

Constructing Call Graphs for WebAssembly using Refined Numeric Type Analysis

ANONYMOUS AUTHOR(S)

WebAssembly has become a widely used compilation target for languages such as C/C++, Rust, and Go, enabling high-performance applications to run in browsers and other sandboxed environments. Static analysis of WebAssembly is important for tasks such as code debloating, performance optimization, and detection of security vulnerabilities. We present *Refined Numeric Type Analysis (RNTA)*, a refinement-type-based technique for constructing WebAssembly call graphs that are more precise than those obtained using previous techniques. WebAssembly's low-level nature creates two significant challenges for static analysis: (i) memory is represented as an unstructured array of bytes, and (ii) indirect calls are encoded as numeric i32 values indexing function tables. Traditional pointer analysis techniques rely on associating sets of objects with pointers and cannot be applied to WebAssembly, where the concepts of pointers and objects do not exist. Existing call graph construction algorithms for WebAssembly associate a set of functions with each `call_indirect` instruction using function signatures, but suffer from the limited nature of WebAssembly's type system, which reduces opportunities for ruling out functions as call targets. The foundation of RNTA is a refinement type system that distinguishes i32 values that are used to index memory—i.e., pointers—from other i32 values and accordingly infers a refinement signature for each function. RNTA computes call graphs on-the-fly by using the refinement types inferred at indirect calls to approximate possible call targets. We prove the type system sound and show that RNTA produces call graphs that are significantly more precise than those computed by the previous state-of-the-art. Across a corpus of 20 real-world WebAssembly binaries, RNTA reduces reachable functions by 9.9% and call graph edges by 14.6% while increasing indirect calls resolved to a unique target by 12.6%.

CCS Concepts: • **Software and its engineering** → **Automated static analysis**.

Additional Key Words and Phrases: WebAssembly, Static Analysis, Call Graph Analysis, Pointer Analysis

1 INTRODUCTION

WebAssembly [24] is a portable, low-level bytecode format that was designed to run computationally intensive tasks in browsers alongside JavaScript host code. It has become widely used as a compilation target for C/C++, Rust, and Go, so that libraries written in those languages can be invoked directly from within a web application. Our long-term goal is to develop static analysis techniques for WebAssembly that can be used for code-size reduction, optimization, and detection of security vulnerabilities. As a first step towards this goal, this paper presents a *Refined Numeric Type Analysis (RNTA)*, a refinement-type-based technique for constructing WebAssembly call graphs that are more precise than those obtained using previous techniques.

During the past 30+ years, a vast literature of algorithms for pointer analysis and call graph construction for high-level object-oriented and imperative programming languages has been developed (see, e.g., [5–7, 10, 17, 22, 23, 38, 40, 45]). In essence, these techniques involve associating sets of abstract objects with abstract pointers. Unfortunately, since WebAssembly is a low-level language that features neither objects nor pointers, these traditional pointer-analysis algorithms cannot be used. Furthermore, analysis at the source level is often insufficient: WebAssembly modules are frequently compiled from libraries, and many optimization opportunities only become apparent once it is known how a host application uses the library's features. Additionally, many web applications only include the binaries of the WebAssembly libraries that they use, leaving out their source code and the provenance of the binaries is often unclear.

WebAssembly's low-level type system comprises only¹ 32- and 64-bit integer (i32, i64) and floating-point (f32, f64) types, presenting significant challenges to call graph construction. Crucially,

¹In this paper, we focus on WebAssembly 1.0 as the first and most widely deployed version of WebAssembly[25].

i32 values serve multiple purposes: indexing memory, controlling conditionals, indexing function tables, and performing arithmetic operations. As a result of this impoverished type system, many functions share the same type, so type-based filtering rarely identifies “monomorphic” call sites that have a single target, limiting key optimizations like call devirtualization and inlining.

In WebAssembly, indirect calls are mediated through a function table, where an integer index is used to select a specific function from the table. The reader might therefore think that a value analysis could be used to determine the index used in indirect calls. However, Lehmann et al. [29] identify several unique challenges to WebAssembly call graph analysis, validated on 8,461 WebAssembly binaries [25]. They find that, in nearly all cases, the dispatch of function calls involves non-constant values with complex data-flow (interprocedural data-flow or loads from memory addresses determined at runtime), complicating call graph construction. As a result, existing static analyses for WebAssembly have made limited progress beyond leveraging WebAssembly’s type system. WASMA [11] uses type-based filtering. WASSAIL [42] additionally restricts its analysis to constant indices, which empirical evidence shows rarely occur in practice [29]. STURDY [28] employs abstract interpretation to track the data flow of abstract integers inter-procedurally as well as through constant memory addresses. However, integer values frequently flow through dynamically assigned memory addresses, leading to significant loss of precision with this approach.

In this paper, we present a refinement type system that refines i32 types to either a singleton pointer type $\text{ptr}(l, n)$ (i.e., a pointer to symbolic location l at abstract offset n) or a singleton number $\text{num}(n)$. Additionally, we type WebAssembly memory, which the original type system leaves untyped, enabling precise tracking of integer values as they flow through memory. Our approach is inspired by recent work on inferring richer types for LLVM IR [31, 32] and binary code [8]. Like those approaches, we enrich the type information of a low-level language. However, our refinement type system is designed specifically for WebAssembly and to enable more precise resolution of indirect calls. This refinement yields more precise function signatures: functions accepting pointers are now distinguished from those accepting numbers. As a result, type-based filtering at indirect call sites becomes significantly more effective because fewer functions share identical signatures. We prove the refinement type system to be type-safe, another distinction from LLVM type analyses.

Our new type system forms the foundation for *Refined Numeric Type Analysis*, an efficient call graph construction algorithm for WebAssembly, which we implement in a tool called RNTA and evaluate on 20 real-world WebAssembly binaries. The results of our experiments show that, on average, when compared against a baseline analysis that uses WebAssembly’s standard type system, RNTA reduces reachable functions by 9.9% and call graph edges by 14.6% while increasing indirect calls that are resolved to a unique target by 12.6%.

In summary, this paper makes the following contributions:

- (1) a refinement type system for WebAssembly that refines the value type i32 to distinguish between numbers and pointers to WebAssembly memory,
- (2) RNTA, a call graph construction algorithm that leverages this refinement type system,
- (3) an empirical evaluation of RNTA on 20 binaries showing, on average, reductions in the number of reachable functions and call graph edges by 9.9% and 14.6%, respectively, while increasing indirect calls that are resolved to a unique target by 12.6%, and
- (4) a proof that the Refinement Type System is type-safe.

The rest of this paper is organized as follows. Section 2 provides background and conveys the intuition behind our approach. Section 3, 5 and 4 cover the refinement type system. An empirical evaluation of RNTA is presented in Section 6. Section 7 presents related work and Section 8 concludes.

99

100

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

145

146

147



Fig. 1. A Rust Library compiled to WebAssembly.

2 BACKGROUND AND MAIN IDEAS

2.1 Call Graph Analysis Over the WebAssembly Binary

A call graph analysis over the source code of the Rust library in Figure 1 will resolve the targets for the indirect calls at R13 and R32 as `num_add`, `num_sub` and `arr_add`, `arr_sub`, respectively. Unfortunately, a call graph analysis over the corresponding WebAssembly binary produced by the compiler is not as precise. Indirect calls in WebAssembly make use of the `call_indirect` instruction and are mediated through a WebAssembly table that contains a list of functions (line W31). WebAssembly is a stack machine and instructions push and pop values from the stack. At runtime, the `call_indirect` instruction pops an `i32` value from the stack which it uses to index into the table, thus determining which function to call. The precision of a call graph analysis depends on how indirect calls are handled. Current state-of-the-art analysis tools handle indirect calls in one of three ways:

- (1) *Naive Analysis*: The target of an indirect call could be any function in the WebAssembly function table. Industry tools like WASM-OPT and WASM-METADCE [48] take this approach.
- (2) *Type-based Analysis*: The syntax of a `call_indirect` instruction contains a function type annotation (line W23) which is guaranteed to match the function type of the indirect call target. State-of-the-art analysis tools like WASSAIL [42] and WASMA [11] restrict the set of possible targets to be the functions in the table whose type matches the type annotation.
- (3) *Index Analysis*: Since an `i32` value is used to index into the function table, tools like WASILLY [33] and STURDY [28] perform a value analysis to determine the set of possible targets.

Unfortunately, for the example in Figure 1, all three strategies yield the same result. The indirect calls at W23 and W28 have the same type annotation: `type0`, a function type that takes two `i32`s as arguments and returns an `i32`. The number and array add and subtract operations all compile down to have the same function type `type0` and the WebAssembly table contains functions for all of these operations. Additionally, the index into the table is provided by the user at runtime, and so is unknown. Hence, a static analysis over this binary would resolve all four operations to be potential targets for both indirect calls. In WebAssembly, the type system is not rich enough to distinguish between different uses of `i32`s. In particular, there is no way to distinguish between `i32`s

that represent numbers and `i32`s that represent pointers to the WebAssembly memory and so, there is no way to distinguish `num_add` from `arr_add`.

2.2 Refining the WebAssembly Type System With Pointers

WebAssembly does not distinguish between numbers and pointers. All `i32` values (32-bit integers) serve multiple purposes: indexing memory, controlling conditionals, indexing the function table, and arithmetic operations. However, we hypothesize that `i32` values used as pointers are consistently used to load from and store to memory. For example, in Figure 1, function `arr_add` treats its first parameter (accessed using `local.get 0`) as a pointer—adding an offset of 12 and loading from that memory location at `W15 – W18`. The second parameter (accessed using `local.get 1`) acts as an array size, triggering an early return at `W08 – W13` when zero. This corresponds to how `WASM-PACK` compiles Rust arrays as a pointer-size pair.

Correspondingly, we refine `i32` to two new types, `ptr` and `num`, both subtypes of `i32`, where `ptr` represents pointers to the WebAssembly memory. We also change the type system to be a refinement type system [14, 19, 21, 35] and type the WebAssembly memory to be able to track reads and writes to memory. The changes to the type system are described in Section 3.2 and are not extensive. We only refine `i32`'s and do not refine the other base WebAssembly types, `i64`, `f32` and `f64`. Our hypothesis is that in working with even a simple set of refined types for WebAssembly, we are able to be more precise in our estimation of indirect call targets simply by virtue of richer type information.

To discover these refinement types, we first generate constraints over the WebAssembly binary under inspection. These constraints encode how data is used. For instance, if a WebAssembly instruction loads from or stores to an `i32` then that type should be refined to a pointer. For simplicity, we generate constraints over `WAFFLE` [13], an intermediate representation for WebAssembly that uses Static Single Assignment Form [15]. We then solve the constraints, while performing a value and pointer analysis, to obtain refinement types over the Wasm binary. We then apply a type-based call graph analysis over the refinement type system. We explain the details of this process in Section 3 and 5. Note that we only generate and solve constraints over WebAssembly 1.0.

3 A SYMBOLIC REFINEMENT TYPE SYSTEM FOR WEBASSEMBLY

3.1 Background on WebAssembly Instructions and Typing

Below we discuss a subset of WebAssembly instructions and their typing. For a description of the complete instruction set and module typing, please refer to Haas et. al.[24] and the WebAssembly 1.0 specification [47]. Instructions in WebAssembly operate on an implicit value stack by popping argument values and pushing computed results. As shown in Figure 2, the instruction typing judgment in WebAssembly has the form $C^w \vdash e^* : \tau_w^m \rightarrow \tau_w^n$ which says that in a context C^w , the one-or-more instructions e^* expect a sequence of m values of types τ_w^i (for $i \in \{1, \dots, m\}$) at the top of the stack and replace these with a sequence of n values of types τ_w^j (for $j \in \{1, \dots, n\}$) at the top of the stack. Note that here we use subscripts and superscripts w for grammars of original WebAssembly elements, while later we will use r for our refinement-typed WebAssembly.

3.1.1 Values and Arithmetic Instructions. WebAssembly has four value types: `i32`, `i64`, `f32`, `f64`. They represent 32- and 64-bit integers and 32- and 64-bit floating-point numbers. Values in WebAssembly are numeric constants tagged with the appropriate value type, e.g., the 32-bit integer 42 is represented as `i32.const 42`. The `i32.const 42` instruction pops nothing from the stack and pushes the `i32.const 42` value onto the stack. This is reflected in the typing rule for constants, as shown in Figure 2. The figure also contains the typing rule for binary operations, which expect two values of

$C^w ::= \{\text{func } tf_w^*, \text{ local } \tau_w^*, \text{ global } \tau_w^*, \text{ table } n^?, \text{ memory } n^?, \text{ label } (\tau_w^*)^?, \text{ return } (\tau_w^*)^?\}$

Typing Instructions With the Original WebAssembly Value Types.

$C^w \vdash e^* : tf_w$

$\tau_w ::= i32 \mid i64 \mid f32 \mid f64$
 $tf_w ::= \tau_w^* \rightarrow \tau_w^*$

$$\begin{array}{c}
 \frac{}{C^w \vdash \tau_w.\text{const } c : \epsilon \rightarrow \tau_w} \text{CONSTANT} \qquad \frac{}{C^w \vdash \tau_w.\text{binop} : \tau_w \tau_w \rightarrow \tau_w} \text{BINARY OPS} \\
 \\
 \frac{tf_w = \tau_w^m \rightarrow \tau_w^n \quad C^w, \text{label}(\tau_w^n) \vdash e^* : tf_w}{C^w \vdash \text{block } tf_w e^* \text{ end} : tf_w} \text{BLOCK} \qquad \frac{C_{\text{label}}^w(i) = \tau_w^m}{C^w \vdash \text{br } i : \tau_w^* \tau_w^m \rightarrow \tau_w^*} \text{BREAK} \\
 \\
 \frac{C_{\text{func}}^w(i) = tf_w}{C^w \vdash \text{call } i : tf_w} \text{CALL} \qquad \frac{C_{\text{memory}}^w = n \quad 2^a \leq (|tp| <) ? |\tau_w| \quad (tp_sz)^? = \epsilon \vee \tau_w = \text{im}}{C^w \vdash \tau_w.\text{load}(tp_sx)^? a o : i32 \rightarrow \tau_w} \text{LOAD}
 \end{array}$$

Fig. 2. Typing of a subset of WebAssembly Instructions in the Original Type System.

the type they are annotated with on the stack and produce a value of the same type. For example, `i32.add` expects two `i32`'s on the stack and pushes a `i32` onto the stack.

3.1.2 Control Constructs and Breaks. WebAssembly has control constructs such as blocks and loops and provides structured control flow with break instructions that are annotated with an index: `br i`. Here, i is a de-bruijn index [16] out of n labels that are associated with n enclosing control constructs. The target construct at i determines how the `br` instruction behaves — if it is a block, control jumps to the *end* of the block and if it is a loop, control jumps to the *start* of a loop. This is evident in the typing of the block and `br` instructions. Blocks are annotated with a signature tf_w that expects τ_w^m types on the top of the stack at the start of the block, and expects τ_w^n types on the top of the stack at the end of a block. The sequence of instructions in a block is type checked under a label that states that this enclosing block expects τ_w^n values on the stack. A `br` instruction finds the label associated with the i^{th} enclosing target control construct and before jumping to the end or start of the target, ensures it has the right types on the stack. If the target construct is a block, it would expect τ_w^n on the top of the stack.

3.1.3 Loads and Stores to Linear Memory. Loads and stores to WebAssembly memory are done with load and store instructions that are annotated with the type of data loaded from memory and stored in memory respectively. For example, `i64.load` loads a `i64` value from memory. WebAssembly memory is only indexed by `i32` values and so the `i64.load` expects an `i32` value on the stack and pushes an `i64` value to the stack.

3.1.4 Function Calls. Functions in WebAssembly are referenced using an immediate index into the function section of a module. Hence, a call instruction is annotated with a index i , identifying the function to be called. This index is used to look up the function type of the called function in the context C^w . The function type specifies the types at the top of the stack before and after the call.

3.2 Symbolic Refinement Type System

We present a refinement over the original WebAssembly (version 1.0) type system in Figure 3. We discuss each change to the type system below:

- (1) *Symbolic Refinement Types, $\hat{\tau}_r$:* We introduce refinements to each base WebAssembly value type τ_w . For example, if the top of the stack is typed to be τ_w in WebAssembly, it is ascribed the refinement type $\tau_w(n)$ in our type system, where n is an abstract number. A refinement

246

Refinement Type System

247

248

Abstract Number n

249

Symbolic Location l

250

Abstract Refinement Type α

251

Symbolic Refinement Type $\hat{\tau}_r ::= i32(n) \mid ptr(l, n) \mid num(n) \mid i64(n) \mid f32(n) \mid f64(n)$

252

Type Tag $\tau ::= i32 \mid ptr \mid num \mid i64 \mid f32 \mid f64$

253

Original Value Types $\tau_w ::= i32 \mid i64 \mid f32 \mid f64$

254

Refinement Function Type $tf_{\hat{\tau}_r} ::= \hat{\tau}_r^* \rightarrow \hat{\tau}_r^*$

255

Symbolic Function Type $tf_{\alpha} ::= \alpha^* \rightarrow \alpha^*$

256

Symbolic Global Type $tg_{\alpha} ::= mut^? \alpha^*$

257

Abstract Memory Typing ζ

258

Symbolic Memory Typing $\Sigma ::= \cdot \mid \Sigma, (l, n) \mapsto \hat{\tau}_r$

259

Constraint System

260

261

Operations $op ::= unop_{iN} \mid unop_{fN} \mid cvtop \mid binop_{iN} \mid binop_{fN} \mid testop_{iN} \mid$

262

 $relop_{iN} \mid relop_{fN}$

263

Type Constraint $t ::= \alpha \doteq \hat{\tau}_r \mid \alpha <: \tau \mid \alpha \doteq \alpha \mid \alpha \doteq \sqcup \alpha^+ \mid \alpha \doteq \zeta[\alpha]_{(tp, sx^2, a, o)} \mid \alpha \doteq op((\alpha^2)^+)$

264

Memory Constraint $m ::= \zeta \doteq \Sigma \mid \zeta \doteq \zeta \mid \zeta \doteq \sqcup \zeta^+ \mid \zeta \doteq \zeta[\alpha \mapsto \alpha]_{(tp^2, a, o)}$

265

Constraint Set $S ::= \cdot \mid S, t \mid S, m$

266

Fig. 3. Syntax for Symbolic Refinement Types and Constraints Generated Over WebAssembly.

267

type such as $i32(42)$, in a more traditional presentation of refinement types, might be written as $\{n: \beta \mid \beta <: i32 \wedge n = 42\}$. For types $i32(n)$ and $f32(n)$, $n \in \{\mathbb{N}_{32} \cup \{\top, \perp\}\}$, while for $i64(n)$ and $f64(n)$, $n \in \{\mathbb{N}_{64} \cup \{\top, \perp\}\}$.

270

(2) *Pointers and Numbers*: Since $i32$ s are used both as pointers into the WebAssembly memory and numbers, we add two new types ptr and num , where ptr and num are subtypes of $i32$. Pointers are represented symbolically as $ptr(l, n)$ with a symbolic base address l and symbolic offset n , allowing us to track pointer arithmetic without knowing concrete addresses. The symbolic address l_c is used for constant memory addresses. Numbers are represented as $num(n)$ with a symbolic value. The types $i32(n)$, $ptr(n)$, and $num(n)$ form a lattice, the subtyping relation and meet and join operations for which are discussed in Section 3.4. We only refine $i32$ because the other WebAssembly value types ($i64$, $f32$, $f64$) are not used to index memory.

271

272

273

274

275

276

277

278

279

(3) *Symbolic Memory Typing*, Σ : Unlike the original WebAssembly type system, we type the flat contiguous array of bytes that serves as the WebAssembly linear memory. Since memory operations are ubiquitous in WebAssembly, we need to be able to recover the (refinement) types of values that are stored in memory at a specific address. A memory typing Σ is a mapping from symbolic location, offset pairs (l, n) to symbolic refinement types $\hat{\tau}_r$.

280

281

282

283

284

285

286

287

288

289

290

291

292

293

294

(4) *Refinement Function Types*, $tf_{\hat{\tau}_r}$: Function types tf_w in WebAssembly are used to encode the types of functions, blocks, loops, and calls. They specify the (input and output) base WebAssembly types τ_w for these constructs. In our refinement type system, a refinement function type $tf_{\hat{\tau}_r}$ specifies not only the lists of (input and output) refinement types τ_r , but also the shape of the (input and output) memory typing Σ . We must track how the memory typing changes over the course of execution since WebAssembly programs perform strong updates that can change the type of data stored in memory. For instance, a function that expects a $i32$ at location $(l, 0)$, because it reads from $(l, 0)$ using a $i32.load$ instruction, might later store an $f32$ at the same location.

- (5) *Label Typing*: In WebAssembly, labels of control constructs contain the shape of the stack expected when there is a jump to the control construct. In our refinement type system, we must specify not only the expected refinement types $\hat{\tau}_r$ on the stack but also the expected memory typing Σ for the control construct in the label.

3.3 Constraint Language to Discover Refinement Types

In this section, we explain how we discover symbolic refinement types for WebAssembly code. WebAssembly’s type validation algorithm, described in the specification [47], performs a single forward pass over the instructions in a function body to validate that WebAssembly function. Unfortunately, a forward-only analysis isn’t possible when we wish to discover refinement types. Instead, we first generate constraints over a WebAssembly function in order to later discover refinement types for each function, instruction, and variable via constraint solving. Accumulating constraints is necessary because it is often not immediately obvious if an i32 value is a pointer or a number. For example, in function `arr_add` (Figure 1), the parameters’ types are not immediately clear. Only at line `W18`, where the first parameter accesses memory, do we discover it is a pointer. We must then propagate this type backward to the function signature and other uses. Thus, a forward-only analysis is insufficient—we need types to propagate backwards too.

We generate constraints over WebAssembly functions and solve them in a fixed-point algorithm to get *symbolic* refinement types $\hat{\tau}_r$ for every stack slot in the value stack, local and global variables, and *symbolic* memory typing Σ . Our call graph analysis is over the symbolic refinement type system. We will later use the symbolic refinement types we discover for WebAssembly code e^* to guide the construction and typing of a concrete-refinement-typed WebAssembly program e^*_{τ} for which we prove type safety in Section 4.3.

To generate constraints, we first associate every stack slot in the value stack and local and global variables, with a unique abstract refinement type variable α and every memory state with a unique abstract memory variable ζ . We generate constraints over `WAFFLE` [13], an SSA IR over WebAssembly, since SSA form allows us to type different updates to a variable with distinct refinement types. For each instruction, we generate *type constraints* over α s and *memory constraints* over ζ s, the syntax of which can be seen in Figure 3. A subset of the constraint-generation rules for WebAssembly instructions are described in Figure 4, with all the rules in Appendix D. We discuss the interesting constraint-generation cases below.

3.3.1 Arithmetic Instructions. Let us consider the example of `i32.add`, which expects two values on the stack, of types α_1 and α_2 and produces a single value of type α_3 , i.e., $\alpha_3 = \alpha_1 + \alpha_2$. We know that the only disallowed case for this operation is the addition of two pointers. This has interesting implications for constraint solving, since if $\alpha_1 = \text{ptr}(l, n_1)$ and $\alpha_2 = \text{i32}(n_2)$, α_2 can be refined to be $\text{num}(n_2)$, since the alternative is not permitted. In fact, the types of each argument and result can affect each other. If $\alpha_1 = \text{num}(n_1)$, $\alpha_2 = \text{i32}(n_2)$ and $\alpha_3 = \text{ptr}(l, n_3)$, α_2 is refined to be ptr . Meanwhile, if $\alpha_3 = \text{num}(n_3)$, α_2 is refined to num . We leave discovering the appropriate refinements to constraint solving and at the time of constraint generation, generate constraints to tie each α to each other α , with the relevant arithmetic instruction. For `i32.add`, we generate the constraints, $\alpha_1 \doteq \text{i32.add}(_, \alpha_2, \alpha_3)$, $\alpha_2 \doteq \text{i32.add}(\alpha_1, _, \alpha_3)$ and $\alpha_3 \doteq \text{i32.add}(\alpha_1, \alpha_2, _)$. The underscore in the constraint holds the place of the α current being refined. A similar strategy is followed for all arithmetic instructions.

3.3.2 Memory Instructions. The `i64.store` instruction expects a pointer type α_{ptr} on the stack and some data type α_{data} . α_{ptr} is constrained to be a pointer by generating a subtyping constraint $\alpha_{\text{ptr}} <: \text{ptr}$, while the data is constrained to be a subtype of the type expected by the store instruction, $\alpha_{\text{data}} <: \text{i64}$. The abstract memory typing represented by ζ , an input to the rule, is updated to record

$C^\alpha ::= \{\text{func } tf_w^*, \text{ local } \alpha^*, \text{ global } tg_\alpha^*, \text{ table } n^?, \text{ memory } n^?, \text{ label } (\alpha^*, \zeta)^*, \text{ return } (\alpha^*, \zeta)^?\}$

Constraint Generation for Instructions

$$S; \zeta; C^\alpha \vdash e : \alpha^* \rightarrow \alpha^*; S'; \zeta'$$

$$\frac{\alpha \text{ fresh} \quad S' = S :: [\alpha \doteq \tau_w(c)]}{S; \zeta; C^\alpha \vdash \tau_w.\text{const } c : \epsilon \rightarrow \alpha; S'; \zeta} \text{CONSTANT}$$

$$\frac{\alpha_1 \alpha_2 \in \text{dom}(S) \quad \alpha_3 \text{ fresh} \quad S' = S :: [\alpha_1 \doteq \text{binop}(_, \alpha_2, \alpha_3) \wedge \alpha_2 \doteq \text{binop}(\alpha_1, _, \alpha_3) \wedge \alpha_3 \doteq \text{binop}(\alpha_1, \alpha_2, _)]}{S; \zeta; C^\alpha \vdash \tau_w.\text{binop} : \alpha_1 \alpha_2 \rightarrow \alpha_3; S'; \zeta} \text{BINARY OPS}$$

$$\frac{tf_w = \tau_w^n \rightarrow \tau_w^m \quad \alpha^n \in \text{dom}(S) \quad \alpha^m, \zeta' \text{ fresh} \quad S; \zeta; C, \text{label}(\alpha^m, \zeta') \vdash e^* : \alpha^n \rightarrow (\alpha')^m; S'; \zeta'' \quad S'' = S' :: [(\alpha \doteq \alpha')^m \wedge \zeta' \doteq \zeta'']}{S; \zeta; C^\alpha \vdash \text{block } tf_w e^* \text{ end} : \alpha^n \rightarrow \alpha^m; S''; \zeta'} \text{BLOCK} \quad \frac{\alpha^* \alpha^n \in \text{dom}(S) \quad C_{\text{label}}^\alpha(i) = \alpha'_n, \zeta' \quad S' = S :: [\zeta \doteq \sqcup \zeta \zeta' \wedge (\alpha \doteq \sqcup \alpha \alpha')^n]}{S; \zeta; C^\alpha \vdash \text{br } i : \alpha^* \alpha^n \rightarrow \alpha^*; S'; \zeta} \text{BREAK}$$

$$\frac{C_{\text{func}}^\alpha = \tau_w^m \rightarrow \tau_w^n \quad \alpha^n \in \text{dom}(S) \quad S' = S :: [(\alpha <: \tau_w)^m \wedge (\alpha' <: \tau_w)^n]}{S; \zeta; C^\alpha \vdash \text{call} : \alpha^m \rightarrow (\alpha')^n; S'; \zeta} \text{CALL} \quad \frac{\alpha_1 \in \text{dom}(S) \quad S' = S :: [\alpha_1 <: \text{ptr} \wedge \alpha_2 \doteq \zeta[\alpha_1](tp_sx^?, a, o) \wedge \alpha_2 <: \tau_w]}{S; \zeta; C^\alpha \vdash \tau_w.\text{load}(tp_sx)^? a o : \alpha_1 \rightarrow \alpha_2; S'; \zeta} \text{LOAD}$$

Fig. 4. Constraint-Generation Rules for a Subset of WebAssembly Instructions.

a mapping from $\alpha_{\text{ptr}} \mapsto \alpha_{\text{data}}$, and is copied into a new ζ' , which is returned by the rule. The constraint on the abstract memory typing is $\zeta' = \zeta[\alpha_{\text{ptr}} \mapsto \alpha_{\text{data}}]$. The `i64.load` instruction is similarly constrained to expect a pointer on the stack and a subtype of `i64` as the result. Here, the load constraint is represented as $\alpha_{\text{data}} \doteq \zeta[\alpha_{\text{ptr}}]$. Similar constraints are generated for all $\tau_w.\text{load}$ and $\tau_w.\text{store}$ instructions.

3.3.3 Blocks, Loops. Let us consider the sequence of instructions, `block` $tf_w e^* \text{end}$. Let the block expect n values on the stack and return m values. m α_1 's are freshly generated and added to the context with a label. This is done so that break instructions out of blocks know the number of values required on the output stack of the target block. We generate constraints for the body of the block, with this new context, which gives us m α_2 s. The m α_1 and α_2 s are now equated with equality constraints, and m α_1 s are returned by the instruction on the result stack. A similar scenario arises for loop instructions. For loops, the n input α s are added to the context instead of the m result α s that are added to the context for blocks. This is because break instructions in loops restart the loop and have to know the number of values required on the input stack of the loop.

3.3.4 Break Instructions. `br` instructions are used to break out of blocks and restart loops. In the case of blocks, the values on the stack at the `br` instruction are returned by the target block. Since our analysis is flow-insensitive, we join all possible result stacks for a block. The context carries the expected value stack at the target block, so we generate join constraints for the α s on the stack and the α s expected by the target block. In the case of loops, the values on the stack at the `br` instruction are used to restart the loop. Constraint generation is the same regardless of whether the target of a `br` instruction is a loop or a block. Note that generating join constraints in this fashion embeds the looping structure of the program into the constraints. For example, consider the following sequence of instructions: `loop` $tf_w e^* \text{end}$. If $tf_w = \text{i32} \rightarrow ()$, we would generate constraints for the loop inputs as $\alpha_0 <: \text{i32}$. We would then generate constraints for the body of the loop with $\text{label}(\alpha_0)$

added to the context. If we came across the instruction `br 0`, in the loop body, with α_1 on the stack, we would generate the constraint $\alpha_0 \doteq \sqcup \alpha_0 \alpha_1$.

3.4 Constraint Solving

We solve a set of constraints generated over a WebAssembly function using a standard fixpoint algorithm, described in detail in Figure 1. At the end of constraint generation, we have a set S of constraints over a function body. We create a mapping *Constraint*, that maps each α in S to the type constraints on α and similarly, maps each ζ in S to the memory constraints on ζ . We resolve the equality constraints for α 's and ζ 's when creating this map. We also record dependencies between α 's and ζ 's using the influence vector *infl*. This is used to recompute the constraints of a certain α or ζ when its dependencies have changed. We pass the *Constraint* and *infl* maps as inputs to constraint solving and compute a maximum fixed point solution for every α and ζ in S using a standard worklist algorithm. The constraints on α 's are solved to get a symbolic refinement type, $\hat{\tau}_r$, and the constraints on ζ 's are solved to get a memory state, Σ . If there exists a $\alpha \doteq \hat{\tau}_r$ in the constraints for α , we initialize the solution for the α to be $\hat{\tau}_r$. Otherwise, we initialize the solution to be the \top type. We initialize the solution of all ζ 's to be an empty memory state. We initialize the worklist with all α 's and ζ 's in the domain of the *Constraint* map, and continue solving till the worklist is empty. If the solution of an α or ζ is found to be different from its solution in the previous iteration, we add all α 's and ζ 's that depend on it into the worklist. We discuss solving type and memory constraints below.

3.4.1 Solving Type Constraints. To solve the type constraints for a given α , we iterate over each constraint in *Constraint*[\(\alpha\)] and progressively narrow the current solution by taking its meet with the type computed from each constraint.

- (1) For a *subtype constraint* $\alpha <: \tau$, the solution is met with $\tau(\top)$, restricting α to the correct base type while keeping its value refinement as precise as possible.
- (2) For an *operation constraint* $\alpha \doteq op(\alpha'_1, \dots, \alpha'_n)$, we evaluate *op* on the current solutions of the α'_i and meet the result into α 's solution.
- (3) For a *join constraint* (generated over blocks and loops) $\alpha \doteq \sqcup \alpha'_1 \dots \alpha'_m$, we join the current solutions of all α'_i and meet that join into α 's solution, capturing merge points in control flow.
- (4) The *memory load constraint* $\alpha \doteq \zeta[\alpha']$ is the most complex case due to potential aliasing. We resolve the current memory state $\Sigma = \rho[\zeta]$ and pointer type $\hat{\tau}_r^{\text{ptr}} = \rho[\alpha']$. If $\hat{\tau}_r^{\text{ptr}} = \text{ptr}(l, n)$ and $(l, n) \mapsto \tau_r \in \Sigma$, we read τ_r as the primary result. We then conservatively identify all potential aliases in Σ : (1) $\text{ptr}(\top, \top)$, a fully-unknown pointer that could alias anything; (2) $\text{ptr}(l, \top)$, a pointer to the same allocation but an unknown offset; and (3) any other entry $\text{ptr}(l', n')$, which may alias $\text{ptr}(l, n)$ since the relationship between allocation labels cannot be determined statically. The final solution for α is the meet of its current value with the join of all these contributions, ensuring that any type stored at a possible alias location is accounted for.

3.4.2 Solving Memory Constraints. Solving the memory constraints for a given ζ iterates over each constraint in *Constraint*[\(\zeta\)], progressively updating the current memory state $\Sigma = \rho[\zeta]$ via meet operations.

- (1) For a *memory equality constraint* $\zeta \doteq \Sigma'$, we meet the given concrete state Σ' into Σ .
- (2) For a *join constraint* $\zeta \doteq \sqcup \zeta'_1 \dots \zeta'_m$, we join the current solutions of all ζ'_i and meet the result into Σ , capturing merge points in control flow.

- (3) The *memory store constraint* $\zeta \doteq \zeta'[\alpha_1 \mapsto \alpha_2]$ requires the most care. We resolve the base state $\Sigma' = \rho[\zeta']$, the pointer $\hat{\tau}_r^{\text{ptr}} = \rho[\alpha_1]$, and the data type $\hat{\tau}_r^{\text{data}} = \rho[\alpha_2]$. If the pointer is a concrete $\text{ptr}(l, n)$, we perform a **strong update**, meeting Σ with Σ' modified to map (l, n) to $\hat{\tau}_r^{\text{data}}$; otherwise we meet Σ with the unmodified Σ' . Note that a concrete pointer might be a constant memory address with a statically known offset or a symbolic address with a statically known offset. We then handle aliasing conservatively: if Σ already contains an entry for $\text{ptr}(\tau, \tau)$, $\text{ptr}(l, \tau)$, or any other $\text{ptr}(l', n')$ whose relationship to $\text{ptr}(l, n)$ cannot be determined statically, we join $\hat{\tau}_r^{\text{data}}$ into the type at that alias entry, reflecting the possibility that the store may have updated it. This ensures that a store through an ambiguous pointer is reflected at all locations it might affect.

The precision of our memory analysis depends on how much is statically known about the pointer used in a load or store. When a function takes a pointer as a parameter and accesses memory exclusively through statically known immediate offsets (encoded directly in the instruction, eg, $\text{i32.load offset} = 8$), the pointer is represented as $\text{ptr}(l, \text{i32}(8))$, i.e., a pointer with a symbolic location l and concrete offset $\text{i32}(8)$, enabling precise, strong-update semantics for every load and store in the function body. Precision degrades when a dynamic offset is used, or when a function accesses memory through two or more pointers whose relationship cannot be determined statically: since no information is available to distinguish their allocation labels, our analysis must conservatively assume they may alias. Precisely characterising aliasing among dynamically computed addresses is undecidable in general [26], so our analysis must overapproximate when aliases cannot be resolved statically. In such cases, the result type of a load, or the updated memory state after a store, is widened to account for all potential aliases.

3.4.3 Refinement Type Lattice and Soundness. The precision and soundness of the analysis depends on the Refinement Type Lattice, which we now explain. The WebAssembly base types that have not been refined beyond the addition of an abstract value (i64 , f32 , f64) have sub-lattices defined as follows: $\mathcal{L}_{\tau_w} = \tau_w(\mathbb{N}_{|\tau_w|}) \cup \{\tau_w(\top), \tau_w(\perp)\}$. The i32 sub-lattice is more complex, containing ptr and num as shown in Figure 5 with a portion of its meet and join operations, which we discuss in detail shortly. The entire type lattice is defined as $\mathcal{L} = \{\top, \perp\} \cup \mathcal{L}_{\text{i32}} \cup \mathcal{L}_{\text{i64}} \cup \mathcal{L}_{\text{f32}} \cup \mathcal{L}_{\text{f64}}$. The lattice has finite height, ensuring that constraint solving terminates. Our constraint generation rules guarantee that every α has a valid base WebAssembly type. Since constraint solving computes the greatest fixed point, the analysis is sound: in the worst case, conflicting constraints cause the refinement to collapse to an unrefined base type with no value refinement. For example, if α is constrained to be both ptr and num , their meet loses the refinement but retains the base i32 type (Figure 5, last rule). We discuss this case and other interesting meet and join cases of the i32 sub-lattice below.

Join of two pointers. The canonical example for joins in static analysis is the if-then-else expression. Let us suppose that the then branch of such an expression returns $\text{ptr}(l_1, n_1)$, and the else branch returns $\text{ptr}(l_2, n_2)$. What pointer does the if-then-else expression return? Since it could be either of the two, we do not presume to know what the pointer could be and instead return $\text{ptr}(\top, \top)$. Instead, if the then branch returned $\text{ptr}(l, n_1)$ and the else branch returned $\text{ptr}(l, n_2)$, we say that the if-then-else expression returns $\text{ptr}(l, n_1 \sqcap n_2)$.

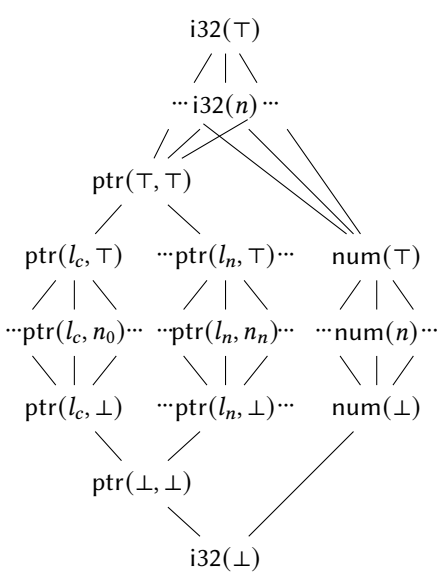
Meet of two pointers. The canonical example for meet operations in static analysis is values during several loop iterations. Say, at loop iteration n , the type of a stack slot is $\text{ptr}(l_1, n_1)$, a pointer with a symbolic address, and at loop iteration $n + 1$, its type is $\text{ptr}(l_c, n_0)$, a constant pointer. Since it is impossible for a stack slot to have a symbolic and constant pointer, we return $\text{ptr}(\perp, \perp)$. On the other hand, say that at loop iteration n , the type of a stack slot is $\text{ptr}(l_1, n_1)$ and at loop iteration $n + 1$,

Algorithm 1 Algorithm for solving constraints described in Figure 3

Require: $Constraint = [\alpha \mapsto t^+; \zeta \mapsto m^+]$, $infl = [\alpha \mapsto \{\alpha^*, \zeta^*\}; \zeta \mapsto \{\alpha^*, \zeta^*\}]$
Ensure: The least Solution $\rho = [\alpha \mapsto \hat{\tau}_\tau; \zeta \mapsto \Sigma]$
 $Worklist \leftarrow dom(Constraint)$
for all $\alpha \in dom(Constraint)$ **do**
if $\exists \alpha \doteq \hat{\tau}_\tau \in Constraint[\alpha]$ **then** $\rho[\alpha] \doteq \hat{\tau}_\tau$ **else** $\rho[\alpha] = \top$
for all $\zeta \in dom(Constraint)$ **do** $\rho[\zeta] = \cdot$
while $Worklist \neq \emptyset$ **do**
 $v := Worklist.pop()$
 $Old_v := \rho[v]$
if $v = \alpha$ **then** $EVAL_\alpha(Constraint, \rho, \alpha)$
if $v = \zeta$ **then** $EVAL_\zeta(Constraint, \rho, \zeta)$
if $Old_v \neq \rho[v]$ **then** $Worklist.append(infl[v])$
procedure $EVAL_\alpha(Constraint, \rho, \alpha)$
 $\hat{\tau}_\tau^\alpha := \rho[\alpha]$
for all $t \in Constraint[\alpha]$ **do**
if $t = \alpha <: \tau$ **then** $\hat{\tau}_\tau^\alpha := \hat{\tau}_\tau^\alpha \sqcap \tau(\top)$
if $t = \alpha \doteq op(\alpha'_1, \dots, \alpha'_n)$ **then** $\hat{\tau}_\tau^\alpha := \hat{\tau}_\tau^\alpha \sqcap eval_{op}(\rho[\alpha'_1], \dots, \rho[\alpha'_n])$
if $t = \alpha \doteq \sqcup \alpha'_1 \dots \alpha'_m$ **then** $\hat{\tau}_\tau^\alpha := \hat{\tau}_\tau^\alpha \sqcap (\rho[\alpha'_1] \sqcup \dots \sqcup \rho[\alpha'_m])$
if $t = \alpha \doteq \zeta[\alpha']$ **then** $\Sigma := \rho[\zeta] \wedge \hat{\tau}_\tau^{ptr} \doteq \rho[\alpha']$
if $\hat{\tau}_\tau^{ptr} = ptr(l, n) \wedge (l, n) \mapsto \hat{\tau}_\tau \in \Sigma$ **then** $\hat{\tau}_\tau^{(l,n)} := \hat{\tau}_\tau$ **else** $\hat{\tau}_\tau^{(l,n)} := \top$
if $\exists ptr(\top, \top) \mapsto \hat{\tau}_\tau \in \Sigma$ **then** $\hat{\tau}_\tau^{(\top, \top)} := \hat{\tau}_\tau$ **else** $\hat{\tau}_\tau^{(\top, \top)} := \perp$
if $\exists ptr(l, \top) \mapsto \hat{\tau}_\tau \in \Sigma$ **then** $\hat{\tau}_\tau^{(l, \top)} := \hat{\tau}_\tau$ **else** $\hat{\tau}_\tau^{(l, \top)} := \perp$
if $\exists ptr(l', n') \mapsto \hat{\tau}_\tau \in \Sigma$ **then** $\hat{\tau}_\tau^{(l', n')} := \hat{\tau}_\tau$ **else** $\hat{\tau}_\tau^{(l', n')} := \perp$
 $\hat{\tau}_\tau^\alpha := \hat{\tau}_\tau^\alpha \sqcap \hat{\tau}_\tau^{(l,n)} \sqcup \hat{\tau}_\tau^{(\top, \top)} \sqcup \hat{\tau}_\tau^{(l, \top)} \sqcup \hat{\tau}_\tau^{(l', n')}$
 $\rho[\alpha] \doteq \hat{\tau}_\tau^\alpha$
procedure $EVAL_\zeta(Constraint, \rho, \zeta)$
 $\Sigma \doteq \rho[\zeta]$
for all $m \in Constraint[\zeta]$ **do**
if $m = \zeta \doteq \Sigma'$ **then** $\Sigma := \Sigma' \sqcap \Sigma$
if $m = \zeta \doteq \sqcup \zeta'_1 \dots \zeta'_m$ **then** $\Sigma := (\rho[\zeta'_1] \sqcup \dots \sqcup \rho[\zeta'_m]) \sqcap \Sigma$
if $m = \zeta \doteq \zeta'[\alpha_1 \mapsto \alpha_2]$ **then** $\Sigma' := \rho[\zeta'] \wedge \hat{\tau}_\tau^{ptr} := \rho[\alpha_1] \wedge \hat{\tau}_\tau^{data} := \rho[\alpha_2]$
if $\hat{\tau}_\tau^{ptr} = ptr(l, n)$ **then** $\Sigma := \Sigma \sqcap \Sigma'[(l, n) \mapsto \hat{\tau}_\tau^{data}]$ **else** $\Sigma := \Sigma \sqcap \Sigma'$
if $\exists ptr(\top, \top) \mapsto \hat{\tau}'_\tau \in \Sigma$ **then** $\Sigma := \Sigma[(\top, \top) \mapsto \hat{\tau}'_\tau \sqcup \hat{\tau}_\tau^{data}]$
if $\exists ptr(l, \top) \mapsto \hat{\tau}'_\tau \in \Sigma$ **then** $\Sigma := \Sigma[(l, \top) \mapsto \hat{\tau}'_\tau \sqcup \hat{\tau}_\tau^{data}]$
if $\exists ptr(l', n') \mapsto \hat{\tau}'_\tau \in \Sigma$ **then** $\Sigma := \Sigma[(l', n') \mapsto \hat{\tau}'_\tau \sqcup \hat{\tau}_\tau^{data}]$
 $\rho[\zeta] \doteq \Sigma$

its type is $ptr(l_2, n_2)$. This means that the type at this stack slot is both $ptr(l_1, n_1)$ and $ptr(l_2, n_2)$. We equate these two pointers to a third pointer $ptr(l_3, 0)$, where l_3 is a fresh symbolic location and $ptr(l_1, n_1) = ptr(l_3, 0) \wedge ptr(l_2, n_2) = ptr(l_3, 0)$. Pointers can be written as polynomials (since the offset is usually added into the base address) and we get that $l_1 + n_1 = l_3$ or that, $l_1 = l_3 - n_1$ and $l_2 = l_3 - n_2$. We say that the type of this stack slot is $ptr(l_3, 0)$.

Join and Meet of a pointer and a number. The join of a pointer with a number results in an $i32(\top)$, except in the case of constant pointers $ptr(l_c, n_0)$. For a constant pointer, we know that the base address l_c equates to 0 and so the operation $ptr(l_c, n_0) \sqcup num(n_1)$, produces $i32(n_0 \sqcup n_1)$ as



$i32(n_1)$	\sqcup	$i32(n_1)$	$=$	$i32(n_1 \sqcup n_2)$
$i32(n_1)$	\sqcup	$\text{ptr}(l_c, n_0)$	$=$	$i32(n_1 \sqcup n_2)$
$i32(n_1)$	\sqcup	$\text{ptr}(l_2, n_2)$	$=$	$i32(\top)$
$i32(n_1)$	\sqcup	$\text{num}(n_2)$	$=$	$i32(n_1 \sqcup n_2)$
$\text{ptr}(l_1, n_1)$	\sqcup	$\text{ptr}(l_1, n_2)$	$=$	$\text{ptr}(l_1, n_1 \sqcup n_2)$
$\text{ptr}(l_1, n_1)$	\sqcup	$\text{ptr}(l_2, n_2)$	$=$	$\text{ptr}(\top, \top)$
$\text{num}(n_1)$	\sqcup	$\text{num}(n_2)$	$=$	$\text{num}(n_1 \sqcup n_2)$
$\text{num}(n_1)$	\sqcup	$\text{ptr}(l_c, n_0)$	$=$	$i32(n_1 \sqcup n_0)$
$\text{num}(n_1)$	\sqcup	$\text{ptr}(l_2, n_2)$	$=$	$i32(\top)$
$i32(n_1)$	\sqcap	$i32(n_1)$	$=$	$i32(n_1 \sqcap n_2)$
$i32(n_1)$	\sqcap	$\text{ptr}(l_c, n_0)$	$=$	$\text{ptr}(l_c, n_1 \sqcap n_0)$
$i32(n_1)$	\sqcap	$\text{ptr}(l_2, n_2)$	$=$	$\text{ptr}(l_c, n_1) \sqcap \text{ptr}(l_2, n_2)$
$i32(n_1)$	\sqcap	$\text{num}(n_2)$	$=$	$\text{num}(n_1) \sqcup \text{num}(n_2)$
$\text{ptr}(l_1, n_1)$	\sqcap	$\text{ptr}(l_c, n_0)$	$=$	$\text{ptr}(\perp, \perp)$
$\text{ptr}(l_1, n_1)$	\sqcap	$\text{ptr}(l_1, n_2)$	$=$	$\text{ptr}(l_1, n_1 \sqcap n_2)$
$\text{ptr}(l_1, n_1)$	\sqcap	$\text{ptr}(l_2, n_2)$	$=$	$\text{ptr}(l_3, 0)$
		where l_3 fresh $\wedge l_1 = l_3 - n_1 \wedge l_2 = l_3 - n_2$		
$\text{num}(n_1)$	\sqcap	$\text{num}(n_2)$	$=$	$\text{num}(n_1 \sqcap n_2)$
$\text{num}(n_1)$	\sqcap	$\text{ptr}(l_c, n_0)$	$=$	$i32(\perp)$
$\text{num}(n_1)$	\sqcap	$\text{ptr}(l_1, n_1)$	$=$	$i32(\perp)$

Fig. 5. The $i32$ Sub-Lattice With a Subset of Meet and Join Operations Defined.

its result. The meet of a pointer and number unequivocally results in a $i32(\perp)$, as per our relation definition shown in Figure 5.

Join and meet of a $i32$ with a pointer. If the then branch of a if-then-else expression returned $i32(n_1)$ and the else branch returned a constant pointer $\text{ptr}(l_c, n_0)$, the if-then-else expression would return $i32(n_0 \sqcup n_1)$. If the else branch returned a pointer with a symbolic address $\text{ptr}(l_2, n_2)$ instead, the if-then-else expression would return $i32(\top)$, since the base address l_1 is unknown. For the meet operation between $i32$'s and ptr 's, let us imagine that the type of a stack slot at loop iteration n is $i32(n_1)$ and that at loop iteration $n + 1$, the type of the stack slot is a constant pointer $\text{ptr}(l_c, n_0)$. The type of the stack slot is then both these types and so, $\text{ptr}(l_c, n_0 \sqcap n_1)$. If instead, at loop iteration $n + 1$, the type of the stack slot is a symbolic pointer $\text{ptr}(l_2, n_2)$, the type of the stack slot would be $\text{ptr}(l_c, n_1) \sqcap \text{ptr}(l_2, n_2)$, which results in $\text{ptr}(\perp, \perp)$.

Join and meet of a $i32$ with a number. If the then branch of a if-then-else expression returned a $i32(n_0)$ and the else branch returned $\text{num}(n_1)$, the if-then-else expression would return $i32(n_0 \sqcup n_1)$. For the meet operation between $i32$'s and num 's, let us imagine that the type of a stack slot at loop iteration n is $i32(n_0)$ and that at loop iteration $n + 1$, the type of the stack slot is a constant pointer $\text{num}(n_1)$. The type of the stack slot is then both these types and so, $\text{num}(n_0 \sqcap n_1)$.

Note that we never explicitly join or meet locations separate from their offsets. When the meet operation is performed over two pointers with the same symbolic base address l , the result is a pointer with that base address l and the join of their offsets. When the meet operation is performed on two pointers with different symbolic locations, we produce a pointer with a \top location and \top offset. This is analogous to $\text{ptr}(\top)$, since we never construct $\text{ptr}(\top, n)$. Similarly, for join operations, we never construct $\text{ptr}(\perp, n)$ and $\text{ptr}(\perp, \perp)$ is analogous to $\text{ptr}(\perp)$. However, in both cases, we do construct $\text{ptr}(l, \top)$ and $\text{ptr}(l, \perp)$.

$R ::= \{\text{inst } C^*, \text{tab } n^*, \text{mem } \Psi^*\}$
 $C^r ::= \{\text{func } tf_r^*, \text{local } \tau_r^*, \text{global } \tau_r^*, \text{table } n^?, \text{memory } n^?, \text{label } (\tau_r^*, \Psi)^*, \text{return } (\tau_r^*, \Psi)^?\}$

Typing WebAssembly Instructions with Refinement Types

$$R; C^r \vdash e_r^* : tf_r$$

$\tau_{r32} ::= \text{i32}(n) \mid \text{ptr}(n) \mid \text{num}(n) \mid \text{f32}(n)$
 $\tau_r ::= \tau_{r32} \mid \text{i64}(n) \mid \text{f64}(n)$
 $\tau ::= \text{i32} \mid \text{ptr} \mid \text{num} \mid \text{i64} \mid \text{f32} \mid \text{f64}$
 $\Psi ::= \{(\text{ptr}(n) \mapsto \tau_{r32})^*\}$
 $tf_r ::= \tau_r^*, \Psi \rightarrow \tau_r^*, \Psi'$

$$\frac{}{R; C^r \vdash \tau.\text{const } c : \tau(c)} \text{CONSTANT} \qquad \frac{\tau_r^3 = \text{binop}(\tau_r^1, \tau_r^2) \quad \tau_r^3 <: \tau(\top)}{R; C^r \vdash \tau.\text{binop} : \tau_r^1 \tau_r^2 \rightarrow \tau_r^3} \text{BINARY OPS}$$

$$\frac{R; C^r, \text{label}(\tau_r^n, \Psi_{\text{post}}) \vdash e^* : \tau_r^m, \Psi_{\text{pre}} \rightarrow (\tau_r')^n, \Psi'_{\text{post}} \quad R\Psi = \Psi_{\text{pre}} \quad (\tau_r' <: \tau_r)^n \quad \Psi'_{\text{post}} <: \Psi_{\text{post}}}{R; C^r \vdash \text{block } tf_r e^* \text{ end} : tf_r} \text{BLOCK} \qquad \frac{C^r_{\text{label}}(i) = (\tau_r')^m, \Psi' \quad R\Psi <: \Psi'}{R; C^r \vdash \text{br } i : \tau_r^* \tau_r^m \rightarrow \tau_r^*} \text{BREAK}$$

$$\frac{C^r_{\text{func}}(i) = (\tau_r')^m, \Psi'_{\text{pre}} \rightarrow (\tau_r')^n, \Psi'_{\text{post}} \quad tf_r = \tau_r^m, \Psi_{\text{pre}} \rightarrow \tau_r^n, \Psi_{\text{post}} \quad R\Psi = \Psi_{\text{pre}} \quad (\tau_r <: \tau_r')^m \quad (\tau_r' <: \tau_r)^n \quad \Psi_{\text{pre}} <: \Psi'_{\text{pre}} \quad \Psi'_{\text{post}} <: \Psi_{\text{post}}}{R; C^r \vdash \text{call } i \text{ } tf_r : \tau_r^m, \Psi_{\text{pre}} \rightarrow \tau_r^n} \text{CALL} \qquad \frac{C^r_{\text{memory}} = n \quad \tau_r = \text{load_and_extend}(c + o, \tau, (tp_sx)^?, R\Psi) \quad 2^a \leq (|tp| <)^? |\tau_r| \quad (tp_sz)^? = \epsilon \vee \tau_r = \text{im}}{\Psi; C^r \vdash \tau.\text{load}(tp_sx)^? a o : \text{ptr}(c) \rightarrow \tau_r} \text{LOAD}$$

Fig. 6. Typing of a subset of WebAssembly Instructions in the Concrete Refinement Type System.

4 CONCRETE REFINEMENT TYPES AND TYPE SAFETY

We have seen how to generate constraints over a given WebAssembly function to discover *symbolic* refinement types $\hat{\tau}_r$. To prove this refinement type system sound, we must derive *concrete* refinement types from the symbolic ones. Progress requires demonstrating that well-typed instructions can execute according to the operational semantics, which operates on concrete addresses and values. Unlike the original WebAssembly type system, which does not type memory and relies on type annotations to ensure bytes are interpreted correctly, our refinement type system explicitly types memory contents. This requires concrete types for the progress proof: symbolic types like $\text{ptr}(l, n)$ suffice for static analysis but cannot represent actual runtime values needed to show memory operations execute correctly.

4.1 Concrete Refinement Type System

To prove type safety, we transform the symbolic refinement types $\hat{\tau}_r$ discovered through constraint solving into concrete refinement types τ_r . We posit that there exists a mapping η from symbolic addresses to concrete addresses for every Σ , where, for a Σ_{pre} before and Σ_{post} after an instruction, $\eta_{\text{post}} \supseteq \eta_{\text{pre}}$. Hence, for every α and ζ , we recover a concrete refinement type τ_r and concrete memory typing Ψ , defined in Figure 6. Since Ψ is only typed with 32-bit types, we split each 64-bit sized type $\text{i64}(n)$ or $\text{f64}(n)$, into two i32 values from the high and low 32 bits of n . We now make several transformations to the WebAssembly function, in order to type check it:

- Instructions tagged with a WebAssembly type annotation τ_w are transformed to have a refined type tag τ . E.g. $\text{i32.const } 42 \rightsquigarrow \text{num.const } 42$, where the latter expects a num on the stack rather than a i32 .

- Instructions annotated with WebAssembly function type annotations tf_w are instead annotated with concrete refinement function types tf_r that include Ψ_{pre} and Ψ_{post} memory typing before and after the instruction.
- Function types are also transformed to have concrete refinement function types tf_r .
- The call instruction is typed to have a tf_r annotation (shown in Figure 6) to denote the shape of the stack before and after the call.

4.2 Typing Rules

Next, we describe the refinement typing rules for a subset of WebAssembly instructions (shown in Figure 6). Here, R is a refined store context and contains the WebAssembly module instances C^r , a shared table of functions and the memory typing Ψ . R is analogous to the store context S in the standard WebAssembly type system, except that S had a linear sequence of bytes as its memory. Refinement typing rules for the entire instruction set can be found in Appendix A.

4.2.1 Arithmetic Instructions. We restrict the refinement typing of arithmetic instructions to only allow certain combinations of $i32$, ptr and num . For instance, in the BINARY OPS rule in Figure 6, the result type is computed using the function $binop$ which takes two input refinement types and computes a result refinement type, but only considers certain pairs of input types valid for each specific binary instruction. The instruction will fail to type check if a invalid pair of input types is provided for a given binary operation. We discuss a few specific binary operations below.

- *Comparison:* Comparison of $i32$ s and all its subtypes result in a num . All combinations of $i32$, ptr and num can be compared.
- *Addition:* WebAssembly allows addition of two $i32$ s, either of which, or the result of which, can be used to index into the WebAssembly memory. We do not allow addition of two pointers, but allow addition of every other combination of $i32$, ptr and num . Moreover, addition of a ptr and a num yields a ptr .
- *Subtraction:* We do not allow subtraction of ptr from a num , but we do allow all other $i32$, ptr and num combinations for the subtraction operation. Subtracting a ptr from a ptr yields a num (an offset).

4.2.2 Control Constructs and Breaks. The BLOCK rule is annotated with a refined function type tf_r which specifies the stack shape and memory typing expected before the block instruction and at the end of the block instruction. Since the end of a block is the meet of several breaks out of the block, we ensure that the shape of the stack and memory after type checking the instructions in the block should be subtypes of the types expected on the top of the stack and of the expected memory typing. This is necessary since it is valid in WebAssembly for one control path to the end of the block to return a ptr and another to return a $i32$, since they would both be integers in the WebAssembly type system. br instructions are typed similarly.

4.2.3 Loads from Linear Memory. While in WebAssembly, $i32$ s are used to index into memory, our refinement type system requires that only values of ptr type be used to index into memory. We modify the typing rule for the LOAD instruction to expect a $ptr(n)$ on the stack, as shown in Figure 6. The load instruction is optionally annotated with a packed type and size, tp_sx , which is used to pack and sign extend the data stored in memory. Additionally, our memory typing Ψ only stores types of 32-bit size, τ_{i32} . If the load instruction expects a type of size 64 bits, we load the sequence of bytes from address $n + o$, where o is the offset provided to the load instruction, and address $n + o + 4$ as a 64-bit value as directed by the τ annotation on the load instruction which specifies the type of data we want to load. All this is done by the `load_and_extend` function, in Appendix A.

687 **4.2.4 Function Calls.** Typing a call instruction is more complex than in WebAssembly since our
688 refinement type system has subtyping and since it keeps track of the expected memory typing at
689 various points in the program. We change the WebAssembly `call i` instruction to `call i tf`, where
690 t_f specifies the stack shape and memory typing expected before and after the call. When typing
691 the call instruction, we need to ensure that the stack shape and memory typing before the call are
692 subtypes of the stack shape and memory typing expected by the function being called, and that the
693 stack shape and memory typing when the function returns are subtypes of the stack shape and
694 memory typing expected after the call by the callee.
695

696 4.3 Type Safety

697 The operational semantics of this transformed WebAssembly program e_r^* remains largely un-
698 changed from the original WebAssembly operational semantics, but there are a few changes. For
699 example, for the binary operation rule, the type annotation on the τ .*binop* instruction no longer tells
700 us the types expected on the WebAssembly stack. Instead of, $(\tau_w.\text{const } c_1)(\tau_w.\text{const } c_2) \tau_w.\text{binop} \hookrightarrow$
701 $\tau_w.\text{const}(\text{binop}(c_1, c_2))$, in the refinement type system, the operational semantics of *binop* steps
702 as follows, $(\hat{\tau}_r^1.\text{const } c_1) (\hat{\tau}_r^2.\text{const } c_2) \hat{\tau}_r^3.\text{binop} \hookrightarrow \text{rt_binop}((\hat{\tau}_r^1.\text{const } c_1), (\hat{\tau}_r^2.\text{const } c_2))$, where
703 rt_binop calculates the binary operation over refinement types.
704

704 We now prove type safety as the standard progress and preservation theorems.

705 **THEOREM 1. Progress:** *If $R; C^r \vdash_i e_r^* : \tau_r^1 \cdots \tau_r^m \rightarrow \tau_r^* \wedge (v : \tau_r)^m \wedge \vdash_i s : R$ then $e_r^* = v^*$ or*
706 *$e_r^* = \text{trap}$ or $s; v_r^*, e_r^* \hookrightarrow s'; v'^*, e_r'^*$*

707 **THEOREM 2. Preservation:** *If $R; C^r \vdash_i e_r^* : \tau_r^1 \cdots \tau_r^m \rightarrow \tau_r^* \wedge \vdash_i (v : \tau_r)^m \wedge \vdash_i s : R \wedge s, v^*, e_r^* \hookrightarrow$*
708 *$s', v'^*, e_r'^*$ then $\exists R'$ such that $\text{dom}(R_\Psi) \subseteq \text{dom}(R'_\Psi) \wedge \vdash_i s' : R' \wedge R'; C^r \vdash_i e_r'^* : \tau_r^*$*

709 Proofs of the above theorems are very similar to the WebAssembly progress and preservation
710 proofs except for minor changes to account for the refined types. The main interesting aspect of
711 our refinement type system is that, unlike WebAssembly, we also type the linear memory, keeping
712 track of the memory typing Ψ in the context R . We present the proof of progress and preservation
713 for the load instruction in Appendix B.
714
715

716 5 CALL GRAPH ANALYSIS OVER SYMBOLIC REFINEMENT TYPES

717 To determine valid targets for an indirect call, RNTA checks type compatibility between the inferred
718 callsite type and each candidate function's signature using standard function-type subtyping: the
719 arrow type is *contravariant* in parameters and *covariant* in return types. A function f is a valid
720 target if (i) each callsite argument type is a subtype of the corresponding parameter type of f ,
721 and (ii) the return type of f is a subtype of the return type expected at the callsite. When both
722 the callsite and a candidate have pointer-typed arguments, comparing symbolic pointer values
723 $\text{ptr}(l_1, n_1)$ and $\text{ptr}(l_2, n_2)$ directly is not meaningful, as their labels and offsets are placeholders for
724 statically unknown addresses. Instead, the check is lifted to the data types stored at those pointers:
725 given $\Sigma_1 = \{(l_1, n_1) \mapsto \hat{\tau}_r^1\}$ and $\Sigma_2 = \{(l_2, n_2) \mapsto \hat{\tau}_r^2\}$, we check $\hat{\tau}_r^1 <: \hat{\tau}_r^2$.

726 Figure 7 shows the refinement types discovered for the motivating example from Section 2.
727 `num_add` and `num_sub` remain unrefined at $(i32(\top), i32(\top)) \rightarrow i32(\top)$, as their bodies contain no
728 memory accesses. `arr_add` and `arr_sub` are refined to $(\text{ptr}(l_1, 0), i32(\top)) \rightarrow i32(\top)$ with entry
729 memory state $\{(l_1, 0) \mapsto i32(\top)\}$, since both load from memory through their first parameter.
730 The callsite in `indirect_call_num` has no memory-accessing parameters, giving caller stack
731 $(i32(\top), i32(\top)) \rightarrow i32(\top)$; the callsite in `indirect_call_arr` refines its first parameter to `ptr`,
732 giving $(\text{ptr}(l_1, 0), i32(\top)) \rightarrow i32(\top)$.
733

734 Under base WebAssembly type-filtering, both callsites resolve to all four functions. With refined
735 types, `indirect_call_num`'s $i32(\top)$ first argument cannot satisfy condition (i) for `arr_add` and

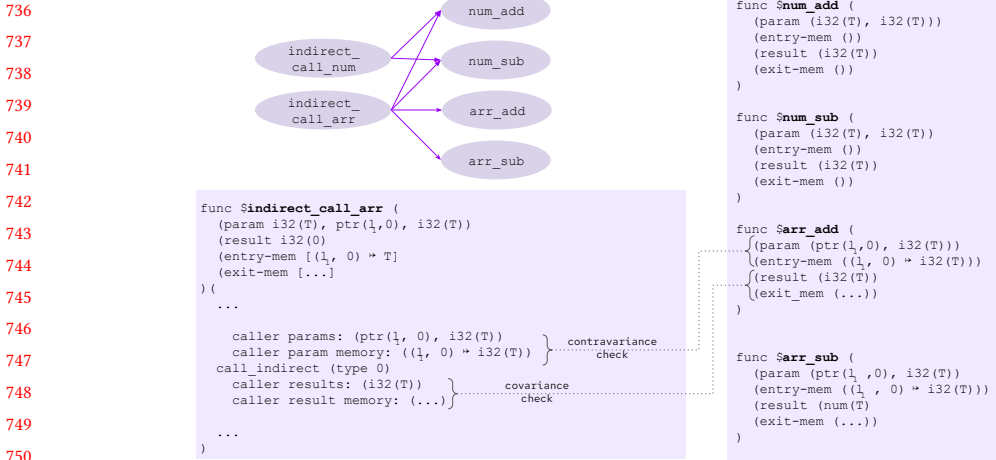


Fig. 7. Call graph over discovered refinement types for the motivating example.

`arr_sub`, since `i32` $\not\prec$ `ptr`, leaving only `num_add` and `num_sub` as valid targets. At `indirect_call_arr`, since `ptr` \prec `i32`, all four candidates remain valid. The asymmetry is fundamental: a `ptr` callsite is compatible with callees expecting the more general `i32`, but an `i32` callsite cannot satisfy a callee requiring the stricter `ptr`.

5.1 Applications to Other Analyses

A more precise call graph is a foundational improvement: virtually every interprocedural static analysis is parameterised by a call graph, and spurious edges propagate imprecision into every downstream client. Beyond this, the `ptr/num` distinction directly benefits analyses that reason about memory by reducing the aliasing overapproximation inherent in WebAssembly's flat `i32` type: a `num`-typed value cannot alias any `ptr`-accessed memory location. In the case where a function accesses memory through a single pointer with statically known offsets, each access resolves to a distinct cell `ptr(l, n)` with no aliasing ambiguity at all. For taint and dataflow analyses, stores through `ptr` need not affect `num`-typed values, and `num` arithmetic introduces no memory side effects to track. For program slicing, `num`-typed values can be excluded as candidate memory addresses, substantially reducing slice size.

6 EVALUATION ON REAL-WORLD WEBASSEMBLY BINARIES

We evaluate our call graph analysis on 20 real-world WebAssembly binaries from the NoWASET dataset [44], a dataset of NPM packages that depend on WebAssembly that has been compiled from diverse source language libraries including Rust, C/C++, Go, etc. NoWASET has a collection of WebAssembly binaries that is instantiated and executed by JavaScript clients, which they refer to by their module hash. We select 20 Wasm 1.0 binaries with a high number of indirect calls and following NoWASET's convention, refer to them by their hash, as shown in Table 1. Our evaluation compares a type-based call graph analysis over base WebAssembly types against the same analysis over refinement types inferred by RNTA. We assess how RNTA improves call graph precision, identifies a smaller set of reachable functions, discovers monomorphic callsites that are candidates for further optimizations, and how it compares against other state-of-the-art call graph analysis tools.

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

Binary	Size (KB)	Function Types	Total Functions	Exported Functions	Imported Functions	Direct Calls	Indirect Calls	Table Funcs
d128e52a5d	360.61	15	544	161	3	2521	5265	160
fd885c2d12	455.67	22	213	13	4	1529	515	50
0f9cbc41a9	666.18	30	143	5	19	1594	169	31
424d9e12e7	1366.61	58	1567	74	9	11939	162	178
04b0c8ef0f	123.52	36	235	49	6	922	116	62
fad0f2d451	57.16	18	167	7	2	463	108	7
b5047ed15e	331.51	31	442	41	8	2936	103	74
47fb4f81d2	86.58	52	248	6	29	761	98	200
f1fb556163	185.58	15	97	22	2	2027	78	6
d1d3df1b33	185.23	38	629	17	9	1177	65	46
e55ee035c7	60.45	23	172	5	0	701	65	51
517f31132d	207.99	15	99	5	18	335	60	35
704bd59a4e	164.80	39	276	58	20	1409	49	30
96e42b64f2	15.84	11	15	7	0	32	31	2
e6171e3491	78.04	18	111	8	0	834	29	23
7c74a56b87	85.72	18	106	8	0	667	26	20
058ca2a32b	65.68	16	99	7	0	606	25	20
1fc3c27d95	19.32	15	73	6	0	132	17	15
b2da8a7b24	17.34	11	66	3	0	289	17	17
477d02ba72	17.03	11	69	3	0	294	17	17

Table 1. Real-World WebAssembly Binary Statistics

RQ1: How much more precise is the call graph over the refinement type system when compared to the original WebAssembly type system and at what cost?

To evaluate the precision improvement from refinement types, we compare a type-based call graph analysis over the refinement type system against the same analysis over the original WebAssembly type system. We use a type-based analysis as it represents the standard approach employed by other WebAssembly analysis tools. A type-based analysis over base WebAssembly types restricts indirect call targets using only the type annotation at the call_indirect instruction and function signatures. By contrast, RNTA infers refinement function signatures and memory states, constructs a refined caller value stack and memory state, which we use to check contravariance for arguments and covariance for return types. Table 2 shows that refinement types yield significantly more precise call graphs.

Unique Function Types. Refinement types substantially increase the number of distinct function signatures. Across all 20 binaries, RNTA infers on average 3.6× more unique function types than the base WebAssembly type system. The largest increases occur in binaries whose functions make heavy use of i32 parameters that can be split into ptr and num: fd885c2d12 sees a 7.0× increase (from 22 to 153 types) and d128e52a5d a 6.9× increase (from 15 to 103). Conversely, binaries with few functions or limited use of i32 see smaller gains—96e42b64f2 increases only 1.3× (from 11 to 14) because it contains just 15 total functions and 2 table entries.

Edge Reduction. Edge reduction depends on how effectively refinement splits groups of functions in the WebAssembly table that share a base type. Binaries where refinement produces the greatest increase in type discrimination benefit most: e55ee035c7 achieves the highest reduction (43.6%) with a 3.2× increase in unique function types (23 base to 73 refined), while 47fb4f81d2 sees only 7.2% with a more modest 2.1× increase (43 to 91). The 4 binaries with no edge reduction either have very few table functions or have refinement types that only distinguish non-table functions. On average, the edge reduction is 14.6%. The primary source of precision improvement is distinguishing pointer-typed arguments from numeric ones: across all 20 binaries, 57.7% of eliminated indirect call edges result from a mismatch where the caller passes a numeric constant

	Binary	Type-Based				RNTA			
		Function Types	Reachable Nodes	Edges	Time(s)	Function Types	Reachable Nodes (%Red)	Edges (%Red)	Time(s)
834	d128e52a5d	15	544	17187	0.17	103	544 (0.0%)	14099 (18.0%)	3.37
835	fd885c2d12	22	213	1682	0.05	153	209 (1.9%)	1114 (33.8%)	10.99
836	0f9cbc41a9	30	39	325	0.10	82	39 (0.0%)	325 (0.0%)	21.60
837	424d9e12e7	57	1566	9474	0.12	370	1459 (6.8%)	6184 (34.7%)	9.16
838	04b0c8ef0f	34	235	882	0.02	130	234 (0.4%)	681 (22.8%)	1.13
839	fad0f2d451	18	167	344	0.02	107	148 (11.4%)	306 (11.0%)	0.20
840	b5047ed15e	31	439	1761	0.02	127	397 (9.6%)	1506 (14.5%)	0.31
841	47fb4f81d2	43	169	1402	0.05	91	152 (10.1%)	1301 (7.2%)	0.84
842	f1fb556163	15	91	317	0.02	52	90 (1.1%)	314 (0.9%)	0.51
843	d1d3df1b33	34	556	1207	0.02	75	519 (6.7%)	1073 (11.1%)	0.18
844	e55ee035c7	23	165	686	0.02	73	118 (28.5%)	387 (43.6%)	0.16
845	517f31132d	15	39	162	0.02	58	39 (0.0%)	162 (0.0%)	4.52
846	704bd59a4e	38	276	797	0.03	97	273 (1.1%)	778 (2.4%)	3.55
847	96e42b64f2	11	15	16	0.01	14	15 (0.0%)	16 (0.0%)	0.09
848	e6171e3491	18	105	311	0.01	52	81 (22.9%)	244 (21.5%)	1.09
849	7c74a56b87	18	100	254	0.01	49	80 (20.0%)	208 (18.1%)	0.35
850	058ca2a32b	16	93	243	0.01	48	73 (21.5%)	197 (18.9%)	0.21
851	1fc3c27d95	15	23	92	0.00	34	23 (0.0%)	92 (0.0%)	0.04
852	b2da8a7b24	11	60	128	0.01	32	43 (28.3%)	106 (17.2%)	0.04
853	477d02ba72	11	63	130	0.01	32	46 (27.0%)	108 (16.9%)	0.05
854	Average				0.04		9.9%	14.6%	2.92

Table 2. Comparing the reachable nodes and edges in a type-based call graph analysis over base WebAssembly types versus refinement types.

but the candidate target function expects a pointer argument, or vice versa. The remaining 41.3% of eliminations arise from cases where the source callsite receives a more specific refinement type that excludes certain targets.

Unreachable Function Detection. Removing spurious edges can render functions unreachable from exported function entry points. Of the 16 binaries with edge reduction, 15 also exhibit node reduction, averaging 9.9%. However, the correspondence between edge reduction and node reduction is not uniform. d128e52a5d achieves 18.0% edge reduction but no node reduction because this binary has 160 table functions sharing only 15 base type signatures, creating a dense web of indirect call edges. Although refinement types eliminate 3,088 spurious edges, every target of a removed edge retains other incoming edges, and 86 of the 88 affected target functions are also reached via direct calls. Thus, no function becomes unreachable despite the substantial edge reduction. By contrast, 47fb4f81d2 achieves 10.1% node reduction from only 7.2% edge reduction, indicating that the removed edges were the sole incoming paths to those functions.

Analysis Cost. The baseline type-based analysis averages 0.04s, while RNTA averages 2.92s (median 0.43s, maximum 21.60s). The overhead stems from constraint generation over every instruction in every function body and constraint solving, so analysis time scales with binary size and function body complexity rather than with the number of indirect calls. For instance, 0f9cbc41a9 takes 21.60s, while d128e52a5d has 5,265 indirect calls but completes in 3.37s. Even in the worst case, the analysis finishes in under 22 seconds, making it practical for real-world binaries.

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

Binary	Type-Based Polymorphic → RNTA Polymorphic	Type-Based Polymorphic → RNTA Monomorphic	Type-Based Polymorphic → RNTA Unreachable	Type-Based Monomorphic → RNTA Monomorphic	Type-Based Monomorphic → RNTA Unreachable	Both Unreachable
d128e52a5d	5267 → 4884	5267 → 0	5267 → 383	0 → 0	0 → 0	0 → 0
fd885c2d12	501 → 479	501 → 5	501 → 17	13 → 10	13 → 3	1 → 1
0f9cbc41a9	1 → 1	1 → 0	1 → 0	0 → 0	0 → 0	168 → 168
424d9e12e7	151 → 128	151 → 17	151 → 6	2 → 0	2 → 2	0 → 0
04b0c8ef0f	103 → 87	103 → 0	103 → 16	7 → 4	7 → 3	5 → 5
fad0f2d451	38 → 0	38 → 22	38 → 16	15 → 0	15 → 15	55 → 55
b5047ed15e	102 → 56	102 → 6	102 → 40	1 → 1	1 → 0	0 → 0
47fb4f81d2	77 → 75	77 → 1	77 → 1	4 → 3	4 → 1	17 → 17
f1fb556163	3 → 0	3 → 1	3 → 2	63 → 63	63 → 0	8 → 8
d1d3df1b33	38 → 1	38 → 2	38 → 35	0 → 0	0 → 0	19 → 19
e55ee035c7	64 → 22	64 → 5	64 → 37	0 → 0	0 → 0	1 → 1
517f31132d	17 → 17	17 → 0	17 → 0	0 → 0	0 → 0	42 → 42
704bd59a4e	39 → 28	39 → 9	39 → 2	8 → 7	8 → 1	2 → 2
96e42b64f2	0 → 0	0 → 0	0 → 0	31 → 31	31 → 0	0 → 0
e6171e3491	16 → 1	16 → 2	16 → 13	13 → 0	13 → 13	0 → 0
7c74a56b87	15 → 1	15 → 2	15 → 12	10 → 0	10 → 10	0 → 0
058ca2a32b	15 → 1	15 → 2	15 → 12	10 → 0	10 → 10	0 → 0
1fc3c27d95	0 → 0	0 → 0	0 → 0	0 → 0	0 → 0	17 → 17
b2da8a7b24	10 → 1	10 → 2	10 → 7	7 → 0	7 → 7	0 → 0
477d02ba72	10 → 1	10 → 2	10 → 7	7 → 0	7 → 7	0 → 0

Table 3. Showcasing the improvement in indirect call resolution for type based call graph analysis on refinement types over an analysis over base WebAssembly types.

RQ2: How does call graph precision improve across callsites?

Table 3 shows how individual indirect callsites transition across three categories when moving from a type-based analysis to RNTA: polymorphic (multiple targets), monomorphic (a single target), and unreachable (no targets).

Monomorphic Callsite Discovery. A key result is the discovery of new monomorphic callsites. Across 14 of the 20 binaries, RNTA identifies 78 indirect callsites as monomorphic that the base type system treats as polymorphic. Among binaries with polymorphic callsites, an average of 12.6% of those callsites are promoted to monomorphic by RNTA. These are direct candidates for call devirtualization (replacing an indirect call with a direct call) which can enable further optimizations such as inlining and specialization. The most notable case is fad0f2d451, where all 38 type-based polymorphic callsites are fully resolved: 22 become monomorphic and 16 are identified as unreachable. 424d9e12e7 discovers 17 new monomorphic callsites from 151 polymorphic ones, and 704bd59a4e resolves 9 of 39.

Polymorphic Callsite Resolution. On a per-binary basis, RNTA resolves a median of 55.4% of type-based polymorphic callsites. Even binaries with modest overall edge reduction benefit from refinement at the callsite level: f1fb556163 had only 0.9% edge reduction, yet RNTA resolves all 3 of its polymorphic callsites—1 becomes monomorphic and 2 become unreachable—leaving all 64 remaining reachable callsites monomorphic.

RQ3: How does RNTA compare to other state-of-the-art call graph analysis tools?

We compare our approach with other state-of-the-art call graph analysis tools, WASM-OPT [48], WASSAIL [43] and WASMA [11], and with a baseline analysis (Naive) that resolves every indirect call target to be every function in the WebAssembly table. Several tools are excluded from our

Binary	Naive		RNTA		WASM-OPT		WASMA		WASSAIL	
	Nodes	Edges	Nodes (Red%)	Edges (Red%)	Nodes (Red%)	Edges (Red%)	Nodes (Red%)	Edges (Red%)	Nodes (Red%)	Edges (Red%)
d128e52a5d	544	32780	544(0.0%)	14099(57.0%)	540(0.7%)	32890(−0.3%)	544(0.0%)	17187(47.6%)	544(0.0%)	17187(47.6%)
fd885c2d12	213	3422	209(1.9%)	1114(67.4%)	208(2.3%)	3418(0.1%)	213(0.0%)	1682(50.8%)	213(0.0%)	1682(50.8%)
0f9cbc41a9	143	776	39(72.7%)	358(53.9%)	123(14.0%)	743(4.3%)	39(72.7%)	375(51.7%)	39(72.7%)	375(51.7%)
424d9e12e7	1567	18759	1459(6.9%)	6184(67.0%)	1558(6.6%)	20044(−6.9%)	1566(0.1%)	9474(49.5%)	1566(0.1%)	9474(49.5%)
04b0c8ef0f	235	2697	234(0.4%)	681(74.7%)	228(3.0%)	2674(0.9%)	235(0.0%)	882(67.3%)	235(0.0%)	882(67.3%)
fad0f2d451	167	459	148(11.4%)	306(33.3%)	164(1.8%)	457(0.4%)	167(0.0%)	344(25.1%)	167(0.0%)	1182(−157.5%)
b5047ed15e	442	3476	397(10.2%)	1508(56.6%)	433(2.0%)	3468(0.2%)	439(0.7%)	1761(49.3%)	439(0.7%)	1761(49.3%)
47fb4f81d2	248	14712	152(38.7%)	1360(90.8%)	217(12.5%)	14653(0.4%)	169(31.9%)	1454(90.1%)	169(31.9%)	1454(90.1%)
fffb556163	97	461	90(7.2%)	314(31.9%)	94(3.1%)	440(4.6%)	91(6.2%)	321(30.4%)	91(6.2%)	321(30.4%)
d1d3df1b33	629	2538	519(17.5%)	1131(55.4%)	FAIL	FAIL	556(11.6%)	1236(51.3%)	556(11.6%)	1236(51.3%)
e55ee035c7	172	1342	118(31.4%)	388(71.1%)	171(0.6%)	1342(0.0%)	165(4.1%)	690(48.6%)	165(4.1%)	690(48.6%)
517f31132d	99	555	39(60.6%)	194(65.0%)	80(19.2%)	555(0.0%)	39(60.6%)	211(62.0%)	39(60.6%)	211(62.0%)
704bd59a4e	276	1403	273(1.1%)	778(44.5%)	255(7.6%)	1361(3.0%)	276(0.0%)	797(43.2%)	276(0.0%)	797(43.2%)
96e42b64f2	15	16	15(0.0%)	16(0.0%)	14(6.7%)	16(0.0%)	15(0.0%)	16(0.0%)	15(0.0%)	16(0.0%)
e6171e3491	111	451	81(27.0%)	253(43.9%)	111(0.0%)	461(−2.2%)	105(5.4%)	311(31.0%)	105(5.4%)	311(31.0%)
7c74a56b87	106	338	80(24.5%)	211(37.6%)	105(0.9%)	357(−5.6%)	100(5.7%)	254(24.9%)	100(5.7%)	254(24.9%)
058ca2a32b	99	327	73(26.3%)	200(38.8%)	98(1.0%)	327(0.0%)	93(6.1%)	243(25.7%)	93(6.1%)	243(25.7%)
1fc3e27d95	23	167	23(0.0%)	102(38.9%)	23(0.0%)	167(0.0%)	23(0.0%)	122(26.9%)	23(0.0%)	122(26.9%)
b2da8a7b24	66	183	43(34.8%)	106(42.1%)	65(1.5%)	183(0.0%)	60(9.1%)	128(30.1%)	60(9.1%)	128(30.1%)
477d02ba72	69	185	46(33.3%)	108(41.6%)	68(1.4%)	185(0.0%)	63(8.7%)	130(29.7%)	63(8.7%)	130(29.7%)
Average			20.3%	50.6%	4.2%	−0.1%	11.1%	41.8%	11.1%	32.6%

Table 4. Comparison of RNTA against state-of-the-art WebAssembly call graph construction tools. Node and edge reduction percentages are relative to the Naive baseline that considers all functions in the WebAssembly function table to be reachable from an indirect call.

evaluation: WASSILLY [33] timed out on each binary while generating a call graph, STURDY [28] does not support analysis over binaries that interoperate with JavaScript (the framework throws an error), and WASMATI [12] does not support WebAssembly functions that return more than one result (multi-value proposal). Table 4 compares the different analysis tools. We are unable to report on the number of monomorphic callsites discovered by WASM-OPT, WASSAIL or WASMA because they generate coarse-grained call graphs without callsite-specific information.

WASM-OPT. WASM-OPT’s generated call graph does not report the resolution of indirect calls, but rather just direct calls. We estimate the resolution of indirect calls by looking through their codebase and find that WASM-OPT performs a naive call graph analysis: at every indirect call site, it assumes that all functions in the function table may be called. As a result, WASM-OPT achieves virtually no edge reduction. They show a node reduction of 4% because they do not report imported functions in the generated call graph. WASM-OPT also fails to run on d1d3df1b33.

WASMA. WASMA performs a type-based call graph analysis, considering all functions in the table with a matching type annotation as potential indirect call targets. This yields substantially better precision than the naive analysis, with 42% average edge reduction and 11% average node reduction. However, RNTA consistently outperforms WASMA, achieving 51% average edge reduction (9 percentage points more) and 20% average node reduction (9 percentage points more). The improvement is most pronounced on binaries where refinement types effectively split function groups that share a base type: e55ee035c7 (71% vs. 49%), fd885c2d12 (67% vs. 51%), 424d9e12e7 (67% vs. 50%), and 04b0c8ef0f (75% vs. 67%).

WASSAIL. WASSAIL performs a type-based call graph analysis like WASMA. On 19 of the 20 binaries, WASSAIL produces identical results to WASMA. The exception is fad0f2d451, where WASSAIL reports 1,182 edges—158% more than the naive baseline’s 459 edges—while WASMA reports 344 edges (25% reduction) and RNTA achieves 306 edges (33% reduction). This is likely a bug in WASSAIL’s source code. Overall, WASSAIL averages 33% edge reduction, 18 percentage points less than RNTA’s 51%.

Note that if an indirect callsite uses a constant value to index into the function table, Wassail's analysis statically resolves the target to a single function. However, this pattern is not common in real-world binaries and does not improve Wassail precision in our experiments.

7 RELATED WORK

Below, we discuss relevant work in the following four categories: (i) call graph construction techniques for imperative and object-oriented programming languages, (ii) call graph construction for WebAssembly, (iii) call graph construction for other low-level languages, and (iv) work on refinement type systems. Below, we discuss relevant work in each of these categories.

7.1 Call Graph Construction for Imperative and Object-Oriented Languages.

In high-level programming languages, indirect calls are typically performed using function pointers or dynamically dispatched (virtual) method calls. The key challenge in constructing call graphs is to approximate the behavior of indirect calls, which in high-level languages involves reasoning about pointers to objects. The concepts of pointer analysis and call graph construction are closely related because determining the objects that a pointer may point to depends on what functions are reachable, and determining the functions that are reachable depends on what functions a function pointer may point to, or on the type of a reference in which a dynamically dispatched method is invoked.

During the past 30+ years, a vast literature on algorithms for approximating the behavior of indirect calls has been developed, which in essence involve associating sets of abstract objects with abstract pointers. This includes techniques based on: class hierarchies and types [7, 17, 45], context-sensitive algorithms [5, 22, 23, 38, 40], declarative techniques [6, 10], demand-driven and selective algorithms [30, 41], and frameworks for static analysis [2–4]. Overviews of pointer analysis techniques can be found in [26, 27, 36, 39]. Unfortunately, since WebAssembly is a low-level language that features neither objects nor pointers, these traditional pointer-analysis algorithms cannot be used.

7.2 WebAssembly Static Program Analysis

There are several static analysis frameworks for WebAssembly. We compare against WASSAIL [42], WASMA [11] and WASM-OPT [48]. WASSILLY [33] and STURDY [28] improve the precision of indirect calls by relying on a better value analysis to determine the index into the function table. They are both abstract interpretation frameworks. We do not compare against them because the former times out on all our subjects and the latter does not support analysis of binaries that interoperate with JavaScript. There are also other analysis tools that have a different focus than that of our work. WASMATI [12] uses code property graphs to statically detect vulnerabilities in WebAssembly binaries and use a type-based call graph analysis. WASMCHECKER [50] detects bugs in WebAssembly programs using abstract interpretation, instantiating the analysis with multiple abstract domains to find null dereferences, out-of-bounds memory accesses, and division-by-zero errors. WANILLA [37] is a static noninterference analysis for WebAssembly that tracks taints with value-sensitive relational reasoning to capture information flows between a Wasm module and its JavaScript host as well as memory integrity within the module. Like other state-of-the-art tools, they resolve indirect calls via type-based filtering. They also track taints through memory but loose precision in the face of runtime addresses, unlike us.

7.3 Inferring Richer Types for Binary Analysis

In general, inferring higher-level types for LLVM IR and binary code is a popular research direction. LLVM IR encodes some structural type information directly in getelementptr (GEP) instructions:

1030 an instruction such as `getelementptr %struct.Foo, ptr %p, i32 0, i32 1` reveals that %p
1031 points into a Foo, providing structural evidence that downstream analyses can exploit. In practice,
1032 however, this information is often incomplete or broken. Liu et al. [31] present TFA, which improves
1033 indirect-call resolution in LLVM IR for C/C++ programs by first consulting debug metadata nodes
1034 embedded in the IR to reconstruct missing struct names and fields, then applying multi-layer type
1035 and data-flow analysis. WebAssembly binaries, by contrast, are routinely shipped at minimum size
1036 and DWARF debug information is rarely retained, so this recovery strategy is unavailable to us; our
1037 refinement type system instead recovers the ptr/num distinction purely from program behaviour.
1038 Concurrent work by Nicolosi et al. [32] also uses GEP-encoded structural types to reconstruct
1039 transparent pointee types after LLVM 17 collapsed all pointer types to the single opaque type ptr.
1040 They show that richer type information improves the precision of downstream analyses, something
1041 that we believe is true for RNTA as well.

1042 The binary setting is perhaps more closely matched to WebAssembly. Concurrent work by
1043 Bosamiya et al. [9] present TRES, which reconstructs C-like types (structs, arrays, unions) from
1044 stripped x86/x64 binaries to aid human reverse engineers, explicitly trading soundness for utility
1045 via *conditional conservativeness* with toggleable opportunistic inferences. Struct types are inferred
1046 via colocation: x86 instructions encode constant-offset memory accesses explicitly (e.g., `mov eax,`
1047 `[rbx+8]`), and a finite set of named registers with well-defined calling conventions means the
1048 analysis can attribute semantic roles (pointer, integer, return value) to variables from the outset.
1049 WebAssembly is a stack machine with anonymous stack slots and local variables that carry no
1050 semantic meaning beyond their base type, so an equivalent colocation type reconstruction does
1051 not exist until pointer-typed values have first been identified, which is precisely what our work
1052 provides. Our work instead prioritises soundness over human-readable output: by computing a
1053 greatest fixpoint over a finite lattice, every inferred refinement type is a safe overapproximation,
1054 making our results suitable for downstream static analyses.

1055 7.4 Refinement Type Systems

1056 Refinement types have long been used to enhance type systems with logical predicates that constrain
1057 the set of values described by a type [18, 19, 34, 46, 49], allowing programmers to express more
1058 precise program properties and enabling program verification. The most closely related work is
1059 WasmPrecheck [20], a richer type system for WebAssembly that uses *indexed types* to express static
1060 constraints that enable safe removal of dynamic checks for type and memory safety. WasmPrecheck
1061 supports general constraints on index terms, while we support only singleton refinements and
1062 refining `i32` to `ptr(a)` or `num(n)`. Another point of comparison is that both WasmPrecheck and our
1063 work shows how to embed a Wasm module into a refined type system. However, Geller et al. give
1064 only a naive embedding — they show that a Wasm module can be (automatically) embedded into
1065 WashPrecheck by replacing all type annotations in the Wasm module with indexed types that have
1066 no constraints.

1068 8 CONCLUSIONS AND FUTURE WORK

1069 We have presented RNTA, a refinement type system and call graph analysis for WebAssembly that
1070 refines `i32` types to either a singleton pointer type `ptr(l, n)` (i.e., a pointer to symbolic location *l*
1071 at abstract offset *n*) or a singleton number `num(n)`. Additionally, we type WebAssembly memory,
1072 which the original type system leaves untyped, enabling precise tracking of integer values as they
1073 flow through memory. The results of our experiments show that, when compared against a baseline
1074 analysis and other state-of-the-art WebAssembly static analysis tools that use WebAssembly's
1075 standard type system, RNTA reduces reachable functions by 9.9% and call graph edges by 14.6% while
1076 increasing indirect calls that are resolved to a unique target by 12.6%. At present, our analysis only
1077

1078

1079 supports WebAssembly version 1.0. We plan to extend our analysis to support subsequent versions
1080 of WebAssembly, and to improve precision through standard analysis techniques as well as through
1081 a multi-language analysis with JavaScript, WebAssembly’s most popular host language.
1082

1083 8.1 Supporting WebAssembly 2.0 and 3.0

1084 WebAssembly’s first version update introduced vector instructions, to support 128-bit wide SIMD
1085 functionality; bulk memory instructions, to support copying and initializing regions of memory;
1086 multi-value results, to support returning more than one value by functions, blocks and instructions;
1087 reference types to support opaque first-class references to functions or pointers; non-trapping con-
1088 versions to allow seamless conversions from float to integer types; and sign-extension instructions
1089 that allow directly extending the width of signed integer values. We already support multi-value
1090 results and do not anticipate difficulty extending our refinement type system to include these new
1091 instructions.
1092

1093 For WebAssembly 3.0, the WebAssembly working group has already implemented other language
1094 extensions [1], the most interesting of which is the Garbage Collection (GC) proposal which
1095 introduces aggregate types like structures and arrays with optional mutable fields and packed size
1096 fields, reference types, imported types, host types and address types, which are used as offsets
1097 into memories and tables. Additionally, linear memory is now 32- and 64-bit addressable, that
1098 is, linear memory can be accessed with $i32$ and $i64$ values. To add support for the GC proposal,
1099 we would first refine $i64$ types to have a similar refinement to the one presented in the paper
1100 for $i32$, such that, $\{ptr32(n), num32(n)\} <: i32(n) \wedge \{ptr64(n), num64(n)\} <: i64(n)$. Address
1101 types ($addrtype ::= i32 | i64$) are a *subset* of number types ($numtype ::= i32 | i64 | f32 | f64$). Since
1102 our refinement introduces *subtype* constraints between pointers and integers, we would still see
1103 precision improvements.
1104

1105 8.2 JavaScript Client Analysis

1106 WebAssembly was originally designed to interoperate with JavaScript and JavaScript remains its
1107 most popular host. Our long-term goal is to develop static-analysis techniques for WebAssembly
1108 binaries used in JavaScript clients for code-size reduction (specializing a WebAssembly library to a
1109 client’s use case), optimization (module splitting to improve run-time performance) and detection
1110 of security vulnerabilities. RNTA is a closed-world analysis over WebAssembly, in that it makes
1111 worst-case assumptions about the behaviour of a JavaScript client. WebAssembly interoperates
1112 with JavaScript through imported and exported functions and a closed-world analysis makes the
1113 assumption that all exported functions in the binary can be called by the client. However, Thalakottur
1114 et al. [44] have shown that clients often only call a subset of the binary’s exported functions. An
1115 open-world analysis of WebAssembly and JavaScript that detects the exported functions called
1116 by the JavaScript client and can propagate values from JavaScript into the WebAssembly analysis
1117 would further improve the precision of the WebAssembly analysis.
1118

1119 8.3 Improvements to Precision

1120 Our analysis discovers refinement types that are flow-insensitive, context-insensitive and intra-
1121 procedural. Additionally, the domain of values is fairly simple — a flat lattice of natural numbers.
1122 Considering the relative simplicity of our analysis we plan on employing several traditional tech-
1123 niques to increase precision of the analysis. For example, we plan to change the domain of values to
1124 be the power-set lattice over natural numbers. We also plan to make our analysis flow and context
1125 sensitive and propagate types through direct calls to increase precision at callsites.
1126
1127

9 DATA-AVAILABILITY STATEMENT

RNTA along with the binaries and scripts used for evaluation is available as an artifact at <https://anonymous.4open.science/r/RNTA>. We intend to submit the artifact for Artifact Evaluation.

REFERENCES

- [1] Andreas Rossberg (Ed.). [n. d.]. *Current up-to-date WebAssembly Specification*. <https://webassembly.github.io/memory64/core/index.html>
- [2] 2026. Doop: Framework for Java Pointer and Taint Analysis (using P/Taint). <https://github.com/plast-lab/doop>. Accessed: 2026-03-16.
- [3] 2026. SOOT: A framework for analyzing and transforming Java and Android applications. <https://soot-oss.github.io/soot/>. Accessed: 2026-03-16.
- [4] 2026. WALA. <https://github.com/wala/WALA>. Accessed: 2026-03-16.
- [5] Ole Agesen. 1995. The Cartesian Product Algorithm: Simple and Precise Type Inference Of Parametric Polymorphism. In *ECOOP'95 - Object-Oriented Programming, 9th European Conference, Århus, Denmark, August 7-11, 1995, Proceedings (Lecture Notes in Computer Science, Vol. 952)*, Walter G. Olthoff (Ed.). Springer, 2–26. https://doi.org/10.1007/3-540-49538-X_2
- [6] Pavel Avgustinov, Oege de Moor, Michael Peyton Jones, and Max Schäfer. 2016. QL: Object-oriented Queries on Relational Data. In *30th European Conference on Object-Oriented Programming, ECOOP 2016, Rome, Italy, July 18-22, 2016 (LIPIcs, Vol. 56)*, Shriram Krishnamurthi and Benjamin S. Lerner (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2:1–2:25. <https://doi.org/10.4230/LIPICS.ECOOP.2016.2>
- [7] David F. Bacon and Peter F. Sweeney. 1996. Fast Static Analysis of C++ Virtual Function Calls. In *Proceedings of the 1996 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications, OOPSLA 1996, San Jose, California, USA, October 6-10, 1996*, Lougie Anderson and James Coplien (Eds.). ACM, 324–341. <https://doi.org/10.1145/236337.236371>
- [8] Jay Bosamiya, Maverick Woo, and Bryan Parno. 2025. {TRex}: Practical Type Reconstruction for Binary Code. In *34th USENIX Security Symposium (USENIX Security 25)*. 6897–6915.
- [9] Jay Bosamiya, Maverick Woo, and Bryan Parno. 2025. TRex: practical type reconstruction for binary code. In *Proceedings of the 34th USENIX Conference on Security Symposium (Seattle, WA, USA) (SEC '25)*. USENIX Association, USA, Article 354, 19 pages.
- [10] Martin Bravenboer and Yannis Smaragdakis. 2009. Strictly declarative specification of sophisticated points-to analyses. In *Proceedings of the 24th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2009, October 25-29, 2009, Orlando, Florida, USA*, Shail Arora and Gary T. Leavens (Eds.). ACM, 243–262. <https://doi.org/10.1145/1640089.1640108>
- [11] Florian Breithfelder, Tobias Roth, Lars Baumgärtner, and Mira Mezini. 2023. WasmA: A Static WebAssembly Analysis Framework for Everyone. In *2023 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. 753–757. <https://doi.org/10.1109/SANER56733.2023.00085>
- [12] Tiago Brito, Pedro Lopes, Nuno Santos, and José Frago Santos. [n. d.]. Wasmati: An Efficient Static Vulnerability Scanner for WebAssembly. ([n. d.]), 102745. <https://doi.org/10.1016/j.cose.2022.102745>
- [13] Chris Fallin. 2025. *Waffle: WebAssembly Analysis Framework for Lightweight Experimentation*. Bytecode Alliance. <https://github.com/bytecodealliance/waffle>
- [14] Ravi Chugh, David Herman, and Ranjit Jhala. 2012. Dependent types for JavaScript. In *Proceedings of the ACM international conference on Object oriented programming systems languages and applications*. 587–606.
- [15] Ron Cytron, Jeanne Ferrante, Barry K Rosen, Mark N Wegman, and F Kenneth Zadeck. 1989. An efficient method of computing static single assignment form. In *Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. 25–35.
- [16] Nicolaas Govert De Bruijn. 1972. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. In *Indagationes mathematicae (proceedings)*, Vol. 75. Elsevier, 381–392.
- [17] Jeffrey Dean, David Grove, and Craig Chambers. 1995. Optimization of Object-Oriented Programs Using Static Class Hierarchy Analysis. In *ECOOP'95 - Object-Oriented Programming, 9th European Conference, Århus, Denmark, August 7-11, 1995, Proceedings (Lecture Notes in Computer Science, Vol. 952)*, Walter G. Olthoff (Ed.). Springer, 77–101. https://doi.org/10.1007/3-540-49538-X_5
- [18] Jeremy Freeman, Catalin Hritcu, Marco Gaboardi, Vincent Laporte, Sergey Firsov, Gregory Malecha, Dustin Swasey, Conor McBride Watt, and Benjamin C. Pierce. 2020. RefinedC: Automating the Foundational Verification of C Code with Refined Types. In *Proceedings of the ACM on Programming Languages (OOPSLA)*, Vol. 4. 1–29. <https://doi.org/10.1145/3428198>

- 1177 [19] Tim Freeman and Frank Pfenning. 1991. Refinement types for ML. In *Proceedings of the ACM SIGPLAN 1991 conference*
1178 *on Programming language design and implementation*. 268–277.
- 1179 [20] Adam T. Geller, Justin P. Frank, and William J. Bowman. 2024. Indexed Types for a Statically Safe WebAssembly.
1180 *Proceedings of the ACM on Programming Languages* 8, POPL (2024), 2395–2424. <https://doi.org/10.1145/3632922>
- 1181 [21] Jessica Gronski, Kenneth Knowles, Aaron Tomb, Stephen N Freund, and Cormac Flanagan. 2006. Sage: Hybrid checking
1182 for flexible specifications. In *Scheme and Functional Programming Workshop*, Vol. 6. 93–104.
- 1183 [22] David Grove and Craig Chambers. 2001. A framework for call graph construction algorithms. *ACM Transactions on*
1184 *Programming Languages and Systems (TOPLAS)* 23, 6 (2001), 685–746.
- 1185 [23] David Grove, Greg DeFouw, Jeffrey Dean, and Craig Chambers. 1997. Call Graph Construction in Object-Oriented
1186 Languages. In *Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*.
- 1187 [24] Andreas Haas, Andreas Rossberg, Derek L Schuff, Ben L Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon
1188 Zakai, and JF Bastien. 2017. Bringing the Web up to Speed with WebAssembly. In *Proceedings of the 38th ACM SIGPLAN*
Conference on Programming Language Design and Implementation (Barcelona, Spain) (PLDI 2017). ACM, New York, NY,
1189 USA, 185–200. <https://doi.org/10.1145/3062341.3062363>
- 1190 [25] Aaron Hilbig, Daniel Lehmann, and Michael Pradel. 2021. An Empirical Study of Real-World WebAssembly Binaries:
1191 Security, Languages, Use Cases. In *Proceedings of the Web Conference 2021 (Ljubljana, Slovenia) (WWW '21)*. Association
1192 for Computing Machinery, New York, NY, USA, 2696–2708. <https://doi.org/10.1145/3442381.3450138>
- 1193 [26] Michael Hind. 2001. Pointer analysis: haven't we solved this problem yet?. In *Proceedings of the 2001 ACM SIGPLAN-*
SIGSOFT Workshop on Program Analysis For Software Tools and Engineering, PASTE'01, Snowbird, Utah, USA, June 18-19,
1194 *2001*, John Field and Gregor Snelting (Eds.). ACM, 54–61. <https://doi.org/10.1145/379605.379665>
- 1195 [27] Michael Hind and Anthony Pioli. 2000. Which pointer analysis should I use?. In *Proceedings of the International*
Symposium on Software Testing and Analysis, ISSTA 2000, Portland, OR, USA, August 21-24, 2000, Debra J. Richardson
1196 and Mary Jean Harold (Eds.). ACM, 113–123. <https://doi.org/10.1145/347324.348916>
- 1197 [28] Sven Keidel and Sebastian Erdweg. 2019. Sound and reusable components for abstract interpretation. *Proc. ACM*
Program. Lang. 3, OOPSLA, Article 176 (Oct. 2019), 28 pages. <https://doi.org/10.1145/3360602>
- 1198 [29] Daniel Lehmann, Michelle Thalakkottur, Frank Tip, and Michael Pradel. 2023. That's a Tough Call: Studying the
1199 Challenges of Call Graph Construction for WebAssembly. In *Proceedings of the 32nd ACM SIGSOFT International*
Symposium on Software Testing and Analysis (ISSTA '23). Association for Computing Machinery, New York, NY, USA.
1200 <https://doi.org/10.1145/3597926.3598104>
- 1201 [30] Yue Li, Tian Tan, Anders Møller, and Yannis Smaragdakis. 2020. A Principled Approach to Selective Context Sensitivity
1202 for Pointer Analysis. *ACM Trans. Program. Lang. Syst.* 42, 2 (2020), 10:1–10:40. <https://doi.org/10.1145/3381915>
- 1203 [31] Dinghao Liu, Shouling Ji, Kangjie Lu, and Qinming He. 2024. Improving {Indirect-Call} analysis in {LLVM} with
1204 type and {Data-Flow}{Co-Analysis}. In *33rd USENIX Security Symposium (USENIX Security 24)*. 5895–5912.
- 1205 [32] Niccolò Nicolosi, Gabriele Magnani, Emilio Corigliano, Davide Baroffio, Federico Reghenzani, and Giovanni Agosta.
1206 2026. Type Deduction Analysis: Reconstructing Transparent Pointer Types in LLVM-IR. In *Proceedings of the 35th*
ACM SIGPLAN International Conference on Compiler Construction. 194–204.
- 1207 [33] Mattia Paccamiccio, Franco Raimondi, and Michele Loreti. 2024. Building Call Graph of WebAssembly Programs via
1208 Abstract Semantics. arXiv:2407.14527 [cs.SE] <https://arxiv.org/abs/2407.14527>
- 1209 [34] Patrick M. Rondon, Kohei Kawaguchi, and Ranjit Jhala. 2008. Liquid Types. In *Proceedings of the ACM SIGPLAN*
Conference on Programming Language Design and Implementation (PLDI). 159–169.
- 1210 [35] Patrick Maxim Rondon, Ming Kawaguchi, and Ranjit Jhala. 2010. Low-level liquid types. *ACM Sigplan Notices* 45, 1
1211 (2010), 131–144.
- 1212 [36] Barbara G. Ryder. 2003. Dimensions of Precision in Reference Analysis of Object-Oriented Programming Languages.
1213 In *Compiler Construction, 12th International Conference, CC 2003, Held as Part of the Joint European Conferences on*
Theory and Practice of Software, ETAPS 2003, Warsaw, Poland, April 7-11, 2003, Proceedings (Lecture Notes in Computer
1214 *Science, Vol. 2622)*, Görel Hedin (Ed.). Springer, 126–137. https://doi.org/10.1007/3-540-36579-6_10
- 1215 [37] Markus Scherer, Jeppe Fredsgaard Blaabjerg, Alexander Sjösten, and Matteo Maffei. 2025. Wanilla: Sound Noninter-
1216 ference Analysis for WebAssembly. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Commu-*
nications Security (Taipei, Taiwan) (CCS '25). Association for Computing Machinery, New York, NY, USA, 126–140.
1217 <https://doi.org/10.1145/3719027.3765156>
- 1218 [38] Olin Shivers. 1991. *Control-Flow Analysis of Higher-Order Languages*. Ph. D. Dissertation. Carnegie-Mellon University.
1219 CMU-CS-91-145.
- 1220 [39] Yannis Smaragdakis and George Balatsouras. 2015. Pointer Analysis. *Found. Trends Program. Lang.* 2, 1 (2015), 1–69.
1221 <https://doi.org/10.1561/2500000014>
- 1222 [40] Yannis Smaragdakis, Martin Bravenboer, and Ondrej Lhoták. 2011. Pick your contexts well: understanding object-
1223 sensitivity. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages,*
POPL 2011, Austin, TX, USA, January 26-28, 2011, Thomas Ball and Mooly Sagiv (Eds.). ACM, 17–30. <https://doi.org/10.1145/1860295.1860300>
- 1224
1225

- 1226 1145/1926385.1926390
- 1227 [41] Manu Sridharan, Denis Gopan, Lexin Shan, and Rastislav Bodik. 2005. Demand-driven points-to analysis for Java. In
- 1228 *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and*
- 1229 *Applications, OOPSLA 2005, October 16-20, 2005, San Diego, CA, USA*, Ralph E. Johnson and Richard P. Gabriel (Eds.).
- 1230 ACM, 59–76. <https://doi.org/10.1145/1094811.1094817>
- 1231 [42] Quentin Stiévenart and Coen De Roover. 2021. Wassail: a WebAssembly Static Analysis Library (*ProWeb21*). <https://2021.programming-conference.org/home/proweb-2021> Fifth International Workshop on Programming Technology
- 1232 for the Future Web.
- 1233 [43] Quentin Stiévenart and Coen De Roover. 2020. Compositional Information Flow Analysis for WebAssembly Programs.
- 1234 In *2020 IEEE 20th International Working Conference on Source Code Analysis and Manipulation (SCAM)*. 13–24. <https://doi.org/10.1109/SCAM51674.2020.00007>
- 1235 [44] Michelle Thalakkottur, Max Bernstein, Daniel Lehmann, Michael Pradel, and Frank Tip. 2026. An Empirical Study of
- 1236 WebAssembly Usage in Node.js. In *48rd IEEE/ACM International Conference on Software Engineering, ICSE 2026*.
- 1237 [45] Frank Tip and Jens Palsberg. 2000. Scalable propagation-based call graph construction algorithms. In *Proceedings*
- 1238 *of the 2000 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications, OOPSLA*
- 1239 *2000, Minneapolis, Minnesota, USA, October 15-19, 2000*, Mary Beth Rosson and Doug Lea (Eds.). ACM, 281–293.
- 1240 <https://doi.org/10.1145/353171.353190>
- 1241 [46] Niki Vazou, Eric Seidel, and Ranjit Jhala. 2014. Refinement Types for Haskell. In *Proceedings of the ACM SIGPLAN*
- 1242 *International Conference on Functional Programming (ICFP)*. 269–282.
- 1243 [47] Andreas Rossberg (Ed.). 2022-03-31. *WebAssembly 1.0 Core Specification*. World Wide Web Consortium (W3C).
- 1244 <https://www.w3.org/TR/wasm-core-1/>
- 1245 [48] WebAssembly Community Group. 2025. Binaryen: Optimizer and compiler/toolchain library for WebAssembly.
- 1246 <https://github.com/WebAssembly/binaryen>.
- 1247 [49] Hongwei Xi and Frank Pfenning. 1999. Dependent Types in Practical Programming. In *Proceedings of the ACM SIGPLAN*
- 1248 *Symposium on Principles of Programming Languages (POPL)*. 214–227.
- 1249 [50] Junjie Zhuang, Hao Jiang, and Baojian Hua. 2024. WasmChecker: Effectively Detecting WebAssembly Bugs via Static
- 1250 Program Analysis. In *2024 8th International Conference on Electrical, Mechanical and Computer Engineering (ICEMCE)*.
- 1251 1669–1678. <https://doi.org/10.1109/ICEMCE64157.2024.10861926>
- 1252
- 1253
- 1254
- 1255
- 1256
- 1257
- 1258
- 1259
- 1260
- 1261
- 1262
- 1263
- 1264
- 1265
- 1266
- 1267
- 1268
- 1269
- 1270
- 1271
- 1272
- 1273
- 1274

Appendices

A INSTRUCTION TYPING FOR THE REFINEMENT TYPE SYSTEM

$C^r ::= \{\text{func } tf_r^*, \text{ local } \tau_r^*, \text{ global } \tau_r^*, \text{ table } n^?, \text{ memory } n^?, \text{ label } (\tau_r^*, \Psi)^*, \text{ return } (\tau_r^*, \Psi)^?\}$

Typing WebAssembly Instructions with Refinement Types

$R; C^r \vdash e_r^* : tf_r$

$$\frac{R = \{\text{inst } C^{r*}, \text{tab } n^*, \text{mem } \Psi^*\} \quad (R \vdash \text{inst} : C^r)^* \quad (n \leq |cl|^*)^* \quad \vdash \text{s.mem} : R_\Psi}{\vdash \{\text{inst } \text{inst}^*, \text{tab } (cl^*)^*, \text{mem } (b^*)^*\} : R}$$

$$\frac{\forall a \in 0 \dots \text{size}(\text{s.mem}) \quad \text{s.mem}(a) = R_\Psi(a)}{\vdash \text{s.mem} : R_\Psi}$$

$$\zeta \text{ fresh} \quad \zeta; \cdot; C^\alpha \vdash ex^* \text{ func } tf_w^l \text{ local } \tau_w^l e^* : \alpha^m \rightarrow \alpha^n; S'; \zeta'$$

$$\rho = \text{solve}(S')$$

$$\tau_r^m = \rho[\alpha^m] \quad \tau_r^n = \rho[\alpha^n] \quad \Sigma_{pre} = \rho[C] \quad \Sigma_{post} = \rho[C']$$

$$\frac{\exists \Psi_{pre}, \Psi_{post}, \exists \eta_{pre}, \eta_{post} \cdot \text{Bij}(\eta_{pre}) \wedge \text{Bij}(\eta_{post}) \wedge \eta_{pre}(\Sigma_{pre}) = \Psi_{pre} \wedge \eta_{post}(\Sigma_{post}) = \Psi_{post} \wedge \eta_{post} \supseteq \eta_{pre} \quad tf_r = \tau_r^m, \Psi_{pre} \rightarrow \tau_r^n, \Psi_{post}}{C^\alpha \vdash ex^* \text{ func } tf_w^l \text{ local } \tau_w^l e^* : ex^* tf_r; \rho}$$

$$\frac{\begin{array}{l} \alpha_w^g \text{ fresh} \\ C^\alpha = \{\text{func } tf_w^g, \text{global } \alpha_w^g, \text{table } n^?, \text{memory } n^?\} \\ (C^\alpha \vdash ex^* \text{ func } tf_w^l \text{ local } \tau_w^l e^* : ex^* tf_r; \rho)^* \\ C^r = \{\text{func } tf_r^*, \text{global } \tau_r^g, \text{table } n^?, \text{memory } n^?\} \\ (\text{module } f^* \text{ global } \tau_w^g \text{ tab}^? \text{ mem}^?) \rightsquigarrow (\text{module } f^* \text{ global } \tau_r^g \text{ tab}^? \text{ mem}^?) \\ (R; C^r \vdash f : ex^* tf_r)^* \end{array}}{R \vdash \text{module } f^* \text{ global } \tau_w^g \text{ tab}^? \text{ mem}^?} \text{MODULE}$$

$$\frac{R, C^r, \text{local}(\tau_r^m, \tau_r^l), \text{label}(\tau_r^n, \Psi_{post}), \text{return}(\tau_r^n, \Psi_{post}) \vdash e^* : \epsilon \rightarrow \tau_r^n}{R; C^r \vdash ex^* \text{ func } tf_r^l \text{ local } \tau_r^l e^* : ex^* tf_r} \text{FUNCTION}$$

$$\frac{R; C^r, \text{label}(\tau_r^n, \Psi_{post}) \vdash e^* : \tau_r^m, \Psi_{pre} \rightarrow (\tau_r')^n, \Psi'_{post} \quad (\tau_r' <: \tau_r)^n \quad \Psi'_{post} <: \Psi_{post}}{R; C^r \vdash \text{block } tf_r e^* \text{ end} : tf_r} \text{BLOCK}$$

$$\frac{R; C^r, \text{label}(\tau_r^m, \Psi_{pre}) \vdash e^* : \tau_r^m, \Psi_{pre} \rightarrow (\tau_r')^n, \Psi'_{post} \quad (\tau_r' <: \tau_r)^n \quad \Psi'_{post} <: \Psi_{post}}{R; C^r \vdash \text{loop } tf_r e^* \text{ end} : tf_r} \text{LOOP}$$

Fig. 8. Typing Rules for the Refinement Type System

```

1324 def load_and_extend( $a, \tau, tp\_sz^?, \Psi$ ) :=
1325   let  $\tau_r(n_1) = \Psi(a)$ 
1326   let  $\tau_r(n_2) = \Psi(a - 4)$ 
1327   if  $tp$  given then let  $N = |tp|$  else let  $N = |\tau|$ 
1328   if  $wasm\_type(\tau_r(n_1)) == \tau$                                then let  $\tau_r(n) = \tau_r(n_1)$ 
1329   if  $|\tau_r(n_1)| < N$                                          then let  $\tau_r(n) = \tau(\text{val}(\text{bits}(n_2) \text{ bits}(n_1)))$ 
1330   if  $|\tau_r(n_1)| > N$                                          then let  $\tau_r(n) = \tau(\text{val}(\text{bits}(n_1)_{0..N}))$ 
1331   if  $(\tau_r(n_1) = (\text{i32}(n_1) \vee \text{ptr}(n_1) \vee \text{num}(n_1)) \wedge \tau = \text{f32})$  then let  $\tau_r(n) = \text{f32}(\text{bits}(n_1) \text{ as } \text{f32})$ 
1332   if  $(\tau_r(n_1) = \text{f32}(n_1) \wedge \tau = (\text{i32} \vee \text{ptr} \vee \text{num}))$  then let  $\tau_r(n) = \text{i32}(\text{f32}(n_1) \text{ as } \text{i32})$ 
1333   if  $tp$  given
1334   then return  $\tau_r(\text{extend\_sxn}_{|\tau_w|}(n))$ 
1335   else return  $\tau_r(n)$ 

```

$$\frac{C_{\text{memory}}^r = n \quad \tau_r^1 <: \text{ptr}(c) \quad \tau_r^2 = \text{load_and_extend}(c + o, \tau, (tp_sx)^?, R_\Psi) \quad 2^a \leq (|tp| <)^? |\tau_r^2| \quad (tp_sz)^? = \epsilon \vee \tau_r^2 = \text{im}}{R; C^r \vdash \tau.\text{load}(tp_sx)^? a o : \tau_r^1 \rightarrow \tau_r^2} \text{LOAD}$$

```

1345 def wrap_and_store( $p, \tau_r(n), \tau, tp^?, R$ ) :=
1346   if  $tp$  then let  $N = |tp|$  else let  $N = |\tau|$ 
1347   if  $|\tau| == 64$  then  $R_\Psi[p \mapsto \text{i32}(\text{val}(\text{bits}(n)_{0..32})) \wedge (p - 4) \mapsto \text{i32}(\text{val}(\text{bits}(n)_{32..64}))]$ 
1348   if  $|\tau_r(n)| > N$  then
1349      $R_\Psi[p \mapsto \tau_r(\text{wrap}_{|\tau|, N}(n))]$ 
1350   else  $R_\Psi[p \mapsto \tau_r(n)]$ 

```

$$\frac{C_{\text{memory}}^r = n \quad 2^a \leq (|tp| <)^? |\tau_r^2| \quad tp^? = \epsilon \vee \tau_r^2 = \text{im} \quad \tau_r^1 <: \text{ptr32}(c) \quad \text{wrap_and_store}(c + o, \tau_r^2, \tau_w, tp^?, R) \quad \tau_r^2 <: \tau(\top)}{R, C^r \vdash \tau.\text{store } tp^? a o : \tau_r^1 \tau_r^2 \rightarrow \epsilon} \text{STORE}$$

$$\frac{C_{\text{memory}}^r = n \quad \text{num32}(n_1) <: \tau_r^1 \quad \text{num32}(n_2) <: \tau_r^2}{R; C^r \vdash \text{grow_memory} : \tau_r^1 \rightarrow \tau_r^2} \text{GROW-MEMORY}$$

$$\frac{C_{\text{memory}}^r = n \quad \text{num32}(n_1) <: \tau_r \quad |R_\Psi| = n}{R; C^r \vdash \text{current_memory} : \epsilon \rightarrow \tau_r} \text{CURRENT-MEMORY}$$

$$\frac{C_{\text{func}}^r(i) = (\tau_r')^m, \Psi'_{\text{pre}} \rightarrow (\tau_r')^n, \Psi'_{\text{post}} \quad t_{f_r} = \tau_r^m, \Psi_{\text{pre}} \rightarrow \tau_r^n, \Psi_{\text{post}} \quad (\tau_r' <: \tau_r')^m \quad (\tau_r' <: \tau_r)^n \quad \Psi'_{\text{pre}} <: \Psi_{\text{pre}} \quad \Psi'_{\text{post}} <: \Psi'_{\text{post}}}{R; C^r \vdash \text{call } i t_{f_r} : t_{f_r}} \text{CALL}$$

$$\frac{t_{f_r} = \tau_r^m, \Psi_{\text{pre}} \rightarrow \tau_r^n, \Psi_{\text{post}} \quad C_{\text{table}}^r = n \quad \text{num32}(c) <: \tau_r^1}{R; C^r \vdash \text{call_indirect } t_{f_r} : \tau_r^m, \tau_r^1 \rightarrow \tau_r^n, \Psi_{\text{post}}} \text{CALL-INDIRECT}$$

Fig. 9. Typing Rules for Refinement Type System continued.

1372

$$\begin{array}{c}
 1373 \\
 1374 \\
 1375 \\
 1376 \\
 1377 \\
 1378 \\
 1379 \\
 1380 \\
 1381 \\
 1382 \\
 1383 \\
 1384 \\
 1385 \\
 1386 \\
 1387 \\
 1388 \\
 1389 \\
 1390 \\
 1391 \\
 1392 \\
 1393 \\
 1394 \\
 1395 \\
 1396 \\
 1397 \\
 1398 \\
 1399 \\
 1400 \\
 1401 \\
 1402 \\
 1403 \\
 1404 \\
 1405 \\
 1406 \\
 1407 \\
 1408 \\
 1409 \\
 1410 \\
 1411 \\
 1412 \\
 1413 \\
 1414 \\
 1415 \\
 1416 \\
 1417 \\
 1418 \\
 1419 \\
 1420 \\
 1421
 \end{array}$$

$$\begin{array}{c}
 \frac{}{R; C^r \vdash \tau.\mathbf{const} \ c : \epsilon \rightarrow \tau(c)}^{\text{CONST}} \qquad \frac{\tau_r^2 = \mathbf{unop}(\tau_r^1)}{R; C^r \vdash \tau.\mathbf{unop} : \tau_r^1 \rightarrow \tau_r^2}^{\text{UNOP}} \\
 \\
 \frac{\tau_r^3 = \mathbf{binop}(\tau_r^1, \tau_r^2) \quad \tau_r^3 <: \tau(\top)}{R; C^r \vdash \tau.\mathbf{binop} : \tau_r^1 \tau_r^2 \rightarrow \tau_r^3}^{\text{BINOP}} \qquad \frac{\tau_r^2 = \mathbf{testop}(\tau_r^1) \quad \tau_r^2 <: \mathbf{num}(\top)}{R; C^r \vdash \tau.\mathbf{testop} : \tau_r^1 \rightarrow \tau_r^2}^{\text{TESTOP}} \\
 \\
 \frac{\tau_r^3 = \mathbf{relop}(\tau_r^1, \tau_r^2) \quad \tau_r^3 <: \mathbf{num}(\top)}{R; C^r \vdash \tau.\mathbf{relop} : \tau_r^1 \tau_r^2 \rightarrow \tau_r^3}^{\text{RELOP}} \qquad \frac{}{R; C^r \vdash \mathbf{unreachable} : \tau_r^n \rightarrow \tau_r^m}^{\text{UNREACHABLE}} \\
 \\
 \frac{}{R; C^r \vdash \mathbf{nop} : \epsilon \rightarrow \epsilon}^{\text{NOP}} \qquad \frac{}{R; C^r \vdash \mathbf{drop} : \tau_r \rightarrow \epsilon}^{\text{DROP}} \\
 \\
 \frac{}{R; C^r \vdash \mathbf{select} : \tau_r^1 \tau_r^2 \mathbf{num32}(n) \rightarrow \tau_r^3}^{\text{SELECT}} \\
 \\
 \frac{\begin{array}{l} t_f^r = \tau_r^m, \Psi_{\text{pre}} \rightarrow \tau_r^n, \Psi_{\text{post}} \quad R_{\Psi} = \Psi_{\text{pre}} \\ R; C^r, \mathbf{label}(\tau_r^n) \vdash e_1^* : \tau_r^m, \Psi_{\text{pre}} \rightarrow (\tau_r')^n, \Psi'_{\text{post}} \\ R; C^r, \mathbf{label}(\tau_r^n) \vdash e_2^* : \tau_r^m, \Psi_{\text{pre}} \rightarrow (\hat{\tau}_r'')^n, \Psi''_{\text{post}} \\ (\tau_r' <: \tau_r)^n \quad (\hat{\tau}_r'' <: \tau_r)^n \quad \Psi'_{\text{post}} <: \Psi_{\text{post}} \quad \Psi''_{\text{post}} <: \Psi_{\text{post}} \end{array}}{R; C^r \vdash \mathbf{if} \ t_f^r \ \mathbf{then} \ e_1^* \ \mathbf{else} \ e_2^* \ \mathbf{end} : t_f^r}^{\text{IF-ELSE}} \\
 \\
 \frac{C^r_{\mathbf{label}}(i) = (\tau_r')^m, \Psi' \quad (\tau_r <: \tau_r')^m \quad R_{\Psi} <: \Psi'}{R; C^r \vdash \mathbf{br} \ i : \tau_r^* \tau_r^m \rightarrow \tau_r^*}^{\text{BR}} \\
 \\
 \frac{\begin{array}{l} C^r_{\mathbf{label}}(i) = (\tau_r')^m, \Psi' \quad \mathbf{num32}(n) <: \tau_r^1 \\ (\tau_r <: \tau_r')^m \quad R_{\Psi} <: \Psi' \end{array}}{R; C^r \vdash \mathbf{br_if} \ i : \tau_r^m \tau_r^1 \rightarrow \tau_r^m}^{\text{BR-IF}} \qquad \frac{\begin{array}{l} C^r_{\mathbf{label}}(i) = ((\tau_r')^m, \Psi)^+ \\ ((\tau_r <: \tau_r')^m)^+ \quad (R_{\Psi} <: \Psi)^+ \end{array}}{R; C^r \vdash \mathbf{br_table} \ i^+ : \tau_r^* \tau_r^m \tau_r^1 \rightarrow \tau_r^*}^{\text{BR-TABLE}} \\
 \\
 \frac{\begin{array}{l} C^r_{\mathbf{label}}(i) = (\tau_r')^m, \Psi \\ (\tau_r <: \tau_r')^m \quad R_{\Psi} <: \Psi \end{array}}{R; C^r \vdash \mathbf{return} : \tau_r^* \tau_r^m \rightarrow \tau_r^*}^{\text{RETURN}} \qquad \frac{C^r_{\mathbf{local}}(i) = \tau_r}{R; C^r \vdash \mathbf{get_local} \ i : \epsilon \rightarrow \tau_r}^{\text{GET-LOCAL}} \\
 \\
 \frac{C^r_{\mathbf{local}}(i) = \tau_r^2 \quad \tau_r^1 <: \tau_r^2}{R; C^r \vdash \mathbf{set_local} \ i : \tau_r^1 \rightarrow \epsilon}^{\text{SET-LOCAL}} \qquad \frac{C^r_{\mathbf{local}}(i) = \tau_r^2 \quad \tau_r^1 <: \tau_r^2}{R; C^r \vdash \mathbf{tee_local} \ i : \tau_r^1 \rightarrow \tau_r^1}^{\text{TEE-LOCAL}} \\
 \\
 \frac{C^r_{\mathbf{global}}(i) = \tau_r}{R; C^r \vdash \mathbf{get_global} \ i : \epsilon \rightarrow \tau_r}^{\text{GET-GLOBAL}} \qquad \frac{C^r_{\mathbf{global}}(i) = \tau_r^2 \quad \tau_r^1 <: \tau_r^2}{R; C^r \vdash \mathbf{set_global} \ i : \tau_r^1 \rightarrow \epsilon}^{\text{SET-GLOBAL}} \\
 \\
 \frac{}{R; C^r \vdash \epsilon : \epsilon \rightarrow \epsilon}^{\text{EMPTYSTACK}} \qquad \frac{R; C^r \vdash e_1^* : \tau_r^a \rightarrow \tau_r^b \quad R; C^r \vdash e_2 : \tau_r^b \rightarrow \tau_r^c}{R; C^r \vdash e_1^* e_2 : \tau_r^a \rightarrow \tau_r^c}^{\text{SEQUENCING}} \\
 \\
 \frac{R; C^r \vdash e^* : \tau_r^b \rightarrow \tau_r^c}{R; C^r \vdash e^* : \tau_r^a \tau_r^b \rightarrow \tau_r^a \tau_r^c}^{\text{TOPOFSTACK}} \qquad \frac{}{R; C^r \vdash \mathbf{trap} : t_f^r}^{\text{TRAP}}
 \end{array}$$

Fig. 10. Typing Rules for the Refinement Type System continued.

B TYPE SAFETY PROOF

THEOREM 3. *Progress*: *If* $R; C^r \vdash_i \hat{e}^* : \tau_r^1 \cdots \tau_r^m \rightarrow \tau_r^* \wedge (v : \tau_r)^m \wedge \vdash_i s : R$ *then* $e^* = v^*$ *or* $e^* = \text{trap}$ *or* $s; v^*; e^* \hookrightarrow s'; v'^*; e'^*$

PROOF. Proof by induction of the derivation of $R; C^r \vdash_i \hat{e}^* : \tau_r^1 \cdots \tau_r^m \rightarrow \tau_r^*$

CASE LOAD:

$$\frac{C_{\text{memory}}^r = n \quad \tau_r^1 <: \text{ptr}(c) \quad \tau_r^2 = \text{load_and_extend}(c + o, \tau_w, (tp_sx)^?, R_\Psi) \quad 2^a \leq (|tp| <)^? |\tau_r^2| \quad (tp_sz)^? = \epsilon \vee \tau_r^2 = \text{im}}{R; C^r \vdash \tau.\text{load}(tp_sx)^? a o : \tau_r^1 \rightarrow \tau_r^2} \text{LOAD}$$

The load instruction has to take a step since only end instructions signify the end of a function. We now case on the shape of a well typed stack. The only possible case is that the load instruction has a $\text{ptr}(a)$ on the stack, since this is a premise of our typing rule. The operational semantics of the load instruction matches this stack shape. If the side conditions are not met, it steps to a trap. \square

THEOREM 4. *Preservation*: *If* $R; C^r \vdash_i \hat{e}^* : \tau_r^1 \cdots \tau_r^m \rightarrow \tau_r^* \wedge \vdash_i (v : \tau_r)^m \wedge \vdash_i s : R \wedge s, v^*, e^* \hookrightarrow s', v'^*, e'^*$ *then* $\exists R'$ *such that* $\text{dom}(R_\Psi) \subseteq \text{dom}(R'_\Psi) \wedge \vdash_i s' : R' \wedge R'; C^r \vdash_i e'^* : \tau_r^*$

PROOF. Proof by induction of the derivation of $R; C^r \vdash_i \hat{e}^* : \tau_r^1 \cdots \tau_r^m \rightarrow \tau_r^*$

CASE LOAD:

$$\frac{C_{\text{memory}}^r = n \quad \tau_r^1 <: \text{ptr}(c) \quad \tau_r^2 = \text{load_and_extend}(c + o, \tau_w, (tp_sx)^?, R_\Psi) \quad 2^a \leq (|tp| <)^? |\tau_r^2| \quad (tp_sz)^? = \epsilon \vee \tau_r^2 = \text{im}}{R; C^r \vdash \tau.\text{load}(tp_sx)^? a o : \tau_r^1 \rightarrow \tau_r^2} \text{LOAD}$$

From our typing rule for the load instruction we have that it expects τ_r^1 on the stack and pushes τ_r^2 on the stack. From the operational semantics of the load instruction we know that there are three possible cases:

CASE 1: $s; \text{ptr.const}(c); \tau.\text{load } a o \hookrightarrow \tau.\text{const}(b^*)$ if $s_{\text{mem}}(i, c + o, |\tau|) = b^*$

From $\vdash_i s : R$, we know that $\vdash s_{\text{mem}} : R_\Psi$, from which we know that $s_{\text{mem}}(c + o) : R_\Psi(c + o)$. Since our memory typing R_Ψ maps addresses to types of size 32, we have to now case on the data expected after the load instruction. This is annotated on the load instruction itself as τ .

CASE 1.1: $\tau.\text{load } a o \wedge R_\Psi[(c + o) \mapsto \tau_{r_{32}}(n)]$

This is the case where the data in memory is of the type that the load expects to produce on the stack. We now have to show that $b^* : \tau(n)$. From the typing rule, we have that load expects τ_r^2 . On inspecting the `load_and_extend` function, we see that for the case where the data in memory is the same type as is expected on the stack, $\tau_r^2 = \tau_r(n_1)$ where $\tau_r(n_1) = \Psi(c + o)$. Since $s_{\text{mem}}(c + o) : R_\Psi(c + o)$, we know that $s_{\text{mem}}(c + o) : \tau_r^2$ and since $s_{\text{mem}}(c + o) = b^*$, $b^* : \tau_r^2$.

CASE 1.2: $i64.\text{load } a o \wedge R_\Psi[(c + o) \mapsto \tau_{r_{32}}(n_1), (c + o - 4) \mapsto \tau_{r_{32}}(n_2)]$

In this case, the load expects a `i64` on the stack and the memory has $\tau_{r_{32}}(n_1)$ at $(c + o)$ and $\tau_{r_{32}}(n_2)$ at $(c + o - 4)$. We now have to show that $b^* : i64(n)$. Note that b^* is the byte sequence from $s_{\text{mem}}[c + o : 8]$, as specified in the detailed Wasm specification [47]. From the typing rule, we know that load expects τ_r^2 on the stack. On inspecting the `load_and_extend` function, we see that this case is the one that matches $|\tau_r(n_1)| < N$, where $N = |\tau| = 64$. We see then that $\tau_r^2 = \tau(\text{val}(\text{bit}(n_1), \text{bit}(n_2)))$. n_1 and n_2 are the values in memory at $c + o$ and $c + o - 4$. Since $s_{\text{mem}}(c + o) : R_\Psi(c + o) \wedge s_{\text{mem}}(c + o - 4) : R_\Psi(c + o - 4)$, we know that $s_{\text{mem}}(c + o : 8) : \tau_r^2$ and since $s_{\text{mem}}(c + o) = b^*$, $b^* : \tau_r^2$.

1471 CASE 1.3: $f64.load\ a\ o \wedge R_{\Psi}[(c + o) \mapsto \tau_{r_{32}}(n_1), (c + o - 4) \mapsto \tau_{r_{32}}(n_2)]$
 1472 In this case, the load expects a f64 on the stack and the memory has $\tau_{r_{32}}(n_1)$ at $(c + o)$ and $\tau_{r_{32}}(n_2)$
 1473 at $(c + o - 4)$. We now have to show that $b^* : f64(n)$. The proof proceeds exactly as the case above.
 1474

1475 CASE 1.4: $ptr.load\ a\ o \wedge R_{\Psi}[(c + o) \mapsto f32(n)]$

1476 CASE 1.5: $num.load\ a\ o \wedge R_{\Psi}[(c + o) \mapsto f32(n)]$

1477 CASE 1.6: $i32.load\ a\ o \wedge R_{\Psi}[(c + o) \mapsto f32(n)]$

1478 In this case, the load expects a i32 on the stack and Ψ has $f32(n)$ at $c + o$. We now have to show that
 1479 $b^* : f32(n)$. From the typing rule, we have that load expects τ_r^2 . On inspecting the `load_and_extend`
 1480 function, we see that for this case, $\tau_r^2 = f32(\text{bits}(n_1 \text{ as } f32))$ where $\tau_r(n_1) = \Psi(c + o)$. Since
 1481 $s_{\text{mem}}(c + o) : R_{\Psi}(c + o)$, we know that $s_{\text{mem}}(c + o) : \tau_r^2$ and since $s_{\text{mem}}(c + o) = b^*$, $b^* : \tau_r^2$.

1482 CASE 1.7: $f32.load\ a\ o \wedge R_{\Psi}[(c + o) \mapsto ptr(n)]$

1483 CASE 1.8: $f32.load\ a\ o \wedge R_{\Psi}[(c + o) \mapsto num(n)]$

1484 CASE 1.9: $f32.load\ a\ o \wedge R_{\Psi}[(c + o) \mapsto i32(n)]$

1485 In this case, the load expects a f32 on the stack and Ψ has $i32(n)$ at $c + o$. We now have to
 1486 show that $b^* : i32(n)$. From the typing rule, we have that load expects τ_r^2 . On inspecting the
 1487 `load_and_extend` function, we see that for this case, $\tau_r^2 = i32(f32(n_1) \text{ as } i32)$ where $\tau_r(n_1) = \Psi(c + o)$.
 1488 Since $s_{\text{mem}}(c + o) : R_{\Psi}(c + o)$, we know that $s_{\text{mem}}(c + o) : \tau_r^2$ and since $s_{\text{mem}}(c + o) = b^*$, $b^* : \tau_r^2$.
 1489

1490 **CASE 2:** $s; ptr.const(c); \tau.load\ tp_sx\ a\ o \hookrightarrow \tau.const(b^*)$ if $s_{\text{mem}}(i, c + o, |tp|) = b^*$

1491 From $\vdash_i s : R$, we know that $\vdash s_{\text{mem}} : R_{\Psi}$, from which we know that $s_{\text{mem}}(c + o) : R_{\Psi}(c + o)$. Since
 1492 our memory typing R_{Ψ} maps addresses to types of size 32, we have to now case on the data expected
 1493 after the load instruction. The proof proceeds as above except with a small change introduced by
 1494 the `tp_sx` annotation on the load. In the WebAssembly operational semantics [47], when a load
 1495 has this packed type annotation, the bits are read from memory upto $|tp|$ and the value read from
 1496 memory is size extended with the `extend_sxN,|\tau|` function. Both of these cases are handled in the
 1497 `load_and_extend` function, in case $|\tau_r(n_1)| > N$ and size extending function `extend_sxN,|\tau|`.
 1498

1499 **CASE 3:** $s; ptr.const(c); \tau.load\ (tp_sx)^? a\ o \hookrightarrow \text{trap}$ otherwise

1500 The trap instruction is well typed. \square

1501

1502 C MEET AND JOIN OPERATIONS OVER THE INTEGER SUB-LATTICE

1503

1504

1505

1506

1507

	\sqcup	$i32(\top)$	$i32(n)$	$ptr(\top, \top)$	$ptr(l_0, n_0)$	$ptr(l_1, n_1)$	$ptr(\perp, \perp)$	$num(\top)$	$num(n)$	$num(\perp)$	$i32(\perp)$
1508	$i32(\top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$
1509	$i32(n')$	$i32(\top)$	$i32(n \sqcup n')$	$i32(n')$	$i32(n' \sqcup n_0)$	$i32(\top)$	$i32(n')$	$i32(\top)$	$i32(n' \sqcup n)$	$i32(n')$	$i32(n')$
1510	$ptr(\top, \top)$	$i32(\top)$	$i32(n)$	$ptr(\top, \top)$	$ptr(\top, \top)$	$ptr(\top, \top)$	$ptr(\top, \top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$ptr(\top, \top)$
1511	$ptr(l_0, n'_0)$	$i32(\top)$	$i32(n \sqcup n'_0)$	$ptr(\top, \top)$	$ptr(l_0, n'_0 \sqcup n_0)$	$ptr(\top, \top)$	$ptr(l_0, n'_0)$	$i32(\top)$	$i32(\top)$	$i32(n'_0)$	$ptr(l_0, n'_0)$
1512	$ptr(l_1, n'_1)$	$i32(\top)$	$i32(\top)$	$ptr(\top, \top)$	$ptr(\top, \top)$	$ptr(l_1, n_1 \sqcup n'_1)$	$ptr(l_1, n'_1)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$ptr(l_1, n'_1)$
1513	$ptr(l_2, n'_2)$	$i32(\top)$	$i32(\top)$	$ptr(\top, \top)$	$ptr(\top, \top)$	$ptr(\top, \top)$	$ptr(l_2, n'_2)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$ptr(l_2, n'_2)$
1514	$ptr(\perp, \perp)$	$i32(\top)$	$i32(\top)$	$ptr(\top, \top)$	$ptr(l_0, n_0)$	$ptr(l_1, n_1)$	$ptr(\perp, \perp)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$ptr(\perp, \perp)$
1515	$num(\top)$	$i32(\top)$	$i32(n)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$i32(\top)$	$num(\top)$	$num(\top)$	$num(\top)$	$num(\top)$
1516	$num(n')$	$i32(\top)$	$i32(n \sqcup n')$	$i32(\top)$	$i32(n_0 \sqcup n')$	$i32(\top)$	$i32(\top)$	$num(n')$	$num(n' \sqcup n)$	$num(n')$	$num(n')$
1517	$num(\perp)$	$i32(\top)$	$i32(n)$	$i32(\top)$	$i32(n_0)$	$i32(\top)$	$i32(\top)$	$num(\top)$	$num(n)$	$num(\perp)$	$num(\perp)$
1518	$i32(\perp)$	$i32(\top)$	$i32(n)$	$ptr(\top, \top)$	$ptr(l_0, n_0)$	$ptr(l_1, n_1)$	$ptr(\perp, \perp)$	$num(\top)$	$num(n)$	$num(\perp)$	$i32(\perp)$

Table 5. Join \sqcup operation for the i32 sub-lattice.

1518

1519

	\sqcap	i32(\top)	i32(n)	ptr(\top, \top)	ptr(l_0, n_0)	ptr(l_1, n_1)	ptr(\perp, \perp)	num(\top)	num(n)	num(\perp)	i32(\perp)
1520	i32(\top)	i32(\top)	i32(n)	ptr(\top, \top)	ptr(l_0, n_0)	ptr(l_1, n_1)	ptr(\perp, \perp)	num(\top)	num(n)	num(\perp)	i32(\perp)
1521	i32(n')	i32(n')	i32($n \sqcap n'$)	ptr(\top, \top)	ptr($l_0, n_0 \sqcap n'_0$)	ptr(\perp, \perp)	ptr(\perp, \perp)	num(n')	num($n \sqcap n'$)	num(\perp)	i32(\perp)
1522	ptr(\top, \top)	ptr(\top, \top)	ptr(\top, \top)	ptr(\top, \top)	ptr(l_0, n_0)	ptr(l_1, n_1)	ptr(\perp, \perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)
1523	ptr(l_0, n'_0)	ptr(l_0, n'_0)	ptr($l_0, n'_0 \sqcap n$)	ptr(l_0, n'_0)	ptr($l_0, n_0 \sqcap n'_0$)	ptr(\perp, \perp)	ptr(\perp, \perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)
1524	ptr(l_1, n'_1)	ptr(l_1, n'_1)	ptr(\perp, \perp)	ptr(l_1, n'_1)	ptr(\perp, \perp)	ptr($l_1, n_1 \sqcap n'_1$)	ptr(\perp, \perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)
1525	ptr(l_2, n'_2)	ptr(l_2, n'_2)	ptr(\perp, \perp)	ptr(l_2, n'_2)	ptr(\perp, \perp)	ptr($l_3, 0$)	ptr(\perp, \perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)
1526	ptr(\perp, \perp)	ptr(\perp, \perp)	ptr(\perp, \perp)	ptr(\perp, \perp)	ptr(\perp, \perp)	ptr(\perp, \perp)	ptr(\perp, \perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)
1527	num(\top)	num(\top)	num(\top)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	num(\top)	num(n)	num(\perp)	i32(\perp)
1528	num(n')	num(n')	num($n \sqcap n'$)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	num(n')	num($n \sqcap n'$)	num(\perp)	i32(\perp)
1529	num(\perp)	num(\perp)	num(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	num(\perp)	num(\perp)	num(\perp)	i32(\perp)
1530	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)	i32(\perp)

$$*l_3 \text{ fresh} \wedge l_1 = l_3 - n_1 \wedge l_2 = l_3 - n'_2$$

Table 6. Meet \sqcap operation for the i32 sub-lattice.

D CONSTRAINT GENERATION FOR ALL WEBASSEMBLY INSTRUCTIONS

context $C^\alpha ::= \{\text{func } tf^*, \text{ local } \alpha^*, \text{ global } \alpha^*, \text{ table } n^?, \text{ memory } n^?, \text{ label } (\alpha^*, \zeta)^*, \text{ return } (\alpha^*, \zeta)^*\}$

Constraints for Instructions

$$S; \zeta; C^\alpha \vdash e : \alpha^* \rightarrow \alpha^*; S'; \zeta'$$

$$\frac{\alpha \text{ fresh} \quad S' = S :: [\alpha \doteq \tau_w(c)]}{S; \zeta; C^\alpha \vdash \tau_w.\text{const } c : \epsilon \rightarrow \alpha; S'; \zeta} \text{CONSTANT} \quad \frac{\alpha_1 \in \text{dom}(S) \quad \alpha_2 \text{ fresh} \quad S' = S :: [\alpha_1 \doteq \text{unop}(\alpha_2)]}{S; \zeta; C^\alpha \vdash \tau_w.\text{unop} : \alpha_1 \rightarrow \alpha_2; S'; \zeta} \text{UNARY-OPS}$$

$$\frac{\alpha_1 \alpha_2 \in \text{dom}(S) \quad \alpha_3 \text{ fresh} \quad S' = S :: [\alpha_1 \doteq \text{binop}(_, \alpha_2, \alpha_3) \wedge \alpha_2 \doteq \text{binop}(\alpha_1, _, \alpha_3) \wedge \alpha_3 \doteq \text{binop}(\alpha_1, \alpha_2, _)]}{S; \zeta; C^\alpha \vdash \tau_w.\text{binop} : \alpha_1 \alpha_2 \rightarrow \alpha_3; S'; \zeta} \text{BINARY-OPS}$$

$$\frac{\alpha_1 \alpha_2 \in \text{dom}(S) \quad \alpha_3 \text{ fresh} \quad S' = S :: [\alpha_1 \doteq \text{testop}(_, \alpha_2, \alpha_3) \wedge \alpha_2 \doteq \text{testop}(\alpha_1, _, \alpha_3) \wedge \alpha_3 \doteq \text{testop}(\alpha_1, \alpha_2, _)]}{S; \zeta; C^\alpha \vdash \tau_w.\text{testop} : \alpha_1 \alpha_2 \rightarrow \alpha_3; S'; \zeta} \text{TEST-OPS}$$

$$\frac{\alpha_1 \alpha_2 \in \text{dom}(S) \quad \alpha_3 \text{ fresh} \quad S' = S :: [\alpha_1 \doteq \text{relop}(_, \alpha_2, \alpha_3) \wedge \alpha_2 \doteq \text{relop}(\alpha_1, _, \alpha_3) \wedge \alpha_3 \doteq \text{relop}(\alpha_1, \alpha_2, _)]}{S; \zeta; C^\alpha \vdash \tau_w.\text{relop} : \alpha_1 \alpha_2 \rightarrow \alpha_3; S'; \zeta} \text{REL-OPS}$$

$$\frac{\alpha_1 \in \text{dom}(S) \quad \alpha_2 \text{ fresh} \quad S' = S :: [\alpha_1 \doteq \text{ctop } \tau_w.\text{sx}^2(_, \alpha_2) \wedge \alpha_2 \doteq \text{ctop } \tau_w.\text{sx}^2(\alpha_1, _)]}{S; \zeta; C^\alpha \vdash \tau_w.\text{ctop } \tau_w.\text{sx}^2 : \alpha_1 \rightarrow \alpha_2; S'; \zeta} \text{CONVERT-OPS}$$

$$\frac{\alpha_1 \in \text{dom}(S) \quad S' = S :: [\alpha_1 <: \text{ptr} \wedge \alpha_2 \doteq \zeta[\alpha_1](tf.\text{sx}^2, a, o) \wedge \alpha_2 <: \tau_w]}{S; \zeta; C^\alpha \vdash \tau_w.\text{load}(tp.\text{sx})^? a o : \alpha_1 \rightarrow \alpha_2; S'; \zeta} \text{LOAD}$$

$$\frac{\alpha_1 \alpha_2 \in \text{dom}(S) \quad S' = S :: [\alpha_1 <: \text{ptr} \wedge \alpha_2 <: \tau_w \wedge \zeta' \doteq \zeta[\alpha_1 \mapsto \alpha_2](tp^?, a, o)]}{S; \zeta; C^\alpha \vdash \tau_w.\text{store } tp^? a o : \alpha_1 \alpha_2 \rightarrow \epsilon; S'; \zeta'} \text{STORE}$$

Fig. 11. Constraint Generation Rules.

$$\begin{array}{c}
 1569 \\
 1570 \\
 1571 \\
 1572 \\
 1573 \\
 1574 \\
 1575 \\
 1576 \\
 1577 \\
 1578 \\
 1579 \\
 1580 \\
 1581 \\
 1582 \\
 1583 \\
 1584 \\
 1585 \\
 1586 \\
 1587 \\
 1588 \\
 1589 \\
 1590 \\
 1591 \\
 1592 \\
 1593 \\
 1594 \\
 1595 \\
 1596 \\
 1597 \\
 1598 \\
 1599 \\
 1600 \\
 1601 \\
 1602 \\
 1603 \\
 1604 \\
 1605 \\
 1606 \\
 1607 \\
 1608 \\
 1609 \\
 1610 \\
 1611 \\
 1612 \\
 1613 \\
 1614 \\
 1615 \\
 1616 \\
 1617
 \end{array}$$

$$\begin{array}{c}
 \frac{tf_w = \tau_w^m \rightarrow \tau_w^n \quad \alpha^m \in \text{dom}(S) \quad \alpha^n \text{ fresh}}{S; \zeta; C^\alpha, \text{label}(\alpha^n) \vdash e^* : \alpha^m \rightarrow (\alpha')^n; S'; \zeta' \quad S'' = S' :: [(\alpha \doteq \alpha')^n]}{S; \zeta; C^\alpha \vdash \mathbf{block} \text{ } tf \ e^* \ \mathbf{end} : \alpha_1^n \dots \alpha_n^n \rightarrow \alpha_1^m \dots \alpha_m^m; S''; \zeta'} \text{BLOCK} \\
 \\
 \frac{tf_w = \tau_w^m \rightarrow \tau_w^n \quad \alpha^m \in \text{dom}(S) \quad \alpha^n \text{ fresh}}{S; \zeta; C^\alpha, \text{label}(\alpha^m) \vdash e^* : \alpha^m \rightarrow (\alpha')^n; S'; \zeta' \quad S'' = S' :: [(\alpha \doteq \alpha')^n]}{S; \zeta; C^\alpha \vdash \mathbf{loop} \text{ } tf_w \ e^* \ \mathbf{end} : \alpha^m \rightarrow \alpha^n; S''; \zeta'} \text{LOOP} \\
 \\
 \frac{\begin{array}{c} tf_w = \tau_w^m \rightarrow \tau_w^n \quad \alpha^m \alpha \in \text{dom}(S) \quad \alpha_t^n, \alpha_e^n, \alpha^n \text{ fresh} \quad S' = S :: [\alpha <: \text{num}] \\ S'; \zeta; C^\alpha, \text{label}(\alpha_t^n) \vdash e_t^* : \alpha_m \rightarrow \alpha_t^n; S_e; \zeta_e \quad S'; \zeta; C^\alpha, \text{label}(\alpha_e^n) \vdash e_e^* : \alpha_m \rightarrow \alpha_e^n; S_t; \zeta_t \\ S'' = S' :: S_e :: S_t :: [\zeta' \doteq \bigsqcup \zeta_e \zeta_t \wedge (\alpha' \doteq \bigsqcup \alpha_t \alpha_e)^n] \end{array}}{S; \zeta; C^\alpha \vdash \mathbf{if} \text{ } tf \ e_t^* \ \mathbf{else} \ e_e^* \ \mathbf{end} : \alpha^m \alpha \rightarrow \alpha^n; S''; \zeta'} \text{IF-ELSE} \\
 \\
 \frac{\alpha^* \alpha^n \in \text{dom}(S) \quad C_{\text{label}}^\alpha(i) = (\alpha')^n, \zeta' \quad S' = S :: [\zeta \doteq \bigsqcup \zeta \zeta' \wedge (\alpha \doteq \bigsqcup \alpha \alpha')^n]}{S; \zeta; C^\alpha \vdash \mathbf{br} \ i : \alpha^* \alpha_1 \dots \alpha_n \rightarrow \alpha^*; S'; \zeta} \text{BR} \\
 \\
 \frac{\alpha^m \alpha \in \text{dom}(S) \quad C_{\text{label}}^\alpha(i) = (\alpha')^n, \zeta' \quad S' = S :: [\zeta \doteq \bigsqcup \zeta \zeta' \wedge \alpha <: \text{num} \wedge (\alpha \doteq \bigsqcup \alpha \alpha')^n]}{S; \zeta; C^\alpha \vdash \mathbf{br_if} \ i : \alpha^m \alpha \rightarrow \alpha^n; S'; \zeta} \text{BR-IF} \\
 \\
 \frac{\begin{array}{c} \alpha^* \alpha^m \alpha \in \text{dom}(S) \quad (C_{\text{label}}^\alpha(i) = (\alpha')^n)^+, \zeta' \\ S' = S :: ([\zeta \doteq \bigsqcup \zeta \zeta' \wedge \alpha_{n+1} <: \text{num} \wedge \alpha_1 \doteq \bigsqcup \alpha_1 \alpha'_1 \wedge \dots \wedge \alpha_n \doteq \bigsqcup \alpha_n \alpha'_n])^+ \end{array}}{S; \zeta; C^\alpha \vdash \mathbf{br_table} \ i^+ : \alpha^* \alpha_1 \dots \alpha_n \alpha_{n+1} \rightarrow \alpha^*; S'; \zeta} \text{BR-TABLE} \\
 \\
 \frac{\begin{array}{c} \alpha^* \alpha_1 \dots \alpha_n \in \text{dom}(S) \quad C_{\text{return}}(i) = \alpha'_1 \dots \alpha'_n \\ S' = S :: [\alpha_1 \doteq \bigsqcup \alpha_1 \alpha'_1 \wedge \dots \wedge \alpha_n \doteq \bigsqcup \alpha_n \alpha'_n] \end{array}}{S; \zeta; C \vdash \mathbf{return} : \alpha^* \alpha_1 \dots \alpha_n \rightarrow \alpha^*; S'; \zeta} \\
 \\
 \frac{\begin{array}{c} C_{\text{func}} = tf \quad tf = \tau_w^1 \dots \tau_w^n \rightarrow \tau_w^1 \dots \tau_w^m \\ \alpha_1 \dots \alpha_n \in \text{dom}(S) \quad S' = S :: [\alpha'_1 <: \tau_1 \wedge \dots \wedge \alpha'_m <: \tau_m] \end{array}}{S; \zeta; C \vdash \mathbf{call} : \alpha_1 \dots \alpha_n \rightarrow \alpha'_1 \dots \alpha'_m; S'; \zeta} \\
 \\
 \frac{\begin{array}{c} C_{\text{table}} = n \quad tf = \tau_w^1 \dots \tau_w^n \rightarrow \tau_w^1 \dots \tau_w^m \quad \alpha_1 \dots \alpha_n \alpha_{n+1} \in \text{dom}(S) \\ S' = S :: [\alpha_{n+1} <: \text{num} \wedge \alpha'_1 <: \tau_1 \wedge \dots \wedge \alpha'_m <: \tau_m] \end{array}}{S; \zeta; C \vdash \mathbf{call_indirect} \text{ } tf : \alpha_1 \dots \alpha_n \alpha_{n+1} \rightarrow \alpha'_1 \dots \alpha'_m; S'; \zeta} \\
 \\
 \frac{\alpha \text{ fresh} \quad S' = S :: [\alpha <: \text{num}]}{S; \zeta; C^\alpha \vdash \mathbf{current_memory} : \epsilon \rightarrow \alpha; S'; \zeta} \text{CURRENT-MEMORY} \\
 \\
 \frac{\alpha_1 \in \text{dom}(S) \quad \alpha_2 \text{ fresh} \quad S' = S :: [\alpha_1 <: \text{num} \wedge \alpha_2 <: \text{num}]}{S; \zeta; C^\alpha \vdash \mathbf{grow_memory} : \alpha_1 \rightarrow \alpha_2; S'; \zeta} \text{GROW-MEMORY} \\
 \\
 \frac{\alpha_1 \ \alpha_2; \alpha_3 \in \text{dom}(S) \quad \alpha_4 \text{ fresh} \quad S' = S :: [\alpha_3 <: \text{num} \wedge \alpha_4 \doteq \bigsqcup \alpha_1 \alpha_2]}{S; \zeta; C^\alpha \vdash \mathbf{select} : \alpha_1 \ \alpha_2 \ \alpha_3 \rightarrow \alpha_4; S'; \zeta} \text{SELECT}
 \end{array}$$

Fig. 12. Constraint Generation Rules cont.

$$\begin{array}{c}
1618 \\
1619 \\
1620 \\
1621 \\
1622 \\
1623 \\
1624 \\
1625 \\
1626 \\
1627 \\
1628 \\
1629 \\
1630 \\
1631 \\
1632 \\
1633 \\
1634 \\
1635 \\
1636 \\
1637 \\
1638 \\
1639 \\
1640 \\
1641 \\
1642 \\
1643 \\
1644 \\
1645 \\
1646 \\
1647 \\
1648 \\
1649 \\
1650 \\
1651 \\
1652 \\
1653 \\
1654 \\
1655 \\
1656 \\
1657 \\
1658 \\
1659 \\
1660 \\
1661 \\
1662 \\
1663 \\
1664 \\
1665 \\
1666
\end{array}$$

$$\begin{array}{c}
\frac{\alpha_1^* \in \text{dom}(S) \quad \alpha_2^* \text{ fresh}}{S; \zeta; C^\alpha \vdash \text{unreachable} : \alpha_1^* \rightarrow \alpha_2^*; S; \zeta} \text{UNREACHABLE} \qquad \frac{}{S; \zeta; C \vdash \text{nop} : \epsilon \rightarrow \epsilon; S; \zeta} \text{NOP} \\
\\
\frac{\alpha \in \text{dom}(S)}{S; \zeta; C \vdash \text{drop} : \alpha \rightarrow \epsilon; S; \zeta} \text{DROP} \qquad \frac{\alpha \text{ fresh} \quad S' = S :: [\alpha \doteq C_{\text{local}}(i)]}{S; \zeta; C \vdash \text{get_local } i : \epsilon \rightarrow \alpha; S'; \zeta} \text{GET-LOCAL} \\
\\
\frac{\alpha \in \text{dom}(S) \quad C_{\text{local}}(i) = \alpha' \quad S' = S :: [\alpha' \doteq \alpha]}{S; \zeta; C \vdash \text{set_local } i : \alpha \rightarrow \epsilon; S'; \zeta} \text{SET-LOCAL} \qquad \frac{\alpha \in \text{dom}(S) \quad C_{\text{local}}(i) = \alpha' \quad S' = S :: [\alpha' \doteq \alpha]}{S; \zeta; C \vdash \text{tee_local } i : \alpha \rightarrow \alpha; S'; \zeta} \text{TEE-LOCAL} \\
\\
\frac{\alpha \text{ fresh} \quad S' = S :: [\alpha \doteq C_{\text{global}}(i)]}{S; \zeta; C \vdash \text{get_global } i : \epsilon \rightarrow \alpha; S'; \zeta} \text{GET-GLOBAL} \qquad \frac{\alpha \in \text{dom}(S) \quad C_{\text{global}}(i) = \alpha' \quad S' = S :: [\alpha' \doteq \alpha]}{S; \zeta; C \vdash \text{set_global } i : \alpha \rightarrow \epsilon; S'; \zeta} \text{SET-GLOBAL} \\
\\
\frac{tf = \tau_w^1 \dots \tau_w^n \rightarrow \tau_w^1 \dots \tau_w^m \quad \tau^f = \alpha_{n_1} \dots \alpha_{n_n} \rightarrow \alpha_{m_1} \dots \alpha_{m_m} \quad \alpha_{n_1} \dots \alpha_{n_n}, \alpha_{m_1} \dots \alpha_{m_m}, \alpha_{l_1} \dots \alpha_{l_l} \text{ fresh}}{S = [\zeta \doteq \cdot \wedge \alpha_{n_1} <: \tau_w^1 \wedge \dots \wedge \alpha_{n_n} <: \tau_w^n \wedge \alpha_{m_1} <: \tau_w^1 \wedge \dots \wedge \alpha_{m_m} <: \tau_w^m \wedge \alpha_{l_1} <: \tau_w^1 \wedge \dots \wedge \alpha_{l_l} <: \tau_w^l]} \\
\frac{S; \zeta; C, \text{local}(\alpha_{n_1}, \dots, \alpha_{n_n}, \alpha_{l_1}, \dots, \alpha_{l_l}), \text{label}(\alpha_{m_1}, \dots, \alpha_{m_m}), \text{return}(\alpha_{m_1}, \dots, \alpha_{m_m}) \vdash e^* : \tau^f; S'; \zeta'}{[]; \zeta; C \vdash \text{ex}^* \text{ func } tf \text{ local } \tau_w^1 \dots \tau_w^l e^* : \tau^f; S'; \zeta'}
\end{array}$$

Fig. 13. Constraint Generation Rules cont.