


①

Total Store Ordering:



In a perfect world...

- CPU writes $x=1$ and directly stores this in the Memory / RAM.
- every other CPU sees $x=1$ immediately.
- Sequential consistency - Strong Memory Model

The truth:

- memory accesses are slow: $\text{clock-rate}(\text{CPU}) \geq 10 \cdot \text{clock-rate}(\text{RAM})$.
- do not immediately write to memory!
 - ↳ if we would, parallel programming very slow!
- allow more behavior - Weak Memory Model.

On x86 architecture:

- Writes of CPU are put into a buffer.
- If there are some writes, flush to memory
- other CPUs may see changes to variables only later.

Parallel Programs:

Definition: Parallel programs with shared memory are defined by the following grammar.

$\langle \text{prog} \rangle ::= \text{program } \langle \text{name} \rangle \langle \text{thread} \rangle^*$ // name and finite list of threads

$\langle \text{thread} \rangle ::= \text{thread } \langle \text{threadid} \rangle$ // identifier
 $\text{regs } \langle \text{reg} \rangle^*$ // list of local registers
 $\text{init } \langle \text{label} \rangle$ // label of initial instruction
 $\text{begin } \langle \text{inst} \rangle^* \text{ end}$ // list of labeled instructions

$\langle \text{inst} \rangle ::= \langle \text{label} \rangle : \langle \text{inst} \rangle ; \text{goto } \langle \text{label} \rangle$

$\langle \text{inst} \rangle ::= \langle \text{reg} \rangle \leftarrow \text{mem}[\langle \text{reg} \rangle]$

// load from memory into local register

$|\text{mem}[\langle \text{reg} \rangle] \leftarrow \langle \text{reg} \rangle$

// store

$|\text{mfence}$

// memory fence

$|\langle \text{reg} \rangle \leftarrow \langle \text{expr} \rangle$

// local assignment

$|\text{assert } \langle \text{expr} \rangle$

// assertion

$\langle \text{expr} \rangle ::= \text{fun}(\langle \text{reg} \rangle^*)$

// function on fin. many registers.

Programs have a finite data domain DOM ,
and a fin. function domain FUN .

We assume:

- The thread identifiers $\langle \text{threadid} \rangle$ are distinct numbers.
- The registers $\langle \text{reg} \rangle$ are from a finite set of names (x, y, z, \dots) and they are not shared among threads.
- The labels $\langle \text{label} \rangle$ are strings and each command $\langle \text{inst} \rangle$ has a distinct label.
- The elements of DOM can be used as register content and memory addresses.
- $0 \in \text{DOM}$.
- FUN are functions over DOM , with multiple input parameters. Functions are computable.
- Each thread only jumps to its own labels and only accesses its own registers.

ISO Semantics:

Example: Dekker's mutex

$l_0: \text{mem}[x] \leftarrow 1; \text{goto } l'_1;$

$l'_1: r \leftarrow \text{mem}[y]; \text{goto } l_2;$

$l_2: \text{assert } r \neq 0; \text{goto } l_3;$

$l_3: \dots // \text{critical section}$

$l'_0: \text{mem}[y] \leftarrow 1; \text{goto } l'_1;$

$l'_1: r' \leftarrow \text{mem}[x]; \text{goto } l'_2;$

$l'_2: \text{assert } r' \neq 0; \text{goto } l'_3;$

$l'_3: \dots // \text{critical section}$


→ Simplified notation. Can easily be transformed to above grammar.

→ Under sequential consistency: at least one thread can enter critical section.

③

idea of TSO:

- Each thread has a store buffer. Stores are buffered locally and later flushed to the memory in FIFO manner.
- Stores in buffer are not visible to other threads.
- A thread can read ~~the~~ stores from its buffer; early read.



Example: Reconsider Dekker's mutex.


This time, we run it under TSO.

⇒ Get unwanted behavior:

• Both threads issue the store and put it into their buffer.

• Both threads load value 0 in r, r' .
Taken from memory.

• Both asserts go through ⇒ both threads in critical section!



TSO Semantics:

We define configurations and the TSO-transition relation.

Definition: TSO configurations

Fix program P with threads t_1, \dots, t_n .

Assume thread t_i has

- identifier i ,
- initial label $l_{0,i}$,
- registers R_i .

Use $TID := \{1, \dots, n\}$ for all thread identifiers,

$LAB :=$ all labels used by all threads,

$VAR := DOM \cup \bigcup_{i=1}^n R_i$; for all locations (addresses and registers).

A configuration of P is an element from

(4)

$$CF = LAB^{TID} \times DOM^{VAR} \times (DOM \times DOM)^{*TID}$$

it is a tuple $cf = (pc, val, buf)$, where

- $pc: TID \rightarrow LAB$ is the program counter,
- $val: VAR \rightarrow DOM$ is the valuation, and
- $buf: TID \rightarrow (DOM \times DOM)^{*}$ a collection of local buffers.

$$buf(i) = (a_1, v_1) \dots (a_n, v_n)$$

means: value v_1 is stored in address a_1 .

The initial configuration is $cf_0 = (pc_0, val_0, buf_0)$ with

$$pc_0(i) = l_0, i \in TID,$$

$$val_0(x) = 0 \quad \forall x \in VAR,$$

$$buf_0(i) = \epsilon \quad \forall i \in TID.$$

Definition: TSO transition relation.

Assume we are in configuration $cf = (pc, val, buf)$ and $pc(i) = l$ for some thread i .
We want to execute $l \langle inst \rangle; goto l'$.

Define $pc' := pc [i := l']$

The TSO transition relation $\rightarrow_{TSO} \subseteq CF \times CF$ is the smallest relation that satisfies the following rules:

$$(EAR(Y)) \frac{\langle inst \rangle = r \leftarrow mem[r], a = val(r), buf(i)_a = (a := v). \beta}{(pc, val, buf) \xrightarrow{TSO} (pc', val[r := v], buf)}$$

- $buf(i)_a$ is the projection of $buf(i)$ to stores to address a .
- $(a := v)$ means that v is stored to address a .

$$(LOAD) \frac{\langle inst \rangle = r \leftarrow mem[r], a = val(r), buf(i)_a = \epsilon, v = val(a)}{(pc, val, buf) \xrightarrow{TSO} (pc', val[r = v], buf)}$$

- $buf(i)$ contains no store on address $a \rightarrow$ read from memory.

$$(STORE) \frac{\langle inst \rangle = mem[r] \leftarrow r', a = val(r'), v = val(r')}{(pc, val, buf) \xrightarrow{TSO} (pc', val, buf [i := (a = v). buf(i)])}$$

- issue a store \rightarrow append to buffer

⑤

(UPDATE) $\frac{\text{buf}(i) = \beta \cdot (a=v)}{(pc, \text{val}, \text{buf}) \xrightarrow{\text{TSO}} (pc, \text{val}[a:=v], \text{buf}[i:=\beta])}$

- Move earliest write in buffer to memory

(MFENCE) $\frac{\langle \text{inst} \rangle = \text{mfence}, \text{buf}(i) = \epsilon}{(pc, \text{val}, \text{buf}) \xrightarrow{\text{TSO}} (pc', \text{val}, \text{buf})}$

- Blocks thread until buffer is flushed to memory completely.

(ASSERT) $\frac{\langle \text{inst} \rangle = \text{assert } e, \llbracket e \rrbracket \neq 0}{(pc, \text{val}, \text{buf}) \xrightarrow{\text{TSO}} (pc', \text{val}, \text{buf})}$

- assert only continues execution if $\llbracket e \rrbracket \neq 0$.
- $\llbracket e \rrbracket$ is the evaluated expression, where registers r are replaced by $\text{val}(r)$.

~~(FENCE)~~

(ASSIGN) $\frac{\langle \text{inst} \rangle = r \leftarrow e, \llbracket e \rrbracket = v}{(pc, \text{val}, \text{buf}) \xrightarrow{\text{TSO}} (pc', \text{val}[r:=v], \text{buf})}$

- local assignment changes register content.

TSO Reachability:

We want to solve the following problem:

TSO Reachability

Given: Program P , program counter p

Quest: Is there a computation $c_0 \xrightarrow{\text{TSO}}^* (pc, \text{val}, \text{buf})$ for some val and buf ?

Goal: Show that problem is decidable

Idea: Construct LCS L_p simulating P and use that reachability on LCS is decidable.

Approach: Construct LCS $L_p^0, L_p^1, L_p^2, L_p^3, \dots, L_p$ and L_p .

Each L_p^i fixes problems from L_p^{i-1} .

STEP 0: From P to L_p^0 :

→ Shared memory communication is like gossip:
A store might get overwritten before it is seen by another thread.

→ Understand TSO buffers as gossip channels.

\Rightarrow Construct L_p^1 with: states $LAB^{TIO} \times DOM^{VAR}$

⑥

a channels for each thread over $DOM \times DOM$
 transition relation induced by \rightarrow_{TSO} .

STEP 1: Towards L_p^1 :

Problem: Well quasi ordering for LCS is \leq^* Higman's ordering.

\rightarrow Not a simulation relation for TSO.

$l_0: r \leftarrow \text{mem}[x];$ $l'_0: \text{mem}[y] \leftarrow 1;$
 $l_1: \text{assert}(x=1);$ $l'_1: \text{mem}[x] \leftarrow 1;$
 $l_2: r \leftarrow \text{mem}[y];$ $l'_2: \dots$
 $l_3: \text{assert}(r=0);$ $l'_3: \dots$
 $l_4: \dots;$

configuration of L_p^1 (also a TSO config.)

$\cdot cf = (\underbrace{l_0, l_2}_{pc}, \underbrace{x=0, y=0}_{val}, \underbrace{(\epsilon, x=1, y=1)}_{buf})$

$\cdot cf' = (l_0, l_2, x=0, (\epsilon, x=1))$

Then $cf \xrightarrow{v_1} \text{Conf. } cf'$ can reach $(l_4, l_2, x=1, \epsilon) = \gamma$

But cf' cannot reach a conf. with $(l_4, l_2, y=0, \epsilon)$
 $(l_4, l_2, x=1, \epsilon)$

\Rightarrow No sim. relation: $cf \not\xrightarrow{v_1} \gamma'$
 $cf' \xrightarrow{v_1} \gamma$

\Rightarrow Lossiness gives inconsistent memory configurations. $(x=1, y=0)$.

Fix: In L_p^1 , ~~the~~ an issuing sends a memory snapshot via the channel. contains values for all memory addresses.

Problem vanishes: $(\underbrace{l_0, l_2}_{pc}, \underbrace{(x=0, y=0)}_{val}, (\epsilon, (x=1, y=1)))$
 and $(l_0, l_2, (x=0, y=0), (\epsilon, (x=1, y=0)))$

corresponding to cf and cf' are incomparable now.

⑦

STEP 2: Towards L_p^2 :

Problem: L_p^1 allows behavior not possible under TSO.

$l_0: \text{mem}[y] \leftarrow 0;$	$l'_0: \text{mem}[x] \leftarrow 1;$
$l_1: \dots$	$l'_1: r \leftarrow \text{mem}[x];$
	$l'_2: \text{assert}(r == 0);$
	$l'_3: \dots$

→ (l_2, l'_3) is not reachable under TSO;
We can only load 1 from x, since store in l'_0 has been performed.

→ But a configuration with (l_2, l'_3, \dots) is reachable in L_p^1 :

$(l_0, l'_0, (x=0, y=0), (\epsilon, \epsilon))$

→_{TSO} $(\{l_2, l'_2\}, (x=0, y=0), ((x=0, y=0), (x=2, y=0)))$ // buffer stores

→_{TSO} $(\{l_2, l'_2\}, (x=1, y=0), ((x=0, y=0), \epsilon))$ // update

→_{TSO} $(l_2, l'_2, (x=0, y=0), (\epsilon, \epsilon))$ // update

→_{TSO} $(l_2, l'_3, (x=0, y=0), (\epsilon, \epsilon))$

⇒ Threads do not synchronize on memory updates.
Use values that are no longer in mem.

Fix: All threads share the same buffer.

In our example we would get:

$(l_2, l'_2, (x=0, y=0), (x=2, y=0))$

store of
y takes x=2
into account.

STEP 3: Towards L_p^3 :

Problem: L_p^2 allows less behavior than TSO.

$\text{mem}[x] \leftarrow 1;$ ① ②	$r_2 \leftarrow \text{mem}[y];$ ⑫	$\text{mem}[y] \leftarrow 1;$ ③ ④	$r_4 \leftarrow \text{mem}[x];$ ⑧
$\text{mem}[x] \leftarrow 2;$ ⑥ ⑦	$\text{assert } r_2 == 2;$ ⑬	$r_3 \leftarrow \text{mem}[x];$ ④	$\text{assert } r_4 == 2;$ ⑨
	$r_2 \leftarrow \text{mem}[y];$ ⑮	$\text{assert } r_3 == 1;$ ⑤	$\text{mem}[y] \leftarrow 2;$ ⑩
	$\text{assert } r_2 == 1;$ ⑯		

This behavior is not possible in L_p^2 :


- All threads use one channel / buffer
- So, $y=1$ will be sent to memory before $y=2$.

→ In L_p^2 , memory updates in the buffer are delivered to the memory in the order in which they ~~occur~~ entered the buffer.


⇒ In t_2 , it is not possible to load $y=2$ before $y=1$.

→ In TSO, memory updates can have any order if they stem from different threads.

Fix: Add for each thread, a pointer to a position inside the buffer.




Then, updates are simulated by moving the pointer to the left. Looks like:



STEP 4: Towards L_p^3 :

Problem: In L_p^3 , early reads are not modeled.

Fix: Remember the last write of a thread to an address.



Formal Construction: Let P be a program.

The LCS with strong symbols L_p is defined by

$$L_p = (Q, q_0, C, M, S, \rightarrow)$$

cannot be lost but are bounded.

9

where $Q = \underbrace{LAB}_{pc}^{TID} \times \underbrace{DOM}_{val}^{VAR}$

$C = \{buf\}$

$M = \underbrace{DOM}_{DOM} \rightarrow \text{NOT } \underbrace{DOM}_{VAR}$ as local registers are not in the shared memory.

$S = \left(\underbrace{DOM}_{DOM} \times \underbrace{(TID \times \underbrace{DOM}_{last\ write} \cup \{E\})}_{\text{threads pointing to snapshot } E} \times \underbrace{P(TID)}_{\text{threads pointing to snapshot } E} \right) / \left(\underbrace{DOM}_{DOM} \times \{E\} \times \{\emptyset\} \right)$
 normal mem. states/updates can get lost.

Remarks:

The number of string symbols that can occur is bounded by $|TID| \cdot |DOM| \cdot |TID|$.

→ For each $a \in DOM$, there are at most $|TID|$ last writes to it, and for each thread, we need one pointer.

Definition: Transition relation of L_p :

In a state $(pc, val) \in Q$, with $pc(t) = l$.
 valuation of local registers and copy of last mem. valuation sent to buffer. } memory content when complete buffer was flushed

We describe transitions that are realized by using sequences of transitions.

Store: If $l: mem[r] \leftarrow r'; goto l'$

Let $val(r) = a, val(r') = v$.
 $(pc, val) \xrightarrow{buf!(mval, E, tidset) / (mval, (t, a), tidset)} *$

realizable with LCS transitions } means: check if buf contains string symbol ~~with~~ $(mval, (t, a), tidset)$ with last write (t, a) .
 → Replace it by $(mval, E, tidset)$. Last write to a will be new after transition. of t

* $\xrightarrow{buf!(val[DOM[a]=v], (a, t), \emptyset)} (pc[t=l'], val[a=v])$.
 memory snapshot with updated a

Load: If $l: r \leftarrow \text{mem}[r]; \text{goto } l'; \text{ and } \text{val}(r') = a.$

Two cases:

Early read:

$(pc, \text{val}) \xrightarrow{\text{assert } (\text{mval}, (t, a), \text{tidset}) \in \text{buf}} (pc[t=l'], \text{val}[r = \text{mval}(a)])$

Load from memory:

$(pc, \text{val}) \xrightarrow{\substack{\text{assert } (\text{mval}', (t, a), \text{tidset}') \notin \text{buf} \\ \text{and } (\text{mval}, *, \{t\} \cup \text{tidset}) \in \text{buf}}} (pc[t=l'], \text{val}[r = \text{mval}(a)])$

fence: If $l: \text{mfence}; \text{goto } l';$

$(pc, \text{val}) \xrightarrow{\text{assert } (\text{mval}, *, \{t\} \cup \text{tidset}) \neq \text{head}(\text{buf})} (pc[t=l'], \text{val}).$

Update:

$(pc, \text{val}) \xrightarrow{\text{assert } \text{buf} = w_1 \cdot m_1 \cdot m_2 \cdot w_2} (pc, \text{val}).$
 modify $\text{buf} = w_1 \cdot m_2' \cdot m_1' \cdot w_2$

with: $m_1 = (\text{mval}_1, \text{last}_1, \text{tidset}_1),$

$m_2 = (\text{mval}_2, \text{last}_2, \{t\} \cup \text{tidset}_2)$ // current mem. state of t

$m_1' = (\text{mval}_1, \text{last}_1 \setminus \{t, *\}, \text{tidset}_1 \cup \{t\})$

$m_2' = (\text{mval}_2, \text{last}_2, \text{tidset}_2).$

Assert: If $l: \text{assert } e; \text{goto } l';$

$(pc, \text{val}) \xrightarrow{\text{assert } [e] \neq 0} (pc[t=l'], \text{val})$

- does not involve buffer
- $[e]$ is computed from val

Assign: If $l: r \leftarrow e; \text{goto } l';$

$(pc, \text{val}) \rightarrow (pc[t=l'], \text{val}[r = [e]])$

Theorem: Atig, Bouajjani, Brockschardt, Musuvathi 2010/12/16

Consider a program P . One can construct an LCS (with strong symbols) L_P such that a control state

pc is reachable in P , $\Leftrightarrow pc$ is reachable executed under TSO in L_P .

Hence, control state reachability for TSO is decidable.

11

Why is this correct? We give an idea.

Let A_1, A_2 be NFAs over disjoint alphabets $\Sigma_1, \Sigma_2: \Sigma_1 \cap \Sigma_2 = \emptyset$.

The shuffle is the language of all interleavings:

$$L(A_1) \sqcup L(A_2) = \{w \in (\Sigma_1 \cup \Sigma_2)^* \mid w|_{\Sigma_1} \in L(A_1) \wedge w|_{\Sigma_2} \in L(A_2)\}$$

We construct an automaton for the shuffle:

- Add to A_1 the following transitions

For each $q \in Q_{A_1}, a \in \Sigma_2$ add $q \xrightarrow{a} q$.

- Similarly for A_2 , where we add loops over Σ_1 .


→ We obtain A_1' and A_2' over $\Sigma_1 \cup \Sigma_2$ and

we have:


$$L(A_1') \cap L(A_2') = L(A_1) \sqcup L(A_2)$$

So, the product $A_1' \times A_2'$ accepts the shuffle.

Example: $\{ab\} \sqcup \{xyz\}$



Apply trick to TSO:




Memory sees a shuffle of the stores of both threads.

The stores by t_i arrive in order, but are interleaved with stores from other threads.



Use above trick: assume each thread has its own memory.


Add loops to each that produces arbitrary writes of other threads. Then intersect



The channels/buffers are FIFO.


⇒ The stores leave the buffer in the order in which they were put in.

Instead of guessing the stores of the other thread at the memory, we let each thread guess these stores when inserting commands in the buffer:



Now the content of both buffers is the same.

But Mem1 and Mem2 may process at different speeds:



This can be understood as one buffer with two pointers

