

Correctness of compilation under weak memory models

Ori Lahav

Viktor Vafeiadis

29 August 2017

What does “correct compilation” mean?

Compilation correctness

$$\llbracket \text{compile}(P) \rrbracket_{\text{target memory model}} \subseteq \llbracket P \rrbracket_{\text{source memory model}}$$

i.e.,

every outcome that is allowed $\text{compile}(P)$ under the target memory model is also an allowed outcome for P under the source memory model

- ▶ The compiler is allowed to “lose” behaviors.
- ▶ Focusing on the weak memory concurrency aspect, we use high-level programming languages in both sides.
- ▶ We are also ignoring non-terminating programs.

C/C++11 to various architectures:

<https://www.cl.cam.ac.uk/~pes20/cpp/cpp0xmappings.html>

Compilation correctness

$$\llbracket \text{compile}_{\text{SC} \rightarrow \text{TSO}}(P) \rrbracket_{\text{TSO}} \subseteq \llbracket P \rrbracket_{\text{SC}}$$

Two possible compilation schemes:

- ▶ place a TSO fence after every write; **or**
- ▶ place a TSO fence before every read

Both schemes ensure there is a fence between a write and a subsequent read.


Example: Compilation of SB

$$\text{compile}_{\text{SC} \rightarrow \text{TSO}} \left(x := 1; \parallel y := 1; \right. \\ \left. a := y \parallel b := x \right) = \begin{array}{l} x := 1; \\ \mathbf{fence}; \\ a := y \end{array} \parallel \begin{array}{l} y := 1; \\ \mathbf{fence}; \\ b := x \end{array}$$

Correctness of compilation of SC to TSO

Consider the first scheme (**fence after writes**).

- ▶ We show that every outcome of $\text{compile}_{\text{SC} \rightarrow \text{TSO}}(P)$ under TSO is an outcome of P under SC.
- ▶ We will use the operational semantics of SC and TSO.



Correctness of compilation of SC to TSO

- ▶ To show: $\llbracket \text{compile}_{\text{SC} \rightarrow \text{TSO}}(P) \rrbracket_{\text{TSO}} \subseteq \llbracket P \rrbracket_{\text{SC}}$
- ▶ To simplify the argument, we split the proof in two steps.
- ▶ First, just insert a fence after each write and show:

$$\llbracket \text{compile}_{\text{SC} \rightarrow \text{TSO}}(P) \rrbracket_{\text{SC}} \subseteq \llbracket P \rrbracket_{\text{SC}}$$

This is trivial because the SC machine simply ignores fences:

$$M \xrightarrow{i:F} M$$

- ▶ Then, we show that $\llbracket P_0 \rrbracket_{\text{TSO}} \subseteq \llbracket P_0 \rrbracket_{\text{SC}}$ for every program P_0 that has a fence after every write.

Simulation for the second step

Assumption on the program

For every $i \in \text{Tid}$, if $P_0(i), s_0 \xrightarrow{*} \xrightarrow{W(x,v)} c, s$, then $c, s \xrightarrow{\varepsilon} \xrightarrow{*} \xrightarrow{F} _$.
(notation: $q \xrightarrow{!} _ \triangleq \exists q'. q \xrightarrow{!} q'$)

A simulation relation:

$\langle P, S, M, B \rangle \mathcal{R} \langle P_{SC}, S_{SC}, M_{SC} \rangle$ if the following hold:

- ▶ for every $i \in \text{Tid}$, one of the following holds:
 - ▶ $P(i) = P_{SC}(i)$, $S(i) = S_{SC}(i)$, and $B(i) = \epsilon$
 - ▶ $P(i), S(i) \xrightarrow{\varepsilon} \xrightarrow{*} \xrightarrow{F} _$ and there exist x, v such that $B(i) = \langle x, v \rangle$ and $P_{SC}(i), S_{SC}(i) \xrightarrow{W(x,v)} \xrightarrow{\varepsilon} \xrightarrow{*} P(i), S(i)$
 - ▶ $M = M_{SC}$
-
- ▶ Show that \mathcal{R} is a simulation relation.
 - ▶ Deduce that every outcome of P_0 under TSO is also an outcome under SC.

The correctness of the alternative compilation scheme (**fence before reads**) can be proved similarly.

In fact, it suffices to ensure the following property:

Assumption on the program

$\neg(P_0(i), s_0 \rightarrow^* \xrightarrow{W(x,v)} \xrightarrow{\varepsilon}^* \xrightarrow{R(x',v')} _))$ for every i, x, v, x', v' .

Exercise

Let P_0 be a program satisfying the assumption above.

- ▶ Devise a simulation relation relating executions of P_0 under TSO with those under SC.
- ▶ Show that every outcome of P_0 under TSO is also an outcome under SC.

Compilation correctness

$$\llbracket \text{compile}(P) \rrbracket_{\text{target memory model}} \subseteq \llbracket P \rrbracket_{\text{source memory model}}$$

Alternatively, compilation correctness can be proven using declarative semantics.

Definition (Allowed outcome under a declarative model)

An outcome O is *allowed* for a program P under X if there exists an execution graph G such that:

- ▶ G is an execution graph of P with outcome O .
- ▶ G is X -consistent.

Given an execution graph of $\text{compile}(P)$ which is consistent according to the target memory model, we have to show an execution graph of P with the same outcome which is consistent according to the source memory model.


Declarative semantics for TSO

Owens et al. [TPHOLs '09] provided a declarative semantics for TSO, which is equivalent to the following:

Definition

An execution graph is *TSO-consistent* if it is complete,

- ▶ $(po \cup rf)^+$ is irreflexive, and
- ▶ there exists a total order *mo* on $W \cup F$, such that *none* of the following patterns occurs:



Theorem

An outcome O is allowed for a program P under TSO if there exists a TSO-consistent execution graph of P with outcome O .

Compilation of SC to TSO via the declarative semantics

- ▶ As before we may assume that the source program has fences and the SC declarative semantics simply ignores them.

Assumption on the program

For every c, s such that $P_0(i), S_0 \rightarrow^* c, s$:

if $c, s \xrightarrow{W(x,v)} \rightarrow^* \xrightarrow{R(x',v')} c', s'$ then

$c, s \xrightarrow{W(x,v)} \rightarrow^* \xrightarrow{F} \rightarrow^* \xrightarrow{R(x',v')} c', s'$

- ▶ Given a TSO-consistent execution graph G of P_0 , we have to construct an SC-consistent execution graph of P_0 with the same outcome.
- ▶ We will take the same execution graph G .
- ▶ We have $[W]; G.po; [R] \subseteq G.po; [F]; G.po$.
- ▶ It remains to show that G is SC-consistent.

Definition (Alternative SC definition)

An execution graph G is SC-consistent if the following hold:

- ▶ G is complete
- ▶ There exists a modification order mo for G such that $G.\text{po} \cup G.\text{rf} \cup \text{mo} \cup \text{rb}$ is acyclic where:
 - ▶ $\text{rb} \triangleq G.\text{rf}^{-1}; \text{mo} \setminus \text{id}$ (from-reads / reads-before)
- ▶ Take $\text{mo}_{\text{SC}} = \bigcup_{x \in \text{Loc}} [W_x]; \text{mo}_{\text{TSO}}; [W_x]$.
- ▶ Suppose that $G.\text{po} \cup G.\text{rf} \cup \text{mo}_{\text{SC}} \cup \text{rb}$ is cyclic.
- ▶ Then, $G.\text{po} \cup G.\text{rf} \cup \text{mo}_{\text{TSO}} \cup \text{rb}$ is cyclic.
- ▶ Consider a cycle of minimal length.
- ▶ In a minimal cycle, there are at most two events in $W \cup F$.
- ▶ Analyze all such cycles. . .

Compilation correctness

$$\llbracket \text{compile}_{\text{RA} \rightarrow \text{TSO}}(P) \rrbracket_{\text{TSO}} \subseteq \llbracket P \rrbracket_{\text{RA}}$$

Trivial:

- ▶ Compilation is identity.
- ▶ It suffices to note that TSO-consistency implies RA-consistency!

Theorem (Reminder)

Let mo be a modification order for an execution graph G . G is **RA-consistent** wrt mo iff the following hold:

- ▶ $(\text{po} \cup \text{rf})^+$ is irreflexive. *(no-future-read)*
- ▶ $\text{mo}; (\text{po} \cup \text{rf})^+$ is irreflexive. *(coherence-ww)*
- ▶ $\text{rf}^{-1}; \text{mo}; (\text{po} \cup \text{rf})^+$ is irreflexive. *(coherence-wr)*
- ▶ $\text{rf}^{-1}; \text{mo}; \text{mo}$ is irreflexive. *(rmw-atomicity)*

Power and ARM are given by both declarative and operational models. For the compilation proofs, the declarative models are more convenient.

Power **Herding cats: modelling, simulation, testing, and data mining for weak memory.** Jade Alglave, Luc Maranget, Michael Tautschnig, TOPLAS 2017.

ARMv8.2 **aarch64.cat on GitHub**

Since, however, the models are much more complex, compilation proofs are error-prone. See, e.g.,

- ▶ **Synchronising C/C++ and POWER.** Susmit Sarkar, Kayvan Memarian, Scott Owens, Mark Batty, Peter Sewell, Luc Maranget, Jade Alglave, Derek Williams, PLDI 2012.
- ▶ **Repairing sequential consistency in C/C++11.** Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, Derek Dreyer, PLDI 2017.

Power declarative model

An execution G is **Power-consistent** if there exists a modification order mo for G such that the following hold:

1. hb is acyclic. (*no-thin-air*)
2. $\text{po}|_{\text{loc}} \cup \text{rf} \cup \text{rb} \cup \text{mo}$ is acyclic. (*SC-per-loc*)
3. $(\text{rb} \setminus \text{po}); \text{prop}; \text{hb}^*$ is irreflexive. (*observation*)
4. $\text{mo} \cup \text{prop}$ is acyclic. (*propagation*)
5. $\text{mo}; [\text{RMW}]; \text{po}; [\text{RMW}]$ is acyclic.

where:

- $\text{sync} = \text{po}; [\text{F}^{\text{sync}}]; \text{po}$ and $\text{lwsync} = \text{po}; [\text{F}^{\text{lwsync}}]; \text{po}$
- $\text{fence} = \text{sync} \cup (([\text{R}]; \text{lwsync}; [\text{RW}] \cup ([\text{W}]; \text{lwsync}; [\text{W}]))$ (*fence order*)
- $\text{rb} = (\text{rf}^{-1}; \text{mo}) \setminus [\text{E}]$ (*read before*)
- $\text{rfe} = \text{rf} \setminus \text{po}$ (*external reads-from*)
- $\text{ppo} = \dots$ (*preserved program order*)
- $\text{hb} = \text{ppo} \cup \text{fence} \cup \text{rfe}$ (*happens-before*)
- $\text{prop}_1 = [\text{W}]; \text{rfe}^?; \text{fence}; \text{hb}^*; [\text{W}]$
- $\text{prop}_2 = ((\text{mo} \cup \text{rb}) \setminus \text{po})^?; \text{rfe}^?; (\text{fence}; \text{hb}^*)^?; \text{sync}; \text{hb}^*$
- $\text{prop} = \text{prop}_1 \cup \text{prop}_2$ (*propagation relation*)

Exercise: Correctness of compilation of SC to TSO

Use a simulation argument to prove the correctness of compilation from SC to TSO, for a compilation scheme that places a TSO fence between every write and read.

That is, let P_0 be a program such that

$\neg(P_0(i), s_0 \rightarrow^* \xrightarrow{W(x,v)} \xrightarrow{\varepsilon}^* \xrightarrow{R(x',v')} _)$ for every i, x, v, x', v' .

- ▶ Devise a simulation relation relating executions of P_0 under TSO with those under SC.
- ▶ Conclude that every outcome of P_0 under TSO is also an outcome under SC.

Exercise: Correctness of compilation of RA to Power

The compilation scheme of RA to Power that places a *lightweight fence* before every write **and** after every read.

Assuming no RMW's, prove the correctness of this compilation scheme.

You may prefer to use the following declarative model for WeakPower (which is clearly weaker than the model for Power presented before):

Definition

An execution G is *WeakPower-consistent* if there exists a modification order mo for G such that the following hold:

1. hb is acyclic. (*no-thin-air*)
2. $po|_{loc} \cup rf \cup rb \cup mo$ is acyclic. (*SC-per-loc*)
3. $(rb \setminus po); prop; hb^*$ is irreflexive. (*observation*)
4. $mo \cup prop$ is acyclic. (*propagation*)

where:

- $fence = [R]; po; [F^{lwsync}]; po; [RW] \cup [W]; po; [F^{lwsync}]; po; [W]$ (*fence order*)
- $rb = (rf^{-1}; mo) \setminus id$ (*reads before*)
- $rfe = rf \setminus po$ (*external reads-from*)
- $hb = fence \cup rfe$ (*(weak) happens-before*)
- $prop = [W]; rfe^?; fence; hb^*; [W]$ (*(weak) propagation relation*)