

Programming Languages

Assignment #7

December 2, 2007

1 Introduction

This assignment has 20 points total.

In this assignment, you will write a type-checker for the PolyMinML language (a language that is pretty similar to the polymorphic lambda calculus). This type checker will perform *bi-directional type inference*, allowing many type annotations to be inferred automatically from the context of a term. Unlike ML-style type inference, which only works for languages with predicative prenex-polymorphism, bi-directional type inference scales to much richer languages, including languages with general polymorphism (like in System F/PolyMinML) and subtyping. The downside of bi-directional type inference is that programmers must write down some typing annotations in their programs.

2 Type inference for PolyMinML

Consider the following example:

```
fun f(x : int) : unit -> unit =  
  (fun g(y : unit) : unit = y)  
end
```

In this example, the type annotations for the argument and return type of `g` are redundant; because `f` specifies its return type, we know what the type of `g` must be. This frequent repetition of type information makes using plain PolyMinML tedious. Take a look at how we might write and use the polymorphic identity function.

```
if ((typefun t. fun f(x : t) : t = x) [ bool ] true) then 1 else 0
```

Here, the instantiation of the polymorphic function at type `bool` is redundant; we know that the result must be of type `bool`, since it is the conditional in an `if` expression. Thus, some expressions need type annotations, and some do not. The essence of our type-checking algorithm will be to make a syntactic distinction between expressions whose type can be inferred, and those whose types can be checked (given some context about what their type *should* be).

The intuition is that *checkable* expressions appear only where their type can be determined by context, and so they don't need type annotations. *Inferable* expressions are what we've primarily used in previous type checkers. These have a unique type which can be determined in a particular variable context Γ .

With two different kinds of expressions, we have two forms of typing judgments:

$$\overline{\Delta, \Gamma \vdash c \downarrow \tau}$$

$$\overline{\Delta, \Gamma \vdash i \uparrow \tau}$$

The first says that given a type variable context Δ and an expression variable context Γ , *as well as the target type* τ , we can verify that the checkable expression c has the given type τ .

The second says that given similar contexts Δ and Γ , we can deduce the unique type τ for the inferable expression i .

To recap: We separate expressions into two categories, checkable and inferable. Then we have two separate judgments: one which says that given a type, we can verify a checkable expression has that type, and one which allows us to deduce the type of an inferable expression. The next two sections give a precise description of the abstract syntax and the algorithm we use to check checkable expressions and infer the types of inferable expressions.

2.1 PolyMinML Abstract Syntax

```
inferable  $i ::= ()$ 
  |  $x$ 
  |  $n$ 
  | true
  | false
  | primop ( $c_1, c_2, \dots$ )
  |  $i\ c$ 
  | fun  $f\ (x : \tau_d) : \tau_r = c$ 
  | if  $c_b$  then  $i_t$  else  $i_f$  fi
  |  $i\ [\ \tau\ ]$ 
  | typefun  $t.\ i$ 

checkable  $c ::= I\ (i)$ 
  | cfun  $f(x) = c$ 
  | cif  $c_b$  then  $c_t$  else  $c_f$  fi
  |  $i\ []$ 
  |  $c\ i$ 
  | ctypefun  $t.\ c$ 
```

All of the inferable expressions should be familiar. For each one, we can determine its unique type using the standard typing rules. Note the places where checkable expressions c appear. If we know the type of a function (because it is inferable), we know what the type of the argument must be, so the argument need only be a *checkable* expression. The body of a function is also checkable, since its type is given in the function expression. Since we know the types of the arguments to primops, those expressions are also *cs*. Finally, we know that the conditional expression in an **if** must be a boolean, so that too is checkable.

Any inferable expression is checkable, since we can just ignore the context and get the type using the inference algorithm (however, it is still necessary to distinguish an inferable expression from an inferable expression *as a checkable expression*). The form $I\ (i)$ means an i converted into a c . The **cfun** expression is checkable: given the type it is supposed to have (which must be a function type), we can type-check the function as we would for a standard annotated **fun**. The **cif** expression can be checked as well, since we know the condition must be of type **BOOL**, and we know that the types of the branches must match the type of the whole expression. A checkable expression applied to an inferable expression is checkable: we know the return type of the function (it is the type we are checking against), and we know the domain type (it is the type of the argument). **ctypefun** is checkable if its body is checkable; we simply wrap the type of the body in a \forall type. Finally, we are able to instantiate a polymorphic inferable expression at an unspecified type, provided we know what the type of the instantiation is supposed to be. All of these strategies are made precise in the next section.

2.2 PolyMinML Typing Rules

Inferable Expressions

$\overline{\Delta, \Gamma \vdash () \uparrow \text{unit}}$	<i>(unit)</i>
$\overline{\Delta, \Gamma \vdash n \uparrow \text{int}}$	<i>(int)</i>
$\overline{\Delta, \Gamma \vdash \text{true} \uparrow \text{bool}}$	<i>(true)</i>
$\overline{\Delta, \Gamma \vdash \text{false} \uparrow \text{bool}}$	<i>(false)</i>
$\frac{\Gamma(x) = \tau}{\Delta, \Gamma \vdash x \uparrow \tau}$	<i>(var)</i>
$\frac{\Delta, \Gamma \vdash c_1 \downarrow \text{int} \quad \Delta, \Gamma \vdash c_2 \downarrow \text{int}}{\Delta, \Gamma \vdash +(c_1, c_2) \uparrow \text{int}}$	<i>(plus)</i> etc.
$\frac{\Delta, \Gamma \vdash i \uparrow \tau' \rightarrow \tau \quad \Delta, \Gamma \vdash c \downarrow \tau'}{\Delta, \Gamma \vdash i c \uparrow \tau}$	<i>(app)</i>
$\frac{\Delta \vdash \tau_1 \text{ ok} \quad \Delta \vdash \tau_2 \text{ ok} \quad \Delta, \Gamma[f : \tau_1 \rightarrow \tau_2][x : \tau_1] \vdash c \downarrow \tau_2}{\Delta, \Gamma \vdash \text{fun } f(x : \tau_1) : \tau_2 = c}$	<i>(fun)</i>
$\frac{\Delta, \Gamma \vdash c \downarrow \text{bool} \quad \Delta, \Gamma \vdash i_t \uparrow \tau \quad \Delta, \Gamma \vdash i_f \uparrow \tau}{\Delta, \Gamma \vdash \text{if } c \text{ then } i_t \text{ else } i_f \text{ fi} \uparrow \tau}$	<i>(if)</i>
$\frac{\Delta \vdash \tau \text{ ok} \quad \Delta, \Gamma \vdash i \uparrow \forall t. \tau'}{\Delta, \Gamma \vdash i[\tau] \uparrow [\tau/t]\tau'}$	<i>(instantiate)</i>
$\frac{\Delta[t], \Gamma \vdash i \uparrow \tau}{\Delta, \Gamma \vdash \text{typefun } t.i \uparrow \forall t. \tau}$	<i>(forall)</i>

Checkable Expressions

$\frac{\Delta, \Gamma \vdash i \uparrow \tau' \quad \tau = \tau'}{\Delta, \Gamma \vdash \text{I}(i) \downarrow \tau}$	<i>(inferable)</i>
$\frac{\Delta, \Gamma[f : \tau_1 \rightarrow \tau_2][x : \tau_1] \vdash c \downarrow \tau_2}{\Delta, \Gamma \vdash \text{cfun } f(x) = c \downarrow \tau_1 \rightarrow \tau_2}$	<i>(cfun)</i>
$\frac{\Delta, \Gamma \vdash c \downarrow \tau' \rightarrow \tau \quad \Delta, \Gamma \vdash i \uparrow \tau'}{\Delta, \Gamma \vdash c i \downarrow \tau}$	<i>(app_c)</i>
$\frac{\Delta[t], \Gamma \vdash c \downarrow \tau}{\Delta, \Gamma \vdash \text{ctypefun } t.c \downarrow \forall t. \tau}$	<i>(forall_c)</i>
$\frac{\Delta, \Gamma \vdash i \uparrow \forall t. \tau' \quad \exists \sigma. (\tau = [\sigma/t]\tau')}{\Delta, \Gamma \vdash i[] \downarrow \tau}$	<i>(instantiate_c)</i>

These rules reflect the observations we made in the last section. The trickiest one is for the checkable polymorphic instantiation (*instantiate_c*). Here, we require the existence of a type σ which “fills in the holes” in τ' to get τ . Algorithmically, we can do this by structurally comparing τ and τ' , and verifying that any occurrence of the type variable t in τ' is matched up with the same type σ in τ .

2.3 Your Job: Implement Bi-directional Type Checking

The rules given in the previous section are suggestive of an algorithm, consisting of two mutually recursive functions, `check` and `infer`. Both take a set of type variables (Δ), and a variable context (Γ). `infer` then takes an i and returns its unique type τ in that context (if it exists). `check` takes a c and a type τ and verifies that it is compatible with that type.

Question 1. [7 points]

Complete the PolyMinML type-checker by implementing the functions `check` and `infer` in the file `typing.sml`. These functions should meet the following specification:

$$\begin{aligned} \text{check} &: (\text{var} \rightarrow \text{bool}) \rightarrow (\text{var} \rightarrow \text{typ}) \rightarrow \\ &\quad \text{cexp} \rightarrow \text{typ} \rightarrow \text{unit} \end{aligned}$$

Such that `check Δ Γ c τ` returns `unit` if the type is consistent, or raises the exception `TypeError` with an informative message otherwise.

$$\begin{aligned} \text{infer} &: (\text{var} \rightarrow \text{bool}) \rightarrow (\text{var} \rightarrow \text{typ}) \rightarrow \\ &\quad \text{iexp} \rightarrow \text{unit} \end{aligned}$$

Such that `infer Δ Γ i` returns the unique type τ of i in that context, or raises the exception `TypeError` with an informative message if no such τ exists.

Hints

- Avoid unnecessarily inefficient or convoluted code, and document non-trivial invariants/preconditions.
- Make sure to be careful about type variables, especially concerning capture and alpha-equivalence ($\forall t.t \rightarrow t$ and $\forall u.u \rightarrow u$ are the *same type*).
- You will probably want to make use of the functions in `Types` and `PolyMinML` when implementing `check` and `infer`. However, you only need to modify `typing.sml` for your solution to this part of the assignment.
- In order to test your type checker, you need to construct test expressions directly in ML using the datatype declared in the `polyminml.sml` file. We have provided a number of test cases in the file `testcases.sml`. Tests `t1`, `t2`, ... should type-check, and tests `f1`, `f2`, ... should fail. Test one of these expressions with, for example:

```
Typing.typecheck Test.t1;
```

- Make sure that you test your code, and since our test cases are not at all exhaustive, give yourself time to develop your own.

Question 2. [7 points]

Extend the type checking algorithm and implementation to handle existential types. Be sure to:

- Extend the syntax to define the checkable and inferable existential terms.
- Define the bi-directional type checking inference rules for existentials.
- Extend the `PolyMinML` type checker with existential types. This will involve extending the files `polyminml.sml` (be sure to extend the abstract syntax of `PolyMinML` as well as the substitution functions defined in this file), `types.sml` and `typing.sml`.
- Suitable test cases that test both well-typed and ill-typed expressions.

Question 3. [6 points]

Extend the inference rules for type checking given in question 1 so they handle subtyping. In this question, make your typing rules as general as possible. You do not need to define the subtyping relation, just the checkable and inferable rules. You may assume that given two types τ_1 and τ_2 that the relation $\tau_1 \leq \tau_2$ is decidable. You may NOT assume that there is a join function $join(\tau_1, \tau_2)$. You must complete your solution in the most general way you can without using such a relation.

You do NOT have to implement this question! Just define the rules.