

10.4 Instrumentation

Goal: Consider program P with abstract $\mathcal{A} = (t_{\mathcal{A}}, \text{start}, \text{end})$.
Characterise ISO 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.

Why: • Unbounded store bytes
• 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 load.
- ↳ Indeed, in $\bar{C}_{\mathcal{A}}$ the helpers are no longer active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 $\bar{t}_A = (t_A, \text{stinst}, \text{ldinst})$ be the attack.

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

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

begin

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

// source code

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

// Move to copy of source code

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

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

// copy of source code.

end.

with

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

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

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

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

$\llbracket l_1: \text{local goto } l_2 \rrbracket_{\mathcal{H}_2} := \tilde{\tau}_1: \text{local goto } \tilde{l}_2;$
 $\llbracket l_1: \text{inference goto } l_2 \rrbracket_{\mathcal{H}_2} := \checkmark$

Comment:

- Note that the instrumentation $\llbracket \text{store} \rrbracket_{\mathcal{H}_2}$ 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 $\llbracket \text{load} \rrbracket_{\mathcal{H}_2}$ 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 $\tilde{\tau}_3$.

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:

↳ Has the thread of act already contributed an action act' to $\tilde{\tau}_3$?

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

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

↳ Does $\tilde{\tau}_3$ 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)
po 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
wrt. 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
 $\llbracket \langle linst \rangle \rrbracket_{H_0}^* \llbracket \langle ldstinst \rangle \rrbracket_{H_2}^* \llbracket \langle linst \rangle \rrbracket_{H_2}^* \llbracket \langle l \rangle \rrbracket_{H_3}^*$

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:

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

$$l_x: \text{instr } \underline{\text{goto}} \ l_2;$$

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

$$\tilde{l}_x: r \leftarrow \text{mem}[e] \ \underline{\text{goto}} \ l_2;$$

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

$$\tilde{l}_{x_1}: \text{mem}[e_1] \leftarrow e_2 \ \underline{\text{goto}} \ \tilde{l}_{x_2};$$

$$\tilde{l}_{x_2}: \text{mem}[(e_1, hb)] \leftarrow \text{sta } \underline{\text{goto}} \ l_2;$$

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

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

$$\tilde{l}_x: \text{mem}[(e_1, hb)] \leftarrow \text{sta } \underline{\text{goto}} \ l_2;$$

$$\llbracket l_1: r \leftarrow \text{mem}[e] \ \underline{\text{goto}} \ l_2 \rrbracket_{H_2} := \tilde{l}_1: \tilde{r} \leftarrow e \ \underline{\text{goto}} \ \tilde{l}_{x_1};$$

$$\tilde{l}_{x_1}: r \leftarrow \text{mem}[\tilde{r}] \ \underline{\text{goto}} \ \tilde{l}_{x_2};$$

$$\tilde{l}_{x_2}: \text{mem}[(\tilde{r}, hb)] \leftarrow \max\{\text{lda}, \text{mem}[(\tilde{r}, hb)]\} \ \underline{\text{goto}} \ l_2;$$

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

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 $i \in co-Prob$ 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