

Full Abstraction for Functional Languages with Control

James Laird
LFCS, Department of Computer Science
University of Edinburgh

Abstract

This paper considers the consequences of relaxing the bracketing condition on ‘dialogue games’, showing that this leads to a category of games which can be ‘factorized’ into a well-bracketed substructure, and a set of classically typed morphisms. These are shown to be sound denotations for control operators, allowing the factorization to be used to extend the definability result for PCF to one for PCF with control operators at atomic types. Thus we define a fully abstract and effectively presentable model of a functional language with non-local control as part of a modular approach to modelling non-functional features using games.

1. Introduction

The simultaneous elegance and inefficiency of pure functional languages is manifest in the fact that whilst many practical languages are described as ‘functional’, they contain so-called ‘imperative features’. The object of this approach is to exploit the nice structural properties of functional programs, whilst avoiding, at least, the obvious inefficiencies which they impose. However, the addition of imperative features is a radical alteration to the previously pure functional idiom, as an increase in expressive power testifies. If it is to be possible to describe such a language as ‘*essentially* functional with imperative features’, it will be necessary to have a clear perspective on the extent of its non-functional behaviour. One way to do so is to give a semantics in which the functional and non-functional elements of the model are clearly and naturally delineated.

The object of this paper is to study the effects of adding operators which give access to the flow of control to call-by-name functional languages using games semantics. Control operators allow programs to use intensional information about their arguments[4], and PCF augmented with such operators is (arguably) a basis for a notion of sequen-

tial computation, as well as a toy ‘functional language with control’. The extent to which control operators alter the sorts of programs which can be written is underlined by their types, which are derivable in classical but not intuitionistic logic. This extension to the Curry-Howard correspondence between proofs and programs has been formalised elsewhere [5], [14]; here it suggests that the same basic games model will give a semantics to PCF with a simple escape mechanism, (which can be extended to the higher order control operators) and to proofs of classical logic. The fact that features of classical reasoning can be identified in games, and are the same elements which model jumps in the flow of control also gives a way of understanding the computational content of classical reasoning: a similar distinction is used in [6] to extend a games semantics for intuitionistic arithmetic to the classical case by allowing backtracking.

The prototypical language PCF has provided a framework in which to study the essential features of functional languages. So a reasonable place to examine the functional/non-functional boundary is in a version of PCF incremented with a simple non-functional feature. The power of game semantics to describe the intensional features of functional languages and calculi has already been established in providing a fully abstract semantics for PCF. Full abstraction provides a formal test of the level of detail at which a semantics describes the dynamics of a language. But the definition of a semantics which passes this formal test is necessarily preceded by a detailed intuitive correspondence between aspects of functional behaviour and the conditions imposed upon games to model it. This correspondence can be exploited when adding non-functional features to PCF by relaxing precisely the restrictions which exclude this behaviour to get a fully abstract semantics. It can also be formalized, as a ‘factorization’ of the denotations in the model as the composition of a morphism from the PCF model and one from a restricted set with essentially non-functional behaviour. Provided these are definable (they are the denotations of appropriately typed control operators), a full abstraction theorem for the extended model is a simple consequence of the PCF definability of its functional sub-

structure.

Abramsky has suggested (by analogy with Barendeg’s λ -cube), that this intensional hierarchy can be seen as a ‘Games’ cube, with a category of dialogue games at each vertex, partially ordered according to the presence or absence of the three constraints which jointly give purely functional behaviour: innocence, well-bracketing and determinacy. A vertex of this cube, at which innocence is lost, allowing Idealized Algol (‘PCF with State’) to be modelled, is studied in [1], and the fact that a similar methodology is used there to prove full abstraction (factorization through a morphism which records a history of the play), suggests that this is a fruitful paradigm for intensional semantics.

However, this is not the first fully abstract semantics to be given to PCF with control operators. Before the emergence of the fully abstract games models of PCF, Cartwright, Curien and Felleisen [4] had used the fact that such operators allow one to distinguish between, for instance, $+_l$ and $+_r$, to achieve this result using Sequential Algorithms. These represent morphisms intensionally as trees, with obvious parallels with various categories of games [7]. However, it was not apparent how the Sequential Algorithms model could be restricted to reflect pure functional behaviour, and hence yield a fully abstract model of PCF itself. Dialogue (and AJM) games represent a move towards a more detailed intensional description of computation, by incorporating a notion corresponding to stack discipline, (the bracketing condition), which can be weakened, allowing operators which pop the stack to be modelled. The possibility of mixing and matching the conditions corresponding to the non-functional features - not available in the sequential algorithms model, supports this view, and the claim that dropping the bracketing condition gives more than ‘just another fully abstract model for sPCF’.

Extensions of dialogue games to model the $\lambda\mu$ -calculus have been considered elsewhere: Ong [13] adds a notion of state, whilst Herbelin [9] suggests the adaptation to the Hyland-Ong framework which is used here: i.e. the relaxation of the bracketing condition, although in the context of a more syntax-oriented study, in contrast to the present work which aims to place these models in context.

2. Weakly bracketed games

The framework proposed here is a variant of the ‘dialogue games and innocent strategies’ setting, which gives fully abstract models of the simply-typed λ -calculus and PCF; namely, ‘weakly-well-bracketed’ games and innocent strategies. Hence it is easily accommodated within a setting which has been well described elsewhere, particularly

in [12], in which most of these definitions are more comprehensively studied.

A game A is a specification of a set of possible plays, i.e. sequences of moves M_A , which are alternately made by Opponent and Player, and can be either questions or answers. Thus the specification consists of:

- a labelling function $\lambda_A: M_A \rightarrow \{P, O\} \times \{Q, A\}$, which indicates the nature of each move (who made it, and whether it is a question or an answer). λ_A^{PO} and λ_A^{QA} denote the composition of λ_A with the first and second projections, respectively.
- a prefix-closed set of legal sequences of moves, or ‘plays’: $P_A \subseteq M_A^*$.
- a justification relation; $\vdash_A \subseteq M_A \times M_A$ associating to each move a set of ‘justifying’ moves. (‘ a justifies b ’ is denoted $a \vdash b$.) If this set is empty, then the move is initial (this fact is denoted $\vdash a$; if it is non-empty then one of the justifying moves must appear before it in any play.

Justification also satisfies the following properties:

- If $a \vdash b$ then $\lambda^{PO}(a) = \overline{\lambda^{PO}(b)}$: i.e. Opponent’s moves are justified by Player, and vice-versa. (Where $\overline{P} = O$, and vice-versa.)
- No move is justified by an answer:
If $a \vdash b$ then $\lambda^{QA}(a) = Q$.
- Initial moves are Opponent questions:
If $\vdash a$ then $\lambda(a) = \langle O, Q \rangle$.

For any play $s \in P_A$ there is a function ($\phi_s: |s| \rightarrow N$) such that $\phi_s(n) \leq n$, giving for each non-initial n th move s_n of s , the index $\phi_s(n)$ of the unique move justifying s_n . As each answer is non-initial, and justified by a unique question, it is reasonable to say that it answers that question.

Plays are subject to these general conditions:

P1 Moves are made alternately by Player and Opponent:

$$sab \in P_A \Rightarrow \lambda_A^{QA}(a) = \overline{\lambda_A^{QA}(b)}.$$

P2 Plays satisfy the ‘visibility condition’: for every non-initial move a , if $sa \in P_A$ then the unique move in s justifying a is in the view of s (defined below).

Definition 2.1 (Player and Opponent views.)

Notions of Player and Opponent view $\lceil s \rceil$ and $\lfloor s \rfloor$ can be defined on any justified sequence s of moves, independently of the question/answer distinction, and,

in particular, any bracketing condition [12]:

$$\begin{aligned} \lceil \epsilon \rceil &= \epsilon, \\ \lceil sm \rceil &= \lceil s \rceil m, \quad \text{if } m \text{ is a P-move.} \\ \lceil sm \rceil &= m, \quad \text{if } \vdash m. \\ \lceil smtn \rceil &= \lceil s \rceil mn, \text{ if } n \text{ is an O-move justified by } m. \end{aligned}$$

$$\begin{aligned} \lfloor \epsilon \rfloor &= \epsilon, \\ \lfloor sm \rfloor &= \lfloor s \rfloor m, \quad \text{if } m \text{ is an O-move.} \\ \lfloor smtn \rfloor &= \lfloor s \rfloor mn, \text{ if } n \text{ is a P-move justified by } m. \end{aligned}$$

In the absence of an explicit bracketing condition, imposing the visibility condition entails that plays satisfy a weaker condition (provided that that answers cannot justify any moves), because all moves between a question and answer disappear from view.

Lemma 2.2 *The Player view of any play s in which Player is to move is well-bracketed.*

Proof: - by induction on the length of s

The view of a single move is trivially well-bracketed. Suppose $s = s'mtn$, where m justifies n (as Player is to move, n is an O-move) then $\lceil smtn \rceil = \lceil s \rceil mn$. By inductive hypothesis, $\lceil s \rceil$ is well-bracketed, and as m justifies another move (n), it cannot be an answer move. However the only move which n can answer is its justifying move m , so $\lceil smtn \rceil$ is also well-bracketed.

Dually, unanswered questions disappear from Opponent's view as soon as a question asked earlier is answered by either participant. Consequently the following 'Weak Bracketing condition' holds:

- Only open questions may be answered: a question is open if it is unanswered, and no questions asked before it have been answered since it was asked:

$$\begin{aligned} smtn \in P_A \wedge \lambda_A^{QA}(m) = \lambda_A^{QA}(n) = A \Rightarrow \\ \phi_{smtn}(n) \leq \phi_{smtn}(m). \end{aligned}$$

This contrasts with the 'well-bracketing condition' applied in [2], [10], which requires that only the most recently asked open question can be answered.

In weakly bracketed games, answering a question closes all of the more recently asked questions as well. This is in accord with the intuition that we can keep track of dialogues between Player and Opponent (and hence compute the result of a composition of strategies) by keeping unanswered questions in a 'stack'. Moves can be characterized either as the pushing of a new question onto the top of the stack, or as the popping of a question from the stack by answering

it. Imposing the original bracketing condition corresponds to rigidly maintaining stack discipline; only the question at the top of the stack can be popped. Relaxation to the weak bracketing condition allows for the stack to be popped to an earlier point, causing every question pushed onto the stack since that point to be lost.

3. Innocent strategies

A (deterministic) strategy over a game A is a specification of Player's responses to (some of) the plays of A in which she is to move, and is therefore generally expressed as an even-prefix-closed subset $\sigma \subseteq P_A^{even}$, such that $sa, sb \in \sigma \Rightarrow a = b$.

Thus a strategy is determined by a partial function from odd length sequences of moves in P_A (situations with player to move) to legal moves. It is innocent if its moves are only determined by the view of the game so far: i.e. it is a prefix closed subset $\sigma \subseteq P_A^{even}$, such that

$$sa, tb \in \sigma \wedge \lceil s \rceil = \lceil t \rceil \Rightarrow a = b.$$

Hence it is determined by a partial function from Player views to legal moves. The representations of a strategy as a tree and as an 'innocent function' will frequently be interchanged.

In order to characterise the substructure of this games model which is restricted to purely functional behaviour (i.e. 'the well-bracketed part'), it is necessary to describe the strategies which correspond to functions without control operators. These will be the strategies which do obey the last-asked-first-answer condition; the question is, how do these strategies correspond to strategies on well-bracketed games, given that the latter only have responses to well-bracketed plays? As views are well-bracketed, there is a simple answer.

Definition 3.1 *A well-bracketed strategy is an innocent strategy which always answers the most recently asked open question, (which is always visible).*

A well-bracketed strategy may contain plays which are not well-bracketed, but the lemma below shows that these are just the plays which are forced to be there by innocence.

Lemma 3.2 *Every well-bracketed strategy corresponds to a unique strategy on well-bracketed games with the same innocent function.*

Proof: Suppose σ is a well-bracketed strategy. Then it is necessary to show that for every play $sa \in \sigma$, there is a well-bracketed play $s'a \in \sigma$, such that $\ulcorner s \urcorner = \ulcorner s' \urcorner$. (From this fact it follows immediately that the innocent function representation of σ is an innocent function in the weakly bracketed setting.) Proof is by induction on the length of sa : - the base case is trivial, so we prove the inductive step.

As s ends with an opponent move, $\ulcorner s \urcorner = \ulcorner t \urcorner bc$, for some prefix t of s , and moves b and c such that b justifies c .

But by inductive hypothesis, there is a well bracketed play $t'b \in \sigma$ such that $\ulcorner t \urcorner = \ulcorner t' \urcorner$. Then $\ulcorner s \urcorner = \ulcorner t'bc \urcorner$, and $t'bc$ is well-bracketed, as move b is made by a well-bracketed strategy, and if c is an answer, it is justified by b anyway.

The functor which embeds the category of well-bracketed games in the category of weakly-bracketed games (described below) by taking the ‘innocent closure’ of each strategy is therefore faithful.

4. A symmetric monoidal closed category of games

Like well-bracketed games and innocent strategies, the non-well-bracketed games form an affine autonomous category \mathcal{G} , in which the objects are games, and morphisms are strategies for the exponential.

Given games A and B :

the game $A \multimap B$ is defined:

- $M_{A \multimap B} = M_A + M_B$.
- $\lambda_{A \multimap B} = [\overline{\lambda_A}, \lambda_B]$
- $P_{A \multimap B} = \{s \in M_{A \multimap B}^* \mid s \upharpoonright M_A \in P_A, s \upharpoonright M_B \in P_B\}$.
- $\forall b \in M_B, m \vdash_{A \multimap B} b$ iff $m \vdash_B b$
for a initial in A , $m \vdash_{A \multimap B} a$ iff m is initial in B .
 $\forall a(\text{non-initial}) \in M_A$ $m \vdash_{A \multimap B} a$ iff $m \vdash_A a$.

The game $A \otimes B$ is defined as for $A \multimap B$, except that $\lambda_{A \otimes B} = [\lambda_A, \lambda_B]$. It has a unit, the game with no moves (which is also a terminal object as this is an affine category). For any strategy $\sigma:A$, the strategy $\ulcorner \sigma \urcorner : 1 \multimap A$, the ‘name of sigma’ has the same graph as σ , allowing the systematic confusion between them adopted here.

The game $!A$ is defined:

- $M_{!A} = M_A$
- $\lambda_{!A} = \lambda_A$

- $P_{!A} = \{s \in M_{!A}^* \mid \text{if } m \in M_A \text{ is initial then the subsequence of } s \text{ hereditarily justified by } m \text{ is in } P_A\}$

$$\vdash_{!A} = \vdash_A$$

For each object A , the *identity* morphism $id_A : A \rightarrow A$ is the copycat strategy, which copies Opponent moves on either side of the arrow as Player moves on the other side.

Composition of morphisms is interaction of strategies:

Given $\sigma : A \rightarrow B$ and $\tau : B \rightarrow C$

$\sigma; \tau = (\sigma \parallel \tau) / B = \{s \upharpoonright A, C : s \in \sigma \parallel \tau\}$, where

$\sigma \parallel \tau = \{s \in (M_A + M_B + M_C)^* \mid s \upharpoonright A, B \in \sigma \wedge s \upharpoonright B, C \in \tau\}$.

It is necessary to check that if σ and τ are innocent, then $\sigma; \tau$ is too. The proof of this fact for well-bracketed games appears in [10], but carries through directly to the weakly-bracketed case. If both σ and τ are well-bracketed, then so is $\sigma \parallel \tau$, and hence also $\sigma; \tau$.

Now define the cartesian closed category \mathcal{C} of weakly bracketed games to be the the co-Kleisli construction of the comonad $!$ on the autonomous category \mathcal{G} . The canonical morphisms for a CCC are then various forms of copycat strategy.

Define an *atomic* game α to be the game consisting of a single (Opponent) question $[\alpha]$ and a set of answers, $\{!a : a \in \alpha\}$. These correspond to the flat domains which interpret ground types in PCF. If \mathcal{S} is any collection of such atomic games, $\mathcal{C}_{\mathcal{S}}$ denotes the CCC of weakly bracketed games freely generated from \mathcal{S} .

For each object $A \in \mathcal{C}_{\mathcal{S}}$, let $At(A)$ denote the set of atomic subgames of A .

Example 4.1 (Peirce’s Law)

$((\alpha \rightarrow \beta) \rightarrow \alpha) \rightarrow \alpha$ (Peirce’s law), is the simplest type which can be given to control operators but not (generally) to functions; it is also the simplest example of a classical tautology which is not provable in minimal intuitionistic logic. This translates to the games semantics framework: this is the simplest construction from atomic games α and β upon which there is a weakly bracketed strategy, independent of α and β which is total (in the obvious sense), but no total well-bracketed strategy. Moreover, this strategy is (in a sense made formal by the factorization theorem) a canonical example of a non-well-bracketed strategy, and will be the denotation of the *call/cc*-like control operators which will be added to PCF.

Definition 4.2 $Peirce_{\beta}^{\alpha}$

$$(((\alpha \rightarrow \beta) \rightarrow \alpha) \rightarrow \alpha) \rightarrow A$$




Figure 2. Elimination of a violation of the bracketing condition.

- If τ is compact (its innocent function has a finite graph) then so is τ' .

Proof:

Define τ' as described by Figure 2. i.e. play as τ , but jump into the premise $((\alpha \rightarrow \beta) \rightarrow \alpha) \rightarrow \alpha$ as soon as an α question is asked by Opponent, and copy it as the initial question. Then jump back and play as τ until a point where τ would close a β -question prematurely by answering the α -question. Instead τ' returns to the premise, and copies the β -question, forcing Opponent to answer a question first, providing either a β -answer to copy, or closing all questions back to the original α -question, allowing it to be answered. (The translation from τ to τ' can be described in a more formal but less perspicuous way by defining the innocent function f_τ .)

First note that τ' is a well defined innocent strategy, i.e. τ' is prefix closed, as its restrictions to the premise $((\alpha \rightarrow \beta) \rightarrow \alpha) \rightarrow \alpha$ and to the conclusion B both are. Plays according to τ' obey the visibility condition, as if τ' is able to close an α question then the play succeeding in in the premise is visible to τ' , and elsewhere τ' plays as τ . It remains to verify the following claims:

- $WB_{\alpha,\beta}(\tau')$: τ' will not give an α -answer when the last-asked question is of type β .
- $Peirce_\beta^\alpha; \tau' = \tau$: if Opponent plays according to $Peirce_\beta^\alpha$ in the premise $((\alpha \rightarrow \beta) \rightarrow \alpha) \rightarrow \alpha$, then for any $s \in \tau'$, $s \upharpoonright B \in \tau$ is played according to the same Opponent strategy in B .
- $\forall \gamma, \delta \in At(B) : WB_{\gamma,\delta}(\tau) \Rightarrow WB_{\gamma,\delta}(\tau')$: τ' only changes the order in which answers are given in the case where the last asked question is of type β , and the answer closing it is of type α .
- If τ is compact then so is τ' : τ' is defined on the same views as τ , except that each visible Opponent α -question is now followed by two moves in

the premise $((\alpha \rightarrow \beta) \rightarrow \alpha) \rightarrow \alpha$, and each Player answer of type α may be preceded by up to four moves in the premise.

Corollary 5.4 *If $\sigma : A$ is compact then so is σ' .*

Proof: $At(A) \times At(A)$ is finite, therefore the lemma need only be applied finitely many times, each yielding a compact strategy, to factorize σ .

6. Augmenting PCF with control operators

Now the Factorization Theorem can be applied to prove that the games and innocent well-bracketed strategies model for call-by-name PCF extends to a fully-abstract model of PCF with control operators $call/cc_\beta^\alpha$ of type $((\alpha \rightarrow \beta) \rightarrow \alpha) \rightarrow \alpha$, for $\alpha, \beta \in \{o, \iota\}$ by dropping the bracketing condition. However, to maintain a concise account, I will consider a version of PCF with only the ground-type ι , allowing $call/cc_\iota^\iota$ to be abbreviated to $call/cc$. (This is described as ‘call with current downward continuation’ in [4], although a less expressive ‘linear’ version is defined there.)

Hyland and Ong proved that the well-bracketed model is fully abstract with respect to a call-by-name operational semantics for PCF in terms of a Martin-Löf style evaluation relation [10]: this can be naturally extended to include $call/cc$ using the notion of evaluation context [8].

Definition 6.1 (Evaluation contexts.)

$$E[] ::= [] \mid E[]M \mid IFOE[]MN \mid call/cc \lambda f. E[] \mid const E[]$$

where $const$ is any constant of PCF.

Evaluation contexts serve to pick out subterms which must be evaluated in any reduction sequence for the whole term, a fact expressed by the Unique Evaluation Contexts lemma [4]:

Any application of PCF + $call/cc$ can be uniquely represented as $E[t]$, where t is a variable or Ω , and $E[]$ is an evaluation context.

Assuming \Downarrow is the evaluation relation for PCF, extend it with the following reductions (and allowing evaluation inside a $call/cc$), where $E[]$ ranges over applicative contexts:

$$\frac{t \Downarrow n}{call/cc \lambda f. t \Downarrow n}$$

$$\frac{\text{call/cc } \lambda f.t \Downarrow n}{\text{call/cc } \lambda f.E[ft] \Downarrow n}$$

The denotational semantics for PCF given in [10] can be extended to PCF + *call/cc* using the faithful embedding of the category of well-bracketed games and innocent strategies into the category of weakly bracketed games and innocent strategies given by mapping strategies to their well-bracketed counterparts, and letting *call/cc* denote $Peirce_{\iota}^t$, which is abbreviated to *Peirce*. ($\llbracket \iota \rrbracket$ is an atomic game as required, with a single question $\llbracket N \rrbracket$ and the answers $\llbracket n \rrbracket | n \in N$, whilst $\bar{n} : \iota$ denotes the strategy $\{\llbracket N \rrbracket_n\}$).

Informally, it is clear that *call/cc* should denote the strategy *Peirce*, as *call/cc* $\lambda f.M$ returns the first argument to *f* that is encountered whilst evaluating *M*, or if $\lambda f.M$ is non-strict, it returns the value of *M*.

Theorem 6.2 (Adequacy) : *The proposed denotational semantics is computationally adequate with respect to the operational semantics; i.e. for all closed terms s of type ι , $s \Downarrow n$ if and only if $\llbracket s \rrbracket = \bar{n}$.*

Proof: \mathcal{C}_{ι} is a ‘standard model’ for PCF [10],[2] (just using the well-bracketed substructure). To extend this adequacy result to PCF + *call/cc*, it is sufficient to extend the reducibility style proof of the adequacy of standard models [15],[3] to the interpretation of *call/cc* as *Peirce* in \mathcal{C}_{ι} : i.e. to show, in the absence of the **Y** combinator, that $\text{call/cc } \lambda f.M \Downarrow m \Leftrightarrow \llbracket \lambda f.M \rrbracket; Peirce = \bar{m}$.

Lemma 6.3 (Evaluation contexts) : *Let $E[\]$ be any evaluation context such that*

$$\Gamma \vdash \lambda f.E[ft] : (\iota \rightarrow \iota) \rightarrow \iota.$$

The strategy $\llbracket \Gamma \vdash \lambda f.E[ft] \rrbracket : \llbracket \Gamma \rrbracket \rightarrow (\iota \rightarrow \iota) \rightarrow \iota$ responds to Opponent’s initial question by asking the initial question in $\iota \rightarrow \iota$.

Proof is by a structural induction over the definition of evaluation context.

Proof of Adequacy:

In the direction of soundness it is sufficient to note that:

- $\llbracket \text{call/cc } \lambda f.n \rrbracket = \llbracket \lambda f.n \rrbracket; Peirce = \llbracket n \rrbracket$
- $\llbracket \text{call/cc } \lambda f.E[ft] \rrbracket = \llbracket \lambda f.E[ft] \rrbracket; Peirce = \llbracket \text{call/cc } \lambda f.t \rrbracket$.

(using the lemma). To establish completeness, suppose that $\llbracket \text{call/cc } \lambda f.M \rrbracket = \bar{n}$, where *M* is in head-normal form. As it is of ground type it must be an application or a value. Applying the Unique Evaluation Contexts Lemma, $M = E[t]$, where $t = f$, or $t = n'$ or $t = \Omega$, and $E[\]$ is an evaluation context. By soundness, in the first two cases, $t \Downarrow n$, whilst the latter contradicts the assumption that $\llbracket \text{call/cc } \lambda f.M \rrbracket = \bar{n}$. The proof can be extended to PCF including the **Y** combinator, as described in [3], by replacing it in all terminating programs by a finite approximant.

Now the factorization of strategies, and the definability of well-bracketed strategies by PCF terms can be exploited to give an easy proof of:

Theorem 6.4 (Definability lemma): *Any compact innocent strategy $\sigma : A$ is the denotation of a term of PCF + *call/cc*.*

Proof:

By the factorization theorem, there is a well-bracketed compact strategy $\sigma' : (((\iota \rightarrow \iota) \rightarrow \iota) \rightarrow \iota) \rightarrow A$ such that $Peirce_{\iota, \iota}; \sigma' = \sigma$.

By the definability theorem for PCF [10], there is a term *s* of PCF denoting σ' , and thus the term $s (\lambda x.\text{call/cc } x)$ denotes σ .

To get a fully abstract model it is now necessary to quotient by the intrinsic preorder, \lesssim , as in [10], [2]. Note, however, that intrinsic equivalence is no longer extensional equivalence, as we can define the strategies $catch^k$, of [4] which distinguish between different sequentializations of the same function.

Definition 6.5

$$\sigma \lesssim_A \tau \iff \forall \rho : A \rightarrow N(\sigma; \rho = \bar{n}) \Rightarrow (\tau; \rho = \bar{n}).$$

$$\sigma \equiv_A \tau \iff \sigma \lesssim_A \tau \wedge \tau \lesssim_A \sigma.$$

The semantics of PCF + *call/cc* can be fully and faithfully translated to \mathcal{C}/\equiv , which is poset-enriched by \lesssim ([2]).

Define the relation of observational approximation, \sqsubseteq^{OBS} , between type-compatible terms of PCF + *call/cc* thus:

$$s \sqsubseteq^{OBS} t \iff C[s] \Downarrow v \Rightarrow C[t] \Downarrow v$$

where $C[\]$ ranges over all type compatible program contexts. Then a denotational model for PCF + *call/cc* is fully abstract if for all type-compatible pairs of terms *s*, *t*:

$$\llbracket s \rrbracket \lesssim \llbracket t \rrbracket \iff s \sqsubseteq^{OBS} t$$

Theorem 6.6 (Full abstraction:) \mathcal{C}/\lesssim is a fully abstract model for PCF + call/cc.

Proof: It is sufficient to prove soundness and completeness for *closed* terms, as given any terms s and t with free variables $x_1 : A_1, x_2 : A_2, \dots, x_n : A_n$, then

$\llbracket s \rrbracket \lesssim \llbracket t \rrbracket$ if and only if $\llbracket \lambda \bar{x}.s \rrbracket \lesssim \llbracket \lambda \bar{x}.t \rrbracket$. Whilst $s \sqsubseteq^{OBS} t$ if $\lambda \bar{x}.s \sqsubseteq^{OBS} \lambda \bar{x}.t$,

as if $C[\]$ is any type-compatible program context, then $C[s] \Downarrow v \iff D[\lambda \bar{x}.s] \Downarrow v$, where $D[\] = C[(\] \bar{x})$.

Soundness: if $s \not\sqsubseteq^{OBS} t$, then as the denotational semantics is adequate, there is some $n \in N$ and type-compatible context $C[\]$, such that $C[s] \Downarrow n$, but $C[t] \not\Downarrow n$.

Then if $\rho = \lambda x.C[x]$ (where x is free in $C[x]$): $\rho; \llbracket s \rrbracket = \bar{n}$, but $\rho; \llbracket t \rrbracket \neq \bar{n}$, and hence $\llbracket s \rrbracket \not\lesssim \llbracket t \rrbracket$.

Completeness: suppose $\llbracket s \rrbracket \not\lesssim \llbracket t \rrbracket$. Then there is a strategy ρ , and number n such that $\llbracket s \rrbracket; \rho = \bar{n}$ and $\llbracket t \rrbracket; \rho \neq \bar{n}$. As $\llbracket s \rrbracket; \rho$ returns a result in N , the uncovering of this play with respect to $\llbracket s \rrbracket$ and ρ is finite, consequently, ρ has a finite approximant ρ' such that $\llbracket s \rrbracket; \rho' = \bar{n}$ and $\llbracket t \rrbracket; \rho' \neq \bar{n}$.

By the definability lemma, there is a closed term r denoting ρ' , thus $\llbracket rs \rrbracket = \bar{n}$ and $\llbracket rt \rrbracket \neq \bar{n}$ so $(rs) \Downarrow \bar{n}$ and $(rt) \not\Downarrow \bar{n}$ by adequacy. So $s \not\sqsubseteq^{OBS} t$ as required.

7. Effective presentability

An interesting question concerns the effective presentability of the extensionally fully abstract domain \mathcal{C}/\equiv : The sequential algorithms model of sPCF does not require quotienting (the intensional preorder is just the subset relation on strategies) - consequently ‘the fully abstract model of sPCF’ is effectively presentable. This contrasts with the extensionally fully abstract model of PCF, which is known not to be effectively presentable, as finitary PCF is undecidable [11]. Adding control operators allows intensionally distinct versions of extensionally equal functions to be distinguished - one of the reasons for introducing them in sPCF [4]. Not all of the strategies in the games model can be distinguished (partly a consequence of the greater intensional precision with which it models the process of computation - there is no operational distinction between e.g. $\lambda x.If\ x = x\ then\ x\ else\ 0 : N \rightarrow N$ and $\lambda x.x$, although it is reasonable to allow them to represent different computations). However the intrinsic preorder for compact strategies is decidable, and thus the quotiented model is effectively presentable as required.

Essentially this is because in the absence of the bracketing condition, any intensional differences between two strategies can be exposed by supplying the same inputs and making a sharp exit as soon as they make different moves *provided* that this does not violate innocence.

Definition 7.1 An innocent play according to an (innocent) strategy σ is a finite play $s \in \sigma$ such that Opponent also plays innocently: i.e.

$\forall s', t' \sqsubseteq_{even} s, \sqcup s' \sqcup = \sqcup t' \sqcup \wedge s'a \sqsubseteq s \implies t'a \sqsubseteq s$.

Given $s \in \text{Innocent plays}(\sigma)$, $O - \text{strategy}(s)$ is the innocent function from Opponent views to moves which is manifested in s .

Proposition 7.2 For any strategies σ, τ in \mathcal{C}_l , $\sigma \lesssim \tau$ if and only if

$\forall s \in \text{Innocent plays}(\sigma), \forall t \in \text{Innocent plays}(\tau)$
 $O - \text{strategy}(\sigma) \subseteq O - \text{strategy}(\tau) \implies \sqcup s \sqcup = \sqcup t' \sqcup$
for some $t' \sqsubseteq t$.

Proof:

Suppose $\sigma \not\lesssim \tau$, i.e. there is some (innocent) $\rho : A \rightarrow N$ such that $\sigma; \rho = \bar{n}$ and $\tau; \rho \neq \bar{n}$. Then let s be the uncovering of (the sequence consisting of) $\llbracket N$, with respect to σ, ρ , and let t be the uncovering of $\llbracket N$, with respect to τ, ρ .

Let s' be the maximal innocent prefix of $s \upharpoonright A$ such that $O - \text{strategy}(s') \subseteq O - \text{strategy}(t)$. Then $s', t \upharpoonright A$ are even-length innocent plays, but $\sqcup s' \sqcup \neq \sqcup t' \sqcup$ for any $t' \sqsubseteq t$, as Opponent plays according to ρ in both s and t , and $O - \text{strategy}(s') \subseteq O - \text{strategy}(t)$ and $\sqcup s' \sqcup \neq \sqcup t \sqcup$.

To show the converse, suppose that there is some $s \in \text{Innocent plays}(\sigma)$ and $t \in \text{Innocent plays}(\tau)$, such that $O - \text{strategy}(s) \subseteq O - \text{strategy}(t)$ but $\sqcup s \sqcup \neq \sqcup t' \sqcup$.

Define the following strategy ρ on $A \rightarrow N$:

$$\rho = Pref^{even}(\llbracket Ns \rrbracket_0)$$

Then ρ is innocent, as s is an innocent play (so the last move does not violate innocence as there is no prefix of $\llbracket Ns$ with the same P-view). And $\sigma; \rho = \bar{0}$, but $\tau; \rho$ is undefined, so $\sigma \not\lesssim \tau$.

Another characterization of the intrinsic preorder is to consider it as being induced by mapping to the sequential algorithm which denotes the same term in sPCF.

The intrinsic preorder on compact strategies is decidable, because we can give a bound (of $2|\text{dom}(f_\sigma)|$) for the length of the innocent plays which can differentiate σ from τ by the above criterion.

8. Further directions and conclusions

- The version of *call/cc* defined here at atomic types is a natural choice of control operator with which to increment PCF: it denotes the morphism used in the factorization theorem, it can be typed with PCF types, and it has no additional variable bindings to deal with. However, as noted in [16], adding double negation elimination, (or, equivalently Peirce’s law) at atomic types to intuitionistic natural deduction is sufficient to derive classical laws for all arrow types. (The factorization theorem is really just a semantic version of this bit of elementary proof theory.) Moreover, the derived higher-order classical terms behave appropriately, extending Curry-Howard: the derived term of type $((A \rightarrow \perp) \rightarrow \perp) \rightarrow A$ obeys the equalities for the operator \mathcal{C} [8], for instance. This offers the possibility of using the well/weakly bracketed distinction to investigate proof theory semantically, in particular using factorizations to extract computational content from classical proofs in a more direct and efficient fashion than by using existing methods of double negation translation.
- The construction corresponding to the factorization theorem appears to be quite general, and to have some nice properties, such as commuting with composition, posing the the question: what is its categorical characterization, and can it be expressed more generally, for instance in terms of the Geometry of Interaction?
- As already noted, the Sequential Algorithms model of sPCF is isomorphic to the quotient of the dialogue games model. The requirement of innocence in the latter corresponds to the characterization of plays in the sequential data structure $!A$ as *non-repetitive* sequences of plays in A . A semantic characterisation of the isomorphism would clarify this distinction between sharing and copying of information.
- Control operators are better understood in the context of call-by-value languages. Games semantics for such languages are currently under investigation, and it appears that innocent, well-bracketed, strategies yield fully abstract models of purely functional languages, and that violating the bracketing condition corresponds to the use of control operators, and that such violations can be factored out, as described here.
- In the call-by-name and call-by-value settings one can combine the bracketing factorization

with the factorization of knowing strategies described in [1] to get a definability result for PCF with state and control mechanisms - seemingly a foundation for a broad notion of sequential non-functional computation.

9. Acknowledgements

I am indebted to Samson Abramsky, and others in the Edinburgh interaction group, particularly Paul-André Mellies and Juliusz Chroboczek, for valuable discussions and suggestions. Graham Clark suggested many corrections to the layout and clarifications of the language which I have adopted. Thanks are also due to Pierre-Louis Curien for a different perspective on control operators, and to suggestions made in reports on an initial version of this paper. Financial support in the form of an EPSRC studentship is also gratefully acknowledged.

References

- [1] S. Abramsky and G. McCusker. Linearity Sharing and state: a fully abstract game semantics for idealized algol with active expressions (extended abstract). *Elsevier Electronic notes in Theoretical Computer Science*, (3), 1996.
- [2] S. Abramsky, R. Jagadeesan and P. Malacaria. Full abstraction for PCF. Submitted for publication.
- [3] G. Berry, P.-L. Curien, J.-J. Lévy. Full abstraction for sequential languages: The state of the art. In J. Reynolds, editor, *Algebraic methods in Semantics*. Cambridge university press, 1985.
- [4] R. Cartwright, P.-L. Curien and M. Felleisen. Fully abstract semantics for observably sequential languages. *Information and Computation*, 1994.
- [5] C. Murthy. *Extracting constructive content from classical proofs*. PhD thesis, Cornell University, 1992.
- [6] T. Coquand. A semantics of evidence for classical arithmetic. *Journal of Symbolic Logic*, 1995.
- [7] P.-L. Curien. On the symmetry of sequentiality. In *Mathematical Foundations of Computer Science*, number 802 in LNCS. Springer, 1993.
- [8] M. Felleisen, D. P. Friedman, E. Kohlbecker and B. Duba. A syntactic theory of sequential control. In *LICS 1986*, pages 131 – 141. IEEE Computer Society Press, 1986.
- [9] H. Herbelin. Games and weak-head reduction for classical logic. Manuscript, 1996.
- [10] J. M. E. Hyland and C.-H. L. Ong. On full abstraction for PCF: I, II and III. to appear, 1995.
- [11] R. Loader. Finitary PCF is undecidable. Manuscript, 1996.
- [12] G. McCusker. *Games and full abstraction for a functional metalanguage with recursive types*. PhD thesis, Imperial College London, 1996.

- [13] C.-H. L. Ong. A semantic view of classical proofs: type-theoretic, categorical and denotational characterizations. In *LICS 96*. IEEE Computer Society Press, 1996.
- [14] C.-H. L. Ong and C. A. Stewart. A Curry-Howard foundation for functional computation with control. to appear in proceedings POPL, 1997.
- [15] G. D. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5:223 – 255, 1977.
- [16] D. Prawitz. *Natural Deduction*. Number 3 in Stockholm studies in Philosophy. Almqvist and Wiksell, 1965.