World Age in Julia

Optimizing Method Dispatch in the Presence of Eval

JULIA BELYAKOVA, Northeastern University, USA

BENJAMIN CHUNG, Northeastern University, USA

JACK GELINAS, Northeastern University, USA

JAMESON NASH, Julia Computing, USA

ROSS TATE, Cornell University, USA

JAN VITEK, Northeastern University, USA and Czech Technical University in Prague, Czech Republic

Dynamic programming languages face semantic and performance challenges in the presence of features, such as `eval`, that can inject new code into a running program. The Julia programming language introduces the novel concept of world age to insulate optimized code from one of the most disruptive side-effects of `eval`: changes to the definition of an existing function. This paper provides the first formal semantics of world age in a core calculus named JULIETTE, and shows how world age enables compiler optimizations, such as inlining, in the presence of `eval`. While Julia also provides programmers with the means to bypass world age, we found that this mechanism is not used extensively: a static analysis of over 4,000 registered Julia packages shows that only 4–9% of packages bypass world age. This suggests that Julia’s semantics aligns with programmer expectations.

CCS Concepts: • **Software and its engineering** → **Language features**; *General programming languages*.

Additional Key Words and Phrases: `eval`, method dispatch, compilation, dynamic languages

ACM Reference Format:

Julia Belyakova, Benjamin Chung, Jack Gelinas, Jameson Nash, Ross Tate, and Jan Vitek. 2020. World Age in Julia: Optimizing Method Dispatch in the Presence of Eval. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 207 (November 2020), 26 pages. <https://doi.org/10.1145/3428275>

1 INTRODUCTION

The Julia programming language [Bezanson et al. 2017] aims to decrease the gap between productivity and performance languages in scientific computing. While Julia provides productivity features such as dynamic types, optional type annotations, reflection, garbage collection, symmetric multiple dispatch, and dynamic code loading, its designers carefully arranged those features to allow for heavy compiler optimization. The key to performance lies in the synergy between language design, language-implementation techniques, and programming style [Bezanson et al. 2018].

The goal of this paper is to shed light on one particular design challenge: how to support `eval` that enables dynamic code loading, *and* achieve good performance. The `eval` construct comes from Lisp [McCarthy 1978] and is found in most dynamic languages, but its expressive power varies from one language to another. Usually `eval` takes a string or a syntax tree as an argument and executes

Authors’ addresses: Julia Belyakova, Northeastern University, USA; Benjamin Chung, Northeastern University, USA; Jack Gelinas, Northeastern University, USA; Jameson Nash, Julia Computing, USA; Ross Tate, Cornell University, USA; Jan Vitek, Northeastern University, USA, Czech Technical University in Prague, Czech Republic.

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

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/11-ART207

<https://doi.org/10.1145/3428275>

```
> x = 3.14
> f(x) = (
    eval(:(x = 0));
    x * 2)

> f(42) # 84
> x     # 0
```

Fig. 1. Scope of eval in Julia

```
> (defn g [] 2)
> (defn f [x]
  (eval `(defn g [] ~x))
  (* x (g)))
> (f 42) ; 1764
> (g)    ; 42
> (f 42) ; 1764
```

Fig. 2. Eval in Clojure

```
> g() = 2
> f(x) = (
    eval(:(g()=$x));
    x * g())
> f(42) # 84
> g()   # 42
> f(42) # 1764
```

Fig. 3. Eval in Julia

it in some environment. In JavaScript and R, `eval` may execute in the current lexical environment; in Lisp and Clojure, it is limited to the “top level”. On this spectrum, Julia takes the latter approach, which enables compiler optimizations that would otherwise be unsound. For example, in a program in Fig. 1, multiplication `x*2` in the body of function `f` can be safely optimized to an efficient integer multiplication for the call `f(42)`. This is because `eval` only accesses the top-level environment and thus cannot change the value of a local parameter `x`, which is known to be the integer 42. For a global `x`, such an optimization would be unsound.

What is unique about the design of `eval` in Julia is the treatment of function definitions. Many compilers rely on the information about functions for optimizations, but those optimizations can be jeopardized by the presence of `eval`. To explain how Julia handles the interaction of `eval` and functions, we contrast it with the Clojure language. Fig. 2 shows a Clojure program with a call to a function `f` which, within its body, updates function `g` by invoking `eval`. Then, the call to `f` returns 1764 because the new definition of `g` is used. Fig. 3 shows a Julia equivalent of the same program. Here, the second call to `f` returns 1764 just like in Clojure, but *the first call returns 84*. This is because, while the first invocation of `f` is running, it does not see the redefinition of `g` made by `eval`: the redefinition becomes visible only after the first call (to `f(42)`) returns to the top level. From the compiler’s point of view, this means that calling `eval` does not force recompilation of any methods that are “in-flight.” Thus, it is safe to devirtualize, specialize, and inline functions in the presence of `eval` without the need for deoptimization. For example, `x*g()` can be safely replaced with `x*2` in `f` for the first call `f(42)`.

Julia made the choice to restrict access to newly defined methods due to pressing performance concerns. Julia heavily relies on symmetric multiple dispatch [Bobrow et al. 1986], which allows a function to have multiple implementations, called methods, distinguished by their parameter type annotations. At run time, a call is dispatched to the most specific method applicable to the types

```
*{x::Int, y::Int} = mul_int(x, y)
*{x::Float64, y::Float64} = mul_float(x, y)
*{a::Number, b::AbstractVector} = ...
*{x::Bool, y::Bool} = x & y
```

Fig. 4. Multiple definitions of *

of its arguments. While some functions might have only one method, plenty have dozens or even hundreds of them. For example, the multiplication function alone has 357 standard methods (see an excerpt in Fig. 4). If Julia were to always use generic method invocation to dispatch `*`, programs would become unbearably slow. By constraining `eval`, the compiler can avoid generic invocations. In Fig. 1, the compiler can pick and inline the right definition of `*` when compiling `f(42)`. It should be noted that this optimization-friendly `eval` semantics does not apply to data. Function definitions are treated differently from variables, and changes to global variables (such as in Fig. 1) can be observed immediately.

Arguably, despite being unusual, Julia’s semantics is easy to understand for programmers. There is always a clear point where new definitions become visible—at the top level—and thus, users can avoid surprises and dependence on the exact position of `eval` in the code. However, in case the default semantics is not desirable, Julia also provides an escape hatch: the built-in function `invokelatest(f)`, which forces the implementation to invoke the most recent definition of `f`. A slower alternative to `invokelatest` is to call `f` within `eval`, which always executes in the top level.

This language mechanism that delays the effect of `eval` on function definitions is called *world age*. In the Julia documentation, world age is described operationally [Bezanson et al. 2018]: every method defined in a program is associated with an age, and for each function call, Julia ensures that the current age is larger than the age of the method about to be invoked. One can think of the world age as a counter that allows the implementation to ignore all methods that were born after the last top-level call started. Much of its specification is tied to implementation details and efficiency considerations. Our contributions are as follows:

- *A core calculus for world age*: We introduce JULIETTE, a calculus that models the notion of world age abstractly. In the calculus, the implementation-oriented world-age counters are replaced with method tables that are explicitly copied at the top level, and `eval` is simplified down to an operation that evaluates its argument in a specific method table.
- *Formalization of optimizations*: We formalize and prove correct three compiler optimizations, namely inlining, devirtualization, and specialization.
- *Corpus analysis*: We analyze Julia packages to understand how `eval` is used, and estimate the potential impact of world age on library code. We also identify a number of programming patterns by manual inspection of selected packages.
- *Testing the semantics*: We develop a Redex model of our calculus and optimizations to allow rapid experimentation and testing.

The corpus analysis and the Redex model are publicly available.¹ The formalization with detailed proofs can be found in the extended version of the paper [Belyakova et al. 2020].

2 BACKGROUND

We start with an overview of the features of Julia relevant to our work, and review related work.

2.1 Julia Overview

Despite the extensive use of types and type annotations for dispatch and compiler optimizations, Julia is not statically typed. A formalization of types and subtyping is provided by Zappa Nardelli et al. [2018], and a general introduction to the language is given by Bezanson et al. [2017].

Values. Values are either instances of *primitive types*—sequences of bits—or *composite types*—collection of fields holding values. Every value has a concrete type (or tag). This tag is either inferred statically or stored in the boxed value. Tags are used to resolve multiple dispatch semantically and can be queried with `typeof`.

Types. Programmers can declare three kinds of user-defined types: *abstract types*, *primitive types*, and *composite types*. Abstract types cannot be instantiated, while concrete types can. For example, `Float64` is concrete, and is a subtype of abstract type `Number`. Concrete types have no subtypes. Additionally, user-defined type constructors can have bounded type parameters and can declare up to a single supertype.

¹<https://github.com/julbinb/juliette-wa>

Annotations. Type annotations include a number of built-in type constructors, such as union and tuple types. Tuple types, written `Tuple{A, ...}`, describe immutable values that have a special role in the language: every method takes a single tuple argument. The `::` operator ascribes a type to a definition. We will use τ to denote annotations.

Subtyping. The subtyping relation, $<$, is used in run-time casts and multiple dispatch. Julia combines nominal subtyping, union types, iterated union types, covariant and invariant constructors, and singleton types. Tuple types are covariant in their parameters, so, for instance, `Tuple{Float64,Float64}` is a subtype of `Tuple{Number,Number}`.

Multiple dispatch. A function can have multiple methods where each method declares what argument types it can handle; an unspecified type defaults to `Any`. At run time, dispatching a call `f(v)` amounts to picking the best applicable method from all the methods of function `f`. For this, the dispatch mechanism first filters out methods whose type annotations τ are not a supertype of the type tag of `v`. Then it takes the method whose type annotation τ_i is the most specific of the remaining ones. If the set of applicable methods is empty, or there is no single best method, a run-time error is raised.

Reflection. Julia provides a number of built-in functions for run-time introspection and meta-programming. For instance, the methods of any function `f` may be listed using `methods(f)`. All the methods are stored in a special data structure, called the *method table*. It is possible to search the method table for methods accepting a given type: for instance, `methods(*, (Int,Float64))` will show methods of `*` that accept an integer-float pair. The `eval` function takes an expression object and evaluates it in the global environment of a specified module. For example, `eval(:(1+2))` will take the expression `:(1+2)` and return `3`.

`Eval` is frequently used for meta-programming as part of code generation. For example, Fig. 5 generalizes some of the basic binary operators to three arguments, generating four new methods. Instead of building expressions explicitly, one can also invoke the parser on a string. For instance, `eval(Meta.parse("id(x) = x"))` creates an identity method.

```
for op in (:+, :*, :&, :|)
    eval(:($op(a,b,c) = $op($op(a,b),c)))
end
```

Fig. 5. Code generation

2.2 Related Work

This paper is concerned with controlling the visibility of function definitions. Most programming languages control *where* definitions are visible, as part of their scoping mechanisms. Controlling *when* function definitions become visible is less common.

Languages with an interactive development environment had to deal with the addition of new definitions for functions from the start [McCarthy 1978]. Originally, these languages were interpreted. In that setting, allowing new functions to become visible immediately was both easy to implement and did not incur any performance overhead.

Just-in-time compilation changed the performance landscape, allowing dynamic languages to have competitive performance. However, this meant that to generate efficient code, compilers had to commit to particular versions of functions. If any function is redefined, all code that depends on that function must be recompiled; furthermore, any function currently executing has to be deoptimized using mechanisms such as on-stack-replacement [Hölzle et al. 1992]. The drawback of deoptimization is that it makes the compiler more complex and hinders some optimizations. For

example, a special `assume` instruction is introduced as a barrier to optimizations by Flückiger et al. [2018], who formalized the speculation and deoptimization happening in a model compiler.

Java allows for dynamic loading of new classes and provides sophisticated controls for where those classes are visible. This is done by the class-loading framework that is part of the virtual machine [Liang and Bracha 1998]. Much research happened in that context to allow the Java compiler to optimize code in the presence of dynamic loading. Detlefs and Agesen [1999] describe a technique, which they call preexistence, that can devirtualize a method call when the receiver object predates the introduction of a new class. Further research looked at performing dependency analysis to identify which methods are affected by the newly added definitions, to be then recompiled on demand [Nguyen and Xue 2005]. Glew [2005] describes a type-safe means of inlining and devirtualization: when newly loaded code is reachable from previously optimized code, these optimizations must be rechecked.

Controlling *when* definitions take effect is important in dynamic software updating, where running systems are updated with new code [Cook and Lee 1983]. Stoye et al. [2007] introduce a calculus for reasoning about representation-consistent dynamic software updating in C-like languages. One of the key elements for their result is the presence of an `update` instruction that specifies when an update is allowed to happen. This has similarities to the world-age mechanism described here.

Substantial amounts of effort have been put into building calculi that support `eval` and similar constructs. For example, Politz et al. [2012] described the ECMAScript 5.1 semantics for `eval`, among other features. Glew [2005] formalized dynamic class loading in the framework of Featherweight Java, and Matthews and Findler [2008] developed a calculus for `eval` in Scheme. These works formalize the semantics of dynamically modifiable code in their respective languages, but, unlike Julia, the languages formalized do not have features explicitly designed to support efficient implementation.

3 WORLD AGE IN JULIA

The world-age mechanism in Julia limits the set of methods that can be invoked from a given call site. World age fixes the set of method definitions reachable from the currently executing method, isolating it from dynamically generated ones. In turn, this allows the compiler to optimize code without need for deoptimization, and limits the number of required synchronization points in a multi-threaded program. If full access to methods is required, however, Julia provides escape hatches to bypass world age by sacrificing performance.

3.1 Defining World Age

The primary goal of the world-age mechanism is to align the language's semantics with the assumptions made by the Julia just-in-time compiler's optimizations. Semantically, newly added methods (i.e. ones defined using `eval`) only become visible when execution returns to the top level, and the set of callable methods for an execution is fixed when it leaves the top level. Compilation of methods is triggered—only at the top level—when one of the following holds: (1) a function is called with previously unobserved types of arguments, or (2) a previously compiled function needs to be recompiled due to a change in its own definition or one of its dependencies. Since the set of visible methods gets fixed at a top-level call, and compilation only occurs from the top level, the compiler may assume that the currently known set of methods is complete and can optimize accordingly.

For performance reasons, the world-age mechanism is implemented by a simple monotonic counter. The counter is incremented every time a method is defined, and its value becomes the method's "birth age". Every method also can store a "death age" (that is initially infinity), which is set when it is replaced or deleted. Methods with their ages are stored in a global data structure

Fig. 6. Age ranges

```

> function ntl()
  eval:(f() = 1))
  f()
end
> ntl()
ERROR: MethodError: no method matching f()
The applicable method may be too new:
running in world age 26806, while current
world is 26807. Closest candidates are:
  f() at REPL[10]:2

```

Fig. 7. Age error

called a method table. The birth and death ages of a method determine the minimum and maximum world age from which the method can be invoked. This is illustrated in Fig. 6. Here, we define functions f and g , where f has only one method, whereas g has two methods, one replacing the other. Let us observe what definitions are visible at each step. Since f is defined once at time t_1 (i.e. when the world-age counter is equal to t_1) and never redefined, it has a birth age of t_1 and a death age of ∞ ; thus, it can be used anytime after t_1 and forevermore. In contrast, method g_1 , created at time t_2 , is redefined at time t_3 , so its birth age is t_2 and death age is t_3 ; therefore, g_1 can be called from world age t where $t_2 \leq t \leq t_3$. Finally, g_2 is never redefined, so its age ranges from t_3 to ∞ .

This language design can be restrictive. The easiest way to run afoul of world age is to attempt to define a method with `eval` and call it immediately thereafter. In Fig. 7, function `ntl` creates a new function f using `eval` and attempts to call it immediately without returning to the top level. Let us dissect the example and its error message. Function `ntl` was invoked from the top level with age of 26806, thus limiting the set of visible methods to the ones born by this time. Then, `ntl` used `eval` to define f , giving it a birth age of 26807. Finally, to call f , `ntl` needs a method of f that was born by 26806, but none exists. Since the only method of f was created at 26807, a `MethodError` is raised indicating that no method was found.

3.2 Breaking the Age Barrier

There are situations in which the world-age mechanism is too restrictive: for example, when a program wishes to programmatically generate code and then use it immediately. To accommodate these circumstances, Julia provides two ways for programmers to execute code ahead of its birth age. The first is `eval` itself, which executes its arguments as if they were at the top level, thus allowing any existing method to be called. However, `eval` needs to interpret arbitrary ASTs and is rather slow. Luckily, in many circumstances, the program only wishes to bypass the world-age restriction for a single function call. For this, one can use `invokelatest`, a built-in function that calls its argument from the latest world age. While substantially slower than a normal call, `invokelatest` is faster than `eval`. Moreover, `invokelatest` is passed arguments directly from the calling context so that values do not need to be inserted into `eval`'s AST. Both `eval` and `invokelatest` can be used to amend the example shown in Fig. 7 to call f in the latest world age. If we replace the bare call to f with a call to `eval:(f())`, then the call to `ntl` will produce 1. Similarly, `Base.invokelatest(f)` will get the same result.

Using either of these mechanisms, programmers can opt out from the limitation imposed by world age, but this comes with performance implications. Since neither `invokelatest` nor `eval` can be optimized, and both can kick off additional JIT compilation, they can have substantial performance impact. However, this impact is limited to only these explicitly impacted call sites.

As a result, programmers can carefully design their programs to minimize the number of broken barriers, thus minimizing the performance impact of the dynamism.

3.3 World Age in Practice

We have argued that world age is useful for performance of Julia programs, but does its semantics match programmers' expectations? We propose a slightly indirect answer to this, analyzing a corpus of programs and observing how often they use the two escape hatches mentioned above. Of those, `invoke_latest` is the clearest indicator, as there is no reason to use it except to bypass the world age. Similarly, `eval`'ing a function call means evaluating that call in the latest world age, thereby allowing it to see the latest method definitions. The `eval` indicator is imprecise, however, as there are uses of `eval` that are not impacted by world age.

We take as our corpus all 4,011 registered Julia packages as of August 2020. The results of statically analyzing the code base are shown in Fig. 8. The analysis shows that 2,846 of the 4,011 packages used neither `eval` nor `invoke_latest`, and thus are definitely age agnostic. Of the remaining packages, 1,094 used `eval` only, and so *could* be impacted. 15 packages used `invoke_latest` only, and some 56 used both. We can reasonably presume that at least these latter 71 packages are impacted by world age because they bypass it using `invoke_latest`. Drawing conclusions about `eval`-using packages requires further analysis.

Fig. 8. `eval` and `invoke_latest` use by package Fig. 9. Static use of AST forms in all packages

To understand if packages that only use `eval` are impacted by world age, we statically analyzed the location of calls to `eval` and their arguments by parsing files that contain `eval`. For each call, we classify the argument ASTs, recursively traversing them and counting occurrences of relevant nodes. The analysis is conservative: it assumes that an AST that is not statically obvious (such as a variable) could contain anything. Fig. 9 shows how many packages use world-age relevant AST forms. Only uses of `eval` from within functions—where world-age could be relevant—are shown; top-level uses of `eval`—which cannot be affected by world age—are filtered out. The “all-others” category encompasses all AST forms not relevant to world age. While this aggregate is, taken as a whole, more common than any other single AST form, none of the constituent AST forms is more prevalent than function calls. Therefore, most common arguments to `eval` are function definitions, followed by function calls and loading of modules and other files.

Using the results of the static analysis, we estimate that about 4–9% of the 4,011 packages might be affected by world age. The upper bound (360 packages) is a conservative estimate, which includes 289 packages with potentially world-age-related calls to `eval` but without calls to `invoke_latest`, and the 71 packages that use `invoke_latest`. The lower bound (186 packages) includes 115 `invoke_latest`-free packages that call `eval` with both function definitions and function calls, and the 71 packages with `invoke_latest`.

To validate our static results, we dynamically analyzed 32 packages out of the 186 identified as possibly affected by world age. These packages were selected by randomly sampling a subset of 49 packages, which was then further reduced by removing packages that did not run, whose tests failed, or that did not call `eval` or `invoke_latest` at least once. Over this corpus, the dynamic

analysis was implemented by adding instrumentation to record calls to `eval` and `invokelatest`, recording the ASTs and functions, respectively, as well as the stack traces for each invocation.

Fig. 10. Static use of AST forms used by package

Fig. 11. Dynamic use of AST forms used by package

The results of the static and dynamic analysis of the 32 packages are given in Fig. 10 and Fig. 11, respectively. Both analysis methods agreed that the most common world-age-relevant use of `eval` was to define functions, followed by making function calls and importing other packages. In general, the dynamic analysis was able to identify more packages that used each AST form, as it can examine every AST ran through `eval`, not only statically declared ones. However, this accuracy is dependent on test coverage.

3.4 Programming with World Age

We now turn to common patterns found by manual inspection of select packages of the corpus.

Boilerplating. The most common use of `eval` is to automatically generate code for boilerplate functions. These generated functions are typically created at the top-level so that they can be used by the rest of the program. Consider the `DualNumbers.jl` package, which provides a common dual number representation for automatic differentiation. A dual number, which is a pair of the normal value and an “epsilon”, which represents the derivative of the value, should support the same operations as any number does and mostly defers to the standard operations. For example, the `real` function, which gets the real component of a number when applied to a dual number should recurse into both the actual and epsilon value. `Eval` can generate all of the needed implementations at package load time (`@eval` is a macro that passes its argument to `eval` as an AST).

```
for op in (:real, :imag, :conj, :float, :complex)
    @eval Base.$op(z::Dual) = Dual($op(value(z)), $op(epsilon(z)))
end
```

A common sub-pattern is to generate proxies for interfaces defined by an external system. For this purpose, the `CxxWrap.jl` library uses `eval` at the top level to generate (with the aid of a helper method that generates the ASTs) proxies for arbitrary C++ libraries.

```
eval(build_function_expression(func, funcidx, julia_mod))
```

Defensive callbacks. The most widely used pattern for `invokelatest` deals with function values of unknown age. For example, when invoking a callback provided by a client, a library may protect itself against the case where the provided function was defined after the library was loaded. There are two forms of this pattern. The simplest uses `invokelatest` for all callbacks, such as the library `Symata.jl`:

```
for hook in preexecute_hooks
    invokelatest(hook)
end
```


Every hook in `preexecute_hooks` is protected against world-age errors (at the cost of slower function calls). To avoid this slowdown, the second common pattern catches world-age exceptions and falls back to `invokelatest` such as in from the `Genie.jl` web server:

```
fr::String = try
    f()::String
catch
    Base.invokelatest(f)::String
end
```

This may cause surprises, however. If a sufficiently old method exists, the call may succeed but invoke the wrong method.² This pattern may also catch unwanted exceptions and execute `f` twice, including its side-effects.

Domain-specific generation. As a language targeting scientific computing, Julia has a large number of packages that do various symbolic domain reasoning. Examples include symbolic math libraries, such as `Symata` and `GAP`, which have the functionality to generate executable code for symbolic expressions. `Symata` provides the following method to convert an internal expression (a `Mxpr`) into a callable function. Here, `Symata` uses a translation function `mxpr_to_expr` to convert the `Symata` `mxpr` into a Julia `Expr`, then wraps it in a function definition (written using explicit AST forms), before passing it to `eval`.

```
function Compile(a::Mxpr{:List}, body)
    aux = MtoECompile()
    jexpr = Expr(:function,
        Expr(:tuple, [mxpr_to_expr(x, aux) for x in margs(a)]...),
        mxpr_to_expr(body, aux))
    Core.eval(Main, jexpr)
end
```

Bottleneck. Generated code is commonly used in Julia as a way to mediate between a high-level DSL and a numerical library. Compilation from the DSL to executable code can dramatically improve efficiency while still retaining a high-level representation. However, functions generated thusly cannot be called from the code that generated them, since they are too new. Furthermore, this code is expected to be high-performance, so using `invokelatest` for every call is not acceptable. The bottleneck pattern overcomes these issues. The idea is to split the program into two parts: one that generates code, and another that runs it. The two parts are bridged with a single `invokelatest` call (the “bottleneck”), allowing the second part to call the generated code efficiently. The pattern is used in the `DiffEqBase` library, part of the `DifferentialEquations` family of libraries that provides numerical differential equation solvers.

```
if hasfield(typeof(_prob),:f) && hasfield(typeof(_prob.f),:f) &&
    typeof(_prob.f.f) <: EvalFunc
    Base.invokelatest(__solve,_prob,args...; kwargs...)
else
    __solve(_prob,args...;kwargs...)
end
```

²In Julia, higher-order functions are passed by name as generic functions, so a callback will be subject to multiple dispatch.

Here, if `_prob` has a field `f`, which has another field `f`, and the type of said inner-inner `f` is an `EvalFunc` (an internally-defined wrapper around any function that was generated with `eval`), then it will invoke the `__solve` function using `invoke_latest`, thus allowing `__solve` to call said method. Otherwise, it will do the invocation normally.

Superfluous eval. This is a rare anti-pattern, probably indicating a misunderstanding of world age by some Julia programmers. For example, `Alpine.jl` package has the following call to `eval`:

```
if isa(m.disc_var_pick, Function)
    eval(m.disc_var_pick)(m)
```

Here, `eval(m.disc_var_pick)` does nothing useful but imposes a performance overhead. Because `m.disc_var_pick` is already a function value, calling `eval` on it is similar to using `eval(42)` instead of `42` directly; this neither bypasses the world age nor even interprets an AST.

Name-based dispatch. Another anti-pattern uses `eval` to convert function names to functions. For example, `ClassImbalance.jl` package chooses a function to call, using its uninterpreted name:

```
func = (labeltype == :majority) ? :argmax : :argmin
indx = eval(func)(counts)
```

It would be more efficient to operate with function values directly, i.e. `func = ... : argmin` and then call it with `func(counts)`. Similarly, when a symbol being looked up is generated dynamically, as it is in the following example from `TextAnalysis.jl`, the use of `eval` could be avoided.

```
newscheme = uppercase(newscheme)
if !in(newscheme, available_schemes) ...
newscheme = eval(Symbol(newscheme))()
```

This pattern could be replaced with a call `getfield(TextAnalysis, Symbol(newscheme))`, where `getfield` is a special built-in function that finds a value in the environment by its name. Using `getfield` would be more efficient than `eval`.

4 JULIETTE, A WORLD AGE CALCULUS

To formally study world age, we propose a core calculus, named JULIETTE, that captures the essence of Julia’s semantics and permits us to reason about the correctness of some of the optimizations performed by the compiler.

Designing such a calculus is always an exercise in parsimony, balancing the need to highlight principles while avoiding entanglements with particular implementation choices. The first decision to grapple with is how to represent world age. While efficient, counters are also pervasive and cause confusion.³ Furthermore, they obscure reasoning about program-state equivalence; two programs with different initial counter values could, if care is not taken, appear different. Dispensing with the counters used by Julia’s compiler is appealing.

An alternative that we chose is a more abstract representation of world age, one that captures its intent: control over method visibility. JULIETTE uses *method tables* to represent sets of methods available for dispatch. The *global table* is the method table that records all definitions and always reflects the “true age” of the world; the global table is part of JULIETTE program state. *Local tables*

³Although Julia’s documentation attempts to explain world age [v1 2020], questions such as [this one](#) pop up periodically.

are method tables used to resolve method dispatch during execution and may lag behind the global table when new functions are introduced. Local tables are then baked into program syntax to make them explicit during execution. As in Julia, JULIETTE separates method tables (which represent code) from data: as mentioned in Sec. 1, the world-age semantics only applies to code. As global variables interact with `eval` in the standard way, we omit them from the calculus.

The treatment of methods is similar in both JULIETTE and Julia up to (lexically) local method definitions. In both systems, a generic function is defined by the set of methods with the same name. In Julia, local methods are syntactic sugar for global methods with fresh names. For simplicity, we do not model this aspect of Julia: JULIETTE methods are always added to the global method table. All function calls are resolved using the set of methods found in the current local table. A function value m denotes the name of a function and is not itself a method definition. Then, since JULIETTE omits global variables, its global environment is entirely captured by the global method table.

Although in Julia `eval` incorporates two features—top-level evaluation and quotation⁴—only top-level evaluation is relevant to world age, and this is what we model in JULIETTE. Instead of an `eval` construct, the calculus has operations for evaluating expressions in different method-table contexts. In particular, JULIETTE offers a *global evaluation construct* $\langle e \rangle$ (pronounced “banana brackets”) that accesses the most recent set of methods. This is equivalent to `eval`’s behavior, which evaluates in the latest world age. Since JULIETTE does not have global variables, $\langle e \rangle$ reads from the local environment directly instead of using quotation.

Every function call $m(\bar{v})$ in JULIETTE gets resolved in the closest enclosing local method table M by using an *evaluation-in-a-table* construct $\langle m(\bar{v}) \rangle_M$. Any top-level function call first takes a snippet of the current global table and then evaluates the call in that *frozen* snippet. That is, $\langle m(\bar{v}) \rangle$ steps to $\langle m(\bar{v}) \rangle_M$ where M is the current global table. Thus, once a snippet of the global table becomes local table, all inner function calls of $m(\bar{v})$ will be resolved using this table, reflecting the fact that a currently executing top-level function call does not see updates to the global table.

To focus on world age, JULIETTE omits irrelevant features such as loops or mutable variables. Furthermore, the calculus is parameterized over values, types, type annotations, a subtyping relation, and primitive operations. For the purposes of this paper, only minimal assumptions are needed about those.

4.1 Syntax

The surface syntax of JULIETTE is given in Fig. 12. It includes method definitions md , function calls $e(\bar{e})$, sequencing $e_1 ; e_2$, global evaluation $\langle e \rangle$, evaluation in a table $\langle e \rangle_M$, variables x , values v , primitive calls $\delta_l(\bar{e})$, type tags σ , and type annotations τ . Values v include `unit` (unit value, called `nothing` in Julia) and m (generic function value). Primitive operators δ_l represent built-in functions such as `Base.mul_int`. Type tags σ include $\mathbb{1}$ (unit type, called `Nothing` in Julia) and $\mathbb{f}_m$ (tag of function value m). Type annotations τ include $\top \in \tau$ ($\top$ is the top type, called `Any` in Julia) and $\sigma \subseteq \tau$ (all type tags serve as valid type annotations).

4.2 Semantics

The internal syntax of JULIETTE is given in the top of Fig. 13. It includes evaluation result r (either value or error), method table M , and two evaluation contexts, X and C , which are used to define small-step operational semantics of JULIETTE. Evaluation contexts X are responsible for simple sequencing, such as the order of argument evaluation; these contexts never contain global/table evaluation expressions $\langle \cdot \rangle$ and $\langle \cdot \rangle_M$. World evaluation contexts C , on the other hand, capture the full grammar of expressions.

⁴Represented with the `$` operator in Julia, as in `eval(: (g() = $x))` in Fig. 3.

e	$::=$	<i>Expression</i>			
	$ $	v	value	v	$::=$... <i>Value</i>
	$ $	x	variable		$ $ unit unit value
	$ $	$e_1 ; e_2$	sequencing		$ $ m generic function
	$ $	$\delta_l(\bar{e})$	primop call		
	$ $	$e(\bar{e})$	function call	σ	$::=$... <i>Type tag</i>
	$ $	md	method definition		$ $ $\mathbb{1}$ unit type
	$ $	$\langle e \rangle$	global evaluation		$ $ $\mathbb{f}_m$ type tag of function m
	$ $	$\langle e \rangle_M$	evaluation in a table		
p	$::=$	$\langle e \rangle$	<i>Program</i>	τ	$::=$... <i>Type annotation</i>
					$ $ $\top$ top type
md	$::=$	$\langle m(\bar{x} :: \bar{\tau}) = e \rangle$	<i>Method definition</i>		

Fig. 12. Surface syntax

Program state is a pair $\langle M, \mathbf{C}[e] \rangle$ of a global method table M and an expression $\mathbf{C}[e]$. We define the semantics of the calculus using two judgments: a normal small-step evaluation denoted by $\langle M, \mathbf{C}[e] \rangle \rightarrow \langle M', \mathbf{C}[e'] \rangle$, and a step to an error $M \vdash \mathbf{C}[e] \rightarrow \text{error}$. The $\text{typeof}(v) \in \sigma$ operator returns the tag of a value. We require that $\text{typeof}(\text{unit}) = \mathbb{1}$ and $\text{typeof}(m) = \mathbb{f}_m$. We write $\text{typeof}(\bar{v})$ as a shorthand for $\overline{\text{typeof}(\bar{v})}$. Function $\Delta(l, \bar{v}) \in \tau$ computes primop calls, and function $\Psi(l, \bar{\sigma}) \in \sigma$ indicates the tag of l 's return value when called with arguments of types $\bar{\sigma}$. These functions have to agree, i.e. $\forall \bar{v}, \bar{\sigma}. (\text{typeof}(\bar{v}) = \bar{\sigma} \wedge \Delta(l, \bar{v}) = v' \implies \text{typeof}(v') = \Psi(l, \bar{\sigma}))$. The subtyping relation $\tau_1 <: \tau_2$ is used for multiple dispatch. We require that $\tau <: \top$ ($\top$ is indeed the top type) and $\sigma_1 <: \sigma_2 \iff \sigma_1 \equiv \sigma_2$ (tags are final, i.e. do not have subtypes).

Normal Evaluation. These rules capture successful program executions. Rule E-SEQ is completely standard: it throws away the evaluated part of a sequencing expression. Rules E-VALGLOBAL and E-VALLOCAL pass value v to the outer context. This is similar to Julia where `eval` returns the result of evaluating the argument to its caller. Rule E-MD is responsible for updating the global table: a method definition md will extend the current global table M into $M \bullet md$, and itself evaluate to m , which is a function value. Note that E-MD only extends the method table and leaves existing definitions in place. If the table contains multiple definitions of a method with the same signature, it is then the dispatcher's responsibility to select the right method; this mechanism is described below in more detail.

The two call forms E-CALLGLOBAL and E-CALLLOCAL form the core of the calculus. The rule E-CALLGLOBAL describes the case where a method is called directly from a global evaluation expression. In Julia, this means either a top-level call, an `invoke_late` call, or a call within `eval` such as `eval(: (g(...)))`. The “direct” part is encoded with the use of a simple evaluation context X . In this global-call case, we need to save the current method table into the evaluation context for a subsequent use by E-CALLLOCAL. To do this, we annotate the call $m(\bar{v})$ with a copy of the current global method table M , producing $\langle m(\bar{v}) \rangle_M$.

To perform a local call—or, equivalently, a call after the invocation has been wrapped in an annotation specifying the current global table—E-CALLLOCAL is used. This rule resolves the call according to the tag-based multiple-dispatch semantics in the “deepest” method table M' (the use of X makes sure there are no method tables between M' and the call). Once an appropriate method has been found, it proceeds as a normal invocation rule would, replacing the method invocation with the substituted-for-arguments method body. Note that the body of the method is still wrapped

			$X ::=$	<i>Simple evaluation context</i>
			$\square$	hole
			$X; e$	sequence
			$\delta_l(\bar{v} \ X \ \bar{e})$	primop call (argument)
			$X(\bar{e})$	function call (callee)
			$v(\bar{v} \ X \ \bar{e})$	function call (argument)
$r ::=$	<i>Result</i>			
v	value			
error	error			
$M ::=$	<i>Method table</i>			
$\emptyset$	empty table	$C ::=$	<i>World evaluation context</i>	
$M \bullet md$	table extension	X	simple context	
		$X \llbracket C \rrbracket$	global evaluation	
		$X \llbracket C \rrbracket_M$	evaluation in a table M	
E-SEQ			E-PRIMOP	E-MD
$\frac{}{\langle M, C \llbracket v; e \rrbracket \rangle \rightarrow \langle M, C \llbracket e \rrbracket \rangle}$			$\frac{\Delta(l, \bar{v}) = v'}{\langle M, C \llbracket \delta_l(\bar{v}) \rrbracket \rangle \rightarrow \langle M, C \llbracket v' \rrbracket \rangle}$	$\frac{md \equiv \triangleleft m(\bar{x} :: \bar{\tau}) = e \triangleright}{\langle M, C \llbracket md \rrbracket \rangle \rightarrow \langle M \bullet md, C \llbracket m \rrbracket \rangle}$
E-CALLGLOBAL			E-CALLLOCAL	
$\frac{}{\langle M, C \llbracket X[m(\bar{v})] \rrbracket \rangle \rightarrow \langle M, C \llbracket X \llbracket m(\bar{v}) \rrbracket_M \rrbracket \rangle}$			$\frac{\text{typeof}(\bar{v}) = \bar{\sigma} \quad \text{getmd}(M', m, \bar{\sigma}) = \triangleleft m(\bar{x} :: \bar{\tau}) = e \triangleright}{\langle M, C \llbracket X[m(\bar{v})] \rrbracket_{M'} \rangle \rightarrow \langle M, C \llbracket X[e[\bar{x} \mapsto \bar{v}]] \rrbracket_{M'} \rangle}$	
E-VALGLOBAL			E-VALLOCAL	
$\frac{}{\langle M, C \llbracket v \rrbracket \rangle \rightarrow \langle M, C \llbracket v \rrbracket \rangle}$			$\frac{}{\langle M, C \llbracket v \rrbracket_{M'} \rangle \rightarrow \langle M, C \llbracket v \rrbracket \rangle}$	
E-VARERR			E-PRIMOPERR	E-CALLEERR
$\frac{}{M \vdash C \llbracket x \rrbracket \rightarrow \text{error}}$			$\frac{\Delta(l, \bar{v}) = \text{error}}{M \vdash C \llbracket \delta_l(\bar{v}) \rrbracket \rightarrow \text{error}}$	$\frac{v_e \neq m}{M \vdash C \llbracket v_e(\bar{v}) \rrbracket \rightarrow \text{error}}$
			E-CALLERR	
			$\frac{\text{typeof}(\bar{v}) = \bar{\sigma} \quad \text{getmd}(M', m, \bar{\sigma}) = \text{error}}{M \vdash C \llbracket X[m(\bar{v})] \rrbracket_{M'} \rightarrow \text{error}}$	
$\text{getmd}(M, m, \bar{\sigma}) = \min(\text{applicable}(\text{latest}(M), m, \bar{\sigma}))$				
$\text{latest}(M) = \text{latest}(\emptyset, M)$				
$\text{latest}(mds, \emptyset) = mds$				
$\text{latest}(mds, M \bullet md) = \text{latest}(mds \cup md, M) \text{ if } \neg \text{contains}(mds, md)$				
$\text{latest}(mds, M \bullet md) = \text{latest}(mds, M) \text{ if } \text{contains}(mds, md)$				
$\text{applicable}(mds, m, \bar{\sigma}) = \{\triangleleft m(\bar{x} :: \bar{\tau}) = e \triangleright \in mds \mid \bar{\sigma} <: \bar{\tau}\}$				
$\min(mds) = \triangleleft m(\bar{x} :: \bar{\tau}) = e \triangleright \in mds \text{ such that } \forall \triangleleft m(\bar{x} :: \bar{\tau}') = e' \triangleright \in mds. \bar{\tau} <: \bar{\tau}'$				
$\min(mds) = \text{error otherwise}$				
$\text{contains}(mds, md) = \exists md' \in mds \text{ such that}$				
$(md \equiv \triangleleft m(\bar{x} :: \bar{\tau}) = e \triangleright) \wedge (md' \equiv \triangleleft m(\bar{x} :: \bar{\tau}') = e' \triangleright) \wedge \bar{\tau} <: \bar{\tau}' \wedge \bar{\tau}' <: \bar{\tau}$				

Fig. 13. Internal Syntax and Semantics

in the $\langle \langle \rangle_M \rangle$ context. This ensures that nested calls will be resolved in the same table (unless they are more deeply wrapped in a global evaluation $\langle \langle \rangle \rangle$).

An auxiliary meta-function $\text{getmd}(M, m, \bar{\sigma})$, which is used to resolve multiple dispatch, is defined in the bottom of Fig. 13. This function returns the most specific method applicable to arguments with type tags $\bar{\sigma}$, or errs if such a method does not exist. If the method table contains multiple equivalent methods, older ones are ignored. For example, for the program

$$\langle \langle \triangleleft g() = 2 \triangleright ; \triangleleft g() = 42 \triangleright ; g() \rangle \rangle,$$

function call $g()$ is going to be resolved in the table $(\emptyset \bullet \triangleleft g() = 2 \triangleright) \bullet \triangleleft g() = 42 \triangleright$, which contains two equivalent methods (we call methods equivalent if they have the same name and their argument type annotations are equivalent with respect to subtyping). In this case, the function getmd will return method $\triangleleft g() = 42 \triangleright$ because it is the newest method out of the two.

Note that functions can be mutually recursive because of the dynamic nature of function call resolution.

Error Evaluation. These rules capture all possible error states of JULIETTE. Rule E-VarErr covers the case of a free variable, an `UndefVarError` in Julia. E-PrimopErr accounts for errors in primitive operations such as `DivideError`. E-CallErr fires when a non-function value is called. Finally, E-CallErr accounts for multiple-dispatch resolution errors, e.g. when the set of applicable methods is empty (no method found), and when there is no best method (ambiguous method).

4.3 Example

Fig. 14 shows a translation of the program from Fig. 3 to JULIETTE. First note that, as part of the translation, we wrap the entire program in $\langle \langle \rangle \rangle$, indicating that the outermost scope is the top level. Translation of method calls and definitions then proceeds, using $\triangleleft m(\bar{x}) = e \triangleright$ as a shorthand for $\triangleleft m(\bar{x} :: \top) = e \triangleright$ where $\top$ is the top type. Method bodies are converted by replacing `eval` invocations with their expressions wrapped in $\langle \langle \rangle \rangle$. The $\langle \langle \rangle \rangle$ context of e in Juliette effectively acts the same way that `eval` of e does in Julia, but evaluates variables in e using local, rather than global, scope.

```
g() = 2
f(x) = (eval(:(g() = $x))); x * g()
f(42)
```

```
 $\langle \langle \triangleleft g() = 2 \triangleright ;$   

 $\triangleleft f(x) = (\langle \triangleleft g() = x \triangleright ; x * g() \rangle) \triangleright ;$   

 $f(42) \rangle \rangle$ 
```

Fig. 14. From Julia (left) to JULIETTE (right)

Now we will show the execution of this translated program according to our small-step semantics. The initial state is $\langle \emptyset, p \rangle$ where p is the program on the right of Fig. 14 (and the $*$ operator is a primop). The first several steps of evaluation use rules E-MD and E-SEQ to add the definitions of g and f to the global table. This produces the state

$$\langle M_0, \langle \langle f(42) \rangle \rangle \rangle,$$

where

$$M_0 = (\emptyset \bullet \triangleleft g() = 2 \triangleright) \bullet \triangleleft f(x) = (\langle \triangleleft g() = x \triangleright ; x * g() \rangle) \triangleright$$

Next, using the E-CALLGLOBAL rule, the top-level call $f(42)$ steps to $\langle \langle f(42) \rangle \rangle_{M_0}$. This then produces the state

$$\langle M_0, \langle \langle \langle f(42) \rangle \rangle_{M_0} \rangle \rangle,$$

copying the global table into the context $\langle \cdot \rangle_{M_0}$. Now, rule E-CALLLOCAL can be used to resolve the call $f(42)$ in the table M_0 . Method $\triangleleft f(x) = \dots \triangleright$ is the only method of f and it is applicable to the integer argument ($\text{typeof}(42) = \text{Int} <: \top$), so the program steps to:

$$\langle M_0, \langle \langle \triangleleft g() = 42 \triangleright \rangle; 42 * g() \rangle_{M_0} \rangle.$$

The next expression to evaluate is the new g definition, $\triangleleft g() = 42 \triangleright$. Rule E-MD fires and the program steps to

$$\langle M_1, \langle \langle \triangleleft g \rangle \rangle; 42 * g() \rangle_{M_0} \rangle,$$

where

$$\begin{aligned} M_1 &= M_0 &= & ((\emptyset \bullet \triangleleft g() = 2 \triangleright) \\ &\bullet \triangleleft g() = 42 \triangleright &\bullet \triangleleft f(x) = (\langle \triangleleft g() = x \triangleright \rangle; x * g()) \triangleright) \\ &&\bullet \triangleleft g() = 42 \triangleright. \end{aligned}$$

The next two steps are:

$$\langle M_1, \langle \langle \triangleleft g \rangle \rangle; 42 * g() \rangle_{M_0} \rangle \xrightarrow{\text{E-VALGLOBAL}} \langle M_1, \langle \triangleleft g; 42 * g() \rangle_{M_0} \rangle \quad (1)$$

$$\xrightarrow{\text{E-SEQ}} \langle M_1, \langle \triangleleft 42 * g() \rangle_{M_0} \rangle. \quad (2)$$

Note that the last program state is represented by $\langle M_1, C [\triangleleft X [g()] \rangle_{M_0}] \rangle$, where $C = \langle \square \rangle$ and $X = 42 * \square$. So we have to use E-CALLLOCAL again to resolve $g()$ in the M_0 that is fixed in the context. Table M_0 has only one definition of g , the one that returns 2, so the program steps to:

$$\langle M_1, \langle \triangleleft 42 * 2 \rangle_{M_0} \rangle.$$

Finally, the application of E-PRIMOP, E-VALLOCAL, and E-VALGLOBAL leads to the final state:

$$\langle M_1, 84 \rangle. \quad (3)$$

Now, consider a modification of the original program where in the definition of f , the call $g()$ is wrapped into a global evaluation $\langle g() \rangle$:

```
g() = 2
f(x) = (eval{:(g() = $x)}; x * eval{:(g())})
f(42)
```

```
 $\triangleleft g() = 2 \triangleright$ ;
 $\triangleleft f(x) = (\langle \triangleleft g() = x \triangleright \rangle; x * \langle g() \rangle) \triangleright$ ;
 $f(42)$ 
```

At the beginning, the modified program will run similarly to the original one, and with step (2), it will reach the state:

$$\langle M_1, \langle \triangleleft 42 * \langle g() \rangle \rangle_{M_0} \rangle.$$

Here, $\langle \triangleleft 42 * \langle g() \rangle \rangle_{M_0}$ is represented by $C [\triangleleft X [g()] \rangle]$, where $C = \langle \triangleleft 42 * \square \rangle_{M_0}$ and $X = \square$. Therefore, the call $g()$ is back at the top level. With E-CALLGLOBAL rule, the call steps to $\langle g() \rangle_{M_1}$ because M_1 is the *current global* table, thus producing the state:

$$\langle M_1, \langle \triangleleft 42 * \langle \langle g() \rangle_{M_1} \rangle \rangle_{M_0} \rangle.$$

Resolved in M_1 , call $g()$ returns 42, and thus the whole program ends in the final state:

$$\langle M_1, 1764 \rangle.$$

Note that the resulting global table is the same as in (3), but the return value is different.

Fig. 15. Redex bases

4.4 Properties

JULIETTE operational semantics is deterministic, and all failure states are captured by error evaluation (meaning that a JULIETTE program never gets stuck).

LEMMA 4.1 (UNIQUE FORM OF EXPRESSION). *Any expression e can be uniquely represented in one of the following ways:*

- (a) $e = v$; or
- (b) $e = X[m(\bar{v})]$; or
- (c) $e = C[\text{rdx}]$,

where rdx (shown in Fig. 15) is a subset of expressions driving the reduction.

PROOF. By induction on e . □

THEOREM 4.2 (PROGRESS). *For any program p and method table M_g , the program either reduces to a value, or it makes a step to another program, or it errs. That is, one of the following holds:*

- (a) $\langle M_g, p \rangle \rightarrow \langle M'_g, v \rangle$; or
- (b) $\langle M_g, p \rangle \rightarrow \langle M'_g, p' \rangle$; or
- (c) $M_g \vdash p \rightarrow \text{error}$.

PROOF. By case analysis on $p = \langle e \rangle$, using Lemma 4.1. □

THEOREM 4.3 (DETERMINISM). *JULIETTE semantics is deterministic.*

PROOF. The proof relies on the fact that (1) any expression that steps can be represented as $C[\text{rdx}]$, and (2) such a representation is unique by Lemma 4.1. By case analysis on rdx , we can see that for all redex bases except $\delta_l(\bar{v})$ and $\langle X[m(\bar{v})] \rangle_M$, there is exactly one (normal- or error-evaluation) rule applicable. For $\delta_l(\bar{v})$ and $\langle X[m(\bar{v})] \rangle_M$, there are two rules for each, but their premises are incompatible. Thus, for any expression $C[\text{rdx}]$, exactly one rule is applicable. □

4.5 Optimizations

The world-age semantics allows function-call optimization even in the presence of `eval`. Recall how an `eval`d or top-level function call $\langle m(\bar{v}) \rangle$ steps. First, rule `E-CALLGLOBAL` is applied: it fixes the current state of the global table M in the call's context, stepping the call to $\langle m(\bar{v}) \rangle_M$. Then, the call $m(\bar{v})$ itself, and all of its nested calls (unless they are additionally wrapped into $\langle \rangle$), are resolved using the now-local table M . Therefore, M provides all necessary information for the resolution of such calls, and they can be optimized based on the method table M .

Next, we will focus on three generic-call optimizations: inlining, specialization, and transforming generic calls into direct calls (devirtualization). Namely, we provide formal definitions of these optimizations and show them correct.

<pre> g(x::Any) = x + x # g1 g(x::Bool) = x # g2 f(x::Int) = x * g(x) f(5) </pre>	<pre> g(x::Any) = x + x # g1 g(x::Bool) = x # g2 f(x::Int) = x * g((println(x);x)) f(5) </pre>
--	---

Fig. 16. Candidate programs for inlining (on the left) and direct call optimization (on the right)

Inlining. If a function call is known to dispatch to a certain method using a *fixed* method table, it might be possible to *inline* the body of the method in place of the call. For example, consider a program on the left of Fig. 16. The call `f(5)` has no choice but to dispatch to the only definition of `f`. Because the call `g(x)` in `f(5)` is not wrapped in an `eval`, it is known that the call to `g` is going to be dispatched in the context with exactly two methods of `g`: `g1` and `g2`. Furthermore, since `x` is known to be of type (tag) `Int` inside `f`, we know that `g(x)` has to dispatch to the method `g1` (because `Int <: Any` but `Int </: Bool`). Thus, it is possible to optimize method `f` for the call `f(5)` by inlining `g(x)`, which yields the following optimized definition of `f`:

```
f(x::Int) = x * (x + x)
```

Direct-call optimization. When inlining is not possible or desirable, but it is clear which method is going to be invoked, a function call can be replaced by a direct invocation. Consider the example on the right of Fig. 16. The only difference from the previous example is that the argument of `g` inside `f` is not a variable but an expression `(println(x);x)`. This expression always returns an integer, so we know that at run time, that `g` will be dispatch to method `g1`. However, unlike previously, the call to `g` cannot be inlined using direct syntactic substitution. In that case, the value of `x` would be printed twice instead of just once, because inlining would transform `g((println(x);x))` into `(println(x);x)+(println(x);x)` and thus change the observable behavior of the program. It is still possible to optimize `f`, by replacing the generic call to `g` with a *direct call* to the method `g1`. In pseudo-code, this can be written as:

```
f(x::Int) = x * g@gl((println(x); x))
```

In the calculus, we model a direct call as a call to a new function with a single method such that the name of the function is not used anywhere in the original method table or expression. For example, for the program above, we can add function `h` with only one method `h(x::Int)=x+x`, allowing `f` to be optimized to:

```
f(x::Int) = x * h(println(x); x)
```

Specialization. The final optimization we consider is specialization of methods for argument types. In Fig. 16, method `g1` is defined for `x` of type `Any`, meaning that the call `x+x` can be dispatched to any of at least 166 standard methods. But because, within `f`, `g` is known to be called with an argument of type `Int` (due to `x` in `f` having that type), it is possible to generate a new implementation of `g` *specialized* for this argument type. The advantage is that the specialized implementation can directly use efficient integer addition. Thus, combined with the direct call, we have:

```

g(x::Int) = Base.add_int(x, x) # g3
f(x::Int) = x * g@g3((println(x); x))

```

In the calculus, specialization is modeled similarly to direct calls: as a function with a fresh name.

$$\begin{array}{c}
\Gamma ::= \\
\quad | \quad \emptyset \\
\quad | \quad \Gamma, x : \tau
\end{array}
\quad \text{Typing environment}$$

$$\begin{array}{c}
\text{T-VAR} \\
\frac{\Gamma(x) = \sigma}{\Gamma \vdash x : \sigma}
\end{array}
\quad
\begin{array}{c}
\text{T-VAL} \\
\frac{\text{typeof}(v) = \sigma}{\Gamma \vdash v : \sigma}
\end{array}
\quad
\begin{array}{c}
\text{T-MD} \\
\frac{}{\Gamma \vdash \llbracket m(\bar{x} :: \bar{\tau}) = e \rrbracket : \mathbb{f}_m}
\end{array}
\quad
\begin{array}{c}
\text{T-SEQ} \\
\frac{\Gamma \vdash e_2 : \sigma}{\Gamma \vdash (e_1 ; e_2) : \sigma}
\end{array}$$

$$\begin{array}{c}
\text{T-PRIMOP} \\
\frac{\Psi(l, \bar{\sigma}) = \sigma' \quad \Gamma \vdash e_i : \sigma_i}{\Gamma \vdash \delta_l(\bar{e}) : \sigma'}
\end{array}
\quad
\begin{array}{c}
\text{T-EVALGLOBAL} \\
\frac{\Gamma \vdash e : \sigma}{\Gamma \vdash \llbracket e \rrbracket : \sigma}
\end{array}
\quad
\begin{array}{c}
\text{T-EVALLOCAL} \\
\frac{\Gamma \vdash e : \sigma}{\Gamma \vdash \llbracket e \rrbracket_M : \sigma}
\end{array}$$

Fig. 17. Concrete-typing judgment

4.6 Optimization Correctness

In this section, we present a formal definition of optimizations and state the main theorem about their correctness. The general idea of optimizations is as follows: if an expression e is going to be executed in a fixed method table M , it is safe to instead execute e in a table M' obtained by optimizing method definitions of M (like we did with the definition of f in the examples above).

As demonstrated by the examples, the first ingredient of optimizations is type information, which is necessary to “statically” resolve function calls; for this, we use a simple concrete-typing relation defined in Fig. 17. The relation $\Gamma \vdash e : \sigma$ propagates information about variables and type tags of values, and succeeds only if the expression would always reduce to a value of concrete type σ if it reduces to any value. This is because to resolve a function call, we need to know the type tags of its arguments. A typing relation can be more complex to enable further optimization opportunities (and it is much more complex in Julia), but typing of Julia is a separate topic that is out of scope of this paper: here, we focus on compiler optimization and use concrete typing only as a tool.

Fig. 18 shows the judgments related to method-table optimization. The rule OT-METHODTABLE says that an optimized version M' of table M (1) has to have all the methods of M , although they can be optimized, and (2) can have more methods given that their names do not appear in the original table M . The latter enables adding new methods that model direct calls and specializations, and the former allows for optimization of existing methods. According to the rule OD-MD, a method in M' optimizes a method in M if it has the same signature (i.e. name and argument types), and its body is an optimization of the original body of the method being optimized.

The method-optimization environment Φ tracks direct calls and specializations: $m(\bar{\sigma}) \rightsquigarrow m'$ tells that when arguments of m have type tags $\bar{\sigma}$, a call to m in M can be replaced by a call to m' in M' . Note that all entries of Φ need to be valid according to METHODOPT-VALID: assuming that the methods are in the optimization relation, their bodies indeed have to be in that relation (the assumption is needed to handle recursion). Both OD-MD and METHODOPT-VALID rely on expression optimization to relate method bodies.

Finally, the expression-optimization relation is shown in Fig. 19. Note that the rules do not allow for function-call optimizations inside the global-evaluation construct $\llbracket \cdot \rrbracket$: the only applicable rule in that case is OE-GLOBAL. Function calls can only be optimized if they are fixed-table calls. Rules

$$\begin{array}{c}
\Phi ::= \\
| \quad \emptyset \\
| \quad \Phi, m(\bar{\sigma}) \rightsquigarrow m'
\end{array}
\quad \text{Method-optimization environment}$$

$$\begin{array}{c}
\text{METHODOPT-VALID} \\
\text{getmd}(M, m, \bar{\sigma}) = \triangleleft m(\bar{x} :: \bar{\tau}) = e_b \triangleright \\
\text{getmd}(M', m', \bar{\sigma}) = \triangleleft m'(\bar{x}' :: \bar{\tau}') = e'_b \triangleright \\
\frac{\bar{x} : \bar{\sigma} \vdash_{M \rightsquigarrow M'}^{\Phi} e_b \rightsquigarrow e'_b[\bar{x}' \mapsto \bar{x}]}{\vdash_{M \rightsquigarrow M'}^{\Phi} m(\bar{\sigma}) \rightsquigarrow m'}
\end{array}
\quad
\begin{array}{c}
\text{OD-MD} \\
\frac{\bar{x} : \bar{\tau} \vdash_{M \rightsquigarrow M'}^{\Phi} e \rightsquigarrow e'[x' \mapsto x]}{\vdash_{M \rightsquigarrow M'}^{\Phi} \triangleleft m(\bar{x} :: \bar{\tau}) = e \triangleright \rightsquigarrow \triangleleft m(\bar{x}' :: \bar{\tau}') = e' \triangleright}
\end{array}$$

$$\begin{array}{c}
\text{OT-METHODTABLE} \\
M = \text{md}_1 \bullet \dots \bullet \text{md}_n \quad M' = \text{md}'_1 \bullet \dots \bullet \text{md}'_n \bullet \text{md}'_{n+1} \bullet \dots \bullet \text{md}'_k \quad \forall 1 \leq i \leq n. \vdash_{M \rightsquigarrow M'}^{\Phi} \text{md}_i \rightsquigarrow \text{md}'_i \\
\frac{\forall (m(\bar{\sigma}) \rightsquigarrow m') \in \Phi. \vdash_{M \rightsquigarrow M'}^{\Phi} m(\bar{\sigma}) \rightsquigarrow m' \quad \forall n+1 \leq j \leq k. \text{name}(\text{md}'_j) \text{ does not occur in } M}{\vdash^{\Phi} M \rightsquigarrow M'}
\end{array}$$

$$\begin{array}{c}
\text{MNamesCompat} \\
\frac{\forall m \text{ referenced by } e. m \in \text{dom}(M) \iff m \in \text{dom}(M')}{\vdash_{M \rightsquigarrow M'}^{\Phi} e}
\end{array}
\quad
\begin{array}{c}
\text{OT-METHODTABLE-EXPR} \\
\frac{\vdash^{\Phi} M \rightsquigarrow M' \quad \vdash_{M \rightsquigarrow M'}^{\Phi} e}{\vdash_e^{\Phi} M \rightsquigarrow M'}
\end{array}$$

Fig. 18. Method table & definition optimization

$$\begin{array}{c}
\nu ::= \quad \nu \mid x \quad \text{Near-value}
\end{array}$$

$$\begin{array}{c}
\text{OE-VAL} \\
\frac{\nu \neq m}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} \nu \rightsquigarrow \nu}
\end{array}
\quad
\begin{array}{c}
\text{OE-VALFUN} \\
\frac{\vdash_{M \rightsquigarrow M'}^{\Phi} m}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} m \rightsquigarrow m}
\end{array}
\quad
\begin{array}{c}
\text{OE-VAR} \\
\frac{}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} x \rightsquigarrow x}
\end{array}$$

$$\begin{array}{c}
\text{OE-GLOBAL} \\
\frac{\vdash_{M \rightsquigarrow M'}^{\Phi} e}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} (\langle e \rangle) \rightsquigarrow (\langle e \rangle)}
\end{array}
\quad
\begin{array}{c}
\text{OE-LOCAL} \\
\frac{\vdash_{M \rightsquigarrow M'}^{\Phi} e}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} (\langle e \rangle)_{M_I} \rightsquigarrow (\langle e \rangle)_{M_I}}
\end{array}
\quad
\begin{array}{c}
\text{OE-MD} \\
\frac{\vdash_{M \rightsquigarrow M'}^{\Phi} \text{name}(\text{md})}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} \text{md} \rightsquigarrow \text{md}}
\end{array}$$

$$\begin{array}{c}
\text{OE-SEQ} \\
\frac{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} e_1 \rightsquigarrow e'_1 \quad \Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} e_2 \rightsquigarrow e'_2}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} e_1 ; e_2 \rightsquigarrow e'_1 ; e'_2}
\end{array}
\quad
\begin{array}{c}
\text{OE-PRIMOP} \\
\frac{\forall i. \Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} e_i \rightsquigarrow e'_i}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} \delta_I(\bar{e}) \rightsquigarrow \delta_I(\bar{e}')}
\end{array}$$

$$\begin{array}{c}
\text{OE-CALL} \\
\frac{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} e_c \rightsquigarrow e'_c \quad \forall i. \Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} e_i \rightsquigarrow e'_i}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} e_c(\bar{e}) \rightsquigarrow e'_c(\bar{e}')}
\end{array}$$

$$\begin{array}{c}
\text{OE-INLINE} \\
\frac{\forall i. \Gamma \vdash \nu_i : \sigma_i \quad \text{getmd}(M, m, \bar{\sigma}) = \triangleleft m(\bar{x} :: \bar{\tau}) = e_b \triangleright \quad \Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} e_b[\bar{x} \mapsto \bar{\nu}] \rightsquigarrow e'}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} m(\bar{\nu}) \rightsquigarrow \text{unit}; e'}
\end{array}$$

$$\begin{array}{c}
\text{OE-DIRECT} \\
\frac{\forall i. \Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} e_i \rightsquigarrow e'_i \quad \Gamma \vdash e'_i : \sigma_i \quad (m(\bar{\sigma}) \rightsquigarrow m') \in \Phi}{\Gamma \vdash_{M \rightsquigarrow M'}^{\Phi} m(\bar{e}) \rightsquigarrow m'(\bar{e}')}
\end{array}$$

Fig. 19. Expression optimization

OE-INLINE and OE-DIRECT correspond to the inlining and the direct-call/specialization optimizations, respectively. As discussed earlier, inlining cannot be done if a function is called with expression arguments. Therefore, in OE-INLINE we use an auxiliary definition ν , “near-value”, which is either a

$$\begin{array}{c}
\gamma ::= \overline{x \mapsto v} \text{ where } \forall i, j. x_i \neq x_j \quad \text{Value substitution} \\
\frac{\text{dom}(\Gamma) = \text{dom}(\gamma) \quad \forall x \in \text{dom}(\gamma). (\Gamma(x) = \sigma \iff \text{typeof}(\gamma(x)) = \sigma)}{\Gamma \vdash \gamma} \quad \gamma\text{-Ok}
\end{array}$$

Fig. 20. Value substitution

value or a variable. Because we model direct-call and specialization optimizations as calls to freshly-named methods, the main job is done in the table-optimization rule OT-METHODTABLE; rule OE-DIRECT only records the fact of invoking a specific method. If the method definition of m' from $m(\bar{\sigma}) \rightsquigarrow m'$ has the same parameter-type annotations as m , it represents a direct call to an original method of m ; otherwise, it represents a specialized method. Note that for all optimizations, function-call arguments have to be concretely typed. Otherwise, we do not know definitively how a function call is going to be dispatched at run time.

The optimizations defined in Fig. 18–19 are sound. That is, the evaluation of the original and optimized programs yield the same result. To show this, we establish a bisimulation relation between original and optimized expressions (after the following auxiliary lemmas):

LEMMA 4.4 (CONTEXT IRRELEVANCE). *For all $C, C', \text{rdx}, e', M_g, M'_g$, the following holds:*

$$\langle M_g, C[\text{rdx}] \rangle \rightarrow \langle M'_g, C[e'] \rangle \iff \langle M_g, C'[\text{rdx}] \rangle \rightarrow \langle M'_g, C'[e'] \rangle.$$

PROOF. By analyzing normal-evaluation steps, we can see that only rdx matters for the reduction. Formally, the proof goes by inspecting a reduction step for $C[\text{rdx}]$ ($C'[\text{rdx}]$) and building a corresponding step for $C'[\text{rdx}]$ ($C[\text{rdx}]$). $\square$

LEMMA 4.5 (SIMPLE-CONTEXT IRRELEVANCE). *For all $M, C, e, M_g, e', M'_g, X$, the following holds:*

$$\langle M_g, C[\langle e \rangle_M] \rangle \rightarrow \langle M'_g, C[\langle e' \rangle_M] \rangle \implies \langle M_g, C[\langle X[e] \rangle_M] \rangle \rightarrow \langle M'_g, C[\langle X[e'] \rangle_M] \rangle.$$

PROOF. By Lemma 4.1, e is either v or $X_e[m(\bar{v})]$ or $C_e[\text{rdx}]$. If e is v , the assumption of the lemma does not hold ($C[\langle v \rangle_M]$ would step to $C[v]$), so only $X_e[m(\bar{v})]$ and $C_e[\text{rdx}]$ cases are possible.

- When e is $X_e[m(\bar{v})]$, $\langle e \rangle_M = \langle X_e[m(\bar{v})] \rangle_M$ is a redex, and $C[\langle e \rangle_M]$ steps by rule E-CALLLOCAL. But $\langle X[e] \rangle_M = \langle X[X_e[m(\bar{v})]] \rangle_M$ is also a redex, and $C[\langle X[e] \rangle_M]$ steps by rule E-CALLLOCAL similarly to $C[\langle e \rangle_M]$.
- When e is $C_e[\text{rdx}]$, $\langle e \rangle_M = \langle C_e[\text{rdx}] \rangle_M$ and $C[\langle e \rangle_M] = C'[\text{rdx}]$ where $C' = C[\langle C_e \rangle_M]$. Since $C[\langle X[e] \rangle_M] = C''[\text{rdx}]$ for $C'' = C[\langle X[C_e] \rangle_M]$, $C[\langle e \rangle_M]$ and $C[\langle X[e] \rangle_M]$ step similarly by Lemma 4.4. $\square$

LEMMA 4.6 (OPTIMIZATION PRESERVES VALUES). *For all $\Phi, M, M', \Gamma, e, v$, the following hold:*

$$\Gamma \vdash_{M \rightsquigarrow M'}^\Phi v \rightsquigarrow e \implies e = v \quad \text{and} \quad \Gamma \vdash_{M \rightsquigarrow M'}^\Phi e \rightsquigarrow v \implies e = v.$$

PROOF. By case analysis on the optimization relation. $\square$

LEMMA 4.7 (VALUE SUBSTITUTION PRESERVES OPTIMIZATION). *For all $\Phi, \Gamma, e, M, e', M', \gamma$, such that $\forall v \in \gamma. \vdash_{M \rightsquigarrow M'}^\Phi v$, the following holds:*

$$(\Gamma \vdash_{M \rightsquigarrow M'}^\Phi e \rightsquigarrow e' \wedge \Gamma \vdash \gamma) \implies \vdash_{M \rightsquigarrow M'}^\Phi \gamma(e) \rightsquigarrow \gamma(e').$$

PROOF. By induction on the derivation of $\Gamma \vdash_{M \rightsquigarrow M'}^\Phi e \rightsquigarrow e'$. $\square$

LEMMA 4.8 (BISIMULATION). *For all method tables M, M' , method-optimization environment Φ , and expressions e_1, e'_1 , such that*

$$\vdash^\Phi M \rightsquigarrow M' \quad \text{and} \quad \vdash_{M \rightsquigarrow M'}^\Phi e_1 \rightsquigarrow e'_1,$$

for all global tables M_g, M'_g and world context C , the following hold:

(1) *Forward direction:*

$$\begin{aligned} \forall e_2. \quad & \langle M_g, C [\langle e_1 \rangle_M] \rangle \rightarrow \langle M'_g, C [\langle e_2 \rangle_M] \rangle \\ \implies & \\ \exists e'_2. \quad & \langle M_g, C [\langle e'_1 \rangle_{M'}] \rangle \rightarrow \langle M'_g, C [\langle e'_2 \rangle_{M'}] \rangle \quad \wedge \quad \vdash_{M \rightsquigarrow M'}^\Phi e_2 \rightsquigarrow e'_2. \end{aligned}$$

(2) *Backward direction:*

$$\begin{aligned} \forall e'_2. \quad & \langle M_g, C [\langle e'_1 \rangle_{M'}] \rangle \rightarrow \langle M'_g, C [\langle e'_2 \rangle_{M'}] \rangle \\ \implies & \\ \exists e_2. \quad & \langle M_g, C [\langle e_1 \rangle_M] \rangle \rightarrow \langle M'_g, C [\langle e_2 \rangle_M] \rangle \quad \wedge \quad \vdash_{M \rightsquigarrow M'}^\Phi e_2 \rightsquigarrow e'_2. \end{aligned}$$

PROOF. The proof goes by induction on the derivation of optimization $\vdash_{M \rightsquigarrow M'}^\Phi e_1 \rightsquigarrow e'_1$. For each case, both directions are proved by analyzing possible normal-evaluation steps. More specifically, the forward-direction proof strategy is as follows (the backward direction is similar):

- (1) Observe that to make the required step, e_1 should have a certain representation. Consider all possible representations that satisfy this requirement.
- (2) For each representation, analyze the suitable normal-evaluation rule (recall that the semantics is deterministic, so there will be just one such rule).
- (3) If e_1 represents an immediate redex (e.g. $v_{11} ; e_{12}$), the optimized expression will be an immediate redex too (possibly, of a different form). Otherwise, use induction hypothesis and auxiliary facts about contexts and evaluation to show that the optimized expression steps in a similar fashion, in particular, facts from Lemma 4.4 and Lemma 4.5.
- (4) Finally, show that the resulting expressions are in the optimization relation. This will follow from the assumptions and induction.

As an example, consider the proof of the forward direction for the sequence case OE-SEQ. By assumption, we have $e_1 = e_{11} ; e_{12}$ and $e'_1 = e'_{11} ; e'_{12}$ where

$$\frac{\vdash_{M \rightsquigarrow M'}^\Phi e_{11} \rightsquigarrow e'_{11} \quad \Gamma \vdash_{M \rightsquigarrow M'}^\Phi e_{12} \rightsquigarrow e'_{12}}{\vdash_{M \rightsquigarrow M'}^\Phi e_{11} ; e_{12} \rightsquigarrow e'_{11} ; e'_{12}} \text{OE-SEQ}.$$

For $C [\langle e_{11} ; e_{12} \rangle_M]$ to reduce, by case analysis, we know there are three possibilities.

- (1) $e_{11} = v_{11}$ and $\langle M_g, C [\langle v_{11} ; e_{12} \rangle_M] \rangle \rightarrow \langle M_g, C [\langle e_{12} \rangle_M] \rangle$ by rule E-SEQ. Then by Lemma 4.6, $e'_{11} = v_{11}$ and the optimized expression steps by the same rule:

$$\langle M_g, C [\langle v_{11} ; e'_{12} \rangle_{M'}] \rangle \rightarrow \langle M_g, C [\langle e'_{12} \rangle_{M'}] \rangle.$$

The desired optimization relation holds by one of the assumptions: $\Gamma \vdash_{M \rightsquigarrow M'}^\Phi e_{12} \rightsquigarrow e'_{12}$.

- (2) $e_{11} = X_1 [m(\bar{v})]$ and the original expression steps by E-CALLLOCAL:

$$\langle M_g, C [\langle X_1 [m(\bar{v})] ; e_{12} \rangle_M] \rangle \rightarrow \langle M_g, C [\langle X_1 [e_b[\bar{x} \mapsto \bar{v}]] ; e_{12} \rangle_M] \rangle.$$

Since $C [\langle X_1 [m(\bar{v})] \rangle_M]$ reduces similarly, by the induction hypothesis, $\exists e'_{21}$ such that

$$\langle M_g, C [\langle e'_{11} \rangle_{M'}] \rangle \rightarrow \langle M_g, C [\langle e'_{21} \rangle_{M'}] \rangle \quad \text{and} \quad \Gamma \vdash_{M \rightsquigarrow M'}^\Phi X_1 [e_b[\bar{x} \mapsto \bar{v}]] \rightsquigarrow e'_{21}.$$

But then, by Lemma 4.5, the entire optimized expression $\mathcal{C} [\llbracket e'_{11} ; e'_{12} \rrbracket_M]$ steps too, and the desired optimization relation holds:

$$\frac{\vdash_{M \rightsquigarrow M'}^\Phi \lambda_1 [e_b[\bar{x} \mapsto \bar{v}]] \rightsquigarrow e'_{21} \quad \vdash_{M \rightsquigarrow M'}^\Phi e_{12} \rightsquigarrow e'_{12}}{\vdash_{M \rightsquigarrow M'}^\Phi \lambda_1 [e_b[\bar{x} \mapsto \bar{v}]] ; e_{12} \rightsquigarrow e'_{21} ; e'_{12}} \text{OE-SEQ}.$$

(3) $e_{11} = \mathcal{C}_1 [\text{rdx}]$ and by Lemma 4.4:

$$\begin{aligned} \langle M_g, \mathcal{C} [\llbracket \mathcal{C}_1 [\text{rdx}] \rrbracket_M] \rangle &\rightarrow \langle M'_g, \mathcal{C} [\llbracket \mathcal{C}_1 [e'] \rrbracket_M] \rangle \\ &\iff \\ \langle M_g, \mathcal{C} [\llbracket \mathcal{C}_1 [\text{rdx}] \rrbracket_M] \rangle &\rightarrow \langle M'_g, \mathcal{C} [\llbracket \mathcal{C}_1 [e'] \rrbracket_M] \rangle. \end{aligned}$$

Since $\mathcal{C} [\llbracket \mathcal{C}_1 [\text{rdx}] \rrbracket_M]$ reduces, by the induction hypothesis, $\exists e'_{21}$ such that

$$\langle M_g, \mathcal{C} [\llbracket \mathcal{C}_1 [e'_{11}] \rrbracket_{M'}] \rangle \rightarrow \langle M_g, \mathcal{C} [\llbracket \mathcal{C}_1 [e'_{21}] \rrbracket_{M'}] \rangle \quad \text{and} \quad \Gamma \vdash_{M \rightsquigarrow M'}^\Phi \mathcal{C}_1 [e'] \rightsquigarrow e'_{21}.$$

Similarly to the previous case, the entire $\mathcal{C} [\llbracket e'_{11} ; e'_{12} \rrbracket_M]$ steps, and the desired optimization relation holds. $\square$

LEMMA 4.9 (REFLEXIVITY OF OPTIMIZATION). *For all M, M', Φ, Γ, e , the following holds:*

$$\vdash_e^\Phi M \rightsquigarrow M' \implies \Gamma \vdash_{M \rightsquigarrow M'}^\Phi e \rightsquigarrow e.$$

PROOF. By induction on e . The only interesting cases are m , md , $\llbracket e' \rrbracket$, and $\llbracket e' \rrbracket_{M'}$. For example, consider the case of m (others are similar). Rule OE-VALFUN requires a method named m to either exist in both tables or do not appear in either (this rules out the case where $\llbracket m() \rrbracket_M$ would err but $\llbracket m() \rrbracket_{M'}$ succeed). This requirement is guaranteed by the assumption that $\vdash_e^\Phi M \rightsquigarrow M'$, which by inversion, gives the necessary $\vdash_{M \rightsquigarrow M'}^\Phi e$. $\square$

The main result, Theorem 4.10, is a corollary of Lemma 4.8. It states that a fixed-table expression can be soundly evaluated in an optimized table.

THEOREM 4.10 (CORRECTNESS OF TABLE OPTIMIZATION). *For all M, M', Φ, e satisfying $\vdash_e^\Phi M \rightsquigarrow M'$, for all $M_g, M'_g, \mathcal{C}, v$, the following holds:*

$$\langle M_g, \mathcal{C} [\llbracket e \rrbracket_M] \rangle \rightarrow^* \langle M'_g, v \rangle \iff \langle M_g, \mathcal{C} [\llbracket e \rrbracket_{M'}] \rangle \rightarrow^* \langle M'_g, v \rangle.$$

PROOF. First of all, note that $\vdash_{M \rightsquigarrow M'}^\Phi e \rightsquigarrow e$ by Lemma 4.9, and that $\vdash^\Phi M \rightsquigarrow M'$ follows from $\vdash_e^\Phi M \rightsquigarrow M'$. Then, we proceed by induction on $\rightarrow^*$ (reflexive-transitive closure of normal evaluation). In the interesting case of the forward direction, we have:

$$\frac{\langle M_g, \mathcal{C} [\llbracket e \rrbracket_M] \rangle \rightarrow \langle M'_g, \mathcal{C} [\llbracket e'_1 \rrbracket_M] \rangle \quad \langle M'_g, \mathcal{C} [\llbracket e'_1 \rrbracket_M] \rangle \rightarrow^* \langle M'_g, v \rangle}{\langle M_g, \mathcal{C} [\llbracket e \rrbracket_M] \rangle \rightarrow^* \langle M'_g, v \rangle}.$$

By applying Lemma 4.8 to the first premise, we get:

$$\langle M_g, \mathcal{C} [\llbracket e \rrbracket_{M'}] \rangle \rightarrow \langle M'_g, \mathcal{C} [\llbracket e'_2 \rrbracket_{M'}] \rangle \quad \text{and} \quad \vdash_{M \rightsquigarrow M'}^\Phi e'_1 \rightsquigarrow e'_2.$$

By applying the induction hypothesis to the second premise, we get:

$$\langle M'_g, \mathcal{C} [\llbracket e'_2 \rrbracket_{M'}] \rangle \rightarrow^* \langle M'_g, v \rangle.$$

By combining the results, we can get the desired derivation:

$$\frac{\langle M_g, \mathcal{C} [\llbracket e \rrbracket_{M'}] \rangle \rightarrow \langle M'_g, \mathcal{C} [\llbracket e'_2 \rrbracket_{M'}] \rangle \quad \langle M'_g, \mathcal{C} [\llbracket e'_2 \rrbracket_{M'}] \rangle \rightarrow^* \langle M'_g, v \rangle}{\langle M_g, \mathcal{C} [\llbracket e \rrbracket_{M'}] \rangle \rightarrow^* \langle M'_g, v \rangle}.$$

The backward direction proceeds similarly. $\square$

Theorem 4.10, in particular, justifies Julia’s choice to execute top-level calls using optimized methods. Once a top-level call $\langle M, \langle m(\bar{v}) \rangle \rangle$ steps to a fixed-table call $\langle M, \langle m(\bar{v}) \rangle_M \rangle$, it is sound to optimize table M into M' (using inlining, direct calls, and specialization), and evaluate the call in the optimized table $\langle M, \langle m(\bar{v}) \rangle_{M'} \rangle$.

4.7 Testing the Semantics

To check if JULIETTE behaves as we expect, we implemented it in Redex [Felleisen et al. 2009] and ran it along with Julia on a small set of 9 litmus tests (provided in App. A); Julia agrees with JULIETTE on all of them. The tests cover the intersection of the semantics of JULIETTE and Julia, and demonstrate the interaction of `eval`, method definitions, and method calls. In particular, the litmus tests ensure: that the executing semantics prohibits calls to too-new methods, that this restriction can be skipped with `eval` or `invokelatest`, and that the semantics of `eval` executes successive statements in the latest age.

Two of the litmus tests are shown in Fig. 21; each test is made up of a small program and its expected output. The tests examine the case where a method `r2` is placed “in between” the generated method `r1` and an older `m`. In the first test, `m` errs. While `r2` is callable from the age that `m` was called in, `r1` is not. In the second test, we use `invokelatest` to execute `r2` in the latest world age; this allows the invocation of the dynamically generated `r1`.

```
r2() = r1()
m() = (
  eval(:(r1() = 2));
  r2()
)
m() # error
```

```
r2() = r1()
m() = (
  eval(:(r1() = 2));
  Base.invokelatest(r2)
)
m() == 2 # passes
```

Fig. 21. Litmus tests

To use the litmus tests, we need to (1) translate them from Julia into our grammar and (2) implement the semantics of JULIETTE into an executable form. The former is done by translating ASTs. The latter is realized with a Redex mechanization, which is publicly available on GitHub⁵ along with the litmus tests. The model implements the calculus almost literally. Values, tags, and type annotations are instantiated with several concrete examples, such as numbers and strings. Primitive operations include arithmetic and `print`. The only difference between the paper and Redex is handling of function names. Similar to Julia, in the Redex model, a definition of the method named `f` introduces a global constant `f`. When referenced, the constant evaluates to a function value `f`. Thus, instead of a single error evaluation rule $E\text{-VAR}$ from Fig. 13, the Redex model has the following two rules, one for normal evaluation and one for erroneous evaluation:

$$\begin{array}{c}
\text{E-VARMETHOD} \\
\frac{x \in \text{dom}(M)}{\langle M, \mathbf{C}[x] \rangle \rightarrow \langle M, \mathbf{C}[x] \rangle} \\
\\
\text{E-VARERR} \\
\frac{x \notin \text{dom}(M)}{M \vdash \mathbf{C}[x] \rightarrow \text{error.}}
\end{array}$$

The new rules treat the global method table as a global environment: $E\text{-VARMETHOD}$ evaluates a global variable to its underlying function value, and $E\text{-VARERR}$ errs if a variable is not found in the global environment; all local variables should be eliminated by substitution. All paper-style programs can be written in the Redex model, and the extension makes it easier to compare and translate Julia programs to corresponding Redex programs. Thus, the litmus test on the left of Fig. 21 translates to the Redex model as follows (the grammar is written in S-expressions style):

⁵<https://github.com/julbinb/juliette-wa>

```
(evalg (seq (seq
  (mdef "r2" () (mcall r1)) # r2() = r1()
  (mdef "m" () (seq
    (evalg (mdef "r1" () 2)) # eval:(r1() = 2))
    (mcall r2)))) # r2()
(mcall m))) # m()
```

The Redex model also implements the optimization judgments presented in Sec. 4.5, as well as a straightforward optimization algorithm that is checked against the judgments.

Discussion with Julia’s developers confirmed that our understanding of world age is correct, and that the table-based semantics has a correspondence to the age-based implementation. Namely, it is possible to generate JULIETTE method tables from the global data structure used by Julia to store methods.

5 CONCLUSION

Julia’s approach to dynamic code loading is distinct; instead of striving to achieve performance *in spite of* the language’s semantics, the designers of Julia chose to restrict expressiveness so that they could keep their compiler simple *and* generate fast code. World age aligns Julia’s dynamic semantics with its just-in-time compiler’s static approximation. As a result, statically resolved function calls have the same behavior as dynamic invocations.

This equivalence—that statically and dynamically resolved methods behave the same—allows Julia to forsake some of the complexity of modern compilers. Instead of needing deoptimization to handle newly added definitions, Julia simply does not allow running code to see those definitions. Thus, optimizations can rely on the results of static reasoning about the method table, while remaining sound in the presence of `eval`. If necessary, the programmer can explicitly ask for newly defined methods, making the performance penalty explicit and user-controllable.

World age need not be limited to Julia. Any language that supports updating existing function definitions may benefit from such a mechanism, namely control over when those new definitions can be observed and when function calls can be optimized. From Java to languages like R, having a clear semantics for updating code, especially in the presence of concurrency, can be beneficial, as it would improve our ability to reason about programs written in those languages.

Although the world-age semantics presented in the paper follows Julia, a world-age semantics does not have to. For instance, an alternative world-age semantics could pick another point when the age counter is incremented. The notion of top level makes sense in the context of an interactive development environment, but is unclear in, for example, a web server that may receive new code to install from time to time. Such a continuously running system may need a definition of quiescence that is different from the top-level used in Julia. One alternative is to provide an explicit `freeze` construct that allows programmers to opt-in to the world-age system. This would allow existing languages to incorporate world age without affecting existing code.

The calculus we present here is a basic foundation intended to capture the operation of world age. Future work may build on this to formalize the semantics of Julia as a whole, but, notably, the additional semantics will not impact the world-age mechanism itself. Of particular note is mutable state: it is orthogonal to world age because Julia decouples code state from data state by design. This was a pragmatic decision, as the compiler depends on knowing the contents of the method table for its optimization. Optimizations based on global variables are much less frequent.

Fig. 22. Litmus Tests

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful comments and suggestions to improve this paper. This work was supported by Office of Naval Research (ONR) award 503353, the National Science Foundation awards 1759736, 1925644 and 1618732, the Czech Ministry of Education from the Czech Operational Programme Research, Development, and Education, under grant agreement No. CZ.02.1.01/0.0/0.0/15_003/0000421, and the European Research Council under the European Union’s Horizon 2020 research and innovation programme, under grant agreement No. 695412.

A LITMUS TESTS

As a basic test of functionality, we provide 9 litmus tests shown in Fig. 22, written in Julia, that exercise the basic world age semantics as well as key Julia semantics surrounding world age. The tests suffice to identify the following semantic characteristics:

- (a) too-new methods cannot be called using a normal invocation;
- (b) `invokelatest` uses the latest world age;
- (c) `eval` uses the latest world age;
- (d) successive `eval` statements run in the latest world age;
- (e) only age at the top-level is relevant for invocation visibility;
- (f) “latest” calls propagate the new world age;
- (g) `eval` executes in the top-level scope;

- (h) normal invocation uses overridden methods if added method too new;
- (i) eval will use latest definition of an overridden method.

REFERENCES

- Julia Belyakova, Benjamin Chung, Jack Gelinas, Jameson Nash, Ross Tate, and Jan Vitek. 2020. World Age in Julia: Optimizing Method Dispatch in the Presence of Eval (Extended Version). arXiv:2010.07516
- 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). <https://doi.org/10.1145/3276490>
- Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. 2017. Julia: A Fresh Approach to Numerical Computing. *SIAM Rev.* 59, 1 (2017). <https://doi.org/10.1137/141000671>
- Daniel G. Bobrow, Kenneth Kahn, Gregor Kiczales, Larry Masinter, Mark Stefik, and Frank Zdybel. 1986. CommonLoops: Merging Lisp and Object-oriented Programming. In *Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*. <https://doi.org/10.1145/28697.28700>
- Robert P. Cook and Insup Lee. 1983. DYMO: A Dynamic Modification System. In *Proceedings of the Symposium on High-Level Debugging*. <https://doi.org/10.1145/1006147.1006188>
- David Detlefs and Ole Agesen. 1999. Inlining of Virtual Methods. In *European Conference on Object-Oriented Programming (ECOOP)*. <https://doi.org/10.5555/646156.679839>
- Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. 2009. *Semantics Engineering with PLT Redex*. MIT Press. <http://mitpress.mit.edu/catalog/item/default.asp?type=2&tid=11885>
- 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). <https://doi.org/10.1145/3158137>
- Neal Glew. 2005. Method Inlining, Dynamic Class Loading, and Type Soundness. *Journal of Object Technology* 4, 8 (2005). <https://doi.org/10.5381/jot.2005.4.8.a2>
- Urs Hölzle, Craig Chambers, and David Ungar. 1992. Debugging Optimized Code with Dynamic Deoptimization, In *Conference on Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/143103.143114>
- Sheng Liang and Gilad Bracha. 1998. Dynamic class loading in the Java virtual machine. In *Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/286936.286945>
- Jacob Matthews and Robert Bruce Findler. 2008. An operational semantics for Scheme. *Journal of Functional Programming* 18 (2008), Issue 1. <https://doi.org/10.1017/S0956796807006478>
- John McCarthy. 1978. History of LISP. In *History of programming languages (HOPL)*. <https://doi.org/10.1145/960118.808387>
- Phung Hua Nguyen and Jingling Xue. 2005. Interprocedural side-effect analysis and optimisation in the presence of dynamic class loading. In *Australasian Conference on Computer Science (ACSC)*. <https://doi.org/10.5555/1082161.1082163>
- Joe Gibbs Politz, Matthew J. Carroll, Benjamin S. Lerner, Justin Pombrio, and Shriram Krishnamurthi. 2012. A Tested Semantics for Getters, Setters, and Eval in JavaScript. In *Symposium on Dynamic Languages (DLS)*. <https://doi.org/10.1145/2384577.2384579>
- Gareth Stoye, Michael Hicks, Gavin Bierman, Peter Sewell, and Iulian Neamtii. 2007. Mutatis Mutandis: Safe and Predictable Dynamic Software Updating. *ACM Trans. Program. Lang. Syst.* 29, 4 (2007). <https://doi.org/10.1145/1255450.1255455>
- Julia Language Manual v1. 2020. *Redefining Methods*. <https://docs.julialang.org/en/v1/manual/methods/#Redefining-Methods-1>
- Francesco Zappa Nardelli, Julia Belyakova, Artem Pelenitsyn, Benjamin Chung, Jeff Bezanson, and Jan Vitek. 2018. Julia Subtyping: A Rational Reconstruction. *Proc. ACM Program. Lang.* 2, OOPSLA (2018). <https://doi.org/10.1145/3276483>