Comparing Control Constructs by Double-barrelled CPS *

Hayo Thielecke
(h.thielecke@cs.bham.ac.uk)
School of Computer Science, University of Birmingham, Birmingham B15 2TT, United Kingdom

Abstract. We investigate call-by-value continuation-passing style transforms that pass two continuations. Altering a single variable in the translation of λ-abstraction gives rise to different control operators: first-class continuations; dynamic control; and (depending on a further choice of a variable) either the `return` statement of C; or Landin's **J**-operator. In each case there is an associated simple typing. For those constructs that allow upward continuations, the typing is classical, for the others it remains intuitionistic, giving a clean distinction independent of syntactic details. Moreover, those constructs that make the typing classical in the source of the CPS transform break the linearity of continuation use in the target.

Keywords: Continuations, control operators, J-operator, intuitionistic and classical logic

1. Introduction

Control operators come in bewildering variety. Sometimes the same term is used for distinct constructs, as with `catch` in early Scheme or `throw` in Standard ML of New Jersey, which are very unlike the `catch` and `throw` in Lisp whose names they borrow. On the other hand, this Lisp `catch` is fundamentally similar to exceptions despite their dissimilar and much more ornate appearance.

Fortunately it is sometimes possible to glean some high-level "logical" view of a programming language construct by looking only at its type. Recall that under the "formulae as types" correspondence, the types of purely functional programs correspond to formulae provable in intuitionistic logic; for example, the identity $\lambda x.x$ has type $A \to A$, which we can read as "A implies A". As Griffin [4] discovered, this correspondence extends to control, in that control operators for first-class continuations can be ascribed types corresponding to formulae which are provable only in classical, but not in intuitionistic logic, such as Peirce's law $((A \to B) \to A) \to A$. In that sense, the addition of first-class continuations leads to an increase in power of the language that is visible even at the level of the types. This gives us a fundamental distinction between languages that have such classical types and those

* An earlier version appeared in the proceedings of the 3rd ACM Workshop on Continuations [20]

that do not, even though they may still enjoy some form of control. Such an approach based on typing complements comparisons based on contextual equivalences [14, 19].

Such a comparison would be difficult unless we blot out complication. In particular, exceptions are typically tied in with other, fairly complicated features of the language which are not relevant to control as such: in ML with the datatype mechanism, in Java with object-orientation. In order to simplify, we first strip down control operators to the bare essentials of labelling and jumping, so that there are no longer any distracting syntactic differences between them. The grammar of our toy language is uniformly this:

$$M ::= x \mid \lambda x.M \mid MM \mid \texttt{here}\, M \mid \texttt{go}\, M.$$

The intended meaning of `here` is that it labels a "program point" or expression without actually naming any particular label—just uttering the demonstrative "here", as it were. Correspondingly, `go` jumps to a place specified by a `here`, without naming the "to" of a `goto`.

Despite the simplicity of the language, there is still scope for variation: not by adding bells and whistles to `here` and `go`, but by varying the meaning of λ-abstraction. Its impact can be seen quite clearly in the distinction between exceptions and first-class continuations. The difference between them is as much due to the meaning of λ-abstraction as due to the control operators themselves, since λ-abstraction determines what is statically put into a closure and what is passed dynamically. Readers familiar with, say, Scheme implementations will perhaps not be surprised about the impact of what becomes part of a closure. But the point of this paper is twofold:

- small variations in the meaning of λ completely change the meaning of our control operators;
- we can see these differences at an abstract, logical level, without delving into the innards of interpreters.

Overview

We give meaning to the λ-calculus enriched with `here` and `go` by means of continuations in Section 2, examining in Sections 3–5 how variations on λ-abstraction determine what kind of control operations `here` and `go` represent. For each of these variations we present a simple typing, which agrees with the transform (Section 6). By refining the typing of the target λ-calculus of the CPS transform with linearity, we show that those constructs that make the typing classical in the source of

the CPS transform break the linearity of continuation use in the target (Section 7). We conclude by summarising the significance of these typings in terms of classical and intuitionistic logic (Section 8).

The prerequisites of this paper, besides some background knowledge in programming languages, are some familiarity with continuations (in the form of denotational semantics or interpreters), and the most basic facts about intuitionistic logic, as can be found in many logic textbooks [21, 22].

2. Double-barrelled CPS transform

Our starting point is a continuation-passing style (CPS) transform, which transforms λ-terms enriched with the `here` and `go`-operations (the source language) into ordinary λ-calculus without control operations (the target). At first, we will read this target language as untyped λ-calculus, before refining it with types in Sections 6 and 7.

This transform is double-barrelled in the sense that it always passes *two* continuations. Hence the clauses start with $\lambda kq.\ldots$ instead of $\lambda k.\ldots$. Other than that, this CPS transform is in fact a very mild variation on the usual call-by-value one [10] (one could just as well use a slightly different transform, for instance one where the continuation is the first argument to a function). As indicated by the $\boxed{?}$, we leave one variable, the extra continuation passed to the body of a λ-abstraction, unspecified.

$$\begin{aligned}
[\![x]\!] &= \lambda kq.kx \\
[\![\lambda_? x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{?}) \\
[\![MN]\!] &= \lambda kq.[\![M]\!](\lambda m.[\![N]\!](\lambda n.mnkq)q)q \\
[\![\texttt{here}\,M]\!] &= \lambda kq.[\![M]\!]kk \\
[\![\texttt{go}\,M]\!] &= \lambda kq.[\![M]\!]qq
\end{aligned}$$

The extra continuation q may be seen as a jump continuation, in that its manipulation accounts for the labelling and jumping. This is done symmetrically: `here` makes the second continuation the same as the current one k, whereas `go` sets the current continuation of its argument to the jump continuation q. The clauses for variables and applications do not interact with the additional jump continuation: the former ignores it, while the latter merely distributes it into the operator, the operand and the function call.

Only in the clause for λ-abstraction do we face a design decision. Depending on which continuation (static s, dynamic d, or the return

continuation r) we fill in for "?" in the clause for λ, there are three different flavours of λ-abstraction.

$$\begin{aligned}
[\![\lambda_{\mathsf{s}} x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{s}) \\
[\![\lambda_{\mathsf{d}} x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{d}) \\
[\![\lambda_{\mathsf{r}} x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{r})
\end{aligned}$$

The lambdas are subscripted to distinguish them, and the box around the last variable is meant to highlight that this is the crucial difference between the transforms. Formally there is also a fourth possibility, the outer continuation k, but this seems less meaningful and would not fit into simple typing.

For all choices of λ, the operation `go` is always a jump to a place specified by a `here`. For example, for any M, the term `here` $((\lambda x.M)(\texttt{go}\, N))$ should be equivalent to N, as the `go` jumps past the M. But in more involved examples than this, there may be different choices *where* `go` can go to among several occurrences of `here`. In particular, if s is passed as the second continuation argument to M in the transform of $\lambda x.M$, then a `go` in M will refer to the `here` that was in scope at the point of definition (unless there is an intervening `here`, just as one binding of a variable x can shadow another). By contrast, if d is passed to M in $\lambda x.M$, then the `here` that is in scope at the point of definition is forgotten; instead `go` in M will refer to the `here` that is in scope at the point of call when $\lambda x.M$ is applied to an argument. In fact, depending upon the choice of variable in the clause for λ as above, `here` and `go` give rise to different control operations:

- first-class continuations like those given by `call/cc` in Scheme [5];
- dynamic control in the sense of Lisp, and typeable in a way reminiscent of checked exceptions;
- a `return`-operation, which can be refined into the **J**-operator invented by Landin in 1965 and ancestral to `call/cc` [5, 7, 8, 18].

It is perhaps surprising how subtle variations in the transform give rise to such different constructs, each of which has precedents in actual languages. Thus it may be helpful to recall a more traditional analogue of such a situation: consider how variations in the passing of environments can yield either static or dynamic binding (see the textbooks by Friedman, Wand and Haynes [3, Section 5.7], or Schmidt [15, Section 8.2]). Concretely, here is a simple denotational semantics $\mathcal{E}[\![-]\!]$ with environments, which we can equally read as a mathematically condensed

form of a straightforward environment-passing interpreter:

$$\begin{aligned}\mathcal{E}[\![x]\!]\, e &= e(x)\\ \mathcal{E}[\![\lambda x.M]\!]\, s &= \lambda v d.\left(\mathcal{E}[\![M]\!]\,(\boxed{?}[x \mapsto v])\right)\\ \mathcal{E}[\![MN]\!]\, e &= (\mathcal{E}[\![M]\!]\, e)\,(\mathcal{E}[\![N]\!]\, e)\, e\end{aligned}$$

Since the environment is passed along in an application MN, it is up to the clause for λ-abstraction *which* environment is to be extended with the actual argument v for the bound variable x, as indicated by $\boxed{?}$. If we choose the static environment s, the behaviour of variables will be that of static binding; if we choose the dynamic environment (supplied at the point of call), the behaviour will be that of dynamic binding. In this example, it is the meaning of variables which differs with the choice of environment, whereas in the double-barrelled CPS transforms, it is the meaning of `go`. In a sense, we can think of the second continuation as analogous to an environment for the single identifier `go`.

We examine the variations on the double-barrelled CPS transform in turn, giving a simple type system in each case. An unusual feature of these type judgements is that, because we have two continuations, there are two types in the succedent on the right of the turnstile, as in

$$\Gamma \vdash M : A, B.$$

The first type on the right accounts for the case that the term returns a value; it corresponds to the current continuation. The second type accounts for the extra continuation used for jumping. In logical terms, the comma on the right may be read as a disjunction. It makes a big difference whether this disjunction is classical or intuitionistic. That is our main criterion of comparing and contrasting the control constructs.

3. Static semantics and first-class continuations

The first choice of which continuation to pass to the body of a function is arguably the cleanest. Passing the static continuation s gives control the same static binding as ordinary λ-calculus variables. In the static case, the transform is this:

$$\begin{aligned}[\![x]\!] &= \lambda kq.kx\\ [\![\lambda_{\mathtt{s}} x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{s})\\ [\![MN]\!] &= \lambda kq.[\![M]\!](\lambda m.[\![N]\!](\lambda n.mnkq)q)q\\ [\![\mathtt{here}\, M]\!] &= \lambda kq.[\![M]\!]kk\\ [\![\mathtt{go}\, M]\!] &= \lambda kq.[\![M]\!]qq\end{aligned}$$

$$\frac{}{\Gamma, x : A, \Gamma' \vdash_{\mathsf{S}} x : A, C}$$

$$\frac{\Gamma \vdash_{\mathsf{S}} M : B, B}{\Gamma \vdash_{\mathsf{S}} \mathtt{here}\, M : B, C} \qquad \frac{\Gamma \vdash_{\mathsf{S}} M : B, B}{\Gamma \vdash_{\mathsf{S}} \mathtt{go}\, M : C, B}$$

$$\frac{\Gamma, x : A \vdash_{\mathsf{S}} M : B, C}{\Gamma \vdash_{\mathsf{S}} \lambda_{\mathsf{S}} x.M : A \to B, C} \qquad \frac{\Gamma \vdash_{\mathsf{S}} M : A \to B, C \qquad \Gamma \vdash_{\mathsf{S}} N : A, C}{\Gamma \vdash_{\mathsf{S}} MN : B, C}$$

Figure 1. Typing for static `here` and `go`

We type our source language with `here` and `go` as in Figure 1.

In logical terms, both `here` and `go` are a combined right weakening and contraction. By themselves, weakening and contraction do not amount to much; but it is the combination with the rule for $\to$-introduction that makes the calculus "classical", in the sense that there are terms whose types are propositions of classical, but not of intuitionistic, minimal logic.

To see how $\to$-introduction gives classical types, consider λ-abstracting over `go`.

$$\frac{x : A \vdash_{\mathsf{S}} \mathtt{go}\, x : B, A}{\vdash_{\mathsf{S}} \lambda_{\mathsf{S}} x.\mathtt{go}\, x : A \to B, A}$$

If we read the comma as "or", and $A \to B$ for arbitrary B as "not A", then this judgement asserts the classical excluded middle, "not A or A". From a slightly different perspective, we could say that the A-accepting continuation, by occurring under the λ, becomes an upward continuation (a continuation which is part of the result of an expression).

We build on the classical type of $\lambda_{\mathsf{S}} x.\mathtt{go}\, x$ for another canonical example: Scheme's `call-with-current-continuation` (`call/cc` for short) operator [5]. It is syntactic sugar in terms of static `here` and `go`:

$$\mathtt{call/cc} = \lambda_{\mathsf{S}} f.(\mathtt{here}\, (f\, (\lambda_{\mathsf{S}} x.\mathtt{go}\, x))).$$

As one would expect [4], the type of `call/cc` is Peirce's law "if not A implies A, then A". We derive the judgement

$$\vdash_{\mathsf{S}} \lambda_{\mathsf{S}} f.(\mathtt{here}\, (f\, (\lambda_{\mathsf{S}} x.\mathtt{go}\, x))) : ((A \to B) \to A) \to A, C$$

as follows. Let Γ be the context $f : (A \to B) \to A$. Then we derive:

$$\dfrac{\dfrac{}{\Gamma \vdash_{\mathsf{s}} f : (A \to B) \to A, A} \qquad \dfrac{\dfrac{\dfrac{}{\Gamma, x : A \vdash_{\mathsf{s}} x : A, A}}{\Gamma, x : A \vdash_{\mathsf{s}} \mathtt{go}\, x : B, A}}{\Gamma \vdash_{\mathsf{s}} \lambda_{\mathsf{s}} x.\mathtt{go}\, x : A \to B, A}}{\dfrac{\Gamma \vdash_{\mathsf{s}} (f\, (\lambda_{\mathsf{s}} x.\mathtt{go}\, x)) : A, A}{\dfrac{\Gamma \vdash_{\mathsf{s}} \mathtt{here}\, (f\, (\lambda_{\mathsf{s}} x.\mathtt{go}\, x)) : A, C}{\vdash_{\mathsf{s}} \lambda_{\mathsf{s}} f.(\mathtt{here}\, (f\, (\lambda_{\mathsf{s}} x.\mathtt{go}\, x))) : ((A \to B) \to A) \to A, C}}}$$

As another example, let Γ be any context, and assume we have $\Gamma \vdash_{\mathsf{s}} M : A, B$. Right exchange is derivable in that we can also derive $\Gamma \vdash_{\mathsf{s}} M' : B, A$ for some M'. Concretely,

$$M' = (\lambda_{\mathsf{s}} f.\mathtt{here}\, (f\, M))\, (\lambda_{\mathsf{s}} x.\mathtt{go}\, x)$$

Note that the `go` is outside the scope of the `here`.

In the typing of `call/cc`, a `go` is (at least potentially, depending on f) exported from its enclosing `here`. Conversely, in the derivation of right exchange, a `go` is imported into a `here`-construct from the outside of its scope. What makes everything work is the static binding of continuations. (If we were to define an operational semantics for the static version, we would need to make sure that a λ_{s}-abstraction builds a closure when it is evaluated.)

4. Dynamic semantics and exceptions

Next we consider the dynamic version of `here` and `go`. The word "dynamic" is used here in the sense of dynamic binding and dynamic control as found in many dialects of Lisp (such as Common Lisp or Emacs Lisp). Another way of phrasing it is that with a dynamic semantics, the `here` that is in scope at the point where a function is *called* will be used, as opposed to the `here` that was in scope at the point where the function was *defined*—the latter being used for the static semantics.

In the dynamic case, the transform is this:

$$\begin{aligned}
[\![x]\!] &= \lambda kq.kx \\
[\![\lambda_{\mathsf{d}} x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{d}) \\
[\![MN]\!] &= \lambda kq.[\![M]\!](\lambda m.[\![N]\!](\lambda n.mnkq)q)q \\
[\![\mathtt{here}\, M]\!] &= \lambda kq.[\![M]\!]kk \\
[\![\mathtt{go}\, M]\!] &= \lambda kq.[\![M]\!]qq
\end{aligned}$$

$$\frac{}{\Gamma, x : A, \Gamma' \vdash_{\mathsf{d}} x : A, C}$$

$$\frac{\Gamma \vdash_{\mathsf{d}} M : B, B}{\Gamma \vdash_{\mathsf{d}} \texttt{here}\, M : B, C} \qquad \frac{\Gamma \vdash_{\mathsf{d}} M : B, B}{\Gamma \vdash_{\mathsf{d}} \texttt{go}\, M : C, B}$$

$$\frac{\Gamma, x : A \vdash_{\mathsf{d}} M : B, C}{\Gamma \vdash_{\mathsf{d}} \lambda_{\mathsf{d}} x.M : A \to B \vee C, D} \qquad \frac{\Gamma \vdash_{\mathsf{d}} M : A \to B \vee C, C \quad \Gamma \vdash_{\mathsf{d}} N : A, C}{\Gamma \vdash_{\mathsf{d}} MN : B, C}$$

Figure 2. Typing for dynamic `here` and `go`

In this transform, the jump continuation q works like an exception handler; since it is passed as an extra argument on each call, the dynamically enclosing handler is chosen. Hence under the dynamic semantics, `here` and `go` become a stripped-down version of Lisp's `catch` and `throw` with only a single catch tag. These `catch` and `throw` operation are themselves a no-frills version of exceptions with only identity handlers. We can think of `here` and `go` as a special case of these more elaborate constructs:

$$\begin{aligned} \texttt{here}\, M &\equiv \texttt{(catch 'tag } M\texttt{)} \\ \texttt{go}\, M &\equiv \texttt{(throw 'tag } M\texttt{)} \end{aligned}$$

Because the additional continuation is administered dynamically, we cannot fit it into our simple typing without annotating the function type. So for dynamic control, we write the function type as $A \to B \vee C$. Syntactically, this should be read as a single operator with the three arguments in mixfix. We regard the type system as a variant of intuitionistic logic in which $\to$ and $\vee$ always have to be introduced or eliminated together.

This annotated arrow can be seen as an idealisation of the Java `throws` clause in method definitions, in that $A \to B \vee C$ could be written as

$$B\texttt{(}A\texttt{) throws } C$$

in a more Java-like syntax. A function of type $A \to B \vee C$ may throw things of type C, so it may only be called inside a `here` with the same type. Our typing for the language with dynamic `here` and `go` is presented in Figure 2.

We do not attempt to idealise the ML way of typing exceptions because ML uses a universal type `exn` for exceptions, in effect allowing a carefully delimited area of untypedness into the language. The typing

of ML exceptions is therefore much less informative than that of checked exceptions.

Note that `here` and `go` are still the same weakening and contraction hybrid as in the static setting. But here their significance is a completely different one because the $\rightarrow$-introduction is coupled with a sort of $\vee$-introduction. To see the difference, recall that in the static setting λ-abstracting over a `go` reifies the jump continuation and thereby, at the type level, gives rise to classical disjunction. This is not possible with the version of λ that gives `go` the dynamic semantics. Consider the following inference:

$$\frac{x : A \vdash_{\mathsf{d}} \mathtt{go}\, x : B, A}{\vdash_{\mathsf{d}} \lambda_{\mathsf{d}} x.\mathtt{go}\, x : A{\rightarrow}B{\vee}A, C}$$

The C-accepting continuation at the point of definition is not accessible to the `go` inside the λ_{d}. Instead, the `go` refers only to the A-accepting continuation that will be available at the point of call. Far from the excluded middle, the type of $\lambda_{\mathsf{d}} x.\mathtt{go}\, x$ is thus "A implies A or B; or anything". Put differently, because of the dynamic behaviour of `go`, the A-accepting continuation cannot become an upward continuation even if the `go` is wrapped into a λ.

In the same vein, as a further illustration how fundamentally different the dynamic `here` and `go` are from the static variety, we revisit the term that, in the static setting, gave rise to `call/cc` with its classical type:

$$\lambda f.\mathtt{here}\,(f\,(\lambda x.\mathtt{go}\, x)).$$

Now in the dynamic case, we can only derive the intuitionistic formula

$$((A{\rightarrow}B{\vee}A){\rightarrow}A{\vee}A){\rightarrow}A{\vee}C$$

as the type of this term.

Let Γ be the context $f : (A{\rightarrow}B{\vee}A){\rightarrow}A{\vee}A$. Then we have:

$$\dfrac{\dfrac{\Gamma \vdash_{\mathsf{d}} f : (A{\rightarrow}B{\vee}A){\rightarrow}A{\vee}A, A \qquad \dfrac{\dfrac{}{\Gamma, x : A \vdash_{\mathsf{d}} x : A, A}}{\Gamma, x : A \vdash_{\mathsf{d}} \mathtt{go}\, x : B, A}}{\Gamma \vdash_{\mathsf{d}} \lambda_{\mathsf{d}} x.\mathtt{go}\, x : A{\rightarrow}B{\vee}A, A}}{\dfrac{\Gamma \vdash_{\mathsf{d}} (f\,(\lambda_{\mathsf{d}} x.\mathtt{go}\, x)) : A, A}{\Gamma \vdash_{\mathsf{d}} \mathtt{here}\,(f\,(\lambda_{\mathsf{d}} x.\mathtt{go}\, x)) : A, C}}{\vdash_{\mathsf{d}} \lambda_{\mathsf{d}} f.\mathtt{here}\,(f\,(\lambda_{\mathsf{d}} x.\mathtt{go}\, x)) : ((A{\rightarrow}B{\vee}A){\rightarrow}A{\vee}A){\rightarrow}A{\vee}C, D}$$

The type system given by $\vdash_{\mathsf{d}}$ is intuitionistic in the sense that the rules of $\vdash_{\mathsf{d}}$ correspond to derivations in the $\rightarrow,\vee$-fragment of intuitionistic logic. For instance, the $\vdash_{\mathsf{d}}$-rule (simultaneous $\rightarrow$- and $\vee$-introduction)

corresponds to the intuitionistic derivation ($\vee$-introduction first, then $\rightarrow$-introduction, then right weakening) displayed on its right here:

$$\frac{\Gamma, A \vdash_{\mathsf{d}} B, C}{\Gamma \vdash_{\mathsf{d}} A \rightarrow B \vee C, D} \qquad \frac{\dfrac{\dfrac{\Gamma, A \vdash_{\mathsf{i}} B, C}{\Gamma, A \vdash_{\mathsf{i}} B \vee C}}{\Gamma \vdash_{\mathsf{i}} A \rightarrow (B \vee C)}}{\Gamma \vdash_{\mathsf{i}} A \rightarrow (B \vee C), D}$$

We could also use intuitionistic logic with a single formula on the right by disjoining the two formulas from $\vdash_{\mathsf{d}}$, so that $\Gamma \vdash_{\mathsf{d}} A, B$ implies $\Gamma \vdash_{\mathsf{i}} A \vee B$:

$$\frac{\Gamma, A \vdash_{\mathsf{d}} B, C}{\Gamma \vdash_{\mathsf{d}} A \rightarrow B \vee C, D} \qquad \frac{\dfrac{\Gamma, A \vdash_{\mathsf{i}} B \vee C}{\Gamma \vdash_{\mathsf{i}} A \rightarrow (B \vee C)}}{\Gamma \vdash_{\mathsf{i}} (A \rightarrow (B \vee C)) \vee D}$$

At the level of terms, this corresponds to an exception-passing-style transform in which the additional disjunct may hold an "exceptional" value, which is propagated until handled. If the type of the exceptions is always the same, say E, the transform is given by the exception monad $(\_) + E$ [9]. In our setting, however, we do not have such a fixed type E, as the `here`-construct can change that type. Thus the double-barrelled approach to exceptions taken here may correspond to a more complex structure, such as perhaps an indexed monad.

5. Return continuation

Our last choice is passing the return continuation as the extra continuation to the body of a λ-abstraction. So the CPS transform is this:

$$\begin{aligned}
[\![x]\!] &= \lambda kq.kx \\
[\![\lambda_{\mathsf{r}} x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{r}) \\
[\![MN]\!] &= \lambda kq.[\![M]\!](\lambda m.[\![N]\!](\lambda n.mnkq)q)q \\
[\![\mathtt{here}\, M]\!] &= \lambda kq.[\![M]\!]kk \\
[\![\mathsf{go}\, M]\!] &= \lambda kq.[\![M]\!]qq
\end{aligned}$$

This transform grants λ_{r} the additional role of a continuation binder. The original operator for this purpose, `here`, is rendered redundant, since $\mathtt{here}\, M$ is now equivalent to $(\lambda_{\mathsf{r}} x.M)(\lambda_{\mathsf{r}} y.y)$ where x is not free in M. At first sight, binding continuations seems an unusual job for a

$$\frac{}{\Gamma, x : A, \Gamma' \vdash_{\mathsf{r}} x : A, C} \qquad \frac{\Gamma \vdash_{\mathsf{r}} M : B, B}{\Gamma \vdash_{\mathsf{r}} \mathtt{go}\, M : C, B}$$

$$\frac{\Gamma, x : A \vdash_{\mathsf{r}} M : B, B}{\Gamma \vdash_{\mathsf{r}} \lambda_{\mathsf{r}} x.M : A \to B, C} \qquad \frac{\Gamma \vdash_{\mathsf{r}} M : A \to B, C \quad \Gamma \vdash_{\mathsf{r}} N : A, C}{\Gamma \vdash_{\mathsf{r}} MN : B, C}$$

Figure 3. Typing for `go` as a `return`-operation

λ; but it becomes less so if we think of `go` as a `return` statement like those of C or Java.

5.1. Second-class `return`

Because the enclosing λ determines which continuation `go` jumps to with its argument, the `go`-operator has the same effect as a `return` statement. The type of extra continuation assumed by `go` needs to agree with the return type of the nearest enclosing λ:

$$\frac{\Gamma, x : A \vdash_{\mathsf{r}} M : B, B}{\Gamma \vdash_{\mathsf{r}} \lambda_{\mathsf{r}} x.M : A \to B, C}$$

The whole type system for the calculus with λ_{r} is in Figure 3.

The agreement between `go` and the enclosing λ_{r} is comparable with the typing in C, where the expression in a `return` statement must have the return type declared by the enclosing function. For instance, M needs to have type `int` in the definition:

$$\mathtt{int\ f()\{} \ldots \mathtt{return}\ M\mathtt{;} \ldots \mathtt{\}}$$

With λ_{r}, the special form `go` cannot be made into a first-class function. If we try to λ-abstract over $\mathtt{go}\, x$ by writing $\lambda_{\mathsf{r}} x.\mathtt{go}\, x$ then `go` will refer to that λ_{r}.

The failure of λ_{r} to give first-class returning can be seen logically as follows. In order for λ_{r} to be introduced, both types on the right have to be the same:

$$\frac{x : A \vdash_{\mathsf{r}} \mathtt{go}\, x : A, A}{\vdash_{\mathsf{r}} \lambda_{\mathsf{r}} x.\mathtt{go}\, x : A \to A, C}$$

Rather than the classical "not A or A" this asserts merely the intuitionistic "A implies A; or anything".

One has a similar situation in Gnu C, which has both the `return` statement and nested functions, without the ability to refer to the re-

turn address of another function. If we admit go as a first-class function, it becomes a much more powerful form of control, Landin's **JI**-operator.

5.2. The **JI**-operator

Keeping the meaning of λ_{r} as a continuation binder, we now consider a control operator **JI** that always refers to the statically enclosing λ_{r}, but which, unlike the special form go, is a first-class expression, so that we can pass the return continuation to some other function f by writing $f(\mathbf{JI})$. This operator is transformed into CPS as follows:

$$[\![\mathbf{JI}]\!] = \lambda ks.k(\lambda xrd.\boxed{s}x)$$

That is almost, but not quite, the same as if we tried to define **JI** as $\lambda_{\mathsf{r}}x.\mathsf{go}\,x$:

$$\begin{aligned}[\![\mathbf{JI}]\!] &= [\![\lambda_{\mathsf{r}}x.\mathsf{go}\,x]\!] \\ &= \lambda ks.k(\lambda xrd.\boxed{r}x)\end{aligned}$$

We can, however, define **JI** in terms of go if we use the static λ_{s}, that is $\mathbf{JI} = \lambda_{\mathsf{s}}x.\mathsf{go}\,x$, as this does not inadvertently shadow the continuation s that we want **JI** to refer to.

The whole transform for the calculus with **JI** is this:

$$\begin{aligned}[\![x]\!] &= \lambda kq.kx \\ [\![\lambda_{\mathsf{r}}x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{r}) \\ [\![MN]\!] &= \lambda kq.[\![M]\!](\lambda m.[\![N]\!](\lambda n.mnkq)q)q \\ [\![\mathbf{JI}]\!] &= \lambda ks.k(\lambda xrd.\boxed{s}x)\end{aligned}$$

Recall that the role of here has been taken over by λ_{r}, and we replaced go by its first-class cousin **JI**.

In the transform for **JI**, the jump continuation is the current "dump" in the sense of the SECD-machine. The dump in the SECD-machine is a sort of call stack, which holds the return continuation for the procedure whose body is currently being evaluated. Making the dump into a first-class object was precisely how Landin invented first-class control, embodied by the **J**-operator.

The typing for the language with **JI** is given in Figure 4. In particular, the type of **JI** is the classical disjunction

$$\frac{}{\Gamma \vdash_{\mathsf{J}} \mathbf{JI} : B \to C, B}$$

The operator **JI** by itself (without even being applied to an argument) yields an upward continuation in that it wraps the B-accepting continuation to the right of the comma into a non-returning function of type $B \to C$.

$$\frac{}{\Gamma, x : A, \Gamma' \vdash_{\mathsf{j}} x : A, C} \qquad \frac{}{\Gamma \vdash_{\mathsf{j}} \mathbf{JI} : B \to C, B}$$

$$\frac{\Gamma, x : A \vdash_{\mathsf{j}} M : B, B}{\Gamma \vdash_{\mathsf{j}} \lambda_{\mathsf{r}} x.M : A \to B, C} \qquad \frac{\Gamma \vdash_{\mathsf{j}} M : A \to B, C \quad \Gamma \vdash_{\mathsf{j}} N : A, C}{\Gamma \vdash_{\mathsf{j}} MN : B, C}$$

Figure 4. Typing for **JI**

As an example of the type system for the calculus with the **JI**-operator, we see that Reynolds's [12, 13] definition of `call/cc` in terms of **JI** typechecks. (Strictly speaking, Reynolds used `escape`, the binding-form cousin of `call/cc`, but `call/cc` and `escape` are syntactic sugar for each other.) We infer the type of $\texttt{call/cc} \equiv \lambda_{\mathsf{r}} f.((\lambda_{\mathsf{r}} k.f\ k)(\mathbf{JI}))$ to be:

$$((A \to B) \to A) \to A)$$

To write the derivation, we abbreviate some contexts as follows:

$$\begin{aligned} \Gamma_{fk} &\equiv f : (A \to B) \to A, k : (A \to B) \\ \Gamma_{f} &\equiv f : (A \to B) \to A \end{aligned}$$

Then we can derive:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\Gamma_{fk} \vdash_{\mathsf{j}} f : (A \to B) \to A, A} \quad \dfrac{}{\Gamma_{fk} \vdash_{\mathsf{j}} k : (A \to B), A}}{\Gamma_{fk} \vdash_{\mathsf{j}} f\ k : A, A}}{\Gamma_{f} \vdash_{\mathsf{j}} \lambda_{\mathsf{r}} k.fk : (A \to B) \to A, A} \qquad \dfrac{}{\Gamma_{f} \vdash_{\mathsf{j}} \mathbf{JI} : A \to B, A}}{\Gamma_{f} \vdash_{\mathsf{j}} (\lambda_{\mathsf{r}} k.f\ k)(\mathbf{JI}) : A, A}}{\vdash_{\mathsf{j}} \lambda_{\mathsf{r}} f.((\lambda_{\mathsf{r}} k.f\ k)(\mathbf{JI})) : ((A \to B) \to A) \to A), C}$$

Because **JI** has such evident logical meaning as classical disjunction, we have considered it as basic. Landin [7] took another operator, called **J**, as primitive, while **JI** was derived as the special case of **J** applied to the identity combinator:

$$\mathbf{J}\,\mathbf{I} = \mathbf{J}\,(\lambda x.x)$$

This explains the name "**JI**", as "**J**" stands for "jump" and **I** for "identity". We were able to start with **JI**, since (as noted by Landin) the **J**-operator is syntactic sugar for **JI** by virtue of:

$$\mathbf{J} = (\lambda_{\mathsf{r}} r.\lambda_{\mathsf{r}} f.\lambda_{\mathsf{r}} x.r(fx))\,(\mathbf{JI}).$$

To accommodate **J** in our typing, we use this definition in terms of **JI** to derive the following type for **J**:

$$\vdash_{\mathsf{J}} \mathbf{J} : (A \to B) \to (A \to C), B$$

Let Γ be the context $x : A, r : B \to C, f : A \to B$. We derive:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\Gamma \vdash_{\mathsf{J}} r : B \to C, C} \quad \dfrac{\dfrac{}{\Gamma \vdash_{\mathsf{J}} f : A \to B, C} \quad \dfrac{}{\Gamma \vdash_{\mathsf{J}} x : A, C}}{\Gamma \vdash_{\mathsf{J}} fx : B, C}}{\Gamma \vdash_{\mathsf{J}} r(fx) : C, C}}{r : B \to C, f : A \to B \vdash_{\mathsf{J}} \lambda_{\mathsf{r}} x.r(fx) : A \to C, A \to C}}{r : B \to C \vdash_{\mathsf{J}} \lambda_{\mathsf{r}} f.\lambda_{\mathsf{r}} x.r(fx) : (A \to B) \to (A \to C), (A \to B) \to (A \to C)}}{\vdash_{\mathsf{J}} \lambda_{\mathsf{r}} r.\lambda_{\mathsf{r}} f.\lambda_{\mathsf{r}} x.r(fx) : (B \to C) \to (A \to B) \to (A \to C), B}}{\vdash_{\mathsf{J}} (\lambda_{\mathsf{r}} r.\lambda_{\mathsf{r}} f.\lambda_{\mathsf{r}} x.r(fx))\,(\mathbf{JI}) : (A \to B) \to (A \to C), B}$$

This type reflects the behaviour of the **J**-operator in the SECD-machine. When **J** is evaluated, it captures the B-accepting current dump continuation; it can then be applied to a function of type $A \to B$. This function is composed with the captured dump, yielding a non-returning function of type $A \to C$, for arbitrary C. By analogy with `call-with-current-continuation`, we may read the **J**-operator as "`compose-with-current-dump`" [18].

The logical significance, if any, of the extra function types in the general **J** seems unclear. There is a curious, though vague, resemblance to exception handlers in dynamic control, since they too are functions only to be applied on jumping. This feature of **J** may be historical, as it arose in a context where greater emphasis was given to attaching dumps to functions than to dumps as first-class continuations in their own right.

6. Type preservation

The typings agree with the transforms in that they are preserved in the usual way for CPS transforms: we have a "double-negation" transform for types, contexts and judgements. The only slight complication is in typing the dynamic continuation in those transforms that ignore it.

We assume some given answer type $\mathbb{A}$ for continuations. The function type of the form $A \to B \vee C$ for the dynamic semantics is translated as follows:

$$[\![A \to B \vee C]\!] \;=\; [\![A]\!] \to ([\![B]\!] \to \mathbb{A}) \to ([\![C]\!] \to \mathbb{A}) \to \mathbb{A}$$

Each call expects not only the B-accepting return continuation, but also the C-accepting continuation determined by the `here` that encloses the call.

Because we have not varied the transform of application, functions defined with λ_{s} and λ_{r} are also passed this dynamic continuation, even though they ignore it:

$$\begin{aligned}[\![\lambda_{\mathsf{s}} x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{s}) \\ [\![\lambda_{\mathsf{r}} x.M]\!] &= \lambda ks.k(\lambda xrd.[\![M]\!]r\boxed{r})\end{aligned}$$

In both of these cases, the dynamic jump continuation d is fed to each function call, but never needed. Each function definition must expect this argument to be of certain type. Because different calls of the same function may have dynamically enclosing `here` operators with different types, the type ascribed to d should be polymorphic.

The function type of the form $A \to B$ is transformed so as to accept this unwanted argument polymorphically:

$$[\![A \to B]\!] = [\![A]\!] \to ([\![B]\!] \to \mathbb{A}) \to \forall\beta.\beta \to \mathbb{A}$$

That is, a function of type $A \to B$ accepts an argument of type A, a B-accepting return continuation, and the continuation determined by the `here` dynamically enclosing the call.

We will use Curry-style polymorphism in our target language for the CPS transform. ("Curry-style" means that there are no type abstractions and applications in the terms, so that we do not have to add anything to the CPS transforms from Section 2). It is given by the following two rules:

$$\frac{\Gamma \vdash M : \forall\alpha.A}{\Gamma \vdash M : A[\alpha \mapsto B]} \qquad \frac{\Gamma \vdash M : A}{\Gamma \vdash M : \forall\alpha.A}\ \alpha \text{ not free in } \Gamma$$

For all the transforms we have preservation of the respective typing: if $\Gamma \vdash_? M : A, B$ in the source, then in the target of the CPS transform we have

$$[\![\Gamma]\!] \vdash [\![M]\!] : ([\![A]\!] \to \mathbb{A}) \to ([\![B]\!] \to \mathbb{A}) \to \mathbb{A}.$$

The proof is a straightforward induction over the derivation; we sketch some representative cases below.

As a typical example, consider how the classical axiom of excluded middle

$$\vdash_{\mathsf{J}} \mathbf{JI} : A \to B, A$$

is translated to an intuitionistic proof $[\![\mathbf{JI}]\!] = \lambda ks.k(\lambda xrd.sx)$ of the formula

$$(([\![A]\!] \to ([\![B]\!] \to \mathbb{A}) \to \forall\beta.\beta \to \mathbb{A}) \to \mathbb{A}) \to ([\![A]\!] \to \mathbb{A}) \to \mathbb{A}$$

in the target.

6.1. The dynamic continuation

We show the type preservation in some more detail for the rule for λ-abstraction in the dynamic case:

$$\frac{\Gamma, x : A \vdash_{\mathsf{d}} M : B, C}{\Gamma \vdash_{\mathsf{d}} \lambda_{\mathsf{d}} x.M : A \to B \vee C, D}$$

By the induction hypothesis, we conclude from $\Gamma, x : A \vdash_{\mathsf{d}} M : B, C$ that

$$[\![\Gamma]\!], x : [\![A]\!] \vdash [\![M]\!] : ([\![B]\!] \to \mathbb{A}) \to ([\![C]\!] \to \mathbb{A}) \to \mathbb{A}$$

By weakening we also have

$$[\![\Gamma]\!], k, s, x, k', d \vdash [\![M]\!] : ([\![B]\!] \to \mathbb{A}) \to ([\![C]\!] \to \mathbb{A}) \to \mathbb{A}$$

Hence

$$[\![\Gamma]\!], k : [\![A \to B \vee C]\!] \to \mathbb{A}, s : [\![D]\!] \to \mathbb{A} \vdash \lambda x k' d.[\![M]\!] k' d : [\![A \to B \vee C]\!].$$

Thus

$$[\![\Gamma]\!] \vdash [\![\lambda_{\mathsf{d}} x.M]\!] : ([\![A \to B \vee C]\!] \to \mathbb{A}) \to ([\![D]\!] \to \mathbb{A}) \to \mathbb{A}$$

as required.

6.2. Ignoring the dynamic continuation polymorphically

For those transforms that ignore the dynamic jump continuation, we need to introduce polymorphism in the case of λ-abstraction. Consider the static λ-abstraction:

$$\frac{\Gamma, x : A \vdash_{\mathsf{s}} M : B, C}{\Gamma \vdash_{\mathsf{s}} \lambda_{\mathsf{s}} x.M : A \to B, C}$$

By the induction hypothesis, we have

$$[\![\Gamma]\!], x : [\![A]\!] \vdash [\![M]\!] : ([\![B]\!] \to \mathbb{A}) \to ([\![C]\!] \to \mathbb{A}) \to \mathbb{A}$$

Hence

$$[\![\Gamma]\!], k : [\![A \to B]\!] \to \mathbb{A}, s : [\![C]\!] \to \mathbb{A} \vdash \lambda x k' d.[\![M]\!] k' s : [\![A \to B]\!].$$

Thus $[\![\Gamma]\!] \vdash [\![\lambda_{\mathsf{s}} x.M]\!] : ([\![A \to B]\!] \to \mathbb{A}) \to ([\![C]\!] \to \mathbb{A}) \to \mathbb{A}$, as required.

While λ-abstraction abstracts over a type variable, application instantiates it. Consider the rule

$$\frac{\Gamma \vdash_{\mathsf{S}} M : A \to B, C \qquad \Gamma \vdash_{\mathsf{S}} N : A, C}{\Gamma \vdash_{\mathsf{S}} MN : B, C}$$

By the induction hypothesis, we assume

$$\begin{aligned}[\![\Gamma]\!] &\vdash [\![M]\!] : ([\![A \to B]\!] \to \mathbb{A}) \to ([\![C]\!] \to \mathbb{A}) \to \mathbb{A} \\ [\![\Gamma]\!] &\vdash [\![N]\!] : ([\![A]\!] \to \mathbb{A}) \to ([\![C]\!] \to \mathbb{A}) \to \mathbb{A}\end{aligned}$$

We have to show that

$$[\![\Gamma]\!] \vdash [\![MN]\!] : ([\![B]\!]) \to \mathbb{A}) \to ([\![C]\!] \to \mathbb{A}) \to \mathbb{A}$$

where

$$[\![MN]\!] = \lambda kq.[\![M]\!](\lambda m.[\![N]\!](\lambda n.mnkq)q)q$$

The crucial step is to instantiate the type of the ignored dynamic jump continuation argument to that of q:

$$\frac{\dfrac{[\![\Gamma]\!], m, n, k, q \vdash mnk : \forall\beta.\beta \to \mathbb{A}}{[\![\Gamma]\!], m, n, k, q \vdash mnk : ([\![C]\!] \to \mathbb{A}) \to \mathbb{A}} \qquad [\![\Gamma]\!], m, n, k, q \vdash q : [\![C]\!] \to \mathbb{A}}{[\![\Gamma]\!], m, n, k, q \vdash mnkq : \mathbb{A}}$$

7. Double-barrelled CPS and linearly used continuations

In a companion paper [1] we have shown that a wide variety of control constructs use continuations *linearly*. That paper also uses some double-barrelled CPS transforms for the sake of simplicity—some similar to the ones used here, others very different. We refer the reader to it for details and further motivation on linearly used continuations. In this section, we only sketch the connection between linear and non-linear continuation use on the one hand and the contrast between classical and intuitionistic typing for control on the other.

To formalise linear use of continuations, we refine the target language of our CPS transforms with linear functions in Figure 5 (for details of this typing, we refer the reader to the companion paper [1]). This type system uses both a linear and an intuitionistic zone. The former will in fact only contain continuations, as it is their usage that we want to restrict.

To bring out the similarity with the CPS transforms in the previous section, it is convenient to introduce a pattern-matching syntax

$$\frac{}{\Gamma, x : A; \_ \vdash x : A} \qquad \frac{}{\Gamma; x : P \vdash x : P}$$

$$\frac{\Gamma; \Delta, x : P \vdash M : Q}{\Gamma; \Delta \vdash \delta x.\, M : P \multimap Q} \qquad \frac{\Gamma; \Delta_1 \vdash M : P \multimap Q \quad \Gamma; \Delta_2 \vdash N : P}{\Gamma; \Delta_1, \Delta_2 \vdash M \sqcup N : Q}$$

$$\frac{\Gamma, x : A; \Delta \vdash M : P}{\Gamma; \Delta \vdash \lambda x.\, M : A \to P} \qquad \frac{\Gamma; \Delta \vdash M : A \to P \quad \Gamma; \_ \vdash N : A}{\Gamma; \Delta \vdash M\, N : P}$$

$$\frac{\Gamma; \Delta \vdash M : P \quad \Gamma; \Delta \vdash N : Q}{\Gamma; \Delta \vdash \langle M, N \rangle : P \& Q} \qquad \frac{\Gamma; \Delta \vdash M : P_1 \& P_2}{\Gamma; \Delta \vdash \pi_i \sqcup M : P_i}$$

$$\frac{\Gamma; \Delta \vdash M : A}{\Gamma; \Delta \vdash M : \forall \alpha. A}\ \alpha \notin \Gamma; \Delta \qquad \frac{\Gamma; \Delta \vdash M : \forall \alpha. A}{\Gamma; \Delta \vdash M : A[\alpha \mapsto B]}$$

Figure 5. Target language with linear typing

$\delta\langle x_1, x_2\rangle.M$ as syntactic sugar for $\delta p.M[x \mapsto \pi_1 \sqcup p][x_2 \mapsto \pi_2 \sqcup p]$. With this notation, we write a double-barrelled CPS transform that uses both continuation arguments together linearly:

$$\begin{aligned}
[\![x]\!] &= \delta\langle k, q\rangle.kx \\
[\![\lambda_{\mathsf{d}} x.M]\!] &= \delta\langle k, s\rangle.k(\lambda x.\delta\langle r, d\rangle.[\![M]\!] \sqcup \langle r, d\rangle) \\
[\![\lambda_{\mathsf{r}} x.M]\!] &= \delta\langle k, s\rangle.k(\lambda x.\delta\langle r, d\rangle.[\![M]\!] \sqcup \langle r, r\rangle) \\
[\![MN]\!] &= \delta\langle k, q\rangle.[\![M]\!] \sqcup \langle \lambda m.[\![N]\!] \sqcup \langle \lambda n.(mn) \sqcup \langle k, q\rangle, q\rangle, q\rangle \\
[\![\mathtt{here}\, M]\!] &= \delta\langle k, q\rangle.[\![M]\!] \sqcup \langle k, k\rangle \\
[\![\mathtt{go}\, M]\!] &= \delta\langle k, q\rangle.[\![M]\!] \sqcup \langle q, q\rangle
\end{aligned}$$

It is tempting to call this double-barrel, one-shot continuation passing; but one needs to bear in mind that there is one shot for both barrels combined.

This transform works for the dynamic, exception-like semantics from Section 4 and for the `return`-operation from Section 5. The function types need to be refined as follows with linear typing:

$$\begin{aligned}
[\![A \to B \vee C]\!] &= [\![A]\!] \to (([\![B]\!] \to \mathbb{A}) \& ([\![C]\!] \to \mathbb{A})) \multimap \mathbb{A} \\
[\![A \to B]\!] &= [\![A]\!] \to \forall\beta.(([\![B]\!] \to \mathbb{A}) \& \beta) \multimap \mathbb{A}
\end{aligned}$$

By contrast, the static λ as in Section 3 does not allow this linear typing. The following fails because of ill-typed sharing between operator

and operand:

$$\not\vdash [\![\lambda_{\mathsf{S}} x.M]\!] = \delta\langle k, s\rangle.k(\lambda x.\delta\langle r, d\rangle.[\![M]\!]\_\langle r, s\rangle)$$

This becomes clearer if we unsugar the $\delta\langle k, s\rangle$ binding into that of a continuation pair p, where $k = \pi_1 \_ p$ and $s = \pi_2 \_ p$:

$$\not\vdash [\![\lambda_{\mathsf{S}} x.M]\!] = \delta p.(\pi_1 \_ p)(\lambda x.\delta p'.[\![M]\!]\_\langle(\pi_1 \_ p'), (\pi_2 \_ p)\rangle)$$

The **JI**-operator fails for the same reason:

$$\not\vdash [\![\mathbf{JI}]\!] = \delta\langle k, s\rangle.k(\lambda x.\delta\langle r, d\rangle.sx)$$

Those constructs that have an intuitionistic typing at the source admit a typing on the *target* that restricts the use of continuation to be *linear*. Those rules whose addition causes the source language typing to become classical break this linearity of continuation use, forcing the target language typing to become intuitionistic (no longer restricted to linearity) in the use of continuations. In sum, the source and target typings of our four little languages are as follows:

Construct	Source language	Use of continuations in target
Static `here`/`go`	Classical	Intuitionistic
Dynamic `here`/`go`	Intuitionistic	Linear
`return`-operation	Intuitionistic	Linear
JI-operator	Classical	Intuitionistic

8. Conclusions

As logical systems, the typings of the four control operations we have considered may seem a little eccentric, with two succedents that can only be manipulated in a slightly roundabout way. But they are sufficient for our purposes here, which is to illustrate the correspondence of first-class continuations with classical logic and weaker control operation with intuitionistic logic, and the central role of the arrow type in this dichotomy.

Recall the following fact from proof theory (see for example the textbooks by Troelstra and Schwichtenberg [21, Exercise 3.2.1A on page 67] or Troelstra and van Dalen [22, Exercise 10.7.6 on page 568]).

Suppose one starts from a presentation of intuitionistic logic with sequents of the form $\Gamma \vdash \Delta$. If a rule like the following is added that

allows $\rightarrow$-introduction even if there are multiple succedents, the logic becomes classical.

$$\frac{\Gamma, A \vdash B, \Delta}{\Gamma \vdash A \rightarrow B, \Delta}$$

In continuation terms, the significance of this rule is that the function closure of type $A \rightarrow B$ may contain any of the continuations that appear in Δ; to use the jargon, these continuations become "reified". The fact that the logic becomes classical means that once we can have continuations in function closures, we gain first-class continuations and thereby the same power as `call/cc`. We have this form of rule for static `here` and `go`; though not for **JI**, since **JI** as the excluded middle is already blatantly classical by itself.

But the logic remains intuitionistic if the $\rightarrow$-introduction is restricted. The rule for this case typically admits only a single formula on the right:

$$\frac{\Gamma, A \vdash B}{\Gamma \vdash A \rightarrow B, \Delta}$$

Considered as a restriction on control operators, this rule prohibits λ-abstraction for terms that contain free continuation variables. There are clearly other possibilities how we can prevent assumptions from Δ to become hidden (in that they can be used in the derivation of $A \rightarrow B$ without showing up in this type itself). We could require these assumptions to remain explicit in the arrow type, by making Δ a singleton that either coincides with the B on the right of the arrow, or is added to it:

$$\frac{\Gamma, A \vdash_{\mathsf{r}} B, B}{\Gamma \vdash_{\mathsf{r}} A \rightarrow B, C} \qquad \frac{\Gamma, A \vdash_{\mathsf{d}} B, C}{\Gamma \vdash_{\mathsf{d}} A \rightarrow B \vee C, D}$$

These are the rules for $\rightarrow$-introduction in connection with the `return`-operation, and dynamic `here` and `go`, respectively. Neither of which gives rise to first-class continuations, corresponding to the fact that with these restrictions on $\rightarrow$-introduction the logics remain intuitionistic.

When the double-barrelled typing is intuitionistic, we can read the comma on the right as an intuitionistic disjunction in the sense that the term produces a result of either the one or the other type, rather like the disjunctive property in intuitionistic logic [21, Theorem 4.2.3]. Moreover, on the level of the target of the CPS transform, this means that the two continuations are joined by an & and are jointly used linearly, so that we can never use both.

The distinction between static and dynamic control in logical terms appears to be new, as is the logical explanation of Landin's **JI**-operator.

It would be natural to add an empty type $\bot$, whose logical meaning is falsity. Then $\Gamma \vdash_? M : A, \bot$ in the double-barrelled systems would correspond to ordinary judgements $\Gamma \vdash M : A$ for intuitionistic or classical logic. At the top level, one could restrict to such judgements of the form $\Gamma \vdash_? M : A, \bot$. Moreover, for the annotated function types, we could then express that a function cannot raise exceptions if it has a type of the form $A \to B \vee \bot$.

Related work

Following Griffin [4], there has been a great deal of work on classical types for control operators, mainly on `call/cc` or minor variants thereof. A similar CPS transform for dynamic control (exceptions) has been used by Kim, Yi and Danvy [6], albeit for a very different purpose. Felleisen describes the **J**-operator by way of a CPS transform, but since his transform is not double-barrelled, **J** means something different in each λ-abstraction [2]. Variants of the `here` and `go` operators are even older than the notion of continuation itself: the operations `valof` and `resultis` from CPL later appeared in Strachey and Wadsworth's report on continuations [16, 17]. These operators led to the modern `return` in C. As we have shown here, they lead to much else besides if combined with different flavours of λ.

Acknowledgements

Thanks to Peter O'Hearn, Olivier Danvy and the anonymous referees.

References

1. J. Berdine, P. W. O'Hearn, U. Reddy, and H. Thielecke. Linearly used continuations. In A. Sabry, editor, *Proceedings of the 3rd ACM SIGPLAN Workshop on Continuations*, Indiana University Technical Report No 545, pages 47–54, 2001.
2. M. Felleisen. Reflections on Landin's J operator: a partly historical note. *Computer Languages*, 12(3/4):197–207, 1987.
3. D. P. Friedman, M. Wand, and C. T. Haynes. *Essentials of Programming Languages*. MIT Press, 1992.
4. T. G. Griffin. A formulae-as-types notion of control. In *Proc. 17th ACM Symposium on Principles of Programming Languages*, pages 47–58, San Francisco, CA USA, 1990.
5. R. Kelsey, W. Clinger, and J. Rees, editors. Revised^5 report on the algorithmic language Scheme. *Higher-Order and Symbolic Computation*, 11(1):7–105, 1998.
6. J. Kim, K. Yi, and O. Danvy. Assessing the overhead of ML exceptions by selective CPS transformation. In *The 1998 ACM SIGPLAN Workshop on ML*, pages 103–114, 1998.

7. P. J. Landin. A generalization of jumps and labels. Report, UNIVAC Systems Programming Research, Aug. 1965.
8. P. J. Landin. A generalization of jumps and labels. *Higher-Order and Symbolic Computation*, 11(2), 1998. Reprint of [7].
9. E. Moggi. Computational lambda calculus and monads. In *Proceedings, Fourth Annual Symposium on Logic in Computer Science*, pages 14–23, 1989.
10. G. D. Plotkin. Call-by-name, call-by-value, and the λ-calculus. *Theoretical Computer Science*, 1(2):125–159, 1975.
11. J. C. Reynolds. Definitional interpreters for higher-order programming languages. In *Proceedings of the 25th ACM National Conference*, pages 717–740. ACM, Aug. 1972.
12. J. C. Reynolds. Definitional interpreters for higher-order programming languages. *Higher-Order and Symbolic Computation*, 11(4):363–397, 1998. Reprint of a conference paper [11].
13. J. C. Reynolds. Definitional interpreters revisited. *Higher-Order and Symbolic Computation*, 11(4):355–361, 1998.
14. J. G. Riecke and H. Thielecke. Typed exceptions and continuations cannot macro-express each other. In J. Wiedermann, P. van Emde Boas, and M. Nielsen, editors, *Proceedings 26th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 1644 of *LNCS*, pages 635–644. Springer Verlag, 1999.
15. D. A. Schmidt. *Denotational Semantics*. Allyn and Bacon, Boston, 1986.
16. C. Strachey and C. P. Wadsworth. Continuations: A mathematical semantics for handling full jumps. Monograph PRG-11, Oxford University Computing Laboratory, Programming Research Group, Oxford, UK, 1974.
17. C. Strachey and C. P. Wadsworth. Continuations: A mathematical semantics for handling full jumps. *Higher-Order and Symbolic Computation*, 13(1/2):135–152, April 2000. Reprint of a technical report [16].
18. H. Thielecke. An introduction to Landin's "A generalization of jumps and labels". *Higher-Order and Symbolic Computation*, 11(2):117–124, 1998.
19. H. Thielecke. On exceptions versus continuations in the presence of state. In G. Smolka, editor, *Programming Languages and Systems, 9th European Symposium on Programming, ESOP 2000*, number 1782 in LNCS, pages 397–411. Springer Verlag, 2000.
20. H. Thielecke. Comparing control constructs by typing double-barrelled CPS transforms. In A. Sabry, editor, *Proceedings of the 3rd ACM SIGPLAN Workshop on Continuations*, Indiana University Technical Report No 545, pages 17–25, 2001.
21. A. S. Troelstra and H. Schwichtenberg. *Basic Proof Theory*. Cambridge University Press, 2001.
22. A. S. Troelstra and D. van Dalen. *Constructivism in Mathematics*, volume 2. North Holland, 1988.