Bipolar Proof Nets for MALL*

Proceedings of the PCC12 Conference – 17-18 August 2012, University of Copenhagen, Denmark

Roberto Maieli
Università degli Studi "Roma Tre"
maieli@uniroma3.it

In this work we present a computation paradigm based on a concurrent and incremental construction of proof nets (de-sequentialized or graphical proofs) of the pure multiplicative and additive fragment of Linear Logic, a resources conscious refinement of Classical Logic. Moreover, we set a correspondence between this paradigm and those more pragmatic ones inspired to transactional or distributed systems. In particular we show that the construction of additive proof nets can be interpreted as a model for *super-ACID* (or *co-operative*) transactions over distributed transactional systems (typically, multi-databases).

Keywords: linear logic, proof nets, transactional systems.

1 Introduction

This work takes a further step towards the development of an ambitious research programme, firstly started by Andreoli in [1], which aims at a theoretical foundation of a *computational programming paradigm based on the construction of proofs of linear logic* (LL, [4]). Naively, this paradigm relies on the following isomorphism: "proof"="state" and "construction step (or inference)"="state transition".
While the view of proof construction is well adapted to theorem proving, it is inadequate when we want to model the execution of widely distributed applications (typically over the Internet) which are designed with very flexible, concurrent and modular approaches. Due to their artificial sequential nature, sequent proofs are difficult to cut into composable (reusable) concurrent modules. A much more appealing solution consists in 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 (constrained) juxtaposition, obeying to some correctness criteria.
Actually, the proof net construction, as well as the proof net cut reduction, can be performed in parallel (concurrently), but despite from the cut reduction, there may not exist executable (sequentializable) construction steps: in other words, construction steps must satisfy an "efficient" correction criterion. The resulting paradigm is very close to more pragmatic ones, like those ones coming from transactional or distributed systems.
Concretely, here, we present a model for the incremental construction of proof nets of the *pure multiplicative and additive fragment of linear logic* (MALL, [5]). This model extends the previous one, given in [2], for the *pure multiplicative fragment of linear logic* (MLL). In particular, we give a syntax for *bipolar focussing proof-structures* that are de-sequentialized (geometrical) representations of possibly *incomplete* (*open* or *with proper axioms*) proofs of the *bipolar focussing sequent calculus* [1]. This calculus has the following properties:

*Research supported by the PRIN Project CONCERTO.

Preliminary Report. Final version to appear in:

1. the possibly incomplete (open) focussing proofs are strictly isomorphic to the possibly open proofs of the bipolar focussing sequent calculus;
2. the complete (closed or with logical axioms) focussing proofs are fully representative of all the closed proofs of linear logic.

Hence by 1 and 2, proof construction can be performed equivalently in these three proof systems of LL: *sequent calculus*, *focussing sequent calculus* and *bipolar focussing sequent calculus*. Bipolarity and focussing properties ensure more compact proofs since they get rid of some irrelevant intermediate steps in the construction.

In [2, 3], the concurrent construction of open (transitory) MLL proof nets was interpreted as an incremental juxtaposition of link modules (agents) that allows to model the behavior of *ACID transactions* over strongly distributed systems. Here the proof construction of transitory MALL proof nets is interpreted as an additive (super) juxtaposition of interacting *slices* (multiplicative transitory proof nets). Locally the concurrent construction of MALL proof nets can be viewed as an incremental juxtaposition of hyper-links (a disjoint sum of multiplicative links) that, like co-operative agents, allow to model some kinds of (non-deterministic) co-operation among ACID transactions.

2 Bipolar Focussing Sequent Calculus

We recall some basic definitions of the *standard sequent calculus of MALL*, then we introduce the related *bipolar focussing sequent calculus*, based on the crucial properties of *focussing* and *bipolarity* (find more in [1], [6] and [7]). We, arbitrarily assume literals $a, a^{\perp}, b, b^{\perp}, \ldots$ with a polarity: *negative* for atoms and *positive* for their duals, then given a set $\mathscr{A}$ of atoms, an $\mathscr{A}$-formula is a formula built from atoms and their duals, using the (two groups of) connectives of MALL: *negative*, $⅋$ ("par") and & ("with") and *positive*, $\otimes$ ("tensor") and $\oplus$ ("plus"). Finally, a proof of MALL is build by means of the following (groups of) inferences:

$$\textit{identity}: \quad \frac{}{A, A^{\perp}}\,\mathrm{ax} \qquad\qquad \textit{multiplicatives}: \quad \frac{\Gamma, A \qquad \Delta, B}{\Gamma, \Delta, A \otimes B}\,\otimes \qquad \frac{\Gamma, A, B}{\Gamma, A ⅋ B}\,⅋$$

$$\textit{additives}: \quad \frac{\Gamma, A \qquad \Gamma, B}{\Gamma, A \& B}\,\& \qquad \frac{\Gamma, A}{\Gamma, A \oplus_1 B}\,\oplus_1 \qquad \frac{\Gamma, B}{\Gamma, A \oplus_2 B}\,\oplus_2$$

The *focussing property* states that, in the proof search (or proof construction), we can build (bottom up) a sequent proof by alternating clusters of negative inferences followed by clusters of positive inferences. As consequence of this bipolar alternation we obtain more compact proofs in which we get rid of the most part of all the bureaucracy hidden in sequential proofs (as, for instance, irrelevant permutation of rules): what remains is a focussing bipolar proof. Remind that w.r.t. proof search negative (resp., positive) connectives involve a kind of *don't care non-determinism* (resp., *true non-determinism*).

An $\mathscr{A}$-*monopole* is an $\mathscr{A}$-formula built on negative $\mathscr{A}$-atoms using only the negative connectives; an $\mathscr{A}$-*bipole* is an $\mathscr{A}$-formula built from $\mathscr{A}$-monopoles and positive $\mathscr{A}$-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 $\mathscr{F}$ of $\mathscr{A}$-bipoles, the *bipolar focussing sequent calculus* $\Sigma[\mathscr{A}, \mathscr{F}]$ is a set of inferences of the form:

$$\frac{\Gamma_1 \qquad \ldots \qquad \Gamma_n}{\Gamma}\,F$$

where the conclusion Γ is a sequent made by only of negative $\mathscr{A}$-atoms and the premises $\Gamma_1,...,\Gamma_n$ are obtained by fully focussing decomposition of some bipole $F \in \mathscr{F}$ in the the context Γ (a multiset of negative atoms). More precisely, due to the presence of additives (in particular the $\oplus$ connectives) a bipole F is naturally associated to a set of inferences $F_1,...,F_{m+1}$, where m is the number of $\oplus$ connectives presents in F. For instance, in the purely multiplicative fragment of LL, the bipole $F = a^{\perp} \otimes b^{\perp} \otimes (c \parr d) \otimes e$, where a,b,c,d,e are (negative) $\mathscr{A}$-atoms, yields the inference on the left-hand side (more compact w.r.t. the explicit one on the right hand side):

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

where Γ, Δ rage over a multiset of negative $\mathscr{A}$-atoms. Note that the identity axioms $a, a^{\perp}$ and $b, b^{\perp}$ are omitted in the bipolar sequent proof for simplicity sake. The couple a, b here plays the role of a *trigger* or *mutlifocus* of the F-inference; more generally, a trigger of a bipole is a multiset of duals of the positive atoms which occurs in it. The main feature of the bipolar focussing sequent calculus is that its inferences are triggered by multiple focus (like in Forum [8]).
The bipolar focussing sequent calculus is proved (Theorem 1, see [1]) to be isomorphic to the focussing sequent calculus, so that proof construction can be performed indifferently in the two systems. The main idea exploited in the proof of Theorem 1 is the *bipolarisation technique*, that is a simple procedure that allows to transform a provable formula F in the LL sequent calculus into a set of bipoles (belonging to an "universal program" in the bipolar sequent calculus). For our purpose, we briefly illustrate this technique only for the MALL fragment, with an instance given in the Example1.

A *naming scheme* is a triple $\langle \mathscr{A}, \mathscr{A}', \eta \rangle$ where $\mathscr{A} \subset \mathscr{A}'$ are sets of negative atoms and η is a bijection from the $\mathscr{A}$-formulas into $\mathscr{A}'$ such that $\eta_a = a$ for all $a \in \mathscr{A}$. The *universal program* for a naming scheme $\langle \mathscr{A}, \mathscr{A}', \eta \rangle$ is the set of $\mathscr{A}'$-bipoles of the form $\nu(F)$ where F ranges over the $\mathscr{A}$-formulas not reduced to a negative atom. The ν-mapping on $\mathscr{A}$-formulas is defined in three steps as follows:

1. (*negative layer*) mapping $\nu^{\uparrow}$ from $\mathscr{A}$-formulas to $\mathscr{A}'$-monopoles

$$\begin{array}{rcl} \nu^{\uparrow}(F_1 \parr F_2) & = & \nu^{\uparrow}(F_1) \parr \nu^{\uparrow}(F_2) \\ \nu^{\uparrow}(F_1 \& F_2) & = & \nu^{\uparrow}(F_1) \& \nu^{\uparrow}(F_2) \\ \nu^{\uparrow}(F) & = & \eta_F \text{ in all the other cases;} \end{array}$$

2. (*positive layer*) mapping $\nu^{\downarrow}$ from $\mathscr{A}$-formulas to $\mathscr{A}'$-bipoles or monopoles

$$\begin{array}{rcl} \nu^{\downarrow}(F_1 \otimes F_2) & = & \nu^{\downarrow}(F_1) \otimes \nu^{\downarrow}(F_2) \\ \nu^{\downarrow}(F_1 \oplus F_2) & = & \nu^{\downarrow}(F_1) \oplus \nu^{\downarrow}(F_2) \\ \nu^{\downarrow}(a^{\perp}) & = & a^{\perp} \text{ if } a \text{ is a negative atom} \\ \nu^{\downarrow}(F) & = & \nu^{\uparrow}(F) \text{ in all the other cases;} \end{array}$$

3. mapping ν from $\mathscr{A}$-formulas to $\mathscr{A}'$-bipoles

$$\nu(F) \quad = \quad \eta_F^{\perp} \otimes \nu^{\downarrow}(F).$$

Theorem 1 (universal program) *Given a naming scheme $\langle \mathscr{A}, \mathscr{A}', \eta \rangle$, let $\mathscr{U}$ be its universal program. For any $\mathscr{A}$-formula F there is an isomorphism between the focussing proofs of F in linear logic and the proofs of η_F in the bipolar focussing sequent calculus $\Sigma[\mathscr{A}', \mathscr{U}]$.*

Example 1 *Assume an $\mathscr{A}$-formula $F = (a\&b)⅋((a^\perp \oplus b^\perp) \otimes c^\perp)⅋(c \otimes (d^\perp \oplus e^\perp))⅋(d\&e)$ with subformulas $G = (a^\perp \oplus b^\perp) \otimes c^\perp$ and $H = c \otimes (d^\perp \oplus e^\perp)$ and negative atoms a,b,c,d,e. After bipolarisation of F we get the following bipoles of the universal program $\mathscr{U}$:*

$$\begin{array}{rcl} \nu(F) & = & \eta_F^\perp \otimes ((a\&b)⅋\eta_G⅋\eta_H⅋(d\&e)) \\ \nu(G) & = & \eta_G^\perp \otimes ((a^\perp \oplus b^\perp) \otimes c^\perp) \\ \nu(H) & = & \eta_H^\perp \otimes c \otimes (d^\perp \oplus e^\perp); \end{array}$$

– the bipole $\nu(F)$ corresponds to the unique inference $\nu(F)$:

$$\frac{\Gamma,\eta_G,\eta_H,a,d \qquad \Gamma,\eta_G,\eta_H,a,e \qquad \Gamma,\eta_G,\eta_H,b,d \qquad \eta_G,\eta_H,b,e}{\Gamma,\eta_F}\,\nu(F)$$

– the bipole $\nu(G)$ is associated to a pair of inferences:

$$\frac{\Gamma}{\Gamma,\eta_G,a,c}\,\nu(G)_1 \qquad \textit{and} \qquad \frac{\Gamma}{\Gamma,\eta_G,b,c}\,\nu(G)_2$$

– similarly, the bipole $\nu(H)$ is associated to a pair of inferences:

$$\frac{\Gamma,c}{\Gamma,\eta_H,d}\,\nu(H)_1 \qquad \textit{and} \qquad \frac{\Gamma,c}{\Gamma,\eta_H,e}\,\nu(H)_2$$

Finally, here is the complete bipolar focussing proof of η_F that is isomorphic, by Theorem 1, to the (omitted) proof of F in the LL focussing sequent calculus:

$$\frac{\dfrac{\dfrac{}{\eta_G,a,c}\,\nu(G)_1}{\eta_G,\eta_H,a,d}\,\nu(H)_1 \quad \dfrac{\dfrac{}{\eta_G,a,c}\,\nu(G)_1}{\eta_G,\eta_H,a,e}\,\nu(H)_2 \quad \dfrac{\dfrac{}{\eta_G,b,c}\,\nu(G)_2}{\eta_G,\eta_H,b,d}\,\nu(H)_1 \quad \dfrac{\dfrac{}{\eta_G,b,c}\,\nu(G)_2}{\eta_G,\eta_H,b,e}\,\nu(H)_2}{\eta_F}\,\nu(F)$$

Observe that while the above derivation is quite compact, it still presents a lot of structural inconvenient such as duplications of sub-trees; phenomena like these are crucial when we want to modelize the behavior of distributed systems. For these reasons, in the next section, we move to more flexible (geometrical indeed) proof structures.

3 Bipolar Focussing Proof Structures

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

Definition 1 (links) *Assume an infinite set $\mathscr{L}$ of* resource places $l_1, l_2, \ldots$ *(also* addresses *or* loci *like in Ludics [6]); the special untyped place $\star$ is called* jump place*. A* link *consists in two disjoint sets of loci, top and bottom, together with a polarity p, positive or negative, and with the conditions that:*

- *a positive link must have* at least one *bottom place; it may contain no more than one jump place among its bottom places;*
- *a negative link must have* exactly one *bottom place; it may contain no more than one jump place among its top places.*

If the set of top places is not empty, then a link is said transitional.

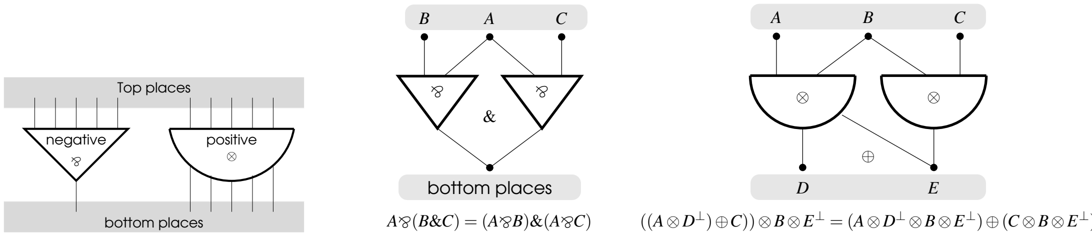

Figure 1: *links* and *hyperlinks*

Graphically links are represented like in the left hand side picture of Figure 1 and distinguished by their shape: triangular for negative and round for positive links. We use variables $x^p, y^p, z^p, \ldots$ with polarity $p \in \{+,-\}$ for links.
Intuitively, negative links correspond to generalized (n-ary) ⅋-links while positive links correspond to generalized $\otimes$-links. In order to capture the additive behavior (a non-deterministic "sharing nature") we need to allow superposition of links; this will naturally bring us to the next notion of hyperlinks.

Definition 2 (hyperlinks) *An* hyperlink *is a set of links that share some (at least one) places as follows:*

- *a* negative hyperlink *contains only negative links and an unique bottom place; all its jump places must be distinguished (i.e., its negative links have no jump places in common).*
- *a* positive hyperlink *contains only positive links and* at least one *bottom place; all its jump places must be distinguished (i.e., its positive links have no jump places in common).*

Analogously to (multiplicative) links, negative hyperlinks correspond to generalized &-links (additive conjunction) while positive links correspond to generalized $\oplus$-links (additive sum). Recall that in linear logic the additive connectives capture non deterministic computational phenomena (typically of distributed middleware systems). An example of negative (resp. positive) hyperlink is depicted in the middle (resp., right) hand side of Figure 1. Observe that these links represent, graphically, the distributive law of negative ⅋/& (resp., positive $\otimes/\oplus$) connectives. The notation X^+ (resp., X^-) denotes a positive (resp., a negative) hyperlink. Moreover we say that:

- an edge is called a *jump edge* (simply *jump*) when it goes from a positive jump place to a negative jump place;
- a (positive) link x^+ *depends on* a (negative) link y^- when there exists a jump edge that goes from x^+ to y^-;
- a pair of positive links x_1^+ and x_2^+ belonging to a same $+$hyperlink X^+ is *toggled* by a negative hyperlink Y^-, if there exist two negative links y_1^-, y_2^- in Y^-, s.t. there is a jump from x_1^+ to y_1^- and a jump from x_2^+ to y_2^-.
  A graphical interpretation of the toggling condition with jump edges is then given in the picture on the left hand side of Figure 2.

Observe that jumps play here the same role (dependency) *eigen weights* play in [5].

Definition 3 (bipolar focussing proof structures) *A MALL focussing proof structure (shortly, BPS) is a set π of hyperlinks satisfying the following conditions:*

1. *the sets of top (bottom) places of any pair of hyperlinks are disjoint;*

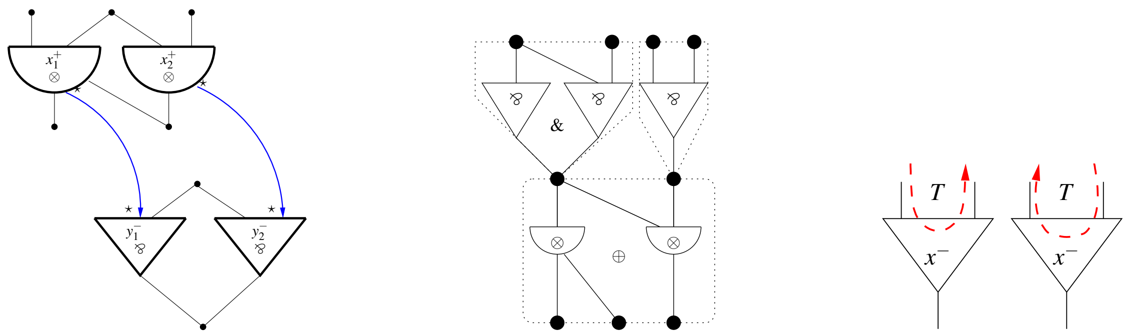

Figure 2: *toggling* (left side), *bipole* (middle side) and *singularities* (right side)

2. *if two hyperlinks are adjacent, then they have opposite polarity;*
3. *in any $+$hyperlink every pair of links is toggled by a $-$hyperlink;*
4. *jump places are distinguished (links do not share jump places).*

Finally, π is said to be *elementary* if it is bipolar and contains exactly a positive hyperlink: each elementary focussing proof structure corresponds to a bipole (see the picture in the middle side of Figure 2).
We are interested on those (correct) proof structures that correspond to bipolar focussing sequent proofs: these are called bipolar proof nets. Before introducing these, we need some technical stuff.
A hyperlink X (or simply, a link) is said to be *just below* (resp., *just above*) an hyperlink Y if there exists a place that is both at the top (resp., bottom) of X and at the bottom (resp., top) of Y. Two hyperlinks are said *adjacent* if one is just below (resp., just above) the other. Then, fixed a BPS π:

- a *&-resolution* is a choice of exactly one negative link for each negative hyperlink (all the other negative links will be erased);
- a *slice* $S(\pi)$ for π is the graph obtained from π after the *erasing* induced by a &-resolution, as follow: *(i)* a place is erased if all the top (bottom) links sharing it are erased; *(ii)* a link is erased when at least one of its places is erased.
- a *trip* T in a slice $S(\pi)$ for π is a non-empty binary relation on $|S|$ (the set of link of S) which is finite, connected and s.t. any link $x \in |S|$ has at most one successor (resp., one predecessor), if it exists. Then, a *negative middle link* x (with a predecessor and a successor) of a *proper trip* T (not reduced to a loop with only two links) is a *singularity* for T iff T *enters* x *downwards* and *exists* x *upwards* (graphically, T bounces on x, like in the right hand side picture of Figure 2).

Definition 4 (bipolar focussing proof net) *A BPS π of MALL is* correct*, i.e., it is a* bipolar proof net *(BPN) iff any proper loop trip in any slice $S(\pi)$ contains at least a singularity.*

An instance of BPN is given in the left hand side picture of Figure 3. It is not difficult to check that any proper loop trip in any slice contains at least a singularity, in particular that is true for the slice depicted in the right hand side of Figure 3. In order to simplify the reading of these pictures, jumps from positive to negative links are drawn as oriented (colored) curved edges.
We can set a precise correspondence between sequent proofs and proof nets: in the literature this correspondence is called ”sequentialization”.

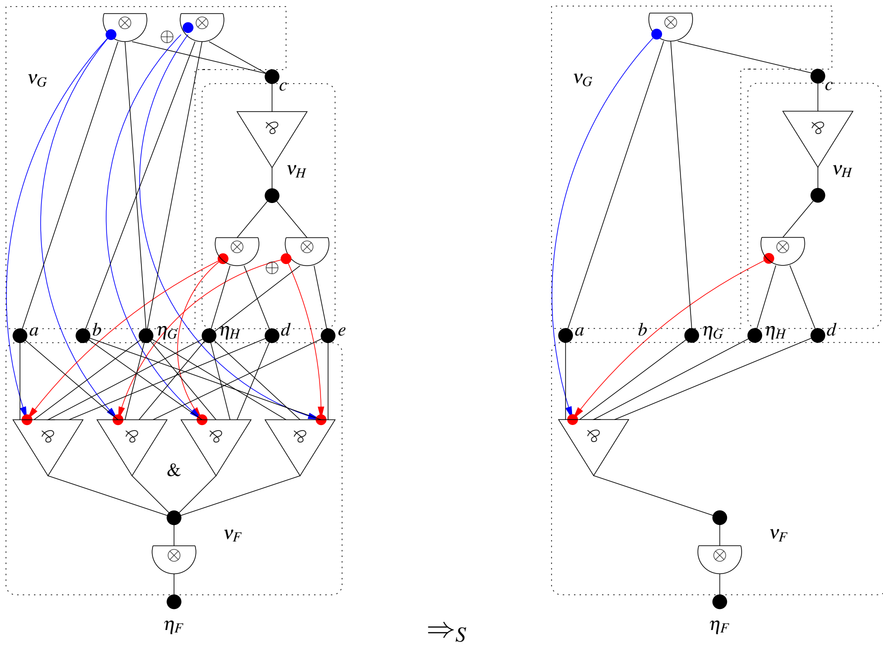

Figure 3: a *bipolar poof net* (left hand side) with a *slice* (right hand side)

Theorem 2 ((de-)sequentialization) *A bipolar focussing sequent proof* Π *of* Γ *can be* de-sequentialized *in a canonical way into a bipolar focussing proof net* π *with same conclusion* Γ *and vice versa, a proof net* π *can be* sequentialized *into a sequent proof* Π *with same conclusions.*

The de-sequentialization part of Theorem 2 is proved by induction on the size of the given sequent proof (i.e. the number of bipoles). For the base of the induction, there exists precise correspondence between a bipole and an elementary proof proof structure which is trivially correct (i.e. a proof net). As an instance, observe the focussing bipolar sequent proof of η_F of the Example 1 de-sequentializes into the bipolar proof net drawn in the left hand side picture of Figure 3 and vice versa. Actually, in order for the bipole $\nu(H)$ to correspond to an elementary focussing proof structure, there is need to introduce a dummy negative link with one top place for c. This could have been avoided by explicitly introducing a *polarity inverter*, as usually done in strictly polarized syntax (see [7]).

The sequentialization part is proved by induction on the number of slices of π; observe that a BPN reduced to a single slice is trivially a MLL BPN which can be shown that sequentializes into a MLL sequential proof. The crucial task is to show how to gluing the multiplicative (MLL) sequential slices into an additive (MALL) sequential proof.

In the next section we study the problem of constructing a proof net by a juxtaposition of concurrent bipoles (agents). This proof net construction can be viewed as a computational paradigm for middleware (infrastructure) programming.

4 Proof Net Construction as a Middleware Paradigm

In building a proof net, places (except $\star$) are decorated by type informations (occurrences of negative atoms); each bipole is viewed as a disjoint sum of collaborative agents which continuously attempt to

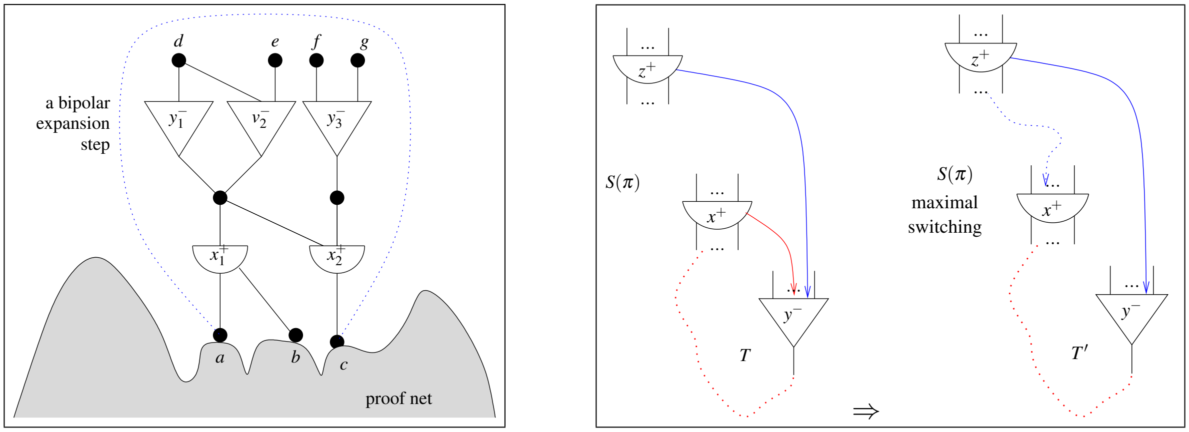

Figure 4: *expansion step* (left hand side) and *maximal switching* (right hand side)

perform a construction step, that is, an expansion of the proof net obtained by adding an elementary bipolar proof structure (a bipole) from the places whose types match the trigger of the given bipole. Bipoles always drive the construction bottom-up like in the left hand picture of Figure 4.

An expansion step is correct if it preserves the property of being a proof net. Checking correctness (singularity-free trips) is a task which may involve *visiting a large portion* of the expanded proof structure. Now, since this construction is performed collaboratively and concurrently by a cluster of bipoles for true concurrency we need to:

1) restrict the traveling region (reducing possible conflicts among agents);
2) protect (lock) the gathered information against attempts of other concurrent agents;
3) increment/update, in case of success, the locked information for transition.

Good bounds for these tasks are necessary; however, in the following two sub-sections 4.1 and 4.2 we mainly focus on the task 1.

4.1 Maximal switchings

First we show that, in order to detect singularity-free trips we may restrict us to consider only particular subgraphs of switchings, these are called maximal switchings.

Definition 5 (maximal switching) *A jump edge from x^+ to y^- is said* maximal *in a switching $S(\pi)$ if there not exists in such a switching a positive link z^+ such that it depends on y^- too and it is above x^+ in π; then a* maximal switching *is a switching containing only maximal jump edges.*

Lemma 1 (maximal switchings) *A proof structure π is correct iff any proper loop of any maximal switching S for π contains at least a singularity.*

By Definition 5, if there exist in $S(\pi)$ two positive links, x^+ and z^+, both depending on y^- and with z^+ above x^+, then there must exist in $S(\pi)$ a path going from x^+ upwards to z^+; clearly, if there exists a singularity-free loop T in $S(\pi)$ containing a jump from x^+ to y^-, there will also exist a singularity-free loop T' in $S(\pi)$ containing a jump from z^+ to y^- (see the right hand side of picture of Figure 4).

4.2 Expansion under domination

We fix once for all a maximal switching S for π, then we show (Lemma 2) that, w.r.t. a candidate expansion, only certain negative links must be explored (isolated and locked); the other ones are available for other possible expansions (or transactions).

Definition 6 (domination order) *Assume x,y are two negative links in a switching S for π; a* root *of S is any (positive) link of S that has no link below it. Then, $x \leq y$ (x* dominates *y) if any singularity-free trip starting at a root and stopping upwards at y visits x upwards.*

Proposition 1 (forest order) *The relation $\leq$ on negative links of S is a* forest order*; it is reflexive, anti-symmetric, transitive and it satisfies the following property on negative links:*

$$\forall x,y,z \ \ if \ \ (x \leq z \wedge y \leq z) \ \ then \ \ (x \leq y \vee y \leq x).$$

The *joint dominator of N*, $\bigwedge(N)$, is the *greatest lower bound* (g.l.b., when it exists), by $\leq$, of a set of negative links N.
If the set of the predecessor by $<$ of a negative link x is not empty, then it has a greatest element, by $\leq$, called the *immediate dominator $D(x)$* of a negative link x.

Lemma 2 (isolation property) *Let x,y be two negative links and T be a singularity-free trip of a switching S for π starting downwards at x and stopping upwards at y; then, any negative link z visited by T is* strictly dominated *by the joint dominator $\bigwedge\{\lambda_1,\lambda_2\}$ (if defined):*

$$\forall z^- \in |T|, \ \bigwedge\{x,y\} < z.$$

Clearly, w.r.t. an expansion of a proof net π by a bipole β, Lemma 2 gives a good (lower) bound to the region to be explored in order to detect a singularity-free trip in a switching $S(\pi * \beta)$ (where $\pi * \beta$ denotes the juxtaposition of β over π). An instance of a candidate multiplicative expansion that is not correct is given in Figure 5: grey (or red) negative links denote all those negative links that (according to Lemma 2) must be visited in order to look for a singularity-free loop inside a top switching; while the light grey (or green) ones are unexplored and so available for other transactions.
We propose in the next (last) subsection some applications of Lemmas 1 and 2 to the theoretical interpretation of distributed transactional systems.

4.3 Transactional Systems

A transaction combines a group of independent actions into a single action with a set of predictable outcomes. Traditionally, transactions are required to adhere to the *ACID* properties of *Atomicity* (ensuring that all actions in the transaction either complete successfully, or revert to a state where none of them were run), *Consistency* (ensuring that the system is not put into an illegal state), *Isolation* (letting concurrent transactions run as if they were the only transaction being processed), and *Durability* (ensuring that any completed transaction has its stable outcome and cannot be undone, even by accidental hardware or software failure). However, while transaction management in traditional systems typically offers an acceptable level of service, the same cannot be said for transactions achieved by combining services offered by multiple systems. Such multi-databases transactions often run for much longer periods of time than traditional transactions, so locking any data may block other transactions for an unacceptable length of time. Because of this, the traditional ACID properties are typically reduced in strength, helping to

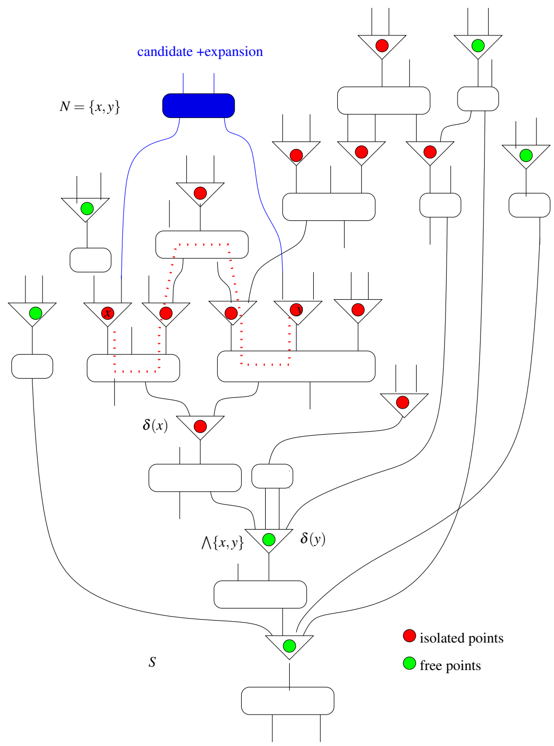

Figure 5: proof net interpretation of the *isolation property*

ensure that the entire system maintains an acceptable level of service. Typically, in the Web Services environment, traditional ACID transactions are not always sufficient to support the activities that businesses would like to process. Transactions that involve multiple service providers can run for long periods of time. This can result in negative side-effects when combined with traditional transaction-based concurrency control mechanisms. While Web Services transactions standards do exist, it is still difficult (e.g., for an end-user) to combine services from loosely-coupled providers so that they are used as a single *co-operative transaction* ([9]).

Under this respect, the paradigm of proof net construction can be put in correspondence with transactional systems paradigms. That can be seen as an analogous of the well known *Curry-Howard correspondence* between the cut-reduction paradigm and the functional programming paradigm. Any correct expansion step can be seen as a transaction; more precisely:

- Lemma 2 captures the *multiplicative behavior* of proof nets and corresponds to the *isolation property* of ACID transactions. For instance, Figure 5 can be interpreted as a candidate multiplicative expansion that is not an ACID transaction (it is not correct);

- Lemmas 1 captures the *additive behavior* of proof nets and corresponds to *co-operative transactions*: actually, we can additively ”slice” a transaction in to a sum of interacting (or cooperative) ACID transactions; maximality of switching guarantees that only certain resources will be locked.

References

[1] J.-M. Andreoli. Focussing and Proof Construction. *Annals of Pure and Applied Logic* 107(1), pp 131–163, 2001.

[2] J.-M. Andreoli. Focussing proof-nets construction as a middleware paradigm. In A. Voronkov, ed., *Proc. of the 18th Int'l Conference on Automated Deduction. Lecture Notes in Computer Science, pp. 501-516.* Denmark, 2002. Springer Verlag.

[3] J.-M. Andreoli and L. Mazarè. Concurrent Construction of Proof-Nets. *In proc. of Computer Science Logic (CSL)*, Wien, Austria, 2003.

[4] J.-Y. Girard. Linear Logic. *Theoretical Computer Science*, 50:1–102, 1987.

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

[6] J.-Y. Girard. Locus Solum. *Mathematical Structures in Computer Science* 11, pp. 301-506, 2001.

[7] Laurent, O. *Polarized Proof-Nets: Proof-Nets for LC (Extended Abstract)*. In J.-Y. Girard, editor, Typed Lambda Calculi and Applications 1999, LNCS 1581, pp. 213-227. Springer-Verlag. Avril 1999.

[8] D. Miller. Forum: a multiple-conclusion specification logic. *Theoretical Computer Science* 165, pp. 201-232, 1996.

[9] D. Paul, M. Wallis, F. Henskens and M. Hannaford. Transaction support for interactive web applications. *Proceedings of the 4th International Conference on Web Information Systems and Technologies (WEBIST 2008).* Funchal-Madeira, Portugal 4-7 May, 2008.