Continuation models are universal for $\lambda\mu$-calculus

Martin Hofmann and Thomas Streicher*

Abstract

We show that a certain simple call-by-name continuation semantics of Parigot's λ_μ-calculus is complete. More precisely, for every λ_μ-theory we construct a cartesian closed category such that the ensuing continuation-style interpretation of λ_μ, which maps terms to functions sending abstract continuations to responses, is full and faithful.

Thus, any λ_μ-category in the sense of is isomorphic to a continuation model [4] derived from a cartesian-closed category of continuations.

1 Introduction and Summary

Parigot's λ_μ-calculus [7] is a proof term assignment system for propositional classical logic and can at the same time be considered as a prototype for a call-by-name functional programming language incorporating explicit handling of continuations. The original motivation for this calculus was to give a functional interpretation for proofs in classical AF_2—a certain system of second-order arithmetic [3].

Ong [6] has defined a categorical notion of model for this calculus for which the usual categorical completeness theorem holds. In this sense Ong's semantics can be seen as a variable-free reformulation of the syntax of λ_μ. On the other hand, there exists a class of rather concrete continuation models for λ_μ where terms are interpreted as functions[1] mapping abstract continuations to answers. We prove in this paper that every λ_μ-theory (thus every model in the sense of Ong) is induced by a particular continuation model.

A similar result for call-by-value lambda-calculus with control operators has been obtained in [2] by category-theoretic means and independently by Felleisen and Sabry using syntactic back-and-forth translations [10]. The technique we use here is inspired by the method used in [2] in the sense that the morphisms of the continuation category to be constructed arise as special terms of a λ_μ-theory. Whereas in loc. cit. these special terms are defined by their syntactic form we use an equational description involving quantification over all observations.

Unlike in the case of [2] or [10] the equational axiomatisation of λ_μ under consideration was not specially tailored towards completeness for continuation models, which were apparently not known to Parigot at the time, but rather arose from syntactic considerations. For instance, it gives rise to a confluent and strongly normalising rewrite system [7].

The fact that by our result this axiomatisation is complete for continuation models thus provides evidence that these models are a very natural semantics for "proof-relevant" classical logic.

A consequence of our result is that λ_μ-equality without nonlogical axioms can be reduced to equality of terms of simply-typed lambda calculus with products via a certain CPS translation derived from our semantics.

2 The λ_μ-calculus

The presentation of λ_μ we use follows Ong's account in [6]. It differs from Parigot's original formulation only in the aspect that we omit continuation variables of type $\bot$. See loc. cit. for a more detailed comparison.

Assume a set $\mathcal{B}$ of base types. The types of λ_μ are the simple types over $\mathcal{B} \cup \{\bot\}$, i.e. every base type is a type, $\bot$ is a type, and if A, B are types so is $A \Rightarrow B$.

There are two sorts of variables. Object variables ranged over by Roman letters $x, y, z, \ldots$ and continuation variables ranged over by Greek letters $\alpha, \beta, \ldots$. An object context is a partial type assignment to object variables of the form $x_1{:}A_1, \ldots, x_n{:}A_n$. A continuation context is a partial type assignment to continuation variables of the form $\alpha_1{:}A_1, \ldots, \alpha_n{:}A_n$ where all the A_i are different from $\bot$.

Assume a set $\mathcal{C}$ of typed constants. The typing judgements of λ_μ take the form $\Gamma; \Delta \vdash t : A$ where

*TH Darmstadt, FB 4, Schlossgartenstr. 7, 64289 Darmstadt, Germany, phone: +49-6151-162794, fax: +49-6151-164011, e-mail:{mh|streicher}@mathematik.th-darmstadt.de

[1] In the sense of cartesian-closed categories

Γ is an object context, Δ is a continuation context, A is a type, and t is a term. The precise form of the terms is given implicitly together with the following rules defining the typing judgement.

$$(\mathsf{Axiom}) \quad \frac{}{\Gamma;\Delta \vdash x : A} \quad \text{if } x{:}A \in \Gamma$$

$$(\mathsf{Const}) \quad \frac{}{\Gamma;\Delta \vdash c : A} \quad \text{if } c{:}A \in \mathcal{C}$$

$$(\Rightarrow\text{-}\mathsf{intro}) \quad \frac{\Gamma, x{:}A;\Delta \vdash t : B}{\Gamma;\Delta \vdash \lambda x{:}A.t : A \Rightarrow B}$$

$$(\Rightarrow\text{-}\mathsf{elim}) \quad \frac{\Gamma;\Delta \vdash t : A \Rightarrow B \qquad \Gamma;\Delta \vdash s : A}{\Gamma;\Delta \vdash ts : B}$$

$$(\bot\text{-}\mathsf{elim}) \quad \frac{\Gamma;\Delta, \alpha{:}A \vdash t : \bot}{\Gamma;\Delta \vdash \mu\alpha{:}A.t : A} \quad \text{if } A \neq \bot$$

$$(\bot\text{-}\mathsf{intro}) \quad \frac{\Gamma;\Delta \vdash t : A}{\Gamma;\Delta \vdash [\alpha]t : \bot} \quad \text{if } \alpha{:}A \in \Delta$$

As usual we identify terms up to renaming of both object and continuation variables. Notice that λ *and* μ bind variables as indicated, but that continuation variable α occurs free in a term of the form $[\alpha]t$.

The typing rules are such that we have $x_1{:}A_1,\ldots,x_n{:}A_n;\alpha_1{:}B_1,\ldots\alpha_m{:}B_m \vdash t : A$ for some t iff $A_1 \wedge \cdots \wedge A_n \wedge \neg B_1 \wedge \cdots \wedge \neg B_m \to A$ is a tautology of classical propositional logic. Under this analogy the rule ($\bot$-elim) corresponds to proof by contradiction: in order to prove A it suffices to deduce a contradiction ($\bot$) from the assumption that A is false ($\alpha{:}A$). Rule ($\bot$-intro), on the other hand, is the canonical way of constructing contradictions: from a proof of A and an assumption that A is false ($\beta{:}A$).

We can also relate λ_μ to classical sequent calculus as follows. We have $x_1{:}A_1,\ldots,x_n{:}A_n;\alpha_1{:}B_1,\ldots\alpha_m{:}B_m \vdash t : A$ for some t iff the sequent $A_1,\ldots,A_n \vdash B_1,\ldots,B_m,A$ is derivable in Gentzen's sequent calculus LK. Under this analogy the two rules ($\bot$-intro) and ($\bot$-elim) correspond to addition and removal of $\bot$ on the right hand side of the "turnstile".

The rules of λ_μ are such that logical rules always have the last conclusion as main formula. Permuting a conclusion into this "active" position is recorded by an instance of ($\bot$-elim) (preceded by an instance of ($\bot$-intro) unless the formula in the active position is $\bot$). Thus, the purpose of the $\bot$-rules is to display (the otherwise implicit) switching of focus, where "being in focus" means to be the main formula of the next logical rule.

We differ from Ong's presentation in that we allow for side context in rule (Axiom), i.e., the fact that variables other than x may be declared in Γ, and that α may appear free in t in rule ($\bot$-intro). Our (stronger) rules are derivable in Ong's system using the structural rules (weakening and contraction) which in turn are admissible in our system. Therefore, the same sequents are derivable in either system.

Definition 2.1 A λ_μ-theory (over a signature $(\mathcal{B},\mathcal{C})$) is a set $\mathcal{E}$ of typed equations of the form $\Gamma;\Delta \vdash s = t : A$ where $\Gamma;\Delta \vdash s : A$ and $\Gamma;\Delta \vdash t : A$ such that

- $\mathcal{E}$ is a congruence stable under weakening,
- $\mathcal{E}$ contains all well-typed instances of the following basic equality laws.

(β) $\quad (\lambda x{:}A.t)s = t[s/x]$.

(η) $\quad \lambda x{:}A.tx = t$,
when x is not free in t.

(μ-β) $\quad [\alpha]\mu\gamma{:}A.t = t[\alpha/\gamma]$.

(μ-η) $\quad \mu\alpha{:}A.[\alpha]t = t$,
α not free in t.

(μ-ζ) $\quad \mu\alpha{:}A \Rightarrow B.t = \lambda x{:}A.\mu\beta{:}B.t[x{::}\beta\,/\,\alpha]$,
when $B \not\equiv \bot$.
$\mu\alpha{:}A \Rightarrow \bot.t = \lambda x{:}A.t[x{::}\star\,/\,\alpha]$.

$\Diamond$

The notation used in the equations deserves some explanation. The term $t[s/x]$ in rule (β) denotes the capture-free substitution of s for x in t and $t[\alpha/\gamma]$ denotes the capture-free substitution of continuation variable α for γ.

The term $t[s{::}\beta\,/\,\alpha]$ called *mixed substitution* of s and β for continuation variable $\alpha : A \Rightarrow B$ (where $B \not\equiv \bot$) is defined inductively by the clauses in Figure 1 $\alpha{:}A \Rightarrow B$, $B \not\equiv \bot$, $\beta{:}B$, $s{:}A$. Mixed substitution of continuation variables of type $A \Rightarrow \bot$ is defined analogously the key clause being

$$([\alpha]t)[s{::}\star\,/\,\alpha] = (t[s{::}\star\,/\,\alpha])s$$

Notice that $\star$ is part of the operation symbol.

Mixed substitution of continuation variables of type $A \Rightarrow \bot$ is defined analogously the key clause being

$$([\alpha]t)[s{::}\star\,/\,\alpha] = (t[s{::}\star\,/\,\alpha])s$$

$$
\begin{array}{rcll}
x[s{::}\beta \,/\, \alpha] & = & x & \\
(tt')[s{::}\beta \,/\, \alpha] & = & (t[s{::}\beta \,/\, \alpha])(t'[s{::}\beta \,/\, \alpha]) & \\
(\lambda y{:}\, C.t)[s{::}\beta \,/\, \alpha] & = & \lambda y{:}\, C.t[s{::}\beta \,/\, \alpha] & y \text{ not free in } s \\
(\mu\gamma{:}\, C.t)[s{::}\beta \,/\, \alpha] & = & \mu\gamma{:}\, C.(t[s{::}\beta \,/\, \alpha]) & \gamma \not\equiv \beta \text{ and not free in } s \\
([\gamma]t)[s{::}\beta \,/\, \alpha] & = & [\gamma](t[s{::}\beta \,/\, \alpha]) & \gamma \not\equiv \alpha \\
([\alpha]t)[s{::}\beta \,/\, \alpha] & = & [\beta](t[s{::}\beta \,/\, \alpha])s &
\end{array}
$$

Figure 1. Definition of mixed substitution

Notice that $\star$ is part of the operation symbol.

The idea behind this so-called *mixed substitution* is that a continuation for a function of type $A \Rightarrow B$ can be understood as an argument s and a continuation β for the ensuing result. The substitution operation $t[s{::}\beta \,/\, \alpha]$ allows one to substitute such an "intended" continuation for a continuation variable.

3 Continuation models of λ_μ

The λ_μ-calculus admits a simple and intuitive continuation semantics in an arbitrary category with enough products and exponentials, in particular in any cartesian closed category with a distinguished object R of responses.

Definition 3.1 [Category of continuations] A *category of continuations* is given by the following data:

1. A category $\mathbf{C}$ with a distinguished class $\mathbf{T}$ of objects of $\mathbf{C}$ called type objects.
2. A distinguished type object R of *responses*.
3. For every object Γ and type object A a chosen cartesian product $\Gamma \cdot A$.
4. A chosen terminal object $[]$ (for the empty context).
5. A chosen terminal object $1 \in \mathbf{T}$ (to interpret $\perp$).
6. For every type object A a chosen exponential $R^A \in \mathbf{T}$ of R by A.
7. For any two type objects A and B a chosen cartesian product $R^A \times B \in \mathbf{T}$ of R^A and B.

$\Diamond$

Clearly, $[] \cong 1$ and $R^A \times B \cong R^A \cdot B$. The presence of these isomorphic copies of terminal objects and cartesian products is not strictly necessary, for instance, we could postulate that a product $R^A \cdot B$ must be a type object if A and B are. However, they reflect syntactic distinctions and facilitate the formulation of term models.

Particular examples of continuation categories are the category of sets and various categories of domains where a natural choice for R is the set (domain) of screen outputs or alternatively booleans (for sets) and the Sierpinski space (two element poset) in the case of domains. A further important example is furnished by the term model of a simply-typed lambda calculus together with a distinguished base type R. This model is generic in the sense that if a certain equation holds in it then it must hold in any other continuation category.

Assume for the rest of this section a fixed category of continuations. Any assignment of type objects $[\![B]\!]$ to base types B extends to an assignment of type objects to all types by the following two clauses.

$$
\begin{array}{rcl}
[\![\perp]\!] & = & 1 \\
[\![A \Rightarrow B]\!] & = & R^{[\![A]\!]} \times [\![B]\!]
\end{array}
$$

The intention is that $[\![A]\!]$ is the space of abstract continuations of type A. Accordingly, we call $R^{[\![A]\!]}$ the (type) object of *denotations* of type A. This may explain the definition of $[\![A \Rightarrow B]\!]$: A continuation for a function is given by an argument (a denotation of type A) and a continuation for the result.

Let $\Gamma \equiv x_1{:}A_1, \ldots, x_n{:}A_n$ be an object context and $\Delta \equiv \alpha_1{:}\, B_1, \ldots, \alpha_m{:}\, B_m$ be a continuation context. We use the notation $R^{[\![\Gamma]\!]} \cdot [\![\Delta]\!]$ for the object $[] \cdot R^{[\![A_1]\!]} \cdot \ldots \cdot R^{[\![A_n]\!]} \cdot [\![B_1]\!] \cdot \ldots \cdot [\![B_m]\!]$ The subsequent interpretation of λ_μ is motivated by the following two natural isomorphisms familiar from the more special case of CCC's. The reader is invited to keep those in mind when going through the semantic clauses below.

Proposition 3.2 *Let A, B be type objects and X any object of a continuation category $\mathbf{C}$. We have the fol-*

lowing two isomorphisms natural in X.

$$\begin{array}{rcl}\mathbf{C}(X \cdot R^A, R^B) & \cong & \mathbf{C}(X, R^{R^A \times B}) \\ \mathbf{C}(X \cdot A, R^1) & \cong & \mathbf{C}(X, R^A)\end{array}$$

Assume an assignment of denotations to the constants, i.e. a morphism $[\![c]\!] : [] \to R^{[\![A]\!]}$ if $c: A$ is in $\mathcal{C}$. To each sequent $\Gamma; \Delta \vdash t : A$ we associate an arrow

$$[\![\Gamma; \Delta \vdash t : A]\!] : R^{[\![\Gamma]\!]} \cdot [\![\Delta]\!] \to R^{[\![A]\!]}$$

by the clauses in Figure 2 where the internal language of **C** is employed to simplify the notation. Notice that we use pattern matching for abstractions over product and unit types. For the case of untyped λ_μ such semantics has been defined in [9]. The crucial observation that $R^A \times B$ is an exponential of B by A in the category $\mathbf{C}_R$ which has the same objects as **C** and homsets given by $\mathbf{C}_R(X, Y) = \mathbf{C}(R^X, R^Y)$ has been made by several people around 1990, including [1] and [4].

Theorem 3.3 (Soundness) *The λ_μ-calculus is sound with respect to this interpretation in the sense that the set of equations $\Gamma; \Delta \vdash t_1 = t_2 : A$ where t_1, t_2 are appropriately typed terms and $[\![\Gamma; \Delta \vdash t_1 : A]\!] = [\![\Gamma; \Delta \vdash t_2 : A]\!]$ is a λ_μ-theory.*

Proof. Induction on derivations using appropriate substitution lemmas. □

4 Completeness of continuation semantics

Our aim in this section is to establish the following completeness result for the continuation semantics.

Theorem 4.1 *For every λ_μ-theory $\mathcal{E}$ over a signature $(\mathcal{B}, \mathcal{C})$ there exists a continuation category $(\mathbf{C}, R)$ with the following two properties.*

1. *$\mathcal{E}$ is the theory induced by this continuation model $(\mathbf{C}, R)$.*
2. *Let Γ be an object context, Δ be a continuation context and A be a type. Every **C**-morphism $f : R^{[\![\Gamma]\!]} \cdot [\![\Delta]\!] \to R^{[\![A]\!]}$ arises as the interpretation of some λ_μ-term $\Gamma; \Delta \vdash t : A$.*

For the rest of this section assume a fixed signature $(\mathcal{B}, \mathcal{C})$ and a theory $\mathcal{E}$. When we refer to continuation contexts or terms it will always be relative to this signature. We will use the notation from Def. 2.1 to refer to the continuation category which we are going to construct from these data.

Before embarking on the actual construction let us first provide some intuition. Assume for the moment that $\mathcal{E}$ happens to be induced by a hypothetical continuation category **C**. It suggests itself to recover **C** from the λ_μ-theory by using the continuation contexts as objects. Unfortunately, morphisms from $[\![\Delta]\!]$ to $[\![A]\!]$ do not arise as meanings of terms so that there is no straightforward way to recover the **C**-morphisms from terms. However, the meaning of a term $; \Delta \vdash t : A \Rightarrow \bot$ is a **C**-morphism from $[\![\Delta]\!]$ to $R^{R^{[\![A]\!] \cdot 1}} \cong R^{R^{[\![A]\!]}}$. Now every morphism $f : [\![\Delta]\!] \to [\![A]\!]$ induces a morphism from $[\![\Delta]\!]$ to $R^{R^{[\![A]\!]}}$ by composition with the curried evaluation map $\eta_{[\![A]\!]}(a{:}\, [\![A]\!]) = \underline{\lambda} p{:}\, R^{[\![A]\!]}.pa$. Alas, in general, it seems to be impossible to tell whether a given morphism $h : [\![\Delta]\!] \to R^{R^{[\![A]\!]}}$ factors through η. However, those h which do, satisfy a certain equation. Namely[2], let $\eta_{R^2([\![A]\!])}$ and $R^2(\eta_{[\![A]\!]})$ be the two canonical maps from $R^2([\![A]\!]) \to R^4([\![A]\!])$, i.e., (in λ-calculus notation)

$$\eta_{R^2([\![A]\!])}(\varphi) = \lambda\Phi{:}\, R^3([\![A]\!]).\Phi(\varphi)$$

$$R^2(\eta_{[\![A]\!]})(\varphi) = \lambda\Phi{:}\, R^3([\![A]\!]).\varphi(\lambda x{:}\, [\![A]\!].\Phi(\lambda k{:}\, R([\![A]\!]).kx))$$

Then, since $\eta_{R^2([\![A]\!])} \circ \eta_{[\![A]\!]} = R^2(\eta_{[\![A]\!]}) \circ \eta_{[\![A]\!]}$ by λ-calculus, we have $\eta_{R^2([\![A]\!])} \circ h = R^2(\eta_{[\![A]\!]}) \circ h$, whenever h factors through $\eta_{[\![A]\!]}$, i.e., can be written in the form $\eta_{[\![A]\!]} \circ h'$. We can also argue element-wise and conclude that if $F : R^2([\![A]\!])$ is of the form $\eta_{[\![A]\!]}(a)$ for some $a : A$, i.e., $F = \lambda k.ka$, then

$$(\dagger) \qquad \Phi(F) = F(\lambda x{:}\, [\![A]\!].\Phi(\eta_{[\![A]\!]}(x))) \quad \text{for all } \Phi{:}\, R^3([\![A]\!])$$

So we are led to decree that a morphism from Δ to A is a λ_μ-term $; \Delta \vdash t : A \Rightarrow \bot$ satisfying this property (which, of course, has to be translated into λ_μ-equations). The morphisms into contexts rather than merely types are then constructed as tuples of these.

Unfortunately, since $\eta_{[\![A]\!]}$ is in general not the equaliser of $\eta_{R^2([\![A]\!])}, R^2(\eta_{[\![A]\!]})$ this equational condition is necessary, but not sufficient for factoring through $\eta_{[\![A]\!]}$. Therefore, the category thus obtained need not be equivalent to a possibly already existing (or hypothetically assumed) category **C**, but fortunately does the job, nevertheless.

In the sequel we will carry out the construction of this syntactic category explicitly and demonstrate that

[2] We follow Paul Taylor in using the notation $R^n(X)$ for an R-tower of height n, so $R^2(X)$ stands for R^{R^X}.

$$
\begin{array}{lcl}
[\![\Gamma;\Delta\vdash x_i : A]\!](\vec{x};\vec{\alpha}) & = & x_i \\
[\![\Gamma;\Delta\vdash \lambda x{:}A.t : A\Rightarrow B]\!](\vec{x};\vec{\alpha}) & = & \underline{\lambda}\langle x,\beta\rangle{:}\, R^{[\![A]\!]}\times[\![B]\!].[\![\Gamma,x{:}A;\Delta\vdash t : B]\!](\vec{x},x;\vec{\alpha})\beta \\
[\![\Gamma;\Delta\vdash ts : B]\!](\vec{x};\vec{\alpha}) & = & \underline{\lambda}\beta{:}\,[\![B]\!].[\![\Gamma;\Delta\vdash t : A\Rightarrow B]\!](\vec{x};\vec{\alpha})\langle[\![\Gamma;\Delta\vdash s : A]\!](\vec{x};\vec{\alpha}),\beta\rangle \\
[\![\Gamma;\Delta\vdash \mu\alpha{:}A.t : A]\!](\vec{x};\vec{\alpha}) & = & \underline{\lambda}\alpha{:}\,[\![A]\!].[\![\Gamma;\Delta,\alpha{:}A\vdash t : \bot]\!](\vec{x};\vec{\alpha},\alpha)\star \\
[\![\Gamma;\Delta\vdash [\alpha_i]t : \bot]\!](\vec{x};\vec{\alpha}) & = & \underline{\lambda}\star{:}\,1.[\![\Gamma;\Delta\vdash t : A]\!](\vec{x};\vec{\alpha})\alpha_i
\end{array}
$$

Figure 2. Interpretation of λ_μ in a λ_μ-category

it meets the requirement of the main theorem. Surprisingly, the main effort consists of showing that the C-maps (as we will call them) compose and thus form a category at all. The reason is that composition cannot be defined as syntactic substitution.

Convention 4.2 *In order to avoid messy case analyses we shall henceforth adopt the convention that lower case Greek letters range over continuation variables as well as $\bot$. We extend μ-abstraction and μ-application by the settings $\mu\bot.M := M$ and $[\bot]M = M$. It is clear that this preserves typing and that equations involving μ and $[\bot]\bot$ generalise accordingly.*

With this convention the two parts of rule (μ-ζ) can be subsumed under the first one.

4.1 The generic continuation category

Definition 4.3 Let Δ be a continuation context and A be a type. A *continuation term* (*C-term* for short) in context Δ and of type A is a λ_μ-term $;\Delta\vdash t : A\Rightarrow\bot$ such that for every λ_μ-term $\Gamma;\Delta,\Delta'\vdash o : (A\Rightarrow\bot)\Rightarrow\bot$ (subsequently called an *observer*) we have

$$\Gamma;\Delta,\Delta'\vdash ot = t(\mu\alpha{:}A.o(\lambda x{:}A.[\alpha]x)) : \bot$$

More generally, a *continuation map* (*C-map*) from Δ to $\Theta\equiv\alpha_1{:}A_1,\dots,\alpha_n{:}A_n$ is an n-tuple $(t_1,\dots,t_n)$ such that t_i is a C-term of type A_i in context Δ. ◊

We remark that t is a C-term of type A iff for each $\vec{\alpha}{:}\,[\![\Delta]\!]$ the element $\lambda d{:}\,R^{[\![A]\!]}.[\![t]\!](;\vec{\alpha})\langle d,\star\rangle$ (that is $[\![t]\!](;\vec{\alpha})$ transported along the isomorphism $R^{R^{[\![A]\!]}\times 1}\cong R^2([\![A]\!])$) satisfies condition (†) above.

Notice that by Convention 4.2 a term $t : \bot\Rightarrow\bot$ is a C-term of type $\bot$ if for

$$\Gamma;\Delta,\Delta'\vdash ot = t(o(\lambda x{:}\bot.x)) : \bot$$

for every observer $\Gamma;\Delta,\Delta'\vdash o : (\bot\Rightarrow\bot)\Rightarrow\bot$.

The desired continuation category **C** will have the continuation contexts as objects and the C-maps as morphisms. In order to define composition of C-maps we need to develop some machinery first.

Proposition 4.4 *If $\alpha{:}A$ is a continuation variable (declared in Δ) then the term $\mathsf{v}_\alpha := \lambda x{:}A.[\alpha]x$ is a C-term of type A.*

Proof. Assume $\Gamma;\Delta,\Delta'\vdash o : (A\Rightarrow\bot)\Rightarrow\bot$. We must show that (in context $\Gamma;\Delta,\Delta'$) we have $o\mathsf{v}_\alpha = \mathsf{v}_\alpha(\mu\alpha{:}A.o(\mathsf{v}_\alpha))$. By ($\beta$) the right hand side equals $[\alpha](\mu\alpha{:}A.o(\mathsf{v}_\alpha))$ which equals the left hand side by (μ-β). □

Henceforth, we will use capital letters M,N,U,V to range over C-terms. The following characterises C-terms of negated type.

Proposition 4.5 *The C-terms of type $(A\Rightarrow\bot)$ are precisely those of the form $\lambda f{:}A\Rightarrow\bot.ft$ for arbitrary term $t : A$. In particular, if M is a C-term of type $A\Rightarrow\bot$ then*

$$M = \lambda f{:}A\Rightarrow\bot.f(\mu\alpha{:}A.M(\mathsf{v}_\alpha))$$

Proof. To prove the equation let $f{:}A\Rightarrow\bot$ be a fresh variable. We calculate as follows. Employing the observer $o = \lambda m.f(\mu\alpha{:}A.m(\mathsf{v}_\alpha))$ we can rewrite the body of the left-hand side as $M(\mu\varphi{:}A\Rightarrow\bot.o(\mathsf{v}_\varphi))$ which equals $M(\lambda x{:}A.f(\mu\alpha.[\alpha]x))$ by (μ-ζ) and finally Mf by (μ-η) and (η). We conclude by abstracting from f.

The proof that $\lambda f.ft$ is a C-map is similar to the proof of Prop. 4.4 □

Definition 4.6 Let t be a term of type $B\Rightarrow\bot$ possibly containing a free continuation variable $\alpha : A$ and s be a term of type $A\Rightarrow\bot$. The term $t[\alpha{:=}s]$ of type $B\Rightarrow\bot$ is defined as $\lambda x{:}B.s(\mu\alpha{:}A.tx)$. ◊

Proposition 4.7 *If M is a C-term of type B (possibly containing $\alpha{:}A$) and N is a C-term of type A then $M[\alpha{:=}N]$ is a C-term of type B.*

Proof. Let $o : (B \Rightarrow \bot) \Rightarrow \bot$ be an observer. We demonstrate the required identity $o(M[\alpha{:=}N]) = M[\alpha{:=}N](\mu\beta{:}\,B.o(\mathsf{v}_\beta))$ by showing that both sides equal $N(\mu\alpha.o(M))$.

Using the observer

$$o' = \lambda n{:}\,A \Rightarrow \bot.o(\lambda x{:}\,B.n(\mu\alpha{:}\,A.Mx))$$

we obtain

$$o(M[\alpha{:=}N]) = N(\mu\alpha{:}\,A.o'(\mathsf{v}_\alpha))$$

Notice that the left-hand side is $o'(N)$. By (β),$(\mu\text{-}\beta)$, and (η) this equals $N(\mu\alpha.o(M))$.

On the other hand, $M[\alpha{:=}N](\mu\beta{:}\,B.o(\mathsf{v}_\beta))$ equals $N(\mu\alpha.o(M))$ using the observer o applied to M. □

We will now show that the C-terms of the form v_- and $\bot[\bot{:=}\bot]$ behave like variables and substitution in ordinary λ-calculus.

Lemma 4.8 *The following equations hold whenever they are well-typed and M, N, U, V are C-terms as indicated.*

1. $\mathsf{v}_\alpha[\alpha{:=}M] = M$.
2. $\mathsf{v}_\beta[\alpha{:=}N] = \mathsf{v}_\beta$, *if* $\alpha \not\equiv \beta$.
3. $M[\alpha{:=}\mathsf{v}_\alpha] = M$.
4. $M[\alpha{:=}U] = M$, *if α not free in M.*
5. $M[\alpha{:=}U][\beta{:=}V] = M[\beta{:=}V][\alpha{:=}U[\beta{:=}V]]$, *if α is not free in V.*
6. $M[\alpha{:=}U][\beta{:=}V] = M[\beta{:=}V][\alpha{:=}U]$ *if α, β are not free in U, V.*
7. $M[\alpha{:=}U][\beta{:=}V] = M[\alpha{:=}U[\beta{:=}V]]$, *if β not free in M.*

Proof. First, observe that if $\alpha{:}\,A$ is a continuation variable and M is any term of type $\bot$ then $\mathsf{v}_\alpha(\mu\alpha{:}\,A.M) = M$ by (β) and $(\mu\text{-}\beta)$.

We will omit type annotations as they depend upon the unspecified typing of the equations.

Ad 1. $\mathsf{v}_\alpha[\alpha{:=}M] = \lambda x.M(\mu\alpha.\mathsf{v}_\alpha x) = \lambda x.Mx = M$ by $(\mu\text{-}\eta)$ and (η).

Ad 3. $M[\alpha{:=}\mathsf{v}_\alpha] = \lambda x.\mathsf{v}_\alpha(\mu\alpha.Mx) = M$ by the above observation followed by $(\mu\text{-}\beta)$ and (η).

Ad 5. The right hand side expands to $\lambda x.o(V)$ where

$$o = \lambda v.v(\mu\beta.U(\mu\alpha.o(\mu\beta.Mx)))$$

Since V is a C-term this equals

$$\lambda x.V(\mu\beta.\mathsf{v}_\beta(\mu\beta.U(\mu\alpha.\mathsf{v}_\beta(\mu\beta.Mx))))$$

By the above observation this equals $\lambda x.V(\mu\beta.U(\mu\alpha.Mx))$ which is a (β)-contraction of the left hand side.

Ad 4. By definition $M[\alpha{:=}U] = \lambda x.U(\mu\alpha.Mx)$. By Lemma 4.5 applied to U and $t := Mx$ this equals $\lambda x.Mx$. The conclusion follows by (η).

The remaining parts 2, 6, 7 are immediate consequences of 5 and 4. □

This allows us to carry out the usual construction [8] of a category with finite products from a substitution calculus. Let M be a C-term of type A in context $\Delta \equiv \alpha_1{:}\,A_1, \dots, \alpha_n{:}\,A_n$. Furthermore, let $f \equiv (N_1, \dots, N_n)$ be a C-map from Θ to Δ. We define the C-term $M \circ f$ of type A in context Θ as $M[\alpha_1{:=}N_1]\dots[\alpha_n{:=}N_n]$.

More generally, if $g = (M_1, \dots, M_m)$ is a C-map from Δ to Ψ, we define the composition $g \circ f$ as $(M_1 \circ f, \dots, M_m \circ f)$. Finally, the identity C-map $\mathrm{id}_\Delta : \Delta \to \Delta$ is defined as $(\mathsf{v}_{\alpha_1}, \dots, \mathsf{v}_{\alpha_n})$.

The proof [8] that the term model of an equational theory forms a category can now be copied word for word so as to demonstrate that the continuation contexts and C-maps form a category **C**. The continuation contexts of length one which we will henceforth identify with types form the subset **T** of **C**. It also follows from this proof that the empty context [] forms a terminal object and that the extended context $\Delta, \alpha{:}\,A$ forms a cartesian product of Δ and type A if $A \neq \bot$. We define the cartesian product $\Delta \cdot \bot$ as Δ. The projection on $\bot$ is given by v_-. By abuse of notation we decree that $\Delta, \alpha{:}\,\bot$ means Δ.

The type object of responses R is defined as the type $\bot \Rightarrow \bot$. This setting is motivated by the observation that the meaning of type $\bot \Rightarrow \bot$ in an arbitrary continuation category equals $R^1 \times 1$ which is isomorphic to R.

The following proposition exhibits the remaining structure required to deomnstrate that **C** is a category of continuations. We omit its proof by straightforward equality reasoning.

Proposition 4.9 *Let A, B be types,* i.e. *type objects of* **C**. *The exponential R^A in* **C** *can be given by $A \Rightarrow \bot$; the product $R^A \times B$ in* **C** *can be given by $A \Rightarrow B$. The type $\bot$ is a terminal object. The associated morphisms* *and operators are given as follows.*

$$app(M,N) = \lambda f.N(\mu\alpha.M(\lambda x.f([\alpha]x)))$$

$$cur_{\alpha:A}(M) = \lambda f\colon A \Rightarrow \bot.f(\mu\alpha\colon A.M(\mathsf{v}_{-}))$$

$$\langle \lambda f\colon R^A.ft, N\rangle = \lambda f\colon A \Rightarrow B.N(ft)$$

$$M.1 = \pi[\varphi{:=}M] = \lambda f\colon R^A.M(\lambda x\colon A.\mu\beta\colon B.fx)$$

$$M.2 = \lambda b\colon B.M(\lambda x\colon A.b)$$

$$\star = \mathsf{v}_{-}$$

Notice that the definition of pairing is exhaustive in view of Prop. 4.5.

4.2 Interpretation of λ_μ

We have seen that **C** with the described settings furnishes a continuation category. Interpreting base types by themselves we obtain immediately

Proposition 4.10 *For any λ_μ-type A the interpretation $[\![A]\!]$ in* **C** *equals A.*

Accordingly, the semantics of a continuation context Δ can be chosen as Δ itself (the choice only affects the names of continuation variables). Next, we examine the interpretation of combined contexts $\Gamma;\Delta$ where $\Gamma \equiv x_1\colon A_1,\ldots,x_n\colon A_n$ is an object context. Its interpretation $R^{[\![\Gamma]\!]}\cdot[\![\Delta]\!]$ is the continuation context $\varphi_1\colon A_1 \Rightarrow \bot,\ldots,\varphi_n\colon A_n \Rightarrow \bot,\Delta$ where the φ_i are freshly chosen continuation variables. We introduce the notation φ_x for the continuation variable in $R^{[\![\Gamma]\!]}\cdot[\![\Delta]\!]$ corresponding to object variable x in Γ, in other words we have $\varphi_{x_i} = \varphi_i$. If $\vec{x}$ is a sequence of object variables we write $\varphi_{\vec{x}}$ for the corresponding sequence of continuation variables.

If $c\colon A$ is a constant then $[\![c]\!] := \lambda f\colon A \Rightarrow \bot.fc$ is a C-term of type $R^{[\![A]\!]}$ yielding an interpretation for the constants. We thus obtain an interpretation of our λ_μ-calculus in the continuation category **C** which associates with every λ_μ-term $\Gamma;\Delta \vdash t : A$ a C-term $[\![t]\!]$ of type $R^{[\![A]\!]} = A \Rightarrow \bot$ in context $R^{[\![\Gamma]\!]}\cdot[\![\Delta]\!] = \varphi_1\colon A_1 \Rightarrow \bot,\ldots,\varphi_n\colon A_n \Rightarrow \bot,\Delta$. The key to the completeness result is the following direct characterisation of this interpretation.

Theorem 4.11 *Whenever $\Gamma;\Delta \vdash t : U$ then $[\![\Gamma;\Delta \vdash t : U]\!]$ is equal, w.r.t. $\mathcal{E}$, to the following term*

$$R^{[\![\Gamma]\!]}\cdot[\![\Delta]\!] \vdash \lambda f\colon U \Rightarrow \bot.f(t^\star) : R^U \Rightarrow \bot$$

where $t^\star$ is obtained from t by replacing each object variable $x_i\colon A_i$ by the term $\overline{\varphi_x} := \mu\alpha\colon A_i.[\varphi_x]\mathsf{v}_\alpha$.

Proof. By induction on derivations. As a representative example we treat the case of rule ($\Rightarrow$-intro), i.e., where $U = A \Rightarrow B$ and $t = \lambda x\colon A.s$. Let us write $\tilde{s}$ for s with every free object variables y of s except x replaced by $\overline{\varphi_y}$. That is to say, we have $s^\star = \tilde{s}[\overline{\varphi_x}/x]$. The equational calculation displayed in Figure 3 where we have used the following sublemma.

Sublemma *Let t be a λ_μ-term of type $\bot$ containing object variable $x\colon A$ and let $\varphi\colon A \Rightarrow \bot$ be a fresh continuation variable. Then $\mu\varphi\colon A \Rightarrow \bot.t[\overline{\varphi}/x] = \lambda x\colon A.t$.*

To see this, we use (μ-ζ) to expand $\mu\varphi.t[\overline{\varphi}/x]$ as $\lambda x\colon A.t[\overline{\varphi}/x][x{::}\star/\varphi]$ which is identical to $\lambda x.t[(\overline{\varphi}[x{::}\star/\varphi])/x]$ by definition of substitution. Now $\overline{\varphi}[x{::}\star/\varphi] = (\mu\alpha.[\varphi]\mathsf{v}_\alpha)[x{::}\star/\varphi] = \mu\alpha.[\alpha]x = x$ by (μ-η). Hence the result. □

Proof of Theorem 4.1. For faithfulness it is sufficient to consider the special case where Γ is empty. So let that $;\Delta \vdash t_i : A$ for $i = 1,2$ be two terms of λ_μ with equal semantics in the generic continuation category **C** constructed from $\mathcal{E}$. By Theorem 4.11 this implies that we have $\lambda f\colon A \Rightarrow \bot.ft_1 = \lambda f\colon A \Rightarrow \bot.ft_2$. Notice that $t_i^\star = t_i$. Application to v_α where $\alpha\colon A$ is a fresh continuation variable gives $[\alpha]t_1 = [\alpha]t_2$ from which we get $t_1 = t_2$ by μ-abstraction of α followed by (μ-η).

For the second part we may also assume that Γ is empty because of the isomorphism from Lemma 3.2. So let F be a **C**-morphism from $[\![\Delta]\!]$ to $R^{[\![A]\!]}$. By definition of **C** this means that F is a C-term of type $A \Rightarrow \bot$ in context Δ. In view of Prop. 4.5 this means that F takes the form

$$\Delta \vdash \lambda f\colon A \Rightarrow \bot.ft : (A \Rightarrow \bot) \Rightarrow \bot$$

for some λ_μ-term t of type A. But by Theorem 4.11 this implies that F is the interpretation of t. □

4.3 Universality of continuation models

In [6] is defined a categorical semantics of λ_μ called λ_μ-categories. By the categorical completeness result of loc. cit. these models are in 1-1-correspondence with λ_μ-theories. Thus, our results can be stated in category-theoretic terms as follows. The soundness theorem (Thm. 3.3) says that every category of continuations can be organised into a λ_μ-category with the same objects and whose homsets $\mathbb{E}_\Delta((A_1,\ldots,A_n), B)$ are given by $\mathbf{C}(\Delta\cdot R^{A_1}\cdot\ldots\cdot R^{A_n}, R^B)$. The completeness theorem (Thm. 4.1), on the other hand, expresses that every λ_μ-category is (up to isomorphism) of this

$$
\begin{array}{lll}
& [\![\lambda x\colon A.s]\!] & \\
= & cur_{\psi:R^A\times B}(app([\![s]\!][\varphi_x := \mathsf{v}_\psi.1]\,,\,\mathsf{v}_\psi.2)) & \text{by definition of } [\![\lambda x.s]\!] \\
= & \lambda f\colon (A \Rightarrow B) \Rightarrow \bot.f(\mu\psi\colon A \Rightarrow B.app([\![s]\!][\varphi_x := \mathsf{v}_\psi.1]\,,\,\mathsf{v}_\psi.2)\mathsf{v}_-) & \text{expansion of } cur \\
= & \lambda f.f(\mu\psi.\mathsf{v}_\psi.2(\mu\beta\colon B.[\![s]\!][\varphi_x := \mathsf{v}_\psi.1]\mathsf{v}_\beta)) & \text{expansion of } app \\
= & \lambda f.f(\mu\psi.\mathsf{v}_\psi.2(\mu\beta.\mathsf{v}_\psi.1(\mu\varphi_x.[\![s]\!]\mathsf{v}_\beta))) & \text{expansion of } [\![s]\!][\varphi_x := \mathsf{v}_\psi.1] \\
= & \lambda f.f(\mu\psi.\mathsf{v}_\psi.2(\mu\beta.\mathsf{v}_\psi.1(\mu\varphi_x.[\beta]\tilde{s}[\overline{\varphi_x}/x]))) & \text{induction hypothesis} \\
= & \lambda f.f(\mu\psi.\mathsf{v}_\psi.2(\mu\beta.\mathsf{v}_\psi.1(\lambda x.[\beta]\tilde{s}))) & \text{by sublemma} \\
= & \lambda f.f(\mu\psi.\mathsf{v}_\psi.2(\mu\beta.[\psi]\lambda x.\mu\beta'.[\beta]\tilde{s})) & \text{expansion of } \mathsf{v}_\psi.1 \\
= & \lambda f.f(\mu\psi.[\psi](\lambda x\colon A.\mu\beta.[\psi]\lambda x.\mu\beta'.[\beta]\tilde{s})) & \text{expansion of } \mathsf{v}_\psi.2 \\
= & \lambda f.f(\lambda x\colon A.\mu\beta\colon B.[\beta]\tilde{s}) & (\mu\text{-}\zeta) \text{ on } \psi,\ (\mu\text{-}\beta) \\
= & \lambda f.f(\lambda x.\tilde{s}) & (\mu\text{-}\eta) \\
= & \lambda f.f(\lambda x.s[\overline{\varphi_{\vec{x}}}/\vec{x}]) & \\
= & \lambda f.f(t[\overline{\varphi_{\vec{x}}}/\vec{x}]) &
\end{array}
$$

Figure 3. Calculation in the proof of Theorem 4.11

form. Namely, by Ong's categorical completeness result every λ_μ-category is isomorphic to the term model of its theory which in turn (by our main result) is isomorphic to the λ_μ- category induced by the associated syntactic category of continuations.

Thus, continuation models are universal in the sense that every model of λ_μ, i.e. any λ_μ-category is isomorphic to a continuation model. We have preferred to stick to the more traditional syntactic presentation of our results as this simplifies the calculations. It might, however, be instructive to explicitly compute the continuation model which induces the game-theoretic λ_μ-category described in loc. cit.

4.4 Completeness of λ_μ for CPS-translation

Of particular interest is the "free" λ_μ-calculus over some signature $(\mathcal{B}, \mathcal{C})$ without non-logical axioms. It can be interpreted in the (cartesian closed category associated with) simply-typed lambda calculus with products and terminal object over base types $\mathcal{B} \cup \{R\}$ and constants of appropriate type. In this particular case the interpretation gives rise to the following CPS translation of λ_μ.

$$
\begin{array}{l}
[\![x]\!] = x \\
[\![c]\!] = c \\
[\![\lambda x.t]\!] = \lambda p.(\lambda x.[\![t]\!]\,p.1)p.2 \\
[\![ts]\!] = \lambda\beta.[\![t]\!]\langle[\![s]\!], \beta\rangle \\
[\![\mu\alpha.t]\!] = \lambda\alpha.[\![t]\!]\star \\
[\![[\alpha]t]\!] = \lambda x\colon 1.[\![t]\!]\alpha
\end{array}
$$

These clauses are derived by instantiating the defining clauses in Section 3 with the term model of simply typed lambda calculus where we keep using the same name for a syntactic and semantic variables. This particular semantics is already complete for free λ_μ.

Proposition 4.12 *We have* $\Gamma; \Delta \vdash t_1 = t_2 : A$ *if and only if* $[\![t_1]\!] = [\![t_2]\!]$.

Proof. If $[\![t_1]\!] = [\![t_2]\!]$ then t_1 and t_2 have equal interpretation in any continuation category because simply-typed lambda calculus is initial for those. Thus $\Gamma; \Delta \vdash t_1 = t_2 : A$ by Thm. 4.1. The other direction is immediate from soundness (Thm. 3.3). □

This result means that the free λ_μ-theory can be decided by way of the above CPS translation.

5 Conclusions and further work

We have presented a natural continuation style interpretation of call-by-name λ_μ-calculus and demonstrated that the equational theory of λ_μ is complete with respect to this interpretation. Unlike in [10] or [2] the equational theory predated the continuation style interpretation which provides evidence for its canonicity. Indeed, although known for a while, the present CPS translation using pairs does not seem to have received the attention it deserves. We hope that this paper will help to popularize it.

It should be stressed that the precise formulation of our notion of model, e.g. the restriction of exponentiation to certain rather peculiar objects and the duplication of products and terminal object, is not essential for completeness and universality.

As mentioned above, every cartesian-closed category together with a distinguished object R defines a category of continuations. Conversely, given an arbitrary category of continuations $(\mathbf{C}, R)$ the category of presheaves $\hat{\mathbf{C}}$ is cartesian closed and the full and faithful Yoneda embedding $\mathcal{Y} : \mathbf{C} \to \hat{\mathbf{C}}$ preserves existing products and exponentials. Therefore, the continuation semantics with respect to $\hat{\mathbf{C}}$ and $\mathcal{Y}(R)$ agrees up to isomorphism with the interpretation in $(\mathbf{C}, R)$ composed with $\mathcal{Y}$. Thus, the completeness result Thm. 4.1 continues to hold for the restricted class of cartesian closed λ_μ-categories.

Some directions for further work suggest themselves. Although our formulation of λ_μ-syntax is very general and emcompasses, e.g., fixpoint operators and recursive types (via fold/unfold constants), it does not immediately extend to an untyped formulation of λ_μ. The reason is the special role of the type $\perp$. However, it seems that our constructions can also be applied to an appopriate untyped version in which we would have two judgements $\Gamma; \Delta \vdash t \bullet$ (to mean that t is a "computation") and $\Gamma; \Delta \vdash t$ (to mean that t is an "observation"). The details remain to be worked out.

We want to stress that although our completeness result encompasses arbitrary λ_μ-theories it does not say anything about their consistency. It is easy to see that the abovementioned extension by general recursion and recursive types is indeed consistent, e.g., using the cartesian closed category of dcpos with the Sierpinski space as object of responses. Moreover, somewhat more interestingly, the category of continuations $\mathbf{C} = \mathbf{Set}$ with $R = \mathbb{N}$ gives rise to a model of λ_μ with a base type of natural numbers N where $[\![N]\!] = [\![\perp]\!] = 1$. This model supports system T style iterators

$$\mathcal{R}_A : A \Rightarrow (N \Rightarrow A \Rightarrow A) \Rightarrow N \Rightarrow A$$

together with the equations

$(\mathcal{R}\text{-}0)$ $\quad \mathcal{R}_A z f 0 = z$

$(\mathcal{R}\text{-}succ)$ $\quad \mathcal{R}_A z f(succ(x)) = f x(\mathcal{R}_A z f x)$

as well as all instances of the induction principle for equations.

However, under the more canonical interpretation $[\![N]\!] = R^{\mathbb{N}}$ not even the computation rules for the iterator are valid in general. That is, rule $\mathcal{R} \perp succ$ must be confined to the case where x is a numeral, i.e., of the form $succ^n(0)$. Given this difficulty with the interpretation of inductive data types a call-by-value formulation of λ_μ is perhaps more appealing from a practical point of view.

Indeed, very recently, Ong and Stewart [5] have come up with a call-by-value variant of the equational theory of λ_μ. Preliminary investigations have shown that this theory is also complete with respect to a more traditional continuation semantics encompassing values. To achieve this result we have to combine the present approach with Hofmann's earlier results in [2]. A joint paper with Ong on this subject is in preparation.

References

[1] B. Agapiev and E. Moggi. Declarative Continuations and Monads. unpublished draft, July 1991.

[2] M. Hofmann. Sound and complete axiomatisations of call-by-value control operators. *Math. Struct. Comp. Sci*, 5:461–482, 1995.

[3] J.-L. Krivine. Classical logic, storage operators, and second-order lambda calculus. *Annals of Pure and Applied Logic*, 68:53–78, 1994.

[4] Y. Lafont, B. Reus, and T. Streicher. Continuation semantics or expressing implication by negation. Technical Report 93-21, University of Munich, 1993.

[5] C.-H. L. Ong and C. A. Stewart. A curry-howard foundation for functional computation with control. In *Proceedings of ACM SIGPLAN-SIGACT Symposium on Principle of Programming Languages, Paris, January 1997*. ACM Press, 1997.

[6] L. Ong. A semantic view of classical proofs. In *Proceedings of LICS '96*. IEEE Press, 1996.

[7] M. Parigot. $\lambda\mu$-calculus: an algorithmic interpretation of classical matural deduction. In *Proc. LPAR '92, vol. 624 of LNCS*. Springer, 1992.

[8] A. Pitts. Categorical logic. In *Handbook of Logic in Computer Science (Vol. VI)*. Oxford University Press, 1997. To appear.

[9] B. Reus and T. Streicher. Continuation Semantics : Abstract Machines and Control Operators, 1996. submitted.

[10] A. Sabry and M. Felleisen. Reasoning about programs in continuation-passing style. *Lisp and Symbolic Computation*, 6(3/4):287–358, 1993. Preliminary version in Proceedings of the ACM Conference on Lisp and Functional Programming, 1992. Also Technical Report 92-180, Rice University.