

Foundations for Virtual Types

Atsushi Igarashi and Benjamin C. Pierce

Department of Computer & Information Science
University of Pennsylvania
200 South 33rd St.
Philadelphia, PA 19104, USA
{igarasha,bcpierce}@saul.cis.upenn.edu

Abstract. *Virtual types* have been proposed as a notation for generic programming in object-oriented languages—an alternative to the more familiar mechanism of *parametric classes*. The tradeoffs between the two mechanisms are a matter of current debate: for many examples, both appear to offer convenient (indeed almost interchangeable) solutions; in other situations, one or the other seems to be more satisfactory. However, it has proved difficult to draw rigorous comparisons between the two approaches, partly because current proposals for virtual types vary considerably in their details, and partly because the proposals themselves are described rather informally, usually in the complicating context of full-scale language designs.

Work on the foundations of object-oriented languages has already established a clear connection between parametric classes and the polymorphic functions found in familiar typed lambda-calculi. Our aim here is to explore a similar connection between virtual types and dependent records. We present, by means of examples, a straightforward model of objects with embedded type fields in a typed lambda-calculus with subtyping, type operators, fixed points, dependent functions, and dependent records with both “bounded” and “manifest” type fields (this combination of features can be viewed as a measure of the inherent complexity of virtual types). Using this model, we then discuss some of the major differences between previous proposals and show why some can be checked statically while others require run-time checks. We also investigate how the partial “duality” of virtual types and parametric classes can be understood in terms of translations between universal and (dependent) existential types.

1 Introduction

Language support for generic programming plays an important role in the development of reusable libraries. In object-oriented languages, two different approaches to genericity have been considered. The more familiar one—based closely on the classical *parametric polymorphism* of functional languages such as ML and Haskell—can be found, for example, in the *template* mechanism of C++ [32] and the *parametric classes* in a number of proposed extensions to

Java [26, 25, 2, 3, 12, etc.]. An alternative approach, commonly called *virtual types* (or *virtual classes*), allows classes and objects to contain types as members, along with the usual fields and methods.¹ Virtual types were originally developed in Beta [23] and have recently been proposed for Java [33].

The static typing of virtual types is not yet clearly understood. Indeed, early proposals were statically unsafe, requiring extra runtime checks; more recent work has produced several proposals for type-safe variants [35, 5]. These proposals vary substantially in their details, and have generally been presented in rather informal terms—and in the complicating context of full-scale language designs—making them difficult to evaluate and compare.

Our goal in this paper is to establish a rigorous setting in which to understand and discuss the basic mechanisms of virtual types. Following a long line of past work on foundations for object-oriented programming (see [4] for history and citations), we model objects and classes with virtual types as a particular style of programming in a fairly standard typed lambda-calculus. On this basis, we examine (1) the type-theoretic features that seem to be required for modeling virtual types, (2) the similarities and differences between existing proposals, and (3) the type-theoretic intuitions behind the much-discussed “overlap” between virtual types and parametric classes in practice.

The rest of the paper is organized as follows. Section 2 reviews the idea of virtual types by means of a standard example, the animal/cow class hierarchy of Shang [31]. Section 3 sketches the main features of the typed lambda-calculus that forms the setting for our model. (The calculus is defined in full in Appendix A, for expert readers.) Section 4 develops the encoding of the Animal/Cow example in detail. Section 5 discusses the relation between virtual types and parametric classes as mechanisms for generic programming. Section 6 reviews previous work on virtual types in the light of our model. Section 7 sketches some directions for future work.

Our presentation is self-contained, but somewhat technical at times. Familiarity with past work on modeling objects in typed lambda-calculi (e.g., [29], [19], [4], or Chapter 18 of [1]) will help the reader interested in following in detail. Another useful source of background is Harper and Lillibridge [18, 21] and Leroy’s [20] papers on modeling module systems using dependent records with “manifest” bindings.

2 Virtual Types

We begin by reviewing the notion of virtual types through an example. This example, used throughout the paper, is a variant of the animal/cow example of Shang [31]. (Our notation is Java-like, but does not exactly correspond to any of the existing proposals for virtual types in Java.)

We begin by defining a generic class of animals, along with its interface.

¹ Referring to this approach with the phrase “virtual types” is somewhat confusing, since—as we will see—these type members may or may not be “virtual” in the sense of *virtual* or *abstract methods*. But the terminology is standard.

```

interface AnimalI {
    type FoodType <: Food;
    void eat (FoodType f);
    void eatALot (FoodType f); }
virtual class Animal
    implements AnimalI {
    virtual type FoodType <: Food;
    virtual void eat (FoodType f);
    void eatALot (FoodType f) {
        eat(f); eat(f); }}

```

Every animal has methods `eat` and `eatALot`, both accepting some food as an argument. The body of the `eat` method, which is specific to particular kinds of animals, is omitted; the `virtual` marker defers the responsibility of providing an implementation to subclasses. (We use the C++ keyword `virtual` in preference to Java’s `abstract` to avoid terminological confusion: locutions like “abstract type” already have a well-established meaning.) The calls to `eat` from the body of the `eatALot` method will call whatever body is provided by the subclass.

Similarly, the class `Animal` defers specifying exactly what kind of food a given kind of animal likes to eat. The virtual member `FoodType` acts as placeholder for this type, allowing it to be mentioned in the types of `eat` and `eatALot`, just as the declaration of `eat` provides a placeholder for its eventual implementation, allowing it to be referred to from the body of `eatALot`. Classes with virtual members (either types or methods) cannot be instantiated, since they are incomplete: they can only be subclassed.

The interface `AnimalI` specifies that every animal object has three members: a type `FoodType` and methods `eat` and `eatALot`. The `FoodType` member of every animal is known to be some kind of `Food` (`FoodType<:Food`), but, since different animals eat different kinds of food, the exact identity of this type is not visible. It follows immediately that it is not possible to feed an animal without knowing what kind of animal it is: if `a` is an object of type `AnimalI`, then `a`’s `eat` method requires an argument of type `a.FoodType`; but there is no way to obtain a value of this type (except, perhaps, by building a nutrient-free empty value using `new`).

Specific kinds of animals are modeled by classes inheriting from `Animal`. For example, here is a `Cow` class and its interface:

```

interface CowI
    extends AnimalI {
    type FoodType ↔ Grass; }
class Cow extends Animal
    implements CowI {
    final type FoodType ↔ Grass;
    void eat (FoodType f) { ... }}

```

In `Cow`, the virtual method `eat` is given a concrete implementation (shown as “...”). Similarly, the virtual type member `FoodType` is given a concrete value, `Grass`. The annotation `final` on the `FoodType` member means that it cannot be redefined by subclasses: every subclass of `Cow` is guaranteed to have `Grass` as its `FoodType`. The interface `CowI` reflects the fact that `FoodType` is `final`: in effect, it tells the world that every cow eats food whose type is *equal* to `Grass`. Thus, given an object `a` of type `CowI`, we may validly obtain some grass from any source and pass it to the `eat` or `eatALot` methods.

Virtual types are also useful in more standard examples of generic programming. For example, a generic `Bag` class can be defined with a virtual type

`ElementType`. Then classes `NatBag`, `StringBag`, etc. can be defined by inheriting from `Bag` and giving `ElementType` a `final` binding to `Nat` or `String`. Other examples of generic programming with virtual types can be found in [23, 33].

3 Summary of Type System

It is well understood [29, 6, etc.] how parametric classes—classes abstracted on type parameters—can be understood as polymorphic functions in a typed lambda-calculus. By analogy, objects with type members should clearly be modeled as some kind of records with type fields. Fortunately, such records have been studied extensively in the type-theory literature (e.g. [9]). Indeed, even the constraints on type members appearing in the interfaces `AnimalI (FoodType < Food)` and `CowI (FoodType ↔ Grass)` correspond to well-known constructions in the typed lambda-calculi used by Harper and Lillibridge [18, 21] and Leroy [20] to model module systems. Records with type fields constrained by `<` are a generalization of *partially abstract types* [11]; records with type fields constrained by `↔` correspond to *translucent* or *manifest* sums.

The typed lambda-calculus sketched in this section is based directly on these intuitions. In essence, it can be described as System F_{ω}^{\leq} (the omega-order polymorphic lambda-calculus with subtyping [8, 10, 27, 14]) plus dependent records with both “bounded” [11] and “manifest” [18, 21, 20] type fields, plus dependent functions.² We begin by briefly reviewing the features of System F_{ω}^{\leq} (Sections 3.1 and 3.2); we then concentrate on explaining records with type fields (Section 3.3) and dependent functions (Section 3.4), which are less familiar. Appendix A gives a more formal summary of the whole system.

3.1 Functions, Polymorphism, and Parameterized Types

The core of the system is Girard’s System F^{ω} [17]. This calculus can be viewed as a simple functional programming language with three distinct forms of abstraction: (1) *ordinary functions* (i.e., terms abstracted over terms); (2) *polymorphic functions* (i.e., terms abstracted over types); and (3) *parametric types* (i.e., types abstracted over types). We write all three forms with similar concrete syntax. For example,

$$\text{plustwo} = \lambda[x:\text{Nat}] \text{succ}(\text{succ}(x));$$

is an ordinary function that adds two to its argument. Similarly,

$$\text{id} = \lambda[X:*\] \lambda[x:X] x;$$

is the polymorphic identity function, and

² The system can also be described as an extension of Coquand and Huet’s *Calculus of Constructions* [15] with subtyping and dependent records [13]. Experts will note that we use an unstratified presentation of dependent records; this renders the system inconsistent as a logic, but the `fix` operator does that anyway.

```
double = λ[X:*] λ[f:X→X] λ[x:X] f(f(x));
```

is a polymorphic function that accepts a type X , a function f (of type $X \rightarrow X$), and an argument x (of type X), and applies f twice to x . (The annotation $X:*$ indicates that X is a type parameter.) Thus,

```
plusfour = double Nat plustwo;
```

is a fancy way of writing the function that adds four to its (numeric) argument. Parametric types are written in a similar style. For example,

```
Pair = λ[A:*] λ[B:*] {fst:A, snd:B};
```

is a convenient abbreviation for the parametric type of pairs, and

```
PairNatNat = Pair Nat Nat;
```

is the concrete type of pairs of numbers. The usual (polymorphic) operations on pairs can be defined as follows:

```
fst = λ[A:*] λ[B:*] λ[p: Pair A B] p.fst;
snd = λ[A:*] λ[B:*] λ[p: Pair A B] p.snd;
pair = λ[A:*] λ[B:*] λ[a:A] λ[b:B] ({fst=a, snd=b} :: Pair A B);
```

The types of these operations are:

```
fst : ∀[A:*] ∀[B:*] Pair A B → A
snd : ∀[A:*] ∀[B:*] Pair A B → B
pair : ∀[A:*] ∀[B:*] A → B → Pair A B
```

(In the following, we will often display defined terms together with their types.) Note that the definition of `pair` uses an explicit coercion (`:: Pair A B`) to control how its type is printed by the typechecker. Leaving it off results in a definition with exactly the same behavior

```
pair = λ[A:*] λ[B:*] λ[a:A] λ[b:B] {fst=a, snd=b};
pair : ∀[A:*] ∀[B:*] A → B → {fst:A, snd:B}
```

(since we have defined `Pair A B` to be interchangeable with `{fst:A, snd:B}`), but less intuitive for the reader.

To ensure their well-formedness, types and type operators are assigned *kinds*, K , which have the form $*$ or $K \rightarrow K$. Type expressions of kind $*$ (pronounced “type”) are ordinary types; type expressions of kind $* \rightarrow *$ are functions from types to types; etc.

It is sometimes useful to write *higher-order* type operators—that is, type operators whose arguments are type operators. For example,

```
BothBool = λ[F:*→*→*] F Bool Bool;
```

is higher-order type operator that, when applied to any operator O , yields the type $O \text{ Bool Bool}$. Thus:

```
mypair = pair Bool Bool true false :: BothBool Pair;
```

A more natural example of higher-order type operators will be seen later in the `Object` type constructor: its argument `I` is itself an operator abstracted over the “self type” `Rep`.

For constructing objects, we shall also need a fixed-point constructor. If `t` is a function from `T` to `T`, then `fix T t` is its fixed point. (Writing `T` explicitly simplifies the typechecking of `fix` in the presence of dependent types.)

$$\frac{\Gamma \vdash t : T \rightarrow T}{\Gamma \vdash \text{fix } T \ t : T}$$

For example, here is how `fix` is used to construct a factorial function:

```
fact = fix (Nat → Nat) λ[f:Nat → Nat] λ[n:Nat]
      if eq n 0 then 1 else times n (f (pred n)) :: Nat → Nat;
```

3.2 Subtyping

Next, we add the familiar notion of subtyping. For example, subtyping of function types is contravariant on the left and covariant on the right. The subtype relation has a maximal element, called `Top`. Constraining a type variable to be a subtype of `Top` is actually no constraint at all, so we can recover unbounded quantification from bounded, writing `λ[X<:Top]t` in place of `λ[X:*]t`. (We will continue to write `λ[X:*]t` in what follows, for readability.)

Subtyping is extended *pointwise* to type operators: `λ[X:K]S` is a subtype of `λ[X:K]T` if `S` is a subtype of `T` under all legal substitutions for `X`.

$$\frac{\Gamma, X:K \vdash S <: T}{\Gamma \vdash \lambda[X:K]S <: \lambda[X:K]T}$$

For example, `λ[T:*] Top → T` is a subtype of `λ[T:*] Nat → T` since `Nat` is a subtype of `Top`.

3.3 Records with Type Fields

To support records with type fields, a bit of machinery is required. First, we must deal with the fact that later fields in a record may refer to earlier fields by name—e.g., the type of the `eat` field must refer to the `FoodType` field. (Thus, in particular, the order of fields is significant in dependent records.) Second, we must be able to deal with record-projection expressions like `a.FoodType` appearing in the types of values (e.g., `a.eat`). The second requirement in particular goes somewhat beyond what can be expressed using ordinary existential types, taking us into the realm of *dependent records*.

In general, a dependent record has the form $\{\beta_i^{i \in 1 \dots n}\}$, where each β_i is a *field* of one of two forms: either a term field $x_i = t_i$ or a type field $X_i = T_i$. The name x_i or X_i is not only used to project a record from outside but also is a

binder whose scope is the rest of the fields in the record.³ For example, in the record value $\mathbf{r} = \{X=\text{Nat}, x=\lambda[y:X]y+1\}$, X in the second field is bound by the first occurrence of X .

A record type has the form $\Downarrow B_i^{i \in 1 \dots n} \Downarrow$, where B_i is a *binding* of one of three forms: a *term binding* $x:T$, a *bounded type binding* $X<:T$, or a *manifest type binding* $X\leftrightarrow T$. (In examples, we will also use type bindings of the form $X:*$ as an abbreviation for $X<:\text{Top}$.) For example, the record \mathbf{r} above has type $\Downarrow X\leftrightarrow\text{Nat}, x:X\rightarrow X \Downarrow$. A less informative type also possessed by \mathbf{r} is $\Downarrow X<:\text{Top}, x:X\rightarrow X \Downarrow$, which hides the representation of X and corresponds to the usual existential type $\exists X.X\rightarrow X$. In order to remind us of a connection to existential types, we sometimes write \exists before a field name in records or record types, like $\Downarrow \exists X<:\text{Top}, x:X\rightarrow X \Downarrow$, although \exists itself doesn't have a significant meaning. Formally, the typing rule for record introduction is:

$$\frac{\Gamma, B_1, \dots, B_{j-1} \vdash \beta_j : B_j \quad \Gamma \vdash \Downarrow B_i^{i \in 1 \dots n} \Downarrow : *}{\Gamma \vdash \{\beta_i^{i \in 1 \dots n}\} : \Downarrow B_i^{i \in 1 \dots n} \Downarrow}$$

Each field definition β_i must satisfy the corresponding binding B_i under a context augmented with the information of the preceding fields $(\Gamma, B_1, \dots, B_{i-1})$. Term fields $x_i=t_i$ satisfy bindings of the form $x_i:T_i$; type fields $X_i=T_i$ satisfy manifest type bindings $X_i\leftrightarrow T_i$. (Note that we cannot directly derive a record type with a bounded type binding using the rule above. For example, the type given to \mathbf{r} above is $\Downarrow \exists X\leftrightarrow\text{Nat}, x:X\rightarrow X \Downarrow$. If we want to hide the identity of X and give \mathbf{r} the abstract type $\Downarrow \exists X:*, x:X\rightarrow X \Downarrow$, we must use the usual subsumption rule plus the record subtyping rules discussed below.)

The rule for record projections is basically the same as the standard record elimination rule: if a field x of \mathbf{t} has binding $x:T$, then $\mathbf{t}.1$ has type T . If \mathbf{t} depends on other fields—that is, if the name X_i (or x_i) occurs free in T —then the corresponding record projection $\mathbf{t}.X_i$ (or $\mathbf{t}.x_i$, resp.) should be substituted for X_i (or x_i , resp.) to prevent the field name from escaping its scope.⁴

$$\frac{\Gamma \vdash \mathbf{t} : \Downarrow B_i^{i \in 1 \dots n} \Downarrow \quad B_j = x : T}{\Gamma \vdash \mathbf{t}.1_j : \{BV(B_i) \mapsto \mathbf{t}.1_i^{i \in 1 \dots j-1}\}T}$$

³ Strictly speaking, these two mechanisms should be kept separate. In the full typing rules in Appendix A, each field is given two names: an *external* name, which can be used for projections, and an *internal* name, which binds the subsequent occurrences in the record. The simplified syntax presented in the body of the paper corresponds to the special case where the external and internal names are identical.)

⁴ Experts will note that we give a somewhat simpler version of this rule than Harper and Lillibridge [18, 21] or Leroy's [20] formulations. The reason we can do this is that we are not—at this stage—considering computational effects such as references or exceptions. If any “effectful” constructs are added to the system, our projection rule needs to be refined to ensure soundness. This can be done in different ways, but the basic intuition is that a dependent projection $\mathbf{t}.1_j$ should be allowed only if the expression \mathbf{t} is *pure*. Similar comments apply to application rule below.

We write $BV(\mathbf{B})$ for the bound variable of the binding \mathbf{B} ; that is, $BV(\mathbf{x:T}) = \mathbf{x}$, $BV(\mathbf{X}\leftrightarrow\mathbf{T}) = \mathbf{X}$, and $BV(\mathbf{X}<\mathbf{T}) = \mathbf{X}$. We also write $\{\mathbf{X} \mapsto \mathbf{T}\}$ for capture-avoiding substitution of \mathbf{T} for \mathbf{X} .

The subtyping rule for record types is:

$$\frac{\vdash \Gamma, \mathbf{B}_1, \dots, \mathbf{B}_{n+k} \text{ ok} \quad \vdash \Gamma, \mathbf{B}'_1, \dots, \mathbf{B}'_n \text{ ok} \quad \Gamma, \mathbf{B}_1, \dots, \mathbf{B}_{j-1} \vdash \mathbf{B}_j <: \mathbf{B}'_j \quad j \in 1 \dots n}{\Gamma \vdash \{\mathbf{B}_i^{i \in 1 \dots n+k}\} <: \{\mathbf{B}'_i^{i \in 1 \dots n}\}}$$

As usual for ordinary (non-dependent) records, “width subtyping” is allowed: extra fields (the $n+1$ -st to $n+k$ -th fields) can be dropped. Also, corresponding bindings \mathbf{B}_i and \mathbf{B}'_i are compared using a *sub-binding* relation. When both are term bindings—i.e., \mathbf{B}_i and \mathbf{B}'_i are of the form $\mathbf{x:S}$ and $\mathbf{x:T} \text{---} \mathbf{S}$ should be a subtype of \mathbf{T} : this captures ordinary “depth subtyping.” For type bindings, we have $(\mathbf{X}\leftrightarrow\mathbf{T}) <: (\mathbf{X}<\mathbf{S}) <: (\mathbf{X}<\mathbf{U})$ if $\mathbf{T} <: \mathbf{S} <: \mathbf{U}$; the first clause ($(\mathbf{X}\leftrightarrow\mathbf{T}) <: (\mathbf{X}<\mathbf{S})$) allows the exact identity of a type field to be replaced with an upper bound; the second ($(\mathbf{X}<\mathbf{S}) <: (\mathbf{X}<\mathbf{U})$), corresponding to subtyping of bounded existential types, allows us to loosen the bound of \mathbf{X} . For example, we can derive $\{\exists \mathbf{X}\leftrightarrow\text{Nat}, \mathbf{x:X} \rightarrow \mathbf{X}\} <: \{\exists \mathbf{X} : *, \mathbf{x:X} \rightarrow \mathbf{X}\}$. (As usual, this rule leads to an undecidable subtyping relation [28, 21].)

3.4 Dependent Functions

For the encoding of classes, we will need to be able to give quite precise types to functions, showing the dependency of the *type* of the result on the *value* of the argument.

In outline, the intuition is this. Suppose we write a function

```
c1 = λ[self : {∃T : *, x:T, f:T→T}]
      {T=self.T, x=self.f(self.x), f=self.f};
```

whose argument is a record containing a type, a value (of that type), and a function (on that type), and whose result is a record with a similar shape, but where the value field is calculated by applying the argument’s function field to the argument’s value field. The type of this function

```
c1 : Π[self : {∃T : *, x:T, f:T→T}] {∃T↔self.T, x:T, f:T→T}
```

expresses the fact that the \mathbf{T} field of the result is identical to the \mathbf{T} field of the argument. Next, suppose we create a record containing these three items

```
r1 = {T=Nat, x=3, f=plusfour} :: {∃T : *, x:T, f:T→T};
```

```
r1 : {∃T : *, x:T, f:T→T}
```

and use the function `c1` to obtain another record of the same shape:

```
r2 = c1 r1;
```

```
r2 : {∃T↔r1.T, x:T, f:T→T}
```

Notice that, because of the dependent typing of `c1`, the type of `r2` exposes the fact that it was built from `r1`—in particular, that their type components are equal. Hence, it is legal to project the function field from `r2` and apply it to the value field from `r1`:

```
i = r2.f r1.x;
```

```
i : r2.T
```

In the absence of dependent functions, the best type we could have given to `c1` would be:

```
c1 : {∃T:*, x:T, f:T→T} → {∃T:*, x:T, f:T→T}
```

If we build `r2` from `r1` using this less refined type for `c1`,

```
r2 = c1 r1;
```

```
r2 : {∃T:*, x:T, f:T→T}
```

we obtain no information about the relation between `r1`'s `T` field and `r2`'s, and the application `r2.f r1.x` is not allowed.

In general, a function $\lambda[x:S]t$ has type $\Pi[x:S]T$, where x is allowed to appear in T . (When x does *not* appear in T , we write $\Pi[x:S]T$ as $S \rightarrow T$, recovering the usual notation for function types as a special case of dependent function types.) The rules for function abstraction and application are generalized accordingly:

$$\frac{\Gamma, x:S \text{ ok} \quad \Gamma, x:S \vdash t : T}{\Gamma \vdash \lambda[x:S]t : \Pi[x:S]T} \qquad \frac{\Gamma \vdash t : \Pi[x:S]T \quad \Gamma \vdash s : S}{\Gamma \vdash t \ s : \{x \mapsto s\}T}$$

4 Encoding Virtual Types

With the formalities of our typed lambda-calculus now in hand, we can proceed to the technical heart of the paper: a straightforward encoding of the animal example from Section 2 in terms of records with type fields. For the sake of concreteness, we extend the familiar *existential encoding* of objects [29, 19].

To avoid introducing additional complexities in the type theory, we give an encoding of *purely functional* objects; for example, we assume that an animal's `eat` method returns a new, satiated animal rather than side-effecting the internals of the receiving animal.

4.1 Interfaces

To get warmed up, let's begin with an example that does *not* involve virtual types: one-dimensional point objects with methods `get` to retrieve a current coordinate, `set` to move to a new coordinate, and `bump` to move a little from the present position.

In the simple existential encoding, the *interface* of an object is represented as a type operator of the form $\lambda[\text{Rep}:*] \llbracket m_i : T_i^{i \in 1 \dots n} \rrbracket$, where the bound variable Rep stands for the hidden type of the object's internal state, and where each T_i is type of the corresponding method m_i . Each method takes the internal state of the object as an explicit argument and, if appropriate, returns a new internal state as its result. For example, the interface `PointI` of point objects is represented as

```
PointI =  $\lambda[\text{Rep}:*] \llbracket \text{get}:\text{Rep} \rightarrow \text{Nat}, \text{set}:\text{Rep} \rightarrow \text{Nat} \rightarrow \text{Rep}, \text{bump}:\text{Rep} \rightarrow \text{Rep} \rrbracket;$ 

PointI :  $* \rightarrow *$ 
```

Interfaces for objects with virtual types may include not only methods but also type fields, which declare the bounds of the virtual types. The interface `AnimalI` is represented as

```
AnimalI =  $\lambda[\text{Rep}:*] \llbracket \exists \text{FT} < \text{Food}, \text{eat}:\text{Rep} \rightarrow \text{FT} \rightarrow \text{Rep}, \text{eatALot}:\text{Rep} \rightarrow \text{FT} \rightarrow \text{Rep} \rrbracket;$ 
```

The binding $\text{FT} < \text{Food}$ is a direct transliteration of the constraint on FT in Section 2. Similarly, the interface `CowI` is represented as

```
CowI =  $\lambda[\text{Rep}:*] \llbracket \exists \text{FT} \leftrightarrow \text{Grass}, \text{eat}:\text{Rep} \rightarrow \text{FT} \rightarrow \text{Rep}, \text{eatALot}:\text{Rep} \rightarrow \text{FT} \rightarrow \text{Rep} \rrbracket;$ 
```

where the binding of FT is now manifest. Note that `CowI` is a subtype of `AnimalI`; this will later allow `Cow` objects to be regarded as animals.

4.2 Objects

Intuitively, an object with interface I comprises some hidden internal state, some methods (described by I) that can manipulate that state, and some mechanism for hiding the type of the state from outside view. In the simple existential encoding, an existential quantifier is used to achieve this hiding (it can also be done with recursive types), so the type of our point objects is:

```
Point =  $\llbracket \exists \text{Rep}:*, \text{state}:\text{Rep},$ 
       $\text{meth}:\llbracket \text{get}:\text{Rep} \rightarrow \text{Nat}, \text{set}:\text{Rep} \rightarrow \text{Nat} \rightarrow \text{Rep}, \text{bump}:\text{Rep} \rightarrow \text{Rep} \rrbracket \rrbracket;$ 
```

More generally, the type of objects with interface I is a record type including a representation type Rep , a method vector field containing a record of type $I \text{ Rep}$, and a state field of type Rep . We can capture this structure uniformly by defining a (higher-order) type operator `Object` that takes I as a parameter:

```
Object =  $\lambda[I:*\rightarrow*] \llbracket \exists \text{Rep}:*, \text{meth}:I \text{ Rep}, \text{state}:\text{Rep} \rrbracket;$ 

Object :  $(*\rightarrow*) \rightarrow *$ 
```

The type `Point` is now expressed concisely as:

```
Point = Object PointI;
```

A point object—i.e., an element of type `Point`—can be constructed “from scratch” as follows (we will see how to create points from classes in Section 4.3):

```
PointR = {x:Nat};
point = {∃Rep=PointR,
  meth= fix (PointI Rep) λ[self:PointI Rep]
    {get= λ[s:Rep]s.x, set= λ[s:Rep]λ[n:Nat]{x=n},
      bump= λ[s:Rep] self.set s (plus 1 (self.get s))},
  state= {x=0}} :: Point;
```

`PointR` is the concrete representation type of the internal state. The method `get` just returns the `x` field of `state`, while `set` returns a new state with the `x` field set to its second argument, `n`. The method `bump` is defined in terms of the other methods `get` and `set`. In order to access other methods, the record of methods is abstracted on a parameter `self` of type `PointI Rep`; the fixed-point operator is used to “tie the knot,” making `self` refer to the record itself.

Invocation of the `get` method of a `Point` object requires simply extracting the `get` field of the object’s methods and applying it to the `state` field:

```
x = point.meth.get point.state :: Nat;
```

More generally, we can write

```
get = λ[p:Point] p.meth.get p.state :: Point → Nat;
```

for the function that “sends the `get` message” to an arbitrary point object `p`.

To send the `set` and `bump` messages to point objects, we need to do a little more work: the implementations of these methods return updated copies of just the internal representation, which must then be repackaged with the original methods into complete objects:

```
bump = λ[p:Point]{∃Rep=p.Rep, meth= p.meth,
  state= p.meth.bump p.state} :: Point → Point;
```

The construction of a `Cow` object is similar. The only significant difference is that the record of methods includes a type field `FT`, which should be given a concrete definition of food for a cow. Furthermore, methods taking arguments of `FT` can do grass-specific operation (such as `enoughGrass`) to the argument. Choosing the simple representation

```
CowR = {hungry:Bool};
```

for the internal state of cows, we can define an element of the type `Object CowI` as follows:

```
cow = {∃Rep=CowR,
  meth= fix (CowI Rep) λ[self:CowI Rep]
    {∃FT=Grass,
      eat=λ[s:Rep]λ[f:FT]
        if enoughGrass f then {hungry=false} else s,
      eatALot=λ[s:Rep]λ[f:FT](self.eat (self.eat s f) f)},
  state= {hungry=true}} :: Object CowI;
```

Like the `bump` method of point objects, the `eatALot` method of cows is defined by invoking the `eat` method via the `self` parameter.

Since we know `FT` is equal to `Grass` (by the definition of `CowI`), we can feed grass to our cow:

```
feed = λ[c:Object CowI] λ[g:Grass]
      {∃Rep=c.Rep, meth=c.meth, state=c.meth.eatALot c.state g}
      :: Object CowI → Grass → Object CowI;
satisfiedCow = feed cow grass :: Object CowI;
```

4.3 Classes

So far, virtual types have presented no special difficulties: the encodings of points and cows have been essentially identical. For encoding classes, however, the virtual types lead to some extra complications.

A *class* is a data structure providing implementations for a collection of methods and abstracted on a `self`-parameter. Concretely, a class whose instances are objects with interface `I` is represented as a function taking `self` as an argument and returning a record of methods of type `I R`, where `R` is the representation type of the state. For example, a class of point objects can be defined as follows:

```
pointClass = λ[self: PointI PointR]
            {get=λ[s:PointR]s.x, set=λ[s:PointR]λ[n:Nat]{x=n},
            bump=λ[s:PointR]self.set s (plus 1 (self.get s))}
            :: PointI PointR→PointI PointR;
```

To build a point object from the point class, we choose some particular representation (some element of type `PointR`) and calculate its record of methods by taking the fixed point of the class:

```
point = {∃Rep=PointR, meth=fix (PointI Rep) pointClass, state={x=0}}
       :: Object PointI;
```

The fact that the methods of `pointClass` are abstracted on `self` allows us to define new subclasses of `pointClass` that *inherit* some of its behavior. For example, here is a class of colored point objects:

```
CPointI = λ[Rep:*] {get:Rep→Nat, set:Rep→Nat→Rep,
                  bump:Rep→Rep, color:Rep→Color};
cpointClass = λ[self: CPointI PointR]
             let super = pointClass self in
             {get=super.get, set=super.set, bump=super.bump,
             color=λ[s:PointR] red}
             :: CPointI PointR → CPointI PointR;

cpoint = {∃Rep=PointR, meth=fix (CPointI Rep) cpointClass,
         state={x=0}} :: Object CPointI;
```

The superclass's method suite `super` is obtained by application of `pointClass` to (`cpointClass`'s) `self`. Note that, for brevity, we choose the same representation

type for both `pointClass` and `cpointClass`; it is easy to generalize this so that `cpointClass` can add new instance variables (such as a `color` field), but the extra mechanism would make the examples harder to read.

When virtual types are involved, we need to be a little more precise about the typing of classes. Here, for example, is the definition of a generic `animalClass`. (Again, for brevity we use the same representation type (`AnimalR`) for both `animalClass` and `cowClass`.)

```
AnimalR = {hungry:Bool};
animalClass = λ[self:AnimalI AnimalR]
  {∃FT=self.FT, eat=self.eat,
   eatALot=λ[s:AnimalR]λ[f:FT]self.eat (self.eat s f) f}
  :: Π[self: AnimalI AnimalR]
    {∃FT↔self.FT, eat: AnimalR→FT→AnimalR,
     eatALot: AnimalR→FT→AnimalR};
```

This definition involves a few subtle points. First, since the type `FT` and the method `eat` are virtual, their concrete definitions cannot be provided. Instead of concrete definitions, the corresponding fields of `self` are used. Second, type of `animalClass` is not `AnimalI AnimalR→AnimalI AnimalR`, but a dependent function type (a more refined subtype of `AnimalI AnimalR→AnimalI AnimalR`). This typing is essential when we derive `cowClass` from `animalClass`, as we will see below.

In the definition of `cowClass`, the `FT` and `eat` fields are filled with their concrete definitions and the `eatALot` method is inherited from `animalClass`. Since `cowClass`'s `self` is passed to `animalClass`, `self.eat` in method `eatALot` refer to the `eat` method of `cowClass` (not the virtual `eat` method of `animalClass`.) Now, since `FT` is not derived from `self`, the type of `cowClass` is just a (non-dependent) function type.

```
cowClass = λ[self:CowI CowR]
  let super = animalClass self in
  {∃FT=Grass,
   eat=λ[s:CowR]λ[f:FT]
     if enoughGrass f then {hungry=false} else s,
   eatALot=super.eatALot}
  :: CowI CowR → CowI CowR;
```

The dependent function type of `animalClass` is critical for `cowClass` to be well-typed: if `animalClass` had only type `AnimalI AnimalR→AnimalI AnimalR`, `cowClass` would be ill-typed since `super.eatALot` has type `CowR→FT→CowR` where `FT <: Food`, which is *not* a subtype of `CowR→Grass→CowR`. Thanks to the dependent function type of `cowClass`, the projection `super.eatALot` has type `CowR→self.FT→CowR`, which is exactly equal to `CowR→Grass→CowR`.

Finally, a `cow` object can be created by instantiating `cowClass` in the usual way:

```
cow = {∃Rep=CowR, meth=fix (CowI Rep) cowClass,
       state={hungry=true}} :: Object CowI;
```

5 Generic Programming with Virtual Types

The “overlap” between virtual types and parametric classes as alternative mechanisms for achieving similar kinds of genericity has been remarked by several authors [5, 34, etc.]. To build a generic `Bag` class, for example, one can proceed in two ways. On one hand, we can make the type of the bag’s elements a (virtual) field of the `Bag` class and obtain concrete instances by *subclassing* the generic `Bag` class, overriding the member type field with the actual member type. On the other hand, we can make the element type a parameter to the class definition, essentially making the class into a polymorphic function, and obtain concrete instances by *instantiating* this polymorphic function with the actual member type. In this section, we first compare these two styles by means of a fully worked example, then comment on the general case. The overlap between the styles can be viewed, in terms of our encoding, as a corollary of the inter-definability of universal and existential polymorphism in the presence of dependent records.

Generic programming was one of the first applications of virtual types. The typical pattern proceeds in two steps: (1) a generic class with a virtual type is defined, with generic implementations of its operations in terms of the virtual type; (2) this class is then specialized, overriding the virtual type to some concrete instance. For example, suppose we want to program with homogeneous collections (bags) of objects of some type T . We start by building a generic `Bag` class with a virtual type E (which stands for type of elements) and implementations of the bag methods (`put`, `get`, etc.). Since the representation type of state of bags is parameterized by E , the interface of bags takes a type operator `Rep` of kind $* \rightarrow *$, and the type of the state is actually represented as `Rep E`.

```
BagI = λ[Rep:*→*] {∃E:*, put:(Rep E)→E→(Rep E), get:(Rep E)→E};
```

Choosing lists of elements as our internal representation,

```
BagR = λ[E:*] {elts>List(E)};
```

we can define a generic bag class as follows:

```
bagClass = λ[self:BagI BagR]
  {∃E=self.E,
   put=λ[s:BagR E]λ[e:E]({elts= cons E e s.elts} ::BagR E),
   get=λ[s:BagR E] car E s.elts}
  :: Π[self:BagI BagR]
  {∃E↔self.E, put:BagR E→E→BagR E, get:BagR E→E};
```

The next step is to make a subclass with a concrete definition for the element type. The class `natBagClass` is defined by giving the concrete value `Nat` to the virtual type E and by inheriting all methods from `bagClass`.

```
NatBagI = λ[Rep:*→*]{∃E↔Nat,
  put:(Rep E)→E→(Rep E), get:(Rep E)→E};
natBagClass = λ[self:NatBagI BagR]
  let super = bagClass self in
  {∃E=Nat, put= super.put, get= super.get}
  :: NatBagI BagR → NatBagI BagR;
```

The interfaces and classes here are fairly similar to the examples we saw in Section 4 (modulo the fact that the representation type here is a type operator); the construction of bag *objects*, however, requires a little explanation.

```
NatBag = {∃Rep:*→*, ∃meth:NatBagI Rep, state:Rep meth.E};
natBag = {∃Rep=BagR, meth=fix (NatBagI Rep) natBagClass,
          state={elts=(nil Nat)}} :: NatBag;
```

The first observation is that the hidden state type is now a type operator. (Intuitively, we “see” that the representation of the object may involve the virtual type field `E`, but that is all we are allowed to know about the representation.) The second is that the order of the `state` field and the `meth` field is essential, since the type of the state depends both on `Rep` and on the `E` component of the `meth`. The code for invoking operations on bag objects is adjusted accordingly:

```
sendget = λ[b:NatBag] b.meth.get b.state :: NatBag→Nat;
sendput = λ[b:NatBag] λ[e:Nat]
          {∃Rep=b.Rep, meth= b.meth,
           state= b.meth.put b.state e} :: NatBag→Nat→NatBag;
```

By contrast, let’s look at how bags can be modeled in terms of parametric classes. Instead of the element type being a *member* of the bag class, it will be a *parameter* to the class. Similarly, the interface `BagI` is parameterized by `E`:

```
BagI = λ[E:*] λ[Rep:*] {put:Rep→E→Rep, get:Rep→E};
bagClass = λ[E:*] λ[self:BagI E (BagR E)]
           {put= λ[s:BagR E] λ[e:E] {elts= cons E e s.elts},
            get= λ[s:BagR E] car E s.elts}
           :: ∀[E:*] BagI E (BagR E)→BagI E (BagR E);
```

Note that `bagClass` has a polymorphic function type. (Also, note that `Rep` has kind `*` now, not `*→*`, since it is being supplied from the outside and there is no need to apply it to anything in this definition.)

The concrete instance `natBagClass` is now defined by *instantiating* `bagClass` with the type parameter `Nat`.

```
NatBagI = λ[Rep:*]{put:Rep→Nat→Rep, get:Rep→Nat};
natBagClass = bagClass Nat;
```

A bag object is defined by instantiating the class in the usual way. (Here there are no subtle dependencies between the type of the `meth` and `state` fields.)

```
NatBag = {∃Rep:*, meth:NatBagI Rep, state:Rep};
natBag = {∃Rep=BagR Nat,
          meth=fix (NatBagI Rep) natBagClass,
          state={elts=(nil Nat)}} :: NatBag;

sendget = λ[b:NatBag] b.meth.get b.state :: NatBag→Nat;
sendput = λ[b:NatBag] λ[e:Nat]
          {∃Rep=b.Rep, meth= b.meth,
           state= b.meth.put b.state e} :: NatBag→Nat→NatBag;
```

These examples illustrate the basic difference between virtual types and parametric classes as mechanisms for generic programming. A parametric class is instantiated by type application, taking the element type directly as an argument. With virtual types, on the other hand, type parameterization is realized by a dependent function whose argument has a type field in it. Since the `get` field depends on `self.E`, it will have type $(\text{List Nat}) \rightarrow \text{Nat}$ when the `E` field of the supplied `self` record has been set to `Nat`.

This correspondence can be viewed as an instance of a more general observation: polymorphic functions can be encoded in terms of dependent functions on dependent records. A polymorphic abstraction $\lambda[X<:S]t$ of type $\forall[X<:S]T$ can be represented as the dependent function $\lambda[x:\{ \exists X<:S \}] (\tau\{X \mapsto x.X\})$ of type $\Pi[x:\{ \exists X<:S \}] (T\{X \mapsto x.X\})$; it takes an argument $\{ \exists X \leftrightarrow U \}$ where U is some subtype of S and behaves as a term of type $T\{X \mapsto \{ \exists X \leftrightarrow U \}.X\}$, which is equal to the type $T\{X \mapsto U\}$ of the corresponding polymorphic application $(\lambda[X<:S]t) U$. The table below summarizes this encoding:

	\forall	$\exists + \Pi$
type	$\forall[X<:S] T$	$\Pi[x:\{ \exists X<:S \}] (T\{X \mapsto x.X\})$
abstraction	$\lambda[X<:S] t$	$\lambda[x:\{ \exists X<:S \}] (\tau\{X \mapsto x.X\})$
application	$t T$	$t \{ \exists X=T \}$

6 Comparisons

Virtual types (called *virtual classes* in the original proposal) were first introduced in Beta [24] by Madsen and Møller-Pedersen [23] as a mechanism to achieve genericity in object-oriented languages. Later, Thorup [33] introduced virtual types as an extension for Java. In all of this work, virtual types in classes are in fact *not* actually virtual in our sense: the interface of animal objects, according to their view, would better be modeled by

```
AnimalI = λ[Rep:*] { ∃ FT ↔ Food, eat : Rep → FT → Rep, eatALot : Rep → FT → Rep };
```

where `FT` is declared equal to `Food`. However, they also allow type fields to be specialized, so that

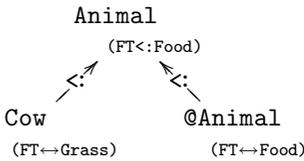
```
CowI = λ[Rep:*] { ∃ FT ↔ Grass, eat : Rep → FT → Rep, eatALot : Rep → FT → Rep };
```

as before. Finally, they want to regard cows as animals, i.e., `CowI <: AnimalI` and `Object CowI <: Object AnimalI`. Taken together, these properties (specifically, the inclusion `CowI <: AnimalI`) yield a statically unsafe type system: we can take a cow, regard it as an animal, and feed it some meat (which has type `Meat`, a subtype of `Food`, and hence an acceptable argument to an `Animal`'s `eat` method).

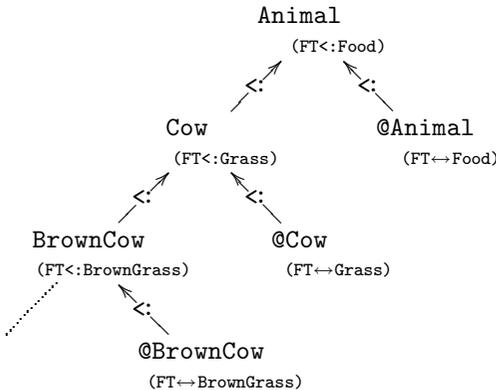
Various approaches have been suggested to remedy this unsoundness. In Beta and in Thorup's proposed Java extension, run-time checks are added to methods like `eat` to make sure that their arguments are actually acceptable. (In [22], it is observed that these checks can be omitted in the case where a type binding is marked `final`.)

Torgersen [35] proposed a statically typesafe variant of virtual types, introducing essentially the same distinction as we have made between virtual type bindings (which may be specialized in subclasses, but which block instantiation of the classes containing them) and final ones (which allow instantiation but block further specialization in subclasses). Our model of objects with virtual types corresponds closely to his proposal.

A possible criticism of Torgersen’s idea is that, in general, it may lead to duplication of the class hierarchy. For one thing, if the class `Animal` contains virtual types but no virtual methods (i.e., if `eat` is given a concrete generic implementation), then we may want to instantiate the class `Animal` itself. This requires making an explicit subclass (let’s call it `@Animal`) of `Animal` in which `FT` is equal to `Food`.



Also, rather than making `Cow` a leaf of the subclass hierarchy, we may wish to allow further specialization in subclasses. In this case, we should change the constraint on `FT` to `<:Grass`, make `Cow` a virtual class, and introduce another leaf class `@Cow` in which `FT↔Grass`.



Fortunately, the `@` variants can be derived mechanically from the other classes, as Torgersen himself pointed out in his original paper. More recently, Bruce, Odersky and Wadler [5] have proposed another statically safe variant of virtual types, which can be viewed as making this idea explicit. (They do not present their proposal in this light, but we find this to be a helpful way of understanding what they did.) In their system, virtual types are always introduced with `<:` constraints (they write “`FT as Food`”); for each class `C`, the “exact” class `@C` is automatically provided. The `new` operator generates instance of exact classes, so that the expression `new Cow ()` yields an object of type `@Cow`, which can be regarded as a `Cow` by forgetting its “exactness,” and further regarded as

an `Animal` (but not an `@Animal`) by ordinary subtyping. Note that their type system does not allow `@` types to have non-trivial subtypes, whereas `@Animal` here has many subtypes, which can be obtained by adding extra fields to `@Animal`. Their restriction becomes crucial when binary methods are involved, since an object type with binary methods will be expressed with recursive types where the recursion variable appears in contra-variant positions, which do not have any non-trivial subtypes.

Bruce, Odersky, and Wadler also pointed out that virtual types have an advantage over parametric classes in defining mutually recursive classes such as alternating lists or the Subject/Observer pattern [16]. In the Subject/Observer pattern, a group of objects (called subjects) has a reference to another group of objects (called observers) and reports their own behavior to observers, which will send back messages to subjects according to the reported behavior. Typically, a subject is realized by a class which has a virtual type bound to corresponding observers and vice versa. Then, generic subject (resp., observer) classes are extended to more specific classes, for example, window subject (resp., window observer) class by overriding virtual types with window observer (resp., window subject) and by implementing specific behavior of them. In [5], they used an extension of inner classes of Java to define mutually recursive classes, extensions (window subject/observer) had to be defined simultaneously.

Recently, Bruce and Vanderwaart [7] also used virtual types as a convenient device to define mutually recursive object types “incrementally”—just as extending an interface of Java, object types can be extended by adding specifications of new methods. Since their language can define object types separately from classes, a subject class and its corresponding observer class do not have to be defined simultaneously: virtual types will refer not to class names, but to object types. Rémy and Vouillon [30] showed programming with virtual types can be expressed in terms of parametric classes with mutually recursive types. Since their language has not only separate notion of object types but type reconstruction, programmers do not even need to write object types. As we discussed in Section 5, it is not so surprising that classes involving virtual types can be expressed in terms of parametric classes: an animal class would be just a parametric class which has a `FT` as a type parameter and there is no generic animal object types. However, they did not take into account type abstraction nature of virtual types. As for object types, our dependent record formulation seems to be essential, especially in order for cows to be animals.

7 Conclusions and Future Work

We have presented a straightforward encoding of objects with virtual types in a fairly standard (though quite powerful) type theory. In our model, objects are expressed as dependent records with manifest and/or bounded type fields; classes are modeled as dependent functions. The overlap between parametric classes and virtual types can then be viewed as a consequence of the encodability of universal

polymorphism in terms of existential polymorphism with dependent functions. We are working to extend this encoding in two main directions:

- Imperative variants of the encoding, where methods like `eat` work by side-effecting mutable instance variables.
- Recursive and mutually recursive classes involving virtual types, such as the well-known subject-observer example.

The second of these seems relatively straightforward. The first, somewhat surprisingly (and disappointingly) does not—the technicalities of the underlying type theory required to achieve soundness when imperative features are combined with dependent types become astonishingly subtle.

An obvious question is whether other type-theoretic encodings of simple objects—for example, the standard recursive-records encoding [4]—could be used instead of the existential encoding presented here. Surprisingly, we have *not* been able to extend a naive recursive-records encoding to include virtual types. Intuitively, the problem is that `Animal` in this encoding would be a recursive type whose body is a dependent record type with an FT field. But now every unfolding of the recursive type produces a *different* FT field, whose (abstract) type is incomparable with all the others.

Another interesting question is whether the type theory in which we are working here is the simplest possible for the task. All of the features described in Section 3—in particular, both dependent records and dependent functions—are used by our encoding, but it is possible that a different encoding could get by with less.

Finally, it would be worthwhile to formalize the translation from a high-level language with virtual types into low-level structures like the ones we have explored here. The interesting point is to see how much of of the type-theoretic complexity of the target language will also show up in the typing rules of the high-level language.

Acknowledgments

This work was supported by Indiana University, the University of Pennsylvania, and the National Science Foundation under grant CCR-9701826, *Principled Foundations for Programming with Objects*. Igarashi is a research fellow of the Japan Society for the Promotion of Science.

Discussions with Kim Bruce, Bob Harper, Didier Rémy, and Philip Wadler deepened our understanding of this material. Comments from the FOOL and ECOOP referees helped us improve the final presentation.

References

- [1] Martín Abadi and Luca Cardelli. *A Theory of Objects*. Springer-Verlag, 1996.
- [2] Ole Agesen, Stephen N. Freund, and John C. Mitchell. Adding type parameterization to the Java language. In *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, October 1997.
- [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 Craig Chambers, editor, *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, ACM SIGPLAN Notices volume 33 number 10, pages 183–200, Vancouver, BC, October 1998.
- [4] Kim B. Bruce, Luca Cardelli, and Benjamin C. Pierce. Comparing object encodings. *Information and Computation*, 1998. To appear in a special issue with papers from *Theoretical Aspects of Computer Software (TACS)*, September, 1997. An earlier version appeared as an invited lecture in the Third International Workshop on Foundations of Object Oriented Languages (FOOL 3), July 1996.
- [5] Kim B. Bruce, Martin Odersky, and Philip Wadler. A statically safe alternative to virtual types. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, 1998.
- [6] Kim B. Bruce, Angela Schuett, and Robert van Gent. PolyTOIL: A type-safe polymorphic object-oriented language. In W. Olthoff, editor, *Proceedings of ECOOP '95*, LNCS 952, pages 27–51, Aarhus, Denmark, August 1995. Springer-Verlag.
- [7] Kim B. Bruce and Joseph C. Vanderwaart. Semantics-driven language design: Statically type-safe virtual types in object-oriented languages. In *Fifteenth Conference on the Mathematical Foundations of Programming Semantics*, April 1999.
- [8] Luca Cardelli. Notes about F_{\leq}^{ω} . Unpublished manuscript, October 1990.
- [9] Luca Cardelli. Typeful programming. In E. J. Neuhold and M. Paul, editors, *Formal Description of Programming Concepts*. Springer-Verlag, 1991. An earlier version appeared as DEC Systems Research Center Research Report #45, February 1989.
- [10] Luca Cardelli and Giuseppe Longo. A semantic basis for Quest. *Journal of Functional Programming*, 1(4):417–458, October 1991. Preliminary version in ACM Conference on Lisp and Functional Programming, June 1990. Also available as DEC SRC Research Report 55, Feb. 1990.
- [11] Luca Cardelli and Peter Wegner. On understanding types, data abstraction, and polymorphism. *Computing Surveys*, 17(4):471–522, December 1985.
- [12] Robert Cartwright and Guy L. Steele Jr. Compatible genericity with run-time types for the Java programming language. In Craig Chambers, editor, *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, SIGPLAN Notices volume 33 number 10, pages 201–215, Vancouver, BC, October 1998. ACM.
- [13] Gang Chen and Giuseppe Longo. Subtyping parametric and dependent types. In Kamareddine et al., editor, *Type Theory and Term Rewriting*, September 1996. Invited lecture.
- [14] Adriana B. Compagnoni. Decidability of higher-order subtyping with intersection types. In *Computer Science Logic*, September 1994. Kazimierz, Poland. Springer *Lecture Notes in Computer Science* 933, June 1995. Also available as University of Edinburgh, LFCS technical report ECS-LFCS-94-281, titled “Subtyping in F_{λ}^{ω} is decidable”.
- [15] Thierry Coquand and Gérard Huet. The Calculus of Constructions. *Information and Computation*, 76(2/3):95–120, February/March 1988.
- [16] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison Wesley, Massachusetts, 1994.

- [17] Jean-Yves Girard. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. PhD thesis, Université Paris VII, 1972. A summary appeared in the Proceedings of the Second Scandinavian Logic Symposium (J.E. Fenstad, editor), North-Holland, 1971 (pp. 63–92).
- [18] Robert Harper and Mark Lillibridge. A type-theoretic approach to higher-order modules with sharing. In *Proceedings of the Twenty-First ACM Symposium on Principles of Programming Languages (POPL), Portland, Oregon*, pages 123–137, Portland, OR, January 1994.
- [19] Martin Hofmann and Benjamin Pierce. A unifying type-theoretic framework for objects. *Journal of Functional Programming*, 5(4):593–635, October 1995. Previous versions appeared in the Symposium on Theoretical Aspects of Computer Science, 1994, (pages 251–262) and, under the title “An Abstract View of Objects and Subtyping (Preliminary Report),” as University of Edinburgh, LFCS technical report ECS-LFCS-92-226, 1992.
- [20] Xavier Leroy. Manifest types, modules and separate compilation. In *Conference record of POPL '94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 109–122, Portland, OR, January 1994.
- [21] Mark Lillibridge. *Translucent Sums: A Foundation for Higher-Order Module Systems*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, May 1997.
- [22] Ole Lehrmann Madsen, Boris Magnusson, and Birger Møller-Pedersen. Strong typing of object-oriented languages revisited. In *Proceedings of the Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA) and European Conference on Object-Oriented Programming (ECOOP)*, pages 140–150, Ottawa, ON Canada, October 1990. ACM Press, New York, NY, USA. Published as SIGPLAN Notices, volume 25, number 10.
- [23] Ole Lehrmann Madsen and Birger Møller-Pedersen. Virtual classes: A powerful mechanism in object-oriented programming. In *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, 1989.
- [24] Ole Lehrmann Madsen, Birger Møller-Pedersen, and Kristen Nygaard. *Object-Oriented Programming in the Beta Programming Language*. Addison-Wesley, 1993.
- [25] Andrew C. Myers, Joseph A. Bank, and Barbara Liskov. Parametrized types for Java. In *ACM Symposium on Principles of Programming Languages (POPL)*, January 1997.
- [26] Martin Odersky and Philip Wadler. Pizza into Java: Translating theory into practice. In *ACM Symposium on Principles of Programming Languages (POPL)*, 1997.
- [27] Benjamin Pierce and Martin Steffen. Higher-order subtyping. In *IFIP Working Conference on Programming Concepts, Methods and Calculi (PROCOMET)*, 1994. Full version in *Theoretical Computer Science*, vol. 176, no. 1–2, pp. 235–282, 1997 (corrigendum in TCS vol. 184 (1997), p. 247).
- [28] Benjamin C. Pierce. Bounded quantification is undecidable. *Information and Computation*, 112(1):131–165, July 1994. Also in Carl A. Gunter and John C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design* (MIT Press, 1994). A preliminary version appeared in POPL '92.
- [29] Benjamin C. Pierce and David N. Turner. Simple type-theoretic foundations for object-oriented programming. *Journal of Functional Programming*, 4(2):207–247, April 1994. A preliminary version appeared in Principles of Programming Languages, 1993, and as University of Edinburgh technical report ECS-LFCS-92-225, under the title “Object-Oriented Programming Without Recursive Types”.
- [30] Didier Rémy and Jérôme Vouillon. On the (un)reality of virtual types, November 1998. manuscript.

- [31] David Shang. Are cows animals? *Object Currents 1*, 1996. <http://www.sigs.com/objectcurrents/>.
- [32] Bjarne Stroustrup. *The C++ Programming Language*. Addison Wesley Longman, Reading, MA, third edition, 1997.
- [33] Kresten Krab Thorup. Genericity in Java with virtual types. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, volume 1241 of *Lecture Notes in Computer Science*, pages 444–471. Springer-Verlag, 1997.
- [34] Kresten Krab Thorup and Mads Torgersen. Unifying genericity: Combining the benefits of virtual types and parametrized classes. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*. Springer-Verlag, June 1999.
- [35] Mads Torgersen. Virtual types are statically safe. In *5th Workshop on Foundations of Object-Oriented Languages (FOOL)*, January 1998.

A Summary of Calculus

This appendix summarizes the syntax, reduction rules, and typing rules of the type system used in this paper. All of the individual features of this calculus have been studied carefully in the literature, and their combination here is believed (but has not been proved!) to be sound. All of the examples in the paper have been checked mechanically by our implementation.

A.1 Syntax

The sets of kinds K , types T , bindings B , terms t , fields β , and contexts Γ are defined by the following grammar:

$K ::=$		$t ::=$	
$*$	proper types	x	variable
$K \rightarrow K$	operator kind	$\lambda [X <: T] t$	polymorphic abstraction
		$\lambda [x : T] t$	abstraction
$T ::=$		$t T$	polymorphic application
X	type variable	$t t$	application
Top	Top type	$\{l_i / \beta_i^{i \in 1 \dots n}\}$	record intro
$\forall [X <: T] T$	polymorphic func. type	$t.l$	record projection
$\Pi [x : T] T$	dependent func. type	$\text{fix } T t$	fixed-point operator
$\lambda [X : K] T$	type operator		
$T T$	type op. application	$\beta ::=$	
$\{l_i / B_i^{i \in 1 \dots n}\}$	type of records	$X = T$	type field
$t.l$	type field projection	$x = t$	term field
$B ::=$		$\Gamma ::=$	
$X <: T$	bounded type binding	\bullet	empty context
$X \leftrightarrow T$	manifest type binding	Γ, B	extended by binding
$x : T$	term binding	$\Gamma, X : K$	extended by type binding

A.2 Reduction

$$\begin{array}{c}
 (\lambda[x:T]s) \ t \longrightarrow \{x \mapsto t\}s \\
 \frac{s = \{l_i / \beta_i^{i \in 1 \dots n}\} \quad \beta_j = x = t}{s.l_j \longrightarrow \{BV(\beta_i) \mapsto s.l_i^{i \in 1 \dots j-1}\}t}
 \end{array}
 \qquad
 \begin{array}{c}
 (\lambda[X<:S]t) \ T \longrightarrow \{X \mapsto T\}t \\
 \text{fix } T \ s \longrightarrow s \ (\text{fix } T \ s)
 \end{array}$$

A.3 Judgments

$$\begin{array}{ll}
 \Gamma \vdash \text{ok} & \Gamma \text{ is a well-formed context} \\
 \Gamma \vdash T : K & T \text{ is a type constructor of kind } K \\
 \Gamma \vdash t : T & t \text{ is a term of type } T \\
 \Gamma \vdash \beta : B & \beta \text{ is a field of binding } B
 \end{array}
 \qquad
 \begin{array}{ll}
 \Gamma \vdash S <: T & S \text{ is a subtype of } T \\
 \Gamma \vdash B_1 <: B_2 & B_1 \text{ is a subbinding of } B_2 \\
 \Gamma \vdash S \leftrightarrow T & S \text{ and } T \text{ are equivalent} \\
 \Gamma \vdash B_1 \leftrightarrow B_2 & B_1 \text{ and } B_2 \text{ are equivalent}
 \end{array}$$

Context Well-formedness

$$\begin{array}{c}
 \vdash \bullet \text{ ok} \\
 \frac{\Gamma \vdash \text{ok} \quad X \notin \text{dom}(\Gamma) \quad \Gamma \vdash T : K}{\vdash \Gamma, X <: T \text{ ok}} \\
 \frac{\Gamma \vdash \text{ok} \quad x \notin \text{dom}(\Gamma) \quad \Gamma \vdash T : *}{\vdash \Gamma, x : T \text{ ok}}
 \end{array}
 \qquad
 \begin{array}{c}
 \frac{\vdash \Gamma \text{ ok} \quad X \notin \text{dom}(\Gamma)}{\vdash \Gamma, X : K \text{ ok}} \\
 \frac{\Gamma \vdash \text{ok} \quad X \notin \text{dom}(\Gamma) \quad \Gamma \vdash T : K}{\vdash \Gamma, X \leftrightarrow T \text{ ok}}
 \end{array}$$

Kinding Rules

$$\begin{array}{c}
 \Gamma \vdash \text{Top} : * \\
 \frac{X \leftrightarrow T \in \Gamma \quad \Gamma \vdash T : K}{\Gamma \vdash X : K} \\
 \frac{\vdash \Gamma, X <: S \text{ ok} \quad \Gamma, X <: S \vdash T : *}{\Gamma \vdash \forall [X <: S] T : *} \\
 \frac{\vdash \Gamma, X : K_1 \text{ ok} \quad \Gamma, X : K_1 \vdash T : K_2}{\Gamma \vdash \lambda [X : K_1] T : K_1 \rightarrow K_2} \\
 \frac{\vdash \Gamma, B_1, \dots, B_n \text{ ok}}{\Gamma \vdash \downarrow l_i / B_i^{i \in 1 \dots n} \uparrow : *}
 \end{array}
 \qquad
 \begin{array}{c}
 \frac{X : K \in \Gamma}{\Gamma \vdash X : K} \\
 \frac{X <: T \in \Gamma \quad \Gamma \vdash T : K}{\Gamma \vdash X : K} \\
 \frac{\vdash \Gamma, x : S \text{ ok} \quad \Gamma, x : S \vdash T : *}{\Gamma \vdash \Pi [x : S] T : *} \\
 \frac{\Gamma \vdash T : K_1 \rightarrow K_2 \quad \Gamma \vdash U : K_1}{\Gamma \vdash T \ U : K_2} \\
 \frac{\Gamma \vdash t : \downarrow l_i / B_i^{i \in 1 \dots n} \uparrow \quad B_j = X <: T \quad \Gamma, B_1, \dots, B_{j-1} \vdash T : K}{\Gamma \vdash t.l_j : K}
 \end{array}$$

Typing Rules

$$\begin{array}{c}
 \frac{x : T \in \Gamma}{\Gamma \vdash x : T} \\
 \frac{\Gamma \vdash t : \forall [X <: T] U \quad \Gamma \vdash S <: T}{\Gamma \vdash t \ S : \{X \mapsto S\}U}
 \end{array}
 \qquad
 \begin{array}{c}
 \frac{\vdash \Gamma, X <: S \text{ ok} \quad \Gamma, X <: S \vdash t : T}{\Gamma \vdash \lambda [X <: S] t : \forall [X <: S] T} \\
 \frac{\vdash \Gamma, x : S \text{ ok} \quad \Gamma, x : S \vdash t : T}{\Gamma \vdash \lambda [x : S] t : \Pi [x : S] T}
 \end{array}$$

$$\begin{array}{c}
\frac{\Gamma \vdash \mathbf{t} : \Pi[x:S]T \quad \Gamma \vdash \mathbf{s} : S}{\Gamma \vdash \mathbf{t} \ \mathbf{s} : \{x \mapsto \mathbf{s}\}T} \\
\frac{\Gamma \vdash \mathbf{t} : \{\llbracket l_i/B_i \rrbracket^{i \in 1 \dots n}\} \quad B_j = \mathbf{x} : T}{\Gamma \vdash \mathbf{t} . l_j : \{BV(B_i) \mapsto \mathbf{t} . l_i \rrbracket^{i \in 1 \dots j-1}\}T} \\
\frac{\Gamma \vdash \mathbf{t} : S \quad \Gamma \vdash S <: T}{\Gamma \vdash \mathbf{t} : T}
\end{array}
\qquad
\begin{array}{c}
\frac{\Gamma, B_1, \dots, B_{j-1} \vdash \beta_j : B_j \quad j \in 1 \dots n}{\Gamma \vdash \{\llbracket l_i/B_i \rrbracket^{i \in 1 \dots n}\} : *}}{\Gamma \vdash \{\llbracket l_i/\beta_i \rrbracket^{i \in 1 \dots n}\} : \{\llbracket l_i/B_i \rrbracket^{i \in 1 \dots n}\}T} \\
\frac{\Gamma \vdash \mathbf{t} : T \rightarrow T}{\Gamma \vdash \mathbf{fix} \ T \ \mathbf{t} : T}
\end{array}$$

Typing Rules for Record Fields

$$\frac{\Gamma \vdash T : K}{\Gamma \vdash (X=T) : (X \leftrightarrow T)}
\qquad
\frac{\Gamma \vdash \mathbf{t} : T}{\Gamma \vdash (\mathbf{x}=\mathbf{t}) : (\mathbf{x}:T)}$$

Subtyping Rules

$$\begin{array}{c}
\frac{\Gamma \vdash S \leftrightarrow T}{\Gamma \vdash S <: T} \\
\frac{X <: T \in \Gamma}{\Gamma \vdash X <: T} \\
\frac{\vdash \Gamma, X <: S_1 \text{ ok} \quad \vdash \Gamma, X <: S_2 \text{ ok}}{\Gamma \vdash S_2 <: S_1 \quad \Gamma, X <: S_2 \vdash T_1 <: T_2} \\
\frac{\Gamma \vdash \forall[X <: S_1]T_1 <: \forall[X <: S_2]T_2}{\Gamma \vdash \lambda[X : K]S <: \lambda[X : K]T} \\
\frac{\vdash \Gamma, B_1, \dots, B_{n+k} \text{ ok} \quad \vdash \Gamma, B'_1, \dots, B'_n \text{ ok}}{\Gamma, B_1, \dots, B_{j-1} \vdash B_j <: B'_j \quad j \in 1 \dots n} \\
\frac{\Gamma \vdash \{\llbracket l_i/B_i \rrbracket^{i \in 1 \dots n+k}\} <: \{\llbracket l_i/B'_i \rrbracket^{i \in 1 \dots n}\}}{\Gamma \vdash S <: T}
\end{array}
\qquad
\begin{array}{c}
\frac{\Gamma \vdash S <: T \quad \Gamma \vdash T <: U}{\Gamma \vdash S <: U} \\
\frac{\Gamma \vdash T : *}{\Gamma \vdash T <: \text{Top}} \\
\frac{\vdash \Gamma, \mathbf{x} : S_1 \text{ ok} \quad \vdash \Gamma, \mathbf{x} : S_2 \text{ ok}}{\Gamma \vdash S_2 <: S_1 \quad \Gamma, \mathbf{x} : S_2 \vdash T_1 <: T_2} \\
\frac{\Gamma \vdash \Pi[x : S_1]T_1 <: \Pi[x : S_2]T_2}{\Gamma \vdash S <: T \quad \Gamma \vdash S \ U : K} \\
\frac{\Gamma \vdash S <: T \quad \Gamma \vdash S \ U : K}{\Gamma \vdash S \ U <: T \ U}
\end{array}$$

Sub-binding Rules

$$\frac{\Gamma \vdash S <: T}{\Gamma \vdash (X <: S) <: (X <: T)}
\qquad
\frac{\Gamma \vdash S <: T}{\Gamma \vdash (X \leftrightarrow S) <: (X <: T)} \\
\frac{\Gamma \vdash S \leftrightarrow T}{\Gamma \vdash (X \leftrightarrow S) <: (X \leftrightarrow T)}
\qquad
\frac{\Gamma \vdash S <: T}{\Gamma \vdash (\mathbf{x} : S) <: (\mathbf{x} : T)}$$

Type Equivalence Rules

$$\frac{\Gamma \vdash T : K}{\Gamma \vdash T \leftrightarrow T}
\qquad
\frac{\Gamma \vdash S \leftrightarrow T}{\Gamma \vdash T \leftrightarrow S} \\
\frac{\Gamma \vdash S \leftrightarrow T \quad \Gamma \vdash T \leftrightarrow U}{\Gamma \vdash S \leftrightarrow U}
\qquad
\frac{X \leftrightarrow T \in \Gamma}{\Gamma \vdash X \leftrightarrow T}$$

$$\begin{array}{c}
 \frac{\frac{\frac{\frac{\vdash \Gamma, X \prec S_1 \text{ ok}}{\Gamma \vdash S_1 \leftrightarrow S_2} \quad \Gamma, X \prec S_1 \vdash T_1 \leftrightarrow T_2}{\Gamma \vdash \forall [X \prec S_1] T_1 \leftrightarrow \forall [X \prec S_2] T_2}}{\Gamma \vdash \lambda [X : K] S \leftrightarrow \lambda [X : K] T}}{\Gamma \vdash (\lambda [X : K_1] S) T : K_2}}{\Gamma \vdash (\lambda [X : K_1] S) T \leftrightarrow \{X \mapsto T\} S} \\
 \frac{\frac{\frac{\frac{\frac{\vdash \Gamma, X \prec S_1 \text{ ok}}{\Gamma \vdash S_1 \leftrightarrow S_2} \quad \Gamma, x : S_1 \vdash T_1 \leftrightarrow T_2}{\Gamma \vdash \Pi [x : S_1] T_1 \leftrightarrow \Pi [x : S_2] T_2}}{\Gamma \vdash S_1 \leftrightarrow S_2} \quad \frac{\frac{\frac{\frac{\vdash \Gamma, x : S_1 \text{ ok}}{\Gamma \vdash S_1 \leftrightarrow S_2} \quad \Gamma \vdash T_1 : K}{\Gamma \vdash T_1 \leftrightarrow T_2}}{\Gamma \vdash S_1 \leftrightarrow S_2} \quad \Gamma \vdash T_1 \leftrightarrow T_2}{\Gamma \vdash S_1 \leftrightarrow S_2} \quad T_2}}{\Gamma \vdash \mathbf{t} : \{\!| \mathbf{l}_i / B_i^{i \in 1 \dots n} \!| \} \quad B_j = X \leftrightarrow T}}{\Gamma \vdash \mathbf{t} . \mathbf{l}_j \leftrightarrow \{BV(B_i) \mapsto \mathbf{t} . \mathbf{l}_i^{i \in 1 \dots j-1}\} T} \\
 \frac{\frac{\frac{\frac{\frac{\vdash \Gamma, B_1, \dots, B_n \text{ ok}}{\Gamma, B_1, \dots, B_{j-1} \vdash B_j \leftrightarrow B'_j} \quad \frac{\vdash \Gamma, B'_1, \dots, B'_n \text{ ok}}{\Gamma, B_1, \dots, B_{j-1} \vdash B_j \leftrightarrow B'_j} \quad j \in 1 \dots n}}{\Gamma \vdash \{\!| \mathbf{l}_i / B_i^{i \in 1 \dots n} \!| \} \leftrightarrow \{\!| \mathbf{l}_i / B'_i^{i \in 1 \dots n} \!| \}}}{\Gamma \vdash \{\!| \mathbf{l}_i / B_i^{i \in 1 \dots n} \!| \} \leftrightarrow \{\!| \mathbf{l}_i / B'_i^{i \in 1 \dots n} \!| \}}
 \end{array}$$

Binding Equivalence Rules

$$\begin{array}{c}
 \frac{\Gamma \vdash S \leftrightarrow T}{\Gamma \vdash (X \prec S) \leftrightarrow (X \prec T)} \\
 \frac{\Gamma \vdash S \leftrightarrow T}{\Gamma \vdash (x : S) \leftrightarrow (x : T)} \\
 \frac{\Gamma \vdash S \leftrightarrow T}{\Gamma \vdash (X \leftrightarrow S) \leftrightarrow (X \leftrightarrow T)}
 \end{array}$$