Infinitary Proof Theory: the Multiplicative Additive Case

David Baelde[1], Amina Doumane[2], and Alexis Saurin[3]

1 **LSV, CNRS & ENS Cachan, Université Paris-Saclay, France**
david.baelde@lsv.fr
2 **IRIF, PPS, CNRS, Université Paris Diderot & Inria, France**
amina.doumane@pps.univ-paris-diderot.fr
3 **IRIF, PPS, CNRS, Université Paris Diderot & Inria, France**
alexis.saurin@pps.univ-paris-diderot.fr

Abstract

Infinitary and regular proofs are commonly used in fixed point logics. Being natural intermediate devices between semantics and traditional finitary proof systems, they are commonly found in completeness arguments, automated deduction, verification, etc. However, their proof theory is surprisingly underdeveloped. In particular, very little is known about the computational behavior of such proofs through cut elimination. Taking such aspects into account has unlocked rich developments at the intersection of proof theory and programming language theory. One would hope that extending this to infinitary calculi would lead, *e.g.*, to a better understanding of recursion and corecursion in programming languages. Structural proof theory is notably based on two fundamental properties of a proof system: cut elimination and focalization. The first one is only known to hold for restricted (purely additive) infinitary calculi, thanks to the work of Santocanale and Fortier; the second one has never been studied in infinitary systems. In this paper, we consider the infinitary proof system $\mu\mathsf{MALL}^\infty$ for multiplicative and additive linear logic extended with least and greatest fixed points, and prove these two key results. We thus establish $\mu\mathsf{MALL}^\infty$ as a satisfying computational proof system in itself, rather than just an intermediate device in the study of finitary proof systems.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases Infinitary proofs, linear logic

Digital Object Identifier 10.4230/LIPIcs.CSL.2016.42

1 Introduction

Proof systems based on non-well-founded derivation trees arise naturally in logic, even more so in logics featuring fixed points. A prominent example is the long line of work on tableaux systems for modal μ-calculi, *e.g.*, [16, 24, 14, 11], which have served as the basis for analysing the complexity of the satisfiability problem, as well as devising practical algorithms for solving it. One key observation in such a setting, and many others, is that one needs not consider arbitrary infinite derivations but can restrict to *regular* derivation trees (also known as *circular* proofs) which are finitely representable and amenable to algorithmic manipulation. Because infinitary systems are easier to work with than the finitary proof systems (or axiomatizations) based on Kozen-Park (co)induction schemes, they are often found in completeness arguments for such finitary systems [16, 26, 27, 28, 15, 12]. We should note, however, that those arguments are far from being limited to translations from (regular) infinitary to finitary proofs, since such translations are very complex and only known to work in limited cases.

25th EACSL Annual Conference on Computer Science Logic (CSL 2016).
Editors: Jean-Marc Talbot and Laurent Regnier; Article No. 42; pp. 42:1–42:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

There are many other uses of infinite (or regular) derivations, *e.g.*, to study the relationship between induction and infinite descent in first-order arithmetic [9], to generate invariants for program verification in separation logic [8], or as an intermediate between ludics' designs and proofs in linear logic with fixed points [5]. Last but not least, Santocanale introduced circular proofs [22] as a system for representing morphisms in μ-bicomplete categories [21, 23], corresponding to simple computations on (co)inductive data.

Surprisingly, despite the elegance and usefulness of infinitary proof systems, few proof theoretical studies are directly targeting these objects. More precisely, we are concerned with an analysis of proofs that takes into account their computational behaviour in terms of cut elimination. In other words, we would hope that the Curry-Howard correspondence extends nicely to infinitary proofs. In this line of proof-theoretical study, two main properties stand out: cut elimination and focalization; we shall see that they have been barely addressed in infinitary proof systems. The idea of cut elimination is as old as sequent calculus, and at the heart of the proof-as-program viewpoint, where the process of eliminating cuts in proofs is seen as computation. Considering logics with least and greatest fixed points, the computational behavior of induction and coinduction is recursion and corecursion respectively, two important and complex programming principles that would a logical understanding. Note that the many completeness results for infinitary proof systems (*e.g.*, for modal μ-calculi) only imply cut admissibility, but say nothing about the computational process of cut elimination. To our knowledge, leaving aside an early and very restrictive result of Santocanale [22], cut elimination has only been studied by Fortier and Santocanale [13] who considered an infinitary sequent calculus for lattice logic (purely additive linear logic with least and greatest fixed points) and showed that certain cut reductions converge to a limit cut-free derivation. Their proof involves a mix of combinatorial and topological arguments. So far, it has resisted attempts to extend it beyond the purely additive case. The second key property, much more recently identified than cut elimination, is focalization. It has appeared in the work of [3] on proof search and logic programming in linear logic, and is now recognized as one of the deep outcomes of linear logic, putting to the foreground the role of *polarity* in logic. In a way, focalization generalizes the reversibility results that are notably behind most deductive systems for classical μ-calculi, by bringing some key observations about non-reversible connectives. Besides its deep impact on proof search and logical frameworks, focalization resulted in important advances in all aspects of computational proof theory: in the game-semantical analysis of logic [17, 19], the understanding of evaluation order of programming languages, CPS translations, or semantics of pattern matching [10, 29], the space compression in computational complexity [25, 7], etc. Briefly, one can say that while proof nets have led to a better understanding of phenomena related to parallelism with proof-theoretical methods, polarities and focalization have led to a fine-grained understanding of sequentiality in proofs and programs. To the best of our knowledge, while reversibility has since long been a key-ingredient in completeness arguments based on infinitary proof systems, focalization has simply never been studied in such settings.

Organization and contributions of the paper. In this paper, we consider the logic μMALL, that is multiplicative additive linear logic extended with least and greatest fixed point operators. It has been studied in finitary sequent calculus [4]: it notably enjoys free-cut elimination, and focalization has been shown to extend nicely (though not obviously) to it. We give in Section 2 a natural infinitary proof system for μMALL, called $\mu\mathsf{MALL}^\infty$, which notably extends that of Santocanale and Fortier [13]. The system $\mu\mathsf{MALL}^\infty$ is also related to μMALL in the sense that any μMALL derivation can be turned into a $\mu\mathsf{MALL}^\infty$ proof, with cuts. We study the focalization of μMALL$^\infty$ in Section 3. We find out that, even though fixed point polarities are not forced in the finitary sequent calculus for μMALL, they are uniquely determined in μMALL$^\infty$. Despite some novel aspects due to the infinitary nature of our calculus, we are able to re-use the generic *focalization graph* argument [20] to prove that focalized proofs are complete. We then turn to cut elimination in Section 4 and show that (fair) cut reductions converge to an infinitary cut free derivation. We could not apply any standard cut elimination technique (*e.g.*, induction on formulas and proofs, reducibility arguments, topological arguments as in [13]) and propose instead an unusual argument in which a coarse truth semantics is used to show that the cut elimination process cannot go wrong. We also note here that, even for the regular fragment of μMALL$^\infty$, it would be highly non-trivial to obtain cut elimination from the result for μMALL, since it is not known whether regular μMALL$^\infty$ derivations can be translated to μMALL derivations (even without requiring that this translation preserves the computational behaviour of proofs). We conclude in Section 5 with directions for future work. Technical details, proofs, and additional background material can be found in the long version of this paper [6].

2 μMALL and its infinitary proof system μMALL$^\infty$

In this section we introduce multiplicative additive linear logic extended with least and greatest fixed point operators, and an infinitary proof system for it.

▶ **Definition 1.** Given an infinite set of propositional variables $\mathcal{V} = \{X, Y, \dots\}$, μMALL$^\infty$ *pre-formulas* are built over the following syntax:

$$\varphi, \psi ::= \mathbf{0} \mid \top \mid \varphi \oplus \psi \mid \varphi \& \psi \mid \bot \mid \mathbf{1} \mid \varphi ⅋ \psi \mid \varphi \otimes \psi \mid \mu X.\varphi \mid \nu X.\varphi \mid X \quad \text{with} \quad X \in \mathcal{V}.$$

The connectives μ and ν bind the variable X in φ. From there, bound variables, free variables and capture-avoiding substitution are defined in a standard way. The subformula ordering is denoted $\leq$ and fv($\bullet$) denotes free variables. Closed pre-formulas are simply called *formulas*. Note that negation is not part of the syntax, so that we do not need any positivity condition on fixed point expressions.

▶ **Definition 2.** *Negation* is the involution on pre-formulas written $\varphi^\perp$ and satisfying $(\varphi ⅋ \psi)^\perp = \psi^\perp \otimes \varphi^\perp$, $(\varphi \oplus \psi)^\perp = \psi^\perp \& \varphi^\perp$, $\bot^\perp = \mathbf{1}$, $\mathbf{0}^\perp = \top$, $(\nu X.\varphi)^\perp = \mu X.\varphi^\perp$, $X^\perp = X$.

Having $X^\perp = X$ might be surprising, but it is harmless since our proof system will only deal with closed pre-formulas. Our definition yields, *e.g.*, $(\mu X.X)^\perp = (\nu X.X)$ and $(\mu X.\mathbf{1} \oplus X)^\perp = (\nu X.X \& \bot)$, as expected [4]. Note that we also have $(\varphi[\psi/X])^\perp = \varphi^\perp[\psi^\perp/X]$.

Sequent calculi are sometimes presented with sequents as sets or multisets of formulas, but most proof theoretical observations actually hold in a stronger setting where one distinguishes between several *occurrences* of a formula in a sequent, which gives the ability to precisely *trace* the provenance of each occurrence. This more precise viewpoint is necessary, in particular, when one views proofs as programs. In this work, due to the nature of our proof system and because of the operations that we perform on proofs and formulas, it is also crucial to work with occurrences. There are several ways to formally treat occurrences; for the sake of clarity, we provide below a concrete presentation of that notion which is well suited for our needs.

▶ **Definition 3.** An *address* is a word over $\Sigma = \{l, r, i\}$, which stands for left, right and inside. We define a *duality* over Σ^* as the morphism satisfying $l^\perp = r$, $r^\perp = l$ and $i^\perp = i$. We say that α' is a *sub-address* of α when α is a prefix of α', written $\alpha \sqsubseteq \alpha'$. We say that α and β are *disjoint* when α and β have no upper bound wrt. $\sqsubseteq$.

$$\frac{\vdash F,\Gamma \quad \vdash G,\Gamma}{\vdash F \& G,\Gamma}(\&) \qquad \frac{\vdash F,G,\Gamma}{\vdash F ⅋ G,\Gamma}(⅋) \qquad \frac{\vdash F_i,\Gamma}{\vdash F_1\oplus F_2,\Gamma}(\oplus_i) \qquad \frac{\vdash F,\Gamma \quad \vdash G,\Delta}{\vdash F\otimes G,\Gamma,\Delta}(\otimes)$$

$$\frac{}{\vdash \top,\Gamma}(\top) \qquad \frac{\vdash \Gamma}{\vdash \bot,\Gamma}(\bot) \qquad \text{(no rule for } \mathbf{0}\text{)} \qquad \frac{}{\vdash \mathbf{1}}(\mathbf{1})$$

$$\frac{\vdash F[\mu X.F/X],\Gamma}{\vdash \mu X.F,\Gamma}(\mu) \qquad \frac{\vdash G[\nu X.G/X],\Gamma}{\vdash \nu X.G,\Gamma}(\nu) \qquad \frac{F\equiv G}{\vdash F,G^\perp}(\mathsf{Ax}) \qquad \frac{\vdash \Gamma,F \quad \vdash F^\perp,\Delta}{\vdash \Gamma,\Delta}(\mathsf{Cut})$$

Figure 1 Rules of the proof system $\mu\mathsf{MALL}^\infty$.

▶ **Definition 4.** A *(pre)formula occurrence* (denoted by F, G, H) is given by a (pre)formula φ and an address α, and written φ_α. We say that occurrences are *disjoint* when their addresses are. The occurrences φ_α and ψ_β are *structurally equivalent*, written $\varphi_\alpha \equiv \psi_\beta$, if $\varphi = \psi$. Operations on formulas are extended to occurrences as follows: $(\varphi_\alpha)^\perp = (\varphi^\perp)_{\alpha^\perp}$; for any $\star \in \{⅋,\otimes,\oplus,\&\}$, $F \star G = (\varphi \star \psi)_\alpha$ if $F = \varphi_{\alpha l}$ and $G = \psi_{\alpha r}$; for any $\sigma \in \{\mu,\nu\}$, $\sigma X.F = (\sigma X.\varphi)_\alpha$ if $F = \varphi_{\alpha i}$; we also allow ourselves to write units as formula occurrences without specifying their address, which can be chosen arbitrarily. Finally, *substitution of occurrences* forgets addresses: $(\varphi_\alpha)[\psi_\beta/X] = (\varphi[\psi/X])_\alpha$.

▶ **Example 5.** Let $F = \varphi_{\alpha l}$ and $G = \psi_{\alpha r}$. We have, on the one hand, $(F\otimes G)^\perp = ((\varphi\otimes\psi)_\alpha)^\perp = (\psi^\perp ⅋ \varphi^\perp)_{\alpha^\perp}$ and, on the other hand, $G^\perp ⅋ F^\perp = (\psi^\perp)_{\alpha^\perp l} ⅋ (\varphi^\perp)_{\alpha^\perp r} = (\psi^\perp ⅋ \varphi^\perp)_{\alpha^\perp}$. Thus, $(F\otimes G)^\perp = G^\perp ⅋ F^\perp$. We could have designed our system to obtain $(F\otimes G)^\perp = F^\perp ⅋ G^\perp$ instead; this choice is inessential for the present work but makes our definitions suitable, in principle, for a treatment of non-commutative logic.

▶ **Definition 6.** The *Fischer-Ladner closure* of a formula occurrence F, denoted by $\mathrm{FL}(F)$, is the least set of formula occurrences such that $F \in \mathrm{FL}(F)$ and, whenever $G \in \mathrm{FL}(F)$,

- $G_1, G_2 \in \mathrm{FL}(F)$ if $G = G_1 \star G_2$ for any $\star \in \{\oplus, \&, ⅋, \otimes\}$;
- $B[G/X] \in \mathrm{FL}(F)$ if $G = \sigma X.B$ for $\sigma \in \{\nu,\mu\}$.

We say that G is a *sub-occurrence* of F if $G \in \mathrm{FL}(F)$. Note that, for any F and α, there is at most one φ such that φ_α is a sub-occurrence of F.

We are now ready to introduce our infinitary sequent calculus. Details regarding formula occurrences can be ignored at first read, and will only make full sense when one starts permuting inferences and eliminating cuts.

▶ **Definition 7.** A *sequent*, written $\vdash \Gamma$, is a finite set of pairwise disjoint formula occurrences. A *pre-proof* of $\mu\mathsf{MALL}^\infty$ is a possibly infinite tree, coinductively generated by the rules of Figure 1, subject to the following conditions: any two formulas occurrences appearing in different branches must be disjoint except if the branches first differ right after a $(\&)$ inference; if φ_α and $\psi_{\alpha^\perp}$ occur in a pre-proof, they must be the respective sub-occurrences of the formula occurrences F and $F^\perp$ introduced by a (Cut) rule.

The disjointness condition on sequents ensures that two formula occurrences from the same sequent will never engender a common sub-occurrence, *i.e.*, we can define traces uniquely. The disjointness condition on pre-proofs is there to ensure that the proof transformations used in focusing and cut elimination preserve the disjointness condition on sequents. Note that these conditions are not restrictive. Clearly, the condition on sequents never prevents the (backwards) application of a propositional rule. Moreover, there is an infinite supply of disjoint addresses, *e.g.*, $\{\, r^n l \;:\; n > 0 \,\}$. One may thus pick addresses from that supply for the conclusion sequent of the derivation, and then carry the remaining supply along proof branches, splitting it on branching rules, and consuming a new address for cut rules.

Pre-proofs are obviously unsound: the pre-proof schema shown below allows one to derive any formula. In order to obtain proper proofs from pre-proofs, we will add a validity condition. This condition will reflect the nature of our two fixed point connectives.

$$\dfrac{\dfrac{\vdots}{\vdash \mu X.X}\,(\mu) \qquad \dfrac{\vdots}{\vdash \nu X.X, F}\,(\nu)}{\vdash F}\,(\mathsf{Cut})$$

▶ **Definition 8.** Let $\gamma = (s_i)_{i\in\omega}$ be an infinite branch in a pre-proof of $\mu\mathsf{MALL}^\infty$. A *thread* t in γ is a sequence of formula occurrences $(F_i)_{i\in\omega}$ with $F_i \in s_i$ and $F_i \sqsubseteq F_{i+1}$. The set of formulas that occur infinitely often in $(F_i)_{i\in\omega}$ (when forgetting addresses) admits a minimum wrt. the subformula ordering, denoted by $\min(t)$. A thread t is *valid* if $\min(t)$ is a ν formula and the thread is not eventually constant, *i.e.*, the formulas F_i are always eventually principal.

▶ **Definition 9.** The *proofs* of $\mu\mathsf{MALL}^\infty$ are those pre-proofs in which every infinite branch contains a valid thread.

This validity condition has its roots in parity games and is very natural for infinitary proof systems with fixed points. It is somehow independent of the ambient logic, and only deals with fixed points. It is commonly found in deductive systems for modal μ-calculi: see [11] for a closely related presentation, which yields a sound and complete sequent calculus for linear time μ-calculus. The validity conditions of Santocanale's circular proofs [22, 13], with and without cut, are also instances of the above notion.

In the rest of the paper, we work mostly with formula occurrences and will often simply call them formulas when it is not ambiguous. As usual in sequent calculus, (Ax) on a formula F can be expanded into axioms on its immediate subformulas. Repeating this process, one obtains an axiom-free and valid proof of the original sequent. In fact, this construction yields a *regular* derivation tree, the simplest kind of finitely representable infinite derivation.

▶ **Proposition 10.** *Rule* (Ax) *is admissible in* $\mu\mathsf{MALL}^\infty$.

This basic observation, proved in [6], justifies that the (Ax) rule will be ignored in the rest of the paper. In particular, we consider that axioms are expanded away before dealing with cut elimination. Our system $\mu\mathsf{MALL}^\infty$ is naturally equipped with the cut elimination rules of MALL, extended with the obvious principal and auxiliary rules for fixed point connectives (we do not show symmetric cases):

$$\dfrac{\dfrac{\vdash \Gamma, F[\mu X.F/X]}{\vdash \Gamma, \mu X.F}\,(\mu) \qquad \dfrac{\vdash F^\perp[\nu X.F^\perp/X], \Delta}{\vdash \nu X.F^\perp, \Delta}\,(\nu)}{\vdash \Gamma, \Delta}\,(\mathsf{Cut})$$
$$\downarrow$$
$$\dfrac{\vdash \Gamma, F[\mu X.F/X] \qquad \vdash F^\perp[\nu X.F^\perp/X], \Delta}{\vdash \Gamma, \Delta}\,(\mathsf{Cut})$$

$$\dfrac{\dfrac{\vdash \Gamma, F[\mu X.F/X], G}{\vdash \Gamma, \mu X.F, G}\,(\mu) \qquad \vdash G^\perp, \Delta}{\vdash \Gamma, \mu X.F, \Delta}\,(\mathsf{Cut})$$
$$\downarrow$$
$$\dfrac{\dfrac{\vdash \Gamma, F[\mu X.F/X], G \qquad \vdash G^\perp, \Delta}{\vdash \Gamma, F[\mu X.F/X], \Delta}\,(\mathsf{Cut})}{\vdash \Gamma, \mu X.F, \Delta}\,(\mu)$$

Natural numbers may be expressed as $\varphi_{\mathsf{nat}} := \mu X.\mathbf{1}\oplus X$. Occurrences of that formula will be denoted N, N', etc. We give below a few examples of proofs/computations on natural numbers, shown using two sided sequents for clarity: $F_1, \ldots, F_n \vdash \Gamma$ should be read as

$\vdash \Gamma, F_1^\perp, \ldots, F_n^\perp$ as usual. The proof π_{succ}, shown below, computes the successor on natural numbers: if we cut it against a (necessarily finite) cut-free proof of N we obtain after a finite number of cut elimination steps a proof of N' which is the right injection (rule (μ) followed by $(\oplus_2)$, which represents the successor) of the original proof of N, relocated at the address of N''.

$$\dfrac{\dfrac{\dfrac{}{N \vdash N''}\,(\mathsf{Ax})}{N \vdash \mathbf{1} \oplus N''}\,(\oplus_2)}{N \vdash N'}\,(\mu)$$

Consider now the following pre-proof, called π_{dup}:

$$\dfrac{\dfrac{\dfrac{}{\vdash N_1}\,(\mu),(\oplus_1),(\mathbf{1}) \qquad \dfrac{}{\vdash N_2}\,(\mu),(\oplus_1),(\mathbf{1})}{\mathbf{1} \vdash N_1 \otimes N_2}\,(\bot),(\otimes) \qquad \dfrac{\dfrac{(\star)}{N' \vdash N_1' \otimes N_2'} \qquad \dfrac{\pi_{\mathsf{succ}} \quad \pi_{\mathsf{succ}}}{N_1' \otimes N_2' \vdash N_1 \otimes N_2}\,(\parr),(\otimes)}{N' \vdash N_1 \otimes N_2}\,(\mathsf{Cut})}{(\star) \qquad N \vdash N_1 \otimes N_2}\,(\nu),(\&)$$

Here, $(\star)$ represents the cyclic repetition of the same proof, on a structurally equivalent sequent (same formulas, new addresses). The resulting pre-proof has exactly one infinite branch, validated by the thread starting with N. If we cut that proof against an arbitrary cut-free proof of N, and perform cut elimination steps, we obtain in finite time a cut-free proof of $N_1 \otimes N_2$ which consists of two copies (up-to addresses) of the original proof of N.

Now let $\varphi_{\mathsf{stream}} = \nu X.\varphi_{\mathsf{nat}} \otimes X$ be the formula representing infinite streams of natural numbers, whose occurrences will be denoted by S, S', etc. Let us consider the derivation shown below, where F is an arbitrary, useless formula occurrence for illustrative purposes.

$$\dfrac{\dfrac{\dfrac{\pi_{\mathsf{dup}}}{N \vdash N_1 \otimes N_2} \qquad \dfrac{\dfrac{\dfrac{}{N_1 \vdash N'}\,(\mathsf{Ax}) \qquad \dfrac{\dfrac{\pi_{\mathsf{succ}}}{N_2 \vdash N''} \qquad \dfrac{(\star)}{N'', F \vdash S'}}{N_2, F \vdash S'}\,(\mathsf{Cut})}{N_1, N_2, F \vdash N' \otimes S'}}{N_1 \otimes N_2, F \vdash N' \otimes S'}}{N, F \vdash N' \otimes S'}\,(\mathsf{Cut})}{(\star) \qquad N, F \vdash S}$$

It is a valid proof thanks to the thread on S. By cut elimination, the computational behaviour of that proof is to take a natural number n, and some irrelevant f, and compute the stream $n :: (n+1) :: (n+2) :: \ldots$. However, unlike in the two previous examples, the result of the computation is not obtained in finite time; instead, we are faced with a productive process which will produce any finite prefix of the stream when given enough time. The presence of the useless formula F illustrates here that weakening may be admissible in $\mu\mathsf{MALL}^\infty$ under some circumstances, and that cutting against some formulas (F in this case) will form a redex that will be delayed forever. These subtleties will show up in the next two sections, devoted to showing our two main results.

3 Focalization

Focalization in linear logic. MALL connectives can be split in two classes: *positive* $(\otimes, \oplus, \mathbf{0}, \mathbf{1})$ and *negative* $(\parr, \&, \top, \bot)$ connectives. The distinction can be easily understood in terms of proof search: negative inferences $(\parr), (\&), (\top)$ and $(\bot)$ are *reversible* (meaning that provability of the conclusion transfers to the premisses) while positive inferences require choices (splitting the context in $(\otimes)$ or choosing between $(\oplus_1)$ and $(\oplus_2)$ rules) resulting in a possible of loss of provability. Still, positive inferences satisfy the *focalization* property [3]: in any provable sequent containing no negative formula, some formula can be chosen as a *focus*, hereditarily selecting its positive subformulas as principal formulas until a negative subformula is reached. It induces the following complete proof search strategy:

Sequent Γ *contains a negative* formula	Sequent Γ *contains no negative* formula
Choose any negative formula (*e.g.*, the leftmost one) and decompose it using the only possible negative rule.	Choose some positive formula and decompose it (and its subformulas) hereditarily until we get to atoms or negative subformulas.

Focalization graphs. Focused proofs are complete for proofs, not only provability: any linear proof is equivalent to a focused proof, up to cut-elimination. Indeed, focalization can be proved by means of proof transformations [18, 20, 7] preserving the denotation of the proof. A flexible, modular method for proving focalization that we shall apply in the next sections has been introduced by Miller and the third author [20] and relies on *focalization graphs*. The heart of the focalization graph proof technique relies on the fact the positive inference, while not reversible, all permute with each other. As a consequence, if the positive layer of some positive formula is completely decomposed within the lowest part of the proof, below any negative inference, then it can be taken as a focus. Focalization graphs ensure that it is always possible: their acyclicity provides a source which can be taken as a focus.

Focusing infinitary proofs. The infinitary nature of our proofs interferes with focalization in several ways. First, while in $\mu\mathsf{MALL}$ μ and ν can be set to have an arbitrary polarity, we will see that in $\mu\mathsf{MALL}^\infty$, ν must be negative. Second, permutation properties of the negative inferences, which can be treated locally in $\mu\mathsf{MALL}$, now require a global treatment due to infinite branches. Last, focalization graphs strongly rely on the finiteness of maximal positive subtrees of a proof: this invariant must be preserved in $\mu\mathsf{MALL}^\infty$. For simplicity reasons, we restrict our attention to cut-free proofs in the rest of this section. The result holds for proofs with cuts thanks to the usual trick of viewing cuts as $\otimes$.

3.1 Polarity of connectives

Let us first consider the question of polarizing $\mu\mathsf{MALL}^\infty$ connectives. Unlike in $\mu\mathsf{MALL}$, we are not free to set the polarity of fixed points formulas: consider the proof π of sequent $\vdash \mu X.X, \nu Y.Y$ which alternates inferences (ν) and (μ). Assigning opposite polarities to dual formulas (an invariant necessary to define properly cut-elimination in focused proof systems), this sequent contains a negative formula; each polarization of fixed points induces one focused pre-proof, either π_μ which always unrolls μ or π_ν which repeatedly unrolls ν. Only π_ν happens to be valid, leaving but one possible choice, $\nu X.F$ negative and $\mu X.F$ positive, resulting in the following polarization:

▶ **Definition 11.** *Negative formulas* are formulas of the form $\nu X.F$, $F \parr G$, $F \& G$, $\bot$ and $\top$, *positive formulas* are formulas of the form $\mu X.F$, $F \otimes G$, $F \oplus G$, $\mathbf{1}$ and $\mathbf{0}$. A $\mu\mathsf{MALL}^\infty$ sequent containing only positive formulas is said to be *positive*. Otherwise, it is *negative*.

The following proposition will be useful in the following:

▶ **Proposition 12.** *An infinite branch of a pre-proof containing only negative (resp. positive) rules is always valid (resp. invalid).*

3.2 Reversibility of negative inferences

The example shown below with $F = \nu X.(X \& X) \oplus \mathbf{0}$ shows that, unlike in (MA)LL, negative inferences cannot be permuted down locally: no occurrence of a negative inference ($\parr$) on $P \parr Q$ can be permuted below a ($\&$) since it is never available in the left premise.

$$\dfrac{\dfrac{\dfrac{\dfrac{(\star)}{\vdash F, P \parr Q} \qquad \dfrac{\dfrac{\pi'}{\vdash F, P, Q}}{\vdash F, P \parr Q}\,(\parr)}{\vdash F \& F, P \parr Q}\,(\&)}{\vdash (F \& F) \oplus \mathbf{0}, P \parr Q}\,(\oplus_1)}{(\star)\ \vdash F, P \parr Q}\,(\nu)$$

We shall thus introduce a global proof transformation (which could be realized by means of cut, as is usual). In order to define this transformation at once for all negative connectives, we rely on the uniform structure of negative inferences, which can be written as follows:

$$\dfrac{(\vdash \Gamma, \mathcal{N}_i^N)_{1 \leq i \leq n}}{\vdash \Gamma, N}\,(\mathsf{r_N})$$

Sub-occurrence families of N are then defined as $\mathcal{N}(N) = (\mathcal{N}_i^N)_{1 \leq i \leq n}$, its *slicing index* being $\mathsf{sl}(N) = \#\mathcal{N}(N)$.

N	$F_1 \parr F_2$	$\bot$	$F_1 \& F_2$	$\top$	$\nu X.F$
$\mathcal{N}(N)$	$\{1 \mapsto \{F_1, F_2\}\}$	$\{1 \mapsto \emptyset\}$	$\{1 \mapsto \{F_1\}, 2 \mapsto \{F_2\}\}$	$\emptyset$	$\{1 \mapsto \{F[\nu X.F/X]\}\}$

We can now define, in two steps, how to transform any proof π into a proof $\mathsf{rev}(\pi)$ where all negative inferences are reversed.

▶ **Definition 13** ($\pi(i, N)$)**.** Let π be a proof of $\vdash \Gamma$ of last rule (r) and premises $\pi_1, \ldots, \pi_n$. If $1 \leq i \leq \mathsf{sl}(N)$, we define $\pi(i, N)$ coinductively:

- if N does not occur in $\vdash \Gamma$, then $\pi(i, N) = \pi$;
- if r is the inference on N, then $\pi(i, N) = \pi_i$ (which is legal since in this case $n = \mathsf{sl}(N)$);
- if r is not the inference on N, then

$$\pi(i, N) = \dfrac{\pi_1(i, N) \qquad \ldots \qquad \pi_n(i, N)}{\vdash \Gamma, \mathcal{N}_i^N}\,(\mathsf{r}).$$

▶ **Definition 14** ($\mathsf{rev}(\pi)$)**.** Let π be a $\mu\mathsf{MALL}^\infty$ proof of $\vdash \Gamma$. Then $\mathsf{rev}(\pi)$ is a pre-proof non-deterministically defined as π if $\vdash \Gamma$ is positive and, otherwise, when $N \in \Gamma$ and $n = \mathsf{sl}(N)$, as

$$\mathsf{rev}(\pi) = \dfrac{\mathsf{rev}(\pi(1, N)) \qquad \ldots \qquad \mathsf{rev}(\pi(n, N))}{\vdash \Gamma}\,(\mathsf{r_N}).$$

Reversed proofs formalize the requirement for the whole negative layer to be reversed:

▶ **Definition 15.** *Reversed pre-proofs* are defined to be the largest set of pre-proofs such that: (i) every pre-proof of a positive sequent is reversed; (ii) a pre-proof of a negative sequent is reversed if it ends with a negative inference and if each of its premises is reversed.

▶ **Example 16.** We illustrate rev on the proof π starting this subsection. We have $\mathsf{sl}(P \parr Q) = 1$ and:

$$\mathsf{rev}(\pi) \quad = \quad \dfrac{\pi(1, P \parr Q)}{\vdash F, P \parr Q}\,(\parr) \quad = \quad \dfrac{\dfrac{\dfrac{\dfrac{\dfrac{(\star)}{\vdash F, P, Q} \qquad \dfrac{\pi'}{\vdash F, P, Q}}{\vdash F \& F, P, Q}\,(\&)}{\vdash (F \& F) \oplus \mathbf{0}, P, Q}\,(\oplus_1)}{(\star)\ \vdash F, P, Q}\,(\nu)}{\vdash F, P \parr Q}\,(\parr)$$

▶ **Theorem 17.** *If π is a $\mu\mathsf{MALL}^\infty$ proof, then $\mathsf{rev}(\pi)$ is a reversed proof of the same sequent.*

3.3 Focalization Graph

In this section, we adapt the focalization graphs introduced in [20] to our setting. Considering the permutability properties of positive inferences in $\mu\mathsf{MALL}^\infty$, finiteness of positive trunks and acyclicity of focalization graphs will be sufficient to make the proof technique of [20] applicable. In order to illustrate this subsection, an example is fully explained in [6].

▶ **Definition 18** (Positive trunk, positive border, active formulas)**.** Let π be a $\mu\mathsf{MALL}^\infty$ proof of $\mathcal{S}$. The *positive trunk* π^+ of π is the tree obtained by cutting (finite or infinite) branches of π at the first occurrence of a negative rule. The *positive border* of π is the collection of lowest sequents in π which are conclusions of negative rules. *P-active* formulas of π are those formulas of $\mathcal{S}$ which are principal formulas of an inference in π^+.

▶ **Proposition 19.** *The positive trunk of a $\mu\mathsf{MALL}^\infty$ proof is always finite.*

▶ **Definition 20** (Focalization graph)**.** Given a $\mu\mathsf{MALL}^\infty$ proof π, we define its *focalization graph* $\mathcal{G}(\pi)$ to be the graph whose vertices are the P-active formulas of π and such that there is an edge from F to G iff there is a sequent $\mathcal{S}'$ in the positive border containing a negative sub-occurrence F' of F and a positive sub-occurrence G' of G.

$\mu\mathsf{MALL}^\infty$ positive inferences are those of MALL extended with (μ) which is not branching: this ensures both that any two positive inferences permute and that the proof of acyclicity of MALL focalization graphs can easily be adapted, from which we conclude that:

▶ **Proposition 21.** *Focalization graphs are acyclic.*

Acyclicity of the focalization graph implies in particular that it has a source, that is a formula P of the conclusion sequent such that whenever one of its subformulas F appears in a border sequent, F is negative. This remark, together with the fact that the trunk is finite ensures that the positive layer of P is completely decomposed in the positive trunk.

▶ **Definition 22** ($\mathsf{foc}(\pi, P)$)**.** Let π be a $\mu\mathsf{MALL}^\infty$ proof of $\vdash \Gamma, P$ with P a source of π's focalization graph. One defines $\mathsf{foc}(\pi, P)$ as the $\mu\mathsf{MALL}^\infty$ proof obtained by permuting down all the positive inferences on P and its positive subformulas (all occurring in π^+).

▶ **Proposition 23.** *Let $\mathcal{S}$ be a lowest sequent of $\mathsf{foc}(\pi, P)$ which is not the conclusion of a rule on a positive subformula of P. Then $\mathcal{S}$ contains exactly one subformula of P, which is negative.*

3.4 Productivity and validity of the focalization process

Reversibility of the negative inferences and focalization of the positive inferences allows one to consider the following (non-deterministic) proof transformation process:

Focalization Process: Let π be a μMALL$^\infty$ proof of $\mathcal{S}$. Define $\mathsf{Foc}(\pi)$ as follows:

- **Asynchronous phase:** If $\mathcal{S}$ is negative, transform π into $\mathsf{rev}(\pi)$ which is reversed. At least one negative inference has been brought to the root of the proof. Apply (corecursively) the synchronous phase to the proofs rooted in the lowest positive sequents of $\mathsf{rev}(\pi)$.
- **Synchronous phase:** If $\mathcal{S}$ is positive, let $P \in \mathcal{S}$ be a source of the associated focalization graph. Transform π into a proof $\mathsf{foc}(\pi, P)$. At least one positive inference on P has been brought to the root of the proof. Apply (corecursively) the asynchronous phase to the proofs rooted in the lowest negative sequents of $\mathsf{foc}(\pi, P)$.

Each of the above phases produces one non-empty phase, the above process is thus productive. It is actually a pre-proof thanks to Theorem 17 and by definition of $\mathsf{foc}(\pi, P)$. It remains to show that the resulting pre-proof is actually a proof. The following property is easily seen to be preserved by both transformations foc and rev and thus holds for $\mathsf{Foc}(\pi)$:

▶ **Proposition 24.** *Let π be a μMALL$^\infty$ proof, r a positive rule occurring in π and r' be a negative rule occurring below r in π. If r occurs in $\mathsf{Foc}(\pi)$, then r' occurs in $\mathsf{Foc}(\pi)$, below r.*

▶ **Lemma 25.** *For any infinite branch γ of $\mathsf{Foc}(\pi)$ containing an infinite number of positive rules, there exists an infinite branch in π containing infinitely many positive rules of γ.*

▶ **Theorem 26.** *If π is a μMALL$^\infty$ proof then $\mathsf{Foc}(\pi)$ is also a μMALL$^\infty$ proof.*

Proof sketch. An infinite branch γ of $\mathsf{Foc}(\pi)$ may either be obtained by reversibility only after a certain point, or by alternating infinitely often synchronous and asynchronous phases. In the first case it is valid by Proposition 12 while in the latter case, Lemma 25 ensures the existence of a branch δ of π containing infinitely many positive rules of γ, with a valid thread t of minimal formula F_m: every rule r of δ in which F_m is principal is below a positive rule occurring in γ. Thus r occurs in γ, which is therefore valid. ◀

4 Cut elimination

In this section, we show that any μMALL$^\infty$ proof can be transformed into an equivalent cut-free derivation. This is done by applying the cut reduction rules described in Section 2, possibly in infinite reductions converging to cut-free proofs. As usual with infinitary reductions it is not the case that any reduction sequence converges: for instance, one could reduce only deep cuts in a proof, leaving a cut untouched at the root. We avoid this problem by considering a form of head reduction where we only reduce cuts at the root.

Cut reduction rules are of two kinds, *principal* reductions and *auxiliary* ones. In the infinitary setting, principal cut reductions do not immediately contribute to producing a cut-free pre-proof. On the contrary, auxiliary cut reductions are productive in that sense. In other words, principal rules are seen as internal computations of the cut elimination process, while auxiliary rules are seen as a partial output of that process. Accordingly, the former will be called *internal rules* and the latter *external rules*.

When analyzing cut reductions, cut commutations can be troublesome. A common way to avoid this technicality [13], which we shall follow, is to introduce a *multicut* rule which merges multiple cuts, avoiding cut commutations.

$$\frac{s_1 \dots s_n}{s}\ (\mathsf{mcut})$$

▶ **Definition 27.** Given two sequents s and s', we say that they are cut-connected on a formula occurrence F when $F \in s$ and $F^\perp \in s'$. We say that they are cut-connected when they are cut-connected for some F. We define the *multicut* rule as shown above with

$$\dfrac{\dfrac{\vdash \Gamma, F \qquad \vdash F^{\perp}, \Delta}{\vdash \Gamma, \Delta}\,(\mathsf{Cut}) \qquad \dots}{\vdash \Sigma}\,(\mathsf{mcut}) \quad\longrightarrow\quad \dfrac{\vdash \Gamma, F \qquad \vdash F^{\perp}, \Delta \qquad \dots}{\vdash \Sigma}\,(\mathsf{mcut})$$

$$\dfrac{\dfrac{\vdash \Gamma, F}{\vdash \Gamma, F \oplus G} \qquad \dfrac{\vdash G^{\perp}, \Delta \qquad \vdash F^{\perp}, \Delta}{\vdash G^{\perp} \& F^{\perp}, \Delta} \qquad \dots}{\vdash \Sigma}\,(\mathsf{mcut}) \quad\longrightarrow\quad \dfrac{\vdash \Gamma, F \qquad \vdash F^{\perp}, \Delta \qquad \dots}{\vdash \Sigma}\,(\mathsf{mcut})$$

$$\dfrac{s_1 \dots s_n \qquad \dfrac{\vdash \Gamma, F \qquad \vdash \Gamma, G}{\vdash \Gamma, F \& G}\,(\&)}{\vdash \Sigma, F \& G}\,(\mathsf{mcut}) \quad\longrightarrow\quad \dfrac{\dfrac{s_1 \dots s_n \qquad \vdash \Gamma, F}{\vdash \Sigma, F}\,(\mathsf{mcut}) \qquad \dfrac{s_1 \dots s_n \qquad \vdash \Gamma, G}{\vdash \Sigma, G}\,(\mathsf{mcut})}{\vdash \Sigma, F \& G}\,(\&)$$

Figure 2 (Cut)/(mcut) and $(\oplus_1)/(\&)$ internal reductions and $(\&)$/(mcut) external reduction.

conclusion s and premisses $\{s_i\}_i$, where the set $\{s_i\}_i$ is connected and acyclic with respect to the cut-connection relation, and s is the set of all formula occurrences F that appear in some s_i but such that no s_j is cut-connected to s_j on F.

From now on we shall work with $\mu\mathsf{MALL}^{\infty}_{\mathsf{m}}$ derivations, which are $\mu\mathsf{MALL}^{\infty}$ derivations in which the multicut rule may occur, though only at most once per branch. The notions of thread and validity are unchanged. In $\mu\mathsf{MALL}^{\infty}_{\mathsf{m}}$ we only reduce multicuts, in a way that is naturally obtained from the cut reductions of $\mu\mathsf{MALL}^{\infty}$. A complete description of the rules is given in [6]; only the (Cut)/(mcut) and $(\oplus_1)/(\&)$ internal reduction cases and the $(\&)$/(mcut) external reduction case are shown in Figure 2. As is visible in the last reduction, applying an external rule on a multicut may yield multiple multicuts, though always on disjoint subtrees.

We will be interested in a particular kind of multicut reduction sequences, the *fair* ones, which are such that any redex which is available at some point of the sequence will eventually have disappeared from the sequence (being reduced or erased), details are provided in [6]. We will establish that these reductions eliminate multicuts:

▶ **Theorem 28.** *Fair multicut reduction sequences on $\mu\mathsf{MALL}^{\infty}_{\mathsf{m}}$ proofs produce $\mu\mathsf{MALL}^{\infty}$ proofs.*

Additionally, if all cuts in the initial derivation are above multicuts, the resulting $\mu\mathsf{MALL}^{\infty}$ derivation must actually be cut-free: indeed, multicut reductions never produce a cut. Thus Theorem 28 gives a way to eliminate cuts from any $\mu\mathsf{MALL}^{\infty}$ proof π of $\vdash \Gamma$ by forming a multicut with conclusion $\vdash \Gamma$ and π as unique subderivation, and eliminating multicuts (and cuts) from that $\mu\mathsf{MALL}^{\infty}_{\mathsf{m}}$ proof. The proof of Theorem 28 is in two parts. We first prove that fair internal multicut reductions cannot diverge (Proposition 40), hence fair multicut reductions are productive, *i.e.*, reductions of $\mu\mathsf{MALL}^{\infty}_{\mathsf{m}}$ proofs converge to $\mu\mathsf{MALL}^{\infty}$ pre-proofs. We then establish that the obtained pre-proof is a valid proof (Proposition 41).

Regarding productivity, assuming that there exists an infinite sequence σ of internal cut-reductions from a given proof π of Γ, we obtain a contradiction by extracting from π a proof of the empty sequent in a suitably defined proof-system. More specifically, we observe that no formula of Γ is principal in the subtree π_σ of π visited by σ. Hence, by erasing every formula of Γ from π_σ, local correctness of the proof is preserved, resulting in a tree deriving the empty sequent. This tree can be viewed as a proof in a new proof-system $\mu\mathsf{MALL}^{\infty}_{\tau}$ which is shown to be sound (Proposition 37) with respect to the traditional Boolean semantics of the μ-calculus, thus the contradiction. The proof of validity of the produced pre-proof is similar: instead of extracting a proof of the empty sequent from π we will extract, for each invalid branch of π, a $\mu\mathsf{MALL}^{\infty}_{\tau}$ proof of a formula containing neither $\mathbf{1}$, $\top$, nor ν formulas, contradicting soundness again.

4.1 Extracting proofs from reduction paths

We define now a key notion to analyze the behaviour of multicut-elimination: given a multicut reduction starting from π, we extract a (slightly modified) subderivation of π which corresponds to the part of the derivation that has been explored by the reduction. More precisely, we are interested in *reduction paths* which are sequences of proofs that end with a multicut rule, obtained by tracing one multicut through its evolution, selecting only one sibling in the case of $(\&)$ and $(\otimes)$ external reductions. Given such a reduction path starting with π, we consider the subtree of π whose sequents occur in the reduction path as premises of some multicut. This subtree is obviously not always a $\mu\mathsf{MALL}^{\infty}$ derivation since some of its nodes may have missing premises. We will provide an extension of $\mu\mathsf{MALL}^{\infty}$ where these trees can be viewed as proper derivations by first characterizing when this situation arises.

▶ **Definition 29** (Useless sequents, distinguished formula)**.** Let $\mathcal{R}$ be a reduction path starting with π. A sequent $s = (\vdash \Gamma, F)$ of π is said to be *useless* with *distinguished formula* F when in one of the following cases:

1. The sequent eventually occurs as a premise of all multicuts of $\mathcal{R}$ and F is the principal formula of s in π. (Note that the distinguished formula F of a useless sequent s of sort (1) must be a sub-occurrence of a cut formula in π. Otherwise, the fair reduction path $\mathcal{R}$ would eventually have applied an external rule on s. Moreover, $F^{\perp}$ never becomes principal in the reduction path, otherwise by fairness the internal rule reducing F and $F^{\perp}$ would have been applied.)
2. At some point in the reduction, the sequent is a premise of $(\&)$ on $F \& F'$ or $F' \& F$ which is erased in an internal $(\&)/(\oplus)$ multicut reduction. (In the $(\oplus_1)/(\&)$ internal reduction of Figure 2, the sequent $\vdash G^{\perp}, \Delta$ is useless of sort (2).)
3. The sequent is ignored at some point in the reduction path because it is not present in the selected multicut after a branching external reduction on $F \star F'$ or $F' \star F$, for $\star \in \{\otimes, \&\}$. (In the $(\&)/(\mathsf{mcut})$ external reduction of Figure 2, if one is considering a reduction path that follows the multicut having $\vdash \Gamma, F$ as a premise, then the sequent $\vdash \Gamma, G$ is useless of sort (3), and vice versa.)
4. The sequent is ignored at some point in the reduction path because a $(\otimes)/(\mathsf{mcut})$ external reduction distributes s to the multicut that is not selected in the path. This case will be illustrated next, and is described in full details in [6].

Note that, although the external reduction for $\top$ erases sequents, we do not need to consider such sequents as useless: indeed, we will only need to work with useless sequents in infinite reduction paths, and the external reduction associated to $\top$ terminates a path.

▶ **Example 30.** Consider a multicut composed of the last example of Section 2 and an arbitrary proof of $\vdash F, \Delta$ where F is principal. In the reduction paths which always select the right premise of an external $(\otimes)/(\mathsf{mcut})$ corresponding to the $N' \otimes S'$ formulas, the sequent $\vdash F, \Delta$ will always be present and thus useless by case (1). In the reduction paths which eventually select a left premise, the sequent $N_2, F \vdash S'$ is useless of sort (3) with S' distinguished, and $\vdash F, \Delta$ is useless of sort (4) with F distinguished.

In order to obtain a proper pre-proof from the sequents occurring in a reduction path, we need to close the derivation on useless sequents. This is done by replacing distinguished formulas by $\top$ formulas. However, a usual substitution is not appropriate here as we are really replacing formula occurrence, which may be distributed in arbitrarily complex ways among sub-occurrences.

▶ **Definition 31.** A *truncation* τ is a partial function from Σ^* to $\{\top, \mathbf{0}\}$ such that:

- For any $\alpha \in \Sigma^*$, if $\alpha \in \texttt{Dom}(\tau)$, then $\alpha^\perp \in \texttt{Dom}(\tau)$ and $\tau(\alpha) = \tau(\alpha^\perp)^\perp$.
- If $\alpha \in \texttt{Dom}(\tau)$ then for any $\beta \in \Sigma^+$, $\alpha.\beta \notin \texttt{Dom}(\tau)$.

▶ **Definition 32** (Truncation of a reduction path)**.** Let $\mathcal{R}$ be a reduction path. The truncation τ associated to $\mathcal{R}$ is defined by setting $\tau(\alpha) = \top$ and $\tau(\alpha^\perp) = \mathbf{0}$ for every formula occurrence φ_α that is distinguished in some useless sequent of $\mathcal{R}$.

The above definition is justified because F and $F^\perp$ cannot both be distinguished, by fairness of $\mathcal{R}$. We can finally obtain the pre-proof associated to a reduction path, in a proof system slightly modified to take truncations into account.

▶ **Definition 33** (Truncated proof system)**.** Given a truncation τ, the infinitary proof system $\mu\mathsf{MALL}^\infty_\tau$ is obtained by taking all the rules of $\mu\mathsf{MALL}^\infty$, with the proviso that they only apply when the address of their principal formula is not in the domain of τ, with the following extra rule:

$$\frac{\vdash \tau(\alpha)_{\alpha i}, \Delta}{\vdash F, \Delta}\ (\tau) \qquad \text{if } \alpha \in \texttt{Dom}(\tau)$$

The address $\alpha.i$ associated with $\tau(\alpha)$ in the rule (τ) forbids loops on a (τ) rule. Indeed if $\alpha \in \texttt{Dom}(\tau)$ then $\alpha.i \notin \texttt{Dom}(\tau)$.

▶ **Definition 34** (Truncated proof associated to a reduction path)**.** Let $\mathcal{R}$ be a fair infinite reduction path starting with π and τ be the truncation associated to it. We define $TR(\mathcal{R})$ to be the $\mu\mathsf{MALL}^\infty_\tau$ proof obtained from π by keeping only sequents that occur as premise of some multicut in $\mathcal{R}$, using the same rules as in π whenever possible, and deriving useless sequents by rules (τ) and $(\top)$.

This definition is justified by definition of τ and because only useless sequents may be selected without their premises (in π) being also selected. Notice that the dual $F^\perp$ of a distinguished formula F may only occur in $\mathcal{R}$ for distinguished formulas of type (1) and (4); in these cases $F^\perp$ is never principal in $\mathcal{R}$ by fairness. Thus, there is no difficulty in constructing $TR(\mathcal{R})$ with a truncation defined on the address of $F^\perp$. Finally, note that $TR(\mathcal{R})$ is indeed a valid $\mu\mathsf{MALL}^\infty_\tau$ pre-proof, because its infinite branches are infinite branches of π.

▶ **Example 35.** Continuing the previous example, we consider the path where the left premise of the tensor is selected immediately. The associated truncation is such that $\tau(S') = \top$ and $\tau(F) = \top$ by (3) and (4) respectively. The derivation $TR(\mathcal{R})$ is shown below, where Π_{ax} denotes the expansion of the axiom given by Proposition 10:

$$\dfrac{\dfrac{}{\vdash F, \Delta}\,(\tau),(\top) \qquad \dfrac{\dfrac{\dfrac{\Pi_{\mathsf{dup}}}{N \vdash N_1 \otimes N_2} \qquad \dfrac{\dfrac{\dfrac{\Pi_{\mathsf{ax}}}{N_1 \vdash N'} \qquad \dfrac{}{N_2, F \vdash S'}\,(\tau),(\top)}{N_1, N_2, F \vdash N' \otimes S'}}{N_1 \otimes N_2, F \vdash N' \otimes S'}}{N, F \vdash N' \otimes S'}\,(\mathsf{Cut})}{N, F \vdash S}}{N \vdash S, \Delta}\,(\mathsf{mcut})$$

4.2 Truncated truth semantics

We fix a truncation τ and define a truth semantics with respect to which $\mu\mathsf{MALL}^{\infty}_{\tau}$ will be sound. The semantics is classical, assigning a Boolean value to formula occurrences. For convenience, we take $\mathcal{B} = \{\mathbf{0}, \top\}$ as our Boolean lattice, with $\wedge$ and $\vee$ being the usual meet and join operations on it. The following definition provides an interpretation of $\mu\mathsf{MALL}$ formulas which consists in the composition of the standard interpretation of μ-calculus formulas with the obvious linearity-forgetting translation from $\mu\mathsf{MALL}$ to classical μ-calculus.

▶ **Definition 36.** Let φ_α be a pre-formula occurrence. We call *environment* any function $\mathcal{E}$ mapping free variables of φ to (total) functions of $E := \Sigma^* \to \mathcal{B}$. We define $[\varphi_\alpha]^{\mathcal{E}} \in \mathcal{B}$, the *interpretation* of φ_α in the environment $\mathcal{E}$, by $[\varphi_\alpha]^{\mathcal{E}} = \tau(\alpha)$ if $\alpha \in \mathtt{Dom}(\tau)$, and otherwise:

- $[X_\alpha]^{\mathcal{E}} = \mathcal{E}(X)(\alpha)$, $[\top_\alpha]^{\mathcal{E}} = [\mathbf{1}_\alpha]^{\mathcal{E}} = \top$ and $[\mathbf{0}_\alpha]^{\mathcal{E}} = [\bot_\alpha]^{\mathcal{E}} = \mathbf{0}$.
- $[(\varphi \owedge \psi)_\alpha]^{\mathcal{E}} = [\varphi_{\alpha.l}]^{\mathcal{E}} \wedge [\psi_{\alpha.r}]^{\mathcal{E}}$, for $\owedge \in \{\&, \otimes\}$.
- $[(\varphi \ovee \psi)_\alpha]^{\mathcal{E}} = [\varphi_{\alpha.l}]^{\mathcal{E}} \vee [\psi_{\alpha.r}]^{\mathcal{E}}$, for $\ovee \in \{\oplus, \parr\}$.
- $[(\mu X.\varphi)_\alpha]^{\mathcal{E}} = \mathrm{lfp}(f)(\alpha)$ and $[(\nu X.\varphi)_\alpha]^{\mathcal{E}} = \mathrm{gfp}(f)(\alpha)$ where $f : E \to E$ is given by $f : h \mapsto \beta \mapsto \big(\tau(\beta)$ if $\beta \in \mathtt{Dom}(\tau)$ and $[\varphi_{\beta.i}]^{\mathcal{E}::X\mapsto h}$ otherwise$\big)$.

When F is closed, we simply write $[F]$ for $[F]^{\emptyset}$.

We refer the reader to the long version [6] for details on the construction of the interpretation. We simply state here the main result about it.

▶ **Proposition 37.** *If* $\vdash \Gamma$ *is provable in* $\mu\mathsf{MALL}^{\infty}_{\tau}$*, then* $[F] = \top$ *for some* $F \in \Gamma$.

We only sketch the soundness proof (see [6] for details) which proceeds by contradiction. Assuming we are given a proof π of a formula F such that $[F] = \mathbf{0}$, we exhibit a branch β of π containing only formulas interpreted by $\mathbf{0}$. A validating thread of β unfolds infinitely often some formula $\nu X.\varphi$. Since the interpretation of $\nu X.\varphi$ is defined as the gfp of a monotonic operator f we have, for each occurrence $(\nu X.\varphi)_\alpha$ in β, an ordinal λ such that $[(\nu X.\varphi)_\alpha] = f^\lambda(\bigvee E)(\alpha)$, where $\bigvee E$ is the supremum of the complete lattice E. We show that this ordinal can be forced to decrease along β at each fixed point unfolding, contradicting the well-foundedness of the class of ordinals.

▶ **Definition 38.** A truncation τ is *compatible* with a formula φ_α if $\alpha \notin dom(\tau)$ and, for any $\alpha \sqsubseteq \beta.d \in \mathtt{Dom}(\tau)$ where $d \in \{l, r, i\}$, we have that φ_α admits a sub-occurrence ψ_β with $\otimes$ or $\&$ as the toplevel connective of ψ, $d \in \{l, r\}$, and $\alpha.d' \notin \mathtt{Dom}(\tau)$ for any $d \neq d'$.

In other words, a truncation τ is compatible with a formula F if it truncates only sons of $\otimes$ or $\&$ nodes in the tree of the formula F and at most one son of each such node.

▶ **Proposition 39.** *If* F *is a formula compatible with* τ *and containing no* ν *binders, no* $\top$ *and no* $\mathbf{1}$*, then* $[F] = \mathbf{0}$.

4.3 Proof of cut elimination

We first show that multicut reduction is productive, then that the resulting (cut-free) pre-proof is actually a valid proof.

▶ **Proposition 40.** *Any fair reduction sequence produces a* $\mu\mathsf{MALL}^{\infty}$ *pre-proof.*

Proof. By contradiction, consider a fair infinite sequence of internal multicut reductions. This sequence is a fair reduction path $\mathcal{R}$. Let τ and $TR(\mathcal{R})$ be the associated truncations and truncated proof. Since no external reduction occurs, it means that conclusion formulas of $TR(\mathcal{R})$ are never principal in the proof, thus we can transform it into a proof of the empty sequent, which contradicts soundness of $\mu\mathsf{MALL}^{\infty}_{\tau}$. ◀

▶ **Proposition 41.** *Any fair mcut-reduction produces a μMALL$^\infty$ proof.*

Proof. Let π be a $\mu\mathsf{MALL}^\infty_\mathsf{m}$ proof of conclusion $\vdash \Gamma$, and π' the cut-free pre-proof obtained by Proposition 40, *i.e.*, the limit of the multicut reduction process. Any branch of π' corresponds to a multicut reduction path. For the sake of contradiction, assume that π' is invalid. It must thus have an invalid infinite branch, corresponding to an infinite reduction path $\mathcal{R}$. Let τ and $\theta := TR(\mathcal{R})$ be the associated truncation and truncated proof in $\mu\mathsf{MALL}^\infty_\tau$.

We first observe that formulas of Γ cannot have suboccurrences of the form $\mathbf{1}_\alpha$ or $\top_\alpha$ that are principal in π'. Indeed, this could only be produced by an external rule $(\top)/(\mathsf{mcut})$ in the reduction path $\mathcal{R}$, but that would terminate the path, contradicting its infiniteness.

Next, we claim that all threads starting from formulas in Γ are invalid. Indeed, all rules applied to those formulas are transferred (by means of external rules) to the branch produced by the reduction path. The existence of a valid thread starting from the conclusion sequent in θ would thus imply the existence of a valid thread in our branch of π'.

By the first observation, we can replace all $\mathbf{1}$ and $\top$ subformulas of Γ by $\mathbf{0}$ without changing the derivation, and obviously without breaking its validity. By the second observation, we can further modify Γ by changing all ν combinators into μ combinators. The derivation is easily adapted (using rule (μ) instead of (ν)) and it remains valid, since the validity of θ could not have been caused by a valid thread starting from the root. We thus obtain a valid pre-proof θ' of $\vdash \Gamma'$ in $\mu\mathsf{MALL}^\infty_\tau$, where Γ' contains no ν, $\mathbf{1}$ and $\top$.

We finally show that τ is compatible with any formula occurrence from Γ. Indeed, if $\tau(\beta)$ is defined for some suboccurrence ψ_β of a formula $\varphi_\alpha \in \Gamma$, then it can only be because of a useless sequent of sort (3), *i.e.*, a truncation due to the fact that the reduction path has selected only one sibling after a branching external rule. We thus conclude, by Proposition 39, that all formulas of Γ are interpreted as $\mathbf{0}$ in the truncated semantics associated to τ, which contradicts the validity of θ' and Proposition 37. ◀

5 Conclusion

We have established focalization and cut elimination for μMALL$^\infty$, the infinitary sequent calculus corresponding to μMALL. Our cut elimination result extends that of Santocanale and Fortier [13], but this extension has required the elaboration of a radically different proof technique.

An obvious direction for future work is now to go beyond linear logic, and notably handle structural rules in infinitary cut elimination. But many interesting questions are also left in the linear case. First, it will be natural to relax the hypothesis on fairness in the cut-elimination result. Other than cut elimination, the other long standing problem regarding μMALL$^\infty$ and similar proof systems is whether regular proofs can be translated, in general, to finitary proofs. Further, one can ask the same question, requiring in addition that the computational content of proofs is preserved in the translation. It may well be that regular μMALL$^\infty$ contains more computations than μMALL; even more so if one considers other classes of finitely representable infinitary proofs. It would be interesting to study how this could impact the study of programming languages for (co)recursion, and understanding links with other approaches to this question [1, 2]. In this direction, we will be interested in studying the computational interpretation of focused cut-elimination, providing a logical basis for inductive and coinductive matching in regular and infinitary proof systems.

References

1 Andreas Abel and Brigitte Pientka. Wellfounded recursion with copatterns: a unified approach to termination and productivity. In Greg Morrisett and Tarmo Uustalu, editors, *ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA – September 25-27, 2013*, pages 185–196. ACM, 2013. doi:10.1145/2500365.2500591.

2 Andreas Abel, Brigitte Pientka, David Thibodeau, and Anton Setzer. Copatterns: programming infinite structures by observations. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL'13, Rome, Italy – January 23-25, 2013*, pages 27–38. ACM, 2013. doi:10.1145/2429069.2429075.

3 Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, 2(3):297–347, 1992.

4 David Baelde. Least and greatest fixed points in linear logic. *ACM Transactions on Computational Logic (TOCL)*, 13(1):2, 2012.

5 David Baelde, Amina Doumane, and Alexis Saurin. Least and greatest fixed points in ludics. In Stephan Kreutzer, editor, *24th EACSL Annual Conference on Computer Science Logic, CSL 2015, September 7-10, 2015, Berlin, Germany*, volume 41 of *LIPIcs*, pages 549–566. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2015. doi:10.4230/LIPIcs.CSL.2015.549.

6 David Baelde, Amina Doumane, and Alexis Saurin. Infinitary proof theory : the multiplicative additive case. *HAL*, hal-01339037, June 2016. URL: https://hal.archives-ouvertes.fr/hal-01339037.

7 Michele Basaldella, Alexis Saurin, and Kazushige Terui. On the meaning of focalization. In Alain Lecomte and Samuel Tronçon, editors, *Ludics, Dialogue and Interaction – PRELUDE Project – 2006-2009. Revised Selected Papers*, volume 6505 of *Lecture Notes in Computer Science*, pages 78–87. Springer, 2011. doi:10.1007/978-3-642-19211-1_5.

8 James Brotherston and Nikos Gorogiannis. Cyclic abduction of inductively defined safety and termination preconditions. In Markus Müller-Olm and Helmut Seidl, editors, *Static Analysis – 21st International Symposium, SAS 2014, Munich, Germany, September 11-13, 2014. Proceedings*, volume 8723 of *Lecture Notes in Computer Science*, pages 68–84. Springer, 2014. doi:10.1007/978-3-319-10936-7_5.

9 James Brotherston and Alex Simpson. Sequent calculi for induction and infinite descent. *Journal of Logic and Computation*, 21(6):1177–1216, December 2011.

10 Pierre-Louis Curien and Guillaume Munch-Maccagnoni. The duality of computation under focus. In Cristian S. Calude and Vladimiro Sassone, editors, *Theoretical Computer Science – 6th IFIP TC 1/WG 2.2 International Conference, TCS 2010, Held as Part of WCC 2010, Brisbane, Australia, September 20-23, 2010. Proceedings*, volume 323 of *IFIP Advances in Information and Communication Technology*, pages 165–181. Springer, 2010. doi:10.1007/978-3-642-15240-5_13.

11 Christian Dax, Martin Hofmann, and Martin Lange. A proof system for the linear time μ-calculus. In *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science, 26th International Conference, Kolkata, India, December 13-15, 2006, Proceedings*, pages 273–284, 2006. doi:10.1007/11944836_26.

12 Amina Doumane, David Baelde, Lucca Hirschi, and Alexis Saurin. Towards Completeness via Proof Search in the Linear Time mu-Calculus. Accepted for publication at LICS, January 2016. URL: https://hal.archives-ouvertes.fr/hal-01275289.

13 Jérôme Fortier and Luigi Santocanale. Cuts for circular proofs: semantics and cut-elimination. In Simona Ronchi Della Rocca, editor, *Computer Science Logic 2013 (CSL*

2013), CSL 2013, September 2-5, 2013, Torino, Italy, volume 23 of *LIPIcs*, pages 248–262. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2013. doi:10.4230/LIPIcs.CSL.2013.248.

14 David Janin and Igor Walukiewicz. Automata for the modal mu-calculus and related results. In Jirí Wiedermann and Petr Hájek, editors, *Mathematical Foundations of Computer Science 1995, 20th International Symposium, MFCS'95, Prague, Czech Republic, August 28 – September 1, 1995, Proceedings*, volume 969 of *Lecture Notes in Computer Science*, pages 552–562. Springer, 1995. doi:10.1007/3-540-60246-1_160.

15 Roope Kaivola. Axiomatising linear time mu-calculus. In *CONCUR'95: Concurrency Theory, 6th International Conference, Philadelphia, PA, USA, August 21-24, 1995, Proceedings*, pages 423–437, 1995. doi:10.1007/3-540-60218-6_32.

16 Dexter Kozen. Results on the propositional mu-calculus. *Theoretical Computer Science*, 27:333–354, 1983. doi:10.1016/0304-3975(82)90125-6.

17 Olivier Laurent. Polarized games. *Annals of Pure and Applied Logic*, 130(1-3):79–123, 2004. doi:10.1016/j.apal.2004.04.006.

18 Olivier Laurent. A proof of the focalization property of linear logic. Unpublished note, May 2004.

19 Paul-André Melliès and Nicolas Tabareau. Resource modalities in tensor logic. *Annals of Pure and Applied Logic*, 161(5):632–653, 2010. doi:10.1016/j.apal.2009.07.018.

20 Dale Miller and Alexis Saurin. From proofs to focused proofs: A modular proof of focalization in linear logic. In Jacques Duparc and Thomas A. Henzinger, editors, *Computer Science Logic, 21st International Workshop, CSL 2007, 16th Annual Conference of the EACSL, Lausanne, Switzerland, September 11-15, 2007, Proceedings*, volume 4646 of *Lecture Notes in Computer Science*, pages 405–419. Springer, 2007. doi:10.1007/978-3-540-74915-8_31.

21 Luigi Santocanale. μ-bicomplete categories and parity games. *ITA*, 36(2):195–227, 2002. doi:10.1051/ita:2002010.

22 Luigi Santocanale. A calculus of circular proofs and its categorical semantics. In Mogens Nielsen and Uffe Engberg, editors, *Foundations of Software Science and Computation Structures*, volume 2303 of *Lecture Notes in Computer Science*, pages 357–371. Springer, 2002. doi:10.1007/3-540-45931-6_25.

23 Luigi Santocanale. Free μ-lattices. *Journal of Pure and Applied Algebra*, 168(2–3):227–264, March 2002. URL: https://hal.archives-ouvertes.fr/hal-01261049, doi:10.1016/S0022-4049(01)00098-6.

24 Robert S. Streett and E. Allen Emerson. An automata theoretic decision procedure for the propositional mu-calculus. *Information and Computation*, 81(3):249–264, 1989. doi:10.1016/0890-5401(89)90031-X.

25 Kazushige Terui. Computational ludics. *Theoretical Computer Science*, 412(20):2048–2071, 2011. doi:10.1016/j.tcs.2010.12.026.

26 Igor Walukiewicz. On completeness of the mu-calculus. In *Proceedings of the Eighth Annual Symposium on Logic in Computer Science (LICS'93), Montreal, Canada, June 19-23, 1993*, pages 136–146. IEEE Computer Society, 1993. doi:10.1109/LICS.1993.287593.

27 Igor Walukiewicz. Completeness of Kozen's axiomatisation of the propositional mu-calculus. In *Proceedings, 10th Annual IEEE Symposium on Logic in Computer Science, San Diego, California, USA, June 26-29, 1995*, pages 14–24. IEEE Computer Society, 1995. doi:10.1109/LICS.1995.523240.

28 Igor Walukiewicz. Completeness of Kozen's axiomatisation of the propositional mu-calculus. *Information and Computation*, 157(1-2):142–182, 2000. doi:10.1006/inco.1999.2836.

29 Noam Zeilberger. *The logical basis of evaluation order and pattern matching.* PhD thesis, The logical basis of evaluation order and pattern matching, 2009.