

11. Robustness against Partial Store Ordering

Goal: • Solve robustness against partial store ordering (PSO), a memory model weaker than TSO.

- PSO does not exist in practice
 - ↳ PGHS - clusters at Fraunhofer ITWM (partitioned global address space) behave similarly
 - ↳ The approach to show decidability carries over to Power.

Problem: PSO is not local

Approach: • Show a weaker normal form result

- Develop an automaton model stronger than SC reachability that is able to detect such normal form computations.

11.1 Partial Store Ordering and Robustness

Syntax of programs: It's before.

PSO - Semantics of programs:

- Idea:
- There is a buffer per address (per thread)
 - Hence, stores to different addresses can overwrite each other.
 - Loads read their values early from the corresponding buffer
 - It fence fences all buffers (heavy-weight synchronization)

Technically: • Configurations take the form (pc, mem, buf)

with

$$pc: TID \rightarrow LOC$$

$$mem: Var \rightarrow DOM, \text{ here } VAR := DOM \cup \bigcup_{t \in TID} \epsilon^*$$

$$buf: TID \times DOM \rightarrow DOM^*$$

- The set of actions that serve as transition labels is

$$ACT := TID \times (\{local, refence\} \cup \{isu, st\} \times DOM)$$

- The definition of the (labelled) transition relation for PSO is left as homework.

- We use

$$C_{PSO}(P) := \{ \tau \in ACT^* \mid \text{so } \xrightarrow{PSO} s \text{ with } buf(t, a) = \epsilon \text{ for all } t \in TID, a \in DOM \}$$

for the set of all PSO computations.

- As in the case of TSO, we associate with a computation $\tau \in C_{PSO}(P)$ a trace $Tr(\tau)$ that only keeps \rightarrow_{po} , \rightarrow_{st} , and \rightarrow_{src} .

Moreover, \rightarrow_{cf} is again a derived relation.

We use

$$Tr_{PSO}(P) := Tr(C_{PSO}(P))$$

for the set of all PSO traces.

- The robustness problem against PSO is defined as follows

Given: A parallel program P .

Question: Does $Tr_{PSO}(P) \subseteq Tr_{sc}(P)$ hold?

- To find a way to handle PSO robustness, we again phrase robustness in terms of (absence of) happens-before cycles.
- Recall that \rightarrow_{src} and \rightarrow_{st} define \rightarrow_{cf} , and that the happens-before relation of $Tr(\tau)$ is

$$\rightarrow_{hb} := \rightarrow_{po} \cup \rightarrow_{src} \cup \rightarrow_{st} \cup \rightarrow_{cf}.$$

Now the Shasha & Snir lemma still holds for PSO.

Lemma (Shasha & Snir 1988):

Consider $Tr(\tau) \in Tr_{PSO}(P)$.

Then $Tr(\tau) \in Tr_{sc}(P)$ iff \rightarrow_{hb} is acyclic.

Again, we call a computation $\tau \in Cr_{PSO}(P)$ violating if $Tr(\tau)$ is cyclic (in terms of happens-before).

11.2 Non-Locality and Normal-Form Violations

Goal: Show that we can assume a particular shape of violating computations.

Proposition Conjecture:

PSO is not local.

Proof:

- In a simplified programming model, the following program disproves locality:

$$\begin{array}{l} w(z, 1); \\ w(x, 1); \end{array} \parallel \begin{array}{l} r(x, 1); \\ r(y, 1); \\ r(z, 0); \end{array} \parallel \begin{array}{l} w(z, 1); \\ w(y, 1); \end{array}$$

- When translating this program to our assembly model, it no longer serves as a counterexample to locality:

$$\begin{array}{l}
 l_1: \text{mem}[z] \leftarrow 1 \text{ goto } l_2; \\
 l_2: \text{mem}[x] \leftarrow 1 \text{ goto } l_3; \\
 \end{array}
 \left| \begin{array}{l}
 l'_1: r \leftarrow \text{mem}[x] \text{ goto } l'_2; \\
 l'_2: \text{assert } r = 1 \text{ goto } l'_3; \\
 l'_3: r \leftarrow \text{mem}[y] \text{ goto } l'_4; \\
 l'_4: \text{assert } r = 1 \text{ goto } l'_5; \\
 l'_5: r \leftarrow \text{mem}[z] \text{ goto } l'_6; \\
 l'_6: \text{assert } r = 0 \text{ goto } l'_7; \\
 \end{array}
 \right|
 \begin{array}{l}
 l''_1: \text{mem}[z] \leftarrow 1 \\
 \text{goto } l''_2; \\
 l''_2: \text{mem}[y] \leftarrow 1 \\
 \text{goto } l''_3; \\
 \end{array}$$

Problem: The assert at l'_6 is not needed to find a cycle.

- If the assert is missing, there is no need to reorder l''_1 and l''_2 .

Lesson learned:

- ↳ Commands $r(x, v)$ are too coarse an abstraction.
- ↳ The combination of load and assert is stronger than what can be found in practice.

Whether PSD is local is open.

Definition (Normal-Form Computation):

A computation $\tau \in \text{PSO}(P)$ is in normal form if it consists of at most three parts

$$\tau = \tau_1 \cdot \tau_2 \cdot \tau_3$$


so that

(1) τ_2, τ_3 only contain store actions that were delayed, (without issues).

(2) there are no delays within τ_1 , and


(3) for all $is_{u_1}, is_{u_2} \in \tau_1$ with associated $st_1, st_2 \in \tau_i, i=2,3$, we have $is_{u_1} < is_{u_2}$ implies $st_1 < st_2$.

Illustration:



↳ τ is in normal form.

↳ The following computation τ' is not in normal form:



What is cool about this normal form?

If τ_1 , τ_2 , and τ_3 we processed simultaneously, they can be accepted in streaming-fashion, without need for unbounded storage.

Theorem (Normal-Form Result, Călin, Derevenek, Rajumder, M. '13)

Program P is robust iff there is no violation $\tau \in \text{PSO}(P)$ in normal form.

The direction from left to right is by Shasha & Snir.

For the reverse implication, we need combinatorics:

↳ For TSO we studied a violation with fewest delays.

↳ For PSO we now concentrate on a shortest violation.

Need an interesting property of PSO.

Lemma (Cancellation):

Let $E \neq \tau \in \text{PSO}(P)$.

There is an action $a \in \tau$ so that

$\tau \setminus a \in \text{PSO}(P)$.

Here,

$\tau \setminus a := \tau$ with a removed.

If a is an issue, we also remove the corresponding store action.

Proof (Cancellation):

Consider the last action a in τ that is not a store.

Then

- a is the program-order last action of a thread and
- the following actions are executable unconditionally (stores that leave buffers).

Hence, a can be removed while preserving feasibility, which means

$\tau \setminus a \in \text{PSO}(P)$. □

Proof (Normal-Form Result):

← Consider a shortest violation $\tau \in \text{PSO}(P)$.

With the cancellation lemma,

τ contains an action a that can be removed.

There are two cases:

$a \neq \text{isu}$:

$\tau = \tau_1 \cdot a \cdot \tau_2$ and $\tau \setminus a := \tau_1 \cdot \tau_2$

$a = \text{isu}$:

$\tau = \tau_1 \cdot \text{isu} \cdot \tau_2 \cdot \text{st} \cdot \tau_3$ and $\tau \setminus a := \tau_1 \cdot \tau_2 \cdot \tau_3$

We focus on the latter case.

We have

$$\bar{\tau} \upharpoonright a := \bar{\tau}_1 \cdot \bar{\tau}_2 \cdot \bar{\tau}_3$$

is shorter than $\bar{\tau}$.

By minimality of $\bar{\tau}$,

$\bar{\tau} \upharpoonright a$ cannot be violating.

Hence, there is an SC computation $\sigma \in \mathcal{L}_{SC}(P)$ with

$$\text{Tr}(\bar{\tau} \upharpoonright a) = \text{Tr}(\sigma).$$

We project σ to $\bar{\tau}_1$, $\bar{\tau}_2$, and $\bar{\tau}_3$,
and reinsert a :

$$\bar{\tau}' := (\sigma \downarrow \bar{\tau}_1) \cdot \text{isu} \cdot (\sigma \downarrow \bar{\tau}_2) \cdot \text{st} \cdot (\sigma \downarrow \bar{\tau}_3).$$

It can be checked that

$$\bar{\tau}' \in \mathcal{L}_{PSO}(P),$$

$$\text{Tr}(\bar{\tau}') = \text{Tr}(\bar{\tau}), \text{ and}$$

$\bar{\tau}'$ is in normal form.

□