

Synchronising C/C++ and POWER


Susmit Sarkar¹ Kayvan Memarian¹ Scott Owens¹ Mark Batty¹
Peter Sewell¹ Luc Maranget² Jade Alglave^{3,4} Derek Williams⁵

¹University of Cambridge ²INRIA ³Oxford University
⁴Queen Mary London ⁵IBM Austin

June 2012

Relaxed Memory Concurrency

- Concurrency on modern hardware/compiler:
Relaxed Memory, not Sequential Consistency (SC)



- Semantics of concurrent programming languages
- ISO C/C++: introduces a new concurrency model
- Hardware: very different concurrency models
 - ▶ Different between x86, Power, ARM
 - ▶ Different from C/C++

Correct implementations of C/C++ on hardware

- Can it be done?
 - ▶ ... on highly relaxed hardware?
- What is involved?
 - ▶ Mapping new constructs to assembly
 - ▶ Optimizations: which ones legal?

Correct implementations of C/C++ on hardware

- Can it be done?
 - ▶ ... on highly relaxed hardware? e.g. Power
- What is involved?
 - ▶ Mapping new constructs to assembly
 - ▶ Optimizations: which ones legal?

Implementing C/C++11 on POWER: Pointwise Mapping

C/C++11 Operation	POWER Implementation
Store (non-atomic) Load (non-atomic)	st ld

(From Paul McKenney and Raul Silveira)

Implementing C/C++11 on POWER: Pointwise Mapping

C/C++11 Operation	POWER Implementation
Store (non-atomic)	st
Load (non-atomic)	ld
Store relaxed	st
Store release	lwsync; st
Store seq-cst	lwsync; st
Load relaxed	ld
Load consume	ld (and preserve dependency)
Load acquire	ld; cmp; bc; isync
Load seq-cst	hwsync; ld; cmp; bc; isync

(From Paul McKenney and Raul Silveira)

Implementing C/C++11 on POWER: Pointwise Mapping

C/C++11 Operation	POWER Implementation
Store (non-atomic)	st
Load (non-atomic)	ld
Store relaxed	st
Store release	lwsync; st
Store seq-cst	lwsync; st
Load relaxed	ld
Load consume	ld (and preserve dependency)
Load acquire	ld; cmp; bc; isync
Load seq-cst	hwsync; ld; cmp; bc; isync
Fence acquire	lwsync
Fence release	lwsync
Fence seq-cst	hwsync

(From Paul McKenney and Raul Silveira)

Implementing C/C++11 on POWER: Pointwise Mapping

C/C++11 Operation	POWER Implementation
Store (non-atomic) Load (non-atomic)	st ld
Store relaxed Store release Store seq-cst	st lwsync; st lwsync; st
Load relaxed Load consume Load acquire Load seq-cst	ld ld (and preserve dependency) ld; cmp; bc; isync hwsync; ld; cmp; bc; isync
Fence acquire Fence release Fence seq-cst	lwsync lwsync hwsync
CAS relaxed	_loop: lwarx; cmp; bc _exit; stwcx.; bc _loop; _exit:
CAS seq-cst	hwsync; _loop: lwarx; cmp; bc _exit; stwcx.; bc _loop; isync; _exit:
...	...

(From Paul McKenney and Raul Silveira)

Implementing C/C++11 on POWER: Pointwise Mapping

C/C++11 Operation	POWER Implementation
Store (non-atomic)	st
Load (non-atomic)	ld
Store relaxed	st
Store release	lwsync; st
Store seq-cst	lwsync; st
Load relaxed	ld
Load consume	ld (and preserve dependency)
Load acquire	
Load seq-cst	
Fence acquire	lwsync
Fence release	lwsync
Fence seq-cst	hwsync
CAS relaxed	<code>_loop: lwarx; cmp; bc _exit; stwcx.; bc _loop; _exit:</code>
CAS seq-cst	<code>hwsync; _loop: lwarx; cmp; bc _exit; stwcx.; bc _loop; isync; _exit:</code>
...	...

Is that mapping correct?

(From Paul McKenney and Raul Silveira)

Implementing C/C++11 on POWER: Pointwise Mapping

C/C++11 Operation	POWER Implementation
Store (non-atomic) Load (non-atomic)	st ld
Store relaxed Store release Store seq-cst	st lwsync; st lwsync; hwsync; st
Load relaxed Load consume Load acquire Load seq-cst	ld ld (and preserve dependency) ld; cmp; bc; isync hwsync; ld; cmp; bc; isync
Fence acquire Fence release Fence seq-cst	lwsync lwsync hwsync
CAS relaxed	Answer: No!
CAS seq-cst	hwsync; _loop: lwarx; cmp; bc _exit; stwcx.; bc _loop; isync; _exit;
...	...

(From Paul McKenney and Raul Silveira)

Implementing C/C++11 on POWER: Pointwise Mapping

C/C++11 Operation	POWER Implementation
Store (non-atomic)	st
Load (non-atomic)	ld
Store relaxed	st
Store release	lwsync; st
Store seq-cst	hwsync; st
Load relaxed	ld
Load consume	ld (and preserve dependency)
Load acquire	
Load seq-cst	
Fence acquire	lwsync
Fence release	lwsync
Fence seq-cst	hwsync
CAS relaxed	
CAS seq-cst	hwsync; _loop: lwarx; cmp; bc _exit; stwcx.; bc _loop; isync; _exit;
...	...

Is that mapping correct?

Answer: Yes!

(From Paul McKenney and Raul Silveira)

Implementing C/C++11 on POWER: Pointwise Mapping

C/C++11 Operation	POWER Implementation
Store (non-atomic)	st
Load (non-atomic)	ld
Store relaxed	st
Store release	lwsync; st
Store seq-cst	hwsync; st
Load relaxed	ld
Load consume	ld (and preserve dependency)
Fence acquire	lwsync
Fence release	lwsync
Fence seq-cst	hwsync
CAS relaxed	Answer: No!
CAS seq-cst	hwsync; _loop: lwarx; cmp; bc _exit; stwcx.; bc _loop; isync; _exit;
...	...

(From Paul McKenney and Raul Silveira)

Implementing C/C++11 on POWER: Pointwise Mapping

C/C++11 Operation	POWER Implementation
Store (non-atomic) Load (non-atomic)	st ld
Store relaxed Store release Store seq-cst	st lwsync; st hwsync; st
Load relaxed Load consume Load acquire Load seq-cst	ld ld (and preserve dependency) ld; cmp; bc; isync hwsync; ld; cmp; bc; isync
Fence acquire Fence release Fence seq-cst	lwsync lwsync hwsync
CAS relaxed	_loop: lwarx; cmp; bc _exit; stwcx.; bc _loop; _exit:
CAS seq-cst	hwsync; _loop: lwarx; cmp; bc _exit; stwcx.; bc _loop; isync; _exit:
...	...

Alternative

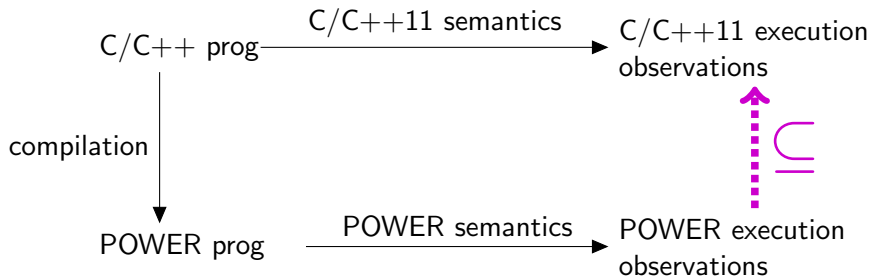
hwsync; st; hwsync;

ld; hwsync

All compilers must agree for separate compilation

Implementing C/C++11 on POWER *correctly*

Theorem: For **any** sane, non-optimising compiler following the mapping:



- Showed previous mapping incorrect
- Easily adapt proof for an alternative mapping

Reasoning about industrial-strength concurrency

Enables:

- Confidence in C/C++ and Power concurrency models
- Confidence in compiler implementations [gcc]
- Reasoning about C/C++ and Power
- (Path to) Reasoning about ARM ??

Context of This Paper

Before [POPL'12]: just loads and stores

- Power concurrency model (of loads and stores) [PLDI'11]
- C++11 concurrency model [POPL'11]
- Proof:
 - ▶ some concepts correspond (e.g. coherence \rightarrow modification order)
 - ▶ others depend on key properties of abstract machine

This paper: also with synchronisation constructs

- Power: load-reserve and store-conditional
- C++11: locks, read-modify-writes, fences
- Proof:
 - ▶ extends smoothly (new cases to be checked)
 - ▶ points out interesting features of the models

Outline

- 1 Introduction
- 2 Relaxed Memory Behaviour (examples)
- 3 Reasoning about Synchronising Operations
- 4 Proof Outline; and What We Learned

Example: Message Passing

Initially: <code>d = 0; f = 0;</code>	
Thread 0	Thread 1
<code>d = 1;</code> <code>f = 1;</code>	<code>while (f == 0)</code> <code> {};</code> <code>r = d;</code>
Finally: <code>r = 0 ??</code>	

- Forbidden on SC

Example: Message Passing (racy)

Initially: <code>d = 0; f = 0;</code>	
Thread 0	Thread 1
<code>d = 1;</code> <code>f = 1;</code>	<code>while (f == 0)</code> <code> {};</code> <code>r = d;</code>
Finally: <code>r = 0 ??</code>	

- Forbidden on SC
- In C/C++11, this has **undefined** semantics
- Data race on `d` and `f` variables

Example (contd.): mark atomics

Mark atomic variables (accesses have memory order parameter)

Initially: <code>d = 0; f = 0;</code>	
Thread 0	Thread 1
<code>d.store(1,rlx);</code> <code>f.store(1,rlx);</code>	<code>while (f.load(rlx) == 0)</code> <code>}</code> ; <code>r = d.load(rlx);</code>
Finally: <code>r = 0 ??</code>	

- (Forbidden on SC)

Example (contd.): mark atomics

Mark atomic variables (accesses have memory order parameter)

Initially:	<code>d = 0; f = 0;</code>	
Thread 0	Thread 1	
<code>d.store(1,rlx);</code> <code>f.store(1,rlx);</code>	<code>while (f.load(rlx) == 0)</code> <code>{</code> <code>r = d.load(rlx);</code> <code>}</code>	
Finally: <code>r = 0</code> ??		

- (Forbidden on SC)
- Defined, and possible, in C/C++11
- Allows for hardware (and compiler) optimisations

Example (contd.): release-acquire synchronization

Mark release stores and acquire loads

Initially:	<code>d = 0; f = 0;</code>	
Thread 0	Thread 1	
<code>d.store(1,rlx);</code> <code>f.store(1,rel);</code>	<code>while (f.load(acq) == 0)</code> <code>{</code> <code>r = d.load(rlx);</code> <code>}</code>	
Finally: <code>r = 0 ??</code>		

- (Forbidden on SC)
- Forbidden in C/C++11 due to release-acquire synchronization
- Implementation must ensure result not observed

Implementation of acquire/release on POWER

Initially: $d = 0; f = 0;$	
Thread 0	Thread 1
<pre>st d 1; lwsync; st f 1;</pre>	<pre>loop: ld f rtmp; cmp rtmp 0; beq loop; isync; ld d r;</pre>
Finally: $r = 0$??	

- Forbidden (and not observed) on POWER7, and ARM
- `lwsync` prevents write reordering
- `control dependency` with `isync` prevents read speculation

Outline

- 1 Introduction
- 2 Relaxed Memory Behaviour (examples)
- 3 Reasoning about Synchronising Operations
- 4 Proof Outline; and What We Learned

What about Synchronising (Atomic) Operations?

- Synchronization operations, e.g. “atomic add”, “CAS”, .. .
- RISC-friendly alternative: Load-reserve/Store-conditional

What about Synchronising (Atomic) Operations?

- Synchronization operations, e.g. “atomic add”, “CAS”, ...
- RISC-friendly alternative: Load-reserve/Store-conditional
- Can be used to implement CAS, spinlocks, ...
- Universal (like CAS) [Herlihy'93], but no ABA problem

Atomic Addition
<pre>loop: lwarx r, d; add r,v,r; stwcx r, d; bne loop;</pre>

- Informally, stwcx succeeds only if no other write to the same address since last lwarx

What *is* no write since ... ?

- In machine time?
 - ▶ Neither necessary, nor sufficient

What *is* no write since ... ?

- In machine time?
 - ▶ Neither necessary, nor sufficient
- Microarchitecturally (simplified): if cache-line ownership not lost since last `lwarx`

Modeling “not lost since”

- Abstractly: ownership chain modeled by building up coherence order
- Coherence: order relating stores to the same location (eventually linear)
- A `stwx` succeeds only if it is (becomes) coherence-next-to the write read from by `lwarx`
- ... and no other write can later come in between

Modeling “not lost since”

- Abstractly: ownership chain modeled by building up coherence order
- Coherence: order relating stores to the same location (eventually linear)
- A `stwx` succeeds only if it is (becomes) coherence-next-to the write read from by `lwarx`
- ... and no other write can later come in between
- Isolate key concept: **write reaching coherence point** —
 - ▶ coherence is linear below this write, and no new edges will be added below

Load-reserve/store-conditional and ordering

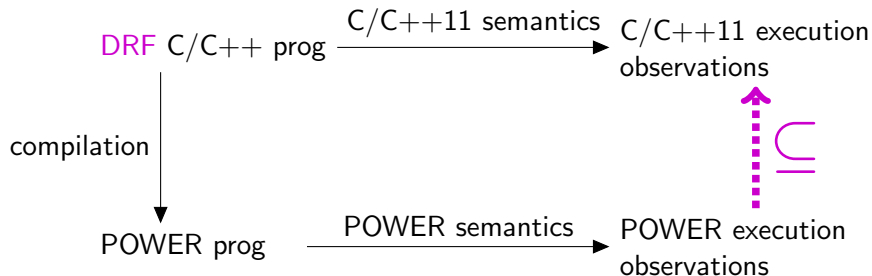
- Same-thread load-reserve/store-conditionals ordered by program order
 - If **all** memory accesses are atomic sequences
 - Then: only SC behaviour
- **But . . .** normal loads/stores (to different addresses) not ordered
 - Confusion here led to Linux bug
 - . . . bad barrier placement in atomic-add-return

Outline

- 1 Introduction
- 2 Relaxed Memory Behaviour (examples)
- 3 Reasoning about Synchronising Operations
- 4 Proof Outline; and What We Learned

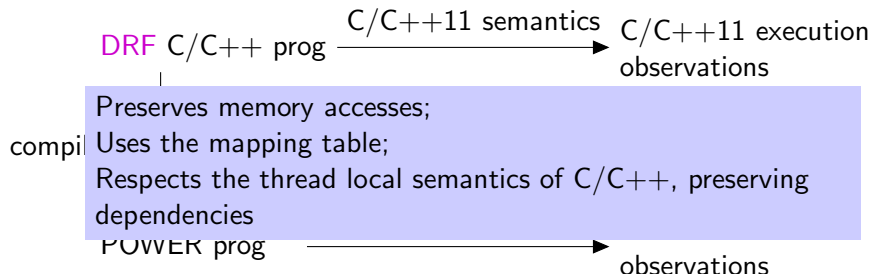
Proof outline

Theorem: For any sane, non-optimising compiler following the mapping:



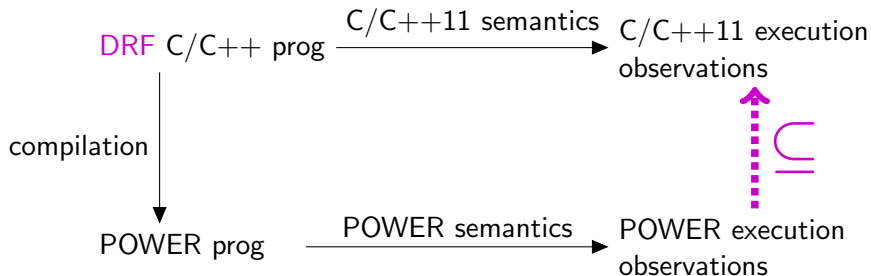
Proof outline

Theorem: For any **sane, non-optimising compiler** following the mapping:



Proof outline

Theorem: For any sane, non-optimising compiler following the mapping:



From POWER trace, build key relations (happens-before, SC order)

Required properties from abs. machine properties

If trace looks like it produces data race, build the C/C++ data race

Also in the paper

- A formal model of load-reserve/store-conditional (in Lem)
- An executable model with exploration tool (ppcmem)
- Simplifications to the C/C++11 lock model
- Models “tight” against each other: relaxing the Power model would make C/C++11 unimplementable

Reasoning about industrial-strength concurrency

- Correct compilation of C/C++ concurrency primitives on Power
 - Formal relaxed-memory semantics of load-reserve/store-conditional
 - Allow proof of SC via atomic RMW sequences
 - Technical simplifications to the C/C++ lock model
- Confidence in both models
- Compiler implementation relevance
- Reasoning about machine code at C/C++ level

Thank You!

More details at:

<http://www.cl.cam.ac.uk/~pes20/cppppc>

Store-conditional speculation?

- Power allows stores to **forward** value to same thread **speculatively**
- Can (and should) stwcx be allowed to be speculated (even before the lwarx) ?

Initially: <code>d = 0 f = 0;</code>	
Thread 0	Thread 1
<code>d = 1; # d.store(1,rlx)</code> <code>lwsync; # f.store(1,rel)</code> <code>f = 1;</code>	<code>loop: lwarx f, r1; # CAS (f,1,2)</code> <code>cmp r1 1; bne exit;</code> <code>stwcx f 2; bne loop;exit:</code>
	<code>ld r1 f; # r1 = f.load(con)</code> <code>xor r2, r1,r1; # r2 = r1 \oplus r1</code> <code>ld [d + r2] r; # r = d[r2]</code>
Finally: <code>r = 0 ??</code>	

Store-conditional speculation?

- Can (and should) stwcx be allowed to be speculated (even before the lwarx) ?

Initially: <code>d = 0 f = 0;</code>	
Thread 0	Thread 1
<code>d = 1; # d.store(1,rlx)</code> <code>lwsync; # f.store(1,rel)</code> <code>f = 1;</code>	<code>loop: lwarx f, r1; # CAS (f,1,2)</code> <code>cmp r1 1; bne exit;</code> <code>stwcx f 2; bne loop;exit:</code>
	<code>ld r1 f; # r1 = f.load(con)</code> <code>xor r2, r1,r1; # r2 = r1 \oplus r1</code> <code>ld [d + r2] r; # r = d[r2]</code>
Finally: <code>r = 0 ??</code>	

- C/C++11 mapping would break (and no good way of fixing)
- Fortunately, current hardware does not do this
- ... and now we know why future hardware should not

60 second pitch

Hi, I am Susmit Sarkar, and I am going to be speaking about shared-memory concurrency not as we would like it to be, but as it actually is in the real world, on mainstream hardware such as PowerPC or ARM and on software such as the new C and C++ concurrency model. These two models are quite strange, and quite different from each other so it is a real question whether you can even compile from one to the other. Yes you can, and we prove this. This explains how these very different models really work. Come to Room B, just after lunch