

1 Introduction

Type checking is a lightweight technique for proving simple properties of programs. Unlike theorem-proving techniques based on axiomatic semantics, type checking usually cannot determine if a program will produce the correct output. Instead, it is a way to test whether a program is *well-formed*, with the idea that a well-formed program satisfies certain desirable properties. The traditional application of type checking is to show that a well-formed program cannot get stuck; that is, that a type-correct program will never reach a non-final configuration in its operational semantics from which its behavior is undefined. This is a weak notion of program correctness, but nevertheless very useful in practice for catching bugs.

Type systems are a powerful technique. In the past couple of decades, researchers have discovered how to use type systems for a variety of different verification tasks.

2 λ^{\rightarrow}

We have already seen some typed languages in class this semester. For example, OCaml and the metalanguage used in class for denotational semantics are both typed.

To explore the idea of type checking itself, we introduce λ^{\rightarrow} , a typed variant of the λ -calculus in which we assign types to certain λ -terms according to some typing rules. A λ -term is considered to be well-formed if a type can be derived for it using the rules. We will give operational and denotational semantics for this language. Along the way, we will discover some interesting properties that give insight about typing in more complex languages.

3 Syntax

The syntax of λ^{\rightarrow} is similar to that of untyped λ -calculus, with some notable differences. There are two kinds of inductively-defined expressions, *terms* and *types*:

$$\begin{array}{l} \text{terms } e ::= n \mid \text{true} \mid \text{false} \mid \text{null} \mid x \mid e_1 e_2 \mid \lambda x : \tau. e \\ \text{types } \tau ::= \text{int} \mid \text{bool} \mid \text{unit} \mid \tau_1 \rightarrow \tau_2 \end{array}$$

One difference is that the natural numbers, Boolean constants, and `null` are taken to be primitive symbols and not defined as λ -terms. For this reason, we no longer need to distinguish between the syntactic objects `\bar{n}` and `true` and their semantic counterparts n and *true*, so we might as well just identify them. Another difference is that a λ -abstraction explicitly mentions the type of its argument. Thus $\lambda x : \tau. e$ represents a function that takes an input value of type τ and evaluates e .

A *value* is either a number, a Boolean constant, `null`, or a closed abstraction $\lambda x : \tau. e$. The set of values is denoted *Val*. The set of types is denoted *Type*.

A *type* represents a collection of related values. The types `int`, `bool`, and `unit` represent integers, Booleans, and `null`, respectively. The type $\tau_1 \rightarrow \tau_2$ represents functions that take inputs of type τ_1 and produce outputs of type τ_2 .

There are some *typing rules*, given below, that can be used to associate a type with a term. The rules are used to derive *type judgments* of the form $\Gamma \vdash e : \tau$, which means intuitively that under the assumptions Γ ,

the term e has type τ . The assumptions Γ specify types for the free variables of e . We write $\vdash e : \tau$ when e is a closed term and Γ is empty.

For example, every number has type `int`, thus $\vdash 3 : \text{int}$. The type `unit` is the type of the value `null`, and nothing else has this type. The function `trueint` = $\lambda x : \text{int}. \lambda y : \text{int}. x$ has the type `int` \rightarrow `int` \rightarrow `int`. By convention, the type constructor \rightarrow associates to the right, so `int` \rightarrow `int` \rightarrow `int` is the same as `int` \rightarrow (`int` \rightarrow `int`). Thus we can write

$$\text{true}_{\text{int}} \triangleq (\lambda x : \text{int}. \lambda y : \text{int}. x) : \text{int} \rightarrow \text{int} \rightarrow \text{int}.$$

Not all λ -terms can be typed, for instance $\lambda x : \text{int}. xx$ or `true 3`. These expressions are considered ill-formed and nonsensical.

Right now, we cannot do anything interesting with integers or Booleans, because we do not have any operations on them. Later on we will be adding other typed constants such as `plus` : `int` \rightarrow `int` \rightarrow `int` and `equal` : `int` \rightarrow `int` \rightarrow `bool`, but for now they are just there to set the stage.

4 Small-Step Operational Semantics and Type Correctness

The small-step CBV operational semantics of λ^{\rightarrow} is the same as that of the untyped λ -calculus. The presence of types does not alter the evaluation rules for expressions.

$$E ::= Ee \mid vE \mid [\cdot] \qquad (\lambda x : \tau. e)v \rightarrow e\{v/x\}$$

We will eventually show that these reduction rules preserve typing in the sense that if e has type τ and ee' , then e' also has type τ . Thus “well-typedness” of programs is invariant under the evaluation rules of the language. This property is known as *type preservation* or *subject reduction*. Another important property is *progress*, which says that a well-typed program is never stuck; that is, it is either a value or a further transition is possible. These two properties together imply that no well-typed term can ever become stuck. Thus the typing rules can be used in place of runtime type checking to ensure strong typing.

It is natural to ask what a type-incorrect term might look like and how it could get stuck. Recall our function definition for `trueint` above, and consider the following additional definition:

$$\text{if}_{\text{int}} \triangleq \lambda t : \text{int} \rightarrow \text{int} \rightarrow \text{int}. \lambda a : \text{int}. \lambda b : \text{int}. tab.$$

Clearly, `ifint trueint 2 3` evaluates to 2. However, `ifint true 2 3` \rightarrow `true 2 3`, and this expression is meaningless, since `true` is not a function and cannot be applied to anything. Therefore, the program would be stuck at this point.

5 Typing Rules

The typing rules will determine which terms are well-formed λ^{\rightarrow} programs. They are a set of rules that allow the derivation of *type judgments* of the form $\Gamma \vdash e : \tau$. Here Γ is a *type environment*, a partial map from variables to types used to determine the types of the free variables in e . The domain of Γ as a partial function $\text{Var} \rightarrow \text{Type}$ is denoted $\text{dom } \Gamma$.

The environment $\Gamma[\tau/x]$ is obtained by rebinding x to τ (or creating the binding anew if $x \notin \text{dom } \Gamma$):

$$\Gamma[\tau/x](y) \triangleq \begin{cases} \Gamma(y), & \text{if } y \neq x \text{ and } y \in \text{dom } \Gamma, \\ \tau, & \text{if } y = x, \\ \text{undefined}, & \text{otherwise.} \end{cases}$$

The notation $\Gamma, x : \tau$ is synonymous with $\Gamma[\tau/x]$. The former is standard notation in the literature. Also, one often sees $x : \tau \in \Gamma$, which means just $\Gamma(x) = \tau$.

We also write $\Gamma \vdash e : \tau$ as a metaexpression to mean that the type judgment $\Gamma \vdash e : \tau$ is derivable from the typing rules. The environment \emptyset is the empty environment, and the judgment $\vdash e : \tau$ is short for $\emptyset \vdash e : \tau$.

Here are the typing rules:

$$\Gamma \vdash n : \text{int} \quad \Gamma \vdash \text{true} : \text{bool} \quad \Gamma \vdash \text{false} : \text{bool} \quad \Gamma \vdash \text{null} : \text{unit} \quad \Gamma, x : \tau \vdash x : \tau$$

$$\frac{\Gamma \vdash e_0 e_1 : \tau}{\Gamma \vdash e_0 : \sigma \rightarrow \tau \Gamma \vdash e_1 : \sigma} \quad \frac{\Gamma \vdash (\lambda x : \sigma. e) : \sigma \rightarrow \tau}{\Gamma, x : \sigma \vdash e : \tau}$$

Let us explain these rules in more detail.

- The first four rules just say that all the base values have their corresponding base types.
- For a variable x , $\Gamma \vdash x : \tau$ holds if the binding $x : \tau$ appears in the type environment Γ ; that is, if $\Gamma(x) = \tau$.
- An application expression $e_0 e_1$ represents the result of applying the function represented by e_0 to the argument represented by e_1 . For this to have type τ , e_0 must be a function of type $\sigma \rightarrow \tau$ for some σ , and its argument e_1 must have type σ . This is captured in the typing rule for $e_0 e_1$.
- Finally, a λ -abstraction $\lambda x : \sigma. e$ is supposed to represent a function. The type of the input should match the annotation in the term, thus the type of the function must be $\sigma \rightarrow \tau$ for some τ . The type τ of the result is the type of the body under the extra type assumption $x : \sigma$. This idea is captured in the typing rule for λ -abstractions.

Every well-typed λ^{\rightarrow} term has a proof tree consisting of applications of the typing rules to derive a type for the term. We can type-check a term by constructing this proof tree. For example, consider the program $(\lambda x : \text{int}. \lambda y : \text{bool}. x) 2 \text{true}$, which evaluates to 2. Since $\vdash 2 : \text{int}$, we expect $\vdash ((\lambda x : \text{int}. \lambda y : \text{bool}. x) 2 \text{true}) : \text{int}$ as well. Here is a proof of that fact:

$$\frac{\frac{\frac{\vdash ((\lambda x : \text{int}. \lambda y : \text{bool}. x) 2 \text{true}) : \text{int}}{\vdash ((\lambda x : \text{int}. \lambda y : \text{bool}. x) 2) : \text{bool} \rightarrow \text{int}} \vdash \text{true} : \text{bool}}{\vdash (\lambda x : \text{int}. \lambda y : \text{bool}. x) : \text{int} \rightarrow \text{bool} \rightarrow \text{int}} \vdash 2 : \text{int}}{\frac{x : \text{int} \vdash (\lambda y : \text{bool}. x) : \text{bool} \rightarrow \text{int}}{x : \text{int}, y : \text{bool} \vdash x : \text{int}}}$$

An automated type checker can effectively construct proof trees like this in order to test whether a program is type-correct.

Note that types, if they exist, are unique. That is, if $\Gamma \vdash e : \tau$ and $\Gamma \vdash e : \tau'$, then $\tau = \tau'$. This can be proved easily by structural induction on e , using the fact that there is exactly one typing rule that applies in each case, depending on the form of e .

6 Expressive Power

By now you may be wondering if we have lost any of the expressive power of λ -calculus by introducing types. The answer to this question is a resounding *yes*. For example, we can no longer compose arbitrary functions, since they may have mismatched types.

More importantly, we have lost the ability to write loops. Recall the paradoxical combinator

$$\Omega \triangleq (\lambda x. xx)(\lambda x. xx).$$

Let us show that any attempt to derive a typing for the term $\lambda x : \sigma. xx$ must fail:

$$\frac{\frac{\Gamma \vdash (\lambda x : \sigma. xx) : \sigma \rightarrow \tau}{\Gamma, x : \sigma \vdash xx : \tau}}{\Gamma, x : \sigma \vdash x : \sigma \rightarrow \tau, x : \sigma \vdash x : \sigma}$$

We see that we must have both $\Gamma, x : \sigma \vdash x : \sigma \rightarrow \tau$ and $\Gamma, x : \sigma \vdash x : \sigma$. However, since types are unique, this is impossible; we cannot have $\sigma = \sigma \rightarrow \tau$, since no type expression can be a subexpression of itself. We conclude that the term $\lambda x : \sigma. xx$ cannot be typed.

In fact, we will see later that we cannot write down *any* nonterminating program in λ^\rightarrow . This will turn out be true from an operational perspective as well. In later lectures we will show how to extend the type system to allow loops and nonterminating programs.

7 Denotational Semantics

Before we can give the denotational semantics for λ^\rightarrow , we need to define a new meaning function $\mathcal{T}[\cdot]$ that maps each type to a domain associated with that type. For this type system, the definition of $\mathcal{T}[\cdot]$ is straightforward:

$$\mathcal{T}[\text{int}] \triangleq \mathbb{Z} \quad \mathcal{T}[\text{bool}] \triangleq \mathbb{2} \quad \mathcal{T}[\text{unit}] \triangleq \{\text{null}\} \quad \mathcal{T}[\sigma \rightarrow \tau] \triangleq \mathcal{T}[\sigma] \rightarrow \mathcal{T}[\tau].$$

In the last equation, note that the \rightarrow on the left-hand side is just a symbol in the language of types, a syntactic object, whereas the \rightarrow on the right-hand side is a semantic object, namely the operator that constructs the space of functions between a given domain and range. For now, these domains need not have any ordering properties; they are just sets. So we have $\mathcal{T}[\cdot] : \text{Type} \rightarrow \text{Set}$.

For any closed term e , if $\vdash e : \tau$, then we expect the denotation of e to be an element of $\mathcal{T}[\tau]$. More generally, for a term e possibly containing free variables, if there is a type environment Γ and a value environment ρ such that ρ and Γ are defined on all the free variables of e and $\rho(x) \in \mathcal{T}[\Gamma(x)]$ for all $x \in FV(e)$, and if $\Gamma \vdash e : \tau$, then we expect the denotation of e in environment ρ to be an element of $\mathcal{T}[\tau]$.

We say that the value environment ρ *satisfies* a type environment Γ and write $\rho \models \Gamma$ if $\text{dom } \Gamma \subseteq \text{dom } \rho$ and for every $x \in \text{dom } \Gamma$, $\rho(x) \in \mathcal{T}[\Gamma(x)]$.

We are now ready to give the denotational semantics for typed λ -terms. We only define the meaning function for well-typed expressions. The following function is defined only for e , Γ , and ρ such that $\rho \models \Gamma$ and e is well-typed under Γ ; that is, $\Gamma \vdash e : \tau$ for some type τ .

$$\begin{aligned} \mathcal{C}[n]\Gamma\rho &\triangleq n \\ \mathcal{C}[\text{true}]\Gamma\rho &\triangleq \text{true} \\ \mathcal{C}[\text{false}]\Gamma\rho &\triangleq \text{false} \\ \mathcal{C}[\text{null}]\Gamma\rho &\triangleq \text{null} \\ \mathcal{C}[x]\Gamma\rho &\triangleq \rho(x), \quad x \in \text{dom } \Gamma \\ \mathcal{C}[e_0 e_1]\Gamma\rho &\triangleq (\mathcal{C}[e_0]\Gamma\rho) (\mathcal{C}[e_1]\Gamma\rho) \\ \mathcal{C}[\lambda x : \tau. e]\Gamma\rho &\triangleq \lambda v \in \mathcal{T}[\tau]. \mathcal{C}[e]\Gamma[\tau/x]\rho[v/x] \end{aligned}$$

8 Soundness

It is possible to show that $\mathcal{C}[\cdot]$ satisfies the following soundness condition: for all ρ , Γ , e , and τ ,

$$\rho \models \Gamma \wedge \Gamma \vdash e : \tau \Rightarrow \mathcal{C}[e]\Gamma\rho \in \mathcal{T}[\tau].$$

The proof is by induction on the structure of e . The only interesting step is the case of a λ -term, which requires the observation

$$\rho \models \Gamma \wedge v \in \mathcal{T}[\tau] \Rightarrow \rho[v/x] \models \Gamma[\tau/x].$$

Note that \perp does not appear anywhere in this semantics. We did not need it because we never took a fixpoint. We can conclude that all well-typed λ -terms represent total functions.

Of course, we will also want to see that the operational semantics is adequate with respect to this model to ensure that our evaluation relation actually produces the values of the correct type that the denotational model advertises are there. For example, for a closed term e such that $\vdash e : \text{int}$, the adequacy condition would say that

$$en \Leftrightarrow \mathcal{C}[e]\emptyset\emptyset = n.$$