



HAL
open science

Proofs as Executions

Emmanuel Beffara, Virgile Mogbil

► **To cite this version:**

Emmanuel Beffara, Virgile Mogbil. Proofs as Executions. 7th International Conference on Theoretical Computer Science (TCS), Sep 2012, Amsterdam, Netherlands. pp.280-294, 10.1007/978-3-642-33475-7_20 . hal-00586459v2

HAL Id: hal-00586459

<https://hal.science/hal-00586459v2>

Submitted on 25 May 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Proofs as executions

(rapport interne LIPN - Mai 2012)*

Emmanuel Beffara¹ and Virgile Mogbil²

¹ IML – FRE3529, CNRS – Université d’Aix-Marseille

² LIPN – UMR7030, CNRS – Université Paris 13

Abstract. This paper proposes a new interpretation of the logical contents of programs in the context of concurrent interaction, wherein proofs correspond to valid executions of a processes. A type system based on linear logic is used, in which a given process has many different types, each typing corresponding to a particular way of interacting with its environment and cut elimination corresponds to executing the process in a given interaction scenario. A completeness result is established, stating that every lock-avoiding execution of a process in some environment corresponds to a particular typing. Besides traces, types contain precise information about the flow of control between a process and its environment, and proofs are interpreted as composable schedulings of processes. In this interpretation, logic appears as a way of making explicit the flow of causality between interacting processes.

1 Introduction

The extension of the familiar Curry-Howard correspondence to interactive models of computation has been an active research topic for several decades. Several systems were proposed based on linear logic [9], following the fundamental observation that it is a logic of interaction. Interpretations of proofs as processes, first formalized by Abramsky [1], later refined by various people including the first author [2], stressed that proof nets [10] and process calculi have significant similarities in dynamics. At the same time, type systems for concurrency [20] revealed to be equivalent to variants of linear logic [12, 4]. These approaches successfully stress the fact that concurrent calculi are very expressive and versatile models of interactive behaviour, however they are not satisfactory yet as a proof-theoretical account of concurrency, because they tend to impose determinism in execution, effectively constraining processes to functional behaviour.

Several approaches to the question of non-determinism in logic have been proposed. The use of the additive connectives of linear logic as a proof-theoretic representation of it was for instance explored by Mairson and Terui [14] to provide a notion of non-deterministic cut-elimination, or by Maurel [15] or the second author [18] who used it to represent the kind of non-determinism familiar

* Work partially supported by the french project Complice (ANR-08-BLANC-0211-01)

in complexity theory. In a different style, differential logic was recently developed by Ehrhard and Regnier [8] and its untyped proof formalism was shown expressive enough to represent the π -calculus [7].

The present work proposes a different approach to the topic, by questioning the “proofs-as-programs” paradigm. Proof theory wants cut elimination to be confluent, because the meaning of proofs lies in their normal forms. On the other hand, reduction in process calculi is execution: the meaning of a term is not its final irreducible form but what happens to get there, as interaction with other processes. Hence we propose to match proofs with executions rather than terms. But this raises the new question of what is the logical meaning of an execution. Here we must remember that cut elimination is a process of *explicitation* and cut-free proofs are explicit, direct reasonings justifying some fact. In our case, the fact is the interaction, which is a scheduling of a set of events in a system. The justification, then, is the *control flow* through the system, specifying when actions happen and when execution jumps from one process to another.

Technically, we illustrate this idea in the very simple setting of finitary CCS with no choice operator, in order to focus on the novel ideas of our approach, but ways to extend these techniques to a larger class of processes are sketched in the perspectives. The corresponding logic is multiplicative linear logic, with a family of modalities à la Hennessy-Milner [11] representing actions.

In our type system, multiplicatives represent causality and independence between parts of a run, using connectedness/acyclicity arguments to describe avoidance of deadlocks. Modalities represent observable transitions, with explicit scheduling constraints using the well-known stratifying effect of boxes in proof nets. Axiom rules have an unusual interpretation: they are void of interactive content (no forwarding or copycat behaviour), but they logically implement the transfer of control flow between different parts of a running process.

This handling of control flow using the symmetries of linear logic is reminiscent of the work of Mazurak and Zdancewicz [16] who use linear negation as an explicit scheduling operation. Our work differs from theirs and other works on typing for concurrency, in that we proceed “backwards”: while Curry-Howard systems for concurrency embed logical systems into concurrent calculi, we embed executions of processes into a logical system.

Outline The paper is organized as follows: Section 2 introduces a logic of schedulings based on linear logic and illustrates our interpretation. Section 3 defines a simple fragment of CCS and a notion of determinisation, used to represent executions as terms. Section 4 presents the proof nets for the logic of schedulings and its cut-elimination. Section 5 shows the typing of executions and the associated subject reduction property, and Section 6 establishes the completeness property that all lock-avoiding executions are typable.

2 A logic of schedulings

We first present the logic we use to describe interactions and schedulings. It corresponds to the multiplicative fragment of linear logic [9], augmented with a family of modalities that describe actions.

$$\begin{array}{c}
\frac{P \vdash \Gamma, A, B}{P \vdash \Gamma, A \wp B} \text{ (}\wp\text{)} \quad \frac{P \vdash \Gamma, A \quad Q \vdash B, \Delta}{P \mid Q \vdash \Gamma, A \otimes B, \Delta} \text{ (}\otimes\text{)} \quad \frac{P \vdash \Gamma, A \quad Q \vdash A^\perp, \Delta}{P \mid Q \vdash \Gamma, \Delta} \text{ (cut)} \\
\frac{}{1 \vdash A, A^\perp} \text{ (ax)} \quad \frac{P \vdash \Gamma, A}{a.P \vdash \Gamma, \langle a \rangle A} \text{ (act)} \quad \frac{P \vdash \Gamma \quad a \notin \Gamma}{(\nu a)P \vdash \Gamma} \text{ (new)} \\
\text{Derived rules: } \left\{ \begin{array}{ll} \frac{P : \Gamma, A \vdash B}{P : \Gamma \vdash A \multimap B} \text{ (}\multimap\text{R)} & \frac{P : \Gamma \vdash A \quad Q : \Delta, B \vdash C}{P \mid Q : \Gamma, \Delta, A \multimap B \vdash C} \text{ (}\multimap\text{L)} \\ \frac{P : \Gamma \vdash A \multimap B \quad Q : \Delta \vdash A}{P \mid Q : \Gamma, \Delta \vdash B} \text{ (mp)} & \frac{P : \Gamma, A \vdash B}{a.P : \Gamma, \langle \bar{a} \rangle A \vdash B} \text{ (act)} \end{array} \right.
\end{array}$$

Table 1. Inference rules for MLL with action modalities (MLL_a)

Definition 1 (MLL_a). *The formulas of MLL_a are built by the grammar*

$$A, B ::= \alpha \mid \alpha^\perp \mid A \otimes B \mid A \wp B \mid \langle a \rangle A \mid \langle \bar{a} \rangle A$$

where the α are literals and the a are CCS names. The negation A^\perp of a non-literal formula A is defined by de Morgan duality as $(A \otimes B)^\perp = A^\perp \wp B^\perp$ and $(\langle a \rangle A)^\perp = \langle \bar{a} \rangle A^\perp$. A type $(\Gamma, \Delta \dots)$ is a finite multiset of formulas. Derivations are built from the rules of table 1, where the left side of \vdash is a CCS term up to structural congruence (as of section 3). The same rule (act) applies for names of both polarities. In (new), $a \notin \Gamma$ means that neither $\langle a \rangle$ nor $\langle \bar{a} \rangle$ occurs in Γ .

Although it is formulated as a type system for processes, this logic should be interpreted as a calculus for building schedulings. To explain this interpretation, we adopt a few notations that stress the functional aspect of the system: $P : A_1, \dots, A_n \vdash B$ represents the judgement $P \vdash A_1^\perp, \dots, A_n^\perp, B$ and the binary connective $A \multimap B$ stands for $A^\perp \wp B$. We easily get the derived rules of table 1: (\multimap R) and (\multimap L) are respectively a reformulation of (\wp) and (\otimes), and (mp) is *modus ponens* for linear implication, obtained with (ax), (\otimes) and (cut). This is an intuitionistic or implicative formulation, but we do need the full expressiveness of the MLL_a for the developments of the following sections.

A formula specifies a way for a process to interact with its environment and a proof provides a way to justify this interaction. A judgement $P : A_1, \dots, A_n \vdash B$ then denotes a function that combines n interactions of types A_i for independent processes Q_i into an interaction of type B of the process $Q_1 \mid \dots \mid Q_n \mid P$.

- A modality $\langle a \rangle A$ means doing the action a and then acting according to A . To lighten notations, we will represent successive modalities as a single one: $\langle abc \rangle \alpha$ means $\langle a \rangle \langle b \rangle \langle c \rangle \alpha$.
- Implication $A \multimap B$ is an interaction that provides a behaviour B expecting A from the environment, as made explicit by the rule (mp). The rule (\multimap R) means that some context may actually be provided by the environment.
- An variable α is a behaviour not known from the considered term. An interaction of type α means jumping to a continuation of type α , necessarily provided by the context: indeed, since a scheduling of this type may not provide any behaviour, it effectively gives control to some other process.

As we will formalize later on, the term P on the left side of a judgement is guaranteed to be able to provide the behaviour computed by the proof, and this behaviour will consume all the actions of P . Reciprocally, all the behaviours that consume all actions of P correspond to some proof.

Let us illustrate this by examining the possible ways of typing a term like $a.b.1 \mid c.1$. This term has three possible ways of interacting: each interleaving of the sequence (a, b) with the sequence (c) is a valid trace. A simple interleaving is the sequential execution of one part followed by the other, as (a, b, c) . E.g.

$$\frac{\frac{\frac{\overline{1 : C \vdash C} \text{ (ax)}}{b.1 : C \vdash \langle b \rangle C} \text{ (act)}}{a.b.1 : C \vdash \langle ab \rangle C} \text{ (act)} \quad \frac{\overline{1 : \alpha \vdash \alpha} \text{ (ax)}}{c.1 : \alpha \vdash \langle c \rangle \alpha} \text{ (act)}}{a.b.1 \mid c.d.1 : \alpha \vdash \langle abc \rangle \alpha} \text{ (cut)} \quad \text{with } C = \langle c \rangle \alpha.$$

The important point is the choice of the axiom on C : it stands for the fact the $a.b.1$ finally hands control to $c.1$ for which we have type C .

The interleaving (a, c, b) is more subtle: now $c.1$ will have to get control from $a.b.1$ after a and give back control to it after doing c . We can write this as

$$\frac{\frac{\frac{\overline{1 : \alpha \vdash \alpha} \text{ (ax)}}{b.1 : \alpha \vdash \langle b \rangle \alpha} \text{ (act)} \quad \frac{\overline{1 : C \vdash C} \text{ (ax)}}{b.1 : \alpha, \langle b \rangle \alpha \multimap C \vdash C} \text{ (}\multimap\text{L)}}{a.b.1 : \alpha, \langle b \rangle \alpha \multimap C \vdash \langle a \rangle C} \text{ (act)}}{a.b.1 : \alpha \vdash (\langle b \rangle \alpha \multimap C) \multimap \langle a \rangle C} \text{ (}\multimap\text{R)} \quad \frac{\frac{\overline{1 : B \vdash B} \text{ (ax)}}{c.1 : B \vdash \langle c \rangle B} \text{ (act)}}{c.1 \vdash B \multimap \langle c \rangle B} \text{ (}\multimap\text{R)}}{a.b.1 \mid c.1 : \alpha \vdash \langle acb \rangle \alpha} \text{ (mp)} \quad \text{with } \begin{cases} B = \langle b \rangle \alpha \\ C = \langle cb \rangle \alpha \end{cases}$$

Again, the choice of the right types for the axioms is crucial because it depends on the continuation in interaction. Indeed, we have three steps (a, c, b) and as many types for continuations: $\langle cb \rangle \alpha$, $\langle b \rangle \alpha$ and α .

The other crucial point is the introduction of a in front of $b.1$, as the succession of rules (ax), (\multimap L), (act), (\multimap R). The conclusion type reads as “if using $\langle b \rangle \alpha$ the environment can do C , then, by combining with it, $a.b.1$ can do a then C ”. Operationally, $a.b.1$ starts by doing a , then jumps to C (the behaviour of the environment), and at some point the environment will give control back from C (that is the negative occurrence of C) and $b.1$ will then perform $\langle b \rangle \alpha$. This part is generic in C : we could use the same reasoning for any type C , including a type variable γ . In a more concise way, $(B \multimap \gamma) \multimap \langle a \rangle \gamma$ is an interruptible version of the modality $\langle a \rangle B$. Similarly, the typing of $c.1$ is generic in B . We only need to choose B and C appropriately for the (mp) rule, so that types unify properly.

Another aspect is when parallel composition is typed by a cut which means that a synchronisation (send/receive) happens between the composed processes:

$$\frac{\frac{\overline{1 : \varepsilon \vdash \varepsilon}}{e.1 : \varepsilon \vdash \langle e \rangle \varepsilon} \text{ (act)}}{\frac{\frac{\overline{1 : \alpha \vdash \alpha}}{\bar{e}.1 : \langle e \rangle \alpha \vdash \alpha} \text{ (act)}}{d.\bar{e}.1 : \langle e \rangle \alpha \vdash \langle d \rangle \alpha} \text{ (act)} \quad \frac{\overline{1 : \delta \vdash \delta}}{\bar{d}.1 : \langle d \rangle \delta \vdash \delta} \text{ (act)}}{d.\bar{e}.1 \mid \bar{d}.1 : \langle e \rangle \alpha \vdash \alpha} \text{ (cut)}}{e.1 \mid d.\bar{e}.1 \mid \bar{d}.1 : \alpha \vdash \alpha} \text{ (cut)} \quad \text{with } \begin{cases} \delta = \alpha \\ \varepsilon = \alpha \end{cases}$$

Here the conclusion type is a simple interaction with the environment. This term has different proofs providing the same type, e.g. using an intermediate trace for $e.1 \mid d.\bar{e}.1$ instead of $d.\bar{e}.1 \mid \bar{d}.1$ as in the proof above. Such variants are irrelevant in scheduling and will be removed by switching to proof nets in the next sections.

3 CCS runs as pairings

We consider processes of the standard language CCS [17]. The general language is defined by the following grammar. Note that we use 1 for the inactive process instead of the usual 0 because it is the neutral element of \mid which is a multiplicative operation. Moreover, actions a are decorated by *locations* ℓ :

$$P, Q ::= a^\ell.P \mid \bar{a}^\ell.P \mid 1 \mid (P \mid Q) \mid P + Q \mid *P \mid (\nu a)P$$

where a is taken from an infinite set \mathcal{N} of names and ℓ is taken from an infinite set \mathcal{L} of locations. Each location is used at most once in any term. The main source of non-determinism is the fact that a given action name may occur several times in a given term, and locations are used to name the different occurrences.

For the purpose of the present study, we actually restrict to the following fragment. The reason for this will be explained in the following development.

Definition 2 (MCCS). *Multiplicative CCS is the fragment of CCS using neither choice (+) nor replication (*). Structural congruence is the smallest congruence \equiv that makes parallel composition associative commutative and 1 neutral.*

The set of locations occurring in P is written $\mathcal{L}(P)$. Given $\ell \in \mathcal{L}(P)$, the *subject* of ℓ is the name tagged by ℓ , written $\text{subj}_P \ell$. The *polarity* of ℓ is that of the action tagged by its subject, written $\text{pol}_P \ell$, element of $\{\pm 1\}$. Intuitively, a negative action \bar{a} represents the sending of a signal on a channel a , and a positive action a represents the reception of such a signal.

Definition 3 (execution). *Execution is the relation over structural congruence classes, labelled by partial involutions over \mathcal{L} , defined by the rule*

$$\bar{a}^\ell.P \mid a^m.Q \mid R \rightarrow_{ex}^{\{(\ell, m)\}} P \mid Q \mid R$$

Let \rightarrow_{ex^*} be the reflexive transitive closure of \rightarrow_{ex} , with the annotations defined as $P \rightarrow_{ex^*}^0 P$ and if $P \rightarrow_{ex^*}^c Q \rightarrow_{ex^*}^d R$ then $P \rightarrow_{ex^*}^{c \cup d} R$.

The annotation c in $P \rightarrow_{ex}^c Q$ describes which occurrences interact in the execution step, we write $P \rightarrow_{ex} Q$ if c is unimportant. Similarly, we keep locations implicit when they do not matter. Remark that, for a given P and c , there is at most one Q such that $P \rightarrow_{ex}^c Q$, since c describes the interaction completely.

Definition 4 (pairing). A pairing of a term P is a partial involution c over $\mathcal{L}(P)$ such that for all $\ell \in \text{dom } c$, $\text{subj } c(\ell) = \text{subj } \ell$ and $\text{pol } c(\ell) = -\text{pol } \ell$.

Let \sim_c be the smallest equivalence that contains c . Let \leq_P be the partial order over $\mathcal{L}(P)$ such that $\ell <_P m$ for every subterm $x^\ell.Q$ of P with $m \in \mathcal{L}(Q)$. c is consistent if $\text{dom } c$ is downward closed for \leq_P and $\sim_c <_P \sim_c$ is acyclic.

Example 1. The total pairings of $P = a^1.c^2 \mid b^3.\bar{a}^4 \mid \bar{b}^5.\bar{c}^6 \mid a^7.\bar{b}^8 \mid b^9 \mid \bar{a}^0$ are $c_1 = \{(9, 5), (1, 0), (2, 6), (3, 8), (4, 7)\}$, $c_2 = \{(3, 5), (1, 4), (2, 6), (7, 0), (9, 8)\}$, $c_3 = \{(1, 4), (3, 8), (7, 0), (9, 5), (2, 6)\}$, $c_4 = \{(1, 0), (3, 5), (7, 4), (9, 8), (2, 6)\}$. Only c_1 is inconsistent as there is a cycle induced by $\{(3, 8), (4, 7)\}$. The maximal consistent pairing included in c_1 is $\{(9, 5), (1, 0), (2, 6)\}$.

Observe that pairings and consistency are preserved by structural congruence, as a direct consequence of the fact that subjects, polarities and prefixing are preserved by structural congruence.

Proposition 1. A pairing c of a term P is consistent if and only if there is a term Q such that $P \rightarrow_{ex^*}^c Q$.

Proof (sketch). In an execution $P_0 \rightarrow_{ex}^{c_1} P_1 \rightarrow_{ex}^{c_2} \dots \rightarrow_{ex}^{c_n} P_n$, the c_i are disjoint, so their union is a pairing, and consistency is ensured by the fact that executions respect prefixing. Conversely, write $c = c_1 \uplus \dots \uplus c_n$ with the c_i atomic. By definition, if c is consistent then \leq_P induces a partial order over the domains of the c_i . Assume that the considered enumeration respects this order, then we can prove by recurrence that there is an execution sequence $P = P_0 \rightarrow_{ex}^{c_1} P_1 \dots \rightarrow_{ex}^{c_n} P_n$, since each c_i joins two actions of P_{i-1} that are minimal for $\leq_{P_{i-1}}$.

We easily get the following (for proof see appendix B.1).

Proposition 2. Let P be a term. Any two executions $P \rightarrow_{ex^*}^c Q$ and $P \rightarrow_{ex^*}^c R$ with the same pairing are permutations of each other, and in this case $Q \equiv R$.

We will thus consider consistent pairings as the proper notion of execution for CCS terms. Maximal consistent pairings represent executions of processes until a state where no more execution is possible.

A useful tool in the study of pairings is the following notion of determinisation, by which we can turn a pairing of a term into a term that has no other pairing. In other words, determinisation is a way to represent a run of a term in the language of M CCS itself.

Definition 5 (deterministic term). A term P is deterministic if it has at most one occurrence of each action.

The pairings of a deterministic term form a lattice, consistent pairings too, so there is a unique maximal consistent pairing for any deterministic term.

The restriction operator (νa) serves two purposes: it limits the scope of a name, and it makes it possible to have names local to each copy of a subterm in the presence of replication; both these features are useless in the deterministic case, hence we leave it out on determinisation. We abide by Barendregt's convention that each bound channel is named distinctly from each other channel.

Definition 6 (determinisation). Assume an injective map $\delta : \mathcal{N} \times \{\pm 1\} \times \mathcal{L} \rightarrow \mathcal{N}$. Given a partial involution c , determinisation along c is the operator ∂_c which commutes with parallel composition such that $\partial_c((\nu a)P) = \partial_c(P)$ and

$$\partial_c(a^\ell.P) = \delta(a, +1, \ell)^\ell.\partial_c(P), \quad \partial_c(\bar{a}^\ell.P) = \begin{cases} \overline{\delta(a, +1, \ell)^\ell}.\partial_c(P) & \text{if } \ell \in \text{dom } c, \\ \delta(a, -1, \ell)^\ell.\partial_c(P) & \text{otherwise.} \end{cases}$$

By construction, $\partial_c(P)$ is deterministic, the pairings of $\partial_c(P)$ are the restrictions of c , consistency preserved, so c is the unique maximal pairing of $\partial_c(P)$.

Example 2. For the term P and the pairings of example 1, we obtain the following determinisations (with $\delta(a, +1, 7) = d$ and $\delta(b, +1, 9) = e$):

$$c_3 = \{(1, 4), (3, 8), (7, 0), (9, 5), (2, 6)\} \text{ induces } \partial_{c_3}(P) = a.c \mid b.\bar{a} \mid \bar{e}.\bar{c} \mid d.\bar{b} \mid e \mid \bar{d},$$

$$c_4 = \{(1, 0), (3, 5), (7, 4), (9, 8), (2, 6)\} \text{ induces } \partial_{c_4}(P) = a.c \mid b.\bar{d} \mid \bar{b}.\bar{c} \mid d.\bar{e} \mid e \mid \bar{a}.$$

If we extended our study to the whole of CCS, determinisations would still be in MCCS, but the theory of pairings would have to be refined: external choice requires a notion of conflict in the space of locations (as in event structures [19]), replications requires the introduction of indices to distinguish copies.

4 Proof nets for MLL with action modalities

Proofs in sequent calculus are well suited to inductive reasoning, however their use in proof theory is uneasy because their rigid structure obscures many arguments, like those below in particular. For this reason, we will turn to proof nets, using the standard machinery of linear logic [10, 5]. Modality rules are represented using boxes (like promotions in standard linear logic, but with different typing rules). The only extra information we add to standard proof structures is the location of each box, to reflect the use of locations in CCS terms in the sequel. For readers not familiar with the standard definitions of proof structures and proof nets, these are put in appendix. We detail here specificities of MLL_a .

Definition 7 (proof structure). A proof structure consists of an ordered forest of nodes labelled by formulas, denoted x^A , with a set Ax of axiom links (pairs of leaves), a set Cut of cuts (pairs of roots) and a set Box of modality boxes, labelled by action modalities, such that each box β has a unique location $\ell(\beta)$. The roots that are not part of a cut are called the conclusion nodes. The conclusion type is the multiset of the labels of the conclusion nodes.

A modality box β is a set of nodes (the ports) associated to a proof structure S whose conclusions are in bijection with the ports. If the modality of β is $\langle a \rangle$, then the *principal port* is labelled $\langle a \rangle A$ and matches a conclusion of S labelled A , while *auxiliary ports* have the same label as their matching conclusion in S .

The graphical notation of proof structures is presented in figure 1. By definition there are arcs only to multiplicative nodes, moreover proof structures can be drawn considering the top-bottom orientation of arcs, so we keep arc orientation




Fig. 1. Representation of proof structures: axiom link, \wp node, \otimes node, boxes, cut.

implicit by this convention. Arcs to a \wp node are joint by a circle on the side of this node. By construction, the conclusion labels suffice to deduce all labels, so we keep most of this information implicit.

Definition 8 (proof net). A proof net is a proof structure built following MLL_a sequent calculus rules. An immediate subnet of a proof net π is an induced subgraph of π that is itself a proof net. A subnet of π is either an immediate subnet of π or (inductively) a subnet of a box of π .

Well known correctness criteria [5, 10, 6] apply to characterise proof nets among proof structures by combinatorial means like acyclicity and connectedness.

Definition 9 (cut elimination). Annotated cut elimination is the relation \rightarrow_{ce}^c over proof structures, labelled by partial involutions c over \mathcal{L} , that is the reflexive transitive closure of the rules below (such that if $\pi \rightarrow_{ce}^c \pi' \rightarrow_{ce}^d \pi''$ then $\pi \rightarrow_{ce}^{c \cup d} \pi''$). We have $\pi \rightarrow_{ce}^c \pi'$ if π contains a cut $\kappa = \{x, y\}$ either at top level or inside a box and one of the following cases occurs:

- Multiplicative step and Axiom step: standard definition, with $c = \emptyset$.
- Modality step: If x and y are principal ports of two boxes β, β' , then c permutes $\ell(\beta)$ and $\ell(\beta')$ and π' is obtained by replacing each box with its associated proof structure.
- Commutation step: If x is the auxiliary port of a box β , then $c = \emptyset$, and the cut and a subnet of π that contains y are moved inside β .

Our proof system enjoys a standard cut-elimination theorem using this definition: if $\pi \rightarrow_{ce}^c \pi'$ and π is a proof net, then π' is a proof net with the same conclusion (this is proved by standard arguments using correctness criteria, hence we will not develop this point); if a proof π is irreducible by \rightarrow_{ce} , then it has no cut link (this is an immediate case analysis). Note however that \rightarrow_{ce} is not confluent, because of commutation steps.

Definition 10 (head reduction). Head reduction is the annotated relation \rightarrow_h^c over proof structures defined as the restriction of \rightarrow_{ce}^c that only applies at top level and does not use the commutation step of cut elimination.

This particular strategy is relevant because it does not reduce inside boxes, that is under prefixes, it only affects cuts in active position (from the point of view of processes). However, this strategy does not eliminate all cuts in general.

In the analysis of proofs, the following notion of path will be useful. It describes a way to traverse arcs and axioms/cuts in a proof structure while respecting the logical meaning of formulas.

Definition 11 (path). *A path in a proof structure S is an alternating path in the underlying graph of S , such that alternations occur only at axioms, cuts and boxes. Each move between ports x and y of a box β must be associated with a path between the corresponding conclusions in β . We further require a typing constraint: a path can only move up a left (resp. right) branch if has moved down a left (resp. right) branch before, with a natural well-bracketing condition.*

For instance, a path starting from an axiom with type α may move down the tree of nodes, reach a cut, move up the other side of the cut, always in the branches that contain α , reach an axiom, and so on.

5 Typing executions of MCCS terms

Proofs in MLL_a will serve as a type system. Although this can be formulated in usual sequent style (as in table 1), the natural notion rather relates proof nets and structural congruence classes of terms.

Definition 12 (term assignment). *Let S be a proof structure. The MCCS term $\llbracket S \rrbracket$ assigned to π is the parallel composition of the $\llbracket \beta \rrbracket$ for each box β in S . In turn, for a box β with location ℓ and associated structure S_β , the term $\llbracket \beta \rrbracket$ is $a^\ell \cdot \llbracket S_\beta \rrbracket$ if the principal port of β has modality $\langle a \rangle$ and $\bar{a}^\ell \cdot \llbracket S_\beta \rrbracket$ if the principal port of β has modality $\langle \bar{a} \rangle$. A term P is said to have type Γ if there is a proof net π of conclusion Γ such that $\llbracket \pi \rrbracket \equiv P$. In this case we write $\pi : P \vdash \Gamma$.*


A proof net is a proof structure that is built using the rules of table 1, ignoring the terms on the left of the \vdash symbols. It is obvious that these terms do reflect the definition of term assignment: A term P has type Γ if and only if there is a type derivation with conclusion $P \vdash \Gamma$ using the rules of table 1.

We now establish the correspondence between cut elimination in a proof and execution steps in the assigned terms. The first result justifies head reduction:

Proposition 3. *Let π be a proof structure. For every head reduction $\pi \rightarrow_h^c \pi'$ there is an execution $\llbracket \pi \rrbracket \rightarrow_{ex^*}^c \llbracket \pi' \rrbracket$.*

Proof (sketch). Axiom and multiplicative cut elimination steps do not affect the assigned terms, besides their annotation is empty, so the result holds immediately for them. When a modality step applies, it reduces a cut between boxes with dual modalities (because of typing), hence the associated terms are ready to interact; the reduct is easily seen to be the assigned term of the reduct proof.

Example 3. Let π be the following proof net.



We have $[\pi] = a.c | b.\bar{a} | \bar{e}.\bar{c} | d.\bar{b} | e | \bar{d}$. (It is $\partial_{c_3}(P)$ of previous examples). As it is deterministic term, we abusively identify locations with names. We consider the head reduction sequence $\pi \rightarrow_h^z \pi'$ (where π' is an axiom link) for $z = \{(d, \bar{d}), (b, \bar{b}), (a, \bar{a}), (e, \bar{e}), (c, \bar{c})\}$. We have $[\pi] \rightarrow_{ex^*}^z [\pi'] \equiv 1$.

Subject reduction does not hold in general. Indeed, a given proof may hold several occurrences of a given modality, corresponding to different occurrences of an action in the term, and the structure of cuts may not match a given execution step. This is not a defect, since we actually intend to type pairings rather than processes: we do get subject reduction if we restrict to proofs that describe deterministic terms.

Definition 13 (linear proof). A proof structure S is called linear if


- S contains at most one box for each modality,
- for each a , all occurrences of $\langle \bar{a} \rangle A$ in the labels in S have the same immediate subformula A , and if $\langle \bar{a} \rangle A$ and $\langle a \rangle B$ occur then A and B are dual,
- if S contains a box for both $\langle a \rangle A$ and $\langle \bar{a} \rangle A^\perp$, then neither formula occurs in the conclusion type of S .

The essence of the linearity condition is the first constraint. Intuitively, the second and third constraints serve to guarantee that the property is preserved by composition. Indeed, if a formula $\langle a \rangle A$ occurs in the conclusion of a proof π , then the proof may be cut against a proof that contains a modality box for $\langle \bar{a} \rangle A^\perp$, which breaks linearity if π already contains a box for some $\langle \bar{a} \rangle B$. Note that the fact of being a linear proof is preserved by cut elimination.

Theorem 1 (subject reduction). Let π be a linear proof of conclusion $P \vdash \Gamma$. For every execution $P \rightarrow_{ex^*}^c P'$ there is a linear proof $\pi' : P' \vdash \Gamma$.

Proof (sketch). An execution step $[\pi] \rightarrow_{ex}^{(\ell, m)} P$ involves immediate subterms $a^\ell.Q$ and $\bar{a}^m.R$ for $a \in \mathcal{N}$. Then π contains two top level boxes with respective principal ports $x^{\langle a \rangle A}$ and $y^{\langle \bar{a} \rangle A^\perp}$, for $A \in MLL_a$. Since π is linear, x and y are elimination boxes for each other, ending a path ρ (as of definition 11) whose axioms contain modalities of x and y in their types. Let π' be the rewriting of π where such modalities are removed (boxes are replaced by their contents, axioms on $\langle a \rangle A$ by axioms on A). Clearly π' is a linear proof of conclusion $P' \vdash \Gamma$.

This theorem states that types are preserved by execution in deterministic terms. However, the proof uses a rewriting of the typing proofs that does not correspond to cut elimination in general. Indeed, consider the following example of typing, call π the l.h.s.:



Then the proof is linear, irreducible by head cut elimination, but the assigned term $[\pi] = \bar{a} \mid \bar{b} \mid a$ does execute into \bar{b} . In π , this involves a cut on the axiom inside the middle box. As done in theorem 1 the rewriting of π in a linear proof π' assigned to \bar{b} is the r.h.s..

We can get a precise correspondence between execution and head cut elimination by imposing an additional constraint on the shape of proofs. In the statement below, an axiom is *immediately contained* in a box if it is an immediate subnet of the structure associated with this box.

Definition 14 (regular proof). *An axiom link immediately contained in a box β is anchored if there is a path from one of its conclusions to an auxiliary port of β and a path from its other conclusion to the principal port. A proof structure π is regular if all its axioms are anchored and for every pair of boxes with dual modalities, one of the boxes does not immediately contain any axiom.*

Theorem 2 (strong subject reduction). *Let π be a regular linear proof net. For every execution $[\pi] \rightarrow_{ex}^c P$ there is a regular linear proof π' such that $\pi \rightarrow_h^c \pi'$ and $[\pi'] = P$.*

Proof (sketch). Consider an execution step $[\pi] \rightarrow_{ex}^{(\ell, m)} P$. As in the proof of theorem 1, linearity implies that there are boxes at top level and a path ρ between their principal ports $x^{(a)A}$ and $y^{(\bar{a})A^\perp}$ for immediate subterms $a^\ell.Q$ and $\bar{a}^m.R$ of $[\pi]$. Since x is cut at top level, ρ traverses no box, otherwise linearity or regularity would be contradicted. Then ρ is a multiplicative cut path whose cut elimination \rightarrow_h^\emptyset until x and y preserves $[\pi]$ as well as regularity and linearity.

6 Anti-execution and completeness

In this section we establish our correspondence theorem relating typings and executions. To achieve this goal we first provide a kind of reciprocal statement for subject reduction: if a term T can reduce into a typed term T' , then we can type T with a proof that reduces to the typing of T' . Because we want logically correct proof structures, this operation requires some care.

Example 4. Consider the term $P := a.\bar{b} \mid b.\bar{c} \mid \bar{a}.c$. We cannot type each thread with a simple type like $\langle a \rangle \alpha$, $\langle \bar{b} \rangle \alpha^\perp$ and then introduce a cut for each interaction, since we would get a cyclic proof structure, what is incorrect.


We now describe a general method for deducing a typing by “anti-execution” of a proof. We stay at a partly informal level for clarity, all formal statements are detailed in appendix B.3.

Consider a generic execution step $P \mid a.Q \mid \bar{a}.R \rightarrow_{ex} P \mid Q \mid R$. Assume the reduct is typed by some proof π . We want to put the parts of π that correspond to Q and R into boxes, with a cut between them, while rewriting the proof to avoid cycles. For this purpose, we proceed in four steps:


Selection consists in moving each box belonging to Q or R away from the main proof, by means of an axiom/cut pair, so that Q and R are represented by simple sets of boxes, cut with the main proof (which corresponds to P), with no multiplicative connectives:



Chaining consists in introducing an extra axiom/cut pair in the middle of each cut between P and R , so that there are cuts only between P and Q or Q and R , and not between P and R directly:




Simplification consists in making sure that there is actually exactly one cut between P and Q and one between Q and R , by multiplexing multiple cuts through multiplicative:



Correctness criteria guarantee that we can always find two cuts for which there is one connected component on one side, two on the other.


Boxing consists in putting Q and R into boxes, cut together, so that Q has one auxiliary port to P and R has no auxiliary port:



Following this method, we prove the following statement:

Proposition 4 (anti-execution). *Let $T_1 \rightarrow_{ex}^c T_2$ be an execution step and let $\pi_2 : T_2 \vdash \Gamma$ be a typing. There exists a typing $\pi_1 : T_1 \vdash \Gamma$ such that $\pi_1 \rightarrow_h^c \pi_2$.*

Example 5. Consider the term of P of example 1. We consider the execution $e = (a, \bar{a})(b, \bar{b})(c, \bar{c})(d, \bar{d})(e, \bar{e})$ of the determinized term $\partial_{c_4}(P) = a.c|b.\bar{d}|\bar{b}.\bar{c}|d.\bar{e}|e|\bar{a}$ for the (total and consistent) pairing $c_4 = \{(1, 0), (3, 5), (7, 4), (9, 8), (2, 6)\}$. A typing synthesized by the construction of proposition 4 is the following.



Lemma 1 (preserved regularity). *In the construction of proposition 4, if π_2 is regular, then so is π_1 . If π_2 is linear and T_2 is deterministic, then π_1 is linear.*

Proof (sketch). Let $T_2 = P \mid Q \mid R$. If an axiom is introduced by anti-execution rewrite steps, used in proposition 4 then: i) it is added to P by selection and it will not be boxed, or ii) it is added to Q by chaining and becomes anchored by simplification and boxing. No axiom is introduced on the side of R , Q only contains chaining axioms, so regularity is satisfied for the new axioms. Besides, regularity is not broken for axioms previously present in the proof.

Example 6. In the previous example 5, one can also start execution by $(b, \bar{b})(a, \bar{a})$ as seen in the typing. All execution permutation of $\partial_{c_4}(P)$ in the pairing c_4 is allowed by the typing proof synthesized from the execution e .

We now summarize the previous results, about subject reduction and the reverse operation, into a precise statement relating typings and execution.

Lemma 2 (initial typing). *Every linear MCCS term where no name occurs with both modalities is typable by a cut-free regular proof.*

Proof. We simply build a proof of $T \vdash A_T, B_T$ with A_T non-modal by induction on T . For $T = 1$, use the axiom rule to get $1 \vdash \alpha^\perp, \alpha$. For $T = P \mid Q$, deduce $T \vdash A_P \wp A_Q, B_P \otimes B_Q$ by the tensor rule. For $T = a.P$, deduce $T \vdash A_P, \langle a \rangle B_P$ by the action rule, similarly for $\bar{a}.P$. The proof thus built is obviously regular since every axiom is at top level or anchored, and there are no pairs of boxes with dual modalities.

Theorem 3 (completeness). *For every execution $P \rightarrow_{ex}^c Q$ there are typings $\pi_P : P \vdash \Gamma$ and $\pi_Q : Q \vdash \Gamma$ such that $\pi_P \rightarrow_h^c \pi_Q$. Moreover, for every execution sequence $P \rightarrow_{ex}^{c_1} P_1 \cdots \rightarrow_{ex}^{c_n} P_n = Q$ with $c_1 \cup \cdots \cup c_n = c$, there is a cut elimination sequence $\pi_P \rightarrow_h^{c_1} \pi_1 \cdots \rightarrow_h^{c_n} \pi_n = \pi_Q$, with $[\pi_i] = P_i$ for all i .*

Proof. By definition, the term $\partial_c(Q)$ is linear and has no dual actions, so by lemma 2 we can find a cut-free regular proof $\pi'_Q : \partial_c(Q) \vdash \Gamma$. If we apply proposition 4 repeatedly to π'_Q with the steps of the considered execution $\partial_c(P) \rightarrow_{ex}^c \partial_c(Q)$, we get a proof $\pi'_P : \partial_c(P) \vdash \Gamma$ that reduces to π'_Q by a head reduction sequence labelled c . Let π_P and π_Q be the relabellings of π'_P and π'_Q by the inverse of ∂_c , then we have $\pi_P : P \vdash \Gamma$, $\pi_Q : Q \vdash \Gamma$ and $\pi_P \rightarrow_h^c \pi_Q$.

Every execution sequence of P with label c is an execution sequence of $\partial_c(P)$ with the same label. By lemma 1, π'_P enjoys strong subject reduction as of theorem 2, hence every run of $\partial_c(P)$ labelled by c corresponds to a head reduction sequence in π'_P labelled by c . By relabelling with ∂_c^{-1} , every run of P labelled by c corresponds to a head reduction sequence $\pi_P \rightarrow_h^c \pi_Q$.

In other words, every execution of a term can be exactly characterized up to permutation by typing, in the sense that the execution sequences of the term within the same pairing will be exactly the head reduction sequences of the associated typing proof. By combining determinisation (definition 6) and strong subject reduction (theorem 2) we get that, conversely, each regular typing of a term defines a set of executions stable by permutation.

7 Conclusion and further works

In this work we have developed, in the simple framework of multiplicative CCS, a precise logical description of executions of processes. A key technical tool is the use of pairings, by which we separate non-determinism in communication from the multiplicity of equivalent schedulings; this technique extends well to more expressive frameworks (full CCS, π -calculus, etc.). The logical interpretation we propose moves beyond the traditional Curry-Howard for concurrency by accepting non-deterministic terms, albeit with a change of interpretation in the correspondence. Indeed, the logic we use is well studied and has a wide range of existing tools (efficient correctness criteria, proof search, etc.) but its interpretation in our paradigm of proof-as-executions is new.

Logical expressiveness The restriction to purely multiplicative objects, in M CCS and MLL, lets us concentrate on the precise role of multiplicatives and axioms as descriptions of how a process interacts with its environment, but it hides the complexity inherent to the other defining features of concurrent systems like choice, recursion, name passing, etc. It should be stressed that extending the calculus or the logic are two different things.

On the one hand, extending the calculus enriches the set of possible executions, by introducing more subtle synchronization possibilities: choice allows for conflict between actions, replication allows for arbitrarily large runs with some uniformity, value passing allows for communication of ground values, name passing allows the set of synchronizable pairs to evolve along execution. After determinisation, all these features essentially disappear and deterministic runs can still be formulated in M CCS. On the other hand, enriching the logic leads to richer descriptions of the control flow in processes, for instance using a first order language with predicates to describe properties of continuations. Furthermore, new connectives allow a given type to correspond to more distinct executions.

These two kinds of extensions are not independent, however, since each feature of the calculus can be usefully described using a feature of the logic. Let us illustrate this in the case of choice and additives. The technique of pairings still works, consistency simply needs to take into account a notion of conflict as in event structures [19]. The type system is naturally extended by additive rules:

$$\frac{P \vdash \Gamma, A \quad Q \vdash \Gamma, B}{P + Q \vdash \Gamma, A \& B} \quad \frac{P \vdash \Gamma, A}{P \vdash \Gamma, A \oplus B}$$

possibly with the restriction that A and B are modal. A type $A \& B$ for a behavior means either behaving as A or as B , according to what the environment provides, while $A \oplus B$ means behaving as A or B depending on one's own choice.

With this we can type useful processes that use choice. For instance, describe a boolean on names t, f as some process that will send a signal on one of the channels t, f . This can be materialized by the type $B(t, f) := \alpha^\perp, \langle \bar{t} \rangle \alpha \oplus \langle \bar{f} \rangle \alpha$ which reads like “give me control (using α), I will terminate by a signal on t or f ”. Then consider a negation function: $N := t.\bar{f}' + f.\bar{t}'$. By studying its

interactions with the environments $E_1 := \bar{t} \mid f'$ and $E_2 := \bar{f} \mid t'$, we see that both the $N \mid E_i$ have complete consistent pairings, hence we can type them as α^\perp, α . Extracting the types of N we get $\langle t \rangle \alpha^\perp, \langle \bar{f}' \rangle \alpha$ and $\langle f \rangle \alpha^\perp, \langle \bar{t}' \rangle \alpha$, which we combine by additives into a unique type $\langle t \rangle \alpha^\perp \& \langle f \rangle \alpha^\perp, \langle \bar{t}' \rangle \alpha \oplus \langle \bar{f}' \rangle \alpha$. This way we get a possible specification for N .

Causality A crucial feature of our work is the interpretation of axioms as a way to transfer causality. An effect is that, most of the time, the type of a term will contain modalities for actions that it does not contain by itself. For instance, $a.\bar{b}$ may have type $\langle a \rangle \langle c \rangle \alpha^\perp, \langle \bar{b} \rangle \langle \bar{c} \rangle \alpha$, which can be read “give me a signal on a with the promise of a signal on c , and I will answer with a signal on b and the promise of a signal c ”. This makes it explicit that this part of interaction will be involved in the triggering of interaction on c , but only indirectly by allowing bearers of c to get active. This idea suggests new ways of analyzing causality in interactive systems, and the fact that the flow of causality is often as complicated as the flow of information. Besides, a similar fact is illustrated by the expressiveness of solos [13, 3], where communication is used to carry all prefixing information in processes. Our interpretation may provide a logical insight on this matter.

References

1. S. Abramsky. Proofs as processes. *TCS*, 135(1):5–9, 1994.
2. E. Beffara. A concurrent model for linear logic. *ENTCS*, 155:147–168, 2006.
3. E. Beffara and F. Maurel. Concurrent nets: a study of prefixing in process calculi. *TCS*, 356(3):356–373, 2006.
4. L. Caires and F. Pfenning. Session types as intuitionistic linear propositions. In P. Gastin and F. Laroussinie, editors, *CONCUR*, volume 6269 of *LNCS*, pages 222–236. Springer, 2010.
5. V. Danos and L. Regnier. The structure of multiplicatives. *Archive for Math. Logic*, 28(3):181–203, 1989.
6. P. Jacobé de Naurois and V. Mogbil. Correctness of linear logic proof structures is NL-complete. *TCS*, 412(20):1941–1957, 2011.
7. T. Ehrhard and O. Laurent. Interpreting a finitary π -calculus in differential interaction nets. In L. Caires and V. Vasconcelos, editors, *CONCUR*, volume 4703 of *LNCS*, pages 333–348. Springer, 2007.
8. T. Ehrhard and L. Regnier. Differential interaction nets. *TCS*, 364(2):166–195, 2006.
9. J.-Y. Girard. Linear logic. *TCS*, 50(1):1–102, 1987.
10. J.-Y. Girard. Proof-nets : the parallel syntax for proof theory. *Logic and Algebra*, 180, 1996.
11. M. Hennessey and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the ACM*, 32(1):137–161, jan 1985.
12. K. Honda and O. Laurent. An exact correspondence between a typed π -calculus and polarised proof-nets. *TCS*, 411(22–24):2223–2238, 2010.
13. C. Laneve and B. Victor. Solos in concert. In J. Wiederman, P. Boas, and M. Nielsen, editors, *ICALP*, volume 1644, pages 513–523. Springer Verlag, 1999.
14. H. Mairson and K. Terui. On the computational complexity of cut-elimination in linear logic. In C. Blundo and C. Laneve, editors, *ICTCS*, volume 2841 of *LNCS*, pages 23–36. Springer, 2003.

15. F. Maurel. Nondeterministic light logics and NP time. In M. Hofmann, editor, *TLCA*, volume 2701 of *LNCS*, pages 241–255. Springer, 2003.
16. K. Mazurak and S. Zdancewic. Lollipop: to concurrency from classical linear logic via curry-howard and control. In *ICFP*, pages 39–50, 2010.
17. R. Milner. *Communication and concurrency*. Prentice Hall, 1989.
18. V. Mogbil. Non-deterministic boolean proof nets. In M. van Eekelen and O. Shkaravska, editors, *FOPARA*, volume 6324 of *LNCS*, pages 131–145. Springer, 2010.
19. G. Winskel. Event structures. In *Advances in Petri nets: applications and relationships to other models of concurrency*, pages 325–392. Springer Verlag, 1987.
20. N. Yoshida, M. Berger, and K. Honda. Strong normalisation in the π -calculus. In *LICS*, pages 311–322, 2001.

A MLL_a and restricted type system

A.1 MLL with action modalities

Proof structures of MLL_a are formally defined as follows:

- The forest is seen as an acyclic graph (V_S, E_S) oriented from leaves to roots, with a total order on the ingoing arcs of each node.
 - We write v^A to signify that vertex v has label A . We impose that if a vertex v^A is not a leaf, then it has two children x^B and y^C (in this order) and either $A = B \otimes C$ or $A = B \wp C$; the main connective of A is the *sort* of v .
 - Axioms are pairwise disjoint pairs of leaves labelled by dual formulas.
 - Cuts are disjoint pairs of roots labelled by dual formulas.
 - Boxes are disjoint non-empty sets of leaves,
 - $Ax \cup Box$ forms a partition of the set of leaves.
 - Each box β has a distinguished element called its *principal port*, the others are *auxiliary ports*. Each box is associated (inductively) to a proof structure S_β with a bijection ϕ from the elements of β to the conclusions of S_β .
 - For the principal port x^A of β we have $A = \langle a \rangle B$ or $A = \langle \bar{a} \rangle B$ where B is the label of $\phi(x)$;
 - for each other element y of β , the labels of y and $\phi(y)$ are the same.
- To each box β we also associate a location $\ell(\beta)$, with the constraint that each location is used at most once in any given structure.

The roots that are not part of a cut are called the *conclusion nodes* of S . The *conclusion type* of S is the multiset of the labels of its conclusion nodes.

Proof nets (PN) are formally defined as follows:

- (ax): $(\{u_{ax}^A, v_{ax}^{A^\perp}\}, \{uv\}, \emptyset, \emptyset, \{u, v\})$ is a PN.
- (\wp): If $G = (V, E, A, P, C)$ is a PN and u^A, v^B are two conclusions of G , then $(V \uplus \{w_{\wp}^{A \wp B}\}, E, A \cup \{uw, vw\}, P \cup \{\{uw, vw\}\}, C \setminus \{u, v\} \cup \{w\})$ is a PN.
- (\otimes): If $G = (V, E, A, P, C)$ and $G' = (V', E', A', P', C')$ are disjoint PNs, u^A is a conclusion of G and v^B is a conclusion of G' , then $(V \uplus V' \uplus \{w_{\otimes}^{A \otimes B}\}, E \uplus E', A \uplus A' \uplus \{uw, vw\}, P \uplus P', (C \setminus \{u\}) \uplus (C' \setminus \{v\}) \uplus \{w\})$ is a PN.
- (act): If $G = (V, E, A, P, C)$ is a PN with conclusions set $C = \{u^A, v_1^{B_1}, \dots, v_k^{B_k}\}$, then $(\{x_{\langle a \rangle}^{\langle a \rangle A}(G), y_1^{B_1}, \dots, y_k^{B_k}\}, \{uv_1, \dots, uv_k\}, \emptyset, \emptyset, \{u, v_1, \dots, v_k\})$ is a PN. Also when changing $\langle a \rangle A$ by $\langle \bar{a} \rangle A$.
- (cut): If $G = (V, E, A, P, C)$ and $G' = (V', E', A', P', C')$ are disjoint PNs, u^A is a conclusion of G and v^{A^\perp} is a conclusion of G' , then $(V \uplus V', E \uplus E' \cup \{uv\}, A \uplus A', P \uplus P', (C \setminus \{u\}) \uplus (C' \setminus \{v\}))$ is a PN.

Cut elimination in MLL_a is formally defined as follows. Let π and π' be a proof structure and c be a partial involution on \mathcal{L} . *Annotated cut elimination* is the relation \rightarrow_{ce}^c over proof structures, labelled by partial involutions c over \mathcal{L} , that is the reflexive transitive closure of the rules below (such that if $\pi \rightarrow_{ce}^c \pi'$ then $\pi \rightarrow_{ce}^d \pi''$ then $\pi \rightarrow_{ce}^{c \cup d} \pi''$). We have $\pi \rightarrow_{ce}^c \pi'$ if π contains a cut $\kappa = \{x, y\}$ either at top level or inside a box and one of the following cases occurs (note that x and y may be freely exchanged):

Axiom and cut rules (A is a literal α or a MLL_a formula):

$$\frac{}{1 \vdash A^\circ, A^\perp} \text{ (ax)} \qquad \frac{P \vdash \Gamma, A^\circ \quad Q \vdash A^\perp, \Delta}{P \mid Q \vdash \Gamma, \Delta} \text{ (cut)}$$

Multiplicative rules:

$$\frac{P \vdash \Gamma, A^p, B^p}{P \vdash \Gamma, (A \wp B)^p} \text{ (\wp)} \qquad \frac{P \vdash \Gamma, A^p \quad Q \vdash B^p, \Delta}{P \mid Q \vdash \Gamma, (A \otimes B)^p, \Delta} \text{ (\otimes)}$$

Modality and New rules:

$$\frac{P \vdash \Gamma, A^\circ}{a^\ell . P \vdash \Gamma^\gamma, (\langle a \rangle A)^p} \text{ (act)} \qquad \frac{P \vdash \Gamma \quad a \notin \Gamma}{(\nu a) P \vdash \Gamma} \text{ (new)}$$

Table 2. Inference rules in MLL_a^p

- Multiplicative step: If x and y have respective sorts \otimes and \wp , then each has two premises, call them respectively $x_1^A, x_2^B, y_1^{A^\perp}, y_2^{B^\perp}$. Then $c = \emptyset$ and π' is obtained by removing κ and the nodes x and y and adding the cuts $\{x_1, y_1\}$ and $\{x_2, y_2\}$.
- Axiom step: If y is a leaf node and it is part of an axiom $\alpha = \{y, z\}$ with $x \neq z$, then $c = \emptyset$ and π' is obtained removing α, κ, y and z and rewriting any outgoing arc of z into an outgoing arc of x .
- Modality step: If x and y are principal ports of two boxes β, β' , then c permutes $\ell(\beta)$ and $\ell(\beta')$ and π' is obtained by replacing each box with its associated proof structure, identifying the conclusions of this structure with the ports of the box.
- Commutation step: If x is the auxiliary port of a box β , call T the smallest subnet of π that contains y . Then $c = \emptyset$ and π' is obtained by moving T and κ inside β , replacing the auxiliary port x by one auxiliary port for each conclusion of T .

A.2 MLL_a^p , a restriction of MLL_a

The restriction to anchored proofs may be designed with a restriction of our type system as in Table 2. The idea is to enforce in each box an orientation from auxiliary ports to principal port, for axioms on modalities. We simply use a decoration on formulas when needed. The letters p, q, \dots indicate the o decoration (output) or no decoration. The sequence of formulas Γ^γ is a sequence of formulas decorated with letters.

B Detailed proofs

B.1 Runs and pairings

Lemma 3. *Let $P \rightarrow_{ex}^c Q$ be an execution and d be a pairing of Q , then $\text{dom } c \cap \text{dom } d = \emptyset$ and $c \cup d$ is a pairing of P . If d is consistent, then so is $c \cup d$.*

Proof. First remark that, by definition of execution, we have $\mathcal{L}(P) = \mathcal{L}(Q) \uplus \text{dom } c$, besides $\text{dom } d \subset \mathcal{L}(Q)$ so the domains of c and d are disjoint. We can thus define the involution $c' = c \cup d$, and check that it is indeed a pairing of P .

Let $\ell \in \text{dom } c'$. If $\ell \in \text{dom } c$ then ℓ is the location of an action involved in the execution step, so $\text{subj}_P c(\ell) = \text{subj}_P \ell$ and $\text{pol}_P c(\ell) = -\text{pol}_P \ell$ by definition of execution, and subsequently the property holds for c' . Otherwise ℓ is in the domain of d , then the same property holds for d in Q since d is a pairing of Q , and again we get it for c' in P . Hence c' is a pairing of P .

Now suppose d that is consistent but c' is not. Write $\ell \prec_P^c m$ if there is a location n such that $\ell <_P n \sim_c m$. Then there exists a cycle $\ell_0 \prec_P^c \ell_1 \prec_P^c \dots \prec_P^c \ell_k = \ell_0$. If all the ℓ_i are in Q then this cycle exists in \prec_Q^d , which cannot be since d is consistent. Since c annotates a reduction of P , all elements of $\text{dom } c$ are minimal for \leq_P , so the cycle cannot consist only of elements of $\text{dom } c$. So we may assume $\ell_0 \in \mathcal{L}(Q)$ and $\ell_1 \in \text{dom } c$. This means either $\ell_0 <_P \ell_1$ or $\ell_0 <_P c(\ell_1)$, in each case this implies that some location in $\text{dom } c$ is prefixed in P , which is impossible. Hence c' is consistent.


Proof (proposition 1). By iterating lemma 3, from an execution $P_0 \rightarrow_{ex}^{c_1} P_1 \rightarrow_{ex}^{c_2} \dots \rightarrow_{ex}^{c_n} P_n$ we can deduce a pairing $c = c_1 \cup \dots \cup c_n$ of P_0 . This pairing represents the execution above, because it contains all the choices made during this execution. Indeed we can prove that executions that yield the same pairing are equivalent. The converse is detailed in the main text.

Lemma 4. *Let $P \rightarrow_{ex}^{c_1} Q_1$ and $P \rightarrow_{ex}^{c_2} Q_2$ be two executions with $\text{dom } c_1 \cap \text{dom } c_2 = \emptyset$. Then there is a unique R such that $Q_1 \rightarrow_{ex}^{c_2} R$ and $Q_2 \rightarrow_{ex}^{c_1} R$.*

Proof. By the existence of the execution step $P \rightarrow_{ex}^{c_1} Q_1$, we know that P can be written $P \equiv \bar{a}^\ell . S | a^m . T | P'$ for some name a and with $\text{dom } c_1 = \{\ell, m\}$. Similarly, the term P can be decomposed as $P \equiv \bar{b}^{\ell'} . S' | b^{m'} . T' | P''$ for some name b and with $\text{dom } c_2 = \{\ell', m'\}$ to justify the execution step $P \rightarrow_{ex}^{c_2} Q_2$. By hypothesis the domains of c_1 and c_2 are disjoint, so ℓ and m are distinct from ℓ' and m' . As a consequence the term P can be decomposed as $P \equiv \bar{a}^\ell . S | a^m . T | \bar{b}^{\ell'} . S' | b^{m'} . T' | P'''$ and the terms Q_1 and Q_2 have execution steps with the expected annotations, with the common reduct $R = S | T | S' | T' | P'''$. Unicity of R up to structural congruence is a consequence of the fact the c_1 and c_2 completely describe which subterms of P, Q_1, Q_2 interact and in which way.

Proof (proposition 2). The pairing c is the disjoint union of the atomic involutions of each step, so clearly the execution sequences are permutations of each other. Write them as $P = P_0 \rightarrow_{ex}^{c_1} P_1 \dots \rightarrow_{ex}^{c_n} P_n$ and $P = P_0 \rightarrow_{ex}^{c_{\sigma(1)}} P'_1 \dots \rightarrow_{ex}^{c_{\sigma(n)}} P'_n$. We now prove that the final terms P_n and P'_n are equal up to structural congruence. Call $d(\sigma)$ the number of pairs (i, j) such that $i < j$ and $\sigma(i) > \sigma(j)$. We proceed by induction on $d(\sigma)$. If this number is 0, then σ is the identity function and the sequences match, so obviously we have $P_n \equiv P'_n$. Otherwise, consider a minimal i such that $\sigma(i) \neq i$, hence $\sigma^{-1}(i) \neq i$ and $\sigma^{-1}(i) > i$, and let $j = \sigma^{-1}(i) - 1$. The reduction sequences match in their first $i - 1$ steps, then one has a reduction labelled c_i while the other has a reduction labelled

$c_{\sigma(i)}$. By repeated applications of lemma 4, we can deduce that for each $k \geq i$ there is a term Q_k such that $P'_k \rightarrow_{ex}^{c_i} Q_k$ and $Q_{k-1} \rightarrow_{ex}^{c_{\sigma(k)}} Q_k$ if $k > i$:



Moreover, by construction $c_{\sigma(j+1)} = c_i$, so $P'_{j+1} \equiv Q_j$, because there is at most one possible reduction for a given annotation. Hence we can deduce a pair of reduction steps $P'_{j-1} \rightarrow_{ex}^{c_{\sigma(j+1)}} Q_{j-1} \rightarrow_{ex}^{c_{\sigma(j)}} P'_{j+1}$. This yields a new reduction sequence from P_0 to P'_n that corresponds to a new permutation σ' of the sequence (c_i) , and σ' is σ where $\sigma(j)$ and $\sigma(j+1)$ are swapped. By definition of j we have $\sigma(j) > \sigma(j+1)$ so $\sigma'(j) < \sigma'(j+1)$. For any $a \notin \{j, j+1\}$ we have $\sigma(a) < \sigma(j)$ if and only if $\sigma'(a) < \sigma'(j+1)$, and the same exchanging j and $j+1$, so we have $d(\sigma') = d(\sigma) - 1$, and we can conclude by induction hypothesis.

B.2 Typing

Proof (proposition 3). Clearly we can deduce the general result from the case of each individual rule. Cut elimination steps for multiplicatives and axioms do not affect the nesting of boxes, which is the only part of proofs used in term assignment, so for each such step $\pi \rightarrow_h^\emptyset \pi'$ we have $[\pi] = [\pi']$, hence $[\pi] \rightarrow_{ex*}^\emptyset [\pi']$ by reflexivity. For an elimination step for modalities, we have $\pi \rightarrow_h^{(\ell, m)} \pi'$ where ℓ and m are the locations of two boxes β and β' . By definition there is a cut between their principal ports x and y , so these ports must have dual types $\langle a \rangle A$ and $\langle \bar{a} \rangle A^\perp$. Call π_1 and π_2 the proofs associated to β and β' , then we have $[\beta] = a^\ell \cdot [\pi_1]$ and $[\beta'] = \bar{a}^m \cdot [\pi_2]$. Moreover, there is a term P such that $[\pi] = [\beta] \mid [\beta'] \mid P$ so we have $[\pi] \rightarrow_{ex}^{(\ell, m)} [\pi_1] \mid [\pi_2] \mid P$, and the latter is equal to $[\pi']$ by definition of the cut elimination step for modalities.

Proof (subject reduction theorem 1). Consider an execution step $[\pi] \rightarrow_{ex}^{(\ell, m)} P$. This step involves immediate subterms $a^\ell \cdot Q$ and $\bar{a}^m \cdot R$ for some name a , hence π must contain a box at top level with principal port $x^{(a)A}$ and one with principal port $y^{(\bar{a})A^\perp}$, for some formula A . Since π is linear, $\langle a \rangle A$ and $\langle \bar{a} \rangle A^\perp$ do not occur in the conclusion type, so they are cut. Since π is linear, no other boxes introducing these modalities can be in π . So x and y are elimination boxes for each other, and there is a path (as of definition 11) ρ from x to y in π . Remark that ends of ρ are modality rules on $\langle a \rangle A$ and $\langle \bar{a} \rangle A^\perp$ whereas all axioms along ρ contain these modalities in their types. Let π' be the rewriting of π where such modalities are removed by rewriting axioms on $\langle a \rangle A / \langle \bar{a} \rangle A^\perp$ in axioms on A / A^\perp , and by rewriting the end boxes by their contents. Clearly π' is a linear proof that infers P .

Proof (strong subject reduction theorem 2). Consider an execution step $[\pi] \xrightarrow{\varepsilon x}^{(\ell, m)}$ P . This step involves immediate subterms $a^\ell.Q$ and $\bar{a}^m.R$ for some name a , hence π must contain a box at top level with principal port $x^{\langle a \rangle A}$ and one with principal port $y^{\langle \bar{a} \rangle A^\perp}$, for some formula A . Since π is linear, $\langle a \rangle A$ and $\langle \bar{a} \rangle A^\perp$ do not occur in the conclusion type, so they are cut. Since π is linear, no other boxes introducing these modalities can be in π . So x and y are elimination boxes for each other, and there is a path (as of definition 11) ρ from x to y in π . For simplification we consider that boxes are replaced by their associated proof net keeping the information of auxiliary and principal ports (this is more like sequent calculus derivations).

Since x is at top level and cut, suppose that along ρ we get through a box β from x . By duality it is only with axioms on modality formulas $\langle a \rangle A / \langle \bar{a} \rangle A^\perp$. Moreover by typing rules going inside β can only be done through an auxiliary port. Since π is regular, ρ go out from β through its principal port. By typing rules it is not possible to reach y without encountering before a principal port of a box eliminating β . Moreover this is only possible using an axiom on modality formulas $\langle a \rangle A / \langle \bar{a} \rangle A^\perp$ in this box. Remark that since π is regular and β has axioms on modalities, corresponding elimination box is simple. Then to use an axiom on modalities contradict that π is regular. Then there is no box traversal along ρ . Then by typing ρ is a multiplicative cut path whose cut elimination \rightarrow_h^\emptyset until x and y preserves $[\pi]$ as well as regularity and linearity of proof.

B.3 Anti-execution by division of proof nets

Definition 15. A 3-partition of a set X is a triple (P, Q, R) of pairwise disjoint, possibly empty subsets of X such that $P \cup Q \cup R = X$.


Definition 16. A division of a proof net π is a 3-partition (P, Q, R) of the set of nodes of π such that the subgraph π_P of π induced by P (resp. Q, R) is a disjoint union of immediate subnets of π . The division is called *proper* if P, Q, R are not empty. It is called *simple* if π_P, π_Q, π_R are immediate subnets of π .

Remark that, because each component must be a proof net, each node in a component must have its ancestors in the same component, so the only edges of π that are not part $\pi_P \cup \pi_Q \cup \pi_R$ must be cut edges.

Let (P, Q, R) be a division of a proof net π , let π' be a reduct of π by some cut elimination step. Call P', Q', R' the restrictions of P, Q, R to the node set of π' . If (P', Q', R') is a division of π' , then we say that (P, Q, R) reduces to (P', Q', R') .

Lemma 5 (selection). Let (B_P, B_Q, B_R) be a 3-partition of the set of boxes of a proof net π . There is a division (P, Q, R) of a proof net π' that reduces to π by axiom steps such that Q' (resp. R') contains exactly the boxes B_Q (resp. B_R).

Proof. For each box $\beta \in B_Q \cup B_R$, for each auxiliary port $x^A \in \beta$, we introduce two new leaves y^{A^\perp}, z^A , a cut $\{x, y\}$ and an axiom $\{y, z\}$, and we turn any outgoing edge of x into an outgoing edge of z . Call π' the resulting proof net.




Let Q be the set of auxiliary ports of boxes in B_Q , R be the set of auxiliary ports of boxes in B_R , and let P be the set of all other nodes. Each box in $B_Q \cup B_R$ is obviously a subnet of π' , moreover P induces a subnet of π' , so (P, Q, R) is a division of π' with the expected properties.

Lemma 6 (chaining). *Let (P, Q, R) be a division of a proof net π . There exists a proper division (P', Q', R') of a proof net π' , with no cut between $\pi_{P'}$ and $\pi_{R'}$, that reduces to (P, Q, R) by axiom steps.*

Proof. Since π is a proof net, it must have at least one conclusion node x^A . If some set among P, Q, R is empty (say Q), extend π by adding an axiom on two fresh nodes $\{y^{A^\perp}, z^A\}$ and a cut $\{x, y\}$, and put y and z into Q . This way we get a proper division.


Now, for each cut $\{x^A, y^{A^\perp}\}$ with $x \in P$ and $y \in R$, introduce a fresh axiom $\{u^{A^\perp}, v^A\}$, put u and v in Q , then remove the cut $\{x, y\}$ and introduce two new cuts $\{x, u\}$ and $\{v, y\}$. This way we make sure there is no cut between P and R .



The only operation we have done here is introducing axiom/cut pairs on existing cut links or conclusions, so the resulting structure is a proof net that reduces to the original one by axiom steps, only removing the newly created nodes.

Lemma 7 (simplification). *Let (P, Q, R) be a proper division of a proof net π with no cut between π_P and π_R . There is a simple division (P', Q', R') of a proof net π' that reduces to (P, Q, R) by multiplicative steps.*

Proof. Since π is a proof net, it must be connected, so there are cuts between π_P and π_Q and between π_Q and π_R . We rewrite π to make sure that in each case there is exactly one cut. Suppose there are two cuts $\{x, x'\}$ and $\{y, y'\}$ between π_P and π_Q , with $x, y \in P$ and $x', y' \in Q$. If x and y are in the same connected component of π_P , then x' and y' must be in two distinct components of π_Q (because of acyclicity of π), so we can replace the two considered cuts by a cut on a new pair of nodes $x \bowtie y$ and $x' \otimes y'$, which we put in P' and Q' respectively:



Otherwise x and y are not in the same connected component of π_P , so x' and y' must be in the same component of π_Q (because of connectedness of π), so this case is symmetric. By construction, the resulting proof net π' reduces to π by multiplicative reduction and it has one less cut between $\pi'_{P'}$ and $\pi'_{Q'}$. Iterating this method, we can make sure that there is exactly one cut between $\pi'_{P'}$ and $\pi'_{Q'}$, and also between $\pi'_{Q'}$ and $\pi'_{R'}$. Subsequently, $\pi'_{P'}$, $\pi'_{Q'}$, $\pi'_{R'}$ are immediate subnets of π' , since they must be acyclic and connected.

Proof (proposition 4). Let $\{\ell, m\}$ be the domain of c . There is a name a such that $T_1 \equiv (P|a^\ell.Q|\bar{a}^m.R) \rightarrow_{ex}^c (P|Q|R) \equiv T_2$, and by hypothesis $[\pi_2] = T_2$. Let (B_P, B_Q, B_R) be the 3-partition of the set of boxes of π_2 such that B_P contains the boxes assigned to the actions in P , and similarly for Q and R .

By lemma 5 there is a division (p, q, r) a proof π that reduces to π_2 by axiom steps and such that q (resp. r) exactly contains the boxes B_Q (resp. B_R). By lemma 6 there is a proper division (p', q', r') of a proof π' that reduces to (p, q, r) by axiom steps, and with no cut between $\pi'_{p'}$ and $\pi'_{r'}$. By lemma 7 there is a simple division (p'', q'', r'') of a proof π'' that reduces to (p', q', r') by multiplicative steps. Then (p'', q'', r'') reduces by multiplicative and axiom steps to π_2 and all the boxes contained in $\pi''_{p''}$ are the ones B_P , and similarly for Q and R . Moreover there is no cut between $\pi''_{p''}$ and $\pi''_{r''}$ and one cut $\kappa = \{x^A, y^{A^\perp}\}$ between $\pi''_{q''}$ and $\pi''_{r''}$.

We define π_1 as the proof net π'' as follows. The subnet $\pi''_{q''}$ is replaced by a box β_q with location ℓ , with $\pi''_{q''}$ as contents, with x as its principal port and one auxiliary port for each conclusion of $\pi''_{q''}$, other than x . The subnet $\pi''_{r''}$ is replaced in the same way, with location m . The cut κ is retyped as $\{x^{(a)A}, y^{(\bar{a})A^\perp}\}$. This is depicted as follows.



By construction there is a cut elimination step $\pi_1 \rightarrow_h^c \pi''$, besides we have $\pi'' \rightarrow_h^\emptyset \pi_2$, so we get the expected reduction.