A neutral approach to proof and refutation in MALL

Olivier Delande and Dale Miller
INRIA Saclay - Île-de France and LIX/École Polytechnique
Route de Saclay, 91128 PALAISEAU Cedex FRANCE
delande at lix.polytechnique.fr, dale.miller at inria.fr

Abstract

We propose a setting in which the search for a proof of B or a refutation of B (a proof of $\neg B$) can be carried out simultaneously: this is in contrast to the usual approach in automated deduction where we need to commit to proving either B or $\neg B$. Our neutral approach to proof and refutation is described as a two player game in which each player follows the same rules. A winning strategy translates to a proof of the formula and a winning counter-strategy translates to a refutation of the formula. The game is described for multiplicative and additive linear logic without atomic formulas. A game theoretic treatment of the multiplicative connectives is intricate and our approach to it involves two important ingredients. First, labeled graph structures are used to represent positions in a game and, second, the game playing must deal with the failure of a given player and with an appropriate resumption of play. This latter ingredient accounts for the fact that neither player might win (that is, neither B nor $\neg B$ might be provable).

1. Introduction

Consider the behavior of an idealized Prolog interpreter given a *noetherian* logic program Δ and query G. We can expect that an attempt to prove G ends with either a *finite success* or a *finite failure*. In the first case, we have a proof of G (from Δ) and in the second case we have a proof of $\neg G$ from (the completion of) Δ [9, 7]. Attempting to capture this simple observation appears difficult in the usual presentation of *proof search* since in that setting, we must first establish what we plan to prove, namely either G or $\neg G$, and then set about to prove that selection. A failure to build a sequent calculus proof of G, for example, may leave little information that helps to build a sequent calculus proof of $\neg G$. Our (idealized) Prolog interpreter, however, does *one* computation from which one constructs a proof of either G or $\neg G$. This example suggests that there might be a *neutral approach* to understanding proof search, at least in certain weak subsets of logic. Ideally, we would only like to organize one computation from which we can extract a proof or a refutation, depending on how the computation terminates.

In this paper, we describe a *neutral approach to proof and refutation* for MALL (multiplicative and additive linear logic) using games in which positions are *neutral graphs* that are composed of *neutral expressions*. This extends previous work by Miller and Saurin in [15], in which the multiplicative part of the logic was strongly restricted to avoid interactions between multiplicative connectives of opposite polarities. These graphs and expressions have two dual translations into logic and they can be seen as describing the frontiers of two derivations (a proof and a refutation) that are being extended simultaneously. In this setting, winning strategies yield proofs: depending on which player has a winning strategy, either the positive or the negative translation into logical formulas has a proof.

Notice that our use of games here is different from the use of games in, say, [1, 11], where games are used to capture the dynamics of cut-elimination for proofs in MALL. Here, instead, games are used to model the construction of *cut-free* proofs: the cut-rule and cut-elimination result are used only to state invariants about how the neutral search for proofs unfolds (see Section 2).

Our choice of MALL is made, in part, because we wish to focus on the essential nature of multiplicative connectives (additive connectives are easy in this setting). An important aspect of MALL without atoms is characterized by the kind of inference rules that are needed to describe proofs. It is common to divide the inference rules of sequent calculus into three groups: the *structural* rules (*e.g.*, weakening and contraction), the *identity* rules (initial and cut), and the *introduction* rules. In MALL without atoms, there is no need for structural rules nor identity rules (since cut and initial can be eliminated): thus, proofs are described entirely using introduction rules and it is these rules that correspond to the moves that propel our game. While MALL without atoms may seem rather weak, the complexity of establishing theoremhood for it is PSPACE-complete [12, 13].

The contributions of this paper are the following.

(1) We present a *neutral approach to proof and refutation* in MALL by presenting a new game that can be seen as an attempt to simultaneously prove and refute a formula. The proof of our main result, Theorem 5.9, translates directly between winning strategies and (focused) proofs.

(2) MALL is not complete in the sense that there are formulas B for which neither B nor $\neg B$ are provable: consider $\bot \otimes \bot$ and $1 \parr 1$. Thus we need to consider games in which play *resumes* after one player loses so we can determine whether or not the game is a win for the other player or a loss for both.

(3) This neutral setting provides an answer to why it is that invertibility/non-invertibility (asynchrony/synchrony) are de Morgan duals of each other: these two qualities are two sides of the same *process*. In our game, both players follow identical rules of play. Invertibility (asynchrony) occurs when a player needs to consider all possible moves of the opponent: one is forced to consider all moves and no choices are considered: the set of all possible moves is part of the definition of the game arena. Non-invertibility (synchrony) occurs when the opponent picks her responding move: here, genuine information is injected into the game and this is expressed in proofs as a path though non-invertible inference rules in a (focused) proof.

Due to space constraints, appendices containing details and the longest proofs are missing from this paper. They are available in an extended report ([5]).

2. Neutral expressions

In this neutral setting, formulas are replaced by neutral expressions, already introduced in [15]. We restrict ourselves to the propositional case here, but we address the complex case of interactions between multiplicatives. There is a neutral connective for each pair of dual connectives or units of the logic. We also define two translations, *i.e.*, two functions mapping each neutral expression to the two dual formulas it represents. Since two dual connectives may appear in a single formula, we need a way to switch to the other translation when translating a neutral expression. We will use the special unary operator $\updownarrow$ to this end.

Definition 2.1. Neutral expressions E *and* guarded neutral expressions G *are defined by the following grammar.*

$$G ::= \mathbb{0} \mid \mathbb{1} \mid E + E \mid E \times E \qquad E ::= G \mid \updownarrow G$$

A guarded neutral expression is therefore a neutral expression which does not begin with $\updownarrow$. The set of the neutral expressions is denoted by $\mathcal{E}$. Notice that $\updownarrow(\updownarrow E)$ is not a subexpression of a neutral expression.

Definition 2.2. *The positive and negative translations of neutral expressions into MALL formulas are defined in Figure 1. Notice that if E is a neutral expression, then $[E]^+$ and $[E]^-$ are de Morgan duals of each other. If E is guarded, then $[E]^+$ is synchronous and $[E]^-$ is asynchronous.*

$$\begin{array}{cc}
[\mathbb{0}]^+ = 0 & [\mathbb{0}]^- = \top \\
[\mathbb{1}]^+ = 1 & [\mathbb{1}]^- = \bot \\
[E+F]^+ = [E]^+ \oplus [F]^+ & [E+F]^- = [E]^- \,\&\, [F]^- \\
[E \times F]^+ = [E]^+ \otimes [F]^+ & [E \times F]^- = [E]^- \parr [F]^- \\
[\updownarrow E]^+ = [E]^- & [\updownarrow E]^- = [E]^+
\end{array}$$

Figure 1. Translations of neutral expressions

3. The additive case

Hintikka (see, for example, [10]) defined a simple game to determine the truth of a formula as follows (the game can also work for quantificational formulas). Two players, P and O, play with a single formula. The player P tries to falsify the formula while O tries to validate the formula. If the formula is a conjunction (&), P must move by choosing one of the conjuncts: in particular, if the formula is the empty conjunction ($\top$), then P can pick nothing and she loses. If the formula is a disjunction ($\oplus$), O must move by choosing one of the disjuncts: in particular, if the formula is the empty disjunction (0), then O can pick nothing and she loses. This game is *determinate* in the sense that one player always has a winning strategy. If P has a winning strategy starting with B then B is false: conversely if O has a winning strategy starting with B then B is true.

This same game can be used to provide a neutral approach to proof and refutation for the additive fragment of linear logic based on just $0, \oplus, \top, \&$. This has been done in [15] and we describe it here as an introduction to the neutral approach. The formulas we consider are exactly those of the form $[E]^-$ or $[E]^+$ where E ranges over the *additive* neutral expressions, namely

$$G ::= \mathbb{0} \mid E + E \qquad E ::= G \mid \updownarrow G$$

Let us define a similar game based on neutral expressions. (Complete details for this example are given in the extended report.) First we define a rewriting relation on neutral expressions:

$$E_1 + E_2 \to E_1 \qquad E_1 + E_2 \to E_2$$

Expressions of the form $\mathbb{0}$ and $\updownarrow E$ do not rewrite. Our game is composed of *positions* that are guarded neutral expressions and a *move* in this game from E to F, denoted by $E \;\rho\; F$, takes place exactly when $E \to^* \updownarrow F$. The first relationship on neutral expressions $\to$ denotes "micro-moves" or "internal moves" while the second relationship ρ denotes "macro-moves" or actual steps in the game. The following theorem is proved in the extended report.

Theorem 3.1. *Let E be a guarded neutral expression. There exists a winning strategy from E iff $\vdash [E]^+$ is provable. There exists a winning counter-strategy from E iff $\vdash [E]^-$ is provable. In either case, the winning strategy or counter-winning strategy provide the corresponding proof.*

This game gives us a neutral approach to proof and refutation as follows: Let B an additive linear logic formula and let E a guarded neutral expression such that $[E]^+$ is B. The move tree from position E is completely neutral and symmetric with respect to the two players: it has a winning strategy if and only if B is provable and has a winning counter-strategy if and only if $\neg B$ is provable.

The rest of this paper addresses the much more complex situation that occurs when we admit the multiplicatives. A step towards accounting for multiplicatives (together with additives) was reported in [15] where *simple neutral expressions* were considered: such expressions allowed for some multiplicative connectives as long as they essentially disappeared during the internal (micro-step) phase. Capturing full MALL in this setting is more involved and is addressed next.

4. Accounting for multiplicatives

Unlike the case for games over purely additive connectives, a game for multiplicatives cannot be determinate: for example, the neutral expression $\updownarrow \mathbb{1} \times \updownarrow \mathbb{1}$ yields the two formulas, $\bot \otimes \bot$ and its negation $1 \parr 1$, neither of which are provable. Thus in the following description of games, we need the possibility that there is a tie in play.

4.1. Two player games with ties

Name two players 0 and 1. For $\sigma \in \{0, 1\}$, we denote by $\overline{\sigma}$ the number $1 - \sigma$. An *arena* is a graph $(\mathcal{P}, \rho)$ where $\mathcal{P}$ is a set of *positions* and ρ is a binary noetherian relation on $\mathcal{P}$ that encode the possible game *moves*: if $p \ \rho \ p'$ then p' is a ρ-successor to p. A position with no ρ-successor is called *final*. All final positions are classified as 0-wins, 1-wins, and ties, and the non-final positions as 0-positions and 1-positions. If p is a position, a *play from* p is a path in the arena starting with p. A play is finite since ρ is noetherian. A play is *won* by player σ iff its last position is a σ-win, and is a *tie* iff its last position is a tie.

Informally, we choose a starting position p and put a token on it. A play from p is a finite sequence of moves of the token starting in p. If the current position of the token is final, then the play ends and we conclude that either player 0 wins the play, player 1 wins the play, or nobody wins the play. If it is a 0-position (resp. 1-position), then player 0 (resp. 1) chooses a ρ-successor of p, moves the token there, and the play continues.

A *σ-strategy for p* is a prefixed closed set $\mathcal{S}$ of plays from p containing (p) and is such that for every $(p_0, \ldots, p_n) \in \mathcal{S}$

- if p_n is a σ-position there exists p_{n+1} such that $p_n \ \rho \ p_{n+1}$ and $(p_0, \ldots, p_{n+1}) \in \mathcal{S}$,
- if p_n is a $\overline{\sigma}$-position then for every p_{n+1} such that $p_n \ \rho \ p_{n+1}$, $(p_0, \ldots, p_{n+1}) \in \mathcal{S}$.

A *winning σ-strategy for p* is a *σ-strategy for p* such that every play in it that ends in a final position is won by σ.

4.2. Focalization

In our neutral approach, we develop two dual derivations simultaneously. When we apply a rule in one, we apply its dual in the other one. However, sequent calculus for MALL lacks the symmetry we need. Consider the following dual derivations:

$$\dfrac{\dfrac{\vdash A \quad \vdash B}{\vdash A \otimes B} \quad \dfrac{\vdash C \quad \vdash D}{\vdash C \otimes D}}{\vdash (A \otimes B) \otimes (C \otimes D)}$$

$$\dfrac{\dfrac{\dfrac{\vdash A^\perp, B^\perp, C^\perp, D^\perp}{\vdash A^\perp, B^\perp, C^\perp \parr D^\perp)}}{\vdash A^\perp \parr B^\perp, C^\perp \parr D^\perp)}}{\vdash (A^\perp \parr B^\perp) \parr (C^\perp \parr D^\perp)}$$

In this example, the first derivation should be seen as a strategy for the player, and the second one as a strategy for the opponent. In the first derivation, $A \otimes B$ and $C \otimes D$ are decomposed in distinct branches, hence the syntax makes it explicit that the order in which these decompositions occur (i.e. in which the player chooses to make them) is irrelevant. In contrast, in the second derivation $A^\perp \parr B^\perp$ is clearly decomposed before $C^\perp \parr D^\perp$. However, those two rules could be trivially permuted and the opponent's strategy should reflect this fact. A *focused* proof system seems more appropriate since it considers derivations modulo such permutations.

Moreover, in the additive game previously, a move from $E \ \rho \ F$ exists iff $E \to^* \updownarrow F$. While each rewriting step $E \to E'$ (which we will refer to as "micro-move") corresponds to the application of an individual rule in the proof system, a move $E \ \rho \ F$ (which we will refer to as "macro-move") corresponds to a maximal sequence of micro-moves. This differs from Hintikka's games, in which each game move corresponds to the application of an individual rule. In Hintikka's setting, the main connective of a formula determines which player makes the next move. In our neutral setting, neutral expressions are decomposed by the same player until $\updownarrow$ is reached.

$$\frac{\vdash \Delta \Uparrow \Gamma}{\vdash \Delta \Uparrow \bot, \Gamma}\,[\bot] \qquad \frac{\vdash \Delta \Uparrow F, G, \Gamma}{\vdash \Delta \Uparrow F \parr G, \Gamma}\,[\parr]$$

$$\frac{}{\vdash \Delta \Uparrow \top, \Gamma}\,[\top] \qquad \frac{\vdash \Delta \Uparrow F, \Gamma \quad \vdash \Delta \Uparrow G, \Gamma}{\vdash \Delta \Uparrow F \,\&\, G, \Gamma}\,[\&]$$

$$\frac{}{\vdash \Downarrow \mathbf{1}}\,[\mathbf{1}] \qquad \frac{\vdash \Delta_1 \Downarrow F, \Gamma_1 \quad \vdash \Delta_2 \Downarrow G, \Gamma_2}{\vdash \Delta_1, \Delta_2 \Downarrow F \otimes G, \Gamma_1, \Gamma_2}\,[\otimes]$$

$$\frac{\vdash \Delta \Downarrow F_1, \Gamma}{\vdash \Delta \Downarrow F_1 \oplus F_2, \Gamma}\,[\oplus_1] \qquad \frac{\vdash \Delta \Downarrow F_2, \Gamma}{\vdash \Delta \Downarrow F_1 \oplus F_2, \Gamma}\,[\oplus_2]$$

$$\frac{\vdash \Delta, F \Uparrow \Gamma}{\vdash \Delta \Uparrow F, \Gamma}\,[R \Uparrow] \qquad \frac{\vdash \Delta \Uparrow \Gamma}{\vdash \Delta \Downarrow \Gamma}\,[R \Downarrow] \qquad \frac{\vdash \Delta \Downarrow \Gamma}{\vdash \Delta, \Gamma \Uparrow}\,[D]$$

Figure 2. The focused proof system $\mathcal{F}$. In the release rules, F is synchronous in $[R \Uparrow]$ and Γ contains only asynchronous formulas in $[R \Downarrow]$. The decide rule $[D]$ requires Γ to be non-empty.

$$\frac{\dfrac{\dfrac{\dfrac{\dfrac{\vdash A, B \Uparrow}{\vdash \Uparrow A, B}\,[R \Uparrow]}{\vdash \Uparrow A \parr B}\,[\parr]}{\vdash \Downarrow A \parr B}\,[R \Downarrow] \quad \dfrac{\vdash \Uparrow C^\perp}{\vdash \Downarrow C^\perp}\,[R \Downarrow]}{\vdash \Downarrow (A \parr B) \otimes C^\perp}\,[\otimes]}{\vdash (A \parr B) \otimes C^\perp \Uparrow}\,[D] \tag{1}$$

$$\frac{\dfrac{\dfrac{\dfrac{\dfrac{\vdash C \Uparrow A^\perp}{\vdash C \Downarrow A^\perp}\,[R \Downarrow] \quad \dfrac{\vdash \Uparrow B^\perp}{\vdash \Downarrow B^\perp}\,[R \Downarrow]}{\vdash C \Downarrow A^\perp \otimes B^\perp}\,[\otimes]}{\vdash A^\perp \otimes B^\perp, C \Uparrow}\,[D]}{\vdash \Uparrow A^\perp \otimes B^\perp, C}\,[R \Uparrow]}{\vdash \Uparrow (A^\perp \otimes B^\perp) \parr C}\,[\parr] \tag{2}$$

Figure 3. Two dual derivations. Here, A, B and C are synchronous formulas.

This division of inference rules into micro- and macro-moves strongly corresponds to what one sees in focused proof systems. Figure 2 contains a focused proof system for MALL, called $\mathcal{F}$, that contains two different *phases* marked by different sequents. The *asynchronous phase* is marked by sequents of the form $\vdash \Delta \Uparrow \Gamma$, while the *synchronous phase* is marked by sequents of the form $\vdash \Delta \Downarrow \Gamma$. In both cases, Δ and Γ are multisets of formulas (the multiset union of these two multisets is written as Γ, Δ). This focusing proof system is a simple variant of Andreoli's Σ_3 proof system in [2]. The main differences are that our system is restricted to MALL without atoms and that the decide rule $[D]$ can decide on more than one formula (*i.e.*, D is not restricted to be a singleton). This extension was introduced in [16]. The soundness and completeness of our system here is a trivial consequence of corresponding result for Σ_3 in [2].

As we shall see, micro-moves correspond to the application of the individual rules in Figure 2 while macro-moves correspond to an entire synchronous or asynchronous phase. Notice that introduction rules in the asynchronous phase are invertible and, as such, proof search in this phase requires no choices. On the other hand, the introduction rules in the synchronous phase are not generally invertible and, as such, proof search in this phase requires choices to be made, *e.g.*, which disjunction to select or how to split the side formulas of a tensor.

In our neutral game setting, every move has two dual readings: one is asynchronous and one is synchronous. When a move is considered from the point-of-view of the player making the move, the interpretation is synchronous: it is in this phase that a player must make choices in how the game should unfold. When a move is considered from the opponent's point-of-view, the interpretation is asynchronous: in this phase a player has no choices since she must accommodate all possible moves of the opponent. This correspondence to games provides an explanation of why the asynchronous connectives of MALL are de Morgan duals of the synchronous connectives.

In Andreoli's Σ_3, *one* formula is selected and decomposed in a synchronous phase, while *all* formulas are decomposed in an asynchronous phase. This asymmetry does not fit well in our neutral setting, which forces formulas to be decomposed simultaneously in two dual derivations. We recover some of the symmetry by allowing several foci to be selected: *some* formulas (read synchronously by the player and asynchronously by the opponent) are decomposed at each move.

4.3. Neutral graphs

In order to account for the complexity and intensional behavior of the multiplicative connectives of MALL, we shall not enrich the structure of arenas and plays (for example, we do not attempt concurrent player games, etc). Instead, we enrich the notion of position by moving from being just simple neutral expressions (as was used in the additive games of Section 3) to labeled graph structures, which we describe next.

Figure 3 shows an example of two dual derivations. It should be noted that at any point in the simultaneous development of those derivations, there are strong relationships between their frontiers. Each formula present in a frontier has its dual in the other frontier. Moreover this is a one-to-one correspondence. For example at the bottom of the derivations the frontier of (1) consists of the sequent $\vdash (A \parr B) \otimes C^\perp \Uparrow$ and the frontier of (2) consists of $\vdash \Uparrow (A^\perp \otimes B^\perp) \parr C$. Clearly there is exactly one formula in each frontier and they are dual. At the top of the two derivations, the frontiers are $\vdash A, B \Uparrow$ and $\vdash\Uparrow C^\perp$ for (1), and $\vdash C \Uparrow A^\perp$ and $\vdash\Uparrow B^\perp$ for (2). Here, the corresponding pairs are $A/A^\perp$, $B/B^\perp$, and $C^\perp/C$.

An even stronger statement can be made about these sequents on the frontier. If we admit the focused cut-rule

$$\frac{\vdash \Delta_1, B \Uparrow \qquad \vdash \Delta_2 \Uparrow B^\perp, \Gamma}{\vdash \Delta_1, \Delta_2 \Uparrow \Gamma}\ \text{cut}$$

then these frontier sequents can be combined to derive (using just this focused cut rule) the empty sequent. In the above example, this cut-derivation would be

$$\cfrac{\cfrac{\cfrac{\vdash A, B \Uparrow \qquad \vdash\Uparrow B^\perp}{\vdash A \Uparrow}\ \text{cut} \qquad \vdash C \Uparrow A^\perp}{\vdash C \Uparrow}\ \text{cut} \qquad \vdash\Uparrow C^\perp}{\vdash\Uparrow}\ \text{cut.}$$

While we do not make explicit use of such cut-derivations of the empty sequent in the sequel, the existence of such derivations provide a useful invariant concerning the evolution of game playing. For example, it immediately follows that at most one player will succeed to win and at least one player must lose.

In [4], the authors analyse the geometry of generalized multiplicative rules and express the duality of two generalized multiplicatives through a graph structure. Following this idea, we define a graph structure to represent links between two frontiers as presented above. In this kind of graph, called *neutral graph*, the vertices represent the sequents of the frontiers. There are two colors of vertices (one for each frontier). As we have seen, there is a one-to-one correspondence between formulas of the two frontiers. We represent each pair of corresponding formulas by an arc between the two vertices representing the sequents in which they appear. The arc

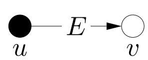

labeled with a guarded neutral expression E means that the formula $[E]^+$ occurs in the sequent represented by u and that the formula $[E]^-$ occurs in the sequent represented by v. A neutral graph is bipartite: recall that we have a color for each frontier and that we do not pair two formulas in the same frontier. For example, two frontiers

$$\vdash [E]^+, [F]^+ \Uparrow \quad \vdash\Uparrow [G]^- \ \mid\ \vdash [G]^+ \Uparrow [E]^- \quad \vdash\Uparrow [F]^-$$

will be represented by the neutral graph

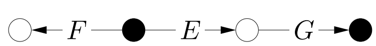

where the black (resp. white) vertices represent the sequents of the left (resp. right) frontier.

We are going to define transition relations on neutral graphs that will correspond to the simultaneous development of the derivations. In the neutral graphs introduced so far, the sequent associated with a vertex v is $\vdash [O_1]^+, \ldots, [O_m]^+ \Uparrow [I_1]^-, \ldots, [I_n]^-$, where $O_1, \ldots, O_m$ (resp. $I_1, \ldots, I_n$) are the neutral expressions labeling the outgoing (resp. incoming) arcs of v. Since all those neutral expressions are guarded, the $[O_i]^+$ are synchronous and the $[I_i]^+$ are asynchronous. Such a neutral graph is called *passive*, because it represents frontiers where no formula is under focus. Consider developing one of the derivations by applying the $[D]$ rule. Some formulas are put under focus. Subsequent development of this derivation will consist in decomposing those formulas. In our neutral setting, this will be matched by a decomposition of their (asynchronous) duals in the other derivation. Just as the $[D]$ rule marks formulas for decomposition at the beginning of a synchronous phase, we need a way to mark arcs of neutral graphs for decomposition of the neutral expressions labeling them. A neutral graph will record which arcs are marked or *focused*. The focused arc

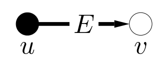

(notice the thicker line) indicates that $[E]^+$ appears under focus in the sequent associated with u. As a formula under focus may be asynchronous, we allow neutral expressions labeling focused arcs not to be guarded. A neutral graph with focused arcs is called *active* because one of the frontiers it represents is in the middle of a synchronous phase.

Definition 4.1. *A* neutral graph *G is a tuple (V, A, p, ϵ, F), where V is a finite set (possibly empty) of* vertices*, $A \subseteq V \times V$ is a set of* arcs*, $p : V \mapsto \{0, 1\}$ associates a* polarity *to each vertex, $\epsilon : A \mapsto \mathcal{E}$ maps each arc to a neutral expression, and $F \subseteq A$. In addition, the following must hold:*

- *The undirected graph based on (V, A) is a set of trees none of which are the degenerate (one-vertex) tree.*
- *For every $a \in A$, if $\epsilon(a)$ is not guarded then $a \in F$.*
- *For every $(u, v) \in A$, $p(u) \neq p(v)$.*

Notice that the definition requires that no vertex be isolated (i.e. without neighbours).

Informally, a move of player σ (for $\sigma \in \{0, 1\}$) corresponds to a synchronous (resp. asynchronous) phase from the frontier consisting of the vertices of polarity σ (resp. $\overline{\sigma}$).

We say that $a \in A$ is *focused* if $a \in F$. A vertex v with polarity σ (*i.e.*, $p(v) = \sigma$) is called a σ-vertex. A neutral graph G is *connected* iff (V, A) is weakly connected, and *disconnected* otherwise. Notice that the polarity assignment p and the restriction above makes the graph (V, A) bipartite.

Definition 4.2. *The neutral graph $(V', A', p', \epsilon', F')$ is a* subgraph *of (V, A, p, ϵ, F) if $V' \subseteq V$, $A' \subseteq A$, $p' = p_{|V'}$,*

$\epsilon' = \epsilon_{|A'}$, *and* $F' = F \cap A'$. *If* $G = (V, A, p, \epsilon, F)$ *is a neutral graph and* $A' \subseteq A$*, we denote by* $G_{|A'}$ *the* maximal subgraph *of* G *whose arcs all belong to* A'. *Formally*

$$G_{|A'} = (V', A', p_{|V'}, \epsilon_{|A'}, F \cap A')$$

where $V' = \{v \in V : \exists u \in V\ (u, v) \in A' \vee (v, u) \in A'\}$.

A *connected component* of a neutral graph is a maximal connected subgraph.

Recall the "cut invariant" presented above. Cutting a formula roughly corresponds to merging the two vertices its corresponding arc connects in a neutral graph. The invariant thus suggests that by repeating this operation the whole neutral graph reduces to a single vertex. This should lead us to require (V, A) to be a tree (see more about this in [4]). However, we allow (V, A) to be disconnected, as long as each connected component is a tree, since the failure of one player may leave the graph disconnected: in that case, the next move typically consists in the player who failed selecting a connected component on which to continue the play.

Definition 4.3. *A neutral graph* G *is* active *when it has at least one focused arc and is* passive *otherwise. A neutral graph is* degenerate *if it has no vertices. There is exactly one such neutral graph, denoted by* δ*, and it is clearly passive.*

Active neutral graphs correspond to states in which a synchronous phase is not finished yet. As such, they will appear between micro-moves, but not between macro-moves. We need to impose restrictions on neutral graphs before they can be introduced in a game. Since our game is not concurrent, we must be able to clearly state whose turn it is. We will assign a *polarity* to suitable neutral graphs. Informally, if G has polarity $\sigma \in \{0, 1\}$, then it is player σ's turn to play in G.

A *source* of a neutral graph is a vertex v such that there is no arc of the form (u, v).

Definition 4.4. *A passive neutral graph* G *is* weakly polarized *if for every connected component* C *of* G *there is a polarity* $\sigma \in \{0, 1\}$ *such that every source of* C *has polarity* σ*. If* G *is weakly polarized and has exactly one connected component, then* G *is* strongly polarized at polarity σ*, where* σ *is the polarity assigned to the sources of* G*.*

Notice that δ is weakly but not strongly polarized. In a passive neutral graph G, all associated sequents are of the form $\vdash \Gamma \Uparrow \Delta$, that is, not in a synchronous phase. Those from which a synchronous phase is *about* to start (via the $[D]$ rule) are those for which Δ is empty, that is, those associated with the sources of G. Requiring all the sources of a connected component to share the same polarity σ ensures that player σ must be the next one to play in that component. As follows from the above explanation, when a move results in a weakly (but not strongly) polarized passive neutral graph, the next move is to select a (strongly polarized) component (if any) to continue from.

Definition 4.5. *An active neutral graph* $G = (V, A, p, \epsilon, F)$ *is* (weakly or strongly) polarized at polarity σ *iff the following hold:*

- *for every* $(u, v) \in F$*,* $p(u) = \sigma$ *and* $p(v) = \overline{\sigma}$ *and*
- *the passive neutral graph* G'*, obtained by reversing and unfocusing* G*'s focused arcs, is weakly polarized.*

Let $\mathcal{N}_a$ (resp. $\mathcal{N}_p$) be the set of the weakly polarized active (resp. passive) neutral graphs. Let $\mathcal{N} = \mathcal{N}_a \cup \mathcal{N}_p$ denote the set of the weakly polarized neutral graphs. In the following, all the neutral graphs we consider are weakly polarized.

If $S \subset \mathcal{N}$, let S^* be the subset of S consisting of its strongly polarized elements. Clearly $\mathcal{N}_a^* = \mathcal{N}_a$.

A notion which will be useful in proving that plays are finite is that of the *size* of a neutral graph $G \in \mathcal{N}$. We define it to be the total number of symbols of the neutral expressions labeling the arcs of G, and denote it by $size(G)$.

4.4. Rewriting neutral graphs

This section describes the transitions on neutral graphs that are the basis of the game. We first introduce six of them, the aforementioned "micro-moves", that should be interpreted as the simultaneous applications of two dual single rules of the proof system. Table 1 lists them along with their interpretations. We subsequently build another transition, which packs a maximal sequence of micro-moves together and should be read as the simultaneous development of two dual phases. Failures may arise in some of these transitions; in that case the transition is labeled with two boolean flags f_0 and f_1, where f_i is $\top$ if and only if player i has encountered a failure. A player who has failed cannot win the play any more but may try to prevent her opponent from winning by making her fail as well (in which case the play ends in a tie).

In the first of the micro-moves of Table 1, G is passive and G' is active: with this transition, we are selecting what neutral expressions should be decomposed. The synchronous reading of this step corresponds to the multifocus inference while the asynchronous reading corresponds to selecting which formulas to use for invertible decomposition. The second transition is the converse: once we have reached the end of a phase (marked by the $\updownarrow$ operator), the designated formulas are released (unfocused). Here, G is active and G' loses one of its foci. The next two transitions result from dealing either with an additive or a multiplicative neutral expression labeling a focused arc. The last two transitions deal with the additive and multiplicative units: it is with these units that failures can arise in game playing.

Transition	Sync reading	Async reading
$G \xrightarrow{D} G'$	$[D]$	none
$G \xrightarrow{R} G'$	$[R \Downarrow]$	$[R \Uparrow]$
$G \xrightarrow{+} G'$	$[\oplus]$	$[\&]$
$G \xrightarrow{\times} G'$	$[\otimes]$	$[\parr]$
$G \xrightarrow{\mathbf{0},f_0,f_1} G'$	none	$[\top]$
$G \xrightarrow{\mathbb{1},f_0,f_1} G'$	$[1]$	$[\bot]$

Table 1. Neutral moves and their two readings

In the following description of the micro-moves we use figures to illustrate the formal definitions. Each micro-move rewrites a strongly polarized neutral graph G. σ- (resp. $\overline{\sigma}$-) vertices are represented in black (resp. white), where σ is the polarity of G. We also refer to player σ (resp. $\overline{\sigma}$) as the black (resp. white) player.

Decision: Let $G = (V, A, p, \epsilon, \emptyset) \in \mathcal{N}_p^*$ and let $v_1, \ldots, v_n$ be the sources of G. For each v_i, let A_i be a non empty subset of $\{(v_i, w) : (v_i, w) \in A\}$. If we then let $G' = (V, A, p, \epsilon, \cup_{i=1}^n A_i)$, we have the labeled transition $G \xrightarrow{D} G'$.

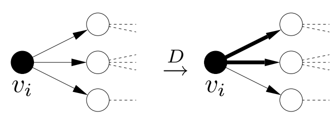

(this figure only shows one source v_i.) Let us give an informal description of this transition $G \xrightarrow{D} G'$. Recall the decision rule ($[D]$ in Figure 2). It is applied to a sequent of the form $\vdash \Gamma \Uparrow$. In G, these sequents exactly correspond to the sources, and the transition corresponds exactly to applying $[D]$ to each one of them. $G \in \mathcal{N}_p^*$ ensures that all sources are σ-vertices, where σ is the polarity of G; this micro-move shall therefore be made by player σ.

To describe the next five labeled transitions, let $G = (V, A, p, \epsilon, F) \in \mathcal{N}_a$ and $a = (v, w) \in F$.

Reaction: If $\epsilon(a)$ is of the form $\updownarrow E$, then one can remove the leading $\updownarrow$, reverse the arc, and unfocus it. Formally, let

$$G' = (V, (A \setminus \{a\}) \cup \{\overline{a}\}, p, \epsilon_{|A\setminus\{a\}} \cup \{(\overline{a}, E)\}, F \setminus \{a\})$$

where $\overline{a} = (w, v)$ is the opposite arc to a. Then we have the transition $G \xrightarrow{R} G'$.

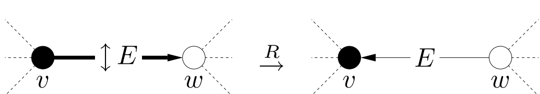

In both interpretations, a formula of the wrong polarity is reclassified.

Additives: If $\epsilon(a)$ is of the form $E_1 + E_2$, then one can replace this expression with one of the operands. Formally, let $G' = (V, A, p, \epsilon', F)$ where ϵ' is the same as ϵ except that $\epsilon'(a) = E_i$ for some $i \in \{1, 2\}$. We then have the labeled transition $G \xrightarrow{+} G'$.

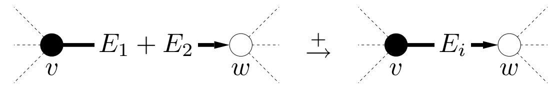

This treatment of $+$ is essentially the same as in the additive game presented before.

If $\epsilon(a) = \mathbb{0}$ (the 0-ary additive), then one can remove w and all its adjacent arcs. Formally, let $G' = G_{|A\cap(V\setminus\{w\})^2}$ and let f_0 and f_1 be the boolean values defined as follows: $f_{p(v)} = \top$ and $f_{p(w)} = \bot$. Then we have the labeled transition $G \xrightarrow{\mathbf{0},f_0,f_1} G'$.

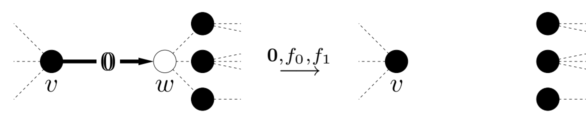

(in the second graph, any isolated vertex shall be removed.) This last transition is particular: on the white player's side we simply remove a sequent of the form $\vdash \Gamma \Uparrow \top, \Delta$, in other words we apply $[\top]$; on the black player's side we are confronted with an unprovable sequent of the form $\vdash \Gamma \Downarrow 0, \Delta$. Consequently the black player fails ($f_{p(v)} = \top$).

Multiplicatives: If $\epsilon(a)$ is of the form $E_1 \times E_2$, then one can split v into two vertices and a into two arcs, labeling each one with an operand. Formally, define two new vertices v_1 and v_2 and for every $b = (t, u) \in A \setminus \{a\}$, define an arc b' as follows: if $t \neq v$ and $u \neq v$, then $b' = b$; if $t = v$, then $b' = (v_i, u)$ for some $i \in \{1, 2\}$; and if $u = v$, then $b' = (t, v_i)$ for some $i \in \{1, 2\}$. Now let $G' = (V', A', p', \epsilon', F')$ where

- $V' = (V \setminus \{v\}) \uplus \{v_1, v_2\}$,
- $A' = \{(v_1, w), (v_2, w)\} \cup \{b' : b \in A \setminus \{a\}\}$,
- $p' = p_{|V\setminus\{v\}} \cup \{(v_1, p(v)), (v_2, p(v))\}$,
- $\epsilon'(v_1, w) = E_1$ and $\epsilon'(v_2, w) = E_2$, and for every $b \in A \setminus \{a\}$, $\epsilon'(b') = \epsilon(b)$,
- $F' = \{(v_1, w), (v_2, w)\} \cup \{b' : b \in F \setminus \{a\}\}$.

We then have the labeled transition $G \xrightarrow{\times} G'$.

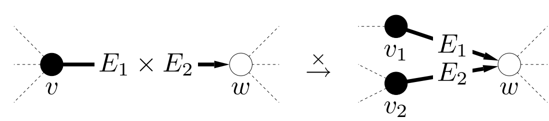

On the black player's side, the splitting corresponds to that of the $[\otimes]$ rule. On the white player's side the invertible $[\parr]$ rule is applied.

If $\epsilon(a) = \mathbb{1}$ (the 0-ary multiplicative), then one can remove a. Formally, let $G' = G_{|A\setminus\{a\}}$ where f_0 and f_1 are boolean values defined as follows: $f_{p(v)} = \top$ iff v is a vertex of G', $f_{p(w)} = \bot$ iff w is a vertex of G'. Then we have the labeled transition $G \xrightarrow{\mathbb{1},f_0,f_1} G' = G_{|A\setminus\{a\}}$.

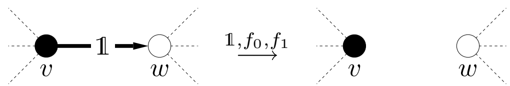

(in the second graph, any isolated vertex shall be removed.) In this transition both players may fail. On the black player's side the transition corresponds to applying [1]. The sequent associated to v should thus be $\vdash\Downarrow 1$, therefore the player fails ($f_{p(v)} = \top$) if 1 is not the only formula of the sequent. On the white player's side $[\bot]$ is applied, and if w is only connected to v then its associated sequent becomes $\vdash\Uparrow$ which is unprovable, and the player fails ($f_{p(w)} = \top$).

Proposition 4.6. *$\xrightarrow{D}$ has the following properties:*

1. *if $G \xrightarrow{D} G'$, then $G' \in \mathcal{N}_a$ and has the same polarity as G;*
2. *if $G \in \mathcal{N}_p^*$, then there is $G' \in \mathcal{N}_a$ such that $G \xrightarrow{D} G'$.*

Proof. Let us show the first property. Since there is at least one source in G, G' has at least one focused arc and is active. Let σ be the polarity of G. All the sources of G are σ-vertices. In G' all focused arcs are of the form (v, w) where v is a source and is therefore a σ-vertex. Let G'' be the passive neutral graph obtained by reversing and unfocusing the focused arcs of G'. Since every source of G' is the origin of at least one focused arc, all the sources of G'' are ends of focused arcs of G' and are therefore $\overline{\sigma}$-vertices. Hence G'' is weakly polarized and G' has polarity σ.

The second property is immediate: since there are no isolated vertices in a neutral graph, every source of G has at least one outgoing arc to focus on. □

We are going to define how to build sequences of micro-moves, which correspond to phases in focused proof search; those sequences are therefore built as follows: a) $\xrightarrow{D}$ is used once to initiate the sequence (thus making the neutral graph active), b) all the other micro-moves are applied until the neutral graph becomes passive again. The following definition introduces a transition representing a generic micro-move occurring in part b).

Definition 4.7. *The relation $G \xrightarrow{f_0,f_1} G'$, where $G \in \mathcal{N}_a$, $G' \in \mathcal{N}$, and f_0, f_1 are boolean values, is defined to hold in the following cases: for every $G \in \mathcal{N}_a$, if $G \xrightarrow{R} G'$, $G \xrightarrow{+} G'$ or $G \xrightarrow{\times} G'$, then $G \xrightarrow{\bot,\bot} G'$; and if $G \xrightarrow{\mathbb{1},f_0,f_1} G'$ or $G \xrightarrow{\mathbf{0},f_0,f_1} G'$, then $G \xrightarrow{f_0,f_1} G'$.*

Proposition 4.8. *The relation $\xrightarrow{-,-}$ is finitely branching. It also satisfies the following properties.*

1. *if $G \xrightarrow{f_0,f_1} G'$, then*
   - *(a) $size(G) > size(G')$;*
   - *(b) if G' is degenerate or disconnected, then G is disconnected or $f_0 \vee f_1 = \top$;*
   - *(c) if G' is active, then it has the same polarity as G;*
2. *if $G \in \mathcal{N}_a$, then there are G', f_0 and f_1 such that $G \xrightarrow{f_0,f_1} G'$;*

Proof. The relation $\xrightarrow{-,-}$ is finitely branching because each of $\xrightarrow{R}$, $\xrightarrow{+}$, $\xrightarrow{\times}$, $\xrightarrow{\mathbf{0},-,-}$, and $\xrightarrow{\mathbb{1},-,-}$ is, as can be seen from their definitions.

Let us show property (1). (a) can be easily seen in each one of the 5 cases. Let us show (b). Assume that G is not disconnected and that G' is degenerate or disconnected. The only possible cases are $G \xrightarrow{\mathbf{0},f_0,f_1} G'$ or $G \xrightarrow{\mathbb{1},f_0,f_1} G'$. In the case $G \xrightarrow{\mathbf{0},f_0,f_1} G'$ we have $f_0 \vee f_1 = \top$. In the case $G \xrightarrow{\mathbb{1},f_0,f_1} G'$, the only situation in which $f_0 \vee f_1 = \bot$ is when, in G, v is only connected to w while w is connected to at least another vertex. This case cannot occur since it would make G' neither degenerate nor disconnected. (c) follows from the observation that each relation preserves the polarities of the origins and ends of the focused arcs. Property (2) is easily observed. If $G \in \mathcal{N}_a$, then it has a focused arc. Depending on the form of the neutral expression labeling it, one of the 5 cases applies. □

We may now pack part b) in a single transition:

Definition 4.9. *The relation $G \xrightarrow{f_0,f_1}_\dagger G'$, where $G \in \mathcal{N}$, $G' \in \mathcal{N}_p$, and f_0, f_1 are boolean values, is the smallest relation such that for every $G \in \mathcal{N}_p$, $G \xrightarrow{\bot,\bot}_\dagger G$ and such that for every $G \in \mathcal{N}_a$, if $G \xrightarrow{f_0,f_1} G' \xrightarrow{f_0',f_1'}_\dagger G''$, then $G \xrightarrow{f_0\vee f_0',f_1\vee f_1'}_\dagger G''$.*

This relation is simply characterized. Let $G \in \mathcal{N}$, $G' \in \mathcal{N}_p$, and f_0, f_1 booleans. $G \xrightarrow{f_0,f_1}_\dagger G'$ iff there are $n \in \mathbb{N}$, $G_0 = G, G_1, \dots, G_n = G' \in \mathcal{N}$ and booleans $f_0^{(i)}, f_1^{(i)}$ for $1 \le i \le n$ such that $G_0 \xrightarrow{f_0^{(1)},f_1^{(1)}} \dots \xrightarrow{f_0^{(n)},f_1^{(n)}} G_n$ and $\forall\sigma \in \{0,1\}, f_\sigma = \bigvee_{1\le i\le n} f_\sigma^{(i)}$. Notice that f_0 and f_1 indicate a failure at some point in the sequence.

A consequence is that $size(G) \ge size(G')$, and the inequality is strict if $n > 0$ (e.g. if $G \in \mathcal{N}_a$).

Composing $\xrightarrow{D}$ and $\xrightarrow{-,-}_\dagger$ yields a transition which represents a phase in the proof system:

Definition 4.10. *The relation $G \overset{f_0,f_1}{\rightsquigarrow} G'$, where $G \in \mathcal{N}_p^*$, $G' \in \mathcal{N}_p$, and f_0, f_1 are boolean values, is defined to hold if $G \xrightarrow{D} G_0 \xrightarrow{f_0,f_1}_\dagger G'$ for some $G_0 \in \mathcal{N}_a$.*

Proposition 4.11. *The relation $\overset{-,-}{\rightsquigarrow}$ is finitely branching. It also satisfies the following properties.*

1. *if $G \overset{f_0,f_1}{\rightsquigarrow} G'$*
   - *(a) $size(G) > size(G')$;*

(b) if G' is degenerate or disconnected, then $f_0 \vee f_1 = \top$;

2. *if $G \in \mathcal{N}_p^*$, then there are G', f_0 and f_1 such that $G \overset{f_0,f_1}{\rightsquigarrow} G'$.*

Proof. Consider $\xrightarrow{-,-}_\dagger$. The $\xrightarrow{-,-}$-sequences starting from some $G \in \mathcal{N}$ have lengths bounded by $size(G)$. Since $\xrightarrow{-,-}$ is finitely branching, so is $\xrightarrow{-,-}_\dagger$. So is $\xrightarrow{D}$ (easily seen), hence so is $\overset{-,-}{\rightsquigarrow}$.

Let us show property (1). Assume we have $G \overset{f_0,f_1}{\rightsquigarrow} G'$. We can write $G \xrightarrow{D} G_0 \xrightarrow{f_0^{(1)},f_0^{(1)}} \ldots \xrightarrow{f_0^{(n)},f_0^{(n)}} G_n = G'$, with $\forall \sigma \in \{0,1\}, f_\sigma = \bigvee_{1 \leq i \leq n} f_\sigma^{(i)}$. We have $size(G) = size(G_0) > size(G')$, which proves (a). Let us show (b). For $H \in \mathcal{N}$, consider the property $P(H) =$ "H is degenerate or disconnected". $G \in \mathcal{N}_p^*$, hence $P(G)$ is false, and then so is $P(G_0)$. Suppose that $P(G')$ is true. Let $k = \min\{1 \leq i \leq n : P(G_i)\}$. By Proposition 4.8 $f_0^{(k)} \vee f_1^{(k)} = \top$, hence $f_0 \vee f_1 = \top$. Property (2) follows from similar results in Propositions 4.6 and 4.8. □

4.5. Positions and moves

We can now define the positions and moves of the game. We must also specify, for each position, whether it is a 0-position, a 1-position, a 0-win, a 1-win or a tie.

Definition 4.12. *A* position *is a tuple of the form (G, f_0, f_1), where $G \in \mathcal{N}_p$ and f_0, f_1 are boolean values, such that if $G \notin \mathcal{N}_p^*$, then $f_0 \vee f_1 = \top$.*

In other words, a game position is some $G \in \mathcal{N}_p$ representing goals to be achieved by the players, along with two boolean values recording which player(s) has(have) failed so far. In addition, if G is degenerate or disconnected then some player must have failed.

Definition 4.13. *We define the terminal game positions and the move relation ρ simultaneously as follows.*

1. *The two positions $(\delta, \top, \bot)$ and $(\delta, \bot, \top)$ are respectively a 1-win and a 0-win;*
2. *the positions of the form $(G, \top, \top)$ are ties;*
3. *a position (G, f_0, f_1), where $G \neq \delta$, $G \notin \mathcal{N}_p^*$, $f_\sigma = \top$ and $f_{\overline{\sigma}} = \bot$, is a σ-position and its ρ-successors are (G', f_0, f_1), one for every connected component G' of G;*
4. *a position (G, f_0, f_1), where $G \in \mathcal{N}_p^*$ and $f_0 \wedge f_1 = \bot$, is a σ-position, where σ is G's polarity. If $G \overset{f_0',f_1'}{\rightsquigarrow} G'$, then $(G, f_0, f_1)\ \rho\ (G', f_0 \vee f_0', f_1 \vee f_1')$.*

Informally, case 1 says that if there are no goals left and only one player has failed, then her opponent wins; case 2 says that as soon as both players fail the play ends in a tie; case 3 says that when a player's failure causes the neutral graph to be disconnected, she should pick a connected component to challenge her opponent to; and case 4 describes a normal move, which corresponds to the simultaneous development of two phases in the proof system.

It can be easily seen (see Proposition 4.11) that every non-final position actually has a ρ-successor.

Proposition 4.14. *The relation ρ is finitely branching and if $(G, f_0, f_1)\ \rho\ (G', f_0', f_1')$, then*
1. *$size(G) > size(G')$;*
2. *if $f_\tau = \top$ for some $\tau \in \{0,1\}$, then $f_\tau' = \top$ (ρ preserves failures).*

Proof. The relation ρ is finitely branching because neutral graphs are finite (in case 3 of definition 4.13) and $\overset{-,-}{\rightsquigarrow}$ is finitely branching (in case 4). Property (1) is a consequence of the fact that every connected component of a neutral graph contains at least one arc (in case 3), and of Proposition 4.11 (in case 4). Lastly, it can be immediately seen from its definition that ρ preserves failures. □

This proposition implies that ρ is noetherian and more: the plays starting from a given position have bounded lengths.

5. Winning strategies as cut-free focused proofs

In this section we relate cut-free proofs (in the proof system) to winning strategies (in the game). Our theorems state the equivalence between provability and the existence of a winning strategy. Their proofs effectively show how to construct a winning strategy from a proof. For the converse to hold, we would need to impose a uniformity condition on strategies like innocence. We leave this as future work.

The operators $[\cdot]^+$ and $[\cdot]^-$ are applied to multisets of neutral expressions in the obvious way. Throughout this paper, we shall not admit atomic formulas (propositional variables) into formulas: formulas will contain no non-logical symbols. Two focused proofs of the same sequent are equivalent iff they differ by the order in which asynchronous rules are applied within asynchronous phases. This is indeed an equivalence relation.

We begin by formally defining two central notions relating concepts of the game to concepts of the proof system: that of sequent associated to a vertex of a neutral graph, and that of σ-provability (for $\sigma \in \{0,1\}$).

Proposition 5.1. *Let $G = (V, A, p, \epsilon, F) \in \mathcal{N}$ and $v \in V$. Consider the multisets of formulas $\mathcal{F}^- = \{[\epsilon(u,v)]^- : (u,v) \in F\}$, $\mathcal{F}^+ = \{[\epsilon(v,w)]^+ : (v,w) \in F\}$, $\mathcal{U}^- = \{[\epsilon(u,v)]^- : (u,v) \in A \setminus F\}$, and $\mathcal{U}^+ = \{[\epsilon(v,w)]^+ : (v,w) \in A \setminus F\}$.*
1. *At least one of $\mathcal{F}^-$ and $\mathcal{F}^+$ is empty.*
2. *The elements of $\mathcal{U}^-$ are asynchronous and those of $\mathcal{U}^+$ are synchronous.*

Proof. Let us show the first property. If $G \in \mathcal{N}_p$, then F is empty, and so are $\mathcal{F}^-$ and $\mathcal{F}^+$. Otherwise, $G \in \mathcal{N}_a$; suppose by contradiction that $\mathcal{F}^-$ and $\mathcal{F}^+$ are both non empty. There exist $(u, v), (v, w) \in F$. Since G is strongly polarized, $p(u) = p(v)$. Yet $(u, v) \in A$ and G is a neutral graph, hence $p(u) \neq p(v)$, which yields a contradiction.

Let us show the second property. For every $a \in A \setminus F$, $\epsilon(a)$ is guarded, hence $[\epsilon(a)]^-$ is asynchronous and $[\epsilon(a)]^+$ is synchronous. □

Definition 5.2 (Sequent associated with a vertex)**.** *Let $G = (V, A, p, \epsilon, F) \in \mathcal{N}$ and $v \in V$. We associate with v a sequent $\Sigma_{G,v}$ defined as follows:*

$$\Sigma_{G,v} = \begin{cases} \vdash \mathcal{U}^+ \Uparrow \mathcal{F}^-, \mathcal{U}^- & \text{if } \mathcal{F}^+ \text{ is empty} \\ \vdash \mathcal{U}^+ \Downarrow \mathcal{F}^+, \mathcal{U}^- & \text{otherwise} \end{cases}$$

where $\mathcal{F}^-$, $\mathcal{F}^+$, $\mathcal{U}^-$, and $\mathcal{U}^+$ are defined as above.

Definition 5.3 (σ-provability)**.** *Let $G \in \mathcal{N}$ and $\sigma \in \{0, 1\}$. G is σ-*provable *iff the sequents associated with its σ-vertices are all provable. A triple (G, f_0, f_1) where $G \in \mathcal{N}$ and f_0, f_1 are boolean values is σ-*provable *iff $f_\sigma = \bot$ and G is σ-provable.*

We relate game moves to derivations by proceeding gradually from small steps (micro-moves and inference rules) to large objects (winning strategies and proofs).

The proofs of the four following propositions are rather long and can be found in the extended report.

Proposition 5.4. *Let $G \in \mathcal{N}_a$ and let σ be G's polarity. Let $S = \{(G', f_0, f_1) : G \xrightarrow{f_0, f_1} G'\}$. G is σ-provable iff there exists $(G', f_0, f_1) \in S$ which is σ-provable.*

Proposition 5.5. *Let $G \in \mathcal{N}_a$ and let σ be G's polarity. Let $S = \{(G', f_0, f_1) : G \xrightarrow{f_0, f_1} G'\}$. G is $\overline{\sigma}$-provable iff every $(G', f_0, f_1) \in S$ is $\overline{\sigma}$-provable.*

Proposition 5.6. *Let $G \in \mathcal{N}_a$ and let σ be G's polarity. Let $S = \{(H, f_0, f_1) : G \xrightarrow{f_0, f_1}_{\dagger} H\}$. G is σ-provable iff there exists $(H, f_0, f_1) \in S$ which is σ-provable. G is $\overline{\sigma}$-provable iff every $(H, f_0, f_1) \in S$ is $\overline{\sigma}$-provable.*

Proposition 5.7. *Let $G \in \mathcal{N}_p^*$ and let σ be G's polarity. Let $S = \{(H, f_0, f_1) : G \overset{f_0, f_1}{\rightsquigarrow} H\}$. G is σ-provable iff there exists $(H, f_0, f_1) \in S$ which is σ-provable. G is $\overline{\sigma}$-provable iff every $(H, f_0, f_1) \in S$ is $\overline{\sigma}$-provable.*

Lemma 5.8. *Let $q = (G, f_0, f_1)$ be a non-final σ-position. Let $S = \{q' : q\ \rho\ q'\}$. q is σ-provable iff there exists $q' \in S$ which is σ-provable. q is $\overline{\sigma}$-provable iff every $q' \in S$ is $\overline{\sigma}$-provable.*

Proof. q is a non-final position; we may be either in case 3 or case 4 of Definition 4.13. Let us prove the result in each case separately.

Case 3 of Definition 4.13. $f_\sigma = \top$ and $f_{\overline{\sigma}} = \bot$ and for every $(G', f'_0, f'_1) \in S$, $f'_\sigma = \top$ and $f'_{\overline{\sigma}} = \bot$. It means that neither q nor any $q' \in S$ is σ-provable. Also, q is $\overline{\sigma}$-provable iff G is $\overline{\sigma}$-provable (since $f_{\overline{\sigma}} = \bot$), iff for every $(G', f'_0, f'_1) \in S$, G' is $\overline{\sigma}$-provable (since the G' are the connected components of G), iff every $q' \in S$ is $\overline{\sigma}$-provable (since the $f'_{\overline{\sigma}}$ all equal $\bot$).

Case 4 of Definition 4.13. $G \in \mathcal{N}_p^*$ and has polarity σ. We may apply Proposition 5.7 to G. Let $S' = \{(G', f'_0, f'_1) : G \overset{f'_0, f'_1}{\rightsquigarrow} G'\}$. G is σ-provable iff there exists $(G', f'_0, f'_1) \in S'$ which is σ-provable. G is $\overline{\sigma}$-provable iff every $(G', f'_0, f'_1) \in S'$ is $\overline{\sigma}$-provable. Therefore

- (G, f_0, f_1) is σ-provable iff there exists $(G', f'_0, f'_1) \in S'$ such that $(G', f_0 \vee f'_0, f_0 \vee f'_1)$ is σ-provable,
- (G, f_0, f_1) is $\overline{\sigma}$-provable iff for every $(G', f'_0, f'_1) \in S'$, $(G', f_0 \vee f'_0, f_0 \vee f'_1)$ is $\overline{\sigma}$-provable.

The result follows from the fact that $S = \{(G', f_0 \vee f'_0, f_0 \vee f'_1) : (G', f'_0, f'_1) \in S'\}$. □

Theorem 5.9. *Let q be a position and $\sigma \in \{0, 1\}$. There is a winning σ-strategy from q iff q is σ-provable.*

Proof. We know that the lengths of the plays from q are bounded. Let us prove the result by induction on the maximal length n_q of a play from q.

If $n_q = 0$, then q is a final position. Player σ has a winning strategy from q iff q is a σ-win, iff q is σ-provable (see Definition 4.13).

Suppose that $n_q > 0$. q is not a final position. Let $S = \{q' : q\ \rho\ q'\}$. There are two cases: either q is a σ-position, or it is a $\overline{\sigma}$-position.

If q is a σ-position, then there is a winning σ-strategy from q iff there is a winning σ-strategy from some $q' \in S$, iff, by induction hypothesis ($n_{q'} < n_q$), there exists $q' \in S$ which is σ-provable, iff, by Lemma 5.8, q is σ-provable.

If q is a $\overline{\sigma}$-position, then there is a winning σ-strategy from q iff there is a winning σ-strategy from every $q' \in S$, iff, by induction hypothesis ($n_{q'} < n_q$), every $q' \in S$ is σ-provable, iff, by Lemma 5.8, q is σ-provable. □

6. Related and future work

There is a great deal of work that address various game-theoretical aspects of logic. Most of the work on using game semantics with linear logic is centered around modeling cut-elimination: in particular, on viewing one player as a processing element and the other player as the environment. Blass introduced a game semantics for linear logic in [3]. From the point of view of modeling cut-free proof search, the most closely related work to that described here is Miller and Saurin's [15]: there the use of games to provide a neutral approach to proof and refutation was applied to additive games weakly extended with some multiplicative aspects: no approach to the full multiplicative setting was considered in that paper. Our work is strongly related, at least in spirit, to a part of Girard's *Ludics* [8]. Less closely related is work by Pym and Ritter [17] where game semantics is proposed as a way to control the search for proofs in intuitionistic and classical logics.

We leave several topics as future work. These include extending the logic to stronger fragments by incorporating, for example, first-order quantification, equality, and fixed points. Developing technical connections to Ludics, in particular works such as [6] would be of particular interest. It would also be of interest to relate proofs and strategies more closely, by switching to asynchronous games and innocent strategies ([14]).

References

[1] S. Abramsky and P.-A. Melliès. Concurrent games and full completeness. In *14th Symp. on Logic in Computer Science*, pages 431–442. IEEE Computer Society Press, 1999.

[2] J.-M. Andreoli. Logic programming with focusing proofs in linear logic. *J. of Logic and Computation*, 2(3):297–347, 1992.

[3] A. Blass. A game semantics for linear logic. *Annals Pure Appl. Logic*, 56:183–220, 1992. Special Volume dedicated to the memory of John Myhill.

[4] V. Danos and L. Regnier. The structure of multiplicatives. *Archive for Mathematical Logic*, 28:181–203, 1989.

[5] O. Delande and D. Miller. A neutral approach to proof and refutation in MALL: extended report. Available via the authors web pages, 2008.

[6] C. Faggian and M. Hyland. Designs, disputes and strategies. In *Computer Science Logic*, volume 2471 of *Lecture Notes in Computer Science*, pages 713–748, 2002.

[7] J.-Y. Girard. A fixpoint theorem in linear logic. An email posting to the mailing list linear@cs.stanford.edu, Feb. 1992.

[8] J.-Y. Girard. Locus solum. *Mathematical Structures in Computer Science*, 11(3):301–506, June 2001.

[9] L. Hallnäs and P. Schroeder-Heister. A proof-theoretic approach to logic programming. II. Programs as definitions. *J. of Logic and Computation*, 1(5):635–660, Oct. 1991.

[10] W. Hodges. Logic and games. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Stanford University, 2004.

[11] J. M. E. Hyland and C.-H. L. Ong. On full abstraction for PCF: I. models, observables and the full abstraction problem, II. dialogue games and innocent strategies, III. A fully abstract and universal game model. *Information and Computation*, 163:285–408, 2000.

[12] M. I. Kanovich. Simulating linear logic with 1-linear logic. Preprint 94-02, Laboratoire de Mathématiques Discrètes, University of Marseille, 1994.

[13] P. Lincoln. Deciding provability of linear logic formulas. In J.-Y. Girard, Y. Lafont, and L. Regnier, editors, *Proceedings of the workshop on Advances in linear logic*, pages 197–210, New York, NY, USA, 1995. Cambridge University Press.

[14] P.-A. Melliès. Asynchronous games 2: The true concurrency of innocence. *Theoretical Computer Science*, 358(2–3):200–228, 2004.

[15] D. Miller and A. Saurin. A game semantics for proof search: Preliminary results. In *Proceedings of the Mathematical Foundations of Programming Semantics (MFPS05)*, number 155 in Electr. Notes Theor. Comput. Sci, pages 543–563, 2006.

[16] D. Miller and A. Saurin. From proofs to focused proofs: a modular proof of focalization in linear logic. In J. Duparc and T. A. Henzinger, editors, *CSL 2007: Computer Science Logic*, volume 4646 of *LNCS*, pages 405–419. Springer-Verlag, 2007.

[17] D. Pym and E. Ritter. A games semantics for reductive logic and proof-search. In D. Ghica and G. McCusker, editors, *GaLoP 2005: Games for Logic and Programming Languages*, pages 107–123, 2005.