

Parametric polymorphism

Lecture 17

Tuesday, March 30, 2010

1 Parametric polymorphism

Polymorph means “many forms”. *Polymorphism* is the ability of code to be used on values of different types. For example, a polymorphic function is one that can be invoked with arguments of different types. A polymorphic datatype is one that can contain elements of different types.

Several kinds of polymorphism are commonly used in modern languages.

- *Subtype polymorphism* gives a single term many types using the subsumption rule. For example, a function with argument τ can operate on any value with a type that is a subtype of τ .
- *Ad-hoc polymorphism* usually refers to code that appears to be polymorphic to the programmer, but the actual implementation is not. A typical example is *overloading*: using the same function name for functions with different kinds of parameters. Although it looks like a polymorphic function to the code that uses it, there are actually multiple function implementations (none being polymorphic) and the compiler invokes the appropriate one. Ad-hoc polymorphism is a dispatch mechanism: the type of the arguments is used to determine (either at compile time or run time) which code to invoke.
- *Parametric polymorphism* refers to code that is written without knowledge of the actual type of the arguments; the code is parametric in the type of the parameters. Examples include polymorphic functions in ML, or generics in Java 5.

We consider parametric polymorphism in more detail. Suppose we are working in the simply-typed lambda calculus, and consider a “doubling” function for integers that takes a function f , and an integer x , applies f to x , and then applies f to the result.

$$\text{doubleInt} \triangleq \lambda f:\mathbf{int} \rightarrow \mathbf{int}. \lambda x:\mathbf{int}. f (f x)$$

We could also write a double function for booleans. Or for functions over integers. Or for any other type...

$$\text{doubleBool} \triangleq \lambda f:\mathbf{bool} \rightarrow \mathbf{bool}. \lambda x:\mathbf{bool}. f (f x)$$

$$\text{doubleFn} \triangleq \lambda f:(\mathbf{int} \rightarrow \mathbf{int}) \rightarrow (\mathbf{int} \rightarrow \mathbf{int}). \lambda x:\mathbf{int} \rightarrow \mathbf{int}. f (f x)$$

$$\vdots$$

In the simply typed lambda calculus, if we want to apply the doubling operation to different types of arguments in the same program, we need to write a new function for each type. This violates the *abstraction principle* of software engineering:

Each significant piece of functionality in a program should be implemented in just one place in the source code. When similar functions are carried out by distinct pieces of code, it is generally beneficial to combine them into one by abstracting out the varying parts.

In the doubling functions above, the varying parts are the types. We need a way to abstract out the type of the doubling operation, and later instantiate this abstract type with different concrete types.

We extend the simply-typed lambda calculus with abstraction over types, giving the *polymorphic lambda calculus*, also called *System F*.

A *type abstraction* is a new expression, written $\Lambda X. e$, where Λ is the upper-case form of the Greek letter lambda, and X is a *type variable*. We also introduce a new form of application, called *type application*, or *instantiation*, written $e_1 [\tau]$.

When a type abstraction meets a type application during evaluation, we substitute the free occurrences of the type variable with the type. Note that instantiation does not require the program to keep run-time type information, or to perform type checks at run-time; it is just used as a way to statically check type safety in the presence of polymorphism.

1.1 Syntax and operational semantics

The new syntax of the language is given by the following grammar.

$$\begin{aligned} e &::= n \mid x \mid \lambda x:\tau. e \mid e_1 e_2 \mid \Lambda X. e \mid e [\tau] \\ v &::= n \mid \lambda x:\tau. e \mid \Lambda X. e \end{aligned}$$

The evaluation rules for the polymorphic lambda calculus are the same as for the simply-typed lambda calculus, augmented with new rules for evaluating type application.

$$E ::= [\cdot] \mid E e \mid v E \mid E [\tau]$$

$$\frac{e \rightarrow e'}{E[e] \rightarrow E[e']} \qquad \beta\text{-REDUCTION} \frac{}{(\lambda x:\tau. e) v \rightarrow e\{v/x\}}$$

$$\text{TYPE-REDUCTION} \frac{}{(\Lambda X. e) [\tau] \rightarrow e\{\tau/X\}}$$

Let's consider an example. In this language, the polymorphic identity function is written as

$$ID \triangleq \Lambda X. \lambda x:X. x$$

We can apply the polymorphic identity function to **int**, producing the identity function on integers.

$$(\Lambda X. \lambda x:X. x) [\mathbf{int}] \rightarrow \lambda x:\mathbf{int}. x$$

We can apply *ID* to other types as easily:

$$(\Lambda X. \lambda x:X. x) [\mathbf{int} \rightarrow \mathbf{int}] \rightarrow \lambda x:\mathbf{int} \rightarrow \mathbf{int}. x$$

1.2 Type system

We also need to provide a type for the new type abstraction. The type of $\Lambda X. e$ is $\forall X. \tau$, where τ is the type of e , and may contain the type variable X . Intuitively, we use this notation because we can instantiate the type expression with any type for X : for any type X , expression e can have the type τ (which may mention X).

$$\tau ::= \mathbf{int} \mid \tau_1 \rightarrow \tau_2 \mid X \mid \forall X. \tau$$

Type checking expressions is slightly different than before. Besides the type environment Γ (which maps variables to types), we also need to keep track of the set of type variables Δ . This is to ensure that a type variable X is only used in the scope of an enclosing type abstraction $\Lambda X. e$. Thus, typing judgments are now of the form $\Delta, \Gamma \vdash e:\tau$, where Δ is a set of type variables, and Γ is a typing context. We also use an additional judgment $\Delta \vdash \tau \text{ ok}$ to ensure that type τ uses only type variables from the set Δ .

$$\frac{}{\Delta, \Gamma \vdash n:\mathbf{int}} \qquad \frac{\Delta \vdash \tau \text{ ok}}{\Delta, \Gamma \vdash x:\tau} \Gamma(x) = \tau \qquad \frac{\Delta, \Gamma, x:\tau \vdash e:\tau' \quad \Delta \vdash \tau \text{ ok}}{\Delta, \Gamma \vdash \lambda x:\tau. e:\tau \rightarrow \tau'}$$

$$\begin{array}{c}
\frac{\Delta, \Gamma \vdash e_1 : \tau \rightarrow \tau' \quad \Delta, \Gamma \vdash e_2 : \tau}{\Delta, \Gamma \vdash e_1 e_2 : \tau'} \quad \frac{\Delta \cup \{X\}, \Gamma \vdash e : \tau}{\Delta, \Gamma \vdash \Lambda X. e : \forall X. \tau} \quad \frac{\Delta, \Gamma \vdash e : \forall X. \tau' \quad \Delta \vdash \tau \text{ ok}}{\Delta, \Gamma \vdash e [\tau] : \tau' \{ \tau / X \}} \\
\\
\frac{}{\Delta \vdash X \text{ ok}} X \in \Delta \quad \frac{}{\Delta \vdash \mathbf{int} \text{ ok}} \quad \frac{\Delta \vdash \tau_1 \text{ ok} \quad \Delta \vdash \tau_2 \text{ ok}}{\Delta \vdash \tau_1 \rightarrow \tau_2 \text{ ok}} \quad \frac{\Delta \vdash \tau_1 \text{ ok} \quad \Delta \cup \{X\} \vdash \tau \text{ ok}}{\Delta \vdash \forall X. \tau \text{ ok}}
\end{array}$$

1.3 Examples

Let's consider the doubling operation again. We can write a polymorphic doubling operation as

$$\mathit{double} \triangleq \Lambda X. \lambda f : X \rightarrow X. \lambda x : X. f (f x).$$

The type of this expression is

$$\forall X. (X \rightarrow X) \rightarrow X \rightarrow X$$

We can instantiate this on a type, and provide arguments. For example,

$$\mathit{double} \ [\mathbf{int}] \ (\lambda n : \mathbf{int}. n + 1) \ 7 \longrightarrow (\lambda f : \mathbf{int} \rightarrow \mathbf{int}. \lambda x : \mathbf{int}. f (f x)) \ (\lambda n : \mathbf{int}. n + 1) \ 7 \longrightarrow^* 9$$

Recall that in the simply-typed lambda calculus, we had no way of typing the expression $\lambda x. x x$. In the polymorphic lambda calculus, however, we can type this expression if we give it a polymorphic type and instantiate it appropriately.

$$\vdash \lambda x : \forall X. X \rightarrow X. x \ [\forall X. X \rightarrow X] \ x \quad : \quad (\forall X. X \rightarrow X) \rightarrow (\forall X. X \rightarrow X)$$

1.4 Polymorphism in ML and Java

In real languages such as ML, programmers don't have to annotate their programs with $\forall X. \tau$ or $e [\tau]$. Both are automatically inferred by the compiler (although the programmer can specify the former if he wishes). For example, we can write `let double f x = f (f x) ; ;` and Ocaml will figure out that the type is $('a \rightarrow 'a) \rightarrow 'a \rightarrow 'a$ (which is roughly equivalent to $\forall A. (A \rightarrow A) \rightarrow A \rightarrow A$).

We can also write `double (fun x -> x+1) 7 ; ;`, and Ocaml will infer that the polymorphic function `double` is instantiated on the type `int`.

The polymorphism in ML is not, however, exactly like the polymorphism in System F. ML restricts what types a type variable may be instantiated with. Specifically, type variables can not be instantiated with polymorphic types. Also, polymorphic types are not allowed to appear on the left-hand side of arrows (i.e., a polymorphic type cannot be the type of a function argument). This form of polymorphism is known as *let-polymorphism* (due to the special role played by `let` in ML), or *prenex polymorphism*. These restrictions ensure that *type inference* is possible.

An example of a term that is typable in System F but not typable in ML is the self-application expression $\lambda x. x x$. (Try typing `fun x -> x x ; ;` in the top-level loop of Ocaml, and see what happens...)

Java, as of version 1.5, provides *generics*, a form of parametric polymorphism. For instance, we can write a class that is parameterized on an unknown reference type `T`:

```

class Pair<T> {
  T x, y;

  Pair(T x, T y) {
    this.x = x;
    this.y = y;
  }

  T fst(Pair<T> p) {
    this.x = p.x;
    return p.x;
  }
}

```

This is a class that contains a pair of two elements of some unknown type T . The parameterization $\forall T$ is implicit around the class declaration. Since Java does not support type inference, type instantiations are required. Type instantiations are done by writing the actual type in angle brackets:

```

Pair<Boolean> p;
p = new Pair<Boolean>(Boolean.TRUE, Boolean.FALSE);
Boolean x = p.fst(p);

```

1.5 Bounded polymorphism

Polymorphism and subtyping can be combined. That is, we can change type abstraction to include an upper bound on the types that can be used to instantiate the type variable: $\forall X \leq \tau_1. \tau_2$. This means that X can only be instantiated with a subtype of τ_1 .

Note that if there is a “top type” \top (i.e., every type τ is a subtype of \top), then $\forall X. \tau$ is equivalent to $\forall X \leq \top. \tau$.

Subtyping and parametric polymorphism are largely orthogonal. We need to change the type variable context Δ so that it contains the bounds on type variables, so Δ is now a sequence of elements of the form $X \leq \tau$.

The syntax of expressions, values and types is now as follows.

$$\begin{aligned}
 e &::= n \mid x \mid \lambda x:\tau. e \mid e_1 e_2 \mid \Lambda X \leq \tau. e \mid e[\tau] \\
 v &::= n \mid \lambda x:\tau. e \mid \Lambda X \leq \tau. e \\
 \tau &::= \mathbf{int} \mid \tau_1 \rightarrow \tau_2 \mid X \mid \forall X. \tau
 \end{aligned}$$

The operational semantics are the same (modulo the minor changes to syntax). Most of the typing rules remain the same. We present the rules for type abstraction, type application, and subsumption. Note that type application requires that the instantiating type is a subtype of the declared bound on the type variable. Note also that the subtyping relation now uses the type variable context Δ .

$$\frac{\Delta, X \leq \tau_1, \Gamma \vdash e:\tau_2}{\Delta, \Gamma \vdash \Lambda X \leq \tau_1. e:\forall X \leq \tau_1. \tau_2} \quad \frac{\Delta, \Gamma \vdash e:\forall X \leq \tau_1. \tau_2 \quad \Delta \vdash \tau \leq \tau_1}{\Delta, \Gamma \vdash e[\tau]:\tau_2\{\tau/X\}} \quad \frac{\Delta, \Gamma \vdash e:\tau \quad \Delta \vdash \tau \leq \tau'}{\Delta, \Gamma \vdash e:\tau'}$$

The rule for subtyping a type variable simply uses the type variable context Δ .

$$\frac{}{\Delta \vdash X \leq \tau} X \leq \tau \in \Delta$$

The subtyping rule for polymorphic types is the most interesting. Going back to the subtyping principle, $\forall X \leq \tau_1. \tau_2$ is a subtype of $\forall X \leq \tau'_1. \tau'_2$ if whenever a value of type $\forall X \leq \tau'_1. \tau'_2$ is expected, a value of type $\forall X \leq \tau_1. \tau_2$ can be used instead. Intuitively, an expression of polymorphic type $\forall X \leq \tau'_1. \tau'_2$ can be thought of as a function from types to terms. That is, it takes a type as an argument, and returns an expression.

(By contrast, functions (also known as term abstractions) that we have seen previously are functions from terms to terms.)

A term of type $\forall X \leq \tau_1'. \tau_2'$ may be instantiated with any type that is a subtype of τ_1' ; for $\forall X \leq \tau_1'. \tau_2'$ to be a subtype, it must also be able to be instantiated with any type that is a subtype of τ_1' , so we want τ_1' to be a subtype of τ_1 . That is, we are contravariant in the bound on the type variable. Also, when instantiated on any type X that satisfies both constraints (i.e., $X \leq \tau_1$ and $X \leq \tau_1'$), a value produced by an expression of type τ_2 should be usable at type τ_2' . So we want $\Delta, X \leq \tau_1' \vdash \tau_2 \leq \tau_2'$. The inference rule for subtyping on polymorphic types is thus the following.

$$\frac{\Delta \vdash \tau_1' \leq \tau_1 \quad \Delta, X \leq \tau_1' \vdash \tau_2 \leq \tau_2'}{\Delta \vdash \forall X \leq \tau_1. \tau_2 \leq \forall X \leq \tau_1'. \tau_2'}$$