An application of parallel cut elimination in unit-free multiplicative linear logic to the Taylor expansion of proof nets

Jules Chouquet, Lionel Vaux Auclair

An application of parallel cut elimination in unit-free multiplicative linear logic to the Taylor expansion of proof nets

Jules Chouquet
IRIF UMR 8243, Université Paris Diderot, Sorbonne Paris Cité, CNRS, France
Jules.Chouquet@irif.fr

Lionel Vaux Auclair
Aix-Marseille Univ, CNRS, Centrale Marseille, I2M, Marseille, France
lionel.vaux@univ-amu.fr
https://orcid.org/0000-0001-9466-418X

Abstract

We examine some combinatorial properties of parallel cut elimination in multiplicative linear logic (MLL) proof nets. We show that, provided we impose some constraint on switching paths, we can bound the size of all the nets satisfying this constraint and reducing to a fixed resultant net. This result gives a sufficient condition for an infinite weighted sum of nets to reduce into another sum of nets, while keeping coefficients finite. We moreover show that our constraints are stable under reduction.

Our approach is motivated by the quantitative semantics of linear logic: many models have been proposed, whose structure reflect the Taylor expansion of multiplicative exponential linear logic (MELL) proof nets into infinite sums of differential nets. In order to simulate one cut elimination step in MELL, it is necessary to reduce an arbitrary number of cuts in the differential nets of its Taylor expansion. It turns out our results apply to differential nets, because their cut elimination is essentially multiplicative. We moreover show that the set of differential nets that occur in the Taylor expansion of an MELL net automatically satisfy our constraints.

In the present work, we stick to the unit-free and weakening-free fragment of linear logic, which is rich enough to showcase our techniques, while allowing for a very simple kind of constraint: a bound on the number of cuts that are crossed by any switching path.

2012 ACM Subject Classification Theory of computation → Linear logic

Keywords and phrases linear logic; proof nets; cut elimination; differential linear logic

Digital Object Identifier 10.4230/LIPIcs.CSL.2018.14

Funding This work was supported by the French-Italian *Groupement de Recherche International* on Linear Logic and by french ANR project Rapido (ANR number: ANR-14-CE25-0007).

1 Introduction

1.1 Context: quantitative semantics and Taylor expansion

Linear logic takes its roots in the denotational semantics of λ-calculus: it is often presented, by Girard himself [15], as the result of a careful investigation of the model of coherence spaces. Since its early days, linear logic has thus generated a rich ecosystem of denotational models, among which we distinguish the family of *quantitative semantics*. Indeed, the first ideas behind linear logic were exposed even before coherence spaces, in the model of normal functors [16], in which Girard proposed to consider analyticity, instead of mere continuity, as

27th EACSL Annual Conference on Computer Science Logic (CSL 2018).
Editors: Dan Ghica and Achim Jung; Article No. 14; pp. 14:1–14:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the key property of the interpretation of λ-terms: in this setting, terms denote power series, representing analytic maps between modules.

This quantitative interpretation reflects precise operational properties of programs: the degree of a monomial in a power series is closely related to the number of times a function uses its argument. Following this framework, various models were considered — among which we shall include the multiset relational model as a degenerate, boolean-valued instance. These models allowed to represent and characterize quantitative properties such as the execution time [5], including best and worst case analysis for non-deterministic programs [18], or the probability of reaching a value [2]. It is notable that this whole approach gained momentum in the early 2000's, after the introduction by Ehrhard of models [7, 8] in which the notion of analytic maps interpreting λ-terms took its usual sense, while Girard's original model involved set-valued formal power series. Indeed, the keystone in the success of this line of work is an analogue of the Taylor expansion formula, that can be established both for λ-terms and for linear logic proofs.

Mimicking this denotational structure, Ehrhard and Regnier introduced the differential λ-calculus [12] and differential linear logic [13], which allow to formulate a syntactic version of Taylor expansion: to a λ-term (resp. to a linear logic proof), we associate an infinite linear combination of approximants [14, 11]. In particular, the dynamics (*i.e.* β-reduction or cut elimination) of those systems is dictated by the identities of quantitative semantics. In turn, Taylor expansion has become a useful device to design and study new models of linear logic, in which morphisms admit a matrix representation: the Taylor expansion formula allows to describe the interpretation of promotion — the operation by which a linear resource becomes freely duplicable — in an explicit, systematic manner. It is in fact possible to show that any model of differential linear logic without promotion gives rise to a model of full linear logic in this way [4]: in some sense, one can simulate cut elimination through Taylor expansion.

1.2 Motivation: reduction in Taylor expansion

There is a difficulty, however: Taylor expansion generates infinite sums and, *a priori*, there is no guarantee that the coefficients in these sums will remain finite under reduction. In previous works [4, 18], it was thus required for coefficients to be taken in a complete semiring: all sums should converge. In order to illustrate this requirement, let us first consider the case of λ-calculus.

The linear fragment of differential λ-calculus, called *resource λ-calculus*, is the target of the syntactical Taylor expansion of λ-terms. In this calculus, the application of a term to another is replaced with a multilinear variant: $\langle s\rangle[t_1,\dots,t_n]$ denotes the n-linear symmetric application of resource term s to the multiset of resource terms $[t_1,\dots,t_n]$. Then, if $x_1,\dots,x_k$ denote the occurrences of x in s, the redex $\langle\lambda x.s\rangle[t_1,\dots,t_n]$ reduces to the sum $\sum_{f:\{1,\dots,k\}\tilde{\to}\{1,\dots,n\}} s[t_{f(1)}/x_1,\dots,t_{f(k)}/x_k]$: here f ranges over all bijections $\{1,\dots,k\}\tilde{\to}\{1,\dots,n\}$ so this sum is zero if $n\neq k$. As sums are generated by reduction, it should be noted that all the syntactic constructs are linear, both in the sense that they commute to sums, and in the sense that, in the elimination of a redex, no subterm is copied nor erased. The key case of Taylor expansion is that of application:

$$\mathcal{T}(MN)=\sum_{n\in N}\frac{1}{n!}\langle\mathcal{T}(M)\rangle\mathcal{T}(N)^n \tag{1}$$

where $\mathcal{T}(N)^n$ is the multiset made of n copies of $\mathcal{T}(N)$ — by n-linearity, $\mathcal{T}(N)^n$ is itself an infinite linear combination of multisets of resource terms appearing in $\mathcal{T}(N)$. Admitting that

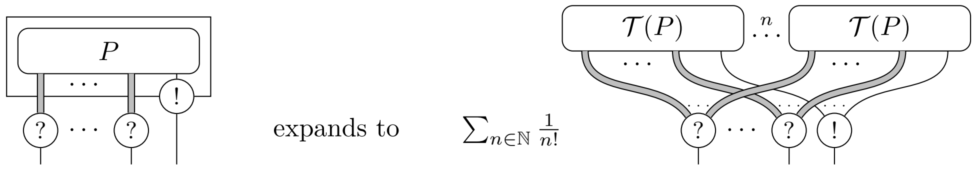

Figure 1 Taylor expansion of a promotion box (thick wires denote an arbitrary number of wires)

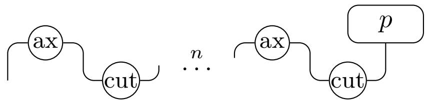

Figure 2 Example of a family of nets, all reducing to a single net

$\langle M\rangle[N_1,\ldots,N_n]$ represents the n-th derivative of M, computed at 0, and n-linearly applied to N_1, …, N_n, one immediately recognizes the usual Taylor expansion formula.

From (1), it is immediately clear that, to simulate one reduction step occurring in N, it is necessary to reduce in parallel in an unbounded number of subterms of each component of the expansion. Unrestricted parallel reduction, however, is ill defined in this setting. Consider the sum $\sum_{n\in\mathbb{N}}\langle\lambda xx\rangle[\cdots\langle\lambda xx\rangle[y]\cdots]$ where each summand consists of n successive linear applications of the identity to the variable y: then by simultaneous reduction of all redexes in each component, each summand yields y, so the result should be $\sum_{n\in\mathbb{N}} y$ which is not defined unless the semiring of coefficients is complete in some sense.

Those considerations apply to linear logic as well as to λ-calculus. We will use proof nets [15] as the syntax for proofs of multiplicative exponential linear logic (MELL). The target of Taylor expansion is then in promotion-free differential nets [13], which we call *resource nets* in the following, by analogy with resource λ-calculus: these form the multilinear fragment of differential linear logic.

In linear logic, Taylor expansion consists in replacing duplicable subnets, embodied by promotion boxes, with explicit copies, as in Fig. 1: if we take n copies of the box, the main port of the box is replaced with an n-ary ! link, while the ? links at the border of the box collect all copies of the corresponding auxiliary ports. Again, to follow a single cut elimination step in P, it is necessary to reduce an arbitrary number of copies. And unrestricted parallel cut elimination in an infinite sum of resource nets is broken, as one can easily construct an infinite family of nets, all reducing to the same resource net p in a single step of parallel cut elimination: see Fig. 2.

1.3 Our approach: taming the combinatorial explosion of antireduction

The problem of convergence of series of linear approximants under reduction was first tackled by Ehrhard and Regnier, for the normalization of Taylor expansion of ordinary λ-terms [14]. Their argument relies on a uniformity property, specific to the pure λ-calculus: the support of the Taylor expansion of a λ-term forms a clique in some fixed coherence space of resource terms. This method cannot be adapted to proof nets: there is no coherence relation on differential nets such that all supports of Taylor expansions are cliques [22, section V.4.1].

An alternative method to ensure convergence without any uniformity hypothesis was first developed by Ehrhard for typed terms in a λ-calculus extended with linear combinations of terms [9]: there, the presence of sums also forbade the existence of a suitable coherence relation. This method can be generalized to strongly normalizable [20], or even weakly normalizable [23] terms. One striking feature of this approach is that it concentrates on the support (*i.e.* the set of terms having non-zero coefficients) of the Taylor expansion. In each case, one shows that, given a normal resource term t and a λ-term M, there are finitely many terms s, such that:

- the coefficient of s in $\mathcal{T}(M)$ is non zero; and
- the coefficient of t in the normal form of s is non zero.

This allows to normalize the Taylor expansion: simply normalize in each component, then compute the sum, which is component-wise finite.

The second author then remarked that the same could be done for β-reduction [23], even without any uniformity, typing or normalizability requirement. Indeed, writing $s \rightrightarrows t$ if s and t are resource terms such that t appears in the support of a parallel reduct of s, the size of s is bounded by a function of the size of t *and* the height of s. So, given that if s appears in $\mathcal{T}(M)$ then its height is bounded by that of M, it follows that, for a fixed resource term t there are finitely many terms s in the support of $\mathcal{T}(M)$ such that $s \rightrightarrows t$: in short, parallel reduction is always well-defined on the Taylor expansion of a λ-term.

Our purpose in the present paper is to develop a similar technique for MELL proof nets: we show that one can bound the size of a resource net p by a function of the size of any of its parallel reducts, and of an additional quantity on p, yet to be defined. The main challenge is indeed to circumvent the lack of inductive structure in proof nets: in such a graphical syntax, there is no structural notion of height.

We claim that a side condition on switching paths, *i.e.* paths in the sense of Danos–Regnier's correctness criterion [3], is an appropriate replacement. Backing this claim, there are first some intuitions:

- the culprits for the unbounded loss of size in reduction are the chains of consecutive cuts, as in Fig. 2;
- we want the validity of our side condition to be stable under reduction so, rather than chains of cuts, we should consider cuts in switching paths;
- indeed, if p reduces to q *via* cut elimination, then the switching paths of q are somehow related with those of p;
- and the switching paths of a resource net in $\mathcal{T}(P)$ are somehow related with those of P.

In the following, we establish this claim up to some technical restrictions, which will allow us to simplify the exposition:

- we use generalized n-ary exponential links rather than separate (co)dereliction and (co)contraction, as this allows to reduce the dynamics of resource nets to that of multiplicative linear logic (MLL) proof nets;[1]
- we limit our study to a *strict* fragment of linear logic, *i.e.* we do not consider multiplicative units, nor the 0-ary exponential links — weakening and coweakening — as dealing with them would require us to introduce much more machinery.

1.4 Outline

In Section 2, we first introduce proof nets formally, in the term-based syntax of Ehrhard [10]. We define the parallel cut elimination relation $\rightrightarrows$ in this setting, that we decompose into multiplicative reduction $\rightrightarrows_m$ and axiom-cut reduction $\rightrightarrows_{ax}$. We also present the notion of

[1] In other words, we adhere to a version of linear logic proof nets and resource nets which is sometimes called *nouvelle syntaxe*, although it dates back to Regnier's PhD thesis [21]. See also the discussion in our conclusion (Section 6).

switching path for this syntax, and introduce the quantity that will be our main object of study in the following: the maximum number $\mathbf{cc}(p)$ of cuts that are crossed by any switching path in the net p. Let us mention that typing plays absolutely no role in our approach, so we do not even consider formulas of linear logic: we will rely only on the acyclicity of nets.

Section 3 is dedicated to the proof that we can bound $\mathbf{cc}(q)$ by a function of $\mathbf{cc}(p)$, whenever $p \rightrightarrows q$: the main case is the multiplicative reduction, as this may create new switching paths in q that we must relate with those in p. In this task, we concentrate on the notion of *slipknot*: a pair of residuals of a cut of p occurring in a path of q. Slipknots are essential in understanding how switching paths are structured after cut elimination.

We show in Section 4 that, if $p \rightrightarrows q$ then the size of p is bounded by a function of $\mathbf{cc}(p)$ and the size of q. Although, as explained in our introduction, this result is motivated by the study of quantitative semantics, it is essentially a theorem about MLL.

We establish the applicability of our approach to the Taylor expansion of MELL proof nets in Section 5: we show that if p is a resource net of $\mathcal{T}(P)$, then the length of switching paths in p is bounded by a function of the size of P — hence so is $\mathbf{cc}(p)$.

Finally, we discuss further work in the concluding Section 6.

2 Definitions

We provide here the minimal definitions necessary for us to work with MLL proof nets. We use a term-based syntax, following Ehrhard [10].

As stated before, let us stress the fact that the choice of MLL is not decisive for the development of Sections 2 to 4. The reader can check that we rely on two ingredients only:

- the definition of switching paths;
- the fact that multiplicative reduction amounts to plug bijectively the premises of a $\otimes$ link with those of $⅋$ link.

The results of those sections are thus directly applicable to resource nets, thanks to our choice of generalized exponential links: this will be done in Section 6.

2.1 Structures

Our nets are finite families of trees and cuts; trees are inductively defined as MLL connectives connecting trees, where the leaves are elements of a countable set of variables V. The duality of two conclusions of an axiom is given by an involution $x \mapsto \overline{x}$ over this set.

Formally, the set T of *raw trees* (denoted by s, t, *etc.*) is generated as follows:

$$t ::= x \mid \otimes(t_1, \dots, t_n) \mid ⅋(t_1, \dots, t_n)$$

where x ranges over a fixed countable set of variables V, endowed with a fixpoint-free involution $x \mapsto \overline{x}$.

We also define the subtrees of a given tree t, written $\mathbf{T}(t)$, in the natural way : if $t \in V$, then $\mathbf{T}(t) = \{t\}$. If $t = \alpha(t_1, \dots, t_n)$, then $\mathbf{T}(t) = \{t\} \cup \bigcup_{i \in \{1,\dots,n\}} \mathbf{T}(t_i)$, for $\alpha \in \{\otimes, ⅋\}$. In particular, we write $\mathbf{V}(t)$ for $\mathbf{T}(t) \cap V$. A *tree* is a raw tree t such that if $\alpha(t_1, \dots, t_n) \in \mathbf{T}(t)$ (with $\alpha = \otimes$ or $⅋$), then the sets $\mathbf{V}(t_i)$ for $1 \leq i \leq n$ are pairwise disjoint: in other words, each variable x occurs at most once in t. A tree t is *strict* if $\{\otimes(), ⅋()\} \cap \mathbf{T}(t) = \emptyset$.

From now on, we will consider strict trees only, *i.e.* we rule out the multiplicative units. This restriction will play a crucial rôle in expressing and establishing the bounds of Sections 3 and 4. It is possible to generalize our results in presence of units: we postpone the discussion on this subject to Section 6.[2]

A *cut* is an unordered pair $c = \langle t|s\rangle$ of trees such that $\mathbf{V}(t) \cap \mathbf{V}(s) = \emptyset$, and then we set $\mathbf{T}(c) = \mathbf{T}(t) \cup \mathbf{T}(s)$. A *reducible cut* is a cut $\langle t|s\rangle$ such that t is a variable and $\overline{t} \notin \mathbf{V}(s)$, or such that we can write $t = \otimes(t_1, \dots, t_n)$ and $s = ⅋(s_1, \dots, s_n)$, or *vice versa*. Note that, in the absence of typing, we do not require all cuts to be reducible, as this would not be stable under cut elimination.

Given a set A, we denote by $\overrightarrow{a}$ any finite family of elements of A. In general, we abusively identify $\overrightarrow{a}$ with any enumeration $(a_1, \dots, a_n) \in A^n$ of its elements, and write $\overrightarrow{a}, \overrightarrow{b}$ for the union of disjoint families $\overrightarrow{a}$ and $\overrightarrow{b}$. If $\overrightarrow{\gamma}$ is a family of trees or cuts, we write $\mathbf{V}(\overrightarrow{\gamma}) = \bigcup_{\gamma \in \overrightarrow{\gamma}} \mathbf{V}(\gamma)$ and $\mathbf{T}(\overrightarrow{\gamma}) = \bigcup_{\gamma \in \overrightarrow{\gamma}} \mathbf{T}(\gamma)$. An MLL *proof net* is a pair $p = (\overrightarrow{c}; \overrightarrow{t})$ of a finite family $\overrightarrow{c}$ of cuts and a finite family $\overrightarrow{t}$ of trees, such that for all cuts or trees $\gamma, \gamma' \in \overrightarrow{c}, \overrightarrow{t}$, $\mathbf{V}(\gamma) \cap \mathbf{V}(\gamma') = \emptyset$, and such that for any $x \in \mathbf{V}(p) = \mathbf{V}(\overrightarrow{c}) \cup \mathbf{V}(\overrightarrow{t})$, we have $\overline{x} \in \mathbf{V}(p)$ too. We then write $\mathbf{C}(p) = \overrightarrow{c}$.

2.2 Cut elimination

The *substitution* $\gamma[t/x]$ of a tree t for a variable x in a tree (or cut, or net) γ is defined in the usual way. By the definition of trees, we notice that this substitution is essentially linear, since each variable x appears at most once in a tree.

There are two basic cut elimination steps, one for each kind of reducible cut:

- the elimination of a connective cut yields a family of cuts: we write $\langle \otimes(t_1, \dots, t_n)| ⅋ (s_1, \dots, s_n)\rangle \to_m (\langle t_i|s_i\rangle)_{i \in \{1,\dots,n\}}$ that we extend to nets by setting $(c, \overrightarrow{c}; \overrightarrow{t}) \to_m (\overrightarrow{c}', \overrightarrow{c}; \overrightarrow{t})$ whenever $c \to_m \overrightarrow{c}'$;
- the elimination of an axiom cut generates a substitution: we write $(\langle x|t\rangle, \overrightarrow{c}; \overrightarrow{t}) \to_{ax} (\overrightarrow{c}; \overrightarrow{t})[t/\overline{x}]$ whenever $\overline{x} \notin \mathbf{V}(t)$.

We are in fact interested in the simultaneous elimination of any number of reducible cuts, that we describe as follows: we write $p \rightrightarrows p'$ if $p = (\langle x_1|t_1\rangle, \dots, \langle x_n|t_n\rangle, c_1, \dots, c_k, \overrightarrow{c}; \overrightarrow{t})$ and $p' = (\overrightarrow{c}'_1, \dots, \overrightarrow{c}'_k, \overrightarrow{c}; \overrightarrow{t})[t_1/\overline{x}_1] \cdots [t_n/\overline{x}_n]$, with $c_i \to_m \overrightarrow{c}'_i$ for $1 \leq i \leq k$, and $\overline{x}_i \notin \mathbf{V}(t_j)$ for $1 \leq i \leq j \leq n$. We moreover write $p \rightrightarrows_m p'$ (resp. $p \rightrightarrows_{ax} p$) in case $n = 0$ (resp. $k = 0$). It is a simple exercise to check that if $p \rightrightarrows p'$ then there exists q such that $p \rightrightarrows_m q \rightrightarrows_{ax} p'$: the converse does not hold, though, as the elimination of connective cuts may generate new axiom cuts.

2.3 Paths

In order to control the effect of parallel reduction on the size of proof nets, we rely on a side condition involving the number of cuts crossed by switching paths, *i.e.* paths in the sense of Danos–Regnier's correctness criterion [3].

In our setting, a *switching* of a net p is a partial map $I : \mathbf{T}(p) \to \mathbf{T}(p)$ such that, for each $t = ⅋(t_1, \dots t_n) \in \mathbf{T}(p)$, $I(t) \in \{t_1, \dots, t_n\}$. Given a net p and a switching I of p, we define adjacency relations between the elements of $\mathbf{T}(p)$, written $\sim_{t,s}$ for $t, s \in \mathbf{T}(p)$ and $\sim_c$ for $c \in \mathbf{C}(p)$, as the least symmetric relations such that:

---

[2] An additional consequence is the fact that, given a (strict) tree t, any other tree u occurs at most once as a subtree of t: *e.g.*, in $⅋^2(t_1, t_2)$, $\mathbf{V}(t_1)$ and $\mathbf{V}(t_2)$ are both non empty and disjoint, so that $t_1 \neq t_2$. In other words, we can identify $\mathbf{T}(t)$ with the positions of subtrees in t, that play the rôle of vertices when considering t as a graphical structure. This will allow us to keep notations concise in our treatment of paths. This trick is of course inessential for our results.

- for any $x \in \mathbf{V}(p)$, $x \sim_{x,\overline{x}} \overline{x}$;
- for any $t = \otimes(t_1, \ldots, t_n) \in \mathbf{T}(p)$, $t \sim_{t,t_i} t_i$ for each $i \in \{1, \ldots, n\}$;
- for any $t = \parr(t_1, \ldots, t_n) \in \mathbf{T}(p)$, $t \sim_{t,I(t)} I(t)$;
- for any $c = \langle t|s\rangle \in \mathbf{C}(p)$, $t \sim_c s$.

Whenever necessary, we may write, e.g., $\sim^p_{t,s}$ or $\sim^{p,I}_{t,s}$ for $\sim_{t,s}$ to make the underlying net and switching explicit. Let l and $m \in (\mathbf{T}(p) \times \mathbf{T}(p)) \cup \mathbf{C}(p)$ be two adjacency labels: we write $l \equiv m$ if $l = m$ or $m = (x, \overline{x})$ and $l = (\overline{x}, x)$ for some $x \in V$.

Given a switching I in p, an *I-path* is a sequence of trees $t_0, \ldots, t_n$ of $\mathbf{T}(p)$ such that there exists a sequence of pairwise $\not\equiv$ labels $l_1, \ldots, l_n$ with, for each $i \in \{1, \ldots, n\}$, $t_{i-1} \sim^{p,I}_{l_i} t_i$.[3] For instance, if $p = (; \otimes(x, y), \parr(\overline{y}, \overline{x}))$ and $I(\parr(\overline{y}, \overline{x})) = \overline{x}$, then the chain of adjacencies $\parr(\overline{x}, \overline{y}) \sim_{\parr(\overline{x},\overline{y}),\overline{x}} \overline{x} \sim_{x,\overline{x}} x \sim_{\otimes(x,y),x} \otimes(x, y) \sim_{\otimes(x,y),y} y \sim_{y,\overline{y}} \overline{y}$ defines an I-path in p, which can be depicted as the dashed line in the following graphical representation of p:

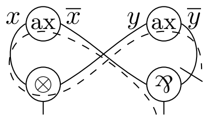

.

We call *path* in p any I-path for I a switching of p, and we write $\mathbf{P}(p)$ for the set of all paths in p. We write $t \rightsquigarrow s$ or $t \rightsquigarrow_p s$ whenever there exists a path from t to s in p. Given $\chi = t_0, \ldots, t_n \in \mathbf{P}(p)$, we call *subpaths* of χ the subsequences of χ: a subpath is either the empty sequence ϵ or a path of p. We moreover write $\overline{\chi}$ for the reverse path: $\overline{\chi} = t_n, \ldots, t_0 \in \mathbf{P}(p)$. We say a net p is *acyclic* if for all $\chi \in \mathbf{P}(p)$ and $t \in \mathbf{T}(p)$, t occurs at most once in χ: in other words, there is no *cycle* t, χ, t. From now on, we consider acyclic nets only: it is well known that if p is acyclic and $p \rightrightarrows q$ then q is acyclic too.

If $c = \langle t|s\rangle \in \mathbf{C}(p)$, we may write χ_1, c, χ_2 for either χ_1, s, t, χ_2 or χ_1, t, s, χ_2: by acyclicity, this notation is unambiguous, unless $\chi_1 = \chi_2 = \epsilon$.

For all $\chi \in \mathbf{P}(p)$, we write $\mathbf{cc}_p(\chi)$, or simply $\mathbf{cc}(\chi)$, for the number of cuts *crossed by* χ: $\mathbf{cc}_p(\chi) = \#\{\langle t|s\rangle \in \mathbf{C}(p) \mid t \in \chi\}$ (recall that cuts are unordered). Observe that, by acyclicity, a path χ crosses each cut $c = \langle t|s\rangle$ at most once: either $\chi = \chi_1, c, \chi_2$, or $\chi = \chi_1, t, \chi_2$, or $\chi = \chi_1, s, \chi_2$, with neither t nor s occurring in χ_1, χ_2. Finally, we write $\mathbf{cc}(p) = \max\{\mathbf{cc}(\chi) \mid \chi \in \mathbf{P}(p)\}$: in the following, we show that the maximal number of cuts crossed by any switching path is a good parameter to limit the decrease in size induced by parallel reduction.

3 Variations of $\mathrm{cc}(p)$ under reduction

Here we establish that the possible increase of $\mathbf{cc}(p)$ under reduction is bounded. It should be clear that if $p \rightrightarrows_{ax} q$ then $\mathbf{cc}(q) \leq \mathbf{cc}(p)$: intuitively, the only effect of $\rightrightarrows_{ax}$ is to straighten some paths, thus decreasing the number of crossed cuts. In the case of connective cuts however, cuts are duplicated and new paths are created.

Consider for instance a net r, as in Fig. 3, obtained from three nets p_1, p_2 and q, by forming the cut $\langle \otimes(t_1, t_2)|\parr(s_1, s_2)\rangle$ where $t_1 \in \mathbf{T}(p_1)$, $t_2 \in \mathbf{T}(p_2)$ and $s_1, s_2 \in \mathbf{T}(q)$. Observe that, in the reduct r' obtained by forming two cuts $\langle t_1|s_1\rangle$ and $\langle t_2|s_2\rangle$, we may very well form a path that travels from p_1 to q then p_2; while in p, this is forbidden by any switching

[3] In standard terminology of graph theory, an I-path in p is a trail in the unoriented graph with vertices in $\mathbf{T}(p)$ and edges given by the sum of adjacency relations defined by I (identifying $\sim_{x,\overline{x}}$ with $\sim_{\overline{x},x}$). The only purpose of our choice of labels for adjacency relations and the definition of $\equiv$ is indeed to capture this notion of path in the unoriented graph of subtrees induced by a switching in a net.

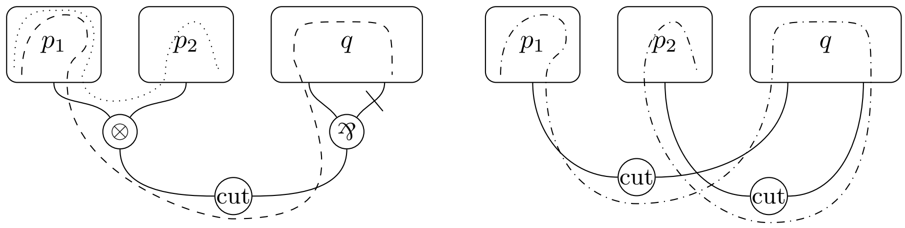

Figure 3 A cut, the resulting slipknot, and examples of paths before and after reduction

of $\parr(s_1, s_2)$. For instance, if we consider $I(\parr(s_1, s_2)) = s_1$, we may only form a path between p_1 and p_2 through $\otimes(t_1, t_2)$, or a path between q and one of the p_i's, through s_1 and the cut.

In the remainder of this section, we fix a reduction step $p \rightrightarrows_m q$, and we show that the previous example describes a general mechanism: if a new path is created in this step $p \rightrightarrows_m q$, it must involve a path ξ between two premises of a $\parr$ involved in a cut c of p, *unfolded* into a path between the residuals of this cut. We call such an intermediate path ξ a *slipknot*.

3.1 Residual cuts and slipknots

Notice that $\mathbf{T}(q) \subseteq \mathbf{T}(p)$. Observe that, given a switching J of q, it is always possible to extend J into a switching I of p, so that, for all $t, s \in \mathbf{T}(q)$:

- if $t \sim^{q,J}_{t,s} s$ then $t \sim^{p,I}_{t,s} s$, and
- if $c \in \mathbf{C}(p)$ and $t \sim^{q,J}_{c} s$ then $t \sim^{p,I}_{c} s$.

To determine I uniquely, is remains only to select a premise for each $\parr$ involved in an eliminated cut. Consider $c = \langle \otimes(t_1, \dots, t_n) | \parr(s_1, \dots, s_n) \rangle \in \mathbf{C}(p)$ and assume c is eliminated in the reduction $p \rightrightarrows_m q$. Then the *residuals* of c in q are the cuts $\langle t_i | s_i \rangle \in \mathbf{C}(q)$ for $1 \leq i \leq n$.

If $\xi \in \mathbf{P}(q)$, a *slipknot* of ξ is any pair (d, d') of (necessarily distinct) residuals in q of a cut in p, such that we can write $\xi = \chi_1, d, \chi_2, d', \chi_3$. We now show that a path in q is necessarily obtained by alternating paths in p and paths between slipknots, that recursively consist of such alternations. This will allow us to bound $\mathbf{cc}(q)$ depending on $\mathbf{cc}(p)$, by reasoning inductively on these paths. The main tool is the following lemma:

▶ **Lemma 1.** *If $\xi \in \mathbf{P}(q)$ then there exists a path $\xi^- \in \mathbf{P}(p)$ with the same endpoints as ξ.*

Proof. Assuming ξ is a J-path of q, we construct an I-path ξ^- in p with the same endpoints as ξ for an extension I of J as above. The definition is by induction on the number of residuals occurring as subpaths of ξ. In the process, we must ensure that the constraints we impose on I in each induction step can be satisfied globally: the trick is that we fix the value of $I(\parr(\overrightarrow{s}))$ only in case exactly one residual of the cut involving $\parr(\overrightarrow{s})$ occurs in ξ.

First consider the case of $\xi = \chi_1, d, \chi_2, d', \chi_3$, for a slipknot (d, d'), where d and d' are residuals of $c \in \mathbf{C}(p)$. We can assume, w.l.o.g, that: (i) no other residual of c occurs in χ_1, nor in χ_3; (ii) no residual of a cut $c' \neq c$ occurs in both χ_1 and χ_3. By the definition of residuals, we can write $c = \langle \otimes(\overrightarrow{t}) | \parr(\overrightarrow{s}) \rangle \in \mathbf{C}(p)$, $d = \langle t | s \rangle$ and $d' = \langle t' | s' \rangle$ with $t, t' \in \overrightarrow{t}$ and $s, s' \in \overrightarrow{s}$. It is then sufficient to prove that $\xi = \chi_1, t, s, \chi_2, s', t', \chi_3$, in which case we can set $\xi^- = \chi_1^-, t, \otimes(\overrightarrow{t}), t', \chi_3^-$, where χ_1^- and χ_3^- are obtained from the induction hypothesis (or by setting $\epsilon^- = \epsilon$ for empty subpaths): by condition (ii), the constraints we impose on I by forming χ_1^- and χ_3^- are independent.

Let us rule out the other three orderings of d and d': (a) $\xi = \chi_1, s, t, \chi_2, t', s', \chi_3$, (b) $\xi = \chi_1, s, t, \chi_2, s', t', \chi_3$ or (c) $\xi = \chi_1, t, s, \chi_2, t', s', \chi_3$. First observe that χ_2 is not empty.

Indeed, if $t \sim_l^q t'$ (or $t \sim_l^q s'$, or $s \sim_l^q t'$) then: l cannot be a cut of q because $\langle t|s\rangle$ and $\langle t'|s'\rangle \in \mathbf{C}(q)$; l cannot be of the form $(\alpha(t_1, \cdots, t_n), t_n)$ because the trees t, t', s, s' are pairwise disjoint; so l must be an axiom and we obtain a cycle in q.

Let u and v be the endpoints of χ_2, and consider $\chi_2^- \in \mathbf{P}(p)$ with the same endpoints, obtained by induction hypothesis. Necessarily, we have $t \sim_l^{q,J} u$ in cases (a) and (b), $s \sim_l^{q,J} u$ in case (c), $t' \sim_m^{q,J} v$ in cases (a) and (c), and $s' \sim_m^{q,J} v$ in case (b), where $l \not\equiv m$, and nor l nor m is a cut: it follows that the same adjacencies hold in p for any extension I of J. Observe that $\otimes(\overrightarrow{t}) \notin \chi_2^-$: otherwise, we would obtain a path $t \rightsquigarrow_p \otimes(\overrightarrow{t})$ (or $\otimes(\overrightarrow{t}) \rightsquigarrow_p t'$) that we could extend into a cycle. Then in case (a), we obtain a cycle in p directly: $t, \chi_2^-, t', \otimes(\overrightarrow{t}), t$. In cases (b) and (c), we deduce that $⅋(\overrightarrow{s}) \notin \chi_2^-$, and we obtain a cycle, e.g. in case (b): $t, \chi_2^-, s', ⅋(\overrightarrow{s}), \otimes(\overrightarrow{t}), t'$, for any I such that $I(⅋(\overrightarrow{s})) = s'$.

We can now assume that each cut of p has at most one residual occurring as a subpath of ξ. If no residual occurs in ξ, then we can set $\xi^- = \xi$. Now fix $c = \langle \otimes(\overrightarrow{t}) | ⅋(\overrightarrow{s})\rangle \in \mathbf{C}(p)$ and assume, w.l.o.g (otherwise, consider $\overline{\xi}$), that $\xi = \chi_1, t, s, \chi_2$ with $t \in \overrightarrow{t}$ and $s \in \overrightarrow{s}$. Then we set $I(⅋(\overrightarrow{s})) = s$ and $\xi^- = \chi_1^-, t, c, s, \chi_2^- \in \mathbf{P}(p)$: this is the only case in which we impose a value for I to construct ξ^-, so this choice, and the choices we make to form χ_1^- and χ_2^- are all independent. ◀

▶ **Lemma 2.** *If $\xi \in \mathbf{P}(q)$ and $c = \langle \otimes(\overrightarrow{t}) | ⅋(\overrightarrow{s})\rangle \in \mathbf{C}(p)$, then at most two residuals of c occur as subpaths of ξ, and then we can write $\xi = \chi_1, t, s, \chi_2, s', t', \chi_3$ with $t, t' \in \overrightarrow{t}$ and $s, s' \in \overrightarrow{s}$.*

Proof. Assume $\xi = \chi_1, d, \chi_2, d', \chi_3$ and $d = \langle t|s\rangle$ and $d' = \langle t'|s'\rangle$ with $t, t' \in \overrightarrow{t}$ and $s, s' \in \overrightarrow{s}$. Using Lemma 1, we establish that $\xi = \chi_1, t, s, \chi_2, s', t', \chi_3$: we can exclude the other cases exactly as in the proof of Lemma 1. Then, as soon as three residuals of c occur in ξ, a contradiction follows. ◀

▶ **Lemma 3.** *Slipknots are well-bracketed in the following sense: there is no path $\xi = d_1, \chi_1, d_2, \chi_2, d_1', \chi_3, d_2' \in \mathbf{P}(q)$ such that both (d_1, d_1') and (d_2, d_2') are slipknots.*

Proof. Assume $c_1 = \langle \otimes(\overrightarrow{t}_1) | ⅋(\overrightarrow{s}_1)\rangle$, $c_2 = \langle \otimes(\overrightarrow{t}_2) | ⅋(\overrightarrow{s}_2)\rangle$, and, for $1 \leq i \leq 2$, $d_i = (t_i, s_i)$ and $d_i' = (t_i', s_i')$, with $t_i, t_i' \in \overrightarrow{t}_i$ and $s_i, s_i' \in \overrightarrow{s}_i$. By the previous lemma, we must have $\xi = t_1, s_1, \chi_1, t_2, s_2, \chi_2, s_1', t_1', \chi_3, s_2', t_2'$. Observe that nor χ_1^- nor χ_3^- can cross c_1 or c_2: otherwise, we obtain a cycle in p. Then $s_1, \chi_1^-, t_2, c_1, s_2', \overline{\chi_3}^-, t_1', c_2, s_1$ is a cycle in p. ◀

▶ **Corollary 4.** *Any path of q is of the form $\zeta_1, c_1, \chi_1, c_1', \zeta_2, \ldots \zeta_n, c_n, \chi_n, c_n', \zeta_{n+1}$ where each subpath ζ_i is without slipknot, and each (c_i, c_i') is a slipknot.*

The previous result describes precisely how paths in q are related with those in p: it will be crucial in the following.

3.2 Bounding the growth of cc

Now we show that we can bound $\mathbf{cc}(q)$ depending only on $\mathbf{cc}(p)$. For each $\xi \in \mathbf{P}(q)$, we define the *width* $w_p(\xi)$ (or just $w(\xi)$): $w_p(\xi) = \max\{\mathbf{cc}_p(\chi^-) | \chi \text{ subpath of } \xi\}$. We have:

▶ **Lemma 5.** *For any path $\zeta \in \mathbf{P}(q)$, $\mathbf{cc}_p(\zeta^-) \leq w_p(\zeta) \leq \mathbf{cc}(p)$ and $w_p(\zeta) \leq \mathbf{cc}_q(\zeta)$. If moreover ζ has no slipknot, then $w_p(\zeta) = \mathbf{cc}_q(\zeta) = \mathbf{cc}_p(\zeta^-)$.*

Defining $\varphi : \mathbb{N} \to \mathbb{N}$ by $\varphi(0) = 0$ and $\varphi(n+1) = 2(n+1) + (n+1)(\varphi(n))$, we obtain:

▶ **Lemma 6.** *If $\xi \in \mathbf{P}(q)$ then $\mathbf{cc}(\xi) \leq \varphi(w_p(\xi))$.*

Proof. The proof is by induction on $w(\xi)$. If $w(\xi) = 0$, then we can easily check that $\mathbf{cc}(\xi) = 0$. Otherwise assume $w(\xi) = n + 1$. Then we set $\xi = \zeta_1, c_1, \chi_1, c'_1, \zeta_2, \dots \zeta_k, c_k, \chi_k, c'_n, \zeta_{k+1}$ as in Corollary 4.

First observe that for all $i \in \{1, \dots, k\}$, $w(\chi_i) \leq w(\xi) - 1$. Indeed, c_i, χ_i is a subpath of ξ and $w(c_i, \chi_i) = w(\chi_i) + 1$ by the definition of width. So, by induction hypothesis, $\mathbf{cc}(\chi_i) \leq \varphi(n)$. We also have that $\sum_{i=1}^{k+1} \mathbf{cc}(\zeta_i) \leq w(\xi) - k$. Observe indeed that $\mathbf{cc}(\xi^-) = \sum_{i=1}^{k+1} \mathbf{cc}(\zeta_i) + k$, because of Lemma 5 applied to ζ_i, and because of the construction of ξ^- that contracts the slipknots c_i, χ_i, c'_i; also recall that $\mathbf{cc}(\xi^-) \leq w(\xi)$.

We obtain:

$$\mathbf{cc}(\xi) = \sum_{1 \leq i \leq k} \mathbf{cc}(\chi_i) + \sum_{1 \leq j \leq k+1} \mathbf{cc}(\zeta_j) + 2k \leq k\varphi(n) + w(\xi) - k + 2k$$

and, since $k \leq \mathbf{cc}(\xi^-) \leq w(\xi) = n+1$, we obtain $\mathbf{cc}(\xi) \leq (n+1)\varphi(n) + 2(n+1) = \varphi(n+1)$. ◀

Using Lemma 5 again, we obtain:

▶ **Corollary 7.** *Let $p \rightrightarrows_m q$. Then, $\mathbf{cc}(q) \leq \varphi(\mathbf{cc}(p))$.*

▶ Remark. It is in fact possible to show that $\mathbf{cc}(q) \leq 2n!\mathbf{cc}(p)$, which is a better bound and closer to the graphical intuition, but the proof is much longer, and we are only interested in the existence of a bound.

4 Bounding the size of antireducts

For any tree, cut or net γ, we define the *size* of γ as $\#\gamma = \mathbf{card}(\mathbf{T}(\gamma))$: graphically, $\#p$ is nothing but the number of wires in p. In this section, we show that the loss of size during parallel reduction is directly controlled by $\mathbf{cc}(p)$ and $\#q$: more precisely, we show that the ratio $\frac{\#p}{\#q}$ is bounded by a function of $\mathbf{cc}(p)$.

First observe that the elimination of multiplicative cuts cannot decrease the size by more than a half:

▶ **Lemma 8.** *If $p \rightrightarrows_m q$ then $\#p \leq 2\#q$.*

Proof. It is sufficient to observe that if $c \to_m \overrightarrow{c}$ then $\#c = 2 + \#\overrightarrow{c} \leq 2\#\overrightarrow{c}$.[4] ◀

4.1 Elimination of axiom cuts

Observe that:

- if $x \in \mathbf{V}(\gamma)$ then $\#\gamma[t/x] = \#\gamma + \#t - 1$;
- if $x \notin \mathbf{V}(\gamma)$ then $\#\gamma[t/x] = \#\gamma$.

It follows that, in the elimination of a single axiom cut $p \to_{ax} q$, we have $\#p = \#q + 1$. But we cannot reproduce the proof of Lemma 8 for $\rightrightarrows_{ax}$: as stated in our introduction, chains of axiom cuts reducing into a single wire are the source of the collapse of size. We can bound the length of those chains by $\mathbf{cc}(p)$, however, and this allows us to bound the loss of size during reduction.

▶ **Lemma 9.** *If $p \rightrightarrows_{ax} q$ then $\#p \leq (2\mathbf{cc}(p) + 1)\#q$.*

[4] This is due to the fact that all the trees are strict, so $\overrightarrow{c}$ is not empty and $\#\overrightarrow{c} \geq 1$. Without the strictness condition, we would have to deal with annihilating reductions $\langle \otimes() \mid \parr () \rangle \to_m \epsilon$: this will be discussed in the conclusion.

Proof. Assume $p = (\langle x_1|t_1\rangle, \ldots, \langle x_n|t_n\rangle, \overrightarrow{c}; \overrightarrow{s})$ and $q = (\overrightarrow{c}; \overrightarrow{s})[t_1/\overline{x}_1] \cdots [t_n/\overline{x}_n]$ with $\overline{x}_i \notin \mathbf{V}(t_j)$ for $1 \leq i \leq j \leq n$. In case $\mathbf{cc}(p) = 0$, we have $n = 0$ and $p = q$ so the result is obvious. We thus assume $\mathbf{cc}(p) > 0$: to establish the result in this case, we make the chains of eliminated axiom cuts explicit.

Due to the condition on free variables, there exists a (necessarily unique) permutation of $\langle x_1|t_1\rangle, \ldots, \langle x_n|t_n\rangle$ yielding a family of the form $\overrightarrow{c}_1, \ldots, \overrightarrow{c}_k$ such that:

- for $1 \leq i \leq k$, we can write $\overrightarrow{c}_i = \langle x_0^i|\overline{x}_1^i\rangle, \ldots, \langle x_{n_i-1}^i|\overline{x}_{n_i}^i\rangle, \langle x_{n_i}^i|t^i\rangle$;
- each $\overrightarrow{c}_i$ is maximal with this shape, *i.e.* $\overline{x}_0^i \notin \{x_1, \ldots, x_n, t_1, \ldots, t_n\}$ and, in case t^i is a variable, $\overline{t}^i \notin \{x_1, \ldots, x_n, t_1, \ldots, t_n\}$;
- if $i < j$, then the cut $\langle x_{n_i}^i|t^i\rangle$ occurs before $\langle x_{n_j}^j|t^j\rangle$ in $\langle x_1|t_1\rangle, \ldots, \langle x_n|t_n\rangle$.

It follows that if $\overline{x}_0^i \in \mathbf{V}(t_j)$ then $j < i$, and then $q = (\overrightarrow{c}; \overrightarrow{s})[t^1/\overline{x}_0^1] \cdots [t^k/\overline{x}_0^k]$, by applying the same permutation to the substitutions as we did to cuts: we can do so because, by a standard argument, if $x \neq y$, $x \notin \mathbf{V}(u)$ and $y \notin \mathbf{V}(u)$ then $\gamma[u/x][v/y] = \gamma[v/y][u/x]$.

For $1 \leq i \leq k$, since $\overrightarrow{c}_i$ is a chain of $n_i + 1$ cuts, it follows that $n_i \leq \mathbf{cc}(p) - 1$. So $\#p = \#\overrightarrow{c} + \#\overrightarrow{s} + \sum_{i=1}^{k}(\#t^i + 2n_i + 1) \leq \#\overrightarrow{c} + \#\overrightarrow{s} + \sum_{i=1}^{k} \#t^i + k(2\mathbf{cc}(p) - 1)$. Moreover $\#q = \#\overrightarrow{c} + \#\overrightarrow{s} + \sum_{i=1}^{k} \#t^i - k$. It follows that $\#p \leq \#q + 2k\mathbf{cc}(p)$ and, to conclude, it will be sufficient to prove that $\#q \geq k$.

For $1 \leq i \leq k$, let $A_i = \{j > i \mid \overline{x}_0^j \in \mathbf{V}(t^i)\}$, and then let $A_0 = \{i \mid \overline{x}_0^i \in \mathbf{V}(\overrightarrow{c}, \overrightarrow{s})\}$. It follows from the construction that $\{A_0, \ldots, A_{k-1}\}$ is a partition (possibly including empty sets) of $\{1, \ldots, k\}$. By construction, $\#t^i > \mathbf{card}(A_i)$. Now consider $q_i = (\overrightarrow{c}; \overrightarrow{s})[t^1/\overline{x}_0^1] \cdots [t^i/\overline{x}_0^i]$ for $0 \leq i \leq k$ so that $q = q_k$. For $1 \leq i \leq k$, we obtain $\#q_i = \#q_{i-1} + \#t^i - 1 \geq \#q_{i-1} + \mathbf{card}(A_i)$. Also observe that $\#q_0 = \#(\overrightarrow{c}; \overrightarrow{s}) \geq \mathbf{card}(A_i)$. We can then conclude: $\#q = \#q_k \geq \sum_{i=0}^{k} \mathbf{card}(A_i) = k$. ◀

4.2 General case

Recall that any parallel cut elimination step $p \rightrightarrows q$ can be decomposed into a multiplicative-then-axiom pair of reductions: $p \rightrightarrows_m q' \rightrightarrows_{ax} q$. This allows us to bound the loss of size in the reduction $p \rightrightarrows q$, using the previous results:

▶ **Theorem 10.** *If $p \rightrightarrows q$ then $\#p \leq 4(\varphi(\mathbf{cc}(p)) + 1)\#q$.*

Proof. Consider first q' such that $p \rightrightarrows_m q'$ and $q' \rightrightarrows_{ax} q$. By Lemma 8, $\#p \leq 2\#q'$. Lemma 9 states that $\#q' \leq (2\mathbf{cc}(q') + 1)\#q$. Finally, Corollary 7, entails that $\mathbf{cc}(q') \leq \varphi(\mathbf{cc}(p))$, and we can conclude: $\#p \leq 2(\varphi(\mathbf{cc}(p) + 1)\#q) \leq 4(\varphi(\mathbf{cc}(p)) + 1)\#q$. ◀

▶ **Corollary 11.** *If q is an MLL net and $n \in \mathbb{N}$, then $\{p \mid p \rightrightarrows q$ and $\mathbf{cc}(p) \leq n\}$ is finite.*

To be precise, due to our term syntax, the previous corollary holds only up to renaming variables in axioms: we keep this precision implicit in the following.

It follows that, given an infinite linear combination of $\sum_{i\in I} a_i.p_i$, such that $\{\mathbf{cc}(p_i) \mid i \in I\}$ is finite, we can always consider an arbitrary family of reductions $p_i \rightrightarrows q_i$ for $i \in I$ and form the sum $\sum_{i\in I} a_i.q_i$: this is always well defined.

5 Taylor expansion

We now show how the previous results apply to Taylor expansion. For that purpose, we must extend our syntax to MELL proof nets. Our presentation departs from Ehrhard's [11] in our treatment of promotion boxes: instead of introducing boxes as tree constructors labelled by nets, with auxiliary ports as inputs, we consider box ports as 0-ary trees, that are related with each other in a *box context*, associating each box with its contents. This is in accordance with the usual presentation of promotion as a black box, and has two motivations:

- In Ehrhard's syntax, the promotion is not a net but an open tree, for which the trees associated with auxiliary ports must be mentioned explicitly: this would complicate the expression of Taylor expansion.
- The *nouvelle syntaxe* imposes constraints on auxiliary ports, that are easier to express when these ports are directly represented in the syntax.

Then we show that if p is a resource net in the support of the Taylor expansion of an MELL proof net P, then $\mathbf{cc}(p)$ (and in fact the length of any path in p) is bounded by a function of P.

Observe that we need only consider the support of Taylor expansion, so we do not formalize the expansion of MELL nets into infinite linear combinations of resource nets: rather, we introduce $\mathcal{T}(P)$ as a set of approximants. Also, as we limit our study to *strict* nets, we will restrict $\mathcal{T}(P)$ to those approximants that take at least one copy of each box of P: this is enough to cover the case of weakening-free MELL.

5.1 MELL nets

In addition to the set of variables, we fix a denumerable set $\mathcal{A}$ of *box ports*: we assume given an enumeration $\mathcal{A} = \{a_i^b \mid i, b \in \mathbb{N}\}$. We call *principal ports* the ports a_0^b and *auxiliary ports* the other ports. In the so-called *nouvelle syntaxe* of MELL, contractions and derelictions are merged together in a generalized contraction cell, and auxiliary ports must be premises of such generalized contractions.

We introduce the corresponding term syntax, as follows. Raw pre-trees (S°, T°, *etc.*) and raw trees (S, T, *etc.*) are defined by mutual induction as follows:

$$T ::= x \mid a_0^b \mid \otimes(T_1, \ldots, T_n) \mid \parr(T_1, \ldots, T_n) \mid ?(T_1^\circ, \ldots, T_n^\circ) \quad \text{and} \quad T^\circ ::= T \mid a_{i+1}^b$$

requiring that each $\otimes$, $\parr$ and $?$ is of arity at least 1. We write $\mathbf{V}(S)$ (resp. $\mathbf{B}(S)$) for the set of variables (resp. of principal and auxiliary ports) occurring in S. A *tree* (resp. a *pre-tree*) is a raw tree (resp. raw pre-tree) in which each variable and port occurs at most once. A *cut* is an unordered pair of trees $C = \langle T|S\rangle$ with disjoints sets of variables and ports.

We now define *box contexts* and *pre-nets* by mutual induction as follows. A box context Θ is the data of a finite set $\mathcal{B}_\Theta \subset \mathbb{N}$, and, for each $b \in \mathcal{B}_\Theta$, a closed pre-net $\Theta(b)$, of the form $(\Theta_b; \overrightarrow{C}_b; T_b, \overrightarrow{S}^\circ_b)$. Then we write $\overrightarrow{S}^\circ_b = S^\circ_{b,1}, \ldots, S^\circ_{b,n_b}$. A pre-net is a triple $P^\circ = (\Theta; \overrightarrow{C}; \overrightarrow{S}^\circ)$ where Θ is a box context, each variable and port occurs at most once in $\overrightarrow{C}$, $\overrightarrow{S}^\circ$, and moreover, if $a_i^b \in \mathbf{B}(\overrightarrow{C}; \overrightarrow{S}^\circ)$ then $b \in \mathcal{B}_\Theta$ and $i \leq n_b$. A closed pre-net is a pre-net $P^\circ = (\Theta; \overrightarrow{C}; \overrightarrow{S}^\circ)$ such that x occurs iff $\overline{x}$ occurs, and moreover, if $b \in \mathcal{B}_\Theta$ then each a_i^b with $0 \leq i \leq n_b$ occurs. Then a *net* is a closed pre-net of the form $P = (\Theta; \overrightarrow{C}; \overrightarrow{S})$.

We write $\mathbf{T}(\gamma)$ for the set of sub-pre-trees of a pre-tree, or cut, or pre-net γ: the definition extends that for subtrees in MLL nets, moreover setting $\mathbf{T}(a) = \{a\}$ for any $a \in \mathcal{A}$ (so we do not look into the content of boxes). As for MLL, we set $\#\gamma = \mathbf{card}(\mathbf{T}(\gamma))$. We write $\mathbf{depth}(P^\circ)$ for the maximum level of nesting of boxes in P°, *i.e.* the inductive depth in the previous definition. Also, the size of MELL pre-nets includes that of their boxes: we set $\mathbf{size}(P^\circ) = \#P^\circ + \sum_{b \in \mathcal{B}_\Theta} \mathbf{size}(\Theta(b))$.

We extend the switching functions of MLL to ? links: for each $T = ?(T_1, \ldots, T_n)$, $I(T) \in \{T_1, \ldots, T_n\}$, which induces a new adjacency relation $T \sim_{T,I(T)} I(T)$. We also consider adjacency relations $\sim_b$ for $b \in \mathcal{B}_\Theta$, setting $a_i^b \sim_b a_j^b$ whenever $0 \leq i < j \leq n_b$: w.r.t. paths, a box be behaves like an $(n_b + 1)$-ary axiom link and the contents is not considered.

We write $\mathbf{P}(P^\circ)$ for the set of paths in P°. We say a pre-net P° is *acyclic* if there is no cycle in $\mathbf{P}(P^\circ)$ and, inductively, each $\Theta(b)$ is acyclic. From now on, we consider acyclic pre-nets only.

5.2 Resource nets and Taylor expansion

The Taylor expansion of a net P will be a set of *resource nets*: these are the same as the multiplicative nets introduced before, except we have two new connectives ! and ?. Raw trees are given as follows:

$$t ::= x \mid \otimes(t_1, \dots, t_n) \mid ⅋(t_1, \dots, t_n) \mid !(t_1, \dots, t_n) \mid ?(t_1, \dots, t_n).$$

Again, we will consider strict trees only: each $\otimes$, ⅋, ! and ? is of arity at least 1. In resource nets, we extend switchings to ? links as in MELL nets, and for each $t = ?(t_1, \dots, t_n)$, we set $t \sim_{t,I(t)} I(t)$. Moreover, for each $t = !(t_1, \dots, t_n)$, we set $t \sim_{t,t_i} t_i$ for $1 \leq i \leq n$.

We are now ready to introduce the expansion of MELL nets. During the construction, we need to track the conclusions of copies of boxes, in order to collect copies of auxiliary ports in the external ? links: this is the rôle of the intermediate notion of pre-Taylor expansion.

▶ **Definition 12.** Taylor expansion is defined by induction on depth as follows. Given a closed pre-net $P^\circ = (\Theta; \overrightarrow{C}; \overrightarrow{S}^\circ)$, a *pre-Taylor expansion* of P° is any pair (p, f) of a resource net $p = (\overrightarrow{c}; \overrightarrow{t})$, together with a function $f : \overrightarrow{t} \to \overrightarrow{S}^\circ$ such that $f^{-1}(T)$ is a singleton whenever $T \in \overrightarrow{S}^\circ$ is a tree, obtained as follows:

- for each $b \in \mathcal{B}_\Theta$, fix a number $k_b > 0$ of copies;
- for $1 \leq j \leq k_b$, fix a pre-Taylor expansion (p^b_j, f^b_j) of $\Theta(b)$, and write $p^b_j = (\overrightarrow{c}^b_j; t^b_j, \overrightarrow{s}^b_j)$ so that $f^b_j(t^b_j) = T_b$;
- up to renaming the variables of the p^b_j's, ensure that the sets $\mathbf{V}(p^b_j)$ are pairwise disjoint, and also disjoint from $\mathbf{V}(\overrightarrow{C}) \cup \mathbf{V}(\overrightarrow{S}^\circ)$;
- $(\overrightarrow{c}; \overrightarrow{t})$ is obtained from $(\overrightarrow{C}; \overrightarrow{S}^\circ)$ by replacing each a^b_0 with $!(t^b_1, \dots, t^b_{k_b})$ and each a^b_{i+1} with an enumeration of $\bigcup_{j=1}^{k_b} (f^b_j)^{-1}(S^\circ_{b,i+1})$ — thus increasing the arity of the ?-connective having a^b_{i+1} as a premise, or increasing the number of trees in $\overrightarrow{t}$ if $a^b_{i+1} \in \overrightarrow{S}^\circ$ — and then concatenating $\overrightarrow{c}^b_j$ for $b \in \mathcal{B}_\Theta$ and $1 \leq j \leq k_b$;
- for $t \in \overrightarrow{t}$, set $f(t) = a^b_{i+1}$ if $f^b_j(t) = S^\circ_{b,i+1}$ for some j, otherwise let $f(t)$ be the only pre-tree of $\overrightarrow{S}^\circ$ such that t is obtained from $f(t)$ by the previous substitution.

The *Taylor expansion*[5] of a net P is then $\mathcal{T}(P) = \{p \mid (p, f) \text{ is a pre-Taylor expansion of } P\}$.

5.3 Paths in Taylor expansion

In the following, we fix a pre-Taylor expansion (p, f) of $P^\circ = (\Theta; \overrightarrow{C}; \overrightarrow{S}^\circ)$, and we describe the structure of paths in p. Observe that if $t \in \mathbf{T}(p)$ then:

- either t is at top level, *i.e.* t is obtained from some $T \in \mathbf{T}(P^\circ) \setminus \mathcal{A}$ by substituting box ports with trees from resource nets, and then we say t is *outer* and write $t^* = T$;
- or t is in a copy of a box, *i.e.* $t \in \mathbf{T}(p^b_j)$ for some $b \in \mathcal{B}_\Theta$ and $1 \leq j \leq k_b$, and then we say t is *inner* and write $\beta(t) = b$ and $\iota(t) = (b, j)$;

[5] More extensive presentations of Taylor expansion of MELL nets exist in the literature, in various styles [19, 17, 6]. Our only purpose here is to introduce sufficient notations to present our analysis of the length of paths in $\mathcal{T}(P)$ by a function of the size of P.

- or t is a *cocontraction*, *i.e* $t = !(t^b_1, \dots, t^b_{k_b})$ for some $b \in \mathcal{B}_\Theta$, and then we write $\beta(t) = b$ and $t = !_b$.

We moreover distinguish the *boundaries*, *i.e.* the cocontractions of p, together with all the elements of the families $\overrightarrow{s}^b_j$ of Definition 12: we write $\lfloor !_b \rfloor = a^b_0$ and $\lfloor s \rfloor = f(s)$ if $s \in \overrightarrow{s}^b_j$.

We say a subpath $\xi = t_1, \dots, t_n$ of $\chi \in \mathbf{P}(p)$ is an *inner subpath* (resp. an *outer subpath*) if each t_i is inner (resp. outer), and ξ is a *box subpath* if each t_i is inner or a cocontraction.

▶ **Lemma 13.** *If $\xi = t_0, \dots, t_n$ is an inner path of p then $\iota(t_i) = \iota(t_j)$ for all i and j. We then write $\beta(\xi) = b$ and $\iota(\xi) = (b, j)$.*

Proof. If $t \sim s$ and t and s are both inner then $\iota(t) = \iota(u)$. ◀

▶ **Lemma 14.** *If ξ is a box path of p then ξ is an inner path or there is $b \in \mathcal{B}_\Theta$ such that $\xi = \chi_1, !_b, \chi_2$ with χ_1 and χ_2 inner subpaths. In the latter case: if $\chi_1 \neq \epsilon$ then $\beta(\chi_1) = b$; if $\chi_2 \neq \epsilon$ then $\beta(\chi_2) = b$; and $\iota(\chi_1) \neq \iota(\chi_2)$ in case both subpaths are non empty.*

Proof. If $t \sim s$ and t and s are both inner then $\iota(t) = \iota(u)$; if $t \sim !_b$ and t is inner then $\beta(t) = b$; and no other adjacency relation can hold between the elements of a box path. ◀

▶ **Lemma 15.** *If $\xi = t_0, \dots, t_n$ is outer then $\xi^* = t^*_0, \dots, t^*_n \in \mathbf{P}(P^\circ)$.*

Proof. If t and s are outer, then $t \sim^{p,I}_l s$ iff $t^* \sim^{P^\circ, I^*}_{l^*} s^*$, where I^* is obtained by restricting I to outer trees and then composing with $-^*$. Moreover, $-^*$ is injective. ◀

▶ **Lemma 16.** *Assume $\xi = \xi_0, \chi_1, \xi_1, \dots, \chi_n, \xi_n \in \mathbf{P}(p)$ where each χ_i is a box path and each ξ_i is outer. Then we can write $\chi_i = u_i, \chi'_i, v_i$ where u_i and v_i are boundaries. Moreover, $\beta(\chi_i) \neq \beta(\chi_j)$ when $i \neq j$, and we obtain $\xi^* = \xi^*_0, \lfloor u_1 \rfloor, \lfloor v_1 \rfloor, \xi^*_1, \dots, \lfloor u_n \rfloor, \lfloor v_n \rfloor, \xi^*_n \in \mathbf{P}(P^\circ)$.*

Proof. The proof is by induction on n. If $n = 0$, *i.e.* ξ is outer, then we conclude by the previous lemma. We can thus assume $n > 0$.

The endpoints of χ_i are boundaries, because χ_i is a box path and the endpoints of ξ_{i-1} and ξ_i are outer. Since each boundary is adjacent to at most one outer tree, of which it is an immediate subtree or against which it is cut, χ_i is not reduced to a single boundary. For $1 \leq i \leq n$, write $\chi_i = (u_i, \chi'_i, v_i)$.

Write $b_i = \beta(\chi_i)$. Observe that, up to $-^*$, the only new adjacency relations in ξ^* are the $\lfloor u_i \rfloor \sim_{b_i} \lfloor v_i \rfloor$ for $1 \leq i \leq n$. Hence, to conclude that ξ^* is indeed a path, it will be sufficient to prove that $b_i \neq b_j$ when $i \neq j$. If $i < j$ then, by applying the induction hypothesis, we obtain $\zeta = \xi^*_i, \dots, \lfloor u_{j-1} \rfloor, \lfloor v_{j-1} \rfloor, \xi^*_{j-1} \in \mathbf{P}(P^\circ)$. Then, if we had $b_i = b_j$, we would obtain a cycle $\lfloor v_i \rfloor, \zeta, \lfloor u_j \rfloor, \lfloor v_i \rfloor$ in P°, which is a contradiction. ◀

From Lemma 16, we can derive that p is acyclic as soon as P° is. Indeed, if ξ is a cycle in p:

- either there is a tree at top level in ξ and we can apply Lemma 16 to obtain a cycle in P°;
- or ξ is an inner path, and we proceed inductively in $\Theta(\beta(\xi))$.

Our final result is a quantitative version of this corollary: not only there is no cycle in p but the length of paths in p is bounded by a function of P°. If $\xi = t_1, \dots, t_n$, we write $|\xi| = n$ for the *length* of ξ.

▶ **Theorem 17.** *If $p \in \mathcal{T}(P^\circ)$ and $\xi \in \mathbf{P}(p)$ then $|\xi| \leq 2^{\mathbf{depth}(P^\circ)} \mathbf{size}(P^\circ)$.*

Proof. Write $\xi = \xi_0, \chi_1, \xi_1, \ldots, \chi_n, \xi_n \in \mathbf{P}(p)$ where each χ_i is a box path and each ξ_i is an outer path.

Write $b_i = \beta(\chi_i)$. By Lemma 14, χ_i is either an inner path or of the form $\zeta_i, !_{b_i}, \zeta'_i$ with ζ_i and ζ'_i inner subpaths in b_i. By induction hypothesis applied to those inner subpaths, we obtain $|\chi_i| \leq 1 + 2 \times 2^{\mathbf{depth}(\Theta(b_i))}\mathbf{size}(\Theta(b_i))$.

Let ξ^* be as in Lemma 16: we have $|\xi^*| = 2n + \sum_{i=0}^{n}|\xi_i^*| \leq \#(P^\circ)$. It follows that $\sum_{i=0}^{n}|\xi_i| \leq \#(P^\circ) - 2n$.

We obtain: $|\xi| = \sum_{i=0}^{n}|\xi_i| + \sum_{i=1}^{n}|\chi_i| \leq \#(P^\circ) - 2n + \sum_{i=1}^{n}(1 + 2^{\mathbf{depth}(\Theta(b_i)+1)}\mathbf{size}(\Theta(b_i)))$ hence $|\xi| \leq \#(P^\circ) + \sum_{i=1}^{n} 2^{\mathbf{depth}(\Theta(b_i)+1)}\mathbf{size}(\Theta(b_i))$ and, since $\mathbf{depth}(\Theta(b_i)) < \mathbf{depth}(P^\circ)$, $|\xi| \leq 2^{\mathbf{depth}(P^\circ)}\left(\#(P^\circ) + \sum_{i=1}^{n}\mathbf{size}(\Theta(b_i))\right)$. We conclude recalling that $\mathbf{size}(P^\circ) = \#(P^\circ) + \sum_{b \in \mathcal{B}_\Theta}\mathbf{size}(\Theta(b))$. ◀

In particular, we obtain $\mathbf{cc}(p) \leq 2^{\mathbf{depth}(P^\circ)}\mathbf{size}(P^\circ)$.

5.4 Cut elimination in Taylor expansion

In resource nets, the elimination of the cut $\langle ?(t_1, \ldots, t_n) | !(s_1, \ldots, s_m) \rangle$ yields the finite sum $\sum_{\sigma : \{1,\ldots,n\} \xrightarrow{\sim} \{1,\ldots,m\}} \langle t_1 | s_{\sigma(1)} \rangle, \ldots, \langle t_n | s_{\sigma(n)} \rangle$. It turns out that the results of Sections 3 and 4 apply directly to resource nets: setting $\langle ?(t_1, \ldots, t_n) | !(s_1, \ldots, s_n) \rangle \to \langle t_1 | s_{\sigma(1)} \rangle, \ldots, \langle t_n | s_{\sigma(n)} \rangle$ for each permutation σ, we obtain an instance of multiplicative reduction, as the order of premises is irrelevant from a combinatorial point of view — this is all the more obvious because no typing constraint was involved in our argument. In other words, Corollary 11 also applies to the parallel reduction of resource nets. With Theorem 17, we obtain:

▶ **Corollary 18.** *If q is a resource net and P is an MELL net, $\{p \in \mathcal{T}(P); p \rightrightarrows q\}$ is finite.*

6 Conclusion

Recall that our original motivation was the definition of a reduction relation on infinite linear combinations of resource nets, simulating cut elimination in MELL through Taylor expansion. We claim that a suitable notion is as follows:

▶ **Definition 19.** Write $\sum_{i\in I} a_i p_i \Rightarrow \sum_{i \in I} a_i q_i$ as soon as:

- for each $i \in I$, the resource net p_i reduces to q_i (which may be a finite sum);
- for any resource net q, there are finitely many $i \in I$ such that q is a summand of q_i.

In particular, if $\sum_{i\in I} a_i p_i$ is a Taylor expansion, then Theorem 18 ensures that the second condition of the definition of $\Rightarrow$ is automatically valid. The details of the simulation in a quantitative setting remain to be worked out, but the main stumbling block is now over: the necessary equations on coefficients are well established, as they have been extensively studied in the various denotational models; it only remained to be able to form the associated sums directly in the syntax.

Let us mention that another important incentive to publish our results is the *normalization-by-evaluation* programme that we develop with Guerrieri, Pellissier and Tortora de Falco [1] — which is limited to strict nets for independent reasons. Indeed, if P is cut-free, the elements of the semantics of P are in one-to-one correspondence with $\mathcal{T}(P)$. Then, given a sequence $P_1, \ldots, P_n$ of MELL nets such that P_i reduces to P_{i+1} by cut elimination and P_n is normal, from $p_n \in \mathcal{T}(P_n)$ we can construct a sequence $p_1, \ldots, p_{n-1}$ of resource nets, such that each $p_i \in \mathcal{T}(P_i)$ and $p_i \rightrightarrows p_{i+1}$. Then our results ensure that $\#p_1$ is bounded by a function of n, $\mathbf{size}(P_1)$ and $\#p_n$, which is a crucial step of our construction.

We finish the paper by reviewing the restrictions that we imposed on our framework. Strictness is not an essential condition for the main results to hold. It is possible to deal with units and weakenings (0-ary ⅋, ⊗ and ? nodes), and then with complete Taylor expansion, including 0-ary developments of boxes (generating weakenings and coweakenings). In this case, we need to introduce additional structure — jumps from weakenings, that can be part of switching paths — and some other constraint — a bound on the number of weakenings that can jump to a given tree. The proof is naturally longer, and the bounds much greater, but the finiteness property still holds. We leave a formal treatment of this extension for further work.

The other notable constraint is the use of the *nouvelle syntaxe*, with generalized exponential links. It is also possible to deal with a standard representation, including separate derelictions and coderelictions, with a finer grained cut elimination procedure. This introduces additional complexity in the formalism but, by contrast with lifting the strictness condition, it essentially requires no new concept or technique: the difficulty in parallel reduction is to control the chains of cuts to be simultaneously eliminated, and decomposing cut elimination into finer reduction steps can only decrease the length of such chains.

References

1. Jules Chouquet, Giulio Guerrieri, Luc Pellissier, and Lionel Vaux. Normalization by evaluation in linear logic. In Stefano Guerrini, editor, *Preproceedings of the International Workshop on Trends in Linear Logic and Applications, TLLA*, September 2017.
2. Vincent Danos and Thomas Ehrhard. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Inf. Comput.*, 209(6):966–991, 2011. URL: `https://doi.org/10.1016/j.ic.2011.02.001`, `doi:10.1016/j.ic.2011.02.001`.
3. Vincent Danos and Laurent Regnier. The structure of multiplicatives. *Arch. Math. Log.*, 28(3):181–203, 1989.
4. Daniel de Carvalho. *Sémantiques de la logique linéaire et temps de calcul*. PhD thesis, Université d'Aix-Marseille II, Marseille, France, 2007.
5. Daniel de Carvalho. Execution time of lambda-terms via denotational semantics and intersection types. *CoRR*, abs/0905.4251, 2009. URL: `http://arxiv.org/abs/0905.4251`, `arXiv:0905.4251`.
6. Daniel de Carvalho. The relational model is injective for multiplicative exponential linear logic. In Jean-Marc Talbot and Laurent Regnier, editors, *25th EACSL Annual Conference on Computer Science Logic, CSL 2016, August 29 - September 1, 2016, Marseille, France*, volume 62 of *LIPIcs*, pages 41:1–41:19. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. URL: `https://doi.org/10.4230/LIPIcs.CSL.2016.41`, `doi:10.4230/LIPIcs.CSL.2016.41`.
7. Thomas Ehrhard. On köthe sequence spaces and linear logic. *Mathematical Structures in Computer Science*, 12(5):579–623, 2002. URL: `https://doi.org/10.1017/S0960129502003729`, `doi:10.1017/S0960129502003729`.
8. Thomas Ehrhard. Finiteness spaces. *Mathematical Structures in Computer Science*, 15(4):615–646, 2005. URL: `https://doi.org/10.1017/S0960129504004645`, `doi:10.1017/S0960129504004645`.
9. Thomas Ehrhard. A finiteness structure on resource terms. In *Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, 11-14 July 2010, Edinburgh, United Kingdom*, pages 402–410, 2010.
10. Thomas Ehrhard. A new correctness criterion for MLL proof nets. In *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the*

Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014, pages 38:1–38:10, 2014.

11 Thomas Ehrhard. An introduction to differential linear logic: proof-nets, models and antiderivatives. *CoRR*, abs/1606.01642, 2016.

12 Thomas Ehrhard and Laurent Regnier. The differential lambda-calculus. *Theor. Comput. Sci.*, 309(1-3):1–41, 2003.

13 Thomas Ehrhard and Laurent Regnier. Differential interaction nets. *Electr. Notes Theor. Comput. Sci.*, 123:35–74, 2005.

14 Thomas Ehrhard and Laurent Regnier. Uniformity and the taylor expansion of ordinary lambda-terms. *Theor. Comput. Sci.*, 403(2-3):347–372, 2008.

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

16 Jean-Yves Girard. Normal functors, power series and lambda-calculus. *Annals of Pure and Applied Logic*, 37(2):129, 1988.

17 Giulio Guerrieri, Luc Pellissier, and Lorenzo Tortora de Falco. Computing connected proof(-structure)s from their taylor expansion. In *1st International Conference on Formal Structures for Computation and Deduction, FSCD 2016, June 22-26, 2016, Porto, Portugal*, pages 20:1–20:18, 2016.

18 Jim Laird, Giulio Manzonetto, Guy McCusker, and Michele Pagani. Weighted relational models of typed lambda-calculi. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28, 2013*, pages 301–310. IEEE Computer Society, 2013. URL: `https://doi.org/10.1109/LICS.2013.36`, `doi:10.1109/LICS.2013.36`.

19 Michele Pagani and Christine Tasson. The inverse taylor expansion problem in linear logic. In *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009, 11-14 August 2009, Los Angeles, CA, USA*, pages 222–231, 2009.

20 Michele Pagani, Christine Tasson, and Lionel Vaux. Strong normalizability as a finiteness structure via the taylor expansion of lambda-terms. In *Foundations of Software Science and Computation Structures - 19th International Conference, FOSSACS 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings*, pages 408–423, 2016.

21 Laurent Regnier. *Lambda-calcul et réseaux*. PhD thesis, Université Paris 7, Paris, France, December 1992.

22 Christine Tasson. *Sémantiques et syntaxes vectorielles de la logique linéaire*. PhD thesis, Université Paris Diderot, Paris, France, December 2009.

23 Lionel Vaux. Taylor expansion, lambda-reduction and normalization. In *26th EACSL Annual Conference on Computer Science Logic, CSL 2017, August 20-24, 2017, Stockholm, Sweden*, pages 39:1–39:16, 2017.