A compositional account of Herbrand's theorem via concurrent games

Aurore Alcolei[1], Pierre Clairambault[1], Martin Hyland[2], and Glynn Winskel[3]

[1] Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP, Lyon, France
[2] Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Cambridge, U.K.
[3] Computer Laboratory, University of Cambridge, Cambridge, U.K.

Abstract

Herbrand's theorem, widely regarded as a cornerstone of proof theory, exposes some of the constructive content of classical logic. In particular, it gives a reduction of first-order validity to propositional validity, by understanding the structure of the assignment of first-order terms to existential quantifiers, and the causal dependency between quantifiers.

Starting with some game semantics intuition already present in the literature, we show how Herbrand's theorem in its general form can be elegantly stated as a theorem in the framework of concurrent games. In this framework, the causal structure of concurrent strategies, paired with annotations by first-order terms, is used to specify the dependencies between quantifiers. Concurrent strategies being composable, we achieve this by interpreting classical sequent proofs within this denotational model. This yields a compositional proof of Herbrand's theorem.

1 Introduction

Classical first order logic is known to be non-constructive. For example, the formula

$$\exists \mathsf{x}\,(\mathsf{P}(\mathsf{x}) \quad\Longrightarrow\quad \mathsf{P}(\mathsf{f}(\mathsf{x}))) \tag{1}$$

is valid (provided the language has some constant symbol c), but there is no first-order term t such that $P(t) \implies P(f(t))$ holds.

In his thesis however, Herbrand proved that although no single closed term can serve as a witness of a formula $\exists x\, \varphi(x)$ (with φ quantifier-free), there always exists finitely many terms $t_1, \ldots, t_n$ such that $\varphi(t_1) \vee \cdots \vee \varphi(t_n)$ holds. The extraction of such t_is is widely regarded as an early account of the computational content of classical proofs. In the example above, one may extract from a proof of (1) the following valid disjunction

$$(\mathsf{P}(\mathsf{c}) \implies \mathsf{P}(\mathsf{f}(\mathsf{c}))) \vee (\mathsf{P}(\mathsf{f}(\mathsf{c})) \implies \mathsf{P}(\mathsf{f}(\mathsf{f}(\mathsf{c}))))$$

where the existential quantifier has been instantiated with two witnesses c and $\mathsf{f}(\mathsf{c})$. We call the disjunction above a Herbrand disjunction and c and $\mathsf{f}(\mathsf{c})$ its witnesses.

In its general form, Herbrand's theorem relates the validity of any first order formula to the validity of such a finite quantifier-free formula where existentially quantified variables have been replaced with finite disjunctions featuring finitely many terms, the witnesses. For an arbitrary formula the data of the witnesses for possibly deep existential quantifiers can be presented as certain trees, Miller's *expansion trees* [15]; which follow the syntax tree of the formula while duplicating (and providing witnesses for) existential sub-formulas.

As they stand, Herbrand witnesses are not composable in general: given witnesses for $\vdash A$ and $\vdash A \implies B$ there is a priori no direct way to deduce witnesses for $\vdash B$ [11]. Understanding how the data of these witnesses can be elaborated to allow such a composition has thus become a folklore question in proof theory [5, 9, 13, 10].

A first approach to this question is functional interpretations proposed by Gerhardy and Kohlenbach with data provided by Shoenfield's version of Gödel's Dialectica Interpretation [5, 18, 8]. Although it interprets cuts, this approach makes no pretence to be faithful to the structure of proofs as encapsulated in systems like the sequent calculus and it is compelling to seek some compositional form of Herbrand's Theorem arising directly from proofs.

Following this desire, another approach has emerged recently: fixing expansion trees as representations of Herbrand witnesses for cut-free first order proofs [15], one seeks generalisations of these trees that support cuts. Works in that direction include Heijltjes' *proof forests* [9], McKinley's *Herbrand nets* [13], and Hetzl and Weller's more recent *expansion trees with cuts* [10]. In all three cases, a generalisation of expansion trees allowing explicit cuts is given along with a weakly normalising cut reduction procedure, proved correct via syntactic means.

Although the present follows the line described above, its novelty comes from the fact that *(1)* it approaches the problem semantically rather than syntactically, and *(2)* it makes explicit the game-theoretic ideas that underlie more or less explicitely the literature on expansion trees. By embedding expansion trees in a realm of strategies, which are *by design* compositional, one obtains a compositional account of Herbrand's theorem. More precisely, our games are based on Rideau and Winskel's *concurrent strategies* [17, 2], extended with annotations for first-order terms. Beyond the term information, the key ingredient of this model is the causal structure of strategies that allows us to represent transparently the dependencies between quantifiers implicitly carried by sequent proofs. Were we interested only in cut-free sequent calculus our strategies would essentially be Miller's expansion trees, but enriched with *explicit acyclicity witnesses*. This additional data makes the process of composition doable.

On the game-theoretic front, our model is closely related to Laurent's model for the first-order $\lambda\mu$-calculus [12], from which we differ by treating a symmetric proof system with an involutive negation, avoiding sequentiality. Also related is Mimram's categorical construction of a games model for a linear first-order logic without propositional connectives [16].

In this extended abstract, we focus on showing how expansion trees can be regarded as concurrent strategies in game semantics. In Section 2, we recall informally Miller's expansion trees. In Section 3, we introduce our notions of games and strategies, and show how expansion trees can be presented as such. In Section 4, we detail the interpretation of formulas and state our version of Herbrand's theorem. For lack of space we do not detail the interpretation of proofs as winning strategies, but we illustrate in Section 5 the core ingredient of the interpretation: the notion of composition of strategies, which allows us to interpret the cut rule.

2 Notations and Miller Expansion Trees

A **signature** is a pair $\Sigma = (\Sigma_f, \Sigma_p)$, with Σ_f a countable set of **function symbols** ($\mathsf{f}, \mathsf{g}, \mathsf{h}$, *etc.* range over function symbols), and Σ_p a countable set of **predicate symbols** (P, Q, *etc.* range over predicate symbols). There is an **arity function** $\mathrm{ar} : \Sigma_f \uplus \Sigma_p \to \mathbb{N}$ where $\uplus$ is the usual set-theoretic union, where the argument sets are disjoint. For a relative gain in simplicity in some arguments and examples, we assume that Σ has at least one constant symbol, *i.e.* a function symbol of arity 0. We use $\mathsf{a}, \mathsf{b}, \mathsf{c}, \dots$ to range over constant symbols.

If $\mathcal{V}$ is a set of **variable names**, we write $\mathsf{Tm}_\Sigma(\mathcal{V})$ for the set of first-order terms on Σ with free variables in $\mathcal{V}$. We use variables $t, s, u, v, \dots$ to range over terms. **Atomic formulas** have the form $\mathsf{P}(t_1, \dots, t_n)$ or $\neg\mathsf{P}(t_1, \dots, t_n)$, where P is a n-ary predicate symbol and the t_is are terms. **Formulas** are also closed under quantifiers, and the connectives $\vee$ and $\wedge$. **Negation**

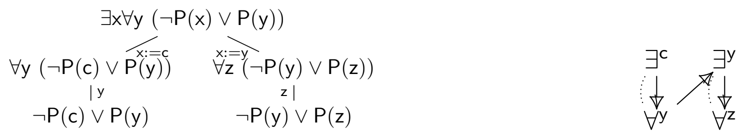

Figure 1: An expansion tree and winning Σ-strategy for *DF*

is not considered a logical connective: the negation $\varphi^\perp$ of φ is obtained by De Morgan rules. We write $\mathsf{Form}_\Sigma(\mathcal{V})$ for the set of **first-order formulas** on Σ with free variables in $\mathcal{V}$, and use $\varphi, \psi, \ldots$ to range over them. We also write $\mathsf{QF}_\Sigma(\mathcal{V})$ for the set of **quantifier-free** formulas. Finally, we write $\mathrm{fv}(\varphi)$ or $\mathrm{fv}(t)$ for the set of free variables in a formula φ or a term t. Formulas are considered up to α-conversion and assumed to satisfy Barendregt's convention.

In [15] Miller introduces expansion trees as compact witnesses of the truth of a first-order formula. Through translations between expansion trees and cut-free sequent proofs, he proves the following theorem, that one may regard as a modern statement of Herbrand's theorem.

Theorem 2.1 ([15])**.** *For any φ, $\models \varphi$ iff φ has an expansion tree.*

Expansion trees can be introduced through a game metaphor, reminiscent of Coquand's game semantics for classical arithmetic [3]. Two players, ∃loïse and ∀bélard, argue about the validity of a formula. On a formula $\forall \mathsf{x}\, \varphi$, ∀bélard provides a fresh variable x and the play keeps going on φ. On a formula $\exists \mathsf{x}\, \varphi$, ∃loïse provides a *term t*, possibly containing variables previously introduced by ∀bélard. ∃loïse, though, has a special power: at any time she can *backtrack* to a previous existential position, and propose a new term. Figure 1 (left) shows an expansion tree the *drinker's formula*:

$$\exists \mathsf{x} \forall \mathsf{y}\ (\neg \mathsf{P}(\mathsf{x}) \vee \mathsf{P}(\mathsf{y})) \qquad (DF)$$

It may be read from top to bottom, and from left to right: ∃loïse plays c, then ∀bélard introduces y, then ∃loïse *backtracks* (we jump to the right branch) and plays y, and ∀bélard introduces z. It is a win for ∃loïse: the disjunction of the leaves is a tautology.

However, the metaphor has limits: the order between two branches of an expansion tree is not part of the structure, but implicit in the term annotations. Yet, this order is crucial to ensure correctness of the trees. Certainly the tree below should not be valid as the formula it plays on is invalid.

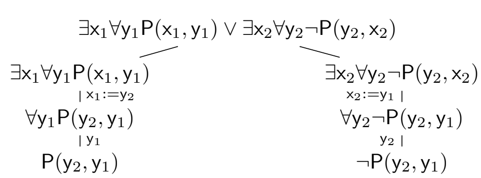

And indeed, the full definition of expansion trees involves a correctness criterion that forbids this: the partial causal relation on nodes resulting from the tree structure and its labelling must be acyclic. As we will see next, in our chosen representations for proofs (called Σ-*strategies*) this causal relation on nodes will be made explicit as a partial order, making the acyclicity correctness criterion redundant.

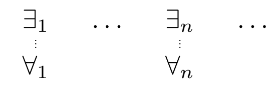

Figure 2: The arena $[\![DF]\!]^{\exists}$

3 Expansion Trees as Winning Σ-Strategies

We give our formulation of expansion trees as Σ-*strategies*. Although our definitions look superficially different from Miller's, the only fundamental difference is the explicit display of the dependency between quantifiers. Σ-strategies will be certain partial orders, with elements either "∀ events" or "∃ events". Events will carry terms, in a way respecting causal dependency. Figure 1 shows, on the right, the representation as a Σ-strategy of the tree on the left.

Σ-strategies will account for first-order *proofs*, and as such, will play on games representing *formulas*. The first component of a game is its *arena*, that specifies the available moves (the quantifiers) and the causal ordering between them.

Definition 3.1. *An **arena** is $A = (|A|, \leq_A, \mathrm{pol}_A)$ where $|A|$ is a set of **events**, $\leq_A$ is a partial order that is* forest-shaped*: (1) if $a_1 \leq_A a$ and $a_2 \leq_A a$, then either $a_1 \leq_A a_2$ or $a_2 \leq_A a_1$, and (2) for all $a \in |A|$, the branch $[a]_A = \{a' \in A \mid a' \leq_A a\}$ is finite. Finally, $\mathrm{pol}_A : |A| \to \{\forall, \exists\}$ is a **polarity function** which expresses if a move belongs to ∃loïse or ∀bélard.*

A **configuration** of an arena (or any partial order) is a down-closed set of events. We write $\mathscr{C}^\infty(A)$ for the set of configurations of A, and $\mathscr{C}(A)$ for the set of *finite* configurations.

The arena only describes the moves available to both players; it says nothing about terms or winning. Similarly to expansion trees where only ∃loïse can replicate her moves ("*backtrack*", although the terminology is imperfect when strategies are not sequential), arenas corresponding to formulas in our interpretation will at first be biased towards ∃loïse: each ∃ move will exist in as many copies as she might desire, whereas ∀ events will not be copied *a priori*. Figure 2 shows the ∃-biased arena $[\![DF]\!]^{\exists}$ for *DF*. The order is drawn from top to bottom, *i.e.* events at the top are minimal. Although only ∃loïse can replicate her moves, the universal quantifier is also copied as it depends on the existential quantifier.

Strategies on arena A will be certain *augmentations* of prefixes of A. They carry causal dependency between quantifiers induced by term annotations, but not the terms themselves.

We introduce the notation $\rightarrowtriangle$, already used implicitly in Figure 1. For A any partial order and $a_1, a_2 \in |A|$, we write $a_1 \rightarrowtriangle_A a_2$ (or $a_1 \rightarrowtriangle a_2$ if A is clear from the context) if $a_1 <_A a_2$ with no other event in between, *i.e.* for any $a \in |A|$ such that $a_1 \leq_A a \leq_A a_2$, then $a_1 = a$ or $a_2 = a$. We call $\rightarrowtriangle$ **immediate causal dependency** in line with event structures where the partial order is that of causal dependency.

Definition 3.2. *A **strategy** σ on arena A, written $\sigma : A$, is a partial order $(|\sigma|, \leq_\sigma)$ with $|\sigma| \subseteq |A|$, such that for all $a \in |\sigma|$, $[a]_\sigma$ is finite (an* elementary event structure*); subject to:*

- *(1)* Arena-respecting. *We have $\mathscr{C}^\infty(\sigma) \subseteq \mathscr{C}^\infty(A)$,*
- *(2)* Receptivity. *If $x \in \mathscr{C}(\sigma)$ such that $x \cup \{a^\forall\} \in \mathscr{C}(A)$, then $a \in |\sigma|$ as well ($a^\forall$ means that $\mathrm{pol}_A(a) = \forall$).*
- *(3)* Courtesy. *If $a_1 \rightarrowtriangle_\sigma a_2$, then either $a_1 \rightarrowtriangle_A a_2$, or $\mathrm{pol}_A(a_1) = \forall$ and $\mathrm{pol}(a_2) = \exists$.*

This is a simplification of Rideau and Winskel's *concurrent strategies* [17] permitted by the purely deterministic setting; also equivalent [17] to Melliès and Mimram's earlier *receptive ingenuous strategies* [14] – though the direct handle on the causal order in the definition above is convenient for our purposes. Receptivity means that ∃loïse cannot refuse to acknowledge a move by ∀bélard, and courtesy that the only new causal constraints that she can enforce with respect to the game is that some existential quantifiers depend on some universal quantifiers. These constraints are consistent with the informal game semantics described in the previous section. Ignoring terms, Figure 1 (on the right) displays a strategy on the arena of Figure 2 – in Figure 1 we also display via dotted lines the immediate dependency of the arena.

We can now add terms, and define Σ-*strategies*.

Definition 3.3. *A* Σ***-strategy*** *on arena* A *is a strategy* $\sigma : A$*, with a* ***labelling function*** $\lambda_\sigma : |\sigma| \to \mathsf{Tm}_\Sigma(|\sigma|)$*, such that:*

$$\begin{array}{rcl} \forall a^\forall \in |\sigma|, \quad \lambda_\sigma(a) & = & a \\ \forall a^\exists \in |\sigma|, \quad \lambda_\sigma(a) & \in & \mathsf{Tm}_\Sigma([a]_\sigma^\forall) \end{array}$$

where $[a]_\sigma^\forall = \{a' \in |\sigma| \mid a' \leq_\sigma a \ \&\ \mathrm{pol}_A(a') = \forall\}$.

Rather than having ∀ moves introduce fresh variables, we find it convenient to consider them as *variables themselves*. Hence, the ∃ moves are annotated by terms having as free variables the ∀ moves in their causal history. For instance, the right of Figure 1 is meant formally to denote the one on the right (where superscripts are the terms given by λ). In the sequel we omit the (redundant) annotation of ∀bélard's events.

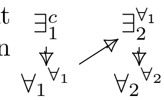

Σ-strategies are more general than expansion trees (besides the fact that they are not assumed finite): they have an explicit causal ordering, which may be more constraining than that given by the terms. A Σ-strategy $\sigma : A$ is **minimal** iff whenever $a_1 \rightarrowtriangle_\sigma a_2$ such that $a_1 \notin \mathrm{fv}(\lambda_\sigma(a_2))$, then $a_1 \rightarrowtriangle_A a_2$ as well. In a minimal Σ-strategy $\sigma : A$, the ordering $\leq_\sigma$ is actually redundant and can be uniquely recovered from λ_σ and $\leq_A$.

In order for the model to be able to discriminate valid from invalid strategies, we lastly need to adjoin *winning conditions* to arenas and define *winning* Σ*-strategies*. As in expansion trees, these amount to the substitution of the expansion of the original formula being a tautology.

Definition 3.4. *A* ***game*** $\mathcal{A}$ *is an arena* A *together with* ***winning conditions****, given as:*

$$\mathcal{W}_\mathcal{A} : (x \in \mathscr{C}^\infty(A)) \to \mathsf{QF}_\Sigma^\infty(x)$$

where $\mathsf{QF}_\Sigma^\infty(x)$ *is the set of* ***infinitary quantifier-free formulas*** *– obtained from* $\mathsf{QF}_\Sigma(x)$ *by adding infinitary connectives* $\bigvee_{i\in I} \varphi_i$ *and* $\bigwedge_{i \in I} \varphi_i$*, where* I *is some countable set.*

We delay until Section 4 the definition of the interpretation of a formula as a game. However, the idea is relatively simple: for the game interpreting φ, the winning conditions associate configurations $x \in \mathscr{C}^\infty(\llbracket\varphi\rrbracket)$ with the propositional part of the sub-formula of φ explored by x, with duplications (as parts of φ may be visited several times). Quantified variables are replaced with the names of the events corresponding to the quantifiers.

$$\begin{array}{rcl} \mathcal{W}_{\llbracket DF \rrbracket^\exists}(\{\exists_3, \forall_3, \exists_6, \forall_6\}) & = & (\neg \mathsf{P}(\exists_3) \vee \mathsf{P}(\forall_3)) \vee \\ & & (\neg \mathsf{P}(\exists_6) \vee \mathsf{P}(\forall_6)) \\ \mathcal{W}_{\llbracket DF \rrbracket^\exists}(\{\exists_3, \forall_3, \exists_6\}) & = & (\neg \mathsf{P}(\exists_3) \vee \mathsf{P}(\forall_3)) \vee \top \end{array}$$

where the arena for *DF* appears in Figure 2. The $\top$ (the true formula) on the second line is due to ∀bélard not having played $\forall_6$ yet, yielding victory to ∃loïse on that configuration. The winning conditions yield syntactic, uninterpreted formulas: we keep the second formula as-is although it is equivalent to $\top$.

Finally, we can define winning strategies.

Definition 3.5. *If $\sigma : A$ is a Σ-strategy and $x \in \mathscr{C}^\infty(\sigma)$, we say that x is **tautological** in σ if the formula $\mathcal{W}_\mathcal{A}(x)[\lambda_\sigma]$ corresponding to the substitution of $\mathcal{W}_\mathcal{A}(x) \in \mathsf{QF}^\infty_\Sigma(x)$ by $\lambda_\sigma : x \to \mathsf{Tm}_\Sigma(x)$, is a (possibly infinite) tautology.*

*A Σ-strategy $\sigma : A$ is **winning** if any $x \in \mathscr{C}^\infty(\sigma)$ that is $\exists$**-maximal** (*i.e.* $x \in \mathscr{C}^\infty(\sigma)$ such that for all $a \in |\sigma|$ with $x \cup \{a\} \in \mathscr{C}^\infty(\sigma)$, $\mathrm{pol}_A(a) = \forall$) is tautological. A Σ-strategy $\sigma : A$ is **top-winning** if $|\sigma| \in \mathscr{C}^\infty(\sigma)$ is tautological.*

A winning Σ-strategy is top-winning, but not always the other way around. The *minimal*, *top-winning* Σ-strategies $\sigma : [\![\varphi]\!]^\exists$ will correspong to expansion trees; but the *winning strategies* will behave better compositionally.

4 Constructions on games and Herbrand's theorem

To complete our statement of Herbrand's theorem, we define the interpretation of formulas.

Arenas. First, we define some operations on arenas. We write $\emptyset$ for the **empty arena**, with no events. If A is an arena, we write $A^\perp$ for the **dual** arena, with the same events and causality but polarity reversed, *i.e.* $\mathrm{pol}_{A^\perp}(a) = \forall$ iff $\mathrm{pol}_A(a) = \exists$. We review some other constructions.

Definition 4.1. *The **simple parallel composition** $A_1 \parallel A_2$ of A_1 and A_2 has as events the tagged disjoint union $\{1\} \times |A_1| \uplus \{2\} \times |A_2|$, causal order given by $(i, a) \leq_{A_1 \parallel A_2} (j, a')$ iff $i = j$ and $a \leq_{A_i} a'$. Polarity is $\mathrm{pol}_{A_1 \parallel A_2}((i, a)) = \mathrm{pol}_{A_i}(a)$.*

Configurations $x \in \mathscr{C}^\infty(A \parallel B)$ have the form $\{1\} \times x_A \cup \{2\} \times x_B$ with $x_A \in \mathscr{C}^\infty(A)$ and $x_B \in \mathscr{C}^\infty(B)$, which we write $x = x_A \parallel x_B$. Binary simple parallel composition has a general counterpart $\parallel_{i \in I} A_i$ with I at most countable, defined likewise. We will use the uniform countably infinite simple parallel composition $\parallel_\omega A$ with ω parallel copies of A.

Another important arena construction is *prefixing*.

Definition 4.2. *For $\alpha \in \{\forall, \exists\}$ and A an arena, the **prefixed arena** $\alpha.A$ has events $\{(1, \alpha)\} \cup \{2\} \times |A|$ and $(i, a) \leq (j, a')$ iff $i = j = 2$ and $a \leq_A a'$, or $(i, a) = (1, \alpha)$; meaning that $(1, \alpha)$ is the unique minimal event in $\alpha.A$. Its polarity is $\mathrm{pol}_{\alpha.A}((1, \alpha)) = \alpha$ and $\mathrm{pol}_{\alpha.A}((2, a)) = \mathrm{pol}_A(a)$.*

Configurations $x \in \mathscr{C}^\infty(\alpha.A)$ are either empty, or of the form $\{(1, \alpha)\} \cup \{2\} \times x_A$ with $x_A \in \mathscr{C}^\infty(A)$, written $\alpha.x_A$.

Winning. To give the inductive interpretation of formulas we have to consider formulas that are not closed, *i.e.* with free variables. For $\mathcal{V}$ a finite set, a $\mathcal{V}$**-game** is defined as a game $\mathcal{A}$ as in Definition 3.4, but with signature Σ extended with $\mathcal{V}$. In other words, for $x \in \mathscr{C}^\infty(A)$,

$$\mathcal{W}_\mathcal{A}(x) \in \mathsf{QF}^\infty_{\Sigma \uplus \mathcal{V}}(x)\,.$$

We now define all our constructions, on $\mathcal{V}$-games rather than on games. The duality operation on arenas $(-)^\perp$ extends to $\mathcal{V}$-games, simply by negating the winning conditions: for all $x \in \mathscr{C}^\infty(A)$, $\mathcal{W}_{\mathcal{A}^\perp}(x) = \mathcal{W}_\mathcal{A}(x)^\perp$. The $\parallel$ of arenas gives rise to *two* constructions on $\mathcal{V}$-games:

$$\begin{array}{rclcrcl}
[\![\top]\!]^{\exists}_{\mathcal{V}} & = & 1 & & [\![\mathsf{P}(t_1,\ldots,t_n)]\!]^{\exists}_{\mathcal{V}} & = & \mathsf{P}(t_1,\ldots,t_n) \\
[\![\bot]\!]^{\exists}_{\mathcal{V}} & = & \bot & & [\![\neg\mathsf{P}(t_1,\ldots,t_n)]\!]^{\exists}_{\mathcal{V}} & = & \neg\mathsf{P}(t_1,\ldots,t_n) \\
[\![\exists \mathsf{x}\varphi]\!]^{\exists}_{\mathcal{V}} & = & ?\exists\mathsf{x}.[\![\varphi]\!]^{\exists}_{\mathcal{V}\uplus\{\mathsf{x}\}} & & [\![\varphi_1 \vee \varphi_2]\!]^{\exists}_{\mathcal{V}} & = & [\![\varphi_1]\!]^{\exists}_{\mathcal{V}} \parr [\![\varphi_2]\!]^{\exists}_{\mathcal{V}} \\
[\![\forall \mathsf{x}\varphi]\!]^{\exists}_{\mathcal{V}} & = & \forall\mathsf{x}.[\![\varphi]\!]^{\exists}_{\mathcal{V}\uplus\{\mathsf{x}\}} & & [\![\varphi_1 \wedge \varphi_2]\!]^{\exists}_{\mathcal{V}} & = & [\![\varphi_1]\!]^{\exists}_{\mathcal{V}} \otimes [\![\varphi_2]\!]^{\exists}_{\mathcal{V}}
\end{array}$$

Figure 3: $\exists$-biased interpretation of formulas

Definition 4.3. *Let $\mathcal{A}$ and $\mathcal{B}$ be $\mathcal{V}$-games. We define two $\mathcal{V}$-games on arena $A \parallel B$, differing by the winning condition:*

$$\begin{array}{rcl}
\mathcal{W}_{\mathcal{A}\otimes\mathcal{B}}(x_A \parallel x_B) & = & \mathcal{W}_{\mathcal{A}}(x_A) \wedge \mathcal{W}_{\mathcal{B}}(x_B) \\
\mathcal{W}_{\mathcal{A}\parr\mathcal{B}}(x_A \parallel x_B) & = & \mathcal{W}_{\mathcal{A}}(x_A) \vee \mathcal{W}_{\mathcal{B}}(x_B)
\end{array}$$

Note the implicit renaming so that $\mathcal{W}_{\mathcal{A}}(x_A), \mathcal{W}_{\mathcal{B}}(x_B)$ are in $\mathsf{QF}^{\infty}_{\Sigma\uplus\mathcal{V}}(x_A \parallel x_B)$ rather than $\mathsf{QF}^{\infty}_{\Sigma\uplus\mathcal{V}}(x_A), \mathsf{QF}^{\infty}_{\Sigma\uplus\mathcal{V}}(x_B)$ respectively – here and in the sequel, we will keep such renamings implicit when we believe it helps readability.

Note that $\otimes$ and $\parr$ are De Morgan duals, *i.e.* $(\mathcal{A}\otimes\mathcal{B})^{\perp} = \mathcal{A}^{\perp} \parr \mathcal{B}^{\perp}$. The reader may wonder why these operations are written $\otimes$ and $\parr$ rather than $\wedge$ and $\vee$. This is because, as we will see, these operations by themselves behave more like the connectives of linear logic [6] than those of classical logic; for each $\mathcal{V}$ the $\otimes$ and $\parr$ will form the basis of a $*$-autonomous structure and hence a model of multiplicative linear logic.

To recover classical logic, we will add *replication* to the interpretation of formulas.

Definition 4.4. *Let $\mathcal{A}$ be a $\mathcal{V}$-game. We define two new $\mathcal{V}$-games $!\mathcal{A}$ and $?\mathcal{A}$ with arena $\parallel_{\omega} A$, and winning conditions:*

$$\begin{array}{rcl}
\mathcal{W}_{!\mathcal{A}}(\parallel_{i\in\omega} x_i) & = & \bigwedge_{i\in\omega} \mathcal{W}_{\mathcal{A}}(x_i) \\
\mathcal{W}_{?\mathcal{A}}(\parallel_{i\in\omega} x_i) & = & \bigvee_{i\in\omega} \mathcal{W}_{\mathcal{A}}(x_i)
\end{array}$$

Although $\mathcal{W}_{!\mathcal{A}}(x)$ (resp. $\mathcal{W}_{?\mathcal{A}}(x)$) is, syntactically, an infinite conjunction (resp. disjunction), we always implicitly simplify it to a finite one when x visits finitely many copies (as we then have infinitely many occurrences of $\mathcal{W}_{\mathcal{A}}(\emptyset)$).

Next we show how $\mathcal{V}$-games support quantifiers.

Definition 4.5. *For $\mathcal{A}$ a $(\mathcal{V} \uplus \{\mathsf{x}\})$-game, the $\mathcal{V}$-game $\forall\mathsf{x}.\mathcal{A}$ and its dual $\exists\mathsf{x}.\mathcal{A}$ have arenas $\forall.A$ and $\exists.A$ respectively, and:*

$$\begin{array}{rclcrcl}
\mathcal{W}_{\forall\mathsf{x}.\mathcal{A}}(\emptyset) & = & \top & & \mathcal{W}_{\forall\mathsf{x}.\mathcal{A}}(\forall.x_A) & = & \mathcal{W}_{\mathcal{A}}(x_A)[\forall/\mathsf{x}] \\
\mathcal{W}_{\exists\mathsf{x}.\mathcal{A}}(\emptyset) & = & \bot & & \mathcal{W}_{\exists\mathsf{x}.\mathcal{A}}(\exists.x_A) & = & \mathcal{W}_{\mathcal{A}}(x_A)[\exists/\mathsf{x}]
\end{array}$$

Finally, we regard an atomic formula φ (*i.e.* $\mathsf{P}(t_1,\ldots,t_n)$ or $\neg\mathsf{P}(t_1,\ldots,t_n)$ with $t_i \in \mathsf{Tm}_{\Sigma}(\mathcal{V})$) as a $\mathcal{V}$-game on arena $\emptyset$, with $\mathcal{W}_{\varphi}(\emptyset) = \varphi$. We write 1 and $\bot$ for the unit $\mathcal{V}$-games on arena $\emptyset$ with winning conditions respectively $\top$ and $\bot$.

Putting all of these together, we give in Figure 3 the general definition of the $\exists$-biased interpretation of a formula $\varphi \in \mathsf{Form}_{\Sigma}(\mathcal{V})$ as a $\mathcal{V}$-game. Note the difference between the case of existential and universal formulas, reflecting the bias towards $\exists$loïse in the interpretation. The reader can check that this is indeed compatible with the examples given previously.

We can now state our formulation of Herbrand's theorem.

Theorem 4.6. *For any closed formula φ, we have $\models \varphi$ iff there exists a finite, top-winning Σ-strategy $\sigma : [\![\varphi]\!]^{\exists}$.*

Though it takes some effort to set up, this is an elegant way of stating Herbrand's theorem, putting the emphasis on causality between quantifiers. But, besides the game-theoretic language, there is nothing fundamentally new or surprising about this statement. Indeed for now Σ-strategies are static objects, alternative bureaucracy-free representations of cut-free proofs. In particular, expansion trees are the *minimal* top-winning Σ-strategies $\sigma : [\![\varphi]\!]^{\exists}$.

5 Compositional Herbrand's theorem

Though Theorem 4.6 above could be deduced via the connection with expansion trees, this would make its validity intrinsically rely on the admissibility of cut in the sequent calculus. But unlike expansion trees, strategies can be *composed*. Our games model allows us to give an alternative proof of Herbrand's theorem where strategies are obtained truly *compositionally* from any sequent proof, without first eliminating cuts. In other words, expansion trees come naturally from an interpretation of the classical sequent calculus in our game model.

To compose Σ-strategies, we have to restore the symmetry between ∃loïse and ∀bélard in the interpretation of formulas so that $[\![\neg\varphi]\!]_{\mathcal{V}} = [\![\varphi]\!]_{\mathcal{V}}^{\perp}$. The *non-biased* interpretation $[\![\varphi]\!]_{\mathcal{V}}$ of $\varphi \in \mathsf{Form}_{\Sigma}(\mathcal{V})$ is defined as for $[\![\varphi]\!]_{\mathcal{V}}^{\exists}$, except for universal formulas, where instead we set $[\![\forall \mathsf{x} \varphi]\!]_{\mathcal{V}} = !\forall \mathsf{x}. [\![\varphi]\!]_{\mathcal{V} \uplus \{\mathsf{x}\}}$. This symmetry means that we lose finiteness, since now ∃loïse must be reactive to the infinite number of copies potentially opened by ∀bélard.

But we can now state:

Theorem 5.1. *For φ closed, the following are equivalent:*

(1) $\models \varphi$,

(2) There exists a finite, top-winning Σ-strategy $\sigma : [\![\varphi]\!]^{\exists}$,

(3) There exists a winning Σ-strategy $\sigma : [\![\varphi]\!]$.

That *(3)* implies *(2)* relies on a compactness argument to extract a *finite* top-winning Σ-strategy; and that *(2)* implies *(1)* can be seen via relatively simple syntactic considerations.

That *(1)* implies *(3)* is where our true contribution lies. It is a *semantic* proof: we construct a model of the classical sequent calculus LK, associating with each rule of LK a corresponding construction on winning Σ-strategies. There are translations between cut-free sequent proofs and expansion trees, and to a large extent our interpretation follows those (modulo a few subtleties due to our non-biased interpretation; for those the interested reader may consult the version of the paper on the authors' web page [1]). In the remainder of this extended abstract, we focus on explaining the interpretation of the cut rule:

$$\textsc{Cut}\ \frac{\vdash \Gamma, \varphi \qquad \vdash \varphi^{\perp}, \Delta}{\vdash \Gamma, \Delta}$$

A proof π of an LK sequent $\vdash \varphi_1, \ldots, \varphi_n$ is interpreted as a winning Σ-strategy $[\![\pi]\!] : [\![\varphi_1]\!] \parr \cdots \parr [\![\varphi_n]\!]$. In order to interpret the cut rule, one needs to make sense of the *composition* $[\![\pi_2]\!] \odot [\![\pi_1]\!]$ between $[\![\pi_1]\!] : [\![\Gamma]\!] \parr [\![\varphi]\!]$ and $[\![\pi_2]\!] : [\![\varphi]\!]^{\perp} \parr [\![\Delta]\!]$ resulting in a strategy over $[\![\Gamma]\!] \parr [\![\Delta]\!]$. This composition is computed in two stages: first, the *interaction* $[\![\pi_2]\!] \circledast [\![\pi_1]\!]$ is obtained as the most general partial-order-with-terms satisfying the constraints given by both $[\![\pi_1]\!]$ and $[\![\pi_2]\!]$. Briefly, $[\![\pi_2]\!] \circledast [\![\pi_1]\!]$ has events those for which the causal constraints imposed by $[\![\pi_1]\!]$ and $[\![\pi_2]\!]$ yield an acyclic relation; and its causal structure is the weakest partial order compatible with $\leq_{[\![\pi_1]\!]}$ and $\leq_{[\![\pi_2]\!]}$ – as this is already the case in the interaction of concurrent strategies without labels described in [2]. On the labelling side, the effect of interaction is to perform substitution of the negative variables in $\lambda_{[\![\pi_1]\!]}$ and $\lambda_{[\![\pi_2]\!]}$ by the terms associated with their corresponding positive variables in the other labelling. The resulting substitution is in fact a most general unifier for the set of equations $\{\lambda_{[\![\pi_1]\!]}(e) \doteq \lambda_{[\![\pi_2]\!]}(e)\}_{e \in [\![\pi_2]\!] \circledast [\![\pi_1]\!]}$. Figure 4 displays such an interaction. Finally, the composition is the resulting Σ-strategies obtained by *hiding* events in B. In the example of Figure 4 we get the single annotated event $\exists_5^{\mathsf{f}(\mathsf{g}(\mathsf{c}),\mathsf{h}(\mathsf{c}))}$.

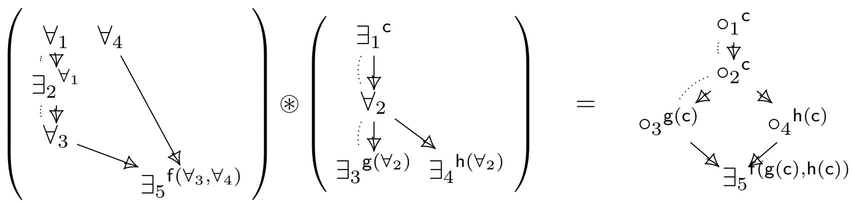

Figure 4: Interaction of $\sigma : 1^{\perp} \parallel (\exists_1 \forall_2 \exists_3 \parallel \exists_4)$ and $\tau : (\exists_1 \forall_2 \exists_3 \parallel \exists_4)^{\perp} \parallel \exists_5$

We can check that composition preserves winning. Furthermore for a game $\mathcal{A}$, we define the winning Σ-strategy $\text{œ}_{\mathcal{A}} : \mathcal{A}^{\perp} \parr \mathcal{A}$, called the *copycat strategy*, to have partial order $\leq_{\text{œ}_{\mathcal{A}}}$ the transitive closure of

$$\leq_{A^{\perp} \parallel A} \cup \{((i,a),(3-i,a)) \mid (i,a)^{\forall} \in |A^{\perp} \parallel A|\}$$

and labelling function $\lambda_{\text{œ}_A}$ defined by

$$\lambda_{\text{œ}_A}((i,a)^{\forall}) = (i,a), \qquad \lambda_{\text{œ}_A}((i,a)^{\exists}) = (3-i,a)$$

With this definition, it turns out that $\mathcal{V}$-games together with winning Σ-strategies $\sigma : \mathcal{A}^{\perp} \parr \mathcal{B}$ viewed as morphism $\sigma : \mathcal{A} \twoheadrightarrow \mathcal{B}$ define a linearly distributive category with negation Games; with identities the copycat strategies and tensor products $\parr$ and $\otimes$. As a consequence,

Theorem 5.2. *The category* Games *of concurrent games and winning* Σ*-strateges is a model of MLL.*

This covers the interpretation of MLL rules; and show that the interpretation is invariant under cut reduction between MLL rules. However, invariance of the interpretation under cut reduction does not hold for full LK (or the model would collapse to a boolean algebra [7]). Likewise, just as classical proofs can lead to arbitrary large cut-free proofs [4], our interpretation may yield *infinite* winning strategies, from which *finite* sub-strategies can nonetheless always be extracted. This seems to relate to the fact that syntactic notions of expansion trees with cuts [9, 13, 10] are in general weakly, rather than strongly, normalizing.

References

[1] A. Alcolei, P. Clairambault, M. Hyland, and G. Winskel. The true concurrency of herbrand's theorem. available at perso.ens-lyon.fr/pierre.clairambault/herbrand.pdf.

[2] Simon Castellan, Pierre Clairambault, Silvain Rideau, and Glynn Winskel. Games and strategies as event structures. 2016. Submitted, https://hal.inria.fr/hal-01302713.

[3] Thierry Coquand. A semantics of evidence for classical arithmetic. *J. Symb. Log.*, 60(1):325–337, 1995.

[4] Vincent Danos, Jean-Baptiste Joinet, and Harold Schellinx. A new deconstructive logic: Linear logic. *J. Symb. Log.*, 62(3):755–807, 1997.

[5] Philipp Gerhardy and Ulrich Kohlenbach. Extracting herbrand disjunctions by functional interpretation. *Arch. Math. Log.*, 44(5):633–644, 2005.

[6] Jean-Yves Girard. Linear logic. *Theor. Comput. Sci.*, 50:1–102, 1987.

[7] Jean-Yves Girard. A new constructive logic: Classical logic. *Mathematical Structures in Computer Science*, 1(3):255–296, 1991.

[8] Kurt Gödel. Über eine bisher noch nicht benützte erweiterung des finiten standpunktes. *dialectica*, 12(3-4):280–287, 1958.

[9] Willem Heijltjes. Classical proof forestry. *Ann. Pure Appl. Logic*, 161(11):1346–1366, 2010.

[10] Stefan Hetzl and Daniel Weller. Expansion trees with cut. *CoRR*, abs/1308.0428, 2013.

[11] Ulrich Kohlenbach. On the no-counterexample interpretation. *J. Symb. Log.*, 64(4):1491–1511, 1999.

[12] Olivier Laurent. Game semantics for first-order logic. *Logical Methods in Computer Science*, 6(4), 2010.

[13] Richard McKinley. Proof nets for Herbrand's theorem. *ACM Trans. Comput. Log.*, 14(1):5, 2013.

[14] Paul-André Melliès and Samuel Mimram. Asynchronous games: Innocence without alternation. In Luís Caires and Vasco Thudichum Vasconcelos, editors, *CONCUR*, volume 4703 of *LNCS*, pages 395–411. Springer, 2007.

[15] Dale A Miller. A compact representation of proofs. *Studia Logica*, 46(4):347–370, 1987.

[16] Samuel Mimram. The structure of first-order causality. *Mathematical Structures in Computer Science*, 21(1):65–110, 2011.

[17] Silvain Rideau and Glynn Winskel. Concurrent strategies. In *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science, LICS 2011, June 21-24, 2011, Toronto, Ontario, Canada*, pages 409–418. IEEE Computer Society, 2011.

[18] J.R. Shoenfield. *Mathematical Logic*. Addison-Wesley, 1967.