

Proving Type Safety

Kim B. Bruce
Pomona College College

October 16, 2008

1 Introduction

The purpose of this brief write-up is to show how one can formally prove (at least partially) the correctness of a type system for a language with respect to its formal semantics. The basic idea is that what the type checker tells you statically should correspond to the value you obtain by actually evaluating the term. In particular, if the type checker tells you that the type of a term is \mathbb{T} , then the value obtained by evaluating that term should also be of type \mathbb{T} . This is usually called a type preservation theorem.

There is a second theorem, usually termed “Progress”, that is generally needed to prove type safety. This theorem states that well typed terms never get “stuck”. That is, if e is a term with type \mathbb{T} , then either it computes to a final legitimate value or it runs forever. More precisely, the computation never gets to a point where no rule is applicable.

It turns out that “progress” theorems are difficult to prove using the kind of natural semantics that we have used here. As a result, we will not attempt to prove that second theorem. Nevertheless we hope it is intuitively clear that the rules that we have written down are sufficient to avoid getting “stuck”.

2 Terms of Typed PCF

We begin by defining a typed version of PCF. The syntax differs from that of untyped PCF by adding a type annotation to bound variables in function and recursive definitions: $\text{fn } (x:\mathbb{T}) \Rightarrow \text{body}$ and $\text{rec } (r:\mathbb{T}) \Rightarrow \text{body}$. The grammar for types is as follows:

$$\mathbb{T} ::= \text{Int} \mid \text{Bool} \mid \mathbb{T} \rightarrow \mathbb{T}$$

Expressions are defined similarly to those in PCF:

$$e ::= x \mid n \mid \text{true} \mid \text{false} \mid \text{succ} \mid \text{pred} \mid \text{iszero} \mid \\ \text{if } e \text{ then } e \text{ else } e \mid (\text{fn } (x:\mathbb{T}) \Rightarrow e) \mid (e \ e) \mid \text{rec } (x:\mathbb{T}) \Rightarrow e$$

In this grammar, n stands for any non-negative integer, x is an identifier, and \mathbb{T} is a type.

3 Type-checking rules

We can specify type-checking rules for this language in a way similar to that for arithmetic expressions. However, type checking will be done with respect to a “type environment,” E , which is an association of strings to types (similar to the regular environment in the environment-based interpreter that was a mapping from strings to values). Thus if $E(x) = \text{Int}$, then identifier x has type Int in E . (See slides 6 and 7 of Lecture 13 for similar rules.)

$$\begin{array}{l}
E \vdash n: \text{Int}, \quad \text{if } n \text{ is an integer} \\
E \vdash \text{true}: \text{Bool}, \quad E \vdash \text{false}: \text{Bool} \\
E \vdash \text{succ}: \text{Int} \rightarrow \text{Int}, \quad E \vdash \text{pred}: \text{Int} \rightarrow \text{Int} \\
E \vdash \text{iszero}: \text{Int} \rightarrow \text{Bool} \\
E \vdash x: T, \quad \text{if } E(x) = T \\
\frac{E \vdash e: \text{Bool} \quad E \vdash e1: T \quad E \vdash e2: T}{E \vdash \text{if } e \text{ then } e1 \text{ else } e2 : T} \\
\frac{E \vdash f: T \rightarrow U \quad E \vdash x: T}{E \vdash (f \ x): U} \\
\frac{E[x:T] \vdash \text{body}: U}{E \vdash (\text{fn } (x:T) \Rightarrow \text{body}): T \rightarrow U}
\end{array}$$

In the last two rules, $E[x:T]$ stands for the type environment obtained by adding to E the binding of identifier x to have type T . In other words, in order to determine the type of $\text{fn } (x:T) \Rightarrow \text{body}$ in the presence of E , determine the type, U , of body in the presence of the updated environment, $E[x:T]$. The type of the function is then $T \rightarrow U$.

Aside from the `if` and `fn` clauses, the other type-checking rules are straightforward. When type-checking `AST_IF`, the type of the result should be the same as for each clause (we must be sure they are the same!).

It is easy to see (formally, by induction on the proof of typing) that if $E \vdash e: T$ then all free variables of e are contained in the domain of E .

4 Natural semantics

The semantic rules for typed PCF are exactly the same as for untyped PCF as the type annotations are simply ignored. Recall from our discussion of the environment-based semantics of PCF, the only non-error *values* obtained by evaluating PCF terms are integers, the boolean values *true* and *false*, the interpretations of primitive functions *succ*, *pred*, and *iszero*, and closures, which we write in the form $\langle \text{fn } x \Rightarrow e3, \text{env} \rangle$.

The function *env* is an assignment of values to identifiers. This makes it quite different from the E above which assigns types to identifiers. Anything not matching a rule returns *error*.

The following axioms and rules define the semantics of type PCF. For $(e, \text{env}) \Rightarrow v$ to be defined, we require that $\text{domain}(\text{env})$ contain all free variables of e .

id $(\text{id}, \text{env}) \Rightarrow \text{env}(\text{id})$

int $(n, \text{env}) \Rightarrow n, \text{ for } n \text{ an integer}$

true $(\text{true}, \text{env}) \Rightarrow \text{true}$

$$false \quad (false, env) \Rightarrow false$$

$$succ \quad (succ, env) \Rightarrow succ$$

$$pred \quad (pred, env) \Rightarrow pred$$

$$iszero \quad (iszero, env) \Rightarrow iszero$$

$$cond1 \quad \frac{(b, env) \Rightarrow true \quad (e1, env) \Rightarrow v}{(if\ b\ then\ e1\ else\ e2, env) \Rightarrow v}$$

$$cond2 \quad \frac{(b, env) \Rightarrow false \quad (e2, env) \Rightarrow v}{(if\ b\ then\ e1\ else\ e2, env) \Rightarrow v}$$

$$succAp \quad \frac{(e1, env) \Rightarrow succ \quad (e2, env) \Rightarrow n}{((e1\ e2), env) \Rightarrow (n+1)}$$

$$predAp1 \quad \frac{(e1, env) \Rightarrow pred \quad (e2, env) \Rightarrow 0}{((e1\ e2), env) \Rightarrow 0}$$

$$predAp2 \quad \frac{(e1, env) \Rightarrow pred \quad (e2, env) \Rightarrow n}{((e1\ e2), env) \Rightarrow (n-1)}$$

$$iszero1 \quad \frac{(e1, env) \Rightarrow iszero \quad e2, env) \Rightarrow 0}{((e1\ e2), env) \Rightarrow true}$$

$$iszero2 \quad \frac{(e1, env) \Rightarrow iszero \quad (e2, env) \Rightarrow n+1}{((e1\ e2), env) \Rightarrow (n-1)}$$

$$userFunc \quad ((fn\ (x:T)\ =>\ e), env) \Rightarrow \langle fn\ (x:T)\ =>\ e, env \rangle$$

$$userAp \quad \frac{(e1, env) \Rightarrow \langle fn\ (x:T)\ =>\ e3, env' \rangle \quad (e2, env) \Rightarrow v2 \quad (e3, env'[v2/x]) \Rightarrow v}{((e1\ e2), env) \Rightarrow v}$$

By our restrictions on pairs (e, env) to be involved in reductions, the closure, $\langle fn\ x\ =>\ e, env \rangle$ resulting from the applicatin of *userFunc* has the property that all free variables of $fn\ x\ =>\ e$ are contained in the domain of *env*.

5 Typing Rules for Values

Because we wish to show that types are preserved under computation, we will need to calculate the types of the values obtained from the semantics rules. It is important to note that these values have no free variables. That is obviously correct for numbers and the values `true` and `false`. Closures, which include the syntax for the function definitions are the only ones that might be able to have free variables. However, the environment paired with the closure is required to have interpretations of all of the free variables in the function definition, so all free variables in the function body are bound in the environment.

The first rules are trivial:

$$\mid -_v n : \text{Int}, \text{ for } n \text{ an integer}$$

$$\mid -_v \text{true} : \text{Bool}$$

$$\mid -_v \text{false} : \text{Bool}$$

$$\mid -_v \text{succ} : \text{Int} \rightarrow \text{Int}$$

$$\mid -_v \text{pred} : \text{Int} \rightarrow \text{Int}$$

$$\mid -_v \text{iszero} : \text{Int} \rightarrow \text{Bool}$$

In order to state the rule for closures we require the definition of when a regular environment is compatible with a type environment.

Definition: We say that a typing environment E and a regular environment env are *compatible* iff $\text{domain}(E) = \text{domain}(env)$, and for all x in $\text{domain}(env)$, $\mid -_v env(x) : E(x)$.

That is, env is compatible with E iff they have the same domain, and the value for an identifier obtained from env has the same type given that identifier by E .

$$\frac{E' \mid - \text{fn } (x:T) \Rightarrow e : T \rightarrow U, \text{ for the } E' \text{ such that } E' \text{ and } env \text{ are compatible}}{\mid -_v \langle \text{fn } (x:T) \Rightarrow e, env \rangle : T \rightarrow U}$$

While this may look a bit confusing, the intuition should be clear. To type a closure, just create a type environment that is compatible with the environment env in the closure and type check with respect to that type environment to get the type of the closure.

6 Type safety

We expect that the type checking rules and semantic rules should have some connection. In particular, if we can prove that a term has a particular type, then we expect that if we evaluate the term then its value should have the same type predicted by the type checker. After all, if that fails to be the case then what is the purpose of the type system!

Theorem 1 (Subject Reduction) *Let E be a type environment assigning types to identifiers and let env be a run-time environment assigning values to identifiers such that E and env are compatible.*

Let e be a term of typed PCF. If $E \mid - e : T$ and $(e, env) \Rightarrow v$, then $\mid -_v v : T$.

Proof The proof is by induction on the derivation of the typing of e . I.e. we are allowed to assume the theorem holds for all e' such that the proof of $E \vdash e: T$ involves an earlier typing of e' .

Identifiers: Let e be an identifier, id , and suppose $E \vdash \text{id}: T$ because $E(x) = T$. By assumption, $\emptyset \vdash \text{env}(x): T$, and we are done.

Numbers: Let e be a number, n . Thus $E \vdash n: \text{Int}$. By the definition of semantics, $(n, \text{env}) \Rightarrow n$ and $\vdash_v n: \text{Int}$, as desired.

true, false, succ, pred, iszero: The proof is similar to that for numbers as they simply evaluate to themselves and contain no identifiers.

Conditional: Suppose $E \vdash \text{if } b \text{ then } e_1 \text{ else } e_2: T$ because $E \vdash b: \text{boolean}$, $E \vdash e_1: T$, and $E \vdash e_2: T$. There are two cases depending on the evaluation of b .

- Suppose that $(b, \text{env}) \Rightarrow \text{true}$. Then if $(e_1, \text{env}) \Rightarrow v$, it follows that $(\text{if } b \text{ then } e_1 \text{ else } e_2, \text{env}) \Rightarrow v$.

By induction and $E \vdash e_1: T$, it follows that $\vdash_v v: T$, which is all we need.

- The case for $(b, \text{env}) \Rightarrow \text{false}$ is similar.

Applications: Suppose that e is $(e_1 \ e_2)$ and $E \vdash (e_1 \ e_2): U$ because $E \vdash e_1: T \rightarrow U$ and $E \vdash e_2: T$. Suppose also that $(e_1, \text{env}) \Rightarrow f$ and $(e_2, \text{env}) \Rightarrow v_2$. By induction, $\vdash_v f: T \rightarrow U$ and $\vdash_v v_2: T$.

The rest of the proof of that the final value has type U depends on the actual value of f .

- f is *succ*. Then $T = U = \text{Int}$. Because $\vdash_v v: \text{Int}$, v must be an integer. By the semantics rule *succAp*, the result is $v+1$, another integer, which has type Int , as desired as $E \vdash (e_1 \ e_2): \text{Int}$.
- f is *pred*. Like the case for *succ*.
- f is *iszero*. Like the case for *succ*.
- f is a closure, $\langle \text{fn } (x:T) \Rightarrow e_3, \text{env}' \rangle$. By induction, $\vdash_v \langle \text{fn } (x:T) \Rightarrow e_3, \text{env}' \rangle: T \rightarrow U$ (note that the T in the closure must match the T in the typing by the typing rule for closures). Because there is only one way to type a closure, there must be an E' such that E' and env' are compatible and $E' \vdash \text{fn } (x:T) \Rightarrow e_3: T \rightarrow U$.

To complete the evaluation of $(e_1 \ e_2)$, via rule *userAp*, suppose $(e_3, \text{env}'[v_2/x]) \Rightarrow v$. We must show that $\vdash_v v: U$.

Because $E' \vdash \text{fn } (x:T) \Rightarrow e_3: T \rightarrow U$, it must have been the case that $E' [x:T] \vdash e_3: U$. However, because E' and env' are compatible, so are $E' [x:T]$ and $\text{env}'[v_2/x]$ (recall that $\vdash_v v_2: T$). Therefore by the induction hypothesis and $(e_3, \text{env}'[v_2/x]) \Rightarrow v$, it follows that $\vdash_v v: U$, as desired.

■

Hopefully this gives you a sense as to how these type-safety results work. Again, we normally also need to prove a progress theorem that shows one never gets to a state where no rule applies. To do this with a natural semantics we must fill in lots of error terms to make sure there is always something to do – and make sure that error terms are never in any of the types that we care about. The most popular alternative operational semantics, structured operational semantics, makes that easier. However, we tend to prefer this sort of “big-step semantics” to the SOS “small-step semantics”.