

Sequential Consistency in C11

October 24, 2022

Duc Than Nguyen

University of Illinois at Chicago

Outline

Introduction

Background

Common OT are invalid

Overhauling SC atomics in C11

Repairing SC in C11

Selected critique

Outline

Introduction

Background

Common OT are invalid

Overhauling SC atomics in C11

Repairing SC in C11

Selected critique

Introduction

- ❁ Programmer's goal—fully understand what they've written
- ❁ Compiler's responsibility—make code as optimized as possible
- ✘ Simple compiler optimizations can generate unexpected behaviors in concurrent setting
- ✔ Memory model—core of concurrent semantics of shared memory
→ alleviate tension between goals of programmer & compiler
- 👉 C11 memory model defines semantics in C programming
- 😊 Batty et al. [2011]¹ formalized concurrency model in C++ standard
- 😞 Several issues discovered with semantics of SC C11 accesses
- 😊 Model has evolved with fixes and revisions (e.g., Vafeiadis et al. [2015], Batty et al. [2016], Lahav et al. [2017])

¹Mathematizing C++ Concurrency. In POPL 2011.

Papers

Authors	Paper	Conference
Vafeiadis et al.	Common compiler optimisations are invalid in the C11 memory model and what we can do about it	POPL 2015
Batty et al.	Overhauling SC atomics in C11 and OpenCL	POPL 2016
Lahav et al.	Repairing sequential consistency in C/C++11	PLDI 2017

Outline

Introduction

Background

Common OT are invalid

Overhauling SC atomics in C11

Repairing SC in C11

Selected critique

Sequential consistency

... 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]²)

$x := 0; \quad y := 0;$

$$\begin{array}{l} x := 1 \\ a := *y; \text{//load } 0 \end{array} \parallel \begin{array}{l} y := 1 \\ b := *x; \text{//load } 0 \end{array}$$

- * Thread-local variables $a = b = 0$?
- * 3 possible outcomes: (1) $a = b = 0$, (2) $a = 0, b = 1$, (3) $a = 1, b = 0$
- * $a = b = 0$ cannot happen (no interleaving yields that result)
- * Modern systems allow its behavior

²How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs. IEEE Trans. Computers 28, 9 (1979)

C11 memory model

$$na \sqsubset rlx \sqsubset \{acq, rel\} \sqsubset sc$$

The full C/C++11 is more general


sc— Sequentially consistent accesses

rel— Release, and acq— Acquire accesses

rlx— Relaxed accesses

na— Non-atomic accesses (a.k.a. normal data accesses)

Declarative/Axiomatic memory model & Semantics



$$x :=_{na} O; \quad y :=_{na} O;$$

$$x :=_{sc} 1; \quad \left\| \begin{array}{l} a :=^{*acq} x \quad //1 \\ c :=^{*sc} y; \quad //0 \end{array} \right\| \left\| \begin{array}{l} b :=^{*acq} y; \quad //1 \\ d :=^{*acq} x; \quad //0 \end{array} \right\| \quad y :=_{sc} 1;$$

events — nodes & relations — edges


$W(x, o)$ event — x set to o , and $R(x, o)$ event — x returns o

sb— Sequenced-before : order of intra-thread memory accesses

rf— Reads-from: connect two write and read accesses

mo— Modification order: total order on all writes to same location

Declarative/Axiomatic memory model & Semantics



$$x :=_{na} 0; \quad y :=_{na} 0;$$

$$x :=_{sc} 1; \quad \left\| \begin{array}{l} a :=^{*acq} x \quad //1 \\ c :=^{*sc} y; \quad //0 \end{array} \right\| \left\| \begin{array}{l} b :=^{*acq} y; \quad //1 \\ d :=^{*acq} x; \quad //0 \end{array} \right\| \quad y :=_{sc} 1;$$

Additional relational types

sw— Synchronized-with: an acq (sc) read reads from a rel (sc) write

hb— Happens-before: $hb = (sb \cup sw)^+$ (transitive closure)

rb— Reads-before: $rb = rf^{-1}$; **mo** $\setminus [E]$ — read access reads from write that is **mo**— before, exclude update event from itself

Semantics

Rules of total order of sequentially consistent operations

1. Total order on SC operations: any two SC operations must be ordered w.r.t. each other
2. Be consistent with **hb** and **mo** restricted to SC atomics
3. SC reads (SCReads axiom) must either
 - read from most recent SC write before them in SC order, or
 - read from non-SC write that does **not** happen-before most recent SC write to that location

Outline

Introduction

Background

Common OT are invalid

Overhauling SC atomics in C11

Repairing SC in C11

Selected critique

Optimization transformations (OT)

- * Compiler converts program to architecture, preserving its semantics
- * Compiler must be aware of memory consistency models of programming language & architecture
- * C11 compilers (GCC, LLVM) rarely perform C code compilation in a single step
- * Front-end translate C11 to intermediate representations (IR), compiler executes several OT on IR then generates target code
- * Source & target code maintains same consistency model
- * OT: reordering independent memory & removing redundant memory accesses
- * $x :=_{\text{rel}} l; y :=_{\text{na}} l; \rightsquigarrow y :=_{\text{na}} l; x :=_{\text{rel}} l;$

Monotonicity

“Adding synchronisation should not introduce new behaviors”

- * Adding a memory fence
- * *Strengthening the access mode of an operation*
- * Reducing parallelism $\mathcal{C}_1 || \mathcal{C}_2 \rightsquigarrow \mathcal{C}_1; \mathcal{C}_2$
- * Roach motel reorderings

Strengthening access mode of an operation

(Recall) SCReads axiom – SC reads must either (1) read from most recent SC write before them in SC order, or (2) read from non-SC write that does **not** happen-before most recent SC write to that location

$x :=_{\text{rlx}} 1;$ $x :=_{\text{sc}} 2;$ $y :=_{\text{sc}} 1;$	$x :=_{\text{rlx}} 3;$ $y :=_{\text{sc}} 2;$	$y :=_{\text{sc}} 3;$ $r :=^{*\text{sc}} x; //1$	$s_1 :=^{*\text{rlx}} x; //1$ $s_2 :=^{*\text{rlx}} x; //2$ $s_3 :=^{*\text{rlx}} x; //3$ $t_1 :=^{*\text{rlx}} y; //1$ $t_2 :=^{*\text{rlx}} y; //2$ $t_3 :=^{*\text{rlx}} y; //3$
--	---	---	--

? Allowed – $r = s_1 = t_1 = 1 \wedge s_2 = t_2 = 2 \wedge s_3 = t_3 = 3$

☞ If it holds, $x :=_{\text{sc}} 2$ – immediate SC prior write w.r.t. $r :=^{*\text{sc}} x$

☞ Reading 1 of $r :=^{*\text{sc}} x$ from $x :=_{\text{rlx}} 1$ is **not** valid

($x :=_{\text{rlx}} 1$ happens before $x :=_{\text{sc}} 2$ via sb)

Strengthening access mode of an operation

(Recall) SCReads axiom – SC reads must either (1) read from most recent SC write before them in SC order, or (2) read from non-SC write that does **not** happen-before most recent SC write to that location

$x :=_{\text{rlx}} 1;$ $x :=_{\text{sc}} 2;$ $y :=_{\text{sc}} 1;$	$x :=_{\text{sc}} 3;$ $y :=_{\text{sc}} 2;$	$y :=_{\text{sc}} 3;$ $r :=^{*\text{sc}} x; //1$	$s_1 :=^{*\text{rlx}} x; //1$ $s_2 :=^{*\text{rlx}} x; //2$ $s_3 :=^{*\text{rlx}} x; //3$ $t_1 :=^{*\text{rlx}} y; //1$ $t_2 :=^{*\text{rlx}} y; //2$ $t_3 :=^{*\text{rlx}} y; //3$
--	--	---	--

? Strengthening $x :=_{\text{rlx}} 3$ into $x :=_{\text{sc}} 3$

☞ Establish SC order from $y :=_{\text{sc}} 1$ to $x :=_{\text{sc}} 3$

☞ $x :=_{\text{sc}} 3$ is immediately prior in SC order of $r :=^{*\text{sc}} x$

☞ Reading 1 of $r :=^{*\text{sc}} x$ is valid (since $x :=_{\text{rlx}} 1$ does not hb $x :=_{\text{sc}} 3$)

✗ New behavior introduced

Correcting the SCReads Axiom

$$\forall a, b. \text{rf}(b) = a \wedge \text{isSC}(b) \Rightarrow \text{imm}(\text{scr}, a, b) \vee \neg \text{isSC}(a) \wedge \nexists x. \text{hb}(a, x) \wedge \text{imm}(\text{scr}, x, b) \quad (1)$$

- ❶ $\text{isSC}(a) = \text{mode}(a) = \text{sc}$
- ❷ $\text{imm}(R, a, b) = R(a, b) \wedge \nexists c. R(a, c) \wedge R(c, b)$
- ❸ $\text{scr}(a, b) = \text{sc}(a, b) \wedge \text{iswrite}_{\text{loc}(b)}(a)$
- ❹ $\text{iswrite}_l(a) = \exists v. (\exists X, v'. \text{lab}(a) \in \{W^X(l, v), U^X(l, v', v)\})$

- ? Problem: disallows a non-SC write that **hb** another SC-write that is immediately prior in SC order to same location
- ✓ Fix: change $\text{imm}(\text{scr}, x, b)$ into $\text{scr}(x, b)$
constrains all SC prior same-location writes instead of only immediate prior write on same location
- ✓ No **hb** between write $\text{rf}(b)$ and any same location write sc -prior to read $\rightarrow r :=^{*\text{sc}} x$ reading from $x :=_{\text{rlx}} 1$ not valid

Outline

Introduction

Background

Common OT are invalid

Overhauling SC atomics in C11

Repairing SC in C11

Selected critique

Derived sets and relations

They are formalized as C11 axioms by referencing the C11 standard

- | | | |
|---|---|---|
| ❶ | $\text{irr}(S; r_1)$ | where $r_1 = \text{hb}$ |
| ❷ | $\text{irr}(S; r_2)$ | where $r_2 = ([F^{\text{SC}}]; \text{sb})^?; \text{mo}; (\text{sb}; [F^{\text{SC}}])^?$ |
| ❸ | $\text{irr}(S; r_3)$ | where $r_3 = \text{rf}^{-1}; [E^{\text{SC}}]; \text{mo}$ |
| ❹ | $\text{irr}((S \setminus (\text{mo}; S)); r_4)$ | where $r_4 = \text{rf}^{-1}; \text{hbl}; [E^{\text{W}}]$,
and $\text{hbl} - \text{hb}$ to events on same location |
| ❺ | $\text{irr}(S; r_5)$ | where $r_5 = ([F^{\text{SC}}]; \text{sb}); \text{rb}$ |
| ❻ | $\text{irr}(S; r_6)$ | where $r_6 = \text{rb}; (\text{sb}; [F^{\text{SC}}])$ |
| ❼ | $\text{irr}(S; r_7)$ | where $r_7 = ([F^{\text{SC}}]; \text{sb}); \text{rb}; (\text{sb}; [F^{\text{SC}}])$ |

S— Sequential consistency order

Axioms ❸ ❹ refer to as SCReads axiom

Axioms ❺ ❻ ❼ govern SC fences

* All axioms express as $\text{irr}(S; r)$ except ❹

Construction partial order on SC operations

④ $\text{irr}((S \setminus (\text{mo}; S)); r_4)$, where $r_4 = \text{rf}^{-1}; \text{hbl}; [E^W]$, and hbl is hb to events on same location

 Replace $S \setminus (\text{mo}; S)$ with S , obtain new ④✓ $\text{irr}(S; r_4)$

④✓ coincides with revised model offered by Vafeiadis et al. [2015] in “Correcting the SCReads Axiom”

④✓ disallows an SC read to see any write that happens before any SC write in S

$\text{acyclic}([E^{\text{SC}}]; (r_1 \cup r_2 \cup r_3 \cup r_4 \cup r_5 \cup r_6 \cup r_7); [E^{\text{SC}}]) \quad (S_{\text{partial}})$

■ Proved $S_{\text{partial}} = \exists S. \text{wfs} \wedge \textcircled{1} \wedge \textcircled{2} \wedge \textcircled{3} \wedge \textcircled{4}\checkmark \wedge \textcircled{5} \wedge \textcircled{6} \wedge \textcircled{7}$

✿ S_{partial} does not require S anymore

✈ This axiom is faster to simulate!

✓ Existing compilation schemes (x86 and Power) remain valid

Stronger & simpler SC axiom

? Question: *Strengthen SC semantics without requiring changes to compilation schemes of C11 target architectures (x86 and Power)?*

- ⑤ ⑥ and ⑦ contain **rb** at begin or finish at a fence
- ④✓ contains **rb** since $r_4 = rf^{-1}; hbl; [E^W]$
 where **hbl** is **hb** to events on same location
 $\rightarrow r_4 = rf^{-1}; mo$ and $rb = rf^{-1}; mo$
- ③ contains $r_3 = rf^{-1}; [E^{SC}]; mo$
- ✂ Remove $[E^{SC}]$, we have $r_3 = rb$ and new ③✓ $irr(S; rb)$
- ✿ Enhance S_{partial} to have
 $S_{\text{simp}} = \text{acyclic}([E^{SC}]; (r_1 \cup r_2 \cup rb \cup r_4 \cup r_5 \cup r_6 \cup r_7); [E^{SC}])$

$$\text{acyclic}([E^{SC}]; ((([F^{SC}]; sb)^?); (hb \cup rb \cup mo); (sb; [F^{SC}])^?); [E^{SC}]) \quad (S_{\text{simp}})$$

Outline

Introduction

Background

Common OT are invalid

Overhauling SC atomics in C11

Repairing SC in C11

Selected critique

Compilation to Power is Broken

❶ $[E^{SC}]; hb; [E^{SC}] \subseteq S$

❷ $[E^{SC}]; mo; [E^{SC}] \subseteq S$

❸ $[E^{SC}]; rb; [E^{SC}] \subseteq S$


❹ ❺ ❻ ❼ indicate that S is required to comply with a few more conditions about SC fences.

Same location: Power & ARM ensure compilation preserves condition

❷ & ❸ force ordering between accesses to *same location*

❶ forces between accesses of *different locations*, requiring insertion of fence instructions

Compilation to Power is Broken




$$x :=_{sc} I; \left\| \begin{array}{l} a :=^{*acq} x \quad //1 \\ c :=^{*sc} y; \quad //0 \end{array} \right\| \left\| \begin{array}{l} b :=^{*acq} y; \quad //1 \\ d :=^{*acq} x; \quad //0 \end{array} \right\| y :=_{sc} I;$$

✱ Independent Reads Independent Writes (IRIW)

✘ Constraints by [Batty et al. 2016] – too strong to preserve compiler correctness from C11 to Power and ARMv7

Compilation to Power is Broken



- * $S(p, k)$ via sc ; $S(k, m)$ via hb (where $hb = (sw \cup sb)^+$);
then $sc(p, m)$ via transitivity.
- * $S(m, p)$ via rb (where $rb = rf^{-1}; mo$) \rightarrow S contains cycles
- ✗ banned by C11
- ✓ permitted by compilation into Power

Fixing the model


- ❶ $[E^{SC}]; hb; [E^{SC}] \subseteq S$ — same or different location
- * same location: hardware maintain the order
- * different location: a sync fence between SC accesses
→ start and ends hb with sb , i.e., $sb; hb; sb$
- * Fix: replace hb with $(sb \cup sb; hb; sb \cup hb)|_{=loc}$

New model ❶✓

$acyclic([E^{SC}]; (sb \cup sb; hb; sb \cup hb)|_{=loc} \cup mo \cup rb); [E^{SC}]$ (❶✓)

❶✓ forbids elimination of SC write immediately followed by SC write to same location, and SC read immediately followed by SC read from same location

Enabling Elimination of SC Accesses



$$a :=^{\text{acq}} x; //2 \quad \parallel \quad x :=_{\text{sc}} 1; \quad \parallel \quad y :=_{\text{sc}} 1;$$

$$b :=^{\text{sc}} y; //2 \quad \parallel \quad x :=_{\text{sc}} 2; \quad \parallel \quad c :=^{\text{sc}} x; //0$$

- ❶✓ forbids elimination of SC write immediately followed by SC write to same location, and SC read immediately followed by SC read from same location
- * Eliminate $x :=_{\text{sc}} 1;$ (event $m : W^{\text{sc}}(x, 1)$)
- * Create cycle $n \rightarrow k \rightarrow l \rightarrow o \rightarrow p \rightarrow n$
- ✗ Violate condition ❶✓ must be acyclic

Fixing the model

acyclic($[E^{SC}]; (sb \cup sb; hb; sb \cup hb|_{=loc} \cup mo \cup rb); [E^{SC}]$) (1✓)

- ✿ Fix: weaken condition by replacing
 $sb; hb; sb$ with $sb|_{\neq loc}; hb; sb|_{\neq loc}$
 where $sb|_{\neq loc}$ – sb edges to different location

New condition (named SC-before) requires acyclicity of $[E^{SC}]; scb; [E^{SC}]$

$$scb = sb \cup sb|_{\neq loc}; hb; sb|_{\neq loc} \cup hb|_{=loc} \cup mo \cup rb$$

SC Fences are Too Weak

? Question: *Shall adding SC fences between every pair of shared memory restore interleaving behavior?*

✘ Original C11 & Batty et al. [2016] do **not** hold

✓ hold for Power & ARM

* Adapted from Batty et al. [2016]

$\text{acyclic}([E^{\text{sc}}] \cup [F^{\text{sc}}]; \text{sb}^?); (\text{hb} \cup \text{mo} \cup \text{rb}); ([E^{\text{sc}}] \cup \text{sb}^?; [F^{\text{sc}}])$


* Replace $\text{hb} \cup \text{mo} \cup \text{rb}$ with scb

where $\text{scb} = \text{sb} \cup \text{sb}|_{\neq \text{loc}}; \text{hb}; \text{sb}|_{\neq \text{loc}} \cup \text{hb}|_{= \text{loc}} \cup \text{mo} \cup \text{rb}$

* Require acyclicity of psc_1

$$\text{psc}_1 = ([E^{\text{sc}}] \cup [F^{\text{sc}}]; \text{sb}^?); \text{scb}; ([E^{\text{sc}}] \cup \text{sb}^?; [F^{\text{sc}}])$$

SC Fences are Too Weak



$$x :=_{rlx} 1; \left\| \begin{array}{l} a :=^{*rlx} x; \quad //1 \\ fence_{sc} \\ b :=^{*rlx} y; \quad //0 \end{array} \right\| \left\| \begin{array}{l} y :=_{rlx} 1; \\ fence_{sc} \\ c :=^{*rlx} x; \quad //0 \end{array} \right.$$

👉 Path f_1 to f_2 via $sb \rightarrow rb \rightarrow sb$

👉 Path f_2 to f_1 via $sb \rightarrow rb \rightarrow rf \rightarrow sb$

Applying $psc_1 = ([E^{sc}] \cup [F^{sc}]; sb^?); scb; ([E^{sc}] \cup sb^?; [F^{sc}])$ is acyclic

✘ Path f_2 to f_1 contributes neither Batty et al. [2016] nor psc_1

Fixing model

- ✿ Extended-coherence-order relation $\text{eco} = (\text{rf} \cup \text{mo} \cup \text{rb})^+$
- ✿ Disallow the weak behavior, it is impossible for either psc_1 or $[\text{F}^{\text{sc}}]; \text{sb}; \text{eco}; \text{sb}; [\text{F}^{\text{sc}}]$ containing cycles

$\text{acyclic}(\text{psc}_1 \cup [\text{F}^{\text{sc}}]; \text{sb}; \text{eco}; \text{sb}; [\text{F}^{\text{sc}}])$

Outline

Introduction

Background

Common OT are invalid

Overhauling SC atomics in C11

Repairing SC in C11

Selected critique

Selected critique

- ❖ (Vafeiadis et al. 2015) – disallows reordering of non-atomic load and store operations, when offered solution to out-of-thin-air reads that restricts ($hb \cup rf$) cycle (source or destination of rf edge is non-atomic)
- ❖ (Batty et al. 2016) – strong to preserve compiler correctness from C11 to Power and ARMv7
- ❖ (Lahav et al. 2017) – RC11 directly map high-level primitive to sequence of machine instructions without code optimizations, then challenging to apply methods to verify optimization passes


Thank you for listening!



Outline

backup slides

Strengthening access mode of an operation


 $x :=_{rlx} 1;$
 $x :=_{sc} 2;$
 $y :=_{sc} 1;$
 $x :=_{rlx} 3;$
 $y :=_{sc} 2;$
 $y :=_{sc} 3;$
 $r :=_{sc} x; //1$
 $s_1 := {}^*rlx x; //1$
 $s_2 := {}^*rlx x; //2$
 $s_3 := {}^*rlx x; //3$
 $t_1 := {}^*rlx y; //1$
 $t_2 := {}^*rlx y; //2$
 $t_3 := {}^*rlx y; //3$