

Asynchronous games 2

The true concurrency of innocence

Paul-André Melliès ¹

Abstract

In game semantics, the higher-order value passing mechanisms of the λ -calculus are decomposed as sequences of atomic actions exchanged by a Player and its Opponent. Seen from this angle, game semantics is reminiscent of trace semantics in concurrency theory, where a process is identified to the sequences of requests it generates in the course of time. Asynchronous game semantics is an attempt to bridge the gap between the two subjects, and to see mainstream game semantics as a refined and interactive form of trace semantics. Asynchronous games are positional games played on Mazurkiewicz traces, which reformulate (and generalize) the familiar notion of arena game. The interleaving semantics of λ -terms, expressed as innocent strategies, may be analyzed in this framework, in the perspective of true concurrency. The analysis reveals that innocent strategies are positional strategies regulated by forward and backward confluence properties. This captures, we believe, the essence of innocence. We conclude the article by defining a non uniform variant of the λ -calculus, in which the game semantics of a λ -term is formulated directly as a trace semantics, performing the syntactic exploration or parsing of that λ -term.

1 Introduction

Game semantics has taught us the art of converting the higher-order value passing mechanisms of the λ -calculus into sequences of atomic actions exchanged by a Player and its Opponent in the course of time. This metamorphosis of higher-order syntax has significantly sharpened our understanding of the simply-typed λ -calculus, either as a pure calculus, or as a calculus extended with programming features like recursion, conditional branching, local control, local states, references, non determinism, probabilistic choice, etc.

Game semantics is reminiscent of trace semantics in concurrency theory. There, a process is described as a symbolic device which interacts with its environment

¹ This work has been partly supported by the FNS ACI Géométrie du Calcul (GEOCAL). Postal address: Equipe PPS, Université Paris VII, 2 place Jussieu, Case 7014, 75251 Paris Cedex 05, FRANCE. Email address: mellies@pps.jussieu.fr

by emitting or receiving requests. A sequence of such requests is called a *trace*. The trace semantics of a process π is defined as the set of traces generated by the process. In many situations, this semantics characterizes the contextual behaviour of the process. In other situations, it is refined into a bisimulation semantics.

Game semantics develops pretty much the same story for the λ -calculus. The terminology changes of course: requests are called *moves*, and traces are called *plays*. But everything works arguably as in trace semantics: the semantics of a λ -term M of type A is the set of plays σ generated by the λ -term M ; and this set of plays σ characterizes the contextual behaviour of the λ -term. The novelty of game semantics, not present in trace semantics, is that the type A defines a *game*, and that the set of plays σ generated by the λ -term M defines a *strategy* of that game.

The main thesis of this work is that game semantics is *really* the trace semantics of the λ -calculus — and even more than that: its Mazurkiewicz trace semantics. The thesis is quite unexpected, since the λ -calculus is often considered as the historical paradigm of sequentiality, whereas Mazurkiewicz traces describe truly concurrent mechanisms. The thesis is also far from immediate. It prescribes to reevaluate a large part of the conceptual and technical choices accepted today in game semantics... in order to bridge the gap with trace semantics and concurrency theory. Three issues are raised here:

- (1) The treatment of duplication in mainstream game semantics (eg. in arena games) distorts the bond with trace semantics — in particular with Mazurkiewicz traces — by adding justification pointers to traces. This prompts us to revisit this specific treatment of duplication in our first article on asynchronous games [30]. We recall below the *group-theoretic* formulation of arena games operated there in order to “eliminate” these justification pointers — or rather, in order to reunderstand them as *copy indices* modulo group action.
- (2) Thirty years ago, Antoni Mazurkiewicz developed a theory of *asynchronous traces* in which the *interleaving* semantics and the *true concurrency* semantics of a concurrent computation are related by permuting the order of *independent* events in sequences of transitions. On the other hand, current game semantics provides an interleaving semantics of the λ -calculus, in which λ -terms are expressed as *innocent* strategies. What is the true concurrency counterpart of this interleaving semantics? The task of this second article on asynchronous games is precisely to answer this question in a satisfactory way.
- (3) Ten years ago, a series of full abstraction theorems for PCF were obtained by characterizing the interactive behaviour of λ -terms as either innocent, or history-free strategies, see [2,18,35]. We believe that the present work is another significant stage in the “full abstraction” program initiated by Robin Milner [34]. For the first time indeed, we do not simply characterize, but also derive the syntax of λ -terms from elementary causality principles, expressed in asynchronous transition systems. This reconstruction requires the mediation of [30] and of its indexed treatment of threads. This leads us to an *in-*

dexed and *non-uniform* λ -calculus, from which the usual λ -calculus follows by group-theoretic principles. In this non-uniform variant of the λ -calculus, the game semantics of a λ -term may be directly formulated as a trace semantics performing the syntactic exploration or parsing of the λ -term.

The treatment of duplication. The language of traces is limited, but sufficient to interpret the *affine* fragment of the λ -calculus, in which every variable occurs at most once in a λ -term. In this fragment, every trace (= play) generated by a λ -term is an alternating sequence of received requests (= Opponent moves) and emitted requests (= Player moves). And a request appears at most once in a trace.

In order to extend the affine fragment to the whole λ -calculus, one needs to handle the duplication mechanisms semantically. This is a delicate matter. Several solutions have been considered in the literature already, and coexist today. By way of illustration, consider the λ -term chosen by Alonzo Church in order to interpret the natural number 2:

$$M = \lambda f.\lambda x.f f x.$$


Placed in front of two λ -terms P and Q , the λ -term M duplicates its first argument P , and applies it twice to its second argument Q . This is performed syntactically by two β -reductions:

$$MPQ \longrightarrow_{\beta} (\lambda x.PPx)Q \longrightarrow_{\beta} PPQ. \quad (1)$$


Obviously, the remainder of the computation depends on the λ -terms P and Q . The game-theoretic interpretation of the λ -term M has to anticipate all cases. This requires the semantics to manipulate several threads of the λ -term P simultaneously — and possibly many more than the two copies $P_{(1)}$ and $P_{(2)}$ appearing in the λ -term $P_{(1)}P_{(2)}Q$, typically when the λ -term $P_{(1)}$ uses its first argument $P_{(2)}$ several times in the remainder of the computation.

Now, the difficulty is that each thread of P should be clearly distinguished. A compact and elegant solution has been devised by Martin Hyland, Luke Ong and Hanno Nickau in the framework of *arena games* [18,35]. We recall that an *arena* is a forest, whose nodes are the *moves* of the game, and whose branches $m \vdash n$ are oriented to express that the move m *justifies* the move n . A move n is called *initial* when it is a root of the forest, or alternatively, when there is no move m such that $m \vdash n$. A *justified play* is then defined as a pair $(m_1 \cdots m_k, \varphi)$ consisting of a sequence of moves $m_1 \cdots m_k$ and a partial function $\varphi : \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ providing the so-called *pointer structure* of the play. The partial function φ associates to every occurrence i of a non-initial move m_i the occurrence $\varphi(i)$ of a move $m_{\varphi(i)}$ such that $m_{\varphi(i)} \vdash m_i$. One requires that $\varphi(i) < i$ in order to ensure that the justifying move $m_{\varphi(i)}$ occurs before the justified move m_i . Finally, the partial function φ is never defined on any occurrence i of an initial move m_i .

The pointer structure φ provides the necessary information to distinguish the several threads of a λ -term in the course of interaction — typically the several threads or copies of P in example (1). The pointer structure φ is conveniently represented by drawing “backward pointers” between occurrences of the sequence $m_1 \cdots m_k$. By way of illustration, consider the arena $m \vdash n \vdash p$ in which the only initial move is m . A typical justified play (s, φ) of this arena is represented graphically as:



Because adding justification pointers distorts the bond with trace semantics, in particular with Mazurkiewicz traces, we shift in [30] to another management principle based on *thread indexing*, already considered in [2,17]. The idea is to assign to each copy of the λ -term P in example (1) a natural number $k \in \mathbb{N}$ (its index) which characterizes the thread among the other copies of P . In the case of the justified play (2), this amounts to (a) adding a dumb move \star in order to justify the initial moves of the sequence, (b) indexing every justification pointer of the resulting sequence with a natural number:



then finally (c) encoding the sequence (3) as the sequence of indexed moves below:

$$m_{17} \cdot n_{17,5} \cdot p_{17,5,69} \cdot n_{17,4} \cdot p_{17,4,20} \cdot n_{17,1} \cdot p_{17,5,7} \cdot m_5 \cdot n_{5,70} \cdot p_{17,4,4}. \quad (4)$$

Obviously, the translation of a justified play (s, φ) depends on the choice of indices put on its justification pointers. Had we not taken sides with trace semantics and concurrency theory, we would be tempted (as most people do in fact) to retract to the notation (2) which is arguably simpler than its translation (4). But we carry on instead, and prompted by our task, decide to regulate the indexing by asking that two justification pointers starting from different occurrences i and j of the same move n , and ending on the same occurrence $\varphi(i) = \varphi(j)$, receive different indices k and k' . This indexing policy ensures that every indexed move occurs at most once in the sequence (4). In this way, we are back to the simplicity of the affine fragment of the λ -calculus.

An interesting point remains to be understood: what can be said about two different encodings of the same justified play? The first article of our series [30] clarifies this

point in the following way. Every game is equipped with a left and a right group action on moves:

$$\begin{aligned} G \times M &\longrightarrow M & (g, m) &\mapsto g \cdot m \\ M \times H &\longrightarrow M & (m, h) &\mapsto m \cdot h \end{aligned} \tag{5}$$

where M denotes the set of indexed moves, and G and H the two groups acting on that set of moves. Intuitively, the left (resp. right) group action operates on a move m_{k_0, \dots, k_j} by altering the indices k_{2i+1} assigned by Player (resp. the indices k_{2i} assigned by Opponent). Consequently, the *orbit* of a move m_{k_0, \dots, k_j} modulo a combination of the left and right group actions is precisely the set of all moves of the form $m_{k'_0, \dots, k'_j}$.

Now, the left and right group action on moves induces a left and a right group action on plays, defined in a pointwise manner:

$$\begin{aligned} g \cdot (m_1 \cdots m_k) &= (g \cdot m_1) \cdots (g \cdot m_k) \\ (m_1 \cdots m_k) \cdot h &= (m_1 \cdot h) \cdots (m_k \cdot h) \end{aligned} \tag{6}$$

It appears that the justified plays of the original arena game coincide precisely with the orbits of plays modulo left and right group action. Typically, the justified play (2) is just the play (4) modulo pointwise group action (6). One significant contribution of the present article is to reveal that the two group actions (5) are inherently *syntactical* group actions on a *non-uniform* variant of the λ -calculus, see Section 6 for details.

Asynchronous traces. After these necessary preliminaries on duplication and thread indexing, we shift to the core of this article: the comparison of true concurrency and interleaving in game semantics. Let us recall first a few principles of trace semantics in concurrency theory. Two requests a and b starting from a process π are called *independent* when they can be emitted or received by the process π in any order, without interference. Independence of the two requests a and b is represented graphically by *tiling* the two sequences $a \cdot b$ and $b \cdot a$ in the 2-dimensional diagram below:

$$\begin{array}{ccc} & \pi' & \\ & \swarrow \quad \searrow & \\ \pi_1 & & \pi_2 \\ & \nwarrow \quad \nearrow & \\ & \pi & \end{array} \tag{7}$$

The *interleaving* semantics of a process π is defined as the set of traces it generates in the course of time. The *true concurrency* semantics of the process is deduced from this by quotienting the traces modulo the *homotopy equivalence* \sim obtained by permuting independent requests. Expressing true concurrency by permuting the order of events in a symbolic trajectory stands among the fundamentals of concurrency theory. The idea originates from the work of Antoni Mazurkiewicz on asynchronous traces over a partially ordered alphabet [25,26] and leads to the notion of asynchronous transition system developed in [36,20,40]. The same idea reappears (independently) in Jean-Jacques Lévy’s description of the λ -calculus [24], and plays a key role in the author’s work on axiomatic rewriting theory [28,29]. The principle may be generalized to n -dimensional transition systems generated by cubical sets — where permutation of events amounts to *directed* homotopy — as advocated by Vaughn Pratt and Eric Goubault in [37,14].

In comparison to concurrency theory and rewriting theory, mainstream game semantics is still very much 1-dimensional. By way of illustration, take the sequential boolean game \mathbb{B} , starting by an Opponent question q followed by a Player answer true or false:



The plays of the tensor product $\mathbb{B} \otimes \mathbb{B}$ are obtained by interleaving the plays of the two instances \mathbb{B}_1 and \mathbb{B}_2 of the boolean game \mathbb{B} . Thus, (a fragment of) the game $\mathbb{B} \otimes \mathbb{B}$ defines a tree which looks exactly like this:



We observe in [31] that the two plays in (9) are different from a *procedural* point of view, but equivalent from an *extensional* point of view — since both of them realize the “extensional value” (*true*, *false*). We thus bend the two paths, and obtain a permutation tile with the shape of a 2-dimensional octagon:



By doing so, we shift from the familiar sequential games played on decision trees, to a new kind of sequential games played on *directed acyclic graphs* (dags). We analyze in this way the extensional content of sequential games, and deliver an alternative (and game-theoretic) proof of Thomas Ehrhard’s collapse theorem [12].

The extensional framework developed in [31] is extremely instructive, but not entirely satisfactory because the permutation tiles are “global” — that is, they involve more than two permuting moves in general. In contrast, the asynchronous games developed in the present article admit only “local” permutation tiles, permuting two moves, and similar to tile (7). By way of illustration, shifting to asynchronous games decomposes the “global” tile (10) into four “local” tiles:



Note that shifting from a directed acyclic graph in Diagram (10) to an asynchronous game in Diagram (11) induces concurrent plays like $q_1 \cdot q_2$ in the model. This indicates that a satisfactory theory of sequentiality requires a truly concurrent background, in which sequential plays like $q_1 \cdot \text{true}_1 \cdot q_2 \cdot \text{false}_2$ coexist with concurrent plays like $q_1 \cdot q_2$ or $q_1 \cdot q_2 \cdot \text{true}_1 \cdot \text{false}_2$.

The non-uniform λ -calculus. Here comes the most surprising, most difficult, and maybe most controversial, part of the paper. An asynchronous game is defined in Section 2 as an *event structure* whose events are polarized $+1$ for Player moves and -1 for Opponent moves. This polarization of events gives rise to a new class of events $m \cdot n$ consisting of an Opponent move m followed by a Player move n . We call *OP-moves* any such pair of moves. Just like ordinary moves, two *OP-moves* $m_1 \cdot n_1$ and $m_2 \cdot n_2$ may be permuted in a play, in the following way:



The permutation diagram (12) induces a homotopy relation \sim_{OP} between plays. The dual relation \sim_{PO} is defined symmetrically, by permuting PO -moves $m \cdot n$ together, where by PO -move $m \cdot n$ we mean a Player move m followed by an Opponent move n , see Section 2 for a formal definition. Note that both \sim_{OP} and \sim_{PO} preserve *alternation* of plays.

Now, there is a well-established theory of *stable* asynchronous transition systems in which every equivalence class modulo homotopy \sim may be represented as an event structure of so-called *canonical representatives*, see for instance [36,20,28]. The canonical representative of a transition a in a given sequence of transitions $s \cdot a$ describes the cascade of transitions necessary in s in order to enable the transition a . More formally, a sequence of transitions $t \cdot a$ is a canonical representative of a sequence of transitions $s \cdot a$ precisely when:

- (1) $s \cdot a \sim t \cdot a \cdot t'$ for some sequence of transitions t' , and
- (2) whenever $t \sim t' \cdot b$, the transition a cannot be permuted before the transition b .

The stability property ensures that this canonical representative $t \cdot a$ of the transition a is *unique*, modulo homotopy equivalence \sim on the sequence t .

Now, the asynchronous transition system with OP -moves as transitions happens to be stable. This implies that every OP -move $m \cdot n$ in an alternating play $s \cdot m \cdot n$ has a *unique* canonical representative of the form $t \cdot m \cdot n$, modulo homotopy equivalence \sim_{OP} on the sequence t . Strikingly, this canonical representative coincides with the so-called *Player view* $\lceil s \cdot m \cdot n \rceil$ of the play $s \cdot m \cdot n$ defined by Martin Hyland, Luke Ong and Hanno Nickau in the framework of arena games [18,35] and adapted to the more “concurrent” framework of asynchronous games in Section 3.

Now, Vincent Danos, Hugo Herbelin and Laurent Regnier observe in their work on arena games that every Player view of a justified play (s, ϕ) corresponds to the branch of an η -long Böhm tree, see [11] for details. The correspondence adapts smoothly to the indexed treatment of threads devised by the author in [30]. In this situation, every Player view of a play s corresponds to the branch of a *non-uniform* η -long Böhm tree. From this results a non-uniform λ -calculus (defined in Section 6) with a remarkable feature: the strategy σ associated to a non-uniform λ -term P may be alternatively formulated as a trace semantics performing the syntactic exploration or parsing of the λ -term P .

In this way, we reconstruct *by rational means* a non-uniform variant of the λ -calculus, starting from purely diagrammatic reflections on Mazurkiewicz traces and two-player games. The simply-typed λ -calculus itself (or more exactly, the familiar notion of η -long Böhm tree) follows by the group-theoretic techniques elaborated in [30] and further studied in Section 6. Hence, a diagrammatic and integrated framework emerges here, liberated from syntax, in which the evaluation of a λ -term P against a context $E[-]$ performs a symbolic trajectory $s : * \rightarrow x$

- whose homotopy class modulo \sim_{OP} expresses the syntactic subterm of P consumed during the evaluation of $E[P]$,
- whose homotopy class modulo \sim_{PO} expresses the syntactic subterm of $E[-]$ consumed during the evaluation of $E[P]$,
- whose homotopy class modulo \sim coincides with the target position x , and provides the type (or formula) of what remains unconsumed after the evaluation.

Related works. The idea of relating a dynamic and a static semantics of linear logic is formulated for the first time by Patrick Baillot, Vincent Danos, Thomas Ehrhard and Laurent Regnier in their early work on “timeless games” [8] and carried on by Patrick Baillot in his PhD thesis [7]. The idea reappears then in the concurrent game model of linear logic introduced by Samson Abramsky and the author [5]. There, concurrent games are defined as *complete lattices* of positions, and concurrent strategies as *closure operators* on these lattices. As a closure operator, every strategy is at the same time an increasing function on positions (the dynamic point of view) and a set of positions (the static point of view). The present paper is the result of a long journey (five years!) to connect this concurrent game semantics to mainstream sequential game semantics. See also the discussion in [1].

Martin Hyland and Andrea Schalk develop in [19] a notion of games on graphs quite similar to the constructions presented here and in [31]. One difference is the treatment of duplication: backtracking in [19,31], repetitive and indexed here. From this choice follows that the permutation tiles are global in [19,31] and local here. Another difference is that our positions are defined as *downward-closed subsets* of moves.

This article reformulates arena games and innocent strategies using concepts imported from concurrency theory. Conversely, much work has been devoted in the process calculus community to clarify the connections between the π -calculus and the λ -calculus — in particular by Martin Berger, Kohei Honda and Nobuko Yoshida in their work on sequentiality [9]. This offers an opportunity for an elegant synthesis of the two subjects, using asynchronous games, which we are currently investigating. Besides, several game models of concurrent programming languages have been already formulated in the interleaving framework of arena games [22,13]. It will be certainly instructive to recast them inside our asynchronous framework.

Outline. In the remainder of the article, we define our notion of asynchronous game (Section 2) and adapt the usual definition of innocent strategy to our setting (Section 3). We then characterize the innocent strategies in two ways: diagrammatically (Section 4) and positionally (Section 5). This leads to a non-uniform variant of the λ -calculus, for which we define a trace semantics, and which we relate to the usual λ -calculus (Section 6). Finally, we describe a series of possible refinements of asynchronous games (Section 7) and conclude (Section 8).

2 Asynchronous games

We choose the simplest possible definition of *asynchronous game*, in which the only relation between moves is an order relation \leq which reformulates the *justification* structure of arena games. This is enough to describe the language PCF, a simply-typed λ -calculus enriched with arithmetic, conditional branching, and recursion. A series of more expressive versions of the semantics are discussed in section 7.

Event structures. An *event structure* is an ordered set (M, \leq) such that every element $m \in M$ defines a *finite* downward-closed subset

$$m \downarrow = \{n \in M \mid n \leq m\}.$$

Asynchronous games. An *asynchronous game* is a triple $A = (M_A, \leq_A, \lambda_A)$ consisting of:

- an event structure (M_A, \leq_A) whose elements are called the *moves* of the game,
- a function $\lambda_A : M_A \rightarrow \{-1, +1\}$ which associates to every move a *polarity* $+1$ (for the Player moves) or -1 (for the Opponent moves).

Positions. A *position* of an asynchronous game A is any *finite* downward closed subset of (M_A, \leq_A) .

The lattice of positions. The set of positions of A is denoted $\mathcal{D}(A)$. Positions are ordered by inclusion, and closed under finite union. The partial order $(\mathcal{D}(A), \subseteq)$ thus defines a sup-lattice. The empty position is the least element of $(\mathcal{D}(A), \subseteq)$. It is denoted $*_A$. Positions are also closed under arbitrary *nonempty* intersection. Adding a top element \top to $(\mathcal{D}(A), \subseteq)$ provides a neutral element to intersection, and induces a *complete* lattice $\mathcal{D}(A)^\top = (\mathcal{D}(A), \subseteq)^\top$. The greatest lower bound and least upper bound of a family $(x_i)_{i \in I}$ of positions in $\mathcal{D}(A)$ are computed respectively as:

$$\bigwedge_{i \in I} x_i = \begin{cases} \top & \text{if } I \text{ is empty,} \\ \bigcap_{i \in I} x_i & \text{otherwise,} \end{cases}$$

$$\bigvee_{i \in I} x_i = \begin{cases} \top & \text{if } \bigcup_{i \in I} x_i \text{ is infinite,} \\ \bigcup_{i \in I} x_i & \text{if } \bigcup_{i \in I} x_i \text{ is finite.} \end{cases}$$

We call $\mathcal{D}(A)^\top$ the *lattice of positions* associated to the game A .

The asynchronous graph. Every asynchronous game A induces a graph $\mathcal{G}(A)$:

- whose nodes are the positions $x, y \in \mathcal{D}(A)$,
- whose edges $m : x \rightarrow y$ are the moves verifying $y = x \uplus \{m\}$, where \uplus denotes disjoint union, or equivalently, that $y = x \cup \{m\}$ and that the move m is not an element of x .

We call this graph $\mathcal{G}(A)$ the *asynchronous graph* of the game A . We write $s : x \rightarrow y$ for a path

$$x \xrightarrow{m_1} x_1 \xrightarrow{m_2} \dots \xrightarrow{m_{k-1}} x_{k-1} \xrightarrow{m_k} y$$

between two positions x and y . Note that there is no repetition of move in the sequence:

$$\forall i, j \in \{1, \dots, k\}, \quad i \neq j \Rightarrow m_i \neq m_j.$$

The target y of the path $s : x \rightarrow y$ may be deduced from the source x and the sequence of moves m_1, \dots, m_k , using the equation:

$$y = x \uplus \bigcup_{1 \leq i \leq k} \{m_i\}.$$

A path of $\mathcal{G}(A)$ is thus characterized by its source (or alternatively, its target) and the sequence of moves $m_1 \dots m_k$.

Homotopy. Given two paths $s, s' : x \rightarrow y$ of length 2 in the asynchronous graph $\mathcal{G}(A)$, we write

$$s \sim^1 s'$$

when

$$s = m \cdot n \quad \text{and} \quad s' = n \cdot m$$

for two moves $m, n \in M_A$. The *homotopy equivalence* \sim between paths is defined as the least equivalence relation containing \sim^1 , and closed under composition; that is, for every four paths $s_1 : x_1 \rightarrow x_2$ and $s, s' : x_2 \rightarrow x_3$ and $s_2 : x_3 \rightarrow x_4$:

$$s \sim s' \Rightarrow s_1 \cdot s \cdot s_2 \sim s_1 \cdot s' \cdot s_2.$$

We also use the notation \sim in our diagrams to indicate that two (necessarily different) moves m and n are permuted:



Note that our current definition of asynchronous game implies that two paths $s_1 : x_1 \rightarrow y_1$ and $s_2 : x_2 \rightarrow y_2$ are homotopic iff $x_1 = x_2$ and $y_1 = y_2$. Thus, homotopy becomes informative only in the presence of an *independence* relation between moves, see Section 7.

Alternating paths. A path $m_1 \cdots m_k : x \rightarrow y$ is *alternating* when:

$$\forall i \in \{1, \dots, k-1\}, \quad \lambda_A(m_{i+1}) = -\lambda_A(m_i).$$

Alternating homotopy. Given two paths $s, s' : x \rightarrow y$ of length 4 in the asynchronous graph $\mathcal{G}(A)$, we write

$$s \sim_{OP}^1 s'$$

when

$$s = m_1 \cdot n_1 \cdot m_2 \cdot n_2 \quad \text{and} \quad s' = m_2 \cdot n_2 \cdot m_1 \cdot n_1$$

for two Opponent moves $m_1, m_2 \in M_A$ and two Player moves $n_1, n_2 \in M_A$. The situation is summarized in diagram (12). The relation \sim_{OP} is defined as the least equivalence relation containing \sim_{OP}^1 and closed under composition.

Note that $s \sim_{OP} s'$ implies $s \sim s'$, but that the converse is not necessarily true, even when the two paths s and s' are alternating. A typical illustration of the phenomenon occurs in diagram (12) when the moves n_1 may be permuted in front of the move m_1 by homotopy:

$$m_1 \cdot n_1 \sim^1 n_1 \cdot m_1.$$

In that case, the sequence of moves $m_2 \cdot n_1 \cdot m_1 \cdot n_2$ defines an alternating path, which is homotopic to the path $m_1 \cdot n_1 \cdot m_2 \cdot n_2$ modulo \sim but not homotopic to that path modulo \sim_{OP} .

Plays. A play is a path starting from the empty position $*_A$:

$$*_A \xrightarrow{m_1} x_1 \xrightarrow{m_2} \cdots \xrightarrow{m_{k-1}} x_{k-1} \xrightarrow{m_k} x_k$$

in the asynchronous graph $\mathcal{G}(A)$. The set of plays is noted P_A .

Equivalently, a play of A is a finite sequence $s = m_1 \cdots m_k$ of moves, without repetition, such that the set $\{m_1, \dots, m_j\}$ is downward closed in (M_A, \leq_A) for every $1 \leq j \leq k$.

Strategy. A strategy σ is a set of alternating plays of even length such that:

- the strategy $s \in \sigma$ contains the empty play,
- every nonempty play $s \in \sigma$ starts with an Opponent move,
- σ is closed by even-length prefix:

$$\forall s \in P_A, \forall m, n \in M_A, \quad s \cdot m \cdot n \in \sigma \Rightarrow s \in \sigma,$$

- σ is deterministic: $\forall s \in P_A, \forall m, n_1, n_2 \in M_A,$

$$s \cdot m \cdot n_1 \in \sigma \text{ and } s \cdot m \cdot n_2 \in \sigma \Rightarrow n_1 = n_2.$$

We write $\sigma : A$ when σ is a strategy of A .

3 Innocent strategies

Ten years ago, Martin Hyland, Luke Ong and Hanno Nickau introduced the notion of *innocent strategy* in the framework of arena games, and solved in this way the Full Abstraction problem for the language PCF, see [18,35] for details. Innocent strategies characterize the interactive behaviour of the simply-typed λ -calculus equipped with a constant Ω for non-termination. This enriched variant of the λ -calculus appears under several guises in the literature: either as a calculus of η -long Böhm trees [11], or as partial proofs of Polarized Linear Logic [23], or (after a continuation-passing style translation) as the language PCF augmented with local control [21,4,15].

The traditional definition of innocence is formulated in two stages. First, a notion of *Player view* of a justified play (s, φ) is computed using the pointer structure φ of the play in the arena game. Then, an innocent strategy is defined as a strategy which plays according to the current Player view.

Here, we recast the definition of innocence in asynchronous games. The resulting definition is simpler than in arena games, for two reasons. First, every move m occurs at most once in a play of an asynchronous game. Consequently, there is no need to distinguish the move m from its occurrences in the play — which is a shallow but irritating difficulty of arena games. Then, every play s comes equipped with an *implicit* pointer structure φ given by the causality relation \leq between moves. Thus, the definition of Player view of a play s does not require any *explicit* pointer structure φ in an asynchronous game. We explain this key point now.

Justification pointers. Suppose that m and n are two different moves of an asynchronous game A . We write $m \vdash_A n$, and say that m justifies n , when:

- $m \leq_A n$, and
- for every move $p \in M_A$ such that $m \leq_A p \leq_A n$, either $m = p$ or $p = n$.

A move m is called *initial* when it has no justifier, or alternatively, when it is minimal in the ordered set (M_A, \leq_A) .

View extraction. We define the binary relation $\overset{\text{OP}}{\rightsquigarrow}$ as the smallest relation between alternating plays such that:

$$s_1 \cdot m \cdot n \cdot s_2 \overset{\text{OP}}{\rightsquigarrow} s_1 \cdot s_2$$

for every alternating play s_1 , every *nonempty* alternating path s_2 , every Opponent move m which does not justify any move in s_2 , and every Player move n which does not justify any move in s_2 .

Player view. The relation $\overset{\text{OP}}{\rightsquigarrow}$ defines a noetherian and locally confluent rewriting system on alternating plays. By Newman's Lemma, the rewriting system is confluent, see [6,10]. Thus, every alternating play $s \in P_A$ induces a unique normal form noted $\lceil s \rceil \in P_A$ and called its *Player view*:

$$s \overset{\text{OP}}{\rightsquigarrow} s_1 \overset{\text{OP}}{\rightsquigarrow} \dots \overset{\text{OP}}{\rightsquigarrow} s_k \overset{\text{OP}}{\rightsquigarrow} \lceil s \rceil.$$

This definition *by extraction* improves in many ways the traditional definition *by induction* formulated in [27,4,15]. The definition by extraction ensures for instance that the Player view $\lceil s \rceil$ of a play s is a play. This is not the case with the inductive definition. We come back to that interesting point later in the section, when we define the notion of *visible play* in an asynchronous game.

Asynchronous innocence. A strategy σ is called *innocent* in an asynchronous game A when for every plays $s, t \in \sigma$, for every Opponent move $m \in M_A$ and Player move $n \in M_A$:

$$s \cdot m \cdot n \in \sigma \text{ and } t \cdot m \in P_A \text{ and } \lceil s \cdot m \rceil \sim_{OP} \lceil t \cdot m \rceil \Rightarrow t \cdot m \cdot n \in \sigma.$$

This definition of innocence is more concise than the familiar one, formulated in [18,35,4]. In particular, it does not require any visibility condition on the strategy. It also generalizes the usual notion of innocence to more “concurrent” arenas, in which several moves m_1, \dots, m_k may justify the same move n — a situation which does not occur in arena games associated to linear or intuitionistic types.

Before carrying on, we establish that in any asynchronous game A ,

Lemma 1 *Every innocent strategy σ is closed under \sim_{OP} -equivalence:*

$$\forall s, t \in P_A, \quad s \in \sigma, s \sim_{OP} t \Rightarrow t \in \sigma.$$

PROOF. It is sufficient to establish the assertion for two plays s and t of the form:

$$s = s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_2, \quad t = s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_2,$$

where s_1 and s_2 are two paths, m_1 and m_2 are two Opponent moves, and n_1 and n_2 are two Player moves. The proof is by induction on the length of the path s_2 . First, we establish the property when the path s_2 is empty. Suppose that

$$s = s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2$$

is a play of the strategy σ , and that

$$t = s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \tag{14}$$

is a play of the game. As prefix of (14) the sequence $s_1 \cdot m_2$ defines a play. This ensures that neither of the two moves m_1 and n_1 justifies the move m_2 , which implies in turn that there exists an extraction step

$$s_1 \cdot m_1 \cdot n_1 \cdot m_2 \overset{\text{OP}}{\rightsquigarrow} s_1 \cdot m_2.$$

By definition of the Player view as the normal form of extraction, the Player views of $s_1 \cdot m_1 \cdot n_1 \cdot m_2$ and of $s_1 \cdot m_2$ coincide. We apply here our hypothesis that the strategy σ is innocent, and deduce from $s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \in \sigma$ that $s_1 \cdot m_2 \cdot n_2 \in \sigma$. We carry on, and establish now that $s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \in \sigma$. As a prefix of (14) the sequence $s_1 \cdot m_2 \cdot n_2 \cdot m_1$ defines a play. Since neither of the moves m_2 and n_2 justifies the move n_1 (as testifies the fact that s is a play), there exists an extraction step

$$s_1 \cdot m_2 \cdot n_2 \cdot m_1 \overset{\text{OP}}{\rightsquigarrow} s_1 \cdot m_1.$$

From this follows that the Player views of $s_1 \cdot m_1$ and of $s_1 \cdot m_2 \cdot n_2 \cdot m_1$ coincide. Again, we apply the hypothesis that the strategy σ is innocent, and deduce from $s_1 \cdot m_1 \cdot n_1 \in \sigma$ that $s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \in \sigma$. This proves the assertion when the path s_2 is empty.

Now, suppose that the path s_2 is not empty, and factors as $s_3 \cdot m \cdot n$. In that case, the two plays s and t factor as:

$$s = s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_3 \cdot m \cdot n, \quad t = s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_3 \cdot m \cdot n.$$

By hypothesis, the play s is an element of the strategy σ . By definition, a strategy is closed by even-length prefix. Thus, the play $s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_3$ is also element of the strategy σ . By induction hypothesis, it follows that the play $s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_3$ is an element of the strategy σ . Now, we observe that two \sim_{OP} -equivalent plays have \sim_{OP} -equivalent Player views: this key property is a simple consequence of the definition by extraction of a Player view. The property ensures that the Player views of $s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_3 \cdot m$ and $s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_3 \cdot m$ are \sim_{OP} -equivalent. We may thus apply the hypothesis that σ is innocent, and deduce from $s \in \sigma$ that $t \in \sigma$. This concludes our argument by induction. ■

Corollary 2 *If an innocent strategy σ contains a play s , it also contains its Player*

view $\lceil s \rceil$. Moreover, if $s \cdot m$ is a play in which m is an Opponent move, then $\lceil s \cdot m \rceil$ factors as $\lceil s \cdot m \rceil = t \cdot m$ where the play t is an element of the strategy σ .

PROOF. By definition of the Player view of the play $s \in \sigma$, there exists a path s' such that $s \sim_{OP} \lceil s \rceil \cdot s'$. We know by Lemma 1 that $\lceil s \rceil \cdot s' \in \sigma$. By definition, the strategy σ is closed under even-length prefix. The Player view of a play s of even-length, is itself of even-length. We conclude that $\lceil s \rceil \in \sigma$. The second assertion is proved in exactly the same way. ■

Intuitionistic games. We conclude this section by showing that our definition of innocence is equivalent to the traditional one when the underlying asynchronous game satisfies:

- every Opponent move n has at most one justifying move m ,
- when it exists, this justifying move m is a Player move.

By convention, we call *intuitionistic* any asynchronous game verifying the two properties. This denomination is justified by the fact that any asynchronous game interpreting an intuitionistic type satisfies the two properties.

Player view (HON). In order to work out the comparison, we recast in our asynchronous framework the original definition of innocence — or more precisely its familiar formulation devised by Guy McCusker in his PhD thesis [27]. We require to that purpose that the underlying asynchronous game is intuitionistic. To every alternating play s of the asynchronous game, we associate its Hyland-Ong-Nickau (HON for short) view $\lceil s \rceil$, defined by induction on the length of the play s , as follows:

$$\begin{aligned} \lceil s \cdot n \rceil &= \lceil s \rceil \cdot n && \text{when the move } n \text{ is Player,} \\ \lceil s \cdot m \cdot t \cdot n \rceil &= \lceil s \cdot m \rceil \cdot n && \text{when the move } n \text{ is Opponent and justified by } m, \\ \lceil s \cdot n \rceil &= n && \text{when the move } n \text{ is Opponent and initial,} \\ \lceil \epsilon \rceil &= \epsilon && \text{where } \epsilon \text{ is the empty play.} \end{aligned}$$

The definition is valid because an Opponent move n has at most one justifying move m in the intuitionistic game. It is worth stressing that the Player HON-view $\lceil s \rceil$ of an alternating play s is not necessarily a play: it is only an alternating sequence of moves. This is not particularly surprising, since the problem is recurrent in arena games. We have just imported it... The bad situation occurs precisely when one applies the first clause:

$$\lceil s \cdot n \rceil = \lceil s \rceil \cdot n \tag{15}$$

to a play $s \cdot n$ in which a move m which justifies the Player move n does not appear in the sequence $\ulcorner s \urcorner$. Note that this is precisely the situation in which the two definitions of Player views (by extraction vs. by induction) differ. In that situation indeed, the equation

$$\ulcorner s \cdot n \urcorner = \ulcorner s \urcorner \cdot n \quad (16)$$

does not hold... since by construction, the Player view $\ulcorner s \cdot n \urcorner$ contains the justifying move m . The following lemma clarifies the situation:

Lemma 3 *Suppose that s is an alternating play in an intuitionistic game. Then, the equality $\ulcorner s \urcorner = \ulcorner s \urcorner$ holds iff the alternating sequence $\ulcorner s \urcorner$ is a play.*

PROOF. The left-to-right implication is immediate, because the Player view $\ulcorner s \urcorner$ of an alternating play s is a play by construction. We prove the right-to-left implication by induction on the length of s . Suppose that the alternating sequence $\ulcorner s \urcorner$ is a play. The assertion of the lemma is immediate when the play s is empty. Otherwise, we proceed by case analysis on its last move. Suppose that the last move of the play s is a Player move n . Then, the play decomposes as $s = t \cdot n$. By hypothesis, the sequence $\ulcorner s \urcorner$ is a play. From the equality

$$\ulcorner s \urcorner = \ulcorner t \cdot n \urcorner = \ulcorner t \urcorner \cdot n \quad (17)$$

follows that $\ulcorner t \urcorner$ is a play, and that every justifying move of the move n appears in the play $\ulcorner t \urcorner$. We apply our induction hypothesis on t , and deduce that

$$\ulcorner t \urcorner = \ulcorner t \urcorner \quad (18)$$

Now, we claim that the equality

$$\ulcorner t \cdot n \urcorner = \ulcorner t \urcorner \cdot n \quad (19)$$

holds. This is established as follows. By definition of the Player view $\ulcorner t \urcorner$, there exists a sequence of extractions:

$$t \overset{\text{OP}}{\rightsquigarrow} \dots \overset{\text{OP}}{\rightsquigarrow} \ulcorner t \urcorner.$$

This sequence induces in turn a sequence

$$t \cdot n \overset{\text{OP}}{\rightsquigarrow} \dots \overset{\text{OP}}{\rightsquigarrow} \ulcorner t \urcorner \cdot n$$

because all the justifying moves of the move n appear in the play $\ulcorner t \urcorner$. Besides, the play $\ulcorner t \urcorner \cdot n$ is a normal form for extraction, since any step

$$\ulcorner t \urcorner \cdot n \overset{\text{OP}}{\rightsquigarrow} u$$

would induce a step

$$\lceil t \rceil \overset{\text{OP}}{\rightsquigarrow} v$$

with $v \cdot n = u$, this contradicting the fact that $\lceil t \rceil$ is a normal form for view extraction. This proves our claim that $\lceil t \rceil \cdot n$ is the Player view $\lceil t \cdot n \rceil$. We conclude from equations (17) and (18) and (19) that

$$\lceil s \rceil = \lceil t \cdot n \rceil = \lceil t \rceil \cdot n = \lceil t \rceil \cdot n = \lceil t \cdot n \rceil = \lceil s \rceil.$$

This concludes our argument by induction when the last move of s is a move by Player.

Now, suppose that the last move of the play s is an Opponent move n . The assertion of the lemma is immediate when n is an initial move: in that case, $\lceil s \rceil = \lceil s \rceil = n$. Otherwise, the play decomposes as $s = t \cdot m \cdot u \cdot n$ where m is the unique move justifying n in the intuitionistic game. This move m is a Player move. By hypothesis, the sequence $\lceil s \rceil$ is a play. From this and the equality

$$\lceil s \rceil = \lceil t \cdot m \cdot u \cdot n \rceil = \lceil t \cdot m \rceil \cdot n \tag{20}$$

follows that $\lceil t \cdot m \rceil$ is a play. We apply our induction hypothesis on $t \cdot m$ and deduce that

$$\lceil t \cdot m \rceil = \lceil t \cdot m \rceil. \tag{21}$$

We claim that the equality

$$\lceil t \cdot m \cdot u \cdot n \rceil = \lceil t \cdot m \rceil \cdot n \tag{22}$$

holds. Note already that there exists a sequence

$$t \cdot m \cdot u \cdot n \overset{\text{OP}}{\rightsquigarrow} \dots \overset{\text{OP}}{\rightsquigarrow} t \cdot m \cdot n$$

which “extracts” the path u from the play $t \cdot m \cdot u \cdot n$. By definition of the Player view as the normal form of the extraction procedure, this implies that

$$\lceil t \cdot m \cdot u \cdot n \rceil = \lceil t \cdot m \cdot n \rceil.$$

There remains to show that

$$\lceil t \cdot m \cdot n \rceil = \lceil t \cdot m \rceil \cdot n.$$

The sequence of extractions

$$t \cdot m \overset{\text{OP}}{\rightsquigarrow} \lceil t \cdot m \rceil$$

induces a sequence of extractions

$$t \cdot m \cdot n \overset{\text{OP}}{\rightsquigarrow} \lceil t \cdot m \rceil \cdot n$$

because m is the only move in $t \cdot m$ justifying the move n . Besides, the resulting play $\lceil t \cdot m \rceil \cdot n$ is a normal form for extraction, because any step

$$\lceil t \cdot m \rceil \cdot n \overset{\text{OP}}{\rightsquigarrow} u$$

would induce a step

$$\lceil t \cdot m \rceil \overset{\text{OP}}{\rightsquigarrow} v$$

with $v \cdot n = u$, this contradicting the fact that $\lceil t \cdot m \rceil$ is a normal form for view extraction. This proves our claim that $\lceil t \cdot m \cdot n \rceil = \lceil t \cdot m \rceil \cdot n$. We deduce from equations (20) and (21) and (22) that

$$\lceil s \rceil = \lceil t \cdot m \cdot u \cdot n \rceil = \lceil t \cdot m \rceil \cdot n = \lceil t \cdot m \rceil \cdot n = \lceil t \cdot u \cdot m \cdot n \rceil = \lceil s \rceil.$$

This concludes our argument by induction when the last move of s is a move by Opponent. This establishes the assertion of Lemma 3. \blacksquare

Visibility. We define a notion of visibility in asynchronous games, similar to the notion of visibility in arena games [3,4]. Consider an asynchronous game, and an alternating play s of that game. We declare that the play s is *P-visible* when the equality

$$\lceil t \cdot m \cdot n \rceil = \lceil t \cdot m \rceil \cdot n$$

holds for every Player move n and prefix $t \cdot m \cdot n$ of the play. This equality formulates in a concise way that every justifying move of n appears in the Player view $\lceil s \cdot m \rceil$. Note that in the particular case of an intuitionistic game, the HON-view $\lceil s \rceil$ of a *P-visible* play s is a play — and not just an alternating sequence of moves. From this follows, by Lemma 3, that the equality $\lceil s \rceil = \lceil s \rceil$ holds for every *P-visible* play s .

We prove that in any asynchronous game A ,

Lemma 4 *An innocent strategy σ contains only P-visible plays.*

PROOF. Suppose that σ is an innocent strategy, and that s is a play of σ . Every prefix $t \cdot m \cdot n$ of the play s in which n is a Player move, is of even-length. By definition of a strategy, the two plays t and $t \cdot m \cdot n$ are elements of σ . By Corollary 2, the Player view $\lceil t \cdot m \rceil$ factors as $u \cdot m$ where $u \in \sigma$. By definition of the Player view, $\lceil t \cdot m \rceil = \lceil u \cdot m \rceil$. We apply here the hypothesis that σ is innocent, and deduce from $u \in \sigma$, $u \cdot m \in P_A$ and $t \cdot m \cdot n \in \sigma$, that

$$\lceil t \cdot m \rceil \cdot n = u \cdot m \cdot n$$

is a play of the strategy σ . The fact that $\lceil t \cdot m \rceil \cdot n$ defines a play ensures that all the justifying moves of n appear in $\lceil t \cdot m \rceil = u \cdot m$. Equivalently, that $\lceil t \cdot m \cdot n \rceil = \lceil t \cdot m \rceil \cdot n$. This being true for every Player move n and prefix $t \cdot m \cdot n$ of the play s , we conclude that the play s is P -visible. ■

Innocence (HON). At this point, we recast the traditional definition of innocence in our asynchronous framework, and show that it coincides with the definition of innocence given previously. For the purpose, we suppose that the underlying asynchronous game is intuitionistic — so that we may speak of the HON-view of a play. In such a game, a strategy σ is called *HON-innocent* when for every plays $s, t \in \sigma$, for every Opponent move $m \in M_A$ and Player move $n \in M_A$:

$$s \cdot m \cdot n \in \sigma \text{ and } t \cdot m \in P_A \text{ and } \lceil s \cdot m \rceil = \lceil t \cdot m \rceil \Rightarrow t \cdot m \cdot n \in \sigma.$$

Besides, one requires that every move justifying the move n in the play $t \cdot m$, appears in the sequence $\lceil t \cdot m \rceil$. This last condition is called the visibility condition, because it is equivalent to requiring that every play of the strategy σ is P -visible.

We prove that in any intuitionistic game A ,

Proposition 5 *A strategy σ is innocent iff it is HON-innocent.*

PROOF. We start by the left-to-right implication. Suppose that the strategy σ is innocent. We establish that the strategy σ is HON-innocent. Suppose that s and t are two plays of the strategy σ , that m is an Opponent move, and n is a Player move such that $s \cdot m \cdot n$ is a play of the strategy σ . Suppose also that $t \cdot m$ is a play of the game, and that $\lceil s \cdot m \rceil = \lceil t \cdot m \rceil$. We show that $t \cdot m \cdot n$ is a play of σ in order to establish that the strategy σ is HON-innocent. By Lemma 4, the play t is P -visible. The move m is an Opponent move. By definition of P -visibility, the play $t \cdot m$ is also P -visible. From this follows that $\lceil t \cdot m \rceil$ is a play. By Lemma 3, the two plays $\lceil t \cdot m \rceil$ and $\lceil t \cdot m \rceil$ coincide. We may reuse the argument to show that $\lceil s \cdot m \rceil$ is a play equal to $\lceil s \cdot m \rceil$. By hypothesis, $\lceil s \cdot m \rceil = \lceil t \cdot m \rceil$. This implies that $\lceil s \cdot m \rceil = \lceil t \cdot m \rceil$. We apply the hypothesis that σ is innocent, and deduce that $t \cdot m \cdot n$ is a play of σ . The proof is nearly finished. There remains to check the visibility condition that every play of the innocent strategy σ is P -visible. This is precisely what Lemma 4 states. We conclude that the strategy σ is HON-innocent.

We establish the right-to-left implication now. Suppose that the strategy σ is HON-innocent, that $s \cdot m \cdot n$ is a play of σ in which s is a play of σ , m is an Opponent move, and n is a Player move. Suppose also that t is a play of σ , that $t \cdot m$ is a play, and that $\lceil s \cdot m \rceil \sim_{OP} \lceil t \cdot m \rceil$. As elements of the HON-innocent strategy σ , the plays s and t are P -visible. By definition of P -visibility, the plays $s \cdot m$ and $t \cdot m$ are also P -visible, because the move m is Opponent. Thus, the sequences $\lceil s \cdot m \rceil$ and $\lceil t \cdot m \rceil$ are plays. By Lemma 3, they are equal to $\lceil s \cdot m \rceil$ and $\lceil t \cdot m \rceil$ respectively. The equivalence $\lceil s \cdot m \rceil \sim_{OP} \lceil t \cdot m \rceil$ follows from that. Here comes

the crux of the proof. Observe that, by definition, every Player move in the HON-view $\lceil s \cdot m \rceil$ justifies the following Opponent move. This ensures that the \sim_{OP} -equivalence class of the play $\lceil s \cdot m \rceil$ is a singleton. The equality $\lceil s \cdot m \rceil = \lceil t \cdot m \rceil$ follows immediately. We apply the hypothesis that the strategy σ is HON-innocent, and deduce that $t \cdot m \cdot n$ is a play of σ . This concludes the proof of the proposition. ■

Proposition 5 ensures that the homotopy relation $\lceil s \cdot m \rceil \sim_{OP} \lceil t \cdot m \rceil$ appearing in our definition of innocence in asynchronous games may be replaced by an equality $\lceil s \cdot m \rceil = \lceil t \cdot m \rceil$ when the asynchronous game is intuitionistic. The homotopy relation $\lceil s \cdot m \rceil \sim_{OP} \lceil t \cdot m \rceil$ appears in the definition only to deal with “concurrent” asynchronous games in which, typically, an Opponent move may be justified by several Player moves of the arena.

4 Diagrammatic innocence

The reformulation of Player views and innocence performed in Section 3 does not really take advantage of the asynchronous structure of our games. It could be easily carried out in arena games. In this section, we shift to a diagrammatic presentation of innocence. This alternative presentation is inherently asynchronous, and could not be formalized properly in arena games. It prepares the positional characterization of innocence delivered in Section 5.

The diagrammatic presentation of innocence devised in this section is inspired by rewriting theory, and more particularly by the diagrammatic approach developed by the author and a few others in that field [28,29,39]. There is a well-established tradition there, initiated by Alonzo Church and Barkley Rosser, to deduce the “global” properties of the rewriting space (like confluence or standardization) from “local” diagrammatic properties satisfied by redexes and residuals. We proceed in a similar way below, and reduce the “global” definition of innocence devised in Section 3 to exactly two “local” diagrammatic properties — called *backward consistency* (see Figure 1) and *forward consistency* (see Figure 2). The two diagrammatic properties should be understood as *interactive* variants of the familiar *local confluence* property in rewriting theory. Each of them captures a particular aspect of innocence, somewhat hidden in the original definition. We show below that, taken together, they characterize innocence. Remarkably, the “global” notion of Player view disappears completely from the resulting presentation.

Let us explain briefly the two diagrammatic properties. Backward consistency expresses that an innocent strategy σ should react *consistently* to a change in the order of Opponent’s inquiries. Consider a sequence of interactions s followed by the strategy σ :

$$s = s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_2$$




Fig. 1. Backward consistency

in which the moves m_1 and m_2 are played by Opponent, and the moves n_1 and n_2 are played by Player. Suppose that the move m_2 is not justified by any of the two moves m_1 and n_1 . In that case, Opponent may permute her order of inquiry, and play the move m_2 before playing the move m_1 . Backward consistency ensures that the strategy σ provides exactly the same answer to each inquiry m_1 and m_2 as in the original play s . That is, Player plays the move n_2 after Opponent has played the move m_2 ; then plays the move n_2 after Opponent has played the move m_1 . The remainder of the interaction (noted s_2) proceeds then as previously. From this follows that the play

$$s^l = s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_2$$

is also element of the strategy σ .

Forward consistency is a kind of converse to backward consistency, which captures the *liveness* aspect of innocence. Consider two sequences of interactions followed by the strategy σ :

$$s = s_1 \cdot m_1 \cdot n_1 \quad \text{and} \quad s^l = s_1 \cdot m_2 \cdot n_2$$

in which the moves m_1 and m_2 are played by Opponent, and the moves n_1 and n_2 are played by Player. Suppose that the two moves m_1 and m_2 do not coincide. In that case, Opponent may extend the play s with the move m_2 . Forward consistency ensures that the strategy σ provides an answer to this inquiry m_2 : this is precisely the liveness property mentioned earlier. Besides, Player answers the move n_2 . Consequently, the play

$$s'' = s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2$$

is an element of the strategy σ , and thus the “local confluence” diagram of Figure 2 may be completed in the same way as in Figure 1.

Backward consistency. A strategy σ is called *backward consistent* (see Figure 1) when every play $s_1 \in \sigma$, every path s_2 , every pair of Opponent moves m_1, m_2 , and




Fig. 2. Forward consistency

every pair of Player moves n_1, n_2 satisfying the properties:

$$s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_2 \in \sigma \text{ and } \neg(n_1 \vdash_A m_2) \text{ and } \neg(m_1 \vdash_A m_2)$$

satisfy also the properties:

$$\neg(n_1 \vdash_A n_2) \text{ and } \neg(m_1 \vdash_A n_2) \text{ and } s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_2 \in \sigma.$$

Forward consistency. A strategy σ is called *forward consistent* (see Figure 2) when every play $s_1 \in \sigma$, every pair of Opponent moves m_1, m_2 , and every pair of Player moves n_1, n_2 satisfying the properties:

$$s_1 \cdot m_1 \cdot n_1 \in \sigma \text{ and } s_1 \cdot m_2 \cdot n_2 \in \sigma \text{ and } m_1 \neq m_2$$

satisfy also the properties:

$$n_1 \neq n_2 \text{ and } s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \in \sigma.$$

We use a diagrammatic proof to establish that, for any strategy σ of an asynchronous game A :

Proposition 6 (diagrammatic characterization) *The strategy σ is innocent iff it is backward and forward consistent.*

PROOF. (\Rightarrow) This direction is the easiest one. Suppose that the strategy σ is innocent. We establish that the strategy σ is backward consistent. Suppose that the sequence of moves $s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_2$ is a play of the strategy σ , in which $\neg(n_1 \vdash_A m_2)$ and $\neg(m_1 \vdash_A m_2)$. From this follows that the sequence $s_1 \cdot m_2$ is a play because the sequence $s_1 \cdot m_1 \cdot n_1 \cdot m_2$ is a play and $\neg(n_1 \vdash_A m_2)$ and $\neg(m_1 \vdash_A m_2)$. Now, the extraction step

$$s_1 \cdot m_2 \overset{\text{OP}}{\rightsquigarrow} s_1 \cdot m_1 \cdot n_1 \cdot m_2$$

implies that the Player views $\lceil s_1 \cdot m_2 \rceil$ and $\lceil s_1 \cdot m_1 \cdot n_1 \cdot m_2 \rceil$ coincide. We apply here the hypothesis that the strategy σ is innocent, and deduce from $s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \in \sigma$ that $s_1 \cdot m_2 \cdot n_2 \in \sigma$. In particular, the sequence $s_1 \cdot m_2 \cdot n_2$ is a play, and $\neg(n_1 \vdash_A n_2)$ and $\neg(m_1 \vdash_A n_2)$. This establishes that the sequence $s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_2$ is a play of the game. We apply Lemma 1 and deduce from $s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_2 \in \sigma$ and from

$$s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_2 \sim_{OP} s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_2$$

that

$$s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_2 \in \sigma.$$

We conclude that the strategy σ is backward consistent. We establish now that the innocent strategy σ is forward consistent. Suppose that $s_1 \cdot m_1 \cdot n_1 \in \sigma$, that $s_1 \cdot m_2 \cdot n_2 \in \sigma$, and that $m_1 \neq m_2$. In that case, the sequence $s \cdot m_1 \cdot n_1 \cdot m_2$ is a play, whose P -view coincides with the P -view of the play $s \cdot m_2$. We apply the hypothesis that the strategy σ is innocent, and deduce from $s_1 \cdot m_2 \cdot n_2 \in \sigma$ that $s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \in \sigma$. We have just established that the innocent strategy σ is backward consistent. From this and $\neg(n_1 \vdash_A m_2)$ and $\neg(m_1 \vdash_A m_2)$ follows that

$$s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \in \sigma.$$

Moreover, the two moves n_1 and n_2 are necessarily different, since they appear in the same play $s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2$. We conclude that every innocent strategy is backward and forward consistent.

(\Leftarrow) This direction is more difficult to establish. Suppose that the strategy σ is backward and forward consistent. Suppose that $s : *_A \twoheadrightarrow x$ is a play of the strategy σ , and that $m : x \longrightarrow y$ is an Opponent move defining a composite play $s \cdot m : *_A \twoheadrightarrow y$. Suppose moreover that n is a Player move. We claim that:

$$s \cdot m \cdot n \in \sigma \iff \lceil s \cdot m \rceil \cdot n \in \sigma. \quad (23)$$

In particular, we claim that each of the two alternating sequences of moves $s \cdot m \cdot n$ and $\lceil s \cdot m \rceil \cdot n$ is a play when one of them is a play of the strategy σ . We prove this claim as follows. The Player view $\lceil s \cdot m \rceil$ is of the form $t \cdot m : *_A \twoheadrightarrow y'$ where $t : *_A \twoheadrightarrow x'$ is a play of even-length, and $m : x' \longrightarrow y'$ is the Opponent move m starting this time from the position x' . By definition of the Player view $\lceil s \cdot m \rceil$, there exists two alternating paths of even-length

$$t_1 = p_1 \cdot p_2 \cdots p_{2k-1} \cdot p_{2k} : x' \twoheadrightarrow x \quad \text{and} \quad t_2 = p_1 \cdot p_2 \cdots p_{2k-1} \cdot p_{2k} : y' \twoheadrightarrow y$$

such that

$$t \cdot t_1 \sim_{OP} s \quad \text{and} \quad t_1 \sim m \cdot t_2.$$

We illustrate the situation in Figure 3 (left) with a diagram for the case $k = 2$. Backward consistency ensures that the strategy σ is closed under \sim_{OP} -equivalence.




Fig. 3. The situation when $k = 2$ (before and after applying the diagrammatic reasoning).

From this follows that the plays t and $t \cdot p_1 \cdot p_2 \cdots p_{2j-1} \cdot p_{2j}$ are elements of the strategy σ , for every index $j \leq k$.

Now, suppose that the sequence $s \cdot m \cdot n$ is a play of the strategy σ . In that case, we apply the backward consistency hypothesis k times on the play

$$t \cdot p_1 \cdot p_2 \cdots p_{2k-1} \cdot p_{2k} \cdot m \cdot n \in \sigma$$

and deduce in this way that $\lceil s \cdot m \rceil \cdot n = t \cdot m \cdot n$ is a play of the strategy σ . Figure 3 (right) describes the situation after applying the backward consistency hypothesis $k = 2$ times. This proves the direction (\Rightarrow) of our claim (23).

Now, suppose that the sequence $\lceil s \cdot m \rceil \cdot n$ is a play of the strategy σ . In that case, we apply the forward consistency hypothesis k times on the play

$$t \cdot p_1 \cdot p_2 \cdots p_{2k-1} \cdot p_{2k} \in \sigma$$

and the play

$$\lceil s \cdot m \rceil \cdot n = t \cdot m \cdot n$$

to deduce that the sequence $t \cdot p_1 \cdot p_2 \cdots p_{2k-1} \cdot p_{2k} \cdot m \cdot n$ is a play of the strategy σ . From this, and the equivalence

$$s \cdot m \cdot n \sim_{OP} t \cdot p_1 \cdot p_2 \cdots p_{2k-1} \cdot p_{2k} \cdot m \cdot n$$

we deduce that the sequence $s \cdot m \cdot n$ is a play of the strategy σ . Again, this last step is justified by the fact that the strategy σ is closed under \sim_{OP} -equivalence, because it satisfies backward consistency. Figure 3 (right) describes the situation after applying the forward consistency hypothesis $k = 2$ times. This proves our claim (23).

After this preliminary result, we establish that the strategy σ is innocent. Suppose that $s \cdot m \cdot n$ and t are two plays of the strategy σ , that $t \cdot m$ is a play, and that $\lceil s \cdot m \rceil \sim_{OP} \lceil t \cdot m \rceil$. In order to establish that the strategy σ is innocent, we want to prove that $t \cdot m \cdot n \in \sigma$. We proceed as follows. We deduce from $s \cdot m \cdot n \in \sigma$ and (23) that $\lceil s \cdot m \rceil \cdot n \in \sigma$. This and $\lceil s \cdot m \rceil \cdot n \sim_{OP} \lceil t \cdot m \rceil \cdot n$ implies

that $\lceil t \cdot m \rceil \cdot n \in \sigma$ because the strategy σ is closed under \sim_{OP} -equivalence. Now, $t \cdot m \cdot n \in \sigma$ follows from $\lceil t \cdot m \rceil \cdot n \in \sigma$ and (23). We conclude that every backward and forward consistent strategy is innocent. ■

5 Positional innocence

We establish the main result of the article in this section. This result states namely that innocent strategies are positional (Theorem 8). We show more precisely that innocent strategies are *relational*², in the sense explained below (Proposition 10). This raises an interesting question. Every relational strategy σ is characterized by the set of positions σ^\bullet it encounters. So, when is a given set of positions X of the form $X = \sigma^\bullet$ for an innocent strategy σ ?

In order to answer that question properly, we introduce the notion of *pure innocence*. A purely innocent strategy is an innocent strategy which satisfies an additional property, a variant of backward consistency, depicted in Figure 4. After showing that innocence and pure innocence coincide in intuitionistic asynchronous games (Lemma 11), we characterize the set X of positions of the form $X = \sigma^\bullet$ for a purely innocent strategy σ (Proposition 12). This characterization demonstrates among other things that innocent (and purely innocent) strategies are *concurrent strategies* in the sense of the *concurrent game model* of linear logic introduced by Samson Abramsky and the author in [5] (Proposition 13).

Positional strategy. A strategy $\sigma : A$ is called *positional* when for every two plays $s_1, s_2 : *A \twoheadrightarrow x$ in the strategy σ , and every path $t : x \twoheadrightarrow y$ of $\mathcal{G}(A)$, one has:

$$(s_1 \sim s_2 \text{ and } s_1 \cdot t \in \sigma) \Rightarrow s_2 \cdot t \in \sigma.$$

We establish below the key lemma to prove Theorem 8. Given two paths s and t , we write $s \lesssim t$ when there exists a path s' such that $s \cdot s' \sim t$. Similarly, we write $s \lesssim_{OP} t$ when there exists a path s' such that $s \cdot s' \sim_{OP} t$.

We prove that

Lemma 7 *For every innocent strategy σ of an asynchronous game A , and for every two plays s and t of the strategy σ ,*

$$s \lesssim t \Rightarrow s \lesssim_{OP} t.$$

PROOF. Consider a play $s_0 : *A \twoheadrightarrow x$ of the innocent strategy σ , and two paths $s_1 : x \twoheadrightarrow y$ and $s_2 : x \twoheadrightarrow z$ such that the two composite plays $s_0 \cdot s_1 : *A \twoheadrightarrow y$ and

² Relationality is called “pure positionality” in the extended abstract [32].

$s_0 \cdot s_2 : *_A \rightarrow z$ are elements of the strategy σ . Suppose moreover that $s_1 \lesssim s_2$. We prove by induction on the length of the path s_1 that $s_1 \lesssim_{OP} s_2$. The assertion is immediate when the path s_1 is empty. Now, suppose that the path s_1 factors as $s_1 = m \cdot n \cdot t_1$ where m is an Opponent move and n is a Player move. The path s_2 decomposes as a sequence

$$s_2 = m_1 \cdot n_1 \cdots m_k \cdot n_k \quad (24)$$

consisting of Opponent moves m_i and Player moves n_i , for $1 \leq i \leq k$. The Opponent move m appears in the play s_2 because $s_1 \lesssim s_2$. From this follows

- (1) that $m = m_j$ for some index $1 \leq j \leq k$, and
- (2) that the move $m = m_j$ is not justified by any Opponent move m_i or Player move n_i for $1 \leq i < j$.

We apply then $j - 1$ times our hypothesis that the strategy σ is backward consistent, and construct in this way a path

$$t_2 = m_1 \cdot n_1 \cdots \widehat{m_j \cdot n_j} \cdots m_k \cdot n_k \quad (25)$$

satisfying

$$m_j \cdot n_j \cdot t_2 \sim_{OP} s_2 \quad \text{and} \quad s_0 \cdot m_j \cdot n_j \cdot t_2 \in \sigma.$$

The notation $\widehat{m_j \cdot n_j}$ used in (25) indicates that the two moves m_j and n_j are removed from the sequence (24). The two plays $s_0 \cdot m \cdot n$ and $s_0 \cdot m_j \cdot n_j$ are elements of the strategy σ because the strategy is closed under even-length prefix. The equality $n = n_j$ follows immediately from the equality $m = m_j$ and from the determinism of the strategy σ . The series

$$m \cdot n \cdot t_1 = s_1 \lesssim s_2 \sim_{OP} m \cdot n \cdot t_2$$

implies that $m \cdot n \cdot t_1 \lesssim m \cdot n \cdot t_2$, which implies in turn that $t_1 \lesssim t_2$ by left-simplification. Left-simplification is justified here by the fact that two paths are homotopic modulo \sim in the asynchronous graph $\mathcal{G}(A)$ if and only if they have the same source and target. Now, we may apply our induction hypothesis to the play $t_0 = s_0 \cdot m \cdot n$ and to the paths t_1 and t_2 — because the length of the path t_1 is strictly less than the length path $s_1 = m \cdot n \cdot t_1$; and because the two plays $t_0 \cdot t_1$ and $t_0 \cdot t_2$ are elements of the strategy σ . We may thus deduce from $t_1 \lesssim t_2$ that $t_1 \lesssim_{OP} t_2$. The series

$$s_1 = m \cdot n \cdot t_1 \lesssim_{OP} m \cdot n \cdot t_2 = m_j \cdot n_j \cdot t_2 \sim_{OP} s_2$$

implies then that $s_1 \lesssim_{OP} s_2$. This concludes our proof by induction of the lemma. \blacksquare

Theorem 8 (positionality) *Every innocent strategy σ is positional.*

PROOF. Suppose that $s_1, s_2 : *A \rightarrow x$ denote two homotopic plays: $s_1 \sim s_2$; and that the two plays are elements of the innocent strategy σ . Suppose now that $t : x \rightarrow y$ denotes a path which may be postcomposed to the plays s_1 and s_2 in order to define composite plays $s_1 \cdot t, s_2 \cdot t : *A \rightarrow y$. Suppose finally that the play $s_1 \cdot t$ is an element of the strategy σ . We deduce from $s_1 \sim s_2$ that $s_1 \lesssim s_2$. We then apply Lemma 7 and deduce that $s_1 \lesssim_{OP} s_2$. This implies in turn that $s_1 \sim_{OP} s_2$ because the two plays s_1 and s_2 have the same length. From this follows that $s_1 \cdot t \sim_{OP} s_2 \cdot t$. We conclude from Lemma 1 and $s_1 \cdot t \in \sigma$ that the play $s_2 \cdot t$ is an element of the strategy σ . ■

Relational strategy. To every strategy σ , we associate the set of positions σ^\bullet played by the strategy in $\mathcal{D}(A)$, defined as:

$$\sigma^\bullet = \{x \in \mathcal{D}(A) \mid \exists s \in \sigma, s : *A \rightarrow x\}.$$

Conversely, to every set of positions $X \subset \mathcal{D}(A)$, we associate the set $X^{\ddagger} \subset P_A$ of alternating plays of even-length

$$*_A = x_0 \xrightarrow{m_1} x_1 \xrightarrow{m_2} x_2 \longrightarrow \cdots \longrightarrow x_{2k-2} \xrightarrow{m_{2k-1}} x_{2k-1} \xrightarrow{m_{2k}} x_{2k}$$

in which

- (1) every move m_{2i+1} is an Opponent move, and
- (2) every move m_{2i+2} is a Player move, and
- (3) every position x_{2j} is an element of X ,

for $0 \leq i \leq k-1$ and $0 \leq j \leq k$.

It is immediate that every strategy σ is included in the set of alternating plays $(\sigma^\bullet)^{\ddagger}$.

A strategy σ is called *relational* when

$$\sigma = (\sigma^\bullet)^{\ddagger}. \tag{26}$$

The terminology is justified by the fact that every relational strategy σ may be alternatively described as the underlying relation σ^\bullet . We prove that:

Lemma 9 *Every relational strategy is positional.*

PROOF. Consider a set X of positions, two plays s_1 and s_2 elements of X^{\ddagger} , and a path t . Suppose that $s_1 \sim s_2$, and that $s_1 \cdot t$ defines a play which is an element of X^{\ddagger} . Every even-length prefix $* \rightarrow x$ of the play $s_2 \cdot t$ is an even-length prefix of the play s_2 , or has the same target x as an even-length prefix of the play $s_1 \cdot t$. From this follows that this target position x is an element of X for every even-length prefix of $s_2 \cdot t$. We conclude that the play $s_2 \cdot t$ is an element of X^{\ddagger} . Now, suppose

that the strategy σ is relational. The property above instantiated at $X = \sigma^\bullet$ implies that the strategy $\sigma = X^\sharp$ is positional. ■

Obviously, every relational strategy σ may be recovered from its set of positions σ^\bullet by using equation (26). This is not necessarily the case for a positional strategy. Consider for instance the asynchronous game $\mathbb{B} \otimes \mathbb{B}$ with two initial Opponent moves q_1, q_2 and four Player moves $\text{false}_1, \text{true}_1, \text{false}_2, \text{true}_2$ justified as expected:

$$q_1 \vdash \text{true}_1, \quad q_1 \vdash \text{false}_1, \quad q_2 \vdash \text{true}_2, \quad q_2 \vdash \text{false}_2.$$

Consider the smallest strategy σ of $\mathbb{B} \otimes \mathbb{B}$ which contains the two plays:

$$q_1 \cdot \text{true}_1 \cdot q_2 \cdot \text{false}_2 \quad \text{and} \quad q_2 \cdot \text{false}_2.$$

The strategy σ is positional, but not relational, because the play

$$s = q_2 \cdot \text{false}_2 \cdot q_1 \cdot \text{true}_1$$

is an element of $(\sigma^\bullet)^\sharp$ but not an element of the strategy σ . For that reason, we strengthen Theorem 8 and establish the following statement:

Proposition 10 (relationality) *Every innocent strategy σ is relational.*

PROOF. Suppose that the strategy σ is innocent, and that s is a play of $(\sigma^\bullet)^\sharp$. We prove that s is a play of the strategy σ by induction on the length of s . The proof is immediate when the play s is empty. Otherwise, by definition of $(\sigma^\bullet)^\sharp$, the play $s : *_A \rightarrow x$ factors as $s = t \cdot m \cdot n$ where t is a play of $(\sigma^\bullet)^\sharp$, where m is an Opponent move, and where n is a Player move. We know by induction hypothesis that the play $t \in (\sigma^\bullet)^\sharp$ is an element of the strategy σ . Besides, the target position x of the play s is an element of σ^\bullet . By definition of σ^\bullet , there exists a play $u \in \sigma$ with the position x as target. In particular, $t \cdot m \cdot n \sim u$, and thus $t \lesssim u$. We deduce from this and Lemma 7 that $t \lesssim_{OP} u$. By definition, there exists an alternating path t' such that $t \cdot t' \sim_{OP} u$. This path t' coincides necessarily with $m \cdot n$. This establishes the equivalence $t \cdot m \cdot n \sim_{OP} u$. From this and Lemma 1, we obtain that $t \cdot m \cdot n$ is a play of the strategy σ . This concludes our proof by induction that $\sigma = (\sigma^\bullet)^\sharp$. ■

Pure innocence. An innocent strategy σ is called *purely innocent* (see Figure 4) when every play $s_1 \in \sigma$, every path s_2 , every pair m_1, m_2 of Opponent moves, and every pair n_1, n_2 of Player moves satisfying the properties:

$$s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_2 \in \sigma \quad \text{and} \quad \neg(n_1 \vdash_A m_2) \quad \text{and} \quad \neg(n_1 \vdash_A n_2)$$

satisfy also the properties:

$$\neg(m_1 \vdash_A m_2) \quad \text{and} \quad \neg(m_1 \vdash_A n_2) \quad \text{and} \quad s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_2 \in \sigma.$$




Fig. 4. Reverse consistency

This additional condition is called *reverse consistency* because it coincides with the backward consistency property (see Figure 1) in which the direction of all moves has been reversed. We establish now that pure innocence coincides with innocence in the particular case of intuitionistic games.

Lemma 11 *In any intuitionistic asynchronous game, a strategy σ is purely innocent iff it is innocent.*

PROOF. The proof is nearly immediate, and works in any asynchronous game in which no Opponent move justifies another Opponent move. It works in particular in any intuitionistic game. Suppose that the strategy σ is innocent, and that we are in the situation of Figure 4 (left) with a play $s_1 \in P_A$, a path s_2 , and moves $m_1, n_1, m_2, n_2 \in M_A$ such that:

$$s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_2 \in \sigma \text{ and } \neg(n_1 \vdash_A m_2) \text{ and } \neg(n_1 \vdash_A n_2).$$

By hypothesis on the underlying asynchronous game, the Opponent move m_1 does not justify the Opponent move m_2 . We are thus in the situation of Figure 1 (left). We may thus apply our hypothesis that the strategy σ satisfies backward consistency, and deduce the properties:

$$\neg(m_1 \vdash_A n_2) \text{ and } s_1 \cdot m_2 \cdot n_2 \cdot m_1 \cdot n_1 \cdot s_2 \in \sigma.$$

We conclude that the strategy σ is purely innocent. ■

We express below our positional characterization of pure innocence (Proposition 12). One intriguing aspect of this characterization is that it is nearly self-dual: the second and fourth clauses are exactly the first and third clauses after reversing the direction and the polarity of the moves. Closure under intersection and union demonstrates that purely innocent strategies are inherently concurrent. We come back to that point in Proposition 13. Forward and backward confluence (together with mutual attraction and the initial condition) ensure that every position $x \in X$ induces an

alternating play of even-length $s \in X^\sharp$ with target the position x . The last clause is called initial condition because it indicates on which position $*_A$ the strategy will start interacting with its environment.

The notion of forward confluence appears in [19] where it is called conflict-freeness, and (independently) in the author's game-theoretic proof [31] of Thomas Ehrhard's collapse theorem — where the notion plays a fundamental role in the dynamic analysis of hypercoherence spaces. To some extent, forward confluence is the positional counterpart of *determinism* in the usual definition of strategy formulated at the end of Section 2. Remarkably, the dual notion of backward confluence offers here a positional account of the fact that plays are *closed under even-length prefix*. This reveals that this familiar condition on strategies (understood as sets of sequences) is a hidden form of *backward* determinism.

Proposition 12 (positional characterization) *A set of positions $X \subset \mathcal{D}(A)$ is of the form $X = \sigma^\bullet$ for a purely innocent strategy σ iff the set X satisfies the following properties:*

- *X is closed under intersection: $x, y \in X \Rightarrow x \cap y \in X$,*
- *X is closed under union: $x, y \in X \Rightarrow x \cup y \in X$,*
- *forward confluence: if $X \ni x \xrightarrow{m} y \rightarrow w \in X$ and m is an Opponent move, then there exists a unique Player move $y \xrightarrow{n} z$ such that $X \ni z \rightarrow w \in X$,*
- *backward confluence: if $X \ni w \rightarrow y \xrightarrow{n} z \in X$ and n is a Player move, then there exists a unique Opponent move $x \xrightarrow{m} y$ such that $X \ni w \rightarrow x \in X$,*
- *mutual attraction: if $X \ni x \rightarrow y \in X$ then either $x = y$, or there exists an Opponent move $x \xrightarrow{m} x'$ and a Player move $y' \xrightarrow{n} y$ such that $x' \rightarrow y'$,*
- *initial condition: the root $*_A$ is an element of X .*

PROOF. Suppose that σ is a purely innocent strategy. We establish that the set of positions σ^\bullet satisfies the six clauses formulated in Proposition 12. We prove first that σ^\bullet is closed under unions and intersections. The proof applies the familiar diagrammatic techniques of rewriting theory, based on local diagram chasing and residuals, see for instance [24,16,10,29]. Suppose that $x \in \sigma^\bullet$ and $y \in \sigma^\bullet$. By definition, there exists two plays $s \in \sigma$ and $t \in \sigma$ such that:

$$s : *_A \rightarrow x \quad \text{and} \quad t : *_A \rightarrow y.$$

The property of forward consistency enables us to apply a series of permutations of OP-moves on s and t , in order to construct two “residual” paths

$$s/t : y \rightarrow x \cup y \quad \text{and} \quad t/s : x \rightarrow x \cup y$$

such that:

$$s \cdot (t/s) \sim_{OP} t \cdot (s/t)$$

and

$$s \cdot (t/s) \in \sigma \quad \text{and} \quad t \cdot (s/t) \in \sigma.$$

This establishes that $x \cup y \in \sigma^\bullet$. The proof that $x \cap y \in \sigma^\bullet$ works in a similar way. The key observation in that respect is that the asynchronous transition system with the elements of σ^\bullet as states, and the OP-moves as transitions, is not only *confluent*: it is also *stable* in the sense of [36,20,28].

We establish now the forward confluence of σ^\bullet . Suppose that two positions x, w are elements of σ^\bullet , and that

$$x \xrightarrow{m} y \twoheadrightarrow w \quad (27)$$

for some position y and Opponent move m . By definition of σ^\bullet , there exists a play $s \in \sigma$ whose target is the position x , and a play $t \in \sigma$ whose target is the positions w . It follows from (27) that $s \lesssim t$, and from Lemma 7 that $s \lesssim_{OP} t$. Thus, there exists a path

$$s' = m_1 \cdot n_1 \cdots m_k \cdot n_k : x \twoheadrightarrow w \quad (28)$$

consisting of Opponent moves m_i and Player moves n_i , for $1 \leq i \leq k$, such that $s \cdot s' \sim_{OP} t$. The Opponent move m is an element of the position w , but not an element of the position x . This implies

- (1) that $m = m_j$ for some index $1 \leq j \leq k$, and
- (2) that the move $m = m_j$ is not justified by any Opponent move m_i or Player move n_i , for $1 \leq i < j$.

We apply $j - 1$ times our hypothesis that the strategy σ is backward consistent, and construct in this way a path

$$s'' = m_1 \cdot n_1 \cdots \widehat{m_j \cdot n_j} \cdots m_k \cdot n_k$$

satisfying

$$m_j \cdot n_j \cdot s'' \sim_{OP} s' \quad \text{and} \quad s_1 \cdot m_j \cdot n_j \cdot s'' \in \sigma.$$

Just as in the proof of Lemma 7, the notation $\widehat{m_j \cdot n_j}$ indicates that the two moves m_j, n_j are removed from the sequence (28).

We claim that the move $n : y \twoheadrightarrow z$ defined as $n = n_j$ is the unique Player move from the position y whose target position z is an element of the set σ^\bullet , which satisfies moreover

$$x \xrightarrow{m} y \xrightarrow{n} z \twoheadrightarrow w. \quad (29)$$

By definition, the position z is the target of the even-length play $s_1 \cdot m_j \cdot n_j = s_1 \cdot m \cdot n$ which is prefix of the play $s_1 \cdot m_j \cdot n_j \cdot s'' \in \sigma$. From this follows that $s_1 \cdot m_j \cdot n_j$ is a play of the strategy σ , and thus, that its target z is an element of the set σ^\bullet . Besides,

the fact that $z \rightarrow w$ follows immediately from the definition of the move n_j . We have established that the Player move $n : y \rightarrow z$ has a position of σ^\bullet as target, and satisfies (29). We prove now that there is a unique such Player move n from the position y . Suppose that another Player move $n' : y \rightarrow z'$ has its target position z' in the set σ^\bullet , and satisfying $z' \rightarrow w$. In that case, the position y coincides with the intersection of the two positions $z = y \uplus \{n\}$ and $z' = y \uplus \{n'\}$. Now, we have just established that the set σ^\bullet is closed under intersection. The position y is thus an element of the set σ^\bullet . This and $y = x \uplus \{m\}$ contradicts the fact that every position of the set σ^\bullet contains as many Opponent moves as Player moves. This concludes the proof that the set σ^\bullet satisfies forward confluence.

The backward confluence property of σ^\bullet is established in the same way, by duality. Reverse consistency replaces backward consistency in the argument to obtain the Opponent move m solution of the confluence problem. Closure under intersection is replaced by closure under union in order to establish the uniqueness of that move m .

The two last assertions are immediate: mutual attraction follows from Lemma 7, and the initial condition that σ^\bullet contains the starting position $*_A$ follows from the fact that the strategy σ contains the empty play ϵ_A . This concludes the proof that the set of positions σ^\bullet satisfies the six assertions of Proposition 12 when the strategy σ is purely innocent.

We establish now the converse property that any set X of positions satisfying the six clauses of Proposition 12 is of the form σ^\bullet for a relational strategy σ . Suppose that we are given such a set X of positions. We define σ as the set of alternating sequences

$$\sigma = X^\natural.$$

We recall that, by definition, the set σ contains the set of alternating plays of even-length

$$*_A = x_0 \xrightarrow{m_1} x_1 \xrightarrow{m_2} x_2 \longrightarrow \cdots \longrightarrow x_{2k-2} \xrightarrow{m_{2k-1}} x_{2k-1} \xrightarrow{m_{2k}} x_{2k}$$

in which (1) every move m_{2i+1} is an Opponent move, (2) every move m_{2i+2} is a Player move, and (3) every position x_{2j} is an element of X , for $0 \leq i \leq k-1$ and $0 \leq j \leq k$.

We show that for every position $x \in X$, there exists a play $s \in \sigma$ whose target is the position x . This is easily established by induction on the size of x . The property is immediate when the position x is empty. Suppose now that the position $x \in X$ is not empty. There exists a path $*_A \rightarrow x$ starting from the position $*_A$. The initial condition ensures that this position $*_A$ is an element of X . By mutual attraction, there exists a Player move $m : y \rightarrow x$. By backward confluence, there exists an Opponent move $n : z \rightarrow y$ such that $z \in X$. By induction hypothesis applied to the position z , there exists a play $s : *_A \rightarrow z$ in the strategy σ . By definition of X^\natural , the play

$$s \cdot m \cdot n : *_A \rightarrow z \xrightarrow{n} y \xrightarrow{m} x$$

is also an element of $\sigma = X^{\downarrow}$. This concludes our proof by induction that every position $x \in X$ is the target of a play $s \in \sigma$.

Now, we show that the set of plays σ defines a strategy in the traditional sense, formulated at the end of Section 2. To that purpose, we check that the four conditions required on the set of plays σ are satisfied:

- the set σ contains the empty play because X contains the empty position,
- by definition of σ as X^{\downarrow} , every nonempty play $s \in \sigma$ starts with an Opponent move, and σ is closed under even-length prefix,
- suppose that $s \cdot m \cdot n_1 \in \sigma$ and $s \cdot m \cdot n_2 \in \sigma$, where $s : *_A \twoheadrightarrow x$ and $m : x \longrightarrow y$ and $n_i : y \longrightarrow z_i$ for $i \in \{1, 2\}$. By definition of σ as X^{\downarrow} , the two positions $z_1 = y \uplus \{n_1\}$ and $z_2 = y \uplus \{n_2\}$ are elements of X . Since the set X is closed under intersection, the position $z_1 \cap z_2$ is also element of X . Suppose that the two moves n_1 and n_2 are different. In that case, $y = z_1 \cap z_2$ is element of X , and thus target of an alternating play $t \in \sigma$. As such, the position y contains as many Opponent moves as Player moves. This contradicts the fact that $y = x \uplus \{m\}$ and that the position x contains as many Opponent moves as Player moves as the target of the alternating play $s \in \sigma$. We conclude that $n_1 = n_2$ and thus, that σ is deterministic.

We have just established that σ defines a strategy. There remains to show that the strategy σ satisfies the three consistency properties of pure innocence (backward, forward, and reverse). We start by establishing the backward consistency property. Suppose that we are in the situation of Figure 1, with a play

$$s_1 : *_A \twoheadrightarrow x$$

four moves

$$x \xrightarrow{m_1} y_1 \xrightarrow{n_1} y_2 \xrightarrow{m_2} y_3 \xrightarrow{n_2} w$$

satisfying

$$\neg(n_1 \vdash_A m_2) \text{ and } \neg(m_1 \vdash_A m_2)$$

and a path

$$s_2 : w \twoheadrightarrow w'$$

satisfying all together

$$s_1 \cdot m_1 \cdot n_1 \cdot m_2 \cdot n_2 \cdot s_2 \in \sigma.$$

By forward confluence applied to the positions $x \in X$ and $w \in X$, and to the Opponent move

$$m_2 : x \longrightarrow z_1$$

there exists a Player move $n : z_1 \longrightarrow z_2$ such that $z_2 \in X$ and:

$$x \xrightarrow{m_2} z_1 \xrightarrow{n} z_2 \twoheadrightarrow w.$$

By forward confluence again, applied to the positions $z_2 \in X$ and $w \in X$, and to the Opponent move:

$$m_1 : z_2 \longrightarrow z_3$$

there exists a Player move $n' : z_3 \longrightarrow z_4$ with $z_4 \rightsquigarrow w$ and thus $z_4 = w$ for cardinality reasons. We conclude that the play

$$*_A \xrightarrow{s_1} x \xrightarrow{m_2} z_1 \xrightarrow{n} z_2 \xrightarrow{m_1} z_3 \xrightarrow{n'} w \xrightarrow{s_2} w'$$

is an element of the strategy σ . There remains to show that $n = n_2$ and $n' = n_1$. The only other possibility is that $n = n_1$ and $n' = n_2$. We claim that this last possibility would contradict the backward confluence of X . In that case, indeed, the two positions y_3 and z_3 are equal, and thus define with the position $x \in X$ and the move $n' = n_2$ a typical *backward confluence* problem:

$$X \ni x \rightsquigarrow y_3 \xrightarrow{n_2} w \in \sigma$$

Now, the Opponent moves m_1 and m_2 provide two *different* solutions to this backward confluence problem:

$$x \rightsquigarrow y_2 \xrightarrow{m_2} y_3 \xrightarrow{n_2} w \in \sigma \quad \text{and} \quad x \rightsquigarrow z_2 \xrightarrow{m_1} z_3 \xrightarrow{n_2} w \in \sigma$$

with the two positions y_2 and z_2 elements of the set X . This contradicts the hypothesis that there exists a *unique* such solution. We conclude that $n = n_2$ and $n' = n_1$, and thus that the play

$$*_A \xrightarrow{s_1} x \xrightarrow{m_2} z_1 \xrightarrow{n_2} z_2 \xrightarrow{m_1} z_3 \xrightarrow{n_1} w \xrightarrow{s_2} w'$$

is an element of the strategy σ . This establishes that the strategy σ satisfies backward consistency. Reverse consistency is established in exactly the same way, but dually, by reversing the direction and the polarity of the moves. Forward consistency is established by reduction to backward consistency, using the fact that the set X is closed under unions. This concludes the proof of Proposition 12. ■

Proposition 13 *Every innocent strategy $\sigma : A$ defines a closure operator σ^\bullet on the complete lattice $\mathcal{D}(A)^\top$ of positions.*

PROOF. By convention, the closure operator σ^\bullet on the lattice $\mathcal{D}(A)^\top$ is denoted in the same way as the set of positions σ^\bullet played by the strategy σ . By definition, the closure operator σ^\bullet associates to every element x of the lattice the element

$$\sigma^\bullet(x) = \bigcap \{z \in \mathcal{D}(A) \mid z \in \sigma^\bullet \text{ and } z \geq x\}. \quad (30)$$

Note that $\sigma^\bullet(x) = \top$ precisely when there exists no position $y \in \sigma^\bullet$ above the element x in the sup-lattice $\mathcal{D}(A)$. Let us check here that (30) defines a closure operator on the lattice $\mathcal{D}(A)^\top$, although the exercise is pretty elementary. By *closure operator* on the lattice $\mathcal{D}(A)^\top$, we mean a monotone, continuous, increasing

and idempotent endofunction of the lattice. Monotonicity means that, for every two elements x, y of the lattice $\mathcal{D}(A)^\top$,

$$x \leq y \quad \Rightarrow \quad \sigma^\bullet(x) \leq \sigma^\bullet(y).$$

This follows immediately from the fact that

$$\{z \in \mathcal{D}(A) \mid z \in \sigma^\bullet \text{ and } z \geq x\} \supseteq \{z \in \mathcal{D}(A) \mid z \in \sigma^\bullet \text{ and } z \geq y\}$$

when $x \leq y$. By increasing, one means that the function σ^\bullet satisfies the inequality

$$x \leq \sigma^\bullet(x) \tag{31}$$

for every element x of the lattice. This inequality follows immediately from the definition of the element $\sigma^\bullet(x)$ as the greatest lower bound of a set of elements greater than the element x . Now, idempotency means that

$$\sigma^\bullet(\sigma^\bullet(x)) = \sigma^\bullet(x)$$

for every element x of the lattice. This follows immediately from the equality:

$$\{z \in \mathcal{D}(A) \mid z \in \sigma^\bullet \text{ and } z \geq x\} = \{z \in \mathcal{D}(A) \mid z \in \sigma^\bullet \text{ and } z \geq \sigma^\bullet(x)\}.$$

Finally, continuity means that

$$\sigma^\bullet\left(\bigvee_{i \in \mathbb{N}} x_i\right) = \bigvee_{i \in \mathbb{N}} \sigma^\bullet(x_i) \tag{32}$$

for every infinite increasing sequence $(x_i)_{i \in \mathbb{N}}$ of elements:

$$x_0 \leq x_1 \leq \cdots \leq x_{i-1} \leq x_i \leq x_{i+1} \leq \cdots \tag{33}$$

in the complete lattice $\mathcal{D}(A)^\top$. At this point, we take advantage of a very particular property of that lattice: every increasing sequence of the form (33) in the lattice $\mathcal{D}(A)^\top$ is either *stationary* — that is, there exists a natural number $N \in \mathbb{N}$ such that:

$$\forall i \in \mathbb{N}, \quad i \geq N \quad \Rightarrow \quad x_i = x_{i+1},$$

or converges to the element \top :

$$\bigvee_{i \in \mathbb{N}} x_i = \top. \tag{34}$$

Equation (32) follows immediately when the sequence (33) is stationary; and it follows from the equality

$$\bigvee_{i \in \mathbb{N}} \sigma^\bullet(x_i) = \top \tag{35}$$

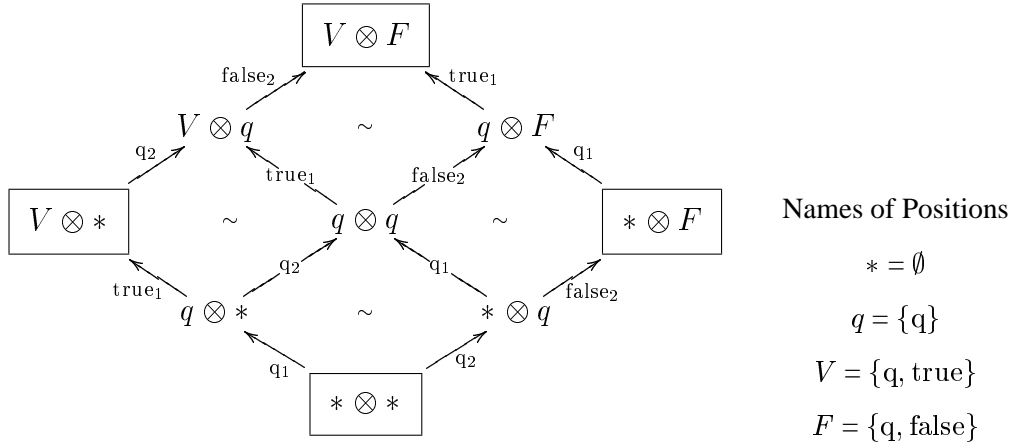


Fig. 5. The innocent strategy σ and its four positions in the game $\mathbb{B} \otimes \mathbb{B}$

when the sequence (33) converges to the element \top — Equation (35) being itself an immediate consequence of Equation (31) and Equation (34). ■

This series of properties explicates the true concurrency nature of innocence. Proposition 13 bridges sequential arena games with concurrent games as they are formulated by Samson Abramsky and the author in [5]. We illustrate this in Figure 5 with the innocent strategy $\sigma : \mathbb{B} \otimes \mathbb{B}$ which answers “true” (Vrai in French) on the left component, and “false” on the right component. The four positions of σ^\bullet are indicated on the graph: $* \otimes *$, $V \otimes *$, $* \otimes F$, and $V \otimes F$. Note that the innocent strategy σ understood as the *closure operator* or *concurrent strategy* σ^\bullet plays *directly* the position $V \otimes F \in \sigma^\bullet$ from the position $q \otimes q$, and thus answers *simultaneously* the two questions q_1 and q_2 asked by Opponent concurrently.

Despite the illustration, the reader may still find the idea of *positionality* difficult to grasp. If this is the case, we hope that the proposition below will clarify the situation. It is quite straightforward to define a notion of *innocent* counter-strategy τ interacting against the strategy σ . The counter-strategy τ may *withdraw* at any stage of the interaction. Every withdrawal of τ is expressed by an even-length play $s : *_{\mathcal{A}} \twoheadrightarrow x$ in the strategy τ , whose target position $x \in \tau^\bullet$ is of even cardinality. Our next result states that the static evaluation (by intersection) of σ^\bullet against τ^\bullet coincides with the dynamic evaluation (by interaction) of σ against τ .

Proposition 14 *For every position $x \in \mathcal{D}(A)$:*

$$\sigma^\bullet \cap \tau^\bullet = \{x\} \iff \sigma \cap \tau = \{s\} \text{ and } s : *_{\mathcal{A}} \twoheadrightarrow x.$$

It is nearly routine to construct from this a category \mathcal{G} with asynchronous games as objects, and innocent strategies as morphisms. The only difficulty is to interpret the exponentials. This is done following the principles of [30]: every game is equipped with a left and right group action, and the exponential $!A$ is interpreted as an infinite tensor product $!A = \bigotimes_{k \in \mathbb{N}} A$. The resulting category \mathcal{G} defines a model of

intuitionistic linear logic without additives. The usual category of arena games and innocent strategies embeds fully and faithfully (as a cartesian closed category) in the Kleisli category associated to the category \mathcal{G} and to its comonad. The interested reader will find the detailed construction in [30].

6 The non uniform λ -calculus

We introduce in this section a non-uniform variant of the λ -calculus. This λ -calculus is called *non-uniform* because the argument of a function $\lambda x.P$ is not a λ -term Q , but a vector \vec{Q} of λ -terms Q_i where $i \in \mathbb{N}$ is an index for each occurrence $x(i)$ (or function call) of the variable x in the λ -term P . The calculus is affine in nature: two occurrences of $x(i)$ never occur in the same term. However, the simply-typed λ -calculus may be encoded in this affine calculus, using the group-theoretic ideas developed in our first article on asynchronous games [30].

Definition of the calculus. The non-uniform λ -terms P and vectors of arguments \vec{Q} are defined by mutual induction:

$$\begin{aligned} P ::= & x(i) \text{ located variable} \\ & | P \vec{Q} \text{ application} \\ & | \lambda x.P \text{ abstraction} \end{aligned}$$

$$\vec{Q} ::= (Q_i)_{i \in \mathbb{N}} \text{ vector of non-uniform } \lambda\text{-terms indexed by an integer } i \in \mathbb{N}$$

where a located variable $x(i)$ consists of a variable x in the usual sense, and an integer $i \in \mathbb{N}$. We require that every located variable $x(i)$ appears at most once in a term. Note that a non-uniform λ -term is generally infinite. The β -reduction is defined as

$$(\lambda x.P) \vec{Q} \longrightarrow_{\beta} P[x(i) := Q_i]$$

where $P[x(i) := Q_i]$ denotes the non-uniform λ -term obtained by replacing each located variable $x(i)$ in P by the non-uniform λ -term Q_i . The non-uniform λ -terms are typed by the simple types of the λ -calculus, built on the base type α :

$$x(i) : A \vdash x(i) : A \quad \frac{\Gamma \vdash P : A \Rightarrow B \quad (\Delta_i \vdash Q_i : A)_{i \in \mathbb{N}}}{\Gamma, \Delta_0, \Delta_1, \Delta_2, \dots \vdash P \vec{Q} : B}$$

$$\frac{\Gamma, x(i_0) : A, x(i_1) : A, x(i_2) : A, \dots \vdash P : B}{\Gamma \vdash \lambda x.P : A \Rightarrow B}$$

Here, a context Γ, Δ, \dots may contain an infinite number of located variables, since the \Rightarrow -elimination rule involves a family of derivation trees $(\Delta_i \vdash Q_i : A)_{i \in \mathbb{N}}$. On

the other hand, the \Rightarrow -introduction rule may migrate an infinite number of located variables $x(i)$ from the context to the λ -term.

Non-uniform η -long Böhm trees. The non-uniform η -long Böhm trees of simple type $A = A_1 \Rightarrow \cdots A_m \Rightarrow \alpha$ are of three kinds:

- (1) $\lambda x_1 \dots \lambda x_m. (y(i) \overrightarrow{Q}_1 \cdots \overrightarrow{Q}_n)$ where
 - every variable x_j is of type A_j for $1 \leq j \leq m$,
 - the located variable $y(i)$ is of type $B = B_1 \Rightarrow \cdots B_n \Rightarrow \alpha$ for some type B ,
 - every non uniform η -long Böhm tree $(Q_k)_i$ is of type B_k , for $1 \leq k \leq n$ and $i \in \mathbb{N}$.
- (2) or Ω_B where Ω_B is a fixed constant of type B ,
- (3) or $\lambda x_1 \dots \lambda x_m. \mathcal{U}$ where \mathcal{U} is a fixed constant of type α , and every variable x_j is of type A_j , for $1 \leq j \leq m$.

Trace semantics. Every non-uniform η -long Böhm tree of simple type

$$A = A_1 \Rightarrow \cdots A_m \Rightarrow \alpha$$

is interpreted by our game model as an innocent strategy in the asynchronous game interpreting A . This game semantics may be formulated as a trace semantics on non-uniform η -long Böhm trees, in the following way.

The Opponent transitions (or moves) are generated by the rule

$$\Omega_A \longrightarrow \lambda x_1 \cdots \lambda x_m. \mathcal{U}$$

where $A = A_1 \Rightarrow \cdots A_m \Rightarrow \alpha$ and the variable x_j is of type A_j for every index $1 \leq j \leq m$.

The Player transitions are generated by the rule

$$\mathcal{U} \longrightarrow x(i) \overrightarrow{\Omega}_{A_1} \cdots \overrightarrow{\Omega}_{A_m}$$

where $x(i)$ is a located variable of type $A = A_1 \Rightarrow \cdots A_m \Rightarrow \alpha$, and $\overrightarrow{\Omega}_{A_j}$ is the vector which associates to every index $i \in \mathbb{N}$ the constant Ω_{A_j} , for every $1 \leq j \leq m$.

Every move from an η -long Böhm tree is then labelled by a subtree of the type A , once translated in linear logic as an infinite formula, using the equation

$$A \Rightarrow B \quad = \quad !A \multimap B$$

and the definition of the exponential modality as an infinite tensor:

$$!A = \bigotimes_{i \in \mathbb{N}} A.$$

Uniformity and bi-invariance. The usual (uniform) η -long Böhm trees of the λ -calculus are extracted from their non-uniform counterpart using the *bi-invariance* principle introduced in [30]. As recalled in the introduction, every game is equipped with a left and a right group action on moves. A strategy σ is called *bi-invariant* when, for every play $s \in \sigma$ and every right action $h \in H$, there exists a left action $g \in G$ such that $(g \cdot s) \cdot h \in \sigma$. This characterizes the strategies which are “blind to thread indexing”, and thus the strategies which behave as if they were defined directly in an arena game. The concept of bi-invariance remains formal and enigmatic in [30]. Here, quite fortunately, the non-uniform λ -calculus provides a syntactical explanation for the concept of bi-invariance, which clarifies its meaning and significance. We discuss that now.

Every intuitionistic type A defines a left and right group action (5) on the asynchronous game $[A]$ interpreting it in the asynchronous game model. These two group actions may be understood syntactically as acting on the non-uniform η -long Böhm trees P of type A , as follows: the effect of a right group action $h \in H$ is to permute the indices inside the vectors of arguments \vec{Q} in P , while the effect of a left group action $g \in G$ is to permute the indices of the located variables $x(i)$ in P .

By analogy with [30], a non-uniform η -long Böhm tree P is called *bi-invariant* when for every permutation $h \in H$, there is a permutation $g \in G$ such that

$$(g \cdot P) \cdot h = P.$$

It is not difficult to see that an η -long Böhm tree in the usual λ -calculus is just a *bi-invariant* η -long Böhm tree in the non-uniform λ -calculus, modulo left group action (that is, permutation of the indices of the located variables.)

For instance, let P_j denote the non-uniform η -long Böhm tree

$$P_j = \lambda x. \lambda y. (x(j) \vec{y})$$

of type $A = (\alpha \Rightarrow \alpha) \Rightarrow (\alpha \Rightarrow \alpha)$ where \vec{y} associates to every index $i \in \mathbb{N}$ the located variable $y(i)$. Obviously, P_j is bi-invariant, and represents the uniform η -long Böhm tree $\lambda x. \lambda y. x y$ of same type A . Note that P_j is equivalent to any P_k modulo left group action. The trace (or game) semantics of P_j is given by:

$$\Omega_A \xrightarrow{m} \lambda x. \lambda y. \mathcal{U} \xrightarrow{n} \lambda x. \lambda y. (x(j) \vec{\Omega}_\alpha) \xrightarrow{m_k} \lambda x. \lambda y. (x(j) \vec{Q}_k) \xrightarrow{n_k} \dots$$

Here, the move m by Opponent (labelled by the type A) asks for the value of the head variable of P_j , and the move n by Player (labelled by the type $(\alpha \Rightarrow \alpha)_j$)

answers $x(j)$; then, the move m_k by Opponent (labelled by α_k in $(\alpha \Rightarrow \alpha)_j$) asks for the value of the head variable of the k -th argument of $x(j)$, inducing the vector of arguments

$$\vec{Q}_k = \begin{cases} (Q_k)_k = \mathcal{U} \\ (Q_k)_i = \Omega_\alpha \text{ when } i \neq k \end{cases}$$

finally the move n_k by Player (labelled by α_k) answers $y(k)$, etc...

This example illustrates the fact that the trace (or game) semantics of a non-uniform η -long Böhm tree is the syntactic exploration or parsing of that tree by the Opponent. At any point of the interaction, the Player view $\lceil s \rceil$ of the play s describes the current branch of the non-uniform η -long Böhm tree.

7 Additional structures


For clarity's sake, we deliver the simplest possible definition of asynchronous game in Section 2. We review below three natural extensions of the definition.

Compatibility. Every asynchronous game may be equipped with an *incompatibility* relation $\#$ between moves, in order to model the *additives* of intuitionistic linear logic. The relation $\#$ indicates when two moves cannot appear in the same position, and thus cannot appear in the same play. The coherence axiom $(m_1 \# m_2 \leq m_3 \Rightarrow m_1 \# m_3)$ is required on every triple of moves m_1, m_2, m_3 , just as in event structures [40].

Internal vs. external positions. We may go further, and assign to every position x of the asynchronous game an integer $\kappa(x) \in \mathbb{Z}$ called its *payoff*. By convention, a position x is called *external* when the payoff $\kappa(x)$ is null, and *internal* otherwise. It is then possible to construct a game model of propositional linear logic, by identifying two strategies playing the same *external* positions. Remarkably, the resulting model incorporates the well-bracketed and the non well-bracketed variants of the original innocent arena game model. We give a detailed account of this construction in [33].

Independence. There is a well-established tradition in trace semantics of describing the *interference* mechanisms between concurrent threads by an independence relation I between events [26]. Similarly, every asynchronous game may be equipped with an independence relation between moves, in order to analyze interference

in imperative programming languages. Consider the game model of Idealized Algol formulated by Samson Abramsky and Guy McCusker in [3]. Suppose that an independence relation indicates that the moves `read` and `write(n)` interfere in the interpretation of the variable type `var`, for every natural number *n*. In that case, the interference between `read` and `write(n)` induces obstructions (“holes”) to the homotopy relation \sim on the game `var`, as indicated below:



Interestingly, the asynchronous definition of innocence adapts smoothly, and remains compositional in the presence of interfering moves (that is, it defines a category). Strategies are not positional anymore, but homotopic: they play according to the homotopy class of the current play. We believe that a geometric account of states and side effects will emerge naturally from this observation. Typically, the “state” of the system would be defined as the homotopy class of the current play; and the analysis of interference between any two such “states” would be resolved topologically. It is encouraging to see that similar intuitions have been already advocated by Uday Reddy in his work on object-based semantics of imperative languages [38].

8 Conclusion

The theory of asynchronous games is designed to bridge the gap between mainstream game semantics and concurrency theory. The preliminary results of this theory (exposed in this article) are extremely encouraging. We establish indeed that the cardinal notion of sequential game semantics — innocence — follows from elementary principles of concurrency theory, formulated in asynchronous transition systems. We introduce on the way a non-uniform λ -calculus, whose game semantics coincides with a trace semantics performing the syntactic exploration or parsing of λ -terms. This provides a concurrency-friendly picture of the λ -calculus, and firm foundations for a diagrammatic investigation of its syntax and semantics.

References

- [1] S. Abramsky. Sequentiality vs. concurrency in games and logic. *Mathematical Structures in Computer Science*, 13:531–565, 2003.

- [2] S. Abramsky, R. Jagadeesan, and P. Malacaria. Full abstraction for PCF. *Information and Computation*, 163(2):409–470, 2000.
- [3] S. Abramsky and G. McCusker. Linearity, sharing and state: a fully abstract game semantics for idealized algol with active expressions. In P. O’Hearn and R. Tennent, editors, *Algol-like Languages, volume 2*, chapter 20, pages 297–329. Birkhäuser, 1997.
- [4] S. Abramsky and G. McCusker. Game semantics. In U. Berger and H. S. eds, editors, *Computational logic*. Springer Verlag, 1999.
- [5] S. Abramsky and P.-A. Melliès. Concurrent games and full completeness. In *Proceedings of the Fourteenth Annual Symposium on Logic in Computer Science, LiCS 1999*, pages 431–442, Trento, July 1999. IEEE Computer Society Press.
- [6] F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.
- [7] P. Baillot. Approches dynamiques en sémantique de la logique linéaire: jeux et géométrie de l’interaction. Phd thesis (in french), Université Aix-Marseille 2, Jan. 1999.
- [8] P. Baillot, V. Danos, T. Ehrhard, and L. Regnier. Timeless games. In M. Nielsen and W. Thomas, editors, *Proceedings of CSL’97*, number 1414 in Lecture Notes in Computer Science, pages 56–77, Aarhus, 1997. Springer Verlag.
- [9] M. Berger, K. Honda, and N. Yoshida. Sequentiality and the pi-calculus (extended abstract). In S. Abramsky, editor, *Proceedings of the 5th International Conference on Typed Lambda Calculi and Applications, TLCA 2001*, volume 2044 of *Lecture Notes in Computer Science*. Springer Verlag, 2001.
- [10] M. Bezem and J. W. Klop. Abstract reduction systems. In Bezem et al. [39], pages 7–23.
- [11] V. Danos, H. Herbelin, and L. Regnier. Games semantics and abstract machines. In *Proceedings of the 11th Symposium on Logic in Computer Science*, pages 394–405, New Brunswick, 1996. IEEE Computer Society Press.
- [12] T. Ehrhard. A relative definability result for strongly stable functions and some corollaries. *Information and Computation*, 1997.
- [13] D. Ghica and A. Murawski. Angelic semantics of fine-grained concurrency. In *Foundations of Software Science and Computation Structures 04*, Barcelona, Mar. 2004.
- [14] E. Goubault. Geometry and concurrency: A user’s guide. *Mathematical Structures in Computer Science*, 10(4), Aug. 2000.
- [15] R. Harmer. Games and full abstraction for nondeterministic languages. PhD thesis, University of London, 2000.
- [16] G. Huet and J.-J. Lévy. Call by need computations in non-ambiguous linear term rewriting systems. Rapport INRIA 359, INRIA, 1979. Reprinted as: Computations in orthogonal rewriting systems. In J.-L. Lassez and G. D. Plotkin, editors, *Computational Logic; Essays in Honor of Alan Robinson*, pages 394–443. MIT Press, 1991.
- [17] M. Hyland. Game semantics. In A. Pitts and P. Dybjer, editors, *Semantics and logics of computation*, Publications of the Newton Institute. Cambridge University Press, 1997.
- [18] M. Hyland and L. Ong. On full abstraction for PCF: I, II and III. *Information and Computation*, 163(2):285–408, Dec. 2000.
- [19] M. Hyland and A. Schalk. Games on graphs and sequentially realizable functionals. In *Logic in Computer Science 02*, pages 257–264, Kopenhagen, July 2002. IEEE Computer Society Press.

- [20] D. Kuske. Non deterministic automata with concurrency relations and domains. In S. Tison, editor, *Proceedings of the Colloquium on Trees in Algebra and Programming, CAAP'94*, volume 787 of *Lecture Notes in Computer Science*. Springer Verlag, 1994.
- [21] J. Laird. Full abstraction for functional languages with control. In *Proceedings of the Twelfth Symposium on Logic in Computer Science, LiCS'97*, pages 58–67, 1997.
- [22] J. Laird. A game semantics of Idealized CSP. In S. Brookes and M. Mislove, editors, *Proceedings of the Seventeenth Conference on Mathematical Foundations of Programming Semantics, MFPS '01*, volume 45 of *Electronic notes in Theoretical Computer Science*. Elsevier, 2001.
- [23] O. Laurent. Polarized games (extended abstract). In *Proceedings of the seventeenth annual symposium on Logic In Computer Science*, pages 265–274, Copenhagen, July 2002. IEEE Computer Society Press.
- [24] J.-J. Lévy. Réductions correctes et optimales dans le lambda-calcul. Thèse de Doctorat d'Etat, Université Paris VII, 1978.
- [25] A. Mazurkiewicz. Concurrent program schemes and their interpretations. Technical Report DAIMI PB 78, Aarhus University, 1977.
- [26] A. Mazurkiewicz. Introduction to trace theory. In V. Diekert and G. Rozenberg, editors, *The book of traces*. World Scientific Publishing, 1995.
- [27] G. McCusker. Games and full abstraction for a functional metalanguage with recursive types. PhD thesis, Imperial College, University of London, 1996. Published in Springer-Verlag's Distinguished Dissertations in Computer Science series, 1998.
- [28] P.-A. Melliès. Axiomatic rewriting 4: a stability theorem in rewriting theory. In *Logic in Computer Science '98*. IEEE Computer Society Press, July 1998.
- [29] P.-A. Melliès. Axiomatic rewriting 6: Residual theory revisited. In S. Tison, editor, *Rewriting Techniques and Applications '02*, volume 2378 of *Lecture Notes in Computer Science*, pages 24–50. Springer Verlag, July 2002.
- [30] P.-A. Melliès. Asynchronous games 1: a group-theoretic formulation of uniformity. Manuscript. Available at <http://www.pps.jussieu.fr/~mellies/papers.html>, 2003.
- [31] P.-A. Melliès. Sequential algorithms and strongly stable functions. Prépublication électronique PPS//03/09//n°23 (pp), Equipe Preuves, Programmes et Systèmes, Apr. 2003. To appear in the special issue "Game Theory Meets Theoretical Computer Science" of *Theoretical Computer Science*.
- [32] P.-A. Melliès. Asynchronous games 2: the true concurrency of innocence. In P. Gardner and N. Yoshida, editors, *Proceedings of the Fifteenth International Conference on Concurrency Theory (CONCUR 2004)*, number 3170 in LNCS. Springer Verlag, 2004.
- [33] P.-A. Melliès. Asynchronous games 3: an innocent model of linear logic. In L. Birkedal, editor, *Proceedings of the Tenth Conference on Category Theory in Computer Science, CTCS 2004*, volume 122 of *Electronic Notes in Theoretical Computer Science*, 2005.
- [34] R. Milner. Fully abstract models of typed lambda-calculi. *Theoretical Computer Science*, 4:1–22, 1977.
- [35] H. Nickau. Hereditarily sequential functionals. In A. Nerode and Y. V. Matiyasevich, editors, *Proceedings of the Symposium on Logical Foundations of Computer Science: Logic at St. Petersburg*, volume 813 of *Lecture Notes in Computer Science*, pages 253–264. Springer Verlag, 1994.
- [36] E. W. S. P. Panangaden, V. Shanbhogue. Stability and sequentiality in data flow networks. In A. Nerode and Y. V. Matiyasevich, editors, *International Conference on*

Automates, Languages and Programming, volume 443 of *Lecture Notes in Computer Science*, pages 253–264. Springer Verlag, 1990.

- [37] V. Pratt. Modeling concurrency with geometry. In *Proceedings of the eighteenth annual symposium on Principles Of Programming Languages*, pages 311–322. ACM, IEEE Computer Society Press, Jan. 1991.
- [38] U. Reddy. Global state considered unnecessary: Introduction to object-based semantics. In P. O’Hearn and R. Tennent, editors, *Algol-like Languages, volume 2*, Progress in Theoretical Computer Science, chapter 19, pages 227–295. Birkhauser, 1997.
- [39] Terese. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2003.
- [40] G. Winskel and M. Nielsen. Models for concurrency. In S. Abramsky, D. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*. Oxford University Press, 1995.