

Multicore Programming

Peter Sewell

Jaroslav Ševčík

Tim Harris

University of Cambridge

MSR

with thanks to

Francesco Zappa Nardelli, Susmit Sarkar, Tom Ridge, Scott Owens, Magnus O. Myreen, Luc Maranget, Mark Batty, Jade Alglave

October – November, 2010

Small language – expressions

location, x address (or pointer value)
variable, r thread-local variable name
integer, n integer
thread_id, t thread id

<i>expression, e</i>	::=	term
	<i>n</i>	integer literal
	<i>*x</i>	read from pointer
	<i>*x = e</i>	write to pointer
	<i>r</i>	read from local variable
	<i>r = e</i>	write to local variable
	<i>e; e'</i>	sequential composition
	if <i>e = e'</i> then <i>e₁</i> else <i>e₂</i>	conditional
	lock <i>x</i>	lock
	unlock <i>x</i>	unlock
	print <i>e</i>	print
	(<i>e</i>)	S

Small language – processes and states

Let memory M be a map from locations to integers and L be a set of locations.

$process, p$	$::=$	process
		expression
		parallel composition

$state, s$	$::=$	state
		$\langle p, M, L \rangle$

$label, l$	$::=$	label
		write
		read
		lock
		unlock
		print
		internal action (tau)

Expression semantics

Let ρ be a map from thread-local variable names to integers.

$$\boxed{\rho_1; e_1 \xrightarrow{l} \rho_2; e_2} \quad e_1 \text{ does } l \text{ to become } e_2$$

$$\frac{}{\rho; * x \xrightarrow{R x=n} \rho; n} \quad \text{READ}$$

$$\frac{}{\rho; * x = n \xrightarrow{W x=n} \rho; n} \quad \text{WRITE}$$

$$\frac{\rho(r) = n}{\rho; r \xrightarrow{\tau} \rho; n} \quad \text{VAR_READ}$$

$$\frac{}{\rho; r = n \xrightarrow{\tau} \rho \oplus (r \mapsto n); n} \quad \text{VAR_WRITE}$$

$$\frac{\rho; e_1 \xrightarrow{l} \rho'; e'_1}{\rho; e_1; e_2 \xrightarrow{l} \rho'; e'_1; e_2} \quad \text{SEQ_CONTEXT}$$

Expression semantics

$$\frac{}{\rho; n; e \xrightarrow{\tau} \rho; e} \quad \text{SEQ}$$

$$\frac{\rho; e_1 \xrightarrow{l} \rho'; e'_1}{\rho; \text{if } e_1 = e_2 \text{ then } e_3 \text{ else } e_4 \xrightarrow{l} \rho'; \text{if } e'_1 = e_2 \text{ then } e_3 \text{ else } e_4} \quad \text{IF_CONTEXT_1}$$

$$\frac{\rho; e_2 \xrightarrow{l} \rho'; e'_2}{\rho; \text{if } n = e_2 \text{ then } e_3 \text{ else } e_4 \xrightarrow{l} \rho'; \text{if } n = e'_2 \text{ then } e_3 \text{ else } e_4} \quad \text{IF_CONTEXT_2}$$

$$\frac{}{\rho; \text{if } n = n \text{ then } e_3 \text{ else } e_4 \xrightarrow{\tau} \rho; e_3} \quad \text{IF_EQ}$$

$$\frac{n \neq n'}{\rho; \text{if } n = n' \text{ then } e_3 \text{ else } e_4 \xrightarrow{\tau} \rho; e_4} \quad \text{IF_NEQ}$$

$$\frac{}{\rho; \text{lock } x \xrightarrow{Lx} \rho; 0} \quad \text{LOCK}$$

Expression semantics

$$\frac{}{\rho; \mathbf{unlock} \ x \xrightarrow{U \ x} \rho; 0} \quad \text{UNLOCK}$$

$$\frac{\rho; e \xrightarrow{l} \rho'; e'}{\rho; \mathbf{print} \ e \xrightarrow{l} \rho'; \mathbf{print} \ e'} \quad \text{PRINT_CONTEXT}$$

$$\frac{}{\rho; \mathbf{print} \ n \xrightarrow{P \ n} \rho; n} \quad \text{PRINT}$$

Process semantics

$p_1 \xrightarrow{t:l} p_2$ p_1 does l to become p_2

$$\frac{\rho;e \xrightarrow{l} \rho';e'}{t:\rho;e \xrightarrow{t:l} t:\rho';e'} \quad \text{THREAD}$$
$$\frac{p_1 \xrightarrow{t:l} p'_1}{p_1|p_2 \xrightarrow{t:l} p'_1|p_2} \quad \text{PAR_CONTEXT_LEFT}$$
$$\frac{p_2 \xrightarrow{t:l} p'_2}{p_1|p_2 \xrightarrow{t:l} p_1|p'_2} \quad \text{PAR_CONTEXT_RIGHT}$$

SC semantics

$s_1 \xrightarrow{t:l} s_2$ s_1 makes a step to become s_2

$$\frac{p \xrightarrow{t:R x=n} p' \quad M(x) = n}{\langle p, M, L \rangle \xrightarrow{t:R x=n} \langle p', M, L \rangle} \text{ SREAD}$$

$$\frac{p \xrightarrow{t:W x=n} p'}{\langle p, M, L \rangle \xrightarrow{t:W x=n} \langle p', M \oplus (x \mapsto n), L \rangle} \text{ SWRITE}$$

$$\frac{p \xrightarrow{t:\tau} p'}{\langle p, M, L \rangle \xrightarrow{t:\tau} \langle p', M, L \rangle} \text{ STAU}$$

$$\frac{p \xrightarrow{t:L x} p' \quad x \notin L}{\langle p, M, L \rangle \xrightarrow{t:L x} \langle p', M, L \cup \{x\} \rangle} \text{ SLOCK}$$

$$\frac{p \xrightarrow{t:U x} p'}{\langle p, M, L \rangle \xrightarrow{t:U x} \langle p', M, L \setminus \{x\} \rangle} \text{ SUNLOCK}$$

$$\frac{p \xrightarrow{t:P n} p'}{\langle p, M, L \rangle \xrightarrow{t:P n} \langle p', M, L \rangle} \text{ SPRINT}$$

Compiler Optimisations

Correctness

Optimisations in SC

Optimisations of Data-race-free Programs

Data race freedom

Limitations of the DRF principle

Correctness under the DRF principle

Out-of-thin-air values

Compiler Optimisations

Correctness

Optimisations in SC

Optimisations of Data-race-free Programs

Data race freedom

Limitations of the DRF principle

Correctness under the DRF principle

Out-of-thin-air values

World of Optimisations

Typical compiler performs many optimisations.

For example, `gcc 4.4.1` with `-O2` option goes through **147** compilation passes (using `-fdump-tree-all` and `-fdump-rtl-all`).

Sun Hotspot Server JVM has 18 high-level passes with each pass composed of one or more smaller passes (<http://www.azulsystems.com/blog/cliff-click/2009-04-14-odds-ends>).

World of Optimisations

Typical compiler performs many optimisations.

- Common subexpression elimination (copy propagation, partial redundancy elimination, value numbering)
- (conditional) constant propagation
- dead code elimination
- loop optimisations (loop invariant code motion, loop splitting/peeling, loop unrolling, etc.)
- vectorisation
- peephole optimisations
- tail duplication removal
- building graph representations/graph linearisation
- register allocation
- call inlining
- local memory to registers promotion
- spilling
- instruction scheduling. . .

Memory Optimisations

Only some optimisations change shared-memory traces.

Memory Optimisations

Only some optimisations change shared-memory traces.

- Common subexpression elimination (copy propagation, partial redundancy elimination, value numbering)
- (conditional) constant propagation
- dead code elimination
- loop optimisations (loop invariant code motion, loop splitting/peeling, loop unrolling, etc.)
- vectorisation
- peephole optimisations
- tail duplication removal
- building graph representations/graph linearisation
- register allocation
- call inlining
- local memory to registers promotion
- spilling
- instruction scheduling. . .

Memory Optimisations

Only some optimisations change shared-memory traces.

The optimisations of shared-memory can perform:

- **Eliminations** (of reads, writes, sometimes synchronisation).
- **Reordering** (of independent non-conflicting memory accesses).
- **Introductions** (of reads – rarely).

Eliminations

This includes common subexpression elimination, dead read elimination, overwritten write elimination, redundant write elimination.

- Irrelevant read elimination:

$$r=*x; C \rightarrow C,$$

where r is not free in C .

- Redundant read after read elimination:

$$r1=*x; r2=*x \rightarrow r1=*x; r2=r1.$$

- Redundant read after write elimination:

$$*x=r1; r2=*x \rightarrow *x=r1; r2=r1.$$

Reordering

Some loop optimisations, code motion.

- Normal memory access **reordering**:

$$\begin{aligned} r1=*x; r2=*y &\rightarrow r2=*y; r1=*x, \\ *x=r1; *y=r2 &\rightarrow *y=r2; *x=r1, \\ r1=*x; *y=r2 &\Leftrightarrow *y=r2; r1=*x. \end{aligned}$$

- **Roach motel** reordering:

$$\begin{aligned} \text{memop}; \text{lock } m &\rightarrow \text{lock } m; \text{memop}, \\ \text{unlock } m; \text{memop} &\rightarrow \text{memop}; \text{unlock } m, \end{aligned}$$

where memop is $*x=r1$ or $r1=*x$.

Memory access introductions

Can an optimisation introduce memory accesses?

Note that the loop body is not executed.

Memory access introductions

Can an optimisation introduce memory accesses?

Yes, but rarely:

```
i = 0;
...
while (i != 0) (
    j = *x + 1;
    i = i - 1 )
→
i = 0;
...
tmp = *x;
while (i != 0) (
    j = tmp + 1;
    i = i - 1 )
```

Note that the loop body is not executed.

Compiler Optimisations

Correctness

Optimisations in SC

Optimisations of Data-race-free Programs

Data race freedom

Limitations of the DRF principle

Correctness under the DRF principle

Out-of-thin-air values

When is an optimisation correct?

Optimisation is correct if any behaviour of the optimised program could be exhibited by the original program.

I.e., for any execution of the optimised program, there is an execution of the original program with the same observable behaviour.

When is an optimisation correct?

Optimisation is correct if any behaviour of the optimised program could be exhibited by the original program.

I.e., for any execution of the optimised program, there is an execution of the original program with the same observable behaviour.

Where the observable behaviour of an execution is the subtrace of external actions (let us ignore termination).

Example

$$P_1 = *x = 1 \quad \left| \quad \begin{array}{l} r1 = *x; \quad r2 = *x; \\ \text{if } r1=r2 \text{ then print 1 else print 2} \end{array} \right.$$
$$P_2 = *x = 1 \quad \left| \quad \begin{array}{l} r1 = *x; \quad r2 = r1; \\ \text{if } r1=r2 \text{ then print 1 else print 2} \end{array} \right.$$

Is the transformation from P_1 to P_2 correct (in our SC semantics)?

Example

$P_1 = *x = 1$		<code>r1 = *x; r2 = *x;</code> <code>if r1=r2 then print 1 else print 2</code>
$P_2 = *x = 1$		<code>r1 = *x; r2 = r1;</code> <code>if r1=r2 then print 1 else print 2</code>

Executions of P_1 :

$W_{t_1} x=1, R_{t_2} x=1, R_{t_2} x=1, P_{t_2} 1$
 $R_{t_2} x=0, W_{t_1} x=1, R_{t_2} x=1, P_{t_2} 2$
 $R_{t_2} x=0, R_{t_2} x=0, W_{t_1} x=1, P_{t_2} 1$
 $R_{t_2} x=0, R_{t_2} x=0, P_{t_2} 1, W_{t_1} x=1$

Example

$P_1 = *x = 1$		<code>r1 = *x; r2 = *x;</code> <code>if r1=r2 then print 1 else print 2</code>
$P_2 = *x = 1$		<code>r1 = *x; r2 = r1;</code> <code>if r1=r2 then print 1 else print 2</code>

Executions of P_2 :

$W_{t_1} x=1, R_{t_2} x=1, P_{t_2} 1$

$R_{t_2} x=0, W_{t_1} x=1, P_{t_2} 1$

$R_{t_2} x=0, P_{t_2} 1, W_{t_1} x=1$

Example

$$P_1 = *x = 1 \quad \left| \quad \begin{array}{l} r1 = *x; \quad r2 = *x; \\ \text{if } r1=r2 \text{ then print 1 else print 2} \end{array} \right.$$
$$P_2 = *x = 1 \quad \left| \quad \begin{array}{l} r1 = *x; \quad r2 = r1; \\ \text{if } r1=r2 \text{ then print 1 else print 2} \end{array} \right.$$

Behaviours of P_1 : $[P_{t_2} 1], [P_{t_2} 2]$.

Behaviours of P_2 : $[P_{t_2} 1]$.

It is correct to rewrite $P_1 \rightarrow P_2$, but not the opposite!

Optimisation Correctness Overview

Transformation	SC	JMM	DRF
Trace-preserving transformations	✓	✓	✓
Reordering normal memory accesses	×	×	✓
Redundant read after read elimination	✓*	×	✓
Redundant read after write elimination	✓*	✓	✓
Irrelevant read elimination	✓	✓	✓
Irrelevant read introduction	✓	×	?
Redundant write before write elimination	✓*	✓	✓
Redundant write after read elimination	✓*	×	✓
Roach-motel reordering	✓	×	✓
External action reordering	×	×	✓

✓ – correct, × – incorrect, ✓* – correct only for adjacent memory accesses.

Correctness proof strategy

Take an arbitrary **execution** of the optimised program.

Correctness proof strategy

Take an arbitrary execution of the optimised program.

Massage it into an execution of the original program...

Correctness proof strategy

Take an arbitrary **execution** of the optimised program.

Massage it into **an execution** of the original program...

... so that the massaged execution has **the same sequence of observable actions**.

Memory actions are **not** observable.

Compiler Optimisations

Correctness

Optimisations in SC

Optimisations of Data-race-free Programs

Data race freedom

Limitations of the DRF principle

Correctness under the DRF principle

Out-of-thin-air values

General CSE incorrect in SC

```
*x = 1;           |   if *y=1 then (  
*y = 1;           |       *x = 2;  
if *y = 2         |       *y = 2  
then print *x    |   )
```

There is only one execution with a printing behaviour:

$W_{t_1} x=1, W_{t_1} y=1, R_{t_2} x=1, W_{t_2} x=2, W_{t_2} y=2, R_{t_1} y=2, R_{t_1} x=2, P_{t_1} 2$

General CSE incorrect in SC

<code>*x = 1;</code>		<code>if *x=1 then (</code>
<code>*y = 1;</code>		<code> *x = 2;</code>
<code>if *y = 2</code>		<code> *y = 2</code>
<code>then print *x</code>		<code>)</code>

But a compiler would optimise to:

<code>*x = 1;</code>		<code>if *x=1 then (</code>
<code>*y = 1;</code>		<code> *x = 2;</code>
<code>if *y = 2</code>		<code> *y = 2</code>
<code>then print 1</code>		<code>)</code>

General CSE incorrect in SC

<pre>*x = 1; *y = 1; if *y = 2 then print 1</pre>		<pre>if *x=1 then (*x = 2; *y = 2)</pre>
---	--	--

The only execution with a printing behaviour in the optimised program is:

$W_{t_1} x=1, W_{t_1} y=1, R_{t_2} x=1, W_{t_2} x=2, W_{t_2} y=2, R_{t_1} y=2, P_{t_1} 1$

So the optimisation is not correct!

General CSE incorrect in SC II

<code>*x = 1;</code>		<code>r = *x;</code>
<code>*y = 1;</code>		<code>print r;</code>
		<code>print *y;</code>
		<code>print *x;</code>

The observable behaviours are:

$[P_{t_2} 1, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 1, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 0]$

Note that $[P_{t_2} 0, P_{t_2} 1, P_{t_2} 0]$ is **not observable**.

General CSE incorrect in SC II

```
*x = 1;      |      r = *x;  
*y = 1;      |      print r;  
              |      print *y;  
              |      print *x;
```

But a compiler would optimise to:

```
*x = 1;      |      r = *x;  
*y = 1;      |      print r;  
              |      print *y;  
              |      print r;
```

General CSE incorrect in SC II

The optimised program

<code>*x = 1;</code>		<code>r = *x;</code>
<code>*y = 1;</code>		<code>print r;</code>
		<code>print *y;</code>
		<code>print r;</code>

has behaviours:

$[P_{t_2} 1, P_{t_2} 1, P_{t_2} 1]$

$[P_{t_2} 1, P_{t_2} 0, P_{t_2} 1]$

$[P_{t_2} 0, P_{t_2} 1, P_{t_2} 0]$

$[P_{t_2} 0, P_{t_2} 0, P_{t_2} 0]$

New behaviour!

Reordering incorrect

<pre>*x = 1; r1 = *y print r1</pre>		<pre>*y = 1; r2 = *x; print r2</pre>	⇒	<pre>r1 = *y *x = 1; print r1</pre>		<pre>*y = 1; r2 = *x; print r2</pre>
---	--	--	---	---	--	--

Behaviours:

$[P_{t_1} 0, P_{t_2} 1]$
 $[P_{t_1} 1, P_{t_2} 0]$
 $[P_{t_1} 1, P_{t_2} 1]$

Behaviours:

$[P_{t_1} 0, P_{t_2} 1]$
 $[P_{t_1} 1, P_{t_2} 0]$
 $[P_{t_1} 1, P_{t_2} 1]$
 $[P_{t_1} 0, P_{t_2} 0]$

This is essentially the same example as in the first x86 lecture.

Elimination of Adjacent Accesses

There are some correct optimisations under SC. For example, it is correct to rewrite

$$r1 = *x; r2 = *x \quad \rightarrow \quad r1 = *x; r2 = r1$$

in any context.

Why?

Elimination of Adjacent Accesses

There are some correct optimisations under SC. For example, it is correct to rewrite

$$r1 = *x; r2 = *x \quad \rightarrow \quad r1 = *x; r2 = r1$$

in **any context**.

Why?

The basic idea is: whenever we perform the read $r1 = *x$ in the transformed program, we perform **both** reads in the original program.

Simulating Elimination

Original program

⋮
↓

$\langle t : \rho; r1=*x; r2=*x; e \mid p, M, L \rangle$

Optimised program

⋮
↓

$\langle t : \rho; r1=*x; r2=r1; e \mid p, M, L \rangle$

Simulating Elimination

Original program

⋮
↓
 $\langle t : \rho; \mathbf{r1=*x}; \mathbf{r2=*x}; e \mid p, M, L \rangle$

Optimised program

⋮
↓
 $\langle t : \rho; \mathbf{r1=*x}; \mathbf{r2=r1}; e \mid p, M, L \rangle$
↓ $R_t x=n$
 $\langle t : \rho(r1 \mapsto n); \mathbf{r2=r1}; e \mid p, M, L \rangle$

Simulating Elimination

Original program

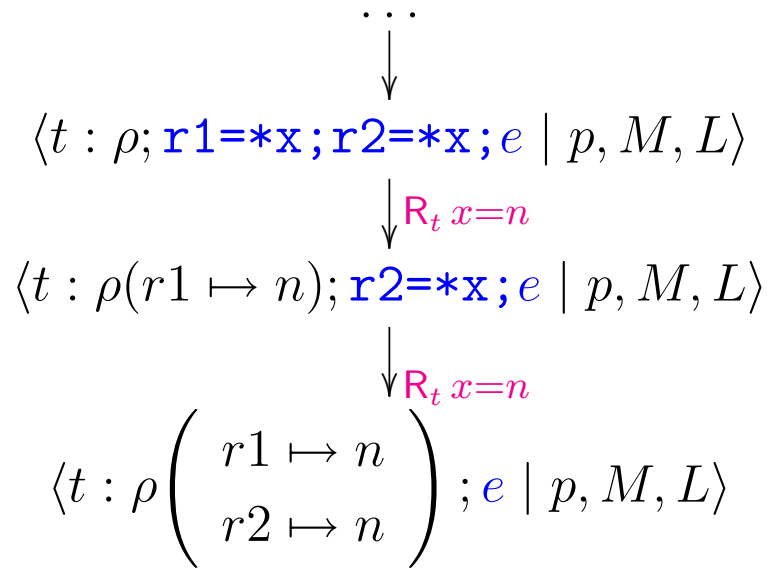
Optimised program

$$\begin{array}{c} \dots \\ \downarrow \\ \langle t : \rho; \mathbf{r1=*x}; \mathbf{r2=*x}; e \mid p, M, L \rangle \\ \downarrow \mathbf{R}_t x=n \\ \langle t : \rho(r1 \mapsto n); \mathbf{r2=*x}; e \mid p, M, L \rangle \\ \downarrow \mathbf{R}_t x=n \\ \langle t : \rho \left(\begin{array}{l} r1 \mapsto n \\ r2 \mapsto n \end{array} \right); e \mid p, M, L \rangle \end{array}$$


$$\begin{array}{c} \dots \\ \downarrow \\ \langle t : \rho; \mathbf{r1=*x}; \mathbf{r2=r1}; e \mid p, M, L \rangle \\ \downarrow \mathbf{R}_t x=n \\ \langle t : \rho(r1 \mapsto n); \mathbf{r2=r1}; e \mid p, M, L \rangle \end{array}$$

Simulating Elimination

Original program




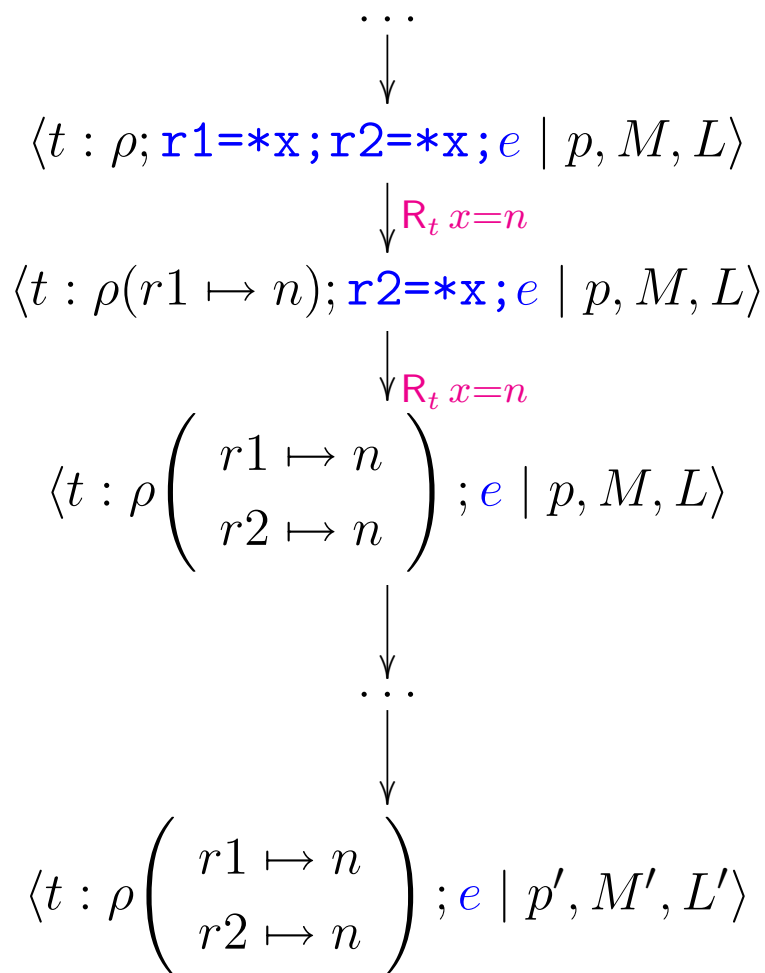
Optimised program



Simulating Elimination

Original program


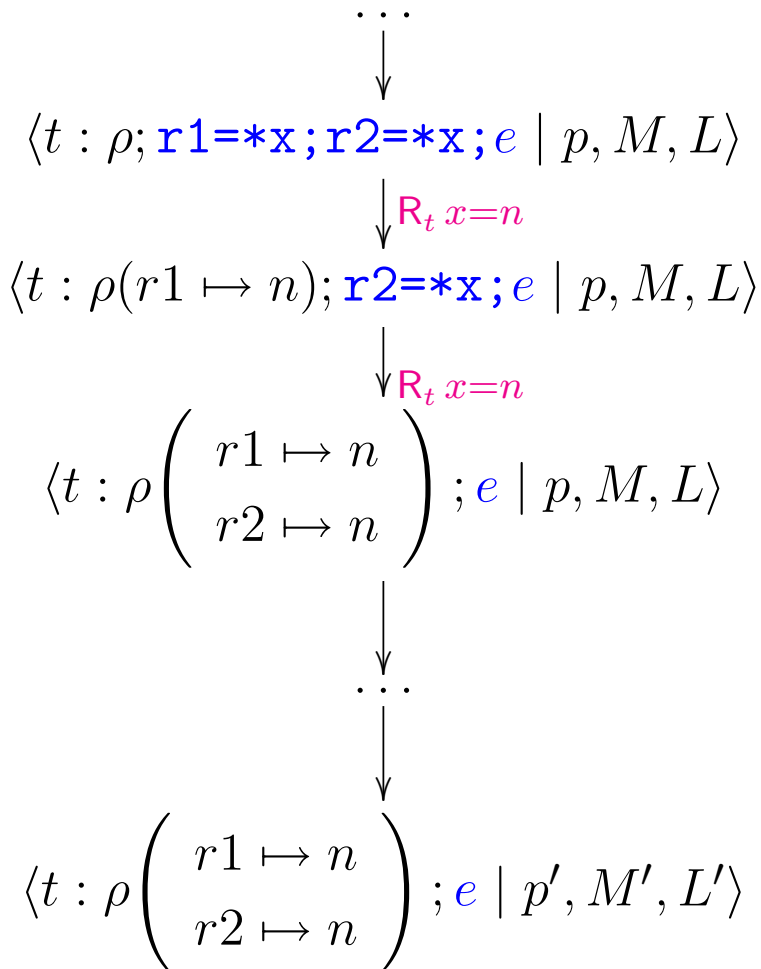
Optimised program



Simulating Elimination

Original program



Optimised program



Simulating Elimination

Original program

Optimised program



Note: there is a bit of cheating w.r.t. the semantics.

Elimination Formally

$e_1 \rightsquigarrow e_2$ e_2 is an elimination of e_1

$$\frac{}{r_1 = *x; r_2 = *x \rightsquigarrow r_1 = *x; r_2 = r_1} \text{EELIM}$$

$$\frac{}{e \rightsquigarrow e} \text{EID}$$

$$\frac{e \rightsquigarrow e'}{r = e \rightsquigarrow r = e'} \text{EASSIGN_CONTEXT}$$

$$\frac{e \rightsquigarrow e'}{*x = e \rightsquigarrow *x = e'} \text{EWRITE_CONTEXT}$$

$$\frac{\begin{array}{l} e_1 \rightsquigarrow e'_1 \\ e_2 \rightsquigarrow e'_2 \end{array}}{e_1; e_2 \rightsquigarrow e'_1; e'_2} \text{ESEQ_CONTEXT}$$

$$\frac{e \rightsquigarrow e'}{\text{print } e \rightsquigarrow \text{print } e'} \text{EPRINT_CONTEXT}$$

Elimination Formally

$$e_1 \rightsquigarrow e'_1$$

$$e_2 \rightsquigarrow e'_2$$

$$e_3 \rightsquigarrow e'_3$$

$$e_4 \rightsquigarrow e'_4$$

$$\frac{}{\text{if } e_1 = e_2 \text{ then } e_3 \text{ else } e_4 \rightsquigarrow \text{if } e'_1 = e'_2 \text{ then } e'_3 \text{ else } e'_4}$$

EIF_CONTEXT

$$\boxed{p_1 \rightsquigarrow p_2} \quad p_2 \text{ is an elimination of } p_1$$

$$\frac{e \rightsquigarrow e'}{t:\rho;e \rightsquigarrow t:\rho;e'} \quad \text{ETHREAD}$$

$$\frac{p_1 \rightsquigarrow p'_1 \quad p_2 \rightsquigarrow p'_2}{p_1|p_2 \rightsquigarrow p'_1|p'_2} \quad \text{EPAR}$$

Proof

We want to show that for every trace of the optimised program, there is a trace of the original program with the same behaviour.

What do we mean by the trace of program $\langle p, M, L \rangle$, exactly?

Proof

We want to show that for every trace of the optimised program, there is a trace of the original program with the same behaviour.

What do we mean by the trace of program $\langle p, M, L \rangle$, exactly?

We write $\langle p, M, L \rangle \xrightarrow{t} \langle p', M', L' \rangle$ if

$$\langle p, M, L \rangle \xrightarrow{l_1} \langle p_1, M_1, L_1 \rangle \xrightarrow{l_2} \dots \xrightarrow{l_n} \langle p', M', L' \rangle$$

for some p_i, M_i, L_i and l_i such that $t = [l_1, \dots, l_n]$.

Proof

We want to show that for every trace of the optimised program, there is a trace of the original program with the same behaviour.

What do we mean by the trace of program $\langle p, M, L \rangle$, exactly?

We write $\langle p, M, L \rangle \xrightarrow{t} \langle p', M', L' \rangle$ if

$$\langle p, M, L \rangle \xrightarrow{l_1} \langle p_1, M_1, L_1 \rangle \xrightarrow{l_2} \dots \xrightarrow{l_n} \langle p', M', L' \rangle$$

for some p_i, M_i, L_i and l_i such that $t = [l_1, \dots, l_n]$.

Then we can say that $\langle p, M, L \rangle$ has trace t (written $\langle p, M, L \rangle \downarrow t$) if $\langle p, M, L \rangle \xrightarrow{t} \langle p', M', L' \rangle$ for some p', M' and L' .

Proof - failed attempt

Let $p \rightsquigarrow q$. We want to show that for any t , if $\langle q, M, L \rangle \downarrow t$, then there is t' such that $\langle p, M, L \rangle \downarrow t'$ and $P(t) = P(t')$, where $P(t)$ is the subsequence of observable events in t .

1. Strategy: by induction on the length of trace t .
2. **Base case** is trivial.

Proof - failed attempt

Let $p \rightsquigarrow q$. We want to show that for any t , if $\langle q, M, L \rangle \downarrow t$, then there is t' such that $\langle p, M, L \rangle \downarrow t'$ and $P(t) = P(t')$, where $P(t)$ is the subsequence of observable events in t .

1. Strategy: by induction on the length of trace t .
2. **Base case** is trivial.
3. **Induction step** is interesting: We have

$\langle q, M, L \rangle \xrightarrow{t} \langle q', M', L' \rangle$ where $0 < |t|$.

Suppose $t = t' ++ [l]$, so we have

$\langle q, M, L \rangle \xrightarrow{t'} \langle \bar{q}, \bar{M}, \bar{L} \rangle \xrightarrow{l} \langle q', M', L' \rangle$ and we can use the

induction hypothesis to get $\langle p, M, L \rangle \xrightarrow{t'} \langle \hat{p}, \hat{M}, \hat{L} \rangle$.

Proof - failed attempt

Let $p \rightsquigarrow q$. We want to show that for any t , if $\langle q, M, L \rangle \downarrow t$, then there is t' such that $\langle p, M, L \rangle \downarrow t'$ and $P(t) = P(t')$, where $P(t)$ is the subsequence of observable events in t .

1. Strategy: by induction on the length of trace t .
2. **Base case** is trivial.
3. **Induction step** is interesting: We have

$\langle q, M, L \rangle \xrightarrow{t} \langle q', M', L' \rangle$ where $0 < |t|$.

Suppose $t = t' ++ [l]$, so we have


$\langle q, M, L \rangle \xrightarrow{t'} \langle \bar{q}, \bar{M}, \bar{L} \rangle \xrightarrow{l} \langle q', M', L' \rangle$ and we can use the induction hypothesis to get $\langle p, M, L \rangle \xrightarrow{t'} \langle \hat{p}, \hat{M}, \hat{L} \rangle$.

Uhm, this is no good because we **do not know anything** about $\hat{p}, \hat{M}, \hat{L}$!

Formal proof - simulation relation

Idea: we need to **relate the state of the optimised and original programs** so that we can **simulate the step and keep being related**. Our simulation relation \sim must:

- Relate a program and its eliminated counterpart, i.e., if $p \rightsquigarrow q$, then $p \sim q$.
- But it must also account for the intermediate step:



Simulation relation formally

Simulation relation for elimination and intermediate states:

$$\frac{e \rightsquigarrow e'}{\rho, e \sim \rho, e'} \quad \text{SEELIM}$$

$$\frac{}{\rho \oplus (r_1 \mapsto n), r_2 = n \sim \rho, r_1 = n; r_2 = r_1} \quad \text{SEASSIGN_N}$$

$$\frac{}{\rho \oplus (r_1 \mapsto n), r_2 = n \sim \rho \oplus (r_1 \mapsto n), n; r_2 = r_1} \quad \text{SESEQ}$$

$$\frac{}{\rho \oplus (r_1 \mapsto n), r_2 = n \sim \rho \oplus (r_1 \mapsto n), r_2 = r_1} \quad \text{SEASSIGN_REG}$$

Simulation relation formally

Simulation relation context rules:

$$\frac{\begin{array}{l} \rho, e_1 \sim \rho', e'_1 \\ e_2 \rightsquigarrow e'_2 \end{array}}{\rho, e_1; e_2 \sim \rho', e'_1; e'_2} \quad \text{SESEQ_CONTEXT}$$

$$\frac{\rho, e \sim \rho', e'}{\rho, \text{print } e \sim \rho', \text{print } e'} \quad \text{SEPRINT_CONTEXT}$$

$$\frac{\rho, e \sim \rho', e'}{\rho, r = e \sim \rho', r = e'} \quad \text{SEASSIGN_CONTEXT}$$

$$\frac{\rho, e \sim \rho', e'}{\rho, *x = e \sim \rho', *x = e'} \quad \text{SEWRITE_CONTEXT}$$

Simulation relation formally

Simulation relation lifting to processes:

$p_1 \sim p_2$ p_1 simulates p_2

$$\frac{\rho, e \sim \rho', e'}{t:\rho;e \sim t:\rho';e'} \quad \text{SP}_{\text{THREAD}}$$

$$\frac{p_1 \sim p'_1 \quad p_2 \sim p'_2}{p_1|p_2 \sim p'_1|p'_2} \quad \text{SP}_{\text{PAR}}$$

Simulation of expressions

First, we should make sure that **a step in an optimised expression can be simulated by step in any related expression**. Formally, we want to establish that if $\rho_1, e_1 \sim \rho'_1, e'_1$, $\rho'_1; e'_1 \xrightarrow{l} \rho'_2; e'_2$ then there are ρ_2, e_2 and trace l' such that $\rho_1; e_1 \xrightarrow{l'} \rho_2; e_2$, $\rho_2, e_2 \sim \rho'_2, e'_2$, $P([l]) = P(l')$ (we should also show that l and l' have the same effect on locks/memory).

Simulation of expressions

First, we should make sure that **a step in an optimised expression can be simulated by step in any related expression**. Formally, we want to establish that if $\rho_1, e_1 \sim \rho'_1, e'_1$, $\rho'_1; e'_1 \xrightarrow{l} \rho'_2; e'_2$ then there are ρ_2, e_2 and trace l' such that $\rho_1; e_1 \xrightarrow{l'} \rho_2; e_2$, $\rho_2, e_2 \sim \rho'_2, e'_2$, $P([l]) = P(l')$ (we should also show that l and l' have the same effect on locks/memory).

More graphically, we want show:

$$\begin{array}{ccc} \rho_1, e_1 & & \rho_1, e_1 \xrightarrow{l'} \rho_2, e_2 \\ \left. \vphantom{\rho_1, e_1} \right\} & \Rightarrow \exists \rho_2, e_2, l'. & \left. \vphantom{\rho_1, e_1} \right\} \wedge P([l]) = P(l') \\ \rho'_1, e'_1 \xrightarrow{l} \rho'_2, e'_2 & & \rho'_1, e'_1 \xrightarrow{l} \rho'_2, e'_2 \end{array}$$

Simulation of expressions

First, we should make sure that **a step in an optimised expression can be simulated by step in any related expression**. Formally, we want to establish that if $\rho_1, e_1 \sim \rho'_1, e'_1$, $\rho'_1; e'_1 \xrightarrow{l} \rho'_2; e'_2$ then there are ρ_2, e_2 and trace l' such that $\rho_1; e_1 \xrightarrow{l'} \rho_2; e_2$, $\rho_2, e_2 \sim \rho'_2, e'_2$, $P([l]) = P(l')$ (we should also show that l and l' have the same effect on locks/memory).

Proof strategy: by induction on the derivation of $\rho'_1; e'_1 \xrightarrow{l} \rho'_2; e'_2$ and then by case analysis on $\rho_1, e_1 \sim \rho'_1, e'_1$.

Most of the induction cases are trivial. **For example**, if

$\rho'_1; e'_1 \xrightarrow{l} \rho'_2; e'_2$ was obtained by the **READ** rule, then $e'_1 = *x$, hence $e_1 = *x$ and $\rho_1 = \rho'_1$ (by definition of \sim). As a result, we can take $\rho_2 = \rho'_2$, $e_2 = e'_2$ and $l' = [l]$.

Simulation of 'SEQ_CONTEXT'

If $\rho'_1; e'_1 \xrightarrow{l} \rho'_2; e'_2$ was derived by SEQ_CONTEXT then $e'_1 = e_1^1; e_1^2$. In this case, the case analysis on $\rho_1, e_1 \sim \rho'_1, e'_1$ admits four rules: SEELIM, SEASSIGN_N and SESEQ.

Consider, for example, SEASSIGN_N. Then

$e'_1 = (r_1 = n; r_2 = r_1)$, $\rho_1 = \rho'_1 \oplus (r_1 \mapsto n)$, $e_1 = (r_2 = n)$. Since only the CONTEXT_SEQ and VAR_WRITE rules apply to e'_1 , we have $\rho'_2 = \rho'_1 \oplus (r_1 \mapsto n)$ and $e'_2 = (n; r_2 = r_1)$. Let us set $\rho_2 = \rho_1$, $e_2 = e_1$ and $l' = []$. Using SEASSIGN_SEQ, we obtain $\rho_1, e_1 \sim \rho'_2, e'_2$. Since $l = \tau$, we have $P(l') = P([l])$.

Simulation of 'SEQ_CONTEXT'

If $\rho'_1; e'_1 \xrightarrow{l} \rho'_2; e'_2$ was derived by SEQ_CONTEXT then $e'_1 = e_1^1; e_1^2$. In this case, the case analysis on $\rho_1, e_1 \sim \rho'_1, e'_1$ admits four rules: SEELIM, SEASSIGN_N and SESEQ.

Consider, for example, SEASSIGN_N. Then $e'_1 = (r_1 = n; r_2 = r_1)$, $\rho_1 = \rho'_1 \oplus (r_1 \mapsto n)$, $e_1 = (r_2 = n)$. Since only the CONTEXT_SEQ and VAR_WRITE rules apply to e'_1 , we have $\rho'_2 = \rho'_1 \oplus (r_1 \mapsto n)$ and $e'_2 = (n; r_2 = r_1)$. Let us set $\rho_2 = \rho_1$, $e_2 = e_1$ and $l' = []$. Using SEASSIGN_SEQ, we obtain $\rho_1, e_1 \sim \rho'_2, e'_2$. Since $l = \tau$, we have $P(l') = P([l])$.

The other cases are similarly tedious. You **do not want to do this kind of reasoning manually!** We advise to use Isabelle or HOL or Coq to check your proofs.

Proof plumbing

We want to show that for any t , if $\langle q, M, L \rangle \xrightarrow{t} \langle q', M', L' \rangle$, then there are t' and p' such that $\langle p, M, L \rangle \xrightarrow{t'} \langle p', M', L' \rangle$, $p' \sim q'$ and $P(t) = P(t')$.

Base case still trivial. **Induction step:**

$$\begin{array}{c} \langle p, M, L \rangle \\ \{ \\ \langle q, M, L \rangle \xrightarrow{\bar{t}} \langle \bar{q}, \bar{M}, \bar{L} \rangle \xrightarrow{l} \langle q', M', L' \rangle \end{array}$$

Use induction hypothesis!

Proof plumbing

We want to show that for any t , if $\langle q, M, L \rangle \xrightarrow{t} \langle q', M', L' \rangle$, then there are t' and p' such that $\langle p, M, L \rangle \xrightarrow{t'} \langle p', M', L' \rangle$, $p' \sim q'$ and $P(t) = P(t')$.

Base case still trivial. **Induction step:**

By induction hypothesis we obtain \bar{t}' and \bar{p} such that $\bar{p} \sim \bar{q}$, $P(\bar{t}) = P(\bar{t}')$ and

$$\begin{array}{ccc} \langle p, M, L \rangle & \xrightarrow{\bar{t}'} & \langle \bar{p}, \bar{M}, \bar{L} \rangle \\ \left. \vphantom{\langle p, M, L \rangle} \right\} & & \left. \vphantom{\langle \bar{p}, \bar{M}, \bar{L} \rangle} \right\} \\ \langle q, M, L \rangle & \xrightarrow{\bar{t}} & \langle \bar{q}, \bar{M}, \bar{L} \rangle \xrightarrow{l} \langle q', M', L' \rangle \end{array}$$

Proof plumbing

We need to find p' and t' such that $P(t') = P([l'])$ and

$$\begin{array}{ccc} \langle \bar{p}, \bar{M}, \bar{L} \rangle & \xrightarrow{t'} & \langle p', M', L' \rangle \\ \{ & & \} \\ \langle \bar{q}, \bar{M}, \bar{L} \rangle & \xrightarrow{l} & \langle q', M', L' \rangle \end{array}$$

This is an easy proof by **induction on the derivation of** $\langle \bar{q}, \bar{M}, \bar{L} \rangle \xrightarrow{l} \langle q', M', L' \rangle$, using our (sketched) expression simulation lemma to do the heavy lifting.

Compiler Optimisations

Correctness

Optimisations in SC

Optimisations of Data-race-free Programs

Data race freedom

Limitations of the DRF principle

Correctness under the DRF principle

Out-of-thin-air values

Data Race Freedom Principle

Since many common optimisations are not valid in SC, language designers specify **relaxed models**.

Insight: **to observe an optimisation** one needs to be able to observe **other thread's memory accesses** at the same time when they are performed — **data race!**

So, in data race free programs, common optimisations are not observable.

DRF principle: each behaviour of a data race free program (after optimisation) **must be the same as the behaviour of some (SC) execution of the program** (before optimisation).

Says nothing about programs with data races!

Compiler Optimisations

Correctness

Optimisations in SC

Optimisations of Data-race-free Programs

Data race freedom

Limitations of the DRF principle

Correctness under the DRF principle

Out-of-thin-air values

What is DRF, exactly?

There is **no execution with adjacent conflicting memory accesses from different threads**, where a pair of memory accesses is **conflicting** if they access the same memory location and at least one of them is a write.

Program with a data race:

```
*y = 1 | if *x = 1
*x = 1 | then print *y
```

A racy execution:

$W_{t_1} y=1, W_{t_1} x=1, R_{t_2} x=1, R_{t_2} y=1, P_{t_2} 1$

What is DRF, exactly?

There is **no execution with adjacent conflicting memory accesses from different threads**, where a pair of memory accesses is **conflicting** if they access the same memory location and at least one of them is a write.

Data race free program:

<code>*y = 1</code>	<code>lock x;</code>
<code>lock x</code>	<code>r = *x;</code>
<code>*x = 1</code>	<code>unlock x;</code>
<code>unlock x</code>	<code>if r = 1</code>
	<code>then print *y</code>

What is DRF, exactly?

There is **no execution with adjacent conflicting memory accesses from different threads**, where a pair of memory accesses is **conflicting** if they access the same memory location and at least one of them is a write.

Data race free program:

```
if *x = 1      |      if *y = 1
then *y = 1    |      then *x = 1
```

Note that the **writes cannot be executed** in any SC execution!
So they cannot participate in a data race...

Another DRF definition

In literature, you will often find a different different definition of data race freedom — **using happens-before relation**.

The intuition:

- happens-before relates events that are ordered in time by **synchronisation** or by **sequencing in a thread**.
- a **data race** is a pair of conflicting memory accesses that is **not ordered by happens-before**.
- a pair of memory accesses is **conflicting** if they access the same memory location and at least one of them is a write.

Happens-before order

Suppose that t is a trace $t = [l_1, \dots, l_n]$.

Let $1 \leq i, j \leq n$. We say that


- i **synchronizes-with** j if $i < j$, $l_i = U_{t_i} x$ and $l_j = L_{t_j} x$.
- i is before j in **program order** if $i < j$ and l_i, l_j belong to the same thread.
- i **happens-before** j if $\langle i, j \rangle$ is in the **transitive closure** of the union of the **synchronizes-with** and **program orders**.

Happens-before order

Suppose that t is a trace $t = [l_1, \dots, l_n]$.

Let $1 \leq i, j \leq n$. We say that

- i **synchronizes-with** j if $i < j$, $l_i = U_{t_i} x$ and $l_j = L_{t_j} x$.
- i is before j in **program order** if $i < j$ and l_i, l_j belong to the same thread.
- i **happens-before** j if $\langle i, j \rangle$ is in the **transitive closure** of the union of the synchronizes-with and program orders.



DRF via happens-before

We say that program is **hb-data-race-free** if in all executions, all different pairs of **conflicting actions** in the executions are **ordered by the happens-before order**. Here, actions are conflicting if they are memory accesses to the same location and at least one of them is a write.

Theorem (cf. Boehm, Adve, PLDI 2008): An execution is **hb-data-race-free if and only if** it is **data-race-free**.

This a very useful theorem to prove correctness of optimisations under the DRF principle.

Compiler Optimisations

Correctness

Optimisations in SC

Optimisations of Data-race-free Programs

Data race freedom

Limitations of the DRF principle

Correctness under the DRF principle

Out-of-thin-air values

Limitations of the DRF

The DRF principle allows most optimisations, but not all.

The **dangerous “optimisations”** are usually the ones that **introduce shared-memory accesses** (this is very rare):

- Write introducing code transformations
 - often speculatively, or
 - writes that would be redundant in sequential code.
- Read introducing optimisations
 - often harmless by themselves, but can be dangerous in combination with optimisations that are correct under DRF.

Loop variable register promotion

Is the following transformation valid (in any context)?

```
while (j != 0) (  
    j = j - 1;  
    *x = *x - 1 )  
→  
r = *x;  
while (j != 0) (  
    j = j - 1;  
    r = r - 1)  
*x = r
```


Loop variable register promotion

Assume the initial state $*x = 0$.

<pre>while (j!=0) (j = j - 1; *x = *x - 1)</pre>		<pre>*x=1 print *x</pre>	→	<pre>r=*x; while (j!=0) (j = j - 1; r = r - 1) *x=r</pre>		<pre>*x=1 print *x</pre>
---	--	--------------------------	---	--	--	--------------------------

Can the process **print 0**?

Aggressive write speculation

Is the following transformation valid (in any context)?

```
if *x = 1
then *y = 1
```

→

```
r = *y
*y = 1;
if *x != 1
then *y = r
```

Aggressive write speculation

Assume initial state $*x = *y = 0$.

<pre>if *x=1 then *y=1</pre>		<pre>if *y=1 then *x=1 print *x</pre>	→	<pre>r=*y *y=1; if *x!=1 then *y=r</pre>		<pre>if *y=1 then *x=1 print *x</pre>
------------------------------	--	---------------------------------------	---	--	--	---------------------------------------

Can the program **print 1**?

Read introduction?

It is correct to introduce a read and then throw away its value.

But...

Read introduction?

It is correct to introduce a read and then throw away its value.

But...

Observe that it is **correct** (under DRF) to do **CSE across lock** (or `unlock` but not both):

```
r1 = *x; lock m; r2 = *x → r1 = *x; lock m; r2 = r1
```

Read introduction?

It is correct to introduce a read and then throw away its value.

But...

Consider the following program and assume that initially $*x = *y = 0$.

```
lock m;  
*x = 1;  
print *y  
unlock m
```

```
lock m;  
*y = 1;  
print *x  
unlock m
```

Can the program **print two zeros**?

Read introduction?

It is correct to introduce a read and then throw away its value.

But...

Consider the following program and assume that initially $*x = *y = 0$.

<code>r1=*y;</code>	<code>r2=*x;</code>
<code>lock m;</code>	<code>lock m;</code>
<code>*x = 1;</code>	<code>*y = 1;</code>
<code>print *y</code>	<code>print *x</code>
<code>unlock m</code>	<code>unlock m</code>

Can the program **print two zeros**?

Read introduction?

It is correct to introduce a read and then throw away its value.

But...

Consider the following program and assume that initially $*x = *y = 0$.

<code>r1=*y;</code>	<code>r2=*x;</code>
<code>lock m;</code>	<code>lock m;</code>
<code>*x = 1;</code>	<code>*y = 1;</code>
<code>print *y</code>	<code>print *x</code>
<code>unlock m</code>	<code>unlock m</code>

Can the program **print two zeros**?

Read introduction?

It is correct to introduce a read and then throw away its value.

But...

Consider the following program and assume that initially $*x = *y = 0$.

<code>r1=*y;</code>	<code>r2=*x;</code>
<code>lock m;</code>	<code>lock m;</code>
<code>*x = 1;</code>	<code>*y = 1;</code>
<code>print r1</code>	<code>print r2</code>
<code>unlock m</code>	<code>unlock m</code>

Now the program **can print two zeros.**

Practical limitations of DRF principle

Data race detection is hard!

- There are **type systems** for DRF, but they are restrictive...
- There are **static analyses** for DRF, but they are imprecise and they must see whole program...
- There are **run-time data race detectors**, but they are slow and incomplete (they cannot guarantee data race freedom)...