIM	MD	CIM	MD	CIM	MD	CIM	MD
straight	LLVM	95	95	0	0	95 †	95 †	0	0
	C11	0	0	0	0	0	0	0	0
& cond	LLVM	66	74	0	6	64 †	77 ††	0	0
	C11	12	39	0	27	15 ††	43 ††	1 †	31 ††
& dead	LLVM	58	74	0	3	57 †	79 ††	0	0
	C11	6	40	0	25	11 †	49 ††	0	26 ††
& loop	LLVM	48	72	0	27	48 †	90 ††	0	0
	C11	6	48	0	32	7 †	54 ††	0	44 ††
small	LLVM	32	38	0	8	32 †	40 ††	0	0
	C11	7	20	5	19	7 ††	24 ††	6 †	23 ††

MD is good stepwise, but not end-to-end.

$$\begin{array}{ccc}
 \text{if}(*)\{ \checkmark R_{NA}(g) !A \} & \overset{(1)}{\rightsquigarrow} & \checkmark R_{NA}(g) !A \\
 \checkmark R_{NA}(g) !B & & \times R_{NA}(g) !B
 \end{array}
 \quad
 \begin{array}{ccc}
 & \overset{(2)}{\rightsquigarrow} & \checkmark R_{NA}(g) !A
 \end{array}$$

Summary

- ▶ C11 and LLVM models differ.
- ▶ This affects allowed compiler optimisations.
- ▶ Validator reveals compilation bugs.

Future directions

- ▶ Extend validation to loop optimisations.
- ▶ Properly formalize LLVM semantics.