

Semantics and tools for low-level concurrent programming





Semantics and tools for low-level concurrent programming



Compilers vs. programmers





Compilers vs. programmers



Constant propagation (an optimising compiler breaks your program)

A simple and innocent looking optimization:

Constant propagation (an optimising compiler breaks your program)

A simple and innocent looking optimization:

int x = 14;
int y = 7 - x / 2; int x = 14;
int y = 7 -
$$\frac{14}{2}$$

Consider the two threads below:

Intuitively, this program always prints 0

Constant propagation (an optimising compiler breaks your program)

A simple and innocent looking optimization:

Consider the two threads below:

Sun HotSpot JVM or GCJ: always prints 1.

Background: lock and unlock

• Suppose that two threads increment a shared memory location:

• If both threads read 0, (even in an ideal world) x = 1 is possible:

Background: lock and unlock

• Lock and unlock are primitives that prevent the two threads from interleaving their actions.

$$\mathbf{x} = \mathbf{0}$$

- lock(); tmp1 = *x; *x = tmp1 + 1; unlock(); lock(); tmp2 = *x; *x = tmp1 + 1; unlock();
- In this case, the interleaving below is forbidden, and we are guaranteed that x = 2 at the end of the execution.

Deferring an object's initialisation util first use: a big win if an object is never used (e.g. device drivers code). Compare:

```
int x = computeInitValue(); // eager initialization
... // clients refer to x
with:
int xValue() {
  static int x = computeInitValue(); // lazy initialization
  return x;
} ... // clients refer to xValue()
```

The singleton pattern

Lazy initialisation is a pattern commonly used. In C++ you would write:

```
Singleton::instance () -> method ();
```

```
But this code is not thread safe! Why?
```

Making the singleton pattern thread safe

A simple thread safe version:

```
class Singleton {
public:
    static Singleton *instance (void) {
        Guard<Mutex> guard (lock_); // only one thread at a time
        if (instance_ == NULL)
            instance_ = new Singleton;
        return instance_;
    }
private:
    static Mutex lock_;
    static Singleton *instance_;
};
```

Every call to instance must acquire and release the lock: excessive overhead.

Obvious (broken) optimisation

```
class Singleton {
public:
  static Singleton *instance (void) {
     if (instance == NULL) {
       Guard<Mutex> guard (lock ); // lock only if unitialised
       instance = new Singleton; }
    return instance ;
  }
private:
  static Mutex lock ;
  static Singleton *instance ;
};
```

```
Exercise: why is it broken?
```

Clever programmers use double-check locking

```
class Singleton {
public:
  static Singleton *instance (void) {
    // First check
    if (instance == NULL) {
       // Ensure serialization
       Guard<Mutex> guard (lock );
       // Double check
       if (instance_ == NULL)
         instance = new Singleton;
     }
    return instance ;
private: [..]
};
```

Idea: re-check that the Singleton has not been created after acquiring the lock.

Double-check locking: clever but broken

The instruction

```
instance_ = new Singleton;
```

does three things:

1) allocate memory

2) construct the object

3) assign to instance the address of the memory

Not necessarily in this order! For example:

```
instance_ = // 3
operator new(sizeof(Singleton)); // 1
new (instance_) Singleton // 2
```

If this code is generated, the order is 1,3,2.

Broken...

```
if (instance_ == NULL) { // Line 1
Guard<Mutex> guard (lock_);
if (instance_ == NULL) {
    instance_ =
        operator new(sizeof(Singleton)); // Line 2
        new (instance_) Singleton; }}
```

Thread 1:

executes through Line 2 and is suspended; at this point, instance_ is non-NULL, but no singleton has been constructed.

Thread 2:

executes Line 1, sees instance_ as non-NULL, returns, and dereferences the pointer returned by Singleton (i.e., instance_).

Thread 2 attempts to reference an object that is not there yet!

Problem: You need a way to specify that step 3 come after steps 1 and 2.

There is no way to specify this in C++

Similar examples can be built for any programming language...

That pesky hardware (1)

Consider misaligned 4-byte accesses

$int32_t a = 0$		
a = 0x44332211	if (a == 0x00002211) print "error"	

(Disclaimer: compiler will normally ensure alignment)

Intel SDM x86 atomic accesses:

- *n*-bytes on an *n*-byte boundary (n = 1, 2, 4, 16)
- P6 or later: ... or if unaligned but within a cache line

Question: what about multi-word high-level language values?

That pesky hardware (1)

Consider misaligned 4-byte accesses

$int32_t a = 0$		
a	= 0x44332211	if (a == 0x00002211) print "error"
Disclaime		
ntel SDN	This is called a <i>o</i>	ut-of-thin air read:
n-byte	the program reads a value	
P6 or	that the programmer never wrote.	
Questior). what about multi-word n	ign-iever language values?

That pesky hardware (2)

Hardware optimisations can be observed by concurrent code:

Thread 0	Thread 1
x = 1	y = 1
print y	print x

At the end of some executions:

0 0

is printed on the screen, both on x86 and Power/ARM).



That pesky hardware (2)

...and differ between architectures...



Compilers vs. programmers





Compilers vs. programmers

Tension:

- the programmer wants to understand the code he writes
- the compiler and the hardware want to optimise it.

Which are the valid optimisations that the compiler or the hardware can perform without breaking the expected semantics of a concurrent program?

Which is the semantics of a concurrent program?

This lecture

Programming language models

- 1) soundness of compiler optimisations
- 2) data-race freedom
- 3) defining the semantics of a concurrent programming language

Tomorrow:

The C11/C++11 model in detail.

effect of VS2005 compiler optimisations on speed



A brief tour of compiler optimisations



effect of additional VS2005 optimisations on speed

World of optimisations

A typical compiler performs many optimisations.

gcc 4.4.1. with –O2 option goes through 147 compilation passes.

computed 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

A 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

World of optimisations

However 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

Optimisations of shared memory can be classified as:

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 \mathbf{r} is not free in \mathbf{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

Common subexpression elimination, some loop optimisations, code motion.

Normal memory access reordering:

Roach motel reordering:

memop; lock m → lock m; memop unlock m; memop → memop; unlock m where memop is *x=r1 or r1=*x

Memory access introduction

Can an optimisation introduce memory accesses?

Yes, but rarely:

Note that the loop body is not executed.

Memory access introduction

Back to our question now:

Which is the semantics of a concurrent program?

TNULE LITAL LITE TOUP DOUY IS HOL EXECULED.

Naive answer: enforce sequential consistency

Multiprocessors have a *sequentially consistent* shared memory when:

...the result of any execution is the same as if the operations of all the processors were executed in some sequential order, and the operations of each individual processor appear in this sequence in the order specified by its program...

Lamport, 1979.





Compilers, programmers & sequential consistency




Compilers, programmers & sequential consistency







Simple and intuitive programming model

Compilers, programmers & sequential consistency







Simple and intuitive programming model

A Case for an SC-Preserving Compiler

Daniel Marino[†] Abhayendra Singh*

Todd Millstein[†]

Madanlal Musuvathi[‡] Satish Narayanasamy*

[†]University of California, Los Angeles

*University of Michigan, Ann Arbor

[‡]Microsoft Research, Redmond

An SC-preserving compiler, obtained by restricting the optimization phases in LLVM, a state-of-the-art C/C++ compiler, incurs an average slowdown of 3.8% and a maximum slowdown of 34% on a set of 30 programs from the SPLASH-2, PARSEC, and SPEC CINT2006 benchmark suites.

This study supposes that the hardware is SC.



A recent paper at ISCA mentions a 6.2% slowdown wrt TSO to enforce end-to-end SC on dedicated hardware.

A Case for an SC-Preserving Compiler



A compiler is correct if any behaviour of the compiled program could be exhibited by the original program.

i.e. for any execution of the compiled program, there is an execution of the source program with the same observable behaviour.

Intuition: we represent programs as sets of memory action traces, where the trace is a sequence of memory actions of a single thread.

Intuition: the observable behaviour of an execution is the subtrace of external actions.

 $P_1 = \mathbf{*x} = \mathbf{1} \begin{vmatrix} \mathbf{r1} = \mathbf{*x}; & \mathbf{r2} = \mathbf{*x}; \\ \text{if } \mathbf{r1} = \mathbf{r2} \text{ then print 1 else print 2} \end{vmatrix}$ $P_2 = \mathbf{*x} = \mathbf{1} \begin{vmatrix} \mathbf{r1} = \mathbf{*x}; & \mathbf{r2} = \mathbf{r1}; \\ \text{if } \mathbf{r1} = \mathbf{r2} \text{ then print 1 else print 2} \end{vmatrix}$

Is the transformation from P1 to P2 correct (in an SC semantics)?

$$P_1 = *x = 1 \begin{vmatrix} r1 = *x; r2 = *x; \\ if r1=r2 \text{ then print 1 else print 2} \end{vmatrix}$$
$$P_2 = *x = 1 \begin{vmatrix} r1 = *x; r2 = r1; \\ if r1=r2 \text{ then print 1 else print 2} \end{vmatrix}$$

 $P_1 = *x = 1 \begin{vmatrix} r1 = *x; r2 = *x; \\ if r1 = r2 \text{ then print 1 else print 2} \end{vmatrix}$ $P_2 = *x = 1 \begin{vmatrix} r1 = *x; r2 = r1; \\ if r1 = r2 \text{ then print 1 else print 2} \end{vmatrix}$

Executions of P1:

$$\begin{split} & \mathsf{W}_{t_1} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{W}_{t_1} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{P}_{t_2} \, 2 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{R}_{t_2} \, x{=}0, \mathsf{W}_{t_1} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{R}_{t_2} \, x{=}0, \mathsf{P}_{t_2} \, 1, \mathsf{W}_{t_1} \, x{=}1 \end{split}$$

 $P_1 = *x = 1 \begin{vmatrix} r1 = *x; r2 = *x; \\ if r1 = r2 \text{ then print 1 else print 2} \end{vmatrix}$ $P_2 = *x = 1 \begin{vmatrix} r1 = *x; r2 = r1; \\ if r1 = r2 \text{ then print 1 else print 2} \end{vmatrix}$

Executions of P1:

Executions of P2:

$$\begin{split} & \mathsf{W}_{t_1} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{W}_{t_1} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{P}_{t_2} \, 2 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{R}_{t_2} \, x{=}0, \mathsf{W}_{t_1} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{R}_{t_2} \, x{=}0, \mathsf{P}_{t_2} \, 1, \mathsf{W}_{t_1} \, x{=}1 \end{split}$$

$$\begin{split} & \mathsf{W}_{t_1} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{W}_{t_1} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{P}_{t_2} \, 1, \mathsf{W}_{t_1} \, x{=}1 \end{split}$$

 $P_1 = *x = 1 \begin{vmatrix} r1 = *x; r2 = *x; \\ if r1 = r2 \text{ then print 1 else print 2} \end{vmatrix}$ $P_2 = *x = 1 \begin{vmatrix} r1 = *x; r2 = r1; \\ if r1 = r2 \text{ then print 1 else print 2} \end{vmatrix}$

Executions of P1:

Executions of P2:

$$\begin{split} & \mathsf{W}_{t_1} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{W}_{t_1} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{P}_{t_2} \, 2 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{R}_{t_2} \, x{=}0, \mathsf{W}_{t_1} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{R}_{t_2} \, x{=}0, \mathsf{P}_{t_2} \, 1, \mathsf{W}_{t_1} \, x{=}1 \end{split}$$

$$\begin{split} & \mathsf{W}_{t_1} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{W}_{t_1} \, x{=}1, \mathsf{P}_{t_2} \, 1 \\ & \mathsf{R}_{t_2} \, x{=}0, \mathsf{P}_{t_2} \, 1, \mathsf{W}_{t_1} \, x{=}1 \end{split}$$

Behaviours of P1: $[P_{t_2} 1], [P_{t_2} 2]$

Behaviours of P2: $[P_{t_2} 1]$

 $P_1 = *x = 1 \begin{vmatrix} r1 = *x; r2 = *x; \\ if r1 = r2 then print 1 else print 2 \end{vmatrix}$ $P_2 = *x = 1$ | r1 = *x; r2 = r1; if r1=r2 then print 1 else print 2 Executions of P2. Executions of D1. W_{t_1} It is correct to rewrite P1 into P2, but not the opposite! R_{t_2} R_{t_2} R_{t_2}

Behaviours of P1: $[P_{t_2} 1], [P_{t_2} 2]$

Behaviours of P2: $[P_{t_2} 1]$

General CSE incorrect in SC

There is only one execution with a printing behaviour:

$$\mathsf{W}_{t_1} x = 1, \mathsf{W}_{t_1} y = 1, \mathsf{R}_{t_2} x = 1, \mathsf{W}_{t_2} x = 2, \mathsf{W}_{t_2} y = 2, \mathsf{R}_{t_1} y = 2, \mathsf{R}_{t_1} x = 2, \mathsf{P}_{t_1} 2 = 2, \mathsf{R}_{t_1} x = 2, \mathsf{R}_{t_1}$$

General CSE incorrect in SC

But a compiler would optimise to:

General CSE incorrect in SC

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

$$\mathsf{W}_{t_1} x \! = \! 1, \mathsf{W}_{t_1} y \! = \! 1, \mathsf{R}_{t_2} x \! = \! 1, \mathsf{W}_{t_2} x \! = \! 2, \mathsf{W}_{t_2} y \! = \! 2, \mathsf{R}_{t_1} y \! = \! 2, \mathsf{P}_{t_1} 1$$

So the optimisation is not correct.

١

Reordering incorrect

*x = 1;	*y = 1;	r1 = *y	*y = 1;
r1 = *y	$r2 = *x; \Rightarrow$	> *x = 1;	r2 = *x;
print r1	print r2	print r1	print r2

Again, the optimised program exhibits a new behaviour:

$$\begin{split} & [\mathsf{P}_{t_1} \, {\color{black} 0}, \mathsf{P}_{t_2} \, {\color{black} 1}] \\ & [\mathsf{P}_{t_1} \, {\color{black} 1}, \mathsf{P}_{t_2} \, {\color{black} 0}] \\ & [\mathsf{P}_{t_1} \, {\color{black} 1}, \mathsf{P}_{t_2} \, {\color{black} 1}] \end{split}$$

$$\begin{split} & [\mathsf{P}_{t_1} \, {}^0, \mathsf{P}_{t_2} \, {}^1] \\ & [\mathsf{P}_{t_1} \, {}^1, \mathsf{P}_{t_2} \, {}^0] \\ & [\mathsf{P}_{t_1} \, {}^1, \mathsf{P}_{t_2} \, {}^1] \\ & [\mathsf{P}_{t_1} \, {}^0, \mathsf{P}_{t_2} \, {}^0] \end{split}$$

Elimination of adjacent accesses

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

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

The basic idea: whenever we perform the read r1 = *x in the optimised program, we perform *both* reads in the source program.

Elimination of adjacent accesses

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

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

Can we define a model that:

1) enables more optimisations than SC, and

2) retains the simplicity of SC?

Alternative answer: data-race freedom

Data-race freedom

Our examples again:

Thread 0	Thread 1	
*y = 1	if *x == 1	
*x = 1	then print *y	

the problematic transformations
 (e.g. swapping the two writes in

Observable behaviour: 0

thread 0) do not change the meaning of single-threaded programs;

 the problematic transformations are detectable only by code that allows two threads to access the same data simultaneously in conflicting ways (e.g. one thread writes the datas read by the other).

Data-race freedom



The basic solution	Thread 0	Thread 1
	*y = 1	if *x == 1
Prohibit data races	*x = 1	then print *y

Observable behaviour: 0

Defined as follows:

- two memory operations **conflict** if they access the same memory location and at least one is a store operation;
- a SC execution (interleaving) contains a data race if two conflicting operations corresponding to different threads are adjacent (maybe executed concurrently).

Example: a data race in the example above:

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

The basic solution	Thread 0	Thread 1
	*y = 1	if *x == 1
Prohibit data races	*x = 1	then print *y

Observable behaviour: 0



Example: a data race in the example above:

$$\mathsf{W}_{t_1} \, y{=}1, \mathsf{W}_{t_1} \, x{=}1, \mathsf{R}_{t_2} \, x{=}1, \mathsf{R}_{t_2} \, y{=}1, \mathsf{P}_{t_2} \, 1$$

Locks

No lock(I) can appear in the interleaving unless prior lock(I) and unlock(I) calls from other threads balance.

Atomic variables

Allow concurrent access "exempt" from data races. Called volatile in Java.

Example:

Thread 0	Thread 1	
*y = 1	lock();	
lock();	tmp = *x;	
*X = 1	unlock();	
unlock();	if tmp = 1	
	then print *y	

How do we avoid data races? (focus on high-level languages)

Thread 0	Thread 1	
<pre>*y = 1 lock();</pre>	lock(); tmp = *x;	
*x = 1 unlock();	unlock(); if tmp = 1	
	then print *y	

This program is data-race free:

*y = 1; lock();*x = 1;unlock(); lock();tmp = *x;unlock(); if tmp=1 then print *y
*y = 1; lock(); tmp = *x; unlock(); lock(); *x = 1; unlock(); if tmp=1
*y = 1; lock(); tmp = *x; unlock(); if tmp=1; lock(); *x = 1; unlock();
lock();tmp = *x;unlock(); *y = 1; lock(); *x = 1; unlock(); if tmp=1
lock(); tmp = *x; unlock(); if tmp=1; *y = 1; lock(); *x = 1; unlock();
lock();tmp = *x;unlock(); *y = 1; if tmp=1; lock(); *x = 1; unlock();

How do we avoid data races? (focus on high-level languages)



How do we avoid data races? (focus on high-level languages)



Another example of DRF program

Exercise: is this program DRF?

Thread 0	Thread 1	
if *x == 1	if *y == 1	
then $*y = 1$	then $*x = 1$	

Another example of DRF program

Exercise: is this program DRF?

Thread 0	Thread 1
if *x == 1	if *y == 1
then $*y = 1$	then $*x = 1$

Answer: yes!

The writes cannot be executed in any SC execution, so they cannot participate in a data race.

Another example of DRF program

Exercise: is this program DRF?

Thread 0	Thread 1	
if *x == 1	if *y == 1	
then $*y = 1$	then $*x = 1$	



Validity of compiler optimisations, summary

Transformation	SC	DRF
Memory trace preserving transformations	\checkmark	\checkmark
Redundant read after read elimination	✓*	\checkmark
Redundant read after write elimination	✓*	\checkmark
Irrelevant read elimination	\checkmark	\checkmark
Redundant write before write elimination	✓*	\checkmark
Redundant write after read elimination	✓*	\checkmark
Irrelevant read introduction	\checkmark	×
Normal memory accesses reordering	×	\checkmark
Roach-motel reordering	\times (\checkmark for locks)	\checkmark
External action reordering	×	\checkmark

* Optimisations legal only on adjacent statements.

Validity of compiler optimisations, summary



Compilers, programmers & data-race freedom





Compilers, programmers & data-race freedom





Compilers, programmers & data-race freedom







Intuitive programming model (but detecting races is tricky!)

Defining programming language memory models



Don't. No concurrency.

Poor match for current trends
Option 2



Don't. No shared memory

A good match for some problems, see e.g. Erlang and MPI.



Don't.

But language ensures data-race freedom

Possible (e.g. by ensuring data accesses protected by associated locks, or fancy effect type systems), but likely to be inflexible.



Don't.

But language ensures data-race freedom

Possible (e.g. by ensuring data accesses protected by associated locks, or fancy effect type systems), but likely to be inflexible.

What about these fancy racy algorithms?



Option 4



Don't.

Leave it (sort of) up to the hardware

Examples: MLton

MLton is a high performance ML-to-x86 compiler, with concurrency extensions. Accesses to ML refs will exhibit the underlying x86-TSO behaviour (compiler guarantees atomicity).

Option 5





Use data race freedom as a definition

1. Programs that race-free have only sequentially consistent behaviours

2. Programs that have a race in some execution can behave in any way

Sarita Adve & Mark Hill, 1990

Data race freedom as a definition

Posix is sort-of DRF



Applications shall ensure that access to any memory location by more than one thread of control (threads or processes) is restricted such that no thread of control can read or modify a memory location while another thread of control may be modifying it. Such access is restricted using functions that synchronize thread execution and also synchronize memory with respect to other threads.

Single Unix SPEC V3 & others

...again, model in informal prose...

Data race freedom as a definition

• Core of the C11/C++11 standard.

Hans Boehm & Sarita Adve, PLDI 2008.

with some escape mechanism called "low level atomics".

Mark Batty & al., POPL 2011.

• Part of the JSR-133 standard.

Jeremy Manson & Bill Pugh & Sarita Adve, PLDI 2008.

DRF gives no guarantees for untrusted code: a disaster for Java, which relies on unforgeable pointers for its security guarantees.

JSR-133 is DRF + some out-of-thin-air guarantees for all code.



Do.

Use data race freedom as a definition

Pro:

- simple

- strong guarantees for most code
- allows lots of freedom for compiler and hardware optimisations

Cons:

- undecidable premise
- can't write racy programs (escape mechanisms?)

Isn't this all obvious?



Isn't this all obvious?



Isn't this all obvious?



1. Uncertainity about details

Initially
$$x = y = 0$$

r1 := [x]; r2 := [y];
if (r1=1) || if (r2=1)
[y] := 1 [x] := 1

Is the outcome r1=r2=1 allowed?

1. Uncertainity about details

Is the outcome r1=r2=1 allowed?

- If the threads speculate that the values of x and y are 1, then each thread writes 1, validating the other thread speculation;
- such execution has a data race on x and y;
- however programmers would not envisage such execution when they check if their program is data-race free...

2. Compiler transformations introduce data races



- Many compilers perform transformations similar to the one above when a is declared as a bit field;
- May be visible to client code since the update to x.b by T2 may be overwritten by the store to the complete structure x.

And many more interesting examples...

2b. Compiler transformations introduce data races

• The vectorisation above might introduce races, but

• most compilers do things along these lines (introduce speculative stores).

3. "escape" mechanisms

Some frequently used idioms (atomic counters, flags, ...) do not require sequentially consistency.

Programmers wants optimal implementations of these idioms.

Speed, much more than safety, makes programmers happier.

A word on JSR-133

Goal 1: data-race free programs are sequentially consistent;

Goal 2: all programs satisfy some memory safety requirements;

Goal 3: common compiler optimisations are sound.

Goal 2: all programs satisfy some memory safety requirements. Programs should never read values that cannot be written by the program:

initially $x = y = 0$			
r1 := x	r2 := y		
y := r1	x := r2		
print r1	print r2		

the only possible result should be printing two zeros because no other value appears in or can be created by the program.

Out-of-thin-air

Under DRF, it is correct to speculate on values of writes:

The transformed program can now print 42. This will be theoretically possible in C++11, but not in Java.

The program above looks benign, why does Java care so much about out-of-thin-air?

Out-of-thin-air

Out-of-thin-air is not so bening for references. Compare:

initially $x = y = 0$			initially x = y = null	
r1 := x	r2 := y	and	r1 := x	r2 := y
y := r1	x := r2	and	y := r1	x := r2
print r1	print r2			r2.run()

What should r2.run() call?

If we allow out-of-thin-air, then it could do anything!



Goal 1: data-race free programs are sequentially consistent;

Goal 2: all programs satisfy some memory safety requirements;

Goal 3: common compiler optimisations are sound.

The model is intricate, and fails to meet goal 3.

An example: should the source program print 1? can the optimised program print 1?

Jaroslav Ševčík, David Aspinall, ECOOP 2008

The end?

/*

C11/C++11 is not yet implemented by mainstream compilers, and low-level atomics are hard to use (just google for low-level atomics).

How are interesting concurrent algorithms currently implemented? Usually C plus asm!

Example: lockfree-lib, by Keir Fraser, starts with some macro definitions...

```
* I. Compare-and-swap.
*/
/*
* This is a strong barrier! Reads cannot be delayed beyond a later store.
* Reads cannot be hoisted beyond a LOCK prefix. Stores always in-order.
*/
#define CAS(_a, _0, _n)
({ __typeof__(_o) __o = _o;
   __asm__ __volatile__(
        "lock cmpxchg %3,%1"
        : "=a" (__o), "=m" (*(volatile unsigned int *)(_a))
        : "0" (__o), "r" (_n) );
___0;
})
```







Next lecture: the C11/C++11 memory model

This afternoon, 2pm: exercices...