

Type Safety via Logical Relations

Andrew Tolmach

May 16, 2024

These notes are based heavily on Amal Ahmed’s lectures at the Oregon Programming Languages Summer School in 2015 (www.cs.uoregon.edu/research/summerschool/summer15/curriculum.html) and Lau Skorstengaard’s notes based on those lectures (*An Introduction to Logical Relations: Proving Program Properties Using Logical Relations*, <https://www.cs.uoregon.edu/research/summerschool/summer16/notes/AhmedLR.pdf>).

We can use the machinery of unary logical relations (or logical “predicates”) to give an alternative proof of type safety that doesn’t rely on the notions of progress and preservation. The basic idea is to define a predicate that captures the property of safe evaluation (i.e. evaluation without getting “stuck”) and show that every (syntactically) well-typed term has this property.

In these notes, we will show type safety for the simply-typed lambda calculus extended with booleans as in Pierce Ch. 9. The proof can be expanded to cover further extensions such as pairs; this is left as an exercise.

Recall that a term \mathfrak{t} is in *normal form*, written $NF(\mathfrak{t})$, iff there is no \mathfrak{t}' such that $\mathfrak{t} \rightarrow \mathfrak{t}'$. We write $VAL(\mathfrak{t})$ iff \mathfrak{t} is a (syntactic) value. Finally, we define

$$SAFE(\mathfrak{t}) \triangleq \forall \mathfrak{t}'. \mathfrak{t} \rightarrow^* \mathfrak{t}' \wedge NF(\mathfrak{t}') \implies VAL(\mathfrak{t}')$$

to capture terms whose evaluation will never lead to a stuck state.

Then we wish to prove

THEOREM (TYPE SAFETY): If $\vdash \mathfrak{t} : T$ then $SAFE(\mathfrak{t})$.

We’ll prove this by defining a type-indexed predicate on terms that holds when the term is *semantically well-typed* (in a sense defined below), and then showing that (A) syntactically well-typed terms are semantically well-typed, and (B) semantically well-typed terms are safe. This is completely analogous to the structure of the normalization proof in Ch. 12.

Naturally, the challenge is to come up with a workable definition of semantic well-typedness for which both (A) and (B) can be shown. Again, as with normalization, we will arrange things so that the proof of (B) is very easy; the effort will be in proving (A).

To characterize semantically well-typed terms, we give a *value interpretation* $\mathcal{V}[[T]]$ of each type T as a set of (closed) values that belong to that type, taking care only to include values that are “well-behaved” with respect to safety. In particular, we must make sure that functions in the value interpretation are not only themselves (trivially) values, but also don’t get stuck when the function is applied. (Similarly, if we extend our language to include values of other compound types, such as pairs, we must ensure that the sub-components of the value are well-behaved.) To do this, we

also need a more general notion of a *term interpretation* $\mathcal{T}[\mathbb{T}]$ of type \mathbb{T} as a set of terms that will safely evaluate to values in the value interpretation.

We give mutually recursive definitions of these predicates as follows:

$$\begin{aligned}\mathcal{V}[\mathbf{Bool}] &\triangleq \{\mathbf{true}, \mathbf{false}\} \\ \mathcal{V}[\mathbb{T}_1 \rightarrow \mathbb{T}_2] &\triangleq \{\lambda \mathbf{x}:\mathbb{T}_1. \mathbf{t}_2 \mid (FV(\mathbf{t}_2) - \{\mathbf{x}\}) = \emptyset \wedge \forall v \in \mathcal{V}[\mathbb{T}_1]. [\mathbf{x} \mapsto v] \mathbf{t}_2 \in \mathcal{T}[\mathbb{T}_2]\} \\ \mathcal{T}[\mathbb{T}] &\triangleq \{\mathbf{t} \mid FV(\mathbf{t}) = \emptyset \wedge \forall \mathbf{t}'. \mathbf{t} \rightarrow^* \mathbf{t}' \wedge NF(\mathbf{t}') \implies \mathbf{t}' \in \mathcal{V}[\mathbb{T}]\}\end{aligned}$$

Here $FV(\mathbf{t})$ is the set of free variables of \mathbf{t} . Observe that these predicates are structurally well-founded, because the recursive mentions are at smaller types.

Here are some simple but useful consequences of these definitions:

- $\mathbf{t} \in \mathcal{V}[\mathbb{T}] \implies VAL(\mathbf{t})$.
- $\mathbf{t} \in \mathcal{V}[\mathbb{T}] \implies \mathbf{t} \in \mathcal{T}[\mathbb{T}]$, i.e., $\mathcal{V}[\mathbb{T}] \subseteq \mathcal{T}[\mathbb{T}]$.
- $\mathbf{t} \in \mathcal{T}[\mathbb{T}] \wedge NF(\mathbf{t}) \implies \mathbf{t} \in \mathcal{V}[\mathbb{T}]$.
- If $\mathbf{t} \in \mathcal{T}[\mathbb{T}]$ then \mathbf{t} is closed.

Notice that we do *not* require values or expressions in the interpretation of type \mathbb{T} to be (syntactically) well-typed, only that they be closed. This is unusual in logical relations proofs, but it makes sense here since we are trying to develop a non-syntactic notion of type safety.

Now we would like to proceed by proving something like:

$$\emptyset \vdash \mathbf{t} : \mathbb{T} \implies \mathbf{t} \in \mathcal{T}[\mathbb{T}]$$

By expanding out the definition of $\mathbf{t} \in \mathcal{T}[\mathbb{T}]$ and comparing it to that of $SAFE(\mathbf{t})$, we can see that this is essentially a strengthening of the original Type Safety theorem statement, with $VAL(\mathbf{t}')$ replaced by the (stronger) $\mathbf{t}' \in \mathcal{V}[\mathbb{T}]$.

Naturally, the proof should be by structural induction on the typing derivation in the hypothesis. But it is clear that we will get stuck in the abstraction case because the typing context will no longer be empty. (Again, this happened in the normalization proof too.) So we need to generalize our theorem by considering how to “close off” terms. To this end, we can define the following *context interpretation* characterizing safe closing substitutions for a context.

$$\mathcal{G}[\Gamma] \triangleq \{\sigma : Var \rightarrow Value \mid Dom(\sigma) = Dom(\Gamma) \wedge \forall \mathbf{x} \in Dom(\Gamma), \sigma(\mathbf{x}) \in \mathcal{V}[\Gamma(\mathbf{x})]\}$$

Using this, we can finally define our actual logical predicate for semantic type safety:

$$\Gamma \vDash \mathbf{t} : \mathbb{T} \triangleq \forall \sigma \in \mathcal{G}[\Gamma]. \sigma(\mathbf{t}) \in \mathcal{T}[\mathbb{T}]$$

and phrase the two lemmas we will actually prove

LEMMA A: If $\Gamma \vdash \mathbf{t} : \mathbb{T}$ then $\Gamma \vDash \mathbf{t} : \mathbb{T}$.

In proof developments based on logical relations, this is usually called the “Fundamental Property” or “Basic Lemma.”

LEMMA B: If $\emptyset \vDash \mathbf{t} : \mathbb{T}$ then $SAFE(\mathbf{t})$.

Combining these lemmas (with $\Gamma = \emptyset$ in Lemma A) directly gives us type safety.

We start by proving Lemma B, which, as promised, is very easy.

Proof. Suppose $\emptyset \models \mathfrak{t} : \mathbb{T}$. Then we can pick σ in the definition of \models to be the empty substitution, so we have $\mathfrak{t} \in \mathcal{T}[\mathbb{T}]$. Thus, $\mathfrak{t} \rightarrow^* \mathfrak{t}' \wedge NF(\mathfrak{t}') \implies \mathfrak{t}' \in \mathcal{V}[\mathbb{T}]$. Since $\mathfrak{t}' \in \mathcal{V}[\mathbb{T}]$ implies $VAL(\mathfrak{t}')$, we have $\mathfrak{t} \rightarrow^* \mathfrak{t}' \wedge NF(\mathfrak{t}') \implies VAL(\mathfrak{t}')$, but this is exactly the definition of $SAFE(\mathfrak{t})$. \square

Now we proceed to prove Lemma A, by induction on the structure of the derivation of $\Gamma \vdash \mathfrak{t} : \mathbb{T}$.

Case T-VAR: $\mathfrak{t} = \mathbf{x} \quad \mathbf{x} : \mathbb{T} \in \Gamma$

Choose any $\sigma \in \mathcal{G}[\Gamma]$. We must show that $\sigma(\mathbf{x}) \in \mathcal{T}[\mathbb{T}]$. By definition of $\mathcal{G}[\Gamma]$, we have $\sigma(\mathbf{x}) \in \mathcal{V}[\mathbb{T}] \subseteq \mathcal{T}[\mathbb{T}]$.

Case T-TRUE: $\mathfrak{t} = \mathbf{true} \quad \Gamma \vdash \mathbf{true} : \mathbf{Bool}$

For any σ , we have $\sigma(\mathbf{true}) = \mathbf{true} \in \mathcal{V}[\mathbf{Bool}] \subseteq \mathcal{T}[\mathbf{Bool}]$.

Case T-FALSE: $\mathfrak{t} = \mathbf{false} \quad \Gamma \vdash \mathbf{false} : \mathbf{Bool}$

Similar.

Case T-IF: $\mathfrak{t} = \mathbf{if} \ \mathfrak{t}_1 \ \mathbf{then} \ \mathfrak{t}_2 \ \mathbf{else} \ \mathfrak{t}_3 \quad \Gamma \vdash \mathfrak{t}_1 : \mathbf{Bool} \quad \Gamma \vdash \mathfrak{t}_2 : \mathbb{T} \quad \Gamma \vdash \mathfrak{t}_3 : \mathbb{T}$

Choose any σ in $\mathcal{G}[\Gamma]$. We must show $\sigma(\mathfrak{t}) = \mathbf{if} \ \sigma(\mathfrak{t}_1) \ \mathbf{then} \ \sigma(\mathfrak{t}_2) \ \mathbf{else} \ \sigma(\mathfrak{t}_3) \in \mathcal{T}[\mathbb{T}]$.

By applying the induction hypothesis to the three sub-derivations and instantiating at σ , we get $\sigma(\mathfrak{t}_1) \in \mathcal{T}[\mathbf{Bool}]$, $\sigma(\mathfrak{t}_2) \in \mathcal{T}[\mathbb{T}]$, and $\sigma(\mathfrak{t}_3) \in \mathcal{T}[\mathbb{T}]$. In particular, this means that $\sigma(\mathfrak{t}_1)$, $\sigma(\mathfrak{t}_2)$, and $\sigma(\mathfrak{t}_3)$ are closed.

Unfolding the definition of $\mathcal{T}[\mathbb{T}]$, we must show

- $\mathbf{if} \ \sigma(\mathfrak{t}_1) \ \mathbf{then} \ \sigma(\mathfrak{t}_2) \ \mathbf{else} \ \sigma(\mathfrak{t}_3)$ is closed. This follows directly from the fact that $\sigma(\mathfrak{t}_1)$, $\sigma(\mathfrak{t}_2)$, and $\sigma(\mathfrak{t}_3)$ are closed.
- $\mathbf{if} \ \sigma(\mathfrak{t}_1) \ \mathbf{then} \ \sigma(\mathfrak{t}_2) \ \mathbf{else} \ \sigma(\mathfrak{t}_3) \rightarrow^* \mathfrak{t}' \wedge NF(\mathfrak{t}') \implies \mathfrak{t}' \in \mathcal{V}[\mathbb{T}]$.

So suppose $\mathbf{if} \ \sigma(\mathfrak{t}_1) \ \mathbf{then} \ \sigma(\mathfrak{t}_2) \ \mathbf{else} \ \sigma(\mathfrak{t}_3) \rightarrow^* \mathfrak{t}'$ and $NF(\mathfrak{t}')$. Analyzing the stepping rules, we can see that this can only happen if, by repeated uses of E-IF, there is a normal form \mathfrak{t}'_1 such that $\sigma(\mathfrak{t}_1) \rightarrow^* \mathfrak{t}'_1$ and

$\mathbf{if} \ \sigma(\mathfrak{t}_1) \ \mathbf{then} \ \sigma(\mathfrak{t}_2) \ \mathbf{else} \ \sigma(\mathfrak{t}_3) \rightarrow^* \mathbf{if} \ \mathfrak{t}'_1 \ \mathbf{then} \ \sigma(\mathfrak{t}_2) \ \mathbf{else} \ \sigma(\mathfrak{t}_3) \rightarrow^* \mathfrak{t}'$

Now, since $\sigma(\mathfrak{t}_1) \in \mathcal{T}[\mathbf{Bool}]$ and $NF(\mathfrak{t}'_1)$, we know $\mathfrak{t}'_1 \in \mathcal{V}[\mathbf{Bool}] = \{\mathbf{true}, \mathbf{false}\}$. So suppose $\mathfrak{t}'_1 = \mathbf{true}$ (the \mathbf{false} case is similar). Then, by E-TRUE,

$\mathbf{if} \ \mathfrak{t}'_1 \ \mathbf{then} \ \sigma(\mathfrak{t}_2) \ \mathbf{else} \ \sigma(\mathfrak{t}_3) \rightarrow \sigma(\mathfrak{t}_2) \rightarrow^* \mathfrak{t}'$

Since $\sigma(\mathfrak{t}_2) \in \mathcal{T}[\mathbb{T}]$ and $NF(\mathfrak{t}')$, we have $\mathfrak{t}' \in \mathcal{V}[\mathbb{T}]$, as required.

Case T-ABS: $\mathfrak{t} = \lambda \mathbf{x} : \mathbb{T}_1 . \mathfrak{t}_2 \quad \mathbb{T} = \mathbb{T}_1 \rightarrow \mathbb{T}_2 \quad \Gamma, \mathbf{x} : \mathbb{T}_1 \vdash \mathfrak{t}_2 : \mathbb{T}_2$

Choose any σ in $\mathcal{G}[\Gamma]$. We must show $\sigma(\mathfrak{t}) = \lambda \mathbf{x} : \mathbb{T}_1 . \sigma(\mathfrak{t}_2) \in \mathcal{T}[\mathbb{T}_1 \rightarrow \mathbb{T}_2] = \mathcal{T}[\mathbb{T}]$. Equivalently, unfolding the definition, we must show that

- $\lambda \mathbf{x} : \mathbb{T}_1 . \sigma(\mathfrak{t}_2)$ is closed. From the sub-derivation, we know that the free variables of \mathfrak{t}_2 are in $Dom(\Gamma) \cup \{\mathbf{x}\}$. Then, by the definition of $\mathcal{G}[\Gamma]$, the only possible free variable in $\sigma(\mathfrak{t}_2)$ is \mathbf{x} , so $\lambda \mathbf{x} : \mathbb{T}_1 . \sigma(\mathfrak{t}_2)$ is indeed closed.

- $\lambda x:T_1.\sigma(\mathbf{t}_2) \rightarrow^* \mathbf{t}' \wedge NF(\mathbf{t}') \implies \mathbf{t}' \in \mathcal{V}[\mathbb{T}_1 \rightarrow \mathbb{T}_2]$.

Suppose $\lambda x:T_1.\sigma(\mathbf{t}_2) \rightarrow^* \mathbf{t}'$. Then, since $\lambda x:T_1.\sigma(\mathbf{t}_2)$ is already a value, we must have $\mathbf{t}' = \lambda x:T_1.\sigma(\mathbf{t}_2)$. To show that $\lambda x:T_1.\sigma(\mathbf{t}_2) \in \mathcal{V}[\mathbb{T}_1 \rightarrow \mathbb{T}_2]$, we pick an arbitrary $v \in \mathcal{V}[\mathbb{T}_1]$ and show that $[x \mapsto v](\sigma(\mathbf{t}_2)) \in \mathcal{T}[\mathbb{T}_2]$. Since σ maps variables to closed terms, $[x \mapsto v](\sigma(\mathbf{t}_2)) = \sigma'(\mathbf{t}_2)$, where $\sigma' = (\sigma, x \mapsto v)$, i.e., the substitution that extends σ with a mapping from x to v . But it is easy to see that $\sigma' \in \mathcal{G}[\Gamma, x : T_1]$. So applying the induction hypothesis to the sub-derivation gives $\sigma'(\mathbf{t}_2) \in \mathcal{T}[\mathbb{T}_2]$, as required.

Case T-APP $\mathbf{t} = \mathbf{t}_1 \ \mathbf{t}_2$ $\Gamma \vdash \mathbf{t}_1 : T_2 \rightarrow T$ $\Gamma \vdash \mathbf{t}_2 : T_2$

Choose any σ in $\mathcal{G}[\Gamma]$. We must show $\sigma(\mathbf{t}_1 \ \mathbf{t}_2) = \sigma(\mathbf{t}_1) \ \sigma(\mathbf{t}_2) \in \mathcal{T}[\mathbb{T}]$. Applying the induction hypothesis on the sub-derivations and instantiating at σ gives us $\sigma(\mathbf{t}_1) \in \mathcal{T}[\mathbb{T}_2 \rightarrow \mathbb{T}]$ and $\sigma(\mathbf{t}_2) \in \mathcal{T}[\mathbb{T}_2]$. In particular, this means that $\sigma(\mathbf{t}_1)$ and $\sigma(\mathbf{t}_2)$ are closed.

Unfolding the definition of $\mathcal{T}[\mathbb{T}]$, we must show that

- $\sigma(\mathbf{t}_1) \ \sigma(\mathbf{t}_2)$ is closed. This follows immediately from the fact that $\sigma(\mathbf{t}_1)$ and $\sigma(\mathbf{t}_2)$ are closed.
- $\sigma(\mathbf{t}_1) \ \sigma(\mathbf{t}_2) \rightarrow^* \mathbf{t}' \wedge NF(\mathbf{t}') \implies \mathbf{t}' \in \mathcal{V}[\mathbb{T}]$.

So suppose $\sigma(\mathbf{t}_1) \ \sigma(\mathbf{t}_2) \rightarrow^* \mathbf{t}'$ and $NF(\mathbf{t}')$. Analyzing the stepping rules, we can see that this can only happen if, by repeated use of E-APP1, there is a normal form \mathbf{t}'_1 such that $\sigma(\mathbf{t}_1) \rightarrow^* \mathbf{t}'_1$ and

$$\sigma(\mathbf{t}_1) \ \sigma(\mathbf{t}_2) \rightarrow^* \mathbf{t}'_1 \ \sigma(\mathbf{t}_2) \rightarrow^* \mathbf{t}'.$$

Now, since $\sigma(\mathbf{t}_1) \in \mathcal{T}[\mathbb{T}_2 \rightarrow \mathbb{T}]$ and $NF(\mathbf{t}'_1)$, we know that $\mathbf{t}'_1 \in \mathcal{V}[\mathbb{T}_2 \rightarrow \mathbb{T}]$, so \mathbf{t}'_1 has the form $\lambda x:T_2.\mathbf{t}_{11}$, where $\forall v \in \mathcal{V}[\mathbb{T}_2], [x \mapsto v]\mathbf{t}_{11} \in \mathcal{T}[\mathbb{T}]$.

So, by repeated use of E-APP2, there is a normal form \mathbf{t}'_2 such that $\sigma(\mathbf{t}_2) \rightarrow^* \mathbf{t}'_2$ and

$$\mathbf{t}'_1 \ \sigma(\mathbf{t}_2) = (\lambda x:T_2.\mathbf{t}_{11}) \ \sigma(\mathbf{t}_2) \rightarrow^* (\lambda x:T_2.\mathbf{t}_{11}) \ \mathbf{t}'_2 \rightarrow^* \mathbf{t}'.$$

Since $\sigma(\mathbf{t}_2) \in \mathcal{T}[\mathbb{T}_2]$ and $NF(\mathbf{t}'_2)$, we know that $\mathbf{t}'_2 \in \mathcal{V}[\mathbb{T}_2]$. Hence \mathbf{t}'_2 is a value, and so, by E-APPABS,

$$(\lambda x:T_2.\mathbf{t}_{11}) \ \mathbf{t}'_2 \rightarrow [x \mapsto \mathbf{t}'_2]\mathbf{t}_{11} \rightarrow^* \mathbf{t}'$$

Finally, since $[x \mapsto \mathbf{t}'_2]\mathbf{t}_{11} \in \mathcal{T}[\mathbb{T}]$ and $NF(\mathbf{t}')$, we have $\mathbf{t}' \in \mathcal{V}[\mathbb{T}]$, as required. \square