Day 19.

1. Types and Type Schemes

Intuitively, we want to think about type schemes as denoting families of (or sets of) types. The type scheme $\forall \alpha.\alpha \to \alpha$ denotes the set $\{\texttt{Int} \to \texttt{Int}, \texttt{Bool} \to \texttt{Bool}, (\texttt{Int} \to \texttt{Int}) \to (\texttt{Int} \to \texttt{Int}), \dots\}$, whereas the scheme $\forall alpha.\alpha$ denotes the set of all types. We can make this idea formal, as follows. We call the set of types denoted by a scheme S the *instances* of S, $\lfloor S \rfloor$, defined as follows:

$$\lfloor T \rfloor = \{T\} \qquad \lfloor \forall \alpha.S \rfloor = \bigcup_{T \in \mathcal{Y}} \lfloor S[T/\alpha] \rfloor$$

So, for example:

$$\lfloor \forall \alpha.\alpha \rfloor = \mathcal{Y}$$
$$\lfloor \alpha \rfloor = \alpha$$
$$\lfloor \forall \alpha.\forall \beta.\alpha \to \beta \rfloor = \{t_1 \to t_2 \mid t_1 \in \mathcal{Y}, t_2 \in \mathcal{Y}\}$$

This idea of instances aligns perfectly with our idea of renaming-equivalence:

Theorem 19.1. $S_1 \equiv_\alpha S_2$ *iff* $\lfloor S_1 \rfloor = \lfloor S_2 \rfloor$.

We can also use instances to order type schemes: we say a type scheme S_1 is *more general* than S_2, $S_1 \geq S_2$, iff $\lfloor S_1 \rfloor \supseteq \lfloor S_2 \rfloor$. For example,

$$\forall \alpha.\alpha \geq \forall \alpha.\alpha \to \alpha \geq \forall \alpha.\forall \beta.(\alpha, \beta) \to (\alpha, \beta) \geq$$
$$\forall \alpha.(\alpha, \texttt{Bool}) \to (\alpha, \texttt{Bool}) \geq (\texttt{Int}, \texttt{Bool}) \to (\texttt{Int}, \texttt{Bool}).$$

We can relate generality to substitution, although with a bit of hair.

Theorem 19.2. $\forall \overline{\alpha}. T_1 \geq \forall \overline{\beta}. T_2$ *iff there are types* $U_1 \dots U_n$ *such that* $T_1[\overline{T}/\overline{\alpha}] = T_2$.

2. Polymorphism in Terms

Let's start with our standard functional term language, and add polymorphism.

$$\mathcal{T} \ni t ::= z \mid t \odot t \mid x \mid \lambda x.t \mid t\, t$$

- We don't get any new term forms for polymorphism—the idea is that $\lambda a.a$ should *implicitly* be polymorphic.
- This is following the *Curry-style* presentation that we've used almost all of the semester. Of course, there are also presentations of polymorphism, which makes polymorphism *explicit*; the most common of these is the *Girard-Reynolds* calculus, also called System F.

We have all the typing rules we're familiar with, and two new ones to introduce and eliminate polymorphism.

$$\frac{}{\Gamma \vdash z : \texttt{Int}} \qquad \frac{\Gamma \vdash t_1 : \texttt{Int} \quad \Gamma \vdash t_2 : \texttt{Int}}{\Gamma \vdash t_1 \odot t_2}$$

$$\frac{}{\Gamma \vdash x : \Gamma(x)} \qquad \frac{\Gamma[x \mapsto T_1] \vdash t : T_2}{\Gamma \vdash \lambda x.t : T_1 \to T_2} \qquad \frac{\Gamma \vdash t_1 : T_1 \to T_2 \quad \Gamma \vdash t_2 : T_1}{\Gamma \vdash t_1\, t_2 : T_2}$$

$$\frac{\Gamma \vdash t : \forall \alpha.S}{\Gamma \vdash t : S[T/\alpha]} \qquad \frac{\Gamma \vdash t : S}{\Gamma \vdash t : \forall \alpha.S}$$

- Intuition: type variables stand in for arbitrary types, so if something produces a result of arbitrary type, it can produce a result of any type we like.
- Quantifiers here are making that assumption *explicit* in types, even though it remains *implicit* in terms.
- Functions restricted to work on types, *not* type schemes, to preserve the stratification of types and schemes.

Does this work as we want? We can derive a polymorphic type for the identity function, as we'd expect.

$$\frac{\dfrac{}{\{a \mapsto \alpha\} \vdash a : \alpha}}{\dfrac{\emptyset \vdash \lambda a.a : \alpha \to \alpha}{\emptyset \vdash \lambda a.a : \forall \alpha.\alpha \to \alpha}}$$

But we can also derive some types that we might not want! For example:

$$(*)\ \frac{\dfrac{\dfrac{\dfrac{}{\{a \mapsto \alpha\} \vdash a : \alpha}}{\{a \mapsto \alpha\} \vdash a : \forall \alpha.\alpha}}{\dfrac{\{a \mapsto \alpha\} \vdash a : \beta}{\dfrac{\emptyset \vdash \lambda a.a : \alpha \to \beta}{\emptyset \vdash \lambda a.a : \forall \beta.\alpha \to \beta}}}}{\emptyset \vdash \lambda a.a : \forall \alpha.\forall \beta.\alpha \to \beta}$$

What's gone wrong? Our intuition about type variables wasn't quite correct.

- Variables in the type of a term may capture dependence of the term on its environment
- Variables only truly represent arbitrary types if they don't appear in the environment. (Called the *eigenvariable* condition.)

We can update our quantification rules to capture the eigenvariable condition:

$$\frac{\Gamma \vdash t : \forall \alpha.S}{\Gamma \vdash t : S[T/\alpha]} \qquad \frac{\Gamma \vdash t : S}{\Gamma \vdash t : \forall \alpha.S}\ (\alpha \notin fv(\Gamma))$$

where $fv(\Gamma)$ is defined in terms of fv on schemes in an entirely boring fashion. Now the bogus generalization step (marked with an $*$) in the derivation above isn't valid—because α is in $fv(\{a \mapsto \alpha\})$.

We have one additional problem: because we've restricted functions we don't have any way to bind polymorphic values. We'll add a new language feature to introduce polymorphism... or actually, reuse an existing language feature.

$$\mathcal{T} \ni t ::= z \mid t \odot t \mid x \mid \lambda x.t \mid t\,t \mid \texttt{let } x = t \texttt{ in } t$$

with the typing rule

$$\frac{\Gamma \vdash t_1 : S \quad \Gamma[x \mapsto S] \vdash t_2 : T}{\Gamma \vdash \texttt{let } x = t_1 \texttt{ in } t_2 : T}$$

This is *not* the way we've been using $\texttt{let}$ so far! We can't interpret $\texttt{let } x = t_1 \texttt{ in } t_2$ as $(\lambda x.t_2)\,t_1$. In the latter, x would have to have a monomorphic type, where our goal here is to make it polymorphic. Now, we can return to our starting example:

$$\frac{\dfrac{\dfrac{\{a \mapsto \alpha\} \vdash a : \alpha}{\emptyset \vdash \lambda a.a : \alpha \to \alpha}}{\emptyset \vdash \lambda a.a : \forall \alpha.\alpha \to \alpha} \quad \dfrac{\dfrac{\dfrac{\Gamma \vdash f : \forall\alpha.\alpha \to \alpha}{\Gamma \vdash f : (\texttt{Int} \to \texttt{Int}) \to (\texttt{Int} \to \texttt{Int})} \quad \dfrac{\Gamma \vdash f : \forall\alpha.\alpha \to \alpha}{\Gamma \vdash f : \texttt{Int} \to \texttt{Int}}}{\Gamma \vdash f\,f : \texttt{Int} \to \texttt{Int}} \quad \dfrac{}{\Gamma \vdash 1 : \texttt{Int}}}{\Gamma \vdash f\,f\,1 : \texttt{Int}}}{\emptyset \vdash \texttt{let } f = \lambda a.a \texttt{ in } f\,f\,1 : \texttt{Int}}$$

where $\Gamma = \{f \mapsto \forall\alpha.\alpha \to \alpha\}$.

3. Principal Types and Type Inference

We'll finish with a brief discussion of why Hindley-Milner polymorphism is restricted the way it is. One of the key properties of the Hindley-Milner system is that it has principal types.

Theorem 19.3. *If $\Gamma \vdash t : S_1$ and $\Gamma \vdash t : S_2$, then there is some S such that $S \geq S_1$, $S \geq S_2$ and $\Gamma \vdash t : S$.*

Intuitively, this means that every typeable term has a most general type, and that if we can figure out that most general type it encompasses all the other types at which that term could be used. The practical consequence is that type inference can be useful. So long as the compiler has some way of computing a term's most general type, then there's no danger that the type the compiler computes isn't as good as the one you might intend.

I'm not going to prove principality for H-M, although I'm happy to talk about how you might outside of class.

Now suppose we relax the Hindley-Milner constraints just a little—we'll allow quantifiers to be nested in types. This gives a much simpler system—no more stratification of types and schemes, and no more special treatment of $\texttt{let}$. And, it's more expressive. But we can show that this system can't have principal types either. Consider these two typings:

$$\emptyset \vdash \lambda f.\lambda a.\lambda b.(f\,a, f\,b) : \forall\alpha.\forall\beta.(\alpha \to \beta) \to \alpha \to \alpha \to (\beta, \beta)$$
$$\emptyset \vdash \lambda f.\lambda a.\lambda b.(f\,a, f\,b) : \forall\alpha.\forall\beta.(\forall\gamma.\gamma \to \gamma) \to \alpha \to \beta \to (\alpha, \beta)$$

These types are fundamentally incomparable: the top one lets you pick a function and arguments that match, while the bottom one lets you pick different types for a and b. There's no more-general type that encompasses both of these (take my word for it). So, such a system, as expressive as it might be, can't have a well-defined type inference algorithm that never makes the wrong guess.