Proof Nets for Unit-free Multiplicative-Additive Linear Logic

DOMINIC J. D. HUGHES and ROB J. VAN GLABBEEK
Stanford University

A cornerstone of the theory of proof nets for unit-free multiplicative linear logic (MLL) is the abstract representation of cut-free proofs modulo inessential rule commutation. The only known extension to additives, based on monomial weights, fails to preserve this key feature: a host of cut-free monomial proof nets can correspond to the same cut-free proof. Thus the problem of finding a satisfactory notion of proof net for unit-free multiplicative-additive linear logic (MALL) has remained open since the inception of linear logic in 1986. We present a new definition of MALL proof net which remains faithful to the cornerstone of the MLL theory.

Key Words and Phrases: Linear logic, proof nets, additives, cut elimination

Contents

An extended abstract of this paper appeared in Proc. 18th Annual IEEE Symp. on *Logic in Computer Science* (LICS 2003), Ottawa, Canada, June 2003, pp. 1-10. A preliminary version of some of the material was presented at *Linear Logic 2002*, Copenhagen, Denmark, July 2002.
When this paper was finalised, the second author was employed at the National ICT Australia, and at the University of New South Wales, Sydney, Australia.

1 Introduction

The beautiful theory of proof nets for unit-free multiplicative linear logic (MLL) appeared alongside the introduction of linear logic [Girard 1987]. A proof net is an abstract representation of a proof: the translation of cut-free proofs into proof nets identifies proofs modulo inessential rule commutation. The identifications have since been verified as canonical from a semantic perspective, with numerous full completeness results for MLL, e.g. [Abramsky and Jagadeesan 1994; Hyland and Ong 1993; Loader 1994; Tan 1997; Blute and Scott 1996; Devarajan, Hughes, Plotkin and Pratt 1999]. Furthermore, the identifications correspond to coherences of free star-autonomous categories [Blute, Cockett, Seely and Trimble 1996].

The problem of finding a satisfactory extension of the theory of proof nets to unit-free multiplicative-additive linear logic (MALL) has remained open since the inception of linear logic [Girard 1987]. Progress towards a solution was made by Girard [1996] with a notion of MALL proof net based on *monomial weights*. Unfortunately, monomial proof nets failed to extend the MLL theory faithfully: a single cut-free proof may correspond to a host of monomial proof nets, and there is no natural map from cut-free proofs onto monomial proof nets. To quote Girard, monomial proof nets are "far from being absolutely satisfactory" [1996]. We illustrate the problems in detail in Section 6.1.

In this paper we propose a new notion of MALL proof net (Section 4) which adheres faithfully to the original MLL theory: we provide a simple function from cut-free proofs to cut-free proof nets, yielding the sought-after abstract representations of cut-free proofs modulo inessential commutation of rules. We define a cut-free proof net on a sequent Γ as a set of linkings on Γ satisfying a geometric correctness criterion[1], and prove that a set of linkings is the translation of a proof if and only if it is a proof net (Theorem 4.18, the cut-free *Sequentialisation Theorem*). The definition of proof net is pleasingly succinct, taking only 11 lines (page 14). The reader can glean an impression of our approach by perusing Figure 1.

In Section 5 we extend our proof nets with cuts, and present a notion of cut elimination (and turbo cut elimination). Cut elimination is simply defined, strongly normalising, and yields a category of cut-free proof nets which is semi (i.e. unit-free) star-autonomous, with products and coproducts. For an impressionistic overview see Figures 2 (cut), 3 (cut elimination), and 4 (composition). After extending to cut, the definition of proof net remains succinct: see the box on page 40. As with Girard's monomial proof nets, in the presence of cuts multiple proof nets may correspond to the same proof. However, from a semantic point of view (viz. full completeness) the provision of abstract representations of MALL proofs modulo rule commutation is crucial only in the cut-free setting.

A crisp notion of cut-free MALL proof net is fully motivated from a proof-theoretic perspective alone. However, just as MLL has blossomed through numerous fully complete semantics via cut-free MLL proof nets, we hope that the new definition of cut-free proof net presented here will lead to a similar blossoming of MALL. Since cut-free monomial proof nets for MALL are unsatisfactory for the reasons outlined earlier (detailed in Section 6.1), any MALL full completeness result based on them (e.g. the concurrent games model [Abramsky and Melliès 1999]

[1] Relaxing the criterion slightly yields a notion of proof net for MALL with mix (Section 4.9).

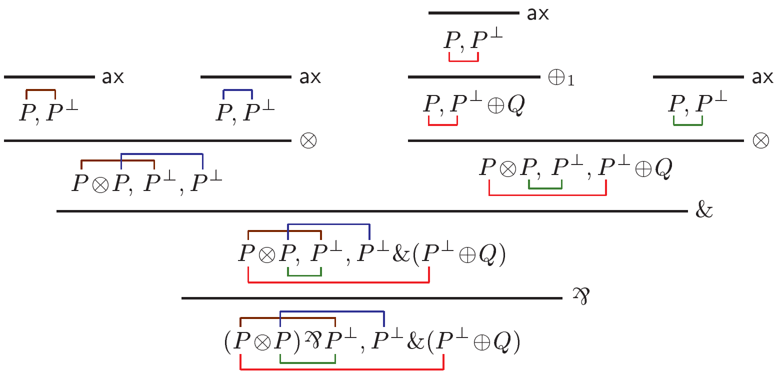

Fig. 1. Example of the inductive translation of a cut-free MALL proof into one of our cut-free proof nets. The concluding proof net has two linkings, one drawn above the sequent, the other below. Each contains two axiom links. The proof nets further up in the derivation have one or two linkings, correspondingly above and/or below the sequent. Had we switched the order of the right-hand tensor rule and the plus rule, we would have obtained exactly the same pair of linkings; thus we identify cut-free proofs modulo a commutation of rules.

---

or the hypercoherence model [Blute, Hamano and Scott 2005]) suffers accordingly, particularly with regard to faithfulness. Our new definition of MALL proof net should yield cleaner and more accessible MALL full completeness results.[2]

Liberation from monomials

The technical starting point for our definition of proof net was Girard's definition of monomial proof net [1996], and we employ variants of Girard's ingenious notions of slice and jump. One of our contributions relative to [Girard 1996] is that we do not partition weights into monomials. In [1996] Girard remarks that he had been trying to circumvent this technical limitation since 1990, and lists three specific problems that must be solved in any attempt to eliminate it, i.e., to define what he calls "more liberal proof-nets", such as ours:

> *Weights must be monomials. However, weights of the form $p \cup q$ will naturally occur if we want to allow more superimpositions. The present state of affairs is as follows:*
> *(1) in spite of years of efforts, I never succeeded in finding the right correctness criterion for these more liberal proof-nets;*
> *(2) general boolean coefficients might be delicate to represent (on the other hand, the case we consider has a natural presentation in terms of coherent spaces);*
> *(3) normalization in the full case might be messy.*
>
> [Girard 1996, Appendix A.1.5]

---

[2]Part of the first author's motivation for finding a satisfactory notion of proof net came from a collaboration with Gordon Plotkin and Vaughan Pratt aiming to extend the Chu space full completeness result [Devarajan, Hughes, Plotkin and Pratt 1999] to MALL: we were initially encumbered by the complexity of monomial proof nets. Ultimately we discovered that full completeness does not extend: the Gustave example (see Section 4.6.1) inhabits the model.

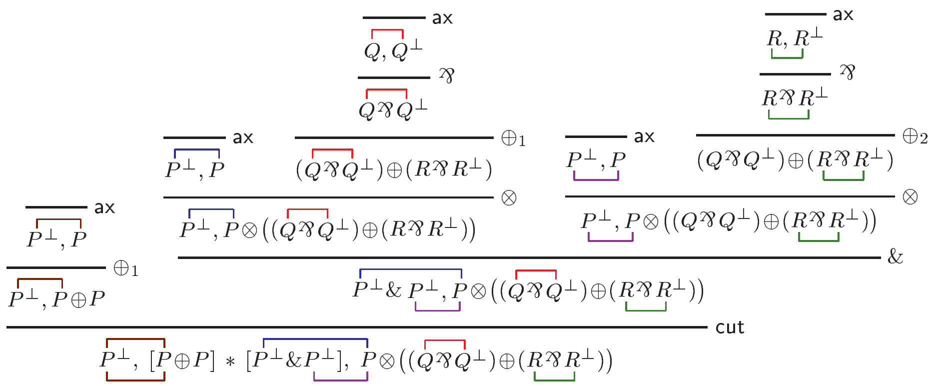

Fig. 2. Example of the translation of a proof with a cut into one of our proof nets. The concluding proof net is on what we call a *cut sequent*: a MALL sequent (the formulas $P^\perp$ and $P \otimes \big((Q ⅋ Q^\perp) \oplus (R ⅋ R^\perp)\big)$) together with a *cut pair* $[P \oplus P] * [P^\perp \& P^\perp]$ formed using the *cut connective* $*$. The concluding proof net comprises two linkings of three axiom links each, one linking drawn above the cut sequent, the other below. When transitioning through the cut rule, the axiom link on $P^\perp, P \oplus P$ on the left becomes duplicated, so that a copy appears in each of the two final linkings; in general, when m linkings pass through the left of a cut rule, and n through the right, we construct all $m \times n$ disjoint unions of the linkings on the conclusion. (Here $m = 1$ and $n = 2$.)

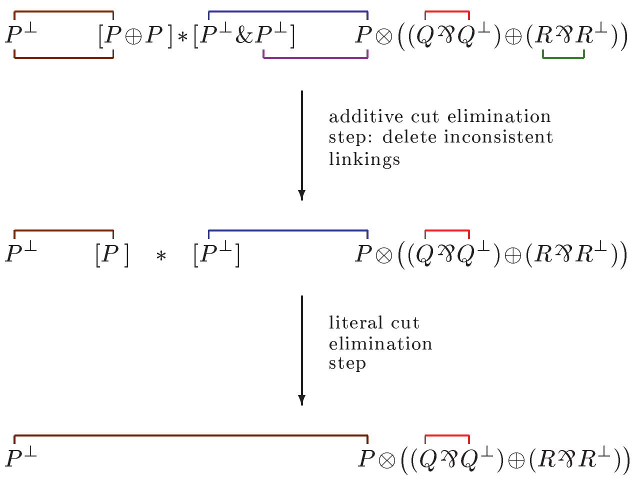

Fig. 3. Example of cut elimination, normalising in two steps. The top proof net, two linkings, was derived in Figure 2. The first elimination step, aside from eliminating the $\oplus$ and $\&$ to leave a literal cut $[P] * [P^\perp]$, deletes the underhanging linking: our rule for additive elimination is simply *delete inconsistent linkings*, where a linking is inconsistent if it chooses opposite arguments for the cut $\oplus$ and $\&$. (Here the underhanging linking chooses $\oplus$-left and $\&$-right, and is therefore inconsistent, hence deleted in the cut elimination step.) Note that the end result is a cut-free proof net: it is the translation of the left branch of the $\&$-rule in Figure 2.

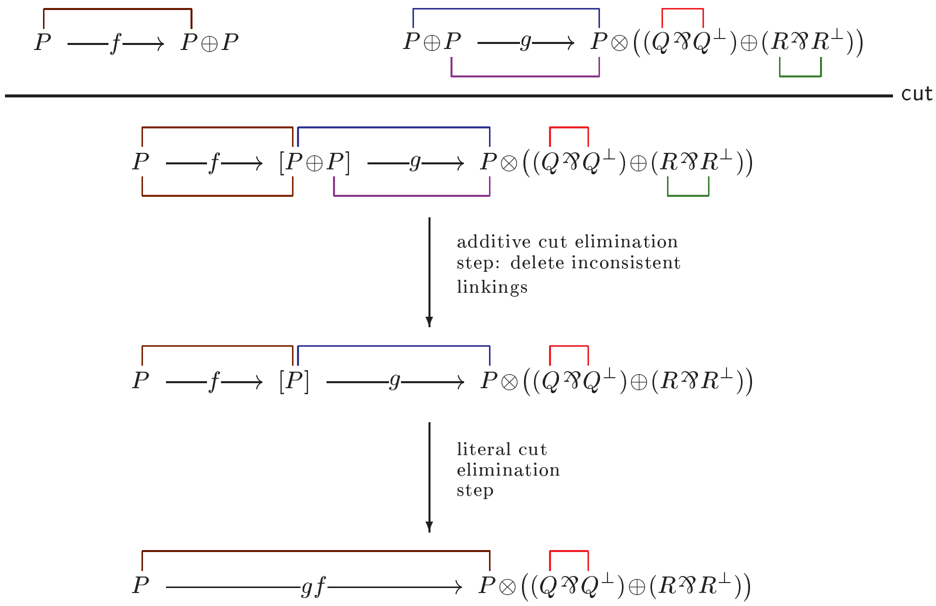

Fig. 4. Example of composition $f, g \mapsto gf$ in our category $\mathcal{N}$ of cut-free proof nets. Objects are MALL formulas, and a morphism $h : A \to B$ is a cut-free proof net on the sequent $A^{\perp}, B$. The morphisms f (top-left) and g (top-right) are the left- and right hypotheses of the cut rule in Figure 2. The first step of composition is to cut the two morphisms; in doing so we are emulating precisely the cut rule of Figure 2. Having negated on the left of the arrow $\to$, the two cut formulas are no longer dual but identical; thus we are afforded the additional economy of superimposing them. The two ensuing computation steps are exactly those of Figure 3, modulo this superposition.

An important stepping-stone towards finding the right criterion to address (1) was to first settle the open problem of whether Girard's criterion becomes insufficient without partitioning weights into monomials. We show that this is indeed the case: in Section 6.2 we present a non-monomial proof structure that does not correspond to any proof (i.e., it is not sequentialisable), yet satisfies Girard's criterion. We address (2) by leaving weights implicit, defining a proof net on a sequent Γ as a set of linkings on an extension of Γ by zero or more cut pairs $A * A^{\perp}$, $B * B^{\perp}$, etc. (See Figure 5 for an example of extracting weights from a proof net.) Issue (3) is addressed by the fact that our definition of cut elimination is very simple: confluence and strong normalisation are immediate.

The proof that our correctness criterion captures proof translations (the *Sequentialisation Theorem*) hinges on an ordering on vertices called *domination*[3]. By introducing domination we avoid the use of empires [Girard 1987; 1996], thereby sidestepping the problem of stability of maximal empires [Girard 1996, Section 1.5.3]—the main technical problem that led Girard to resort to monomials.

In Section 6.4 we define a surjection collapsing Girard's proof nets to ours. There are more Girard proof nets than ours because of the redundancy issues related to monomials (see Section 6.1).

[3]Unrelated to domination in flowgraphs.

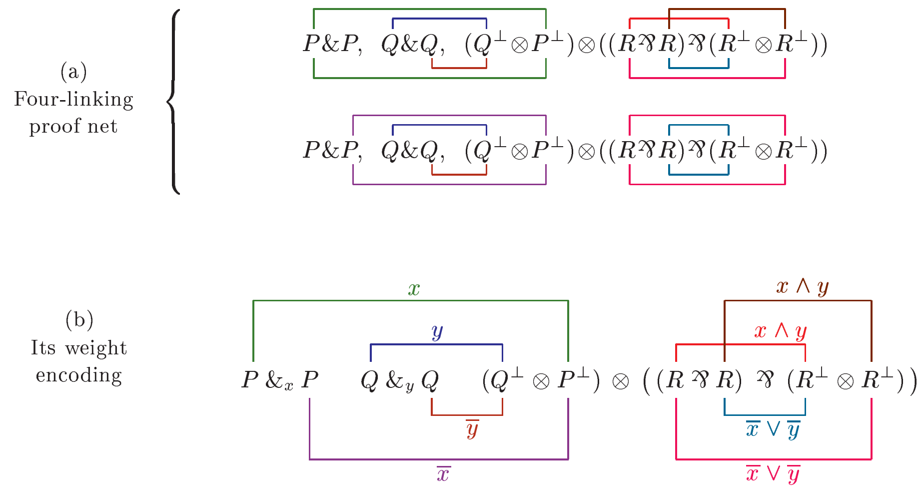

Fig. 5. (a) Four-linking example of one of our proof nets. Rather than draw all four linkings on one sequent, we have drawn two linkings (one above, one below) on each of two copies of the sequent. Section 4.8 shows how to encode a proof net as a collection of axiom links labelled with predicates ('weights', c.f. [Girard 1996]). Subfigure (b) shows the weight encoding of (a). To distinguish the &s, we have subscripted them. Every &-assignment (assignment of *left* or *right* to each of $\&_x$ and $\&_y$) determines a linking by restricting to axiom links whose predicates hold, where we read the predicate x (resp. $\overline{x}$) as "$\&_x$ is assigned *left* (resp. *right*)" (and y analogously), $\wedge$ is *and* and $\vee$ is *or*. We invite the reader to verify that taking each of the four possible &-assignments in turn produces the four original linkings.

2 MALL

By MALL we mean multiplicative-additive linear logic without units [Girard 1987]. Formulas are built from literals (propositional variables $P, Q, \ldots$ and their negations $P^\perp$, $Q^\perp, \ldots$) by the binary connectives ***tensor*** $\otimes$, ***par*** $⅋$, ***with*** & and ***plus*** $\oplus$. Negation $(-)^\perp$ extends to arbitrary formulas with $P^{\perp\perp} = P$ on propositional variables, and de Morgan duality: $(A \otimes B)^\perp = A^\perp ⅋ B^\perp$, $(A ⅋ B)^\perp = A^\perp \otimes B^\perp$, $(A \oplus B)^\perp = A^\perp \& B^\perp$, and $(A \& B)^\perp = A^\perp \oplus B^\perp$. Throughout the paper we shall identify a formula with its parse tree, a tree labelled with literals at the leaves and connectives at internal vertices. A ***sequent*** is a non-empty disjoint union of formulas. Thus a sequent is a particular kind of labelled forest. We write comma for disjoint union. For example,

$$P^\perp, (P \otimes P^\perp) ⅋ P$$

is the graph

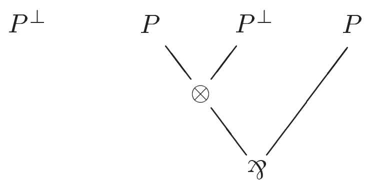

Throughout the paper we adopt the convention of *Mother Nature*, and depict the leaves of a tree above, and the root below. Sequents are proved using the following rules:

$$\frac{}{P, P^{\perp}}\,\mathsf{ax} \qquad \frac{\Gamma, A \quad A^{\perp}, \Delta}{\Gamma, \Delta}\,\mathsf{cut} \qquad \frac{\Gamma, A \quad B, \Delta}{\Gamma, A \otimes B, \Delta}\,\otimes \qquad \frac{\Gamma,\, A,\, B}{\Gamma,\, A \parr B}\,\parr$$

$$\frac{\Gamma, A \quad \Gamma, B}{\Gamma, A \& B}\,\& \qquad \frac{\Gamma,\, A}{\Gamma,\, A \oplus B}\,\oplus_1 \qquad \frac{\Gamma,\, B}{\Gamma,\, A \oplus B}\,\oplus_2$$

Here, and throughout this document, $P, Q, \ldots$ range over propositional variables, $A, B, \ldots$ over formulas, and $\Gamma, \Delta, \ldots$ over (possibly empty) disjoint unions of formulas. Without loss of generality we restrict the axiom rule to literals [Girard 1987].

3 Background: Cut-free MLL proof nets

For didactic purposes, we review the definition of cut-free MLL proof net [Girard 1987; Danos and Regnier 1989]. MLL is the subsystem of MALL obtained by omitting the additive connectives, $\oplus$ and $\&$. Our style of presentation anticipates the subsequent definition of MALL proof net.

An ***axiom link*** or simply ***link*** on an MLL sequent Γ is an edge between complementary leaves in Γ, i.e., between leaves in Γ labelled with complementary literals P and $P^{\perp}$. A ***linking*** on Γ is a partitioning of the leaves of Γ into links, i.e., a set of disjoint links whose union contains every leaf of Γ. A linking on an MLL sequent is also called a ***cut-free MLL proof structure***.

Example 3.1. Two linkings are possible on the sequent $P^{\perp}, (P \otimes P^{\perp}) \parr P$:

$$P^{\perp}, (P \otimes P^{\perp}) \parr P \qquad\qquad P^{\perp}, (P \otimes P^{\perp}) \parr P$$

3.1 A function from cut-free MLL proofs to linkings

Let Π be a cut-free MLL proof of a sequent Γ. By downwards tracking of formula leaves, the axiom rules of Π determine a linking λ_{Π} on Γ. Alternatively, one can define the same function from proofs to linkings by induction. The base case of an axiom rule $\overline{P, P^{\perp}}$ defines the linking $P, P^{\perp}$. Writing $\lambda \rhd \Gamma$ for the judgement "λ is a linking on Γ", the inductive translation is as follows:

$$\frac{\lambda \;\rhd\; \Gamma, A \qquad \lambda' \;\rhd\; B, \Delta}{\lambda \cup \lambda' \;\rhd\; \Gamma, A \otimes B, \Delta}\,\otimes \qquad\qquad \frac{\lambda \;\rhd\; \Gamma, A, B}{\lambda \;\rhd\; \Gamma, A \parr B}\,\parr$$

Here we use the implicit tracking of formula leaves above the line of a rule to leaves below the line. Figure 6 shows an example. Any linking λ which is the image of a proof is ***sequentialisable***, and any such proof is a ***sequentialisation*** of λ. In general a linking has many distinct sequentialisations, corresponding to the fact that MLL proof nets are canonical abstract representations of MLL proofs modulo inessential rule commutation.

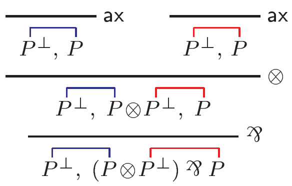

Fig. 6. An example of the translation of a cut-free MLL proof into a linking, i.e., into a cut-free MLL proof structure.

3.2 Geometric characterisation of sequentialisability

Given a linking λ on Γ, the ***graph*** $\mathcal{G}_\lambda$ ***of*** λ is the graph Γ together with the edges λ. A ***⅋-switching*** of a linking λ on Γ is any subgraph of $\mathcal{G}_\lambda$ obtained by deleting one of the two argument edges of each ⅋-vertex.

Example 3.2. One of two possible ⅋-switchings of the first linking of Example 3.1:

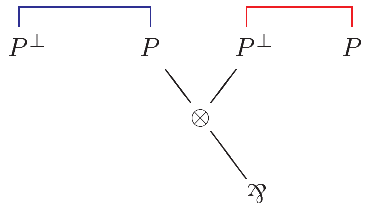

Definition 3.3. A linking on an MLL sequent (i.e., a cut-free MLL proof structure) is a ***cut-free MLL proof net*** if each of its ⅋-switchings is a tree (acyclic and connected).

Example 3.4. The second linking of Example 3.1 fails to be a cut-free MLL proof net. This ⅋-switching is not a tree:

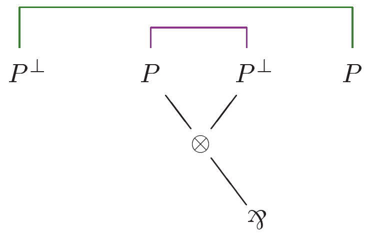

The first linking of Example 3.1 is a proof net: both ⅋-switchings (one of which was depicted in Example 3.2) are trees.

THEOREM 3.5 CUT-FREE MLL SEQUENTIALISATION. *A linking is the translation of a cut-free proof iff it is a cut-free proof net.*

This was proved by Girard, for a different geometric criterion, based on *long trips* [1987]. Danos and Regnier [1989] simplified the criterion to the elegant one above, showing it to be equivalent to Girard's. Several other equivalent formulations will be presented in Sections 4.7.1 and 4.7.2.

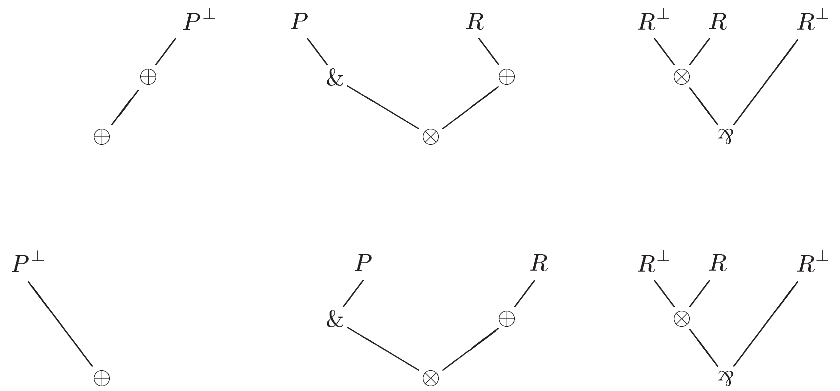

$$P^\perp \oplus (Q \oplus P^\perp),\ (P \& P) \otimes (R \oplus R),\ (R^\perp \otimes R) \parr R^\perp$$

$$P^\perp \oplus (Q \oplus P^\perp),\ (P \& P) \otimes (R \oplus R),\ (R^\perp \otimes R) \parr R^\perp$$

Fig. 7. Top: two additive resolutions of $P^\perp \oplus (Q \oplus P^\perp)$, $(P \& P) \otimes (R \oplus R)$, $(R^\perp \otimes R) \parr R^\perp$. Equivalent compact 'in-line' representations are shown underneath.

4 Cut-free MALL proof nets

We begin by defining a *linking* on a MALL sequent, and a simple function from cut-free MALL proofs to sets of linkings. With such a function in hand, it is natural to ask about its image and kernel:

(I) *Image.* Can one characterise the sound sets of linkings, i.e., those that come from proofs?

(K) *Kernel.* Does the kernel exactly characterise proof equivalence modulo rule commutation?

We answer both in the affirmative. In Section 4.3 we present a geometric characterisation of those sets of linkings that arise as the translations of cut-free MALL proofs, and call them *proof nets.* In a sibling paper we show that any two cut-free MALL proofs are equal modulo rule commutation if and only if they map to the same proof net (see Section 4.11). Thus:

Our cut-free MALL proof nets provide canonical abstract representations of cut-free MALL proofs modulo rule commutation.

4.1 Linkings

An ***additive resolution*** of a MALL sequent Γ is any result of deleting one argument subtree of every additive connective ($\&$ or $\oplus$) of Γ. See Figure 7 for examples. An ***axiom link*** or simply ***link*** on Γ is an edge between complementary leaves in Γ, i.e., between leaves in Γ labelled with complementary literals P and $P^\perp$. A ***linking***

λ on Γ is a set of disjoint links on Γ such that $\cup\lambda$ is the set of leaves of an additive resolution of Γ; this additive resolution is denoted $\Gamma \restriction \lambda$.

Example 4.1. Let Γ be the sequent

$$P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \& P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) ⅋ R^{\perp}.$$

The following set λ of three disjoint links is an example of a linking on Γ:

$$P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \& P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) ⅋ R^{\perp}$$

For λ to be a linking, as opposed to merely an ad hoc collection of disjoint links, it must take the leaves of some additive resolution of Γ. This is indeed the case: the leaves of (the links of) λ are exactly those of the first of the two additive resolutions depicted in Figure 7:

$$P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \& P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) ⅋ R^{\perp}$$

Example 4.2. Multiple linkings can have the same additive resolution. For example, the following linking λ'

$$P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \& P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) ⅋ R^{\perp}$$

has the same additive resolution as the linking λ of Example 4.1, i.e., $\Gamma \restriction \lambda = \Gamma \restriction \lambda'$:

$$P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \& P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) ⅋ R^{\perp}$$

Note that λ and λ' are the only two linkings possible on this additive resolution.

Example 4.3. This pair of disjoint links fails to be a linking:

$$P \,\&\, Q,\ Q^{\perp} \otimes P^{\perp}.$$

It is not a linking because it contains a leaf on each side of the &.

Numerous other examples of linkings can be seen in Figures 1 and 2 (pages 4 and 5). One can easily verify that each of them takes the leaves of an additive resolution. See also Figure 5 (page 7).

4.1.1 *Every linking induces an MLL proof structure.* Every additive connective ($\oplus$/&) remaining in an additive resolution is unary (i.e., has one remaining argument), by construction. One can observe this, for example, in the parse trees in Figure 7. Thus any additive resolution R of a MALL sequent Γ induces an MLL sequent R^{-} by collapsing its additive connectives. A linking λ on Γ, viewed as being on $(\Gamma \restriction \lambda)^{-}$, is a cut-free MLL proof structure (as defined in Section 3), which we call the ***MLL proof structure induced by*** λ.

Example 4.4. The MLL proof structure induced by the linking λ of Example 4.1:

$$P^{\perp},\ P \otimes R,\ (R^{\perp} \otimes R) ⅋ R^{\perp}$$

$$\frac{}{\{\overline{P, P^{\perp}}\} \;\rhd\; P, P^{\perp}}\,\mathsf{ax} \qquad \frac{\theta \;\rhd\; \Gamma, A, B}{\theta \;\rhd\; \Gamma, A \parr B}\,\parr \qquad \frac{\theta \;\rhd\; \Gamma, A \qquad \theta' \;\rhd\; \Gamma, B}{\theta \cup \theta' \;\rhd\; \Gamma, A \& B}\,\&$$

$$\frac{\theta \;\rhd\; \Gamma, A \qquad\qquad \theta' \;\rhd\; B, \Delta}{\{\lambda \cup \lambda' \,:\, \lambda \in \theta, \lambda' \in \theta'\} \;\rhd\; \Gamma, A \otimes B, \Delta}\,\otimes \qquad \frac{\theta \;\rhd\; \Gamma, A}{\theta \;\rhd\; \Gamma, A \oplus B}\,\oplus_1 \qquad \frac{\theta \;\rhd\; \Gamma, B}{\theta \;\rhd\; \Gamma, A \oplus B}\,\oplus_2$$

Table I. Inductive definition of the function from cut-free MALL proofs to sets of linkings. Here $\theta \rhd \Gamma$ is the judgement "θ is a set of linkings on Γ". We use the implicit tracking of formula leaves downwards through rules. The base case ax is a singleton set of linkings whose only linking comprises a single link, between P and $P^{\perp}$.

4.2 A function from cut-free MALL proofs to sets of linkings

Every cut-free MALL proof Π of Γ defines a set θ_Π of linkings on Γ as follows. Define a ***&-resolution*** R of Π to be any result of deleting one branch above each &-rule of Π. By downwards tracking of formula leaves, the axiom rules of R determine a linking λ_R on Γ. Define $\theta_\Pi = \{\lambda_R : R \text{ is a \&-resolution of } \Pi\}$. See Figure 8 (page 13) for an example. Alternatively, Table I defines the same function by induction; see Figure 1 (page 4) for an example.

By structural induction, each linking is well-defined (i.e., takes the leaves of an additive resolution); thus the translation is well-defined. The fact that the above procedures yield the same set of linkings follows from a simple structural induction on proofs. A set of linkings Λ is ***sequentialisable*** if it is the translation of a proof; any such proof is a ***sequentialisation*** of Λ.

4.3 Geometric characterisation of sequentialisability

In this section we define a *proof net* as a set of linkings satisfying three conditions. These conditions characterise the image of the function from cut-free proofs to sets of linkings defined in Section 4.2: in Theorem 4.18 (the cut-free *Sequentialisation Theorem*) we prove that a set of linkings is the translation of a proof if and only if it is a proof net. The definition of proof net is pleasingly succinct, and is given in the box on page 14. In the remainder of this section we clarify the definition and work through examples. As in the standard approach to MLL (and as in [Girard 1996]), we define a *proof structure* as a stepping-stone towards the definition of proof net.

4.3.1 *Resolution condition.* Similar to the definition of additive resolution in Section 4.1, define a ***&-resolution*** of a sequent Γ to be any result of deleting one argument subtree of every & of Γ.

Example 4.5. The two possible &-resolutions of the sequent

$$P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \& P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) \parr R^{\perp}$$

featured in Examples 4.1 and 4.2 are:

$$\Gamma_1^\star : \qquad P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \& P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) \parr R^{\perp}$$

$$\Gamma_2^\star : \qquad P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \& P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) \parr R^{\perp}$$

(Π)
$$\dfrac{\dfrac{\dfrac{\dfrac{}{P^\perp, P}\,\mathsf{ax}}{P^\perp \oplus Q^\perp,\ P}\,\oplus_1}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ P}\,\oplus_1 \qquad \dfrac{\dfrac{\dfrac{\dfrac{}{Q^\perp, Q}\,\mathsf{ax}}{P^\perp \oplus Q^\perp,\ Q}\,\oplus_2}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ Q}\,\oplus_1 \qquad \dfrac{\dfrac{}{R^\perp, R}\,\mathsf{ax}}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ R}\,\oplus_2}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ Q \& R}\,\&}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ P \& (Q \& R)}\,\&$$

(R_1)
$$\dfrac{\dfrac{\dfrac{\dfrac{}{P^\perp, P}\,\mathsf{ax}}{P^\perp \oplus Q^\perp,\ P}\,\oplus_1}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ P}\,\oplus_1}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ P \& (Q \& R)}\,\&$$

(R_2)
$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{Q^\perp, Q}\,\mathsf{ax}}{P^\perp \oplus Q^\perp,\ Q}\,\oplus_2}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ Q}\,\oplus_1}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ Q \& R}\,\&}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ P \& (Q \& R)}\,\&$$

(R_3)
$$\dfrac{\dfrac{\dfrac{\dfrac{}{R^\perp, R}\,\mathsf{ax}}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ R}\,\oplus_2}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ Q \& R}\,\&}{(P^\perp \oplus Q^\perp)\oplus R^\perp,\ P \& (Q \& R)}\,\&$$

λ_1 : $(P^\perp \oplus Q^\perp)\oplus R^\perp,\ P \& (Q \& R)$

λ_2 : $(P^\perp \oplus Q^\perp)\oplus R^\perp,\ P \& (Q \& R)$

λ_3 : $(P^\perp \oplus Q^\perp)\oplus R^\perp,\ P \& (Q \& R)$

Fig. 8. Example of the mapping of a cut-free MALL proof into a set of linkings. At the top is a proof Π, followed by its three possible &-resolutions R_1, R_2, R_3, followed by the corresponding linkings λ_1, λ_2, λ_3. Each linking comprises a single link. Categorically, this example expresses associativity $(P \times Q) \times R \to P \times (Q \times R)$. Note the compactness of the representation as a set of linkings relative to the size of the proof.

Definition: cut-free MALL proof net on a sequent Γ

Additive resolution: deletion of one argument subtree of each $\oplus/\&$; $\&$-***resolution*** analogous.
(***Axiom***) ***link*** on Γ: edge between complementary leaves (literal occurrences) in Γ.
Linking λ on Γ: partitioning of the leaves of an additive resolution $\Gamma\restriction\lambda$ of Γ into links.
A set Λ of linkings on Γ ***toggles*** a $\&$ w if both arguments of w are in $\Gamma\restriction\Lambda \equiv \bigcup_{\lambda\in\Lambda}\Gamma\restriction\lambda$.
Graph $\mathcal{G}_\Lambda$: $\Gamma\restriction\Lambda + \cup\Lambda +$ ***jump*** edges l—w—l' if $\{l,l'\}\in\lambda\setminus\lambda'$ and $\{\lambda,\lambda'\}\subseteq\Lambda$ toggles w only.
$\parr$-***switching*** of λ: any subgraph of $\mathcal{G}_{\{\lambda\}}$ obtained by deleting one argument edge of each $\parr$.
Switching cycle: cycle with ≤ 1 ***switch edge*** (= jump or argument edge) of each $\parr/\&$.

A set θ of linkings on Γ is a ***proof net*** if it satisfies:

RESOLUTION: Exactly one linking of θ is on any given $\&$-resolution of Γ.
MLL: Every $\parr$-switching of every $\lambda\in\theta$ is a tree (i.e., each $\lambda\in\theta$ induces an MLL proof net).[4]
TOGGLING: Every set Λ of ≥ 2 linkings of θ toggles a $\&$ that is in no switching cycle of $\mathcal{G}_\Lambda$.[5]

A linking λ on Γ is ***on*** a $\&$-resolution $\Gamma^\star$ of Γ if every leaf of λ is in $\Gamma^\star$. A set of linkings θ on Γ is a ***cut-free proof structure*** if it satisfies

(P1) RESOLUTION. *For any $\&$-resolution $\Gamma^\star$ of Γ, exactly one linking of θ is on $\Gamma^\star$.*

Example 4.6. Here is a two-linking proof structure $\theta = \{\lambda_1, \lambda_2\}$ on the sequent of Example 4.5, with λ_1 drawn above the sequent and λ_2 drawn below:

λ_1 :
λ_2 :

$$P^\perp \oplus (Q \oplus P^\perp),\ (P \& P) \otimes (R \oplus R),\ (R^\perp \otimes R) \parr R^\perp$$

To verify RESOLUTION, we must check that exactly one of the linkings fits on each of the two $\&$-resolutions of Γ, depicted in Example 4.5. Taking the $\&$-resolution $\Gamma_1^\star$,

λ_1 :
λ_2 :

$$P^\perp \oplus (Q \oplus P^\perp),\ (P \& P) \otimes (R \oplus R),\ (R^\perp \otimes R) \parr R^\perp$$

we see that λ_1 is on $\Gamma_1^\star$ (all six of its leaves are in $\Gamma_1^\star$), but λ_2 is not (its P literal is not in $\Gamma_1^\star$). Similarly, taking the second $\&$-resolution $\Gamma_2^\star$,

λ_1 :
λ_2 :

$$P^\perp \oplus (Q \oplus P^\perp),\ (P \& P) \otimes (R \oplus R),\ (R^\perp \otimes R) \parr R^\perp$$

we see that λ_2 is on $\Gamma_2^\star$ (all six of its leaves are in $\Gamma_2^\star$), but λ_1 is not (its P literal is not in $\Gamma_2^\star$). Hence RESOLUTION is satisfied.

Example 4.7. The pair of linkings

$$P^\perp,\ P \otimes Q^\perp,\ Q \oplus Q$$

[4]Tree = acyclic + connected. Dropping the connectedness requirement in the MLL condition yields a cut-free proof net for MALL augmented with the mix rule. See Section 4.9.
[5]In fact, it suffices to verify TOGGLING merely for ***saturated*** sets of linkings Λ, namely, such that any strictly larger subset of θ toggles more $\&$s than Λ. There is exactly one saturated set of linkings in θ for each ***partial*** $\&$-***resolution*** of Γ, the latter being any result of deleting at most one argument subtree of each $\&$ of Γ.

fails RESOLUTION: any &-free sequent is its own unique &-resolution, and therefore RESOLUTION will hold if and only if there is a single linking.

Example 4.8. The singleton set of linkings

$$\overbrace{P \oplus (Q\&R),\ P}{}^{\perp}$$

(comprising just one link) satisfies RESOLUTION. Note that the sequent has two distinct &-resolutions, but there is only one linking.

Remark 4.9. In the restricted case of an MLL sequent Γ, since there are no &s, a set of linkings satisfies the resolution condition iff it comprises a single MLL linking on Γ (in the sense of Section 3). Thus our cut-free MALL proof structures generalise cut-free MLL proof structures.

Example 4.10. We invite the reader to verify the resolution condition for the sets of linkings in Figures 1, 2 and 5 (pages 4, 5, and 7).

Section 4.4 provides intuition for the resolution condition. The resolution condition, on its own, suffices as a correctness criterion for pure additive proof nets: see Section 4.10. Section 4.8 shows how to encode a proof structure using *weights* (c.f. [Girard 1996]), as illustrated by the example in Figure 5 (page 7). In Section 6.3 we detail the relationship between RESOLUTION and Girard's so-called technical condition.

4.3.2 *MLL condition.* The second requirement for a set of linkings θ to be a proof net is "pointwise MLL correctness":

(P2) MLL. *Every linking of θ induces an MLL proof net.*

In other words, for each linking $\lambda \in \theta$, the MLL proof structure induced by λ (as defined in Section 4.1.1), is an MLL proof net (as defined in Section 3).

Example 4.11. See Figure 9, subfigures (a)–(d).

Example 4.12. The proof structure $\theta = \{\lambda_1, \lambda_2\}$ in Example 4.6 (page 14) satisfies the MLL condition. Both λ_1 and λ_2 induce the same MLL proof net, whose graph is Figure 9(c).

Naturally one need not collapse to an MLL proof structure to check the MLL condition for a linking λ: one can simply leave the unary $\oplus$/&s of the additive resolution in place, and verify that every $\parr$-switching is a tree. For self-containedness of our definition of cut-free MALL proof net, without reference to MLL proof nets, we describe this formally.

Construct the ***graph*** $\mathcal{G}_\lambda$ ***of*** λ from the graph of the additive resolution $\Gamma \upharpoonright \lambda$ (a subgraph of Γ) by adding the edges λ. For example, Figure 9(e) shows the graph $\mathcal{G}_{\lambda_1}$ of the linking λ_1 of Figure 9(a). A ***$\parr$-switching*** of a linking λ on Γ is any subgraph of $\mathcal{G}_\lambda$ obtained by deleting one of the two argument edges of each $\parr$. See Figure 9(f) for an example. Clearly, the induced MLL proof structure of a linking λ is an MLL proof net if and only if every $\parr$-switching of λ (in $\mathcal{G}_\lambda$) is a tree. Thus we can reformulate the MLL condition on a set of linkings θ, without reference to MLL proof nets, as follows:

(a)
The linking
λ_1 on Γ

$$P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \,\&\, P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) \parr R^{\perp}$$

(b)
The linking λ_1
on its additive
resolution $\Gamma \upharpoonright \lambda_1$

$$P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \,\&\, P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) \parr R^{\perp}$$

(c)
The graph of the
MLL proof net
induced by λ_1

(d)
A $\parr$-switching
of (c)

(e)
The graph $\mathcal{G}_{\lambda_1}$

(f)
The corresponding
$\parr$-switching of λ_1

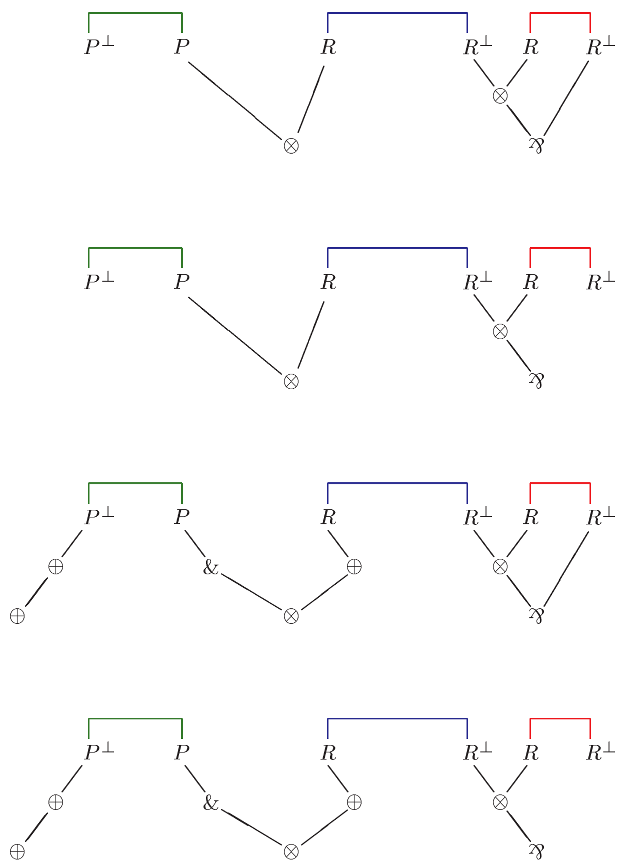

Fig. 9. Subfigure (a) shows a linking λ_1 on a MALL sequent Γ, which is shown on its additive resolution in (b). Subfigure (c) is the MLL proof structure induced by λ_1, which is an MLL proof net since each of its $\parr$-switchings is a tree. Subfigure (d) shows one of its two $\parr$-switchings. Subfigure (e) is the graph $\mathcal{G}_{\lambda_1}$ of λ_1 on Γ, and (f) is the $\parr$-switching of λ_1 in $\mathcal{G}_{\lambda_1}$ corresponding to the $\parr$-switching (d) of the induced MLL proof net (c).

(P2) MLL. *Every ⅋-switching of every linking of θ is a tree (acyclic and connected).*

Relaxing the connectedness requirement yields a notion of cut-free proof net for MALL augmented with the mix rule. See Section 4.9.

4.3.3 *Toggling condition.* We require some auxiliary concepts to state our third and last proof net condition. A set of linkings Λ ***toggles*** a &-vertex w of Γ if both arguments of w are present in $\bigcup_{\lambda\in\Lambda}\Gamma\restriction\lambda$, i.e., there exist $\lambda_l, \lambda_r \in \Lambda$ such that the left argument of w is present in the additive resolution $\Gamma\restriction\lambda_l$ and the right argument of w is present in the additive resolution $\Gamma\restriction\lambda_r$.

Example 4.13. Recall our running example,

$$\begin{array}{ll} \lambda_1 : & \\ & P^\perp \oplus (Q \oplus P^\perp),\ (P \,\&\, P) \otimes (R \oplus R),\ (R^\perp \otimes R) ⅋ R^\perp \\ \lambda_2 : & \end{array}$$

The pair of linkings $\theta = \{\lambda_1, \lambda_2\}$ toggles the & of the underlying sequent Γ because its left argument (the left P) is present in the additive resolution $\Gamma\restriction\lambda_1$, and its right argument (the right P) is present in the additive resolution $\Gamma\restriction\lambda_2$. Neither $\{\lambda_1\}$ nor $\{\lambda_2\}$ toggles the &: a single linking can never toggle a & because all additives are unary in an additive resolution.

Let Λ be a set of linkings. A link a ***depends on*** w ***in*** Λ if, inside Λ, a can be made to vanish by toggling w alone: there exist $\lambda, \lambda' \in \Lambda$ such that $a \in \lambda$, $a \notin \lambda'$, and w is the only & toggled by $\{\lambda, \lambda'\}$.

Example 4.14. In

$$\begin{array}{ll} \lambda_1 : & \\ & P^\perp \oplus (Q \oplus P^\perp),\ (P \,\&\, P) \otimes (R \oplus R),\ (R^\perp \otimes R) ⅋ R^\perp \\ \lambda_2 : & \end{array}$$

let w be the & of the sequent. The link between the left-most R and the left-most $R^\perp$ depends on w in $\Lambda = \{\lambda_1, \lambda_2\}$: it is present in $\lambda_1 \in \Lambda$ but not in $\lambda_2 \in \Lambda$, and w is the only & toggled by $\{\lambda_1, \lambda_2\}$. The link between the right-most R and $R^\perp$ does not depend on w in Λ, since it is present in both λ_1 and λ_2. It is the only one of the five links in Λ (more precisely, in $\bigcup\Lambda$) that does not depend on w in Λ.

We now extend the definition of the graph of a linking to the graph of a set Λ of linkings on Γ. The ***partial additive resolution*** of Λ is the graph $\Gamma\restriction\Lambda = \bigcup_{\lambda\in\Lambda}\Gamma\restriction\lambda$, the union (superposition) of the additive resolutions of the linkings of Λ. Some additives of $\Gamma\restriction\Lambda$ may be unary, some binary. The ***graph*** $\mathcal{G}_\Lambda$ ***of*** Λ is $\Gamma\restriction\Lambda$ together with each edge $\{l, l'\}$ of (a linking of) Λ, and ***jump*** edges from l and l' to any &-vertex on which $\{l, l'\}$ depends in Λ. For example, in Figure 10 (page 18), subfigure (b) shows the graph of the pair of linkings in subfigure (a), and subfigure (f) shows the graph of the pair of linkings in subfigure (e). Note that $\Lambda \subseteq \Lambda'$ implies $\mathcal{G}_\Lambda \subseteq \mathcal{G}_{\Lambda'}$, and that for any linking λ, $\mathcal{G}_{\{\lambda\}} = \mathcal{G}_\lambda$ (the graph of a single linking, defined in Section 4.3.2), because $\mathcal{G}_{\{\lambda\}}$ has no jumps (since a single linking toggles no &s).

A ***switch edge*** of a &- or ⅋-vertex x of $\mathcal{G}_\Lambda$ is an edge between x and one of its arguments, or a jump to x (if x is a &). For example, Figure 9(e) on page 16 has three switch edges, the left argument edge of the &, and both argument edges of

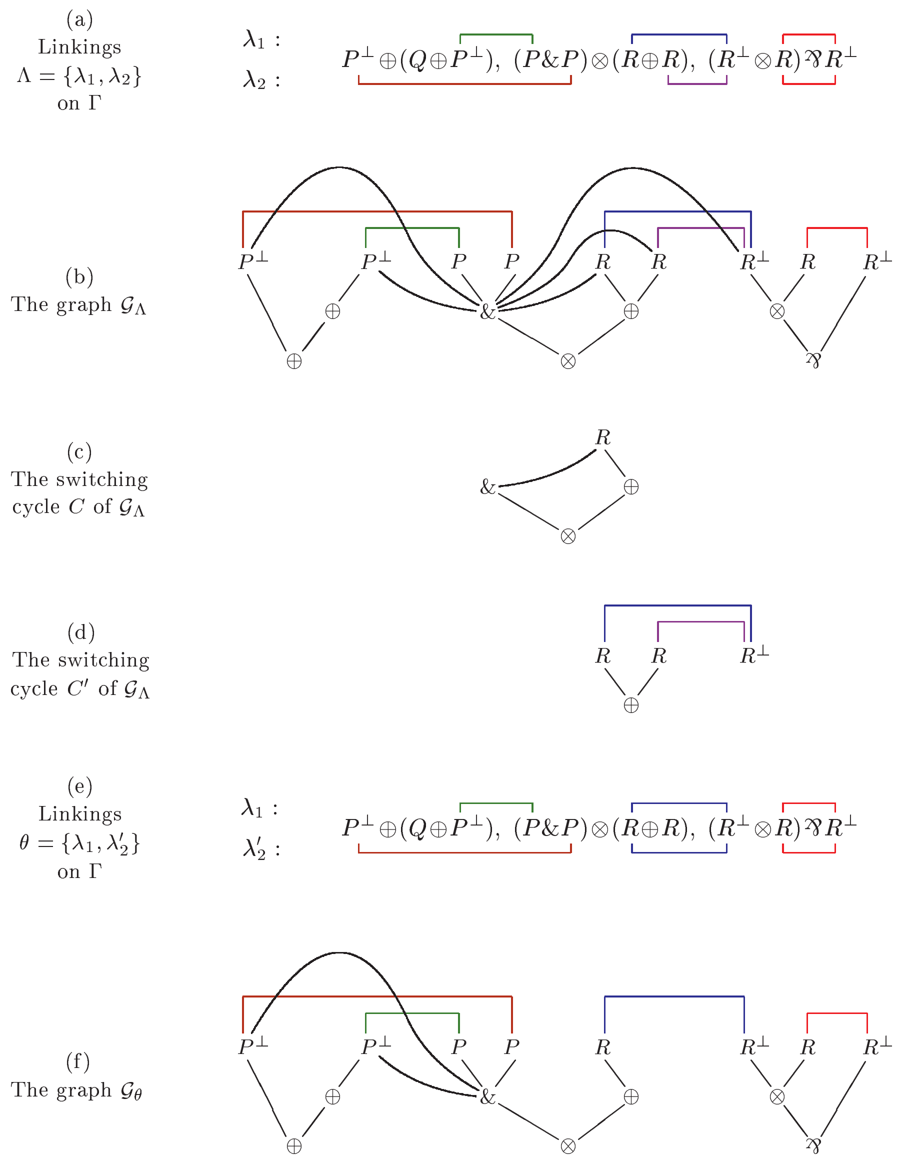

Fig. 10. Subfigure (a) is a pair of linkings $\Lambda = \{\lambda_1, \lambda_2\}$ whose graph $\mathcal{G}_{\Lambda}$ is depicted in subfigure (b). To distinguish jumps, we draw them as curved edges (unless the jump edge was already present as an argument edge, in which case it remains straight). There is no jump to a leaf of the right-most link, since it does not depend on the & in Λ. (This was explained in detail in Example 4.14, page 17.) Subfigures (c) and (d) show switching cycles of $\mathcal{G}_{\Lambda}$. Subfigure (e) is a pair of linkings $\theta = \{\lambda_1, \lambda_2'\}$, whose graph $\mathcal{G}_{\theta}$ is (f). This pair of linkings satisfies the toggling condition: the only subset of θ of two or more linkings is θ itself, so to verify the condition we need only confirm that $\mathcal{G}_{\theta}$ contains no switching cycle; this is apparent from the depiction of $\mathcal{G}_{\theta}$ in subfigure (f).

the ⅋. Figure 10(b) has 9 switch edges, the two argument edges of the ⅋ and the 7 jumps to the & (two of which are argument edges).

A ***cycle*** of $\mathcal{G}_\Lambda$ is a subgraph of $\mathcal{G}_\Lambda$ with vertex set $\{x_1, \dots, x_n\}$ for $n \geq 3$, all x_i distinct, and an edge $x_i — x_{i+1}$ for all i (mod n). A cycle ***switches*** or ***is switching*** if it contains at most one switch edge of each & and ⅋. For example, the graph $\mathcal{G}_\Lambda$ of Figure 10(b) contains the switching cycles C and C' shown below it (subfigures (c) and (d)). Our third and final proof net condition on a set of linkings θ is:

(P3) TOGGLING. *Every set Λ of two or more linkings of θ toggles a & that is not in any switching cycle of $\mathcal{G}_\Lambda$.*

It is clear from the definition of the graph $\mathcal{G}_\Lambda$ that it suffices to verify TOGGLING for ***saturated*** sets of linkings Λ, namely, such that any strictly larger subset of θ toggles more &s than Λ. Note that there is exactly one saturated set of linkings in θ for each ***partial &-resolution*** of Γ, the latter being any result of deleting at most one argument subtree of each & of Γ. We retain the more general quantification over Λ in the formulation of the toggling condition so that the definition of proof net is more succinct.

Example 4.15. The pair of linkings Λ in Figure 10(a) fails the toggling condition, because of the switching cycle C of subfigure (c), which traverses the &.

More generally, whenever every & is in a switching cycle (in the case of Example 4.15, just one &), the toggling condition fails. Another example of this will be given in Section 6.2.

Example 4.16. The pair of linkings θ in Figure 10(e) satisfies the toggling condition. Any switching cycle in the graph $\mathcal{G}_\theta$ (Figure 10(f)) is only permitted to use one switch edge of the &, and therefore to traverse the & it must go via the $\otimes$ immediately below it. Since there is no cycle containing the $\otimes$, there is no switching cycle through the &.

The box on page 14 defines a cut-free MALL proof net as a set of linkings on a MALL sequent satisfying all three conditions introduced above: (P1) RESOLUTION, (P2) MLL, and (P3) TOGGLING. (In other words, a cut-free MALL proof net is a cut-free MALL proof structure satisfying the MLL and toggling conditions.) In the example below, we go through the full process of verifying all three conditions.

Example 4.17. Consider the pair of linkings on the sequent $\Gamma \equiv P^\perp \& P^\perp, P \oplus P$ obtained as follows:

$$\dfrac{\dfrac{\dfrac{}{P^\perp, P}\,\mathsf{ax}}{P^\perp, P\oplus P}\,\oplus_2 \qquad \dfrac{\dfrac{}{P^\perp, P}\,\mathsf{ax}}{P^\perp, P\oplus P}\,\oplus_1}{P^\perp \& P^\perp, P\oplus P}\,\&$$

Let λ_1 and λ_2 be the upper- and lower linking of the concluding sequent, respectively (each having just one link). We shall verify that $\theta = \{\lambda_1, \lambda_2\}$ is a cut-free proof net. Γ has two &-resolutions, $\Gamma_1^\star \equiv P^\perp \& P^\perp, P \oplus P$ and $\Gamma_2^\star \equiv P^\perp \& P^\perp, P \oplus P$. The resolution condition holds, since θ contains exactly one linking on $\Gamma_i^\star$, namely λ_i. Here are the graphs $\mathcal{G}_{\lambda_1}$, $\mathcal{G}_{\lambda_2}$, and $\mathcal{G}_\theta$:

$\mathcal{G}_{\lambda_1}$

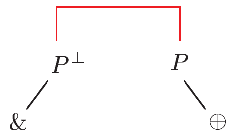

$\mathcal{G}_{\lambda_2}$

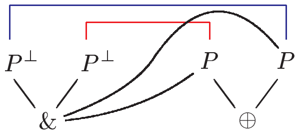

$\mathcal{G}_\theta$

Each λ_i has just one ⅋-switching, namely $\mathcal{G}_{\lambda_i}$; since each $\mathcal{G}_{\lambda_i}$ is a tree, the MLL condition holds. Finally, the toggling condition holds since θ toggles the &, which is not in any switching cycle of $\mathcal{G}_\theta$. (An outermost &, i.e., one that is not an argument of any other connective, can never be in a switching cycle.)[6]

Section 4.6 provides proof-theoretic intuition for the toggling condition.

Theorem 4.18 Cut-free sequentialisation. *A set of linkings is the translation of a cut-free proof iff it is a cut-free proof net.*

By a simple induction, the translation of a cut-free proof is a cut-free proof net. The proof of the converse reduces to a simple induction on the number of ⅋s and &s (Section 4.13) once we prove (Section 4.12):

Lemma 4.19 Separation Lemma. *For any cut-free proof net θ, if $\mathcal{G}_\theta$ has a ⅋ or &, then it has a ⅋ or & that separates.*

Here a ⅋- or &-vertex x ***separates*** if it is not an argument (i.e., is an outermost connective), or it is the argument of y and deleting the edge between x and y disconnects[7] $\mathcal{G}_\theta$. We shall prove the Separation Lemma via an ordering on &s and ⅋s which we call *domination*[8], a concept reminiscent of the ordering induced by the notion of an *empire* of Girard [1996], but different in an essential way.

The remainder of this section is structured as follows. Sections 4.4, 4.5 and 4.6 provide intuition for the resolution, MLL and toggling conditions, respectively. Section 4.7 presents some alternative formulations of the definition of proof net. Section 4.8 describes how to encode a proof structure/net using weights. Section 4.9 defines a *mix net* as the analogue of a proof net in the case of MALL augmented with the mix rule. Section 4.10 notes that the resolution condition, on its own, suffices as

[6] More generally, there are n^m proof nets on the sequent $\&^m P^\perp, \oplus^n P$ (above $m = n = 2$), in bijection with natural transformations $\coprod^m X \to \coprod^n X$ on sets, or equivalently, $\prod^n X \to \prod^m X$.
[7] In the case with mix, read "disconnects" as "increases the number of connected components of".
[8] Unrelated to domination in flowgraphs.

a correctness criterion for additive proof nets. Section 4.11 observes that our cut-free proof nets exactly capture cut-free MALL proofs modulo commutation of rules. We conclude by proving the cut-free sequentialisation theorem in Sections 4.12 and 4.13 (the Separation Lemma and the main induction, respectively).

4.4 Intuition for the resolution condition

Recall from Section 4.2 that the set θ_Π of linkings obtained from a cut-free MALL proof Π comprises one linking $\lambda_R \in \theta_\Pi$ per &-resolution R of Π. This correspondence between proof &-resolutions and linkings is what is captured in the resolution condition. (One can observe this correspondence in Figures 8 (page 13) and 11 (page 22).)

Define a ***&-assignment*** of a sequent Γ to be a choice of left or right for each of its &s, i.e., a function from the set of &-vertices of Γ to $\{l, r\}$ (l=*left*, r=*right*). Every &-assignment φ defines a &-resolution Γ^φ in the obvious way, by restricting each & to the argument dictated by its assignment (i.e., delete the right (resp. left) argument subtree of w iff $\varphi(w) = l$ (resp. r)). In turn, every &-resolution $\Gamma^\star$ of a sequent Γ induces a &-resolution $\Pi \upharpoonright \Gamma^\star$ of a proof Π of Γ: work upwards from the concluding rule of Π and delete branches of &-rules according to which branch of the corresponding &-occurrence is deleted in $\Gamma^\star$. Note that more than one &-assignment can give rise to the same &-resolution of the sequent Γ, and that more than one &-resolution of Γ can give rise to the same &-resolution of a proof Π of Γ: see Figure 11.

4.5 Intuition for the MLL condition

Every &-resolution R of a proof Π has all additive rules unary. (The $\oplus$ rules are unary at the outset, and the & rules become unary upon taking the &-resolution.) Collapsing the unary additive rules of R (and the now-unary connectives in the corresponding formula parse trees) yields an MLL proof. Since every linking of θ_Π comes from a &-resolution of Π, i.e., from a disguised MLL proof, we demand that every linking of a MALL proof net be MLL correct.

4.6 Intuition for the toggling condition

In the preceding subsections we saw how a cut-free MALL proof Π determines a set of cut-free MLL proofs, one per &-resolution of Π. However, Π is more than just a set of non-interacting MLL proofs, as each of them is implicitly embedded inside the tree of Π. Correspondingly, a set of linkings merely satisfying the resolution and MLL conditions need not be sequentialisable, as one must capture the constraint associated with the superposition of branches of the &-resolutions of Π inside the tree structure of Π. We have already seen an example: the pair of linkings $\Lambda = \{\lambda_1, \lambda_2\}$

$$\begin{array}{ll} \lambda_1 : & \\ & P^\perp \oplus (Q \oplus P^\perp),\ (P \& P) \otimes (R \oplus R),\ (R^\perp \otimes R) ⅋ R^\perp \\ \lambda_2 : & \end{array}$$

of Figure 10(a) (page 18) satisfies the resolution condition (verified in Example 4.6, page 14) and the MLL condition (Example 4.12, page 15), but Λ is not sequentialisable. It fails to sequentialise because we cannot write down a rule to introduce the central tensor: its left argument $P \& P$ must go in the left hypothesis of the

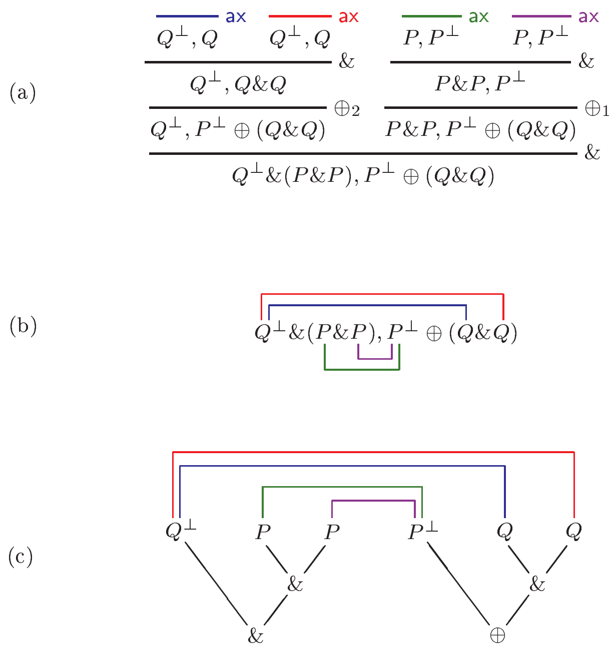

Fig. 11. (a) A proof Π of a sequent Γ illustrating a collapse from &-assignments of Γ to &-resolutions of Γ to &-resolutions of Π. The sequent Γ has $2^3 = 8$ &-assignments, more than its $3 \times 2 = 6$ &-resolutions, more than the 4 &-resolutions of Π. (b) The set of linkings associated with Π, one from each of its &-resolutions. It is convenient to show all four linkings on the same copy of the sequent; no ambiguity arises because every linking has only one link. (c) For additional clarity, we show the same set of four singleton linkings displayed on the parse trees of the two formulas (i.e., we show the union of the graphs $\mathcal{G}_\lambda$ for each of the four linkings λ).

rule, and its right argument $R \oplus R$ in the right hypothesis; but then the & will not be available in the right branch to superimpose a left-$\oplus$ and right-$\oplus$ rule as would be required to obtain λ_1 with the left R of $R \oplus R$ and λ_2 with the right R.

There is a conflict between the central $\otimes$ and the &: the tensor wishes to separate its & argument from its $\oplus$ argument, into distinct non-interacting proofs; meanwhile the & argument interacts with the $\oplus$ argument since in the λ_i the $\oplus$ goes left iff the & goes left, a direct dependency (interaction) across the tensor. Via jumps, the toggling condition captures this kind of dependency, and rules out Λ as a proof net: the graph $\mathcal{G}_\Lambda$ (Figure 10(b), page 18) of Λ contains the switching cycle

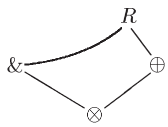

(copied from Figure 10(c)) traversing the only &, and therefore breaking the toggling condition. The jump captures the communication between the & and the $\oplus$.

In contrast, the pair of linkings $\theta = \{\lambda_1, \lambda_2'\}$

$$\begin{array}{ll} \lambda_1 : & \\ & P^{\perp} \oplus (Q \oplus P^{\perp}),\ (P \& P) \otimes (R \oplus R),\ (R^{\perp} \otimes R) \parr R^{\perp} \\ \lambda_2' : & \end{array}$$

of Figure 10(e) (page 18), formed from Λ by shifting the end of just one link in λ_2 to create λ_2', *is* sequentialisable: after writing down the rule introducing the central tensor, choose the left-$\oplus$ rule for $R \oplus R$, since both linkings λ_1 and λ_2' choose the left argument of $R \oplus R$. In contrast to the case $\Lambda = \{\lambda_1, \lambda_2\}$ above, there is no communication across the tensor between the & and the $\oplus$. Thus the graph $\mathcal{G}_\theta$ of θ (Figure 10(f), page 18) does not possess a jump across the tensor, hence θ satisfies the toggling condition, and is a proof net (verified in Example 4.16, page 19).

4.6.1 *The Gustave example.* In the previous section we saw sequentialisation hampered by a dependency across a tensor, and how this was captured by the toggling condition. In this section we examine a case of dependency across the other unswitched connective, the plus, yielding additional intuition for the toggling condition.

The following remarkable non-sequential function γ is due to Gustave, studied by G. Berry in the context of sequential algorithms:

$$\begin{aligned} \gamma(1,0,z) &= u \\ \gamma(0,y,1) &= v \\ \gamma(x,1,0) &= w \end{aligned}$$

for all x, y, and z, possibly divergent/halting. The actual outputs u, v, w are of no concern, so long as they are non-divergent. Both $\gamma(1,1,1)$ and $\gamma(0,0,0)$ are divergent. This partial function cannot be implemented sequentially. For example, suppose our implementation inspects the argument x of $\gamma(x,y,z)$ first. If x diverges, while $y = 1$ and $z = 0$, the equations for γ dictate an output w; however, our implementation would diverge, having become stuck on the divergent x. By symmetry, we cannot choose to inspect y or z first either, hence there is no sequential implementation of γ.

Girard (e.g. [1999, sec. 5.5.4]) and Abramsky and Melliès [1999] have studied a corresponding example in the context of models of linear logic. Analogous to Girard/Abramsky/Melliès, in our setting one can capture the three equations specifying γ as part of a set of five linkings on the sequent

$$(P \& Q) \oplus (P^{\perp} \otimes Q^{\perp}), \quad (Q \& (P^{\perp} \otimes Q^{\perp})) \oplus P, \quad ((P^{\perp} \otimes Q^{\perp}) \& P) \oplus Q,$$

satisfying the resolution and toggling conditions. To emphasise the rotational symmetry, write R for $P^{\perp} \otimes Q^{\perp}$, so that the sequent becomes the more palatable

$$\Gamma \ = \ (P \& Q) \oplus R, \quad (Q \& R) \oplus P, \quad (R \& P) \oplus Q$$

and write the "triplet" linking

$$P \qquad R \qquad Q$$

to abbreviate a pair of links

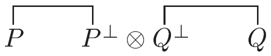

The ***Gustave proof structure***[9] G consists of the following five linkings on Γ, three shown above, and two below.

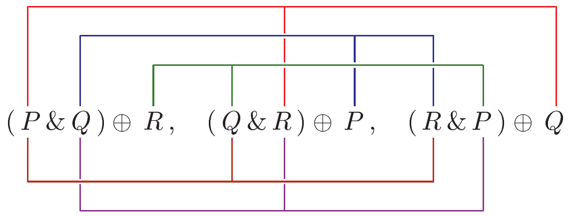

From left to right, the three &s correspond to the arguments x, y, z of the Gustave function γ specified on page 23. Values 1/0 for x, y, z correspond to the &s being left/right, respectively. Thus the eight possible (non-divergent) inputs to the Gustave function correspond to the eight &-resolutions of the sequent. The top three linkings correspond to the three Gustave equations, in order, from top to bottom. For example, the top linking takes the first & left, the second & right, and is ambivalent to the third &; this corresponds to the equation for $\gamma(1, 0, z)$. The two underhanging linkings correspond to the divergent $\gamma(1, 1, 1)$ and $\gamma(0, 0, 0)$, and are added so that the resolution condition holds. (One can readily verify the resolution condition by working through each of the eight &-resolutions and checking that exactly one linking fits in each case.) The MLL property holds since every linking induces the same MLL proof net, the pair of links displayed immediately prior to the five Gustave linkings.

The Gustave proof structure is not the translation of any cut-free proof: any proof of Γ must end in a final $\oplus$-rule (a simple syntactic observation), hence any translation of a proof of Γ has at least one of the six $\oplus$-arguments uninhabited (corresponding to *softness* [Joyal 1995]); G touches all six arguments. Thus, by the sequentialisation theorem, we should be able to witness the failure of the toggling condition. This is indeed the case, since every & is contained in the following switching cycle of the graph of G:

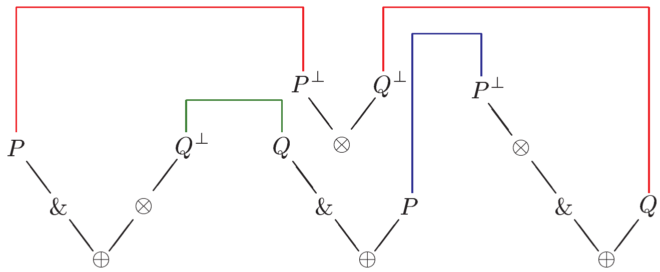

(Note that we did not require jumps to forge this switching cycle.)

[9] The corresponding structure in Girard's setting is not a proof structure. See the end of Section 6.3 (page 62) for a direct verification, or footnote 30 (page 63) which shows that every Girard proof structure must be soft.

4.6.2 *Strong—but not too strong.* We have seen from the previous two sections that the toggling condition captures unwanted dependencies across ⊗s and ⊕s by finding switching cycles through &s. In this section we illustrate a subtle feature of the toggling condition: it is possible to have switching cycles through &s *without* obstructing sequentialisability. One has to be extremely discerning of those switching cycles through &s which are *essential*, in the sense that they represent intrinsic parallelism, versus those which are *harmless*, and do not impair sequentialisation.

We shall define a proof net θ on a sequent Γ of three &s, with $2^3 = 8$ linkings (one per &-assignment of three independent &s). The graph $\mathcal{G}_\theta$ of θ will have the shape

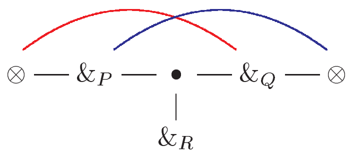

with the portion

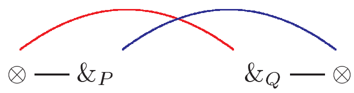

constituting a switching cycle of $\mathcal{G}_\theta$ between the &-occurrences $\&_P$ and $\&_Q$.

The underlying sequent will be

$$(1 \oplus 1) \otimes (P \& P),\ (1 \oplus 1) \otimes (Q \& Q),\ (P^\perp \otimes Q^\perp) \otimes R^\perp,\ R \& R$$

where 1 denotes the tensor unit. (As with MLL proof nets, our MALL proof nets extend trivially to the tensor unit 1: view each occurrence of 1 as $(P \parr P^\perp)$ for a fresh atom P each time.[10]) Links on the atoms P, Q, R are forced, so to determine the eight linkings of θ it suffices to specify how the ⊕s choose their 1s:

(a) in $(1 \oplus 1) \otimes (P \& P)$ choose the right 1 iff $R \& R$ is left and $Q \& Q$ is right;

(b) in $(1 \oplus 1) \otimes (Q \& Q)$ choose the right 1 iff $R \& R$ is right and $P \& P$ is right.

The set of linkings θ is sequentialisable. (Figure 12 shows a sequentialisation.) The graph $\mathcal{G}_\theta$ of θ contains the following switching cycle C:

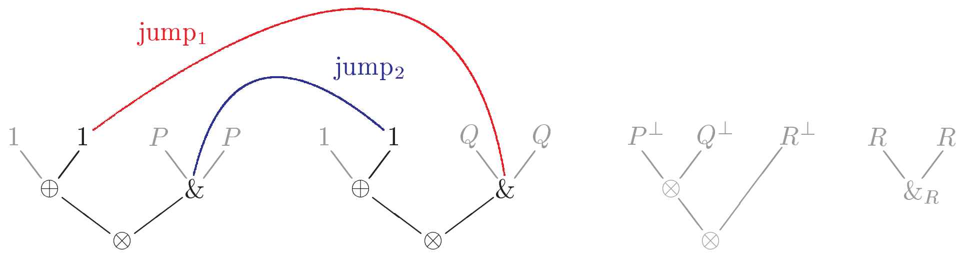

Due to the way we defined the linkings in clauses (a) and (b) above, given a set Λ of two or more linkings of θ:

[10]Similarly, one can define the plus unit 0 as $P^\perp \otimes P$ for a fresh atom P each time. Note the interesting complementarity between 0 and 1, units which are not dual in the logic.

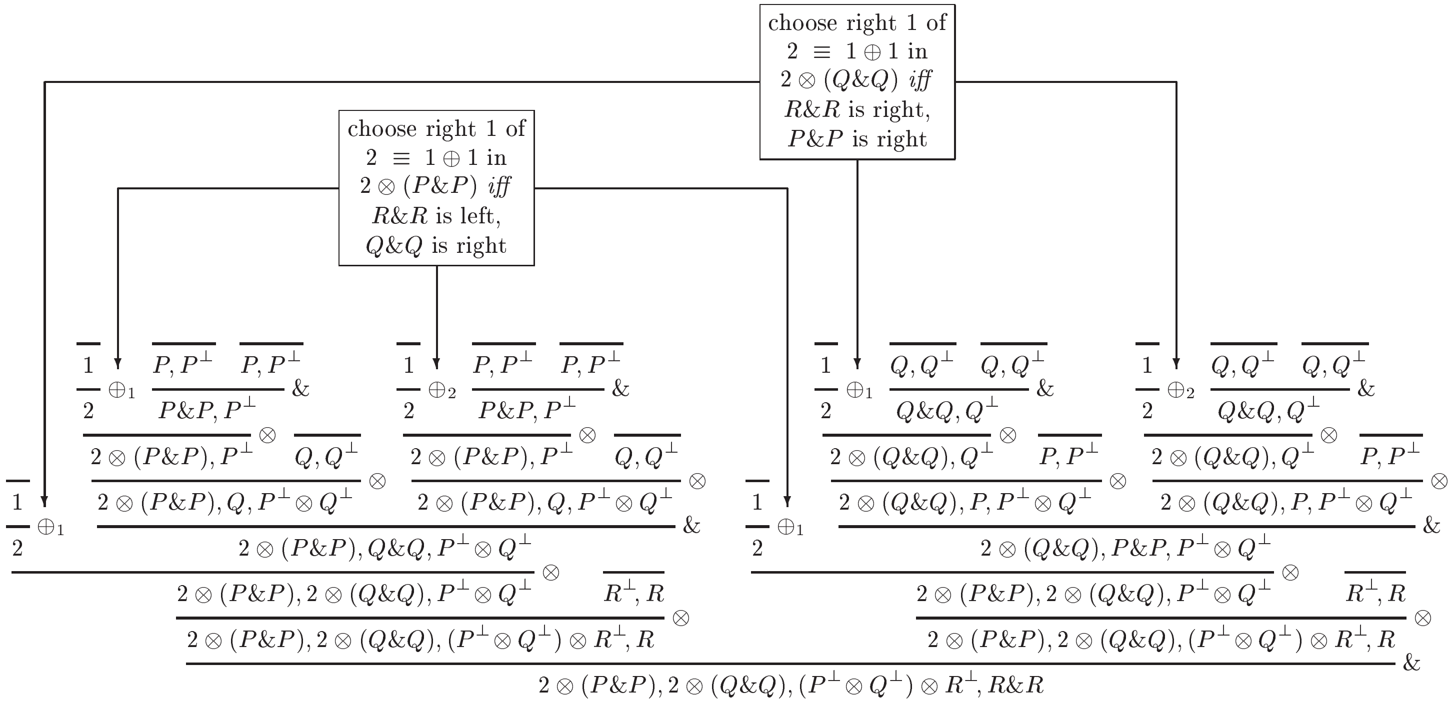

Fig. 12. A sequentialisation of θ.

(1) jump_1 exists in $\mathcal{G}_\Lambda$ only if $\&_R$ possesses its left argument in $\mathcal{G}_\Lambda$, and
(2) jump_2 exists in $\mathcal{G}_\Lambda$ only if $\&_R$ possesses its right argument in $\mathcal{G}_\Lambda$.

Therefore C can be in $\mathcal{G}_\Lambda$ *only if* Λ *toggles* $\&_R$. Since $\&_R$ is outermost, it cannot be in a switching cycle, hence C cannot witness a failure of the toggling condition. We deduce that θ satisfies the toggling condition, and is therefore a proof net. Thus C is *harmless*, in the sense that it does not represent any inherent lack of sequentialisability in θ.

Here we witness the subtlety of the toggling condition at work: it must rule out many switching cycles—but not *too* many.

Proof-theoretic analogue. Additional intuition for the toggling condition follows from analysing the harmless switching cycle above at the proof-theoretic level. The &-rule skeleton of the sequentialisation of θ depicted in Figure 12 is:

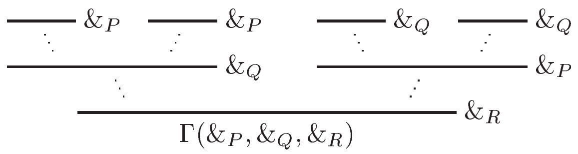

Here each &-rule is marked with the &-vertex it introduces into Γ, e.g., each —— $\&_P$ introduces $\&_P$ (the &-vertex of $P\&P$) into (a subsequent of) Γ.

In the left branch of the proof, $\&_P$ is forced to come above $\&_Q$, and in the right branch, $\&_Q$ is forced to come above $\&_P$, *forced* in the sense that every sequentialisation of θ must have exactly the same &-rule skeleton. The &-rules simply do not commute past each other. Similarly, $\&_R$ is forced to come below $\&_P$ and $\&_Q$. Writing $\&_P \longrightarrow \&_Q$ for "—— $\&_P$ is forced to come above —— $\&_Q$", we derive the following precedence graph:

$$\begin{tikzcd}
\&_P \arrow[rr, shift left] \arrow[dr] & & \&_Q \arrow[ll, shift left] \arrow[dl] \\
& \&_R &
\end{tikzcd}$$

The back-and-forth cycle between $\&_P$ and $\&_Q$ in this graph is the direct analogue of the switching cycle C of $\mathcal{G}_\theta$ analysed earlier. That C is harmless corresponds to the fact that the cycle here is a relic of the superposition of the two branches of the proof: $\&_P \longrightarrow \&_Q$ holds only in the left branch of the proof, and $\&_Q \longrightarrow \&_P$ holds only in the right branch.

4.7 Alternative but equivalent definitions of proof net

This section considers alternative definitions of a proof net obtained by varying (P1) RESOLUTION, (P2) MLL and (P3) TOGGLING.

4.7.1 *Acyclicity, balance, and connectedness.* Say that a linking λ on a MALL sequent Γ is ***balanced*** if $|\mathsf{ax}| = |\otimes| + 1$, where $|\mathsf{ax}|$ denotes the number of links in λ and $|\otimes|$ the number of tensors in the additive resolution $\Gamma \upharpoonright \lambda$. Consider the following properties.

(A) every $⅋$-switching of λ is acyclic	(a) some $⅋$-switching of λ is acyclic
(B) λ is balanced	
(C) every $⅋$-switching of λ is connected	(c) some $⅋$-switching of λ is connected

By definition, the MLL condition (P2) holds for a linking λ precisely when λ satisfies (A)$\wedge$(C).

Proposition 4.20. *The following conditions are all equivalent to the MLL condition* (P2) *on a linking* λ*:* (A)$\wedge$(C)*,* (A)$\wedge$(c)*,* (a)$\wedge$(C)*,* (A)$\wedge$(B) *and* (B)$\wedge$(C)*.*

The proof is essentially due to the simple combinatorial relationship between the number of vertices and the number of edges of a tree. See Appendix B.

4.7.2 *Switching acyclicity and switching connectedness.* It is immediately clear that (A) above is equivalent to $\mathcal{G}_\lambda$ being ***switching acyclic***, that is, containing no switching cycle. In the presence of (A), condition (C) is equivalent to $\mathcal{G}_\lambda$ being ***switching connected***, that is, any two vertices of $\mathcal{G}_\lambda$ are connected by a ***switching path***, a path that does not traverse two switch edges of any given $⅋$. Switching connectedness is clearly implied by (C); the equivalence with (C) follows from the observation that one can carry out sequentialisation (specifically, the MLL restriction of the proof of the sequentialisation theorem) with this condition in place of (C).[11] Thus we have proved:

Proposition 4.21. *The* MLL *condition* (P2) *on a set of linkings* θ *is equivalent to:*

(S) *For every linking* $\lambda \in \theta$*, the graph* $\mathcal{G}_\lambda$ *is switching acyclic and switching connected.*

4.7.3 *Illegal unions of switching cycles.* We provide an alternative formulation of the toggling condition (P3), assuming the MLL condition (P2). Call a union S of switching cycles of $\mathcal{G}_\theta$ ***illegal*** if it is non-empty and for some $\Lambda \subseteq \theta$ with $S \subseteq \mathcal{G}_\Lambda$, every & toggled by Λ is in S.

($\mathrm{P3}^l$) $\mathcal{G}_\theta$ contains no illegal union of switching cycles.

Note that this condition implies condition (A) for each linking (every $⅋$-switching is acyclic). The proof of equivalence with (P3) follows from simple manipulation using Proposition 4.20 above. Details are in Appendix C.

Conjecture 4.22 Single switching cycle conjecture. *Property* ($\mathrm{P3}^l$) *is equivalent to:*

($\mathrm{P3}^{l-}$) $\mathcal{G}_\theta$ *contains no illegal switching cycle.*

In other words, the original toggling condition (P3) *is equivalent to:*

($\mathrm{P3}^-$) *For any set* Λ *of two or more linkings of* θ *and any switching cycle* C *of* $\mathcal{G}_\Lambda$*,* Λ *toggles a & that is not in* C*.*

[11] The three subcases of the primary induction step on page 38 use the fact that [θ satisfies (P2)] implies [θ_i (or θ on Γ', in case (a)) satisfies (P2)]. This implication also holds for the variant of (P2) with switching connectedness instead of (C). There are three other places in the primary and secondary induction of the sequentialisation proof where (C) is used, listed in Footnote 21 on page 39; in each case, the property derived is also a consequence of switching connectedness.

4.7.4 *Additional jumps.* We shall use the following variation of the MLL condition (P2) in comparing Girard's proof nets to ours in Section 6.4. Given a set of linkings θ on a sequent Γ and a subset $\Lambda \subseteq \theta$, let $\mathcal{G}^{\theta}_{\Lambda}$ be defined as $\mathcal{G}_{\Lambda}$ but with jump edges between every &-vertex $w \in \mathcal{G}_{\Lambda}$ and the leaves of every link $a \in \mathcal{G}_{\Lambda}$ depending on w in θ (rather than in Λ, as in the definition of $\mathcal{G}_{\Lambda}$). Note that $\mathcal{G}_{\Lambda} = \mathcal{G}^{\Lambda}_{\Lambda}$. Define the variant (P2*) of (P2) by using $\mathcal{G}^{\theta}_{\{\lambda\}}$ in place of $\mathcal{G}_{\lambda}$ in the definition of a $⅋$-switching of λ, and in taking the switching delete in addition all but one switch edge of each & (i.e., we move from $⅋$-switchings to "$⅋$/&-switchings"). Clearly (P2*) implies (P2), since it involves more switchings. In fact, (P2*) is strictly stronger than (P2): for $\theta = \{\lambda_1, \lambda_2\}$ of Example 4.6 (page 14), the graph $\mathcal{G}^{\theta}_{\lambda_1}$ has a switching cycle (cycle C in Figure 10(c), page 18), whereas $\mathcal{G}_{\lambda_1}$ (Figure 9(e), page 16) does not. However, (P2*) is implied by the MLL condition (P2) and the toggling condition (P3) together:

PROPOSITION 4.23. (P2) $\wedge$ (P3) $\implies$ (P2*).

PROOF. Let θ be a set of linkings satisfying (P2) and (P3), and let $\lambda \in \theta$. By (P2), λ is balanced. It suffices to show that $\mathcal{G}^{\theta}_{\lambda}$ has no switching cycle, for this implies that every $⅋$/&-switching of λ within $\mathcal{G}^{\theta}_{\lambda}$ is acyclic, and hence also connected, by (the proof of) Proposition 4.20.

Towards a contradiction, assume C is a switching cycle of $\mathcal{G}^{\theta}_{\lambda}$. If C does not contain a jump edge, it is a switching cycle of $\mathcal{G}_{\lambda}$, contradicting (P2). Otherwise, let Λ be the largest set of linkings in θ containing λ and toggling only &s occurring in C. For every jump edge in C from a leaf to a &-vertex w, there is a linking $\lambda' \in \theta$ such that w is the only & toggled by $\{\lambda, \lambda'\}$. Hence $\lambda' \in \Lambda$. Thus, all jumps in C are also present in $\mathcal{G}_{\Lambda}$, so C is a switching cycle of $\mathcal{G}_{\Lambda}$ containing all &s toggled by Λ. Since $|\Lambda| \geq 2$, this contradicts (P3). □

We could also define a variant (P3*) of (P3) with more jumps, using $\mathcal{G}^{\theta}_{\Lambda}$ instead of $\mathcal{G}_{\Lambda}$. By an argument similar to the one above, this variant is equivalent to (P3).

4.7.5 *Other variations.* In Section 6.3 we develop a correspondence between the resolution condition (P1) and Girard's so-called technical condition [1996]. We also present alternative formulations of (P1) and the technical condition, and note that, without monomials, the Abramsky-Melliès reformulation [1999] of the technical condition is no longer valid.

4.8 Weights

This section describes how to encode any proof structure (hence any proof net) as a single set of links labelled with predicates, called *weights* (c.f. [Girard 1996]). Figure 5 (page 7) conveys the idea informally with an example.

Recall from Section 4.4 that a &-assignment of a sequent Γ is a function from its &-vertices to $\{l, r\}$ (l=*left*, r=*right*), and that every &-assignment φ defines a &-resolution Γ^{φ} by restricting each & to the argument dictated by φ. Multiple &-assignments can determine the same &-resolution. For example, if $\Gamma = (P \&_1 Q) \&_2 R$, then the assignments $\&_1 \mapsto l$, $\&_2 \mapsto r$ and $\&_1 \mapsto r$, $\&_2 \mapsto r$ both determine the &-resolution $(P \&_1 Q) \&_2 R$ retaining only R. See also Figure 11 (page 22) for more on the relationship between &-assignments and &-resolutions.

Let θ be a proof structure on Γ. Given a &-assignment φ of Γ, write λ_{φ} for the unique linking of θ which is on the &-resolution Γ^{φ} of φ (existence and uniqueness due to the resolution condition (P1)). Every link a of θ determines a predicate on &-assignments, its ***weight*** $\mu(a)$, by $\varphi \in \mu(a)$ *iff* $a \in \lambda_{\varphi}$. One can then represent θ by its links labelled with weights, as in Figure 5 (page 7), for example.

Weights can be expressed succinctly as follows. First, mark each &-vertex with a distinct subscript, x, y, ... Write x as shorthand for $\{\,\varphi : \varphi(\&_x) = l\,\}$ (all &-assignments that take $\&_x$ to the left) and $\overline{x}$ as shorthand for $\{\,\varphi : \varphi(\&_x) = r\,\}$ (all &-assignments that take $\&_x$ to the right); $\vee$ and $\wedge$ are union and intersection, respectively. Again, see Figure 5 for an example.

The set of linkings of a proof structure is recoverable from its weight presentation as follows. Every &-assignment φ determines a linking λ_{φ} by deleting each link a whose predicate does not hold, i.e. $\lambda_{\varphi} = \{a : \varphi \in \mu(a)\}$. Taking each &-assignment in turn produces the full set of linkings.

4.9 Mix nets

Let MALL$^{\mathsf{mix}}$ denote the extension [Girard 1987] of MALL with the additional rule

$$\frac{\Gamma \quad \Delta}{\Gamma, \Delta}\,\mathsf{mix}$$

and define the following variant of the MLL condition on a set of linkings θ by relaxing connectedness:

(P2$^{\mathsf{mix}}$) MLL$^{\mathsf{mix}}$. *Every* $⅋$*-switching of every linking of* θ *is acyclic.*

A ***cut-free mix net*** is a cut-free MALL proof net but for relaxing connectedness of $⅋$-switchings, i.e., a set of linkings satisfying (P1) RESOLUTION, (P2$^{\mathsf{mix}}$) MLL$^{\mathsf{mix}}$, and (P3) TOGGLING.

THEOREM 4.24 CUT-FREE MIX SEQUENTIALISATION. *A set of linkings is the translation of a cut-free MALL*$^{\mathsf{mix}}$ *proof iff it is a cut-free mix net.*

We prove this theorem concurrently with the main sequentialisation theorem. Only very minor modifications are necessary.

Proof nets for MLL with mix and weakening were discovered prior even to linear logic [Ketonen and Weyhrauch 1984]. (Bellin and Ketonen [1992] correct a bug in the proof of the sequentialisation theorem.)

4.10 The resolution condition suffices for pure additive proof nets

The RESOLUTION condition, on its own, suffices as a correctness criterion for pure additive proof nets. Let additive linear logic, ALL, be MALL without $\otimes$ and $⅋$. Every ALL sequent has exactly two formulas. When a cut-free ALL proof translates into a set of linkings, every linking is merely a single link between the two formulas of the sequent. Thus every cut-free ALL proof Π of the sequent $\Gamma = A, B$ translates into a set L of links between A and B, a binary relation between the leaves of A and the leaves of B. In this simple pure additive case, the RESOLUTION condition for L on Γ reduces to:

— RESOLUTION$'$. *For any &-resolution* $\Gamma^{\star}$ *of* Γ, *a unique link of* L *is on* $\Gamma^{\star}$.

$$\dfrac{\dfrac{\Pi_1}{\Gamma, A} \qquad \dfrac{\dfrac{\Pi_2}{B, \Delta, X} \quad \dfrac{\Pi_3}{B, \Delta, Y}}{B, \Delta, X \& Y}\,\&}{\Gamma, A \otimes B, \Delta, X \& Y}\,\otimes \qquad \longleftrightarrow \qquad \dfrac{\dfrac{\dfrac{\Pi_1}{\Gamma, A} \quad \dfrac{\Pi_2}{B, \Delta, X}}{\Gamma, A \otimes B, \Delta, X}\,\otimes \qquad \dfrac{\dfrac{\Pi_1}{\Gamma, A} \quad \dfrac{\Pi_3}{B, \Delta, Y}}{\Gamma, A \otimes B, \Delta, Y}\,\otimes}{\Gamma, A \otimes B, \Delta, X \& Y}\,\&$$

$$\dfrac{\dfrac{\dfrac{\Pi_1}{\Gamma, A, X} \quad \dfrac{\Pi_2}{\Gamma, B, X}}{\Gamma, A \& B, X}\,\& \qquad \dfrac{\dfrac{\Pi_3}{\Gamma, A, Y} \quad \dfrac{\Pi_4}{\Gamma, B, Y}}{\Gamma, A \& B, Y}\,\&}{\Gamma, A \& B, X \& Y}\,\& \qquad \longleftrightarrow \qquad \dfrac{\dfrac{\dfrac{\Pi_1}{\Gamma, A, X} \quad \dfrac{\Pi_3}{\Gamma, A, Y}}{\Gamma, A, X \& Y}\,\& \qquad \dfrac{\dfrac{\Pi_2}{\Gamma, B, X} \quad \dfrac{\Pi_4}{\Gamma, B, Y}}{\Gamma, B, X \& Y}\,\&}{\Gamma, A \& B, X \& Y}\,\&$$

Fig. 13. Two examples of rule commutation. The commutations can be read in either direction.

This yields a proof net for cut-free ALL: by a simple induction, the condition characterises the image of the translation from cut-free ALL proofs.[12] The category of cut-free ALL proof nets is the free (binary) product-sum category generated by the set of literals [Hughes 2002; 2005]. Relaxing uniqueness in RESOLUTION′ characterises free distributive lattice categories[13] [Hughes 2005], and (also relaxing the inter-formula restriction on links) captures the image of proofs in classical propositional sequent calculus with mix (translated in the obvious way) [Lamarche and Straßburger 2005]. For abstract classical proofs with a richer graph-theoretic structure on axiom links, rather than simply a set (or multiset) of axiom links, see [Hughes 2004].

4.11 Representation of cut-free proofs modulo rule commutation

The kernel of our function from cut-free MALL proofs to sets of linkings coincides precisely with equivalence modulo rule commutation. A rule commutation is a local conversion on a proof that retains the subproofs of its hypotheses, with possible duplication/identification. Figure 13 shows two examples of rule commutation.

In a sibling paper we prove that two cut-free MALL proofs translate to the same proof net if and only if they can be converted into each other by a series of rule commutations. The same paper explores other aspects of rule commutation in MALL (with/without the mix rule, with/without the cut rule).

4.12 Proof of the Separation Lemma

This section proves the Separation Lemma (Lemma 4.19, page 20), the key to the Sequentialisation Theorem.

Throughout this section θ is a cut-free proof net on a sequent Γ. For vertices x

[12] Using softness: given an ALL proof net on $A \oplus B, C \oplus D$ one can apply a $\oplus$-rule; otherwise there are edges A–C and B–D (or A–D and B–C), contradicting uniqueness in RESOLUTION′. Composition (see Section 5.2, page 43) is also simple in the special case of ALL proof nets: it reduces to the standard path composition of binary relations.

[13] Došen and Petrić define a distributive lattice category as a product-sum category with a distribution, equipped with certain coherence laws [Došen and Petrić 2004].

and y of the graph $\mathcal{G}_\theta$, write $x — y$ if there is an edge between x and y, and write $x \to y$ iff x is an argument of y, $\{x, y\}$ is a link[14], or there is a jump from x to y (i.e., x is a leaf of a link depending on a &-vertex y of θ).

Henceforth "$\parr$/&" abbreviates "$\parr$ or &". A ***path*** from x_0 to x_n in $\mathcal{G}_\theta$ is a sequence of distinct vertices $x_0 x_1 \ldots x_n$ ($n = 0$ permitted) such that $x_i — x_{i+1}$ for $0 \le i < n$. (Note that a path cannot intersect itself.) A path ***switches*** or ***is switching*** if it does not traverse two switch edges of any $\parr$/& (i.e., $x_{i-1} \to x_i \leftarrow x_{i+1}$ only if x_i is not a $\parr$/&.) A ***strong path*** $x_0 \ldots x_n$ is a switching path which does not start from a $\parr$/& along one of its switch edges (i.e., $x_0 \leftarrow x_1$ only if x_0 is not a $\parr$/&).

Suppose paths $\pi = x_0 \ldots x_n$ and $\pi' = y_0 \ldots y_m$ are disjoint but for $x_n = y_0$, so that the composite $\pi ; \pi' = x_0 \ldots x_n y_1 \ldots y_m$ is a well-defined path (non self-intersecting). If π and π' switch:

- $\pi ; \pi'$ need not switch (namely if $x_n = y_0$ is a &/$\parr$ and $x_{n-1} \to x_n = y_0 \leftarrow y_1$), even if π is strong.
- if π' is strong, then $\pi ; \pi'$ switches.
- if π and π' are strong then $\pi ; \pi'$ is strong.

Let X be a set of vertices in $\mathcal{G}_\theta$. A path is ***in*** X if each of its vertices is in X. Write $x \Rightarrow_X y$ (and/or $y \Leftarrow_X x$) if there is a strong path in X from x to y.

Example 4.25. If C is a switching cycle then $x \Rightarrow_C y$ for all $x, y \in C$ (case $x = y$ included), by going round C one way or the other to avoid departing along a switch edge of x, if x is a $\parr$/&.

Note that the relation $\Rightarrow_X$ is reflexive, but in general not transitive.[15] We shall sometimes overload the notation $x \Rightarrow_X y$, using it to denote a specific choice of strong path in X from x to y. For example, if $x \Rightarrow_X y$ and $y \Rightarrow_Y z$, with X and Y disjoint but for y, then we may speak of *the* strong path $x \Rightarrow_X y \Rightarrow_Y z$ in $X \cup Y$ from x to z.

A set X of vertices in $\mathcal{G}_\theta$ is an ***x-zone*** if, for all $y \in X$, there exists $z \in X$ with $y \Rightarrow_X z \to x$.

Example 4.26. Let x be a vertex in a switching cycle C. Then C is an x-zone: let z be a vertex adjacent to x on C with $z \to x$ (uniquely determined if x is a $\parr$/&, since C switches), then $y \Rightarrow_C z$ for any $y \in C$ (see Example 4.25).

Given a $\parr$/&-vertex x and a vertex y, define x ***dominates*** y, denoted $x \sqsupset y$, if y is in an x-zone. If x is not dominated, it is ***free***.[16]

Lemma 4.27 Properties of domination.

- switch. *If $x \leftarrow y$ is a switch edge then $x \sqsupset y$.*
- transitivity. *Domination is transitive.*

---

[14] Note that if $\{x, y\}$ is a link then $x \leftrightarrow y$, i.e., $x \to y$ and $x \leftarrow y$.

[15] If $x \to p \leftarrow y$ and $p \to t$, p a $\parr$ and t a $\otimes$, and $X = \{x, p, y, t\}$, then $x \Rightarrow_X t \Rightarrow_X y$ yet $x \not\Rightarrow_X y$.

[16] The union of all x-zones is itself an x-zone, which we call the ***realm*** of x, a concept reminiscent of the notion of ***empire*** of [Girard 1996], but different in an essential way. The realm of x is the set of all vertices dominated by x.

– SELF. *A* $⅋/\&$*-vertex dominates itself iff it is in a switching cycle.*
– JUMP-CYCLE. *If* $w \leftarrow l$ *is a jump and* l *is in a switching cycle* C*, then* w *dominates every vertex of* C*.*
– EXTEND. *If* $x \sqsupset y_0$ *and there is a path* $y_0 \ldots y_n$ *which never enters a* $⅋/\&$ *from above (i.e.,* $y_{i-1} \rightarrow y_i$ *only if* y_i *is not a* $⅋/\&$*), then* $x \sqsupset y_n$*.*
– FORK. *Let* x *be a* $⅋/\&$ *and let* $y_0 \ldots y_n$ *be a switching path with* $y_0 \rightarrow x \leftarrow y_n$*. Then* $x \sqsupset y_i$ *for each* i*.*
– MEET. *If* $x \sqsupset y \sqsubset z$ *for distinct free* $⅋/\&$*-vertices* x *and* z*, then there exists a switching path* $xy_0 \ldots y_n z$ *with* $x \leftarrow y_0$ *and* $y_n \rightarrow z$*.*

PROOF. SWITCH. $\{y\}$ is an x-zone.

TRANSITIVITY. We show that if X is an x-zone, $y \in X$ and Y is a y-zone, then $X \cup Y$ is an x-zone. Take $z \in Y \setminus X$. We have $z \Rightarrow_Y y' \rightarrow y \Rightarrow_X x' \rightarrow x$ for some $x' \in X$ and $y' \in Y$. If the strong path $z \Rightarrow_Y y'$ does not intersect X, then $z \Rightarrow_Y y' \rightarrow y \Rightarrow_X x'$ is a strong path, so we are done. Otherwise let y'' be the first vertex along $z \Rightarrow_Y y'$ that is in X. Since $y'' \in X$ we have $y'' \Rightarrow_X x'' \rightarrow x$ for some x'', and the initial sub-path of $z \Rightarrow_Y y'$ from z to y'' is a strong path $z \Rightarrow_Y y''$; the composition of these paths yields $z \Rightarrow_{X \cup Y} x'' \rightarrow x$, since the only common vertex is y''.

SELF. If $x \sqsupset x$ then $x \Rightarrow_X z \rightarrow x$ for some x-zone X, hence x is in a switching cycle. Conversely, every switching cycle containing x is an x-zone (see Example 4.26).

JUMP-CYCLE. C is a w-zone. (See Example 4.25.)

EXTEND. Let X be an x-zone containing y_0, and let y_k be the last vertex of $y_0 \ldots y_n$ in X. Then $y_k \Rightarrow_X z \rightarrow x$ for some z. Now $Y = X \cup \{y_{k+1}, ..., y_n\}$ is an x-zone, since for each $i > k$ the composite $y_i y_{i-1} \ldots y_{k+1} y_k \Rightarrow_X z$ is a strong path in Y.

FORK. $\{y_0, \ldots, y_n\}$ is an x-zone.

MEET. Let X be an x-zone containing y, so there is a strong path $\pi_x = x_0 \ldots x_n$ in X with $x_0 = y$ and $x_n \rightarrow x$. Let x_k be the last vertex of π_x with $z \sqsupset x_k$. Since $z \sqsupset x_k$ there is a strong path π_z in a z-zone from x_k to some z' with $z' \rightarrow z$. Now $x x_n x_{n-1} \ldots x_{k+1} \pi_z z$ is the desired switching path, well-defined because: (a) every vertex is distinct (none of the included x_i is in π_z, since $z \not\sqsupset x_i$ and z dominates all of π_z (because π_z is in a z-zone); none of the x_i equals x or z, since $x \sqsupset x_i$ (because π_x is in an x-zone) and x and z are free; neither x nor z is in π_z, since every vertex of π_z is dominated by z, and x and z are free), and (b) the path $x_n x_{n-1} \ldots x_{k+1} \pi_z$ switches (since $x_n x_{n-1} \ldots x_k$ switches and π_z is strong). □

Figure 14 shows the dependency between the above domination properties and the forthcoming lemmas (and one corollary) en route to the Separation Lemma. We do not use any properties of domination other than the seven shown in the figure (those of Lemma 4.27).

A subset $\Lambda \subseteq \theta$ is ***saturated*** if any strictly larger subset of θ toggles more &s than Λ. Clearly θ itself is saturated. For Λ a set of linkings and w a & of Γ let Λ^w denote the set of all linkings in Λ whose additive resolution does not contain the right argument of w. Write $\lambda \overset{w}{=} \lambda'$ if linkings $\lambda, \lambda' \in \theta$ are either equal or w is the only & toggled by $\{\lambda, \lambda'\}$. It is straightforward to check that:

(S1) If Λ is saturated and toggles w then Λ^w is saturated.

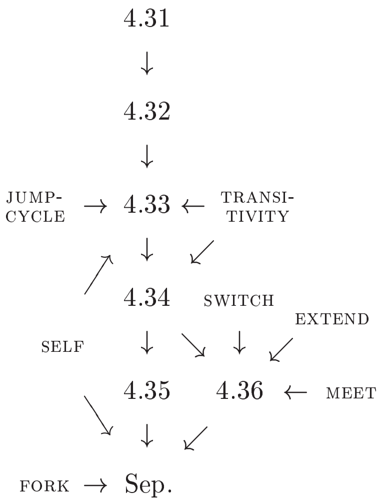

4.31	$\mathcal{G}_\Lambda \ni l \overset{e}{\to} \notin \mathcal{G}_{\Lambda^w} \;\Rightarrow\; l \to w$ in $\mathcal{G}_\Lambda$
4.32	$\exists$ jump out of $\cup$ switching cycles
4.33	$x \sqsupset x \;\Rightarrow\; \exists\, y \not\sqsupset y \sqsupset x$
4.34	every ⅋/& is free or $\sqsubset$ by free
4.35	$\exists$ ⅋/& $\;\Rightarrow\;$ $\exists$ free ⅋/&
4.36	x free, $\neg$ sep. $\;\Rightarrow\; y \to\leftarrow x \leftarrow\to z$
Sep.	$\exists$ ⅋/& $\;\Rightarrow\;$ $\exists$ separating ⅋/&

Fig. 14. Dependency between domination properties and lemmas (and Corollary 4.35) en route to the Separation Lemma, denoted "Sep." above. A rough mnemonic guide is shown to the right of the diagram.

(S2) If Λ is saturated and toggles w and $\lambda \in \Lambda$ then $\lambda \overset{w}{=} \lambda_w$ for some $\lambda_w \in \Lambda^w$.

(S3) If Λ is saturated and toggles w and $\lambda \overset{x}{=} \lambda'$ for $\lambda, \lambda' \in \Lambda$ then

$$\begin{tikzcd} \lambda \arrow[r, equal, "x"] \arrow[d, equal, "w"'] & \lambda' \arrow[d, equal, "w"] \\ \lambda_w \arrow[r, equal, "x"] & \lambda'_w \end{tikzcd}$$

for some $\lambda_w, \lambda'_w \in \Lambda^w$.

Examples below illustrate (S2) and (S3).

Example 4.28. (S2). Let Λ be the following set of three linkings, each having just one link:

$$(P\&P)\&_w P,\;\; P^\perp$$

Two linkings are shown above and one below, and w is the second &. Λ^w is the top pair of linkings. If λ is the bottom linking, either of the top two linkings suffices for λ_w in (S2).

Example 4.29. (S3). Let Λ be the following set of four linkings, each having just one link:

$$P\,\&_x P,\;\; P^\perp \&_w P^\perp$$

Call the linkings rr, rl, ll, lr, from top to bottom. The left & is x, and the right & is w, thus $\Lambda^w = \{rl, ll\}$. Here are two possible instances of the square in (S3):

$$\begin{tikzcd} rr \arrow[r, equal, "x"] \arrow[d, equal, "w"'] & lr \arrow[d, equal, "w"] \\ rl \arrow[r, equal, "x"] & ll \end{tikzcd} \qquad \begin{tikzcd} rr \arrow[r, equal, "x"] \arrow[d, equal, "w"'] & rr \arrow[d, equal, "w"] \\ rl \arrow[r, equal, "x"] & rl \end{tikzcd}$$

Example 4.30. Let Λ be the same set of three linkings as in example 4.28:

$$(P \&_x P) \&_w P, \quad P^\perp$$

Let $\lambda, \lambda', \lambda''$ be the three (single-link) linkings, from top to bottom. Thus $\Lambda^w = \{\lambda, \lambda'\}$. Here is a degenerate instance of the (S3) square, in which the (suppressed) bottom edge is $\lambda'' = \lambda''$:

$$\begin{tikzcd} \lambda \arrow[rr, equal, "x"] \arrow[dr, equal, "w"'] & & \lambda' \arrow[dl, equal, "w"] \\ & \lambda'' & \end{tikzcd}$$

This illustrates why the definition of $\lambda_1 \overset{x}{=} \lambda_2$ includes equality $\lambda_1 = \lambda_2$.

Lemma 4.31. *Let w be a & toggled by a saturated set $\Lambda \subseteq \theta$, and let e be an edge in $\mathcal{G}_\Lambda$ originating from a leaf l, such that $e \notin \mathcal{G}_{\Lambda^w}$. Then the jump $l \to w$ is in $\mathcal{G}_\Lambda$.*

Proof. Let e be $l \to x$. If e is not a jump, $e \notin \mathcal{G}_{\Lambda^w}$ implies $l \notin \mathcal{G}_{\Lambda^w}$. Choose $\lambda \in \Lambda$ with l a leaf of some link $a \in \lambda$. By (S2) $\lambda \overset{w}{=} \lambda_w$ for some $\lambda_w \in \Lambda^w$. Since $a \notin \lambda_w$ (for $l \notin \mathcal{G}_{\Lambda^w}$), the jump $l \to w$ is in $\mathcal{G}_\Lambda$.

If e is a jump, we have $\lambda, \lambda' \in \Lambda$ with $a \in \lambda$, $a \notin \lambda'$, l a leaf of a, and $\lambda \overset{x}{=} \lambda'$. By (S3) $\lambda \overset{w}{=} \lambda_w \overset{x}{=} \lambda'_w \overset{w}{=} \lambda'$ for $\lambda_w, \lambda'_w \in \Lambda^w$. Either $a \notin \lambda_w$ or $a \in \lambda'_w$, else $e \in \mathcal{G}_{\Lambda^w}$; either way, the jump $l \to w$ is in $\mathcal{G}_\Lambda$. $\square$

Lemma 4.32. *Every non-empty union S of switching cycles of $\mathcal{G}_\theta$ has a jump out of it: for some leaf $l \in S$ and &-vertex $w \notin S$, there is a jump $l \to w$ in $\mathcal{G}_\theta$.*

Proof. Let Λ be a minimal saturated subset of θ with $\mathcal{G}_\Lambda$ containing S. By (P2), ⅋-switchings of singleton subsets of θ are acyclic, so Λ contains at least two linkings. Let w be a & toggled by Λ that is not in any switching cycle of $\mathcal{G}_\Lambda$ (existing by (P3)), so $w \notin S$. Since Λ is minimal, $S \not\subseteq \mathcal{G}_{\Lambda^w}$ (using (S1)), so some edge e of S is in $\mathcal{G}_\Lambda$ but not in $\mathcal{G}_{\Lambda^w}$. Without loss of generality e is an edge from a leaf l, because for any other edge $y \to x$ in S we have $l \to z_1 \to \ldots \to z_n = y \to x$ in S for some leaf l, and $y \to x$ is in $\mathcal{G}_{\Lambda^w}$ whenever $l \to z_1$ is in $\mathcal{G}_{\Lambda^w}$. By Lemma 4.31 the jump $l \to w$ is in $\mathcal{G}_\Lambda$, hence also in $\mathcal{G}_\theta$. $\square$

Lemma 4.33. *If $x \sqsupset x$ then $y \sqsupset x$ for some &-vertex $y \not\sqsupset y$.*

Proof. By domination property self, x is in a switching cycle. Iterate Lemma 4.32, adding switching cycles until jumping to a &-vertex y not in a switching cycle. Then $y \sqsupset x$ by jump-cycle and transitivity, and $y \not\sqsupset y$ by self. $\square$

Lemma 4.34. *Every $\parr/\&$ of $\mathcal{G}_\theta$ is either free or is dominated by a free $\parr/\&$.*[17]

Proof. If x_0 is neither free nor dominated by a free $\parr/\&$-vertex, then we can build an infinite chain $x_0 \sqsubset x_1 \sqsubset \ldots$ of distinct vertices with the same property. If $x_i \sqsupset x_i$, obtain $x_{i+1} \not\sqsupset x_{i+1} \sqsupset x_i$ from Lemma 4.33; x_{i+1} is fresh otherwise $x_{i+1} \sqsupset x_{i+1}$ by transitivity. If $x_i \not\sqsupset x_i$, then x_{i+1} exists since x_i is not free; x_{i+1} is fresh otherwise $x_i \sqsupset x_i$ by transitivity. $\square$

Corollary 4.35. *If $\mathcal{G}_\theta$ has a $\parr/\&$ then it has a free $\parr/\&$.*

Distinct $\parr/\&$-vertices x and y or $\mathcal{G}_\theta$ are ***face-to-face***, denoted $x \leftarrow\rightarrow y$, if there is a switching path $xz_0 \ldots z_n y$ in $\mathcal{G}_\theta$ such that $x \leftarrow z_0$ and $z_n \rightarrow y$, and are ***back-to-back***, denoted $x \rightarrow\leftarrow y$, if there exists a path $xz_0 \ldots z_n y$ in $\mathcal{G}_\theta$ such that $x \rightarrow z_0$ and $z_n \leftarrow y$, and none of the z_i are $\parr/\&$-vertices (so in particular $xz_0 \ldots z_n y$ is a strong path).

Recall that a $\parr/\&$-vertex x of $\mathcal{G}_\theta$ ***separates*** if it is not an argument (i.e., is an outermost connective), or it is the argument of y and deleting the edge between x and y disconnects[18] $\mathcal{G}_\theta$.

Lemma 4.36. *If a $\parr/\&$-vertex x is free and does not separate, then $x \rightarrow\leftarrow y$ and $x \leftarrow\rightarrow z$ for free y and z.*

Proof. Since x does not separate, it is in a cycle C (say clockwise) whose first (resp. last) edge is oriented out of (resp. into) x. Take y to be the first $\parr/\&$ reached clockwise along C from x. Then $x \rightarrow\leftarrow y$ (otherwise $y \sqsupset x$ by switch then extend) and y is free since $y' \sqsupset y$ implies $y' \sqsupset x$ by extend, contradicting the freedom of x.

By switch, the anti-clockwise neighbour of x in C is dominated by x. Let v be the first vertex reached anti-clockwise from x that is not dominated by x, and let v' be its predecessor. Since $x \sqsupset v'$, we have v a $\parr/\&$ and $v' \rightarrow v$, otherwise $x \sqsupset v$ by extend. Let $z = v$ if v is free, otherwise let z be a free $\parr/\&$ dominating v provided by Lemma 4.34; in the first case $z \sqsupset v'$ by switch, in the second case by extend.

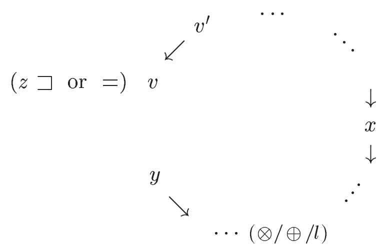

Note that $z \neq x$ since either $v \neq x$ (case $z = v$) or $z \sqsupset v \not\sqsubset x$ (otherwise). Apply meet to $z \sqsupset v' \sqsubset x$. $\square$

[17] This lemma is not specific to proof nets, but is a general observation about binary relations $\succ$. Say that x is $\succ$-dominated if $y \succ x$ for some y, and $\succ$-free otherwise. For any finite transitive binary relation $\succ$ such that $x \succ x$ implies $y \succ x$ for some $y \not\succ y$ (c.f. Lemma 4.33), every x is either $\succ$-free or $y \succ x$ for some $\succ$-free y.

[18] In the case with mix, read "disconnects" as "increases the number of connected components of".

All the auxiliary material is in place for us to prove the Separation Lemma (if $\mathcal{G}_\theta$ has a $\parr$/& then it has a separating[19] $\parr$/&, page 20).

Proof of Lemma 4.19 (Separation Lemma). If $\mathcal{G}_\theta$ had no separating $\parr$/& then $x_0 \longleftrightarrow x_1 \rightarrow\leftarrow x_2 \longleftrightarrow x_3 \rightarrow\leftarrow \ldots$ for free $\parr$/&-vertices x_i with $x_{i+1} \neq x_i$ by Lemma 4.36, and x_0 existing by Corollary 4.35. By finiteness, the composite π of the paths witnessing the $\longleftrightarrow$ and $\rightarrow\leftarrow$ relations eventually intersects itself at a vertex x, yielding a path $\pi' = xy_0 \ldots y_n$ such that $\{x, y_0, \ldots, y_n\}$ is a cycle. Each witness is a switching path, so π' is a switching path (since by design, composition at each x_i avoids introducing consecutive switch edges of x_i). Furthermore, one of the x_i must be among the y_j (since each witness is a path of distinct vertices). Using self if $\{x, y_0, \ldots, y_n\}$ is a switching cycle, and fork otherwise, this x_i is dominated, a contradiction (since x_i is free). □

4.13 Proof of the cut-free sequentialisation theorem

With the Separation Lemma in hand, the proof that every cut-free proof net is the translation of a cut-free proof reduces to simple induction.

Let θ be a proof net on Γ. We proceed by induction on the sum of the number of $\parr$s and &s of $\mathcal{G}_\theta$.

Base case (primary induction) Γ is $\parr$/&-free, hence θ comprises a single linking λ on Γ. We proceed by induction on the number of connectives of Γ.

- *Base case (secondary induction).* Γ contains no connectives, so Γ has the form $P_1, P_1^\perp, \ldots, P_n, P_n^\perp$ for $n \geq 0$ and propositional variables $P_1, \ldots, P_n$, and λ links P_i and $P_i^\perp$ for $i = 1, ..., n$. By (P2) $n = 1$. The axiom rule with conclusion $P_1, P_1^\perp$ is a sequentialisation of θ.
- *Induction step (secondary induction).* With no $\parr$s, $\mathcal{G}_\lambda$ is the only $\parr$-switching of λ, so by (P2) $\mathcal{G}_\lambda$ is a tree.
  - Suppose $\Gamma = \Delta, A \oplus B$, with $\oplus$-vertex $x \in \mathcal{G}_\lambda$ corresponding to $A \oplus B$. Since $\Gamma \upharpoonright \lambda$ is an additive resolution, x is unary in $\mathcal{G}_\lambda$, i.e., there is a unique $y \in \mathcal{G}_\lambda$ with $y \to x$. Depending on whether y is the left/right argument of x, let ρ be a left/right $\oplus$-rule, with conclusion $\Delta, A \oplus B$ and hypothesis $\Gamma' = \Delta, A$ or Δ, B, correspondingly. The linking λ on Γ also constitutes a linking λ' on Γ', since no leaves of the deleted $\oplus$-argument were incident with a link of λ. The graph $\mathcal{G}_{\lambda'}$ is a tree, because $\mathcal{G}_\lambda$ is a tree. Hence $\theta' = \{\lambda'\}$ is a proof net on Γ'. By induction, θ' is the translation of a cut-free MALL proof of Γ', which when followed by ρ constitutes a cut-free MALL proof of Γ whose translation is θ.
  - Suppose $\Gamma = \Delta, A_0 \otimes A_1$, with $\otimes$-vertex $x \in \mathcal{G}_\lambda$ corresponding to $A_0 \otimes A_1$. Deleting x separates the tree $\mathcal{G}_\lambda$ into a left tree T_0 and right tree T_1 whose respective conclusions define sequents Δ_0 and Δ_1, a partitioning of Δ. Let ρ be a $\otimes$-rule with conclusion Γ and hypotheses Δ_0, A_0 and Δ_1, A_1. Since $\mathcal{G}_\lambda$ is a tree, no link of λ goes between Δ_0, A_0 and Δ_1, A_1, hence λ partitions to form linkings λ_0 and λ_1 on Δ_0, A_0 and Δ_1, A_1, respectively. Each $\theta_i = \{\lambda_i\}$ is a proof net on Δ_i, A_i since each $\mathcal{G}_{\lambda_i} = T_i$ is a tree. Appeal to the induction hypothesis with θ_0 and θ_1, in the manner of the $\oplus$ case above.

[19] We actually prove a stronger result, that if $\mathcal{G}_\theta$ has a $\parr$/& then it has a separating *free* $\parr$/&.

Induction step (primary induction) Γ has a $\parr$/&. By (P2) $\mathcal{G}_\theta$ is connected.

(a) Suppose $\Gamma = \Delta, A \parr B$, with $\parr$-vertex $x \in \mathcal{G}_\theta$ corresponding to $A \parr B$. Let ρ be a $\parr$-rule with conclusion Γ and hypothesis $\Gamma' = \Delta, A, B$. The sequents Γ and Γ' have the same leaves and (aside from the presence/absence of x) the same &- and additive resolutions, so θ constitutes a proof structure on Γ'. On Γ', the $\parr$-switchings of the linkings of θ are trees, since they are obtained from those on Γ by deleting x. Any subset $\Lambda \subseteq \theta$ toggles the same &s in Γ' as it does in Γ, and $\mathcal{G}_\Lambda$ has the same switching cycles with respect to Γ' as with respect to Γ. Therefore θ is a proof net on Γ'. Appeal to the induction hypothesis with the proof net θ on Γ'; follow the resulting proof with ρ.

(b) Suppose $\Gamma = \Delta, A_0 \& A_1$, with vertex $w \in \mathcal{G}_\theta$ corresponding to $A_0 \& A_1$. Let ρ be a &-rule with conclusion Γ and hypotheses $\Gamma_0 = \Delta, A_0$ and $\Gamma_1 = \Delta, A_1$. Define the sets of linkings θ_i on Γ_i to comprise those linkings of θ which are on $\Gamma_i \subseteq \Gamma$. Trivially, each θ_i is a proof net. Appeal to the induction hypothesis with each θ_i; combine the resulting proofs with ρ.

(c) Suppose $\mathcal{G}_\theta$ has no $\to$-terminal (i.e. concluding) $\parr$ or &. By the Separation Lemma $\mathcal{G}_\theta$ has a $\parr$/&-vertex x such that the deletion of the edge $x \to y$ disconnects $\mathcal{G}_\theta$ into G_0 and G_1.

Let G_0 be the component containing x, and let Γ_0 comprise the formulas corresponding to the $\to$-terminal vertices of G_0 (some formulas of Γ together with the subformula $A \& B$ corresponding to x). Define[20] $\theta_0 = \{\lambda \restriction \Gamma_0 : \lambda \in \theta\}$ on Γ_0 (each $\lambda \restriction \Gamma_0$ is well-defined since no $a \in \lambda$ goes between G_0 and G_1).

Let Γ_1 be the subsequent of Γ containing the formulas corresponding to the $\to$-terminal vertices of G_1. In G_1, y is $\to$-initial. Form G_1^+ from G_1 by adding literals P and $P^\perp$ with a link edge a between them. Let $\widehat{\Gamma}_1$ be Γ_1 with P substituted for the subformula $A \& B$ corresponding to x, and let $\Gamma_1^+ = \widehat{\Gamma}_1, P^\perp$. Define $\theta_1 = \{\lambda \restriction \widehat{\Gamma}_1 \cup \{a\} : \lambda \in \theta\}$ on Γ_1^+.

Claim: $x \in \Gamma \restriction \lambda$ for all $\lambda \in \theta$.

Proof. If not, there is $\lambda \in \theta$ and a &-vertex w with x in $\Gamma \restriction \lambda$ but not in $\Gamma \restriction \lambda_w$ for some $\lambda_w \in \theta$ such that $\lambda \overset{w}{=} \lambda_w$. Thus there is a jump $l \to w$ in $\mathcal{G}_\theta$ for some $l \in G_0$ with l in a link of $\lambda \setminus \lambda_w$. Since linkings are total on additive resolutions there is a leaf l' in a link of $\lambda_w \setminus \lambda$ connecting to the formula containing x, but not satisfying $l' \to \cdots \to x$, so there is a jump $l' \to w$ in $\mathcal{G}_\theta$. If $w \in G_0$ then $l' \to w$ is a jump from G_1 to G_0, and if $w \in G_1$ then $l \to w$ is a jump from G_0 to G_1; either case violates the disconnectedness of G_0 from G_1. ∎

The claim implies that θ_0 and θ_1 are sets of linkings on Γ_0 and Γ_1^+, respectively. Moreover, $\mathcal{G}_{\theta_0} = G_0$ and $\mathcal{G}_{\theta_1} = G_1^+$. We now check that θ_0 and θ_1 are proof nets, i.e., satisfy (P1)–(P3). Since θ satisfies (P1), θ_0 (resp. θ_1) has at least one linking on every &-resolution of Γ_0 (resp. Γ_1^+). Had θ_i two distinct linkings on the same &-resolution, there would be a jump from a link in G_i to a & in G_{1-i}, violating the disconnectedness of G_0 from G_1. Thus θ_i satisfies (P1). (P2) is trivially inherited from θ. Finally, (P3) holds since any set Λ' of

[20]This instance $\lambda \restriction \Gamma_0$ of restriction is a normal instance of restriction, and should not be confused with the notation $\Gamma \restriction \lambda$ for the additive resolution of a linking λ on Γ.

linkings in θ_0 or θ_1 corresponds to a set Λ of linkings in θ toggling the same &s, such that any switching cycle of $\mathcal{G}_{\Lambda'}$ is a switching cycle of $\mathcal{G}_\Lambda$.

Since Γ_0 has an outermost &, by case (b) above θ_0 is the translation of a cut-free proof Π_0 of Γ_0. Since $\mathcal{G}_{\theta_1}$ has less ⅋s and &s than $\mathcal{G}_\theta$, by induction θ_1 is the translation of a cut-free proof Π_1. Substituting Π_0 for the axiom rule with conclusion $P, P^\perp$ in Π_1 yields a proof whose translation is θ. □

In the case of $\mathrm{MALL}^{\mathsf{mix}}$, the connectedness requirement of (P2) does not apply. In each of the cases (a)–(c) of the primary induction step above we check that θ satisfying (P2) implies θ_i (or θ on Γ', in case (a)) satisfies (P2); note that this also works for ($\mathrm{P2}^{\mathsf{mix}}$). Additionally, connectedness is used three times in the above proof.[21] To prove that a set of linkings is the translation of a cut-free $\mathrm{MALL}^{\mathsf{mix}}$ proof if it is a cut-free mix net, in each part of the inductive proof above, the case that $\mathcal{G}_\theta$ is not connected can be dealt with by partitioning Γ into a number of non-empty subsequents Γ_i, each harbouring a connected component of $\mathcal{G}_\theta$. The mix net θ projects to mix nets θ_i on Γ_i, which by induction are translations of cut-free $\mathrm{MALL}^{\mathsf{mix}}$ proofs Π_i. By the mix rule these combine into a sequentialisation of θ.

5 Cut

This section extends proof nets with cuts. Section 5.1 defines a simple and strongly normalising cut elimination on proof nets, which can be executed in a single step (*turbo cut elimination*, Section 5.1.1). Normalisation yields an associative composition of cut-free MALL proof nets, whence a category $\mathcal{N}$ of cut-free MALL proof nets which is semi (i.e. unit-free) star-autonomous with products and sums (Section 5.2). Section 5.3 defines a translation from MALL proofs to sets of linkings, and proves the *Sequentialisation Theorem*: a set of linkings is a translation of a MALL proof *iff* it is a MALL proof net.

A ***cut pair*** is a formula $A{*}A^\perp$ where A is any MALL formula. The connective $*$ is called ***cut***. By definition, we take $*$ to be unordered, i.e., $A{*}A^\perp = A^\perp{*}A$. This is in contrast to MALL formulas, where connectives are ordered, e.g., $A\otimes B \neq B\otimes A$ when $A \neq B$. We continue to identify a formula with its parse tree (including a cut pair, whose root is a $*$-labelled vertex with two unordered children). A ***cut sequent*** is a disjoint union of a MALL sequent and zero or more cut pairs. (Recall that a MALL sequent is a non-empty disjoint union of MALL formulas.) Given a (possibly empty) disjoint union Σ of cut pairs and a MALL sequent Γ, write $[\Sigma]\,\Gamma$ for the cut sequent which is the disjoint union of Σ and Γ.

A ***cut-additive resolution*** of a cut sequent Δ is any result of deleting zero or more cut pairs from Δ and one argument subtree of every additive connective (& or $\oplus$). Thus every remaining & and $\oplus$ is unary.

Example 5.1. Here is a cut sequent followed by one of its cut-additive resolutions:

$$P\otimes P,\ Q*Q^\perp,\ P^\perp\oplus Q,\ (R\oplus S)*(R^\perp\&S^\perp)$$

$$P\otimes P,\ Q*Q^\perp,\ P^\perp\oplus Q,\ (R\oplus S)*(R^\perp\&S^\perp)$$

[21] In the base case of the secondary induction, to conclude $n = 1$; in the secondary induction step, to conclude that (P2) can be reformulated as $\mathcal{G}_\lambda$ being a tree; and in the primary induction step, to conclude that $\mathcal{G}_\theta$ is connected.

Definition: MALL proof net

Cut pair: formula $A*A^\perp$ $(= A^\perp * A)$ for any MALL formula A.
Cut sequent Δ: disjoint union of a MALL sequent and any number of cut pairs.
&-***resolution***: deletion of one argument subtree of each &.
Cut-additive resolution: deletion of some cuts and one argument subtree of each $\oplus$/&.
(***Axiom***) ***link*** on Δ: edge between complementary leaves (literal occurrences) in Δ.
Linking λ on Δ: partitioning of the leaves of an additive resolution $\Delta\restriction\lambda$ of Δ into links.
A set Λ of linkings on Δ ***toggles*** a & w if both arguments of w are in $\Delta\restriction\Lambda \equiv \bigcup_{\lambda\in\Lambda}\Delta\restriction\lambda$.
Graph $\mathcal{G}_\Lambda$: $\Delta\restriction\Lambda + \cup\Lambda$ + ***jump*** edges l—w—l' if $\{l,l'\}\in\lambda\backslash\lambda'$ and $\{\lambda,\lambda'\}\subseteq\Lambda$ toggles w only.
Switching cycle: cycle with ≤ 1 ***switch edge*** (= jump or argument edge) of each ⅋/&.

A set θ of linkings on Δ is a ***proof net*** if it satisfies:

CUT: Every cut pair has a leaf in θ.
RESOLUTION: Exactly one linking of θ is on any given &-resolution of Δ.
MLL: Every ⅋-switching of every linking in θ is a tree (acyclic and connected).[22]
TOGGLING: Every set Λ of ≥ 2 linkings of θ toggles a & that is in no switching cycle of $\mathcal{G}_\Lambda$.[23]

A ***link*** on a cut sequent Δ is a pair of complementary leaves in Δ, i.e., a pair of leaves in Δ labelled with complementary literals P and $P^\perp$. A ***linking*** λ on Δ is a set of disjoint links on Δ such that $\cup\lambda$ is the set of leaves of a cut-additive resolution of Δ; this cut-additive resolution is denoted $\Delta\restriction\lambda$.

Example 5.2. Here are two examples of sets of linkings on cut sequents:

$$\theta: \quad P,\ P^\perp * P,\ P^\perp * P,\ P^\perp \& (P^\perp \oplus Q)$$

$$\phi: \quad P,\ P^\perp * P,\ P^\perp \& (P^\perp \oplus Q)$$

Each of θ and ϕ has two linkings, one shown above the cut sequent, the other below. Each linking has two links. Note that each linking takes the leaves of a cut-additive resolution.

In the presence of cut, we update all the auxiliary definitions of Section 4 (&-resolution, $\mathcal{G}_\Lambda$, switching cycle, etc.) by substituting *cut sequent* for *sequent* and *cut-additive resolution* for *additive resolution* throughout.

Definition 5.3. A set θ of linkings on a cut sequent Δ is a ***proof net*** if it satisfies:

(P0) CUT. At least one leaf of every cut pair is in θ (i.e., in some link of some linking of θ).
(P1) RESOLUTION. For any &-resolution $\Delta^\star$ of Δ, exactly one linking of θ is on $\Delta^\star$.
(P2) MLL. Every ⅋-switching of every linking in θ is a tree (acyclic and connected).[22]
(P3) TOGGLING. Every set Λ of two or more linkings of θ toggles a & that is not in any switching cycle of $\mathcal{G}_\Lambda$.[23]

[22]By dropping connectedness, we obtain a proof net for MALL augmented by the mix rule.
[23]In fact, it suffices to verify TOGGLING merely for ***saturated*** sets of linkings Λ, namely, such that any strictly larger subset of θ toggles more &s than Λ. There is exactly one saturated set of linkings in θ for each ***partial*** &-***resolution*** of Δ, the latter being any result of deleting at most one argument subtree of each & of Δ.

The definition is summarised in the box on page 40. Note that (P1)–(P3) are inherited from the cut-free case. We say that θ is a ***proof structure*** if it satisfies (P0) and (P1).

Alternative definitions of proof net. The material in Section 4.7 still applies now that we have extended proof nets with cut, with one small change: in the equation defining balance, add the number of cuts to the number of tensors. The equivalence proofs (appendices B and C) extend verbatim once a cut is viewed as a tensor.

5.1 Cut elimination

Let θ be a set of linkings on a cut sequent Δ, and let $A * A^\perp$ be a cut pair in Δ. Define the ***elimination*** of $A * A^\perp$ (or, of the cut $*$ between A and $A^\perp$) as follows.

(a) If A is a literal, delete $A * A^\perp$ from Δ, and replace any pair of links $\{l, A\}, \{A^\perp, l'\}$ in a linking of θ (l and l' being other occurrences of $A^\perp$ and A respectively) with the link $\{l, l'\}$.

(b) If $A = A_1 \otimes A_2$ and $A^\perp = A_1^\perp \parr A_2^\perp$ (or vice versa), replace $A * A^\perp$ with two cut pairs $A_1 * A_1^\perp$ and $A_2 * A_2^\perp$. Retain all the original linkings.

(c) If $A = A_1 \& A_2$ and $A^\perp = A_1^\perp \oplus A_2^\perp$ (or vice versa) replace $A * A^\perp$ with two cut pairs $A_1 * A_1^\perp$ and $A_2 * A_2^\perp$. Delete the ***inconsistent*** linkings, namely those $\lambda \in \theta$ such that in $\Delta \upharpoonright \lambda$ the children $\&$ and $\oplus$ of the cut take opposite arguments (i.e., such that the right argument of the $\&$ is in $\Delta \upharpoonright \lambda$ and the left argument of the $\oplus$ is in $\Delta \upharpoonright \lambda$, or vice versa). Finally, 'garbage collect' by deleting $A_i * A_i^\perp$ if no leaf of $A_i * A_i^\perp$ is in any of the remaining linkings.

An example of cut elimination was presented in Figure 3 (page 5).

Proposition 5.4. *Eliminating a cut from a proof net yields a proof net.*[24]

Proof. Section 5.4. □

Theorem 5.5. *Cut elimination of proof nets is strongly normalising.*[24]

Proof. Confluence is immediate from the definition; cut elimination reduces the size of the cut sequent, and is therefore strongly normalising. □

5.1.1 *Turbo cut elimination.* Cut elimination can be completed in a single step. For l the i^{th} leaf of A in a cut pair $A * A^\perp$, let $l^\perp$ denote the i^{th} leaf of $A^\perp$.[25] A linking λ on a cut sequent Δ ***matches*** if, for every cut pair $A * A^\perp$ in Δ, any given leaf l of $A * A^\perp$ is in $\Delta \upharpoonright \lambda$ *iff* $l^\perp$ is in $\Delta \upharpoonright \lambda$.

Example 5.6. The first linking below matches, the second does not:

$$Q \oplus P, \quad [P^\perp \& (Q^\perp \oplus P^\perp)] * [P \oplus (Q \& P)], \quad P^\perp \oplus Q^\perp$$

$$Q \oplus P, \quad [P^\perp \& (Q^\perp \oplus P^\perp)] * [P \oplus (Q \& P)], \quad P^\perp \oplus Q^\perp$$

Note that, although not matching, this second linking is consistent (the opposite of inconsistent, defined above).

[24] Proposition 5.4 and Theorem 5.5 also hold for mix nets: that elimination preserves (P2$^{\mathsf{mix}}$) is part of the argument (page 53) that it preserves (P2).

[25] Remember that cut $*$ is unordered, i.e., $A * A^\perp = A^\perp * A$; thus $l^{\perp\perp} = l$, as one would expect.

A linking matches iff it is hereditarily/iteratively consistent: when (non-turbo) cut elimination is carried out, the linking never becomes inconsistent in the sense of case (c) in the definition of (non-turbo) cut elimination above.

Suppose a linking λ on a cut sequent Δ matches. The ***reduction*** $\overline{\Delta}$ of Δ is the result of deleting all cut pairs from Δ. The ***reduction*** $\overline{\lambda}$ of λ is the linking on $\overline{\Delta}$ obtained by replacing every set of links $\{l_0, l_1\}, \{l_1^\perp, l_2\}, \{l_2^\perp, l_3\}, \ldots, \{l_{n-1}^\perp, l_n\}$ in λ in which only l_0 and l_n occur in $\overline{\Delta}$ by the single link $\{l_0, l_n\}$.

Example 5.7. Here is an example of the reduction of a matching linking. The informal intermediate step is for visualisation only.

$$P\&Q, \quad [R \oplus (Q^\perp \otimes Q)] * [R^\perp \& (Q \parr Q^\perp)], \quad Q^\perp$$

$$\downarrow$$

$$P\&Q, \qquad\qquad\qquad\qquad\qquad\qquad Q^\perp$$

$$\downarrow$$

$$P\&Q, \qquad\qquad\qquad\qquad\qquad\qquad Q^\perp$$

Let θ be a set of linkings on the cut sequent Δ. The ***normal form*** of θ is the set of linkings $\overline{\theta}$ on $\overline{\Delta}$ obtained from θ by deleting every non-matching linking and reducing every linking which remains. By a simple structural induction on the size of the cut pairs in Δ, the set of linkings $\overline{\theta}$ is precisely the normal form obtained by (non-turbo) cut elimination.

Example 5.8. Let θ be the following proof net with four linkings (two shown on each of two copies of the sequent):

$$Q \oplus P, \quad [P^\perp \& (Q^\perp \oplus P^\perp)] * [P \oplus (Q \& P)], \quad P^\perp \oplus Q^\perp$$

$$Q \oplus P, \quad [P^\perp \& (Q^\perp \oplus P^\perp)] * [P \oplus (Q \& P)], \quad P^\perp \oplus Q^\perp$$

Only the first of the four linkings is consistent:

$$Q \oplus P, \quad [P^\perp \& (Q^\perp \oplus P^\perp)] * [P \oplus (Q \& P)], \quad P^\perp \oplus Q^\perp$$

Reducing this linking yields the following one-linking normal form of θ:

$$Q \oplus P, \qquad\qquad\qquad\qquad\qquad\qquad P^\perp \oplus Q^\perp$$

Note that turbo cut elimination operates *independently* on each linking: a given linking is either deleted (if non-matching) or reduced using path composition (if matching), without reference to any other linking. This is similar to the situation in proof nets for polarised linear logic [Laurent and Tortura de Falco 2004].

5.2 The category of proof nets

Non-categorists can skip to Section 5.3 without loss of continuity.

Cut elimination yields a category $\mathcal{N}$ of MALL proof nets. Objects are MALL formulas, and a morphism $A \to B$ is a cut-free proof net on the sequent $A^\perp, B$. The composition of $\theta : A \to B$ and $\theta' : B \to C$ is the normal form of the proof net $\{\lambda \cup \lambda' : \lambda \in \theta, \lambda' \in \theta'\}$ on $A^\perp, B * B^\perp, C$. See Figure 4 (page 6). Composition is associative, since cut elimination is strongly normalising.[26] The identity morphism $\mathsf{id}_A : A \to A$ is defined as follows. An ***identity link*** on the sequent $A^\perp, A$ is a link between the i^{th} leaf of $A^\perp$ and the i^{th} leaf of A, for some i. An ***identity linking*** is one whose every link is an identity link. The set id_A comprises every identity linking on $A^\perp, A$.

Define a ***semi star-autonomous category*** as a category $\mathbb{C}$ equipped with the following structure of a star-autonomous category [Barr 1979], not involving units:

– Tensor. A functor $- \otimes - : \mathbb{C} \times \mathbb{C} \to \mathbb{C}$.

– Associativity. A natural isomorphism $a_{A,B,C} : (A \otimes B) \otimes C \to A \otimes (B \otimes C)$ natural in objects $A, B, C \in \mathbb{C}$ such that the following pentagon commutes:

$$\begin{tikzcd}
\big((A \otimes B) \otimes C\big) \otimes D \arrow[r, "a"] \arrow[d, "a\otimes\mathsf{id}"'] & (A \otimes B) \otimes (C \otimes D) \arrow[r, "a"] & A \otimes \big(B \otimes (C \otimes D)\big) \\
\big(A \otimes (B \otimes C)\big) \otimes D \arrow[rr, "a"'] & & A \otimes \big((B \otimes C) \otimes D\big) \arrow[u, "\mathsf{id}\otimes a"']
\end{tikzcd}$$

– Symmetry. A natural isomorphism $c_{A,B} : A \otimes B \to B \otimes A$ natural in objects $A, B \in \mathbb{C}$ such that $c_{B,A} \circ c_{A,B} = \mathsf{id}_{A \otimes B}$ and the following hexagon commutes:

$$\begin{tikzcd}
(A \otimes B) \otimes C \arrow[r, "a"] \arrow[d, "c\otimes\mathsf{id}"'] & A \otimes (B \otimes C) \arrow[r, "c"] & (B \otimes C) \otimes A \arrow[d, "a"] \\
(B \otimes A) \otimes C \arrow[r, "a"'] & B \otimes (A \otimes C) \arrow[r, "\mathsf{id}\otimes c"'] & B \otimes (C \otimes A)
\end{tikzcd}$$

– Involution. A functor $(-)^\perp : \mathbb{C}^{\mathsf{op}} \to \mathbb{C}$ with a natural isomorphism $A \to A^{\perp\perp}$.

– An isomorphism $\mathbb{C}(A \otimes B, C^\perp) \to \mathbb{C}(A, (B \otimes C)^\perp)$ natural in all objects A, B, C.

The category $\mathcal{N}$ has a very simple semi star-autonomous structure. Tensor $- \otimes - : \mathcal{N} \times \mathcal{N} \to \mathcal{N}$ acts symbolically on objects (i.e., the tensor of formulas A and B is the formula $A \otimes B$), and the tensor $\theta \otimes \theta' : A \otimes C \to B \otimes D$ of $\theta' : A \to B$ and $\theta' : C \to D$ is obtained as follows, using the notation of Table I (page 12):

$$\dfrac{\dfrac{\theta \ \rhd\ A^\perp, B \qquad\qquad \theta' \ \rhd\ D, C^\perp}{\{\lambda \cup \lambda' : \lambda \in \theta, \lambda' \in \theta'\} \ \rhd\ A^\perp, B \otimes D, C^\perp}\otimes}{\theta \otimes \theta' \ \rhd\ A^\perp ⅋ C^\perp, B \otimes D}⅋$$

Duality/negation $(-)^\perp : \mathcal{N}^{\mathsf{op}} \to \mathcal{N}$ on objects is as already defined on formulas (i.e. $(A \otimes B)^\perp = A^\perp ⅋ B^\perp$ etc., page 7). On morphisms it is trivial, since a proof

[26] Associativity is also straightforward with turbo cut elimination as the primary definition, since linking reduction is path composition.

net on A, B can be read equally well as a morphism $A^\perp \to B$ or $B^\perp \to A$. Tensor associativity is immediate since the formula graphs $A \otimes (B \otimes C)$ and $(A \otimes B) \otimes C$ are topologically equivalent, in particular having the same leaves. Symmetry, and the natural isomorphism $\mathcal{N}(A \otimes B, C^\perp) \cong \mathcal{N}(A, (B \otimes C)^\perp)$, are similarly trivial.

A semi star-autonomous category, as axiomatised above, is but a very rudimentary notion of "unitless" star-autonomous category. For example, the axiomatisation does not appear to provide a map $A \to A \otimes (B \parr B^\perp)$, which is present in the proof net category $\mathcal{N}$.

Products and sums. The category $\mathcal{N}$ of MALL proof nets has products and sums (coproducts). (By duality, the one yields the other.) Product is & and sum is $\oplus$, each acting syntactically on objects and defined on morphisms in a manner analogous to tensor above. The universal property of & holds because it takes the disjoint union of non-empty sets of linkings; $\oplus$ is dual.

Softness. The category $\mathcal{N}$ of proof nets is soft [Joyal 1995], that is, any morphism $\otimes_{1 \le i \le m}(A_i \& A_i') \to \parr_{1 \le j \le n}(B_j \oplus B_j')$ factorises through either a product projection on the left or a coproduct injection on the right. This is immediate, via sequentialisation, from the corresponding observation at the level of proofs. (Alternatively, it is straightforward to verify softness directly.)

5.3 Sequentialisation

This section defines a translation from MALL proofs to sets of linkings, and proves the *Sequentialisation Theorem*: a set of linkings on a cut sequent is a translation of a MALL proof *iff* it is a proof net. The translation goes via a technically convenient variant $\mathrm{MALL}^{\mathsf{cut}}$ of MALL in which cuts are retained in sequents.

5.3.1 *A function from $\mathrm{MALL}^{\mathsf{cut}}$ derivations to sets of linkings.* Cut sequents are derived in $\mathrm{MALL}^{\mathsf{cut}}$ using the rules in Table II. Example $\mathrm{MALL}^{\mathsf{cut}}$ derivations are shown in Figure 15. Every $\mathrm{MALL}^{\mathsf{cut}}$ derivation projects to a MALL proof in the obvious way, by deleting the cut pairs. For example, the $\mathrm{MALL}^{\mathsf{cut}}$ derivations of Figure 15 project to the MALL proofs of Figure 16, as follows:

$$\begin{tikzcd}[row sep=tiny] D_1 \arrow[dr] & \\ & \Pi \\ D_2 \arrow[dr] & \\ & \Pi' \\ D_3 \arrow[ur] & \\ & \Pi'' \\ D_4 \arrow[ur] & \end{tikzcd}$$

The system $\mathrm{MALL}^{\mathsf{cut}}$ is an extension of cut-free MALL. The function taking a cut-free MALL proof to a set of linkings on a MALL sequent (defined in Section 4.2, page 12) extends in the obvious way to a function taking a $\mathrm{MALL}^{\mathsf{cut}}$ derivation D to a set θ_D of linkings on a cut sequent Δ. Define a ***&-resolution*** R of D to be any result of deleting one branch above each &-rule of D. By downwards tracking of formula leaves, the axiom rules of R determine a linking λ_R on Δ. Define $\theta_D = \{\lambda_R : R \text{ is a \&-resolution of } D\}$. Alternatively, Table III defines the same function by induction, a direct extension of the cut-free case in Table I (page 12). Figure 17 (page 48) shows how each derivation D_i in Figure 15 (page 45) translates

$$\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{[\Omega]\,\Gamma, A, B}{[\Omega]\,\Gamma, A \parr B}\,\parr \qquad \dfrac{[\Omega]\ \Gamma, A \qquad [\Omega']\ A^{\perp}, \Delta}{[\Omega, \Omega',\ A * A^{\perp}]\ \Gamma, \Delta}\,\mathsf{cut} \qquad \dfrac{[\Omega]\,\Gamma, A}{[\Omega]\,\Gamma, A \oplus B}\,\oplus_1$$

$$\dfrac{[\Sigma, \Omega]\,\Gamma, A \qquad [\Sigma, \Omega']\,\Gamma, B}{[\Sigma, \Omega, \Omega']\ \Gamma, A \& B}\,\& \qquad \dfrac{[\Omega]\,\Gamma, A \qquad [\Omega']\,B, \Delta}{[\Omega, \Omega']\,\Gamma, A \otimes B, \Delta}\,\otimes \qquad \dfrac{[\Omega]\,\Gamma, B}{[\Omega]\,\Gamma, A \oplus B}\,\oplus_2$$

Table II. Rules for deriving cut sequents in MALL$^{\mathsf{cut}}$. Here P ranges over propositional variables, A, B range over MALL formulas, Γ, Δ range over (possibly empty) disjoint unions of MALL formulas, and Σ, Ω, Ω' range over (possibly empty) disjoint unions of cut pairs. Note that the &-rule may superimpose one or more cut pairs from its two hypotheses (if Σ is non-empty), or may leave all cut pairs separate (if Σ is empty).

(D_1)

$$\dfrac{\dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}{*}P,\ P^{\perp}}\,\mathsf{cut} \qquad \dfrac{\dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}{*}P,\ P^{\perp}}\,\mathsf{cut}}{P,\ P^{\perp}{*}P,\ P^{\perp}{\oplus}Q}\,\oplus_1}{P,\ P^{\perp}{*}P,\ P^{\perp}{*}P,\ P^{\perp}\&(P^{\perp}{\oplus}Q)}\,\&$$

(D_2)

$$\dfrac{\dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}{*}P,\ P^{\perp}}\,\mathsf{cut} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}{\oplus}Q}\,\oplus_1}{P,\ P^{\perp}{*}P,\ (P^{\perp}{\oplus}Q)}\,\mathsf{cut}}{P,\ P^{\perp}{*}P,\ P^{\perp}{*}P,\ P^{\perp}\&(P^{\perp}{\oplus}Q)}\,\&$$

(D_3)

$$\dfrac{\dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}{*}P,\ P^{\perp}}\,\mathsf{cut} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}{\oplus}Q}\,\oplus_1}{P,\ P^{\perp}{*}P,\ (P^{\perp}{\oplus}Q)}\,\mathsf{cut}}{P,\ P^{\perp}{*}P,\ P^{\perp}\&(P^{\perp}{\oplus}Q)}\,\&$$

(D_4)

$$\dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}{\oplus}Q}\,\oplus_1}{P,\ P^{\perp}\&(P^{\perp}{\oplus}Q)}\,\&}{P,\ P^{\perp}{*}P,\ P^{\perp}\&(P^{\perp}{\oplus}Q)}\,\mathsf{cut}$$

Fig. 15. Examples of derivations of cut sequents in MALL$^{\mathsf{cut}}$. The only difference between derivations D_1 and D_2 is a commutation of the cut and $\oplus_1$ rules in the right branch. Both derivations yield the same cut sequent. The only difference between derivations D_2 and D_3 is the final &-rule: the application in D_2 keeps the cut pairs in the hypotheses separate (an instance of the Table II &-rule taking Σ empty and $\Omega = \Omega' = P^{\perp} * P$), whereas the application in D_3 superimposes the two cut pairs ($\Sigma = P^{\perp} * P$ and each of Ω and Ω' empty). Derivation D_4 yields the same cut sequent as D_3, but with the cut and & rules commuted.

$$(\Pi) \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}}\,\mathsf{cut} \qquad \dfrac{\dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}}\,\mathsf{cut}}{P,\ P^{\perp}\oplus Q}\,\oplus_1$$
$$\dfrac{\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad}{P,\ P^{\perp}\&(P^{\perp}\oplus Q)}\,\&$$

$$(\Pi') \qquad \dfrac{\dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}}\,\mathsf{cut} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}\oplus Q}\,\oplus_1}{P,\ P^{\perp}\oplus Q}\,\mathsf{cut}}{P,\ P^{\perp}\&(P^{\perp}\oplus Q)}\,\&$$

$$(\Pi'') \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax} \qquad \dfrac{\dfrac{}{P,\ P^{\perp}}\,\mathsf{ax}}{P,\ P^{\perp}\oplus Q}\,\oplus_1}{P,\ P^{\perp}\&(P^{\perp}\oplus Q)}\,\&}{P,\ P^{\perp}\&(P^{\perp}\oplus Q)}\,\mathsf{cut}$$

Fig. 16. The MALL proofs projected from the MALL$^{\mathsf{cut}}$ derivations D_1, D_2, D_3, D_4 in Figure 15. Derivation D_1 projects to Π, derivations D_2 and D_3 project to Π', and D_4 projects to Π''. All three proofs yield the same MALL sequent. The only difference between proofs Π and Π' is a commutation of the cut and $\oplus_1$ rules in the right branch. The only difference between Π' and Π'' is a commutation of the cut and $\&$ rules.

$$\dfrac{}{\{\overparen{P, P^{\perp}}\} \ \rhd\ P, P^{\perp}}\,\mathsf{ax} \qquad\qquad \dfrac{\theta\ \rhd\ [\Omega]\ \Gamma, A \qquad\qquad \theta'\ \rhd\ [\Omega']\ A^{\perp}, \Delta}{\{\lambda\cup\lambda'\ :\ \lambda\in\theta, \lambda'\in\theta'\}\ \rhd\ [\Omega, \Omega', A * A^{\perp}]\ \Gamma, \Delta}\,\mathsf{cut}$$

$$\dfrac{\theta\ \rhd\ [\Sigma,\Omega]\ \Gamma, A \qquad \theta'\ \rhd\ [\Sigma,\Omega']\ \Gamma, B}{\theta\cup\theta'\ \rhd\ [\Sigma,\Omega,\Omega']\ \Gamma, A\&B}\,\& \qquad \dfrac{\theta\ \rhd\ [\Omega]\ \Gamma, A \qquad\qquad \theta'\ \rhd\ [\Omega']\ B, \Delta}{\{\lambda\cup\lambda' : \lambda\in\theta, \lambda'\in\theta'\}\ \rhd\ [\Omega,\Omega']\ \Gamma, A\otimes B, \Delta}\,\otimes$$

$$\dfrac{\theta\ \rhd\ [\Omega]\ \Gamma, A}{\theta\ \rhd\ [\Omega]\ \Gamma, A\oplus B}\,\oplus_1 \qquad\qquad \dfrac{\theta\ \rhd\ [\Omega]\ \Gamma, B}{\theta\ \rhd\ [\Omega]\ \Gamma, A\oplus B}\,\oplus_2 \qquad\qquad \dfrac{\theta\ \rhd\ [\Omega]\ \Gamma, A, B}{\theta\ \rhd\ [\Omega]\ \Gamma, A \parr B}\,\parr$$

Table III. Inductive definition of the function from MALL$^{\mathsf{cut}}$ derivations to sets of linkings. Here $\theta\ \rhd\ \Delta$ is the judgement "θ is a set of linkings on the cut sequent Δ". We use the implicit tracking of formula leaves downwards through rules. The base case ax is a singleton set of linkings whose only linking comprises a single link, between P and $P^{\perp}$. Here, as in the presentation of the rules of MALL$^{\mathsf{cut}}$ in Table II (page 45), P ranges over propositional variables, A, B range over MALL formulas, Γ, Δ range over (possibly empty) disjoint unions of MALL formulas, and Σ, Ω, Ω' range over (possibly empty) disjoint unions of cut pairs. This table is a direct extension of the inductive translation of cut-free MALL proofs, Table I (page 12); every cut-free MALL proof is in particular a MALL$^{\mathsf{cut}}$ derivation.

into a set of linkings.

By structural induction, each linking is well-defined (i.e., takes the leaves of a cut-additive resolution); thus the translation is well-defined. The fact that the above procedures yield the same set of linkings follows from a simple structural induction on derivations. A set of linkings Λ on a cut sequent Δ is ***cut-sequentialisable*** if it is the translation of a $\mathrm{MALL}^{\mathsf{cut}}$ derivation of Δ; any such derivation is a ***cut-sequentialisation*** of Λ.

5.3.2 *Translating a MALL proof into a set of linkings.* We have seen that every $\mathrm{MALL}^{\mathsf{cut}}$ derivation D of a cut sequent $[\Sigma]\,\Gamma$ projects to a MALL proof Π_D of the underlying MALL sequent Γ, and also translates into a set of linkings θ_D on $[\Sigma]\,\Gamma$. For example, the $\mathrm{MALL}^{\mathsf{cut}}$ derivations D_i in Figure 15 (page 45) project and translate as follows:

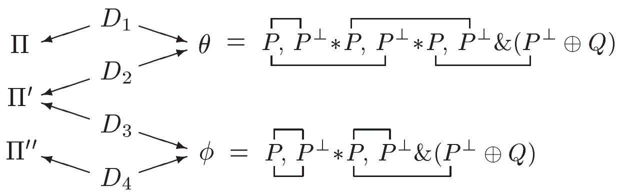

The leftward arrows show projection to the MALL proofs Π_j of Figure 16 (page 46), and the rightward arrows show translation into the sets of linkings θ and ϕ of Example 5.2 (page 40), with translations shown in Figure 17 (page 48).

Let θ be a set of linkings on a cut sequent. A MALL proof Π ***translates*** into θ, or is a ***sequentialisation*** of θ, if Π is the projection of a $\mathrm{MALL}^{\mathsf{cut}}$ derivation translating to θ; we say that θ is ***sequentialisable***[27], and write $\Pi \xrightarrow{\mathsf{s}} \theta$. For example, the projection/translation diagram above yields

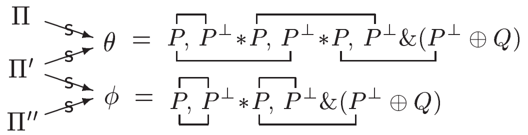

(the composite of the previous diagram of relations: from left to right, the inverse of projection, followed by translation). Restricted to the cut-free case, the sequentialisation relation $\xrightarrow{\mathsf{s}}$ is a function taking a proof to a set of linkings on a MALL sequent, exactly the cut-free translation defined in Table I (page 12). In the presence of cuts, more than one set of linkings on a cut sequent may correspond to the same MALL proof. In the diagram above, the MALL proof Π' of Figure 16 (page 46) is a common sequentialisation of θ and ϕ.

Our definition of proof net (Definition 5.3, page 40) characterises the image of the $\xrightarrow{\mathsf{s}}$ sequentialisation relation on MALL proofs (i.e., the image of the function on $\mathrm{MALL}^{\mathsf{cut}}$ derivations defined in Table III). Section 5.3.4 considers two alternative notions of sequentialisation, one in which the &-rule superimposes no cuts, the other in which it superimposes as many cuts as possible. It finishes with a variation of sets of linkings in which each linking has its own local set of cut pairs.

[27] Thus, by definition, θ is sequentialisable iff it is cut-sequentialisable.

(D_1)

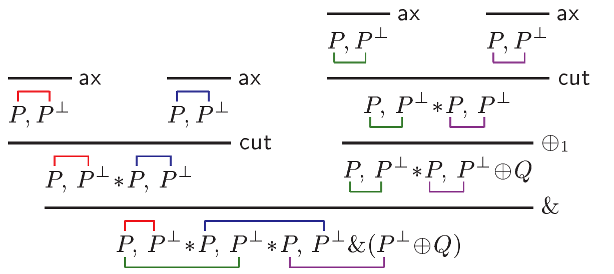

(D_2)

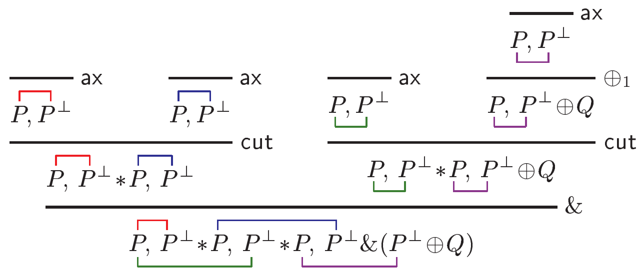

(D_3)

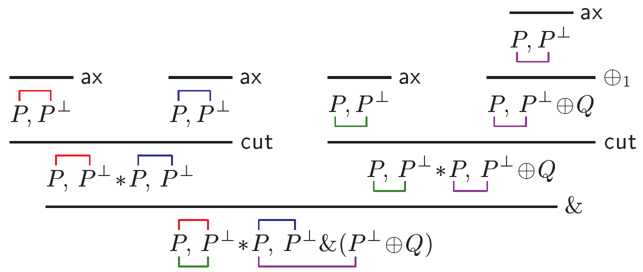

(D_4)

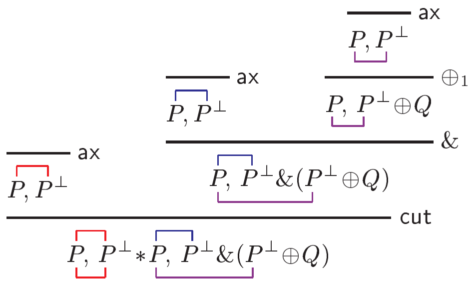

Fig. 17. Translating each MALL$^{\mathsf{cut}}$ derivation D_i of Figure 15 into a set of linkings on a cut sequent, using the function defined inductively in Table III.

5.3.3 *The Sequentialisation Theorem*

THEOREM 5.9 SEQUENTIALISATION. *A set of linkings on a cut sequent is a translation of a MALL proof iff it is a proof net.*

PROOF. Section 4.12, the proof of the Separation Lemma, applies verbatim when θ is a proof net on a cut sequent Γ. We adapt the proof of Theorem 4.18 (Cut-free Sequentialisation) in three places to deal with cut. First, in the base case of the primary induction, treat a cut as an outermost tensor. Second, in the case $\Gamma = \Delta, A_0 \& A_1$, garbage collect to ensure that θ_1 and θ_2 satisfy (P0): delete from Γ_i every cut pair without a leaf in θ_i. Finally, if the appeal to the Separation Lemma in case (c) of the primary inductive step (page 38) yields a separating $⅋/\&$-vertex x inside a cut pair $A{*}A^{\perp}$, immediate separation would destroy the complementarity of the cut (since in G_1 a strict subformula of either A or $A^{\perp}$ will have been removed). The following claim will allow us to substitute a tensor $A \otimes A^{\perp}$ for $A * A^{\perp}$, so that lack of complementary is no longer a problem.

Claim: If a cut pair $A * A^{\perp}$ contains a free $⅋/\&$-vertex y, then every linking in θ visits leaves in $A * A^{\perp}$.

Proof. If not, there is $\lambda \in \theta$ and a &-vertex w with the cut c in $\Gamma \restriction \lambda$ but not in $\Gamma \restriction \lambda_w$ for some $\lambda_w \in \theta$ such that $\lambda \overset{w}{=} \lambda_w$. Thus there are jumps $l \to w$ and $l' \to w$ in $\mathcal{G}_\theta$ for leaves $l \in A$ and $l' \in A^{\perp}$. By domination property FORK (page 33, with $y_0 \dots y_n$ as the path from l down to c and back up to l'), $w \sqsupset c$, and hence, by EXTEND (travelling up from c to y), $w \sqsupset y$, contradicting the freeness of y. ■

The Separation Lemma always yields a separating *free* $⅋/\&$ (see footnote 19, page 37), thus x is free, and the claim with $y = x$ implies every linking in θ visits $A * A^{\perp}$. Therefore θ remains a proof net upon replacing $A * A^{\perp}$ by $A \otimes A^{\perp}$, and the argument in the proof of Theorem 4.18 yields a sequentialisation. This sequentialisation remains a MALL proof upon replacing the tensor by a cut, and that MALL proof is a sequentialisation of θ. □

5.3.4 *Alternative notions of sequentialisation.* This section explores two alternative definitions of sequentialisation. It concludes with a variation on sets of linkings in which each linking has its own local set of cut pairs.

For reference, recall the projection/translation diagram from page 47 involving the MALL proofs of Figure 16 (page 46) and the MALL$^{\mathsf{cut}}$ derivations of Figure 15 (page 45), and the resulting sequentialisation $\overset{s}{\longrightarrow}$ relation:

$$\begin{tikzcd}[row sep=small]
 & D_1 \arrow[dl] \arrow[dr] & \\
\Pi & & \theta \\
 & D_2 \arrow[dl] \arrow[ur] & \\
\Pi' & & \\
 & D_3 \arrow[ul] \arrow[dr] & \\
\Pi'' & & \phi \\
 & D_4 \arrow[ul] \arrow[ur] &
\end{tikzcd}$$

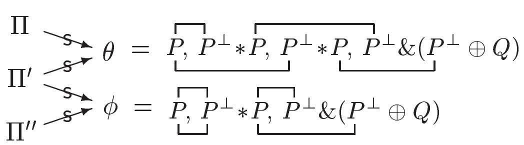

Superimposing no cuts. To obtain a deterministic translation (i.e., a function) from MALL proofs (including the cut rule) to sets of linkings, we can force Σ to be empty in the &-rule in Table II (page 45), i.e., "never superimpose cuts". Let MALL$^{\mathsf{cut}}_{\mathsf{sep}}$ (with sep standing for "keep cuts separate") be the restriction of MALL$^{\mathsf{cut}}$ obtained by replacing the &-rule in Table II with

$$\frac{[\Omega]\,\Gamma, A \qquad [\Omega']\,\Gamma, B}{[\Omega, \Omega']\;\Gamma, A\&B}\,\&$$

The restriction to $\mathrm{MALL}^{\mathsf{cut}}_{\mathsf{sep}}$ of the projection from $\mathrm{MALL}^{\mathsf{cut}}$ derivations to MALL proofs is a bijection: every MALL proof Π is the projection of a unique $\mathrm{MALL}^{\mathsf{cut}}_{\mathsf{sep}}$ derivation. Thus the translation becomes a function. For example, of the derivations D_i in Figure 15, only D_3 is not in $\mathrm{MALL}^{\mathsf{cut}}_{\mathsf{sep}}$, so the projection/translation diagram above restricts to:

$$\begin{tikzcd}[row sep=small]
 & D_1 \arrow[dl] \arrow[dr] & \\
\Pi & & \theta \\
 & D_2 \arrow[dl] \arrow[ur] & \\
\Pi' & & \\
\Pi'' & & \phi \\
 & D_4 \arrow[ul] \arrow[ur] &
\end{tikzcd}$$

yielding a function from MALL proofs to sets of linkings:

$$\begin{tikzcd}[row sep=small]
\Pi \arrow[dr] & \\
 & \theta \\
\Pi' \arrow[ur] & \\
 & \phi \\
\Pi'' \arrow[ur] &
\end{tikzcd}$$

By keeping cuts separate, never superimposing them (i.e., by going via $\mathrm{MALL}^{\mathsf{cut}}_{\mathsf{sep}}$ rather than $\mathrm{MALL}^{\mathsf{cut}}$), we have obtained a function from MALL proofs to sets of linkings. However, the notion of proof net defined above, which was a simple and natural extension of our cut-free definition, producing a nice category of proof nets, does not characterise the image of this function; rather, it characterises the image of the original $\mathrm{MALL}^{\mathsf{cut}}$-based $\xrightarrow{\mathsf{s}}$ relation (or equivalently, the image of the translation function from $\mathrm{MALL}^{\mathsf{cut}}$ derivations to sets of linkings on cut sequents, Table III (page 46)). For example, the following proof net can be derived only via an instance of the &-rule in Table III which superimposes cuts (Σ non-empty in the rule)

$$P \oplus P,\ P^\perp * P,\ P^\perp \& P^\perp$$

and is therefore the translation of a $\mathrm{MALL}^{\mathsf{cut}}$ derivation, but not a translation of a $\mathrm{MALL}^{\mathsf{cut}}_{\mathsf{sep}}$ derivation; thus it is a proof net beyond the image of the cut-separating function defined above.

Characterising the image of the cut-separating function would require additional conditions in the definition of proof net. For example, that every cut must have monomial weight is necessary (though not sufficient, as witnessed by the proof net above, in which the cut has monomial weight). The kernel of the cut-separating function does not include the commutation of the cut rule with the & rule. For example, the function maps the MALL proofs Π' and Π'' (Figure 16, page 46) to distinct proof nets θ and ϕ, yet the proofs differ only by a commutation of cut- and &-rules (the first rule commutation in Figure 13 (page 31), with cut in place of $\otimes$).

$$(D)\qquad \dfrac{\dfrac{\dfrac{\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}{*}P,\,P^{\perp}}\,\mathsf{cut}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}{*}P,\,P^{\perp}\boxed{*}P,\,P^{\perp}}\,\boxed{\mathsf{cut}}\qquad\dfrac{\dfrac{\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}{*}P,\,P^{\perp}}\,\mathsf{cut}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}{*}P,\,P^{\perp}\boxed{*}P,\,P^{\perp}}\,\boxed{\mathsf{cut}}}{P,\,P^{\perp}{*}P,\,P^{\perp}\boxed{*}P,\,P^{\perp}\&P^{\perp}}\,\&$$

$$(D')\qquad \dfrac{\dfrac{\dfrac{\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}{*}P,\,P^{\perp}}\,\mathsf{cut}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}{*}P,\,P^{\perp}\boxed{*}P,\,P^{\perp}}\,\boxed{\mathsf{cut}}\qquad\dfrac{\dfrac{\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}\boxed{*}P,\,P^{\perp}}\,\boxed{\mathsf{cut}}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}{*}P,\,P^{\perp}\boxed{*}P,\,P^{\perp}}\,\mathsf{cut}}{P,\,P^{\perp}{*}P,\,P^{\perp}\boxed{*}P,\,P^{\perp}\&P^{\perp}}\,\&$$

$$(\Pi)\qquad \dfrac{\dfrac{\dfrac{\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}}\,\mathsf{cut}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}}\,\mathsf{cut}\qquad\dfrac{\dfrac{\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}}\,\mathsf{cut}\quad\dfrac{}{P,\,P^{\perp}}\,\mathsf{ax}}{P,\,P^{\perp}}\,\mathsf{cut}}{P\ P\&P^{\perp}}\,\&$$

Fig. 18. Derivations D and D' in $\mathrm{MALL}^{\mathsf{cut}}_{\mathsf{sup}}$, the restriction of $\mathrm{MALL}^{\mathsf{cut}}$ in which the &-rule superimposes as many cuts as possible, followed by the MALL proof Π to which both D and D' project. In each derivation, one cut occurrence, together with the rule that introduces it, has been marked, so that that the two derivations can be distinguished. The difference between the derivations is that in D the marked cut $\boxed{*}$ is the last cut introduced in each branch of the &-rule, whereas in D' the marked cut $\boxed{*}$ is the last cut introduced in the left branch but the first cut introduced in the right branch.

Superimposing as many cuts as possible. We discussed above the possibility of taking Σ in the &-rule of Table II (page 45) minimal, i.e., empty, yielding a function from MALL proofs to sets of linkings. The alternative of taking Σ maximal, i.e., "superimpose as many cuts as possible", does not define a function, since there may be a choice of how to identify cuts.

Let $\mathrm{MALL}^{\mathsf{cut}}_{\mathsf{sup}}$ (with sup standing for "superimpose as many cuts as possible") be the restriction of $\mathrm{MALL}^{\mathsf{cut}}$ obtained by limiting the &-rule in Table II with the side condition that Ω and Ω' have no common cut pair. Two $\mathrm{MALL}^{\mathsf{cut}}_{\mathsf{sup}}$ derivations D and D' are shown in Figure 18, followed by their common projection to a MALL proof Π. The derivations differ only in how they choose to superimpose two equal cuts. Figure 19 (page 52) translates each derivation into a set of linkings on a cut sequent. Let θ and θ' be the translations of D and D' respectively. Then we have the following projection/translation relationship:

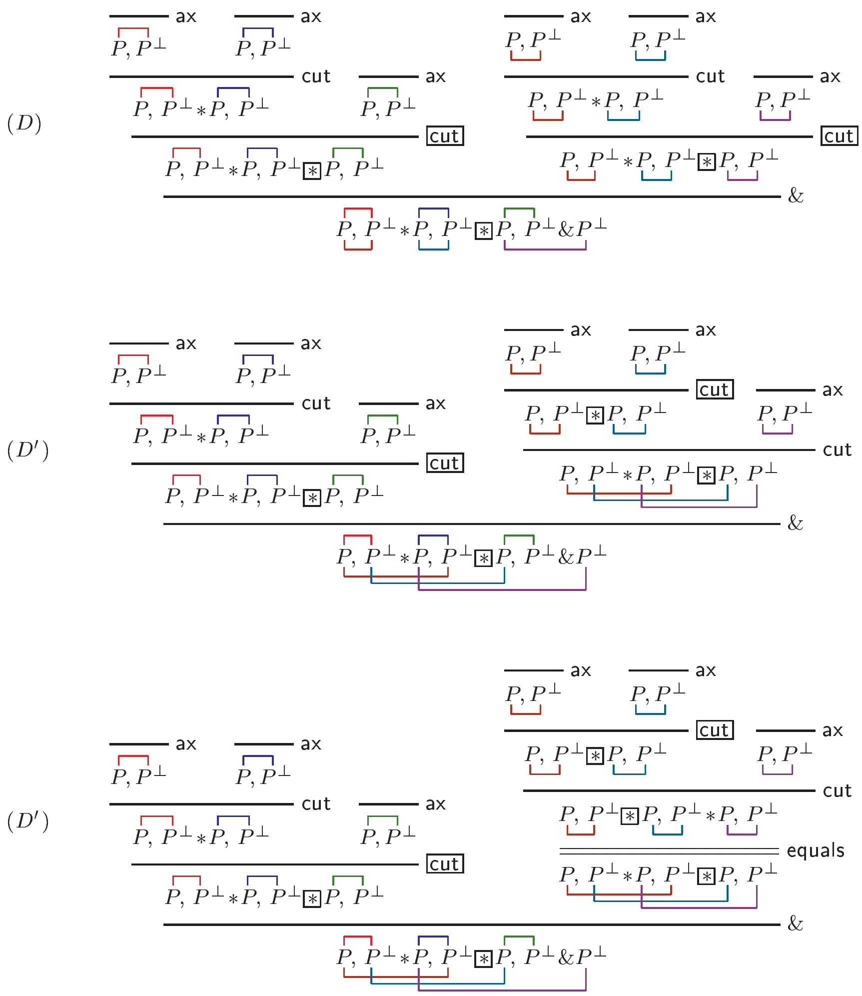

Fig. 19. Translating each $\mathrm{MALL}^{\mathsf{cut}}_{\mathsf{sup}}$ derivation D and D' of Figure 18 (page 51) into a set of linkings on a cut sequent. The second of the two depictions of the inductive translation of D' includes an equality, to help the reader track the superposition of literals and cuts.

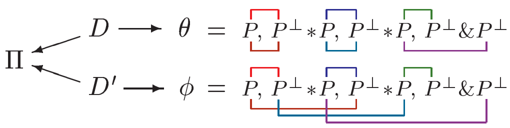

yielding the following relation between the MALL proof and the two sets of linkings:

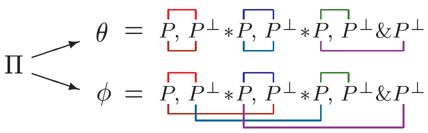

Girard was aware of this issue in the context of monomial proof nets; see Appendix A.1.6 of [1996].

Local cuts. A final variation is to depart from sets of linkings on a fixed cut sequent, and permit each linking its own set of cut pairs. This yields a deterministic translation (function) from MALL proofs. Define a ***cut linking*** on a MALL sequent Γ as a linking on a cut sequent $[\Sigma]\,\Gamma$ with Σ a disjoint union of cut pairs. In order to abstract from the identity of the cut pairs we consider Σ (but not Γ) up to isomorphism.

Every MALL proof Π of Γ yields a set of cut linkings on Γ in the obvious way: each &-resolution R of Π (any result of deleting one branch above each &-rule of Π) yields a cut linking λ_R on Γ by downwards tracking of leaves; the axiom (resp. cut) rules of R are in bijection with the axiom links (resp. cut pairs) of λ_R. This translation identifies more proofs than the translations discussed above. All three MALL proofs in Figure 16 translate to the same set of cut linkings, the pair

$$P,\ P^{\perp} * P,\ P^{\perp} \& (P^{\perp} \oplus Q)$$

$$P,\ P^{\perp} * P,\ P^{\perp} \& (P^{\perp} \oplus Q)$$

Since there is no information indicating how to identify cuts between different cut linkings it is not immediately clear how to define a meaningful correctness criterion to characterise the image of the translation. All we have is that a set of cut linkings is the translation of a proof iff if can be obtained from a proof net as in Definition 5.3 by localising the cut pairs to the linkings in whose cut-additive resolution they occur. Note that this localisation erases the differences between the superposition variants of sequentialisation discussed above. (In other words, translation to a set of cut linkings factorises through any of the translations considered above to a set of linkings on a cut sequent.)

5.4 Proof that eliminating a cut from a proof net yields a proof net

In this section we establish that cut elimination preserves (P0)–(P3). Preservation of (P0) is trivial. Preservation of (P1) for a literal or multiplicative cut is also trivial; for an additive cut it is an immediate consequence of the following lemma.

Lemma 5.10. *Let $A{*}A^{\perp}$ be an additive cut pair in a cut sequent Γ with $A = A_0 \& A_1$ and $A^{\perp} = A_0^{\perp} \oplus A_1^{\perp}$ (or vice versa), and let λ, λ' be linkings of a proof net on Γ such that the cut & is the only & toggled by $\{\lambda, \lambda'\}$. Then λ and λ' take the same argument of $A^{\perp}$, i.e., exactly one of $A_0^{\perp}$ and $A_1^{\perp}$ is in both $\Gamma \restriction \lambda$ and $\Gamma \restriction \lambda'$.*

Proof. If λ and λ' took opposite arguments of $A^{\perp}$, a leaf of $A^{\perp}$ would depend on the cut &. The resulting jump yields a switching cycle of $\mathcal{G}_{\{\lambda,\lambda'\}}$ containing the only & toggled by $\{\lambda, \lambda'\}$, violating (P3). $\square$

Preservation of (P2) is straightforward for a literal or additive cut, since $⅋$-switchings correspond before and after the elimination. For the multiplicative case, consider a linking λ on Γ, and let Γ' be Γ after eliminating a multiplicative cut. Any switching cycle C' of λ on Γ' induces a switching cycle C of λ on Γ: if C' doesn't traverse both new cuts of Γ', obtain C by re-routing a possible passage through a cut of Γ' to go through the cut of Γ instead; if it does, a portion of C' yields a switching cycle via the cut or cut tensor of Γ. Thus switching acyclicity is preserved. Balance (see Section 4.7.1, page 27, but counting a cut as a tensor) is preserved (for we lose a tensor and gain a cut), so (P2) is preserved.

The remainder of this section proves that cut elimination preserves (P3).

Fix a proof net θ on a cut sequent Γ. We localise the notion of domination of Section 4.12 from θ to any saturated set of linkings $\Lambda \subseteq \theta$. Write $x \rightarrow_\Lambda y$ if the edge $x \rightarrow y$ of $\mathcal{G}_\theta$ is in $\mathcal{G}_\Lambda$. A set X of vertices in $\mathcal{G}_\Lambda$ is an ***x-zone under*** Λ if for all $y \in X$ there exists z with $y \Rightarrow_X z \rightarrow_\Lambda x$. Given a $⅋/\&$-vertex $x \in \mathcal{G}_\Lambda$ and a vertex $y \in \mathcal{G}_\Lambda$, define x ***dominates*** y ***in*** Λ, denoted $x \sqsupset_\Lambda y$, if $y \in X$ for some x-zone X under Λ. The domination properties of Lemma 4.27 localise from θ to any saturated set of linkings $\Lambda \subseteq \theta$, as follows:

Localised Lemma 4.27 Properties of localised domination.

- l-switch. *If $x \leftarrow_\Lambda y$ is a switch edge then $x \sqsupset_\Lambda y$.*[28]
- l-transitivity. *Localised domination $\sqsupset_\Lambda$ is transitive.*
- l-self. *Let x be a $⅋/\&$. Then $x \sqsupset_\Lambda x$ iff x is in a switching cycle of $\mathcal{G}_\Lambda$.*
- l-jump-cycle. *If $w \leftarrow l$ is a jump in $\mathcal{G}_\Lambda$ and l is in a switching cycle C of $\mathcal{G}_\Lambda$, then $w \sqsupset_\Lambda y$ for all vertices $y \in C$.*
- l-extend. *If $x \sqsupset_\Lambda y_0$ and there is a path $y_0 \ldots y_n$ in $\mathcal{G}_\Lambda$ which never enters a $⅋/\&$ from above (i.e., $y_{i-1} \rightarrow_\Lambda y_i$ only if y_i is not a $⅋/\&$), then $x \sqsupset_\Lambda y_n$.*
- l-fork. *Let x be a $⅋/\&$ and $y_0 \ldots y_n$ a switching path in $\mathcal{G}_\Lambda$ with $y_0 \rightarrow_\Lambda x \leftarrow_\Lambda y_n$. Then $x \sqsupset_\Lambda y_i$ for each i.*
- l-meet. *If $x \sqsupset_\Lambda y \sqsubset_\Lambda z$ for distinct free $⅋/\&$-vertices x and z, there exists a switching path $xy_0 \ldots y_n z$ in $\mathcal{G}_\Lambda$ with $x \leftarrow_\Lambda y_0$ and $y_n \rightarrow_\Lambda z$.*[28]

Proof. Make the following substitutions in the proofs of the original domination properties in Lemma 4.27 (page 32): Λ for θ, $\sqsupset_\Lambda$ for $\sqsupset$, $\rightarrow_\Lambda$ for $\rightarrow$, and *zone under* Λ for *zone*. $\square$

Lemmas 4.32 and 4.33 of Section 4.12 localise similarly.

Localised Lemma 4.32. *For every non-empty union S of switching cycles of $\mathcal{G}_\Lambda$ there is a jump $l \rightarrow_\Lambda w$ from a leaf $l \in S$ to a $\&$-vertex $w \notin S$ toggled by Λ.*

Proof. A relatively straightforward adaptation of the proof of the original Lemma 4.32 (page 35). Let Λ_m be a minimal saturated subset of Λ with $\mathcal{G}_{\Lambda_m}$ containing S. Switchings of singleton sets of linkings are cycle-free by (P2), so Λ_m contains at least two linkings. Let w be a $\&$ toggled by Λ_m that is not in any switching cycle of $\mathcal{G}_{\Lambda_m}$ (existing by (P3)), so $w \notin S$. Since $\Lambda_m \subseteq \Lambda$, w is certainly toggled by Λ. Since Λ_m is minimal, $S \not\subseteq \mathcal{G}_{\Lambda_m^w}$ (using (S1)), so some edge e of S is in $\mathcal{G}_{\Lambda_m}$ but not in $\mathcal{G}_{\Lambda_m^w}$. Without loss of generality e is an edge from a leaf l, because for any other edge $y \rightarrow x$ in S we have $l \rightarrow z_1 \rightarrow \ldots \rightarrow z_n = y \rightarrow x$ in S for some leaf l, and $y \rightarrow x$ is in $\mathcal{G}_{\Lambda_m^w}$ whenever $l \rightarrow z_1$ is in $\mathcal{G}_{\Lambda_m^w}$. By Lemma 4.31 the jump $l \rightarrow w$ is in $\mathcal{G}_{\Lambda_m}$, hence also in $\mathcal{G}_\Lambda$. $\square$

[28] We shall not actually use this localised property in the proof that cut elimination is well defined on proof nets; we include the property here to maintain the correspondence with the original Lemma 4.27.

Localised Lemma 4.33. *If $x \sqsupset_\Lambda x$ then $y \sqsupset_\Lambda x$ for some &-vertex $y \not\sqsupset_\Lambda y$ toggled by Λ.*

Proof. We essentially repeat the original proof of Lemma 4.33 (page 35). By domination property l-self, x is in a switching cycle of $\mathcal{G}_\Lambda$. Iterate Localised Lemma 4.32, adding switching cycles until jumping to a &-vertex y not in a switching cycle of $\mathcal{G}_\Lambda$. Then $y \sqsupset_\Lambda x$ by l-jump-cycle and l-transitivity, and $y \not\sqsupset_\Lambda y$ by l-self. □

Proof that cut elimination preserves the toggling condition (P3)

Preservation is immediate for the elimination of a literal cut pair $P * P^\perp$, since for every set Λ of linkings on Γ, the &-vertices toggled by Λ and the switching cycles of $\mathcal{G}_\Lambda$ correspond before and after the elimination. Thus consider the elimination of an additive cut pair $(A_0 \& A_1) * (A_0^\perp \oplus A_1^\perp)$ or multiplicative cut pair $(A_0 \parr A_1) * (A_0^\perp \otimes A_1^\perp)$.

Let θ' on the cut sequent Γ' be the result of eliminating $(A_0 \& A_1) * (A_0^\perp \oplus A_1^\perp)$ or $(A_0 \parr A_1) * (A_0^\perp \otimes A_1^\perp)$ from the proof net θ on Γ. Let x be the & or $\parr$ and y the $\oplus$ or $\otimes$ of the cut, let x_0, x_1 and y_0, y_1 be the arguments of x and y respectively, and let c be the cut vertex $*$ between x and y.

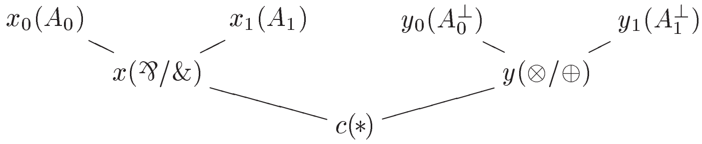

Thus in Γ' each of c, x and y have been deleted, and cut vertices c_0 between x_0 and y_0 and c_1 between x_1 and y_1 have been added,

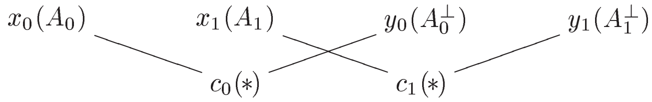

unless one of $A_0, A_0^\perp$ or $A_1, A_1^\perp$ disappeared in the 'garbage collection' phase of additive elimination, in which case only one of c_0 or c_1 is present.

Suppose θ' fails (P3), i.e., there exists a set of two or more linkings $\Lambda' \subseteq \theta'$ such that every & in Γ' toggled by Λ' is in a switching cycle of $\mathcal{G}_{\Lambda'}$.

Lemma 5.11. *There exists a saturated set of linkings $\Lambda \subseteq \theta$ on Γ such that Λ on Γ toggles the same &s as Λ' on Γ', except perhaps x in addition; x is toggled by Λ on Γ iff the cut is additive and there are $\lambda_l, \lambda_r \in \Lambda'$ such that x_0 is present in $\Gamma \restriction \lambda_l$ and x_1 is present in $\Gamma \restriction \lambda_r$.*

Proof. Since eliminating an additive or multiplicative cut at most deletes linkings, Λ' can also be viewed as a set of linkings on Γ, and $\Lambda' \subseteq \theta$. Furthermore, Λ' on Γ toggles exactly the same &s as Λ' on Γ', except perhaps x in addition (in the case indicated in the lemma). Let Λ be a minimal saturated set of linkings of θ on Γ containing Λ'. By minimality, Λ on Γ toggles the same &s as Λ' on Γ. □

Lemma 5.12. *The vertex y is not in a switching cycle of $\mathcal{G}_\Lambda$.*

Proof. If y is in a switching cycle, by l-self then Localised Lemma 4.33, Λ toggles a &-vertex $w \sqsupset_\Lambda y$ with $w \not\sqsupset_\Lambda w$. By l-self w is in no switching cycle of $\mathcal{G}_\Lambda$, and

$w \sqsupset_\Lambda x$ by L-EXTEND. Necessarily $w \neq x$, otherwise $w \sqsupset_\Lambda w$, a contradiction. By Lemma 5.11, w is toggled by Λ' on Γ', hence[29] is in a switching cycle C of $\mathcal{G}_{\Lambda'}$.

Suppose C does not go through both c_0 and c_1. Then C induces a switching cycle of $\mathcal{G}_\Lambda$, still containing w, obtained by re-routing a possible passage through c_0 or c_1 to go through c instead, a contradiction.

Suppose C goes through both c_0 and c_1. Re-routing both passages to go through c instead either yields two switching cycles through c with w in one of them, a contradiction, or yields a switching cycle C_y through y and a switching path $\pi_x = z_0 \ldots z_n$ in $\mathcal{G}_\Lambda$ with $z_0 \to_\Lambda x$ and $z_n \to_\Lambda x$, such that w is either in C_y or π_x. The first possibility immediately yields a contradiction, so assume $w \in \pi_x$. By L-FORK $x \sqsupset_\Lambda w$, so by L-TRANSITIVITY $w \sqsupset_\Lambda w$, a contradiction. □

LEMMA 5.13. *Every &-vertex $v \neq x$ toggled by Λ on Γ is in a switching cycle of $\mathcal{G}_\Lambda$.*

PROOF. By Lemma 5.11, v is toggled by Λ' on Γ', hence[29] is in a switching cycle C of $\mathcal{G}_{\Lambda'}$. Suppose C goes through c_0 and/or c_1. By re-routing the passage(s) through c_0 and/or c_1 to go through c instead, C induces a switching cycle of $\mathcal{G}_\Lambda$ that contains y, contradicting Lemma 5.12. Thus C avoids c_0 and c_1. Hence C is also a switching cycle of $\mathcal{G}_\Lambda$, containing v. □

COROLLARY 5.14. *If the cut is multiplicative, every & toggled by Λ on Γ is in a switching cycle of $\mathcal{G}_\Lambda$.*

Thus if the cut is multiplicative, θ fails to be a proof net, a contradiction. Henceforth we assume the cut is additive.

LEMMA 5.15. *The &-vertex x is the unique & toggled by Λ that is not in any switching cycle of $\mathcal{G}_\Lambda$.*

PROOF. Since θ is a proof net, Λ toggles a &-vertex v in no switching cycle of $\mathcal{G}_\Lambda$. By Lemma 5.13, necessarily $v = x$. □

Since Λ toggles x, by Lemma 5.11 there are $\lambda_l, \lambda_r \in \Lambda'$ such that $x_0 \in \Gamma \upharpoonright \lambda_l$ and $x_1 \in \Gamma \upharpoonright \lambda_r$. On Γ, every linking of Λ' is consistent, so $y_0 \in \Gamma \upharpoonright \lambda_l$ and $y_1 \in \Gamma \upharpoonright \lambda_r$. No linking in Λ has an additive resolution containing both y_0 and y_1, so $y_0 \notin \Gamma \upharpoonright \lambda_r$. Since Λ is saturated on Γ, there must be a &-vertex u in Γ and $\lambda_0, \lambda_1 \in \Lambda$ such that $y_0 \in \Gamma \upharpoonright \lambda_0$, $y_0 \notin \Gamma \upharpoonright \lambda_1$ and u is the only & toggled by $\{\lambda_0, \lambda_1\}$.

If $y_1 \in \Gamma \upharpoonright \lambda_1$ then for $i = 0, 1$ there are leaves l_i above y_i with jumps $l_i \to_\Lambda u$; otherwise $y_1 \notin \Gamma \upharpoonright \lambda_1$ so $c \notin \Gamma \upharpoonright \lambda_1$ and $c \in \Gamma \upharpoonright \lambda_0$ and there are leaves l_0 above y and l_1 above x with jumps $l_i \to_\Lambda u$. In either case y lies on a switching path from l_0 to l_1, so we have $u \sqsupset_\Lambda y$ by L-FORK. Using L-EXTEND we obtain $u \sqsupset_\Lambda x$.

If $u = x$, then by L-SELF x is in a switching cycle in $\mathcal{G}_\Lambda$, a contradiction. Thus $u \neq x$, so by Lemma 5.13 u is in a switching cycle of $\mathcal{G}_\Lambda$, hence $u \sqsupset_\Lambda u$ by L-SELF, so $x \sqsupset_\Lambda u$ by Localised Lemma 4.33, Lemma 5.15 and L-SELF. Thus $x \sqsupset_\Lambda x$ by L-TRANSITIVITY, so by L-SELF x is in a switching cycle of $\mathcal{G}_\Lambda$, a contradiction.

This completes the proof that eliminating a cut preserves the toggling condition (P3), and hence the proof that cut elimination is well-defined on proof nets (Proposition 5.4, page 41).

[29] Recall that Λ' was chosen as a witness to the failure of (P3) for θ': any & in Γ' toggled by Λ' is in a switching cycle of $\mathcal{G}_{\Lambda'}$.

6 Girard's monomial proof nets

In this section we relate our MALL proof nets to the monomial proof nets of Girard [1996], whose definition is recalled in Appendix A. We begin in Section 6.1 with a detailed explanation of why monomial proof nets are unsatisfactory. (The reader unfamiliar with monomial proof nets should be able to follow the general shape of the discussion.) In Section 6.2 we settle the open question of whether Girard's criterion becomes insufficient without partitioning weights into monomials, by presenting a non-monomial proof structure which is not sequentialisable, yet satisfies Girard's criterion. In Section 6.3 we examine the similarities between Girard's proof structures and our own. In particular, we show that the resolution condition (P1) corresponds to Girard's so-called technical condition. We show that this characterisation of Girard's technical condition, and several others, remains valid when not requiring Girard's dependency condition, demanding the partitioning of weights into monomials. By contrast, without partitioning weights into monomials the reformulation of the technical condition by Abramsky and Melliès [1999] is no longer valid. In Section 6.4 we define a map collapsing Girard's proof structures to our own which preserves the property of being a sequentialisation of any given MALL proof. Hence this mapping also collapses Girard's proof nets to ours, in the cut-free case providing a surjection from the former to the latter.

6.1 Why monomial proof nets are unsatisfactory

We give a detailed account of how monomial proof nets [Girard 1996] fail to provide abstract representations of cut-free MALL proofs modulo rule commutation. A single cut-free proof may correspond to a host of monomial proof nets, and there is no natural map from cut-free MALL proofs onto monomial proof nets.

Consider the following pair of cut-free monomial proof nets:

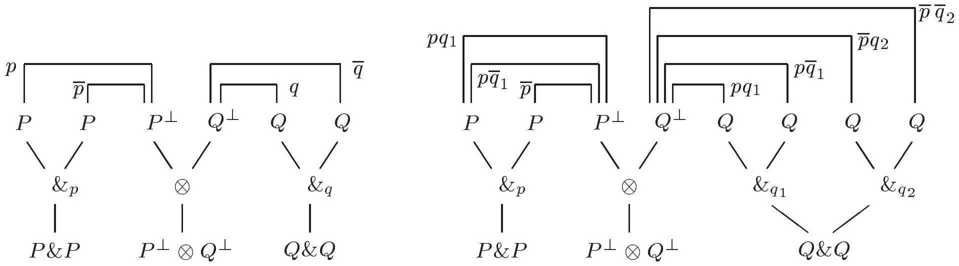

Eigenvariables associated with &s are shown as subscripts; we omit implied weights. These two monomial proof nets correspond to the same proof. The second monomial proof net has two forms of redundancy relative to the first: (i) the & with eigenweight q has been replaced by two similar 'copies', and (ii) the axiom link with weight p has been split into two.

Even if one attempts to fix a choice of representation (*e.g.* favouring the first monomial proof net above over the second), one still runs into difficulty. As a concrete illustration, we exhibit cut-free proofs Π_α and monomial proof nets θ_β for which the binary relation of sequentialisation is the "zig-zag" shown in Figure 20. Define Π_{tqp} to be the proof shown in Figure 21, where $A = (R \parr R) \parr (R^\perp \otimes R^\perp)$, id denotes the identity proof (in which the left R of A descends from the left axiom rule) and tw denotes the twist proof (in which the left R of A descends from the

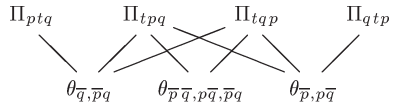

Fig. 20. The "zig-zag" sequentialisation relation between four proofs Π_α and three monomial proof nets θ_β, demonstrating that monomial proof nets fail to provide canonical representations of proofs. The proofs and nets are defined in the main text. (By redundancies of type (i) and (ii) described at the beginning of Section 6.1, there are in addition a host of other monomial proof nets θ_β which parody the three above, and also sequentialise to the Π_α.) By contrast, we represent all the Π_α by the same proof net. Thus the surjection from cut-free Girard proof nets to ours (defined in Section 6.4) is not injective.

$$\dfrac{\dfrac{\dfrac{P,Q,Q^\perp\otimes P^\perp \qquad \overset{\mathsf{id}}{\vdots}\atop A}{P,Q,(Q^\perp\otimes P^\perp)\otimes A}\,t \qquad \dfrac{P,Q,Q^\perp\otimes P^\perp \qquad \overset{\mathsf{tw}}{\vdots}\atop A}{P,Q,(Q^\perp\otimes P^\perp)\otimes A}\,t}{P,Q\&Q,(Q^\perp\otimes P^\perp)\otimes A}\,q \qquad \dfrac{\dfrac{P,Q,Q^\perp\otimes P^\perp \qquad \overset{\mathsf{tw}}{\vdots}\atop A}{P,Q,(Q^\perp\otimes P^\perp)\otimes A}\,t \qquad \dfrac{P,Q,Q^\perp\otimes P^\perp \qquad \overset{\mathsf{tw}}{\vdots}\atop A}{P,Q,(Q^\perp\otimes P^\perp)\otimes A}\,t}{P,Q\&Q,(Q^\perp\otimes P^\perp)\otimes A}\,q}{P\&P,Q\&Q,(Q^\perp\otimes P^\perp)\otimes A}\,p$$

Fig. 21. The proof Π_{tqp}. (We omit the unique cut-free proof of $P,Q,Q^\perp\otimes P^\perp$.)

right axiom rule). Let Π_{tpq} be the result of commuting rules q and p in Π_{tqp}, let Π_{qtp} be the result of commuting t and q in the right half of Π_{tqp}, and let Π_{ptq} be the result of commuting t and p in the right half of Π_{tpq}. Define the monomial proof nets θ_β as follows. To specify θ_β it suffices to present a configuration of weighted axiom links. On P and Q literals, fix the configuration as below-left:

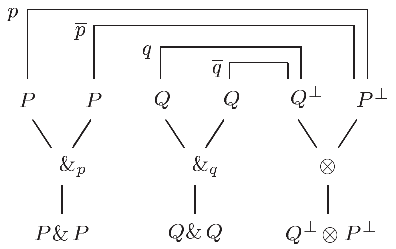

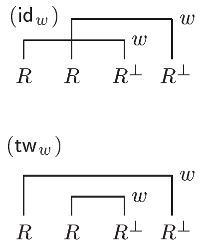

We have taken as eigenweights the labels of the &-rules of the Π_α. The configuration of axiom links on A will be a disjoint union of axiom links in the identity and twist configurations: id_w and tw_w (above-right) denote a pair of axiom links of weight w in the identity and twist configurations, respectively. We specify the θ_β by the following disjoint unions of weighted identity and twist configurations on A:

$$\begin{aligned}
\theta_{\overline{p},p\overline{q}}&\colon\ \mathsf{id}_{pq}\sqcup\mathsf{tw}_{\overline{p}}\sqcup\mathsf{tw}_{p\overline{q}}\\
\theta_{\overline{q},\overline{p}q}&\colon\ \mathsf{id}_{pq}\sqcup\mathsf{tw}_{\overline{q}}\sqcup\mathsf{tw}_{\overline{p}q}\\
\theta_{\overline{p}\,\overline{q},p\overline{q},\overline{p}q}&\colon\ \mathsf{id}_{pq}\sqcup\mathsf{tw}_{\overline{p}\,\overline{q}}\sqcup\mathsf{tw}_{p\overline{q}}\sqcup\mathsf{tw}_{\overline{p}q}
\end{aligned}$$

The sequentialisation relation between the Π_α and the θ_β is as indicated in Figure 20. Since the Π_α are equivalent modulo rule commutation, any satisfactory theory of proof nets should provide a canonical representation uniting all of them.

With monomial proof nets one would have to close under the "zig-zag" sequentialisation relation between proofs and monomial proof nets depicted in Figure 20, thereby matching the set of proofs Π_α and the set of monomial proof nets θ_β, and then artificially choose a representative from amongst the θ_β.

In contrast, in our setting each Π_α maps to the same proof net: the four-linking proof net in Figure 5 (page 7). Thus we preserve the spirit of MLL proof nets by providing an abstract representation of all of the Π_α in one.

6.2 Girard's criterion is insufficient without monomials

A key stepping-stone towards our formulation of a new definition of proof net was to first settle the open question of whether Girard's proof net correctness criterion [1996] becomes insufficient upon relaxing the *dependency condition*, demanding that weights be monomial. The answer is yes: in Figure 22 we present a cut-free non-monomial Girard proof net θ which is not sequentialisable. By ***non-monomial Girard proof net*** we mean a proof net as in [Girard 1996] but for the omission of the dependency condition.

Figure 22 also encodes one of our proof structures θ, via the notion of weight described in Section 4.8. It is not a proof net, since (P3) fails: $\mathcal{G}_\theta$ contains a switching cycle passing through all four &s (follow the four jumps $R_{i+1} \to \&_{p_i}$ (mod 4)).

6.3 The resolution condition is equivalent to Girard's technical condition

Define an ***elementary Girard proof structure*** as a proof structure Θ as in [Girard 1996] (reproduced in Appendix A, Definition 3 on page 66) but for the omission of the dependency condition, the requirement that weights be non-zero, and Girard's technical condition (T):

(T) If v is any element of the boolean algebra generated by the weights occurring in Θ, and x is a &-link, then $v.\neg w(x)$ does not depend on p_x, where $w(x)$ is the weight of x and p_x is the eigenweight of x (see Definition 2 on page 66).

Note that the weights in an elementary Girard proof structure Θ are completely determined by the weights of the axiom links of Θ: every weight occurring in Θ is a (disjoint) sum of weights on axiom links. The weights in Θ are also completely determined by the function allocating a slice of Θ to each valuation of Θ (see Definition 6 in Appendix A). Each slice is completely determined by the set of axiom links of Θ that occur in that slice: it consist of the formulas and links "below" these axiom links. The axiom links of a slice partition the set of literal occurrences in that slice.

Define a ***&-resolution*** of an elementary Girard proof structure Θ as any result of deleting one premise of every &-link of Θ, as well as, recursively, all links with a deleted conclusion and all premises of deleted links. For any Θ-valuation φ (see Definition 6) let Θ^φ be the &-resolution of Θ obtained by deleting, for every &-link x, its left premise if $\varphi(p_x) = 0$ and its right premise if $\varphi(p_x) = 1$. A ***cut-additive resolution*** of Θ is any result of deleting any number of cut links (possibly zero) from a &-resolution of Θ, and for every formula A, all but one link with conclusion A, as well as, recursively, all links with a deleted conclusion and all premises of deleted links. It is not hard to see that the three conditions involving

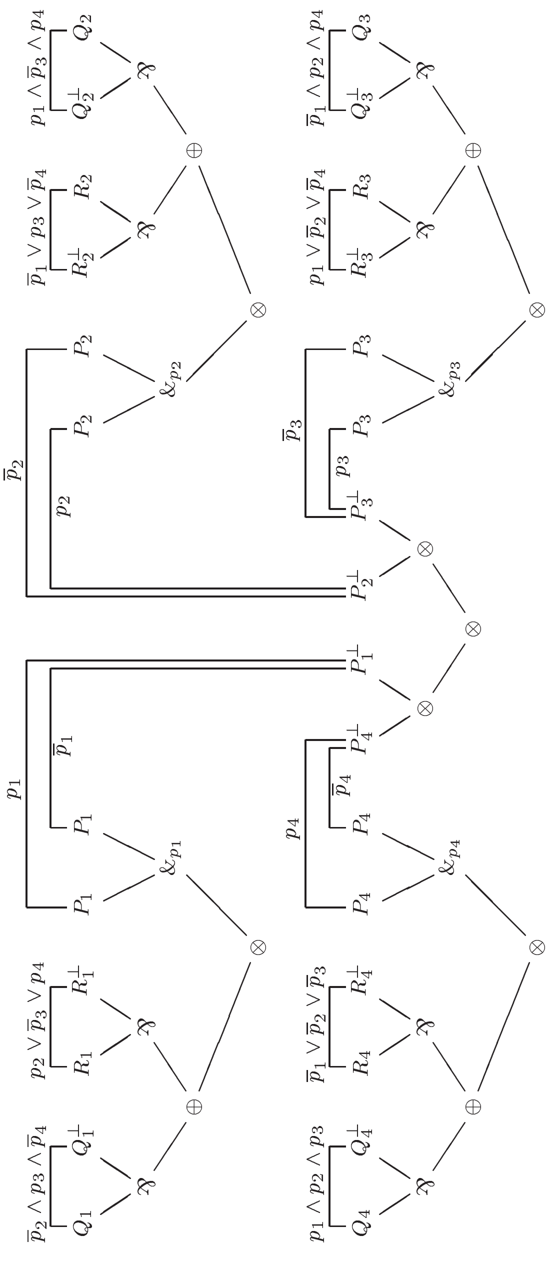

Fig. 22. Girard's correctness criterion is insufficient without monomials: this abbreviated non-monomial Girard proof net is not sequentialisable. Expand the abbreviation as follows: view each p_i as an eigenvariable and split each $\oplus$ into a separate $\oplus_1$ and $\oplus_2$; formulas and remaining weights are implied.

weights in the definition of an elementary Girard proof structure Θ are equivalent to requiring that for every Θ-valuation φ:

- the slice $\varphi(\Theta)$ is a cut-additive resolution of Θ, and
- $\varphi(\Theta) \subseteq \Theta^{\varphi}$, i.e., $\varphi(\Theta)$ resolves &s consistently with φ.

In the context of elementary Girard proof structures Θ we reformulate our conditions (P0) CUT and (P1) RESOLUTION as follows.

$(\mathrm{P0})_G$ Every link of Θ appears in a slice of Θ.

$(\mathrm{P1})_G$ For any &-resolution $\Theta^\star$ of Θ, exactly one slice of Θ is contained in $\Theta^\star$.

Obviously, $(\mathrm{P0})_G$ is equivalent to Girard's condition that weights must be non-zero. Next we will show that Girard's technical condition (T) is equivalent to $(\mathrm{P1})_G$.

For brevity, write $w \not\curvearrowright p_x$ if the weight w does not depend on the eigenweight p_x, i.e., belongs to the boolean algebra generated by the eigenweights distinct from p_x. Write $w \not\curvearrowright_\varphi p_x$ if in the slice $\varphi(\Theta)$ the weight w does not depend on p_x, i.e., if $\varphi(w) = \varphi_x(w)$, where φ_x is the valuation obtained from φ by toggling the 0/1 value of $\varphi(p_x)$ (see Definition 7 in Appendix A, page 68). Thus

$$w \not\curvearrowright p_x \text{ iff } w \not\curvearrowright_\varphi p_x \text{ for all valuations } \varphi. \qquad (*)$$

To prove (T) $\iff$ $(\mathrm{P1})_G$ we shall use the following stepping-stone conditions.

(T^-) If v is any weight occurring in Θ, and x is a &-link, then $v.\neg w(x) \not\curvearrowright p_x$.

This is simply the relaxation of (T) obtained by restricting to weights *in* Θ, rather than ranging over all weights *generated by* weights in Θ.

(SELF-INDEP) No & depends on itself, i.e., for all &-links x of Θ, $w(x) \not\curvearrowright p_x$.

(PRESENCE) For any &-link x and any valuation φ, if $x \notin \varphi(\Theta)$ then $\varphi(\Theta) = \varphi_x(\Theta)$.

PRESENCE says that a link or formula in Θ can depend on a &-link x in a slice $\varphi(\Theta)$ only if $x \in \varphi(\Theta)$. Clearly this is equivalent to

(PRESENCE-AX) An axiom link can depend on a &-link x in a slice $\varphi(\Theta)$ only if $x \in \varphi(\Theta)$.

PROPOSITION 6.1.

$$(\mathrm{T}) \iff (\mathrm{T}^-) \textit{ and } (\textsc{self-indep}) \iff (\textsc{presence}) \iff (\mathrm{P1})_G$$

PROOF. (T) $\implies$ (SELF-INDEP) follows by taking $v = \neg w(x)$ in (T), since $x \not\curvearrowright p_x$ iff $\neg x \not\curvearrowright p_x$; (T) $\implies$ (T^-) is trivial.

(T^-) and (SELF-INDEP) $\implies$ (T) follows because, for all $u \not\curvearrowright p_x$, and any weights w and w':

(1) $[w.u \not\curvearrowright p_x]$ and $[w'.u \not\curvearrowright p_x] \implies (w \vee w').u \not\curvearrowright p_x$;
(2) $[w.u \not\curvearrowright p_x]$ and $[w'.u \not\curvearrowright p_x] \implies (w.w').u \not\curvearrowright p_x$;
(3) $w.u \not\curvearrowright p_x \implies (\neg w).u \not\curvearrowright p_x$.

In turn, (1) and (2) follow from:

(a) $[x \not\curvearrowright p_x]$ and $[y \not\curvearrowright p_x] \implies x \vee y \not\curvearrowright p_x$;

(b) $[x \not\curvearrowright p_x]$ and $[y \not\curvearrowright p_x] \implies x.y \not\curvearrowright p_x$;

respectively, since $(w \vee w').u = (w.u) \vee (w'.u)$ and $(w.w').u = (w.u).(w'.u)$. Finally, for (3): since $u \not\curvearrowright p_x$ (by hypothesis), using (*) above, for any Θ-valuation φ we have $\varphi(u) = \varphi_x(u)$; thus

$$\varphi(w.u) = \varphi_x(w.u) \quad \text{iff} \quad \varphi(w) = \varphi_x(w) \vee \varphi(u) = 0 \quad \text{iff} \quad \varphi(\neg w.u) = \varphi_x(\neg w.u).$$

(PRESENCE) $\implies$ (SELF-INDEP): If (SELF-INDEP) fails, then, using (*), for some valuation φ we have $x \notin \varphi(\Theta)$ and $x \in \varphi_x(\Theta)$, contradicting (PRESENCE).

(T^-) $\iff$ (PRESENCE), given (SELF-INDEP): Below, x ranges over the &-links of Θ, y over all links and formulas in Θ, φ over the Θ-valuations, and p_x is the eigenweight of x.

$$\begin{array}{lll}
(\mathrm{T}^-) & \iff & (\forall x)\ (\forall y)\ w(y).\neg w(x) \not\curvearrowright p_x \\
& \iff & (\forall x)\ (\forall y)\ (\forall \varphi)\ \varphi(w(y).\neg w(x)) = \varphi_x(w(y).\neg w(x)) \qquad \text{(by (*))} \\
& \iff & (\forall x)\ (\forall y)\ (\forall \varphi)\ [y \in \varphi(\Theta) \wedge x \notin \varphi(\Theta)] \text{ iff } [y \in \varphi_x(\Theta) \wedge x \notin \varphi_x(\Theta)] \\
& \iff & (\forall x)\ (\forall y)\ (\forall \varphi)\ x \notin \varphi(\Theta) \text{ implies } [y \in \varphi(\Theta) \text{ iff } y \in \varphi_x(\Theta)] \\
& & \qquad\qquad \text{using } [x \notin \varphi(\Theta) \text{ iff } x \notin \varphi_x(\Theta)] \qquad\qquad \text{(by (SELF-INDEP))} \\
& \iff & (\forall x)\ (\forall \varphi)\ x \notin \varphi(\Theta) \text{ implies } \varphi(\Theta) = \varphi_x(\Theta) \\
& \iff & \text{(PRESENCE)} \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \square
\end{array}$$

(PRESENCE) $\implies$ $(\mathrm{P1})_G$: Existence follows immediately from the observation that $\varphi(\Theta) \subseteq \Theta^{\varphi}$ for each valuation φ. For uniqueness, let φ, φ' be Θ-valuations such that $\varphi(\Theta)$ and $\varphi'(\Theta)$ are contained in the same &-resolution $\Theta^\star$. We must show that $\varphi(\Theta) = \varphi'(\Theta)$. Let x be a &-link for which $\varphi(p_x) \neq \varphi'(p_x)$. We show (below) that $\varphi(\Theta) = \varphi_x(\Theta)$; hence by induction on the number of eigenweights on which φ and φ' differ, we have $\varphi(\Theta) = \varphi'(\Theta)$.

Claim: $x \notin \varphi(\Theta)$ or $x \notin \varphi'(\Theta)$.

Proof: If x is in both $\varphi(\Theta)$ and $\varphi'(\Theta)$, then x is in $\Theta^\star$. This is impossible, since $\Theta^\star$ chooses (say) left for x, and one of $\Theta(\varphi)$ and $\Theta(\varphi')$ chooses right (since $\varphi(p_x) \neq \varphi'(p_x)$). ■

Without loss of generality, $x \notin \varphi(\Theta)$. Thus, by (PRESENCE), $\varphi(\Theta) = \varphi_x(\Theta)$.

$(\mathrm{P1})_G \implies$ (PRESENCE): Suppose $x \notin \varphi(\Theta)$, yet $\varphi(\Theta) \neq \varphi_x(\Theta)$. Now $\varphi_x(\Theta) \subseteq \Theta^{\varphi_x}$ and $\varphi(\Theta) \subseteq \Theta^{\varphi}$, but since $x \notin \varphi(\Theta)$ also $\varphi(\Theta) \subseteq \Theta^{\varphi_x}$, contradicting $(\mathrm{P1})_G$. $\square$

This completes the proof of Proposition 6.1. Thus Girard's technical condition (T) is equivalent to our own resolution condition (P1), expressed in Girard's setting as $(\mathrm{P1})_G$.

The Abramsky-Melliès reformulation of Girard's technical condition is valid only with monomials. Recall Girard's technical condition [1996], and its reformulation by Abramsky and Melliès [1999]:

(T) If v is any element of the boolean algebra generated by the weights occurring in Θ, and x is a &-link, then $v.\neg w(x) \not\curvearrowright p_x$.

(AM) If v is a weight occurring in Θ and x is a &-vertex, if $v \curvearrowright p_x$ then $v < w(x)$.

Here $v \curvearrowright p_x$ denotes that the weight v depends on the eigenweight p_x. By (*) in Section 6.3, this holds iff v depends on p_x in some slice $\varphi(\Theta)$. However, using Girard's dependency (monomial) condition, it also holds iff v depends on p_x in *any* slice $\varphi(\Theta)$ containing a link or formula of weight v in Θ. Using this, it is not hard to see that (AM) is equivalent to (PRESENCE), formulated in Section 6.3.

Unlike the characterisations of (T) presented above, (AM) is an adequate characterisation of (T) only for proof structures in which all weights are monomials. The following is a simple non-monomial example which satisfies (T) but not (AM):

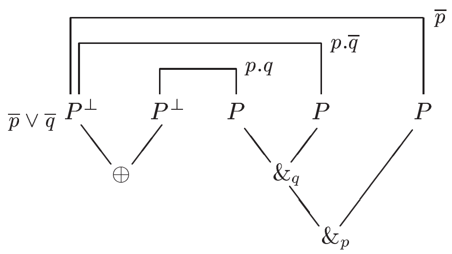

(We draw only the axiom links and literals; compound formulas and remaining weights are implied.) Similarly, the Gustave example (page 24) fails (AM), but satisfies (T). To see that the Gustave example satisfies (T), observe that it satisfies (PRESENCE-AX), formulated on page 61.

Condition (T) implies softness[30] only in conjunction with the monomial condition, since the Gustave proof structure is not soft, yet satisfies (T).

6.4 Collapsing monomial proof structures to our own

Let a ***non-monomial Girard proof structure*** be a proof structure as in [Girard 1996] but for the omission of the dependency condition, i.e., an elementary Girard proof structure satisfying $(\mathrm{P0})_G$ and $(\mathrm{P1})_G$. Define a non-monomial Girard proof structure to be ***compact*** if (a) any non-literal formula occurrence is the conclusion of exactly one link, except that a formula $A \oplus B$ may be the conclusion of both a $\oplus_1$- and a $\oplus_2$-link, and (b) any two literal occurrences constitute the conclusions of at most one axiom link. Each non-monomial Girard proof structure, and thus also each monomial one, can be collapsed into a compact non-monomial Girard proof structure by identifying, along with their premises, links of the same type

[30] Let Θ be a Girard proof structure in which every concluding connective is a $\oplus$, with at least one connective. The dependency (i.e. monomial) and technical conditions together imply that Θ is ***soft***: there exists a concluding formula $C = A \oplus B$ with just one $\oplus$-link as its child (c.f. [Joyal 1995]). *Proof.* Suppose the concluding formulas of Θ, aside from literals, are $C_1, \ldots, C_n$, with $C_i = A_i \oplus B_i$. Suppose Θ is not soft. For each i choose a child $\oplus$-link L_i of C_i. For all i we have $w(L_i) < 1$, so for each i we can choose a &-link $\&_i$ on which L_i depends. Without loss of generality, assume $\&_i$ is above $L_{i+1}[\mathrm{mod}\ n]$. Thus, by the (AM) condition (equivalent to the technical condition (T), given the dependency (monomial) condition), and the fact that $w(\text{parent}) \leq w(\text{child})$, we have $w(L_1) < w(\&_1) \leq w(L_2) < w(\&_2) \leq \ldots \leq w(L_n) < w(\&_n) \leq w(L_1)$, a contradiction (since $w(L_1)$ is at either end). *QED.* See [Hamano 2004] for more on softness of Girard's monomial proof structures.

with the same conclusion(s)[31], and summing the weights of links and formulas so identified. Clearly this collapse does not preserve the dependency condition. A straightforward case distinction shows that if a link in a non-monomial Girard proof structure Θ is terminal, respectively removable, in the sense of Definition 4 in Appendix A (page 67), it retains this property upon collapsing Θ. Moreover, the removal of such links commutes with the collapse. Therefore, collapsing a proof structure preserves the property of having a given sequentialisation.

Compact non-monomial Girard proof structures are in bijection with our proof structures: as explained in Section 6.3, Girard's condition that weights must be non-zero corresponds to our cut condition (P0), and Girard's technical condition (T) corresponds to our resolution condition (P1). The bijection between compact non-monomial Girard proof structures and our proof structures can be further refined: it is not too hard to see that compact non-monomial Girard proof nets are in bijection with sets of linkings in our sense which satisfy (P0), (P1) and (P2*), where (P2*) is the strengthening of (P2) defined in Section 4.7.4 (page 29). (To obtain the correspondence, note that it suffices to jump to axiom links in Girard's setting.)

A straightforward induction on the size of proof structures shows that a compact non-monomial Girard proof structure has a proof Π as a sequentialisation (in the sense of Appendix A.2) if and only if its counterpart as one of our proof structures has Π as a sequentialisation (in the sense of Section 5.3). Thus sequentialisable compact non-monomial Girard proof structures are in bijection with our proof nets.

6.4.1 *Surjection from cut-free monomial proof nets to our own.* The map f from (monomial) Girard proof structures to our proof structures obtained by composing the collapse (to compact non-monomial Girard proof structures) and the bijection (between compact non-monomial Girard proof structures and our proof structures) preserves the property of having a given sequentialisation (since both the collapse and the bijection do), and hence of being a proof net. Thus the restriction of f to cut-free (monomial) Girard proof nets is a surjection onto our cut-free proof nets, and the diagram of binary relations in Figure 23 commutes. On proof nets, the map f is surjective since it is the composite of Girard's surjective sequentialisation relation and the surjective translation function from cut-free MALL proofs to our cut-free proof nets (Section 4.2, page 12).

Although surjective from cut-free (monomial) proof nets, the function f is not surjective from cut-free (monomial) proof structures in general. The Gustave example G on page 24 cannot be the f-image of a monomial proof structure[32]: G is not soft (in the sense that the proof structure inhabits both arguments of each outermost $\oplus$), while all monomial proof structures are soft (footnote 30), and softness is preserved by f (since it is preserved by both the collapse and the bijection).

---

[31] When two &-links x and y with the same conclusion collapse, their eigenweights p_x and p_y are identified. This is done by syntactic substitution of p_y for p_x (or vice versa) in any weight w occurring in the proof structure, which is similar to replacing w by $w \wedge (p_x \Leftrightarrow p_y)$. It amounts to ignoring all slices in the uncollapsed proof structure for which x and y choose opposite arguments. This is unproblematic, since by (PRESENCE) on page 61, given that no slice can contain both x and y, any ignored slice is identical to one that is not ignored.

[32] We are very grateful to Masahiro Hamano for pointing this out.

$$\begin{tikzcd}[column sep=large]
\substack{\text{Cut-free}\\ \text{monomial}\\ \text{proof nets}} \arrow[rr, dash, "\substack{\textit{Girard's}\\ \textit{sequentialisation}}", "\substack{\textit{relation}\\ \text{(surjective)}}"'] \arrow[dr, two heads, "\text{(surjection) } f"'] & \arrow[d, phantom, "="] & \substack{\text{Cut-free}\\ \text{MALL}\\ \text{proofs}} \arrow[dl, two heads, "\substack{\textit{translation function}\\ \text{(surjection)}}"] \\
& \substack{\text{Our}\\ \text{cut-free}\\ \text{proof nets}} &
\end{tikzcd}$$

Fig. 23. Relationship between Girard's monomial proof nets and our own proof nets. This triangle of binary relations commutes in both directions, i.e., both from the top-left vertex and from the top-right vertex to the bottom vertex. Note that Girard's sequentialisation relation is surjective in both directions: every cut-free MALL proof and every cut-free monomial proof net is in the relation.

Appendices

A Girard's definition of monomial proof net (verbatim reproduction)

The following definition of monomial proof net for MALL is copied essentially verbatim from Girard's paper [Girard 1996].

A.1 Proof-structures

Definition 1. A ***link*** L is an expression

$$\frac{P_1, \ldots, P_n}{Q_1, \ldots, Q_m}\, L$$

involving n formulas (the ***premises*** of L) $P_1, \ldots, P_n$ and m formulas (the ***conclusions*** of L) $Q_1, \ldots, Q_m$

ID—links	: 0 premises	2 conclusions	: $A, A^\perp$
CUT—links	: 2 premises : $A, A^\perp$	0 conclusions	
$\otimes$—links	: 2 premises : A, B	1 conclusion	: $A \otimes B$
$\parr$—links	: 2 premises : A, B	1 conclusion	: $A \parr B$
$\oplus_1$—links	: 1 premise : A	1 conclusion	: $A \oplus B$
$\oplus_2$—links	: 1 premise : B	1 conclusion	: $A \oplus B$
&—links	: 2 premises : A, B	1 conclusion	: $A \& B$

The premises of $\otimes$, $\parr$, &-links are ordered: this means that we can distinguish a left premise (here A) and a right premise (here B). On the other hand the premises of a *CUT*-link and the conclusions of an *ID*-link are unordered.

Remark. — It is convenient to consider generalized axioms $\vdash A_1, \ldots, A_n$ $(n > 0)$, which are interpreted by generalized axiom links (no premise, but ordered conclusions $A_1, \ldots, A_n$). Such generalized axioms will occur in the proof of our main theorem 2; they also occur when one wants to accommodate other styles of syntax, which are foreign to the proof-net technology, in which case they are called ***boxes***.

The idea of a box is that from the outside it looks like a generalized axiom, whereas it has an inside which can be in turn another proof-net. A box freezes n formulas, and can therefore be seen as a sequent, the conclusion of a rule, whose premises are proven in the box. Traditional sequent calculus is therefore a system of proof-nets in which the only links are boxes, and all the improvement made in 9 years consist in progressively restricting the use of boxes: in this paper boxes are limited to the exponential connectives (and to the neutral $\top$).

Remark. — One should never speak of formulas, but of *occurrences*, which is extremely awkward. We adopt once and for all the convention that all our formulas are distinct (for instance by adding extra indices). In particular ID, $\otimes$, $\invamp$, $\oplus$ and &-links are determined by their conclusion(s), and a CUT-link is determined by its premises.

Definition 2.

- If L is a &-link, with $A\&B$ its conclusion, we introduce the ***eigenweight*** p_L, which is a boolean variable. The intuitive meaning of p_L is the choice l/r between the two premises A and B of the link, p_L for "left", i.e. A, $\neg p_L$ for "right", i.e. B; we use ϵp_L to speak of p_L or $\neg p_L$.
- If Θ is a structure involving the &-links $L_1, \ldots, L_k$ (with associated eigenweights $p_1, \ldots, p_k$), then a ***weight*** (relative to Θ) is any element of the boolean algebra generated by $p_1, \ldots, p_k$.

Definition 3. A ***proof-structure*** Θ consists of:

- A set of formulas (see the previous remark);
- A set of links; each of these links takes its premise(s) and conclusion(s) among the formulas of Θ;
- For each formula A of Θ, a ***weight*** $w(A)$, i.e. a non-zero element of the boolean algebra generated by the eigenweights $p_1, \ldots, p_n$ of the &-links of Θ;
- For each link L of Θ, a ***weight*** $w(L)$.

satisfying the following conditions:

- Each formula is the premise of *at most* one link and the conclusion of at least one link; the formulas which are not the premises of some link are called the ***conclusions*** of Θ;
- $w(A) = \Sigma w(L)$, the sum being taken over the set of links with conclusion A;
- if A is a conclusion of Θ, then $w(A) = 1$;
- if w is any element of the boolean algebra generated by the weights occurring in Θ, and L is a &-link, then $w.\neg w(L)$ does not depend on p_L, i.e. belongs to the boolean algebra generated by the eigenweights distinct from p_L;
- if w is any weight occurring in Θ, then w is a ***monomial*** $\epsilon_1 p_{L_1} \ldots \epsilon_k p_{L_k}$ of eigenweights and negations of eigenweights[33];
- $w(L) \neq 0$; moreover if L is any non-identity link, with premises A and (or) B then

[33]This is the ***dependency*** condition.

– if L is any of $\otimes$, $\parr$, CUT, then $w(L) = w(A) = w(B)$;
– if L is $\oplus_1$, then $w(L) = w(A)$;
– if L is $\oplus_2$, then $w(L) = w(B)$;
– if L is a &-link, then $w(A) = w(L).p_L$ and $w(B) = w(L).\neg p_L$ (hence $w(L) = w(A) + w(B)$).

Remark. —

▶ Weights are in a boolean algebra, and therefore both algebraic and logical graphism can be used; here we decide to use the product notation (instead of the intersection), but we keep $\neg w$ (instead of $1 - w$); when we use the sum, we of course mean the disjoint union, i.e. when I write $w(L) = w(A) + w(B)$, I implicitly mean that $w(A).w(B) = 0$.

▶ The technical condition "$w.\neg w(L)$ does not depend on p_L" says that the boolean variable p_L has no real meaning "outside $w(L)$"; applying the condition to $\neg w(L)$, we see that $w(L)$ does not depend on p_L, in particular $w(L).\epsilon p_L \neq 0$.

▶ There are two ways to think of the dependency condition: either as a technical restriction needed for the sequentialisation theorem (all our efforts to get rid of it failed) or as a nice companion to the previous condition, since both are very natural when a proof-structure is seen as a coherent space, see A.1.1. [of [Girard 1996]].

A.2 Sequent calculus and proof-nets

Definition 4. Let Θ be a proof-structure and let L be either a CUT-link, or a link with only one conclusion, which is in turn a conclusion of Θ and such that $w(L) = 1$; we say that L is a ***terminal*** link of Θ. Given such a link, we define the removal of L in Θ which consists (provided it makes sense) in one or two proof-structures.

▶ If L is a $\otimes$-link (resp. a CUT-link) with premises A, B, and $\Gamma, A \otimes B$ (resp. Γ) is the set of conclusions of Θ: the removal of L consists in partitioning (if possible) the formulas of Θ distinct from $A \otimes B$ (resp. the formulas of Θ) into two subsets X and Y, one containing A, the other containing B, in such a way that, whenever a link L' distinct from L has a premise or a conclusion in X (resp. in Y), then all other premises and conclusions of L' belong to X (resp. to Y). The restrictions Θ/X and Θ/Y are defined in an obvious way, and are proof-structures with respective conclusions Γ', A and Γ'', B. Observe that $\Gamma = \Gamma', \Gamma''$.

▶ If L is a $\parr$-link with premises A, B, and $\Gamma, A \parr B$ is the set of conclusions of Θ: the removal of L consists in removing the conclusion $A \parr B$ and the link L; this induces a proof-structure with conclusions Γ, A, B.

▶ If L is a $\oplus_1$-link with premises A, B, and $\Gamma, A \oplus B$ is the set of conclusions of Θ: the removal of L consists in removing the conclusion $A \oplus B$ and the link L; this induces a proof-structure with conclusions Γ, A.

▶ If L is a $\oplus_2$-link with premises A, B, and $\Gamma, A \oplus B$ is the set of conclusions of Θ: the removal of L consists in removing the conclusion $A \oplus B$ and the link L; this induces a proof-structure with conclusions Γ, B.

- If L is a &-link with premises A, B, and $\Gamma, A\&B$ is the set of conclusions of Θ: the removal of L consists in first removing the conclusion $A\&B$ and the link L (to get Θ') and then forming two proof-structures Θ_A and Θ_B:
  - In Θ' make the replacement $p_L = 1$, and keep only those links L' whose weight is still non-zero, together with the premises and conclusions of such links: the result is by definition Θ_A, a proof-structure with conclusions Γ, A.
  - In Θ' make the replacement $p_L = 0$, and keep only those links L' whose weight is still non-zero, together with the premises and conclusions of such links: the result is by definition Θ_B, a proof-structure with conclusions Γ, B.

Definition 5. A proof-structure Θ is ***sequentialisable*** when it can be reduced, by iterated removal of terminal rules, to identity links. In more pedantic terms:

- An identity link is sequentialisable;
- If the result of removing the terminal link L in Θ yields sequentialisable proof-structures, then Θ is sequentialisable.

Remark. —

- The removal of a given terminal link is not always possible, and its result is not necessarily unique (however, for *proof-nets*, it would be easy to show, by a connectivity argument, that the removal of a $\otimes$- or CUT-link is unique).
- Each removal step consists in the writing down of a rule of MALL; therefore a sequentialisable proof-structure has a ***sequentialisation***, which consists in a proof in MALL.

[Rest of subsection not reproduced.]

A.3 A wrong answer: slicing

Definition 6. Let φ be a ***valuation*** for Θ, i.e. a function from the set of eigenweights of Θ into the boolean algebra $\{0, 1\}$, which induces a function (still denoted φ) from the weights of Θ to $\{0, 1\}$. The ***slice*** $\varphi(\Theta)$ is obtained by restricting to those formulas A of Θ such that $\varphi(w(A)) = 1$, with an obvious modification for the remaining &-links: only one premise is present.

[Rest of subsection not reproduced.]

A.4 Proof-nets

Our basic idea will be to mimic our criterion of [Girard 1990]; in this paper, certain switchings for $\forall$-links were induced by the dependency of some formula upon the *eigenvariable* of the link.

Definition 7. Let φ be a valuation of Θ, let p_L be an eigenweight; we say that the weight w (in Θ) ***depends on*** p_L (in $\varphi(\Theta)$) iff $\varphi(w) \neq \varphi_L(w)$, where the valuation φ_L is defined by:

- $\varphi_L(p_L) = \neg(\varphi(p_L))$
- $\varphi_L(p_{L'}) = \varphi(p_{L'})$ if $L' \neq L$.

A formula A of Θ is said to ***depend*** on p_L (in $\varphi(\Theta)$), if A is the conclusion of a link L' such that $\varphi(w(L')) = 1$ and $\varphi_L(w(L')) = 0$. This basically means that A and

L' are present in $\varphi(\Theta)$, but that changing the value of the valuation for p_L would make A (or at least L') disappear from the slice.

Definition 8. A ***switching*** $\mathcal{S}$ of a proof-structure Θ consists in:

- The choice of a valuation $\varphi_{\mathcal{S}}$ for Θ;
- The selection of a choice $\mathcal{S}(L) \in l, r$ for all $⅋$-links of $\varphi_{\mathcal{S}}(\Theta)$;
- The selection for each &-link L of $\varphi_{\mathcal{S}}(\Theta)$ of a formula $\mathcal{S}(L)$, the ***jump*** of L, depending on p_L in $\varphi_{\mathcal{S}}(\Theta)$. There is always a ***normal*** choice of jump for L, namely the premise A of L such that $\varphi_{\mathcal{S}}(w(A)) = 1$.

Definition 9. Let $\mathcal{S}$ be a switching of a proof-structure Θ; we define the graph $\Theta_{\mathcal{S}}$ as follows:

- The vertices of $\Theta_{\mathcal{S}}$ are the formulas of $\varphi_{\mathcal{S}}(\Theta)$;
- For all *ID*-links of $\varphi_{\mathcal{S}}(\Theta)$, we draw an edge between the conclusions;
- For all generalized axiom links with conclusions $A_1, \dots, A_n$, we draw an edge between A_1 and A_2, etc., A_{n-1} and A_n;
- For all *CUT*-links of $\varphi_{\mathcal{S}}(\Theta)$, we draw an edge between the premises;
- For all $\oplus$-links of $\varphi_{\mathcal{S}}(\Theta)$, we draw an edge between the conclusion and the premise;
- For all $\otimes$-links of $\varphi_{\mathcal{S}}(\Theta)$, we draw an edge between the left premise and the conclusion, and between the right premise and the conclusion;
- For all $⅋$-links L of $\varphi_{\mathcal{S}}(\Theta)$, we draw an edge between the premise (left or right) selected by $\mathcal{S}(L)$ and the conclusion;
- For all &-links L of $\varphi_{\mathcal{S}}(\Theta)$, we draw an edge between the jump $\mathcal{S}(L)$ of L and the conclusion.

Definition 10. A proof-structure Θ is said to be a ***proof-net*** when for all switchings $\mathcal{S}$, the graph $\Theta_{\mathcal{S}}$ is connected and acyclic.

B Reformulations of the MLL condition

This appendix proves Proposition 4.20 (page 28).

Proof. To streamline the proof, without loss of generality we may consider λ in the statement of the proposition as a linking on an MLL sequent (after collapsing the unary additives).

Define an ***abstract switching*** as the graph $\mathcal{G}_\lambda$ of a linking λ on a sequent in the variant of MLL in which $\otimes$ is a binary connective and $⅋$ is unary. An abstract switching s is ***balanced*** if $|\mathsf{ax}| = |\otimes| + 1$, where $|\mathsf{ax}|$ and $|\otimes|$ are the numbers of links and tensors in s, respectively. Clearly, any $⅋$-switching of a linking on an MLL sequent Γ induces an abstract switching, namely by abstracting from the underlying structure of Γ, regarding any subformula reached by a deleted argument edge of a $⅋$-vertex as a separate formula in the sequent. Moreover, a linking on an MLL sequent is balanced iff any of the induced abstract switchings is balanced. We now show the following:

(i) If an abstract switching is acyclic and connected, it is balanced.
(ii) If an abstract switching is acyclic and balanced, it is connected.

(iii) If an abstract switching is balanced and connected, it is acyclic.

From this we obtain:

(A)$\wedge$(c) $\Rightarrow$ (B)	by (i);
(a)$\wedge$(C) $\Rightarrow$ (B)	by (i);
(A)$\wedge$(B) $\Rightarrow$ (C)	by (ii);
(C)$\wedge$(B) $\Rightarrow$ (A)	by (iii);
(A)$\Rightarrow$(a)	because every linking has at least one $⅋$-switching;
(C)$\Rightarrow$(c)	because every linking has at least one $⅋$-switching.

These six implications yield the statement of the proposition.

Proof of (i), by structural induction on the abstract switching s:

Induction base: If s has no connectives, it must be the disjoint union of n pairs of complementary literals, each connected by a link. Since sequents are non-empty and s is connected, $n = 1$. Hence s is balanced.

Induction step: Suppose s has an outermost $⅋$-vertex p. Deleting p yields an abstract switching s' which is a tree, hence balanced by induction. Thus s is balanced.

Suppose s has an outermost $\otimes$-vertex t. Deleting t yields an abstract switching s' comprising two disjoint abstract switchings s_0 and s_1, both trees. By induction s_0 and s_1 are balanced, so s' satisfies $|\mathsf{ax}| = |\otimes| + 2$, and s is balanced.

Proof of (ii), by structural induction on s:

Induction base: If s has no connectives, $|\otimes| = 0$, so $|\mathsf{ax}| = 1$, and s is connected.

Induction step: Suppose s has an outermost $⅋$-vertex p. Deleting p yields an abstract switching s' which is still acyclic and balanced, and hence connected by induction. Thus s is connected.

Suppose s has an outermost $\otimes$-vertex t with argument vertices t_0 and t_1. Deleting t yields an abstract switching s' satisfying $|\mathsf{ax}| = |\otimes| + 2$. Since s is acyclic, in s' the vertices t_0 and t_1 are not connected. Let s_0 be the part of s' connected to t_0 and let s_1 be the remainder of s'. Because s_0 is acyclic and connected, it must be balanced by (i). Thus s_1 must be balanced, since any surplus in links in s_1 would be compensated by a shortage in s_0 and vice versa. Hence, by induction, s_1 is connected. Thus s is connected.

Proof of (iii): We prove the stronger statement that if an abstract switching s is connected and satisfies $|\mathsf{ax}| \geq |\otimes| + 1$, it is acyclic. We proceed by structural induction on s.

Induction base: If s has no connectives, $|\otimes| = 0$, $|\mathsf{ax}| = 1$, and s is acyclic.

Induction step: Suppose s has an outermost $⅋$-vertex p. Deleting p yields an abstract switching s' which is connected and satisfies $|\mathsf{ax}| \geq |\otimes| + 1$. Hence it is acyclic by induction. Thus s is acyclic.

Suppose s has an outermost $\otimes$-vertex t with argument vertices t_0 and t_1. Deleting t yields an abstract switching s' satisfying $|\mathsf{ax}| \geq |\otimes| + 2$. In s' the vertices t_0 and t_1 are not connected, otherwise s' would be connected, therefore acyclic by induction, hence balanced by (i), a contradiction. Let s_0 be the part of s' that is connected to t_0 and s_1 the remainder, connected to t_1. As least one of s_0 and s_1 must satisfy $|\mathsf{ax}| \geq |\otimes| + 1$. Hence it is acyclic by induction, and thus balanced by (i). This implies that the other also satisfies $|\mathsf{ax}| \geq |\otimes| + 1$, and thus is acyclic by induction. Hence s is acyclic. $\square$

C Illegal unions of switching cycles

This appendix proves that condition ($P3^l$), defined in Section 4.7.3, is equivalent to the toggling condition (P3). Employing Proposition 4.20 (page 28), condition (P2) in our definition of a MALL proof net θ can be partitioned into

(P2a) every $\parr$-switching of every linking $\lambda \in \theta$ is acyclic, and

(P2b) every linking $\lambda \in \theta$ is balanced.

As remarked in Section 4.7.2, the former is equivalent to

(P2a′) for no linking $\lambda \in \theta$ does $\mathcal{G}_\lambda$ contain a switching cycle.

Therefore, conditions (P2a) and (P3) can be combined into

(P2a3) every set $\Lambda \subseteq \theta$ with a switching cycle in $\mathcal{G}_\Lambda$ toggles a & that is not in any switching cycle of $\mathcal{G}_\Lambda$.

This condition is in turn equivalent to:

(P2a3′) for any $\Lambda \subseteq \theta$ and any non-empty union of switching cycles S of $\mathcal{G}_\Lambda$, Λ toggles a & that is not in S.

Now (P2a3′) can be reformulated as:

($P3^l$) $\mathcal{G}_\theta$ contains no illegal union of switching cycles.

Acknowledgements

Many thanks to Vaughan Pratt for invaluable feedback and support during the long development of this work, and to Stanford University for its financial and institutional support under account 1DMA644. Thanks to Paul-André Melliès for suggesting to the first author in 1999 that the search for a satisfactory notion of MALL proof net was an interesting and potentially fruitful research topic. We are extremely grateful to Pierre-Louis Curien, Masahiro Hamano and Robin Houston for their insightful and detailed technical comments. Thanks to Nil Demirçubuk and Selena Clancy for their support and understanding.

References

Abramsky, S. and Jagadeesan, R. 1994. Games and full completeness for multiplicative linear logic. *Journal of Symbolic Logic 59,* 2, 543–574.

Abramsky, S. and Melliès, P.-A. 1999. Concurrent games and full completeness. In Proc. 14th Annual IEEE Symp. on *Logic in Computer Science*, Trento, Italy, July 1999. IEEE Computer Society Press, 431–442.

Barr, M. 1979. *∗-Autonomous categories.* Lecture Notes in Mathematics, vol. 752. Springer-Verlag.

Bellin, G. and Ketonen, J. 1992. A decision procedure revisited: Notes on direct logic, linear logic and its implementation. *Theoretical Computer Science 95*, 115–142.

Blute, R. F., Cockett, J. R. B., Seely, R. A. G. and Trimble, T. H. 1996. Natural deduction and coherence for weakly distributive categories. *Journal of Pure and Applied Algebra 113*, 229–296.

Blute, R., Hamano, M. and Scott, P. 2005. Softness of hypercoherences and MALL full completeness. *Annals of Pure and Applied Logic 131*, 1–63.

Blute, R. F. and Scott, P. J. 1996. Linear Läuchli semantics. *Annals of Pure and Applied Logic 77*, 101–142.

DANOS, V. AND REGNIER, L. 1989. The structure of multiplicatives. *Archive for Mathematical Logic 28*, 181–203.

DEVARAJAN, H., HUGHES, D. J. D., PLOTKIN, G. D. P. AND PRATT, V. R. 1999. Full completeness of the multiplicative linear logic of Chu spaces. In Proc. 14th Annual IEEE Symp. on *Logic in Computer Science*, Trento, Italy, July 1999. IEEE Computer Society Press, 234–245.

DOŠEN, K. AND PETRIĆ, Z. 2004. *Proof-Theoretical Coherence.* Preprint, Mathematical Institute, Belgrade, 2004.

GIRARD, J.-Y. 1987. Linear logic. *Theoretical Computer Science 50*, 1–102.

GIRARD, J.-Y. 1990. Quantifiers in linear logic II. In Proc. *Nuovi problemi della logica e della filosofia della scienze*, Viareggio, Italy, 1990. Clueb, Bologna.

GIRARD, J.-Y. 1996. Proof-nets: the parallel syntax for proof theory. In *Logic and Algebra.* Lecture Notes In Pure and Applied Mathematics, vol. 180. Marcel Dekker, New York.

GIRARD, J.-Y. 1999. On the meaning of logical rules I: syntax vs. semantics. In *Computational Logic*, U. Berger and H. Schwichtenberg, eds., NATO ASI Series 165, vol. 14. Springer-Verlag, 215–272.

HAMANO, M. 2004. Softness of MALL proof-structures and a correctness criterion with Mix. *Archive for Mathematical Logic 43:6*, 751-794.

HUGHES, D. J. D. 2002. A canonical graphical syntax for non-empty finite products and sums. Technical report, `http://boole.stanford.edu/~dominic/papers` .

HUGHES, D. J. D. 2004. Proofs Without Syntax. Submitted for publication, August 2004. Archived: `http://arxiv.org/abs/math/0408282` .

HUGHES, D. J. D. 2005. Logic Without Syntax. Submitted to a conference, January 2005. `http://boole.stanford.edu/~dominic/papers` .

HYLAND, J. M. E. AND ONG, C.-H. L. 1993. Fair games and full completeness for multiplicative linear logic without the MIX-rule. On Ong's web page.

JOYAL, A. 1995. Free bicomplete categories. *Math. Reports XVII*, 219–225.

KETONEN, J. AND WEYHRAUCH, R. 1984. A decidable fragment of predicate calculus. *Theoretical Computer Science 32*, 297–307.

LAMARCHE, F. AND STRASSBURGER, L. 2005. Naming Proofs in Classical Propositional Logic. Accepted for *Typed Lambda Calculus and Applications '05.*

LAURENT, O. AND TORTURA DE FALCO, L. 2004. Slicing polarized additive normalization. In *Linear Logic in Computer Science*, T. Ehrhard, J-Y. Girard, P. Ruet and P. Scott, eds., London Mathematical Society Lecture Note Series 316, Cambridge University Press.

LOADER, R. 1994. Linear logic, totality and full completeness. In Proc. 9th Annual IEEE Symp. on *Logic in Computer Science*, Paris, July 1994. IEEE Computer Society Press, 292–298.

TAN, A. 1997. Full completeness for models of linear logic. Ph.D. thesis, King's College, University of Cambridge.