

A Typed Intermediate Representation for Dynamic Languages

MICKAËL LAURENT, Charles University, Czech Republic

JAKOB HAIN, Purdue University, United States

FILIP KŘÍKAVA, Czech Technical University, Czech Republic

SEBASTIÁN KRYNSKI, Czech Technical University, Czech Republic

JAN VITEK, Charles University, Czech Republic and Czech Technical University, Czech Republic

Dynamic programming languages pose significant challenges for optimizing compilers due to features such as dynamic typing, late binding, reflection, copy-on-write, and delayed evaluation. To generate efficient code, compilers must speculate on which dynamic features will be exercised and produce specialized code based on these assumptions. This paper presents the design of a statically typed, high-level intermediate representation that makes dynamic behaviors explicit and amenable to static analysis. Our IR combines gradual typing with ownership tracking, and explicitly represents promises, multiple function versions, and contextual dispatch. Together, these features directly support optimizations such as specialization, inlining, scope elision, and copy elimination. We formalize a core calculus, called FIR, that captures the essential features required for these optimizations. We provide an operational semantics, a type system, and flow and reflection analyses, and we prove the soundness of the type system.

CCS Concepts: • **Software and its engineering** → **Just-in-time compilers**.

Additional Key Words and Phrases: JIT, compilation, gradual typing, dynamic language

ACM Reference Format:

Mickaël Laurent, Jakob Hain, Filip Kříkava, Sebastián Krynski, and Jan Vitek. 2026. A Typed Intermediate Representation for Dynamic Languages. *ACM Trans. Program. Lang. Syst.* 1, 1 (January 2026), 33 pages. <https://doi.org/10.1145/3811921>

1 Introduction

Optimizing compilers typically use multiple intermediate representations (IRs) to organize the compilation process into modular stages. Each IR highlights specific aspects of programs to enable targeted optimization passes. High-level IRs remain close to the source language, facilitating transformations guided by source-level semantics and supporting the verification of language-level properties and invariants. In contrast, low-level IRs are closer to the target machine, exposing control flow and resource management details required for optimizations such as instruction scheduling and register allocation. This paper focuses on the design of a typed, high-level IR tailored to the needs of just-in-time compilation of dynamic programming languages.

Dynamic languages—such as Python, PHP, or Julia—pose significant challenges for traditional compilation techniques. Features like dynamic typing, late binding, and reflection often prevent the compiler from reliably inferring program behavior through source code analysis alone.

Authors' Contact Information: Mickaël Laurent, mickael.laurent@mffyz.cuni.cz, Charles University, Prague, Czech Republic; Jakob Hain, jakobeha@fastmail.com, Purdue University, West Lafayette, Indiana, United States; Filip Kříkava, filip.krikava@fit.cvut.cz, Czech Technical University, Prague, Czech Republic; Sebastián Krynski, skrynski@gmail.com, Czech Technical University, Prague, Czech Republic; Jan Vitek, vitekj@me.com, Charles University, Prague, Czech Republic and Czech Technical University, Prague, Czech Republic.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

ACM 1558-4593/2026/1-ART

<https://doi.org/10.1145/3811921>

(a) => a + 1	lambda a: a + 1	(lambda (a)(+ a 1))	function(a)a + 1
JavaScript	Python	Scheme	R

Fig. 1. Implementations of “add 1” in four dynamic languages

Consider the archetypal function that takes its argument, adds 1, and returns it (Fig. 1). Even this function is not simple in a dynamic language, and generating efficient code for it is not easy. Though the intended use may be to add 1 to a number, in JavaScript, if the argument *a* is a string, the function actually performs string concatenation. In Python, the `+` operator may be redefined, making the function run arbitrary code. In Scheme, the `lambda` macro may be overridden, changing the entire expression. In R, not only can `+` and `function` be redefined before the expression is evaluated, but *a* may be a delayed thunk that runs arbitrary code as the function is being called.

This example illustrates that in dynamic languages, even simple code can elude static reasoning, rendering most standard compiler optimizations unsound or inapplicable. The key insight is that a function’s behavior may depend on non-local properties of the program—such as whether operators have been redefined or whether functions perform reflection.

Building on this observation, just-in-time (JIT) compilers adopt a speculative strategy: they use observations of past runtime behavior (*feedback*) to predict future behavior, then generate code optimized for the predictions. The predictions for dynamic languages effectively replace the guarantees of static languages: for example, predicting that a variable will always hold a value of a particular runtime type, or that a block of code will not perform reflection. Optimizations that rely on those predictions include, respectively, scalar unboxing and code motion. Since predictions are not guarantees, the compiler inserts before any optimizations that rely on them *guards*—runtime checks that, on failure, divert control flow to the original unoptimized code.

In a compiler, design of the high-level IR is critical, as it determines which optimizations can be expressed and verified. A representative example is the Julia compiler, as described in [3]. According to the authors, the two most impactful optimizations are call inlining and function specialization, as these enable nearly all subsequent optimizations. Accordingly, the compiler is structured to apply these early transformations on a high-level IR, after which lower-level optimizations are delegated to LLVM, a separate compiler with its own IR.

What, then, should a high-level IR for a dynamic language include? As explained above, dynamism renders most optimizations impossible, but a compiler can use guarded predictions to create regions where it is guaranteed absent. The IR must make dynamism—and its absence—explicit, to enable said regions to be expressed, thus enabling optimizations typical for a static compiler. The precise set of features varies by language, but the goal is the same: convert properties that are resolved at runtime into ones that can be reasoned about statically.

Our thesis is that these properties can—and should—be captured in a static type system. Doing so serves two purposes: it enables compile-time verification that optimizations are applied correctly, and it provides a formal specification for compiler writers.

Our work builds on prior results from the development of the $\check{R}$ compiler for the R programming language and its intermediate representation, PIR. We introduce $\check{\text{FIR}}$, a typed intermediate representation designed to support dynamic typing, runtime code loading, delayed evaluation, copy-on-write semantics, and unrestricted reflection.

The key features of $\check{\text{FIR}}$ (Section 4) address these features at two levels—restricting sources of dynamism and enabling targeted optimizations:

- **Variables:** $\check{\text{FIR}}$ distinguishes *named* from *register* variables; the former can be updated non-locally through reflection, while the latter are shielded from it.

- **Types:** Inspired by gradual typing, FiR supports a stratified type system in which variables range from statically typed to entirely dynamic.
- **Ownership:** A limited form of ownership tracks value lifetimes and transfer semantics, enabling optimization of copy-on-write behavior.
- **Specialization and Dispatch:** Functions can be compiled into multiple specialized versions; calls are dispatched to an applicable version at runtime or bound statically.
- **Inlining:** Function inlining exposes optimization opportunities while preserving the scopes of named variables.
- **Promises:** Delayed computations are represented as promises, supporting optimizations that target lazy evaluation.

Although motivated by R, these features are common across dynamic languages (Section 2), and we argue that FiR’s design is broadly applicable. Moreover, while designed for a JIT compiler, none of its features depend on runtime feedback, so it can equally serve ahead-of-time compilers.

FiR is presented as a minimalistic calculus, inspired by Featherweight Java [15]. Validity is guaranteed by a combination of type checking, to ensure that variables and functions are used according to their declarations, and flow analysis, to ensure that variables are initialized before use (Section 5.1). FiR’s reduced complexity enables us to provide a formal operational semantics (Section 5.2) and to prove the soundness of its type system (Section 5.3).

2 Dynamic Languages

Dynamic languages are characterized by idiosyncratic designs; the particular set of features supported in any given language is driven by the needs of the application domain and the whims of the language designer. Table 1 describes which features are provided by six popular languages.

Feature	Julia	JavaScript	Python	PHP	Clojure	R
Dynamic typing	✓	✓	✓	✓	✓	✓
Code loading	✓	✓	✓	✓	✓	✓
Late binding	✓	✓	✓	✓	✓	✓
Reflection	✓	✓	✓	✓	✓	✓
Copy-on-write	—	—	—	✓	—	✓
Delayed evaluation	—	—	—	—	✓	✓

Table 1. Support for dynamic language features across six popular languages

Dynamic Typing. In dynamically typed languages, variables do not have fixed types; instead, type information is associated with values. As a result, a variable can be rebound to values of different types over the course of execution. This flexibility is a hallmark of all the languages discussed above.

Code Loading. In these languages, all code outside of the standard library is loaded during (as opposed to before) execution. New global variables and functions may be added, and most existing ones may be rebound or modified. Many contain “eval”, a function that takes an arbitrary string as input, and parses and then evaluates it.

Late Binding. Late binding refers to the deferral of method resolution until runtime, when the types (or values) of arguments are known. In object-oriented languages, this is typically achieved through dynamic method dispatch; in functional languages, through higher-order functions. Additionally, many allow built-in functions—such as + —to be redefined or overloaded.

Reflection. Reflection is the ability of a program to inspect and modify its own structure and behavior at runtime. All the languages discussed support reflection to varying degrees. Julia is the most restricted, as it allows introspection but not reflective updates. R, on the other hand, is the most permissive—it supports call stack introspection and even allows local variables to be deleted dynamically.

Copy-on-Write. Copy-on-write is a memory management optimization in which multiple references share the same data until one of them attempts a modification, at which point a copy is made. This technique can be used to simulate referential transparency—for example, allowing functions to update arguments internally without those changes being externally visible.

Delayed Evaluation. In R, function arguments are lazily evaluated by default. When a function is called, its arguments are not immediately computed but instead represented as promises—deferred computations that are evaluated only when their values are needed. Clojure supports a similar, though explicit, mechanism via the `delay` construct, which postpones evaluation until the associated value is forced.

To illustrate the challenges compilers face when targeting such languages, consider the following reflective code snippet in R. This example demonstrates several language features that complicate optimization, and highlights the subtleties that arise from their interaction:

```
f = function(p, v) { a=1; p; print(a); v[1]='b' }
g = function(k, v) { assign(k, v, env=sys.frame(-1)) }
v = c(2, 3) # Creation of a vector
f(g('a', 666), v)
print(v[1])
```

When executed, this program prints 666 and 2. The first printed value results from evaluating the argument `p` in function `f`. As arguments are lazily evaluated, `p` is actually a promise containing the code `g('a', 666)`. The expression `p` in the body of `f` forces its evaluation, triggering a call to `g` that uses reflection to update the local variable `a` in the caller's environment. The second value, 2, is printed because although `f` appears to mutate vector `v`, R's copy-on-write semantics ensure that this mutation applies to a duplicate. Additionally, the assignment of the string 'b' into a numeric vector causes implicit coercion of the entire vector to character type.

Each feature presents challenges to a compiler; however, it is often their interaction that makes reasoning and optimization especially difficult. Consider the following interdependencies:

- **Dynamic typing** $\Rightarrow$ **late binding**: Without static type information, it is often impossible to determine which function is invoked at a call site, shifting many decisions to runtime.
- **Late binding** $\Rightarrow$ **reflection**: When a callee is not known, the compiler conservatively assumes that the target function performs reflection—even if such behavior is rare.
- **Delayed evaluation** $\Rightarrow$ **reflection**: Promises delay evaluation; the compiler often cannot predict if simply reading a variable will cause evaluation of reflective operations.
- **Code loading** $\Rightarrow$ **late binding**: The ability to load code at runtime limits opportunities for static binding, as new definitions may change the program's behavior.

In summary, efficient compilation hinges on the compiler's ability to rule out—or at least localize—these dynamic behaviors. Dynamism may originate from the calling context, from function arguments, or from callees and code loaded during execution.

We choose to focus on the R language precisely because it represents a worst-case scenario. Among popular dynamic languages, R combines extensive use of reflection, lazy evaluation, dynamic typing, and mutable environments, making it a particularly demanding target for optimizing compilers.

3 Related Work

3.1 Compiler IRs

Aycock provides an accessible overview of early implementation strategies for dynamic languages [1], covering systems ranging from Smalltalk [7] and Self [4] to more modern environments such as Java and JavaScript [21, 25].

The idea of formally specifying and verifying compiler IRs is well established, though mostly for static languages. Java bytecode is both formally specified and verified [12, 17]; its bytecode verifier—a data-flow analysis ensuring variables are initialized before use—is analogous to FiR’s well-flowedness check, though Java’s static typing makes this substantially simpler. Typed Assembly Language demonstrated that static types can preserve invariants even through low-level compilation [19]. FiR extends this idea to a setting where the source language lacks static types.

Other IR designs address different aspects of compilation. The sea-of-nodes representation makes all dependencies explicit in control and data flow graphs [5], enabling optimizations such as duplication removal and promise elimination, but operates at a lower level than FiR. CacheIR [6] targets JavaScript JIT compilation, organized around inline caching as a single optimization strategy; FiR targets a broader set of concerns, capturing multiple sources of dynamism in a type system. MLIR offers a general-purpose, extensible IR framework [16], but does not itself provide the type-level features FiR needs—gradual types, ownership, and reflection tracking—and formalizing type safety would require building these from scratch within its infrastructure.

3.2 Copy-on-Write, Ownership and Borrowing

Languages such as Swift and PHP implement copy-on-write semantics. Swift includes explicit borrowing and consuming annotations [13]: borrowing a value does not increment its reference count, while consuming a value transfers ownership to the callee. These correspond to FiR’s borrowed and owned parameter modes, respectively. PHP provides operations for assignment and aliasing; prior efforts in compilation have focused on distinguishing between them [14], but not on eliminating redundant copies. In FiR, aliasing cannot occur on owned variables: a variable that is aliased is considered shared, and copies will not be statically eliminated for it.

More broadly, ownership types constrain the use of aliases to heap-allocated objects [20], and borrowing relaxes strict ownership invariants [23]. Rust enforces these ideas via its borrow checker [18]. Immutable borrows in Rust correspond to FiR’s borrowing semantics: they are temporary and can reference owned values, which remain valid after the borrow expires. Rust’s smart pointers resemble FiR’s shared values in that they do not enforce liveness, and any value assigned to them becomes immutable. Ownership in R is more complex than in any of these languages, due to lazy evaluation and first-class mutable environments. To handle this, FiR introduces an explicit **consume** instruction along with a flow-sensitive analysis to detect potential read-after-consume violations.

3.3 Gradual Typing

Gradual typing allows statically typed and dynamically typed code to coexist within the same language [22]. FiR adopts the three-level type stratification introduced in the Thorn programming language [24]: concrete types, whose values are statically guaranteed to conform to a kind; a dynamic type, for values whose type is entirely unknown; and *like types*, which record the compiler’s expectation about a value without providing a static guarantee. In Thorn, for a variable of type `like string`, the compiler would check that it was only operated on as if it were a string, but would allow unchecked coercion from dynamic values. FiR repurposes this stratification for a compiler IR: *like types* record speculative information—such as feedback from previous executions—and an

explicit type-cast is required to convert a like type to a concrete type, providing a natural point for deoptimization on failure (Section 4.6).

3.4 Deoptimization with Sourir and CoreIR

Sourir [11] and CoreIR [2] formalize speculative compilation and deoptimization. Both allow multiple versions of a function to coexist—typically an unoptimized version alongside optimized versions that include speculative assumptions. Sourir encodes speculation via an assume instruction: if the assumption holds, execution continues; otherwise, control transfers to the unoptimized version. CoreIR refines this with separate Anchor and Assume instructions, distinguishing non-deterministic deoptimization points from deterministic guard checks.

FiR shares the multi-version design but deliberately omits deoptimization from its formalization: including it would be straightforward but would obscure the core semantics with support for on-stack replacement (Section 4.10). Instead, FiR’s type-cast mechanism serves an analogous role—a cast from a like type to a concrete type acts as a guard, providing a natural point for deoptimization on failure. Our implementation follows the approach of CoreIR for deoptimization.

3.5 The $\check{R}$ Compiler and PIR

FiR is designed as a replacement for PIR, the high-level intermediate representation of the $\check{R}$ optimizing compiler for R (Fig. 2). The $\check{R}$ compiler reuses the GNU-R front end to parse source code into an AST, and has its own interpreter and bytecode (RIR) as well as PIR for high-level transformations. Code generation is delegated to LLVM. PIR is an SSA-based IR [10]: each instruction has a type, effects, and immediates, where the type combines a set of kinds (*e.g.*, *real*, *int*, *str*) with flags tracking properties such as scalarity (\$), objectness (+), and delayed evaluation (^). Fig. 3 shows the unoptimized identity function: *LdArg* loads an argument from the stack returning *val?^*, *i.e.* any value that might be missing or a promise; *MkEnv* creates an environment and stores the parameter in it; *LdVar* loads the variable into a register (*eR* signals possible failure); and *Force* evaluates the value if it was a promise, with ! indicating possible reflective effects.




Fig. 2. $\check{R}$ ’s existing IRs

The PIR type system is ad hoc. It has evolved into its current shape, rather than being designed. Its idiosyncrasies include the fact that types are implemented with flags, hindering composition. For example, the missing value type is represented as a kind but also as a flag resulting in complex types such as *val?^|?* indicating a promise that can evaluate into a missing value or is a missing value itself. FiR is a formally grounded redesign of PIR.

4 An IR for Dynamic Languages

An IR does not inherently create opportunities for new optimization; any transformation expressible in FiR can be performed on a lower-level IR. Instead, an IR should be evaluated on two criteria:

- (1) does it make certain optimizations easier to express, and
- (2) does it reduce the likelihood of introducing unsound transformations.

```

val?^ %1.0 = LdArg 0
env   e1.1 = MkEnv a=%1.0, parent=G
val?^ %3.0 = LdVar a, e1.1          eR
val?  %3.3 = Force %3.0, e1.1      !
void          Return %3.3

```

Fig. 3. An identity function in PIR

Compilers are filled with subtle invariants that are easy to violate as the code base evolves. FiR is deliberately opinionated in the features it exposes, and includes both a type checker and a flow analysis to detect violations of key invariants early. While the IR is tailored to R and the $\hat{\text{R}}$ compiler, the underlying principles are broadly applicable to dynamic languages—particularly those with fewer dynamic features than R.

4.1 Syntax of FiR

FiR is a core calculus that gives semantics to an intermediate language for compiling dynamic languages such as R. It deliberately abstracts away many of the low-level details required in a full-featured intermediate representation. FiR includes only a minimal set of data types and control-flow constructs: it is an imperative language over integers and mutable integer vectors. Environments are not first-class values: while variables are typed, environments are not accessible as values. Unlike PIR, FiR is not SSA: the same variable can be assigned multiple times.

$F ::= f_{n_0} \dots f_{n_n}$	$v ::= x \mid r$
$fn ::= \text{fun } f\{ver_1 \dots ver_n\}$	$e ::= i \mid \text{vec}(e_1 \dots e_n) \mid v \mid e[e] \mid e\x
$ver ::= (rdefs) \xrightarrow{fx} t\{defs; e\}$	$\mid v = e \mid e[e] = e \mid e\$x = e$
$rdefs ::= r_0 : t_0 \dots r_n : t_n$	$\mid ver \leftarrow (e_0 \dots e_n)$
$defs ::= v_0 : t_0 \dots v_n : t_n$	$\mid f.i(e_0 \dots e_n)$
$t ::= k \text{ ox } cx$	$\mid f <sig> (e_0 \dots e_n)$
$k ::= * \mid V \mid I \mid v(I) \mid p^{fx}(t)$	$\mid e \text{ as } t$
$ox ::= o \mid b \mid s \mid f$	$\mid \text{force } e$
$fx ::= + \mid -$	$\mid \text{consume } r$
$cx ::= ! \mid ?$	$\mid \text{dup } e$
$sig ::= t_0 \dots t_n \xrightarrow{fx} t$	$\mid \text{prom}^{fx} <t>\{e\}$
	$\mid e; e$

Fig. 4. Syntax

The full syntax of the FiR calculus is given in Fig. 4. A program consists of a *function table* F and an *expression* e . A function table consists of one or more *function declarations*, $\text{fun } f\{ver_1 \dots ver_n\}$, where f is the function’s unique name and ver_i is a version. A *version*, $(rdefs) \xrightarrow{fx} t\{defs; e\}$, has typed argument definitions $rdefs$, local variables $defs$, a body expression e , a return type t , and a reflection bit fx . A type t is composed of a *kind* k defining the structure of the value, an *ownership* ox , and a *concreteness* cx indicating whether the type is certain or speculated (cf. Section 4.6). Variables are either *register variables* (r) or *named variables* (x); the distinction, motivated by reflection, is discussed in Section 4.3.

The expressions (Fig. 4) include standard constructs (integer and vector literals, variable reads and writes, sequencing), reflective access to named variables ($e\$x$ and $e\$x = e$), three forms of

function application (static calls, dispatched calls, and inlined versions), and constructs specific to FIR: type-casts ($e \text{ as } t$), promise creation ($\text{prom}^{f^x} \langle t \rangle \{e\}$) and forcing ($\text{force } e$), duplication ($\text{dup } e$), and register consumption ($\text{consume } r$). Each is discussed in the following subsections.

Examples use a slightly beautified syntax: we add **reg** and **var** annotations to distinguish between register and named variables, omit semicolons, allow declarations and assignments to be combined, and use defaults in type signatures (types with no ownership modifier are **s**, and those with no concreteness modifier are **!**).

4.2 Functions and Versions

In FIR, each function can have multiple implementations, or *versions*. Fig. 5 presents a source-language function definition (left) alongside three compiler-generated versions (right): a baseline version that handles any argument, a version specialized for a promise wrapping an integer vector, and a version specialized for a borrowed integer vector. Internally, the compiler maintains a flat namespace, where each unique function in the program has a single entry. Thus, `fn#1` in Fig. 5 contains all the compiled versions of the function expression shown on the left. The mapping between source-level function names and compiler-internal identifiers is an implementation detail. For our purposes, we focus on the internal table of functions maintained by the compiler.

Each version is intended to be semantically equivalent to the source function under some assumptions about its execution context. Verifying this equivalence is beyond our scope; instead, the IR guarantees that each version—independently of its siblings—is well-typed, and satisfies a property we call *well-flowed*, meaning that registers are not used before they are defined (cf. Section 5).

A version has a *signature* that consists of argument types, return type, and reflection bit, as well as a *body* (e.g. $(\text{reg } r:\mathbf{I}) \rightarrow \mathbf{I} \{ r+1 \}$). Versions are created by the compiler to provide optimized implementations of a function. They differ in their signature and the transformations applied to the body. Versions need not be mutually exclusive in their signature, and the compiler is free to select among them—possibly considering additional factors such as the likelihood of deoptimization or resource requirements.




Fig. 5. Source code (left) and FIR code with three versions (right)

4.3 Scopes and Variables

Each version defines a *scope* in which variables can be defined. Our source language allows scopes to be shared with promises and to be inspected reflectively, so variables that appear in the source code must be preserved in case another function looks them up. To allow the compiler to eliminate scopes that are not reflectively inspected, FiR distinguishes two kinds of variables: *named variables*, which are looked up symbolically and are exposed to reflection, and *register variables*, which are not available to reflective operations and can be implemented in any way the compiler sees fit. Parameters are passed as register variables. A common pattern is to store them into named variables before calling a function that may perform reflection:

```
(reg r:*)+→* {
  var a:* = r
  someFun<*>+→*(a)
}
```

While our source language allows shadowing and nested scopes, FiR does not model nesting; any lookup outside the current scope returns an unknown value.

Reflective reads and writes. A promise in FiR captures both an expression—its body—and the environment of the enclosing version. Given a variable r referencing a promise defined in a different version, the following mutates a named variable x in the captured environment:

```
r$x = 42
```

This can occur regardless of the expected type of x , or even its existence. Only named variables can be accessed reflectively; register variables can only be accessed from their definition scope.

Reflective access makes naive inlining unsound: inlining a version's body into the surrounding context merges their environments, altering behavior when reflection is present. Consider:

```
r1 = ((reg r:v(I)o)→I { r[1] = 2 }) <-- (consume r0)
r1 = (r0[1] = 2)
```

These two statements are equivalent since the version's body does not employ reflection. This contrasts with the following two lines, which have different behaviors (we assume $r0$ is a promise capturing the current local environment):

```
r1 = ((reg r:p+(*))+→* { var x:* ; x = 41 ; r$x = 42 ; x }) <-- (consume r0)
var x:* ; r1 = (x = 41 ; r0$x = 42 ; x)
```

The first line yields $r1=41$, while the second yields $r1=42$.

4.4 Promises: Creation and Force

A *promise* represents a delayed computation. In the source language, promises are *forced* implicitly when their value is needed, and their result is cached upon evaluation. In FiR, promises are first-class values with well-defined types that must be explicitly forced:

```
r : p(I)
r = prom<I>{ 42 }
force r
```

Here, r is a promise that, when forced, returns an integer. As a more realistic example, the source expression $f(x)$ may be translated to:

$$f \langle p^+(V) \rangle \xrightarrow{+} V \langle \text{prom}^+ \langle V \rangle \{ \text{force } x \} \rangle$$

This passes to f a promise that, when evaluated, reads and forces the variable x . Promises execute in the lexical scope in which they were created: they can read and update variables in that scope, and may even introduce new bindings—for instance, in $f(z=x+y)$ if z was not previously defined.

In R, promises are automatically forced when assigned to a variable or returned from a function, so they do not outlive the scope in which they were created. Our type system leverages this invariant to simplify reasoning about promise lifetimes (Section 4.6). Because force triggers evaluation of a delayed computation—potentially with side effects—the compiler cannot freely reorder code around it. F_{IR}'s flow analysis ensures that variables are not read before they are written (Section 5.1).

4.5 Copy-on-Write

F_{IR} supports copy-on-write semantics by making duplication explicit through a dedicated `dup` operation. All program points where a copy may be required must be explicitly marked with `dup`. The actual implementation of `dup` is left to the compiler: in a language like R that maintains an over-approximating reference count, `dup` needs to copy only when the count exceeds one. By making duplication explicit, the IR enables optimizations that eliminate unnecessary copies—copies that would otherwise be indistinguishable from semantically required ones when performed implicitly by other operations. The example below illustrates how the compiler inserts a `dup` before a vector is updated:

```
(reg r1: v(I)) : I {
  reg r2: v(I) = dup r1
  r2[1] = 42
}
```

Once duplicated, the vector can be safely mutated multiple times without further copying. As a result, function arguments typically undergo a single duplication before any mutation, enabling efficient updates while preserving copy-on-write semantics. The interaction between duplication and ownership is discussed in Section 4.6.

4.6 Ownership and Gradual Typing

Types in F_{IR}, ranged over by t , serve multiple purposes: they verify that compiler transformations preserve well-formedness, guide optimizations, and record compiler-generated hints. As described in Section 3.3, F_{IR} adopts the three-level type stratification of Thorn [24]—concrete types, a dynamic type, and like types—and repurposes it for a compiler IR. In addition, the type system tracks ownership qualifiers, enabling optimizations such as copy elimination.

Formally, each type is composed of three components: a *kind* (k), defining the structure of the value (e.g. I for integers, $v(I)$ for vectors of integers); an *ownership modifier* (ox); and a *concreteness modifier* (cx), indicating whether the type is certain (!) or speculated (?). A reflection bit (fx) extends the types of functions and promises (Section 4.7).

Expressions of a concrete type are guaranteed by the static type system to reduce to a value that conforms to their kind k . Consequently, no runtime checks or wrappers are needed. Like types are the only types that can be given to named variables, as we cannot statically prevent them from being reflectively reassigned to a value of a different kind. The dynamic type, represented by the kind $*$, captures any value; it can only be associated with the concreteness ? (the ? is elided in our examples).

The specific kinds are: scalar value (I), mutable vector ($v(I)$), and promise ($p^{fx}(t)$); the remaining kinds represent the union of the above: value (V) is the union of scalars and vectors, and any ($*$) is the union of all.

Ownership modifiers can be *f* (fresh), *o* (owned), *s* (shared), or *b* (borrowed). A value is fresh if it is not referenced by any variable (e.g., a newly-allocated value). The rules for copy-on-write are that any shared mutable value must be duplicated before being updated:

```
reg r2:v(I)o! = ...
reg r2:v(I)s! = dup r1
r2[0] = 42
```

The `dup` operation performs a copy (if needed, the source language maintains reference counts) and changes the ownership from *s* to *o* (owned). An owned value is uniquely referenced, even though it can be temporarily borrowed during a function call as long as this function does not mutate it and does not leak it (so that the value remains unique after the call). Borrowed types exist only in function parameters.

An owned value cannot be assigned to a variable, otherwise it would compromise uniqueness as it could be then assigned to multiple live variables. A fresh value can be assigned to an owned or shared variable. A value can be made fresh by copying using `dup`. Furthermore, a value in an owned register variable can become fresh by “consuming” the register (`consume r`); flow analysis ensures a register cannot be accessed after it is consumed.

```
(reg r1:v(I)o!)→v(I)o! { # `r1' is owned: it should not be aliased
  reg r2:v(I)o! ;
  r2 = r1 ; # Not allowed, because both `r1' and `r2' would reference the same value
  r2 = dup r1 ; # Ok, because `r1' points to a separate copy
  r2 = consume r1 ; # Ok, because it disallows any future access to `r1'
}
```

Ownership information can be used to statically eliminate useless copies of vectors. For instance, consider the following code:

```
(reg r1: v(I)o!) : I {
  reg r2: v(I)o! = dup r1
  r2[1] = 42
}
```

As `r1` is owned and is not accessed anymore after the duplication, `dup r1` can be replaced by `consume r1` (which is a no-op at runtime, but prevents `r1` from being accessed after this point):

```
(reg r1: v(I)o!) : I {
  reg r2: v(I)o! = consume r1
  r2[1] = 42
}
```

This code can then be simplified into the following:

```
(reg r1: v(I)o!) : I {
  r1[1] = 42
}
```

The syntax of types is permissive, but we restrict types that may appear in declarations. First, we define a left-associative component setter operator `&` on types:

$$\begin{aligned}
 k \text{ } o x \text{ } c x \quad \& \quad k' &= k' \text{ } o x \text{ } c x \\
 k \text{ } o x \text{ } c x \quad \& \quad o x' &= k \text{ } o x' \text{ } c x \\
 k \text{ } o x \text{ } c x \quad \& \quad c x' &= k \text{ } o x \text{ } c x'
 \end{aligned}$$

We use a number of predicates for convenience:

$$\begin{aligned}
 \text{shared}(t) &\equiv (t \ \& \ \mathbf{s}) = t \\
 \text{owned}(t) &\equiv (t \ \& \ \mathbf{o}) = t \\
 \text{fresh}(t) &\equiv (t \ \& \ \mathbf{f}) = t \\
 \text{borrowed}(t) &\equiv (t \ \& \ \mathbf{b}) = t \\
 \text{like}(t) &\equiv (t \ \& \ ?) = t \\
 \text{vec}(t) &\equiv (t \ \& \ \mathbf{v}(\mathbf{I}) \ \& \ !) = t \\
 \text{prom}(t) &\equiv \exists t', \text{fx. } (t \ \& \ \mathbf{p}^{\text{fx}}(t') \ \& \ !) = t \\
 \text{value}(t) &\equiv (t \ \& \ \mathbf{I} \ \& \ ! = t) \text{ or } (t \ \& \ \mathbf{v}(\mathbf{I}) \ \& \ ! = t) \text{ or } (t \ \& \ \mathbf{V} \ \& \ ! = t)
 \end{aligned}$$

We can now define the meaning of *well-formed* for types:

$$\begin{array}{c}
 \frac{t = t \ \& \ * \ \& \ ?}{\text{wf}(t)} \quad \frac{t = t \ \& \ \mathbf{I}}{\text{wf}(t)} \quad \frac{t = t \ \& \ \mathbf{v}(\mathbf{I})}{\text{wf}(t)} \quad \frac{t = t \ \& \ \mathbf{V}}{\text{wf}(t)} \quad \frac{\text{wf}(t') \quad \frac{t = t \ \& \ \mathbf{p}^{\text{fx}}(t')}{\text{shared}(t')} \quad \text{value}(t')}{\text{wf}(t)}
 \end{array}$$

Each rule defines well-formedness constraints for a specific kind. For instance, the first rule ensures that a type of the any kind ($*$) has a concreteness modifier of $?$ ($t = t \ \& \ * \ \& \ ?$ is true if and only if t is of kind $*$ and of concreteness $?$). The last rule ensures that the type returned by a promise is shared (it is referenced by the promise itself, which will return this same value each time it is forced) and cannot itself be a promise. The well-formedness constraints are extended to variable declarations with:

$$\frac{\text{wf}(t) \quad \text{shared}(t) \quad \text{like}(t)}{\text{wf}(x : t)} \quad \frac{\text{wf}(t) \quad \neg \text{fresh}(t)}{\text{wf}(r : t)}$$

This ensures that named variables are always shared and that their concreteness is $?$, and register variables cannot be fresh.

Subtyping and casts. A type-cast $e \text{ as } t$ checks that e has type t at runtime (or a subtype of it). Subtyping on types has, for instance, type $\text{Io}!$ be a subtype of $\text{Io}?$, itself a subtype of $*$. Two types can only be in subtyping relation if they have the same ownership: the ownership of a variable is invariant, and a change of ownership can only occur through an explicit operation. The role of type-casts is essential in FIR, as they are the only way to convert a like type into a concrete type. Thereby, any named variable must be cast before being used on an operation that requires a concrete type. Type casts can also be used on expressions with a concrete type in order to refine this type (e.g., refining a type $\text{Vo}!$ into $\text{Io}!$).

Type-casts are inserted when some uncertain type information needs to be checked against the expectations of the compiler in order to guarantee a safe execution: reflection may have altered the type of a variable in a way that cannot be predicted statically, the type of an expression may have been speculated by the compiler to enable more optimizations, etc. In a practical implementation, type-casts should provide an alternative execution in the case of failure: for instance, it could trigger some deoptimization mechanism that switches to a more general version of the current function. However, FIR does not assume any specific deoptimization mechanism. A compiler implementation is free to implement speculation and deoptimization (such as CoreIR [2]) or not.

4.7 Reflection in the Type System

Promise types and function signatures carry a *reflection bit*: $+$ indicates that the promise or function may use reflection, while $-$ ensures it does not. For example, the kind of a promise that yields an integer and may perform reflection when forced is written $\mathbf{p}^+(1s!)$. The first time a promise is forced, its result is cached, so the value returned by a promise is considered shared.

Reflection does not directly impact the typeability of an expression, as named variables are already treated pessimistically: they have dynamic or like types. Still, annotating promises and function signatures with their potential for reflection is essential for optimizations—in particular, scope elision is only sound when the compiler can verify that no reflective operation observes a scope.

4.8 Binding Time

When considering a call such as $f(x)$, it may be impossible to determine statically which function is being invoked. In dynamic languages, this ambiguity can arise for several reasons: f may be an argument passed by the caller of the current function; it may refer to a global symbol that can be redefined at runtime; or it may be resolved through dynamic dispatch based on properties of the argument. Our source language supports all of these forms of dynamic binding.

FIR does not have built-in support for dynamic binding mechanisms; both static and dispatched calls point to statically-known functions, though the version may be dynamic. Unknown functions can be called by passing them to an intrinsically-defined “dynamic call” function from the source language’s reference implementation. In practice, we are able to statically resolve most calls using speculation, but that is beyond the scope of this paper.

A function call is a *static call* if the target version of the function is known at compile time. In contrast, a *dispatched call* refers to cases where the compiler selects the most appropriate version based on a combination of static and dynamic information available at the call site. The compiler may restrict the set of candidate versions statically and generate dispatch logic to choose among the remaining versions at runtime. We do not mandate how a compiler selects the “best” version; each implementation may rely on its own heuristics, profiling data, or runtime inspection. A dispatched call to $fn\#1$ is written as:

```
r = fn#1<math>p^+(\star) \rightarrow I>(r0) \# Dispatched call
```

Here, the signature $\langle p^+(\star) \rightarrow I \rangle$ constrains the version of $fn\#1$ that the compiler will call; this signature is compatible with version 1 in Fig. 5. A more sophisticated compiler—such as $\tilde{R}$, which supports contextual dispatch [9]—might dynamically inspect the promise held in $r0$. If it determines that the promise is reflection-free and returns an integer, it could instead select version 2.

Dispatched calls are annotated with static type signatures for two reasons. First, some properties required to select a compatible version at runtime may not be dynamically recoverable: ownership qualifiers are computed statically and cannot be inferred from the program state, and the kind of an argument cannot be determined at runtime if the argument is *undef*. Second, to ensure soundness, the version selected at runtime must return a value that is type-compatible with the version selected statically. Two versions of the same function may produce different results for the same inputs—this is acceptable because each version is optimized under different assumptions—so long as the selected version conforms to the type system. The type system ensures that at least one version is type-compatible with the annotated static signature.

Static calls reference a version directly by index and do not require a signature. At runtime, the version is called without dispatch. A static call to version 3 of $fn\#1$ is written as:

```
r1 = fn#1.3( force r0 ) \# Static call
```

Concrete type information can be used to devirtualize a dispatched call into a static one. For instance, a call to $fn\#1$ (Fig. 5) could be statically bound to version 3 if the argument is known to be a vector of integers ($v(I)o!$). Like types, which do not offer static guarantees, can provide hints for the same purpose. For example, a version typed with $I?$ indicates an expectation that it will be called with an integer:

```

(reg r0:I?)→* {
  reg r1:I = r0 as I # deoptimize if fails, not shown here
  r1 + 1
}

```

The body must cast the like type to a concrete type before use, and should deoptimize on failure. In this context, casts serve as guards.

4.9 Inlining

When the compiler knows which version will be executed—such as in a static call—it may replace the call with the body of the corresponding version. Fig. 6 illustrates the syntax used to represent inlined code. A key point is that inlining preserves the scope of the inlined version: named variables defined within the callee remain in a distinct environment, separate from the caller. This is required because reflection can observe scope boundaries.

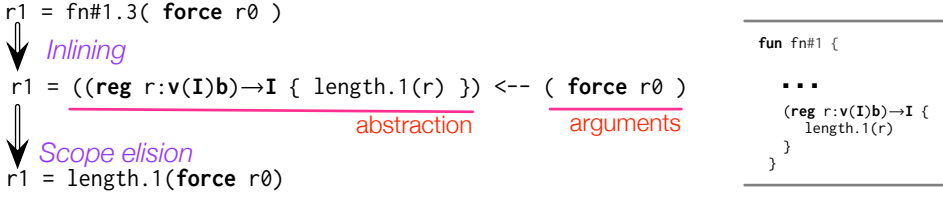


Fig. 6. Inlining and scope elision

If the compiler can determine that no reflective operation will observe a particular scope, it may safely remove it. This optimization is known as scope elision and is separate from inlining. FIR does not define how a scope is represented at runtime, but provides static information that may help optimize the representation, in particular the distinction between variables that can be accessed through reflection and those that cannot.

4.10 Discussion

The design of FIR draws on lessons from prior work on the R compiler infrastructure. This section discusses noteworthy design choices and simplifications.

Deoptimization. Our implementation follows the approach of CoreIR [2] for deoptimization. For simplicity, we omit these constructs from the definition of FIR. Including them would be straightforward but would obscure the core semantics with support for on-stack replacement. The only requirement that deoptimization adds is an identified base version for each function; this version is called when speculation fails.

Reflection. FIR allows named variables to be redefined or deleted via reflection. The specific reflection interface is not specified, but this possibility is why named variables can only have like types—if the source language did not expose local variables to reflection, they could be given stronger types. Consider a version that stores a value in named variable *a*, calls a possibly-reflective function *f*, and returns *a*:

```

(reg r:*)+→* {
  var a:*
  a = r
  f<*+→*>( 42 )
  a
}

```

In R, f can delete a , causing the lookup to return the value of a in an enclosing scope, or to fail.

Promises. In FiR, promises cannot escape their scope of creation. Without this restriction, any register accessed in a promise that outlives its scope would have to be considered shared¹, complicating the ownership analysis. This restriction is acceptable because in R, promises are forced when assigned or returned. It is possible for a promise to escape via a *delayed assignment*², or by returning the environment in which it lives (environments are first-class values in R), but these cases are rare and can be treated as reflection.

Omitted Features. The R language has features that are not modeled in FiR. Most are implemented by calls to native functions and do not need to be reflected in the IR. Two are worth mentioning: name lookups and attributes.

The semantics of name lookup in R is surprisingly rich: a chain of environments is searched, and in some cases promises must be forced to determine which value to return. These complex lookups are limited to named variables and do not affect FiR because the compiler only reasons about names in the current scope; for names outside that scope, FiR makes no assumption and the implementation defaults to the source language lookup mechanism.

In R, any value can have attached attributes—a map from name to value. For example, a `class` attribute gives a value a particular behavior during object-oriented dispatch. The $\check{R}$ compiler tracks only the absence of attributes, which is the overwhelmingly common case. This could be added to FiR as an extra bit on value types; a more fine-grained treatment is left to future work.

With these design choices in place, we now turn to the formal semantics of FiR.

5 Static and Dynamic Semantics of FiR

We first define a static semantics for FiR, composed of a type-checker and a flow analysis (Section 5.1). We then define a dynamic semantics (Section 5.2), and prove that programs that are well-typed and well-flowed cannot get stuck (Section 5.3). Finally, we demonstrate the static and dynamic semantics on a small example (Section 5.4).

5.1 Static Semantics

The static semantics of FiR are defined by two relations over versions, $\llbracket \text{ver} \rrbracket : \text{sig}$ and $F\llbracket \text{ver} \rrbracket$. These relations assert, respectively, that a version ver is well-typed and well-flowed. The function table F of a program is considered well-defined when every version within it satisfies both criteria. The typing relation guarantees adherence to constraints related to types, ownership, concreteness, and reflection. The flow relation ensures that register variables are initialized before they are used and prevents their access after a **consume** operation. Although both properties are required for safety, we define two separate analyses: the typing relation is checked by a flow-insensitive analysis (with an invariant type environment), while the flow relation is checked by a simpler but flow-sensitive analysis—mirroring the separation between type checking and borrow checking in Rust.

For convenience, we consider the function table F to be fixed, and allow the deduction rules for the judgments $\llbracket \text{ver} \rrbracket : \text{sig}$ and $F\llbracket \text{ver} \rrbracket$ to access it, without making this dependency explicit in the judgment.

We establish transitive, reflexive and antisymmetric subtyping relations as follows:

- For reflection: $-\leq_f +$.
- For ownership, subtyping is restricted to equality: $ox \leq_o ox$.

¹We recall that a promise is not evaluated when it is created, but the first time it is forced.

²A delayed assignment is a special instruction that assigns an expression to a variable without forcing the underlying promise.

- For concreteness modifiers: $!\leq_c?$.
- For kinds:

$$\frac{}{\mathbf{I} \leq_k \mathbf{V}} \quad \frac{}{\mathbf{v}(\mathbf{I}) \leq_k \mathbf{V}} \quad \frac{}{k \leq_k *} \quad \frac{t \leq t' \quad f_x \leq_f f_{x'}}{\mathbf{p}^{f_x}(t) \leq_k \mathbf{p}^{f_{x'}}(t')}$$

- For types:

$$\frac{k \leq_k k' \quad o_x \leq_o o_{x'} \quad c_x \leq_c c_{x'}}{k \ o_x \ c_x \leq_k k' \ o_{x'} \ c_{x'}}$$

- For version signatures:

$$\frac{t'_0 \leq t_0 \dots t'_n \leq t_n \quad t \leq t' \quad f_x \leq_f f_{x'}}{t_0 \dots t_n \xrightarrow{f_x} t \leq_{\text{sig}} t'_0 \dots t'_n \xrightarrow{f_{x'}} t'}$$

Type checking. The [VER] rule shown next checks the definition of a version. The variables that can be used in a version are its parameters (all registers), and locals (registers and named variables); these variables and their types make up the type context which is used to type the version's body. All type annotations must be well-formed. The reflection bit and return type of the version must be those inferred for its body, with the ownership component transformed as follows: if the body returns an owned value, the returned value is fresh (the only reference ceases to exist when exiting the version's scope); otherwise ownership is unchanged. A version that returns a borrowed value is untypeable. Versions can only return non-promise values; allowing a promise to escape its scope would force every register it captures to be considered shared (cf. Section 5.3).

$$\begin{array}{c} \Gamma = \text{rdefs}, \text{defs} \quad \text{wf}(\Gamma) \\ \Gamma \vdash \llbracket e \rrbracket : t' \\ \text{types}(\text{rdefs}) = t_0 \dots t_n \\ \Gamma \vdash R\llbracket e \rrbracket = f_x \quad t = \begin{cases} t' & \text{if shared}(t') \text{ or fresh}(t') \\ t' \ \& \ \mathbf{f} & \text{if owned}(t') \end{cases} \\ \text{value}(t) \quad \text{wf}(t) \end{array} \quad \frac{}{[\text{VER}] \quad \llbracket (\text{rdefs}) \xrightarrow{f_x} t \{ \text{defs}; e \} \rrbracket : t_0 \dots t_n \xrightarrow{f_x} t}$$

The reflection analysis $\Gamma \vdash R\llbracket e \rrbracket = f_x$ is described below (Fig. 8). Fig. 7 gives the typing judgments. These rely on the following auxiliary definition, which relates parameter types t' with the argument types t they can receive:

$$\frac{t \ \& \ \mathbf{f} \leq t' \ \& \ \mathbf{f} \quad \text{shared}(t') \Rightarrow (\text{shared}(t) \text{ or fresh}(t)) \quad \text{owned}(t') \Rightarrow \text{fresh}(t)}{\text{match}(t, t')}$$

A parameter of type t' can be given an argument of type t if and only if: (1) the kind and concreteness of t are subtypes of those of t' ; and (2) if t' expects a borrowed value, t can have any ownership; if t' expects a shared value, t must be shared or fresh; if t' expects an owned value, t must be fresh. Note that function parameters must be well-formed, so they cannot have fresh ownership.

The only literals are integers, considered shared by convention (ownership is irrelevant for immutable values). All types are explicit; promise return and call signatures are checked, never inferred. Sequences have the type of their last element. Only vectors can be indexed, and only integers can be used as indices. Variables read via reflection always have unknown types, and variables of any type can be written via reflection; the latter named variable types are never concrete.

$$\begin{array}{c}
\text{[LIT]} \frac{}{\Gamma \vdash \llbracket i \rrbracket : \mathbf{Is!}} \quad \text{[VAR]} \frac{}{\Gamma, v : t \vdash \llbracket v \rrbracket : t} \quad \text{[SUB]} \frac{\Gamma \vdash \llbracket e \rrbracket : t \quad t \leq t' \quad \text{wf}(t')}{\Gamma \vdash \llbracket e \rrbracket : t'} \quad \text{[SEQ]} \frac{\Gamma \vdash \llbracket e \rrbracket : t \quad \Gamma \vdash \llbracket e' \rrbracket : t'}{\Gamma \vdash \llbracket e; e' \rrbracket : t'} \\
\\
\text{[PRO]} \frac{\Gamma \vdash \llbracket e \rrbracket : t' \quad \text{value}(t') \quad \Gamma \vdash R[\llbracket e \rrbracket] = fx \quad t = p^{fx}(t')s!}{\Gamma \vdash \llbracket \text{prom}^{fx} \langle t' \rangle \{e\} \rrbracket : t} \quad \text{[VEC]} \frac{\Gamma \vdash \llbracket e_1 \rrbracket : \mathbf{Is!} \dots \Gamma \vdash \llbracket e_n \rrbracket : \mathbf{Is!} \quad t' = v(\mathbf{I})f!}{\Gamma \vdash \llbracket \text{vec}(e_1 \dots e_n) \rrbracket : t'} \quad \text{[REA]} \frac{\Gamma \vdash \llbracket e \rrbracket : t \quad \Gamma \vdash \llbracket e' \rrbracket : \mathbf{Is!} \quad t \& f \leq v(\mathbf{I})f!}{\Gamma \vdash \llbracket e[e'] \rrbracket : \mathbf{Is!}} \\
\\
\text{[APP]} \frac{\begin{array}{c} \llbracket \text{ver} \rrbracket : t_0 \dots t_n \xrightarrow{fx} t \\ \Gamma \vdash \llbracket e_0 \rrbracket : t'_0 \dots \Gamma \vdash \llbracket e_n \rrbracket : t'_n \\ \text{match}(t'_0, t_0) \dots \text{match}(t'_n, t_n) \end{array}}{\Gamma \vdash \llbracket \text{ver} \leftarrow (e_0 \dots e_n) \rrbracket : t} \quad \text{[CAL]} \frac{\begin{array}{c} \text{sig}(F(f)(i)) = t_0 \dots t_n \xrightarrow{fx} t \\ \Gamma \vdash \llbracket e_0 \rrbracket : t'_0 \dots \Gamma \vdash \llbracket e_n \rrbracket : t'_n \\ \text{match}(t'_0, t_0) \dots \text{match}(t'_n, t_n) \end{array}}{\Gamma \vdash \llbracket f.i(e_0 \dots e_n) \rrbracket : t} \\
\\
\text{[DIS]} \frac{\begin{array}{c} \text{sig} = t_0 \dots t_n \xrightarrow{fx} t \\ \exists i. \text{sig}(F(f)(i)) \leq \text{sig} \\ \Gamma \vdash \llbracket f.i(e_0 \dots e_n) \rrbracket : t' \end{array}}{\Gamma \vdash \llbracket f < \text{sig} \rangle (e_0 \dots e_n) \rrbracket : t} \quad \text{[CAS]} \frac{\Gamma \vdash \llbracket e \rrbracket : t' \quad \text{wf}(t) \quad t' \leq t \text{ or } t \leq t'}{\Gamma \vdash \llbracket e \text{ as } t \rrbracket : t} \quad \text{[FOR]} \frac{\Gamma \vdash \llbracket e \rrbracket : p^{fx}(t)s!}{\Gamma \vdash \llbracket \text{force } e \rrbracket : t} \\
\\
\text{[USE]} \frac{\begin{array}{c} \Gamma \vdash \llbracket r \rrbracket : t \\ \text{owned}(t) \\ t' = t \& f \end{array}}{\Gamma \vdash \llbracket \text{consume } r \rrbracket : t'} \quad \text{[DUP]} \frac{\begin{array}{c} \Gamma \vdash \llbracket e \rrbracket : t \\ \text{vec}(t) \\ t' = t \& f \end{array}}{\Gamma \vdash \llbracket \text{dup } e \rrbracket : t'} \quad \text{[ASS]} \frac{\begin{array}{c} t = \Gamma(v) \quad \Gamma \vdash \llbracket e \rrbracket : t' \\ \text{match}(t', t) \\ \neg \text{borrowed}(t) \end{array}}{\Gamma \vdash \llbracket v = e \rrbracket : t} \\
\\
\text{[WR1]} \frac{\begin{array}{c} \Gamma \vdash \llbracket e \rrbracket : v(\mathbf{I})o! \\ \Gamma \vdash \llbracket e' \rrbracket : t \\ \Gamma \vdash \llbracket e'' \rrbracket : t \\ t = \mathbf{Is!} \end{array}}{\Gamma \vdash \llbracket e[e'] = e'' \rrbracket : t} \quad \text{[REF]} \frac{\Gamma \vdash \llbracket e \rrbracket : t \quad \text{prom}(t)}{\Gamma \vdash \llbracket e\$x \rrbracket : *s?} \quad \text{[RW1]} \frac{\begin{array}{c} \Gamma \vdash \llbracket e \rrbracket : t \quad \text{prom}(t) \\ \Gamma \vdash \llbracket e' \rrbracket : t' \quad \text{value}(t') \quad \text{shared}(t') \end{array}}{\Gamma \vdash \llbracket e\$x = e' \rrbracket : t'}
\end{array}$$

Fig. 7. Types

Promises cannot be reflectively written, preventing them from escaping their scope of definition (cf. Section 5.3).

The type system as defined in Fig. 7 is not syntax-directed because of the subsumption rule SUB. This is a presentation choice for clarity and concision; subsumption can easily be inlined in other rules when needed.

Reflection analysis. Fig. 8 shows the analysis that checks reflection annotations: a function must have the reflection bit (+) if any of its applied versions or forced promises have the reflection bit. Note that the reflection analysis does not impact the typeability directly (reflection has no influence on the applicability of the deduction rules of the type system), but the information it brings is useful for optimizing code, as discussed in Section 1.

- A version is reflective if annotated as such. A version must be annotated as reflective if its body is reflective.

- Force is reflective if the type of the promise being forced has the reflective modifier. This makes the reflection analysis and the type system mutually recursive, though this definition is well-founded as recursive calls only apply on strict sub-expressions.
- An application is reflective if the version being applied is, and a call is reflective if the version being called is.
- Any expression containing reflective sub-expressions is reflective.

We define the reflection union $fx \vee fx'$ to be $+$ if either fx or fx' is, otherwise $-$:

$$\begin{aligned} - \vee - &= - \\ - \vee + &= + \\ + \vee fx &= + \end{aligned}$$

Flow Analysis. A program is well-flowed if register variables are initialized before their first access and never accessed after a **consume**. The challenge is that the point where a promise is evaluated may be difficult to know statically—in the general case it amounts to deciding reachability of every force instruction. A precise solution would involve a data-flow analysis modeling the flow of promises into force operations. We choose to approximate this analysis coarsely for the purpose of simplicity, enabling us to provide a correctness proof. An implementation may replace the analysis presented here with any sound approximation.

$$\begin{array}{c} \frac{\Gamma = rdefs, defs \quad \Gamma \vdash R[e] = fx' \quad fx' \leq_f fx}{[RVER] \quad R[(rdefs) \xrightarrow{fx} t\{defs; e\}] = fx} \quad [RLIT] \quad \frac{}{\Gamma \vdash R[i] = -} \quad [RVAR] \quad \frac{}{\Gamma \vdash R[v] = -} \\ \\ [RSEQ] \quad \frac{\Gamma \vdash R[e] = fx \quad \Gamma \vdash R[e'] = fx' \quad fx'' = fx \vee fx'}{\Gamma \vdash R[e; e'] = fx''} \quad [RVEC] \quad \frac{\Gamma \vdash R[e_0; \dots e_n] = fx}{\Gamma \vdash R[\mathbf{vec}(e_0 \dots e_n)] = fx} \quad [RAPP] \quad \frac{R[ver] = fx \quad \Gamma \vdash R[e_0; \dots e_n] = fx' \quad fx'' = fx \vee fx'}{\Gamma \vdash R[ver \leftarrow (e_0 \dots e_n)] = fx''} \\ \\ [RCAL] \quad \frac{\text{sig}(F(f)(i)) = t_0 \dots t_n \xrightarrow{fx} t \quad \Gamma \vdash R[e_0; \dots e_n] = fx' \quad fx'' = fx \vee fx'}{\Gamma \vdash R[f.i(e_0 \dots e_n)] = fx''} \quad [RDIS] \quad \frac{\text{sig} = t_0 \dots t_n \xrightarrow{fx} t \quad \Gamma \vdash R[e_0; \dots e_n] = fx' \quad fx'' = fx \vee fx'}{\Gamma \vdash R[f < \text{sig} > (e_0 \dots e_n)] = fx''} \\ \\ [RFOR] \quad \frac{\Gamma \vdash R[e] = fx \quad \Gamma \vdash [e] : \mathbf{p}^{fx'}(t) \quad fx'' = fx \vee fx'}{\Gamma \vdash R[\mathbf{force} \ e] = fx''} \quad [RCAS] \quad \frac{\Gamma \vdash R[e] = fx}{\Gamma \vdash R[e \ \mathbf{as} \ t] = fx} \quad [RREFL] \quad \frac{}{\Gamma \vdash R[r\$x] = +} \\ \\ [RREA] \quad \frac{\Gamma \vdash R[e] = fx}{\Gamma \vdash R[r[e]] = fx} \quad [RDUP] \quad \frac{\Gamma \vdash R[e] = fx}{\Gamma \vdash R[\mathbf{dup} \ e] = fx} \quad [RUSE] \quad \frac{\Gamma \vdash R[e] = fx}{\Gamma \vdash R[\mathbf{consume} \ e] = fx} \\ \\ [RPRO] \quad \frac{}{\Gamma \vdash R[\mathbf{prom}^{fx} < t > \{e\}] = -} \quad [RASS] \quad \frac{\Gamma \vdash R[e] = fx}{\Gamma \vdash R[v = e] = fx} \quad [RWRI] \quad \frac{\Gamma \vdash R[e; e'] = fx}{\Gamma \vdash R[r[e] = e'] = fx} \\ \\ [RRWR] \quad \frac{}{\Gamma \vdash R[r\$x = e] = +} \end{array}$$

Fig. 8. Reflection Analysis

An *action* (R, W, U, C) is a tuple consisting of four sets of registers. It describes the registers that an expression may interact with:

- R : The registers that it may read before assigning. This does not include registers that are read after being assigned.
- W : The registers that it assigns.
- U : The registers that it may consume.
- C : The registers that it may capture, *i.e.* the registers that are read/assigned/used in a promise that it creates.

The notable restrictions that our flow analysis enforces are:

- R registers must be assigned and not used before the expression annotated with the action.
- W registers must not be used before. They don't have to be assigned before.
- C registers must not be used before or after.

U doesn't have intrinsic restrictions, although any register in U is guaranteed to be in R (so it must be assigned and not used before). U 's purpose is to restrict *other* actions when it's composed with them.

$$\begin{array}{c}
\begin{array}{ccc}
\text{[FLIT]} \frac{A = (\emptyset, \emptyset, \emptyset, \emptyset)}{F[i] = A} & \text{[FVAR]} \frac{A = (\emptyset, \emptyset, \emptyset, \emptyset)}{F[x] = A} & \text{[FSQ]} \frac{F[e] = A \quad F[e'] = A' \quad A'' = A ; ; A'}{F[e; e'] = A''} \\
\text{[FMUL]} \frac{F[e_0] = A_0 \dots F[e_n] = A_n \quad A' = (\emptyset, \emptyset, \emptyset, \emptyset) ; ; A_0 ; ; \dots A_n}{F[e_0 \dots e_n] = A'} & \text{[FVEC]} \frac{F[e_1 \dots e_n] = A}{F[\text{vec}(e_1 \dots e_n)] = A} & \text{[FAPP]} \frac{F[\text{ver}] \quad F[e_0 \dots e_n] = A}{F[\text{ver} \leftarrow (e_0 \dots e_n)] = A} \\
\text{[FCAL]} \frac{F[e_0 \dots e_n] = A}{F[f.i(e_0 \dots e_n)] = A} & \text{[FDis]} \frac{F[e_0 \dots e_n] = A}{F[f <sig> (e_0 \dots e_n)] = A} & \text{[FCAS]} \frac{F[e] = A}{F[e \text{ as } t] = A} \\
\text{[FFOR]} \frac{F[e] = A}{F[\text{force } e] = A} & \text{[FREG]} \frac{A = \text{read } r}{F[r] = A} & \text{[FREFL]} \frac{A = \text{read } r}{F[r\$x] = A} & \text{[FREA]} \frac{F[e] = A \quad A' = \text{read } r ; ; A}{F[r[e]] = A'} \\
\text{[FDUP]} \frac{F[e] = A}{F[\text{dup } e] = A} & \text{[FUSE]} \frac{A = \text{read } r ; ; \text{use } r}{F[\text{consume } r] = A} & \text{[FPRO]} \frac{F[e] = A \quad A = (R, W, U, C) \quad C' = R \cup W \cup U \cup C}{F[\text{prom}^{fx} <t> \{e\}] = (R, \emptyset, U, C')} \\
\text{[FASS1]} \frac{F[e] = A}{F[x = e] = A} & \text{[FASS2]} \frac{F[e] = A \quad A' = A ; ; \text{write } r}{F[r = e] = A'} & \text{[FWRI]} \frac{F[e] = A \quad F[e'] = A' \quad A'' = \text{read } r ; ; A ; ; A' ; ; \text{write } r}{F[r[e] = e'] = A''} \\
& \text{[FRWR]} \frac{F[e] = A \quad A' = \text{read } r ; ; A ; ; \text{write } r}{F[r\$x = e] = A'}
\end{array}
\end{array}$$

Fig. 9. Flow Analysis

If the following rule holds then the version is well-flowed:

$$\frac{F[e] = (R, W, U, C) \quad R \subseteq \{r_0 \dots r_n\}}{F[(r_0 : t_0 \dots r_n : t_n) \xrightarrow{fx} t\{\text{defs}; e\}]}$$

The judgments of Fig. 9 rely on the following auxiliary definitions:

$$\begin{aligned} \text{read } r &= (\{r\}, \emptyset, \emptyset, \emptyset) \\ \text{write } r &= (\emptyset, \{r\}, \emptyset, \emptyset) \\ \text{use } r &= (\emptyset, \emptyset, \{r\}, \emptyset) \\ \text{capture } r &= (\emptyset, \emptyset, \emptyset, \{r\}) \end{aligned}$$

When two expressions are executed in sequence, the sequence's action is the composition ($;;$) of their actions. Composition is only defined when valid, e.g. there is no way to compose an action that uses a register with an action that reads or assigns it, since a register cannot be read nor assigned after it is used.

$$\frac{\begin{array}{l} A = (R, W, U, C) \\ A' = (R', W', U', C') \\ U \cap (R' \cup W' \cup U' \cup C') = \emptyset \\ C \cap U' = \emptyset \end{array}}{A ;; A' = (R \cup (R' - W), W \cup W', U \cup U', C \cup C')}$$

Our flow analysis runs on each subexpression of the body of a version and composes their actions. Analysis fails if this action cannot be computed, or if its R set is not a subset of the version's parameters. The former means there is a read/assign/capture after consume or a capture before consume; the latter means there is a read before assign.

5.2 Dynamic Semantics

We define a small-step operational semantics for FiR with evaluation contexts [8]. The program state is captured by a *heap* (H), mapping *references* (o) to values (0), written $o_0 \mapsto 0_0 \dots o_n \mapsto 0_n$. References include *undef*, which does not occur on the left of a mapping.

A FiR expression that reduces to *undef* models a source-language error or control-flow diversion. We consider *undef* an instance of every type, so that failable operations like indexing and casting can be typed as if they cannot fail; our type system focuses on preventing stuck states, and *undef* is not stuck.

Values are defined as follows:

$$\begin{array}{ll} 0 ::= & i \quad \text{integers} \\ & | \text{vec}(i_0 \dots i_n) \quad \text{mutable vectors} \\ & | \langle e, k, o \rangle \quad \text{unevaluated promises} \\ & | \langle o_1, k, o_2 \rangle \quad \text{cached promises} \\ & | (v_0 \mapsto o_0 \dots v_n \mapsto o_n) \quad \text{environments} \end{array}$$

The syntax is extended with the following terms:

$$\begin{array}{ll} e ::= & \dots \\ & | o \quad \text{references} \\ & | \{e\}_o \quad \text{version scope} \\ & | \{e\}_{o'} \quad \text{promise scope} \end{array}$$

Evaluation contexts are defined as follows:

$$\begin{aligned} \mathcal{E} ::= & \quad [] \mid \mathbf{vec}(o_1 \dots o_k, \mathcal{E}, e_1 \dots e_n) \mid \mathcal{E}[e] \mid o[\mathcal{E}] \mid \mathcal{E}\$x \\ & \mid v = \mathcal{E} \mid \mathcal{E}[e] = e \mid o[\mathcal{E}] = e \mid o[o] = \mathcal{E} \mid \mathcal{E}\$x = e \mid o\$x = \mathcal{E} \\ & \mid \mathbf{ver} \leftarrow (o_1 \dots o_k, \mathcal{E}, e_1 \dots e_n) \\ & \mid f.i(o_1 \dots o_k, \mathcal{E}, e_1 \dots e_n) \mid f < \mathbf{sig} > (o_1 \dots o_k, \mathcal{E}, e_1 \dots e_n) \\ & \mid \mathcal{E} \mathbf{as} \ t \mid \mathbf{force} \ \mathcal{E} \mid \mathbf{consume} \ r \mid \mathbf{dup} \ \mathcal{E} \mid \mathcal{E}; e \end{aligned}$$

The relation $H[\![e]\!]_o \Rightarrow H'[\![e']\!]_o$ transforms a configuration consisting of a heap H and expression e in an environment referenced by o , to an updated heap H' and expression e' . Reduction proceeds through context:

$$H[\![\mathcal{E}[e]\!]_o \Rightarrow H'[\![\mathcal{E}[e']\!]_o \quad \text{if} \quad H[\![e]\!]_o \Rightarrow H'[\![e']\!]_o$$

The main judgments appear in Fig. 10. The semantics deal with out-of-bounds array indexing and reads of undefined named variables by producing `undef` (the rather mundane rules [DNVARUNDEF], [DWRIUNDEF], [DREAUNDEF] and [DRRDUNDEF] are not shown in the figure). Furthermore, the following rule propagates `undef` to the top level:

$$H[\![\mathcal{E}[e]\!]_o \Rightarrow H'[\![\mathbf{undef}]\!]_o \quad \text{if} \quad H[\![e]\!]_o \Rightarrow H'[\![\mathbf{undef}]\!]_o$$

The above rule is prioritized over the other judgments.

Several expressions allocate on the heap. Literals (DLIT), promise constructors (DPRO), and vector constructors (DVEC) each create a fresh reference and reduce to it. A `dup` expression (DDUP) copies its argument into a fresh reference. An application (DAPP1) allocates a new environment initialized with the parameter values; this environment is released when the application returns (DAPP2).

The remaining rules manipulate existing environments. A variable expression (DVAR) looks up its variable in the enclosing environment; a `consume` expression (DCON) does the same but additionally removes the mapping, since the register cannot be accessed again. Assignment (DASS) maps a variable to a reference in the enclosing environment, and reflective read (DRRD) and write (DRWR) operate on the promise's captured environment instead. In contrast, a vector write (DWRI) mutates the vector in-place rather than remapping a variable; the vector must be referenced at most once in the heap, otherwise the reduction is stuck.

Control flow is handled by calls, dispatch, forcing, and casts. A static call (DCAL) looks up a version by function and index and reduces to an application. A dispatch call (DDIS) selects the version whose signature is a subtype of the annotation at the call site and whose parameter types are satisfied by the runtime argument types. Forcing a cached promise (DFOR1) returns the cached value; forcing an unevaluated promise (DFOR2, DFOR3) evaluates its body in the promise's environment and caches the result. A cast (DCAS) evaluates to its argument if the kind matches, otherwise to `undef`. Sequences (DSEQ) evaluate left to right and reduce to the last element. Rules DCTX1 and DCTX2 are mechanical context switches that evaluate version or promise code in the appropriate scope.

The rules rely on the following notation:

- o fresh denotes a reference o not occurring in the heap.
- The $\mathbf{env}$ function extracts the environment reference captured by a promise ($\mathbf{env}(\langle _, k, o \rangle) = o$).
- $H(o)$ returns the value of o in H ; if that value is an environment, $H(o)(v)$ returns the value mapped to v .
- $(H, o \mapsto o')$ and $(H(o'), x \mapsto o'')$ update a mapping; $H \setminus o$ and $H(o) \setminus v$ remove one.
- $\mathbf{sig}(\mathbf{ver})$ returns the signature of $\mathbf{ver}$.
- $F(f)(i)$ returns version i of function f .
- $\mathbf{refs}(H, o)$ returns the number of occurrences of o on the right-hand side of environment mappings in H .

$$\begin{array}{c}
\text{o' fresh} \\
\frac{H' = H, \text{o}' \mapsto i}{H[\![i]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad [\text{DLIT}]
\end{array}
\quad
\begin{array}{c}
\text{o' fresh} \\
\frac{H' = H, \text{o}' \mapsto \langle e, \mathbf{p}(t \text{ fx}), \text{o} \rangle}{H[\![\mathbf{prom}^{\text{fx}} \langle t \rangle \{e\}]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad [\text{DPRO}]
\end{array}$$

$$\begin{array}{c}
\text{o' fresh} \\
\frac{H(\text{o}_1) = i_1 \dots H(\text{o}_n) = i_n \quad H' = H, \text{o}' \mapsto \mathbf{vec}(i_1 \dots i_n)}{H[\![\mathbf{vec}(\text{o}_1 \dots \text{o}_n)]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad [\text{DVEC}]
\end{array}
\quad
\begin{array}{c}
\text{o' = H(o)(v)} \\
\frac{\text{o' = H(o)(v)}}{H[\![v]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad [\text{DVAR}]
\end{array}
\quad
\begin{array}{c}
0 = H(\text{o}), v \mapsto \text{o}' \\
\frac{H' = H, \text{o} \mapsto 0}{H[\![v = \text{o}']\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad [\text{DASS}]
\end{array}$$

$$\begin{array}{c}
\text{o'' fresh} \\
\frac{H' = H, \text{o}'' \mapsto H(\text{o}')}{H[\![\mathbf{dup} \text{o}']\!]_{\text{o}} \Rightarrow H'[\![\text{o}'']\!]_{\text{o}}} \quad [\text{DDUP}]
\end{array}
\quad
\begin{array}{c}
\text{o' = H(o)(r)} \\
0 = H(\text{o}) \setminus r \\
\frac{H' = H, \text{o} \mapsto 0}{H[\![\mathbf{consume} r]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad [\text{DCON}]
\end{array}$$

$$\begin{array}{c}
\text{smallest } i \text{ such that} \\
\text{sig}(F(f)(i)) = t_0 \dots t_n \xrightarrow{\text{fx}} t \\
\frac{H \vdash \text{o} \leq_{\text{dyn}} t \Rightarrow \text{o}' = \text{o} \quad H \vdash \text{o} \not\leq_{\text{dyn}} t \Rightarrow \text{o}' = \text{undef}}{H[\![\text{o}' \text{ as } t]\!]_{\text{o}} \Rightarrow H[\![\text{o}']\!]_{\text{o}}} \quad [\text{DCAS}]
\end{array}
\quad
\begin{array}{c}
H(\text{o}') = \langle \text{o}'', k, \text{o}'' \rangle \\
\frac{H' = H, \text{o}' \mapsto \langle \text{undef}, k, \text{o}'' \rangle}{H[\![\mathbf{force} \text{o}']\!]_{\text{o}} \Rightarrow H[\![\text{o}'']\!]_{\text{o}}} \quad [\text{DFOR1}]
\end{array}$$

$$\begin{array}{c}
H(\text{o}') = \langle e, k, \text{o}'' \rangle \\
\frac{H' = H, \text{o}' \mapsto \langle \text{undef}, k, \text{o}'' \rangle}{H[\![\mathbf{force} \text{o}']\!]_{\text{o}} \Rightarrow H'[\![\{e\}_{\text{o}''}^{\text{o}'}]\!]_{\text{o}}} \quad [\text{DFOR2}]
\end{array}
\quad
\begin{array}{c}
H' = H, \text{o}''' \mapsto \langle \text{o}', k, \text{o}'' \rangle \\
\frac{H' = H, \text{o}''' \mapsto \langle \text{o}', k, \text{o}'' \rangle}{H[\![\{ \text{o}' \}_{\text{o}''}^{\text{o}'''}]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad [\text{DFOR3}]
\end{array}$$

$$\begin{array}{c}
\text{smallest } i \text{ such that} \\
\text{sig}(F(f)(i)) = t_0 \dots t_n \xrightarrow{\text{fx}} t \\
\frac{H \vdash \text{o}_0 \leq_{\text{dyn}} t_0 \dots H \vdash \text{o}_n \leq_{\text{dyn}} t_n \quad e = f.i(\text{o}_0 \dots \text{o}_n)}{H[\![f \langle \text{sig} \rangle (\text{o}_0 \dots \text{o}_n)]\!]_{\text{o}} \Rightarrow H[\![e]\!]_{\text{o}}} \quad [\text{DDIS}]
\end{array}
\quad
\begin{array}{c}
\text{o' fresh} \\
\text{ver} = (r_0 : t_0 \dots r_n : t_n) \xrightarrow{\text{fx}} t \{ \text{defs}; e \} \\
\frac{\text{ver} = (r_0 : t_0 \dots r_n : t_n) \xrightarrow{\text{fx}} t \{ \text{defs}; e \} \quad 0 = r_0 \mapsto \text{o}_0 \dots r_n \mapsto \text{o}_n \quad H' = H, \text{o}' \mapsto 0}{H[\![\text{ver} \leftarrow (\text{o}_0 \dots \text{o}_n)]\!]_{\text{o}} \Rightarrow H'[\![\{e\}_{\text{o}'}]\!]_{\text{o}}} \quad [\text{DAPP1}]
\end{array}$$

$$\begin{array}{c}
H' = H \setminus \text{o}'' \\
\frac{H' = H \setminus \text{o}''}{H[\![\{ \text{o}' \}_{\text{o}''}]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad [\text{DAPP2}]
\end{array}
\quad
\begin{array}{c}
F(f)(i) = \text{ver} \\
e = \text{ver} \leftarrow (\text{o}_0 \dots \text{o}_n) \\
\frac{F(f)(i) = \text{ver} \quad e = \text{ver} \leftarrow (\text{o}_0 \dots \text{o}_n)}{H[\![f.i(\text{o}_0 \dots \text{o}_n)]\!]_{\text{o}} \Rightarrow H[\![e]\!]_{\text{o}}} \quad [\text{DCAL}]
\end{array}$$

$$\begin{array}{c}
H(\text{o}') = \mathbf{vec}(i_0 \dots i_n) \\
\text{refs}(H, \text{o}') \leq 1 \\
\frac{H(\text{o}'') = j \quad H(\text{o}''') = i \quad 0 = \mathbf{vec}(i_0 \dots i_{j-1}, i, i_{j+1} \dots i_n) \quad H' = H, \text{o}' \mapsto 0}{H[\![\text{o}'[\text{o}''] = \text{o}''']\!]_{\text{o}} \Rightarrow H'[\![\text{o}''']\!]_{\text{o}}} \quad [\text{DWR1}]
\end{array}
\quad
\begin{array}{c}
H(\text{o}') = \mathbf{vec}(i_0 \dots i_n) \\
\text{o}''' \text{ fresh} \quad H(\text{o}'') = j \\
\frac{H' = H, \text{o}'' \mapsto i_j}{H[\![\text{o}'[\text{o}'']]\!]_{\text{o}} \Rightarrow H'[\![\text{o}''']\!]_{\text{o}}} \quad [\text{DREA}]
\end{array}$$

$$\begin{array}{c}
\text{o}'' = \text{env}(H(\text{o}')) \\
\text{o}''' = H(\text{o}'')(\text{x}) \\
\frac{\text{o}'' = \text{env}(H(\text{o}')) \quad \text{o}''' = H(\text{o}'')(\text{x})}{H[\![\text{o}'\$x]\!]_{\text{o}} \Rightarrow H[\![\text{o}''']\!]_{\text{o}}} \quad [\text{DRRD}]
\end{array}
\quad
\begin{array}{c}
\text{o}'' = \text{env}(H, \text{o}') \quad 0 = H(\text{o}'), \text{x} \mapsto \text{o}'' \\
\frac{\text{o}'' = \text{env}(H, \text{o}') \quad 0 = H(\text{o}'), \text{x} \mapsto \text{o}'' \quad H' = H, \text{o}'' \mapsto 0}{H[\![\text{o}'\$x = \text{o}''']\!]_{\text{o}} \Rightarrow H'[\![\text{o}''']\!]_{\text{o}}} \quad [\text{DRWR}]
\end{array}$$

$$\begin{array}{c}
H[\![e]\!]_{\text{o}'} \Rightarrow H'[\![e]\!]_{\text{o}'} \\
\frac{H[\![e]\!]_{\text{o}'} \Rightarrow H'[\![e]\!]_{\text{o}'} \quad H[\![\{e\}_{\text{o}'}]\!]_{\text{o}} \Rightarrow H'[\![\{e'\}_{\text{o}'}]\!]_{\text{o}}}{H[\![\{e\}_{\text{o}'}]\!]_{\text{o}} \Rightarrow H'[\![\{e'\}_{\text{o}'}]\!]_{\text{o}}} \quad [\text{DCRX1}]
\end{array}
\quad
\begin{array}{c}
H[\![e]\!]_{\text{o}'} \Rightarrow H'[\![e]\!]_{\text{o}'} \\
\frac{H[\![e]\!]_{\text{o}'} \Rightarrow H'[\![e]\!]_{\text{o}'} \quad H[\![\{e\}_{\text{o}'}^{\text{o}''}]\!]_{\text{o}} \Rightarrow H'[\![\{e'\}_{\text{o}'}^{\text{o}''}]\!]_{\text{o}}}{H[\![\{e\}_{\text{o}'}^{\text{o}''}]\!]_{\text{o}} \Rightarrow H'[\![\{e'\}_{\text{o}'}^{\text{o}''}]\!]_{\text{o}}} \quad [\text{DCRX2}]
\end{array}$$

Fig. 10. Dynamic Semantics

The $H \vdash o \leq_{dyn} t$ relation is defined as follows:

$$\frac{}{H \vdash \text{undef} \leq_{dyn} t} \quad \frac{H(o) = \mathbf{vec}(_) \quad \mathbf{v(I)}! \leq t \ \& \ \mathbf{f}}{H \vdash o \leq_{dyn} t} \quad \frac{H(o) = i \quad \mathbf{I}! \leq t \ \& \ \mathbf{f}}{H \vdash o \leq_{dyn} t} \quad \frac{H(o) = \langle _, k, _ \rangle \quad \mathbf{ks}! \leq t}{H \vdash o \leq_{dyn} t}$$

5.3 Type Safety

If a program is well-typed and well-flowed, then its reduction cannot get stuck (although it can reduce to `undef`). This section overviews the type safety proof and the invariants preserved throughout execution; a detailed proof appears in Appendix A (Theorem A.29).

THEOREM (TYPE SAFETY). *If every version in the function table F is well-defined (i.e. well-typed and well-flowed), and if $\emptyset \vdash \llbracket f.i() \rrbracket : t$, then either $\emptyset \llbracket f.i() \rrbracket$ diverges by the small step semantics, or it reduces to $H' \llbracket o' \rrbracket$ for some H' and o' such that $H' \vdash o' \leq_{dyn} t$.*

The proof is decomposed into two parts: one for the type checker (Appendix A.2) and one for the flow analysis (Appendix A.3). Combining their safety results yields the theorem above.

Safety of the type checker. The type checker alone does not guarantee type safety: it does not prevent reading or consuming a register before initialization, or after it has been consumed. The flow analysis handles those properties. The safety theorem for the type checker therefore applies to a modified dynamic semantics with two additional rules:

$$\begin{array}{c} \text{[DREGUNDEF]} \quad \frac{r \notin \text{dom}(H(o))}{H \llbracket r \rrbracket_o \implies H \llbracket \text{undef} \rrbracket_o} \quad \text{[DCONUNDEF]} \quad \frac{r \notin \text{dom}(H(o))}{H \llbracket \text{consume } r \rrbracket_o \implies H' \llbracket \text{undef} \rrbracket_o} \end{array}$$

Because `undef` propagates to the top level, these rules treat an invalid register access as an unchecked runtime error rather than a stuck reduction.

The proof is decomposed into two lemmas: *type preservation* (reduction steps preserve typeability) and *progress* (a well-typed expression can always be reduced unless it is already a reference o).

Proving these lemmas requires extending the type system to handle transient expressions that arise during reduction—references o and context switches $\{e\}_{o'}$ and $\{e\}_{o'}^{o''}$. The heap H is given as an additional input so that references can be typed according to it, and the type environment Γ is replaced by a mapping Δ from scopes o to type environments Γ .

Type preservation relies on several invariants that must hold throughout execution:

Typed environments Every environment in the heap H has an associated type environment in Δ .

Valid types Every environment in H is compatible with its associated type environment in Δ (i.e. its bindings have the right types).

Typeable promises For each promise in H , the captured expression is typeable and has the right type.

Valid promise scoping Promises cannot escape their scope of creation (the environment they capture must be on the call stack).

Valid borrowing Any reference in H with an owned type according to Δ cannot appear elsewhere in H , except in a more recent environment (higher on the call stack) where it has a borrowed type.

The first three invariants state that all the variables in any runtime environment must match the types of our typing environment Δ . The two last invariants are used to ensure ownership is respected, so that DWRI will never fail (only vectors with a reference count of at most 1 can be

```

fun main#1 {
  () → I {
    reg t:v(I)
    reg p:p(v(I))
    t = [1, 2, 3]
    p = prom<v(I)>{ t }
    bump#1.1(p)
  }
}

fun bump#1 {
  (reg p:p(v(I))) → I {
    reg v:v(I)o
    reg i:I
    v = dup(force p)
    v[0] = 42
    v[0]
  }
}

```

Fig. 11. A small FIR program

assigned). In particular, the *valid promise scoping* property ensures that the scope of a function is unreachable once this function returns (it cannot be leaked in a promise).

Formal definitions appear in Appendix A.2 (Definition A.7), together with proofs of type preservation and progress.

Safety of the flow analysis. The flow analysis guarantees that registers are neither read nor consumed before initialization, nor accessed after consumption. Equivalently, for any well-flowed expression the rules DREGUNDEF and DCONUNDEF never apply. Combining this with the type checker's guarantees yields type safety for the dynamic semantics of Section 5.2.

As with the type checker, correctness is decomposed into two lemmas: *preservation of well-flowedness* (a well-flowed expression remains well-flowed after reduction) and *correctness of well-flowedness* (DREGUNDEF and DCONUNDEF cannot apply to well-flowed expressions).

As before, the flow analysis must be extended to handle transient expressions. Actions A are replaced by mappings $\mathbb{A}$ from scopes o to actions A , and rules are added for references and context switches. These extensions and their proofs appear in Appendix A.3.

5.4 Example

We now demonstrate how the above rules apply to an example, Fig. 11. The program allocates a vector, wraps it in a promise, then calls a function with the promise. The function forces the promise, duplicates the resulting vector, mutates the copy, then returns the updated element.

Typing. In `main#1`: by LIT and VEC, the expression `[1, 2, 3]` has type $v(I)f!$. The assignment `t = [1, 2, 3]` is accepted by ASS, since a fresh vector may be stored in a shared register. By VAR, `t` has type $v(I)s!$, and by PRO, `prom<v(I)>{ t }` has type $p(v(I))s!$. The static call `bump#1.1(p)` is accepted by CAL.

In `bump#1`: forcing `p` yields a shared vector (via FOR), and duplicating this yields a fresh vector (via DUP). Assigning this fresh value to `v:v(I)o!` is valid because (only) a fresh value may be stored in an owned register. The update `v[0] = 42` is accepted by WRI, since `v` is owned, and the read `v[0]` yields an integer scalar (via REA).

VER validates both versions with signatures $\vec{\rightarrow} Is!$ and $p^-(v(I)s!)s! \vec{\rightarrow} Is!$.

Reflection. Neither version contains any reflective operations, and neither signature nor the promise in `main#1` contains a reflective annotation. If the promise in `main#1` were annotated as reflective, consequently: both register and parameter `p` would need to be annotated as reflective to satisfy type-checking (by PRO, the promise type would be $p^+(v(I))$, which is not a subtype of $p(v(I))$); `bump#1` would need to be reflective to satisfy RFOR, RSEQ, and RVER (the force is reflective, which propagates to the entire body of `bump#1` and its signature); then `main#1` would need to be reflective to call `bump#1` according to RCAL, RSEQ, and RVER (via the same propagation).

Flow analysis. The body of `main#1` has action $(\emptyset, \{t, p\}, \emptyset, \{t\})$. The assignment to `t` writes `t`. The promise constructor reads `t` and captures it, and the assignment to `p` writes `p`. The final call reads `p`. After composition, the read set for the whole body is empty because both registers are assigned before they are read. The capture set records that `t` appears in the body of the promise created in `main#1`. No registers are consumed.

The body of `bump#1` has action $(\{p\}, \{v\}, \emptyset, \emptyset)$. The version reads its parameter `p`, assigns `v`, and neither consumes nor captures any register. Since the read set is included in the parameter set, $F[\llbracket \text{bump\#1} \rrbracket]$ holds.

Reduction. Starting from `main#1.1()`, reduction proceeds as follows. `DCAL`, `DAPP1`, and `DAPP2` allocate an environment $\mathbf{o}_0$ for `main#1`. `DVEC` allocates a fresh vector reference $\mathbf{o}_1$ containing $[1, 2, 3]$. `DASS` extends $\mathbf{o}_0$ with $t \mapsto \mathbf{o}_1$. `DPRO` allocates a promise reference $\mathbf{o}_2$ whose body is `t` and whose captured environment is $\mathbf{o}_0$:

$$H(\mathbf{o}_2) = \langle t, \mathbf{p}^-(\mathbf{v}(\mathbf{I})\mathbf{s}!), \mathbf{o}_0 \rangle.$$

The assignment to `p`, via `DASS`, extends $\mathbf{o}_0$ with $p \mapsto \mathbf{o}_2$.

We omit most reduction in `bump#1.1` for brevity. Notably, the force expression, evaluated via `DFOR2`, evaluates the promise body in its captured environment $\mathbf{o}_0$, *not* in the callee's environment; so the read of `t` within the promise, via `DVAR`, returns $\mathbf{o}_1$.

6 Conclusion

This paper presents `FiR`, a typed high-level intermediate representation for compiling dynamic languages. `FiR` makes dynamic behaviors explicit—reflective accesses, dispatched calls, promise evaluation, and vector copies—while combining gradual typing with ownership tracking. This design enables classical optimizations such as specialization, inlining, scope elision, and copy elimination, allowing compilers to reason locally about transformations without compromising correctness.

We formalized `FiR`'s operational semantics along with a static semantics comprising a type system, flow analysis, and reflection analysis, and proved type safety. Although `FiR` is informed by `R`'s semantics, its core abstractions—tracking reflection, ownership, and delayed evaluation—apply to any dynamic language with similar features.

We are building a new JIT compiler for `R` using an extended variant of `FiR`. Although we do not yet have the necessary perspective to assess the long-term benefits of `FiR`, our implementation validates that `FiR` is effective in a real-world compiler: we can compile `R` into native code, and have straightforwardly implemented optimizations including variable elision and inlining.

Acknowledgments

This work was supported by the Czech Ministry of Education, Youth and Sports under program ERC-CZ, grant agreement LL2325, as well as by the Czech Science Foundation Grant No. 23-07580X.

References

- [1] John Aycock. 2003. A Brief History of Just-In-Time. *ACM Computing Surveys (CSUR)* 35, 2 (2003). doi:10.1145/857076.857077
- [2] Aurele Barriere, Olivier Flückiger, Sandrine Blazy, David Pichardie, and Jan Vitek. 2021. Formally Verified Speculation and Deoptimization in a JIT Compiler. *Proc. ACM Program. Lang.* 5, POPL (2021). doi:10.1145/3434327
- [3] Jeff Bezanson, Jiahao Chen, Ben Chung, Stefan Karpinski, Viral B. Shah, Jan Vitek, and Lionel Zoubritzky. 2018. Julia: Dynamism and Performance Reconciled by Design. *Proc. ACM Program. Lang.* 2, OOPSLA (2018). doi:10.1145/3276490
- [4] Craig Chambers, David Ungar, and Elgin Lee. 1991. An Efficient Implementation of SELF, a Dynamically-Typed Object-Oriented Language Based on Prototypes. In *Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*. doi:10.1145/117954.117959

- [5] Clifford Click. 1995. *Combining Analyses, Combining Optimizations*. Ph.D. Dissertation. Rice University. <https://hdl.handle.net/1911/16837>
- [6] Jan de Mooij, Matthew Gaudet, Iain Ireland, Nathan Henderson, and J. Nelson Amaral. 2023. CacheIR: The Benefits of a Structured Representation for Inline Caches. In *Proceedings of the 20th ACM SIGPLAN International Conference on Managed Programming Languages and Runtimes*. doi:10.1145/3617651.3622979
- [7] L. Peter Deutsch and Allan M. Schiffman. 1984. Efficient Implementation of the Smalltalk-80 System. In *Symposium on Principles of Programming Languages (POPL)*. doi:10.1145/800017.800542
- [8] Matthias Felleisen and Robert Hieb. 1992. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science* 103, 2 (1992). doi:10.1016/0304-3975(92)90014-7
- [9] Olivier Flückiger, Guido Chair, Ming-Ho Yee, Jan Jecmen, Jakob Hain, and Jan Vitek. 2020. Contextual Dispatch for Function Specialization. *Proc. ACM Program. Lang.* 4, OOPSLA (2020). doi:10.1145/3428288
- [10] Olivier Flückiger, Guido Chari, Jan Jecmen, Ming-Ho Yee, Jakob Hain, and Jan Vitek. 2019. R melts brains: an IR for first-class environments and lazy effectful arguments. In *International Symposium on Dynamic Languages (DLS)*. doi:10.1145/3359619.3359744
- [11] Olivier Flückiger, Gabriel Scherer, Ming-Ho Yee, Aviral Goel, Amal Ahmed, and Jan Vitek. 2018. Correctness of speculative optimizations with dynamic deoptimization. *Proc. ACM Program. Lang.* 2, POPL (2018). doi:10.1145/3158137
- [12] Stephen Freund and John Mitchell. 2003. A Type System for the Java Bytecode Language and Verifier. *Journal of Automated Reasoning* 30, 3 (May 2003). doi:10.1023/A:1025011624925
- [13] Michael Gottesman, Joe Groff, and John McCall. 2024. *borrowing and consuming parameter ownership modifiers*. RFC 0377. Apple. <https://github.com/swiftlang/swift-evolution/blob/main/proposals/0377-parameter-ownership-modifiers.md>
- [14] Andrei Homescu and Alex Şuhan. 2011. HappyJIT: a tracing JIT compiler for PHP. In *Symposium on Dynamic Languages (DLS)*. doi:10.1145/2047849.2047854
- [15] Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. 2001. Featherweight Java: a minimal core calculus for Java and GJ. *ACM Trans. Program. Lang. Syst.* (2001). doi:10.1145/503502.503505
- [16] Chris Lattner, Mehdi Amini, Uday Bondhugula, Albert Cohen, Andy Davis, Jacques Pienaar, River Riddle, Tatiana Shpeisman, Nicolas Vasilache, and Oleksandr Zinenko. 2020. MLIR: A Compiler Infrastructure for the End of Moore's Law. arXiv:2002.11054 [cs.PL] <https://arxiv.org/abs/2002.11054>
- [17] Tim Lindholm, Frank Yellin, Gilad Bracha, and Alex Buckley. 2015. The Java® Virtual Machine Specification, Java SE 8 Edition. <https://docs.oracle.com/javase/specs/jvms/se8/html/>
- [18] Niko Matsakis. 2016. Introducing MIR. <https://blog.rust-lang.org/2016/04/19/MIR/>
- [19] Greg Morrisett, David Walker, Karl Cray, and Neal Glew. 1999. From system F to typed assembly language. *ACM Transactions on Programming Languages and Systems* 21 (1999). doi:10.1145/319301.319345
- [20] James Noble, John Potter, and Jan Vitek. 1998. Flexible Alias Protection. In *European Conference on Object-Oriented Programming (ECOOP)*. doi:10.1007/BFb0054091
- [21] Michael Paleczny, Christopher Vick, and Cliff Click. 2001. The Java HotSpot Server Compiler. In *Proceedings of the Java Virtual Machine Research and Technology Symposium (JVM)*. doi:10.5555/1267847.1267848
- [22] Jeremy G. Siek and Walid Taha. 2007. Gradual Typing for Objects. In *European Conference on Object-Oriented Programming (ECOOP)*. doi:10.1007/978-3-540-73589-2_2
- [23] Tobias Wrigstad. 2006. *Ownership-Based Alias Management*. Doctoral Dissertation. Department of Information Technology, Uppsala University. <https://wrigstad.com/papers/thesis.pdf>
- [24] Tobias Wrigstad, Francesco Zappa Nardelli, Sylvain Lebesne, Johan Östlund, and Jan Vitek. 2010. Integrating typed and untyped code in a scripting language. In *Symposium on Principles of Programming Languages (POPL)*. doi:10.1145/1706299.1706343
- [25] Thomas Würthinger, Andreas Wöß, Lukas Stadler, Gilles Duboscq, Doug Simon, and Christian Wimmer. 2012. Self-optimizing AST interpreters. In *Symposium on Dynamic Languages (DLS)*. doi:10.1145/2384577.2384587

A Proofs

A.1 Full Dynamic Semantics

The full small-step semantics is the same as in Section 5.2, with the two extra rules [DREGUNDEF] and [DCONUNDEF] (in blue). Reduction steps are written $H[e]_o \rightsquigarrow H'[e']_o$. We write $H[e]_o \rightsquigarrow^* H'[e']_o$ for a reduction of any number of steps, and $H[e]_o \rightsquigarrow^\infty$ for a diverging reduction.

$$\begin{array}{c}
\text{[DLit]} \frac{o' \text{ fresh} \quad H' = H, o' \mapsto i}{H[i]_o \rightsquigarrow H'[o']_o} \quad \text{[DPRO]} \frac{o' \text{ fresh} \quad H' = H, o' \mapsto \langle e, p(t \text{ fx}, o) \rangle}{H[\text{prom}^{fx} \langle t \rangle \{e\}]_o \rightsquigarrow H'[o']_o} \quad \text{[DVec]} \frac{o' \text{ fresh} \quad H(o_1) = i_1 \dots H(o_n) = i_n}{H' = H, o' \mapsto \text{vec}(i_1 \dots i_n)} \frac{H[\text{vec}(o_1 \dots o_n)]_o \rightsquigarrow H'[o']_o}{} \\
\text{[DVar]} \frac{o' = H(o)(v)}{H[v]_o \rightsquigarrow H'[o']_o} \quad \text{[DNVarUNDEF]} \frac{x \notin \text{dom}(H(o))}{H[x]_o \rightsquigarrow H[\text{undef}]_o} \quad \text{[DCONUNDEF]} \frac{r \notin \text{dom}(H(o))}{H[\text{consume } r]_o \rightsquigarrow H'[\text{undef}]_o} \\
\text{[DDup]} \frac{o'' \text{ fresh} \quad H' = H, o'' \mapsto H(o')}{H[\text{dup } o']_o \rightsquigarrow H'[o'']_o} \quad \text{[DCon]} \frac{o' = H(o)(r) \quad H' = H, o \mapsto H(o) \setminus r}{H[\text{consume } r]_o \rightsquigarrow H'[o']_o} \quad \text{[DREGUNDEF]} \frac{r \notin \text{dom}(H(o))}{H[r]_o \rightsquigarrow H[\text{undef}]_o} \\
\text{[DFor1]} \frac{H(o') = \langle o'', k, o''' \rangle}{H[\text{force } o']_o \rightsquigarrow H[o'']_o} \quad \text{[DFor2]} \frac{H(o') = \langle e, k, o'' \rangle}{H[\text{force } o']_o \rightsquigarrow H'[\{e\}_{o'''}]_o} \quad \text{[DFor3]} \frac{H' = H, o''' \mapsto \langle o', k, o'' \rangle}{H[\{o'\}_{o'''}]_o \rightsquigarrow H'[o']_o} \\
\text{[DApP1]} \frac{o' \text{ fresh} \quad \text{ver} = (r_0 : t_0 \dots r_n : t_n) \xrightarrow{fx} t \{ \text{defs}; e \} \quad 0 = r_0 \mapsto o_0 \dots r_n \mapsto o_n \quad H' = H, o' \mapsto 0}{H[\text{ver} \leftarrow (o_0 \dots o_n)]_o \rightsquigarrow H'[\{e\}_{o'}]_o} \quad \text{[DApP2]} \frac{H' = H \setminus o''}{H[\{o'\}_{o''}]_o \rightsquigarrow H'[o']_o} \\
\text{[DCal]} \frac{F(f)(i) = \text{ver} \quad e = \text{ver} \leftarrow (o_0 \dots o_n)}{H[f.i(o_0 \dots o_n)]_o \rightsquigarrow H[e]_o} \quad \text{[DDis]} \frac{\text{smallest } i \text{ such that} \quad \text{sig}(F(f)(i)) = t_0 \dots t_n \xrightarrow{fx} t \quad t_0 \dots t_n \xrightarrow{fx} t \leq_{\text{sig}} \text{Sig} \quad H \vdash o_0 \leq_{\text{dyn}} t_0 \dots H \vdash o_n \leq_{\text{dyn}} t_n \quad e = f.i(o_0 \dots o_n)}{H[f <\text{sig}> (o_0 \dots o_n)]_o \rightsquigarrow H[e]_o} \\
\text{[DWriUNDEF]} \frac{H(o_1) = \text{vec}(i_0 \dots i_n) \quad H(o_2) = j \quad j < 0 \text{ or } j > n}{H[o_1[o_2]]_o \rightsquigarrow H[\text{undef}]_o} \quad \text{[DWri]} \frac{H(o') = \text{vec}(i_0 \dots i_n) \quad \text{refs}(H, o') \leq 1 \quad H(o'') = j \quad H(o''') = i \quad 0 = \text{vec}(i_0 \dots i_{j-1}, i, i_{j+1} \dots i_n) \quad H' = H, o' \mapsto 0}{H[o'[o'']]_o \rightsquigarrow H'[o''']_o} \\
\text{[DREAUNDEF]} \frac{H(o_1) = \text{vec}(i_0 \dots i_n) \quad H(o_2) = j \quad j < 0 \text{ or } j > n}{H[o_1[o_2]]_o \rightsquigarrow H[\text{undef}]_o} \quad \text{[DREA]} \frac{H(o') = \text{vec}(i_0 \dots i_n) \quad o'' \text{ fresh} \quad H(o'') = j \quad H' = H, o'' \mapsto i_j}{H[o' [o'']]_o \rightsquigarrow H'[o'']_o} \\
\text{[DAss]} \frac{0 = H(o), v \mapsto o' \quad H' = H, o \mapsto 0}{H[v = o']_o \rightsquigarrow H'[o']_o} \quad \text{[DRrd]} \frac{o'' = \text{env}(H(o')) \quad o''' = H(o'')(x)}{H[o' \$ x]_o \rightsquigarrow H[o''']_o} \quad \text{[DRwb]} \frac{o''' = \text{env}(H, o') \quad 0 = H(o'), x \mapsto o'' \quad H' = H, o'' \mapsto 0}{H[o' \$ x = o'']_o \rightsquigarrow H'[o'']_o} \\
\text{[DCas]} \frac{H \vdash o \leq_{\text{dyn}} t \Rightarrow o' = o \quad H \vdash o \not\leq_{\text{dyn}} t \Rightarrow o' = \text{undef}}{H[o' \text{ as } t]_o \rightsquigarrow H[o']_o} \quad \text{[DSeq]} \frac{}{H[o' ; e]_o \rightsquigarrow H[e]_o} \quad \text{[DCTX1]} \frac{H[e]_{o'} \rightsquigarrow H'[e']_{o'}}{H[\{e\}_{o'}]_o \rightsquigarrow H'[\{e'\}_{o'}]_o} \\
\text{[DCTX2]} \frac{H[e]_{o'} \rightsquigarrow H'[e']_{o'}}{H[\{e\}_{o'}]_o \rightsquigarrow H'[\{e'\}_{o'}]_o} \quad \text{[DCTX3]} \frac{H[e]_o \rightsquigarrow H'[\text{undef}]_o}{H[\mathcal{E}[e]]_o \rightsquigarrow H'[\text{undef}]_o} \quad \text{[DCTX4]} \frac{H[e]_o \rightsquigarrow H'[e']_o}{H[\mathcal{E}[e]]_o \rightsquigarrow H'[\mathcal{E}[e']]_o}
\end{array}$$

A.2 Proofs for the Static Semantics

We prove type safety for the small-step semantics defined above. Due to the rules [DREGUNDEF] and [DConUNDEF], this theorem does not ensure correct register initialization; that property is established by the flow analysis (Section A.3).

Definition A.1 (Well-typed version). A version ver is well-typed if and only if there exists $t_0 \dots t_n, t$ such that $\llbracket \text{ver} \rrbracket : t_0 \dots t_n \xrightarrow{fx} t$ holds.

In the following, we assume that every version ver in our function table F is well-typed.

The proof proceeds via a type preservation lemma (Lemma A.9) and a progress lemma (Lemma A.10). Both require extending the type system to handle intermediate results—references o , version scopes $\{e\}_o$, and promise scopes $\{e\}_o^{o'}$ —by taking the current heap and a heap descriptor as additional inputs:

Definition A.2 (Heap descriptor). We call heap descriptor Δ a partial mapping from references o to type environments Γ .

Definition A.3 (Dynamic kind). We define the dynamic kind of a runtime value o as follows:

$$\text{dynkind}(i) = \mathbf{I} \quad \text{dynkind}(\text{vec}(i_0 \dots i_n)) = \mathbf{v}(\mathbf{I}) \quad \text{dynkind}(\langle _, k, _ \rangle) = k$$

$$\begin{array}{c} \text{[VAR]} \frac{t = \Delta(o)(v)}{\Delta, H \vdash_o \llbracket v \rrbracket : t} \quad \text{[UNDEF]} \frac{}{\Delta, H \vdash_o \llbracket \text{undef} \rrbracket : t} \\ \\ \text{[DYN]} \frac{\begin{array}{l} k = \text{dynkind}(H(o')) \quad H(o') = \langle _, _, o'' \rangle \Rightarrow o'' \in \text{dom}(H) \\ ox = \mathbf{s} \Rightarrow \text{owner}(H, \Delta, o') = \{\} \quad ox = \mathbf{o} \Rightarrow \text{owner}(H, \Delta, o') = \{o\} \quad ox = \mathbf{f} \Rightarrow o' \text{ fresh} \end{array}}{\Delta, H \vdash_o \llbracket o' \rrbracket : k \text{ ox} !} \end{array}$$

$$\begin{array}{c} \Delta, H \vdash_{o'} \llbracket e \rrbracket : t' \\ t = \begin{cases} t' & \text{if shared}(t') \text{ or fresh}(t') \\ t' \& \mathbf{f} & \text{if owned}(t') \end{cases} \\ \text{[ASC]} \frac{\text{value}(t) \quad \text{wf}(t)}{\Delta, H \vdash_o \llbracket \{e\}_{o'} \rrbracket : t} \quad \text{[PSC]} \frac{\begin{array}{c} \Delta, H \vdash_{o'} \llbracket e \rrbracket : t \\ \text{value}(t) \quad \text{shared}(t) \end{array}}{\Delta, H \vdash_o \llbracket \{e\}_{o'}^{o''} \rrbracket : t} \end{array}$$

where $\text{owner}(H, \Delta, o) = \{o' \mid \exists r. H(o')(r) = o \text{ and } \text{owned}(\Delta(o')(r))\}$ and o' fresh holds when o' does not appear in H except on the left-hand side of a binding, and has no other occurrence in the expression.

The remaining typing rules carry H and o recursively and are otherwise identical to those in Section 5.1. The rule [VER] needs no additional parameters: since we do not reduce under versions, no reference, version scope, or promise scope can appear in a version body, so [VER] is reused as is:

$$\begin{array}{c} \Gamma = \text{rdefs}, \text{defs} \quad \text{wf}(\Gamma) \quad \Gamma \vdash \llbracket e \rrbracket : t' \quad \text{types}(\text{rdefs}) = t_0 \dots t_n \quad \Gamma \vdash \mathbf{R}\llbracket e \rrbracket = \text{fx} \\ t = \begin{cases} t' & \text{if shared}(t') \text{ or fresh}(t') \\ t' \& \mathbf{f} & \text{if owned}(t') \end{cases} \\ \text{[VER]} \frac{\text{value}(t) \quad \text{wf}(t)}{\llbracket (\text{rdefs}) \xrightarrow{fx} t\{\text{defs}; e\} \rrbracket : t_0 \dots t_n \xrightarrow{fx} t} \end{array}$$

PROPOSITION A.4. *If $\Gamma \vdash \llbracket e \rrbracket : t$ (typing judgement from Section 5.1), then for any o, H, Δ such that $\Delta(o) = \Gamma$, we have $\Delta, H \vdash_o \llbracket e \rrbracket : t$.*

PROOF. Straightforward. $\square$

Definition A.5 (Dynamic/static environment matching). Let H be a heap, o a reference, Γ a type environment. We say that H matches Γ at reference o if and only if:

- $H(o) = E$.
- $\forall r \in \text{dom}(\Gamma) \cap \text{dom}(E). \Gamma(r) \&! = \Gamma(r) \Rightarrow \text{dynkind}(E(r)) \leq \text{kind}(\Gamma(r))$.
- $\forall r \in \text{dom}(\Gamma) \cap \text{dom}(E). \text{owned}(\Gamma(r)) \Rightarrow \text{refs}(H, E(r)) \leq 1$.

We define call contexts C as follows:

$$C ::= \dots \mid \{C\}_o^\circ$$

where $\dots$ includes all the cases of the grammar of $\mathcal{E}$ (defined in Section 5.2).

Definition A.6 (Call stack). Let e be an expression. The call stack of e is the longest sequence of references $o_1, \dots, o_n$ such that $e = C_1[\dots C_n[\{e'\}_{o_n}] \dots]_{o_1}$ for some $C_1, \dots, C_n$ and e' .

A call stack (a sequence of references) is written $\vec{o}$. The concatenation of two call stacks $\vec{o}$ and $\vec{o}'$ is written $\vec{o}, \vec{o}'$.

Definition A.7 (Heap validity). The heap H is valid for the heap descriptor Δ and call stack $o_1, \dots, o_n$ if and only if:

- Typed environments** For every binding $(o \mapsto E) \in H$, there exists a unique i such that $o_i = o$.
- Valid types** For every i , $H \setminus \{o_{i+1}, \dots, o_n\}$ matches $\Delta(o_i)$ at reference o_i .
- Typeable promises** For every $(o' \mapsto \langle e, k, o \rangle) \in H$, if $o \in \text{dom}(H)$ then $k = \mathbf{p}^{fx}(k' \mathbf{s} cx)$ for some k' , cx and fx , and $\Delta, H \vdash_o \llbracket e \rrbracket : k' \mathbf{s} cx$.
- Valid promise scoping** For every i , for every $(v \mapsto o') \in H(o_i)$ with $H(o') = \langle _, k, o'' \rangle$, we have $o'' = o_{i'}$ for some $i' \leq i$.
- Valid borrowing** For every i and $i' > i$, for every $r \in \text{dom}(\Delta(o_i)) \cap \text{dom}(H(o_i))$, for every $v \in \text{dom}(H(o_{i'}))$: $(\text{owned}(\Delta(o_i)(r)) \text{ and } (H(o_i)(r) = H(o_{i'})(v))) \Rightarrow \text{borrowed}(\Delta(o_{i'})(v))$.

Definition A.8 (Quadruplet validity). A quadruplet $(H, \Delta, \vec{o}, e)$ is valid if and only if H is valid for the heap descriptor Δ and call stack $(\vec{o}, \vec{o}')$, with $\vec{o}'$ being the call stack of e .

LEMMA A.9 (PRESERVATION OF TYPE AND VALIDITY). Let $(H, \Delta, \vec{o}, e)$ be a valid quadruplet. If $\Delta, H \vdash_o \llbracket e \rrbracket : t$ and $H \llbracket e \rrbracket_o \rightsquigarrow H' \llbracket e' \rrbracket_o$, then there exists Δ' such that $\Delta', H' \vdash_o \llbracket e' \rrbracket : t$ and $(H', \Delta', \vec{o}, e')$ is valid.

PROOF. We proceed by structural induction on the derivation $\Delta, H \vdash_o \llbracket e \rrbracket : t$.

If the root of the derivation is a [SUB], we can conclude by using the induction hypothesis.

Other rules are structural, thus we can match on the syntax of e (we only include the most interesting cases):

- o' Impossible case as no reduction step can apply on o' .
- $\{o_1\}_{o_2}$ The reduction rule [DAPP2] applies, we thus have $e' = o_1$ and $H' = H \setminus o_2$. We can conclude with a [DYN] rule as $o_1 \neq o_2$ ($H(o_2)$ is an environment E , while $H(o_1)$ is not). Note that we use the *valid promise scoping* property to ensure the new heap $H \setminus o_2$ does not have accessible "orphan" promises.
- $\{e_1\}_{o_2}$ The reduction rule [DCTX1] applies, we conclude by applying the induction hypothesis on e_1 .
- $\{o_1\}_{o_2}^{o_3}$ The reduction rule [DFOR3] applies, we thus have $e' = o_1$ and $H' = H, o_3 \mapsto \langle o_1, k, o_2 \rangle$, we can thus conclude with a [DYN] rule.
- $\{e_1\}_{o_2}^{o_3}$ The reduction rule [DCTX2] applies, we conclude by applying the induction hypothesis on e_1 .

- v If the reduction rule [DVAR1] applies, we can thus conclude with a [DYN] rule (using the *valid types* property and the *valid promise scoping* property). Otherwise, if the rule [DVAR2] or [DVAR3] applies, we conclude with a [UNDEF] rule.
- o' **as** t' The reduction rule [DCAS] applies. If the dynamic cast failed, we get $e' = \text{undef}$, in which case we can conclude with the rule [UNDEF]. Otherwise, from the typing derivation of e , we can extract a typing derivation for o' of a type t'' of the same ownership as t' . As the cast succeeded, we know that $\text{dynkind}(H(o')) \leq \text{kind}(t')$, and we can thus derive the type t' for o' .
- force** o' The reduction rule [DFOR1] or [DFOR2] applies. In both cases, we can conclude using the *typeable promises* property.
- ver** $\leftarrow (o_0 \dots o_n)$ The rule [DAPP1] applies. We must type the body of **ver**, which can easily be done by extracting the subderivation of the body from $\llbracket \text{ver} \rrbracket : t_0 \dots t_n \xrightarrow{fx} t$ and applying Proposition A.4 to type the body under the new heap descriptor $\Delta, o' \mapsto \Gamma$ with $\Gamma = \text{rdefs}, \text{defs}$.
- f.i** $(o_0 \dots o_n)$ The rule [DCAL] applies. We can conclude using the fact that versions in F are well-typed, and using Proposition A.4.
- f <sig>** $(o_0 \dots o_n)$ The rule [DDIS] applies. Using the fact that the version selected by the [DDIS] rule has a smaller signature than the version used in our typing derivation, this case can be concluded similarly to the previous case.
- o' \$x = o''** The rule [DRWR] applies. The typing derivation of e gives the guarantee that o'' is not a promise, which can be used to guarantee the preservation of the *valid promise scoping* property, and that it is shared, which, thanks to the *valid borrowing* property, can be used to guarantee that the *valid types* property is preserved.

□

LEMMA A.10 (PROGRESS). *Let $(H, \Delta, \vec{o}, e)$ be a valid quadruplet. If $\Delta, H \vdash_o \llbracket e \rrbracket : t$, then either e is a reference o' or $H \llbracket e \rrbracket_o \rightsquigarrow H' \llbracket e' \rrbracket_o$ for some H' and e' .*

PROOF. We proceed by structural induction on the derivation $\Delta, H \vdash_o \llbracket e \rrbracket : t$.

If the root of the derivation is a [SUB], we can conclude by using the induction hypothesis.

Other rules are structural, thus we can match on the syntax of e (we only include the most interesting cases):

- o' Trivial.
- $o_1[o_2] = o_3$ Our typing derivation must end with a [WRI] rule, which requires o_1 to have type $\mathbf{v}(\mathbf{I})\mathbf{o}!$. It yields that $\text{dynkind}(H(o_1)) \leq \mathbf{v}(\mathbf{I})$ and thus $H(o_1)$ is a vector. From $\text{owner}(H, \Delta, o') = \{o\}$ and the *valid types* property, we can deduce $\text{refs}(H, o') \leq 1$, which allows to conclude that either the rule [DWRI] or [DWRIUNDEF] applies.
- force** o' The typing derivation gives a $\mathbf{p}^{fx}(t)\mathbf{s}!$ type to o' , which ensures that $H(o')$ is a promise. Thus, either [DFOR1] or [DFOR2] can be applied.
- force** e' The typing derivation gives a $\mathbf{p}^{fx}(t)\mathbf{s}!$ type to e' , we can thus apply the induction hypothesis and conclude with a [DCTX3] or [DCTX4] rule.
- f <sig>** $(o_0 \dots o_n)$ The typing derivation gives the types $t'_0 \dots t'_n$ to the arguments, with $t'_0 \leq t_0, \dots, t'_n \leq t_n$ and $\text{sig}(F(f)(i)) = t_0 \dots t_n \xrightarrow{fx} t'$ for some i and fx . Thus, using the *valid types* property, we know that the rule [DDIS] can be applied on the version i (or on a version at a smaller position if applicable).

□

THEOREM A.11 (TYPE SAFETY).

If $\emptyset \vdash \llbracket f.i() \rrbracket : \mathbf{t}$, then either $\emptyset \llbracket f.i() \rrbracket_o \rightsquigarrow^\infty$ or $\emptyset \llbracket f.i() \rrbracket_o \rightsquigarrow^* H' \llbracket o' \rrbracket_o$ for some H' and o' such that $H' \vdash o' \leq_{\text{dyn}} \mathbf{t}$.

PROOF. Immediate consequence of Proposition A.4, of the preservation (Lemma A.9) and progress (Lemma A.10). $\square$

We recall that this theorem, as well as all the lemmas in the section, assume that every version ver in our function table F is well-typed.

A.3 Proofs for the Flow Analysis

We prove that the flow analysis ensures correct register initialization: if an expression is well-flowed, the rules [DREGUNDEF] and [DCONUNDEF] cannot apply after any number of steps. Combined with the type safety of the previous section, this yields type safety for the dynamic semantics of Section 5.2, where accessing an uninitialized register results in a stuck reduction.

Definition A.12 (Well-flowed version). A version ver is well-flowed if and only if $F \llbracket \text{ver} \rrbracket$ holds.

Definition A.13 (Well-defined version). A version ver is well-defined if and only if ver is well-typed and well-flowed.

In the following, we assume that every version ver in our function table F is well-defined.

We first generalize the flow analysis so that it applies to intermediate results of a reduction, not only source expressions. We then prove that the action returned by this generalized analysis is preserved by reduction (Lemma A.25), and that a well-flowed expression cannot be reduced using [DREGUNDEF] or [DCONUNDEF] (Lemma A.27).

PROPOSITION A.14 (ASSOCIATIVITY OF COMPOSITION). *The composition operator $;;$ is associative.*

PROOF. Let $A_1 = (R_1, W_1, U_1, C_1)$, $A_2 = (R_2, W_2, U_2, C_2)$, and $A_3 = (R_3, W_3, U_3, C_3)$. The expressions $(A_1 ;; A_2) ;; A_3$ and $A_1 ;; (A_2 ;; A_3)$ are both defined if and only if:

- $U_1 \cap (R_2 \cup W_2 \cup U_2 \cup C_2) = \emptyset$, and
- $(U_1 \cup U_2) \cap (R_3 \cup W_3 \cup U_3 \cup C_3) = \emptyset$, and
- $C_1 \cap U_2 = \emptyset$, and
- $(C_1 \cup C_2) \cap U_3 = \emptyset$

and when defined, are equal to the expression (R', W', U', C') where:

- $R' = R_1 \cup (R_2 \setminus W_1) \cup (R_3 \setminus (W_1 \cup W_2))$
- $W' = W_1 \cup W_2 \cup W_3$
- $U' = U_1 \cup U_2 \cup U_3$
- $C' = C_1 \cup C_2 \cup C_3$

$\square$

Definition A.15. We define a partial order $\leq$ over actions as follows:

$$(R, W, U, C) \leq (R', W', U', C') \Leftrightarrow R \subseteq R', W' \subseteq W, U \subseteq U', C \subseteq C'$$

PROPOSITION A.16 (MONOTONICITY OF COMPOSITION). *The composition operator $;;$ is monotonic with respect to the order $\leq$.*

PROOF. Straightforward. $\square$

Definition A.17 (Heap actions). A heap action $\mathbb{A}$ is a total mapping from references o to actions A . We write $\emptyset$ for the heap action mapping every reference to $(\emptyset, \emptyset, \emptyset, \emptyset)$, and $(o \mapsto A)$ for the heap action mapping o to A and every other reference to $(\emptyset, \emptyset, \emptyset, \emptyset)$.

We extend the composition $;;$ to work on heap actions: $\forall o. (\mathbb{A} ;; \mathbb{A}') (o) = \mathbb{A}(o) ;; \mathbb{A}'(o)$. We also extend the order $\leq$: $\mathbb{A} \leq \mathbb{A}' \Leftrightarrow \forall o. \mathbb{A}(o) \leq \mathbb{A}'(o)$. The properties above about $;;$ still hold.

The flow analysis of Section 5.1 is extended. It now takes a reference o as additional input, and returns heap actions:

$$\begin{array}{ccc} \text{[FDyn]} \frac{}{F[o']_o = \emptyset} & \text{[FASc]} \frac{F[e]_{o'} = \mathbb{A}}{F[\{e\}_{o'}]_o = \mathbb{A}} & \text{[FPSc]} \frac{F[e]_{o'} = \mathbb{A}}{F[\{e\}_{o'}^{o''}]_o = \mathbb{A}} \end{array}$$

The remaining rules carry o recursively and lift each action $\mathbb{A}$ to $o \mapsto \mathbb{A}$; for instance:

$$\text{[FReg]} \frac{\mathbb{A} = o \mapsto \text{read } r}{F[r]_o = \mathbb{A}}$$

The rule [FVER] needs no additional parameters: since we do not reduce under versions, no reference, version scope, or promise scope can appear in a version body, so [FVER] from Section 5.1 is reused as is:

$$\text{[FVER]} \frac{F[e] = (R, W, U, C) \quad R \subseteq \{r_0 \dots r_n\}}{F[(r_0 : t_0 \dots r_n : t_n) \xrightarrow{f_x} t\{\text{defs}; e\}]} \quad \square$$

PROPOSITION A.18. *If $F[e] = \mathbb{A}$, then for any reference o , we have $F[e]_o = \mathbb{A}$ with $\mathbb{A} = o \mapsto \mathbb{A}$.*

PROOF. Follows from the definition. $\square$

Definition A.19 (Delayed action). For an action $\mathbb{A} = (R, W, U, C)$, we define the action $\text{delay}(\mathbb{A})$ as follows:

$$\text{delay}(\mathbb{A}) = (R, \emptyset, U, R \cup W \cup U \cup C)$$

PROPOSITION A.20. *For any action $\mathbb{A}$, $\mathbb{A} \leq \text{delay}(\mathbb{A})$.*

PROOF. Straightforward. $\square$

Delaying some actions is a way to make them commutative, as stated by the proposition below.

PROPOSITION A.21 (COMMUTATIVITY OF COMPOSITION ON DELAYED ACTIONS). *For any actions $\mathbb{A}_1$ and $\mathbb{A}_2$, we have $\text{delay}(\mathbb{A}_1) ;; \text{delay}(\mathbb{A}_2) = \text{delay}(\mathbb{A}_2) ;; \text{delay}(\mathbb{A}_1)$.*

PROOF. Straightforward. $\square$

Definition A.22 (Pending actions). For a heap H , the pending action $\text{pending}(H)$ is the heap action recursively defined as follows (cases by order of decreasing priority):

$$\begin{aligned} \text{pending}(\emptyset) &= \emptyset \\ \text{pending}(H, o \mapsto \langle e, k, o' \rangle) &= \text{pending}(H) ;; (o' \mapsto \text{delay}(F[e])) \\ \text{pending}(H, o \mapsto 0) &= \text{pending}(H) \end{aligned}$$

Note that the order of the bindings in the heap does not matter, according to Proposition A.21.

Definition A.23 (Immediate actions). For a heap H , the immediate action $\text{immediate}(H)$ is the heap action defined as follows:

$$\text{immediate}(H)(o) = \begin{cases} (\emptyset, \text{dom}(H(o)), \emptyset, \emptyset) & \text{if } H(o) \text{ is an environment,} \\ (\emptyset, \emptyset, \emptyset, \emptyset) & \text{otherwise} \end{cases}$$

Definition A.24. The action of a heap H , expression e and reference o is defined as follows:

$$\text{action}(H, e, o) = \text{immediate}(H) ;; \text{pending}(H) ;; F[e]_o$$

LEMMA A.25 (PRESERVATION OF WELL-FLOWEDNESS).

If $\text{action}(H, e, o) = \mathbb{A}$ and $H[e]_o \rightsquigarrow H'[e']_o$, then $\text{action}(H', e', o) \leq \mathbb{A}$.

PROOF. We proceed by structural induction on e .

We match on the syntax of e (we only include the most interesting cases):

o' Impossible case as no reduction step can apply on o' .

prom ^{f_x} $\langle t \rangle \{e\}$ The rule [DPRO] applies. By adding a promise capturing the current environment in the heap H , its delayed action gets composed to the pending action $\text{pending}(H)$, but in the same time it gets removed from $F[e]_o$. We can conclude by using the commutativity of delayed actions (Proposition A.21).

$\{o_1\}_{o_2}^{o_3}$ The reduction rule [DFOR3] applies, we thus have $e' = o_1$ and $H' = H, o_3 \mapsto \langle o_1, k, o_2 \rangle$. We can trivially conclude as adding a cached promise to the heap has no effect on the pending actions.

$\{e_1\}_{o_2}^{o_3}$ The reduction rule [DCTX2] applies, we conclude by applying the induction hypothesis on e_1 .

force o' The rule [DFOR1] or [DFOR2] applies. In the first case, we can conclude immediately. Otherwise, we have $e' = \{e''\}_{o''}^{o'}$, and $H' = H, o' \mapsto \langle \text{undef}, k, o'' \rangle$, meaning that the delayed action of e'' gets removed from $\text{pending}(H)$ (as the promise is updated with the cache undef), and the action of e'' is added to $F[e]_o$ instead. We can conclude using Proposition A.20.

f.i $(o_0 \dots o_n)$ The rule [DCAL] applies. We can conclude using the fact that versions in F are well-flowed, and using Proposition A.18.

$o_1 ; e_2$ The rule [DSEQ] applies. We thus have $e' = e_2$. We can trivially conclude this case.

$e_1 ; e_2$ The rule [DCTX3] or [DCTX4] applies. In the first case, we can immediately conclude. Otherwise, we have $e' = e'_1 ; e_2$. We apply the induction hypothesis on e_1 , yielding $\text{action}(H', e'_1, o) \leq \text{action}(H, e_1, o)$. We conclude by using the monotonicity of the composition operator (Proposition A.16).

□

Definition A.26 (Well-flowedness). For a heap H , expression e and reference o , we say that (H, e, o) is well-flowed, written $\text{wfl}(H, e, o)$, if and only if $\mathbb{A} = \text{action}(H, e, o)$ is defined and $\forall (o' \mapsto (R, W, U, C)) \in \mathbb{A}. R = \emptyset$.

LEMMA A.27 (CORRECTNESS OF WELL-FLOWEDNESS). If $\text{wfl}(H, e, o)$, then the reduction rules [DREGUNDEF] and [DCONUNDEF] do not apply.

PROOF. Straightforward structural induction on e .

In order for the reduction rule [DREGUNDEF] to apply, e must be a register r . As $F[r]_o = o \mapsto (\{r\}, \emptyset, \emptyset, \emptyset)$, this means that, in order for $\text{wfl}(H, e, o)$ to hold, we must have $\text{immediate}(H)(o) \leq (\emptyset, \{r\}, \emptyset, \emptyset)$, and thus $r \in \text{dom}(H(o))$. The same applies for the reduction rule [DCONUNDEF]. □

THEOREM A.28 (CORRECTNESS OF THE FLOW ANALYSIS).

If $\emptyset \Vdash [f.i()]_o \rightsquigarrow^* H'[e']_o$, then $\emptyset \Vdash [f.i()]_o \Rightarrow^* H'[e']_o$.

If $\emptyset \Vdash [f.i()]_o \rightsquigarrow^\infty$, then $\emptyset \Vdash [f.i()]_o \Rightarrow^\infty$.

PROOF. Immediate consequence of Proposition A.18, Lemma A.25 and Lemma A.27. □

THEOREM A.29 (TYPE SAFETY).

If $\emptyset \vdash [f.i()] : t$, then either $\emptyset \Vdash [f.i()]_o \Rightarrow^\infty$ or $\emptyset \Vdash [f.i()]_o \Rightarrow^* H'[o']_o$ for some H' and o' such that $H' \vdash o' \leq_{\text{dyn}} t$.

PROOF. Immediate consequence of Theorem A.11 and Theorem A.28. □