

1 Introduction

Type checking as seen earlier in this class is pessimistic by nature: if we cannot prove a program will never run into the problem we are checking for (e.g. getting stuck), it is not well-formed, and will hence be rejected by a type checker. This can be quite onerous, especially in cases where the type system at some point forces one to “forget” information that it will later need (e.g. storing a number in a place that accepts arbitrary values – in many cases, we know that whatever value we retrieve from there is a number, but the type checker does not). *Gradual typing* is the idea that sometimes we might tell the type checker to be optimistic instead, in which case it should only reject a program if it can prove that a program *will* run into the problem we are checking for (e.g. the program will get stuck).

2 $\lambda_{\text{?}}^{\rightarrow}$

We’ll start by adding a special type ? to the simply typed lambda calculus λ^{\rightarrow} .

terms $e ::= n \mid \text{true} \mid \text{false} \mid \text{null} \mid x \mid e_1 e_2 \mid \lambda x : \tau. e$
 types $\tau, \sigma ::= \text{int} \mid \text{bool} \mid \text{unit} \mid \tau_1 \rightarrow \tau_2 \mid \text{?}$

Otherwise, the syntax of $\lambda_{\text{?}}^{\rightarrow}$ is exactly the same as that of λ^{\rightarrow} . The typing rules are also similar, only the rule for function application changes somewhat. It has three new components that we need to introduce. The first is the consistency relation \sim . It is defined as follows:

$$\tau \sim \tau \quad \tau \sim ? \quad ? \sim \tau \quad \frac{\tau \sim \tau' \quad \sigma \sim \sigma'}{\tau \rightarrow \sigma \sim \tau' \rightarrow \sigma'}$$

The consistency relation represents a relaxed version of type equality: two types are consistent with each other if they *could* be equal if the occurrences of ? in them are replaced appropriately. This is the central ingredient to making the type checker optimistic while still being able to warn us if two types do not match up at all. Note that \sim is reflexive and symmetric, but not transitive – if it were, everything would be consistent with everything. The other two components are helper functions mapping function types to their domain and range, lifted to gradual types:

$$\text{dom}_{\text{?}}(\tau) \triangleq \begin{cases} \sigma, & \text{if } \tau = \sigma \rightarrow \sigma', \\ ?, & \text{if } \tau = ?, \\ \text{undefined}, & \text{otherwise.} \end{cases} \quad \text{rng}_{\text{?}}(\tau) \triangleq \begin{cases} \sigma', & \text{if } \tau = \sigma \rightarrow \sigma', \\ ?, & \text{if } \tau = ?, \\ \text{undefined}, & \text{otherwise.} \end{cases}$$

Now, we can define the typing rules:

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

$$\frac{\Gamma \vdash_{\text{?}} e_0 : \tau \quad \Gamma \vdash_{\text{?}} e_1 : \sigma \quad \sigma \sim \text{dom}_{\text{?}}(\tau)}{\Gamma \vdash_{\text{?}} e_0 e_1 : \text{rng}_{\text{?}}(\tau)} \quad \frac{\Gamma, x : \sigma \vdash_{\text{?}} e : \tau}{\Gamma \vdash_{\text{?}} (\lambda x : \sigma. e) : \sigma \rightarrow \tau}$$

There are two big changes for function application. The first is that the type of the function itself might be $?$, so we cannot immediately match it to a function type constructed with \rightarrow . If the type of the function is indeed $?$, we will still optimistically assume that it is a function, and instruct the type checker to again be optimistic when it comes to its domain and range, which is what the helper functions **dom** and **rng** express. The second change is that the type of the argument does not have to be equal to the domain of the function; the two only have to be consistent. Thus, a $?$ argument is accepted by any function, and a function that accepts a $?$ accepts any argument.

3 Relationship to the λ -calculus and λ^{\rightarrow}

With the above rules, we can prove that for all expressions e^{\rightarrow} of λ^{\rightarrow} , $\Gamma \vdash e : \tau \Leftrightarrow \Gamma \vdash_{?} e^{\rightarrow} : \tau$. That is, fully typed programs that do not contain $?$ are well-formed in $\lambda_{?}^{\rightarrow}$ iff they are well-formed in λ^{\rightarrow} . Similarly, for any expression e of the (untyped) λ -calculus, let $e^{?}$ be an expression in $\lambda_{?}^{\rightarrow}$ that is obtained by the following translation:

$$\begin{aligned} \llbracket n \rrbracket &\triangleq n \\ \llbracket \text{true} \rrbracket &\triangleq \text{true} \\ \llbracket \text{false} \rrbracket &\triangleq \text{false} \\ \llbracket \text{null} \rrbracket &\triangleq \text{null} \\ \llbracket \lambda x. e \rrbracket &\triangleq \lambda x : ?. \llbracket e \rrbracket \\ \llbracket e_1 e_2 \rrbracket &\triangleq ((\lambda x : ?. x) \llbracket e_1 \rrbracket) \llbracket e_2 \rrbracket \end{aligned}$$

Every such $e^{?}$ is well-formed in $\lambda_{?}^{\rightarrow}$.

4 Semantics and Soundness

While the previous section suggests that for fully typed or fully untyped programs, the small-step operational semantics for the λ -calculus and λ^{\rightarrow} would work just fine, the resulting language would have neither progress nor preservation. The reason for that is that the optimism we introduced may in some cases be unfounded – a program that *could* be safe does not always have to be safe. In some sense, we accepted that when we introduced optimism in the first place, but as it stands, we have no guarantees about how a program that is neither fully typed or fully untyped might behave. In order to get back some guarantees about the behavior of mixed programs, we need to add some run-time checks to a gradually typed program, creating an intermediate language called $\lambda_{\langle \tau \rangle}^{\rightarrow}$:

terms	$e ::= n \mid \text{true} \mid \text{false} \mid \text{null} \mid x \mid e_1 e_2 \mid \lambda x : \tau. e \mid \langle \tau \Leftarrow \sigma \rangle e \mid \mathbf{box}^{\gamma} v$
values	$v ::= n \mid \text{true} \mid \text{false} \mid \text{null} \mid x \mid \lambda x : \tau. e \mid \langle \tau \rightarrow \sigma \Leftarrow \tau' \rightarrow \sigma' \rangle v \mid \mathbf{box}^{\gamma} v$
types	$\tau, \sigma ::= \text{int} \mid \text{bool} \mid \text{unit} \mid \tau_1 \rightarrow \tau_2 \mid ?$
ground types	$\gamma ::= \text{int} \mid \text{bool} \mid \text{unit} \mid ? \rightarrow ?$

This language adds two new constructs. The first is the cast expression $\langle \tau \Leftarrow \sigma \rangle e$, where the expression e has type σ and is cast to type τ . The second new construct is a boxed value container for when a value is treated as $?$. The box remembers the original ground type γ of the value (functions are always boxed as $? \rightarrow ?$, and $?$ itself is never boxed). Note that we treat casts of function types as values – this is both for ease of notation and to highlight an important point later. With these constructs, we can give $\lambda_{\langle \tau \rangle}^{\rightarrow}$ pessimistic typing rules

that essentially treat $?$ as the type of boxed values and allow casts between types that are consistent with each other:

$$\begin{array}{c}
\Gamma \vdash n : \text{int} \quad \Gamma \vdash \text{true} : \text{bool} \quad \Gamma \vdash \text{false} : \text{bool} \quad \Gamma \vdash \text{null} : \text{unit} \quad \Gamma, x : \tau \vdash x : \tau \\
\\
\frac{\Gamma \vdash e_0 : \tau \rightarrow \sigma \quad \Gamma \vdash e_1 : \tau}{\Gamma \vdash e_0 e_1 : \sigma} \quad \frac{\Gamma, x : \sigma \vdash e : \tau}{\Gamma \vdash (\lambda x : \sigma. e) : \sigma \rightarrow \tau} \quad \boxed{\frac{\Gamma \vdash e : \sigma \quad \tau \sim \sigma}{\Gamma \vdash \langle \tau \Leftarrow \sigma \rangle e : \tau}} \quad \boxed{\frac{\Gamma \vdash v : \gamma}{\Gamma \vdash \text{box}^\gamma v : ?}}
\end{array}$$

This language can now have a small-step CBV operational semantics that we can prove progress and preservation for (with the caveat that there can be run-time errors **error** representing a failed cast). First, evaluation contexts are extended so we can evaluate under casts and propagate errors:

$$\begin{array}{c}
E ::= E e \mid v E \mid \boxed{\langle \tau \Leftarrow \sigma \rangle E} \mid [\cdot] \\
\\
\frac{e_0 \rightarrow e_1}{E[e_0] \rightarrow E[e_1]} \quad \boxed{\frac{e \rightarrow \text{error}}{E[e] \rightarrow \text{error}}}
\end{array}$$

Then, a number of new reduction rules handle casting (first applicable rule applies exclusively):

$$\begin{array}{c}
(\lambda x : \tau. e) v \rightarrow e\{v/x\} \\
\boxed{\langle ? \Leftarrow ? \rangle v} \rightarrow v \\
\boxed{\langle \gamma \Leftarrow \gamma \rangle v} \rightarrow v \\
\boxed{\langle ? \Leftarrow \gamma \rangle v} \rightarrow \text{box}^\gamma v \\
\boxed{\langle \tau \Leftarrow ? \rangle \text{box}^\gamma v} \rightarrow \boxed{\langle \tau \Leftarrow \gamma \rangle v} \quad \text{if } \tau \sim \gamma \\
\boxed{\langle \tau \Leftarrow ? \rangle \text{box}^\gamma v} \rightarrow \text{error} \quad \text{if } \tau \not\sim \gamma \\
\boxed{\langle \langle \tau \rightarrow \sigma \Leftarrow \tau' \rightarrow \sigma' \rangle v_0 \rangle v_1} \rightarrow \boxed{\langle \sigma \Leftarrow \sigma' \rangle (v_0 (\langle \tau' \Leftarrow \tau \rangle v_1))} \\
\boxed{\langle ? \Leftarrow \tau \rightarrow \sigma \rangle v} \rightarrow \boxed{\text{box}^{? \rightarrow ?} \langle ? \rightarrow ? \Leftarrow \tau \rightarrow \sigma \rangle v}
\end{array}$$

The first two rules eliminate superfluous casts from a ground types and $?$ to themselves. The third rule boxes values of ground types. The fourth and fifth rule handle casts from $?$ – this is where run-time type errors are eventually caught. If the actual type γ of the boxed value is not consistent with type that the value is cast to, we raise an error. Such an inconsistency may happen in two cases¹: either γ is $? \rightarrow ?$ and τ is not a function type, or γ is not $? \rightarrow ?$ and τ is not equal to γ . The sixth rule deals with the fact that it is in general impossible to prove something about the output of an arbitrary function, even when given information about the input. Thus, we cannot decide immediately whether a cast from one function type to another succeeds. Instead, we wrap the function (which is why a function cast is a value) and insert casts every time the function is called: one cast to make sure that the input to the function is what the function expected, and one cast to make sure that its output is what the caller expected. The last rule expresses that when we wrap a function, we remember that it was a function that we wrapped, but leave the details of its type to an inner cast.

The above typing and reduction rules give us progress and preservation as follows:

¹ $\tau = ?$ is handled by the first new rule, above

Progress For all e , if $\vdash e : \tau$, either $e = v$ for some v , or $e \rightarrow e'$ for some e' , or $e \rightarrow \mathbf{error}$.

Preservation For all e , if $\Gamma \vdash e : \tau$ and $e \rightarrow e'$ for some e' , $\Gamma \vdash e' : \tau$

Compared to the progress lemma of λ^{\rightarrow} , we just added the case where we might reduce to an error, while preservation stays the same.

4.1 Translation

We can now define the semantics of λ^{\rightarrow} by translation to $\lambda_{\langle\tau\rangle}^{\rightarrow}$, inserting casts whenever the type-checker relies on optimism. The resulting type-directed translation rules are of the form $\Gamma \vdash_{\tau}^{\langle\tau\rangle} e \rightsquigarrow e' : \tau$, where an expression e in λ^{\rightarrow} is translated to an expression e' in $\lambda_{\langle\tau\rangle}^{\rightarrow}$, having type τ under environment Γ in both languages.

$$\begin{array}{c}
\Gamma \vdash_{\tau}^{\langle\tau\rangle} n \rightsquigarrow n : \text{int} \quad \Gamma \vdash_{\tau}^{\langle\tau\rangle} \text{true} \rightsquigarrow \text{true} : \text{bool} \quad \Gamma \vdash_{\tau}^{\langle\tau\rangle} \text{false} \rightsquigarrow \text{false} : \text{bool} \\
\\
\Gamma \vdash_{\tau}^{\langle\tau\rangle} \text{null} \rightsquigarrow \text{null} : \text{unit} \quad \Gamma, x : \tau \vdash_{\tau}^{\langle\tau\rangle} x \rightsquigarrow x : \tau \quad \frac{\Gamma, x : \sigma \vdash_{\tau}^{\langle\tau\rangle} e \rightsquigarrow e' : \tau}{\Gamma \vdash_{\tau}^{\langle\tau\rangle} (\lambda x : \sigma. e) \rightsquigarrow (\lambda x : \sigma. e') : \sigma \rightarrow \tau} \\
\\
\frac{\Gamma \vdash_{\tau}^{\langle\tau\rangle} e_0 \rightsquigarrow e'_0 : \tau \quad \Gamma \vdash_{\tau}^{\langle\tau\rangle} e_1 \rightsquigarrow e'_1 : \sigma \quad \sigma \sim \mathbf{dom}_{\tau}(\tau)}{\Gamma \vdash_{\tau}^{\langle\tau\rangle} e_0 e_1 \rightsquigarrow \langle \mathbf{dom}_{\tau}(\tau) \rightarrow \mathbf{rng}_{\tau}(\tau) \Leftarrow \tau \rangle e'_0 \langle \mathbf{dom}_{\tau}(\tau) \Leftarrow \sigma \rangle e'_1 : \mathbf{rng}_{\tau}(\tau)}
\end{array}$$

In this translation, almost all the work happens for applications (except that λ -abstractions translate their inner expressions recursively). The application rule inserts two casts to reflect the two optimistic assumptions the type checker might make: that the functor e_0 is a function and that the argument e_1 matches the expected argument type.

It follows from this translation that a well-formed program in λ^{\rightarrow} translates to a well-formed program in $\lambda_{\langle\tau\rangle}^{\rightarrow}$, thus progress and preservation apply to λ^{\rightarrow} under this translation semantics, giving us a basic notion of safety.