

Java Wildcards Meet Definition-Site Variance

John Altidor¹, Christoph Reichenbach¹, and Yannis Smaragdakis^{1,2}

¹ University of Massachusetts, Amherst

² University of Athens, Greece

Abstract. Variance is concerned with the interplay of parametric polymorphism (i.e., templates, generics) and subtyping. The study of variance gives answers to the question of when an instantiation of a generic class can be a subtype of another. In this work, we combine the mechanisms of use-site variance (as in Java) and definition-site variance (as in Scala and C#) in a single type system, based on Java. This allows maximum flexibility in both the specification and use of generic types, thus increasing the reusability of code. Our VarJ calculus achieves a safe synergy of def-site and use-site variance, while supporting the full complexities of the Java realization of variance, including F-bounded polymorphism and wildcard capture. We show that the interaction of these features with definition-site variance is non-trivial and offer a full proof of soundness—the first in the literature for an approach combining variance mechanisms.

1 Introduction

Consider a generic type $C\langle X \rangle$. When is a type-instantiation $C\langle Exp1 \rangle$ a subtype of another type instantiation $C\langle Exp2 \rangle$? This is the question that *variance* mechanisms in modern programming languages try to answer. Variance (specifically, subtype variance with respect to generic type parameters) is a key topic in language design because it develops the exact rules governing the interplay of the two major forms of polymorphism: parametric polymorphism (i.e., generics or templates) and subtype (inclusion) polymorphism.

Languages like C# and Scala support a type system with *definition-site variance*: at the point of defining the generic type $C\langle X \rangle$ we state its subtyping policy and the type system attempts to prove that our assertion is statically safe. For instance, a C# definition `class C<out X> ...` means that C is *covariant*: $C\langle S \rangle$ is a subtype of $C\langle T \rangle$ if S is a subtype of T . The type system's obligation is to ensure that type parameter X of C is used in the body of C in a way that guarantees type safety under this subtyping policy. For instance, X cannot appear as the argument type of a public method in C —a rule colloquially summarized as “the argument type of a method is a contravariant position”.

By contrast, the type system of Java employs the concept of *use-site variance* [11]: a class does not itself state its variance when it is defined. Uses of the class, however, can choose to specify that they are referring to a *covariant*, *contravariant*, or *bivariant* version of the class. For instance, a method `void`

`meth(C<? extends T> cx)` can accept arguments of type `C<T>` but also `C<S>` where `S` is a subtype of `T`. An object with type `C<? extends T>` may not offer the full functionality of a `C<T>` object: the type system ensures that the body of method `meth` employs only such a subset of the functionality of `C<T>` that would be safe to use on any `C<S>` object (again, with `S` a subtype of `T`). This can be viewed informally as automatically projecting class `C` and deriving per-use versions.

Each flavor of variance has its own advantages. Use-site variance is arguably a more advanced idea, yet it suffers from specific usability problems because it places the burden on the *user* of a generic type. (Although one should keep in mind that the users of one generic type are often the implementors of another.) Definition-site variance may be less expressive, but leaves the burden of specifying general interfaces with the *implementor* of a generic. A natural idea, therefore, is to combine the two flavors in the same language design and allow full freedom: For instance, when a type is naturally covariant, its definition site can state this property and relieve the user from any further obligation. Conversely, when the definition site does not offer options for fully general treatment of a generic, a sophisticated user can still provide fully general signatures.

This natural combination of the two kinds of variance is complicated especially by the interaction of use-site and definition-site annotations: for example, when does the declared variance of a type variable agree with occurrences of that variable in use-site annotations? We recently proposed a unifying framework for checking and inferring both definition and use-site variance [1]. That proposal was not accompanied by a language operational semantics however—its proof of soundness was expressed as a meta-theorem, i.e., under assumptions over what an imaginary language’s type system should be able to prove about sets of values. This meta-theorem was welcome as an intuition about why it makes sense to combine variances in a certain way, but did not establish a firm connection with any real programming language.

This paper investigates combining of definition- and use-site variance with all relevant language constructs using a new formal model, VarJ. VarJ applies novel ideas and integrates techniques from various formalisms: Java wildcards are a form of use-site variance that was proven sound with the TameFJ [4] calculus. VarJ directly extends TameFJ with definition-site variance. VarJ also employs ideas from our VarLang calculus [1], which introduces a variance transform operator. Finally, VarJ integrates definition-site subtyping rules from the work of Kennedy et al. [13,8]. The result is a language with highly expressive genericity. For instance, given an invariant class `List`, our type system allows defining a (definition-site) covariant class `ROStack` that returns covariant (intuitively: read-only) `Lists` of members:

```
class ROStack<+X> {
  X pop() { ... }
  List<? extends X> toList() { ... }
}
```

Note the simultaneous use of a definition-site variance annotation (+) on `ROStack`, as well as a use-site variance annotation on its `toList` method. The former is not safe without the latter.

Overall our work makes several contributions. At the high level:

- Compared to the type systems of Java, C#, or Scala, our combination of definition-site and use-site variance allows the programmer to pick the best tool for the job. Libraries can avoid offering different flavors of interfaces just to capture the notion of, e.g., “the covariant part of a list” vs. “the contravariant part of a list”. Conversely, users can often use purely-variant types more easily and with less visual clutter if the implementor of that type had the foresight to declare its variance.
- Our approach maintains other features of the Java type system, namely full support for wildcards, which are a mechanism richer than plain use-site variance (e.g., [10]) and allow uses directly inspired by existential types.
- We provide a framework for determining the variance of the various positions in which a type can occur. (For example, why is the upper bound of the type parameter of a polymorphic method a contravariant position?)

Also, at the technical level:

- We show how definition-site variance interacts in interesting ways with advanced typing features, such as existential types, F-bounded polymorphism, and wildcard capture. A naive application of our earlier work [1] to Java would result in unsoundness, as we show with concrete examples. (Our earlier approach avoided unsoundness when applied to actual Java code by making several over-conservative assumptions to completely eliminate any interaction between, e.g., definition-site variance and F-bounded polymorphism.)
- We clarify and extend the TameFJ formalism with definition-site variance. TameFJ is a thorough, highly-detailed formalism and extending it is far from a trivial undertaking. The result is that we offer the first full formal modeling and proof of soundness for a language combining definition- and use-site variance.

2 Background

We next offer a brief background on definition- and use-site variance as well as their relative advantages.

2.1 Definition-Site Variance

Languages supporting definition-site variance [14,9] typically require each type parameter to be declared with a variance annotation. For instance, Scala [14] requires the annotation `+` for covariant type parameters, `-` for contravariant type parameters, and invariance as default. A well-established set of rules can then be used to verify that the use of the type parameter in the generic¹ is consistent with the annotation.

In intuitive terms, we can understand the restrictions on the use of type parameters as applying to “positions”. Each typing position in a generic’s signature

¹ We refer to all generic types (e.g., classes, traits, interfaces) uniformly as “generics”.




Fig. 1. Standard variance lattice

has an associated variance. For instance, method return and exception types, supertypes, and upper bounds of class type parameters are covariant positions; method argument types and class type parameter lower bounds are contravariant positions; field types are both co- and contravariant occurrences, inducing invariance. Type checking the declared variance annotation of a type parameter requires determining the variance of the positions the type parameter occurs in. *The variance of all such positions should be at least the declared variance of the type parameter.* Figure 1 presents the variance lattice. Consider the following templates of Scala classes, where v_X , v_Y , and v_Z stand for variance annotations.

```

abstract class RList[vXX] { def get(i:Int):X }
abstract class WList[vYY] { def set(i:Int, y:Y):Unit }
abstract class IList[vZZ] { def setAndGet(i:Int, z:Z):Z }
    
```

The variance v_X is the declared definition-site variance for type variable X of the Scala class `RList`. If $v_X = +$, the `RList` class type checks because X does not occur in a contravariant position. If $v_Y = +$, the `WList` class does not type check because Y occurs in a contravariant position (second argument type in `set` method) but $v_Y = +$ implies Y should only occur in covariant position. `IList` type checks only if $v_Z = o$ because Z occurs in both a covariant and a contravariant position.

Intuitively, `RList` is a read-only list: it only supports retrieving objects. The return type of a method indicates this “retrieval” capability. Retrieving objects of type T can be safely thought of as retrieving objects of any supertype of T . Thus, a read-only list of T s (`RList[T]`) can always be safely thought of as a read-only list of some supertype of T s (`RList[S]`, where $T <: S$). This is the exact definition of covariant subtyping and the reason why a return type is a covariant position. Thus, `RList` is covariant in X . Similarly, `WList` is a write-only list, and is intuitively *contravariant*. Its definition supports this intuition: Objects of type T can be written to a write-only list of T s (`WList[S]`) and written to a write-only list of S s (`WList[S]`), where $T <: S$, because objects of type T are also objects of type S . Hence, a `WList[S]` can safely be thought of as a `WList[T]`, if $T <: S$.

The variance of type variables is *transformed* by the variance of the context the variables appear in. Covariant positions *preserve* the variance of types that appear in them, whereas contravariant positions *reverse* the variance of the types that appear in them. The “reverse” of covariance is contravariance, and vice versa. The “reverse” of invariance is itself. Thus, we can consider the occurrence

of a type parameter to be initially covariant. For instance, consider again the Scala classes above. In `RList`, `X` only appears as the return type of a method, which preserves the initial covariance of `X`, so `RList` is covariant in `X`. In `WList`, `Y` appears in a contravariant position, which reverses its initial covariance, to contravariance. Thus, `WList` is contravariant.

When a type parameter is used to instantiate a generic, its variance is further transformed by the declared definition-site variance of that generic. For example:

```
class SourceList[+Z] { def copyTo(to:WList[Z]):Unit }
```

Suppose the declared definition-site variance of `WList` (with respect to its single parameter) is contravariance. In `WList[Z]`, the initial covariance of `Z` is transformed by the definition-site variance of `WList` (contravariance). It is then transformed again by the contravariant method argument position. As a result, `Z` appears covariantly in this context, and `SourceList` is covariant in `Z`, as declared. Any variance transformed by invariance becomes invariance. Thus, if `Z` had been used to parameterize an invariant generic, its appearance would have been invariant.

We have so far neglected to discuss *bivariance*: `C<X>` is bivariant implies that `C<S><:C<T>` for any types `S` and `T`. Declaring a bivariant type parameter is not supported by the widely used definition-site variant languages, since designating a type parameter as bivariant typically means it does not appear in the generic's type signature. Nevertheless, the concept is useful in our more general treatment.

2.2 Use-Site Variance

An alternative approach to variance is *use-site variance* [11,18,4]. Instead of declaring the variance of `X` at its definition site, generics are assumed to be *invariant* in their type parameters. However, a type-instantiation of `C<X>` can be made co-, contra-, or bivariant using variance annotations.

For instance, using the Java wildcard syntax, `C<? extends T>` is a *covariant* instantiation of `C`, representing a type “C-of-some-subtype-of-T”. `C<? extends T>` is a supertype of all type-instantiations `C<S>`, or `C<? extends S>`, where `S<:T`. In exchange for such liberal subtyping rules, type `C<? extends T>` can only access fully those methods and fields of `C` in which `X` appears covariantly. (Other methods can be used only with type-neutral values, e.g., called with `null` instead of values of type `X`.) In determining this, use-site variance applies the same set of rules used in definition-site variance, with the additional condition that the upper bound of a wildcard is considered a covariant position, and the lower bound of a wildcard a contravariant position.

For example, consider an invariant generic class `List` that uses its type parameter in both covariant and contravariant positions:

```
class List<X> {
  ... // other members that don't affect variance
  void add(int i, X x) { ... } // requires a List<? super X>
  X get(int i) { ... } // requires a List<? extends X>
  int size() { ... } // requires a List<?>
}
```

`List<? extends T>`, only has access to method “`X get(int i)`”, but not method “`void add(int i, X x)`”. (More precisely, method `add` can only be called with `null` for its second argument.)

Similarly, `C<? super T>` is the contravariant version of `C`, and is a supertype of any `C<S>` and `C<? super S>`, where `T<:S`. Of course, `C<? super T>` has access only to methods and fields in which `X` appears contravariantly or not at all. (The `get` method returns `Object` for a `C<? super T>`.)

Use-site variance also allows the representation of the *bivariant* version of a generic. In Java, this is accomplished through the unbounded wildcard: `C<?>`. Using this notation, `C<S><:C<?>`, for any `S`. The bivariant type, however, only has full access to methods and fields in which the type parameter does not appear at all. In definition-site variance, these methods and fields would have to be factored out into a non-generic class.

2.3 A Comparison

Both approaches to variance have their merits and shortcomings. Definition-site variance enjoys a certain degree of conceptual simplicity: the generic type instantiation rules and subtyping relationships are clear. However, the class or interface designer must pay for such simplicity by splitting the definitions of data types into co-, contra-, and bivariant versions. This can be an unnatural exercise. For example, the data structures library for Scala contains immutable (covariant) and mutable (invariant) versions of almost every data type—and this is not even a complete factoring of the variants, since it does not include contravariant (write-only) versions of the data types.

The situation gets even more complex when a generic has more than one type parameter. In general, a generic with n type parameters needs 3^n (or 4^n if bivariance is allowed as an explicit annotation) interfaces to represent a complete variant factoring of its methods. Arguably, in practice, this is often not necessary.

Use-site variance, on the other hand, allows users of a generic to create co-, contra-, and bivariant versions of the generic on the fly. This flexibility allows class or interface designers to implement their data types in whatever way is natural. The users of these generics must pay the price, by carefully considering the correct use-site variance annotations, so that the type can be as general as possible. This might not seem very difficult for a simple instantiation such as `List<? extends Number>`. However, type signatures can very quickly become complicated. For instance, the following method signature is part of the Apache Commons-Collections Library:

```
Iterator<? extends Map.Entry<? extends K,V>>
    createEntrySetIterator(Iterator<? extends Map.Entry<? extends K,V>>)
```

3 Combining Definition- and Use-Site Variance

Our formalism supports combining definition- and use-site variance in the context of Java. We next see informally some of its main insights and complications.

3.1 Insights for Combining Variances

High-level elements of our approach are inherited from our earlier work [1], in which we presented rules for combining definition- and use-site variance in a type system. These rules are a significant generalization over what has been explored in the past (mainly in Scala) in two ways:²

- The variance of an arbitrary type expression with respect to a type variable is defined for all cases with the help of a “transform” operator, \otimes . The operator determines how variances compose. Given two generic types $A\langle X \rangle$ and $B\langle X \rangle$ with declared variances v_A and v_B for their parameters (i.e., declared as $A\langle v_A X \rangle$ and $B\langle v_B X \rangle$), we can compute the variance of type $A\langle B\langle X \rangle \rangle$ as $v = v_A \otimes v_B$. Variances take values from the lattice of Figure 1.

Figure 2 summarizes the behavior of the transform operator: invariance and bivariance override other variances, while covariance preserves and contravariance reverses variance (with invariance and bivariance being their own reverses, respectively). To sample why the definition of the transform operator makes sense, we derive one case relating the inputs and output. (Remaining cases are derived similarly.)

- **Case $+\otimes - = -$:** This means that type expression $C\langle E \rangle$ is contravariant with respect to type variable X when generic C is covariant in its type parameter and type expression E is contravariant in X . This is true because, for any two types T_1 and T_2 :

$$\begin{aligned}
 T_1 <: T_2 & \\
 \implies E[T_2/X] <: E[T_1/X] & \quad \text{(by contravariance of } E\text{)} \\
 \implies C\langle E[T_2/X] \rangle <: C\langle E[T_1/X] \rangle & \quad \text{(by covariance of } C\text{)} \\
 \implies C\langle E \rangle[T_2/X] <: C\langle E \rangle[T_1/X] &
 \end{aligned}$$

Hence, $C\langle E \rangle$ is contravariant with respect to X .

| | | | |
|---|------------------|------------------|------------------|
| Definition of variance transformation: \otimes | | | |
| $+\otimes + = +$ | $-\otimes + = -$ | $*\otimes + = *$ | $o\otimes + = o$ |
| $+\otimes - = -$ | $-\otimes - = +$ | $*\otimes - = *$ | $o\otimes - = o$ |
| $+\otimes * = *$ | $-\otimes * = *$ | $*\otimes * = *$ | $o\otimes * = *$ |
| $+\otimes o = o$ | $-\otimes o = o$ | $*\otimes o = *$ | $o\otimes o = o$ |

Fig. 2. Variance transform operator

- The interaction of use-site and definition-site variance is expressed as a join operation on the same variance lattice of Figure 1. In the VarLang calculus [1], types have the form $C\langle \overline{vT} \rangle$, where \overline{v} are use-site annotations. Considering the `List` generic from Section 2.2, for example, `List<+T>` passes type `T` to a *covariant*

² There is also a third way: our earlier framework allowed reasoning about unknown variances, represented as variables in recursive constraints, thus enabling variance *inference* instead of *checking*. This capability is not relevant here, however.

version of the `List`, where the `add` method was “removed”³ because it contains the type parameter in a contravariant position. If generic `C<X>` has definition-site variance v_1 with respect to `X`, then the type expression `C<v2X>` has variance $v_1 \sqcup v_2$ with respect to `X`. Consider a covariant class `RList`. If we request a contravariant instantiation, we end up with a bivariate type expression ($+ \sqcup - = *$). That is, a method “`void foo(RList<? super Animal> l) { ... }`” can really accept any `RList`: the method is guaranteed to never use its argument in a way that reveals anything about the type of element in the list. (In practice, this means that the method may only take the size of the list, or only treats its elements as being of the general type `Object`, etc.)

In practice, the ability to combine definition- and use-site variance gives the programmer maximum flexibility. The variance of a generic class does not need to be anticipated at its definition site. Consider the usual invariant `List` class (from Section 2.2). This `List` supports both reading and writing of data, hence it includes both kinds of methods, instead of being split into two types (as, for instance, is common in the Scala libraries). The methods that use `List` can still be made fully general, however, as long as they specify use-site annotations. Generally, allowing both kinds of variance in a single language ensures modularity: parts of the code can be made fully general regardless of how other code is defined.

At the same time, allowing definition-site variance eliminates much of the need for extensive use-site variance annotations and the risk of too-restricted types: purely variant types can be declared up-front without burdening the programmer at the use point.

Of course, combining definition- and use-site variance means more than just using each kind separately, when applicable. It also includes using one kind of annotation when reasoning about the other. For instance, consider the example of a read-only stack type that we briefly saw in the Introduction. The stack refers to the invariant class `List` (defined earlier):

```
class ROStack<+X> {
  X pop() { ... }
  List<? extends X> toList() { ... }
}
```

Note the use of a definition-site variance annotation (+) on `ROStack`, as well as a use-site variance annotation on its `toList` method. The example would not have been safe if the return type of `toList` were merely `List<X>`. With such an invariant use of type parameter `X`, we could use a (dynamic) `List<Dog>` as a (static) `List<Animal>` and thus dynamically add a `Cat` (which is fine to add to a `List<Animal>`) to a `List<Dog>`.

3.2 Realistic Complications

The main contribution of our present work consists of formalizing and proving sound the combination of definition- and use-site variance in the context of Java.

³ Again, method `add` can only be called with `null` for its second argument.

In order to do so, we need to reason about the interaction of definition-site variance with many complex language features, such as F-bounded polymorphism, polymorphic methods, bounds on type parameters, and existential-types (arising in the use of wildcards). This interaction is highly non-trivial, as we see in examples next.

One complication is that Java wildcards are not merely use-site variance, but also include mechanisms inspired by existential typing mechanisms. *Wildcard capture* is the process of passing an unknown type, hidden by a wildcard, as a type parameter in a method invocation. Consider the following method, which swaps the order the two elements at the top of a stack.

```
<E> void swapLastTwo(Stack<E> stack)
{ E elem1 = stack.pop(); E elem2 = stack.pop();
  stack.push(elem2); stack.push(elem1); }
```

Although a programmer may want to pass an object of type `Stack<?>` as a value argument to the `swapLastTwo` method, the type parameter to pass for `E` cannot be manually specified because the type hidden by `?` cannot be named by the programmer. However, passing a `Stack<?>` type checks because Java allows the compiler to automatically generate a name for the unknown type (capture conversion) and use this name in a method invocation. Our formalism has to model wildcard capture and its interaction with definition-site variance. This handling comprises some of the more significant changes of our formalism relative to TameFJ [4].

Another major complication concerns F-bounded polymorphism. Consider the following definition:

```
interface Trouble<P extends List<P>> extends Iterator<P> {}
```

The type `Trouble<P>` extends `Iterator<P>`, which is assumed in the example to be covariant (exporting a method “`P next()`”, per the Java library convention for iterators). It would stand to reason that `Trouble` is also covariant: an object of type `Trouble<P>` does precisely what an `Iterator<P>` object does, since it simply inherits methods. Consider, however, the type `Trouble<? super A>`. This is a contravariant use of a covariant generic. According to our approach for combining variances, this results in a bivariant type (due to the variance joining). For example, we can derive the following subtype relationship even though the types, `MyList` and `YourList`, are not subtype related.

```
Trouble<YourList> <: Trouble<Object> (by covariance assumption of Trouble)
  <: Trouble<? super Object>
  <: Trouble<? super MyList>
```

The problem, however, is that the bounds of type variables (`List<P>` in this case) are preserved in the existential type representing a use of `Trouble` with wildcards. This results in unsoundness because, in F-bounded polymorphism, the bound includes the hidden type, allowing its recovery and use. We can cause

a problem with the following code (`ArrayList` is a standard implementation of the usual Java `List` interface, both invariant types):⁴

```
class MyList extends ArrayList<MyList> {}
class YourList extends ArrayList<YourList> {
    int i = 0;
    public boolean add(YourList list)
    { System.out.println(list.i); return super.add(list); }
}
void foo(Trouble<? super MyList> itr) { itr.next().add(new MyList()); }
Trouble<YourList> preitr = ...;
foo(preitr);
```

Function `foo` type checks because `itr.next()` is guaranteed to return an unknown supertype, X , of `MyList` but also (due to the F-bound on `Trouble`) a subtype of `List<X>`. Thus, X has a method `add` (from `List`) which accepts X instances, and thus also accepts `MyList` instances (since X is a supertype of `MyList`).

The problem arises in the last line. If `Trouble<? super MyList>` were truly bi-variant (as a contravariant use of a covariant generic), then that line would type check, allowing the unsound addition of a `MyList` object to a list of `YourList`s. Thus, the joining of definition- and use-site variances needs to be carefully restricted in the presence of F-bounded polymorphism. We revisit the above example more formally in Section 5.

4 VarJ

We investigate extending Java with definition-site variance by developing the VarJ calculus. VarJ’s syntax found in Figure 3 is a slight extension of the TameFJ syntax. A program that type checks in TameFJ also type checks in VarJ. Significant differences are highlighted using shading. To improve readability, some syntactic categories are overloaded with multiple meta-variables. *Existential types* range over T, U, V , and S . *Type variables* range over X, Y , and Z . *Bounds* range over B and A . *Variances* range over v and w . The bottom type, \perp , is used only as a lower bound. We follow the syntactic conventions of TameFJ: all source-level type expressions are written as existential types, with an empty range for non-wildcard Java type uses and type variables written as $\exists\emptyset.X$; substitution is performed as usual except $[T/X]\exists\emptyset.X = T$; \star is a syntactic marker designating that a method type parameter (i.e., for a polymorphic method) should be inferred. Class type parameters (\overline{X}) now have definition-site variance annotations (\overline{v}) and lower bounds (\overline{B}_L). Method type variables now have lower bounds as well. The remainder of this section focuses on semantic differences between VarJ and TameFJ and new concepts from adding definition-site variance. Sections 4.1 and 4.2 formally present notions of the variance of type expression and the variance of a type position. Section 4.3 covers subtyping with definition-site and use-site variance in VarJ. Section 4.4 discusses the updates made to allow safe interaction between wildcard capture and variant types.

⁴ This example is originally due to Ross Tate.

Syntax:

| | |
|--|------------------------------|
| $e ::= x \mid e.f \mid e.\langle \overline{P} \rangle_m(\overline{e}) \mid \text{new } C\langle \overline{T} \rangle(\overline{e})$ | <i>expressions</i> |
| $s ::= \text{new } C\langle \overline{T} \rangle(\overline{s})$ | <i>values</i> |
| $v ::= + \mid - \mid * \mid o$ | <i>variance</i> |
| $Q ::= \text{class } C\langle \overline{v} X \rightarrow [\overline{B}_L - \overline{B}_U] \rangle \triangleleft N \{ \overline{T} f; \overline{M} \}$ | <i>class declarations</i> |
| $M ::= \langle X \rightarrow [\overline{B}_L - \overline{B}_U] \rangle T m(\overline{T} x) \{ \text{return } e; \}$ | <i>method declarations</i> |
| $N ::= C\langle \overline{T} \rangle$ | <i>non-variable types</i> |
| $R ::= N \mid X$ | <i>non-existential types</i> |
| $T ::= \exists \Delta.N \mid \exists \emptyset.X$ | <i>existential types</i> |
| $B ::= T \mid \perp$ | <i>type bounds</i> |
| $P ::= T \mid \star$ | <i>method type parameter</i> |
| $\Delta ::= \overline{X} \rightarrow [\overline{B}_L - \overline{B}_U]$ | <i>type ranges</i> |
| $\Gamma ::= x : \overline{T}$ | <i>var environments</i> |
| $X ::= \dots$ | <i>type vars</i> |
| $x ::= \dots$ | <i>expr vars</i> |
| $C ::= \dots$ | <i>class names</i> |

Lookup Functions:

Shared premise for lookup rules except F-OBJ:

$$CT(C) = \text{class } C\langle \overline{v} X \rightarrow [\overline{B}_L - \overline{B}_U] \rangle \triangleleft N \{ \overline{S} f; \overline{M} \}$$

$$\text{fields}(\text{Object}) = \emptyset \quad (\text{F-OBJ})$$

$$\text{fields}(C) = \overline{g}, \overline{f}, \text{ if } N = D\langle \overline{U} \rangle \text{ and } \text{fields}(D) = \overline{g} \quad (\text{F-SUPER})$$

$$\text{ftype}(f; C\langle \overline{T} \rangle) = \text{ftype}(f; [\overline{T}/\overline{X}]N), \text{ if } f \notin \overline{f} \quad (\text{FT-SUPER})$$

$$\text{ftype}(f_i; C\langle \overline{T} \rangle) = [\overline{T}/\overline{X}]S_i, \quad (\text{FT-CLASS})$$

$$\text{mtype}(m; C\langle \overline{T} \rangle) = \text{mtype}(m; [\overline{T}/\overline{X}]N), \text{ if } m \notin \overline{M} \quad (\text{MT-SUPER})$$

$$\text{mtype}(m; C\langle \overline{T} \rangle) = [\overline{T}/\overline{X}](\langle \Delta \rangle \overline{U} \rightarrow U), \text{ if } \langle \Delta \rangle \overline{U} m(\overline{U} x) \{ \text{return } e; \} \in \overline{M} \quad (\text{MT-CLASS})$$

$$\text{mbody}(m; C\langle \overline{T} \rangle) = \text{mbody}(m; [\overline{T}/\overline{X}]N), \text{ if } m \notin \overline{M} \quad (\text{MB-SUPER})$$

$$\text{mbody}(m; C\langle \overline{T} \rangle) = \langle \overline{x}. [\overline{T}/\overline{X}]e \rangle, \text{ if } \langle \Delta \rangle \overline{U} m(\overline{U} x) \{ \text{return } e; \} \in \overline{M} \quad (\text{MB-CLASS})$$

Fig. 3. Syntax and Lookup Functions

4.1 Variance of a Type

Before we embark on the specifics of the VarJ formalism, we examine the essence of variance reasoning, i.e., how variances are computed in type expressions. For now, consider the subtyping relation of our formalism as a black box—it will be defined in Section 4.3. When is a type instantiation $\mathbf{C}\langle Exp1 \rangle$ a subtype of another instantiation $\mathbf{C}\langle Exp2 \rangle$? We answer a more general question by defining a general predicate $var(\mathbf{X}; \mathbf{T})$, where \mathbf{X} is a type variable and \mathbf{T} is a type expression. The goal of var is to determine: Given a type variable \mathbf{X} and a type expression \mathbf{T} that can contain \mathbf{X} , what is the subtyping relationship between different “instantiations” of \mathbf{T} with respect to (wrt) \mathbf{X} , where an instantiation of \mathbf{T} wrt to \mathbf{X} is a substitution for \mathbf{X} in \mathbf{T} . For example, we want $var(\mathbf{X}; \mathbf{T}) = +$ to imply $[\mathbf{U}/\mathbf{X}]\mathbf{T} <: [\mathbf{U}'/\mathbf{X}]\mathbf{T}$, if $\mathbf{U} <: \mathbf{U}'$.

To define var , we use predicate $\mathbf{v}(\mathbf{T}; \mathbf{T}')$ as a notational shorthand, denoting the type of subtype relation between \mathbf{T} and \mathbf{T}' :

- $+(\mathbf{T}; \mathbf{T}') \equiv \mathbf{T} <: \mathbf{T}'$ • $-(\mathbf{T}; \mathbf{T}') \equiv \mathbf{T}' <: \mathbf{T}$
- $\mathbf{o}(\mathbf{T}; \mathbf{T}') \equiv +(\mathbf{T}; \mathbf{T}') \wedge -(\mathbf{T}; \mathbf{T}')$ • $\ast(\mathbf{T}; \mathbf{T}') \equiv true$

Note that, by the variance lattice (in Figure 1), we have

$$\mathbf{v} \leq \mathbf{w} \implies \left[\mathbf{v}(\mathbf{T}; \mathbf{T}') \implies \mathbf{w}(\mathbf{T}; \mathbf{T}') \right] \quad (1)$$

In general, we want for var the following property, which is a generalization of the subtype lifting lemma of Emir et al.’s modeling of definition-site variance [8]:

$$var(\mathbf{X}; \mathbf{T}) = \mathbf{v} \implies \left[\mathbf{v}(\mathbf{U}; \mathbf{U}') \implies [\mathbf{U}/\mathbf{X}]\mathbf{T} <: [\mathbf{U}'/\mathbf{X}]\mathbf{T} \right] \quad (2)$$

By (1), (2) entails a more general implication:

$$\mathbf{v} \leq var(\mathbf{X}; \mathbf{T}) \implies \left[\mathbf{v}(\mathbf{U}; \mathbf{U}') \implies [\mathbf{U}/\mathbf{X}]\mathbf{T} <: [\mathbf{U}'/\mathbf{X}]\mathbf{T} \right] \quad (3)$$

We assume there is a usual class table CT that maps class identifiers \mathbf{C} to their definition (i.e. $CT(\mathbf{C}) = \mathbf{class} \ \mathbf{C}\langle \overline{\mathbf{v}\mathbf{X}} \rightarrow [\mathbf{B}_L - \mathbf{B}_U] \rangle \triangleleft \mathbf{N} \{ \dots \}$). Similarly, we define a variance table VT that maps class identifiers to their type parameters with their def-site variances. For example, assuming the class table mapping above, $VT(\mathbf{C}) = \overline{\mathbf{v}\mathbf{X}}$. VT is overloaded to take an extra index parameter i to the i^{th} def-site variance annotation; e.g, if $VT(\mathbf{C}) = \overline{\mathbf{v}\mathbf{X}}$, then $VT(\mathbf{C}, i) = \mathbf{v}_i$.

The expression $var(\mathbf{X}; \mathbf{B})$ computes the variance of type variable \mathbf{X} in type expression \mathbf{B} . Figure 4 contains var ’s definition. var ’s type input is overloaded for non-existential types (\mathbf{R}) and type ranges (Δ). ($var(\overline{\mathbf{X}}; \phi$) is further overloaded in the expected way for computing variances for sequences of type variables.)

The var relation is used in our type system to determine which variance is appropriate for each type expression. Eventually our proof connects it to the subtype relation, in Lemma 1. (Proofs of all key lemmas can be found in a technical report available at <http://people.cs.umass.edu/~jaltidor/ecoop12tr.pdf>.)

| | |
|---|-----------|
| Variance of Types and Ranges: $var(\mathbf{X}; \phi)$, where $\phi : : = \mathbf{B} \mid \mathbf{R} \mid \Delta$ | |
| $var(\mathbf{X}; \mathbf{X}) = +$ | (VAR-XX) |
| $var(\mathbf{X}; \mathbf{Y}) = *$, if $\mathbf{X} \neq \mathbf{Y}$ | (VAR-XY) |
| $var(\mathbf{X}; \mathbf{C} < \overline{\mathbf{T}} >) = \prod_{i=1}^n (v_i \otimes var(\mathbf{X}; \mathbf{T}_i))$, if $VT(\mathbf{C}) = \overline{v\mathbf{X}}$ | (VAR-N) |
| $var(\mathbf{X}; \perp) = *$ | (VAR-B) |
| $var(\mathbf{X}; \exists \Delta. \mathbf{R}) = var(\mathbf{X}; \Delta) \sqcap var(\mathbf{X}; \mathbf{R})$, if $\mathbf{X} \notin dom(\Delta)$ | (VAR-T) |
| $var(\mathbf{X}; \overline{\mathbf{Y} \rightarrow [\mathbf{B}_L - \mathbf{B}_U]}) = \prod_{i=1}^n ((- \otimes var(\mathbf{X}; \mathbf{B}_{L_i})) \sqcap (+ \otimes var(\mathbf{X}; \mathbf{B}_{U_i})))$ | (VAR-R) |
| $[var(\overline{\mathbf{X}}; \phi) = \overline{v}] \equiv [\forall i, var(\mathbf{X}_i; \phi) = v_i]$, where $\phi : : = \mathbf{B} \mid \mathbf{R} \mid \Delta$ | (VAR-SEQ) |

Fig. 4. Variance of types and ranges

Lemma 1 (Subtype Lifting Lemma). If (a) $\overline{v} \leq var(\overline{\mathbf{X}}; \mathbf{B})$ and (b) $\Delta \vdash \overline{v(\mathbf{T}; \mathbf{U})}$ then $[\mathbf{T}/\mathbf{X}]\mathbf{B} <: [\mathbf{U}/\mathbf{X}]\mathbf{B}$.

We provide some intuition on the soundness of var 's definition. One “base case” of var 's definition is the VAR-XX rule. To see why it returns $+$, note that the desired implication from the subtype lifting lemma holds for this case: if $+(\mathbf{T}; \mathbf{U})$, which is equivalent to $\mathbf{T} <: \mathbf{U}$, then $[\mathbf{T}/\mathbf{X}]\mathbf{X} = \mathbf{T} <: \mathbf{U} = [\mathbf{U}/\mathbf{X}]\mathbf{X}$. The VAR-N rule computes the variance in a non-variable type using the \otimes operator, which determines how variances compose, as described in Section 3. VAR-R computes the variance of a type variable in a range. Computing the variance of ranges is necessary for computing the variance of constraints from type bounds on type parameters, which occur in existential types and method signatures. The domains of ranges are ignored by VAR-R. A range becomes more “specialized” as the bounds get “squeezed”. Informally, a range $[\mathbf{B}_L - \mathbf{B}_U]$ is a *subrange* of range $[\mathbf{A}_L - \mathbf{A}_U]$ if $\overline{\mathbf{A}_L} <: \overline{\mathbf{B}_L}$ and $\overline{\mathbf{B}_U} <: \overline{\mathbf{A}_U}$. The variance of the lower bound is transformed by contravariance to “reverse” the subtype relation, since we want the lower bound in the subrange to be a *supertype* of the lower bound in the super-range.⁵ The subtype lifting lemma can be used to entail subrange relationships:

$$\begin{aligned} var(\mathbf{X}; \overline{\mathbf{Y} \rightarrow [\mathbf{B}_L - \mathbf{B}_U]}) = \overline{v} \text{ and } v(\mathbf{T}; \mathbf{U}) \\ \implies \overline{[\mathbf{U}/\mathbf{X}]\mathbf{B}_L} <: \overline{[\mathbf{T}/\mathbf{X}]\mathbf{B}_L} \text{ and } \overline{[\mathbf{T}/\mathbf{X}]\mathbf{B}_U} <: \overline{[\mathbf{U}/\mathbf{X}]\mathbf{B}_U} \end{aligned}$$

The variance of an existential type variable is just the meet of the variances of its range (Δ) and its body (\mathbf{R}). The VAR-T rule has the premise “ $\mathbf{X} \notin dom(\Delta)$ ” to follow *Barendregt's variable convention* [19], as in the TameFJ formalism. (var is undefined when this premise is not satisfied.) The variable convention substantially reduces the number of places requiring alpha-conversion to be applied and allows for more elegant proofs. The rules of the convention basically are: (1) relations in the type system are equivariant (respect alpha-renaming), and (2) no binder (declared variable) in a rule occurs free in the conclusion. For example, condition (1) holds for $var(\mathbf{X}; \mathbf{T})$ because we can rename binders

⁵ Intuitively, the upper/lower bounds are in co-/contravariant positions, respectively.

to fresh names in existential types in T without changing the variance of X in T . Without the premise of rule $\mathsf{VAR}\text{-}\mathsf{T}$, this property would no longer hold. So that the important premises are more apparent, in the remaining rules we skip such “side-conditions” in the text and just mention that the premises for following the variable convention are implicit.

4.2 Variance of a Position

To see how *var* is used in our type system, we need to consider the variance of positions in a class definition. For example, return types are assumed to be in covariant positions while argument types are assumed to be in contravariant positions. These assumptions are used to type check class and method definitions and their def-site variance annotations.

The expressions “ $\overline{\mathsf{v}} \leq \mathit{var}(\overline{\mathsf{X}}; \mathsf{B})$ ” and “ $-\otimes \mathsf{v}$ ” are used frequently in the VarJ formalism. To connect our notation to previous work, we define the following:

$$\left[\overline{\mathsf{v}}\overline{\mathsf{X}} \vdash \mathsf{B} \textit{ mono} \right] \equiv \left[\overline{\mathsf{v}} \leq \mathit{var}(\overline{\mathsf{X}}; \mathsf{B}) \right] \quad (4)$$

$$\left[\neg \mathsf{v} \right] = \left[- \otimes \mathsf{v} \right] \quad (5)$$

A “monotonicity” judgment of the syntactic form “ $\overline{\mathsf{v}}\overline{\mathsf{X}} \vdash \mathsf{T} \textit{ mono}$ ” appears originally in Emir et al.’s definition-site variance treatment [8] and later in Kennedy and Pierce [13] as “ $\overline{\mathsf{v}}\overline{\mathsf{X}} \vdash \mathsf{T} \textit{ ok}$ ”. The semantics of these judgments in the aforementioned sources are similar to its definition here but differs in that they had no function similar to *var* nor a variance lattice. The negation operator \neg also appears in [8] and [13] and is used to transform a variance by contravariance. Using the implications in Section 4.1, it is easy to show the following properties, which are important for type checking class definitions:

$$\mathsf{w} = \neg \mathsf{v} \implies \left[\mathsf{v}(\mathsf{B}, \mathsf{B}') \iff \mathsf{w}(\mathsf{B}', \mathsf{B}) \right] \quad (6)$$

$$\overline{\mathsf{v}}\overline{\mathsf{X}} \vdash \mathsf{B} \textit{ mono} \implies \left[\overline{\mathsf{v}}(\overline{\mathsf{T}}, \overline{\mathsf{U}}) \implies \left[\overline{\mathsf{T}}/\overline{\mathsf{X}} \right] \mathsf{B} <: \left[\overline{\mathsf{U}}/\overline{\mathsf{X}} \right] \mathsf{B} \right] \quad (7)$$

$$\overline{\mathsf{v}}\overline{\mathsf{X}} \vdash \mathsf{B} \textit{ mono} \implies \left[\overline{\mathsf{v}}(\overline{\mathsf{T}}, \overline{\mathsf{U}}) \implies \left[\overline{\mathsf{U}}/\overline{\mathsf{X}} \right] \mathsf{B} <: \left[\overline{\mathsf{T}}/\overline{\mathsf{X}} \right] \mathsf{B} \right] \quad (8)$$

Figure 5 contains rules for checking class and method definitions and the definition of the *override* predicate. Premises related to type checking with definition-site variance are highlighted. Auxiliarily lookup functions are used to compute the types of members (fields and methods) in class definitions. Their definitions are in Figure 3. These lookup functions take in non-variable types (N) instead of existential types. In the expression typing rules (in Figure 7), existential types are implicitly “unpacked” to non-variable types to type some expressions such as a field access. The process for packing and unpacking types is similar to the process performed in the TameFJ formalism. Section 4.4 has a brief overview of this process and an example type derivation.

Class and Method Typing:

$$\begin{array}{c}
\overline{vX} \vdash N, \overline{T} \text{ mono} \quad \Delta = \overline{X} \rightarrow [\overline{B_L - B_U}] \\
\hline
\emptyset \vdash \Delta \text{ OK} \quad \Delta \vdash N, \overline{T} \text{ OK} \quad \vdash \overline{M} \text{ OK in } C \\
\hline
\vdash \text{class } C \langle \overline{vX} \rightarrow [\overline{B_L - B_U}] \rangle \triangleleft N \{ \overline{T} \text{ f}; \overline{M} \} \text{ OK} \\
\text{(W-CLS)} \\
\\
CT(C) = \text{class } C \langle \overline{vX} \rightarrow [\overline{B_L - B_U}] \rangle \triangleleft N \{ \dots \} \\
\Delta = \overline{X} \rightarrow [\overline{B_L - B_U}] \quad \Delta \vdash \Delta' \text{ OK} \quad \Delta, \Delta' \vdash T, \overline{T} \text{ OK} \\
\text{override}(m; N; \langle \Delta' \rangle \overline{T} \rightarrow T) \quad \overline{vX} \vdash \overline{T}, \Delta' \text{ mono} \quad \overline{vX} \vdash T \text{ mono} \\
\Delta, \Delta'; \overline{x} : \overline{T}, \text{this} : \exists \emptyset. C \langle \overline{X} \rangle \vdash e : T \mid \emptyset \\
\hline
\vdash \langle \Delta' \rangle T \text{ m}(\overline{T} \overline{x}) \{ \text{return } e; \} \text{ OK in } C \\
\text{(W-METH)} \\
\\
\frac{mtype(m; N) = \langle \Delta \rangle \overline{T} \rightarrow T}{\text{override}(m; N; \langle \Delta \rangle \overline{T} \rightarrow T)} \quad \frac{mtype(m; N) \text{ is undefined}}{\text{override}(m; N; \langle \Delta \rangle \overline{T} \rightarrow T)} \\
\text{(OVER-DEF)} \quad \text{(OVER-UNDEF)}
\end{array}$$

Wellformed Ranges: $\Delta \vdash \Delta \text{ OK}$

$$\begin{array}{c}
\frac{}{\Delta \vdash \emptyset \text{ OK}} \quad \text{(W-RNG-EMPTY)} \\
\frac{X \notin \text{dom}(\Delta) \quad \Delta, X \rightarrow [\overline{B_L - B_U}], \Delta' \vdash B_L, B_U \text{ OK} \\
\Delta \vdash \text{ubound}_\Delta(B_L) \sqsubseteq: \text{ubound}_\Delta(B_U) \\
\Delta \vdash B_L <: B_U \quad \Delta, X \rightarrow [\overline{B_L - B_U}] \vdash \Delta' \text{ OK}}{\Delta \vdash X \rightarrow [\overline{B_L - B_U}], \Delta' \text{ OK}} \quad \text{(W-RNG)}
\end{array}$$

Non-Variable Upper Bound: $\text{ubound}_\Delta(B)$

$$\text{ubound}_\Delta(B) = \begin{cases} \text{ubound}_\Delta(B_U), & \text{if } B = \exists \emptyset. X, \text{ where } \Delta(X) = [B_L - B_U] \\ B, & \text{if } B = \exists \Delta'. N \end{cases}$$

Wellformed Types: $\Delta \vdash \phi \text{ OK}$, where $\phi ::= B \mid P \mid R$

$$\begin{array}{c}
\frac{}{\Delta \vdash \text{Object} \langle \rangle \text{ OK}} \quad \text{(W-OBJ)} \quad \frac{X \in \text{dom}(\Delta)}{\Delta \vdash X \text{ OK}} \quad \text{(W-X)} \quad \frac{}{\Delta \vdash \perp \text{ OK}} \quad \text{(W-B)} \quad \frac{}{\Delta \vdash \star \text{ OK}} \quad \text{(W-I)} \\
\\
\frac{\text{class } C \langle \overline{vX} \rightarrow [\overline{B_L - B_U}] \rangle \triangleleft N \{ \dots \} \\
\Delta \vdash [\overline{T/X}] B_L <: T \quad \Delta \vdash T <: [\overline{T/X}] B_U \\
\Delta \vdash \overline{T} \text{ OK}}{\Delta \vdash C \langle \overline{T} \rangle \text{ OK}} \quad \text{(W-N)} \quad \frac{\Delta \vdash \Delta' \text{ OK} \quad \Delta, \Delta' \vdash R \text{ OK}}{\Delta \vdash \exists \Delta'. R \text{ OK}} \quad \text{(W-T)}
\end{array}$$

Wellformed Expression Variable Environments $\Delta \vdash \Gamma \text{ OK}$

$$\frac{}{\Delta \vdash \emptyset \text{ OK}} \quad \text{(W-ENV-EMPTY)} \quad \frac{x \notin \text{dom}(\Gamma) \quad \Delta \vdash T \text{ OK} \quad \Delta \vdash \Gamma \text{ OK}}{\Delta \vdash \Gamma, x : T \text{ OK}} \quad \text{(W-ENV)}$$

Fig. 5. Wellformedness Judgments

The definition-site subtyping relation judgment $\Delta \vdash N \prec: N'$ is defined over non-variable types and considers definition-site annotations when concluding subtype relationships. For example, $VT(C) = +x \implies \Delta \vdash C \prec \exists \emptyset.\text{Dog} \prec: C \prec \exists \emptyset.\text{Animal}$, assuming $\Delta \vdash \exists \emptyset.\text{Dog} \prec: \exists \emptyset.\text{Animal}$. This relation is defined in Figure 6.

| Definition-Site Subtyping: $R \prec: R$ | | |
|---|---|---|
| $\frac{\text{class } C \langle \overline{vX} \rightarrow [B_L - B_U] \rangle \triangleleft N \{ \dots \} \quad C \neq D \quad \Delta \vdash [\overline{T/X}]N \prec: D \langle \overline{U} \rangle}{\Delta \vdash C \langle \overline{T} \rangle \prec: D \langle \overline{U} \rangle} \quad \text{(SD-SUPER)}$ | $\frac{VT(C) = \overline{vX} \quad \Delta \vdash \overline{v(T, U)}}{\Delta \vdash C \langle \overline{T} \rangle \prec: C \langle \overline{U} \rangle} \quad \text{(SD-VAR)}$ | $\frac{}{\Delta \vdash X \prec: X} \quad \text{(SD-X)}$ |
| Existential Subtyping: $\Delta \vdash B \sqsubset: B$ | | |
| $\frac{\Delta, \Delta' \vdash N \prec: N'}{\Delta \vdash \exists \Delta'. N \sqsubset: \exists \Delta'. N'} \quad \text{(SE-SD)}$ | $\frac{}{\Delta \vdash B \sqsubset: B} \quad \text{(SE-REFL)}$ | $\frac{\Delta \vdash B \sqsubset: B' \quad \Delta \vdash B' \sqsubset: B''}{\Delta \vdash B \sqsubset: B''} \quad \text{(SE-TRAN)}$ |
| $\text{dom}(\Delta') \cap \text{fv}(\overline{\exists X} \rightarrow [B_L - B_U].N) = \emptyset \quad \text{fv}(\overline{T}) \subseteq \text{dom}(\Delta, \Delta')$ | | |
| $\frac{}{\Delta \vdash \perp \sqsubset: B} \quad \text{(SE-BOT)}$ | $\frac{\Delta, \Delta' \vdash [\overline{T/X}]B_L \prec: T \quad \Delta, \Delta' \vdash T \prec: [\overline{T/X}]B_U}{\Delta \vdash \exists \Delta'. [\overline{T/X}]N \sqsubset: \overline{\exists X} \rightarrow [B_L - B_U].N} \quad \text{(SE-PACK)}$ | |
| Subtyping: $\Delta \vdash B \prec: B$ | | |
| $\frac{\Delta \vdash B \sqsubset: B'}{\Delta \vdash B \prec: B'} \quad \text{(ST-SE)}$ | $\frac{\Delta \vdash B \prec: B' \quad \Delta \vdash B' \prec: B''}{\Delta \vdash B \prec: B''} \quad \text{(ST-TRAN)}$ | $\frac{\Delta(X) = [B_L - B_U]}{\Delta \vdash B_L \prec: \exists \emptyset.X} \quad \text{(ST-LBOUND)}$ $\Delta \vdash \exists \emptyset.X \prec: B_U \quad \text{(ST-UBOUND)}$ |

Fig. 6. Subtyping Relations

The motivation for the assumed variances of positions is to ensure the subsumption principle holds for the subtyping hierarchy. Informally, if $T \prec: U$, then a value of type T may be provided whenever a value of type U is required. In the case of `VarJ`, the subsumption principle is established by showing appropriate subtype relationships between types of members from class definitions. Lemma 2 states a goal subsumption property, which is to have the type of field \mathbf{f} of the supertype N' become a more specific type for the subtype N . Although inherited fields syntactically have the same type as in the superclass definition, definition-site subtyping allows fields to have more specific types in the subtype. Lemma 3 states the goal subsumption property for types in method signatures; the sixth conclusion of this lemma holds because of the *override* predicate.

Lemma 2 (Subtyping Specializes Field Type).

If (a) $\vdash \text{class } C \langle \overline{vX} \rightarrow [\dots] \rangle \triangleleft N \dots$ OK and (b) $\Delta \vdash C \langle \overline{T} \rangle \prec: N'$ and (c) $f\text{type}(f; N') = T$, then $\Delta \vdash f\text{type}(f; C \langle \overline{T} \rangle) \prec: T$.⁶

Lemma 3 (Subtyping Specializes Method Type).

If (a) $\vdash \text{class } C \langle \overline{vX} \rightarrow [\dots] \rangle \triangleleft N \dots$ OK and (b) $\Delta \vdash C \langle \overline{T} \rangle \prec: N'$ and (c) $m\text{type}(m; N') = \langle \overline{Y} \rightarrow [\overline{B_L - B_U}] \rangle \overline{U} \rightarrow U$, then: (1) $m\text{type}(m; C \langle \overline{T} \rangle) = \langle \overline{Y} \rightarrow [\overline{A_L - A_U}] \rangle \overline{V} \rightarrow V$, (2) $\Delta \vdash V \prec: U$, (3) $\Delta \vdash \overline{U} \prec: \overline{V}$, (4) $\Delta \vdash \overline{A_L} \prec: \overline{B_L}$, (5) $\Delta \vdash \overline{B_U} \prec: \overline{A_U}$, and (6) $\text{var}(\overline{Y}; U) = \text{var}(\overline{Y}; V)$.

To satisfy the two lemmas above, we make assumptions about the variance of the positions that types can occur in. To *preserve* the subtype relationship order of a type in a member signature, we assume the type occurs in a *covariant* position (i.e., the subtype needs to have a more specific type appear in such a position). To *reverse* the subtype relationship order of a type in a member signature, we assume the type occurs in a *contravariant* position. The assumptions about the variance of the positions are reflected in the *mono* judgments in the W-CLS and W-METH rules for checking class and method definitions. By (7), *not* negating the def-site variance annotations, \overline{v} , in the judgment “ $\overline{vX} \vdash T$ *mono*” reflects that T is assumed to be in a covariant position. Since covariance, $+$, is the identity element for the \otimes operator ($+ \otimes v = v$), the variances \overline{v} do not need to be transformed by $+$. By (8), negating the def-site variance annotations in the judgment “ $\overline{vX} \vdash T$ *mono*” reflects that T is assumed to be in a contravariant position. We need to reverse the subtype relationship order for argument types and ranges in method type signatures. Negating the variance annotations for the argument types ensures the argument types are more general supertypes for the subtype.⁷

Negating the range of a method type signature ensures the range is *wider* for the subtype. For code examples motivating why ranges need to be widened for the subtype, see Section 2.4 of [8]. More generally, if (1) $e.\langle T \rangle_m()$ type checks implying the type actual T is within the type bounds for m 's type argument and (2) $\text{typeof}(e') \prec: \text{typeof}(e)$, then $e'.\langle T \rangle_m()$ should type check as well even if m is overridden in the subclass. Hence, the subtype's version of m should accept a superset/wider range of types than accepted by the supertype's version of m .

4.3 Subtyping

Subtyping in VarJ is defined similarly to TameFJ. Figure 6 contains the subtyping rules. There are three levels of subtyping in VarJ, as in TameFJ. The first level of subtyping in TameFJ, the subclass relation, has been replaced with the definition-site subtyping relation \prec : defined on non-existential types. Def-site

⁶ If field assignments were allowed, then field types would be in both co- and contravariant positions, and both $f\text{type}(f; C \langle \overline{T} \rangle)$ and $f\text{type}(f; N')$ would be subtypes of each other.

⁷ Bounds on class type parameters may make unrestricted use of type parameters by similar reasoning as in [8, p.7]. Once an object is created, they are forgotten.

subtyping is defined by the SD-^* rules, which are similar to the subtyping rules from [13]. Like the subtype relation from [13], \prec : is defined by syntax-directed rules⁸ and shares the reflexive and transitive properties by similar reasoning as in [13]. The \prec : judgment requires a typing context to check subtyping relationships between pairs of type actuals as done in the SD-VAR rule.

The existential subtyping relation \sqsubset : is defined by the SE-^* rules and is similar to the “Extended subclasses” relation in TameFJ. The XS-ENV rule from TameFJ was renamed to SE-PACK ; it is the only subtyping rule that can pack (and actually also unpack) types into existential type variables. The XS-SUB-CLASS rule was not only renamed to SE-SD , but its premise was updated to use def-site subtyping. SE-SD allows def-site subtyping to be applied to both type variables in the type context Δ and existential type variables in Δ' . As a result, a type packed to an existential type variable may not be in the range of the variable. For example, if `Iterator` is covariant in its type parameter ($VT(\text{Iterator}) = +X$), then the following subtype relationship is derivable: $\exists \emptyset. \text{Iterator} \prec \text{PrettyDog} \prec \exists \emptyset. \text{Iterator} \prec \text{Dog} \prec \exists Y \rightarrow [\text{Dog-Animal}]. \text{Iterator} \prec Y$. Subtyping between two types implies the subsumption principle between the types. Since $\text{Iterator} \prec \text{Dog}$ can be packed to $\exists X \rightarrow [\text{Dog-Animal}]. \text{Iterator} \prec X$ and $\text{Iterator} \prec \text{PrettyDog} \prec \text{Iterator} \prec \text{Dog}$, it must be the case that $\text{Iterator} \prec \text{PrettyDog}$ can also be packed to $\exists X \rightarrow [\text{Dog-Animal}]. \text{Iterator} \prec X$. This intuition is formalized in Lemma 4, which is similar to Lemma 35 from TameFJ, and establishes a relationship between existential subtyping and def-site subtyping.

Lemma 4 (Existential subtyping to def-site subtyping). If (a) $\Delta \vdash \exists \Delta'. R' \sqsubset \exists X \rightarrow [\text{B}_L\text{-B}_U]. R$ and (b) $\emptyset \vdash \Delta \text{ OK}$, then there exists \bar{T} such that: (1) $\Delta, \Delta' \vdash R' \prec: [\bar{T}/X]R$ and (2) $\Delta, \Delta' \vdash [\bar{T}/X]B_L \prec: T$ and (3) $\Delta, \Delta' \vdash T \prec: [\bar{T}/X]B_U$ and (4) $fv(\bar{T}) \subseteq dom(\Delta, \Delta')$.

Existential subtyping does not conclude subtype relationships for type variables except for the reflexive case using SE-REFL . The (all) subtyping relation \prec : allows non-reflexive subtype relationships with type variables by considering their bounds in the typing context. Since T or U may be type variables in a subtype relationship $T \prec: U$, we want a stronger relationship between the non-variable upper bounds of T and U . Lemma 5 formalizes this notion and is similar to lemma 17 from TameFJ. The non-variable upper bound of a type T is $ubound_\Delta(T)$, defined in Figure 5.

Lemma 5 (Subtyping to existential subtyping). If (a) $\Delta \vdash T \prec: T'$ and (b) $\emptyset \vdash \Delta \text{ OK}$ then $\Delta \vdash ubound_\Delta(T) \sqsubset: ubound_\Delta(T')$.

4.4 Typing and Wildcard Capture

The expression typing rules in VarJ are mostly the same as in TameFJ and are given in Figure 7. Unlike TameFJ, VarJ allows method signatures to have lower

⁸ The syntax-directed nature of these rules does not ensure that an algorithmic test of \prec : is straightforward, because the premise of rule SD-VAR appeals to the definition of the full \prec : relation (hidden inside the v shorthand).

bounds. The *sift* function is needed for safe wildcard capture and is applied in the T-INVK rule for typing method invocations. The definition of *sift* required updating because of interaction with variant types. First, we give a brief overview of expression typing; see [4] for more thorough coverage.

Expression Typing. Consider the Java segment below. It type checks because the expression `box.elem` is typed as `String`. The type of `box.elem` is the same as the type actual passed to the `Box` type constructor. In this case, the type actual is “`? extends String`”, which refers to some unknown subtype of `String`. To type `box.elem` with some known/named type, the most specific named type that can be assigned to `box.elem` is chosen, which is `String`.

```
class Box<E> { E elem; Box(E elem) { this.elem = elem; } }
Box<? extends String> box = ...
box.elem.charAt(0);
```

We explain this type derivation through the formal calculus. Types hidden by wildcards such as “`? extends String`” are “captured” as existential type variables. The type `Box<? extends String>` is modeled in VarJ by $\exists X \rightarrow [\perp\text{-String}].\text{Box}\langle X \rangle$. Expression typing judgments have the form $\Delta; \Gamma \vdash e : T \mid \Delta'$. The second type variable environment Δ' is the *guard* of the judgment. It is used to keep track of type variables that have been unpacked from existential types during type checking. Variables in $\text{dom}(\Delta')$ may occur free in T and model hidden types. To type an expression without exposed (free) hidden types (existential type variables), the T-SUBS rule is applied to find a suitable type without free existential type variables. The example typing derivation below illustrates this process on typing the “`box.elem`” expression from the previous code segment, where we assume $\Gamma = \text{box} : \exists X \rightarrow [\perp\text{-String}].\text{Box}\langle X \rangle$.

$$\begin{array}{c}
 \frac{\frac{\emptyset; \Gamma \vdash \text{box} : \exists X \rightarrow [\perp\text{-String}].\text{Box}\langle X \rangle \mid \emptyset \quad \text{ftype}(\text{elem}; \text{Box}\langle X \rangle) = X}{\emptyset; \Gamma \vdash \text{box.elem} : X \mid X \rightarrow [\perp\text{-String}]} \quad \frac{\emptyset, X \rightarrow [\perp\text{-String}] \vdash X <: \text{String} \quad \emptyset \vdash X \rightarrow [\perp\text{-String}] \text{ OK}}{\emptyset \vdash \text{String} \text{ OK}}}{\emptyset; \Gamma \vdash \text{box.elem} : \text{String} \mid \emptyset} \\
 \text{(T-FIELD)} \qquad \qquad \qquad \text{(T-SUBS)}
 \end{array}$$

Matching for Wildcard Capture. The T-INVK rule type checks a method invocation and uses *match* to perform wildcard capture. The definition of *match* is updated to use the definition-site subtyping relation ($<:$). Ignoring return types, consider a polymorphic method declared with type $\langle \bar{Y} \rangle_m(\bar{U})$ and called with types $\langle \bar{P} \rangle_m(\exists \Delta. \bar{R})$. The parameters of *match*($\bar{R}; \bar{U}; \bar{P}; \bar{Y}; \bar{T}$) and their expected conditions are:

1. The bodies of the actual value argument types of a method invocation (\bar{R}).
2. The formal argument value types of a method (\bar{U}).
3. The *specified* type actuals of a method invocation (\bar{P}).
4. The formal type arguments of a method (\bar{Y}).
5. The *inferred* type actuals of a method invocation (\bar{T}).

| | |
|--|---|
| Expression Typing: $\Delta; \Gamma \vdash e : \tau \mid \Delta$ | |
| $\frac{}{\Delta; \Gamma \vdash x : \Gamma(x) \mid \emptyset}$ <p style="text-align: center;">(T-VAR)</p> | $\frac{\Delta \vdash c \langle \overline{\tau} \rangle \text{ OK} \quad \text{fields}(c) = \overline{f}}{\text{ftype}(f, c \langle \overline{\tau} \rangle) = \text{U}}$ $\frac{}{\Delta; \Gamma \vdash e : \text{U} \mid \emptyset}$ <p style="text-align: center;">(T-NEW)</p> |
| $\frac{\Delta; \Gamma \vdash e : \exists \Delta'. N \mid \emptyset \quad \text{ftype}(f; N) = \text{T}}{\Delta; \Gamma \vdash e.f : \text{T} \mid \Delta'}$ <p style="text-align: center;">(T-FIELD)</p> | $\frac{\Delta; \Gamma \vdash e : \text{U} \mid \Delta' \quad \Delta, \Delta' \vdash \text{U} <: \text{T} \quad \Delta \vdash \Delta' \text{ OK} \quad \Delta \vdash \text{T OK}}{\Delta; \Gamma \vdash e : \text{T} \mid \emptyset}$ <p style="text-align: center;">(T-SUBS)</p> |
| $\Delta; \Gamma \vdash e : \exists \Delta'. N \mid \emptyset \quad \text{mtype}(m; N) = \langle \overline{Y} \rightarrow [\text{B}_L - \text{B}_U] \rangle \overline{\text{U}} \rightarrow \text{U}$ | |
| $\frac{\Delta \vdash \overline{P} \text{ OK} \quad \Delta; \Gamma \vdash e : \exists \Delta.R \mid \emptyset \quad \text{sift}(\overline{R}; \overline{\text{U}}; \overline{\text{Y}}) = (\overline{R'}; \overline{\text{U}'}) \quad \text{match}(\overline{R'}; \overline{\text{U}'}; \overline{P}; \overline{\text{Y}}; \overline{\text{T}})}{\Delta'' = \Delta, \Delta', \overline{\Delta} \quad \Delta'' \vdash \exists \emptyset.R <: [\overline{\text{T}}/\overline{\text{Y}}]\text{U} \quad \Delta'' \vdash [\overline{\text{T}}/\overline{\text{Y}}]\text{B}_L <: \text{T} \quad \Delta'' \vdash \text{T} <: [\overline{\text{T}}/\overline{\text{Y}}]\text{B}_U}$ | |
| $\frac{}{\Delta; \Gamma \vdash e. \langle \overline{P} \rangle_m(\overline{e}) : [\overline{\text{T}}/\overline{\text{Y}}]\text{U} \mid \Delta', \overline{\Delta}}$ <p style="text-align: center;">(T-INVK)</p> | |
| Match: | |
| $\frac{\forall j, P_j = \star \implies Y_j \in \text{fv}(\overline{R'}) \quad \forall i, P_i \neq \star \implies T_i = P_i \quad \emptyset \vdash R <: [\overline{\text{T}}/\overline{\text{Y}}, \overline{\text{T}'}/\overline{\text{X}}]\overline{R'} \quad \text{dom}(\overline{\Delta}) = \overline{\text{X}} \quad \text{fv}(\overline{\text{T}}, \overline{\text{T}'}) \cap \overline{\text{Y}}, \overline{\text{X}} = \emptyset}{\text{match}(\overline{R}; \exists \Delta.R'; \overline{P}; \overline{\text{Y}}; \overline{\text{T}})}$ <p style="text-align: center;">(MATCH)</p> | |
| Sift: $\text{sift}(\overline{R}; \overline{\text{U}}; \overline{\text{Y}}) = (\overline{R'}; \overline{\text{U}'})$ | |
| $\frac{}{\text{sift}(\emptyset; \emptyset; \overline{\text{Y}}) = (\emptyset; \emptyset)}$ <p style="text-align: center;">(SIFT-EMPTY)</p> | $\frac{\overline{\text{Y}} \cap \text{fv}(\text{U}) = \overline{\text{X}} \quad \text{var}(\overline{\text{X}}; \text{U}) = \overline{o} \quad \text{sift}(\overline{R}; \overline{\text{U}}; \overline{\text{Y}}) = (\overline{R'}; \overline{\text{U}'})}{\text{sift}((R, \overline{R}); (\text{U}, \overline{\text{U}}); \overline{\text{Y}}) = ((R, \overline{R}'); (\text{U}, \overline{\text{U}'}))}$ <p style="text-align: center;">(SIFT-ADD)</p> |
| $\frac{\overline{\text{Y}} \cap \text{fv}(\text{U}) = \overline{\text{X}} \quad \text{var}(\text{X}_j; \text{U}) \neq o, \text{ for some } \text{X}_j \in \overline{\text{X}} \quad \text{sift}(\overline{R}; \overline{\text{U}}; \overline{\text{Y}}) = (\overline{R'}; \overline{\text{U}'})}{\text{sift}((R, \overline{R}); (\text{U}, \overline{\text{U}}); \overline{\text{Y}}) = (\overline{R'}; \overline{\text{U}'})}$ <p style="text-align: center;">(SIFT-SKIP)</p> | |

Fig. 7. Expression Typing and Auxiliary Functions For Wildcard Capture

Figure 8 contains the reduction rules for performing runtime evaluation. The R-INVK rule also uses *match* to compute inferred type actuals because some of the specified type actuals (\bar{P}) may be the type inference marker \star . Since each occurrence of the \star marker may refer to different types, *match* is needed to compute the concrete types to substitute for the formal type arguments' (\bar{Y}) occurrences in the method body.

Computation Rules: $e \mapsto e$

$$\frac{fields(C) = \bar{f}}{new\ C\langle\bar{T}\rangle(\bar{v}).f_i \mapsto v_i}$$

(R-FIELD)

$$\frac{v = new\ N(\bar{v}') \quad \overline{v = new\ N(\bar{v}'')} \quad mbody(m; N) = \langle\bar{x}.e_0\rangle \quad mtype(m; N) = \langle Y \rightarrow [B_L - B_U] \rangle \bar{U} \rightarrow U \quad sift(\bar{N}; \bar{U}; \bar{Y}) = (\bar{N}'; \bar{U}') \quad match(\bar{N}'; \bar{U}'; \bar{P}; \bar{Y}; \bar{T})}{v.\langle\bar{P}\rangle_m(\bar{v}) \mapsto [v/x, v/this, T/Y]e_0}$$

(R-INVK)

Congruence Rules: $e \mapsto e$

| | |
|--|---|
| $\frac{e \mapsto e'}{e.f \mapsto e'.f}$ <p style="text-align: center;">(RC-FIELD)</p> | $\frac{e_i \mapsto e'_i}{new\ C\langle\bar{T}\rangle(..e_i..) \mapsto new\ C\langle\bar{T}\rangle(..e'_i..)}$ <p style="text-align: center;">(RC-NEW-ARG)</p> |
| $\frac{e \mapsto e'}{e.\langle\bar{P}\rangle_m(\bar{e}) \mapsto e'.\langle\bar{P}\rangle_m(\bar{e})}$ <p style="text-align: center;">(RC-INV-RECV)</p> | $\frac{e_i \mapsto e'_i}{e.\langle\bar{P}\rangle_m(..e_i..) \mapsto e.\langle\bar{P}\rangle_m(..e'_i..)}$ <p style="text-align: center;">(RC-INV-ARG)</p> |

Fig. 8. Reduction Rules

Sifting for Wildcard Capture. The *sift* function is used in VarJ and TameFJ to filter inputs passed to *match* (in the T-INVK and R-INVK rules). The goal of *sift* is to only allow inference from types that are in “fixed” or invariant positions. Without applying *sift*, counter examples to the subject reduction (type preservation) theorem can result. First, note that the following two judgments are derivable.

1. $match(Dog; \exists\emptyset.Y; \star; Y; Dog)$ (mainly) because $Dog \prec: [Dog/Y]Y = Dog$.
2. $match(Dog; \exists\emptyset.Y; \star; Y; Animal)$ (mainly) because $Dog \prec: [Animal/Y]Y = Animal$.

Assume List is invariant and consider the following.

```
<X> List<X> createList(X arg) { return new List<X>(); }
```

```
createList<*>(new Dog()) : List<Animal>
  ↦ new List<Dog>() : List<Dog>
```

The expression `createList<*>(new Dog())` can be typed with `List<Animal>` because `new Dog() : Animal`, and the inferred type actual used for typing the expression can be `Animal`. However, the inferred type used for typing the method invocation is not required to be the same inferred type, computed in the `R-INVK` rule, substituted into the method body. Without *sift*, the above evaluation step is possible, which contradicts the subject reduction theorem, since, by the invariance of `List`, `new List<Dog>()` cannot be typed with `List<Animal>`.

In TameFJ, *sift* filters out a pair of a type actual body \mathbb{R} and a formal type \mathbb{U} , if $\mathbb{U} = \exists \emptyset.x$ and x is one of the formal type arguments (\bar{Y}). Due to *sift*, the two *match* judgments above could never be derived in TameFJ. Moreover, TameFJ allows an existential type variable to be passed as parameter for a formal type variable argument only if the formal type variable is used as a type parameter. Since every type constructor in TameFJ is assumed to be invariant, every type variable used for inference is in an invariant position. This no longer holds in VarJ with variant type constructors. If we assume `Iterator` is covariant, a counter example similar to the previous one can be produced with the following method:

```
<X> List<X> createList2(Iterator<X> arg) { return new List<X>(); }
```

Hence, we update the definition of *sift* to use *var* to check if a method type parameter occurs at most invariantly. We find the restriction of not allowing wildcard capture in variant positions not to be practically restrictive. A wildcard type for a variant type typically has an equivalent non-wildcard type. `Iterator<?>` is equivalent to `Iterator<Object>` by covariance of `Iterator`. `BiGeneric<?>` is equivalent to `BiGeneric<T>`, for any T , if `BiGeneric` is bivariate. In such cases, the need for wildcard capture is eliminated because the required type actuals to specify in a method call can be named. The VarJ grammar does not allow the bottom type \perp to be specified as a type actual. However, we have not found any practical need for wildcard capture with contravariant types.

4.5 Type Soundness

We prove type soundness for VarJ by proving the progress and subject reduction theorems below. As in TameFJ, a non-empty guard is required in the statement of the progress theorem when applying the inductive hypothesis in the proof for the case when the `T-SUBS` rule is applied.

Theorem 1 (Progress). For any Δ, e, T , if $\emptyset; \emptyset \vdash e : T \mid \Delta$, then either $e \mapsto e'$ or there exists a v such that $e = v$.

Theorem 2 (Subject Reduction). For any e, e', T , if $\emptyset; \emptyset \vdash e : T \mid \emptyset$ and $e \mapsto e'$, then $\emptyset; \emptyset \vdash e' : T \mid \emptyset$.

The key difficulty in proving these theorems can be captured by a small number of key lemmas whose proofs are substantially affected by variance reasoning. Lemma 6 is probably the main one, which relates subtyping and wildcard capture, and is similar to lemma 36 from [4]. It states that the method receiver's ability to perform wildcard capture is preserved in subtypes with respect to the

method receiver. (A similar lemma holds for method arguments.) It shows that the subsumption principle holds even under interaction with wildcard capture.

Lemma 6 (Subtyping Preserves *matching* (receiver)). If (a) $\Delta \vdash \exists \Delta_1.N_1 \sqsubset: \exists \Delta_2.N_2$ and (b) $mtype(m; N_2) = \langle \overline{Y_2} \rightarrow [\overline{B_{2L}}-\overline{B_{2U}}] \rangle \overline{U_2} \rightarrow U_2$ and (c) $mtype(m; N_1) = \langle \overline{Y_1} \rightarrow [\overline{B_{1L}}-\overline{B_{1U}}] \rangle \overline{U_1} \rightarrow U_1$ and (d) $sift(\overline{R}; \overline{U_2}; \overline{Y_2}) = (\overline{R'}; \overline{U'_2})$ and (e) $match(\overline{R'}; \overline{U'_2}; \overline{P}; \overline{Y_2}; \overline{T})$ and (f) $\emptyset \vdash \underline{\Delta} OK$ and (g) $\Delta, \Delta' \vdash \overline{T} OK$ then: (1) $sift(\overline{R}; \overline{U_1}; \overline{Y_1}) = (\overline{R'}; \overline{U'_1})$ and (2) $match(\overline{R'}; \overline{U'_1}; \overline{P}; \overline{Y_1}; \overline{T})$.

5 Discussion

Boundary Analysis. Definition-site variance can imply that the variance of a type does not depend on all of the type bounds that occur in the type. Our earlier work [1] presented a definition of $var(X; U)$ that performed a simple *boundary analysis* to compute such irrelevant bounds. As discussed in Section 3.1, if generic $C\langle Y \rangle$ is covariant wrt to Y , then the lower bound of a use-site variant instantiation is ignored, which is sound for the VarLang calculus [1]: $var(X; C\langle T \rangle) = (+ \sqcup -) \otimes var(X; T) = * \otimes var(X; T) = *$. Hence, $var(X; C\langle T \rangle) = *$, even if X occurred in the lower bound, T .

The ability to ignore type bounds is present in a disciplined way in our VarJ formalism, although there is no explicit variance joining mechanism in the definition of var . For example, if `Iterator` is covariant in its type parameter, we can infer the following (where the notation $T \equiv U$ denotes $T <: U \wedge U <: T$): $\exists X \rightarrow [\text{Dog-Animal}].\text{Iterator}\langle X \rangle \equiv \exists X \rightarrow [\perp\text{-Animal}].\text{Iterator}\langle X \rangle$. Clearly, $\exists X \rightarrow [\text{Dog-Animal}].\text{Iterator}\langle X \rangle <: \exists X \rightarrow [\perp\text{-Animal}].\text{Iterator}\langle X \rangle$ because the range of the type variable is wider in the supertype. The inverse is derivable by applying a combination of the SE-SD, SE-PACK, and ST-* rules:

$$\begin{aligned} \exists X \rightarrow [\perp\text{-Animal}].\text{Iterator}\langle X \rangle <: \exists X \rightarrow [\perp\text{-Animal}].\text{Iterator}\langle \text{Animal} \rangle \\ <: \exists \emptyset.\text{Iterator}\langle \text{Animal} \rangle <: \exists X \rightarrow [\text{Dog-Animal}].\text{Iterator}\langle X \rangle \end{aligned}$$

As we saw in Section 3.2, this reasoning is not sound in the presence of F-bounded polymorphism. It is important to realize that the issue with recursive bounds is not specific to the use of bounds in type definitions.⁹ The counterexample of Section 3.2 used `interface Trouble<P> extends List<P>> extends Iterator<P> {}`. However, even if we restrict our attention to a plain `Iterator` (or, equivalently, if the class type bound, `extends List<P>`, of `Trouble` is removed) it is still *not* safe to assume the following subtype relation, by reasoning similar to that used in Section 3.2:

$\exists X \rightarrow [\text{YourList-List}\langle X \rangle].\text{Iterator}\langle X \rangle <: \exists X \rightarrow [\text{MyList-List}\langle X \rangle].\text{Iterator}\langle X \rangle$.
The above subtype relationship would violate the subsumption principle. The

⁹ In our earlier work [1], when we performed an application to Java it sufficed to be overly conservative at this point: the mere appearance of a type variable in the upper bound of a type definition caused us to consider the definition as invariant relative to this type variable. For VarJ, which is richer in terms of where bounds can appear, even this kind of conservatism is not sufficient.

latter type can return a $\exists X \rightarrow [\text{MyList-List}\langle X \rangle].\text{List}\langle X \rangle$ from its `next` method, but the former type cannot because

$\exists X \rightarrow [\text{YourList-List}\langle X \rangle].\text{List}\langle X \rangle \not\leq \exists X \rightarrow [\text{MyList-List}\langle X \rangle].\text{List}\langle X \rangle$, by the invariance of `List`. In contrast to the earlier, correct subtyping, `VarJ` does not support the above erroneous subtyping because it cannot establish that the upper bounds of the two instantiations of `Iterator` are related: we cannot derive that $\exists X \rightarrow [\text{YourList-List}\langle X \rangle].\text{List}\langle X \rangle$ is a subtype of some non-existential type, $\exists \emptyset.\text{List}\langle T \rangle$, which is, in turn, a subtype of $\exists X \rightarrow [\text{MyList-List}\langle X \rangle].\text{List}\langle X \rangle$.

Contrasting the two examples shows that boundary analysis is complex and can be unintuitive to the programmer. Note, however, that the `VarJ` calculus merely tells us what is possible to infer correctly. A practical implementation may choose not to perform all possible inferences. A specific scenario is that of separating boundary analysis from type checking. Useless bounds can be “removed” during a preprocessing step performed *before* type checking. This is analogous to general type inference algorithms relative to type checking algorithms: type checking can be performed independently of the type inference performed to compute type annotations skipped by programmers. Our variance-based type checking can be performed independently of the “useless boundary analysis”. For example, a boundary preprocessing step could transform input type $\exists X \rightarrow [\text{Dog-Animal}].\text{Iterator}\langle X \rangle$ to the equivalent type $\exists X \rightarrow [\perp\text{-Animal}].\text{Iterator}\langle X \rangle$. This opens the door to many practical instantiations—e.g., an optimistic but possibly unsound bound inference inside an IDE (which interacts with the user, offering immediate feedback and suggesting relaxations of expressions the user types in) combined with a simple but sound checking inside the compiler.

Definition-Site Variance and Erasure. A practical issue with definition-site variance concerns its use with an erasure-based translation. Consider a covariant `class A<+X> {...}` and an invariant subtype `class B<oX> extends A<X> {...}`. We can then have:

```
A<Integer> a = new B<Integer>;
A<Object> a2 = a; // fine by covariance of A
B<Object> b = (B<Object>) a2;
```

In a language with an expansion-based translation, such as `C#`, the last line will fail dynamically: an object with dynamic type `B<Integer>` cannot be cast to a `B<Object>`. In an erasure-based translation, however, the cast cannot check the type parameter (which has been erased) and will therefore succeed, causing errors further down the road. (In this case, a runtime error could result from a non-cast expression, thus violating type soundness.) This practical consideration affects all type systems that combine variance, casts, and erasure. For instance, `Scala` already handles such cases with a static type warning. Effectively, no cast to a subtype with tighter variance is safe. The result is somewhat counter-intuitive in that it defies common patterns for safe casting. For instance, the cast could have been performed after an “`a2 instanceof B<Object>`” check to establish that `a2` is indeed of type `B<Object>`. In this case the programmer would think that the cast warning can be ignored, which is not the case. In practice, any deployment

of our type system in an erasure-based setting would have to follow the same policy as Scala regarding cast warnings.

6 Related Work

Definition-site variance was first investigated in the late 80's [7,2,3] when parametric types were incorporated into object-oriented languages. It has recently experienced a resurgence as newer languages such as Scala [14] and C# [9] chose it as means to support variant subtyping. Perhaps surprisingly, with such a long history, it has only recently been formalized and proven sound in a non-toy setting [8].

Use-site variance was introduced by Thorup and Torgersen [17] in response to the rigidity in class definitions imposed by definition-site variance. The concept was later generalized and formalized by Igarashi and Viroli [10]. The elegance and flexibility of the approach evoked a great deal of enthusiasm, and was quickly introduced into Java [18]. The same flexibility also proved challenging to both researchers and practitioners. The soundness of wildcards in Java has only recently been proven [4], and the implementation of wildcards has been mired in issues [5,15,16].

The work of Viroli and Rimassa [20] attempts to clarify when variance is to be used, introducing concepts of produce/consume, which are an improvement over the write/read view. Our approach offers a generalization and a high-level way to reason soundly about the variance of a type. Other recent work discusses the complex relationship between type-erasure and wildcards [6], as well as the concept of variance at the level of tuning access to a path type in tree-like class definitions [12].

7 Conclusion

This paper presented VarJ, the first formal model for Java with definition-site variance, wildcards, and intricate features such as wildcard capture. VarJ gives a framework for reasoning about the variance of various types (e.g., bounded existential types). We presented theory underlying the assumed variances of the positions that types can occur in (e.g., the upper bound of a method type parameter is contravariant). Thus, our calculus resolves questions that are central in the design of any language involving parametric polymorphism and subtyping.

Acknowledgments. We thank the anonymous ECOOP reviewers for their feedback, Ross Tate for providing examples on the complications of wildcards, Nicholas Cameron for discussions on TameFJ and on performing type inference at runtime, Andrew Kennedy for discussions about the C# formalism with definition-site variance, and Christian Urban for clarifying Barendregt's variable convention. This work was funded by the National Science Foundation under grants CCF-0917774 and CCF-0934631.

References

1. Altidor, J., Huang, S.S., Smaragdakis, Y.: Taming the wildcards: Combining definition- and use-site variance. In: Programming Language Design and Implementation, PLDI (2011)

2. America, P., van der Linden, F.: A parallel object-oriented language with inheritance and subtyping. In: European Conf. on Object-Oriented Programming and Object-Oriented Programming Systems, Languages, and Applications, OOPSLA/ECOOP (1990)
3. Bracha, G., Griswold, D.: Strongtalk: typechecking smalltalk in a production environment. In: Object-Oriented Programming Systems, Languages, and Applications, OOPSLA (1993)
4. Cameron, N., Drossopoulou, S., Ernst, E.: A Model for Java with Wildcards. In: Ryan, M. (ed.) ECOOP 2008. LNCS, vol. 5142, pp. 2–26. Springer, Heidelberg (2008)
5. Chin, W.-N., Craciun, F., Khoo, S.-C., Popeea, C.: A flow-based approach for variant parametric types. In: Object-Oriented Programming Systems, Languages, and Applications, OOPSLA (2006)
6. Cimadamore, M., Viroli, M.: Reifying wildcards in Java using the EGO approach. In: SAC 2007: Proceedings of the 2007 ACM Symposium on Applied Computing (2007)
7. Cook, W.: A proposal for making Eiffel type-safe. In: European Conf. on Object-Oriented Programming, ECOOP (1989)
8. Emir, B., Kennedy, A., Russo, C.V., Yu, D.: Variance and Generalized Constraints for C# Generics. In: Hu, Q. (ed.) ECOOP 2006. LNCS, vol. 4067, pp. 279–303. Springer, Heidelberg (2006)
9. Hejlsberg, A., Wiltamuth, S., Golde, P.: C# Language Specification. Addison-Wesley Longman Publishing Co., Inc., Boston (2003)
10. Igarashi, A., Viroli, M.: On Variance-Based Subtyping for Parametric Types. In: Deng, T. (ed.) ECOOP 2002. LNCS, vol. 2374, pp. 441–469. Springer, Heidelberg (2002)
11. Igarashi, A., Viroli, M.: Variant parametric types: A flexible subtyping scheme for generics. *ACM Trans. Program. Lang. Syst.* 28(5), 795–847 (2006)
12. Igarashi, A., Viroli, M.: Variant path types for scalable extensibility. In: Object-Oriented Programming Systems, Languages, and Applications, OOPSLA (2007)
13. Kennedy, A.J., Pierce, B.C.: On decidability of nominal subtyping with variance, 2006. In: FOOL-WOOD 2007 (2007)
14. Odersky, M.: The Scala Language Specification v 2.8 (2010)
15. Smith, D., Cartwright, R.: Java type inference is broken: can we fix it? In: Object-Oriented Programming Systems, Languages, and Applications, OOPSLA (2008)
16. Tate, R., Leung, A., Lerner, S.: Taming wildcards in Java’s type system. In: Programming Language Design and Implementation, PLDI (2011)
17. Thorup, K.K., Torgersen, M.: Unifying Genericity: Combining the Benefits of Virtual Types and Parameterized Classes. In: Guerraoui, R. (ed.) ECOOP 1999. LNCS, vol. 1628, pp. 186–204. Springer, Heidelberg (1999)
18. Torgersen, M., Hansen, C.P., Ernst, E., von der Ahe, P., Bracha, G., Gafter, N.: Adding wildcards to the Java programming language. In: SAC 2004: Proc. of the 2004 Symposium on Applied Computing (2004)
19. Urban, C., Berghofer, S., Norrish, M.: Barendregt’s Variable Convention in Rule Inductions. In: Pfenning, F. (ed.) CADE 2007. LNCS (LNAI), vol. 4603, pp. 35–50. Springer, Heidelberg (2007)
20. Viroli, M., Rimassa, G.: On access restriction with Java wildcards. *Journal of Object Technology* 4(10), 117–139 (2005)