

4.4 Proof Irrelevance

So far, we have carefully designed the proof term language so that we can reconstruct the original deduction, maintaining a bijection. In many contexts this will give us too much information. Returning to one of our motivating examples,

$$\forall x:\text{nat}. \exists y:\text{nat}. y > x \wedge \text{prime}(y)$$

we can see from what we have developed so far, that the computational content of a constructive proof of this proposition will be a function from a natural number x that returns a pair consisting of a witness p and a proof that $p > x \wedge \text{prime}(p)$ which is again a pair of proofs. While it may be important for us to know that these last two proofs exist, in practice we may only be interested in the prime p , and not the proof that it is indeed greater than x or a prime number.

Our objective in this section is to find means to selectively hide portions of a proof. In one direction, this will allow us to extract only part of the computational content of a proof, making the resulting code much more efficient to execute. In the other direction, it will allow us write proof-agnostic functions and reason about them logically.

We will use all concepts and techniques we have developed so to achieve this goal. The basic idea is to introduce a new proposition $[A]$, pronounced “*bracket A*”, for any proposition A . $[A]$ should be true if A is true, except that we intend to erase the proof before using it computationally. A first cut at the introduction rule would be

$$\frac{A \text{ true}}{[A] \text{ true}} \quad [I]?$$

The corresponding elimination rule

$$\frac{A \text{ true}}{[A] \text{ true}} \quad [E]?$$

is sound with respect to the introduction rule, but unfortunately fails to capture the intent of the type $[A]$. It says that if we have a proof of $[A]$ (which will be erased at runtime) then we can get a proof of A (which will be needed at runtime). It is evident that after erasing the proof of $[A]$ we would have to make up a proof of A out of thin air, which is not possible in general.

In order to capture this, we need a new judgment $A \text{ irr}$ (pronounced “*A irrelevant*”) which says that A is true but its proof is not computationally available. The revised elimination rule exploits this judgment to express that if we know $[A] \text{ true}$ we know that A is true, but that we cannot use its proof computationally.

$$\frac{\begin{array}{c} \overline{\quad} \quad u \\ A \text{ irr} \\ \vdots \\ [A] \text{ true} \quad C \text{ true} \end{array}}{C \text{ true}} \quad [E]^u$$

How do we use assumptions $A \text{ irr}$? From the discussion above it should be clear that it is must remain too weak to prove $A \text{ true}$: the latter requires a computationally relevant proof term but the former cannot offer one. However, when we are trying to prove $[C] \text{ true}$, then we should be allowed to use assumptions $A \text{ irr}$ because a proof of $[C]$ will be erased. The corresponding rule is a bit difficult to express in the two-dimensional format. It says that at the point we introduce $[C]$ we transform every assumption $A \text{ irr}$ to a corresponding assumption $A \text{ true}$.

$$\frac{\frac{\overline{A_1 \text{ irr}}^{u_1}}{A_1 \text{ true}} \quad \cdots \quad \frac{\overline{A_n \text{ irr}}^{u_n}}{A_n \text{ true}}}{\vdots} \quad \frac{C \text{ true}}{[C] \text{ true}} \quad [I]$$

Before carrying out some examples, let us make sure that the rules are locally sound and complete. First, the local reduction which witnesses local soundness.

$$\frac{\frac{\frac{\overline{A_1 \text{ irr}}^{u_1}}{A_1 \text{ true}} \quad \cdots \quad \frac{\overline{A_n \text{ irr}}^{u_n}}{A_n \text{ true}}}{\mathcal{D}} \quad \frac{A \text{ true}}{[A] \text{ true}} \quad [I]}{C \text{ true}} \quad \frac{\frac{\overline{A \text{ irr}}^u}{\mathcal{E}}}{C \text{ true}} \quad [E^u]}{\implies_R} \quad \frac{\frac{\frac{\overline{A_1 \text{ irr}}^{u_1}}{A_1 \text{ true}} \quad \cdots \quad \frac{\overline{A_n \text{ irr}}^{u_n}}{A_n \text{ true}}}{\mathcal{D}} \quad \frac{A \text{ true}}{A \text{ irr}} \quad u}{\mathcal{E}}}{C \text{ true}} \quad [E^u]$$

The operation on the right is again not well represented, visually. It requires a new substitution principle to substitute a deduction of $A \text{ true}$ for assumption $A \text{ irr}$. This works because the assumption $A \text{ irr}$ can only be used when it is promoted to an assumption $A \text{ true}$, at which place we can use the ordinary substitution principle. We will see more formally how this works when we have proof terms below, and hypotheses are written in a localized form.

The local expansion is much simpler.

$$[A] \text{ true} \implies_E \frac{\mathcal{D} \quad \frac{[A] \text{ true} \quad \frac{\overline{A \text{ true}}}{[A] \text{ true}} \quad [I]}{\overline{A \text{ irr}}^u} \quad [E^u]}{[A] \text{ true}}$$

Here, $A \text{ true}$ is available to prove the premise of the $[I]$ rule because applying this rule promotes the assumption $A \text{ irr}$.

We consider three examples. The first proves that $A \supset [A]$, that is, we can

forget computational content if we so choose.

$$\frac{\frac{\overline{A \text{ true}} \quad u}{[A] \text{ true}} []I}{A \supset [A] \text{ true}} \supset I^u$$

The second one shows that we cannot prove $[A] \supset A$, that is, we cannot spontaneously create computational content,

$$\frac{\frac{\overline{A \text{ true}} \quad u \quad \vdots \quad A \text{ true}}{A \text{ true}} []E^w}{[A] \supset A \text{ true}} \supset I^u$$

Of course, this is not a convincing argument that we cannot prove $[A] \supset A$, but if we also knew that we can always find a proof by using introduction rules from below and elimination rules from above, then indeed there cannot be any proof.

Finally, we can distribute brackets over implication.

$$\frac{\frac{\frac{\overline{A \supset B \text{ irr}} \quad u' \quad \overline{A \text{ irr}} \quad w'}{A \supset B \text{ true}} \supset E}{\frac{\overline{A \supset B} \text{ true} \quad u}{[A \supset B] \text{ true}} \quad \frac{\frac{\overline{[A] \text{ true}} \quad w \quad \frac{\overline{B \text{ true}} []I}{[B] \text{ true}} []E^{w'}}{[B] \text{ true}} []E^{u'}}{[B] \text{ true}} \supset I^w}{[A] \supset [B] \text{ true}} \supset I^u}{[A \supset B] \supset [A] \supset [B] \text{ true}} \supset I^u$$

Now we move on to proof terms. We write hypotheses in localized form, $\Gamma \vdash M : A$. There are two forms of hypothesis: $u:A$ which labels an assumption $A \text{ true}$, and $[u]:A$, which labels an assumption $A \text{ irr}$. The brackets around u indicate that it cannot be used directly, but only after a $[]$ introduction, which “unlocks” the variable $u:A$. In order to describe this process we define the process of *promotion*, written Γ^\oplus .

$$\begin{aligned} (\cdot)^\oplus &= \cdot \\ (\Gamma, u:A)^\oplus &= \Gamma^\oplus, u:A \\ (\Gamma, [u]:A)^\oplus &= \Gamma^\oplus, u:A \end{aligned}$$

We then have the following two rules.

$$\frac{\Gamma^\oplus \vdash M : A}{\Gamma \vdash [M] : [A]} []I \quad \frac{\Gamma \vdash M : [A] \quad \Gamma, [u]:A \vdash N : C}{\Gamma \vdash \mathbf{let} [u] = M \mathbf{in} N : C} []E^u$$

What the hypothesis promotion achieves is, intuitively, that a variable $[u]:A$ can only be used inside brackets in a term M and not outside. For example:

$$\lambda f. \lambda x. \mathbf{let} [f'] = f \mathbf{in} \mathbf{let} [x'] = x \mathbf{in} [f' x'] : [A \supset B] \supset [A] \supset [B]$$

$$\lambda u. [u] : A \supset [A]$$

are well-typed. The first, because we are free to use u (which stands for A *true*) inside brackets, where it will not be used computationally, and the second because both f' and x' are used only inside brackets, where they will not be used computationally. On the other hand

$$\lambda u. \mathbf{let} [u'] = u \mathbf{in} u' \not\vdash [A] \supset A$$

because u' is incorrectly used *outside* a bracket context.

Local reductions and expansions, as well as the substitution principle, are now much clearer on proof terms.

$$\text{If } \Gamma^{\oplus} \vdash M : A \text{ and } \Gamma, [u] : A, \Gamma' \vdash N : C \text{ then } \Gamma, \Gamma' \vdash [M/u]N : C.$$

This is correct because the only place where $[u]:A$ can be used in the second deduction is underneath a $[\]$ constructor where the context is promoted to $\Gamma^{\oplus}, \Gamma'^{\oplus}$ so that the ordinary substitution principle applies.

$$\begin{array}{l} \mathbf{let} [u] = [M] \mathbf{in} N \quad \Longrightarrow_R \quad [M/u]N \\ M : [A] \quad \quad \quad \quad \quad \Longrightarrow_E \quad \mathbf{let} [u] = M \mathbf{in} [u] \end{array}$$

We now go through a few more examples. In order to write these examples more compactly, we assume that every assumption $v:[A]$ is immediately decomposed into $[u]:A$, and, moreover, u is longer used. In that case, we can write $\lambda[u]. M$ instead of $\lambda v. \mathbf{let} [u] = v \mathbf{in} M$.

First, we consider variants of the following proposition:

$$(\exists x. A(x) \wedge B(x)) \supset (\exists x. B(x) \wedge A(x))$$

The starting proof of this is straightforward:

$$\lambda u. \mathbf{let} \langle y, w \rangle = u \mathbf{in} \langle y, \langle \mathbf{snd} w, \mathbf{fst} w \rangle \rangle$$

Hiding in proof information in the output is consistent.

$$\lambda u. \mathbf{let} \langle y, w \rangle = u \mathbf{in} \langle y, [\langle \mathbf{snd} w, \mathbf{fst} w \rangle] \rangle : (\exists x. A(x) \wedge B(x)) \supset (\exists x. [B(x) \wedge A(x)])$$

We can still hide consistently even if the input proofs (of $A(x)$ and $B(x)$) are not available at runtime.

$$\lambda u. \mathbf{let} \langle y, [w] \rangle = u \mathbf{in} \langle y, [\langle \mathbf{snd} w, \mathbf{fst} w \rangle] \rangle : (\exists x. [A(x) \wedge B(x)]) \supset (\exists x. [B(x) \wedge A(x)])$$

This is correct because the only occurrences of w in its scope are inside brackets. Erasing w will therefore not lead to any dangling variables, that is, variables

that would be incorrectly assumed to be available at runtime. However, if we also hide the witness y then the result is no longer well-formed.

$$\lambda u. \mathbf{let} [\langle y, w \rangle] = u \mathbf{in} \langle y, [\langle \mathbf{snd} w, \mathbf{fst} w \rangle] \rangle \not\vdash (\exists x. A(x) \wedge B(x)) \supset (\exists x. [B(x) \wedge A(x)])$$

The problem here is that y is bound inside the brackets but used outside. If we erased all the content inside the brackets we would have

$$\lambda u. \mathbf{let} [] = u \mathbf{in} \langle y, [] \rangle \not\vdash \supset (\exists x. [])$$

which is not meaningful as a program even if we interpret $[]$ at \top in the proposition and the unit element $\langle \rangle$ in the term.

It is even possible to be more fine-grained. For example, we can hide the proof of $A(x)$ in the antecedent of the implication or we also hide in the succedent.

Another set of examples comes from a specification of graph reachability. We specify that for any two nodes x and y in a graph, either there is a path p connecting the two nodes or not.

$$\forall x. \forall y. (\exists p. \mathbf{path}(p, x, y)) \vee \neg(\exists p. \mathbf{path}(p, x, y))$$

A proof of this, when viewed computationally, will be a function taking nodes x and y as input and returning either $\mathbf{inl} M$ or $\mathbf{inr} N$. In the first case, a path exists, and M is a pair consisting of the path p and a proof that it connects x and y in the graph. In the second case, no path exists, and M is a function which derives a contradiction from any candidate path given to it as an argument.

One natural refinement of this specification is to hide the proof that p is indeed a valid path connecting x and y (which can easily verify this ourselves) and the proof that there is no path connecting x and y . For the latter, we would trust the proof (because, say, have verified it before erasing it). To capture this we would write

$$\forall x. \forall y. (\exists p. [\mathbf{path}(p, x, y)]) \vee [\neg(\exists p. \mathbf{path}(p, x, y))]$$

Now the function extracted from this proof would at runtime return either $\mathbf{inl} \langle p, [] \rangle$ (that is, the path connecting x and y) or $\mathbf{inr} []$ to indicate there is no such path.

We can take this one step further, also hiding the path itself. In this case, the function returns either $\mathbf{inl} []$ if there is a path connecting x and y and $\mathbf{inr} []$ if there is none.

$$\forall x. \forall y. [\exists p. \mathbf{path}(p, x, y)] \vee [\neg(\exists p. \mathbf{path}(p, x, y))]$$

Summary. We present a summary of the rules for proof irrelevance, using the local form for hypothesis. We have a new form of judgment, $A \text{ irr}$ which means that A is true, but the evidence for that is computationally irrelevant, that is,

may be erased before executing the program. We arrange that $A \text{ irr}$ is only used as a hypothesis. We can *promote* such assumptions using the Δ^\oplus operation:

$$\begin{aligned} (\cdot)^\oplus &= \cdot \\ (\Delta, A \text{ true})^\oplus &= \Delta^\oplus, A \text{ true} \\ (\Delta, A \text{ irr})^\oplus &= \Delta^\oplus, A \text{ true} \end{aligned}$$

Then the introduction and elimination rules are

$$\frac{\Delta^\oplus \vdash A \text{ true}}{\Delta \vdash [A] \text{ true}} [I] \quad \frac{\Delta \vdash [A] \text{ true} \quad \Delta, A \text{ irr} \vdash C \text{ true}}{\Delta \vdash C \text{ true}} [E^u]$$

In the presence of proof terms, we write $[u]:A$ as a labeling of the assumption $A \text{ true}$. Then the promotion operation on labeled contexts Γ becomes

$$\begin{aligned} (\cdot)^\oplus &= \cdot \\ (\Gamma, u:A)^\oplus &= \Gamma^\oplus, u:A \\ (\Gamma, [u]:A)^\oplus &= \Gamma^\oplus, u:A \end{aligned}$$

With proof terms, the introduction and elimination rules then read as follows:

$$\frac{\Gamma^\oplus \vdash M : A}{\Gamma \vdash [M] : [A]} [I] \quad \frac{\Gamma \vdash M : [A] \quad \Gamma, [u]:A \vdash N : C}{\Gamma \vdash \mathbf{let} [u] = M \mathbf{in} N : C} [E^u]$$