Game Semantics & Abstract Machines

V. Danos — CNRS-Université Paris 7
H. Herbelin — INRIA-LITP
L. Regnier — IML-CNRS Marseille

September 2, 1996

Abstract

The interaction processes at work in Hyland and Ong (HO) and Abramsky, Jagadeesan and Malacaria (AJM) new game semantics are two preexisting paradigmatic implementations of *linear head reduction*: respectively Krivine's Abstract Machine and Girard's Interaction Abstract Machine.

There is a simple and natural embedding of AJM-games to HO-games, mapping strategies to strategies and reducing AJM definability (or full abstraction) property to HO's one.

1 Introduction

Syntax without semantics is blind, semantics without syntax is empty, once said ... ? Anyway game semantics has recently produced two quite different-looking models, HO-games and AJM-games, the former due to Hyland and Ong, and the latter to Abramsky, Jagadeesan and Malacaria. Both triggered new intuitions about syntax, while offering enough domain-theoretic structure to carry over the basic regulative tasks traditional semantics afford. They aroused considerable interest since they were delivering *concrete* definability results: all strategies were interpreting terms, and these results needed real proofs, which could only mean one was facing two non-syntactic reformulations of syntax. One felt the interesting part of the full abstraction problem was solved. But then two questions popped out: what were the interaction processes at work in these two models; what was the relation between them.

The present paper traces back HO and AJM semantics to two abstract machines: the PAM, a variant of Krivine's Abstract Machine directly formulated in terms of pointers (credible as an abstract machine for pure λ-calculus, elegant, no closures); and the IAM, a machine for Linear Logic that comes from Girard's geometry of interaction, GoI for short.

These two machines themselves implement linear head reduction, a hyper-lazy reduction strategy. Thus we answer the first question. Now our "machine-behind-the-game" theorems not only give concrete counter-weight to abstract game manipulation, but they show both models touch common ground in that they are tied to linear head reduction. So the second question is all the more tempting.

Let's consider it again. Both models bring to the fore the fact that linear head reduction really is an *interaction* between two independent agents: the function and its arguments. But to be able to interact, both agents need their own information. Here the two models depart in the way they let agents handle that information.

In HO games, information of both agents is mixed into a common data base: the history of the reduction. It is then by a pointer-based mechanism, as emphasized in the PAM, that function and arguments recover from the history the data they need. On the opposite, in AJM games, information shared by both agents is just a "token" that passes from one to the other, *but*, in this token, both agents encode also their own private information. This is a "communication without understanding machinery, of which the GoI, as implemented by the IAM, is a typical instance. Our third theorem, that exhibits an embedding of AJM strategies in HO ones, answers the second question. But as said, GoI is just an instance of what one might call an history-free digitalization of pointer-based evaluation mechanisms, and it would be interesting to investigate in general such digitalizations.

The paper sticks to typed λ-calculus for which the question/answers apparatus is superfluous. This helps in making shorter definitions: all the games and machines in a nutshell !

2 Linear head reduction

This section introduces our basic notations on typed λ-calculus together with various concepts linked to linear head reduction.

2.1 λ-calculus ...

NOTATION. We write $(U)V$ for the application of U to V and sometimes $(U)V_1 \ldots V_n$ for $(\ldots(U)V_1\ldots)V_n$. If we are not interested in the actual number of V_i's we write $(U)\vec{V}$ for $(U)V_1\ldots V_n$. Similarly we write $\lambda\vec{x}\,U$ for $\lambda x_1\ldots\lambda x_n\,U$ when we are not interested in the number of λ's.

OCCURRENCES. Let T be a term. We will call *T-occurrences*, or when it is not ambiguous, simply *occurrences* the occurrences of variables in T.

σ-EQUIVALENCE. We will deal with typed and untyped λ-calculus quotiented by the *σ-equivalence*, that is, the smallest congruence on (typed) terms which contains:

$$\begin{array}{rcl} ((\lambda x\,U)V_1)V_2 & \simeq_\sigma & (\lambda x(U)V_2)V_1 \\ (\lambda x_1\lambda x_2\,U)V & \simeq_\sigma & \lambda x_2(\lambda x_1\,U)V \end{array}$$

where x is not a free variable of V_2 in the first clause and x_2 is not a free variable of V in the second clause.

Suppose $T \simeq_\sigma T'$ then both are typable of the same type, furthermore (as proved in [11]) T and T' have the same length of head reduction, leftmost reduction and longest reduction. So there is no significant difference in the dynamical behavior of σ-equivalent terms.

PRIME REDEXES. A *prime redex* in T is given by a λ and/or an argument subterm V (called the argument of the redex) such that there is a term $T_\sigma \simeq_\sigma T$ in which this λ and/or V defines a redex. Typically in $T = ((\lambda x_1\lambda x_2 U)V_1)V_2$ there is one redex $(\lambda x_1, V_1)$ but since $T \simeq_\sigma (\lambda x_1(\lambda x_2 U)V_2)V_1$ (by the first σ-equation), we may as well consider $(\lambda x_2, V_2)$ as a prime redex.

CANONICAL FORMS. It is easy to see that any term T is σ-equivalent to a *canonical form* T_σ (not unique in general):

$$T_\sigma = \lambda y_1\ldots\lambda y_m(\lambda x_1\ldots(\lambda x_n(x)U_1\ldots U_p)V_n\ldots)V_1$$

where x, the leftmost T-occurrence may be free, or bounded by some λy_i or some λx_j. The V_j's and the U_k's are called the *immediate subterms* of T. The U_k's are also called the *arguments* of the T-occurrence x. When T is normal, its immediate subterms correspond to the sons of its head variable in T's Böhm tree.

OUTPUT SUBTERMS. Let T be a term and x be a T-occurrence. The *output subterm* T_x of x in T is the biggest subterm of T whose leftmost occurrence is x. When T is normal this corresponds to the subtree of T's Böhm tree starting with x.

NOTATIONS AND TERMINOLOGY. Up to σ-equivalence, T_x has the form of T_σ above. The $(\lambda x_i, V_i)$'s are called the *spine* redexes of x or of T_x, *i.e.*, prime redexes whose λ's are left of x in T_x. Each λy_j is called the *head* λ of *rank* j of x or of T_x and will be denoted by $H_T(x,j)$. Each U_k's is called the *argument* of rank k of x and will be denoted by $A_T(x,k)$. If z is the leftmost occurrence of some U_k then we will also say that the T-occurrence z is argument of the T-occurrence x. These definitions make sense since the rank of a head λ or of the argument of a T-occurrence is invariant by σ-equivalence. Finally we will denote by ${}_xT$ the maximal subterm of T, if there is one, for which the binder of x is a head λ.

From now on we will consider λ-terms up to σ-equivalence.

η-CONVERSION. Given a term T and a T-occurrence x, $H_T(x,j)$ and $A_T(x,k)$ may be viewed as partial functions on integers. In the sequel we will consider terms up to η-conversion so that these functions will actually be *total*. Indeed, suppose T_x has the form of T_σ above and $q > m$; since

$$\lambda y_1\ldots\lambda y_m\,U \simeq_\eta \lambda y_1\ldots\lambda y_m\lambda z_{m+1}\ldots\lambda z_q(U)z_q\ldots z_{m+1}$$

where $U = (\lambda x_1\ldots(\lambda x_n(x)U_1\ldots U_p)V_n\ldots)V_1$ has no head λ, we will write $H_T(x,q) = \lambda z_q$. Similarly, if $q > p$, since

$$(x)U_1\ldots U_p \simeq_\eta \lambda z_q\ldots\lambda z_{p+1}(x)U_1\ldots U_p z_{p+1}\ldots z_q$$

we will write $A_T(x,q) = z_q$. Note that the performed η-expansions don't add any new prime redex to T.

Be aware that when we speak of a T-occurrence x, it doesn't mean that x occurs in T, but in some η-expansion of T.

LIFT. The *lift* c_x of an occurrence x of a bounded variable is the *Böhm index* (by analogy with de Bruijn indexes) between ${}_xT$ and T_x, *i.e.*, the number of times, starting from ${}_xT$ one has to move in an immediate argument to reach x.

LINEAR SUBSTITUTION. Let T and U be terms, x be the leftmost occurrence of T. We will call *linear substitution* of x by U, denoted $T\,[U/x]$, the substitution of x as a T-occurrence, by U. Note that the linear substitution operates even though x is an occurrence of a bound variable in T, *e.g.*, $\lambda x(x_0)x_1\,[\lambda y(y)y/x_0] = \lambda x(\lambda(y)y)x_1$ where x_0 and x_1 are two occurrences of the variable x.

LINEAR HEAD REDUCTION. Suppose T is σ-equivalent to the T_σ above. Note that T is a head normal form iff $n = 0$. Also, if x is not bounded by a λx_j, *i.e.*, if λx doesn't define a spine redex, then in exactly n steps of beta-reduction we obtain a head normal form for T; in this case we say that T is a *quasi head normal form*. It is equivalent for a term to have a head normal form or to have a quasi head normal form. The *linear head reduction* is a reduction strategy seeking for the "closest" quasi head normal form of a term. If x is bounded by λx_j then the strategy performs the linear substitution of x by V_j producing the term:

$$T' = \lambda y_1 \ldots \lambda y_p(\lambda x_1 \ldots (\lambda x_n(V_j)U_1 \ldots U_p)V_n \ldots)V_1.$$

Take note that the linear head reduction is *not* a β-reduction because, on the one hand at each step only one occurrence of variable is substituted and the fired redex remains, on the other hand it may be defined only up to σ-equivalence.

The head linear reduction may also be defined easily using explicit substitution calculi, *e.g.*, the $\lambda\sigma$-calculus [1].

CORRESPONDENCE WITH PROOF-NETS. In fact the linear head reduction is a strategy for reducing proof-nets of linear logic. In proof-nets σ-equivalence corresponds to equivalence up to multiplicative reductions and the linear substitution corresponds to the firing of an exponential branch [9].

2.2 ... and types

ATOMS. Types are built with the $\to$ connective, starting from a *single* atomic type ι. We will denote types by A, B, C. We will call *atoms* the occurrences of ι in A. If A is a type then the positive atoms in A will be denoted by β and the negative ones by α. We will use γ and δ for atoms in A with unspecified sign.

The type $A = A_1 \to (\ldots \to (A_k \to \gamma)\ldots)$ will be denoted by $A_1 \ldots A_k \to \gamma$ and will be said of *arity* k. The atom γ will be called the *terminal* atom of A and will also be said of arity k. The A_i's are called the *immediate subtypes* of A and we will write $A_i \prec A$. If γ_i is the terminal atom of A_i, then we will also write $\gamma_i \prec \gamma$ and say that γ_i is an immediate *subatom* of γ. The indices i's are the *ranks* of the immediate subtypes of A and we will write $A/A_i = i$ or $\gamma/\gamma_i = i$. Given an atom δ in A, we will denote by A_δ the biggest subtype of A which has δ as terminal atom. Finally we will call *depth of* δ *in* A the number $d(\delta, A)$ inductively given by: if δ is γ then $d(\delta, A) = 0$, otherwise δ is an atom of some A_i and $d(\delta, A) = d(\delta, A_i) + 1$.

In the sequel of this section, $T[x_i : A_i] : A$ will stand for a *normal* term of type A, whose free variables $x_1, \ldots, x_n$ are declared with types $A_1, \ldots, A_n$.

INPUT AND OUTPUT ATOMS. We associate to each T-occurrence x an *input atom* $I_T(x)$ which is a negative atom in $A_1 \ldots A_n \to A$, and an *output atom* $O_T(x)$ which is a positive atom in $A_1 \ldots A_n \to A$. Input and output atoms are defined by induction on T:

- if T is $(x_i)U_1 \ldots U_m$ and x is its leftmost occurrence (thus as a variable x is x_i), then $I_T(x)$ is the terminal atom of A_i and $O_T(x)$ is the terminal atom of A.
- if T is $(x_i)U_1 \ldots U_m$ and x occurs in U_j then A_i must have the form $A_i = B_1 \ldots B_p \to \alpha$ and for each j we have $U_j[x_i : A_i] : B_j$. We define $I_T(x) = I_{U_j}(x)$ and $O_T(x) = O_{U_j}(x)$.
- If T is $\lambda y\, U$, then $A = B \to C$ and $U[x_i : A_i, y : B] : C$; we define $I_T(x) = I_U(x)$ and $O_T(x) = O_U(x)$.

Note the the output atom of the T-occurrence x is the terminal atom of the type of the output subterm T_x.

LINKS. Let x be a T-occurrence and β_x, α_x be respectively its output and input atoms in $A_1 \ldots A_k \to A$. We will say that x *links* β_x *and* α_x. This terminology is reminiscent from proof-nets. Indeed a T-occurrence linking two atoms corresponds to an *axiom link* linking these atoms in the proof-net.

3 A linear head reduction machine: the PAM

We will now design a machine which performs linear head reduction. This machine was inspired by Krivine's environment machine and is called the PAM (Pointer Abstract Machine). For the sake of simplicity we will work with a term T of the form

$$T = (U)V_1 \ldots V_k$$

where U and the V_i's are normal. Since it is easy with a few β-expansions to transform any term in such a form, this restriction is innocuous.

RUNS. Given the term T, the PAM works by building a T-*run*, *i.e.*, a finite sequence $s = x_0, \ldots, x_n$ of T-occurrences together with a *pointer function* θ from $\{1, \ldots, n\}$ into $\{0, \ldots, n-1\}$ such that for each $p > 0$, $\theta(p) < p$ and, denoting by T_p the output subterm of x_p in U (resp. V_i) if x_p occurs in U (resp. V_i):

λ-invariant: the binder of x_p is a head λ of $T_{\theta(p)+1}$, that is:

$$_{x_p}T = T_{\theta(p)+1};$$

Application invariant: T_{p+1} is an argument of the T-occurrence $x_{\theta(p)}$, that is, for some i:

$$T_{p+1} = A_T(x_{\theta(p)}, i).$$

The pair $(x_p, \theta(p))$ is called a *pointing occurrence* and we will denote it simply by x_p. The integer $\theta(p)$ is called the *pointer* of x_p and we say that x_p points to $x_{\theta(p)}$ in the T-run s.

INITIALIZATION OF THE PAM. The machine starts with the sequence x_0, x_1 and $\theta(1) = 0$ where x_1 is the leftmost occurrence of T. The initial x_0 plays a special rôle: we will treat it as a T-occurrence (although it actually appears neither in T, nor in any η-expansion of T), which has no head λ, no spine redex but has the V_i's as arguments, *i.e.*, $A_T(x_0, i) = V_i$. Intuitively, the "T-occurrence" x_0 is defined by $T = (x_0)V_1 \ldots V_k\,[U/x_0]$.

CONSTRUCTION OF x_{n+1}. At step n the PAM proceeds as follows: let m be $\theta(n)$. By the λ-invariant, the binder of x_n is a head λ in T_{m+1} of rank i; put T_{n+1} to be $A_T(x_m, i)$ and x_{n+1} to be the leftmost occurrence of T_{n+1}. In this way we obviously satisfy the application invariant.

If $m = 0$ and i is strictly greater than k then the machines stops. This means that x_{n+1} is bounded by a head λ of T.

CONSTRUCTION OF $\theta(n+1)$. For each $p \le n$, if $\theta(p) \ne 0$ then define $F(p)$ to be $\theta(p) - 1$. We set $\theta(n+1)$ to be $F^c(n)$ where $c = c_{x_{n+1}}$ is the lift of x_{n+1}. If $F^c(n)$ is undefined, then the machine stops. This means that x_{n+1} is a free T-occurrence.

We are to show that $\theta(n+1)$ satisfies the λ-invariant. By the application invariant, note that for each p smaller than n, T_{p+1} is an immediate subterm of $T_{F(p)+1}$. Therefore, if x_{n+1} is bounded in T, its binder must appear as a head λ of $T_{F^i(n)+1}$ for some i. But by definition of the lift of x_{n+1} this happens exactly for $i = c_{x_{n+1}}$.

ALTERNATION. It is easily checked that the successive T-occurrences of a T-run are alternatively occurrences in U and in the V_i's. In other words, if n is even then x_n is a V_i-occurrence for some i and if n is odd then x_n is a U-occurrence. This is one reason why we conventionally assume that x_0 is a V_i-occurrence. Also this alternation property is the first hint that there might be a link between the PAM and games.

THEOREM 1 **(Correction of the PAM)** *Let $s = x_0, \ldots, x_n$ be the T-run of length n produced by the* PAM *on the input T and $T^h_0 = T, T^h_1, \ldots, T^h_n$ be the sequence of terms obtained by head linear reduction from T. Then for each $p > 0$ we have:*

$$T^h_p = T\,[T_2\,[\ldots[T_p\,[T_{p+1}/x_p]/x_{p-1}]\ldots]/x_1]$$

where all the substitutions are linear.

Proof. By induction on p. If $p=1$ then by definition of the PAM, T_1 is U since x_1 is the leftmost occurrence of T, thus of U. Now since $\theta(1)=0$, the binder of x_1 is a head λ of $T_1=U$, thus T_2 is one of the V_i. Therefore, by definition of head linear reduction T_1^h is $T\,[T_2/x_1]$

Suppose $p \geq 1$. By induction x_{p+1}, which is the leftmost occurrence of T_{p+1}, is the leftmost occurrence of T_p^h. Therefore T_{p+1}^h is $T_p^h\,[V/x_{p+1}]$ for some V. We are to show that $V=T_{p+2}$. Let m be $\theta(p+1)$. By the λ-invariant we know that the binder λx_{p+1} of x_{p+1} is a head λ of T_m. By induction, $T_m^h = T_{m-1}^h\,[T_m/x_{m-1}]$ therefore if λx_{p+1} defines a prime redex in T_m^h, thus in T_{p+1}^h, then the argument V of λx_{p+1} is an argument of x_{m-1} in T_{m-1}. Furthermore V is the argument of x_{m-1} of rank the rank of λx_{p+1} among the head λ's of T_m, *i.e.*, by definition of the PAM, V is T_{p+2} and we are done. □

4 HO-games

We will now give an interpretation of terms by *innocent strategies*. This definition is the direct adaptation for typed λ-calculus of Hyland and Ong's definition for PCF. We start by giving a short introduction to games and strategies. We have restricted our presentation of HO-games to what we strictly needed. See [6] for details.

Pointing sequences. A *pointing sequence* in the formula A is a finite sequence $s=\gamma_0,\ldots,\gamma_n$ of atoms in A together with a partial *pointing function* θ_s from $\{1,\ldots,n\}$ into $\{0,\ldots,n-1\}$ satisfying $\theta_s(k)<k$ when $\theta_s(k)$ is defined. When it is not ambiguous, we will simply write $\theta(k)$ for $\theta_s(k)$.

Each pair $(\gamma_k,\theta(k))$ will be called a *move* or a *pointing* atom. We will abusively denote by γ_k the move $(\gamma_k,\theta(k))$. We say that $\theta(k)$ is the *pointer* of the move γ_k and that γ_k *points to* $\gamma_{\theta(k)}$ in s. Note that, considered as a move, γ_0 has no associated pointer.

Playing. Given a pointing sequence $s=\gamma_0,\ldots,\gamma_n$, we say that we *play a new move* γ *w.r.t.* s when we extend s with a new move (γ,k) with $k\leq n$, obtaining a pointing sequence $s^+=\gamma_0,\ldots,\gamma_n,\gamma$ with an extended pointing function θ^+ defined by $\theta^+(i)=\theta(i)$ when $i\leq n$ and $\theta(i)$ is defined, $\theta^+(n+1)=k$.

Subsequences. Let $s=\gamma_0,\ldots,\gamma_n$ be a pointing sequence. A *subsequence* of s is a pointing sequence $s'=\gamma_{i_0},\ldots,\gamma_{i_k}$ where $\theta_{s'}(j)=l$ if $\theta_s(i_j)=i_l$. In other words γ_{i_j} points to γ_{i_l} in s' when it does in s.

Conversely, if s' is a subsequence of s and if we add a new move γ w.r.t. s' then γ points to a move in s', thus in s, so that γ may also be considered as a new move for s. We shall use this fact without further comment in the sequel.

Player and opponent moves. Each pointing atom has a sign which is the sign of its corresponding atom in A. Positive pointing atoms will be called *opponent moves* (O-moves), and negative ones will be called *player moves* (P-moves).

Alternating condition. Let $s=\gamma_0,\ldots,\gamma_n$ be a pointing sequence in A. We say that s satisfies the *alternating condition* if, firstly, γ_0 is the terminal atom of the formula A (in particular γ_0 is an O-move); secondly, any two successive moves in s are of opposite sign.

Justification condition. We say that s satisfies the *justification condition* if for each $k>0$, $\theta(k)$ is defined and γ_k is *justified by* $\gamma_{\theta(k)}$, that is γ_k is an immediate subatom of $\gamma_{\theta(k)}$ in A.

Views. Suppose s satisfies both the alternating and justification conditions. The *P-view* $\mathcal{V}^P(s)$ of s is the subsequence of s given by:

- if s is empty then $\mathcal{V}^P(s)$ is empty; if $s=\beta$ is reduced to a single O-move then $\mathcal{V}^P(s)=\beta$;
- if $s=s_0,\alpha$ ends with a P-move, then $\mathcal{V}^P(s)=\mathcal{V}^P(s_0),\alpha$;
- if $s=s_0,\alpha,s_1,\beta$ ends with an O-move β justified by α then $\mathcal{V}^P(s)=\mathcal{V}^P(s_0),\alpha,\beta$.

The O-view of s is defined symmetrically.

Visibility condition. Let $s=\gamma_1,\ldots,\gamma_n$ be a pointing sequence. The *visibility condition* states that a move has to be justified in the current player's view, that is, for any $0<k\leq n$, if γ_k is an X move ($X=P$ or O) then $\mathcal{V}^X(\gamma_0,\ldots,\gamma_{k-1})$ is required to contain the move $\gamma_{\theta(k)}$.

GAMES. Given a formula A, the *game* A is the set of *plays* in A, that is, pointing sequences satisfying the alternating, justification and visibility conditions. Note that any view of a play in the game A is a play in the game A.

From now on all the pointing sequences that we will consider will be plays in some game.

INNOCENT STRATEGIES. An *innocent strategy* σ for player in the game A is a tree of P-views that is *deterministic*, by which we mean that the tree branches only on O-moves, *i.e.*, O-moves have a unique son.

PLAYING STRATEGIES. We will say that a play s in the game A *belongs* to the strategy σ if for any prefix p of s ending with a P-move we have $\mathcal{V}^P(p) \in \sigma$. If s is a play belonging to σ and ending with an O-move, then we will say that σ *plays* the move α, or that σ *answers* α to s if s, α belongs to σ. Note that by the determinism condition, σ can answer at most one move to s. We will say that σ is *extensionally given* if σ is presented as the tree of its plays.

TOTAL STRATEGIES. We will say that σ is *total* if for any play s belonging to σ and ending with an O-move, there is a move α that σ answers to s. An equivalent (and somehow more concrete) definition is that σ is total if for any P-view $p \in \sigma$ and any O-move β justified by the last move of p, there is a P-move α such that $p, \beta, \alpha \in \sigma$.

4.1 HO-dialogs.

In this section A will stand for the type $A_1 \ldots A_n \to \beta_0$ and for each i we denote by α_i the terminal atom of A_i.

LEMMA 2 **(Switching)** *Let β_0, s be an O-view in A. Then s is a P-view in A_i for some i.*

Proof. By induction on s. In the base case, s is empty and there is nothing to say. Otherwise $s = s', \gamma$ and by induction, s' is a P-view in A_i. Now γ is justified by some move γ' in β_0, s'. If γ' is β_0 then by definition of O-views, $s = \beta_0, \gamma$ and γ is an immediate subatom of β_0, *i.e.*, the terminal atom of some A_i . Therefore $s' = \gamma$ is a P-view consisting of only one move in A_i. Otherwise γ' being in s' belongs to A_i and γ being a subatom of γ' also belongs to A_i. But A_i is a negative subtype of A so that P-moves (resp. O-moves) in A_i are O-moves (resp. P-moves) in A. From the definition of views it is now immediate that s is a P-view in A_i. □

REMARK. A consequence of this lemma is that in a play the opponent can never switch between the A_i's. Indeed, when it is opponent to play, his view contains only atoms in A_i so that, by the visibility condition, he has to play in A_i. This induced rule is known as the *switching convention.*

HO-DIALOGS. Let σ be a strategy in A and for $i = 1, \ldots, n$, let σ_i be a strategy in A_i. We will now build a play in A, called the *HO-dialog* between σ and the σ_i's. The construction is just a particular case of the composition of strategies which however contains all the dynamics of the general case. The idea is that σ plays the P-moves and the σ_i's play the O-moves. Precisely, the HO-dialog s_k of length $k+1$ between σ and the σ_i's is inductively defined by: s_0 is β_0 and for $k \geq 0$:

- if k is even then s_k ends with an O-move; σ plays the move α such that the play $s_{k+1} = s_k, \alpha$ belongs to σ; if there is no such move then s_k is maximal (and σ lost the play);
- if k is odd then s_k ends with a P-move α which therefore cannot be β_0; thus α is an atom in one of the A_i. Let s_k^i be the O-view of $s = s_k/\beta_0$, *i.e.*, s_k from which the first move β_0 has been removed. By the switching lemma s_k^i is a P-view in A_i so that σ_i plays the move β, if there is one such that $s_k^i, \beta \in \sigma_i$ and s_{k+1} is defined to be s_k, β. If there is no such move then s_k is maximal (and the σ_i's lost the play).

5 HO-games and λ-calculus

5.1 Interpretation of terms

NOTATIONS. Let $x_1, \ldots, x_m$ be variables of type $A_1, \ldots, A_m$ and $T = \lambda x_{m+1} \ldots \lambda x_n (x) U_1 \ldots U_p$ be a normal η-long term of type $A_{m+1} \ldots A_n \to \beta$ whose free variables are $x_1, \ldots, x_m$. Denote by α_k the terminal atom of A_k for each $k = 1, \ldots n$. In particular the leftmost occurrence x of T is an occurrence of x_i for some i and its input type is $A_i = B_1 \ldots B_p \to \alpha_i$ where the B_j's are the types of the U_j's. For each $j = 1, \ldots, p$ let β_j be the terminal atom of B_j.

Interpreting terms with strategies. We build by induction on T a set σ_T of pointing sequences in the game $A_1 \ldots A_n \to \beta$. By induction on U_j we may suppose that for each $j = 1, \ldots, p$ we already have built σ_{U_j} a set of pointing sequence in the game $A_1 \ldots A_n \to B_j$. Then σ_T contains β, α_i where α_i points onto β and for each $j = 1, \ldots, p$, all the pointing sequences of the form:

$$\beta, \alpha_i, \beta_j, s$$

where α_i points onto β, β_j points onto α_i, and β_j, s is the *lifting* to the game $A_1 \ldots A_n \to \beta$ of a pointing sequence β_j, s' belonging to σ_{U_j}; by lifting we mean that s is the same sequence of atoms than s' and that all the atoms α_k appearing in s points onto β, the pointers of the other atoms being unchanged between s and s'.

Proposition 3 *The set σ_T is an innocent strategy.*

Proof. By induction on T. If β_j, s' is an element of σ_{U_j} then by induction on U_j, the sequence β_j, s' is a P-view in B_j which ends with a P-move. But B_j being a positive subtype of $A_1 \ldots A_n \to \beta$, a P-move in B_j is also a P-move in $A_1 \ldots A_n \to \beta$ so that $p = \beta, \alpha_i, \beta_j, s$ also ends with a P-move. From the definition we immediately get that p is a P-view in $A_1 \ldots A_n \to \beta$. It is also immediate that, since by induction on U_j, σ_{U_j} is deterministic, so is σ_T. □

5.2 A definability theorem

Proposition 4 *Let $s = \beta_0, \alpha_0, \ldots, \beta_p, \alpha_p$ be a P-view ending with a P-move in the game $A = A_1 \ldots A_n \to \beta_0$. There is a term T_s of type A together with a sequence of T_s-occurrences $x_0, \ldots, x_p$ such that for each k:*

- *x_k links β_k and α_k;*
- *x_{k+1} is argument of x_k;*
- *if α_k points to β_m in s then the binder of x_k is a head λ of x_m.*

Proof. The construction of T_s is by induction on p. If $p = 0$ then $s = \beta_0, \alpha_0$. Note that α_0 must be the terminal atom of A_i for some i. Write $A_i = B_1 \ldots B_l \to \alpha_0$. Then we define $T_s = \lambda z_1 \ldots \lambda z_n (z_i) \omega_1 \ldots \omega_l$ where the ω_j's are free variables of type B_j and the z_k's are declared with type A_k. Also we set x_0 to be the leftmost occurrence of T_s so that the binder of x_0 is indeed a head λ of T_s

Otherwise $s = \beta_0, \alpha_0, \ldots, \beta_p, \alpha_p, \beta_{p+1}, \alpha_{p+1}$. Let T_s^p be the term associated to $\beta_0, \alpha_0, \ldots, \beta_p, \alpha_p$. By induction T_s^p has an occurrence of variable x_p which links β_p and α_p and the output subterm T_p of x_p in T_s^p has the form $\lambda \vec{x}(x_p)\omega_1 \ldots \omega_q$ where q is the arity of α_p. Therefore the input type of x_p has the form $B_1 \ldots B_q \to \alpha_p$. By definition of P-views, β_{p+1} is justified by α_p, thus is a subatom of α_p. Let $j_{p+1} = \alpha_p / \beta_{p+1}$ so that the type $B_{j_{p+1}}$ has the form $B_{j_{p+1}} = C_1 \ldots C_m \to \beta_{p+1}$. Define T'_{p+1} to be the term $\lambda z_1 \ldots \lambda z_m \omega$ where the z_i's are new variables of type C_i and ω is a new variable.

If α_{p+1} points to β_{p+1} in s, then α_{p+1} is an immediate subatom of β_{p+1} so that we may define $k_{p+1} = \beta_{p+1}/\alpha_{p+1}$. We define x_{p+1} to be a new occurrence of $z_{k_{p+1}}$ and $T_{p+1} = \lambda z_1 \ldots \lambda z_m (x_{p+1}) \omega'_1 \ldots \omega'_l$ where l is the arity of α_{p+1} and the ω'_i's are new variables of appropriate types. If α_{p+1} points to β_m in s with $m \leq p$ then by induction, the output subterm of x_m in T_s^p has the form $\lambda z'_1 \ldots \lambda z'_{m'} U$ where U has type β_m and the z'_i's have types B'_i. Since $\alpha_{p+1} \prec \beta_m$, α_{p+1} is the terminal atom of $B'_{k_{p+1}}$ where $k_{p+1} = \beta_m / \alpha_{p+1}$. We define x_{p+1} to be a new occurrence of $z'_{k_{p+1}}$ and $T_{p+1} = \lambda z_1 \ldots \lambda z_m (x_{p+1}) \omega'_1 \ldots \omega'_l$ where l is the arity of α_{p+1} and the ω'_i's are new variables of appropriate types.

Finally we define T_s to be $T_s^p<T_{p+1}/\omega_{j_{p+1}}>$, *i.e.*, the *brutal substitution* of the occurrence $\omega_{j_{p+1}}$ by T_{p+1} in T_s^p; by "brutal" we mean that the substitution performs no α-conversion, so that it allows capture of variables. In particular x_{p+1} falls under the scope of its binder. The other variables of T_{p+1}, namely the ω_i's, being new cannot be captured. Note that the construction is made so that the binder of x_{p+1} is a head λ of the output subterm of x_m in T_s, where m is such that α_{p+1} points to β_m in s. Also, by definition, x_{p+1} is argument of x_p and links β_{p+1} and α_{p+1}. Therefore the proposition is proved. □

RESTRICTION. Let σ be a strategy in $A \to B$. The *restriction* $\sigma|_B$ of σ to B is defined by: s is in $\sigma|_B$ iff $s \in \sigma$ and all the atoms of s belongs to B. It is fairly easy to check that $\sigma|_B$ is a strategy in B.

THEOREM 5 **(Definability)** *Let σ be a finite strategy in A. Then there is a normal term T of type A, whose free variables have type $B_1, \ldots, B_l$ such that $\sigma \subset \sigma_T|_A$. If furthermore σ is total then there is a unique closed term T of type A such that $\sigma = \sigma_T$.*

Proof (sketchy). Say that two terms T_1 and T_2 are *unifiable* w.r.t. $\omega_1, \ldots, \omega_n$ if the ω_i's are occurrences of free variables in the T_i's and there are some terms $U_1, \ldots, U_n$ such that $T_1{<}U_1/\omega_1, \ldots, U_N/\omega_n{>} = T_2{<}U_1/\omega_1, \ldots, U_N/\omega_n{>}$ where the substitutions are brutal. Let $T_1, \ldots T_p$ be the terms associated to each views in σ. The deterministic condition of strategies implies that all the T_i's are unifiable together w.r.t. the ω's. The unification of the T_i's forms a term T which satisfies the theorem. If furthermore σ is total, then all the ω's are instantiated by the unification so that T is closed. By construction T is such that $\sigma_T = \sigma$. But the application $T \mapsto \sigma_T$ is clearly injective on closed terms so that T is unique. □

5.3 HO-games and the PAM

VIEWS AND OCCURRENCES. For each term T and each view s in σ_T we define a T-occurrence $\mathcal{O}_T(s)$ by induction on T:

- if $s = \beta, \alpha$ then by definition of σ_T, α is the input atom of the leftmost occurrence x of T. We set $\mathcal{O}_T(s) = x$;
- if $s = \beta, \alpha, s'$ then by definition of σ_T, T has the form $\lambda\vec{y}(y)U_1 \ldots U_m$ and s' is (the lifting of) a view in σ_{U_i} for some i. By induction on s' we get an U_i-occurrence $x = \mathcal{O}_{U_i}(s')$. But U_i is a subterm of T, therefore x is a T-occurrence and we set $\mathcal{O}_T(s) = x$.

LEMMA 6 *Let $s = \beta_0, \alpha_0, \ldots, \beta_p, \alpha_p$ be a view in σ_T. For $k = 0, \ldots, p$, let $x_k = \mathcal{O}_T(\beta_0, \alpha_0, \ldots, \beta_k, \alpha_k)$. Then for each $k \leq n$ we have:*

- *the T-occurrence x_k links β_k and α_k;*
- *x_0 is the leftmost variable in T and if $k > 0$, the T-occurrence x_k is the ith argument of x_{k-1} where $i = \alpha_{k-1}/\beta_k$*
- *the binder of x_k is the jth head λ of x_l where l is such that β_l is the move justifying α_k in s and $j = \beta_l/\alpha_k$.*

Proof. Take the term T_s defined in proposition 4 and note that T is T_s in which the ω's have been instantiated by some terms. □

HO-DIALOGS AND RUNS. Let now $T = (U)V_1 \ldots V_k$ be an η-long term where U and the V_i's are normal and have type respectively $A = A_1 \ldots A_k \to \beta_0$ and A_i. Suppose that the free variables of T are $x_1, \ldots, x_l$ of type $B_1, \ldots, B_l$. Let σ (resp. σ_i for $i = 1, \ldots, k$) be the restriction of σ_U (resp. σ_{V_i}) to A (resp. to A_i).

We extend the function $\mathcal{O}_T(.)$ to a morphism $\overline{\mathcal{O}}_T(.)$ from HO-dialogs between σ and the σ_i's into sequences of pointing occurrences.

- if s is β_0 then $\overline{\mathcal{O}}_T(s) = x_0$ where x_0 is a fake T-occurrence.
- If s is s_0, β, s_1, α where α is a P-move pointing on the O-move β then, by induction on s, $\overline{\mathcal{O}}_T(s_0, \beta, s_1)$ is a sequence $x_0, \ldots, x_n$. Since $\overline{\mathcal{O}}_T(.)$ is a morphism between pointing sequences, $\overline{\mathcal{O}}_T(s_0, \beta)$ is the sequence $x_0, \ldots, x_m$ for some $m \leq n$. By definition of HO-dialogs the P-view $p = \mathcal{V}^P(s)$ belongs to σ_U. Let x_{n+1} be $\mathcal{O}_U(p)$. Then we define $\overline{\mathcal{O}}_T(s)$ to be the sequence $x_0, \ldots, x_{n+1}$ where x_{n+1} points to x_m.
- If s is s_0, α, s_1, β where β is an O-move pointing on α then by induction on s, $\overline{\mathcal{O}}_T(s_0, \alpha, s_1)$ is a sequence $x_0, \ldots, x_n$ of which $\overline{\mathcal{O}}_T(s_0, \alpha)$ is a subsequence $x_0, \ldots, x_m$. Let p be the O-view of s, so that $p = \beta_0, p'$. By the switching lemma and the definition of HO-dialogs, p' is a P-view in some σ_i so let x_{n+1} be $\mathcal{O}_{V_i}(p')$. Then we define $\overline{\mathcal{O}}_T(s)$ to be the sequence $x_0, \ldots, x_{n+1}$ where x_{n+1} points to x_m.

Conversely, let $s = x_0, x_1, \ldots, x_n$ be a T-run of the PAM. We define a pointing sequence $\mathcal{D}_T(s) = \beta_0, \alpha_1, \ldots, \gamma_n$ of atoms in A by:

- if $s = x_0, x_1$ where, by definition of the PAM, x_1 is the leftmost occurrence of U then $\mathcal{D}_T(s) = \beta_0, \alpha_0$ where α_0 is the input atom of x_1;
- if $s = x_0, \ldots, x_n, x_{n+1}$ then, by induction we have $\mathcal{D}_T(x_0, \ldots, x_n) = \beta_0, \ldots, \gamma_n$. We define the move γ_{n+1} to be the input atom of x_{n+1} in U if x_{n+1} is in U, in V_i if x_{n+1} is in V_i, which points to γ_m in $\mathcal{D}_T(s)$ iff x_{n+1} points to x_m in s. We set $\mathcal{D}_T(s) = \beta_0, \ldots, \gamma_n, \gamma_{n+1}$.

Occurrence views. Recall that in section 3 we defined $F(p) = \theta(p) - 1$ where θ is the pointing function of the T-run. Let $s = x_0, \ldots, x_{n+1}$ be a T-run of the PAM and d be the maximum integer such that $F^d(n)$ is defined. Let $n_0 = F^d(n)$, $n_1 = F^{d-1}(n)$, $\ldots n_d = F^0(n) = n$ so that $x_{n_0+1}, \ldots, x_{n_d+1}$ is a subsequence of s that we will call the *view of* x_{n+1}. Precisely denoting by W the subterm U if x_{n+1} belongs to U, V_i if x_{n+1} belongs to V_i, we have that, according to the pointer invariants, x_{n_0+1} is the leftmost variable of W and for each $j < d$, $x_{n_{j+1}+1}$ is argument of x_{n_j+1}. For each $j = 0, \ldots, d$ define β_j (resp. α_j) to be the output (resp. input) atom of x_{n_j+1} in W. Furthermore, for $j > 0$ set β_j to point on α_{j-1} and α_j to point on β_k iff x_{n_j} points on x_{n_k} in s.

Lemma 7 *The sequence $p = \beta_0, \alpha_0, \ldots, \beta_d, \alpha_d$ is a P-view in σ_W. Moreover $\mathcal{O}_W(p) = x_{n+1}$.*

Proof. From the fact that the x_{n_j}'s form a chain of argument W-occurrences together with the definition of σ_W one gets that p is in σ_W. A reasoning similar to the one of lemma 6 shows that $\mathcal{O}_W(p) = x_{n+1}$. □

Theorem 8 *If s is the HO-dialog of length n between σ and the σ_i's then $\overline{\mathcal{O}}_T(s)$ is the T-run of length n of the* PAM. *Conversely, if s is the T-run of length n then $\mathcal{D}_T(s)$ is the HO-dialog of length n between σ and the σ_i's.*

Proof. Let $s = \beta_0, \alpha_1, \ldots, \gamma_n$ be the HO-dialog of length n between σ and the σ_i's and let $x_0, \ldots x_n$ be $\overline{\mathcal{O}}_T(s)$. We show that $\overline{\mathcal{O}}_T(s)$ is a T-run by induction on n.

If $n = 1$ note that β_0, α_1 is a P-view, which by definition of HO-dialogs belongs to σ, thus to σ_U. By lemma 6 this entails that β_0 and α_1 are respectively the output and input atoms of the leftmost U-occurrence x_1. Thus $\overline{\mathcal{O}}_T(s) = x_0, x_1$ where x_1 points to x_0, which by definition of the PAM is the T-run of length 1.

Suppose $n \geq 1$ and is even, so that γ_n is an opponent move β_n. Let $p = \mathcal{V}^P(s)$ so that σ is now to play the move α_{n+1} such that $p, \alpha_{n+1} \in \sigma$. By induction we may suppose that $\overline{\mathcal{O}}_T(s)$ is a T-run. Let x_{n+1} be the pointing occurrence that the PAM is to play and $x_{n_0+1}, \ldots, x_{n_d+1}$ be the view of x_{n+1} in $\overline{\mathcal{O}}_T(s)$. Now apply lemma 7 to s, x_{n+1} getting a P-view p' in σ_U (since n is even) such that $\mathcal{O}_U(p') = x_{n+1}$. Therefore $p' = p, \alpha_{n+1}$ which was what we needed to show.

The converse is much in the same vein and is left to the reader. □

6 Another linear head reduction machine: the IAM

We give a presentation of the IAM as a *bideterministic automaton* acting on a set of *moves*. The bideterminicity means that the automaton is deterministic and that if one decides to move back at some point then he has no choice but undoing what he had done to reach this point, eventually reaching his starting point. The underlying graph of the automaton is the net corresponding to the term the IAM is executing and the transitions are given as actions on moves.

6.1 Nets.

Nets are oriented graphs with nodes of given *arity* (number of incident edges or premisses), and *co-arity* (number of emergent edges or conclusions) which are a or axiom $(0,2)$, cut $(2,0)$, $\otimes$ and $\wp$ $(2,1)$, weakening $(0,1)$, d or dereliction $(1,1)$, p or auxiliary door $(1,1)$, c or contraction $(2,1)$ and ! or principal door $(1,1)$. A type is attached to each edge of a net in the appropriate way, *e.g.*, the conclusion of a dereliction node has type $?A$ when the premise has type A. We will say that a node of coarity 1 has type A when its conclusion has type A.

Each ! node defines a *box*, that is a subnet of which all emergent edges are premisses of auxiliary doors, except one which points to that ! node. Boxes are either disjoint or one in the other. The *depth* of a node is the number of boxes it is in.

In general a term in net-form will have one dangling edge for each free variable plus one special edge corresponding to its head-variable. The type attached to the head variable edge is by definition the type of the net.

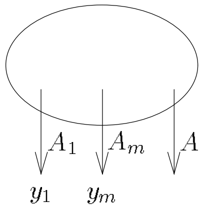

Figure 1: Principle

NORMAL TERMS AS NETS. Let $U = \lambda x_1 ... \lambda x_n\ (x)U_1 ... U_p$ be normal, with head variable x of input and output atoms α and β. Then its net representation is as in figure 2, where edges dangling from auxiliary doors, corresponding to free variables of the U_i's head-bound in U (that is which are among $x_1, \ldots, x_n$) have to be connected through contraction nodes to the nodes above the $\wp$-branch corresponding to their respective binders (*e.g.*, on the figure we supposed that U_p has x_n as free variable and that the leftmost occurrence x of U is x_1). The other dangling edge here, the head variable edge, will always conclude a $\wp$-branch, that is a sequence of connected $\wp$ nodes. There is exactly one $\otimes$ (resp. $\wp$) -branch (the dashed lines on figure 2) associated to each occurrence of variable, corresponding to its arguments (resp. head binders). So that a view uniquely determines an alternate sequence of $\wp$'s and $\otimes$'s, namely the roots of these branches. This remark is one of the key feature that is used by Lamarche in his own definition of games, which is very close to HO one [7].

We only deal now with nets representing terms. As above, each axiom node will be associated to a unique occurrence of variable and be labeled by its input and output atoms, by which one can retrieve a unique formula F for each node. More precision on the relations between nets and lambda-calculus may be found in [4, 10].

6.2 Nets as automatons

T-MOVES. Let T be a net. A *T-move*, in a node $\mathtt{n}$ of T labeled by a formula F, consists in a triple $\langle \gamma, \vec{j}_1, \vec{j}_2 \rangle$ where γ is an atom of F and $\vec{j}_1$, $\vec{j}_2$ are vectors of integers of respective size the depth of $\mathtt{n}$ in T and the depth of γ in F. The T-move is *ascending* (*descending*) if $\gamma > 0$ ($\gamma < 0$). We denote by ϵ the vector of zero size and by $\vec{j}.\vec{j}'$ the vector obtained by concatenation of $\vec{j}$ and $\vec{j}'$. Two noteworthy formats of T-moves are the *high* or *private format* when γ is the terminal atom of F: $\langle \gamma, \vec{j}, \epsilon \rangle$, and the *low* or *public format* when depth of $\mathtt{n}$ is zero: $\langle \gamma, \epsilon, \vec{j} \rangle$.

EDGES AS ACTIONS ON T-MOVES. Let a bijection $\lceil .,. \rceil$ from couples of integers to integers, and injections ρ and σ from integers to integers with disjoint codomains be given.

To each edge of a net one associates a partial action on T-moves depending on the node of which it is a premise. Just as a finite automaton acts on words. All these actions are reversible and the reverse action, is associated to the ... reverse edge.

Premisses of $\otimes$ and $\wp$ nodes leave the T-move intact; crossing an axiom node switches the output atom to the input one; crossing a cut node switches the atom to its dual; the dereliction premise maps $\langle \alpha, \vec{j}_1, \vec{j}_2 \rangle$ to $\langle \alpha, \vec{j}_1, 0.\vec{j}_2 \rangle$; the contraction left (right) premise maps $\langle \alpha, \vec{j}_1, i.\vec{j}_2 \rangle$ to $\langle \alpha, \vec{j}_1, \rho(i).\vec{j}_2 \rangle$ ($\langle \alpha, \vec{j}_1, \sigma(i).\vec{j}_2 \rangle$); the auxiliary door maps $\langle \alpha, \vec{j}_1.i, i'.\vec{j}_2 \rangle$ to $\langle \alpha, \vec{j}_1, \lceil i, i' \rceil.\vec{j}_2 \rangle$; the ! (or principal door) maps $\langle \alpha, \vec{j}_1.i, \vec{j}_2 \rangle$ to $\langle \alpha, \vec{j}_1, i.\vec{j}_2 \rangle$. See figure 3 for a brief summary.

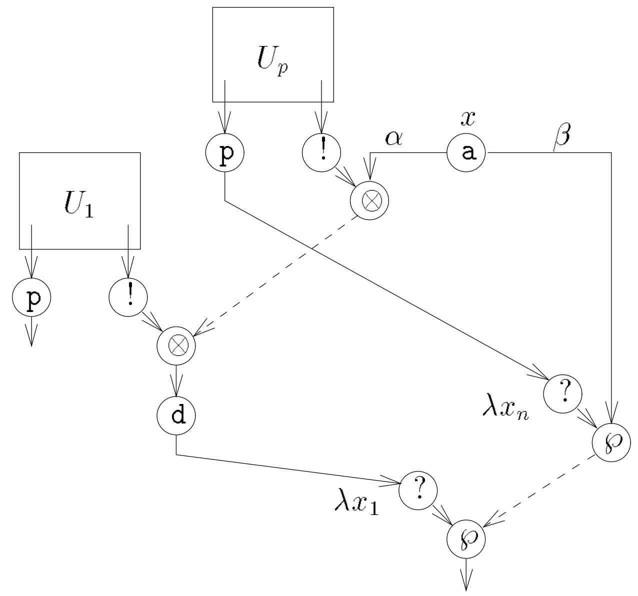

Figure 2: Net form of terms.

Note that the p and ! actions are not defined unless the first component, $\vec{j}_1$, is non-empty. The notation $_{x/y}$ means that only x or y figures, as appropriate, *e.g.*, when moving downard the left edge of a $\otimes$ or $\wp$ link (m on the figure) with a T-move $\langle \gamma, \vec{j}_1, \vec{j}_2 \rangle$ one gets the T-move $\langle \gamma, \vec{j}_1, \vec{j}_2 \rangle$ in the conclusion. Note that since an atom cannot be subatom of both premises of a multiplicative link, the upward move in the m node is deterministic. Also, since we supposed ρ and σ with disjoint codomains, the upward move in a contraction link c is deterministic. This, together with the fact that all transitions are injective on T-moves, shows that the IAM is bideterministic, as announced. Also one easily checks that the actions defined preserve the format of T-moves.

6.3 The Interaction Abstract Machine

IAM RUNS. As in the PAM we will suppose $T = (U)V_1 \ldots V_k$, where U and the V_i's are normal. The net form of T, which is not normal, is obtained by cutting U's net form against x_0's dangling edge in the net-form of $(x_0)V_1 \ldots V_k$, where x_0 is a new occurrence corresponding to the fake occurrence introduced in the PAM (see figure 4).

The *starting T-move* is $\langle \beta, \epsilon, \epsilon \rangle$ in T's head variable edge, where β is x_0's output atom. The *T-run* is by definition the sequence of T-moves obtained by letting edges act upon the starting T-move. This sequence is unique since, as said, T is a deterministic atomaton. Less obviously:

THEOREM 9 **(correction of the IAM)** *Let $s = x_0, \ldots, x_n$ be the sequence of occurrences associated to axiom nodes visited during the T-run produced by the* IAM *on T, then it matches the sequence of linear head reducts of T in the sense that was defined in theorem 1.*

Proof in the last section. That same machine was proved to be a "reversibilization" of both the KAM and the PAM in [5] and was extended to PCF by Mackie in [8].

7 AJM-games

We start by a short introduction to AJM games and strategies adapted to typed λ-calculus, details in [2].

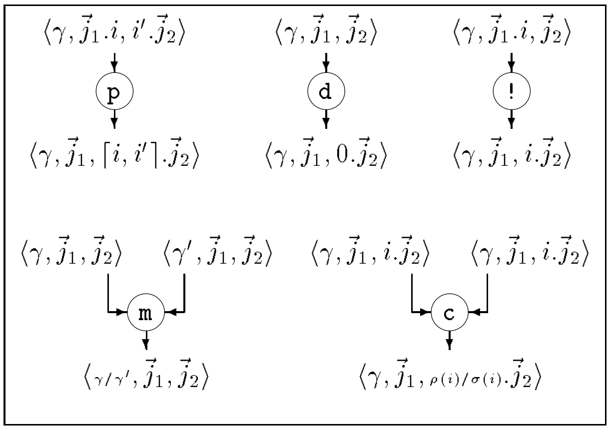

Figure 3: Summary of transitions.

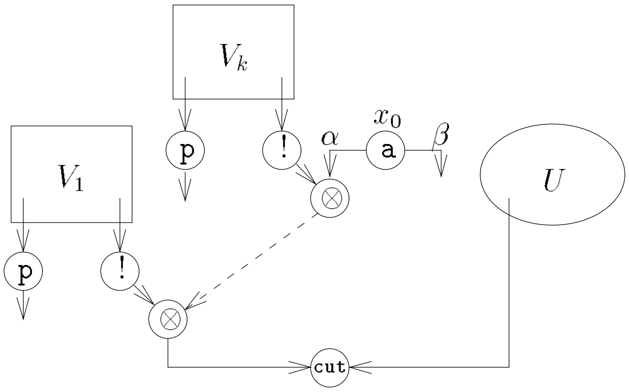

Figure 4: Net form of T

A-Moves and Positions. Let A be a type. An *A-move* is a couple $\langle\gamma,\vec{j}\rangle$, where γ is an atom of A and $\vec{j}$ is a vector of integers of size the *depth* of γ in A. It is said positive (or an O-move) if γ is positive in A, negative (or a P-move) else. Note that a T-move in a node $\mathtt{n}$ of type F in a term T of type A has the format $\langle\gamma,\vec{j},\vec{j}'\rangle$ where $|\vec{j}'|$ is the depth of γ in F and $|\vec{j}|$ is the depth of $\mathtt{n}$ in T. Henceforth $\langle\gamma,\vec{j}'\rangle$ is an F-move. In particular, if one chooses the head variable edge of T which corresponds to the type A, the T-move is in public format $\langle\gamma,\epsilon,\vec{j}\rangle$ and may therefore be identified with the A-move $\langle\gamma,\vec{j}\rangle$.

It is useful to think of A-moves as addresses in the tree of the formula A. More precisely, one expands the formula A with binary nodes in multiplicative connectives and with infinitary (but denumerably) nodes in exponential connectives (as if the ! was an infinite $\otimes$). Then the sequence $\vec{j}$ in the A-move $\langle\gamma,\vec{j}\rangle$ is the sequence of addresses in the exponential nodes, while the atom γ serves to choose the side in binary (multiplicative) nodes.

An *A-position* s is an alternating sequence of A-moves, that is a sequence of A-moves of alternating signs.

Valid Positions. Let F be a subtype of A and $\vec{j}'$ be a vector of integers of size the *depth* of F in A. The *projection* of an A-move $\langle\gamma,\vec{j}\rangle$ along $\langle F,\vec{j}'\rangle$, denoted by $\langle\gamma,\vec{j}\rangle|_{\langle F,\vec{j}'\rangle}$, is defined iff γ is in F and $\vec{j}=\vec{j}'.\vec{j}''$ for some $\vec{j}''$; it then is equal to $\langle\gamma,\vec{j}''\rangle$ and is an F-move (this because $|\vec{j}'|+|\vec{j}''|=|\vec{j}|=d(\gamma,A)=d(\gamma,F)+d(F,A)$, and since $|\vec{j}'|=d(F,A)$ we indeed get $|\vec{j}''|=d(\gamma,F)$). The projection of an A-position s along $\langle F,\vec{j}'\rangle$, likewise denoted by $s|_{\langle F,\vec{j}'\rangle}$, is the subsequence of projected A-moves, it then is an F-position. An A-position is *valid* if all its projections are alternating and begin by a positive A-move. So that projections of a valid position are valid themselves.

POINTIFIXION. Let s be a non-empty valid position, then the first A-move of s is the terminal atom of A, and no A-move can occur twice in s.

For the first point, if the first A-move writes $\langle\gamma, i.\vec{j}\rangle$ it can be projected along $\langle A_\alpha, i\rangle$ where α is the unique immediate subatom of β containing γ (recall that A_α is the biggest subtype of A having α for terminal atom). So that $\langle\gamma, \vec{j}\rangle$ is the first A-move of $s|_{\langle A_\alpha, i\rangle}$, but then it is positive in A_α, hence $\langle\gamma, i.\vec{j}\rangle$ is negative in A, so s is not valid. For the second point, A-moves in $s|_{\langle\gamma,\vec{j}\rangle}$ are but γ's, since this sequence is alternating there is at most one A-move in it, so no $\langle\gamma, \vec{j}\rangle$ occurs twice.

Let now $m = \langle\gamma, \vec{j}.i\rangle$ be in the valid position s, then there is a unique A-move denoted $\theta_s(m)$ which has the form $\theta_s(m) = \langle\gamma', \vec{j}\rangle$ in s, where γ is an immediate subatom of γ'. As said there can be at most one such A-move. Since $s|_{\langle A_{\gamma'}, \vec{j}\rangle}$ contains the projection of $\langle\gamma, \vec{j}.i\rangle$, namely $\langle\gamma, i\rangle$, it begins, as said, with the terminal atom of $A_{\gamma'}$, that is γ'. This γ' is then the projection of the announced $\langle\gamma', \vec{j}\rangle$ in s. Now let $P(s)$, the *pointifixion* of s, be the pointing sequence consisting in the atoms of (the A-move in) s equipped with the partial pointing function θ_s just defined. Clearly,

LEMMA 10 *The pointifixion of s satisfies the alternating and justification condition.*

As this example of McCusker (private communication) shows $P(s)$ may not satisfy the visibility condition: take $s = \langle\beta_0, \epsilon\rangle$, $\langle\alpha_0, 0\rangle$, $\langle\beta_1, 00\rangle$, $\langle\alpha_0, 1\rangle$, $\langle\beta_1, 10\rangle$, $\langle\alpha_1, 000\rangle$, $\langle\beta_1, 01\rangle$, $\langle\alpha_1, 100\rangle$ which is a valid position in $((\alpha_1 \to \beta_1) \to \alpha_0) \to \beta_0$.

EQUIVALENCE. Say two positions are *equivalent* if they map to the same pointing sequence.

AJM STRATEGIES. A negative (or player) pre-strategy is a tree of valid positions, branching on all (validity-preserving) O-moves.

A such negative pre-strategy σ is *self-equivalent* if for any two equivalent positions ending with O-moves it contains, it can extend either none or both, and if both, extended positions are still equivalent.

A such negative pre-strategy σ is *history-free* if for any two positions it contains ending with the same O-move, it can extend either none or both, and if both, by the same move.

A negative *strategy* is a self-equivalent and history-free negative pre-strategy.

Positive (or opponent) pre-strategies and strategies are defined symmetrically.

AJM-DIALOGS. Let σ be a negative strategy and τ be a positive one, both in A, we denote by $\sigma\tau$ the position (or play, or dialog) they generate, *i.e.*, the opening move is A's terminal atom, then σ plays, then τ and so on.

8 AJM-games and λ-calculus

8.1 Interpretation of terms

Let U be a normal term of type A, and f_U be the mapping of O-moves to P-moves defined by $f_U(\langle\beta, \vec{j}\rangle) = \langle\alpha, \vec{j}'\rangle$ iff U's net-form starting with the ascending T-move $\langle\beta, \epsilon, \vec{j}\rangle$ in the head-variable edge ends with the descending T-move $\langle\alpha, \epsilon, \vec{j}'\rangle$ in that same edge. That function clearly generates a set of positions. Better:

THEOREM 11 *For all U, f_U generates a strategy, denoted by σ^d_U.*

Baillot proves this in [3] and also that σ^d_U is the actual interpretation of terms used in AJM model.

8.2 AJM games and the IAM

Obviously, a family σ^d_i of negative strategies in A_i can be resized into a positive strategy in $A_1 \ldots A_k \to \beta$. Given $V_1 : A_1, \ldots, V_k : A_k$, we denote by $\sigma^d_{\vec{V}}$ the corresponding positive strategy. Note that conversely a positive strategy in $A_1 \ldots A_k \to \beta$ may be downsized to a family of negative strategies in A_i's.

THEOREM 12 *For all $T = (U)V_1 \ldots V_k$, where U and the V_i's are normal, $\sigma^d_U \sigma^d_{\vec{V}}$ is the subsequence of* IAM*'s T-run beginning with the starting move and consisting thereafter only in moves in the cut node.*

9 AJM-games and HO-games

In this section σ^d and σ^a will respectively denote AJM and HO strategies (d for "digital", a for "analogic"). Furthemore we shall suppose that HO strategies are given extensionally. Recall that by this we mean that they are presented as set of plays, not as set of views.

Define $P(\sigma^d)$ to be the set of $P(s)$ that satisfy the visibility condition (because, as noticed before, AJM-opponents are wilder than HO-opponents). Thus, the pointifixion map extends to a map from pre-strategies to trees of HO plays. If a pre-strategy σ^d is self-equivalent then $P(\sigma^d)$ is branching only on positive moves. Better, if σ^d is also history-free, then $P(\sigma^d)$ is an HO-strategy given extensionally (as a tree of plays).

THEOREM 13 **(pointifixion 1)** *For all strategy σ^d, $P(\sigma^d)$ is an HO-strategy.*

Proof. To prove this, we first need a lemma. Let s be a valid position of which the last move is positive, and define $\mathcal{V}'(s)$ to be the subsequence that P maps to $\mathcal{V}(P(s))$.

LEMMA 14 **(digital innocence)** *For all strategy σ^d, if $s \in \sigma^d$ then $\mathcal{V}'(s) \in \sigma^d$.*

By induction on s, if the last A-move in s "points" somewhere, then $s = s', \langle\alpha, \vec{j}\rangle, s'', \langle\beta, \vec{j}.i\rangle$ where s' ends with a O-move (hence can't be empty). By induction $\mathcal{V}'(s') \in \sigma^d$, on the other hand $s', \langle\alpha, \vec{j}\rangle \in \sigma^d$, hence $\mathcal{V}'(s')$ ending with the same O-move than s' and σ^d being history-free, we have $\mathcal{V}'(s'), \langle\alpha, \vec{j}\rangle \in \sigma^d$. But then $\mathcal{V}'(s) = \mathcal{V}'(s'), \langle\alpha, \vec{j}\rangle, \langle\beta, \vec{j}.i\rangle$ is valid (and thus belongs to σ^d) since any projection that keeps the last A-move $\langle\beta, \vec{j}.i\rangle$, either also keeps $\langle\alpha, \vec{j}\rangle$ in which case alternation is preserved, or is reduced to β. If the last A-move in s doesn't point anywhere, then since s is valid it has to be reduced to $\langle\beta, \epsilon\rangle$, and $\mathcal{V}'(\langle\beta, \epsilon\rangle) = \langle\beta, \epsilon\rangle$.

Now for the theorem, consider any pointing sequence $P(s)$ in $P(\sigma^d)$ of which the last A-move is positive. Since $\mathcal{V}'(s) \in \sigma^d$ by the lemma, $\mathcal{V}'(s)$'s next A-move in σ^d "points" to some A-move in $\mathcal{V}'(s)$ and because σ^d is history-free, it plays the same next A-move in s, which therefore "points" in $\mathcal{V}'(s)$. So P-moves in any $P(s)$ point in their views. As to whether they are determined by their views, consider any $P(s)$ and $P(s')$ in $P(\sigma^d)$ of which the last A-moves are positive. If those A-moves have the same view, $\mathcal{V}(P(s)) = \mathcal{V}(P(s'))$, then $\mathcal{V}'(s)$ and $\mathcal{V}'(s')$ are equivalent so that, since σ^d is self-equivalent, both A-moves point in $\mathcal{V}(P(s)) = \mathcal{V}(P(s'))$ to the same A-move, and since σ^d is history-free, it is just the same in $P(s)$ and $P(s')$. □

Let T be a term, any view p in σ^a_T is "digitalized" in σ^d_T, that is to say:

LEMMA 15 **(digital view)** *For all view $p \in \sigma^a_T$ there is an $s \in \sigma^d_T$ such that $P(s) = p$.*

Before coming into the proof, let us note that this is the very moment where we have to use the definition of nets and GoI. More precisely, the careful reader may check that all that we need to prove the digital view lemma is that the IAM acts as a bideterministic automaton on T-moves.

Proof. Let $p = \beta_0, \alpha_0, \ldots, \beta_n, \alpha_n$, with by definition $\beta_{i+1} \prec \alpha_i$.

Recall p can be mapped to an alternate sequence of $\wp$'s and $\otimes$'s nodes in T's net-form, say $\wp_0\otimes_0 \ldots \wp_n\otimes_n$. Let φ_i denote the action on $\vec{j}_1, \vec{j}_2$ of the descent from $\otimes_i$ to the conclusion above which $\otimes_i$ stands and let $\vec{0}_i$ denote the vector of zero's of size i. Then we claim that:

$$\begin{aligned} s \quad = \quad & \langle\beta_0, \epsilon\rangle, \langle\alpha_0, \varphi_0(\epsilon, \epsilon)\rangle, \ldots, \langle\alpha_i, \varphi_i(\vec{0}_i, \epsilon)\rangle, \\ & \langle\beta_{i+1}, \varphi_i(\vec{0}_i, \epsilon).0\rangle, \ldots \end{aligned}$$

belongs to σ^d_T; or, in other words, that there is a position in σ^d_T such that successive P-moves read in private format are but $\langle\alpha_i, \vec{0}_i, \epsilon\rangle$. Note that the depth of $\otimes_i$ in T, not to be confused with the depth of α_i in A, is i which is the size of $\vec{0}_i$ so that s is well-defined.

Suppose the last P-move was $\langle\alpha_i, \varphi_i(\vec{0}_i, \epsilon)\rangle$, then opponent may play $\langle\beta_{i+1}, \varphi_i(\vec{0}_i, \epsilon).0\rangle$ since a such move will either be alone in a given projection, or be accompanied therein by its predecessor $\langle\alpha_i, \varphi_i(\vec{0}_i, \epsilon)\rangle$, and thus its addition preserves validity.

Let's see now how this O-move will be acted upon by T's net-form. It is an ascending move, which will climb back up to $\otimes_i$ and there simply be $\langle\beta_{i+1}, \vec{0}_i, 0\rangle$, then higher in $\wp_{i+1}$ it will be $\langle\beta_{i+1}, \vec{0}_{i+1}, \epsilon\rangle$, will go through the axiom and down the $\otimes$-branch and become $\langle\alpha_{i+1}, \vec{0}_{i+1}, \epsilon\rangle$ in $\otimes_{i+1}$.

Whence we see the claim is correct, that is $s \in \sigma^d_T$.

Now let's follow the move further, it goes down the exponential branch below $\otimes_{i+1}$ and connects to say $\wp_l$ with $l \leq i+1$ where it is $\langle \alpha_{i+1}, \vec{0}_l, m \rangle$ for some integer m. The positive move in $\wp_l$ was, as said, $\langle \beta_l, \vec{0}_l, \epsilon \rangle$ so that their images by φ_l will be such that α_{i+1}'s one "points" to β_l's one, whence one sees that $P(s) = p$. $\square$

THEOREM 16 **(pointifixion 2)** *For all T, $\sigma^a_T = P(\sigma^d_T)$.*

We know first (Baillot's theorem) that σ^d_T is a strategy, then by the first pointifixion theorem $P(\sigma^d_T)$ is a strategy, and since by the previous lemma it coincides with σ^a_T on views, it is it.

It is not obvious that for all σ^a there exists a σ^d that pointifies to it, because history-freedom at first sight seems quite a formidable property, which maybe no strategies have. The geometry of interaction provides a solution.

THEOREM 17 **(AJM definability)** *Any strategy σ^d is equivalent to a term-strategy.*

Again a strategy σ^d maps to an HO strategy, by the first pointifixion theorem, which is definable by HO definability result, that is which is a σ^a_T for some possibly partial and infinite T. Whence by the theorem above $P(\sigma^d) = P(\sigma^d_T)$, that is σ^d is equivalent to σ^d_T.

THEOREM 18 **(pointifixion 3)** *Let $(U)\vec{V}$ be a term, then $\sigma^a_U \sigma^a_{\vec{V}} = P(\sigma^d_U \sigma^d_{\vec{V}})$.*

Say U is to play in the AJM play s, by the digital innocence lemma $\mathcal{V}'(s) \in \sigma^d_U$ and σ^d_U moves in $\mathcal{V}'(s)$ as in s. So its next move will digitalize the correct HO-move by the digital view lemma and σ^d_U's self-equivalence. Same reasoning if $\vec{V}$ is to play. From this one also sees the IAM is correct.

References

[1] Martin Abadi, Pierre-Louis Curien, and Jean-Jacques Lévy. Explicit substitutions. In *Seventh Annual ACM Symposium on Principles of Programming Languages*, pages 31–46. ACM Press, 1990.

[2] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. Draft, available by `ftp` on `theory.doc.ic.ac.uk`, 1994.

[3] Patrick Baillot. AJM-strategies and the geometry of interaction. Draft available by `ftp` on `ftp.logique.jussieu.fr`, 1995.

[4] Vincent Danos. *Une Application de la Logique Linéaire à l'Étude des Processus de Normalisation (principalement du λ-calcul).* Thèse de doctorat, Université Paris 7, 1990.

[5] Vincent Danos and Laurent Regnier. Reversible, irreversible and optimal lambda-machines. To appear in ENTCS, 1996.

[6] Martin Hyland and Luke Ong. On full abstraction for PCF. Draft available by `ftp` on `theory.doc.ic.ac.uk`, 1994.

[7] François Lamarche. Games semantics for full propositional linear logic. In *Proceedings of the tenth Annual IEEE Symposium on Logic in Computer Science, San Diego*, pages 464–473, 1995.

[8] Ian Mackie. The geometry of implementation. In *Proceedings of the 22^{th} ACM Symposium on Principles of Programming Languages*, 1995.

[9] Gianfranco Mascari and Marco Pedicini. Head linear reduction and pure proof net extraction. In *Selected papers of the Meeting on the Mathematical Foundations of Programming Semantics (MFPS'92, Oxford)*, volume 135, pages 111–137. Theoretical Computer Science, 1994.

[10] Laurent Regnier. *Lambda-Calcul et Réseaux.* Thèse de doctorat, Université Paris 7, 1992.

[11] Laurent Regnier. Une équivalence sur les lambda-termes. *Theoretical Computer Science*, 126, 1994.