Construction of retractile proof structures
(extended version)

Roberto Maieli

Department of Mathematics and Physics, "Roma Tre" University
Largo San Leonardo Murialdo 1 – 00146 Rome, Italy
roberto.maieli@uniroma3.it

Abstract. In this work we present a paradigm of focusing proof search based on an incremental construction of *retractile* (i.e, *correct* or *sequentializable*) proof structures of the pure (units free) multiplicative and additive fragment of linear logic. The correctness of proof construction steps (or expansion steps) is ensured by means of a system of graph retraction rules; this graph rewriting system is shown to be convergent, that is, terminating and confluent. Moreover, the proposed proof construction follows an optimal (parsimonious, indeed) retraction strategy that, at each expansion step, allows to take into account (abstract) graphs that are "smaller" (w.r.t. the size) than the starting proof structures.

Keywords: linear logic, sequent calculus, focusing proofs, proof search, proof construction, proof nets, proof net retraction, graph rewriting.

1 Introduction

This work aims to make a further step towards the development of a research programme, firstly launched by Andreoli in 2001 (see [1], [2] and [3]), which points to a theoretical foundation of a *computational programming paradigm based on the construction of linear logic proofs* (LL, [8]). Naively, this paradigm relies on the following isomorphism: *proof = state* and *construction* (or *inference*) *step = state transition*. Traditionally, this paradigm is presented as an incremental (bottom-up) construction of possibly *incomplete* (i.e., open or with proper axioms) proofs of the *bipolar focusing sequent calculus* (see Sect. 2 for a brief introduction). This calculus satisfies the property that the complete (i.e., closed or with logical axioms) bipolar focusing proofs are fully representative of all closed proofs of linear logic: this correspondence is, in general, not satisfied by the polarized fragments of linear logic. Bipolarity and focusing properties ensure more compact proofs since they get rid of some irrelevant intermediate steps during proof search (or proof construction).

Now, while the view of sequent proof construction is well adapted to theorem proving, it is inadequate when we want to model some proof-theoretic intuitions behind, e.g., concurrent logic programming which requires very flexible and modular approaches. Due to their artificial sequential nature, sequent proofs are difficult to cut into modular (reusable) concurrent components.

A much more appealing solution consists of using the technology offered by *proof nets* of linear logic or, more precisely, some forms of de-sequentialized (geometrical, indeed) proof structures in which the composition operation is simply given by (possibly, constrained) juxtaposition, obeying some correctness criteria. Actually, the proof net construction, as well the proof net cut reduction, can be performed in parallel (concurrently), but despite the cut reduction, there may not exist executable (i.e., sequentializable) construction steps: in other words, construction steps must satisfy a, possibly efficient, correctness criterion. Here, a proof net is a particular "open" proof structure, called *transitory net* (see Sect. 3), that is incrementally built bottom-up by juxtaposing, via construction steps, simple proof structures or modules, called *bipoles*. Roughly, bipoles correspond to Prolog-like *methods* of Logic Programming Languages: the *head* is represented by a multiple trigger (i.e., a multiset of positive atoms) and the *body* is represented by a layer of negative connectives with negative atoms. We say that a construction step is correct (that is, a *transaction*) when it preserves, after juxtaposition, the property of being a transitory net: that is the case when the given abstract transitory structure *retracts* (after a finite sequence of rewriting steps) to an elementary collapsed graph (i.e., single node with only pending edges). Each retraction step consists of a simple (local) graph deformation or graph rewriting. The resulting rewriting system is shown to be convergent (i.e., terminating and confluent), moreover, it preserves, step by step, the property of being a transitory structure (see Theorem 1 and Lemma 1 in Sect. 3.1). Transitory nets (i.e., retractile structures) correspond to derivations of the focusing bipolar sequent calculus (Sect. 4, Theorem 2).

The first retraction algorithm for checking correctness of the proof structures of the pure multiplicative fragment of linear logic (MLL), was given by Danos in his Thesis ([6]); the complexity of this algorithm was later shown to be linear, in the size of the given proof structure, by Guerrini in [10]. Then, the retraction criterion was extended, respectively, by the author, in [14], to the pure multiplicative and additive (MALL) proof nets with *boolean weights* and then by Fouqueré and Mogbil, in [7], to polarized multiplicative and exponential proof structures.

Traditionally, concerning proof nets of linear logic, the main interest on the retraction system is oriented to study the complexity of correctness criteria or cut reduction. Here, our (original) point of view is rather to exploit retraction systems for incrementally building (correct) proof structures. Indeed, the convergence of our retraction system allows to focus on particular retraction strategies that turn out to be optimal (in the graph size) w.r.t. the problem of incrementally constructing transitory nets. Actually, checking correctness of an expanded proof structure is a task which may involve visiting (i.e., retracting) a large portion of the so obtained net: some good bound for these task would be welcome. Here, we show that checking correctness (retraction) of a MALL transitory net, after a construction attempt, is a task that can be performed by restricting to some "minimal" (i.e., already partially retracted) transitory nets. The reason is that some subgraphs of the given transitory net will not play an active role in the construction process, since they are already correct and encapsulated (i.e., border free): so, their retraction can be performed regardless of the construction process (that is the main content of Corollary 1, in Sect. 3.2).

Finally, we give in Sect. 5 a comparing with some related works concerning:

1. analogous attempts to give a theoretical foundation of computational programming paradigms based on the construction of proofs of intuitionistic or linear logic (notably, some works of Pfenning and co-authors, [4], and some works of Miller and co-authors, [15] and [5]);
2. alternative syntaxes for additive-multiplicative proof structures (mainly, those ones given, respectively, by Girard [8] and Hughes–van Glabbeek in [11]).

2 Construction of bipolar focusing proofs

In this section we give a brief presentation of the *bipolar focusing sequent calculus* introduced by Andreoli; more technical details can be found in [1]. We start with the basic notions of the *MALL fragment* of LL, without units and Mix rule. We arbitrarily assume literals $a, a^{\perp}, b, b^{\perp}, \ldots$ with a polarity: *negative* for atoms and *positive* for their duals. A *formula* is built from literals by means of the two groups of connectives:

negative : $⅋$ ("par") and $\&$ ("with");
positive : $\otimes$ ("tensor") and $\oplus$ ("plus").

A proof is then built by the following rules of the *MALL sequent calculus*:

$$\textit{identity}: \quad \frac{}{A, A^{\perp}}\, id \qquad \frac{\Gamma, A \qquad \Delta, A^{\perp}}{\Gamma, \Delta}\, cut$$

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

$$\textit{additives}: \quad \frac{\Gamma, A \qquad \Gamma, B}{\Gamma, A \& B}\, \& \qquad \frac{\Gamma, A_i}{\Gamma, A_1 \oplus_i A_2}\, \oplus_{i=1,2}$$

The *bipolar focusing sequent calculus* is a refinement of the previous one, based on the crucial properties of *focusing* and *bipolarity* (see, also, [12]). The *focusing property* states that, in proof search (or proof construction), we can build (bottom up) a sequent proof by alternating clusters of negative inferences with clusters of positive ones. As consequence of this bipolar alternation we get more compact proofs in which we get rid of the most part of the bureaucracy hidden in sequential proofs (as, for instance, irrelevant permutations of rules). Remind that, w.r.t. proof search, negative (resp., positive) connectives involve a kind of *don't care non-determinism* (resp., *true non-determinism*).

A *monopole* is a formula built on negative atoms using only the negative connectives, while a *bipole* is a formula built from monopoles and positive atoms, using only positive connectives; moreover, bipoles must contain at least one positive connective or be reduced to a positive atom, so that they are always disjoint from monopoles. Given a set $\mathcal{F}$ of bipoles, the *bipolar focusing sequent calculus* $\Sigma[\mathcal{F}]$ is a set of inferences of the form

$$\frac{\Gamma_1 \quad \dots \quad \Gamma_n}{\Gamma} B$$

where the conclusion Γ is a sequent made by a multiset of negative atoms and the premises $\Gamma_1, ..., \Gamma_n$ are obtained by fully focusing decomposition of some bipole $B \in \mathcal{F}$ in the context Γ (therefore, $\Gamma_1, ..., \Gamma_n$ are mutiset of negative atoms too). More precisely, due to the presence of additives (in particular the sum $\oplus$ connective) a bipole B is naturally associated to a set of inferences $B_1, ..., B_{m+1}$, where m is the number of $\oplus$ connectives present in B. For instance, in the purely multiplicative fragment of LL (i.e, MLL), the bipole $B = a^\perp \otimes b^\perp \otimes (c \parr d) \otimes e$, where a, b, c, d and e are (negative) atoms, yields the inference below (on the left hand side), more compact than the explicit one (on the right hand side):

$$\frac{\Gamma, c, d \qquad \Delta, e}{\Gamma, \Delta, a, b} B \quad \Leftrightarrow \quad \cfrac{\cfrac{\cfrac{\Gamma, c, d}{\Gamma, c \parr d} \parr \qquad \Delta, e}{\Gamma, \Delta, (c \parr d) \otimes e} \otimes \qquad b, b^\perp \qquad a, a^\perp}{\Gamma, \Delta, a, b, \mathbf{a}^\perp \otimes \mathbf{b}^\perp \otimes (\mathbf{c} \parr \mathbf{d}) \otimes \mathbf{e}} \otimes$$

where Γ and Δ range over a multiset of negative atoms; the identity axioms $\overline{a, a^\perp}$ and $\overline{b, b^\perp}$ are omitted in the bipolar sequent proof for simplicity. Observe, the couple a and b plays the role of a *trigger* (or *multi-focus*) of the B inference; more generally, a trigger (of a bipole) is a multi-set of duals of the positive atoms occurring in the bipole. Intuitively, the main feature of the bipolar focusing sequent calculus is that its inferences are triggered by multiple focus, like in [15] and and [5]. Bipoles are clearly inspired by the *methods* used in logic programming languages: the positive layer of a bipole corresponds to the *head*, while the negative layer corresponds to the *body* of a *Prolog-like method*.

The bipolar focusing sequent calculus, with only logical axioms (id), has been proven in [1] to be isomorphic to the focusing sequent calculus, so that (closed) proof construction can be performed indifferently in the two systems. The main idea behind this isomorphism is the *bipolarisation technique*, that is a simple procedure that allows to transform any provable formula F of the LL sequent calculus into a *set* of bipoles, called *universal program* of the bipolar sequent calculus. In Example 1 we give an instance of (closed) bipolar focusing derivation.

Example 1. Assume an universal program with four bipoles as follows:

$$\begin{aligned} B_1 &= f^\perp \otimes (x \parr g \parr h \parr (d \& e)), \\ B_2 &= x^\perp \otimes (a \& b), \\ B_3 &= g^\perp \otimes ((a^\perp \oplus b^\perp) \otimes c^\perp), \\ B_4 &= h^\perp \otimes c \otimes (d^\perp \oplus e^\perp). \end{aligned}$$

Each bipole induces a non empty set of bipolar inferences as follows:

– both bipoles B_1 and B_2 induce a single inference

$$\frac{\Gamma, x, g, h, d \qquad \Gamma, x, g, h, e}{\Gamma, f} B_1 \qquad \text{resp.,} \qquad \frac{\Gamma, a \qquad \Gamma, b}{\Gamma, x} B_2$$

– while both bipoles B_3 and B_4 induce two inferences

$$\frac{\Gamma}{\Gamma, g, a, c} B_3' \text{ and } \frac{\Gamma}{\Gamma, g, b, c} B_3'' \text{ resp., } \frac{\Gamma, c}{\Gamma, h, d} B_4' \text{ and } \frac{\Gamma, c}{\Gamma, h, e} B_4''$$

– then, the resulting bipolar focusing proof Π of f if built as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{}{\mathbf{g}, \mathbf{a}, \mathbf{c}} B_3'}{a, g, \mathbf{h}, \mathbf{d}} B_4' \qquad \dfrac{\dfrac{}{\mathbf{g}, \mathbf{b}, \mathbf{c}} B_3''}{b, g, \mathbf{h}, \mathbf{d}} B_4'}{\mathbf{x}, g, h, d} B_2 \qquad \dfrac{\dfrac{\dfrac{}{\mathbf{g}, \mathbf{a}, \mathbf{c}} B_3'}{a, g, \mathbf{h}, \mathbf{e}} B_4'' \qquad \dfrac{\dfrac{}{\mathbf{g}, \mathbf{b}, \mathbf{c}} B_3''}{b, g, \mathbf{h}, \mathbf{e}} B_4''}{\mathbf{x}, g, h, e} B_2}{\mathbf{f}} B_1$$

Although this derivation is quite compact and abstract, it still presents some structural drawbacks like duplications of some sub-proofs. Therefore, we will move, in the next section, to more flexible proof structures.

3 Bipolar transitory structures

In this section we introduce the de-sequentialized version of the bipolar focusing sequent calculus, i.e. a graphical representation of *bipolar structures* (eventually correct, i.e. *bipolar nets*) which preserves only essential sequentializations.

Definition 1 (links). *Assume an infinite set $\mathcal{L}$ of* resource places $a, b, c, \ldots$ *(also* ports *or* addresses*). A* link *consists of two disjoint sets of places,* top *and* bottom*, together with a polarity,* positive *or* negative*, and s.t. a positive link must have* at least one *bottom place, while a negative link must have* exactly one *bottom place. The* border *or* frontier *of a link is the set of its top and bottom places.*

Graphically, links are represented like in Fig. 1 and distinguished by their shape: triangular for negative and round for positive links. Top (resp., bottom) places are drawn as edges incident to a vertex. We may use variables $x^p, y^p, z^p, \ldots$ for links with a polarity $p \in \{+, -\}$, and the compact expression $link^+$ (resp., $link^-$) for a positive (resp., negative) link. Moreover, we define some relations on the set of links; in particular, given two links, x and y, we say:

– they are *adjacent* if they have (or *share*) a common place;
– x is *just above* (resp., *just below*) y if there exists a place that is both at the bottom (resp., top) of x and at the top (resp., bottom) of y;
– they are *connected* if they belong to the transitive closure of the adjacency relation.

Definition 2 (transitory structure). *A* transitory structure *(TS) is a set π of links satisfying the following conditions:*

1. *if two links are one above the other, then they have opposite polarity;*

2. *if two links have a top (resp. bottom) place in common, then they must have the same polarity;*
3. *if two negative links have a top place in common, then they must share their (unique) bottom place.*

Moreover, a TS π is called:

- bipolar *(BTS), if any place occurring at the top of some positive link of π also occurs at the bottom of some negative link of π and vice-versa (the bottom place of any negative link also occurs at the top of some positive link);*
- negative hyperlink*, if it is a set of, at least two, negative links with same bottom place;*
- positive hyperlink*, if it is a set of connected positive links;*
- bipole*, when it contains exactly one positive link; a bipole is then called* elementary *(or* multiplicative*) when it does not contain any negative hyperlink.*

Finally, in a TS π, the set of bottom (resp., top) places that do not occur at the top (resp., bottom) of any link of π is called the bottom *(resp.,* top*)* border *or* frontier *of π. If the top border of π is empty, then π is called* closed*. A place shared by at least two links of the same polarity is called (*additive*)* multiport*.*

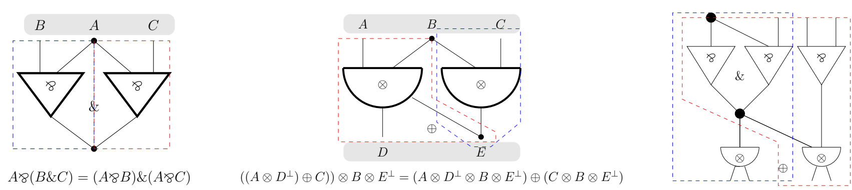

Fig. 1. links, hyperlinks, bipoles and bipolar transitory structures

Intuitively, w.r.t. the standard syntax of proof nets of linear logic, negative (resp., positive) links correspond to generalized (i.e., n-ary) ⅋-links (resp., $\otimes$-links). Similarly, negative (resp., positive) hyperlinks correspond, modulo distributivity and associativity of linear connectives, to generalized & (resp., $\oplus$) of negative (resp., positive) links. Instances of negative and positive hyperlinks are, respectively, given in the leftmost and middle side pictures of Fig. 1, where links are enclosed within dashed lines; graphically, these hyperlinks represent the distributive law of negative (⅋/&), respectively, positive ($\otimes/\oplus$) connectives. An instance of BTS is also given in the rightmost picture of Fig. 1, with two bipoles enclosed within dashed lines (bullets, $\bullet$, graphically represent multiports). Intuitively, bipoles correspond to bipolar inferences of the sequent calculus.

3.1 Retraction of bipolar transitory structures

We are interested in those BTSs that *correspond* to bipolar focusing sequent proofs: these *correct* BTSs will be called *bipolar transitory nets* (BTN). In the following we will give a geometrical way to characterize BTNs: actually we will show that BTNs are those BTSs whose *abstract structures retract*, by means of sequences of rewriting rules (graph deformation steps), to special terminal *collapsed* graphs . This retraction technique was primarily exploited by Danos in his thesis ([6]), limited to the multiplicative proof structures (see *rules* R_1, R_2 *and* R_3 of Definition 5) and then extended by Maieli in [14] to the multiplicative and additive proof structures. The latter work provides a binary version of *rules* R_5 *and* R_6 of Definition 5 that only works with closed proof structures labeled by *boolean monomial weights* (see [9]). Here, we further extend these techniques, by generalizing the rules above, to weightless proof structures that are focusing, bipolar, possibly open and with n-ary links.

Definition 3 (abstract structure). *An* abstract structure *(*AS*) is a undirected graph* π^* *equipped with a set* $\mathcal{C}(\pi^*)$ *of pairs of coincident edges: two edges are* coincident *if they share at least a vertex, called* base *of the pair. Each pair has a* type $\alpha \in \{\parr, \&, C\}$ *(where* C *denotes the* additive contraction*). We call* cluster *of type* α *a tuple of edges that are pairwise pairs of* $\mathcal{C}(\pi^*)$ *with type* α*. A pair (resp., cluster) is graphically denoted by a crossing* arc *close to the base. Some* pending edges *(i.e., edges that are incident to only one node) of an AS are called* conclusions *(resp.* hypotheses*) of the AS. We call* collapsed *any acyclic AS* π^* *with at most a single node and* $\mathcal{C}(\pi^*) = \emptyset$*.*

Notation: a dashed edge incident to a vertex v is a compact representation of possibly several edges (with possibly clusters) incident to v; variables $a, b, c, ...$ denote (dashed) edges; possibly partially dotted arcs with labels $\alpha \in \{\parr, \&, C\}$ are compact representations of pairs (clusters) of type α; vertices may be denoted by naturals inside (dotted) circles ①, ②, ... A cluster of n edges, $a_1, ..., a_n$, with type α, is denoted by $\alpha(a_1, ..., a_n)$ (sometimes, simply α_n).

Definition 4 (abstraction). *We may transform (*abstract*) a given BTS* π*, with bottom border* Γ *and top border* Δ*, in to an AS* π^* *(also* abstraction of π*) with conclusions* Γ *and hypothesis* Δ*, built by applying the following* procedure*:*

1. *a link*$^+$ *with border* $a_1, ..., a_n$ *becomes a vertex with incident edges* $a_1, ..., a_n$*;*
2. *a link*$^-$ *with top places* $a_1, ..., a_n$ *and bottom place* b *becomes a vertex that is base for a cluster* $\parr(a_1, ..., a_n)$ *and with* b *as an additional incident edge;*
3. *a place (multiport)* a *that is bottom (resp., top) place of* n *links*$^-$ *becomes a vertex that it is base of a cluster* $\&(a_1, ..., a_n)$ *(resp.,* $C(a_1, ..., a_n)$*) with* n *copies of* a*, and with an additional incident edge labeled by* a*;*
4. *a place (multiport)* a *that is top (resp., bottom) place of* n *links*$^+$ *becomes a vertex that is base of a cluster* $C(a_1, ..., a_n)$*, with* n *copies of* a*, and with an additional incident edge labeled by* a*;*
5. *we may compact* π^* *by some applications of structural retractions* R_1*,* R_2*.*

Definition 5 (retraction system). *Given an AS* π^**, a* retraction step *is a replacement (also,* deformation *or* rewriting*) of a subgraph* S *(called,* retraction graph*) of* π^* *with a new graph* S' *(called,* retracted graph*), leading to* π'^* *according to one of the following rules* $R_1, ..., R_9$*.*

R_1 **(structural):** *with the condition that, like in Fig. 2, the retraction graph of π^* contains a vertex ① with only two incident edges, a and b, none of them pending; then, this graph is replaced in π'^* by a single new edge c s.t. any pair of $\mathcal{C}(\pi^*)$ containing a or b is replaced in $\mathcal{C}(\pi'^*)$ by a pair of the same type and with c at the place of a or b.*

R_2 **(structural):** *with the condition that, like in Fig. 2, the retraction graph of π^* contains two distinct vertices ① and ② with a common edge c not occurring in any pair of $\mathcal{C}(\pi^*)$; then, one of these two nodes, ① or ②, together with the edge c, does not occur in π'^*; moreover, $\mathcal{C}(\pi^*) = \mathcal{C}(\pi'^*)$.*

R_3 **(multiplicative):** [1] *with the condition that, w.r.t. the retraction graph of π^* in Fig. 2, all vertices are distinct and there exists a cluster $\parr(a_1, ..., a_n)$, with base in ①, whose edges, a_{n-1} and a_n are also incident to vertex ②; moreover, a_{n-1} and a_n do not occur in any pair, except the cluster $\parr(a_1, ..., a_n)$. Then, π'^* (resp. $\mathcal{C}(\pi'^*)$) is obtained from π^* (resp. from $\mathcal{C}(\pi^*)$) by erasing a_n (resp., by replacing $\parr(a_1, ..., a_n)$ with $\parr(a_1, ..., a_{n-1})$).*

R_4 **(associative):** [2] *with the conditions that, w.r.t. the retraction graph of π^* in the Fig. 2 (all vertices are distinct):*

1. *vertex ① is a base for the cluster $\alpha(a_1, ..., a_n)$;*
2. *vertex ② is a base for the cluster $\alpha(b_1, ..., b_m)$;*
3. *$\alpha \in \{\parr, \&, C\}$ and $n, m \geq 2$;*
4. *the only edges incident to the vertex ② are $b_1, ..., b_m, a_n$.*

Then, the edge a_n (resp., vertex ②) does not occur in π'^* and both clusters, $\alpha(a_1, ..., a_n)$ and $\alpha(b_1, ..., b_m)$ of $\mathcal{C}(\pi^*)$, are replaced in $\mathcal{C}(\pi'^*)$ by an unique cluster $\alpha(a_1, ..., a_{n-1}, b_1, ..., b_m)$ with base in vertex ①.

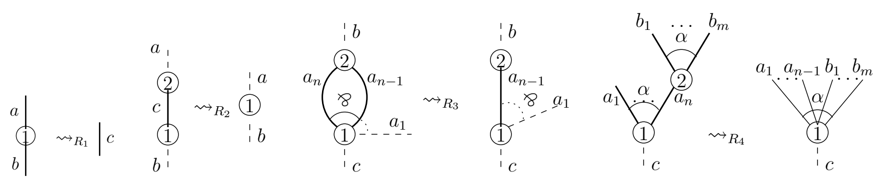

Fig. 2. *structurals* (R_1, R_2), *multiplicative* (R_3) and *associative* (R_4) retractions

R_5 **(distributive):** [3] *with the condition that, w.r.t. the retraction graph of π^* in Fig. 3, all vertices are distinct and each vertex v_i ($1 \leq i \leq n$) has only a_i, b_i and c_i ($1 \leq i \leq n$) as incident edges with the following conditions:*

1. *c_i is an edge occurring in the cluster $\&(c_1, ..., c_n, d)$ with base in v_h;*
2. *b_i is an edge occurring both in the cluster $C(b_1, ..., b_n)$, with base in vertex v_k, and in the cluster $\parr_i(a_i, b_i)$, with base in vertex v_i;*
3. *a_i is a non empty bundle of edges occurring in the cluster $\parr(a_i, b_i)$; moreover, each edge $e \in a_i$ must satisfy one of the following conditions:*

---

[1] Intuitively, this rule corresponds to the replacement of an axiom by its *η-expansion*.
[2] Intuitively, this rule corresponds to the *associativity* of, respectively, $\parr, \&$ and C.
[3] A reminiscence of the *distributivity* $(\&_{i=1}^n (a_i \parr f))\&d \dashv\vdash (\&_{i=1}^n (a_i) \parr f)\&d$ (see [14]).

(a) either it is a pending edge or an edge incident to a vertex with only pending edges not labeled by any conclusion of π^*; in that case, there must exist at least such an analogous edge for each bundle $a_1, ..., a_n$;

(b) or it must occur in a C cluster and, in that case, for each bundle $a_1, ..., a_n$, there must exist exactly one edge that occurs in this C cluster too.

Then, π^* retracts to π'^* like in Fig. 3. Observe that edges $b_1, ..., b_n$, except one, b_i, do not occur in π'^*; similarly, the cluster $C(b_1, ..., b_n) \notin \mathcal{C}(\pi'^*)$. Moreover, new edges g and e are added to π'^* (similarly, new pairs, $\parr(b_i, g)$ and $\&(d, e)$ occur in $\mathcal{C}(\pi'^*)$ with base, respectively, in the new vertex $v_{h'}$ and $v_{h''}$).

R_6 **(semi-distributive):** [4] *with the condition that, w.r.t. the retraction graph of π^* in Fig. 3, all vertices are distinct and each vertex v_i, with $1 \leq i \leq n$, has only a_i, b_i and c_i $(1 \leq i \leq n)$ as incident edges with the following conditions:*

1. *c_i is an edge occurring in the cluster $\&(c_1, ..., c_n, d)$ with base in v_h;*
2. *b_i is an edge occurring in the cluster $C(b_1, ..., b_n)$ with base in v_k;*
3. *a_i is, possibly, a bundle of edges occurring neither in a pair with b_i nor in a pair containing c_i.*

Then, π^* retracts to π'^* like in Fig. 3. Observe, π'^* does not contain any $b_1, ..., b_n$ except one, b_i, (resp., $C(b_1, ..., b_n) \notin \mathcal{C}(\pi'^*)$). Finally, in π'^* we add a new edge g and a new vertex $v_{h'}$ (resp., a, possibly, new cluster $\&(d, g') \in \mathcal{C}(\pi'^*)$ with base $v_{h'}$).

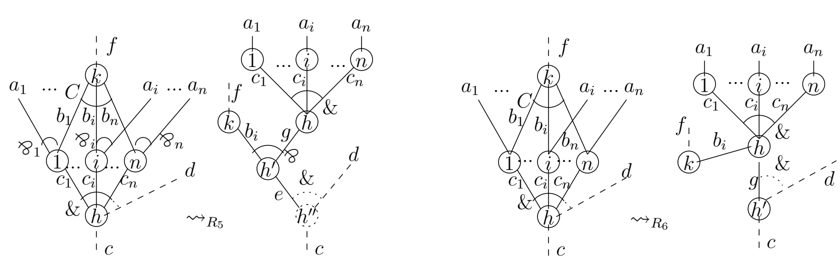

Fig. 3. retraction rules: *distributive* (R_5) and *semi-distributive* (R_6)

R_7 **($\&$-annihilation):** *with the conditions that, w.r.t. the retraction graph of π^* in Fig. 4, all vertices are distinct and each a_i, with $1 \leq i \leq n$, is an edge occurring in the cluster $\&(a_1, ..., a_n)$; moreover, each a_i must belong to a collapsed subgraph of π^* non containing conclusions of π^*, with the condition that any couple a_i, a_j $(1 \leq i, j \leq n)$ cannot belong to the same collapsed graph. Then, in π'^*, each a_i will be disconnected from v_h (so, $\&(a_1, ..., a_n) \notin \mathcal{C}(\pi'^*)$).*

R_8 **($\parr$-annihilation):** *with the condition that, w.r.t. the retraction graph of π^* in Fig. 4, all vertices are distinct and edges $a_1, ..., a_n$ occur in a cluster $\parr(a_1, ..., a_n)$; then, π^* retracts to π'^*, like in Fig. 4, whenever d is:*

[4] Reminiscence of the *semi-distributivity* $(\&_{i=1}^n (f \otimes a_i))\&d \dashv (f \otimes (\&_{i=1}^n (a_i))\&d$ ([14]).

1. either a bundle of pending edges not labeled by any conclusion of π^* and not occurring in a pair with any a_i;

2. or a bundle of pending edges not occurring in any pair with any a_i and e is also a bundle of pending edges with at least one of them labeled by a conclusion of π^* and none of them occurring in a pair with any a_i.

Then, in π'^* the edge a_n will be disconnected form vertex ①; therefore, $\mathcal{C}(\pi'^*)$ will contain all the pairs of $\mathcal{C}(\pi^*)$ except those one containing a_n.

R_9 **(merge):** *with the condition that, w.r.t. the retraction graph of π^* in Fig. 4, all vertices are distinct and χ_1^* and χ_2^* are both collapsed AS made, resp., by a vertex ① and a vertex ②, with, resp., only pending edges $a_1, ..., a_{n\geq 1}$ and $b_1, ..., b_{m\geq 1}$, with b_m that is neither a conclusion nor an hypothesis of π^*. Then, π'^* is obtained by gluing χ_1^* with χ_2^* and erasing ② and b_m.*

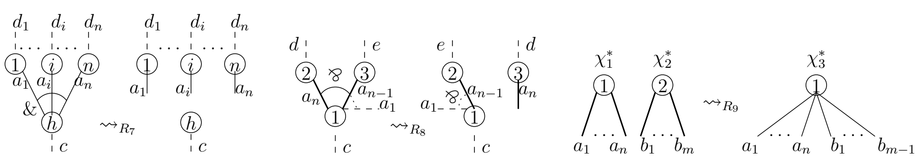

Fig. 4. retraction rules: *annihilations* (R_7 and R_8) and *merge* (R_9)

We say that π^* *retracts to* π'^* when there exists a non empty finite sequence of retraction steps starting at π^* and terminating at π'^*; then, we say that π^* is *retractile* when there exists a $\sigma^* \neq \pi^*$ s.t. π^* retracts to σ^*. A non retractile AS is called *terminal.* A sequence of retraction steps is said *complete* when it ends with a terminal AS. An AS *collapes* when it retracts to a collapsed graph. A pair of possible (or admissible) retraction instances for π^*, R_i and R_j, with $i \neq j$, is called a *critical pair* (denoted by $R_i|R_j$) when the application of R_i inhibits the application of R_j (or vice-versa).

Theorem 1 (convergence of retraction). *If π^* is an AS with conclusions Γ and hypothesis Δ then, any complete retraction sequence starting at π^* ends with a terminal AS χ^*; moreover, if χ^* is collapsed, then any complete retraction sequence starting at π^* ends with χ^*.*

Proof. Termination is proved by (lexicographic) induction on the *complexity degree* of π^*, that is, a triple $\langle \sharp P, \sharp N, \sharp E\rangle$, where "$\sharp P$", "$\sharp N$" and "$\sharp E$" denotes the number of, respectively, pairs, nodes and edges of π^*.

For the confluence, we reason, analogously, by induction on the complexity degree of the starting π^*. The crucial point is to show that for each critical pair, $R_5|R_5, R_5|R_8$ and $R_8|R_8$, we can find, in a few steps, an almost local confluence strategy that allows to apply the hypothesis of induction.

Assume π^* collapses by a retraction sequence S. For simplicity reasons, we use natural numbers to denote the border of the retraction (resp. retracted) graph.

$R_5|R_5$. Assume S starts with an instance of R_5, leading to π_1^*; moreover, assume π^* may also retracts, by a different instance of R_5, to $\pi_2^* \neq \pi_1^*$. We reason by induction on the complexity degree of π^*. Then the proof follows by showing a local confluence by means of two instances of rule R_6, like in Figure 5.

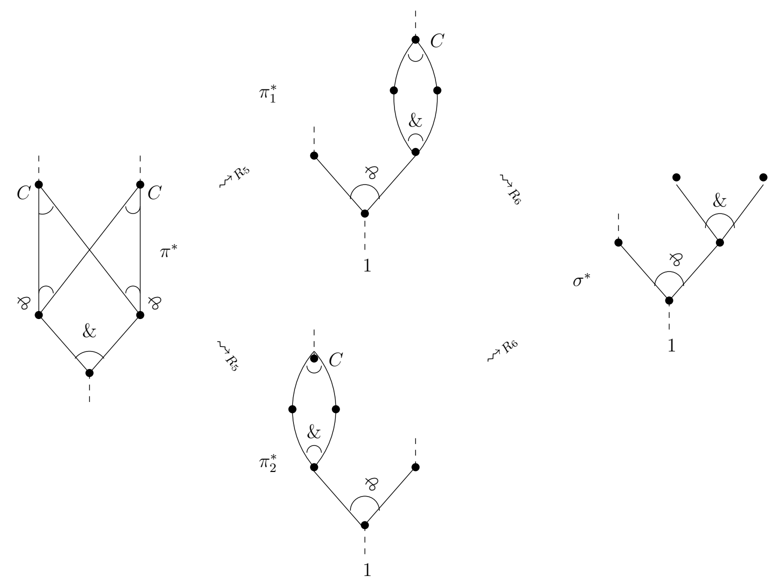

Fig. 5. confluence of a *critical pair* $R_5|R_5$

$R_5|R_8$. Assume S starts with an instance of R_5, leading to π_1^*; moreover, assume π^* may also retract, by an other admissible instance of $⅋$-annihilation R_8 to $\sigma_1^* \neq \pi_1^*$. By definition of R_5 (case 3a), the only admissible case of rule R_8 is case 1. So, we reason by induction on the complexity degree of π^*. Then the proof follows by showing a local confluence, like in Figure 6.

$R_8|R_8$. Assume S starts with an instance of R_8, leading to π_1^*; moreover, assume π^* also retracts, by an other admissible instance of R_8, to $\pi_2^* \neq \pi_1^*$, like in Figure 7. Now observe that we can apply to both π_1^* and π_2^* the same retraction rules, modulo some irrelevant labeling, and finally converge to the same correct by a last application of the merge rule R_9.

Observe that, for the way distributive and semi-distributive retraction rules are defined, it never occurs the case of a non critical pair (R_5, R_6) like in Figure 8, where the C-cluster with base in vertex 2 is partially shared by the two retraction rules R_5 and R_6. Indeed, it never occurs the case, except the multiplicative retraction, that a cluster is partially shared by the retraction graphs of two retraction rules.

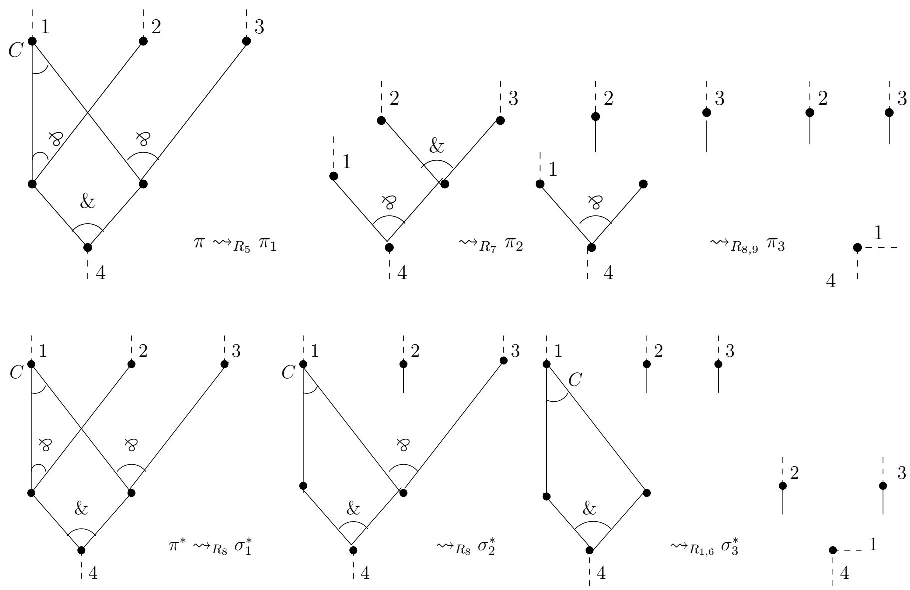

Fig. 6. confluence of a *critical pair* $R_5|R_8$

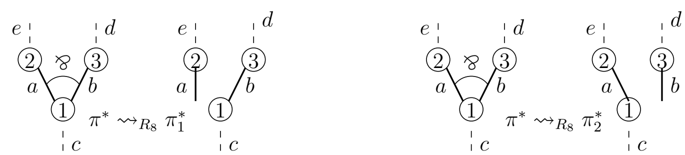

Fig. 7. confluence of a *critical pair* $R_8|R_8$

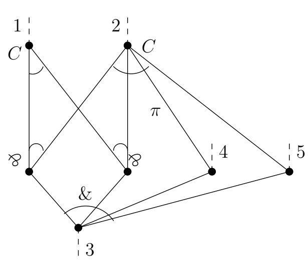

Fig. 8. a pseudo non-critical pair $R_5|R_6$

Next Lemma 1 intuitively says that abstraction commutes under retraction; it will play a crucial role in the sequentialization of BTSs (Theorem 2, Sect. 4).

Lemma 1 (abstraction). *Assume π^* is an AS that retracts to π'^* by an instance of R_i $(i = 1, ..., 9)$ and assume there exists a BTS π that abstracts to π^*; then, we can find a BTS π' whose abstraction is π'^*.*

Proof. For each case $R_{i=1,...,9}$, we show how to locally deform some bipoles of the given BTS π in such a way to get a BTS π' whose abstraction is π'^*.

1. If $i = 1$ (resp., $i = 2$), then the proof follows simply by the definition of the abstraction procedure (point 5 of Definition 4).
2. If $i = 3$, then assume π abstracts to a π^* which retracts to π'^*, by an instance of multiplicative retraction R_3 following the notation given in Figure 2. Assume π is like in the left hand side picture of Figure 9, with bipoles β_1 and β_2 enclosed within dotted lines; then we can build a BTS π' simply by replacing in π bipoles β_1 and β_2 with bipoles β'_1 and β'_2 like in the contiguous picture. Clearly, π' abstracts to the π'^* as given in Figure 2, with the bundle c, possibly, partitioned in two sub sets, c' (on the top) and c'' (on the bottom).

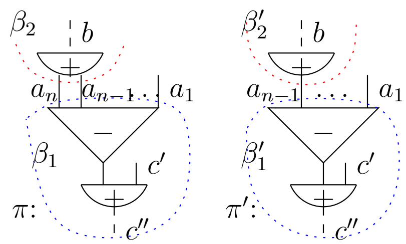

Fig. 9. abstraction under (multiplicative) retraction R_3

3. If $i = 4$, then assume π abstracts to a π^* which retracts, by an instance of associative retraction R_4, to π'^* following the notation given in Figure 2; then we distinguish two cases, according to the type of the cluster, α:
   (a) if $\alpha = \parr$ (resp., $\alpha = \&$), assume π appears like in the *Picture case (a)* of Figure 10. Then, we can build a BTS π' simply by replacing bipoles β_1 and β_2 with the unique bipole β'_1. Clearly, π' abstracts to π'^* like in the retraction rule R_4 of Figure 2.
   (b) otherwise, $\alpha = C$, then we reason like before except for the reference to the *Picture case (b)* of Figure 10. We could also assume a partitioning of c in to a cluster $\parr(c')$ and a bundle of edges c'', respectively, on the top and on the bottom of β_1 and β'_1. Observe that, in this case a partial renaming (a partial substitution) of the additive multiport, corresponding to the cluster $C(a_1, ..., a_{n-1}, b_1, ..., b_m)$, will be necessary in the BTS π', in such

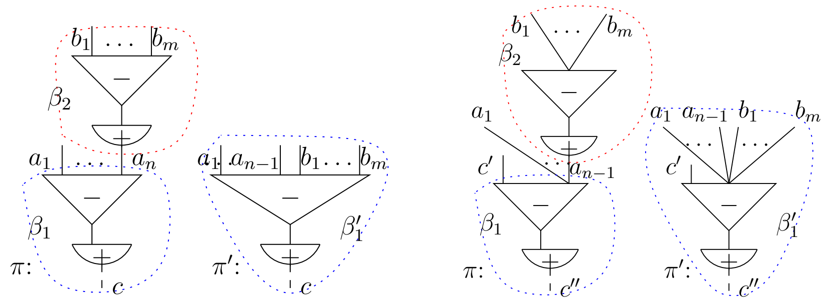

Picture case (a) Picture case (b)

Fig. 10. abstraction under associative retractions R_4

a way that the tuple $(a_1, ..., a_{n-1}, b_1, ..., b_m)$ would represent different copies of the same address (remind, by Definition 1, a link consists of two disjoint sets of ports).

Remark : observe that a positive (resp., negative) link with an unique top and bottom place plays the role of a ”polarity inverter”, widely used in the most part of polarized fragments of LL.

4. If $i = 5$ (resp., $i = 6$), then we reason like in the previous case, except for:
   (a) the reference to the notation of the retraction rule R_5 (resp., R_6) as given in Figure 3;
   (b) the abstraction of both π and π' as given in the two contigous pictures of Figure 11 (resp., Figure 12).

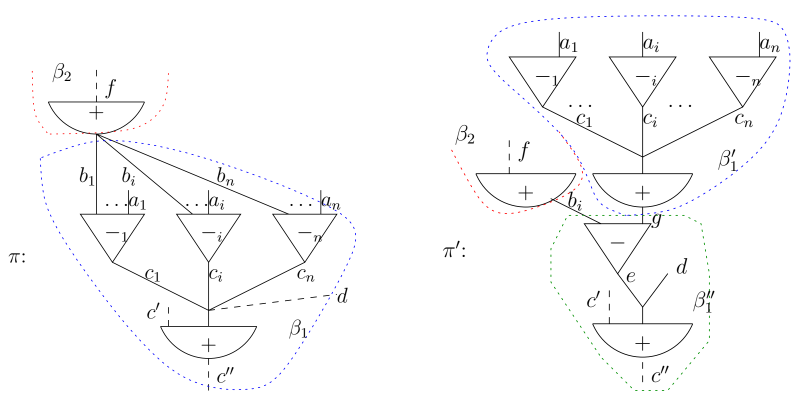

Fig. 11. abstraction under retraction R_5

5. Cases $i = 7, 8, 9$ are similar. We only mention the case $i = 8$. Assume $\pi^* \rightsquigarrow_{R_8} \pi'^*$, w.r.t. the notation given in Figure 4, and assume π is like on left

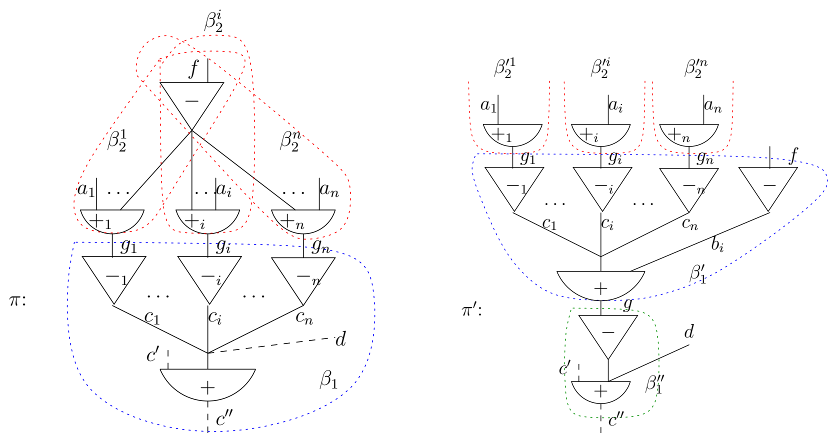

Fig. 12. abstraction under retraction R_6

hand side picture of Figure 13. Then, we can build a π', by simply replacing in π bipoles $\beta_1, \beta_2, \beta_3$ with bipoles β'_1, β_2 and β_3, like on the right hand side picture of Figure 13. Clearly, π' abstracts to the retracted AS π'^* as given in the retraction rule R_8 of Figure 4.

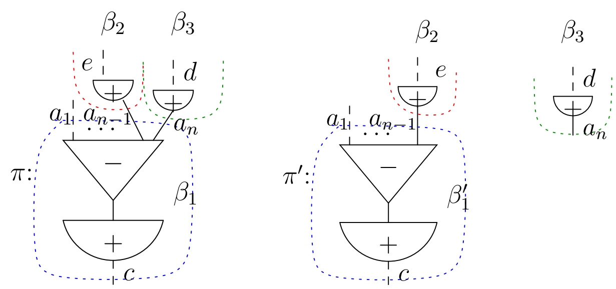

Fig. 13. abstraction under retraction R_8

Definition 6 (bipolar transitory net). *A BTS π with bottom border Γ and top border Δ, is* correct*, that is a* bipolar transitory net *(BTN), when its abstraction π^*, with conclusions Γ and hypothesis Δ,* collapses*.*

Example 2. In Fig. 14 we give an instance of (closed) BTS π (*Pic.* A_1) obtained by juxtaposing bipoles $\beta_1, \beta_2, \beta'_3, \beta''_3, \beta'_4$ and β''_4. Observe, π is correct (it is a BTN) since its abstraction π^* (*Pic.* A_2) collapses after few retraction steps:

1. first we get the AS of *Pic.* A_3 after some instances of distributive retraction R_5 applied to the dotted retraction graph of *Pic.* A_2;
2. then we get the AS of *Pic.* A_4 after a couple of instances of semi-distributive retractions R_6 applied to the dotted retraction graphs of *Pic.* A_3;
3. finally, we get the collapsed graph after three multiplicative retractions instances R_3 applied to the dotted retraction graphs of *Pic.* A_4 (modulo some structural retractions).

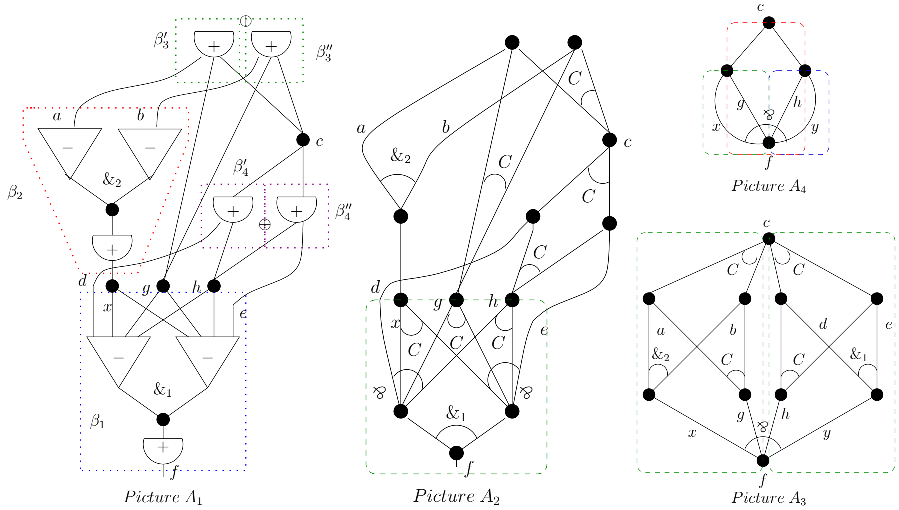

Fig. 14. bipolar net (*Pic.* A_1) with its collapsing abstractions (*Pics.* $A_{2,3,4}$)

3.2 Construction of transitory nets *via* optimal retraction

Analogously to the construction of bipolar focusing sequent proof seen in Sect. 2, in the construction of BTNs, places are decorated by type information, that is, occurrences of negative atoms. A bipole β is viewed as an *agent* which continuously attempts to perform a bottom-up *expansion step* of the given BTN π: this step consists of adding (by a gluing operation "$\star$") a non empty cluster (a sum, indeed) of bipoles from the top border places whose types match the trigger, i.e. the bottom places, of the given bipoles. Not all construction steps are admissible. We will only consider those ones that preserve correctness by retraction. Now, checking correctness of an expansion is a task which, *a-priori*, repeatedly involves *visiting* (i.e., retracting) the whole portion of the expanded BTS. Actually, we could avoid, at each construction step, considering the whole structure built up, by e.g. taking advantage of the incremental construction in such a way to reduce the complexity of the contraction task. That is exactly the content of the next

Corollary 1, immediate consequence of the Convergence Theorem 1. Intuitively, Corollary 1 allows us to incrementally pursue an optimal retraction strategy that manages, when they exist, abstract correction graphs that are strictly smaller (w.r.t. the complexity degree) than the starting ones.

Corollary 1 (optimal retraction). *Let π be a BTN (with a non empty top border) and let β a non empty cluster (a sum) of bipoles, whose bottom border matches some places of the top border of π. Assume π abstracts to π^* and assume η^* is the AS, which π^* retracts to, by only applying those retraction instances whose retraction graph does not contain pending (border) edges. Then, $(\pi \star \beta)^*$ collapses iff $\eta^* \star \beta^*$ collapses too.*

Example 3. We graphically show a reason why Corollary 1 "delays" those retractions that involve the border of the abstraction associated to the BTN to be expanded. Actually, assume π abstracts to an AS π^* with hypothesis a, b, d, e, f and conclusion c, like in the graph below the dotted line of *Pic.* B_1 in Fig. 15. Assume π'^* is obtained from π^* by an instance of distributivity R_5 applied to the couple of ⅋-pairs with, respectively, base ② and ③, like in the graph below the dotted line of *Pic.* B_2: this retraction involves the border a, b and d. Now, if we expand π^* by the (abstract) sum of bipoles $(\beta_1 \oplus \beta_2)^*$, through the border d, e, we get the AS $\pi^* \star (\beta_1 \oplus \beta_2)^*$ (the whole *Pic.* B_1) whose retraction does not collapse[5] while the expanded AS $\pi'^* \star (\beta_1 \oplus \beta_2)^*$ (the whole *Pic.* B_2) collapses.

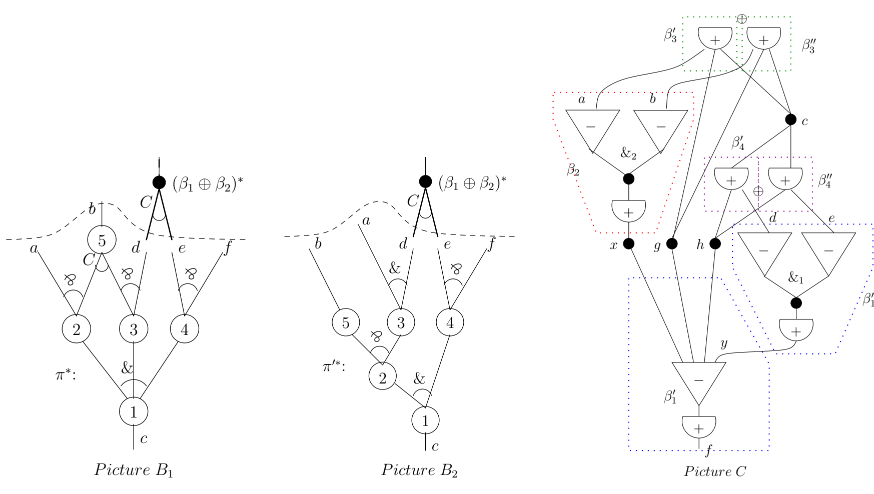

Fig. 15. expansion steps (*Pics.* B_1, B_2) and a BTS (*Pic.* C)

[5] It is no longer possible to apply rule R_5 since *condition 3b* of Definition 5 is violated.

4 Sequentialization of bipolar transitory nets

In this section we show that correct BTSs correspond (sequentialize) to proofs of the bipolar focusing sequent calculus and vice-versa.

There exists an almost direct *correspondence* (modulo associativity and distributivity of linear connectives) between a sequential bipole B and a an additive sum of bipoles $\mathcal{B} = \{\beta_1 \oplus \cdots \oplus \beta_{n \geq 1}\}$, as follows:

1. the positive layer of B corresponds to the positive hyperlink made by the positive links of $\mathcal{B}$ connected through the border (see Definition 1);
2. the negative layer of B corresponds to the set of negative hyperlinks of $\mathcal{B}$;
3. the negative literals (i.e., atoms) of B correspond to the top places of $\mathcal{B}$ while (the dual of) the positive literals of B correspond to the bottom places of $\mathcal{B}$;
4. each bipole β_i corresponds to the i-th bipolar inference induced by the sequential bipole B (see Example 1 in Sect. 2).

In general, ports (resp. multiports) correspond to a single (resp., multiple) occurrence of literals. Then, we say that a bipolar sequential proof Π with hypothesis Δ and conclusions Γ *de-sequentializes to* π, when π is a BTN with top border Δ (resp., bottom border Γ) and each instance of the i-th bipolar inference induced by $B \in \Pi$ corresponds to a bipole $\beta_i \in \pi$. The other way round, from BTNs to bipolar sequential proofs, is called *sequentialization*.

Theorem 2 ((de-)sequentialization). *A sequent proof Π, with conclusions $\Gamma = c_1, ..., c_n$ and hypothesis $\Delta = d_1, ..., d_m$),* de-sequentializes *in to a BTN π, with bottom places Γ and top places Δ.*

Proof. By induction on the number of bipolar inferences of Π^{Δ}_{Γ}. Assume the last bipole of Π is B, like in the derivation of *Picture A* in Figure 16. By the induction hypothesis, we know that each $\Pi^{\Delta_i}_{\Gamma_i}$, with $1 \leq i, j \leq n$, de-sequentializes to a BTN π_i whose corresponding AS π_i^* collapses in to χ_i^*. Moreover, we know that, modulo associativity and distributivity of linear connectives, the bipolar inference B of Π corresponds to the bipole β in the *picture B* of Figure 16. Clearly β abstracts to the obvious β^*. So, we may glue $\chi_1^*, ..., \chi_m^*$ with β^* through the common border and, finally, we may retract the resulting AS until the collapsed graph with pending edges Γ and Δ, like in *Picture C* of Figure 16.

Remark 1. In the proof of the de-sequentialization part we took the same "minimal" decision made by Girard in [9], that is: given a sequential proof with an additive inference

$$\Pi : \dfrac{\begin{array}{c}\Pi_1 \\ \Gamma, A\end{array} \qquad \begin{array}{c}\Pi_2 \\ \Gamma, B\end{array}}{\Gamma, A\&B}\ \&$$

> "*how do we know that a formula or a link X of Π_1 is the same as another formula or link Y of Π_2 ? There is no simple answer* [...] *there is at least the possibility to decide that no identification between Π_1 and Π_2 is made,*

$$\Pi_{\Delta}^{\Gamma}\colon \frac{\Pi_{\Gamma_1}^{\Delta_1}\ \ldots\ \Pi_{\Gamma_i}^{\Delta_i}\ \ldots\ \Pi_{\Gamma_j}^{\Delta_j}\ \ldots\ \Pi_{\Gamma_n}^{\Delta_n}}{\Gamma = \Gamma', \Gamma''}\ B$$

Picture A

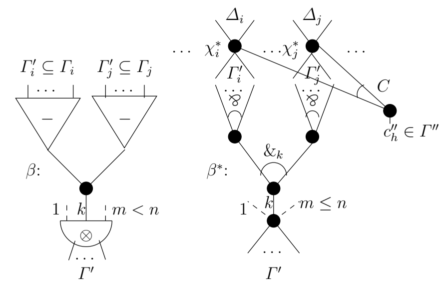

Picture B

Picture C

Fig. 16. a de-sequentialization step

but for the conclusions, i.e. the formulas of Γ; by the way the sequent calculus formulation of the &-rule stipulates that the contexts of the two premises must be equal, hence this is a clear case where there is no doubt as to the identification between a formula in Π_1 and a formula in Π_2." (Girard, [9], pp. 7-8).

The other "maximal" decision might be to follow Hughes and van Glabbeek in [11]: they, roughly, identify everything in a proof structure that is immediately below the axioms links; in other words, they, move all the additive contractions immediately below the axiom links (in a canonical form). This seems an elegant solution for representing closed proof structures but not for juxtaposing open modules with, a-priori, "non canonical" additive contractions (see also Section 5).

"*Anyway, the main problem is to find a sequentialisation theorem ; this means to give an intrinsic characterization of sequentialisable proof-structures.*" (Girard, [9], page 8).

Theorem 3 (sequentialization). *A BTN π, with bottom places Γ and top places Δ,* sequentializes *in to a bipolar sequential proof Π, with conclusions Γ and hypothesis Δ, in such a way that each bipole β of π corresponds to a bipolar inference B in Π.*

Proof. It is given by induction on the complexity degree of the abstraction π^* corresponding to the given π. By the the stability of abstraction under retraction (Lemma 1), we show that at each retraction step $\pi^* \rightsquigarrow_{R_i} \pi'^*$, for $i = 1, ..., 9$, it is possible to recover a BTN π' from the retracted ATS π'^*, with same border. Then, by hypothesis of induction, π' sequentializes to a proof $\Pi'^{\Delta'}_{\Gamma'}$ from which, finally, by deformations of Π', that is, permutations of some bipolar inferences of Π', we get a sequential proof Π_{Γ}^{Δ}.

We reason by cases, according to R_i, with $1 \leq i \leq 9$. We only show few interesting cases (the other ones are similar and so omitted).

Case $i = 6$ **(semi-distributive)** . Assume $\pi^* \rightsquigarrow_{R_6} \pi'^*$, then, by Lemma 1, we know how to build a BTN π' that abstracts to π'^*; so, assume π and π' are like in Figure 12. Then, by hypothesis of induction, π' sequentializes to the following bipolar focusing sequential proof Π':

$$\Pi' : \dfrac{\begin{array}{c}\Sigma_1\\ \vdots\\ \Gamma_1, c'\end{array} \qquad \dfrac{\cdots \quad \begin{array}{c}\Sigma'_{2,i}\\ \vdots\\ \dfrac{\Gamma'_2, a_i}{\Gamma'_2, g_i}\,\beta'^{i}_2\end{array} \quad \cdots \quad \begin{array}{c}\Sigma''_2\\ \vdots\\ \Gamma''_2, f\end{array}}{\Gamma_2, g}\,\beta'_1 \qquad \begin{array}{c}\Sigma_2\\ \vdots\\ \Gamma_2, d\end{array}}{\begin{array}{c}\Gamma, c''\\ \vdots\end{array}}\,\beta''_1$$

Now, simply, by permuting n-copies of the sub-proof Σ''_2 of Π' we can build the proof Π below, according to the notation of Figure 12 (observe: bipoles β'^{i}_2, β'_1 and β''_1 are replaced by bipoles β^i_2 and β_1, for $1 \leq i \leq n$).

$$\Pi : \dfrac{\begin{array}{c}\Sigma_1\\ \vdots\\ \Gamma_1, c'\end{array} \quad \cdots \quad \dfrac{\begin{array}{c}\Sigma'_{2,i}\\ \vdots\\ \Gamma'_2, a_i\end{array} \quad \begin{array}{c}\Sigma''_2\\ \vdots\\ \Gamma''_2, f\end{array}}{\Gamma_2, g_i}\,\beta^i_2 \quad \cdots \quad \begin{array}{c}\Sigma_2\\ \vdots\\ \Gamma_2, d\end{array}}{\begin{array}{c}\Gamma, c''\\ \vdots\end{array}}\,\beta_1$$

The other possible cases, obtained by adding (resp., removing) some top (resp., bottom) places in (resp., from) the involved bipoles, are treated analogously and so omitted.

Remark : observe how, in this case, the induction step of the proof (prove that "if π' sequentializes, then π sequentializes too") follows (indeed, exploits) the correct direction of the semi-distributivity $((f \otimes (\&_{i=1}^{n}(a_i))\&d \vdash (\&_{i=1}^{n}(f \otimes a_i))\&d$ w.r.t. Figure 3), while the retraction rule R_6 follows the opposite direction $((\&_{i=1}^{n}(f \otimes a_i))\&d \nvdash (f \otimes (\&_{i=1}^{n}(a_i))\&d)$.

Case $i = 8$ **($\parr$-annihilation)** . Assume $\pi^* \rightsquigarrow_{R_8} \pi'^*$ so, by Lemma 1, we know how to build a BTN π' that abstracts to π'^*. Assume π and π' are like in the Figure 13. Then, by hypothesis of induction, π' sequentializes in to the following bipolar focusing sequential proof Π':

$$\Pi' : \dfrac{\cdots \quad \begin{array}{c}\Sigma'_i\\ \vdots\\ \dfrac{\dfrac{\dfrac{\Gamma_i, a_1, ..., a_{n-2}, d, e}{\Gamma_i, a_1, ..., a_{n-2}, \mathbf{a_n}, e}\,\beta_3}{\Gamma_i, a_1, ..., \mathbf{a_{n-1}}, a_n}\,\beta_2}{\Gamma, \mathbf{c}, a_n}\,\beta'_1\end{array} \quad \cdots}{\vdots}$$

Now, we can easily build the below proof Π from Π', simply, by replacing bipole β_1' of Π' with bipole β_1 in Π and then by erasing all the occurrences of a_n everywhere in the built Π except those ones occurring above the inference corresponding to β_1, following the notation of Figure 13.

$$\Pi : \dfrac{\dots \qquad \dfrac{\dfrac{\dfrac{\begin{array}{c}\Sigma_i\\ \vdots\\ \Gamma_i, a_1, ..., a_{n-2}, d, e\end{array}}{\Gamma_i, a_1, ..., a_{n-2}, \mathbf{a_n}, e}\,\beta_3}{\Gamma_i, a_1, ..., \mathbf{a_{n-1}}, a_n}\,\beta_2}{\Gamma, \mathbf{c}}\,\beta_1 \qquad \dots}{\vdots}$$

Observe that the erasing of the occurrences of a_n can be done safely[6], since, in the built Π, the only bipolar inferences whose trigger may consume the resource a_n are those ones corresponding to the bipole β_3 (whose abstraction, by definition of R_8, is a collapsed graph): this follows from the fact that, by Definition 1, a_n labels a port that in π' (resp., in π, by the abstraction Lemma 1) may only occur at the bottom of the bipole β_3 (resp., both at the bottom of β_3 and at the top of bipole β_1).
Other possible cases, obtained by adding (resp., removing) some top or bottom places in (resp., from) bipoles (e.g., β_3 could contain pending edges with other conclusions than a_n) are treated analogously and so omitted.

Example 4. Observe, the closed bipolar net given in Example 2 (Fig. 14, *Pic.* A_1), sequentializes in to the bipolar focusing proof Π displayed at the end of Example 1; we illustrates how the sequentialization works in that case. Assume π (*Pic.* A_1, Fig. 14) abstracts to π^* (*Pic.* A_2) and assume π^* retracts to π'^* like in *Pic.* A_3, after a block of distributive retractions (without losing generality, we may treats a sequence of retractions of the same type R_5 as a single generalized retraction R_5). By Abstraction Lemma 1 we may build a BTN π' from π'^* like in *Pic.* C of Fig. 15; then, by hypothesis of induction we know that π' sequentializes to the bipolar sequent proof Π' below:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\mathbf{g}, \mathbf{a}, \mathbf{c}}\,B_3'}{a, g, \mathbf{h}, \mathbf{d}}\,B_4' \quad \dfrac{\dfrac{}{\mathbf{g}, \mathbf{b}, \mathbf{c}}\,B_3''}{b, g, \mathbf{h}, \mathbf{d}}\,B_4'}{\mathbf{x}, g, h, d}\,B_2 \quad \dfrac{\dfrac{\dfrac{}{\mathbf{g}, \mathbf{a}, \mathbf{c}}\,B_3'}{a, g, \mathbf{h}, \mathbf{e}}\,B_4'' \quad \dfrac{\dfrac{}{\mathbf{g}, \mathbf{b}, \mathbf{c}}\,B_3''}{b, g, \mathbf{h}, \mathbf{e}}\,B_4''}{\mathbf{x}, g, h, e}\,B_2}{x, g, h, \mathbf{y}}\,B_1''}{\mathbf{f}}\,B_1'$$

Clearly, π is nothing else that π' in which we replaced bipoles β_1' and β_1'' with the single bipole β_1. Since bipole β_1' (resp., β_1'') corresponds (sequentializes) to the inference B_1' (resp., B_1''), then π sequentializes to Π obtained from Π' by simply replacing the two inferences B_1' and B_1'' with the unique inference B_1 which trivially corresponds to bipole β_1.

[6] In the sense that what we get after the erasing is still a sequent proof.

An other possible sequentialization of π' could have been the following sequent proof (we leave as an exercise to verify the sequentialization of π in that case):

$$
\dfrac{
  \dfrac{
    \dfrac{\dfrac{}{\mathbf{g},\mathbf{a},\mathbf{c}}\,B'_3}{a,g,\mathbf{h},\mathbf{d}}\,B'_4
    \qquad
    \dfrac{\dfrac{}{\mathbf{g},\mathbf{a},\mathbf{c}}\,B'_3}{a,g,\mathbf{h},\mathbf{e}}\,B''_4
  }{a,g,h,\mathbf{y}}\,B_2
  \qquad
  \dfrac{
    \dfrac{\dfrac{}{\mathbf{g},\mathbf{b},\mathbf{c}}\,B''_3}{b,g,\mathbf{h},\mathbf{d}}\,B'_4
    \qquad
    \dfrac{\dfrac{}{\mathbf{g},\mathbf{b},\mathbf{c}}\,B''_3}{b,g,\mathbf{h},\mathbf{e}}\,B''_4
  }{b,g,h,\mathbf{y}}\,B''_1
}{
  \dfrac{\mathbf{x},g,h,y}{\mathbf{f}}\,B'_1
}\,B_2
$$

(Picture H)

5 Conclusions, related and future works

In this work we provided:

1. a very simple syntax for open proof structures (BTSs) that allows to extend the paradigm of proof construction to the MALL fragment of LL. In particular, we set a precise correspondence, called *sequentialization* (Theorem 2) between focusing bipolar sequent proofs and correct BTSs (i.e., BTNs);
2. a convergent retraction system to check correctness of BTNs (Theorem 1);
3. an optimal strategy for incrementally building BTNs (Corollary 1).

Concerning other attempts to give a theoretical foundation of computational paradigms based on sequent proof construction, we only mention:

- some works of Pfenning and co-authors, from 2002 and later (see, e.g., [4]), which rely neither on focusing (or polarities) nor on proof nets but on softer notions of sequent calculus proofs;
- some works of Miller and co-authors which generalize focused sequent proofs to admit multiple "foci": see, e.g., [15] and [5]; the latter also provides a bijection to the unit-free proof nets of the MLL fragment, but it only discusses the possibility of a similar correspondence for larger fragments. At this moment, we are exploring a direct sequentialization from retractile transitory nets to, possibly open, multi-focus sequential calculi.

Concerning the related literature on additive proof nets, although there currently exist several satisfactory syntaxes for MALL proof structures, we briefly discuss some reasons that lead us to avoid most of them (at least in this first approach):

- *Girard*, [8]: requiring boolean (monomial) weights over proof structures is a condition that prevents certain transactional structures: take e.g. a simple BTS containing a single positive hyperlink or the rightmost BTS of Fig. 1;
- *Hughes-van Glabbeek*, [11]: similarly to the previous one, this syntax seems well adapt to take in to account only closed proof structures; actually, it has the inconvenient of allowing additive contractions only immediately below the axiom links; although this canonical form has great advantages for semantical reasons, it does not seem adapted to the composition of arbitrary modules that may require "non canonical" contractions.

Moreover, since these syntaxes make, more or less, explicit reference to graph dependencies (like *jumps*) they, *a-priori*, seem to garble the "principle of locality" required by retraction systems.

Finally, as future works, we aim at investigating:

- the complexity class of the optimal BTNs construction;
- an extension of the retraction system that could preserve BTNs under the (almost local) cut reduction proposed by Laurent and Maieli in [13].

Acknowledgements.

The author thanks the anonymous reviewers and Lorenzo Tortora de Falco for their useful comments and suggestions.

This work was partially supported by the PRIN Research Program *Logical Methods of Information Management*.

References

1. Andreoli, J.-M.: Focussing and Proof Construction. Annals of Pure and Applied Logic 107(1), pp. 131-163 (2001).
2. Andreoli, J.-M.: Focussing proof-nets construction as a middleware paradigm. In: Proc. of the 18th Int'l Conf. on Automated Deduction. LNCS, vol. 2392, pp. 501-516, Springer, Heidelberg (2002).
3. Andreoli, J.-M., Mazarè, L.: Concurrent Construction of Proof-Nets. In: Proc. of Computer Science Logic. LNCS, vol. 2803, pp. 29-42, Springer, Heidelberg (2003).
4. Cervesato, I., Pfenning, F., Walker, D., Watkins, K.: A concurrent logical framework II: Examples and applications. Technical Report CMU-CS-02-102, Department of Computer Science, Carnegie Mellon University (2002).
5. Chaudhuri, K., Miller, D., Saurin, A.: Canonical Sequent Proofs via Multi-Focusing. In: Proc. of International Conference on Theoretical Computer Science (TCS-5), Milan, Italy, pp. 383-396. Springer (2008).
6. Danos., V.: La Logique Linéaire appliquée à l'étude de divers processus de normalisation (principalment du λ-calcul). PhD Thesis, Univ. Paris VII (1990).
7. Fouqueré, C., Mogbil, V.: Rewritings for polarized multiplicative and exponential proof structures. Electronic Notes in Theoretical Computer Science, vol. 203(1), pp. 109-121 (2008).
8. Girard, J.-Y.: Linear Logic. Theoretical Computer Science 50, 1–102 (1987).
9. Girard, J.-Y.: Proof-nets: The parallel syntax for proof-theory Ursini, Agliano (Eds.), Logic and Algebra (1996).
10. Guerrini, S.: A linear algorithm for MLL proof net correctness and sequentialization. Theoretical Computer Science, 412(20): 1958-1978 (2011).
11. Hughes D., van Glabbeek, R.: Proof Nets for Unit-free Multiplicative-Additive Linear Logic. In: Proc. of the 18th IEEE Logic in Computer Science, IEEE Computer Society Press, Los Alamitos (2003).
12. Laurent, O.: Polarized Proof-Nets: Proof-Nets for LC. In: Girard, J.-Y. (ed.) TLCA 1999, LNCS, vol. 1581, pp. 213-227. Springer, Heidelberg (1999).
13. Laurent, O., Maieli, R.: Cut Elimination for Monomial MALL Proof Nets. In: Proc. of the 23rd IEEE Logic in Computer Science, IEEE Computer Society Press, Los Alamitos (2008).

14. Maieli, R.: Retractile Proof Nets of the Purely Multiplicative and Additive Fragment of Linear Logic. In: Proc. of the 14th Int'l Conf. LPAR. LNAI 4790, pp. 363-377, Springer (2007).
15. Miller, D.: A multiple-conclusion specification logic. Theoretical Computer Science 165:201-232 (1996).