

Concurrent Games and Full Completeness

Samson Abramsky
LFCS, Division of Informatics
Univ. of Edinburgh
<samson@dcs.ed.ac.uk>

Paul-André Melliès*
Equipe de Logique Mathématique
CNRS, Univ. Paris 7
<mellies@logique.jussieu.fr>

Abstract

A new concurrent form of game semantics is introduced. This overcomes the problems which had arisen with previous, sequential forms of game semantics in modelling Linear Logic. It also admits an elegant and robust formalization. A Full Completeness Theorem for Multiplicative-Additive Linear Logic is proved for this semantics.

1 Introduction

This paper contains two main contributions:

- the introduction of a new form of game semantics, which we call *concurrent games*.
- a proof of *full completeness* of this semantics for Multiplicative-Additive Linear Logic.

We explain the significance of each of these in turn.

Concurrent games Traditional forms of game semantics which have appeared in logic and computer science have been *sequential* in format: a play of the game is formalized as a sequence of moves. The key feature of this sequential format is the existence of a *global schedule* (or *polarization*): in each (finite) position, it is (exactly) one player’s turn to move¹. This sequential format turns out to have important limitative consequences when we wish to use game semantics to model programs or proofs:

- There is a modelling limitation. Sequential games can be used to model sequential computation, but do not yield models of parallel computation in a natural way.
- There is also a mathematical limitation. Despite the evident inherent duality in games (interchange the rôles of the two players), sequentiality is an obstacle to modelling logics in a classical format,

*Research supported by the EPSRC grant “Foundational Structures”.

¹All the games considered in this paper, and in the relevant literature to date, are two-person games.

such as (Classical) Linear Logic. In fact, sequential games have only yielded satisfactory models of *fragments* of Linear Logic: the Multiplicative fragment in [AJ92b], the Multiplicative-Exponential fragment in [BDER97], and the *negative fragment* in most other work. This has been sufficient to model λ -calculus-based programming languages, but is inadequate as a general account. We will illustrate this problem with respect to Andreas Blass’s pioneering work on game semantics below.

We will solve these problems by making a radical departure from *all* the formal versions of “games in extensive form” used to date in Logic and Computer Science of which we are aware. We shall introduce a “true concurrency” version of games in which the global polarization is abandoned. *Local* decisions are still polarized, in the sense that they are for one player or the other to make, but globally the two players act in a distributed, asynchronous fashion. At any given time, *both* may be active in different parts of the “game board”. Moreover, these concurrent games are a strict generalization of the usual sequential games.

Remarkably, this generalization and apparent complication can be formalized in a simple and robust way, arguably *more* elegant and mathematically tractable than the current formalizations of sequential games. In fact, the key ideas of the formalization were present in Abramsky and Jagadeesan’s “New Foundations for the Geometry of Interaction” [AJ92a]. What was missing in that work was the game-theoretic interpretation of the mathematics, which in turn suggests a more intuitive presentation using closure operators.

This basic formalization, combined with a suitable form of “Classical Linear Realizability”, which also contains a number of important novel ingredients, leads to a model of the whole system of Classical Linear Logic—and hence to models of Intuitionistic and Classical Logic and various λ -calculi, via the by now well-established interpretations of these systems into Linear Logic [Gir87, DJS97]. But how good is this model?

Full Completeness The usual notion of completeness for a logic is with respect to provability; Full Completeness is with respect to *proofs*.

Let \mathcal{M} be a model of the formulas *and* proofs of a logic \mathcal{L} . Typically this means that \mathcal{M} is a category with structure of an appropriate kind, such that the formulas of \mathcal{L} denote objects of \mathcal{M} , proofs Π in \mathcal{L} of entailments $A \vdash B$ denote morphisms

$$\llbracket \Pi \rrbracket : \llbracket A \rrbracket \longrightarrow \llbracket B \rrbracket,$$

and the convertibility of proofs in \mathcal{L} with respect to cut-elimination is soundly modelled by the equations between morphisms holding in \mathcal{M} . We say that \mathcal{M} is *fully complete* for \mathcal{L} if for all formulas A, B of \mathcal{L} , every morphism $f : \llbracket A \rrbracket \rightarrow \llbracket B \rrbracket$ in \mathcal{M} is the denotation of some proof Π of $A \vdash B$ in \mathcal{L} : $f = \llbracket \Pi \rrbracket$. Thus the full completeness of \mathcal{M} means that it characterizes “what it is to be a proof in \mathcal{L} ” in a very strong sense. If \mathcal{M} is defined in a syntax-independent way, this is a true semantic characterization of the “space” of proofs spanned by \mathcal{L} .

The notion of Full Completeness was introduced in [AJ92b], and a Full Completeness theorem was proved for a game semantics of Multiplicative Linear Logic (with the MIX rule). This was followed by a series of papers which established full completeness results for a variety of models with respect to various versions of Multiplicative Linear Logic (MLL) [HO92, BS96, Loa94a, Loa94b]. However, there have been no results for logics beyond the (very weak) multiplicative fragment of Linear Logic. In this paper, we make a first significant extension beyond the multiplicative fragment, by proving that the concurrent games model is fully complete for Multiplicative-Additive Linear Logic (MALL). MALL is already a much richer system than MLL, as shown by the much more sophisticated and complex notion of proof net it requires [Gir95]. Our proof of Full Completeness is correspondingly lengthy and complex. (We can only give an outline in this extended abstract; a detailed account is given in a draft full paper [AM98].) However, we believe that our methods and results will extend to the exponentials as well, thus yielding a complete analysis of Linear Logic.

Independently, Girard has obtained a form of Full Completeness result using a game semantics [Gir98a,b]. His methods, and the details of his results, appear very different to our’s. We are not yet familiar enough with his work to make a detailed comparison.

The structure of the rest of the paper is as follows. In Section 2, we present the concurrent games model. In Section 3 we show how MALL proof structures are constructed from strategies, and in Section 4 we outline the proofs of the correctness criteria for these proof structures. Finally, Section 5 gives the main result.

2 The concurrent games model

As a convenient point of departure, we begin with *Blass games* [Bla92]. These have the form

$$\prod_{i \in I} A_i \quad \text{or} \quad \sum_{i \in I} A_i$$

where each A_i is (co)inductively a Blass game. The idea is that in $\prod_{i \in I} A_i$, Opponent starts by playing some $i \in I$, and play then proceeds as in A_i . $\sum_{i \in I} A_i$ is defined dually, with Player making the initial move $i \in I$. Thus these games are trees; at each stage, it is (exactly) one player’s turn to move—the “global schedule”. Play proceeds as a sequence of moves tracing a path through the tree. This commitment to a sequential format forces an interleaving interpretation of the multiplicatives (analogous to the Expansion Theorem of CCS), which leads inexorably to the failure of composition of strategies to be associative, as shown in [AJ92b].

We begin our development of concurrent games with the idea that game trees can be viewed as partial orders, in which $x \leq y$ means that the position y can be reached from the position x by playing some additional moves. This is a natural “information ordering” as in Domain theory [AJ94]. If we add “limit points” corresponding to the infinite branches in the game tree, we obtain a complete partial order D .

Viewed in these terms, the construction of sums and products of games as in Blass games can be described as *lifted sums* as far as the underlying domains of positions are concerned:

$$D_{\sum_{i \in I} A_i} = D_{\prod_{i \in I} A_i} = \left(\sum_{i \in I} D_{A_i} \right)_{\perp}.$$

We shall represent strategies as functions on these domains of positions: $f : D \rightarrow D$, where $f(x)$ is the position obtained from x by extending it with whatever moves the strategy makes in that position. It is then immediate that $f(x) \sqsupseteq x$. Moreover, those positions where f has no moves to make (e.g. because “it is not its turn”) are exactly the *fix-points* of f . In the usual way, we require computationally reasonable strategies to be monotonic and continuous. Finally, as a useful normalizing condition, we require strategies to be *idempotent*: $f^2 = f$. To understand this, consider f applied to $f(x)$. The only moves made in $f(x)$ which were not already made in x are those made by f itself: $f(x)$ contains no more information supplied by the Opponent (i.e. the environment) than x did. Hence anything f decides to do at $f(x)$ it should have already been able to decide to do at x , and we require that $f(f(x)) = f(x)$. Of course, this allows several moves to be made in a block by a player. This possibility already exists in Blass games, e.g. Opponent must move twice initially in

$$(A \times B) \times (C \times D).$$

An important point is that strategies may not be well-defined at all positions. In general there are some positions that can never be reached by following that strategy. To mesh with the requirement that strategies are increasing functions, we adjoin a top element to the domain of positions D , writing this as D^\top . We represent f being undefined at x by $f(x) = \top$.

In summary, strategies (for either player) are represented as continuous closure operators on D^\top , which under modest assumptions on D (bounded completeness) is a complete lattice. We can completely specify a game as a structure (D, S, S^*) , where D is the domain of positions, S is the set of legal strategies for Player, and S^* is the set of legal counter-strategies, i.e. strategies for Opponent. This strictly generalizes Blass games: for such games, D is the domain of finite and infinite sequences under the prefix ordering corresponding to the paths through the game tree, and the conventions about who is to play are formalized by saying that for all $x \in D$, either $\sigma(x) = x$ for all $\sigma \in S$ (it is Opponent's turn to move), or $\tau(x) = x$ for all $\tau \in S^*$ (it is Player's turn). However, this is a very special case of our general setting; and we will overcome the problems with Blass games precisely by allowing situations in which *both* players can move.

To do this, we shall interpret the game boards for the multiplicatives differently to Blass: by a true concurrency rather than an interleaving representation. In our setting, this is simply a matter of defining

$$D_{A \otimes B} = D_A \wp B = D_A \times D_B,$$

the cartesian product of domains.

How should the Linear Logic connectives be interpreted as acting on sets of strategies? We define $S_{A^\perp} = S_A^*$, $S_{A^\perp}^* = S_A$, which corresponds to interchanging the rôles of Player and Opponent. Clearly then $A^{\perp\perp} = A$. For products, we define

$$S_{\prod_{i \in I} A_i} = \{ \langle \sigma_i | i \in I \rangle \mid \forall i \in I, \sigma_i \in S_{A_i} \}$$

where

$$\begin{aligned} \langle \sigma_i | i \in I \rangle (\perp) &= \perp \\ \langle \sigma_i | i \in I \rangle (\text{in}_i(x)) &= \text{in}_i(\sigma_i(x)). \end{aligned}$$

This corresponds exactly to the idea that Player must first wait for Opponent to choose an $i \in I$, and then plays according to some strategy for A_i .

$$S_{\prod_{i \in I} A_i}^* = \{ \text{in}_i(\sigma) \mid i \in I, \sigma \in S_{A_i}^* \}$$

where

$$\begin{aligned} \text{in}_i(\sigma)(\perp) &= \text{in}_i(\sigma(\perp)) \\ \text{in}_i(\sigma)(\text{in}_i(x)) &= \text{in}_i(\sigma(x)) \\ \text{in}_i(\sigma)(\text{in}_j(x)) &= \top, \quad (i \neq j). \end{aligned}$$

Again, this corresponds to the idea that a strategy for Opponent will firstly choose some i , then play according to a strategy for Opponent in A_i . Note that the last case above covers “unreachable states” for $\text{in}_i(\sigma)$. We can then define sums by De Morgan duality: $\sum_{i \in I} A_i = (\prod_{i \in I} A_i^\perp)^\perp$.

For the strategies for tensor, we define

$$S_{A \otimes B} = \{ \sigma \times \tau \mid \sigma \in S_A, \tau \in S_B \}$$

where $\sigma \times \tau(x, y) = (\sigma(x), \tau(y))$. (This is really smash product with respect to \top ; if either $\sigma(x) = \top$ or $\tau(y) = \top$, then the result is \top .) This exactly captures the idea of informational independence between Player's actions in A and in B (cf. [AJ92b]). How Player moves in A depends only on the information available in A , and similarly for B .

In order to define the counter-strategies for the Tensor (and hence the strategies for Par and Linear implication, and eventually the morphisms in the category of concurrent games), we introduce the most important feature of our formalization: the elegant treatment it affords of composition of strategies. Suppose firstly that $\sigma \in S$ and $\tau \in S^*$ in a game (D, S, S^*) . How do we play σ off against τ ? We define

$$\begin{aligned} \langle \sigma | \tau \rangle &= \mathbf{Y}(\sigma \circ \tau) = \bigvee_{k \in \omega} (\sigma \circ \tau)^k(\perp) \\ &= \bigvee_{k \in \omega} (\tau \circ \sigma)^k(\perp) = \mathbf{Y}(\tau \circ \sigma). \end{aligned}$$

The fact that these two least fixpoints coincide follows easily from the fact that σ and τ are continuous closures; in fact, this is a special case of the construction of the *join* of two closure operators. Thus $\langle \sigma | \tau \rangle \in D$ is the position we reach as a result of playing σ against τ . The equality of the two formulas above also shows that this is independent of all questions about “who starts”.

Now given closure operators σ on $(D \times E)^\top$ and τ on $(E \times F)^\top$, we want to “compose” them to get a closure $\sigma; \tau$ on $(D \times F)^\top$. We define this as follows:

$$\sigma; \tau(x, z) = (\pi_1 \circ \sigma(x, y), \pi_2 \circ \tau(y, z))$$

where

$$y = \langle \pi_2 \circ \sigma(x, -) | \pi_1 \circ \tau(-, z) \rangle.$$

That is, given input in D and F , we play σ and τ off against each other in E relative to this input, and obtain their external response taking into account their interaction with each other.

In particular, if σ is a closure on $(D \times E)^\top$, it induces an “action” taking closures on D^\top to closures on E^\top , $\alpha \mapsto \alpha; \sigma$, and a “coaction” taking closures on E^\top back to closures on D^\top , $\beta \mapsto \sigma; \beta$. E.g. $\sigma; \beta(x) = \pi_1 \circ \sigma(x, y)$, where $y = \langle \pi_2 \circ \sigma(x, -) | \beta \rangle$.

Now we can define

$$S_{A \otimes B}^* = \{ \sigma : (D_A \times D_B)^\top \mid \forall \alpha \in S_A. \alpha; \sigma \in S_B^* \wedge \forall \beta \in S_B. \sigma; \beta \in S_A^* \}.$$

Again, by De Morgan duality,

$$A \wp B = (A^\perp \otimes B^\perp)^\perp, \quad A \multimap B = A^\perp \wp B.$$

In particular,

$$S_{A \multimap B} = \{ \sigma : (D_A \times D_B)^\top \mid \forall \alpha \in S_A. \alpha; \sigma \in S_B \\ \wedge \forall \beta \in S_B^*. \sigma; \beta \in S_A^* \}.$$

This is a symmetric, “classical” version of the familiar logical relations condition: σ ’s action carries Player strategies on A to Player strategies on B , and its coercion carries counter-strategies on B to counter-strategies on A . We round out our account by defining the units, which all have the one-point poset as their underlying domain. For the tensor unit $\mathbf{1}$, $S_1 = \{(\perp \mapsto \perp)\}$, while $S_1^* = \{(\perp \mapsto \perp), (\perp \mapsto \top)\}$. For the unit for Plus (i.e. the initial object) we have $S_0 = \emptyset$, $S_0^* = \{(\perp \mapsto \top)\}$. The exponentials can also be defined in our model (the basic ideas are along the lines of [AJ92a]), but for lack of space we will refrain from doing so. The reader will really understand our model by checking that composition is associative, and seeing how the problems with Blass games simply do not arise in our setting.

We are now almost ready to define our category of concurrent games. Two further refinements are needed.

Stability Rather than taking all continuous closure operators as possible strategies, we will impose the domain-theoretic condition of *stability* [AC98]. (This will turn out to be important for our proof of Full Completeness, although we don’t know if it is a necessary condition; it is deeply related to the “monomial condition” in [Gir95].) But what does it mean for a closure operator to be stable? Define the *domain* $\text{dom}\sigma$ of a closure operator σ on D to be the set of all $x \in D$ such that $\sigma(x) \neq \top$. An *output function* for σ is a continuous function $f : D \rightarrow D$ such that, for all x in the domain of σ , $\sigma(x) = x \vee f(x)$. We say that σ is stable if it has a stable output function. (The link between concurrent games and the NFGOI [AJ92a] is here; the output functions for the strategies interpreting proofs will be exactly the denotations of proofs in [AJ92a]. Moreover, composition of closure operators is “tracked” by the composition of the corresponding output functions, defined as in [AJ92a].) We shall henceforth restrict ourselves to stable closure operators. This needs some conditions on the underlying domains: it suffices to assume that they are bounded complete, and distributive in a suitable sense. For details, see [AC98].

Extensionality To ensure that we get a genuine model of Linear Logic in which all the required equations hold, we adapt classical ideas from realizability to our setting. Rather than taking a game to be (D, S, S^*) , where S, S^* are simply predicates picking out the sets of strategies and counter-strategies, we will take games of the form (D, E, E^*) ,

where E and E^* are *partial equivalence relations* on the (stable) closure operators on D^\top . We simply adapt the definitions given above for the Linear connectives from unary predicates to binary relations, and check that they preserve symmetry and transitivity. If A is a game of this form, we write $\sigma : A$ to mean that $\sigma \in E_A$.

We are now in a position to define the category \mathcal{R} of concurrent games. Objects are structures (D, E, E^*) of a domain and two partial equivalence relations on stable closure operators, as already explained, subject to the following condition: if $\sigma : E$ and $\tau : E$ have the same maximal fixpoints, then $\sigma E \tau$, and similarly for E^* . A morphism from A to B is a closure operator σ such that $\sigma : A \multimap B$. Composition is defined as above. Identities are given by:

$$\text{id}_A(x, y) = (x \vee y, x \vee y).$$

They can be understood as “symmetric, bidirectional copy-cat strategies” as in [AJ92b]. We can then define \mathcal{E} as the “extensional quotient of \mathcal{R} ”, in which a morphism from A to B is a partial equivalence class of $E_{A \multimap B}$.

Proposition 2.1 \mathcal{E} is a \star -autonomous category with all limits and colimits.

Discussion To motivate the passage to the extensional category \mathcal{E} , note that \mathcal{R} only has weak products and coproducts. Indeed, the lifted sum which we used to model the additives is non-associative, and we need to quotient out the behaviour at the partial elements in order to obtain the required structure. This might lead the reader to wonder why we bother with the partial elements at all: why not just work with the maximal elements, which in effect means using a relational model as in [Loa94b]? In fact, *the relational model in [Loa94b] is not fully complete for MALL*—we have an explicit counter-example. Even though the behaviour of strategies on partial elements is factored out in \mathcal{E} , it still plays a crucial role in determining what the strategies are in the first place. We use it to cut the space of strategies down to the *stable* closure operators, and stability allows us to capture the causality between $\&$ -links and their contexts which is the key issue in MALL, and which is represented for example in the boolean weights used in MALL proof nets. In particular, we will find a beautiful correspondence between the *trace* of the stable output function for a strategy, and the monomial weights appearing in the proof net we shall “read back” from the strategy.

The model of MALL The fragment of Linear Logic we will consider in this paper consists of formulas built from propositional atoms X and their negations X^\perp with the binary connectives \otimes (Tensor), \wp (Par), $\&$ (With) and \oplus (Plus). We refer to this fragment as MALL. A sequent $\vdash A_1, \dots, A_k$ will be interpreted as the formula $A_1 \wp \dots \wp A_k$.

Since proofs of propositional formulas should be *uniform* over all substitution instances, we will treat propositional atoms as variables, so that in effect we are viewing propositional logic as the Π^1 fragment of second-order propositional logic. In a by-now standard fashion [BFSS88], a MALL formula in the propositional atoms X_1, \dots, X_n can be interpreted as a mixed-variance functor

$$F : (\mathcal{R}^{\text{op}} \times \mathcal{R})^n \longrightarrow \mathcal{R}.$$

The set of such functors will form the objects of a category $\mathcal{M}[n]$, and collectively we will obtain an indexed \star -autonomous category \mathcal{M} with all limits and colimits. This will provide the right algebraic structure to model propositional MALL.

We define the morphisms in $\mathcal{M}[n]$ in two stages. Firstly, a *dinatural family* for a functor F in $\mathcal{M}[n]$ is a family of strategies $(\sigma_{\vec{A}})_{\vec{A} \in \mathcal{R}^n}$ indexed by n -tuples of games such that, for all $\vec{A}, \sigma_{\vec{A}} E_F(\vec{A}) \sigma_{\vec{A}}$, and moreover for all tuples of strategies $\vec{\tau} : \vec{A} \rightarrow \vec{B}$,

$$\sigma_{\vec{A}}; F(\text{id}_{\vec{A}}, \vec{\tau}) = \sigma_{\vec{B}}; F(\vec{\tau}, \text{id}_{\vec{B}}).$$

We define a per E_F for each such functor by

$$(\sigma_{\vec{A}}) E_F(\tau_{\vec{A}}) \iff \forall \vec{A}. \sigma_{\vec{A}} E_F(\vec{A}) \tau_{\vec{A}},$$

and similarly for E_F^* . We then define morphisms in $\mathcal{M}[n]$ as partial equivalence classes of families of (stable, symmetric²) strategies, generalizing what we did for \mathcal{E} . Defining the various connectives and constructions for the Linear types pointwise on the functors and families in $\mathcal{M}[n]$, we obtain a \star -autonomous category with all limits and colimits.

There is a fine point here; as is well known, dinaturality is not in general preserved by composition. However, *after* we have proved Full Completeness—which will not use the assumption that dinaturality is preserved by composition (!!)—then by pulling back to syntax it will be easy to see that (for definable functors) closure under composition does hold, and hence each $\mathcal{M}[n]$ (restricted to definable functors) is indeed a category. We could avoid this logical detour by using a stronger property than dinaturality, namely Reynolds-style relational parametricity [BFSS88], for which closure under composition can be proved directly; however, this would complicate the description of the model, and dinaturality is sufficient to prove Full Completeness.

Another hypothesis we will need to prove Full Completeness is that the closure operators are *symmetric*, in the sense of having output functions f satisfying $f^3 = f$.

²By symmetric we mean a closure operator having an output function f satisfying $f^3 = f$.

Again, the only problem with this condition is proving closure under composition; and again, once Full Completeness is proved, we obtain closure under composition as a corollary.


Petri games We would like to give a very concrete description of the (dinatural) strategy $[\pi]$ interpreting a cut-free MALL proof π of the formula Γ in our games model. (This can be related formally to our abstract domain-theoretic description using the theory developed in [NPW81].) The first step is to see each position of a game as the state of a Petri net. Of course, this is only possible for very special games, which we call *Petri games* and define below.

Formally, a Petri net is a quadruple $N = (P, T, pre, post)$ where P is the set of *places* and T the set of *transitions*. To every transition t is associated two nonempty sets $pre(t) \subset P$ of *pre-conditions* and $post(t) \subset P$ of *post-conditions*. A *state* in a Petri net N is a subset of P . States are related by transition relations: We write $x \xrightarrow{t} y$ when $x = z + pre(t)$ and $y = z + post(t)$ for some state z and transition t , where $+$ means disjoint union. A state y is accessible from x , which we write $x \leq_N y$, when there exists a sequence of transitions t_1, \dots, t_n such that $x \xrightarrow{t_1} \dots \xrightarrow{t_n} y$.

A Petri net is called *unfolded* when

1. $pre(t)$ is a singleton for every $t \in T$,
2. $post(t_1) \cap post(t_2) \neq \emptyset$ implies $t_1 = t_2$, for every $t_1, t_2 \in T$.


Graphically, this means that the patterns (1) and (2) are forbidden, in other words that the Petri net N looks like a forest.



By restricting to unfolded Petri nets, the accessibility relation becomes an order. We call a *root* any place p of the unfolded Petri net N such that p does not appear in any $post(t)$, $t \in T$. We single out the state $\perp_N \subset P$ consisting of all root places of N , and call a state x of N *accessible* when it is accessible from that special state \perp_N . Now, the accessibility relation \leq_N defines a domain D_N on the set of accessible states of N . We call a *Petri game* any game A whose domain of positions D_A is of the form D_N for an unfolded Petri net N , and call N the *board* of A . Intuitively, a Petri game A with board N has \perp_N as least position, and every state x of N reachable by a sequence of transitions $\perp_N \xrightarrow{t_1} \dots \xrightarrow{t_n} x$ as a position.

Now, we show how to construct the board of $A \otimes B$ and $A \& B$ given the boards N_A and N_B of A and B . For the multiplicatives, $N_{A \otimes B} = N_A \otimes B$ is the Petri net obtained

by juxtaposing the Petri nets N_A and N_B . In particular, $\perp_{N_A \otimes B}$ is the union of \perp_{N_A} and \perp_{N_B} .




For the additives, $N_{A \& B} = N_{A \oplus B}$ is the juxtaposition of N_A and N_B with a new place p_* and two new transitions t_{left} and t_{right} with pre-conditions:

$$\text{pre}(t_{\text{left}}) = \text{pre}(t_{\text{right}}) = \{p_*\}$$


and post-conditions:

$$\text{post}(t_{\text{left}}) = \perp_{N_A} \quad \text{post}(t_{\text{right}}) = \perp_{N_B}$$


Graphically



For instance, given two Petri games A_1 and A_2 with boards N_{A_1} and N_{A_2} :



the Petri game $B = ((A_1 \otimes A_2) \& (A_1 \otimes A_2)) \wp A_1^\perp \wp A_2^\perp$ has board:



where the least position \perp_B is indicated by the three tokens.

The concrete interpretation Every formula Γ on the atoms X_1, \dots, X_n , is interpreted as a mixed-variance functor

$$[\Gamma] : (\mathcal{R}^{\text{op}} \times \mathcal{R})^n \longrightarrow \mathcal{R}.$$

and every proof π of Γ defines a strategy $[\pi]_{\vec{A}}$ of $\Gamma(\vec{A}, \vec{A})$ for every tuple $\vec{A} = A_1, \dots, A_n$ of games in \mathcal{R} . In the particular case of Petri games A_1, \dots, A_n , the strategy $[\pi]_{\vec{A}}$ plays

on the board $N_{\Gamma(\vec{A}, \vec{A})}$ associated to $\Gamma(\vec{A}, \vec{A})$ and defined from the boards N_{A_i} 's of the A_i 's as in the previous paragraph. Let us describe briefly how this strategy plays on $N_{\Gamma(\vec{A}, \vec{A})}$.

The first step is to consider that the proof π is translated as a proof-net with additive boxes, see [Gir87]. Observe that every place p of $N_{\Gamma(\vec{A}, \vec{A})}$ is associated either to an additive link L of Γ , or to the residual of a place among the N_{A_i} 's. In the first case, each of the valuations “left” and “right” of L corresponds to a transition t_{left}^L and t_{right}^L with pre-condition $\{p\}$.

By definition, every position $x \in D_{\Gamma(\vec{A}, \vec{A})}$ induces a transition sequence

$$\perp_{N_{\Gamma(\vec{A}, \vec{A})}} \xrightarrow{t_1} \dots \xrightarrow{t_n} x \quad (1)$$

which may not be unique, yet induces a partial valuation v_x of the $\&$ -links of Γ which does not depend on the sequence (1): The valuation assigns **left** or **right** to a $\&$ -links L when one of the transitions t_{left}^L or t_{right}^L appears among the t_i 's in the sequence (1).

Once this partial valuation v_x of the $\&$ -link of Γ is computed from x , a new proof-net π_x with additive boxes is constructed by removing every additive box of π whose principal door is a $\&$ -formula assigned a value by v_x .

At this point, the strategy $\sigma = [\pi]_{\vec{A}}$ determines its answer $\sigma(x)$ from the information it reads in π_x , considering every remaining additive box in π_x as a “black box” with no possibility to look inside.

Observe that the proof-net π_x verifies the two following fundamental properties:

1. every \oplus -link visible in π_x (*i.e.* not in a black box) is either **left** or **right**,
2. every visible literal of π_x is related to another visible literal with an axiom link.


Now, the position $\sigma(x)$ is defined as the least position y above x in $D_{\Gamma(\vec{A}, \vec{A})}$ such that:


1. the valuation of every \oplus -link in π_x appears in y ,
2. given an axiom link, the position on one side is the same as the position on the other side (concurrent copy-cat).

For instance, the only proof π of formula

$$((X_1 \otimes X_2) \& (X_1 \otimes X_2)) \wp X_1^\perp \wp X_2^\perp$$

is interpreted as a dinatural strategy $[\pi]$ whose instance $\sigma = [\pi]_{\vec{A}}$ on the Petri games $\vec{A} = A_1, A_2$ (see previous paragraph) is a strategy of $B = ((A_1 \otimes A_2) \& (A_1 \otimes A_2)) \wp A_1^\perp \wp A_2^\perp$, which plays as below:





3 Proof-structures from strategies

In this section, we construct the proof-structure associated to a dinatural family σ for Γ . A technical lemma (lemma 3.3) requires to restrict dinaturality to the category of civil games, defined as follows. A strategy $\sigma : A$ is *civil* when for every $\tau : A^\perp$, $\langle \sigma | \tau \rangle$ is not \top . A game $A = (D_A, E_A, E_A^*)$ is civil when it contains a civil strategy and a civil counter-strategy. A civil game is *total* when $\langle \sigma | \tau \rangle$ is maximal for every $\sigma : A$ and $\tau : A^\perp$. The two classes of civil and total games are closed under the MALL constructs \otimes , \wp , \oplus , $\&$ and $(-)^\perp$. Moreover, a strategy of $A \otimes B$ is a strategy of $A \wp B$ when A and B are total.

A \mathcal{C} -dinatural family for a functor Γ in $\mathcal{M}[n]$ is defined as a family of strategies $(\sigma_{\vec{A}})_{\vec{A} \in \mathcal{C}^n}$ indexed by tuples $\vec{A} = A_1, \dots, A_n$ of civil games such that, for all $\vec{A} \in \mathcal{C}^n$, $\sigma_{\vec{A}} : \Gamma(\vec{A}, \vec{A})$, and moreover

$$\sigma_{\vec{A}}; \Gamma(\text{id}_{\vec{A}}, \vec{\tau}) = \sigma_{\vec{B}}; \Gamma(\vec{\tau}, \text{id}_{\vec{B}}).$$

for all tuple $\vec{\tau} = \tau_1, \dots, \tau_n$ of morphisms $\tau_i : A_j \rightarrow B_j$ in \mathcal{C} . Such a \mathcal{C} -dinatural family σ is *uniform* when the closure operator $\sigma_{\vec{A}}$ depends only on the domain D_{A_j} for $1 \leq j \leq n$. The following lemma ensures that every dinatural family σ restricts to a uniform \mathcal{C} -dinatural family.

Lemma 3.1 (uniformity) *Let σ be a dinatural family and \vec{A} be a tuple A_1, \dots, A_n of games. The strategy $\sigma_{\vec{A}}$ does not depend on the pers E_{A_i} and $E_{A_i}^*$, only on the domains D_{A_i} .*

We fix in sections 3.1, 3.2 and 3.4:

- a formula Γ built from X_1, \dots, X_n and $Y_1^\perp, \dots, Y_n^\perp$ with the connectives \otimes , \wp , \oplus , $\&$,
- a uniform \mathcal{C} -dinatural family σ for the functor Γ in $\mathcal{M}[n]$.

Every valuation v of the $\&$ -links of Γ defines a $\text{M}\oplus\text{LL}$ formula Γ_v , also interpreted as a functor in $\mathcal{M}[n]$. We leave to the reader the inductive definition of the associated natural transformation

$$\text{proj}_v^\Gamma : \Gamma \rightarrow \Gamma_v$$

using the projection maps $\pi_1 : A \& B \rightarrow A$ and $\pi_2 : A \& B \rightarrow B$ in \mathcal{R} . Composing the \mathcal{C} -dinatural family σ with proj_v^Γ defines a \mathcal{C} -dinatural family for the $\&$ -free Γ_v . It is natural then to describe the \mathcal{C} -dinatural family σ in the special case of a $\text{M}\oplus\text{LL}$ formula Γ .

We need a few notations. An *annotated* formula is a pair (Γ, index) consisting of a MALL formula Γ and a one-to-one function index associating an integer i to every occurrence of X_j and Y_j^\perp in Γ . An annotated formula (Γ, index)

is best seen as a formula Γ where occurrences of X_j (resp. Y_j^\perp) associated to i are replaced by X_j^i (resp. $(Y_j^\perp)^i$).

Let m be the number of literal occurrences in Γ . All through this section, we consider the formula Γ annotated in $\{1, \dots, m\}$ in such a way that each index i appears to the left of the index $i + 1$. Given this ‘‘canonical’’ annotation, two functions $\psi : \{1, \dots, m\} \rightarrow \{1, \dots, n\}$ and $\epsilon : \{1, \dots, m\} \rightarrow \{+, -\}$ indicate which literal the index i annotates: $X_{\psi(i)}^i$ when $\epsilon(i) = +$, and $(Y_{\psi(i)}^\perp)^i$ when $\epsilon(i) = -$.

3.1 The multiplicative and $\text{M}\oplus\text{LL}$ fragments

Starting with the case of a multiplicative formula Γ , we construct the multiplicative proof-structure Θ_σ associated to the \mathcal{C} -dinatural family σ . This construction does not require any of the game-theoretic properties yet, and we simply follow the steps of R. Loader who carried out the construction in the relational model, see [Loa94b].

Lemma 3.2 *There exists a fixpoint-free involution $\phi : \{1, \dots, m\} \rightarrow \{1, \dots, m\}$ such that for any tuple $\vec{A} = A_1, \dots, A_n$ of civil games, the set of maximal fixpoints of $\sigma_{\vec{A}}$ is*

$$\{(x_1, \dots, x_m) \mid x_i = x_{\phi(i)} \text{ for all } 1 \leq i \leq m\}$$

We suppose now that the formula Γ is constructed from X_1, \dots, X_n and $Y_1^\perp, \dots, Y_n^\perp$ with the connectives \otimes , \wp and \oplus . We associate a $\&$ -free proof-structure to σ by reducing that problem to the multiplicative case.

By duality, a total valuation v of the \oplus -links of Γ is also a total valuation w of the $\&$ -links of Γ^\perp . Thus, the valuation v induces a functor $\Gamma_v : (\mathcal{C}^{\text{op}} \times \mathcal{C})^n \rightarrow \mathcal{C}$ and a natural transformation $\text{inj}_v^\Gamma : \Gamma_v \rightarrow \Gamma$ defined as:

$$\Gamma_v = ((\Gamma^\perp)_w)^\perp \quad \text{inj}_v^\Gamma = (\text{proj}_w^{\Gamma^\perp})^\perp$$

The main result of the section (factorization) implies with lemma 3.2 that the uniform \mathcal{C} -dinatural family σ describes a $\text{M}\oplus\text{LL}$ proof-structure.

Lemma 3.3 (factorization) *Given a \mathcal{C} -dinatural family σ for the $\&$ -free Γ , there exists a valuation v of the \oplus -link of Γ such that the morphism $\sigma_{\vec{A}}$ factors uniquely as*

$$\mathbf{1} \xrightarrow{\tau_{\vec{A}}} \Gamma_v(\vec{A}, \vec{A}) \xrightarrow{(\text{inj}_v^\Gamma)_{\vec{A}, \vec{A}}} \Gamma(\vec{A}, \vec{A}) \quad (2)$$

at every instance \vec{A} of civil games. Moreover, the family τ for Γ_v defined in (2) is uniform and \mathcal{C} -dinatural.

PROOF We take the total game $\Delta = \Delta^\perp$ with domain $\{\perp\}$ and play σ_Δ off against the counterstrategy $(x \mapsto x)$. The maximal element we obtain as a result defines the valuation

v . We then prove factorization (2) at every instance by establishing that $(\text{inj}_v^\Gamma)_{\vec{A}, \vec{B}}$ is a split mono when \vec{A} and \vec{B} are tuples of *civil* games. This is where the restriction from dinaturality to \mathcal{C} -dinaturality appears in the proof. ■

3.2 Notations

Given an instance of X_j^i (resp. Y_j^\perp) in the annotated formula Γ , and tuples \vec{A}, \vec{B} of games, we associate to any element $x \in D_{A_j}$ (resp. $x \in D_{B_j}$) the corresponding element $(X_j^i \mapsto x)_{\vec{A}, \vec{B}}^\Gamma$ (resp. $((Y_j^\perp)^i \mapsto x)_{\vec{A}, \vec{B}}^\Gamma$) of $D_{\Gamma(\vec{A}, \vec{B})}$ as follows, (the definition by induction is similar for $(Y_j^\perp)^i$):

$$(X_j^i \mapsto x)_{\vec{A}, \vec{B}}^{X_{j'}^{i'}} = \begin{cases} x & \text{if } (i, j) = (i', j') \\ \perp & \text{otherwise} \end{cases}$$

$$\begin{aligned} (X_j^i \mapsto x)_{\vec{A}, \vec{B}}^{\Gamma_1 \otimes \Gamma_2} &= (X_j^i \mapsto x)_{\vec{A}, \vec{B}}^{\Gamma_1} \wp \Gamma_2 \\ &= \begin{cases} ((X_j^i \mapsto x)_{\vec{A}, \vec{B}}^{\Gamma_1}, \perp) & \text{if } X_j^i \text{ is in } \Gamma_1 \\ (\perp, (X_j^i \mapsto x)_{\vec{A}, \vec{B}}^{\Gamma_2}) & \text{if } X_j^i \text{ is in } \Gamma_2 \end{cases} \end{aligned}$$

$$\begin{aligned} (X_j^i \mapsto x)_{\vec{A}, \vec{B}}^{\Gamma_1 \oplus \Gamma_2} &= (X_j^i \mapsto x)_{\vec{A}, \vec{B}}^{\Gamma_1 \& \Gamma_2} \\ &= \begin{cases} (\text{inl}(X_j^i \mapsto x)_{\vec{A}, \vec{B}}^{\Gamma_1})_\perp & \text{if } X_j^i \text{ is in } \Gamma_1 \\ (\text{inr}(X_j^i \mapsto x)_{\vec{A}, \vec{B}}^{\Gamma_2})_\perp & \text{if } X_j^i \text{ is in } \Gamma_2 \end{cases} \end{aligned}$$

The definition of $(X_j^i \mapsto x)_{\vec{A}, \vec{B}}^\Gamma$ is easily adapted to associate a prime element $(L \mapsto \text{left})_{\vec{A}, \vec{B}}^\Gamma$ of $D_{\Gamma(\vec{A}, \vec{B})}$ to the valuation $(L \mapsto \text{left})$ of an additive link L of Γ . The base case of the induction is:

$$(L \mapsto \text{left})_{\vec{A}, \vec{B}}^{\Gamma_1 L \Gamma_2} = (\text{inl} \perp)_\perp$$

The prime element $(L \mapsto \text{right})_{\vec{A}, \vec{B}}^\Gamma$ corresponding to $(L \mapsto \text{right})$ is defined similarly.

To every partial valuation v of the $\&$ -links of Γ , to every m -tuple $\vec{\alpha} = \alpha_1, \dots, \alpha_m$ and $\vec{\beta} = \beta_1, \dots, \beta_m$ of closure operators $\alpha_i : D_{A_{\psi(i)}}^\Gamma$ and $\beta_i : D_{B_{\psi(i)}}^\Gamma$, we associate a closure operator

$$[\alpha, \beta]_v^\Gamma : D_{\Gamma(\vec{A}, \vec{B})}^\Gamma$$

as follows:

- $[\alpha, \beta]_\emptyset^{X_j^i} = \alpha_i$ and $[\alpha, \beta]_\emptyset^{(Y_j^\perp)^i} = \beta_i$,
- Letting v_1 and v_2 be the respective restrictions of v to Γ_1 and Γ_2 :

$$[\alpha, \beta]_v^{\Gamma_1 \wp \Gamma_2} = [\alpha, \beta]_v^{\Gamma_1 \otimes \Gamma_2}$$

is equal to the smash product

$$[\alpha, \beta]_{v_1}^{\Gamma_1} \times [\alpha, \beta]_{v_2}^{\Gamma_2}$$

$$[\alpha, \beta]_v^{\Gamma_1 \oplus \Gamma_2} : \begin{aligned} \perp &\mapsto \perp, \text{inl}x \mapsto \text{inl}[\alpha, \beta]_{v_1}^{\Gamma_1}x, \\ \text{inr}x &\mapsto \text{inr}[\alpha, \beta]_{v_2}^{\Gamma_2}x. \end{aligned}$$

When v does not assign a value to the root $\&$ -link:

$$[\alpha, \beta]_v^{\Gamma_1 \& \Gamma_2} : (x, y) \mapsto (x, y)$$

When v assigns the value `left` to the root $\&$ -link:

$$[\alpha, \beta]_v^{\Gamma_1 \& \Gamma_2} : \begin{aligned} \perp &\mapsto \text{inl}[\alpha, \beta]_{v_1}^{\Gamma_1} \perp, \\ \text{inl}x &\mapsto \text{inl}[\alpha, \beta]_{v_1}^{\Gamma_1}x, \\ \text{inr}x &\mapsto \top. \end{aligned}$$

When v assigns the value `right` to the root $\&$ -link:

$$[\alpha, \beta]_v^{\Gamma_1 \& \Gamma_2} : \begin{aligned} \perp &\mapsto \text{inr}[\alpha, \beta]_{v_2}^{\Gamma_2} \perp, \\ \text{inl}x &\mapsto \top, \\ \text{inr}x &\mapsto \text{inr}[\alpha, \beta]_{v_2}^{\Gamma_2}x. \end{aligned}$$

The motivation for this construction is that given n -tuples \vec{A} and \vec{B} of *total* games, and a *total* valuation v , the closure operator $[\alpha, \beta]_v^\Gamma$ is a strategy of $\Gamma(\vec{A}, \vec{B})^\perp$ when the α_i 's and β_i 's are strategies $\alpha_i : A_{\psi(i)}^\perp, \beta_i : B_{\psi(i)}^\perp$ for all $1 \leq i \leq n$.

3.3 MALL proof-structures

We recall the definition of a MALL proof-structure in [Gir95]. A proof-structure Θ consists of:

1. a set of formula occurrences,
2. a set of links; each of these links takes its premise(s) and conclusion(s) among the formula occurrences of Θ ;
3. for each formula occurrence A of Θ , a weight $w(A)$, i.e. a non-zero element of the boolean algebra generated by the eigenweights p_1, \dots, p_n of the $\&$ -links of Θ

satisfying the following conditions:

1. each formula occurrence is the premise of at most one link and the conclusion of at least one link
2. if A is a conclusion of Θ , then $w(A) = 1$,
3. if w is any weight occurring in Θ , then w is a monomial of eigenweights and negations of eigenweights,
4. if $\epsilon p_L w$ is any weight occurring in Θ , then $w \subset w(L)$,
5. $w(A) = \sum w(L)$, the sum being taken over the set of links with conclusion A . It is required here that the sum is disjoint.
6. $w(L) \neq 0$, moreover, if L is any non-axiom link, with premises A (or) B then

- if L is any \otimes or \wp , then $w(L) = w(A) = w(B)$,
- if L is \oplus_1 , then $w(L) = w(A)$,
- if L is \oplus_2 , then $w(L) = w(B)$,
- if L is a $\&$ -link, then $w(A) = w(L).p_L$ and $w(B) = w(L).\neg p_L$

Our condition 4 is equivalent to the requirement appearing in [Gir95] that for any element of the boolean algebra generated by the weights occurring in Θ , and any $\&$ -link L , $w.\neg w(L)$ does not depend on p_L .

3.4 Main construction

We associate an event structure EVENT_Γ to the formula Γ canonically annotated. So, Γ is a tree of $\otimes, \oplus, \wp, \&$ -links and MALL formulas, with the formula Γ at its root, the literals X_j^i and $(Y_j^\perp)^i$ at its leaves, and every path an alternating sequence of formulas and links.

If we replace every subtree $\oplus(A, B)$ by two subtrees $\oplus_1(A)$ and $\oplus_2(B)$ connected to the same ancestor, the resulting tree TREE_Γ is labelled with $\otimes, \wp, \oplus_1, \oplus_2, \&$ -links and MALL formulas (each literal annotated). A *node* in TREE_Γ means either a link or a formula in the tree. Nodes are ordered by the tree-nesting ordering. An *axiom link* in TREE_Γ is a 2-element set $\{X_j^i, (Y_j^\perp)^{i_2}\}$ of annotated literals in TREE_Γ . An additive boundary is a $\&$ -link or a formula $A \oplus B$ in TREE_Γ .

★ The event structure EVENT_Γ is defined as follows:

1. the events are the nodes and axiom links of TREE_Γ ,
2. $e \leq e'$ when e and e' are nodes and e nests e' , or when e is a node and e' is an axiom link containing a literal nested by e , or e and e' are the same axiom link.
3. $e \# e'$ when e and e' are nodes and $e \wedge e'$ is an additive boundary, when e is a node and e' is an axiom link containing a literal X_j^i such that $e \# X_j^i$, when e and e' are axiom links whose intersection $e \cap e'$ is singleton, or when e and e' are axiom links containing incompatible literals.

Observe that every maximal state in EVENT_Γ describes a multiplicative proof-structure.

The equivalence relation \equiv relates two nodes of TREE_Γ when there is a path between them that does not cross an additive boundary. More formally, it is the least equivalence relation on the events of EVENT_Γ such that:

- $L \equiv A$ when A is a formula premise of L not a $\&$ -link,
- $A \equiv L$ when A is the conclusion of a \otimes, \wp or $\&$ -link L .

In particular, every axiom link is the only element of its \equiv -class.

★ Let \mathbb{Z} be the total game associated based on the flat domain of integers with strategies all $(\perp \mapsto n), n \in \mathbb{Z}$. Let \mathbb{Z}_2 be $\mathbb{Z} \wp \mathbb{Z}^\perp$. This instance \mathbb{Z}_2 will play a particular role in the construction of the proof-structure corresponding to σ . Every prime element of the form

$$(X_j^{i_1} \mapsto (i_2, \perp))_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma \quad \text{or} \quad ((Y_j^\perp)^{i_1} \mapsto (\perp, i_2))_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma$$

is called an input of $D_{\Gamma(\mathbb{Z}_2, \mathbb{Z}_2)}$ and every prime element of the form

$$(X_j^{i_1} \mapsto (\perp, i_2))_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma \quad \text{or} \quad ((Y_j^\perp)^{i_1} \mapsto (i_2, \perp))_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma$$

an output of $D_{\Gamma(\mathbb{Z}_2, \mathbb{Z}_2)}$.

Given a valuation v of the $\&$ -links in Γ , the closure operator γ_v on $D_{\Gamma(\mathbb{Z}_2, \mathbb{Z}_2)}^\Gamma$ is defined as

$$\gamma_v = [\alpha, \beta]_v^\Gamma$$

(here Γ is considered canonically annotated) with the strategies $\alpha_i : \mathbb{Z}_2$ and $\beta_i : \mathbb{Z}_2^\perp$ for all $1 \leq i \leq m$:

$$\alpha_i : (x, y) \mapsto (x \vee i, y) \quad \beta_i : (x, y) \mapsto (x, y \vee i)$$

Note that γ_v is a counter-strategy of $\Gamma(\mathbb{Z}_2, \mathbb{Z}_2)$ when the valuation v is total. Moreover,

Lemma 3.4 (stability) *The map $v \mapsto \langle \sigma_{\mathbb{Z}_2} | \gamma_v \rangle$ from valuations to $D_{\Gamma(\mathbb{Z}_2, \mathbb{Z}_2)}$ is stable.*

PROOF The proof relies heavily on the connection between our model and NFGOI, the definition of stable closure operators, the fact that the category of domains and stable maps is cartesian-closed, and finally the choice of the α_i 's and β_i 's. ■

★ We introduce two *linear* functions EV_Γ^+ and EV_Γ^- that associate a state of EVENT_Γ to any element of $D_{\Gamma(\mathbb{Z}_2, \mathbb{Z}_2)}$. These functions are defined as follows:

- \perp is sent to the union of the \equiv -class of the conclusion,
- for $L = \&$, the prime $(L \mapsto \text{left})_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma$ is sent to the \equiv -class of L 's first premise, and the prime $(L \mapsto \text{right})_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma$ to the \equiv -class of L 's second premise,
- for $L = \oplus$, the prime $(L \mapsto \text{left})_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma$ is sent to the \equiv -class of the corresponding \oplus_1 -link, and the prime $(L \mapsto \text{right})_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma$ is sent to the \equiv -class of the corresponding \oplus_2 -link,
- the output prime $(X_j^{i_1} \mapsto (\perp, i_2))_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma$ for an index i_2 of some literal $(Y_j^\perp)^{i_2}$ is sent by EV_Γ^+ to the axiom link $\{X_j^{i_1}, (Y_j^\perp)^{i_2}\}$,

- the output prime $((Y_j^\perp)^{i_1} \mapsto (i_2, \perp))_{\mathbb{Z}_2, \mathbb{Z}_2}^\Gamma$ for an index i_2 of some literal $X_j^{i_2}$ is sent by EV_Γ^- to the axiom link $\{(Y_j^\perp)^{i_1}, X_j^{i_2}\}$,
- any other prime is sent to the empty state.

The following (stable) function PRF_σ^+ associates to any partial valuation v of the $\&$ -links in Γ a state in EVENT_Γ :

$$\text{PRF}_\sigma^+ : v \mapsto \text{EV}_\Gamma^+ \langle \sigma_{\mathbb{Z}_2} | \gamma v \rangle$$

★ The proof-structure Θ_σ is constructed as follows:

1. its formula occurrences (resp. links) are the elements of $\text{tr}(\text{PRF}_\sigma^+)$, the *trace*³ of the stable function PRF_σ^+ , of the form (e, v) where e is a formula occurrence (resp. a link) of EVENT_Γ . Every node (e, v) is labelled as e ,
2. the weight $w(e, v)$ is computed from the valuation v as follows:

$$w(e, v) = \prod p_{(L, v_1)} \times \prod \neg p_{(L, v_2)}$$

where the first product is taken over the $\&$ -links (L, v_1) such that v extends $v_1 + (L \mapsto \text{left})$, and the second product over the $\&$ -links (L, v_2) such that v extends $v_2 + (L \mapsto \text{right})$.

3. a formula occurrence (A, v_1) is the premise of a link (L, v_2) when A is a premise of L in Γ and v_1 extends v_2 ,
4. a formula occurrence (A, v_1) is the conclusion of a link (L, v_2) when A is a conclusion of L in Γ and v_2 extends v_1 .

For lack of space, we omit the lengthy verification that

Lemma 3.5 Θ_σ defines a MALL proof-structure.

We also omit the proof that the alternative definition of Θ_σ as the trace of

$$\text{PRF}_\sigma^- : v \mapsto \text{EV}_\Gamma^- \langle \sigma_{\mathbb{Z}_2} | \gamma v \rangle$$

leads to the same proof-structure when the strategy $\sigma_{\mathbb{Z}_2}$ is symmetric.

4 Correctness criteria

4.1 MALL proof-nets

Let φ be a valuation of the $\&$ -links of a proof-structure Θ . φ induces a function (still denoted φ) from the weights of Θ to $\{0, 1\}$. The *slice* $\varphi(\Theta)$ is obtained by restricting Θ to the formulas A verifying $\varphi(w(A)) = 1$, with the obvious modification for the remaining $\&$ -links: only one premise is present.

A switching of a proof-structure Θ consists in

³It is here that stability is used in our proof.

1. the choice of a valuation φ_S for Θ ,
2. the selection of a choice $\mathcal{S}(L) \in \{l, r\}$ for all \wp -links of $\varphi_S(\Theta)$
3. the selection for each $\&$ -link L of $\varphi_S(\Theta)$ of an occurrence $\mathcal{S}(L)$, the *jump* of L , depending on p_L in $\varphi_S(\Theta)$. There is always a *normal* choice of jump for L , namely the premise A of L such that $\varphi_S(w(A)) = 1$. Any other choice is called *proper*.

A *normal switching* is a switching with no proper jump.

Definition 4.1 Let \mathcal{S} be a switching of a proof-structure Θ . We define the graph $\Theta_{\mathcal{S}}$ as follows:

1. the vertices of $\Theta_{\mathcal{S}}$ are the occurrences and links of $\varphi_S(\Theta)$,
2. for all links of $\varphi_S(\Theta)$, we draw an edge between the link and each of its conclusions,
3. for all \oplus_i -links of $\varphi_S(\Theta)$, we draw an edge between the link and its premise,
4. for all \otimes -links of $\varphi_S(\Theta)$, we draw an edge between the link and its left premise, and between the link and its right premise,
5. for all \wp -links L of $\varphi_S(\Theta)$, we draw an edge between the link and the premise (left or right) selected by $\mathcal{S}(L)$,
6. for all $\&$ -links L of $\varphi_S(\Theta)$, we draw an edge between the link and the jump $\mathcal{S}(L)$ selected by \mathcal{S} .

A *proof-net* is a proof-structure Θ such that for all switchings \mathcal{S} , the induced graph is acyclic and connected.

4.2 Acyclicity

We prove in this section and the next that the proof-structure Θ_σ defined in section 3 satisfies the two conditions to be a proof-net: acyclicity by exhibiting a deadlock on the assumption of a cycle, connectedness by propagating the “error” value \top around the net on the assumption that there is more than one connected component.

We fix in subsections 4.2 and 4.3:

- a formula Γ built from X_1, \dots, X_n and $Y_1^\perp, \dots, Y_n^\perp$ with the connectives $\otimes, \wp, \oplus, \&$,
- a symmetric dinatural family σ for the functor Γ in $\mathcal{M}[n]$.

Given two total games A and B , and two strategies $\sigma : A^\perp$ and $\tau : B^\perp$ such that σ is stable as a function $D_A \rightarrow D_A^\perp$, we construct the strategy $(\sigma \triangleright \tau) : (A \otimes B)^\perp$ as follows:

$$(x, y) \mapsto \begin{cases} \top & \text{if } \sigma x = \top \text{ or } \tau x = \top \\ (\sigma x, \tau y) & \text{if } \sigma x \neq \top \text{ is not maximal} \\ (\sigma x, y) & \text{otherwise} \end{cases}$$

Let T be any set of disjoint \otimes -links in the formula Γ , v a total valuation of the $\&$ -nodes of Γ , strategies $\alpha_i : A_i$ and $\beta_i : B_i^\perp$ stable as functions $D_{A_i} \rightarrow D_{A_i}^\perp$ and $D_{B_i} \rightarrow D_{B_i}^\perp$ for all $1 \leq i \leq m$. The inductive definition of the strategy

$$\lceil \alpha, \beta \rceil_{v, T}^\Gamma : \Gamma(A, B)^\perp$$

is similar to the definition of $\lceil \alpha, \beta \rceil_v^\Gamma : \Gamma(A, B)^\perp$, except for the case of a \otimes -link in T :


• letting v_1 and v_2 (resp. T_1 and T_2) be the respective restrictions of v (resp. of S) to Γ_1 and Γ_2 :

$$\lceil \alpha, \beta \rceil_{v, T}^{\Gamma \otimes \Delta} = \begin{cases} \lceil \alpha, \beta \rceil_{v_1, T_1}^{\Gamma_1} \triangleright \lceil \alpha, \beta \rceil_{v_2, T_2}^{\Gamma_2} & \text{if } \otimes \text{ is in } T \\ \lceil \alpha, \beta \rceil_{v_1, T_1}^{\Gamma_1} \times \lceil \alpha, \beta \rceil_{v_2, T_2}^{\Gamma_2} & \text{otherwise} \end{cases}$$


For instance, $\lceil \alpha, \beta \rceil_v^\Gamma = \lceil \alpha, \beta \rceil_{v, T}^\Gamma$ when T is empty.

Lemma 4.2 *Every switching of Θ_σ is acyclic.*

PROOF (sketch) For lack of space, we omit the nice graph-theoretic reasoning which shows that, given a switching \mathcal{S} of Θ_σ , every potential cycle in the “switching” graph $(\Theta_\sigma)_\mathcal{S}$ may be transformed into an “oriented” cycle \mathcal{C} of the following form (oriented means that all jumps are headed in the same direction):



Let us label \otimes_k each \otimes -link appearing as a “rebounding” tensor in the figure: $\otimes_0, \dots, \otimes_{N-1}$. We have to consider two possible cases for any $k \in \{0, \dots, N-1\}$:



In the first case, each of the literals $X_j^{i_1}$ and $(Y_j^\perp)^{i_2}$ forming the axiom link is nested by one of \otimes_k or \otimes_{k+1} . We define p_{2k} (resp. p_{2k+1}) as the input prime (resp. output prime) of the literal nested by \otimes_k (resp. \otimes_{k+1}). In the second case, \otimes_k nests the $\&$ -link L , \otimes_{k+1} nests one of the literals forming the axiom link, and the jump from L to the axiom link connects the cycle. We define p_{2k+1} as the output prime of the literal nested by \otimes_{k+1} , and p_{2k} as the prime $(L \mapsto \square)_{\vec{z}_2, \vec{z}_2}^\Gamma$, where \square is the valuation of L (left or right) compatible with the valuation v associated to \mathcal{S} .

We prove that the cycle \mathcal{C} does not exist by constructing a strategy $\tau : \Gamma(\vec{z}_2, \vec{z}_2)^\perp$ whose interaction against $\sigma_{\vec{z}_2}$ does

not reach a maximal element (i.e. deadlocks), contradicting the totality of $\Gamma(\vec{z}_2, \vec{z}_2)$. Calling T the set $\{\otimes_0, \dots, \otimes_{N-1}\}$ the strategy τ is defined as $\lceil \alpha, \beta \rceil_{v, T}^\Gamma$ where the α_i 's and β_i 's are defined as in section 3.4. Suppose that x is an element of $D_{\Gamma(\vec{z}_2, \vec{z}_2)}$ verifying

$$\forall k \in \{0, \dots, 2N-1\}, \quad p_k \not\leq x \leq \langle \gamma_v | \sigma_{\vec{z}_2} \rangle.$$

Observe that every p_{2k+1} is an output and every p_{2k} an input or the prime associated to the valuation of a $\&$ -link of Γ . This implies that for all $0 \leq k \leq N-1$, $p_{2k+1} \not\leq \tau x$ and $p_{2k} \not\leq \sigma_{\vec{z}_2} x$. We then deduce from the construction of Θ_σ and the definition of τ that:

$$\forall k \in \{0, \dots, 2N-1\}, \quad \begin{cases} p_k \not\leq \sigma_{\vec{z}_2} x \leq \langle \gamma_v | \sigma_{\vec{z}_2} \rangle, \\ p_k \not\leq \tau x \leq \langle \gamma_v | \sigma_{\vec{z}_2} \rangle. \end{cases}$$

We conclude that $\langle \tau | \sigma_{\vec{z}_2} \rangle$ is not maximal. \blacksquare

4.3 Connectedness

Lemma 4.3 *Every normal switching of Θ_σ is connected.*

PROOF (sketch) By projection, the lemma reduces to proving that every $\mathbf{M}\oplus\mathbf{LL}$ proof-structure defined from the semantics has connected switchings. We restrict ourself to the multiplicative case here, for lack of space. Although the adaptation of the proof to $\mathbf{M}\oplus\mathbf{LL}$ is not entirely straightforward, the essence of the proof appears below.

Let $(\sigma_{\vec{A}})_{\vec{A} \in \mathcal{R}^n}$ be a dinatural family for a multiplicative formula Γ interpreted in $\mathcal{M}[n]$, $\Theta = \Theta_\sigma$ the corresponding proof-structure, and \mathcal{S} a switching such that $\Theta_\mathcal{S}$ is non-connected.

Composing σ with the natural transformations $A \otimes (B \mathfrak{Z} C) \xrightarrow{\cdot} B \mathfrak{Z} (A \otimes C)$ and $A \otimes (B \mathfrak{Z} C) \xrightarrow{\cdot} C \mathfrak{Z} (A \otimes B)$ transform σ into a dinatural family v for a functor Υ corresponding to the switching \mathcal{S} . Here, Υ is a L -fold \mathfrak{Z} -product of \otimes -products of literals, and the proof-structure associated to v by lemma 3.2 is homeomorphic to the “switching” graph $\Theta_\mathcal{S}$.

Define the game $A = A^\perp$ with D_A the flat domain of integers, and strategies the closure operators $(x \mapsto x \vee i)$, for i any integer, with E_A the equality on them. Observe that A is not total, not even civil. We use the notation $A^{\otimes k}$ or $A^{\mathfrak{Z} k}$ to mean respectively the k -fold \otimes or \mathfrak{Z} -product of a game A .

For every pair of integers $k \geq 1$ and N , there is a strategy ρ_N^k of $A^{\mathfrak{Z} k}$ whose set of maximal fixpoints is exactly:

$$\{(x_1, \dots, x_k) \mid \sum_{i=1}^{i=k} x_i = N\}$$

Now, suppose that $\Theta_\mathcal{S}$ is disconnected, and that $\Upsilon(\vec{A}, \vec{A})$ is the \mathfrak{Z} -product of $A^{\otimes k_1} \mathfrak{Z} \dots \mathfrak{Z} A^{\otimes k_L}$. We may assume

without loss of generality that the first l \otimes -powers $A^{\otimes k_i}$ and axiom links form a connected component of Θ_S .

Set the strategy $\rho : (A^{\otimes k_1} \wp \dots \wp A^{\otimes k_{L-1}})^\perp$ as

$$\rho = \rho_{N_1}^{k_1} \times \dots \times \rho_{N_{L-1}}^{k_{L-1}} \quad (4)$$

Here, we are careful to choose N_1, \dots, N_l such that

$$\sum_{i=1}^{i=l} N_i \text{ is odd} \quad (5)$$

Let α be any counter-strategy of $A^{\otimes k_L}$. A simple argument shows that the result of playing $\rho \times \alpha : \Upsilon(\vec{A}_m, \vec{A}_m)^\perp$ against $v_{\vec{A}}$ is either a maximal element (x_1, \dots, x_m) or \top . Assume that it is a maximal element (x_1, \dots, x_m) . By lemma 3.2, there exists an involution ϕ of $\{1, \dots, m\}$ such that $x_i = x_{\phi(i)}$ for every $i \in \{1, \dots, m\}$. The involution ϕ describes the axiom links of Θ_S , therefore restricts by definition of l to an involution of $\{1, \dots, k_{\leq l}\}$ where $k_{\leq l} = \sum_{i=1}^{i=l} k_i$. We obtain that the sum $\sum_{i=1}^{i=k_{\leq l}} x_i$ is even, which forbids by (5) (x_1, \dots, x_m) to be a fixpoint of $\rho \times \alpha$. We conclude that the result of playing $\rho \times \alpha$ against $v_{\vec{A}}$ is necessarily \top .

Let μ be the strategy of $A^{\otimes k_L}$ obtained by composing ρ and $v_{\vec{A}}$. We have just shown that \top is the result of playing μ off against any strategy $\alpha : A^{\wp k_L}$. This implies that μ is $(x \mapsto \top)$, hence not a strategy of $A^{\otimes k_L}$, yielding the required contradiction. ■

Theorem 4.4 Θ_σ is a MALL proof-net.

PROOF Lemmas 4.2 and 4.3 imply by a simple graph-theoretic argument that every (not necessarily normal) switching of Θ_σ is connected. We conclude. ■

5 Main result

Theorem 5.1 (Full Completeness) Let Γ and Δ be MALL formulas. Every morphism $\llbracket \Gamma \rrbracket \longrightarrow \llbracket \Delta \rrbracket$ in $\mathcal{M}[n]$ is the interpretation $\llbracket \Pi \rrbracket$ of a MALL proof Π of the formula $\Gamma \multimap \Delta$.

PROOF By theorem 4.4 and the Sequentialization Theorem in [Gir95], every symmetric dinatural family σ in the per-class $[\sigma] : \llbracket \Gamma \rrbracket \longrightarrow \llbracket \Delta \rrbracket$ defines a proof Π of $\Gamma \multimap \Delta$. In turn, the proof Π is interpreted as a per-class $[\tau] = \llbracket \Pi \rrbracket$ of symmetric dinatural families from $\llbracket \Gamma \rrbracket$ to $\llbracket \Delta \rrbracket$. The proof is reduced to showing that the two classes $[\sigma]$ and $[\tau]$ are equal. A simple argument on the slices of Π shows that σ and τ have the same maximal fixpoints at every instance \vec{A} . We then use the fact that two strategies of a game are equivalent when they have the same maximal fixpoints, and conclude that $[\sigma] = [\tau]$. ■

References

- [AJ92a] S. Abramsky, R. Jagadeesan, New foundations for the geometry of interaction, *Information and Computation*, 111(1):53-119, 1994. Conference version appeared in LiCS '92.
- [AJ92b] S. Abramsky and R. Jagadeesan. Games and Full Completeness for Multiplicative Linear Logic, *Journal of Symbolic Logic*, (1994), vol. 59, no. 2, 543–574. Conference version appeared in FSTTCS '92.
- [AJ94] S. Abramsky and A. Jung. Domain Theory. Handbook of Logic in Computer Science. Volume III: Semantic Structures, pages 1–168. S. Abramsky, D. M. Gabbay and T. S. E. Maibaum, editors. Oxford University Press 1994.
- [AM98] S. Abramsky and P.-A. Mellès. Concurrent games and full completeness. Draft paper, 37 pages, 1998.
- [AC98] R. Amadio, P.-L. Curien, *Domains and Lambda-calculi*. Cambridge University Press, 1998.
- [BDER97] P. Bailloot and V. Danos and T. Ehrhard and L. Regnier. Believe it or not, AJM's games is a Model of Classical Linear Logic. *Proceedings of LiCS '97*, pp. 68–75, 1997.
- [BFSS88] E. S. Bainbridge and P. J. Freyd and A. Scedrov and P. J. Scott. Functorial Polymorphism. *Theoretical Computer Science* 70,1:35–64, 1988.
- [Bla92] A. Blass, A game semantics for linear logic, *Annals of Pure and Applied Logic*, 56:183-220, North-Holland, 1992.
- [BS96] R. Blute and P. J. Scott. Linear Läuchli Semantics. *Annals of Pure and Applied Logic*, 1996.
- [DJS97] V. Danos and J.-B. Joinet and H. Schellinx. A new deconstructive logic: Linear Logic. *Journal of Symbolic Logic*, 1997.
- [Gir87] J.-Y. Girard. Linear Logic. *Theoretical Computer Science*, 50:1–102, 1987.
- [Gir95] J.-Y. Girard. Proof-nets: the parallel syntax for proof theory. 1995.
- [Gir98a,b] J.-Y. Girard. On the Meaning of Logical Rules I, II. Unpublished papers, 1998.
- [HO92] M. Hyland and C.-H. L. Ong. Fair Games and Full Completeness for Multiplicative Linear Logic without the MIX rule. Unpublished Manuscript, 1992.
- [Loa94a] R. Loader. Linear Logic, Totality and Full Completeness. In *Proceedings of LiCS '94*.
- [Loa94b] R. Loader, *Models of lambda calculi and linear logic: structural, equational and proof-theoretic characterisations*, PhD Thesis, Oxford, 1994.
- [NPW81] M. Nielsen and G. Plotkin and G. Winskel. Petri nets, event structures and domains, part 1. *Theoretical Computer Science* 13:85–108, 1981.