

Concurrent Cerise

A program logic for multi-core capability machine

Bastien ROUSSEAU
Supervised by Lars Birkedal and Aïna Linn Georges
LogSem, Aarhus University, Denmark

2021

Formal methods on capability machine

Architecture security-oriented

- ▶ ~70% Microsoft security updates: memory safety
- ▶ Coarse-grained memory compartmentalization
- ▶ Capabilities Hardware-Enhanced RISC Instructions (CHERI)
- ▶ Arm Morello — Prototype January 2022

Formal methods


- ▶ Strong security guarantees
- ▶ Cerise: program logic & logical relation
- ▶ Huge gap between model and real-world machine

Outline

1. Context
 - ▶ Capability machines
 - ▶ Separation logic
 - ▶ Motivations
2. Contributions
 - ▶ Threat model
 - ▶ Program logic
 - ▶ Case study
3. Conclusion

Context

Capabilities machine



Hardware capability


- ▶ Usual CPU: everything is integer
- ▶ Capability machine: integer for arithmetic, capability for pointers
- ▶ Unforgeable token of authority

Representation

Capability (p,b,e,a)


$p \in \{RO, RW, RWX, \dots\}$	permission
$b \in Addr$	base address
$e \in Addr$	end address
$a \in Addr$	current address

Example sub-buffer




```
lea r1 [s-b]
store r1 42
lea r1 -[s-b]
subseg r1 b s
jmp radv
```

Example sub-buffer




```
lea r1 [s-b]      <---  
store r1 42  
lea r1 -[s-b]  
subseg r1 b s  
jmp radv
```

Example sub-buffer




```
lea r1 [s-b]
store r1 42 <---
lea r1 -[s-b]
subseg r1 b s
jmp radv
```


Example sub-buffer




```
lea r1 [s-b]
store r1 42
lea r1 -[s-b] <---
subseg r1 b s
jmp radv
```

Example sub-buffer



```
lea r1 [s-b]
store r1 42
lea r1 -[s-b]
subseg r1 b s <---
jmp radv
```

Example sub-buffer



```
lea r1 [s-b]
store r1 42
lea r1 -[s-b]
subseg r1 b s
jmp radv <---
```

Separation Logic in Iris

Iris: highly expressive framework for separation logic

Separation logic

- ▶ Logic of resources
- ▶ Separation conjunction $*$
 - ▶ exclusive ownership
 - ▶ disjoints resources
- ▶ Magic wand \multimap

$$P \triangleq (r_1 \multimap_r (\text{RWX}, \text{init}, \text{end}, \text{init}) * \text{init} \mapsto_a 0)$$

Cerise

What is Cerise ?

Operational semantic

- ▶ semantic for capability machine ISA
- ▶ security properties

Program logic

- ▶ resource for machine registers and memory addresses
- ▶ program specification *à la* Hoare triples

Logical relation

- ▶ *safe-to-share*: transitive access to *safe-to-share* words
- ▶ *safe-to-execute*: execute safely in any *safe context*

Goal & Motivations

Motivations

- ▶ Single-core machine
- ▶ Concurrency orthogonal with security ?
- ▶ Reduce gap between formal model and real-world machine

Goal — Add concurrency

- ▶ Define threat model
- ▶ Update semantic, program logic, logical relation
- ▶ Case study
- ▶ Synchronization

Contributions

Assumptions

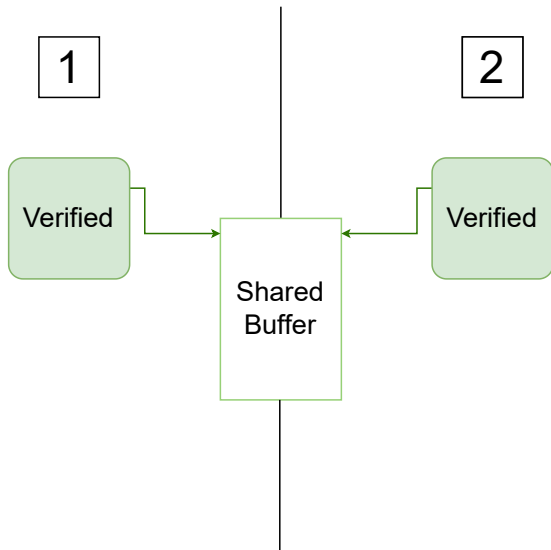
Model

- ▶ small ISA
- ▶ no virtual memory / page table
- ▶ sequentially consistent memory model


Sequentially consistent

- ▶ interleaving instructions
- ▶ non deterministic
- ▶ locally, behaves as sequential execution


Threat model




Threat model



Threat model



Threat model





Sub-buffer — Non concurrent

$$\left\{ \begin{array}{l} (i, r_0) \mapsto w_{adv} * \\ (i, r_1) \mapsto (RW, b, e, b) * \\ (RWX, \text{init}, \text{end}, \text{init}); [b, s) \mapsto [0 \dots 0] * \\ s \mapsto 0 * \\ [\text{init}, \text{end}) \mapsto \text{prog_instrs} \end{array} \right\}$$

$$\rightsquigarrow^i$$

$$\left\{ \begin{array}{l} (i, r_0) \mapsto w_{adv} * \\ (i, r_1) \mapsto (RW, b, s, b) * \\ w_{adv}; [b, s) \mapsto [0 \dots 0] * \\ s \mapsto 42 * \\ [\text{init}, \text{end}) \mapsto \text{prog_instrs} \end{array} \right\}$$





Sub-buffer — Non concurrent

$$\left\{ \begin{array}{l} (i, r_0) \mapsto w_{adv} * \\ (i, r_1) \mapsto (RW, b, e, b) * \\ (RWX, \text{init}, \text{end}, \text{init}); [b, s) \mapsto [0 \dots 0] * \\ s \mapsto 0 * \\ [\text{init}, \text{end}) \mapsto \text{prog_instrs} \end{array} \right\}$$

\rightsquigarrow^i

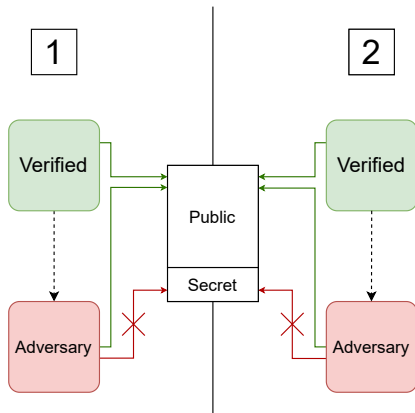
$$\left\{ \begin{array}{l} (i, r_0) \mapsto w_{adv} * \\ (i, r_1) \mapsto (RW, b, s, b) * \\ w_{adv}; [b, s) \mapsto [0 \dots 0] * \\ s \mapsto 42 * \\ [\text{init}, \text{end}) \mapsto \text{prog_instrs} \end{array} \right\}$$



Shared sub-buffer — Invariant

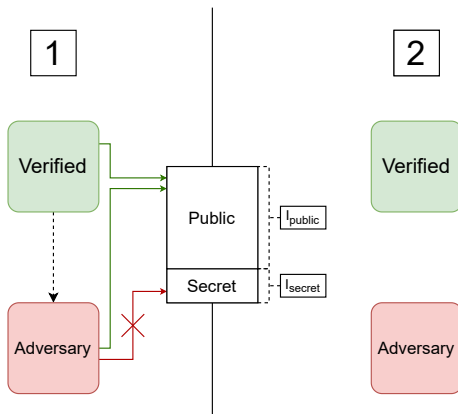
I_{public} — only *safe-to-share* words

I_{secret} — secret data



Shared sub-buffer — Invariant

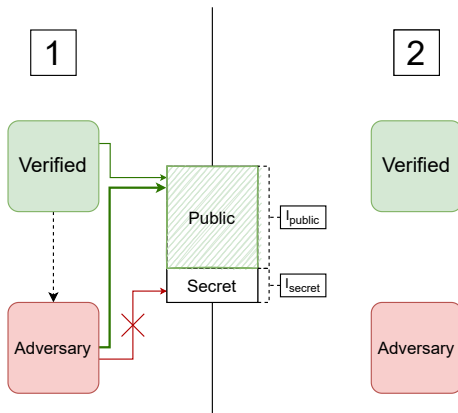
- I_{public} — only *safe-to-share* words
- I_{secret} — secret data



Shared sub-buffer — Invariant

$$I_{public} \triangleq *_{a \in [b,s)} \exists w, a \mapsto w * \mathcal{V}(w)$$

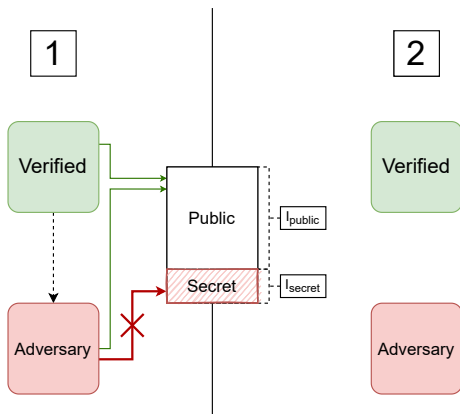
I_{secret} — secret data



Shared sub-buffer — Invariant

$$I_{public} \triangleq *_{a \in [b,s]} \exists w, a \mapsto w * \mathcal{V}(w)$$

$$I_{secret} \triangleq (s1 \mapsto 0 \vee s1 \mapsto secret_1) \wedge (s2 \mapsto 0 \vee s2 \mapsto secret_2)$$



Shared sub-buffer — Full specification

$$\boxed{I_{public}} * \boxed{I_{secret}} \vdash \left\{ \begin{array}{l} (i, r_0) \models w_{adv} * \\ (RWX, \text{init}, \text{end}, \text{init}); (i, r_1) \models (RW, b, e, b) * \\ [\text{init}, \text{end}] \mapsto \text{prog_instrs} \end{array} \right\}$$

\rightsquigarrow

$$\left\{ \begin{array}{l} (i, r_0) \models w_{adv} * \\ w_{adv}; (i, r_1) \models (RW, b, s, b) * \\ [\text{init}, \text{end}] \mapsto \text{prog_instrs} \end{array} \right\}$$

Complete specification

- ▶ assume the previous specification for each core
- ▶ Fundamental Theorem Logical Relation (non-trivial)
- ▶ safely complete execution

Adequacy theorem \rightarrow invariant on operational semantic

Conclusion

Further works

Synchronization

- ▶ Compare and Swap instruction
- ▶ Spinlock library
- ▶ Concurrently safe macro for dynamic allocation
- ▶ Scenario involving malloc

Other

- ▶ Other scenarios
- ▶ Pedagogical exercises to learn Cerise in Coq (WIP)

Summary and Future Work

Reason formally on multi-core capability machines

Summary

- ▶ Extend Cerise with concurrency
- ▶ Define the threat model
- ▶ Add synchronization mechanism
- ▶ Design, specify and prove case studies
- ▶ Mechanized and proved with Iris and Coq

Future work


- ▶ Relaxed memory model
- ▶ Real world ISA
- ▶ Virtual memory / Page tables
- ▶ Security properties, including interrupts/exceptions

Appendix

Hoare triples — Example load

$$\frac{\text{ValidPC}(p_{pc}, b_{pc}, e_{pc}, a_{pc}) \quad \text{ValidLoad}(p, b, e, a) \quad \text{decode}(n) = \text{load } dst \text{ } src}{\left\langle \begin{array}{l} (i, pc) \Rightarrow (p_{pc}, b_{pc}, e_{pc}, a_{pc}) * a_{pc} \mapsto n * \\ (i, dst) \Rightarrow - * (i, src) \Rightarrow (p, b, e, a) * a \mapsto w \end{array} \right\rangle \xrightarrow{i} \left\langle \begin{array}{l} (i, pc) \Rightarrow (p_{pc}, b_{pc}, e_{pc}, a_{pc} + 1) * a_{pc} \mapsto n * \\ (i, dst) \Rightarrow w * (i, src) \Rightarrow (p, b, e, a) * a \mapsto w \end{array} \right\rangle}$$

Shared Sub-buffer — Invariant



$$I_{public} \triangleq *_{a \in [p,s]} \exists w, a \mapsto w * \mathcal{V}(w)$$

$$I_{secret} \triangleq (s1 \mapsto 0 \vee s1 \mapsto secret_1) \wedge (s2 \mapsto 0 \vee s2 \mapsto secret_2)$$

Logical relation — Full definition

$$\mathcal{V}(w) \begin{cases} \mathcal{V}(z), \mathcal{V}(0, -, -, -) \triangleq \text{True} \\ \mathcal{V}(E, b, e, a) \triangleq \triangleright \square \mathcal{E}(\text{RX}, b, e, a) \\ \mathcal{V}(\text{RW}/\text{RWX}, b, e, -) \triangleq *_{a \in [b, e]} \boxed{\exists w, a \mapsto w * \mathcal{V}(w)} \\ \mathcal{V}(\text{RO}/\text{RX}, b, e, -) \triangleq *_{a \in [b, e]} \exists P, \boxed{\exists w, a \mapsto w * P(w)} * \triangleright \square (\forall w, P(w) \multimap \mathcal{V}(w)) \end{cases}$$

$$\mathcal{E}(w) \triangleq \forall i \text{ reg}, \left\{ (i, w); *_{\substack{(j, r, v) \in \text{reg}, \\ i=j \\ r \neq \text{pc}}} (j, r) \mapsto v * \mathcal{V}(v) \right\} \overset{i}{\rightsquigarrow} \bullet$$

Spinlock

Principle

- ▶ prevent concurrent access to a critical section
- ▶ loop and try acquire the lock since available

Acquire

```
; r0 -> capability pointing to the lock state
loop:
mov r3 0
cas r0 r3 1
jnz r3 [loop]
end:
```

Spinlock

Specification — Acquire

$$\boxed{[a_{\text{acquire}}, e_{\text{acquire}}] \mapsto \text{instrs}_{\text{acquire}} * \text{is_lock} \ \gamma \ a_{\text{lock}} \ P}$$
$$\vdash \left\{ \begin{array}{l} (i, PC) \Rightarrow (\text{RWX}, b_{pc}, e_{pc}, a_{\text{acquire}}) * \\ (i, r_0) \Rightarrow (p_{\text{lock}}, b_{\text{lock}}, e_{\text{lock}}, a_{\text{lock}}) * [\dots] \end{array} \right\} \xrightarrow{i} \left\{ \begin{array}{l} (i, PC) \Rightarrow (\text{RWX}, b_{pc}, e_{pc}, e_{\text{acquire}}) * \\ (i, r_0) \Rightarrow (p_{\text{lock}}, b_{\text{lock}}, e_{\text{lock}}, a_{\text{lock}}) * [\dots] * \\ P * \text{locked} \ \gamma \end{array} \right\}$$

Specification — Release

$$\boxed{[a_{\text{release}}, e_{\text{release}}] \mapsto \text{instrs}_{\text{release}} * \text{is_lock} \ \gamma \ a_{\text{lock}} \ P}$$
$$\vdash \left\{ \begin{array}{l} (i, PC) \Rightarrow (\text{RWX}, b_{pc}, e_{pc}, a_{\text{release}}) * \\ (i, r_0) \Rightarrow (p_{\text{lock}}, b_{\text{lock}}, e_{\text{lock}}, a_{\text{lock}}) * \\ P * \text{locked} \ \gamma \end{array} \right\} \xrightarrow{i} \left\{ \begin{array}{l} (i, PC) \Rightarrow (\text{RWX}, b_{pc}, e_{pc}, e_{\text{release}}) * \\ (i, r_0) \Rightarrow (p_{\text{lock}}, b_{\text{lock}}, e_{\text{lock}}, a_{\text{lock}}) \end{array} \right\}$$

Operational Semantic

Model N-cores capability machine

i	$\in \mathcal{N}$	\triangleq	$\{i \mid i \in \mathbb{N} \wedge i < N\}$
m	$\in \text{Mem}$	\triangleq	$\text{Addr} \rightarrow \text{Word}$
reg	$\in \text{Reg}$	\triangleq	$(\mathcal{N} \times \text{RegName}) \rightarrow \text{Word}$
\mathcal{C}	$\in \text{CoreState}$	$::=$	$\text{Running} \mid \text{Halted} \mid \text{Failed}$
\mathcal{E}	$\in \text{ExecState}$	\triangleq	$\mathcal{N} \rightarrow \text{CoreState}$
φ	$\in \text{Conf}$	\triangleq	$\text{ExecState} \times \text{Reg} \times \text{Mem}$

Operational Semantic

Reduction relation

- ▶ per-core reduction:

$$(\text{CoreState} \times \text{Reg} \times \text{Mem}) \xrightarrow[\text{core}]{i} (\text{CoreState} \times \text{Reg} \times \text{Mem})$$

- ▶ configuration reduction

$$\frac{(s, r, m) \xrightarrow[\text{core}]{i} (s', r', m')}{(\mathcal{E}[i \mapsto s], r, m) \xrightarrow[\text{conf}]{} (\mathcal{E}[i \mapsto s'], r', m')}$$

Example: load r_1 r_2

$$\frac{\begin{array}{l} \text{reg}(i, r_2) = (p, b, e, a) \quad \text{ValidLoad}(p, b, e, a) \\ \text{mem}(a) = w \quad \text{reg}' \triangleq \text{updPC}(\text{reg}(i, r_1) \mapsto w) \end{array}}{(\text{Running}, \text{reg}, \text{mem}) \xrightarrow[\text{core}]{i} (\text{Running}, \text{reg}', \text{mem})}$$

Program Logic

Syntax

$P, Q \in iProp ::=$

$\top \mid \perp \mid P \wedge Q \mid \dots$

$\mid P * Q \mid P \multimap Q$

$\mid [\varphi] \mid \square P \mid \triangleright P$

$\mid \boxed{P}$

$\mid a \mapsto w \mid (i, r) \Rrightarrow w$

$\mid \langle P \rangle \xrightarrow{i} \langle s. Q \rangle \mid \{P\} \overset{i}{\rightsquigarrow} \{s. Q\} \mid \{P\} \overset{i}{\rightsquigarrow} \bullet$

HOL

separation logic

iris features

iris invariant

machine resources

program logic

Program Logic

Program specification

$\langle P \rangle \xrightarrow{i} \langle s. Q \rangle$	single instruction
$\{P\} \rightsquigarrow^i \{s. Q\}$	code fragment
$\{P\} \rightsquigarrow^i \bullet$	complete safe execution.

Sequencing rule

Compose the specification together

$$\frac{\{P\} \rightsquigarrow^i \{Q\} \quad \{Q\} \rightsquigarrow^i \{R\}}{\{P\} \rightsquigarrow^i \{R\}}$$

Logical Relation

Capture semantic properties of the language

Overview Definition

$$\mathcal{V}(w) \begin{cases} \mathcal{V}(z), \mathcal{V}(O, -, -, -) \triangleq \text{True} \\ \mathcal{V}(\text{RW/RWX}, b, e, -) \triangleq *_{a \in [b, e]} \boxed{\exists w, a \mapsto w * \mathcal{V}(w)} \\ \mathcal{V}(E, b, e, a) \triangleq \triangleright \square \mathcal{E}(\text{RX}, b, e, a) \\ \mathcal{V}(\text{RO/RX}, b, e, -) \triangleq - \end{cases}$$
$$\mathcal{E}(w) \triangleq \forall i \{w; \text{any safe context}\} \rightsquigarrow^i \bullet$$

FTLR

Fundamental Theorem of Logical Relation

$$\forall w, \mathcal{V}(w) \Rightarrow \mathcal{E}(w)$$