CS510
Types and Semantics
Fall'05 Lecture Notes
Type Inference

Type Inference

Type inference is the problem of producing a well-typed term from an **untyped term**. Having a type inference mechanism for a language can be very useful, since it saves the programmer from having to declare the types of all identifiers (though it does not **prevent** such declarations).

Type inference is **not** possible for every language and typing system, but there **is** a type inference algorithm relating the untyped and simply-typed λ-calculi.

Formally, we define a function **erase**, which takes a typed term, removes all the type annotations from the λ-bindings, and returns the resulting untyped term.

An untyped term M is **typable** if there exists a (simply-)typed term M', a type τ and an environment E (with FV(M) $\in$ Dom(E)) such that

$E \vdash M' : \tau$ and $M = \text{erase}(M')$

A (typable) untyped term may be the erasure of more than one typed term. Example: $\lambda\texttt{x}.\texttt{x}$ is the erasure of both $\lambda\texttt{x}:\texttt{int}.\texttt{x}$ and $\lambda\texttt{x}:\texttt{bool}.\texttt{x}$.

Type Schemes

To make the type inference algorithm precise, it is helpful to extend the language of type expressions to include **type variables**, denoted by $\{\alpha, \beta, \gamma, \ldots\}$.

The resulting grammar for **type schemes** is:

$$\tau ::= b \mid \alpha \mid \tau_1 \rightarrow \tau_2 \qquad (b \in B)$$

They are called schemes because each expression can be thought of as a schematic representation of a whole family of (ordinary) types, produced by **instantiating** the type variables with ordinary types.

Examples

α
$\alpha \rightarrow \beta$
$\alpha \rightarrow (\texttt{int} \rightarrow \alpha)$
`bool`

An ordinary type is just a type scheme with no type variables. They are also called **ground types** or **monotypes**.

We modify our typing system so that typing rules, environments, λ-binding annotations, assertions, etc., all refer to type **schemes** where they used to refer to ground types. All this requires is a suitable change in the interpretation of the metavariables τ.

Type Inference

Our type inference algorithm will produce a type **scheme** for a given expression (in a given environment). Working with type schemes allows us to concentrate on the **structure** of the term being typed (and the corresponding typing derivation), which is the same regardless of the specific types in question.

Recall that the form of a type derivation for a given term is a tree isomorphic to the syntax tree for the term. To infer a type for the term, we start with a derivation tree of the correct form in which

- each λ-binding type scheme annotation is a fresh type variable

- the type scheme of each term is a fresh type variable.

Example $\quad \lambda \mathtt{y.y} \ 3$

$$
\cfrac{\cfrac{}{\{\mathtt{y}:\alpha_0\} \ \vdash \ \mathtt{y}:\alpha_3} VAR \quad \cfrac{}{\{\mathtt{y}:\alpha_0\} \ \vdash \ 3:\alpha_4} CONST}{\cfrac{\{\mathtt{y}:\alpha_0\} \ \vdash \ (\mathtt{y} \ 3):\alpha_2}{\emptyset \ \vdash \ (\lambda \mathtt{y}:\alpha_0.\mathtt{y} \ 3):\alpha_1} ABS} APP
$$

To make this a **valid** derivation, certain relationships must hold between the type variables mentioned in adjacent rules. We can express these in the form of **equations between type schemes**. These equations represent **constraints** on the possible (valid) instantiations of the type variables.

Example continued

$$\frac{\overline{\{\texttt{y}:\alpha_0\} \ \vdash \ \texttt{y}:\alpha_3}}{} VAR \quad \frac{\overline{\{\texttt{y}:\alpha_0\} \ \vdash \ \texttt{3}:\alpha_4}}{} \begin{matrix} CONST \\ APP \end{matrix}$$
$$\frac{\{\texttt{y}:\alpha_0\} \ \vdash \ (\texttt{y} \ \texttt{3}):\alpha_2}{\emptyset \ \vdash \ (\lambda\texttt{y}:\alpha_0.\texttt{y} \ \texttt{3}):\alpha_1} ABS$$

Each use of an axiom or rule produces one constraint equation. In this example, we get the following

ABS rule:	$\alpha_1 = \alpha_0 \rightarrow \alpha_2$
APP rule:	$\alpha_3 = \alpha_4 \rightarrow \alpha_2$
VAR rule:	$\alpha_3 = \alpha_0$
CONST rule:	$\alpha_4 = \texttt{int}$

To produce a valid type derivation, we must solve all the constraint equations **simultaneously**.

Here's one way to do this. We eliminate α_3 and α_4 using the (CONST) and (VAR) constraints, leaving

$\alpha_1 = \alpha_0 \rightarrow \alpha_2$
$\alpha_0 = \texttt{int} \rightarrow \alpha_2$

Substituting the second equation into the first, we get:

$\alpha_1 = (\texttt{int} \rightarrow \alpha_2) \rightarrow \alpha_2$

Example continued again

Instantiating the original typing tree with this solution, we have the following derivation:

$$\dfrac{\dfrac{}{\{\texttt{y:int} \to \alpha_2\} \,\vdash\, \texttt{y:int} \to \alpha_2} \, VAR \quad \dfrac{}{\{\texttt{y:int} \to \alpha_2\} \,\vdash\, \texttt{3:int}} \, CONST}{\dfrac{\{\texttt{y:int} \to \alpha_2\} \,\vdash\, \texttt{(y 3):}\alpha_2}{\emptyset \,\vdash\, (\lambda\texttt{y:int} \to \alpha_2.\texttt{y 3}):(\texttt{int} \to \alpha_2) \to \alpha_2} \, ABS} \, APP$$

It is easy to see that every **instance** of

$$\emptyset \,\vdash\, (\lambda\texttt{y:int} \to \alpha_2.\texttt{y 3}):(\texttt{int} \to \alpha_2) \to \alpha_2$$
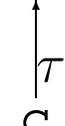
i.e., every assertion resulting from instantiation of the remaining type variable (α_2), will also be derivable. Moreover, since this scheme was constrained as minimally as possible, **every** derivable scheme must be an instance of it.
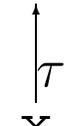
Such a type scheme is called **principal**. The approach outlined here always generates a principal type scheme, (if the given term is typable).
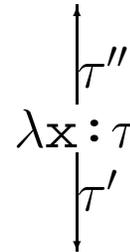
The inference algorithm consists of two steps:
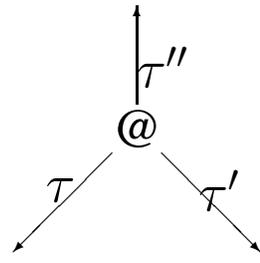
• The term's structure is analyzed to derive a set of constraint equations. This can be described simply in terms of the syntax tree.

• The equations are solved simultaneously. This is an instance of the **unification** problem.
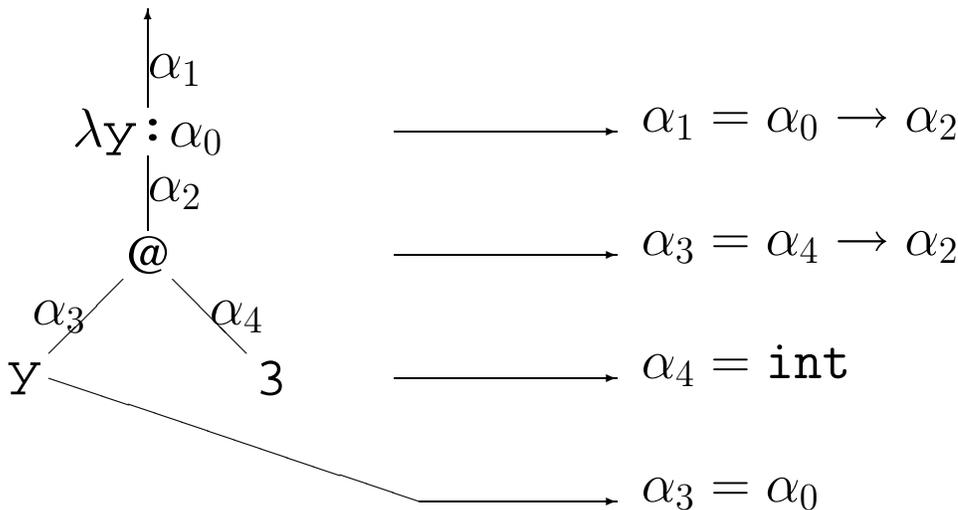
Constraints

(CONST)   E⊢c:τ (TypeOf(c) $=\tau$)

$$\begin{array}{c} \uparrow \\ |\tau' \\ \texttt{c} \end{array} \qquad \tau = \tau'$$

(VAR)   E⊢x:τ (E(x) $=\tau$)

$$\begin{array}{c} \uparrow \\ |\tau' \\ \texttt{x} \end{array} \qquad \tau = \tau'$$

(ABS) $\dfrac{\texttt{E + \{x:}\tau\texttt{\}} \vdash \texttt{M:}\tau'}{\texttt{E} \vdash (\lambda\texttt{x:}\tau\texttt{.M):}\tau''}$

$$\begin{array}{c} \uparrow \\ |\tau'' \\ \lambda\texttt{x:}\tau \\ |\tau' \\ \downarrow \end{array} \qquad \tau'' = \tau \rightarrow \tau'$$

(APP) $\dfrac{\texttt{E} \vdash \texttt{M:}\tau \qquad \texttt{E} \vdash \texttt{N:}\tau'}{\texttt{E} \vdash (\texttt{MN):}\tau''}$

$$\begin{array}{c} \uparrow \\ |\tau'' \\ @ \\ \swarrow \tau \quad \tau' \searrow \end{array} \qquad \tau = \tau' \rightarrow \tau''$$

Given a complete expression, we patch the patterns together to get constraint equations, e.g.:

$$\begin{array}{c} \uparrow \\ |\alpha_1 \\ \lambda\texttt{y:}\alpha_0 \\ |\alpha_2 \\ @ \\ \alpha_3 \swarrow \quad \searrow \alpha_4 \\ \texttt{y} \qquad \texttt{3} \end{array}$$

$$\xrightarrow{\hspace{2cm}} \alpha_1 = \alpha_0 \rightarrow \alpha_2$$

$$\xrightarrow{\hspace{2cm}} \alpha_3 = \alpha_4 \rightarrow \alpha_2$$

$$\xrightarrow{\hspace{2cm}} \alpha_4 = \texttt{int}$$

$$\xrightarrow{\hspace{2cm}} \alpha_3 = \alpha_0$$

Constraint Equations

We now have a collection of equations between type schemes

$\tau_{11} = \tau_{12}$
$\tau_{21} = \tau_{22}$
$\ldots$
$\tau_{n1} = \tau_{n2}$

to be solved simultaneously.

This is a special case of the general problem called **unification**. A solution to a unification problem is a **substitution** from variables to terms that makes all the equations true. Such a substitution is called a **unifier**.

Example: the set of equations

$\alpha_1 = \alpha_0 \rightarrow \alpha_2$
$\alpha_3 = \alpha_4 \rightarrow \alpha_2$
$\alpha_4 = \mathtt{int}$
$\alpha_3 = \alpha_0$

has the unifier

$$S = \{ \; \alpha_1 \mapsto (\mathtt{int} \rightarrow \alpha_2) \rightarrow \alpha_2,$$
$$\alpha_0 = \alpha_3 \mapsto \mathtt{int} \rightarrow \alpha_2,$$
$$\alpha_4 \mapsto \mathtt{int}\}$$

Moreover, S is a **most general unifier**, meaning that every unifying substitution S' can be obtained from S by a further substitution.

Unification

In general, the unification problem is as follows:

Given a set of equations between terms built from variables and constructors of (fixed) arity ≥ 0, either:

- compute a most general unifier that solves the equations; or

- report that the equations have no solution.

(It can be proved that if the set of equations have any unifier, they have a most general unifier.)

Unification is a widely applicable technique; e.g., it is the fundamental computational process in logic programming. There are several efficient algorithms known for computing unifiers.

For type inference, we specialize the problem as follows:

- the variables are type variables;

- there is one constructor ($\rightarrow$) of arity 2;

- each base type is treated as a constructor of arity 0.

If there is a most general unifier for the equations, it yields a **principal** schematic typing for the original λ-expression. If there is no solution, the original expression is not typable.

A Unification Algorithm

The essence of unification is the recursive traversal of two terms to be equated ("unified").

We denote a substitution by S, and extend its domain from variables to arbitrary terms in the obvious way. Substitutions can be **composed**; we write $S_1 \circ S_2$ to mean the result of applying S_2 and then S_1.

Function $\mathcal{U}$ takes an existing substitution and two terms, and returns an extension of the original substitution that unifies the terms (if possible):

$$\mathcal{U}(S, t_1, t_2) =$$

$$\text{case } (S(t_1), S(t_2)) \text{ of}$$

$$(v_1, v_2) : (v_1 \mapsto v_2) \circ S$$

$$(v_1, t_2') : \text{if } v_1 \text{ occurs in } t_2' \text{ then error ``No Unifier''}$$

$$\text{else } (v_1 \mapsto t_2') \circ S$$

$$(t_1', v_2) : \text{if } v_2 \text{ occurs in } t_1' \text{ then error ``No Unifier''}$$

$$\text{else } (v_2 \mapsto t_1') \circ S$$

$$(C_1(t_{11}, t_{12}, \ldots, t_{1m}), C_2(t_{21}, t_{22}, \ldots, t_{2n})) :$$

$$\text{if } C_1 \neq C_2$$

$$\text{or } m \neq n \text{ then error ``No Unifier''}$$

$$\text{else } \mathcal{U}(\ldots \mathcal{U}(\mathcal{U}(S, t_{11}, t_{21}), t_{12}, t_{22}) \ldots, t_{1m}, t_{2m})$$

To find the m.g.u. of a set of equations

$$t_{11} = t_{21}, t_{12} = t_{22}, \ldots, t_{1n} = t_{2n}$$

we calculate $\mathcal{U}(\ldots \mathcal{U}(\mathcal{U}(\emptyset, t_{11}, t_{21}), t_{12}, t_{22}) \ldots, t_{1n}, t_{2n})$

Example Again

Use the unification algorithm to solve the set of constraints generated from $\lambda \texttt{x}.\texttt{x}\ \ 3$. To make the role and arity of the constructors clearer, we rewrite the infix $\rightarrow$ constructor as the prefix arity-2 constructor $C_\rightarrow$, and the base type $\texttt{int}$ as the arity-0 constructor $C_{\texttt{int}}$.

$$\alpha_1 = C_\rightarrow(\alpha_0, \alpha_2)$$
$$\alpha_3 = C_\rightarrow(\alpha_4, \alpha_2)$$
$$\alpha_4 = C_{\texttt{int}}$$
$$\alpha_3 = \alpha_0$$

We calculate the m.g.u. by calling $\mathcal{U}$ on each equation in turn, passing the result substitution of each call as the initial substitution for the next:

$$S_1 = \mathcal{U}(\emptyset, \alpha_1, C_\rightarrow(\alpha_0, \alpha_2))$$
$$S_2 = \mathcal{U}(S_1, \alpha_3, C_\rightarrow(\alpha_4, \alpha_2))$$
$$S_3 = \mathcal{U}(S_2, \alpha_4, C_{\texttt{int}})$$
$$\text{m.g.u.} = S_4 = \mathcal{U}(S_3, \alpha_3, \alpha_0)$$

(Notice that the result substitution grows steadily; the algorithm never backtracks. An efficient implementation can therefore maintain a single substitution in a global variable rather than threading it through each call to $\mathcal{U}$.)

Example continued

To pursue the calculation:

$$
\begin{aligned}
S_1 \;=\; \{ &\; \alpha_1 \mapsto C_\rightarrow(\alpha_0, \alpha_2)\} \\
S_2 \;=\; \{ &\; \alpha_3 \mapsto C_\rightarrow(\alpha_4, \alpha_2)\} \circ S_1 \\
\;=\; \{ &\; \alpha_1 \mapsto C_\rightarrow(\alpha_0, \alpha_2), \\
&\; \alpha_3 \mapsto C_\rightarrow(\alpha_4, \alpha_2)\} \\
S_3 \;=\; \{ &\; \alpha_4 \mapsto C_{\texttt{int}}\} \circ S_2 \\
\;=\; \{ &\; \alpha_1 \mapsto C_\rightarrow(\alpha_0, \alpha_2), \\
&\; \alpha_3 \mapsto C_\rightarrow(C_{\texttt{int}}, \alpha_2), \\
&\; \alpha_4 \mapsto C_{\texttt{int}}\} \\
S_4 \;=\; \{ &\; \alpha_0 \mapsto C_\rightarrow(C_{\texttt{int}}, \alpha_2)\} \circ S_3 \\
\;=\; \{ &\; \alpha_1 \mapsto C_\rightarrow(C_\rightarrow(C_{\texttt{int}}, \alpha_2), \alpha_2), \\
&\; \alpha_3 \mapsto C_\rightarrow(C_{\texttt{int}}, \alpha_2), \\
&\; \alpha_4 \mapsto C_{\texttt{int}}\} \\
&\; \alpha_0 \mapsto C_\rightarrow(C_{\texttt{int}}, \alpha_2)\}
\end{aligned}
$$

This is the same m.g.u. we derived before by ad-hoc substitution.

Note that the compositions performed in the calculation of S_1 and S_2 are trivial, whereas those in S_3 and S_4 are less so.

It is crucial to remember that the substitution argument to $\mathcal{U}$ must be applied to both the term arguments **before** performing the case dispatch. This is illustrated in the calculation of the new substitution component of S_4.

Polymorphism

Polymorphism means using the same code to operate on values of different types at different points in the same program. We are mainly interested in so-called **parametric polymorphism**, which refers to functions that are "oblivious" to the types of their arguments. An example is the function

```
pair ≡ λx.(x,x)
```

which forms pairs of any type. Obviously, we need some conventions for uniform data representation in order to actually compile such functions.

We distinguish this from **ad-hoc polymorphism**, a.k.a. **overloading**, where the same source code function must have different runtime behavior depending on its argument's types, e.g.,

```
double ≡ λx.x + x
```

which has different behavior for `x:real` and `x:int`.

We want a type system for polymorphism that is (efficiently) decideable, sound, and has an (efficient decideable) inference algorithm.

Limitations of schemes

We already are part-way there through our use of type schemes.

When we derive a schematic type for `pair`, e.g.,

$$\emptyset \vdash \lambda \texttt{x.(x,x)} \;:\; \alpha \to (\alpha \times \alpha)$$

this implies that the application of `pair` to an expression of **any** type τ will type-check, with α being unified to τ.

Unfortunately, we have no way to express the idea that `pair` may be used with **different** types in different parts of the same expression, because this would require unifying α with two conflicting types.

For example, the following will not type-check:

$$(\lambda \texttt{f.(f 3, f true)})(\lambda \texttt{x.(x,x)})$$

since α would need to unify with both `int` and `bool`.

What we need is a way to **quantify** type variables over sub-expressions, so they can be **instantiated** to different types at different points in the expression.

2nd-order Polymorphic λ-calculus

One way to express polymorphic programs of this kind is to use an **explicitly-typed** calculus which allows **abstraction** and **application** of type variables as well as ordinary variables.

Examples

```
(λf:∀α.α → (α × α).(f [int] 3, f [bool] true))
            (Λα.λx:α.(x,x))
```

```
(λa:int.λb:bool.λf:∀α.α → α.(f [int] a,f [bool] b))
            1 true (Λα.λx:α.x)
```

In general, expressions are:

$$
\begin{array}{lll}
M & := \ c & \text{Built-in constants} \\
& | \quad v & \text{Variable names} \\
& | \quad (M_1 M_2) & \text{Function applications} \\
& | \quad (\lambda v : \sigma.M) & \text{Function abstractions} \\
& | \quad (\Lambda \alpha.M) & \text{Type abstractions} \\
& | \quad (M[\sigma]) & \text{Type applications}
\end{array}
$$

where polymorphic type schemes are:

$$\sigma := b \mid \alpha \mid \sigma_1 \to \sigma_2 \mid \forall \alpha.\sigma$$

ML Polymorphism

It is perfectly possible to give (sound) typing rules for the 2nd-order polymoprhic calculus. We'll investigate this in the next chapter. But meanwhile we want a type inference mechanism, so that the user doesn't have to write down these complex types.

Unfortunately, this calculus has no principal types, and typability of untyped terms is actually undecidable.

So nearly all FL's use a more restrictive style of polymorphism (called **ML-polymorphism**, `let`-**style polymorphism**, **Hindley-Milner polymorphism**, etc.) for which there **is** an inference algorithm.

The key idea is that we can do inference successfully if we have a visible **definition** for every polymorphic function.

Syntax

We use the following **untyped** syntax for expressions:

$$
\begin{array}{lll}
M & := & c & \text{Built-in constants} \\
& | & v & \text{Variable names} \\
& | & (M_1 M_2) & \text{Function applications} \\
& | & (\lambda v.M) & \text{Function abstractions} \\
& | & (\texttt{let } v = M_1 \texttt{ in } M_2) & \texttt{let} \text{ abstractions}
\end{array}
$$

Note that `let` is now a basic expression form in its own right. `let x = a in b` still has the same **operational** semantics as `(`λ`x.b)a`, but the two expressions will be **different** for typing purposes. Specifically, `let` is used to define expressions that are to be used **polymorphically**.

We would rewrite our first earlier example as:

`let f = `λ`x.(x,x) in (f 3, f true)`

The use of `let` is essential; we **can't** write the body here as an ordinary λ-abstraction

`$\lambda$f.(f 3, f true)`

as this will still not type check.

We cannot express the second example at all, however; this style of polymorphism is strictly less powerful.

Top-level Expressions

Note that expressions entered at the "top-level" in CAML are treated like `let`-bound expressions; that is, for typing purposes the sequence

```
let a = exp₁;;
let b = exp₂;;
...;;
```

is equivalent to

```
let a = exp₁ in let b = exp₂ in ...;;
```

In particular, library functions are treated as if they were `let`-bound in the scope of ordinary programs.

ML Polymorphism Type System

We can give rules similar to the ones we introduced earlier for the simply-typed λ-calculus.

In these rules, we use ordinary type schemes just as before

$$\tau ::= b \mid \alpha \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \times \tau_2 \mid \ldots \qquad (\, b \in B \,)$$

and also **polymorphic type schemes** of the following form

$$\sigma ::= \tau \mid \forall \alpha.\sigma$$

Note that in these polymorphic schemes any quantifiers always come **first**. For example,

$$\forall \alpha.(\alpha \rightarrow (\alpha \times \alpha))$$

is a legitimate polymorphic scheme;

$$(\forall \alpha.\alpha) \rightarrow (\forall \alpha.\alpha)$$

is not.

Instantiation

To get an (ordinary) type scheme from a polymorphic scheme $\sigma \equiv \forall \alpha_1 \forall \alpha_2 \ldots \forall \alpha_n.\tau$ we instantiate each of the quantified type variables with a type. The resulting scheme will be of the form $\tau' \equiv S(\tau)$, where S is substitution with domain $\{\alpha_1, \ldots, \alpha_n\}$. In this case we write

$$\sigma > \tau'$$

Examples

$$\forall \alpha.\forall \beta.\alpha \rightarrow \gamma \rightarrow \beta > \texttt{int} \rightarrow \gamma \rightarrow \texttt{bool}$$
$$\text{with } S = \{\alpha \mapsto \texttt{int}, \beta \mapsto \texttt{bool}\}$$
$$\forall \alpha.\forall \beta.\alpha \rightarrow \gamma \rightarrow \beta > \texttt{int} \rightarrow \gamma \rightarrow (\delta \rightarrow \delta)$$
$$\text{with } S = \{\alpha \mapsto \texttt{int}, \beta \mapsto (\delta \rightarrow \delta)\}$$
$$\forall \alpha.\alpha \rightarrow \beta \rightarrow \alpha \not> \texttt{int} \rightarrow \texttt{bool} \rightarrow \texttt{int}$$

Typing Rules

In these rules, environments E map (expression) variables to
polymorphic type schemes, and `TypeOf` also returns a
polymorphic scheme.

$$\frac{}{E \vdash \texttt{c:}\tau \ (\forall\texttt{c},\tau \ \texttt{s.t.} \ \ \texttt{Typeof(c)} \ > \ \tau)} CONST$$

$$\frac{}{E \vdash \texttt{x:}\tau \ \ (\texttt{if} \ \texttt{E(x)} \ > \ \tau)} VAR$$

$$\frac{E \ + \ \{\texttt{x:}\tau_1\} \ \vdash \ M \ : \ \ \tau_2}{E \vdash \ (\lambda\texttt{x:}\tau_1.M) \ : \ \ \tau_1 \to \tau_2} ABS$$

$$\frac{E \ \vdash \ M \ : \ \ \tau_1 \to \tau_2 \quad E \ \vdash \ N \ : \ \ \tau_1}{E \ \vdash \ (MN) \ : \ \ \tau_2} APP$$

$$\frac{E \ \vdash \ M \ : \ \ \tau \quad E \ + \ \{x : \mathrm{Clos}(\tau, E)\} \ \vdash \ N \ : \ \ \tau_1}{E \vdash \texttt{let} \ \texttt{x} \ = \ M \ \texttt{in} \ N \ : \ \ \tau_1} LET$$

Here $\mathrm{Clos}(\tau, E)$ denotes the "closure" of the type τ formed by
abstracting over all the free type variables in τ that are **not** free in
E.

Formally

$$\mathrm{Clos}(\tau, E) = \forall\alpha_1\forall\alpha_2\ldots\forall\alpha_n\tau$$

where $\{\alpha_1, \alpha_2, \ldots, \alpha_n\} = FV(\tau) - FV(E)$.

Example Derivation

We show a typing derivation for our previous example

```
let f = λx.(x,x) in (f 3, f true)
```

To handle this example, we need to add a typing rule to deal with the `pair` type constructor.

$$\frac{E \ \vdash \ M_1 \ : \ \tau_1 \quad E \ \vdash \ M_2 \ : \ \tau_2}{E \ \vdash \ (M_1, M_2) \ : \ \tau_1 \times \tau_2} \, PAIR$$

The derivation is then as follows (use xdvi to read this!)

$$\frac{E_1 \vdash \text{x} : \alpha}{} VAR \quad \frac{E_1 \vdash \text{x} : \alpha}{} VAR \quad \frac{E_2 \vdash \text{f} : \text{int} \rightarrow (\text{int} \times \text{i}}{E_2 \vdash \text{f 3} :}$$
$$\frac{}{E_1 \vdash (\text{x,x}) : (\alpha \times \alpha)} PAIR$$
$$\frac{\emptyset \vdash \lambda\text{x.(x,x)}}{\emptyset \vdash \text{let f} = \lambda\text{x.(x,x) in (f 3, f t}}$$

where:

$$E_1 = \{\text{x} : \alpha\}$$
$$E_2 = \{\text{f} : \forall\alpha.\alpha \rightarrow (\alpha \times \alpha)\}$$

The Closure Rule

In the previous example, the Clos rule generalized **all** the type variables in the type derived for the `let`-defining expression. But in general, the rule acts to avoid abstracting over type variables that are really acting as type constants in the current environment.

Consider this example:

λx.let f = λy.x in ((f 3) + 1, not(f 3))

Obviously, this should not be typable, since no x can be a legitimate argument to both + and `not`. But without the restrictions on closures, for we could derive its type as

$\alpha \rightarrow (\text{int} \times \text{bool})$.

The problem would arise because we would type the pair expression in the environment:

$\{\text{x} : \alpha, \text{f} : \forall \alpha \forall \beta.\beta \rightarrow \alpha\}$

But the α in the type scheme for x and the α in the type scheme for f must be the **same** α, so the polymorphic scheme for f should really be quantified **only** over β.

Type Inference

Notice that type abstraction **always** occurs as part of the LET rule and type application occurs as part of the VAR rule. Because of this, polymorphic schemes appear only in the environment; they are never attached directly to expressions. Sometimes a similar system is given with type abstraction and application split out as separate rules; the systems are equivalent.

The advantage of our approach is that, as in the simply-typed case, there is at most one rule applicable to any expression, which can be chosen based completely on its syntactical form. As before, this gives an immediate **type inference** algorithm for this system: use the syntactical structure to generate a schematic derivation and a set of constraints, and use unification to solve the constraints.

In practice, the derivation and solution of the constraints is normally combined into a one-pass algorithm over the structure. For historical reasons, the inference algorithm for ML-style polymorphism is called **Algorithm W**; see text.

The key point about this algorithm is that a polymorphic schematic type of a `let`-bound variable is always deduced **before** the variable's uses are inspected; this is what makes inference feasible.

Efficiency of Algorithm W

Interestingly, this is an example of an algorithm that is efficient in practice but very inefficient (deterministic exponential time complete) in theory. For example, the following nasty example

```
let pair = λx.λy.λz.zxy
  in let x₁ = λy.pair y y
     in let x₂ = λy.x₁(x₁(y))
        in ...
                 in let xₙ = λy.xₙ₋₁(xₙ₋₁(y))
                    in xₙ(λz.z)
```

produces a principal type of length $2^{2^{cn}}$ (as a string). For $n = 5$ this produces 173 printed pages of output after several hours.

Recursion

As before, we can't define the Y combinator directly in an ML-polymorphism type system, so we need to add explicit support for recursion. Here's a possible definition for a LETREC rule:

$$\frac{E \,+\, \{f : \tau_2 \to \tau_1, x : \tau_2\} \,\vdash\, M \,:\, \tau_1 \quad E \,+\, \{f : \mathrm{Clos}(\tau_2 \to \tau_1, E)\} \,\vdash\, N \,:\, \tau}{E \,\vdash\, \texttt{letrec f = } \lambda\texttt{x}.M \texttt{ in } N \,:\, \tau}$$

Example

```
letrec map = λf.λx.if null x
      then nil
      else cons(f (hd x),map f (tail x))
  in map negate [1,2,3]; map not [true,false]
```

will typecheck correctly as expected, with `map` being given a polymorphic type scheme

$\forall\alpha.\forall\beta.(\alpha \to \beta) \to \alpha \texttt{ list} \to \beta \texttt{ list}.$

Note that this typing rule doesn't allow us to define a recursive function that invokes itself in a polymorphic way:

```
letrec f = λn.λx.if n = 0
      then x
      else f (n-1) (x,x)
```

Unfortunately, if we extend the typing system to permit expressions of this kind to be typable, type inference becomes undecideable!

Side-effects

We have to be very careful in extending these typing rules to programs with side-effects. For example, we might naively add the following polymorphic type schemes for ML-style references to `TypeOf`:

$$\{ \ \mathtt{ref} : \forall \alpha.\alpha \rightarrow \alpha \mathtt{ref},$$
$$\mathtt{:=} \ : \forall \alpha.\alpha \ \mathtt{ref} \rightarrow \alpha \rightarrow \mathtt{unit},$$
$$\mathtt{!} : \forall \alpha.\alpha \ \mathtt{ref} \rightarrow \alpha\}$$

where `ref` is a new type constructor.

Unfortunately, under these rules the following expression is typable, though it surely shouldn't be!

```
let r = ref(λx.x) in
    (r := λx.x+1; (!r) true)
```

Intuitively, the problem is that we infer the type $(\alpha \rightarrow \alpha) \ \mathtt{ref}$ for `r`'s defining expression, and then **abstract** over α to form a polymorphic scheme for `r`. This abstraction is bogus, since α really refers to the particular type of the value currently held by the reference.

Similar problems arise with other imperative features, such as exceptions. Various solutions to these problems have been proposed and implemented. The cleanest is to put a restriction on the form of `let` definitions, namely the defining expression must be a λ-abstraction, a constant, or a simple variable. In all these cases, the **cbv evaluation** of the expression is essentially a no-op, which turns out to avoid the typing problems.