

Least and Greatest Fixpoints in Game Semantics

Pierre CLAIRAMBAULT
pclairam@pps.jussieu.fr

PPS — Université Paris 7

Abstract. We show how solutions to many recursive arena equations can be computed in a natural way by allowing loops in arenas. We then equip arenas with winning functions and total winning strategies. We present two natural winning conditions compatible with the loop construction which respectively provide initial algebras and terminal coalgebras for a large class of continuous functors. Finally, we introduce an intuitionistic sequent calculus, extended with syntactic constructions for least and greatest fixed points, and prove it has a sound and (in a certain weak sense) complete interpretation in our game model.

1 Introduction

The idea to model logic by game-theoretic tools can be traced back to the work of Lorenzen [21]. The idea is to interpret a formula by a game between two players O and P, O trying to refute the formula and P trying to prove it. The formula A is then valid if P has a *winning strategy* on the interpretation of A . Later, Joyal remarked [17] that it is possible to compose strategies in Conway games [8] in an associative way, thus giving rise to the first category of games and strategies. This, along with parallel developments in Linear Logic and Geometry of Interaction, led to the more recent construction of compositional game models for a large variety of logics [5,23,18] and programming languages [15,3,22,4].

We aim here to use these tools to model an intuitionistic logic with induction and coinduction. Inductive/coinductive definitions in syntax have been defined and studied in a large variety of settings, such as linear logic [6], λ -calculus [1] or Martin-Löf's type theory [9]. Motivations are multiple, but generally amount to increasing the expressive power of a language without paying the price of exponential modalities (as in [6]) or impredicativity (as in [1] or [9]). However, less work has been carried out when it comes to the semantics of such constructions. Of course we have the famous order-theoretic Knaster-Tarski fixed point theorem [25], the nice categorical theory due to Freyd [11], set-theoretic models [9] (for the *strictly positive* fragment) or PER-models [20], but it seems they have gone through the current trend of intensional models without notice. We fix this issue here, showing that (co)induction admits a nice game-theoretic model which arises naturally if one enriches McCusker's [22] work on recursive types with winning functions inspired by parity games [24].

In Section 2, we first recall the basic definitions of the Hyland-Ong-Nickau setting of game semantics. Then we sketch McCusker’s interpretation of recursive types, and show how most of these recursive types can be modelled by means of loops in the arenas. For this purpose, we define a class of functors called *open functors*, including in particular all the endofunctors built out of the basic type constructors. We also present a mechanism of winning functions inspired by [16], allowing us to build a category **Gam** of games and total winning strategies. In section 3, we present μLJ , the intuitionistic sequent calculus with least and greatest fixpoints that we aim to model. We briefly discuss its proof-theoretic properties, then present its semantic counterpart: we show how to build initial algebras and terminal coalgebras to most positive open functors. Finally, we use this semantic account of (co)induction to give a sound and (weakly) complete interpretation of μLJ in **Gam**.

2 Arena Games

2.1 Arenas and Plays

We recall briefly the now usual definitions of arena games, introduced in [15]. More detailed accounts can be found in [22,13]. We are interested in games with two participants: Opponent (O, the *environment*) and Player (P, the *program*). Possible plays are generated by directed graphs called *arenas*, which are semantic versions of *types* or *formulas*. Hence, a play is a sequence of *moves* of the ambient arena, each of them being annotated by a *pointer* to an earlier move — these pointers being required to comply with the structure of the arena. Formally, an **arena** is a structure $A = (M_A, \lambda_A, \vdash_A)$ where:

- M_A is a set of **moves**,
- $\lambda_A : M_A \rightarrow \{O, P\} \times \{Q, A\}$ is a **labelling** function indicating whether a move is an Opponent or Player move, and whether it is a question (Q) or an answer (A). We write λ_A^{OP} for the projection of λ_A to $\{O, P\}$ and λ_A^{QA} for its projection on $\{Q, A\}$. $\overline{\lambda_A}$ will denote λ_A where the $\{O, P\}$ part has been reversed.
- \vdash_A is a relation between $M_A + \{\star\}$ to M_A , called **enabling**, satisfying:
 - $\star \vdash m \implies \lambda_A(m) = OQ$;
 - $m \vdash_A n \wedge \lambda_A^{QA}(n) = A \implies \lambda_A^{QA}(m) = Q$;
 - $m \vdash_A n \wedge m \neq \star \implies \lambda_A^{OP}(m) \neq \lambda_A^{OP}(n)$.

In other terms, an arena is a directed bipartite graph, with a set of distinguished **initial** moves (m such that $\star \vdash_A m$) and a distinguished set of **answers** (m such that $\lambda_A^{QA} = A$) such that no answer points to another answer. We now define plays as **justified sequences** over A : these are sequences s of moves of A , each non-initial move m in s being equipped with a pointer to an earlier move n in s , satisfying $n \vdash_A m$. In other words, a justified sequence s over A is such that each reversed pointer chain $s_{\phi(0)} \leftarrow s_{\phi(1)} \leftarrow \dots \leftarrow s_{\phi(n)}$ is a path on A , viewed as a directed bipartite graph.

The role of pointers is to allow *reopenings* in plays. Indeed, a path on A may be (slightly naively) understood as a linear play on A , and a justified sequence as an interleaving of paths, with possible duplications of some of them. This intuition is made precise in [14]. When writing justified sequences, we will often omit the justification information if this does not cause any ambiguity. \sqsubseteq will denote the prefix ordering on justified sequences. If s is a justified sequence on A , $|s|$ will denote its length.

Given a justified sequence s on A , it has two subsequences of particular interest: the P-view and O-view. The view for P (resp. O) may be understood as the subsequence of the play where P (resp. O) only sees his own duplications. In a P-view, O never points more than once to a given P-move, thus he must always point to the previous move. Concretely, P-views correspond to branches of Böhm trees [15]. Practically, the P-view $\lceil s \rceil$ of s is computed by forgetting everything under Opponent's pointers, in the following recursive way:

- $\lceil sm \rceil = \lceil s \rceil m$ if $\lambda_A^{OP}(m) = P$;
- $\lceil sm \rceil = m$ if $\star \vdash_A m$ and m has no justification pointer;
- $\lceil s_1 m s_2 n \rceil = \lceil s \rceil mn$ if $\lambda_A^{OP}(n) = O$ and n points to m .

The O-view $\lfloor s \rfloor$ of s is defined dually. Note that in some cases — in fact if s does not satisfies the *visibility condition* introduced below — $\lceil s \rceil$ and $\lfloor s \rfloor$ may not be correct justified sequences, since some moves may have pointed to erased parts of the play. However, we will restrict to plays where this does not happen. The **legal sequences** over A , denoted by \mathcal{L}_A , are the justified sequences s on A satisfying the following conditions:

- **Alternation.** If $tmn \sqsubseteq s$, then $\lambda_A^{OP}(m) \neq \lambda_A^{OP}(n)$;
- **Bracketing.** A question q is **answered** by a if a is an answer and a points to q . A question q is **open** in s if it has not yet been answered. We require that each answer points to the *pending* question, *i.e.* the last open question.
- **Visibility.** If $tm \sqsubseteq s$ and m is not initial, then if $\lambda_A^{OP}(m) = P$ the justifier of m appears in $\lceil t \rceil$, otherwise its justifier appears in $\lfloor t \rfloor$.

2.2 The cartesian closed category of Innocent strategies

A **strategy** σ on A is a prefix-closed set of even-length legal plays on A . A strategy is **deterministic** if only Opponent branches, *i.e.* $\forall smn, smn' \in \sigma, n = n'$. Of course, if A represents a type (or formula), there are often many more strategies on A than programs (or proofs) on this type. To address this issue we need **innocence**. An innocent strategy is a strategy σ such that

$$sab \in \sigma \wedge t \in \sigma \wedge ta \in \mathcal{L}_A \wedge \lceil sa \rceil = \lceil ta \rceil \implies tab \in \sigma$$

We now recall how arenas and innocent strategies organize themselves into a cartesian closed category. First, we build the **product** $A \times B$ of two arenas A and B :

$$\begin{aligned} M_{A \times B} &= M_A + M_B \\ \lambda_{A \times B} &= [\lambda_A, \lambda_B] \\ \vdash_{A \times B} &= \vdash_A + \vdash_B \end{aligned}$$

We mention the empty arena $I = (\emptyset, \emptyset, \emptyset)$, which will be terminal for the category of arenas and innocent strategies. We mention as well the arena $\perp = (\bullet, \bullet \mapsto OQ, (\star, \bullet))$ with only one initial move, which will be a weak initial object. We define the **arrow** $A \Rightarrow B$ as follows:

$$\begin{aligned} M_{A \Rightarrow B} &= M_A + M_B \\ \lambda_{A \Rightarrow B} &= [\overline{\lambda_A}, \lambda_B] \\ m \vdash_{A \Rightarrow B} n &\Leftrightarrow \begin{cases} m \neq \star \wedge m \vdash_A n \\ m \neq \star \wedge m \vdash_B n \\ \star \vdash_B m \wedge \star \vdash_A n \\ m = \star \wedge \star \vdash_B n \end{cases} \end{aligned}$$

We define composition of strategies by the usual parallel interaction plus hiding mechanism. If A, B and C are arenas, we define the set of **interactions** $I(A, B, C)$ as the set of justified sequences u over A, B and C such that $u_{\uparrow A, B} \in \mathcal{L}_{A \Rightarrow B}$, $u_{\uparrow B, C} \in \mathcal{L}_{B \Rightarrow C}$ and $u_{\uparrow A, C} \in \mathcal{L}_{A \Rightarrow C}$. Then, if $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$, we define parallel interaction:

$$\sigma || \tau = \{u \in I(A, B, C) \mid u_{\uparrow A, B} \in \sigma \wedge u_{\uparrow B, C} \in \tau\}$$

Composition is then defined as $\sigma; \tau = \{u_{\uparrow A, C} \mid u \in \sigma || \tau\}$. It is associative and preserves innocence (a proof of these facts can be found in [15] or [13]). We also define the identity on A as the copycat strategy (see [22] or [13] for a definition) on $A \Rightarrow A$. Thus, there is a category **Inn** which has arenas as objects and innocent strategies on $A \Rightarrow B$ as morphisms from A to B . In fact, this category is cartesian closed, the cartesian structure given by the arena product above and the exponential closure given by the arrow construction. This category is also equipped with a weak coproduct $A + B$, which is constructed as follows:

$$\begin{aligned} M_{A+B} &= M_A + M_B + \{q, L, R\} \\ \lambda_{A+B} &= [\lambda_A, \lambda_B, q \mapsto OQ, L \mapsto PA, R \mapsto PA] \\ m \vdash_{A+B} n &\Leftrightarrow \begin{cases} m, n \in M_A \wedge m \vdash_A n \\ m, n \in M_B \wedge m \vdash_B n \\ m = \star \wedge n = q \\ (m = q \wedge n = L) \vee (m = q \wedge n = R) \\ (m = L \wedge \star \vdash_A n) \vee (m = R \wedge \star \vdash_B n) \end{cases} \end{aligned}$$

2.3 Recursive types and Loops

Let us recall briefly the interpretation of recursive types in game semantics, due to McCusker [22]. Following [22], we first define an ordering \leq on arenas as follows. For two arenas A and B , $A \leq B$ iff

$$\begin{aligned} M_A &\subseteq M_B \\ \lambda_A &= \lambda_B \upharpoonright_{M_A} \\ \vdash_A &= \vdash_B \cap (M_A + \{\star\} \times M_A) \end{aligned}$$

This defines a (large) dcpo, with least element I and directed sups given by the componentwise union. If $F : \mathbf{Inn} \rightarrow \mathbf{Inn}$ is a functor which is continuous with respect to \sqsubseteq , we can find an arena D such that $D = F(D)$ in the usual way by setting $D = \bigsqcup_{n=0}^{\infty} F^n(I)$. McCusker showed[22] that when the functors are *closed* (*i.e.* their action can be internalized as a morphism $(A \Rightarrow B) \rightarrow (FA \Rightarrow FB)$), and when they preserve inclusion and projection morphisms (*i.e.* partial copycat strategies) corresponding to \sqsubseteq , this construction defines *parametrized minimal invariants* [11]. Note that the crucial cases of these constructions are the functors built out of the product, sum and function space constructions.

We give now a concrete and new (up to the author's knowledge) description of a large class of continuous functors, that we call **open functors**. These include all the functors built out of the basic constructions, and allow a rereading of recursive types, leading to the model of (co)induction.

Open arenas. Let \mathbf{T} be a countable set of names. An **open arena** is an arena A with distinguished question moves called **holes**, each of them labelled with an element of \mathbf{T} . We denote by \square_X the holes annotated by $X \in \mathbf{T}$. We will sometimes write \square_X^+ to denote a hole of Player polarity, or \square_X^- to denote a hole of Opponent polarity. If A has holes labelled by X_1, \dots, X_n , we denote it by $A[X_1, \dots, X_n]$. By abuse of notation, the corresponding open functor we are going to build will be also denoted by $A[X_1, \dots, X_n] : (\mathbf{Inn} \times \mathbf{Inn}^{op})^n \rightarrow \mathbf{Inn}$.

Image of arenas. If $A[X_1, \dots, X_n]$ is an open arena, If B_1, \dots, B_n and B'_1, \dots, B'_n are arenas (possibly open as well), we build a new arena $A(B_1, B'_1, \dots, B_n, B'_n)$ by replacing each occurrence of $\square_{X_i}^+$ by B_i and each occurrence of $\square_{X_i}^-$ by B'_i . More formally:

$$\begin{aligned}
M_{A(B_1, B'_1, \dots, B_n, B'_n)} &= (M_A \setminus \{\square_{X_1}, \dots, \square_{X_n}\}) + \sum_{i=1}^n (M_{B_i} + M_{B'_i}) \\
\lambda_{A(B_1, B'_1, \dots, B_n, B'_n)} &= [\lambda_A, \lambda_{B_1}, \overline{\lambda_{B'_1}}, \dots, \lambda_{B_n}, \overline{\lambda_{B'_n}}] \\
m \vdash_{A(B_1, B'_1, \dots, B_n, B'_n)} p &\Leftrightarrow \begin{cases} m \vdash_A \square_{X_i}^+ \wedge \star \vdash_{B_i} p \\ m \vdash_A \square_{X_i}^- \wedge \star \vdash_{B'_i} p \\ \star \vdash_{B_i} m \wedge \square_{X_i}^+ \vdash_A p \\ \star \vdash_{B'_i} m \wedge \square_{X_i}^- \vdash_A p \\ m \vdash_{B_i} p \\ m \vdash_{B'_i} p \\ m \vdash_A p \end{cases}
\end{aligned}$$

Note that in this definition, we assimilate all the moves sharing the same hole *label* \square_{X_i} and with the same polarity. This helps to clarify notations, and is justified by the fact that we never need to distinguish moves with the same hole label, apart from when they have different polarity.

Image of strategies. If A is an arena, we will, by abuse of notation, denote by I_A both the set of initial moves of A and the subarena of A with only these moves.

Let $A[X_1, \dots, X_n]$ be an open arena, $B'_1, B_1, \dots, B'_n, B_n$ and $C'_1, C_1, \dots, C'_n, C_n$ be arenas. Consider the application ξ defined on moves as follows:

$$\xi(x) = \begin{cases} \square_{X_i} & \text{if } x \in \bigcup_{i \in \{1, \dots, n\}} (I_{B'_i} \cup I_{B_i} \cup I_{C'_i} \cup I_{C_i}) \\ x & \text{otherwise} \end{cases}$$

and then extended recursively to an application ξ^* on legal plays as follows:

$$\xi^*(sa) = \begin{cases} \xi^*(s) & \text{if } a \text{ is a non-initial move of } B_i, B'_i, C_i \text{ or } C'_i \\ \xi^*(s)\xi(a) & \text{otherwise} \end{cases}$$

ξ^* erases moves in the inner parts of B'_i, B_i, C'_i, C_i and agglomerates all the initial moves back to the holes. This way we will be able to compare the resulting play with the identity on $A[X_1, \dots, X_n]$. Now, if $\sigma_i : B_i \rightarrow C_i$ and $\tau_i : C'_i \rightarrow B'_i$ are strategies, we can now define the action of open functors on them by stating:

$$s \in A(\sigma_1, \tau_1, \dots, \sigma_n, \tau_n) \Leftrightarrow \begin{cases} \forall i \in \{1, \dots, n\}, s \upharpoonright_{B_i \Rightarrow C_i} \in \sigma_i \\ \forall i \in \{1, \dots, n\}, s \upharpoonright_{C'_i \Rightarrow B'_i} \in \tau_i \\ \xi^*(s) \in id_{A[X_1, \dots, X_n]} \end{cases}$$

Proposition 1 *For any $A[X_1, \dots, X_n]$, this defines a functor $A[X_1, \dots, X_n] : (\mathbf{Inn} \times \mathbf{Inn}^{op})^n \rightarrow \mathbf{Inn}$, which is continuous with respect to \trianglelefteq .*

Proof (Proof sketch). Preservation of identities and composition are rather direct. A little care is needed to show that the resulting strategy is innocent: this relies on two facts: First, for each Player move the three definition cases are mutually exclusive. Second, a P-view of $s \in A(\sigma_1, \tau_1, \dots, \sigma_n, \tau_n)$ is (essentially) an initial copycat appended with a P-view of one of σ_i or τ_i , hence the P-view of s determines uniquely the P-view presented to one of σ_i, τ_i or $id_{A[X_1, \dots, X_n]}$.

Example. Consider the open arena $A[X] = \square_X \Rightarrow \square_X$. For any arena B , we have $A(B) = B \Rightarrow B$ and for any $\sigma : B_1 \rightarrow C_1$ and $\tau : C_2 \rightarrow B_2$, we have $A(\sigma, \tau) = \tau \Rightarrow \sigma : (B_2 \Rightarrow B_1) \rightarrow (C_2 \Rightarrow C_1)$, the strategy which precomposes its argument by τ and postcomposes it by σ .

Loops for recursive types. All these open functors are continuous, therefore having a parametrized minimal invariant obtained by the infinite expansion of the arena. However, for a large subclass of the open functors, this parametrized minimal invariant can be expressed in a simple way by replacing holes with a loop up to the initial moves. Suppose $A[X_1, \dots, X_n]$ is an open functor, and i is such that \square_{X_i} appears only in non-initial, positive positions in A . Then we define an arena $\mu_{X_i}.A$ as follows:

$$\begin{aligned} M_{\mu_{X_i}.A} &= (M_A \setminus \square_{X_i}) \\ \lambda_{\mu_{X_i}.A} &= \lambda_A \upharpoonright_{M_{\mu_{X_i}.A}} \\ m \vdash_{\mu_{X_i}.A} n &\Leftrightarrow \begin{cases} m \vdash_A n \\ m \vdash_A \square_{X_i} \wedge \star \vdash_A n \end{cases} \end{aligned}$$

A simple argument ensures that the obtained arena is isomorphic to the one obtained by iteration of the functor. For this issue we take inspiration from Laurent [19] and prove a theorem stating that two arenas are isomorphic in the categorical sense if and only if their set of paths are isomorphic. A **path** in A is a sequence of moves a_1, \dots, a_n such that for all $i \in \{1, \dots, n-1\}$ we have $a_i \vdash_A a_{i+1}$. A **path isomorphism** between A and B is a bijection ϕ between the set of paths of A and the set of paths on B such that for any non-empty path p on A , $\phi(ip(p)) = ip(\phi(p))$ (where $ip(p)$ denotes the immediate prefix of p). We have then the theorem:

Theorem 1 *Let A and B be two arenas. They are categorically isomorphic if and only if there is a path isomorphism between their respective sets of paths.*

Now, it is clear by construction that, if $A[X]$ is an open functor, the set of paths of $D = \bigsqcup_{n=0}^{\infty} A^n(I)$ and of $\mu X.A$ are isomorphic, therefore when $A[X]$ is closed and preserves inclusions and projections, $\mu X.A$ defines as well a parametrized minimal invariant for $A[X]$. But in fact, we have the following fact:

Proposition 2 *If $A[X]$ is an open functor, then it is closed and preserves inclusions and projections. Hence $\mu X.A$ is a parametrized minimal invariant for $A[X]$.*

This interpretation of recursive types as loops preserves finiteness of the arena, and as we shall see, allows to easily express the winning conditions necessary to model induction and coinduction.

2.4 Winning and Totality

A **total strategy** on A is a strategy $\sigma : A$ such that for all $s \in \sigma$, if there is a such that $sa \in \mathcal{L}_A$, then there is b such that $sab \in \sigma$. In other words, σ has a response to any legal Opponent move. This is crucial to interpret logic because the interpretation of proofs in game semantics always gives total strategies: this is a counterpart in semantics to the cut elimination property in syntax. To model induction and coinduction in logic, we must therefore restrict to total strategies. However, it is well-known that the class of total strategies is not closed under composition, because an infinite chattering can occur in the hidden part of the interaction. This is analogous to the fact that in λ -calculus, the class of strongly normalizing terms is not closed under application: $\delta = \lambda x.xx$ is a normal form, however $\delta\delta$ is certainly not normalizable. This problem is discussed in [2,16] and more recently in [7]. We take here the solution of [16], and equip arenas with winning functions: for every infinite play we choose a loser, hence restricting to winning strategies has the effect of blocking infinite chattering.

The definition of legal plays extends smoothly to infinite plays. Let \mathcal{L}_A^ω denote the set of infinite legal plays over A . If $\bar{s} \in \mathcal{L}_A^\omega$, we say that $\bar{s} \in \sigma$ when for all $s \sqsubset \bar{s}$, $s \in \sigma$. We write $\overline{\mathcal{L}}_A = \mathcal{L}_A + \mathcal{L}_A^\omega$. A **game** will be a pair $\mathbb{A} = (A, \mathcal{G}_A)$ where A is an arena, and \mathcal{G}_A is a function from infinite **threads** on A (*i.e.* infinite legal

plays with exactly one initial move) to $\{W, L\}$. The winning function \mathcal{G}_A extends naturally to potentially finite threads by setting, for each finite s :

$$\mathcal{G}_A(s) = \begin{cases} W & \text{if } |s| \text{ is even;} \\ L & \text{otherwise.} \end{cases}$$

Finally, \mathcal{G}_A extends to legal plays by saying that $\mathcal{G}_A(s) = W$ iff $\mathcal{G}_A(t) = W$ for every thread t of s . By abuse of notation, we keep the same notation for this extended function. The constructions on arenas presented in section 2.2 extend to constructions on games as follows:

- $\mathcal{G}_{A \times B}(s) = [\mathcal{G}_A, \mathcal{G}_B]$ (indeed, a thread on $A \times B$ is either a thread on A or a thread on B);
- $\mathcal{G}_{A+B}(s) = W$ iff all threads of $s \upharpoonright_A$ are winning for \mathcal{G}_A and all threads of $s \upharpoonright_B$ are winning for \mathcal{G}_B .
- $\mathcal{G}_{A \Rightarrow B}(s) = W$ iff if all threads of $s \upharpoonright_A$ are winning for \mathcal{G}_A , then $\mathcal{G}_B(s \upharpoonright_B) = W$.

It is straightforward to check that these constructions commute with the extension of winning functions from infinite threads to potentially infinite legal plays. We now define **winning strategies** $\sigma : \mathbb{A}$ as innocent strategies $\sigma : A$ such that for all $s \in \sigma$, $\mathcal{G}_A(s) = W$. Now, the following proposition is satisfied:

Proposition 3 *Let $\sigma : \mathbb{A} \Rightarrow \mathbb{B}$ and $\tau : \mathbb{B} \Rightarrow \mathbb{C}$ be two total winning strategies. Then $\sigma; \tau$ is total winning.*

Proof (Proof sketch). If $\sigma; \tau$ is not total, there must be infinite s in their parallel interaction $\sigma || \tau$, such that $s \upharpoonright_{A,C}$ is finite. By switching, we have in fact $|s \upharpoonright_A|$ even and $|s \upharpoonright_C|$ odd. Thus $\mathcal{G}_A(s \upharpoonright_A) = W$ and $\mathcal{G}_C(s \upharpoonright_C) = L$. We reason then by disjunction of cases. Either $\mathcal{G}_B(s \upharpoonright_B) = W$ in which case $\mathcal{G}_{B \Rightarrow C}(s \upharpoonright_{B,C}) = L$ and τ cannot be winning, or $\mathcal{G}_B(s \upharpoonright_B) = L$ in which case $\mathcal{G}_{A \Rightarrow B}(s \upharpoonright_{A,B}) = L$ and σ cannot be winning. Therefore $\sigma; \tau$ is total.

$\sigma; \tau$ must be winning as well. Suppose there is $s \in \sigma; \tau$ such that $\mathcal{G}_{A \Rightarrow C}(s) = L$. By definition of $\mathcal{G}_{A \Rightarrow C}$, this means that $\mathcal{G}_A(s \upharpoonright_A) = W$ and $\mathcal{G}_C(s \upharpoonright_C) = L$. By definition of composition, there is $u \in \sigma || \tau$ such that $s = u \upharpoonright_{A,C}$. But whatever the value of $\mathcal{G}_B(u \upharpoonright_B)$ is, one of σ or τ is losing. Therefore $\sigma; \tau$ is winning.

It is clear from the definitions that all plays in the identity are winning. It is also clear that all the structural morphisms of the cartesian closed structure of **Inn** are winning (they are essentially copycat strategies), thus this defines a cartesian closed category **Gam** of games and innocent total winning strategies.

3 Fixpoints

3.1 μLJ : an intuitionistic sequent calculus with fixpoints

Formulas. $S ::= S \Rightarrow T \mid S \vee T \mid S \wedge T \mid \mu X.T \mid \nu X.T \mid X \mid \top \mid \perp$
A formula F is **valid** if for any subformula of F of the form $\mu X.F'$,

- (1) X appears only positively in F' ,
(2) X does not appear at the root of F' (i.e. X appears at least under a \vee or a \Rightarrow in the abstract syntax tree of F').

(2) corresponds to the restriction to arenas where loops allow to express recursive types, whereas (1) is the usual positivity condition. We could of course hack the definition to get rid of these restrictions, but we choose not to obfuscate the treatment for an extra generality which is neither often considered in the literature, nor useful in practical examples of (co)induction.

Derivation rules. We present the rules with the usual dichotomy.

Identity group	
$\frac{}{A \vdash A} ax$	$\frac{\Gamma \vdash A \quad \Delta, A \vdash B}{\Gamma, \Delta \vdash B} Cut$

Structural group		
$\frac{\Gamma, A, A \vdash B}{\Gamma, A \vdash B} C$	$\frac{\Gamma \vdash B}{\Gamma, A \vdash B} W$	$\frac{\Gamma, A, B, \Delta \vdash C}{\Gamma, B, A, \Delta \vdash C} \gamma$

Logical group			
$\frac{\Gamma, A \vdash B}{\Gamma \vdash A \Rightarrow B} \Rightarrow_r$	$\frac{\Gamma \vdash A \quad \Delta, B \vdash C}{\Gamma, \Delta, A \Rightarrow B \vdash C} \Rightarrow_l$	$\frac{}{\Gamma, \perp \vdash A} \perp_l$	$\frac{}{\Gamma \vdash \top} \top_r$
$\frac{\Gamma \vdash A \quad \Gamma \vdash B}{\Gamma \vdash A \wedge B} \wedge_r$	$\frac{\Gamma, A \vdash C}{\Gamma, A \wedge B \vdash C} \wedge_l$	$\frac{\Gamma, B \vdash C}{\Gamma, A \wedge B \vdash C} \wedge_l$	
$\frac{\Gamma \vdash A}{\Gamma \vdash A \vee B} \vee_r$	$\frac{\Gamma \vdash B}{\Gamma \vdash A \vee B} \vee_r$	$\frac{\Gamma, A \vdash C \quad \Delta, B \vdash C}{\Gamma, \Delta, A \vee B \vdash C} \vee_l$	

Fixpoints			
$\frac{\Gamma \vdash T[\mu X.T/X]}{\Gamma \vdash \mu X.T} \mu_r$	$\frac{T[A/X] \vdash A}{\mu X.T \vdash A} \mu_l$	$\frac{T[\nu X.T/X] \vdash B}{\nu X.T \vdash B} \nu_l$	$\frac{A \vdash T[A/X]}{A \vdash \nu X.T} \nu_r$

Note that the μ_l , ν_l and ν_r rules are not relative to any context. In fact, the general rules with a context Γ at the left of the sequent are derivable from these ones (even if, for μ_l and ν_r , the construction of the derivation requires an induction on T), and we stick with the present ones to clarify the game model. Cut elimination on the \Rightarrow , \wedge , \vee fragment is the same as usual. For the reduction

of μ and ν , we need an additional rule to handle the unfolding of formulas. For this purpose, we add a new rule $[T]$ for each type T with free variables. This method can already be found in [1] for strictly positive functors: no type variable appears on the left of an implication. From now on, $T[A/X]$ will be abbreviated $T(A)$. This notation implies that, unless otherwise stated, X will be the variable name for which T is viewed as a functor. In the following rules, X appears only positively in T and only negatively in N :

Functors	
$\frac{A \vdash B}{T(A) \vdash T(B)} [T]$	$\frac{A \vdash B}{N(B) \vdash N(A)} [N]$

The dynamic behaviour of this rule is to locally perform the unfolding. We give some of the reduction rules. These are of two kinds: the rules for the elimination of $[T]$, and the cut elimination rules. Here are the main cases:

$$\frac{\pi}{A \vdash B} \frac{A \vdash B}{T \vdash T} [T] (X \notin FV(T)) \rightsquigarrow \frac{\pi}{T \vdash T} \text{ax}$$

$$\frac{\pi}{A \vdash B} [X] \rightsquigarrow \frac{\pi}{A \vdash B}$$

$$\frac{\pi}{A \vdash B} [N \Rightarrow T] \rightsquigarrow \frac{\frac{\pi}{A \vdash B} [N] \quad \frac{\pi}{T(A) \vdash T(B)} [T]}{N(A) \Rightarrow T(A), N(B) \vdash T(B)} \Rightarrow_l \frac{}{N(A) \Rightarrow T(A) \vdash N(B) \Rightarrow T(B)} \Rightarrow_r$$

$$\frac{\pi}{A \vdash B} [\mu Y.T] \rightsquigarrow \frac{\frac{\pi}{A \vdash B} [T[\mu Y.T(B)/Y]]}{\frac{T(A)[\mu Y.T(B)/Y] \vdash T(B)[\mu Y.T(B)/Y]}{T(A)[\mu Y.T(B)/Y] \vdash \mu Y.T(B)} \mu_r} {\mu Y.T(A) \vdash \mu Y.T(B)} \mu_l$$

We omit the rule for ν , which is dual, and for \wedge and \vee , which are simple pairing and case manipulations. Note also that most of these cases have a counterpart where T is replaced by negative N , which has the sole effect of π being a proof of $B \vdash A$ instead of $A \vdash B$ in the expansion rules. With that, we can

express the cut elimination rule for fixpoints:

$$\frac{\frac{\frac{\pi_1}{\Gamma \vdash T[\mu X.T/X]}}{\Gamma \vdash \mu X.T} \mu_r \quad \frac{\frac{\pi_2}{T[A/X] \vdash A}}{\mu X.T \vdash A} \mu_l}{\Gamma \vdash A} \text{Cut} \rightsquigarrow \frac{\frac{\frac{\pi_1}{\Gamma \vdash T[\mu X.T/X]} \quad \frac{\frac{\frac{\pi_2}{T[A/X] \vdash A}}{\mu X.T \vdash A} \mu_l}{T[\mu X.T/X] \vdash T[A/X]} [T]}{\Gamma \vdash T[A/X]} \text{Cut}}{\Gamma \vdash A} \text{Cut}$$

We skip once again the rule for ν , which is dual to μ . We choose consciously not to recall the usual cut elimination rules nor the associated commutation rules, since they are not central to our goals. μLJ , as presented above, does not formally eliminate cuts since there is no rule to reduce the following (and its dual with ν):

$$\frac{\frac{\frac{\pi_1}{T(A) \vdash A}}{\mu X.T \vdash A} \mu_l \quad \frac{\pi_2}{\Gamma, A \vdash B}}{\Gamma, \mu X.T \vdash B} \text{Cut}$$

This cannot be reduced without some prior unfolding of the $\mu X.T$ on the left. This issue is often solved [6] by replacing the rule for μ presented here above by the following:

$$\frac{T(A) \vdash A \quad \Gamma, A \vdash B}{\Gamma, \mu X.T \vdash B} \mu'$$

With the corresponding reduction rule, and analogously for ν . We choose here not to do this, first because our game model will prove consistency without the need to prove cut elimination, and second because we want to preserve the proximity with the categorical structure of initial algebras / terminal coalgebras.

3.2 The games model

We present the game model for fixpoints. We wish to model a proof system, therefore we need our strategies to be total. The base arenas of the interpretation of fixpoints will be the arenas with loops presented in section 2.3, to which we will adjoin a winning function. While the base arenas will be the same for greatest and least fixpoints, they will be distinguished by the winning function: intuitively, Player loses if a play grows infinite in a least fixpoint (inductive) game, and Opponent loses if this happens in a greatest fixpoint (coinductive) game. The winning functions we are going to present are strongly influenced by Santocanale's work on games for μ -lattices [24]. A **win open functor** is a functor $\mathbb{T} : (\mathbf{Gam} \times \mathbf{Gam}^{op})^n \rightarrow \mathbf{Gam}$ such that there is an open functor $T[X_1, \dots, X_n]$ such that for all games $\mathbb{A}_1, \dots, \mathbb{A}_{2n}$ of base arenas A_1, \dots, A_{2n} , the base arena of $\mathbb{T}(\mathbb{A}_1, \dots, \mathbb{A}_{2n})$ is $T(A_1, \dots, A_n)$. In other terms, it is the natural lifting of open functors to the category of games. By abuse of notation, we denote this by $\mathbb{T}[X_1, \dots, X_n]$, and $T[X_1, \dots, X_n]$ will denote its underlying open functor.

Least fixed point. Let $\mathbb{T}[X_1, \dots, X_n]$ be a win open functor such that \square_{X_1} appears only positively and at depth higher than 0 in $T[X_1, \dots, X_n]$. Then we define a new win open functor $\mu X_1. \mathbb{T}[X_2, \dots, X_n]$ as follows:

- Its base arena is $\mu X_1. T[X_2, \dots, X_n]$;
- If $\mathbb{A}_3, \dots, \mathbb{A}_{2n} \in \mathbf{Gam}$, $\mathcal{G}_{\mu X_1. \mathbb{T}(\mathbb{A}_3, \dots, \mathbb{A}_{2n})}(s) = W$ iff
 - There is $N \in \mathbb{N}$ such that no path of s takes the external loop more than N times, and ;
 - s is winning in the subgame inside the loop, or more formally:
 $\mathcal{G}_{\mathbb{T}(\mathbb{I}, \mathbb{I}, \mathbb{A}_3, \dots, \mathbb{A}_{2n})}(s \upharpoonright_{\mathbb{T}(\mathbb{I}, \mathbb{I}, \mathbb{A}_3, \dots, \mathbb{A}_{2n})}) = W$.

Greatest fixed point. Dually, if the same conditions are satisfied, we define the win open functor $\nu X_1. \mathbb{T}[X_1, \dots, X_n]$ as follows:

- Its base arena is $\mu X_1. T[X_2, \dots, X_n]$;
- If $\mathbb{A}_3, \dots, \mathbb{A}_{2n} \in \mathbf{Gam}$, $\mathcal{G}_{\nu X_1. \mathbb{T}(\mathbb{A}_3, \dots, \mathbb{A}_{2n})}(s) = W$ iff
 - For any $N \in \mathbb{N}$, there is a path of s crossing the external loop more than N times, or ;
 - s is winning in the subgame inside the loop, or more formally:
 $\mathcal{G}_{\mathbb{T}(\mathbb{I}, \mathbb{I}, \mathbb{A}_3, \dots, \mathbb{A}_{2n})}(s \upharpoonright_{\mathbb{T}(\mathbb{I}, \mathbb{I}, \mathbb{A}_3, \dots, \mathbb{A}_{2n})}) = W$.

It is straightforward to check that these are still functors, and in particular win open functors. There is one particular case that is worth noticing: if $\mathbb{T}[X]$ has only one hole which appears only in positive position and at depth greater than 0, then $\mu X. \mathbb{T}$ is a constant functor, *i.e.* a game. Moreover, theorem 1 implies that it is isomorphic in \mathbf{Inn} to $\mathbb{T}(\mu X. \mathbb{T})$. It is straightforward to check that this isomorphism $i_{\mathbb{T}} : \mathbb{T}(\mu X. \mathbb{T}) \rightarrow \mu X. \mathbb{T}$ is winning (it is nothing but the identity strategy), which shows that they are in fact isomorphic in \mathbf{Gam} . Then, one can prove the following theorem:

Theorem 2 *If $\mathbb{T}[X]$ has only one hole which appears only in positive position and at depth greater than 0, then the pair $(\mu X. \mathbb{T}, i_{\mathbb{T}})$ defines an **initial algebra** for $\mathbb{T}[X]$ and $(\nu X. \mathbb{T}, i_{\mathbb{T}}^{-1})$ defines a **terminal coalgebra** for $\mathbb{T}[X]$.*

Proof. We give the proof for initial alebras, the second part being dual. Let (\mathbb{A}, σ) another algebra of $\mathbb{T}[X]$. We need to show that there is a unique $\sigma^\dagger : \mu X. \mathbb{T} \Rightarrow \mathbb{B}$ such that

$$\begin{array}{ccc} \mathbb{T}(\mu X. \mathbb{T}) \xrightarrow{\mathbb{T}(\sigma^\dagger)} & \mathbb{T}(\mathbb{B}) \\ i_{\mathbb{T}} \downarrow & \downarrow \sigma \\ \mu X. \mathbb{T} \xrightarrow{\sigma^\dagger} & \mathbb{B} \end{array}$$

commutes. The idea is to iterate σ :

$$\dots \xrightarrow{\mathbb{T}^3(\sigma)} \mathbb{T}^3(\mathbb{B}) \xrightarrow{\mathbb{T}^2(\sigma)} \mathbb{T}^2(\mathbb{B}) \xrightarrow{\mathbb{T}(\sigma)} \mathbb{T}(\mathbb{B}) \xrightarrow{\sigma} \mathbb{B}$$

and to somehow take the limit. In fact we can give a direct definition of σ^\dagger :

$$\begin{aligned}\sigma^{(1)} &= \sigma \\ \sigma^{(n+1)} &= \mathbb{T}^n(\sigma); \sigma^{(n)} \\ \sigma^\dagger &= \{s \in \mathcal{L}_{\mu X. \mathbb{T} \Rightarrow \mathbb{B}} \mid \exists n \in \mathbb{N}^*, s \in \sigma^{(n)}\}\end{aligned}$$

This defines an innocent strategy, since when restricted to plays of $\mu X. \mathbb{T}$, these strategies agree on their common domain. This strategy is winning. Indeed, take an infinite play $\bar{s} \in \sigma^\dagger$. Suppose $\bar{s}_{\upharpoonright_{\mu X. \mathbb{T}}}$ is winning. By definition of $\mathcal{G}_{\mu X. \mathbb{T}}$, this means that there is $N \in \mathbb{N}$ such that no path of $\bar{s}_{\upharpoonright_{\mu X. \mathbb{T}}}$ takes the external loop more than N times. Thus, $\bar{s} \in \overline{\mathcal{L}_{\mathbb{T}^n(\mathbb{I}) \Rightarrow \mathbb{B}}}$. But this implies that $\bar{s} \in \sigma^{(n)}$, and $\sigma^{(n)}$ is a composition of winning strategies thus winning, therefore \bar{s} is winning. Moreover, σ^\dagger is the unique innocent strategy making the diagram commute: suppose there is another f making this square commute. Since $\mathbb{T}(\mu X. \mathbb{T})$ and $\mu X. \mathbb{T}$ have the same set of paths, $i_{\mathbb{T}}$ is in fact the identity, thus we have $\mathbb{T}(f); \sigma = f$. By applying T and post-composing by σ , we get:

$$\mathbb{T}^2(f); \mathbb{T}(\sigma); \sigma = \mathbb{T}(f); \sigma = f$$

And by iterating this process, we get for all $n \in \mathbb{N}$:

$$\mathbb{T}^{n+1}(f); \mathbb{T}^n(\sigma); \dots; \mathbb{T}(\sigma); \sigma = f$$

Thus:

$$\mathbb{T}^{n+1}(f); \sigma^{(n)} = f$$

Now take $s \in f$, and let n be the length of the longest path in s . Since $\mathbb{T}[X]$ has no hole at the root, no path of length n can reach B in $\mathbb{T}^{n+1}(B)$, thus $s \in \sigma^{(n)}$, therefore $s \in \sigma^\dagger$. The same reasoning also works for the other inclusion. Likewise, if $\sigma : \mathbb{B} \rightarrow \mathbb{T}(\mathbb{B})$, we build a unique $\sigma^\ddagger : \mathbb{B} \rightarrow \nu X. \mathbb{T}$ making the coalgebra diagram commute.

3.3 Interpretation of μLJ

Interpretation of Formulas. As expected, we give the interpretation of valid formulas.

$$\begin{aligned}\llbracket \top \rrbracket &= \mathbb{I} & \llbracket A \Rightarrow B \rrbracket &= \llbracket A \rrbracket \Rightarrow \llbracket B \rrbracket \\ \llbracket \perp \rrbracket &= \perp & \llbracket X \rrbracket &= \square_X \\ \llbracket A \vee B \rrbracket &= \llbracket A \rrbracket + \llbracket B \rrbracket & \llbracket \mu X. T \rrbracket &= \mu X. \llbracket T \rrbracket \\ \llbracket A \wedge B \rrbracket &= \llbracket A \rrbracket \times \llbracket B \rrbracket & \llbracket \nu X. T \rrbracket &= \nu X. \llbracket T \rrbracket\end{aligned}$$

Interpretation of Proofs. As usual, the interpretation of a proof π of a sequent $A_1, \dots, A_n \vdash B$ will be a morphism $\llbracket \pi \rrbracket : \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket \longrightarrow \llbracket B \rrbracket$. The interpretation is computed by induction on the proof tree. The interpretation of the rules of LJ is standard and its correctness follows from the cartesian closed structure of **Gam**. Here are the interpretations for the fixpoint and functor rules:

$$\left[\frac{\frac{\pi}{\Gamma \vdash T[\mu X. T/X]} \mu_r}{\Gamma \vdash \mu X. T} \right] = \llbracket \pi \rrbracket; i_{\llbracket T \rrbracket} \qquad \left[\frac{\frac{\pi}{T[A/X] \vdash A} \mu_l}{\mu X. T \vdash A} \right] = \llbracket \pi \rrbracket^\dagger$$

$$\begin{aligned}
\left[\frac{\pi}{\frac{T[\nu X.T/X] \vdash B}{\nu X.T \vdash B}} \nu_l \right] &= i_{[[T]]}^{-1}; [[\pi]] & \left[\frac{\pi}{\frac{A \vdash T[A/X]}{A \vdash \nu X.T}} \nu_r \right] &= [[\pi]]^\ddagger \\
\left[\frac{\pi}{\frac{A \vdash B}{T(A) \vdash T(B)}} [T] \right] &= [[T]]([[\pi]])
\end{aligned}$$

We do not give the details of the proof that this defines an invariant of reduction. The main technical point is the validity of the interpretation of the functor rule; more precisely when the functor is a (least or greatest) fixpoint. Given that, we get the following theorem.

Theorem 3 *If $\pi \rightsquigarrow \pi'$, then $[[\pi]] = [[\pi']$.*

In particular, this proves the following theorem which is certainly worth noticing, because μLJ has large expressive power. In particular, it contains Gödel's system T [12].

Theorem 4 *μLJ is consistent: there is no proof of \perp .*

Proof. There is no total strategy on the game \perp .

Completeness. When it comes to completeness, we run into the issue that the total winning innocent strategies are not necessarily finite, hence the usual definability process does not terminate. Nonetheless, we get a definability theorem in an infinitary version of μLJ . Whether a more precise completeness theorem is possible is a subtle point. First, we would need to restrict to an adequate subclass of the recursive total winning strategies (for example, the Ackermann function is definable in μLJ). Then again, the problem to find a proof whose interpretation is *exactly* the original strategy would be highly non-trivial: if $\sigma : \mu X.T \Rightarrow A$, we have to *guess* an invariant B , a proof π_1 of $T(B) \vdash B$ and a proof π_2 of $B \vdash A$ such that $[[\pi_1]]^\ddagger; [[\pi_2]] = \sigma$. Perhaps it would be more feasible to look for a proof whose interpretation is *observationally equivalent* to the original strategy, which would be very similar to the universality result in [15].

4 Conclusion and Future Work

We have successfully constructed a games model of a propositional intuitionistic sequent calculus μLJ with inductive and coinductive types. It is striking that the adequate winning conditions on legal plays to model (co)induction are almost identical to those used in parity games to model least and greatest fixpoints, to the extent that the restriction of our winning condition to paths coincides *exactly* with the winning condition used in [24]. It would be worthwhile to investigate this connection further: given a game viewed as a bipartite graph along with winning conditions for infinite plays, under which assumptions can these winning conditions be canonically lifted to the set of legal plays on this graph, viewed as an arena? Results in this direction might prove useful, since they would allow to import many game-theoretic results into game semantics, and thus programming languages.

This work is part of a larger project to provide game-theoretic models to total programming languages with dependent types, such as COQ or Agda. In those settings, (co)induction is crucial, since they deliberately lack general recursion. We believe that in the appropriate games setting, we can push the present results further and model Dybjer's Inductive-Recursive[10] definitions.

Acknowledgements. We would like to thank Russ Harmer, Stephane Gimenez and David Baelde for stimulating discussions, and the anonymous referees for useful comments and suggestions.

References

1. Andreas Abel and Thorsten Altenkirch. A predicative strong normalisation proof for a lambda-calculus with interleaving inductive types. In *TYPES*, 1991.
2. S. Abramsky. Semantics of interaction: an introduction to game semantics. *Semantics and Logics of Computation*, pages 1–31, 1996.
3. S. Abramsky, R. Jagadeesan, and P. Malacaria. Full Abstraction for PCF. *Info. & Comp.*, 2000.
4. S. Abramsky, H. Kohei, and G. McCusker. A fully abstract game semantics for general references. In *LICS*, pages 334–344, 1998.
5. Samson Abramsky and Radha Jagadeesan. Games and full completeness for multiplicative linear logic. *J. Symb. Log.*, 59(2):543–574, 1994.
6. David Baelde and Dale Miller. Least and greatest fixed points in linear logic. In *LPAR*, pages 92–106, 2007.
7. P. Clairambault and R. Harmer. Totality in Arena Games. *Submitted.*, 2008.
8. J.H. Conway. *On Numbers and Games*. AK Peters, Ltd., 2001.
9. P. Dybjer. Inductive sets and families in Martin-Löfs Type Theory and their set-theoretic semantics: An inversion principle for Martin-Löfs type theory. *Logical Frameworks*, 14:59–79, 1991.
10. Peter Dybjer. A general formulation of simultaneous inductive-recursive definitions in type theory. *J. Symb. Log.*, 65(2):525–549, 2000.
11. P. Freyd. Algebraically complete categories. In *Proc. 1990 Como Category Theory Conference*, volume 1488, pages 95–104. Springer, 1990.
12. K. Gödel. Über eine bisher noch nicht benutzte Erweiterung des finiten Standpunktes. *Dialectica*, 1958.
13. R. Harmer. Innocent game semantics. *Lecture notes*, 2004.
14. Russ Harmer, Martin Hyland, and Paul-André Mellies. Categorical combinatorics for innocent strategies. In *LICS*, pages 379–388, 2007.
15. J. M. E. Hyland and C.-H. Luke Ong. On full abstraction for PCF: I, II, and III. *Inf. Comput.*, 163(2):285–408, 2000.
16. M. Hyland. Game semantics. *Semantics and Logics of Computation*, 1996.
17. A. Joyal. Remarques sur la théorie des jeux à deux personnes. *Gaz. Sc. Math. Qu.*, 1977.
18. Joachim De Lataillade. Second-order type isomorphisms through game semantics. *Ann. Pure Appl. Logic*, 151(2-3):115–150, 2008.
19. Olivier Laurent. Classical isomorphisms of types. *Mathematical Structures in Computer Science*, 15(5):969–1004, 2005.
20. Ralph Loader. Equational theories for inductive types. *Ann. Pure Appl. Logic*, 84(2):175–217, 1997.
21. P. Lorenzen. Logik und Agon. *Atti Congr. Internat. di Filosofia*, 1960.
22. Guy McCusker. Games and full abstraction for FPC. *Inf. Comput.*, 160(1-2):1–61, 2000.
23. Paul-André Mellies. Asynchronous games 4: A fully complete model of propositional linear logic. In *LICS*, pages 386–395, 2005.
24. L. Santocanale. Free μ -lattices. *J. Pure Appl. Algebra*, 168(2-3):227–264, 2002.
25. A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955.