

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49

MSWasm: Soundly Enforcing Memory-Safe Execution of Unsafe Code

ALEXANDRA E. MICHAEL*, UCSD, USA
ANITHA GOLLAMUDI*, University of Massachusetts, Lowell, USA
JAY BOSAMIYA, CMU, USA
EVAN JOHNSON, UCSD, USA
AIDAN DENLINGER, UCSD, USA
CRAIG DISSELKOEN, UCSD, USA
CONRAD WATT, University of Cambridge, UK
BRYAN PARNO, CMU, USA
MARCO PATRIGNANI, University of Trento, Italy
MARCO VASSENA, Utrecht University, Netherlands
DEIAN STEFAN, UCSD, USA

Most programs compiled to WebAssembly (Wasm) today are written in unsafe languages like C and C++. Unfortunately, memory-unsafe C code remains unsafe when compiled to Wasm—and attackers can exploit buffer overflows and use-after-frees in Wasm almost as easily as they can on native platforms. Memory-Safe WebAssembly (MSWasm) proposes to extend Wasm with language-level memory-safety abstractions to precisely address this problem. In this paper, we build on the original MSWasm position paper to realize this vision. We give a precise and formal semantics of MSWasm, and prove that well-typed MSWasm programs are, by construction, robustly memory safe. To this end, we develop a novel, language-independent memory-safety property based on *colored* memory locations and pointers. This property also lets us reason about the security guarantees of a formal C-to-MSWasm compiler—and prove that it always produces memory-safe programs (and preserves the semantics of safe programs). We use these formal results to then guide several implementations: Two compilers of MSWasm to native code, and a C-to-MSWasm compiler (that extends Clang). Our MSWasm compilers support different enforcement mechanisms, allowing developers to make security-performance trade-offs according to their needs. Our evaluation shows that on the PolyBenchC suite, the overhead of enforcing memory safety in software ranges from 22% (enforcing spatial safety alone) to 198% (enforcing full memory safety), and 51.7% % when using hardware memory capabilities for spatial safety and pointer integrity.

More importantly, MSWasm’s design makes it easy to swap between enforcement mechanisms; as fast (especially hardware-based) enforcement techniques become available, MSWasm will be able to take advantage of these advances almost for free.

In the following, we use syntax highlighting accessible to both colourblind and black & white readers [Patrignani 2020]. Specifically, we use a blue, sans-serif font for C and a red, bold font for MSWasm.

1 INTRODUCTION

WebAssembly (Wasm) is a new bytecode designed to run native applications—e.g., applications written in C/C++ and Rust—at native speeds, everywhere—from the Web, to edge clouds, and IoT

*These authors contributed equally to this work

Authors’ addresses: Alexandra E. Michael*, UCSD, USA; Anitha Gollamudi*, University of Massachusetts, Lowell, USA; Jay Bosamiya, CMU, USA; Evan Johnson, UCSD, USA; Aidan Denlinger, UCSD, USA; Craig Disselkoen, UCSD, USA; Conrad Watt, University of Cambridge, UK; Bryan Parno, CMU, USA; Marco Patrignani, University of Trento, Italy; Marco Vassena, Utrecht University, Netherlands; Deian Stefan, UCSD, USA;.

50 platforms. Unlike most industrial bytecode and compiler intermediate representations, Wasm was
51 designed with safety in mind: Wasm programs run in an isolated sandbox *by construction*. On the
52 Web, this means that Wasm programs cannot read or corrupt the browser’s memory [Haas et al.
53 2017a]. On edge clouds, where Wasm programs written by different clients run in a single process,
54 this means that one client cannot interfere with another [McMullen 2020].

55 Within the sandbox, however, Wasm offers little protection. Programs written in unsafe languages—
56 and two thirds of existing Wasm programs are compiled from C/C++ [Hilbig et al. 2021]—remain
57 unsafe when compiled to Wasm [Lehmann et al. 2020]. Indeed, buffer overflows and use-after-free
58 vulnerabilities are as easy to exploit in Wasm as they are natively; sometimes even *easier* (e.g.,
59 because Wasm lacks abstractions like read-only memory). Worse, attackers can use such exploits
60 to confuse the code hosting Wasm into performing unsafe actions—to effectively bypass the Wasm
61 sandbox. [Lehmann et al. 2020], for example, show how attackers can turn a buffer overflow vul-
62 nerability in the Libpng image processing library (executing in a Wasm sandbox) into a cross-site
63 scripting (XSS) attack.

64 To prevent such attacks, C/C++ compilers would have to insert memory-safety checks *before*
65 compiling to Wasm—e.g., to ensure that pointers are valid, within bounds, and point to memory
66 that has not been freed [Nagarakatte et al. 2009, 2010; Necula et al. 2005]. Industrial compilers like
67 Emscripten and Clang do not. Also, they *should not*. Retrofitting programs to enforce memory safety
68 gives up on *robustness*, i.e., preserving memory safety when linking a (retrofitted) memory-safe
69 module with a potentially memory-unsafe module. It gives up on *performance*: efficient memory-
70 safety enforcement techniques rely on operating system abstractions (e.g., virtual memory [Dang
71 et al. 2017]), abuse platform-specific details (e.g., encoding bounds information in the (unused)
72 upper bits of an address [Akritidis et al. 2009]), and take advantage of hardware extensions (e.g.,
73 Arm’s pointer authentication and memory tagging extensions [Arm 2019; Liljestrang et al. 2019]).
74 Finally, it also makes it harder to prove that memory safety is preserved end-to-end.

75 With *Memory-Safe WebAssembly*, [Disselkoen et al. 2019] propose to bridge this gap by extending
76 Wasm with language-level memory-safety abstractions. In particular, MSWasm extends Wasm with
77 *segments*, i.e., linear regions of memory that can only be accessed using *handles*. Handles, like
78 CHERI capabilities [Watson et al. 2015], are unforgeable, well-typed pointers—they encapsulate
79 information that make it possible for MSWasm compilers to ensure that each memory access is valid
80 and within the segment bounds. Alas, the MSWasm position paper only outlines this design—they
81 do not give a precise semantics for MSWasm, nor implement or evaluate MSWasm as a memory-safe
82 intermediate representation.

83 This paper builds on this work to realize the vision of MSWasm. We do this via five contributions:

84 **1. Semantics and Memory Safety for MSWasm (Section 3).** Our first contribution is a formal
85 specification of MSWasm as an extension of the Wasm language, type system, and operational
86 semantics. Our semantics give precise meaning to the previous informal design [Disselkoen et al.
87 2019]. Moreover, these semantics allow us to prove that all well-typed MSWasm programs are
88 *robustly memory safe*; i.e., MSWasm programs are memory safe when linked against arbitrary code.

89 **2. Color-based Memory-Safety Monitor (Section 4).** We develop a novel, abstract memory-
90 safety monitor based on *colored* memory locations and pointers, which we use to show that
91 MSWasm is memory safe. Colors abstract away specific mechanisms that MSWasm backends can
92 employ to enforce memory safety. Additionally, they enable reasoning about spatial as well as
93 temporal memory safety, both at the granularity of individual memory objects and within structured
94 objects as well. Furthermore, since our memory-safety monitor is language-independent, we can
95 reason about memory-safety across compilation and establish the soundness of our compiler-based
96 memory-safety enforcement in our next contribution.

99 **3. Sound Compilation from C to MSWasm (Section 5).** Like Wasm, MSWasm is intended to be
100 used as a compilation target from higher-level languages. Hence, our third contribution is a formal
101 C-to-MSWasm compiler, which guarantees memory-safe execution of unsafe code. In particular, we
102 formalize a compiler from a subset of C to MSWasm and prove that the compiler soundly *enforces*
103 memory-safety. Intuitively, this result ensures that memory-safe C programs when compiled to
104 MSWasm remain safe and preserve their semantics, while memory-unsafe C programs trap at the
105 first memory violation (and are thus safe too).

106 **4. Implementations of MSWasm (Section 6).** Our next contribution is the implementation
107 of three MSWasm-related compilers. First, we implement an ahead-of-time (AOT) MSWasm-to-
108 machine code compiler by extending the rWasm [Bosamiya et al. 2022] compiler with 1900 lines of
109 code (LOC). Our extension of rWasm supports multiple options for enforcing memory safety, with
110 tradeoffs between performance and differing levels of memory safety (spatial and temporal safety,
111 and handle integrity). Our second compiler is a just-in-time (JIT) MSWasm-to-JVM compiler (1200
112 LOC), which uses the GraalVM Truffle framework [Oracle 2021b]. Finally, our third compiler is
113 an LLVM-to-MSWasm compiler (1600 LOC) created as an extension of the Cheri Clang compiler
114 toolchain [CTSRD-CHERI 2022].

115 **5. Evaluation of MSWasm (Section 7).** Our final contribution is an empirical evaluation of
116 MSWasm. We benchmark MSWasm on PolyBenchC, the de-facto Wasm benchmarking suite [Pouchet
117 2011]. We find that, on (geomean) average, MSWasm when enforced in software using our AOT
118 compiler imposes an overhead of 197.5%, which is comparable with prior work on enforcing
119 memory safety for C [Nagarakatte et al. 2010]. MSWasm, however, makes it easy to change the
120 underlying enforcement mechanism (e.g., to boost performance), without changing the application.
121 To this end, we find that enforcing just spatial and temporal safety imposes a 52.2% overhead, and
122 enforcing spatial safety alone using a technique similar to Baggy Bounds [Akritidis et al. 2009],
123 is even cheaper—21.4%. Our JIT compiler, which enforces spatial and temporal safety, but not
124 handle integrity, has an overhead of 42.3%. While these overheads are relatively large on today’s
125 hardware, upcoming hardware features explicitly designed for memory-safety enforcement can
126 reduce these overheads (e.g., Arm’s PAC can be used to reduce pointer integrity enforcement to
127 under 20% [Liljestrang et al. 2019], while Arm’s Cheri [Grisenthwaite 2019] or Intel’s CCC [LeMay
128 et al. 2021] can also reduce the cost of enforcing temporal and spatial safety). MSWasm will be
129 able to take advantage of these features as soon they become available, as illustrated by the ease of
130 swapping memory-safety enforcement techniques within our AOT compiler.

131 **Open Source & Technical Report.** Our technical report, implementations, benchmarks, and data
132 sets are available as supplementary material and will be made open source.
133

134 2 BACKGROUND AND MOTIVATION

135 We now give a brief introduction to Wasm (Section 2.1), its attacker model (Section 2.2), and the
136 implications of memory unsafety within the Wasm sandbox (Section 2.3). Then we give a brief
137 introduction to MSWasm and to the open challenges we address in this work (Section 2.4).
138

139 2.1 WebAssembly

140 Wasm is a low-level bytecode, designed as a safe compilation target for higher-level languages like
141 C/C++ and Rust [Haas et al. 2017b]. Wasm bytecode is executed in a sandboxed environment by a
142 stack-based virtual machine. Prior to execution, the virtual machine type-checks the bytecode to
143 ensure that each instruction finds the appropriate operands on the stack. Wasm’s type system is
144 extremely simple; the language has four primitive types—32- and 64-bit integers and floats (**i32** and
145 **i64**, and **f32** and **f64** respectively)—and only structured control flow constructs (i.e., no **gotos**) which
146
147

simplify type checking. The Wasm heap (or *linear memory*), however, is an untyped contiguous linear array of bytes. Instructions `τ.load` and `τ.store` allow values of the four primitive types to be read from and written to the memory at arbitrary integer offsets. At runtime, Wasm ensures that these accesses are in bounds (and *traps* when they are not).

This simple design makes whole classes of attacks impossible by design. For example, the type-system ensures that well-typed bytecode cannot hijack the virtual machine’s control flow via stack-smashing attacks [One 1996]. The coarse-grained bounds-checks on memory accesses, together with structured control flow, confine Wasm to a sandbox [Tan 2017; Wahbe et al. 1993]—and thus prevent Wasm from harming its host environment.

This simple design has a trade-off: We necessarily lose information when compiling programs written in high-level languages to Wasm. Clang, for example, compiles complex source-level values (e.g., structs and arrays) into “bags of bytes” in the untyped linear memory and compiles pointers to Wasm 32-bit integers, offsets in the linear memory where values are layed out. This, unfortunately, means that misusing C/C++ pointers is as simple and severe in Wasm as it is for native platforms.

2.2 Threat Model

In this work, we consider a Wasm-level attacker who attempts to exploit a memory vulnerability present in a C program compiled to Wasm. We consider vulnerabilities that can be triggered by *spatial* memory errors (e.g., buffer overflows), *temporal* memory errors (e.g., use-after-free and double-free vulnerabilities), and *pointer integrity* violations (e.g., corrupting function pointers to bend control flow). We assume the vulnerable program is linked with arbitrary code written by the attacker, which can interact with the program in any way allowed by Wasm semantics. To exploit a vulnerability, the attacker code can supply malicious inputs to the program and abuse values (including pointers) returned by or passed to the program. We leave memory unsafety of C++ programs and type confusion vulnerabilities [Haller et al. 2016] for future work.

2.3 Sandboxing Without Memory Safety

Memory unsafe C programs, when compiled to Wasm, largely remain unsafe: They can run uninterrupted as long as their reads and writes stay within the bounds of the entire linear memory. Unfortunately, Wasm also lacks most mitigations we rely on today to deal with memory unsafety (e.g., memory protection bits and ASLR), so a program compiled to run within Wasm’s sandbox may be more vulnerable than if it were running on bare metal [Lehmann et al. 2020].

```
1 char *trim_token( char *token ){
2   char *trimmed = malloc( 1024 *sizeof(char) );
3   int i = 0, j = 0;
4   while (token[i++] == '\ ');
5   char next = token[--i];
6   while (next != '\0') {
7     trimmed[j++] = next; // Possible buffer overflow
8     next = token[++i];
9   }
10  trimmed[i] = '\0';
11  return trimmed;
12 }
```

Listing 1. Vulnerable code adapted from libpng 1.6.37

To understand how source-level memory vulnerabilities persist across compilation, consider the C code snippet in Listing 1 from Libpng 1.6.37. Function `trim_token` takes a pointer to a null-terminated string as input and returns a pointer to a dynamically-allocated copy of the string,

197 trimmed of the leading whitespace characters. The first loop (Line 4) simply scans the string `token`
 198 and skips all the whitespace characters, while the second loop (Line 6–Line 9) copies the rest of the
 199 string into `trimmed` one character at the time, until it finds the null terminator. The vulnerability is
 200 on Line 7: the length of the string `token` after trimming may exceed the size allocated for buffer
 201 `trimmed`. To exploit this vulnerability, an attacker only needs to call this function on a sufficiently
 202 long string (longer than 1024 characters after trimming). This will cause the function to write past
 203 the bounds of `trimmed`, thus corrupting the memory of the program with the payload supplied by
 204 the attacker. This vulnerability remains in the code obtained by compiling function `trim_token`
 205 with existing Wasm compilers (e.g., Emscripten and Clang). In particular, Line 7 gets translated
 206 into the Wasm instructions in Listing 2.

207 The first three Wasm instructions compute the address (a 32-bit integer) where the next character gets copied, by adding index `$j` to address
 208 `$Strimmed`. Then, instruction `get $next` pushes the value of the next
 209 character on the stack and `i32.store` writes it to the address computed
 210 before. As long as this address is within the linear memory region, the
 211 store instruction succeeds—even if the address does not belong to the
 212 buffer allocated for `$Strimmed`.

```

1 get $Strimmed
2 get $j
3 i32.add
4 get $next
5 i32.store
6 ... ;; increment $j
  
```

Listing 2. "Compilation of Line 7 into Wasm."

214 Although an attacker could not use this memory-safety vulnerability
 215 to escape Wasm’s sandbox, they could use it to corrupt and steal data
 216 (e.g., private keys) sensitive to the Wasm program itself. Wasm programs
 217 on the Web already handle sensitive data, and as Wasm’s adoption expands beyond the Web,
 218 addressing memory safety within the sandbox is crucial.

2.4 The MSWasm Proposal

221 Memory-Safe WebAssembly (MSWasm) addresses these challenges by extending Wasm with ab-
 222 stractions for enforcing memory safety [Disselkoen et al. 2019]. Specifically, MSWasm introduces a
 223 new memory region called *segment memory*. The segment memory consists of individual *segments*,
 224 which are linearly addressable, bounded regions of memory representing dynamic memory alloca-
 225 tions. Unlike Wasm’s linear memory, the segment memory cannot be accessed at arbitrary offsets
 226 through standard `load` and `store` instructions. Instead, MSWasm provides new types, values, and
 227 instructions to regulate access to segments and enforce per-allocation memory safety. Segments can
 228 only be accessed through *handles*, unforgeable memory capabilities that model pointers bounded
 229 to a particular allocation of the segment memory. MSWasm adopts this low-level memory model
 230 since an object-based model (like that of the JVM) would be an inefficient (due to garbage collection
 231 overhead) and overly restrictive (due to the constraints of an object-based type-system) compilation
 232 target for C code deployed to Wasm.

233 Handles are tuples: $\langle \text{base}, \text{offset}, \text{bound}, \text{isCorrupted}, \text{id} \rangle$, where `base` represents the beginning
 234 of the segment in segment memory, `offset` is the handle’s offset *within* the segment, i.e., within
 235 the `bound`, that the handle points to. Thus, a handle points to the address given by `base + offset`.
 236 MSWasm guarantees *handle integrity* using the `isCorrupted` flag. Intuitively, attempts to forge
 237 handles (e.g., by casting an integer, or altering the bitstring representation of an existing handle
 238 in memory) result in a corrupted handle. MSWasm traps only when an out-of-bounds or corrupt
 239 handle is *used*, not when it is created. This improves both performance, by eliminating checks on
 240 every pointer-arithmetic operation, and compatibility, since many C idioms create benign out-of-
 241 bound pointers [Memarian et al. 2019a, 2016; Ruef et al. 2019]. Finally, MSWasm associates each
 242 segment allocation with a unique identifier `id`, which is used to enforce *temporal* memory safety.

243 MSWasm provides new instructions to create and manipulate handles, and to access segments
 244 safely through them. Instructions `τ.segload` and `τ.segstore` are analogous to `τ.load` and `τ.store`,

245

but operate on handles and trap if the handle is corrupted or points outside the segment bounds, or if the segment has been freed. Instruction `segalloc` allocates a segment of the desired size in a free region of segment memory and returns a handle to it. Instruction `segfree` frees the segment associated with a valid handle, thus making that region of segment memory available for new allocations. Lastly, instruction `handle.add` is for pointer arithmetic and modifies the handle offset, without changing the base or bound.

Eliminating unsafety by compiling C to MSWasm. With MSWasm we can eliminate potential memory vulnerabilities automatically, via compilation. For example, a C to MSWasm compiler would emit the instructions in Listing 3 for the code snippet from Listing 1. This code allocates a new 1024-byte segment and stores the handle for it in variable `$trimmed`. Then, the `handle.add` instruction increments the offset of `$trimmed` with index `$j` and instruction `i32.segment_store` writes `$next` in the segment. Since MSWasm instructions enforce memory safety, this code is safe to execute even with malicious inputs. In particular, if the offset of `$trimmed` is incremented past the bound of the handle, the store instruction simply traps, thus preventing the buffer overflow.

```
1 i32.const 1024
2 segalloc
3 set $trimmed
4 ...
5 get $trimmed
6 get $j
7 handle.add
8 get $next
9 i32.segment_store
10 ...
```

Listing 3. "Compilation of Listing 1 into Wasm."

Enforcing Intra-Object Memory Safety. Through the abstractions described above, MSWasm enforces *inter-object* memory safety, i.e., at the granularity of individual allocations. Unfortunately, this alone is insufficient to prevent memory-safety violations within composite data types (e.g., structs), in which a pointer to a field overflows (or overruns) an adjacent field.

```
1 struct User {char name[32], int id };
2 struct User *my_user = malloc(sizeof(struct User));
3 char *my_name = my_user->name;
4 ...
```

Listing 4. Intra-object memory safety vulnerability.

Consider the code snippet in Listing 4, which defines a struct object containing a fixed-length string `name` and an integer user `id`. When compiled to MSWasm, this code allocates a single segment for the `User` structure; thus the handle corresponding to `my_name` and derived from `my_user` via pointer arithmetic can also access field `id` without trapping. Therefore, an attacker could exploit a memory vulnerability in the code that manipulates `my_name` to corrupt the user `id` and impersonate another user.

Hence, to enforce *intra-object* memory safety, MSWasm provides an additional instruction called `slice`. Instruction `slice` shrinks the portion of the segment that a handle can access by growing its `base` and reducing its `bound` field by a given offset. By emitting a `slice` instruction with appropriate offsets for expression `my_user->name`, a compiler can generate a sliced handle that includes only the `name` field. As a result, if the attacker later tries to overflow `my_name`, the safety checks of the sliced handle will detect a violation and trap the execution, thus preventing the program from corrupting the user `id`.

The missing pieces. The original MSWasm position paper [Disselkoen et al. 2019] only outlines the basic abstractions we describe above. The position paper does not give a formal (or even informal) semantics for the proposed language extensions. They do not describe compilation techniques—how one would compile C code to MSWasm or how MSWasm would be compiled to native code—nor an implementation (and thus evaluation) of MSWasm. In this paper we address these limitations

```

295     Modules  $M ::= \{\text{funcs } \Phi^*, \text{imports } \rho^*, \text{heap } n_H, \text{segment } n_S\}$ 
296
297     Fun. Defs  $\Phi ::= \{\text{var } \tau^*, \text{body } i^*\} : \rho$    Instructions Types  $\rho ::= \tau^* \rightarrow \tau^*$ 
298
299     Value Types  $\tau ::= i32 \mid i64 \mid f32 \mid f64 \mid \text{handle}$ 
300
301     Instructions  $i ::= \text{nop} \mid \text{trap} \mid \tau.\text{const } c \mid \tau.\otimes \mid \text{get } n \mid \text{set } n \mid \tau.\text{load} \mid \tau.\text{store} \mid \text{branch } i^* \ i^* \mid \text{call } n$ 
302      $\mid \text{return} \mid \tau.\text{segload} \mid \tau.\text{segstore} \mid \text{slice} \mid \text{segalloc} \mid \text{handle.add} \mid \text{segfree}$ 
303
304
305
306
307
308
309
310
311

```

Fig. 1. Syntax of MSWasm with extensions to Wasm highlighted.

and, for the first time, provide an end-to-end, robust, memory-safe C-to-MSWasm compiler that is rooted in formal methods.

3 THE MSWASM LANGUAGE

This section develops a formal model of the design of MSWasm described above. The model includes syntax, typing (Section 3.1), and operational semantics for MSWasm (Section 3.2) and it serves as a specification for different low-level mechanisms (bounds checks, segment identifiers, integrity tags, etc.) needed to enforce memory safety in MSWasm. We present the properties of MSWasm in the next section, after formally defining memory safety.

Due to space constraints, we present a selection of the formalization, and elide proofs and auxiliary lemmas. The interested reader can find these omissions in the supplementary material.

3.1 MSWasm Syntax

The syntax of MSWasm is defined in Figure 1. MSWasm programs are modules M , which specify a list of function definitions Φ^* , the type of imported functions ρ^* , and the size of the linear and segment memory (n_H and $n_S \in \mathbb{N}$).¹ Syntax $\{\text{var } \tau^*, \text{body } i^*\} : \rho$ defines a function with local variables of types τ^* , body i^* , and function type ρ . Instructions i manage the operand stack and are mostly standard. Variables are referred to through numeric indices n , which are statically validated during type-checking. For example, instructions `get n` and `set n` retrieve and update the value of the n -th local variable, respectively. Function calls are similar, i.e., instruction `call n` calls the n -th function (either defined or imported) in the scope of the module. We describe MSWasm's instructions on segments and handles below.

Typing. The type system of MSWasm is a straightforward extension of Wasm's, and it similarly guarantees type safety (i.e., well-typed modules satisfy progress and preservation). Instructions are typed by the judgment $\Gamma \vdash i : \tau_1^* \rightarrow \tau_2^*$, where τ_1^* and τ_2^* are the types of the values that i pops and pushes on the stack, respectively, and the typing context Γ tracks the type of the variables and functions in scope. Compared to Wasm, the only restriction imposed by the type system of MSWasm is that it prevents programs from forging handles by reading raw bytes from the unmanaged linear memory, i.e., $\Gamma \vdash \tau.\text{load} : [i32] \rightarrow [\tau]$ iff $\tau \neq \text{handle}$.

¹We use e^* to denote a list of e elements, and e^n for a list of length n . We write $[e_0, e_1, \dots]$ for finite lists, $[]$ for the empty list, $e : e^*$ to add e in front of e^* , and $e_1^* \# e_2^*$ to append e_2^* to e_1^* . Notation $e^*[i]$ looks up the i -th element of e^* and $e^*[i \mapsto v]$ replaces the i -th element of e^* with v .

344 Store $\Sigma ::= (\mathbf{H}, \mathbf{T}, \mathbf{A})$ Heaps $\mathbf{H} ::= \mathbf{b}^*$ Bytes $\mathbf{b} \in \{0..2^8-1\}$
 345 Constants \mathbf{c} Segments $\mathbf{T} ::= (\mathbf{b}, \mathbf{t})^*$ Allocators \mathbf{A}
 346 Local Frames $\mathbf{F} ::= (\theta, \mathbf{i}^*, \mathbf{v}^*)$ Locals $\theta ::= (\mathbf{n} \mapsto \mathbf{v})^*$ Segment Tags $\mathbf{t} ::= \circ \mid \square$
 347 Values $\mathbf{v} ::= \mathbf{c} \mid \mathbf{h}$ Handles $\mathbf{h} ::= \langle \mathbf{n}_{\text{base}}, \mathbf{n}_{\text{offset}}, \mathbf{n}_{\text{bound}}, \mathbf{b}_{\text{valid}}, \mathbf{n}_{\text{id}} \rangle$
 348 Memory Events $\alpha ::= \epsilon \mid \text{read}_\tau(\mathbf{h}) \mid \text{write}_\tau(\mathbf{h}) \mid \text{salloc}(\mathbf{h}) \mid \text{sfree}(\mathbf{h}) \mid \text{trap}$
 349
 350
 351
 352
 353

Fig. 2. Wasm and MSWasm runtime structures.

$$\begin{array}{c}
 \text{(Stack-Top)} \\
 \Phi^* \vdash \langle \Sigma, (\theta, \mathbf{i}, \mathbf{v}^*) \rangle \xrightarrow{\alpha} \langle \Sigma', (\theta', \mathbf{i}', \mathbf{v}'^*) \rangle \\
 \hline
 \Phi^* \vdash \langle \Sigma, (\theta, \mathbf{i} : \mathbf{i}^*, \mathbf{v}^* + \mathbf{v}_{\mathbf{b}}^*) \rangle \xrightarrow{\alpha} \langle \Sigma, (\theta, \mathbf{i}'^* + \mathbf{i}^*, \mathbf{v}'^* + \mathbf{v}_{\mathbf{b}}^*) \rangle \\
 \text{(Get)} \\
 \theta[\mathbf{n}] = \mathbf{v} \\
 \hline
 \Phi^* \vdash \langle \Sigma, (\theta, \text{get } \mathbf{n}, []) \rangle \rightarrow \langle \Sigma, (\theta, [], [\mathbf{v}]) \rangle \\
 \text{(If-T)} \\
 \mathbf{n} \neq \mathbf{0} \\
 \hline
 \Phi^* \vdash \langle \Sigma, (\theta, \text{branch } \mathbf{i}_1^* \ \mathbf{i}_2^*, [\mathbf{n}]) \rangle \rightarrow \langle \Sigma, (\theta, \mathbf{i}_1^*, []) \rangle \\
 \text{(Load)} \\
 \mathbf{0} \leq \mathbf{n} \quad \mathbf{n} + |\tau| < |\mathbf{H}| \quad \mathbf{b}^{|\tau|} = [\Sigma.\mathbf{H}[\mathbf{n} + \mathbf{j}] \mid \mathbf{j} \in \{0..|\tau| - 1\}] \quad \mathbf{v} = \tau.\text{unpack}(\mathbf{b}^{|\tau|}) \\
 \hline
 \Phi^* \vdash \langle \Sigma, (\theta, \tau.\text{load}, [\mathbf{n}]) \rangle \rightarrow \langle \Sigma, (\theta, [], [\mathbf{v}]) \rangle
 \end{array}$$

Fig. 3. Semantics of Wasm (excerpts).

3.2 MSWasm Operational Semantics

To reason about the memory-safety guarantees of MSWasm, we define a small-step labeled operational semantics, which generates events for memory-relevant operations such as segment allocations and accesses.

3.2.1 Semantics of Wasm. Figure 2 defines the runtime structures used in the semantics judgment. A local configuration $\langle \Sigma, \mathbf{F} \rangle$ consists of the store Σ and the stack frame \mathbf{F} of the function currently executing. In Wasm, the store Σ contains only the unmanaged linear memory \mathbf{H} , which is a list of bytes \mathbf{b}^* of fixed length. The local stack frame \mathbf{F} maintains the environment θ for variable bindings (mapping from variable indices to values), a list of instructions \mathbf{i}^* to be executed, and the operand stack \mathbf{v}^* for the values produced (and consumed) by those instructions. Values include constants \mathbf{c} and integers \mathbf{n} .

The semantics judgment $\Phi^* \vdash \langle \Sigma, \mathbf{F} \rangle \xrightarrow{\alpha} \langle \Sigma', \mathbf{F}' \rangle$ indicates that under function definitions Φ^* , local configuration $\langle \Sigma, \mathbf{F} \rangle$ executes a single instruction and steps to $\langle \Sigma', \mathbf{F}' \rangle$, generating event α (explained below). The semantics features also a separate judgment for function calls and returns, which is standard and omitted.

Figure 3 presents a selection of rules that MSWasm inherits from Wasm. Auxiliary rule (Stack-Top) extracts the first instruction and its operands from the list of instructions and the stack, respectively, and executes the instruction using the rules for individual instructions. Rule (Get) executes instruction `get n`, which looks up the value of variable \mathbf{n} in the environment θ , i.e.,

393 $\theta[\mathbf{n}] = \mathbf{v}$, and pushes \mathbf{v} on the stack. Instruction **branch** $i_1^* i_2^*$ pops the integer condition \mathbf{n} from
 394 the stack and returns instructions i_1^* if \mathbf{n} is non-zero via rule (If-T).² Rule (Load) loads a value of
 395 type τ from address \mathbf{n} in linear memory. Since the linear memory consists of plain bytes, the rule
 396 reads $|\tau|$ bytes at address \mathbf{n} into byte string $\mathbf{b}^{|\tau|}$ and converts them into a value of type τ , i.e., $\mathbf{v} =$
 397 $\tau.\text{unpack}(\mathbf{b}^{|\tau|})$, which is then pushed on the stack. In the rule, premises $0 \leq \mathbf{n}$ and $\mathbf{n} + |\tau| < |\mathbf{H}|$
 398 ensure that the load instruction does not read outside the bounds of the linear memory, but do not
 399 enforce memory safety, as explained above.

400
 401 **3.2.2 Semantics of MSWasm.** MSWasm extends the runtime structures of Wasm with a managed
 402 segment memory, handle values, and a memory allocator (see Figure 2). The segment memory \mathbf{T} is
 403 a fixed-length list $\langle \mathbf{b}, \mathbf{t} \rangle^*$ of *tagged* bytes, where each tag \mathbf{t} indicates whether the corresponding
 404 byte is part of a numeric value ($\mathbf{t} = \circ$) or a handle ($\mathbf{t} = \square$). These tags are used to detect forged or
 405 corrupted handles stored in segment memory and thus ensure handle integrity.

406 Handles $\langle \mathbf{n}_{\text{base}}, \mathbf{n}_{\text{offset}}, \mathbf{n}_{\text{bound}}, \mathbf{b}_{\text{valid}}, \mathbf{n}_{\text{id}} \rangle$ contain the base address \mathbf{n}_{base} of the memory region
 407 they span, length $\mathbf{n}_{\text{bound}}$, offset $\mathbf{n}_{\text{offset}}$ from the base, integrity flag $\mathbf{b}_{\text{valid}}$ which indicates whether
 408 the handle is authentic ($\mathbf{b}_{\text{valid}} = \text{true}$) or corrupted ($\mathbf{b}_{\text{valid}} = \text{false}$), and segment identifier \mathbf{n}_{id} . Finally,
 409 MSWasm instructions generate memory events α , which include the silent event ϵ , reading and
 410 writing values of type τ through a handle \mathbf{h} (i.e., $\text{read}_\tau(\mathbf{h})$ and $\text{write}_\tau(\mathbf{h})$), segment allocations
 411 $\text{salloc}(\mathbf{h})$, segment free $\text{sfree}(\mathbf{h})$, and **trap** which is raised in response to a memory violation.

412 **Memory Allocator.** The MSWasm runtime system is responsible for providing a memory allocator
 413 to serve memory allocations of compiled programs. In our model, we represent the state of the
 414 memory allocator and its semantics explicitly, as this simplifies reasoning about memory safety.
 415 The allocator state \mathbf{A} keeps track of free and used regions of segment memory and their identifiers,
 416 i.e., $\mathbf{A}.\text{free}$ and $\mathbf{A}.\text{allocated}$, respectively. The allocator serves allocation requests via reductions
 417 of the form $\langle \mathbf{T}, \mathbf{A} \rangle \xrightarrow{\text{salloc}(\mathbf{a}, \mathbf{n}, \mathbf{n}_{\text{id}})} \langle \mathbf{T}', \mathbf{A}' \rangle$, which allocates and initializes a free segment of \mathbf{n} bytes,
 418 which starts at address \mathbf{a} in segment memory and can be identified by fresh identifier \mathbf{n}_{id} . Dually,
 419 reductions of the form $\langle \mathbf{T}, \mathbf{A} \rangle \xrightarrow{\text{sfree}(\mathbf{a}, \mathbf{n}_{\text{id}})} \langle \mathbf{T}', \mathbf{A}' \rangle$ free the segment identified by \mathbf{n}_{id} and allocated
 420 at address \mathbf{a} , or traps, if no such segment is currently allocated at that address. We omit further
 421 details about the allocator state and semantics—the memory-safety guarantees of MSWasm do not
 422 depend on the concrete allocation strategy.

423
 424 **MSWasm Rules.** Figure 4 gives some important rules for the new instructions of MSWasm.
 425 Rule (H-Load) loads a non-handle value ($\tau \neq \text{handle}$) from segment memory through a *valid* handle
 426 $\langle \mathbf{n}_1, \mathbf{o}, \mathbf{n}_2, \text{true}, \mathbf{n}_{\text{id}} \rangle$. Rule (H-Load) reads bytes \mathbf{b}^* from the address pointed to by the handle, i.e.,
 427 $\mathbf{n} = \mathbf{n}_1 + \mathbf{o}$, and converts them into a value of type τ , i.e., $\mathbf{v}_2 = \tau.\text{unpack}(\mathbf{b}^*)$.³ The rule enforces
 428 memory safety by checking that (1) the handle is *not* corrupted, (2) the load does not read bytes
 429 outside the bounds of the segment, i.e., $0 \leq \mathbf{o}$ and $\mathbf{o} + |\tau| < \mathbf{n}_2$, and (3) the segment is still allocated,
 430 i.e., $\mathbf{n}_{\text{id}} \in \mathbf{A}.\text{allocated}$. Rule (H-Load-Handle) is similar, but for loading values of type **handle**;
 431 therefore it includes additional checks (highlighted in gray), to enforce handle integrity. First, the
 432 rule checks that all the bytes read from memory are tagged as handle bytes, i.e., $\mathbf{b}_c = \bigwedge_{t \in \tau} (t = \square)$,
 433 and then combines this flag with the flag \mathbf{b}'_c obtained from the raw bytes of the segment; i.e., it
 434 returns handle $\langle \mathbf{n}'_1, \mathbf{a}', \mathbf{n}'_2, \mathbf{b}_c \wedge \mathbf{b}'_c, \mathbf{n}_{\text{id}} \rangle$. The combined flag invalidates handles obtained from bytes
 435 tagged as data, thus preventing programs from forging handles by altering their byte representation
 436

437 ²Wasm does not provide instructions for *unstructured control-flow*, common on native architectures (e.g., **JMP** on x86). Wasm
 438 code can define and jump to typed *labeled* blocks. Since these features do not affect the memory safety guarantees of
 439 MSWasm, we omit them from our model.

440 ³Total function $\tau.\text{unpack}$ converts $|\tau|$ bytes (the number of bytes needed to represent a value of type τ) into a value of type
 441 τ . The inverse function $\tau.\text{pack}$ converts values to their byte representation.

$$\begin{array}{c}
 \text{(H-Load)} \\
 \tau \neq \text{handle} \quad v_1 = \langle n_1, o, n_2, \text{true}, n_{\text{id}} \rangle \quad 0 \leq o \quad (b^*, _) = [\Sigma.T[n+j] \mid j \in \{0..|\tau|-1\}] \\
 o + |\tau| < n_2 \quad n = n_1 + o \quad n_{\text{id}} \in \Sigma.A.\text{allocated} \quad v_2 = \tau.\text{unpack}(b^*) \\
 \hline
 \Phi^* \vdash \langle \Sigma, (\theta, \tau.\text{segload}, [v_1]) \rangle \xrightarrow{\text{read}_\tau(v_1)} \langle \Sigma, (\theta, [], [v_2]) \rangle \\
 \text{(H-Load-Handle)} \\
 \tau = \text{handle} \quad v_1 = \langle n_1, o, n_2, \text{true}, n_{\text{id}} \rangle \quad 0 \leq o \quad (b^*, t^*) = [\Sigma.T[n+j] \mid j \in \{0..|\tau|-1\}] \\
 o + |\tau| < n_2 \quad n = n_1 + o \quad n_{\text{id}} \in \Sigma.A.\text{allocated} \quad n \% |\text{handle}| = 0 \\
 b_c = \bigwedge_{t \in \tau^*} (t = \square) \quad \tau.\text{unpack}(b^*) = \langle n'_1, a', n'_2, b'_c, n'_{\text{id}} \rangle \quad v_2 = \langle n'_1, a', n'_2, b_c \wedge b'_c, n'_{\text{id}} \rangle \\
 \hline
 \Phi^* \vdash \langle \Sigma, (\theta, \tau.\text{segload}, [v_1]) \rangle \xrightarrow{\text{read}_\tau(v_1)} \langle \Sigma, (\theta, [], [v_2]) \rangle \\
 \text{(H-Store)} \\
 \tau \neq \text{handle} \quad v_1 = \langle n_1, o, n_2, \text{true}, n_{\text{id}} \rangle \quad o \geq 0 \quad o + |\tau| < n_2 \quad n_{\text{id}} \in \Sigma.A.\text{allocated} \\
 b^* = \tau.\text{pack}(v_2) \quad t = \square \quad a = (n_1 + o) \quad \Sigma'.T = \Sigma.T[a+j \mapsto (b_j, t) \mid j \in \{0..|\tau|-1\}] \\
 \hline
 \Phi^* \vdash \langle \Sigma, \theta, \tau.\text{segstore}, [v_2, v_1] \rangle \xrightarrow{\text{write}_\tau(v_1, v_2)} \langle \Sigma', \theta, [], [] \rangle \\
 \text{(H-Store-Handle)} \\
 \tau = \text{handle} \quad v_1 = \langle n_1, o, n_2, \text{true}, n_{\text{id}} \rangle \quad o \geq 0 \quad o + |\tau| < n_2 \\
 n_{\text{id}} \in \Sigma.A.\text{allocated} \quad b^* = \tau.\text{pack}(v_2) \quad t = \square \\
 a = n_1 + o \quad a \% |\text{handle}| = 0 \quad \Sigma'.T = \Sigma.T[a+j \mapsto (b_j, t) \mid j \in \{0..|\tau|-1\}] \\
 \hline
 \Phi^* \vdash \langle \Sigma, \theta, \tau.\text{segstore}, [v_2, v_1] \rangle \xrightarrow{\text{write}_\tau(v_1, v_2)} \langle \Sigma', \theta, [], [] \rangle \\
 \text{(H-Alloc)} \\
 \Sigma = (H, T, A) \quad \langle T, A \rangle \xrightarrow{\text{salloc}(a, n, n_{\text{id}})} \langle T', A' \rangle \quad v = \langle a, 0, n, \text{true}, n_{\text{id}} \rangle \quad \Sigma' = (H, T', A') \\
 \hline
 \Phi^* \vdash \langle \Sigma, (\theta, \text{segalloc}, [n]) \rangle \xrightarrow{\text{salloc}(v)} \langle \Sigma', (\theta, [], [v]) \rangle \\
 \text{(H-Free)} \\
 \Sigma = (H, T, A) \quad h = \langle a, 0, _, \text{true}, n_{\text{id}} \rangle \quad \langle T, A \rangle \xrightarrow{\text{sfree}(a, n_{\text{id}})} \langle T', A' \rangle \quad \Sigma' = (H, T', A') \\
 \hline
 \Phi^* \vdash \langle \Sigma, \theta, \text{segfree}, [h] \rangle \xrightarrow{\text{sfree}(h)} \langle \Sigma', \theta, [], [] \rangle \\
 \text{(Handle-Add)} \\
 v = \langle n_1, o, n_2, b, n_{\text{id}} \rangle \quad v' = \langle n_1, o + n, n_2, b, n_{\text{id}} \rangle \\
 \hline
 \Phi^* \vdash \langle \Sigma, (\theta, \text{handle.add}, [n, v]) \rangle \rightarrow \langle \Sigma, (\theta, [], [v']) \rangle \\
 \text{(Slice)} \\
 v = \langle n_1, o, n_2, b, n_{\text{id}} \rangle \quad 0 \leq o_1 < n_2 \quad 0 \leq o_2 \quad v' = \langle n_1 + o_1, o, n_2 - o_2, b, n_{\text{id}} \rangle \\
 \hline
 \Phi^* \vdash \langle \Sigma, (\theta, \text{slice}, [v, o_2, o_1]) \rangle \rightarrow \langle \Sigma, (\theta, [], [v']) \rangle
 \end{array}$$

Fig. 4. Semantics of MSWasm (excerpts). The premises that ensure handle integrity are highlighted.

in memory. Furthermore, to enforce handle integrity, the rule allows loading handle values only from `|handle|`-aligned memory addresses, i.e., $(n_1 + o) \% |\text{handle}| = 0$. The alignment requirement is needed to avoid crafting fake handles. In fact, if one were to store two handles next to each other and then load from an address *within* the first one, the load would succeed and load bytes that all have the capability tag. However, the loaded value would be a fake capability, since the loaded bytes would be part of the first capability, and part of the second. Loading and storing at aligned addresses prevents this issue. The rules for `τ .segstore` are analogous—they include similar

491 bounds checks and alignment restrictions for handles—and additionally set the tag of the bytes that
 492 they write in memory according to τ . For example, Rule (H-Store) applies to values whose type
 493 are not handle, therefore it tags the bytes of the value written to memory as data (\circ). In contrast,
 494 Rule (H-Store-Handle) writes a **handle** to memory and so it tags its bytes accordingly (i.e., \square).

495 Rule (H-Alloc) invokes the allocator to allocate and initialize a new segment of n bytes at address
 496 a in segment memory, and returns a handle to it. Rule (H-Free) invokes the allocator to free the
 497 segment bound to the given *valid* handle. Rule (Handle-Add) increments the offset of a handle
 498 v , without changing the other fields. Notice that this rule allows programs to create handles that
 499 point out of bounds; out-of-bounds handles only cause a trap when they are used to access memory.
 500 Rule (Slice) creates a sliced handle $\langle n_1 + o_1, o, n_2 - o_2, b, n_{id} \rangle$, where the base is increased by offset
 501 o_1 and the bound is reduced by offset o_2 . Premises $0 \leq o_1 < n_2$ and $0 \leq o_2$ ensure that the handle
 502 obtained after slicing can only access a subset of the segment accessible from the original handle.

503 Whenever a τ .segload, a τ .segstore, a *segfree*, or a *slice* do not match their premise, the semantics
 504 traps, emitting a **trap** action and halting the execution immediately, with no values on the operand
 505 stack (omitted for brevity).

506

507 4 ABSTRACT MEMORY-SAFETY MONITOR

508 This section presents an abstract notion of memory safety that is based on *colored* memory locations
 509 (Section 4.1). Colors soundly abstract away many implementation details, which in turn let us
 510 formalize memory safety compactly as a trace property checked by a corresponding monitor
 511 (Section 4.2). We use this monitor to establish the spatial and temporal memory-safety guarantees
 512 of MSWasm (Section 4.3). Since our monitor is language-independent, we will reuse it to prove our
 513 C-to-MSWasm secure compiler enforces memory safety (Section 5).

514

515 4.1 Color-based Memory Safety

516 Our notion of memory safety associates pointers and memory locations with *colors* (which represent
 517 pointer provenance [Memarian et al. 2019b]), shades, and allocation tags. Intuitively, each memory
 518 allocation generates a pointer annotated with a *unique* color (and shades as described below) and
 519 assigns the same color to each location in the allocated region of memory. Then, we consider a
 520 memory access *spatially* safe if the color of the pointer corresponds to the color of the memory
 521 location it points to. To account for *temporal* safety, memory locations are tagged as free or allocated
 522 and we enforce that accessed locations are tagged as allocated.

523 Colors are suitable to reason about memory safety at the granularity of individual memory
 524 objects. In particular, this simple model is sufficient for low-level languages that do not natively
 525 support composite data types (e.g., Wasm and MSWasm). However, colors alone cannot capture
 526 *intra-object* memory violations (e.g., the vulnerability in Listing 4). Intuitively, this is because the
 527 simple model assigns the *same* color to all the fields of a struct object. To reason about intra-object
 528 safety, we thus extend colors with *shades* and use a different shade to decorate the memory locations
 529 of each field in a struct. As a result, a pointer to a struct field cannot be used to access another field
 530 of the same struct, as their shades do not match.

531 As explained above, our definition of memory safety is intentionally minimal and language-
 532 agnostic: it does not specify other operations on colored pointers, e.g., pointer arithmetic, and how
 533 they propagate colors. This lets us reuse this definition of memory safety for different languages
 534 and reason about enforcing memory-safety via compilation in Section 5.

535

536 4.2 Memory-Safety Monitor

537 We formalize our notion of memory safety by constructing a safety monitor [Schneider 2000], i.e., a
 538 state machine that checks whether a trace satisfies memory safety. Intuitively, the monitor consumes
 539

539

$$\begin{array}{c}
 \text{540} \\
 \text{541} \\
 \text{542} \\
 \text{543} \\
 \text{544} \\
 \text{545} \\
 \text{546} \\
 \text{547} \\
 \text{548} \\
 \text{549} \\
 \text{550} \\
 \text{551} \\
 \text{552} \\
 \text{553}
 \end{array}
 \frac{
 \begin{array}{c}
 \text{(MS-Read)} \\
 \hline
 T(a) = A(c, s) \\
 \hline
 \alpha^* \vdash T \xrightarrow{\text{read}(a^{(c,s)})} T
 \end{array}
 \quad
 \begin{array}{c}
 \text{(MS-Write)} \\
 \hline
 T(a) = A(c, s) \\
 \hline
 \alpha^* \vdash T \xrightarrow{\text{write}(a^{(c,s)})} T
 \end{array}
 }{
 \begin{array}{c}
 \text{fresh}(c) \quad \forall j \in \{0..n-1\}. T(a+j) = F(_, _) \quad T' = T[a+i \mapsto A(c, \phi(i)) \mid i \in \{0..n-1\}] \\
 \hline
 \alpha^* \vdash T \xrightarrow{\text{alloc}(n, a^c, \phi)} T' \\
 \text{(MS-Alloc)} \\
 \hline
 \text{sfree}(a^c) \notin \alpha_2^* \quad T' = T[i \mapsto F(c, s_i) \mid i \mapsto A(c, s_i) \in T] \\
 \text{(MS-Free)} \\
 \hline
 \alpha_1^* \cdot \text{alloc}(n, a^c, \phi) \cdot \alpha_2^* \vdash T \xrightarrow{\text{sfree}(a^c)} T'
 \end{array}
 }$$

Fig. 5. Trace-based definition of memory safety.

554 a trace of memory events and gets stuck when it encounters a memory violation. We assume an
 555 infinite set of colors C , shades \mathcal{S} , and define a *colored shadow memory* $T \in \mathbb{N} \rightarrow \{A(c, s), F(c, s)\}$,
 556 i.e., a finite partial map from addresses $a \in \mathbb{N}$ to tagged colors $c \in C$ and shades $s \in \mathcal{S}$, where tags
 557 A and F denote whether a memory location is allocated or free, respectively. Then, we define an
 558 *abstract trace model* of memory events α , which include read and write operations with colored
 559 pointers, i.e., $\text{read}(a^{(c,s)})$ and $\text{write}(a^{(c,s)})$, memory allocations, i.e., $\text{alloc}(n, a^c, \phi)$ denoting a
 560 n -sized c -colored allocation starting at address a , in which sub-regions are shaded according to
 561 function $\phi : \{0, \dots, n-1\} \rightarrow \mathcal{S}$, and free operations, i.e., $\text{sfree}(a^c)$ which frees the c -colored
 562 memory region allocated at address a . Lastly, we define the transition system of the monitor over
 563 shadow memories and event history α^* through the judgment $\alpha^* \vdash T \xrightarrow{\alpha} T'$ (Fig. 5).

564 Rules **(MS-Read)** and **(MS-Write)** consume events $\text{read}(a^{(c,s)})$ and $\text{write}(a^{(c,s)})$, respectively,
 565 provided that the color and the shade are equal to those stored at location a in shadow memory
 566 and that location a is allocated, i.e., $T(a) = A(c, s)$. If the colors or the shades do not match, or the
 567 memory location is free, the state machine simply gets stuck, thus detecting a memory violation.
 568 To consume event $\text{alloc}(n, a^c, \phi)$, rule **(MS-Alloc)** allocates n contiguous, currently *free* locations
 569 in shadow memory, starting at address a , and assigns fresh color c and the shade given by ϕ to
 570 them. In response to event $\text{sfree}(a^c)$, the monitor frees the c -colored region of memory previously
 571 allocated at address a through rule **(MS-Free)**. First, the rule checks that a matching allocation
 572 event is present in the history, i.e., $\alpha_1^* \cdot \text{alloc}(n, a^c, \phi) \cdot \alpha_2^*$ for some size n and shading function
 573 ϕ , and that region has not already been freed, i.e., $\text{sfree}(a^c) \notin \alpha_2^*$, and then sets the tag of the
 574 memory locations colored c as free.

575 We say a trace is memory safe, written $\text{MS}(\alpha^*)$, if and only if the state machine does not get
 576 stuck while processing the trace starting from the empty shadow memory \emptyset and empty history ϵ .

577 In the definition below, we write $\xrightarrow{\alpha^*}$ for the reflexive transitive closure of $\xrightarrow{\alpha}$, which accumulates
 578 single events into a trace and records the event history.

580 **Definition 1** (Memory Safety). $\text{MS}(\alpha^*) \stackrel{\text{def}}{=} \exists T. \epsilon \vdash \emptyset \xrightarrow{\alpha^*} T$

583 4.3 Memory Safety of MSWasm.

584 In order to establish memory safety for MSWasm, we first need to map the trace model of MSWasm
 585 to the abstract trace model of Section 4.1. The main difference between the two is that the abstract
 586 model identifies safe memory accesses using colors and shades, while MSWasm relies on bounds
 587

$$\begin{array}{c}
589 \\
590 \\
591 \\
592 \\
593 \\
594 \\
595 \\
596 \\
597 \\
598 \\
599 \\
600 \\
601 \\
602 \\
603 \\
604 \\
605 \\
606 \\
607 \\
608 \\
609 \\
610 \\
611 \\
612 \\
613 \\
614 \\
615 \\
616 \\
617 \\
618 \\
619 \\
620 \\
621 \\
622 \\
623 \\
624 \\
625 \\
626 \\
627 \\
628 \\
629 \\
630 \\
631 \\
632 \\
633 \\
634 \\
635 \\
636 \\
637
\end{array}$$

$$\begin{array}{c}
\text{(Trace-Read)} \\
\hline
\mathbf{h} = \langle \mathbf{n}_b, \mathbf{n}_o, _, _, \mathbf{n}_{id} \rangle \quad \delta(\mathbf{n}_b, \mathbf{n}_{id}) = b^{(c,s)} \quad a = b + \mathbf{n}_o \quad n = |\tau| \\
\text{read}_\tau(\mathbf{h}) =_\delta \text{read}(a^{(c,s)}) \cdots \text{read}((a+n-1)^{(c,s)}) \\
\text{(Trace-Write)} \\
\hline
\mathbf{h} = \langle \mathbf{n}_b, \mathbf{n}_o, _, _, \mathbf{n}_{id} \rangle \quad \delta(\mathbf{n}_b, \mathbf{n}_{id}) = b^{(c,s)} \quad a = b + \mathbf{n}_o \quad n = |\tau| \\
\text{write}_\tau(\mathbf{h}) =_\delta \text{write}(a^{(c,s)}) \cdots \text{write}((a+n-1)^{(c,s)}) \\
\text{(Trace-SAlloc)} \\
\hline
\mathbf{h} = \langle \mathbf{n}_b, \mathbf{0}, \mathbf{n}_o, _, \mathbf{n}_{id} \rangle \quad n = \mathbf{n}_o \quad \forall i \in \{0..n-1\}. \delta(\mathbf{n}_b + i, \mathbf{n}_{id}) = (a+i)^{(c,\phi(i))} \\
\text{salloc}(\mathbf{h}) =_\delta \text{alloc}(n, a^c, \phi) \\
\text{(Tr-Sfree)} \\
\hline
\mathbf{h} = \langle \mathbf{n}_b, _, \mathbf{n}, _, \mathbf{n}_{id} \rangle \quad \delta(\mathbf{n}_b, \mathbf{n}_{id}) = a^{(c,_)} \\
\text{sfree}(\mathbf{h}, \mathbf{n}_{id}) =_\delta \text{sfree}(a^c) \\
\text{(Trace-Trap)} \\
\hline
\text{trap} =_\delta \epsilon
\end{array}$$

Fig. 6. Relation between MSWasm and abstract events (excerpts).

checks and segment identifiers. Furthermore, individual $\text{read}_\tau(\mathbf{h})$ and $\text{write}_\tau(\mathbf{h})$ events correspond to multiple memory accesses in the abstract trace model, as these operations read and write byte sequences in MSWasm. We reconcile these differences between the two trace models with the relation $\alpha =_\delta \alpha^*$ whose most relevant rules are defined in Fig. 6. The relation is parametrized by a partial bijection $\delta : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N} \times \mathcal{C} \times \mathcal{S}$, which maps pairs $(\mathbf{a}, \mathbf{n}_{id})$, consisting of an allocated segment memory address \mathbf{a} and a segment identifier \mathbf{n}_{id} , into corresponding shadow memory addresses $a^{(c,s)}$, decorated with colors and shades. Intuitively, we can construct a suitable bijection δ from the MSWasm allocator, which has information about what is allocated in segment memory.

Rules **(Trace-Read)** and **(Trace-Write)** relate single MSWasm $\text{read}_\tau(\mathbf{h})$ and $\text{write}_\tau(\mathbf{h})$ events to a sequence of $|\tau|$ contiguous abstract read and write events, respectively. The rules convert the handle base address and the segment identifier into the corresponding colored base address, i.e., $\delta(\mathbf{n}_b, \mathbf{n}_{id}) = b^{(c,s)}$, which is then incremented with the offset of the handle to obtain the first abstract location accessed, i.e., $a = b + \mathbf{n}_o$, similar to MSWasm semantics. Since these abstract events originate from the same handle, the rule labels their address with the same color c and shade s obtained from the base address of the handle to reflect their provenance. If we computed the colors for these addresses using the bijection δ , then they would automatically match the color stored in the shadow memory and memory safety would hold trivially. Instead, these addresses are tagged with the provenance color, and therefore proving memory safety (i.e., stepping using the rules of Figure 5) requires showing that this color matches the color found in shadow memory, which in turn requires reasoning about the integrity of the handle and the bounds checks performed by MSWasm. A final subtlety of these rules is that they seem to ignore the integrity flag of the handle. This is because in MSWasm, only authentic handles can generate read and write events—reading and writing memory via corrupted handles results in a **trap** event.

Rule **(Trace-SAlloc)** relates the allocation of a \mathbf{n}_o -byte segment in MSWasm to a corresponding abstract allocation of the same size, i.e., event $\text{alloc}(n, a^c, \phi)$ where $\mathbf{n}_o = n$. In the rule, premise $\forall i \in \{0, \dots, n-1\}. \delta(\mathbf{a} + i) = (\mathbf{n}_b + i, \mathbf{n}_{id})^{(c,\phi(i))}$ ensure that (i) all the abstract addresses share the same color c , and (ii) the bijection δ and the shading function ϕ agree on the shades used for the segment. In general, function ϕ can be a constant function, when we reason about memory safety for native MSWasm programs, e.g., to prove that MSWasm is memory safe in Theorem 1 below. Intuitively, MSWasm does not provide an explicit representation for structured data, therefore it is sufficient to assign the same shade to all locations of a segment to prove memory safety. When we

use MSWasm as a compilation target however, segments can store also structured objects (e.g., a struct) in addition to flat objects (e.g., an array). In this scenario, we instantiate ϕ according to the source type of the object, which let us show that compiled C/C++ programs achieve intra-object memory safety later (Theorem 2).

To relate free events, rule (Tr-Sfree) requires the bijection δ to match the base \mathbf{n}_b and identifier \mathbf{n}_{id} of the segment pointed to by the handle to the colored address a^c freed by the monitor, i.e., $\delta(\mathbf{n}_b, \mathbf{n}_{id}) = a^{(c, _)}$. Because identifiers and colors are never reused, freed segments and regions can be reused for other allocations, while keeping dangling handles and colored pointers related by the bijection. For example, if a segment is later allocated at address \mathbf{n}_b , it will be associated with a *unique* identifier $\mathbf{n}'_{id} \neq \mathbf{n}_{id}$, which can be related to some shadow address a' and *fresh* color $c' \neq c$ through an *extended* bijection $\delta' \supseteq \delta$.⁴

Lastly, rule (Trace-Trap) relates event **trap** in MSWasm to the empty trace ϵ , since **trap** simply stops the program and thus cannot cause a memory safety violation.

We can now state memory safety for MSWasm traces in terms of memory safety of a δ -related abstract trace for the state machine defined above.

Definition 2 (Memory Safety for MSWasm Traces). $\text{MS}(\alpha^*) \stackrel{\text{def}}{=} \exists \alpha^*, \delta. \alpha^* =_{\delta} \alpha^*$ and $\text{MS}(\alpha^*)$

We define memory safety for MSWasm modules if the trace generated during execution is memory safe. In the following, we write $\mathbf{M} \rightarrow \alpha^*$ for the trace generated by module \mathbf{M} with the semantics of Section 3.2.

Definition 3 (Memory Safety for MSWasm Modules). $\vdash \text{MS}(\mathbf{M}) \stackrel{\text{def}}{=} \mathbf{M} \rightarrow \alpha^*$ and $\text{MS}(\alpha^*)$

A module \mathbf{M} achieves robust memory safety if, given any valid attacker \mathbf{C} (denoted as $\mathbf{M} \vdash \mathbf{C} : \text{attacker}$, in the sense of Section 2.2), linking \mathbf{M} with \mathbf{C} produces a memory safe module. In the following, we write $\mathbf{M} \circ \mathbf{C}$ for the module obtained by linking \mathbf{M} with \mathbf{C} , i.e., the module obtained by instantiating the functions imported by \mathbf{M} with those of \mathbf{C} .

Definition 4 (Robust Memory Safety for MSWasm Modules). $\vdash \text{RMS}(\mathbf{M}) \stackrel{\text{def}}{=} \forall \mathbf{C} \text{ s.t. } \mathbf{M} \vdash \mathbf{C} : \text{attacker}. \vdash \text{MS}(\mathbf{M} \circ \mathbf{C})$

The main result for MSWasm is that any well-typed module ($\vdash \mathbf{M} : \text{wt}$) is memory-safe, robustly.

Theorem 1 (Robust Memory Safety for MSWasm). If $\vdash \mathbf{M} : \text{wt}$ then $\vdash \text{RMS}(\mathbf{M})$

Proof (Sketch). Intuitively, the type system of MSWasm ensures that well-typed modules can access segment memory only through **handle** values and safe instructions. Programs that accesses memory via *invalid* handles **trap** and so trivially respect memory safety (Rule (Trace-Trap)). When accessing segments via *valid* handles, MSWasm performs memory safety checks using their metadata, so the rest of the proof requires showing an invariant about handle integrity. Intuitively, this invariant guarantees that valid handles (whether proper values, or stored in segment memory), correspond to allocated segments in memory. Then, using this invariant, we can show that programs that access segment memory without trapping, pass the memory safety checks, and thus are memory safe. \square

In the next section, we leverage the memory-safety abstractions of MSWasm to develop a formal C compiler that provably enforces memory safety.

5 MEMORY SAFETY THROUGH COMPILATION

This section shows how a C compiler targeting MSWasm can enforce memory safety. Thus, we formalize a simplified version of C (Section 5.1), as a memory-unsafe source language for our compiler, and the compiler itself (Section 5.2). We then prove that the compiler enforces memory safety

⁴In technical terms the bijection grows *monotonically*, which provides a suitable inductive principle for our formal results.


```

687 Programs  $M ::= l^*, D^*, F^*, n_{hs}$  Imports  $l ::= \tau g(x : \tau)$  Structs  $D ::= s \mapsto (f : \tau)^*$ 
688
689 Values  $v ::= n \mid n^{(n_1, n_2, w, n_{id})}$  Functions  $F ::= \tau g(x : \tau) \{ \text{var } (y : \tau)^*, e \}$ 
690
691 Word Types  $w ::= \tau \mid \text{struct } s \mid \text{array } \tau$  Expr. Types  $\tau ::= \text{int} \mid \text{ptr } w$ 
692
693 Expr.  $e ::= v \mid x \mid e; e \mid e \oplus e \mid x := e \mid x := \text{call } g(e) \mid *e \mid e[e]$ 
694
695  $\mid *e := e \mid e[e] := e \mid \text{if } e \text{ then } e \text{ else } e \mid \&e \rightarrow f$ 
696
697  $\mid \text{malloc}(\tau, e) \mid \text{malloc}(w) \mid \text{free}(e)$ 
698
699 Stores  $\Sigma ::= (H, A)$  Heaps  $H ::= [] \mid v : H$  Local Env.  $\theta ::= (x \mapsto v)^*$  Allocators  $A$ 
700
701 Events  $\alpha ::= \epsilon \mid \text{alloc}(v) \mid \text{free}(v) \mid \text{read}_\tau(v) \mid \text{write}_\tau(v)$ 

```

Fig. 7. C syntax and runtime structures (excerpts).

(Section 5.3), i.e., memory-safe programs compiled to MSWasm execute unchanged (Theorem 2), while memory-unsafe programs abort at the first memory violation (Theorem 3).

5.1 The Source Language C

Figure 7 presents the syntax of our source language, a subset of C, inspired by previous work [Ruef et al. 2019]. Source programs M specify the type of imported functions l^* , struct definitions D^* , function definitions F^* , and the heap size n_{hs} . Struct definitions map struct names s to a list of field names f and their types τ . Types are mutually defined by expression types τ , i.e., integers (int) and pointers ($\text{ptr } w$), and word types w , which, in addition to τ , include also multi-word values, i.e., structs ($\text{struct } s$) and arrays ($\text{ptr array } \tau$). Note that arrays are always typed as pointers. Syntax $\tau g(x : \tau) \{ \text{var } (y : \tau)^*, e \}$ defines function g , its argument and return type, and declares local variables $(y : \tau)^*$ in scope of the body e . Expressions are standard and include reading and writing memory via pointers (i.e., $*e$ and $*e := e$) and accessing struct fields (i.e., $\&e \rightarrow f$). An array e_1 at index e_2 is read and written via $e_1[e_2]$ and $e_1[e_2] := e_3$.

Expression $\text{malloc}(\tau, e)$ allocates an array containing e elements of type τ , while $\text{malloc}(w)$ allocates a buffer to store a single element of type $w \neq \text{array } \tau$. Values include integers n and annotated pointers, i.e., $n^{(n_1, n_2, w, n_{id})}$, where n is the address pointed to by the pointer, and (n_1, n_2, w, n_{id}) indicates that the pointer refers to a buffer allocated at address n_1 , containing n_2 elements of type w , and identified by n_{id} . These annotations are inspired by previous work on pointer provenance [Memarian et al. 2019a] and are only needed to reason about memory safety of source programs, i.e., the source semantics does *not* enforce memory safety and ignores them.

Typing. The type system for the source language is mostly standard and defined by judgment $F^*, \Gamma \vdash e : \tau$, which indicates that expression e has type τ under functions F^* and typing context Γ (which binds variables to types). The type system allows typing integers as pointers and restricts function type signatures to expression types for simplicity.

Semantics. We define a small-step contextual semantics for C with the following judgment, $M \vdash \langle \Sigma, \theta, e \rangle \xrightarrow{\alpha} \langle \Sigma', \theta', e' \rangle$, in which local configuration $\langle \Sigma, \theta, e \rangle$ steps and produces event α , under program definition M . Local configurations contain the store Σ , the local variable environment θ mapping named variables to values, and an expression e to be evaluated. The store Σ contains the heap H , a list of values, and the allocator state A . The heap abstracts away low-level details about the memory layout and the byte representation of values (e.g., we store structs and arrays simply as a flattened sequence of single-word values). Similar to MSWasm, events α record memory relevant

$$\begin{array}{c}
 \text{(Ptr-Arith)} \\
 \hline
 \text{736 } M \vdash \langle \Sigma, \theta, a^{(b,\ell,w,n_{id})} \oplus n \rangle \mapsto \langle \Sigma, \theta, (a+n)^{(b,\ell,w,n_{id})} \rangle \\
 \text{737 } \\
 \text{738 } \Sigma = \langle H, A \rangle \quad H' = H[a \mapsto v] \quad \Sigma' = \langle H', A \rangle \quad \vdash v : \tau \\
 \text{739 } \hline
 \text{740 } M \vdash \langle \Sigma, \theta, *a^{(b,\ell,w,n_{id})} := v \rangle \xrightarrow{\text{write}_\tau(a^{(b,\ell,w,n_{id})})} \langle \Sigma', \theta, 0 \rangle \\
 \text{741 } \\
 \text{742 } \Sigma = \langle H, A \rangle \quad H' = H[a \mapsto v] \quad \Sigma' = \langle H', A \rangle \quad \vdash v : \tau \\
 \text{743 } \hline
 \text{744 } M \vdash \langle \Sigma, \theta, *a := v \rangle \xrightarrow{\text{write}_\tau(a)} \langle \Sigma', \theta, 0 \rangle \\
 \text{745 } \text{(Malloc-Single)} \\
 \text{746 } \Sigma = \langle H, A \rangle \xrightarrow{\text{alloc}(a^{(a,1,w,n_{id})})} \langle H', A' \rangle = \Sigma' \quad v = a^{(a,1,w,n_{id})} \\
 \text{747 } \hline
 \text{748 } M \vdash \langle \Sigma, \theta, \text{malloc}(w) \rangle \xrightarrow{\text{alloc}(v)} \langle \Sigma', \theta, v \rangle \\
 \text{749 } \text{(Malloc-Array)} \\
 \text{750 } \Sigma = \langle H, A \rangle \xrightarrow{\text{alloc}(a^{(a,n,\tau,n_{id})})} \langle H', A' \rangle = \Sigma' \quad v = a^{(a,n,\tau,n_{id})} \\
 \text{751 } \hline
 \text{752 } M \vdash \langle \Sigma, \theta, \text{malloc}(\tau, n) \rangle \xrightarrow{\text{alloc}(v)} \langle \Sigma', \theta, v \rangle \\
 \text{753 } \\
 \text{754 } \\
 \text{755 } \\
 \text{756 } \\
 \text{757 } \\
 \text{758 } \\
 \text{759 } \\
 \text{760 } \\
 \text{761 } \\
 \text{762 } \\
 \text{763 } \\
 \text{764 } \\
 \text{765 } \\
 \text{766 } \\
 \text{767 } \\
 \text{768 } \\
 \text{769 } \\
 \text{770 } \\
 \text{771 } \\
 \text{772 } \\
 \text{773 } \\
 \text{774 } \\
 \text{775 } \\
 \text{776 } \\
 \text{777 } \\
 \text{778 } \\
 \text{779 } \\
 \text{780 } \\
 \text{781 } \\
 \text{782 } \\
 \text{783 } \\
 \text{784 }
 \end{array}$$

Fig. 8. Semantics of C (excerpts).

operations, including silent events ϵ , allocating and releasing memory, i.e., $\text{alloc}(v)$ and $\text{free}(v)$, and reading and writing values of type τ with a pointer v , i.e., $\text{read}_\tau(v)$ and $\text{write}_\tau(v)$.

Figure 8 presents some of the semantics rules of the source language. Rule **(Ptr-Arith)** performs pointer arithmetic by incrementing the address of the pointer, without changing the metadata. Rules **(Write-Ptr)** and **(Write-Int)** write a value v in the heap through a pointer and a raw integer address, respectively. As explained above, these rules do not check that the write operation is safe, but only record the pointer and the type τ of the value that gets stored in the generated event, i.e., $\text{write}_\tau(a^{(b,\ell,w,n_{id})})$ and $\text{write}_\tau(a)$. Rules **(Malloc-Single)** and **(Malloc-Array)** allocate a buffer for a single object of type w and a n -elements array, respectively, and return a pointer value annotated with appropriate metadata. Similar to the MSWasm semantics, the source language invokes the allocator to serve allocation and free requests ($\langle H, A \rangle \xrightarrow{\alpha} \langle H', A' \rangle$). In contrast to the safe allocator of MSWasm however, the source allocator does not trap upon an invalid free request, i.e., a free of an unallocated memory region, but silently drops the request.⁵

The source language uses a separate semantic judgment for function calls and returns (omitted), and a top-level judgment $M \rightarrow \alpha^*$, which collects the trace generated by program M . We define memory safety for source traces using the general abstract monitor from Section 4.1:

Definition 5 (Memory-Safety for C). $\text{MS}(\alpha^*) \stackrel{\text{def}}{=} \exists \alpha^*, \delta. \alpha^* =_\delta \alpha^*$ and $\text{MS}(\alpha^*)$

This definition is analogous to Definition 2 for MSWasm: it relies on a bijection δ to map source addresses into corresponding colored abstract addresses, and a relation $\alpha^* =_\delta \alpha^*$ to connect source and abstract traces through the bijection. The trace relation is defined similarly to the relation given in Figure 6 for MSWasm and additionally constructs appropriate shading functions for $\text{alloc}(v)$

⁵Invalid free requests cause *undefined behavior* in C and usually result in the corruption of memory objects or the allocator state. Since we represent the allocator state explicitly and separately from the program memory, free requests cannot cause such specific behaviors in our model.

$$\begin{array}{c}
785 \\
786 \\
787 \\
788 \\
789 \\
790 \\
791 \\
792 \\
793 \\
794 \\
795 \\
796 \\
797 \\
798 \\
799 \\
800 \\
801 \\
802 \\
803 \\
804 \\
805 \\
806 \\
807 \\
808 \\
809 \\
810 \\
811 \\
812 \\
813 \\
814 \\
815 \\
816 \\
817 \\
818 \\
819 \\
820 \\
821 \\
822 \\
823 \\
824 \\
825 \\
826 \\
827 \\
828 \\
829 \\
830 \\
831 \\
832 \\
833
\end{array}$$

$$\begin{array}{c}
\text{(C-Ptr-Arith)} \\
\frac{\llbracket P, \Gamma \vdash e_1 : \text{array } \tau \rrbracket^{\text{exp}} = \mathbf{i}_1^* \quad \llbracket P, \Gamma \vdash e_2 : \text{int} \rrbracket^{\text{exp}} = \mathbf{i}_2^* \quad \mathbf{n} = \text{sz}(\tau)}{\llbracket P, \Gamma \vdash e_1 \oplus e_2 : \text{array } \tau \rrbracket^{\text{exp}} = \mathbf{i}_1^*; \mathbf{i}_2^*; \mathbf{i32.const } \mathbf{n}; \mathbf{i32.}\times; \mathbf{handle.add}} \\
\text{(C-BinOp)} \\
\frac{\llbracket P, \Gamma \vdash e_1 : w \rrbracket^{\text{exp}} = \mathbf{i}_1^* \quad \llbracket P, \Gamma \vdash e_2 : w \rrbracket^{\text{exp}} = \mathbf{i}_2^* \quad \llbracket \tau \rrbracket = \tau}{\llbracket P, \Gamma \vdash e_1 \oplus e_2 : \tau \rrbracket^{\text{exp}} = \mathbf{i}_1^*; \mathbf{i}_2^*; \tau.\otimes} \\
\text{(C-Malloc-Array)} \\
\frac{\llbracket P, \Gamma \vdash e : \text{int} \rrbracket^{\text{exp}} = \mathbf{i}_1^* \quad \mathbf{n} = \text{sz}(\tau)}{\llbracket \Gamma \vdash \text{malloc}(\tau, e) : \text{ptr}(\text{array } \tau) \rrbracket^{\text{exp}} = \mathbf{i}_1^*; \mathbf{i32.const } \mathbf{n}; \mathbf{i32.}\otimes; \mathbf{segalloc}} \\
\text{(C-Malloc-Single)} \\
\frac{\mathbf{n} = \text{sz}(w)}{\llbracket P, \Gamma \vdash \text{malloc}(w) : \text{ptr } w \rrbracket^{\text{exp}} = \mathbf{i32.const } \mathbf{n}; \mathbf{segalloc}} \\
\text{(C-Deref)} \\
\frac{\llbracket P, \Gamma \vdash e : \text{ptr } \tau \rrbracket^{\text{exp}} = \mathbf{i}^* \quad \llbracket \tau \rrbracket = \tau}{\llbracket P, \Gamma \vdash *e : \tau \rrbracket^{\text{exp}} = \mathbf{i}^* + [\tau.\text{segload}]} \\
\text{(C-Struct-field)} \\
\frac{\llbracket P, \Gamma \vdash e : \text{ptr}(\text{struct } s) \rrbracket^{\text{exp}} = \mathbf{i}^* \quad (\mathbf{o}_1, \mathbf{o}_2) = \text{offset}(s, f)}{\llbracket P, \Gamma \vdash \&e \rightarrow f : \text{ptr } \tau \rrbracket^{\text{exp}} = \mathbf{i}^* + [\mathbf{i32.const } \mathbf{o}_1, \mathbf{i32.const } \mathbf{o}_2, \text{slice}]}
\end{array}$$

Fig. 9. Compiler from C to MSWasm (excerpts).

events according to the type of the allocated object (e.g., an array or a struct). Accesses via raw addresses \mathbf{n} are excluded from the relation, i.e., $\text{write}_\tau(\mathbf{a}) \neq_s \alpha^*$ and $\text{read}_\tau(\mathbf{a}) \neq_s \alpha^*$ for any abstract trace α^* . Omitting them from the relation captures the fact that memory accesses with forged pointers violate memory safety, as the provenance of these pointers is undefined.

5.2 The Compiler

We define the compiler $\llbracket \cdot \rrbracket$ from C to MSWasm inductively on the type derivation of C modules, functions and expressions (Figure 9). To prevent untrusted code from violating memory safety, our compiler translates pointers to handles and only uses MSWasm segment memory. Thus, we translate source types τ into MSWasm types τ as $\llbracket \text{int} \rrbracket = \mathbf{i32}$ and $\llbracket \text{ptr } w \rrbracket = \mathbf{handle}$. The compiler relies on source types to emit MSWasm instructions with appropriate byte sizes (calculated with function $\text{sz}(\cdot) : \tau \rightarrow \mathbf{n}$) and offsets for expressions that involve pointer arithmetic, struct accesses and memory allocations. For example, a binary operation (\oplus) whose first operand is an array ($\text{array } \tau$) needs to be compiled in a $\mathbf{handle.add}$, as in Rule (C-Ptr-Arith). On the other hand, a binary operation on naturals needs to be compiled in the related MSWasm binary operation, as in Rule (C-BinOp). Another example of the way source types guide the compilation is for the compilation of expression $\text{malloc}(\tau, e)$. Here, if the resulting type is a pointer to an array ($\text{ptr}(\text{array } \tau)$), the compiler must first emit instructions to compute the size of a segment large enough for an array containing e elements of type τ , and then instruction $\mathbf{segalloc}$ to invoke the allocator and generate the corresponding handle. Therefore, rule (C-Malloc-Array) recursively compiles the array length e , i.e., $\llbracket P, \Gamma \vdash e : \text{int} \rrbracket^{\text{exp}} = \mathbf{i}^*$, which then gets multiplied by $|\tau|$, i.e., the size in bytes of a value of type τ , via instruction $\mathbf{i32.}\otimes$, and finally passed to $\mathbf{segalloc}$. On the other hand, if the return type is a pointer to any other type ($\text{ptr } w$), the compiler needs to calculate its size (\mathbf{n}) and allocate enough memory (Rule (C-Malloc-Single)) Since expression $*e$ reads a pointer to a value of type τ , rule (C-Deref)

emits instruction to first evaluate the corresponding handle, i.e., $\llbracket P, \Gamma \vdash e : \text{ptr } \tau \rrbracket^{\text{exp}} = i^*$, followed
834 by instruction $\tau.\text{segload}$, whose compiled type $\llbracket \tau \rrbracket = \tau$ ensures that the generated code reads
835 the right number of bytes and interprets them at the corresponding target type. Lastly, rule (C-
836 Struct-field) translates a struct field access $\&e \rightarrow f$ by slicing the handle obtained from pointer e ,
837 thus enforcing intra-object safety in the generated code. To this end, the rule emits instructions
838 $[i32.\text{const } o_1, i32.\text{const } o_2, \text{slice}]$, where offsets o_1 and o_2 are obtained from function $\text{offset}(s, f)$,
839 which statically computes the offsets necessary to select field f in the byte representation of struct
840 s .

842 5.3 Properties of the Compiler

843 We establish two properties for our compiler. The first (Theorem 2) shows that the compiler is
844 functionally correct and preserves memory safety for memory-safe source programs. The second
845 (Theorem 3) shows that *memory-unsafe* programs compiled to MSWasm abort at the first memory
846 violation. Together, these results show that our compiler enforces memory-safety (Corollary 1).

847 **Cross-Language Equivalence Relation.** Since our notion of memory safety is defined over traces,
848 and the source and target languages have different trace models, the formal results of the compiler
849 rely on a *cross-language* equivalence relation to show functional correctness and memory-safety
850 preservation [Leroy 2009]. Figure 10 (top) defines this relation for pointer values up to a partial
851 bijection $\delta : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N} \times \mathbb{N}$, which maps addresses and identifiers from source to target. Rule (Val-
852 Rel-Ptr) relates an annotated pointer $a^{(b, \ell, w, n_{\text{id}})}$ to a *valid* handle $\langle b, o, \ell, \text{true}, n_{\text{id}} \rangle$ as long as their
853 base address and identifier are matched by the bijection, i.e., $\delta(b, n_{\text{id}}) = b, n_{\text{id}}$, and the length and
854 offset fields match, taking into account the byte-size representation of w , i.e., $\ell \times |w| = \ell$ and
855 $(a - b) \times |w| = o$. In contrast, Rule (Val-Rel-Int) relates integer pointers to arbitrary *invalid* handles.
856

857 The relation between source and target events $\alpha =_{\delta} \alpha'$, relates the same single events (Figure 10,
858 bottom). When relating reads, writes, allocates, and frees, we insist that source pointers and target
859 handles are related (according to the cross-language value relation) and the handles are valid (these
860 rules have the validity bit set to **true**). Additionally, for reading and writing, their values being read
861 or written must be of related types, i.e., $\llbracket \tau \rrbracket = \tau$.

862 For memory-safe, well-typed source programs ($\vdash M : \text{wt}$), Theorem 2 states that the compiler
863 produces equivalent memory-safe target programs; i.e., the compiled program emits a memory-safe
864 trace that is related to the source trace.

865 **Theorem 2** (Memory-Safety Preservation).

866 If $\vdash M : \text{wt}$ and $M \rightarrow \alpha^*$ and $\text{MS}(\alpha^*)$ then $\exists \delta, \alpha^*. \llbracket M \rrbracket \rightarrow \alpha^*$ and $\alpha^* =_{\delta} \alpha^*$ and $\text{MS}(\alpha^*)$

867 In contrast, Theorem 3 states that memory-unsafe programs compiled to MSWasm abort at the
868 first memory violation.

869 **Theorem 3** (Memory Violations Trap).

870 If $\vdash M : \text{wt}$ and $M \rightarrow \alpha^* \# [\alpha] \# \alpha'^*$ and $\text{MS}(\alpha^*)$ and $\neg \text{MS}(\alpha^* \# [\alpha])$
871 then $\exists \delta, \alpha^*. \alpha^* =_{\delta} \alpha^*$ and $\llbracket M \rrbracket \rightarrow \alpha^* \# [\text{trap}]$

872 Together these theorems characterize the scope of our compiler-based memory-safety *enforcement*:

873 **Corollary 1** (Memory-Safety Enforcement). If $\llbracket M \rrbracket \rightarrow \alpha^*$ then $\text{MS}(\alpha^*)$

874 Figure 11 shows the essence of the proof technique for Theorem 2, in the diagram, full arrows
875 represent hypotheses and dashed arrows represent conclusions.

876 In the theorem statement, judgements of the form $M \rightarrow \alpha^*$ unfold to the reflexive-transitive
877 closure of a single semantics step (i.e., the rules presented in Figure 4 for MSWasm and in Figure 8
878

$$\begin{array}{c}
883 \\
884 \\
885 \\
886 \\
887 \\
888 \\
889 \\
890 \\
891 \\
892 \\
893 \\
894 \\
895 \\
896 \\
897 \\
898 \\
899 \\
900 \\
901 \\
902 \\
903 \\
904 \\
905 \\
906 \\
907 \\
908 \\
909 \\
910 \\
911 \\
912 \\
913 \\
914 \\
915 \\
916 \\
917 \\
918 \\
919 \\
920 \\
921 \\
922 \\
923 \\
924 \\
925 \\
926 \\
927 \\
928 \\
929 \\
930 \\
931
\end{array}$$

$$\begin{array}{c}
\text{(Val-Rel-Ptr)} \\
\delta(\mathbf{b}, \mathbf{n}_{id}) = \mathbf{b}, \mathbf{n}_{id} \quad \ell \times |\mathbf{w}| = \ell \quad (\mathbf{a} - \mathbf{b}) \times |\mathbf{w}| = \mathbf{o} \\
\hline
\mathbf{a}^{(\mathbf{b}, \ell, \mathbf{w}, \mathbf{n}_{id})} \sim_{\delta} \langle \mathbf{b}, \mathbf{o}, \ell, \text{true}, \mathbf{n}_{id} \rangle \\
\hline
\text{(Tr-Rel-Read-Ptr)} \\
\llbracket \tau \rrbracket = \tau \quad \mathbf{a}^{(\mathbf{b}, \ell, \mathbf{w}, \mathbf{n}_{id})} \sim_{\delta} \langle \mathbf{b}, \mathbf{o}, \ell, \text{true}, \mathbf{n}_{id} \rangle \\
\hline
\text{read}_{\tau}(\mathbf{a}^{(\mathbf{b}, \ell, \mathbf{w}, \mathbf{n}_{id})}) =_{\delta} \text{read}_{\tau}(\langle \mathbf{b}, \mathbf{o}, \ell, \text{true}, \mathbf{n}_{id} \rangle) \\
\text{(Tr-Rel-Write-Ptr)} \\
\llbracket \tau \rrbracket = \tau \quad \mathbf{a}^{(\mathbf{b}, \ell, \mathbf{w}, \mathbf{n}_{id})} \sim_{\delta} \langle \mathbf{b}, \mathbf{o}, \ell, \text{true}, \mathbf{n}_{id} \rangle \\
\hline
\text{write}_{\tau}(\mathbf{a}^{(\mathbf{b}, \ell, \mathbf{w}, \mathbf{n}_{id})}) =_{\delta} \text{write}_{\tau}(\langle \mathbf{b}, \mathbf{o}, \ell, \text{true} \rangle) \\
\text{(Tr-Rel-Allocate)} \\
\mathbf{a}^{(\mathbf{b}, \ell, \mathbf{w}, \mathbf{n}_{id})} \sim_{\delta} \langle \mathbf{b}, \mathbf{o}, \ell, \text{true}, \mathbf{n}_{id} \rangle \\
\hline
\text{salloc}(\mathbf{a}^{(\mathbf{b}, \ell, \mathbf{w}, \mathbf{n}_{id})}) =_{\delta} \text{salloc}(\langle \mathbf{b}, \mathbf{o}, \ell, \text{true}, \mathbf{n}_{id} \rangle) \\
\text{(Tr-Rel-Free)} \\
\mathbf{a}^{(\mathbf{b}, \ell, \mathbf{w}, \mathbf{n}_{id})} \sim_{\delta} \langle \mathbf{b}, \mathbf{o}, \ell, \text{true}, \mathbf{n}_{id} \rangle \\
\hline
\text{sfree}(\mathbf{a}^{(\mathbf{b}, \ell, \mathbf{w}, \mathbf{n}_{id})}) =_{\delta} \text{free}(\langle \mathbf{b}, \mathbf{o}, \ell, \text{true}, \mathbf{n}_{id} \rangle) \\
\hline
\text{(Val-Rel-Int)} \\
\mathbf{n} \sim_{\delta} \langle \mathbf{b}, \mathbf{o}, \ell, \text{false}, \mathbf{n}_{id} \rangle
\end{array}$$

Fig. 10. Cross-language equivalence relation: values (top) and events (bottom).




Fig. 11. Proof diagram for Theorem 2: functional correctness (left) and memory-safety preservation (right).

for C). The proof then proceeds unsurprisingly by induction over the reflexive-transitive reductions that generate the source trace, the figure shows the single-step case. We use metavariable Ω to indicate program states, which are the tuples presented in the semantics rules of each language.

We first describe the most interesting case of the functional correctness part of Theorem 2, i.e., the left of Figure 11. There, we need to show how one single source step ($\Omega \xrightarrow{\alpha} \Omega'$) that triggers a change in the source allocator ($A \xrightarrow{\alpha} A'$)⁶, causes a series of ‘related’ target steps ($\Omega' \xrightarrow{\alpha}^* \Omega'$) that change the target allocator accordingly ($A \xrightarrow{\alpha} A'$). Essentially, target steps are related when they generate actions that are related (as per Figure 10), and they take related states ($\Omega \sim_{\delta} \Omega$) into still-related states ($\Omega' \sim_{\delta} \Omega'$). We do not present the formalisation of the state relation, intuitively it

⁶For ease of reading, we massage the allocator reduction judgement $\langle H, A \rangle \xrightarrow{\alpha} \langle H', A' \rangle$ to only contain the allocator.

932 just lifts the value relation of Figure 10 to all elements of a program state. Proving that the allocators
933 step using related actions ensures that the source and target allocators are in related, consistent
934 states. This is key to the memory safety preservation part of the theorem, i.e., the right of Figure 11.

935 To prove memory safety preservation for Theorem 2, we start from the functional correctness
936 square between allocators (i.e., the base of the prism on the right in Figure 11). Then, we assume
937 that the C step is memory-safe, this is represented by the blue side of the prism. Technically, this
938 relies on another omitted piece of formalisation that relates source allocator states A and shadow
939 memories T , a relation that holds when the addresses tracked in A and T are the same up to a
940 bijection δ . The source memory safety assumption tells that the states of the initial C allocator
941 and of the initial shadow memory are related ($A =_{\delta} T$), that they take a related step ($\alpha =_{\delta'} \alpha'$), and
942 that leads to related final states ($A' =_{\delta'} T'$). The goal of the memory safety part of the proof is
943 depicted as the corresponding red side of the prism: there is a relation between the states of the
944 initial MSWasm allocator and of the initial shadow memory ($A =_{\delta} T$), the states take a related step
945 ($\alpha =_{\delta'} \alpha'$) and that leads to related final states ($A' =_{\delta'} T'$). To construct the relations in the red
946 square, we need to derive the dashed edges of the vertical triangles according to the correct relation
947 with the correct bijection. This relation we obtain by combining the corresponding source-to-target
948 relations (i.e., $A =_{\delta} A'$) and the source-to-monitor relation (i.e., $A =_{\delta} T$), and compose their bijections
949 to relate abstract and target locations. That is, we obtain $A =_{\delta} T$, where δ (relating MSWasm and
950 abstract addresses) is the composition of δ (relating MSWasm and C addresses) with δ (relating
951 C and abstract addresses). Importantly, the triangle of relations guarantees that the C notion of
952 memory safety is preserved *exactly* in MSWasm. Since in C we instantiate our abstract notion of
953 memory safety to account for intra-object safety, we get the same fine-grained memory safety
954 notion preserved in MSWasm.

955 The proof of Theorem 3 is analogous. There, we use the same intuition presented above to
956 simulate all actions of the memory-safe trace α^* starting from their memory-safe counterparts in
957 α^* . Then, at some point, starting from related states, C performs a memory-unsafe action α and
958 MSWasm emits a **trap**. This proof is by case analysis over C memory safety violations, which we
959 identify by the related abstract monitor getting stuck. In the proof, we relate these violations to a
960 *failing* memory safety check in MSWasm, which causes the compiled program to **trap**, as expected.

961 6 IMPLEMENTING MSWASM

962 In this section we describe our prototype MSWasm compilation framework (Figure 12). We im-
963 plement two compilers of MSWasm following the language semantics of Section 3. Our first
964 compiler is an ahead-of-time (AOT) compiler from Wasm to executable machine code (Section 6.1);
965 it demonstrates MSWasm's flexibility in employing different enforcement techniques, including
966 both software-based enforcement and hardware-accelerated enforcement. Our second compiler is a
967 compiler from Wasm to Java bytecode (Section 6.2); it demonstrates MSWasm's compatibility with
968 just-in-time (JIT) compilation. We also implement a compiler from C to MSWasm (Section 6.3),
969 following the formal compiler model of Section 5. We describe these prototypes next.

970 Our prototype implementation of MSWasm extends the bytecode of Wasm with instructions to
971 manipulate the segment memory as well as handles. In doing so, it takes a few shortcuts in the name
972 of expediency—most notably, it replaces the existing Wasm opcodes for τ .load and τ .store with
973 τ .segload and τ .segstore. A production MSWasm implementation would support both segment-
974 based and linear-memory-based operations simultaneously, by using two-byte opcode sequences
975 for τ .segload and τ .segstore.
976

977 6.1 Ahead of Time Compilation of MSWasm

981 To compile MSWasm bytecode to machine
 982 code, we build on the rWasm
 983 compiler [Bosamiya et al. 2022]. rWasm
 984 is a provably-safe sandboxing compiler
 985 from Wasm to Rust, and thus to high-
 986 performance machine code.⁷ We ex-
 987 tended rWasm to support MSWasm as fol-
 988 lows. We modified rWasm’s frontend to
 989 parse MSWasm instructions and propa-
 990 gate them through to later phases. We up-
 991 dated rWasm’s stack analysis to account
 992 for MSWasm’s new types and instructions
 993 (e.g., `τ.segload` and `τ.segstore`, which take
 994 a `handle` as argument). Finally, we up-
 995 dated rWasm’s backend—the code gener-
 996 ator, specifically—to implement MSWasm’s
 997 instructions and segment memory.

998 One of the benefits of MSWasm is that it
 999 gives Wasm compilers and runtimes flexi-
 1000 bility in how to best enforce memory safety. This is especially important today: memory-safety
 1001 hardware support is only starting to see deployment and applications have different security-
 1002 performance requirements—we cannot realistically expect everyone to pay the cost of software-
 1003 based memory safety. When hardware becomes available, MSWasm programs can take advantage
 1004 of hardware acceleration almost trivially: in our AOT compiler, for example, we only need to tweak
 1005 the codegen stage. We demonstrate this flexibility by prototyping two different software techniques,
 1006 and one hardware-accelerated technique that have different safety and performance characteristics.

1007 **Segments as Vectors.** Our default technique for memory-safety enforcement closely matches
 1008 Section 3.2.1, and enforces spatial safety, temporal safety, and handle integrity (rWasm_{STH} in
 1009 Section 7). We implement the segment memory as a vector (Vec) of segments. Each segment is
 1010 a pair composed of a Vec of bytes (giving us spatial safety) and a Vec of tags, which is used to
 1011 enforce handle integrity. Handles themselves are implemented using an enum (i.e., a tagged union).
 1012 To enforce temporal safety we clear free segments from memory and use sentinel value to prevent
 1013 the reuse of segment indexes. A slight variation of this technique (rWasm_{ST} in Section 7) gives up
 1014 on handle integrity (we remove the Vec of tags and related checks) for performance.

1015 **Segments with Baggy Bounds.** Our second technique is inspired by Baggy Bounds checking [Akri-
 1016 tidis et al. 2009], which is a technique that performs fast checks at each handle-modifying operation
 1017 and elides checks at loads and stores, enabled by expanding buffers to the next power of two at
 1018 the point of allocation. This technique gives up on handle integrity and temporal safety, since
 1019 accesses are not checked, but is considerably faster (rWasm_S in Section 7). To implement this
 1020 technique, our compiler uses a single growable Vec of bytes, within which a binary buddy allocator
 1021 allocates implicit segment boundaries. We implement the handles as 64-bit values storing an offset
 1022 in memory and the log of the segment size (rounded up to nearest power of two at allocation). We
 1023 emit bounds checks for each operation that might modify handles, ensuring that handles remain
 1024 within the (baggy) bounds of their corresponding segment. Specifically, when handles stray a short
 1025

1026 ⁷In modifying rWasm, we were careful to ensure that we preserve its previously-established sandboxing/isolation guarantees.
 1027 These guarantees, together with the internal memory-safety guarantees from MSWasm, increases the level of protection for
 1028 native code generated by rWasm.

1029




Fig. 12. End-to-end compilation pipeline. We first compile C to MSWasm (via LLVM), and then compile MSWasm to machine code using either our modified rWasm AOT compiler (which supports different notions of safety) or our modified GraalWasm JIT compiler.

Alexandra E. Michael*, Anitha Gollamudi*, Jay Bosamiya, Evan Johnson, Aidan Denlinger, Craig 22 Disselkoe, Conrad Watt, Bryan Parno, Marco Patrignani, Marco Vassena, and Deian Stefan

distance outside their segment, we mark them as such (and they can safely return back), but we trap when they (try to) stray too far.

Hardware acceleration using CHERI. Our third technique is to implement segment memory using CHERI capabilities [Watson et al. 2015]. Each handle is represented as a CHERI capability: a 128-bit register that consists of a capability base, offset, size, and permissions. Each capability also contains a tag bit that tracks the integrity of the capability—that the capability has not been expanded, has not been corrupted in memory, etc. By implementing handles using capabilities, we can enforce handle and spatial safety without software checks. This is because CHERI bounds checks memory accesses (spatial safety) and checks the validity of the tag bit (handle integrity) in hardware. Temporal safety for CHERI code is still a work in progress [Filardo et al. 2020; Xia et al. 2019], and consequently, our MSWasm-on-CHERI prototype does not yet support it.

To implement MSWasm using CHERI, we add a new rWasm backend that emits a subset of CHERI-compatible C code which implements handles as capabilities. Since CHERI capabilities and MSWasm handles are conceptually very similar, this compilation step is straightforward: `segalloc(size)` becomes `calloc(size, 1)`, `segfree(handle)` becomes `free(handle)`, `handle.add` becomes standard pointer addition, etc. We then compile this CHERI-C code to CHERI-Aarch64 using the already mature CHERI LLVM [CTSRD-CHERI 2022] compilation pipeline.

Implementation Effort. Our modifications to rWasm, for both software-only memory enforcement techniques, comprise roughly 1900 lines of additional code. The implementation of these two techniques comprise approximately 500 lines of code each in rWasm’s codegen, and share the rest of rWasm’s codebase. The hardware-accelerated technique also uses our modified rWasm frontend, but could not share much of rWasm’s codegen with the other backends (which target Rust), instead requiring code to support a new target language—CHERI-C (even for the regular non-MSWasm-specific components of Wasm); our modifications for this backend comprise approximately 3000 lines of code. The relative ease of these modifications, both for software- and hardware-based techniques illustrates how MSWasm provides a fertile ground for experimenting with new techniques for providing performant memory safety.

6.2 Just in Time Compilation of MSWasm

Our second prototype is a just-in-time compiler of MSWasm built on top of GraalWasm [Prokopec 2019]. GraalWasm is a Wasm frontend for GraalVM [Oracle 2021a], a JVM-based JIT compiler capable of compiling a wide range of languages through the Truffle framework [Oracle 2021b]. We extend GraalWasm to support MSWasm. Our modifications mirror those we made to rWasm: We modified the GraalWasm frontend to parse MSWasm and the backend—the GraalWasm interpreter in this case—to support MSWasm’s instructions and segment memory model. We were able to reuse the GraalVM JIT compiler unmodified, as it automatically optimizes the AST generated by Truffle from the interpreter.

Segments as Objects. Unlike our rWasm implementation, we only consider one enforcement technique. We pick a middle ground between safety and performance: We enforce spatial and temporal safety, but not handle integrity (Gaal_{ST} in Section 7). Our implementation of memory segments in GraalWasm is similar to our first rWasm technique (but does not track handle-integrity tags). We implement the segment memory as a Java object, `SegmentMemory`, which tracks a list of segments. `SegmentMemory` is backed by Java’s `Unsafe` memory manager, an internal framework that facilitates manual memory management. Unlike objects created on the Java heap, memory allocated through `Unsafe` is not garbage-collected and is accessed directly by pointer addresses. Using Java `Unsafe`, `SegmentMemory` manually allocates a new chunk of memory for each new segment, which lets us avoid the overhead of Java objects in exchange for explicitly tracking

1079 the allocated memory. A segment is represented by a `Segment` object, which contains an address
 1080 within the `Unsafe` memory, the (inclusive) upper bound of the segment in memory, and a randomly
 1081 generated key. To ensure temporal safety, free segments are removed from the list of segments in
 1082 `SegmentMemory`, leaving no way to reference them.

1083 **Implementation Effort.** We added roughly 1200 lines of code to GraalWasm. Our prototype is
 1084 relatively simple and not yet tuned to take full advantage of GraalVM’s optimizations. We leave
 1085 this to future work.
 1086

1087 6.3 Compiling C to MSWasm

1088 MSWasm, like Wasm, is intended to be a compilation target from higher level languages. We
 1089 implement a compiler from C to MSWasm by extending the CHERI fork of Clang and LLVM [CTSRD-
 1090 CHERI 2022]. CHERI modified LLVM to support fat pointers, which share many characteristics
 1091 with MSWasm handles, and is thus a good starting point for MSWasm.

1092 CHERI represents fat pointers at the LLVM IR level as 64- to 512-bit pointers in a special,
 1093 distinguished “address space”; pointers in this address space are lowered to CHERI capabilities
 1094 in the appropriate LLVM backends. CHERI today only targets MIPS and RISC-V (with CHERI
 1095 hardware extensions) backends; other backends, including the Wasm backend, are incompatible
 1096 with CHERI’s fat pointers. We modified the Wasm backend to emit MSWasm bytecode, lowering
 1097 64-bit fat-pointer abstractions to MSWasm abstractions. Since most of the implementation details
 1098 follow from Section 5, we focus on details not captured by our formal model.
 1099

1100 **Global and Static Data.** Our C-to-MSWasm compiler only emits `handle`-based load and store
 1101 operations, resulting in MSWasm programs which do not use the linear memory at all. This provides
 1102 additional safety guarantees (and implementation expediency) at the expense of some flexibility
 1103 (e.g., we do not support integer-to-pointer casts, except for a few special cases like constant 0). One
 1104 consequence of this is that even global variables and static data need to be accessed via `handles`,
 1105 and thus placed in the segment memory.⁸ Our compiler emits instructions to allocate a segment for
 1106 each LLVM global variable and store the corresponding `handle` in a Wasm global variable. When
 1107 the target program needs a pointer to the global array, it simply retrieves the `handle` from the
 1108 appropriate Wasm global variable.

1109 Some global variables in C are themselves pointers, initialized via initialization expressions, and
 1110 need to be pointing to valid, initialized memory at the beginning of the program. Our compiler
 1111 generates the necessary information in the output `.wasm` file to instruct MSWasm compilers and
 1112 runtimes (e.g., `rWasm` and `GraalWasm`) to initialize certain segments at module initialization time.

1113 **C Stack.** We compile part of the C stack to the segment memory. Specifically, stack variables
 1114 whose address-of are taken and stack-allocated arrays cannot be placed on the (simple and safe)
 1115 Wasm stack. Compilers from C to ordinary Wasm place these variables in the linear memory; our
 1116 compiler places them in the segment memory.⁹ We allocate a single large segment to represent
 1117 stack memory for all of the variables which must be allocated in the segment memory; this means
 1118 we have a single stack pointer, which we store in a dedicated Wasm global variable of type `handle`.
 1119 Compared to using a separate segment for each stack allocation, our single-segment implementation
 1120

1121 ⁸More precisely, global variables which the program never takes the address of, do not need this treatment, as we can
 1122 compile them into Wasm globals; but global variables which the program does take the address of, such as global arrays,
 1123 are accessed via pointers and thus must be located in the segment memory.

1124 ⁹Stack variables which the program never takes the address of can be compiled to Wasm local variables, and data such as
 1125 return addresses are never placed in the linear memory at all; Wasm implementations place them on a safe stack which is
 1126 inaccessible to Wasm load and store instructions. The only stack variables which need to be placed in the linear memory, or
 1127 for us the segment memory, are those we need pointers to.

is simpler (and faster) but trades-off some safety, e.g., we cannot prevent a stack buffer overflow from corrupting another stack-allocated buffer.

Standard library. Wasm programs which depend on Libc need a Wasm-compatible implementation of Libc. We modified WASI [WebAssembly [n.d.]] to be compatible with MSWasm to the extent necessary for our benchmarks. Most importantly, we fully recompiled the WASI Libc using our MSWasm compiler, in order to generate Libc bytecode compatible with MSWasm. In our MSWasm version of the WASI Libc, the implementations of malloc and free are completely replaced by trivial implementations consisting of the segalloc and segfree MSWasm instructions.

Implementation Effort. Our CHERI LLVM additions (in particular to its Wasm backend) and the WASI Libc, amounted to approximately 1600 lines of code. While our compiler can target any MSWasm backend, compiling general, real-world applications would likely require additional changes to WASI Libc. We leave this to future work.

7 PERFORMANCE EVALUATION

In this section we describe our performance evaluation of the MSWasm compiler. We use the PolyBenchC benchmarking suite [Pouchet 2011] since PolyBenchC has become the de-facto suite used by almost all Wasm compilers (although limitations to PolyBenchC are noted by [Jangda et al. 2019]). We compare the performance of MSWasm to the performance of the same benchmarks compiled to normal Wasm, on each of our implementations.

Machine setup. We compile all benchmarks from C to Wasm using Clang, and from C to MSWasm using our modified CHERI Clang compiler; in both cases we set the optimization level to -O3. We run all our software-based enforcement benchmarks on a single core on a Linux-based system with an Intel Xeon 8160, and our hardware-accelerated enforcement benchmarks on the ARM Morello platform [ARM 2022].




Fig. 13. Performance of our implementations of MSWasm compared to normal Wasm, normalized against native (non-Wasm) execution on benchmarks from PolyBenchC

Results. Figure 13 summarizes our measurements (see Appendix A for a detailed breakdown), normalized against the execution time of native (non-Wasm) execution. In this figure, rWasm_{Wasm} and Graal_{Wasm} refer to execution of normal Wasm. We distinguish the different MSWasm compilers according to their enforcement techniques: rWasm_{STH} enforces spatial safety, temporal safety, and handle integrity; rWasm_{ST} and Graal_{ST} only enforce spatial and temporal safety; and, rWasm_S only enforces spatial safety (in the style of baggy bounds).

1177 As expected, and in line with prior work [Nagarakatte et al. 2009, 2010], each safety enforcement
 1178 techniques comes with a performance cost—handle integrity being the most expensive. For the
 1179 AOT compiler, we observe that enforcing spatial safety alone $rWasm_S$ has a geomean overhead of
 1180 21.4% over $rWasm_{Wasm}$; additionally enforcing temporal safety ($rWasm_{ST}$) results in an overhead
 1181 of 52.2% over $rWasm_{Wasm}$; and, finally, further enforcing handle integrity ($rWasm_{STH}$) increases
 1182 the end-to-end overhead to 197.5%. For the JIT compiler, enforcing spatial and temporal safety
 1183 results in an overhead comparable to that of the AOT compiler: $Graal_{ST}$ imposes a 42.3% geomean
 1184 overhead. The JIT approach is much slower than the AOT approach though—the overheads of
 1185 $rWasm_{Wasm}$ and $Graal_{Wasm}$ over native (non-Wasm) execution are 71.8% and 3230.0% respectively.
 1186 We also note that with increasing iterations of the GraalVM JIT, $Graal_{Wasm}$'s performance improves
 1187 more rapidly than $Graal_{ST}$'s, which suggests that our implementation still has potential to make
 1188 better use of GraalVM's optimizer.

1189 Our hardware-accelerated approach, which enforces spatial safety and handle integrity (but not
 1190 temporal safety), runs on an entirely distinct architecture and platform. Thus, a direct comparison
 1191 against the same native code baseline used for the other techniques would not be particularly
 1192 instructive. Instead, we evaluate our MSWasm-CHERI backend, against baseline native CHERI code
 1193 (pure capability mode) on the Morello platform, and find an overhead of 51.7%.

1194 Since normal Wasm and MSWasm have different bytecode formats, our evaluation of MSWasm
 1195 performance necessarily includes slowdowns caused by inefficiencies in our compilation from
 1196 C to MSWasm. But because MSWasm decouples memory safety enforcement from the genera-
 1197 tion of MSWasm bytecode, both parts of this pipeline (C-to-MSWasm compilation, and MSWasm
 1198 to machine code) can be independently optimized, with MSWasm performance benefiting from
 1199 improvements on both sides.

1200

1201 8 RELATED WORK

1202

1203 **Comparison with [Disselkoen et al. 2019].** Their work introduces MSWasm and provides
 1204 an informal design; however it only conjectures the memory-safety guarantees. In contrast, our
 1205 current work specifies the design of MSWasm using formal semantics, which makes it possible to
 1206 establish precise memory-safety guarantees, and to provide a specification for a variety of MSWasm
 1207 implementations. Furthermore, we provide the first implementation and evaluation of MSWasm as
 1208 well as a C-to-MSWasm compiler.

1209 **Memory safety for C-like languages.** Despite a tremendous amount of work on memory-safety
 1210 protection mechanisms [Szekeres et al. 2013], researchers still struggle to agree on a common
 1211 definition for *memory safety* [Hicks 2014]. Azevedo de [Azevedo de Amorim et al. 2018] characterize
 1212 memory safety as a 2-hypersafety property, similar to non-interference. Their definition belongs to
 1213 a richer class of security properties, which are harder to enforce and to preserve robustly through
 1214 compilation [Abate et al. 2019].

1215 Many compiler-based instrumentations have been proposed to enforce memory safety in C
 1216 programs via software-based checks attached to pointer and memory operations [Akritidis et al.
 1217 2009; Austin et al. 1994; Jim et al. 2002; Nagarakatte et al. 2009; Nacula et al. 2005; Patil and Fischer
 1218 1997; Ruef et al. 2019; Xu et al. 2004]. Some of these solutions are also supported by formal memory-
 1219 safety guarantees [Nagarakatte et al. 2009, 2010; Ruef et al. 2019]. These formal results however, are
 1220 not *robust*, i.e., they do not guarantee memory safety when linking with arbitrary adversarial code.
 1221 Moreover, these formalizations do not actually include the instrumentation pass of the compiler,
 1222 but prove memory safety via *type safety* of an instrumented C-like language, where pointers are
 1223 annotated with bounds metadata. Unlike MSWasm, these languages adopt a high-level memory
 1224 model, which implicitly provides pointer integrity.

1225

Our color-based memory-safety monitor and similarly our notion of authentic pointers and handles are inspired by previous work on *pointer provenance* in C [Memarian et al. 2019a]. Some of the C semantics proposed in that work track pointer provenance also through integer and pointer casts, which we do not consider in this work, also given that MSWasm has no native notion of casts. Our definition of memory-safety is also related to the micro-policies [d. Amorim et al. 2015]. The main difference is that they use (finite number of) color tags to enforce memory safety whereas we use (possibly infinite) colors to develop a general language-independent definition of memory safety.

Efficient memory-safety implementations. Unlike compiler-based instrumentations, compiling to MSWasm does not commit to a particular concrete strategy for enforcing memory safety: Different implementations of MSWasm can use different enforcement approaches. In particular, MSWasm enables backends compilers and runtimes to leverage efficient software- and hardware-based mechanisms, independently proposed to enforce pointer integrity [Liljestrand et al. 2019], spatial [Akritidis et al. 2009; Arm 2019; Kroes et al. 2018], and temporal [Lee et al. 2015; Parkinson et al. 2017] safety, to create new practical memory-safety enforcement schemes. Because MSWasm is platform-agnostic, we expect that implementations will be able to opportunistically take advantage of hardware memory protection mechanisms on individual platforms [Arm 2019; Devietti et al. 2008; Kwon et al. 2013; Oleksenko et al. 2018] (current and proposed) to efficiently implement handles.

Software isolation via Wasm. Wasm abstractions provide an efficient software-isolation mechanism, which has been applied in many different domains. For example, using Wasm, the RLBox framework [Narayan et al. 2020] retrofits isolation into the Firefox browser; Sledge [Gadepalli et al. 2020] enables lightweight serverless-first computing on the Edge; and eWASM [Peach et al. 2020] demonstrates practical software fault isolation for resource-constrained embedded platforms. These use cases already rely on both the performance and the sandboxing safety of Wasm, and stand to benefit from MSWasm’s focus on memory safety.

[Bosamiya et al. 2022] use formal methods and non-traditional techniques respectively to provide provable isolation between the Wasm module, running as a native library, and the host process executing it. Their focus is on provable module–host isolation, and module-internal memory safety is explicitly left out of scope. As shown by [Lehmann et al. 2020], Wasm lacks many common defenses (e.g., stack canaries, guard pages, ASLR) against classic memory safety vulnerabilities, such as buffer overflows.

[Jangda et al. 2019] perform a large-scale performance evaluation of browser Wasm runtimes, comparing to native code. Our evaluation of MSWasm’s performance (Section 7) shows that adding memory-safety protections does not fundamentally change Wasm’s performance story. In particular, adding spatial and temporal safety imposes less overhead on Wasm than the overhead Wasm already incurs vs native code.

9 CONCLUSION

This paper realised the MSWasm proposal to extend Wasm with language-level memory-safety abstractions, giving it a formal semantics, proving that its programs are all memory safe and implementing the MSWasm language runtime. Like Wasm, MSWasm is intended to be used as a compilation target, so this paper formalised a C-to-MSWasm compiler, proved that it enforces memory safety, and implemented variations of said compiler with different tradeoffs between speed and security. Our PolyBenchC-based evaluation shows that MSWasm introduces an overhead ranging from 22% (enforcing spatial safety alone) to 198% (enforcing full memory safety). Our software-based implementations only serve to highlight that enforcing memory safety for Wasm

1275 is possible and, moreover, that MSWasm makes it easy to change the underlying enforcement
1276 mechanism without modifying application code. This means MSWasm engines will be able to take
1277 advantage of clever memory safety enforcement techniques today and hardware extensions in the
1278 near future, progressively (and transparently) improving the safety of the applications they run.

1279 **ACKNOWLEDGMENTS**

1281 This work was partially supported: by the German Federal Ministry of Education and Research
1282 (BMBF) through funding for the CISPA-Stanford Center for Cybersecurity (FKZ: 13N1S0762),
1283 by the Italian Ministry of Education through funding for the Rita Levi Montalcini grant (call of
1284 2019) by a fellowship from the Alfred P. Sloan Foundation, and by the CONIX Research Center,
1285 one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored
1286 by DARPA. Gollamudi was supported in part through a generous gift to support research on
1287 applied cryptography and society in the Center for Research on Computation and Society (Harvard
1288 University) and Computing Innovating Fellowship (2021). We would also like to thank Reesha
1289 Rajen for their help on earlier stages of this work.

1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323

Alexandra E. Michael*, Anitha Gollamudi*, Jay Bosamiya, Evan Johnson, Aidan Denlinger, Craig 28 Disselkoen, Conrad Watt, Bryan Parno, Marco Patrignani, Marco Vassena, and Deian Stefan

REFERENCES

- 1324 Carmine Abate, Roberto Blanco, Deepak Garg, Cătălin Hrițcu, Marco Patrignani, and Jérémy Thibault. 2019. Journey Beyond
1325 Full Abstraction: Exploring Robust Property Preservation for Secure Compilation. In *2019 IEEE 32th Computer Security*
1326 *Foundations Symposium (CSF 2019)*.
- 1327 Periklis Akrkitidis, Manuel Costa, Miguel Castro, and Steven Hand. 2009. Baggy Bounds Checking: An Efficient and
1328 Backwards-Compatible Defense against Out-of-Bounds Errors. In *USENIX Security Symposium*, Vol. 10.
- 1329 Arm. 2019. Armv8.5-A Memory Tagging Extension. *White Paper* (2019).
- 1330 ARM 2022. Arm Morello Program. <https://www.arm.com/architecture/cpu/morello>.
- 1331 Todd M. Austin, Scott E. Breach, and Gurindar S. Sohi. 1994. Efficient Detection of All Pointer and Array Access Errors.
1332 In *Proceedings of the ACM SIGPLAN 1994 Conference on Programming Language Design and Implementation* (Orlando,
1333 Florida, USA) (*PLDI '94*). Association for Computing Machinery, New York, NY, USA, 290–301. [https://doi.org/10.1145/](https://doi.org/10.1145/178243.178446)
1334 [178243.178446](https://doi.org/10.1145/178243.178446)
- 1335 Arthur Azevedo de Amorim, Cătălin Hrițcu, and Benjamin C. Pierce. 2018. The Meaning of Memory Safety. In *Principles of*
1336 *Security and Trust*, Lujo Bauer and Ralf Küsters (Eds.). Springer International Publishing, Cham, 79–105.
- 1337 Jay Bosamiya, Wen Shih Lim, and Bryan Parno. 2022. Provably-Safe Multilingual Software Sandboxing using WebAssembly.
1338 In *USENIX Security Symposium*.
- 1339 CTSRD-CHERI. 2022. The CHERI LLVM Compiler Infrastructure. <https://github.com/CTSRD-CHERI/llvm-project>.
- 1340 A. A. d. Amorim, M. Dènès, N. Giannarakis, C. Hritcu, B. C. Pierce, A. Spector-Zabusky, and A. Tolmach. 2015. Micro-
1341 Policies: Formally Verified, Tag-Based Security Monitors. In *2015 IEEE Symposium on Security and Privacy*. 813–830.
1342 <https://doi.org/10.1109/SP.2015.55>
- 1343 Thurston HY Dang, Petros Maniatis, and David Wagner. 2017. Oscar: A Practical page-permissions-based scheme for
1344 thwarting dangling pointers. In *USENIX Security*.
- 1345 Joe Devietti, Colin Blundell, Milo M. K. Martin, and Steve Zdancewic. 2008. Hardbound: Architectural Support for Spatial
1346 Safety of the C Programming Language. In *Proceedings of the 13th International Conference on Architectural Support for*
1347 *Programming Languages and Operating Systems*. ACM, New York, NY, USA.
- 1348 Craig Disselkoen, John Renner, Conrad Watt, Tal Garfinkel, Amit Levy, and Deian Stefan. 2019. Position Paper: Progressive
1349 Memory Safety for WebAssembly. In *Proceedings of the 8th International Workshop on Hardware and Architectural Support*
1350 *for Security and Privacy*. ACM.
- 1351 Nathaniel Wesley Filardo, Brett F Gutstein, Jonathan Woodruff, Sam Ainsworth, Lucian Paul-Trifu, Brooks Davis, Hongyan
1352 Xia, Edward Tomasz Napierala, Alexander Richardson, John Baldwin, et al. 2020. Cornucopia: Temporal safety for CHERI
1353 heaps. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 608–625.
- 1354 Phani Kishore Gadepalli, Sean McBride, Gregor Peach, Ludmila Cherkasova, and Gabriel Parmer. 2020. Sledge: A Serverless-
1355 First, Light-Weight Wasm Runtime for the Edge. In *Proceedings of the 21st International Middleware Conference (Middleware*
1356 *'20)*. ACM.
- 1357 Richard Grisenthwaite. 2019. Supporting the UK in becoming a leading global player in cybersecurity. [https://community.](https://community.arm.com/blog/company/b/blog/posts/supporting-the-uk-in-becoming-a-leading-global-player-in-cybersecurity)
1358 [arm.com/blog/company/b/blog/posts/supporting-the-uk-in-becoming-a-leading-global-player-in-cybersecurity](https://community.arm.com/blog/company/b/blog/posts/supporting-the-uk-in-becoming-a-leading-global-player-in-cybersecurity).
- 1359 Andreas Haas, Andreas Rossberg, Derek L. Schuff, Ben L. Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon
1360 Zakai, and JF Bastien. 2017a. Bringing the Web up to Speed with WebAssembly. In *Proceedings of the 38th ACM*
1361 *SIGPLAN Conference on Programming Language Design and Implementation* (Barcelona, Spain) (*PLDI 2017*). Association
1362 for Computing Machinery, New York, NY, USA, 185–200. <https://doi.org/10.1145/3062341.3062363>
- 1363 Andreas Haas, Andreas Rossberg, Derek L. Schuff, Ben L. Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon
1364 Zakai, and JF Bastien. 2017b. Bringing the Web up to Speed with WebAssembly. *SIGPLAN Not.* 52, 6 (June 2017), 185–200.
1365 <https://doi.org/10.1145/3140587.3062363>
- 1366 Istvan Haller, Yuseok Jeon, Hui Peng, Mathias Payer, Cristiano Giuffrida, Herbert Bos, and Erik Van Der Kouwe. 2016.
1367 TypeSan: Practical type confusion detection. In *Conference on Computer and Communications Security*. ACM.
- 1368 Michael Hicks. 2014. What is memory safety? <http://www.pl-enthusiast.net/2014/07/21/memory-safety/>.
- 1369 Aaron Hilbig, Daniel Lehmann, and Michael Pradel. 2021. An Empirical Study of Real-World WebAssembly Binaries: Security,
1370 Languages, Use Cases. ACM.
- 1371 Abhinav Jangda, Bobby Powers, Emery D. Berger, and Arjun Guha. 2019. Not So Fast: Analyzing the Performance of
1372 WebAssembly vs. Native Code. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*. USENIX Association.
- 1373 Trevor Jim, Greg Morrisett, Dan Grossman, Michael Hicks, James Cheney, and Yanling Wang. 2002. Cyclone: A Safe Dialect
1374 of C. In *2002 USENIX Annual Technical Conference (USENIX ATC 02)*. USENIX Association.
- 1375 Taddeus Kroes, Koen Koning, Erik van der Kouwe, Herbert Bos, and Cristiano Giuffrida. 2018. Delta Pointers: Buffer
1376 Overflow Checks without the Checks. In *Proceedings of the Thirteenth EuroSys Conference*. ACM.
- 1377 Albert Kwon, Udit Dhawan, Jonathan M. Smith, Thomas F. Knight, and Andre DeHon. 2013. Low-Fat Pointers: Compact
1378 Encoding and Efficient Gate-Level Implementation of Fat Pointers for Spatial Safety and Capability-Based Security. In
1379 *Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security*. ACM.

- 1373 Byoungyoung Lee, Chengyu Song, Yeongjin Jang, Tielei Wang, Taesoo Kim, Long Lu, and Wenke Lee. 2015. Preventing
1374 Use-after-free with Dangling Pointers Nullification.. In *NDSS. The Internet Society*. [http://dblp.uni-trier.de/db/conf/](http://dblp.uni-trier.de/db/conf/ndss/ndss2015.html#LeeSJWKL15)
1375 [ndss/ndss2015.html#LeeSJWKL15](http://dblp.uni-trier.de/db/conf/ndss/ndss2015.html#LeeSJWKL15)
- 1376 Daniel Lehmann, Johannes Kinder, and Michael Pradel. 2020. Everything Old is New Again: Binary Security of WebAssembly.
1377 In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association.
- 1378 Michael LeMay, Joydeep Rakshit, Sergej Deutsch, David M Durham, Santosh Ghosh, Anant Nori, Jayesh Gaur, Andrew
1379 Weiler, Salmin Sultana, Karanvir Grewal, et al. 2021. Cryptographic Capability Computing. In *IEEE/ACM International*
1380 *Symposium on Microarchitecture*. ACM.
- 1381 Xavier Leroy. 2009. A Formally Verified Compiler Back-end. *Journal of Automated Reasoning* 43, 4 (2009), 363–446.
1382 <http://dx.doi.org/10.1007/s10817-009-9155-4>
- 1383 Hans Liljestrand, Thomas Nyman, Kui Wang, Carlos Chinae Perez, Jan-Erik Ekberg, and N. Asokan. 2019. PAC it up:
1384 Towards Pointer Integrity using ARM Pointer Authentication. In *28th USENIX Security Symposium (USENIX Security 19)*.
1385 USENIX Association.
- 1386 Kayvan Memarian, Victor B. F. Gomes, Brooks Davis, Stephen Kell, Alexander Richardson, Robert N. M. Watson, and Peter
1387 Sewell. 2019a. Exploring C Semantics and Pointer Provenance. *Proc. ACM Program. Lang.* 3, POPL, Article 67 (Jan. 2019),
1388 32 pages. <https://doi.org/10.1145/3290380>
- 1389 Kayvan Memarian, Victor B. F. Gomes, Brooks Davis, Stephen Kell, Alexander Richardson, Robert N. M. Watson, and Peter
1390 Sewell. 2019b. Exploring C Semantics and Pointer Provenance. *Proc. ACM Program. Lang.* 3, POPL, Article 67 (jan 2019),
1391 32 pages. <https://doi.org/10.1145/3290380>
- 1392 Kayvan Memarian, Justus Matthiesen, James Lingard, Kyndylan Nienhuis, David Chisnall, Robert N. M. Watson, and Peter
1393 Sewell. 2016. Into the Depths of C: Elaborating the de Facto Standards. *SIGPLAN Not.* 51, 6 (June 2016), 1–15.
- 1394 Santosh Nagarakatte, Jianzhou Zhao, Milo M.K. Martin, and Steve Zdancewic. 2009. SoftBound: Highly Compatible and
1395 Complete Spatial Memory Safety for C. *SIGPLAN Not.* 44, 6 (June 2009), 245–258. <https://doi.org/10.1145/1543135.1542504>
- 1396 Santosh Nagarakatte, Jianzhou Zhao, Milo M.K. Martin, and Steve Zdancewic. 2010. CETS: Compiler Enforced Temporal
1397 Safety for C. *SIGPLAN Not.* 45, 8 (June 2010), 31–40. <https://doi.org/10.1145/1837855.1806657>
- 1398 Shrayan Narayan, Craig Disselkoen, Tal Garfinkel, Nathan Froyd, Eric Rahm, Sorin Lerner, Hovav Shacham, and Deian
1399 Stefan. 2020. Retrofitting Fine Grain Isolation in the Firefox Renderer. In *29th USENIX Security Symposium (USENIX*
1400 *Security 20)*. USENIX Association.
- 1401 George C. Necula, Jeremy Condit, Matthew Harren, Scott McPeak, and Westley Weimer. 2005. CCured: Type-safe Retrofitting
1402 of Legacy Software. *ACM Trans. Program. Lang. Syst.* 27, 3 (May 2005), 477–526. <https://doi.org/10.1145/1065887.1065892>
- 1403 Oleksii Oleksenko, Dmitrii Kuvaishii, Pramod Bhatotia, Pascal Felber, and Christof Fetzer. 2018. Intel MPX Explained: A
1404 Cross-Layer Analysis of the Intel MPX System Stack. *Proc. ACM Meas. Anal. Comput. Syst.* 2, 2, Article 28 (jun 2018),
1405 30 pages. <https://doi.org/10.1145/3224423>
- 1406 Aleph One. 1996. Smashing the stack for fun and profit. *Phrack magazine* 7, 49 (1996), 14–16.
- 1407 Oracle. 2021a. GraalVM. <https://www.graalvm.org/>.
- 1408 Oracle. 2021b. Truffle Language Implementation Framework. [https://www.graalvm.org/22.0/graalvm-as-a-platform/](https://www.graalvm.org/22.0/graalvm-as-a-platform/language-implementation-framework/)
1409 [language-implementation-framework/](https://www.graalvm.org/22.0/graalvm-as-a-platform/language-implementation-framework/).
- 1410 Matthew Parkinson, Kapil Vaswani, Dimitrios Vytiniotis, Manuel Costa, Pantazis Deligiannis, Aaron Blankstein, Dy-
1411 lan McDermott, and Jonathan Balkind. 2017. *Project Snowflake: Non-blocking safe manual memory management*
1412 *in .NET*. Technical Report MSR-TR-2017-32. Microsoft. [https://www.microsoft.com/en-us/research/publication/](https://www.microsoft.com/en-us/research/publication/project-snowflake-non-blocking-safe-manual-memory-management-net/)
1413 [project-snowflake-non-blocking-safe-manual-memory-management-net/](https://www.microsoft.com/en-us/research/publication/project-snowflake-non-blocking-safe-manual-memory-management-net/)
- 1414 Harish Patil and Charles Fischer. 1997. Low-Cost, Concurrent Checking of Pointer and Array Accesses in C Programs.
1415 *Softw. Pract. Exper.* 27, 1 (jan 1997), 87–110.
- 1416 Marco Patrignani. 2020. Why Should Anyone use Colours? or, Syntax Highlighting Beyond Code Snippets. *CoRR*
1417 [abs/2001.11334](https://arxiv.org/abs/2001.11334) (2020). [arXiv:2001.11334](https://arxiv.org/abs/2001.11334) <https://arxiv.org/abs/2001.11334>
- 1418 Gregor Peach, Runyu Pan, Zhuoyi Wu, Gabriel Parmer, Christopher Haster, and Ludmila Cherkasova. 2020. eWASM: Practical
1419 Software Fault Isolation for Reliable Embedded Devices. *IEEE Transactions on Computer-Aided Design of Integrated*
1420 *Circuits and Systems* 39, 11 (2020), 3492–3505. <https://doi.org/10.1109/TCAD.2020.3012647>
- 1421 Louis-Noel Pouchet. 2011. PolyBench-C: the Polyhedral Benchmark suite. [https://web.cs.ucla.edu/~pouchet/software/](https://web.cs.ucla.edu/~pouchet/software/polybench/)
1422 [polybench/](https://web.cs.ucla.edu/~pouchet/software/polybench/). Accessed: March 2022.
- 1423 Aleksandar Prokopec. 2019. Announcing GraalWasm – a WebAssembly engine in GraalVM. [https://medium.com/graalvm/](https://medium.com/graalvm/announcing-graalwasm-a-webassembly-engine-in-graalvm-25cd0400a7f2)
1424 [announcing-graalwasm-a-webassembly-engine-in-graalvm-25cd0400a7f2](https://medium.com/graalvm/announcing-graalwasm-a-webassembly-engine-in-graalvm-25cd0400a7f2).
- 1425 Andrew Ruef, Leonidas Lampropoulos, Ian Sweet, David Tarditi, and Michael Hicks. 2019. Achieving Safety Incrementally
1426 with Checked C. In *Principles of Security and Trust*. Springer.

Alexandra E. Michael*, Anitha Gollamudi*, Jay Bosamiya, Evan Johnson, Aidan Denlinger, Craig
30 Dasselkoen, Conrad Watt, Bryan Parno, Marco Patrignani, Marco Vassena, and Deian Stefan

1422 Fred B. Schneider. 2000. Enforceable Security Policies. *ACM Trans. Inf. Syst. Secur.* 3, 1 (feb 2000), 30–50. <https://doi.org/10.1145/353323.353382>

1423 Laszlo Szekeres, Mathias Payer, Tao Wei, and Dawn Song. 2013. SoK: Eternal War in Memory. In *Proceedings of the 2013*
1424 *IEEE Symposium on Security and Privacy (SP '13)*. IEEE Computer Society.

1425 Gang Tan. 2017. *Principles and Implementation Techniques of Software-Based Fault Isolation*. Vol. 1. Now Publishers Inc.
1426 137–198 pages.

1427 Robert Wahbe, Steven Lucco, Thomas E. Anderson, and Susan L. Graham. 1993. Efficient Software-based Fault Isolation. In
1428 *Proceedings of the Fourteenth ACM Symposium on Operating Systems Principles (SOSP '93)*. ACM.

1429 Robert N.M. Watson, Jonathan Woodruff, Peter G. Neumann, Simon W. Moore, Jonathan Anderson, David Chisnall, Nirav
1430 Dave, Brooks Davis, Khilan Gudka, Ben Laurie, Steven J. Murdoch, Robert Norton, Michael Roe, Stacey Son, and Munraj
1431 Vadera. 2015. CHERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization. In *2015*
1432 *IEEE Symposium on Security and Privacy*. 20–37. <https://doi.org/10.1109/SP.2015.9>

1433 WebAssembly. [n.d.]. WebAssembly System Interface. <https://github.com/WebAssembly/wasi>.

1434 Hongyan Xia, Jonathan Woodruff, Sam Ainsworth, Nathaniel W. Filardo, Michael Roe, Alexander Richardson, Peter Rugg,
1435 Peter G. Neumann, Simon W. Moore, Robert N. M. Watson, and Timothy M. Jones. 2019. CHERIvoke: Characterising
1436 Pointer Revocation Using CHERI Capabilities for Temporal Memory Safety. In *Proceedings of the 52nd Annual IEEE/ACM*
1437 *International Symposium on Microarchitecture (Columbus, OH, USA) (MICRO '52)*. Association for Computing Machinery,
1438 New York, NY, USA, 545–557. <https://doi.org/10.1145/3352460.3358288>

1439 Wei Xu, Daniel C. DuVarney, and R. Sekar. 2004. An Efficient and Backwards-Compatible Transformation to Ensure Memory
1440 Safety of C Programs. In *Proceedings of the 12th ACM SIGSOFT Twelfth International Symposium on Foundations of Software*
1441 *Engineering*. ACM.

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504

A DETAILED EVALUATION BREAKDOWN OF OUR IMPLEMENTATIONS OF MSWASM

Fig. 14. A detailed per-program breakdown of the performance of our implementations of MSWasm compared to normal Wasm, normalized against native (non-Wasm) execution on benchmarks from PolyBenchC. Rather than relying on less accurate external measurements using `time(1)`, we use PolyBenchC's own internal execution time reporting (i.e., `-DPOLYBENCH_TIME`).