Linear Logic by Levels and Bounded Time Complexity

Patrick Baillot[a,1], Damiano Mazza[b,1]

[a]LIP - UMR 5668 CNRS - ENS Lyon, F-69364 Lyon, France
[b]LIPN - UMR 7030 CNRS - Université Paris 13, F-93430 Villetaneuse, France

Abstract

We give a new characterization of elementary and deterministic polynomial time computation in linear logic through the proofs-as-programs correspondence. Girard's seminal results, concerning elementary and light linear logic, achieve this characterization by enforcing a *stratification principle* on proofs, using the notion of *depth* in proof nets. Here, we propose a more general form of stratification, based on inducing *levels* in proof nets by means of indexes, which allows us to extend Girard's systems while keeping the same complexity properties. In particular, it turns out that Girard's systems may be recovered by forcing depth and level to coincide. A consequence of the higher flexibility of levels with respect to depth is the absence of boxes for handling the paragraph modality. We use this fact to propose a variant of our polytime system in which the paragraph modality is only allowed on atoms, and which may thus serve as a basis for developing lambda-calculus type assignment systems with more efficient typing algorithms than existing ones.

Key words: Implicit computational complexity, light linear logics, type systems for polynomial time.

Introduction

Linear logic and implicit computational complexity. The intersection between logic and implicit computational complexity is at least twofold, as there are at least two alternative views on logic itself: a first possibility is to see it as a *descriptive language*, i.e., as a language for expressing properties of mathematical objects; a second possibility is to see it, via the Curry-Howard isomorphism, as a *programming language*, i.e., a tool for computing functions. These two views closely correspond to two fundamental branches of mathematical logic: model theory, and proof theory, respectively. The first approach has been taken quite successfully by what is known as *descriptive computational complexity*. The idea of exploring the second approach is more recent: the first results of this kind

Email addresses: `patrick.baillot@ens-lyon.fr` (Patrick Baillot), `damiano.mazza@lipn.univ-paris13.fr` (Damiano Mazza)

[1]Partially supported by project NOCoST (ANR, JC05_43380).

can be found in Leivant (1994), Leivant and Marion (1993), and in the work of Girard et al. (1992), to which the present work is more closely related.

As mentioned above, the use of logic as a programming language capturing certain complexity classes passes through the Curry-Howard isomorphism: a proof is a program, whose execution is given by cut-elimination; therefore, the idea is to define a logical system whose cut-elimination procedure has a bounded complexity, so that the algorithms programmable in this logical system intrinsically have that complexity, i.e., the system is *sound* w.r.t. a complexity class.

Due to its "resource awareness", linear logic (Girard, 1987) is the ideal setting to attempt this. In fact, linear logic brings to light the logical primitives which are responsible for the complexity of cut-elimination, under the form of modalities, called *exponential*. These are in control of duplication during the cut-elimination process; by restraining the rules for these modalities, one achieves the desired goal. Of course one has to make sure that the resulting system is also *complete*, i.e., that *all* functions of the given complexity can be programmed in it. This methodology has been successfully followed to characterize complexity classes like deterministic polynomial time (Girard et al., 1992; Girard, 1998; Asperti and Roversi, 2002; Lafont, 2004), elementary time (Girard, 1998; Danos and Joinet, 2003), deterministic logarithmic space (Schöpp, 2007), and, very recently, polynomial space (Gaboardi et al., 2008).

Stratification. In this work we focus on Girard's (1998) elementary linear logic (**ELL**) and light linear logic (**LLL**), corresponding to elementary time and deterministic polynomial time, respectively.

The complexity bound on the cut-elimination procedure of these systems relies on a principle called *stratification*, which is also at the base of other approaches to implicit computational complexity, both related to logic and not.

Stratification can be interpreted in at least three informal ways. The first, which is where Girard (1998) originally drew inspiration from, comes from a sharp analysis of Russel's paradox in naive set theory (Curry and Feys, 1958), and was first considered by Leivant (1994). Unrestricted comprehension can be obtained as a theorem in first order classical logic plus the following two rules:

$$\frac{\vdash \Gamma, A[t/x]}{\vdash \Gamma, t \in \{x \mid A\}} \qquad\qquad \frac{\vdash \Gamma, \neg A[t/x]}{\vdash \Gamma, t \notin \{x \mid A\}}$$

where $\{x \mid A\}$ is the standard set-builder notation for the set containing all and only the elements satisfying the formula A. Russel's antinomy is obtained by considering the term $r = \{x \mid x \notin x\}$, from which we build the formula $R = r \in r$. One can see that R is a fixpoint of negation, i.e., R is provably equivalent to $\neg R$. In fact, one can obtain $\vdash \Gamma, R$ from $\vdash \Gamma, \neg R$ by applying the rule above on the left, and $\vdash \Gamma, \neg R$ from $\vdash \Gamma, R$ by applying the rule above on the right. The empty sequent, i.e., a contradiction, can then be derived as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{}{\vdash \neg R, R}}{\vdash \neg R, \neg R}}{\vdash \neg R} \qquad \dfrac{\dfrac{\dfrac{}{\vdash \neg R, R}}{\vdash R, R}}{\vdash R}}{\vdash}$$

Remark that contraction is necessary: in multiplicative linear logic, where contraction is forbidden, the empty sequent cannot be derived even in presence of the self-contradicting formula R (this was first observed by Grishin (1982)).

Another setting in which stratification can be applied is the λ-calculus, where Russel's paradox corresponds to the diverging term Ω. The fundamental construct behind this term is self-application, which, from the logical point of view, also needs contraction.

A third intuition comes from recursion theory, where more and more complex functions can be obtained by diagonalization. For instance, if $P_m(n)$ is a sequence of polynomial functions of degree m in n (for example, $P_m(n) = n^m$), the function $P_n(n)$ is super-exponential, i.e., elementary; if $\theta_m(n)$ is a sequence of elementary functions in n whose complexity rises with m (for example, $\theta_m(n) = 2^n_m$, i.e., a tower of exponentials of height m in n), then $\theta_n(n)$ is hyper-exponential, i.e., non-elementary.

In all of these incarnations, stratification can be seen as a way of forbidding the identification of two variables, or the contraction of two formulas, because they belong to two morally different "levels": the occurrence of R coming from the axiom and that coming from the application of the $\in$-rule in the derivation of Russel's paradox; the occurrence of x in function position and that in argument position in the self application $\lambda x.xx$; the index of the sequence and the argument of the members of the sequence in the diagonalization examples.

Note that stratification is reminiscent of the notion of *ramification*, or its variants like *safe recursion*, used for restricting primitive recursion in implicit computational complexity (Bellantoni and Cook, 1992; Leivant and Marion, 1993). The relation between safe recursion and light linear logic was investigated in Murawski and Ong (2004), while a study on diagonalization and complexity was recently carried out by Marion (2007).

Proof nets, boxes, and stratification. The bound on the cut-elimination procedure for **ELL** and **LLL** is proved using *proof nets*, a graphical representation of proofs (Girard, 1996). These are a crucial tool for applying linear logic to implicit computational complexity: they allow a fine-grained analysis of cut-elimination, the definition of adequate measures and invariants, and the introduction of adapted reduction strategies. In particular, the fundamental stratification property of **ELL** and **LLL** is defined and enforced through *boxes*, a construct in the syntax of proof nets corresponding to the rules for exponential modalities. Boxes have been around since the introduction of proof nets (Girard, 1987) and can be understood intuitively in two ways:

(i) *logically*: they correspond to sequentiality information;
(ii) *operationally*: they mark subgraphs (i.e., subproofs) that can be duplicated.

Boxes can be nested; as a consequence, a node in a proof net (corresponding to a logical rule) may be assigned an *exponential depth*, which is the number of nested boxes containing that node. Stratification is achieved precisely on the base of the exponential depth: in full linear logic, two occurrences of the same formula introduced at different exponential depths may eventually be contracted; in **ELL** and **LLL**, they cannot. From the operational point of view, boxes therefore assume a twofold role in **ELL** and **LLL**: they serve for the purpose (ii) explained above, and they enforce stratification.

A new stratification. The main contribution of this work is the proposal of an alternative way to achieve stratification, which is orthogonal to boxes. It is a direct application of the intuitions concerning stratification given above: occurrences of formulas in a proof net are "tested" by assigning to them an *index*, which must satisfy certain constraints; in particular, if two occurrences of the same formula are contracted, then they must have the same index. If the proof net "passes the test", i.e., if there is a way of assigning indexes to its formulas in a way which is compatible with the constraints, then the proof net is accepted.

The assignment of indexes naturally determines the stratification of a proof net into *levels*, which need not match exponential depths. We obtain in this way a logical system, called *linear logic by levels* ($\mathbf{L^3}$), which can be proved to correspond to elementary time in the same sense as **ELL**. There is actually more: **ELL** turns out to be the subsystem of $\mathbf{L^3}$ in which levels and depths coincide, so Girard's approach to stratification can be seen as a special case of our own.

As said above, the main novelty of $\mathbf{L^3}$ is that it shows how stratification and exponential depths must not necessarily be related. This is, in our opinion, an important contribution to the understanding of the principles underlying light logics. It may also be a starting point for finding new kinds of denotational semantics for bounded time computation, extending the ideas of Baillot (2004) and Laurent and Tortora de Falco (2006).

Removing useless boxes. In **LLL**, along the exponential modalities of linear logic, an additional exponential modality, the *paragraph* §, must be added in order to reach the desired expressive power, i.e., programming all polytime functions. Since stratification is linked to exponential depth, the paragraph modality too is handled in proof nets by means of boxes; however, §-boxes cannot be duplicated, so they lose their original function (ii), and their existence is only justified by stratification.

By imposing on our $\mathbf{L^3}$ the same kind of constraints that define **LLL** from **ELL**, we obtain *light linear logic by levels* ($\mathbf{L^4}$), which, as expected, characterizes deterministic polynomial time. This system offers an additional advantage with respect to **LLL**: since our stratification is orthogonal to boxes, and since §-boxes exist only to enforce stratification, these are no longer needed in $\mathbf{L^4}$.

Improving type systems. In several cases, the characterization of complexity classes with subsystems of linear logic has allowed, in a second step, to define type systems for the λ-calculus statically ensuring complexity properties (Baillot and Terui, 2004; Gaboardi and Ronchi Della Rocca, 2007): if a λ-term, expecting for instance a binary list argument, is well typed, then it admits a complexity bound w.r.t. the size of the input. Such results naturally call for type inference procedures (Coppola and Martini, 2006; Atassi et al., 2007), which can be seen as tests for sufficient conditions for a program to admit a complexity bound.

From this point of view, the presence of §-boxes in **LLL** is a heavy drawback: in fact, a large part of the work needed to perform type inference in **LLL**, or subsystems like **DLAL** (Atassi et al., 2007), comes from the problem of placing correctly §-boxes, in particular in such a way that they are compatible with other rules, or with λ bindings in the λ-calculus (remember that boxes also carry sequentialization information, cf. point (i) above). A system like $\mathbf{L^4}$ clearly offers the possibility of overcoming these problems: the absence of §-boxes may yield major simplifications in the development of type systems for polynomial time.

A further contribution of this paper is making a first step in that direction: exploiting the lack of sequentiality constraints on the paragraph modality, we devise a variant of $\mathbf{L^4}$ in which the paragraph modality is hidden in atomic formulas; as a consequence, the paragraph modality completely disappears from this system, and there is no need for a rule handling it. This may turn out to be extremely helpful for designing a type system out of our work.

Plan of the paper. Sect. 1 contains a sort of mini-crash-course on linear logic and its light subsystems **ELL** and **LLL**. Apart from introducing the material necessary to our work, this (quite lengthy) section should make the paper as self-contained as possible, and hopefully accessible to the reader previously unfamiliar with these topics. The systems $\mathbf{L^3}$ and $\mathbf{L^4}$ are introduced in Sect. 2, and their relationship with **ELL** and **LLL** is spelled out. Sect. 3 is the technical core of the paper: it contains the proof of the complexity bounds for $\mathbf{L^3}$ (Theorem 17) and $\mathbf{L^4}$ (Theorem 24), from which the characterization result follows (Theorem 26). Sect. 4 introduces the variant of $\mathbf{L^4}$ without paragraph modality; the main result of this section is Theorem 37. In Sect. 5 we conclude the paper with a discussion about open questions and future work.

Acknowledgments. The authors would like to thank Daniel De Carvalho for his useful comments and suggestions on the subject of this paper.

1. Multiplicative Exponential Linear Logic

1.1. Formulas

The formulas of second order unit-free multiplicative exponential linear logic (**meLL**) are generated by the following grammar, where $X, X^\perp$ range over a denumerable set of propositional variables:

$$A, B ::= X \mid X^\perp \mid A \otimes B \mid A \parr B \mid !A \mid ?A \mid \exists X.A \mid \forall X.A \mid \S A.$$

Linear negation is defined through De Morgan laws:

$$
\begin{array}{rclcrcl}
(X)^\perp & = & X^\perp & \qquad & (X^\perp)^\perp & = & X \\
(A \otimes B)^\perp & = & B^\perp \parr A^\perp & & (A \parr B)^\perp & = & B^\perp \otimes A^\perp \\
(!A)^\perp & = & ?A^\perp & & (?A)^\perp & = & !A^\perp \\
(\exists X.A)^\perp & = & \forall X.A^\perp & & (\forall X.A)^\perp & = & \exists X.A^\perp
\end{array}
$$

$$(\S A)^\perp = \S A^\perp$$

Note that the paragraph modality is not present in the standard definition of **meLL** (Girard, 1987); we include it here for convenience. Also observe that full linear logic has a further pair of dual binary connectives, called *additive* (denoted by & and $\oplus$), which we shall briefly discuss in Sect. 5. They are not strictly needed for our purposes, hence we restrict to **meLL** in the paper.

Linear implication is defined as $A \multimap B = A^\perp \parr B$. Multisets of formulas will be ranged over by $\Gamma, \Delta, \ldots$

For technical reasons, it is also useful to consider *discharged formulas*, which will be denoted by $\flat A$, where A is a formula.

1.2. Proofs

Sequent calculus and cut-elimination. The proof theory of **meLL** can be formulated using the sequent calculus of Table 1. This calculus, which can be shown to enjoy cut-elimination, differs from the one originally given by Girard (1987) because of the addition of the last three rules. All of them are added for convenience. The paragraph rule actually makes this modality trivial, as expressed by the following:

Proposition 1. *For any A, $\S A$ is provably isomorphic to A in* **meLL**.

Proof. It is not hard to see that there are two derivations D_1, D_2 of $\vdash \S A^\perp, A$ and $\vdash A^\perp, \S A$, from which one can obtain two derivations of $\vdash \S A \multimap A$ and $\vdash A \multimap \S A$, respectively. Moreover, the derivations obtained by cutting D_1 with D_2 in the two possible ways both reduce to the identity (i.e., an axiom modulo η-expansion) after cut-elimination. □

Nevertheless, we shall consider subsystems of **meLL** in which the paragraph modality is not trivial, and this is why we find it convenient to include it right from the start. The mix rule, and its nullary version (here called the daimon rule), are discussed more thoroughly at the end of this section. Basically, their presence simplifies the presentation of proof nets.

This last point is very important to us. In fact, the backbone of our work is a detailed analysis, in terms of computational complexity, of the cut-elimination procedure of **meLL**. In sequent calculus, this is composed of rules which are suitable reformulations of those originally given by Gentzen (1934) to prove his *Hauptsatz* for classical logic (the calculus **LK**). As a consequence, most of them are commutations, i.e., rules permuting a cut with another inference rule; only a few of them act on derivations in a non-trivial way. This is why we consider proof nets, an alternative presentation of the proof theory of **meLL** offering, among other things, the advantage of formulating cut-elimination without commutations: only the "interesting" rules are left.

$$\frac{}{\vdash A^\perp, A}\ \text{Axiom} \qquad \frac{\vdash \Gamma, A \quad \vdash \Delta, A^\perp}{\vdash \Gamma, \Delta}\ \text{Cut}$$

$$\frac{\vdash \Gamma, A \quad \vdash \Delta, B}{\vdash \Gamma, \Delta, A \otimes B}\ \text{Tensor} \qquad \frac{\vdash \Gamma, A, B}{\vdash \Gamma, A \parr B}\ \text{Par}$$

$$\frac{\vdash \Gamma, A}{\vdash \Gamma, \forall X.A}\ \text{For all } (X \text{ not free in } \Gamma) \qquad \frac{\vdash \Gamma, A[B/X]}{\vdash \Gamma, \exists X.A}\ \text{Exists}$$

$$\frac{\vdash ?\Gamma, A}{\vdash ?\Gamma, !A}\ \text{Promotion} \qquad \frac{\vdash \Gamma, A}{\vdash \Gamma, ?A}\ \text{Dereliction}$$

$$\frac{\vdash \Gamma}{\vdash \Gamma, ?A}\ \text{Weakening} \qquad \frac{\vdash \Gamma, ?A, ?A}{\vdash \Gamma, ?A}\ \text{Contraction}$$

$$\frac{\vdash \Gamma, A}{\vdash \Gamma, \S A}\ \text{Paragraph}$$

$$\frac{}{\vdash}\ \text{Daimon} \qquad \frac{\vdash \Gamma \quad \vdash \Delta}{\vdash \Gamma, \Delta}\ \text{Mix}$$

Table 1: The rules for **meLL** sequent calculus.

Proof nets. The proof net formalism was introduced by Girard (1987, 1996), and subsequently reformulated by other authors using slightly different syntactical definitions. In this paper, we use a combination of the presentations given by Danos and Regnier (1995) and Tortora de Falco (2003), with a slight change in the terminology: the term "proof structure", introduced by Girard (1987) and traditionally used in the literature, is here dismissed in favor of the term *net*. On the contrary, the term proof net, i.e., a net satisfying certain structural conditions (the correctness criterion), retains its usual meaning.

In the following definition, and throughout the rest of the paper, unless explicitly stated we shall make no distinction between the concepts of *formula* and *occurrence of formula*. The same will be done for what we call *links* and their occurrences.

Definition 1 (Net). A *pre-net* is a pair $(\mathcal{G}, \mathsf{B})$, where $\mathcal{G}$ is a finite graph-like object whose nodes are occurrences of what we call *links*, and whose edges are directed and labelled by formulas or discharged formulas of **meLL**; and B is a set of subgraphs of $\mathcal{G}$ called *boxes*.

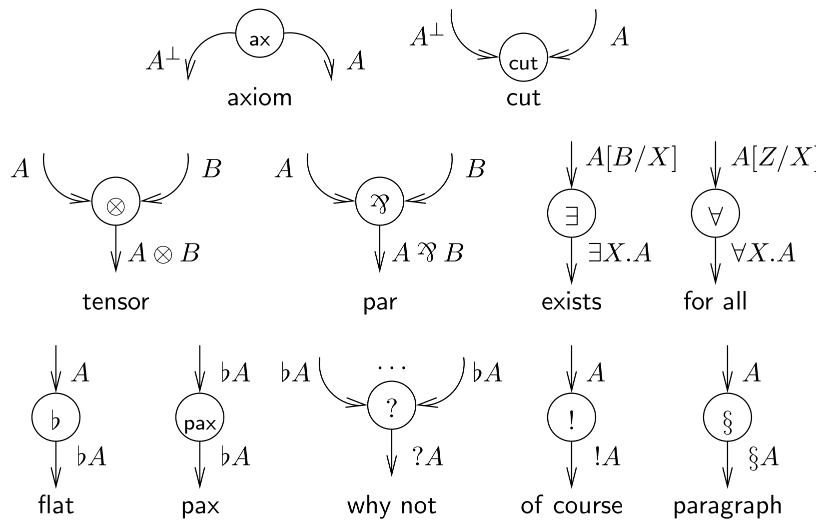

Figure 1: Links.

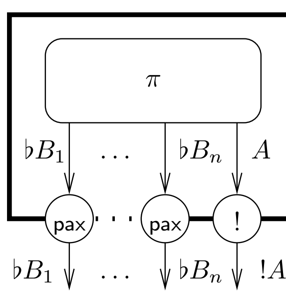

Figure 2: A box.

- Each link (Fig. 1) has an arity and co-arity, which are resp. the number of its incoming and outgoing edges. The arity and co-arity is fixed for all links except why not links, which have co-arity 1 and arbitrary arity. A nullary why not link is also referred to as a weakening link. Par and for all links are called *jumping* links.
- The incoming edges of a link (and the formulas that label them) are referred to as its *premises*, and are assumed to be ordered, with the exception of cut and why not links; the outgoing edges of a link (and the formulas that label them) are referred to as its *conclusions*.
- Premises and conclusions of links must respect a precise labeling (which depends on the link itself), given in Fig. 1. In particular:

– edges labelled by discharged formulas can only be premises of pax and why not links;

– in a for all link l, the variable Z in its premise $A[Z/X]$ is called the *eigenvariable* of l. Each for all link is assumed to have a different eigenvariable.

– in an exists link l, the formula B in its premise $A[B/X]$ is said to be *associated* with l.

- Each edge must be the conclusion of exactly one link, and the premise of at most one link. The edges that are not premises of any link (and the formulas that label them) are deemed *conclusions* of the pre-net. (Note that the presence of these "pending" edges, together with the fact that some premises are ordered, is why pre-nets are not exactly graphs).

- A box is depicted as in Fig. 2, in which π is a pre-net, said to be *contained* in the box. The links that are explicitly represented in Fig. 2 (i.e., the pax links and the of course link) form the *border* of the box. The unique of course link in the border is called the *principal port* of the box, while the pax links are called *auxiliary ports*. We have the following conditions concerning boxes:

  a. each of course link is the principal port of exactly one box;
  b. each pax link is in the border of exactly one box;
  c. any two distinct boxes are either disjoint or included in one another.

A *net* is a pre-net such that in its conclusions there is no discharged formula, nor any formula containing an eigenvariable.

Definition 2 (Depth, size). Let σ be a pre-net.

- A link (or edge) of σ is said to have *depth* d if it is contained in d (necessarily nested) boxes. The depth of a box of σ is the depth of the links forming its border. The depth of a link l, edge e, or box $\mathcal{B}$ are denoted resp. by $\mathrm{d}(l)$, $\mathrm{d}(e)$ and $\mathrm{d}(\mathcal{B})$. The depth of σ, denoted by $\mathrm{d}(\sigma)$, is the maximum depth of its links.

- The *size* of σ, denoted by $|\sigma|$, is the number of links contained in σ, excluding auxiliary ports.

Definition 3 (Switching). Let σ be a pre-net. For each jumping link l of σ, we define the set of *jumps* of l, denoted by $J(l)$, as follows:

par: $J(l)$ is the set containing the links whose conclusions are the premises of l.

for all: if Z is the eigenvariable of l, $J(l)$ is the set containing:

- the link whose conclusion is the premise of l;
- any link whose conclusion is labelled by a formula containing Z;

- any exists link whose associated formula contains Z.

A *switching* of σ is an undirected graph built as follows:

- the conclusions of σ are erased, and its edges considered as undirected;
- for each jumping link l, the premises of l (if any) are erased, exactly one node $m \in J(l)$ is chosen and a new edge between m and l is added.
- the boxes at depth zero of σ are collapsed into single nodes, i.e., if $\mathcal{B}$ is a box at depth zero of σ, it is erased together with all the edges connecting its links to the rest of the graph, and replaced with a new node l; then, for any link m of depth zero which was connected to a link of $\mathcal{B}$, a new edge between m and l is added.

Definition 4 (Proof net). A pre-net $(\mathcal{G}, \mathsf{B})$ is *correct* iff:

- all of its switchings are acyclic;
- for all $\mathcal{B} \in \mathsf{B}$, the pre-net contained in $\mathcal{B}$ is correct.

A *proof net* is a correct net.

Sequent calculus and proof nets. The relationship between sequent calculus and proof nets is clarified by the notion of *sequentializable* net, whose definition mimics the rules of sequent calculus:

Definition 5 (Sequentializable net). We define the set of *sequentializable nets* inductively: the empty net and the net consisting of a single axiom link are sequentializable (daimon and axiom); the juxtaposition of two sequentializable nets is sequentializable (mix); if σ, σ_1, σ_2 are sequentializable nets of suitable conclusions, the nets of Fig. 3 are sequentializable; if

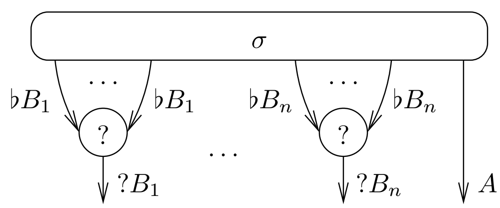

is a sequentializable net, then the net

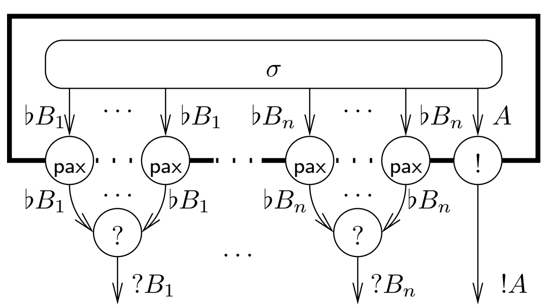

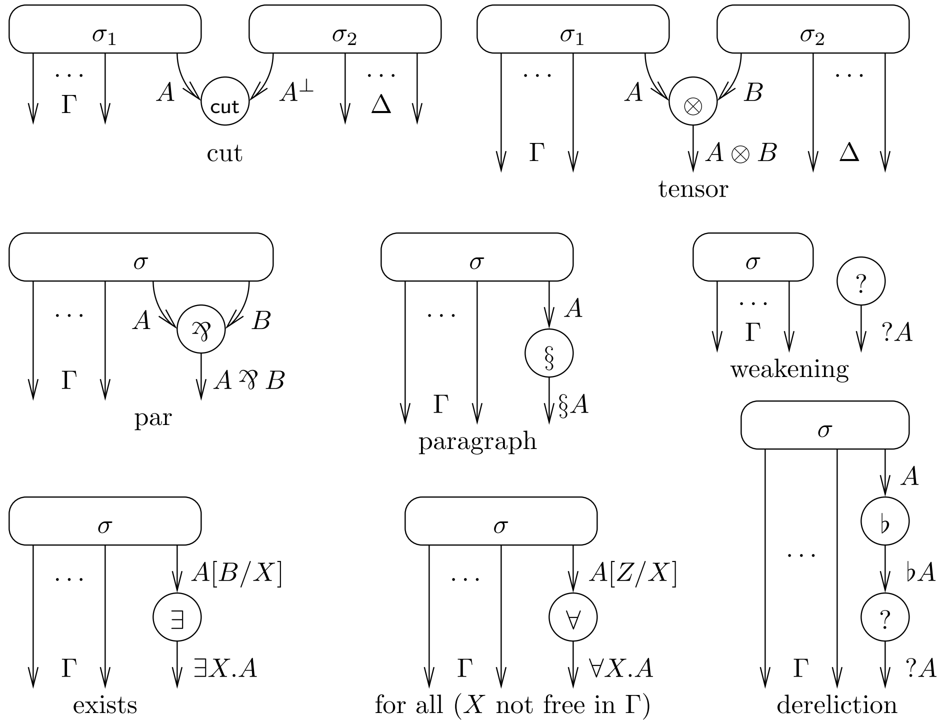

Figure 3: Rules for building sequentializable nets.

is sequentializable (promotion); if

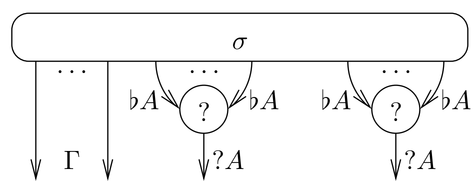

is a sequentializable net, then the net

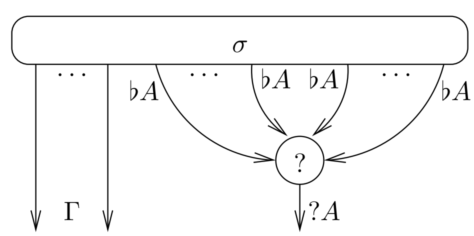

is sequentializable (contraction).

Proposition 2 (Girard (1996)). *A net is sequentializable iff it is a proof net.*

The above result, combined with Definition 5, gives a simple intuition for looking at proof nets: they can be seen as a sort of "graphical sequent calculus".

Cut-elimination. As anticipated above, formulating the cut-elimination procedure in proof nets is quite simple: there are only five rules (or *steps*, as we shall more often call them), taking the form of the graph-rewriting rules given in Figures 4 through 8. When a net π is transformed into π' by the application of one cut-elimination step, we write $\pi \to \pi'$, and we say that π *reduces* to π'. Of course, in that case, if π is a proof net, then π' is also a proof net, i.e., cut-elimination preserves correctness.

The following notions, taken from Tortora de Falco (2003), are needed to analyze the dynamics of proof nets under cut-elimination, and will prove to be quite useful in the sequel:

Definition 6 (Lift, residue). Whenever $\pi \to \pi'$, by simple inspection of the cut-elimination rules it is clear that any link l' of π' different from a cut comes from a unique ("the same") link l of π; we say that l is the *lift* of l', and that l' is a *residue* of l. We define the lift and residues of a box in the same way.

Untyped proof nets. We shall also use an untyped version of proof nets:

Definition 7 (Untyped proof net). An untyped pre-net is a directed graph with boxes built using the links of Fig. 1 as in Definition 1, but without any labels on edges, or any constraint induced by such labels. An untyped net is an untyped pre-net such that:

- the conclusion of a flat link must be the premise of a pax or $\mathsf{why\ not}$ link;
- the premise of a pax link must be the conclusion of a flat or pax link, and the conclusion of pax link must be the premise of a pax or $\mathsf{why\ not}$ link;
- the premises of a $\mathsf{why\ not}$ link must be conclusions of flat or auxiliary port links.

The notions of switching can be applied to untyped pre-nets without any change, and hence the notion of correctness. We then define an untyped proof net as a correct untyped net.

Observe that the cut-elimination steps of Figures 4 through 8 do not use formulas, i.e., they make sense also in an untyped framework. Therefore, cut-elimination can be defined also for untyped nets. Obviously, in the untyped case there may be "clashes", i.e., cut links connecting the conclusions of two non-dual links. In that case, the cut link is said to be *irreducible*; otherwise, we call it *reducible*. Hence, untyped proof nets may admit normal forms which are not cut-free.

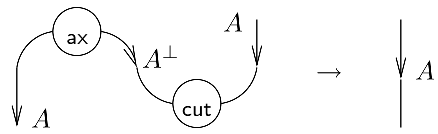

Figure 4: Axiom step.

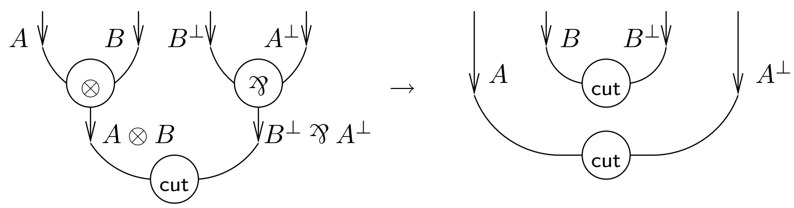

Figure 5: Multiplicative step.

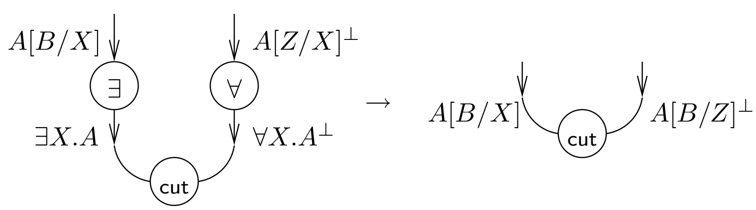

Figure 6: Quantifier step; the substitution is performed on the whole net.

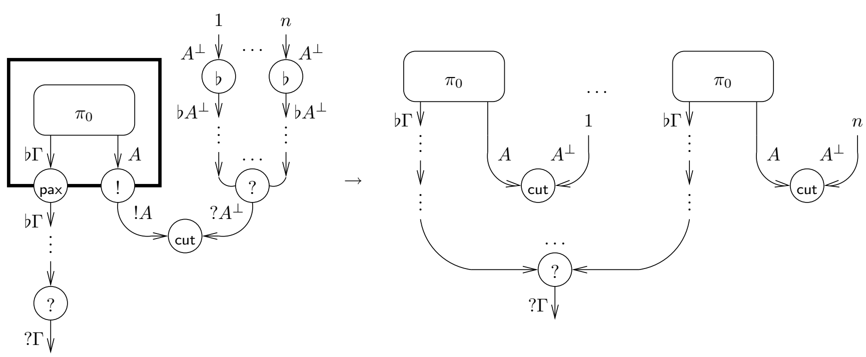

Figure 7: Exponential step; $\flat\Gamma$ is a multiset of discharged formulas, so one pax link, why not link, or wire in the picture may in some case stand for several (including zero) pax links, why not links, or wires.

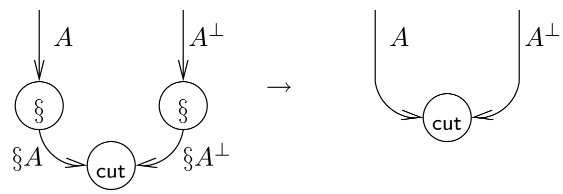

Figure 8: Paragraph step.

Remarks on mix and daimon. We mentioned above that admitting the mix and daimon rules makes the definition of proof nets simpler. The latter rule may actually be excluded quite easily: it is enough to modify Definition 4 by saying that switchings must be acyclic and non-empty. On the contrary, excluding mix has a much higher price: at present, all known solutions are quite cumbersome and bring up issues which are morally unproblematic but technically disturbing (see Tortora de Falco, 2003).

The status of the mix rule in the proof theory of linear logic is somewhat controversial (Girard, 2007). Its computational meaning is not clear, and no complexity-related subsystem of linear logic makes use of it. Its presence is harmless though: as a matter fact, while we shall explicitly rely on the acyclicity condition of Definition 4 in one crucial occasion (Lemma 12), the soundness of our systems (Theorems 17 and 24) holds without requesting any further condition on switchings which would exclude daimon or mix. Nevertheless, the completeness results (Sect. 3.4) hold for much smaller subsystems, using none of the debated rules (see Sect. 1.3 below). For this reason, the reader who is puzzled by daimon and mix (in particular the former, which makes the empty sequent provable in **meLL**, and with it all formulas of the form $?A$) may consider them to be mere technical tricks to obtain simpler proof nets.

1.3. Computational interpretation

The most direct computational interpretation of **meLL** can be given by considering its intuitionistic subsystem. The intuitionistic sequent calculus of **meLL** is obtained from that of Table 1 in the same way one obtains **LJ** from **LK** (Gentzen, 1934). The interest of the intuitionistic sequent calculus for **meLL** is that its derivations can be decorated with λ-terms in such a way that cut-elimination in proofs is consistent with β-reduction in the λ-calculus.

The calculus is given in Table 2, directly with the decorations. Note that, as expected, the morphological constraints imposed on intuitionistic sequents force linear implication to be treated as a primitive connective, and eliminate the par connective. For the same reason, the daimon and mix rules are excluded.

By translating $A \multimap B$ as $A^\perp \parr B$, and by converting an intuitionistic sequent $\Gamma \vdash A$ into $\vdash \Gamma^\perp, A$, one can define *intuitionistic proof nets* as nets which can be built mimicking the rules of Table 2, in the spirit of Definition 5. Intuitionistic proof nets are of course proof nets, but the decoration of Table 2 attaches a

$$\frac{}{x : A \vdash x : A}\ \text{Axiom} \qquad \frac{\Gamma \vdash t : A \quad \Delta, x : A \vdash u : B}{\Gamma \vdash u[t/x] : B}\ \text{Cut}$$

$$\frac{\Gamma, x : A \vdash u : B}{\Gamma \vdash \lambda x.u : A \multimap B}\ \text{R}\multimap \qquad \frac{\Gamma \vdash t : A \quad \Delta, y : B \vdash v : C}{\Gamma, \Delta, z : A \multimap B \vdash v[zt/y]}\ \text{L}\multimap$$

$$\frac{\Gamma, x : A[B/X] \vdash u : B}{\Gamma, x : \exists X.A \vdash u : B}\ \text{L}\forall \qquad \frac{\Gamma \vdash t : A}{\Gamma \vdash t : \forall X.A}\ \text{R}\forall\ (X \text{ not free in } \Gamma)$$

$$\frac{\Gamma, x : A \vdash u : B}{\Gamma, x : !A \vdash u : B}\ \text{D} \qquad \frac{!\Gamma \vdash t : A}{!\Gamma \vdash t : !A}\ \text{P}$$

$$\frac{\Gamma \vdash u : B}{\Gamma, x : !A \vdash u : B}\ \text{W} \qquad \frac{\Gamma, x : !A, y : !A \vdash u : B}{\Gamma, z : !A \vdash u[z/x, z/y] : B}\ \text{C}\ (z \text{ fresh})$$

$$\frac{\Gamma, x : A \vdash u : B}{\Gamma, x : \S A \vdash u : B}\ \text{L}\S \qquad \frac{\Gamma \vdash t : A}{\Gamma \vdash t : \S A}\ \text{R}\S$$

Table 2: The rules for **meLL** intuitionistic sequent calculus, and their attached λ-terms.

λ-term to them. As anticipated above, this turns into a concrete computational semantics, thanks to the following:

Proposition 3. *Let π be an intuitionistic proof net, and let $\pi \to \pi'$. Then:*

1. *π' is intuitionistic;*
2. *if t, t' are the λ-terms attached to π, π', respectively, then $t \to_\beta^* t'$.*

Proposition 3 is a useful guideline for programming with **meLL** proof nets: if one sticks to the intuitionistic subsystem, it is possible to use the λ-calculus as a target language into which proof nets can be "compiled". All complexity-related subsystems of **meLL** exploit this; as a matter of fact, the completeness with respect to the complexity classes they characterize is always proved within their intuitionistic subsystem. This will be the case for our systems too.

1.4. Elementary and light linear logic

The logical systems which are the main objects of this paper are extensions of the multiplicative fragments of elementary linear logic (**ELL**) and light linear logic (**LLL**), both introduced by Girard (1998). These two systems characterize, in a sense which will be made precise at the end of the section, the complexity classes **FE** and **FP**, respectively: the former is the class of functions computable by a Turing machine whose runtime is bounded by a tower of exponentials of fixed height (also known as *elementary functions*); the latter is the class of functions computable in polynomial time by a deterministic Turing machine. In this section, we briefly recall the definition of these two systems.

The stratification condition. The multiplicative fragment of **ELL** can be defined in our proof net syntax by using the notion of *exponential branch*, as in Danos and Joinet (2003):

Definition 8 (Exponential branch). Let σ be a (typed or untyped) **meLL** net, and let b be a flat link of σ. The *exponential branch* of b is the directed path starting from the conclusion of b, crossing a number (maybe null) of auxiliary ports and ending in the premise of a $\mathsf{why\ not}$ link (which must exist by Definition 1, or Definition 7 in the untyped case).

Definition 9 (Multiplicative elementary linear logic). Multiplicative elementary linear logic (**mELL**) is the subsystem of **meLL** composed of all proof nets satisfying the following condition:

Stratification: Each exponential branch of π crosses exactly one auxiliary port.

Note once again that the paragraph modality is absent in original definition of **mELL**, but including it is harmless (Proposition 1 still holds).

Of course the stratification condition is preserved by cut-elimination: if π is in **mELL**, and $\pi \to \pi'$, then π' is also in **mELL**. As suggested by its name, the fundamental purpose of this condition is to assure a *stratification property*, which can be formally stated as follows: whenever $\pi \to \pi'$, if l is a link of π different from a cut and l' is a residue of l in π', we have $\mathrm{d}(l') = \mathrm{d}(l)$. By contrast, in a generic **meLL** proof net a residue of a link l may also have depth smaller (by one) or greater (by any number) than l itself. In other words, depths can "communicate" in **meLL**, but are "separated worlds" in **mELL**.

Round-by-round cut-elimination. The essential property of a **mELL** proof net π is that its cuts can be eliminated so that the size of all proof nets obtained during cut-elimination is bounded by a tower of exponentials of fixed height, in the size of π itself. This is a consequence of the following facts:

F1. reducing a cut at depth i does not affect depth $j < i$;
F2. cut-elimination does not increase the depth of proof nets;
F3. reducing a cut at depth i strictly decreases the size at depth i.

F1 is true for all **meLL** proof nets; F2 and F3 are consequences of the stratification property.

Now, the idea of Girard (1998) is to eliminate cuts by operating at increasingly higher depths: if we have a **mELL** proof net of depth d, we start with a first "round" at depth 0, which will eliminate all cuts at that depth in a finite amount of time because of F3; then, we proceed with a second round at depth 1,

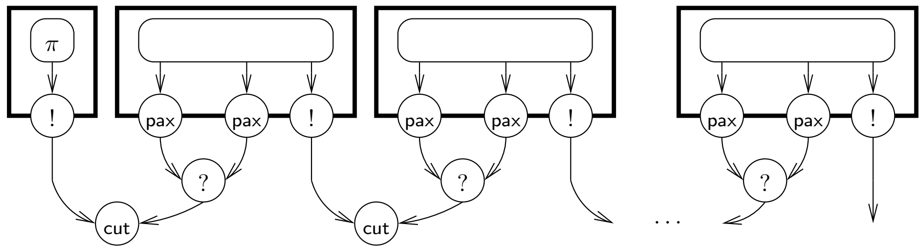

Figure 9: A chain of boxes causing an exponential blow-up in the size during cut-elimination.

which, for the same reason, will eliminate all cuts at that depth, and will not create new cuts at depth 0 because of F1; and we keep going on like this for all depths. By F2, this whole "round by round" procedure is guaranteed to terminate in at most $d+1$ rounds. After showing that the size of a proof net at the end of each round is at most $s^{s+1} < 2^{2^s}$, where s is the size of the proof net at the beginning of the round, one easily obtains an elementary bound in the size of the initial proof net, with the height of the tower of exponentials being at most twice the depth of the proof net itself. It is important to remark that the above argument makes no use of types: normalization in elementary size is possible even for untyped **mELL** proof nets.

Box chains and light linear logic. The reason for the superexponential blow-up in the size of **mELL** proof nets after each round can be understood intuitively by considering the "chain" of boxes of Fig. 9. If the number of boxes with two auxiliary ports in the chain is n, a simple calculation shows that there will be 2^n copies of π when all cuts shown are reduced. In general, the why not links involved in a chain need to be binary; but their arity can be (very roughly) bounded by the size of the proof net containing the chain, and since the length of a chain can also be subjected to a similar bound, we end up obtaining the superexponential blow-up mentioned above.

If we want to moderate the increment of the size of proof nets under cut-elimination, by naïvely looking at Fig. 9 we are led to think of a simple method: impose that boxes have at most one auxiliary port. This actually turns out to work, and is the idea underlying Girard's (1998) definition of light linear logic. Unfortunately though, this restriction is quite heavy in terms of expressive power: in fact, while normalizable in polynomial time, **mELL** proof nets using boxes with at most one auxiliary port are not able to compute all polytime functions. This is the original reason behind the introduction of the paragraph modality.

However, using the paragraph modality as we introduced it in **meLL** is not compatible with the stratification property: the paragraph too must be linked to the depth, and in order to do so we must introduce a further kind of boxes, called §-*boxes* (Fig. 10). In presence of these boxes, the usual ones are called !-boxes, and the word "box" refers to any of the two kinds.

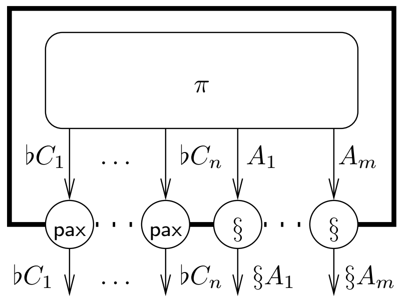

Figure 10: A §-box.

Definition 10 ($\mathbf{meLL}_{\S\mathrm{box}}$). The pre-nets and nets of $\mathbf{meLL}_{\S\mathrm{box}}$ are defined as in Definition 1, with the following modifications on the requirements concerning boxes:

a′. each of course link is the principal port of exactly one !-box;

b. each pax link is in the border of exactly one box;

c. any two distinct boxes are either disjoint or included in one another;

d. each paragraph link is in the border of exactly one §-box.

The size of a $\mathbf{meLL}_{\S\mathrm{box}}$ pre-net is defined just as in Definition 2, while the depth also takes into account §-boxes, i.e., the depth of a link is the number of nested !- and §-boxes containing it.

The proof nets of $\mathbf{meLL}_{\S\mathrm{box}}$ are defined as in Definition 4, with §-boxes being treated exactly as !-boxes.

In terms of sequent calculus, a §-box corresponds to the following rule:

$$\frac{\vdash ?\Gamma, \Delta}{\vdash ?\Gamma, \S\Delta}$$

After adapting Definition 5 to this rule, Proposition 2 extends to $\mathbf{meLL}_{\S\mathrm{box}}$.

To define cut-elimination inside $\mathbf{meLL}_{\S}$, one need only establish what the reduction of two §-boxes looks like: informally, the two §-boxes are "merged" into one, and the cut link "enters" into this new §-box. No detailed description is needed for our purposes; we refer the reader to Mazza (2006).

Multiplicative **LLL** can be defined as a subsystem of $\mathbf{meLL}_{\S}$:

Definition 11 (Multiplicative light linear logic). Multiplicative light linear logic (**mLLL**) is composed of all $\mathbf{meLL}_{\S\mathrm{box}}$ proof nets π satisfying the following conditions:

Stratification: Each exponential branch of π crosses exactly one auxiliary port.

Lightness: Each !-box of π has at most one auxiliary port.

Observe that, in the stratification condition, the auxiliary ports of §-boxes count just as the auxiliary ports of !-boxes.

In the case of **mLLL**, a round starting with a proof net of size s can be shown to lead to a proof net of size at most s^2, so that the round-by-round procedure applied to a proof net of size s and depth d terminates with a proof net of size at most s^{2^d}.

From size to time. For the moment, we have only spoken of *size* bounds to cut-elimination, whereas we started by claiming that **mELL** and **mLLL** characterize *time* complexity classes. The first step is transforming these size bounds into time bounds, which is done as follows. We consider the case of **mLLL**, the case of **mELL** being analogous. Let π be a **mLLL** proof net of size s and depth d. We know that we can eliminate all of its cuts in at most $d+1$ rounds, each operating on a proof net of size at most s^{2^d}. By F3, each round takes a linear number of steps in the size of the proof net from which the round itself starts; then, the round-by-round procedure for π terminates in at most $(d+1)s^{2^d}$ steps.

Observe now that a single cut-elimination step can at most square the size of a proof net; then, with a reasonable representation of proof nets, we are able to simulate a cut-elimination step on a Turing machine with a quadratic cost, in the size of the proof net under reduction. Assuming that all proof nets during the reduction of π have the maximum size possible, we have $(d+1)s^{2^d}$ cut-elimination steps taking each $s^{2^{d+1}}$ Turing machine steps, which means that we can compute the result of the round-by-round procedure on π in at most $(d+1)s^{2^{d+1}+2^d}$ Turing machine steps, which is polynomial in the size, and doubly-exponential in the depth. Similarly, computing the result of the round-by-round procedure for a **mELL** proof net takes a number of Turing machine steps which is elementary in the size, and hyperexponential in the depth.

Representing functions. To state precisely what it means for a logical system like **mELL** or **mLLL** to characterize a complexity class, we first need to formulate a notion of *representability* of functions from binary strings to binary strings. This is done by resorting to a formula (i.e., a type), which we may denote by $\mathbf{S}$, such that there is an infinite number of proof nets of conclusion $\mathbf{S}$, each representing a different binary string. It is very convenient at this point to operate within the intuitionistic subsystems of **mELL** and **mLLL**, and to choose $\mathbf{S}$ so that the proof nets of type $\mathbf{S}$ correspond, via the computational interpretation discussed in Sect. 1.3, to the usual λ-terms representing binary strings.

Then, we say that a function f from binary strings to binary strings is representable in **mELL** or **mLLL** just if there exists an intuitionistic proof net φ of conclusions $\mathbf{S}^\perp, \mathbf{S}$ computing f via cut-elimination, that is, $f(x) = y$ iff, whenever ξ is the proof net representing x, the proof net $\varphi(\xi)$ obtained by cutting the conclusion (of type $\mathbf{S}$) of ξ to the dual conclusion (of type $\mathbf{S}^\perp$) of φ reduces to υ, where υ is the proof net representing y. (Actually, it is necessary to allow representations of functions to be more generally of conclusions $\mathbf{S}^\perp, \mathbf{S}'$, where $\mathbf{S}'$ is the formula $\mathbf{S}$ with a number of suitable modalities prepended to it; but this is not essential at this level of detail).

Characterizing complexity classes. Characterizing a complexity class $\mathcal{C}$ now simply means that $f \in \mathcal{C}$ iff f is representable in our logical system. The forward implication is usually called the *completeness* of the system, while the backward implication is its *soundness*.

Proving the completeness of **mELL** and **mLLL** with respect to **FE** and **FP**, respectively, is a sort of (quite difficult) programming exercise, which is carried on with varying degrees of detail in Girard (1998), Roversi (1999), Danos and Joinet (2003), and Mairson and Terui (2003); we shall not discuss this here.

On the other hand, the soundness of these two systems is a consequence of the results mentioned above, plus the following crucial remark: *all proof nets of type* $\mathbf{S}$ *have constant depth* 1*, and size linear in the length of the string they represent*. Thanks to this, we see that if φ is a proof net of **mLLL** of size s and depth d representing the function f, and if ξ represents the string x, then computing the representation of $f(x)$ can be done by applying the round-by-round cut-elimination procedure to the proof net $\varphi(\xi)$, whose size is $c_1|x|+c_2+s$ (where c_1 and c_2 are suitable constants), and whose depth is $\max(d,1)$, which *does not depend on* x, but solely on φ, and thus, ultimately, on f. Hence, $f(x)$ can be computed on a Turing machine in time $\mathcal{O}(P(|x|))$, where P is a polynomial whose degree depends on f. We therefore have $f \in \mathbf{FP}$. Similarly, one can prove that if f is representable in **mELL**, then $f \in \mathbf{EF}$.

2. Linear Logic by Levels

2.1. Indexings

In **meLL** proof nets there is an asymmetry between the behavior of the two kinds of exponential links (of course and why not) with respect to the depth. More precisely, let us say that a link l is "above" an of course link o if one of the conclusions of l is the premise of o, and, similarly, let us say that l is "above" a why not link w if one of its conclusions is the premise of a flat link whose exponential branch (Definition 8) ends in w. Then, we see that if a link l is above an of course link o, we have $\mathrm{d}(l) = \mathrm{d}(o) + 1$; on the contrary, if l is above a why not link w, all we can say is that $\mathrm{d}(l) \geq \mathrm{d}(w)$.

The situation changes in **mELL**. In fact, the stratification condition guarantees that the behavior is perfectly symmetric: if a link l is above a why not link w, we have $\mathrm{d}(l) = \mathrm{d}(w) + 1$. This is true also in **mLLL**, and for paragraph links as well, because of §-boxes (remember that, in **mLLL**, the depth takes into account these boxes too).

The idea is then to take a **meLL** proof net and to try assigning to its links an index which behaves as the depth would behave in elementary and light linear logic:

Definition 12 (Indexing). Let π be a **meLL** net. An *indexing* for π is a function I from the edges of π to $\mathbb{Z}$ satisfying the constraints given in Fig. 11 and

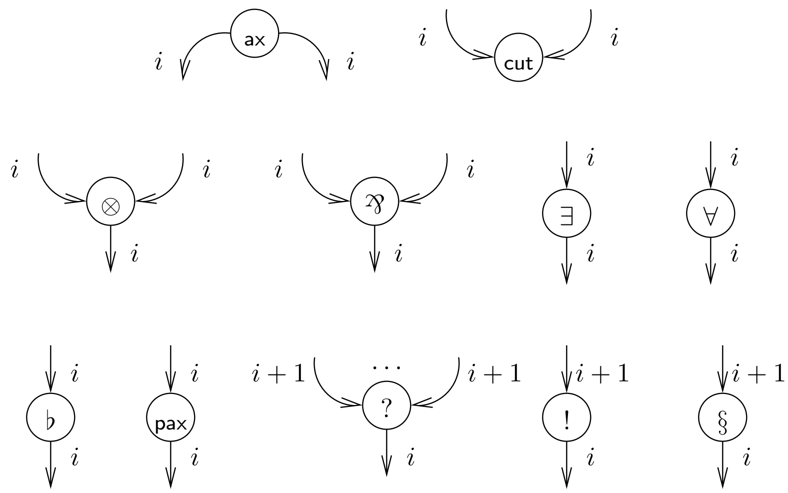

Figure 11: Constraints for indexing **meLL** proof nets. Next to each edge we represent the integer assigned by the indexing; formulas are omitted, because irrelevant to the indexing.

such that, for all conclusions e, e' of π, $I(e) = I(e')$. An assignment satisfying the constraints of Fig. 11 but not meeting the requirement on conclusions is said to be a *weak indexing*.

Note that indexings do not use formulas in any way, so the notion can be applied to untyped nets without any change.

Not all **meLL** nets admit an indexing. An example is the proof net in Fig. 12, which is the cut-free proof of the dereliction principle $!A \multimap A$ (a key principle excluded in **ELL** and **LLL**). An analogous example is given by the two proof nets corresponding to the derivations mentioned in the proof of Proposition 1, i.e., the ones asserting the isomorphism between A and $\S A$, although these do admit a weak indexing, contrarily to the proof net of Fig. 12.

Observe that (weak) indexings are transparent to connection: if π_1, π_2 are two nets admitting (weak) indexings I_1, I_2, respectively, then the net obtained by juxtaposing π_1 and π_2 admits as (weak) indexing the "disjoint union" of I_1 and I_2, which we denote by $I_1 \uplus I_2$. Likewise, if π is net whose connected components are $\pi_1, \ldots, \pi_n$, every (weak) indexing of π can be written as $\biguplus I_k$, where I_k is a (weak) indexing for π_k, for all $1 \leq k \leq n$. We use this fact to state the following:

Proposition 4 (Rigidity). *Let π be a* **meLL** *net whose connected components are $\pi_1, \ldots, \pi_n$, and let $I = \biguplus I_k$ be a (weak) indexing for π. Then, for all $p_1, \ldots, p_n \in \mathbb{Z}$, $\biguplus I_k + p_k$ is also a (weak) indexing for π. Conversely, given another (weak) indexing I' for π, there exist $p_1, \ldots, p_n \in \mathbb{Z}$ such that $I' = \biguplus I_k + p_k$.*

Proof. The first implication is trivial, so let us concentrate on the second. Let

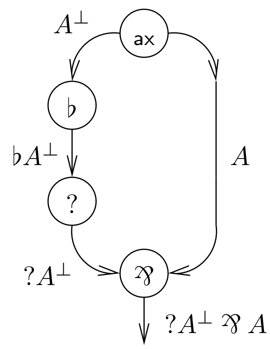

Figure 12: A **meLL** proof net admitting no (weak) indexing.

I, I' be two indexings for π, and set, for each edge e of π, $\Delta(e) = I(e) - I'(e)$. Now, observing Fig. 11, we see that differences in indexing propagate across any path in π; more precisely, whenever e_1, e_2 are both conclusions, both premises, or one conclusion and one premise of a link of π, then $\Delta(e_1) = \Delta(e_2)$. Hence, for any two edges e, e' in the same connected component of π, we have $\Delta(e) = \Delta(e')$, which is enough to prove the result. □

The following is a simple corollary of the first part of Proposition 4:

Proposition 5 (Composition). *Let π, π' be two proof nets of resp. conclusions Γ, A and $\Delta, A^\perp$, and let π'' be the proof net obtained by adding a cut link whose premises are the conclusions of π and π' labelled resp. by A and $A^\perp$. Then, if π and π' both admit an indexing, so does π''.*

As a simple case-by-case inspection shows, indexings also have the fundamental property of being preserved under cut-elimination:

Proposition 6 (Stability). *Let π be a* **meLL** *proof net such that $\pi \to \pi'$. Then, if there exists an indexing for π, there exists an indexing for π' as well. More precisely, if I is an indexing for π, there exists an indexing I' of π' such that, if e, e' are conclusions of two links l, l' of resp. π, π' such that l' is a residue of l, then $I'(e') = I(e)$. In other words, I' is "the same" indexing as I, modulo the erasures/duplications possibly induced by the cut-elimination step.*

We can therefore make the following definition:

Definition 13 (Multiplicative linear logic by levels). Multiplicative linear logic by levels ($\mathbf{mL^3}$) is the logical system defined by taking all **meLL** proof nets admitting an indexing.

The fact that an $\mathbf{mL^3}$ proof net has several (in fact, an infinity of) indexings may seem inconvenient; however, Proposition 4 settles this problem, by giving us a way to choose a *canonical indexing*:

Definition 14 (Canonical indexing). Let π be an $\mathbf{mL^3}$ proof net, and let I be an indexing for π. We say that I is *canonical* if each connected component of π has an edge e_0 such that $I(e_0) = 0$, and $I(e) \geq 0$ for all edges e of π.

Proposition 7. *Every* $\mathbf{mL^3}$ *proof net admits a unique canonical indexing.*

Proof. Let π be an $\mathbf{mL^3}$ proof net, let $\pi_1, \ldots, \pi_n$ be the connected components of π, and let k range over $\{1, \ldots, n\}$. By definition, there exists an indexing $\biguplus I_k$ for π, where I_k is an indexing for π_k. Let $m_k = \min_e I_k(e)$, where e ranges over the edges of π_k. Then, by Proposition 4, $\biguplus I_k - m_k$ is still an indexing for π, which is clearly canonical. Suppose now there exist two canonical indexes $I = \biguplus I_k$ and $I' = \biguplus I'_k$ for π. By the fact that I and I' are canonical, we know that for all k there exist e_k, e'_k in π_k such that $I(e_k) = I'(e'_k) = 0$. By Proposition 4, we also know that there exists $p_k \in \mathbb{Z}$ such that $I'_k = I_k + p_k$. Suppose $p_k > 0$; then, we would have $I(e'_k) < 0$. On the other hand, if $p_k < 0$, we would have $I'(e_k) < 0$. In both cases, we would be in contradiction with the fact that I and I' are canonical, hence we must have $p_k = 0$, and $I = I'$. □

Definition 15 (Level). Let π be an $\mathbf{mL^3}$ proof net, and let I_0 be its canonical indexing. The *level* of π, denoted by $\ell(\pi)$, is the maximum integer assigned by I_0 to the edges of π. If l is a link of π of conclusion e (or of conclusions e_1, e_2 in the case of an axiom link), and if $\mathcal{B}$ is a box of π whose principal port has conclusion e', we say that the level of l, denoted by $\ell(l)$, is $I_0(e)$ (or $I_0(e_1) = I_0(e_2)$ in the case of an axiom), and that the level of $\mathcal{B}$, denoted by $\ell(\mathcal{B})$, is $I_0(e')$.

From now on, when we speak of an $\mathbf{mL^3}$ proof net π, we shall always refer to its canonical indexing. The reader may wonder why we did not use $\mathbb{N}$ instead of $\mathbb{Z}$ as the range of our indexes in the first place; we simply believe $\mathbb{Z}$ to be a more natural choice, as the set of indexes need not be well-founded. Moreover, using $\mathbb{N}$ would be awkward in the sequent calculus formulation of $\mathbf{mL^3}$ (cf. Table 3 below): it would force to impose a restriction on exponential rules, an unnecessary complication.

Recall that levels are conceived to behave like depths in **mELL**; then, it is not surprising that **mELL** is exactly the (proper) subsystem of $\mathbf{mL^3}$ in which levels and depths coincide:

Proposition 8. *Let* π *be a* **meLL** *proof net. Then,* π *is in* **mELL** *iff* π *is in* $\mathbf{mL^3}$ *and, for every link* l *of* π *whose conclusion is not a discharged formula, we have* $\ell(l) = \mathrm{d}(l)$.

Now to help relating proof nets to the intuitions coming from the λ-calculus, we give an example of a λ-term and a corresponding proof net of $\mathbf{mL^3}$. The following term is the Church representation of the binary list 101, and its syntactic tree is given in Fig. 13:

$$t_{101} = \lambda s_0.\lambda s_1.\lambda z.(s_1\ (s_0\ (s_1\ z))).$$

An $\mathbf{mL^3}$ proof net corresponding to this term, according to Proposition 3, is given in Fig. 14. Note that nodes λ (resp. @) of the syntactic tree correspond to nodes $\parr$ (resp. $\otimes$) of the proof net.

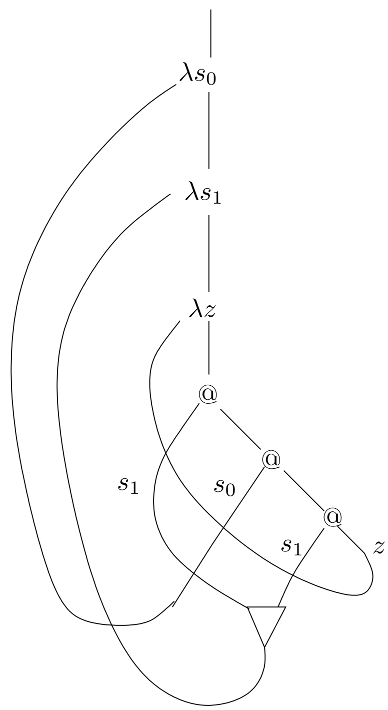

Figure 13: Syntactic tree for the λ-term t_{101}.

2.2. Light linear logic by levels

Chains of boxes like that of Fig. 9 may be built in $\mathbf{mL^3}$, so there is no hope of finding sub-exponential bounds for the size of $\mathbf{mL^3}$ proof nets under cut-elimination. We then follow the same idea as light linear logic:

Definition 16 (Multiplicative light linear logic by levels). Multiplicative light linear logic by levels ($\mathbf{mL^4}$) is the logical system composed of all $\mathbf{mL^3}$ proof nets π satisfying the following conditions:

(Weak) Stratification: Each exponential branch (Definition 8) of π crosses at most one auxiliary port.

Lightness: Each box of π has at most one auxiliary port.

It is not hard to see that $\mathbf{mL^4}$ is stable under cut-elimination, i.e., that a suitable version of Proposition 6 holds. Indeed, the stratification condition is needed precisely for that purpose: in its absence, one can find an $\mathbf{mL^3}$ proof net satisfying the lightness condition which reduces to a proof net no longer satisfying it.

As expected, $\mathbf{mL^4}$ is related to **mLLL**. To see how, we consider the forgetful embedding of **mLLL** into **meLL** which simply removes paragraph boxes,

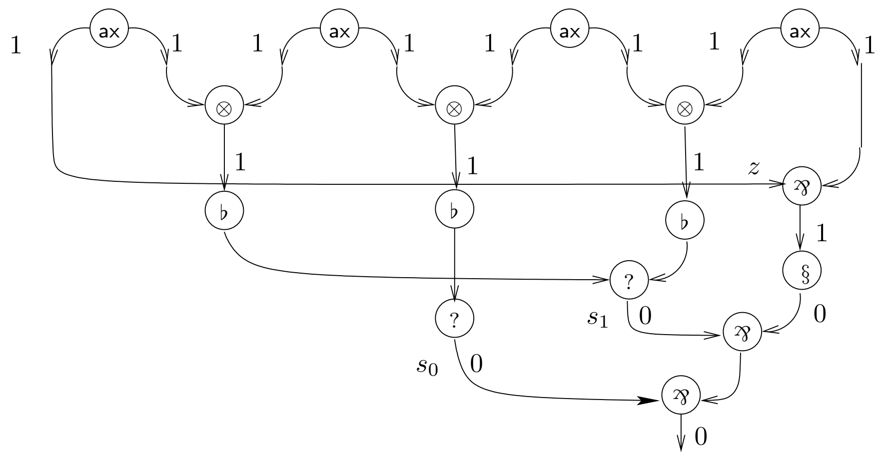

Figure 14: An $\mathbf{mL^3}$ proof-net corresponding to t_{101}.

retaining only the corresponding paragraph links (recall that our definition of **meLL** includes the paragraph modality). Observe that this embedding is compatible with cut-elimination: if $\pi_1 \to \pi_2$, then $\pi_1^+ \to \pi_2^+$ (see Mazza (2006) for the details on cut-elimination with §-boxes). We can then see **mLLL** as a subsystem of $\mathbf{mL^4}$, in the following sense:

Proposition 9. *Let π be a* **mLLL** *proof net, and let π^+ be its forgetful image in* **meLL***. Then, π^+ is in* $\mathbf{mL^4}$ *and, for every link l^+ of π^+ whose conclusion is not a discharged formula and which corresponds to a link l of π, we have $\ell(l^+) = \mathrm{d}(l)$ (we remind that in* **mLLL** *proof nets the depth also takes into account paragraph boxes, see Definition 10).*

As already observed above, $\S A$ is not isomorphic to A in $\mathbf{mL^3}$ (or $\mathbf{mL^4}$). However, it is not hard to check that in both systems the paragraph modality commutes with all connectives: for all A, B, $\S(A \otimes B)$, $\S!A$, and $\S\forall X.A$ are all provably isomorphic (in the same sense as that of Proposition 1) to $\S A \otimes \S B$, $!\S A$, and $\forall X.\S A$, respectively (and, by duality, similar isomorphisms hold for the connectives $⅋$, ?, and $\exists$).

None of the above isomorphisms holds in **LLL**, and this is why it does not make much sense to establish a converse of Proposition 9. We therefore obtained a system in which the paragraph modality, like **LLL**, is not trivial, but, unlike **LLL**, enjoys more flexible principles. In Sect. 3 we shall see that $\mathbf{mL^3}$ and $\mathbf{mL^4}$ have also interesting properties with respect to the complexity of their cut-elimination procedure.

2.3. Linear logic by levels as a sequent calculus

It is possible to formulate $\mathbf{mL^3}$ and $\mathbf{mL^4}$ as sequent calculi, which may be useful for having a clearer correspondence with λ-terms, as in Sect. 1.3. In

$$\frac{}{\vdash A^{\perp i}, A^i}\ \text{Axiom} \qquad \frac{\vdash \Gamma, A^i \quad \vdash \Delta, A^{\perp i}}{\vdash \Gamma, \Delta}\ \text{Cut}$$

$$\frac{\vdash \Gamma, A^i \quad \vdash \Delta, B^i}{\vdash \Gamma, \Delta, A \otimes B^i}\ \text{Tensor} \qquad \frac{\vdash \Gamma, A^i, B^i}{\vdash \Gamma, A \parr B^i}\ \text{Par}$$

$$\frac{\vdash \Gamma, A^i}{\vdash \Gamma, \forall X.A^i}\ \text{For all } (X \text{ not free in } \Gamma) \qquad \frac{\vdash \Gamma, A[B/X]^i}{\vdash \Gamma, \exists X.A^i}\ \text{Exists}$$

$$\frac{\vdash ?\Gamma, A^{i+1}}{\vdash ?\Gamma, !A^i}\ \text{Promotion} \qquad \frac{\vdash \Gamma, A^{i+1}}{\vdash \Gamma, ?A^i}\ \text{Dereliction}$$

$$\frac{\vdash \Gamma}{\vdash \Gamma, ?A^i}\ \text{Weakening} \qquad \frac{\vdash \Gamma, ?A^i, ?A^i}{\vdash \Gamma, ?A^i}\ \text{Contraction}$$

$$\frac{\vdash \Gamma, A^{i+1}}{\vdash \Gamma, \S A^i}\ \text{Paragraph}$$

Table 3: The rules for **mL³** 2-sequent calculus. Daimon and mix are omitted.

doing this, one immediately realizes that *2-sequents*, rather than sequents, are the natural syntax for this purpose. Calculi for 2-sequents have been extensively studied by Masini (1992) and have been found to be quite useful for the proof-theory of modal logics. In particular, linear logic and its elementary and light variants can all be formulated as 2-sequent calculi (Guerrini et al., 1998).

A **meLL** *2-sequent* M is a function from $\mathbb{Z}$ to **meLL** sequents such that $M(i)$ is the empty sequent for all but finitely many i. 2-sequents can be succinctly represented as standard sequents by decorating formulas with an integer index: $\vdash A_1^{i_1}, \ldots, A_n^{i_n}$ represents the 2-sequent M such that $M(i) = \vdash \Gamma$, where Γ contains all and only the occurrences of formulas $A_j^{i_j}$ such that $i_j = i$.

The 2-sequent calculus for **mL³** is given in Table 3, where Γ, Δ stand for multisets of **meLL** formulas decorated with an integer. The daimon and mix rules are omitted, because identical to those in Table 1.

We say that a derivation of $\vdash \Gamma$ in the calculus of Table 3 is *proper* if all the formulas in Γ have the same index, i.e., the derived 2-sequent is indeed a sequent; moreover, we say that a *weak* **mL³** net is a net admitting a weak indexing. By Proposition 2, it is more or less evident that a sequentializable weak **mL³** net is a weak **mL³** proof net. Hence, we see that **mL³** proof nets exactly correspond to the proper derivations of the calculus of Table 3.

We remark that the calculus of Table 3 is very similar to Guerrini, Martini, and Masini's **2ELL** (Guerrini et al., 1998), without additive connectives. In their work, the authors prove that **2ELL** is an equivalent formulation of **ELL** with respect to provability, and claim that such formulation satisfies cut-elimination (Girard's original definition of **ELL** did not (Kanovich et al., 2003)). Our calculus eliminates all of the constraints of the multiplicative fragment of **2ELL**, and in fact $\mathbf{mL^3}$ is a proper extension of **mELL**—preserving, however, its complexity properties, as we shall see below.

The system $\mathbf{mL^4}$ is obtained in terms of sequent calculus by replacing the promotion rule with the following one:

$$\frac{\vdash B^{j+1}, A^{i+1}}{\vdash ?B^{j}, !A^{i}}\ \text{Light promotion}$$

where the formula B may not be present.

3. Complexity Bounds

To establish the complexity bounds for $\mathbf{mL^3}$ and $\mathbf{mL^4}$, we shall try to adapt the arguments originally given by Girard (1998) for **ELL** and **LLL**. Let us then go back to Sect. 1.4 and consider again the three facts about cut-elimination in **mELL** which are at the base of its elementary size bound:

F1. reducing a cut at depth i does not affect depth $j < i$;
F2. cut-elimination does not increase the depth of proof nets;
F3. reducing a cut at depth i strictly decreases the size at depth i.

We know that F1 is true in general in **meLL**, and hence in $\mathbf{mL^3}$ too; it is not hard to see that F2 and F3 instead fail altogether in $\mathbf{mL^3}$ and $\mathbf{mL^4}$. Nevertheless, in the light of Propositions 8 and 9, we may expect those facts to hold in our systems provided we replace the word "depth" with "level". Indeed, this works for F1 and F2:

Lemma 10. *Let π be an $\mathbf{mL^3}$ proof net such that $\pi \to \pi'$ by reducing a* cut *link at level i. Then, for all $j < i$, every link of π of level j has exactly one residue of level j in π'; moreover, if c' is a* cut *link of π' at level j, then c' is a residue of a* cut *link c of π.*

Lemma 11. *Let π be an $\mathbf{mL^3}$ proof net such that $\pi \to \pi'$. Then, $\ell(\pi') \leq \ell(\pi)$.*

In other words, reducing a cut at level i does not affect lower levels (in particular, no cut link at level $j < i$ is created), and the overall level does not increase.

On the contrary, even the "corrected" version of F3 fails for $\mathbf{mL^3}$ and $\mathbf{mL^4}$, because a box of level i may contain links of *any* level, in particular i itself. Fig. 15 gives an example of this: reducing a cut at level i ($i = 0$ in this case) may duplicate cuts at the same level. Therefore, a straightforward adaptation of Girard's "round-by-round" procedure, which trades depths for levels, will not work. There is a workaround though: in fact, there are cuts for which the corrected version of F3 holds; our solution will consist in showing that these can be reduced first.

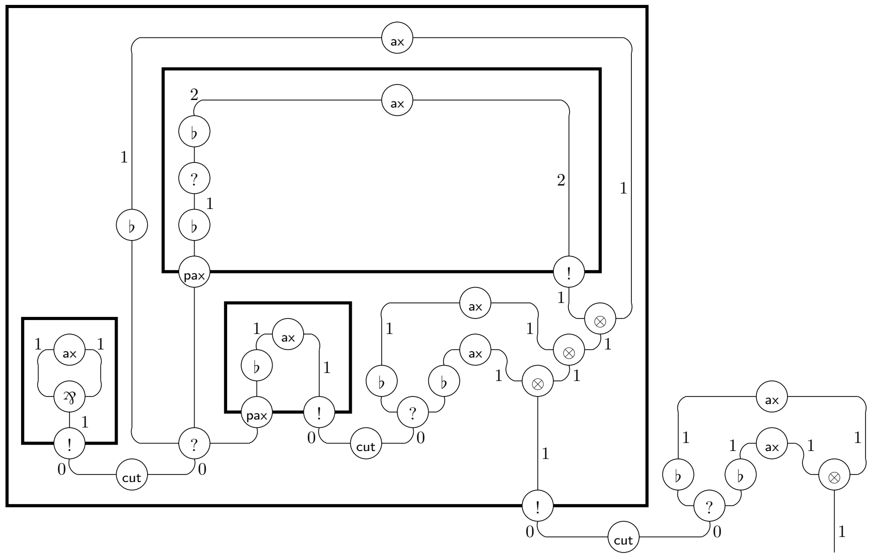

Figure 15: An example of nested boxes of identical level (much smaller examples exist; we gave this one because we shall re-use it later on for different purposes).

3.1. Termination

First of all, we prove that reduction of $\mathbf{mL^3}$ proof nets always terminates, even in the untyped version of the system. From this moment on, by "**meLL** proof net" we shall mean "untyped **meLL** proof net", and by "$\mathbf{mL^3}$ (resp. $\mathbf{mL^4}$) proof net" we shall mean "untyped **meLL** proof net admitting an indexing (resp. admitting an indexing and satisfying the structural conditions of Definition 16)".

Definition 17 (Isolevel tree). Let π be a **meLL** proof net, and let e be an edge of π which is the conclusion of a link l different from flat or pax. The *isolevel tree* of e is defined by induction as follows:

- if l is an axiom, why not, of course, or paragraph link, then the isolevel tree of e consists of the link l alone;
- otherwise, let $e_1, \ldots, e_k$ (with $k \in \{1, 2\}$) be the premises of l; then, the isolevel tree of e is the tree whose root is l and whose immediate subtrees are the isolevel trees of $e_1, \ldots, e_k$.

Definition 18 (Complexity of reducible cuts). Let π be a **meLL** proof net, and let c be a reducible cut link of π, whose premises are e_1, e_2. The *complexity* of c, denoted by $\sharp c$, is the sum of the number of nodes contained in the isolevel trees of e_1 and e_2. (Note that the isolevel trees of e_1, e_2 are always defined because the premises of a cut can never be conclusions of flat or pax links).

Definition 19 (Weight of an $\mathbf{mL^3}$ proof net). Let π be an $\mathbf{mL^3}$ proof net of level l. If $k \in \mathbb{Z}$, we denote by $\mathsf{cuts}_k(\pi)$ the set of reducible cut links of π at level k. The *weight* of π, denoted by α_π, is the function from $\mathbb{N}$ to $\mathbb{N}$ defined as follows:

$$\alpha_\pi(i) = \sum_{c \in \mathsf{cuts}_{l-i}(\pi)} \sharp c.$$

Note that, if π has level l, then for all $i > l$, we have $\alpha_\pi(i) = 0$. Weights are therefore almost everywhere null, and the set of all weights can be well-ordered so as to be isomorphic to ω^ω.

We recall that, concretely, this order is a variant of the lexicographical order, and is defined as follows. Let α, β be two almost-everywhere-null functions from $\mathbb{N}$ to $\mathbb{N}$. We put $C_{\alpha,\beta} = \{i \in \mathbb{N} \ ; \ \alpha(i) \neq \beta(i)\}$. Observe that $C_{\alpha,\beta}$ is finite, because α and β are almost everywhere null. Moreover, $C_{\alpha,\beta}$ is non-empty iff $\alpha \neq \beta$; in this case, let $m = \max C_{\alpha,\beta}$, and we set $\alpha < \beta$ iff $\alpha(m) < \beta(m)$.

So for all π, α_π can be seen as an ordinal strictly smaller than ω^ω. Our cut-elimination proof will simply show that, whenever an $\mathbf{mL^3}$ proof net π is not normal, there always exists π' such that $\pi \to \pi'$ and $\alpha_{\pi'} < \alpha_\pi$.

Below, we say that a flat link b is *above* a $\mathsf{why\ not}$ link w iff the exponential branch of b ends in w.

Definition 20 (Contractive order). Let π be an $\mathbf{mL^3}$ proof net, and let $\mathcal{B}, \mathcal{C}$ be two boxes of π. We write $\mathcal{B} \prec_1 \mathcal{C}$ iff $\mathcal{B}$ and $\mathcal{C}$ are at the same level, $\mathcal{B}$ is cut with a $\mathsf{why\ not}$ link w, and $\mathcal{C}$ contains a flat link above w. We denote by $\preceq$ the reflexive-transitive closure of $\prec_1$.

Lemma 12. *The relation $\preceq$ is a partial order.*

Proof. Suppose there is a cycle in $\prec_1$, i.e., there exist $n \geq 1$ different boxes $\mathcal{B}_1, \ldots, \mathcal{B}_n$ such that $\mathcal{B}_1 \prec_1 \cdots \prec_1 \mathcal{B}_n \prec_1 \mathcal{B}_1$. We say that such a cycle has a *lump* iff there exist $i \neq j$ such that $\mathcal{B}_i \prec_1 \mathcal{B}_j$ and $\mathcal{B}_i$ is contained in $\mathcal{B}_j$. Let k be the number of lumps in the cycle; we shall prove a contradiction by induction on k. If $k = 0$, then all boxes are disjoint. In this case, it is easy to build, by induction on n, a cyclic switching of π (or of the contents of the minimal box containing the whole chain), which is impossible, since π is supposed to be a proof net. If $k > 0$, let $\mathcal{B}_i, \mathcal{B}_j$ be a pair of boxes inducing a lump. Since we have a cycle, there certainly exists p such that $\mathcal{B}_p \prec_1 \mathcal{B}_i$. If $p = j$, then there is obviously a cyclic switching around $\mathcal{B}_j$, yielding again a contradiction. Otherwise, by definition, $\mathcal{B}_p \prec_1 \mathcal{B}_i$ means that there is a flat link inside $\mathcal{B}_i$ which is above the $\mathsf{why\ not}$ link to which $\mathcal{B}_p$ is cut. But $\mathcal{B}_i$ is contained in $\mathcal{B}_j$, so this flat link is also in $\mathcal{B}_j$, which means that $\mathcal{B}_p \prec_1 \mathcal{B}_j$ as well. Independently of whether $\mathcal{B}_p$ is included in $\mathcal{B}_j$ or not, the cycle obtained by removing $\mathcal{B}_i$ from the original one necessarily has $k - 1$ lumps, and the induction hypothesis applies. Therefore, $\prec_1$ is acyclic, and its reflexive-transitive closure is a partial order. □

In the following, we deem a cut link *contractive* iff its premises are the conclusions of an $\mathsf{of\ course}$ link and a $\mathsf{why\ not}$ link of arity strictly greater than zero. All other reducible cut links are called *non-contractive.*

Definition 21 (Cut order). Let π be an $\mathbf{mL^3}$ proof net, and let $\mathsf{cuts}(\pi)$ be the set of reducible cut links of π. We turn $\mathsf{cuts}(\pi)$ into a partially ordered set by posing, for $c, c' \in \mathsf{cuts}(\pi)$, $c \leq c'$ iff one of the following holds:

- $\ell(c) < \ell(c')$;
- c is non-contractive and c' is contractive;
- c and c' are both contractive, involving resp. the boxes $\mathcal{B}$ and $\mathcal{B}'$, and $\mathcal{B} \preceq \mathcal{B}'$.

That the above relation is indeed a partial order follows easily from the definition and Lemma 12.

The weak normalization of untyped $\mathbf{mL^3}$ is a trivial corollary of the following result, as anticipated above:

Lemma 13. *Let π be an $\mathbf{mL^3}$ proof net which is not normal. Then, there exists π' such that $\pi \to \pi'$ and $\alpha_{\pi'} < \alpha_\pi$.*

Proof. By hypothesis, $\mathsf{cuts}(\pi) \neq \emptyset$; of course $\mathsf{cuts}(\pi)$ is also finite, so there is at least one minimal element w.r.t. the cut order. Take any one of them (call it c), and reduce it, obtaining π'. Let M (resp. M') be the maximum k such that $\alpha_\pi(k) > 0$ (resp. $\alpha_{\pi'}(k) > 0$). First of all, by Lemmas 11 and 10 we have resp. that $\ell(\pi') \leq \ell(\pi)$ and that $M' \leq M$. If any of the two inequalities is strict, we immediately have $\alpha_{\pi'} < \alpha_\pi$. Therefore, we may assume $\ell(\pi') = \ell(\pi) = l$ and $M' = M$. By the minimality hypothesis, we see that the level of c must be $i = l - M$, and that π contains no reducible cut at level $j < i$. At this point, whatever happens in reducing c, by Lemma 10 we know that $\alpha_{\pi'}(n) = \alpha_\pi(n) = 0$ for all $n > M$, so it is enough to check that something decreases at level i, i.e., that $\alpha_{\pi'}(M) < \alpha_\pi(M)$. The proof now splits into five cases, depending on the nature of c. If c is not an exponential cut, or if it is a weakening cut, we leave it to the reader to verify that the condition holds.

So let c be contractive, and let $\mathcal{B}$ be the box involved. We claim that the content of $\mathcal{B}$ contains no reducible cut links at level i. As a matter of fact, suppose for the sake of contradiction that $\mathcal{B}$ contains a reducible cut c' of level i (which is necessarily different from c). Because of the second clause of Definition 21, c' must be contractive, otherwise we would contradict the minimality of c. But in this case, let $\mathcal{B}'$ and w be resp. the box and the $\mathsf{why\ not}$ link involved in c'. Since c' is contractive, there is at least one flat link above w, which entails $\mathcal{B}' \preceq \mathcal{B}$; by the third clause of Definition 21, we would thus obtain a second, definitive contradiction.

Now that we know that $\mathcal{B}$ is normal at level i, it is not hard to verify that the thesis holds: π' contains at least one copy of the content of $\mathcal{B}$, but none of these copies contributes to the value of $\alpha_{\pi'}(M)$. Moreover, the new cuts contained in π' are all at level $i + 1$, whereas one reducible cut at level i (c itself) has disappeared. Therefore, $\alpha_{\pi'}(M) < \alpha_\pi(M)$, as desired. □

Proposition 14 (Untyped weak normalization). *Untyped $\mathbf{mL^3}$ proof nets are weakly normalizable.*

Proof. By transfinite induction up to ω^ω. Let $\beta < \omega^\omega$, and suppose that for all $\alpha < \beta$, $\alpha_\pi = \alpha$ implies that π is weakly normalizable. Take a proof net π such that $\alpha_\pi = \beta$; π is either normal, hence weakly normalizable, or, by Lemma 13 and by the above induction hypothesis, it reduces to a weakly normalizable proof net. But any proof net reducing to a weakly normalizable proof net is also weakly normalizable. □

3.2. Elementary bound for $\mathbf{mL^3}$

From now on, we shall only consider the cut-elimination procedure given by the proof of Lemma 13, i.e., the one reducing only minimal cuts in the cut order. More concretely, given an $\mathbf{mL^3}$ proof net π, this procedure chooses a cut to be reduced in the following way:

1. find the lowest level at which reducible cuts are present in π, say i;
2. if non-contractive cuts are present at level i, choose any of them and reduce it;
3. if only contractive cuts are left, chose one involving a minimal box in the contractive order.

This is nothing but Girard's "round by round" procedure, modulo two modifications: we use levels instead of depths, and we are more restrictive on which contractive cuts can be reduced (in Girard's procedure for **mLLL**, *any* contractive cut may be reduced once all non-contractive cuts at the same depth are reduced). This last point is strictly technical: it is required because of configurations such as the one shown in Fig. 15, as discussed above. What is really fundamental is the shift from depth to level, which is indeed the key novelty of our work.

Let us start with a few useful definitions:

Definition 22. Let π be an $\mathbf{mL^3}$ proof net.

1. The *size of level* i of π, denoted by $|\pi|_i$, is the number of links at level i of π different from auxiliary ports.
2. π is *i-normal* iff it contains no reducible cut link at all levels $j \leq i$.
3. π is *i-contractive* iff it is $(i-1)$-normal and contains only contractive cut links at level i.

Lemma 15. *Let π be an $(i-1)$-normal proof net. Then, the round-by-round procedure reaches an i-normal proof net in at most $|\pi|_i$ steps.*

Proof. Let $\pi = \pi_0 \to \pi_1 \to \cdots \to \pi_n$ be reduction sequence generated by our procedure, with π_n i-normal. By what we have seen in the proof of Lemma 13, if we put $M = \ell(\pi) - i$, we have that $\alpha_{\pi_{j+1}}(M) < \alpha_{\pi_j}(M)$ for all $0 \leq j \leq n-1$. Therefore, $n \leq \alpha_\pi(M)$. But by definition $\alpha_\pi(M) \leq |\pi|_i$, hence the thesis. □

Below, we use the notation 2^n_k with the following meaning: for all n, $2^n_0 = n$, and $2^n_{k+1} = 2^{2^n_k}$.

Lemma 16. *Let π be an i-contractive proof net, such that $\pi \to^* \pi'$ under the round-by-round procedure, with π' i-normal. Then, $|\pi'| \leq 2_2^{|\pi|}$.*

Proof. In the proof, we shall say that the arity of a contractive cut link c is the arity of the why not link whose conclusion is premise of c. Let π_0 be an i-contractive proof net, such that $\pi_0 \to \pi_1$ by reducing a minimal cut c at level i. We have that, for all $k \neq i$, $|\pi_0|_k = B_k + C_k$, while $|\pi_0|_i = B_i + C_i + 3$, where B_k is the size of level k of the content of the box $\mathcal{B}$ whose principal port's conclusion is premise of c, and C_k is a suitable non-negative integer. It is enough to inspect Fig. 7 to see that, if the arity of c is A, we have $|\pi_1|_k = AB_k + C_k$, for all k. Now, since the step is contractive, $A \geq 1$, so that $|\pi_1|_k \leq A(B+C) = A|\pi_0|_k$.

We now make the following claims:

1. π_1 is i-contractive;
2. if c_1 is cut link of π_1 at level i, and c_0 is its lift in π_0, then the arities of c_0 and c_1 coincide.

The first fact can be checked by simply looking at Fig. 7. For what concerns the second, let $w_0, \mathcal{B}_0$ and $w_1, \mathcal{B}_1$ be resp. the why not link and box cut by resp. c_0 and c_1. Note that, by hypothesis, w_0 and $\mathcal{B}_0$ are the lifts of resp. w_1 and $\mathcal{B}_1$. Now suppose, for the sake of contradiction, that the arity of w_1 is different than that of w_0. Another simple inspection of Fig. 7 shows that this may be the case only if an exponential branch of π_0 ending in w_0 crosses the border of $\mathcal{B}$ (the box involved in the reduction leading from π_0 to π_1). But if it is so, then there is a flat link above w_0 which is inside $\mathcal{B}$, which implies that $\mathcal{B}_0 \preceq \mathcal{B}$. By Definition 21, we have $c_0 < c$, contradicting the minimality of c. Therefore, the maximum arity of all cuts of π_1 at level i cannot exceed the maximum arity of all cuts of π_0 at level i.

Let now $\pi = \pi_0 \to \cdots \to \pi_n = \pi'$ be the reduction sequence generated by the round-by-round procedure. If $A_1, \ldots, A_n$ are the arities of the cut links reduced at each step, we have, for all k,

$$|\pi'|_k \leq |\pi|_k \prod_{j=1}^{n} A_j.$$

But, by the above claim, each A_j cannot be greater than the greatest arity of why not links present in π. This is of course bounded by $|\pi|_{i+1}$ (a contraction of arity A at level i needs the presence of A flat links at level $i+1$), so we can conclude that

$$|\pi'|_k \leq |\pi|_k |\pi|_{i+1}^n \leq |\pi|_k |\pi|_{i+1}^{|\pi|_i},$$

where we have used Lemma 15, which tells us that $n \leq |\pi|_i$. Now, if put $l = \ell(\pi') = \ell(\pi)$, we have

$$|\pi'| = \sum_{k=0}^{l} |\pi'|_k \leq \sum_{k=0}^{l} |\pi|_k |\pi|_{i+1}^{|\pi|_i} = |\pi||\pi|_{i+1}^{|\pi|_i} \leq |\pi|^{|\pi|+1} \leq 2^{2^{|\pi|}},$$

as stated in our thesis. □

Theorem 17 (Elementary bound for $\mathbf{mL^3}$). *Let π be an $\mathbf{mL^3}$ proof net of size s and level l. Then, the round-by-round procedure reaches a normal form in at most $(l+1)2^s_{2l}$ steps.*

Proof. We can decompose the reduction from π to its normal form π_l as follows: $\pi = \pi_{-1} \to^* \pi_0 \cdots \to^* \pi_l$, where each π_i is i-normal. By Lemma 15, if we call the length of the whole reduction sequence L, we have

$$L \leq \sum_{i=0}^{l} |\pi_{i-1}|_i \leq \sum_{i=0}^{l} |\pi_{i-1}|.$$

The reductions leading from π_i to π_{i+1} can be further decomposed as $\pi_i \to^* \pi'_i \to^* \pi_{i+1}$, where π'_i is the first i-contractive proof net obtained in the reduction sequence. Observe now that the size of proof nets does not grow under non-contractive steps; therefore, for all i, $|\pi'_i| \leq |\pi_i|$. From this, if we apply Lemma 16, we have that, for all i, $|\pi_{i+1}| \leq 2_2^{|\pi_i|}$.

It can now be proved by a straightforward induction that, for all $i \geq 0$, we have $|\pi_{i-1}| \leq 2^s_{2i}$. Hence, we obtain

$$L \leq \sum_{i=0}^{l} |\pi_{i-1}| \leq \sum_{i=0}^{l} 2^s_{2i} \leq (l+1)2^s_{2l},$$

as desired. □

Note that, in case we have a $\mathbf{mELL}$ proof net π of size s and depth d, by Proposition 8 depth and level coincide, so the above results tells us that π can be reduced in at most $(d+1)2^s_{2d}$ steps, which is the bound found by Danos and Joinet (2003). However, in $\mathbf{mL^3}$ it is in general the level that controls the complexity, not the depth. Fig. 16 gives a clear example of this. It uses the fact that, following again Danos and Joinet (2003), in $\mathbf{mELL}$ the exponential function $\exp(n) = 2^n$ can be programmed as a proof net of conclusions $\mathbf{N}^\perp, !\mathbf{N}$, where $\mathbf{N}$ is a suitable type of natural numbers, the cut-free proof nets of conclusion $\mathbf{N}$ corresponding to Church integers, in analogy with the example given in Fig. 14. Then, the cut-free form of the proof net θ_n of Fig. 16 is the proof net representing the number 2_n, i.e., a tower of powers of 2 of height n. Hence, the size of θ_n is linear in n, but the size of its cut-free form is hyperexponential in n. This is in accordance with Theorem 17, because the level of θ_n turns out to be n. And yet, the depth of each θ_n is constant, indeed merely equal to 1.

3.3. Polynomial bound for $\mathbf{mL^4}$

In the case of $\mathbf{mL^4}$, a finer analysis leads to a substantial improvement of Theorem 17. In the following, if a box $\mathcal{C}$ contains a box $\mathcal{B}$, we shall write $\mathcal{B} \subseteq \mathcal{C}$. The relation $\subseteq$ is obviously a finite, downward-arborescent partial order.

Definition 23 (Light contractive order). Let π be an $\mathbf{mL^3}$ proof net, and let $\mathcal{B}, \mathcal{C}$ be boxes of π. We put $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}$ iff $\mathcal{B} \prec_1 \mathcal{C}$ and $\mathcal{B} \not\subseteq \mathcal{C}$. We denote by $\preceq^{\mathbf{L}}$ the reflexive transitive closure of $\prec_1^{\mathbf{L}}$, or, equivalently, we put $\mathcal{B} \preceq^{\mathbf{L}} \mathcal{C}$ iff $\mathcal{B} = \mathcal{C}$, or $\mathcal{B} \preceq \mathcal{C}$ and $\mathcal{B} \not\subseteq \mathcal{C}$.

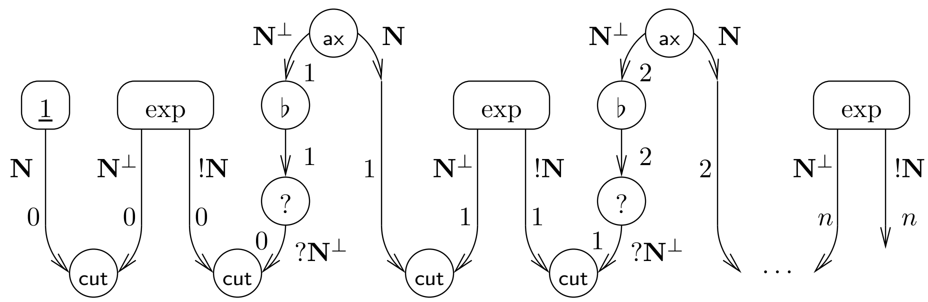

Figure 16: The proof net θ_n, an iteration of n proof nets computing the exponential function.

Lemma 18. *In* $\mathbf{mL^4}$, *the relation* $\preceq^{\mathbf{L}}$ *is an upward-arborescent partial order.*

Proof. The fact that it is a partial order follows trivially from its definition and from Lemma 12, and indeed this is true for $\mathbf{mL^3}$ as well. For what concerns its arborescence, simply observe that, by the lightness condition of Definition 16, for each box $\mathcal{C}$ of an $\mathbf{mL^4}$ proof net there may be at most one $\mathcal{B}$ such that $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}$. □

Observe that, if $\mathcal{B}, \mathcal{C}$ are two boxes of an $\mathbf{mL^4}$ proof net, thanks to the stratification condition $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}$ implies $\mathrm{d}(\mathcal{B}) = \mathrm{d}(\mathcal{C})$. In fact, in $\mathbf{mL^4}$ the light contractive order is simply a "depth-wise slicing" of the contractive order.

For example, if we take the proof net of Fig. 17, we see that the contractive order at level 0 is linear, i.e., $\mathcal{B} \preceq \mathcal{C} \preceq \mathcal{B}_0$, while in the light contractive order we only have $\mathcal{B} \preceq^{\mathbf{L}} \mathcal{C}$, and $\mathcal{B}_0$ is incomparable with both $\mathcal{B}$ and $\mathcal{C}$, because it is not at the same depth.

Definition 24 (Arity of a box). Let π be an $\mathbf{mL^3}$ proof net, and let $\mathcal{B}$ be a box of π. The *arity* of $\mathcal{B}$, denoted by $\nabla(\mathcal{B})$, is defined as follows:

- if the principal port of $\mathcal{B}$ is premise of a cut link whose other premise is the conclusion of a why not link w, then $\nabla(\mathcal{B})$ is equal to the arity of w minus the number of flat links above w which are inside a box $\mathcal{C}$ such that $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}$;
- otherwise, $\nabla(\mathcal{B}) = 1$.

Concretely, the arity of a box at level i and depth d is the number of copies that will be made of its content and that will not be subjected to further duplications by reducing cuts at level i and depth d.

In the example of Fig. 17, the why not link w to which $\mathcal{B}$ is cut has arity 3, but one of the flat links above it is inside a box $\mathcal{C}$ such that $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}$, hence $\nabla(\mathcal{B}) = 2$ (note that we do not have $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{D}$ because $\mathcal{D}$ is not at the same level as $\mathcal{B}$). On the other hand, the arities of the other two boxes at level 0 are equal to the arities of their corresponding why not links: $\nabla(\mathcal{C}) = 2$ and $\nabla(\mathcal{B}_0) = 2$. Instead, since $\mathcal{D}$ is not involved in a cut, $\nabla(\mathcal{D}) = 1$.

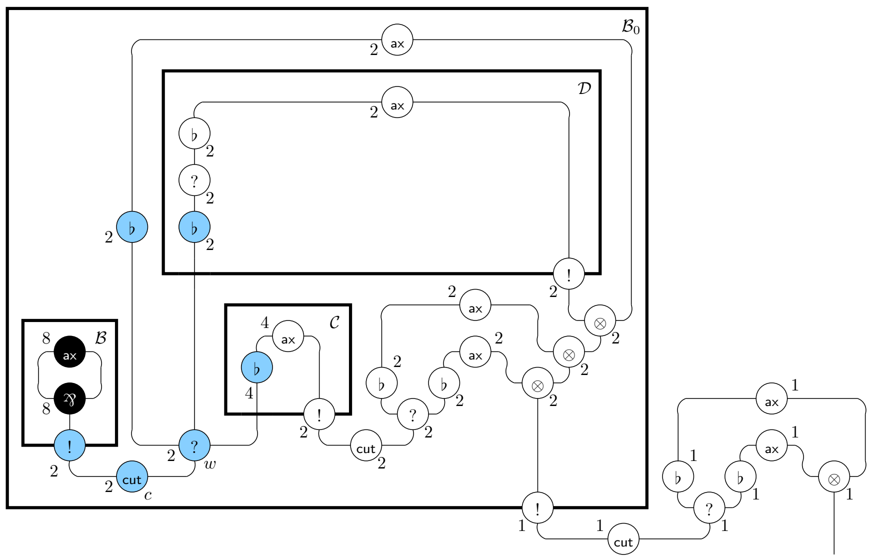

Figure 17: The proof net of Fig. 15 (auxiliary ports are not drawn because irrelevant to the discussion of this section). Levels are omitted, since they are the same as those of Fig. 15. Instead, each link has its potential size relative to level 0 (see Definition 27) annotated beside it.

Definition 25 (Contractive factor). Let π be an $\mathbf{mL^3}$ proof net, and let $\mathcal{B}$ be a box of π. The *contractive factor* of $\mathcal{B}$, denoted by $\mu(\mathcal{B})$, is then defined as follows:

$$\mu(\mathcal{B}) = \sum_{\mathcal{B} \preceq^{\mathbf{L}} \mathcal{C}} \nabla(\mathcal{C}).$$

Lemma 19. *Let π be an $\mathbf{mL^4}$ proof net, and let $\mathcal{B}$ be a box of π. Then,*

$$\mu(\mathcal{B}) = \nabla(\mathcal{B}) + \sum_{\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}} \mu(\mathcal{C}).$$

Proof. Simply observe that, by Lemma 18, the set $\{\mathcal{C} \; ; \; \mathcal{B} \preceq^{\mathbf{L}} \mathcal{C}\}$ can be partitioned into $\{\mathcal{B}\} \cup \bigcup_{\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}} \{\mathcal{D} \; ; \; \mathcal{C} \preceq^{\mathbf{L}} \mathcal{D}\}$. □

Definition 26 (Duplication factor). Let π be an $\mathbf{mL^3}$ proof net, and let $\mathcal{B}$ be a box of π. The *duplication factor* of $\mathcal{B}$, denoted by $\delta(\mathcal{B})$, is the following non-negative integer:

$$\delta(\mathcal{B}) = \prod_{\mathcal{B} \subseteq \mathcal{C}} \mu(\mathcal{C}),$$

where only boxes at the same level as $\mathcal{B}$ are considered in the product.

Still referring to Fig. 17, we have $\mu(\mathcal{B}) = \nabla(\mathcal{B}) + \nabla(\mathcal{C}) = 4$, while the contractive factors of $\mathcal{C}$ and $\mathcal{B}_0$ are equal to their arities, because these boxes are maximal in the light contractive order. This gives $\delta(\mathcal{B}) = \mu(\mathcal{B})\mu(\mathcal{B}_0) = 8$, $\delta(\mathcal{C}) = \mu(\mathcal{C})\mu(\mathcal{B}_0) = 4$, while $\mathcal{B}_0$ is maximal w.r.t. $\subseteq$ and so $\delta(\mathcal{B}_0) = \mu(\mathcal{B}_0) = 2$.

Intuitively, the duplication factor of a box $\mathcal{B}$ at level i says how many copies of the content of $\mathcal{B}$ will be present at the end of the round at level i of our cut-elimination procedure. In fact, the contractive factor takes into account the duplications originating from "chains" of boxes at the same depth; to obtain the duplication factor of a box $\mathcal{B}$, one must multiply the contractive factors of all boxes containing $\mathcal{B}$.

This is well shown in Fig. 17: when one reduces the cut link c, 3 copies of the content of $\mathcal{B}$ are made, but one of them will be duplicated again when the cut concerning $\mathcal{C}$ is reduced, so $4 = \mu(\mathcal{B})$ copies are actually produced. We are not quite done though: the reduction of the cut concerning $\mathcal{B}_0$ yields a further duplication of (the residues of) the content of $\mathcal{B}$. Indeed, we invite the reader to check that exactly $8 = \delta(\mathcal{B})$ residues of the content of $\mathcal{B}$ are present in the normal form of the proof net of Fig. 17.

This motivates the following definition:

Definition 27 (Potential size). Let π be an $\mathbf{mL^3}$ proof net, and $k \in \mathbb{Z}$. The *potential size* relative to k of a link a of π, denoted by $[a]_k$, is defined as follows: let $\mathcal{B}$ be the minimal box w.r.t. $\subseteq$ of level k containing a; if $\mathcal{B}$ exists, we set $[a]_k = \delta(\mathcal{B})$, otherwise $[a]_k = 1$. The *potential size* relative to k of π is simply the sum of the potential sizes of its links:

$$[\pi]_k = \sum_a [a]_k,$$

where a ranges over all links of π which are not auxiliary ports.

As suggested above, $[\pi]_i$ is intended to give an estimate of the size of the proof net obtained by executing the round-by-round procedure at level i. This intuition is formalized by the following result:

Lemma 20. *Let π be an i-contractive $\mathbf{mL^4}$ proof net. Then:*

1. *if π is i-normal, then $[\pi]_i = |\pi|$;*
2. *if $\pi \to \pi'$ by reducing a minimal* cut *link (in the cut order) at level i, then $[\pi']_i < [\pi]_i$.*

Proof. Part 1 is easy: simply observe that, if there is no reducible cut link at level i, then for all $\mathcal{B}$ at level i, by definition we have $\nabla(\mathcal{B}) = 1$. From this, since every box is maximal in the contractive order, we deduce $\mu(\mathcal{B}) = \nabla(\mathcal{B}) = 1$ for all $\mathcal{B}$ at level i, and similarly $\delta(\mathcal{B}) = 1$. This implies $[a]_i = 1$ for any link a of π, which proves the result.

The proof of part 2 is based on a careful inspection of Fig. 7. We call the $\mathsf{why\ not}$ link and the box reduced by the step resp. w and $\mathcal{B}$. We also follow the convention that all links/boxes of π will be denoted by "simple" letters $(a, \mathcal{C}, \ldots)$, while the links/boxes of π' will be denoted by letters with a "prime" $(a', \mathcal{C}', \ldots)$; it shall be assumed that if the names of two links/boxes of resp.

π, π' differ only because of the absence/presence of a "prime", then one is the lift/residue of the other. For example, a is the lift of a', $\mathcal{C}$ is the lift of $\mathcal{C}'$, etc. The links of π are partitioned into three classes (we ignore auxiliary ports because they are not taken into account by the potential size):

C_1: links represented in Fig. 7 having a residue in π'; these are exactly the content of $\mathcal{B}$ (i.e., the links contained in the pre-net called π_0 in the picture), and, if present, the why not link of conclusion $?\Gamma$ (recall that, by the lightness condition, Γ is at most one formula; if Γ is empty, this link is not present);

C_2: links represented in Fig. 7 having no residue in π'; these are exactly w, the principal port of $\mathcal{B}$, the cut link reduced by the step, and all of the flat links shown;

C_3: all other links of π, i.e., those "outside of the picture" in Fig. 7. These links have exactly one residue in π'.

Similarly, the links of π' can be partitioned into the following three classes:

C_1': links having a lift of class 1 in π; these are exactly the links contained in one of the copies of π_0, and (if present) the why not link of conclusion $?\Gamma$;

C_2': links having no lift in π; these are exactly all of the cut links represented in the right member of Fig. 7;

C_3': links having a lift of class 3 in π.

The class of a box of π or π' will be the one of its principal port.

Intuitively, in π (resp. π'), a link of class 1 is a link which will be (resp. has been) duplicated or altered by the execution of the step; a link of class 2 is a link that disappears during (resp. is created by) the execution of the step; and a link of class 3 is a link to which "nothing will happen" (resp. "nothing has happened") during the execution of the step.

Before continuing with the proof, we invite the reader to pause a moment and look again at Fig. 17. The proof net in the picture, which we denote by π, is readily seen to be 0-contractive. As already noted above, the contractive order at level 0 is $\mathcal{B} \preceq \mathcal{C} \preceq \mathcal{B}_0$, so the minimal cut in the cut order is the one denoted by c. After reducing it, we obtain the proof net π' given in Fig. 18. In both figures, links filled with a dark shade are of class 1, those filled with a light shade are of class 2, and unfilled links are of class 3.

We shall now verify part 2 of the lemma on this concrete example, by counting the links in each class and their potential sizes. We start with class 1 (dark-filled links). There are only 2 such links in π: the par and axiom link inside $\mathcal{B}$. The deepest box of level 0 containing them is precisely $\mathcal{B}$, so their potential size is $\delta(\mathcal{B}) = 8$. Therefore, the potential size of class 1 links of π is 16. For what concerns π', we find 3 copies of these two links: one inside $\mathcal{C}'$, one inside $\mathcal{D}'$, and one strictly inside $\mathcal{B}_0'$. The first ones have potential size $\delta(\mathcal{C}') = 4$, and the last ones $\delta(\mathcal{B}_0') = 2$. For concerns the remaining copy, although it is contained in $\mathcal{D}'$, this box has level 1, so the potential size is again $\delta(\mathcal{B}_0') = 2$. Hence, the total potential size is $8 + 4 + 4 = 16$, i.e., identical to that of the links of class 1 of π.

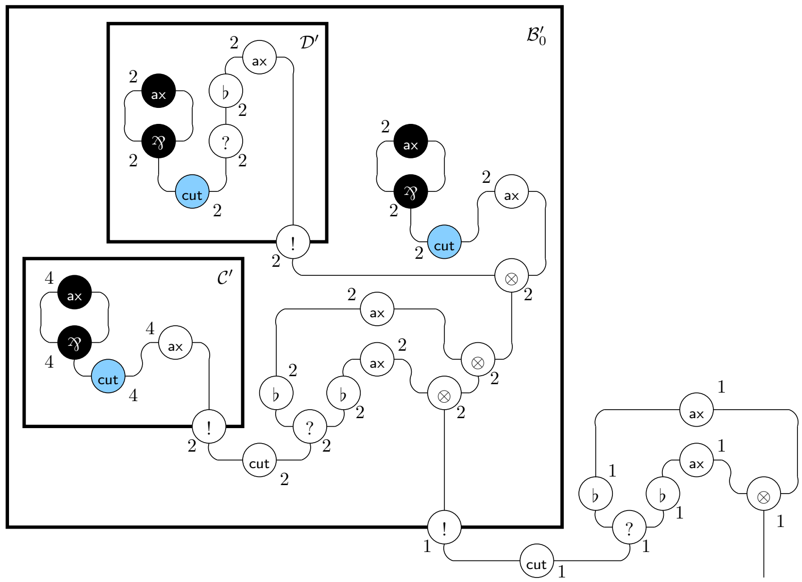

Figure 18: The result of reducing the cut link c in the proof net of Fig. 17.

We may now turn to the links of class 2 (light-filled links). In π, there are 6 of these, all of potential size 2 except the flat link inside $\mathcal{C}$, which has potential size $\delta(\mathcal{C}) = 4$. The overall contribution to the potential size of π from the links of class 2 is therefore 14. In π', all of these links have disappeared, and have been replaced by 3 cuts at level 1. Just as the flat links of class 2 in π, two of these cut links have potential weight 2, and one 4, giving a total of $8 < 14$. Hence, in going from π to π' we have lost the potential size of the three links of class 2 of π directly involved in the cut, i.e., the principal port of $\mathcal{B}$, the $\mathsf{why\ not}$ link w, and the cut link c itself.

Finally, we consider the links of class 3 (unfilled links). We invite the reader to check that, for each link a of class 3 in π, there is exactly one residue a' in π', and $[a]_0 = [a']_0$. Therefore, the contribution to the potential size of the links in this class is preserved under reduction, and in the end we get $[\pi']_0 < [\pi]_0$, as stated in the lemma.

We may now resume the proof. First of all, we recall the following fundamental fact, which holds by the minimality of the cut under reduction:

Fact *If $\mathcal{B}_1$ is a box of level i such that $\mathcal{B}_1 \subseteq \mathcal{B}$, then $\mathcal{B}_1$ is not involved in a reducible cut.*

The above fact can be used to infer the following series of preliminary results (before even reading the proofs, we strongly invite the reader to verify each one of them on the examples of Fig. 17 and 18):

Claim 1. *Let $\mathcal{B}'_1, \mathcal{B}'_2$ be two boxes of level i. Then, $\mathcal{B}'_1 \preceq^{\mathbf{L}} \mathcal{B}'_2$ iff $\mathcal{B}_1 \preceq^{\mathbf{L}} \mathcal{B}_2$.*

Proof. Start by supposing that $\mathcal{B}_1' \prec_1^{\mathbf{L}} \mathcal{B}_2'$. By definition, $\mathcal{B}_1'$ is cut, by means of a cut link c', with a why not link above which there is exactly one (by the lightness condition) flat link inside $\mathcal{B}_2'$. Observe that there are no cut links of class 1 in π', so c' must be either of class 2 or 3. In the second case, obviously $\mathcal{B}_1'$ and $\mathcal{B}_2'$ are also of class 33, so $\mathcal{B}_1 \preceq^{\mathbf{L}} \mathcal{B}_2$. The first case is actually impossible, because the premises of c' would be of level $i+1$, hence none of them could be conclusion of the principal port of $\mathcal{B}_1'$.

Suppose now that $\mathcal{B}_1 \prec_1^{\mathbf{L}} \mathcal{B}_2$. Note firstly that we are supposing $\mathcal{B}_1, \mathcal{B}_2$ to be the lifts of resp. $\mathcal{B}_1'$ and $\mathcal{B}_2'$, so neither of $\mathcal{B}_1, \mathcal{B}_2$ can be equal to $\mathcal{B}$. If they are both of class 3, we immediately have $\mathcal{B}_1' \prec_1^{\mathbf{L}} \mathcal{B}_2'$. Suppose now that $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{B}_1$. We cannot have $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{B}_2$, because this would contradict Lemma 18. Therefore, $\mathcal{B}_2$ is of class 3, and again obviously $\mathcal{B}_1' \prec_1^{\mathbf{L}} \mathcal{B}_2'$. We are left with the case in which $\mathcal{B}_1$ is of class 1 and $\mathcal{B} \not\prec_1^{\mathbf{L}} \mathcal{B}_1$. The only possibility would be that $\mathcal{B}_1 \subseteq \mathcal{B}$, but this is excluded by the above Fact, since we have supposed that $\mathcal{B}_1$ is involved in a reducible cut. We have thus shown that $\mathcal{B}_1' \prec_1^{\mathbf{L}} \mathcal{B}_2'$ iff $\mathcal{B}_1 \prec_1^{\mathbf{L}} \mathcal{B}_2$, which obviously implies our claim. □

Claim 2. *Let $\mathcal{C}'$ be a box of level i. Then, $\nabla(\mathcal{C}') = \nabla(\mathcal{C})$.*

Proof. If $\mathcal{C}'$ is not involved in a cut, then neither is $\mathcal{C}$, so in this case the statement is obvious. In case $\mathcal{C}'$ is involved in a cut c', this cannot be one of the links of class 2 of π', because they are all at level $i+1$. Therefore, $\mathcal{C}$ is also involved in a cut, with a why not link that we may call u. Now notice that, if u is of class 3, then the arities of u and u' coincide, and everything "above" u is also of class 3, so the statement holds. But this is actually the only possibility: in fact, if u were of class 1, it is easy to see that u would have to be the unique (by the lightness condition) why not link such that, among its premises, there is (by the stratification condition) the conclusion of the auxiliary port of $\mathcal{B}$. In this case, we would obtain $\mathcal{C} \prec_1 \mathcal{B}$, contradicting the minimality of the cut under reduction. □

Claim 3. *If $\mathcal{B}_1$ is a box of level i such that $\mathcal{B}_1 \subseteq \mathcal{B}$, then $\mu(\mathcal{B}_1) = 1$.*

Proof. In fact, $\mu(\mathcal{B}_1) > 1$ would imply, by definition, that $\mathcal{B}_1$ is involved in a contractive cut, which is impossible by the above Fact. □

Claims 1 and 2 have the following fundamental corollary:

Claim 4. *If $\mathcal{C}$ is a box of class 3 of π at level i, then $\delta(\mathcal{C}') = \delta(\mathcal{C})$.*

Proof. Claims 1 and 2 immediately imply that, whenever $\mathcal{D}$ is of class 3, $\mu(\mathcal{D}') = \mu(\mathcal{D})$. Now, any box containing a box of class 3 in π is also of class 3, so if $\mathcal{D}_1, \ldots, \mathcal{D}_n$ are the nested boxes of level i surrounding $\mathcal{C}$ in π, then in π' we have boxes $\mathcal{D}_1', \ldots, \mathcal{D}_n'$ of level i containing $\mathcal{C}'$, with $\mu(\mathcal{D}_j') = \mu(\mathcal{D}_j)$ for all $1 \leq j \leq n$, which proves the claim. □

Let now $a_3 \in C_3$, and let a_3' be its unique residue. It is not hard to see that, if a_3 is not contained in any box at level i, then neither is a_3', in which case $[a_3]_i = [a'_3]_i = 1$. Otherwise, let $\mathcal{B}_0$ be the minimal box (w.r.t. $\subseteq$) of level i containing a_3. Observe that $\mathcal{B}_0 \not\subseteq \mathcal{B}$, because otherwise a_3 would not be of class 3. Therefore, $\mathcal{B}_0$ has a unique residue $\mathcal{B}'_0$, and both are of class 3. By Claim 4, $\delta(\mathcal{B}_0) = \delta(\mathcal{B}'_0)$, so again $[a_3] = [a'_3]$. Recalling that every link of class 3 of π has exactly one residue in π', this shows that

$$\sum_{a_3 \in C_3} [a_3]_i = \sum_{a'_3 \in C'_3} [a'_3]_i.$$

Let instead $a_1 \in C_1$. If a_1 is the why not link of conclusion $?\Gamma$, then it has a unique residue a'_1; in this case, by the same reasoning given above for links of class 3, we can easily infer that $[a_1]_i = [a'_1]_i$. Otherwise, a_1 is a link belonging to the pre-net called π_0 in Fig. 7. In this case, a_1 is contained in a box $\mathcal{B}_1 \subseteq \mathcal{B}$ at level i; more precisely, there are n boxes $\mathcal{B}_1, \ldots, \mathcal{B}_n$, all at level i, such that a_1 is in $\mathcal{B}_1$ and $\mathcal{B}_1 \subseteq \cdots \subseteq \mathcal{B}_n \subseteq \mathcal{B}$, where each inclusion is immediate, i.e., there is no box at level i between $\mathcal{B}_j, \mathcal{B}_{j+1}$ and $\mathcal{B}_n, \mathcal{B}$. Now, let $\Delta = \delta(\mathcal{B}_0)$, where $\mathcal{B}_0$ is the minimal (w.r.t. $\subseteq$) box of level i containing $\mathcal{B}$, or let $\Delta = 1$ if no such box exists. By Claim 3, we have $[a_1]_i = \delta(\mathcal{B}_1) = \Delta\mu(\mathcal{B})$.

Consider now a residue a'_2 of a_2. Each of the $\mathcal{B}_j$ above has a corresponding residue $\mathcal{B}'_j$ at level i containing a'_2, such that $\mathcal{B}'_1 \subseteq \cdots \subseteq \mathcal{B}'_n$. Since the structure of π_0 is not changed in the duplication, each $\mathcal{B}'_j$ is maximal in the light contractive order and is not involved in a reducible cut, so $\mu(\mathcal{B}'_j) = 1$ for all j. There are now two cases:

1. $\mathcal{B}'_n$ is not contained in any box of level i, or the minimal (w.r.t. $\subseteq$) box containing it is $\mathcal{B}'_0$. Then, $[a'_1]_i = \Delta$. In fact, in case it exists, $\mathcal{B}_0$ is of class 3, so by Claim 4, $\delta(\mathcal{B}'_0) = \delta(\mathcal{B}_0) = \Delta$;
2. There is a box $\mathcal{C}'$ of level i strictly contained in $\mathcal{B}'_0$ and containing $\mathcal{B}'_n$. In this case, by inspecting Fig. 7 under the stratification condition, it is not hard to see that $\mathcal{C}'$ is the unique residue of a box $\mathcal{C}$ such that $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}$. Observe that $\mathcal{C}$ is of class 3, so by Claim 4 we have $[a'_1]_i = \delta(\mathcal{C}') = \delta(\mathcal{C}) = \Delta\mu(\mathcal{C})$.

If the arity of w is $k \geq 1$, there are k residues of a_1. Observe that case 1 applies to exactly $\nabla(\mathcal{B})$ of them, while case 2 applies to all other residues, and, because of the lightness condition, there is exactly one residue of this latter kind for each $\mathcal{C}$ such that $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}$. So, if we denote by A'_1 the set of all residues of a_1, we have, using Lemma 19,

$$\sum_{a'_1 \in A'_1} [a'_1]_i = \Delta\nabla(\mathcal{B}) + \sum_{\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}} \Delta\mu(\mathcal{C}) = \Delta\mu(\mathcal{B}) = [a_1]_i.$$

If we put together what we have said up to now, we obtain an identical result for the links of class 1 as the one obtained above for the links of class 3:

$$\sum_{a_1 \in C_1} [a_1]_i = \sum_{a'_1 \in C'_1} [a'_1]_i.$$

We now get to the links of class 2, starting with those of π. The principal port of $\mathcal{B}$, w, and c, have all potential size Δ, where Δ is the same quantity introduced above. For what concerns the flat links shown in the picture, $\nabla(\mathcal{B})$ of them have again potential weight Δ, while the others are each immediately (by the stratification condition) contained in a different (by the lightness condition) box $\mathcal{C}$ such that $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}$, in which case the potential size is $\Delta\mu(\mathcal{C})$. Therefore, we have

$$\sum_{a_2 \in C_2} [a_2]_i = 3\Delta + \Delta\nabla(\mathcal{B}) + \sum_{\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}} \Delta\mu(\mathcal{C}) = \Delta(3 + \mu(\mathcal{B})).$$

On the other hand, the only links of class 2 of π' are the cut links shown in the picture. Exactly $\nabla(\mathcal{B})$ of these have potential size Δ, while the rest have each potential size $\delta(\mathcal{C}')$, where $\mathcal{C}$ is a box such that $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}$ (of course we are implicitly using the above Claims to infer these facts). But, using Claim 4, we have that $\delta(\mathcal{C}') = \delta(\mathcal{C}) = \Delta\mu(\mathcal{C})$, for all $\mathcal{C}$ as above. Therefore, remembering that $\Delta \geq 1$, we obtain

$$\sum_{a_2' \in C_2'} [a_2']_i = \Delta\nabla(\mathcal{B}) + \sum_{\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{C}} \Delta\mu(\mathcal{C}) = \Delta\mu(\mathcal{B}) < \sum_{a_2 \in C_2} [a_2]_i,$$

which concludes the proof of part 2. □

We remark that the strict inequality of part 2 of Lemma 20 is a sort of an "accident", and is of no real technical value: what matters in the statement is that $[\pi]_i$ linearly bounds $[\pi']_i$. Lemma 23 below, which crucially uses Lemma 20, would hold even if we only had $[\pi']_i = [\pi]_i$, and indeed this is true at all levels except level i itself, where the three links directly involved in the cut "disappear", and with them their potential size. More precisely, if we define the quantity $[\pi]_i^j$ as the potential size relative to i of all links of π of level j, then point 2 of Lemma 20 can be replaced by $[\pi']_i^j = [\pi]_i^j$ for all $i \neq j$ and $[\pi']_i^i < [\pi]_i^i$.

As already noted above, the duplication factor of a box $\mathcal{B}$ is influenced not only by the boxes $\mathcal{C}$ at the same depth as $\mathcal{B}$ such that $\mathcal{B} \preceq^{\mathbf{L}} \mathcal{C}$, but also by the boxes at the same *level* as $\mathcal{B}$ which contain it. To quantify this phenomenon, we define the notion of *relative depth*, which will be useful in bounding the potential size of a proof net (Lemma 21) and will be proved to have the same behavior as the level with respect to reduction, i.e., it is non-increasing (Lemma 22).

Definition 28 (Relative depth). Let π be an $\mathbf{mL^3}$ proof net, and let $\mathcal{B}$ be a box of π. We denote by $\widehat{\mathcal{B}}$ the maximal (w.r.t. $\subseteq$) box of π at the same level as $\mathcal{B}$ such that $\mathcal{B} \subseteq \widehat{\mathcal{B}}$. The *relative depth* of $\mathcal{B}$, denoted by $\rho(\mathcal{B})$, is the following non-negative integer:

$$\rho(\mathcal{B}) = \mathrm{d}(\mathcal{B}) - \mathrm{d}(\widehat{\mathcal{B}}).$$

The relative depth of π, also denoted by $\rho(\pi)$, is the maximum relative depth of its boxes.

Observe that, because $\subseteq$ is downward-arborescent, the relative depth of a box $\mathcal{B}$ can be equivalently defined as the number of boxes $\mathcal{C}$ at the same level as $\mathcal{B}$ such that $\mathcal{B} \subseteq \mathcal{C}$, minus one.

Lemma 21. *Let π be an* $\mathbf{mL^3}$ *proof net. Then,* $[\pi]_i \leq |\pi|^{\rho(\pi)+2}$ *for all* $i \in \mathbb{Z}$.

PROOF. Recall from the definition that $[\pi]_i = \sum_a [a]_i$, where the sum ranges over all links of π other than auxiliary ports. Now let $M = \max\{[a]_i\ ;\ a \in \pi\}$. Clearly we have that $[\pi]_i \leq M|\pi|$. Now M must be the duplication factor of a box $\mathcal{B}$ of level i of π. For any such box, we have $\mu(\mathcal{B}) = \sum_{\mathcal{B} \preceq^{\mathbf{L}} \mathcal{C}} \nabla(\mathcal{C})$. Observe that a flat link contributing to the arity of a box cannot contribute to the arity of another box; therefore, even if the sum defining $\mu(\mathcal{B})$ ranged over *every* box of π, we would still have $\mu(\mathcal{B}) \leq |\pi|$. From this, recalling that the relative depth of a box $\mathcal{B}$ of level i is the number of boxes $\mathcal{C}$ of level i such that $\mathcal{B} \subseteq \mathcal{C}$, minus one, we have

$$\delta(\mathcal{B}) = \prod_{\substack{\mathcal{B} \subseteq \mathcal{C} \\ \ell(\mathcal{C}) = i}} \mu(\mathcal{C}) \leq \prod_{\substack{\mathcal{B} \subseteq \mathcal{C} \\ \ell(\mathcal{C}) = i}} |\pi| \leq |\pi|^{\rho(\pi)+1},$$

which concludes the proof. □

Lemma 22. *Let π be an* $\mathbf{mL^4}$ *proof net such that $\pi \to \pi'$. Then, $\rho(\pi') \leq \rho(\pi)$.*

PROOF. The depth of a box $\mathcal{C}$ can only be affected during an exponential step, and only if it is contained in the pre-net called π_0 in Fig. 7. Then, if $\mathcal{C}'$ is a residue of $\mathcal{C}$ in π', by the stratification condition we either have $\mathrm{d}(\mathcal{C}') = \mathrm{d}(\mathcal{C})$ or $\mathrm{d}(\mathcal{C}') = \mathrm{d}(\mathcal{C}) - 1$, so in general $\mathrm{d}(\mathcal{C}') \leq \mathrm{d}(\mathcal{C})$.

Now, call the box under reduction $\mathcal{B}$; observe that $\mathcal{C} \subseteq \mathcal{B}$, so $\mathcal{B}$ and $\widehat{\mathcal{C}}$ cannot be disjoint. If we write $\mathcal{B}_1 \subset \mathcal{B}_2$ for $\mathcal{B}_1 \subseteq \mathcal{B}_2$ and $\mathcal{B}_1 \neq \mathcal{B}_2$, then we can distinguish three cases: either $\widehat{\mathcal{C}} \subset \mathcal{B}$, or $\mathcal{B} \subset \widehat{\mathcal{C}}$, or $\widehat{\mathcal{C}} = \mathcal{B}$. In all cases, we put $\mathcal{D}' = \widehat{\mathcal{C}'}$.

- In the first case, the depth of $\mathcal{D}'$ varies w.r.t. the depth of $\widehat{\mathcal{C}}$ just as the depth of $\mathcal{C}'$ varies w.r.t. the depth of $\mathcal{C}$, so $\rho(\mathcal{C}') = \rho(\mathcal{C})$.

- In the second case, $\mathcal{D}'$ is the unique residue of $\widehat{\mathcal{C}}$, and $\mathrm{d}(\mathcal{D}') = \mathrm{d}(\widehat{\mathcal{C}})$, so

$$\rho(\mathcal{C}') = \mathrm{d}(\mathcal{C}') - \mathrm{d}(\mathcal{D}') \leq \mathrm{d}(\mathcal{C}) - \mathrm{d}(\widehat{\mathcal{C}}) = \rho(\mathcal{C}).$$

- In the third case, we start by supposing that the lift $\mathcal{D}$ of $\mathcal{D}'$ is disjoint from $\mathcal{B}$. Then, the stratification condition gives us that $\mathcal{B} \prec_1^{\mathbf{L}} \mathcal{D}$ and $\mathrm{d}(\mathcal{D}') = \mathrm{d}(\mathcal{D}) = \mathrm{d}(\mathcal{B})$, so that $\rho(\mathcal{C}') = \rho(\mathcal{C})$. Suppose now that $\mathcal{D}$ and $\mathcal{B}$ are not disjoint. Since $\mathcal{B}$ has no residue in π', we have either $\mathcal{B} \subset \mathcal{D}$ or $\mathcal{D} \subset \mathcal{B}$. But the first case is actually impossible, because it would contradict the fact that $\mathcal{B} = \widehat{\mathcal{C}}$, since $\mathcal{D}$ is at the same level as $\mathcal{C}$. Therefore, we must have $\mathcal{D} \subset \mathcal{B}$, so that $\mathrm{d}(\mathcal{B}) < \mathrm{d}(\mathcal{D})$. Now, as in the first case,

$$\rho(\mathcal{C}') = \mathrm{d}(\mathcal{C}') - \mathrm{d}(\mathcal{D}') = \mathrm{d}(\mathcal{C}) - \mathrm{d}(\mathcal{D}) < \mathrm{d}(\mathcal{C}) - \mathrm{d}(\mathcal{B}) = \rho(\mathcal{C}).$$

□

The technical machinery we have been building up through the section will now be used to finally infer our polynomial bound on the reduction of $\mathbf{mL^4}$ proof nets.

Lemma 23. *Let π be an $(i-1)$-normal $\mathbf{mL^4}$ proof net, and let π' be the i-normal proof net obtained from π by applying the round-by-round procedure at level i. Then, $|\pi'| \leq |\pi|^{\rho(\pi)+2}$.*

Proof. We can decompose the reduction from π to π' into $\pi \rightarrow^* \pi_0 \rightarrow^* \pi'$, where π_0 is the first i-contractive proof net obtained during the reduction. Now, applying, in the order, points 1 and 2 of Lemma 20, Lemma 21, Lemma 22, and the well known fact that $|\pi_0| \leq |\pi|$, we obtain

$$|\pi'| = [\pi']_i \leq [\pi_0]_i \leq |\pi_0|^{\rho(\pi_0)+2} \leq |\pi_0|^{\rho(\pi)+2} \leq |\pi|^{\rho(\pi)+2},$$

as desired. □

Theorem 24 (Polynomial bound for $\mathbf{mL^4}$). *Let π be an $\mathbf{mL^4}$ proof net of size s, level l, and relative depth r. Then, the round-by-round procedure reaches a normal form in at most $(l+1)s^{(r+2)^l}$ steps.*

Proof. We start by applying the same arguments used in the beginning of the proof of Theorem 17: we decompose the reduction from π to its normal form π_l into $\pi = \pi_{-1} \rightarrow^* \pi_0 \cdots \rightarrow^* \pi_l$, where each π_i is i-normal; then, using Lemma 15 (which is valid because $\mathbf{mL^4}$ is a subsystem of $\mathbf{mL^3}$), if we call the length of the whole reduction sequence L, we can write

$$L \leq \sum_{i=0}^{l} |\pi_{i-1}|.$$

Now, using Lemma 23, we have, for all $0 \leq i \leq l$, $|\pi_i| \leq |\pi_{i-1}|^{\rho(\pi_{i-1})+2}$. But, by Lemma 22, for all $0 \leq i \leq l$, we have $\rho(\pi_i) \leq \rho(\pi)$, so we can actually write

$$|\pi_i| \leq |\pi_{i-1}|^{r+2}.$$

From this, it can be proved by a straightforward induction that, for all $i \geq 0$, we have $|\pi_{i-1}| \leq s^{(r+2)^i}$. Hence, we obtain

$$L \leq \sum_{i=0}^{l} |\pi_{i-1}| \leq \sum_{i=0}^{l} s^{(r+2)^i} \leq (l+1)s^{(r+2)^l},$$

which is the bound stated in the thesis. □

Observe that, by Proposition 9, if π^+ is the $\mathbf{mL^4}$ embedding of an $\mathbf{mLLL}$ proof net π of size s and depth d, then $|\pi^+| = s$, $\ell(\pi^+) = d$, and $\rho(\pi^+) = 0$, so that normalizing π^+ takes at most $(d+1)s^{2^d}$ steps, which is the same bound given by Girard (1998).

3.4. Characterization of **FE** *and* **FP**

Propositions 8 and 9 tell us that $\mathbf{mL^3}$ and $\mathbf{mL^4}$ are conservative extensions of **mELL** and **mLLL**, so programming in the former systems can be done using the same types and proofs as in the latter. In particular, the type of finite binary strings in $\mathbf{mL^3}$ and $\mathbf{mL^4}$ are respectively

$$\begin{aligned} \mathbf{S_E} &= \forall X.(?(X^\perp \otimes X) \parr ?(X^\perp \otimes X) \parr !(X^\perp \parr X)), \\ \mathbf{S_P} &= \forall X.(?(X^\perp \otimes X) \parr ?(X^\perp \otimes X) \parr \S(X^\perp \parr X)). \end{aligned}$$

Then, one can represent binary strings as in Girard (1998) and Danos and Joinet (2003). In the following, we write $!^k A$ (resp. $\S^k A$) for the formula A preceded by k of course (resp. paragraph) modalities, and if φ and ξ are two proof nets of respective conclusions $A^\perp, B$ and A, we denote by $\varphi(\xi)$ the proof net of conclusion B obtained from φ and ξ by adding a cut link whose premises are the conclusions of type $A^\perp, A$ of resp. φ and ξ.

Definition 29 (Representation). A function $f : \{0,1\}^* \to \{0,1\}^*$ is *representable* in $\mathbf{mL^3}$ (resp. $\mathbf{mL^4}$) iff there exists $k \in \mathbb{N}$ and a proof net φ of conclusions $\mathbf{S_E^\perp}, !^k\mathbf{S_E}$ (resp. $\mathbf{S_P^\perp}, \S^k\mathbf{S_P}$) such that $f(x) = y$ iff $\varphi(\xi) \to^* \upsilon$, where ξ is the proof net of conclusion $\mathbf{S_E}$ (resp. $\mathbf{S_P}$) representing x, and υ is the proof net of conclusion $!^k\mathbf{S_E}$ (resp. $\S^k\mathbf{S_P}$) which is the representation of y enclosed in k boxes (resp. followed by k paragraph links). We denote by $\mathbf{FmL^3}$ (resp. $\mathbf{FmL^4}$) the class of functions representable in $\mathbf{mL^3}$ (resp. $\mathbf{mL^4}$).

A fundamental remark now is that the level and relative depth of the representation of a datum do not depend on the datum itself: all cut-free proof nets of type $\mathbf{S_E}$ representing binary strings in $\mathbf{mL^3}$ have level 1, and all cut-free proof nets of type $\mathbf{S_P}$ representing binary strings in $\mathbf{mL^4}$ have level 1 and relative depth 0. In both cases, the size of the proof net is equal to $3n + 6$, where n is the length of the string represented.

Thanks to the above, the soundness of $\mathbf{mL^3}$ and $\mathbf{mL^4}$ with respect to **FE** and **FP**, respectively, is a consequence of Theorems 17 and 24, modulo the arguments given at the end of Sect. 1.4. For the completeness side we have:

Proposition 25. *Any function $f : \{0,1\}^* \to \{0,1\}^*$ computable on a Turing machine in time $\mathcal{O}(2^n_d)$ can be represented in $\mathbf{mL^3}$ by a proof net of level d and of conclusions $\mathbf{S_E^\perp}, !^d\mathbf{S_E}$.*

Any function $f : \{0,1\}^* \to \{0,1\}^*$ computable on a Turing machine in time $\mathcal{O}(n^{2^d})$ can be represented in $\mathbf{mL^4}$ by a proof net of level d and of conclusions $\mathbf{S_P^\perp}, \S^d\mathbf{S_P}$.

Proof. Let us start with the second statement. First, Mairson and Terui (2003) show that a $\mathcal{O}(2^n_d)$ function can be represented in **mLLL** by a proof net of depth d and of conclusions $\mathbf{S_P^\perp}, \S^d\mathbf{S_P}$. Now we can obtain our statement by using the fact that any **mLLL** proof net of depth d gives an $\mathbf{mL^4}$ proof net of level d (Proposition 9).

As to the first statement, we have already recalled in the discussion after Theorem 17 that Danos and Joinet (2003) give an encoding of the function $n \mapsto 2^n_d$ in **mELL** as a proof net of depth d of conclusions $\mathbf{N}^\perp, !^d\mathbf{N}$, where $\mathbf{N}$ is a type for tally integers. Using this fact and the encoding of Turing machines in **mELL** following the one from Mairson and Terui (2003), we obtain that a function of $\mathcal{O}(2^n_d)$ can be represented in **mELL** by a proof net of depth d and of conclusions $\mathbf{S}_\mathbf{E}^\perp, !^d\mathbf{S}_\mathbf{E}$. We then conclude as above, recalling that any **mELL** proof net of depth d gives an $\mathbf{mL^3}$ proof net of level d (Proposition 8). □

Hence, we finally have:

Theorem 26 (Characterization of FE and FP). $\mathbf{FmL^3}$ *and* $\mathbf{FmL^4}$ *coincide respectively with* **FE** *and* **FP**.

Observe that, due to the isomorphism $\S(A \parr B) \cong \S A \parr \S B$, in $\mathbf{mL^4}$ one may use the type $\mathbf{S'_P} = \forall X.(?(X^\perp \otimes X) \parr ?(X^\perp \otimes X) \parr (\S X^\perp \parr \S X))$ with virtually no difference, i.e., Theorem 26 still holds if we represent binary strings with this modified type.

4. Restricting the Language of Formulas

We have already observed that in $\mathbf{mL^4}$ there are the following isomorphisms:

$$\S(A \otimes B) \cong \S A \otimes \S B \qquad \S!A \cong !\S A \qquad \S\forall X.A \cong \forall X.\S A.$$

(Of course these isomorphisms hold in $\mathbf{mL^3}$ too, but we shall only deal with the polytime system in this section, since the paragraph modality is not really needed in $\mathbf{mL^3}$). More generally, given a formula A containing §, we may find several isomorphic formulas by commuting § connectives with other connectives. This implies that given a proof π of conclusion A, there are several computationally equivalent proofs that are obtained by composing π with isomorphisms.

Hence, if we want to use $\mathbf{mL^4}$, or a fragment of it, as a type system for λ-terms, we will have for each term the choice between several types which carry essentially the same information.

A natural idea at this point is to choose a representative of each equivalence class of formulas, so as to obtain a "canonical" syntax. Given an $\mathbf{mL^4}$ formula A, the obvious candidates to represent the equivalence class of A are the formula in which all paragraphs have been pulled as close as possible to the root, and the formula in which all paragraphs have been pushed to the atoms. Clearly, only this latter choice is stable under composition of formulas (or prefixing with quantifiers and modalities); therefore, we shall draw our attention to the sublanguage of $\mathbf{mL^4}$ in which § *connectives are only applied to atoms*, and we shall define a logical system, called $\mathbf{mL^4_0}$, which uses such sublanguage.

To simplify the notations we shall replace $\S^p X$ by the notation pX and let p range over $\mathbb{N}$. Thus, the language of formulas of $\mathbf{mL^4_0}$, denoted by $Form_0$, will be generated by the following grammar:

$$A, B ::= pX \mid pX^\perp \mid A \otimes B \mid A \parr B \mid !A \mid ?A \mid \exists X.A \mid \forall X.A,$$

where $p \in \mathbb{N}$. Linear negation is defined as expected: $(pX)^\perp = pX^\perp$, $(pX^\perp)^\perp = pX$, and $(\cdot)^\perp$ commutes with all connectives, replacing the given connective with its dual.

Given $p \in \mathbb{N}$ and a formula $A \in \mathit{Form}_0$, we define $p \cdot A$ by induction on A as follows:

$$\begin{array}{rcl} p \cdot (qX) & = & (p+q)X \\ p \cdot (qX^\perp) & = & (p+q)X^\perp \\ p \cdot (A \bullet B) & = & (p \cdot A) \bullet (p \cdot B), \text{ where } \bullet \in \{\otimes, \parr\} \\ p \cdot \dagger A & = & \dagger(p \cdot A), \text{ where } \dagger \in \{!, ?\} \\ p \cdot \nabla X.A & = & \nabla X.(p \cdot A), \text{ where } \nabla \in \{\forall, \exists\}. \end{array}$$

Lemma 27. *For any $p, q \in \mathbb{N}$ and $A \in \mathit{Form}_0$, we have*

$$\begin{array}{rcl} p \cdot (q \cdot A) & = & (p+q) \cdot A, \\ 0 \cdot A & = & A. \end{array}$$

Therefore, $\cdot$ is a monoid action on Form_0.

It is a straightforward consequence of the definition that whenever a formula $A \in \mathit{Form}_0$ is equal to $p \cdot B$ for some B, then all subformulas of A are also of the form $p \cdot B'$ for some subformula B' of B. Also, it is easy to check that $(p \cdot A)^\perp = p \cdot A^\perp$.

In the language of formulas we could actually let p range over $\mathbb{Z}$ instead of $\mathbb{N}$, and define a group action. We would then keep the same properties, but here we stick to $\mathbb{N}$ in order to have a clearer correspondence with $\mathbf{mL^4}$ (that will be described below).

We now introduce a notion of substitution adapted to the formulas of Form_0:

Definition 30. For $A, B \in \mathit{Form}_0$ we define $A\{B/X\}$ by induction on A:

- if $A = pX$: $pX\{B/X\} = p \cdot B$,
- if $A = pX^\perp$: $pX^\perp\{B/X\} = p \cdot B^\perp$,
- and $\{B/X\}$ commutes with all connectives; for instance,

$$(A_1 \otimes A_2)\{B/X\} = A_1\{B/X\} \otimes A_2\{B/X\}.$$

We may now proceed to introducing the system $\mathbf{mL_0^4}$. For this, we first need to define a suitable class of proof nets using the formulas of Form_0.

Definition 31 ($\mathbf{meLL_0}$ proof nets). The nets of $\mathbf{meLL_0}$ are defined as in Definition 1, but for the following modifications (w.r.t. Fig. 1):

- edges are labelled by formulas in Form_0;

- there is no paragraph link;
- axiom links may have conclusions $p \cdot A^\perp$, A, for any $p \in \mathbb{N}$;
- exists links have premise and conclusion with resp. types $A\{B/X\}$ and $\exists X.A$.

The proof nets of $\mathbf{meLL_0}$ are defined from these nets as in Definition 4.

The intuition behind the unusual typing of the axiom link is that it corresponds in $\mathbf{L^4}$ to a proof of $\S^k A^\perp, A$, so an axiom followed by a series of paragraph links. However in $\mathbf{mL^4_0}$ paragraphs are only on atoms, and this is why we have a conclusion $p \cdot A^\perp$ instead of $\S^k A^\perp$.

Cut-elimination for $\mathbf{meLL_0}$ proof nets is defined as in $\mathbf{meLL}$ (Figures 4 through 8), except for the quantifier step (Fig. 6), which uses the substitution $A\{B/X\}$ instead of $A[B/X]$, and for the axiom step (Fig. 4), which is treated as follows.

Let π be a $\mathbf{meLL_0}$ proof net, and let e be an edge of π. We say that a link l of π is *above* e if there exists a directed path from the conclusion of l to e. We define the *tree* of e, denoted by $\mathcal{T}(e)$, as the tree (ignoring boxes) whose root is e and whose leaves are the conclusions of all the axiom and weakening links above e. The axiom links above e are partitioned into three classes:

- a *neutral axiom* is an axiom link such that both of its conclusions are leaves of $\mathcal{T}(e)$;
- a *negative axiom* is an axiom link whose conclusions are labelled by $p \cdot A^\perp$, A and such that only the conclusion labelled by $p \cdot A^\perp$ is a leaf of $\mathcal{T}(e)$;
- a *positive axiom* is an axiom link whose conclusions are labelled by $p \cdot A^\perp$, A and such that only the conclusion labelled by A is a leaf of $\mathcal{T}(e)$.

If, in the negative or positive case, $p = 0$, then the axiom may be considered as either positive or negative.

Now, suppose that π contains a cut link such that one premise is e and the other premise is the conclusion e' of an axiom link a. The reduction of such a cut depends on whether a is positive or negative with respect to e' (it cannot be neutral, because $\mathcal{T}(e')$ has only one leaf, e' itself):

negative: we may assume that e' is labelled by $p \cdot A^\perp$, so that e is labelled by $p \cdot A$ and the other conclusion e'' of a is labelled by A. In this case, π reduces to the proof net π' obtained as follows:

  - remove a, and make e coincide with e'';
  - since e is labelled by $p \cdot A$, all formulas labelling the edges of $\mathcal{T}(e)$ must be of the form $p \cdot B$ (cf. the remark after Lemma 27); then, in π' replace each $p \cdot B$ with B. It is easy to see that such a tree will have conclusion A;

- after this relabeling, if an axiom is neutral w.r.t. e, its conclusions will change from $p \cdot B, q \cdot p \cdot B^\perp$ to $B, q \cdot B^\perp$, so its residue is a valid axiom of $\mathbf{meLL_0}$; if an axiom is positive or negative w.r.t. e, there is nothing to check because only one of its conclusions has been affected.

positive: we may assume that e' is labelled by $A^\perp$, so that e is labelled by A and the other conclusion e'' of a is labelled by $p \cdot A$. In this case, π reduces to the proof net π' obtained as follows:

- remove a, and make e coincide with e'';
- for each formula B labelling an edge of $\mathcal{T}(e)$, in π' label the corresponding edge with $p \cdot B$; it is easy to see that such a tree will have conclusion $p \cdot A$;
- it is also easy to check that all axioms in π' are still correctly labelled, just as in the negative case.

Definition 32 (Indexing). An indexing I for a $\mathbf{meLL_0}$ proof net is defined as in Definition 12 but for the following modification: if e, e' are the conclusions of an axiom link with respective types $p \cdot A^\perp$ and A, then I should satisfy $I(e') = I(e) + p$.

Definition 33 ($\mathbf{mL_0^4}$). The system $\mathbf{mL_0^4}$ is composed of all the proof nets of $\mathbf{meLL_0}$ admitting an indexing as in Definition 32 and satisfying the Stratification and Lightness conditions of Definition 16.

It only takes a (tedious) case-by-case inspection to check that the above definition is sound, i.e., that $\mathbf{mL_0^4}$ is stable under cut-elimination.

Note that, because of the constraint on axiom links (Definition 32), the possibility of assigning an indexing to a $\mathbf{meLL_0}$ proof net depends on the typing, in sharp contrast with the case of **meLL** proof nets. Because of this, defining an untyped version of $\mathbf{mL_0^4}$ cannot be done as easily as for $\mathbf{mL^4}$ (i.e., just forgetting the formulas).

A possible solution is the following. Consider a family of "p-links", with $p \in \mathbb{N}^*$, to be added to the usual links of untyped **meLL** proof nets. The effect of a p-link is to "change the level by p", i.e., a p-link has one premise and one conclusion, whose levels must be resp. $i + p$ and i (if typed, a p-link would have premise A and conclusion $p \cdot A$). We add the restriction that the premise of a p-link must be the conclusion of an axiom link, and that each axiom has at most one p-link "below". Cut-elimination handles p-links by suitably adapting the axiom steps to an untyped framework. We shall not give any detail of this; the informal sketch we just gave is enough for our purposes.

Surprisingly, normalization fails in this system: there are untyped $\mathbf{mL^4}$ proof nets whose reduction goes on forever. Perhaps this is not so strange after all: these p-links basically add the possibility of "changing the level at will", hence they completely break the fundamental invariant of $\mathbf{mL^3}$ and $\mathbf{mL^4}$ proof nets (in fact, the level of an untyped $\mathbf{mL_0^4}$ proof net may increase under reduction).

The above discussion implies that it is impossible to adapt the arguments of Theorem 24 to prove a complexity bound for $\mathbf{mL^4_0}$. Nonetheless, in the rest of the section we shall argue that this system still characterizes deterministic polytime computation.

In what follows, we denote by *Form* the set of **meLL** formulas as defined in Sect. 1.1, i.e., including the paragraph modality. We shall now introduce two translations between our two systems:

$$\mathbf{mL^4} \xrightarrow{(\cdot)_\mathbf{0}} \mathbf{mL^4_0}$$
$$\mathbf{mL^4} \xleftarrow{(\cdot)_\mathbf{1}} \mathbf{mL^4_0}$$

We first define them on formulas; this is done by induction on the argument formula:

$$\begin{array}{rcl} X_\mathbf{0} & = & 0X \\ (X^\perp)_\mathbf{0} & = & 0X^\perp \\ (\S A)_\mathbf{0} & = & 1 \cdot A_\mathbf{0} \end{array}$$

and $(\cdot)_\mathbf{0}$ commutes with the other connectives, e.g.

$$(A \otimes B)_\mathbf{0} \quad = \quad A_\mathbf{0} \otimes B_\mathbf{0}$$

Similarly,

$$\begin{array}{rcl} (pX)_\mathbf{1} & = & \S^p X \\ (pX^\perp)_\mathbf{1} & = & \S^p X^\perp \end{array}$$

and $(\cdot)_\mathbf{1}$ commutes with all connectives, e.g.

$$(A \otimes B)_\mathbf{1} \quad = \quad A_\mathbf{1} \otimes B_\mathbf{1}$$

Observe that $(\cdot)_\mathbf{0} \circ (\cdot)_\mathbf{1}$ is the identity on Form_0, while $(\cdot)_\mathbf{1} \circ (\cdot)_\mathbf{0}$ sends $A \in \mathit{Form}$ to the "canonical" representative of its equivalence class, i.e., the formula with all § pushed to the atoms.

We shall now define how $(\cdot)_\mathbf{0}$ and $(\cdot)_\mathbf{1}$ behave on proofs. Let π be an $\mathbf{mL^4}$ proof net. We say that a link l is *below* an edge e or, equivalently, that e is *above* l if in π there is a directed path from e to the premise of l. We then define $\pi_\mathbf{0}$ as follows:

- replace each axiom of conclusions $A^\perp, A$ by an axiom of conclusions $q \cdot A^\perp$, $p \cdot A$ where q (resp. p) is the number of paragraph links below $A^\perp$ (resp. A) in π;
- remove paragraph links, and label each edge according to the relabeling of the axioms.

Informally speaking, $\pi_\mathbf{0}$ is obtained from π by pushing paragraph connectives upwards in the proof net, and "absorbing" them into the axioms. We have:

Proposition 28. *Let π be an* $\mathbf{mL^4}$ *proof net of conclusions* Γ*; then* $\pi_{\mathbf{0}}$ *is an* $\mathbf{mL_0^4}$ *proof net of conclusions* $\Gamma_{\mathbf{0}}$*.*

Proof. Since π is an $\mathbf{mL^4}$ proof net it can be given an indexing I. To define an indexing I_0 on $\pi_{\mathbf{0}}$ it is sufficient to define it on the conclusions of axioms. Each axiom link a' in $\pi_{\mathbf{0}}$ has conclusions e'_1, e'_2 with respective types of the form $q \cdot A^{\perp}, p \cdot A$ and comes from an axiom a of π of conclusions e_1, e_2 with respective types $A^{\perp}$, A. W.l.o.g. we can assume $q \geq p$. Let $i = I(e_1) = I(e_2)$. Then set $I_0(e'_1) = i - q$, $I_0(e'_2) = i - p$. Note that we have $q \cdot A^{\perp} = (q-p) \cdot (p \cdot A)^{\perp}$ and $I_0(e'_2) = I_0(e'_1) + (q-p)$, so I_0 satisfies the condition on axioms, and is indeed an indexing. One can verify that $\pi_{\mathbf{0}}$ is well-typed; a fundamental remark for this is that $(\cdot)_{\mathbf{0}}$ preserves duality, i.e., $(A^{\perp})_{\mathbf{0}} = A_{\mathbf{0}}^{\perp}$. To conclude, observe that the structure of π and $\pi_{\mathbf{0}}$ are basically identical: the only difference is the absence of paragraph links in $\pi_{\mathbf{0}}$. But these are completely transparent to both the connected-acyclic condition (Definition 4) and the Stratification and Lightness conditions (Definition 16). Hence, since π satisfies these conditions, so does $\pi_{\mathbf{0}}$, which means that this latter is an $\mathbf{mL_0^4}$ proof net. □

The translation $(\cdot)_{\mathbf{1}}$ requires a few preliminary definitions:

Definition 34. Let $A \in$ *Form* and $p \in \mathbb{N}$; the net R_A^p is defined as follows:

- let S_A be the $\mathbf{mL^4}$ proof net of conclusions $A^{\perp}, A$, representing the η-expansion of the axiom of conclusions $A^{\perp}, A$;
- R_A^p is obtained from S_A by replacing each axiom link of conclusion $X^{\perp}, X$, where $X^{\perp}$ is the type of the edge above the conclusion $A^{\perp}$, by the same link followed by p paragraph links below $X^{\perp}$.

In the following, a *weak* $\mathbf{mL^4}$ proof net is a **meLL** proof net satisfying the Stratification and Lightness conditions (Definition 16) and admitting a weak indexing.

Lemma 29. *For all $A \in$ Form and $p \in \mathbb{N}$, R_A^p is a weak* $\mathbf{mL^4}$ *proof net.*

Proof. A straightforward induction on A. □

Let now π be an $\mathbf{mL_0^4}$ proof net of conclusions Γ. Then, $\pi_{\mathbf{1}}$ is obtained by replacing each axiom of conclusions $p \cdot A^{\perp}, A$ in π by R_A^p, and typing the rest of the edges accordingly.

Proposition 30. *Let π be an* $\mathbf{mL_0^4}$ *proof net of conclusions* Γ*; then* $\pi_{\mathbf{1}}$ *is an* $\mathbf{mL^4}$ *proof net of conclusions* $\Gamma_{\mathbf{1}}$*.*

Proof. A more or less obvious corollary of Lemma 29. □

Observe that $(\cdot)_{\mathbf{0}} \circ (\cdot)_{\mathbf{1}}$ does not act exactly as identity on $\mathbf{mL_0^4}$ proof nets, but performs an η-expansion. On the other hand, $(\cdot)_{\mathbf{1}} \circ (\cdot)_{\mathbf{0}}$ behaves just like its counterpart on *Form*: given π, it gives the isomorphic proof net in which all paragraph links have been pushed to the axioms.

Both $\mathbf{mL_0^4}$ and $\mathbf{mL^4}$ can be embedded in **meLL**. For the first system, there is clearly a forgetful embedding U which simply erases the integers from atoms, both in formulas and proofs: $U(pX) = X$, $U(pX^\perp) = X^\perp$, and U commutes with all connectives. The second system is by definition a subsystem of **meLL**, so the embedding would be trivial (the identity!); however, we are interested here in the following translation $(.)^-$:

- given a formula $A \in \mathit{Form}$, A^- is A in which all § have been removed;
- given an $\mathbf{mL^4}$ proof net π, π^- is π in which all paragraph links have been removed, and types have been changed accordingly.

Clearly, both U and $(.)^-$ embed resp. $\mathbf{mL_0^4}$ and $\mathbf{mL^4}$ in "standard" **meLL**, i.e., multiplicative exponential linear logic *without the paragraph modality* (actually, the embedding takes place in **mELL**). These two embeddings preserve cut-elimination:

Lemma 31. *Let π be an $\mathbf{mL_0^4}$ proof net. Then, $\pi \to \pi'$ iff $U(\pi) \to U(\pi')$.*

Proof. Simply observe that the untyped structure of π and $U(\pi)$ is identical, and cuts are reduced regardless of types (except quantifier cuts, but these are easily seen to be reciprocally simulated in one step). □

Lemma 32. *Let π be an $\mathbf{mL^4}$ proof net. Then, $\pi \to \pi'$ iff $\pi^- \to^* (\pi')^-$ in at most one step.*

Proof. If $\pi \to \pi'$, and the step applied is not a paragraph step, then clearly $\pi^- \to (\pi')^-$. If it is a paragraph step, then it easy to see that $(\pi')^- = \pi^-$. For the converse, one reduction step in π^- is always simulated by exactly one reduction step in π. □

An important corollary of Lemma 31 is the confluence and strong normalization of $\mathbf{mL_0^4}$, which follows from the similar properties of **meLL** (Girard, 1987).

We also have a useful result relating the two embeddings:

Lemma 33. *Let π be an $\mathbf{mL^4}$ proof net. Then, $U(\pi_\mathbf{0}) = \pi^-$.*

Proof. As noted above, the translation $(\cdot)_\mathbf{0}$ pushes paragraph links to the axioms, and then "absorbs" them into the formulas; then U forgets the annotations concerning paragraphs. But this amounts to simply removing the § modality from both π and its formulas. □

In the sequel, we denote by $\to_\eta$ the application of one η-expansion step to an $\mathbf{mL_0^4}$ proof net. One η-expansion step replaces a non-atomic axiom of conclusions $p \cdot C, C^\perp$ with axioms introducing the immediate subformulas of C. Figures 19 and 20 give the definition for the cases $C = A \otimes B$ and $C = ?A$; the other cases are treated similarly, as the reader may expect.

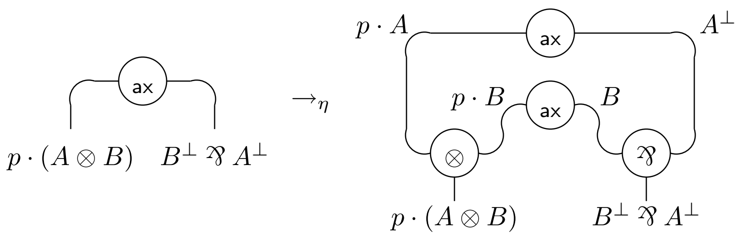

Figure 19: Multiplicative η-expansion step.

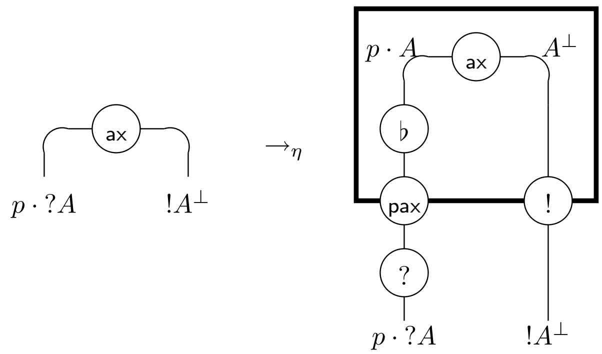

Figure 20: Exponential η-expansion step.

Lemma 34. *Let π be an $\mathbf{mL_0^4}$ proof net such that $\pi \to_\eta \pi_1 \to \pi_2$. Then, there exist π_1', π_2' such that $\pi \to \pi_1' \to_\eta^* \pi_2''$ and π_2' is β-equivalent to π_2, i.e., they have a common reduct through cut-elimination.*

Proof. If the cut-elimination step applied in $\pi_1 \to \pi_2$ is "far" from the axioms, then the result is obvious. We can thus concentrate on the *critical pairs*, i.e., the situations in which the axiom which is expanded in going from π to π_1 is involved in a cut, and (the residue of) this cut is exactly the one reduced in going from π_1 to π_2. We check the only interesting case, leaving the others to the reader. Suppose that π contains an axiom a of conclusions $p \cdot ?A, !A^\perp$, and the conclusion of type $!A^\perp$ is the premise of a cut c, whose other premise is the conclusion of a why not link w. We shall assume $p = 0$; the general case is entirely similar. The η-expansion replaces a with a box containing a pre-net ι consisting of an axiom of conclusions $A, A^\perp$ and a flat link just below A. The cut-elimination step makes n copies of ι, and cuts them to the appropriate links. If we reduce these n cuts, we obtain a proof net that we call π_2'. Now, if we take π and reduce c right away, it is immediate to see that we obtain exactly π_2', and η-expansion is not even needed. □

If π is a **meLL** or $\mathbf{mL_0^4}$ proof net, we denote by $\mathrm{NF}(\pi)$ its normal form, and by $\xrightarrow{\mathrm{NF}}$ reduction to the normal form. Then, we have:

Lemma 35. *The following diagrams commute:*

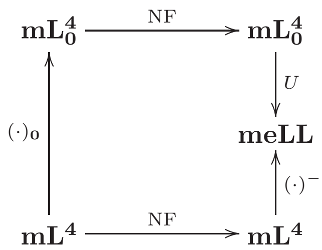

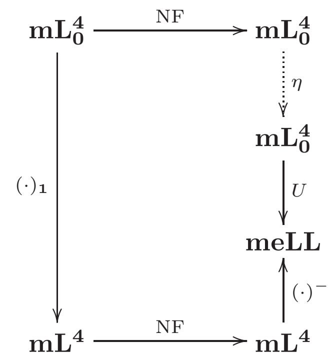

where the dotted arrow means that one may need to η-expand some axioms to close the second diagram.

Proof. For the first diagram, it is enough to prove that the three subdiagrams of the following diagram commute:

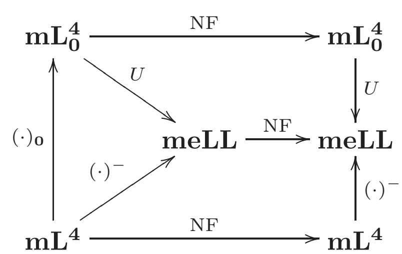

These are consequences of Lemmas 31, 32, and 33. For what concerns the second diagram, it is enough to prove that the three subdiagrams of the following diagram commute:

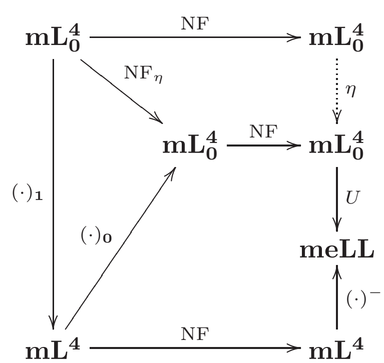

where NF_η is the function associating with a proof net π its η-expanded form, i.e., the proof net obtained by η-expanding all axioms of π until only atomic axioms are left. Now, the commutation of the triangle on the left is simply the remark we made after Proposition 30, while the bottom subdiagram is nothing but the first diagram of this lemma. Hence, all that is left to prove is the commutation of the top subdiagram. This is a consequence of Lemma 34. In fact, let π be an $\mathbf{mL_0^4}$ proof net, and let $\pi' = \mathrm{NF}_\eta(\pi)$ and $\pi'' = \mathrm{NF}(\pi')$. By definition, we have $\pi \to_\eta^* \pi' \to^* \pi''$. We shall prove by induction on the length of the reduction $\pi \to_\eta^* \pi'$ that $\mathrm{NF}(\pi) \to_\eta^* \pi''$. If $\pi' = \pi$, then clearly $\mathrm{NF}(\pi) = \pi''$. If $\pi \to_\eta^* \pi_1 \to_\eta \pi'$, then, using Lemma 34, by a further induction on the length of the reduction $\pi' \to^* \pi''$ we can prove that $\pi_1 \to^* \pi_2 \to_\eta^* \pi_3$, and π_3 is β-equivalent to π''. But π'' is a normal form, so $\pi_2 \to_\eta^* \pi''$. Composing the reductions, we have $\pi \to_\eta^* \pi_1 \to^* \pi_2 \to_\eta^* \pi''$. Now the induction hypothesis applies, because the reduction $\pi \to_\eta^* \pi_1$ is strictly shorter than $\pi \to_\eta^* \pi'$. This gives us $\mathrm{NF}(\pi) \to_\eta^* \pi_2 \to_\eta^* \pi''$, as desired. □

Note that from the first diagram and Lemma 33 we can infer that, for every $\mathbf{mL^4}$ proof net π, $U(\mathrm{NF}(\pi_\mathbf{0})) = U((\mathrm{NF}(\pi))_\mathbf{0})$. However, U is not injective, so we cannot conclude that the translation $(\cdot)_\mathbf{0}$ commutes with reduction. The situation for the translation $(\cdot)_\mathbf{1}$ is even worse: $(\mathrm{NF}(\pi_\mathbf{1}))^- = ((\mathrm{NF}(\pi))_\mathbf{1})^-$ holds only up to η-equivalence.

We now proceed to argument how $\mathbf{mL_0^4}$ characterizes $\mathbf{FP}$ (Theorem 37). First of all, we define the $\mathbf{mL_0^4}$ type of finite binary strings as follows:

$$\mathbf{S_0} = \forall X.(?(0X^\perp \otimes 0X) \parr ?(0X^\perp \otimes 0X) \parr (1X^\perp \parr 1X)).$$

The reader can check that $\mathbf{S_0} = (\mathbf{S_P})_\mathbf{0} = (\mathbf{S'_P})_\mathbf{0}$, where $\mathbf{S_P}$ and $\mathbf{S'_P}$ are the two isomorphic types that can be used for representing binary strings in $\mathbf{mL^4}$ (cf. Sect. 3.4). Hence, by Proposition 28, if $\overline{x}$ is the $\mathbf{mL^4}$ proof net of conclusion $\mathbf{S_P}$ (or $\mathbf{S'_P}$) representing the string x, the same string can be represented in $\mathbf{mL_0^4}$ by the proof net $(\overline{x})_\mathbf{0}$.

Lemma 36. *Let ξ, ξ' be two cut-free proof nets of resp. $\mathbf{mL^4}$ and $\mathbf{mL_0^4}$, of resp. conclusion $\S^p\mathbf{S_P}$ (or $((\S^p\mathbf{S'_P})_\mathbf{0})_\mathbf{1}$) and $p\cdot\mathbf{S_0}$, such that $U(\xi') = \xi^-$. Then, ξ and ξ' represent the same binary string.*

Proof. The fact that $U(\xi') = \xi^-$ implies that ξ and ξ' have the same untyped structure modulo the presence of paragraph links in ξ; then the lemma is a consequence of the types of the two proof nets, and of the fact that they are cut-free. □

Given a non-negative integer p and an $\mathbf{mL_0^4}$ proof net π not containing existential links, we denote by $p\cdot\pi$ the proof net obtained by replacing all atoms A appearing in the types of π with $p\cdot A$. It is easy to check that if π is of conclusions Γ, then $p\cdot\pi$ is a well-typed $\mathbf{mL_0^4}$ proof net of conclusions $p\cdot\Gamma$. Moreover, if π contains only atomic axioms, then so does $p\cdot\pi$.

In the following, if φ is a proof net of conclusions $A^\perp, B$ and ξ a proof net of conclusion A, we use the notation $\varphi(\xi)$ as introduced in Sect. 3.4. Observe that both $(\cdot)_\mathbf{0}$ and $(\cdot)_\mathbf{1}$ are modular with respect to this notation, i.e., $(\varphi(\xi))_\mathbf{0} = \varphi_\mathbf{0}(\xi_\mathbf{0})$ and $(\varphi(\xi))_\mathbf{1} = \varphi_\mathbf{1}(\xi_\mathbf{1})$.

Definition 35 (Representation). Let $f : \{0,1\}^* \to \{0,1\}^*$. We say that f is *representable* in $\mathbf{mL^4_0}$ if there exists $p \in \mathbb{N}$ and an $\mathbf{mL^4_0}$ proof net φ of conclusions ${\mathbf{S_0}}^\perp, p \cdot \mathbf{S_0}$ such that, whenever ξ is a proof net of conclusion $\mathbf{S_0}$ representing the string x, we have $f(x) = y$ iff $\mathrm{NF}(\varphi(x)) = p \cdot \upsilon$, where υ represents y.

Theorem 37. *Let $f : \{0,1\}^* \to \{0,1\}^*$. Then, $f \in \mathbf{FP}$ iff f is representable in $\mathbf{mL^4_0}$.*

Proof. Let us start with the completeness of $\mathbf{mL^4_0}$ w.r.t. $\mathbf{FP}$. Let $f \in \mathbf{FP}$. By Theorem 26 there exist $p \in \mathbb{N}$ and an $\mathbf{mL^4}$ proof net φ such that, for all $x \in \{0,1\}^*$, $f(x) = y$ iff $\mathrm{NF}(\varphi(\xi)) = \upsilon$, where υ is the representation of y with p paragraph links added to its conclusion. Let $\upsilon' = \mathrm{NF}((\varphi(\xi))_\mathbf{0}) = \mathrm{NF}(\varphi_\mathbf{0}(\xi_\mathbf{0}))$. By the first diagram of Lemma 35, $\upsilon^- = U(\upsilon')$, so by Lemma 36 $\varphi_\mathbf{0}$ represents f.

For what concerns soundness, let φ be an $\mathbf{mL^4_0}$ proof net of conclusions ${\mathbf{S_0}}^\perp, p \cdot \mathbf{S_0}$ representing the function f. For all $x \in \{0,1\}$, if ξ is the $\mathbf{mL^4_0}$ representation of x, we have $f(x) = y$ iff $\mathrm{NF}(\varphi(\xi)) = \upsilon'$, where $\upsilon' = p \cdot \upsilon$ and υ represents y. Now, observe that the representations of binary strings are all η-expanded, which means that $\upsilon' \to^*_\eta \upsilon''$ implies $\upsilon'' = \upsilon'$. Hence, in the second diagram of Lemma 35 we can replace the dotted arrow with the identity, and obtain $U(\upsilon') = (\mathrm{NF}((\varphi(\xi))_\mathbf{1}))^- = (\mathrm{NF}(\varphi_\mathbf{1}(\xi_\mathbf{1})))^-$. The proof net $\mathrm{NF}(\varphi_\mathbf{1}(\xi_\mathbf{1}))$ is a normal form of type $(p \cdot \mathbf{S_0})_\mathbf{1} = ((\S^p \mathbf{S'_P})_\mathbf{0})_\mathbf{1}$, so Lemma 36 applies, and $\varphi_\mathbf{1}$ represents f in $\mathbf{mL^4}$ according to the alternative definition which uses the type $\mathbf{S'_P}$ for binary strings. But, as we pointed out in Sect. 3.4, Theorem 26 is still valid in this case, so $f \in \mathbf{FP}$. □

4.1. Sequent calculus for $\mathbf{mL^4_0}$

It may be interesting to consider a sequent calculus formulation of $\mathbf{mL^4_0}$, especially if one seeks to derive from it a type assignment system for the λ-calculus, to be used to infer complexity properties about λ-terms (in the style, for example, of **DLAL** (Baillot and Terui, 2004)). Starting from the 2-sequent calculus for $\mathbf{mL^4}$ (Sect. 2.3), we end up with the rules given in Table 4 (daimon and mix are again omitted, because identical to Table 1). As expected, weak $\mathbf{mL^4_0}$ proof nets correspond to derivations in this calculus, and $\mathbf{mL^4_0}$ proof nets to proper derivations. Observe the complete absence of a paragraph rule.

5. Concluding Remarks and Further Work

We may perhaps summarize the fundamental contribution of the present work in one sentence: *in linear-logical characterizations of complexity classes, exponential boxes and stratification levels are two different things.* From this fact, we have seen how one can define an elementary system extending **ELL**, and a polynomial system extending **LLL**. The main novelty of this latter, which is in direct connection with the above fact, is the absence of §-boxes. This implies that the paragraph modality commutes with all connectives; these commutations can be exploited to devise a polynomial system with a simpler class of formulas and fewer typing rules, which may be of interest for type

$$\frac{}{\vdash p \cdot A^{\perp i}, A^{i+p}}\ \text{Axiom} \qquad \frac{\vdash \Gamma, A^i \quad \vdash \Delta, A^{\perp i}}{\vdash \Gamma, \Delta}\ \text{Cut}$$

$$\frac{\vdash \Gamma, A^i \quad \vdash \Delta, B^i}{\vdash \Gamma, \Delta, A \otimes B^i}\ \text{Tensor} \qquad \frac{\vdash \Gamma, A^i, B^i}{\vdash \Gamma, A \parr B^i}\ \text{Par}$$

$$\frac{\vdash \Gamma, A^i}{\vdash \Gamma, \forall X.A^i}\ \text{For all } (X \text{ not free in } \Gamma) \qquad \frac{\vdash \Gamma, A\{B/X\}^i}{\vdash \Gamma, \exists X.A^i}\ \text{Exists}$$

$$\frac{\vdash B^{j+1}, A^{i+1}}{\vdash ?B^j, !A^i}\ \text{Light promotion} \qquad \frac{\vdash \Gamma, A^{i+1}}{\vdash \Gamma, ?A^i}\ \text{Dereliction}$$

$$\frac{\vdash \Gamma}{\vdash \Gamma, ?A^i}\ \text{Weakening} \qquad \frac{\vdash \Gamma, ?A^i, ?A^i}{\vdash \Gamma, ?A^i}\ \text{Contraction}$$

Table 4: Rules for $\mathbf{mL^4_0}$ 2-sequent calculus. Daimon and mix are omitted.

assignment purposes. This is probably the most obvious direction of further research given by this work; in the sequel, we discuss other remarks and open questions.

Indexes and tiers. We already mentioned in the introduction how our form of stratification reminds of ramification, a technique devised by Leivant and Marion (1993) to characterize complexity classes within the λ-calculus. Ramification is enforced by so-called *tiers*, which are integers assigned to subterms of a λ-term, in close analogy with our indexes. However, we have not been able so far to understand the formal relationship between the two, and we suspect this may be an interesting subject for further work.

Intensionality. Concretely, the fact that $\mathbf{mL^3}$ and $\mathbf{mL^4}$ extend resp. **mELL** and **mLLL** means that the first two systems have "more proofs" that the latter two. Through the Curry-Howard looking glass, this means that $\mathbf{mL^3}$ and $\mathbf{mL^4}$ are intensionally more expressive than Girard's corresponding systems, i.e., they admit "more programs". How many and which is still not clear though: we do have examples of λ-terms which are not typable in multiplicative **ELL** and yet are typable in $\mathbf{mL^3}$ (or even in $\mathbf{mL^4}$!), but none of these corresponds to any "interesting" algorithm. So the question of whether our systems actually improve on the intensionality of **ELL** and **LLL** remains open.

Naive set theory. Proposition 14 states that, if we take an untyped $\mathbf{mL^3}$ proof net and start reducing its cuts, after a finite number of steps we either reach a cut-free form or a deadlock, i.e., a proof net whose all cuts are ill-formed. Now, the preservation of typing under reduction guarantees that, if the starting proof net is typed, then the latter case never happens; hence, $\mathbf{mL^3}$ satisfies cut-elimination.

This sharply contrasts with the situation one has in **meLL**: weak normalization blatantly fails in untyped **meLL** proof nets (the pure λ-calculus can be translated in the system), and the proof of cut-elimination in the typed case is highly complex, because of the presence of second order quantification. Indeed, cut-elimination of second-order **meLL** proof nets is known to be equivalent to the consistency of $\mathbf{PA}_2$ (Girard, 1987), for which no inductive proof has ever been given (in other words, no-one knows what ordinal should replace ω^ω in a proof like that of Proposition 14).

Following Girard (1998) and Terui (2004), one can build two naive set theories out of $\mathbf{mL^3}$ and $\mathbf{mL^4}$, which can still be proved to be consistent, i.e., to satisfy cut-elimination. In spite of their low logical complexity (as in the proof of Proposition 14, the consistency of these theories can be proved by an induction up to ω^ω), these set theories are particularly interesting because they are conservative extensions of the set theories based on elementary and light linear logic: they still use unrestricted comprehension, and thus allow arbitrary fixpoints of formulas, but they have more flexible logical principles, i.e., they admit more proofs. Asking how many more is of course another way of posing the above question about intensionality.

Additives. The additive connectives of linear logic (& and $\oplus$) have been excluded from this work; this is only a convenient choice, justified by the fact that some proofs (in particular those of Proposition 14 and Theorems 17 and 24) become simpler. There is no technical problem in adding them to our systems, thus defining what we would call $\mathbf{L^3}$ and $\mathbf{L^4}$, which we still believe to exactly characterize resp. elementary and deterministic polytime computation.

There is however one point worth mentioning. The most natural definition of $\mathbf{L^4}$ extends the commutation of the paragraph modality to additive connectives as well; in particular, the isomorphism $\S(A\oplus B) \cong \S A\oplus\S B$ holds. Girard (2007) has a nice argument against this being possible in **LLL**, which goes as follows. For the sake of contradiction, suppose we can prove $\S(A\oplus B) \multimap \S A\oplus\S B$ in **LLL**, and hence $\S^p(A\oplus B) \multimap \S^p A\oplus\S^p B$ for any $p \in \mathbb{N}$. Booleans can be easily encoded using the type $V_1 \oplus V_2$, where V_1 and V_2 are two formulas admitting exactly one proof (for example $V_1 = V_2 = \forall X.(X^\perp \parr X)$). By similar definitions and arguments to those of Definition 29 and Theorem 26, any language in **P** can be represented by an **LLL** proof net φ of conclusions $\mathbf{S}_\mathbf{P}^\perp, \S^p(V_1 \oplus V_2)$ for a suitable value of p depending on the language itself. Now, using the commutation of the paragraph modality, we can transform φ into a proof net φ' of $\mathbf{S}_\mathbf{P}^\perp, \S^p V_1 \oplus \S^p V_2$. If we want to know whether the string x belongs to our language or not, we may simply take the proof net ξ representing x and normalize $\varphi'(\xi)$ (we are using the notation of Sect. 3.4), which has conclusion $\S^p V_1 \oplus \S^p V_2$. Observe that the main connective of this formula is $\oplus$, hence the plus link introducing it must be at depth zero, i.e., it is not contained in any exponential box. Observe also that the result of the computation is known as soon as the nature of this link is known, i.e., whether $\S^p V_1 \oplus \S^p V_2$ is introduced from $\S^p V_1$ or $\S^p V_2$. But then, to have our answer, it is enough to stop the "round-by-round" cut-elimination procedure right after depth zero. In **LLL**, normalizing just one depth is linear in the size of the proof net, so we can solve any deterministic polytime problem in linear time, which is obviously false.

This argument however does not apply to $\mathbf{L^4}$ because of the crucial difference between *depth* and *level*. A language in **P** may as well be represented in $\mathbf{L^4}$ by a proof net φ' of conclusions $\mathbf{S}_{\mathbf{P}}^{\perp}, \S^p V_1 \oplus \S^p V_2$, and it remains true that it is enough to normalize depth zero of $\varphi'(\xi)$ to know whether the string represented by ξ is in the language or not; however, the "round-by-round" cut-elimination procedure for $\mathbf{L^4}$ goes *level by level*, and depth zero may contain arbitrary many levels (in this case, p levels is a good guess). Hence, normalizing just one depth may take a number of steps far from being linear in the size of the proof net, as we already showed in the example of Fig. 16.

Denotational semantics. Recently, Laurent and Tortora de Falco (2006) have proposed a denotational semantics for Girard's **ELL** and Lafont's **SLL**. Together with stratified coherence spaces (Baillot, 2004), these are very interesting attempts at giving a completely semantic definition of complexity classes.

The present paper offers a new and arguably novel starting point in this perspective: ongoing work with Tortora de Falco seems to be yielding promising results in the direction of finding a denotational semantics for $\mathbf{mL^3}$, which, like that for **ELL**, is still based on the relational semantics for linear logic, but is of a rather different nature.

References

Asperti, A., Roversi, L., 2002. Intuitionistic light affine logic. ACM Transactions on Computational Logic 3 (1), 1–39.

Atassi, V., Baillot, P., Terui, K., 2007. Verification of Ptime Reducibility for system F Terms: Type Inference in Dual Light Affine Logic. Logical Methods in Computer ScienceTo appear. 32 pages.

Baillot, P., 2004. Stratified coherence spaces: a denotational semantics for Light Linear Logic. Theoretical Computer Science 318 (1-2), 29–55.

Baillot, P., Terui, K., 2004. Light types for polynomial time computation in lambda-calculus. In: Proceedings of LICS'04. IEEE Computer Society Press, pp. 266–275.

Bellantoni, S., Cook, S., 1992. New recursion-theoretic characterization of the polytime functions. Computational Complexity 2, 97–110.

Coppola, P., Martini, S., 2006. Optimizing optimal reduction. a type inference algorithm for elementary affine logic. ACM Transactions on Computational Logic 7 (2), 219–260.

Curry, H., Feys, R., 1958. Combinatory Logic. North Holland.

Danos, V., Joinet, J.-B., 2003. Linear logic & elementary time. Inform. and Comput. 183, 123–137.

Danos, V., Regnier, L., 1995. Proof nets and the Hilbert space. In: Girard, J.-Y., Lafont, Y., Regnier, L. (Eds.), Advances in Linear Logic. Cambridge University Press, pp. 307–328.

Gaboardi, M., Marion, J.-Y., Rocca, S. R. D., 2008. A Logical Account of Pspace. In: Proceedings of POPL 2008. ACM Press, to appear.

Gaboardi, M., Ronchi Della Rocca, S., 2007. A soft type assignment system for *lambda*-calculus. In: Proceedings of Computer Science Logic (CSL'07). Vol. 4646 of Lecture Notes in Computer Science. Springer, pp. 253–267.

Gentzen, G., 1934. Investigations into logical deductions. In: Szabo, M. E. (Ed.), The Collected Papers of Gerhard Gentzen. North-Holland, Amsterdam, pp. 68–131.

Girard, J.-Y., 1987. Linear logic. Theoretical Computer Science 50 (1), 1–102.

Girard, J.-Y., 1996. Proof-nets: the parallel syntax for proof-theory. In: Agliano, P., Ursini, A. (Eds.), Logic and Algebra. Marcel Dekker, pp. 97–124.

Girard, J.-Y., 1998. Light linear logic. Inform. and Comput. 14 (3), 175–204.

Girard, J.-Y., 2007. Le Point Aveugle. Hermann.

Girard, J.-Y., Scedrov, A., Scott, P., 1992. Bounded linear logic: A modular approach to polynomial time computability. Theoretical Computer Science 97, 1–66.

Grishin, V. N., 1982. Predicate and set-theoretic calculi based on a logic without contractions. Math. of USSR, Izvestiya 18 (1), 41–59.

Guerrini, S., Martini, S., Masini, A., 1998. An Analysis of (Linear) Exponentials Based on Extended Sequents. Logic Journal of the IGPL 6 (5), 735–753.

Kanovich, M., Okada, M., Scedrov, A., 2003. Phase semantics for light linear logic. Theoret. Comput. Sci. 294 (3), 525–549.

Lafont, Y., 2004. Soft linear logic and polynomial time. Theoret. Comput. Sci. 318 (1–2), 163–180.

Laurent, O., Tortora de Falco, L., 2006. Obsessional cliques: a semantic characterization of bounded time complexity. In: Proceedings of LICS'06. IEEE Computer Society Press, pp. 179–188.

Leivant, D., 1994. Predicative recurrence and computational complexity I: word recurrence and poly-time. In: Feasible Mathematics II. Birkhauser, pp. 320–343.

Leivant, D., Marion, J.-Y., 1993. Lambda-calculus characterisations of polytime. Fundamenta Informaticae 19, 167–184.

Mairson, H., Terui, K., 2003. On the computational complexity of cut-elimination in linear logic. In: Proceedings of ICTCS 2003. Vol. 2841 of LNCS. Springer-Verlag, pp. 23–36.

Marion, J.-Y., 2007. Predicative analysis of feasibility and diagonalization. In: Rocca, S. R. D. (Ed.), Proceedings of TLCA 2007. Vol. 4583 of Lecture Notes in Computer Science. Springer, pp. 290–304.

Masini, A., 1992. 2-sequent calculus: A proof-theory of modalities. Annals of Pure and Applied Logic 58, 229–246.

Mazza, D., 2006. Linear logic and polynomial time. Mathematical Structures in Computer Science 16 (6), 947–988.

Murawski, A. S., Ong, C.-H. L., 2004. On an interpretation of safe recursion in light affine logic. Theor. Comput. Sci. 318 (1-2), 197–223.

Roversi, L., 1999. A P-time completeness proof for light logics. In: Flum, J., Rodríguez-Artalejo, M. (Eds.), Proceedings of CSL 99. Vol. 1683 of LNCS. Springer, pp. 469–483.

Schöpp, U., 2007. Stratified bounded affine logic for logarithmic space. In: Proceedings of Symposium on Logic in Computer Science (LICS'07). IEEE Computer Society, pp. 411–420.

Terui, K., 2004. Light affine set theory: a naive set theory of polynomial time. Studia Logica 77, 9–40.

Tortora de Falco, L., 2003. Additives of linear logic and normalization – Part I: a (restricted) Church-Rosser property. Theoret. Comput. Sci. 294 (3), 489–524.