

The $\lambda\mu$ -calculus: Function and Control

G.M. Bierman
Gonville and Caius College, Cambridge.

December 10, 1998

Abstract

Parigot's $\lambda\mu$ -calculus is an intriguing extension of the typed λ -calculus which corresponds via the Curry-Howard correspondence to classical logic. Following the seminal work of Griffin, it is known that certain control operators can be given types which, when viewed as formulae, are classical but not intuitionistic tautologies. Previous computational explanations of the $\lambda\mu$ -calculus have simply translated terms into an existing control calculus, or presented an operational semantics from which it is hard to determine what is going on. In particular the treatment of a call-by-value strategy has appeared problematic.

In this paper I give a very simple computational interpretation of the $\lambda\mu$ -calculus: it can be thought of as a typed λ -calculus with operators to save and restore the runtime environment. This interpretation can be given precisely using a single-step semantics. In particular both call-by-value and call-by-name strategies can be handled quite easily. The $\lambda\mu$ -calculus can be extended with natural numbers, a conditional, pairs and recursion to give a simple programming language which Ong and Stewart call μ PCF. A natural notion of program equivalence, contextual equivalence, is presented. A significantly simpler relation, based on transitions in an abstract machine, is given and proven to coincide with contextual equivalence. This provides a new method for reasoning about programs with control operators.

1 Introduction

It is well-known that the typed λ -calculus can be viewed as a term assignment for natural deduction proofs in intuitionistic logic (**IL**). Consequently the set of types of all closed λ -terms enumerates all intuitionistic tautologies. This is known as the Curry-Howard correspondence, or the formulae-as-types principle. Thus one can talk of a computational interpretation of **IL**. A natural question is whether there is such a computational interpretation of classical logic (**CL**). A first step is to devise a well behaved natural deduction formulation for **CL** and give a term assignment. A number of proposals have been made but recently Parigot (1992) introduced an extension of the typed λ -calculus, which he called the $\lambda\mu$ -calculus. The set of types of all closed $\lambda\mu$ -terms enumerates all classical tautologies and the calculus is amazingly well behaved, satisfying both strong normalisation and confluence.

However two questions remain. First, what does this extension to the λ -calculus mean computationally? Secondly, if the $\lambda\mu$ -calculus is extended in much the same way as the λ -calculus is extended to yield PCF, what is its operational theory? Of course the answer to the second question is heavily dependent upon the answer to the first. In this paper I suggest that the $\lambda\mu$ -calculus has a natural computational reading: it is a λ -calculus which has operators

to save and restore the runtime environment. This can easily be expressed using evaluation contexts which are common in work on control operators.

Morris-style contextual equivalence is commonly accepted as the natural notion of equivalence for functional languages. There has been significant effort in devising alternative characterisations of contextual equivalence which are more amenable for constructing proofs. For PCF the common solution is to use some form of (applicative) bisimilarity (Gordon 1995). However these techniques do not often extend to (call-by-value) languages with control. In §7 I give a simple notion of program equivalence, based on transitions in an abstract machine, which coincides with contextual equivalence.

2 Parigot's $\lambda\mu$ -calculus

In his seminal paper Parigot introduced an extension of the typed λ -calculus, which he called the $\lambda\mu$ -calculus. The extension is such that terms no longer have a single type but a *sequence* of types, one of which is designated to be the active type and the rest are said to be passive.

Types are given by the grammar

$$\phi ::= \perp \mid \phi \rightarrow \phi$$

and raw $\lambda\mu$ -terms are given by the grammar

$$\begin{array}{l} M ::= x \quad \text{Variable} \\ \quad | \lambda x: \phi. M \quad \text{Abstraction} \\ \quad | MM \quad \text{Application} \\ \quad | [a: \phi]M \quad \text{Passivate} \\ \quad | \mu a: \phi. M \quad \text{Activate;} \end{array}$$

where x is taken from a countable set of λ -variables, ϕ is a well-formed type (formula) and a is taken from another countable set of μ -variables. (As is usual, the types will sometimes be omitted from terms for brevity.)

Typing judgements are of the form, $\Gamma \triangleright M: \phi, \Sigma$, where Γ is a set of pairs of λ -variables and types written $x: \psi$, M is a term from the above grammar and Σ denotes a set of pairs of μ -variables and (passive) types written $a: \varphi$ (thus ϕ is the active type). The typing rules are as follows.

$$\begin{array}{c} \frac{}{\Gamma, x: \phi \triangleright x: \phi, \Sigma} \text{Identity} \\ \\ \frac{\Gamma, x: \phi \triangleright M: \psi, \Sigma}{\Gamma \triangleright \lambda x: \phi. M: \phi \rightarrow \psi, \Sigma} \rightarrow_{\mathcal{I}} \quad \frac{\Gamma \triangleright M: \phi \rightarrow \psi, \Sigma \quad \Gamma \triangleright N: \phi, \Sigma}{\Gamma \triangleright MN: \psi, \Sigma} \rightarrow_{\mathcal{E}} \\ \\ \frac{\Gamma \triangleright M: \phi, \Sigma}{\Gamma \triangleright [a: \phi]M: \perp, a: \phi, \Sigma} \text{Passivate} \quad \frac{\Gamma \triangleright M: \perp, a: \phi, \Sigma}{\Gamma \triangleright \mu a: \phi. M: \phi, \Sigma} \text{Activate} \end{array}$$

The new rules are called Passivate and Activate. The former takes a term whose active type is ϕ (where ϕ is not \perp) and passivates it, i.e. ϕ becomes a passive type (and is hence labelled with a). The resulting term has an active type of \perp .¹ The Activate rule works similarly but in the reverse direction.

¹This ensures that every term has an active type. It is possible to give a formulation where terms need not have an active type.

As an example the following derivation gives a (well-typed) term whose type is the *Peirce formula*.

$$\begin{array}{c}
\frac{}{x: \phi, y: (\phi \rightarrow \psi) \rightarrow \phi \triangleright x: \phi, a: \psi} \text{Identity} \\
\frac{}{x: \phi, y: (\phi \rightarrow \psi) \rightarrow \phi \triangleright [b: \phi]x: \perp, b: \phi, a: \psi} \text{Passivate} \\
\frac{}{x: \phi, y: (\phi \rightarrow \psi) \rightarrow \phi \triangleright \mu a: \psi.[b: \phi]x: \psi, b: \phi} \text{Activate} \\
\frac{}{y: (\phi \rightarrow \psi) \rightarrow \phi \triangleright y: (\phi \rightarrow \psi) \rightarrow \phi, b: \phi} \rightarrow_{\mathcal{I}} \\
\frac{}{y: (\phi \rightarrow \psi) \rightarrow \phi \triangleright \lambda x: \phi. \mu a: \psi.[b: \phi]x: \phi \rightarrow \psi, b: \phi} \rightarrow_{\mathcal{E}} \\
\frac{}{y: (\phi \rightarrow \psi) \rightarrow \phi \triangleright y(\lambda x: \phi. \mu a: \psi.[b: \phi]x): \phi, b: \phi} \text{Passivate} \\
\frac{}{y: (\phi \rightarrow \psi) \rightarrow \phi \triangleright [b: \phi]y(\lambda x: \phi. \mu a: \psi.[b: \phi]x): \perp, b: \phi} \text{Activate} \\
\frac{}{y: (\phi \rightarrow \psi) \rightarrow \phi \triangleright \mu b: \phi.[b: \phi]y(\lambda x: \phi. \mu a: \psi.[b: \phi]x): \phi} \rightarrow_{\mathcal{I}} \\
\frac{}{\triangleright \lambda y. \mu b: \phi.[b: \phi]y(\lambda x: \phi. \mu a: \psi.[b: \phi]x): ((\phi \rightarrow \psi) \rightarrow \phi) \rightarrow \phi} \rightarrow_{\mathcal{E}}
\end{array}$$

There are a number of reduction rules associated with the $\lambda\mu$ -calculus. In full they are as follows.

$$\begin{array}{l}
(\lambda x: \phi. M)N \rightsquigarrow_{\beta} M[x := N] \\
\mu a: \phi. [a: \phi]M \rightsquigarrow_s M \quad \text{where } a \notin \mu\text{FV}(M) \\
[a: \phi]\mu b: \phi. M \rightsquigarrow_s M[a/b] \\
(\mu a: \phi \rightarrow \psi. M)N \rightsquigarrow_c \mu b: \psi. M[a: \phi \rightarrow \psi \Leftarrow [b: \psi] \bullet N]
\end{array}$$

The β -rule is familiar from the λ -calculus. The $\lambda\mu$ -calculus introduces three new reduction rules. Two are known as *simplification rules* (Parigot 1997) and are written \rightsquigarrow_s . In the first simplification rule, $\mu\text{FV}(M)$ denotes the set of free μ -variables in M , which is defined as follows.

$$\begin{array}{l}
\mu\text{FV}(x) \stackrel{\text{def}}{=} \emptyset \\
\mu\text{FV}(\lambda x. M) \stackrel{\text{def}}{=} \mu\text{FV}(M) \\
\mu\text{FV}(MN) \stackrel{\text{def}}{=} \mu\text{FV}(M) \cup \mu\text{FV}(N) \\
\mu\text{FV}([a]M) \stackrel{\text{def}}{=} \mu\text{FV}(M) \cup \{a\} \\
\mu\text{FV}(\mu a. M) \stackrel{\text{def}}{=} \mu\text{FV}(M) - \{a\}
\end{array}$$

A term is said to be λ -closed if it has no free λ -variables; it is said to be μ -closed if it has no free μ -variables, and *closed* if it is both λ -closed and μ -closed. A λ -closed term will often be referred to as a program. In the second simplification rule, $M[a/b]$ denotes the term M where all free occurrences of the μ -variable b are replaced with a .

The third new reduction rule is essentially a commuting conversion, and is written \rightsquigarrow_c . I have used the notation $M[a \Leftarrow P[\bullet]]$ to denote the term M where *all* occurrences of the subterm $[a]N$ have been replaced by the term $P[N]$ (where $P[\bullet]$ is a term with a single hole in it, and $P[N]$ is the result of replacing the hole with N). This complicated commuting conversion is often glossed over in other papers. In truth there are actually two different rules depending on the type of the μ -variable a . In detail, the commuting conversion is defined as follows.

$$(\mu a: \phi \rightarrow \psi. M)N \rightsquigarrow_c \begin{cases} \mu b: \psi. M[a: \phi \rightarrow \psi \Leftarrow [b: \psi] \bullet N] & \text{where } \psi \neq \perp \\ M[a: \phi \rightarrow \psi \Leftarrow \bullet N] & \text{where } \psi = \perp \end{cases}$$

where

$$\begin{aligned}
x[a: \phi \Leftarrow P[\bullet]] &\stackrel{\text{def}}{=} x \\
(\lambda x.M)[a: \phi \Leftarrow P[\bullet]] &\stackrel{\text{def}}{=} \lambda x.(M[a: \phi \Leftarrow P[\bullet]]) \\
(MN)[a: \phi \Leftarrow P[\bullet]] &\stackrel{\text{def}}{=} (M[a: \phi \Leftarrow P[\bullet]])(N[a: \phi \Leftarrow P[\bullet]]) \\
(\mu b: \psi.M)[a: \phi \Leftarrow P[\bullet]] &\stackrel{\text{def}}{=} \mu b: \psi.(M[a: \phi \Leftarrow P[\bullet]]) \\
([b: \phi]M)[a: \phi \Leftarrow P[\bullet]] &\stackrel{\text{def}}{=} \begin{cases} P[M[a: \phi \Leftarrow P[\bullet]]] & \text{if } a = b \\ [b: \phi](M[a: \phi \Leftarrow P[\bullet]]) & \text{o'wise} \end{cases}
\end{aligned}$$

As a rewriting step, this commuting conversion is a highly unusual rule, the substitution involves a replacement of a term for a term, rather than the more familiar substitution of a term for a variable. Certainly one would hope not to implement this operation in practice. Fortunately the treatment given in later sections removes the need to implement this substitution, in contrast to the framework given by Ong and Stewart (1997).

Term will be assumed to have been written so that all forms of substitution are non-capturing.

3 A Computational Interpretation

In contrast to the situation for the λ -calculus, there is little attention in the literature to computational aspects of the $\lambda\mu$ -calculus. How do programs execute? How do we handle different evaluation orders? What is the computational significance of having two distinct variable spaces? How can we reason about programs? At the time of writing only the paper by Ong and Stewart (1997) addresses these sorts of questions. This paper is an attempt to provide an alternative, and hopefully simpler, answer to these questions.

Before presenting this approach I need first to introduce some standard terminology from work on control operators. To formalise the notion of an evaluation order, Felleisen and Friedman (1986) defined an *evaluation context*. This is essentially a term with a single ‘hole’ in it, written $E[\bullet]$ (this will be defined formally in the next section). The result of placing a term, M , in that hole is written $E[M]$. The idea is that the hole sits at the place where reduction will next occur. In other words, evaluation contexts are devised so that every closed term, M , is either a value (canonical) or can be written *uniquely* as $E[R]$, where R is a redex. The context $E[\bullet]$ essentially represents the rest of the computation that remains to be done after R has been reduced. In this sense it can be seen as the *continuation* of R and is often referred to as the *current continuation*.

Evaluation is then written as

$$(E[R], \mathcal{E}) \Rightarrow (M', \mathcal{E}')$$

where \mathcal{E} is a function from μ -variables to evaluation contexts—the need for this will become clear. The important evaluation rules are

$$\begin{aligned}
(E[\mu a.M], \mathcal{E}) &\Rightarrow (M, \mathcal{E} \uplus \{a \mapsto E[\bullet]\}) \\
(E[[a]M], \mathcal{E} \uplus \{a \mapsto E'[\bullet]\}) &\Rightarrow (E'[M], \mathcal{E} \uplus \{a \mapsto E'[\bullet]\});
\end{aligned}$$

where $\mathcal{E} \uplus \{a \mapsto E[\bullet]\}$ denotes the extension of the function \mathcal{E} with the mapping $a \mapsto E[\bullet]$. Thus in the first reduction rule the current continuation is ‘saved’ by adding it to \mathcal{E} , indexed with a . In the second reduction rule the current continuation is thrown away and the appropriate indexed continuation is restored from \mathcal{E} . In summary, the Activate and Passivate rules are interpreted as (indexed) *save* and *restore* operators, respectively.²

²The reader familiar with control operators will recognise the save operation as a form of ‘catch’ and restore

4 μ PCF

Rather than develop an operational theory for the $\lambda\mu$ -calculus, I shall first enrich it with natural numbers, a conditional, pairs and recursion. This is essentially what Ong and Stewart (1997) call μ PCF. Thus μ PCF types are given by the grammar

$$\begin{array}{l} \phi ::= \mathbf{nat} \\ | \perp \\ | \phi \rightarrow \phi \\ | \phi \times \phi \end{array}$$

and the additional typing rules are as follows.

$$\begin{array}{c} \frac{}{\Gamma \triangleright \underline{n}: \mathbf{nat}, \Sigma} \text{Nat} \quad \frac{\Gamma \triangleright M: \mathbf{nat}, \Sigma}{\Gamma \triangleright \text{suc}(M): \mathbf{nat}, \Sigma} \text{Suc} \\ \\ \frac{\Gamma \triangleright M: \mathbf{nat}, \Sigma \quad \Gamma \triangleright N: \phi, \Sigma \quad \Gamma \triangleright P: \phi, \Sigma}{\Gamma \triangleright \text{ifz } M \text{ then } N \text{ else } P: \phi, \Sigma} \text{Conditional} \\ \\ \frac{\Gamma \triangleright M: \phi, \Sigma \quad \Gamma \triangleright N: \psi, \Sigma}{\Gamma \triangleright \langle M, N \rangle: \phi \times \psi, \Sigma} \times_{\mathcal{I}} \quad \frac{\Gamma \triangleright M: \phi \times \psi, \Sigma}{\Gamma \triangleright \text{fst}(M): \phi, \Sigma} \times_{\mathcal{E}} \quad \frac{\Gamma \triangleright M: \phi \times \psi, \Sigma}{\Gamma \triangleright \text{snd}(M): \psi, \Sigma} \times_{\mathcal{E}} \\ \\ \frac{\Gamma, f: \phi \rightarrow \phi, x: \phi \triangleright M: \phi, \Sigma \quad \Gamma, f: \phi \rightarrow \phi \triangleright N: \psi, \Sigma}{\Gamma \triangleright \text{letrec } f = \lambda x.M \text{ in } N: \psi, \Sigma} \text{Recursion} \end{array}$$

The next step is to choose an evaluation strategy. Most work on control operators has considered a *call-by-value* strategy and to aid comparison I shall adopt the same. (This means that recursion is restricted to function definitions in the usual way (Winskel 1993, §11.1).) It is important to note that what is developed in this and following sections can easily be adjusted to reflect a call-by-name strategy; some details are given in §9. This is in contrast with Ong and Stewart's framework, which requires significant changes to move from call-by-name to call-by-value (some details are given in their paper (Ong and Stewart 1997)).

The syntactic classes of values, evaluation contexts and redexes are defined as follows.

$$\begin{array}{ll} \text{Values} & v ::= \underline{n} \mid \lambda x.M \mid \langle v, v \rangle \\ \\ \text{Evaluation Contexts} & E ::= \bullet \\ & | vE \mid EM \\ & | \langle E, M \rangle \mid \langle v, E \rangle \\ & | \text{fst}(E) \mid \text{snd}(E) \\ & | \text{suc}(E) \\ & | \text{ifz } E \text{ then } M \text{ else } M \\ \\ \text{Redexes} & R ::= vv \\ & | \text{fst}(v) \mid \text{snd}(v) \\ & | \text{suc}(v) \\ & | \text{ifz } v \text{ then } M \text{ else } M \\ & | \text{letrec } f = \lambda x.M \text{ in } N \\ & | [a]M \mid \mu a.M \end{array}$$

The fundamental property of evaluation contexts is the following.

as a form of 'throw'.

Lemma 1 *Every closed term, M , is either a value, v , or is uniquely of the form $E[R]$, where $E[\bullet]$ is an evaluation context and R is a redex.*

We can now write out the (single-step) reduction rules in full, which are as follows.

$$\begin{aligned}
(E[(\lambda x.M)v], \mathcal{E}) &\Rightarrow (E[M[x := v]], \mathcal{E}) \\
(E[\text{fst}(\langle v, w \rangle)], \mathcal{E}) &\Rightarrow (E[v], \mathcal{E}) \\
(E[\text{snd}(\langle v, w \rangle)], \mathcal{E}) &\Rightarrow (E[w], \mathcal{E}) \\
(E[\text{suc}(\underline{n})], \mathcal{E}) &\Rightarrow (E[\underline{n+1}], \mathcal{E}) \\
(E[\text{ifz } \underline{0} \text{ then } M \text{ else } N], \mathcal{E}) &\Rightarrow (E[M], \mathcal{E}) \\
(E[\text{ifz } (\underline{n+1}) \text{ then } M \text{ else } N], \mathcal{E}) &\Rightarrow (E[N], \mathcal{E}) \\
(E[\text{letrec } f = \lambda x.M \text{ in } N], \mathcal{E}) &\Rightarrow (E[N[f := \lambda x.\text{letrec } f = \lambda x.M \text{ in } M]], \mathcal{E}) \\
(E[\mu a.M], \mathcal{E}) &\Rightarrow (M, \mathcal{E} \uplus \{a \mapsto E[\bullet]\}) \\
(E[[a]M], \mathcal{E} \uplus \{a \mapsto E'[\bullet]\}) &\Rightarrow (E'[M], \mathcal{E} \uplus \{a \mapsto E'[\bullet]\})
\end{aligned}$$

5 Examples

In this section I give a number of examples of μ PCF-programs to give the reader a feel for the computational power of the language. In particular, in §5.4, I shall reconsider the examples of encodings given by Ong and Stewart (1997).

5.1 Idealised Scheme

Felleisen and Friedman (1986) presented an extension to the (untyped) call-by-value λ -calculus, called *Idealised Scheme*. Two new operators are added, written $\mathcal{A}(M)$ and $\mathcal{C}(M)$, which are called abort and control, respectively. Both forms are considered to be redexes and their reduction behaviour is given by the following rules.

$$\begin{aligned}
E[\mathcal{A}(M)] &\rightsquigarrow_{\mathcal{A}} M \\
E[\mathcal{C}(M)] &\rightsquigarrow_{\mathcal{C}} M(\lambda z.\mathcal{A}(E[z]))
\end{aligned}$$

Informally $\mathcal{A}(M)$ abandons the current continuation, $E[\bullet]$. $\mathcal{C}(M)$ also abandons the current continuation and M is applied to the abstraction of the current continuation. If this abstraction is invoked with the value v in an evaluation context $E_1[\bullet]$, then $E_1[\bullet]$ will be abandoned and evaluation will continue with $E[v]$.

More concretely, consider the term

$$E_0[\mathcal{C}(\lambda j.M)].$$

Evaluation of this term can be thought of as a ‘catch’ which labels the current continuation $E_0[\bullet]$ with j . If j doesn’t get used in the evaluation of M then $E_0[\bullet]$ is effectively garbaged. If an application of j occurs in the evaluation of M , e.g. $E_1[jv]$, then the computation is ‘thrown’ back to the evaluation context labelled with j , along with the value v . Thus computation continues with the term $E_0[v]$.

Griffin (1990) showed that control and abort can be typed as follows.

$$\frac{\Gamma \triangleright M: (\phi \rightarrow \perp) \rightarrow \perp}{\Gamma \triangleright \mathcal{C}_\phi(M): \phi} \text{Control} \quad \frac{\Gamma \triangleright M: \perp}{\Gamma \triangleright \mathcal{A}_\phi(M): \phi} \text{Abort}$$

Thus control corresponds to the double negation elimination rule proposed by Gentzen (1969) for classical logic, and abort corresponds to the (intuitionistic) \perp -elimination rule.

With these typings in mind, these operators can be encoded in μ PCF as follows.

$$\begin{aligned} \mathcal{C}_\phi(M) &\stackrel{\text{def}}{=} \mu a: \phi. M (\lambda z: \phi. [a: \phi]z) \\ \mathcal{A}_\phi(M) &\stackrel{\text{def}}{=} \mu b: \phi. M \end{aligned}$$

(In both encodings the μ -variables are assumed to be fresh.) Consider an evaluation of the encoding of an Abort in μ PCF.

$$\begin{aligned} &(E_0[\mathcal{A}_\phi(M)], \mathcal{E}) \\ &\stackrel{\text{def}}{=} (E_0[\mu b: \phi. M], \mathcal{E}) \\ &\Rightarrow (M, \mathcal{E} \uplus \{b \mapsto E_0[\bullet]\}) \end{aligned}$$

As b is assumed to be fresh, then it is easy to see that the evaluation context $E_0[\bullet]$ is effectively abandoned.

Consider an evaluation of the encoding of an application of Control in μ PCF.

$$\begin{aligned} &(E_0[\mathcal{C}_\phi(\lambda j. M)], \mathcal{E}) \\ &\stackrel{\text{def}}{=} (E_0[\mu a: \phi. (\lambda j. M) (\lambda z: \phi. [a: \phi]z)], \mathcal{E}) \\ &\Rightarrow ((\lambda j. M) (\lambda z: \phi. [a: \phi]z), \mathcal{E} \uplus \{a \mapsto E_0[\bullet]\}) \\ &\Rightarrow (M [j := (\lambda z: \phi. [a: \phi]z)], \mathcal{E} \uplus \{a \mapsto E_0[\bullet]\}) \\ &\vdots \\ &(E_1[(jv) [j := (\lambda z: \phi. [a: \phi]z)]], \mathcal{E} \uplus \{a \mapsto E_0[\bullet]\}) \quad (\dagger) \\ &\stackrel{\text{def}}{=} (E_1[(\lambda z: \phi. [a: \phi]z)v], \mathcal{E} \uplus \{a \mapsto E_0[\bullet]\}) \\ &\Rightarrow (E_1[[a: \phi]v], \mathcal{E} \uplus \{a \mapsto E_0[\bullet]\}) \\ &\Rightarrow (E_0[v], \mathcal{E} \uplus \{a \mapsto E_0[\bullet]\}) \end{aligned}$$

At stage (\dagger) the substitution has been left explicit to aid comparison with the earlier discussion of the control operator, \mathcal{C} , in Idealised Scheme. Hopefully it is clear that the effect of this encoding is that evaluation has been thrown back to E_0 , along with the value v .

Remark. It is also possible to encode the Scheme variant of the Control operator, which does not initially abandon the current continuation (Clinger and Rees 1986). This rule, as observed by Griffin, is typed using the Peirce formula, i.e.

$$\frac{\Gamma \triangleright M: (\phi \rightarrow \psi) \rightarrow \phi}{\Gamma \triangleright \mathcal{P}_{\phi, \psi}(M): \phi}$$

Its reduction rule is

$$E[\mathcal{P}_{\phi, \psi}(M)] \rightsquigarrow_{\mathcal{P}} E[M (\lambda z: \phi. \mathcal{A}_\psi(E[z]))]$$

and it can be encoded into μ PCF as follows.

$$\mathcal{P}_{\phi, \psi}(M) \stackrel{\text{def}}{=} \mu b: \phi. [b: \phi]M (\lambda x: \phi. \mu a: \psi. [b: \phi]x)$$

It is left as an exercise to the reader to verify that this encoding has the expected operational behaviour.

5.2 de Groote's Exception Handling Calculus

de Groote (1995) presented a simply-typed λ -calculus extended with an ML-like exception handling mechanism. A new class of exception variable is introduced and typing judgements

are of the form $\Gamma; \Delta \triangleright M: \phi$ where Γ is the typing environment for normal λ -variables and Δ the typing environment for exception variables. Two new term constructors are introduced whose typing rules are as follows.

$$\frac{\Gamma; \Delta \triangleright M: \phi}{\Gamma; \Delta, e: \phi \rightarrow \perp \triangleright \text{raise}(e, M): \psi} \text{Raise} \quad \frac{\Gamma; \Delta, e: \neg\phi \triangleright M: \psi \quad \Gamma, x: \phi; \Delta \triangleright N: \psi}{\Gamma; \Delta \triangleright \text{let } e \text{ in } M \text{ handle } e x \Rightarrow N: \psi} \text{Handle}$$

(where $\neg\phi \stackrel{\text{def}}{=} \phi \rightarrow \perp$). As de Groote notes, the typing rule *Raise* corresponds to the standard (intuitionistic) rules for falsity. The typing rule *Handle* is clearly related to classical logic: it corresponds to the rule of the excluded middle.

Associated with these rules are a number of reduction rules.³

$$\begin{aligned} v(\text{raise}(e, w)) &\rightsquigarrow \text{raise}(e, w) \\ (\text{raise}(e, v))M &\rightsquigarrow \text{raise}(e, v) \\ \text{raise}(e, \text{raise}(e', v)) &\rightsquigarrow \text{raise}(e', v) \\ \text{let } e \text{ in } v \text{ handle } e x \Rightarrow N &\rightsquigarrow v && (e \notin \text{FV}(v)) \\ \text{let } e \text{ in } \text{raise}(e, v) \text{ handle } e x \Rightarrow N &\rightsquigarrow N[x := v] && (e \notin \text{FV}(v, N)) \\ \text{let } e \text{ in } \text{raise}(e', v) \text{ handle } e x \Rightarrow N &\rightsquigarrow \text{raise}(e', v) && (e \notin \text{FV}(v)) \end{aligned}$$

This exception handling mechanism can be encoded into μ PCF quite simply, as follows.

$$\begin{aligned} \llbracket \text{raise}(e, M) \rrbracket &\stackrel{\text{def}}{=} (\lambda x. \mu b. [e]x) \llbracket M \rrbracket \\ \llbracket \text{let } e \text{ in } M \text{ handle } e x \Rightarrow N \rrbracket &\stackrel{\text{def}}{=} \mu b. [b](\lambda x. \llbracket N \rrbracket)(\mu e. [b] \llbracket M \rrbracket) \end{aligned}$$

It is quite easy to verify that this translation preserves the expected operational behaviour. Here are two examples.

$$\begin{aligned} &(E[\llbracket \text{let } e \text{ in } v \text{ handle } e x \Rightarrow N \rrbracket], \mathcal{E}) \\ \stackrel{\text{def}}{=} &(E[\mu b. [b](\lambda x. \llbracket N \rrbracket)(\mu e. [b] \llbracket v \rrbracket)], \mathcal{E}) \\ \Rightarrow^2 &(E[(\lambda x. \llbracket N \rrbracket)(\mu e. [b] \llbracket v \rrbracket)], \mathcal{E} \uplus \{b \mapsto E[\bullet]\}) \\ \Rightarrow &([b] \llbracket v \rrbracket, \mathcal{E} \uplus \{b \mapsto E[\bullet], e \mapsto E[(\lambda x. \llbracket N \rrbracket) \bullet]\}) \\ \Rightarrow &(E[\llbracket v \rrbracket], \mathcal{E} \uplus \{b \mapsto E[\bullet], e \mapsto E[(\lambda x. \llbracket N \rrbracket) \bullet]\}) \end{aligned}$$

$$\begin{aligned} &(E[\llbracket \text{let } e \text{ in } \text{raise}(e, v) \text{ handle } e x \Rightarrow N \rrbracket], \mathcal{E}) \\ \stackrel{\text{def}}{=} &(E[\mu b. [b](\lambda x. \llbracket N \rrbracket)(\mu e. [b](\lambda x. \mu c. [e]x) \llbracket v \rrbracket)], \mathcal{E}) \\ \Rightarrow^2 &(E[(\lambda x. \llbracket N \rrbracket)(\mu e. [b](\lambda x. \mu c. [e]x) \llbracket v \rrbracket)], \mathcal{E} \uplus \{b \mapsto E[\bullet]\}) \\ \Rightarrow &([b](\lambda x. \mu c. [e]x) \llbracket v \rrbracket, \mathcal{E} \uplus \{b \mapsto E[\bullet], e \mapsto E[(\lambda x. \llbracket N \rrbracket) \bullet]\}) \\ \Rightarrow &(E[(\lambda x. \mu c. [e]x) \llbracket v \rrbracket], \mathcal{E} \uplus \{b \mapsto E[\bullet], e \mapsto E[(\lambda x. \llbracket N \rrbracket) \bullet]\}) \\ \Rightarrow &(E[\mu c. [e] \llbracket v \rrbracket], \mathcal{E} \uplus \{b \mapsto E[\bullet], e \mapsto E[(\lambda x. \llbracket N \rrbracket) \bullet]\}) \\ \Rightarrow &([e] \llbracket v \rrbracket, \mathcal{E} \uplus \{b \mapsto E[\bullet], e \mapsto E[(\lambda x. \llbracket N \rrbracket) \bullet]c \mapsto E[\bullet]\}) \\ \Rightarrow &(E[(\lambda x. \llbracket N \rrbracket) \llbracket v \rrbracket], \mathcal{E} \uplus \{b \mapsto E[\bullet], e \mapsto E[(\lambda x. \llbracket N \rrbracket) \bullet]c \mapsto E[\bullet]\}) \\ \Rightarrow &(E[\llbracket N \rrbracket][x := \llbracket v \rrbracket], \mathcal{E} \uplus \{b \mapsto E[\bullet], e \mapsto E[(\lambda x. \llbracket N \rrbracket) \bullet]c \mapsto E[\bullet]\}) \end{aligned}$$

Remark. The reader familiar with SML will recognise that this exception handling mechanism is less powerful than that in SML. This is to be expected as it known that the typed λ -calculus extended with SML exception handling can encode the untyped λ -calculus, and thus is not strongly normalising (Lillibridge 1995). As the typed $\lambda\mu$ -calculus is strongly normalising, it clearly can not encode true SML exception handling.

³de Groote also gives a second, more complicated, set of reduction rules.

5.3 Pairing

It is easy to verify that $\phi \times \psi \equiv \neg(\phi \rightarrow \neg\psi)$ in **CL**. This logical equivalence can be used to simulate pairing in μ PCF. The constructor and deconstructors are encoded as follows.

$$\begin{aligned} \text{pair} &\stackrel{\text{def}}{=} \lambda m: \phi. \lambda n: \psi. \lambda f: (\phi \rightarrow (\psi \rightarrow \perp)). f \ m \ n \\ \text{fst} &\stackrel{\text{def}}{=} \lambda p. \mu a. p(\lambda x. \mu b. [a]x) \\ \text{snd} &\stackrel{\text{def}}{=} \lambda p. \mu a. p(\lambda y. \lambda x. [a]x) \end{aligned}$$

It is simple to see that these encodings satisfy the expected (call-by-value) behaviour, e.g.

$$\begin{aligned} &(\text{fst} (\text{pair} \ v \ w), \mathcal{E}) \\ \stackrel{\text{def}}{=} &(\text{fst} ((\lambda m n f. f \ m \ n) \ v \ w), \mathcal{E}) \\ \Rightarrow^2 &(\text{fst} (\lambda f. f \ v \ w), \mathcal{E}) \\ \Rightarrow &(\mu a. ((\lambda f. f \ v \ w) (\lambda x. \mu b. [a]x)), \mathcal{E}) \\ \Rightarrow &((\lambda f. f \ v \ w) (\lambda x. \mu b. [a]x), \mathcal{E} \uplus \{a \mapsto \bullet\}) \\ \Rightarrow &((\lambda x. \mu b. [a]x) \ v \ w, \mathcal{E} \uplus \{a \mapsto \bullet\}) \\ \Rightarrow &((\mu b. [a]v) \ w, \mathcal{E} \uplus \{a \mapsto \bullet\}) \\ \Rightarrow &([a]v, \mathcal{E} \uplus \{a \mapsto \bullet, b \mapsto (\bullet w)\}) \\ \Rightarrow &(v, \mathcal{E} \uplus \{a \mapsto \bullet, b \mapsto (\bullet w)\}) \end{aligned}$$

5.4 Ong/Stewart Encodings

In their paper, Ong and Stewart (1997) give variants of callcc and exceptions and show how they can be encoded in μ PCF. I shall reconsider these encodings in the light of the simpler operational treatment offered by this paper.

5.4.1 Exceptions

Ong and Stewart considered PCF extended with a simple system of exception handling, which was inspired by (but less powerful than) work by Gunter, Rémy and Riecke (1995). (Their system is actually quite similar to de Groote's system, considered in §5.2.) Typed exceptions are identified with names, thus typing judgements are now of the form $\Gamma; \Delta \triangleright M: \phi$ where Γ is the usual typing environment and Δ is the typing environment for the exception names. Two new operators are added to PCF whose typing rules are as follows.

$$\frac{\Gamma; \Delta \triangleright M: \phi}{\Gamma; \Delta, a: \phi \triangleright \text{raise}(a, M): \psi} \quad \frac{\Gamma; \Delta, a: \phi \triangleright M: \phi \rightarrow \psi \quad \Gamma; \Delta, a: \phi \triangleright N: \psi}{\Gamma; \Delta \triangleright \text{handle}(a, M, N): \psi}$$

The intended interpretation is that the term $\text{raise}(a, M)$ first evaluates M to a value v and then raises an exception named a associated with v . The term $\text{handle}(a, M, N)$ evaluates M to a value (say v) and then evaluates N . If N evaluates to a value w then this is the overall result, but if it raises an exception named a with a value u , then this is applied to v . Given as reduction rules the intended interpretation is as follows.

$$\begin{aligned} \text{handle}(a, v, w) &\rightsquigarrow w && (a \notin \text{FN}(w)) \\ \text{handle}(a, v, E[\text{raise}(a, u)]) &\rightsquigarrow vu && (a \notin \text{FN}(v, u)) \end{aligned}$$

These operators can be translated into μ PCF as follows (where b is a fresh μ -variable).

$$\begin{aligned} \llbracket \text{raise}(a, M) \rrbracket &\stackrel{\text{def}}{=} (\lambda x. \mu b. [a]x) \llbracket M \rrbracket \\ \llbracket \text{handle}(a, M, N) \rrbracket &\stackrel{\text{def}}{=} \mu b. [b] \llbracket M \rrbracket (\mu a. [b] \llbracket N \rrbracket) \end{aligned}$$

It is relatively simple to show that this translation preserves the operational behaviour, e.g. (where $\llbracket M \rrbracket \Rightarrow^* v$ and $\llbracket N \rrbracket \Rightarrow^* u$)

$$\begin{aligned}
& (\llbracket \text{handle}(a, M, E[\text{raise}(a, N)]) \rrbracket, \mathcal{E}) \\
\stackrel{\text{def}}{=} & (\mu b.[b]\llbracket M \rrbracket (\mu a.[b]E[(\lambda x.\mu c.[a]x)\llbracket N \rrbracket]), \mathcal{E}) \\
\Rightarrow^2 & (\llbracket M \rrbracket (\mu a.[b]E[(\lambda x.\mu c.[a]x)\llbracket N \rrbracket]), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\
\Rightarrow^* & (v(\mu a.[b]E[(\lambda x.\mu c.[a]x)\llbracket N \rrbracket]), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\
\Rightarrow & ([b]E[(\lambda x.\mu c.[a]x)\llbracket N \rrbracket], \mathcal{E} \uplus \{a \mapsto (v\bullet), b \mapsto \bullet\}) \\
\Rightarrow & (E[(\lambda x.\mu c.[a]x)\llbracket N \rrbracket], \mathcal{E} \uplus \{a \mapsto (v\bullet), b \mapsto \bullet\}) \\
\Rightarrow^+ & (E[\mu c.[a]u], \mathcal{E} \uplus \{a \mapsto (v\bullet), b \mapsto \bullet\}) \\
\Rightarrow & ([a]u, \mathcal{E} \uplus \{a \mapsto (v\bullet), b \mapsto \bullet, c \mapsto E[\bullet]\}) \\
\Rightarrow & (vu, \mathcal{E} \uplus \{a \mapsto (v\bullet), b \mapsto \bullet, c \mapsto E[\bullet]\})
\end{aligned}$$

5.4.2 Callcc

Ong and Stewart also considered PCF extended with a variant of callcc (again inspired by work by Gunter et al. (1995)). Continuations are both typed and associated with names, and so typing judgements are of the form $\Gamma; \Delta \triangleright M: \phi$, where Δ is the typing environment for continuation names. Three new operators are added to PCF, whose typing rules are as follows.

$$\frac{\Gamma; \Delta \triangleright M: (\phi \rightarrow \psi) \rightarrow \phi}{\Gamma; \Delta \triangleright \text{callcc}(M): \phi} \quad \frac{\Gamma; \Delta \triangleright M: \phi}{\Gamma; \Delta, a: \phi \triangleright \text{abort}(a, M): \psi} \quad \frac{\Gamma; \Delta, a: \phi \triangleright M: \phi}{\Gamma; \Delta \triangleright \text{set}(a, M): \phi}$$

The *callcc* operator applies the term M to an abstraction of the current continuation. The *set* serves as a delimiter for continuations, and the *abort* discards the current continuation (delimited by a). Given as reduction rules their intended operational behaviour is as follows.

$$\begin{aligned}
\text{set}(a, E[\text{abort}(a, M)]) & \rightsquigarrow M & (a \notin \text{FN}(M)) \\
\text{set}(a, v) & \rightsquigarrow v & (a \notin \text{FN}(v)) \\
E[\text{callcc}(M)] & \rightsquigarrow \text{set}(a, E[M(\lambda x.\text{abort}(a, E[x]))])
\end{aligned}$$

These operators can be translated into μ PCF as follows.

$$\begin{aligned}
\llbracket \text{callcc}(M) \rrbracket & \stackrel{\text{def}}{=} \mu a.[a](\llbracket M \rrbracket (\lambda x.\mu b.[a]x)) \\
\llbracket \text{abort}(a, M) \rrbracket & \stackrel{\text{def}}{=} \mu b.[a]\llbracket M \rrbracket \quad \text{where } b \notin \text{FN}(\llbracket M \rrbracket) \\
\llbracket \text{set}(a, M) \rrbracket & \stackrel{\text{def}}{=} \mu a.[a]\llbracket M \rrbracket
\end{aligned}$$

Again it is simple to check that this translation preserves the operational behaviour, e.g.

$$\begin{aligned}
& (\llbracket \text{set}(a, E[\text{abort}(a, M)]) \rrbracket, \mathcal{E}) \\
\stackrel{\text{def}}{=} & (\mu a.[a]E[\mu b.[a]\llbracket M \rrbracket], \mathcal{E}) \\
\Rightarrow^2 & (E[\mu b.[a]\llbracket M \rrbracket], \mathcal{E} \uplus \{a \mapsto \bullet\}) \\
\Rightarrow & ([a]\llbracket M \rrbracket, \mathcal{E} \uplus \{a \mapsto \bullet, b \mapsto E[\bullet]\}) \\
\Rightarrow & (\llbracket M \rrbracket, \mathcal{E} \uplus \{a \mapsto \bullet, b \mapsto E[\bullet]\})
\end{aligned}$$

6 A Transition System

An implementation based on the reduction rules given in §4 would work as follows. Take a term M : if it is a value then we are done; if not, it can be given uniquely as $E[R]$. One takes the relevant reduction step (determined by R)—the resulting term is either a value, in

which case we are done, or it has to be re-written again as an evaluation context and a redex. This process is repeated until a value is reached. The continual intermediate step of rewriting a term into an evaluation context and a redex would be inefficient in practice and is quite cumbersome theoretically. Consequently I shall give a new set of reduction rules where the evaluation context and the redex are actually separated. Reduction rules are now of the form

$$(S, M, \mathcal{E}) \longrightarrow (S', M', \mathcal{E}')$$

where S is a stack of *evaluation frames*, which are defined as follows.

$$F ::= \begin{array}{l|l} \bullet M & v \bullet \\ \langle \bullet, M \rangle & \langle v, \bullet \rangle \\ \text{fst}(\bullet) & \text{snd}(\bullet) \\ \text{suc}(\bullet) & \text{ifz } \bullet \text{ then } M \text{ else } M \end{array}$$

Stacks are just compositions of frames— \circ will denote (right-to-left) composition and id denotes the empty stack. Clearly \mathcal{E} is now a function from μ -variables to stacks. The reduction rules essentially describe the transitions of a simple abstract machine. Harper and Stone (1997) use similar transition rules in their analysis of SML and Pitts (1997a) has also used similar rules in work on functional languages with dynamic allocation of store and in work on parametric polymorphism (Pitts 1998). In full the transition rules are as follows.

$$\begin{array}{lll} (S \circ F[\bullet], v, \mathcal{E}) & \longrightarrow & (S, F[v], \mathcal{E}) \\ (S, MN, \mathcal{E}) & \longrightarrow & (S \circ (\bullet N), M, \mathcal{E}) & M \text{ not a value} \\ (S, vN, \mathcal{E}) & \longrightarrow & (S \circ (v \bullet), N, \mathcal{E}) & N \text{ not a value} \\ (S, (\lambda x.M)v, \mathcal{E}) & \longrightarrow & (S, M[x := v], \mathcal{E}) \\ (S, \langle M, N \rangle, \mathcal{E}) & \longrightarrow & (S \circ \langle \bullet, N \rangle, M, \mathcal{E}) & M \text{ not a value} \\ (S, \langle v, N \rangle, \mathcal{E}) & \longrightarrow & (S \circ \langle v, \bullet \rangle, N, \mathcal{E}) & N \text{ not a value} \\ (S, \text{fst}(M), \mathcal{E}) & \longrightarrow & (S \circ \text{fst}(\bullet), M, \mathcal{E}) & M \text{ not a value} \\ (S, \text{fst}(\langle v, w \rangle), \mathcal{E}) & \longrightarrow & (S, v, \mathcal{E}) \\ (S, \text{snd}(M), \mathcal{E}) & \longrightarrow & (S \circ \text{snd}(\bullet), M, \mathcal{E}) & M \text{ not a value} \\ (S, \text{snd}(\langle v, w \rangle), \mathcal{E}) & \longrightarrow & (S, w, \mathcal{E}) \\ (S, \text{suc}(M), \mathcal{E}) & \longrightarrow & (S \circ \text{suc}(\bullet), M, \mathcal{E}) & M \text{ not a value} \\ (S, \text{suc}(\underline{n}), \mathcal{E}) & \longrightarrow & (S, \underline{n+1}, \mathcal{E}) \\ (S, \text{ifz } M \text{ then } N \text{ else } P, \mathcal{E}) & \longrightarrow & (S \circ (\text{ifz } \bullet \text{ then } N \text{ else } P), M, \mathcal{E}) & M \text{ not a value} \\ (S, \text{ifz } \underline{0} \text{ then } M \text{ else } N, \mathcal{E}) & \longrightarrow & (S, M, \mathcal{E}) \\ (S, \text{ifz } \underline{n+1} \text{ then } M \text{ else } N, \mathcal{E}) & \longrightarrow & (S, N, \mathcal{E}) \\ (S, \text{letrec } f = \lambda x.M \text{ in } N, \mathcal{E}) & \longrightarrow & (S, N[f := \lambda x.\text{letrec } f = \lambda x.M \text{ in } M], \mathcal{E}) \\ (S, \mu a.M, \mathcal{E}) & \longrightarrow & (id, M, \mathcal{E} \uplus \{a \mapsto S\}) \\ (S, [a]M, \mathcal{E} \uplus \{a \mapsto T\}) & \longrightarrow & (T, M, \mathcal{E} \uplus \{a \mapsto T\}) \end{array}$$

An example may make these reduction rules clearer. Consider an instance of the Ong/Stewart ‘callcc’ reduction rule given in §5.4.2.

$$\text{set}(a, (\lambda x.N)(\text{abort}(a, M))) \rightsquigarrow M$$

The left hand term is translated into the following μ PCF-term.

$$\mu a.[a](\lambda x. \llbracket N \rrbracket)(\mu b.[a] \llbracket M \rrbracket)$$

The reduction of this term is as follows.

$$\begin{array}{l}
\rightarrow (S, \mu a.[a](\lambda x. \llbracket N \rrbracket)(\mu b.[a] \llbracket M \rrbracket), \mathcal{E}) \\
\rightarrow (id, [a](\lambda x. \llbracket N \rrbracket)(\mu b.[a] \llbracket M \rrbracket), \mathcal{E} \uplus \{a \mapsto S\}) \\
\rightarrow (S, (\lambda x. \llbracket N \rrbracket)(\mu b.[a] \llbracket M \rrbracket), \mathcal{E} \uplus \{a \mapsto S\}) \\
\rightarrow (S \circ ((\lambda x. \llbracket N \rrbracket) \bullet), \mu b.[a] \llbracket M \rrbracket, \mathcal{E} \uplus \{a \mapsto S\}) \\
\rightarrow (id, [a] \llbracket M \rrbracket, \mathcal{E} \uplus \{a \mapsto S, b \mapsto S \circ ((\lambda x. \llbracket N \rrbracket) \bullet)\}) \\
\rightarrow (S, \llbracket M \rrbracket, \mathcal{E} \uplus \{a \mapsto S, b \mapsto S \circ ((\lambda x. \llbracket N \rrbracket) \bullet)\})
\end{array}$$

To compare these two formulations I shall make use of a function $\lceil E \rceil$ which converts an evaluation context, E , to a stack of frames, and a function $S @ M$ which takes a stack of frames, S , and a term, M , and converts the stack back to an evaluation context before inserting M . The definitions are routine and omitted. Here are a couple of examples (where the identity $id \circ F \equiv F$ has been used).

$$\begin{aligned}
\lceil ((\lambda x.M) \bullet) P \rceil &\stackrel{\text{def}}{=} ((\bullet Q) \circ (\bullet P)) \circ ((\lambda x.M) \bullet) \\
((\bullet Q) \circ (\bullet P)) \circ ((\lambda x.M) \bullet) @ N &\stackrel{\text{def}}{=} ((\lambda x.M) N) P Q
\end{aligned}$$

The two sets of reduction rules can be related in the following sense.

Proposition 1 $(S @ M, \mathcal{E}) \Rightarrow (N, \mathcal{E}')$ iff $\exists S', M'. N = S' @ M', (S, M, \lceil \mathcal{E} \rceil) \rightarrow^* (S', M', \lceil \mathcal{E}' \rceil)$

A useful fact is that the set

$$\searrow \stackrel{\text{def}}{=} \{(S, M, \mathcal{E}) \mid \exists v, \mathcal{E}'. (S, M, \mathcal{E}) \rightarrow^* (id, v, \mathcal{E}')\}$$

has a direct, inductive definition which is as follows.

$$\begin{array}{l}
(id, v, \mathcal{E}) \searrow \frac{(S, F[v], \mathcal{E}) \searrow}{(S \circ F[\bullet], v, \mathcal{E}) \searrow} \\
\frac{(S \circ (\bullet N), M, \mathcal{E}) \searrow}{(S, MN, \mathcal{E}) \searrow} M \text{ not a value} \quad \frac{(S \circ (v \bullet), N, \mathcal{E}) \searrow}{(S, vN, \mathcal{E}) \searrow} N \text{ not a value} \\
\frac{(S, M[x := v], \mathcal{E}) \searrow}{(S, (\lambda x.M)v, \mathcal{E}) \searrow} \quad \frac{(S, N[f := \lambda x. \text{letrec } f = \lambda x.M \text{ in } M], \mathcal{E}) \searrow}{(S, \text{letrec } f = \lambda x.M \text{ in } N, \mathcal{E}) \searrow} \\
\frac{(S \circ \langle \bullet, N \rangle, M, \mathcal{E}) \searrow}{(S, \langle M, N \rangle, \mathcal{E}) \searrow} M \text{ not a value} \quad \frac{(S \circ \langle v, \bullet \rangle, N, \mathcal{E}) \searrow}{(S, \langle v, N \rangle, \mathcal{E}) \searrow} N \text{ not a value} \\
\frac{(S \circ \text{fst}(\bullet), M, \mathcal{E}) \searrow}{(S, \text{fst}(M), \mathcal{E}) \searrow} M \text{ not a value} \quad \frac{(S, v, \mathcal{E}) \searrow}{(S, \text{fst}(\langle v, w \rangle), \mathcal{E}) \searrow} \\
\frac{(S \circ \text{snd}(\bullet), M, \mathcal{E}) \searrow}{(S, \text{snd}(M), \mathcal{E}) \searrow} M \text{ not a value} \quad \frac{(S, w, \mathcal{E}) \searrow}{(S, \text{snd}(\langle v, w \rangle), \mathcal{E}) \searrow} \\
\frac{(T, M, \mathcal{E} \uplus \{a \mapsto T\}) \searrow}{(S, [a]M, \mathcal{E} \uplus \{a \mapsto T\}) \searrow} \quad \frac{(id, M, \mathcal{E} \uplus \{a \mapsto S\}) \searrow}{(S, \mu a.M, \mathcal{E}) \searrow}
\end{array}$$

These rules will form the basis of a notion of program equivalence given in the next section.

7 Program Equivalence

μ PCF is a simple functional programming language for recursively defined, higher order functions along with facilities to save and restore the runtime environment. As for pure functional languages, we should like to develop methods for reasoning about properties of μ PCF programs. One possibility is to devise a set of equations and reason equationally about programs. Such equations have been considered for Idealised Scheme by Felleisen and Hieb (1992), Griffin (1990) and Hofmann (1995). Another possibility is to develop a denotational model and reason about programs via their denotations. Cartwright, Curien and Felleisen (1994) give a (fully abstract) model of Idealised Scheme and some models of the $\lambda\mu$ -calculus have been considered recently by Ong (1996), Hofmann and Streicher (1997) and Selinger (1998a).

In this paper I shall rather consider techniques based on the *operational behaviour* of μ PCF programs. One advantage of such operationally based techniques is that they require relatively little mathematical overhead. (Further arguments in favour of operationally based techniques can be found in the literature, e.g. (Gordon and Pitts 1998).)

Morris-style, contextual equivalence is accepted as a natural notion of equivalence for sequential languages. Essentially two programs are considered contextually equivalent if interchanging one for the other in any larger program does not affect the result. Before the formal definition, here are a couple of simple definitions.

$$\begin{aligned} (M, \mathcal{E}) \Downarrow (v, \mathcal{E}') &\stackrel{\text{def}}{=} (M, \mathcal{E}) \Rightarrow^* (v, \mathcal{E}') \text{ and } (v, \mathcal{E}') \not\Leftarrow \\ (M, \mathcal{E}) \Downarrow v &\stackrel{\text{def}}{=} \exists \mathcal{E}'. (M, \mathcal{E}) \Downarrow (v, \mathcal{E}') \\ (M, \mathcal{E}) \Downarrow &\stackrel{\text{def}}{=} \exists v. (M, \mathcal{E}) \Downarrow v \\ (M, \mathcal{E}) \Uparrow &\stackrel{\text{def}}{=} \neg(\exists v. (M, \mathcal{E}) \Downarrow v) \end{aligned}$$

A context, written \mathcal{C} , is a μ PCF-term with possibly many holes, written \bullet . $\mathcal{C}[M]$ denotes the term (or, more generally, context) obtained by replacing all occurrences of the hole in context \mathcal{C} with the term M . Unlike normal substitution this process may well capture free λ - and μ -variables in M . A context is said to be λ -closing (with respect to a term M) if it captures all the free λ -variables of M . (Clearly the evaluation contexts of §4 are very simple contexts—they are non-capturing and have exactly one hole.)

Definition 1 *Let M and N be terms and \mathcal{C} a λ -closing context. M is said to contextually refine N , written $\Gamma \triangleright M \sqsubseteq N: \phi, \Sigma$, when $\forall \mathcal{C}, \mathcal{E}$. if $(\mathcal{C}[M], \mathcal{E}) \Downarrow$ then $(\mathcal{C}[N], \mathcal{E}) \Downarrow$. They are said to be contextually equivalent, written $\Gamma \triangleright M \approx N: \phi, \Sigma$, just when $\Gamma \triangleright M \sqsubseteq N: \phi, \Sigma$ iff $\Gamma \triangleright N \sqsubseteq M: \phi, \Sigma$.*

Reasoning about languages with control is difficult. One reason for this is that, in terms of contextual equivalence, control is invariably a non-conservative extension. By this, I mean that terms which are contextually equivalent in the call-by-value λ -calculus, may no longer be contextually equivalent in a call-by-value λ -calculus extended with control. For example, consider the following PCF-terms (this example is due to Meyer and Riecke (1988)).

$$\begin{aligned} M_1 &\stackrel{\text{def}}{=} \lambda x. \lambda y. \lambda z. (\lambda w. (y \ x) \ w)(z \ x): \phi \rightarrow (\phi \rightarrow \psi \rightarrow \varphi) \rightarrow (\phi \rightarrow \psi) \rightarrow \varphi \\ M_2 &\stackrel{\text{def}}{=} \lambda x. \lambda y. \lambda z. (y \ x)(z \ x): \phi \rightarrow (\phi \rightarrow \psi \rightarrow \varphi) \rightarrow (\phi \rightarrow \psi) \rightarrow \varphi \end{aligned}$$

It can be seen that these terms are PCF-contextually equivalent: Meyer and Riecke claim an complicated inductive argument can be used, although a more modern approach would be to

show that they are applicatively bisimilar (some details of the relevant definition are given by Pitts (1997b, §6)) and use the fact that applicative bisimilarity coincides with contextual equivalence.

Of course, M_1 and M_2 are valid μ PCF-terms. However they are *not* μ PCF-contextually equivalent: they can be distinguished by the context

$$\mathcal{C} \stackrel{\text{def}}{=} \mu b: \mathbf{nat}. [b: \mathbf{nat}]((\bullet \underline{1})(\lambda u. \Omega^{\mathbf{nat} \rightarrow \mathbf{nat}})(\lambda v. \mu a. [b] \underline{1}))$$

where $\Omega^{\mathbf{nat} \rightarrow \mathbf{nat}}$ is a non-terminating term.⁴ Thus $\mathcal{C}[M_1]$ terminates by the following steps (where, for clarity, I have in places underlined the redex).

$$\begin{aligned} & (\mathcal{C}[M_1], \mathcal{E}) \\ \stackrel{\text{def}}{=} & (\mu b: \mathbf{nat}. [b: \mathbf{nat}]((\underline{M_1} \underline{1})(\lambda u. \Omega)(\lambda v. \mu a. [b] \underline{1})), \mathcal{E}) \\ \Rightarrow^2 & ((\underline{M_1} \underline{1})(\lambda u. \Omega)(\lambda v. \mu a. [b] \underline{1}), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\ \stackrel{\text{def}}{=} & (((\lambda x. \lambda y. \lambda z. (\lambda w. (y \ x) \ w)(z \ x)) \underline{1})(\lambda u. \Omega)(\lambda v. \mu a. [b] \underline{1}), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\ \Rightarrow^3 & ((\lambda w. ((\lambda u. \Omega) \underline{1}) \ w)((\lambda v. \mu a. [b] \underline{1}) \underline{1}), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\ \Rightarrow & ((\lambda w. ((\lambda u. \Omega) \underline{1}) \ w)(\mu a. [b] \underline{1}), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\ \Rightarrow & ([b] \underline{1}, \mathcal{E} \uplus \{b \mapsto \bullet, a \mapsto (\lambda w. ((\lambda u. \Omega) \underline{1}) \ w) \bullet\}) \\ \Rightarrow & (\underline{1}, \mathcal{E} \uplus \{b \mapsto \bullet, a \mapsto (\lambda w. ((\lambda u. \Omega) \underline{1}) \ w) \bullet\}) \end{aligned}$$

Unfortunately $\mathcal{C}[M_2]$ does not terminate. It evaluates as follows.

$$\begin{aligned} & (\mathcal{C}[M_2], \mathcal{E}) \\ \stackrel{\text{def}}{=} & (\mu b: \mathbf{nat}. [b: \mathbf{nat}]((\underline{M_2} \underline{1})(\lambda u. \Omega)(\lambda v. \mu a. [b] \underline{1})), \mathcal{E}) \\ \Rightarrow^2 & ((\underline{M_2} \underline{1})(\lambda u. \Omega)(\lambda v. \mu a. [b] \underline{1}), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\ \stackrel{\text{def}}{=} & (((\lambda x. \lambda y. \lambda z. (y \ x)(z \ x)) \underline{1})(\lambda u. \Omega)(\lambda v. \mu a. [b] \underline{1}), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\ \Rightarrow^3 & (((\lambda u. \Omega) \underline{1})((\lambda v. \mu a. [b] \underline{1}) \underline{1}), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\ \Rightarrow & (\underline{\Omega} ((\lambda v. \mu a. [b] \underline{1}) \underline{1}), \mathcal{E} \uplus \{b \mapsto \bullet\}) \\ \uparrow & \end{aligned}$$

This counter-example is quite serious. Both M_1 and M_2 are closed terms and so they are not distinguished by binding some free μ -variable in some complicated way. In fact they differ by having a slightly different order of evaluation in their bodies. In μ PCF there is enough computational power to differentiate between these different orders of evaluation (a similar thing happens with languages with state (Pitts and Stark 1998)).

Here's another problem when reasoning about μ PCF programs. Typically given two programs of function type, we might think that if they behave in the same way for all given arguments, then they can safely be thought of as indistinguishable, i.e. contextually equivalent. For example, notions of applicative bisimilarity of PCF-programs use this idea in their definition (Abramsky 1990, Pitts 1997b, Gordon 1995). As one may have feared, things are more complicated for μ PCF. Consider the closed programs

$$\begin{aligned} T_1 & \stackrel{\text{def}}{=} \mu a. [a](\lambda y. \mu c. [a](\lambda x. \text{ifz } y \text{ then } \Omega \text{ else } \underline{0})): \mathbf{nat} \rightarrow \mathbf{nat} \\ T_2 & \stackrel{\text{def}}{=} \lambda z. \mu b. [b]((\lambda y. \mu c. [b]((\lambda x. \text{ifz } y \text{ then } \Omega \text{ else } \underline{0})z))z): \mathbf{nat} \rightarrow \mathbf{nat} \end{aligned}$$

where, again, Ω is a non-terminating term. It is easy to verify that for all natural numbers n and \mathcal{E}

$$(T_1 n, \mathcal{E}) \Downarrow \underline{0} \iff (T_2 n, \mathcal{E}) \Downarrow \underline{0}$$

⁴For example $\Omega^{\mathbf{nat} \rightarrow \mathbf{nat}} \stackrel{\text{def}}{=} \text{letrec } f = \lambda x: \mathbf{nat} \rightarrow \mathbf{nat}. f x \text{ in } f(\lambda y: \mathbf{nat}. y)$.

Thus given the argument above, one may conclude that T_1 and T_2 should be thought of as contextually equivalent. Unfortunately they are not, as the context

$$\mathcal{C} \stackrel{\text{def}}{=} (\lambda s.s(\underline{1}))\bullet$$

can distinguish between T_1 and T_2 ! Indeed $\mathcal{C}[T_1]$ terminates in the following way.

$$\begin{aligned} & (\mathcal{C}[T_1], \mathcal{E}) \\ \stackrel{\text{def}}{=} & ((\lambda s.s(\underline{1}))(\mu a.[a](\lambda y.\mu c.[a](\lambda x.\text{ifz } y \text{ then } \Omega \text{ else } \underline{0}))), \mathcal{E}) \\ \Rightarrow^2 & ((\lambda s.s(\underline{1}))(\lambda y.\mu c.[a](\lambda x.\text{ifz } y \text{ then } \Omega \text{ else } \underline{0})), \mathcal{E} \uplus \{a \mapsto (\lambda s.s(\underline{1}))\bullet\}) \\ \Rightarrow & (U(U\underline{1}), \mathcal{E} \uplus \{a \mapsto (\lambda s.s(\underline{1}))\bullet\}) \\ \Rightarrow & (U(\mu c.[a](\lambda x.\text{ifz } \underline{1} \text{ then } \Omega \text{ else } \underline{0})), \mathcal{E} \uplus \{a \mapsto (\lambda s.s(\underline{1}))\bullet\}) \\ \Rightarrow^2 & ((\lambda s.s(\underline{1}))(\lambda x.\text{ifz } \underline{1} \text{ then } \Omega \text{ else } \underline{0}), \mathcal{E} \uplus \{a \mapsto (\lambda s.s(\underline{1}))\bullet, c \mapsto U\bullet\}) \\ \Rightarrow^3 & (\underline{0}, \mathcal{E} \uplus \{a \mapsto (\lambda s.s(\underline{1}))\bullet, c \mapsto U\bullet\}) \end{aligned}$$

(U is used as shorthand for the term $\lambda y.\mu c.[a](\lambda x.\text{ifz } y \text{ then } \Omega \text{ else } \underline{0})$.) However $\mathcal{C}[T_2]$ loops in the following way.

$$\begin{aligned} & (\mathcal{C}[T_2], \mathcal{E}) \\ \stackrel{\text{def}}{=} & ((\lambda s.s(\underline{1}))T_2, \mathcal{E}) \\ \Rightarrow & (T_2(T_2\underline{1}), \mathcal{E}) \\ \Rightarrow & (T_2(\mu b.[b](\lambda y.\mu c.[b](\lambda x.\text{ifz } y \text{ then } \Omega \text{ else } \underline{0}\underline{1}\underline{1}))), \mathcal{E}) \\ \Rightarrow^2 & (T_2((\lambda y.\mu c.[b](\lambda x.\text{ifz } y \text{ then } \Omega \text{ else } \underline{0}\underline{1}\underline{1}))), \mathcal{E} \uplus \{b \mapsto T_2\bullet\}) \\ \Rightarrow & (T_2(\mu c.[b](\lambda x.\text{ifz } \underline{1} \text{ then } \Omega \text{ else } \underline{0}\underline{1}\underline{1})), \mathcal{E} \uplus \{b \mapsto T_2\bullet\}) \\ \Rightarrow^2 & (T_2((\lambda x.\text{ifz } \underline{1} \text{ then } \Omega \text{ else } \underline{0}\underline{1}\underline{1})), \mathcal{E} \uplus \{b \mapsto T_2\bullet, c \mapsto T_2\bullet\}) \\ \Rightarrow & (T_2\underline{0}, \mathcal{E} \uplus \{b \mapsto T_2\bullet, c \mapsto T_2\bullet\}) \\ \Rightarrow & (\mu b'.[b']((\lambda y.\mu c'.[b'](\lambda x.\text{ifz } y \text{ then } \Omega \text{ else } \underline{0}\underline{0}\underline{0}))), \mathcal{E} \uplus \{b \mapsto T_2\bullet, c \mapsto T_2\bullet\}) \\ \Rightarrow^2 & ((\lambda y.\mu c'.[b'](\lambda x.\text{ifz } y \text{ then } \Omega \text{ else } \underline{0}\underline{0}\underline{0}))), \mathcal{E} \uplus \{b \mapsto T_2\bullet, c \mapsto T_2\bullet, b' \mapsto \bullet\}) \\ \Rightarrow & (\mu c'.[b'](\lambda x.\text{ifz } \underline{0} \text{ then } \Omega \text{ else } \underline{0}\underline{0}\underline{0})), \mathcal{E} \uplus \{b \mapsto T_2\bullet, c \mapsto T_2\bullet, b' \mapsto \bullet\}) \\ \Rightarrow^2 & ((\lambda x.\text{ifz } \underline{0} \text{ then } \Omega \text{ else } \underline{0}\underline{0}\underline{0})), \mathcal{E} \uplus \{b \mapsto T_2\bullet, c \mapsto T_2\bullet, b' \mapsto \bullet, c' \mapsto \bullet\}) \\ \Rightarrow & (\text{ifz } \underline{0} \text{ then } \Omega \text{ else } \underline{0}, \mathcal{E} \uplus \{b \mapsto T_2\bullet, c \mapsto T_2\bullet, b' \mapsto \bullet, c' \mapsto \bullet\}) \\ \uparrow & \end{aligned}$$

A consequence of this example is the fact that simple definitions of applicative bisimilarity are unlikely to coincide with contextual equivalence. (Indeed as Ong and Stewart point out, the most obvious definition of applicative bisimilarity fails even to be a congruence!)

The problem with contextual equivalence is that it is very hard to reason about. As demonstrated above, it is easy to see when two programs are *not* contextually equivalent—one simply finds a context which distinguishes them. To show that two programs *are* contextually equivalent is much harder: the problem lies in the quantification over all program contexts. Indeed, the discussion above demonstrates how μ PCF's save and restore features also complicate the problem.

Despite these difficulties let us consider a very simple notion of program equivalence which uses the notion of termination given at the end of §6. Informally, the idea is that two programs are considered equivalent if they have the same termination properties. The formal definition is as follows.

Definition 2 *Given two programs M and N , M is said to ciu-refine N , written $M \leq N: \phi, \Sigma$, just when $\forall S, \mathcal{E}$. if $(S, M, \mathcal{E}) \searrow$ then $(S, N, \mathcal{E}) \searrow$. They are said to be ciu-equivalent, written $M \simeq N: \phi, \Sigma$ just when $M \leq N: \phi, \Sigma$ and $N \leq M: \phi, \Sigma$.*

These can be extended to open terms as follows

$$\begin{aligned}\vec{x}: \Gamma \triangleright M \leq^\circ N: \phi, \Sigma &\stackrel{\text{def}}{=} \forall \vec{v}. M[\vec{x} := \vec{v}] \leq N[\vec{x} := \vec{v}]: \phi, \Sigma \\ \vec{x}: \Gamma \triangleright M \simeq^\circ N: \phi, \Sigma &\stackrel{\text{def}}{=} \forall \vec{v}. M[\vec{x} := \vec{v}] \simeq N[\vec{x} := \vec{v}]: \phi, \Sigma\end{aligned}$$

The obvious question is in what way are ciu-equivalence and contextual equivalence related? It is possible to show that they coincide, or, put in another way, ciu-equivalence is an alternative (but relatively simpler) characterisation of contextual equivalence.

Theorem 1 (ciu theorem) $\Gamma \triangleright M \approx N: \phi, \Sigma$ iff $\Gamma \triangleright M \simeq^\circ N: \phi, \Sigma$.

Proof. Some details are given in Appendix A. ■

This means that to prove two terms contextually equivalent we need only to show that they are ciu-equivalent, which is often significantly easier. For example, it can be shown that the (call-by-value variants of the) reduction rules of the $\lambda\mu$ -calculus given in §2 are contextual equivalences, by showing that they are, in fact, ciu-equivalences.

$$(\lambda x: \phi. M)v \approx M[x := v]: \psi, \Sigma \tag{1}$$

$$\mu a: \phi. [a: \phi]M \approx M: \phi, \Sigma \tag{2} \quad (a \notin \mu\text{FV}(M))$$

$$[a: \phi]\mu b: \phi. [c: \psi]M \approx ([c: \psi]M)[a/b]: \perp, a: \phi, \Sigma \tag{3}$$

$$(\mu a: \phi \rightarrow \psi. M)N \approx \mu b: \psi. M[a \leftarrow [b: \psi] \bullet N]: \psi, \Sigma \tag{4}$$

For example, the equivalence (2) holds by observing

$$\frac{\frac{(S, M, \mathcal{E} \uplus \{a \mapsto S\}) \searrow}{(id, [a]M, \mathcal{E} \uplus \{a \mapsto S\}) \searrow}}{(S, \mu a. [a]M, \mathcal{E}) \searrow}$$

and the fact that $a \notin \mu\text{FV}(M)$ by assumption.

A number of other equivalences can be derived. For example, consider a term M such that $\Gamma, x: \phi \triangleright M: \psi, \Sigma$. An instance of equivalence (1) is

$$\Gamma, x: \phi \triangleright (\lambda x: \phi. M)x \simeq^\circ M: \psi, \Sigma \tag{5}$$

In Appendix A it is shown that ciu-equivalence is a congruence. In particular, this means that the following rule is valid.

$$\frac{\Gamma, x: \phi \triangleright M \simeq^\circ N: \psi, \Sigma}{\Gamma \triangleright \lambda x: \phi. M \simeq^\circ \lambda x: \phi. N: \phi \rightarrow \psi, \Sigma}$$

Applying this to equivalence 5 gives

$$\Gamma \triangleright \lambda x: \phi. (\lambda x: \phi. M)x \simeq^\circ \lambda x: \phi. M: \phi \rightarrow \psi, \Sigma.$$

Setting v for the term $\lambda x: \phi. M$, this implies the following contextual equivalence.

$$\Gamma \triangleright \lambda x: \phi. vx \approx v: \phi \rightarrow \psi, \Sigma \tag{6}$$

(where $x \notin \text{FV}(v)$). Hence the call-by-value η -rule is a contextual equivalence. Here are a couple more examples of contextual equivalences.

$$E[\mu a.[b]M] \approx [b]M: \perp \quad (a \notin \mu\text{FV}(M)) \quad (7)$$

$$\mu a.[b](\text{ifz } M \text{ then } N \text{ else } P) \approx \text{ifz } M \text{ then } \mu a.[b]N \text{ else } \mu a.[b]P: \phi \quad (a \notin \mu\text{FV}(M)) \quad (8)$$

The first equivalence is a sort of ‘abort’ axiom, the second is a typical compiler optimisation. Further study of provable contextual equivalences is important future work.

8 Implementation

In this section I shall give details of two implementations. The first is simply the SML code for a tail recursive interpreter based on the transition system of §6. The second is a simple abstract machine, which is based on the framework proposed by Curien (1991).

8.1 Interpreter

As mentioned in §6, the transition system essentially describes a tail recursive interpreter for μPCF . In this section I shall give details for such an interpreter, written in SML. Stacks are implemented as lists of frames. The function mapping μ -variables to stacks is implemented as a list of pairs. The main datatypes are as follows.

```
datatype 'a exp = const    of int
                | suc      of 'a exp
                | var      of 'a
                | abs      of 'a * ('a exp)
                | app      of ('a exp) * ('a exp)
                | pair     of ('a exp) * ('a exp)
                | fst      of ('a exp)
                | snd      of ('a exp)
                | ifz      of ('a exp) * ('a exp) * ('a exp)
                | letrec   of 'a * 'a * ('a exp) * ('a exp)
                | save     of 'a * ('a exp)
                | restore  of 'a * ('a exp);

type term = string exp;
(* type typed_term = (string*formula) exp *)

datatype hole = Bullet;

datatype frame = csuc    of hole
                | appl   of hole*term
                | appr   of term*hole
                | pairl  of hole*term
                | pairr  of term*hole
                | cfst   of hole
                | csnd   of hole
                | cifz   of hole*term*term;

type stack = frame list;
```

```

type stack_function = (string*stack) list;
(* type typed_stack_function = ((string*formula)*stack) list *)

```

The main interpreter consists of two mutually recursive routines `eval` and `unwind`. They call three subsidiary functions: `insert` and `get`, which handle the function update and access respectively, and `subs` which performs substitution of a term for a variable. The code for the interpreter is then as follows.

```

fun eval S (const i)          C = unwind S (const i) C
  | eval S (suc(e))           C = eval (csuc(Bullet)::S) e C
  | eval S (abs(x,e))        C = unwind S (abs(x,e)) C
  | eval S (app(e,f))        C = eval (appl(Bullet,f)::S) e C
  | eval S (fst(e))          C = eval (cfst(Bullet)::S) e C
  | eval S (snd(e))          C = eval (csnd(Bullet)::S) e C
  | eval S (pair(e,f))       C = eval (pairl(Bullet,f)::S) e C
  | eval S (ifz(e,f,g))      C = eval (cifz(Bullet,f,g)::S) e C
  | eval S (letrec(f,x,e,g)) C = unwind
                                S
                                subs(abs(x,letrec(f,x,e,e)),f,g)
                                C
  | eval S (restore(n,e))    C = let val T=get n C
                                in eval T e C
                                end
  | eval S (save(n,e))       C = eval [] e (insert (n,S) C)
and
unwind S v C = case S of
  [] => v
| (csuc(Bullet)::S) => let val (const i) = v
                      in unwind S (const (i+1)) C
                      end
| (appl(Bullet,e)::S) => eval (appr(v,Bullet)::S) e C
| (appr(abs(x,e),Bullet)::S) => eval S (subs(v,x,e)) C
| (pairl(Bullet,e)::S) => eval (pairr(v,Bullet)::S) e C
| (pairr(w,Bullet)::S) => unwind S (pair(w,v)) C
| (cfst(Bullet)::S) => let val pair(u,w)=v
                      in unwind S u C
                      end
| (csnd(Bullet)::S) => let val pair(u,w)=v
                      in unwind S w C
                      end
| (cifz(Bullet,f,g)::S) => let val (const i) = v
                      in if (i=0) then eval S f C
                          else eval S g C
                      end;
end;

```

8.2 Towards an Abstract Machine

In his paper, Curien (1991) describes a simple framework for abstract machines which evaluate functional programs using environments and closures. Both the CAM and Krivine's machine

arise naturally as instances of this abstract framework. In this section, I shall extend Curien’s “Eager Machine” to μ PCF (actually without recursion).

The machine essentially consists of a triple. The first part is an *environment*. For simplicity I shall consider this to be a pair of symbol tables (functions), the first maps λ -variables to values and the second mapping μ -variables to stacks. In a realistic machine this indirection would be removed by compiling both λ - and μ -variables into de Bruijn indices. I have chosen not to do this for clarity. The second part of the triple is the term. The third part is a stack of values (I use $::$ as an infix ‘push’ operator). The evaluation rules for the abstract machine are as follows.

$$\begin{array}{ll}
\langle \Gamma, \mathcal{E} \rangle, \underline{n}, S & \gg (-, -, \underline{n} :: S) \\
\langle \Gamma, \mathcal{E} \rangle, \text{succ}(M), S & \gg \langle \Gamma, \mathcal{E} \rangle, M, \mathbf{SUC} :: S \\
(-, -, \underline{n} :: \mathbf{SUC} :: S) & \gg (-, -, \underline{n+1} :: S) \\
\langle \Gamma \uplus \{x \mapsto v\}, \mathcal{E} \rangle, x, S & \gg (-, -, v :: S) \\
\langle \Gamma, \mathcal{E} \rangle, \lambda x.M, S & \gg (-, -, \mathbf{cl}(\langle \Gamma, \mathcal{E} \rangle, \lambda x.M) :: S) \\
\langle \Gamma, \mathcal{E} \rangle, MN, S & \gg \langle \Gamma, \mathcal{E} \rangle, M, \mathbf{L} :: \mathbf{cl}(\langle \Gamma, \mathcal{E} \rangle, N) :: S \\
(-, -, v :: \mathbf{L} :: \mathbf{cl}(\langle \Gamma, \mathcal{E} \rangle, M) :: S) & \gg \langle \Gamma, \mathcal{E} \rangle, M, \mathbf{R} :: v :: S \\
(-, -, v :: \mathbf{R} :: \mathbf{cl}(\langle \Gamma, \mathcal{E} \rangle, \lambda x.M) :: S) & \gg \langle \Gamma \uplus \{x \mapsto v\}, \mathcal{E} \rangle, M, S \\
\langle \Gamma, \mathcal{E} \rangle, \mu a.M, S & \gg \langle \Gamma, \mathcal{E} \uplus \{a \mapsto S\} \rangle, M, [] \\
\langle \Gamma, \mathcal{E} \uplus \{a \mapsto S\} \rangle, [a]M, T & \gg \langle \Gamma, \mathcal{E} \uplus \{a \mapsto S\} \rangle, M, S \\
\langle \Gamma, \mathcal{E} \rangle, \text{ifz } M \text{ then } N \text{ else } P, S & \gg \langle \Gamma, \mathcal{E} \rangle, M, \mathbf{IFZ} :: \mathbf{cl}(\langle \Gamma, \mathcal{E} \rangle, \langle N, P \rangle) :: S \\
(-, -, \underline{0} :: \mathbf{IFZ} :: \mathbf{cl}(\langle \Gamma, \mathcal{E} \rangle, \langle N, P \rangle) :: S) & \gg \langle \Gamma, \mathcal{E} \rangle, N, S \\
(-, -, \underline{n+1} :: \mathbf{IFZ} :: \mathbf{cl}(\langle \Gamma, \mathcal{E} \rangle, \langle N, P \rangle) :: S) & \gg \langle \Gamma, \mathcal{E} \rangle, P, S
\end{array}$$

As the environment is separate from the term, the value of a function is naturally a closure, which has a code part and an environment part. Closures are written $\mathbf{cl}(\langle \Gamma, \mathcal{E} \rangle, M)$.

The feature of this machine (and related environment machines) is that the recursive calls are implemented using the stack and the “markers” (\mathbf{L} , \mathbf{R} , \mathbf{SUC} and \mathbf{IFZ}).⁵ Consequently the stack contains either natural numbers, closures or markers.

Further details of such environment machines and more generally about calculi of closures can be found in Curien’s paper (Curien 1991). Low level details of implementing other related control operators are given by Hieb, Dybvig and Bruggeman (1990).

9 Call-by-Name

This paper has so far considered only call-by-value computation. However it is very simple to provide a computational interpretation for a call-by-name evaluation strategy. The main difference is in the (new) definition of values, evaluation contexts and redexes, which are as follows. (As normal the call-by-value function restriction on recursion can be relaxed (Winskel 1993, §11.5).)

⁵The \mathbf{L} and \mathbf{R} markers are not required in a call-by-name setting. The resulting machine in that case is essentially a Krivine machine—a related machine has been given by Selinger (1998b).

Values	$v ::= \underline{n} \mid \lambda x.M \mid \langle M, M \rangle$
Evaluation Contexts	$E ::= \bullet$ $\mid EM$ $\mid \text{fst}(E) \mid \text{snd}(E)$ $\mid \text{suc}(E)$ $\mid \text{ifz } E \text{ then } M \text{ else } M$
Redexes	$R ::= vM$ $\mid \text{fst}(v) \mid \text{snd}(v)$ $\mid \text{suc}(v)$ $\mid \text{ifz } v \text{ then } M \text{ else } M$ $\mid \text{rec } x.M$ $\mid [a]M \mid \mu a.M$

The evaluation rules are essentially as before.

$$\begin{aligned}
(E[(\lambda x.M)N], \mathcal{E}) &\Rightarrow (E[M[x := N]], \mathcal{E}) \\
(E[\text{fst}(\langle M, N \rangle)], \mathcal{E}) &\Rightarrow (E[M], \mathcal{E}) \\
(E[\text{snd}(\langle M, N \rangle)], \mathcal{E}) &\Rightarrow (E[N], \mathcal{E}) \\
(E[\text{suc}(\underline{n})], \mathcal{E}) &\Rightarrow (E[\underline{n+1}], \mathcal{E}) \\
(E[\text{ifz } \underline{0} \text{ then } M \text{ else } N], \mathcal{E}) &\Rightarrow (E[M], \mathcal{E}) \\
(E[\text{ifz } \underline{n+1} \text{ then } M \text{ else } N], \mathcal{E}) &\Rightarrow (E[N], \mathcal{E}) \\
(E[\text{rec } x.M], \mathcal{E}) &\Rightarrow (E[M[x := (\text{rec } x.M)]], \mathcal{E}) \\
(E[\mu a.M], \mathcal{E}) &\Rightarrow (M, \mathcal{E} \uplus \{a \mapsto E[\bullet]\}) \\
(E[[a]M], \mathcal{E} \uplus \{a \mapsto E'[\bullet]\}) &\Rightarrow (E'[M], \mathcal{E} \uplus \{a \mapsto E'[\bullet]\})
\end{aligned}$$

The development of the corresponding operational theory follows closely that outlined in §7. The differs from the treatment given by Ong and Stewart (1997) who have to introduce completely new reduction rules to move from a call-by-name to a call-by-value setting.

10 Conclusion

In this paper I have given a simple computation interpretation of the $\lambda\mu$ -calculus: it is a λ -calculus which is extended with indexed operators to save and restore the runtime environment. This is maybe not too surprising as Griffin (1990) has shown the close relationship between classical logic and other languages with control. This interpretation can be expressed as a single-step reduction semantics using environment contexts. In turn I gave an equivalent semantics expressed as steps of a simple transition system, which eliminated the need for the evaluation contexts. Using this simple transition system it is possible to define a notion of program equivalence based on a termination relation which can be proved to be equivalent to a natural definition of contextual equivalence.

Clearly the work by Ong and Stewart (1997) is most closely related to that reported here. Their thesis is that μ PCF is a foundational language for call-by-value functional computation with control and this paper can be seen as further evidence to that claim. However I would suggest that the operational treatment given here is more intuitive, more flexible (in that different calling mechanisms can be handled easily) and leads to a more refined notion of program equivalence.

This paper has given a computational interpretation directly to the $\lambda\mu$ -calculus but de Groote (1994) has worked in the other direction, using existing work on continuation passing to give an interpretation for the (call-by-name) $\lambda\mu$ -calculus. A fuller comparison of these approaches is important future work.

Another interesting development would be to relate the operational notions of program equivalence from this paper, to various denotational models. It should certainly be possible to show that ciu-equivalence is adequately modelled in the (appropriate extensions of) models of Hofmann and Streicher (1997) and Selinger (1998a). Recently Abramsky and McCusker (1998) have extended the fully abstract game models of PCF to handle various imperative and control-like features. It would be interesting to see if their techniques can be used to define a fully abstract (game) model of μ PCF.

After completing the first draft of this paper, the work of Talcott (1998) was brought to my attention. Talcott also proves a ciu-theorem for an untyped variant of Idealised Scheme (Felleisen and Friedman 1986) as well as for a language with explicit memory effects. Interestingly she does not use a variant of Howe's method, but rather uses a notion of "uniform computation". The reader is invited to compare this technique with that outlined in Appendix A.

The techniques used in this paper to define ciu-equivalence and prove a ciu-theorem can be applied to a number of complicated languages. They have been used by Pitts in recent unpublished work on various calculi with explicit state. In other unpublished work, the author and A.M. Pitts have applied these techniques to both Idealised Scheme and the generalised control language of Gunter et al. (1995).

Acknowledgements

I am currently supported by EPSRC Grant GR/M04716 and Gonville and Caius College, Cambridge. I am grateful to Nick Benton, Søren Lassen, Luke Ong, Andrew Pitts and Peter Selinger for helpful comments and discussions.

A preliminary version of this paper was presented at MFCS'98 (Bierman 1998).

References

- Abramsky, S. and McCusker, G. (1998). Game semantics, Lecture notes from 1997 Marktoberdorf summer school.
- Abramsky, S. (1990). The lazy lambda calculus, in D. Turner (ed.), *Research Topics in Functional Programming*, Addison-Wesley, chapter 4, pp. 65–116.
- Bierman, G.M. (1998). A computational interpretation of the $\lambda\mu$ -calculus, in L. Brim, J. Gruska and J. Zlatuška (eds), *Proceedings of Symposium on Mathematical Foundations of Computer Science*, Vol. 1450 of *Lecture Notes in Computer Science*, pp. 336–345.
- Cartwright, R., Curien, P.-L. and Felleisen, M. (1994). Fully abstract semantics for observably sequential languages, *Information and Computation* **111**(2): 297–401.
- Clinger, W. and Rees, J. (1986). The revised³ report on the algorithmic language Scheme, *ACM SIGPLAN Notices* **21**(12): 37–79.

- Curien, P.-L. (1991). An abstract framework for environment machines, *Theoretical Computer Science* **82**(2): 389–402.
- de Groote, P. (1994). On the relation between the $\lambda\mu$ -calculus and the syntactic theory of sequential control, *Proceedings of Conference on Logic Programming and Automated Reasoning*, Vol. 822 of *Lecture Notes in Computer Science*, pp. 31–43.
- de Groote, P. (1995). A simple calculus of exception handling, *Proceedings of Second International Conference on Typed λ -calculi and applications*, Vol. 902 of *Lecture Notes in Computer Science*, pp. 201–215.
- Felleisen, M. and Friedman, D. (1986). Control operators, the SECD-machine and the λ -calculus, *Formal Description of Programming Concepts III*, North-Holland, pp. 131–141.
- Felleisen, M. and Hieb, R. (1992). The revised report on the syntactic theories of sequential control and state, *Theoretical Computer Science* **103**(2): 235–271.
- Gentzen, G. (1969). Investigations into logical deduction, in M. Szabo (ed.), *The Collected Papers of Gerhard Gentzen*, North-Holland, pp. 68–131. English Translation of 1935 German original.
- Gordon, A.D. (1995). Bisimilarity as a theory of functional programming: Mini-course, *Technical Report NS-95-2*, BRICS, Department of Computer Science, University of Århus.
- Gordon, A.D. and Pitts, A.M. (eds) (1998). *Higher Order Operational Techniques in Semantics*, Publications of the Newton Institute, Cambridge University Press.
- Griffin, T. (1990). A formulae-as-types notion of control, *Proceedings of Symposium on Principles of Programming Languages*, pp. 47–58.
- Gunter, C., Rémy, D. and Riecke, J. (1995). A generalisation of exceptions and control in ML-like languages, *Proceedings of Conference on Functional Programming Languages and Computer Architecture*, pp. 12–23.
- Harper, R. and Stone, C. (1997). An interpretation of Standard ML in type theory, *Technical Report CMU-CS-97-147*, School of Computer Science, Carnegie Mellon University.
- Hieb, R., Dybvig, R. and Bruggeman, C. (1990). Representing control in the presence of first-class continuations, *Proceedings of the Conference on Programming Language Design and Implementation*, pp. 66–77.
- Hofmann, M. (1995). Sound and complete axiomatisations of call-by-value control operators, *Mathematical Structures in Computer Science* **5**: 461–482.
- Hofmann, M. and Streicher, T. (1997). Continuation models are universal for $\lambda\mu$ -calculus, *Proceedings of Symposium on Logic in Computer Science*, pp. 387–397.
- Howe, D. (1989). Equality in lazy computation systems, *Proceedings of Symposium on Logic in Computer Science*, pp. 198–203.
- Lillibridge, M. (1995). Exceptions are strictly more powerful than call/cc, *Technical Report CMU-CS-95-178*, School of Computer Science, Carnegie Mellon University.

- Meyer, A. and Riecke, J. (1988). Continuations may be unreasonable (Preliminary Report), *Proceedings of the 1988 ACM Conference on Lisp and Functional Programming*, pp. 63–71.
- Ong, C.-H.L. (1996). A semantic view of classical proofs: type-theoretic, categorical and denotational characterizations, *Proceedings of Symposium on Logic in Computer Science*, pp. 230–241.
- Ong, C.-H.L. and Stewart, C. (1997). A Curry-Howard foundation for functional computation with control, *Proceedings of Symposium on Principles of Programming Languages*, pp. 215–227.
- Parigot, M. (1992). $\lambda\mu$ -calculus: an algorithmic interpretation of classical natural deduction, *Proceedings of Conference on Logic Programming and Automated Reasoning*, Vol. 624 of *Lecture Notes in Computer Science*, pp. 190–201.
- Parigot, M. (1997). Proofs of strong normalisation for second order classical natural deduction, *Journal of Symbolic Logic* **62**(4): 1461–1479.
- Pitts, A.M. (1997a). Operational semantics for program equivalence, Slides from talk given at MFPS.
- Pitts, A.M. (1997b). Operationally-based theories of program equivalence, in P. Dybjer and A.M. Pitts (eds), *Semantics and Logics of Computation*, Publications of the Newton Institute, Cambridge University Press, pp. 241–298.
- Pitts, A.M. (1998). Parametric polymorphism and operational equivalence (preliminary version), Vol. 10 of *Electronic Notes in Theoretical Computer Science*.
- Pitts, A.M. and Stark, I.D.B. (1998). Operational reasoning for functions with local state, in A.D. Gordon and A.M. Pitts (eds), *Higher Order Operational Techniques in Semantics*, Publications of the Newton Institute, Cambridge University Press, pp. 227–273.
- Selinger, P. (1998a). Control categories: an axiomatic approach to the semantics of control in functional languages, Unpublished manuscript.
- Selinger, P. (1998b). An implementation of the call-by-name $\lambda\mu\nu$ -calculus, Unpublished manuscript.
- Talcott, C. (1998). Reasoning about functions with effects, in A.D. Gordon and A.M. Pitts (eds), *Higher Order Operational Techniques in Semantics*, Publications of the Newton Institute, Cambridge University Press, pp. 347–390.
- Winskel, G. (1993). *The Formal Semantics of Programming Languages: An Introduction*, MIT Press.

A The ciu-theorem

In this appendix, I shall give details of the proof of the ‘ciu-theorem’: the proof that contextual equivalence coincides with ciu-equivalence. First recall these two definitions of program equivalence.

Definition 3 *Let M and N be terms and \mathcal{C} a λ -closing context. M is said to contextually refine N , written $\Gamma \triangleright M \sqsubseteq N: \phi, \Sigma$, when $\forall \mathcal{C}, \mathcal{E}$. if $(\mathcal{C}[M], \mathcal{E}) \Downarrow$ then $(\mathcal{C}[N], \mathcal{E}) \Downarrow$. They are said to be contextually equivalent, written $\Gamma \triangleright M \approx N: \phi, \Sigma$, just when $\Gamma \triangleright M \sqsubseteq N: \phi, \Sigma$ iff $\Gamma \triangleright N \sqsubseteq M: \phi, \Sigma$.*

Definition 4 *Given two programs M and N , M is said to ciu-refine N , written $M \leq N: \phi, \Sigma$, just when $\forall S, \mathcal{E}$. if $(S, M, \mathcal{E}) \searrow$ then $(S, N, \mathcal{E}) \searrow$. They are said to be ciu-equivalent, written $M \simeq N: \phi, \Sigma$ just when $M \leq N: \phi, \Sigma$ and $N \leq M: \phi, \Sigma$.*

These can be extended to open terms as follows

$$\begin{aligned} \vec{x}: \Gamma \triangleright M \leq^\circ N: \phi, \Sigma &\stackrel{\text{def}}{=} \forall \vec{v}. M[\vec{x} := \vec{v}] \leq N[\vec{x} := \vec{v}]: \phi, \Sigma \\ \vec{x}: \Gamma \triangleright M \simeq^\circ N: \phi, \Sigma &\stackrel{\text{def}}{=} \forall \vec{v}. M[\vec{x} := \vec{v}] \simeq N[\vec{x} := \vec{v}]: \phi, \Sigma \end{aligned}$$

Two facts are almost immediate from the definition of ciu-refinement.

Lemma 2 1. $\forall M. M \leq M: \phi, \Sigma$.

2. If $M \leq M': \phi, \Sigma$ and $M' \leq M'': \phi, \Sigma$ then $M \leq M'': \phi, \Sigma$.

The following properties of relations between terms will be useful.

Definition 5 *A relation, \mathcal{R} , is said to be compatible if it satisfies the following rules.*

$$\begin{array}{c} \frac{}{\Gamma, x: \phi \triangleright x \mathcal{R} x: \phi, \Sigma} \qquad \frac{}{\Gamma \triangleright \underline{n} \mathcal{R} \underline{n}: \text{nat}, \Sigma} \\ \frac{\Gamma \triangleright M \mathcal{R} N: \phi, \Sigma}{\Gamma, x: \psi \triangleright M \mathcal{R} N: \phi, \Sigma} \text{Weakening}_{\mathcal{L}} \qquad \frac{\Gamma \triangleright M \mathcal{R} N: \phi, \Sigma}{\Gamma \triangleright M \mathcal{R} N: \phi, \Sigma, a: \psi} \text{Weakening}_{\mathcal{R}} \\ \frac{\Gamma, x: \phi \triangleright M \mathcal{R} M': \psi, \Sigma}{\Gamma \triangleright \lambda x. M \mathcal{R} \lambda x. M': \phi \rightarrow \psi, \Sigma} \rightarrow_{\mathcal{I}} \\ \frac{\Gamma \triangleright M \mathcal{R} N: \phi \rightarrow \psi, \Sigma \quad \Gamma \triangleright M' \mathcal{R} N': \phi, \Sigma}{\Gamma \triangleright MM' \mathcal{R} NN': \psi, \Sigma} \rightarrow_{\mathcal{E}} \\ \frac{\Gamma \triangleright M \mathcal{R} M': \phi, \Sigma \quad \Gamma \triangleright N \mathcal{R} N': \psi, \Sigma}{\Gamma \triangleright \langle M, N \rangle \mathcal{R} \langle M', N' \rangle: \phi \times \psi, \Sigma} \times_{\mathcal{I}} \\ \frac{\Gamma \triangleright M \mathcal{R} M': \phi \times \psi, \Sigma}{\Gamma \triangleright \text{fst}(M) \mathcal{R} \text{fst}(M'): \phi, \Sigma} \times_{\mathcal{E}} \qquad \frac{\Gamma \triangleright M \mathcal{R} M': \phi \times \psi, \Sigma}{\Gamma \triangleright \text{snd}(M) \mathcal{R} \text{snd}(M'): \psi, \Sigma} \times_{\mathcal{E}} \end{array}$$

$$\begin{array}{l}
\Gamma \triangleright \langle M, M' \rangle \leq^* N: \phi \times \psi, \Sigma \stackrel{\text{def}}{=} \exists P, P'. \Gamma \triangleright M \leq^* P: \phi, \Sigma \\
\Gamma \triangleright M' \leq^* P': \psi, \Sigma \\
\Gamma \triangleright \langle P, P' \rangle \leq^\circ N: \phi \times \psi, \Sigma \\
\Gamma \triangleright \text{fst}(M) \leq^* N: \phi, \Sigma \stackrel{\text{def}}{=} \exists P. \Gamma \triangleright M \leq^* P: \phi \times \psi, \Sigma \\
\Gamma \triangleright \text{fst}(P) \leq^\circ N: \phi, \Sigma \\
\Gamma \triangleright \text{snd}(M) \leq^* N: \psi, \Sigma \stackrel{\text{def}}{=} \exists P. \Gamma \triangleright M \leq^* P: \phi \times \psi, \Sigma \\
\Gamma \triangleright \text{snd}(P) \leq^\circ N: \psi, \Sigma \\
\Gamma \triangleright \text{letrec } f = \lambda x. M \text{ in } M' \leq^* N: \psi, \Sigma \stackrel{\text{def}}{=} \exists P, P'. \Gamma, f: \phi \rightarrow \phi, x: \phi \triangleright M \leq^* P: \phi, \Sigma \\
\Gamma, f: \phi \rightarrow \phi \triangleright M' \leq^* P': \psi, \Sigma \\
\Gamma \triangleright \text{letrec } f = \lambda x. P \text{ in } P' \leq^\circ N: \psi, \Sigma \\
\Gamma \triangleright \mu a: \phi. M \leq^* N: \phi, \Sigma \stackrel{\text{def}}{=} \exists P. \Gamma \triangleright M \leq^* P: \perp, \Sigma, a: \phi \\
\Gamma \triangleright \mu a: \phi. P \leq^\circ N: \phi, \Sigma \\
\Gamma \triangleright [a: \phi]M \leq^* N: \perp, a: \phi, \Sigma \stackrel{\text{def}}{=} \exists P. \Gamma \triangleright M \leq^* P: \phi, \Sigma \\
\Gamma \triangleright [a: \phi]P \leq^\circ N: \perp, a: \phi, \Sigma
\end{array}$$

This relation can be extended to frames, stacks of frames and functions from μ -variables to stacks. This is important for technical reasons. The definitions are as follows.

Definition 7 1.

$$\begin{array}{l}
\bullet M \leq^* F[\bullet]: \psi, \Sigma \stackrel{\text{def}}{=} \exists N. M \leq^* N: \phi, \Sigma \\
\forall P. PN \leq F[P]: \psi, \Sigma \\
v \bullet \leq^* F[\bullet]: \psi, \Sigma \stackrel{\text{def}}{=} \exists N. v \leq^* N: \phi \rightarrow \psi, \Sigma \\
\forall P. NP \leq F[P]: \psi, \Sigma \\
\langle \bullet, M \rangle \leq^* F[\bullet]: \phi \times \psi, \Sigma \stackrel{\text{def}}{=} \exists N. M \leq^* N: \psi, \Sigma \\
\forall P. \langle P, N \rangle \leq F[P]: \phi \times \psi, \Sigma \\
\langle v, \bullet \rangle \leq^* F[\bullet]: \phi \times \psi, \Sigma \stackrel{\text{def}}{=} \exists N. v \leq^* N: \phi, \Sigma \\
\forall P. \langle v, P \rangle \leq F[P]: \phi \times \psi, \Sigma \\
\text{fst}(\bullet) \leq^* F[\bullet]: \psi, \Sigma \stackrel{\text{def}}{=} \forall P. \text{fst}(P) \leq F[P]: \psi, \Sigma \\
\text{snd}(\bullet) \leq^* F[\bullet]: \psi, \Sigma \stackrel{\text{def}}{=} \forall P. \text{snd}(P) \leq F[P]: \psi, \Sigma \\
\text{suc}(\bullet) \leq^* F[\bullet]: \text{nat}, \Sigma \stackrel{\text{def}}{=} \forall P. \text{suc}(P) \leq F[P]: \text{nat}, \Sigma \\
\text{ifz } \bullet \text{ then } M \text{ else } M' \leq^* F[\bullet]: \psi, \Sigma \stackrel{\text{def}}{=} \exists N, N'. M \leq^* N: \psi, \Sigma \\
M' \leq^* N': \psi, \Sigma \\
\forall P. \text{ifz } P \text{ then } N \text{ else } N' \leq F[P]: \psi, \Sigma
\end{array}$$

2.

$$\text{id} \leq^* \text{id} \quad \frac{S \leq^* S' \quad F[\bullet] \leq^* F'[\bullet]}{(S \circ F[\bullet]) \leq^* (S' \circ F'[\bullet])}$$

3.

$$\mathcal{E} \leq^* \mathcal{E}' \text{ iff } \forall a. \mathcal{E}a \leq^* \mathcal{E}'a.$$

It is fairly easy to verify that all these relations are reflexive.

Lemma 4 (Reflexivity) 1. $\forall M. \Gamma \triangleright M \leq^* M: \phi, \Sigma.$

2. $\forall F[\bullet]. F[\bullet] \leq^* F[\bullet]: \phi, \Sigma.$

3. $\forall S. S \leq^* S.$

4. $\forall \mathcal{E}. \mathcal{E} \leq^* \mathcal{E}$.

It is hard to see a direct proof that the \leq^* relation is transitive. However the following property will suffice.

Lemma 5 *If $M \leq^* M': \phi, \Sigma$ and $M' \leq M'': \phi, \Sigma$ then $M \leq^* M'': \phi, \Sigma$.*

Proof. By induction on $M \leq^* M': \phi, \Sigma$. ■

One half of the coincidence of the \leq^* relation and ciu-refinement is now immediate.

Proposition 2 *If $M \leq N: \phi, \Sigma$ then $M \leq^* N: \phi, \Sigma$.*

Proof. We know that $M \leq^* M: \phi, \Sigma$ and have that $M \leq N: \phi, \Sigma$. Hence by lemma 5 we conclude that $M \leq^* N: \phi, \Sigma$. ■

The following two properties will be useful.

Lemma 6 1. *If $v \leq^* v': \phi, \Sigma$ and $\Gamma, x: \phi \triangleright M \leq^* M': \psi, \Sigma$ then $\Gamma \triangleright M[x := v] \leq^* M'[x := v']: \psi, \Sigma$.*

2. *If $v \leq^* M: \phi, \Sigma$ then $\exists w. v \leq^* w: \phi, \Sigma$ and $w \leq M: \phi, \Sigma$.*

Proof. Part 1 follows by induction on M . Part 2 by induction on v . ■

The following property relates the \leq^* relation and the termination relation.

Proposition 3 *If $S \leq^* S', M \leq^* M': \phi, \Sigma, \mathcal{E} \leq^* \mathcal{E}'$ and $(S, M, \mathcal{E}) \searrow_{\downarrow}$ then $(S', M', \mathcal{E}') \searrow_{\downarrow}$.*

Proof. By induction on depth of $(S, M, \mathcal{E}) \searrow_{\downarrow}$. ■

Corollary 1 *If $(S, M, \mathcal{E}) \searrow_{\downarrow}$ and $M \leq^* N: \phi, \Sigma$ then $(S, N, \mathcal{E}) \searrow_{\downarrow}$.*

We can now show the other direction of the coincidence of the \leq^* relation and ciu-refinement (cf. Proposition 2)

Proposition 4 *If $M \leq^* N: \phi, \Sigma$ then $M \leq N: \phi, \Sigma$.*

Proof. Assume that $M \leq^* N: \phi, \Sigma$ and that $(S, M, \mathcal{E}) \searrow_{\downarrow}$. Then by corollary 1, it follows that $(S, N, \mathcal{E}) \searrow_{\downarrow}$. ■

As the \leq^* relation is trivially compatible, Propositions 2 and 4 allow one to conclude that ciu-refinement is also compatible. In addition, ciu-refinement is transitive (Lemma 2), which allows one to conclude that it is a precongruence.

Proposition 5 *\leq is a precongruence.*

This now allows us to conclude that ciu-refinement is included in contextual refinement.

Proposition 6 *If $M \leq N: \phi, \Sigma$ then $M \sqsubseteq N: \phi, \Sigma$.*

Proof. By Proposition 5 we have that $\mathcal{C}[M] \leq \mathcal{C}[N]: \psi$ which is sufficient. ■

An important property of the termination relation and evaluation contexts is the following.

Lemma 7 $(S, E[M], \mathcal{E}) \searrow$ iff $(S \circ^{\lceil} E[\bullet]^{\rceil}, M, \mathcal{E}) \searrow$.

Proof. By induction on the structure of $E[\bullet]$. ■

Evaluation contexts and frame stacks can be seen to be in one-to-one correspondence as follows.

Lemma 8 1. $\forall E[\bullet], \exists S$ such that $S =^{\lceil} E[\bullet]^{\rceil}$.

2. $\forall S, \exists E[\bullet]$ such that $S =^{\lceil} E[\bullet]^{\rceil}$.

We can now conclude that contextual refinement is included in ciu-refinement.

Proposition 7 If $M \sqsubseteq N: \phi, \Sigma$ then $M \leq N: \phi, \Sigma$.

Proof. By assumption we have that if $(id, \mathcal{C}[M], \mathcal{E}) \searrow$ then $(id, \mathcal{C}[N], \mathcal{E}) \searrow$. In particular we can take only the contexts which are evaluation contexts; thus, if $(id, E[M], \mathcal{E}) \searrow$ then $(id, E[N], \mathcal{E}) \searrow$. From Lemma 7 we have that if $(E[\bullet]^{\lceil}, M, \mathcal{E}) \searrow$ then $(E[\bullet]^{\lceil}, N, \mathcal{E}) \searrow$. From Lemma 8 we are done. ■

Thus we have proved the ciu-theorem.

Theorem 2 $\Gamma \triangleright M \sqsubseteq N: \phi, \Sigma$ iff $\Gamma \triangleright M \leq^{\circ} N: \phi, \Sigma$.

Proof. The λ -closed instances follow from Propositions 6 and 7. It is relatively straightforward to extend these to the open versions. ■