Last time we saw how to infer types for STLC terms, and we noticed that Robinson's unification algorithm would return the most general type, which sometimes had some type variables in it. For example, if we tried to infer the type for $\mathsf{id} = \lambda x.\, x$, we would end up with $\alpha \to \alpha$ for the free type variable α. Using STLC with type inference, that type means that we can type id at $\tau \to \tau$ for any type τ. Intuitively, that seems to mean that we should be able to apply id to a value of any type, and we can, but *only if we only ever apply it to values of one type*.

As an example, consider the following program.

$$\begin{array}{l}\mathsf{let}\ f = \lambda x.\, x\ \mathsf{in} \\ \mathsf{if}\ (f\ \mathsf{true})\ \mathsf{then}\ (f\ 3)\ \mathsf{else}\ (f\ 4)\end{array}$$

The type inference algorithm will correctly infer the type $\alpha \to \alpha$ for f, but then as soon as it sees the application $f\ \mathsf{true}$ in the condition, it will unify $\alpha = \mathsf{bool}$. Then when it sees the applications in the branches, it will attempt to unify $\mathsf{bool} = \mathsf{int}$, at which point it will fail; those are disparate types that cannot be unified.

To solve this problem, we need to somehow capture the fact that f doesn't have type $\tau \to \tau$ for an unspecified by specific τ, but rather f needs to be able to have a type that can behave like $\mathsf{bool} \to \mathsf{bool}$ in some places and $\mathsf{int} \to \mathsf{int}$ in others. That is, we want the type to support *polymorphism*.[1]

1 Polymorphic λ-Calculus

To add polymorphism to our types and support the above example, we make two changes to our grammar of types that introduce concepts we have not seen before. The first is *type variables*, which we will denote α (and sometimes β). Like program variables, these type variables are stand-ins for something we do not yet have, just now that something is a type, not a program term.

The second addition is a new constructor that *universally quantifies over types*, denoted $\forall\alpha.\, \tau$. This universal quantifier serves as a binder for type variables. In the type $\forall\alpha.\, \tau$, the variable α is bound in τ. These universal quantifiers function almost exactly as λ-abstractions, but they work at the level of types; $\forall\alpha.\, \tau$ expects to be provided with a type τ', which it will substitute for free occurrences of the bound variable α in τ.

These variables and binders have the same notions of scope, free and bound variables, renaming, and safe substitution as variables in λ-calculus. Together, these change the type grammar as follows.

$$\tau \quad ::= \quad \cdots \mid \alpha \mid \forall\alpha.\, \tau.$$

The type $\forall\alpha.\, \tau$ is a *polymorphic type*, or a *type schema*, which is a pattern with type variables that can be instantiated to obtain actual types. For example, the polymorphic type of id would be

$$\mathsf{id} : \forall\alpha.\, \alpha \to \alpha.$$

The language that results from adding these types is known as the *polymorphic λ-calculus*. It has the same terms and evaluation rules as STLC, but with these extra polymorphic types. All terms that were well-typed before will still be well-typed, but now more terms will be typable as well.

[1] Greek for "many forms"

1.1 Typing Rules

In addition to the old typing rules, polymorphic λ-calculus adds two new rules, called the *generalization* and *instantiation* rules, that introduce and eliminate these quantifiers, respectively. The full type system is then as follows.

$$[\textsc{Unit}]\ \frac{}{\Gamma \vdash () : \mathsf{unit}} \qquad [\textsc{Var}]\ \frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau} \qquad [\textsc{Abs}]\ \frac{\Gamma, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \lambda x.\, e : \tau_1 \to \tau_2} \qquad [\textsc{App}]\ \frac{\Gamma \vdash e_1 : \tau_1 \to \tau_2 \quad \Gamma \vdash e_2 : \tau_1}{\Gamma \vdash e_1\ e_2 : \tau_2}$$

$$[\textsc{Generalize}]\ \frac{\Gamma \vdash e : \tau \qquad \alpha \notin \mathrm{FV}(\Gamma)}{\Gamma \vdash e : \forall \alpha.\, \tau} \qquad [\textsc{Instantiate}]\ \frac{\Gamma \vdash e : \forall \alpha.\, \tau}{\Gamma \vdash e : \tau[\alpha \mapsto \tau']}$$

One notable change here is that the types themselves are *open*—that is, they may contain free variables. The Generalize rule—which is only interesting when $\alpha \in \mathrm{FV}(\tau)$—says that, if $\Gamma \vdash e : \tau$ holds and there are no assumptions involving α, then we can safely put in any type for α and the typing proof will still work. The constraint that $\alpha \notin \mathrm{FV}(\Gamma)$ says that no *assumptions* reference α, which is exactly what we need to ensure it is unconstrained. We can then conclude that e has type τ for all values of α, or, succinctly, $\Gamma \vdash e : \forall \alpha.\, \tau$. The Instantiate rule is the other side of this concept. It says that, if e has a quantified type, then anything we put in for that type variable will work.

1.1.1 Examples

To see how this works, we can look at a couple of proof trees.

First, we see how to show that $\vdash \lambda x.\, x : \forall \alpha.\, \alpha \to \alpha$.

$$[\textsc{Generalize}]\ \frac{[\textsc{Abs}]\ \dfrac{[\textsc{Var}]\ \dfrac{}{x : \alpha \vdash x : \alpha}}{\vdash \lambda x.\, x : \alpha \to \alpha}}{\vdash \lambda x.\, x : \forall \alpha.\, \alpha \to \alpha}$$

Notably, some terms are typable that were not in Stlc. For example, $\omega = \lambda x.\, x\ x$ can now have a type!

$$[\textsc{Generalize}]\ \frac{[\textsc{Abs}]\ \dfrac{[\textsc{App}]\ \dfrac{[\textsc{Instantiate}]\ \dfrac{[\textsc{Var}]\ \dfrac{}{x : (\forall \alpha.\, \alpha) \vdash x : \forall \alpha.\, \alpha}}{x : (\forall \alpha.\, \alpha) \vdash x : \gamma \to \beta} \qquad \dfrac{[\textsc{Instantiate}]\ \dfrac{[\textsc{Var}]\ \dfrac{}{x : (\forall \alpha.\, \alpha) \vdash x : \forall \alpha.\, \alpha}}{x : (\forall \alpha.\, \alpha) \vdash x : \gamma}}{x : (\forall \alpha.\, \alpha) \vdash x\ x : \beta}}{\vdash \lambda x.\, x\ x : (\forall \alpha.\, \alpha) \to \beta}}{\vdash \lambda x.\, x\ x : \forall \beta.\, (\forall \alpha.\, \alpha) \to \beta}$$

Unfortunately, this type is not particularly meaningful because, without adding extra primitives to the language, *nothing* has the type $\forall \alpha.\, \alpha$. It is said to be *uninhabited*, and we can give it the name void. By a nearly identical argument, however, we could also have proven that ω has type $\forall \beta.\, (\forall \alpha.\, \alpha \to \alpha) \to \beta \to \beta$, which is a meaningful type.

Interestingly, while ω is typable, polymorphism is not enough by itself to allow us to type $\Omega = \omega\ \omega$. Indeed, polymorphic λ-calculus on its own (without something more powerful like recursive functions) still has the same property of Stlc that every well-typed program terminates. The proof of that fact, however, is very complicated and is beyond the scope of this course.

2 Using Polymorphism

While polymorphic λ-calculus allows for reuse of programs in very convenient ways, it has a major downside: type inference is now undecidable. A compiler could try its best and ask programmers to insert type annotations when it fails, but we can also restrict the use of polymorphism to regain decidability.

2.1 Let-Polymorphism

A simple approach taken by languages like Haskell and OCaml is two restrict polymorphism in two ways.

1. Type quantification can only appear at the top level of a type. That is, we only allow polymorphic expressions of the form $\forall\alpha_1, \ldots, \forall\alpha_n.\, \tau$ where τ is quantifier-free.
2. Polymorphism can only be introduced in the context of a let expression, not arbitrary variable bindings.

Together, these restrictions result in the following modifications to the language.

$$\begin{array}{rrcl} \text{Quantifier-Free Terms} & \tau & ::= & \mathsf{unit} \mid \alpha \mid \tau_1 \to \tau_2 \\ \text{Primitive let} & e & ::= & \cdots \mid \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 \end{array}$$

$$[\textsc{PolyLet}]\ \frac{\begin{array}{c}\Gamma \vdash e_1 : \tau_1 \qquad \Gamma, x{:}\forall\alpha_1, \ldots, \forall\alpha_n.\, \tau_1 \vdash e_2 : \tau_2 \\ \{\alpha_1, \ldots, \alpha_n\} = \mathrm{FV}(\tau_1) - \mathrm{FV}(\Gamma)\end{array}}{\Gamma \vdash \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 : \tau_2}$$

Note that we retain the Instantiate rule from Section 1.1, but we remove that Generalize rule.

We previously considered the term $\mathsf{let}\ x = e_1\ \mathsf{in}\ e_2$ to be syntactic sugar for $(\lambda x.\, e_2)\ e_1$. While the two still have the same semantic rules, they are no longer equivalent in the type system. Some terms are well-typed using let, but not using function application, including the example from the beginning of this lecture.

The fact that polymorphism can only be introduced with let expressions is why this is known as *let-polymorphism*. Both Haskell and OCaml use let-polymorphism. In theory, this could cause the type checker to require exponential time, but in practice it is not a problem.

2.2 System F

When we first introduced Stlc, we used Church-style terms with types explicitly annotated on function arguments. The corresponding version of polymorphic λ-calculus is called *System F*. In System F, we explicitly abstract terms with respect to types and explicitly instantiate those types before using the term. We therefore augment the syntax of Stlc with new types and terms as follows.

$$\begin{array}{rcl} \tau & ::= & \cdots \mid \alpha \mid \forall\alpha.\, \tau \\ e & ::= & \cdots \mid \Lambda\alpha.\, e \mid e\ \tau \end{array}$$

These new terms are known as *type abstraction* and *type application*, respectively. Operationally, we can add the following semantic rule

$$\frac{}{(\Lambda\alpha.\, e)\ \tau \longrightarrow e[\alpha \mapsto \tau]}.$$

This just gives a rule for instantiating a polymorphic type. Since these reductions only involve types, they can be performed at compile time.

Unlike in polymorphic λ-calculus, in System F we only want to introduce new type variables if they are going to be bound by a Λ-expression. To check this, the type system needs to keep track of which type variables are bound, and make sure that any type that appears in a type application is *well-formed*, meaning all

of its free variables have been bound in the surrounding context. To do this, we introduce a new *type variable context* Δ to keep track of these variables, in addition to the existing typing (regular variable) context Γ. In more complicated type systems, Δ would be a partial map from variables to *kinds*, that tell us the space a type lives in (kinds are to types as types are to terms), so it is sometimes called a *kinding context*. Right now we only have one kind, so we will simply consider Δ to be a set.

The type system then has two classes of judgements:

$$\Delta \vdash \tau \qquad\qquad \Delta; \Gamma \vdash e : \tau$$

The first says that τ is well-formed in type variable context Δ, while the second says that we can prove e has type τ in type variable context Δ and (regular) variable context Γ. The rules for the well-formed type judgment are as follows.

$$\frac{}{\Delta \vdash \mathsf{unit}} \qquad \frac{\alpha \in \Delta}{\Delta \vdash \alpha} \qquad \frac{\Delta \vdash \tau_1 \qquad \Delta \vdash \tau_2}{\Delta \vdash \tau_1 \to \tau_2} \qquad \frac{\Delta, \alpha \vdash \tau}{\Delta \vdash \forall \alpha.\, \tau}$$

Right now, since these rules do nothing more than keep track of the type variables bound in the surrounding context, one can show that $\Delta \vdash \tau$ if and only if $\mathrm{FV}(\tau) \subseteq \Delta$.

The typing rules for terms are as follows, with differences from the Curry-style polymorphic λ-calculus defined above highlighted in red.

$$\frac{}{\Delta; \Gamma \vdash () : \mathsf{unit}} \qquad \frac{\Gamma(x) = \tau \qquad \Delta \vdash \tau}{\Delta; \Gamma \vdash x : \tau} \qquad \frac{\Delta; \Gamma, x : \tau_1 \vdash e : \tau_2 \qquad \Delta \vdash \tau_1}{\Delta; \Gamma \vdash \lambda x : \tau_1.\, e : \tau_1 \to \tau_2} \qquad \frac{\Delta; \Gamma \vdash e_1 : \tau_1 \to \tau_2 \qquad \Delta; \Gamma \vdash e_2 : \tau_1}{\Delta; \Gamma \vdash e_1\ e_2 : \tau_2}$$

$$\frac{\Delta, \alpha; \Gamma \vdash e : \tau \qquad \alpha \notin \mathrm{FV}(\Gamma)}{\Delta; \Gamma \vdash \Lambda \alpha.\, e : \forall \alpha.\, \tau} \qquad \frac{\Delta; \Gamma \vdash e : \forall \alpha.\, \tau \qquad \Delta \vdash \tau'}{\Delta; \Gamma \vdash e\ \tau' : \tau[\alpha \mapsto \tau']}$$

The first four rules are the same as the rules for STLC, but with Δ included and a requirement that any types that appear be well-formed. The last two rules specify the types for type abstractions and applications, also requiring that types be well-formed. In fact, one can show that if $\Delta; \Gamma \vdash e : \tau$, then τ and all type annotations occurring in e must be well-formed in Δ. In particular $\vdash e : \tau$ is only possible when e and τ are both closed.

To see how to use this, we can look at the polymorphic identity function: $\Lambda \alpha.\, \lambda x : \alpha.\, x$ which has type $\forall \alpha.\, \alpha \to \alpha$ according to the following proof.

$$\dfrac{\dfrac{\dfrac{\alpha \vdash \alpha}{\alpha; x : \alpha \vdash x : \alpha} \qquad \alpha \vdash \alpha}{\alpha; \cdot \vdash \lambda x : \alpha.\, x : \alpha \to \alpha}}{\vdash \Lambda \alpha.\, \lambda x : \alpha.\, x : \forall \alpha.\, \alpha \to \alpha}$$

To apply this function, we must specify what type we are applying it to. So, the correct version of our original example written in System F (with booleans and integers) would be:

$$\begin{array}{l} \mathsf{let}\ f = \Lambda \alpha.\, \lambda x : \alpha.\, x \\ \mathsf{in\ if}\ (f\ \mathsf{bool\ true}) \\ \quad \mathsf{then}\ (f\ \mathsf{int}\ 3) \\ \quad \mathsf{else}\ (f\ \mathsf{int}\ 4) \end{array}$$