

A logical view on scheduling in concurrency

Emmanuel Beffara

I2M, Université d'Aix-Marseille & CNRS

This paper elaborates on a new approach of the question of the proof-theoretic study of concurrent interaction called “proofs as schedules”. Observing that proof theory is well suited to the description of confluent systems while concurrency has non-determinism as a fundamental feature, we develop a correspondence where proofs provide what is needed to make concurrent systems confluent, namely scheduling. In our core logical system (a variant of multiplicative linear logic), processes and schedulers appear explicitly as proofs in different fragments of the proof language and cut elimination between them does correspond to execution of a concurrent system. This separation of roles suggests new insights for the denotational semantics of processes and new methods for the translation of π -calculi into prefix-less formalisms (like solos), as the operational counterpart of a translation between logical systems.

1 Introduction

The extension of the familiar Curry-Howard correspondence outside the intuitionistic and functional worlds has been an active topic for decades, with a variety of approaches to the question of determinism and confluence, or lack thereof. The interactive nature of cut elimination procedures suggested relationships with actually interactive models of computation like games or process algebras. Several systems were proposed based on linear logic [13], following the intuition that it is a logic of interaction. Interpretations of proofs as processes, first sketched by Abramsky [1] and formalized by Bellin and Scott [6], later refined by various people including the author of this paper [3], stressed that proof nets [14] and process calculi have significant similarities in dynamics. At the same time, type systems for concurrency [26] revealed to be equivalent to variants of linear logic [16, 8]. 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 essentially functional behaviour.

Getting out of this restriction requires to actually handle the inherent non-determinism in concurrent systems. A natural option is of course to relate processes with proofs in a non-deterministic system. The prototypical such system is the original classical sequent calculus LK, however because of its structural rules it feels too remote from actual process dynamics to offer a direct correspondence (although encodings through intermediate systems have been proposed [2]). Several more direct approaches have been proposed based on linear logic: non-determinism in the style of complexity theory has been modelled using the additives of linear logic [18, 22, 23]; differential logic was recently developed by Ehrhard and Regnier [12] and features a structured form of non-determinism in its cut elimination; its untyped proof formalism was even shown expressive enough to represent the π -calculus [11].

The present work elaborates on a different approach to the topic, first sketched in a previous paper called *Proofs as executions* [5]. Our point of view is that cut elimination and general process execution should not be made to match directly, because their meaning is on different levels. On the one hand, the meaning of proofs lies in their normal forms, hence cut elimination should be confluent in order to preserve meaning. On the other hand, the meaning of a process is not its final irreducible form but what

happens to get there, as interaction with other processes (hence execution of interactions should definitely not preserve meaning). We thus establish a correspondence between proofs and interaction plans for processes, hereafter called *schedules* (rather than executions). These schedules are what provides necessary information to make a system deterministic, i.e. to make its execution confluent.

Concretely, we develop this idea in a simple framework, relating schedules of finite processes with proofs in multiplicative linear logic; this is to be considered as a first step towards a wider correspondence, illustrating the idea of our interpretation. The process calculus we use is introduced in Section 2, the logical system is introduced in Section 3. Section 4 presents two logical translations of processes for which the correspondence is proved, once in a synchronous form in Section 4.1 and once in an asynchronous form in Section 4.2, in which logic can abstract away from execution order. Section 5 discusses extensions and directions for future developments of the proofs-as-schedules paradigm, related to semantics of processes (Section 5.1), the status of name hiding and passing (Section 5.2) and the relationships with CPS translations and determinisation in classical logic (Section 5.3).

Note. The previous paper [5] formalizes this idea by defining a logical system, which appears as a kind of type system for processes, with the crucial property that for every lock-avoiding execution of a process P one can deduce a typing of P , hence a proof, whose cut elimination does correspond to the considered execution. A limitation of this result was that the “type” of a process could be very different depending on its particular environment and execution, which made it difficult to deduce an interpretation of processes in logical terms. The present work improves the previous results in several respects: Firstly, the typing of processes is uniform, and really independent of the particular executions it might exhibit. Subsequently, the contributions of the process and the scheduler in the proof corresponding to a scheduling are clearly identified, as simple fragments of the proof language. This in turn suggests new approaches to the logical study of denotational models of processes. Despite these references to a previous work, this paper aims to be self-contained.

Related work. The idea of matching proofs with executions is reminiscent of the proof search approach to computation. Indeed, the relationship between logical linearity and interaction has been explored (for instance by Miller and Tiu [19, 24]). Our approach has different roots than this, in particular because of the status of cut-elimination in our work, nevertheless proof search in our context looks like the building of a schedule for a fixed process, as illustrated in the technical proofs in this paper. In particular, Bruscoli establishes a correspondence between proof search in a variant of MLL and execution of a minimal CCS, in a context of deep-inference [7]. Although our translation of processes into formulas is different than that used by the above works, connections can certainly be established, but we defer this to later developments.

2 Multiplicative CCS

We consider processes in a fragment of the standard language CCS [21]. The fragment we use, hereafter called *multiplicative CCS* (or *MCCS*), is defined by the following grammar:

$$\begin{array}{ll}
 P, Q := 1 & \text{inactive process} \\
 P \mid Q & \text{parallel composition} \\
 a^\ell.P & \text{positive action prefix} \\
 \bar{a}^\ell.P & \text{negative action prefix}
 \end{array}$$

where a is taken from a given set of *channel names* and ℓ is taken from a given set $\mathcal{L}$ of *locations*. We impose that each location is used at most once in any term: locations are used to identify occurrences of

actions in a process. Note that we use 1 for the inactive process instead of the usual 0 because it is the neutral element of $|$ which is a multiplicative operation.

Definition 1. *Structural congruence is the smallest congruence $\equiv$ that makes parallel composition commutative and associative with 1 as neutral element.*

Definition 2 (execution). *Execution is the relation over structural congruence classes, labelled by partial involutions over $\mathcal{L}$, defined by the rule*

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

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

Locations will be used in Section 4.1 for the correspondence between execution and cut elimination and in Section 5.1 for the discussion of the semantic interpretation of our results. We introduce a few useful notations regarding locations:

Definition 3. *Let P be an M CCS term.*

- *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 $\{\pm 1\}$.*
- *The action order of P is the partial order $\leq_P$ over $\mathcal{L}(P)$ such that $\ell <_P m$ for every sub-term $x^\ell.Q$ of P with $m \in \mathcal{L}(Q)$.*

Remark that the information of locations, subjects, polarities and action order completely characterize a class of structural congruence of M CCS terms, so we can consider that providing such information does define a process without ambiguity.

When locations are unimportant for a given statement, they will be kept implicit. With this convention, the definition of the reduction relation is the standard one:

$$a.P | \bar{a}.Q | R \rightarrow P | Q | R$$

This language is very minimalistic and its operational theory is actually very simple (one can prove that its bisimilarity is completely axiomatized by a single rule scheme $(a.P)^{n+1} \simeq a.(P | (a.P)^n)$ plus structural congruence). The results of this paper extend smoothly to a framework with replication and sum. A much more subtle point is that of name hiding, which is discussed in Section 5.2.

3 MLL with actions

Formulas of propositional multiplicative linear logic with actions (in short MLLa) are generated from the following grammar:

$$\begin{array}{ll} A, B := \alpha, \alpha^\perp & \text{propositional variable} \\ A \otimes B, A \wp B & \text{conjunction, disjunction} \\ \langle a \rangle A, \langle \bar{a} \rangle A & \text{action modalities} \end{array}$$

where α is taken from a given set of propositional variables and a is taken from the set of channel names. Negation is defined inductively on formulas in the standard way, with $(\langle a \rangle A)^\perp = \langle \bar{a} \rangle (A^\perp)$; linear implication $A \multimap B$ is defined as a shorthand for $A^\perp \wp B$.




$\frac{}{\vdash A^\perp, A} \text{ ax}$		$\frac{\vdash \Gamma, A \quad \vdash \Delta, A^\perp}{\vdash \Gamma, \Delta} \text{ cut}$	
$\frac{\vdash \Gamma, A \quad \vdash \Delta, B}{\vdash \Gamma, \Delta, A \otimes B} \otimes$		$\frac{\vdash \Gamma, A, B}{\vdash \Gamma, \Delta, A \wp B} \wp$	
$\frac{\vdash \Gamma, A}{\vdash \Gamma, \langle a \rangle A} \langle a \rangle$		$\frac{\vdash \Gamma, A}{\vdash \Gamma, \langle \bar{a} \rangle A} \langle \bar{a} \rangle$	

Table 1: Proof rules and proof net syntax for MLLa.

The modalities are reminiscent of those of Hennessy-Milner [15] logic, however the logic itself and is very different, in particular because of linearity. While Hennessy-Milner modalities $\langle a \rangle A$ and $[a]A$ mean, respectively, “I can do a and then satisfy A ” and “whenever I do a , I will satisfy A ”, in our logic the formula $\langle a \rangle A$ means something like “I will do a and then exhibit behaviour A ”. There is no analogue of the modality $[a]$ because all the information we provide is positive. The duality induced by linear negation $(\cdot)^\perp$ is not a classical negation but a change of roles: $A^\perp$ is the type of behaviours that interact correctly with behaviours of type A . The connectives $\otimes$ and $\wp$ are not classical conjunction and disjunction either, but spatial ones representing causality and independence between parts of a run, using connectedness/acyclicity arguments to describe avoidance of deadlocks.


Proof rules are the standard rules of linear logic, extended with modality, as shown in Table 1. We use the language of proof nets, with the notations from the same table. As far as proof theory is concerned, these modalities are innocuous: they commute with all other rules, their cancellation rule in cut elimination is the obvious one, the types A and $\langle a \rangle A$ are isomorphic. Indeed, they are nothing more than markers in formulas. For this reasons, standard theory for multiplicative proof nets [10] applies to proof nets of MLLa. In order to ease the formulation of operational correspondence proofs, we assume that each proof structure comes with an injective labelling of its modality links with locations in $\mathcal{L}$.

We restrict ourselves to propositional logic, however our construction heavily relies on the instantiation of propositional variables with particular formulas: this substitution mechanism is a key part of interaction as we model it. The same work could be carried out, possibly in a cleaner way, using second-order quantification, however we decided not to use this quantification explicitly because it makes proof theory more complicated and we do not need its full power in the present work. A similar approach was taken by Terui in his translation of boolean circuits into MLL proofs [23].

Definition 4. Let π be an MLLa proof structure (possibly with cuts).

- The set of locations occurring in π is written $\mathcal{L}(\pi)$.
- Given $\ell \in \mathcal{L}(\pi)$, the subject of ℓ is the name in the modality tagged by ℓ , written $\text{subj}_\pi \ell$. The polarity of ℓ is that of the modality, written $\text{pol}_\pi \ell$, element of $\{\pm 1\}$.
- The proof order of π is the partial order $\leq_\pi$ over $\mathcal{L}(\pi)$ such that $\ell <_\pi m$ when the link tagged m appears in the tree of premisses of the link tagged by ℓ .

These definitions are similar to those for MCCS terms. The partial order $\leq_\pi$ is a bit more complicated to define but can be explained in a simple way. Observe that a proof structure π can always be decomposed as a family of trees of links together with a set of axioms between leafs and possibly a set of cuts between roots:



Then the partial order $\leq_\pi$ over $\mathcal{L}(\pi)$ such that $\ell \leq_\pi m$ when the modality link labelled ℓ occurs below that labelled m in one of the trees π_i .

4 Execution as implication

In this section, we formally develop the correspondence between operational semantics of processes and proofs in MLLa:

- Each process term P is translated into a cut-free proof in MLLa, where instances of modality rules correspond to action prefixes in P ; the conclusion $\llbracket P \rrbracket$ of this proof (which can be considered as the type of P) is deduced syntactically from P .
- A schedule for reducing P to Q is a *multiplicative* proof of the implication $\llbracket P \rrbracket \multimap \llbracket Q \rrbracket$, i.e. this proof may not contain modality rules, which stresses the fact that it is only allowed to relate actions in a given process without introducing new actions.
- An execution of P corresponds to a cut elimination sequence of P cut against a schedule for P ; various cut elimination sequences correspond to different linear orderings of independent events in a given execution.

Hence the concurrent aspect of execution is represented by the variety of proofs for a given implication, while different choices in cut elimination correspond to unimportant ordering decisions, which is consistent with the fact that cut elimination is confluent.

We implement this correspondence in two different ways. The first version is *synchronous* in that it exactly relates provability of implication and cut elimination with step-by-step execution of processes. The second version is *asynchronous* in that it allows scheduling decisions to be made in advance of execution, which leads to a more flexible system, albeit with a more intricate interpretation.

4.1 Synchronous translation

In this translation, each term P is mapped to a formula $\llbracket P \rrbracket_S$ that has a unique cut-free proof, which follows the syntactic structure of P . We then prove, in Theorem 1, that $\llbracket P \rrbracket_S \multimap \llbracket Q \rrbracket_S$ is provable if and only if there is an execution $P \rightarrow^* Q$.

Definition 5 (type assignment). *Terms of MCCS are translated into MLLa formulas as follows, where α is a fresh propositional variable:*

$$\begin{aligned}
 \llbracket 1 \rrbracket_S &:= \alpha^\perp \wp \alpha \\
 \llbracket P \mid Q \rrbracket_S &:= \llbracket P \rrbracket_S \otimes \llbracket Q \rrbracket_S \\
 \llbracket a.P \rrbracket_S &:= \langle a \rangle (\alpha^\perp \wp (\llbracket P \rrbracket_S \otimes \alpha)) &= \langle a \rangle (\alpha \multimap (\llbracket P \rrbracket_S \otimes \alpha)) \\
 \llbracket \bar{a}.P \rrbracket_S &:= \langle \bar{a} \rangle (\llbracket P \rrbracket_S \otimes \alpha^\perp) \wp \alpha &= \langle a \rangle (\llbracket P \rrbracket_S \multimap \alpha) \multimap \alpha
 \end{aligned}$$

The freshness condition is a way to enforce polymorphism in our translation: each proposition variable occurs exactly twice, once in each polarity. As observed in Section 3, the intended meaning is

indeed a universal quantification over these variables, in a context where we chose not to use second order quantification for simplicity. For the same reason, the type for 1 is the type of an identity, where we could have used the multiplicative unit $\mathbf{1}$ of linear logic (indeed, $\mathbf{1}$ and $\forall \alpha (\alpha^\perp \multimap \alpha)$ are equivalent formulas): we deliberately restrict ourselves to a unit-free logic in order to avoid the slight proof-theoretic complications of the units.

Proposition 1 (proof assignment). *For every M CCS term P , there is a unique cut free proof of $\llbracket P \rrbracket_S$. This proof will be denoted as $\llbracket P \rrbracket_S$.*

Proof. By straightforward induction on terms. The case of action prefixes is as follows:



The axioms are necessarily in this position because the variables are fresh: no other occurrences appear anywhere else so they must be introduced by the same axiom rule. $\square$

This proof assignment property is formulated in the absence of locations. We enrich it in the obvious way with location information: the modality link $\langle a \rangle$ that corresponds to a prefix $a^\ell.P$ gets labelled by the location ℓ .

Definition 6 (term extraction). *Let P be a term of M CCS. An MLLa proof structure π is said to be compatible with P if $\mathcal{L}(\pi) \subseteq \mathcal{L}(P)$, subjects and polarities coincide between P and π and the order $\leq_P$ is included in the order $\leq_\pi$.*

If π is compatible with P then define the term $P \upharpoonright \pi$ to be the M CCS term with locations $\mathcal{L}(\pi)$, subjects and polarities as in P and π and action order as in P .

Remark that for all term P , by construction $\llbracket P \rrbracket_S$ is compatible with P and we have $P \upharpoonright \llbracket P \rrbracket_S \equiv P$. Note also that compatibility implies that the process is *less* constrained than the proof: given a proof π , the set of processes compatible with π contains terms that are more parallel than the structure of π (and in particular the processes that is just a parallel composition of all actions that correspond to modality links in π). The stricter ordering in the proof π effectively means that π has made decisions on how P will run, as illustrated by the lemma below and Theorem 1.

Lemma 1. *Assume P is a term and π be a proof structure compatible with P . For any cut elimination step $\pi \rightarrow \pi'$, the structure π' is compatible with P and one of the two following situations occurs:*

- *either $\pi \rightarrow \pi'$ is an elimination of modality links $\langle a \rangle$ and $\langle \bar{a} \rangle$ at some locations ℓ and m , then $P \upharpoonright \pi \rightarrow_{\{(\ell, m)\}} P \upharpoonright \pi'$ is a valid execution step,*
- *or it is another elimination step and $P \upharpoonright \pi \equiv P \upharpoonright \pi'$.*

Proof. The case of an elimination of modality links follows from the definition of term extraction: the actions in the extracted term are necessarily at top level since the modalities are premisses of a cut so their locations are minimal with respect to $\leq_\pi$ hence with respect to $\leq_P$ too. In π' the location set is the same with ℓ and m removed and the other data (subjects, polarities and proof order) are simply restricted to this new set, so compatibility with P is preserved.

The other cut elimination steps can be either multiplicative eliminations or axiom eliminations. In these cases the location set is unchanged. For multiplicative steps the proof order is unchanged too, while in the case of axiom eliminations the proof order is possibly strengthened, since one tree of links in the proof is merged on top of another. In both cases, this preserves compatibility with P . $\square$


Remark that the strengthening of the proof order in the case of axiom elimination is the very reason why we need term extraction: in this case the syntactic structure of the proof is more constrained than that of the process term P we are observing, so we cannot read back a reduct of P without extra information.

Theorem 1. *Let P and Q be two MCCS terms. There is an execution $P \rightarrow^* Q$ if and only if $\lceil P \rceil_S \multimap \lceil Q \rceil_S$ is provable in MLL (without modality rules) for some instantiation of the propositional variables of $\lceil P \rceil_S$.*

Proof. For the direct implication, remark that for a structural congruence $P \equiv Q$ at top level (i.e. not under prefixes), an implication $\lceil P \rceil_S \multimap \lceil Q \rceil_S$ is provided by a standard isomorphism for the associativity and commutativity of the tensor, hence it is enough to handle the case of a reduction $(a.P \mid \bar{a}.Q) \mid R \rightarrow (P \mid Q) \mid R$. The expected type is

$$\begin{aligned} \lceil (a.P \mid \bar{a}.Q) \mid R \rceil_S \multimap \lceil (P \mid Q) \mid R \rceil_S = \\ ((\langle \bar{a} \rangle (\alpha \otimes (\lceil P \rceil_S^\perp \wp \alpha^\perp))) \wp (\langle a \rangle (\lceil Q \rceil_S^\perp \wp \beta) \otimes \beta^\perp)) \wp \lceil R \rceil_S^\perp \wp ((\lceil P \rceil_S \otimes \lceil Q \rceil_S) \otimes \lceil R \rceil_S) \end{aligned}$$

so we get a proof as follows, if we instantiate α as $\lceil Q \rceil_S$ and β as $\lceil P \rceil_S \otimes \lceil Q \rceil_S$:



For the reverse implication, consider an MLL proof of some $\lceil P \rceil_S \multimap \lceil Q \rceil_S$. If we cut this proof against $\lceil P \rceil_S$, we get a proof of $\lceil Q \rceil_S$. Cut elimination reaches the unique cut free proof of $\lceil Q \rceil_S$, namely $\lceil Q \rceil_S$. By Lemma 1, we know that each cut elimination step induces either a structural congruence or an execution step in terms extracted from P by intermediate proofs, so from a cut elimination sequence $\lceil P \rceil_S \rightarrow^* \lceil Q \rceil_S$ we can extract a CCS execution sequence $P \rightarrow^* Q$, and incidentally we also deduce that $P \upharpoonright \lceil Q \rceil_S = Q$. $\square$

4.2 Asynchronous executions

The tight correspondence of Theorem 1 is obtained thanks to the fact that the action order $\leq_P$ of processes is mapped to proof order $\leq_{\lceil P \rceil_S}$ because of syntactic structure of the type $\lceil P \rceil_S$, in which modalities are nested as in P . The downside of this correspondence is that it does not leave much space for modular reasoning about the behaviour on different channels in processes. Besides, it does not account for proofs extracted from executions as in our previous work [5].

In this section, in order to generalise this result we, provide an asynchronous variant of the correspondence. We discuss below the relevance of this variant.

Definition 7 (asynchronous type assignment). *Terms of M CCS are translated into MLLa formulas as follows, where α and β are fresh propositional variables:*

$$\begin{aligned}
 [1]_A &:= \alpha^\perp \wp \alpha \\
 [P \mid Q]_A &:= [P]_A \otimes [Q]_A \\
 [a.P]_A &:= \langle a \rangle \alpha^\perp \wp ([P]_A \otimes \alpha) &= \langle a \rangle \alpha \multimap ([P]_A \otimes \alpha) \\
 [\bar{a}.P]_A &:= ([P]_A \otimes \beta^\perp) \wp \langle \bar{a} \rangle \beta &= ([P]_A \multimap \beta) \multimap \langle \bar{a} \rangle \beta
 \end{aligned}$$

Note that this translation is the same as the synchronous translation of Definition 5, except for the position of modalities.

Proposition 2 (asynchronous proof assignment). *For every M CCS term P , there is a unique cut free proof of $[P]_A$. This proof will be denoted as $\langle P \rangle_A$.*

Proof. By straightforward induction on terms. The case of action prefixes is as follows:




□

Proposition 3. *For all M CCS reduction $P \rightarrow Q$, the formula $[P]_A \multimap [Q]_A$ is provable in MLL for some instantiation of the variables of $[P]_A$.*

Proof. For any structural congruence $P \equiv Q$, the implication $[P]_A \rightarrow [Q]_A$ is proved by a standard isomorphism for the monoidal structure of the tensor. For an execution step $(a.P \mid \bar{a}.Q) \mid R \rightarrow (P \mid Q) \mid R$, we prove the implication:

$$\begin{aligned}
 [(a.P \mid \bar{a}.Q) \mid R]_A \multimap [(P \mid Q) \mid R]_A = \\
 ((\langle \bar{a} \rangle \alpha \otimes ([P]_A^\perp \wp \alpha^\perp) \wp (([Q]_A^\perp \wp \beta) \otimes \langle a \rangle \beta^\perp)) \wp [R]_A^\perp) \wp (([P]_A \otimes [Q]_A) \otimes [R]_A)
 \end{aligned}$$

with the following proof:




□

Lemma 2. *Let P be an MCCS term with at least one action prefix. If $\lceil P \rceil_A \multimap \lceil 1 \rceil_A$ is provable in MLL for some instantiation of the propositional variables, then there exists a reduction $P \rightarrow Q$ such that $\lceil Q \rceil_A \multimap \lceil 1 \rceil_A$ is also provable.*


Proof. Consider a proof π of $\lceil P \rceil_A^\perp \wp \lceil 1 \rceil_A$ for some instantiation of the propositional variables of $\lceil P \rceil_A$. MLL admits cut elimination and expansion of axiom links, so we may assume without loss of generality that π is cut-free and that all multiplicative connectives in $\lceil P \rceil_A^\perp$ are introduced by π .

The term P can be written as a composition of action prefixes $a_1.P_1 \mid \dots \mid a_n.P_n$ (where the a_i are actions of arbitrary polarities) so π is decomposed as follows, with some tensors possibly reversed depending on the polarities of the a_i :




with axiom links on top (the $\wp$ link of arity $n+2$ in the picture stands for a tree of binary $\wp$ links, since our language only includes binary connectives). Note that, since π does not use modality rules, all premisses $\langle \bar{a}_i \rangle \alpha_i$ must be introduced by axiom links.

We now argue that two of these are actually connected by an axiom rule. For contradiction, assume that it is not the case: every $\langle \bar{a}_i \rangle \alpha_i$ is introduced by an axiom link whose other conclusion is a leaf in one of the subproofs $\pi_{f(i)}$. Since there are finitely many subproofs, the sequence $1, f(1), f^2(1), \dots$ is eventually periodic. Consider a minimal cycle $i, f(i), f^2(i), \dots, f^k(i) = i$. Up to the reordering of sub-terms of P , let us assume this sequence is $1, 2, \dots, k, 1$. Then we have the following situation:



Call t_i the tensor link under the node $\langle \bar{a}_i \rangle \alpha_i$, and x_i the axiom link that has $\langle \bar{a}_i \rangle \alpha_i$ as one of its premisses. Then we have a cycle that goes from t_1 up to x_1 into π_2 , down to t_2 , up to x_2 into π_3 and so on until t_k , up to x_k and back into π_1 , down to t_1 . This cycle traverses each tree π_i straight from a leaf to the root, so it only goes through at most one premiss of each $\wp$ link. Hence there is a $\wp$ -switching of π that has a cycle, which violates the Danos-Regnier correctness criterion for multiplicative proof nets [10].

Hence there is an axiom link between two of the $\langle \bar{a}_i \rangle \alpha_i$. For simplicity, assume that the indices are 1 and 2. Then a_1 and a_2 are dual, so P can be written $\bar{a}.P_1 \mid a.P_2 \mid P'$, and we have $P \rightarrow P_1 \mid P_2 \mid P'$. Moreover the cut between $\lceil P \rceil_A$ and π has the following shape:



After a few steps of cut elimination (five pairs of multiplicatives, three axioms and a pair of modalities) this proof eventually reduces into π' cut against $(P_1)_A$, $(P_2)_A$ and $(P')_A$ plus a cut between the two remaining ports of π' . This reduct is also a reduct of $(P_1 | P_2 | P')_A$ against π' plus two $\mathcal{J}$ rules. The latter is thus an MLL proof of $[P_1 | P_2 | P']_A \multimap [1]_A$. $\square$

Theorem 2. *Let P be an MCCS term. There is an execution $P \rightarrow^* 1$ if and only if $[P]_A \multimap [1]_A$ is provable in MLL (without modality rules) for some instantiation of the propositional variables.*

Proof. For the direct implication, Proposition 3 provides the implication for one step of execution. The result follows immediately since linear implications properly compose and the correctness of proofs is preserved by instantiation of propositional variables.

For the reverse implication, we can reason by induction on the number of action prefixes in P . The base case of a term with no action is vacuous since the neutral process 1 is the only such term, up to structural congruence. Lemma 2 provides the induction step. $\square$

The interpretation of Theorem 2 is more subtle than for Theorem 1, because in this variant there is no step-by-step correspondence between execution and cut elimination. Of course, Lemma 2 does prove that there is a *strategy* for eliminating cuts that corresponds to an execution, but in general the normalisation of $(P)_A$ against a particular proof of $[P]_A \multimap [1]_A$ may consume modalities in P in an arbitrary order. This is why Theorem 2 requires the final type to be $[1]_A$: the statement like of Theorem 1 is actually false for the asynchronous variant.

Definition 8. *Let $\Rightarrow$ be the relation over processes such that $P \Rightarrow Q$ holds whenever $[P]_A \multimap [Q]_A$ is provable in MLL for some instantiation of the propositional variables of $[P]_A$.*

By Proposition 3 we know that $P \rightarrow^* Q$ implies $P \Rightarrow Q$, but other rules are easily provable, like the following:

$$\frac{P \rightarrow Q}{P \Rightarrow Q} \qquad \frac{P \Rightarrow Q}{a.P \Rightarrow a.Q} \qquad u.a.P | v.\bar{a}.Q \Rightarrow u.P | v.Q$$

Therefore $\Rightarrow$ can make decisions in advance about the execution of a process, and update the visible part of the process accordingly. In other words, $\Rightarrow$ is an fully asynchronous form of execution: it allows to perform an interaction as soon as this interaction may be part of some full execution.

5 Discussion

We thus have defined two type assignment systems for MCCS processes into MLLa that enjoy good properties relating cut-elimination and execution. It appears that processes correspond to cut-free proofs using modality links, while schedules correspond to pure MLL proofs without modality links.

Although the technical details of the paper are limited to multiplicative CCS, it is clear that the same approach extends to the full calculus, including replication and sum, with a type assignment like

$$\lceil !P \rceil_A = ! \lceil P \rceil_A \qquad \lceil P + Q \rceil_A = \lceil P \rceil_A \& \lceil Q \rceil_A$$

so that, for replication, a schedule has to decide the number of copies of P it will use, and for the sum it will have to decide which side of the additive conjunction will be effectively used. This leads to a correspondence using full propositional linear logic; the theory of proof nets is harder in this case but the principles of the correspondence remain the same.

We now conclude with ideas for research directions open by the present work.

5.1 Semantics of processes

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.

The annotation c in an execution step $P \rightarrow_c Q$ describes which occurrences interact. Remark that, for a given P and c , there is at most one Q such that $P \rightarrow_c Q$, since c describes the interaction completely.

Definition 9 (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 . 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 $P = a^1.c^2 | b^3.\bar{a}^4 | \bar{b}^5.c^6 | a^7.\bar{b}^8 | b^9 | \bar{a}^0$ 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. In our previous work [5], we established the following precise relationship between pairings and executions:

Theorem 3. *Let P be an MCCS term and c a pairing of P .*

- *c is consistent if and only if there is a term Q such that $P \rightarrow_c^* Q$,*
- *any two executions $P \rightarrow_c^* Q$ and $P \rightarrow_c^* R$ with the same pairing are permutations of each other, and in this case $Q \equiv R$.*

Maximal consistent pairings represent executions of processes until a state where no more execution is possible.

We can actually relate pairings and logical schedules in a precise way. Consider a term P and proof π of $\lceil P \rceil_A \multimap \lceil 1 \rceil_A$. In a cut elimination process of $\langle P \rangle_A$ against π , eventually each modality link in $\langle P \rangle_A$ will be eliminated by a dual link also from $\langle P \rangle_A$, and this association between modality links is of course independent of the cut elimination sequence since proof normalisation is confluent. Hence π induces a pairing between modality links in $\langle P \rangle_A$, and since modality links are in bijection with locations in P , this in turn induces a pairing of P . This pairing is maximal since it reaches all locations in P , and by Theorem 2 it is consistent. Reciprocally, every maximal consistent pairing of P induces a class of proofs of $\lceil P \rceil_A \multimap \lceil 1 \rceil_A$.

Moreover, this relationship between pairings and proofs is actually a correspondence between pairings of P and possible ways of positioning axiom links on top of a canonical multiplicative proof structure

for $\lceil P \rceil_A \multimap \lceil 1 \rceil_A$, as observed in the proof of Lemma 2. However, this correspondence is not straightforward (like one axiom for each pair) because there are more leaves in proofs of this type than there are actions in P , moreover the instantiation of propositional variables is far from obvious since, as in the proof of Proposition 3, it may include the whole type of a term involved in the execution. We defer the precise study of this relationship to future work.

The next step in semantics is the study of denotational models of proofs in our system, which should provide new insights for the denotational semantics of processes. For instance, it is a trivial remark that coherence spaces [13] have a formal similarity with event structures [25], although a precise relationship is hard to formulate, especially in a proofs-as-processes approach. Since our constructions maps processes to types, i.e. spaces, and schedules to proofs i.e. cliques or configurations in spaces, it should be possible to formalize such a relationship in our setting in a meaningful way. A difficulty in this respect is of course to define a non-trivial interpretation of modalities: this is needed in order to identify cliques that are interpretations of multiplicative proofs, as opposed to general proofs with modalities.

5.2 Name hiding and passing

The question of how to integrate name hiding in our system is crucial as it is necessary for expressiveness and modularity, however Theorems 1 and 2, which characterize execution in the proof system do not extend to a calculus with name hiding. Indeed, a reasonable type assignment for a term $(\nu a)P$ should not let a appear in the formula $\lceil (\nu a)P \rceil$, yet actions on this bound a should be scheduled.

A syntactic way of handling name hiding could be to allow quantification over channel names in the logic, so that (νa) in a process would become a quantification in the type. A schedule would then have to handle actions on this name without knowing anything about it, hence considering it effectively as a fresh name. However, the nature of such a quantifier is unclear:

- A universal $\forall a$ is not an option, since it would allow the scheduler to provide any name for a in $(\nu a)P$, including one that P or its environment already knows, which would produce new possible interactions.
- An existential $\exists a$ is better suited since it forces the scheduler to behave independently of the actual name used, effectively treating a as fresh. However, a proper correspondence can only be obtained only if, by construction, $\exists a$ is only introduced with a premiss where a is fresh, and such a restriction is very unnatural as it does not respect the meaning of connectives in the target logic.
- Using a specific quantifier for freshness, like Miller and Tiu's ∇ [20], would probably solve this issue but at the price of introducing an ad-hoc construct with its particular theory in an otherwise well-studied logic, with the consequence that it would make the semantic explorations mentioned above more difficult.

A very different approach, more semantic than the use of a quantifier, would be to decide that placing a (νa) in front of a process would mean something like making decisions about the scheduling that will happen on name a , by cutting the relevant conclusions against a partial schedule (i.e. a proof without modality rules).

- Only the asynchronous translation of section 4.2 can handle this approach, since Theorem 1 proves that any partial schedule in the synchronous variant is the beginning of an execution. This is consistent with the fact that it amounts to scheduling a channel in advance of actual execution.


- It blurs the distinction between a process and a scheduler, since now a process does include scheduling decisions. This issue could be solved by typing a process $(\nu a)P$ with something like $S \multimap [P']_A$ where S is the type of a scheduler for a channel (depending on how a is used in P) and P' is a type for P where a is hidden.

The latter approach would most likely require second-order quantification in order to specify S . Further study of this point should also be the way to approach the problem of extending our work to name passing calculi: once we know how to properly specify what a channel is for restricting it, we can investigate how to communicate it through another channel.

5.3 CPS translations and determinisation

Both our type assignments actually translate actions and processes into types of *intuitionistic* MLL formulas, therefore proofs extracted from processes could be interpreted as linear functions. Moreover, the structure of the translations, up to the position of modalities, is $[\bar{a}.P] = \forall \alpha. ([P] \multimap \alpha) \multimap \alpha$, which is a weak form of double-negation, so weak in fact that this formula, without modalities, is isomorphic to $[P]$. This suggests a kind of (linear) continuation-passing-style translation of processes into a linear λ -calculus, which could be another setting for the semantic and operational study of processes and schedules. In relation with other forms of proofs-as-processes correspondences, this might suggest an approach to the question of desequentialisation of processes (like translations into solos [4, 17]) as a translation between logical systems, in a way similar to CPS translations from classical logic into intuitionistic logic.

This idea of CPS-like translations indeed suggest a step in a wider programme for the approach of non-determinism in logic. It is well known that the classical sequent calculus LK is non confluent, to the point that any two proofs of the same formula can be identified by conversion through cut-elimination, which makes any direct denotational semantics of proofs degenerate, and that interesting semantics can be obtained by means of double-negation translations into intuitionistic logic. Such translations are effective because they impose constraints on proof reduction, which on the computational side corresponds to fixing reduction strategies like call-by-name or call-by-value in languages with control. It is a form of determinisation, and such translations have been extensively studied. The systematic study of decompositions of LK into linear logic by Danos, Joinet and Schellinx in the system LKtq [9] is of particular interest with respect to the present work since it combines determinisation and linearisation. Reading these translations from the computational point of view provides translations of classical proofs into a deterministic process calculus.



Our approach of non-determinism by means of scheduling suggest a different take on these questions: would it be possible to provide a for; of linearisation of classical proofs that would preserve the intrinsic

non-determinism of LK, so that various evaluations strategies could be obtained by effectively choosing different types of schedulers? The semantic study of our proofs-as-schedules correspondence could be an approach to this kind of question by means of a study of the semantics of processes.

References

- [1] Samson Abramsky (1994): *Proofs as processes*. *Theoretical Computer Science* 135(1), pp. 5–9, doi:10.1016/0304-3975(94)00103-0. Available at <http://www.sciencedirect.com/science/article/pii/0304397594001030>.
- [2] Steffen van Bakel, Luca Cardelli & Maria Grazia Vigliotti (2014): *Classical cut-elimination in the π -calculus*.
- [3] Emmanuel Beffara (2006): *A concurrent model for linear logic*. In: *21st International Conference on Mathematical Foundations of Programming Semantics (MFPS)*, 155, pp. 147–168.
- [4] Emmanuel Beffara & François Maurel (2006): *Concurrent nets: a study of prefixing in process calculi*. *Theoretical Computer Science* 356(3), pp. 356–373.
- [5] Emmanuel Beffara & Virgile Mogbil (2012): *Proofs as executions*. In Jos C. M. Baeten, Tom Ball & Frank S. de Boer, editors: *Proceedings of IFIP TCS, Lecture Notes in Computer Science* 7604, Springer, pp. 280–294. Available at <http://hal.archives-ouvertes.fr/hal-00586459/>.
- [6] Gianluigi Bellin & Philip J. Scott (1994): *On the π -calculus and linear logic*. *Theoretical Computer Science* 135(1), pp. 11–65.
- [7] Paola Bruscoli (2002): *A purely logical account of sequentiality in proof search*. In Peter J. Stuckey, editor: *International Conference on Logic Programming (ICLP), Lecture Notes in Computer Science* 2401, Springer, pp. 302–316.
- [8] Luís Caires & Frank Pfenning (2010): *Session types as intuitionistic linear propositions*. In: *Proceedings of the 21st International Conference on Concurrency Theory*, Springer, Paris, France, p. 222–236. Available at http://link.springer.com/chapter/10.1007/978-3-642-15375-4_16.
- [9] Vincent Danos, Jean-Baptiste Joinet & Harold Schellinx (1997): *A new deconstructive logic: linear logic*. *The Journal of Symbolic Logic* 62(3), pp. 755–807, doi:10.2307/2275572. Available at <http://www.jstor.org/stable/2275572>. ArticleType: research-article / Full publication date: Sep., 1997 / Copyright © 1997 Association for Symbolic Logic.
- [10] Vincent Danos & Laurent Regnier (1989): *The structure of multiplicatives*. *Archive for Mathematical Logic* 28, pp. 181–203.
- [11] Thomas Ehrhard & Olivier Laurent (2010): *Interpreting a finitary pi-calculus in differential interaction nets*. *Information and Computation* 208(6), pp. 606–633, doi:10.1016/j.ic.2009.06.005. Available at <http://www.sciencedirect.com/science/article/pii/S0890540110000155>.
- [12] Thomas Ehrhard & Laurent Regnier (2006): *Differential interaction nets*. *Theoretical Computer Science* 364(2), p. 166–195, doi:DOI: 10.1016/j.tcs.2006.08.003.
- [13] Jean-Yves Girard (1987): *Linear Logic*. *Theoretical Computer Science* 50, pp. 1–102.
- [14] Jean-Yves Girard (1996): *Proof-nets: the parallel syntax for proof-theory*. In Paolo Agliano & Aldo Ursini, editors: *Logic and Algebra*, M. Dekker, New York.
- [15] Matthew Hennessy & Robin Milner (1985): *Algebraic laws for nondeterminism and concurrency*. *Journal of the ACM* 32(1), pp. 137–161.
- [16] Kohei Honda & Olivier Laurent (2010): *An exact correspondence between a typed π -calculus and polarised proof-nets*. *Theoretical computer science* 411(22–24), p. 2223–2238.
- [17] Cosimo Laneve & Björn Victor (2003): *Solos in Concert*. *Mathematical Structures in Computer Science* 13(5), pp. 657–683.

- [18] François Maurel (2003): *Nondeterministic light logics and NP time*. In Martin Hofmann, editor: *Typed Lambda Calculi and Applications (TLCA)*, 6th International Conference, LNCS, Springer, pp. 241–255.
- [19] Dale Miller (1992): *The π -calculus as a theory in linear logic: preliminary results*. In Evelina Lamma & Paola Mello, editors: *Extensions of Logic Programming, Third International Workshop (WELP'92)*, *Lecture Notes in Computer Science* 660, Springer, pp. 242–264.
- [20] Dale Miller & Alwen Tiu (2005): *A proof theory for generic judgments*. *ACM Transactions on Computational Logic (TOCL)* 6(4), p. 749–783. Available at <http://dl.acm.org/citation.cfm?id=1094628>.
- [21] Robin Milner (1989): *Communication and concurrency*. Prentice Hall.
- [22] Virgile Mogbil (2010): *Non-deterministic Boolean Proof Nets*. In Marko van Eekelen & Olha Shkaravska, editors: *Foundational and Practical Aspects of Resource Analysis*, *Lecture Notes in Computer Science* 6324, Springer Berlin Heidelberg, pp. 131–145. Available at http://link.springer.com/chapter/10.1007/978-3-642-15331-0_9.
- [23] Kazushige Terui (2004): *Proof nets and boolean circuits*. In: *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science*, LICS '04, IEEE Computer Society, Washington, DC, USA, p. 182–191, doi:10.1109/LICS.2004.34. Available at <http://dx.doi.org/10.1109/LICS.2004.34>.
- [24] Alwen Tiu & Dale Miller (2005): *A proof search specification of the π -calculus*. *ENTCS* 138(1), pp. 79–101.
- [25] Glynn Winskel (1987): *Event structures*. In W. Brauer, W. Reisig & G. Rozenberg, editors: *Petri Nets: Applications and Relationships to Other Models of Concurrency*, *Lecture Notes in Computer Science* 255, Springer Berlin Heidelberg, pp. 325–392. Available at http://link.springer.com/chapter/10.1007/3-540-17906-2_31.
- [26] Nobuko Yoshida, Martin Berger & Kohei Honda (2001): *Strong Normalisation in the π -Calculus*. In: *16th IEEE Symposium on Logic in Computer Science (LICS)*, pp. 311–322.