December 17, 2007

“Introduction to Deep Inference and Proof Nets”

Paola Bruscoli
Department of Computer Science
University of Bath
Bath BA2 7AY
United Kingdom
http://www.cs.bath.ac.uk/pb/

Lutz Straßburger
INRIA Futurs — Parsifal
École Polytechnique — LIX
Rue de Saclay, 91128 Palaiseau
France
http://www.lix.polytechnique.fr/~lutz/

These are the notes for the fourth lecture (written by Lutz Straßburger).

---

1 Unit-free multiplicative linear logic

Today we learn about proof nets. For this, we will consider only multiplicative linear logic without units. The sequent calculus system MLL^- was shown in the previous lecture. In the calculus of structures, The unit-free system looks as follows:

$$\mathsf{ai}\!\downarrow \frac{}{[a \parr a^\perp]} \qquad \mathsf{ai}\!\downarrow \frac{S\{B\}}{S\{B \otimes [a \parr a^\perp]\}} \qquad \mathsf{s}\,\frac{S\{[A \parr B] \otimes C\}}{S\{A \parr (B \otimes C)\}} \tag{1}$$

Because there are no units present, the $\mathsf{ai}\!\downarrow$ rule looks slightly different from what you have seen before, and there are two versions of it. We use the name MLS^- for the system in (1). We can drop the equations for the units

$$\begin{array}{rclcrcl} (A \otimes (B \otimes C)) &=& ((A \otimes B) \otimes C) & \qquad & (A \otimes B) &=& (B \otimes A) \\ [A \parr [B \parr C]] &=& [[A \parr B] \parr C] & & [A \parr B] &=& [B \parr A] \end{array} \tag{2}$$

The cut rule is

$$\mathsf{ai}\!\uparrow \frac{S\{B \parr (a \otimes a^\perp)\}}{S\{B\}} \tag{3}$$

For simplicity, we will also use the general versions of the interaction rules:

$$\mathsf{ai}\!\downarrow \frac{}{[A \parr A^\perp]} \qquad \mathsf{ai}\!\downarrow \frac{S\{B\}}{S\{B \otimes [A \parr A^\perp]\}} \qquad \mathsf{ai}\!\uparrow \frac{S\{B \parr (A \otimes A^\perp)\}}{S\{B\}} \tag{4}$$

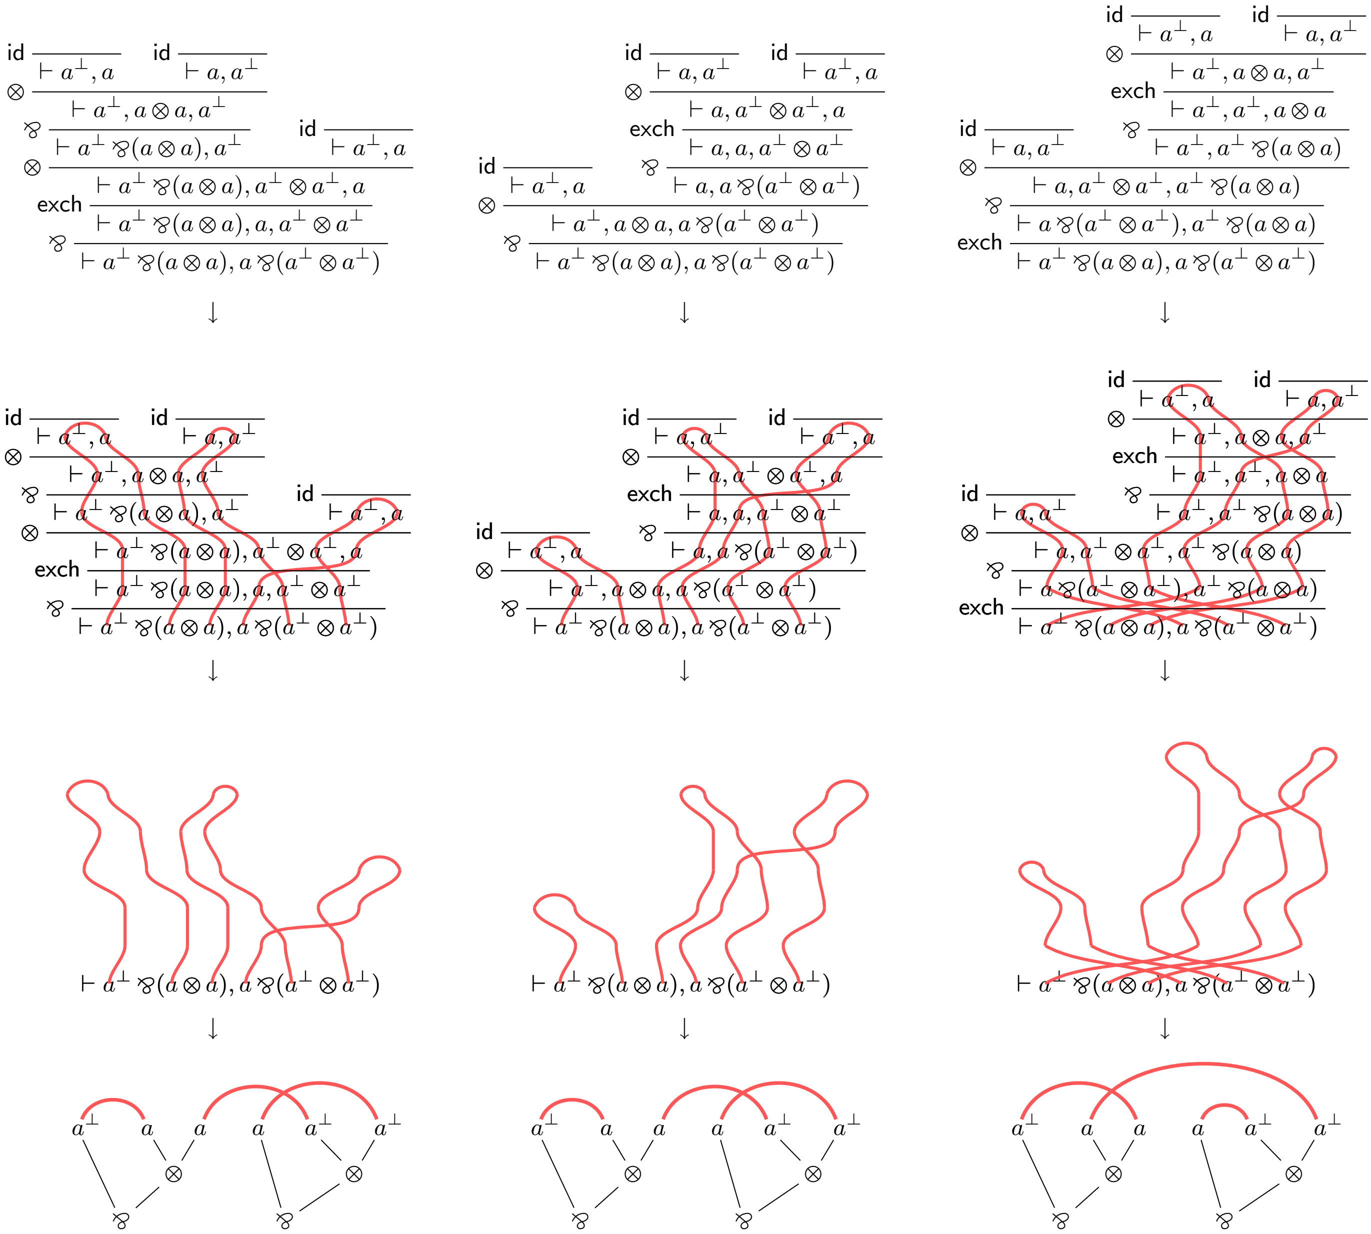

Figure 1: From sequent calculus to proof nets via coherence graphs

2 From sequent calculus to proof nets

It turns out that for MLL^- the "essence" of a proof is captured by the axiom links. More precisely, the proof net is obtained by drawing the "flow-graph" (or "coherence-graph") through the sequent calculus proof. This means that we trace all atom occurrences through the proof. The idea is quite simple, but the formal definitions tend to be messy. In these lecture notes,

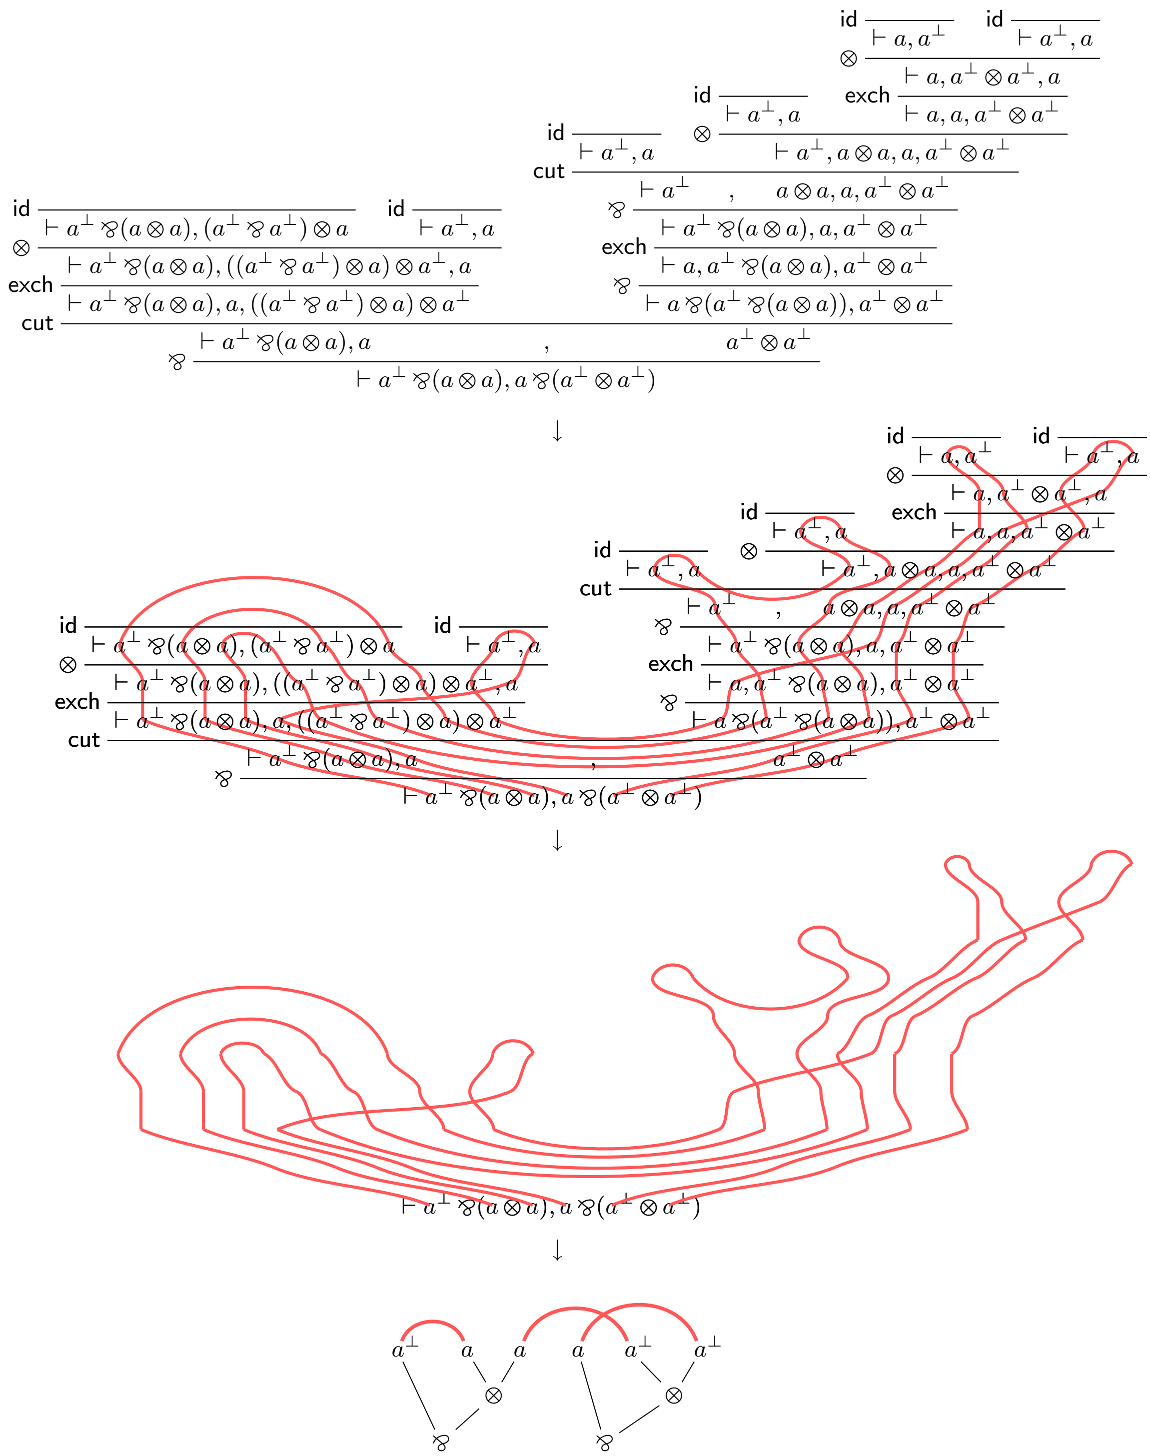

Figure 2: From sequent calculus to proof nets via coherence graphs

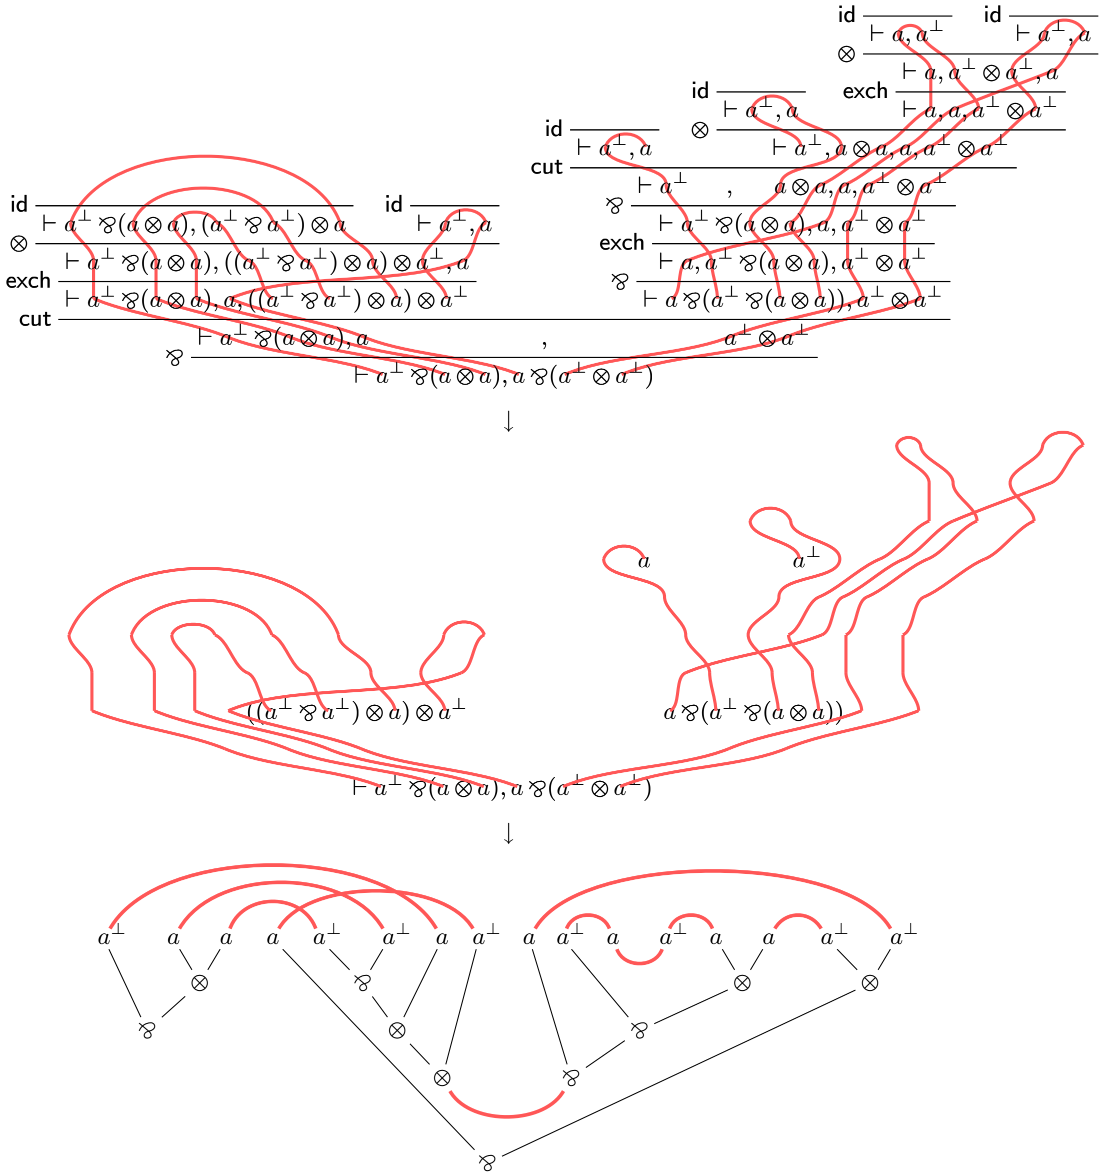

Figure 3: From sequent calculus to proof nets with cuts via coherence graphs

we show the idea via examples in Figure 1.

In Figure 2 we convert an example with cut into a proof net via the flow-graph method.

For dealing with cuts (without forgetting them!), we can prevent the flow-graph from flowing through the cut, i.e., by keeping the information that there is a cut. What is meant by this is shown in Figure 3.

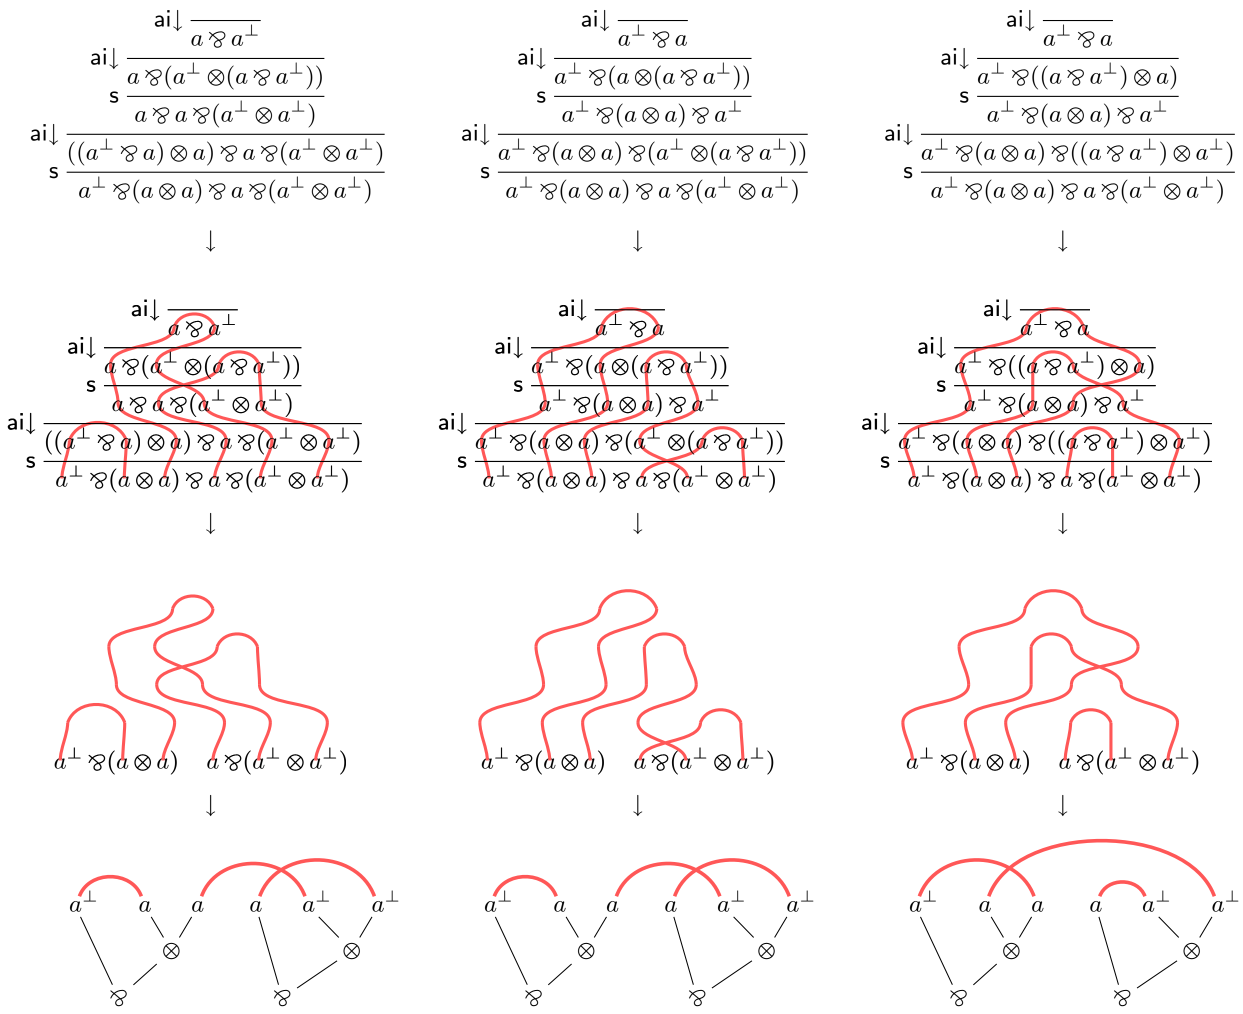

Figure 4: From calculus of structures to proof nets

3 From the calculus of structures to proof nets

In this section we do the same as in the previous section. But this time, we start from MLS^- instead if MLL^-. But the result is exactly the same.

We simply trace the atoms through the derivation. Figures 4–6 show the calculus of structures version of Figures 1–3.

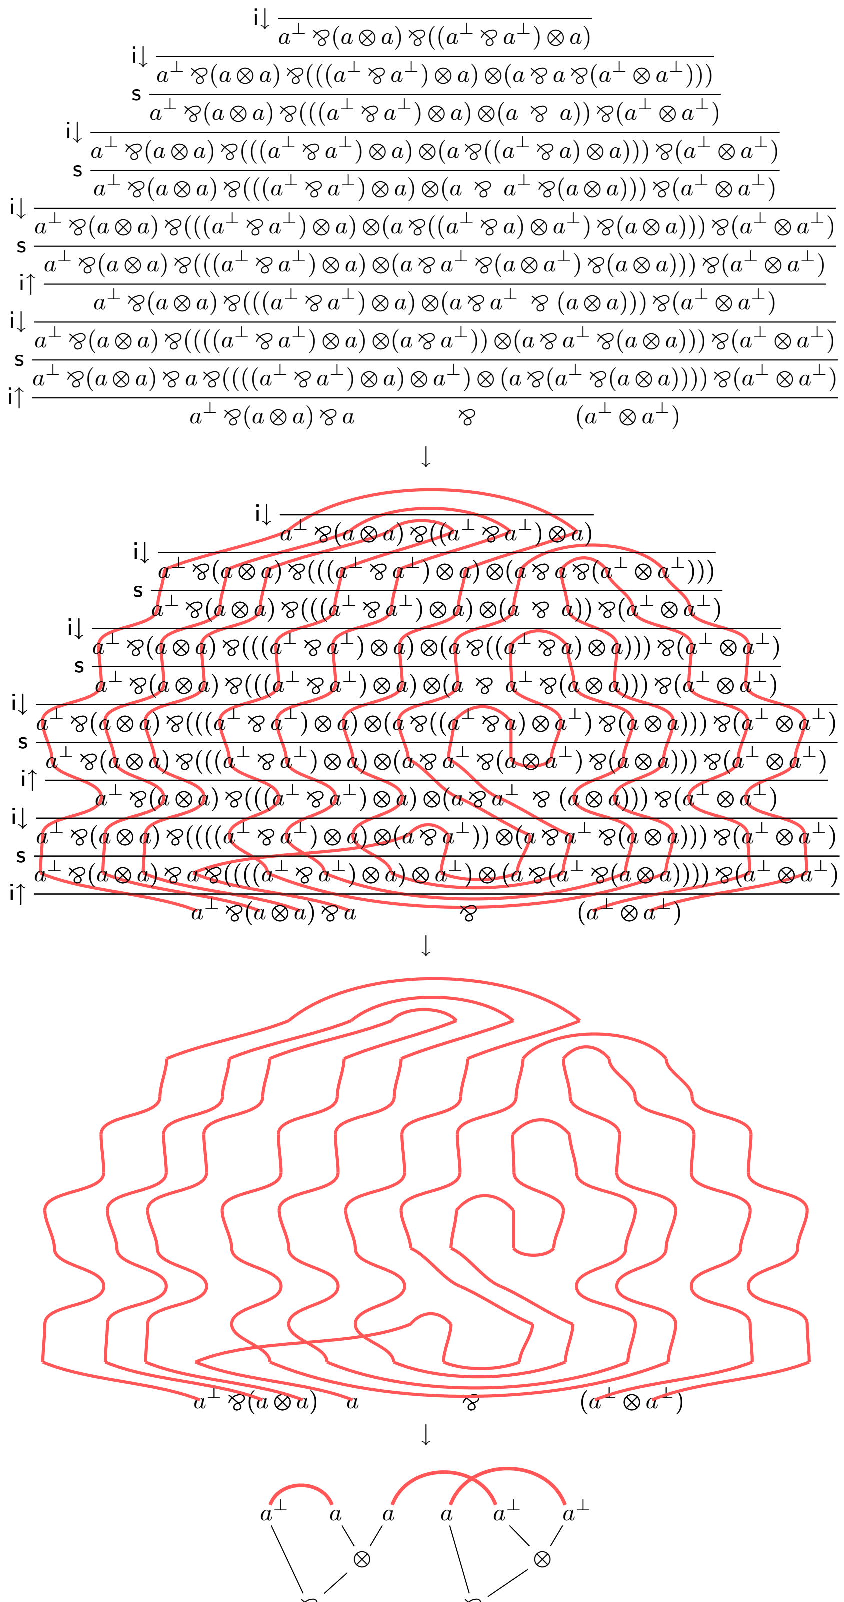

Figure 5: From calculus of structures to proof nets

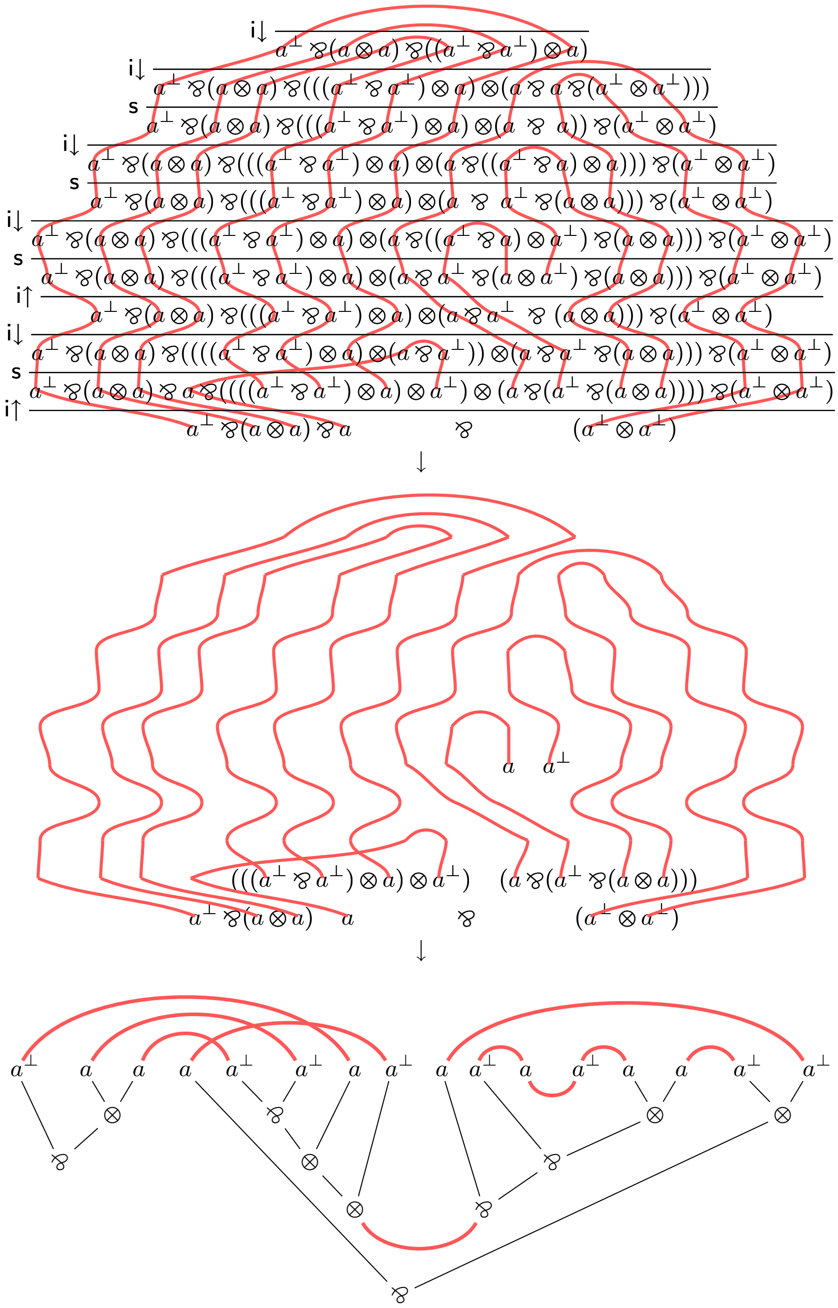

Figure 6: From calculus of structures to proof nets

4 Correctness criteria

We have seen how we can obtain a proof net out of a formal proof in some deductive system. But what about the other way around? Suppose we have such a graph that looks like a proof net. Can we decide whether it really comes from a proof, and if so, can we recover this proof? Of course the answer is trivially yes because the graph is finite and we just need to check all proofs of that size. The interesting question is therefore, whether we can do it efficiently.

The answer is still yes, and it is done via so-called *correctness criteria*. For introducing the idea, we take the following graphs as running examples

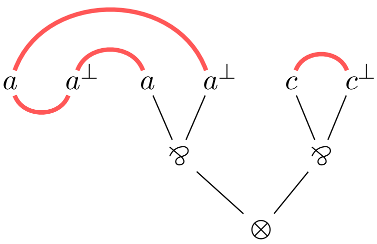

(5)

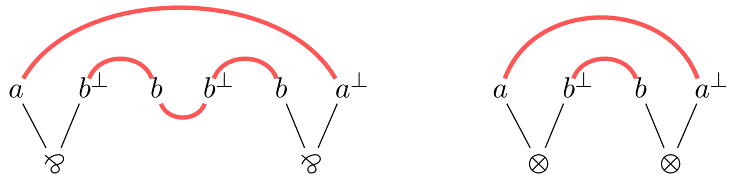

(6)

By playing around, you will notice that it is quite easy to find a proof (in sequent calculus or calculus of structures) that translates into the net in (5), but it seems impossible to find such proofs for the two examples in (6). We are now going to show that this is indeed impossible. For doing so, we need some formal definitions.

4.1 Definition A *pre-proof net* is a sequent forest Γ, possibly with cuts, together with a perfect matching of the set of leaves (i.e., the set of occurrences of propositional variables and their duals), such that only dual pairs are matched.

In this context, a cut must be seen as a special kind of formula $A \varobar A^\perp$, where $\varobar$ is a special connective which may occur only at the root of a formula tree in which the two direct subformulas are dual to each other. For example, (5) should be read as

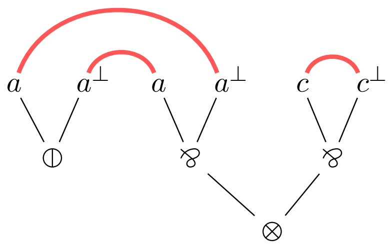

(5')

Clearly, the examples in (5) and (6) are all pre-proof nets. In the following, we will think of an inner node (i.e., a non-leaf node) of the sequent forest labeled not only by the connective but by the whole subformula rooted by that connective. Our favorite example (5) should then be read as

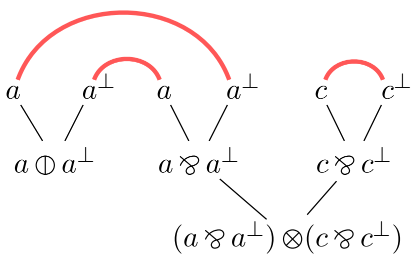

(5")

Although sometimes we think of pre-proof nets to be written as in (5"), we will keep writing them as in (5) for better readability.

4.2 Definition A pre-proof net π is called *sequentializable* iff there is a proof in the sequent calculus or in the calculus of structures that translates into π.

Originally, the term “sequentializable” was motivated by the name “sequent calculus”. But we use it here also if the “sequentialization” is done in the calculus of structures.

4.3 Definition Let π be a pre-proof net. A *DR-switching* for π is a graph obtained from π by removing for every ⅋-node one of the two edges connecting it to its children.

Clearly, if a pre-proof net contains n ⅋-nodes, then there are 2^n switchings. Here are all 4 switchings for the example in (5):

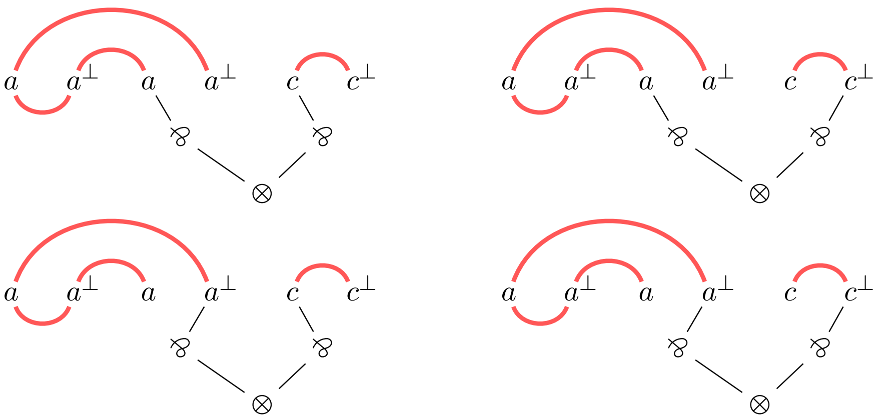

(7)

4.4 Definition A pre-proof net *obeys the DR-switching criterion* (or, shortly, is *correct*) iff all its switchings are connected and acyclic.

As (7) shows, the pre-proof net in (5) is correct. The two pre-proof nets in (6) are not, as the following switchings show:

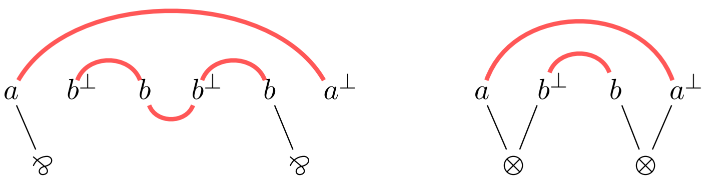

(8)

The first is not connected, and the second is cyclic.

In the following, we use the term *proof net* for those pre-proof nets which are correct, i.e., obey the switching criterion. The following theorem says that the proof nets are exactly those pre-proof nets that represent an actual proof.

4.5 Theorem *A pre-proof net is correct if and only if it is sequentializable.*

We will give two proofs of this theorem. The first uses the sequent calculus, and the second the calculus of structures. For the first proof, we need the following lemma:

4.6 Lemma *Let π be a proof net with conclusions $A_1, \ldots, A_n$. If all A_i have a $\otimes$ or a cut as root, then one of them is splitting, i.e., by removing that $\otimes$ (or ⦶), the net becomes disconnected.*

For proving this lemma, we need some more concepts.

4.7 Definition Let σ and π be pre-proof nets. We say σ is a *subprenet* of π, written as $\sigma \subseteq \pi$ if all formulas/cuts appearing in σ are subformulas of the formulas/cuts appearing in π, and the linking of σ is the restriction of the linking of π to the formulas/cuts in σ. We say σ is a *subnet* of π if $\sigma \subseteq \pi$, and σ and π are both correct. A *door* of σ is any formula that appears as conclusion of σ.

4.8 Example Consider the following three graphs:

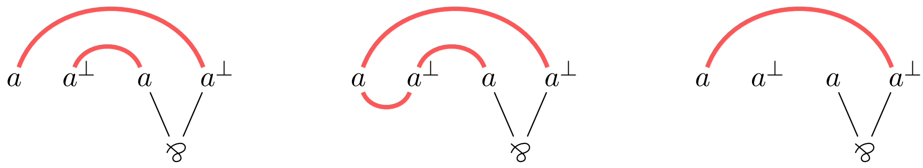

The first two are subprenets of (5), the third one is not (because a link is missing). The second one is in fact a subnet of (5), but the first one is not (because it is not correct). The doors of the first example are a, $a^\perp$, and $a \bindnasrepma a^\perp$. The doors of the second example are $a \oplus a^\perp$ and $a \bindnasrepma a^\perp$.

4.9 Lemma *Let σ and ρ be subnets of some proof net π.*

(i) *The subprenet $\sigma \cup \rho$ is a subnet of π if and only if $\sigma \cap \rho \neq \emptyset$.*

(ii) *If $\sigma \cap \rho \neq \emptyset$ then $\sigma \cap \rho$ is a subnet of π.*

Proof: Intersection and union in the statement of that lemma have to be understood in the canonical sense: An edge/node/link appears in in $\sigma \cap \rho$ (resp. $\sigma \cup \rho$) if it appears in both, σ and ρ (resp. in at least one of σ or ρ). For giving the proof, let us first note that because in π every switching is acyclic, also in every subprenet of π every switching is acyclic, in particular also in $\sigma \cup \rho$ and $\sigma \cap \rho$. Therefore, we need only to consider the connectedness condition.

(i) If $\sigma \cap \rho = \emptyset$ then every switching of $\sigma \cup \rho$ must be disconnected. Conversely, if $\sigma \cap \rho \neq \emptyset$, then every switching of $\sigma \cup \rho$ must be connected (in every switching of $\sigma \cup \rho$ every node in $\sigma \cap \rho$ must be connected to every node in σ and to every node in ρ, because σ and ρ are both correct).

(ii) Let $\sigma \cap \rho \neq \emptyset$ and let s be a switching for $\sigma \cup \rho$. Then s is connected and acyclic by (i). Let s_σ, s_ρ, and $s_{\sigma\cap\rho}$, be the restrictions of s to σ, ρ, and $\sigma \cap \rho$, respectively. Now let A and B be two vertices in $\sigma \cap \rho$. Then A and B are connected by a path in s_σ because σ is correct, and by a path in s_ρ because ρ is correct. Since s is acyclic, the two paths must be the same and therefore be contained in $s_{\sigma\cap\rho}$. $\square$

4.10 Lemma *Let π be a proof net, and let A be a subformula of some formula/cut appearing in π. Then there is a subnet σ of π, that has A as a door.*

Proof: For proving this lemma, we need the following notation. Let π be a proof net, let A be some formula occurrence in π, and let s be a switching for π. Then we write $s(\pi, A)$ for the graph obtained as follows:

- If A is an immediate subformula of a formula occurrence B in π, and there is an edge from B to A in s, then remove that edge and let $s(\pi, A)$ be the connected component of (the remainder of) s that contains A.

- Otherwise let $s(\pi, A)$ be just s.

Now let

$$\sigma = \bigcap_s s(\pi, A)$$

where s ranges over all possible switchings of π. (Note that it could happen that formally σ is not a subprenet because some edges in the formula trees might be missing. We graciously add these missing edges to σ such that it becomes a subprenet.) Clearly, A is in σ. We are now going to show that A is a door of σ. By way of contradiction, assume it is not. This means there is ancestor B of A that is in $\bigcap_s s(\pi, A)$. Now choose a switching $\hat{s}$ such that whenever there is a $\parr$ node between A and B, i.e.,

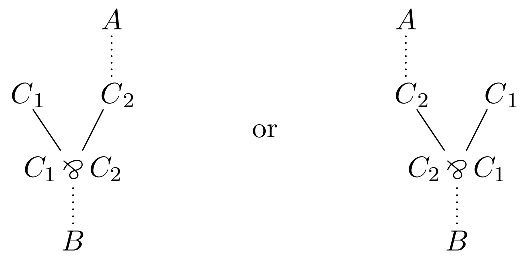

then $\hat{s}$ chooses C_1 (i.e., removes the edge between C_2 and its parent). Then there must be a $\otimes$ between A and B:

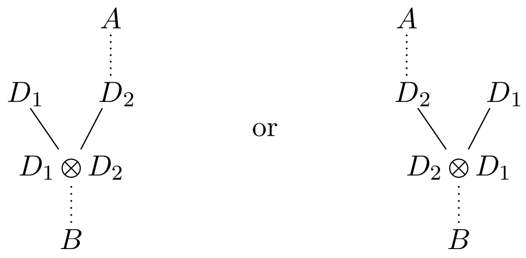

Otherwise B would not be in σ. Now suppose we have chosen the uppermost such $\otimes$. Then the path connecting A and D_1 in $\hat{s}(\pi, A)$ cannot pass through D_2 (by the definition of $\hat{s}(\pi, A)$). But this means that in $\hat{s}$ there are two distinct paths connecting A and D_1, which contradicts the acyclicity of $\hat{s}$.

Now we have to show that σ is a subnet. Let s be a switching for σ. Since σ is a subprenet of π, we have that s is acyclic. Now let $\tilde{s}$ be an extension of s to π. Then s is the restriction of $\tilde{s}(\pi, A)$ to σ, and hence connected. $\square$

4.11 Definition Let π be a proof net, and let A be a subformula of some formula/cut appearing in π. The *kingdom of A in π*, denoted by kA, is the smallest subnet of π, that has A as a door. Similarly, the *empire of A in π*, denoted by eA, is the largest subnet of π, that has A as a door. We define $A \ll B$ iff $A \in kB$, where A and B can be any (sub)formula/cut occurrences in π.

An immediate consequence of Lemmas 4.9 and 4.10 is that kingdom and empire always exist.

4.12 Exercise Why?

4.13 Remark The subnet σ constructed in the proof of Lemma 4.10 is in fact the empire of A. But we will not need this fact later and will not prove it here.

4.14 Lemma *Let π be a proof net, and let A, A', B, and B' be subformula occurrences appearing in π, such that A and B are distinct, A' is immediate subformula of A, and B' is immediate subformula of B. Now suppose that $B' \in eA'$. Then we have that $B \notin eA'$ if and only if $A \in kB$.*

Proof: We have $B' \in eA' \cap kB$. Hence, $\sigma = eA' \cap kB$ and $\rho = eA' \cup kB$ are subnets of π. By way of contradiction, let $B \notin eA'$ and $A \notin kB$. Then ρ has A' as door and is larger than eA' because it contains B. This contradicts the definition of eA'. On the other hand, if $B \in eA'$ and $A \in kB$ then σ has B as door and is smaller than kB because it does not contain A. This contradicts the definition of kB. □

4.15 Lemma *Let π be a proof net, and let A and B be subformulas appearing in π. If $A \ll B$ and $B \ll A$, then either A and B are the same occurrence or they are dual atoms connected via an identity link.*

Proof: If a and $a^\perp$ are two dual atom occurrences connected by a link, then clearly $ka = ka^\perp$. Now let A and B be two distinct non-atomic formula occurrences in π with $A \in kB$ and $B \in kA$. Then $kA \cap kB$ is a subnet and hence $kA = kA \cap kB = kB$. We have two cases:

- If $A = A' \parr A''$ then the result of removing A from kB is still a subnet, contradicting the minimality of kB.
- If $A = A' \otimes A''$ then $kA = kA' \cup kA'' \cup \{A' \otimes A''\}$. Hence $B \in kA'$ or $B \in kA''$. This contradicts Lemma 4.14, which says that $B \notin eA'$ and $B \notin eA''$. □

From Lemma 4.15 it immediately follows that $\ll$ is a partial order on the non-atomic subformula occurrences in π. We make crucial use of this fact in the following:

Proof of Lemma 4.6: Choose among the conclusions $A_1, \ldots, A_n$ (including the cuts) of π one which is maximal w.r.t. $\ll$. Without loss of generality, assume it is $A_i = A'_i \otimes A''_i$. We will now show that it is splitting, i.e., $\pi = \{A'_i \otimes A''_i\} \cup eA'_i \cup eA''_i$. By way of contradiction, assume $A'_i \otimes A''_i$ is not splitting. This means we have somewhere in π a formula occurrence B with immediate subformula B' such that (without loss of generality) $B' \in eA'_i$ and $B \notin eA'_i$. We also know that B must occur at or above some other conclusion, say $A_j = A'_j \otimes A''_j$. Hence $B \in kA_j$ and therefore $kB \subseteq kA_j$. But by Lemma 4.14 we have $A_i \in kB$ and therefore $A_i \in kA_j$, which contradicts the maximality of A_i w.r.t. $\ll$. □

Finally, we can prove Theorem 4.5.

First Proof of Theorem 4.5: Let us first show that the (in the sequent calculus) sequentializable pre-proof nets are indeed correct. This is done by verifying that the id-rule yields correct nets (which is obvious) and that all other inference rules preserve correctness. For the exch-rule this is obvious. Let us now consider the $\otimes$-rule. By way of contradiction, assume that

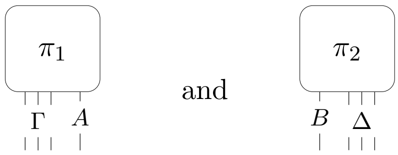

are correct, but

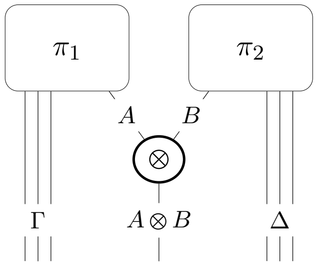

is not correct. This means there is a switching that is either disconnected or contains a cycle. Since a $\otimes$-node does not affect switchings, we conclude that the property of being disconnected or cyclic must hold for the same switching in one of π_1 or π_2. But this is a contradiction to the correctness of π_1 and π_2. For the the $\parr$-rule and the cut-rule we proceed similarly.

Conversely, let π be a correct pre-proof net. We proceed by induction on the size of π, i.e., the number n of $\parr$-, $\otimes$-, and cut-nodes in π, to construct a sequent calculus proof Π, that translates into π. If n is 0, then π must be of the shape

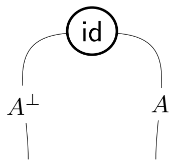

and we can apply the id-rule. Now let $n > 0$. If one of the conclusion formulas of π has a $\parr$-root, we can apply the $\parr$-rule and proceed by induction hypothesis. Now suppose all roots are $\otimes$ or cuts. Then we apply Lemma 4.6, which tells us, that there is one of them which splits the net. Assume, without loss of generality, that it is a $\otimes$-root, say $A_i = A'_i \otimes A''_i$. This means, the net is of the shape

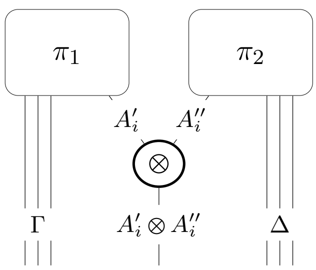

and we can apply the $\otimes$-rule and proceed by induction hypothesis for π_1 and π_2. If the splitting root is a cut, we apply the cut-rule instead. $\square$

Let us now see the second proof. For this, we need the following lemma:

4.16 Lemma *Let π be a proof net with conclusion*

$$S\{(A \otimes B\{a\}) \parr (C\{a^\perp\} \otimes D)\} \quad ,$$

such that the a and the $a^\perp$ are paired up in the linking. Let π' and π'' be pre-proof nets with conclusions

$$S\{A \otimes [B\{a\} \parr (C\{a^\perp\} \otimes D)]\} \qquad \textit{and} \qquad S\{[(A \otimes B\{a\}) \parr C\{a^\perp\}] \otimes D\}$$

respectively, such that the linkings of π' and π'' (i.e., the pairing of dual atoms) are the same as the linking of π. Then at least one of π' and π'' is also correct.

Proof: Let us visualize the information we have about π, π', and π'' as follows:

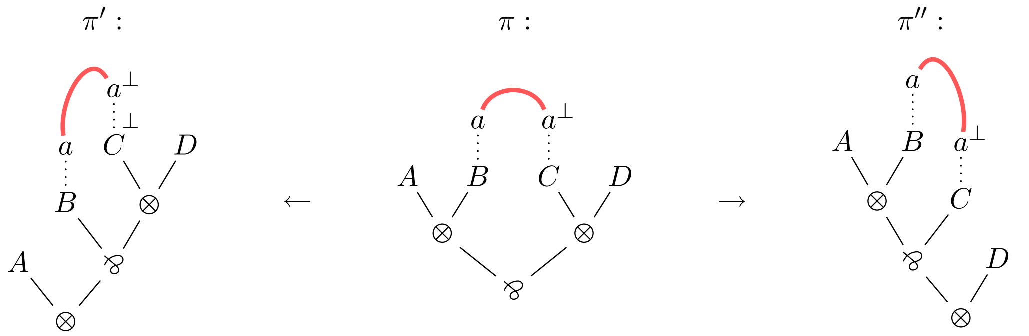

We proceed by way of contradiction, and assume that π is correct and that π' and π'' are both incorrect. If there is a switching s for π' (or π'') that is disconnected, then the same switching is also disconnected in π. Hence, we need to consider only the acyclicity condition. Suppose that there is a switching s' for π' that is cyclic. Then, in s' the $\parr$ below B must be switched to the right, and the cycle must pass through A, the root $\otimes$ and the $\parr$ as follows:

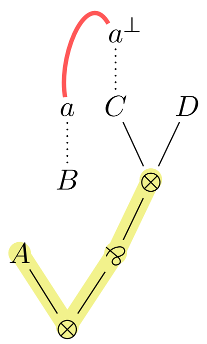

Otherwise we could construct a switching with the same cycle in π. If our cycle continues through D, i.e.,

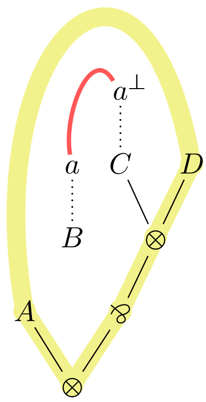

(9)

then we can use the path from A to D (that does not go through B or C, see Exercise 4.17) to construct a cyclic switching s in π as follows:

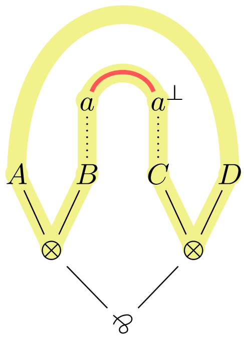

Hence, the cycle in s' goes through C, giving us a path from A to C, not passing through B (see Exercise 4.17):

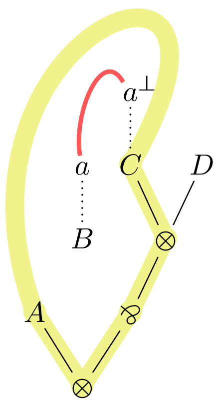

(10)

By the same argumentation we get a switching s'' in π'' with a path from B to D, not going through C. From s' and s'', we can now construct a switching s for π with a cycle as follows:

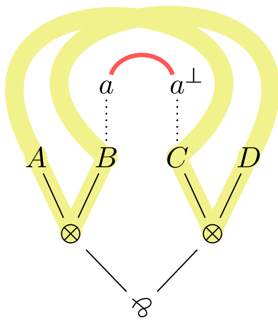

which contradicts the correctness of π. □

4.17 Exercise Explain why we can in (9) assume that the cycle does not go through B or C, and in (10) not through B.

In our second proof of Theorem 4.5 we will also need the following concept:

4.18 Definition Let A be a formula. The *relation web* of A is the complete graph, whose vertices are the atom occurrences in A. An edge between two atom occurrences a and b is colored red, if the first common ancestor of them in the formula tree is a $\otimes$, and green if it is a $\parr$.

4.19 Example Consider the formula $[[a^\perp \parr (a \otimes a)] \parr [a \parr (a^\perp \otimes a^\perp)]]$. Its formula tree is the following:

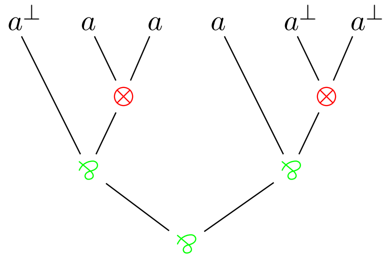

The relation web is therefore

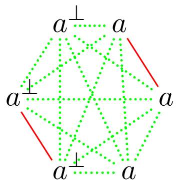

where we use regular edges for red and dotted edges for green.

4.20 Definition The *degree of freedom* of a formula A, is the number of green edges in its relation web.

Second Proof of Theorem 4.5: Again, we start by showing that all rules preserve correctness. Here, the only interesting case is the switch rule (all others being trivial), which does the following transformation somewhere inside the net:

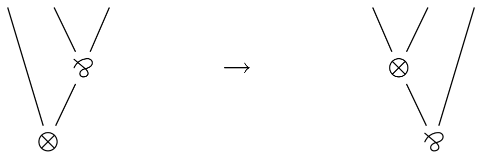

(11)

By way of contradiction, assume the net on the left is correct, and the one on the right is not. First, suppose there is a switching for the second net that is cyclic. If that cycle does not contain the $\otimes$-node shown on the right in (11), then this cycle is also present in the net on the left in (11). If our cycle contains the $\otimes$-node, then we can make the same cycle be present in the net on the left by switching the $\parr$-node to the left (i.e., removing the edge to the right). Now assume we have a disconnected switching for the net on the right. Then the same switching also disconnects the net on the left. Contradiction.

Conversely, assume we have a correct net π with conclusion F. For the time being, assume that π is cut-free. We proceed by induction on the degree of freedom of F. Pick inside F any pair of atoms that are linked together, say a and $a^\perp$. Then $F = S\{S_1\{a\} \parr S_2\{a^\perp\}\}$. Without loss of generality, we can assume that $S_1\{\ \}$ and $S_2\{\ \}$ are not $\parr$-contexts. We have the following cases:

- If $S_1\{\ \} = S_2\{\ \} = \{\ \}$, we can apply the rule $\mathsf{i}{\downarrow}$, and proceed by induction hypothesis.
- If $S_1\{\ \} \neq \{\ \}$ and $S_2\{\ \} = \{\ \}$, then $F = S\{(A \otimes B\{a\}) \parr a^\perp\}$ for some A and $B\{\ \}$. We can apply the switch rule to get $S\{A \otimes [B\{a\} \parr a^\perp]\}$, which is still correct (with the same linking as for F), but has smaller degree of freedom than F. The case where $S_1\{\ \} = \{\ \}$ and $S_2\{\ \} \neq \{\ \}$ is similar.
- If $S_1\{\ \} \neq \{\ \}$ and $S_2\{\ \} \neq \{\ \}$, then, without loss of generality, $F = S\{(A \otimes B\{a\}) \parr (C\{a^\perp\} \otimes D)\}$, for some A, $B\{\ \}$, $C\{\ \}$, D. By Lemma 4.16, we can apply the switch rule, since one of

  $$S\{A \otimes [B\{a\} \parr (C\{a^\perp\} \otimes D)]\} \quad \text{and} \quad S\{[(A \otimes B\{a\}) \parr C\{a^\perp\}] \otimes D\}$$

  is still correct. Since both of them have smaller degree of freedom than F, we can proceed by induction hypothesis.

If π contains cuts, we can replace inside π all cuts with $\otimes$, to get a formula F' such that there is a derivation

$$\mathsf{i}{\uparrow}\,\begin{array}{c} F' \\ \Vert \\ F \end{array}$$

Then π becomes a cut-free net with conclusion F', and we can proceed as above. $\square$

Note that the two different proofs of Theorem 4.5 yield a stronger version of the equivalence between MLL^- and MLS^- that we established in the previous lecture.

4.21 Theorem *For every sequent calculus proof of* $\vdash A_1, A_2, \ldots, A_n$ *in* MLS^- *there is a proof in the calculus of structures in system* MLS^- *of* $[A_1 \mathbin{\invamp} A_2 \mathbin{\invamp} \cdots \mathbin{\invamp} A_n]$ *yielding the same proof net, and vice versa.*

A geometric or graph-theoretic criterion like the one in Definition 4.4 and Theorem 4.5 is called a *correctness criterion*. The desired properties are soundness and completeness, as stated in Theorem 4.5. For MLL^-, the literature contains quite a lot of such criteria, and it would go far beyond the scope of this lecture notes to attempt to give a complete survey. But nonetheless, we will show here two other correctness criteria.

For the next one, we write the pre-proof nets in a different way:

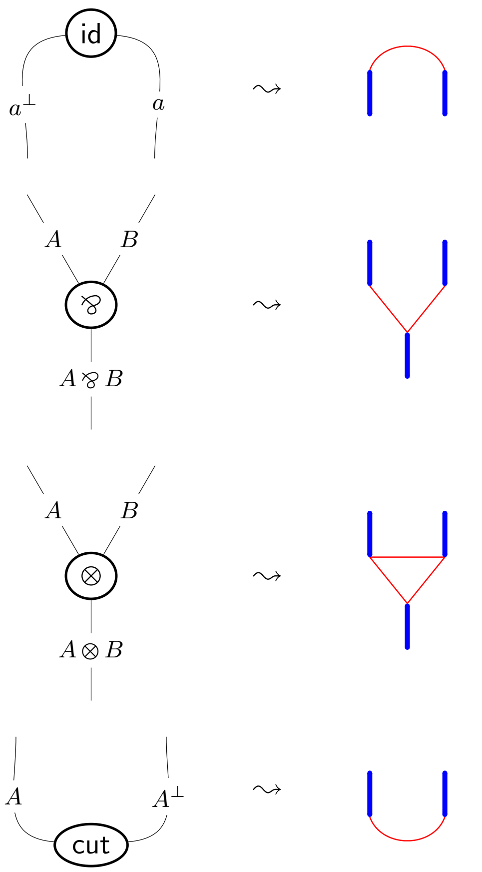

(12)

We call the resulting graphs *RB-graphs*. The R and B stand for Regular/Red and Bold/Blue. The main property of these graphs is that the blue/bold edges (in the following called *B-edges*) provide a bipartition of the set of vertices, i.e., every vertex in the RB-graph is connected to exactly one other vertex via a B-edge. The red/regular edges are in the following called *R-edges*.

Here are the examples from (5) and (6) written as RB-graphs:

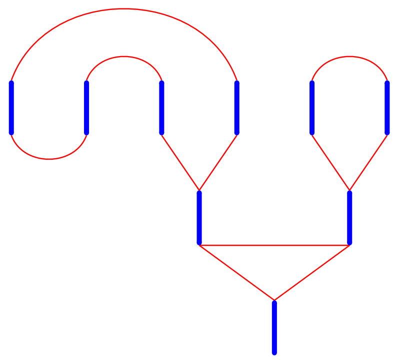

(13)

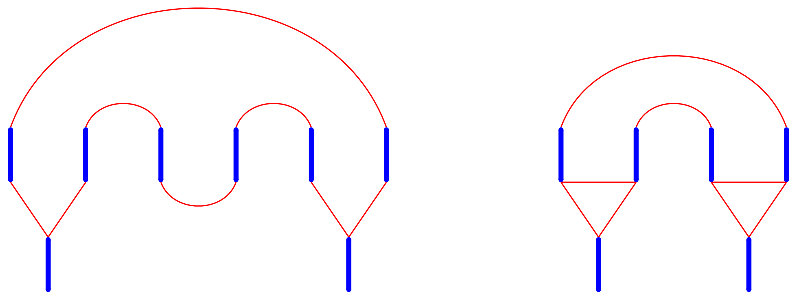

(14)

4.22 Definition Let G be an RB-graph. An *Æ-path* in G is a path whose edges are alternating R- and B-edges, and that does not touch any vertex more than once. An *Æ-cycle* in G is a Æ-path from a vertex to itself, starting with a B-edge and ending with an R-edge.

The A and E stand for "alternating" and "elementary". The meaning of "alternating" should be clear, and the meaning of "elementary" is that the path or cycle must not cross itself.

4.23 Definition A pre-proof net π *obeys the RB-criterion* (or shortly, is *RB-correct*) iff its RB-graph G_π contains no Æ-cycle and every pair of vertices in G_π is connected via an Æ-path.

4.24 Theorem *A pre-proof net is RB-correct if and only if it is a proof net.*

Proof: We show that a pre-proof net is RB-correct iff it obeys the switching criterion, which is easy: If there are two vertices in the RB-graph not connected by an Æ-path, then there is a switching yielding a disconnected graph, and vice versa. Similarly, the RB-graph contains an Æ-cycle if and only if we can provide a switching with a cycle. □

4.25 Exercise Work out the details of the previous proof.

For the third correctness criterion, we write our nets in yet another way:

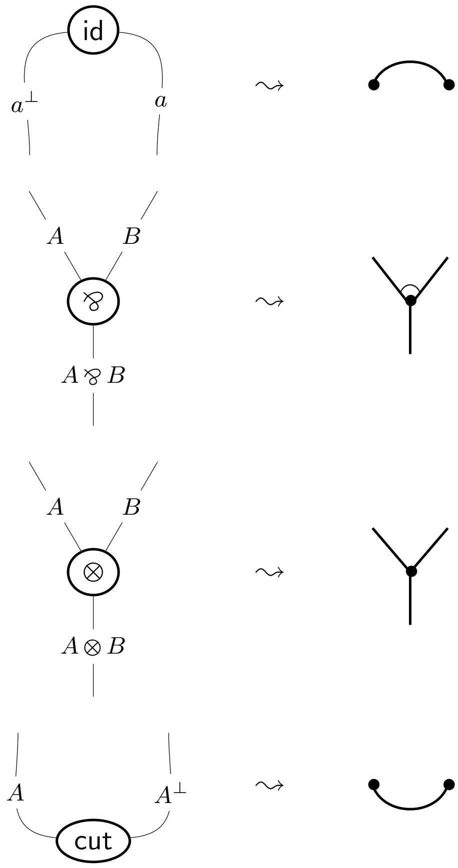

(15)

Now consider the following two rewriting rules on these graphs:

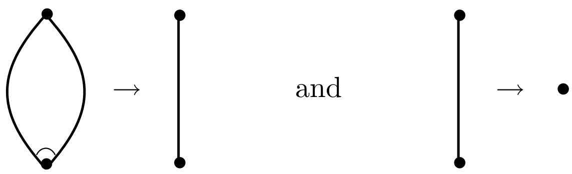

(16)

It is important to note that in the first rule the two edges are between the same pair of vertices and are connected by an arc at exactly one of the two vertices. The second rule only applies if the two vertices on the lefthand side are distinct, and the edge is not connected to another edge by an arc.

4.26 Theorem *The reduction relation induced by the rules in* (16) *is terminating and confluent.*

Proof: Termination is obvious because at each step the size of the graph is reduced. Hence, it suffices to show local confluence to get confluence. But this is easy since there are no (proper) critical pairs. □

This means that for each pre-proof net we get a uniquely defined reduced graph, and the question is now how this graph looks like.

4.27 Exercise Apply the reduction relation defined in (16) to the nets in (5) and (6).

4.28 Definition A pre-proof net *obeys the contraction criterion* if its normal form according to the reduction relation defined in (16) is

$$\bullet$$

i.e., a single vertex without edges.

At this point rather unsurprisingly, we get:

4.29 Theorem *A pre-proof net obeys the contraction criterion if and only if it is a proof net.*

Proof: As before, we show this by showing the equivalence of the switching criterion and the contraction criterion. This is easy to see since both reductions in (16) preserve and reflect correctness according to the switching criterion. □

Before we leave the topic of correctness criteria, let us make some important observations on their complexity. The naive implementation of checking the switching criterion needs exponential time: if there are n par-links in the net, then there are 2^n switchings to check. However, checking the RB-criterion needs only quadratic runtime. To verify this is an easy graph-theoretic exercise. It is also easy to see that checking the contraction criterion can be done in quadratic time. But it is rather surprising that it can be done in linear time in the size of the net. This means that (in the case of MLL^-) when we go from a formal proof in a deductive system like the sequent calculus or the calculus of structures (whose correctness can trivially be checked in linear lime in the size of the proof) to the proof net, we do not lose any information. The proof net contains the *essence* of the proof, including the "deductive information". Unfortunately, MLL^- is (so far) the only logic (except some variants of it), for which this ideal of proof nets is reached. We might come back to this in the last lecture.

5 Cut elimination

In the previous lecture you have already seen two different proofs of cut elimination: one using the sequent calculus, and one using the calculus of structures. In this section, you will see yet another one, using proof nets.

Consider the following reduction rules on pre-proof nets with cuts:

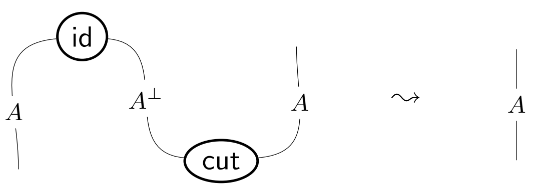

(17)

and

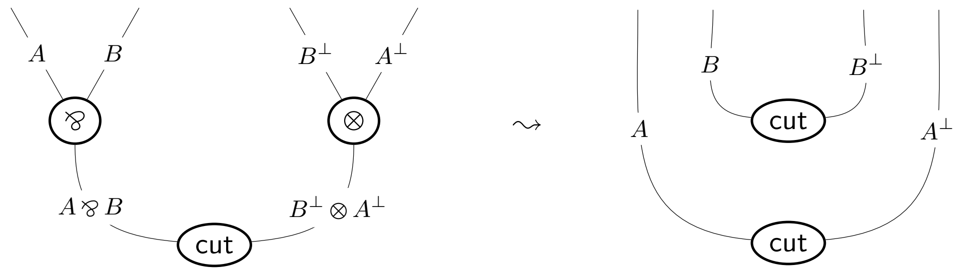

(18)

5.1 Theorem *The cut reduction relation defined by* (17) *and* (18) *terminates and is confluent.*

Proof: Showing termination is trivial because in every reduction step the size of the net decreases. For showing confluence, note that the only possibility for making a critical pair is when two cuts want to reduce with the same identity link. Then the situation must be of the shape:

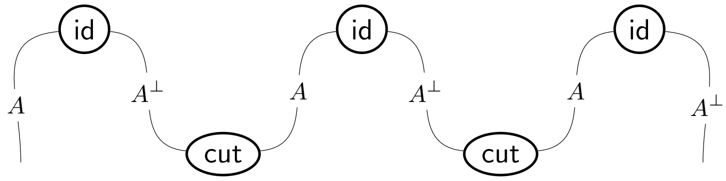

But no matter in which order and with which identity we reduce the cuts, the final result will always be

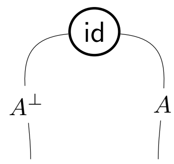

Hence we also have confluence. □

However, it could happen, that we end up in a situation like

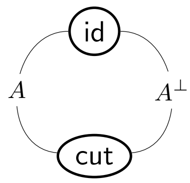

where we cannot reduce any further. That something like this cannot happen if we start out with a correct net is ensured by the following theorem, which says that the cut reduction preserves correctness.

5.2 Theorem *Let π and π' be pre-proof nets such that π reduces to π' via the reductions* (17) *and* (18)*. If π is correct, then so is π'.*

Proof: For proving this, let us use the RB-correctness criterion. Written in terms of RB-graphs, the two reduction rules look as follows:

(19)

and

(20)

That the first rule preserves RB-correctness is obvious because it just shortens an existing path. For the second rule, we proceed by way of contradiction. First, assume that the graph on the right contains an Æ-cycle, while the one on the left does not. There are three possibilities:

1. The Æ-cycle does not contain one of the new B-R-B-paths. Then the same cycle is also present on the left. Contradiction.

2. The Æ-cycle contains exactly one of the new B-R-B-paths. Then, as before, the same cycle is also present on the left. Contradiction.

3. The Æ-cycle contains both of the new B-R-B-paths. Then we can construct an Æ-cycle on the left that comes in at the upper left corner, goes down through the $\otimes$-link, and goes out at the lower left corner. Again, we get a contradiction.

That Æ-path connectedness is preserved is shown in a similar way. □

5.3 Exercise Complete the proof of Theorem 5.2, i.e., show that if we apply (20) to an RB-correct net, then in the result every pair of vertices is connected by an Æ-path. Hint 1: Note that the two rightmost vertices in (20) must be connected by an Æ-path that does not touch the new B-R-B-paths (why?). Hint 2: You will need the fact that the first net is also Æ-cycle free.

The important point of Theorem 5.2 is that it allows us to give a short proof of the cut elimination theorem for MLL^- and for MLS^-: Let Π be a proof with cuts in MLL^- given in the sequent calculus or the calculus of structures. We can translate Π into a proof net π, as described in Sections 2 and 3 and remove the cuts from the proof net as described above. This gives us a proof net π', which we can translate back to the sequent calculus or the calculus of structures. This works because removing the cuts from the proof net preserves the property of being correct (i.e., being a proof net), and translating back does not introduce any new cuts.

This raises an important question: Suppose we start out with a proof Π with cuts in MLL^- (given in sequent calculus or the calculus of structures). Now we could first remove the cuts as shown in the previous lecture for the sequent calculus and for the calculus of structures, and then translate the resulting cut-free proof Π' into a proof net π'_1. Alternatively, we could first translate Π into a proof net π, and then remove the cuts from π, to obtain the cut-free proof net π'_2. Do we get the same result? Is $\pi'_1 = \pi'_2$?

The answer is of course **yes**. To see this, note that the cut reduction steps in the sequent calculus either preserve the proof net (if the cut is just permuted up via a trivial rule permutation) or do exactly the same as the cut reduction steps for proof nets.

The same is true for the calculus of structures. The proof of the splitting lemma is designed such that it preserves the net. To make this formally precise would go beyond the scope of these lecture notes, but by comparing Figures 5 and 6 you should get the idea.

We can summarize this by the following commuting diagram:

$$
\begin{tikzcd}
\begin{array}{l}\text{proof with cuts}\\ \text{(in } \mathsf{MLL}^- \text{ or } \mathsf{MLS}^-\text{)}\end{array} \arrow[r] \arrow[d, "\begin{array}{l}\text{cut elimination}\\ \text{(in SC or CoS)}\end{array}"'] & \text{proof net with cuts} \arrow[d, "\begin{array}{l}\text{cut elimination}\\ \text{(in proof nets)}\end{array}"] \\
\begin{array}{l}\text{cut-free proof}\\ \text{(in } \mathsf{MLL}^- \text{ or } \mathsf{MLS}^-\text{)}\end{array} \arrow[r] & \text{cut-free proof net}
\end{tikzcd}
\tag{21}
$$

Our basic introduction into the theory of proof nets for unit-free multiplicative linear logic is now finished. However, a very important and fundamental question has not yet been mentioned:

5.4 Big Question Let π and π' be two proof nets such that π' is obtained from π by applying some cut reduction steps. Do π and π' represent the *same* proof?

One can safely say that in the simple case of unit-free multiplicative linear logic the answer is **yes**. However, when it comes to richer fragments of linear logic, or classical logic, the answer for this question is far from clear.