CS 4110 – Programming Languages and Logics
Lecture #24: Polymorphism

1 Parametric polymorphism

Polymorphism (Greek for "many forms") is the ability for code to be used with 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.

There are several different kinds of polymorphism that are commonly used in modern programming languages.

- *Subtype polymorphism* allows a term to have many types using the subsumption rule, which allows a value of type τ to masquerade as a value of type τ' provided that τ 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. For example, languages with *overloading* allow the same function name to be used with functions that take different types 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 and Java generics.

In this lecture we will consider parametric polymorphism in detail. As a motivating example, suppose we are working in the simply-typed λ-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.

$$\mathsf{doubleInt} \triangleq \lambda f : \mathbf{int} \to \mathbf{int}.\, \lambda x : \mathbf{int}.\, f\ (f\ x)$$

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

$$\mathsf{doubleBool} \triangleq \lambda f : \mathbf{bool} \to \mathbf{bool}.\, \lambda x : \mathbf{bool}.\, f\ (f\ x)$$
$$\mathsf{doubleFn} \triangleq \lambda f : (\mathbf{int} \to \mathbf{int}) \to (\mathbf{int} \to \mathbf{int}).\, \lambda x : \mathbf{int} \to \mathbf{int}.\, f\ (f\ x)$$
$$\vdots$$

In the simply-typed λ-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 a fundamental principle of software engineering:

Abstraction Principle: *Every major piece of functionality in a program should be implemented in just one place in the code. When similar functionality is provided by distinct pieces of code, the two should be combined 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 it with different concrete types.

1.1 Polymorphic λ-calculus

We can extend the simply-typed λ-calculus with abstraction over types. The resulting system is known by two names: *polymorphic λ-calculus* and *System F*.

A *type abstraction* is a new expression, written $\Lambda \alpha.\, e$, where Λ is the upper-case form of the Greek letter lambda, and α is a *type variable*. A *type application*, written $e_1\, [\tau]$, *instantiates* a type application at a particular type.

When a type abstraction meets a type application during evaluation, we substitute the free occurrences of the type variable with the type. Importantly, 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.

1.2 Syntax and operational semantics

The syntax of the polymorphic λ-calculus is given by the following grammar.

$$e ::= n \mid x \mid \lambda x : \tau.\, e \mid e_1\, e_2 \mid \Lambda \alpha.\, e \mid e\, [\tau]$$
$$v ::= n \mid \lambda x : \tau.\, e \mid \Lambda \alpha.\, e$$

The evaluation rules are the same as for the simply-typed λ-calculus, as well as two new rules for evaluating type abstractions and applications.

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

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

$$\text{TYPE-REDUCTION } \frac{}{(\Lambda \alpha.\, e)\, [\tau] \to e\{\tau/\alpha\}}$$

To illustrate, consider a simple example. In this language, the polymorphic identity function is written as

$$\mathsf{ID} \triangleq \Lambda \alpha.\, \lambda x : \alpha.\, x$$

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

$$(\Lambda \alpha.\, \lambda x : \alpha.\, x)\, [\textbf{int}] \to \lambda x : \textbf{int}.\, x$$

We can apply ID to other types as well:

$$(\Lambda \alpha.\, \lambda x : \alpha.\, x)\, [\textbf{int} \to \textbf{int}] \to \lambda x : \textbf{int} \to \textbf{int}.\, x$$

1.3 Type system

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

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 α is only used in the scope of an enclosing type abstraction $\Lambda\alpha.\,e$. 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 $\Delta \vdash \tau$ ok to ensure that type τ uses only type variables from the set Δ.

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

$$\frac{\Delta, \Gamma \vdash e_1 : \tau \to \tau' \quad \Delta, \Gamma \vdash e_2 : \tau}{\Delta, \Gamma \vdash e_1\,e_2 : \tau'} \qquad \frac{\Delta \cup \{\alpha\}, \Gamma \vdash e : \tau}{\Delta, \Gamma \vdash \Lambda\alpha.\,e : \forall\alpha.\,\tau} \qquad \frac{\Delta, \Gamma \vdash e : \forall\alpha.\,\tau' \quad \Delta \vdash \tau \text{ ok}}{\Delta, \Gamma \vdash e\,[\tau] : \tau'\{\tau/\alpha\}}$$

$$\frac{}{\Delta \vdash \alpha \text{ ok}}\,\alpha \in \Delta \qquad \frac{}{\Delta \vdash \mathbf{int} \text{ ok}} \qquad \frac{\Delta \vdash \tau_1 \text{ ok} \quad \Delta \vdash \tau_2 \text{ ok}}{\Delta \vdash \tau_1 \to \tau_2 \text{ ok}} \qquad \frac{\Delta \cup \{\alpha\} \vdash \tau \text{ ok}}{\Delta \vdash \forall\alpha.\,\tau \text{ ok}}$$

2 Programming in Polymorphic λ-Calculus

Now we consider a number of examples of programming in the polymorphic λ-calculus.

2.1 Doubling

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

$$\text{double} \triangleq \Lambda\alpha.\,\lambda f : \alpha \to \alpha.\,\lambda x : \alpha.\,f\,(f\,x).$$

The type of this expression is

$$\forall\alpha.\,(\alpha \to \alpha) \to \alpha \to \alpha$$

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

$$\text{double}\,[\mathbf{int}]\,(\lambda n : \mathbf{int}.\,n + 1)\,7 \to (\lambda f : \mathbf{int} \to \mathbf{int}.\,\lambda x : \mathbf{int}.\,f\,(f\,x))\,(\lambda n : \mathbf{int}.\,n + 1)\,7$$
$$\to^* 9$$

2.2 Self Application

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

$$\vdash \quad \lambda x : \forall \alpha.\, \alpha \to \alpha.\, x\, [\forall \alpha.\, \alpha \to \alpha]\, x \quad : \quad (\forall \alpha.\, \alpha \to \alpha) \to (\forall \alpha.\, \alpha \to \alpha)$$

2.3 Sums and Products

We can encode sums and products in polymorphic λ-calculus without adding any additional types! The encodings are based on the Church encodings from untyped λ-calculus.

$$
\begin{aligned}
\tau_1 \times \tau_2 \quad &\triangleq \quad \forall R.\, (\tau_1 \to \tau_2 \to R) \to R \\
(\cdot, \cdot) \quad &\triangleq \quad \Lambda T_1.\, \Lambda T_2.\, \lambda v_1 : T_1,\, \lambda v_2 : T_2.\, \Lambda R.\, \lambda p : (T_1 \to T_2 \to R).\, p\, v_1\, v\, 2 \\
\#1 \quad &\triangleq \quad \Lambda T_1.\, \Lambda T_2.\, \lambda v : T_1 \times T_2.\, v\, [T_1]\, (\lambda x : T_1.\, \lambda y : T_2.\, x) \\
\#2 \quad &\triangleq \quad \Lambda T_1.\, \Lambda T_2.\, \lambda v : T_1 \times T_2.\, v\, [T_2]\, (\lambda x : T_1.\, \lambda y : T_2.\, y) \\
\\
\mathsf{unit} \quad &\triangleq \quad \forall R.\, R \to R \\
() \quad &\triangleq \quad \Lambda R.\, \lambda x : R.\, x \\
\\
\tau_1 + \tau_2 \quad &\triangleq \quad \forall R.(\tau_1 \to R) \to (\tau_2 \to R) \to R \\
\mathsf{inl} \quad &\triangleq \quad \Lambda T_1.\, \Lambda T_2.\, \lambda v_1 : T_1.\, \Lambda R.\, \lambda b_1 : T_1 \to R.\, \lambda b_2 : T_2 \to R.\, b_1\, v_1 \\
\mathsf{inr} \quad &\triangleq \quad \Lambda T_1.\, \Lambda T_2.\, \lambda v_2 : T_2.\, \Lambda R.\, \lambda b_1 : T_1 \to R.\, \lambda b_2 : T_2 \to R.\, b_2\, v_2 \\
\mathsf{case} \quad &\triangleq \quad \Lambda T_1.\, \Lambda T_2.\, \Lambda R.\, \lambda v : T_1 + T_2.\, \lambda b_1 : T_1 \to R.\, \lambda b_2 : T_2 \to R.\, v\, [R]\, b_1\, b_2 \\
\\
\mathsf{void} \quad &\triangleq \quad \forall R.\, R
\end{aligned}
$$

3 Type Erasure

The semantics presented above explicitly passes type. In an implementation, one often wants to eliminate types for efficiency. The following translation "erases" the types from a polymorphic λ-calculus expression.

$$
\begin{aligned}
erase(x) \quad &= \quad x \\
erase(\lambda x : \tau.\, e) \quad &= \quad \lambda x.\, erase(e) \\
erase(e_1\, e_2) \quad &= \quad erase(e_1)\, erase(e_2) \\
erase(\Lambda \alpha.\, e) \quad &= \quad \lambda z.\, erase(e) \qquad \text{where } z \text{ is fresh for } e \\
erase(e\, [\tau]) \quad &= \quad erase(e)\, (\lambda x.\, x)
\end{aligned}
$$

The following theorem states that the translation is adequate.

Theorem (Adequacy). *For all expressions e and e', we have $e \to e'$ iff $erase(e) \to erase(e')$.*

4 Type Inference

The type reconstruction problem asks whether, for a given untyped λ-calculus expression e' there exists a well-typed System F expression e such that $erase(e) = e'$. It was shown to be undecidable

by Wells in 1994. See Chapter 23 of Pierce for further discussion, as well as restrictions for which type reconstruction is decidable.