

A Flow-Based Approach for Variant Parametric Types

Wei-Ngan Chin Florin Craciun Siau-Cheng Khoo Corneliu Popeea

Department of Computer Science, National University of Singapore

{chinwn,craciunm,khoosc,corneliu}@comp.nus.edu.sg

Abstract

A promising approach for type-safe generic codes in the object-oriented paradigm is *variant parametric type*, which allows covariant and contravariant subtyping on fields where appropriate. Previous approaches formalise variant type as a special case of the existential type system. In this paper, we present a new framework based on *flow analysis* and *modular type checking* to provide a simple but accurate model for capturing generic types. Our scheme stands to benefit from past (and future) advances in flow analysis and subtyping constraints. Furthermore, it fully supports casting for variant types with a special reflection mechanism, called *cast capture*, to handle objects with unknown types. We have built a constraint-based type checker and have proven its soundness. We have also successfully annotated a suite of Java libraries and client code with our flow-based variant type system.

Categories and Subject Descriptors D.3.3 [Programming Languages]: Language Constructs and Features—Classes and objects; Polymorphism; Constraints; D.3.1 [Programming Languages]: Formal Definitions and Theory; D.3.2 [Programming Languages]: Language Classifications—Object-oriented languages; F.3.3 [Logics and Meanings of Programs]: Studies of Program Constructs—Object-oriented constructs; Type structure

General Terms Design, Languages, Theory, Verification

Keywords Genericity, Flow Analysis, Variant Parametric Types, Subtyping, Constraints

1. Introduction

Software reuse is an important aspect of software engineering. Traditionally, most mainstream object-oriented (OO) languages, such as Java, C++ and C#, have relied on class subtyping to support reuse (or genericity) via inclusion polymorphism. While this mechanism allows the convenient storage of objects via safe upcast into generic data structure, the converse process of retrieving objects from the same data structure requires downcast testing, which incurs runtime overheads and is possibly unsafe.

To address the shortcomings of inclusion polymorphism, there have been several recent proposals (amongst the Java [3] and C# [19] communities) for parametric types to be supported. Here, each class c is allowed to carry a list of type parameters for its fields, e.g., $c(t_1, \dots, t_n)$, whereby the type of each field can either be

instantiated or left as a type variable. Below are two classes whose fields have been parameterised:

```
class Cell<A> {
  A fst; ... }
class Pair<A,B> extends Cell<A> {
  B snd; ... }
```

With such parameterised class declarations, we may then define specialised instances, such as `Cell<Int>`, `Cell<Float>` or `Pair<Int,Num>`, which contain more specific type information for the fields of each class instance. Though parametric types can co-exist with class subtyping, pointwise matching of the respective fields is required. For example, the subtyping relation (denoted by $<:$) `Pair<t1,t2> <: Cell<t3>` is allowed only when `Pair<t1,t2> <: Cell<t3>` and `t1=t3`. The latter condition is for pointwise matching of the common field. Similarly, `Pair<t1,t2> <: Pair<t3,t4>` holds, provided `t1=t3` and `t2=t4`. Pointwise matching (invariant subtyping) is required because field reading and field writing are based on opposite flows that change the directions of subtyping. This requirement limits the reusability of programs based on parametric types.

To address this shortcoming, Igarashi and Viroli [17] developed a new variant parametric type system (or variant type, in short) to improve the subtyping of generic structures, depending on how the fields are being accessed. Let c denote a class with one type parameter. Let o denote an object of variant type $c\langle\alpha_1 t_1\rangle$ while v denotes a location of variant type $c\langle\alpha_2 t_2\rangle$, into which o is to pass. Each variant type $c\langle\alpha t\rangle$ has a variance α (see Section 3.1) attached to its field to indicate how the field is to be accessed. A field that is subject to read-only access via reference of v (denoted by $\alpha_2 = \oplus$) may be supported by covariant subtyping. That is, $c\langle\alpha_1 t_1\rangle <: c\langle\oplus t_2\rangle$ if $\alpha_1 <: \oplus$ and $t_1 <: t_2$. Conversely, a field that is subject to write-only access via reference of v (denoted by $\alpha_2 = \ominus$) may be supported by contravariant subtyping. That is, $c\langle\alpha_1 t_1\rangle <: c\langle\ominus t_2\rangle$ if $\alpha_1 <: \ominus$ and $t_2 <: t_1$. Also, a field that is subject to both read and write accesses via reference of v (denoted by $\alpha_2 = \odot$) must be supported by invariant subtyping. That is, $c\langle\alpha_1 t_1\rangle <: c\langle\odot t_2\rangle$ if $\alpha_1 <: \odot$ and $t_1 <: t_2 \wedge t_2 <: t_1$. Lastly, if a field is not accessed via reference of v (denoted by $\alpha_2 = \otimes$), we can use bivariate subtyping. That is, we support $c\langle\alpha_1 t_1\rangle <: c\langle\otimes t_2\rangle$ for any t_1 and t_2 .

Variant types give a much richer subtyping hierarchy than parameterised types do. Figure 1 illustrates some variant types for `Cell` objects and their places in the subtyping hierarchy. Note that \rightarrow denotes a subtyping relation in the graph. Also, `Cell<⊗t>`, `Cell<⊕Object>` and `Cell<⊖⊥>` are equivalent to each other while `Cell<⊙Num>`, `Cell<⊙Float>` and `Cell<⊙Int>` are unrelated. Note that \perp denotes the type of *null* values which can be assigned into any class type. However, each `Cell<⊙t>` is a subtype of both `Cell<⊕t>` and `Cell<⊖t>`. Also, types of the form `Cell<⊕t>` and `Cell<⊖t>` have a subtyping hierarchy based on covariance and contravariance, respectively.

The benefits of variant typing have been known for some time. However, early proposals have attached access rights to the fields

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

OOPSLA'06 October 22–26, 2006, Portland, Oregon, USA.
Copyright © 2006 ACM 1-59593-348-4/06/0010...\$5.00.

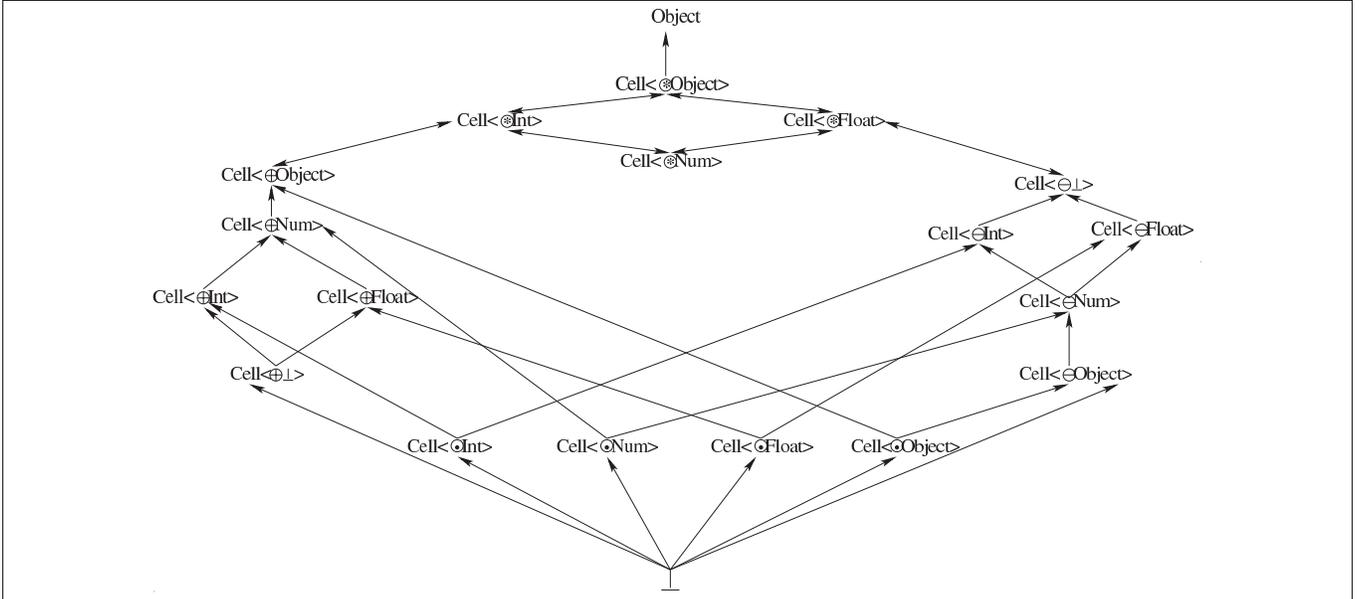


Figure 1. A Rich Subtyping Hierarchy

of each class declaration. This mechanism is known as *declaration-site variance* and is shown in the following example:

```
class DSCell(A) {
  ⊕A fst;
  A getFst() { return fst; }
  void setFst(A x) { fst=x; } }
```

The field `fst` is declared read only using the variance \oplus . Consequently, the method `setFst` cannot be invoked. Using the concept of structured virtual type, Thorup and Torgersen [35] were the first to link access rights and covariant subtyping to the fields of each *use of a class* rather than the class declaration itself. This *use-site variance* mechanism is much more flexible than previous mechanisms based on *declaration-site variance*. In the following example, the access to the field `fst` is governed by the variance variable α . A reference of type `USCell<⊕Int>` allows read-only access, while a reference of type `USCell<⊙Int>` allows read-write access to the field `fst`.

```
class USCell(αB) {
  αB fst;
  B getFst() { return fst; }
  void setFst(B x) { fst=x; } }
```

Later, Igarashi and Viroli extended this concept to support contra- and bi-variance [17]. They formalised the variant type system by mapping it into a corresponding *existential type* system [7, 21, 22]. A recent proposal by Sun Microsystems for generics in Java 1.5 [37] supports *wildcard type* based on an improvement of Igarashi and Viroli’s variant type system, but it is still viewed as a special case of the existential type system with subtyping. However, a more general version of existential type system, called System F_{\leq} , has undecidable subtyping [28], while the decidability of Igarashi and Viroli’s variant type system, remains an open problem [18].

In this paper, we propose a new approach for the variant parametric type system that is based on the mechanism of flow analysis. Our flow analysis captures value flows via subtyping constraints. A major benefit of this approach is the considerable knowledge in flow analysis that has been accumulated in the recent past

[25, 32, 38, 15, 16, 34, 26]. In particular, to support modular type-checking, we require non-structural subtype entailment of the form $\forall \bar{v}(C_1 \implies \exists \bar{w}C_2)$, where C_1, C_2 are subtyping constraints while \bar{v}, \bar{w} are sets of type variables. These constraints are non-structural as we use $\perp <: t <: Object$, to support the OO class inheritance mechanism. While the decidability of non-structural subtype entailment remains an open problem, there exist sound approximations that use constraint simplification and induction techniques [32, 38]. Our work is built on top of sound but practical solutions to subtyping (flow) constraints, and we have developed a systematic framework for the variant parametric type system with the following new features:

- Our framework is based on *flow analysis* which can concisely and intuitively capture flow of values on a per method basis (Section 3). We use variance annotations primarily to predict the flows of values, and not for access control. We also provide special considerations for two type values. A value of `Object` type can always flow out from any location while a null value of \perp type can always flow into any location.
- We augment our generic type system with *intersection type* to help capture information flow more accurately. An intersection type $\tau_1 \& \tau_2$ denotes a type with both the properties of τ_1 and τ_2 . Such types are important for languages with multiple inheritance (such as Java via its interface mechanism), and can accurately capture the flow of objects with their expected field accesses.
- Our approach is based on *modular type checking* (Section 5). Each method is specified with a flow constraint (and variant types) that is used to predict the value flows that may occur in the method’s body. We verify each method separately to ensure that the predicted accesses, flow constraint and variant typings are efficiently and safely checked.
- We advocate the support of *downcast to arbitrary variant types* (Section 6). With this mechanism is a novel *cast capture* that uses a reflection technique to deal with values of unknown type. Cast capture has helped improve the generic implementation of several JDK 1.5 libraries.

- We present a soundness theorem and a *variant type checker*. We have successfully applied our prototype to a suite of Java libraries and client codes (Section 8). On average, we are able to eliminate 87.9% of the casts from non-generic Java 1.4 application code, that means 12.9% more casts than wildcard-generic Java 1.5 application code.

Our goal is to strive for type-safe OO programs with better genericity via a modular flow-based approach to variant parametric type system. Next, we explain our approach with the help of some examples.

2. Better Genericity

The main goal of genericity is to support highly reusable software components. To allow this to happen in a type-safe way, we should strive to provide type descriptions that are concise, understandable, general and accurate. Specifically, each well-typed generic program should be accurately identified where possible. As a side benefit, we are able to track type information in a precise manner, allowing redundant cast operations to be eliminated where possible. In this section, we examine the key aspects for which our approach based on flow analysis makes improvements over existing approaches based on existential types. Some of these improvements may not be peculiar to the flow-based approach, but they were gradually developed starting from a different view point.

2.1 Intersection Type

Parametric type systems use number of cast operations eliminated as a measure of accuracy [10, 14]. As it turns out, there may be competing decisions on what types to use for certain cast operations to be eliminated. The following example from [10] illustrates:

```
class B1 extends A implements I { ... }
class B2 extends A implements I { ... }
void foo(Boolean b) {
  Cell cb1 = new Cell(new B1());
  Cell cb2 = new Cell(new B2());
  Cell c = b ? cb1 : cb2;
  A a = (A) c.get();
  I i = (I) c.get();
  B1 b1 = (B1) cb1.get();
  B2 b2 = (B2) cb2.get(); }

```

This program contains four cast operations. With the help of parametric types, Donovan et al. [10] suggested three sets of possible types, each with a different subset of casts eliminated, as summarised below:

Types of Variables			Casts Eliminated			
cb1	cb2	c	(A)	(I)	(B1)	(B2)
Cell⟨B1⟩	Cell⟨B2⟩	Cell			√	√
Cell⟨A⟩	Cell⟨A⟩	Cell⟨A⟩	√			
Cell⟨I⟩	Cell⟨I⟩	Cell⟨I⟩		√		

Note that `Cell` denotes a raw type where nothing is known of its components. Hence, only `Object` values are statically retrievable from it. Raw type was originally proposed in [3] for backwards compatibility, and it is the basis for generic typing through inclusion polymorphism. However, none of the three proposed solutions are able to eliminate all four casts. This indicates that parametric typing is not expressive enough to capture generic type for such programs. There are two possible improvements. First, note that the fields of `cb1`, `cb2` and `c` are subject to read-only accesses, and not modified in the program fragment. We can therefore provide covariant annotations to the fields of these variables, and obtain two possible outcomes, each with three casts eliminated:

cb1	cb2	c	(A)	(I)	(B1)	(B2)
Cell⟨⊕B1⟩	Cell⟨⊕B2⟩	Cell⟨⊕A⟩	√		√	√
Cell⟨⊕B1⟩	Cell⟨⊕B2⟩	Cell⟨⊕I⟩		√	√	√

Second, both classes `B1` and `B2` have supertypes `A` and `I` in common. To exploit this, we can use an intersection type parameter in `Cell⟨⊕(A&I)⟩` to describe the variable `c`. In a lattice of type values, an intersection type `A&I` essentially defines the greatest lower bound of `A` and `I`. With this, all four casts can now be eliminated in our new solution to genericity, as shown below:

cb1	cb2	c	(A)	(I)	(B1)	(B2)
Cell⟨⊕B1⟩	Cell⟨⊕B2⟩	Cell⟨⊕A&I⟩	√	√	√	√

Note that the above example cannot be coded in Java 1.5 syntax. Java 1.5 does not allow the use of intersection types for local variable declaration, field declaration or method argument/return types. Intersection types can be used only as upper bounds for a method type parameter.

2.2 Modular Flow Specification

Another important principle for better genericity is that type description should be designed in a *modular* fashion (on a per method basis). Type annotations appearing in the method header should depend only on the method body while each call site should be a specific instance of the method's type declaration. This principle is important for efficient type checking and ease of type annotation. Specifically, for each instance method, we provide the following method declaration:

$$t \mid t_0 \text{ mn}(t_1 v_1, \dots, t_n v_n) \text{ where } \psi \{ \dots \}$$

A separate annotation “ $t \mid$ ” is added at the beginning of each method's declaration to capture the variance of the implicit `this` parameter. This separate annotation (omitted in previous works, such as [17, 37]) allows us to capture the behaviour of each method, independent of its class declaration. Note that ψ captures the expected value flows of each method's body in terms of type of the parameters (t_1, \dots, t_n) , result (t_0) , and receiver (t) . We support modular type checking by localising type variables which are not present in the type of parameters, result and receiver. A previous approach [17] relies on the existential open/close mechanism for the receiver parameter to determine if the receiver parameter is of suitable variance while other parameters are checked via subtyping. In contrast, we achieve uniform treatment for all parameters.

To illustrate the modular type annotation mechanism, consider three method declarations for the `Pair` class:

```
class Pair<A,B> extends Cell<A> {
  B snd;
  Pair<⊗,⊕Y> | Y getSnd()
  {return this.snd;}
  Pair<⊗,⊖Y> | void setSnd(Y v)
  {this.snd=v;}
  Pair<⊗,⊗&W> | Pair<⊖W,⊖W> dup()
  {return new Pair<W,W>(this,this);} }

```

First, note that `getSnd` will read the second field while `setSnd` will write to it. Because of these effects, we may apply covariant (\oplus) and contravariant (\ominus) subtypings to the second component of the `Pair` object for `getSnd` and `setSnd`, respectively. Second, bivariant (\otimes) subtyping is allowed for the unaccessed component of the `Pair` object for both methods. As a shorthand, we may write \otimes to denote $\otimes t$ since all bivariant types are equivalent. Note that Y from `getSnd` and Y from `setSnd` denote different type variables treated independently by our modular type checker.

The third method is an interesting application of intersection type. The method itself does not access the fields of the `this` parameter, which escapes into the two fields of the method's `Pair` result. To capture this value flow, we declare an intersection type `Pair(⊗, ⊗)&W` for the `this` parameter. The type `Pair(⊗, ⊗)` is to acknowledge that we have a `Pair` object whose fields are not accessed by the current `dup` method. A type variable `w` helps indicate that this parameter will escape into the fields of the result with type `Pair(⊙w, ⊙w)`. This flow allows the variant type of `w` to flow into the two fields of the output `Pair`. Hence, for a given receiver of type `t`, we have `t<:Pair(⊗, ⊗)` and `t<:w`. Possible candidates for the type `t` are `Pair(⊕X, ⊕Y)` or `Pair(⊕X, ⊖Y)`, etc. In contrast, if we use the following type suggested in [17]:

```
Pair(⊙X, ⊙Y) | Pair(⊙Pair(⊙X, ⊙Y), ⊙Pair(⊙X, ⊙Y)) dup()
```

we require `t=Pair(⊙X, ⊙Y)` or `t=⊥`, which restricts the possible uses of the method. One way to improve this situation is to duplicate the `dup` method for different scenarios, as shown below:

```
Pair(⊕X, ⊕Y) | Pair(⊙Pair(⊕X, ⊕Y), ⊙Pair(⊕X, ⊕Y)) dup()
Pair(⊕X, ⊖Y) | Pair(⊙Pair(⊕X, ⊖Y), ⊙Pair(⊕X, ⊖Y)) dup()
Pair(⊙X, ⊖Y) | Pair(⊙Pair(⊙X, ⊖Y), ⊙Pair(⊙X, ⊖Y)) dup()
```

However, such duplications go against the goal of genericity. On the other hand, our solution with intersection types can improve genericity by allowing value flows to be accurately captured.

2.3 Avoiding F-Bounds where Possible

One feature that adds to the expressivity of bounded existential type is the use of F-bounds [5] which effectively capture recursive constraints of the form $T <: C \langle \dots, T, \dots \rangle$ where T is a type variable and C is a class name. While the designers of Java 1.5 consider this feature to be significant and useful [37], it is also a source of complication as reported recently in [20]. In particular, F-bound together with existential type is a source of undecidability for System F_{\leq} which caused an earlier implementation of Java 1.5 to fail in accepting some programs with F-bounds that were actually type-safe (as first reported in [20]). Subsequent improvements in Java 1.6 have removed the reported errors, but the decidability of its type system remains an open problem.

While the flow-based approach that we advocate also supports recursive flow constraints (if the inductive mechanism of [32, 38] is used), our pragmatic philosophy is to avoid F-bounds whenever it is possible to do so.

As an example of F-bound, consider the following definition of the `Comparable` interface for Java 1.5:

```
interface Comparable(T) {
    int compareTo(T o);
}
```

Here, class parameter `T` is being used to capture the parameter of the method `compareTo`. As this parameter is required to be a subtype of `Comparable` itself, F-bound of the form $T <: Comparable(\ominus T)$ is usually needed when `Comparable` is used, as shown in the next example:

```
class Collections {
    <T extends Comparable(? super T)> static T max
    (Collection(? extends T) col) { ... }
}
```

In our flow-based approach, the current philosophy is to capture the value flows of each method independently. Hence, we have chosen to capture the value flow and subtyping relation directly for each method instead, as shown below for our definition of `Comparable`:

```
interface Comparable(A) {
    Comparable(\ominus T) | int compareTo(T o);
}
```

Based on this definition, we can write the `max` method, as follows:

```
class Collections {
    static T max(Collection(\oplus T) col)
    where T<:Comparable(\ominus T) { ... }
}
```

This alternative is equivalent to the earlier Java 1.5 definition.

We also support a simpler way, to express `Comparable` interface, as follows:

```
interface Comparable {
    S & Comparable | int compareTo(T o)
    where T<:Comparable & T<:S ;
}
```

The use of this definition does not require any F-bound, but it is more restrictive than Java 1.5 definition of `Comparable` interface.

Another potential use of F-bound occurs for recursive fields of class declarations. An example is the following recursive `List` class:

```
class List(A,B) extends Object
    where B<:List(A,B) {
    A val;
    B next; ...
}
```

This solution uses an F-bound $B <: List(A, B)$ that makes constraint solving more complex [32]. However, in our system we may choose to avoid the recursive constraint from the invariant of the class `List` by leaving the recursive `next` field with an incomplete variance \ominus , as follows:

```
class List(A) extends Object {
    A val;
    \ominus List(A) next; ...
}
```

The variance of the `next` field is incomplete at its declaration site and can be promoted to either \ominus or \oplus , depending on how its underlying type parameter `List(A)` is being instantiated at the use site. This type promotion process is elaborated later in Section 4.2, and can be used to avoid F-bound, where possible.

2.4 Avoiding Existential Type Always

It has been generally acknowledged that existential type is useful for describing data types whose implementation details can be made abstract. This aspect is closely related to the use of bivariate type $\otimes t$ where the underlying type t is unknown and may be assumed to be of any type. While no-access is one way to enforce bivariate type, it is also possible to use the open/close mechanism of existential type system to describe situations where implementation details can be made abstract. A typical example is the `copy` operation on two elements of a vector that was highlighted in [18], and reproduced below:

```
void copy(Vector(\otimes) x, int i, int j) {
    open x as [Y,y] in
    y.setElementAt(y.elementAt(i), j)
}
```

The above code opens the bivariate type of `x` as an object bound to variable `y` with an abstract type `Y`. As all elements of each vector are of the same `Y` type, we may safely copy a value from one position of the vector into another position, without knowing the

actual underlying type. The close correspondence between existential type and bivariate type is a primary reason why Igarashi and Viroli considered existential type system as the underlying model for their variant parametric type system.

However, the designers of Java 1.5 considered the open/close mechanism of existential type system to be somewhat restrictive [36]. They have therefore proposed a relaxation to open each expression as an existential type by associating it with a global type variable *without* a corresponding close operation. This use is similar to the flow-based approach where each parameter (or local variable) is regarded as a location where values may flow in and/or out. Nevertheless, in the context of existential type system, such relaxation might possibly be unsound since each existential type may in fact correspond to contradicting type values. This is possibly why correctness proof is yet to be completed (as of [36]), even though a full-scale implementation for wildcard type system has already been released for public use.

Furthermore, Java 1.5 relied on polymorphic (generic) type system for selected methods to capture situations where invariant type appears necessary, as shown by the following example:

```
<T> void docopy(Vector<T> x, int i, int j) {
    T tmp = x.elementAt(i);
    x.setElementAt(tmp, j)
}
```

Through a wildcard capture mechanism, it is possible to provide a method with bivariate parameter, as shown below:

```
void copy(Vector<?> x, int i, int j) {
    docopy(x, i, j);
}
```

Note that wildcard type of x has been captured by the global T type variable. Again, the open/close mechanism is averted, even though the underlying system is still based on bounded existential type system.

Our current philosophy is to avoid existential type system altogether. To capture the effect of an unknown abstract type, we have introduced a casting mechanism that is able to capture the underlying type of an object via a fresh type variable. We refer to this as a *cast capture* technique which is elaborated in more details in Section 6. The same `copy` method can be re-written with a casting of the x parameter from bivariate type $\text{Vector}(\otimes)$ to an invariant type $\text{Vector}(\odot T)$. In the process, T is used to capture the unknown type, as shown below:

```
void copy(Vector<\otimes> x, int i, int j) {
    Vector<\odot S> w;
    {w = (Vector<\odot T>) x;
      w.setElementAt(w.elementAt(i), j) }
}
```

While this cast capture construct may look like a syntactic sugar for the open/close mechanism, we stress that it is part of a more general mechanism that can take an arbitrary type as source (instead of a bivariate type) for casting into another arbitrary type as target (instead of an invariant type). A cast for a c_1 -object into an invariant type of the form $c_2((\odot t)^*)$ where $c_1 <: c_2$ is always safe since every object is built using an invariant type. Furthermore, cast-capture is a runtime mechanism while open-close is a type-related operation to expose an obtained type at compile-time. Our cast capture mechanism using reflection is more general as it can capture type values at runtime, and also support a mix of cast capture and cast testing. In our formulation of variant parametric type system, the flow-based approach with casting has therefore avoided the need for existential type system altogether.

Some readers may contend that the casting mechanism is the prerogative of programmers and may be too burdensome to write.

While this is so, we believe that there is still scope for automatic insertion of safe casts to invariant type (in a spirit similar to automatic type coercion) that is consistent with each user program.

3. Variance via Flow Analysis

A central feature of our proposed approach is the focus on flow analysis. Variance annotations are used to support the analysis of value flows to capture more accurate generic types, whereby suitable field subtypings (covariance and contravariance) are facilitated where possible.

We highlight the expressiveness of variant types through some more examples in Figure 2. Apart from a generic $\text{Vector}(\text{A})$ class declaration, we provide a number of static methods to highlight how flow analysis may assist in the formulation of generic types. In the `copyVec` method, the elements from a first vector $\text{Vector}(\oplus X)$ are copied into a second vector $\text{Vector}(\ominus Y)$, while a constraint $X <: Y$ captures the direction of the value flow.

```
class Vector<A> extends Collection<A> {
    Vector<\otimes> | int size() {...}
    Vector<\oplus X> | X elementAt(int i) {...}
    Vector<\ominus X> | void setElementAt(X v, int i) {...}
}
void copyVec(Vector<\oplus X> v, Vector<\ominus Y> w,
    int start) where X<:Y {
    for(int i=0; i<v.size(); i++)
        w.setElementAt(v.elementAt(i), i+start);
}
void copyNestVec(Vector<\oplus Vector<\oplus X>>v,
    Vector<\ominus Y> w) where X<:Y {
    int pos=0;
    for(int i=0; i<v.size(); i++) {
        Vector<\oplus Z> s=v.elementAt(i);
        if (pos+s.size()<w.size())
            {copyVec(s,w,pos); pos +=s.size(); }
    }
}
void clearVec(Vector<\ominus \perp> v) {
    for(int i=0; i<v.size(); i++)
        v.setElementAt(null, i);
}
Vector<\odot Z> merge(Vector<\oplus X> v, Vector<\oplus Y> w)
    where X<:Z^Y<:Z
{...}
Vector<\odot Pair<\odot X, \odot Z>> join(Vector<\oplus Pair<\oplus X, \oplus Y>> v,
    Vector<\oplus Pair<\oplus Y, \oplus Z>> w)
{...}
void swap(Pair<\odot X, \odot Y> p) where X<:Y^Y<:X {
    T t=p.fst; p.fst=p.snd; p.snd=t;
}
}
```

Figure 2. Examples with Variant Types

Method `copyNestVec` copies from a nested vector of type $\text{Vector}(\oplus \text{Vector}(\oplus X))$ into a second vector $\text{Vector}(\ominus Y)$ with flow constraint $X <: Y$. This code remains highly generic as it uses covariant and contravariant subtypings. The next example shows how we use a special type \perp to indicate that null values will be written into the vector. Given that $\text{Vector}(\ominus \perp)$ is high up in the class hierarchy, this method is rather generic as we can supply *any* vector as its argument.

We also provide method headers for `merge` and `join`. From the type annotation of `merge`, we can tell that values from the first two vectors are retrieved, and then they flow into a new result vector. For the `join` method, we retrieve values from the two vectors $\text{Vector}(\oplus \text{Pair}(\oplus X, \oplus Y))$ and $\text{Vector}(\oplus \text{Pair}(\oplus Y, \oplus Z))$ before building a new vector $\text{Vector}(\odot \text{Pair}(\odot X, \odot Z))$ that is joined

on the Y type. The result's invariant type offers a strong post-condition with read/write capability.

For the `swap` method, the two fields of a `Pair` object are swapped. Due to both reading and writing, we require the invariant type `Pair`($\odot X, \odot Y$) and the expected value flow: $X <: Y \wedge Y <: X$. Based on the flows from the three assignments of the `swap` body, we may obtain the following constraints: $\odot X <: \oplus T$, $\odot Y <: \oplus X$ and $\odot T <: \oplus Y$, where T is a local type variable (using type rules in Section 5.1). These constraints are simplified to obtain the following collected flow for the method body: $X <: T \wedge Y <: X \wedge T <: Y$. The `swap` method type checks as the expected flow implies the collected flow:

$$\forall X, Y. (X <: Y \wedge Y <: X \implies \exists T. (X <: T \wedge Y <: X \wedge T <: Y))$$

Note that the local type variable T is existentially quantified, while type variables X, Y from method parameters are universally quantified.

3.1 Improved Variant Subtyping

Variant parametric type τ consists of a variance α and a type t . Its grammar is introduced in Figure 4. We use variance annotations \odot, \oplus, \ominus and \otimes , which correspond to read-write access, read-only access, write-only access, and no-access, respectively. These annotations are ordered by the following relation that is denoted by $<:_{\alpha}$ but abbreviated to $<:$ below:

$$\begin{array}{c} \odot <: \oplus \quad \odot <: \ominus \quad \oplus <: \otimes \quad \ominus <: \otimes \\ \hline \frac{\alpha_1 <: \alpha_2 \quad \alpha_2 <: \alpha_3}{\alpha_1 <: \alpha_3} \quad \alpha <: \alpha \end{array}$$

A type t is either a type variable v_t , a variant parametric class $c(\tau_1, \dots, \tau_n)$, the bottom type \perp or an intersection type $t_1 \& t_2$. The bottom type is used to hold the `null` value.

We allow finite intersections of types through the type operator $\&$. Semantically, $t_1 \& t_2$ denotes the set of objects satisfying the interface specification of both t_1 and t_2 . In a lattice of type values with partial order defined by class inheritance (through `extends`) and interface mechanism (through `implements`), $t_1 \& t_2$ defines the greatest lower bound of t_1 and t_2 . Our intersection types are similar to the compound types proposed in [4]. Specifically, they can be of the form $[t_1 \&] t_2 \& \dots \& t_n [\& W]$, where t_1 is a class, t_2, \dots, t_n are interfaces, and W is a type variable.

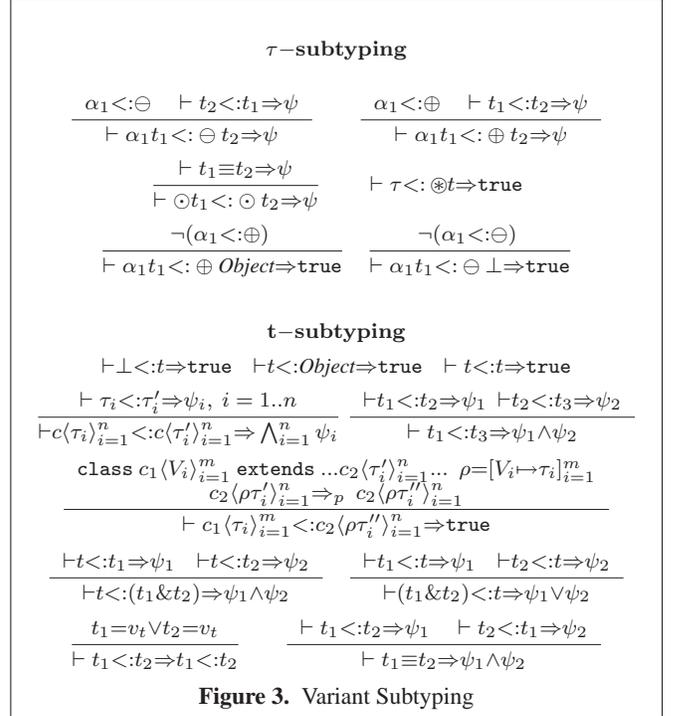
In our system, variant parametric types are used to support flow analysis rather than access controls. As we focus on value flows at each method boundary, we apply variance annotations primarily to fields. The outermost variance of local variables is always \odot . For fields, variance annotations are used to support covariant or contravariant subtyping where possible.

The subtyping relations are denoted by $<:_{\tau}$ and $<:_t$, both abbreviated to $<:$ as follows:

$$\vdash \tau_1 <: \tau_2 \Rightarrow \psi \quad \vdash t_1 <: t_2 \Rightarrow \psi$$

The resulting constraints ψ (see Figure 4 for their grammar) are kept in a disjunctive normal form. Instead of proving each subtyping directly, we collect a set of subtyping constraints ψ via τ -subtyping and t -subtyping in Figure 3.

The first four τ -subtyping rules support contravariance, covariance, invariance and bivariate, respectively. The invariant case generates a constraint from the semantical equivalence of the two types ($t_1 \equiv t_2$). Unlike the subtyping rule of Igarashi and Viroli [17], our improved mechanism handles two special values in the subtyping hierarchy, namely \perp (for type of null) and `Object` (for top of class hierarchy). These two types are special in that it is always safe to write a null (of \perp type) into any location (even if it has been marked for read-only access), and it is safe to read an `Object` value from any location (even if it has been marked for write-only



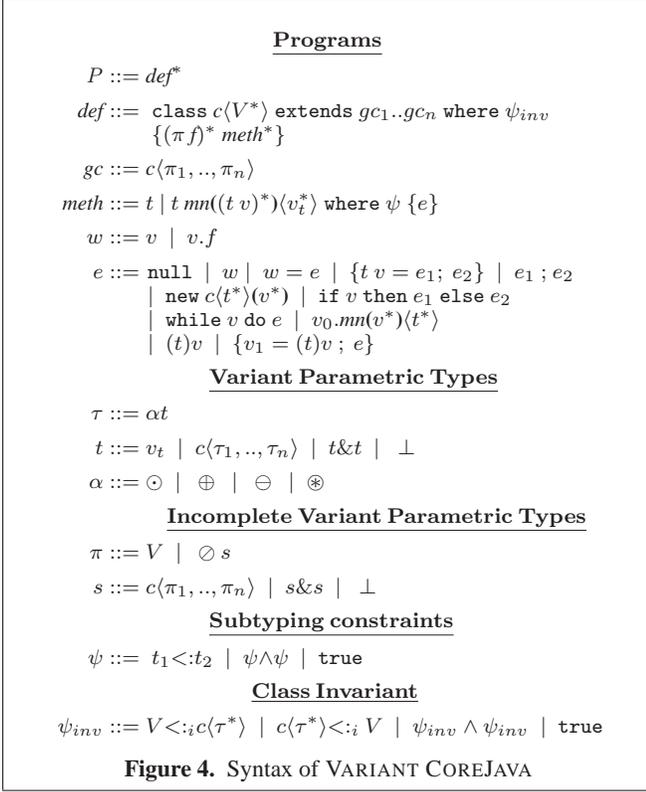
access). We may also cast any type τ to either $\oplus \text{Object}$ or $\ominus \perp$ as it is always safe to read an object or write a null value. This mechanism is implemented by the last two τ -subtyping rules.

In the second part of Figure 3, the first two t -subtyping rules handle the bottom and top of the hierarchy. Subtyping between types of the same class is decomposed structurally by the fourth rule. The next two rules describe transitivity and the class inheritance relation. The class inheritance rule uses type promotion mechanism that is described later in Section 4.1 Intersection types satisfy the subtyping relations as in [29]. Subtyping relations that contain type variables are not simplified further and preserved in the resulting constraint. Semantic equality ($t_1 \equiv t_2$) is given by the last t -subtyping rule. In summary, from the subtyping relations between types, we generate a set of subtyping constraints (on type variables). Note that in the following sections, we will use $\tau_1 <: \tau_2$ as an abbreviation for ψ , where $\vdash \tau_1 <: \tau_2 \Rightarrow \psi$.

4. Core Language

We introduce a core language to ease the formulation of static and dynamic semantics. This language can be viewed as a result of translation from full Java language prior to type checking. For ease of presentation, we omit features that are related to static methods, exception handling, concurrency and inner classes. (Our implementation handles all features of the Java language.)

Our core language is named Variant CoreJava, and summarised in Figure 4. We use the suffix notation g^* to denote a list of (zero or more) distinct syntactic terms that are suitably separated. Both class and interface declarations are supported using the same syntactic grammar term `def`. As with Java, the main difference is that interface definitions do not have fields, and are defined using abstract methods (without body). Furthermore, while we support multiple inheritance, it is of the same restricted kind as that supported by the Java language. Each class may extend from only a single superclass but may implement multiple interfaces. In our language, the declaration `class c(V*) extends g1..gn` assumes that g_{c_1} is a



class while $gc_2..gc_n$ are essentially interfaces (implements is also represented by `extends` for easy presentation). Each class declaration captures a class invariant ψ_{inv} that is expected to hold for all newly constructed objects of the class. This is being used to specify suitable class lower and/or upper bounds for type variables. Since our system is based on use-site variance, the class fields types and the arguments of class inheritance have incomplete variance at declaration-site (denoted by π and V). Section 4.1 describes the annotations of class declarations with incomplete variant parametric types.

Each method declaration $meth$ contains a constraint ψ which captures the expected value flows for its type variables. It also specifies method type parameters $\langle v_t^* \rangle$ in order to support modular type checking. This set of type variables is automatically inserted by our compiler.

We use an expression-oriented language, where method body is denoted by e . Local variable declaration is supported by block structure of the form: $\{t v = e_1; e_2\}$, with e_2 denoting its result. Each object is always built with an invariant type $c\langle \odot t^* \rangle$ via the construct `new c⟨t*⟩(v*)`. Our core language also supports a full casting mechanism via $(t)v$, where t can be an arbitrary variant type. In addition, we support a novel cast capture mechanism via $\{v_1 = (t)v; e\}$, where t is an invariant type with unknown type variables that may be captured at runtime and used in e . These special features will be described in more detail in Section 6.

For simplicity of presentation, our core language represents primitive types (such as `void`, `bool`) by their corresponding classes (such as `Void`, `Bool`). In our implementation, we handle primitive types directly, as elaborated in Section 9. For soundness reasons, we treat arrays in the same way as other classes (unlike Java 1.5, which assumes arrays to be covariant).

In the subtyping constraints, disjunction is supported internally as it may be generated by subtyping relation for intersection types.

4.1 Class Parameterisation and Inheritance

For class declarations, an important decision is which fields are to be parameterised and how the class inheritance mechanism is to be supported. In general, each class declaration should be written in the following manner:

```

class c1⟨V1..Vn⟩ extends c2⟨π1..πs⟩ where ψinv {
  π1 f1;
  ...
  πm fm; ...
}

```

where each $\{V_i\}_{i=1}^n$ originates either from the fields of the current class $\{\pi_i\}_{i=1}^m$ or from the arguments of its superclass, $\{\hat{\pi}_i\}_{i=1}^s$. $\{V_i\}_{i=1}^n$ are variables corresponding to types with variance. For instance, the following non-generic declarations of `Cell` and `Pair` classes:

```

class Cell {
  Object fst; ... }
class Pair extends Cell {
  Object snd; ...
}

```

can be parameterized as:

```

class Cell⟨A⟩ {
  A fst; ... }
class Pair⟨A,B⟩ extends Cell⟨A⟩ {
  B snd; ...
}

```

The variance of the fields `fst` and `snd` is governed by the variables A and B . Given the type `Pair⟨⊕Int, ⊖Int⟩`, the field `fst` is covariant and the field `snd` is contravariant.

4.2 Type Promotion

There are some situations where the variance of a class field cannot be specified at use site. In the following example, the variance of the field `sndP` does not have any correspondence in the class parameters A, B, C and remains unknown after instantiation of these parameters.

```

class Triple⟨A,B,C⟩ extends Cell⟨A⟩ {
  Pair⟨B,C⟩ sndP; ...
}

```

The compiler inserts a special variance marker \odot to represent the unknown variance of field `sndP`:

```

class Triple⟨A,B,C⟩ extends Cell⟨A⟩ {
  ∅Pair⟨B,C⟩ sndP; ...
}

```

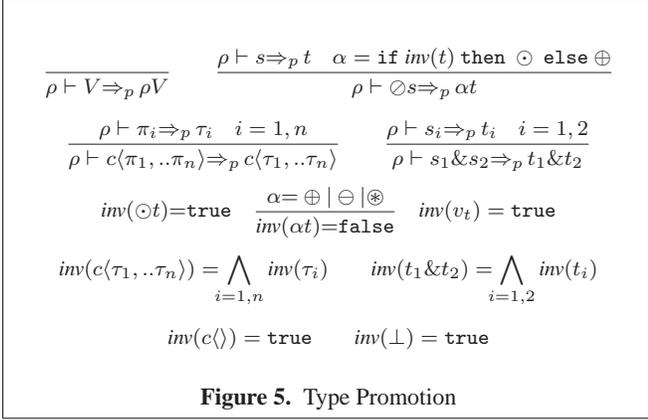
Note that the source program does not contain any variance markers. We use them to explain how incomplete (or unknown) variance of variant parametric types are computed to either \oplus or \odot . This process is known as *type promotion* and can be used for incomplete variant parametric types from field declarations and arguments of class inheritance.

The type promotion is defined using the relations

$$\rho \vdash \pi \Rightarrow_p \tau \quad \rho \vdash s \Rightarrow_p t$$

where ρ is a substitution $[V \mapsto \tau]$ from class declaration parameters V to variant parametric types τ . The types π and s may contain unknown variance \odot . The rules are described in Figure 5.

The second rule promotes the unknown variance \odot to either \oplus or \odot depending on the predicate $inv(t)$ where t is the type obtained after substitution. Predicate $inv(t)$ returns `true`, when all



variances from τ (if any) are \odot and false otherwise. Given $\text{Triple}(\oplus \text{Int}, \oplus \text{Int}, \oplus \text{Int})$, the type of field `sndP` is computed as follows: $\rho \vdash \odot \text{Pair}(\text{B}, \text{C}) \Rightarrow_p \oplus \text{Pair}(\oplus \text{Int}, \oplus \text{Int})$ where $\rho = [\text{A} \mapsto \oplus \text{Int}, \text{B} \mapsto \oplus \text{Int}, \text{C} \mapsto \oplus \text{Int}]$. As another example, given $\text{Triple}(\oplus \text{Int}, \odot \text{Int}, \odot \text{Int})$, the type of field `sndP` is computed as follows: $\rho \vdash \odot \text{Pair}(\text{B}, \text{C}) \Rightarrow_p \odot \text{Pair}(\odot \text{Int}, \odot \text{Int})$ where $\rho = [\text{A} \mapsto \oplus \text{Int}, \text{B} \mapsto \odot \text{Int}, \text{C} \mapsto \odot \text{Int}]$.

Another application of type promotion is for recursive fields of a class. The recursive field `next` of the class `List` has an incomplete variance \odot as follows:

```

class List(A) extends Object {
  A val;
  \odot List(A) next; ...
}

```

The variance of the field `next` is incomplete at its declaration site and can be promoted to either \odot or \oplus , depending on how its underlying type parameter `List(A)` is being instantiated at the use site. For example, when `A` is instantiated to $\odot \text{X}$, the variance of the `next` field will be promoted to \oplus via $\rho \vdash \odot \text{List}(\text{A}) \Rightarrow_p \oplus \text{List}(\odot \text{X})$, where $\rho = [\text{A} \mapsto \odot \text{X}]$. On the other hand, if `A` is instantiated to $\odot \text{X}$, then $\rho = [\text{A} \mapsto \odot \text{X}]$ and the variance of the `next` fields is instantiated to $\odot \text{X}$ as follows: $\rho \vdash \odot \text{List}(\text{A}) \Rightarrow_p \odot \text{List}(\odot \text{X})$.

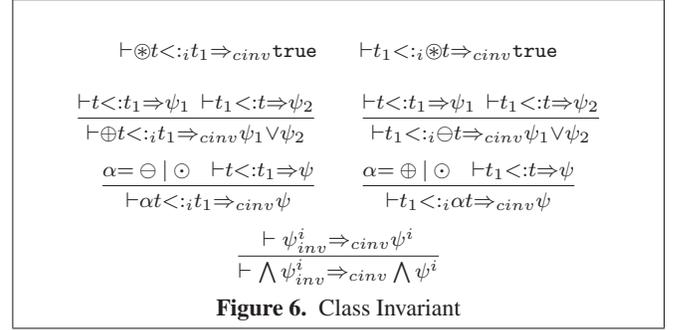
Our type promotion is a refinement of that proposed in [17]. First, we allow promotion to \odot whenever possible while Igarashi and Viroli considered mainly the promotion of nested types with \oplus . Second, we consider type promotion for only field access and class inheritance where the outer variance is dependent on the variance of the underlying type. In contrast, Igarashi and Viroli focused on the promotion of nested types of arguments/result for method declarations, which need not be handled in our approach as these types are fully specified in our method declarations.

4.3 Class Invariant

The class invariant ψ_{inv} is used to capture the lower and upper bounds for the parameterised fields of each newly created object of the class. These bounds are of the form $\bigwedge c_1 \langle \tau^* \rangle <: V <: c_2 \langle \tau^* \rangle$. Class invariant may also support F-bounds when variable V occurs in the parameters of classes c_1 and c_2 . If unspecified, the default lower and upper bounds are \perp and `Object`, respectively. An upper bound invariant on a write-only field restricts the class of the object that can be written to the field to be subclasses of the bound, and a lower bound invariant on a read-only field restricts the class of the object that can be read from the field to be superclass of the bound.

We use the relation $\Rightarrow_{\text{cinv}}$ to reduce bounds from the class invariant to a constraint form: $\vdash [V_i \mapsto \tau_i] \psi_{\text{inv}} \Rightarrow_{\text{cinv}} \psi$, where τ_i

are the current variant types for the class fields. The relation $\Rightarrow_{\text{cinv}}$ is defined in Figure 6. Note that this relation invokes the subtyping relations defined in Figure 3.



To illustrate the use of these bounded invariants, consider a class declaration for `Cell(X)` with an upper bound $\text{X} <: \text{Num}$. For declarations of the form `Cell($\odot \text{Int}$)` and `Cell($\odot \text{T}$)`, the relation $\Rightarrow_{\text{cinv}}$ generates the `Int <: Num` and `T <: Num`, respectively. The first constraint reduces to `true`, while the second constraint contains a type variable and will be checked later for satisfiability. As another example, for `Cell($\odot \text{Object}$)` the relation $\Rightarrow_{\text{cinv}}$ fails as the upper bound is violated. Correspondingly, for read access, we support `Cell($\oplus \text{Int}$)` and `Cell($\oplus \text{Object}$)`, but not `Cell($\oplus \text{String}$)` since no `String` objects can be read from the `Num`-bounded field.

The class invariant is accumulated recursively from all the superclasses, as shown below:

$$\frac{[\text{CINV}] \quad \text{class } c(V_i)_{i=1}^m \text{ extends } (c_k \langle \pi_{ik} \rangle_{i=1}^{n_{k=1}})^s \text{ where } \psi_{\text{inv}} \{ \dots \} \in P}{\rho = [V_i \mapsto \tau_i]_{i=1}^m \quad \rho \vdash c_1 \langle \pi_{i1} \rangle_{i=1}^{n_1} \Rightarrow_p t \quad \vdash \rho \psi_{\text{inv}} \Rightarrow_{\text{cinv}} \psi} \quad \text{cinv}(c_1 \langle \tau_i \rangle_{i=1}^m) = \psi \wedge \text{cinv}(t)$$

5. Variant Type System

Variance annotations of programs are used to support flow analysis for more accurate generic types. We verify the flow of values through the following typing relation:

$$\Gamma; Q \vdash e :: \alpha t, \psi$$

The relation is for type checking, and assumes that Γ (type environment), Q (type variables in scope) and αt (type with expected variance) are given while ψ is the collected flow constraint. Syntax-directed rules for various language constructs are given in Figure 7.

Our type system is flow-insensitive as every location (variable, parameter and field) is given a type that never changes. In our type system, each object of type t_1 can be placed in a location of type t_2 , provided $t_1 <: t_2$. The type of a location is therefore a particular *type view* of its object. This type view of an object may be changed by upcasting (via assignment or parameter passing) or by downcast operation that is checkable at runtime. The following rule shows how to type check an assignment expression:

$$\frac{[\text{ASSIGN}] \quad \alpha t = \text{GetType}(\Gamma, w) \quad \alpha <: \ominus \quad \Gamma; Q \vdash e :: \oplus t, \psi}{\Gamma; Q \vdash w = e :: \oplus \text{Void}, \psi}$$

Flow-in or write-only \ominus is mandated on the left-hand side (w) while flow-out or read-only \oplus is mandated on the right-hand side (e). To highlight how these flows are enforced, we present the rule for variable and field access (w stands for either v or $v.f$):

$\frac{}{\Gamma; Q \vdash \text{null} :: \tau, \oplus \perp <: \tau}$	$\frac{\Gamma' = \Gamma + \{v :: \odot t\} \quad \Gamma; Q \vdash e_1 :: \oplus t, \psi_1 \quad \Gamma'; Q \vdash e_2 :: \tau, \psi_2}{\Gamma; Q \vdash \{t v = e_1; e_2\} :: \tau, \psi_1 \wedge \psi_2}$	$\frac{\Gamma; Q \vdash e_1 :: \oplus t, \psi_1 \quad \Gamma; Q \vdash e_2 :: \tau, \psi_2}{\Gamma; Q \vdash e_1; e_2 :: \tau, \psi_1 \wedge \psi_2}$
$\frac{\Gamma(v) <: \oplus \text{Bool} \quad \Gamma; Q \vdash e_1 :: \tau, \psi_1 \quad \Gamma; Q \vdash e_2 :: \tau, \psi_2}{\Gamma; Q \vdash \text{if } v \text{ then } e_1 \text{ else } e_2 :: \tau, \psi_1 \wedge \psi_2}$	$\frac{\Gamma(v) <: \oplus \text{Bool} \quad \Gamma; Q \vdash e :: \tau, \psi}{\Gamma; Q \vdash \text{while } v \text{ do } e :: \oplus \text{Void}, \psi}$	$\frac{\vdash_{def} \text{InheritanceOK}(def_i), i = 1..n \quad \vdash_{def} def_i, i = 1..n}{\vdash_{prg} def_{i=1..n}}$
$\frac{\rho = [V_j \mapsto t_j]_{j=1}^k \quad \tau'_i = \Gamma(v'_i)_{i=0}^q \quad \hat{t}_0 \mid t \text{ mn}(\hat{t}_i v_i)_{i=1}^q \langle V_{1..k} \rangle \text{ where } \psi.. \in \tau'_0 \quad \psi_1 = \bigwedge_{i=0}^q \tau'_i <: \rho(\oplus \hat{t}_i) \wedge \rho(\oplus t) <: \tau}{\Gamma; Q \vdash v'_0 \text{ mn}(v'_1, \dots, v'_q) \langle t_{1..k} \rangle :: \tau, \psi_1 \wedge \rho \psi}$	$\frac{c_1 \vdash_{meth} meth_i, i = 1..q \quad \text{vars}\{\pi_i\}_{i=1}^n \cup (\text{vars} \bigcup_{k=1}^s \{\hat{\pi}_{ik}\}_{i=1}^{n_k}) \subseteq \{X_i\}_{i=1}^m}{\vdash_{def} \text{class } c_1 \langle X_i \rangle_{i=1}^m \text{ extends } (\hat{c}_k \langle \hat{\pi}_{ik} \rangle_{i=1}^{n_k})_{k=1}^s \text{ where } \psi_{in v} \{(\pi_i f_i)_{i=1}^n \text{ meth}_{i=1..q}\}}$	
$\frac{\text{def} = \text{class } c_1 \langle V_i \rangle_{i=1}^p \text{ extends } c_2 \langle \hat{\pi}_i \rangle_{i=1}^q \text{ where } \{fd^* \text{ meth}_{1..p}\} \quad (\exists \text{meth} \cdot \text{meth} \in c_2 \langle \hat{\pi}_i \rangle_{i=1}^q \wedge \text{name}(\text{meth}) = \text{name}(\text{meth}_i)) \Rightarrow \vdash \text{OverridesOK}(\text{meth}_i, \text{meth}) \quad i \in 1..p}{\vdash \text{InheritanceOK}(\text{def})}$	$\frac{\text{meth}_1 = t_0 \mid t \text{ mn}(\langle t_i v_i \rangle_{i=1}^p) \langle V^* \rangle \text{ where } \psi_1 \{e_1\} \quad \text{meth}_2 = \hat{t}_0 \mid t \text{ mn}(\langle t_i v_i \rangle_{i=1}^p) \langle V^* \rangle \text{ where } \psi_2 \{e_2\} \quad V_L = \text{vars}(\hat{t}_0) - \text{vars}(t_0) \quad \vdash t_0 <: \hat{t}_0 \Rightarrow \psi \exists V_L \cdot (\psi \wedge \psi_2 \Rightarrow \psi_1)}{\vdash \text{OverridesOK}(\text{meth}_1, \text{meth}_2)}$	
$\frac{\text{class } c_1 \langle V_i \rangle_{i=1}^n \text{ extends } c_2 \langle \hat{\pi}_i \rangle_{i=1}^r \dots \{(\pi'_i f_i)_{i=1}^m\} \quad \rho = [V_i \mapsto \tau_i]_{i=1}^n \quad \rho \vdash \pi'_i \Rightarrow_p \tau'_i, i \in 1..m \quad \rho \vdash \hat{\pi}_i \Rightarrow_p \hat{\tau}'_i, i \in 1..r}{\text{fields}(c_1 \langle \tau_i \rangle_{i=1}^n) = [(\tau'_i f_i)_{i=1}^m + \text{fields}(c_2 \langle \hat{\tau}'_i \rangle_{i=1}^r)}$	$\frac{\tau = \Gamma(v)}{\tau = \text{GetType}(\Gamma, v)}$	$\frac{\alpha t = \text{GetType}(\Gamma, v) \quad t = c \langle \tau_i \rangle_{i=1}^n \quad (\tau f) \in \text{fields}(c \langle \tau_i \rangle_{i=1}^n)}{\tau = \text{GetType}(\Gamma, v.f)}$

Figure 7. Variant Type Rules

$$\frac{\tau_1 = \text{GetType}(\Gamma, w) \quad \vdash \tau_1 <: \tau \Rightarrow \psi}{\Gamma; Q \vdash w :: \tau, \psi}$$

To retrieve the types of the variables and class fields, we use the auxiliary [GetType] rules from Figure 7. The current type τ_1 of w is retrieved from the type environment Γ . Further, the rule checks that τ_1 is a subtype of the expected variant type τ . This supports a flow-out from the variable w .

For object creation, we ensure that each object is constructed with an invariant type using $c \langle \odot t_i \rangle_{i=1}^q$. A type is said to be *invariant* if each variance on its immediate fields is marked with \odot . Note that the views of nested fields, namely t_1, \dots, t_q from $c \langle \odot t_i \rangle_{i=1}^q$, may still be of variant types. Note that the variance of all class fields (including those which require type promotion) returned by fields is \odot .

$$\frac{\text{vars}\{t_i\}_{i=1}^q \subseteq Q \quad t_0 = c \langle \odot t_i \rangle_{i=1}^q \quad (\odot t'_i f_i)_{i=1}^p = \text{fields}(t_0) \quad \vdash \oplus t_0 <: \tau \Rightarrow \psi_0 \quad \Gamma; Q \vdash v_i :: \oplus t'_i, \psi_i \quad i = 1..p}{\Gamma; Q \vdash \text{new } c \langle t_i \rangle_{i=1}^q (v_1, \dots, v_p) :: \tau, \bigwedge_{i=0}^p \psi_i \wedge \text{cinv}(t_0)}$$

For the purpose of constructing invariant types, the type variables in $\{t_i\}_{i=1}^q$ must be instantiated from Q . The class invariant $\text{cinv}(t_0)$ captures the specified upper/lower bounds on fields that must be satisfied for every object of the class. When such fields are updated, we statically ensure that their bounds are never violated. Given an instantiated class type, the rule [FIELDS] returns the variant types of the class fields using type promotion if necessary.

Local variable declaration v is marked for read-write access via $v :: \odot t$ as shown in the rule [LOCAL]. The rule for method call [Call] collects the flow-in for receiver and arguments, flow-out for the result and the method precondition.

5.1 Modular Flow Verification

We design a variant type system that can be verified in a modular fashion. Each method declaration is given suitable variant type annotations for its parameters, result and receiver. A “may” flow constraint ψ is specified at the header of each method declaration. This *may-flow* specification captures all possible flows that may occur in the method’s body e . The type checking rule for a method is formalised as follows:

$$\frac{\text{chkRecv}(\text{cn}, t_0) \quad \Gamma = \{v_i :: \oplus t_i\}_{i=1}^p + \{\text{this} :: \oplus t_0\} \quad \psi_1 = \psi \wedge \bigwedge_{i=0}^p \text{cinv}(t_i) \wedge \text{cinv}(t) \quad \psi_1 \neq \text{false} \quad Q = \{V^*\} \quad \text{vars}(\psi) \subseteq Q \quad \text{vars}(\Gamma, t) \subseteq Q \quad \Gamma; Q \vdash e :: \oplus t, \psi_2 \quad V_I = \text{vars}(\psi_2) - Q \quad \psi_1 \Rightarrow \exists V_I \cdot \psi_2}{\text{cn} \vdash_{meth} t_0 \mid t \text{ mn}(\langle t_i v_i \rangle_{i=1}^p) \langle V^* \rangle \text{ where } \psi \{e\}}$$

We first construct an initial assumed flow constraint ψ_1 that is derived from the declared may-flow specification ψ , class invariants for each parameter, and result $\bigwedge_{i=0}^p \text{cinv}(t_i) \wedge \text{cinv}(t)$. The initial assumed flow constraint must be satisfiable, that is, $\psi_1 \neq \text{false}$. Furthermore, we collect the flow constraint of the method body using $\Gamma; Q \vdash e :: \oplus t, \psi_2$, where ψ_2 captures all flows that may occur in the method body e . To prove the correctness of each declared flow constraint, we perform a subtype entailment on the flow constraint with V_I as local type variables using: $\psi_1 \Rightarrow \exists V_I \cdot \psi_2$. If this entailment holds, we have successfully verified the flow specification of a given method declaration. We also check if t_0 , the given type of this , is compatible (no stupid cast) with the current class via the predicate $\text{chkRecv}(\text{cn}, t_0) = \text{cn} \langle \odot t^* \rangle <: t_0$.

Method overriding is checked by the [Override] rule. For safe function subtyping, we require each overriding method to have *weaker or equal* flow specification compared to the overwritten method.

5.2 Soundness

The soundness of our type system can be proven by relating to dynamic evaluation semantics of the form:

$$\langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \Pi', \varpi' \rangle [e']$$

where Π and ϖ denote runtime stack and heap, respectively. This evaluation may yield three possible runtime errors, namely $\mathbf{E} = \mathbf{Error-Null} \mid \mathbf{Error-Cast} \mid \mathbf{Error-Type}$. The second error is due to cast operations guarded by runtime checks inserted by the compiler. The third error is due to an object of the wrong type being written into a location with some expected static type. For well-typed programs, this last error can never happen. The progress theorem states that **Error-Type** cannot occur while the type preservation theorem shows that the type of an expression is preserved with each reduction step. We outline the two theorems below; details of proof may be found in Appendices A, B and C.

THEOREM 1 (Progress). *Let Γ be the environment mapping program variables to ground types. If $\Gamma; \Sigma; Q \vdash e :: \tau, \psi$ and $\Gamma; \Sigma; Q; \psi \models \Pi, \varpi$, then either:*

- e is a value, or
- $\langle \Pi, \varpi \rangle [e] \hookrightarrow \mathbf{Error-Null} \mid \mathbf{Error-Cast}$, or
- there exist Π', ϖ', e' such that $\langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \Pi', \varpi' \rangle [e']$.

Note that the type rules are extended to include store typing Σ . $\Gamma; \Sigma; \psi \models \Pi, \varpi$ denotes a consistency relation that relates static and dynamic semantics. The following theorem states the preservation of type during dynamic evaluation.

THEOREM 2 (Preservation). *Let Γ be an environment mapping program variables to ground types. If*

$$\begin{aligned} \Gamma; \Sigma; Q \vdash e :: \tau, \psi \\ \Gamma; \Sigma; Q; \psi \models \Pi, \varpi \\ \langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \hat{\Pi}, \hat{\varpi} \rangle [\hat{e}] \end{aligned}$$

then there exists $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} such that

$$\begin{aligned} \Gamma - \text{diff}(e, \hat{e}) = \hat{\Gamma} - \text{diff}(\hat{e}, e) \\ \hat{\Sigma} \supseteq \Sigma \\ \hat{\Gamma}; \hat{\Sigma}; \hat{Q} \vdash \hat{e} :: \tau, \hat{\psi} \\ \hat{\Gamma}; \hat{\Sigma}; \hat{Q}; \hat{\psi} \wedge \psi \models \hat{\Pi}, \hat{\varpi}. \end{aligned}$$

Function $\text{diff}(e, e')$ returns a list of local variables that appear in e but not e' .

6. Casting and Cast Capture

While a key goal of a generic type system is to provide precise information to eliminate unnecessary downcasts, there remains always the need for cast operations to support the class subtyping mechanism. Furthermore, the introduction of generics and variance has complicated type casting as these operations must handle type variables and nested variant types. For example, cast operations may target nested types, such as $\text{Vector}(\odot \text{Vector}(\oplus \text{Num}))$, or those with type variables, such as $\text{Vector}(\oplus X)$.

However, existing solutions that support casting in Java 1.5 are restricted in that they use cast checks on the outermost type constructor only [37], and rely on unchecked warnings that may cause runtime errors (e.g., when a cast to type variable occurs). The only system that supports cast operations fully (but for parametric types) was proposed by Viroli and Natali [39]. Their technique can be adapted to handle arbitrary variant types.

In the presence of single inheritance, we can classify each casting relation from t_0 to t into three categories: (1) safe upcast if $\neg(t_0 <: t)$, (2) downcast with runtime check if $\vdash t <: t_0$, and (3) stupid

cast if $\neg(\vdash t_0 <: t \vee \vdash t <: t_0)$. However, in the presence of multiple inheritance with interfaces, a class and an interface may be unrelated but a valid downcast is still possible if the actual type is a subtype of the two. Though it is possible to identify stupid cast with a more complex test, namely $\neg(\exists X. X \neq \perp \wedge X <: t \wedge X <: t_0)$, we avoid it for simplicity. Instead, we only check to ensure that the type variables used in t have come from Q . Our type rule to support a variant cast operation is given below:

$$\frac{[\mathbf{CAST}] \quad \alpha t_0 = \Gamma(v) \quad \alpha <: \oplus \quad \text{vars}(t) \subseteq Q}{\Gamma; Q \vdash (t)v :: \tau, \oplus t <: \tau}$$

While casting is used to check a specific type for an object, we often have to deal with objects of unknown types. For example, we may have an object with a static type $\text{List}(\oplus A)$, and we may be interested to know its actual invariant type $\text{List}(\odot T)$, where T is an unknown type. To help identify the invariant type of a given object, we introduce a *cast capture* construct based on reflection mechanism: $\{v_1 = (t)v; e\}$ The following rule shows how to type check the capture construct:

$$\frac{[\mathbf{CAPTURE}] \quad \alpha c(\tau_i)_{i=1}^n = \Gamma(v) \quad \alpha <: \oplus \quad t = c(\odot V_i)_{i=1}^n \quad \{V_i\}_{i=1}^n \cap Q = \{\} \\ \Gamma; Q \vdash v_1 :: \odot t, \psi_1 \quad Q_1 = Q \cup \{V_i\}_{i=1}^n \quad \Gamma; Q_1 \vdash e :: \tau, \psi_2}{\Gamma; Q \vdash \{v_1 = (t)v; e\} :: \tau, \psi_1 \wedge \psi_2}$$

Note that t is an invariant type of the form $c(\odot V_i)$: c should have the same class type as v , while the *captured type variables* V_i stand for unknown types. Each V_i can be used in the expression e with its flow captured by the collected flow $(\psi_1 \wedge \psi_2)$.

The flow of captured type variables should not cause additional restriction or generalization at the method boundary. We next present how the type system ensures the correct use of captured type variables.

The actual type t obtained via reflection is guaranteed to be more precise than v 's static type, $\Gamma(v)$. We call this guarantee *reflection assumption*. For each method, a relation $\Gamma \vdash e \Rightarrow_C V_C, \psi_C$ collects captured type variables, V_C , and their reflection assumptions, ψ_C as follows:

$$\frac{\alpha t_0 = \Gamma(v) \quad \vdash t <: t_0 \Rightarrow \psi \\ X = \text{vars}(t) \quad \Gamma \vdash e \Rightarrow_C V, \psi_1}{\Gamma \vdash \{v_1 = (t)v; e\} \Rightarrow_C V \cup X, \psi \wedge \psi_1}$$

The method judgement is modified to exclude captured type variables V_C from the local type variables V_I . Additionally, the expected flow ψ_1 is strengthened with reflection assumptions ψ_C .

$$\frac{[\mathbf{METHOD-WITH-CAPTURE}] \quad \text{chkRecv}(cn, t_0) \quad \Gamma = \{v_i :: \oplus t_i\}_{i=1}^p + \{\text{this} :: \oplus t_0\} \\ \Gamma \vdash e \Rightarrow_C V_C, \psi_C \quad \psi_1 = \psi \wedge \bigwedge_{i=0}^p \text{cinv}(t_i) \wedge \text{cinv}(t) \wedge \psi_C \\ Q = \{V^*\} \quad \text{vars}(\psi) \subseteq Q \quad \text{vars}(\Gamma, t) \subseteq Q \quad \psi_1 \neq \text{false} \\ \Gamma; Q \vdash e :: \oplus t, \psi_2 \quad V_I = \text{vars}(\psi_2) - Q - V_C \quad \psi_1 \Rightarrow \exists V_I. \psi_2 \\ cn \vdash_{\text{meth}} t_0 \mid t \text{ mn}((t_i v_i)_{i=1}^p)(V^*) \text{ where } \psi \{e\}}{\Gamma \vdash e \Rightarrow_C V_C \cup V_I, \psi_1 \wedge \psi_2}$$

The proper flow of captured type variables is then ensured by the entailment from the above rule.

6.1 Cast Capture Examples

The cast capture mechanism can also be viewed as a downcast to the object's invariant type. Unknown types that are captured (via reflection) may be used in the program code, as shown in the example below:

```

void addNode(List<⊖A> y, B z) where B<:A {
  List<⊖S> v; List<⊖S> w;
  {v = (List<⊖T>) y; w = new List<T>();
  w.val = z ; w.next = v.next ; v.next = w; } }

```

Though we do not know the exact type of y , we can use a cast capture on $(\text{List}\langle\ominus T\rangle)$ to obtain its invariant type. Correspondingly, the reflection assumption is $A<:T$. We use the captured type T to build a $\text{List}\langle\ominus T\rangle$ node, write z to $w.\text{val}$, and also reconstruct pointers for the linked list in a type-safe and yet generic way. For this example, the initial assumed flow is $\psi_1 \equiv (B<:A \wedge A<:T)$, whereby $B<:A$ is from the flow specification and $A<:T$ is the reflection assumption. This initial assumed flow implies the collected flow constraint $\exists S \cdot \psi_2$, where $\psi_2 \equiv (S<:T \wedge T<:S \wedge B<:S)$. Hence, modular verification holds for this example.

The same cast capture mechanism may also be used to capture an unknown invariant type, enabling a swap of elements within the same collection – without knowledge of its type. Consider:

```

void swapVec(Vector<⊕> v, int i, int j) {
  Vector<⊖S> w;
  {w = (Vector<⊖T>) v;
  S v1 = w.elementAt(i);
  S v2 = w.elementAt(j);
  w.setElementAt(v2, i); w.setElementAt(v1, j); } }

```

Note that input parameter v uses a bivariant type $\text{Vector}\langle\oplus\rangle$, which can be used to support an argument of an arbitrary Vector object. The initial assumed flow is $\psi_1 \equiv \text{true}$, while the collected flow is $\exists S \cdot \psi_2$, where $\psi_2 \equiv (S<:T \wedge T<:S)$. Hence, the entailment $\psi_1 \implies \exists S \cdot \psi_2$ holds.

An example of a method that does not type check is presented below:

```

Vector<⊕Y> foo1(Vector<⊕> v) {
  Vector<⊖S> w; {w = (Vector<⊖T>) v; w } }

```

The initial assumed flow is $\psi_1 \equiv \text{true}$ while the collected flow is $\psi_2 \equiv T<:Y$. Note that neither T (captured type variable) nor Y (global type variable) are existentially quantified from ψ_2 . The entailment $\psi_1 \implies \exists S \cdot \psi_2$ does not hold, since the captured type variable T introduces an additional flow at method boundary. As another example, the following definition type checks as the collected flow from the method's body (after elimination of the local type variable S) is $\psi_2 \equiv \text{true}$:

```

Vector<⊕> foo2(Vector<⊕> v) {
  Vector<⊖S> w; {w = (Vector<⊖T>) v; w } }

```

7. Implementation

We built a prototype for our variant parametric type system and carried out initial experiments to validate its feasibility. Our system was built using the Glasgow Haskell compiler [27], with a constraint solver (for handling subtyping constraints) implemented using Constraint Handling Rules (CHR) [13].

Our constraint solver employs a two-step algorithm to prove the non-structural subtype entailment of the form $\forall V_G \cdot (\psi_1 \implies \exists V_I \cdot \psi_2)$. Note that ψ_1, ψ_2 are conjunctions of subtyping constraints, while V_G and V_I are sets of type variables. Even though the entailment from the [METHOD] rule may contain disjunctions, it can be reduced to entailments of the above form.

1. We eliminate the local type variables V_I (based on their upper and lower bounds) from ψ_2 to obtain $\psi'_2 = \bigwedge_{i=1}^n X_i <: Y_i$ using techniques similar to [32, 38].

To support the language's semantics a local type inference similar to [30, 23] is employed to identify appropriate instantiated types for local type variables or type parameters.

2. The resulting entailment $\forall V_G \cdot (\psi_1 \implies \bigwedge_{i=1}^n X_i <: Y_i)$ is equivalent to $\bigwedge_{i=1}^n (\forall V_G \cdot (\psi_1 \implies X_i <: Y_i))$. Each entailment can be proven by contradiction using the falsity of the formula $\forall V_G \cdot (\psi_1 \wedge \text{notsub}(X_i, Y_i))$, where $\text{notsub}(t_1, t_2)$ represents negation of subtyping relation.

Our constraint solver implements the variant subtyping rules (from Figure 3). Its deduction mechanism detects falsity based on pair of constraints of the form $t_1 <: t_2$ and $\text{notsub}(t_1, t_2)$. Our algorithm is a sound approximation of the subtype entailment problem.

The deduction mechanism can be further extended by the techniques of *case analysis* and *inductive proving*. However, from our experience working with large sets of Java library and application codes that have been annotated and checked with variant parametric types, we have yet to encounter real examples which require such extensions.

8. Experimental Results

To test the utility of our flow-based variant type system, we evaluated our prototype on a set of Java applications¹ as used in [10, 14]. These applications make use of library classes from package `java.util`, which we annotated with our variant parametric types. We counted each method declaration with flow specification, each class declaration with type parameters and each cast capture as a line of annotation. On average, these annotations constituted about 5.5% of the source code, which may be considered a reasonable price to pay for better reuse of type safe generic code. Due to modular type checking, the time needed to verify type-safe generic code was less than one second for each library code and less than 30 seconds for each application code. We expect that the time can be reduced by using a specialised constraint solver. Currently, our prototype is based on a meta constraint handling system written in CHR (which compiled to a Prolog program under IC-Parc's ECLiPSe system [2]).

Library	Prog. Lines	Java 1.4			VPT	
		Casts	Casts	Warnings	Casts	Warnings
AbstractList	909	1	1	0	0	0
AbstractSet	162	1	1	0	0	0
ArrayList	623	2	8	9	1	0
HashMap	1103	7	9	20	3	0
HashSet	231	2	4	3	1	0
Hashtable	1154	10	14	31	7	0
LinkedList	814	2	4	5	2	0
Properties	925	8	8	1	0	0
Vector	1062	2	9	9	0	0
Total	6983	35	58	78	14	0

Figure 8. Results for Library Code

Figures 8 and 9 show the experimental results for representative classes from the `java.util` package and application code (in terms of remaining casts). We counted the number of casts in Java 1.4 code (non-generic), Java 1.5 (annotated with wildcards) and our system (VPT - annotated with variant parametric types). The Java 1.5 compiler could not statically check some operations (especially those related to raw types and casts to type variables), and issued unchecked warnings since these operations cannot be verified by JVM runtime. Therefore, it is the programmer's responsibility to ensure that all unchecked operations are in fact safe.

¹For more details: www.junit.org, www.cs.princeton.edu/~appel/modern/java/JLex/, www.cs.princeton.edu/~appel/modern/java/CUP/, www.spec.org/osg/jvm98/, v poker.sourceforge.net, telnetd.sourceforge.net.

Application	Prog. Lines	Java 1.4		Java 1.5		VPT	
		Casts	Casts	Warnings	Casts	Warnings	
DB	842	19	1	0	0	0	
JUnit	5886	54	30	1	15	0	
VPoker	6792	36	8	0	6	0	
JLex	7260	69	12	3	0	0	
Jess	10639	95	34	0	12	0	
TelnetD	11314	46	8	0	6	0	
JavaCup	11468	543	98	2	65	0	
Total	54201	862	191	6	104	0	

Figure 9. Results for Application Code

To summarize, our method can eliminate a significant portion (on average 87.9%) of the casts from non-generic Java 1.4 application code and 45.5% of the casts from wildcard-generic Java 1.5 application code. We have also made improvements for library code by eliminating about 60% casts from non-generic Java 1.4 code and about 75.8% casts from the wildcard-generic Java 1.5 code. Since our system fully supports casting for variant types, we can verify the unsafe operations for which the Java 1.5 compiler generates unchecked warnings. Note that Java 1.5 libraries contain more casts than Java 1.4 libraries do, since Java 1.4 containers are based on `Object` type instead of generic types. As expected, Java 1.4 application code requires more downcasts compared to Java 1.5 code.

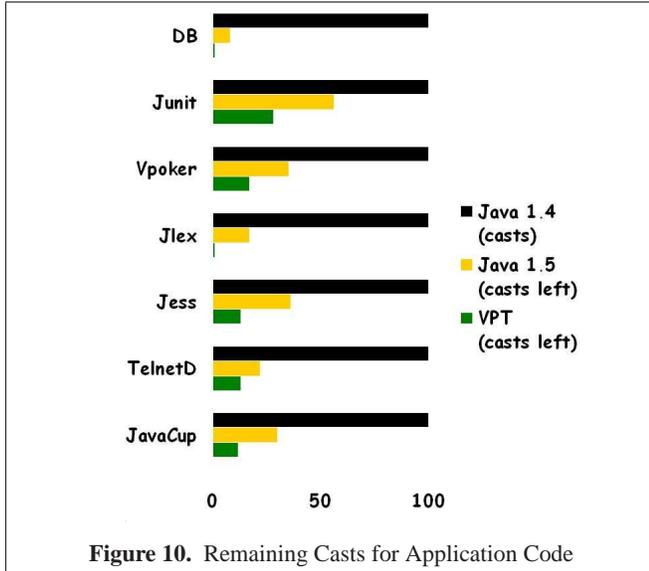


Figure 10. Remaining Casts for Application Code

Figure 10 shows a chart that visualises the percentage of remaining casts in each benchmark written in Java 1.4, Java 1.5 and our VPT. Java 1.4 which contains the casts from the original application code serves as reference.

Note that the casts eliminated using our type system measure the improvement in program safety. Current Java implementation (which translates parametric programs via *type erasure*) would reintroduce casts at the bytecode level. While such re-admitted casts may cause runtime overheads, they are known to be type safe and will never fail at runtime. Obviously, a better solution is to support variant parametric type at the bytecode level and we look forward to this scenario.

9. Extensions

In this section, we present some features omitted in the main presentation for brevity.

The hierarchy of primitive types forms a separate lattice from reference types. Furthermore, it is *not* the case that $\perp <: p <: Object$ for each primitive type p . Due to such differences, primitives are excluded from use as type arguments for generic classes in Java 1.5. Furthermore, the type erasure algorithm for the parametric program will transform each parametric field into an `Object` type for backwards compatibility. This is invalid if primitive types are used as type arguments.

We now show how primitive types can be used as type arguments for generic classes in our system. First, we need to add two constraints to distinguish reference and primitive types, as shown below:

$$\psi ::= \dots \mid \text{ref}(t) \mid \text{prim}(t)$$

As these two families of types are disjoint, we add the following CHR irrevocable rule:

$$\text{ref}(t) \wedge \text{prim}(t) \Leftrightarrow \text{false}$$

Second, we need to consider primitive types in the new variant subtyping mechanism. In the new subtyping hierarchy, $\otimes t$ denotes any type (reference or primitive) while $\oplus Object$ and $\ominus \perp$ denote only reference types (that are still equivalent, namely $\oplus Object \equiv \ominus \perp$). The subtyping relation is changed accordingly: $\oplus Object <: \otimes t$ still holds while $\otimes t <: \oplus Object$ does not hold anymore. Furthermore, we allow $\perp <: t$ and $t <: Object$ if and only if t is not a primitive type. To support these changes, we modify the main variant subtyping rules from Figure 3 to the following:

$$\frac{\alpha \neq \otimes}{\vdash \alpha t <: \oplus Object \Rightarrow \text{ref}(t)} \quad \frac{\alpha \neq \otimes}{\vdash \alpha t <: \ominus \perp \Rightarrow \text{ref}(t)}$$

$$\frac{\alpha_1 \neq \otimes \quad \neg(\alpha_1 <: \oplus)}{\vdash Object <: t_2 \Rightarrow \psi} \quad \frac{\alpha_1 \neq \otimes \quad \neg(\alpha_1 <: \ominus)}{\vdash t_2 <: \perp \Rightarrow \psi}$$

$$\frac{\vdash \alpha_1 t_1 <: \oplus t_2 \Rightarrow \psi \wedge \text{ref}(t_1)}{\vdash \perp <: t \Rightarrow \text{ref}(t)} \quad \frac{\vdash \alpha_1 t_1 <: \ominus t_2 \Rightarrow \psi \wedge \text{ref}(t_1)}{\vdash t <: Object \Rightarrow \text{ref}(t)}$$

$$\frac{}{\vdash t <: \perp \Rightarrow t <: \perp \wedge \text{ref}(t)} \quad \frac{}{\vdash Object <: t \Rightarrow Object <: t \wedge \text{ref}(t)}$$

Programmers often make use of the `instanceof` test on the runtime type of an object prior to some operations. Due to flow and path insensitivity, the type system is currently unable to take advantage of such runtime testing of types. To help eliminate more cast operations, our compiler translates each program fragment of the form:

```
if v instanceof(t) then e1 else e2
```

to use a special program construct with fresh v_0 variable:

```
if v instanceof(t) then let v0 :: t = v in [v → v0] e1
else e2
```

This construct is part of our core intermediate language, and it is generated prior to type checking. It is valid on the proviso that any assignment into v is a subtype of the more specific t . A type rule corresponding to the new language construct is shown below:

$$\frac{[\text{LET-INSTANCEOF}]}{\Gamma' = \Gamma + \{v_0 :: \odot t\} \quad \Gamma'; Q \vdash e :: \tau, \psi_1 \quad \Gamma; Q \vdash e_2 :: \tau, \psi_2}{\Gamma; Q \vdash \text{if } v.\text{instanceof}(t) \text{ then } e_1 \text{ else } e_2 :: \tau, \psi_1 \wedge \psi_2}$$

Flow-sensitivity may also cause some loss in type precision (such that some downcasts cannot be statically verified) when the same local variable is used for objects with different variant parametric types. To rectify this, we could use Static Single Assignment (SSA) intermediate form [8] which is known to give better flow-sensitive analysis results. Conversion of programs to SSA form can be handled in a preprocessing step, prior to type checking.

These techniques for incorporating path and flow sensitivity are quite standard, and were also explored in [41].

10. Conclusion

Software reuse has received much research interest for its boost to software development and maintenance productivities. Recently, generic type has become a main thrust in supporting software reuse. In reusing Java code, several works have proposed for refactoring legacy Java programs into those that use generic versions of popular container classes [10, 11, 14, 40].

Other works try to achieve precise Java type inference results in the presence of parametric polymorphism and data polymorphism in order to reduce the redundant cast operations [31, 1, 41]. The precision typically comes at the price of a whole program analysis. Every time when an application code is analysed, the reachable library code must also be re-analysed.

Variation parametric types attempt to increase language expressivity and code reuse by introducing another subtyping scheme, based on the notion of variance. This idea originated from the structured virtual types proposed by Thorup and Torgersen [35]. Their work is the first to link access rights and covariant subtyping to the fields of each use of a class rather than the class itself. Igarashi and Viroli extended this concept to support contra- and bi-variance [17]. They also formalised the variant type system by mapping it into a corresponding *existential type* system [17, 18] for Featherweight Java. While Igarashi and Viroli's design faithfully models the existential type system, it has been found to be too restrictive by the designers of Java 1.5. One improvement made in Java 1.5 is to allow each wildcard type to be opened without a corresponding close operation. This provides more flexibility for writing generic code, but weakens the link to the traditional pack/unpack mechanism of the existential type system. Hence, even though a full-scale language system has been implemented, the soundness of the wildcard type system is still under development (as of [36]). Other than Java, a recently developed language Scala [24] supports variance for parametric polymorphism. In contrast with our approach, Scala uses variance at *declaration-site*. However, an earlier version of Scala has experimented with the use-site variance mechanism that is consistent with the original system of Igarashi and Viroli but without the flexibility of the wildcard capture. This was considered to be too restrictive before the authors abandoned the approach. Recently, generic types of C# [12] were extended with *declaration-site* variance following the design adopted for the language Scala.

Theoretical foundations of the variance have also been studied in the context of typed λ -calculi, where type operators are equipped with a polarity property [6, 33, 9]. These foundations have even been extended to handle higher-order functions, but are closer in nature to declaration-site variance, and have mostly been formalised in only a theoretical setting, without practical implementations.

In our paper, we have proposed a new approach based on flow analysis to support the variant parametric type system. We leverage prior knowledge that has been accumulated for flow analysis and entailment for non-structural subtyping constraints. Palsberg and O'Keefe [25] show the equivalence of flow analysis and non-structural subtyping. Subtype entailment is known to be hard even for simple subtyping constraints. Rehof and Henglein determined the complexity of structural subtype entailment: for

simple types, it is coNP-complete [15] and for recursive types it is PSPACE-complete [16]. Furthermore, they showed that non-structural subtype entailment is PSPACE-hard and is conjectured PSPACE-complete [16]. Su et al. [34] show the decidability of the first-order theory of non-structural subtyping with unary function symbols. Algorithms for non-structural subtype entailment (sound, but incomplete) were developed in Pottier [32], Trifonov and Smith [38]. While the decidability of non-structural subtype entailment remains an open problem, we use sound techniques based on these previous algorithms.

Our new approach is practically driven and can give better generic types. We have also augmented it with intersection types to support Java-like multiple (interface) inheritance. We have built a prototype system based on a set of syntax-directed type rules. This prototype is supported by a constraint-solver for variant subtyping, customised using CHR. Furthermore, our system supports full casting for variant types. Through a new cast capture mechanism, we can use reflection to handle objects with unknown types in a type-safe way. Experimental evaluation indicates that more downcasts can be eliminated by our approach, even when it is compared against the state-of-the-art type system from Java 1.5. Our flow-based approach to variant parametric type system is another step towards better genericity for type-safe OO programs.

Acknowledgments

We are grateful to Atsushi Igarashi for clarifying many questions regarding the variant type system. We also thank Alex Aiken, Greg Morrisett and Martin Rinard for providing useful feedback on this work. Shengchao provided useful technical comments on the paper while Hong Yaw implemented the Java-to-CoreJava translator and hand-annotated a suite of Java libraries and application code. We also thank Alexia Leong for her proof-reading efforts. This work is supported by research grant R-252-000-151-112 and a gift from Microsoft Singapore.

References

- [1] Ole Agesen. The cartesian product algorithm: Simple and precise type inference of parametric polymorphism. In *ECOOP '95: Proceedings of the 9th European Conference on Object-Oriented Programming*, pages 2–26, London, UK, 1995. Springer-Verlag.
- [2] IC-Parc at Imperial College. ECLiPSe Constraint Logic Programming. <http://www.icparc.ic.ac.uk/eclipse/>.
- [3] Gilad Bracha, Martin Odersky, David Stoutamire, and Philip Wadler. Making the future safe for the past: adding genericity to the Java programming language. In *OOPSLA '98: Proceedings of the 13th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 183–200, New York, NY, USA, 1998. ACM Press.
- [4] Martin Buchi and Wolfgang Weck. Compound types for Java. In *OOPSLA '98: Proceedings of the 13th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 362–373, New York, NY, USA, 1998. ACM Press.
- [5] Peter S. Canning, William R. Cook, Walter L. Hill, Walter G. Olthoff, and John C. Mitchell. F-Bounded polymorphism for object-oriented programming. In *Conference on Functional Programming Languages and Computer Architecture*, pages 273–280, 1989.
- [6] Luca Cardelli. Notes about F_{\leq}^{ω} . 1994. Available at <http://research.microsoft.com/Users/luca/Notes/FwSub.ps>.
- [7] Luca Cardelli and Peter Wegner. On understanding types, data abstraction, and polymorphism. *ACM Comput. Surv.*, 17(4):471–522, 1985.
- [8] Ron Cytron, Jeanne Ferrante, Barry K. Rosen, Mark N. Wegman, and F. Kenneth Zadeck. Efficiently computing static single assignment form and the control dependence graph. *ACM Trans. Program. Lang. Syst.*, 13(4):451–490, 1991.

- [9] Adriana Compagnoni Dominic Duggan. Subtyping for object type constructors. In *Foundations of Object-Oriented Languages (FOOL 1999)*, 1999.
- [10] Alan Donovan, Adam Kiezun, Matthew S. Tschantz, and Michael D. Ernst. Converting Java programs to use generic libraries. In *OOPSLA '04: Proceedings of the 19th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 15–34, New York, NY, USA, 2004. ACM Press.
- [11] Dominic Duggan. Modular type-based reverse engineering of parameterized types in Java code. In *Proceedings of the 1999 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA '99)*, pages 97–113, 1999.
- [12] Burak Emir, Andrew J. Kennedy, Claudio Russo, and Dachuan Yu. Variance and generalized constraints for C# generics. In *Proceedings of the 20th European Conference on Object-Oriented Programming (ECOOP)*, July 2006.
- [13] Thom Fruhwirth and et al. Constraint Handling Rules. <http://www.cs.kuleuven.ac.be/~dtai/projects/CHR/>.
- [14] Robert Fuhrer, Frank Tip, Adam Kiezun, Julian Dolby, and Markus Keller. Efficiently refactoring Java applications to use generic libraries. In *ECOOP '05: Proceedings of the 19th European Conference on Object-Oriented Programming*, July 2005.
- [15] Fritz Henglein and Jakob Rehof. The complexity of subtype entailment for simple types. In *Proceedings of 12th Symposium on Logic in Computer Science (LICS '97)*, pages 352–361, June 1997.
- [16] Fritz Henglein and Jakob Rehof. Constraint automata and the complexity of recursive subtype entailment for simple type. In *Automata, Languages and Programming, 25th International Colloquium, ICALP'98*, pages 616–627, 1998.
- [17] Atsushi Igarashi and Mirko Viroli. On variance-based subtyping for parametric types. In *ECOOP '02: Proceedings of the 16th European Conference on Object-Oriented Programming*, pages 441–469, 2002.
- [18] Atsushi Igarashi and Mirko Viroli. Variant parametric types: A flexible subtyping scheme for generics. *ACM Trans. Program. Lang. Syst.*, 2006.
- [19] Andrew Kennedy and Don Syme. Design and implementation of generics for the .NET common language runtime. In *Proceedings of the 2001 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 1–12, 2001.
- [20] Karl Mazurak and Steve Zdancewic. Type inference for Java 5: Wildcards, F-Bounds, and Undecidability. 2006. A note available at <http://www.cis.upenn.edu/~stevez/note.html>.
- [21] John C. Mitchell and Gordon D. Plotkin. Abstract types have existential type. *ACM Trans. Program. Lang. Syst.*, 10(3):470–502, 1988.
- [22] Bengt Nordstrom, Kent Petersson, and Jan M. Smith. *Programming in Martin-Lof's Type Theory*. Oxford University Press, 1990.
- [23] Martin Odersky, Christoph Zenger, and Matthias Zenger. Colored local type inference. In *The 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 41–53, 2001.
- [24] Martin Odersky and Matthias Zenger. Scalable component abstractions. In *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2005*, pages 41–57, 2005.
- [25] Jens Palsberg and Patrick O'Keefe. A type system equivalent to flow analysis. In *22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 367–378, 1995.
- [26] Jens Palsberg and Christina Pavlopoulou. From polyvariant flow information to intersection and union types. In *The 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 197–208, 1998.
- [27] Simon Peyton-Jones and et al. Glasgow Haskell Compiler. <http://www.haskell.org/ghc>.
- [28] Benjamin C. Pierce. Bounded quantification is undecidable. *Information and Computation*, 112(1):131–165, July 1994.
- [29] Benjamin C. Pierce. *Types and Programming Languages*. The MIT Press, 2002.
- [30] Benjamin C. Pierce and David N. Turner. Local type inference. In *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 252–265, 1998.
- [31] John Plevyak and Andrew A. Chien. Precise concrete type inference for object-oriented languages. In *OOPSLA*, pages 324–340, 1994.
- [32] François Pottier. Simplifying subtyping constraints. In *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming (ICFP '96)*, pages 122–133, 1996.
- [33] Martin Steffen. *Polarized Higher-Order Subtyping*. PhD thesis, Universitat Erlangen-Nurnberg, 1997.
- [34] Zhendong Su, Alexander Aiken, Joachim Niehren, Tim Priesnitz, and Ralf Treinen. The first-order theory of subtyping constraints. In *The 29th SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 203–216, 2002.
- [35] Kresten Krab Thorup and Mads Torgersen. Unifying genericity - combining the benefits of virtual types and parameterized classes. In *ECOOP '99 - Object-Oriented Programming, 13th European Conference*, pages 186–204, 1999.
- [36] Mads Torgersen, Erik Ernst, and Christian Plesner Hansen. Wild FJ. In *Foundations of Object-Oriented Languages (FOOL 2005)*, Long Beach, CA, January 2005.
- [37] Mads Torgersen, Erik Ernst, Christian Plesner Hansen, Peter von der Ahé, Gilad Bracha, and Neal M. Gafter. Adding Wildcards to the Java programming language. *Journal of Object Technology*, 3(11):97–116, 2004.
- [38] Valery Trifonov and Scott F. Smith. Subtyping constrained types. In *Static Analysis, Third International Symposium, SAS'96*, pages 349–365, 1996.
- [39] Mirko Viroli and Antonio Natali. Parametric polymorphism in java: an approach to translation based on reflective features. In *Proceedings of the 2000 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA 2000)*, pages 146–165, 2000.
- [40] Daniel von Dincklage and Amer Diwan. Converting Java classes to use generics. In *Proceedings of the 19th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 1–14, 2004.
- [41] Tiejun Wang and Scott F. Smith. Precise constraint-based type inference for java. In *ECOOP '01: Proceedings of the 15th European Conference on Object-Oriented Programming*, pages 99–117, 2001.

A. Dynamic Semantics

The dynamic operational semantics of Variant CoreJava is described in small steps. Notations used are defined as follows.

<i>Locations</i> :	$\iota \in$	<i>Location</i>
<i>Primitives</i> :	$k \in$	$\text{prim} = \text{int} \uplus \text{bool} \uplus \text{float}$ $\uplus \text{null} \uplus \text{void}$
<i>Values</i> :	$\delta, \nu \in$	$\text{Value} = (\text{TyPrim} \times \text{prim}) \uplus \text{Location}$
<i>Subs</i> :	$\mu, \rho \in$	$\text{Subs} = \text{TVar} \rightarrow_{\text{fin}} \text{Type}$
<i>Store</i> :	$\varpi \in$	$\text{Store} = \text{Location} \rightarrow_{\text{fin}} \text{ObjVal}$
<i>Variable Env</i> :	$\Pi \in$	$\text{VEnv} = \text{Var} \rightarrow_{\text{fin}} \text{Value}$
<i>Object values</i> :	$\eta \in$	$\text{ObjVal} = \text{Type} \times (\text{Fd} \rightarrow_{\text{fin}} \text{Value})$
<i>Type</i> :	$t \in$	<i>Type</i>

TyPrim consists of primitive types. A type t maintained at runtime does not contain any variant information. If need be, it will be treated as one with invariant annotation \odot . A runtime environment Π is a finite map from program variables to their associated values.

A value can be a location referencing an object or a pair containing a primitive value and a primitive type.

A runtime store ϖ is a finite map from locations to object values. An object value is comprised of its type and its field values. We write $\eta.f$ to denote the value of the field f of an object η . When the object is referred by its location ι , we also write $\iota.f$ to refer to the value of its field f .

We overload the function *type* to accept (1) primitive value and return the primitive type; (2) location and return the type of the dereferenced object; (3) object and return the object type; and (4) object field and return the field type.

The variable environment Π is such a stackable mapping. We write $\Pi[\nu/v]$ to denote an update of the value of the latest variable v in Π to ν . We write $\Pi + \{v \mapsto \nu\}$ to denote an extension of Π to include a binding of ν to v , while $\Pi - \{v^*\}$ removes a subset of the mappings. Similar notations are used for the update and enhancement of object values and stores. In the case of store, we also provide an abbreviated notation $\varpi[\nu/\iota.f] =_{df} \text{let } (t, \xi) = \varpi(\iota) \text{ in } \varpi[(t, \xi[\nu/f])/\iota]$. Given an object value, $\eta = (t, \xi)$, we have $\text{Flds}(\eta) =_{df} \xi$.

We require some intermediate expressions for the dynamic semantics to follow through. Our syntax is thus extended from the original expression syntax as follows:

$$e ::= \dots \mid \eta \mid \iota \mid \nu \mid \text{ret}_d(v^*, e) \mid \text{ret}_m(Q, v^*, \tau, e)$$

The expression $\text{ret}_d(v^*, e)$ is used to capture the result of evaluating a local block, and $\text{ret}_m(Q, v^*, \tau, e)$ captures the result of method invocation. The set of variables v^* occurring in both result structures contain the local names and method parameters when entering local body and method body respectively. They are dropped at the end of the local/method body's evaluation. τ captures the type of the result of method invocation, whereas Q captures the set of type variables declared in the method header. Q is an instrument used to facilitate our soundness proof.

The dynamic evaluation rules are of the following form.

$$\langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \Pi', \varpi' \rangle [e']$$

We shall formulate the rules using an exception-style semantics with three possible errors, namely

E = Error-Null | Error-Cast | Error-Type.

Whenever one such error is raised, the evaluation aborts. This error occurrence can be stated using $\langle \Pi, \varpi \rangle [e] \hookrightarrow \mathbf{E}$. The small-step dynamic call-by-name semantics is formalised in Fig 11, together with some auxiliary functions in Fig 12.

B. Soundness of Type System

Before formulating the soundness, we extend the static semantics of the language to include those intermediate expressions given in Sec A. In the process, we require introduction of a *store typing* to describe the type of each location. This ensures that objects created in the store during run-time are type-wise consistent with that captured by the static semantics. Store typing is conventionally used to link static and dynamic semantics. In our case, it is denoted by:

$$\Sigma \in \text{StoreType} = \text{Location} \rightarrow_{\text{fin}} \text{Type}$$

Judgements in the static semantics will be extended with store typing, as follows:

$$\Gamma; \Sigma; Q \vdash e :: \tau, \psi$$

$$\boxed{\begin{array}{c} \text{[ELF}_m\text{]} \\ \frac{v^* \subseteq \text{dom}(\Gamma) \quad Q' \subseteq Q}{\Gamma; \Sigma; Q \vdash e :: \tau, \psi \quad \vdash \tau <: \tau_1 \Rightarrow \psi_1} \\ \Gamma; \Sigma; Q \vdash \text{ret}_m(Q', v^*, \tau, e) :: \tau_1, \psi \wedge \psi_1 \\ \\ \text{[ELF}_d\text{]} \qquad \text{[LOC]} \\ \frac{\Gamma; \Sigma; Q \vdash e :: \tau, \psi}{\Gamma; \Sigma; Q \vdash \text{ret}_d(v^*, e) :: \tau, \psi} \quad \frac{\tau = \Sigma(\iota) \quad \vdash \tau <: \tau_1 \Rightarrow \psi}{\Gamma; \Sigma; Q \vdash \iota :: \tau_1, \psi} \\ \\ \text{[OBJ]} \qquad \text{[VALUE]} \\ \frac{(t, \xi) = \eta \quad \vdash \odot t <: \tau \Rightarrow \psi}{\Gamma; \Sigma; Q \vdash \eta :: \tau, \psi} \quad \frac{\vdash \odot t <: \tau \Rightarrow \psi}{\Gamma; \Sigma; Q \vdash (t, \delta) :: \tau, \psi} \end{array}}$$

Figure 13. Type Rules for Intermediates

The static semantics for these intermediate expressions is shown in Figure 13.

The soundness of our static semantics relies on the following consistency relationship between the static and dynamic semantics, defined as follows:

$$\begin{array}{l} \text{dom}(\Pi) = \text{dom}(\Gamma) \quad \text{dom}(\varpi) = \text{dom}(\Sigma) \quad V_L = \text{vars}(\psi) - Q \\ \forall v \in \text{dom}(\Pi) \cdot \forall \rho_1 \in \text{Subs} \cdot \exists \rho_L \in \text{Subs} \cdot \\ (\text{dom}(\rho_L) = V_L \wedge \rho = \rho_1 \circ \rho_L \wedge (\rho(\psi) \Rightarrow \\ (\Pi(v) \in \text{prim} \Rightarrow \text{type}(\Pi(v)) <: \rho(\Gamma(v))) \wedge \\ (\Pi(v) \in \text{Location} \Rightarrow \\ \text{type}(\varpi(\Pi(v))) <: \rho(\Gamma(v)))))) \\ \hline \Gamma; \Sigma; Q; \psi \models \Pi, \varpi \end{array}$$

In the above relation, ρ_L is a ground substitution of local type variables occurring in the constraint ψ , and the composition of substitutions is recursively defined as: $(\rho_1 \circ \rho_2)(v) = \text{if } (v \in \text{dom}) \text{ then } \rho_2(v) \text{ else } \rho_1(v)$.

The following theorem states the progress of well-typed expressions.

THEOREM 1 (PROGRESS) *Let Γ be an environment mapping program variables to ground types. If $\Gamma; \Sigma; Q \vdash e :: \tau, \psi$ and $\Gamma; \Sigma; Q; \psi \models \Pi, \varpi$, then either*

- e is a value, or
- $\langle \Pi, \varpi \rangle [e] \hookrightarrow \text{Error-Null} \mid \text{Error-Cast}$, or
- there exist Π', ϖ', e' such that $\langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \Pi', \varpi' \rangle [e']$.

A proof of Theorem 1 can be found in Appendix C.1.

The next theorem states that each well-typed expression preserves its type under reduction with a runtime environment and a store which are consistent with the compile-time counterparts.

THEOREM 2 (PRESERVATION) *Let Γ be an environment mapping program variables to ground types. If*

$$\begin{array}{l} \Gamma; \Sigma; Q \vdash e :: \tau, \psi \\ \Gamma; \Sigma; Q; \psi \models \Pi, \varpi \\ \langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \hat{\Pi}, \hat{\varpi} \rangle [e] \end{array}$$

then there exists $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} such that

$$\begin{array}{l} \Gamma - \text{diff}(e, \hat{e}) = \hat{\Gamma} - \text{diff}(\hat{e}, e) \\ \hat{\Sigma} \supseteq \Sigma \\ \hat{\Gamma}; \hat{\Sigma}; \hat{Q} \vdash \hat{e} :: \tau, \hat{\psi} \\ \hat{\Gamma}; \hat{\Sigma}; \hat{Q}; \hat{\psi} \wedge \psi \models \hat{\Pi}, \hat{\varpi}. \end{array}$$

$\frac{\text{[D-Const]}}{k \text{ has type } t} \frac{}{\langle \Pi, \varpi \rangle [k] \hookrightarrow \langle \Pi, \varpi \rangle [(t, k)]}$	$\frac{\text{[D-Var-FD]}}{w = v v.f \quad \nu = \text{read}(\Pi, \varpi, w)} \frac{}{\langle \Pi, \varpi \rangle [w] \hookrightarrow \langle \Pi, \varpi \rangle [\nu]}$	$\frac{\text{[D-If-true]}}{\Pi(v) = (\text{Bool}, \text{true})} \frac{}{\langle \Pi, \varpi \rangle [\text{if } v \text{ then } e_1 \text{ else } e_2] \hookrightarrow \langle \Pi, \varpi \rangle [e_1]}$
$\frac{\text{[D-Assign-1]}}{\langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \Pi', \varpi' \rangle [e']} \frac{}{\langle \Pi, \varpi \rangle [w = e] \hookrightarrow \langle \Pi', \varpi' \rangle [w = e']}$	$\frac{\text{[D-Assign-2]}}{(\Pi', \varpi') = \text{upd}(\Pi, \varpi, w, \nu)} \frac{}{\langle \Pi, \varpi \rangle [w = \nu] \hookrightarrow \langle \Pi', \varpi' \rangle [(\text{void}, ())]}$	$\frac{\text{[D-If-false]}}{\Pi(v) = (\text{Bool}, \text{false})} \frac{}{\langle \Pi, \varpi \rangle [\text{if } v \text{ then } e_1 \text{ else } e_2] \hookrightarrow \langle \Pi, \varpi \rangle [e_2]}$
$\frac{\text{[D-Blk-1]}}{\langle \Pi, \varpi \rangle [e_1] \hookrightarrow \langle \Pi', \varpi' \rangle [e'_1]} \frac{}{\langle \Pi, \varpi \rangle [\{t v = e_1; e_2\}] \hookrightarrow \langle \Pi', \varpi' \rangle [\{t v = e'_1; e_2\}]}$	$\frac{\text{[D-Blk-2]}}{\text{subType}(\text{type}(\nu), t) \quad \Pi' = \Pi + \{v \mapsto \nu\}} \frac{}{\langle \Pi, \varpi \rangle [\{t v = \nu; e_2\}] \hookrightarrow \langle \Pi', \varpi \rangle [\text{ret}_d(v, e_2)]}$	
$\frac{\text{[D-While-true]}}{\Pi(v) = (\text{Bool}, \text{true})} \frac{}{\langle \Pi, \varpi \rangle [\text{while } v \text{ do } e] \hookrightarrow \langle \Pi, \varpi \rangle [e; \text{while } v \text{ do } e]}$	$\frac{\text{[D-While-false]}}{\Pi(v) = (\text{Bool}, \text{false})} \frac{}{\langle \Pi, \varpi \rangle [\text{while } v \text{ do } e] \hookrightarrow \langle \Pi, \varpi \rangle [(\text{void}, ())]}$	
$\frac{\text{[D-Ret-d-1]}}{\langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \Pi', \varpi' \rangle [e']} \frac{}{\langle \Pi, \varpi \rangle [\text{ret}_d(v^*, e)] \hookrightarrow \langle \Pi', \varpi' \rangle [\text{ret}_d(v^*, e')]}$	$\frac{\text{[D-Ret-d-2]}}{\Pi' = \Pi - (v^*)} \frac{}{\langle \Pi, \varpi \rangle [\text{ret}_d(v^*, \nu)] \hookrightarrow \langle \Pi', \varpi \rangle [\nu]}$	
$\frac{\text{[D-Ret-m-1]}}{\langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \Pi', \varpi' \rangle [e']} \frac{}{\langle \Pi, \varpi \rangle [\text{ret}_m(Q, v^*, t, e)] \hookrightarrow \langle \Pi', \varpi' \rangle [\text{ret}_m(Q, v^*, t, e')]}$	$\frac{\text{[D-Ret-m-2]}}{\text{subType}(\text{type}(\nu), t) \quad \Pi' = \Pi - (v^*)} \frac{}{\langle \Pi, \varpi \rangle [\text{ret}_m(Q, v^*, t, \nu)] \hookrightarrow \langle \Pi', \varpi \rangle [\nu]}$	$\frac{\text{[D-Seq-1]}}{\langle \Pi, \varpi \rangle [e_1] \hookrightarrow \langle \Pi', \varpi' \rangle [e'_1]} \frac{}{\langle \Pi, \varpi \rangle [e_1; e_2] \hookrightarrow \langle \Pi', \varpi' \rangle [e'_1; e_2]}$
$\frac{\text{[D-Cast]}}{\langle \Pi, \varpi \rangle [v] \hookrightarrow \langle \Pi, \varpi \rangle [\nu]} \frac{\text{chkCast}(\text{type}(\nu), t)}{\langle \Pi, \varpi \rangle [(t) v] \hookrightarrow \langle \Pi, \varpi \rangle [\nu]}$	$\frac{\text{[D-Capture]}}{\langle \Pi, \varpi \rangle [v] \hookrightarrow \langle \Pi, \varpi \rangle [\nu] \quad t_0 = \text{type}(\nu)} \frac{\rho = \text{match}(t, t_0) \quad (\Pi', \varpi') = \text{upd}(\Pi, \varpi, v_1, \nu)}{\langle \Pi, \varpi \rangle [\{v_1 = (t) v; e\}] \hookrightarrow \langle \Pi', \varpi' \rangle [\rho(e)]}$	$\frac{\text{[D-Seq-2]}}{\langle \Pi, \varpi \rangle [\delta; e_2] \hookrightarrow \langle \Pi, \varpi \rangle [e_2]}$
$\frac{\text{[D-New]}}{\text{class } c(X_i)_{i=1}^q \dots \text{ where } \psi \{ \dots \} \in P \quad \iota = \text{fresh}()}{\mu = [t_i / X_i]_{i=1}^q \quad \nu_i = \text{read}(\Pi, \varpi, v_i) \quad \forall i \in \{1..p\} \\ \text{chk}(\mu(\psi)) \quad t'_i = \text{type}(\nu_i) \quad \forall i \in \{1..p\} \\ \text{subType}(\text{type}(c(t'_i)_{i=1}^q), c(t_i)_{i=1}^q)} \frac{}{\eta = (c(t_i)_{i=1}^q, \{f_i \mapsto \nu_i\}_{i=1}^p), \varpi' = \varpi + \{\iota \mapsto \eta\}} \frac{}{\langle \Pi, \varpi \rangle [\text{new } c(t_i)_{i=1}^q (v_{1..p})] \hookrightarrow \langle \Pi, \varpi' \rangle [\iota]}$	$\frac{\text{[D-Call]}}{\nu_i = \Pi(v'_i) \quad \forall i \in \{0..q\} \quad c(t'_i)_{i=1}^m = \text{type}(\nu_0)} \frac{t_0 \mid t \text{ mn}((t_i v_i)_{i=1..q})(V^*) \text{ where } \psi \text{ eb} \in \text{mtds}(c)}{\mu = [t^* / V^*] \quad \text{chk}(\mu(\psi)) \quad \Pi' = \Pi + [\nu_0 / \text{this}] [\nu_i / v_i]_{i=1}^q \\ \text{subType}(\text{type}(\nu_i), \mu(t_i)) \quad \forall i \in \{0..q\}} \frac{V' = \{\text{this}\} \cup \{v_i\}_{i=1}^q \quad e = \text{ret}_m(V^*, V', \mu(t), \mu(\text{eb}))}{\langle \Pi, \varpi \rangle [v'_0.\text{mn}(v'_1, \dots, v'_q)(t^*)] \hookrightarrow \langle \Pi', \varpi \rangle [e]}$	

Figure 11. Dynamic Semantics

$\text{read}(\Pi, \varpi, v) = \Pi(v);$ $\text{read}(\Pi, \varpi, v.f) =$ $\iota = \Pi(v);$ $\text{if } \varpi(\iota) = \text{null} \text{ throw Error-Null};$ $\varpi(\iota).f;$ $\text{chk}(\phi) =$ $\text{if } \neg(\vdash \phi) \text{ throw Error-Type};$ $\text{true};$ $\text{chkCast}(t_1, t_2) =$ $\text{if } \neg(\vdash t_1 <: t_2) \text{ throw Error-Cast};$ $\text{true};$	$\text{upd}(\Pi, \varpi, v, \nu_s) =$ $\nu = \Pi(v);$ $\text{if } \neg(\vdash \text{type}(\nu_s) <: \text{type}(\nu))$ $\text{throw Error-Type};$ $(\Pi[\nu_s / v], \varpi);$ $\text{upd}(\Pi, \varpi, v.f, \nu_s) =$ $\iota = \Pi(v);$ $\text{if } \varpi(\iota) = \text{null} \text{ throw Error-Null};$ $\nu_f = \varpi(\iota).f;$ $\text{if } \neg(\vdash \text{type}(\nu_s) <: \text{type}(\nu_f)) \text{ throw Error-Type};$ $(\Pi, \varpi[\nu_s / \varpi(\iota).f]);$	$\text{subType}(t_1, t_2) =$ $\text{if } \neg(\vdash t_1 <: t_2) \text{ throw Error-Type};$ $\text{true};$ $\text{match}(t_v, t) = [t / t_v];$ $\text{match}(c(t_v^*), c(t^*)) = [t^* / t_v^*];$ $\text{match}(t', t) = \text{throw Error-Type};$
---	--	---

Figure 12. Dynamic Semantics: Auxiliary Definitions

Function $\text{diff}(e, e')$ returns a list of local variables that appears in e but not e' :

$$\begin{aligned} \text{diff}(e, e') &=_{df} \text{let } \text{lst} = \text{local}(e) \\ &\quad \text{lst1} = \text{local}(e') \\ &\quad n = \text{length}(\text{lst}) - \text{length}(\text{lst1}) \\ &\quad \text{in } (\text{take}(n, \text{lst}) \triangleleft n \geq 0 \triangleright []) \\ \text{take}(n, \text{lst}) &=_{df} ([] \triangleleft n \leq 0 \triangleright [\text{head}(\text{lst})] ++ \text{take}(n-1, \text{tail}(\text{lst}))) \\ x \triangleleft b \triangleright y &=_{df} \text{if } b \text{ then } x \text{ else } y \end{aligned}$$

Function $\text{local}(e)$ returns a list of sets of local variables. It is defined as follows:

$$\begin{aligned} \text{local}(e) &=_{df} \text{case } e \text{ of} \\ \text{ret}_a(Q, v^*, \tau, e) &\rightarrow \text{local}(e) ++ \{v^*\} \\ \text{ret}_a(v^*, e) &\rightarrow \text{local}(e) ++ \{v^*\} \\ w = e &\rightarrow \text{local}(e) \\ (t v = e_1; e_2) &\rightarrow \text{local}(e_1) \\ \text{otherwise} &\rightarrow \emptyset \end{aligned}$$

Note that $\Gamma - [] =_{df} \Gamma$, $\Gamma - ([s] ++ S) =_{df} (\Gamma - s) - S$. The proof is simple: by induction over the depth of type derivation of expression e . A proof of Theorem 2 can be found in Appendix C.2.

C. Proofs of Theorems

C.1 Proof of Theorem 1 (Progress)

By induction over the depth of type derivation for expression e .

Cases [NULL, VOID, VALUE, LOC, OBJ]. Trivial.

Case [VAR-FIELD]. We deal with expression w . As $w = v \mid v.f$ is well-typed, the evaluation rule [D-Var-FD] is followed, the evaluation either reports an **Error-Null** or advances one step yielding a value.

Case [ASSIGN]. We deal with expression $w = e$. From type rule, we have $\Gamma; \Sigma; Q \vdash e :: \oplus t, \psi$. By induction hypothesis, either (i) e is a value ν , or (ii) $\langle \Pi, \varpi \rangle [e] \hookrightarrow \text{Error}$, or (iii) $\langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \Pi', \varpi' \rangle [e']$.

Subcase (i): Let the runtime type of ν be \hat{t} , and that of w be t_1 . Then, we have $\odot \hat{t} < \oplus t$ and $\odot t_1 < \ominus t$, which implies $\hat{t} < t < t_1$. Hence, the upd function at [D-Assign-2] will not throw **Error-Type** exception, and proceed to update the runtime environment Π or the runtime store, as described in [D-Assign-2].

Subcase (ii): This will result in the execution of $\langle \Pi, \varpi \rangle [w = e]$ aborted with **Error**.

Subcase (iii): This will result in the execution of the assignment to reach $\langle \Pi, \varpi \rangle [w = e']$, via [D-Assign-1].

Case [SEQ]. We have $\Gamma; \Sigma; Q \vdash e_1 :: \otimes t, \psi$. By induction hypothesis, either (i) e_1 is a value ν , or (ii) $\langle \Pi, \varpi \rangle [e_1] \hookrightarrow \text{Error}$, or (iii) $\langle \Pi, \varpi \rangle [e_1] \hookrightarrow \langle \Pi', \varpi' \rangle [e'_1]$.

Subcase (i): The execution proceeds to reach $\langle \Pi, \varpi \rangle [e_2]$ unconditionally, according to [D-Seq-2].

Subcase (ii): The execution will be aborted with **Error** exception.

Subcase (iii): The execution proceeds to reach $\langle \Pi', \varpi' \rangle [e'_1; e_2]$, according to [D-Seq-1].

Case [LOCAL]. Given that $\Gamma; \Sigma; Q \vdash \{t v = e_1; e_2\} :: \tau, \psi_1 \wedge \psi_2$. We have $\Gamma; \Sigma; Q \vdash e_1 :: \oplus t, \psi_1$. By induction hypothesis, either (i) e_1 is a value ν , or (ii) $\langle \Pi, \varpi \rangle [e_1] \hookrightarrow \text{Error}$, or (iii) $\langle \Pi, \varpi \rangle [e_1] \hookrightarrow \langle \Pi', \varpi' \rangle [e'_1]$.

Subcase (i): Let the runtime type of ν be \hat{t}_0 and the runtime type of v be \hat{t} . As the consistency relation holds between the static and the dynamic semantics, we have for all ground substitution ρ , $\vdash \rho(\psi_1) \Rightarrow \hat{t} = \rho(t)$. Since $\vdash \rho(\psi_1) \Rightarrow \odot \hat{t}_0 < \oplus t$, $\text{subType}(\text{type}(\nu), \hat{t}) = \text{subType}(\hat{t}_0, \hat{t}) = \text{true}$. Hence, the execu-

tion will proceed to the state $\langle \Pi', \varpi \rangle [\text{ret}_a(v, e_2)]$ according to [D-Blk-2].

Subcase (ii). The execution will throw the corresponding **Error** exception.

Subcase (iii). The execution will proceed to $\langle \Pi', \varpi' \rangle [\{t v = e'_1; e_2\}]$ according to [D-Blk-1].

Case [NEW]. Given $\Gamma; \Sigma; Q \vdash \text{new } c \langle t_i \rangle_{i=1}^q (v_1, \dots, v_p) :: \tau, \psi$, let \hat{t}_i (for all $i = 1..q$) and \hat{t}_{v_i} (for all $i = 1..p$) be the runtime types of type arguments and value arguments to **new**. Then we have, for all ground substitution ρ , $\vdash \rho(\psi) \Rightarrow \bigwedge_{i=1}^q (\hat{t}_i = \rho(t_i))$ and $\vdash \rho(\psi) \Rightarrow \bigwedge_{i=1}^p (\hat{t}_{v_i} < \rho(\Gamma(v_i)))$. Furthermore, $\vdash \rho(\psi) \Rightarrow \rho(\Gamma(v_i)) < t'_i$, for all i . Hence, both calls to chk and subType at runtime do not fail, and the execution proceeds to the state $\langle \Pi, \varpi' \rangle [l]$, where l is the location referencing the new object.

Case [COND]. Given $\Gamma; \Sigma; Q \vdash \text{if } v \text{ then } e_1 \text{ else } e_2 :: \tau, \psi$ and $\Gamma(v) < \oplus \text{Bool}$, the runtime value of v will either be **true**, **false**, or **null** (). In the first two subcases, the execution proceeds to next state according to the rules [D-If-true] and [D-If-false] respectively. In the last subcase, there is no corresponding dynamic rule, and exception **Error-Null** will be thrown.

Case [WHILE]. Given $\Gamma; \Sigma; Q \vdash \text{while } v \text{ do } e :: \tau, \psi$ and $\Gamma(v) < \oplus \text{Bool}$, the runtime value of v will either be **true**, **false**, or **null** (). In the first two subcases, the execution proceeds to next state according to the rules [D-While-true] and [D-While-false] respectively. In the last subcase, there is no corresponding dynamic rule, and exception **Error-Null** will be thrown.

Case [ELF_d, ELF_m]. We have $\Gamma; \Sigma; Q \vdash e :: \tau, \psi$ as the premise of the static semantics. By induction hypothesis, either (i) e is a value ν , or (ii) $\langle \Pi, \varpi \rangle [e]$ produces **Error**, or (iii) $\langle \Pi, \varpi \rangle [e] \hookrightarrow \langle \Pi', \varpi' \rangle [e']$.

Subcase (i): Let the runtime type of ν be \hat{t}_ν and that of return value be \hat{t} then for all ground substitution ρ we have $\vdash \rho(\psi) \Rightarrow \rho(\tau) = \odot \hat{t}$. Also, we have $\vdash \rho(\psi) \Rightarrow \hat{t}_\nu < \rho(\tau)$. Hence, the call to subType in the rule [D-Ret-2] returns **true**, and the execution proceeds to $\langle \Pi', \varpi \rangle [\nu]$ accordingly.

Subcase (ii): The execution will throw the corresponding **Error** exception, as no rule applies.

Subcase (iii): The execution step to the new state following rule [D-Ret-1].

Case [CAST]. Any type mismatch during cast will be captured by chkCast and **Error-Cast** exception will be thrown. Otherwise, casting will succeed and the execution proceeds to the next state $\langle \Pi, \varpi \rangle [(t, \nu)]$.

Case [CAPTURE]. We have $\Gamma; \Sigma; Q \vdash \{v_1 = (t)v; e\} :: \tau, \psi_1 \wedge \psi_2$. From its premise, we have $t = c(\odot V_i)_{i=1}^n$. Executing the expression v either yields an **Error** exception or returns a value ν . We consider the case where ν is returned. Let t_0 be the type of ν as declared in the runtime environment. The use of flow symbol \odot in t implies that $\text{match}(t, t_0)$ succeeds and produces ρ only when $\rho(t) = t_0$. Hence, by rule [D-Capture], the execution proceeds to the state $\langle \Pi', \varpi' \rangle [\rho e]$. Updating of v_1 does not fail, similar with [ASSIGN].

Case [CALL]. Given $\Gamma; \Sigma; Q \vdash v'_0.mn(v'_1, \dots, v'_q) \langle t^* \rangle : \tau, \psi$. Let the runtime type arguments be \hat{t}^* and the value arguments have type $\hat{t}_{v'_i}$ for $i = 0..q$. Also, the ground substitution μ in [D-Call] is an instance of ρ in [CALL], which makes ψ true. Thus, we have, $\vdash \mu(\psi) \Rightarrow \hat{t}_{v'_i} < \mu(\tau_i)$, $i = 0..q$, and $\vdash \hat{t}_0 < \mu(t_0)$. Hence, the call to subType in [D-Call] yields true,

and the execution proceeds to the state $(\Pi, \varpi)[e]$ according to [D-Call]. \square

C.2 Proof of Theorem 2 (Preservation)

The proof for Theorem 2 requires several lemmas.

LEMMA 3 (Type Substitution). *If $\Gamma; \Sigma; Q \vdash e :: \tau, \psi$, then for all substitution ρ such that $\vdash \rho(\psi)$, we have $\rho(\Gamma); \rho(\Sigma); Q \vdash \rho(e) :: \rho(\tau), \rho(\psi)$.*

The proof is by induction on a derivation of $\Gamma; \Sigma; Q \vdash e :: \tau, \psi$.

The next lemma, called *assumption weakening lemma*, states that the static judgment remains valid despite a variation of its assumption. This assumes the store type Σ to have unbounded mapping of locations to types. However, the type environment Γ takes the form of stackable mapping between variables and types, and is allowed to grow (by pushing in new mappings) and shrink (by popping out mappings from stack). The lemma states that such change to type environment preserves the type judgment, if the change are properly constrained.

LEMMA 4 (Assumption Weakening). *Given that the following judgment holds:*

$$\Gamma; \Sigma; Q \vdash e :: \tau, \psi$$

Let $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} be such that:

$$\begin{aligned} \text{vars}(e) &\subseteq \text{dom}(\Gamma) \cap \text{dom}(\hat{\Gamma}) \\ Q &\subseteq \hat{Q} \vee \hat{Q} \subseteq Q \\ \text{vars}(\psi) - Q &= \text{vars}(\psi) - \hat{Q} \\ \exists v^* \cdot (\Gamma - \{v^*\} = \hat{\Gamma}) \vee (\hat{\Gamma} - \{v^*\} = \Gamma) \\ \hat{\Sigma} &\supseteq \Sigma \end{aligned}$$

Then, there exists $\hat{\psi}$ such that $\vdash \hat{\psi} \Leftrightarrow \psi$ and

$$\hat{\Gamma}; \hat{\Sigma}; \hat{Q} \vdash e :: \tau, \hat{\psi}$$

The call $\text{vars}(e)$ returns all program variables occurring in e , whereas $\text{vars}(\psi)$ returns all (type) variables occurring in ψ .

Proof of Lemma 4: By structural induction on the static semantics of the form $\Gamma; \Sigma; Q \vdash e :: \tau, \psi$. For any $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} , we say that they satisfy the premises of the Lemma if the following holds:

$$\begin{aligned} \text{vars}(e) &\subseteq \text{dom}(\Gamma) \cap \text{dom}(\hat{\Gamma}) \\ Q &\subseteq \hat{Q} \vee \hat{Q} \subseteq Q \\ \text{vars}(\psi) - Q &= \text{vars}(\psi) - \hat{Q} \\ \exists v^* \cdot (\Gamma - \{v^*\} = \hat{\Gamma}) \vee (\hat{\Gamma} - \{v^*\} = \Gamma) \\ \hat{\Sigma} &\supseteq \Sigma \end{aligned}$$

Cases [NULL, VOID, LOC, OBJ, VALUE]. Trivial.

Case [VAR-FIELD]. We deal with expression w , where $w = v \mid v.f$. For any $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} satisfying the premise of the lemma, we see that $\Gamma(v) = \hat{\Gamma}(v)$. Hence, $\hat{\Gamma}; \hat{\Sigma}; \hat{Q} \vdash w : \tau, \psi$.

Case [ASSIGN]. We deal with expression $w = e$. We have $\hat{\Gamma}; \hat{\Sigma}; \hat{Q} \vdash e :: \oplus t, \hat{\psi}$ for $\alpha t = \text{GetType}(\hat{\Gamma}, w) = \text{GetType}(\Gamma, w)$. The desired result is then derived by induction hypothesis.

Cases [LOCAL, SEQ, COND, WHILE, CAST, CAPTURE, ELF_d, ELF_m]. By induction hypothesis.

Case [NEW]. The result holds because $\Gamma(v_i) = \hat{\Gamma}(v_i)$, for all $i = 1..p$.

Cases [CALL]. The result holds because $\Gamma(v'_i) = \hat{\Gamma}(v'_i)$ for all $i = 1..q$. \square

Proof of Theorem 2 This can be proven by induction over the depth of type derivation of expression e .

Cases [NULL, VOID, LOC, OBJ, VALUE]. Vacuously true.

Case [VAR-FD]. This can be proven by setting $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} to Γ, Σ and Q respectively.

Case [ASSIGN]. There are two rules by which one-step derivation can be applied:

Subcase [D-Assign-1]: By induction hypothesis, there exists Γ', Σ' and Q' such that $\Gamma'; \Sigma'; Q' \vdash e' :: \oplus t', \psi'$ and which satisfies the premise of the theorem. Since $\oplus t' <: \oplus t$, we thus have $\Gamma'; \Sigma'; Q' \vdash e' :: \oplus t, \psi' \wedge \psi''$, where $\vdash \oplus t' <: \oplus t \Rightarrow \psi''$. The desired result can then be proven by setting $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} to Γ', Σ' and Q' respectively.

Subcase [D-Assign-2]: Consider the assignment to a variable v . Given that $\text{upd}(\Pi, \varpi, w, \nu)$ returns successfully (Π', ϖ') , it must be the case that $\text{type}(\nu) <: \text{type}(\Pi'(v))$. The desired result can then be proven by setting $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} to Γ, Σ and Q respectively. Similar argument applies to the assignment to a field.

Case [SEQ]. There are two rules by which one-step derivation can be applied:

Subcase [D-Seq-1]: By induction hypothesis, there exists Γ', Σ' and Q' that establishes the consistency relation at the hypothesis. We elect $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} to be Γ', Σ' and Q' respectively to obtain the desired result.

Subcase [D-Seq-2]: By setting $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} to be Γ, Σ and Q respectively.

Case [COND]. There are two rules by which one-step derivation can be applied: [D-If-True], [D-If-False]. Both can be proven by setting $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} to Γ, Σ and Q respectively.

Case [WHILE]. Similar as the argument for case [COND].

Case [LOCAL]. There are two rules to consider:

Subcase [D-Blk-1]: By induction hypothesis.

Subcase [D-Blk-2]: Since $\text{subType}(\text{type}(\nu), t)$, Γ' and Σ used in [LOCAL] remain consistent with Π' and ϖ in this subcase. We set $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} to Γ', Σ and Q respectively.

Case [CAST]. This can be proven by setting $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} to Γ, Σ and Q respectively.

Case [CAPTURE]. The argument for one-step derivation [D-Capture] is similar to that for case [D-Assign-2], except for the assignment of runtime type information of ν to the type variables occurring in t . This assignment proceeds successfully because of the premise of [CAPTURE]. We set $\hat{\Gamma}, \hat{\Sigma}$ and \hat{Q} to Γ, Σ and Q respectively. It suffices to show that $\hat{\Gamma}; \hat{\Sigma}; \hat{Q} \vdash \rho(e) :: \tau, \hat{\psi}$. This is true by applying Type Substitution Lemma to the following premise of [CAPTURE]: $\Gamma; \Sigma; Q \vdash e :: \tau, \psi_2$.

Case [NEW]. We set $\hat{\Gamma} = \Gamma, \hat{\Sigma} = \Sigma + \{t \mapsto \odot c(\odot t_i)_{i=1}^q\}$ and $\hat{Q} = Q$.

Case [CALL]. The fact that $\hat{\tau}$, as obtained from [ELF_m], is a subtype of τ obtained from [CALL], is established from the result of [ELF_m] and the constraint $\rho(\oplus t) <: \tau$ occurred in ψ in the premise of [CALL]. Finally, by assumption weakening rule, we set $\hat{\Gamma} = \Gamma + \{v_i :: \oplus t_i\}_{i=1}^q + \{\text{this} :: \oplus t_0, \hat{\Sigma} = \Sigma, \hat{Q} = Q \cup \{V^*\}$.

Case [ELF_d, ELF_m]. There are two subcases for consideration:

Subcase [D-Ret-d-1, D-Ret-m-1]: By induction hypothesis.

Subcase [D-Ret-d-2, D-Ret-m-2]: By induction hypothesis and the Assumption Weakening Lemma. \square