

Capability Machines

The future of secure processors

Aina Linn Georges - Friday 24 February

Who am I?

PhD student in the Logic and Semantics group

Mechanised Reasoning about Capability Machines

```
lemma wp_invoke_native (s : stuckness) (E : coPset) (ϕ : val → iProp Z) (ves vcs t1s t2s ts a es i m f0 :
  iris.to_val ves = Some (immV vcs) →
  length vcs = length t1s →
  length t2s = m →
  ⌊frame⌋ f0 →
  (N.of_nat a) → [wf] (FC_func_native i (Tf t1s t2s) ts es) →
  ▷ (⌊frame⌋ f0 × (N.of_nat a) → [wf] (FC_func_native i (Tf t1s t2s) ts es) →
  WP [AI_local m (Build_frame (vcs ++ (n_zeros ts)) i) [AI_basic (BI_block (Tf [] t2s) es)]] @ s; E { (v, ϕ v × ⌊frame⌋ f0) }) →
  WP ves ++ [AI_invoke a] @ s; E { (v, ϕ v × ⌊frame⌋ f0) }) .
Proof.
iIntros (Hparams Hlen Hret) "Hf H1 H0".
iApply wp_lift_step.
{ apply to_val_cat_None2. auto. }
iIntros ([[[? ?] ?] ?] ns k ks nt) "(H01H02H03H04H05H06)".
iApply fupd_frame.L.
iDestruct (gen_heap_valid with "H01 H1") as Hlook.
set (σ := (sg, s1, l, i0)).
assert (reduce (host_instance:=host_instance) sg s1 { f_locs := l; f_inst := i0 })
  (ves ++ [AI_invoke a])%list sg s1 { f_locs := l; f_inst := i0 })
  { AI_local m { f_locs := vcs ++ n_zeros ts; f_inst := i } [AI_basic (BI_block (Tf [] t2s) es)] as Hred.
  eapply r_invoke_native with (ts:=ts); eauto.
  { rewrite gmap_of_list_lookup Nat2N.id in Hlook. rewrite /=- nth_error_lookup //. }
  { symmetry. apply v_to_e_list_to_val. auto. }
iSplit.
= iPureIntro.
  destruct s = //.
  unfold language.reduceable, language.prim_step = /.
  eexists [], _, σ, [].
  unfold iris.prim_step = /.
  repeat split = //.
= iApply fupd_mask_intro; solve_ndisj.
iIntros "Hcls" (es1 es2 efs HStep).
iMod "Hcls". iModIntro.
  destruct es as [[hs' ws'] locs' inst'].
  destruct HStep as (H & → & ←).
  assert (first_instr (ves ++ [AI_invoke a]) = Some (AI_invoke a)) as Hf.
  { apply first_instr_const. eapply to_val_const_list. eauto. }
  eapply reduce_det in H as HH; [[apply Hred].
  destruct HH as [H1 | Hstart | [((?&?&?&?&?&?&?&?) | (Hstart & Hstart1 & Hstart2 & H0) ]]]; try done.
  simplify_eq. iApply bi.sep_exist.L. iExists f0. iFrame.
  iSplit => //. iIntros "Hf".
  iSpecialize ("H0" with "[s]"). iFrame.
  rewrite Hf in Hstart. done.
  rewrite Hf in H0. simplify_eq. rewrite /=- nth_error_lookup // in H1.
  rewrite gmap_of_list_lookup Nat2N.id in Hlook.
  congruence.
Qed.

(* ----- Host invocations ----- *)
(* ----- *)
(* ----- *)

(* TODO0: discuss what we want host invocations, in case the host invokes a wasm function *)

(* In the version below, host invocations are handled with an imported host program logic,
  as described in iris_wp_def. *)
```

```
iInstance
m : N
f0 : frame
Hparams : to_val ves = Some (immV vcs)
Hlen : length vcs = length t1s
Hret : length t2s = m
sg : host_state host_instance
s1 : store_record
l : seq.seq value
i0 : instance
ns : N
ks : seq.seq (language.observation iris_wp_def.wasm_lang)
nt : N
Hlook : gmap_of_list (s_funcs s1) !! N.of_nat a =
  Some (FC_func_native i (Tf t1s t2s) ts es)
σ := (sg, s1, l, i0)
: host_state host_instance * store_record * seq.seq value * instance
Hred : reduce (host_instance:=host_instance) sg s1
  { f_locs := l; f_inst := i0 } (ves ++ [AI_invoke a])%list
  { f_locs := l; f_inst := i0 }
  [AI_local m { f_locs := vcs ++ n_zeros ts; f_inst := i }
  [AI_basic (BI_block (Tf [] t2s) es)]]
es1 : language.expr iris_wp_def.wasm_lang
hs' : host_state host_instance
ws' : store_record
locs' : seq.seq value
inst' : instance
efs : seq.seq (language.expr iris_wp_def.wasm_lang)
HStep : language.prim_step (ves ++ [AI_invoke a]) σ k es1
  (hs', ws', locs', inst') efs

"Hf" : ⌊frame⌋ f0
"H1" : N.of_nat a → [wf] FC_func_native i (Tf t1s t2s) ts es
"H0" : ⌊frame⌋ f0 × N.of_nat a → [wf] FC_func_native i (Tf t1s t2s) ts es
  WP [AI_local m { f_locs := vcs ++ n_zeros ts; f_inst := i }
  [AI_basic (BI_block (Tf [] t2s) es)]]
  @ s; E
  { (v, ϕ v × ⌊frame⌋ f0) }
"H01" : gen_heap_interp (gmap_of_list (s_funcs s1))
"H02" : gen_heap_interp (gmap_of_table (s_tables s1))
"H03" : gen_heap_interp (gmap_of_memory (s_mems s1))
"H04" : gen_heap_interp (gmap_of_list (s_globals s1))
"H05" : ghost_map_auth frameName 1
  (<[[]] := [ f_locs := l; f_inst := i0 ] >> σ)
"H06" : gen_heap_interp (gmap_of_list (mem_length <=> s_mems s1))
"Hcls" : emp
```

$$\begin{aligned} \mathcal{E}(W)(v) &\triangleq \forall \text{reg}. \mathcal{R}(W)(\text{reg}) * \text{sharedResources}(W) * \text{stsCollection}(W) * \text{pc} \mapsto v \\ &\quad * *_{(r,w) \in \text{reg}/\text{pc}} r \mapsto w \text{ ---} \\ \text{wp Executable} &\begin{cases} v, v = \text{Halted} \rightarrow \exists W' \text{ reg}', W' \sqsupseteq^{\text{prio}} W \\ * \text{sharedResources}(W') * \text{stsCollection}(W') \\ * *_{(r,w) \in \text{reg}} r \mapsto w \end{cases} \\ \mathcal{R}(W)(\text{reg}) &\triangleq *_{(r,w) \in \text{reg}/\text{pc}} \mathcal{V}(W)(w) \\ \mathcal{V}(W)(z) &\triangleq \top \\ \mathcal{V}(W)(o, -) &\triangleq \top \\ \mathcal{V}(W)(p, g, b, e, a) &\triangleq *_{a' \in [b, e]} \begin{cases} \mathcal{S}^u(W)(a', g, p, a) & \text{if } p = u - \\ \mathcal{S}(W)(a', g, p) & \text{otherwise} \end{cases} \\ &\quad \wedge \begin{cases} \exists P, \text{rel}(a', P) * \text{rcond}(P) & \text{if } p \in \{\text{RO}, \text{RX}\} \\ \text{rel}(a', \mathcal{V}) & \text{otherwise} \end{cases} \\ \mathcal{V}(W)(\epsilon, g, b, e, a) &\triangleq \square \forall W' \sqsupseteq^g W, \triangleright \mathcal{E}(W')(RX, g, b, e, a) \\ \text{rcond}(P) &\triangleq \triangleright \square \forall W, w, P(W)(w) \text{ ---} *_{\triangleright} \mathcal{V}(W)(w) \\ &\quad * \triangleright \square \forall W_1, W_2, z, P(W_1)(z) \text{ ---} *_{\triangleright} P(W_2)(z) \end{aligned}$$

State relation

$$\mathcal{S}(W)(a, g, p) \triangleq \begin{cases} W^{\text{std}}(a) \in \{\text{Temporary}, \text{Permanent}\} & \text{if } \neg \text{write-local}(p) \wedge g = \text{DIRECTED} \\ W^{\text{std}}(a) = \text{Temporary} & \text{if } \text{write-local}(p) \wedge g = \text{DIRECTED} \\ W^{\text{std}}(a) = \text{Permanent} & \text{if } g \neq \text{DIRECTED} \end{cases}$$
$$\mathcal{S}^u(W)(a, g, p, mid) \triangleq \begin{cases} \mathcal{S}(W)(a, g, p) \vee \exists w, W^{\text{std}}(a) = \text{Uninitialized}(w) & \text{if } a \geq \text{mid} \\ \wedge g = \text{DIRECTED} & \\ \mathcal{S}(W)(a, g, p) & \text{if } a < \text{mid} \\ \vee g \neq \text{DIRECTED} & \end{cases}$$

Fig. 7. A Logical Relation with Support for Temporal Stack Safety. \sqsupseteq^g equals $\sqsupseteq^{\text{prio}}$ whenever g is GLOBAL or LOCAL, and \sqsupseteq^f whenever g is DIRECTED

can be used to access the ghost state of a , while guaranteeing that ϕ holds at the current physical state of a in the current world W . Normally, the predicate we associate with such an address a is \mathcal{V} . However, we distinguish between a read-only and a read-write permission by the predicate of an address within a read-only region to be stronger than \mathcal{V} .



Capability Machines

A capability at the lowest level of a machine:

Hardware capabilities

- Motivation and Design Principles
- Hardware Capability Operations
- Beyond Memory Safety

History

The concept of machine capabilities is old!

[Dennis & Van Horn - 1966]

Programming Semantics for Multiprogrammed Computations

Jack B. Dennis and Earl C. Van Horn
Massachusetts Institute of Technology, Cambridge, Massachusetts

The semantics are defined for a number of meta-instructions which perform operations essential to the writing of programs in multiprogrammed computer systems. These meta-instructions relate to parallel processing, protection of separate computations, program debugging, and the sharing among users of memory segments and other computing objects, the names of which are hierarchically structured. The language sophistication contemplated is midway between an assembly language and an advanced algebraic language.

Introduction

An increasing percentage of computation activity will be carried out by multiprogrammed computer systems. Such systems are characterized by the application of computation resources (processing capacity, main memory, file storage, peripheral equipment) to many separate but concurrently operating computations.

We can cite three quite different examples of multiprogrammed computer systems to illustrate their diversity of application. The American Airlines SABRE passenger record system couples ticketing agents at dispersed offices to a central data file [1]. The computer support systems of NASA provide real time control and monitoring of manned space flights [2]. The Project MAC time-sharing system permits research workers closer interaction with the powers of automatic computation [3]. Although these are all on-

Presented at an ACM Programming Languages and Pragmatics Conference, San Dimas, California, August 1965.

Work reported herein was supported by Project MAC, an MIT research program sponsored by the Advanced Research Projects Agency, Department of Defense, under Office of Naval Research Contract Number Nonr-4102(01). Reproduction in whole or in part is permitted for any purpose of the United States Govern-

CHERI: Capability Hardware Enhanced RISC Instructions



[Woodruff et.al. - 2014]

[Watson et.al. - 2015]

The industry's growing interest in capability machines or: "why this is an excellent time to be excited about capabilities"

Microsoft

2019




Figure 1: ~70% of the vulnerabilities Microsoft assigns a CVE each year continue to be memory safety issues

2020

We conservatively assessed the percentage of vulnerabilities reported to the Microsoft Security Response Center (MSRC) in 2019 and found that approximately 31% would no longer pose a risk to customers and therefore would not require addressing through a security update on a CHERI system based on the default configuration of the CheriBSD operating system. If we also assume that automatic initialization of stack variables ([InitAll](#)) and of heap allocations (e.g. [pool zeroing](#)) is present, the total number of vulnerabilities deterministically mitigated exceeds 43%. With additional features such as [Cornucopia](#) that help prevent temporal safety issues such as use after free, and assuming that it would cover 80% of all the UAFs, **the number of deterministically mitigated vulnerabilities would be at least 67%**. There is additional work that needs to be done to protect the stack and add fine grained CFI, but this combination means CHERI looks very promising in its early stages.

Google - Chromium

2020

Around 70% of our high severity security bugs are memory unsafety problems (that is, mistakes with C/C++ pointers). Half of those are use-after-free bugs.


High+, impacting stable

Security-related assert
7.1%

Other
23.9%

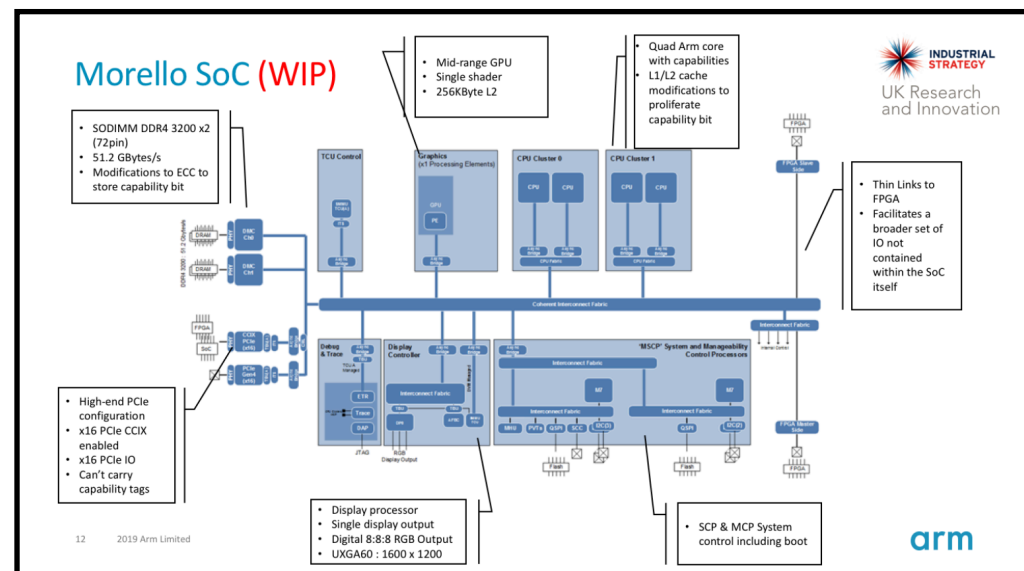
Other memory unsafety
32.9%

Use-after-free
36.1%



(Analysis based on 912 high or critical [severity](#) security bugs since 2015, affecting the Stable channel.)

arm : Morello



arm Arm[®] Architecture Reference Manual Supplement Morello for A-profile Architecture

Document number DDI0606
 Document version A.k
 Document confidentiality Non-confidential

Copyright © 2019-2022 Arm Limited or its affiliates. All rights reserved.

Important message

Morello is a prototype architecture, which has a particular meaning to Arm of which the recipient must be aware as follows:

- Subject to change without consent of all parties, and it is not committed for product development.
- Includes the majority of expected features.
- Includes detail on the majority of expected features.
- Includes some necessary information from documentation relating to earlier architectures, but some cross-referencing might be necessary.
- See the architecture release notes for more detail.
- No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document unless specifically stated.

September 2020

Verified security for the Morello capability-enhanced prototype Arm architecture

- THOMAS BAUEREISS, University of Cambridge, UK
- BRIAN CAMPBELL, University of Edinburgh, UK
- THOMAS SEWELL, University of Cambridge, UK
- ALASDAIR ARMSTRONG, University of Cambridge, UK
- LAWRENCE ESSWOOD, University of Cambridge, UK
- IAN STARK, University of Edinburgh, UK
- GRAEME BARNES, Arm Ltd., UK
- ROBERT N. M. WATSON, University of Cambridge, UK
- PETER SEWELL, University of Cambridge, UK

Memory safety bugs continue to be a major source of security vulnerabilities in our critical infrastructure. The CHERI project has proposed extending conventional architectures with hardware-supported *capabilities* to enable fine-grained memory protection and scalable compartmentalisation, allowing historically memory-unsafe C and C++ to be adapted to deterministically mitigate large classes of vulnerabilities, while requiring only minor changes to existing system software sources. Arm is currently designing and building Morello, a CHERI-enabled prototype architecture, processor, SoC, and board, extending the high-performance Neoverse N1, to enable industrial evaluation of CHERI and pave the way for potential mass-market adoption. However, for such a major new security-oriented architecture feature, it is important to establish high confidence that it does provide the protections it intends to, and that cannot be done with conventional engineering techniques.

In this paper we put the Morello architecture on a solid mathematical footing from the outset. We define the fundamental security property that Morello aims to provide, reachable capability monotonicity, and prove that the architecture definition satisfies it. This proof is mechanised in Isabelle/HOL, and applies to a translation of the official Arm Morello specification into Isabelle. The main challenge is handling the complexity and scale of a production architecture: 62,000 lines of specification, translated to 210,000 lines of Isabelle. We do so by factoring the proof via a narrow abstraction capturing the essential properties of instruction execution in an arbitrary CHERI ISA, expressed above a monadic intra-instruction semantics. We also develop a model-based test generator, which generates instruction-sequence tests that give good specification coverage, used in early testing of the Morello implementation and in Morello QEMU development. We also use Arm's internal test suite to validate our internal model.

This gives us machine-checked mathematical proofs of whole-ISA security properties of a full-scale industry architecture, at design-time. To the best of our knowledge, this is the first demonstration that that is feasible, and it significantly increases confidence in Morello.

2022

“Memory Safety”

Design Principles behind CHERI

Spatial memory safety:

- Absence of certain undefined behaviours via memory accesses

```
int arr[4] = {0,1,2,3};  
int *p = arr + 5;  
p = 0;  
int a = *p;
```

undefined behaviour for indexing out of bounds

undefined behaviour for dereferencing a null pointer

undefined behaviour of dereferencing an arbitrary cast

```
long f(long i){  
    int* p = (int *)i;  
    *p = 4;  
    return ((long) p);  
}
```


Design Principles behind CHERI

Spacial memory safety: **pointers as capabilities principle**

- Absence of the kinds undefined behaviours via memory accesses that pose a risk, i.e. an exploitable memory vulnerability
 - Dereferencing out of bounds
 - Dereferencing arbitrary pointer casts

What about temporal memory safety?

- use-after-free

interesting topic, but outside the scope of this lecture! (see Cornucopia [Filardo et.al. 2020])

Hardware Capabilities

A memory pointer should not appear out of thin air (no casting!)

unforgeable token of authority

```
graph TD; A["unforgeable token of authority"] --> B["A memory pointer should not appear out of thin air (no casting!)"]; A --> C["A memory pointer should grant access to a specific segment of memory (no out of bounds!)"]; A --> D["If you want to access a memory access, you need to prove you have the authority to do so (no confused deputy)"];
```

A memory pointer should grant access to a specific segment of memory (no out of bounds!)

If you want to access a memory access, you need to prove you have the authority to do so (no confused deputy)

Hardware Capabilities


unforgeable token of authority

Authority over memory cannot be artificially increased or created

Hardware Capabilities


Pointers as Capabilities

Pointers are replaced by hardware capabilities



- Bounds of authority
- Permission: RO/RW/etc
- Address

A Lattice of Permissions




- RO : read-only
- RW : read-write

What about the program counter?

- RX : read execute
- RWX : read-write execute

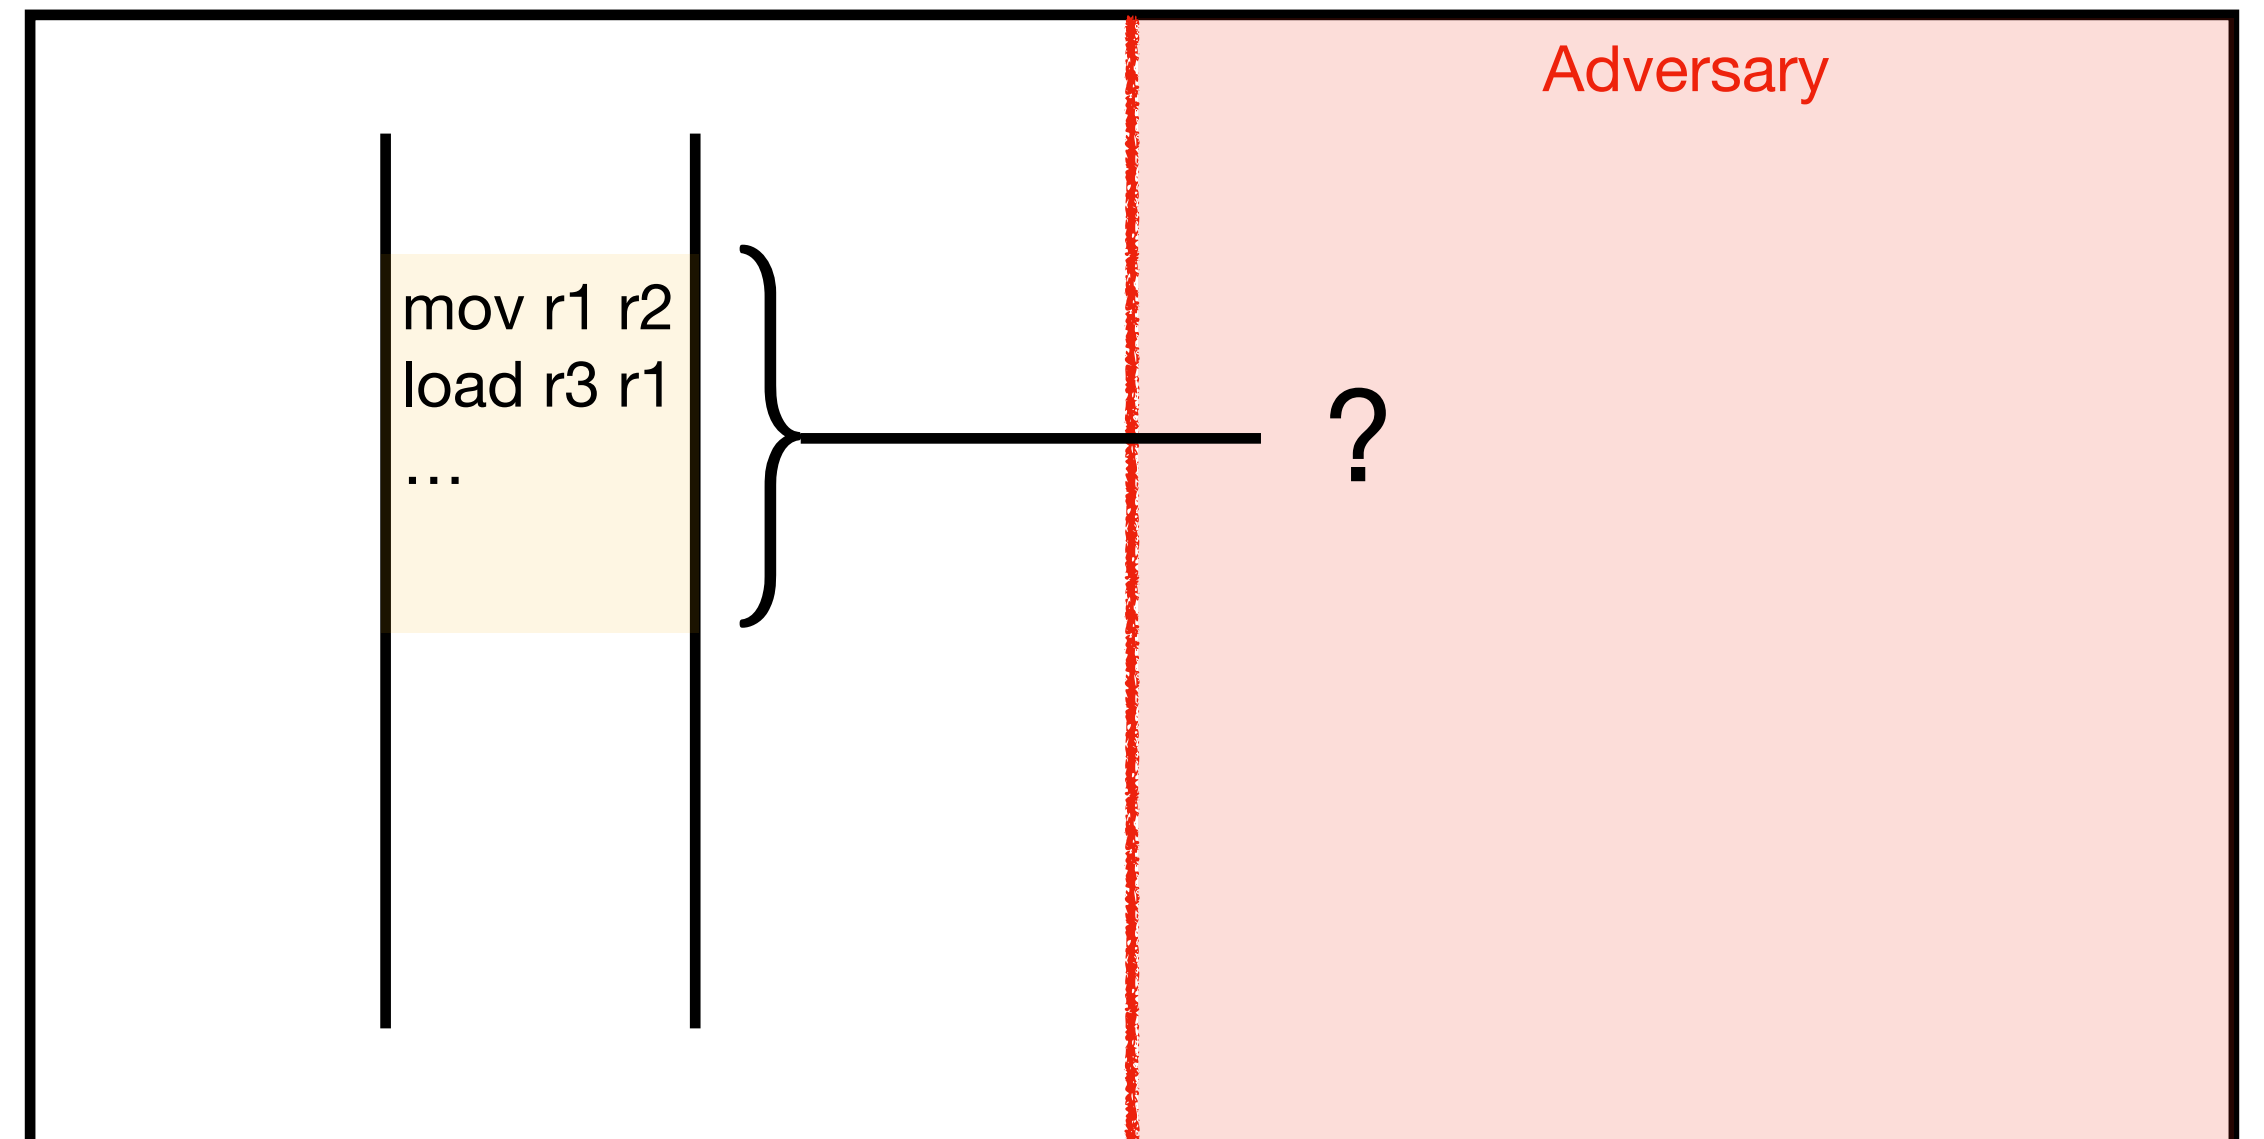
A Lattice of Permissions



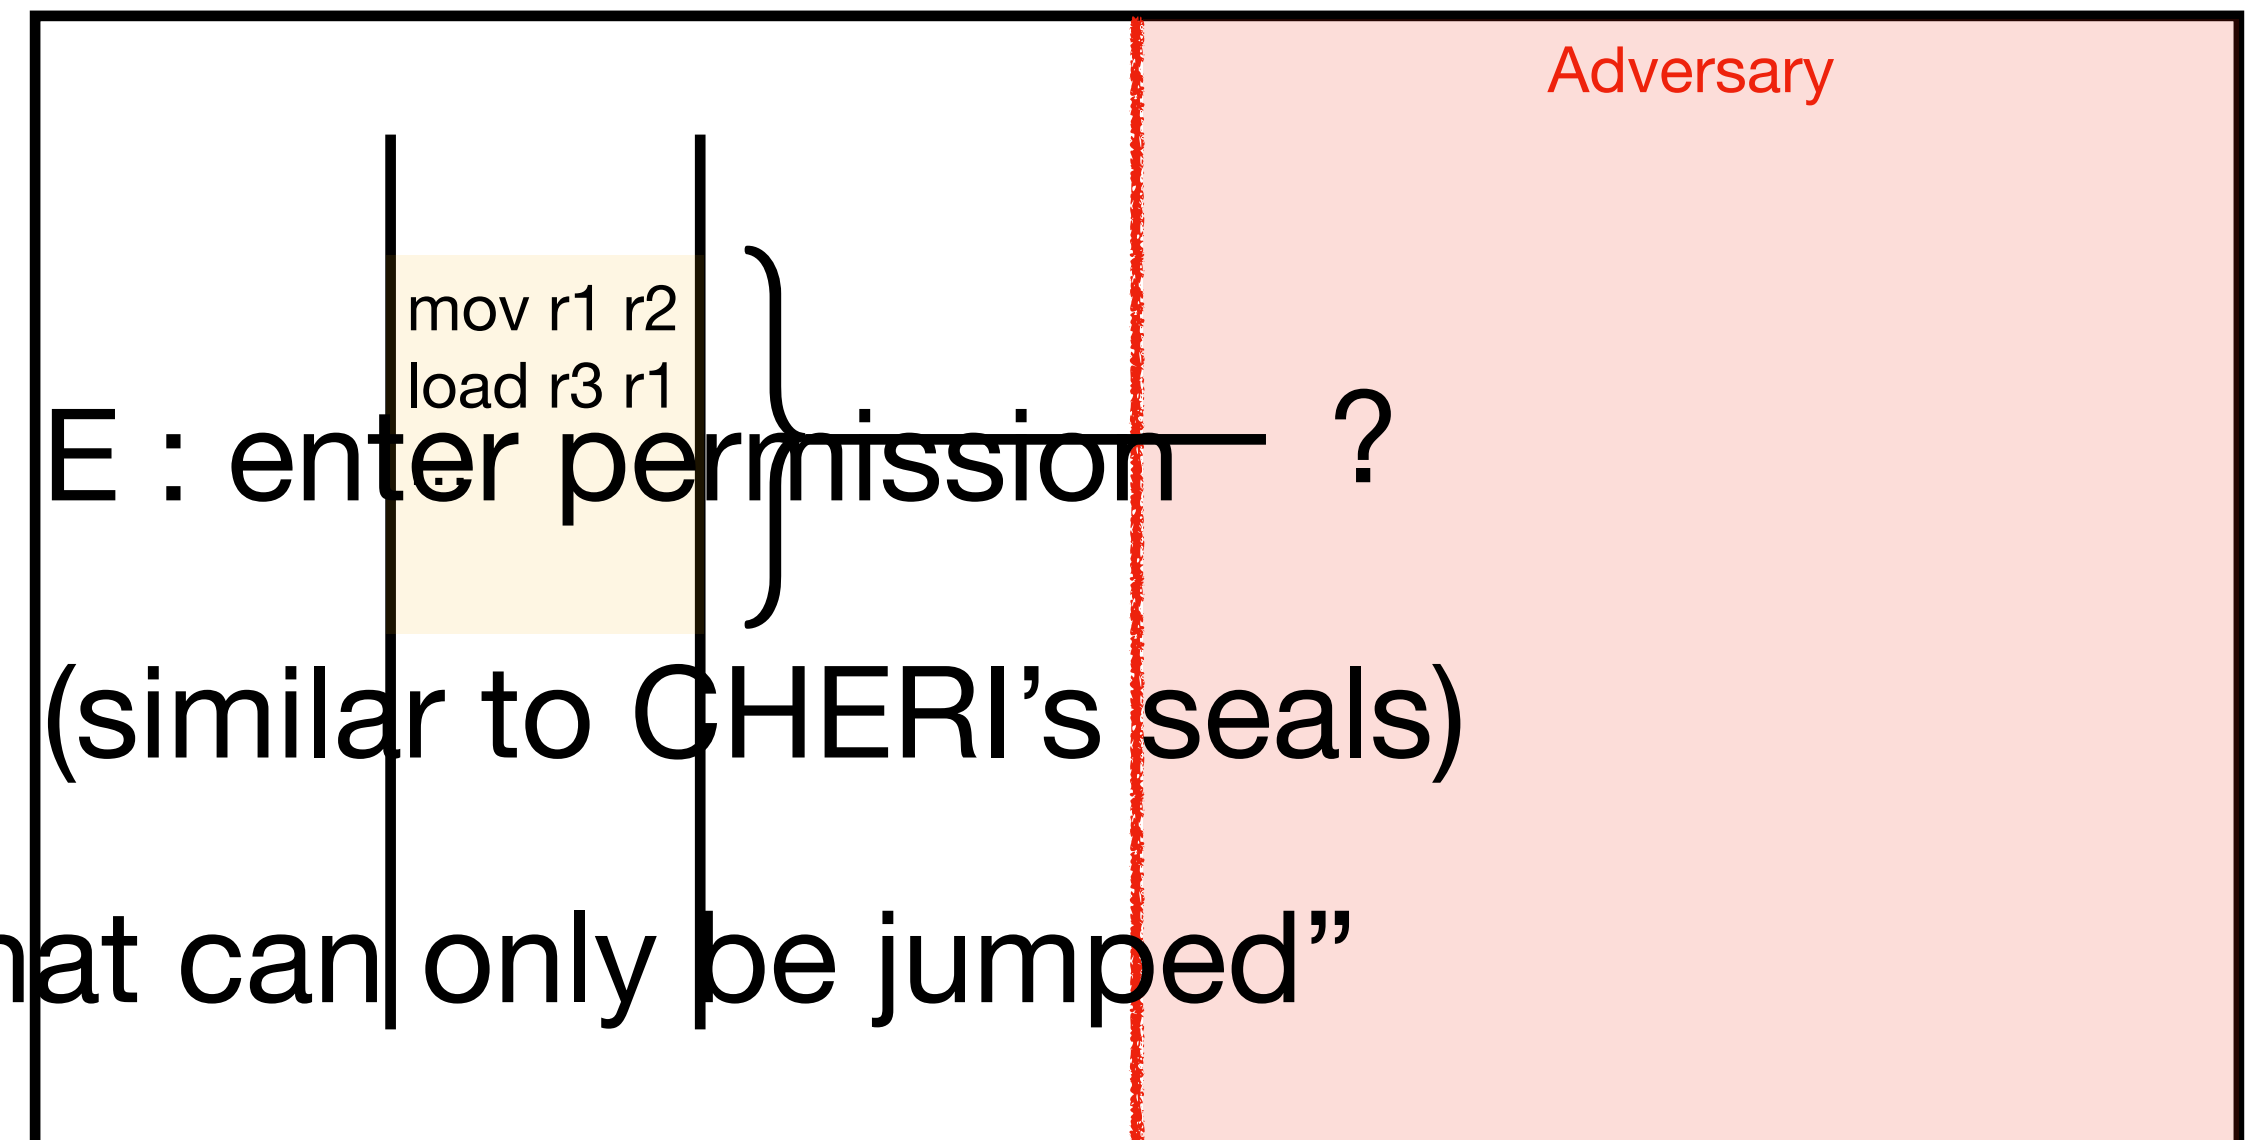

What about function pointers?

What authority does a function pointer grant?

Attacker model: linking with arbitrary capability machine code



A Lattice of Permissions



We need a permission for “a capability that can only be jumped”

Operations

Operations

State machine

$$(reg, mem) \longrightarrow (reg', mem')$$

$reg : \text{RegName} \rightarrow \text{Word}$
 $mem : \text{Addr} \rightarrow \text{Word}$

$\text{RegName} = \{\text{PC}, r1, r2, r3\}$

$\text{Addr} = [0, \text{AddrMax})$

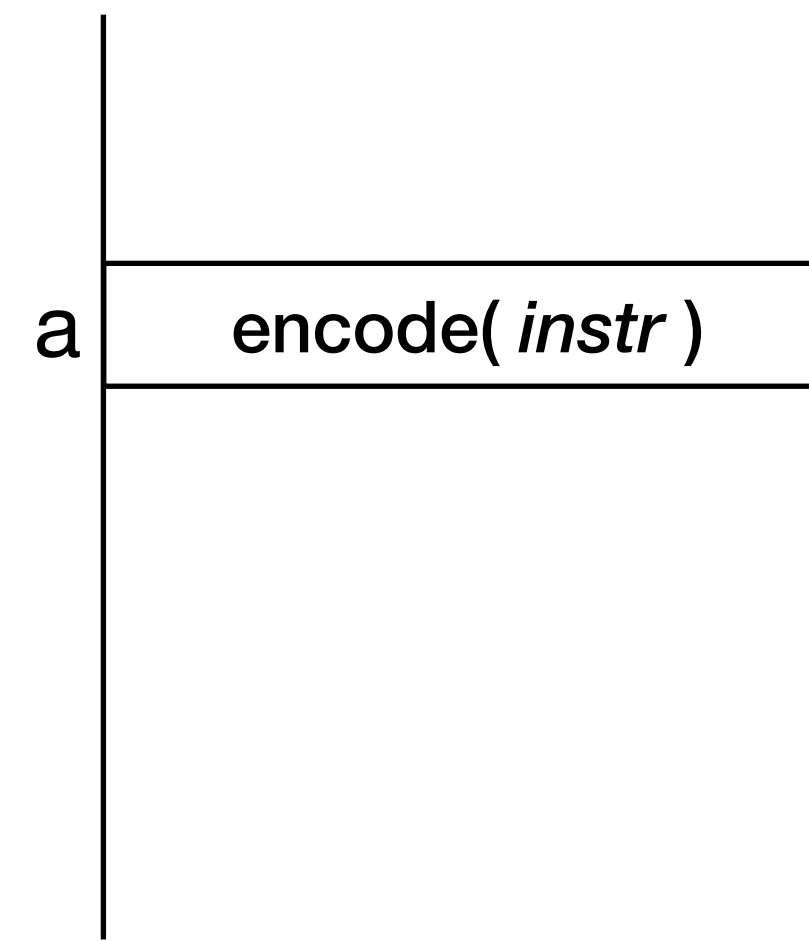

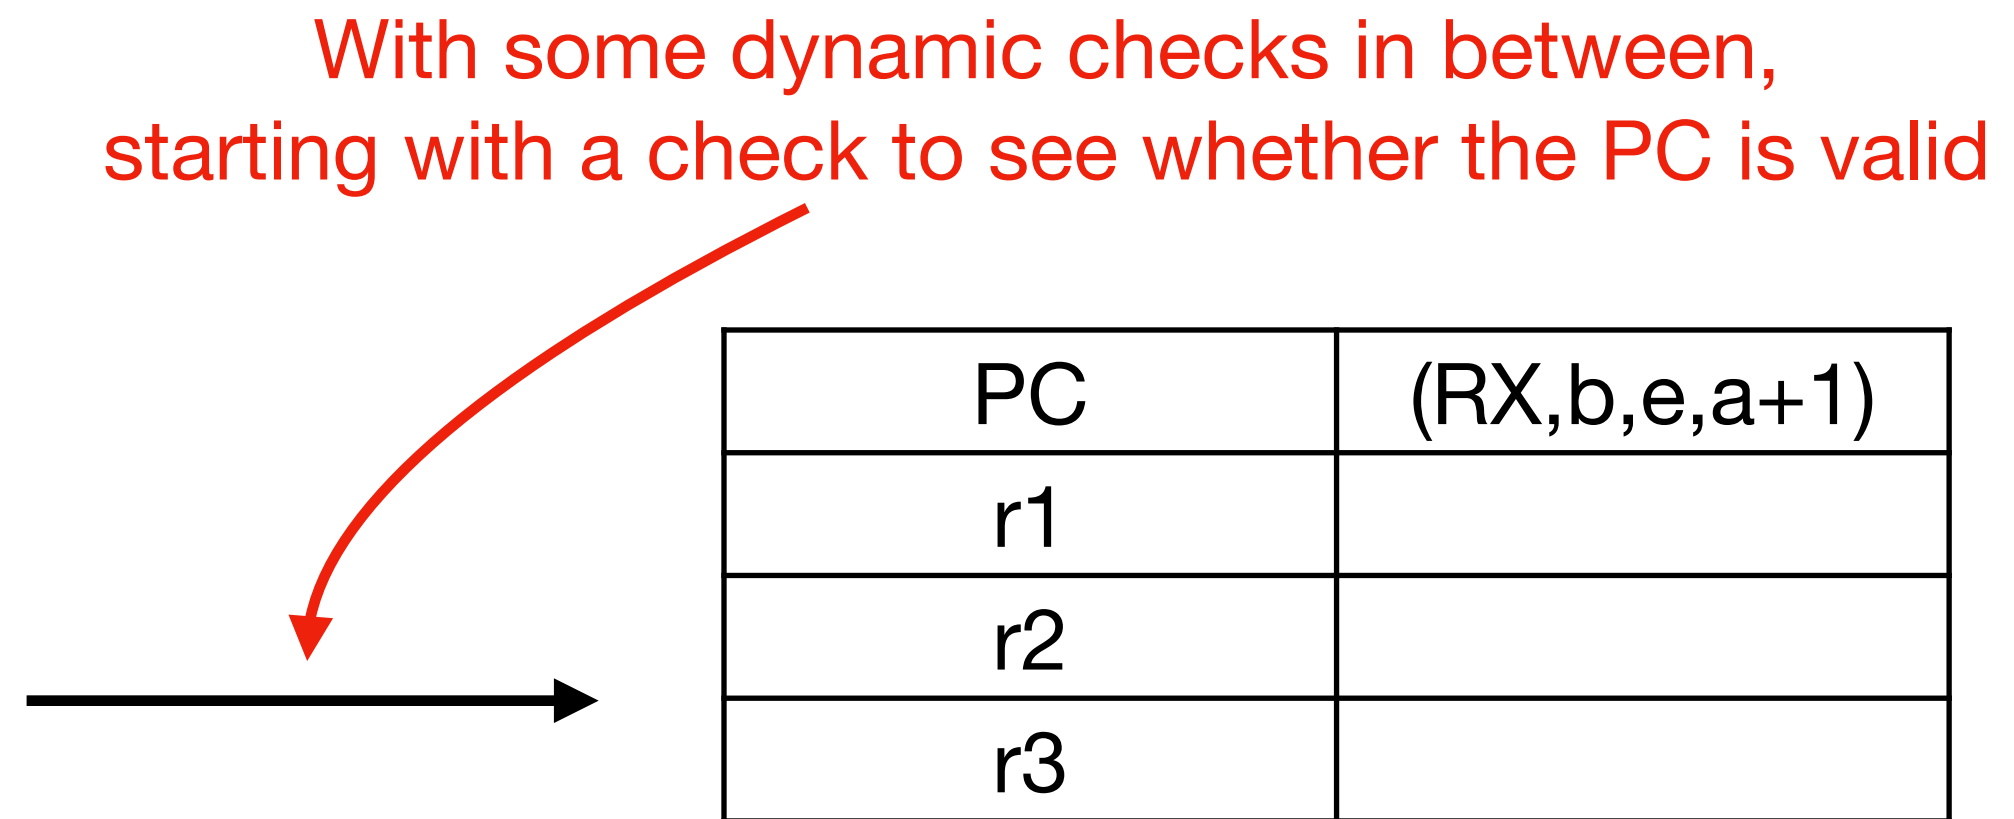
a word is either a
capability or an integer



Operations

General shape of an instruction

PC	(RX,b,e,a)
r1	
r2	
r3	



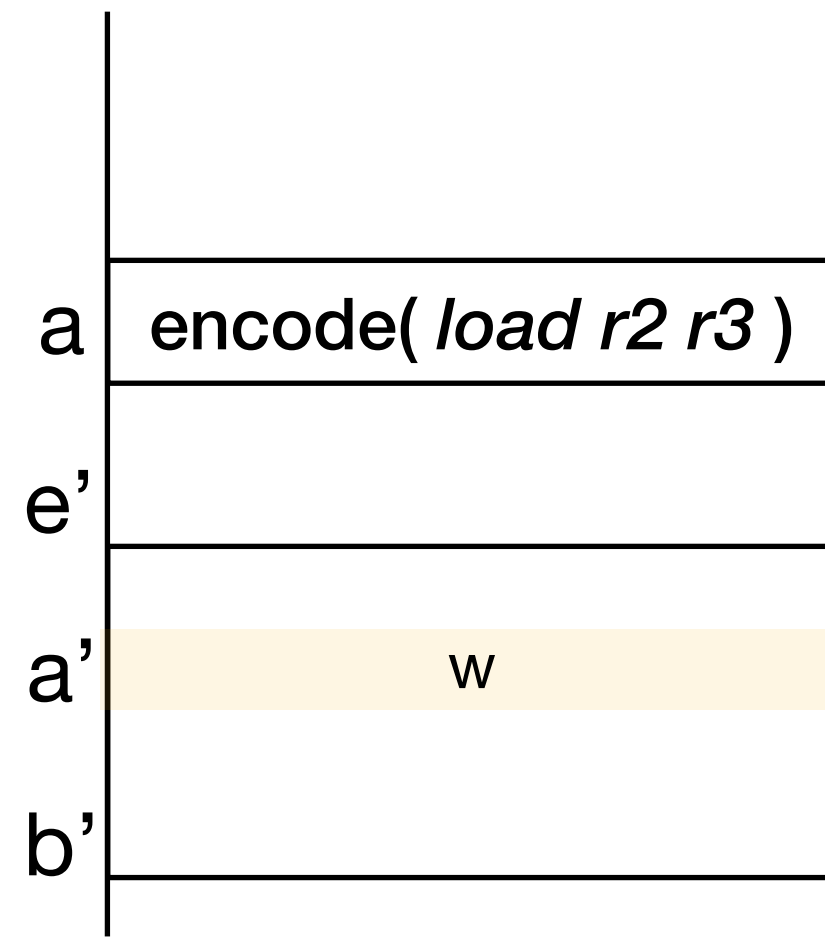
Operations

Load

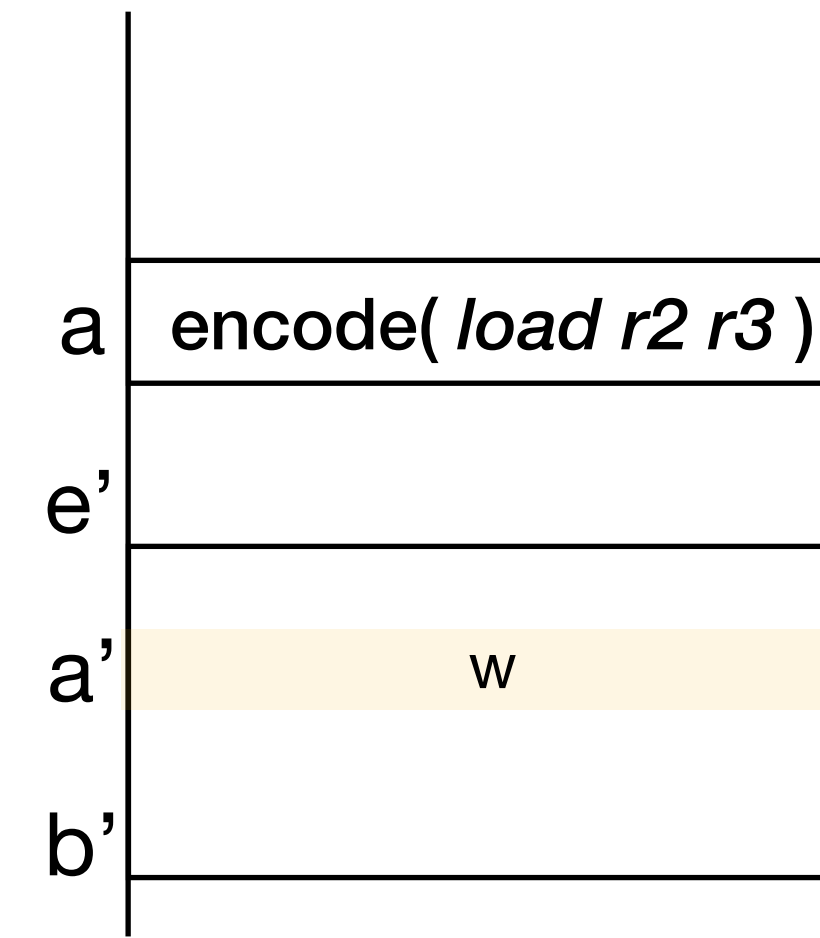
PC	(RX,b,e,a)
r1	
r2	
r3	(p,b',e',a')



PC	(RX,b,e,a+1)
r1	
r2	w
r3	(p,b',e',a')



$b' \leq a' < e'$
p is at least RO



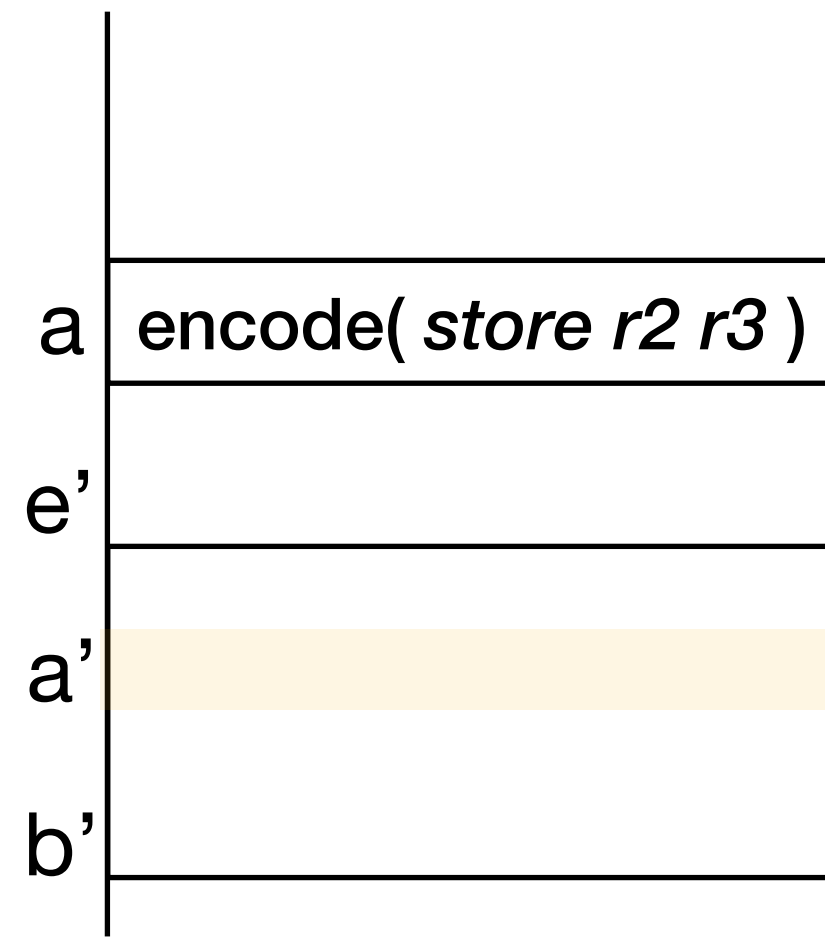
Operations

Store

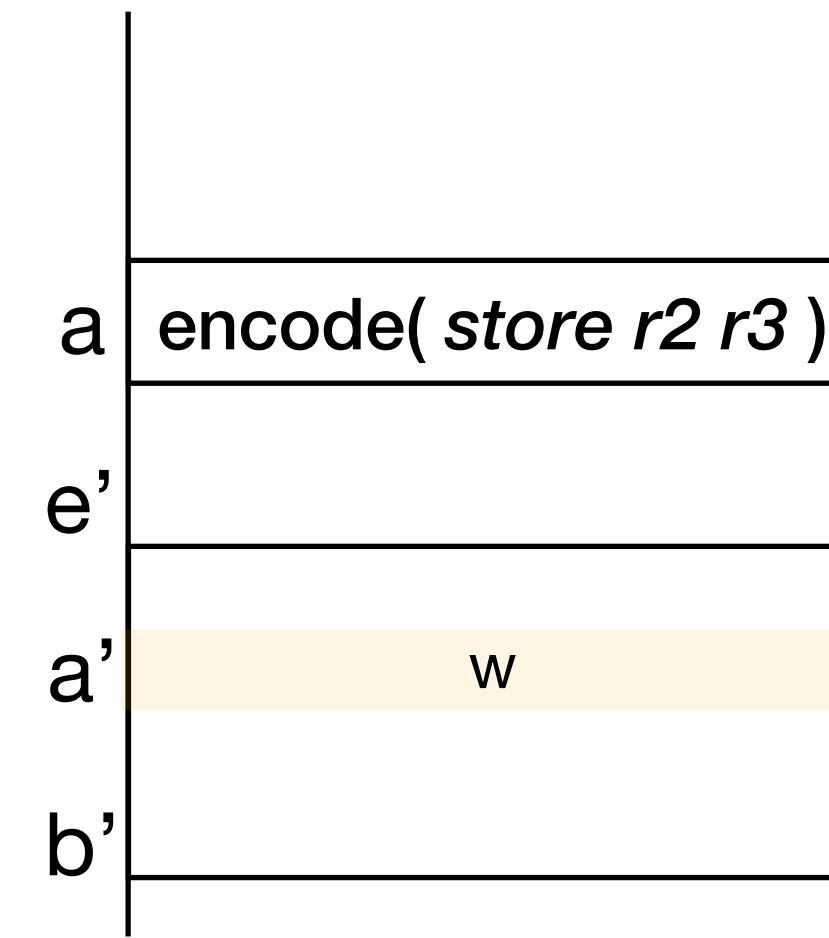
PC	(RX,b,e,a)
r1	
r2	(p,b',e',a')
r3	w



PC	(RX,b,e,a+1)
r1	
r2	(p,b',e',a')
r3	w



$b' \leq a' < e'$
p is at least RW




Operations

Jmp


PC	(RX,b,e,a)
r1	
r2	
r3	w



PC	w
r1	
r2	
r3	w



?




Operations

Jmp to an enter capability


PC	(RX,b,e,a)
r1	
r2	
r3	(E,b',e',a')



PC	
r1	
r2	
r3	(E,b',e',a')



?




Operations

Jmp to an enter capability


PC	(RX,b,e,a)
r1	
r2	
r3	(E,b',e',a')



PC	(RX,b',e',a')
r1	
r2	
r3	(E,b',e',a')




?



Operations

Jmp to an enter capability



Operations

Is it safe to have instructions that can change capabilities?

- Restrict : **lowers** the permission of a capability
- Subseg : **tightens** the bounds of a capability

What about pointer arithmetic?

- Lea : produces pointer arithmetic on the address of a **non-enter** capability

Operations

What should happen if a dynamic check fails?

Operational Semantics

Exercise: symbolically execute the following program

```

0:  mov r1 PC
    lea r1 5
    load r1 r1
    subseg r1 10 11
    jmp r2
5:  (RW,9,12,9)
6:  lea r1 1
7:  store r1 42
8:  halt
9:  0
10: 0
11: 0

```

From the following starting register state:

PC	(RWX,0,6,0)
r1	0
r2	(E,6,9,0)
r3	0


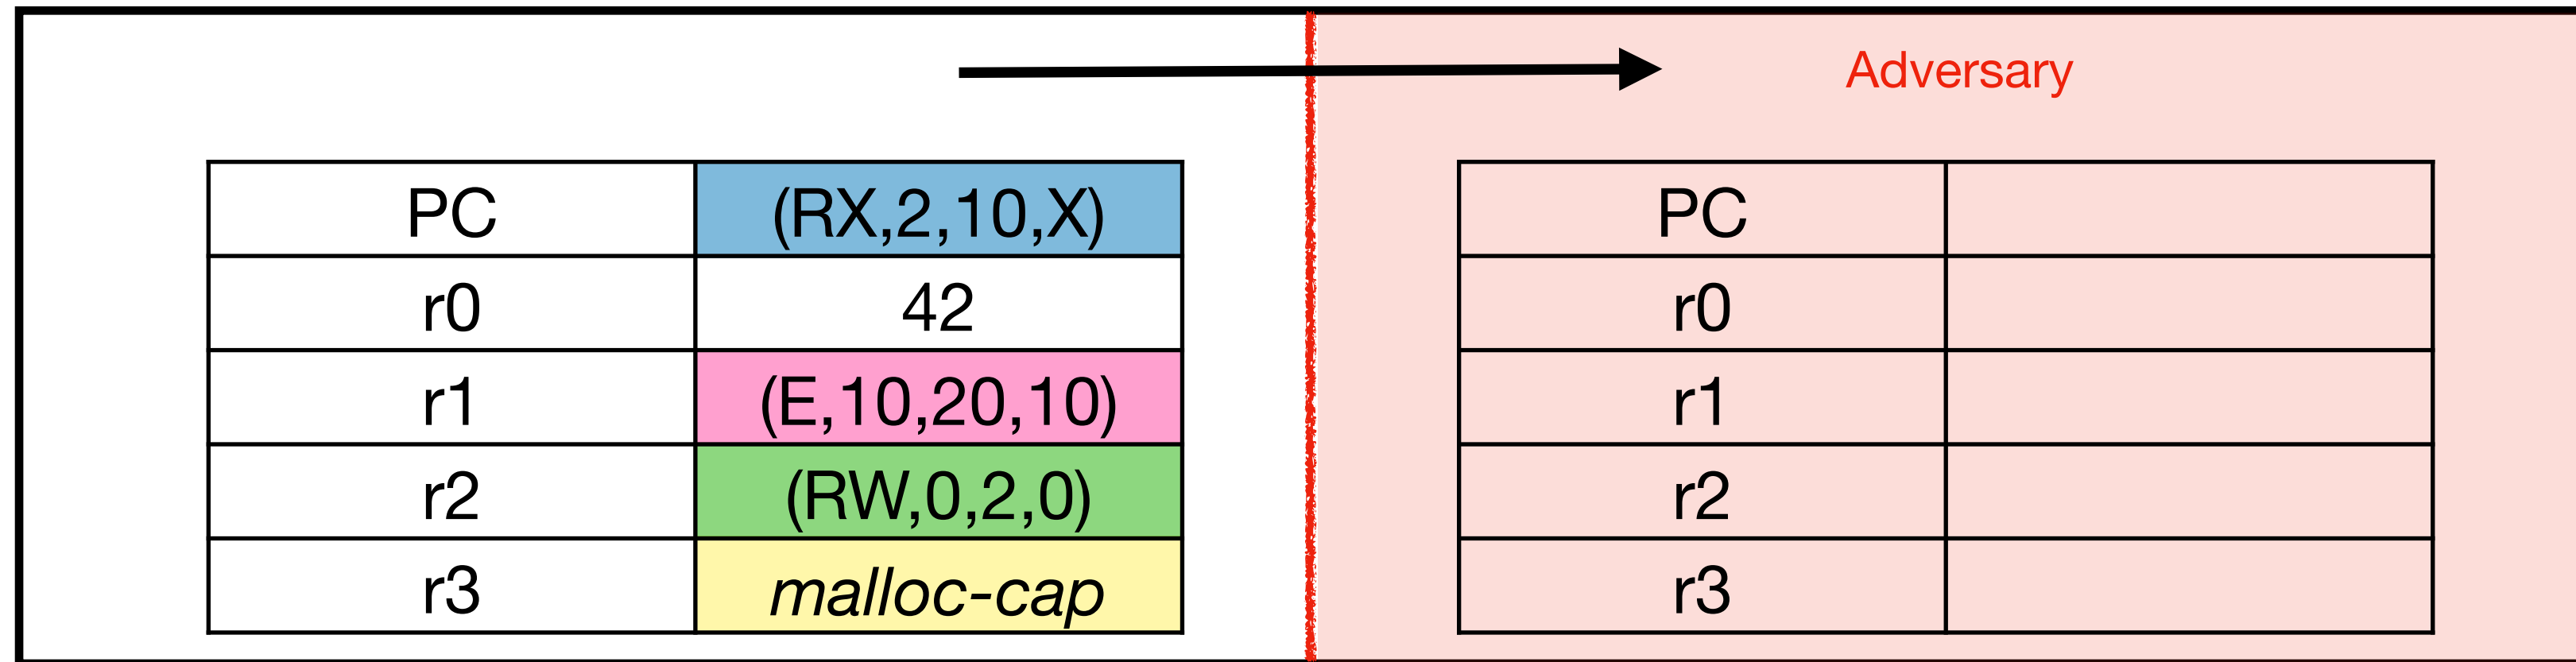
instr	dst	src	PC	(RX,6,9,8)
		r1		(RW,10,11,10)

lea dst z adds z to the address of the capability in dst

				0
--	--	--	--	---



subseg dst z1 z2 changes the bounds of the capability in dst to [z1,z2)

Secure Domain Change



Secure Domain Change

		Adversary	
PC	(RX,2,10,X)	PC	(RX,10,20,10)
r0	42	r0	(E,b,e,b)
r1	(E,10,20,10)	r1	(E,10,20,10)
r2	(RW,0,2,0)	r2	0
r3	<i>malloc-cap</i>	r3	<i>malloc-cap</i>



**Capability Safety
(and going further)**

Capability Safety

- Often we talk about memory safety, but as discussed, memory safety is not exactly a precise definition
- Instead, we can formally talk about “capability safety”

- What is capability safety?

“A single call must end with less (or equal) authority than it started out with”

- **Capability monotonicity**
- **Software compartmentalisation**

Rigorous engineering for hardware security:
Formal modelling and proof in the CHERI design
and implementation process

Kyndylan Nienhuis*, Alexandre Joannou*, Thomas Bauereiss*, Anthony Fox†, Michael Roe*, Brian Campbell‡, Matthew Naylor*, Robert M. Norton*, Simon W. Moore*, Peter G. Neumann§, Ian Stark‡, Robert N. M. Watson*, and Peter Sewell*

*University of Cambridge

†ARM Limited

‡University of Edinburgh

§SRI International

Capability Safety and Security

What other security properties can we enforce using capabilities?

In particular, what properties of **high level languages** can we now enforce?

let $x := \text{ref}(0)$ *in*

$\lambda f : () \rightarrow (), x \leftarrow 42; f(); \text{assert}(!x = 42)$

Local state encapsulation

Capability Safety and Security

What security properties can enforce using capabilities?

In particular, what properties of **high level languages** can we now enforce?

let $x := \text{ref}(0)$ *in*

$\lambda f : () \rightarrow (), x \leftarrow 42; f(); x \leftarrow 43; f(); \text{assert}(!x = 43)$



Well bracketed control flow

Well bracketed control flow

What is missing from our simple machine to be able to reason about well bracketed calls?

What is missing from our machine if you were to implement a compiler that targets a capability machine?

A stack!


Difficulties of a stack

granted
authority



lifetime
properties of
the stack

Difficulties of a stack



$\cdots; f(x); \cdots; f(y); \cdots;$

Difficulties of a stack

- This showed an example of local state encapsulation
- Other properties include:
 - Well bracketed control flow (frames are pushed and popped in a first in first out order)
 - Temporal stack safety (popped frames cannot be read by caller)

Difficulties of a stack

- How can we solve this issue?
 - Using a calling convention
- How can we solve this fast?
 - Using a calling convention and FANCY capabilities! (aka my research!)