

Thus, we can think of 0 as the type of a term that does not actually return to its surrounding context.

3 Continuations and Negation

What is the significance of negation? We know that logically $\neg\varphi$ is equivalent to $\varphi \Rightarrow \perp$, which suggests that we can think of $\neg\varphi$ as corresponding to a function $\tau \rightarrow 0$. We have seen functions that accept a type and do not return a value before: continuations have that behavior. If φ corresponds to τ , a reasonable interpretation of $\neg\varphi$ is as a continuation expecting a τ . Negation corresponds to turning outputs into inputs.

As we saw above with currying and uncurrying, meaning-preserving program transformations can have interesting logical interpretations. What about conversion to continuation-passing style? We represent a continuation k expecting a value of type τ as a function with type $\tau \rightarrow 0$.

We can then define CPS conversion as a type-preserving translation $\llbracket \Gamma \vdash e : \tau \rrbracket$. Here we include the entire type derivation $\Gamma \vdash e : \tau$ inside the $\llbracket \cdot \rrbracket$ because types are not unique and the translation depends on the typing. The translation is type-preserving in the sense that a well-typed source term $(\Gamma \vdash e : \tau)$ translates to a well-typed target term:

$$\mathcal{G}[\Gamma] \vdash \llbracket \Gamma \vdash e : \tau \rrbracket : \mathcal{T}[\tau].$$

The translation of the typing context $\mathcal{G}[\Gamma]$ simply translates all the contained variables:

$$\mathcal{G}[x_1 : \tau_1, \dots, x_n : \tau_n] = x_1 : \mathcal{T}[\tau_1], \dots, x_n : \mathcal{T}[\tau_n].$$

The soundness of the translation can be seen by induction on the typing derivation.

$$\begin{aligned} \llbracket \Gamma, x : \tau \vdash x : \tau \rrbracket &= \lambda k : \mathcal{T}[\tau] \rightarrow 0. k x \\ \llbracket \Gamma \vdash (\lambda x : \tau. e) : \tau \rightarrow \tau' \rrbracket &= \lambda k : \mathcal{T}[\tau \rightarrow \tau'] \rightarrow 0. k (\lambda k' : \mathcal{T}[\tau'] \rightarrow 0. \lambda x : \mathcal{T}[\tau]. \llbracket \Gamma, x : \tau \vdash e : \tau' \rrbracket k') \\ \llbracket \Gamma \vdash (e_0 e_1) : \tau' \rrbracket &= \lambda k : \mathcal{T}[\tau'] \rightarrow 0. \llbracket \Gamma \vdash e_0 : \tau \rightarrow \tau' \rrbracket (\lambda f : \mathcal{T}[\tau \rightarrow \tau']. \llbracket \Gamma \vdash e_1 : \tau \rrbracket (\lambda v : \mathcal{T}[\tau]. f k v)) \end{aligned}$$

To make this type-check, we define the type translation $\mathcal{T}[\cdot]$ as follows:

$$\begin{aligned} \mathcal{T}[B] &= B \\ \mathcal{T}[\tau \rightarrow \tau'] &= (\mathcal{T}[\tau'] \rightarrow 0) \rightarrow (\mathcal{T}[\tau] \rightarrow 0) \end{aligned}$$

Note that the logical interpretation of the translation of a function type corresponds to the use of the contrapositive: $(\varphi \Rightarrow \psi) \Rightarrow (\neg\psi \Rightarrow \neg\varphi)$.

By induction on the typing derivation, we can see that CPS conversion converts a source term of type τ into a target term of type $(\mathcal{T}[\tau] \rightarrow 0) \rightarrow 0$. Since programs correspond to proofs, CPS conversion shows how to convert a proof of proposition φ into a proof of proposition $\neg\neg\varphi$. In other words, CPS conversion proves the admissibility in constructive logic of the rule for introducing double negation:

$$\frac{\varphi}{\neg\neg\varphi}$$

However, we are unable to invert CPS translation, and similarly we are unable (constructively) to *remove* double negation.

4 Extracting Computational Content

Many automated deduction systems, such as NuPrl and Coq, are based on constructive logic. Automatic programming was a significant research direction that motivated the development of these systems. The idea

was that a constructive proof of the existence of a function would automatically yield a program to compute it: the statement asserting the existence of the function is a type, and a constructive proof yields a λ -term inhabiting that type. For example, to obtain a program computing square roots, one merely has to give a constructive proof of the statement $\forall x \geq 0 \exists y \ y^2 = x$.

5 Other Directions

If second-order constructive predicate logic corresponds to System F, do other logics correspond to new kinds of programming language features? This has been an avenue of fruitful exploration over the last couple of decades, with programming-language researchers deriving insights from classical logic, higher-order, and linear logics that help guide the design of useful language features.

For example, *linear logic* is a logic that keeps track of resources. One may only use an assumption in the application of a rule once; the assumption is consumed and may not be reused. This corresponds to functions that consume their arguments, and hence is a possible model for systems with bounded resources.