

A Typed Intermediate Representation for Dynamic Languages

MICKAËL LAURENT, Charles University, Czech Republic

JAKOB HAIN, Purdue University, United States

FILIP KRIKAVA, Czech Technical University, Czech Republic

SEBASTIÁN KRYNSKI, Czech Technical University, Czech Republic

JAN VITEK, Northeastern University, United States

Dynamic programming languages are characterized by features such as dynamic typing, late binding, runtime code loading, and reflection. To generate efficient code for such languages, implementations must speculate on which of these dynamic features will actually be exercised during the execution of a given function. Based on these assumptions, optimized code can be produced. A common approach is to compile source programs into explicitly typed intermediate representations that make relevant behaviors manifest and amenable to static analysis. This paper presents key design decisions for a statically typed, high-level intermediate representation that directly supports optimizations such as specialization, devirtualization, inlining, and copy elimination. We propose a core calculus, called FIR, that captures essential features required for these optimizations. We give a formal semantics, a type system and flow analysis rules, 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 Krikava, Sebastián Krynski, and Jan Vitek. 2025. A Typed Intermediate Representation for Dynamic Languages. 1, 1 (July 2025), 34 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

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, operator overloading, 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 Krikava, 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, Northeastern University, Boston, Massachusetts, United States.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM XXXX-XXXX/2025/7-ART

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

Consider the function of Fig. 1, written in a hypothetical dynamic language. The ease of generating efficient code for it varies dramatically depending on the language’s degree of dynamism. At first glance, the function appears to return its first argument x if it is a scalar numeric value, the j -th element of x if it is a vector, and \emptyset otherwise.

```
function(x, j){
  r =  $\emptyset$ 
  if (scalar(x)){ r = x+r }
  else if (length(x) > j){ r = x[j] }
  return(r)
}
```

Fig. 1. A simple function

However, our expectations can be subverted in multiple ways. Focusing on the case where `scalar(x)` holds, we might expect the function to behave as the identity on numeric values. Yet even this seemingly simple case is fraught. If x is not a number, implicit string coercions could be triggered. If the `+` operator is overloaded, the expression `x+r` could invoke arbitrary effectful code and return unexpected values. Worse still, if `scalar` performs reflection or invokes native APIs, it might alter the program state—e.g., by mutating or even removing variable r from scope. Of course, the `if else` construct itself may be redefined to evaluate its arguments in non-standard ways.

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 invoke reflective mechanisms.

Building on this observation, just-in-time (JIT) compilers adopt a speculative strategy: they treat past program behavior as a predictor of future behavior and generate code optimized for the patterns observed at runtime. The compilation process is typically organized around key transformations that reduce—or at least expose—the sources of dynamism, enabling the application of standard optimizations such as constant propagation, code motion, and dead code elimination.

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? Primarily, it should provide features that allow the absence of dynamism to be explicitly marked—features that grant the compiler permission to perform optimizations. For example, the IR might encode knowledge of an object’s shape, the target of a function call, or guarantee that reflective features will not be invoked in certain regions of code.

The precise set of such features will vary depending on the target language, but the design process follows a common pattern: identify the aspects of the language that hinder compilation and find ways to restrict or make them explicit. In essence, the goal is to convert properties that are typically resolved at runtime into ones that can be known and reasoned about statically.

The central thesis underlying this paper is that incorporating as many of these features as possible into a static type system enables static verification of compiler optimizations and provides a clear, formal specification for compiler writers.

Our work builds on prior results from the development of the \tilde{R} compiler for the R programming language and its intermediate representation, PIR. In this paper, we introduce a new core intermediate representation, $\text{Fi}\tilde{R}$, designed to support features including dynamic typing, runtime code loading, delayed evaluation, copy-on-write semantics, and unrestricted reflection.

Although $\text{Fi}\tilde{R}$ is motivated by the specifics of the R language, we argue that its principles and design choices are broadly applicable to other dynamic languages. $\text{Fi}\tilde{R}$ is presented as a minimalistic calculus, inspired by core languages such as Featherweight Java [14]. Its reduced complexity enables

us to provide a formal operational semantics (Section 5.3) and to prove the soundness of its type system (also in Section 5.4).

The key features of our IR include (see Section 4):

- **Variables:** FIR distinguishes *named* from *register* variables; the former can be updated non-locally through reflection, while the latter are restricted to the current lexical scope and shielded from reflection.
- **Types:** Inspired by gradually typed languages, FIR supports a stratified type system: some variables are statically typed, others are gradually typed, and some remain entirely dynamic.
- **Ownership:** A limited form of ownership is supported, enabling basic tracking of value lifetimes and transfer semantics, particularly for optimizing copy-on-write behavior.
- **Specialization:** Functions in FIR can be compiled into multiple specialized versions, each associated with a context that determines its applicability.
- **Binding:** FIR has dynamic binding, with calls matched to an applicable version of a function at runtime, and static binding.
- **Inlining:** Function inlining is supported to expose opportunities for further optimizations while preserving the scopes of named variables.
- **Promises:** FIR includes a notion of promises to represent delayed computations, enabling support for optimizations that target lazy evaluation.

Validity of FIR programs is guaranteed by a combination of type checking to ensure that variables and functions are used according to their declaration and flow analysis to make sure that variables are initialized before they are accessed (Section 5.2). Section 3 describes related work including prior that supported speculation and deoptimization which are omitted from our IR.

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 popular dynamic 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. Dynamic languages support the ability to load and execute new code at runtime. This can be achieved either by programmatically including external libraries or by reflectively transforming code represented as text into executable form.

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 dynamic languages allow built-in functions—such as `+`—to be redefined or overloaded at runtime.

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. 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 dynamic 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 value, results from evaluating the argument `p` in function `f`. Since `p` is a promise, its evaluation triggers a call to `g`, which 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 of dynamic languages hinges on the compiler’s ability to rule out—or at least localize—the effects of dynamic features. These sources of dynamism may arise from the calling context, function arguments, or the transitive closure of reachable code.

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

This section reviews prior work related to intermediate representations in compilers, particularly in the context of dynamic languages.

3.1 Compiler IRs

Aycock provides an accessible overview of early implementation strategies for dynamic languages [1], covering systems ranging from Smalltalk [6] and Self [4] to more modern environments such as Java and JavaScript [20, 24].

While most compilers employ multiple intermediate representations to modularize the compilation process, these IRs are often treated as internal implementation details and are seldom documented comprehensively. One notable exception is the Java bytecode format, which is both clearly specified and formally studied [11, 16]. Since Java is statically typed and offers relatively limited dynamic behavior, the bytecode closely reflects the structure of the source language. A key addition, however, is the bytecode verifier—a simple form of data-flow analysis that ensures, among other things, that variables are initialized before use. The sea-of-nodes representation is a lower-level IR that makes all dependencies explicit in the control and data flow graphs [5]. Its structure is well-suited for optimizations such as duplication removal and promise elimination. Typed Assembly Language further demonstrated that static types can be applied to low-level IRs to preserve invariants during optimization—such as verifying that operands to an integer addition are indeed integers [18]. TAL assumes a statically typed source language, whereas our work addresses dynamically typed languages. Finally, MLIR offers a flexible, extensible IR framework capable of expressing a wide range of language features and analyses [15]. Although MLIR could support our target language, we chose to design an IR that more directly capture the language features of interest.

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 [12], that resemble F_{IR}. In Swift, borrowing a value does not increment its reference count. Consuming parameters are mutable; a value passed to a consuming function cannot be reused by the caller afterward. PHP provides operations for assignment and aliasing; prior efforts in compilation have focused on distinguishing between them [13], but have not focused on eliminating redundant copies. In F_{IR}, aliasing cannot occur on owned variables: a variable that is aliased is considered shared, and copies will not be statically eliminated for it. However, our approach could be adapted to a more precise treatment of aliasing.

Ownership types refers to systems that constrain the use of aliases to heap allocated objects [19]. Wrigstad introduced the notion of **borrowing** as a way to relax strict ownership invariants [22]. Rust enforces ownership semantics via its borrow checker [17]. Immutable borrows in Rust correspond to F_{IR}’s borrowing semantics: they are temporary and can reference owned values, which remain valid after the borrow expires. Rust’s smart pointers resemble F_{IR}’s shared values in that they do not enforce liveness, and any value assigned to them becomes immutable. Ownership in R

is more complex due to the language’s lazy evaluation model and the presence of first-class, mutable environments. To handle this, we introduce an explicit use instruction in `FiR`, along with a flow-sensitive analysis to detect potential read-after-use violations.

3.3 Gradual types and Like types

Gradual typing is a type system framework that allows both statically typed and dynamically typed code to coexist and interoperate within the same language. The gradual guarantee is a formal property that states that adding more precise types to a program should not change its behavior, except possibly to raise type errors earlier [21]. `FiR` follows the design of the Thorn programming language which, for performance reasons, eschews the gradual guarantee [23]. Thorn had concrete types (traditionally statically checked types), dynamic types (unchecked) and **like types**. The latter were checked for consistency but allowed to be assigned dynamic values. For a variable of type `like string`, the compiler would check that it was only operated on as if it was a string, but would allow unchecked coercion from dynamic types.

3.4 Deoptimization with Sourir and CoreIR

Prior work on formalizing just-in-time (JIT) compilation has introduced intermediate representations (IRs) that support both speculative compilation and deoptimization: `Sourir` [10] and `CoreIR` [2]. Both IRs adopt a similar design principle: they allow multiple versions of a function to coexist—typically at least one unoptimized version and one or more optimized versions that include speculative assumptions. **Sourir** extends a simple imperative intermediate language with a special `assume` instruction. This instruction encodes a speculative assumption: if the assumption holds at runtime, execution continues; if it fails, execution transfers to the corresponding unoptimized version of the function. This mechanism provides a foundation for deoptimization while preserving soundness of speculative optimizations. **CoreIR** takes a similar approach but introduces two distinct instructions: `Anchor deoptInfo label` and `Assume guardExp label`. The `Anchor` instruction is non-deterministic: it allows execution to either proceed normally or transition (via deoptimization) to a designated unoptimized code path. Anchors are primarily used during speculative optimization passes (e.g., speculation insertion) and are eliminated after optimization is complete. The `Assume` instruction is deterministic: it performs a deoptimization only if its guard expression evaluates to false. Unlike `FiR`, `CoreIR` does not index function versions with type information.

3.5 The \check{R} compiler and PIR

The \check{R} compiler is an optimizing compiler for the R language. Our work is done in the context of an extension of \check{R} , and specifically of its PIR intermediate representation. Fig. 2 illustrates the compiler’s various code representations. The original GNU-R code base is used to parse R source code into an AST. Then \check{R} has its own interpreter and bytecode, known as RIR, as well as an IR designed explicitly for high-level code transformations. Code generation is delegated to LLVM.




Fig. 2. \check{R} ’s existing IRs

PIR is an SSA-based IR [9]. An instruction has a type, effects, and immediates. The type is a combination of a kind and flags. A kind can be either an R type (e.g., real, int, env), a union (e.g., real | int), a missing value (?), or any value (val). Flags track relevant properties such as dimension (\$ is a scalar), objectness (+ is not an object), delayed evaluation (^ may be a promise).

```
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

Fig. 3 shows the unoptimized identity function.

The `LdArg` loads an argument from stack returning `val?^`, i.e. any value that might be missing or a promise. The next instruction creates an environment, and stores the function parameter `a` in it. The `LdVar` loads `a` from the environment into a new register. The `eR` effects signal that it can possibly fail. `Force` evaluates the loaded value if it was a promise. The return type `%3.3` is `val?`, i.e. it can no longer be a promise. Lastly, the symbol `!` indicates that forcing the promise can have reflective effects. 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. Our work is an attempt to present formally grounded redesign of that IR.

4 An IR for Dynamic Languages

This section presents the features of `FiR` relevant to dynamic languages. Examples are given in slightly beautified syntax¹, one that we use in our practical implementation, however all features discussed in this section are formalized in the next section.

It is important to emphasize that 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 key criteria:

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

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. These help ensure that optimizations remain sound and that the IR remains a stable and reliable target for compiler transformations. While the IR is tailored to a specific source language and compiler infrastructure, we believe that the underlying principles are broadly applicable to dynamic languages—particularly those with fewer dynamic features than R.

We begin with an example. Fig. 4 presents a source-language function definition (left) alongside the corresponding code emitted by the compiler (right). A distinctive characteristic of our IR is that it allows each function to be compiled into a set of variants, each indexed by a signature. Each variant 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 variant—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. Next, we review the features of `FiR`.

¹Specifically, we add `reg` and `var` annotations to distinguish between register and named variables, omit semi-colons, allow declarations and assignments to be combined, and use defaults in type signatures (types with no ownership modifier are `s`, and no concreteness are `!`).




Fig. 4. Source code (left) and FIR code with three versions (right)

4.1 Functions and Abstractions

A FIR *function* is composed of one or more *abstractions*. We refer to these abstractions as variants or versions of the function. An abstraction has a *signature* that lists argument types and return type as well as the potential effects of the abstraction, and a *body*.

```
(reg r:I)→I { r+1 }
```

Variants are created by the compiler to provide optimized version of a function. They differ in their signature and the transformations applied to the body. Variants 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.

Internally, the compiler maintains a flat namespace, where each unique function in the program has a single entry. Thus, `fn#1` in Fig. 4 corresponds to the compiled version of the function expression shown on the left. The mapping between names is an implementation detail. For our purposes, we focus on the internal table of functions maintained by the compiler.

4.2 Scopes and Variables

Each abstraction defines a *scope* in which variables can be defined. Our source language allows scopes to be shared with promises and to be inspected reflectively. Thus, variables that appear in the source code should be preserved in case another function decide to look them up reflectively.

One effective optimization is to eliminate a scope and all of its variables when it can be shown that reflection cannot inspect it. To enable this optimization, FIR supports two kinds of variables, *named variables* which are looked up symbolically and are exposed to reflection, and *register variables* which can be implemented in anyway the compiler sees fit and are not available to reflective operations.

Our source language allows shadowing of named variables and has nested scopes. FIR does not model nesting, so any lookup outside the current scope returns an unknown value.

The arguments of an abstraction are passed by register. A common pattern is to store them into named variables before calling a functions that may perform reflection. For instance, the source declaration `function(a) someFun(a)` can be compiled as follows:

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

where the register argument is explicitly copied to a named variable before calling a function that may perform reflection.

4.3 Promises: Creation and force

A *promise* represents a delayed computation. For example, the expression `f(x+y)` passes to `f` a promise that, when evaluated, computes `x+y`. Notably, promises execute in the lexical scope in which they were created. As a result, they can read and update register and named variables in that scope, and may even introduce new bindings—for instance, in `f(z=x+y)` if `z` was not previously defined. Promises, in the source language, are *forced* implicitly when their value is needed, and their result is cached upon evaluation.

In FiR, promises are represented explicitly. They are treated as first-class values with well-defined types and must be explicitly forced before use:

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

In this example, `r` is a promise that, when forced, returns an integer. The promise body here is a simple expression that evaluates to 42.

A subtle but important language specific detail is that, in our source language, promises are automatically forced when assigned to a variable or returned from a function: thus promises do not outlive the scope in which they were created. Our type system leverages this invariant to simplify reasoning about promise lifetimes and side effects.

The `force` instruction creates some indeterminacy regarding when promises are evaluated and when their side effects are observed. FiR incorporates flow analysis that prevents read before write errors.

4.4 Copy on write

FiR 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. For example, in a language like R that maintains an approximate reference count, `dup` needs to copy only when the reference count exceeds one. By making duplication explicit, the IR enables optimizations that eliminate unnecessary copies—something that is difficult or impossible when copies are implicitly performed 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)o = dup r1
  r2[1] = 42
}
```

Once duplicated, the vector can be safely mutated multiple times without further copying. As a result, mutable function arguments typically undergo a single duplication before any mutation, enabling efficient updates while preserving copy-on-write semantics.

4.5 Types: ownership, graduality and effects

Types in F_{IR} serve multiple purposes: they are used to verify that compiler transformations preserve well-formedness, to enable and guide optimizations, and to record compiler-generated hints. To support this range of uses, F_{IR} adopts a gradual type system, building on the work of Wrigstad *et al.* [23].

At a high level, the type system distinguishes between three kinds of types: *Concrete types*, which describe values whose types are known precisely; a *dynamic type*, used when the type of a value is entirely unknown; and *Like types*, which represent values the compiler expects to conform to a certain type, but for which it lacks static guarantees. In addition, the type system tracks ownership qualifiers for values, enabling optimizations such as copy elimination and enforcing safe behavior in the presence of reflection and delayed evaluation. This information is particularly important when reasoning about functions and promises that may observe or mutate shared state. Formally, each type in F_{IR} is composed of three components: a *kind* k , defining the structure of the values (e.g. I for integers, $v(I)$ for vectors of integers); an *ownership modifier*, and a *concreteness modifier* cx , indicating whether the type is certain (!) or speculated (?).

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).

Information about the kind of an expression can be used by the compiler to devirtualize a call, or to statically select the best version of a function. For instance, a call to $fn\#1$ (Fig. 4) could be statically bound to the version 3 if the argument is known to be a vector of integers (e.g. $v(I)o!$). Although they do not offer static guarantees, like types are used to provide hints to the compiler. For example, one could type an abstraction with a $I?$ to indicate that we have some reason to expect 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, if it is to use that information, will cast (and should deoptimize in case of failure or have a branch).

Ownership modifiers can be either f (fresh), o (owned), s (shared), and 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 is 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 change 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 “using” the register (use *r*); flow analysis ensures a register cannot be accessed after it is used.

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

This code can then be simplified into the following:

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

Subtyping on types has, for instance, type *Io!* be a subtype of *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.

Promise types and function signatures also hold information about their *effects*. In *FIR*, effects hold a single bit of information: *+* indicates that the promise or function may make use of reflection, while *-* ensures they do not. As such, the kind of a promise yielding an integer and that may perform reflection when forced is noted *p⁺(Is!)*. Values returned by a promise are shared as they can be accessed reflectively. Although effects only store one bit of information (can/cannot reflect), one could easily extend them to distinguish reflective operations (e.g. reflective read versus writes). Effects do not directly impact the typeability of an expression, as named variables are already treated pessimistically: they are either dynamic or like types. Still, annotating promises and function signatures with their potential effects is essential for optimizations.

4.6 Binding Time and Inlining

When considering a call such as *f(x)*, it is 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. However, *FIR* does not have built-in support for dynamic binding mechanisms. Instead, in the case of object-oriented dispatch—which can be highly complex—it is delegated to user-level functions from the source language’s reference implementation. *FIR* focuses on the binding decisions made by the compiler when selecting a version of a known function.

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. For dispatched calls, we do not mandate how a compiler selects the "best" version of a function; 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\ ) \# \text{ 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. 4. A more sophisticated compiler—such as \tilde{R} , which supports contextual dispatch [8]—might dynamically inspect the promise held in `r0`. If it determines that the promise is effect-free and returns an integer, it could instead select version 2.

Dispatched calls are annotated with static type signatures for two reasons. Firstly, some of the properties required to select a compatible version at runtime may not be dynamically recoverable. For example, ownership qualifiers are computed statically and cannot be inferred from the program state. Similarly, the kind of an argument cannot be determined at runtime if the argument is `undef`. Secondly, to ensure soundness, the version selected at runtime must return a value that is type-compatible with the version selected statically. In $\text{Fi}\tilde{R}$, it is permissible for two versions of the same function to produce different results for the same inputs, so long as the selected version conforms to the type system.

Static calls do not require a signature as they specify the version number explicitly; the call is checked statically, and the abstraction is called directly without searching for an applicable version, as in:

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

Finally, $\text{Fi}\tilde{R}$ supports *inlining*. When the compiler knows which abstraction will be executed—such as in a static call—it may choose to replace the call with the body of the corresponding abstraction. Fig. 5 illustrates the syntax used to represent inlined code. A key point is that inlining preserves the scope of the inlined abstraction. That is, named variables defined within the callee remain scoped to a distinct environment, separate from the caller. This design is required by the semantics of the source language, which allows reflection to observe the presence and boundaries of variable scopes. In other words, scope boundaries are semantically visible and must be respected.




Fig. 5. Inlining and scope elision

The removal of such scopes—known as scope elision—is typically performed as a separate optimization pass. If the compiler can determine that no reflective operation will observe a particular scope, it may safely remove it. Importantly, this transformation does not require explicit representation or support in the IR itself.

4.7 Discussion

The design of FİR draws on lessons from prior work on the R compiler infrastructure. Here are some noteworthy design choices and simplifications:

Deoptimization. Our implementation uses the deoptimization construct of CoreIR unchanged [2]. For simplicity we have chosen to omit them from the definition of FİR. Including them would be relatively straightforward, but would obscure the semantics of the IR with support for on stack replacement. The only added requirement that deoptimization places on the IR is the need for an identified base (or unoptimized) version for each function, it is that version that will be called when speculation fails.

Reflection. The specific reflection interface is not specified. However, implicit to FİR's design is the possibility for named variables to be redefined or even deleted. If the source language did not expose local variables to reflection, it would be possible to give them stronger types. Consider the following abstraction, it take a value, stores it in named variable *a*, calls a possibly-reflective function *f* and then returns *a*:

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

In R, *f* can delete the named variable which results in the lookup of *a* to return either the value of *a* in an enclosing scope (or at the top level), or fail.

Promises. In FİR, promises cannot escape their scope of creation. This simplifies the ownership analysis as without this restriction, any register accessed in a promise that outlives its scope of creation should be considered shared. This restriction is acceptable since in R, promises are forced when they are assigned and returned. Still, it is possible in R for a promise to escape its scope of creation by using a *delayed assignment* or by returning the environment in which it lives. However, these cases are rare and can be treated as reflection.

Omitted Features. The R language has a number of features that are not presented here. Most of them are implemented by calls to native functions and do not deserve to be reflected in the IR. Two of those are worth mentioning: name lookups and attributes.

The semantics of name lookup in R is surprisingly rich, a chain of environment is searched and in some cases it is necessary to force promises to know which value to return. These complex lookups are limited to named variables and do not affect FİR because the compiler only reasons about names in the current scope; for names outside of that scope, the implementation defaults the source language lookup mechanism.

In R, any value can have a set of attached attributes; the equivalent of a map from name to value. So, for example, the `class` attribute can be attached to any value to give that value a particular behavior during object-oriented dispatch. The R compiler only supports attributes in as much as being able to note their absence which is the overwhelmingly common case. This could be easily added to FİR as an extra bit on the type of values. A more fine grained treatment is left to future work.

5 Syntax and Semantics of FiR

FiR is a core calculus that gives semantics to an intermediate language for compiling dynamic languages such as R. As such, it includes the features described in the previous section and deliberately abstracts away many of the low-level details required in a full-featured intermediate representation.

Reflecting this design philosophy, 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.

5.1 Syntax

$\text{fn} ::= \text{fun } f\{\text{abs}_1 \dots \text{abs}_n\}$	$e ::= i \mid \text{vec}(e_1 \dots e_n) \mid v \mid e[e] \mid e\x
$\text{abs} ::= (\text{rdefs}) \xrightarrow{f_x} t\{\text{defs}; e\}$	$\mid v = e \mid e[e] = e \mid e\$x = e$
$\text{rdefs} ::= r_0 : t_0 \dots r_n : t_n$	$\mid \text{abs} \leftarrow (e_0 \dots e_n)$
$\text{defs} ::= v_0 : t_0 \dots v_n : t_n$	$\mid f.i(e_0 \dots e_n)$
$t ::= k \text{ ox } cx$	$\mid f \langle \text{sig} \rangle (e_0 \dots e_n)$
$k ::= * \mid \mathbf{V} \mid \mathbf{I} \mid \mathbf{v}(\mathbf{I}) \mid \mathbf{p}^{f_x}(t)$	$\mid e \text{ as } t$
$\text{ox} ::= \mathbf{o} \mid \mathbf{b} \mid \mathbf{s} \mid \mathbf{f}$	$\mid \text{force } e$
$f_x ::= + \mid -$	$\mid \text{use } r$
$cx ::= ! \mid ?$	$\mid \text{dup } e$
$\text{sig} ::= t_0 \dots t_n \xrightarrow{f_x} t$	$\mid \text{prom}^{f_x} \langle t \rangle \{e\}$
$v ::= x \mid r$	$\mid e ; e$

Fig. 6. Syntax

The full syntax of the FiR calculus is given in Fig. 6. A program consists of a *function table* F and an *expression* e . A function table consists of one or more *function* declarations, $\text{fun } f\{\text{abs}_1 \dots \text{abs}_n\}$, where f is the function's unique name, and abs_i is an *abstraction*. An abstraction, $(\text{rdefs}) \xrightarrow{f_x} t\{\text{defs}; e\}$, has typed argument definitions rdefs , local variables defs , a body expression e , a return type t and an effect modifier f_x . FiR distinguishes between *register variables*, ranged over by r , and *named variables*, ranged over by x . Parameters are restricted to register variables, while the body can define both named and register variables.

FiR's expressions are: integer literals (i), vector literals ($\text{vec}(e_1 \dots e_n)$), variables (v) which can be either register variables (r) or named variables (x), indexed vector reads ($e[e]$), reflective named variable reads ($e\$x$), variable writes ($v = e$), indexed vector writes ($e[e] = e$), reflective named variable writes ($e\$x = e$), inlined abstraction applications ($\text{abs} \leftarrow (e_0 \dots e_n)$), static function calls ($f.i(e_0 \dots e_n)$), dispatched function calls ($f \langle \text{sig} \rangle (e_0 \dots e_n)$), type-casts ($e \text{ as } t$), promise forces ($\text{force } e$), register variable uses ($\text{use } r$), duplications ($\text{dup } e$), promise literals ($\text{prom}^{f_x} \langle t \rangle \{e\}$), and sequencing ($e ; e$).

Reflective reads and writes. To model the challenges posed by dynamic features, we include a limited form of reflection in the calculus. This restricted reflective capability is sufficient to capture the difficulties it presents to an optimizing compiler when one function can modify the state of another.

Creating a promise in FiR captures both an expression—its body—and the environment of the enclosing abstraction. Given a variable r referencing a promise defined in a different abstraction, the following code demonstrates how to mutate a named variable x in the captured environment:

```
r$x = 42
```


This operation can occur regardless of the expected type of x , or even its existence. Note that only named variables can be accessed and assigned reflectively, register variables can only be accessed directly from their definition scope.

This ability to access named variables reflectively may leak the structure of the environments. In particular, inlining only the body of an abstraction is unsound precisely because it would merge the abstraction's environment with the surrounding context, potentially altering behavior in the presence of reflection. Consider the following example:

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

These two statements have equivalent meanings since the abstraction's body does not employ reflection. This contrasts with the following two lines that have different behaviors due to the use of reflection (we assume that $r0$ is a promise capturing the current local environment):

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

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

Reflection introduces a compelling challenge in language design, as any variable subject to reflection inherently possesses uncertain types. Operations that induce reflective effects, indicated by a $+$ in their signatures, can alter the types of all involved variables.

Types. In FiR , types, ranged over by t , are composed of: a *kind* (k), an *ownership modifier* (ox), and a *concreteness modifier* (cx). Effect modifiers, fx , extend the types of functions and promises: $+$ indicates that reflective operations may be performed and $-$ guarantees otherwise. 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. The ownership modifiers are: owned (o), borrowed (b), shared (s), and fresh (f). Finally, the possible concreteness modifiers are: concrete ($!$), or like ($?$).

The syntax of types is permissive, but we restrict types that may appear in declaration. First, we define a component setter operator $\&$ on types:

$$\begin{aligned} k \text{ } ox \text{ } cx \quad \& \quad k' &= k' \text{ } ox \text{ } cx \\ k \text{ } ox \text{ } cx \quad \& \quad ox' &= k \text{ } ox' \text{ } cx \\ k \text{ } ox \text{ } cx \quad \& \quad cx' &= k \text{ } ox \text{ } cx' \end{aligned}$$

We use a number of predicates for convenience:

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

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

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

These rules ensure that the any kind $(*)$ has a concreteness modifier of $?$, and that the type returned by a promise is shared and cannot itself be a promise. The well-formedness constraints are extended

to variable declarations with:

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

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

Casts. A type-cast $e \text{ as } t$ checks that e has type t at runtime. 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. Types casts can also be used on expressions with a concrete type in order to refine this type (e.g., refining a type Vo! into 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 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.

5.2 Static Semantics

The static semantics of FiR are defined by two relations over abstractions, $\llbracket \text{abs} \rrbracket : \text{sig}$ and $F \llbracket \text{abs} \rrbracket$. These relations assert, respectively, that an abstraction abs is well-typed and well-flowed. The function table F of a program is considered well-defined when every abstraction within it satisfies both criteria. The typing relation guarantees adherence to constraints related to types, ownership, concreteness, and effects. The flow relation ensures that register variables are initialized before they are used and prevents their access after a use operation. For convenience, we consider the function table F to be fixed, and allow the deduction rules for the judgments $\llbracket \text{abs} \rrbracket : \text{sig}$ and $F \llbracket \text{abs} \rrbracket$ to access it, without making this dependency explicit in the judgment.

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

- For effect modifiers: $-\leq_f +$.
- For ownership, subtyping is restricted to equality: $\text{ox} \leq_o \text{ox}$.
- For concreteness modifiers: $!\leq_c ?$.
- For kinds:

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

- For types:

$$\frac{k \leq_k k' \quad \text{ox} \leq_o \text{ox}' \quad \text{cx} \leq_c \text{cx}'}{k \text{ ox cx} \leq_k \text{ox}' \text{cx}'}$$

- For abstraction 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 [Abs] rule shown next checks the definition of an abstraction. The variables that can be used in an abstraction 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 abstraction's body. All type annotations must be well-formed. The effect modifier and return type of the abstraction must be those inferred for its body, with the ownership component transformed as follows: if the body returns an owned value, then the value returned by the abstraction is considered fresh (because the only reference to that value does not exist anymore when exiting the scope of the abstraction); otherwise, if the body returns a fresh or shared value, the ownership is unchanged. Note that the abstraction body should not return a borrowed value (in such a case, the abstraction is untypeable). Also, abstractions can only return non-promise values, otherwise it would allow promises to escape their scope of definition, which is undesirable as discussed below.

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

The effect analysis $\Gamma \vdash E\llbracket e \rrbracket = f_x$ 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:

$$\begin{array}{c}
 t \& f \leq t' \& f \\
 \text{shared}(t') \Rightarrow (\text{shared}(t) \text{ or fresh}(t)) \\
 \text{owned}(t') \Rightarrow \text{fresh}(t) \\
 \hline
 \text{match}(t, t')
 \end{array}$$

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 the kind and concreteness of t' ; and (2i) if t' expects a borrowed value, then t can be of any ownership, (2ii) if t' expects a shared value, then t can be either shared or fresh, or (2iii) if t' expects an owned value, then t must be fresh. Note that function parameters must be well-formed, and thus they cannot be given types with a fresh ownership.

The only literals are integers, and integers are always shared. All types are explicit; promise return and call signatures are checked, never inferred. Sequences have the type of their last element, the initial elements' types are checked but otherwise ignored. 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. Lastly, promises cannot be reflectively written, which prevents them from escaping their scope of definition: this is required for the type safety to hold, as explained in Section 5.4. Note that the type system as defined in Fig. 7 is not syntax-directed because of the subsumption rule SUB. However, this is only a presentation choice we made for clarity and concision, as subsumption can easily be inlined in other rules whenever needed.

Effect analysis. Fig. 8 shows the analysis that checks effect annotations: a function must have the reflection modifier (+) if any of its applied abstractions or forced promises have the reflection modifier. Note that the effect analysis does not impact the typeability directly (effects have 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.

$$\begin{array}{c}
\text{834} \quad \Gamma \vdash \llbracket e \rrbracket : t \quad \Gamma \vdash \llbracket e \rrbracket : t \\
\text{835} \quad \text{836} \quad \frac{}{\Gamma \vdash \llbracket i \rrbracket : \mathbf{Is}!} \text{[LIT]} \quad \frac{}{\Gamma, v : t \vdash \llbracket v \rrbracket : t} \text{[VAR]} \quad \frac{\Gamma \vdash \llbracket e \rrbracket : t}{\Gamma \vdash \llbracket e \rrbracket : t'} \text{[SUB]} \quad \frac{\Gamma \vdash \llbracket e \rrbracket : t}{\Gamma \vdash \llbracket e[e'] \rrbracket : t'} \text{[SEQ]} \\
\text{837} \quad \text{838} \quad \Gamma \vdash \llbracket e \rrbracket : t' \quad \Gamma \vdash \llbracket e \rrbracket : t \\
\text{839} \quad \Gamma \vdash E[\llbracket e \rrbracket] = fx \quad \Gamma \vdash \llbracket e_1 \rrbracket : \mathbf{Is}! \dots \Gamma \vdash \llbracket e_n \rrbracket : \mathbf{Is}! \quad \Gamma \vdash \llbracket e \rrbracket : t \\
\text{840} \quad t = \mathbf{p}^{fx}(t')\mathbf{s}! \quad wf(t) \quad t' = v(\mathbf{I})\mathbf{f}! \quad \Gamma \vdash \llbracket e \rrbracket : \mathbf{Is}! \\
\text{841} \quad \text{842} \quad \frac{}{\Gamma \vdash \llbracket \mathbf{prom}^{fx} \langle t' \rangle \{e\} \rrbracket : t} \text{[PRO]} \quad \frac{}{\Gamma \vdash \llbracket \mathbf{vec}(e_1 \dots e_n) \rrbracket : t'} \text{[VEC]} \quad \frac{}{\Gamma \vdash \llbracket e[e'] \rrbracket : \mathbf{Is}!} \text{[REA]} \\
\text{843} \quad \text{844} \quad \frac{\llbracket \mathbf{abs} \rrbracket : t_0 \dots t_n \xrightarrow{fx} t}{\Gamma \vdash \llbracket \mathbf{abs} \leftarrow (e_0 \dots e_n) \rrbracket : t} \text{[APP]} \quad \frac{\text{sig}(F(f)(i)) = t_0 \dots t_n \xrightarrow{fx} t}{\Gamma \vdash \llbracket f.i(e_0 \dots e_n) \rrbracket : t} \text{[CAL]} \\
\text{845} \quad \Gamma \vdash \llbracket e_0 \rrbracket : t'_0 \dots \Gamma \vdash \llbracket e_n \rrbracket : t'_n \quad \Gamma \vdash \llbracket e_0 \rrbracket : t'_0 \dots \Gamma \vdash \llbracket e_n \rrbracket : t'_n \\
\text{846} \quad \text{847} \quad \text{match}(t'_0, t_0) \dots \text{match}(t'_n, t_n) \quad \text{match}(t'_0, t_0) \dots \text{match}(t'_n, t_n) \\
\text{848} \quad \text{849} \quad \frac{\text{sig} = t_0 \dots t_n \xrightarrow{fx} t}{\Gamma \vdash \llbracket f \langle \text{sig} \rangle (e_0 \dots e_n) \rrbracket : t} \text{[DIS]} \quad \frac{\Gamma \vdash \llbracket e \rrbracket : t' \quad wf(t)}{\Gamma \vdash \llbracket e \text{ as } t \rrbracket : t} \text{[CAS]} \quad \frac{\Gamma \vdash \llbracket e \rrbracket : \mathbf{p}^{fx}(t)\mathbf{s}!}{\Gamma \vdash \llbracket \mathbf{force } e \rrbracket : t} \text{[FOR]} \\
\text{850} \quad \exists i. \text{sig}(F(f)(i)) \leq \text{sig} \quad \text{851} \quad \Gamma \vdash \llbracket f.i(e_0 \dots e_n) \rrbracket : t' \quad \text{852} \quad \Gamma \vdash \llbracket f \langle \text{sig} \rangle (e_0 \dots e_n) \rrbracket : t \\
\text{853} \quad \text{854} \quad \Gamma \vdash \llbracket r \rrbracket : t \quad \Gamma \vdash \llbracket e \rrbracket : t \quad t = \Gamma(v) \quad \Gamma \vdash \llbracket e \rrbracket : t' \\
\text{855} \quad \text{856} \quad \text{owned}(t) \quad \text{vec}(t) \quad \text{match}(t', t) \\
\text{857} \quad \text{858} \quad t' = t \ \& \ \mathbf{f} \quad t' = t \ \& \ \mathbf{f} \quad \neg \text{borrowed}(t) \\
\text{859} \quad \text{860} \quad \frac{}{\Gamma \vdash \llbracket \mathbf{use } r \rrbracket : t'} \text{[USE]} \quad \frac{}{\Gamma \vdash \llbracket \mathbf{dup } e \rrbracket : t'} \text{[DUP]} \quad \frac{}{\Gamma \vdash \llbracket v = e \rrbracket : t} \text{[ASS]} \\
\text{861} \quad \Gamma \vdash \llbracket e \rrbracket : v(\mathbf{I})\mathbf{o}! \quad \Gamma \vdash \llbracket e \rrbracket : t \quad \Gamma \vdash \llbracket e \rrbracket : t \quad \text{prom}(t) \\
\text{862} \quad \Gamma \vdash \llbracket e' \rrbracket : t \quad \Gamma \vdash \llbracket e'' \rrbracket : t \quad \Gamma \vdash \llbracket e \rrbracket : t \quad \text{value}(t') \quad \text{shared}(t') \\
\text{863} \quad \text{864} \quad t = \mathbf{Is}! \quad \text{prom}(t) \quad \Gamma \vdash \llbracket e' \rrbracket : t' \quad \Gamma \vdash \llbracket e \rrbracket : t \quad \text{prom}(t) \\
\text{865} \quad \text{866} \quad \frac{}{\Gamma \vdash \llbracket e[e'] = e'' \rrbracket : t} \text{[WR1]} \quad \frac{}{\Gamma \vdash \llbracket e\$x \rrbracket : \mathbf{*s?}} \text{[REF]} \quad \frac{}{\Gamma \vdash \llbracket e\$x = e' \rrbracket : t'} \text{[RW1]}
\end{array}$$

Fig. 7. Types

- An abstraction is reflective if annotated as such. An abstraction 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 effect 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 abstraction being applied is, and a call is reflective if the abstraction being called is.
- Any expression containing reflective sub-expressions is reflective.

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

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

Flow Analysis. A program is well-flowed if register variables are initialized before their first access and never accessed after a use. The challenge here is related to promises, as the point in

the program where they are evaluated may be difficult (or even, in the general case, impossible) to know statically. A precise solution would involve a data-flow analysis that would model the flow of promise 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.

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 use (in the sense of *use r*).
- 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.

$$\begin{array}{c}
 \Gamma = r\text{defs}, \text{defs} \\
 \frac{\Gamma \vdash E[e] = fx' \quad fx' \leq_f fx}{\text{[EAbs]} \quad E[(r\text{defs}) \xrightarrow{fx} t\{\text{defs}; e\}] = fx} \quad \text{[ELit]} \quad \frac{}{\Gamma \vdash E[i] = -} \quad \text{[EVar]} \quad \frac{}{\Gamma \vdash E[v] = -} \\
 \\
 \frac{\Gamma \vdash E[e] = fx \quad \Gamma \vdash E[e'] = fx' \quad fx'' = fx \vee fx'}{\text{[ESeq]} \quad \Gamma \vdash E[e; e'] = fx''} \quad \frac{\Gamma \vdash E[e_0; \dots e_n] = fx}{\text{[EVec]} \quad \Gamma \vdash E[\text{vec}(e_1 \dots e_n)] = fx} \quad \frac{\Gamma \vdash E[e_0; \dots e_n] = fx' \quad fx'' = fx \vee fx'}{\text{[EApp]} \quad \Gamma \vdash E[\text{abs} \leftarrow (e_0 \dots e_n)] = fx''} \\
 \\
 \frac{\text{sig}(F(f)(i)) = t_0 \dots t_n \xrightarrow{fx} t \quad \Gamma \vdash E[e_0; \dots e_n] = fx' \quad fx'' = fx \vee fx'}{\text{[ECAL]} \quad \Gamma \vdash E[f.i(e_0 \dots e_n)] = fx''} \quad \frac{\text{sig} = t_0 \dots t_n \xrightarrow{fx} t \quad \Gamma \vdash E[e_0; \dots e_n] = fx' \quad fx'' = fx \vee fx'}{\text{[EDis]} \quad \Gamma \vdash E[f <\text{sig}> (e_0 \dots e_n)] = fx''} \\
 \\
 \frac{\Gamma \vdash E[e] = fx \quad \Gamma \vdash [e] : \mathbf{p}^{fx'}(t) \quad fx'' = fx \vee fx'}{\text{[EFor]} \quad \Gamma \vdash E[\text{force } e] = fx''} \quad \frac{\Gamma \vdash E[e] = fx}{\text{[ECas]} \quad \Gamma \vdash E[e \text{ as } t] = fx} \quad \text{[ERefl]} \quad \Gamma \vdash E[r\$x] = + \\
 \\
 \frac{\Gamma \vdash E[e] = fx}{\text{[ERea]} \quad \Gamma \vdash E[r[e]] = fx} \quad \frac{\Gamma \vdash E[e] = fx}{\text{[EDup]} \quad \Gamma \vdash E[\text{dup } e] = fx} \quad \frac{\Gamma \vdash E[e] = fx}{\text{[EUse]} \quad \Gamma \vdash E[\text{use } e] = fx} \\
 \\
 \frac{}{\text{[EPro]} \quad \Gamma \vdash E[\text{prom}^{fx} <t>\{e\}] = -} \quad \frac{\Gamma \vdash E[e] = fx}{\text{[EAss]} \quad \Gamma \vdash E[v = e] = fx} \quad \frac{\Gamma \vdash E[e; e'] = fx}{\text{[EWri]} \quad \Gamma \vdash E[r[e] = e'] = fx} \\
 \\
 \text{[ERwr]} \quad \Gamma \vdash E[r\$x = e] = +
 \end{array}$$

Fig. 8. Effect Analysis

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.

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

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

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.

$$\begin{array}{lll} \text{[FLIT]} \frac{A = (\emptyset, \emptyset, \emptyset, \emptyset)}{F[i] = A} & \text{[FVAR]} \frac{A = (\emptyset, \emptyset, \emptyset, \emptyset)}{F[x] = A} & \text{[FSEQ]} \frac{F[e] = A \quad F[e'] = A' \quad A'' = A ;; A'}{F[e; e'] = A''} \\ \text{[FMUL]} \frac{F[e_0] = A_0 \dots F[e_n] = A_n \quad A' = (\emptyset, \emptyset, \emptyset, \emptyset) ;; A_0 ;; \dots A_n}{F[e_0 \dots e_n] = A'} & \text{[FVEC]} \frac{F[e_1 \dots e_n] = A}{F[\text{vec}(e_1 \dots e_n)] = A} & \text{[FAPP]} \frac{F[\text{abs}]}{F[\text{abs} \leftarrow (e_0 \dots e_n)] = A} \\ \text{[FCAL]} \frac{F[e_0 \dots e_n] = A}{F[f.i(e_0 \dots e_n)] = A} & \text{[FDis]} \frac{F[e_0 \dots e_n] = A}{F[f <sig> (e_0 \dots e_n)] = A} & \text{[FCAS]} \frac{F[e] = A}{F[e \text{ as } t] = A} \\ \text{[FFOR]} \frac{F[e] = A}{F[\text{force } e] = A} & \text{[FREG]} \frac{A = \text{read } r}{F[r] = A} & \text{[FREFL]} \frac{A = \text{read } r}{F[r\$x] = A} & \text{[FREa]} \frac{F[e] = A \quad A' = \text{read } r ;; A}{F[r[e]] = A'} \\ \text{[FDUP]} \frac{F[e] = A}{F[\text{dup } e] = A} & \text{[FUSE]} \frac{A = \text{read } r ;; \text{use } r}{F[\text{use } r] = A} & \text{[FPRO]} \frac{F[e] = A \quad A = (R, W, U, C) \quad C' = R \cup W \cup U \cup C}{F[\text{prom}^{f_x} \langle t \rangle \{e\}] = (R, \emptyset, U, C')} \\ \text{[FASs1]} \frac{F[e] = A}{F[x = e] = A} & \text{[FASs2]} \frac{F[e] = A \quad A' = A ;; \text{write } r}{F[r = e] = A'} & \text{[FWRI]} \frac{F[e] = A \quad F[e'] = A' \quad A'' = \text{read } r ;; A ;; A' ;; \text{write } r}{F[r[e] = e'] = A''} \\ & & \text{[FRWR]} \frac{F[e] = A \quad A' = \text{read } r ;; A ;; \text{write } r}{F[r\$x = e] = A'} \end{array}$$

Fig. 9. Flow Analysis

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

Our flow analysis runs on each subexpression of the body of an abstraction and composes their actions. Analysis fails if this action cannot be computed, or if its R set is not a subset of the abstraction parameters. The former means there is a use invariant violation (read/assign/capture after use, or capture before use), and the latter means there is a read before assign.

5.3 Dynamic Semantics

We endow FiR with a small-step operational semantics with evaluation contexts [7]. The program state is captured by a *heap* (H), mapping *references* (o) to values (o), written $o_0 \mapsto O_0 \dots o_n \mapsto O_n$. References include *undef*, which does not occur on the left of a mapping.

A FiR expression that reduces to *undef* is analogous to an expression of the source language that raises an error or otherwise diverts control-flow. Hence, *undef* is considered an instance of every type, because code where it shows up would not be evaluated at all. This allows us to type failable operations, like indexing and casting, the same as if they cannot fail.

Values are defined as follows:

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

The syntax is extended with the following terms:

$e ::= \dots$	
$\quad \mid o$	references
$\quad \mid \{e\}_o$	abstraction scope
$\quad \mid \{e\}_o'$	promise scope

Evaluation contexts are defined as follows:

$\mathcal{E} ::= [] \mid \text{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	
$\quad \mid v = \mathcal{E} \mid \mathcal{E}[e] = e \mid o[\mathcal{E}] = e \mid o[o] = \mathcal{E} \mid \mathcal{E}\$x = e \mid o\$x = \mathcal{E}$	
$\quad \mid \text{abs} \leftarrow (o_1 \dots o_k, \mathcal{E}, e_1 \dots e_n)$	
$\quad \mid f.i(o_1 \dots o_k, \mathcal{E}, e_1 \dots e_n) \mid f \langle \text{sig} \rangle (o_1 \dots o_k, \mathcal{E}, e_1 \dots e_n)$	
$\quad \mid \mathcal{E} \text{ as } t \mid \text{force } \mathcal{E} \mid \text{use } r \mid \text{dup } \mathcal{E} \mid \mathcal{E}; e$	

The relation $H \llbracket e \rrbracket_o \Rightarrow H' \llbracket e' \rrbracket_o$ 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:

$$H \llbracket \mathcal{E}[e] \rrbracket_o \Rightarrow H' \llbracket \mathcal{E}[e'] \rrbracket_o \quad \text{if} \quad H \llbracket e \rrbracket_o \Rightarrow H' \llbracket e' \rrbracket_o$$

The main judgments appear in Fig. 10. The semantics deal with out of bound arrays 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:

$$H \llbracket \mathcal{E}[e] \rrbracket_o \Rightarrow H' \llbracket \mathcal{E}[\text{undef}] \rrbracket_o \quad \text{if} \quad H \llbracket e \rrbracket_o \Rightarrow H' \llbracket \text{undef} \rrbracket_o$$

The above rule is prioritized over the other judgments. This treatment of `undef` simplifies the proofs of type safety: as usual, our type system focuses on preventing stuck states, and thus failable operations that are not caught by our type system should not lead to a stuck state, but instead should evaluate to `undef`.

A literal (`DLIT`), promise constructor (`DPRO`), or vector constructor (`DVEC`) creates a fresh reference within the heap, then reduces to that reference. The other expressions that create fresh references are `dup` expressions (`DDUP`) and applications (`DAPP1`); a `dup` expression copies its argument into a fresh reference, while an application references a new environment (corresponding to the scope of the abstraction being called), initialized with the parameter values. This environment is unreferenced when the application returns (`DAPP2`).

An assignment (`DASS`) maps a variable to a reference in the enclosing environment, while a reflective write (`DRWR`) maps a variable to a reference in the provided promise's environment. Notably, a vector write (`DWRI`) does not remap the variable containing the vector reference, instead it mutates the vector in-place. This vector value must be referenced at most once in the heap, otherwise the rule does not apply and the reduction is stuck.

A variable expression (`DVAR`) and use expression (`DUSE`) looks up its variable in the enclosing environment, and a use expression additionally removes the variable mapping (since the used register variable cannot be accessed again, its mapping is redundant). A reflective read (`DRRD`) looks up its variable in the provided promise's environment.

A force expression whose argument is a cached promise (`DFor1`) reduces to the cached value. A force expression whose argument is an unevaluated promise (`DFor2` and `DFor3`), reduces to the result of evaluating the promise's body in the promise's environment, and caches the result for future forces.

A static call reduces to an application, where the abstraction is looked up by function and version, and the application arguments are the call arguments. A dispatch call reduces to a static call, where the version to call is selected according to the signature annotation at the call site and the argument types; specifically, the arguments are evaluated to get their runtime types, and the selected version in the function is the first whose signature is a subtype of the annotated signature, and whose parameter types are satisfied by all of said argument types.

A cast expression (`DCAS`) evaluates to its argument if the argument is of the correct kind, otherwise it reduces to `undef`. A sequence (`DSEQ`) evaluates each expression in order and reduces to the last one. Lastly, `DCTX1` and `DCTX2` are both mechanical rules to evaluate promise or abstraction code in the appropriate context.

Some additional notation is described. o fresh denotes a reference o that does not occur in the heap; the `env` function extracts a reference to the environment captured by a promise ($\text{env}(\langle _, k, o \rangle) = o$); $H(o)$ returns the value of o in H ; if a value is an environment then $H(o)(v)$ 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 a mapping; $\text{sig}(\text{abs})$ returns the signature of abs ; $F(f)(i)$ returns the version i of function f ; $\text{refs}(H, o)$ returns the number of occurrences of o on the right hand-side of a mapping in the environments in H . Lastly, the $H \vdash o \leq_{\text{dyn}} t$ relation is defined as follows:

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

5.4 Type Safety

The static semantics (in particular the type checker and flow analysis described in Section 5.2) ensure the safety of the execution for the dynamic semantics: if a program is well-typed and

$$\begin{array}{c}
1079 \quad \text{o' fresh} \\
1080 \quad \frac{H' = H, \text{o}' \mapsto i}{\text{[DLIT]}} \frac{}{H[\![i]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \\
1081 \\
1082 \\
1083 \quad \text{o' fresh} \\
1084 \quad \frac{H(\text{o}_1) = i_1 \dots H(\text{o}_n) = i_n}{\text{[DVEC]}} \frac{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}}} \\
1085 \\
1086 \quad \frac{\text{o}' = H(\text{o})(v)}{\text{[DVAR]}} \frac{}{H[\![v]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad \frac{0 = H(\text{o}), v \mapsto \text{o}'}{\text{[DASS]}} \frac{H' = H, \text{o} \mapsto 0}{H[\![v = \text{o}']\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \\
1087 \\
1088 \\
1089 \quad \text{o'' fresh} \\
1090 \quad \frac{H' = H, \text{o}'' \mapsto H(\text{o}')}{\text{[DSEQ]}} \frac{}{H[\![\text{o}'; e]\!]_{\text{o}} \Rightarrow H[\![e]\!]_{\text{o}}} \quad \frac{\text{o'' fresh}}{\text{[DDUP]}} \frac{H' = H, \text{o}'' \mapsto H(\text{o}')}{H[\![\mathbf{dup} \text{o}']\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad \frac{\text{o}' = H(\text{o})(r)}{\text{[DUSE]}} \frac{0 = H(\text{o}) \setminus r}{H' = H, \text{o} \mapsto 0} \frac{}{H[\![\mathbf{use} r]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \\
1091 \\
1092 \\
1093 \quad \frac{H \vdash \text{o} \leq_{\text{dyn}} t \Rightarrow \text{o}' = \text{o}}{\text{[DCAS]}} \frac{H \vdash \text{o} \not\leq_{\text{dyn}} t \Rightarrow \text{o}' = \mathbf{undef}}{H[\![\text{o}' \mathbf{as} t]\!]_{\text{o}} \Rightarrow H[\![\text{o}']\!]_{\text{o}}} \quad \frac{H(\text{o}') = \langle \text{o}'', k, \text{o}''' \rangle}{\text{[DFOR1]}} \frac{}{H[\![\mathbf{force} \text{o}']\!]_{\text{o}} \Rightarrow H[\![\text{o}']\!]_{\text{o}}} \\
1094 \\
1095 \\
1096 \\
1097 \quad \frac{H(\text{o}') = \langle e, k, \text{o}'' \rangle}{\text{[DFOR2]}} \frac{H' = H, \text{o}' \mapsto \langle \mathbf{undef}, k, \text{o}'' \rangle}{H[\![\mathbf{force} \text{o}']\!]_{\text{o}} \Rightarrow H'[\![\{e\}_{\text{o}''}]\!]_{\text{o}}} \quad \frac{H' = H, \text{o}''' \mapsto \langle \text{o}', k, \text{o}'' \rangle}{\text{[DFOR3]}} \frac{}{H[\![\{ \text{o}' \}_{\text{o}''}]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \\
1098 \\
1099 \\
1100 \\
1101 \quad \text{smallest } i \text{ such that} \\
1102 \quad \text{sig}(F(f)(i)) = t_0 \dots t_n \xrightarrow{f_x} t \\
1103 \quad t_0 \dots t_n \xrightarrow{f_x} t \leq_{\text{sig}} \text{sig} \\
1104 \quad H \vdash \text{o}_0 \leq_{\text{dyn}} t_0 \dots H \vdash \text{o}_n \leq_{\text{dyn}} t_n \\
1105 \quad e = f.i(\text{o}_0 \dots \text{o}_n) \\
1106 \quad \frac{}{\text{[DDIS]}} \frac{}{H[\![f \langle \text{sig} \rangle (\text{o}_0 \dots \text{o}_n)]\!]_{\text{o}} \Rightarrow H[\![e]\!]_{\text{o}}} \\
1107 \\
1108 \\
1109 \quad \text{o' fresh} \\
1110 \quad \frac{H' = H \setminus \text{o}'}{\text{[DAPP2]}} \frac{}{H[\![\{ \text{o}' \}_{\text{o}''}]\!]_{\text{o}} \Rightarrow H'[\![\text{o}']\!]_{\text{o}}} \quad \frac{\text{o' fresh}}{\text{[DAPP1]}} \frac{\text{abs} = (r_0 : t_0 \dots r_n : t_n) \xrightarrow{f_x} t \{ \text{defs}; e \}}{0 = r_0 \mapsto \text{o}_0 \dots r_n \mapsto \text{o}_n} \frac{H' = H, \text{o}' \mapsto 0}{H[\![\text{abs} \leftarrow (\text{o}_0 \dots \text{o}_n)]\!]_{\text{o}} \Rightarrow H'[\![\{e\}_{\text{o}'}]\!]_{\text{o}}} \\
1111 \\
1112 \quad \frac{F(f)(i) = \text{abs}}{\text{[DCAL]}} \frac{e = \text{abs} \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}}} \\
1113 \\
1114 \quad \frac{H(\text{o}') = \mathbf{vec}(i_0 \dots i_n)}{\text{[DWR1]}} \frac{\text{refs}(H, \text{o}') \leq 1}{H(\text{o}'') = j \quad H(\text{o}''') = i} \frac{0 = \mathbf{vec}(i_0 \dots i_{j-1}, i, i_{j+1} \dots i_n)}{H' = H, \text{o}' \mapsto 0} \frac{}{H[\![\text{o}'[\text{o}']] = \text{o}''']_{\text{o}} \Rightarrow H'[\![\text{o}''']_{\text{o}}} \\
1115 \\
1116 \quad \frac{H(\text{o}') = \mathbf{vec}(i_0 \dots i_n)}{\text{[DREA]}} \frac{\text{o}''' \text{ fresh} \quad H(\text{o}'') = j}{H' = H, \text{o}'' \mapsto i_j} \frac{}{H[\![\text{o}'[\text{o}']]_{\text{o}} \Rightarrow H'[\![\text{o}''']_{\text{o}}} \\
1117 \\
1118 \\
1119 \quad \frac{\text{o}'' = \text{env}(H(\text{o}'))}{\text{[DRRD]}} \frac{\text{o}''' = H(\text{o}'')(x)}{H[\![\text{o}'\$x]\!]_{\text{o}} \Rightarrow H[\![\text{o}''']_{\text{o}}} \quad \frac{\text{o}'' = \text{env}(H, \text{o}') \quad 0 = H(\text{o}'), x \mapsto \text{o}''}{\text{[DRWR]}} \frac{H' = H, \text{o}'' \mapsto 0}{H[\![\text{o}'\$x = \text{o}''']_{\text{o}} \Rightarrow H'[\![\text{o}''']_{\text{o}}} \\
1120 \\
1121 \\
1122 \\
1123 \quad \frac{H[\![e]\!]_{\text{o}'} \Rightarrow H'[\![e]\!]_{\text{o}'}}{\text{[DCTX1]}} \frac{}{H[\![\{e\}_{\text{o}'}]\!]_{\text{o}} \Rightarrow H'[\![\{e'\}_{\text{o}'}]\!]_{\text{o}}} \quad \frac{H[\![e]\!]_{\text{o}'} \Rightarrow H'[\![e]\!]_{\text{o}'}}{\text{[DCTX2]}} \frac{}{H[\![\{e\}_{\text{o}'}]\!]_{\text{o}} \Rightarrow H'[\![\{e'\}_{\text{o}'}]\!]_{\text{o}}} \\
1124 \\
1125 \\
1126 \\
1127
\end{array}$$

Fig. 10. Dynamic Semantics

well-flowed, then its reduction cannot get stuck (although it can reduce to `undef`). In this section, we provide an overview of the type safety proofs and the invariants that are preserved throughout execution. A more detailed proof is available in Appendix A (Theorem A.29).

THEOREM (TYPE SAFETY). *If every abstraction in the function table F is well-defined (i.e. well-typed and well-flowed), and if $\emptyset \vdash \llbracket f.i() \rrbracket : \mathfrak{t}$, 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 $H' \vdash o' \leq_{\text{dyn}} \mathfrak{t}$.*

The proof is decomposed in two parts. The first part (Appendix A.2) focuses on the type checker, while the second part (Appendix A.3) focuses on the flow analysis.

Safety of the type checker. The type checker alone does not guarantee type safety for the dynamic semantics described in Section 5.3. In particular, it does not ensure that registers are not read or **used** before having been initialized (or after having been **used**). Consequently, the type safety theorem we prove for the type checker applies to a modified version of the dynamic semantics, in which those two reduction rules are added:

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

As `undef` propagates to the top level, the purpose of these two rules is to treat an invalid register read or **use** as an unchecked runtime error instead of a stuck reduction.

The proof of type safety is then decomposed into two lemmas: *type preservation*, that states that reduction steps preserve typeability, and *progress*, that states that a well-typed expression can be reduced (unless it is already a final result, that is, a reference `o`).

Proving these two lemmas requires extending our type system so that it can not only type expressions of the source language, but also transient expressions that we obtain during reduction and that may feature references `o` and context switches $\{e\}_o$ and $\{e\}_{o'}^{o''}$. The idea is to give the current heap H as an additional input to the type system, and to type references according to this heap. Additionally, in order to type context switches, the type environment Γ for the current scope is replaced by a mapping Δ from scopes `o` to type environments Γ .

The type preservation can then be proved, relying on several invariants that are preserved throughout execution and are required for the type preservation lemma to be inductive:

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

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

Valid borrowing Any reference in H that has an owned type according to Δ cannot be used anywhere else in H , except in a more recent environment (i.e. an environment above in the call stack) in which it has a borrowed type according to Δ , and

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

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

A formal description of these invariants is given in Appendix A.2 (Definition A.7), as well as a proof of type preservation and progress.

Safety of the flow analysis. The flow analysis guarantees that registers are not read or **used** before having been initialized or after having been **used**. In other terms, it guarantees that, for any well-flowed expression, the two reduction rules `DREGUNDEF` and `DUSEUNDEF` are never used.

Consequently, combining the guarantees of the type checker with those of the flow analysis gives type safety for the dynamic semantics of Section 5.3.

Similarly to the proofs for the type checker, we decompose the correctness of the flow analysis into two lemmas: *preservation of well-flowedness* and *correctness of well-flowedness*. The first one ensures that a well-flowed expression stays well-flowed after a reduction step, and the second ensures that DREGUNDEF and DUSEUNDEF rules cannot apply on well-flowed expressions.

As for the type-checker, in order to satisfy these lemmas, the flow analysis has to be extended to support transient expressions that we obtain during reduction. For that, we reuse similar ideas, replacing actions A by mappings \mathbb{A} from scopes \mathfrak{o} to actions A , and adding rules to handle the case of references and context switches. These additional rules are formalized in Appendix A.3, together with the associated proofs.

6 Conclusion

Dynamic languages pose many challenges for compiler writers due to features like dynamic typing, reflection, late binding, copy-on-write, and delayed evaluation. These features make it difficult to reason statically about program behavior.

This paper introduces F $\dot{\text{R}}$, a statically typed high-level intermediate representation designed to address these challenges. By exposing key dynamic behaviors explicitly (reflective reads and writes, dispatched calls, inlining, promise creations and evaluations, vector copies) and providing gradual types that combine information about the structure of values and their ownership, F $\dot{\text{R}}$ simplifies the implementation of classical compiler optimizations such as specialization, inlining, scope elision, and copy elimination. It allows compiler writers to reason locally about code transformations without compromising global semantic correctness.

We have formalized the operational semantics of F $\dot{\text{R}}$, as well as a static semantics composed of a type system, a flow analysis, and an effect analysis. We proved that together, they ensure the safety of the execution (in the usual sense of type safety). Although the design of F $\dot{\text{R}}$ is informed by the semantics of R, we believe it generalizes to a wide range of dynamic languages that exhibit similar behaviors.

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] 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
- [7] 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
- [8] 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
- [9] 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

- [10] 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
- [11] 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
- [12] Michael Gottesman, Joe Groff, and John McCall. 2024. borrowing and consuming parameter ownership modifiers. RFC 0377. Apple. <https://github.com/swiflang/swift-evolution/blob/main/proposals/0377-parameter-ownership-modifiers.md>
- [13] Andrei Homescu and Alex Şuhan. 2011. HappyJIT: a tracing JIT compiler for PHP. In *Symposium on Dynamic Languages (DLS)*. doi:10.1145/2047849.2047854
- [14] 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
- [15] 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>
- [16] 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/>
- [17] Niko Matsakis. 2016. Introducing MIR. <https://blog.rust-lang.org/2016/04/19/MIR/>
- [18] Greg Morrisett, David Walker, Karl Crary, 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
- [19] James Noble, John Potter, and Jan Vitek. 1998. Flexible Alias Protection. In *European Conference on Object-Oriented Programming (ECOOP)*. doi:10.1007/BFb0054091
- [20] 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
- [21] 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
- [22] Tobias Wrigstad. 2006. *Ownership-Based Alias Management*. Doctoral Dissertation. Department of Information Technology, Uppsala University. <https://wrigstad.com/papers/thesis.pdf>
- [23] 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
- [24] 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

We start by giving the full set of rules for a small step semantics. This semantics is the same as the one given in Section 5.3, except for the two extra rules [DREGUNDEF] and [DUSEUNDEF] (in blue). Reduction steps for this semantics are noted $H[e]_o \rightsquigarrow H'[e']_o$. We note $H[e]_o \rightsquigarrow^* H'[e']_o$ a reduction of any number of steps, and $H[e]_o \rightsquigarrow^\infty$ a diverging reduction.

$$\begin{array}{c}
\text{1275} \\
\text{1276} \\
\text{1277} \\
\text{1278} \\
\text{1279} \\
\text{1280} \\
\text{1281} \\
\text{1282} \\
\text{1283} \\
\text{1284} \\
\text{1285} \\
\text{1286} \\
\text{1287} \\
\text{1288} \\
\text{1289} \\
\text{1290} \\
\text{1291} \\
\text{1292} \\
\text{1293} \\
\text{1294} \\
\text{1295} \\
\text{1296} \\
\text{1297} \\
\text{1298} \\
\text{1299} \\
\text{1300} \\
\text{1301} \\
\text{1302} \\
\text{1303} \\
\text{1304} \\
\text{1305} \\
\text{1306} \\
\text{1307} \\
\text{1308} \\
\text{1309} \\
\text{1310} \\
\text{1311} \\
\text{1312} \\
\text{1313} \\
\text{1314} \\
\text{1315} \\
\text{1316} \\
\text{1317} \\
\text{1318} \\
\text{1319} \\
\text{1320} \\
\text{1321} \\
\text{1322} \\
\text{1323}
\end{array}
\begin{array}{l}
\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(\text{t fx}), o \rangle}{H[\text{prom}^{\text{fx}} \langle \text{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 \quad H' = H, o' \mapsto \text{vec}(i_1 \dots i_n)}{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{[DREGUNDEF]} \frac{r \notin \text{dom}(H(o))}{H[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{[DUSE]} \frac{o' = H(o)(r) \quad H' = H, o \mapsto H(o) \setminus r}{H[\text{use } r]_o \rightsquigarrow H'[o']_o} \quad \text{[DUSEUNDEF]} \frac{r \notin \text{dom}(H(o))}{H[\text{use } 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 \quad H' = H, o' \mapsto \langle \text{undef}, k, o'' \rangle}{H[\text{force } o']_o \rightsquigarrow H'[\{e\}_{o''}^o]_o} \quad \text{[DFOR3]} \frac{H' = H, o''' \mapsto \langle o', k, o'' \rangle}{H[\{o'\}_{o''}^o]_o \rightsquigarrow H'[o'']_o} \\
\text{[DAPP1]} \frac{o' \text{ fresh} \quad \text{abs} = (r_0 : t_0 \dots r_n : t_n) \xrightarrow{\text{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{abs} \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{abs} \quad e = \text{abs} \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{\text{fx}} t \quad t_0 \dots t_n \xrightarrow{\text{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} \rangle (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_3]_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''']_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{[DRWR]} \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]_o \rightsquigarrow H'[\{e'\}_{o'}^o]_o} \quad \text{[DCTX3]} \frac{H[e]_o \rightsquigarrow H'[\text{undef}]_o}{H[\mathcal{E}[e]_o]_o \rightsquigarrow H'[\text{undef}]_o} \quad \text{[DCTX4]} \frac{H[e]_o \rightsquigarrow H'[e']_o}{H[\mathcal{E}[e]_o]_o \rightsquigarrow H'[\mathcal{E}[e']_o]_o}
\end{array}$$

A.2 Proofs for the static semantics

In this section, we prove that the type system defined in Section 5.2 ensures type safety (“well-typed programs cannot go wrong”) for the small-step semantics defined above. However, due to the presence of the two reduction rules [DREGUNDEF] and [DUSEUNDEF], this type safety theorem does not ensure that registers are correctly initialized. Ensuring a correct initialization of the registers is done by the flow analysis, and the associated proofs are available in the next section.

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

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

Our type safety theorem will be proved by first proving a type preservation lemma (Lemma A.9), stating that typeability is preserved by reduction, and then a progress lemma (Lemma A.10), stating that a well-typed expression is either a reference o or can be reduced. However, before proving these lemmas, we have to extend the type system so that it can type intermediate results: in particular, we must be able to type references o , abstraction scopes $\{e\}_o$, and promise scopes $\{e\}_o^{o'}$. For that, our extended type system needs to take as input the current heap and a heap descriptor:

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:

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

$$\begin{array}{c} \text{[VAR]} \frac{t = \Delta(o)(v)}{\Delta, H \vdash_o \llbracket v \rrbracket : t} \quad \text{[UNDEF]} \frac{}{\Delta, H \vdash_o \llbracket \text{undef} \rrbracket : t} \\ \\ \text{[DYN]} \frac{\begin{array}{l} k = \text{dynkind}(H(o')) \quad H(o') = \langle _ , _ , o'' \rangle \Rightarrow o'' \in \text{dom}(H) \\ ox = \mathbf{s} \Rightarrow \text{owner}(H, \Delta, o') = \{\} \quad ox = \mathbf{o} \Rightarrow \text{owner}(H, \Delta, o') = \{o\} \quad ox = \mathbf{f} \Rightarrow o' \text{ fresh} \end{array}}{\Delta, H \vdash_o \llbracket o \rrbracket : k \text{ ox} !} \\ \\ \text{[ASC]} \frac{\begin{array}{l} \Delta, H \vdash_{o'} \llbracket e \rrbracket : t' \\ t = \begin{cases} t' & \text{if shared}(t') \text{ or fresh}(t') \\ t' \& \mathbf{f} & \text{if owned}(t') \end{cases} \\ \text{value}(t) \quad \text{wf}(t) \end{array}}{\Delta, H \vdash_o \llbracket \{e\}_{o'} \rrbracket : t} \quad \text{[PSC]} \frac{\begin{array}{l} \Delta, H \vdash_{o'} \llbracket e \rrbracket : t \\ \text{value}(t) \quad \text{shared}(t) \end{array}}{\Delta, H \vdash_o \llbracket \{e\}_{o'}^{o''} \rrbracket : t} \end{array}$$

where $\text{owner}(H, \Delta, o) = \{o' \mid \exists r. H(o')(r) = o \text{ and } \text{owned}(\Delta(o')(r))\}$ and o' fresh is true if and only if o' does not appear in H except at the left-hand side of a binding, and does not have any other occurrence in the whole expression.

We omit the other typing rules as they can trivially be derived from the typing rules defined in Section 5.2. The additional inputs H and o are just passed recursively through the derivation. Note that the rule [ABS] (of judgement $\llbracket \text{abs} \rrbracket : t_0 \dots t_n \xrightarrow{f_x} t$) does not require any additional parameters: as we do not reduce under abstractions, no reference o , abstraction scope or promise scope can appear in the body of an abstraction, and thus the rule [ABS] defined in Section 5.2 can be reused

as is (typing the body with the non-extended version of the type system):

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

PROPOSITION A.4. *If $\Gamma \vdash \llbracket e \rrbracket : t$ (typing judgement from Section 5.2), 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:

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

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 5.3).

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 (i.e. a sequence of references) can be noted \vec{o} . The concatenation of two call stacks \vec{o} and \vec{o}' is noted \vec{o}, \vec{o}' .

Definition A.7 (Heap validity). We say that the heap H is valid for the heap descriptor Δ and for the call stack o_1, \dots, o_n if and only if we have:

Typed environments For every binding $(o \mapsto E) \in H$, there exists a unique i such that $o_i = o$, and

Valid types for every $i, H \setminus \{o_{i+1}, \dots, o_n\}$ matches $\Delta(o_i)$ at reference o_i , and

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 $v \in \text{dom}(H(o_{i'}))$, we have $(\text{owned}(\Delta(o_i)(r)) \text{ and } (H(o_i)(r) = H(o_{i'})(v))) \Rightarrow \text{borrowed}(\Delta(o_{i'})(v))$, and

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 $i' \leq i$, and

Typeable promises for every $(o'' \mapsto \langle e, k, o \rangle) \in H$, if $o \in \text{dom}(H)$ then we have $k = p^{fx}(k' \mathbf{s} cx)$ for some k', cx and fx , and we have $\Delta, H \vdash_o \llbracket e \rrbracket : k' \mathbf{s} cx$.

Definition A.8 (Quadruplet validity). A quadruplet (H, Δ, \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, Δ, \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'.
- {o₁}_{o₂} The reduction rule [DAPP2] applies, we thus have e' = o₁ and H' = H \ o₂. We can conclude with a [DYN] rule as o₁ ≠ o₂ (H(o₂) is an environment E, while H(o₁) is not). Note that we use the *valid promise scoping* property to ensure the new heap H \ o₂ does not have accessible "orphan" promises.
- {e₁}_{o₂} The reduction rule [DCTX1] applies, we conclude by applying the induction hypothesis on e₁.
- {o₁}_{o₂}^{o₃} The reduction rule [DFOR3] applies, we thus have e' = o₁ and H' = H, o₃ ↦ ⟨o₁, k, o₂⟩, we can thus conclude with a [DYN] rule.
- {e₁}_{o₂}^{o₃} The reduction rule [DCTX2] applies, we conclude by applying the induction hypothesis on e₁.
- 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' = 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 dynkind(H(o')) ≤ 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.
- abs ← (o₀ ... o_n) The rule [DAPP1] applies. We must type the body of abs, which can easily be done by extracting the subderivation of the body from [abs] : t₀ ... t_n $\xrightarrow{f_x}$ t and applying Proposition A.4 to type the body under the new heap descriptor Δ, o' ↦ Γ with Γ = rdefs, defs.
- f.i(o₀ ... o_n) The rule [DCAL] applies. We can conclude using the fact that abstractions in F are well-typed, and using Proposition A.4.
- f <sig> (o₀ ... o_n) The rule [DDIS] applies. Using the fact that the abstraction selected by the [DDIS] rule has a smaller signature than the abstraction 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, Δ, \vec{o} , e) be a valid quadruplet. If Δ, H ⊢_o [e] : t, then either e is a reference o' or H[e]_o \rightsquigarrow H'[e']_o for some H' and e'.*

PROOF. We proceed by structural induction on the derivation Δ, H ⊢_o [e] : 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.
- o₁[o₂] = o₃ Our typing derivation must end with a [WRI] rule, which requires o₁ to have type v(I)o!. It yields that dynkind(H(o₁)) ≤ v(I) and thus H(o₁) is a vector. From owner(H, Δ, o') = {o} and the *valid types* property, we can deduce refs(H, o') ≤ 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.

$f <\text{sig}> (o_0 \dots o_n)$ 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 $\emptyset \vdash \llbracket f.i() \rrbracket : t$, 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 $H' \vdash o' \leq_{dyn} t$.

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

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

A.3 Proofs for the flow analysis

In this section, we prove that the flow analysis defined in Section 5.2 ensures a correct initialization of the registers. More precisely, what we want to show is that if an expression is well-flowed, then the two reduction rules [DREGUNDEF] and [DUSEUNDEF] cannot apply after any number of steps. By combining this result with the type safety of the previous section, it yields a type safety theorem for the dynamic semantics of the paper (Section 5.3), where an access to an uninitialized register results in the reduction being stuck.

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

Definition A.13 (Well-defined abstraction). An abstraction abs is well-defined if and only if abs is well-typed and well-flowed.

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

The correctness of the flow analysis will be proved by first introducing a more general version of it that can be applied not only to the source expression but also to the intermediate results of a reduction. We will then prove that the action returned by this generalized flow analysis is preserved by reduction (Lemma A.25), and that an expression that is well-flowed according to this flow analysis cannot be reduced using the rule [DREGUNDEF] or [DUSEUNDEF] (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:

- $R' = R_1 \cup (R_2 \setminus W_1) \cup (R_3 \setminus (W_1 \cup W_2))$

- $W' = W_1 \cup W_2 \cup W_3$
- $U' = U_1 \cup U_2 \cup U_3$
- $C' = C_1 \cup C_2 \cup C_3$

□

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 respects to the order \leq .*

PROOF. Straightforward. □

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

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 5.2 is extended. It now takes a reference o as additional input, and returns heap actions:

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

We omit the other rules as they can trivially be derived from the rules defined in Section 5.2: the additional input o is just passed recursively through the derivation, and each action A returned becomes $o \mapsto A$, for instance:

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

Note that the rule [FABS] (of judgement $F[\text{abs}]$) does not require any additional parameters: as we do not reduce under abstractions, no reference o , abstraction scope or promise scope can appear in the body of an abstraction, and thus the rule [FABS] defined in Section 5.2 can be reused as is (checking the body with the non-extended version of the flow analysis):

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

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

PROOF. Follow from the definition. □

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

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

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

PROOF. Straightforward. □

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 A_1 and A_2 , we have $\text{delay}(A_1) ;; \text{delay}(A_2) = \text{delay}(A_2) ;; \text{delay}(A_1)$.*

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):

$$\text{pending}(\emptyset) = \emptyset$$

$$\text{pending}(H, o \mapsto \langle e, k, o' \rangle) = \text{pending}(H) ;; (o' \mapsto \text{delay}(F[e]))$$

$$\text{pending}(H, o \mapsto 0) = \text{pending}(H)$$

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^{fx}<t>{ 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 abstractions 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).

\square

Definition A.26 (Well-flowedness). For a heap H , expression e and reference o , we say that (H, e, o) is well-flowed, noted $wfl(H, e, o)$, 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 $wfl(H, e, o)$, then the reduction rules $[DREGUNDEF]$ and $[DUSEUNDEF]$ 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 $wfl(H, e, o)$ 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 $[DUSEUNDEF]$. \square

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

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

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

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

We recall that this last theorem, as well as all the lemmas in the section, assume that every abstraction abs in our function table F is well-defined.

THEOREM A.29 (TYPE SAFETY).

If $\emptyset \vdash [f.i()] : t$, then either $\emptyset[f.i()]_o \Rightarrow^\infty$ or $\emptyset[f.i()]_o \Rightarrow^ H'[o']_o$ 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. \square