

TSO Robustness 1	2
TSO Robustness 2	10
TSO Robustness 3	15
TSO Robustness 4	21
TSO Robustness 5	25

## 10. Robustness against Total Store Ordering-

(Bouajjan, M., Möhlmann, ICFLP 2011)

(Bouajjan, Derevenete, M., ESOP 2013)

Claim:

- Programmers think in terms of sequential consistency
- Hence, behaviour that deviates from SC should be considered a programming error.

Correctness criterion:

- The program behaviour under TSO should coincide with the SC behaviour:

$$B_{TSO}(P) = B_{SC}(P).$$

- In this case, the program is called robust against (execution under) TSO.

### Notion of behaviour:

#### Trade-off:

- ↳ Want a weak comparison between TSO and SC to maximise the benefits of relaxed executions
- ↳ The weaker the correspondence between TSO and SC, the harder it is to check the correspondence.

(Analogue:

Language equivalence  $=_L$  between finite automata is PSPACE-complete

while

the stronger bisimulation equivalence  $=_B$  is in PTIME



## Options:

- State-based robustness:  $\text{Reach}_{TSO}(P) = \text{Reach}_{SC}(P)$

↳ very good: weak notion that allows for quite relaxed executions.

↳ bad: as hard to check as reachability under TSO.

- (Happens-before) Trace-based robustness:  $\text{Traces}_{TSO}(P) = \text{Traces}_{SC}(P)$

↳ good: still quite relaxed executions possible

↳ very good: only PSPFLE-complete.

- Goal:
- (1) Show that trace-based robustness is decidable, actually PSPFLE-complete.
    - Reduce trace-based robustness to classic SC reachability.
  - (2) Develop an algorithm that turns a non-robust program into a robust one.

## 10.1 Traces and Robustness

- To define robustness, we define (SC-) happens-before traces
- To define happens-before traces, we define computations under SC and under TSO.
- To define computations, we label the transition relation by actions from

$$\text{ACT} := \text{TID} \times (\{\text{isu}, \text{loc}\} \cup \{\text{ld}, \text{st}\} \times \text{DOM} \times \text{DOM})$$

### Definition:

The labeled TSO transition relation

$$\rightarrow_{TSO} \subseteq \text{CF} \times \text{ACT} \times \text{CF}$$

is defined by the following rules, which assume  $cf = (pc, val, buf)$   
 with  $pc(t) = l$  with  $l: \langle inst \rangle \xrightarrow{go to} l'$   
 and  $pc' := pc[t := l']$ :

(EARLY)

$$\frac{\langle inst \rangle = r \leftarrow mem[r'], a = val(r') \quad buf(t) \downarrow (a = *) = (a = v). \beta}{cf \xrightarrow{(t, ld, a, v)}_{TSO} (pc', val[r := v], buf)}$$

(LOAD)

$$\frac{\langle inst \rangle = r \leftarrow mem[r'], a = val(r'), v = val(a) \quad buf(t) \downarrow (a = *) = \epsilon}{cf \xrightarrow{(t, ld, a, v)}_{TSO} (pc', val[r := v], buf)}$$

(STORE)

$$\frac{\langle inst \rangle = mem[r] \leftarrow r', a = val(r), v = val(r')}{cf \xrightarrow{(t, st, a, v)}_{TSO} (pc', val, buf[t := (a = v). buf(t)])}$$

(UPDATE)

$$\frac{buf(t) = \beta. (a = v)}{cf \xrightarrow{(t, st, a, v)}_{TSO} (pc, val[a := v], buf[t := \beta])}$$

The remaining transitions are labelled by  $(t, loc)$ .

The set of TSO computations is

$$C_{TSO}(P) := \{ \tau \in ACT^* \mid s_0 \xrightarrow{\tau}_{TSO} s \text{ for some } s = (pc, val, buf) \text{ with } buf(t) = \epsilon \text{ for all } t \in TID \}$$

For sequential consistency (SC),

stores are not buffered — to be precise, buffered and immediately flushed.

So  $C_{SC}(P)$  is a special case of the TSO computations.

## Example (Decker = Store buffering (SB))

$l_0: \text{mem}[x] \leftarrow 1; \text{goto } l_1; \quad \parallel \quad l'_0: \text{mem}[y] \leftarrow 1; \text{goto } l'_1;$   
 $l_1: r_1 \leftarrow \text{mem}[y]; \text{goto } l_2; \quad \parallel \quad l'_1: r_2 \leftarrow \text{mem}[x]; \text{goto } l'_2;$

is TSO computation:

$\tilde{c} = (t_1, \text{isa}), (t_1, \text{ld}, y, 0), (t_2, \text{isa}), (t_2, \text{st}, y, 1), (t_2, \text{ld}, x, 0), (t_1, \text{st}, x, 1)$

## Definition (Trace):

Consider  $\tilde{c} \in (\text{TSO}(P))$ .

The trace  $\text{Tr}(\tilde{c})$  is a node-labelled graph

$\text{Tr}(\tilde{c}) := (N, \lambda, \rightarrow_{po}, \rightarrow_{st}, \rightarrow_{src})$

with

- set of nodes  $N$
- labelling  $\lambda: N \rightarrow \text{ACT}$
- $\rightarrow_{po}, \rightarrow_{st}, \rightarrow_{src} \subseteq N \times N$

program-order, store-order (coherence relation),  
and source-relation (reads-from)

The definition is by induction on the length of the computation:

↳ Empty word  $\epsilon \rightsquigarrow$  empty trace

↳ Assume  $\text{Tr}(\tilde{c}) = (N, \lambda, \rightarrow_{po}, \rightarrow_{st}, \rightarrow_{src})$ .

Then  $\text{Tr}(\tilde{c}.act) := (N \cup \{n\}, \lambda', \rightarrow'_{po}, \rightarrow'_{st}, \rightarrow'_{src})$ ,

where the choice of node  $n$  depends on the type of act:

act =  $(t, \text{st}, a, v)$ :

Let  $n$  be the minimal node in  $\rightarrow_{po}^+$  labelled by  $\lambda(n) = \text{isa}$ .

Set  $\lambda' := \lambda[n := act]$  and  $\rightarrow_{p0}' := \rightarrow_{p0}$ .

// If we have a store, we use the moment the action was issued.

act  $\neq (t, st, a, v)$ :

Add a fresh node  $n \notin N$  to the trace.

Set  $\lambda' := \lambda \cup \{(n, act)\}$ ,

$\rightarrow_{p0}' := \rightarrow_{p0} \cup \{(\max(\rightarrow_{p0}^t), n)\}$ .

Store order  $\rightarrow_{st}'$ :

Only updated for stores  $(t, st, a, v)$ :

$\rightarrow_{st}' := \rightarrow_{st} \cup \{(\max(\rightarrow_{st}^a), n)\}$

Not changed otherwise.

Source relation  $\rightarrow_{src}'$ :

Only updated for loads and stores:

Load  $(t, ld, a, v)$ :

$\rightarrow_{src}' := \rightarrow_{src} \cup \{(\max(\rightarrow_{st}^a), n)\}$ .

Store  $(t, st, a, v)$ :

Update the source relation for all loads that read early from this store:

$\forall m \in N$  with  $n \rightarrow_{p0}' m$  and  $\lambda(m) = (t, ld, a, v)$ :

$\rightarrow_{src}' := (\rightarrow_{src} \setminus \{(*, m)\}) \cup \{(n, m)\}$ .

Consider  $Tr(\tau) = (N, \lambda, \rightarrow_{p0}, \rightarrow_{st}, \rightarrow_{src})$ .

The conflict relation (from-read)


is derived from  $\rightarrow_{src}$  and  $\rightarrow_{st}$ .

We define

$ld \rightarrow_{cf} st$  if  $\exists st' : st' \rightarrow_{src} ld$  and  $st' \rightarrow_{st} st$

as  $ld$  loads the initial value and  $st$  is the first store on the address

Illustration:



The (SC-) happens-before relation of the trace

$$\text{is } \rightarrow_{hb} := \rightarrow_{po} \cup \rightarrow_{it} \cup \rightarrow_{src} \cup \rightarrow_{cf}.$$


Let  $Tr_{Tso/Sc}(P) := Tr(C_{Tso/Sc}(P))$  denote the set of all Tso/Sc traces of P.

Example:

Consider

$$\tau = (t_1, isu), (t_1, ld, y, 0), (t_2, isu), (t_2, st, y, 1), (t_2, ld, x, 0), (t_1, st, x, 1).$$

The corresponding trace + conflict relation is



The decision problem we tackle is the following:

Definition (Robustness):

Given: A parallel program P.

Problem: Does  $Tr_{Tso}(P) = Tr_{Sc}(P)$  hold?

Note:

This notion of trace-based robustness is really stronger than state-based robustness:

Lemma:

If  $Tr_{Tso}(P) = Tr_{Sc}(P)$  then  $Reach_{Tso}(P) = Reach_{Sc}(P)$ .

The reverse implication does not hold.

Note:

- ↳ Inclusion  $Tr_{SC}(P) \subseteq Tr_{TSO}(P)$  always holds, have to check the reverse inclusion.
- ↳ How to check  $Tr_{TSO}(P) \subseteq Tr_{SC}(P)$ ?  
Cannot complement an automaton.

Characterisation of robustness:

- ↳ Computation  $\bar{c} \in C_{TSO}(P)$  is called violating if  $Tr(\bar{c}) \notin Tr_{SC}(P)$ .
- ↳ Observe that violating computations employ cyclic accesses to addresses that SC is unable to serialize.
- ↳ These cyclic accesses are made visible by the conflict relation.

Lemma (Shasha & Smir, TOPLAS 1988):

Consider  $Tr(\bar{c}) \in Tr_{TSO}(P)$ .

Then  $Tr(\bar{c}) \in Tr_{SC}(P)$  iff  $\rightarrow_{hb}$  is acyclic.

Proof:

$\Rightarrow$  " SC only generates computations with acyclic  $hb$ -relation.

$\Leftarrow$  " Every partial order can be extended to a total order.

This total order is an SC computation. □

Comment:

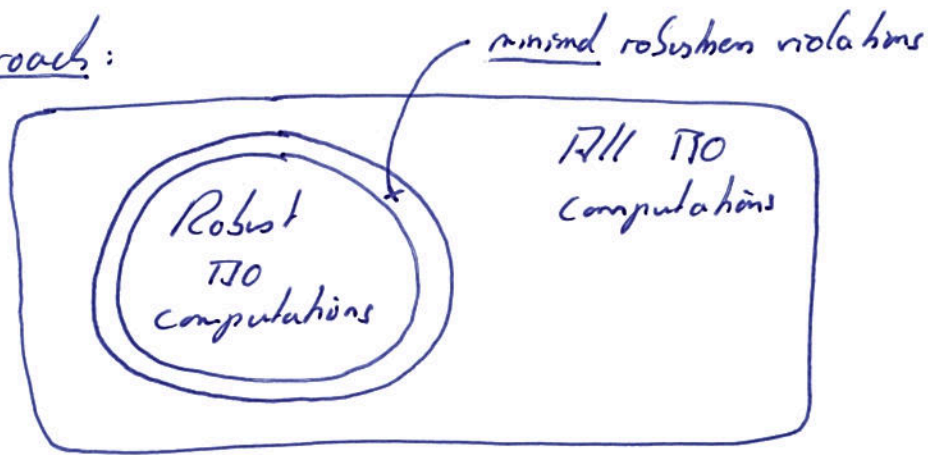
↳ Result of Shasha & Smir is semantical, does not provide an algorithm to find cyclic traces.

↳ Our main goal is to show

that cyclic traces can be found in PSPACE.



Approach:



- (1) Understand shape of minimal robustness violations  
(locality result for TSO: only a single thread has to delay stores).
- (2) Devise an algorithm to detect violations of this shape.

Combines

- combinatorial reasoning for (1)
- with
- algorithm design for (2).

## 10.2 Minimal Violations and Locality:

Goal: Show that in a minimal violation only a single thread records its actions.

Definition (Minimal violation):

Consider  $\tau = \alpha . a . \beta . b . \gamma \in C_{TSO}(P)$   
with  $\text{thread}(a) = t = \text{thread}(b)$ .

• Then the distance of a and b in  $\tau$   
is  $d_{\tau}(a, b) := |\beta \downarrow t|$ .

• The number of delays in  $\tau$  is

$$\#(\tau) := \sum_{\substack{\text{corresponding} \\ \text{isu, st} \in \tau}} d_{\tau}(\text{isu}, \text{st}).$$

• A violating computation  $\tau$  is minimal  
if  $\#(\tau)$  is minimal among all violating computations.

Note:

Program  $P$  is not robust iff it has a minimal violation.

Lemma 1 (Delays in minimal violations are required):

Consider  $\tau = \alpha . \text{isu} . \beta . \text{st} . \gamma \in C_{TSO}(P)$  a minimal violation  
with  $\text{isu}, \text{st}$  from the same instruction of thread  $t$ .

Then  $\beta \downarrow t = \epsilon$

or  $\beta \downarrow t = \beta' . \text{ld} . \beta''$  with  $\text{addr}(\text{ld}) \neq \text{addr}(\text{st})$   
and  $\beta''$  contains only stores.

Proof: Suppose  $\beta$  contains one or more actions of thread  $t$ .

• If all actions of  $t$  in  $\beta$  are stores,

then also  $\tau' := \alpha . \beta . \text{isu} . \text{st} . \gamma \in C_{TSO}(P)$ .

The computation has the same trace as  $\tau$ ,  
but

$$\#(\tau') < \#(\tau). \quad \text{by Minimality.}$$

• Let  $a$  be the last non-store action in  $\beta \downarrow t$ :

$$\beta = \beta_1 \cdot a \cdot \beta_2.$$

This means all actions of  $t$  in  $\beta_2$  are stores,  
the remaining actions belong to other threads.

• Since store actions cannot be delayed past a fence,  
 $a$  is <sup>(1)</sup> issue, <sup>(2)</sup> local action, <sup>(3)</sup> or load.

In case (1), (2), and (3) with early load  
( $\text{addr}(a) = \text{addr}(st)$ ),  
delaying  $st$  past  $a$  can be avoided:

$$\tau' := \alpha \cdot \text{isu} \cdot \beta_1 \cdot \beta_2 \cdot st \cdot a \cdot \gamma \in \text{TSO}(P).$$

Again  $\text{Tr}(\tau') = \text{Tr}(\tau)$  and  $\#(\tau') < \#(\tau)$ .  $\square$

Goal: Detect happens-before cycles in a trace (graph structure)  
on the computation (linear structure).  $\square$

Definition (Happens before through):

Let  $\tau = \alpha \cdot a \cdot \beta \cdot b \cdot \gamma \in \text{TSO}(P)$ .

Then  $a$  happens before  $b$  through  $\beta$

if there is a subsequence  $c_1, \dots, c_n$  of  $\beta$

with  $c_i \rightarrow_{hb} c_{i+1}$  or  $c_i \xrightarrow{po} c_{i+1}$  for  $0 \leq i < n$   
with  $c_0 := a$   
and  $c_{n+1} := b$ .

Lemma 2 (Happens before through is stable under insertion):

Consider  $\tau = \alpha \cdot a \cdot \beta \cdot b \cdot \gamma$

and  $\tau' = \alpha' \cdot a \cdot \beta' \cdot b \cdot \gamma' \in (TSO(P))$

so that  $\tau \downarrow t = \tau' \downarrow t$  for all  $t \in TID$ .

Moreover, assume  $\beta$  is a subsequence of  $\beta'$ .

If  $a \rightarrow_{hb}^+ b$  through  $\beta$ , then  $a \rightarrow_{hb}^+ b$  through  $\beta'$ .

Proposition (Dichotomy):

Consider a minimal violation  $\tau = \alpha \cdot a \cdot \beta \cdot b \cdot \gamma \in (TSO(P))$ .

Then

(1)  $a \rightarrow_{hb}^+ b$  through  $\beta$

or (2)  $\exists \tau' = \alpha \cdot \beta_1 \cdot b \cdot \beta_2 \cdot \gamma \in (TSO(P))$

so that

$$Tr(\tau') = Tr(\tau)$$

and  $\tau' \downarrow t = \tau \downarrow t$  for all  $t \in TID$ .

Proof:

Showing (1) or (2) is equivalent to  $\neg(1) \Rightarrow (2)$ , which is what we prove.

We proceed by induction on the length of  $\beta$  and strengthen the hypothesis:

We additionally show that  $\beta_2$  is a subsequence of  $\beta$ .

II.7: Then  $\tau = \alpha a b \gamma$  and  $a \rightarrow_{hb}^+ b$ .

$|\beta| = 0$  • If  $\text{hread}(a) = \text{hread}(b)$ , then  $b \rightarrow_{po}^+ a$ .

Therefore,  $b$  is a store action which has been delayed past  $a$ .

Swapping  $a$  and  $b$  will save the delay

without changing the hacc. by minimality.

- If  $\text{Thread}(a) \neq \text{Thread}(b)$ , then either
  - $\hookrightarrow$  one of the actions is local,
  - $\hookrightarrow$  the actions access different addresses, or
  - $\hookrightarrow$  both are loads.

In all three cases, swapping the actions produces  $\tau'$  as required.

IS: Assume the statement holds for all  $\beta'$  with  $|\beta'| \leq n$ .  
Consider  $\tau = \alpha.a.\beta.c.b.\delta$  with  $|\beta.c| = n+1$ .

Since we assume  $a \not\rightarrow_{hs}^+ b$  through  $\beta.c$ ,  
we have  $a \rightarrow_{hs}^+ c$  through  $\beta$  or  $c \rightarrow_{hs}^+ b$ .

Let  $a \rightarrow_{hs}^+ c$  through  $\beta$ :

We apply the induction hypothesis to  $a$  and  $c$ .

This gives  $\tau' = \alpha.\beta_1.c.a.\beta_2.b.\delta$

with  $\text{Tr}(\tau') = \text{Tr}(\tau)$  and  $\beta_2$  a subsequence of  $\beta$ .

and  $\tau' \downarrow t = \tau \downarrow t$  for all  $t \in \text{TD}$

If we had  $a \rightarrow_{hs}^+ b$  through  $\beta_2$  in  $\tau'$ ,

then also  $a \rightarrow_{hs}^+ b$  through  $\beta.c$  in  $\tau$ .

This holds by the induction hypothesis  
(that  $\beta_2$  is a subsequence of  $\beta$ )

together with Lemma 2,

and contradicts the assumption  $a \not\rightarrow_{hs}^+ b$  through  $\beta.c$ .

Hence, we can apply the hypothesis to  $a$  and  $b$  in  $\tau'$ .

This yields

$\tau'' = \alpha.\beta_1.c.\beta_{21}b.\beta_{22}\delta$ .

Again

$$\text{Tr}(\tau'') = \text{Tr}(\tau')$$

and  $\tau'' \downarrow t = \tau' \downarrow t$  for all  $t \in \text{TID}$

and  $\beta_{22}$  a subsequence of  $\beta_2$ .

Together:

- $\text{Tr}(\tau'') = \text{Tr}(\tau)$
- $\tau'' \downarrow t = \tau \downarrow t$  for all  $t \in \text{TID}$
- $\beta_{22}$  is a subsequence of  $\beta_2$ ,  
which is a subsequence of  $\beta$ ,  
which is a subsequence of A.c.,  
so  $\beta_{22}$  is a subsequence of A.c.

Let  $c \rightarrow b$ :

We apply the induction hypothesis to  $b$  and  $c$ .

This yields

$$\tau' = \alpha a \beta . b . c \delta$$

with  $\text{Tr}(\tau') = \text{Tr}(\tau)$  and  $\tau' \downarrow t = \tau \downarrow t$  for all  $t \in \text{TID}$ .

We now apply the hypothesis to  $a$  and  $b$  in  $\tau'$   
and get

$$\tau'' = \alpha . \beta_1 . b . \beta_2 . c \delta$$

with  $\text{Tr}(\tau'') = \text{Tr}(\tau')$ ,  $\tau'' \downarrow t = \tau' \downarrow t$  for all  $t \in \text{TID}$ ,  
and  $\beta_2$  a subsequence of  $\beta$ .

Together,  $\text{Tr}(\tau'') = \text{Tr}(\tau)$ ,  $\tau'' \downarrow t = \tau \downarrow t$  for all  $t \in \text{TID}$ ,  
and  $\beta_2 . c$  is a subsequence of P.c. □

## Recapitulation:

Goal: Establish locality

In a minimal violation,  
only a single thread delays stores.

Tool: Happens-before through

Consider  $\tau_1 a \tau_2 b \tau_3 \in \text{TSO}(P)$ .


Then  $a$  happens-before  $b$  through  $\tau_2$

if there is a subsequence  $c_1, \dots, c_n$  of  $\tau_2$   
so that

$$c_i \xrightarrow[\text{st}]{\text{src, cf.}} c_{i+1} \quad \text{or} \quad c_i \xrightarrow{\text{po}} c_{i+1}$$

for all  $0 \leq i \leq n$  with  $c_0 := a$  and  $c_{n+1} := b$ .

## Illustration:



## Proposition (Dichotomy):

Consider a minimal violation  $\tau = \tau_1 a \tau_2 b \tau_3 \in \text{TSO}(P)$

Then

(1)  $a \xrightarrow{\text{hb}} b$  through  $\tau_2$

or

(2) There is  $\tau' = \tau_1 \tau_{21} b a \tau_{22} \tau_3 \in \text{TSO}(P)$

with  $\text{Tr}(\tau) = \text{Tr}(\tau')$

and  $\tau \downarrow t = \tau' \downarrow t$  for all  $t \in \text{TID}$ .

The proposition gives rise to the following  
principle for proving cycles:

"In a minimal violation, whenever a store is delayed past a load there is a happens-before cycle (that includes the two)."

Corollary:

Consider a minimal violation  $\tau = \tau_1.\text{isu}.\tau_2.\text{ld}.\tau_3.\text{st}.\tau_4 \in \text{LSD}(P)$ , where  $\text{st}$  is the store corresponding to  $\text{isu}$ .

Then  $\text{Tr}(\tau)$  is cyclic.

Proof:

We show  $\text{st} \xrightarrow{+p_0} \text{ld} \xrightarrow{+h_3} \text{st}$ .

- The same dependence holds since  $\text{isu}$  is issued before  $\text{ld}$ .
- For  $\text{ld} \xrightarrow{+h_3} \text{st}$  we argue that  $\text{ld} \xrightarrow{+h_3} \text{st}$  through  $\tau_3$ .

Indeed, with dichotomy we have

(1)  $\text{ld} \xrightarrow{+h_3} \text{st}$  through  $\tau_3$

or (2) a reordering of  $\text{ld}$  and  $\text{st}$

that does not change the trace and the threads' computations.

If we reordered  $\text{ld}$  and  $\text{st}$  using (2),

then we would change  $\tau \downarrow t$  with  $t = \text{thread}(\text{ld}) = \text{thread}(\text{st})$ .

Hence, case (1) applies and gives hb-through.  $\square$

Theorem (Locality, Bouajjani, M., Möhlmann, ICFLP '11):

In a minimal violation, only a single thread reorders its actions.



Proof:

Consider a minimal violation  $\tau \in C_{TSO}(P)$ .

Towards a contradiction,

suppose at least two threads delay stores.

- By lemma from last lecture, each store is delayed past a load of the same thread.

Let  $st_2$  be the overall least store that was delayed in  $\tau$ .

Let  $st_2$  be from thread  $t_2$ .

Let  $ld_2$  be the overall least load that was overtaken by  $st_2$ .

Similarly, let  $st_1$  be the overall least store delayed in a thread  $t_1 \neq t_2$ .

Let  $ld_1$  be the least load overtaken by  $st_1$ .

- There are the following three situations:

$$(1) \quad \tau = \tau_1 \cdot ld_1 \cdot \tau_2 \cdot st_1 \cdot \tau_3 \cdot ld_2 \cdot \tau_4 \cdot st_2 \cdot \tau_5 \in C_{TSO}(P)$$

$$(2) \quad \tau = \tau_1 \cdot ld_2 \cdot \tau_2 \cdot ld_1 \cdot \tau_3 \cdot st_1 \cdot \tau_4 \cdot st_2 \cdot \tau_5 \in C_{TSO}(P)$$

$$(3) \quad \tau = \tau_1 \cdot ld_1 \cdot \tau_2 \cdot ld_2 \cdot \tau_3 \cdot st_1 \cdot \tau_4 \cdot st_2 \cdot \tau_5 \in C_{TSO}(P)$$

Case (1):

We argue that in this case  $\tau$  is not minimal and therefore the case does not apply.

Indeed, consider

$$\tau' := \tau_1 \cdot ld_1 \cdot \tau_2 \cdot st_2 \cdot \tau_5 \in C_{TSO}(P)$$

Here,  $\tau_{st}$  contains stores of thread  $t_2$  that were issued before  $st_2$ .

• Then:

$\#(\tau') < \#(\tau)$  as the delay of  $st_2$  over  $ld_2$  is missing.

Moreover,

$Tr(\tau')$  is cyclic.

The reason is that

$ld_1 \xrightarrow{hb} st_2$  through  $\tau_2$

continues to hold.

The only critical check here is

$ci \xrightarrow{po} ci_2$ .

One can show that

- as long as  $\tau_2$  is not changed
  - and the resulting computation  $\tau'$  is feasible ( $\in C_{TSO}(P)$ ),
- then  $hb$ -through continues to hold.

Together,  $\#(\tau') < \#(\tau)$

and  $Tr(\tau')$  cyclic contradicts minimality of  $\tau$ .

Case (2):

Again, the case would imply that  $\tau$  is not minimal and can therefore not occur.

- Starting from  $ld_2$ , thread  $t_2$  does not do actions except delayed stores, until  $st_2$  (this was a lemma last time).

Therefore,  $ld_2$  and all program-order later actions of  $t_2$

can be removed without affecting feasibility of the computation.

To be precise, we also have to remove  $\tau_5$ ,  
because other threads may load from stores of  $t_2$   
that are program-ordered later than  $ld_2$ :

$$\tau' := \tau_1. \tau_2. ld_2. \tau_3. st_2. \tau_4. st_2 \in (TSO(P)).$$

• Then

$$\#(\tau') < \#(\tau).$$

Moreover,

$Tr(\tau')$  is cyclic

due to  $ld_1 \xrightarrow{+hs} st_1$  through  $\tau_3$ .

$\leq$  minimality of  $\tau$ .

Case (3):

Again we derive a contradiction to minimality of  $\tau$ .

Remember

$$\tau = \tau_1. ld_2. \tau_2. ld_2. \tau_3 \xrightarrow{+hs} st_1 \tau_4 \xrightarrow{+hs} st_2 \tau_5 \in (TSO(P)).$$

• First, we delete  $\tau_5$ .

Then we delete all actions from  $\tau_4$

that do not belong to thread  $t_2$ :

$$\tau' := \tau_1. ld_2. \tau_2. ld_2. \tau_3. st_2. (\tau_4 \downarrow t_2). st_2 \in (TSO(P)).$$

As this does not change  $\tau_2. ld_2. \tau_3$ ,

we still have

$ld_1 \xrightarrow{+hs} st_1$  through  $\tau_2. ld_2. \tau_3$

and  $st_1 \xrightarrow{+po} ld_2$ .

So

$Tr(\tau')$  is cyclic.

Moreover, deleting actions does not introduce reorderings.

Therefore,

$$\#(\bar{\tau}') \leq \#(\bar{\tau}).$$

Together, also  $\bar{\tau}'$  is a minimal violation.

We apply dichotomy again and get

$$ld_2 \xrightarrow{+}_{hb} st_2 \text{ through } \bar{\tau}_3.st_2. (\bar{\tau}_4 \downarrow t_2).$$

Moreover,

$$st_2 \xrightarrow{+}_{po} ld_2.$$

Now,  $ld_2$  is the program-order least action of thread  $t_1$  in  $\bar{\tau}'$ .

We delete it and get

$$\bar{\tau}'' := \bar{\tau}_1 \bar{\tau}_2 ld_2 \bar{\tau}_3 st_2 (\bar{\tau}_4 \downarrow t_2) st_2 \in G_{\Pi_0}(P).$$

Then

$$\#(\bar{\tau}'') < \#(\bar{\tau}') \leq \#(\bar{\tau}).$$

Moreover,

$\text{Tr}(\bar{\tau}'')$  is cyclic as the cycle

$$st_2 \xrightarrow{+}_{po} ld_2 \xrightarrow{+}_{hb} st_2 \text{ remains.}$$

by minimality of  $\bar{\tau}$ .

□

### 10.3 Attacks on Robustness

Know: If robustness breaks, then one thread delays actions.


Show: There are two instructions (a store and a load) where a delay yields a cycle.

## 10.3 Attacks on Robustness

Goal: Rephrase robustness in terms of a simple problem:  
absence of sensible attacks.

Definition (Attacks):

- An attack is a triple  $A = (t_A, st_{inst}, ld_{inst})$  consisting of
  - ↳ a thread  $t_A \in TID$ , called attacker,
  - ↳ a store instruction  $st_{inst}$  in  $t_A$ , and
  - ↳ a load instruction  $ld_{inst}$  in  $t_A$ .
- A TSO witness of  $A$  is a computation of the form



so that

- (W1) Only the attacker delays stores
- (W2) Store  $st_{inst}$  is an instance of  $st_{inst}$ .  
It is the first store of the attacker that is delayed.
- Load  $ld_{inst}$  is an instance of  $ld_{inst}$ .  
It is the last action of the attacker overtaken by  $st_{inst}$ .

- ↳ So  $\tau_2$  contains loads, assignments, asserts, and issues, but no fences and stores of the attacker.
- ↳ It may contain arbitrary actions from the other threads, called helpers.

- (W3) For all actions  $act$  in  $ld_{inst} \tau_3 st_{inst}$  we have  $ld \rightarrow_{hb}^* act$ .  
Here, an  $isu + st$  of a helper is counted as one action  $act$ .

(W4) Sequence  $\tau_4$  only consists of stores of the addresser that were issued before  $ld_{17}$  and have been delayed.

(W5)  $\forall$  these stores  $st$  satisfy  $addr(st) \neq addr(ld_{17})$ , which means  $ld_{17}$  has not read its value early.

If a TSO witness for  $\mathcal{A}$  exists, the attack is called feasible.

Example (Decker = SB):

$l_0: mem[x] \leftarrow 1$  goto  $l_1$ ;     $l'_0: mem[y] \leftarrow 1$  goto  $l'_1$ ;  
 $l_1: r_1 \leftarrow mem[y]$  goto  $l_2$ ;     $l'_1: r_2 \leftarrow mem[x]$  goto  $l'_2$ ;

There is an attack  $\mathcal{A} = (t_1, st_{inst}, ld_{inst})$

with  $st_{inst} =$  the store at  $l_0$   
 $ld_{inst} =$  the load at  $l_2$ .

$\mathcal{A}$  TSO witness of the attack is

$\tau = (t_1, is_u) \cdot (t_1, ld, y, 0) \cdot \underbrace{(t_2, is_u) \cdot (t_2, st, y, 1) \cdot (t_2, ld, x, 0)}_{\tau_3} \cdot (t_1, st, x, 1)$   
is<sub>u</sub> at  $t_1$     ld at  $t_1$      $\tau_3$     st at  $t_1$

So attack  $\mathcal{A}$  is feasible.

The program contains a symmetric attack  $\mathcal{A}'$  with  $t_2$  as the addresser.

Theorem (Characterisation of robustness with attacks):

Program  $P$  is robust iff no attack is feasible.

Proof:  
 $\Rightarrow$  If TSO witness comes with a happens-before cycle  
 $st_{\mathcal{A}} \rightarrow_{p_0}^+ ld_{\mathcal{A}} \rightarrow_{h_0}^+ st_{\mathcal{A}}$ .

$\Leftarrow$  Show that if  $P$  is not robust,  
then there is a feasible attack.

- Among the violating computations,  
we select  $\tau \in C_{TSO}(P)$  with  $\#(\tau)$  minimal.  
By locality, only one thread  $t_A$  uses its buffer.  
Hence, (W1) holds.

- Initially, the attacker  $t_A$  executes under SC,  
so stores immediately follow their issues.


- $\hookrightarrow$  This computation is embedded into  $\tau_1$ .  
Eventually, the attacker starts delaying stores.

- $\hookrightarrow$  Let  $st_{iA}$  be the first store that is delayed.

- $\hookrightarrow$  It gets delayed past several loads,  
the last being  $ld_{iA}$ .

Together, we get (W2).

Illustration:



- For the helpers, note that  $ld_{iA} \rightarrow_{hb}^+ st_{iA}$  through  $\tau_3$   
by dichotomy.

Consider the actions act with

$ld_{iA} \rightarrow_{hb}^+ act$  through the intermediary computation.

Then act can be moved before  $ld_{iA}$  using dichotomy.

Hence, (W3)

- With cycle  $st_{iA} \rightarrow_{po}^+ ld_{iA} \rightarrow_{hb}^+ st_{iA}$ .

$\tau_4$  only needs to contain stores of the attacker  
that have been delayed past  $ld_{iA}$ .

These stores are non-blocking, so the helpers can stop with the last action of  $\tau_3$ .

We can moreover assume  $ldw$  to be the program-order last action of the attacker.

(W4) holds.

- The early read could be avoided by moving  $ldw$  to  $\tau_4$ . This would save a delay and hence contradict minimality.

Together, we have a TSO witness of the attack

$(\tau_1, stinst, ldinst)$

where  $stinst$  is the instruction of  $stw$  and  $ldinst$  is the instruction of  $ldw$ .

□

Note:

- The number of attacks is quadratic in the size of the program.
- Enumerate them and check each for a TSO witness.
- How to? Instrumentation - next.



## 10.4 Instrumentation

Goal: Consider program  $P$  with abstract  $\mathcal{A} = (t_{\mathcal{A}}, \text{start}, \text{end})$ .  
Characterise TSO witnesses of  $\mathcal{A}$  in  $P$   
by SC computations in a program  $P_{\mathcal{A}}$   
that is instrumented for attack  $\mathcal{A}$ .

By instrumentation we mean we replace every thread in  $P$  by a modified version.

Crux: • Unbounded store buffers  
• Unbounded happens-before dependencies.

### Instrumenting the attacker:

Idea: • Emulate store buffering under SC  
using auxiliary addresses

- ↳ When the attacker executes the delayed store  $st_{\mathcal{A}}$ ,  
under SC it is done right behind the issue action.
- ↳ To mimic store buffering,  $st_{\mathcal{A}}$  now accesses  
an auxiliary address that helps do not leak.
- ↳ Indeed, in  $\bar{C}_{\mathcal{A}}$  the helpers are no longer active  
and hence do not access the delayed stores.
- How many auxiliary addresses?  
↳ One per address in the program (last store).

Technically: • Starting from  $st_{\mathcal{A}}$  to address  $a$ ,  
stores are replaced by  $st_{\mathcal{A}}^{\text{aux}}$  to addresses  $(a, d)$ .  $d = \text{delay}$ .

- As long as address  $a$  has not been written,  
 $(a, d)$  holds the initial value 0.
- When the attacker stores  $v$  to address  $a$ ,  
we set  $\text{mem}[(a, d)] = (v, d)$ .
- Hence,  $(a, d)$  always holds the most recent store  
to address  $a$ .

- A load  $r \leftarrow \text{mem}[a]$  of the attacker reaches value  $v$  from the buffer whenever  $\text{mem}[a, d] = (v, d)$ .  
Otherwise  $\text{mem}[a, d] = 0$ , and the load obtains  $v = \text{mem}[a]$  from memory.

Definition (Instrumentation of the attacker):

Consider thread  $t_A$  regs  $r^*$  init to begin  $\langle \text{inst} \rangle^*$  end.

Let  $\tilde{t}_A = (t_A, \text{stinst}, \text{ldinst})$  be the attack.

The instrumentation of  $t_A$  for attack  $\tilde{t}_A$  is the thread:

$\llbracket t_A \rrbracket :=$  Thread  $\tilde{t}_A$  regs  $r^*$  init to

begin

$\langle \text{inst} \rangle^*$

// source code

$\llbracket \text{stinst} \rrbracket_{\tilde{t}_A}$

// Move to copy of source code

$\llbracket \text{ldinst} \rrbracket_{\tilde{t}_A}$

$\llbracket \langle \text{inst} \rangle^* \rrbracket_{\tilde{t}_A}$

// copy of source code.

end.

with

$\llbracket l_1 : \text{mem}[e_1] \leftarrow e_2 \text{ goto } l_2 \rrbracket_{\tilde{t}_A} :=$   $l_1 : \text{mem}[e_1, d] \leftarrow (e_2, d) \text{ goto } \tilde{l}_1;$   
 $\tilde{l}_1 : \text{mem}[e_1, d] \leftarrow e_1 \text{ goto } \tilde{l}_2;$

$\llbracket l_1 : r \leftarrow \text{mem}[e] \text{ goto } l_2 \rrbracket_{\tilde{t}_A} :=$   $\tilde{l}_1 : \text{assert } \text{mem}[e, d] = 0 \text{ goto } \tilde{l}_1;$   
 $\tilde{l}_{x1} : \text{mem}[hb] \leftarrow \text{true} \text{ goto } \tilde{l}_{x2};$   
 $\tilde{l}_{x2} : \text{mem}[e, hb] \leftarrow \text{lda} \text{ goto } \tilde{l}_{x3};$

$\llbracket l_1 : \text{mem}[e_1] \leftarrow e_2 \text{ goto } l_2 \rrbracket_{\tilde{t}_A} :=$   $\tilde{l}_1 : \text{mem}[e_1, d] \leftarrow (e_2, d) \text{ goto } \tilde{l}_2;$

$\llbracket l_1 : r \leftarrow \text{mem}[e] \text{ goto } l_2 \rrbracket_{\tilde{t}_A} :=$   $\tilde{l}_1 : \text{assert } \text{mem}[e, d] = 0 \text{ goto } \tilde{l}_{x1};$   
 $\tilde{l}_{x1} : r \leftarrow \text{mem}[e] \text{ goto } \tilde{l}_2;$   
 $\tilde{l}_1 : \text{assert } \text{mem}[e, d] \neq 0 \text{ goto } \tilde{l}_1;$   
 $\tilde{l}_{x2} : (r, d) \leftarrow \text{mem}[e, d]; \text{ goto } \tilde{l}_2;$

$\llbracket l_1: \text{local goto } l_2 \rrbracket_{\mathbb{R}_2} := \tau_1: \text{local goto } \tilde{l}_2;$   
 $\llbracket l_1: \text{infinite goto } l_2 \rrbracket_{\mathbb{R}_2} := \checkmark$

### Comment:

- Note that the instrumentation  $\llbracket \text{store} \rrbracket_{\mathbb{R}_2}$  keeps the address used in the store in a fresh address  $a_{st}$ .
- The instrumentation deletes fences as they forbid to delay  $st$  over  $ld$ .
- The instrumentation  $\llbracket \text{load} \rrbracket_{\mathbb{R}_2}$  checks the value is not read early. Moreover, it sets a happens-before address  $(a, hb)$  to access level load,  $lda$ . It also sets a flag  $hb$  to forbid helper actions that do not contribute to happens-before path  $\tau_3$ .

### Instrumenting helpers:

Idea: • How to decide whether a new action  $act$  is in happens-before relation with an earlier action  $act'$  so that  $act' \xrightarrow{hb} act$ ?

Need to know two facts:

$\hookrightarrow$  Has the thread of  $act$  already contributed an action  $act'$  to  $\tau_3$ ?

In this case,  $act' \xrightarrow{po} act$ .

The information about such a contribution can be kept in the control-flow of the helper.

$\hookrightarrow$  Does  $\tau_3$  contain a load or store access to  $addr(act)$ ?

- If there was a load  $act' = ld$ , we can add a store  $act = st$  and get  $ld \xrightarrow{hb} st$ .
- If there was a store, we are free to add a load or a store.

↳ Need one auxiliary address  $(a, hb)$   
 pu address  $a$  in the program.

The addresses  $(a, hb)$  range over the domain  
 $\{0, lba, sta\}$

of access types.

It is sufficient to store the maximal access type  
 w.r.t. the ordering:

$0$  (no access)  $<$   $lba$  (load access)  $<$   $sta$  (store access).

Technically: The augmentation on a thread's contribution to  $\tilde{\tau}_3$   
 + access types is based on the following lemma.

Lemma:

Consider  $\tilde{\tau} = \tilde{\tau}_1 \cdot act_1 \cdot \tilde{\tau}_2 \in (sc(P))$

where for all  $act_2 \in \tilde{\tau}_2$  we have  $act_1 \rightarrow_{hb}^* act_2$ .

Then  $\tilde{\tau} \cdot act$  subspes  $act_1 \rightarrow_{hb}^* act$

- $\exists!$
- (1)  $\exists act_2 \in act_1 \cdot \tilde{\tau}_2 : thread(act_2) = thread(act)$ ,
  - (2)  $act$  is a load whose address is stored in  $act_2 \cdot \tilde{\tau}_2$ , or
  - (3)  $act$  is a store (with issue) whose address is loaded or stored in  $act_2 \cdot \tilde{\tau}_2$ .

With this, the instrumentation of helpers is as follows.

Definition (Instrumentation of helpers):

Consider thread  $t$  regs  $r^*$  init to begin  $\langle linst \rangle^*$  end.

The instrumentation of  $t$  is

$\llbracket t \rrbracket :=$  thread  $\tilde{t}$  regs  $\tilde{r}, r^*$  init to

begin  
 $\llbracket \langle linst \rangle \rrbracket_{H_0}^* \llbracket \langle ldstinst \rangle \rrbracket_{H_2}^* \llbracket \langle linst \rangle \rrbracket_{H_2}^* \llbracket \langle linst \rangle \rrbracket_{H_3}^*$

end

• Here,  $\langle \text{ldstinst} \rangle^*$  is the subsequence of all load and store instructions.

The instrumentation  $\llbracket \langle \text{ldstinst} \rangle \rrbracket_{H_2}^*$  is used to move to the code copy  $\llbracket \langle \text{list} \rangle \rrbracket_{H_2}^*$

• Let  $\langle l \rangle^*$  be all labels used by the thread.

The instructions  $\llbracket \langle l \rangle \rrbracket_{H_3}^*$  raise a success flag when a TSO witness has been found.

• The instrumentation  $\llbracket \langle \text{list} \rangle \rrbracket_{H_0}^*$  of the original source code forces the helper to either enter the code copy  $\llbracket \langle \text{list} \rangle \rrbracket_{H_2}^*$  or stop when the hb-flag is raised.

The functions are as follows:

$\llbracket l_1: \text{instr goto } l_2 \rrbracket_{H_0} := l_1: \text{assert mem}[hb] = 0 \text{ goto } l_x;$   
 $l_x: \text{instr goto } l_2;$

$\llbracket l_1: r \leftarrow \text{mem}[e] \text{ goto } l_2 \rrbracket_{H_2} := l_1: \text{assert mem}[(e, hb)] = \text{sta goto } \tilde{l}_x;$   
 $\tilde{l}_x: r \leftarrow \text{mem}[e] \text{ goto } l_2;$

$\llbracket l_1: \text{mem}[e_1] \leftarrow e_2 \text{ goto } l_2 \rrbracket_{H_1} := l_1: \text{assert mem}[(e_1, hb)] \neq \text{lda goto } \tilde{l}_{x_1};$   
 $\tilde{l}_{x_1}: \text{mem}[e_1] \leftarrow e_2 \text{ goto } \tilde{l}_{x_2};$   
 $\tilde{l}_{x_2}: \text{mem}[(e_1, hb)] \leftarrow \text{sta goto } l_2;$

$\llbracket l_1: \text{local / fence goto } l_2 \rrbracket_{H_2} := \tilde{l}_1: \text{local / fence goto } l_2;$

$\llbracket l_1: \text{mem}[e_1] \leftarrow e_2 \text{ goto } l_2 \rrbracket_{H_2} := \tilde{l}_1: \text{mem}[e_1] \leftarrow e_2 \text{ goto } \tilde{l}_x;$   
 $\tilde{l}_x: \text{mem}[(e_1, hb)] \leftarrow \text{sta goto } l_2;$

$\llbracket l_1: r \leftarrow \text{mem}[e] \text{ goto } l_2 \rrbracket_{H_2} := \tilde{l}_1: \tilde{r} \leftarrow e \text{ goto } \tilde{l}_{x_1};$   
 $\tilde{l}_{x_1}: r \leftarrow \text{mem}[\tilde{r}] \text{ goto } \tilde{l}_{x_2};$   
 $\tilde{l}_{x_2}: \text{mem}[(\tilde{r}, hb)] \leftarrow \text{max}\{\text{lda}, \text{mem}[(\tilde{r}, hb)]\} \text{ goto } l_2;$

$$\begin{aligned} \llbracket l \rrbracket_{H_3} := & \tilde{r} : \tilde{r} \leftarrow \text{mem}[a_{st_H}] \text{ goto } \tilde{r}_{x_1}; \\ & \tilde{r}_{x_1} : \tilde{r} \leftarrow \text{mem}[(\tilde{r}, hb)] \text{ goto } \tilde{r}_{x_2}; \\ & \tilde{r}_{x_2} : \text{assert } \tilde{r} \neq 0 \text{ goto } \tilde{r}_{x_3}; \\ & \tilde{r}_{x_3} : \text{mem}[suc] \leftarrow \text{true} \text{ goto } \tilde{r}_{x_4}; \end{aligned}$$

Note:

In the instrumentation of loads,  $\llbracket l_1 : r \leftarrow \text{mem}[c] \text{ goto } l_2 \rrbracket_{H_2}$ , auxiliary register  $\tilde{r}$  ensures that we do not overwrite the addresses given by  $c$  when modifying  $r$  (may be used within  $c$ ).

Theorem (Soundness and completeness of instrumentation):

$\exists$  Hack  $\tilde{r}$  is feasible in program  $P$   
 iff  $P_{\tilde{r}}$  reaches a goal configuration under SC.

$\exists$  goal configuration is a pair  $(pc, val)$   
 with  $val(suc) = \text{true}$ .

Theorem:

- Program  $P$  is robust iff no instrumentation  $P_{\tilde{r}}$  reaches a goal configuration under SC
- If the data domain is finite and given as input, robustness is PSPITLE-complete.

Proof:

Upper bound: We show that the complement of robustness, the non-robustness problem (given a program  $P$ , check that  $P$  is not robust) can be solved in non-deterministic polynomial space (NPSPACE).  
 By Savitch's theorem NPSPACE = PSPACE,  
 and hence non-robustness  $\in$  PSPACE

we negate the answer and get robustness  $\in$  PSPACE.

Essentially, we use that

$$\text{co-NPSPACE} = \text{co-PSPACE} = \text{PSPACE} (= \text{NPSPACE}).$$

To solve non-robustness in NPSPACE,

we guess a suitable attack  $A$

and compute the linear-site instrumentation  $P_A$ .

Then we guess a reaching path in  $P_A$ .

For the path, we only need to store

- the current configuration (works in linear space)
- the number of steps taken (works in linear space as well).

Return yes, if a goal configuration is reached.

Return no, if the search for a path deadlocks

or the number of steps exceeds the number of configurations

For the latter, note that

there are  $2^n$  configurations with  $n$  bits.

We need another  $n$ -bits to count to  $2^n$ .

Lower bound: We first give a reduction of SC-reachability

in Boolean programs to non-robustness.

• Consider a single-threaded Boolean program  $P$  with control-location  $l$



• Using a second thread and fresh variables, we append a Dekker cycle to  $l$ :

• Then  $l$  is SC-reachable in  $P$

iff  $P'$  is not robust.

• Since control-state reachability

is PSPACE-hard, so is non-robustness.



To see that also robustness is PSPACE-hard,  
consider a problem  $Prob \in PSPACE$   
that we want to reduce to robustness.

Since  $PSPACE$  is closed under complement,  
we have

$$co-Prob \in PSPACE.$$

We just showed that non-robustness is PSPACE-hard.

Hence there is a reduction

$$f: co-Prob \rightarrow non-Prob$$

so that

$$instance i \in co-Prob \text{ iff } f(i) \text{ is not robust.}$$

Since  $i \in co-Prob \text{ iff } i \notin Prob$ ,

we have

$$i \in Prob \text{ iff } f(i) \text{ is robust.}$$

So function  $f$  is also a reduction of  $Prob$  to robustness.  $\square$