

Julia Subtyping: A Rational Reconstruction

FRANCESCO ZAPPA NARDELLI, Inria and Northeastern U.

JULIA BELYAKOVA, Czech Technical U. in Prague

ARTEM PELENITSYN, Czech Technical U. in Prague

BENJAMIN CHUNG, Northeastern U.

JEFF BEZANSON, Julia Computing

JAN VITEK, Northeastern U. and Czech Technical U. in Prague

Programming languages that support multiple dispatch rely on an expressive notion of subtyping to specify method applicability. In these languages, type annotations on method declarations are used to select, out of a potentially large set of methods, the one that is most appropriate for a particular tuple of arguments. Julia is a language for scientific computing built around multiple dispatch and an expressive subtyping relation. This paper provides the first formal definition of Julia’s subtype relation and motivates its design. We validate our specification empirically with an implementation of our definition that we compare against the existing Julia implementation on a collection of real-world programs. Our subtype implementation differs on 122 subtype tests out of 6,014,476. The first 120 differences are due to a bug in Julia that was fixed once reported; the remaining 2 are under discussion.

1 INTRODUCTION

Multiple dispatch is used in languages such as CLOS [DeMichiel and Gabriel 1987], Perl [Randal et al. 2003], R [Chambers 2014], Fortress [Allen et al. 2011], and Julia [Bezanson 2015]. It allows programmers to overload a generic function with multiple methods that implement the function for different type signatures; invocation of the function is resolved at run-time depending on the actual types of the arguments. The expressive power of multiple dispatch stems from the way it constrains the applicability of a method to a particular set of values. With it, programmers can write code that is concise and clear, as special cases, such as optimized versions of matrix multiplication, can be relegated to dedicated methods. The inset shows three of the 181 methods implementing multiplication in Julia’s standard library. The first method implements the case where a range is multiplied by a number. The second method is invoked on generic numbers: it explicitly converts the arguments to a common type via the `promote` function. The last method invokes native multiplication; its signature has a type variable T that can be instantiated to any integer type.

```
* (x::Number, r::Range) = range(x*first(r),...)
* (x::Number, y::Number) = *(promote(x,y)...)
* (x::T, y::T)
  where T <: Union{Signed,Unsigned} =
  mul_int(x,y)
```

For programmers, understanding multiple dispatch requires reasoning about the subtype relation. Consider the infix call `3 * x`. If x is bound to a float, only the second method is applicable. If, instead, x is an integer, then two methods are applicable and Julia’s runtime must identify the *most specific* one. Now, consider `3 * 4`, with argument type `Tuple{Int, Int}`. The signature of the first method is `Tuple{Number, Range}`. Tuples are covariant, so the runtime checks that `Int <: Number` and `Int <: Range`. Integers are subtypes of numbers, but not of ranges, so the first method is not applicable, but the second is, as `Tuple{Int, Int} <: Tuple{Number, Number}`. The third method is also applicable, as `Tuple{Int, Int}` is a subtype of `Tuple{T, T}` where `T <: Union{Signed, Unsigned}`; because there *exists* an instance of the variable T (namely `Int`) for which subtyping holds. As multiple methods are applicable, subtyping is used to compare their signatures; it holds that `Tuple{T, T}` where `T <: Union{Signed, Unsigned}` is a subtype of `Tuple{Number, Number}` because this holds *for all* instances of the variable T . The call will be dispatched, as expected, to the third method.

50 Subtyping can surprise programmers. For instance, is type `Tuple{String, Int}` a subtype of type
 51 `Tuple{Union{Bool, T}, T}` where `T`? One could choose to instantiate `T` with `Union{String, Int}`, and,
 52 in a system with union and tuple types such as [Vouillon 2004], subtyping would hold. In Julia
 53 this is not the case because of the *diagonal rule*. This rule requires that if a type variable appears
 54 more than once in covariant position, it can be instantiated only with a *concrete* type (e.g. `Int`). A
 55 `Union` is not concrete and thus cannot be used to instantiate `T`. The diagonal rule is used to restrict
 56 applicability of methods to values that have the same representation, which enables expressing
 57 common scientific computing idioms: it correctly prevents `3 * 0x4`, whose type is `Tuple{Int, UInt8}`,
 58 to dispatch to the third method above. However, the rule's interaction with other features can be
 59 complex. Consider `Tuple{Bool, Int}`; it is a subtype of `Tuple{Union{Bool, T}, T}` where `T` because `T`
 60 can be instantiated to `Int` and the union type matches with `Bool`, which lets us build a derivation.

61 Our goal is to provide an account of Julia's subtype relation that allows programmers to reason
 62 about their code, Julia implementors to evaluate the correctness of the compiler, and language de-
 63 signers to study an interesting point in the language design space. This has proved to be surprisingly
 64 difficult for the following three reasons. *Dynamic typing*: Julia does not have a static type system,
 65 so subtyping is only needed for multiple dispatch. Properties one would expect from such a relation
 66 may not hold. For instance, while working on
 67 this paper we discovered that, in the produc-
 68 tion implementation of subtyping, reflexivity
 69 did not hold. It was an implementation mis-
 70 take that was promptly fixed, but it is telling
 71 that it went undiscovered. *No formal specifica-
 72 tion*: apart from a partial description in prose
 73 in Bezanson [2015], the only specification of
 74 subtyping is 2,800 lines of heavily optimized,
 75 undocumented C code (a snippet is shown in
 76 Fig. 1 for your enjoyment). Inspection of Julia's
 77 2017 commit log shows that only three out of
 78 over 600 contributors made substantial edits to
 79 `subtype.c`, the file that implements it. Anecdotal
 80 evidence, based on discussion with users,
 81 suggests that the subtype relation is perceived
 82 as a black box that behaves mysteriously.

83 *Unique combination of features*: Julia's type language features an original combination of *nominal
 84 single subtyping*, *union types*, *existential types*, *covariant tuples*, *invariant parametric datatypes*,
 85 *distributivity*, and *singleton types*, as well as the *diagonal rule*. One source of inspiration for the design
 86 of subtyping in Julia was semantic subtyping [Frisch et al. 2002, 2008], but practical considerations
 87 caused the language to evolve in a unique direction. Table 1 illustrates Julia's unique combination
 88 of features; further discussion is in the related work section.

89 Given the absence of a denotational model of subtyping, it was clear from the outset that we would
 90 not be able to prove our specification correct. Instead, we provide an implementation of the subtype
 91 relation that mirrors the specification, and then validate empirically that our specification-based
 92 implementation agrees with the existing implementation. Our contributions are the following:

- 93 (1) The first specification of Julia subtyping, covering all features of Julia except `Vararg` (omitted
 94 as it would decrease readability for little conceptual benefit).
- 95 (2) An implementation of our specification and a validation of that implementation against the
 96 reference implementation on a suite of real-world packages.
 97
 98

```

int forall_exists_subtype(jl_value_t *x,
jl_value_t *y, jl_stenv_t *e, int param) {
    save_env(e, &saved, &se);
    memset(e->Lunions.stack, 0,
          sizeof(e->Lunions.stack));
    int lastset = 0; int sub;
    while (1) {
        sub = exists_subtype(x, y, e, saved, &se, param);
        int set = e->Lunions.more;
        if (!sub || !set) break;
        save_env(e, &saved, &se);
        for (int i = set; i <= lastset; i++)
            statestack_set(&e->Lunions, i, 0);
        lastset = set - 1;
        statestack_set(&e->Lunions, lastset, 1);
    }
    free(se.buf);
    return sub;
}

```

Fig. 1. Julia `subtype.c` extracted verbatim.

99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147

	[Chambets ea. 94]	[Bourdoncle ea. 97]	[Litvinov 03]	[Vouillon 04]	[Frisch ea. 08]	[Cameron ea. 08]	[Smith ea. 08]	[Allen ea. 11]	Julia
Nominal Subtyping	●	◐	●	○	○	●	●	●	●
Invariant Type Constructors	○	●	●	○	◐	●	●	●	●
Covariant Type Constructors	◐ ²	●	●	●	◐ ⁴	○	○	○	◐ ⁴
Explicit Union Types	●	◐ ³	●	●	●	○	●	●	●
Distributivity ¹	○	○	○	○	●	○	○	○	●
Bounded Existential Types	○	○	○	○	○	●	◐	◐	●
Diagonal Rule	○	○	○	○	○	○	○	○	●

(1) Union/tuple or union/intersection distributivity visible to users.
 (2) Built-in covariant vectors used internally for arguments but not available to users.
 (3) Constraints on type parameters seen as union/intersection types.
 (4) Only built-in covariant tuples.

Table 1. Julia subtyping compared.

(3) Identification of problems with Julia’s design and implementation. Four bugs have been fixed and one proposal was accepted for the next revision of Julia.

Non-results. We do not provide a proof of soundness, as there is no formal semantics of Julia. We do not compare performance between our implementation and the Julia subtype code as our code is written so as to mirror our rules one to one, whereas the Julia implementation is written in C and is heavily optimized. We do not attempt to provide a “better” definition for subtyping; we leave that to future work. And, lastly, we do not prove decidability of Julia’s subtyping or of its underlying algorithm.

Artifacts. Our implementation of subtyping is available in the supplementary material. All our infrastructure has been submitted for artifact evaluation. The paper web-page [Zappa Nardelli et al. 2018] complements the paper with additional data and resources.

2 BACKGROUND: JULIA

Julia is a language designed for scientific computing, released in 2012, which has achieved a degree of success — as evidenced by over 6,000 independently developed packages hosted on GitHub. Julia is a high-level, dynamic, memory-safe language without a type system but with user-defined type declarations and a rich type annotation sublanguage. Its design, therefore, reflects the tension between supporting dynamic features and ensuring efficient native code generation. As with other dynamic languages, the implementation executes any grammatically correct program and can load new code with `eval`. This is challenging for a compiler, yet Julia’s LLVM-based back-end can be competitive with C [Bezanson et al. 2017].

While Julia has a rich type annotation language, we emphasize its lack of a static type system. The first method for function `f`, shown here, does not have type annotations on its argument and will work as long as there is an addition method for the actual value of `x`. The second method is specific to strings, but invocations will fail at run-time unless a multiplication method is provided between a string and an integer. There is no notion of soundness in Julia, even for fully type-annotated

```
f(x) = x + 1
f(x::String) = x * 3
```

programs. If a method call does not have a most specific method, a runtime error will be reported. Ambiguity in dispatch is always resolved dynamically.

Julia types are *nominal*: the hierarchical relationship between types is specified explicitly by the programmer rather than inferred from their structure. This enables a function to behave differently on different types even if they have the same representation. Julia types are *parametric*: user-defined types can be parametrized by other types (and by values of primitive types as integers and booleans).

Top and Bottom. The abstract type `Any` is the type of all values and is the default when type annotations are omitted. The empty union `Union{}` is a subtype of all types; it is not inhabited by any value. Unlike many common languages, Julia does not have a null value or a null type that is a subtype of all types.

Datatypes. Datatypes can be *abstract* or *concrete*. Abstract datatypes may have subtypes but cannot have fields. Concrete datatypes have fields but cannot have declared subtypes. Every value is an instance of a concrete `DataType` that has a size, storage layout, supertype (`Any` if not otherwise declared), and, optionally, field names. Consider the inset definitions.

The first declaration introduces `Integer` as a subtype of `Real`. The type is abstract; as such it cannot be instantiated.

The second declaration introduces a

concrete, primitive, type for boolean values and specifies that its size is 8 bits; this type cannot be further subtyped. The last declaration introduces a concrete, mutable structure `PointRB` with two fields, `x` of abstract type `Real` and `y` of concrete type `Bool`. Abstract types are always stored as references, while concrete types are unboxed.

```
abstract type Integer <: Real end
primitive type Bool <: Integer 8 end
mutable struct PointRB <: Any x::Real y::Bool end
```

Type Unions. A union is an abstract type which includes, as values, all instances of any of its argument types. Thus the type `Union{Integer,AbstractString}` denotes any values from the set of `Integer` and `AbstractString` values.

Parametric Datatypes. The following defines an immutable, parametrized, concrete type.

`Rational`, with no argument, is a valid type, containing all instances `Rational{Int}`, `Rational{UInt}`, `Rational{Int8}`, etc. Thus, the following holds: `Rational{Int} <: Rational`. Type parameters are *invariant*, thus the following does not hold:

`Rational{Int} <: Rational{Integer}`. This restriction stems from practical considerations: the memory layout of abstract types (`Integer`) and concrete types (`Int`) is different and can impact the representation of the parametric type. In a type declaration, parameters can be used to instantiate the supertype. This allows the declaration of an `AbstractVector` of as a mono-dimensional `AbstractArray` of values of type `T`:

```
struct Rational{T<:Integer} <: Real
    num::T
    den::T
end
```

```
abstract type AbstractVector{T} <: AbstractArray{T,1} end
```

Tuple types. Tuples are an abstraction of the arguments of a function; a tuple type is a parametrized immutable type where each parameter is the type of one field. Tuple types may have any number of parameters, and they are *covariant* in their parameters: `Tuple{Int}` is a subtype of `Tuple{Any}`. `Tuple{Any}` is considered an abstract type; tuple types are only concrete if their parameters are.

UnionAll. A parametric type without argument like `Rational` acts as a supertype of all its instances (`Rational{Int}` etc.) because it is a different kind of type called a *UnionAll* type. Julia documentation

describes UnionAll types as “the iterated union of types for all values of some parameter”; a more accurate way to write such type is `Rational{T}` where `Union{ } <: T <: Any`, meaning all values whose type is `Rational{T}` for some value of `T`. UnionAll types correspond to bounded existential types in the literature, and a more usual notation for the type above would be $\exists T. \text{Rational}\{T\}$. Julia does not have explicit *pack/unpack* operations; UnionAll types are abstract. Each `where` introduces a single type variable. The combination of parametric and existential types is expressive: the type of 1-dimensional arrays can be simply specified by `Array{T, 1}` where `T`. Type variable bounds can refer to outer type variables. For example, `Tuple{T, Array{S}}` where `S <: AbstractArray{T}` where `T <: Real` refers to 2-tuples whose first element is some `Real`, and whose second element is an `Array` of any kind of array whose element type contains the type of the first tuple element. The `where` keyword itself can be nested. Consider the types `Array{Array{T, 1} where T, 1}` and `Array{Array{T, 1}, 1}` where `T`. The former defines a 1-dimensional array of 1-dimensional arrays; each of the inner arrays consists of objects of the same type, but this type may vary from one inner array to the next. The latter type instead defines a 1-dimensional array of 1-dimensional arrays all of whose inner arrays must have the same type. UnionAll types can be explicitly instantiated with the type application syntax `(t where T){t'}`; partial instantiation is supported, and for instance `Array{Int}` denotes arrays of integers of arbitrary dimension.

Singleton Types. There are two special abstract parametric types. For an arbitrary type `t`, `Type{t}` defines a type whose only instance is `t` itself; similarly `Val{3}` is used to create the singleton type for integer 3.

3 SUBTYPING IN JULIA

We focus on the following grammar of types, denoted by `t`:

$$\begin{aligned}
 t & ::= \text{Any} \mid \text{Union}\{t_1, \dots, t_n\} \mid \text{Tuple}\{a_1, \dots, a_n\} \mid \text{name}\{a_1, \dots, a_n\} \mid t \text{ where } t_1 <: T <: t_2 \\
 & \quad \mid T \mid \text{Type}\{a\} \mid \text{DataType} \mid \text{Union} \mid \text{UnionAll} \\
 a & ::= t \mid v
 \end{aligned}$$

The variable `v` ranges over *plain-bit* values: in addition to types, plain-bit values can be used to instantiate all parametric types. Our only omission is the `Vararg` construct, discussed at the end of this section. We follow Julia’s conventions. We write type variables in big-caps. Given `t where t1 <: T <: t2`, the variable `T` binds in the type `t`, but not in `t1` or `t2`. We abbreviate with `Bot` the empty union type `Union{}`, the subtype of all types. In the `where` construct, omitted lower bounds (resp. upper bounds) for type variables default to `Bot` (resp. `Any`); the notation `t where T` is thus a shorthand for `t where Bot <: T <: Any`. We also remove empty applications and denote `name{ }` simply with `name`. We assume that all user-defined types are recorded in a global environment `tds` which for each type stores its name, attribute, type parameters with bounds, and the declared supertype. A supertype can refer to the parameters of the type being defined. Searching a type name, e.g. `name` in `tds`, returns either its definition, denoted:

$$\text{attr name}\{t_1 <: T_1 <: t'_1, \dots, t_m <: T_m <: t'_m\} <: t'' \in \text{tds}$$

or fails. The attribute, denoted `attr`, records whether the defined type is abstract or concrete. When not relevant, we omit the lower and upper bounds of the binding type variables.

Julia’s frontend simplifies types written by the programmer e.g. by removing redundant unions or parameters. We choose to formalize the subtype relation over the source syntax of types, rather than the internal Julia representation. Our approach enables reasoning about the type simplification phase itself: it is arguable that, to prevent unexpected behaviors, all frontend type transformations

ought to be correct with respect to the type equivalence induced by the subtype relation. For instance this allowed us to identify a surprising behavior, discussed in Sec. 4.4.

Julia defines a *typeof* function that returns the concrete type of a value. Since types are themselves values, it is legitimate to invoke *typeof* on them, and the types `DataType`, `Union`, and `UnionAll` play the role of kinds. Intuitively, *typeof*(*t*) analyses the top-level constructor of *t* and returns `UnionAll` if it is a `where` construct, `Union` if it is a `Union` construct, and `DataType` otherwise. The *typeof* function plays a role in the subtyping rule for the `Type{a}` constructor, and we additionally rely on it to rule out badly formed types. A precise formalization of *typeof* is reported in Appendix A.

3.1 Understanding Subtyping

The literature never studied a subtype system with all the features of Julia. Unexpected, subtle, interactions between existential types and distributivity of union/tuple types forced us to depart from established approaches. We give an informal overview of the subtype relation, pointing out where, and why, standard rules fall short.

Building intuition. Two subtyping rules follow naturally from Julia’s design: parametric types are *invariant* in their parameters, while tuples are *covariant*. The former follows immediately from Julia’s memory representation of values. An array of dissimilar values is represented as a list of pointers to the boxed values, under type `Vector{Any}`. However, if all the values are primitive, then an unboxed representation is used. For instance, a vector of 32-bit integers is represented as an array of machine integers, under type `Vector{Int32}`. It would be wrong to treat `Vector{Int32}` as a subtype of `Vector{Any}`, as pointers can require more than 32 bits. This is *incompatible* with a semantic subtyping interpretation of the subtype relation [Frisch et al. 2002]. Invariance of type application is enforced via $name\{t_1, \dots, t_n\} <: name\{t'_1, \dots, t'_n\}$ iff for all i , $t_i <: t'_i$ and $t'_i <: t_i$. Tuples are an abstraction of the arguments of a function: covariance enables dispatch to succeed when the function arguments are a subtype of a more general function signature. Covariance of tuple types is usually enforced via $Tuple\{t_1, \dots, t_n\} <: Tuple\{t'_1, \dots, t'_n\}$ iff for all i , $t_i <: t'_i$.

Subtyping union types follows instead the semantic subtyping intuition, of Vouillon [2004] or Frisch et al. [2002]. Subtyping union types is asymmetrical but intuitive. Whenever a union type appears on the left hand side of a subtyping judgment, as in $Union\{t_1, \dots, t_n\} <: t$, all the types $t_1 \dots t_n$ must be subtypes of t . In contrast, if a union type appears on the right-hand side of a judgment instead, as in $t <: Union\{t_1, \dots, t_n\}$, then there needs to be only one type t_i in $t_1 \dots t_n$ that is a supertype of t . Combining the two, a judgment $Union\{t_1, \dots, t_n\} <: Union\{t'_1, \dots, t'_n\}$ thus reads as: *forall* types t_i , there *exists* a type t'_j such that $t_i <: t'_j$.

These rules are simple in isolation, but their interaction with other Julia features is not.

Unions and Tuples. Covariant tuples should be *distributive* with respect to unions. In particular, it should hold that:

$$Tuple\{Union\{t_1, t_2\}, t\} <: Union\{Tuple\{t_1, t\}, Tuple\{t_2, t\}\}$$

but it is known since Vouillon [2004] that the judgment cannot be derived from the above rules. The rule for tuples does not apply, while decomposing immediately the union on the right, picking up either `Tuple{t1, t}` or `Tuple{t2, t}` does not allow to conclude. Indeed, if a derivation commits to an element of a union type in the right-hand side before having explored all the (possibly nested) union types in the left-hand side, the derivation has effectively performed an exist/forall search, rather than a forall/exist one, losing the option to choose a different matching type for all the types in the union on the left-hand side.

295 A standard approach, relied upon e.g. by the CDuce language, solves this conundrum by
 296 rewriting types into their disjunctive normal form, that is, as unions of intersections of liter-
 297 als, *ahead-of-time*. Rewriting types as top-level unions of other types is correct in CDuce se-
 298 mantic model, but it is unsound in Julia, due to invariant constructors. For example, in Julia,
 299 the type `Vector{Union{Int,Bool}}`, denoting the set of vectors whose elements are integers or
 300 booleans, cannot be expressed with a top-level union. It would be incorrect to rewrite it as
 301 `Union{Vector{Int}, Vector{Bool}}`, the set of vectors whose elements are all integers or all booleans.
 302 Despite this, the distributivity law above holds in Julia and the subtype relation must successfully
 303 derive similar judgments. Julia's implementation thus relies on an efficient, but complex and fragile,
 304 backtracking mechanism to keep the forall/exist ordering of quantifications correct independently
 305 of the syntactic structure of types. This algorithm is hard to formalize and to reason about.

306 It is however possible to formalize an exhaustive search on top of the aforementioned rules for
 307 tuples and unions. The key intuition is that rewriting a `Tuple{Union{t1, t2}, t}` type that occurs on
 308 the left-hand side of a subtyping judgment into the equivalent `Union{Tuple{t1, t}, Tuple{t2, t}}` has
 309 the effect of syntactically anticipating the union types (and thus the induced forall quantifications)
 310 as much as possible in a derivation. Similarly, performing the opposite rewriting whenever a `Union`
 311 type occurs on the right-hand side of a subtyping judgment delays the existential quantifications.
 312 Care is required to lift, or unlift, `where` constructs correctly, but by adding rules that apply these
 313 rewritings *dynamically*, while building the derivation tree, we define a complete subtype relation
 314 on top of the intuitive subtype rules for tuples, invariant constructors, and union types above.

315 *UnionAll, environments, and bounds.* Julia's type system features bounded existential types,
 316 denoted $t \text{ where } t_1 <: \top <: t_2$, and (confusingly) referred to as *iterated unions* or *UnionAll* types.
 317 Analogously to union types, the subtyping rules for bounded existentials must have either a forall
 318 or an exist semantics according to whether the existential appears on the left or right of the
 319 subtyping judgment. So

$$320 \quad (t \text{ where } t_1 <: \top <: t_2) <: t'$$

321 is satisfied if *forall* types t_3 supertype of t_1 and subtype of t_2 it holds that $t[t_3/\top] <: t'$. Conversely,

$$322 \quad t' <: (t \text{ where } t_1 <: \top <: t_2)$$

323 is satisfied if *there exists* a type t_3 supertype of t_1 and subtype of t_2 such that $t[t_3/\top] <: t'$. The
 324 correct quantification of a variable is specified by the position of the `where` construct that introduced
 325 it, not by where the variable occurs in the judgment. Intuitively, when checking if:

$$326 \quad \text{Ref}\{\text{Int}\} <: \text{Ref}\{\top\} \text{ where } \top$$

327 the invariance of `Ref` will force us to check both $\text{Int} <: \top$ and $\top <: \text{Int}$. In both cases, the subtyping
 328 check must be performed assuming \top has an *exist* semantics; in this case both constraints can be
 329 satisfied by picking \top to be `Int`.

330 To keep track of the semantics of each variable, we record them in an environment E . A variable \top
 331 introduced by a `where` on the left of the subtyping judgment is recorded as L_{\top}^{ub} , a variable introduced
 332 on the right as R_{\top}^{ub} : lb and ub are the lower bound and upper bound types for the variable. The
 333 judgments we consider thus have the form $E \vdash t_1 <: t_2$. Given an environment in which \top has
 334 a forall (that is, L) semantics, we distinguish two additional cases. If \top appears on the left of the
 335 judgment, then the judgment can be satisfied only if the upper-bound for \top is smaller than t :

$$336 \quad L_{\top}^{ub} \vdash \top <: t \quad \text{if} \quad L_{\top}^{ub} \vdash ub <: t .$$

337 If instead \top appears on the right, then it is the lower bound for \top that must be a supertype of t :

$$338 \quad L_{\top}^{ub} \vdash t <: \top \quad \text{if} \quad L_{\top}^{ub} \vdash t <: lb .$$

Right-introduced variables have exist semantics, so the least constraining bound can be chosen to satisfy a judgment, resulting in:

$$R_{lb}^{ub} \vdash t <: \top \text{ if } R_{lb}^{ub} \vdash t <: ub \qquad R_{lb}^{ub} \vdash \top <: t \text{ if } R_{lb}^{ub} \vdash lb <: t$$

It might be surprising that variables introduced by a `where` on the left of the judgment suddenly appear on its right, but this is a consequence of the invariance check for type application. For instance, when checking $(\text{Ref}\{\top\} \text{ where } \text{Int} <: \top <: \text{Int}) <: \text{Ref}\{\text{Int}\}$, the \top variable is introduced on the left, but we must prove both $L_{\text{Int}}^{\text{Int}} \vdash \top <: \text{Int}$ and $L_{\text{Int}}^{\text{Int}} \vdash \text{Int} <: \top$.

Matching right-introduced variables requires extra care because these types are not in subtype relation:

$$\text{Tuple}\{\text{Ref}\{\text{Int}\}, \text{Ref}\{\text{Bool}\}\} \not<: \text{Tuple}\{\text{Ref}\{\top\}, \text{Ref}\{\top\}\} \text{ where } \top \quad (1)$$

Indeed, there does not exist a type t for \top that satisfies both the constraints $\text{Int} <: t <: \text{Int}$ and $\text{Bool} <: t <: \text{Bool}$. To account for this, whenever an existential variable is matched against a type, its bounds are updated to handle the new hypotheses on the variable, and the updated environments are propagated across the subtyping derivation tree. The actual subtyping judgment thus has the form:

$$E \vdash t_1 <: t_2 \vdash E'$$

and should be read as: *in the environment E , type t_1 is a subtype of t_2 , with updated constraints E' .*

For instance, the judgment:

$$R_{\text{Bot}}^{\text{Any}} \vdash \text{Ref}\{\text{Int}\} <: \text{Ref}\{\top\} \vdash R_{\text{Int}}^{\text{Int}}$$

states that if \top has exist semantics and no bounds, then $\text{Ref}\{\text{Int}\} <: \text{Ref}\{\top\}$ holds, and later uses of \top must satisfy the updated bounds $\text{Int} <: \top <: \text{Int}$. The subtyping rule for tuples thus chains the environments across subtype tests of tuple elements. In the judgment (1) the second element $\text{Ref}\{\text{Bool}\} <: \text{Ref}\{\top\}$ is thus checked assuming $R_{\text{Int}}^{\text{Int}}$ and the derivation fails accordingly.

Environment structure. The environment itself has a non-trivial structure. First, an environment E is composed of two lists, denoted by $E.\text{curr}$ and $E.\text{past}$. The former, $E.\text{curr}$, is a stack of the variables currently in scope (growing on the right), reflecting the order in which variables have been added to the scope. In addition to variables, $E.\text{curr}$ records *barriers*: tags pushed in the environment whenever the subtype check encounters an invariant constructor. Barriers will be discussed later. The second list, $E.\text{past}$, keeps track of variables which are not any longer in scope. Consider the judgment:

$$\text{Tuple}\{\text{Ref}\{S\} \text{ where } S <: \text{Int}\} <: \text{Tuple}\{\text{Ref}\{\top\}\} \text{ where } \top$$

In the derivation the variable \top is introduced before the variable S and the judgment

$$R_{\text{Bot}}^{\text{Any}}, L_{\text{Bot}}^{\text{Int}} \vdash \text{Ref}\{S\} <: \text{Ref}\{\top\} \vdash R_{\text{S}}^{\text{S}}, L_{\text{Bot}}^{\text{Int}}$$

thus appears in the derivation tree. A naive rule for `where` would discharge the variable S from the environment, obtaining:

$$R_{\text{Bot}}^{\text{Any}} \vdash (\text{Ref}\{S\} \text{ where } S <: \text{Int}) <: \text{Ref}\{\top\} \vdash R_{\text{S}}^{\text{S}}$$

The type variable S is still mentioned in constraints of variables in scope, but it is not any longer defined by the environment. If the variable \top is subsequently matched against other types, the subtyping algorithm cannot know if the variable S appearing in the bounds of \top has a forall or exist semantics, nor which are its bounds. Discharged variables are thus stored in $E.\text{past}$ and accessed whenever required. The subtyping rules guarantee that it is never necessary to update the bounds of a no-longer-in-scope variable. Relying on a separate $E.\text{past}$ environment avoids confusion when rules must determine precisely the scope of each variable, as motivated in the next paragraph.

Variables can be subject to unsatisfiable constraints. For instance, the subtype relation

$$\text{Tuple}\{\text{Real}, \text{Ref}\{\text{Int}\}\} \not\prec: \text{Tuple}\{\text{S}, \text{Ref}\{\text{T}\}\} \text{ where } \text{S} <: \text{T} \text{ where } \text{T}$$

does not hold because the type variables are subject to the three unsatisfiable constraints below:

$$\text{Real} <: \text{S} \quad \text{Int} <: \text{T} <: \text{Int} \quad \text{S} <: \text{T}$$

and in Julia, $\text{Real} \not\prec: \text{Int}$. The subtype algorithm records these constraints in the environment as $R_{\text{T}}^{\text{Int}}, R_{\text{Int}}^{\text{T}}, R_{\text{Real}}^{\text{T}}$, and whenever a right-variable is discharged, it checks that its lower bound is a subtype of its upper bound. In the example above, the derivation is invalidated by the failure of the consistency check for S:

$$R_{\text{T}}^{\text{Int}}, R_{\text{Int}}^{\text{T}}, R_{\text{Real}}^{\text{T}} \vdash \text{Real} \not\prec: \text{T}$$

From forall/exist to exist/forall. In some cases enforcing the correct ordering of type variable quantifications requires extra care. Consider the judgment:

$$\text{Vector}\{\text{Vector}\{\text{T}\} \text{ where } \text{T}\} \not\prec: \text{Vector}\{\text{Vector}\{\text{S}\}\} \text{ where } \text{S}$$

The type on the left denotes the set of all the vectors of vectors of elements of some type; the type on the right requires a common type for all the inner vectors. For instance the value $[[1, 2], [\text{True}, \text{False}]]$ belongs to the first, but not the second, type. Unfortunately, the rules sketched so far let us build a successful subtype derivation. The variables S and T are introduced in the environment, and then the left-to-right and right-to-left checks

$$R_{\text{Bot}}^{\text{Any}}, L_{\text{Bot}}^{\text{Any}} \vdash \text{T} <: \text{S} \quad \text{and} \quad R_{\text{T}}^{\text{Any}}, L_{\text{Bot}}^{\text{Any}} \vdash \text{S} <: \text{T}$$

are performed. These trivially succeed because for all instances of T there is a matching type for S.

Let us focus on the quantification order of the variables in the above judgment. It is still true that variables introduced on the left have a forall semantics, and variables introduced on the right have exist semantics. However, here we must find an instance of S such that forall T the judgment holds: perhaps surprisingly, the outer invariant construct `Vector` forces the inversion of the order of quantifications. Instead of a forall/exist query we must solve an *exist/forall* one. To correctly account for inversion in the order of quantifications, derivations must keep track of the relative ordering of variable introductions and invariant constructors. For this, the environment $E.\text{curr}$ is kept ordered, and *barrier* tags are pushed in $E.\text{curr}$ whenever the derivation goes through an invariant constructor.

We say that a variable S is *outside* a variable T in an environment E if S precedes T in $E.\text{curr}$ and they are separated by a barrier tag in $E.\text{curr}$. In our running example, the first check thus becomes:

$$R_{\text{Bot}}^{\text{Any}}, \text{Barrier}, L_{\text{Bot}}^{\text{Any}} \vdash \text{T} <: \text{S}$$

The environment correctly identifies the variable S as outside T and the judgment should thus be interpreted as there exists an instance of S such that, forall instances of T, $\text{T} <: \text{S}$ holds. The variable S must thus be compared with the upper bound of T, deriving `Any` as lower bound:

$$R_{\text{Bot}}^{\text{Any}}, \text{Barrier}, L_{\text{Bot}}^{\text{Any}} \vdash \text{Any} <: \text{S} \vdash R_{\text{Any}}^{\text{Any}}$$

Again, given S outside T, the right-to-left check must now prove

$$R_{\text{Any}}^{\text{Any}}, \text{Barrier}, L_{\text{Bot}}^{\text{Any}} \vdash \text{S} <: \text{T}$$

that is, it must conclude that there exists an instance of S such that, forall instances of T, $\text{S} <: \text{T}$ holds. In other terms the variable S must be a subtype of the lower bound of T. This fails, as expected.

442 A subtlety: whenever the forall variable is constrained tightly and quantifies over only one type,
 443 the exist/forall quantification can still correctly succeed, as in the valid judgment below:

```
444 Vector{Vector{T} where Int <: T <: Int} <: Vector{Vector{S}} where S
```

446 *The diagonal rule.* Consider the Julia code in the inset that defines equality in terms of equal-
 447 ity of representations (computed by `===`) for all numerical types. This is correct provided that
 448 only values of the same type are compared, as in Julia `Int` and `Float` have different rep-
 449 resentations. The type of the `==` method is

```
449 ==(x::T, y::T) where T<:Number = x === y
```

450 `Tuple{T, T} where T <: Number`, and the usual interpretation of `UnionAll` allows `T` to range over all
 451 allowed types, including `Number`. This type is thus equivalent to `Tuple{Number, Number}` and would
 452 match values as `(3, 3.0)`, where the types of components of the tuples are different.

454 Being able to dispatch on whether two values have the same type is useful in practice, and the
 455 Julia subtype algorithm is extended with the so-called *diagonal rule* [The Julia Language 2018]. A
 456 variable is said to be in *covariant position* if only `Tuple`, `Union`, and `UnionAll` type constructors occur
 457 between an occurrence of the variable and the `where` construct that introduces it. The *diagonal rule*
 458 states that if a variable occurs more than once in covariant position, and never in invariant position,
 459 then it is restricted to ranging over only concrete types. In the type `Tuple{T, T} where T <: Number` the
 460 variable `T` is diagonal: this precludes it getting assigned the type `Union{Int, Float}` and matching
 461 the value `(3, 3.0)`. Observe that in the type `Tuple{Ref{T}, T, T} where T` the variable `T` occurs twice
 462 in covariant position, but also occurs in invariant position inside `Ref{T}`; the variable `T` is not
 463 considered diagonal because it is unambiguously determined by the subtyping algorithm. Albeit
 464 this design might appear arbitrary, it is informed by pragmatic considerations; the C# language
 465 implements similar constraints (the paper web-page has an example).

466 Enforcing the diagonal rule involves two distinct parts: counting the occurrences of covariant
 467 and invariant variables, and checking that diagonal variables are only instantiated with concrete
 468 types. Formalizing faithfully these tasks requires some additional boilerplate. The variable entries
 469 in the subtyping environment are extended with two counters to keep track of the number of
 470 occurrences encountered in the current subtyping derivation. These counters must be updated
 471 while the derivation is built. Consider again these judgments from Sec. 1:

```
472 Tuple{Bool, Int} <: Tuple{Union{Bool, T}, T} where T
473 Tuple{String, Int}  $\not<$ : Tuple{Union{Bool, T}, T} where T
```

475 The former holds because, even if in the right-hand side the variable `T` appears syntactically twice,
 476 it is possible to build a valid derivation that matches `T` only once. The variable `T` is not diagonal in
 477 the former judgment. In a valid derivation for the latter judgment, the variable `T` must occur twice
 478 in covariant position and its final lower bound is `Union{String, Int}`, which is not a concrete type.
 479 This is only one example but, in general, subtle interactions between union and existential types do
 480 not allow counting occurrences to be correctly performed statically; it must be a *dynamic* process.
 481 The check that diagonal variables are bound only to concrete types is then performed during the
 482 validation of the consistency of the environment.

483 3.2 Specification of the subtyping algorithm

484 Our formalization of Julia subtyping is reported in Fig. 2. It closely follows the intuitions presented
 485 in the previous section.

487 A *variable definition*, denoted $L_{lb}^T{}^{ub}$ or $R_{lb}^T{}^{ub}(co, io)$, specifies a variable name `T`, its lower bound `lb` and
 488 upper bound `ub`, and if it has forall (*L*) or exist (*R*) semantics. To model the diagonal rule, variable
 489 definitions for *R*-variables additionally keep two counters: *co* for covariant occurrences and *io* for

invariant occurrences. Our notation systematically omits the counters as they are only accessed and modified by the auxiliary functions *add*, *upd* and *consistent*. A *barrier* is a tag, denoted `Barrier`. An *environment*, denoted by E , is composed by two stacks, denoted $E.\text{curr}$ and $E.\text{past}$, of variable definitions and barriers. The following operations are defined on environments, where v ranges over variable definitions and barriers:

$add(v, E)$: push v at top of $E.\text{curr}$, with occurrence counters initialised to 0;

$del(\top, E)$: pop v from $E.\text{curr}$, check that it defines the variable \top , and push v at top of $E.\text{past}$;

$del(\text{Barrier}, E)$: pop v from $E.\text{curr}$ and check that it is a barrier tag;

$search(\top, E)$: return the variable definition found for \top in $E.\text{curr}$ or $E.\text{past}$; fail if the variable definition is not found;

$upd(\overset{R}{\top}_{lb}^{ub}, E)$: update the lower and upper bounds of the variable definition \top in $E.\text{curr}$; if the variable is found in E after a barrier then increase the invariant occurrence counter, and the covariant occurrence counter otherwise. Fail if the variable definition is not found;

$reset_occ_{E'}(E)$ and $max_occ_{E_1..E_n}(E)$: for each variable \top defined in E , $reset_occ_{E'}(E)$ updates its occurrence counting with the occurrence counting for \top in E' , while $max_occ_{E_1..E_n}(E)$ updates its occurrence counting with the max of the occurrence counting for \top in $E_1..E_n$;

$consistent(\top, E)$: search \top in E . If the search returns $\overset{L}{\top}_{lb}^{ub}$, then return true if $E \vdash lb <: ub$ and false otherwise; while building this judgment recursive consistency checks are disabled. If the search returns $\overset{R}{\top}_{lb}^{ub}(co, io)$, then check if $E \vdash lb <: ub$ is derivable. If not, return false. If yes, additionally check the diagonal rule: if $co > 1$ and $io = 0$ then its lower-bound lb must be a concrete type, as checked by the $is_concrete(lb)$ function. The definition of this function is non-trivial as a lower bound might depend on the values of other type variable bounds. For example, $\text{Vector}\{\top\}$ is equivalent to a concrete type $\text{Vector}\{\text{Int}\}$ only if both the upper and lower bounds of \top equal Int . At the time of writing, Julia's implementation of $is_concrete$ is heuristic and does not catch all possible concrete types. We omit its formalisation but our artifact includes a simple implementation. The shorthand $consistent(E)$ checks the consistency of all variables in the environment E .

We assume that types appearing in a judgment are well-formed, as enforced by the *typeof* relation. We comment the subtyping rules. The rule `ANY` states that `Any` is the super-type of all types. The rule `TUPLE_LIFT_UNION` rewrites tuple types on the left-hand-side of the judgment in disjunctive normal forms, making the distributivity of unions with respect to tuples derivable. This rule can be invoked multiple times in a subtype derivation, enabling rewriting tuples in disjunctive normal form even inside invariant constructors. Rewriting is performed by the auxiliary function $lift_union(t)$, which pulls union and where types out of tuples, anticipating syntactically the forall quantifications in a derivation. Symmetrically, the rule `TUPLE_UNLIFT_UNION` performs the opposite rewriting, delaying syntactically the exist quantifications, on union types appearing on the right-hand side of a judgment. The auxiliary function $unlift_union(t)$ returns a type t' such that $t = lift_union(t')$. Finally, the rule `TUPLE` checks covariant subtyping of the tuple elements. The constraints generated by subtyping each element are assumed by subsequent checks, consistency is verified at the end.

The, perhaps surprising, need for the `TUPLE_UNLIFT_UNION` rule is due to the *complex interaction between invariant constructors, union types, and existentials*. The following judgment:

$$\text{Ref}\{\text{Union}\{\text{Tuple}\{\text{Int}\}, \text{Tuple}\{\text{Bool}\}\}\} <: \text{Ref}\{\text{Tuple}\{\top\}\} \text{ where } \top$$

is valid because \top can be instantiated with $\text{Union}\{\text{Int}, \text{Bool}\}$. However building a derivation without the `TUPLE_UNLIFT_UNION` rule fails. Initially the left-to-right check for invariant application generates the constraint $\top >: \text{Union}\{\text{Int}, \text{Bool}\}$. Given the environment $\overset{R}{\top}_{\text{Union}\{\text{Int}, \text{Bool}\}}^{\text{Any}}$, the right-to-left check $\text{Tuple}\{\top\} <: \text{Union}\{\text{Tuple}\{\text{Int}\}, \text{Tuple}\{\text{Bool}\}\}$ gets stuck trying to prove

540			[TUPLE]
541	[TOP]	[REFL]	
542	$\frac{}{E \vdash t <: \text{Any} \vdash E}$	$\frac{}{E \vdash a <: a \vdash E}$	$\frac{E \vdash a_1 <: a'_1 \vdash E_1 \dots E_{n-1} \vdash a_n <: a'_n \vdash E_n \text{ consistent}(E_n)}{E \vdash \text{Tuple}\{a_1, \dots, a_n\} <: \text{Tuple}\{a'_1, \dots, a'_n\} \vdash E_n}$
543			
544			
545	[TUPLE_LIFT_UNION]	[TUPLE_UNLIFT_UNION]	
546	$\frac{t' = \text{lift_union}(\text{Tuple}\{a_1, \dots, a_n\})}{E \vdash t' <: t \vdash E'}$	$\frac{t' = \text{unlift_union}(\text{Union}\{t_1, \dots, t_n\})}{E \vdash t <: t' \vdash E'}$	
547			
548	$\frac{}{E \vdash \text{Tuple}\{a_1, \dots, a_n\} <: t \vdash E'}$	$\frac{}{E \vdash t <: \text{Union}\{t_1, \dots, t_n\} \vdash E'}$	
549			
550		[UNION_LEFT]	[UNION_RIGHT]
551	$\frac{E \vdash t_1 <: t \vdash E_1 \dots \text{reset_occ}_E(E_{n-1}) \vdash t_n <: t \vdash E_n}{E \vdash \text{Union}\{t_1, \dots, t_n\} <: t \vdash \text{max_occ}_{E_1..E_n}(E_n)}$		$\frac{\exists j. E \vdash t <: t_j \vdash E'}{E \vdash t <: \text{Union}\{t_1, \dots, t_n\} \vdash E'}$
552			
553			
554		[APP_INV]	[APP_SUPER]
555	$\frac{n \leq m \quad E_0 = \text{add}(\text{Barrier}, E)}{\forall 0 < i \leq n. E_{i-1} \vdash a_i <: a'_i \vdash E'_i \wedge E'_i \vdash a'_i <: a_i \vdash E_i}$		$\frac{\text{name}\{\top_1, \dots, \top_m, \dots\} <: t'' \in \text{tds}}{E \vdash t''[a_1/\top_1 \dots a_m/\top_m] <: t' \vdash E'}$
556			
557	$\frac{}{E \vdash \text{name}\{a_1, \dots, a_m\} <: \text{name}\{a'_1, \dots, a'_n\} \vdash \text{del}(\text{Barrier}, E_n)}$		$\frac{}{E \vdash \text{name}\{a_1, \dots, a_m\} <: t' \vdash E'}$
558			
559			[R_INTRO]
560	[L_INTRO]		
561	$\frac{\text{add}({}^L\top_{l_1}^{l_2}, E) \vdash t <: t' \vdash E'}{E \vdash t \text{ where } t_1 <: \top <: t_2 <: t' \vdash \text{del}(\top, E')}$	$\frac{\text{add}({}^R\top_{l_1}^{l_2}, E) \vdash t <: t' \vdash E' \text{ consistent}(\top, E')}{E \vdash t <: t' \text{ where } t_1 <: \top <: t_2 \vdash \text{del}(\top, E')}$	
562			
563			
564			[R_L]
565	[L_LEFT]	[L_RIGHT]	
566	$\frac{\text{search}(\top, E) = {}^L\top_l^u}{E \vdash u <: t \vdash E'}$	$\frac{\text{search}(\top, E) = {}^L\top_l^u}{E \vdash t <: l \vdash E'}$	$\frac{\text{search}(\top_1, E) = {}^R\top_{l_1}^{u_1} \quad \text{search}(\top_2, E) = {}^L\top_{l_2}^{u_2} \quad \text{outside}(\top_1, \top_2, E) \Rightarrow E \vdash u_2 <: l_2 \vdash E' \quad E \vdash u_1 <: l_2 \vdash E''}{E \vdash \top_1 <: \top_2 \vdash \text{upd}({}^R\top_{l_1}^{u_1} \text{Union}\{\top_1, l_2\}, E')}$
567			
568			
569			
570			
571			[R_RIGHT]
572	[R_LEFT]		
573	$\frac{\text{search}(\top, E) = {}^R\top_l^u}{E \vdash l <: t \vdash E'}$	$\frac{\text{search}(\top, E) = {}^R\top_l^u \quad (\text{is_var}(t) \wedge \text{search}(t, E) = {}^L\top_{l_1}^{u_1}) \Rightarrow \neg \text{outside}(\top, S, E)}{E \vdash t <: u \vdash E'}$	
574			
575	$\frac{}{E \vdash \top <: t \vdash \text{upd}({}^R\top_l^u, E')}$	$\frac{}{E \vdash t <: \top \vdash \text{upd}({}^R\top_{l_1}^{u_1} \text{Union}\{l, t\}, E')}$	
576			
577			
578	[TYPE_LEFT]	[TYPE_RIGHT]	
579	$\frac{\neg \text{is_var}(a_1)}{E \vdash \text{typeof}(a_1) <: t_2 \vdash E'}$	$\frac{\text{is_kind}(t_1) \quad \text{is_var}(t_2)}{E \vdash \text{Type}\{\top\} \text{ where } \top <: \text{Type}\{t_2\} \vdash E'}$	
580			
581	$\frac{}{E \vdash \text{Type}\{a_1\} <: t_2 \vdash E'}$	$\frac{}{E \vdash t_1 <: \text{Type}\{t_2\} \vdash E'}$	
582			
583			[TYPE_TYPE]
584	$\frac{\text{add}(\text{Barrier}, E) \vdash a_1 <: a_2 \vdash E' \quad E' \vdash a_2 <: a_1 \vdash E''}{E \vdash \text{Type}\{a_1\} <: \text{Type}\{a_2\} \vdash \text{del}(\text{Barrier}, E')}$		
585			
586			
587			
588			

Fig. 2. The subtype relation.

589 $\top <: \text{Int}$ or $\top <: \text{Bool}$. Rule `TUPLE_UNLIFT_UNION` enables rewriting the right-to-left check into
 590 `Tuple{T} <: Tuple{Union{Int, Bool}}`, which is provable because the existential quantification due
 591 to the Union in the right-hand side is syntactically delayed.

592 Rules `UNION_LEFT` and `UNION_RIGHT` implement the forall and exist semantics for union types on
 593 the left and on the right of the subtyping judgment. In rule `UNION_LEFT`, the constraints generated
 594 by subtyping each element are assumed by each subsequent check and thus propagated into the
 595 final constraints. Discarding these constraints would allow proving that `Pair{Union{Int, Bool}, Int}`
 596 is a subtype of `Pair{T, T}` where \top , which is incorrect. However, to count correctly the occurrences
 597 of variables for the diagonal rule, each forall subderivation must reset the dynamic counting of the
 598 occurrences to that of its initial state, while the occurrences of the variables in the final state must
 599 be updated with the max of their occurrences in each intermediary state. From `UNION_LEFT` we
 600 immediately derive that `Union{ }` is subtype of all types, because its hypothesis is trivially validated
 601 by the forall quantification over an empty set. We conjecture that, given a type `Union{t1, ..., tn}`,
 602 the order of the types t_i is irrelevant for subtyping, but a formal proof is non-trivial.

603 Type application is governed by `APP_INV` and `APP_SUPER`. When subtyping type applications with
 604 the same callee, the rule `APP_INV` pushes a barrier onto the environment and checks the invariance
 605 of the actual type parameters. Constraints are all propagated across all subtype checks. If all checks
 606 succeed, the latest barrier is deleted from the environment and the final constraints are passed on.
 607 A subtlety: the number of actual parameters on the right-hand side can be smaller than that on the
 608 left-hand side. It is indeed always the case that partial application gives rise to super-types; for
 609 example `Dict{Int, String} <: Dict{Int}` holds because `Dict{Int, String}` denotes all dictionaries
 610 associating integers to strings, while `Dict{Int}` denotes all dictionaries associating integers to
 611 arbitrary values: it is natural to consider the latter a supertype of the former. Rule `APP_SUPER`
 612 enables replacing a user-defined type by its supertype in the left-hand side of a judgment; while
 613 doing so, the rule also appropriately instantiates the type variables of the supertype.

614 Rules `L_INTRO` and `R_INTRO` add a where introduced variable to the current environment, spec-
 615 ifying the relevant forall (L) or exist (R) semantics, and attempt to build a subtype derivation in
 616 this extended environment. Finally, since it gets out of scope, the introduced variable is deleted
 617 from the `curr` list and added to the `past` list. Variables with exist semantics might have had their
 618 bounds updated in unsatisfiable way; before discarding them, the consistency of their bounds is
 619 checked by the `consistent(\top , E)` auxiliary function.

620 Subtyping for type variables is governed by rules `L_LEFT`, `L_RIGHT`, `R_LEFT` and `R_RIGHT`. Type
 621 variables with forall semantics are replaced with the hardest-to-satisfy bound: the upper bound
 622 if the variable is on the left of the judgment, and the lower bound if the variable is on the right.
 623 Variables with exist semantics are instead replaced with their easiest-to-satisfy bound, and, to
 624 keep track of the match, bounds of these variables are updated if a successful derivation is found,
 625 reflecting their new bound. By symmetry one would expect the rule `R_LEFT` to update \top upper
 626 bound with $t \cap u$. Until recently, it was believed that, because of invariance, the explicit ordering of
 627 the checks performed by rule `APP_INV` or `TYPE_TYPE` would ensure that $t <: u$ had already been
 628 checked by rule `R_RIGHT`. Therefore it would always hold that $t = t \cap u$, avoiding the need to
 629 compute intersections of Julia types. To everybody surprise this turned out to be false. Consider:

630
 631 `Vector{Vector{Number}} <: Vector{Union{Vector{Number}, Vector{S}}}` where `Int <: S <: Signed`
 632

633 This judgment contradicts the idea that `Vector{S}` can be subtype of `Vector{Number}` only if S is
 634 equivalent to `Number`, which is not possible here. However both Julia and our formalization can
 635 build a derivation for it: due to the existential on the right-hand side, the check that ought to ensure
 636 $t <: u$, that is `Number <: Signed`, is skipped when performing the left-to-right subtype check of the
 637

invariant constructor `Vector`. In the spirit of this work, our formalization faithfully mimic Julia behaviour. Consequences and possible mitigations to this design issue are discussed in Section 5.3.

To account for the exist/forall quantification inversion, the `R_RIGHT` does not apply if the type on the left is a variable with forall (that is, L) semantics and the variables are in the exists/forall quantification (the check $\neg \text{outside}(\top, S, E)$ is responsible for this). Matching R-L variables is specially dealt by the `R_L` rule, which also performs the necessary outside check: if the R -variable is outside, then the bounds on the L -variable must constraint it to only one type. For this the check $ub <: lb$ is sufficient, as the other direction is later verified by the environment consistency check.

Subtyping the `Type` construct is more subtle than expected. Recall that for each type or plain-bit value a , the singleton type `Type{a}` is an abstract type whose only instance is the object a . Subtyping two `Type{a}` is analogous to check invariance of constructors, as done by rule `TYPE_TYPE`. But there are additional cases to be considered. `Type{a}` is subtype of the type of a (e.g. `Type{1} <: Int`), as enforced by the rule `TYPE_LEFT`. Conversely, `Type{t}` has subtypes only if t is a type variable, and the only subtypes are kinds; the recursive check updates correctly the constraints for t .

Interestingly, reflexivity of subtyping is not derivable from these rules, due to the asymmetric treatment of L variables. Consider for instance $L_{\text{Bot}}^{\text{Any}} \vdash \top <: \top$: the judgment ought to be true, but the subtyping rules will independently replace the left and right occurrences of \top by upper and lower bounds, ignoring that the same variable was thus attempting to prove $L_{\text{Bot}}^{\text{Any}} \vdash \text{Any} <: \text{Bot}$. Julia 0.6.2 subtype algorithm systematically performs reflexivity checks on the fast path; reflexivity ought to hold. This is solved by explicitly adding the `REFL` rule to the system. Plain-bit values behave as singleton types; as such, the rule `REFL` is the only one that applies on plain-bit values.

We made the explicit choice of not baking transitivity into the subtype rules, expecting it to be derivable. This design choice allowed us to identify a bug in Julia 0.6.2, discussed in Sec. 4.4. More interestingly, it turned out that by exploiting empty tuples it is possible to build judgements for which transitivity does not hold, as discussed in Sec. 5.1. Although surprising, the programming practice is not affected because empty tuple types are not inhabited.

Unprovable judgments. Julia’s subtype algorithm, and in turn our formalization, cannot prove all judgments expected to hold. For instance it cannot prove:

```
(Tuple{ $\top$ } where  $\text{String} <: \top <: \text{Int}$ ) <: Union{ }    or    Tuple{Union{ }} <: Union{ }
```

despite all these types having no elements (the type on the left-hand side being a valid Julia type). Additionally, constraints on type variables that are declared in the type definitions, such as in `struct Foo{T<: Integer} end`, are not relied upon by subtyping; therefore it is not possible to prove judgments as `(Foo{T} where \top) <: Foo{T} where $\top <: \text{Integer}$` . For dispatch these are not issues, as similar examples do not occur in the programming practice.

Omitted features. Our work omits the `Vararg{T, N}` construct. This can be used as last parameter of a tuple to denote N trailing arguments of type T . Supporting it would add considerable boilerplate to the formalization to distinguish the case where a concrete integer has been supplied for N from the general case where it is left parametric, without adding interesting interactions between the type system features.

We mentioned that Julia type syntax allows to instantiate explicitly existential types, via the syntax `(t where $t_1 <: \top <: t_2$){ a }`. These types are immediately rewritten by Julia frontend into their equivalent “beta-reduced” type `t[a/ \top]`; this behavior can be modeled by a simple ahead-of-time rewrite step, which we omit for simplicity from our formalization, although it is performed by our reference implementation.

4 EMPIRICAL VALIDATION

Is the complexity of Julia’s subtype relation motivated by real-world requirements? If not, then a simpler notion of subtyping may be devised. Is our specification a model of the reference implementation? Perhaps we have over-engineered a specification with unnecessary rules or missed some important corner cases.

To answer both questions we present data obtained by considering 100 popular Julia projects from GitHub. We show through static analysis that developers avail themselves to the full power of the Julia type annotation sublanguage, and dynamic analysis allows us to answer whether ours is a faithful model.

4.1 Type annotations in the real-world

The need for an expressive subtype relation must be motivated by real-world use cases. We analyzed a corpus of projects and extracted statistics on all type declarations. In Fig. 3(a) each bar depicts the total number of types declared by a project, how many of those type declarations have at least one type parameter, and how many of those apply a bound to one of their parameters (package names are reported in the paper web-page). The total number of types declared in the corpus is 2 717, with the Merlin package defining the per-package maximum of 204 types. The median number of types declared in a package is 15, with 60% of the packages defining at least 10 types. The total number of parametric types is 815 and the number of bound parametric types (where the bound is not trivial, i.e. Any) is 341.

Fig. 3(b) depicts four statistics for each project regarding type annotations on methods: of methods with at least one argument of a type other than Any, methods with a union or a where clause,

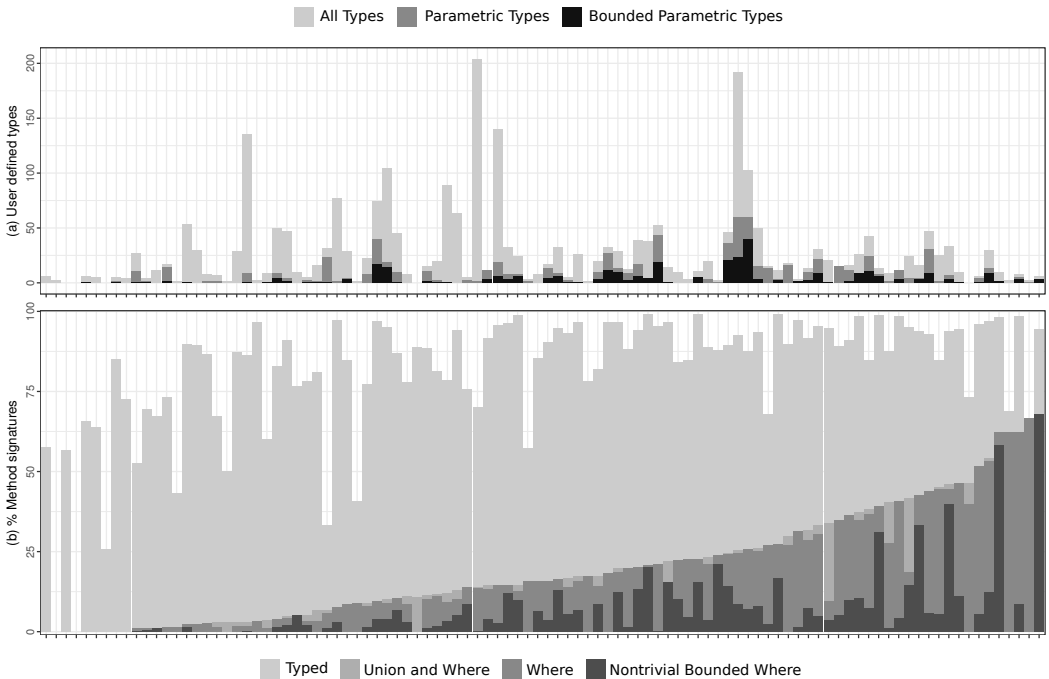


Fig. 3. (a) Type declarations (upper). (b) Proportion of complex types in methods (lower).

736 methods with a where clause, and methods with a where clause with at least one explicit bound.
 737 The mean proportion of union or where clauses is 18%, the mean proportion of where clauses
 738 is 16%, and the mean proportion of nontrivial where clauses is 6%. These numbers exclude the
 739 standard library: language implementors are more likely to use advanced features than end-users.

740 Overall these numbers suggest that programmers use the type declaration features of Julia and
 741 even the more complex where-bounded type annotations occur in most projects.
 742

743 4.2 Re-implementing Subtyping

744 We wrote our own implementation of Julia subtyping in Julia. Our development comprises about
 745 1,000 lines of code. For parsing types, we rely on Julia's parser but prevent front-end simplification.
 746 Our subtyping rules do not naturally define an algorithm. For instance, in the simple judgment
 747 `Union{Int, String} <: Union{String, Int}` two rules apply, namely `UNION_LEFT` and `UNION_RIGHT`,
 748 but a derivation requires `UNION_LEFT` to be used first. The challenge is thus to direct the search in
 749 the proof space. Our implementation applies the following strategy:
 750

- 751 (1) if the type on the right-hand side is `Any`, return true;
- 752 (2) if `UNION_LEFT` applies, then it has higher priority than all other rules (including `UNION_RIGHT`);
- 753 (3) if the left-hand side and right-hand side are both type variables, analyze their `L` or `R` annota-
 754 tions to check if the `R_L` rule applies;
- 755 (4) `TUPLE_LIFT_UNION` and `TUPLE_UNLIFT_UNION` have higher priority than `TUPLE`;
- 756 (5) replacing a type variable with its lower or upper bound is priority over decomposing `where`
 757 constructs;
- 758 (6) rule `APP_INV` has higher priority than rule `APP_SUPER`.

759 Additionally, in rules `L_LEFT`, `L_RIGHT`, and `R_L`, we substitute all occurrences of the left variable \top ;
 760 this simplifies the search of the states where the `REFL` rule must be applied. As we mentioned,
 761 when checking consistency of the environment, nested calls to the function `consistent` are disabled
 762 and assumed to succeed. Our implementation can exhaustively explore the search space due to
 763 the `UNION_RIGHT` existential quantification; this is the only source of backtracking. A complete
 764 implementation of the auxiliary function `unlift_union` is complex, our implementation is heuristic.

765 Our implementation outputs XML traces of derivations, useful for manual inspection, and collects
 766 statistics about the rule usage. Comforting the above claims, rule usage statistics confirm that
 767 all rules are needed to validate real-world subtype judgments, including the perhaps surprising
 768 `UNLIFT_UNION` and `R_L`, used respectively 27 and 1163 times on a benchmark of 6 millions tests
 769 (full numbers on the paper web-page).
 770

771 4.3 Subtyping validated

772 Our first benchmark is the test suites for the Julia subtype implementation internally used by Julia
 773 developers (`test/subtype.jl`): about 160 hand-picked tests inspired by bug-reports, and 335 097
 774 subtype queries over a list of 150 types for properties of the subtype relation such as `Union{T, S} ==`
 775 `T` implies `S <: T`. Our reference implementation passes all the tests from both test suites.
 776

777 To further explore corner cases, we developed a fuzzer that generates pairs of types; it builds on
 778 the approach pioneered by Claessen and Hughes [2000] to fuzz-test a unification algorithm. The
 779 key idea is to randomly generate one term, and derive a second term by *mutation* of the first one.
 780 Our fuzzer relies on the FEAT library by Duregård et al. [2012] to enumerate exhaustively up to a
 781 certain size *pre-types* over a simplified algebra of type constructors:
 782

$$783 p ::= \square \mid \text{Union}\{p_1, p_2\} \mid \text{Tuple}\{p_1, p_2\} \mid \Delta\{p\}$$

785 Every pre-type is then mutated by replacing the instances of the placeholder \square with either the
 786 concrete type `Int`, or the abstract type `Number`, or a type variable T . If a type variable is used, it is
 787 bound by a `where` clause at the top level. Additionally, the placeholder Δ is instantiated by either the
 788 concrete type constructor `Val`, or by the abstract type constructor `Ref`. Mutating from a simplified
 789 algebra ensures that generated types satisfy the well-formedness conditions imposed by Julia.
 790 Pairs of mutated types from the same pre-type are then passed to Julia and compared with our
 791 implementation. The generated types explore many corner cases of the algorithm (e.g. type variables
 792 appearing in covariant/invariant positions, tuple and union types, and various combinations of
 793 concrete/abstract matching). The fuzzer discovered three previously unknown issues in the Julia
 794 implementation (reported in Appendix B).

795 Finally, to stress test our implementation on realistic workloads, we instrumented the Julia C
 796 runtime to log all the calls to the subtype function. We traced the execution of the commands
 797 `using PkgName` and `Pkg.test(PkgName)` for 100 packages. The former builds the dispatch table for
 798 the methods defined in the imported package, calling `subtype` to sort their types. The latter executes
 799 tests, allowing us to explore the calls to `subtype` performed during execution. To reduce noise in
 800 the logs, we filter duplicates and precompile all dependencies of packages before logging anything.
 801 Our implementations of `typeof` and `subtype` require that all the declarations of user-defined types
 802 (denoted *tds* in the formalization) are passed as an argument. We wrote a tool that dumps the whole
 803 subtype hierarchy loaded in a Julia instance, by starting from the type `Any` and recursively walking
 804 the type hierarchy. We compare the outcome of each logged subtype test (ignoring those for which
 805 at least one type is not well-formed) with the result returned by our implementation.

806 Our subtype implementation differs from Julia's algorithm on 122 tests out of 7,612,469 (of which
 807 6,014,476 are non-trivial, that is they are not instances of the TOP or REFL rule, or subtype tuples
 808 with different number of arguments). Per-package numbers are reported on the paper web-page.
 809 We have manually inspected and analyzed the failures: 120 are instances a Julia 0.6.2 bug described
 810 below, which we reported and has been acknowledged and fixed. The remaining 2 cases are also
 811 suspected to be Julia bugs and are under examination.

812

813 4.4 Julia Bugs

814 Since the inception of this project, we have discovered and reported several issues affecting Julia 0.6.2
 815 subtyping. Appendix B lists all our bug reports with links to GitHub issues. Here is a discussion of
 816 the most interesting ones.

817 The majority of discrepancies on the realistic workload of the previous section, 120 differences
 818 out of 122, can be reduced to judgments of this form: $\text{Tuple}\{\text{Ref}\{\text{Ref}\{T\}\}\ \text{where } T, \text{Ref}\{T\}\ \text{where } T\} <:
 819 \text{Tuple}\{\text{Ref}\{S\}, S\}\ \text{where } S$. Such judgments hold in Julia 0.6.2, though they should not. We reported
 820 this issue to the Julia developers and it has been promptly fixed; it was due to the incorrect treatment
 821 of variables that go out of scope.

822 While developing our system, we also identified corner cases of Julia subtype design which
 823 were not covered by the reference test suite. These include subtle counterexamples to reflexivity
 824 and transitivity of subtyping in Julia 0.6.2 design. The transitivity one is interesting. These two
 825 judgments hold:

826

$$827 \quad \text{Tuple}\{\text{Number}, \text{Number}, \text{Ref}\{\text{Number}\}\} <: \text{Tuple}\{T, T, \text{Ref}\{T\}\}\ \text{where } T \quad (2)$$

$$828 \quad (\text{Tuple}\{T, T, \text{Ref}\{T\}\}\ \text{where } T) <: \text{Tuple}\{S, S, \text{Ref}\{Q\}\}\ \text{where } Q\ \text{where } S \quad (3)$$

829

830 but their transitive closure does not hold:

831

$$832 \quad \text{Tuple}\{\text{Number}, \text{Number}, \text{Ref}\{\text{Number}\}\} \not<: \text{Tuple}\{S, S, \text{Ref}\{Q\}\}\ \text{where } Q\ \text{where } S$$

833

834 Type variable S is diagonal, so its lower bound cannot be the abstract type `Number`. This is a design
 835 issue. In the judgment (3), it is incorrect to allow the diagonal variable S to be instantiated with the
 836 variable T because there is no guarantee that T itself is instantiated with a concrete type. Indeed in
 837 judgment (2) the variable T is not diagonal, and gets instantiated with the abstract type `Number`. Our
 838 formalization makes the problem evident: occurrence counters to identify diagonal variables are
 839 only kept and updated for right variables. The issue will be fixed in the next revision of Julia by
 840 keeping occurrence counters for left variables too.

841 One more surprising issue affects Julia's frontend. We mentioned that the frontend implicitly
 842 performs several simplifications when parsing and building the internal representation of types;
 843 for instance, T where T is replaced by `Any`, or `Union{Int, Int}` is replaced by `Int`. In general these
 844 transformations are sound, with one notable exception: simplification of types under explicit type
 845 instantiation is incorrect. In our formalization we have:

846 $(\text{Vector}\{T\} \text{ where } T \text{ where } S)\{Int\} \not\prec: \text{Vector}\{Int\}$

847 Julia incorrectly simplifies the type on the left-hand side as `Vector{Int}` (while it would be correct
 848 to rewrite it as `Vector{T} where T`) and concludes that subtyping holds. We have reported this
 849 arguably incorrect behavior.

851 5 DESIGN TRADEOFFS

852 In this work we set out to formalize the subtype relation as implemented by Julia 0.6.2; while
 853 doing so we have contributed to identifying both bugs in the current implementation and issues
 854 in the design. Most of these have been addressed by developers and are fixed in Julia 0.7dev, the
 855 development branch of Julia. In this section we briefly review alternative proposals to address
 856 potentially unsatisfactory design points, and discuss their implementation drawbacks.

858 5.1 Transitivity and uniqueness of the bottom type

859 In Julia, types `Tuple{Vector{T}, Union{ }}` where T and `Tuple{Vector{T} where T, Union{ }}` are equiv-
 860 alent, and the following judgments hold:

862 $\text{Tuple}\{\text{Vector}\{T\} \text{ where } T, \text{Union}\{ \}\} <: \text{Tuple}\{\text{Vector}\{S\}, \text{Vector}\{S\}\} \text{ where } S$
 863 $(\text{Tuple}\{\text{Vector}\{S\}, \text{Vector}\{S\}\} \text{ where } S) <: \text{Tuple}\{Q, Q\} \text{ where } Q$ (4)

864 However, in Julia 0.6.2 their transitive closure does not, as we have:

866 $\text{Tuple}\{\text{Vector}\{T\} \text{ where } T, \text{Union}\{ \}\} \not\prec: \text{Tuple}\{Q, Q\} \text{ where } Q$ (5)

867 The judgment (4) holds because for all instantiations of S the type `Vector{S}` is concrete and `Union{ }`
 868 is a subtype of any type, while in judgment (5) the diagonal rule prevents subtyping because the
 869 type `Vector{T} where T` is abstract. This has been fixed in Julia 0.7dev by making (4) false. We argue
 870 that this solution is unsatisfactory. It is odd to have

872 $(\text{Tuple}\{\text{Vector}\{S\}, \text{Vector}\{S\}\} \text{ where } S) \not\prec: \text{Tuple}\{Q, Q\} \text{ where } Q,$

873 while any instantiation of S with a concrete type, e.g. `Int`, leads to valid subtyping, e.g.

875 $\text{Tuple}\{\text{Vector}\{Int\}, \text{Vector}\{Int\}\} <: \text{Tuple}\{Q, Q\} \text{ where } Q.$

876 We propose an alternative design. The type `Tuple{Union{ }}` (or, more generally, any tuple type
 877 containing `Union{ }` as one of its components) is not inhabited by any value, and dispatch-wise
 878 it behaves as `Union{ }`. However, neither Julia 0.6.2 nor our formalization can prove it equivalent
 879 to `Union{ }` because the judgment $\text{Tuple}\{\text{Union}\{ \}\} <: \text{Union}\{ \}$ is not derivable: following Julia 0.6.2
 880 semantics, the `lift_union` function does not lift empty unions out of tuples. Extending `lift_union`
 881 to lift empty unions, thus rewriting types such as `Tuple{Union{ }}` into `Union{ }`, is straightforward;

882

883 the resulting subtype relation is not affected by the transitivity problem described above. We have
 884 modified our reference implementation along these lines. Testing over the real-world workloads
 885 does not highlight differences with the standard subtype relation, suggesting that this change
 886 does not impact the programming practice. However, this alternative design has implementation
 887 drawbacks. Assuming that a `Tuple{t}` type in a program where t is unknown yields a 1-element
 888 tuple type becomes incorrect, making dataflow analyses more complex. Also, uniqueness of the
 889 bottom type is lost, and testing if a type is bottom (a test that occurs surprisingly often in Julia
 890 code-base) becomes slower. These tradeoffs are being investigated by Julia developers.

891

892 5.2 Ahead-of-time normalization and static detection of diagonal variables

893 We have seen in Sec. 3.1 that the interactions between subtyping invariant constructors, union
 894 types, and existential types make dynamic lifting and unlifting of union and existential types with
 895 respect to tuples necessary to specify a complete subtype relation. It is, however, interesting to
 896 explore if an ahead-of-time normalization phase has any benefit.

897 Since lifting unions across invariant constructors is unsound, our *normalization* phase rewrites
 898 tuples of unions into unions of tuples, pulling wheres out of tuples and pushing wheres inside
 899 unions, both at top-level and inside all invariant constructors. Additionally, it rewrites tuples with
 900 a bottom element into the bottom type, as suggested in Sec. 5.1.

901 Search over normalized types does not require the rule `TUPLE_LIFT_UNION` anymore, but rule `TU-`
 902 `TUPLE_UNLIFT_UNION` is still crucial (even more so) for completeness. Despite this, ahead-of-time
 903 normalization may have benefits. At the end of Sec. 3.1 we explained that in a type such as

904
$$\text{Tuple}\{\text{Union}\{\text{Bool}, T\}, T\} \text{ where } T$$

 905

906 it is not possible to determine statically if the variable T is diagonal as this depends on the type on
 907 the left-hand side of the judgment and on the derivation. However, if this type is normalized into
 908 the equivalent type

909
$$\text{Union}\{\text{Tuple}\{\text{Bool}, T_1\} \text{ where } T_1, \text{Tuple}\{T_2, T_2\} \text{ where } T_2\},$$

 910

911 the confusion about the variable T is removed: the variable T_2 is diagonal, while the variable T_1 is not.
 912 It is then straightforward to write a function *mark_diagonal* that marks variables over normalized
 913 types as diagonal whenever they occur more than once in a covariant context and never in an
 914 invariant context. In the general case, static marking can only *approximate* the dynamic counting
 915 of occurrences, for variables that appear in bounds get expanded only while building a complete
 916 derivation. However, the static counting has some nice properties.

917 First, a syntactic separation between diagonal and non-diagonal variables avoids subtle interac-
 918 tions of unions and type variables. Both Julia 0.6.2 and Julia 0.7dev, as well as our formalisation,
 919 state that the judgment below is correct:

920
$$\langle \text{Tuple}\{Q, \text{Int}\} \text{ where } Q <: \text{Union}\{\text{Bool}, S\} \text{ where } S \rangle <: \text{Tuple}\{\text{Union}\{\text{Bool}, T\}, T\} \text{ where } T.$$

 921

922 We argue that this judgment should not hold. The variable T , when matched with S and `Int`, should
 923 be considered diagonal. This becomes explicit if the right-hand side is normalized into the type
 924 $\text{Union}\{\text{Tuple}\{\text{Bool}, T_1\} \text{ where } T_1, \text{Tuple}\{T_2, T_2\} \text{ where } T_2\}$; Julia 0.6.2 and Julia 0.7dev confirm that

925
$$\langle \text{Tuple}\{Q, \text{Int}\} \text{ where } Q <: \text{Union}\{\text{Bool}, S\} \text{ where } S \rangle \quad \not<:$$

 926
$$\text{Union}\{\text{Tuple}\{\text{Bool}, T_1\} \text{ where } T_1, \text{Tuple}\{T_2, T_2\} \text{ where } T_2\}.$$

 927

928 Additionally, it might be argued that static marking of diagonal variables makes subtyping more
 929 *predictable*. As we briefly mentioned, to address the transitivity problem of Sec. 4.4, Julia 0.7dev
 930 identifies covariant and invariant occurrences of each variable also on types that appear on the

931

left-hand side of the judgment. Diagonal variables are not allowed to have non-diagonal variables as lower-bounds. With this in mind, consider the judgments below:

$$\text{(Tuple}\{T, S, S\} \text{ where } T <: \text{Ref}\{S\} \text{ where } S) <: \text{Tuple}\{\text{Any}, Q, Q\} \text{ where } Q \quad (6)$$

$$\text{(Tuple}\{T, S, S\} \text{ where } T <: \text{Ref}\{S\} \text{ where } S) \not<: \text{Tuple}\{\text{Ref}\{R\}, Q, Q\} \text{ where } R \text{ where } Q \quad (7)$$

Both judgments exhibit the same type on the left-hand side, matched against different types. In the former the variable S is diagonal: it occurs twice in covariant position, and, since the subtype derivation does not use the constraint $T <: \text{Ref}\{S\}$, its invariant occurrence in $\text{Ref}\{S\}$ is not counted. In the latter the derivation does use the information on the upper bound of T ; the variable S is no longer diagonal, and Q (which is diagonal) cannot be instantiated with S . Programmers implement the methods of a function one at a time, possibly in different files; the lack of *predictability* of which variables are diagonal might lead to confusing dispatch behaviors. Static diagonal variable marking identifies the variable S as non-diagonal in both judgments: judgment (6) no longer holds and the behavior of the type on the left-hand-side thus becomes consistent with (7). We have modified our reference implementation of the subtyping algorithm to support ahead-of-time normalization and static marking of diagonal variables. This version passes the Julia regression test suites, and comparing the two algorithms over the real-world workload highlights only 41 differences. Of these, 35 of them reduce to 3 cases in which our modified algorithm cannot prove judgments due to our incomplete implementation of the *unlift_union* function. Since Julia relies on a different search mechanism, it would not be affected by these. The remaining 6 are interesting: the *typeof* function sometimes behaves differently on a normalized type, affecting the subtyping of `Type` types. For instance, *typeof*(`Tuple{Ref{T}} where T`) returns `DataType`, but if the argument is normalized, *typeof*(`Tuple{Ref{T}}` where T) returns `UnionAll`.

Summarizing: in theory subtyping based on ahead-of-time normalization and static marking of diagonal variables might constitute an alternative design of Julia subtyping. In practice the tradeoff is less clear: normalization may result in explosion of types' size, which is unacceptable for the actual implementation of Julia. It is an open question, and a future research project, to determine if a space-efficient subtyping algorithm can implement this revised relation.

5.3 Intersection types and symmetric UNION_LEFT

We have seen in Sec. 3.2 that rule `R_LEFT` allows arguably incorrect judgments to be derived because it propagates the new upper bound for the right variable, instead of the intersection of the old and new upper bounds. An hypothetical correct rule appears in the inset. Without native support for intersection types, computing the intersection of two arbitrary Julia types is a hard problem in itself. Julia code-base already includes a complex algorithm that computes an approximation of intersection of two types, which is used internally to compute dataflow informations, but this algorithm is too slow (and bug-ridden) to be integrated in the subtype algorithm. It should be noted that our counterexample is artificial and is unlikely to appear in the programming practice (e.g. it did not appear in the subtype calls we logged on real-world workloads, and it was not reported before), so there is a tradeoff between the extra complexity added to the implementation and the benefit of a more correct relation. In reply to our call, Julia developers have introduced a `simple_meet` function which computes intersections for simple cases; our counterexample has not been addressed yet.

An ambitious redesign of Julia's *internal* type language, that would include native intersection types, has been considered, but no steps have been undertaken in this direction. This is an ambitious research project on its own.

As an aside note, we remark that support for intersection types would enable the alternative

$$\frac{\text{search}(T, E) = R_T^u \quad E \vdash l <: t \vdash E'}{E \vdash T <: t \vdash \text{upd}(R_T^{t \cap u}, E')}$$

981 formulation of rule `UNION_LEFT` in the inset.

982 The *merge* function returns an environment in

983 which, for each variable, the lower bound is

984 the union of all the lower bounds for the variable in $E_1..E_n$, and the upper bound is the intersection

985 of all the upper bounds for the variable in $E_1..E_n$. In this formulation the order of the types in

986 the list $t_1..t_n$ is obviously irrelevant for subtyping, a property non-trivial to prove in the current

987 formulation.

988

989

6 RELATED WORK

990 Surprisingly for a dynamic language, Julia’s subtype relation is defined over a rich grammar of
991 types, which often is the prerogative of statically-typed programming languages.

992 Languages with multimethods differ on whether parametric polymorphism is supported or not.

993 Most previous efforts focused on non-polymorphic types, such as Cecil [Chambers and Leavens

994 1994], Typed Clojure [Bonnaire-Sergeant et al. 2016], and MultiJava [Clifton et al. 2000]. Subtyping

995 is used to check that classes implement all of the required methods of their supertypes. The subtype

996 relations are defined over covariant tuples and discrete unions. Approaches that combine multimeth-

997 ods with parametric polymorphism are more involved. The earliest work, ML<: [Bourdoncle and

998 Merz 1997], extends ML with subtyping and multimethods and shows that the type system is sound

999 and decidable by showing that the constraint solving system that it uses to handle both subtyping

1000 and pattern matching is decidable. Following the ML polymorphism, types have only top-level

1001 quantifiers (for example, $\forall\alpha.\text{list}[\alpha]$ is allowed but not $\text{list}[\forall\alpha.\text{list}[\alpha]]$), with subtyping being

1002 defined over monotypes. Constraints on type variables partially model union types: for instance,

1003 the type $\forall\alpha.(\text{int}<:\alpha, \text{bool}<:\alpha).\alpha$ can be understood as the set union of `int` and `bool`. Due to the

1004 lack of nesting quantifiers, this does not equate to Julia’s union types.

1005 Universal polymorphism and parametric multimethods have been proposed in Mini-Cecil [Litvinov

1006 1998, 2003]. Similar to ML<:, universal types have only top-level quantifiers. Fortress [Allen

1007 et al. 2011], in addition, supports arrow types, and internally uses both universal and existential

1008 types, with top-level quantifiers. Mini-Cecil and Fortress both use a constraint generation strategy

1009 to resolve subtyping; they support union and intersection types but do not provide distributivity

1010 “in order to simplify constraints solving” [Litvinov 2003]. For Mini-Cecil typechecking is decidable.

1011 Fortress argued decidability based on [Castagna and Pierce 1994], though no proof is provided.

1012 Frisch et al. [2002, 2008] studies the semantic interpretation of subtyping with union, intersection,

1013 negation types, and function types. Types are interpreted as sets of values; base types have their own

1014 denotation, and all the operators on types correspond to the equivalent set theoretical operations.

1015 Subtyping is defined semantically as inclusion of the sets denoted by types. The main challenge

1016 of the approach is due to the function types. However, only one type operator, namely, reference

1017 types, is described as an extension. An important contribution was a sound and complete algorithm

1018 to decide semantic subtyping. The algorithm crucially relies on semantic properties of their domain,

1019 in particular that types can be rewritten in disjunctive-normal form. As we have seen, Julia does

1020 not fully embrace semantic subtyping, and due to interactions between union types, invariant

1021 constructors, and existential types, search is considerably more complex. [Castagna et al. 2015,

1022 2014] extended their system with parametric polymorphism: terms with type variables are first

1023 compiled away to a variable-free core language with a type-case construct. Similar to [Frisch et al.

1024 2002, 2008], their interpretation of types differs from Julia’s.

1025 Subtyping of union types in Julia builds on [Vouillon 2004], which proposes an algorithm to

1026 decide subtyping of union types in a statically typed language with functions and pairs but without

1027 union/pair distributivity. The same paper also considers an extension of the language with ML-style

1028 polymorphism and co-/contravariant type constructors, but not invariant ones.

1029

1030 Bounded existential types have been used to model Java wildcards [Cameron et al. 2008; Tate et al.
 1031 2011]: for instance, the wildcard type $\text{List}\langle ? \rangle$ can be represented as an existential type $\exists T. \text{List}\langle T \rangle$.
 1032 Wildcards are less expressive than where-types in Julia, because they cannot denote types such as
 1033 $\exists T. \text{List}\langle \text{List}\langle T \rangle \rangle$ while $\text{List}\langle \text{List}\langle ? \rangle \rangle$ corresponds to $\text{List}\langle \exists T. \text{List}\langle T \rangle \rangle$. Nevertheless, inexpressible
 1034 types may appear during typechecking, and therefore the formal models use the full grammar
 1035 of existential types, and so does subtyping. Java wildcards do not have to deal with structural type
 1036 constructors of Julia, such as union types, covariant tuples, and their distributivity. This allows for
 1037 simpler subtyping rules for existential types that rely on well-developed machinery of unification
 1038 and explicit substitution for type variables. Subtyping of wildcards with union and intersection
 1039 types are studied in [Smith and Cartwright 2008]. Though the paper mentions that a distributivity
 1040 rule is a desired extension of subtyping, the rule is omitted due to a “tedious normalization steps”
 1041 that would have been needed. As our experience shows, in presence of invariant constructors
 1042 normalization does not solve all the problems, and should be accompanied by “unlifting unions”
 1043 (recall the example $\text{Ref}\{\text{Union}\{\text{Tuple}\langle \text{Int} \rangle, \text{Tuple}\langle \text{Bool} \rangle\}\} <: \text{Ref}\{\text{Tuple}\langle T \rangle\}$ where T).

1044 In type systems with bounded existential types, as well as type systems with nominal subtyping
 1045 and variance, decidability of subtyping has been a major concern [Kennedy and Pierce 2007; Wehr
 1046 and Thiemann 2009]. By design, Julia lacks language features that are known to cause undecidability.
 1047 Firstly, unlike in traditional existential types [Pierce 1992], types such as $\exists T_{t_1}^{t_2}. T$ are instantaneously
 1048 rewritten into the upper bound of T by the frontend and do not appear in subtyping. Secondly, unlike
 1049 in Java, where subtyping has been proved undecidable [Grigore 2017], neither of the following
 1050 is allowed in Julia: recursive bounds on type variables (e.g. $\text{Ref}\{T\}$ where $T <: \text{Foo}\{T\}$) [Wehr and
 1051 Thiemann 2009], contravariant type constructors [Kennedy and Pierce 2007], existential types in
 1052 type definitions (e.g. $\text{struct Foo}\{T\} <: \text{Bar}\{S>:T$ where S) [Tate et al. 2011]. An unpublished
 1053 manuscript on decidability of type checking for Java wildcards [Mazurak and Zdancewic 2005],
 1054 while failing on modeling of a particular language feature [Zdancewic 2006], develops a formal
 1055 machinery for updating environments which resembles ours.

1056 A simplified kernel of the subtype relation was documented in Bezanson PhD thesis [Bezanson
 1057 2015], together with a minimal Julia implementation of the algorithm. This effort introduced some
 1058 ideas: for example, it sketches the strategy to update the bounds on type variables. But it was
 1059 neither complete nor correct, and reflected an older version of Julia’s type system. In particular,
 1060 it ignored the subtle rules that govern propagation of the constraints, and the exist/forall quantifier
 1061 inversion; it did not support user-defined parametric types or the diagonal rule.

1062

1063

1064

1065 7 CONCLUSION

1065 We have provided a specification of the subtype relation of the latest release of the Julia language.
 1066 In many systems answering the question whether $t_1 <: t_2$ is an easy part of the development. It
 1067 was certainly not our expectation, approaching Julia, that reverse engineering and formalizing the
 1068 subtype relation would prove to be the challenge on which we would spend our time and energy.
 1069 As we kept uncovering layers of complexity, the question whether all of this was warranted kept us
 1070 looking for ways to simplify the subtype relation. We did not find any major feature that could be
 1071 dropped. Indeed, we carried out static and dynamic analysis of the core libraries, as well as of 100
 1072 popular Julia packages, and found that both language implementors and end-users avail themselves
 1073 of the full power of the type sublanguage. The usage of the advanced features of Julia type system
 1074 is widespread in both groups. Our formalization enables the study of the metatheory of the subtype
 1075 relation; the system is intricate and even simple properties require complex reasoning. Additionally,
 1076 it is not *a priori* clear if it is possible to define the subtype relation in a more declarative style.
 1077 Arguably these would be research contributions in their own right.

1078

1079

1079 As a separate contribution, to validate our formalization and to explore the implementation
1080 challenges of the subtype algorithm, we have provided a proof-of-concept implementation that
1081 closely mirrors our specification and relies on a simple search strategy. Our experimental results
1082 show that this implementation captures the real subtype relation, and is useful in identifying issues
1083 in the Julia implementation.

1084 At the outset, our ambition was to simplify the subtype relation, as we felt that it had too many
1085 corner cases. While our specification is more elegant and easier to understand than C code, we
1086 were not able to identify major features of the algorithm that could be eliminated. The design of
1087 the relation is motivated by seemingly reasonable end-user expectations and requirements.
1088

1089 REFERENCES

- 1090 Eric Allen, Justin Hilburn, Scott Kilpatrick, Victor Luchangco, Sukeyoung Ryu, David Chase, and Guy Steele. 2011. Type
1091 Checking Modular Multiple Dispatch with Parametric Polymorphism and Multiple Inheritance. In *Conference on Object*
1092 *Oriented Programming Systems, Languages and Applications (OOPSLA)*. <https://doi.org/10.1145/2048066.2048140>
- 1093 Jeff Bezanson. 2015. *Abstraction in technical computing*. Ph.D. Dissertation. Massachusetts Institute of Technology.
1094 <http://hdl.handle.net/1721.1/99811>
- 1095 Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. 2017. Julia: A Fresh Approach to Numerical Computing.
1096 *SIAM Rev.* 59, 1 (2017). <https://doi.org/10.1137/141000671>
- 1097 Ambrose Bonnaire-Sergeant, Rowan Davies, and Sam Tobin-Hochstadt. 2016. Practical optional types for Clojure. In
1098 *European Symposium on Programming (ESOP)*. https://doi.org/10.1007/978-3-662-49498-1_4
- 1099 François Bourdoncle and Stephan Merz. 1997. Type Checking Higher-order Polymorphic Multi-methods. In *Symposium on*
1100 *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/263699.263743>
- 1101 Nicholas Cameron, Sophia Drossopoulou, and Erik Ernst. 2008. A Model for Java with Wildcards. In *European Conference on*
1102 *Object-Oriented Programming (ECOOP)*. https://doi.org/10.1007/978-3-540-70592-5_2
- 1103 Giuseppe Castagna, Kim Nguyen, Zhiwu Xu, and Pietro Abate. 2015. Polymorphic Functions with Set-Theoretic Types:
1104 Part 2: Local Type Inference and Type Reconstruction. In *Symposium on Principles of Programming Languages (POPL)*.
1105 <https://doi.org/10.1145/2676726.2676991>
- 1106 Giuseppe Castagna, Kim Nguyen, Zhiwu Xu, Hyeonseung Im, Sergueï Lenglet, and Luca Padovani. 2014. Polymorphic
1107 Functions with Set-theoretic Types: Part 1: Syntax, Semantics, and Evaluation. In *Symposium on Principles of Programming*
1108 *Languages (POPL)*. <https://doi.org/10.1145/2535838.2535840>
- 1109 Giuseppe Castagna and Benjamin C. Pierce. 1994. Decidable Bounded Quantification. In *Symposium on Principles of*
1110 *Programming Languages (POPL)*. <https://doi.org/10.1145/174675.177844>
- 1111 Craig Chambers and Gary T. Leavens. 1994. Typechecking and Modules for Multi-methods. In *Conference on Object-oriented*
1112 *Programming Systems, Languages and Applications (OOPSLA)*. <https://doi.org/10.1145/191080.191083>
- 1113 John Chambers. 2014. Object-Oriented Programming, Functional Programming and R. *Statist. Sci.* 2 (2014). Issue 29.
1114 <https://doi.org/10.1214/13-ST5452>
- 1115 Koen Claessen and John Hughes. 2000. QuickCheck: A Lightweight Tool for Random Testing of Haskell Programs. In
1116 *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00)*. ACM, New York,
1117 NY, USA, 268–279. <https://doi.org/10.1145/351240.351266>
- 1118 Curtis Clifton, Gary T. Leavens, Craig Chambers, and Todd Millstein. 2000. Multijava: Modular Open Classes and Symmetric
1119 Multiple Dispatch for Java. In *Conference on Object-oriented Programming, Systems, Languages, and Applications (OOPSLA)*.
1120 <https://doi.org/10.1145/353171.353181>
- 1121 Linda DeMichiel and Richard Gabriel. 1987. The Common Lisp Object System: An Overview. In *European Conference on*
1122 *Object-Oriented Programming (ECOOP)*. https://doi.org/10.1007/3-540-47891-4_15
- 1123 Jonas Duregård, Patrik Jansson, and Meng Wang. 2012. Feat: Functional Enumeration of Algebraic Types. In *Proceedings of*
1124 *the 2012 Haskell Symposium (Haskell '12)*. ACM, New York, NY, USA, 61–72. <https://doi.org/10.1145/2364506.2364515>
- 1125 Alain Frisch, Giuseppe Castagna, and Véronique Benzaken. 2002. Semantic Subtyping. In *Symposium on Logic in Computer*
1126 *Science (LICS)*. <https://doi.org/10.1109/LICS.2002.1029823>
- 1127 Alain Frisch, Giuseppe Castagna, and Véronique Benzaken. 2008. Semantic subtyping: Dealing set-theoretically with
1128 function, union, intersection, and negation types. *J. ACM* 55, 4 (2008). <https://doi.org/10.1145/1391289.1391293>
- 1129 Radu Grigore. 2017. Java Generics Are Turing Complete. In *Symposium on Principles of Programming Languages (POPL)*.
1130 <https://doi.org/10.1145/3009837.3009871>
- 1131 Andrew Kennedy and Benjamin C. Pierce. 2007. On Decidability of Nominal Subtyping with Variance. In *Workshop on*
1132 *Foundations and Developments of Object-Oriented Languages (FOOL/WOOD)*. [https://www.microsoft.com/en-us/research/
1133 publication/on-decidability-of-nominal-subtyping-with-variance/](https://www.microsoft.com/en-us/research/publication/on-decidability-of-nominal-subtyping-with-variance/)

- 1128 Vassily Litvinov. 1998. Constraint-based Polymorphism in Cecil: Towards a Practical and Static Type System. In *Addendum*
 1129 *to the Conference on Object-oriented Programming, Systems, Languages, and Applications*. [https://doi.org/10.1145/346852.](https://doi.org/10.1145/346852.346948)
 1130 346948
- 1131 Vassily Litvinov. 2003. *Constraint-Bounded Polymorphism: an Expressive and Practical Type System for Object-Oriented*
 1132 *Languages*. Ph.D. Dissertation. University of Washington.
- 1133 Karl Mazurak and Steve Zdancewic. 2005. Type Inference for Java 5: Wildcards, F-Bounds, and Undecidability. (2005).
 1134 <https://pdfs.semanticscholar.org/a73a/aace3ecafb9f8f4f627705647115c29ef63e.pdf> unpublished.
- 1135 Benjamin C. Pierce. 1992. Bounded Quantification is Undecidable. In *Symposium on Principles of Programming Languages*
 1136 *(POPL)*. <https://doi.org/10.1145/143165.143228>
- 1137 Allison Randal, Dan Sugalski, and Leopold Toetsch. 2003. *Perl 6 and Parrot Essentials*. O'Reilly.
- 1138 Daniel Smith and Robert Cartwright. 2008. Java Type Inference is Broken: Can We Fix It?. In *Conference on Object-oriented*
 1139 *Programming Systems, Languages and Applications (OOPSLA)*. <https://doi.org/10.1145/1449764.1449804>
- 1140 Ross Tate, Alan Leung, and Sorin Lerner. 2011. Taming Wildcards in Java's Type System. In *Conference on Programming*
 1141 *Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/1993498.1993570>
- 1142 The Julia Language. 2018. Manual: Diagonal Types. (2018). Retrieved 2018-07-24 from [https://docs.julialang.org/en/v0.6.1/](https://docs.julialang.org/en/v0.6.1/devdocs/types/#Diagonal-types-1)
 1143 [devdocs/types/#Diagonal-types-1](https://docs.julialang.org/en/v0.6.1/devdocs/types/#Diagonal-types-1)
- 1144 Jerome Vouillon. 2004. Subtyping Union Types. In *Computer Science Logic (CSL)*. [https://doi.org/10.1007/978-3-540-30124-0_](https://doi.org/10.1007/978-3-540-30124-0_32)
 1145 32
- 1146 Stefan Wehr and Peter Thiemann. 2009. On the Decidability of Subtyping with Bounded Existential Types. In *Programming*
 1147 *Languages and Systems (ESOP)*.
- 1148 Francesco Zappa Nardelli, Julia Belyakova, Artem Pelenitsyn, Benjamin Chung, Jeff Bezanson, and Jan Vitek. 2018. Julia
 1149 Subtyping: a Rational Reconstruction — Project Web-Page. (2018). Retrieved 2018-07-24 from [https://www.di.ens.fr/](https://www.di.ens.fr/~zappa/projects/lambdajulia/)
 1150 [~zappa/projects/lambdajulia/](https://www.di.ens.fr/~zappa/projects/lambdajulia/)
- 1151 Steve Zdancewic. 2006. A Note on "Type Inference for Java 5". (2006). [https://web.archive.org/web/20060920024504/http://](https://web.archive.org/web/20060920024504/http://www.cis.upenn.edu/~stevez/note.html)
 1152 [www.cis.upenn.edu/~stevez/note.html](https://web.archive.org/web/20060920024504/http://www.cis.upenn.edu/~stevez/note.html)

1153 A THE *typeof*(*t*) FUNCTION

1154 Julia's *typeof* function returns the concrete type of each value. Since types are themselves values, it
 1155 is legitimate to invoke *typeof* on them, and the types `DataType`, `Union`, and `UnionAll` play the role of
 1156 kinds. Indeed, the auxiliary function *is_kind*(*t*) returns true if *t* is `DataType`, or `Union`, or `UnionAll`.
 1157 Since the *typeof* function plays a role in the definition of subtyping, we provide its formalization
 1158 in Fig. 4. We use an environment *G* to store the type variables in scope. We write *typeof*(*t*, *G*) as a
 1159 shorthand for there exists *t'* such that *typeof*(*t*, *G*) = *t'*.

1160 We have seen that Julia's frontend implicitly performs several simplifications when processing
 1161 types; for instance, `⊔ where T` where `T` is replaced by `Any`, or `Union{Int, Int}` is replaced by `Int`; these simpli-
 1162 fications must be taken explicitly into account when formalizing the *typeof* relation. The auxiliary
 1163 function *simplify*(*t*) performs the following simplification steps:

- 1164 • simplify trivial where constructs, e.g.: replace `⊔ where T <: t2` by `t2` and replace `t where T` by *t*
 1165 whenever `T ∉ fv(t)`;
- 1166 • remove redundant `Union` types, e.g.: replace `Union{t}` by *t*;
- 1167 • remove duplicate and obviously redundant types from `Union{t1, ..., tn}` types.

1168 Given a list of types *t*₁, ..., *t*_{*n*}, a type *t*_{*i*} is obviously redundant whenever there exists a type *t*_{*j*} which is
 1169 its supertype given an empty variable environment. These simplifications are guided by pragmatic
 1170 considerations rather than semantic issues. As such they tend to vary between Julia versions,
 1171 and we do not explicitly formalize them; our reference implementation mimics the simplification
 1172 behavior of Julia 0.6.2, apart for the issue described in Sec. 4.4.

1173 The function *typeof*(*t*) returns `Union` if the type *t*, after simplification or instantiation of trailing
 1174 where constructs, is a union. The case for `UnionAll` is similar, except that trailing where constructs
 1175 and the instantiation of user defined parametric types must be taken into account. In all other cases,
 1176 a type has the `DataType` kind.

```

1177  typeof(a) = typeof(simplify(a), ∅)
1178
1179  typeof(t, G) = DataType  if  is_kind(t)
1180                          or  t = Any
1181                          or  t = Tuple{t1, ..., tn} and ∀i, typeof(ti, G)
1182                          or  t = T and T ∈ G
1183                          or  t = (t where t1 <: T <: t2){t'} and typeof(t', G)
1184                              and typeof(t1, G) and typeof(t2, G) and t1 <: t' <: t2
1185                              and typeof(t[t'/T], G) = DataType
1186                          or  t = name and attr name{ } <: t' ∈ tds
1187                          or  t = name{t1, ..., tn} and ∀i, typeof(ti, G)
1188                              and attr name{t'1 <: T1 <: t'1', ..., t'n <: Tn <: t'n'} <: t''' ∈ tds
1189                              and ∀i, fV(t'i, ti, t'i') = ∅ ⇒ t'i <: ti <: t'i'
1190
1191  typeof(t, G) = Union  if  t = Union{t1, ..., tn} and ∀i, typeof(ti, G)
1192                          or  t = (t where t1 <: T <: t2){t'} and typeof(t', G)
1193                              and typeof(t1, G) and typeof(t2, G) and t1 <: t' <: t2
1194                              and typeof(t[t'/T], G) = Union
1195
1196  typeof(t, G) = UnionAll  if  t = t1 where T and typeof(t1, T; G)
1197                          or  t = (t where t1 <: T <: t2){t'} and typeof(t', G)
1198                              and typeof(t1, G) and typeof(t2, G) and t1 <: t' <: t2
1199                              and typeof(t[t'/T], G) = UnionAll
1200                          or  t = name and attr name{T1, ..., Tn} ∈ tds
1201                          or  t = name{t1, ..., tn} and ∀i, typeof(ti, G)
1202                              and name{t'1 <: T1 <: t'1', ..., t'n <: Tn <: t'n', .. } <: t''' ∈ tds
1203                              and ∀i, fV(t'i, ti, t'i') = ∅ ⇒ t'i <: ti <: t'i'
1204
1205  typeof(v, G) = ...      return the type tag of the value v

```

Fig. 4. The *typeof*(*t*) function.

The *typeof* function additionally checks that a type is well-formed, and in particular that, for all type variable instantiations, all actual types respect the bounds of binding variable. This check cannot be performed if the bounds or the variable itself have free variables, and in some cases Julia allows unsatisfiable bounds to be defined. For instance, the type `Foo{S}` where `S >: String` is accepted even assuming that `Foo` is declared as abstract type `Foo{T <: Int}` end. Since there is no type that is subtype of `Int` and supertype of `String`, this type denotes an empty set of values. Well-formedness of type-definitions can be checked along similar lines, keeping in mind that all parameters bind in the supertype, and each parameter binds in all bounds that follow.

B ISSUES REPORTED TO JULIA DEVELOPERS

The complete list of the issues we reported to the Julia bug tracker since starting this project follows. For each report, the number (and active link) in parentheses is the issue id's in Julia's github database. We distinguish between bug reports that have been fixed, bug reports that have been acknowledged and for which a solution is currently being investigated, and other design improvement proposals.

B.1 Fixed Bugs

(1) Reflexivity and transitivity broken due to buggy diagonal rule (#24166)

Flaws in the implementation of the diagonal rule check lead invalidate expected properties of the subtype relation, as discussed in Sec. 4.4. These flaws are observable in Julia 0.6.2 but have been fixed in the development version.

(2) Propagation of constraints when subtype unions (#26654)

The order of types inside a Union constructor should not affect the subtype relation (a property we call symmetry of Union). The subtype algorithm however traverses the types inside a Union constructor in a precise order. Incorrect propagation of constraints during subtyping made subtyping dependent on the order of types inside a Union constructor, as highlighted by the Julia 0.6.2 behavior below:

```
julia> Ref{Union{Int, Ref{Number}}} <: Ref{Union{Ref{T}, T}} where T
true
```

This issue was found by our fuzz tester. It has been fixed in the development version.

(3) Union{Ref{T}, Ref{T}} and Ref{T} behave differently (#26180)

This bug was introduced after the Julia 0.6.2 release:

```
julia> Ref{Union{Ref{Int}, Ref{Number}}} <: Ref{Ref{T}} where T
false

julia> Ref{Union{Ref{Int}, Ref{Number}}} <: Ref{Union{Ref{T}, Ref{T}}} where T
true
```

The second check should return `false`, as the first one, because the two types on the right-hand side are equivalent. This bug was found by our fuzz tester. It has been fixed in the development version (with the same commit that fixes the previous bug report).

B.2 Open Issues

(1) Missing intersection types (#26131)

```
julia> Vector{Vector{Number}} <:
  Vector{Union{Vector{Number}, Vector{S}}} where S<:Integer
true
```

As discussed in Sec. 3.2, this query should return `false` because `Vector{S}` is not a subtype of `Vector{Number}` when `Vector{S}<:Integer`. To correctly derive similar judgments, the subtype algorithm must be able to compute the intersection of types. This is a hard problem in itself. As a temporary band-aid, in reply to our call, Julia developers have introduced a `simple_meet` function which computes intersections for simple cases. The current implementation is still too weak to handle this particular case. The fact that not computing the intersection of the upper bounds in rule `R_LEFT` might be source of problems in presence of union types was suggested by an anonymous reviewer; our example is built on top of reviewer's remark.

(2) Stack overflows / Loops in subtype.c subtype_unionall

Unexpected looping inside the subtype algorithm, or large computer-generated types, can make Julia subtype algorithm to exceed the space allocated for the recursion stack. We reported this issue on a large computer-generated type (#26065). We discovered later that other reports address a similar issue; some are referenced in the ticket above, but some are more recent (#26487).

1275 (3) **Inconsistent constraints are ignored (#24179)**

1276 Frontend simplification rewrites types of the form \top where $lb <: T <: ub$ into the upper bound
 1277 ub , without checking first if the user-specified bounds are inconsistent, as in:

1278

```
1279 julia>  $\top$  where  $String <: T <: Signed$   

1280 Signed
```

1281

This may lead to unexpected results in subtype queries, and the type above is not considered
 1282 equivalent to the `Union{}`. Julia developers agree this behavior is incorrect.

1283 (4) **Diagonality is ignored and constraints are missing when matching with union (#26716)**

1284 Both Julia 0.6.2 and 0.7-dev incorrectly return `true` on these judgments (on the left types are
 1285 equivalent, on the right it is the same type):

1286

```
1287 julia>  $(Tuple\{Q, Bool\} \text{ where } Q <: Union\{Int, P\} \text{ where } P) <: Tuple\{Union\{T, Int\}, T\} \text{ where } T$   

1288 true  

1289 julia>  $(Tuple\{Union\{Int, P\}, Bool\} \text{ where } P) <: Tuple\{Union\{T, Int\}, T\} \text{ where } T$   

1290 true  

1291 julia>  $(Union\{Tuple\{Int, Bool\}, Tuple\{P, Bool\}\} \text{ where } P) <: Tuple\{Union\{T, Int\}, T\} \text{ where } T$   

1292 true
```

1293

1294 The correct answer is `false` because the variable T should be considered diagonal and gets
 1295 matched both with P and `Bool`, and as such it cannot be concrete. This is confirmed by rewriting
 1296 into an equivalent type by the `lift_union` function, thus making the diagonal variable explicit.
 1297 In this case Julia returns the correct answer.

1298

1299 **B.3 Proposals we made**1300 (1) **Interaction of diagonal rule and lower bounds (#26453)**

1301 Whenever the programmer specifies explicitly a lower bound for a type-variable, as in
 1302 `Tuple{T, T} where T >: t`, it is not always easy to decide if T should be considered diagonal
 1303 or not. This depends on whether the lower bound, t , is concrete, but in general deciding
 1304 concreteness is hard and Julia implementation approximates it with an heuristic. We proposed
 1305 that the variables should be considered diagonal only if their concreteness is obvious. The
 1306 proposal was approved, implemented and merged into the master branch.

1307 (2) **Another approach to fix problem with concreteness of `Vector{T}` / transitivity (com-
 1308 ment #372746252).**

1309 A subtle interaction between the bottom type and the diagonal rule can break transitivity of
 1310 the subtype relation. We propose an alternative approach to fix the issue, as the solution to
 1311 the problem applied in Julia seems unsatisfactory.

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323