Full Abstraction and the Context Lemma[1]

Trevor Jim[2] and Albert R. Meyer[3]
MIT Laboratory for Computer Science

December 20, 1991

[1]Extended version of the paper appearing in the proceedings of TACS '91 (Meyer and Ito, Eds., volume 526 of *Lecture Notes in Computer Science*, pages 131–151, Springer-Verlag, 1991).

[2]E-mail: `trevor@theory.lcs.mit.edu`. Work supported by ARO grant DAAL03-89-G-0071.

[3]E-mail: `meyer@theory.lcs.mit.edu`. Work supported by ONR grant N00014-89-J-1988 and NSF grant 8819761-CCR.

Abstract

It is impossible to add a combinator to PCF to achieve full abstraction for models such as Berry's stable domains in a way analogous to the addition of the "parallel-or" combinator that achieves full abstraction for the familiar cpo model. In particular, we define a general notion of rewriting system of the kind used for evaluating simply typed λ-terms in Scott's PCF. Any simply typed λ-calculus with such a "PCF-like" rewriting semantics is shown necessarily to satisfy Milner's Context Lemma. A simple argument demonstrates that any denotational semantics that is adequate for PCF, and in which certain simple Boolean functionals exist, cannot be fully abstract for *any* extension of PCF satisfying the Context Lemma. An immediate corollary is that stable domains cannot be fully abstract for any extension of PCF definable by PCF-like rules.

Keywords: Stable functions, full abstraction, context lemma, PCF, standardization.

AMS(MOS) subject classifications: 03B40 68N15 68Q40 68Q50 68Q55

1 Introduction

A paradigmatic example of a functional programming language is PCF, Scott's simply typed λ-calculus for recursive functions on the integers [32]. Many categories of denotational meaning are known to *adequately* reflect the computational behavior of PCF in a precise technical sense, namely, a PCF term evaluates to the numeral $\underline{n}$ iff it means the integer n. But typically there are pairs of terms with distinct meanings that nevertheless are computationally indistinguishable in PCF. For example, with the semantics based on cpo's, PCF must be extended with a "parallel-or" combinator in order to express enough computations to be *fully abstract*, *i.e.*, semantical distinctions and computational distinctions between terms coincide [31, 30].

The problem of characterizing a fully abstract model of unextended PCF remains open after nearly two decades, *cf.* [27, 8, 28, 36]. Efforts to construct spaces of "sequential" functions corresponding to those definable in the original PCF without parallelism have led to the discovery of a number of new domains suitable for denotational semantics. Although none are fully abstract for PCF, one motivation for the development of spaces such as the *stable* functions, *bistable* functions, *sequential algorithms* [5, 4, 8, 7, 15], and most recently the *strongly stable* functions [13] was that they captured various aspects of sequentiality and so seemed "closer" to full abstraction for unextended PCF than the popular cpo model.

The stable function model in particular has a simple definition and attractive category-theoretic properties. Its only apparent technical peculiarity is that stable domains of functions are not partially ordered pointwise; in general, the stable ordering strictly refines the pointwise ordering. Nevertheless, just as for the cpo model, the elements of stable domains of type $\sigma \to \tau$ are actually total functions from elements of type σ to elements of type τ. Likewise, there is a natural notion of finite and effective elements of stable domains, and these domains yield an adequate least fixed-point model for PCF. Further, they form a Cartesian Closed Category with solutions for domain equations [5]. This category was also independently discovered and used in constructing a model of polymorphic λ-calculus [16]. So the stable domains seem to offer a setting for a theory for higher-order recursive computation with many of the attractions of the cpo category.

However, one important result about cpo's is not known for stable domains, namely, full abstraction with respect to some extension of PCF analogous to the parallel-or extension which Plotkin and Sazonov provided for the cpo model. What might a symbolic-evaluator for an extended PCF look like if it was well matched—fully abstract—with the stable model? We conclude that such an evaluator will have to be unusual looking: it cannot be specified by the kind of term-rewriting based evaluation rules known for PCF and its extensions.

The significance of this negative result hinges heavily on how drastic we judge it to go beyond the scope of PCF-like rules. It is of course possible that some operational behavior that we declare to be non-PCF-like, in our technical sense, will nevertheless

offer a useful extension of PCF for which stable domains are fully abstract. For example, Bloom [10] provides such an extension for complete lattice models, though he goes on to criticize the rather complex algorithmic specification of the combinators in his extension. (The general benefits of structured approaches to operational semantics and connections to full abstraction are discussed in [26, 11].)

To illustrate the generality of our notion of PCF-like rules, we note that the standard extensions of PCF by parallel-or and existential combinators are easily seen to be PCF-like. For example, we can define an evaluator for Plotkin's $\exists$ constant [30] while remaining within a term rewriting discipline, as follows. Let $p : \iota \to o$ be an "integer predicate" variable, and use the rules:

$$\begin{aligned} \exists p &\to \mathtt{cond}\,(p\underline{n})\,\mathtt{tt}\,\Omega, \\ \exists p &\to \mathtt{cond}\,(p\Omega)\,\Omega\,\mathtt{ff}. \end{aligned}$$

The resulting PCF-like language no longer has a confluent rewriting system, though it remains single-valued, *viz.*, every term rewrites to at most one numeral. In general, our PCF-like rules need not even be single valued.

A substantial technical contribution of this paper is a simple, modest restriction on the format of rewrite rules which is sufficient to guarantee Milner's Context Lemma [27] for languages defined by such rules. Informally, this "Approximation" Context Lemma requires that if two phrases M, N of the same syntactic functional type yield visibly distinct computational outcomes when used in *some* language context, then there are actual parameters of appropriate argument type, such that M and N each simply *applied* to these arguments, yield visibly distinct computational outcomes. This property, more perspicuously dubbed *operational extensionality* by Bloom [9, 10], has been identified by many authors as technically significant in program semantics [37, 29, 24, 1, 18, 2, 35]. The key to the proof of the Context Lemma is a new Standard Reduction Theorem 25 for PCF-like rewrite systems.

Our work borrows much from Bloom [9, 10]. The second author raised the question of whether there is a "reasonable" extension of PCF that would yield a fully abstract evaluator for lattice models [33, 34]. In answering this question, Bloom emphasized how the Context Lemma and full abstraction were incompatible with *single-valued* evaluators for the lattice model. He also characterized a general class of *consistent* rewrite rules that ensured the soundness of the Context Lemma. However, in order to encompass the computational behavior of the $\exists$ combinator, Bloom needed to develop an auxiliary notion of "observation calculi".

Our PCF-like rules are, in an appropriate sense, as powerful as Bloom's observational calculi, and strictly subsume the class of consistent rules. In particular, consistent rules are necessarily confluent and hence single-valued; as Bloom remarks [9], introducing a `join` combinator with simple multiple-valued rewrite rules yields a PCF extension both fully abstract for the lattice model and also satisfying the Context Lemma. Our wish

to simplify Bloom's criteria while dealing with nonconfluent rewriting systems forced us, however, to a rather elaborate theory of standard reductions.

As an aside, we also point out that it is questionable whether the (bi)stable and similar domains are closer to full abstraction for PCF. In particular, although some operationally valid equations that fail in the cpo model do hold, for example, in the stable model, we note in Corollary 15 that the converse also happens: some equations that hold in the cpo model fail in the stable model. The cpo, stable and likewise the bistable models thus offer information about the operational behavior of PCF terms that is not apparently comparable, and it is hard to see how to judge which is a more accurate model.

The outline of our argument is as follows: in Section 2 we formulate the key concepts of observational approximation, adequacy, and full abstraction in a fairly general setting. Then in Section 3, Theorem 14, we give a short proof that any denotational semantics that is adequate for PCF, and in which a certain simple Boolean functional exists, cannot be fully abstract for extensions of PCF satisfying the Context Lemma. The Boolean functional is obviously not continuous in Scott's sense, but it is stably continuous, and so does appear in the stable model. We also formulate a Comparability Context Lemma which applies to the bistable domains. Section 4 gives our general notion of term rewriting systems of the kind used for symbolic evaluation of PCF terms. Then in Section 5, we show that any such system defines an observational approximation relation that must satisfy the Context Lemma [27]. An immediate corollary is Theorem 30 that there is no extension of PCF defined by PCF-like rewriting rules for which the stable domain semantics is fully abstract. A similar result for the bistable domains is announced but not proved.

2 Adequacy and Full Abstraction

Concepts concerning program behavior, such as observational congruence, adequacy, and full abstraction, can usefully be defined in a general setting consisting of:

- an arbitrary set $\mathcal{L}$, called a *language*, whose elements, $M, N, \ldots$, are called *terms*;
- partial operators $C[\cdot]$ on terms called *contexts*; and
- an arbitrary set $\mathcal{O}$, called a *notion of observation*, whose elements are predicates on terms called *observations*. When an observation is true of a term, the term is said to *yield* the observation.

We will work with languages whose operational behavior is specified by (possibly nondeterministic) symbolic evaluation of terms, so we further assume a binary relation, "evaluates to", on terms. For such languages, $\mathcal{O}_{\text{eval}}$ captures the familiar notion of observing the final output of an evaluation:

$$\mathcal{O}_{\text{eval}} \quad = \quad \{\text{“evaluates to } O\text{”} \mid O \text{ is an output term}\}.$$

Here the output terms are those terms regarded as observable "output values". These typically include the ground constants (integers, truth values, ...); λ-abstractions and finite lists of output values might also be included.

There are other notions of observation based on evaluation. For instance, $\mathcal{O}_{\text{lazy}}$ consists of the single predicate true of exactly those terms whose evaluation can terminate. And notions of observation can be based on semantics of terms, *e.g.*,

$$\mathcal{O}_{\text{int}} \quad = \quad \{\text{"has the meaning of } O\text{"} \mid O \text{ is an output term}\}.$$

In this paper, however, we will be mainly concerned with $\mathcal{O}_{\text{eval}}$.

Any notion of observation induces a preordering on terms called *observational approximation*. Intuitively, one term approximates another if, according to the chosen notion of observation, the approximated term exhibits at least as much observable behavior when used in any program as the approximating term.

Definition 1 Let $\mathcal{L}$ be a language with a notion of observation $\mathcal{O}$. A term M *observationally approximates* a term N, written $M \sqsubseteq_{obs} N$, if for all contexts $C[\cdot]$, whenever $C[M]$ is a term yielding an observation from $\mathcal{O}$, then $C[N]$ is a term yielding it as well. M and N are *observationally congruent*, written $M \equiv_{obs} N$, iff $M \sqsubseteq_{obs} N$ and $N \sqsubseteq_{obs} M$.

Observational approximation provides precise meaning for questions such as, "Does my code meet a specification?" or "Will my new implementation of a module change the behavior of the program?"

In languages like PCF with applicative syntax and a suitable notion of closed terms, analysis of observational approximation can be simplified by appealing to a *Context Lemma*:

Definition 2 Let $\mathcal{L}$ be a language with a notion of observation $\mathcal{O}$. We say a term M *applicatively approximates* a term N, written $M \sqsubseteq_{app} N$, iff for all vectors of closed terms, $\vec{P}$, whenever $M\vec{P}$ is a term yielding an observation, $N\vec{P}$ is a term yielding it as well. *The Approximation Context Lemma*[1] holds if for all closed terms M and N,

$$M \sqsubseteq_{app} N \quad \text{iff} \quad M \sqsubseteq_{obs} N.$$

A fundamental result of Milner [27] is that under $\mathcal{O}_{\text{eval}}$ with numerals taken as the output terms, PCF itself, as well as its extension with parallel-or, satisfies the Approximation Context Lemma. We will see later that the Approximation Context Lemma holds for *all* languages defined in a "PCF-like" operational discipline, including, of course, PCF and its familiar extensions.

One method for proving observational approximations is by developing an abstract meaning, $[\![M]\!]$, of a term M that is adequate to determine its observations.

[1] In particular when $\mathcal{O}$ is $\mathcal{O}_{\text{eval}}$, Bloom [9] calls this "operational extensionality" while Milner [27] uses simply "the Context Lemma". We use the more descriptive "Approximation Context Lemma" because we will later consider Context Lemma's that are not based on approximation.

Definition 3 A *meaning function* for a language $\mathcal{L}$ is a function $[\![\cdot]\!]$ from terms M to values $[\![M]\!]$ in some set, partially ordered by a relation $\sqsubseteq$. A meaning function is *compositional* iff for all terms M, N and contexts $C[\cdot]$, if $[\![M]\!] \sqsubseteq [\![N]\!]$ and $C[M]$ is a term, then $C[N]$ is a term and $[\![C[M]]\!] \sqsubseteq [\![C[N]]\!]$.

A meaning function is *adequate*[2] for a notion of observation $\mathcal{O}$ iff for all terms M, N and all observations $obs \in \mathcal{O}$,

$$\Big([\![M]\!] \sqsubseteq [\![N]\!] \text{ and } obs(M)\Big) \quad \text{implies} \quad obs(N).$$

Adequacy and compositionality guarantee that the meanings accurately predict observational approximation.

Lemma 4 *A compositional meaning function $[\![\cdot]\!]$ is adequate for a notion of observation iff for all terms M and N,*

$$[\![M]\!] \sqsubseteq [\![N]\!] \quad \text{implies} \quad M \sqsubseteq_{obs} N.$$

The ordering on adequate meanings may be strictly finer than observational approximation. In the ideal situation, known as *full abstraction*, the two orderings coincide:

Definition 5 Let $[\![\cdot]\!]$ be a meaning function for a language $\mathcal{L}$ with a notion of observation $\mathcal{O}$. We say $[\![\cdot]\!]$ is *approximation fully abstract*[3] if for all terms M and N,

$$[\![M]\!] \sqsubseteq [\![N]\!] \quad \text{iff} \quad M \sqsubseteq_{obs} N.$$

It is *equationally fully abstract* if for all M and N,

$$[\![M]\!] = [\![N]\!] \quad \text{iff} \quad M \equiv_{obs} N.$$

Approximation full abstraction trivially implies adequacy for compositional meaning functions. Assuming that each output term evaluates to itself, it follows immediately that if $[\![\cdot]\!]$ is adequate for $\mathcal{O}_{\text{eval}}$ and $[\![O]\!] \sqsubseteq [\![M]\!]$, then M evaluates to O, for any output term O. If, in addition, the meaning function is *sound* for the evaluator, we easily obtain a familiar (*cf.* [26]) alternate characterization of adequacy:

Definition 6 A meaning function $[\![\cdot]\!]$ is *sound* for an "evaluates to" relation if for all terms M and N,

$$M \text{ evaluates to } N \quad \text{implies} \quad [\![M]\!] = [\![N]\!].$$

[2] As with the Context Lemma, we might more descriptively call this "approximation adequate"; but we will use only the version of adequacy based on approximation, and call it simply adequacy for brevity.

[3] Stoughton [36] calls this "inequationally fully abstract".

Lemma 7 *A sound, compositional meaning function* $[\![\cdot]\!]$ *is adequate for* $\mathcal{O}_{\text{eval}}$ *iff*

$$[\![O]\!] = [\![M]\!] \quad \text{iff} \quad M \text{ evaluates to } O,$$

for all terms M *and output terms* O.

This paper focuses specifically on the language PCF and its extensions. The precise (usual) definitions of PCF syntax and semantics appear in Appendix A, and we provide only a quick review here.

PCF is a simply typed λ-calculus with Boolean and natural number ground types, numerals $\underline{n}$ for $n \geq 0$, Boolean constants `tt` and `ff`, and simple arithmetic, recursion, and conditional operators. The evaluation relation $\twoheadrightarrow$ of the language is given by term rewriting rules.

Definition 8 An *extension of PCF* is a simply typed language together with a set of rewrite rules. The types, typed constants, and rewrite rules of the extension must include those of PCF. The extension is *conservative* iff for all PCF terms M, and all terms N in the extension,

$$M \twoheadrightarrow_{\text{extended}} N \quad \text{iff} \quad M \twoheadrightarrow_{\text{PCF}} N.$$

Observational congruence, adequacy, *etc.*, for PCF and its extensions will be defined with respect to $\mathcal{O}_{\text{eval}}$, where we take the rewriting relation $\twoheadrightarrow$ as the "evaluates to" relation, and the output terms are the ground constants `tt`, `ff`, and $\underline{n}$ for $n \geq 0$.

The results of the next section, which examines full abstraction for models of extensions of PCF, require that we prove facts about the meanings of terms while knowing very little about the extensions or the models. We will only have adequacy, conservativity, and a few other assumptions to work with. The following lemma shows that this gives us enough to reason about the unextended terms of the language.

Lemma 9 *If a model is adequate for a conservative extension of PCF, then it is also adequate for PCF.*

Proof: Suppose a model $[\![\cdot]\!]$ is adequate for a conservative extension of PCF, and $[\![M]\!] \sqsubseteq [\![N]\!]$ for some PCF terms M, N. All models are compositional, so $[\![C[M]]\!] \sqsubseteq [\![C[N]]\!]$ for any PCF context $C[\cdot]$. So for any ground PCF constant c, if $C[M] \twoheadrightarrow_{\text{extended}} c$, then $C[N] \twoheadrightarrow_{\text{extended}} c$ by adequacy. And then by conservativity, if $C[M] \twoheadrightarrow_{\text{PCF}} c$, then $C[N] \twoheadrightarrow_{\text{PCF}} c$. Hence, $M \sqsubseteq_{obs}^{\text{PCF}} N$. ■

We will further require that our models be sound, and that the ground types o and ι be interpreted as the flat cpos $\{tt, ff\}_\perp$ and $\{0, 1, \ldots\}_\perp$, with the standard interpretation of `tt`, `ff`, and the numerals $\underline{n}$. Such models will be called *models with Booleans* (though they are indeed also models with integers). Two models with Booleans of particular

interest are the cpo model $\mathcal{C}[\![\cdot]\!]$ and the stable model $\mathcal{S}[\![\cdot]\!]$. Both models are adequate but not fully abstract for PCF.

The additional information about the ground types of models with Booleans is in fact enough to determine the meanings of ground PCF terms.

Lemma 10 *The meaning of any closed PCF term of ground type is the same in all models with Booleans that are adequate for PCF.*

Proof: Let M be a closed PCF term of type o (the case $M : \iota$ is similar). In PCF, exactly one of the following holds: (1) $M \twoheadrightarrow_{\rm PCF} \mathtt{tt}$; (2) $M \twoheadrightarrow_{\rm PCF} \mathtt{ff}$; or (3) neither (1) nor (2) holds. And by Lemma 7, $M \twoheadrightarrow_{\rm PCF} \mathtt{tt}$ iff $[\![M]\!] = [\![\mathtt{tt}]\!] = tt$ for *any* model with Booleans $[\![\cdot]\!]$ adequate for PCF. Similarly, cases (2) and (3) imply $[\![M]\!] = ff$ and $[\![M]\!] = \bot$ respectively. ■

Thus we can use any particular adequate model with Booleans, like the familiar cpo model, to discover the meaning of ground PCF terms for arbitrary adequate models with Booleans. We have less to say about terms of higher type. But the following notions are useful:

Definition 11 Let τ be a first-order type, that is, a type of the form $\sigma_1 \to \cdots \to \sigma_n$, where σ_j is a ground type for $1 \leq j \leq n$. Let $[\![\cdot]\!]_i$ for $i = 1, 2$ be type frames such that $\sqsubseteq_1$ on $[\![\sigma_j]\!]_1$ equals $\sqsubseteq_2$ on $[\![\sigma_j]\!]_2$, and let $f_i \in [\![\tau]\!]_i$. Then f_1 *pointwise approximates* f_2, written $f_1 \sqsubseteq_{\rm pnt} f_2$, iff for all $d_j \in [\![\sigma_j]\!]_1$,

$$f_1(d_1)\cdots(d_n) \quad \sqsubseteq_1 \quad f_2(d_1)\cdots(d_n).$$

It follows immediately from Lemma 10 that the functions that are the meanings of a PCF term of first-order type agree *pointwise* in all models with Booleans that are adequate for PCF. So we can use the meaning of a first-order PCF term in some particular model to reason about its meaning in any adequate model with Booleans.

However, pointwise equality is not quite the same as equality of functions. For example, consider the conditional constant $\mathtt{cond}_o : o \to o \to o \to o$. Now $\mathcal{S}[\![\mathtt{cond}_o]\!] \equiv_{\rm pnt} \mathcal{C}[\![\mathtt{cond}_o]\!]$. But the stable domain does not contain parallel-or, so the stable and cpo meanings of $o \to o \to o$ are different. Thus, $\mathcal{S}[\![\mathtt{cond}_o]\!] \neq \mathcal{C}[\![\mathtt{cond}_o]\!]$ since the two functions have different codomains.

Nevertheless, it follows immediately from the definitions that pointwise approximation has the following useful property:

Lemma 12 *Let $[\![\cdot]\!]$ be a model with Booleans that is adequate for PCF, and let M and N be closed PCF terms of first-order type. Then*

$$[\![M]\!] \sqsubseteq_{\rm pnt} [\![N]\!] \quad \text{implies} \quad M \sqsubseteq_{app} N.$$

3 Failures of Full Abstraction

Our first theorem hinges on the presence of certain simple functionals over the Booleans.

Definition 13 Let *True* be the constant *tt* function on the flat Booleans, and *True!* be the strict constant *tt* function. A *true-separator* is a function f satisfying

$$\begin{aligned} f(\mathit{True}) &= \mathit{tt}, \\ f(\mathit{True!}) &= \mathit{ff}. \end{aligned}$$

Theorem 14 *Let $[\![\cdot]\!]$ be a model with Booleans that is adequate for some conservative extension of PCF satisfying the Approximation Context Lemma. If $[\![\cdot]\!]$ contains a true-separator, it is not equationally fully abstract.*

Proof: Define the terms

$$\begin{aligned} \texttt{True} &\stackrel{\text{def}}{\equiv} \lambda x.\texttt{tt}, \\ \texttt{True!} &\stackrel{\text{def}}{\equiv} \lambda x.\texttt{cond}\, x\, \texttt{tt}\, \texttt{tt}. \end{aligned}$$

By the definition of model with Booleans, we have $[\![\texttt{True}]\!] = \mathit{True}$. And by Lemma 10, $[\![\texttt{cond}]\!] \equiv_{\text{pnt}} \mathcal{C}[\![\texttt{cond}]\!]$, so by definition of model with Booleans, we have $[\![\texttt{True!}]\!] = \mathit{True!}$. Then $\texttt{True!} \sqsubseteq_{app} \texttt{True}$ by Lemmas 9 and 12. So by the Approximation Context Lemma, $\texttt{True!} \sqsubseteq_{obs} \texttt{True}$.

We conclude that there is no term P defining a true-separator; otherwise True! and True yield distinct observations in the context $(P\ [\cdot])$, contradicting the fact that $\texttt{True!} \sqsubseteq_{obs} \texttt{True}$.

However, we can define a true-separator *detector*, D, as follows:

$$D \stackrel{\text{def}}{\equiv} \lambda x.\texttt{cond}\,(x\,\texttt{True})\,(\texttt{cond}\,(x\,\texttt{True!})\,\Omega^o\,\texttt{tt})\,\Omega^o,$$

where Ω^o is the divergent term $(\texttt{Y}_o(\lambda z^o.z))$. By Lemma 10, $[\![\Omega^o]\!] = \mathcal{C}[\![\Omega^o]\!] = \bot$, and so

$$[\![D]\!](f) = \begin{cases} \mathit{tt} & \text{if } f \text{ is a true-separator,} \\ \bot & \text{otherwise.} \end{cases}$$

Now $[\![\lambda x.\Omega^o]\!]$ is the constant $\bot$ function, so $[\![D]\!] \neq [\![\lambda x.\Omega^o]\!]$, since they differ exactly on arguments that are true-separators. But since true-separators are not definable by terms, D and $\lambda x.\Omega^o$ are applicatively congruent. Then by the Approximation Context Lemma, they are observationally congruent, contradicting equational full abstraction. ■

Corollary 15 *If a stable function model with Booleans is adequate for a conservative extension of PCF that satisfies the Approximation Context Lemma, then the model is not equationally fully abstract.*

Proof: Every stable function model with Booleans contains a true-separator *truesep*, defined as follows:

$$truesep(g) \;=\; \begin{cases} tt & \text{if } g = True, \\ ff & \text{if } g = True!, \\ \bot & \text{otherwise.} \end{cases}$$

■

Corollary 16 *The PCF equations valid in the stable model do not include those valid in the cpo model.*

Proof: Just note that $\mathcal{C}[\![D]\!] = \mathcal{C}[\![\lambda x.\Omega^o]\!]$, but $\mathcal{S}[\![D]\!] \neq \mathcal{S}[\![\lambda x.\Omega^o]\!]$. ■

Our proof of Corollary 15 of course takes advantage of the notable fact that the stable ordering of functions differs from the pointwise ordering, *e.g.*, the pair of functions *True* and *True!* are ordered pointwise but are stable-incomparable. In fact, the first few lines of the proof of Theorem 14 already show that *in*equational full abstraction is incompatible with the Approximation Context Lemma for any model in which *True* and *True!* are incomparable; the rest of the proof justifies the stronger conclusion that *equational* full abstraction fails as well.

We remark that the authors of [13] have informed us that their *strongly stable* models are adequate models with Booleans for PCF, and that *truesep* is strongly stable, so Theorem 15 and Corollary 16 hold for strongly stable models.

Berry realized that altering the pointwise ordering of functions caused difficulties, and he proposed from the start an additional *bistable* model which combines stability with the pointwise ordering. Since the counterexample of Corollary 15 relies on the non-pointwise stable ordering, it does not apply to the bistable model.

There is, however, an interesting counterexample to the full abstraction of the bistable model that provides a starting point for extending our results. The counterexample, noted in [15], has its roots in the fundamental motivation behind stable models, *viz.*, to eliminate elements like parallel-or. Consider the following definition:

Definition 17 Let *lor* be the or-function that is strict in its left argument, and *ror* be the or-function that is strict in its right argument. An *or-separator* is a function f satisfying

$$\begin{aligned} f(lor) &= tt, \\ f(ror) &= ff. \end{aligned}$$

The cpo model contains a parallel-or function which bounds the left- and right-strict or-functions, and thus, by monotonicity, cannot contain an or-separator. Since the cpo model is adequate for PCF, an or-separator is not definable in PCF. On the other hand, the stable and bistable models do not contain parallel-or, and in fact, both contain or-separators.

Thus in extending the results to the bistable model, one might try to use an or-separator in the role played by the true-separator in the stable case. Since neither *lor* nor *ror* applicatively approximates the other, an argument based on the Approximation Context Lemma will not work; but a similar argument based on a notion of *observational comparability* does apply:

Definition 18 Let $\mathcal{L}$ be a language with a notion of observation $\mathcal{O}$. Terms M and N are *directly comparable* provided the set of observations yielded by M is setwise comparable to that yielded by N. The terms are *observationally comparable*, written $M \sim_{obs} N$, if for all contexts $C[\cdot]$, the terms $C[M]$ and $C[N]$ are directly comparable. They are *applicatively comparable*, written $M \sim_{app} N$, if for all vectors $\vec{P}$ of closed terms, $M\vec{P}$ and $N\vec{P}$ are directly comparable. $\mathcal{L}$ with $\mathcal{O}$ is said to *satisfy the Comparability Context Lemma* if for all closed terms M and N,

$$M \sim_{app} N \quad \text{iff} \quad M \sim_{obs} N.$$

Theorem 19 *Let $[\![\cdot]\!]$ be a model with Booleans that is adequate for some conservative extension of PCF satisfying the Comparability and Approximation Context Lemmas. If $[\![\cdot]\!]$ contains an or-separator, it is not equationally fully abstract.*

Proof: Consider the terms

$$\begin{aligned} \mathtt{lor} &\stackrel{\text{def}}{\equiv} \lambda xy.\mathtt{cond}\, x\, \mathtt{tt}\, (\mathtt{cond}\, y\, \mathtt{tt}\, \mathtt{ff}), \\ \mathtt{ror} &\stackrel{\text{def}}{\equiv} \lambda xy.\mathtt{cond}\, y\, \mathtt{tt}\, (\mathtt{cond}\, x\, \mathtt{tt}\, \mathtt{ff}). \end{aligned}$$

By Lemmas 9, 10 and 12, we have $[\![\mathtt{lor}]\!] = \mathit{lor}$, $[\![\mathtt{ror}]\!] = \mathit{ror}$, and $\mathtt{lor} \sim_{app} \mathtt{ror}$. So by the Comparability Context Lemma, $\mathtt{lor} \sim_{obs} \mathtt{ror}$.

We conclude that there is no term P defining an or-separator; otherwise $\mathtt{lor}$ and $\mathtt{ror}$ yield distinct observations in the context $(P\ [\cdot])$, contradicting the fact that $\mathtt{lor} \sim_{obs} \mathtt{ror}$.

However, we can define an or-separator *detector* as follows:

$$D \stackrel{\text{def}}{\equiv} \lambda x.\mathtt{cond}\, (x\ \mathtt{lor})\, (\mathtt{cond}\, (x\ \mathtt{ror})\, \Omega^o\, \mathtt{tt})\, \Omega^o.$$

By Lemma 10,

$$[\![D]\!](f) \;=\; \begin{cases} \mathtt{tt} & \text{if } f \text{ is an or-separator,} \\ \bot & \text{otherwise.} \end{cases}$$

Now $[\![D]\!] \neq [\![\lambda x.\Omega^o]\!]$, since they differ exactly on arguments that are or-separators. But since or-separators are not definable by terms, D and $[\![\lambda x.\Omega^o]\!]$ are applicatively congruent. Then by the Approximation Context Lemma, they are observationally congruent, contradicting equational full abstraction. ■

Corollary 20 *If a bistable model with Booleans is adequate for a conservative extension of PCF that satisfies the Comparability and Approximation Context Lemmas, then the model is not equationally fully abstract.*

Proof: Every bistable model with Booleans contains an or-separator *orsep*, defined as follows:

$$orsep(g) \quad = \quad \begin{cases} tt & \text{if } g = lor, \\ ff & \text{if } g = ror, \\ \bot & \text{otherwise.} \end{cases}$$

■

Corollary 21 ([21]) *The PCF equations valid in the bistable model do not include those valid in the cpo model.*

Proof: Just note that $\mathcal{C}[\![D]\!] = \mathcal{C}[\![\lambda x.\Omega^o]\!]$, but $\mathcal{B}[\![D]\!] \neq \mathcal{B}[\![\lambda x.\Omega^o]\!]$, where $\mathcal{B}[\![\cdot]\!]$ is the bistable model of [5]. ■

The PCF-like languages, defined in the next section, do not satisfy the Comparability Context Lemma. In fact, an or-separator constant can defined through the following PCF-like rules:

$$\begin{aligned} \mathtt{orsep}(f) \quad &\to \quad \mathtt{cond}\,(f\ \mathtt{tt}\ \Omega^o)\,(\mathtt{cond}\,(f\ \mathtt{ff}\ \mathtt{tt})\,(\mathtt{cond}\,(f\ \mathtt{ff}\ \mathtt{ff})\ \mathtt{tt}\ \Omega^o)\,\Omega^o)\,\Omega^o, \\ \mathtt{orsep}(f) \quad &\to \quad \mathtt{cond}\,(f\ \Omega^o\ \mathtt{tt})\,(\mathtt{cond}\,(f\ \mathtt{tt}\ \mathtt{ff})\,(\mathtt{cond}\,(f\ \mathtt{ff}\ \mathtt{ff})\ \mathtt{ff}\ \Omega^o)\,\Omega^o)\,\Omega^o. \end{aligned}$$

Thus we will have to restrict the class of rules we consider if we wish to apply Theorem 19. The *consistent* rules of Bloom [10] are an important, natural candidate for the restricted class. We do not know whether the Comparability Context Lemma holds for them. However, we can prove that an or-separator is not definable in consistent systems by a method involving a notion of comparability based on logical relations, as we indicate at the end of the next section.

4 PCF-like rewrite systems

Symbolic evaluators for PCF terms are often presented as term rewriting systems. In this section, we give the basic definitions for such systems, and give our criteria for calling such a system "PCF-like". Our evaluator for PCF is given in Appendix A.

A *rewrite rule* is a pair $l \to r$ of terms of the same type, such that the free variables of the right-hand side r are included in those of the left-hand side l. We write $M \xrightarrow{\Delta}_{\pi} N$ if for some subterm Δ of M, $\Delta \to \Delta'$ is an instance of the rule π, and N is obtained from M by replacing Δ with Δ'. We will omit Δ or π as convenient.

Since all of our languages are simply typed λ-calculi, we will always include β-reduction in the rewrite rules of the language. Additionally, we may specify some set Θ of *δ-rules* defining the behavior of the constants. Together, Θ and β define the *rewriting relation* $\rightarrow_{\Theta,\beta}$ on the language $\mathcal{L}$. We omit Θ and β when they can be recovered from context.

The δ-rules of PCF have a particularly simple form:

Definition 22 A *linear ground δ-rule* is a rewrite rule of the form

$$\delta m_1 m_2 \cdots m_n \rightarrow P,$$

where each m_i is either a ground constant c_i or a variable x_i. The variables x_i must be distinct. A *PCF-like rewrite system* is a language $\mathcal{L}$ together with a set Θ of linear ground δ-rules on the constants of $\mathcal{L}$.

Note that this definition of "PCF-like" is meant to be generous. In particular, although the system for pure, unextended PCF is both single-valued—every term reduces to at most one constant—and confluent, PCF-like systems in general may be multiple-valued and nonconfluent.

An interesting example of a multiple-valued PCF-like system arises in [9]. There, Bloom defines an extension of PCF that is both fully abstract and denotationally universal for the lattice model of PCF. The key to the construction amounts to the addition of operators $\top : o$ and $\mathtt{join} : o \rightarrow o \rightarrow o$ with rules

$$\begin{array}{rcll} \mathtt{join}\, x\, y & \rightarrow & x, & \\ \mathtt{join}\, x\, y & \rightarrow & y, & \\ \mathtt{join}\, \underline{n_1}\, \underline{n_2} & \rightarrow & \top, & \underline{n_1} \neq \underline{n_2}, \\ \top & \rightarrow & \underline{n}, & n \geq 0. \end{array}$$

Nonconfluent but single-valued systems are also of interest. For example, [30] extends parallel PCF by an existential operator, $\exists : (\iota \rightarrow o) \rightarrow o$, to achieve a language that is fully abstract and denotationally universal for the cpo model. There, $\exists$ is defined by the deductive rules

$$\frac{p\underline{n} \twoheadrightarrow \mathtt{tt}}{\exists p \rightarrow \mathtt{tt}}, \qquad \frac{p\Omega \twoheadrightarrow \mathtt{ff}}{\exists p \rightarrow \mathtt{ff}},$$

where $\twoheadrightarrow$ is the reflexive transitive closure of $\rightarrow$. The resulting language is indeed confluent, but goes beyond mere term rewriting. Because he wanted to be able to specify constants like $\exists$, Bloom [10] introduced *observation calculi* as a definition of "PCF-like" deductive rules.

But note that if we give up confluence, it is possible to define an $\exists$ constant while remaining in a term rewriting discipline. One such definition was given in the introduction;

we provide here a second implementation, which uses the parallel-or combinator $\mathtt{por}$.

$$\begin{aligned} \exists p &\rightarrow \mathtt{por}\,(p0)\,\Big(\exists(\lambda x.p(\mathtt{succ}\,x))\Big), \\ \exists p &\rightarrow \mathtt{cond}\,(p\Omega)\,\mathtt{tt}\,\mathtt{ff}. \end{aligned}$$

This kind of rewriting is more straightforward, but actually as powerful as the deductive discipline.

Since PCF-like systems are not confluent in general, we will not be able to use confluence in our proof of the Context Lemma. Instead we will rely on a *standardization theorem*, which states that if a term M rewrites to a term N, then there is a "standard" reduction from M to N. Thus we only need consider these standard reductions in our proof.

Typically, the standard reductions are a class of reductions with a particularly nice structure. For instance, in the pure, typed λ-calculus, a standard reduction is one in which redexes are contracted from left to right.

The definition of standard reductions in PCF-like rewrite systems is more complicated because they admit the upwards creation of redexes, *cf.* [19]. However, there is a simple inductive characterization of those standard reductions that end at a ground constant. This will be sufficient to follow the proof of the Context Lemma given in the next section, so we defer the general definition of standard reductions, and the proof of the Standardization Theorem, to Appendix C.

Before defining the standard reductions to ground constants, we introduce some useful notation. Consider the set of indices

$$\{\, i \mid m_i \text{ is a constant } c_i \text{ in rule } \theta : \delta\vec{m} \rightarrow P \,\}.$$

These indices identify what we call the *critical* arguments of θ, since the rule θ applies to a term $\delta\vec{Q}$ iff $Q_i \equiv c_i$ for i in the set. For expository purposes it will be convenient to separate the critical and non-critical arguments of a constant δ (relative to some linear ground δ-rule θ).

Notation 23 Let $\theta : \delta\vec{m} \rightarrow P$ be a linear ground δ-rule with j critical arguments and k non-critical arguments. Then for vectors $\vec{A} \equiv A_1 \cdots A_j$ and $\vec{B} \equiv B_1 \cdots B_k$, we let

$$\delta_\theta\langle \vec{A}, \vec{B}\rangle \stackrel{\text{def}}{\equiv} \delta\vec{Q},$$

where $\vec{Q}$ is the interleaving of $\vec{A}$ and $\vec{B}$ such that the A_i's appear at the critical indices of $\vec{Q}$. We drop the subscript θ when it can be recovered from context.

Note that we do not require that $\delta\vec{Q}$ be an instance of $\delta\vec{m}$; we will want to use the $\delta\langle\cdot,\cdot\rangle$ notation on terms that we anticipate becoming θ-redexes over the course of a reduction.

In this notation, we write linear ground δ-rules as

$$\theta \quad : \quad \delta\langle \vec{c}, \vec{x} \rangle \to P$$

or even

$$\theta \quad : \quad \delta\langle \vec{c}, \vec{x} \rangle \to P(\vec{x})$$

when we wish to make the dependence of P on $\vec{x}$ explicit.

Definition 24 The *standard reductions to ground constants* in a PCF-like rewrite system are defined inductively as follows. We will write $M \twoheadrightarrow_s c$ for a standard reduction of a term M to a ground constant c.

- If c is a ground constant, then the 0-step reduction $c \twoheadrightarrow c$ is standard.
- If $M_1, M_2, \ldots, M_n$ are terms, and c is a ground constant, then a reduction

$$\begin{array}{rcl} (\lambda x M_1) M_2 M_3 \cdots M_n & \to_\beta & M_1[x := M_2] M_3 \cdots M_n \\ & \twoheadrightarrow_s & c \end{array}$$

  is standard.
- If $C_1, C_2, \ldots, C_n, \vec{D}, \vec{E}$ are terms, and $c, c_1, c_2, \ldots, c_n$ are ground constants, then a reduction of the following form is standard:

$$\begin{array}{lrcl} \boldsymbol{\sigma}_1 : & \delta_\theta\langle C_1 C_2 \cdots C_n, \vec{D} \rangle \vec{E} & \twoheadrightarrow & \delta_\theta\langle c_1 C_2 \cdots C_n, \vec{D} \rangle \vec{E} \\ \boldsymbol{\sigma}_2 : & & \twoheadrightarrow & \delta_\theta\langle c_1 c_2 \cdots C_n, \vec{D} \rangle \vec{E} \\ \vdots & & \twoheadrightarrow & \cdots \\ \boldsymbol{\sigma}_n : & & \twoheadrightarrow & \delta_\theta\langle c_1 c_2 \cdots c_n, \vec{D} \rangle \vec{E} \\ & & \to_\theta & P_\theta(\vec{D}) \vec{E} \\ & & \twoheadrightarrow_s & c, \end{array}$$

  where for $1 \leq i \leq n$, the subreduction $\boldsymbol{\sigma}_i$ consists of a standard reduction from the subterm C_i to the ground constant c_i.

Theorem 25 (Standardization) *For any PCF-like rewrite system, if $M \twoheadrightarrow N$, then there is a standard reduction $M \twoheadrightarrow_s N$.*

Note that if we require our rules to be *non-overlapping*, then they are a special case of *orthogonal* rewrite systems, for which both confluence and standardization have been known for some time [19]. Similarly, confluence and standardization have been known for the systems of Bloom [10], which restrict our systems by allowing only so-called *consistent* overlaps at the root. However, it is not clear whether $\exists$ can be defined in such systems, and we certainly lose the ability to define interesting non-confluent systems, such as PCF extended with `join`.

5 The Context Lemma

Once standardization is known, the Context Lemma can be proved by a straightforward adaptation of Bloom's proof for his observation calculi [10]. First we recall the following basic facts about substitutions.

Lemma 26 (Substitution Lemma) *If $x \not\equiv y$ and $y \notin \mathrm{FV}(L)$, then*

$$M[x := L]\Big[y := N[x := L]\Big] \quad \equiv \quad M[y := N][x := L].$$

Lemma 27 *If $x \notin \mathrm{FV}(P)$, then*

$$P\Big[\vec{y} := \vec{N}[x := M]\Big] \quad \equiv \quad (P[\vec{y} := \vec{N}])[x := M].$$

The Context Lemma will follow immediately from this next result.

Lemma 28 *Suppose C is a ground term, c is a ground constant, M and N are closed terms of the same type, and $M \sqsubseteq_{app} N$. If $C[x := M] \twoheadrightarrow c$, then $C[x := N] \twoheadrightarrow c$.*

Proof: By Standardization, $C[x := M] \twoheadrightarrow_s c$. We show $C[x := N] \twoheadrightarrow c$ by induction on the length of the reduction $C[x := M] \twoheadrightarrow_s c$.

1. The only reduction $C[x := M] \twoheadrightarrow_s c$ of length zero is $c \twoheadrightarrow c$. Then one of the following holds:

   (a) $C \equiv c$. Then clearly $C[x := N] \equiv c \twoheadrightarrow c$.

   (b) $C \equiv x$ and $M \equiv c$. Here $C[x := N] \twoheadrightarrow c$ because $M \sqsubseteq_{app} N$.

For the induction, we consider subcases on the form of C.

2. $C \equiv (\lambda y C_1) C_2 \cdots C_n$. Assume $x \not\equiv y$ (the case $x \equiv y$ is similar). Since M is closed, we have

   $$C[x := M] \quad \equiv \quad \Big(\lambda y (C_1[x := M])\Big) C_2[x := M] \cdots C_n[x := M].$$

   Then the reduction $C[x := M] \twoheadrightarrow_s c$ is of the form

   $$\begin{array}{rcl} C[x := M] & \equiv & \Big(\lambda y (C_1[x := M])\Big) C_2[x := M] \cdots C_n[x := M] \\ & \to_\beta & (C_1[x := M])\Big[y := C_2[x := M]\Big] C_3[x := M] \cdots C_n[x := M] \\ & \twoheadrightarrow_s & c. \end{array}$$

   By the Substitution Lemma,

   $$(C_1[x := M])\Big[y := C_2[x := M]\Big] \quad \equiv \quad (C_1[y := C_2])[x := M],$$

so our reduction can be rewritten

$$\begin{array}{rcl} C[x := M] & \equiv & \Big((\lambda y C_1)C_2 \cdots C_n\Big)[x := M] \\ & \to_\beta & \Big((C_1[y := C_2])C_3 \cdots C_n\Big)[x := M] \\ & \twoheadrightarrow_s & c. \end{array}$$

Now by β-reduction, the fact that N is closed, and the Substitution Lemma,

$$\begin{array}{rcl} C[x := N] & \equiv & \Big((\lambda y C_1)C_2 \cdots C_n\Big)[x := N] \\ & \to_\beta & \Big((C_1[y := C_2])C_3 \cdots C_n\Big)[x := N]. \end{array}$$

And by induction,

$$\Big((C_1[y := C_2])C_3 \cdots C_n\Big)[x := N] \quad \twoheadrightarrow \quad c.$$

Thus we have a reduction $C[x := N] \twoheadrightarrow c$ as desired.

3. $C \equiv \delta C_1 \cdots C_n$. Then the reduction $C[x := M] \twoheadrightarrow_s c$ must contract the head δ by some rule $\theta : \delta_\theta\langle \vec{d}, \vec{y}\rangle \to P(\vec{y})$ (where each d_i is a ground constant). Accordingly, we rewrite C as

$$C \quad \equiv \quad \delta_\theta\langle \vec{D}, \vec{E}\rangle \vec{F}.$$

Then the reduction $C[x := M] \twoheadrightarrow_s c$ is of the form

$$\begin{array}{rcl} C[x := M] & \equiv & \delta_\theta\langle \vec{D}[x := M], \vec{E}[x := M]\rangle \vec{F}[x := M] \\ & \twoheadrightarrow & \delta_\theta\langle \vec{d}, \vec{E}[x := M]\rangle \vec{F}[x := M] \\ & \to_\theta & P(\vec{E}[x := M])\vec{F}[x := M] \\ & \twoheadrightarrow_s & c, \end{array}$$

where each $D_i[x := M] \twoheadrightarrow_s d_i$ in turn. By Lemma 27,

$$P(\vec{E}[x := M]) \quad \equiv \quad P(\vec{E})[x := M],$$

so the reduction can be rewritten

$$\begin{array}{rcl} C[x := M] & \equiv & (\delta_\theta\langle \vec{D}, \vec{E}\rangle \vec{F})[x := M] \\ & \twoheadrightarrow & (\delta_\theta\langle \vec{d}, \vec{E}\rangle \vec{F})[x := M] \\ & \to_\theta & (P(\vec{E})\vec{F})[x := M] \\ & \twoheadrightarrow_s & c. \end{array}$$

Again by Lemma 27,

$$P(\vec{E}[x := N])\vec{F}[x := N] \quad \equiv \quad (P(\vec{E})\vec{F})[x := N].$$

And by induction, $(P(\vec{E})\vec{F})[x := N] \twoheadrightarrow c$, and $D_i[x := N] \twoheadrightarrow d_i$. Thus we have found a reduction

$$\begin{array}{rcl} C[x := N] & \equiv & (\delta_\theta\langle \vec{D}, \vec{E}\rangle\vec{F})[x := N] \\ & \twoheadrightarrow & (\delta_\theta\langle \vec{d}, \vec{E}\rangle\vec{F})[x := N] \\ & \rightarrow_\theta & (P(\vec{E})\vec{F})[x := N] \\ & \twoheadrightarrow & c. \end{array}$$

4. $C \equiv xC_1 \cdots C_n$. Then consider the term

$$C' \quad \stackrel{\text{def}}{\equiv} \quad MC_1 \cdots C_n.$$

Note that $C[x := M] \equiv C'[x := M]$, so $C'[x := M] \twoheadrightarrow_s c$. Moreover C' must be of a form considered in the two previous cases, and so by the previous argument we conclude $C'[x := N] \twoheadrightarrow c$. Now consider the applicative *context*

$$C''[\cdot] \quad \stackrel{\text{def}}{\equiv} \quad [\cdot]C_1[x := N] \cdots C_n[x := N].$$

Since $C''[M] \equiv C'[x := N]$, we have $C''[M] \twoheadrightarrow c$. Finally, $M \sqsubseteq_{app} N$ implies $C''[N] \twoheadrightarrow c$; and

$$\begin{array}{rcl} C''[N] & \equiv & NC_1[x := N] \cdots C_n[x := N] \\ & \equiv & C[x := N], \end{array}$$

so $C[x := N] \twoheadrightarrow c$.

Note that we need not consider the case $C \equiv yC_1 \cdots C_n$, where $y \not\equiv x$, since then $C[x := M]$ can never reduce to a ground constant. ■

Theorem 29 (Approximation Context Lemma) *In any PCF-like rewrite system,*

$$M \sqsubseteq_{obs} N \quad \text{iff} \quad M \sqsubseteq_{app} N$$

for all closed terms M and N.

Proof:

($\Longrightarrow$) Trivial.

($\Longleftarrow$) It is sufficient to show the following: for all ground contexts $C[\cdot]$ and ground constants c, if $C[M] \twoheadrightarrow c$, then $C[N] \twoheadrightarrow c$.

Remember that the action of placing a term into the "holes" of a context differs from substitution only in that free variables of the term can be captured. But M and N are closed, with no free variables to capture; so for any context $C[\cdot]$,

$$\begin{array}{rcl} C[M] & \equiv & (C[x])[x := M], \\ \text{and} \quad C[N] & \equiv & (C[x])[x := N], \end{array}$$

where x is a fresh variable. So by Lemma 28, if $C[M] \twoheadrightarrow c$, then $C[N] \twoheadrightarrow c$ as well. ■

We now have immediately from Corollary 15:

Theorem 30 *Every stable function model with Booleans that is adequate for a conservative extension of PCF defined by PCF-like rewrite rules is not equationally fully abstract.*

We remark that a simple sufficient condition to ensure that an extension of PCF by PCF-like rules is conservative is that δ-rules whose left-hand sides involve no new (non-PCF) constants must be exactly the rules of PCF.

Because we are unable to prove a Comparability Context Lemma for consistent PCF-like rewrite rules, Corollary 20 cannot be applied. Nevertheless, our analysis of comparability can be extended to show:

Theorem 31 *Every bistable model with Booleans that is adequate for a conservative extension of PCF defined by consistent PCF-like rewrite rules is not equationally fully abstract.*

This will be proved in a forthcoming paper.

6 Conclusions and Future Work

We have extended the metatheory of term rewriting semantics for simply typed λ-calculi and have shown that certain denotational models, in particular those based on stable and strongly stable domains, cannot be fully abstract for such operational semantics. Our proof exploits the lack of order-extensionality in these domains, but an extension of our results to certain order-extensional domains such as the bistable domains is possible and will be the subject of a forthcoming paper.

The category of sequential algorithms [6] is technically not a model in our sense, but is like the stable model in that it is a Cartesian Closed Category with partially ordered function objects that are not pointwise ordered. We believe that with some minor modifications our results will apply to it as well. (This claim stands in apparent contradiction to the results of [6], which shows that the language CDS, based on concrete data structures [22], is fully abstract for the sequential algorithm model. However, it seems questionable to us to call a language such as CDS "PCF-like", since it does not have λ-abstraction or even variables.)

We conjecture that our methods and results will extend to untyped versions of PCF-like languages. Extensions to lazy and call-by-value languages also seem plausible, though with more difficulties, since higher order terms now yield observations and the notion of lazy model is more technical.

A particular open problem that we have not yet resolved is the case when the definition of model with Booleans is relaxed to allow "extra" Boolean elements, *e.g.*, if the Boolean type is interpreted as $\{tt, ff, error\}_\perp$. Finally, although we are able to show the failures of some order-extensional models, like the bistable models, the extensional embedding methods of [12] offer a more sophisticated way to restore order-extensionality which, for example, guarantees that the theory of the extensionally embedded models includes that of cpo's. We do not know whether these models can avoid the kind of failure of full abstraction that we have identified.

How great a failing of, for example, the stable domains, is lack of full abstraction? The category of stable domains is mathematically rich and offers a plausible formulation of higher-order effective computability. We have shown that stable computability cannot be captured precisely in the familiar rewriting style of operational semantics which works for the cpo or even the lattice models. But as we observed in the introduction, the failures of full abstraction we have shown might be avoidable by some other attractive, as yet undeveloped, operational semantics. Such an operational semantics would be interesting to see; and indeed, some recent work of Cartwright and Felleisen [14] suggests a fruitful development in this direction.

Acknowledgments

We are grateful to G. Berry, B. Bloom, P.-L. Curien, J.-J. Lévy, G.D. Plotkin, and Scott Smith for helpful discussions.

References

[1] S. Abramsky. The lazy lambda calculus. In D. L. Turner, editor, *Research Topics in Functional Programming*. Addison-Wesley Publishing Co., 1989.

[2] S. Abramsky. Domain theory in logical form. *Ann. Pure Appl. Logic*, 51:1–77, 1991.

[3] H. P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, second edition, 1984.

[4] G. Berry. Séquentialité de l'evaluation formelle des lambda-expressions. In B. Robinet, editor, *Program Transformations, 3^{eme} Colloque International sur la programmation*, pages 67–80, 1978.

[5] G. Berry. Stable models of typed lambda-calculi. In G. Ausiello and C. Böhm, editors, *Automata, Languages and Programming: Fifth Colloquium*, volume 62 of *Lecture Notes in Computer Science*, pages 72–89. Springer-Verlag, July 1978.

[6] G. Berry and P.-L. Curien. Sequential algorithms on concrete data structures. *Theor. Comput. Sci.*, 20(3):265–321, July 1982.

[7] G. Berry and P.-L. Curien. Theory and practice of sequential algorithms: the kernel of the programming language CDS. In M. Nivat and J. C. Reynolds, editors, *Algebraic Methods in Semantics*, chapter 2, pages 35–87. Cambridge Univ. Press, 1985.

[8] G. Berry, P.-L. Curien, and J.-J. Lévy. Full abstraction for sequential languages: the state of the art. In M. Nivat and J. C. Reynolds, editors, *Algebraic Methods in Semantics*, chapter 3, pages 89–132. Cambridge Univ. Press, 1985.

[9] B. Bloom. Can LCF be topped? Flat lattice models of typed lambda calculus (preliminary report). In *Third Annual Symposium on Logic in Computer Science* [20], pages 282–295.

[10] B. Bloom. Can LCF be topped? Flat lattice models of typed λ-calculus. *Information and Computation*, 87(1/2):263–300, July/Aug. 1990.

[11] B. Bloom, S. Istrail, and A. R. Meyer. Bisimulation can't be traced (preliminary report). In *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages*, pages 229–239, 1988. Also appears as MIT Technical Memo MIT/LCS/TM-345; submitted for journal publication.

[12] A. Bucciarelli and T. Ehrard. Extensional embedding of a strongly stable model of PCF. In *Automata, Languages and Programming: Eighteenth Colloquium*, July 1991.

[13] A. Bucciarelli and T. Ehrard. Sequentiality and strong stability. In *Proceedings, Sixth Annual IEEE Symposium on Logic in Computer Science*, 1991.

[14] R. Cartwright and M. Felleisen. Observable sequentiality and full abstraction. Draft of October 1, 1991; to appear in POPL '92.

[15] P.-L. Curien. *Categorical Combinators, Sequential Algorithms and Functional Programming.* John Wiley and Sons, 1986.

[16] J.-Y. Girard. The system F of variable types, fifteen years later. *Theor. Comput. Sci.*, 45:152–192, 1986.

[17] J. R. Hindley and J. P. Seldin. *Introduction to Combinators and λ-calculus*, volume 1 of *London Math. Soc. Student Texts*. Cambridge Univ. Press, 1986.

[18] D. J. Howe. Equality in lazy computation systems. In *Proceedings, Fourth Annual Symposium on Logic in Computer Science*, pages 198–203. IEEE, 1989.

[19] G. Huet and J.-J. Lévy. Computations in nonambiguous term rewriting systems. Technical Report 359, INRIA, Rocquencourt, France, 1979.

[20] IEEE. *Third Annual Symposium on Logic in Computer Science*, 1988.

[21] T. Jim and A. R. Meyer. Communication in the TYPES electronic forum (types@theory.lcs.mit.edu). June 17^{th}, 1989.

[22] G. Kahn and D. B. MacQueen. Coroutines and networks of parallel processes. In B. Gilchrist, editor, *Information Processing '77*, pages 993–998. North-Holland Publishing Co., 1977.

[23] J. W. Klop. Combinatory reduction systems. Tract 127, Mathematisch Centrum, Amsterdam, 1980.

[24] I. Mason and C. Talcott. Programming, transforming, and proving with function abstractions and memories. In G. Ausiello, M. Dezani-Ciancaglini, and S. R. D. Rocca, editors, *Automata, Languages and Programming: 16^{th} International Colloquium*, volume 372 of *Lecture Notes in Computer Science*. Springer-Verlag, 1989.

[25] A. R. Meyer. What is a model of the lambda calculus? *Information and Control*, 52(1):87–122, Jan. 1982.

[26] A. R. Meyer. Semantical paradigms: Notes for an invited lecture, with two appendices by Stavros Cosmadakis. In *Third Annual Symposium on Logic in Computer Science* [20], pages 236–253.

[27] R. Milner. Fully abstract models of the typed lambda calculus. *Theor. Comput. Sci.*, 4:1–22, 1977.

[28] K. Mulmuley. *Full Abstraction and Semantic Equivalence*. ACM Doctoral Dissertation Award 1986. MIT Press, 1987.

[29] C.-H. L. Ong. *The Lazy Lambda Calculus: An Investigation into the Foundations of Functional Programming*. PhD thesis, Imperial College, University of London, 1988.

[30] G. D. Plotkin. LCF considered as a programming language. *Theor. Comput. Sci.*, 5(3):223–256, Dec. 1977.

[31] V. Sazonov. Expressibility of functions in D. Scott's LCF language. *Algebra i Logika*, 15:308–330, 1976. (Russian).

[32] D. S. Scott. A type theoretical alternative to CUCH, ISWIM, OWHY. Manuscript, Oxford Univ., 1969.

[33] D. S. Scott. Continuous lattices. In F. W. Lawvere, editor, *Toposes, Algebraic Geometry and Logic*, volume 274 of *Lecture Notes in Mathematics*, pages 97–136. Springer-Verlag, 1972.

[34] D. S. Scott. Data types as lattices. *SIAM J. Comput.*, 5:522–587, 1976.

[35] S. Smith. From operational to denotational semantics. In *Mathematical Foundations of Programming Semantics*, 1991. To appear.

[36] A. Stoughton. *Fully Abstract Models of Progamming Languages*. Research Notes in Theoretical Computer Science. Pitman/Wiley, 1988. Revision of Ph.D thesis, Dept. of Computer Science, Univ. Edinburgh, Report No. CST-40-86, 1986.

[37] C. Talcott. Programming and proving with function and control abstractions. Technical Report STAN-CS-89-1288, Stanford Univ., 1988.

A PCF

Because we will work with both PCF and its extensions, we give the general definitions for simply typed λ-calculi. A language is parameterized by its ground types and typed constants; for instance, PCF's ground types are the Booleans o and the numerals ι, and its constants are listed in Figure 1.

The set of *types* of the language is the least set containing the ground types and $(\sigma \to \tau)$ for types σ and τ. The set of *first-order* types is the least set containing the ground types and $(\sigma \to \tau)$ for ground types σ and first-order types τ.

The typed *terms* of the language are defined inductively:

- A constant δ^σ is a term of type σ.
- A variable x^σ is a term of type σ.
- If M is a term of type $(\sigma \to \tau)$ and N is a term of type σ, then (MN) is a term of type τ.
- If M is a term of type τ, then $(\lambda x^\sigma M)$ is a term of type $(\sigma \to \tau)$.

We omit types and parentheses whenever possible, adopting the standard conventions of association: application associates to the left, and types associate to the right. We will use $M, N, P, \ldots$ to denote arbitrary terms; $x, y, z, \ldots$ to denote arbitrary variables; and $\sigma, \tau, \gamma, \ldots$ to denote arbitrary types. δ will always denote a constant, and c will always be a ground constant. The binary relation symbol $\equiv$ denotes syntactic equality.

Free and bound variables are defined as usual, and we consider terms that are identical modulo a change of bound variables to be syntactically identical. A term is *closed* if it has no free variables; otherwise it is *open*. A *program* is a closed term of ground type.

A *substitution* is a typed-respecting mapping of variables to terms. Substitutions are extended to terms as usual (taking care to avoid capture of free variables), and are written postfix, so that $M\rho$ is the application of the substitution ρ to the term M. We call $M\rho$ an *instance* of M. If $\vec{x} \equiv x_1, \ldots, x_n$ and $\vec{N} \equiv N_1, \ldots, N_n$, then $[\vec{x} := \vec{N}]$ is the

$\mathtt{tt}, \mathtt{ff}$	:	o	
$\underline{n}$	:	ι	for each integer $n \geq 0$
$\mathtt{succ}, \mathtt{pred}$	:	$\iota \to \iota$	
$\mathtt{zero?}$	:	$\iota \to o$	
$\mathtt{cond}_o$	:	$o \to o \to o \to o$	
$\mathtt{cond}_\iota$	:	$o \to \iota \to \iota \to \iota$	
$\mathtt{Y}_\sigma$	:	$(\sigma \to \sigma) \to \sigma$	for each type σ

Figure 1: Constants of PCF

$$
\begin{array}{rcl}
\mathtt{cond}\,\mathtt{tt}\;x\;y & \to & x \\
\mathtt{cond}\,\mathtt{ff}\;x\;y & \to & y \\
\\
\mathtt{zero?}\,\underline{0} & \to & \mathtt{tt} \\
\mathtt{zero?}\,\underline{n+1} & \to & \mathtt{ff} \\
\\
\mathtt{succ}\,\underline{n} & \to & \underline{n+1} \\
\\
\mathtt{pred}\,\underline{0} & \to & \underline{0} \\
\mathtt{pred}\,\underline{n+1} & \to & \underline{n} \\
\\
\mathtt{Y}\,f & \to & f(\mathtt{Y}\,f)
\end{array}
$$

Figure 2: Rewrite rules for PCF

substitution that maps each x_i to N_i (simultaneously), and is the identity otherwise. A special case is $[x := N]$, so that $M[x := N]$ is the result of substituting N for x in M. Sometimes we write $M \equiv M(\vec{x})$, with the intent that $M(\vec{N}) \equiv M[\vec{x} := \vec{N}]$.

A *context* $C[\cdot]$ is a term with some "holes". $C[M]$ denotes the result of putting M into the holes of $C[\cdot]$, which may cause free variables of M to become bound. We say $C[\cdot]$ is a *program context* for M if $C[M]$ is a closed term of ground type.

The interpreter of the language is defined via a rewrite system; any set of δ-rules, together with the classical rule (β), induces the one-step reduction relation $\to$. The relation $\twoheadrightarrow$ is the reflexive transitive closure of $\to$. Figure 2 gives the δ-rules for PCF.

B Simply Typed Models

Here we develop the general framework for function-based models of simply typed λ-calculi.

A *type frame* $\{[\![\sigma]\!]\}$ is collection of sets indexed by type such that $[\![\sigma \to \tau]\!]$ is a set of functions from $[\![\sigma]\!]$ to $[\![\tau]\!]$. The sets $[\![\sigma]\!]$ are called *domains*, and the elements of each $[\![\sigma]\!]$ are called *meanings* or *values* of type σ.

Since our discussion focuses on issues of adequacy and full abstraction, we also require the following:

- there is a partial order $\sqsubseteq_\sigma$ associated with each domain $[\![\sigma]\!]$;
- the functions of $[\![\sigma \to \tau]\!]$ are monotone with respect to the orderings $\sqsubseteq_\sigma$ and $\sqsubseteq_\tau$;

- and the relation $\sqsubseteq_{\sigma\to\tau}$ refines the pointwise relation on functions $f, g \in [\![\sigma \to \tau]\!]$, *i.e.*,

$$f \sqsubseteq_{\sigma\to\tau} g \quad \text{implies} \quad f(d) \sqsubseteq_\tau g(d) \text{ for all } d \in [\![\sigma]\!].$$

The last two conditions say that function application is monotone in both arguments; this implies that models, defined below, are compositional.

An *environment* is a type-respecting mapping from variables to values. If ρ is an environment, then the environment $\rho[x := d]$ is ρ with the value of x updated to d:

$$\rho[x := d](y) \;=\; \begin{cases} d & \text{if } y \equiv x, \\ \rho(y) & \text{otherwise.} \end{cases}$$

An *interpretation* is a type-respecting mapping from constants to values. For a given type frame $\{[\![\sigma]\!]\}$ and interpretation $\mathcal{I}$ we can try to define a *model*, $[\![\cdot]\!]$, that is a mapping from each term to a meaning with respect to an environment, satisfying the following conditions:

$$[\![\delta]\!]\rho \;=\; \mathcal{I}(\delta) \tag{1}$$
$$[\![x]\!]\rho \;=\; \rho(x) \tag{2}$$
$$[\![(MN)]\!]\rho \;=\; ([\![M]\!]\rho)([\![N]\!]\rho) \tag{3}$$
$$([\![\lambda x M]\!]\rho)(d) \;=\; [\![M]\!]\rho[x := d] \tag{4}$$

Implicit in condition (4) is the requirement that the function defined to be $([\![\lambda x M]\!]\rho)$ must be an element of the type frame. In other words, a model is a type frame that is closed under lambda-definability. Such closure certainly does not hold for all type frames (*cf.* [25]).

The meaning of a closed term is the same in any environment:

$$[\![M]\!]\rho \;=\; [\![M]\!]\rho'$$

for all closed M and arbitrary ρ, ρ'. Therefore we sometimes write $[\![M]\!]$ for the meaning of a closed term M, omitting the environment.

Continuity

We give the standard definitions for cpo's and continuous functions, then define the cpo model of PCF.

A *partial order* or *poset* is a set D together with a binary relation $\sqsubseteq$ that is reflexive, transitive, and anti-symmetric. We will refer to the partial order $\langle D, \sqsubseteq\rangle$ as just D. A subset $X \subseteq D$ is *directed* if every finite subset of X has an upper bound in X. A partial order D is a *complete partial order* or *cpo* if it has a least element $\bot_D$ and every directed

subset $X \subseteq D$ has a least upper bound $\sqcup X$. We omit the subscript D in $\bot_D$ when it can be recovered from context. For any set X we define the cpo $X_\bot$, with elements $X \cup \{\bot_X\}$, ordered $x \sqsubseteq y$ iff $x = y$ or $x = \bot_X$.

A function $f : D \to E$ between posets is *monotone* if $f(x) \sqsubseteq_E f(y)$ whenever $x \sqsubseteq_D y$. We say f is *continuous* if it is monotone and $f(\sqcup X) = \sqcup f(X)$ for every directed $X \subseteq D$.

The set $D \to_c E$ of continuous functions from cpo D to cpo E is a cpo under the pointwise order $\sqsubseteq_p$, defined as follows:

$$f \sqsubseteq_p g \quad \text{iff} \quad f(x) \sqsubseteq_E g(x) \text{ for all } x \in D.$$

If D is a cpo and $f : D \to D$ is continuous, then f has a least fixed point $\mathit{fix}(f)$. The function fix itself is continuous, which will allow us to interpret the recursion operator $\mathtt{Y}$.

Now we define the cpo model $\mathcal{C}[\![\cdot]\!]$ of PCF, based on continuous functions and cpos. First we construct a type frame with ground domains $\mathcal{C}[\![o]\!] = \{tt, ff\}_\bot$ and $\mathcal{C}[\![\iota]\!] = \{0, 1, 2, \ldots\}_\bot$, and higher-order domains $\mathcal{C}[\![\sigma \to \tau]\!] = \mathcal{C}[\![\sigma]\!] \to_c \mathcal{C}[\![\tau]\!]$. The cpo model of PCF is then the model $\mathcal{C}[\![\cdot]\!]$ associated with $\{\mathcal{C}[\![\sigma]\!]\}$ and the *standard interpretation*: the ground constants are interpreted in the obvious way; the constants $\mathtt{Y}_\sigma$ are interpreted as *least* fixed-point operators; and the interpretation of the remaining function constants is determined by the condition that the rewrite rules of Figure 2 be valid as equations.

Theorem 32 (Plotkin[30], Sazonov[31]) *The cpo model $\mathcal{C}[\![\cdot]\!]$ is adequate but not fully abstract for PCF.*

Stability

If D is a partial order and $X \subseteq D$, then X is *bounded* or *consistent* if there is an element $y \in D$ such that $x \sqsubseteq y$ for all $x \in X$. If elements x and y are consistent we will write $x \uparrow y$. We say D is *bounded complete* if every bounded subset $X \subseteq D$ has a least upper bound $\sqcup X$.

An element $a \in D$ is *compact* if, for every directed $X \subseteq D$ with $a \sqsubseteq \sqcup X$, there is some $x \in X$ such that $a \sqsubseteq x$. We define $\mathbf{K}D$, the *kernel* of D, to be the set of compact elements of D. The cpo D is *algebraic* if, for every $x \in D$, the set $\downarrow x = \{\, a \in \mathbf{K}D \mid a \sqsubseteq x \,\}$ is directed and $\sqcup \downarrow x = x$.

The greatest lower bound of a set X is denoted $\sqcap X$. A cpo is *distributive* if $x \sqcap (y \sqcup z) = (x \sqcap y) \sqcup (x \sqcap z)$ whenever y and z are consistent. An algebraic cpo D *has property I* if $\downarrow a$ is finite for each $a \in \mathbf{K}D$. A *dI-domain* is a distributive, bounded complete cpo that has property I.

A continuous function f between dI-domains is *stable* if whenever $x \uparrow y$, we have that $f(x \sqcap y) = f(x) \sqcap f(y)$. We let $D \to_s E$ be the set of stable functions between dI-domains D and E. As noted in [5], $D \to_s E$ ordered pointwise is *not* a dI-domain; accordingly we define the stable ordering $\sqsubseteq_s$:

$$f \sqsubseteq_s g \qquad \text{iff} \qquad f(x) = f(y) \sqcap g(x) \text{ whenever } x \sqsubseteq y.$$

Function	tt	ff	$\bot$
True	tt	tt	tt
False	ff	ff	ff
True!	tt	tt	$\bot$
False!	ff	ff	$\bot$
Id	tt	ff	$\bot$
Not	ff	tt	$\bot$
$(tt \Rightarrow tt)$	tt	$\bot$	$\bot$
$(tt \Rightarrow ff)$	ff	$\bot$	$\bot$
$(ff \Rightarrow tt)$	$\bot$	tt	$\bot$
$(ff \Rightarrow ff)$	$\bot$	ff	$\bot$
Bot	$\bot$	$\bot$	$\bot$

Figure 3: Boolean functions

If D and E are dI-domains, then $D \to_s E$ is a dI-domain under the stable order.

It must be noted that the stable order is quite different from the pointwise order. For instance, consider the monotone Boolean functions, listed in Figure 3. These functions are both continuous and stable, and so they are elements of both the continuous and stable type frames. However, the stable ordering of $o \to o$ (Figure 5) is different from its pointwise ordering (Figure 4). In particular, consider *True*, the constant tt function, and *True!*, the strict constant tt function. Although $True! \sqsubseteq_p True$, we have $True! \not\sqsubseteq_s True$, since $\bot \sqsubseteq_s tt$ but

$$True!(\bot) = \bot \quad \neq \quad tt = (True!(tt) \sqcap True(\bot)).$$

(It is this that permits the existence of the function *truesep* that was needed in Corollary 15.)

Nevertheless, a stable model $\mathcal{S}[\![\cdot]\!]$ of PCF, based on dI-domains and stable functions, can be defined in much the same way as the cpo model. The ground domains $\mathcal{S}[\![o]\!]$ and $\mathcal{S}[\![\iota]\!]$ of the stable type frame are identical to the ground domains of the cpo model. At higher types, however, we use stable functions: $\mathcal{S}[\![\sigma \to \tau]\!] = \mathcal{S}[\![\sigma]\!] \to_s \mathcal{S}[\![\tau]\!]$. Then we let $\mathcal{S}[\![\cdot]\!]$ be the model associated with the stable type frame and the (stable) standard interpretation (*cf.* the interpretation of the cpo model).

Theorem 33 (Berry[5]) *The stable model $\mathcal{S}[\![\cdot]\!]$ is adequate but not fully abstract for PCF.*

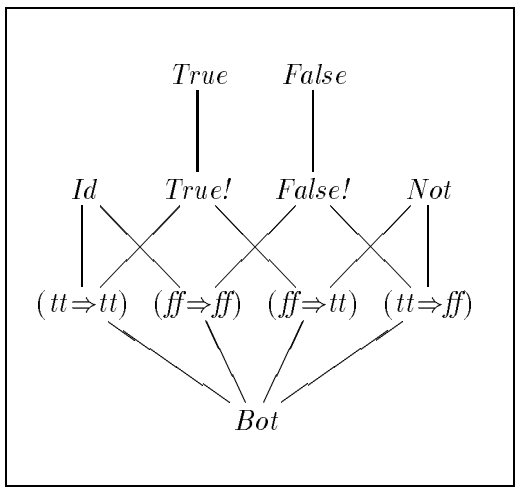

Figure 4: Pointwise ordering of $o \to o$

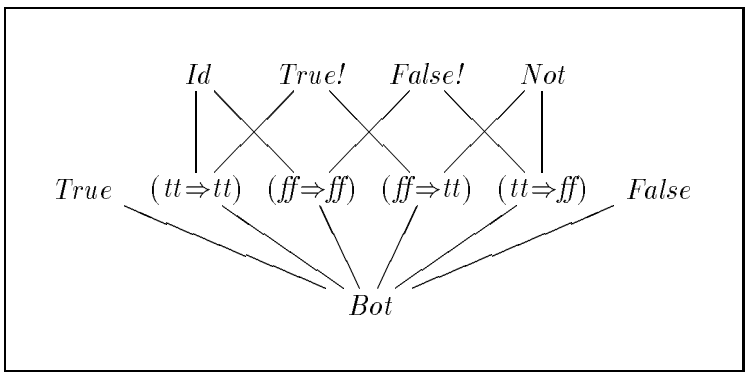

Figure 5: Stable ordering of $o \to o$

C Standard reductions in PCF-like rewrite systems

C.1 Preliminaries

This appendix gives a full definition of standard reductions and proof of the Standardization Theorem. In this section we sketch out some of the basic terminology of rewriting systems. Section C.2 introduces descendants, which allow us to trace subterms from step to step in a reduction. In Section C.3 we show that a very weak form of confluence holds for PCF-like systems; this property will be essential in proving the Standardization Theorem. Section C.4 introduces labelled rewrite systems, and proves that they are strongly normalizing. The labelled systems will be used in the proof of Standardization. The standard reductions are defined in Section C.5, and Standardization is proved in Section C.6. The proof is a variation of Klop's proof for the pure λ-calculus [23], and involves a rewriting system on reductions. The system successively rewrites non-standard reduction paths to "more standard" paths; Standardization is proved by showing that the system is strongly normalizing, and that normal forms are standard reductions.

Our presentation of the machinery used to state and prove Standardization is necessarily brief. Much of the material is covered in more depth in standard references [3, 23]. Throughout we will work with a PCF-like rewrite system given by a language, $\mathcal{L}$, and set, Θ, of linear ground δ-rules.

We assume that the reader is familiar with the following terminology. The notation $M \subset N$ denotes that M is a *subterm* of N. A subterm may appear several times in a term; multiple *occurrences* of a subterm can be distinguished by their *paths*, which specify the exact position of a subterm inside the term. When we speak of a subterm $M \subset N$ we implicitly mean a particular occurrence of M in N; the disambiguating paths are omitted.

Note that $M \to N$ iff there is an instance $\Delta \to \Delta'$ of a rule π such that $\Delta \subset M$, and N is obtained from M by replacing Δ with Δ'. We will write $M \xrightarrow{\Delta}_{\pi} N$ in this case, and we call Δ a *(π)-redex* and Δ' its *(π)-contractum.*

A *reduction (path)* $\boldsymbol{\sigma}$ is a sequence

$$\boldsymbol{\sigma} : M_1 \xrightarrow{\Delta_1}_{\pi_1} M_2 \xrightarrow{\Delta_2}_{\pi_2} M_3 \xrightarrow{\Delta_3}_{\pi_3} \cdots .$$

We will use $\boldsymbol{\sigma}, \boldsymbol{\tau}, \ldots$ to refer to reduction paths. Two reductions are *coinitial* if they start in the same term, and *cofinal* if they end in the same term.

C.2 Descendants

Consider some possible effects of a reduction $M \to N$ on a subterm $\Delta \subset M$:

- Δ could be erased, as in $(\lambda x.y)\Delta \to y$.

- Δ could be copied to some instances in N, as in $(\lambda x.\delta xx)\Delta \to \delta\Delta\Delta$.
- Δ could be left untouched and in its original position, as in $\Delta((\lambda x.x)y) \to \Delta y$.
- The contracted redex might occur within Δ, transforming it into a syntactically different subterm in the same position.

In order to define and prove standardization, we will need to speak precisely about these cases, so we introduce *descendants*, which let us track a subterm throughout a reduction. We will not define descendants in their full generality, but only for certain subterms of interest. Our definition is equivalent to the standard definition [23] on those subterms.

Descendants are introduced via an annotated rewrite system derived from $\mathcal{L}$ and Θ, in which some λ's and δ's are marked with a $*$. Thus we define the language $\mathcal{L}_*$, whose symbols are those of $\mathcal{L}$, with the addition of λ_*, and δ^σ_* for each constant δ^σ of $\mathcal{L}$. The terms of $\mathcal{L}_*$ are defined inductively:

- A constant δ^σ or δ^σ_* is a term of type σ.
- A variable x^σ is a term of type σ.
- If M is a term of type $(\sigma \to \tau)$ and N is a term of type σ, then (MN) is a term of type τ.
- If M is a term of type τ, then $(\lambda x^\sigma M)$ and $(\lambda_* x^\sigma M)$ are terms of type $(\sigma \to \tau)$.

The *erasure* $|M| \in \mathcal{L}$ of $M \in \mathcal{L}_*$ is obtained from M by leaving out the $*$'s. Substitution for the language is defined in the obvious way (with λ_*'s binding variables just as λ's). The rules of the new system include β and the rule scheme β_*:

$$\beta_* : (\lambda_* xM)N \to M[x := N].$$

Similarly, the δ-rules Θ_* of the system are derived from the rules Θ. If θ is a rule of Θ,

$$\theta : \delta\langle \vec{c}, \vec{x}\rangle \to P(\vec{x}),$$

then Θ_* contains all rules of the form θ' and θ_*:

$$\begin{aligned} \theta' : \delta\langle \vec{c}\,', \vec{x}\rangle &\to P(\vec{x}), \\ \theta_* : \delta_*\langle \vec{c}\,', \vec{x}\rangle &\to P(\vec{x}), \end{aligned}$$

where $\vec{c}\,'$ is any vector of $\mathcal{L}_*$ ground constants such that $|\vec{c}\,'| \equiv \vec{c}$.

There is a strong connection between the systems. Any Θ_*-reduction path $\boldsymbol{\sigma}$,

$$\boldsymbol{\sigma} : M_1 \overset{\Delta_1}{\to}_{\pi_1} M_2 \overset{\Delta_2}{\to}_{\pi_2} M_3 \overset{\Delta_3}{\to}_{\pi_3} \cdots,$$

projects to a Θ-reduction path $|\boldsymbol{\sigma}|$:

$$|\boldsymbol{\sigma}| : |M_1| \stackrel{|\Delta_1|}{\rightarrow}_{|\pi_1|} |M_2| \stackrel{|\Delta_2|}{\rightarrow}_{|\pi_2|} |M_3| \stackrel{|\Delta_3|}{\rightarrow}_{|\pi_3|} \cdots .$$

Conversely, for any $M \in \mathcal{L}_*$ and Θ-reduction path $\boldsymbol{\sigma} : |M| \rightarrow \cdots$, there is a unique *lift* of $\boldsymbol{\sigma}$ to a Θ_*-reduction path $\boldsymbol{\sigma}' : M \rightarrow \cdots$ such that $\boldsymbol{\sigma} \equiv |\boldsymbol{\sigma}'|$.

We will be interested in tracing subterms of the form $(\lambda x.M_1)M_2$ or $\delta M_1 \cdots M_n$ throughout a reduction; that is, β-redexes and possible δ-redexes. Accordingly, we introduce the following terminology. A subterm $(\lambda x.M_1)M_2$ or $\delta M_1 \cdots M_n$ of M is called a *predescendant of* M. If $\mathcal{F}$ is a set of predescendants of $M \in \mathcal{L}$, we write $(M, \mathcal{F})$ for the $\mathcal{L}_*$ term derived from M by marking the head λ or δ of each predescendant in $\mathcal{F}$ with a $*$.

Definition 34 Suppose $\boldsymbol{\sigma} : M \rightarrow \cdots \rightarrow N$ is a Θ-reduction path.

(i) If Δ is a predescendant of M, its set of *descendants* in N relative to $\boldsymbol{\sigma}$, written $(\Delta/\boldsymbol{\sigma})$, is defined as follows.

Let $M' \equiv (M, \{\Delta\})$ and lift $\boldsymbol{\sigma}$ to $\boldsymbol{\sigma}' : M' \rightarrow \cdots \rightarrow N'$. If $\Delta \equiv (\lambda x M_1)M_2$ (resp. $\Delta \equiv \delta M_1 \cdots M_n$), then $(\Delta/\boldsymbol{\sigma}) \stackrel{\text{def}}{=} \mathcal{F}$, where $\mathcal{F}$ is the unique set of subterms of N of the form $(\lambda x M_1')M_2'$ (resp. $\delta M_1' \cdots M_n'$), such that $N' \equiv (N, \mathcal{F})$.

(ii) If $\mathcal{F}$ is a set of predescendants of M, its descendants $\mathcal{F}/\boldsymbol{\sigma}$ are defined

$$\mathcal{F}/\boldsymbol{\sigma} \stackrel{\text{def}}{=} \bigcup \{ \Delta/\boldsymbol{\sigma} \mid \Delta \in \mathcal{F} \}.$$

(iii) $\Delta \subset M$ is an *ancestor* of $\Delta' \subset N$ if $\Delta' \in \Delta/\boldsymbol{\sigma}$.

For a given reduction $M_1 \rightarrow M_2 \rightarrow M_3 \rightarrow \cdots$, we will sometimes speak of descendants and ancestors for subterms of terms M_i and M_j, where i and j are any indices such that $j \geq i$. We do not specify the reduction from M_i to M_j, as it can be recovered from context.

Note 35

(i) If $M \stackrel{\Delta}{\rightarrow} N$, then Δ has no descendants in N.

(ii) If $M \stackrel{\Delta}{\rightarrow}_\theta N$, where $\Delta \equiv \delta\langle \vec{c}, \vec{B} \rangle$, then no c_i has a descendant in N.

We mention that the following important property holds for our PCF-like systems, since it does not hold for all rewrite systems [23].

Note 36 If $\Delta \subset M$ and $M \rightarrow N$, then descendants of Δ in N are disjoint.

Disjointness of descendants does not extend to $\twoheadrightarrow$, as we indicate here:

$$\begin{aligned}(\lambda y.(\lambda x.yx)y)(\lambda z.\delta_* z) &\to_\beta (\lambda x.(\lambda z.\delta_* z)x)(\lambda z.\delta_* z)\\ &\to_\beta (\lambda x.\delta_* x)(\lambda z.\delta_* z)\\ &\to_\beta \delta_*(\lambda z.\delta_* z).\end{aligned}$$

Definition 37 Suppose M_i is a term in a reduction $\boldsymbol{\sigma}$,

$$\boldsymbol{\sigma} : M_1 \overset{\Delta_1}{\to}_{\pi_1} M_2 \overset{\Delta_2}{\to}_{\pi_2} M_3 \overset{\Delta_3}{\to}_{\pi_3} \cdots.$$

(i) We say $\Delta \subset M_i$ is *(π)-contracted (in $\boldsymbol{\sigma}$)* if for some $j \geq i$, Δ_j is a descendant of Δ and $\pi_j = \pi$.

(ii) We say $\Delta \subset M_i$ is *active (in $\boldsymbol{\sigma}$)* if there is a $\Delta' \subset \Delta$ that is contracted in $\boldsymbol{\sigma}$.

Sometimes it will be useful to specify a set of subterms of some term M, and consider reductions from M in which only those subterms are contracted. Such reductions are called *developments*. Because we work with systems in which a subterm can contract by more than one rule, our definition of developments extends the standard definition by specifying a rule for each redex contracted in a development.

Definition 38 Suppose the following: $\boldsymbol{\sigma}$ is a reduction from M to N; $\mathcal{F}$ is a set of subterms of M; and Π is a mapping that takes each $\Delta \in \mathcal{F}$ to a rule π_Δ.

(i) We call $\boldsymbol{\sigma}$ a *development* of $\mathcal{F}$ from M by Π, written $\boldsymbol{\sigma} : (M, \mathcal{F}) \overset{\Pi}{\twoheadrightarrow} N$, if each redex Δ' contracted in $\boldsymbol{\sigma}$ is a descendant of some $\Delta \in \mathcal{F}$, and Δ' is contracted by rule π_Δ.

(ii) We say a development $\boldsymbol{\sigma}$ is a *complete development*, written $\boldsymbol{\sigma} : (M, \mathcal{F}) \underset{\text{cpl}}{\overset{\Pi}{\twoheadrightarrow}} N$, if $\mathcal{F}/\boldsymbol{\sigma} = \emptyset$.

When Π is evident from context, we will omit mention of it.

Note 39 If $\mathcal{F}$ is a set of n disjoint redexes of M, then clearly all complete developments of $\mathcal{F}$ from M are of length n and are cofinal.

C.3 Properties related to confluence

Note 39 is a special case of a much stronger theorem, the Finite Developments theorem. We will not need to prove the Finite Developments theorem in its full generality; this section proves a weaker result that will be sufficient for our application.

Definition 40 We say two δ-redexes Δ_1 and Δ_2 *overlap* if either

(i) they share the same head δ, or

(ii) one Δ_i appears as a critical argument of the other.

Note that in case (ii), the Δ_i must be a ground constant.

Often, rewrite systems are constrained to avoid overlapping redexes; such systems are guaranteed to be confluent. Because we allow overlapping rules, our systems are not confluent in general. However, they do satisfy the following much weaker property, which will be essential in our proof of standardization.

Lemma 41 *Suppose* $\boldsymbol{\sigma}_1 : M_0 \stackrel{\Delta_1}{\to} M_1$ *and* $\boldsymbol{\sigma}_2 : M_0 \stackrel{\Delta_2}{\to} M_2$, *where* Δ_1 *and* Δ_2 *do not overlap. Then complete developments of* $\Delta_2/\boldsymbol{\sigma}_1$ *from* M_1 *and* $\Delta_1/\boldsymbol{\sigma}_2$ *from* M_2 *are finite and cofinal.*

Proof: For each of the various cases on the relative positions of Δ_1 and Δ_2 in M_0, we find a term M_3 that is the final term of every complete development of $\Delta_1/\boldsymbol{\sigma}_2$ and $\Delta_2/\boldsymbol{\sigma}_1$:

$$\begin{array}{ccc} M_0 & \xrightarrow{\Delta_1} & M_1 \\ \Big\downarrow{\scriptstyle \Delta_2} & & \Big\downarrow{\scriptstyle \Delta_2/\sigma_1} \\ M_2 & \xrightarrow{\Delta_1/\sigma_2} & M_3 \end{array}$$

1. Δ_1 and Δ_2 are disjoint. Then M_0, M_1, and M_2 can be written

$$\begin{aligned} M_0 &\equiv \cdots\Delta_1\cdots\Delta_2\cdots, \\ M_1 &\equiv \cdots\Delta_1'\cdots\Delta_2\cdots, \\ M_2 &\equiv \cdots\Delta_1\cdots\Delta_2'\cdots, \end{aligned}$$

   where Δ_1' and Δ_2' are the respective contractums of Δ_1 and Δ_2. Now defining

$$M_3 \stackrel{\text{def}}{\equiv} \cdots\Delta_1'\cdots\Delta_2'\cdots,$$

   we see that the only complete development of $\Delta_2/\boldsymbol{\sigma}_1$ is $M_1 \stackrel{\Delta_2}{\to} M_3$, and the only complete development of $\Delta_1/\boldsymbol{\sigma}_2$ is $M_2 \stackrel{\Delta_1}{\to} M_3$, as desired.

2. $\Delta_1 \subset \Delta_2$. Then there is a unique descendant Δ_2' of Δ_2 in M_1, and we consider three subcases.

(a) $\Delta_2 \equiv (\lambda x.\cdots\Delta_1\cdots)N$. Then we can write M_0, M_1, and M_2 as

$$\begin{aligned} M_0 &\equiv \cdots((\lambda x.\cdots\Delta_1\cdots)N)\cdots, \\ M_1 &\equiv \cdots((\lambda x.\cdots\Delta_1'\cdots)N)\cdots, \\ M_2 &\equiv \cdots((\cdots\Delta_1\cdots)[x := N])\cdots, \end{aligned}$$

where Δ_1' is the contractum of Δ_1, and $\Delta_2' \equiv (\lambda x.\cdots\Delta_1'\cdots)N$. If we take

$$M_3 \stackrel{\text{def}}{\equiv} \cdots((\cdots\Delta_1'\cdots)[x := N])\cdots,$$

then the only complete development of $\Delta_2/\boldsymbol{\sigma}_1$ is $M_1 \stackrel{\Delta_2'}{\to}_\beta M_3$. Furthermore, substitutivity holds for PCF-like rewrite systems; that is,

$$M \stackrel{\Delta}{\to} M' \implies M[x := N] \stackrel{\Delta'}{\to} M'[x := N],$$

where Δ' is Δ with any free occurrences of x replaced by N. Thus the only complete development of $\Delta_1/\boldsymbol{\sigma}_2$ is $M_2 \to M_3$.

(b) $\Delta_2 \equiv (\lambda x.N)(\cdots\Delta_1\cdots)$. Then M_0, M_1, and M_2 can be written

$$\begin{aligned} M_0 &\equiv \cdots((\lambda x.N)(\cdots\Delta_1\cdots))\cdots, \\ M_1 &\equiv \cdots((\lambda x.N)(\cdots\Delta_1'\cdots))\cdots, \\ M_2 &\equiv \cdots(N[x := (\cdots\Delta_1\cdots)])\cdots, \end{aligned}$$

where Δ_1' is the contractum of Δ_1, and $\Delta_2' \equiv (\lambda x.N)(\cdots\Delta_1'\cdots)$. Defining

$$M_3 \stackrel{\text{def}}{\equiv} \cdots(N[x := (\cdots\Delta_1'\cdots)])\cdots,$$

we see that the only complete development of $\Delta_2/\boldsymbol{\sigma}_1$ is $M_1 \stackrel{\Delta_2'}{\to}_\beta M_3$. Furthermore, descendants of Δ_1 in M_2 are disjoint, and any contraction of them in turn is a reduction $M_2 \stackrel{\Delta_1}{\to} \cdots \stackrel{\Delta_1}{\to} M_3$.

(c) $\Delta_2 \equiv \delta_\theta\langle\cdots,\cdots(\cdots\Delta_1\cdots)\cdots\rangle$. Then we write M_0, M_1, and M_2 as

$$\begin{aligned} M_0 &\equiv \cdots(\delta_\theta\langle\cdots,\cdots(\cdots\Delta_1\cdots)\cdots\rangle)\cdots, \\ M_1 &\equiv \cdots(\delta_\theta\langle\cdots,\cdots(\cdots\Delta_1'\cdots)\cdots\rangle)\cdots, \\ M_2 &\equiv \cdots(P_\theta(\cdots(\cdots\Delta_1\cdots)\cdots))\cdots, \end{aligned}$$

where Δ_1' is the contractum of Δ_1, and $\Delta_2' \equiv \delta_\theta\langle\cdots,\cdots(\cdots\Delta_1'\cdots)\cdots\rangle$. Defining

$$M_3 \stackrel{\text{def}}{\equiv} \cdots(P_\theta(\cdots(\cdots\Delta_1'\cdots)\cdots))\cdots,$$

we see that the only complete development of $\Delta_2/\boldsymbol{\sigma}_1$ is $M_1 \stackrel{\Delta_2'}{\to}_\theta M_3$. And just as in case 2b, the descendants of Δ_1 in M_2 are disjoint, so by contracting them in turn we find a reduction $M_2 \stackrel{\Delta_1}{\to} \cdots \stackrel{\Delta_1}{\to} M_3$.

3. $\Delta_2 \subset \Delta_1$. This case is handled exactly as case 2.

■

C.4 A labelled λ-calculus

For any PCF-like rewrite system, there is a corresponding *labelled* PCF-like system that is strongly normalizing. The labelling technique has led to some of the simplest proofs for many syntactic properties, and we will use it in our proof of standardization. This section introduces labelled calculi and proves that they are strongly normalizing.

The labelled system is similar to the system that we introduced earlier to define descendants. However, the systems are also different in important ways, since they are intended for different purposes. In the labelled system, we will mark δ's with nonnegative integers instead of $*$'s, and we will not need to mark λ's. Furthermore, we do not allow unmarked δ's. The reasons for this will become apparent in what follows.

For any PCF-like language $\mathcal{L}$, the language $\mathcal{L}_{\mathbb{N}}$ is just the PCF-like language with constants δ_n^σ for each constant δ^σ of $\mathcal{L}$ and each $n \in \mathbb{N}$.

Notation 42

(i) If $M \in \mathcal{L}_{\mathbb{N}}$, then $|M| \in \mathcal{L}$ is the term derived from M by erasing the labels on the constants.

(ii) If $M \in \mathcal{L}$, then $M^n \in \mathcal{L}_{\mathbb{N}}$ is the term derived from M by labelling each constant with n.

The δ-rules $\Theta_{\mathbb{N}}$ of the labelled calculus are defined as follows. If θ is a rule of Θ,

$$\theta : \delta\langle \vec{c}, \vec{x}\rangle \to P(\vec{x}),$$

then $\Theta_{\mathbb{N}}$ contains all rules of the form $\theta_{\mathbb{N}}$:

$$\theta_{\mathbb{N}} : \delta_{n+1}\langle \vec{c}', \vec{x}\rangle \to P^n(\vec{x}),$$

where $\vec{c}'$ is a vector of $\mathcal{L}_{\mathbb{N}}$ ground constants such that $|\vec{c}'| \equiv \vec{c}$. Note that there is no rule for any δ_0.

The projection $|\boldsymbol{\sigma}|$ of a $\Theta_{\mathbb{N}}$-reduction path $\boldsymbol{\sigma}$ is defined in the obvious way. And any Θ-reduction $\boldsymbol{\sigma}$ can be lifted to a $\Theta_{\mathbb{N}}$-reduction $\boldsymbol{\sigma}'$ such that $\boldsymbol{\sigma} \equiv |\boldsymbol{\sigma}'|$ (*e.g.*, label each constant in the first term of $\boldsymbol{\sigma}$ by the length of $\boldsymbol{\sigma}$).

Definition 43 A term M is *strongly normalizable* (SN) if all reductions starting at M are finite.

Theorem 44 (Strong Normalization) *Every $\mathcal{L}_{\mathbb{N}}$ term is strongly normalizable.*

The rest of this section lays out the proof of strong normalization. We use a straightforward extension of the method of [17].

Definition 45 The notion of *strong computability* (SC) of a term is defined by induction as follows:

(i) A term of ground type is SC iff it is SN

(ii) A term $M^{(\sigma\to\tau)}$ is SC iff, for every SC term N^σ, the term $(MN)^\tau$ is SC

Note 46 By definition 45(ii) a term M is SC iff, for all vectors $\vec{N}$ of SC terms driving M to ground type, the term $M\vec{N}$ is SC. And by definition 45(i), such an $M\vec{N}$ is SC iff it is SN.

Definition 47 An *atom* is a variable or a constant δ_n with no rule.

Lemma 48

(i) *If a is an atom and $\vec{N}$ is a vector of SN terms, then the term $a\vec{N}$ is SC.*

(ii) *Every SC term M is SN.*

Proof: By induction on the type of $a\vec{N}$ and M.

1. Basis: $a\vec{N}$ and M have ground type.

   (i) Since each N_i is SN, $a\vec{N}$ must be SN, and therefore SC by definition 45(i).

   (ii) By definition 45(i).

2. Induction: $a\vec{N}$ and M have type $\sigma\to\tau$.

   (i) Let P^σ be SC. By the induction hypothesis (ii), P is SN. Then by induction, the term $(a\vec{N}P)^\tau$ is SC. Therefore so is $a\vec{N}$ by definition 45(ii).

   (ii) Let x^σ be a variable not occurring in M. By the induction hypothesis (i), x is SC. Then $(Mx)^\tau$ is SC, and therefore SN by induction. But any subterm of an SN term is SN, so M is SN as well.

■

Lemma 49 *If N is SC and $M[x := N]$ is SC, then so is $(\lambda xM)N$.*

Proof: Let $\vec{P} \equiv P_1, \ldots, P_n$ be a vector of SC terms driving M to ground type. Since $M[x := N]$ is SC, the term

$$(M[x := N])\vec{P} \tag{5}$$

is SN by Note 46. The lemma follows from Note 46 if we can prove that

$$(\lambda x M) N \vec{P} \tag{6}$$

is SN.

Now since (5) is SN, all of its subterms are SN, including $M[x := N], \vec{P}$. Furthermore by hypothesis and the preceding lemma, N is SN. Therefore an infinite reduction from (6) cannot consist entirely of contractions in $M, N, P_1, \ldots, P_n$. So an infinite reduction of (6) must have the form

$$\begin{array}{rcl} (\lambda x M) N P_1 \cdots P_n & \twoheadrightarrow & (\lambda x M') N' P_1' \cdots P_n' \\ & \rightarrow & M'[x := N'] P_1' \cdots P_n' \\ & \twoheadrightarrow & \ldots \end{array}$$

(where $M \twoheadrightarrow M'$, etc.) From the reductions $M \twoheadrightarrow M'$ and $N \twoheadrightarrow N'$ we have

$$M[x := N] \twoheadrightarrow M'[x := N']$$

Then we can construct an infinite reduction from (5) as follows:

$$\begin{array}{rcl} M[x := N] P_1 \cdots P_n & \twoheadrightarrow & M'[x := N'] P_1' \cdots P_n' \\ & \twoheadrightarrow & \ldots \end{array}$$

But this contradicts the fact that (5) is SN. Therefore there is no infinite reduction from (6); it must be SN. ■

Lemma 50 *Consider a constant δ and a vector $\vec{N}$ of SC terms driving δ to ground type. If for each rule θ on δ,*

$$\theta : \delta_\theta \langle \vec{c}, \vec{x} \rangle \rightarrow P_\theta(\vec{x}),$$

where $\delta \vec{N} \equiv \delta_\theta \langle \vec{N}_1, \vec{N}_2 \rangle \vec{N}_3$, we have that

$$P_\theta(\vec{N}_2) \vec{N}_3 \tag{7}$$

is SC, then $\delta \vec{N}$ is SC.

Proof: We must show that $\delta \vec{N}$ is SN. Since the $\vec{N}$ are SC, by Lemma 48 they are SN. Therefore any infinite reduction from $\delta \vec{N}$ must look like

$$\begin{array}{rcl} \delta_\theta \langle \vec{N}_1, \vec{N}_2 \rangle \vec{N}_3 & \twoheadrightarrow & \delta_\theta \langle \vec{c}, \vec{N}_2{}' \rangle \vec{N}_3{}' \\ & \rightarrow & P_\theta(\vec{N}_2{}') \vec{N}_3{}' \\ & \twoheadrightarrow & \cdots \end{array}$$

where $\vec{N}_1 \twoheadrightarrow \vec{c}$, $\vec{N}_2 \twoheadrightarrow \vec{N}_2{}'$, etc. But then we can construct an infinite reduction from (7) as follows:

$$\begin{aligned} P_\theta(\vec{N}_2)\vec{N}_3 &\twoheadrightarrow P_\theta(\vec{N}_2{}')\vec{N}_3{}' \\ &\twoheadrightarrow \cdots \end{aligned}$$

But as (7) is SC, by Lemma 48 it is SN, a contradiction. Therefore $\delta\vec{N}$ is SN. ■

Lemma 51 *For any term M and substitution $\rho \equiv [\vec{x} := \vec{N}]$, where each N_i is SC, the term $M\rho$ is SC.*

Proof: The proof is by induction on the lexicographic ordering of (m, M), where m is the maximum δ-index appearing in M.

1. M is a variable x_i. Then $M\rho$ is N_i and the result follows.

2. M is an atom distinct from $x_1, \ldots, x_n$. Then $M\rho \equiv M$ which is SC by Lemma 48. Note that this includes all constants δ_0.

3. $M \equiv \delta_{m+1}$. Then $M\rho \equiv \delta_{m+1}$. Thus it is sufficient to show that for any vector $\vec{N}'$ of SC terms driving δ_{m+1} to ground type, the term $\delta_{m+1}\vec{N}'$ is SC.

   Consider any rule θ on δ_{m+1}:

   $$\theta : \delta_{m+1}\langle\vec{c}, \vec{x}\rangle \to P(\vec{x}).$$

   By construction of the labelled calculus, no constants in P are labelled with an index greater than m. Thus we can apply the induction hypothesis to P.

   If we rewrite $\delta_{m+1}\vec{N}'$ as $\delta_{m+1}\langle\vec{N}_1{}', \vec{N}_2{}'\rangle\vec{N}_3{}'$, by induction $P(\vec{N}_2{}')$ is SC. Then by the definition of SC, the term $P(\vec{N}_2{}')\vec{N}_3{}'$ is SC. Therefore by Lemma 50, $\delta_{m+1}\vec{N}'$ is SC.

4. $M \equiv \lambda y^\sigma M_1$. Then $M\rho \equiv \lambda y(M_1\rho)$, neglecting changes in bound variables.

   To show that $M\rho$ is SC we must show that for all SC terms N^σ, the term $(M\rho)N$ is SC. But $(M\rho)N \equiv (\lambda y(M_1\rho))N$, and

   $$(M_1\rho)[y := N] \equiv M_1[x_1 := N_1]\cdots[x_n := N_n][y := N]$$

   which is SC by induction. Therefore $(\lambda y(M_1\rho))N$ is SC by Lemma 49.

5. $M \equiv M_1M_2$. Then $M\rho \equiv (M_1\rho)(M_2\rho)$, and $M_1\rho$ and $M_2\rho$ are SC by induction. Therefore $M\rho$ is SC by definition 45(ii).

■

Proof of Theorem 44 (Strong Normalization): By Lemma 51, every term M is SC (just let $\vec{x}$ be empty). Then by Lemma 48, M is strongly normalizing. ■

C.5 Standard Reductions

Our definition of standard reductions is similar to that of [19], with a few important differences. The "linear ground" restriction imposed on our systems gives us a particularly simple class of rewrite rules, and this simplicity carries over to the definition of standard reductions. On the other hand, the systems of [19] do not include λ-abstraction, and forbid overlapping rewrite rules, which we allow.

Overlapping rules do not add much complication to the definition of standard reductions, but they are more of an obstacle in the proof of standardization. Overlapping systems are not confluent in general, so we cannot use confluence and related properties in our proof. This is offset by the fact that we consider only typed systems.

The standard reductions of [19] are based on "outside-in" reductions. Informally, outside-in reductions are reductions in which no subterm of a term reduces before the term itself contracts, unless the subterm reduces outside-in and contributes towards making the term a redex. For example, consider the PCF reduction

$$\begin{aligned} \mathtt{cond}\,(\mathtt{zero?}\,0)\;M\;N &\rightarrow \mathtt{cond}\,\mathtt{tt}\;M\;N \\ &\rightarrow M. \end{aligned}$$

The reduction is standard, even though the term $\mathtt{cond}\,(\mathtt{zero?}\,0)\;M\;N$ contracts after its subterm $(\mathtt{zero?}\,0)$, because it is the contraction of $(\mathtt{zero?}\,0)$ that turns the `cond` term into a redex.

There is a natural way of testing whether or not a reduction is outside-in: first, identify "outermost" subterms that contract; each of these identifies subterms that must reduce before the outer subterm itself contracts. By iterating the process, we can identify a subterm or subterms that must reduce before any others, if the reduction is to be outside-in. This idea is the basis of our definition of standard reductions.

For each term in a reduction, we identify a *principal redex*, and call a reduction standard if the redex contracted at each step is the principal redex. For the pure λ-calculus, the principal redex for some M_i will simply be the leftmost redex of M_i contracted in the reduction.

For systems with constants, we must allow reductions to take place in the critical arguments of some δ-terms. To find the principal redex, then, we start by considering the leftmost contracted subterm; if it is a δ-term, we then consider critical arguments in which contractions take place, etc. Eventually, consideration of these *preprincipal* subterms leads to the principal redex.

Definition 52 Let M_i be a term in a reduction path $\boldsymbol{\sigma}$,

$$\boldsymbol{\sigma} : M_1 \stackrel{\Delta_1}{\rightarrow} M_2 \stackrel{\Delta_2}{\rightarrow} M_3 \stackrel{\Delta_3}{\rightarrow} \cdots.$$

A contracted subterm Δ of M_i is *preprincipal in* $\boldsymbol{\sigma}$ if

(i) Δ is the leftmost subterm of M_i contracted in $\boldsymbol{\sigma}$; or

(ii) there is a subterm Δ' of M_i such that:

- Δ' is θ-contracted in $\boldsymbol{\sigma}$;
- Δ' is of the form $\delta_\theta\langle \vec{A}, \vec{B}\rangle$, where the leftmost active critical argument, A_k, is of the form $\Delta\vec{N}$; and
- Δ' is preprincipal in $\boldsymbol{\sigma}$.

We write $\mathrm{pp}_\sigma(\Delta)$ if Δ is preprincipal in $\boldsymbol{\sigma}$.

This next lemma is essential in showing an important property of the preprincipal subterms: they are linearly ordered by $\subset$ (see the following note):

Lemma 53 *Let M_i be a term in a reduction path $\boldsymbol{\sigma}$,*

$$\boldsymbol{\sigma} : M_1 \stackrel{\Delta_1}{\to} M_2 \stackrel{\Delta_2}{\to} M_3 \stackrel{\Delta_3}{\to} \cdots,$$

and let Δ be a preprincipal subterm of M_i. If $\Delta \not\equiv \Delta_i$, then Δ has a unique, preprincipal descendant $\Delta' \subset M_{i+1}$.

Proof: By induction on how $\mathrm{pp}_\sigma(\Delta)$.

(i) $\mathrm{pp}_\sigma(\Delta)$ because Δ is the leftmost contracted subterm of M_i. Then clearly Δ has some unique descendant Δ' in M_{i+1}. Furthermore Δ' is the leftmost contracted subterm of M_{i+1}, as the contraction of Δ_i can only introduce terms to the right of Δ'. Thus $\mathrm{pp}_\sigma(\Delta')$.

(ii) $\mathrm{pp}_\sigma(\Delta)$ because M_i contains a preprincipal, θ-contracted subterm, $\delta_\theta\langle \vec{A}, \vec{B}\rangle$, whose leftmost active critical argument, A_k, is of the form $\Delta\vec{N}$.

Now $\Delta_i \not\equiv \delta_\theta\langle \vec{A}, \vec{B}\rangle$, else by Note 35(ii), Δ would have no descendant in M_{i+1}, contradicting the fact that it is contracted in $\boldsymbol{\sigma}$.

So by induction, $\delta_\theta\langle \vec{A}, \vec{B}\rangle$ has a unique, preprincipal descendant, which must be of the form $\delta_\theta\langle \vec{A}', \vec{B}'\rangle$. But then $A'_k \equiv \Delta'\vec{N}'$, where Δ' is the unique descendant of Δ, and furthermore $\mathrm{pp}_\sigma(\Delta')$.

■

Note 54

(i) By Lemma 53, every preprincipal subterm contracts exactly once in $\boldsymbol{\sigma}$. Thus the θ and A_k of Definition 52(ii) are unique.

(ii) By (i), we conclude that if Δ_1 and Δ_2 are distinct, preprincipal subterms of M_i, then either $\Delta_1 \subset \Delta_2$ or $\Delta_2 \subset \Delta_1$.

Definition 55 Suppose $\boldsymbol{\sigma}$ is a reduction path,

$$\boldsymbol{\sigma} : M_1 \stackrel{\Delta_1}{\to} M_2 \stackrel{\Delta_2}{\to} M_3 \stackrel{\Delta_3}{\to} \cdots.$$

(i) We define the *principal redex* $\mathrm{pr}_\sigma(M_i)$ to be the innermost preprincipal subterm of M_i. By Note 54(ii), this is well defined.

(ii) We say $\boldsymbol{\sigma}$ is a *standard reduction* if for all i, $\Delta_i \equiv \mathrm{pr}_\sigma(M_i)$.

The following theorem is the main result of this appendix.

Theorem 56 (Standardization) *If $M \twoheadrightarrow N$ is a finite reduction in a PCF-like rewrite system, then there is a standard reduction from M to N.*

C.6 Path-reduction

This section gives our proof of Standardization. It is based on a proof in [23] for the pure λ-calculus, which introduced a sort of meta-reduction: a reduction relation on reduction paths. This *path-reduction* rewrites non-standard reductions into "more standard" reductions. The following results motivate the definition of path-reduction.

Lemma 57 *Let $\boldsymbol{\sigma}$ be a reduction path,*

$$\boldsymbol{\sigma} : M_1 \stackrel{\Delta_1}{\to} M_2 \stackrel{\Delta_2}{\to} M_3 \stackrel{\Delta_3}{\to} \cdots,$$

and let $\Delta \equiv \mathrm{pr}_\sigma(M_i)$. If $\Delta_i \not\equiv \Delta$, then Δ has a unique descendant $\Delta' \subset M_{i+1}$, and $\Delta' \equiv \mathrm{pr}_\sigma(M_{i+1})$.

Proof: Lemma 53 proves uniqueness. To show $\Delta' \equiv \mathrm{pr}_\sigma(M_{i+1})$, by the definition of pr_σ and Lemma 53 it suffices to note the following: if $\Delta_1 \subset \Delta_2 \subset M$ have unique descendants $\Delta_1', \Delta_2' \subset M'$, where $M \to M'$, then $\Delta_1' \subset \Delta_2'$. ■

Corollary 58 *Suppose $\boldsymbol{\sigma}$ is a reduction path,*

$$\boldsymbol{\sigma} : M_1 \stackrel{\Delta_1}{\to} M_2 \stackrel{\Delta_2}{\to} \cdots \stackrel{\Delta_{n-1}}{\to} M_n.$$

Then $\boldsymbol{\sigma}$ is standard iff there is no j such that Δ_j is the descendant of $\mathrm{pr}_\sigma(M_{j-1})$.

The corollary suggests a possible way to transform a non-standard reduction into a standard reduction: successively "swap" the contraction of a principal redex with the contraction of a non-principal redex at the previous step. If we reach a reduction in which each principal redex contracts as soon as it becomes principal, we will have found a standard reduction.

Definition 59 Suppose $\boldsymbol{\sigma}$ is a non-standard reduction, that is, there is some j such that

$$\boldsymbol{\sigma} : \cdots \to M_{j-1} \overset{\Delta_{j-1}}{\to} M_j \overset{\Delta_j}{\to} M_{j+1} \to \cdots$$

where Δ_j is the descendant of $\Delta'_j \equiv \mathrm{pr}_\sigma(M_{j-1})$. The subpath

$$M_{j-1} \overset{\Delta_{j-1}}{\to} M_j \overset{\Delta_j}{\to} M_{j+1}$$

is called the *path-redex at step* j. Note that Δ'_j and Δ_{j-1} do not overlap, and furthermore, by Lemma 57, Δ_j is the unique descendant of Δ'_j. Therefore by Lemma 41, we can find a sequence

$$M_{j-1} \overset{\Delta'_j}{\to} M'_j \overset{\Delta'_{j-1}}{\to} \cdots \overset{\Delta'_{j-1}}{\to} M_{j+1},$$

where the Δ'_{j-1} are the descendants of Δ_{j-1}. Such a sequence is call a *path-contractum*. Finally, we define *path-reduction*: $\boldsymbol{\sigma} \underset{\text{path}}{\overset{j}{\to}} \boldsymbol{\sigma'}$ if $\boldsymbol{\sigma'}$ is obtained from $\boldsymbol{\sigma}$ by replacing the path-redex at step j by a corresponding path-contractum. We will drop the index j when convenient.

Clearly, path-reduction preserves initial and final terms, and any path-reduction normal form is a standard reduction. Moreover, the next two lemmas show that path-reduction is strongly normalizing.

Lemma 60 *Suppose* $\boldsymbol{\sigma} \underset{\text{path}}{\overset{j}{\to}} \boldsymbol{\sigma'}$, *where*

$$\begin{aligned} \boldsymbol{\sigma} &: M_1 \to \cdots \to M_{j-1} \overset{\Delta_{j-1}}{\to} M_j \overset{\Delta_j}{\to} M_{j+1} \to \cdots, \\ \boldsymbol{\sigma'} &: M_1 \to \cdots \to M_{j-1} \overset{\Delta'_j}{\to} M'_j \overset{\Delta'_{j-1}}{\to} \cdots \overset{\Delta'_{j-1}}{\to} M_{j+1} \to \cdots. \end{aligned}$$

Then for $i \neq j$, *the following hold:*

(i) *If* $\Delta \subset M_i$ *is not contracted in* $\boldsymbol{\sigma}$, *then it is not contracted in* $\boldsymbol{\sigma'}$.

(ii) *If* $\Delta \subset M_i$ *is contracted in* $\boldsymbol{\sigma}$ *and* $\mathrm{pp}_\sigma(\Delta)$, *then* Δ *is contracted in* $\boldsymbol{\sigma'}$.

(iii) *If* $\Delta \subset M_i$ *is preprincipal in* $\boldsymbol{\sigma}$, *then it is preprincipal in* $\boldsymbol{\sigma'}$.

(iv) $\mathrm{pr}_\sigma(M_i) \equiv \mathrm{pr}_{\sigma'}(M_i)$.

Proof:

(i) Just note that path-reduction only permutes the order of contraction of subterms; it does not introduce new contractions.

(ii) It is clear that if Δ contracts in $\boldsymbol{\sigma}$ and does not contract in $\boldsymbol{\sigma}'$, then Δ is either Δ_{j-1} or one of its ancestors. Thus we only need consider Δ_{j-1}.

If Δ_{j-1} does not contract in $\boldsymbol{\sigma}'$, then it must be contained in Δ'_j. But Δ'_j is the principal redex of M_{j-1}, that is, the innermost preprincipal subterm of M_{j-1}. So if Δ_{j-1} is not contracted in $\boldsymbol{\sigma}'$, it is not preprincipal in $\boldsymbol{\sigma}$.

(iii) We use induction on how $\mathrm{pp}_\sigma(\Delta)$.

1. $\mathrm{pp}_\sigma(\Delta)$ because Δ is the leftmost contracted subterm of M_i. By (ii), Δ is contracted in $\boldsymbol{\sigma}'$, and by (i), it is the leftmost contracted subterm of M_i in $\boldsymbol{\sigma}'$. Therefore $\mathrm{pp}_{\sigma'}(\Delta)$.
2. $\mathrm{pp}_\sigma(\Delta)$ because $\mathrm{pp}_\sigma(\delta_\theta\langle\vec{A},\vec{B}\rangle)$, and the leftmost active critical argument, A_k, is of the form $\Delta\vec{N}$. By induction, $\mathrm{pp}_{\sigma'}(\delta_\theta\langle\vec{A},\vec{B}\rangle)$, and by (ii), Δ is contracted in $\boldsymbol{\sigma}'$. So A_k is active in $\boldsymbol{\sigma}'$, and by (i), it is the leftmost active critical argument. Therefore $\mathrm{pp}_{\sigma'}(\Delta)$.

(iv) This follows from (i), (iii), and the definition of pr_σ.

■

Lemma 61 *If $\boldsymbol{\sigma}$ is a finite reduction, then there is no infinite path-reduction starting from $\boldsymbol{\sigma}$.*

Proof: Consider a path-reduction

$$\boldsymbol{\sigma} \equiv \boldsymbol{\sigma}_1 \underset{\text{path}}{\to} \boldsymbol{\sigma}_2 \underset{\text{path}}{\to} \boldsymbol{\sigma}_3 \underset{\text{path}}{\to} \cdots.$$

It is not hard to see that the reduction could have been carried out in the labelled system; that is, if $\boldsymbol{\sigma}'_i$ is a labelled reduction such that $|\boldsymbol{\sigma}'_i| \equiv \boldsymbol{\sigma}_i$, and $\boldsymbol{\sigma}_i \underset{\text{path}}{\overset{j}{\to}} \boldsymbol{\sigma}_{i+1}$, then there is a labelled reduction $\boldsymbol{\sigma}'_{i+1}$ such that $|\boldsymbol{\sigma}'_{i+1}| \equiv \boldsymbol{\sigma}_{i+1}$, and $\boldsymbol{\sigma}'_i \underset{\text{path}}{\overset{j}{\to}} \boldsymbol{\sigma}'_{i+1}$. Thus we can find labelled reductions $\boldsymbol{\sigma}'_1, \boldsymbol{\sigma}'_2, \boldsymbol{\sigma}'_3, \ldots$ such that $|\boldsymbol{\sigma}'_i| \equiv \boldsymbol{\sigma}_i$, and

$$\boldsymbol{\sigma}'_1 \underset{\text{path}}{\to} \boldsymbol{\sigma}'_2 \underset{\text{path}}{\to} \boldsymbol{\sigma}'_3 \underset{\text{path}}{\to} \cdots.$$

And because labelled reduction is strongly normalizing, and each $\boldsymbol{\sigma}'_i$ begins with the same $\mathcal{L}_{\mathbb{N}}$ term, each $\boldsymbol{\sigma}_i$ is finite.

Furthermore, the path-reduction can be thought of as constructing a tree of terms, with each path from root to leaf corresponding to a reduction $\boldsymbol{\sigma}_i$. Each contracted path-redex introduces a branching in the tree. For example, if $\boldsymbol{\sigma}_i \underset{\text{path}}{\overset{j}{\to}} \boldsymbol{\sigma}_{i+1}$, then the root-to-leaf path corresponding to $\boldsymbol{\sigma}_{i+1}$ is obtained by branching off of the root-to-leaf

path of $\boldsymbol{\sigma}_i$ at depth $j-1$. The situation is depicted in the following figure, where the root of the tree is displayed at the left and the leaves are displayed at the right:

$$\begin{array}{l} M_1 \longrightarrow \cdots \longrightarrow M_{j-1} \stackrel{\Delta_{j-1}}{\longrightarrow} M_j \stackrel{\Delta_j}{\longrightarrow} M_{j+1} \longrightarrow \cdots \longrightarrow M_n : \boldsymbol{\sigma}_i \\ \qquad\qquad\qquad\quad \downarrow \Delta'_j \\ \qquad\qquad\qquad\quad M'_j \stackrel{\Delta'_{j-1}}{\longrightarrow} \cdots \stackrel{\Delta'_{j-1}}{\longrightarrow} M_{j+1} \longrightarrow \cdots \longrightarrow M_n : \boldsymbol{\sigma}_{i+1} \end{array}$$

By Lemma 60(iv), the tree is a binary tree, and we have just seen that there is no infinite path from the root. Then by König's Lemma, the tree is finite, so the number of different reductions given by the tree must be finite. ■

Proof of Theorem 56 (Standardization): If $\boldsymbol{\sigma} : M \twoheadrightarrow N$ is a finite reduction in a PCF-like system, we can obtain a standard reduction from M to N just by finding a path-reduction normal form of $\boldsymbol{\sigma}$. ■

Note that we have not shown that path-reduction normal forms are unique: that is, if

$$\begin{array}{rl} & \boldsymbol{\sigma} \underset{\text{path}}{\rightarrow} \cdots \underset{\text{path}}{\rightarrow} \boldsymbol{\sigma}_1, \\ \text{and} & \boldsymbol{\sigma} \underset{\text{path}}{\rightarrow} \cdots \underset{\text{path}}{\rightarrow} \boldsymbol{\sigma}_2, \end{array}$$

where $\boldsymbol{\sigma}_1$ and $\boldsymbol{\sigma}_2$ are normal forms, we are not guaranteed that $\boldsymbol{\sigma}_1 \equiv \boldsymbol{\sigma}_2$. We expect that the property holds, but haven't tried to verify that it does, since it is not needed to prove Standardization.