# WasmRef-Isabelle or how to formally verify a not-slow interpreter with not-insane amount of effort Maja Trela<sup>1</sup>

<sup>1</sup>affiliation note: worked on this at University of Cambridge, now employed at Jane Street

Lambda Days 2024

28 May 2024

#### RESEARCH-ARTICLE OPEN ACCESS • X in **the** f WasmRef-Isabelle: A Verified Monadic Interpreter and Industrial Fuzzing Oracle for WebAssembly

Authors: Sconrad Watt, Maja Trela, Peter Lammich, Florian Märkl Authors Info & Claims

Proceedings of the ACM on Programming Languages, Volume 7, Issue PLDI • Article No.: 110, pp 100-123 • https:// doi.org/10.1145/3591224

Published: 06 June 2023 Publication History



## WebAssembly (Wasm)



- Executable bytecode format
- Precise small-step semantics
- Strongly-typed

### Wasmtime

- Cross-platform WebAssembly runtime
- Focused on correctness and security

#### Features

- Fast. Wasmtime is built on the optimizing Cranelift code generator to quickly generate high-quality machine code either at runtime or ahead-of-time. Wasmtime is optimized for efficient instantiation, low-overhead calls between the embedder and wasm, and scalability of concurrent instances.
- Secure. Wasmtime's development is strongly focused on correctness and security. Building on top of Rust's runtime safety guarantees, each Wasmtime feature goes through careful review and consideration via an RFC process. Once features are designed and implemented, they undergo 24/7 fuzzing donated by Google's OSS Fuzz. As features stabilize they become part of a release, and when things go wrong we have a well-defined security policy in place to quickly mitigate and patch any issues. We follow best practices for defense-in-depth and integrate protections and mitigations for issues like Spectre. Finally, we're working to push the state-of-the-art by collaborating with academic researchers to formally verify critical parts of Wasmtime and Cranelift.

## Wasmtime fuzzing

### Fuzzing: comparing Wasmtime with a reference implementation on random input

Use the official reference interpreter?

README.md

#### WebAssembly Reference Interpreter

This repository implements an interpreter for WebAssembly. It is written for clarity and simplicity, not speed. It is intended as a playground for trying out ideas and a device for nailing down their exact semantics. For that purpose, the code is written in a fairly declarative, "speccy" way.

### Reference interpreter problem

### Evaluation contexts are inefficient

 Execution time scales quadratically



takes many dozens of seconds to run in the interpreter. If the block ... end is all removed though it finishes near instantaneously.

### Interpreter tradeoffs



#### Can't move towards one without giving up on the other!



#### Could we do something like this?



### WasmRef-Isabelle

#### This is exactly WasmRef-Isabelle:

- Verified with respect to WebAssembly semantics
- ► Fast enough for use in fuzzing; adopted by Wasmtime team
  - Running in their CI infrastructure!

### WasmCert-Isabelle

Building on WasmCert-Isabelle:



- Mechanization of WebAssembly semantics in Isabelle/HOL
- Notably, found type safety bugs before WebAssembly release
- Interpreter written in Isabelle/HOL, extracted to OCaml

# What are we verifying?

• We prove *soundness*:

for every initial configuration, if the interpreter terminates without crashing, the result is consistent with the given specification

Note: says nothing about it not crashing

theorem intermediate\_interpreter\_sound: assumes "interpreter start\_config = (new\_store, RValue values)" shows "accepted\_by\_semantics start\_config new\_store values"

### Interpreters

### Pure interpreter:

- Sound wrt. spec
- Tweaked WasmCert-Isabelle interpreter
- No mutable state, uses lists (O(n) memory access)

### Monadic interpreter:

- ► Sound wrt. pure interpreter
- Mutable arrays (O(1) memory access)
- Needs state monad

### Comparison

```
definition load
:: "mem ⇒ nat ⇒ off ⇒ nat ⇒ bytes option"
where
"load m n off l =
   (if (mem_length m ≥ (n+off+l))
        then Some (read_bytes m (n+off) l)
        else None)"
```

```
definition load_bytes_m_v
:: "mem_m ⇒ nat ⇒ off ⇒ nat ⇒ (bytes option) Heap"
where
"load_bytes_m_v m n off l =
    do {
        m_len ← len_byte_array (fst m);
        let ind = n+off;
        (if (m_len ≥ (ind+l)) then do {
            bs ← load_uint8_list (fst m) ind l;
            return (Some bs) }
    else
            return None)
}"
```

```
'a Heap:
text <<u>Monadic</u> heap actions either produce values
and transform the heap, or fail>
datatype 'a Heap = Heap "heap ⇒ ('a × heap) option"
```

### Proof

- Relate two implementations with Hoare triples
- Separation logic: Sepref (modified)
- Automation: sep\_auto, Sledgehammer

### Sledgehammer!

```
proof (prove)
goal (4 subgoals):
1. \bigwedgexl x2 xla a b.
    m_m = (xla, x2) \implies
    n + off + l \leq mem_length (x1, x2) \implies
    n + off + l \leq length (Rep_mem_rep x1) \implies
    0 < l \implies
    (a, b) \models xla \mapsto Rep_mem_rep x1 \implies
    take l (drop (n + off) (Rep_mem_rep x1)) = read_bytes (x1, x2) (n + off) l
```



"z3": Try this: apply (metis mem\_rep\_read\_bytes.rep\_eq\_prod.sel(1) read\_bytes\_def) (83 ms) "vampire": Try this: apply (metis mem rep\_read\_bytes.rep\_eq\_prod.sel(1) read\_bytes\_def) (104 r

### Not always so smooth

emma store vec triple: shows "<mems m assn ms ms m \* inst m assn (f inst f) f inst2> app s f v s store vec m sv off ms m f inst2 v s  $<\lambda r$ . let (ms', v s', res) = app s f v s store vec sv off ms f v s in  $\uparrow$ (r = (v s', res)) \* mems m assn ms' ms m \* inst m assn (f inst f) f inst2>t" unfolding app s f v s store vec m def app s f v s store vec def inst m assn def list assn conv idx mems m assn def apply(sep auto split: v num.splits v.splits v vec.splits prod.splits) apply(extract reinsert list idx "inst.mems (f inst f) ! 0") apply(sep auto) apply(extract list idx "inst.mems (f inst f) ! 0") apply(sep auto heap:store uint8 list triple) apply(sep auto) apply(sep auto simp:smem ind def split:list.splits prod.splits) apply(sep auto simp:store def let def) apply(sep auto simp:store def Let def) apply(rule listI assn reinsert upd, frame inference, simp, simp) apply(sep auto simp:store def Let def bytes takefill def) apply(sep auto simp:smem ind def store def split:list.splits) apply(sep auto)+ done

#### Issues

- Higher-order functions
- Shared mutable states
- "UX" problems

## **Proof summary**

#### Two-step refinement proof



### Interpreter performance

#### Good worst-case performance!



Fig. 12. Graphing execution times for the function  $f(0) \triangleq 0$ ;  $f(n + 1) \triangleq f(n) + n^5$ (note the log scale for runtime).



## Interpreter performance (cont.)

#### No longer the bottleneck in fuzzing!

oracle	total tests	slowest test (secs)
Reference Interpreter	1900091	48
WasmRef-Isabelle	2740908	< 1
Wasmi (release)	2793736	< 1

## Conclusions

- Monadic soundness proof done in a few months as a 4th year project
- How to verify a not-slow interpreter with a not-insane amount of effort?
  - Refinement proof + decent automation!
- Only possible due to good tools
  - Neel Krishnaswami, "The Golden Age of PL Research"
- As costs fall we should expect more projects!

# Thank you for listening!

#### Email: maja@majatrela.com References:

- Paper dl.acm.org/doi/abs/10.1145/3591224
- Wasmtime fuzzing bytecodealliance.org/articles/security-and-correctness-in-wasmtime
- Proof repo github.com/WasmCert/WasmCert-Isabelle/tree/new\_interp\_2021-1
- The Golden Age of PL Research semantic-domain.blogspot.com/2022/09/the-golden-age-of-pl-research.html
- Personal (still incomplete) website: majatrela.com