WebAssembly

history, internals, security and more

Mingshen Sun, April 16, 2020

Background History all starts from the Web

- Text
- HTML
- Flash, ActionScript, Adobe Air, Silverlight, Java, ActiveX, JavaScript







Background History all starts from the Web

- Text
- HTML
- Flash, ActionScript, Adobe Air, Silverlight, Java, ActiveX, JavaScript

Background JavaScript engine

- SpiderMonkey
- V8
- WebKit/JavaScriptCore
- Chakra

Background JavaScript is ...

	Console	Search	Emula
(97	<top fr<="" th=""><th>ame></th></top>	ame>
	> <mark>0 > nu</mark> false	11	
	> 0 >= n true	ull	
	> 0 == n false	ull	
	> 0 <= n true	ull	
	> <mark>0 < nu</mark> false	11	
1	> FML		

undefined==false	false
null==false	false
undefined==0	false
null==0	false
undefined==''	false
null==''	false
false==''	true
''==0	true
false==0	true
null==undefined	true

>	<pre>let objOne = [] undefined</pre>
> .:	<pre>let objTwo = [] undefined</pre>
> <·	objOne == objTwo false
> <:	<pre>let arr = [1, 2, 3, 4] undefined</pre>
> 	<pre>let arrDup = [arr] undefined</pre>
>	console.log(arrDup)
	▶ (4) [1, 2, 3, 4]
<-	undefined
> <:	arr == arrDup false
> <:	objOne === objTwo false
> <·	arr === arrDup false

Background

JavaScript is confusing.



https://stackoverflow.com/questions/359494/which-equals-operator-vs-should-be-used-in-javascript-comparisons



Background TypeScript



JavaScript that scales.

TypeScript is a <u>typed superset of</u> <u>JavaScript</u> that compiles to plain JavaScript.



Background asm.js

- asm.js is a <u>subset</u> of JavaScript
- allow other languages such as C to be run as web applications while maintaining performance characteristics considerably better than standard JavaScript,



Introduction

WebAssembly to save the world?

 WebAssembly (abbreviated Wasm) is a safe, portable, low-level code format designed for efficient execution and compact representation.



Introduction

Where does WebAssembly Fit?



https://hacks.mozilla.org/2017/02/creating-and-working-with-webassembly-modules/

Introduction

Where does WebAssembly Fit?



https://hacks.mozilla.org/2017/02/creating-and-working-with-webassembly-modules/

wat/wasm

Binary representation (.wasm)

wat2wasm demo

WebAssembly has a text format and a binary format. This demo converts from the text format to the binary format.

Enter WebAssembly text in the textarea on the left. The right side will either show an error, or will show a log with a description of the generated binary file.

WAT example: simple 🗘 Do	wnload BUILD LOG	
1 (module 2 (func \$addTwo (param i32 i32) (result i32) 3 rot logal 0	0000000: 0061 736d 0000004: 0100 0000	; WASM_BINARY_MAGIC ; WASM_BINARY_VERSION
3 get_local 0	; section type (1)	, sostion and
4 get_local I		; section code
6 (export "addTwo" (func \$addTwo)))	0000003: 00	, section size (guess)
7		, num cypes
	, cype 0	• func
	0000002: 02	, Tunc
	000000C. 02	, india params
	000000d: 71	; 132
	000000E. 01	, 152
	0000001: 01	; num results
		; 132
		; FIXUP section size
	; section "Function" (3)	
		; section code
(tual raprocentation (wat)	0000012: 00	; section size (guess)
(luar representation (.wat)	0000013: 01	; num functions
	0000014: 00	; function 0 signature index
	0000012: 02	; FIXUP section size
	; section "Export" (7)	
	0000015: 07	; section code
	0000016: 00	; section size (quess)
JS	JS LOG	
<pre>1 const wasmInstance = 2 new WebAssembly.Instance(wasmModule, {});</pre>	0 2	
<pre>3 const { addTwo } = wasmInstance.exports;</pre>	4	
4 for (let i = 0; i < 10; i++) {	6	
<pre>5 console.log(addTwo(i, i));</pre>	8	
6 }	10	
7		
	JS APIs	

JavaScript APIs

```
WebAssembly.compile(new Uint8Array(
  `00 61 73 6d 0d 00 00 00 01 09 02 60
                                         00 00 60 01
7f 01 7f 03 03 02 00 01 05 04 01 00 80 01 07 07
01 03 66 6f 6f 00 01 08 01 00 0a 3a 02 02 00 0b
35 01 01 7f 20 00 41 04 6c 21 01 03 40 01 01 01
0b 03 7f 41 01 0b 20 01 41 e4 00 6c 41 cd 02 20
01 1b 21 01 41 00 20 01 36 02 00 41 00 21 01 41
00 28 02 00 0f 0b 0b 0e 01 00 41 00 0b 08 00 00
00 00 2c 00 00 00`.split(/[\s\r\n]+/g).map(v => parseInt(v, 16))
)).then(mod => {
 let m = new WebAssembly.Instance(mod)
 console.log('foo(1) =>', m.exports.foo(1))
 console.log('foo(2) =>', m.exports.foo(2))
 console.log('foo(3) =>', m.exports.foo(3))
})
```

wat/wasm

; function body 0	
0000024: 00	; func body size (guess)
0000025: 00	; local decl count
0000026: 20	; local.get
0000027: 00	; local index
0000028: 20	; local.get
0000029: 01	; local index
000002a: 6a	; i32.add
000002b: 0b	; end
0000024: 07	; FIXUP func body size
0000022: 09	; FIXUP section size

S-expression

C++11 -Os	COMPILE Wat	ASSEMBLE	DOWNLOAD	Firefox x86 Assembly	<
<pre>1- int add(int a, int b) { 2 return a + b; 3 }</pre>	ped Myhy? h blob/ma	nc) (" (memory \$0)) (; 0;) (param \$0 i32) (param) achine, not an based ttps://github.ce ster/Rationale	am \$1 i32) AST, o I byteco om/We .md#w	 wasm-function[0]: sub rsp, 8 mov ecx, esi mov eax, ecx add eax, edi nop add rsp, 8 ret ret sadd rsp, 8 ret 	; 0x000000 48 83 ec ; 0x000004 8b ce ; 0x000006 8b c1 ; 0x000008 03 c7 ; 0x000000 66 90 ; 0x00000c 48 83 c4 ; 0x000010 c3
				0	
✓ Console					75

https://mbebenita.github.io/WasmExplorer/

Wasm Anatomy of a WebAssembly program

module "exam	nple1"						
version	1						
type section			start section				
#0	#0 (i32, i32) → (i32)			function #0			
#1	() → ()						
			code section				
<pre>import section #0 "half" from "example2" of type 0</pre>			#0	i32.const 2 i32.load 2 0 0 call 1			
function section #1 type 1 #2 type 0			#1	load_local 0 load_local 1 i32.mul			
export section		(data section				
"double" of type 0			0x5 0x0 0x0 0x0				

...

Wasm Anatomy of a WebAssembly program

module "example1"				
version 1				
type section		start section		
#0	Types	¢0		
#1	.,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,			
import section	WebAssembly has the f	following value types:		
#0 "half" f	 i32: 32-bit integer i64: 64-bit integer 			
function section	f32 : 32-bit floating point			
#1	g point			
		data soction		
export section				
"doub	le" of type 0	0x5 0x0 0x0 0x0		

...

Wasm	A module contains the following sections:	
Anatomy o	import	
nodule "example1"	exportstart	
version 1	• global	
type section	memory	
#0 (i3	• data	
#1	• table	
	elements	
import section	 function and code 	
#0 "half" from	A module also defines several index spaces which are	
function section	section fields in the module:	
#1 #2	 the function index space the global index space	
export section	 the linear memory index space 	
"double	 the table index space 	¢0

...

Wasm Stack Machine



Wasm Memory Model Linear memory mordel

- Memory is linear
- memory-page size: 64KiB

Memory being "linear" means that there's no random allocator operators available — all memory addresses used in a module's code are expressed in terms of byte offsets from the beginning of a memory segment.

WASM being a low-level format, this makes a lot of sense. It's up to the higher-level language that targets WASM to provide memory management on top of this linear memory space, if needed.

Wasm Memory Model Load/Store

16-bit

32-bit



i64.store32 2 3 // align = 2 = 32-bit, offset = 3 => addr = 6

64-bit



Wasm Memory Model

Load/Store



The **effective address** is the offset in bytes measured from the beginning of the memory. The effective address is the sum of the address operand and the offset immediate.

```
effective-address = address-operand + offset-immediate
```

https://rsms.me/wasm-intro

WASI

a system interface for WebAssembly to run outside the web

- POSIX (The Portable Operating System Interface)
 - Glibc, musl libc
- MSVC
- CloudABI
- Android

WASI

a system interface for WebAssembly to run outside the web

• POSIX (The Portable Operating System Interface)

Capability-based security means that processes can only perform actions

- that have no global impact. Processes cannot open files by their absolute
- path, cannot open network
 connections, and cannot observe
 global system state such as the
- process table.

CloudABI is what you get if you take POSIX, add capability-based security, and remove everything that's incompatible with that. The result is a minimal ABI consisting of only 49 syscalls.

WASI API

For example, instead of a typical open system call, WASI provides an openatlike system call, requiring the calling process to have a file descriptor for a directory that contains the file, representing the capability to open files within that directory.



Example Hello, World!

```
fn main() {
    println!("Hello, world!");
}
```

```
$ rustup target add wasm32-wasi
$ rustc hello.rs --target wasm32-wasi
$ wasmtime hello.wasm
Hello, world!
```

Example Hello, World







Example Hello, World!

(module

(type \$t0 (func)) Results (type **\$t1** (func (param i32))) (type **\$t2** (func (param i32) (result i64))) (type \$t3 (func (param i32 i32) (result i32))) error: errno (type \$t4 (func (param i32 i32 i32 i32))) (type \$t5 (func (param i32 i32))) • nwritten : size The number of bytes written. (type **\$t6** (func (param i32 i32 i32))) (type **\$t7** (func (param i32) (result i32))) (type **\$t8** (func (param i32 i32 i32) (result i32))) (type **\$t9** (func (param i32 i32 i32 i32) (result i32))) (type **\$t10** (func (result i32))) (type \$t11 (func (param i32 i32 i32 i32 i32))) (type **\$t12** (func (param i32 i32 i32 i32 i32) (result i32))) (type **\$t13** (func (param i32 i32 i32 i32 i32 i32) (result i32))) (type **\$t14** (func (param i64 i32 i32) (result i32))) (import "wasi_snapshot_preview1" "proc_exit" (func \$__wasi_proc_exit (type \$t1))) (import "wasi_snapshot_preview1" "fd_write" (func \$_ZN4wasi13lib_generated22wasi_snapshot_preview18fd_write17h647699597ede499dE (type \$t9))) (import "wasi_snapshot_preview1" "fd_prestat_get" (func \$__wasi_fd_prestat_get (type \$t3))) (import "wasi_snapshot_preview1" "fd_prestat_dir_name" (func **\$__wasi_fd_prestat_dir_name** (type **\$t8**))) (import "wasi_snapshot_preview1" "environ_sizes_get" (func \$__wasi_environ_sizes_get (type \$t3))) (import "wasi_snapshot_preview1" "environ_get" (func \$__wasi_environ_get (type \$t3))) (func **\$__wasm_call_ctors** (type **\$t0**) (call \$__wasilibc_populate_environ) (call \$__wasilibc_populate_libpreopen))

Params

• fd : fd

fd_write(fd: fd, iovs: ciovec_array) -> (errno, size)

• iovs : ciovec array List of scatter/gather vectors from which to retrieve data.

Write to a file descriptor. Note: This is similar to writev in POSIX.





Example Hello, World

```
fn fd_write(&self, fd: types::Fd, ciovs: &types::CiovecArray<'_>) -> Result<types::Size> {
424
425
              let mut bc = GuestBorrows::new();
426
             let mut slices = Vec::new();
             bc.borrow_slice(&ciovs)?;
427
             for ciov_ptr in ciovs.iter() {
428
                 let ciov_ptr = ciov_ptr?;
429
                 let ciov: types::Ciovec = ciov_ptr.read()?;
430
                 let slice = unsafe {
431
                     let buf = ciov.buf.as_array(ciov.buf_len);
432
                      let raw = buf.as raw(&mut bc)?;
433
434
                     &∗raw
435
                 };
                 slices.push(io::IoSlice::new(slice));
436
             }
437
             let required_rights = EntryRights::from_base(types::Rights::FD_WRITE);
438
             let entry = self.get_entry(fd)?;
439
440
             let isatty = entry.isatty();
             let host_nwritten = entry
441
                                                         OsHandle
                  .as_handle(&required_rights)?
442
                  .write_vectored(&slices, isatty)?
443
444
                  .try_into()?;
             Ok(host nwritten)
445
                                                           io::stdout();
         }
446
```

/crates/wasi-common/src/snapshots/wasi_snapshot_preview1.rs

XXX to/in Wasm

- Rust: rustc
- C/C++
 - Emscripten an LLVM-to-JavaScript/Webassembly compiler. It takes LLVM bitcode which can be generated from C/C++, using llvm-gcc (DragonEgg) or clang, or any other language that can be converted into LLVM - and compiles that into JavaScript or wasm.
 - Cheerp an open-source, commercial C/C++ compiler for Web applications. It can compile virtually any C/C++ code (up to C++14) to WebAssembly, JavaScript, asm.js or a combination thereof.
- Python
 - Pyodide a port of Python to WebAssembly that includes the core packages of the scientific Python stack (Numpy, Pandas, matplotlib). Objects transparently convert and share between Python and Javascript.
 - MicroPython a lean and efficient Python implementation for microcontrollers and constrained systems.
 - Batavia Batavia is an implementation of the Python virtual machine, written in JavaScript. With Batavia, you can run Python bytecode in your browser.

Wasm Runtime

- V8: <u>https://chromium.googlesource.com/v8/v8/+/refs/heads/</u> master/src/wasm/
- Wastime: <u>https://github.com/bytecodealliance/wasmtime</u>
- Lucet: <u>https://github.com/bytecodealliance/lucet</u>
- wasm-micro-runtime: <u>https://github.com/bytecodealliance/</u> wasm-micro-runtime
- Wasmer: <u>https://github.com/wasmerio/wasmer</u>
- Wasmi: <u>https://github.com/paritytech/wasmi</u>

Security

- WebAssembly design doc: <u>https://webassembly.org/docs/</u> <u>security/</u>
- Users
 - Protect <u>users</u> from buggy or malicious modules.
- Developers
 - Provide <u>developers</u> with useful primitives and mitigations for developing safe applications, within the constraints of above one.

Wasm Security Users

- Sandbox (with fault isolation)
- Deterministic execution (with limited exceptions)
- Each module is subject to the security policies of its embedding: same-origin policy or POSIX

- Modules
 - explicit function declaration allow CFI enforcement
 - code is immutable and not observable at runtime
- Function calls must specify the index of a target that corresponds to a valid entry in the function index space or table index space.
- Indirect function calls are subject to a type signature check at runtime; the type signature of the selected indirect function must match the type signature specified at the call site.
- A protected call stack that is invulnerable to buffer overflows in the module heap ensures safe function returns.
- Branches must point to valid destinations within the enclosing function.

- Varibles
 - local variables: zero by default, protected call stack
 - global variables: global index space
 - unclear static scope (address-of op): separate useraddressable stack linear memory, zero by default
 - References to this memory are computed with infinite precision to avoid wrapping and simplify bounds checking

Developers

Traps: terminate execution and signal abnormal behavior to the execution environment

- specifying an invalid index in any index space,
- performing an indirect function call with a mismatched signature,
- exceeding the maximum size of the protected call stack,
- accessing out-of-bounds addresses in linear memory,
- executing an illegal arithmetic operations (e.g. division or remainder by zero, signed division overflow, etc).

- Memory safety (overflows)
 - Buffer overflows <u>cannot affect local or global variables</u> stored in index space, they are fixed-size and addressed by index.
 - Data stored in linear memory <u>can overwrite adjacent objects</u>, since bounds checking is performed at linear memory region granularity and is not context-sensitive.
 - However, the presence of control-flow integrity and <u>protected</u> <u>call stacks prevents direct code injection attacks</u>. Thus, common mitigations such as data execution prevention (DEP) and stack smashing protection (SSP) are not needed.

- Memory safety (use-after-free)
 - the semantics of <u>pointers have been eliminated</u> for function calls and variables with fixed static scope, allowing references to invalid indexes in any index space to trigger a validation error at load time, or at worst a trap at runtime.
 - Accesses to <u>linear memory are bounds-checked</u> at the region level, potentially resulting in a <u>trap</u> at runtime. These memory region(s) are isolated from the internal memory of the runtime, and are set to zero by default unless otherwise initialized.

Developers

- Memory safety (others)
 - possible to hijack the control flow of a module using code reuse attacks against indirect calls
 - ROP attacks are not possible, because control-flow integrity ensures that call targets are valid functions declared at load time.
 - race conditions, such as time of check to time of use (TOCTOU) vulnerabilities, are possible, since no execution or scheduling guarantees are provided beyond in-order execution and atomic memory primitives

side channel attacks

 In the future, additional protections may be provided by runtimes or the toolchain, such as code diversification or memory randomization (ASLR), or bounded pointers ("fat" pointers).

- Control-flow integrity
 - Runtime instrumentation: During compilation, the compiler generates an expected <u>control flow graph</u> of program execution, and <u>inserts runtime instrumentation</u> at each call site to verify that the transition is safe.
 - Sets of <u>expected call targets</u> are constructed from the set of all possible call targets in the program, unique identifiers are assigned to each set, and the instrumentation checks whether the current call target is a member of the expected call target set.
 - If this check succeeds, then the original call is allowed to proceed, otherwise a failure handler is executed, which typically terminates the program.

- 1. Direct function calls,
- 2. Indirect function calls,
- 3. Returns.

- Control-flow integrity
 - In WebAssembly, the execution semantics implicitly guarantee the safety of (1) through usage of <u>explicit function section</u> <u>indexes</u>, and (3) through a <u>protected call stack</u>. Additionally, the <u>type signature</u> of indirect function calls is already checked at runtime, effectively implementing coarse-grained typebased control-flow integrity for (2). All of this is achieved without <u>explicit runtime instrumentation</u> in the module.
 - However, as discussed previously, this protection <u>does not</u> <u>prevent code reuse attacks</u> with function-level granularity against indirect calls.

- 1. Direct function calls,
- 2. Indirect function calls,
- 3. Returns.

- Control-flow integrity
 - Clang/LLVM -fsanitize=cfgi
 - better defend against code reuse attacks that leverage indirect function calls (2)
 - enhances the <u>built-in function signature checks</u> by operating at the C/C++ type level, which is semantically richer that the WebAssembly type level, which consists of only four value types.
 - Currently, enabling this feature has a <u>small performance cost</u> for each indirect call, because an integer range check is used to verify that the target index is trusted, but this will be eliminated in the future by leveraging built-in support for multiple indirect tables with homogeneous type in WebAssembly.

Reference

- Frontend revolution: <u>https://github.com/ManzDev/frontend-evolution</u>
- State of JS: <u>https://2019.stateofjs.com/overview/</u>
- Creating and working with WebAssembly modules: <u>https://hacks.mozilla.org/2017/02/</u> creating-and-working-with-webassembly-modules/
- https://github.com/mbasso/awesome-wasm
- Bringing the Web up to Speed with WebAssembly: https://people.mpi-sws.org/ ~rossberg/papers/Haas,%20Rossberg,%20Schuff,%20Titzer,%20Gohman, %20Wagner,%20Zakai,%20Bastien,%20Holman%20- %20Wagner,%20Gohman, %20Wagner,%20Zakai,%20Bastien,%20Holman%20- %20Wagner,%20Zakai,%20Bastien,%20Holman%20- %20Wagner,%20Zakai,%20Bastien,%20Holman%20- %20Bringing%20the%20Web%20up%20to%20Speed%20with%20WebAssembly.pd f
- Awesome WebAssembly Languages: <u>https://github.com/appcypher/awesome-wasm-langs</u>
- Introduction to WebAssembly: <u>https://rsms.me/wasm-intro</u>
- WebAssembly Out of Bounds Trap Handling: <u>https://docs.google.com/document/d/</u> <u>17y4kxuHFrVxAiuCP_FFtFA2HP5sNPsCD10KEx17Hz6M/edit#</u>