Focusing strategies in the sequent calculus of synthetic connectives

Kaustuv Chaudhuri

INRIA Saclay - Île-de-France
Kaustuv.Chaudhuri@inria.fr

Abstract. It is well-known that focusing striates a sequent derivation into phases of like polarity where each phase can be seen as inferring a synthetic connective. We present a sequent calculus of synthetic connectives based on neutral proof patterns, which are a syntactic normal form for such connectives. Different focusing strategies arise from different polarisations and arrangements of synthetic inference rules, which are shown to be complete by synthetic rule permutations. A simple generic cut-elimination procedure for synthetic connectives respects both the ordinary focusing and the maximally multi-focusing strategies, answering the open question of cut-admissibility for maximally multi-focused proofs.

1 Introduction

The story of focusing has been told several times since Andreoli [3] with essentially the same construction. The inference rules of the sequent calculus divide into two groups: invertible and non-invertible, and, during proof search, the invertible rules can be eagerly applied until only non-invertible rules remain. Then, one connective is selected for *focus* and a maximal sequence of non-invertible rules are applied until invertible rules become available again. Focused proof search alternates between these two phases, invertible and non-invertible, or *negative* and *positive*, until no unproven goals remain. Furthermore, the principal formulas of the positive and negative rules, themselves called positive and negative, are perfectly dual, evoking a number of dualities that have recently been explained via focusing; a short list includes: call-by-value (positive) dual to call-by-name (negative) [16]; the Q-protocol (positive) dual to the T-protocol (negative) [14]; the proponent (positive) dual to the opponent (negative) [2]; and forward-chaining (positive) dual to backward-chaining (negative) [7] (see [18] for a survey).

Proof theoretically, the innovation of focusing is not its operational interpretation, however, but the derived notion of a *synthetic connective*. If the operand of a positive (resp. negative) connective is itself positive (resp. negative), then the two connectives fuse into a larger positive (resp. negative) connective; this fusion eventually synthesises connectives whose operands have the opposite polarity and whose internal structure does not matter; thus, $- \otimes (- \otimes -)$ and $(- \otimes -) \otimes -$ are essentially the same ternary positive synthetic connective when applied to three negative operands. A focused derivation amounts to a derivation using synthetic inference rules for such synthetic connectives.

But are synthetic connectives true connectives, and can one construct a traditional sequent calculus based on synthetic inferences? A first approximation to an answer is to note that the focused sequent calculus admits identity and cut-elimination [7,14], so *provability*-wise there is no problem with synthetic inferences. However, if one looks at the *proofs* themselves, the question gets more clouded. Ordinary unfocused rules of like polarity freely permute with each other (as long as not prohibited by the subformula relation), but synthetic positive rules never do. Indeed, the question of positive-positive permutation is meaningless in the standard presentation of focusing because the neighbours of a positive phase are negative. Yet, a positive-positive permutation is not a dubious concept: consider the sequent $\vdash a^\perp \otimes \perp, b^\perp \otimes \perp, a, b, \mathbf{1}$, for instance, where indeed the two $\otimes$ rules are non-interfering and permute. In game-theoretic terms, the equivalent $\otimes$ moves are *truly concurrent*; however, the focused sequent calculus can only represent a serialisation, which has been a long standing criticism of focusing *qua* syntax [1]. This limitation was partially removed in [5] by the use of *multi-focusing* to represent the truly concurrent foci, but in that work the question of synthetic permutation was answered by discarding one half of the synthetic connectives (the negatives) in their entirety– very unsatisfactory!

The second break in the proof theory of synthetic derivations comes from the explicit polarity switches or *delays*. These switching connectives are commonly used to define a so-called *polarised linear logic* [10,12]; in fact, careful use of switches allows one to mimic "strongly focusing" calculi like LJQ and LJT in a general focusing framework [14]. This connective, however, has an incarnation only in the polarised world; the unfocused calculus cannot detect it (*e.g.*, by cut-reduction). Why should one countenance the invention of new connectives from one's choice of proofs?

In this paper we propose an answer to both questions by paring focusing down into three orthogonal concepts, each involving choices specific forms, that are usually confused together: *neutral expressions*, *proof patterns*, and *focusing strategies*. To be concrete, we limit ourselves to propositional MALL, although the construction itself is as general as focusing. The central component, proof patterns, is a technical device used to construct static synthetic inference rules by explicitly representing a normal form of the branching *search tree* for a synthetic connective. Synthetic identity, synthetic cut-elimination, and synthetic permutations can be defined by simple analysis of these static proof patterns. The other two concepts are dynamic. Neutral expressions represent the polarisation of a connective by recording the phase changes with the switch; however, the switches themselves are not connectives but part of the dynamics of the proof: prolonging phases as long as possible, as in ordinary focusing, is a *maximal* polarisation, while switching phases always, as in unfocused calculi, is a *minimal polarisation*. The third concept, focusing strategies, defines recipes for applying the synthetic inference rules; ordinary focusing is a negative-eager strategy, and maximal multi-focusing [5] is a positive-eager strategy. The generic synthetic cut elimination preserves both these strategies using a priority assignment to cut permutations.

Our main departure from ordinary focusing is an abolishment of the regimented phase alternation. Relaxed phase alternation is a recent innovation in focusing, but it has appeared at least three times before: first in [15] where the *focalisation graph* exposed several roots that could be simultaneously or sequentially focused on, second in [9] where a neutral game is forced to pick multiple positive foci to maintain parity with the multiple negative "foci", and third in [5] which recovers a limited form of permutative canonicity by requiring the foci to be as large as possible. Our presentation is reminiscent of several related exegeses of focusing: the generic cut-elimination is present in ludics [10], even though it is phase alternating (and monistic, which our presentation decidedly isn't); proof patterns are a generalisation of a similar construct in the CU calculus [18]; neutral expressions are present in [9], although their use there was to define a neutral game that grows a dual pair of mutually normalising derivations.

This paper is organised as follows. In sec. 2, neutral expressions and proof patterns are formally introduced and the sequent calculus of synthetic inferences is defined. The key identity (theorem. 11) and cut (theorem. 12) theorems are proved. In sec. 3, permutation is defined for synthetic connectives. Sec. 4 introduces focusing strategies and sketches the two main variants: ordinary and maximally multi-focused.

2 A sequent calculus of synthetic connectives

This section will reconstruct a cut-free sequent calculus, called MALL-S, of synthetic connectives and inference rules for propositional multiplicative-additive linear logic. MALL is selected for simplicity and clarity; it contains the important features of focusing without the distraction of polarising the exponentials. We shall adopt a polarised syntax similar to [13], but polarised propositions will be seen as dual interpretations of neutral expressions.

Definition 1 Neutral expressions*, written E, F, etc., have the following syntax.*

$$E, F, \ldots ::= a \mid E \times F \mid 1 \mid E + F \mid 0 \mid \updownarrow E$$

Here, a represents an atomic proposition with unassigned polarity. The $\updownarrow$ operator represents an explicit switch of polarities.

Definition 2 (polarisation) *A polarised proposition is defined as either a positive or a negative polarisation of a neutral expression, given respectively by the polarisation functions $\langle - \rangle^+$ and $\langle - \rangle^-$:*

$$\langle a \rangle^+ = a^+ \quad \langle a \rangle^- = a^-$$

$$\langle E \times F \rangle^+ = \langle E \rangle^+ \otimes \langle F \rangle^+ \quad \langle 1 \rangle^+ = \mathbf{1} \quad \langle E \times F \rangle^- = \langle E \rangle^- \parr \langle F \rangle^- \quad \langle 1 \rangle^- = \bot$$

$$\langle E + F \rangle^+ = \langle E \rangle^+ \oplus \langle F \rangle^+ \quad \langle 0 \rangle^+ = \mathbf{0} \quad \langle E + F \rangle^- = \langle E \rangle^- \mathbin{\&} \langle F \rangle^- \quad \langle 0 \rangle^- = \top$$

$$\langle \updownarrow E \rangle^+ = \langle E \rangle^- \quad \langle \updownarrow E \rangle^- = \langle E \rangle^+$$

Here, a^+ (resp. a^-) refers to a positively (resp. negatively) polarised atom. We write $\langle E \rangle^\pm$ to refer to either $\langle E \rangle^+$ or $\langle E \rangle^-$, and $\langle E \rangle^\mp$ to refer to its dual polarisation.

Note that polarisation is an injection: infinitely many expressions can polarise to the same proposition by means of repeated "administrative" $\updownarrow$s. Nevertheless, all MALL propositions are either positive or negative polarisations of some expression. This restriction will need to be relaxed when moving to non-linear logics such as classical logic where the propositional connectives have ambiguous polarities. Note also that the representation of & as a polarisation of a sum (+) differs starkly from the popular view of & as a "conjunction", possibly because the two polarised interpretations of the classically ambiguous $\wedge$ are $\otimes$ and & [14]. Our view of & as a "sum" (and $\parr$ as a "product") is supported by distributivity: $A \parr (B \mathbin{\&} C) \equiv (A \parr B) \mathbin{\&} (A \parr C)$.

Indeed, the rules of the calculus will be given not for polarised neutral expressions but for an associated unique *proof pattern* that reorganises the expression into a disjunctive normal form up to the atoms or the polarity switches, using distributivity to move the $+$s to the surface through the $\times$s. This reorganisation will generally repeat a sub-expression—for example, $a \times (b + c)$ is reorganised to $(a \times b) + (a \times c)$, repeating a—but *will not* duplicate any sub-derivations because of the $\updownarrow$ guards. For instance, the duplication of the $\vdash B$ in the following derivation is syntactically prohibited.

$$\dfrac{\dfrac{\dfrac{\vdash A, C \qquad \vdash A, D}{\vdash A, C \mathbin{\&} D}\,\& \qquad \vdash B}{\vdash A \otimes B, C \mathbin{\&} D}\,\otimes}{\vdash (A \otimes B) \parr (C \mathbin{\&} D)}\,\parr \quad\equiv\quad \dfrac{\dfrac{\dfrac{\vdash A, C \qquad \vdash B}{\vdash A \otimes B, C}\,\otimes \qquad \dfrac{\vdash A, D \qquad \vdash B}{\vdash A \otimes B, D}\,\otimes}{\vdash A \otimes B, C \mathbin{\&} D}\,\&}{\vdash (A \otimes B) \parr (C \mathbin{\&} D)}\,\parr$$

The propositions $A \otimes B$ and $C \mathbin{\&} D$ are of opposite polarities, so we only observe the equivalence $\updownarrow(E \times F) \times (G + H) \equiv (\updownarrow(E \times F) \times G) + (\updownarrow(E \times F) \times H)$ (for suitable E, F, G and H), both of which have identical sub-derivations after the outer negative phase.

Definition 3 (proof patterns) Product patterns *(π) and* sum patterns *(σ) are generated by the following grammars:*

$$\pi ::= \hat{E} \mid 1 \mid \pi_1 \cdot \pi_2 \qquad\qquad \sigma ::= \pi \mid 0 \mid \sigma_1 + \sigma_2$$

where $\hat{E}$ is either an atom or of the form $\updownarrow E$. The structures $\langle \Pi, \cdot, 1\rangle$ and $\langle \Xi, +, 0\rangle$ are commutative monoids, where Π is the set of product patterns and Ξ the set of sum patterns. A product pattern will always be depicted in its normal form $\prod_{i \in I} \hat{E}_i$, and a sum pattern similarly as $\sum_{i \in I} \pi_i$, where I is a finite index set. The unqualified term "proof pattern" will refer to sum patterns.

Notation 4 *We write $\Downarrow\!\left\langle \hat{E} \right\rangle^{\pm}$ for $(a)^{\pm}$ if $\hat{E} = a$ and for $\langle F\rangle^{\mp}$ if $\hat{E} = \updownarrow F$.*

Abstractly, a proof pattern $\sum_{i \in I} \pi_i$ represents the proof search tree for a synthetic connective. In the positive interpretation of the connective, the outer sum represents the disjunctive ($\oplus$) choices, while each π_i represents the multiplicative ($\otimes$) structure. In the negative interpretation, the outer sum represents the alternatives (&), while the inner product represents the sequent structure ($\parr$).

Definition 5 *Given two patterns $\sigma = \sum_{i \in I} \pi_i$ and $\sigma' = \sum_{j \in J} \pi'_j$, their* product, *written $\sigma \times \sigma'$ is the pattern $\sum_{i \in I} \sum_{j \in J} \pi_i \cdot \pi'_j$.*

Fact 6 *$\langle \Xi, +, 0, \times, 1\rangle$ is a commutative semiring.* □

Every neutral expression has a corresponding pattern derived simply by treating the expression constructors as these semiring operators. Precisely, we can define a judgement $\sigma \Vdash E$ with the following rules:

$$\dfrac{}{a \Vdash a} \qquad \dfrac{}{\updownarrow E \Vdash \updownarrow E} \qquad \dfrac{\sigma \Vdash E \qquad \sigma' \Vdash F}{\sigma \times \sigma' \Vdash E \times F} \qquad \dfrac{}{1 \Vdash 1} \qquad \dfrac{\sigma \Vdash E \qquad \sigma' \Vdash F}{\sigma + \sigma' \Vdash E + F} \qquad \dfrac{}{0 \Vdash 0}$$

For example, $(a \cdot c) + (a \cdot \updownarrow E) + c + \updownarrow E \Vdash (a + 1) \times (c + \updownarrow E)$.

Before proceeding further, we note that a similar notion of "pattern" has been developed in the realm of focused lambda calculi in the system CU [18,17]. The differences are as follows: we define patterns for both positive and negative propositions and furthermore represent both the product and the sum structure; patterns in CU, on the other hand, are defined only for the positive propositions and keep just the products, forgetting one half of each sum. CU patterns are therefore a slice of the disjunctive normal form—indeed, the disjunctive normal forms cannot be computed at all unless the sums are represented—which necessitates quantification over (their equivalent of) the $\Vdash$ judgement to recover the full sum. This quantification makes the synthetic rules in CU higher-order. (This is most likely by design, as CU is intended as a logical explanation for higher-order abstract syntax.)

Definition 7 (contexts and sequents)

- *A* context *is a finite multiset of expressions annotated with polarities,* $(E_1)^\pm, \ldots, (E_n)^\pm$, *written* $\varDelta$. *Two contexts that differ only on* $(E)^\pm$ *and* $(F)^\pm$ *with* $E \neq F$ *are considered different even if* $(E)^\pm = (F)^\pm$ *(defn. 2).*
- *A* sequent $\vdash \varDelta$, *where* $\varDelta$ *is a context, is a judgement that the context* $\varDelta$ *is linearly contradictory. The form* $\vdash_\xi \varDelta$ *is used to indicate that* ξ *is a derivation of* $\vdash \varDelta$.
- *A* focused sequent *is a structure of the form* $\vdash \varDelta \,;\, (\pi)^\pm$ *where* $\varDelta$ *is a context and* $(\pi)^\pm$ *is a product pattern* π *annotated with a polarisation.*

Definition 8 *Let* $\mathbf{D} = \langle \varDelta_i \rangle_{i \in 1..n}$ *be a list of contexts. Then,*

1. $\bigotimes \mathbf{D}$ *stands for the list of sequents* $\vdash \varDelta_1, \ldots, \vdash \varDelta_n$.
2. $⅋\, \mathbf{D}$ *stands for the sequent* $\vdash \varDelta_1, \ldots, \varDelta_n$.

For non-atomic principal formulas, we define the synthetic rules on the proof pattern of the underlying neutral expression. The outer sum in the pattern represents an enumeration of choices. If the corresponding proposition were positively polarised, then it represents a disjunction of choices of which only one needs to be taken. For a negatively polarised proposition it represents an alternation of choices, *all* of which must be taken. The *bottom half* of the two synthetic rules thus looks as follows:

$$\frac{\sum_{i \in I} \pi_i \Vdash E \qquad \exists u \in I.\ \vdash \varDelta \,;\, (\pi_u)^+}{\vdash \varDelta, (E)^+}\ \mathrm{P}{\downarrow} \qquad \frac{\sum_{i \in I} \pi_i \Vdash E \qquad \forall u \in I.\ \vdash \varDelta \,;\, (\pi_u)^-}{\vdash \varDelta, (E)^-}\ \mathrm{N}{\downarrow}$$

In each case, we obtain a number of focused sequents of the form $\vdash \varDelta \,;\, (\pi)^\pm$. If the polarisation is positive, i.e., the product pattern represents a $\otimes$, then the context $\varDelta$ must be distributed into the components of the product. This *demultiplexion* operation is defined by a ternary relation.

Definition 9 (demultiplexion) *Given a context* $\varDelta$ *and a product pattern* $\pi = \prod_{i \in 1..k} \hat{E}_i$, *a* demultiplexion *of* $\varDelta$ along $(\pi)^\pm$ *produces a list of contexts* $\mathbf{D} = \varDelta_1 \,; \cdots ; \varDelta_k$, *written* $\varDelta \,;\, (\pi)^+ \gg \mathbf{D}$, *generated by the following rules:*

$$\frac{\varDelta \,;\, (\pi)^+ \gg \mathbf{D}}{\varDelta, (a)^- \,;\, (\pi \cdot a)^+ \gg \mathbf{D} \,;\, (a)^-, (a)^+}$$

$$\frac{\mathit{\Delta}\,;(\pi)^- \gg \mathbf{D}}{\mathit{\Delta},\mathit{\Delta}'\,;(\pi\cdot a)^- \gg \mathbf{D}\,;\mathit{\Delta}',(a)^-} \qquad \frac{\mathit{\Delta}\,;(\pi)^\pm \gg \mathbf{D}}{\mathit{\Delta},\mathit{\Delta}'\,;(\pi\cdot \updownarrow E)^\pm \gg \mathbf{D}\,;\mathit{\Delta}',(E)^\mp}$$

The upper half of the positive rule, P↑, is then obvious: we select a demultiplexion of the context along the positive pattern and interpret every context in it as a sequent.

$$\frac{\exists \mathbf{D} : (\mathit{\Delta}\,;(\pi)^+ \gg \mathbf{D}).\ \bigotimes \mathbf{D}}{\vdash \mathit{\Delta}\,;(\pi)^+}\ \mathrm{P}\!\uparrow$$

Somewhat surprisingly, the upper half of the negative rule, N↑, can be written analogously:

$$\frac{\forall \mathbf{D} : (\mathit{\Delta}\,;(\pi)^- \gg \mathbf{D}).\ \bigparr \mathbf{D}}{\vdash \mathit{\Delta}\,;(\pi)^-}\ \mathrm{N}\!\uparrow$$

The demultiplexion used to construct **D** is simply undone by the $\bigparr$ operator. It therefore does not matter how the demultiplexion is done, and there is always a way to demultiplex along a negative pattern (for example, all of the context can be "sent" to the first element of the product pattern, if one exists). The premise of the N↑ rule is therefore uniquely determined.

We shall henceforth ignore the halves of the rules and just consider the combined rules P and N, each of which is a polarisation of the following neutral rule (where E is non-atomic and $\sum_{i\in I} \pi_i \Vdash E$):

$$\frac{Qu \in I.\ Q\mathbf{D} : (\mathit{\Delta}\,;(\pi_u)^p \gg \mathbf{D}).\ \bigcirc \mathbf{D}}{\vdash \mathit{\Delta},(E)^p}\ \mathrm{R}(p, Q, \bigcirc)$$

In the positive interpretation $\mathrm{P} = \mathrm{R}(+, \exists, \bigotimes)$, there is one premise corresponding to each element of a demultiplexion of the context along a product pattern; in the negative interpretation $\mathrm{N} = \mathrm{R}(-, \forall, \bigparr)$, there is one premise for each element of the outer sum pattern. The rules have been written in this way to highlight the precise duality of their premises.

It is important to note that the use of the ∀ and ∃ in the rules is merely a notational device. As the pattern that corresponds to an expression is statically known, we actually have instances of the P and N rules specialised to the index sets of these statically known patterns. Thus, the P and N rules are "first-order": they do not depend on reasoning in the meta-language.

There is one additional synthetic rule for atomic propositions:

$$\frac{}{\vdash (a)^+, (a)^-}\ \mathrm{I}$$

This rule is neither positive, nor negative, and can only be applied at the leaves of a derivation. As usual, it is only defined for atoms, but it can be proved for arbitrary expressions. This is a syntactic completeness theorem that is usually called the *identity principle*. Its proof is almost immediate in MALL-S.

Notation 10 *We shall adopt the following notational shorthands:* $\{\mathit{\Delta}_i\}_{i\in I}$ *for the multiset union of the* $\mathit{\Delta}_i$, $\langle \mathit{\Gamma}_i\rangle_{i\in I}$ *for the list of contexts* $\mathit{\Gamma}_i$, *and* $\langle\vdash \mathit{\Gamma}_i\rangle_{i\in I}$ *for the list of premises* $\vdash \mathit{\Gamma}_i$.

Theorem 11 (identity principle) *The sequent* $\vdash (E)^+, (E)^-$ *is derivable for any* E.

Proof. We reason by induction on the structure of the proof pattern for E. The atomic case follows simply by I. For the non-atomic cases, suppose $\sum_{i\in I}\prod_{j\in J_i}\hat{E}_{ij} \Vdash E$. We have:

$$\dfrac{\forall u \in I. \quad \dfrac{\left\langle \vdash \Downarrow\left(\hat{E}_{uj}\right)^+, \Downarrow\left(\hat{E}_{uj}\right)^- \right\rangle_{j\in J_u}}{\vdash (E)^+, \left\{\Downarrow\left(\hat{E}_{uj}\right)^-\right\}_{j\in J_u}}\ \mathrm{P}}{\vdash (E)^+, (E)^-}\ \mathrm{N}$$

because

$$\left\{\Downarrow\left(\hat{E}_{uj}\right)^-\right\}_{j\in J_u} ; \left(\textstyle\prod_{j\in J_u}\hat{E}_{uj}\right)^+ \Vdash \left\langle \Downarrow\left(\hat{E}_{uj}\right)^+, \Downarrow\left(\hat{E}_{uj}\right)^-\right\rangle_{j\in J_u}.$$

We then use the induction hypothesis on the sequents of the form $\vdash \Downarrow\left(\hat{E}_{uj}\right)^+, \Downarrow\left(\hat{E}_{uj}\right)^-$, which contain strictly smaller expressions. □

For syntactic soundness, we turn to admissibility of the following *cut* rule:

$$\dfrac{\vdash \varDelta, (E)^+ \qquad \vdash \varGamma, (E)^-}{\vdash \varDelta, \varGamma}\ \mathrm{C}$$

Theorem 12 (cut elimination) *The C rule is admissible*

Proof. We shall prove this by first admitting C as an inference rule and then eliminating it by (non-deterministically) rewriting it out of a proof that uses it. This rewrite $\longrightarrow$ is generated from the following cases.

- Initial cuts, where one of the premises is derived from I. In these cases, the elimination is trivial because we can just drop the C and the initial premise.
- Principal cuts, where the cut-expression is principal in P and N in the two premises. Suppose $\sum_{i\in I}\pi_i \Vdash E$, $\pi_i = \prod_{j\in J_i}\hat{E}_{ij}$ and $u \in I$, such that:

$$\dfrac{\dfrac{\left\langle \vdash_{\xi(j)} \varGamma_j, \Downarrow\left(\hat{E}_{uj}\right)^+ \right\rangle_{j\in J_u}}{\vdash \{\varGamma_j\}_{j\in J_u}, (E)^+}\ \mathrm{P} \qquad \dfrac{\forall i \in I.\ \vdash_{\zeta(i)} \varDelta, \left\{\Downarrow\left(\hat{E}_{ij}\right)^-\right\}_{j\in J_i}}{\vdash \varDelta, (E)^-}\ \mathrm{N}}{\vdash \{\varGamma_j\}_{j\in J_u}, \varDelta}\ \mathrm{C}$$

  Let $\phi : 1..n \to J_u$ be a bijection. We rewrite the above cut as follows:

$$\dfrac{\vdash_{\xi(\phi(1))} \varGamma_{\phi(1)}, \Downarrow\left(\hat{E}_{u\phi(1)}\right)^+ \qquad \dfrac{\vdash_{\xi(\phi(2))} \varGamma_{\phi(2)}, \Downarrow\left(\hat{E}_{u\phi(2)}\right)^+ \qquad \begin{array}{c}\dfrac{\vdash_{\xi(\phi(n))} \varGamma_{\phi(n)}, \Downarrow\left(\hat{E}_{u\phi(n)}\right)^+ \qquad \vdash_{\zeta(u)} \varDelta, \left\{\Downarrow\left(\hat{E}_{u\phi(k)}\right)^-\right\}_{k\in 1..n}}{\vdash \varGamma_{\phi(n)}, \varDelta, \left\{\Downarrow\left(\hat{E}_{u\phi(k)}\right)^-\right\}_{k\in 1..n-1}}\ \mathrm{C} \\ \vdots\end{array}}{\vdash \{\varGamma_{\phi(k)}\}_{k\in 2..n}, \varDelta, \Downarrow\left(\hat{E}_{i\phi(1)}\right)^-}\ \mathrm{C}}{\vdash \{\varGamma_{\phi(k)}\}_{k\in 2..n}, \varDelta}\ \mathrm{C}$$

  Each instance of C is now on a strictly smaller cut expression.

– Commutative cuts, where the cut-expression is not principal in one derivation. In each of the following two cases, we suppose that $\sum_{i\in I}\prod_{j\in J_i}\hat{F}_{ij} \Vdash F$, $u \in I$ and $v \in J_u$. The two cases of the rewrite are named [PCC] and [NCC] for positive and negative commutative cuts respectively.

$$\dfrac{\dfrac{\bigotimes \mathbf{D} \qquad \vdash_\xi \Delta_v, (E)^\pm, \Downarrow\!\left(\hat{F}_{uv}\right)^+ \qquad \bigotimes \mathbf{D}'}{\vdash \{\Delta_j\}_{j\in J_u}, (F)^+, (E)^\pm}\;\mathrm{P} \qquad \vdash_\zeta \Gamma, (E)^\mp}{\vdash \{\Delta_j\}_{j\in J_u}, (F)^+, \Gamma}\;\mathrm{C} \quad \longrightarrow_{[\mathrm{PCC}]}$$

$$\dfrac{\bigotimes \mathbf{D} \qquad \dfrac{\vdash_\xi \Delta_v, (E)^\pm, \Downarrow\!\left(\hat{F}_{uv}\right)^+ \qquad \vdash_\zeta \Gamma, (E)^\mp}{\vdash \Delta_v, (E)^\pm, \Downarrow\!\left(\hat{F}_{uv}\right)^+, \Gamma}\;\mathrm{C} \qquad \bigotimes \mathbf{D}'}{\vdash \{\Delta_j\}_{j\in J_u}, (F)^+, \Gamma}\;\mathrm{P}$$

$$\dfrac{\dfrac{\forall i \in I.\ \vdash_\xi \Delta, (E)^\pm, \left\{\Downarrow\!\left(\hat{F}_{ij}\right)^-\right\}_{j\in J_i}}{\vdash \Delta, (F)^-, (E)^\pm}\;\mathrm{P} \qquad \vdash_\zeta \Gamma, (E)^\mp}{\vdash \Delta, (F)^-, \Gamma}\;\mathrm{C} \quad \longrightarrow_{[\mathrm{NCC}]}$$

$$\dfrac{\forall i \in I. \qquad \dfrac{\vdash_\xi \Delta, (E)^\pm, \left\{\Downarrow\!\left(\hat{F}_{ij}\right)^-\right\}_{j\in J_i} \qquad \vdash_\zeta \Gamma, (E)^\mp}{\vdash \Delta, \Gamma, \left\{\Downarrow\!\left(\hat{F}_{ij}\right)^-\right\}_{j\in J_i}}\;\mathrm{C}}{\vdash \Delta, (F)^-, (E)^\pm}\;\mathrm{P}$$

In these cases the instance of C on the right is on a smaller derivation. □

The cut-elimination proof above is remarkable for several reasons. First, it is a generic argument that is independent of any logical connective. Second, it is obviously correct for each cut can be seen to be smaller by inspection.[1] Lastly, it is compact: there is no important detail missing in the proof. To be sure, these remarkable properties are also observed in CU [18,17], but in MALL-S, because the rules are first-order, we do not need to depend on the meta language for the proof of coverage.

As already mentioned, a key distinguishing feature of MALL-S from other focusing systems such as LLF [3] or CU is that the positive and negative rules do not alternate. In this sense, it would be a mistake to call MALL-S a "focusing" system, so we cannot prove MALL-S complete with respect to the unfocused MALL (rules in fig. 1) by citation. Fortunately, we can easily recover the unfocused rules by selecting a suitably minimal polarisation.

Definition 13 *The* minimal polarisation $\lfloor - \rfloor$ *is given inductively as follows:*

$$\lfloor a \rfloor = a \qquad \lfloor a^\perp \rfloor = a$$

$$\lfloor A \otimes B \rfloor = \lfloor A \rfloor^- \times \lfloor B \rfloor^- \qquad \lfloor \mathbf{1} \rfloor = 1 \qquad \lfloor A \parr B \rfloor = \lfloor A \rfloor^+ \times \lfloor B \rfloor^+ \qquad \lfloor \perp \rfloor = 1$$

$$\lfloor A \oplus B \rfloor = \lfloor A \rfloor^- + \lfloor B \rfloor^- \qquad \lfloor \mathbf{0} \rfloor = 0 \qquad \lfloor A \,\&\, B \rfloor = \lfloor A \rfloor^+ + \lfloor B \rfloor^+ \qquad \lfloor \top \rfloor = 0$$

[1] This proof does get more complex if, in some extension of MALL-S, the index sets are infinite, because in that case the result of eliminating a principal cut would be a derivation of infinite depth.

$$\frac{}{\vdash a, a^{\perp}}\ \mathrm{I} \quad \frac{\vdash \Gamma, A \qquad \vdash \Delta, B}{\vdash \Gamma, \Delta, A \otimes B}\ \otimes \quad \frac{}{\vdash \mathbf{1}}\ \mathbf{1} \quad \frac{\vdash \Delta, A, B}{\vdash \Delta, A \mathbin{⅋} B}\ ⅋ \quad \frac{\vdash \Delta}{\vdash \Delta, \bot}\ \bot$$

$$\frac{\vdash \Delta, A_i}{\vdash \Delta, A_1 \oplus A_2}\ \oplus \quad \text{no } \mathbf{0} \quad \frac{\vdash \Delta, A \qquad \vdash \Delta, B}{\vdash \Delta, A \,\&\, B}\ \& \quad \frac{}{\vdash \Delta, \top}\ \top$$

Fig. 1. MALL rules

$$\lfloor P \rfloor^{-} = \updownarrow\updownarrow\lfloor P \rfloor \quad \lfloor N \rfloor^{-} = \updownarrow\lfloor N \rfloor \quad \lfloor P \rfloor^{+} = \updownarrow\lfloor P \rfloor \quad \lfloor N \rfloor^{+} = \updownarrow\updownarrow\lfloor N \rfloor$$

Here P (resp. N) refers to any positive (resp. negative) proposition.

Note that the polarisation of every strict subformula uses at least one administrative $\updownarrow$ switch, and that $(\lfloor P \rfloor)^{+} = P$ and $(\lfloor N \rfloor)^{-} = N$. The rules of the ordinary MALL calculus then reappear as MALL-S rules for these minimal polarisations.

Theorem 14 (completeness of MALL-S) *If Δ is a context of unpolarised propositions, let $\lfloor \Delta \rfloor$ represent that context that replaces every positive $P \in \Delta$ with $(\lfloor P \rfloor)^{+}$ and every negative $N \in \Delta$ with $(\lfloor N \rfloor)^{-}$. If $\vdash \Delta$ in MALL, then $\vdash \lfloor \Delta \rfloor$ in MALL-S.*

Proof (sketch). By induction on the structure of the given MALL derivation.

Case of I: the I rules in MALL and MALL-S are identical.

Case of $\otimes$: Consider $\lfloor P \otimes N \rfloor = \lfloor P \rfloor^{-} \times \lfloor N \rfloor^{-} = \updownarrow\updownarrow\lfloor P \rfloor \times \updownarrow\lfloor N \rfloor$. Its proof pattern is just $\updownarrow\updownarrow\lfloor P \rfloor \cdot \updownarrow\lfloor N \rfloor$, so its derivation is:

$$\cfrac{\cfrac{\cfrac{}{\vdash \Delta_1, (\lfloor P \rfloor)^{+}}\ \text{i.h.}}{\vdash \Delta_1, (\updownarrow\lfloor P \rfloor)^{-}}\ \mathrm{N} \qquad \cfrac{}{\vdash \Delta_2, (\lfloor N \rfloor)^{-}}\ \text{i.h.}}{\cfrac{\vdash \Delta_1, \Delta_2, (\updownarrow\updownarrow\lfloor P \rfloor \times \updownarrow\lfloor N \rfloor)^{+}}{\vdash \Delta_1, \Delta_2, (\lfloor P \otimes N \rfloor)^{+}}\ =}\ \mathrm{P}$$

because $\Delta_1, \Delta_2 \,;\, (\updownarrow\updownarrow\lfloor P \rfloor \times \updownarrow\lfloor N \rfloor)^{+} \Vdash \Delta_1, (\updownarrow\lfloor P \rfloor)^{-} \,;\, \Delta_2, (\lfloor N \rfloor)^{-}$. We thus obtain the $\otimes$ rule for $P \otimes N$ in MALL. This characteristic case shows the way the induction works for subformulas of the same and opposite polarities, and the remaining cases are similar. □

Theorem 14 shows that the ordinary unfocused MALL is recoverable in MALL-S by picking specific polarisations. This is a strong indication that synthetic rules with delays are a more primitive notion than the usual binary connectives, an observation already made in the genesis of ludics [10], but not well appreciated outside a certain section of the proof theory community.

On the other hand, theorem. 14 does not show completeness for other polarisations. In ordinary (unpolarised) focusing [3] the polarity of the rules matches the natural polarity of the principal propositions. This suggests a maximal polarisation of the ordinary connectives that contains no administrative switches.

Definition 15 *E is a* maximal polarisation *for A if $A = (E)^{\pm}$ and E contains no subexpressions of the form $\updownarrow\updownarrow E'$. We write $\lceil A \rceil$ to refer to the unique maximal polarisation of A because of the following trivial fact.*

Fact 16 *If E and F are maximal polarisations for A, then $E = F$.* ☐

Lemma 17 $\vdash (\lfloor P \rfloor)^-, (\lceil P \rceil)^+$ *and* $\vdash (\lfloor N \rfloor)^+, (\lceil N \rceil)^-$ *are derivable in* MALL-S.

Proof (Sketch). Replay the proof of theorem. 11, but in one half of the reduction replay the maximally polarised synthetic rule with many minimally polarised synthetic rules. ☐

Minimal polarisations can thus be used to simulate maximal polarisations, directly giving the key focalisation result by an appeal to synthetic cuts.

Corollary 18 (Focalisation) *If Δ is a context of unpolarised propositions, let $\lceil \Delta \rceil$ represent that context that replaces every positive $P \in \Delta$ with $(\lceil P \rceil)^+$ and every negative $N \in \Delta$ with $(\lceil N \rceil)^-$. If $\vdash \Delta$ in* MALL*, then $\vdash \lceil \Delta \rceil$ in* MALL-S*.*

Proof. By theorem. 14, $\vdash \lfloor \Delta \rfloor$ is provable in MALL-S. Cut (theorem. 12) every $(\lfloor P \rfloor)^+$ and $(\lfloor N \rfloor)^-$ in $\vdash \lfloor \Delta \rfloor$ with a proof of $\vdash (\lfloor P \rfloor)^-, (\lceil P \rceil)^+$ or $\vdash (\lfloor N \rfloor)^+, (\lceil N \rceil)^-$ (lem. 17). ☐

Proofs of the focalisation result using cut-elimination have also been attempted in [7,14], but their proofs tend to be considerably more complex than the one presented here, partly because they are based on a more traditional formulation of focusing, but also because their proofs attempt to simulate the unfocused rules directly with cut. Our approach of simulating the unfocused rules with a different polarisation and then cutting them out *ex post facto* is a more perspicuous decomposition of the focalisation result.

3 Synthetic permutations

In this section, we investigate the matter of permutations of the synthetic inference rules P and N. If the synthetic sequent calculus is to be seen as a generalisation of the ordinary calculus, then it is essential to define the corresponding generalisations of the binary permutations [11]. We write r_1/r_2 as a type of permutation where the rule(s) r_1 is (are) used immediately above r_2 and the result of the permutation moves (possibly with replication) r_2 above a single instance of r_1 without affecting the rest of the derivation. Such permutations are familiar from the ordinary (unfocused) logic; for instance the permutation $\otimes/\&$ in MALL is the following reordering:

$$\dfrac{\dfrac{\vdash_\xi \Gamma, A \qquad \vdash_\zeta \Delta, B, C}{\vdash \Gamma, \Delta, A \otimes B, C} \qquad \dfrac{\vdash_\xi \Gamma, A \qquad \vdash_\varphi \Delta, B, D}{\vdash \Gamma, \Delta, A \otimes B, D}}{\vdash \Gamma, \Delta, A \otimes B, C \mathbin{\&} D} \quad \longrightarrow \quad \dfrac{\vdash_\xi \Gamma, A \qquad \dfrac{\vdash_\zeta \Delta, B, C \qquad \vdash_\varphi \Delta, B, D}{\vdash \Delta, B, C \mathbin{\&} D}}{\vdash \Gamma, \Delta, A \otimes B, C \mathbin{\&} D}$$

Unsurprisingly, the N/N and P/P permutations are freely allowed in MALL-S.

Definition 19 (equipollent permutation)
Suppose $\sum_{i \in I} \prod_{j \in J_i} \hat{E}_{ij} \Vdash E$ and $\sum_{k \in K} \prod_{l \in L_k} \hat{F}_{kl} \Vdash F$

1. *A P/P permutation (for any $u \in I$, $v \in J_u$, $w \in K$ and $x \in L_w$) is as follows:*

$$\dfrac{\left\langle \vdash_{\xi(j)} \varDelta_j, \Downarrow\left(\hat{E}_{uj}\right)^+ \right\rangle_{j\in J_u\setminus v} \qquad \dfrac{\left\langle \vdash_{\zeta(l)} \varGamma_l, \Downarrow\left(\hat{F}_{wl}\right)^+ \right\rangle_{l\in L_w\setminus x} \quad \vdash_{\varphi} \varOmega, \Downarrow\left(\hat{E}_{uv}\right)^+, \Downarrow\left(\hat{F}_{wx}\right)^+}{\vdash \varOmega, \{\varGamma_l\}_{l\in L_w\setminus x}, (F)^+, \Downarrow\left(\hat{E}_{uv}\right)^+}\ \mathrm{P}}{\vdash \varOmega, \{\varDelta_j\}_{j\in J_u\setminus v}, \{\varGamma_l\}_{l\in L_w\setminus x}, (E)^+, (F)^+}\ \mathrm{P}$$

$$\longrightarrow \quad \dfrac{\left\langle \vdash_{\zeta(l)} \varGamma_l, \Downarrow\left(\hat{F}_{wl}\right)^+ \right\rangle_{l\in L_w\setminus x} \qquad \dfrac{\left\langle \vdash_{\xi(j)} \varDelta_j, \Downarrow\left(\hat{E}_{uj}\right)^+ \right\rangle_{j\in J_u\setminus v} \quad \vdash_{\varphi} \varOmega, \Downarrow\left(\hat{E}_{uv}\right)^+, \Downarrow\left(\hat{F}_{wx}\right)^+}{\vdash \varOmega, \{\varDelta_j\}_{j\in J_u\setminus v}, \Downarrow\left(\hat{F}_{wx}\right)^+, (E)^+}\ \mathrm{P}}{\vdash \varOmega, \{\varDelta_j\}_{j\in J_u\setminus v}, \{\varGamma_l\}_{l\in L_w\setminus x}, (E)^+, (F)^+}\ \mathrm{P}$$

2. *An N/N permutation is as follows:*

$$\dfrac{\forall i \in I. \quad \dfrac{\forall k \in K. \quad \vdash_{\xi(i,k)} \varDelta, \left\{\Downarrow\left(\hat{F}_{kl}\right)^-\right\}_{l\in L_k}, \left\{\Downarrow\left(\hat{E}_{ij}\right)^-\right\}_{j\in J_i}}{\vdash \varDelta, (F)^-, \left\{\Downarrow\left(\hat{E}_{ij}\right)^-\right\}_{j\in J_i}}\ \mathrm{N}}{\vdash \varDelta, (E)^-, (F)^-}\ \mathrm{N}$$

$$\longrightarrow \quad \dfrac{\forall k \in K. \quad \dfrac{\forall i \in I. \quad \vdash_{\xi(i,k)} \varDelta, \left\{\Downarrow\left(\hat{F}_{kl}\right)^-\right\}_{l\in L_k}, \left\{\Downarrow\left(\hat{E}_{ij}\right)^-\right\}_{j\in J_i}}{\vdash \varDelta, (E)^-, \left\{\Downarrow\left(\hat{F}_{kl}\right)^-\right\}_{l\in L_k}}\ \mathrm{N}}{\vdash \varDelta, (E)^-, (F)^-}\ \mathrm{N}$$

Equipollent permutations have no restrictions on the form of the left of $\longrightarrow$, so these permutations are always allowed. Of the two remaining permutation forms, the N/P permutation is also always valid and readily defined.

Definition 20 (N/P permutation)
Suppose $\sum_{i\in I} \prod_{j\in J_i} \hat{E}_{ij} \Vdash E$ and $\sum_{k\in K} \prod_{l\in L_k} \hat{F}_{kl} \Vdash F$. An N/P permutation is of the following form (for any $u \in I$ and $v \in J_u$)

$$\dfrac{\left\langle \vdash_{\xi(j)} \varDelta_j, \Downarrow\left(\hat{E}_{uj}\right)^+ \right\rangle_{j\in J_u\setminus v} \qquad \dfrac{\forall k \in K. \quad \vdash_{\zeta(k)} \varGamma, \Downarrow\left(\hat{E}_{uv}\right)^+, \left\{\Downarrow\left(\hat{F}_{kl}\right)^-\right\}_{l\in L_k}}{\vdash \varGamma, (F)^-, \Downarrow\left(\hat{E}_{uv}\right)^+}\ \mathrm{N}}{\vdash \{\varDelta\}_{j\in J_u}, \varGamma, (E)^+, (F)^-}\ \mathrm{P}$$

$$\longrightarrow \quad \dfrac{\forall k \in K. \quad \dfrac{\left\langle \vdash_{\xi(j)} \varDelta_j, \Downarrow\left(\hat{E}_{uj}\right)^+ \right\rangle_{j\in J_u\setminus v} \quad \vdash_{\zeta(k)} \varGamma, \Downarrow\left(\hat{E}_{uv}\right)^+, \left\{\Downarrow\left(\hat{F}_{kl}\right)^-\right\}_{l\in L_k}}{\vdash \varGamma, \{\varDelta\}_{j\in J_u}, (E)^+, \left\{\Downarrow\left(\hat{F}_{kl}\right)^-\right\}_{l\in L_k}}\ \mathrm{P}}{\vdash \{\varDelta\}_{j\in J_u}, \varGamma, (E)^+, (F)^-}\ \mathrm{N}$$

The final permutations are the P/N permutations, which are not generally permissible. In fact, writing this permutation type as P/N is somewhat misleading because actually several *coherent* instances of P in the premises of the bottom N rule will be merged.

Two P instances are coherent if, essentially, they make the same disjunctive and multiplicative choices. However, they cannot be exactly identical because they have different conclusions.

Definition 21 (P/N permutation)
Suppose $\sum_{i\in I}\prod_{j\in J_i}\hat{E}_{ij} \Vdash E$ and $\sum_{k\in K}\prod_{l\in L_k}\hat{F}_{kl} \Vdash F$. A P/N permutation is of the following form (for any $u \in K$ and $v \in L_u$)

$$\dfrac{\forall i \in I. \quad \dfrac{\left\langle \vdash_{\xi(l)} \varDelta_l, \Downarrow\!\left(\hat{F}_{ul}\right)^{+} \right\rangle_{l\in L_u\setminus v} \qquad \vdash_{\zeta(i)} \varGamma, \Downarrow\!\left(\hat{F}_{uv}\right)^{+}, \left\{\Downarrow\!\left(\hat{E}_{ij}\right)^{-}\right\}_{j\in J_i}}{\vdash \varGamma, \{\varDelta_l\}_{l\in L_u}, \left\{\Downarrow\!\left(\hat{E}_{ij}\right)^{-}\right\}_{j\in J_i}, (F)^{+}}\,\mathrm{P}}{\vdash \varGamma, \{\varDelta_l\}_{l\in L_u}, (E)^{-}, (F)^{+}}\,\mathrm{N}$$

$$\longrightarrow \quad \dfrac{\left\langle \vdash_{\xi(l)} \varDelta_l, \Downarrow\!\left(\hat{F}_{ul}\right)^{+} \right\rangle_{l\in L_u\setminus v} \qquad \dfrac{\forall i \in I. \quad \vdash_{\zeta(i)} \varGamma, \Downarrow\!\left(\hat{F}_{uv}\right)^{+}, \left\{\Downarrow\!\left(\hat{E}_{ij}\right)^{-}\right\}_{j\in J_i}}{\vdash \varGamma, \Downarrow\!\left(\hat{F}_{uv}\right)^{+}, (E)^{-}}\,\mathrm{N}}{\vdash \varGamma, \{\varDelta_l\}_{l\in L_u}, (E)^{-}, (F)^{+}}\,\mathrm{P}$$

Observe how this permutation is restricted: all instances of P above the N must pick the same term in the sum pattern of $(F)^+$, must have all premises but one (the $\xi(l)$) exactly identical and independent of E, and the remaining premise (the $\zeta(i)$) in each case must contain all the subexpressions of E in that position in its sum pattern.

All the permutations defined are quite obviously sound (each application of P and N is correct), so we state the following lemma without proof.

Lemma 22 *The equipollent, P/N and N/P permutations are sound, i.e., if the left then the right hand side of $\longrightarrow$.* □

We end this section with a sketched proof that the N rule is invertible, using only synthetic permutations, instead of proving it in the usual way using cuts.

Theorem 23 *The N rule is invertible.*

Proof (Sketch). Since both N/N and N/P permutations are always allowed, every N rule can be permuted repeatedly towards the goal. Hence, for any derivation of $\vdash \varGamma, (E)^-$ that contains an instance of N for $(E)^-$, there is an equivalent derivation that begins with that N rule. If there are no instances of N in the proof of $\vdash \varGamma, (E)^-$, then there must be a subderivation that proves $\vdash \varGamma', (E)^-, (0)^-$ with N (because E is non-atomic). An instance of N for $(E)^-$ can be inserted here; each of its premises will contain $(0)^-$ and will therefore be provable. This instance of N can now be permuted to the goal. □

4 Strategies

In this section we shall assume that all polarisations are maximal (defn. 15).

As already seen, the MALL-S calculus, despite being a calculus of synthetic connectives, is more permissive in the order of synthetic rules than other focusing calculi such as LLF. However, LLF is recoverable in MALL-S as a *strategy* of applying inference rules to refine the goal sequent. By theorem. 23, the N rule is invertible, so it can always be applied to remove negatively polarised expressions from the context. Such propositions are only introduced to the context by the P rule. Therefore, ordinary focusing is a strategy of eagerly applying the N rules.

Definition 24 (ordinary focusing strategy) *The* focused *proofs of* $\vdash \Delta$ *are those that:*

1. *end in I; or*
2. *end with N if there are any negatively polarised propositions in* Δ *with the premises of the rule also focused; or*
3. *end with P if there are no negatively polarised propositions in* Δ *and all the premises of that rule are also focused.*

It is easy to see that this strategy degenerates to the familiar phase alternation after the pre-existing negatively polarised propositions in the goal sequent are removed, for each P step produces at-most one negatively polarised proposition in each premise, which upon decomposition produces only positively polarised premises. The completeness of this strategy is immediate from the invertibility of N (theorem. 23) and focalisation (cor. 18). We also state the following rather obvious instance of the cut-elimination algorithm (which is a synthetic restatement of the T-permutation [8]); we omit the proof.

Theorem 25 (focused cut-elimination) *The cut-elimination rewrite* $\longrightarrow$ *of theorem. 12 preserves focused proofs if the [NCC] case is given a higher precedence than the [PCC] case. That is, given focused proofs of* $\vdash \Gamma, (A)^{+}$ *and* $\vdash \Delta, (A)^{-}$*, the cut on A is eliminated to give a focused proof of* $\vdash \Gamma, \Delta$*.* ◻

As expected, this is not the sole interesting strategy for MALL-S proofs. In [5], a notion of *maximally multi-focused* proofs is introduced, which aims to equate all permutatively isomorphic MALL proofs in a unique syntax; such a proof exhibits the "true concurrency" inherent in the selection of foci. Maximality was defined there as a terminating permutative rewrite enlarging the principal formulas in a focusing calculus with multiple foci. The recipe in [5] for generating maximally multi-focused proofs from complete proofs can easily be repeated for MALL-S, but we do not pursue that direction here; instead, we characterise them here in terms of a strategy. In these maximal proofs the P rule rather than the N rule is eagerly applied, giving a tantalisingly dual picture from the strategy that generates ordinary LLF proofs.

Definition 26 (maximal multi-focusing strategy) *The* maximally multi-focused *(abbreviated as* maximal*) proofs of* $\vdash \Delta$ *are those proofs that:*

1. *end in I; or*
2. *end in P if there are any positively polarised propositions in* Δ *such that applying P (reading backwards) leads to a proof, and all the premises of this rule are also maximal; or*
3. *end in N if the situation for (2) does not apply, and the premises of this rule are also maximal.*

We state without the rather technical proof that the maximally multi-focused proofs in the sense of [5] are exactly those proofs in the above class up to equipollent permutations. Instead, we consider the question that was left open in [5] with regard to cut-elimination on maximal proofs.

Theorem 27 (maximal cut-elimination) *The cut-elimination rewrite $\longrightarrow$ in theorem. 12 preserves maximality if the [PCC] case is given a higher precedence than the [NCC] case. That is, given maximal proofs of $\vdash \Gamma, (A)^{+}$ and $\vdash \Delta, (A)^{-}$, the cut on A is eliminated to give a maximal proof of $\vdash \Gamma, \Delta$.*

Proof (Sketch). We reason by induction on the structure of the two input derivations for the cut being eliminated. The initial and principal cases are a straightforward application of the induction hypothesis. For the commutative cases, the result of the cut-elimination can only be maximal if the instances of P are kept closer to the root of the derivation, which requires prioritising a [PCC] rewrite over an [NCC] rewrite. □

The duality of theorems 25 and 27 is surprisingly clean, which gives further credence to the notion of maximally multi-focused proof. Of course, the above strategy is not implementable in a purely backwards reasoning (goal upwards) algorithm as it quantifies over proofs of the conclusion, which will not be available until the search completes. However, proofs in this class can be generated by saturation-based forward reasoning (axioms downwards) algorithms, such as in the inverse method [6]. Such search algorithms incrementally build a database of proved facts, so whenever a rule is applied resulting in a new fact, the N rules are permuted upwards in the proof as much as possible.

5 Conclusions and future work

We have presented a reconstruction of the calculus of synthetic connectives by means of polarisations of neutral proof patterns. Among the key technical merits of this presentation are simple and obviously correct proofs of cut, identity, focalisation, and synthetic permutations. We have shown how, among the proofs using synthetic inferences, the ordinary and maximally multi-focused proofs can be seen as diametrically opposite strategies, and demonstrated that the generic synthetic cut-elimination with a priority assignment preserves maximality.

The obvious next step is to extend the development to interesting fragments larger than MALL. We have already extended it to the exponentials and first-order quantifiers, but have left them out of this paper for presentational clarity and lack of space. Extending to infinite proof patterns (both sums and products) would also be interesting, requiring finite presentations of infinite operations; of particular interest is the question of proof patterns corresponding to the (co-)inductive connectives of μMALL [4]. Yet another important extension would be to a proof of synthetic cut-elimination for second-order MALL. Precise comparisons of synthetic derivations to game semantics would also be instructive.

References

1. Samson Abramsky. Sequentiality vs. concurrency in games and logic. *Mathematical Structures in Computer Science*, 13(4):531–565, August 2003.
2. Samson Abramsky and Paul-André Melliès. Concurrent games and full completeness. In *14th Symp. on Logic in Computer Science*, pages 431–442. IEEE Computer Society Press, 1999.
3. Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *J. of Logic and Computation*, 2(3):297–347, 1992.
4. David Baelde and Dale Miller. Least and greatest fixed points in linear logic. In N. Dershowitz and A. Voronkov, editors, *International Conference on Logic for Programming and Automated Reasoning (LPAR)*, volume 4790 of *LNCS*, pages 92–106, 2007.
5. Kaustuv Chaudhuri, Dale Miller, and Alexis Saurin. Canonical sequent proofs via multi-focusing. In Giorgio Ausiello, Juhani Karhumäki, Giancarlo Mauri, and Luke Ong, editors, *Fifth IFIP International Conference on Theoretical Computer Science*, volume 273 of *IFIP International Federation for Information Processing*, pages 383–396. Boston: Springer, September 2008.
6. Kaustuv Chaudhuri and Frank Pfenning. Focusing the inverse method for linear logic. In C.-H. Luke Ong, editor, *CSL 2005: Computer Science Logic*, volume 3634 of *LNCS*, pages 200–215. Springer, 2005.
7. Kaustuv Chaudhuri, Frank Pfenning, and Greg Price. A logical characterization of forward and backward chaining in the inverse method. *J. of Automated Reasoning*, 40(2-3):133–177, March 2008.
8. V. Danos, J.-B. Joinet, and H. Schellinx. LKT and LKQ: sequent calculi for second order logic based upon dual linear decompositions of classical implication. In J.-Y. Girard, Y. Lafont, and L. Regnier, editors, *Advances in Linear Logic*, number 222 in London Mathematical Society Lecture Note Series, pages 211–224. Cambridge University Press, 1995.
9. Olivier Delande and Dale Miller. A neutral approach to proof and refutation in MALL. In F. Pfenning, editor, *23th Symp. on Logic in Computer Science*, pages 498–508. IEEE Computer Society Press, 2008.
10. Jean-Yves Girard. Locus solum. *Mathematical Structures in Computer Science*, 11(3):301–506, June 2001.
11. Stephen Cole Kleene. Permutabilities of inferences in Gentzen's calculi LK and LJ. *Memoirs of the American Mathematical Society*, 10:1–26, 1952.
12. Olivier Laurent. *Etude de la polarisation en logique*. Thèse de doctorat, Université Aix-Marseille II, March 2002.
13. Olivier Laurent. Syntax vs. semantics: a polarized approach. Prépublication électronique PPS//03/04//no 17 (pp), Laboratoire Preuves, Programmes et Systèmes, March 2003. Submitted for publication.
14. Chuck Liang and Dale Miller. Focusing and polarization in linear, intuitionistic, and classical logics. Accepted to *Theoretical Computer Science*, 2008.
15. Dale Miller and Alexis Saurin. From proofs to focused proofs: a modular proof of focalization in linear logic. In J. Duparc and T. A. Henzinger, editors, *CSL 2007: Computer Science Logic*, volume 4646 of *LNCS*, pages 405–419. Springer, 2007.
16. Philip Wadler. Call-by-value is dual to call-by-name. In *8th Int. Conf. on Functional Programming*, pages 189–201, 2003.
17. Noam Zeilberger. Focusing and higher-order abstract syntax. In George C. Necula and Philip Wadler, editors, *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008*, pages 359–369. ACM, 2008.
18. Noam Zeilberger. On the unity of duality. *Ann. of Pure and Applied Logic*, 2008. To appear.