In Proceedings of IEEE Logic in Computer Science, 2003

Proof Nets for Unit-free Multiplicative-Additive Linear Logic
(Extended abstract)*

DOMINIC HUGHES
Stanford University
dominic@cs.stanford.edu

ROB VAN GLABBEEK
Stanford University
rvg@cs.stanford.edu

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 commutations of rules. 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.

1 Introduction

The beautiful theory of proof nets for unit-free multiplicative linear logic (MLL) appeared alongside the introduction of linear logic [Gir87]. A proof net is an abstract representation of a proof: the translation of cut-free proofs into proof nets identifies proofs modulo inessential commutations of rules. The identifications have since been verified as canonical from a semantic perspective, with numerous full completeness results for MLL, *e.g.* [AJ94, HO93, Loa94, Tan97, BS96, DHPP99]. Furthermore, the identifications correspond to coherences of free star-autonomous categories [BCST96].

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 [Gir87]. Progress towards a solution was made by Girard [Gir96] with a notion of MALL proof net based on monomial weights. Unfortunately, monomial proof nets fail 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 translation of cut-free proofs into monomial proof nets. Indeed, to quote Girard, monomial proof nets are "far from being absolutely satisfactory" [Gir96]. We illustrate the problems in detail in Section 4.1.

In this paper we propose a new notion of MALL proof net (Section 2) which adheres faithfully to the original MLL theory: we provide a simple inductive translation of cut-free proofs into cut-free proof nets, yielding the sought-after abstract representations of cut-free proofs modulo inessential commutations 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 (Definition 1), and prove that a set of linkings is the translation of a proof if and only it is a proof net (Theorem 1).

In Section 3 we extend our proof nets to include the cut rule, and present a notion of cut elimination. Our approach to cut suffers from the same problem as 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 inessential rule commutations is crucial only in the cut-free setting. Moreover, our notion of cut elimination is simply defined, strongly normalising, and yields a category of proof nets in which $\&$ and $\oplus$ are product and coproduct.

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, hopefully 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 mentioned earlier (and detailed in Section 4.1), any MALL full completeness result[1] based on them (*e.g.* [AM99], and the work in progress of Blute, Hamano and Scott on hypercoherence spaces) suffers accordingly, particularly with regard to faithfulness. We anticipate that our new definition of MALL proof net will yield cleaner and more accessible MALL full completeness results.

Relationship with Girard's monomial proof nets. The technical starting point for our definition of proof net was Girard's definition of monomial proof net [Gir96], and indeed we employ variants of Girard's ingenious notions of slice and jump. Each of our proof nets translates natu-

*A preliminary version of some of the material in this paper was presented in a talk at the workshop *Linear Logic 2002*, Copenhagen.

[1] The original motivation for this work came as part of a project by the first author, Gordon Plotkin and Vaughan Pratt aiming to extend the full completeness of Chu spaces for MLL [DHPP99] to MALL. We have since discovered that the result does not extend.

rally into a non-monomial Girard proof net, *i.e.*, a Girard proof net without the condition demanding that weights must be monomials. Thus one of our contributions relative to [Gir96] is the successful elimination of the monomial condition. In [Gir96] Girard remarks that he had been trying to circumvent this technical limitation since 1990, and Appendix A.1.5 of [Gir96] lists three specific problems that must be solved in any attempt to eliminate the monomial condition, *i.e.*, in any attempt 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.*

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 upon relaxing the monomial condition. We show that this is indeed the case: in Section 4.2 we present a non-monomial proof structure that does not correspond to any proof, yet satisfies Girard's criterion. We address (2) by leaving weights implied, defining a proof net on a sequent Γ as a set of axiom linkings on an extension Γ^+ of Γ with complementary pairs of cut formula occurrences. Point (3) is addressed by the fact that our definition of cut elimination is sufficiently simple that confluence and strong normalisation are immediate.

The proof that our correctness criterion captures proof translations hinges on an ordering of $\&$'s and $⅋$'s which we call *domination*. By introducing domination we avoid the use of empires [Gir87, Gir96], thereby sidestepping the problem of stability of maximal empires ([Gir96], section 1.5.3)—the main technical problem that led Girard to resort to monomials in the first place.

MALL. By MALL we mean multiplicative-additive linear logic without units [Gir87]. 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 by de Morgan duality. For technical convenience we take sequents to be unordered, *i.e.*, a sequent is a non-empty set of formula occurrences $A_1, \ldots, A_n$. We omit turnstiles, which are redundant since all sequents are right-sided. Sequents are proved using the following rules:

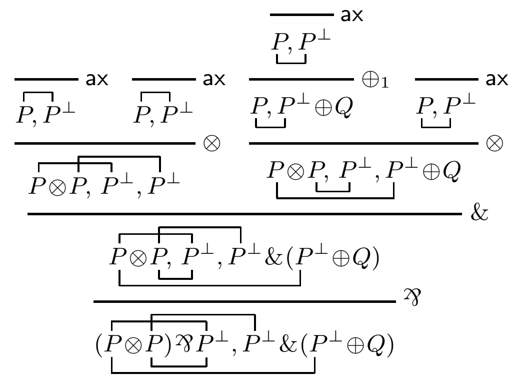

Figure 1. Example of the inductive translation of a cut-free proof into a cut-free proof net.

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

$$\frac{\Gamma, A \quad B, \Delta}{\Gamma, A \otimes B, \Delta}\ \otimes \qquad \frac{\Gamma, A, B}{\Gamma, A ⅋ B}\ ⅋$$

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

Here, and throughout this document, $P, Q, \ldots$ range over propositional variables, $A, B, \ldots$ over formulas, and $\Gamma, \Delta, \ldots$ over sets of formula occurrences. In eliminating the permutation rule, we assume an implicit tracking of formula occurrences above the line of a rule to formula occurrences below the line. Without loss of generality (see [Gir87]) we restrict the axiom rule to literals.

Flavour of our approach. To give a flavour of our approach, Figure 1 shows an example of the inductive translation of a cut-free proof into one of our proof nets. The concluding proof net consists of 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. Two additional translations are shown in Figure 2.

Here is an example of a proof net with four linkings:

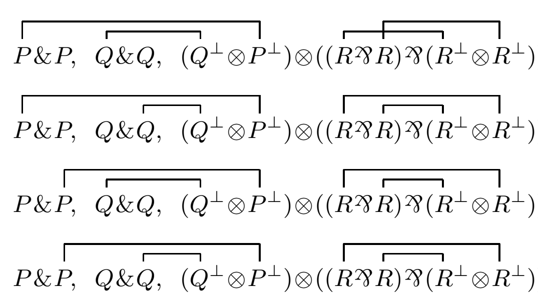

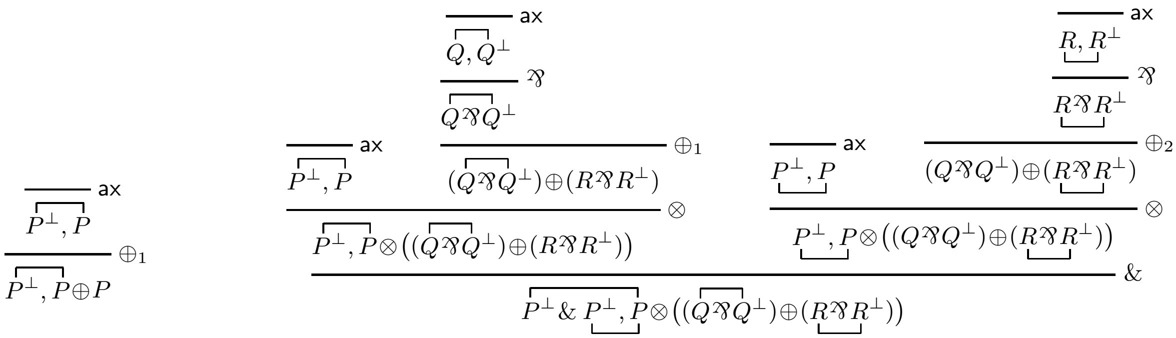

Figure 2. Deriving the proof nets used in the example of cut elimination.

(To aid presentation, we duplicated the underlying sequent.) Our proof nets can be encoded compactly as collections of axiom links labelled with predicates ('weights'). For example, the four-linking proof net above can be represented as follows:

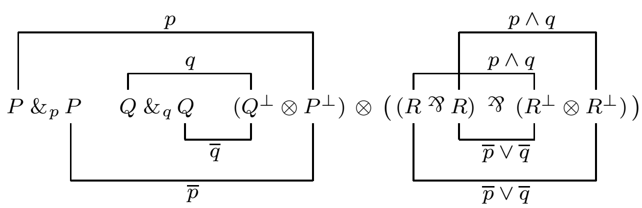

To distinguish the $\&$'s, we have subscripted them. Every $\&$-assignment (assignment of 'left' or 'right' to each of $\&_p$ and $\&_q$) determines a linking as follows: delete each axiom link whose predicate does not hold, where we read p (resp. $\overline{p}$) as "$\&_p$ is assigned 'left' (resp. 'right')" (and q analogously). The reader can check that taking each of the four possible $\&$-assignments in turn produces the four original linkings.

We sketch the idea behind our approach to cut elimination with an example. Consider the proof nets derived in Figure 2. Viewing MALL formulas as objects[2], and a proof net on $A^\perp, B$ as a morphism $A \to B$, the left proof net is a morphism $P \to P \oplus P$ and the right proof net is a morphism $P \oplus P \to P \otimes \big((Q ⅋ Q^\perp) \oplus (R ⅋ R^\perp)\big)$. Composition, yielding a morphism $P \to P \otimes \big((Q ⅋ Q^\perp) \oplus (R ⅋ R^\perp)\big)$, in other words, a proof net on the sequent $P^\perp, P \otimes \big((Q ⅋ Q^\perp) \oplus (R ⅋ R^\perp)\big)$, proceeds as follows. First, we concatenate the two sequents into a combined sequent of four formulas (omitting commas):

$$P^\perp \quad [P \oplus P] \cdots [P^\perp \& P^\perp] \quad P \otimes \big((Q ⅋ Q^\perp) \oplus (R ⅋ R^\perp)\big)$$

The cut formulas are annotated with square brackets, and the dotted line represents the cut. Next, we merge the two original proof nets into a proof net on the combined sequent:

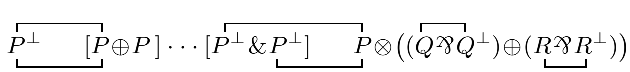

[2]The example should be accessible to readers with no knowledge of category theory: focus on the underlying cut elimination.

This new proof net has two linkings, one drawn above the sequent and one drawn below.[3] Each linking consists of three axiom links. Cut elimination proceeds as follows:

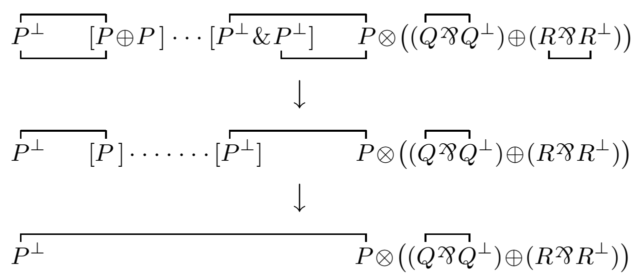

The first step, aside from eliminating the $\oplus$ and $\&$ to leave a cut pair of literals $[P] \cdots [P^\perp]$, retains only one of the two original linkings. The underhanging linking is deleted because it is 'inconsistent': it chooses opposite arguments for the cut $\oplus$ and $\&$ (left and right, respectively). This is an instance of our general rule for additive elimination: *retain precisely the consistent linkings*, those which choose the same argument for the cut $\oplus$ and $\&$ (in the example above, the upper linking, which chooses left for both). The second step is the usual MLL elimination of a cut pair of literals, which we include in order to frame the example in a familiar context. Note that the end result really is a cut-free proof net, *i.e.*, the translation of a cut-free proof: its witness is a subproof of the right-hand proof of Figure 2 (the left branch of the final $\&$-rule).

2 Cut-free MALL proof nets

In this section we introduce our definition of cut-free MALL proof net. As an *aide mémoire* we provide a summary of the definition in the box ahead. We treat cut in Section 3.

An ***additive resolution*** of a MALL sequent Γ is any result of deleting one argument subtree of every additive connective ($\&$ or $\oplus$) of Γ. Thus every remaining $\&$ and $\oplus$ is

[3]The link of the original proof net on $P^\perp, P \oplus P$ appears in both the upper linking and the lower linking of the new proof net; it is duplicated in the merging process, to match the fact that the other proof net had two linkings.

$$\overline{\{\{\{P, P^\perp\}\}\} \;\rhd\; P, P^\perp}\;\mathsf{ax}$$

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

$$\frac{\theta \;\rhd\; \Gamma, A, B}{\theta \;\rhd\; \Gamma, A ⅋ B}⅋ \qquad \frac{\theta \;\rhd\; \Gamma, A \quad \theta' \;\rhd\; \Gamma, B}{\theta \cup \theta' \;\rhd\; \Gamma, A \& B}\&$$

$$\frac{\theta \;\rhd\; \Gamma, A}{\theta \;\rhd\; \Gamma, A \oplus B}\oplus_1 \qquad \frac{\theta \;\rhd\; \Gamma, B}{\theta \;\rhd\; \Gamma, A \oplus B}\oplus_2$$

Figure 3. The inductive translation of cut-free MALL proofs into sets of linkings.

unary. For example, two of 12 possible additive resolutions of the sequent

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

are

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

Let Γ^* be an additive resolution of Γ. An ***axiom link*** on Γ^* is a pair of complementary literal occurrences of Γ^*. A ***linking*** on Γ^* is a partitioning of the set of literal occurrences of Γ^* into axiom links, *i.e.*, a set of disjoint axiom links whose union contains every literal occurrence of Γ^*. For example, there are two linkings possible on the first of the two additive resolutions depicted above:

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

Every additive resolution Γ^* of Γ induces an MLL sequent, namely by collapsing its additive connectives, which are unary in Γ^*. A linking λ on Γ^*, viewed as being on the induced MLL sequent, is exactly an MLL proof structure in the standard sense [Gir87], which we call the ***MLL proof structure induced by*** λ. For example, the MLL proof structure induced by the first of the two linkings above is:

$$P^\perp,\; P \otimes R,\; (R^\perp \otimes R) ⅋ R^\perp$$

A ***linking on*** a MALL sequent Γ is a linking on an additive resolution of Γ. Write $\Gamma \restriction \lambda$ for the additive resolution associated with a linking λ. Every cut-free MALL proof of Γ defines a set of linkings on Γ by a simple induction, as in Figure 3, where $\theta \;\rhd\; \Gamma$ is the judgement "θ is a set of linkings on Γ". (We use the implicit tracking of literal occurrences downwards through rules.) The base case ax is a singleton set of linkings whose only linking comprises a single axiom link, between P and $P^\perp$. Examples of the inductive translation of cut-free proofs into sets of linkings were presented in Figures 1 and 2. Note that if a cut-free proof Π' can be obtained from Π by a series of rule commutations, then Π and Π' translate to the same set of linkings.

Geometric characterisation of proof translations. 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*. Analogous to [Gir96], as a stepping stone to the definition of a proof net, we introduce the less restrictive notion of a *proof structure*.

A $\&$***-resolution*** $\Gamma^\star$ of a sequent Γ is any result of deleting one argument subtree of every $\&$ of Γ. A linking λ on Γ is ***on*** $\Gamma^\star$ if every literal occurrence of λ is in $\Gamma^\star$. A set of linkings θ on Γ is a ***proof structure*** on Γ if it satisfies

(P1) *For every $\&$-resolution $\Gamma^\star$ of Γ, exactly one linking of θ is on $\Gamma^\star$.*[4]

We invite the reader to verify (P1) for the sets of linkings in Figures 1 and 2. Any proof structure can be represented compactly as a set of axiom links labelled with predicates ('weights'), using the encoding described on the third page of the Introduction. In Section 4 we relate our proof structures to those of Girard.

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

(P2) *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 is an MLL proof net, in the usual sense [Gir87, DR89]. To be self-contained, we characterise (P2) explicitly below.

Henceforth view a sequent Γ as a graph: a disjoint union of parse trees, with literals above. For a linking λ on Γ obtain the ***graph*** $\mathcal{G}_\lambda$ ***of*** λ from the additive resolution $\Gamma \restriction \lambda$ (a subgraph of Γ) by adding each axiom link a of λ as a vertex above $\Gamma \restriction \lambda$, with edges from a down to its two literal occurrences. 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. (P2) holds if and only if every switching of every linking of θ is a tree (acyclic and connected).

We require some auxiliary concepts to state our third and last requirement for a set of linkings to be a proof net. A set of linkings Λ ***toggles*** a $\&$-occurrence w of Γ if both arguments of w are present in $\bigcup_{\lambda \in \Lambda} \Gamma \restriction \lambda$. An axiom link a ***depends on*** w ***within*** Λ if, within Λ, a can be made to vanish by toggling w alone: there are $\lambda, \lambda' \in \Lambda$ such that $a \in \lambda$, $a \notin \lambda'$, and w is the only $\&$ toggled by $\{\lambda, \lambda'\}$.

[4] Therefore a proof structure on Γ is a maximal clique in the coherence space of linkings on Γ with incoherence $\lambda \asymp \lambda'$ iff there exists a $\&$-resolution $\Gamma^\star$ of Γ such that both λ and λ' are on $\Gamma^\star$.

Definition of cut-free MALL proof net — ***aide mémoire***

Additive resolution of Γ: any result of deleting one argument subtree of every $\&$ or $\oplus$ of Γ. ($\&$-***resolution*** analogously.)
Axiom link: pair of complementary literal occurrences.
Linking λ on Γ: partitioning of the set of literal occurrences in an additive resolution $\Gamma \restriction \lambda$ of Γ into axiom links.

Graph $\mathcal{G}_\lambda$: $\Gamma \restriction \lambda$ + λ + edges from each axiom link in λ to its two literal occurrences in $\Gamma \restriction \lambda$.
Switching of a linking λ: any subgraph of $\mathcal{G}_\lambda$ obtained by deleting one of the two argument edges of each $⅋$-vertex.

A set of linkings Λ ***toggles*** a $\&$-occurrence w of Γ if both arguments of w are present in $\bigcup_{\lambda\in\Lambda} \Gamma \restriction \lambda$.
An axiom link a ***depends on*** w ***within*** Λ if $\exists \lambda, \lambda' \in \Lambda$ such that $a \in \lambda$, $a \notin \lambda'$, and w is the only $\&$ toggled by $\{\lambda, \lambda'\}$.
Graph $\mathcal{G}_\Lambda$: $\bigcup_{\lambda\in\Lambda} \mathcal{G}_\lambda$ + ***jump*** edges between each axiom link in Λ and any $\&$ on which it depends within Λ.
Switch edge of a $\&$- or $⅋$-vertex x in $\mathcal{G}_\Lambda$: any argument or jump edge of x.
Switching cycle of Λ: a (non self-intersecting) cycle in $\mathcal{G}_\Lambda$ containing at most one switch edge of each $\&$ and $⅋$.

A set of linkings θ is a ***proof net*** if it satisfies
(P1) For every $\&$-resolution $\Gamma^\star$ of Γ, exactly one linking of θ is on $\Gamma^\star$.
(P2) Every switching of every linking of θ is a tree (acyclic and connected).
(P3) Every set Λ of two or more linkings of θ toggles a $\&$ that is not in any switching cycle of Λ.[6]

EXAMPLE 1 Consider the two linkings

$\lambda_1 :$ / $\lambda_2 :$ $\quad P^\perp \oplus (Q \oplus P^\perp),\ (P \& P) \otimes (R \oplus R),\ (R^\perp \otimes R) ⅋ R^\perp$

Here are λ_1 and λ_2 on their respective additive resolutions:

$\lambda_1 :\quad P^\perp \oplus (Q \oplus P^\perp),\ (P \& P) \otimes (R \oplus R),\ (R^\perp \otimes R) ⅋ R^\perp$

$\lambda_2 :\quad P^\perp \oplus (Q \oplus P^\perp),\ (P \& P) \otimes (R \oplus R),\ (R^\perp \otimes R) ⅋ R^\perp$

Let w be the $\&$ of the sequent, and let $\Lambda = \{\lambda_1, \lambda_2\}$. The axiom link between the left-most $R^\perp$ and the left-most R depends on w within Λ: 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 axiom link between the right-most R and $R^\perp$ does not depend on w within Λ, since it is present in both λ_1 and λ_2. It is the only one of the 5 axiom links in Λ (more precisely, in $\bigcup \Lambda$) that does not depend on w within Λ.

We now extend the definition of the graph of a linking to the graph of a set of linkings. Given a set Λ of linkings on Γ, obtain the ***graph*** $\mathcal{G}_\Lambda$ ***of*** Λ from $\bigcup_{\lambda\in\Lambda} \mathcal{G}_\lambda$ by adding, for every $\&$-vertex w and every axiom link a depending on w within Λ, an edge between w and a. Each edge of the latter form, between a $\&$-vertex w and an axiom link, is called a ***jump*** of w. Figure 4 shows $\mathcal{G}_{\{\lambda_1,\lambda_2\}}$ for λ_1 and λ_2 of Example 1, with four jumps (the curved edges). In drawing an axiom link $\sqcap$, we view the horizontal section as a vertex, and the two verticals as edges. We overlap edges from axiom links coming down into the same literal occurrence (*i.e.*, ⅂ means ⅂|). There is no jump to the right-most axiom link, since it does not depend on the $\&$ within $\{\lambda_1, \lambda_2\}$. Note that if $\Lambda \subseteq \Lambda'$, then $\mathcal{G}_\Lambda$ is a subgraph of $\mathcal{G}_{\Lambda'}$, and that for any linking λ, $\mathcal{G}_{\{\lambda\}}$ is precisely $\mathcal{G}_\lambda$ defined on the previous page. ($\mathcal{G}_{\{\lambda\}}$ has no jumps, since no $\&$ is toggled.)

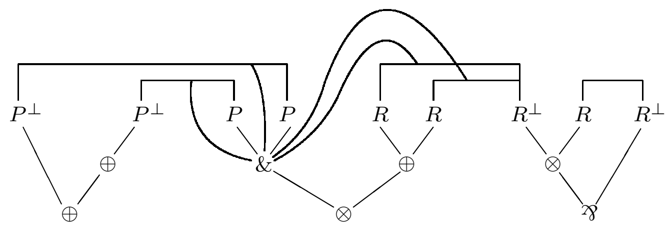

Figure 4. The graph $\mathcal{G}_{\{\lambda_1,\lambda_2\}}$ of Example 1.

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 of x (if x is a $\&$). A ***switching cycle*** of a set of linkings Λ is a cycle in $\mathcal{G}_\Lambda$ containing at most one switch edge of each $\&$ and $⅋$. (We do not permit a cycle to intersect itself.) For example, in $\mathcal{G}_{\{\lambda_1,\lambda_2\}}$ of Figure 4, the cycle "$\& \to \otimes \to \oplus \to$ left-$R \to$ left-$\{R, R^\perp\} \xrightarrow{\text{jump}} \&$" contains only one switch edge of the $\&$, and is therefore a switching cycle of $\{\lambda_1, \lambda_2\}$ of Example 1.

DEFINITION 1 *A set θ of linkings on a MALL sequent Γ is a* ***cut-free proof net*** *if it satisfies (P1), (P2)*[5] *and:*

(P3) *Every set Λ of two or more linkings of θ toggles a $\&$ that is not in any switching cycle of Λ.*[6]

EXAMPLE 2 The set of linkings $\{\lambda_1, \lambda_2\}$ of Example 1 is not a proof net. It fails (P3) since $\mathcal{G}_{\{\lambda_1,\lambda_2\}}$ (Figure 4) contains a switching cycle through the $\&$.

[5] By dropping connectedness from (P2) we obtain a cut-free proof net for MALL with the MIX rule (hypotheses Γ and Δ, conclusion Γ, Δ).

[6] In fact, one need only verify (P3) for those Λ which are ***saturated***, 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 one argument subtree of some of the $\&$'s of Γ.

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

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

$$\frac{\theta \;\rhd\; [\Omega]\; \Gamma, A}{\theta \;\rhd\; [\Omega]\; \Gamma, A \oplus B}\,\oplus_1 \qquad \frac{\theta \;\rhd\; [\Omega]\; \Gamma, B}{\theta \;\rhd\; [\Omega]\; \Gamma, A \oplus B}\,\oplus_2 \qquad \frac{\theta \;\rhd\; [\Omega]\; \Gamma, A, B}{\theta \;\rhd\; [\Omega]\; \Gamma, A ⅋ B}\,⅋$$

Figure 5. Rules for deriving sequentialisable sets of linkings on MALL cut sequents.

EXAMPLE 3 Consider the pair of linkings on the sequent $\Gamma \equiv P^{\perp} \& P^{\perp}, P \oplus P$ obtained as follows:

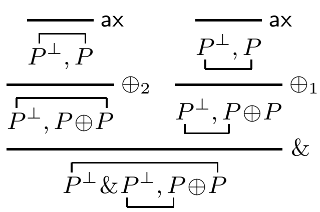

Let λ_1 and λ_2 be the upper- and lower linking respectively (each having just one axiom link). We shall verify that $\{\lambda_1, \lambda_2\}$ is a cut-free proof net. Γ has two $\&$-resolutions, $\Gamma_1^* \equiv P^{\perp} \& \cancel{P^{\perp}}, P \oplus P$ and $\Gamma_2^* \equiv \cancel{P^{\perp}} \& P^{\perp}, P \oplus P$. (P1) holds, since $\{\lambda_1, \lambda_2\}$ contains exactly one linking on $\Gamma_i^{\star}$, namely λ_i. Here are the graphs $\mathcal{G}_{\lambda_1}$, $\mathcal{G}_{\lambda_2}$, and $\mathcal{G}_{\{\lambda_1, \lambda_2\}}$:

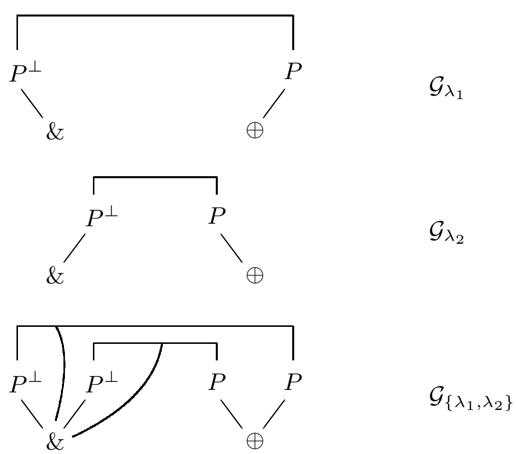

Each λ_i has just one switching, namely $\mathcal{G}_{\lambda_i}$. Since each $\mathcal{G}_{\lambda_i}$ is a tree, (P2) holds. Finally, (P3) holds since $\{\lambda_1, \lambda_2\}$ toggles the $\&$, and the $\&$ is not in any switching cycle of $\{\lambda_1, \lambda_2\}$.

THEOREM 1 *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, once we prove the *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 $\mathcal{G}_\theta$. We prove the Separation Lemma via an ordering on $\&$'s and ⅋'s which we call *domination*[7], a concept reminiscent of the ordering induced by the notion of an *empire* of [Gir96], but different in an essential way. The details are in the full paper. The proof in the case of MIX (see footnote 5) requires only minor changes.

3 Cut

A ***cut*** is a pair $\{A, A^{\perp}\}$ of complementary MALL formulas. We write $A * A^{\perp}$ for $\{A, A^{\perp}\}$, and treat $A * A^{\perp}$ akin to a MALL formula, referring to $*$ as the ***cut connective***. (In the cut elimination example in the Introduction we drew a cut $A * A^{\perp}$ informally as $[A] \cdots [A^{\perp}]$.) A ***cut sequent*** is a non-empty set of occurrences of MALL formulas and cuts. A ***cut-additive resolution*** of a cut sequent Δ is any result of deleting some cuts from Δ and one argument subtree of every remaining additive connective ($\&$ or $\oplus$). Thus every remaining $\&$ and $\oplus$ is unary. For example, here is a cut sequent followed by one of its cut-additive resolutions:

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

$$P \otimes P,\; \cancel{Q * Q^{\perp}},\; P^{\perp} \oplus \cancel{Q},\; (\cancel{R} \oplus S) * (R^{\perp} \& \cancel{S^{\perp}})$$

An ***axiom link*** on a cut-additive resolution Δ^* of a cut sequent Δ is a pair of complementary literal occurrences of Δ^*. A ***linking*** on Δ^* is a partitioning of the literal occurrences of Δ^* into axiom links, *i.e.*, a set of disjoint axiom links on Δ^* whose union contains every literal occurrence of Δ^*. A ***linking on*** Δ is a linking on a cut-additive resolution of Δ. We write $\Delta \restriction \lambda$ for the cut-additive resolution associated with a linking λ.

Write $[\Omega]\,\Gamma$ for the cut sequent obtained by taking the disjoint union of a set Ω of cut occurrences and a MALL sequent Γ. A set of linkings on $[\Omega]\,\Gamma$ is ***sequentialisable*** if it can be derived from the rules in Figure 5, in which $\theta \;\rhd\; [\Omega]\; \Gamma$ is the judgement "θ is a sequentialisable set of linkings on the cut sequent $[\Omega]\,\Gamma$". (We once again use the

[7] Unrelated to domination in flowgraphs.

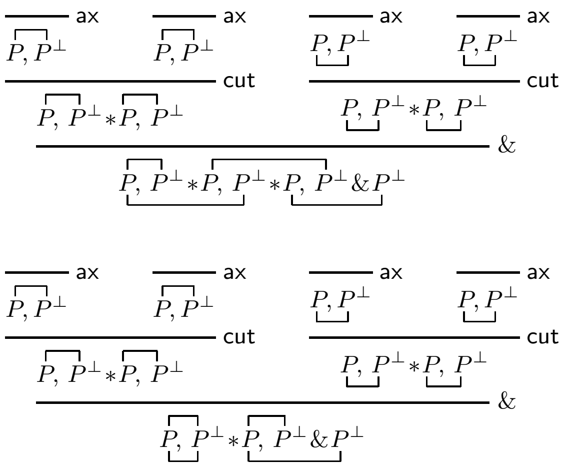

Figure 6. Examples of the translation of a proof with cuts.

implicit tracking of literal occurrences downwards through rules.) The base case ax is a single linking with a single axiom link and no cuts. Figure 6 shows two examples. Each of the conclusions is a set of two linkings, one drawn above the cut sequent and one drawn below. The only difference between the derivations is the final &-rule. The left application keeps the cuts in the hypotheses separate (an instance of the &-rule taking Σ empty and $\Omega = \Omega' = P^\perp * P$), whereas the right application superimposes the two cuts ($\Sigma = P^\perp * P$ and Ω, Ω' empty).

Any derivation of a set of linkings using the rules of Figure 5 projects in an obvious way to a MALL proof, namely, by restricting to the underlying sequents (*viz.*, read Γ for $\theta \rhd [\Omega]\ \Gamma$). For example, the two derivations of Figure 6 each yield the same MALL proof of $P, P^\perp \& P^\perp$.

Write $\Pi \rightsquigarrow \theta$ if Π is the MALL proof obtained from a derivation of a set of linkings θ, and say that Π is a ***sequentialisation*** of θ. If a MALL proof Π' can be obtained from Π by a series of rule commutations in which no &-rules are moved upwards, then Π and Π' are sequentialisations of the same set of linkings. In the cut-free case, $\rightsquigarrow$ is a function from proofs to sets of linkings, exactly the translation defined in Figure 3. In the presence of cuts, more than one set of linkings may correspond to the same proof. For example, since the two derivations in Figure 6 have the same underlying MALL proof (say Π), the concluding sets of linkings (say θ and θ') have a common sequentialisation: $\Pi \rightsquigarrow \theta$, $\Pi \rightsquigarrow \theta'$, and $\theta \neq \theta'$.

We can of course extend the cut-free translation of proofs by always choosing Σ to be empty in the &-rule (*i.e.*, "never superimpose cuts"). However, our notion of proof net defined below, which characterises sequentialisability, does not characterise the image of this translation, since there would exist sequentialisable sets of linkings that are not proof translations, such as

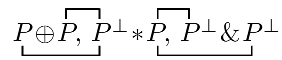

Moreover, under this convention two proofs that differ only in a commutation of cut and &-rules would be translated to different sets of linkings.

Note that the alternative of taking Σ maximal (*i.e.*, "superimpose as many cuts as possible") does not define a canonical function from proofs to sets of linkings, since there may be a choice of how to make the identifications. The following two &-rules illustrate such a choice.

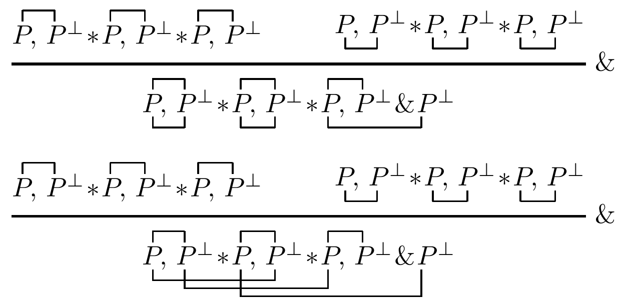

Girard was aware of this issue in the context of monomial proof nets; see Appendix A.1.6 of [Gir96].

Geometric characterisation of sequentialisability. In the presence of cut, we update all the auxiliary definitions of Section 2 (&-resolution, $\mathcal{G}_\Lambda$, switching cycle, *etc.*) by substituting "cut sequent" for "sequent" and "cut-additive resolution" for "additive resolution" throughout.

DEFINITION 2 *A set θ of linkings on a cut sequent Δ is a* ***proof net*** *if:*

(P0) *At least one literal occurrence of every cut is in θ (i.e., in some axiom link of some linking of θ).*

(P1) *For every &-resolution $\Delta^\star$ of Δ, exactly one linking of θ is on $\Delta^\star$.*

(P2) *Every switching of every linking of θ is a tree (acyclic and connected[8]).*

(P3) *Every set Λ of two or more linkings of θ toggles a & that is not in any switching cycle of Λ.*

θ is a ***proof structure*** *if it satisfies (P0) and (P1).*

Note that (P1)–(P3) are inherited from the cut-free case.

THEOREM 2 (SEQUENTIALISATION) *A set of linkings is sequentialisable iff it is a proof net.*

The proof is essentially the same as the proof of Theorem 1; the cut connective $*$ is akin to an outermost $\otimes$.

[8] By dropping connectedness, we obtain a proof net for MALL augmented by the MIX rule.

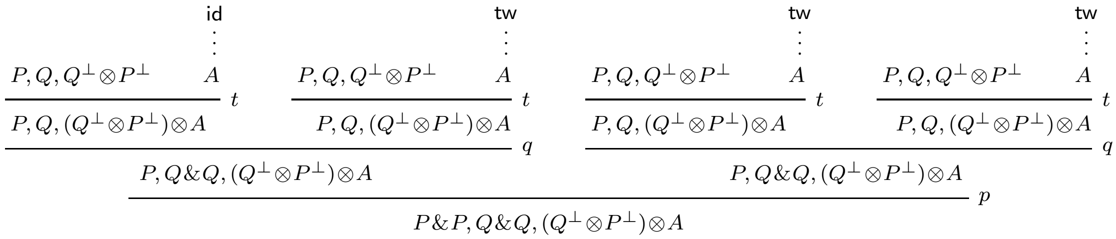

Figure 7. The proof Π_{tqp}. (We omit the unique cut-free proof of $P, Q, Q^\perp \otimes P^\perp$.)

Cut elimination. Let θ be a set of linkings on the cut sequent Δ, and let $A * A^\perp$ be a cut of Δ. Define the ***elimination*** of $A * A^\perp$ as follows.

- If A is a literal, delete $A * A^\perp$ from Δ, and replace any pair of axiom 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 axiom link $\{l, l'\}$.
- 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 cuts $A_1 * A_1^\perp$ and $A_2 * A_2^\perp$. Retain all the original linkings.
- 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 cuts $A_1 * A_1^\perp$ and $A_2 * A_2^\perp$. Retain precisely the 'consistent' linkings: delete any linkings λ such that in $\Delta \restriction \lambda$ the (now unary) $\&$ and $\oplus$ take opposite arguments (*i.e.*, such that the right argument of the $\&$ is in $\Delta \restriction \lambda$ and the left argument of the $\oplus$ is in $\Delta \restriction \lambda$, or vice versa). Finally, 'garbage collect' by deleting $A_i * A_i^\perp$ if no literal occurrence of $A_i * A_i^\perp$ is in any of the remaining linkings.

An example of cut elimination was presented in the Introduction.

PROPOSITION 1 *Eliminating a cut from a proof net yields a proof net.*

Proposition 1 is proved in the full version of this paper.

THEOREM 3 *Cut elimination of proof nets is strongly normalising.*

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

A category of proof nets. Our cut elimination allows us to define a category 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 $A^\perp, B * B^\perp, C$. Composition is associative, since cut elimination is strongly normalising. The identity $A \to A$ contains a linking λ on $A^\perp, A$ iff λ matches the i^{th} literal of $A^\perp$ with the i^{th} literal of A. This category has the structure of a star-autonomous category minus the units; $\&$ is product, and $\oplus$ is coproduct.

4 Girard's monomial proof nets

In this section we relate our MALL proof nets to the monomial proof nets of Girard [Gir96].

4.1 Monomial proof nets are unsatisfactory

We give a detailed account of how monomial proof nets [Gir96] fail to provide abstract representations of cut-free MALL proofs modulo inessential commutations of rules. A single cut-free proof may correspond to a host of monomial proof nets, and there is no natural translation of cut-free MALL proofs into monomial proof nets. (The reader unfamiliar with monomial proof nets should be able to follow the general shape of the argument.)

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

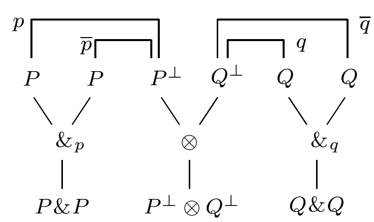

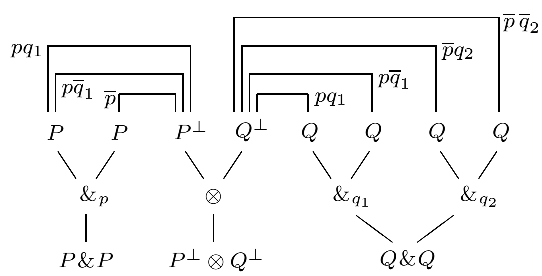

(Eigenvariables associated with $\&$'s are shown as subscripts; we omit implied weights.) These two monomial proof nets correspond to the same proof. The second monomial proof net has two forms of redundancy relative to the first: (i) the $\&$ with eigenweight q has been replaced by two similar 'copies',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

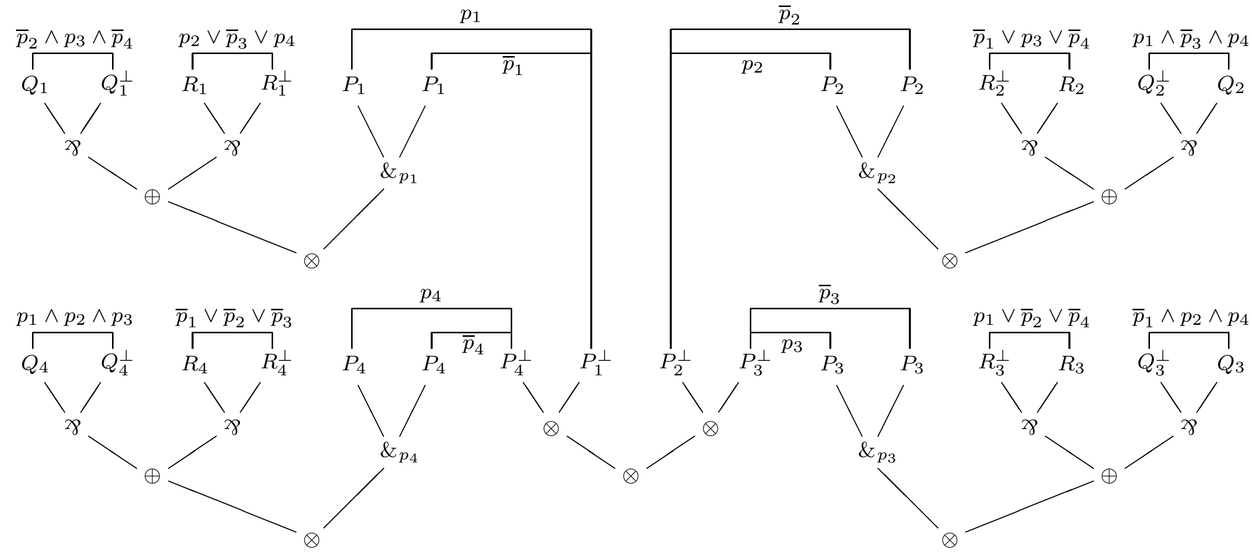

Figure 8. Girard's correctness criterion is insufficient without monomials: this (abbreviated) non-monomial Girard proof net is not sequentialisable.

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

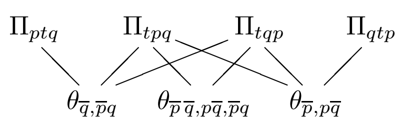

Define Π_{tqp} to be the proof shown in Figure 7, where $A = (R \parr R) \parr (R^\perp \otimes R^\perp)$, id denotes the identity proof and tw denotes the twist proof. Let Π_{tpq} be the result of commuting rules q and p in Π_{tqp}, let Π_{qtp} be the result of commuting t and q in the right half of 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:

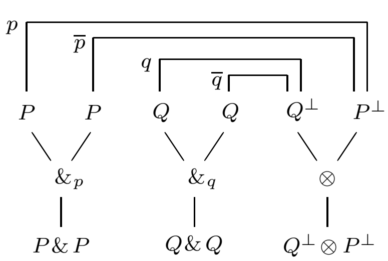

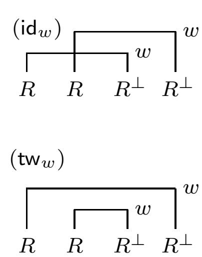

We have taken as eigenweights the labels of the $\&$-rules of the Π_α. The configuration of axiom links on A will be a disjoint union of axiom links in the identity and twist configurations: id_w and tw_w (above-right) denote a pair of axiom links of weight w in the identity and twist configurations,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

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

(By redundancies of type (i) and (ii) illustrated earlier, there are of course a host of other monomial proof nets θ_β which are parodies of the three above.) Since the Π_α are equivalent modulo inessential rule commutations, 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 sequentialisation relation between proofs and monomial proof nets depicted earlier, thereby creating a matching between 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.

By contrast, in our setting we associate the same proof net with each Π_α: the four-linking proof net on the second page of the Introduction. Thus we preserve the spirit of MLL proof nets by providing an abstract representation of all of the Π_α in one.

4.2 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 problem of whether Girard's proof net correctness criterion [Gir96] becomes insufficient upon relaxing the dependency condition, which is the requirement that every weight be a monomial. The answer is yes: in Figure 8 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 [Gir96] but for the omission of the dependency condition.

Strictly speaking θ is an abbreviation of a non-monomial Girard proof net: 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.

Figure 8 also encodes one of our proof structures θ, via the notion of weight described on the third page of the Introduction. 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 $\&_{p_i}$ to the axiom link $\{R_{i+1}^\perp, R_{i+1}\}$ (mod 4)).

4.3 Mapping monomial proof structures to ours

Let a ***non-monomial Girard proof structure*** be a proof structure as in [Gir96] but for the omission of the dependency condition. Define a non-monomial Girard proof structure to be ***compact*** if (a) every 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 with the same conclusion(s), and summing the weights of links and formulas so identified. This collapse does not preserve the dependency condition. Any compact non-monomial Girard proof structure can be obtained as the collapse of a (monomial) Girard proof structure.

Compact non-monomial Girard proof structures are in bijection with our proof structures. The counterpart of Girard's "technical condition" is implied by (P1) and our definition of a set of linkings in terms of additive resolutions. The surjective map from (monomial) Girard proof structures to our proof structures obtained by composing the collapse and the bijection preserves the property of being a sequentialisation of a particular MALL proof.

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 every axiom link $a \in \mathcal{G}_\Lambda$ depending on w *within* θ. 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 (also deleting all but one switch edge of each $\&$). (P2*) clearly implies (P2), since it involves more switchings. In fact, (P2*) is strictly stronger than (P2): for $\theta = \{\lambda_1, \lambda_2\}$ of Example 1, the graph $\mathcal{G}^\theta_{\lambda_1}$ has a switching cycle (the one presented below Figure 4), whereas $\mathcal{G}_{\lambda_1}$ does not. However, it is not hard to check that (P2*) is implied by (P2) and (P3) together.

The bijection between compact non-monomial Girard proof structures and our proof structures can now be further refined: compact non-monomial Girard proof nets are in bijection with sets of linkings in our sense which satisfy (P1) and (P2*).

5 Work in progress

The equivalence relation on cut-free MLL proofs induced by their translation into cut-free MLL proof nets is canonical in sense that the equivalence corresponds to coherence in a star-autonomous category [BCST96]. We conjecture that the equivalence on cut-free MALL proofs induced by our translation into proof nets corresponds to coherence in a star-autonomous category with products (hence sums).

We are seeking a reformulation of cut that preserves the elegance of the cut-free definition, in the sense of retaining a natural translation from proofs to proof nets.

We are investigating whether the following variant of (P3) yields an alternative definition of proof net: for any switching cycle S of a set of linkings Λ, at least one $\&$ toggled in Λ is not in S.

Acknowledgements. Vaughan Pratt, for invaluable feedback during the development of this work. Paul-André Melliès, for suggesting to the first author that the search for a satisfactory notion of MALL proof net was an interesting and potentially fruitful research topic.

References

[AJ94] S. Abramsky & R. Jagadeesan (1994): *Games and full completeness for multiplicative linear logic. Journal of Symbolic Logic* 59(2), pp. 543–574.

[AM99] S. Abramsky & P.-A. Melliès (1999): *Concurrent games and full completenes.* In *Proc. LICS'99*, IEEE, pp. 431–442.

[BCST96] R. F. Blute, J. R. B. Cockett, R. A. G. Seely & T. H. Trimble (1996): *Natural deduction and coherence for weakly distributive categories. Journal of Pure and Applied Algebra* 113, pp. 229–296.

[BS96] R. F. Blute & P. J. Scott (1996): *Linear Läuchli semantics. Annnals of Pure & Applied Logic* 77, pp. 101–142.

[DHPP99] H. Devarajan, D. Hughes, G. Plotkin & V. R. Pratt (1999): *Full completeness of the multiplicative linear logic of chu spaces.* In *Proc. LICS'99*, IEEE, pp. 234–245.

[DR89] V. Danos & L. Regnier (1989): *The structure of multiplicatives. Arch. Math. Logic* 28, pp. 181–203.

[Gir87] J.-Y. Girard (1987): *Linear logic. Theoretical Computer Science* 50, pp. 1–102.

[Gir96] J.-Y. Girard (1996): *Proof-nets: the parallel syntax for proof theory. Logic and Algebra.* Marcel Dekker.

[HO93] J. M. E. Hyland & C.-H. L. Ong (1993): *Fair games and full completeness for multiplicative linear logic without the* MIX*-Rule.* On Ong's web page.

[Loa94] R. Loader (1994): *Linear logic, totality and full completeness.* In *Proc. LICS'94*, IEEE, pp. 292–298.

[Tan97] A. Tan (1997): *Full completeness for models of linear logic*. PhD thesis, University of Cambridge.