

Proof and refutation in MALL as a game

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

17 April 2009

Abstract

We present a setting in which the search for a proof of B or a refutation of B (*i.e.*, a proof of $\neg B$) can be carried out simultaneously: in contrast, the usual approach in automated deduction views proving B or proving $\neg B$ as two, possibly unrelated, activities. Our 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 counter-winning strategy translates to a refutation of the formula. The game is described for multiplicative and additive linear logic (MALL). 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

The connections between games and logic are numerous. For example, in the general area of the semantics of logic, games have played a significant role: descriptive set theorists make use of Banach-Mazur (forcing) games to build infinite structures with prescribed organization and model theorists use Ehrenfeucht-Fraïssé (back-and-forth) games to compare infinite structures. In the area of proof theory, proofs are occasionally used as winning strategies in “dialog games”: for example, if one has a proof of a formula, one should be able to defend against an opponent who might be skeptical of the truth of that formula.

Another possible connection between logic and games can be motivated by considering a common approach to proving the completeness of first-order logic, following, say, Smullyan [32]. To prove completeness, one attempts to build a tableau or sequent calculus proof of a given formula, say B , trying to complete an

incomplete proof. Such an attempt grows the incomplete derivation by adding inference rules at any unfinished leaf of the derivation. If one does so in a systematic fashion, one either succeeds to prove all remaining incomplete leaves or one is left with a tree of inferences with a (possibly infinite) branch that was never closed. In the first case, there is a proof of B and in the second case, the never-closed branch provides a falsifying model of B . This process can be viewed as the interaction of two agents working on this derivation tree. One agent attempts to finish the incomplete proof tree and the another agent attempts to find a path that is not completed. Clearly, only one of these agents eventually succeeds at their task.

Consider now two aspects of this outline. First, this completeness result relates, of course, provability and truth: a semantic notion of truth here seems forced since a set-theoretic model is a convincing means of structuring the information in an infinite path. If we restrict ourselves to a decidable logic (such as a propositional logic), then such non-closed paths can be expected to be always finite. In that case, one might be able to convert an open path into a proof of $\neg B$ (that is, a refutation of B) instead of a counter-model. In that case, both agents are attempting to construct proofs. Notice also that if we can view both agents as attempting to build proofs (one for B and the other for $\neg B$), the steps taken by these two agents appear to be rather different and there is no obvious reason to expect the kinds of proofs to be built by these two agents to be of the same style. One proof would be based on a tree and the other on a path.

In this paper, we describe a two-player game in which both players play exactly the same rules. If a player has a winning strategy, that player is able to construct a proof: for one of the players, this would be a proof of B while dually for the other player, this would be a proof of $\neg B$. Furthermore, the resulting proofs for both players are described using the same, simple sequent calculus proof system.

An interesting, degenerate version of such a game can be seen in the behavior of an idealized Prolog interpreter given a *noetherian* logic program Δ and query G . Here, the interpreter loads the program Δ and then becomes the first and only player to actually make moves: that is, the rules of the game will allow the interpreter to move repeatedly. The restriction that Δ is a Horn clause program means that there is no need to switch players (equivalent to switching phases in a later proof system) and the restriction that Δ is noetherian means that the interpreter will either end in a *finite success* or a *finite failure*: recall that a Prolog-like interpreter will explore all of the finitely many, finitely long, branches required to build a possible proof of G by repeatedly backtracking. In the first case, there is a proof of G from Δ and in the second case, there is a proof of $\neg G$ from (the if-and-only-if completion of) Δ [17, 13]. In either case, the first step is all that needs to be considered in order to determine the winner of the game.

In this paper, we present a two-person game in which one player is attempting to prove and the other is attempting to refute a formula from multiplicative and additive linear logic (MALL). The logic MALL is decidable (in fact, PSPACE-complete [21, 23]) and not complete in the sense that there are for-

mulas B such that neither B nor $\neg B$ are provable (consider, for example, the pair of de Morgan dual formulas $\perp \otimes \perp$ and $1 \wp 1$). This incompleteness makes games in this setting non-determinate (neither player might win) and, as a result, games in this setting need to be able to continue play after one player has failed. Positions in a game are graphs involving “neutral expressions”: the graph structure accounts for the multiplicative aspects of MALL while the neutral expressions are mapped into MALL formulas using either a positive or negative translation. Winning strategies yield proofs: depending on which player has a winning strategy, either the positive or the negative translation of the graph into logical formulas and sequents has a proof.

Games have been successfully used to study programming languages. In particular, in functional programming games allowed to solve long-standing problems such as the full-abstraction problem for PCF [2, 20]. This approach also provides models for logics capturing the dynamics of cut-elimination [1, 3].

Game-theoretical studies of logic programming are less common. Interestingly, those works can also be classified into two groups: (i) games modelling Prolog engines and (ii) games used to model the proof-theoretical foundations of logic programming, namely proof-search. In the first line of research, van Emden [34] provided the first game-theoretic interpretation of logic programming, connecting Prolog computations and two-person games using the $\alpha\beta$ -algorithm. Loddo *et al.* [7, 24] developed this approach and considered constraint logic programming [25]. Recently, Galanaki *et al.* [11] generalized van Emden’s games for logic programs with (well-founded) negation. In the second direction, Pym and Ritter [29] proposed games for uniform proofs and backtracking by relating intuitionistic and classical provability. Our present work can also be connected to recent research by the third author [31, 30] on modelling of proof-search in Ludics [16] as a process of interacting with tests. Compared to these works, the present paper and our previous works [28, 9] are guided by the so called “neutral approach.” In particular, while [30] deliberately chooses one player to be opposed to tests, we develop a framework in which both players have the same status, both attempting to prove a formula and to refute its negation. However, both approaches are inspired by the monistic program introduced in [15].

The contributions of this paper are the following.

(1) We present the *neutral approach to proof and refutation* and use it to describe a new game that can be seen as an attempt to simultaneously prove and refute a MALL formula.

(2) In order to deal with the fact that there are formulas B such that neither B nor $\neg B$ are provable, the game we describe *resumes* play after one player loses so that one 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. 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 through non-invertible inference rules in a proof.

(4) It is a common observation within the proof theory of sequent calculus that one can explain variations of logics by varying the structural rules. A similar observation is true for *focusing proof systems* [4]. Across a range of focusing proof systems for, say, MALL, the introduction rules remain the same while the structural rules can vary. For example, one can have systems that focus on a unique formula or on multiple formulas [6]. One can insist that an asynchronous phase terminates when all asynchronous formulas are removed but one can also allow for the phase to end before they are all removed. We have found that our use of games implies a particular set of choices that are natural and different from those in, say, Andreoli's original set of rules [4].

This paper is the result of merging and extending two previous conference papers by the authors [9, 28].

2 Logical Preliminaries

The formulas of MALL are built from literals (*i.e.*, atoms or negated atoms), the binary connectives \otimes , \oplus , \wp , $\&$, and their respective units 1 , 0 , \perp , \top . Figure 1 contains the rules for (cut-free) proofs in MALL. The connectives and units \otimes , \wp , 1 , \perp are *multiplicative*, while \oplus , $\&$, 0 , \top are *additive*. The exponentials $!$ and $?$ of linear logic and the quantifiers \forall and \exists of first-order or second-order logics are not considered in this paper.

The additive fragment of MALL is composed of those formulas that do not contain multiplicative connectives. Notice that if an additive formula has a proof Ξ using the proof rules in Figure 1 then all sequents in Ξ contain exactly one formula. Thus, Ξ does not contain any instances of the initial rule and, as a result, no role is played by atomic formulas in such formulas. Thus, we shall identify the *purely additive* fragment of linear logic as those formulas containing the additive connectives only (with no occurrences of atomic formulas).

$$\begin{array}{c}
 \frac{\vdash \Delta}{\vdash \Delta, \perp} \quad \frac{\vdash \Delta, F, G}{\vdash \Delta, F \wp G} \quad \frac{}{\vdash \Delta, \top} \quad \frac{\vdash \Delta, F \quad \vdash \Delta, G}{\vdash \Delta, F \& G} \\
 \frac{}{\vdash 1} \quad \frac{\vdash \Delta_1, F \quad \vdash \Delta_2, G}{\vdash \Delta_1, \Delta_2, F \otimes G} \quad \frac{\vdash \Delta, F_1}{\vdash \Delta, F_1 \oplus F_2} \quad \frac{\vdash \Delta, F_2}{\vdash \Delta, F_1 \oplus F_2} \quad \frac{}{\vdash F^\perp, F}
 \end{array}$$

Figure 1: A proof system for MALL. The last rule is called the *initial rule* and can be restricted so that F is atomic.

Based on the invertibility of introduction rules in Figure 1, we introduce a second classification of the logical connectives. An introduction rule is invertible when its conclusion is provable if and only if its premises are. The connectives \wp , $\&$, \perp , \top are called *asynchronous*: these are the connectives for which the

corresponding introduction rules in Figure 1 are invertible. The connectives $\otimes, \oplus, 1, 0$ are called *synchronous*: these are the connectives for which the corresponding introduction rules in Figure 1 are not necessarily invertible.

The proof system in Figure 2 is a restriction of the focusing proof system of Andreoli [4] to MALL. This focusing proof system is organized around groupings of introduction rules for asynchronous and for synchronous connectives. A stronger grouping is possible if we also classify literals as belonging to either the “asynchronous phase” or the “synchronous phase”. Since literals do not have introduction rules it seems better to introduce yet another dichotomy of logical formulas. A formula is *positive* if it is either synchronous or an atom and is *negative* if it is either asynchronous or the negation of an atom. The sequents used in Figure 2 are of two kinds. The *asynchronous* sequent $\Delta \uparrow \Gamma$ contains two multiset contexts Δ and Γ , where Δ contains only positive formulas or negated atoms. The *synchronous* sequent $\Delta \downarrow F$ contains one multiset context Δ containing only positive formulas or negated atoms. In this sequent, the formula F is called the *focus* of the sequent: this proof system provides a positive sequent with a unique focus. Later it will be natural to consider a generalization to this focusing proof systems that allows for *multiple foci* (see Figure 5). Notice that all the introduction rules from Figure 1 now have all their premises and conclusions labeled either with \uparrow or with \downarrow . The initial rule is classified as part of the positive phase since it is annotated with a \downarrow . To describe the remaining three “structural” rules, we read them bottom-up. For example, the $[R \uparrow]$ rule moves a positive formula or a negated atom out of the right context since they are not addressed in the asynchronous phase. The $[R \downarrow]$ stops the focusing (synchronous) phase when the focused formula becomes negative. Finally, the *decide rule* $[D]$ picks a positive formula on which to focus once the right context is empty (that is, all asynchronous connectives have been decomposed).

$$\begin{array}{c}
\frac{\vdash \Delta \uparrow \Gamma}{\vdash \Delta \uparrow \perp, \Gamma} [1] \quad \frac{\vdash \Delta \uparrow F, G, \Gamma}{\vdash \Delta \uparrow F \wp G, \Gamma} [\wp] \quad \frac{}{\vdash \Delta \uparrow \top, \Gamma} [\top] \quad \frac{\vdash \Delta \uparrow F, \Gamma \quad \vdash \Delta \uparrow G, \Gamma}{\vdash \Delta \uparrow F \& G, \Gamma} [\&] \\
\\
\frac{}{\vdash \cdot \downarrow 1} [1] \quad \frac{\vdash \Delta_1 \downarrow F \quad \vdash \Delta_2 \downarrow G}{\vdash \Delta_1, \Delta_2 \downarrow F \otimes G} [\otimes] \quad \frac{\vdash \Delta \downarrow F_1}{\vdash \Delta \downarrow F_1 \oplus F_2} [\oplus_l] \quad \frac{\vdash \Delta \downarrow F_2}{\vdash \Delta \downarrow F_1 \oplus F_2} [\oplus_r] \\
\\
\frac{}{\vdash A^\perp \downarrow A} [I] \quad \frac{\vdash \Delta \downarrow P}{\vdash \Delta, P \uparrow} [D] \quad \frac{\vdash \Delta, P_a \uparrow L}{\vdash \Delta \uparrow P_a, L} [R \uparrow] \quad \frac{\vdash \Delta \uparrow N}{\vdash \Delta \downarrow N} [R \downarrow]
\end{array}$$

Figure 2: A focused proof system for MALL. Here, A is an atomic formula, P is positive, P_a is positive or a negative atom, and N is negative.

It is worth noting that if we limit ourselves to the purely additive fragment of MALL, then any cut-free MALL proof of it is already focused. To show this, simply annotate the unique formula in all sequents of a MALL proof as follows: if the top-level connective is \top or $\&$ then use the arrow \uparrow and if the top-level connective is \oplus then use the \downarrow . Extra inference rules must be added to switch between these two arrows: use either $[R \downarrow]$ to switch from a synchronous

conclusion to an asynchronous premise or $[R \uparrow]$ and $[D]$ to switch from an asynchronous conclusion to a synchronous premise. In Section 3, we shall consider “additive games” that are based on purely additive formulas: in that discussion, there is little need to mention focusing since the focused and unfocused proofs systems are isomorphic. In Section 4, we consider “simple games”: these games correspond the situation where the sequents in the decide rule $[D]$ have a unique formula (that is, Δ is empty). In such situations, the decide rule is forced in its selection of what formula to consider next. In Section 5, we finally deal with full MALL logic (including atomic formulas).

3 Additive games

Hintikka (see, for example, [19]) defined a simple game to determine the truth of a formula as follows (the game can also work for quantificational formulas). Two players, A and E , play with a single formula. The player A tries to falsify the formula while E tries to validate the formula. If the formula is a conjunction ($\&$), A must move by choosing one of the conjuncts: in particular, if the formula is the empty conjunction (\top), then A can pick nothing and she loses. If the formula is a disjunction (\oplus), E must move by choosing one of the disjuncts: in particular, if the formula is the empty disjunction (0), then E can pick nothing and she loses. This game is *determinate* in the sense that one player always has a winning strategy. If A has a winning strategy starting with ϕ then ϕ is false; conversely if E has a winning strategy starting with ϕ then ϕ 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 $0, \oplus, \top, \&$. We describe this game in some detail here as an introduction to the neutral approach.

3.1 Neutral expressions

In the game presented above, each logical connective is assigned to a player who decomposes it, and the main connective of a formula determines whose turn it is. The game is symmetric in the sense that A treats a conjunction in the same way that E treats a disjunction. In our neutral approach, we make this observation clearer by introducing *neutral expressions*: these expressions represents a pair of dual formulas. In particular, the syntax of neutral expressions contains a single constructor for each pair of dual connectives and units. 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 use the special unary operator \updownarrow to this end.

Definition 3.1. *The (additive) neutral expressions E and guarded neutral expressions G are defined as follows:*

$$G ::= \mathbf{0} \mid E + E \quad E ::= G \mid \updownarrow G$$

$$\begin{array}{ll}
[\mathbf{0}]^+ = 0 & [\mathbf{0}]^- = \top \\
[E + F]^+ = [E]^+ \oplus [F]^+ & [E + F]^- = [E]^- \& [F]^- \\
[\Downarrow E]^+ = [E]^- & [\Downarrow E]^- = [E]^+
\end{array}$$

Figure 3: Translations of additive neutral expressions into formulas

A guarded neutral expression is therefore a neutral expression which does not begin with \Downarrow . Notice that $\Downarrow(\Downarrow E)$ is not a subexpression of a neutral expression.

We define two *translations* (functions) that map a neutral expression into the dual formulas it represents.

Definition 3.2. *The positive and negative translations of neutral expressions into formulas of the additive fragment of linear logic are defined in Figure 3. 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.*

3.2 A game based on neutral expressions

For determinate games, a simple notion of game and arena is appropriate. For this section, we take an *arena* to be a directed graph (\mathcal{P}, ρ) , where \mathcal{P} is a set of *positions* and $\rho \subseteq \mathcal{P} \times \mathcal{P}$ is the *move* relation. A *final* position is a sink, *i.e.*, a position with no ρ -successor. A *play* from a position P is a (finite or infinite) sequence of ρ -related positions starting in P , that is, a sequence $(P = P_0, P_1, \dots)$ such that for every $i \geq 0$, $P_i \rho P_{i+1}$. We will suppose that ρ is *noetherian*, *i.e.*, that there are no infinite plays. The *length* of a play (P_0, \dots, P_n) is n .

A *winning strategy* \mathcal{S} from a position P is a set of plays from P with the following properties:

- $(P) \in \mathcal{S}$;
- \mathcal{S} is prefix closed: that is, for every $\pi \in \mathcal{S}$ and every prefix π' of π , $\pi' \in \mathcal{S}$;
- for every play $(P_0, \dots, P_n) \in \mathcal{S}$ such that n is even, there exists $P_{n+1} \in \mathcal{P}$ such that $P_n \rho P_{n+1}$ and $(P_0, \dots, P_{n+1}) \in \mathcal{S}$;
- for every play $(P_0, \dots, P_n) \in \mathcal{S}$ such that n is odd, for every $P_{n+1} \in \mathcal{P}$ such that $P_n \rho P_{n+1}$, we have $(P_0, \dots, P_{n+1}) \in \mathcal{S}$;

The definition of a *winning counter-strategy* \mathcal{S} from a position P is obtained by swapping the words “even” and “odd” in the two last properties.

In order to define the arena of the particular game we present here, we first introduce a rewriting relation on neutral expressions:

$$E_1 + E_2 \rightarrow E_1 \quad E_1 + E_2 \rightarrow E_2$$

Expressions of the form $\mathbf{0}$ and $\Downarrow E$ do not rewrite. The reflective and transitive closure of \rightarrow is written as \rightarrow^* .

In our game the positions are the guarded neutral expressions and a move from E to F takes place exactly when $E \rightarrow^* \downarrow F$. To establish a correspondence between winning strategies and proofs, we first need the following lemma.

Lemma 3.3. *Let E be a neutral expression. Let $S = \{\downarrow F : E \rightarrow^* \downarrow F\}$.*

- *For every $F \in S$, the sequent $\vdash [E]^+$ derives from $\vdash [F]^+$.*
- *The sequent $\vdash [E]^-$ derives from the sequents $\{\vdash [F]^- : F \in S\}$.*

Proof. Let us show these properties by induction on E .

If $E = \mathbf{0}$, then S is empty and the first property is true. The second property expresses that $\vdash [E]^-$ (i.e., $\vdash \top$) is provable, which immediately follows from the introduction rule for \top .

If E is of the form $\downarrow F$, then $S = \{E\}$ and the two properties are trivial.

Suppose now that E is of the form $E_1 + E_2$. Then $S = S_1 \cup S_2$, where $S_i = \{\downarrow F : E_i \rightarrow^* \downarrow F\}$ for every $i \in \{1, 2\}$. By induction hypothesis, the properties hold for E_i with respect to S_i , for every $i \in \{1, 2\}$. Let us prove the first property. Let $F \in S$. There is some $i \in \{1, 2\}$ such that $F \in S_i$ and, hence, $\vdash [E_i]^+$ derives from $\vdash [F]^+$. All we have to do is show that $\vdash [E]^+$ (i.e., $\vdash [E_1]^+ \oplus [E_2]^+$) derives from $\vdash [E_i]^+$, which is easily seen by applying the introduction rule for \oplus . Let us now prove the second property. The sequents $\vdash [E_1]^-$ and $\vdash [E_2]^-$ derive from $\{\vdash [F]^- : F \in S_1\}$ and $\{\vdash [F]^- : F \in S_2\}$, respectively, therefore both derive from $\{\vdash [F]^- : F \in S\}$. All we need to do is show that $\vdash [E]^-$ (i.e., $\vdash [E_1]^- \& [E_2]^-$) derives from $\vdash [E_1]^-$ and $\vdash [E_2]^-$, which follows by applying the introduction rule for $\&$. \square

There is a converse to this lemma. Namely, every cut-free proof of $\vdash [E]^+$ derives its conclusion from $\vdash [F]^+$, for some $F \in S$; and every cut-free proof of $\vdash [E]^-$ derives its conclusion from the sequents in $\{\vdash [F]^- : F \in S\}$. We do not give the proof here for the sake of brevity, but it is easily obtained by reversing the above arguments.

We can now prove the following theorem, which establishes the correspondence between winning strategies and proofs in this purely additive setting.

Theorem 3.4. *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 provides the corresponding proof.*

Proof. We know that $\vdash [E]^+$ and $\vdash [E]^-$ cannot be both provable and that there cannot exist both a winning strategy and a winning counter-strategy from E . It is therefore enough to show that either $\vdash [E]^+$ is provable and there exists a winning strategy from E , or that $\vdash [E]^-$ is provable and there exists a winning counter-strategy from E .

The length of a play from E is bounded by the maximal number of nested \downarrow in E . Let us prove our claim by induction on the maximal length of a play $n(E)$. If $n(E) = 0$, then E is a final position and there is a winning counter-strategy.

In the above lemma S is empty and there is a proof of $\vdash [E]^-$. Suppose now that $n(E) > 0$. In the above lemma S is not empty. We consider two cases. First case: there exists a ρ -successor F of E such that there is a winning counter-strategy from F . Then there is a winning strategy from E (just prepend E to all plays). We have $E \rightarrow^* \downarrow F$ and $n(F) < n(E)$. By the induction hypothesis, $\vdash [F]^-$ is provable. By the previous lemma $\vdash [E]^+$ derives from $\vdash [\downarrow F]^+$, which is $\vdash [F]^-$. Therefore $\vdash [E]^+$ is provable. Second case: there is a winning strategy from every ρ -successor F of E . Then there is a winning counter-strategy from E (just take the union of those strategies and prepend E to all plays). Define S as in the above lemma. For every $\downarrow F \in S$, $n(F) < n(E)$; therefore, by induction hypothesis, $\vdash [F]^+$ (i.e., $\vdash [\downarrow F]^-$) is provable. By the previous lemma $\vdash [E]^-$ derives from those sequents, hence, it is provable. \square

In Hintikka's game, the same player might move several times in a row, as long as the principal connective of the formula remains the same. In our game these moves correspond to individual rewrite steps and a move is a maximal sequence of such steps, which ensures a strict alternation of players. From the point of view of the proof objects, a rewrite step (a.k.a. a *micro-move*) corresponds the introduction of an individual connective or unit, while a move (a.k.a. a *macro-move*) corresponds to a full phase. The games considered in the next sections are more complex but their moves still have those two levels.

Another important remark is that this game may be seen as a process accounting for the simultaneous development of two dual derivations: starting from a neutral expression E , the player who begins sees the game as an attempt to derive $\vdash \downarrow [E]^+$, while her opponent sees it as an attempt to derive $\vdash \uparrow [E]^-$. With this in mind, each player develops what she sees as synchronous phases during her turn and leaves it to her opponent to decompose asynchronous phases.

4 Simple games

In this section we present an extension of the additive games that incorporates some multiplicative behavior. This first extension yields the so-called *simple* games: as we shall see, these games are too "simple" to properly handle the full range of multiplicative connectives. As a result, the simple games help to illustrate the need for the more sophisticated neutral graph structures presented in Section 5.3. There are, however, a number of examples of games that are "simple-like": see [28] for examples of simple games related to computing on finite sets and with determining bisimulation in labelled transition systems. We do not present those examples here since they involve adding operators for quantification and fixed points to neutral expressions and these extensions are a significant departure from MALL.

Definition 4.1 (Neutral expressions). *We extend the language of neutral expressions by introducing a multiplicative connective and its unit.*

$$G ::= \mathbf{0} \mid \mathbf{1} \mid E + E \mid E \times E \quad E ::= G \mid \uparrow G$$

$$\begin{array}{ll}
[\mathbf{0}]^+ = 0 & [\mathbf{0}]^- = \top \\
[\mathbf{1}]^+ = 1 & [\mathbf{1}]^- = \perp \\
[E + F]^+ = [E]^+ \oplus [F]^+ & [E + F]^- = [E]^- \& [F]^- \\
[E \times F]^+ = [E]^+ \otimes [F]^+ & [E \times F]^- = [E]^- \wp [F]^- \\
[\Downarrow E]^+ = [E]^- & [\Downarrow E]^- = [E]^+
\end{array}$$

Figure 4: Translations of neutral expressions into MALL

Micro-dynamics of neutral expressions. Multisets of neutral expressions can be rewritten non-deterministically as follows:

Definition 4.2 (Rewriting for neutral expressions). *The binary relation \mapsto between finite multisets of neutral expressions is given as follows:*

$$\begin{array}{ll}
\mathbf{1}, \Gamma \mapsto \Gamma & E \times F, \Gamma \mapsto E, F, \Gamma \\
E + F, \Gamma \mapsto E, \Gamma & E + F, \Gamma \mapsto F, \Gamma
\end{array}$$

Let \mapsto^* be the reflective and transitive closure of \mapsto .

If we consider the *size of a multiset of neutral expressions* to be the total number of occurrences of constructors in expressions in that multiset, then the size of multisets decreases as they are rewritten. As a result, rewriting always terminates. Notice that an expression of the form $\Downarrow E$ is not rewritten: it represents a formula that is left by one player for the other player.

We shall be interested in whether or not an expression E (considered as a singleton multiset) rewrites (via \mapsto^*) to $\{\Downarrow E_1, \dots, \Downarrow E_n\}$ ($n \geq 0$). This last type of multiset will be written $\Downarrow \Gamma$ if Γ is the multiset $\{E_1, \dots, E_n\}$. Notice that if $\mathbf{0} \in \Gamma$, then Γ cannot reduce to $\Downarrow \Gamma'$.

In Figure 4, the updated positive and negative translations of neutral expressions into MALL are provided. As before, layers of neutral connectives between \Downarrow translate to phases.

A useful measure of a neutral expression is the maximum number of expressions beginning with \Downarrow that it can yield on some non-deterministic rewriting.

Definition 4.3 ($\natural(E)$). $\natural(\cdot)$ is defined to assign a natural number to a neutral expression in the following way:

- $\natural(\mathbf{0}) = \natural(\mathbf{1}) = 0$;
- $\natural(\Downarrow E) = 1$;
- $\natural(E_1 + E_2) = \max(\natural(E_1), \natural(E_2))$ and
- $\natural(E_1 \times E_2) = \natural(E_1) + \natural(E_2)$.

Clearly, for any neutral expression E , $\natural(E) = 0$ if and only if E does not contain a \Downarrow .

Proposition 4.4. *Let E be a neutral expression containing no \downarrow . If $E \mapsto^* \{\}$ then $\vdash\downarrow [E]^+$. If E cannot be rewritten to $\{\}$ then $\vdash\uparrow [E]^-$.*

Proof. Let $k \geq 0$ and define \mapsto^k to be the k -fold join of \mapsto (in particular, \mapsto^0 is multiset equality). We prove by induction on k that if $n \geq 0$ and $\{E_1, \dots, E_n\} \mapsto^k \{\}$ then for all $j \in \{1, \dots, n\}$, $\vdash\downarrow [E_j]^+$. If $k = 0$ then $n = 0$ then the conclusion is immediate. Assume that $k > 0$ and that $\{E_1, \dots, E_n\} \mapsto^k \{\}$. Consider the cases for the first step of this rewriting. The result follows easily by noticing that the rewriting rules for neutral expressions correspond to introduction rules for their positive translations.

Next we show that by induction on the size of multisets of neutral expressions (where one counts the number of occurrences of constructors of such expressions as their size) that if $\{E_1, \dots, E_n\}$ is a multiset of neutral expressions which does not rewrite to $\{\}$ then the sequent $\vdash\uparrow [E_1]^-, \dots, [E_n]^-$ is provable. It is clearly the case if $n = 0$. If $n > 0$, then it can be easily seen by examining the cases for E_n that $\vdash\uparrow [E_1]^-, \dots, [E_n]^-$ is either immediately provable (case $E_n = \mathbf{0}$) or that rewriting the multiset by decomposing E_n corresponds exactly to introducing the principal connective of $[E_n]^-$. \square

A class of neutral expressions that will be of particular interest in the rest of this section is the class of *simple expressions*:

Definition 4.5 (Simple expression). *An expression E is simple if $\mathfrak{h}(E) \leq 1$ and for every subexpression $\uparrow E'$ of E , E' is simple.*

A multiset $\{E_1, \dots, E_k\}$ (where $k \geq 0$) of expressions is simple if $E_1 \times \dots \times E_k$ is simple.

Simple expressions can alternatively be defined by the following grammar:

$$\begin{aligned} Z &::= \mathbf{0} \mid \mathbf{1} \mid Z + Z \mid Z \times Z \\ S &::= Z \mid S + S \mid Z \times S \mid S \times Z \mid \uparrow S, \end{aligned}$$

where S and Z are syntactic variables ranging over simple expressions and expressions without occurrences of \downarrow , respectively.

We now describe the arena of games involving simple expressions. Let the set of positions be the set of neutral expressions. The move relation, ρ , is defined as the smallest relation such that $E \rho \mathbf{0}$ if $E \mapsto^* \{\}$ and $E \rho F$ if $E \mapsto^* \{\uparrow F\}$. The fact that there are only these two cases possible is the key feature of simple games: that is, if we restrict E to be a simple expressions, then it is not possible for $E \mapsto^* \{\uparrow F_1, \dots, \uparrow F_n\}$ and for $n > 1$. That is, while multiplicative expressions can be treated internally to one player, the multiset of expressions must degenerate to leave at most a single expression: thus, when the players switch, the game must appear to be, essentially, additive.

We now consider the nature of winning strategies and winning counter-strategies based on this move relation.

If E does not contain \uparrow then there is either no move from E or the only move possible is to $\mathbf{0}$. In the first case, there exists a winning counter-strategy for E .

In the second case, there is a winning strategy for E (since the second player can make no move from $\mathbf{0}$).

Since all plays are finite and all final positions for a game are classified as a win for one player or the other, games for simple expressions are *determinate*: that is, given a simple expression E , there is either a winning strategy or a winning counter-strategy for E .

Proposition 4.6. *Let E be a simple expression. There is a winning strategy for E iff $\vdash\Downarrow [E]^+$. There is a winning counter-strategy for E iff $\vdash\Uparrow [E]^-$.*

Proof. Given a winning (counter-)strategy we build by induction a proof for either $\vdash\Downarrow [E]^+$ or $\vdash\Uparrow [E]^-$ depending on the sort of strategy we have. Notice that in winning strategies all the branches have same parity. We thus reason on the length of the largest branch in the strategy, which we refer to as the *size* of the strategy.

Base case: strategy σ has size 0. We are trying to build a proof for $\vdash\Uparrow [E]^-$ given that σ is a winning counter-strategy. Thus E is such that $E \not\mapsto^* \Downarrow F$ nor $E \not\mapsto^* \{\}$, that is any maximal internal derivation from E ends up with a multiset of expressions that contains some expressions not beginning with \Downarrow and which cannot be rewritten (by maximality). These multisets are made of possibly one expression beginning with \Downarrow (but not more because E is simple) and of at least one $\mathbf{0}$: they are the only kind of multiset that cannot be rewritten and that are not legal positions: $\{\{\Downarrow F\}, \mathbf{0}, \dots, \mathbf{0}\}$.

Since $\vdash\Uparrow [E]^-$ is a negative sequent, it is possible to generate part of a proof tree by applying negative logical rules in any order up to a point where no negative rule can be applied (with the additional constraint that the \top rule is applied only when all the negative formulas have been decomposed). Such a tree is actually a proof. Indeed if any branch of the tree leads to a non justified sequent, that means that either a sequent made of exactly one positive formula or an empty sequent is reached. In any other situation, the branch would be extendable. These two situations contradict the fact that no internal derivation from E can lead to a legal position: it is straightforward to see that any branch of the proof tree from $\vdash\Uparrow [E]^-$ to one of the two kinds of sequents just mentioned can be transformed into an internal derivation from E to a legal position. Thus the negative tree is a proof.

Base case: strategy σ has size 1. We are now building a proof for $\vdash\Downarrow [E]^+$ given that σ is a winning strategy. Thus σ starts with a move: $E \rho F$. There are two cases: either F is $\mathbf{0}$ and $E \mapsto^* \{\}$ or F is any expression and $E \mapsto^* \Downarrow F$: in this case, the move is extended by a winning counter-strategy σ' of length 0 for F . In both cases we pick any internal derivation justifying the move and use it to inductively build a proof for $\vdash\Downarrow [E]^+$. In the second case we will additionally refer to the base case for size 0 in order to have a proof of $\vdash\Uparrow [F]^-$.

We reason by induction on the length of the considered internal derivation maintaining the following invariant: if the derivation starts with E_1, \dots, E_k then the sequents $\vdash\Downarrow [E_1]^+, \dots, \vdash\Downarrow [E_k]^+$ are all provable. The invariant is true for empty derivations: by hypothesis, we have either an empty multiset, and the assertion is trivial, or the multiset consists in a singleton $\Downarrow F$ and the

invariant requires $\vdash\Downarrow [\Downarrow F]^+$ to be provable and base case for size 0 on σ' tells us that $\vdash\Uparrow [F]^-$ is provable. If the induction hypothesis is true for a derivation of length n it is also true for a derivation of length $n + 1$. Indeed let ρ be an internal derivation of length $n + 1$ from multiset E_1, \dots, E_k that is $\rho : E_1, \dots, E_k \mapsto_r E'_1, \dots, E'_l \mapsto^* \dots$. Induction hypothesis for the derivation starting with E'_1, \dots, E'_l ensures that $\vdash\Downarrow [E'_1]^+, \dots, \vdash\Downarrow [E'_l]^+$ are all provable. Finally, r corresponds to introducing the main connective of $[E_i]^+$, where E_i is the neutral expression of the multiset being decomposed by r .

Induction case. Two cases need to be considered: σ has even size $n + 1$ (it is winning counter-strategy) or σ has odd size $n + 1$ (it is a winning strategy).

1. σ is of odd size. This case is similar to the base case for strategies of size 1 when σ starts with a move $E \rho F$ and yields σ' as a winning counter-strategy for F . Using the induction hypothesis on σ' we obtain a proof

$$\frac{\Pi}{\vdash\Uparrow [F]^-}$$

and by picking some internal derivation $\rho : E \mapsto_\rho^* \Downarrow F$ we build an open positive derivation rooted in $\vdash\Downarrow [E]^+$:

$$\frac{\frac{\vdash\Uparrow [F]^-}{\vdash\Downarrow [\Downarrow F]^+}}{\vdash\Downarrow [E]^+}$$

that we can close with Π .

2. σ is of even size. Once more, the proof is close to the base case. Considering a negative tree built thanks to the same process as the one described previously, the same three cases might occur: either the branch is closed, that is the corresponding internal derivation does not lead to a legal position (\Downarrow s appear in the multiset and on the proof-theoretical side a \top rule is used), or we get to a sequent made of a positive formula A (there is a move $E \rho F$ with $[F]^+ = A$, the induction hypothesis allows us to proceed), or we get to an empty sequent (but this case is impossible for parity reasons on the length of the strategy which is required to be winning).

We give only a sketch of the proof for the second part of the proposition which essentially relies on the completeness of focused proofs [4]. We additionally require that the rule for \top is used when no other rule can be applied. We use such a proof to build a winning strategy. Each positive layer of the proof results in a first move of a winning strategy (the detail of this positive part of the proof can be mapped to an internal derivation justifying the move). Each negative layer of the proof results in a \forall -branching in the strategy. More precisely an additive slice of a negative layer results in a move of a counter-strategy (or an absence of move, see below). It is easy to check that all these moves are

justified by internal derivations and that the branching is complete. Concerning the parity condition, it is due to the fact that the proof can end with either a \perp rule or a \top rule. In the first case, the corresponding position in the game is an empty position ($\{\}$ also denoted by $\mathbf{0}$) immediately following a move by the player, that is a position from which no move is possible. In the second case the last multiset of a maximal internal derivation corresponding to this branch of the proof contains a $\mathbf{0}$ so that it cannot justify a move (it cannot be of the form $\uparrow\Gamma$). \square

Since all terminal positions for a game are classified as a win for one player or the other, games for simple expressions are determinate: that is, given a simple expression E , there is either a winning strategy or a winning counter-strategy for E . Thus, the excluded middle $\vdash [E]^- \oplus [E]^+$ is provable for all E . The next corollary actually follows immediately from the preceding Proposition:

Corollary 4.7. *Let E be a simple expression. Either $\vdash [E]^-$ or $\vdash [E]^+$ is provable.*

Interests and limitations of the simple fragment

Before moving to the more general case of games for MALL, let us restate the main characteristics of simple games, emphasizing both what is interesting in these games and why a more elaborate framework is required:

- *simple games provide a simple, though expressive framework.* As already mentioned at the beginning of this section, we emphasize that despite being very simple, the fragment of simple expressions is already fairly expressive. We show in [28] that a large number of games are actually “simple”. Moreover, this fragment is flexible enough to allow for an easy extension with quantification and fixed points;
- *simple games characterize a complete fragment of MALL.* Simple games are determinate. As a consequence, corollary 4.7 states that simple expressions characterize a complete fragment of MALL. Other complete fragments of logic might be looked for by building corresponding determinate games;
- *simple games are too restricted to capture MALL.* For the very reason they are determinate, there is no hope that simple games can capture MALL. Moreover, no determinate game can achieve this. Extending our game-theoretical framework to capture MALL will thus require to modify significantly the structure of our games, in order to lose determinacy for instance.

5 Games for MALL

In this section, we finally develop a game for full MALL by allowing both non-simple neutral expressions as well as allowing atomic formulas (propositional

variables).

5.1 Two-player games with ties

The fragments of MALL considered in the previous sections were complete (the games were determinate). The full logic of MALL is no longer complete so we must modify the description of games to account for the possibility of ties.

Name two-players 0 and 1. For $\sigma \in \{0, 1\}$, set $\bar{\sigma} = 1 - \sigma$. We still base our game on an arena (\mathcal{P}, ρ) . A position from which no move is possible 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 . We usually assume ρ to be noetherian, so that all plays are finite. 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 $(P_0, \dots, P_n) \in \mathcal{S}$

- if P_n is a σ -position, there exists P_{n+1} s.t. $P_n \rho P_{n+1}$, $(P_0, \dots, P_{n+1}) \in \mathcal{S}$,
- if P_n is a $\bar{\sigma}$ -position, for every P_{n+1} s.t. $P_n \rho P_{n+1}$, $(P_0, \dots, P_{n+1}) \in \mathcal{S}$.

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 player σ .

5.2 Proof system

In our neutral approach, a single (neutral) object accounts for two dual derivations being developed simultaneously, each player “viewing” one of them. As we saw in section 4, focalization plays an important role and brings some symmetry. However, the proof system used in the previous games lacks some important features and this motivates a change. Specifically, we need to address two main shortcomings.

1. In the previous games, a play would end as soon as the current player could not develop her derivation any more, making her opponent win immediately. We now have games based on neutral expressions E such that neither $[E]^+$ nor $[E]^-$ are provable. When a player σ fails to develop her derivation, we cannot conclude that $\bar{\sigma}$ wins as we did before. The play must continue until $\bar{\sigma}$ completes her derivation (thus winning the play) or fails as well (thus ending the play in a tie). The derivations must therefore

be richer objects which leave some room for failure. For example, consider the following dual derivations:

$$\frac{\frac{\frac{\vdash \uparrow A \quad \overline{\vdash \uparrow \top}}{\vdash \downarrow A \otimes \top}}{\vdash \downarrow A \otimes \top}}{\vdash \downarrow A \otimes \top} \quad \frac{\frac{\frac{\vdash A^\perp \downarrow 0}{\vdash A^\perp, 0 \uparrow}}{\vdash \uparrow A^\perp \wp 0}}{\vdash \uparrow A^\perp \wp 0}}$$

The second derivation may not be developed, but we still need a way to challenge the first derivation, which may or may not be completed into a proof, depending on A . This is done by adding a special inference rule called “daimon” (by analogy with Ludics [16]) to the proof system. A proof is now a closed derivation which does not use the daimon rule. A player who uses daimon cannot win, but may try to make her opponent fail as well.

2. In Andreoli’s Σ_3 [4], *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 in a synchronous phase, and *some*, not necessarily all, asynchronous formulas to be decomposed in an asynchronous phase.

Figure 5 shows our proof system. Sequents are of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \downarrow \mathcal{F}$ or $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}$, where \mathcal{L} is a multiset of literals, \mathcal{P} is a multiset of positive formulas, \mathcal{N} is a multiset of negative formulas, and \mathcal{F} is a multiset of formulas. The sequents $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \downarrow \cdot$ and $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \cdot$ are identified and also denoted to by $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N}$.

5.3 Neutral expressions

Definition 5.1. *We extend the syntax of neutral expressions and guarded neutral expressions as follows:*

$$G ::= k \mid \mathbf{0} \mid \mathbf{1} \mid E + E \mid E \times E \quad E ::= G \mid \downarrow G$$

where k denotes a neutral atom. There is one neutral atom for each pair of dual literals in the logic.

In the rest of the paper, the set of the neutral expressions is denoted by \mathcal{E} .

Definition 5.2. *The updated positive and negative translations of neutral expressions into MALL formulas are defined in Figure 6.*

5.4 Neutral graphs

In order to account for the complexity and intensional behavior of the multiplicative connectives and atoms of MALL, we shall not enrich the structure of

Additives

$$\frac{\frac{\frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow F_i, \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow F_1 \oplus F_2, \mathcal{F}} [\oplus_i]}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow F_1, \mathcal{F}} \quad \frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow F_2, \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow F_1 \& F_2, \mathcal{F}} [\&]}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow \top, \mathcal{F}} [\top]$$

Multiplicatives

$$\frac{\frac{\frac{\vdash \mathcal{L}_1; \mathcal{P}_1; \mathcal{N}_1 \Downarrow F_1, \mathcal{F}_1}{\vdash \mathcal{L}_1, \mathcal{L}_2; \mathcal{P}_1, \mathcal{P}_2; \mathcal{N}_1, \mathcal{N}_2 \Downarrow F_1 \otimes F_2, \mathcal{F}_1, \mathcal{F}_2} [\otimes]}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow F_1, F_2, \mathcal{F}} [\wp]} \quad \frac{\frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow \perp, \mathcal{F}} [\perp]}{\vdash \cdot; \cdot; \cdot \Downarrow 1} [1]}$$

Literals

$$\frac{}{\vdash K^\perp; \cdot; \cdot \Downarrow K} [init] \quad \frac{\vdash K^\perp; \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow K^\perp, \mathcal{F}} [R \Uparrow \text{atomic}]$$

Daimon

$$\frac{}{\Sigma} [\boxtimes] \quad \frac{\Sigma'}{\Sigma} [\boxtimes]$$

Phase changes

$$\frac{\frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N}, \mathcal{N} \Downarrow \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow N, \mathcal{F}} [R \Downarrow]}{\vdash \mathcal{L}; \mathcal{P}_1; \cdot \Downarrow \mathcal{P}_2} [D \Downarrow]} \quad \frac{\frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Uparrow P, \mathcal{F}} [R \Uparrow]}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N}_1 \Uparrow \mathcal{N}_2, \mathcal{F}} [D \Uparrow]}$$

K denotes a positive literal, N a negative formula and P a positive formula. In $[D \Downarrow]$ (resp. $[D \Uparrow]$), \mathcal{P}_2 (resp. \mathcal{N}_2) is not empty. A proof is a closed derivation which does not use \boxtimes .

Figure 5: The proof system used in the game for MALL

$$\begin{array}{ll} [k]^+ = K & [k]^- = K^\perp \\ [\mathbf{0}]^+ = 0 & [\mathbf{0}]^- = \top \\ [\mathbf{1}]^+ = 1 & [\mathbf{1}]^- = \perp \\ [E + F]^+ = [E]^+ \oplus [F]^+ & [E + F]^- = [E]^- \& [F]^- \\ [E \times F]^+ = [E]^+ \otimes [F]^+ & [E \times F]^- = [E]^- \wp [F]^- \\ [\Downarrow E]^+ = [E]^- & [\Downarrow E]^- = [E]^+ \end{array}$$

Figure 6: Translations of neutral expressions

$$\begin{array}{c}
\frac{\frac{\frac{\vdash A, B \uparrow}{\vdash \uparrow A, B} [R \uparrow]}{\vdash \uparrow A \wp B} [\wp]}{\vdash \downarrow A \wp B} [R \downarrow] \quad \frac{\vdash \uparrow C^\perp}{\vdash \downarrow C^\perp} [R \downarrow]}{\frac{\vdash \downarrow (A \wp B) \otimes C^\perp}{\vdash (A \wp B) \otimes C^\perp \uparrow} [D]} [\otimes] \tag{1}
\end{array}$$

$$\begin{array}{c}
\frac{\frac{\frac{\vdash C \uparrow A^\perp}{\vdash C \downarrow A^\perp} [R \downarrow] \quad \frac{\vdash \uparrow B^\perp}{\vdash \downarrow B^\perp} [R \downarrow]}{\vdash C \downarrow A^\perp \otimes B^\perp} [\otimes]}{\frac{\vdash A^\perp \otimes B^\perp, C \uparrow}{\vdash \uparrow A^\perp \otimes B^\perp, C} [R \uparrow]} [D] \quad \frac{\vdash \uparrow A^\perp \otimes B^\perp, C}{\vdash \uparrow (A^\perp \otimes B^\perp) \wp C} [\wp]}{\vdash \uparrow (A^\perp \otimes B^\perp) \wp C} [R \uparrow] \tag{2}
\end{array}$$

A , B and C are synchronous formulas.


Figure 7: Two dual derivations in Andreoli's Σ_3 .

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 previous games) to labeled graph structures, which we describe next.

Figure 7 shows an example of two dual derivations in Andreoli's Σ_3 . 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 $\vdash (A \wp B) \otimes C^\perp \uparrow$ and the frontier of (2) consists of $\vdash \uparrow (A^\perp \otimes B^\perp) \wp C$. 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 .

This tight correspondence is best seen by cutting two dual derivations together and applying the cut elimination procedure. At each step, the current state consists of an upper layer of subderivations of the original derivations, whose conclusions are fed to a lower layer of cuts inferring the empty sequent. Each cut effectively pairs two dual formulas from those conclusions. The order in which those cuts are applied is irrelevant. One way to abstract away from this order it to pack the full layer of cuts in a synthetic “multicut” inference rule. Another way is to consider proof structures and cut links between them [12, 14] or cut-net as introduced in Ludics [16]. The state of our game will be a representation of those cut-nets by means of a graph structure called a *neutral graph*.


This approach is inspired by [8], in which Danos and Regnier analyze the geometry of generalized multiplicative rules and express the duality of two generalized multiplicatives through a graph structure. Following this idea, we define neutral graphs to represent cut links between two (slices of) frontiers as presented above. The vertices represent the sequents of the slices. There are two colors of vertices (one for each frontier). 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

$$\vdash \cdot; [E]^+, [F]^+; \cdot \quad \vdash \cdot; \cdot; [G]^- \quad | \quad \vdash \cdot; [G]^+; [E]^- \quad \vdash \cdot; \cdot; [F]^-$$

will be represented by the neutral graph



where the black (resp. white) vertices represent the sequents of the left (resp. right) frontier. In the examples given so far, all the sequents are of the form $\vdash \cdot; \mathcal{P}; \mathcal{N}$, but we also need to represent sequents of the forms $\vdash \cdot; \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}$ and $\vdash \cdot; \mathcal{L}; \mathcal{P}; \mathcal{N} \downarrow \mathcal{F}$. We will use several types of arcs to this end.

It should be noted that the graphs themselves are not proof structures. They merely represent cut links between goals for the players to prove, which are placeholders for proof structures.

Definition 5.3. A neutral graph G is a tuple (V, A, p, t, ϵ) , 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 player to each vertex, $t : A \mapsto \{\text{atomic}, \text{normal}, \text{focused}\}$ associates a type to each arc, and $\epsilon : A \mapsto \mathcal{E}$ associates a neutral expression to each arc. 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.
- The graph is bipartite, i.e., for every $(u, v) \in A$, $p(u) \neq p(v)$.
- For every normal arc $a \in A$, $\epsilon(a)$ is guarded.
- For every atomic arc $a \in A$, $\epsilon(a)$ is an atom.
- The origins of the focused arcs all belong to the same player (and similarly for the ends, since the graph is bipartite).

Notice that the definition requires that no vertex be isolated (*i.e.*, without neighbors). A vertex with player σ is said to *belong* to σ or to be a σ -*vertex*.

Let us now relate neutral graphs to frontiers. Each vertex has an associated sequent, whose formulas are the translations of the neutral expressions labeling the arcs connected to it. The direction of an arc determines which translation (positive or negative) to consider, while its type (atomic, normal or focused) determines in which part of the sequent the formula occurs.

Definition 5.4 (Sequent associated with a vertex). *Let $G = (V, A, p, t, \epsilon)$ be a neutral graph and $v \in V$. Let*

- $\mathcal{L}^+ = \{[\epsilon(v, w)]^+ : (v, w) \in A, t(v, w) = \mathbf{atomic}\},$
- $\mathcal{L}^- = \{[\epsilon(u, v)]^- : (u, v) \in A, t(u, v) = \mathbf{atomic}\},$
- $\mathcal{U}^+ = \{[\epsilon(v, w)]^+ : (v, w) \in A, t(v, w) = \mathbf{normal}\},$
- $\mathcal{U}^- = \{[\epsilon(u, v)]^- : (u, v) \in A, t(u, v) = \mathbf{normal}\},$
- $\mathcal{F}^+ = \{[\epsilon(v, w)]^+ : (v, w) \in A, t(v, w) = \mathbf{focused}\},$
- $\mathcal{F}^- = \{[\epsilon(u, v)]^- : (u, v) \in A, t(u, v) = \mathbf{focused}\}.$

Since G is a neutral graph, at least one of \mathcal{F}^+ and \mathcal{F}^- is empty. The sequent $\Sigma_{G,v}$ associated with v is $\vdash \mathcal{L}^+, \mathcal{L}^-; \mathcal{U}^+; \mathcal{U}^- \uparrow \mathcal{F}^-$ if \mathcal{F}^+ is empty, and $\vdash \mathcal{L}^+, \mathcal{L}^-; \mathcal{U}^+; \mathcal{U}^- \Downarrow \mathcal{F}^+$ otherwise.

Definition 5.5 (Source). *A source of a neutral graph is a vertex which is not the end of a normal or focused arc, but is the origin of some normal arc.*

Sources play a significant role, because they are precisely the vertices associated with sequents of the form $\vdash \mathcal{L}; \mathcal{P}; \cdot$ with $\mathcal{P} \neq \emptyset$, *i.e.*, those which may appear as the conclusion of the $[D \Downarrow]$ rule, which marks the beginning of a synchronous phase.

5.5 Positions and moves

5.5.1 Positions

As in the previous games, we define positions and moves on two levels. A first level is made of *micro-positions* and *micro-moves* between them. A second level is made of *macro-positions* and *macro-moves* between them.

We first introduce a basic notion which will be used to define micro-positions and macro-positions.

Definition 5.6 (Position). *A position is a triple (G, f_0, f_1) where G is a neutral graph and f_0 and f_1 are Boolean values, satisfying the two following properties:*

- *if G is empty, then at least one of f_0 and f_1 is true;*
- *if there is some atomic arc (u, v) in G , then f_σ must be true, where σ is the player associated with u .*

Informally, the flag f_0 (resp. f_1) indicates whether player 0 (resp. 1) has failed and cannot win the play. If G is empty, the play ends, and at least one of the players must have failed. The origin of an atomic arc has an associated sequent which will only be obtained through an application of the daimon $[\boxtimes]$. Consequently, the corresponding player will have failed.

Definition 5.7 (Macro-position). *A macro-position is a position (G, f_0, f_1) such that G has no focused arc and all its sources (possibly zero) belong to the same player.*

Informally, a macro-move is seen as a synchronous phase by the player and as an asynchronous phase by the opponent. At a macro-position, there must be at most one player ready to start a synchronous phase, *i.e.*, ready to move.

Definition 5.8 (Final macro-position, σ -macro-position). *Let $P = (G, f_0, f_1)$ be a macro-position. If G has no source, then P is final. Otherwise, the sources of G belong to some player σ and P is a σ -macro-position.*

Definition 5.9 (Tie, σ -win). *Let $P = (G, f_0, f_1)$ be a final macro-position. Then at least one of f_0 and f_1 is true. If they are both true, P is a tie. Otherwise P is a σ -win, where $\sigma \in \{0, 1\}$ is such that f_σ is false.*

Proof. By definition 5.7, G has no focused arc. Since G is acyclic and has no sources, G has no normal arc. All the arcs of G are therefore atomic. Whether G is empty or not, one of f_0 and f_1 is true by definition 5.6. \square

The arena of the game consists of the macro-positions and the macro-moves, which will be defined later. Player σ plays at a σ -macro-position and wins at a σ -win.

Definition 5.10 (σ -micro-position). *A σ -micro-position is a position (G, f_0, f_1) such that all the origins of the focused arcs of G belong to player σ .*

Informally, a σ -micro-position is an intermediate step which may appear during player σ 's macro-moves. The origins of the focused arcs are the vertices associated to sequents in the middle of a synchronous phase.

Definition 5.11 (Playable σ -micro-position). *A σ -micro-position (G, f_0, f_1) is playable if G has at least one focused arc or one source belonging to player σ .*

Informally, a player continues her turn as long as some of her sequents are in the middle of a synchronous phase (origins of focused arcs) or are ready to start one (sources). Note that every σ -macro-position is a playable σ -micro-position. Informally, player σ can always play at a σ -macro-position.

5.5.2 Micro-moves

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

Transition	Sync reading	Async reading
$p \xrightarrow{D} p'$	$[D \Downarrow]$	$[D \Uparrow]$
$p \xrightarrow{R} p'$	$[R \Downarrow]$	$[R \Uparrow]$
$p \xrightarrow{\oplus} p'$	$[\oplus]$	$[\&]$
$p \xrightarrow{\otimes} p'$	$[\otimes]$	$[\wp]$
$p \xrightarrow{0} p'$	$[\boxtimes]$	$[\top]$
$p \xrightarrow{1} p'$	$[1]$ or $[\boxtimes]$	$[\perp]$ or $[\boxtimes]$
$p \xrightarrow{\text{at}} p'$	$[init]$ or $[\boxtimes]$	$[R \Uparrow \text{ atomic}]$ or $[\boxtimes]$

Table 1: Neutral moves and their two readings


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. Essentially, micro-moves are cut reduction rules for proof structures, with focalization. A notable difference is that cut reduction operates on readily available proof structures, while our micro-moves can be seen as an attempt to develop those structures as cut reduction goes. It is not always possible to develop them. As a result, failures may arise in some of these transitions; in that case the transition makes the relevant flags f_0 and/or f_1 true.

In the following description of the micro-moves we use figures to illustrate the formal definitions. Each micro-move rewrites a playable σ -micro-position $p = (G, f_0, f_1)$. σ - (resp. $\bar{\sigma}$ -) vertices are represented in black (resp. white). We also refer to player σ (resp. $\bar{\sigma}$) as the black (resp. white) player. Arcs are represented using the following convention:



To describe the transitions, let $G = (V, A, p, t, \epsilon)$.

Decision: Assume G has a σ -source v . Let F be a non empty subset of $\{(v, w) : (v, w) \in A\}$. If we then let $G' = (V, A, p, t', \epsilon)$, where t' is the same as t except that $t'(a) = \text{focused}$ if $a \in F$, we have the labeled transition $(G, f_0, f_1) \xrightarrow{D} (G', f_0, f_1)$.



Let us give an informal description of this transition. Recall the decision rules ($[D \Downarrow]$ and $[D \Uparrow]$ in Figure 5). $[D \Downarrow]$ is applied to a sequent of the form $\vdash \mathcal{L}; \mathcal{P}; \cdot$. In G , these sequents exactly correspond to the sources, and the


transition corresponds exactly to applying $[D \Downarrow]$ to one of them, while applying $[D \Uparrow]$ to some of its neighbors.

To describe the next five labeled transitions, assume G has a focused arc $a = (v, w)$. Since p is a σ -micro-position, v belongs to player σ and w belongs to player $\bar{\sigma}$.

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


$$G' = (V, (A \setminus \{a\}) \cup \{\bar{a}\}, p, t_{|A \setminus \{a\}} \cup \{(\bar{a}, \mathbf{normal})\}, \epsilon_{|A \setminus \{a\}} \cup \{(\bar{a}, E)\})$$

where $\bar{a} = (w, v)$ is the opposite arc to a . Then we have the transition $(G, f_0, f_1) \xrightarrow{R} (G', f_0, f_1)$.




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 $G' = (V, A, p, t, \epsilon')$ 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 $(G, f_0, f_1) \xrightarrow{+} (G', f_0, f_1)$.



This treatment of $+$ is essentially the same as in the additive game presented before. It corresponds exactly to the reduction of a cut link between two MALL proof structures with boxes with conclusions $[E_1]^+ \oplus [E_2]^+$ and $[E_1]^- \& [E_2]^-$: the black player's choice between E_1 and E_2 corresponds to that between $[E_1]^+$ and $[E_2]^+$.

If $\epsilon(a) = \mathbf{0}$ (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: $f_{\sigma}' = \top$ and $f_{\bar{\sigma}}' = f_{\bar{\sigma}}$. Then we have the labeled transition $(G, f_0, f_1) \xrightarrow{0} (G', f_0', f_1')$.




(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

$\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \top, \mathcal{F}$, in other words we apply $[\top]$; on the black player's side we face an unprovable sequent of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow 0, \mathcal{F}$ and we must apply $[\boxtimes]$. Consequently the black player fails ($f_{\sigma'} = \top$).

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 $G' = (V', A', p', t', \epsilon')$ where

- $V' = (V \setminus \{v\}) \uplus \{v_1, v_2\}$,
- $A' = \{(v_1, w), (v_2, w)\} \cup \{b' : b \in A \setminus \{a\}\}$,
- $p' = p|_{V \setminus \{v\}} \cup \{(v_1, \sigma), (v_2, \sigma)\}$,
- $t'(v_1, w) = t'(v_2, w) = \text{focused}$, and for every $b \in A \setminus \{a\}$, $t'(b') = t(b)$,
- $\epsilon'(v_1, w) = E_1$ and $\epsilon'(v_2, w) = E_2$, and for every $b \in A \setminus \{a\}$, $\epsilon'(b') = \epsilon(b)$,

We then have the labeled transition $(G, f_0, f_1) \xrightarrow{\boxtimes} (G', f_0', f_1')$.




On the black player's side, the splitting corresponds to that of the $[\otimes]$ rule. On the white player's side the invertible $[\wp]$ rule is applied. Here again, this transition is exactly the reduction of a cut link between two proof structures with conclusions $[E_1]^+ \otimes [E_2]^+$ and $[E_1]^- \wp [E_2]^-$.

If $\epsilon(a) = \mathbb{1}$ (the 0-ary multiplicative), then one can remove a . Formally, let $G' = G|_{A \setminus \{a\}}$ and f_0' and f_1' be Boolean values defined as follows:

$$f_{\sigma'} = \begin{cases} \top & \text{if } v \text{ is a vertex of } G' \\ f_{\sigma} & \text{otherwise} \end{cases} \quad f_{\bar{\sigma}'} = \begin{cases} f_{\bar{\sigma}} & \text{if } w \text{ is a vertex of } G' \\ \top & \text{otherwise} \end{cases}$$

Then we have the labeled transition $(G, f_0, f_1) \xrightarrow{\mathbb{1}} (G', f_0', f_1')$.




(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 $\vdash \cdot; \cdot; \cdot \Downarrow 1$, therefore the player fails ($f_{\sigma'} = \top$) if 1 is not the only formula of the sequent. On the white player's side $\lceil \perp \rceil$ is applied, and if w is only connected to v then its associated sequent becomes $\vdash \cdot; \cdot; \cdot \Uparrow \perp$ which is unprovable, and the player fails ($f_{\sigma'} = \top$).


If $\epsilon(a)$ is an atom k , then the sequent associated to v is of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow K, \mathcal{F}$. There are two cases, depending on whether the rule $[init]$ may be applied or not.

First case: if there is exactly one arc b connected to v beside a , and b is of the form (u, v) , with $t(b) = \mathbf{atomic}$ and $\epsilon(b) = k$, then one may remove v , a and b . Formally, let $G' = G_{|A \setminus \{a, b\}}$. Then we have the labeled transition $(G, f_0, f_1) \xrightarrow{\text{at}} (G', f_0, f_1)$.



(in the second graph, any isolated vertex shall be removed.) In this case the sequent associated to v is $\vdash K^\perp; \cdot; \cdot \Downarrow K$ and $[init]$ may be applied, which is reflected by the transition. On the white player's side the transition corresponds to applying $\lceil \boxtimes \rceil$ to the sequents associated with u and w . This can be done safely as the white player has already failed ($f_{\sigma} = \top$) since u is the origin of an atomic arc.

Second case: if the first case does not apply, then one may make a atomic and make the player fail in the process. Formally, let $G' = (V, A, p, t', \epsilon)$, where t' is the same as t except that $t'(a) = \mathbf{atomic}$, and f_0' and f_1' be Boolean values defined as follows: $f_{\sigma'} = \top$ and $f_{\bar{\sigma}'} = f_{\bar{\sigma}}$. Then we have the labeled transition $(G, f_0, f_1) \xrightarrow{\text{at}} (G', f_0, f_1)$.



In this case the sequent associated to v is not $\vdash K^\perp; \cdot; \cdot \Downarrow K$, $[init]$ may not be applied, and the player applies $\lceil \boxtimes \rceil$ and fails. On the white player's side the transition corresponds to applying $[R \Uparrow \mathbf{atomic}]$ to the sequent associated with w .

Definition 5.12 (σ -micro-move). *Let p, p' be σ -micro-positions. There is a σ -micro-move from p to p' (notation $p \mapsto_{\sigma} p'$) iff one of the following holds: $p \xrightarrow{D} p'$, $p \xrightarrow{R} p'$, $p \xrightarrow{\perp} p'$, $p \xrightarrow{0} p'$, $p \xrightarrow{\times} p'$, $p \xrightarrow{1} p'$, or $p \xrightarrow{\text{at}} p'$.*

Proposition 5.13. *There is a σ -micro-move from a σ -micro-position iff it is playable.*

Proof. A σ -micro-position is playable iff its neutral graph has a source or a focused arc. A σ -micro-move \xrightarrow{D} is possible iff the neutral graph has a source. A σ -micro-move \xrightarrow{R} , $\xrightarrow{+}$, $\xrightarrow{0}$, $\xrightarrow{\times}$, $\xrightarrow{1}$, or $\xrightarrow{\text{at}}$ is possible iff the neutral graph has a focused arc (those moves cover all the cases for the neutral expression labeling the arc). \square

5.5.3 Macro-moves

We proceed to define the macro-moves, which are the actual moves of the game, as maximal sequences of micro-moves.

Proposition 5.14. *The length of the sequences of micro-moves starting in a fixed micro-position is bounded.*

Proof. Associate a triple (s, n, f) with each micro-position (G, f_0, f_1) , where s is the total number of symbols of the neutral expressions labeling the arcs of G , n is the number of normal arcs of G , and f is the number of focused arcs of G . Every micro-move decreases this triple for the lexicographical ordering. Moreover, these triples verify $n \leq s$ and $f \leq s$. \square

Definition 5.15 (Macro-move). *Let p be a σ -macro-position and $p \mapsto_{\sigma}^* p'$ a maximal sequence of σ -micro-moves from p . Then p' is either a $\bar{\sigma}$ -macro-position or a final macro-position. We say that there is a macro-move from p to p' and denote it by $p \rho p'$.*

Proof. p' is a σ -micro-position, but it is not playable by Proposition 5.13. p' has no focused arcs and all its sources belong to player $\bar{\sigma}$, which makes it a $\bar{\sigma}$ -macro-position or a final macro-position. \square

Note that there is a macro-move from every non-final macro-position.

Proposition 5.16. *The length of the plays starting in a fixed macro-position is bounded.*

Proof. Every macro-move expands to a sequence of micro-moves which is non-empty since every non-final σ -macro-position is a playable σ -micro-position. The result thus follows from Proposition 5.14. \square

5.6 Winning strategies as cut-free 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. Our proofs effectively show how to construct a winning strategy from a proof. They also show how to construct a proof from a winning strategy, but this construction is not unique in general. To recover unicity, thus having a one-to-one correspondence between winning strategies and proofs, 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. 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 a central notion relating concepts of the game to concepts of the proof system: that of σ -provability (for $\sigma \in \{0, 1\}$).

Definition 5.17 (σ -provability). *Let G be a neutral graph 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 is a neutral graph and f_0, f_1 are Boolean values is σ -provable iff $f_\sigma = \perp$ 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).

Proposition 5.18. *Let P be a playable σ -micro-position. Let $S = \{P' : P \mapsto_\sigma P'\}$. P is σ -provable iff there exists $P' \in S$ which is σ -provable.*

Proof. Let us write $P = (G, f_0, f_1)$ and $G = (V, A, p, t, \epsilon)$. We prove this result in two parts: (1) the “if” part, (2) the “only if” part.

(1) Suppose that there exists $P' = (G', f_0', f_1') \in S$ which is σ -provable. Let us show that P is σ -provable. We have $f_{\sigma'} = \perp$ and G' is σ -provable. Let us write $G' = (V', A', p', t', \epsilon')$. We examine the cases for $P \mapsto_\sigma P'$.

Case $P \xrightarrow{D} P'$. Thus $f_\sigma = f_{\sigma'} = \perp$ and, moreover, the only σ -sequent affected by the move is the one associated with the source v selected for the move. In G' , this sequent is of the form $\vdash \mathcal{L}; \mathcal{P}_1; \cdot \Downarrow \mathcal{P}_2$, and in G it is $\vdash \mathcal{L}; \mathcal{P}_1, \mathcal{P}_2; \cdot$. The result follows from the derivation

$$\frac{\vdash \mathcal{L}; \mathcal{P}_1; \cdot \Downarrow \mathcal{P}_2}{\vdash \mathcal{L}; \mathcal{P}_1, \mathcal{P}_2; \cdot} [D \Downarrow]$$

In each one of the other cases we consider the focused arc $a = (v, w)$ from the definition of the corresponding transition. First of all, the case $P \xrightarrow{0} P'$ does not happen, since $f_{\sigma'} = \perp$. In the other cases v is the only σ -vertex of G affected by the transition, hence all we need to show is that $f_\sigma = \perp$ and $\Sigma_{G,v}$ is provable. For the latter, showing that $\Sigma_{G,v}$ derives from σ -sequents of G' (which are provable) is enough.

Case $P \xrightarrow{R} P'$. Thus $f_\sigma = f_{\sigma'} = \perp$ and, moreover, $\Sigma_{G,v}$ is of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow \phi, \mathcal{F}$, where $\phi = [\epsilon(a)]^+$. Remark that $[\epsilon'(w, v)]^- = \phi$ and that ϕ is negative; it is then clear that $\Sigma_{G',v} = \vdash \mathcal{L}; \mathcal{P}; \phi, \mathcal{N} \Downarrow \mathcal{F}$ and the result follows from the derivation

$$\frac{\vdash \mathcal{L}; \mathcal{P}; \phi, \mathcal{N} \Downarrow \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow \phi, \mathcal{F}} [R \Downarrow]$$

Case $P \xrightarrow{\pm} P'$. Thus $f_\sigma = f_{\sigma'} = \perp$ and, moreover, $[\epsilon(a)]^+$ is of the form $\phi_1 \oplus \phi_2$. $\Sigma_{G,v}$ is of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow \phi_1 \oplus \phi_2, \mathcal{F}$ and $\Sigma_{G',v} = \vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow$

ϕ_i, \mathcal{F} , for some $i \in \{1, 2\}$. The result follows from the derivation

$$\frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow \phi_i, \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow \phi_1 \oplus \phi_2, \mathcal{F}} [\oplus_i]$$

Case $P \overset{\times}{\mapsto} P'$. Thus, $f_\sigma = f_{\sigma'} = \perp$ and, moreover, $[\epsilon(a)]^+$ is of the form $\phi_1 \otimes \phi_2$. The vertex v is replaced in G' with two σ -vertices v_1 and v_2 . Σ_{G', v_1} and Σ_{G', v_2} are of the forms $\vdash \mathcal{L}_1; \mathcal{P}_1; \mathcal{N}_1 \Downarrow \phi_1, \mathcal{F}_1$ and $\vdash \mathcal{L}_2; \mathcal{P}_2; \mathcal{N}_2 \Downarrow \phi_2, \mathcal{F}_2$, and we have $\Sigma_{G, v} = \vdash \mathcal{L}_1, \mathcal{L}_2; \mathcal{P}_1, \mathcal{P}_2; \mathcal{N}_1, \mathcal{N}_2 \Downarrow \phi_1 \otimes \phi_2, \mathcal{F}_1, \mathcal{F}_2$. The result follows from the derivation

$$\frac{\vdash \mathcal{L}_1; \mathcal{P}_1; \mathcal{N}_1 \Downarrow \phi_1, \mathcal{F}_1 \quad \vdash \mathcal{L}_2; \mathcal{P}_2; \mathcal{N}_2 \Downarrow \phi_2, \mathcal{F}_2}{\vdash \mathcal{L}_1, \mathcal{L}_2; \mathcal{P}_1, \mathcal{P}_2; \mathcal{N}_1, \mathcal{N}_2 \Downarrow \phi_1 \otimes \phi_2, \mathcal{F}_1, \mathcal{F}_2} [\otimes]$$

Case $P \overset{1}{\mapsto} P'$. Thus, $[\epsilon(a)]^+ = 1$. Since $f_{\sigma'} = \perp$, $f_\sigma = \perp$ and v is not a vertex of G' . a is therefore the only arc connected to it in G and $\Sigma_{G, v} = \vdash \cdot; \cdot \Downarrow 1$ is provable:

$$\overline{\vdash \cdot; \cdot \Downarrow 1} [1]$$

Case $P \overset{\text{at}}{\mapsto} P'$. Thus, $[\epsilon(a)]^+ = K$. Since $f_{\sigma'} = \perp$, the first case in the definition of $\overset{\text{at}}{\mapsto}$ applies, $f_\sigma = f_{\sigma'} = \perp$ and $\Sigma_{G, v} = \vdash K^\perp; \cdot; \cdot \Downarrow K$ is provable:

$$\overline{\vdash K^\perp; \cdot; \cdot \Downarrow K} [\text{init}]$$

(2) This is the converse to the previous part. Suppose that $P = (G, f_0, f_1)$ is σ -provable. We have $f_\sigma = \perp$ and G is σ -provable. We will consider several cases, and in each one we will show that there is a σ -provable $P' = (G', f_0', f_1') \in S$. There will be few σ -vertices affected by the transition $P \mapsto_\sigma P'$, hence it will be enough to show that $f_{\sigma'} = \perp$ and that the sequents associated to those vertices are provable. Since P is a playable σ -micro-position, G has a source (belonging to player σ) or a focused arc (whose origin belongs to player σ).

First case. Suppose that G has a σ -source v . $\Sigma_{G, v}$ is of the form $\vdash \mathcal{L}; \mathcal{P}; \cdot$ with $\mathcal{P} \neq \emptyset$ and it is provable by assumption, and a proof must end with the $[D \Downarrow]$ rule:

$$\frac{\vdash \mathcal{L}; \mathcal{P}_1; \cdot \Downarrow \mathcal{P}_2}{\vdash \mathcal{L}; \mathcal{P}_1, \mathcal{P}_2; \cdot} [D \Downarrow]$$

where \mathcal{P}_1 and \mathcal{P}_2 partition \mathcal{P} , \mathcal{P}_2 is not empty and $\vdash \mathcal{L}; \mathcal{P}_1; \cdot \Downarrow \mathcal{P}_2$ is provable. This corresponds to a transition $P \overset{D}{\mapsto} P' = (G', f_0', f_1')$ in which the outgoing normal arcs of v corresponding to \mathcal{P}_2 become focused. The only affected σ -vertex is v and $\Sigma_{G', v}$ is precisely $\vdash \mathcal{L}; \mathcal{P}_1; \cdot \Downarrow \mathcal{P}_2$. Moreover $f_{\sigma'} = f_\sigma = \perp$, therefore P' is σ -provable.

Second case. Suppose that $G = (V, A, p, t, \epsilon)$ has a focused arc, and let v be its origin. $\Sigma_{G, v}$ is provable. Consider the last rule R of such a proof. $\Sigma_{G, v}$ is of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \Downarrow \phi, \mathcal{F}$, where ϕ is the principal formula of R . There is a focused arc $a = (v, w)$ for some w such that $[\epsilon(a)]^+ = \phi$. We are

going to consider all the cases for $\epsilon(a)$. In each case we show that there is $P' = (G', f_0', f_1') \in S$ which is σ -provable. More specifically, we show that $f_{\sigma'} = \perp$ and that the sequents associated to the σ -vertices of G' affected by the transition are in fact the premises of R , and are therefore provable.

Case $\epsilon(a) = \downarrow E$. Thus ϕ is negative and R is $[R \Downarrow]$. As we did in part (1), we match this rule with the transition $P \xrightarrow{R} (G', f_0, f_1)$ in which a is reversed. The only affected σ -vertex is v . Since $\phi = [E]^-$, $\Sigma_{G',v}$ is precisely the premise of R .

Case $\epsilon(a) = E_1 + E_2$. Thus $\phi = [E_1]^+ \oplus [E_2]^+$ and R is $[\oplus_i]$ for some $i \in \{1, 2\}$. As we did in part (1), we match this rule with the transition $P \xrightarrow{\oplus} (G', f_0, f_1)$ in which $\epsilon(a)$ is replaced with E_i . The only affected σ -vertex is v , and $\Sigma_{G',v}$ is precisely the premise of R .

Case $\epsilon(a) = \mathbf{0}$. Thus, $\phi = 0$. This case does not happen, since there is no introduction rule for 0.

Case $\epsilon(a) = E_1 \times E_2$. Thus $\phi = [E_1]^+ \otimes [E_2]^+$ and R is $[\otimes]$. As we did in part (1), we match this rule with a transition $P \xrightarrow{\times} (G', f_0, f_1)$. The only affected σ -vertex is v which is split into v_1 and v_2 , and Σ_{G',v_1} and Σ_{G',v_2} are precisely the premises of R .

Case $\epsilon(a) = \mathbf{1}$. Thus $\phi = 1$ and R is $[1]$. It means that $\Sigma_{G,v} = \vdash \cdot; \cdot \Downarrow 1$. As we did in part (1), we match this rule with the transition $P \xrightarrow{\mathbf{1}} (G', f_0', f_1')$ in which a is removed, and it is clear that v is not a vertex of G' , therefore $f_{\sigma'} = f_{\sigma} = \perp$ as needed. There are no new/affected σ -vertices in G' .

Case $\epsilon(a) = k$. Thus $\phi = K$ and R is $[init]$. It means that $\Sigma_{G,v} = \vdash K^\perp; \cdot; \cdot \Downarrow K$. As we did in part (1), we match this rule with the transition $P \xrightarrow{\text{at}} (G', f_0', f_1')$ in which v is removed (first case of the definition of $\xrightarrow{\text{at}}$). Then $f_{\sigma'} = f_{\sigma} = \perp$ as needed. There are no new/affected σ -vertices in G' . \square

Proposition 5.19. *Let P be a playable σ -micro-position. Let $S = \{P' : P \mapsto_{\sigma} P'\}$. P is $\bar{\sigma}$ -provable iff every $P' \in S$ is $\bar{\sigma}$ -provable.*

Proof. Let us write $P = (G, f_0, f_1)$ and $G = (V, A, p, t, \epsilon)$. We are going to prove both directions simultaneously. Let us give names to the two hypotheses: (A) P is $\bar{\sigma}$ -provable, and (B) every $P' \in S$ is $\bar{\sigma}$ -provable. Since P is playable, S is not empty. Let $P' = (G', f_0', f_1') \in S$. We examine the cases for $P \mapsto_{\sigma} P'$. In each case we show that (A) implies that P' is $\bar{\sigma}$ -provable, and that (B) implies (A). This will prove the proposition, since we chose $P' \in S$ arbitrarily and we cover all the cases for $P \mapsto_{\sigma} P'$. Moreover, it will be enough to consider only the $\bar{\sigma}$ -vertices affected by the transition.

Case $P \xrightarrow{D} P'$. Let v be the source of G which is selected for the move. Every $\bar{\sigma}$ -vertex w affected by the move is such that the arc (v, w) is normal in G and becomes focused in the transition. In other words, $\Sigma_{G,w}$ and $\Sigma_{G',w}$ are of the form $\vdash \mathcal{L}; \mathcal{P}; \phi, \mathcal{N} \uparrow \mathcal{F}$ and $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi, \mathcal{F}$. The derivation

$$\frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi, \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \phi, \mathcal{N} \uparrow \mathcal{F}} [D \uparrow]$$

shows that if (B), then each $\Sigma_{G',w}$ has a proof, hence so does each $\Sigma_{G,w}$; moreover $f_{\bar{\sigma}}' = \perp$, hence $f_{\bar{\sigma}} = \perp$ and (A). Conversely, if (A), then each $\Sigma_{G,w}$ has a proof, which is equivalent to a proof ending with the above derivation, hence each $\Sigma_{G',w}$ is provable; moreover $f_{\bar{\sigma}} = \perp$, hence $f_{\bar{\sigma}}' = \perp$ and P' is $\bar{\sigma}$ -provable.

For each one of the other cases we consider the arc $a = (v, w)$ from the definition of the corresponding transition.

Case $P \xrightarrow{\text{at}} P'$. There are two cases in the definition of this transition. In the first case, we must have $f_{\bar{\sigma}} = \top$ since the $\bar{\sigma}$ -vertex u is the origin of the atomic arc (u, v) . Then $f_{\bar{\sigma}}' = \top$ and (A) and (B) are both false, which concludes this case. In the second case, the only affected $\bar{\sigma}$ -vertex is w , and $\Sigma_{G,w}$ and $\Sigma_{G',w}$ are of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow K^\perp, \mathcal{F}$ and $\vdash K^\perp, \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}$. The derivation

$$\frac{\vdash K^\perp, \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow K^\perp, \mathcal{F}} [R \uparrow \text{atomic}]$$

shows that if (B), then $\Sigma_{G',w}$ has a proof, hence so does $\Sigma_{G,w}$; moreover $f_{\bar{\sigma}}' = \perp$, hence $f_{\bar{\sigma}} = \perp$ and (A). Conversely, if (A), then $\Sigma_{G,w}$ has a proof, which is equivalent to a proof ending with the above derivation, hence $\Sigma_{G',w}$ is provable; moreover $f_{\bar{\sigma}} = \perp$, hence $f_{\bar{\sigma}}' = \perp$ and P' is $\bar{\sigma}$ -provable.

In all the remaining cases the only $\bar{\sigma}$ -vertex affected by the transition is w .

Case $P \xrightarrow{R} P'$. $\epsilon(a)$ is of the form $\Downarrow E$. $\Sigma_{G,w}$ and $\Sigma_{G',w}$ are of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi, \mathcal{F}$ and $\vdash \mathcal{L}; \phi, \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}$ with $\phi = [\Downarrow E]^- = [E]^+$. The derivation

$$\frac{\vdash \mathcal{L}; \phi, \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi, \mathcal{F}} [R \uparrow]$$

allows us to conclude as in the previous cases.

Case $P \xrightarrow{\pm} P'$. $\epsilon(a)$ is of the form $E_1 + E_2$ and a choice is made between E_1 and E_2 in P' . Consider both choices. They lead to two elements of S , (G_1, f_0, f_1) and (G_2, f_0, f_1) (one of them is P'). Let us write $G_i = (V', A', p', t', \epsilon'_i)$ for each $i \in \{1, 2\}$. $[\epsilon(a)]^-$ is of the form $\phi_1 \& \phi_2$ and $[\epsilon'_i(a)]^- = \phi_i$. $\Sigma_{G,w}$ is of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi_1 \& \phi_2, \mathcal{F}$ and $\Sigma_{G_i,w} = \vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi_i, \mathcal{F}$. The derivation

$$\frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi_1, \mathcal{F} \quad \vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi_2, \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi_1 \& \phi_2, \mathcal{F}} [\&]$$

shows that if (B), then $\Sigma_{G_1,w}$ and $\Sigma_{G_2,w}$ have proofs, hence so does $\Sigma_{G,w}$; moreover $f_{\bar{\sigma}} = \perp$, hence (A). Conversely, if (A), then $\Sigma_{G,w}$ has a proof, which is equivalent to a proof ending with the above derivation, hence $\Sigma_{G_1,w}$ and $\Sigma_{G_2,w}$ are provable; moreover $f_{\bar{\sigma}} = \perp$, hence P' is $\bar{\sigma}$ -provable.

Case $P \xrightarrow{0} P'$. $\Sigma_{G,w}$ is of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \top, \mathcal{F}$ and it is provable:

$$\frac{}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \top, \mathcal{F}} [\top]$$

w is not a vertex of G' and $f_{\bar{\sigma}}' = f_{\bar{\sigma}}$, hence P is $\bar{\sigma}$ -provable iff P' is. If (B), then (A). Conversely, if (A), then P' is $\bar{\sigma}$ -provable.

Case $P \overset{\times}{\mapsto} P'$. $\Sigma_{G,w}$ and $\Sigma_{G',w}$ are of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi_1 \wp \phi_2, \mathcal{F}$ and $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi_1, \phi_2, \mathcal{F}$. The derivation

$$\frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi_1, \phi_2, \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \phi_1 \wp \phi_2, \mathcal{F}} [\wp]$$

allows us to conclude as in the previous cases.

Case $P \overset{\perp}{\mapsto} P'$. $\Sigma_{G,w}$ is of the form $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \perp, \mathcal{F}$. Consider the derivation

$$\frac{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}}{\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \perp, \mathcal{F}} [\perp]$$

If (B), then $f_{\bar{\sigma}'} = \perp$ and w is a vertex of G' . Then $\Sigma_{G',w} \Vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}$, it has a proof, hence so does $\Sigma_{G,w}$ by the above derivation; moreover $f_{\bar{\sigma}} = \perp$, hence (A). Conversely, if (A), then $f_{\bar{\sigma}} = \perp$ and $\Sigma_{G,w}$ has a proof, which is equivalent to a proof ending with the above derivation, hence $\vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}$ is provable; $\mathcal{L}, \mathcal{P}, \mathcal{N}$ and \mathcal{F} cannot all be empty, therefore w is a vertex of G' , and $\Sigma_{G',w} \Vdash \mathcal{L}; \mathcal{P}; \mathcal{N} \uparrow \mathcal{F}$ is provable; moreover $f_{\bar{\sigma}'} = \perp$, hence (A). \square

Lemma 5.20. *Let P be a σ -macro-position. Let $S = \{P' : P \rho P'\}$. P is σ -provable iff there exists $P' \in S$ which is σ -provable. P is $\bar{\sigma}$ -provable iff every $P' \in S$ is $\bar{\sigma}$ -provable.*

Proof. We show a more general result. For every σ -micro-position P , let $S_P = \{P' : P \mapsto_{\sigma}^* P' \text{ and this sequence is maximal}\}$. Let us show that for every σ -micro-position P , P is σ -provable iff there exists $P' \in S_P$ which is σ -provable, and P is $\bar{\sigma}$ -provable iff every $P' \in S_P$ is $\bar{\sigma}$ -provable. By proposition 5.14, we may show it by induction on the maximal length l_P of the sequences $P \mapsto_{\sigma}^* P'$ for $P' \in S_P$. If $l_P = 0$, then $S = \{P\}$ and the result is trivial. Now suppose that $l_P > 0$ and let $S' = \{P' : P \mapsto_{\sigma} P'\}$. By Propositions 5.18 and 5.19, P is σ -provable iff there exists $P' \in S'$ which is σ -provable, and P is $\bar{\sigma}$ -provable iff every $P' \in S'$ is $\bar{\sigma}$ -provable. For every $P' \in S'$, $l_{P'} < l_P$ and the induction hypothesis applies. The result follows from the fact that $S_P = \bigcup_{P' \in S'} S_{P'}$. \square

Theorem 5.21. *Let P be a position and $\sigma \in \{0, 1\}$. There is a winning σ -strategy from P iff P is σ -provable.*

Proof. We prove the result by induction on the maximal length l_P of the plays starting in P (see Proposition 5.16). If $l_P = 0$, P is final and there is a winning σ -strategy from P iff P is a win, iff P is σ -provable. Suppose that $l_P > 0$. Let $S = \{P' : P \rho P'\}$. There are two cases: P is either a σ -macro-position or a $\bar{\sigma}$ -macro-position.

If P is a σ -macro-position, then there is a winning σ -strategy from P iff there is a winning σ -strategy from some $P' \in S$, iff, by induction hypothesis ($l_{P'} < l_P$), there exists $P' \in S$ which is σ -provable, iff, by lemma 5.20, P is σ -provable.

If P is a $\bar{\sigma}$ -macro-position, then there is a winning σ -strategy from P iff there is a winning σ -strategy from every $P' \in S$, iff, by induction hypothesis ($l_{P'} < l_P$), every $P' \in S$ is σ -provable, iff, by lemma 5.20, P is σ -provable. \square

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 [5] introduced a game semantics for linear logic along these lines: this approach was extended by Abramsky and Jagadeesan [1], Abramsky and Melliès [3], and Hyland and Ong [20].

The games described here, however, deal with the question about what is actually achieved when a particular inference rule is used in a proof. In that sense, our work is more closely related to the “dialog games” inspired by the early work of Lorenzen [26] (see also [18]).

It is natural to consider extending the games described here for more aspects of logic, such as the exponentials (! and ?) of linear logic as well as first-order and second-order quantification. The earlier paper [28] consider adding first-order quantification, the equality predicate on terms, and least and greatest fixed points to MALL. That paper also contained numerous examples that exploited such extensions to neutral strategies and to (simple) games: it was shown that a natural approach to fixed points immediately lead to viewing the usual logical specification of bisimulation as the usual game theoretic description of bisimulation [33].

Developing technical connections to Girard’s *Ludics* [16] and Faggian and Hyland’s treatment of parts of Ludics [10] would be of particular interest. Relating proofs and strategies more closely would be another worthy development: here we will probably need to switch to a framework similar to asynchronous games and innocent strategies [27].

Finally, Lakatos [22] famously describes the value of a dialog between proof and refutation for uncovering mathematical truths. It would be interesting to see if our work here might be seen as a modest start on a logician’s formalization of his ideas.

7 Conclusion

We have presented a two-person game in which one player is attempting to build a (cut-free) proof of a formula while the other player, who is working with exactly the same set of rules, is attempting to build a refutation of that formula. In order to capture the sameness of the two players, the game makes use of a neutral expression language that can be mapped into logical expressions using two, dual mappings. Restricting the interaction between multiplicatives of opposite polarities yields a simple notion of game with considerable expressiveness. Removing that restriction requires a treatment of the multiplicatives using a graph structure for game “positions” which essentially represents cut links between two dual proof structures and can also be mapped to sequents using two, dual mappings. The computational content of the game can be seen

as cut reduction and simultaneous dual proof search at the same time. The most natural proof system that results from this game-theoretic consideration is a focusing proof system that allows for multifocusing.

Acknowledgements. We would like to thank Claudia Faggian and the reviewers of an earlier version of this paper for their comments and suggestions for improving the presentation of this work.

References

- [1] Samson Abramsky and Radha Jagadeesan. Games and full completeness for multiplicative linear logic. *J. of Symbolic Logic*, 59(2):543–574, 1994.
- [2] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Information and Computation*, 163(2):409–470, 2000.
- [3] Samson Abramsky and Paul-André Melliès. Concurrent games and full completeness. In *14th Symp. on Logic in Computer Science*, pages 431–442. IEEE Computer Society Press, 1999.
- [4] Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *J. of Logic and Computation*, 2(3):297–347, 1992.
- [5] 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.
- [6] Kaustuv Chaudhuri, Dale Miller, and Alexis Saurin. Canonical sequent proofs via multi-focusing. In Giorgio Ausiello, Juhani Karhumäki, Giancarlo Mauri, and Luke Ong, editors, *Fifth IFIP International Conference on Theoretical Computer Science*, volume 273 of *IFIP International Federation for Information Processing*, pages 383–396. Boston: Springer, September 2008.
- [7] Roberto Di Cosmo, Jean-Vincent Loddo, and Stephane Nicolet. A game semantics foundation for logic programming (extended abstract). In *PLILP/ALP*, pages 355–373, 1998.
- [8] V. Danos and L. Regnier. The structure of multiplicatives. *Archive for Mathematical Logic*, 28:181–203, 1989.
- [9] Olivier Deland and Dale Miller. A neutral approach to proof and refutation in MALL. In F. Pfenning, editor, *23th Symp. on Logic in Computer Science*, pages 498–508. IEEE Computer Society Press, 2008.
- [10] Claudia Faggian and Martin Hyland. Designs, disputes and strategies. In *Computer Science Logic*, volume 2471 of *Lecture Notes in Computer Science*, pages 713–748, 2002.

- [11] Chryside Galanaki, Panos Rondogiannis, and William W. Wadge. An infinite-game semantics for well-founded negation in logic programming. *Annals of Pure and Applied Logic*, 151(2):70–88, 2008.
- [12] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- [13] Jean-Yves Girard. A fixpoint theorem in linear logic. An email posting to the mailing list linear@cs.stanford.edu, February 1992.
- [14] Jean-Yves Girard. Proof-nets: the parallel syntax for proof-theory. In Aldo Ursini and Paolo Agliano, editors, *Logic and Algebra*, volume 180 of *Lecture Notes In Pure and Applied Mathematics*, pages 97–124, New York, 1996. Marcel Dekker.
- [15] Jean-Yves Girard. On the meaning of logical rules I: syntax vs. semantics. In Berger and Schwichtenberg, editors, *Computational Logic*, pages 215–272. SV, Heidelberg, 1999.
- [16] Jean-Yves Girard. Locus solum: From the rules of logic to the logic of rules. *Mathematical Structures in Computer Science*, 11(3):301–506, June 2001.
- [17] Lars Hallnäs and Peter Schroeder-Heister. A proof-theoretic approach to logic programming. II. Programs as definitions. *J. of Logic and Computation*, 1(5):635–660, October 1991.
- [18] Hugo Herbelin. *Séquents qu'on calcule: de l'interprétation du calcul des séquents comme calcul de lambda-termes et comme calcul de stratégies gagnantes*. PhD thesis, Université Paris 7, 1995.
- [19] Wilfrid Hodges. Logic and games. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Stanford University, 2004.
- [20] 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.
- [21] Max I. Kanovich. Simulating linear logic with 1-linear logic. Preprint 94-02, Laboratoire de Mathématiques Discrètes, University of Marseille, 1994.
- [22] Imre Lakatos. *Proofs and Refutations*. Cambridge University Press, 1976.
- [23] P. Lincoln, J. Mitchell, A. Scedrov, and N. Shankar. Decision problems for propositional linear logic. *Annals Pure Applied Logic*, 56:239–311, 1992.
- [24] Jean-Vincent Loddo. *Généralisation des Jeux Combinatoires et Applications aux Langages Logiques*. PhD thesis, Université Paris VII, 2002.

- [25] Jean-Vincent Loddo and Roberto Di Cosmo. Playing logic programs with the alpha-beta algorithm. In *LPAR*, pages 207–224, 2000.
- [26] Paul Lorenzen. Ein dialogisches konstruktivitätskriterium. In *Infinitistic Methods: Proceed. Symp. Foundations of Math.*, pages 193–200. PWN, 1961.
- [27] Paul-André Melliès. Asynchronous games 2: The true concurrency of innocence. *Theoretical Computer Science*, 358(2–3):200–228, 2004.
- [28] Dale Miller and Alexis 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.
- [29] David Pym and Eike Ritter. *Reductive Logic and Proof-search: proof theory, semantics, and control*, volume 45. Oxford Logic Guides, Oxford, 2004.
- [30] Alexis Saurin. Towards ludics programming: Interactive proof search. In *Logic Programming, 24th International Conference*, volume 5366 of *LNCS*, pages 253–268, December 2008.
- [31] Alexis Saurin. *Une étude logique du contrôle (appliquée à la programmation fonctionnelle et logique)*. PhD thesis, Ecole Polytechnique, September 2008.
- [32] Raymond M. Smullyan. *First-Order Logic*. Dover Publications, 1995. Second edition.
- [33] Colin Stirling. Bisimulation, modal logic and model checking games. *Logic Journal of the IGPL*, 7(1):103–124, 1999.
- [34] Maarten H. van Emden. Quantitative deduction and its fixpoint theory. *Journal of Logic Programming*, 3(1):37–53, 1986.