Bounding Data Races in Space and Time

KC Sivaramakrishnan

University of Cambridge



OCaml Labs



Darwin College, Cambridge



1851 Royal Commission



- OCaml is an industrial-strength, functional programming language
 - ★ Projects: MirageOS unikernel, Coq proof assistant, F* programming language
 - Companies: Facebook (Hack, Flow, Infer, Reason), Microsoft (Everest, F*), JaneStreet (all trading & support systems), Docker (Docker for Mac & Windows), Citrix (XenStore)

- OCaml is an industrial-strength, functional programming language
 - ★ Projects: MirageOS unikernel, Coq proof assistant, F* programming language
 - Companies: Facebook (Hack, Flow, Infer, Reason), Microsoft (Everest, F*), JaneStreet (all trading & support systems), Docker (Docker for Mac & Windows), Citrix (XenStore)
- No multicore support!

- OCaml is an industrial-strength, functional programming language
 - ★ Projects: MirageOS unikernel, Coq proof assistant, F* programming language
 - Companies: Facebook (Hack, Flow, Infer, Reason), Microsoft (Everest, F*), JaneStreet (all trading & support systems), Docker (Docker for Mac & Windows), Citrix (XenStore)
- No multicore support!
- Multicore OCaml
 - ★ Native support for *concurrency* and *parallelism* in OCaml
 - ★ Lead from OCaml Labs + (JaneStreet, Microsoft Research, INRIA).

• How do you reason about access to memory?

- How do you reason about access to memory?
 - ★ Spoiler: No single global sequentially consistent memory

- How do you reason about access to memory?
 - ★ Spoiler: No single global sequentially consistent memory
- Modern multicore processors reorder instructions for performance

- How do you reason about access to memory?
 - ★ Spoiler: No single global sequentially consistent memory
- Modern multicore processors reorder instructions for performance

Initially a = 0 && b = 0

Thread 2
b = 1
r2 = a

r1 == 0 && r2 ==0 ???

- How do you reason about access to memory?
 - ★ Spoiler: No single global sequentially consistent memory
- Modern multicore processors reorder instructions for performance

Initially a = 0 && b = 0

Thread 1	Thread 2	
a = 1	b = 1	
r1 = b	r2 = a	
r1 == 0 && r2 ==0 ???		

Allowed under x86, ARM, POWER

- How do you reason about access to memory?
 - ★ Spoiler: No single global sequentially consistent memory
- Modern multicore processors reorder instructions for performance

Initially a = 0 && b = 0

Thread 1	Thread 2	
a = 1	b = 1	
r1 = b	r2 = a	
r1 == 0 &&	r2 ==0 ???	Write buffering

Allowed under x86, ARM, POWER

- How do you reason about access to memory?
 - ★ Spoiler: No single global sequentially consistent memory
- Modern multicore processors reorder instructions for performance

Initially a = 0 && b = 0

Thread 1	Thread 2	
r1 = b	r2 = a	
a = 1	b = 1	
r1 == 0 && r	2 ==0 ???	Write buffering

Allowed under x86, ARM, POWER













Memory Model

- Unambiguous specification of program outcomes
 - ★ More than just thread interleavings



Memory Model



- Unambiguous specification of program outcomes
 - \star More than just thread interleavings
- Memory Model Desiderata
 - ★ Not too weak (good for programmers)
 - ★ Not too strong (good for hardware)
 - ★ Admits optimisations (good for compilers)
 - ★ Mathematically rigorous (good for verification)

Memory Model



- Unambiguous specification of program outcomes
 - \star More than just thread interleavings
- Memory Model Desiderata
 - ★ Not too weak (good for programmers)
 - ★ Not too strong (good for hardware)
 - ★ Admits optimisations (good for compilers)
 - ★ Mathematically rigorous (good for verification)

Difficult to get right

- ★ C/C++II memory model is flawed
- \star Java memory model is flawed
- Several papers every year in top PL conferences proposing / fixing models

- Data race
 - ★ Concurrent access to memory location, one of which is a write

- Data race
 - ★ Concurrent access to memory location, one of which is a write
- Sequential consistency (SC)
 - * No intra-thread *reordering*, only inter-thread *interleaving*

- Data race
 - ★ Concurrent access to memory location, one of which is a write
- Sequential consistency (SC)
 - * No intra-thread *reordering*, only inter-thread *interleaving*
- **DRF-SC**: primary tool in concurrent programmers arsenal
 - ★ If a program has no races (under SC semantics), then the program has SC semantics
 - ★ Well-synchronised programs do not have surprising behaviours

- Data race
 - ★ Concurrent access to memory location, one of which is a write
- Sequential consistency (SC)
 - * No intra-thread *reordering*, only inter-thread *interleaving*
- **DRF-SC**: primary tool in concurrent programmers arsenal
 - ★ If a program has no races (under SC semantics), then the program has SC semantics
 - ★ Well-synchronised programs do not have surprising behaviours
- Our observation: DRF-SC is too weak for programmers

• C/C++ (CII) memory model offers DRF-SC, but..

- C/C++ (CII) memory model offers DRF-SC, but..
 - ★ If a program has races (even benign), then the behaviour is undefined!

- C/C++ (CII) memory model offers DRF-SC, but..
 - ★ If a program has races (even benign), then the behaviour is undefined!
 - ★ Most C/C++ programs have races => most C/C++ programs are allowed to crash and burn

- C/C++ (CII) memory model offers DRF-SC, but..
 - ★ If a program has races (even benign), then the behaviour is undefined!
 - ★ Most C/C++ programs have races => most C/C++ programs are allowed to crash and burn
- Races on unrelated locations can affect behaviour

- C/C++ (CII) memory model offers DRF-SC, but..
 - ★ If a program has races (even benign), then the behaviour is undefined!
 - ★ Most C/C++ programs have races => most C/C++ programs are allowed to crash and burn
- Races on unrelated locations can affect behaviour
 - We would like a memory model where data races are bounded in space

Java Memory Model

- Java also offers DRF-SC
 - ★ Unlike C++, type safety necessitates defined behaviour under races

Java Memory Model

- Java also offers DRF-SC
 - ★ Unlike C++, type safety necessitates defined behaviour under races
 - ★ No data races in **space**, but allows races in **time**...

Java Memory Model

- Java also offers DRF-SC
 - ★ Unlike C++, type safety necessitates defined behaviour under races
 - ★ No data races in **space**, but allows races in **time**...

int a; volatile bool flag;
- Java also offers DRF-SC
 - ★ Unlike C++, type safety necessitates defined behaviour under races
 - ★ No data races in **space**, but allows races in **time**...

```
int a;
volatile bool flag;
Thread 1
a = 1;
```

flag = true;

- Java also offers DRF-SC
 - ★ Unlike C++, type safety necessitates defined behaviour under races
 - ★ No data races in **space**, but allows races in **time**...

- Java also offers DRF-SC
 - ★ Unlike C++, type safety necessitates defined behaviour under races
 - ★ No data races in **space**, but allows races in **time**...

- Java also offers DRF-SC
 - ★ Unlike C++, type safety necessitates defined behaviour under races
 - ★ No data races in **space**, but allows races in **time**...



• Future data races can affect the past

• Future data races can affect the past

Class C { int x; }

• Future data races can affect the past

Class C { int x; }

Thread 1 C c = new C(); c.x = 42; r1 = c.x;

Can assert (r1 == 42) fail?

• Future data races can affect the past

```
Class C { int x; }
C g;
```

Thread 1

C c = new C(); c.x = 42; r1 = c.x; g = c;

```
Thread 2
g.x = 7;
```

Can assert (r1 == 42) fail?

• Future data races can affect the past

```
Class C { int x; }
C g;
```

Thread 1

C c = new C(); c.x = 42; g = c; r1 = c.x;

```
Thread 2
g.x = 7;
```

• Future data races can affect the past



assert (r1 == 42) fails

• Future data races can affect the past



assert (r1 == 42) fails

• We would like a memory model that bounds data races in *time*

Language memory models should specify behaviours under data races

- Language memory models should specify behaviours under data races
 - \star Not because they are useful

- Language memory models should specify behaviours under data races
 - \star Not because they are useful
 - ★ But to limit their damage

- Language memory models should specify behaviours under data races
 - \star Not because they are useful
 - ★ But to limit their damage

If I read a variable twice and there are no concurrent writes, then both reads return the same value

OCaml MM: Contributions

- Memory Model Desiderata
 - Not too weak (good for programmers)
 - Not too strong (good for hardware)
 - ★ Admits optimisations (good for compilers)
 - Mathematically rigorous (good for verification)

- OCaml Memory model
 - ★ Local version of DRF-SC key discovery
 - Free on x86, 0.6% overhead on ARM, 2.6% overhead on POWER
 - ★ Allows most common compiler optimisations
 - Simple operational and axiomatic semantics + proved soundness (optimization + to-hardware)

• If there are no data races,

- If there are no data races,
 - ★ on some variables (space)

- If there are no data races,
 - ★ on some variables (space)
 - ★ in some interval (time)

- If there are no data races,
 - ★ on some variables (space)
 - ★ in some interval (time)
 - * then the program has SC behaviour on those variables in that time interval

- If there are no data races,
 - ★ on some variables (space)
 - ★ in some interval (time)
 - * then the program has SC behaviour on those variables in that time interval
- Space = {all variables} && Time = whole execution => DRF-SC

- If there are no data races,
 - ★ on some variables (space)
 - ★ in some interval (time)
 - * then the program has SC behaviour on those variables in that time interval
- Space = {all variables} && Time = whole execution => DRF-SC

Flag is atomic

Thread 1	Thread 2
msg = 1;	b = 1;
b = 0;	if (Flag) {
Flag = 1;	r = msg;
	}

- If there are no data races,
 - ★ on some variables (space)
 - ★ in some interval (time)
 - * then the program has SC behaviour on those variables in that time interval
- Space = {all variables} && Time = whole execution => DRF-SC

Flag is atomic

Thread 1	Thread 2
msg = 1;	b = 1;
b = 0;	if (Flag) {
Flag = 1;	r = msg;
	}

- If there are no data races,
 - ★ on some variables (space)
 - ★ in some interval (time)
 - * then the program has SC behaviour on those variables in that time interval
- Space = {all variables} && Time = whole execution => DRF-SC

Flag is atomic



Due to local DRF, despite the race on **b**, message-passing idiom still works!

- Most programmers can live with local DRF
 - ★ Experts demand more (concurrency libraries, high-performance code, etc.)

- Most programmers can live with local DRF
 - ★ Experts demand more (concurrency libraries, high-performance code, etc.)
- Simple operational semantics that captures all of the allowed behaviours

- Most programmers can live with local DRF
 - * Experts demand more (concurrency libraries, high-performance code, etc.)

a read H(t)

• Simple operational semantics that captures all of the allowed behaviours

$$(SILENT) \qquad \frac{e \stackrel{e}{\hookrightarrow} e'}{\langle S, P[i \mapsto (F, e)] \rangle \to \langle S, P[i \mapsto (F, e')] \rangle} \qquad (READ-NA) \qquad H; F \qquad \frac{a \cdot read H(t)}{if F(a) \leq t, t \in dom(H)} \\ (WRITE-NA) \qquad H; F \qquad \frac{a \cdot write x}{if F(a) < t, t \notin dom(H)} \\ (WRITE-NA) \qquad H; F \qquad \frac{a \cdot write x}{if F(a) < t, t \notin dom(H)} \\ (MEMORY) \qquad \frac{e \stackrel{\ell:\phi}{\longleftrightarrow} e' \qquad S(\ell); F \stackrel{\ell:\phi}{\longrightarrow} C'; F'}{\langle S, P[i \mapsto (F, e)] \rangle \to \langle S[\ell \mapsto C'], P[i \mapsto (F', e')] \rangle} \qquad (READ-AT) \qquad (F_A, x); F \qquad \frac{A \cdot read x}{\longrightarrow} \qquad (F_A, x); F_A \sqcup F \\ (WRITE-AT) \qquad (F_A, y); F \qquad \frac{A \cdot write x}{\longrightarrow} \qquad (F_A \sqcup F, x); F_A \sqcup F \\ (C) Memory operations \end{cases}$$







Non atomic



read(b) -> 3/4/5








Non atomic



Non atomic



Non atomic



Non atomic



Non atomic



Non atomic



Trace
$$\Sigma = M_0 \xrightarrow{T_1} M_1 \xrightarrow{T_2} \dots \xrightarrow{T_n} M_n$$







• Pick a set of L of locations





• Pick a machine state M where there are no ongoing races in L

 $[\]star$ M is said to be L-stable



- Pick a machine state M where there are no ongoing races in L
 - \star M is said to be L-stable
- Local DRF Theorem
 - ★ Starting from an L-stable state M, until the next race on any location in L under SC semantics, the program has SC semantics



• Pick a machine state M where there are no ongoing races in L



under SC semantics, the program has SC semantics

- Local DRF prohibits certain hardware and software optimisations
 - ★ Preserve load-to-store ordering

- Local DRF prohibits certain hardware and software optimisations
 - ★ Preserve load-to-store ordering
- No compiler optimisation that reorders load-to-store ordering is allowed

- Local DRF prohibits certain hardware and software optimisations
 - ★ Preserve load-to-store ordering
- No compiler optimisation that reorders load-to-store ordering is allowed

$$\begin{array}{cccc}
r1 &= a; \\
b &= c; \\
a &= r1; \\
\end{array} \xrightarrow{\text{Redundant store elimination}} & r1 &= a; \\
b &= c; \\
; \\
\end{array}$$

- Local DRF prohibits certain hardware and software optimisations
 - ★ Preserve load-to-store ordering
- No compiler optimisation that reorders load-to-store ordering is allowed

$$\begin{array}{c} r1 = a; \\ b = c; \\ a = r1; \end{array} \xrightarrow{\text{Redundant store elimination}} \begin{array}{c} r1 = a; \\ b = c; \\ ; \end{array}$$

- Local DRF prohibits certain hardware and software optimisations
 - ★ Preserve load-to-store ordering
- No compiler optimisation that reorders load-to-store ordering is allowed

$$\begin{array}{c} r1 = a; \\ b = c; \\ a = r1; \end{array} \xrightarrow{\text{Redundant store elimination}} \begin{array}{c} r1 = a; \\ b = c; \\ ; \end{array}$$

- Local DRF prohibits certain hardware and software optimisations
 - ★ Preserve load-to-store ordering
- No compiler optimisation that reorders load-to-store ordering is allowed

$$r1 = a;$$

$$b = c;$$

$$a = r1;$$

$$Redundatt store elimination$$

$$b = c;$$

$$i$$

- Local DRF prohibits certain hardware and software optimisations
 - ★ Preserve load-to-store ordering
- No compiler optimisation that reorders load-to-store ordering is allowed

$$r1 = a;$$

$$b = c;$$

$$a = r1;$$

$$Redundant store elimination$$

$$b = c;$$

$$i$$

- ARM & POWER do not preserve load-to-store ordering
 - ★ Insert necessary synchronisation between every mutable load and store
 - ★ What is the performance cost?

Performance



Performance



0.6% overhead on AArch64 (ARMv8)

Performance



0.6% overhead on AArch64 (ARMv8) Free on x86, 2.6% on POWER

Summary

- OCaml memory model
 - ★ Balances comprehensibility (Local DRF theorem) and Performance (free on x86, 0.6% on ARMv8, 2.6% on POWER)
 - ★ Allows common compiler optimisations
 - ★ Compilation + Optimisations proved sound

Summary

- OCaml memory model
 - ★ Balances comprehensibility (Local DRF theorem) and Performance (free on x86, 0.6% on ARMv8, 2.6% on POWER)
 - ★ Allows common compiler optimisations
 - ★ Compilation + Optimisations proved sound
- Proposed as the memory model for OCaml
 - ★ Also suitable for other safe languages (Swift, WebAssembly, JavaScript)