



Citation for published version:

Heijltjes, W, Hughes, D & Strassburger, L 2019, Proof nets for first-order additive linear logic. in H Geuvers & H Geuvers (eds), *Fourth International Conference on Formal Structures for Computation and Deduction (FSCD 2019)*. vol. 131, 22, Leibniz International Proceedings in Informatics, vol. 2019, Schloss Dagstuhl- Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, pp. 22:1-22:22, Fourth International Conference on Formal Structures for Computation and Deduction, Dortmund, Germany, 24/06/19.
<https://doi.org/10.4230/LIPIcs.FSCD.2019.22>

DOI:

[10.4230/LIPIcs.FSCD.2019.22](https://doi.org/10.4230/LIPIcs.FSCD.2019.22)

Publication date:

2019

Document Version

Publisher's PDF, also known as Version of record

[Link to publication](#)

Publisher Rights

CC BY

University of Bath

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Proof Nets for First-Order Additive Linear Logic

Willem B. Heijltjes

University of Bath, United Kingdom
<http://willem.heijltj.es>

Dominic J. D. Hughes

Logic Group, UC Berkeley, USA
<http://boole.stanford.edu/~dominic>

Lutz Straßburger

Inria Saclay, Palaiseau, France
LIX, École Polytechnique, Palaiseau, France
<http://www.lix.polytechnique.fr/Labo/Lutz.Strassburger>

Abstract

We present canonical proof nets for first-order additive linear logic, the fragment of linear logic with sum, product, and first-order universal and existential quantification. We present two versions of our proof nets. One, *witness nets*, retains explicit witnessing information to existential quantification. For the other, *unification nets*, this information is absent but can be reconstructed through unification. Unification nets embody a central contribution of the paper: first-order witness information can be left implicit, and reconstructed as needed. Witness nets are canonical for first-order additive sequent calculus. Unification nets in addition factor out any inessential choice for existential witnesses. Both notions of proof net are defined through coalescence, an additive counterpart to multiplicative contractibility, and for witness nets an additional geometric correctness criterion is provided. Both capture sequent calculus cut-elimination as a one-step global composition operation.

2012 ACM Subject Classification Theory of computation → Proof theory; Theory of computation → Linear logic

Keywords and phrases linear logic, first-order logic, proof nets, Herbrand's theorem

Digital Object Identifier 10.4230/LIPIcs.FSCD.2019.22

Related Version A full version [13] is available at <http://hal.inria.fr/hal-01867625/>.

Funding *Willem B. Heijltjes*: was supported by EPSRC Project EP/R029121/1 *Typed Lambda-Calculi with Sharing and Unsharing*.

Lutz Straßburger: was supported by the ANR-FWF international grant ANR-5-CE25-0014 *The Fine Structure of Formal Proof Systems and their Computational Interpretations (FISP)*.

Acknowledgements We would like to thank the anonymous referees for their constructive feedback. Dominic Hughes thanks his hosts, Wes Holliday and Dana Scott, at the UC Berkeley Logic Group.

1 Introduction

Additive linear logic (ALL) is the logic of sum, product, and their canonical morphisms: projections, injections, and diagonals. Semantically, the logic represents parallel communication between two parties, with sum and product as respectively the sending and receiving of a binary choice [22, 3]. As such it is a core part of *session types* for process calculi [16, 2, 28].

A microcosm of parallelism, ALL already demonstrates the *Blass problem* of game semantics [1], that sequential strategies do not in general have associative composition. This is resolved by *proof nets* [7, 21], which are a canonical, *true-concurrency* presentation of ALL.

Here, we extend proof nets to first-order additive linear logic (ALL1). Beyond the solution to the proof-net problem, a main contribution is the (further) development of the two techniques we consider: *explicit substitutions* for witness assignment, and reconstruction of



© Willem B. Heijltjes, Dominic J. D. Hughes, and Lutz Straßburger;
licensed under Creative Commons License CC-BY

4th International Conference on Formal Structures for Computation and Deduction (FSCD 2019).

Editor: Herman Geuvers; Article No. 22; pp. 22:1–22:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany


22:2 Proof Nets for First-Order Additive Linear Logic

witness information through *unification*, as pioneered for MLL by the second author [18]. We expect to apply these to first-order logics more generally.

Proofs have been relegated to the appendix – for a full version see the report [13].

Additive proof nets

ALL proof nets [21, Section 4.10] represent a morphism $A \rightarrow B$ by a sequent $\vdash \bar{A}, B$ plus a **linking**, a relation between the propositional atoms of \bar{A} (the dual of A) and those of B . They are **canonical**: they factor out the permutations of sequent calculus, and correspond 1–1 to morphisms of the free category with binary sums and products. **Composition**, of proof nets over $\vdash \bar{A}, B$ and $\vdash \bar{B}, C$ to one over $\vdash \bar{A}, C$, is by the relational composition of their linkings along the dual formulas B and \bar{B} and captures sequent-calculus cut-elimination. Below are examples of proof nets and their composition.



We extend additive proof nets with first-order quantification. Our central challenge is to incorporate the essential content of first-order proof, the **witness assignment** to existential quantifiers. Commonly, as in the sequent calculus rule below left, a witness to $\exists x.B$ is given by an immediate substitution $B[t/x]$. To assign different witnesses in different branches of a proof, the subformula is duplicated first, giving $B[s/x]$ and $B[t/x]$, as below right.

$$\frac{\vdash A, B[t/x]}{\vdash A, \exists x.B} \exists R, t \qquad \frac{\frac{\frac{\vdash \bar{P}(s), P(s)}{\vdash \exists x.\bar{P}(x), P(s)} \exists R, s \quad \frac{\vdash \bar{P}(t), P(t)}{\vdash \exists x.\bar{P}(x), P(t)} \exists R, t}}{\vdash \exists x.\bar{P}(x), P(s) \times P(t)} \times R$$

This is incompatible with a **sequent + links** proof net design, where the conclusion sequent remains intact, and a subformula B cannot be the subject of substitution or duplication. Instead, we propose two alternative treatments of witnessing terms, embodied in two notions of proof net: **witness nets** and **unification nets**. Our solutions are based on the second author's recent **unification nets** for first-order multiplicative linear logic [18]. Their main feature is to omit existential witnesses altogether, and reconstruct them by unification.

Witness nets record witness assignment in substitution maps attached to each link. The example below left shows the proof net for the sequent proof earlier. (We will assume a different variable for each quantifier, and we attach links to predicate letters, as the root connective of an atomic proposition.) Witness nets are canonical for ALL1 sequent calculus permutations. Composition is direct, where the witness assignments of links are composed through a simple process of *interaction + hiding* similar to that of game semantics [26].

Unification nets omit any witness information, as illustrated below right. In addition to canonicity, they embody a notion of **generality**: where more than one witness could be assigned, unification nets do not require a definite choice, while witness nets do. Composition is direct, by relational composition. We compare further properties in Figure 11 in the conclusion, where we also discuss related work and Lambek's notion of **generality** [23].



Background

Additive linear logic is combinatorially rich, yet well-behaved and tractable: proof search [6] and proof net correctness [12] for a net over $\vdash A, B$ are linear in $|A| \times |B|$ (with $|A|$ the size of the syntax tree of A). Proof nets remain canonical and equally tractable when extended with the two units [11, 12], and the first-order case is merely NP-complete [12].

ALL is of course part of MALL (multiplicative-additive linear logic), and its lessons are clearly visible in the second author's canonical MALL proof nets [21], as well as the first and second author's locally canonical *conflict nets* [20]. Its proof nets also appear in the third author's study of the *medial rule* for classical logic [27], and as the *skew fibrations* in the second author's *combinatorial proofs* for classical logic [17]. To prepare the ground for cut elimination in first-order combinatorial proofs [19] is further motivation for the present work.

Proof identity

At the heart of a theory of proof nets is the question of *proof identity*: when are two proofs equivalent? The answer determines which proofs should map onto the same proof net. The introduction of quantifiers creates an interesting issue: if two proofs differ by an immaterial choice of existential witness, should they be equivalent? For example, to prove the sequent $\vdash \exists x.P(x), \exists y.\overline{P}(y)$ both quantifiers must receive *the same* witness, as in the following two proofs, but *any* witness will do.

$$\frac{\overline{\vdash P(s), \overline{P}(s)}}{\vdash \exists x.P(x), \exists y.\overline{P}(y)} \stackrel{?}{\equiv} \frac{\overline{\vdash P(t), \overline{P}(t)}}{\vdash \exists x.P(x), \exists y.\overline{P}(y)}$$


The issue is more pronounced where quantifiers are *vacuous*, $\exists x.A$ with x not free in A . The proofs below left can only be distinguished even syntactically because the $\exists R$ rule makes the instantiating witness explicit. Below right is an interesting intermediate variant: the witness s or t can be observed without explicit annotation in the $\exists R$ rule, but the choice is equally immaterial to the content of the proof as when the quantifier were vacuous.

$$\frac{\overline{\vdash P, \overline{P}}}{\vdash \exists x.P, \overline{P}} \stackrel{?}{\equiv} \frac{\overline{\vdash P, \overline{P}}}{\vdash \exists x.P, \overline{P}} \stackrel{\exists R, t}{\equiv} \frac{\overline{\vdash P, \overline{P}}}{\vdash P + Q(s), \overline{P}} \stackrel{+R, 1}{\equiv} \frac{\overline{\vdash P, \overline{P}}}{\vdash \exists x.P + Q(x), \overline{P}} \stackrel{\exists R, s}{\equiv} \frac{\overline{\vdash P, \overline{P}}}{\vdash P + Q(t), \overline{P}} \stackrel{+R, 1}{\equiv} \frac{\overline{\vdash P, \overline{P}}}{\vdash \exists x.P + Q(x), \overline{P}} \stackrel{\exists R, t}{\equiv}$$

In this paper we will not attempt to settle the question of proof identity. Rather, our two notions of proof net each represent a natural and coherent perspective, at either end of the spectrum. *Witness nets* make all existential witnesses explicit, including those to vacuous quantifiers, rejecting all three equivalences above. *Unification nets* leave all witnesses implicit, thus identifying all proofs modulo witness assignment, and validating all three equivalences.


Correctness: coalescence and slicing

Additive proof nets have two natural correctness criteria. **Coalescence** [12, 20], a counterpart to multiplicative *contractibility* [4, 9], provides efficient correctness and sequentialization via local rewriting: it asks that the steps below left result in a single link, connecting both formulas. **Slicing** [21] is a global, geometric criterion: it asks that each **slice**, a choice to remove one subformula of each product along with all connected links, retains a single link.




We extend coalescence to both witness nets and unification nets, and slicing to witness nets. Here, we illustrate the former, and leave a discussion of slicing to the conclusion of the paper.

We will distinguish **strict coalescence** (\rightarrow) for witness nets and **unifying coalescence** (\rightsquigarrow) for unification nets. We first give an example of the former (abbreviating $P(x, y)$ to Pxy). In the initial witness net, below left, each link corresponds to a sequent calculus axiom between both linked subformulas, with the substitutions applied.



The first two steps (above middle and right) move both links from the subformula $\overline{P}xy$ to $\exists y. \overline{P}xy$, removing the substitutions $[s/y]$ and $[t/y]$ on y , and corresponding to sequent rules $\exists R, s$ and $\exists R, t$. Observe that we maintain the domain of the substitutions on a link as the variables of those existential quantifiers that either linked subformula is (strictly) in scope of.

The next step, from above right to below left, combines both links, and corresponds to an additive conjunction rule. We require that both substitutions agree (their domains are the same by the above observation); hence this step could not have been performed before the previous two, corresponding to the non-permutability of the generated inference rules.



The final steps introduce an inference $\exists R, x$ and one $\forall R$. For the latter, we require that the *eigenvariable* x of the universal quantification does not occur in the range of the substitution of the link, as in $[x/z]$ in the net above left – hence the two steps could not be interchanged. It corresponds to the *eigenvariable condition* on the $\forall R$ rule that x is not free in the context.

Unifying coalescence, for unification nets, is similar to strict coalescence; two differences allow it to reconstruct witnesses by unification, which we illustrate. To initialize coalescence, the links in the net below left are given, as a substitution map, the **most general unifier** of the two propositions connected by the link. Both links now correspond to sequent axioms.

$$\begin{array}{c} \overline{P}xx \\ \swarrow \quad \searrow \\ Pys \times Ptz \end{array} \rightsquigarrow \begin{array}{c} \overline{P}xx \\ \swarrow [s/y, s/x] \quad \searrow [t/z, t/x] \\ Pys \times Ptz \end{array} \rightsquigarrow \begin{array}{c} \overline{P}xx \\ \searrow [u/z, u/y, u/x] \\ Pys \times Ptz \end{array}$$

$$\begin{array}{c} \overline{\vdash Pss, Pss} \quad \overline{\vdash Ptt, Ptt} \\ \hline \overline{\vdash Puu, Puu} \quad \overline{\vdash Puu, Puu} \\ \hline \overline{\vdash Puu, Puu \times Puu} \times R \end{array}$$

The second difference is the coalescence step for additive conjunction, above right. Where strict coalescence requires both links to carry identical substitution maps, here we require both maps to be *unifiable*, in the sense that both must have a common, more special (less general) substitution map. This is then applied to the new link. In the example, the terms s and t are unified to u , i.e. u is a most general term that specializes both s and t . Note that in the sequentialization, both subproofs also need to be specialized, from s and t to u .

2 Proof nets for first-order additive linear logic

First-order terms and the formulas of first-order ALL are generated by the following grammars.

$$\begin{aligned} t &::= x \mid f(t_1, \dots, t_n) \\ a &::= P(t_1, \dots, t_n) \mid \overline{P}(t_1, \dots, t_n) \\ A &::= a \mid A + A \mid A \times A \mid \exists x.A \mid \forall x.A \end{aligned}$$

Negation ($\overline{}$) is applied to predicate symbols, \overline{P} as a matter of convenience. The **dual** \overline{A} of an arbitrary formula A is given by De Morgan. We use the following notational conventions:

x, y, z	$\in \text{VAR}$	first-order variables
f, g, h	$\in \Sigma_f$	n -ary ($n \geq 0$) function symbols from a fixed alphabet Σ_f
P, Q, R	$\in \Sigma_p$	n -ary ($n \geq 0$) predicate symbols from a fixed alphabet Σ_p
s, t, u	$\in \text{TERM}$	first-order terms over VAR and Σ_f
a, b, c	$\in \text{ATOM}$	atomic propositions
A, B, C	$\in \text{FORM}$	ALL1 formulas

A **sequent** $\vdash A, B$ is a pair of formulas A and B . A sequent calculus for ALL1 is given in Figure 1, where each rule has a symmetric counterpart for the first formula in the sequent. We write $\pi \vdash A, B$ for a proof π with conclusion sequent $\vdash A, B$. Two proofs are **equivalent** $\pi \sim \pi'$ if one is obtained from the other by rule permutations, given in Figure 2.

By a **subformula** we will mean a subformula **occurrence**. For instance, a formula $A \times A$ has two subformulas A , one on the left and one on the right. The **subformulas** $\text{SUB}(A)$ of a formula are defined as follows; we write $B \leq A$ if B is a subformula of A , i.e. if $B \in \text{SUB}(A)$.

$$\text{SUB}(A) = \{A\} \cup \begin{cases} \text{SUB}(B) \uplus \text{SUB}(C) & \text{if } A = B + C \text{ or } A = B \times C \\ \text{SUB}(B) & \text{if } A = \exists x.B \text{ or } A = \forall x.B \end{cases}$$

22:6 Proof Nets for First-Order Additive Linear Logic

$$\frac{}{\vdash a, \bar{a}}^{\text{ax}} \quad \frac{\vdash A, B_i}{\vdash A, B_1 + B_2}^{+\text{R},i} \quad \frac{\vdash A, B \quad \vdash A, C}{\vdash A, B \times C}^{\times\text{R}} \quad \frac{\vdash A, B[t/x]}{\vdash A, \exists x.B}^{\exists\text{R},t} \quad \frac{\vdash A, B}{\vdash A, \forall x.B}^{\forall\text{R}} \quad (x \notin \text{FV}(A))$$

■ **Figure 1** A sequent calculus for ALL1.

$$\begin{array}{cccc} \frac{\vdash A, B}{\vdash A, \forall y.B}^{\forall\text{R}} & \frac{\vdash A, B[t/y]}{\vdash A, \exists y.B}^{\exists\text{R},t} & \frac{\vdash A, B_i}{\vdash A, B_1 + B_2}^{+\text{R},i} & \frac{\vdash A, B \quad \vdash A, C}{\vdash A, B \times C}^{\times\text{R}} \\ \frac{\vdash \forall x.A, \forall y.B}{\vdash \forall x.A, \forall y.B}^{\forall\text{R}} & \frac{\vdash \forall x.A, \exists y.B}{\vdash \forall x.A, \exists y.B}^{\forall\text{R}} & \frac{\vdash \forall x.A, B_1 + B_2}{\vdash \forall x.A, B_1 + B_2}^{\forall\text{R}} & \frac{\vdash \forall x.A, B \times C}{\vdash \forall x.A, B \times C}^{\forall\text{R}} \\ \sim & \sim & \sim & \sim \\ \frac{\vdash A, B}{\vdash \forall x.A, B}^{\forall\text{R}} & \frac{\vdash A, B[t/y]}{\vdash \forall x.A, B[t/y]}^{\forall\text{R}} & \frac{\vdash A, B_i}{\vdash \forall x.A, B_i}^{\forall\text{R}} & \frac{\vdash A, B \quad \vdash A, C}{\vdash \forall x.A, B \times C}^{\forall\text{R}} \\ \frac{\vdash \forall x.A, B}{\vdash \forall x.A, \forall y.B}^{\forall\text{R}} & \frac{\vdash \forall x.A, B[t/y]}{\vdash \forall x.A, \exists y.B}^{\exists\text{R},t} & \frac{\vdash \forall x.A, B_i}{\vdash \forall x.A, B_1 + B_2}^{+\text{R},i} & \frac{\vdash \forall x.A, B \times C}{\vdash \forall x.A, B \times C}^{\times\text{R}} \\ \\ \frac{\vdash A[s/x], B[t/y]}{\vdash A[s/x], \exists y.B}^{\exists\text{R},t} & \frac{\vdash A[t/x], B_i}{\vdash A[t/x], B_1 + B_2}^{+\text{R},i} & \frac{\vdash A[t/x], B \quad \vdash A[t/x], C}{\vdash A[t/x], B \times C}^{\times\text{R}} \\ \frac{\vdash A[s/x], \exists y.B}{\vdash \exists x.A, \exists y.B}^{\exists\text{R},s} & \frac{\vdash \exists x.A, B_1 + B_2}{\vdash \exists x.A, B_1 + B_2}^{\exists\text{R},t} & \frac{\vdash \exists x.A, B \times C}{\vdash \exists x.A, B \times C}^{\exists\text{R},t} \\ \sim & \sim & \sim \\ \frac{\vdash A[s/x], B[t/y]}{\vdash \exists x.A, B[t/y]}^{\exists\text{R},t} & \frac{\vdash A[t/x], B_i}{\vdash \exists x.A, B_i}^{\exists\text{R},t} & \frac{\vdash A[t/x], B \quad \vdash A[t/x], C}{\vdash \exists x.A, B \times C}^{\times\text{R}} \\ \frac{\vdash \exists x.A, B[t/y]}{\vdash \exists x.A, \exists y.B}^{\exists\text{R},s} & \frac{\vdash \exists x.A, B_1 + B_2}{\vdash \exists x.A, B_1 + B_2}^{+\text{R},i} & \frac{\vdash \exists x.A, B \times C}{\vdash \exists x.A, B \times C}^{\times\text{R}} \\ \\ \frac{\vdash A_i, B_j}{\vdash A_i, B_1 + B_2}^{+\text{R},j} & \frac{\vdash A_i, B \quad \vdash A_i, C}{\vdash A_i, B \times C}^{\times\text{R}} \\ \frac{\vdash A_1 + A_2, B_1 + B_2}{\vdash A_1 + A_2, B_1 + B_2}^{+\text{R},i} & \frac{\vdash A_1 + A_2, B \times C}{\vdash A_1 + A_2, B \times C}^{+\text{R},i} \\ \sim & \sim \\ \frac{\vdash A_i, B_j}{\vdash A_1 + A_2, B_j}^{+\text{R},i} & \frac{\vdash A_i, B \quad \vdash A_i, C}{\vdash A_1 + A_2, B \times C}^{\times\text{R}} \\ \frac{\vdash A_1 + A_2, B_1 + B_2}{\vdash A_1 + A_2, B_1 + B_2}^{+\text{R},j} & \frac{\vdash A_1 + A_2, B \quad \vdash A_1 + A_2, C}{\vdash A_1 + A_2, B \times C}^{\times\text{R}} \\ \\ \frac{\vdash A, C \quad \vdash A, D}{\vdash A, C \times D}^{\times\text{R}} & \frac{\vdash B, C \quad \vdash B, D}{\vdash B, C \times D}^{\times\text{R}} \\ \frac{\vdash A \times B, C \times D}{\vdash A \times B, C \times D}^{\times\text{R}} \\ \sim \\ \frac{\vdash A, C \quad \vdash B, C}{\vdash A \times B, C}^{\times\text{R}} & \frac{\vdash A, D \quad \vdash B, D}{\vdash A \times B, D}^{\times\text{R}} \\ \frac{\vdash A \times B, C \times D}{\vdash A \times B, C \times D}^{\times\text{R}} \end{array}$$

■ **Figure 2** Cut-free rule permutations.

Since we will be working with a graphical representation, we will adopt **Barendregt's convention**, that bound variable names are globally unique identifiers, in the following form. In a sequent $\vdash A, B$ we assume all quantifiers have a unique binding variable, distinct from any free variable. In a proof π over $\vdash A, B$, a variable x that is universally quantified as $\forall x.C$ in $\vdash A, B$ is an **eigenvariable**. A $\forall R$ rule on $\forall x.C$ is considered to bind all free occurrences of the eigenvariable x in its direct subproof. Accordingly, we assume that x does not occur free outside these subproofs (which can be guaranteed by globally renaming x). We take variable names to be persistent throughout a proof, in the sense that we don't admit alpha-conversion (renaming of bound variables) between proof rules. We thus have unique bound variable names in $\vdash A, B$, but in π all $\forall R$ rules on the subformula $\forall x.C$ share the same eigenvariable x .

A **link** (C, D) on a sequent $\vdash A, B$ is a pair of subformulas $C \leq A$ and $D \leq B$. A **linking** λ on the sequent $\vdash A, B$ is a set of links on $\vdash A, B$.

► **Definition 1.** A *pre-net* $\lambda \triangleright A, B$ is a sequent $\vdash A, B$ with a linking λ on it.

Witness maps

We will record the witnessing terms to existential quantifiers as (explicit) substitutions at each link. A **witness map** $\sigma: \text{VAR} \rightarrow \text{TERM}$ is a substitution map which assigns terms to variables, given as a (finite) partial function $\sigma = [t_1/x_1, \dots, t_n/x_n]$. Its **domain** $\text{DOM}(\sigma)$ is $\{x_1, \dots, x_n\}$. We abbreviate by $y \in \sigma$ that a variable y occurs free in the range of σ ($y \in \text{FV}(t_i)$ for some $i \leq n$). The map $\sigma - x$ is undefined on x and as σ otherwise; $\sigma|_V$ is the restriction of σ to a set of variables V , and \emptyset is the empty witness map. We write $A\sigma$ for the application of the substitutions in σ to the formula A , and $\sigma\tau$ is the **composition** of two maps, where $A(\sigma\tau) = (A\sigma)\tau$. We apply σ to a proof π , written $\pi\sigma$, by applying it to each formula in the proof and to each existential witness t recorded with a rule $\exists R, t$.

A **witness linking** λ_Σ is a linking λ with a **witness labelling** $\Sigma: \lambda \rightarrow \text{VAR} \rightarrow \text{TERM}$ that assigns each link (C, D) a witness map. We may use and define λ_Σ as a set of **witness links** $(C, D)_\sigma$ where $(C, D) \in \lambda$ and $\Sigma(C, D) = \sigma$. A witness link $(a, b)_\sigma$ on atomic formulas is an **axiom link** if $\bar{a}\sigma = b\sigma$. An **axiom** witness linking is one comprising axiom links.

► **Definition 2.** A *witness pre-net* $\lambda_\Sigma \triangleright A, B$ is a sequent $\vdash A, B$ with a witness linking λ_Σ .

► **Definition 3.** The *de-sequentialization* $[\pi]$ of a sequent proof $\pi \vdash A, B$ is the witness pre-net $[\pi]_{\emptyset}^{A, B} \triangleright A, B$ where the function $[-]_{\sigma}^{A, B}$ is defined inductively in Figure 3 (symmetric cases are omitted).

A function call $[\pi \vdash A, B]_{\sigma}^{A', B'}$ expects that $A = A'\sigma$ and $B = B'\sigma$: the translation separates a sequent $\vdash A, B$ into subformulas A', B' of the ultimate conclusion of the proof, and the accumulated existential witnesses σ . For an example, we refer to the introduction.

Correctness and sequentialization by coalescence

For sequentialization, the links in a pre-net will be labelled with a sequent proof. An axiom link will carry an axiom, and each coalescence step will introduce one proof rule. Formalizing this, a **proof linking** λ_Π is a witness linking λ_Σ with a **proof labelling** $\Pi: \lambda \rightarrow \text{PROOF}$ assigning a sequent proof to each link. We will use and define λ_Π as a set of **proof links** $(C, D)_\sigma^\pi$, where we require that $\pi \vdash C\sigma, D\sigma$, i.e. that π proves the conclusion $\vdash C\sigma, D\sigma$. A **labelled pre-net** $\lambda_\Pi \triangleright A, B$ is a witness pre-net $\lambda_\Sigma \triangleright A, B$ with a proof labelling Π on λ_Σ .

$$\begin{aligned}
 \left[\frac{}{\vdash a, \bar{a}}^{\text{ax}} \right]_{\sigma}^{b, c} &= \{(b, c)_{\sigma}\} \\
 \left[\frac{\pi}{\vdash A, B_i} \right]_{\sigma}^{A', B'_1 + B'_2} &= \left[\pi \right]_{\sigma}^{A', B'_i} \\
 \left[\frac{\pi \quad \pi'}{\vdash A, B \times C} \right]_{\sigma}^{A', B' \times C'} &= \left[\pi \right]_{\sigma}^{A', B'} \cup \left[\pi' \right]_{\sigma}^{A', C'} \\
 \left[\frac{\pi}{\vdash A, B[t/x]} \right]_{\sigma}^{A', \exists x. B'} &= \left[\pi \right]_{\sigma[t/x]}^{A', B'} \\
 \left[\frac{\pi}{\vdash A, \forall x. B} \right]_{\sigma}^{A', \forall x. B'} &= \left[\pi \right]_{\sigma}^{A', B'}
 \end{aligned}$$

■ **Figure 3** De-sequentialization.

If λ_{Σ} is an axiom linking, we assign an **initial proof labelling** λ_{Σ}^* as follows.

$$\lambda_{\Sigma}^* = \{ (a, b)_{\sigma}^{\pi} \mid (a, b)_{\sigma} \in \lambda_{\Sigma}, \pi = \overline{\vdash a\sigma, b\sigma} \}$$

For correctness we may coalesce a pre-net directly, without constructing a proof. So far we have accumulated the following notational conventions:

π, ϕ, ψ	\in	PROOF	ALL1 sequent proofs
κ, λ	\subset	FORM \times FORM	linkings (sets of pairs of formulas)
ρ, σ, τ	:	VAR \rightarrow TERM	witness maps
Σ, Θ	:	$\lambda \rightarrow$ VAR \rightarrow TERM	witness labellings on a linking λ
Π, Φ, Ψ	:	$\lambda \rightarrow$ PROOF	proof labellings on a linking λ


► **Definition 4. Strict sequentialization** (\rightarrow) is the rewrite relation on labelled pre-nets generated by the rules in Figure 4, which replace one or two links by another in a pre-net $\lambda_{\Sigma}^{\Pi} \triangleright A, B$ where B has a subformula $D_1 + D_2$, $D_1 \times D_2$, $\exists x.D$, and $\forall x.D$ respectively. Symmetric cases are omitted. **Strict coalescence** is the same relation on witness pre-nets, ignoring proof labels, illustrated in Figure 5. A witness pre-net $\lambda_{\Sigma} \triangleright A, B$ **strict-coalesces** if it reduces to $\{(A, B)_{\emptyset}\} \triangleright A, B$. It **strongly** strict-coalesces if any coalescence path terminates at $\{(A, B)_{\emptyset}\} \triangleright A, B$.

For an example of coalescence, see the introduction.

► **Definition 5. An ALL1 witness proof net or witness net** is a witness pre-net $\lambda_{\Sigma} \triangleright A, B$ with λ_{Σ} an axiom linking, that strict-coalesces. It **sequentializes** to a proof π if its initial labelling $\lambda_{\Sigma}^* \triangleright A, B$ reduces in (\rightarrow) to $\{(A, B)_{\emptyset}^{\pi}\} \triangleright A, B$.

$$\begin{array}{l}
(C, D_i)_\sigma^\pi \rightarrow (C, D_1 + D_2)_\sigma^\psi \qquad \psi = \frac{\pi}{\frac{\vdash C\sigma, D_i\sigma}{\vdash C\sigma, D_1\sigma + D_2\sigma} +R, i} \quad (+S, i) \\
\left. \begin{array}{l} (C, D_1)_\sigma^\pi \\ (C, D_2)_\sigma^\phi \end{array} \right\} \rightarrow (C, D_1 \times D_2)_\sigma^\psi \qquad \psi = \frac{\pi \quad \phi}{\frac{\vdash C\sigma, D_1\sigma \quad \vdash C\sigma, D_2\sigma}{\vdash C\sigma, D_1\sigma \times D_2\sigma} \times R} \quad (\times S) \\
(C, D)_\sigma^\pi \rightarrow (C, \exists x.D)_{\sigma-x}^\psi \quad (x \in \text{DOM}(\sigma)) \qquad \psi = \frac{\pi}{\frac{\vdash C\sigma, D\sigma}{\vdash C(\sigma-x), \exists x.D(\sigma-x)} \exists R, \sigma(x)} \quad (\exists S) \\
(C, D)_\sigma^\pi \rightarrow (C, \forall x.D)_\sigma^\psi \quad (x \notin \sigma) \qquad \psi = \frac{\pi}{\frac{\vdash C\sigma, D\sigma}{\vdash C\sigma, \forall x.D\sigma} \forall R} \quad (\forall S)
\end{array}$$

■ **Figure 4** Strict coalescence.



■ **Figure 5** Coalescence rules.

We conclude this section by establishing that sequentialization and de-sequentialization for witness nets are inverses, and that witness nets are canonical.

- ▶ **Theorem 6.** *For any ALL1 proof π , the witness net $[\pi]$ sequentializes to π .*
- ▶ **Theorem 7.** *If $\lambda_\Sigma \triangleright A, B$ sequentializes to π , then $[\pi]$ is $\lambda_\Sigma \triangleright A, B$.*
- ▶ **Theorem 8.** *Witness nets are canonical: $[\pi] = [\phi]$ if and only if $\pi \sim \phi$.*

3 Geometric correctness

We will first identify two aspects of sequent proofs, arising from the local nature of the rules, which need to be enforced explicitly in a geometric correctness condition.

Local eigenvariables. The side-condition on the $\forall R$ -rule, that the eigenvariable is not free in the context, means that eigenvariables are *local* to the subproof of the $\forall R$ -rule. Correspondingly, a link $(C, D)_\sigma$ on $\vdash A, B$ has **local eigenvariables** if for any variable $x \in \sigma$, if x is an eigenvariable quantified as $\forall x.X$ in $\vdash A, B$, then $C \leq X$ or $D \leq X$. A witness linking or pre-net has **local eigenvariables** if all its links do.

22:10 Proof Nets for First-Order Additive Linear Logic

Exact coverage. The local witness substitution $[t/x]$ in a rule instance $\exists R, t$ will have been applied exactly to the axioms $\vdash a, \bar{a}$ in the subproof of that rule. Correspondingly, for a link $(C, D)_\sigma$ on $\vdash A, B$ we expect the domain of σ to be exactly the existential variables in A and B that (could) occur free in C and D . For a subformula C of A , let the **free existential variables** of C in A be the set $\text{EV}_A(C) = \{x \mid C < \exists x.X \leq A\}$. A link $(C, D)_\sigma$ on $\vdash A, B$ then has **exact coverage** if $\text{DOM}(\sigma) = \text{EV}_A(C) \cup \text{EV}_B(D)$. If it does, σ consists of two components, $\sigma|_{\text{EV}_A(C)}$ and $\sigma|_{\text{EV}_B(D)}$, which we abbreviate as σ_C and σ_D respectively. A witness linking or pre-net has **exact coverage** if all its links do.

Both conditions are captured naturally by coalescence, as can be observed from the rules.

Slices

A **slice** is the fraction of a proof that depends on a given choice of one branch (or projection) on each product formula $A \times B$. Important to additive proof theory is that many operations can be performed on a per-slice basis, such as normalization, or proof net correctness. We will here use slices for the latter purpose. As in the propositional case [21], we define a **slice** of a sequent $\vdash A, B$ as a set of potential links, of which exactly one must be realized in a proof net $\lambda_\Sigma \triangleright A, B$. We extend the propositional criterion in two ways:

Expansion. When defining slices, we interpret an existential quantification $\exists x.A$ as a sum over all witnesses t_i to x that occur in the pre-net, $A[t_1/x] + \dots + A[t_n/x]$. This captures the non-permutability of a product rule over distinct instantiations, as below left. (Technically, a slice of $\exists x.A$ will correspond to an infinite sum over $A[t/x]$ for every term t , but only the actually occurring terms t_i will ever be relevant.)

Dependency. We define a **dependency** relation between a universal quantification $\forall x.A$ and an instantiation $B[t/y]$ of $\exists y.B$ where the eigenvariable x occurs free in t (a standard approach to first-order quantification [25, 8, 10, 26]). For each link, we will require this dependency relation to be *acyclic*, which amounts to *slice-wise* first-order correctness. It captures the non-permutability of universal and existential sequent rules due to the *eigenvariable condition* of the former, as below right.

$$\frac{\frac{\vdash A[s/x], B}{\vdash \exists x.A, B} \exists R, s \quad \frac{\vdash A[t/x], C}{\vdash \exists x.A, C} \exists R, t}{\vdash \exists x.A, B \times C} \times R \qquad \frac{\frac{\vdash A, B[t/y]}{\vdash A, \exists y.B} \exists R, t}{\vdash \forall x.A, \exists y.B} \forall R \quad \text{where } x \in \text{FV}(t)$$

Expansion covers the interaction between existential quantifiers and products, and *dependency* that between existential and universal quantifiers. Note that in proof nets that include the multiplicatives (e.g. [8, 18]), instead of a dependency as a *partial order* it is common to consider the underlying graph, created by adding a *jump* or *leap* edge when an existential instantiation $B[t/y]$ of $\exists y.B$ contains the eigenvariable x of $\forall x.A$. These edges participate in the *switching* condition [5], to capture the interaction between quantifiers and tensors. Without the tensor, here, the dependency as a partial order suffices.

► **Definition 9 (Slice).** A **slice** S of a formula A and a witness map σ is a set of pairs (A', σ') , where $A' \leq A$ and $\sigma' \supseteq \sigma$, given by $S = \{(A, \sigma)\} \cup S'$ where:

- If $A = a$ then $S' = \emptyset$.
- If $A = B + C$ then $S' = S_B \uplus S_C$ with S_B a slice of B and σ , and S_C one of C and σ .
- If $A = B \times C$ then S' is a slice of B and σ or a slice of C and σ .
- If $A = \exists x.B$ then $S' = \uplus_{t \in \text{TERM}} S_t$ where each S_t is a slice of B and $\sigma[t/x]$.
- If $A = \forall x.B$ then S' is a slice of B and σ .

A *slice* of a sequent $\vdash A, B$ is a set of links

$$\{ (C, D)_{\sigma \cup \tau} \mid (C, \sigma) \in S_A, (D, \tau) \in S_B \}$$

where S_A is a slice of A and \emptyset , and S_B a slice of B and \emptyset . A *slice* of a witness pre-net $\lambda_\Sigma \triangleright A, B$ is the intersection $\lambda_\Sigma \cap S$ of λ_Σ with a slice S of $\vdash A, B$.

As in the propositional case, for correctness we will require that each slice is a singleton. We will further define a **dependency** condition to ensure that the order in which quantifiers are instantiated is sound, corresponding to the *eigenvariable condition* on the \forall R-rule of sequent calculus. For simplicity, we define the condition on individual links rather than slices.

► **Definition 10.** In a pre-net $\lambda_\Sigma \triangleright A, B$, let the **column** of a link $(C, D)_\sigma$ be the set of pairs

$$\{ (X, \sigma|_{\text{EV}_A(X)}) \mid C \leq X \leq A \} \cup \{ (Y, \sigma|_{\text{EV}_B(Y)}) \mid D \leq Y \leq B \},$$

with a **dependency** relation (\preceq): $(X, \rho) \preceq (Y, \tau)$ if $X \leq Y$ or Y occurs as $\forall x.Y$ and $x \in \rho$.

► **Definition 11.** A witness pre-net is **correct** if:

- it has local eigenvariables and exact coverage,
- it is **slice-correct**: every slice is a singleton, and
- it is **dependency-correct**: every column is a partial order (i.e. is acyclic/antisymmetric).

To conclude this section we will establish that both correctness conditions, by coalescence and by slicing, are equivalent. Moreover, both are equivalent to *strong* coalescence.

► **Theorem 12.** A witness pre-net that strict-coalesces is correct, and a correct witness pre-net strongly strict-coalesces.

► **Corollary 13.** A correct witness pre-net with axiom linking is a witness proof net.

4 Composition

We will describe the composition of two witness nets by a global operation. It consists of the relational composition of both linkings, as in the propositional case, where for each pair of links that are being connected, their witness maps are composed. As links correspond to slices, the operation is effectively first-order composition [26] applied slice-wise.

Cut-elimination rules for ALL1 are given in Figure 6; the requisite permutations are in Figure 7.

$$\begin{array}{ccc}
 \exists v. \overline{P} & & \exists v. \overline{P} \\
 \begin{array}{|c} \hline [f(x)/y, z/v] \\ \hline \end{array} & & \\
 \forall x. \exists y. \forall z. P & \Rightarrow & [g(f(t))/v] \\
 \exists x. \forall y. \exists z. \overline{P} & & \\
 \begin{array}{|c} \hline [t/x, g(y)/z] \\ \hline \end{array} & & \\
 P & & P
 \end{array}$$

We use the example above to illustrate the composition of links. To eliminate the central cut, on $\forall x. \exists y. \forall z. P$ and $\exists x. \forall y. \exists z. \overline{P}$, the explicit substitutions for both formulas must be effectuated. An inductive procedure, as in sequent calculus, could apply them from outside in: first $[t/x]$, then $[f(t)/y]$ (previously $[f(x)/y]$), then $[g(f(t))/z]$ (previously $[g(y)/z]$).

For a direct definition, to compose two links $(a, b)_\sigma$ and $(\overline{b}, c)_\tau$, the substitutions into the cut-formula σ_b and $\tau_{\overline{b}}$ must be applied as often as needed, up to the depth of quantifiers

22:12 Proof Nets for First-Order Additive Linear Logic

$$\begin{array}{c}
 \frac{\frac{\pi_1}{\vdash A, B_1} \quad \frac{\pi_2}{\vdash A, B_2}}{\vdash A, B_1 \times B_2} \times R \quad \frac{\frac{\phi}{\vdash \overline{B}_i, C}}{\vdash \overline{B}_1 + \overline{B}_2, C} +R, i}{\vdash A, C} \text{cut} \Rightarrow \frac{\frac{\pi_i}{\vdash A, B_i} \quad \frac{\phi}{\vdash \overline{B}_i, C}}{\vdash A, C} \text{cut} \\
 \\
 \frac{\frac{\pi}{\vdash A, B[t/x]} \exists R, t \quad \frac{\frac{\phi}{\vdash \overline{B}, C}}{\vdash \forall x. \overline{B}, C} \forall R}{\vdash A, C} \text{cut} \Rightarrow \frac{\frac{\pi}{\vdash A, B[t/x]} \quad \frac{\phi[t/x]}{\vdash \overline{B}[t/x], C}}{\vdash A, C} \text{cut}
 \end{array}$$

■ **Figure 6** ALL1 cut-elimination steps.

$$\begin{array}{c}
 \frac{\frac{\vdash A, B}{\vdash \forall x. A, B} \forall R \quad \frac{}{\vdash \overline{B}, C} \text{cut}}{\vdash \forall x. A, C} \text{cut} \quad \frac{\frac{\vdash A[t/x], B}{\vdash \exists x. A, B} \exists R, t \quad \frac{}{\vdash \overline{B}, C} \text{cut}}{\vdash \exists x. A, C} \text{cut} \\
 \sim \\
 \frac{\frac{\vdash A, B \quad \vdash \overline{B}, C}{\vdash A, C} \text{cut}}{\vdash \forall x. A, C} \forall R \quad \frac{\frac{\vdash A[t/x], B \quad \vdash \overline{B}, C}{\vdash A[t/x], C} \text{cut}}{\vdash \exists x. A, C} \exists R, t \quad \frac{\frac{\vdash A, B \quad \vdash \overline{B}, C}{\vdash A, C} \text{cut} \quad \frac{}{\vdash \overline{C}, D} \text{cut}}{\vdash A, D} \text{cut} \\
 \sim \\
 \frac{\frac{\vdash A_i, B}{\vdash A_1 + A_2, B} +R, i \quad \frac{}{\vdash \overline{B}, C} \text{cut}}{\vdash A_1 + A_2, C} \text{cut} \quad \frac{\frac{\vdash A_1, B \quad \vdash A_2, B}{\vdash A_1 \times A_2, B} \times R \quad \frac{}{\vdash \overline{B}, C} \text{cut}}{\vdash A_1 \times A_2, C} \text{cut} \quad \frac{\frac{\vdash \overline{B}, C \quad \vdash \overline{C}, D}{\vdash \overline{B}, D} \text{cut}}{\vdash A, D} \text{cut} \\
 \sim \\
 \frac{\frac{\vdash A_i, B \quad \vdash \overline{B}, C}{\vdash A_i, C} \text{cut}}{\vdash A_1 + A_2, C} +R, i \quad \frac{\frac{\vdash A_1, B \quad \vdash \overline{B}, C}{\vdash A_1, C} \text{cut} \quad \frac{\vdash A_2, B \quad \vdash \overline{B}, C}{\vdash A_2, C} \text{cut}}{\vdash A_1 \times A_2, C} \times R
 \end{array}$$

■ **Figure 7** Cut-permutations.

above b , to the terms in the range of the remaining substitutions, σ_a and τ_c . To formalize this, we will use the following notions:

- The **domain-preserving composition** of two witness maps $\sigma \cdot \tau$ is the map $(\sigma\tau)|_{\text{DOM}(\sigma)}$.
- The **least fixed point** $\vec{\sigma}$ of a witness map σ is the least map ρ satisfying $\rho = \rho\sigma$.

The latter is the shortest sequence $\vec{\sigma} = \sigma\sigma\dots\sigma$ such that no variable is both in the domain and range of $\vec{\sigma}$. This is not necessarily finite; in our composition operations, finiteness is ensured by the correctness conditions on proof nets (see Theorem 15).

► **Definition 14.** The **composition** $(A, B)_\sigma^\pi; (\overline{B}, C)_\tau^\phi$ of two proof links is $(A, C)_\rho^\psi$ where

$$\rho = \sigma_A \tau_C \cdot \overrightarrow{\sigma_B \tau_B} \quad \text{and} \quad \psi = \frac{\left(\frac{\pi}{\vdash A\sigma, B\sigma} \right) \overrightarrow{\sigma_B \tau_B} \quad \left(\frac{\phi}{\vdash \overline{B}\tau, C\tau} \right) \overrightarrow{\sigma_B \tau_B}}{\vdash A\rho, C\rho} \text{cut} .$$

The **composition** $\lambda_{\Sigma}^{\Pi}; \kappa_{\Theta}^{\Phi}$ of two linkings is the linking

$$\{ (X, Y)_{\sigma}^{\pi}; (\bar{Y}, Z)_{\tau}^{\phi} \mid (X, Y)_{\sigma}^{\pi} \in \lambda_{\Sigma}^{\Pi}, (\bar{Y}, Z)_{\tau}^{\phi} \in \kappa_{\Theta}^{\Phi} \}$$

The **composition** $(\lambda_{\Sigma}^{\Pi} \triangleright A, B); (\kappa_{\Theta}^{\Phi} \triangleright \bar{B}, C)$ of two pre-nets is the pre-net $(\lambda_{\Sigma}^{\Pi}; \kappa_{\Theta}^{\Phi}) \triangleright A, C$. These compositions may omit proof annotations and witness annotations.

The composition of two links is strongly related to composition of strategies in game semantics. There, two strategies on $\vdash A, B$ and $\vdash \bar{B}, C$ are composed by *interaction* on the interface of B and \bar{B} , and subsequently *hiding* that interaction.

In the following we will demonstrate that composition gives the desired result: if a net L sequentializes to π and R to ϕ , then $L; R$ sequentializes to a normal form of the composition of π and ϕ with a cut. To this end we will explore how composition and sequentialization interact. We will consider the critical pairs of sequentialization (\rightarrow) with composition (\Rightarrow) given in Figures 8–10, and demonstrate how they are resolved.

- $\vdash A, B_1 \times B_2; \vdash \bar{B}_1 + \bar{B}_2, C$ (Figure 8)
Since the free existential variables of B and B_1 are the same, $\sigma_B \tau_{\bar{B}} = \sigma_{B_1} \tau_{\bar{B}_1}$ and $\rho = \rho'$. It then follows that ψ' cut-eliminates in one step to ψ .
- $\vdash A, \exists x. B; \vdash \forall x. \bar{B}, C$ (Figure 9)
Since x is not free in the range of τ , nor in the range of σ (by Barendregt's convention), we have that $\overline{\sigma_B \tau_{\bar{B}}}$ is $\overline{(\sigma_B - x) \tau_{\bar{B}}}$ plus the substitution $[\sigma(x)/x]$. Then $\rho = \rho'$ (as x does not occur in the range of $\sigma_A \tau_C$) and ψ' reduces to ψ in a single cut-elimination step.
- $\vdash A, B; \vdash \bar{B}, \exists x. C$ (Figure 10)
Observe that since x occurs in C but not B , it is not in the domain of τ_B , so that $\tau_B - x$ is just τ_B . Then $\rho' = \rho - x$, and the diagram is closed by a sequentialization step (from left to right) which extends ψ with an existential introduction rule, to a proof equivalent to ψ' :

$$\frac{\psi \quad \vdash A\rho, C\rho}{\vdash A\rho', \exists x. C\rho'} \exists R, \rho(x)$$

There are three further critical pairs, for a proof net on $\vdash A, B$ composed with one on $\vdash \bar{B}, C_1 + C_2$, one on $\vdash \bar{B}, C_1 \times C_2$, and one on $\vdash \bar{B}, \forall x. C$. These converge like the one above. Resolving these critical pairs gives the soundness of the composition operation, per the following theorem. We abbreviate a cut on proofs $\pi \vdash A, B$ and $\phi \vdash \bar{B}, C$ by $\pi; \phi$.

► **Theorem 15.** *If proof nets $\lambda_{\Sigma} \triangleright A, B$ and $\kappa_{\Theta} \triangleright \bar{B}, C$ sequentialize to π and ϕ respectively, then their composition $(\lambda_{\Sigma} \triangleright A, B); (\kappa_{\Theta} \triangleright \bar{B}, C)$ is well-defined (i.e. all fixed points are finite) and sequentializes to a normal form ψ of $\pi; \phi$.*

5 Unification nets

In this final section we explore a second notion of ALL1 proof net: **unification nets** omit any witness information, which is then reconstructed by coalescence. This yields a natural notion of *most general* proof net, where every other proof net is obtained by introducing more witness information. Conversely, every witness net has an underlying unification net, that sequentializes to a *most general* proof.

We consider a proof $\pi \vdash A, B$ **more general** than $\pi' \vdash A, B$, written $\pi \leq \pi'$, if there is a substitution map ρ such that $\pi\rho = \pi'$. Unlike for proof nets, this notion is not so natural for

22:14 Proof Nets for First-Order Additive Linear Logic

$$\begin{array}{ccc}
 \begin{array}{c} A \\ \pi, \sigma \swarrow \searrow \pi', \sigma \\ B_1 \times B_2 \\ \phi, \tau \swarrow \searrow \\ \bar{B}_1 + \bar{B}_2 \\ C \end{array} & \rightarrow & \begin{array}{c} A \\ \pi'', \sigma \\ B_1 \times B_2 \\ \bar{B}_1 + \bar{B}_2 \\ \phi', \tau \\ C \end{array} \\
 \Downarrow & & \Downarrow \\
 \begin{array}{c} A \\ \psi, \rho \\ C \end{array} & & \begin{array}{c} A \\ \psi', \rho' \\ C \end{array}
 \end{array}
 \quad
 \begin{array}{l}
 \rho = \sigma_{ATC} \cdot \overrightarrow{\sigma_B \tau \bar{B}} \\
 \rho' = \sigma_{ATC} \cdot \overrightarrow{\sigma_{B_1} \tau \bar{B}_1} \\
 \psi = \frac{\left(\frac{\pi}{\vdash A \sigma, B_1 \sigma} \right) \overrightarrow{\sigma_{B_1} \tau \bar{B}_1} \quad \left(\frac{\phi}{\vdash \bar{B}_1 \tau, C \tau} \right) \overrightarrow{\sigma_{B_1} \tau \bar{B}_1}}{\vdash A \rho, C \rho} \text{cut} \\
 \psi' = \frac{\left(\frac{\pi \quad \pi'}{\vdash A \sigma, B_1 \sigma \quad \vdash A \sigma, B_2 \sigma} \right) \overrightarrow{\sigma_B \tau \bar{B}} \quad \left(\frac{\phi}{\vdash \bar{B}_1 \tau, C \tau} \right) \overrightarrow{\sigma_{B_1} \tau \bar{B}_1}}{\vdash A \rho', C \rho'} \text{cut}
 \end{array}$$

■ **Figure 8** The critical pair $\vdash A, B_1 \times B_2 ; \vdash \bar{B}_1 + \bar{B}_2, C$.

$$\begin{array}{ccc}
 \begin{array}{c} A \\ \pi, \sigma \swarrow \searrow \\ \exists x. B \\ \forall x. \bar{B} \\ \phi, \tau \swarrow \searrow \\ C \end{array} & \xrightarrow{x \notin \tau} & \begin{array}{c} A \\ \pi', \sigma - x \swarrow \searrow \\ \exists x. B \\ \forall x. \bar{B} \\ \phi', \tau \swarrow \searrow \\ C \end{array} \\
 \Downarrow & & \Downarrow \\
 \begin{array}{c} A \\ \psi, \rho \\ C \end{array} & & \begin{array}{c} A \\ \psi', \rho' \\ C \end{array}
 \end{array}
 \quad
 \begin{array}{l}
 \rho = \sigma_{ATC} \cdot \overrightarrow{\sigma_B \tau \bar{B}} \\
 \rho' = \sigma_{ATC} \cdot \overrightarrow{(\sigma_B - x) \tau \bar{B}} \\
 \psi = \frac{\left(\frac{\pi}{\vdash A \sigma, B \sigma} \right) \overrightarrow{\sigma_B \tau \bar{B}} \quad \left(\frac{\phi}{\vdash \bar{B} \tau, C \tau} \right) \overrightarrow{\sigma_B \tau \bar{B}}}{\vdash A \rho, C \rho} \text{cut} \\
 \psi' = \frac{\left(\frac{\pi}{\vdash A \sigma, B \sigma} \right) \overrightarrow{(\sigma_B - x) \tau \bar{B}} \quad \left(\frac{\phi}{\vdash \bar{B} \tau, C \tau} \right) \overrightarrow{(\sigma_B - x) \tau \bar{B}}}{\vdash A \rho', C \rho'} \text{cut}
 \end{array}$$

■ **Figure 9** The critical pair $\vdash A, \exists x. B ; \vdash \forall x. \bar{B}, C$.

$$\begin{array}{ccc}
 \begin{array}{c} A \\ \pi, \sigma \\ B \\ \bar{B} \\ \phi, \tau \swarrow \searrow \\ \exists x. C \end{array} & \rightarrow & \begin{array}{c} A \\ \pi, \sigma \\ B \\ \bar{B} \\ \phi', \tau - x \swarrow \searrow \\ \exists x. C \end{array} \\
 \Downarrow & & \Downarrow \\
 \begin{array}{c} A \\ \psi, \rho \\ \exists x. C \end{array} & & \begin{array}{c} A \\ \psi', \rho' \\ \exists x. C \end{array}
 \end{array}
 \quad
 \begin{array}{l}
 \rho = \sigma_{ATC} \cdot \overrightarrow{\sigma_B \tau \bar{B}} \\
 \rho' = \sigma_A(\tau C - x) \cdot \overrightarrow{\sigma_B \tau \bar{B}} \\
 \psi = \frac{\left(\frac{\pi}{\vdash A \sigma, B \sigma} \right) \overrightarrow{\sigma_B \tau \bar{B}} \quad \left(\frac{\phi}{\vdash \bar{B} \tau, C \tau} \right) \overrightarrow{\sigma_B \tau \bar{B}}}{\vdash A \rho, C \rho} \text{cut} \\
 \psi' = \frac{\left(\frac{\pi}{\vdash A \sigma, B \sigma} \right) \overrightarrow{\sigma_B \tau \bar{B}} \quad \left(\frac{\phi}{\vdash \bar{B} \tau, C \tau} \right) \overrightarrow{\sigma_B \tau \bar{B}}}{\vdash A \rho', \exists x. C \rho'} \text{cut}
 \end{array}$$

■ **Figure 10** The critical pair $\vdash A, B ; \vdash \bar{B}, \exists x. C$.

sequent proofs: in the permutation of existential and product rules below, from left to right u must be generated as the least term more general than s and t ; from right to left, s and t cannot be reconstructed from u , and must be retrieved from their respective subproofs.

$$\frac{\frac{\frac{\vdash A, C}{\vdash A, \exists x.C} \exists R, s \quad \frac{\vdash B, C}{\vdash B, \exists x.C} \exists R, t}{\vdash A \times B, \exists x.C} \times R}{\vdash A, C \quad \vdash B, C} \times R \sim \frac{\frac{\vdash A, C \quad \vdash B, C}{\vdash A \times B, C} \times R}{\vdash A \times B, \exists x.C} \exists R, u$$

To reconstruct witnesses by unification, we define the following operations.

$\sigma \leq \tau$: A witness map σ is **more general** than τ if there is a map ρ such that $\sigma\rho = \tau$.

$\sigma \frown \tau$: Two witness maps σ and τ are **coherent** if there is a map ρ such that $\sigma\rho = \tau\rho$.

$\sigma \vee \tau$: The **join** of coherent witness maps is the least map ρ such that $\sigma \leq \rho$ and $\tau \leq \rho$.

A link (a, b) on two atomic formulas is an **axiom** link if there exists a witness map σ such that $\bar{a}\sigma = b\sigma$. To an axiom link (a, b) over $\vdash A, B$ we assign an **initial witness map**, which is the least witness map σ over the domain $\text{EV}_A(a) \cup \text{EV}_B(b)$ such that $\bar{a}\sigma = b\sigma$. In other words, σ is the most general unifier of \bar{a} and b , over the given domain, written $\text{MGU}(\bar{a}, b)$. For an axiom linking λ over $\vdash A, B$ the **initial witness pre-net** $\lambda_* \triangleright A, B$ is given by

$$\lambda_* = \{ (a, b)_\sigma \mid (a, b) \in \lambda, \sigma = \text{MGU}(\bar{a}, b) \}.$$

An initial witness pre-net has *exact coverage*, while coalescence will give *local eigenvariables*. Note that eigenvariables are constants for the purpose of unification (they are not substituted into). For Barendregt's convention, that free variables have distinct names from bound ones, we should assume that variables in the range of a witness map are fresh; for example, the most general unifier of existential variables x and y should be $[z/x, z/y]$ for a fresh variable z , and not $[y/x]$ or $[x/y]$.

► **Definition 16.** *Unifying sequentialization* (\rightsquigarrow) is the rewrite relation on labelled pre-nets generated by the rules $(+U, i)$, $(\exists U)$, $(\forall U)$, which are respectively as $(+S, i)$, $(\exists S)$, and $(\forall S)$, and the rule

$$\left. \begin{array}{l} (C, D_1)_\sigma^\pi \\ (C, D_2)_\tau^\phi \end{array} \right\} \rightsquigarrow (C, D_1 \times D_2)_{\sigma \vee \tau = \sigma\rho = \tau\rho}^\psi \quad (\sigma \frown \tau) \quad \psi = \frac{\left(\frac{\pi}{\vdash A\sigma, B\sigma} \right)^\rho \quad \left(\frac{\phi}{\vdash A\tau, C\tau} \right)^\rho}{\vdash A(\sigma \vee \tau), B\sigma \times C\tau\rho} \quad (\times U)$$

Unifying coalescence is the relation (\rightsquigarrow) on witness pre-nets, ignoring proof labels. A witness pre-net $\lambda_\Sigma \triangleright A, B$ **unifying-coalesces** if it reduces to $\{(A, B)_\emptyset\} \triangleright A, B$ and **strongly unifying-coalesces** if any coalescence path terminates at $\{(A, B)_\emptyset\} \triangleright A, B$.

► **Definition 17.** An ALL1 **unification proof net** or **unification net** is a pre-net $\lambda \triangleright A, B$ with axiom linking λ such that the initial witness pre-net $\lambda_* \triangleright A, B$ unifying-coalesces. It **sequentializes** to π if $\lambda_*^* \triangleright A, B$ reduces in (\rightsquigarrow) to $\{(A, B)_\emptyset^\pi\} \triangleright A, B$.

In the above definition, note that $\lambda_*^* = (\lambda_*)^*$ is the initial proof labelling of λ_* , which assigns an axiom rule to each axiom link. For a minimal example, see the Introduction. Observe also that unifying coalescence includes strict coalescence, $(\rightarrow) \subseteq (\rightsquigarrow)$. The following two lemmata relate sequentialization for witness nets and unification nets.

► **Lemma 18.** In (\rightsquigarrow), if $\lambda_\Sigma \triangleright A, B$ sequentializes to π then $\lambda_* \triangleright A, B$ sequentializes to $\pi' \leq \pi$.

► **Lemma 19.** If $\lambda_* \triangleright A, B$ unifying-sequentializes to π then there exists a witness assignment Σ and substitution ρ such that $\lambda_\Sigma \triangleright A, B$ strict-sequentializes to π and $\lambda_\Sigma = \lambda_*\rho$.

	Monomial nets	Witness nets	Unification nets
Canonicity	?	✓	✓
Generality	✗	✗	✓
Direct composition	✗	✓	✓
Coalescence	?	✓	✓
Slicing	✓	✓	?

■ **Figure 11** Comparison of different notions of proof nets for quantifiers and additive connectives.

We can then show that sequentialization and de-sequentialization for unification nets are inverses up to generality, and that composition is sound.

► **Theorem 20.** *If $[\pi \vdash A, B]$ is $\lambda_\Sigma \triangleright A, B$ then $\lambda \triangleright A, B$ unifying-sequentializes to $\pi' \leq \pi$.*

► **Theorem 21.** *If $\lambda \triangleright A, B$ sequentializes to π , then $[\pi] = \lambda_\Sigma \triangleright A, B$ for some Σ .*

► **Theorem 22.** *If $\lambda \triangleright A, B$ sequentializes to π and $\kappa \triangleright \overline{B}, C$ to ϕ then their composition $\lambda; \kappa \triangleright A, C$ sequentializes to a proof $\psi' \leq \psi$ where ψ is a normal form of $\pi; \phi$.*

6 Conclusion and related work

We have presented two notions of first-order additive proof net, *witness nets* and *unification nets*, to capture canonically two natural notions of proof identity for first-order additive linear logic. Figure 11 summarizes our results, along with some observations we make below.

Proof nets with additives and quantifiers existed before as *monomial nets* [8]. These are not generally canonical: they admit the permutation (and duplication) of proof rules past implicit *contractions* (that is, the shared context of the additive conjunction rule). However, it might be possible to restrict additive monomial nets to some notion of canonical form. Likewise, coalescence (or contractibility) has not been studied for first-order monomial nets, though it has been extended to a related form of MALL proof nets [24].

Our slicing condition is loosely related to a number of approaches to first-order *classical* logic. A formula $\exists x.A$ is interpreted as the sum (or classically, the *disjunction*) over a fixed number of instantiations $A[t_1/x] + \dots + A[t_n/x]$. This can be traced to Herbrand's Theorem [14]: $\exists x.A$ is equivalent to the infinite sum over $A[t/x]$ for all terms t in the language, but for any given proof a finite set of terms suffices. *Expansion tree proofs* [25, 10] are a graphical proof formalism based on this idea. In our slicing condition, the interpretation of $\exists x.A$ is an infinite sum of which only a finite part $A[t_1/x] + \dots + A[t_n/x]$ is relevant, over the witnesses t_1, \dots, t_n actually assigned to x in the proof net. An interesting alternative approach to additive proof nets, which we may explore in future work, is to take the expansion of $\exists x.A$ to $A[t_1/x] + \dots + A[t_n/x]$ as primary, and record it explicitly in the syntax, as expansion tree proofs do for classical logic. It is expected, however, that this would sacrifice *generality* and *direct composition*.

Hetzl [15] explores *explicit substitution* for first-order classical sequent calculus.

Lambek observed that the two canonical proofs of $\vdash A+A, \overline{A}$ can be distinguished by casting each as a specialization (by substituting into propositional variables) of the more general proofs of $\vdash a+b, \overline{a}$ and $\vdash a+b, \overline{b}$ (corresponding to the two injections of a sum). He proposed to use this idea of *generality* as the basis for a notion of *proof identity*: two proofs are equivalent if their most general forms are isomorphic [23]. How this extends to first order is not obvious. The natural first-order analogue of $\vdash A+A, \overline{A}$ would be $\vdash \exists x.A, \overline{A}$ where the quantifier is vacuous, as $\exists x.A$ represents the infinite sum over A (for all terms t). Where

Lambek’s generality *distinguishes* the two proofs of $\vdash A+A, \bar{A}$, ours *identifies* the proofs of $\vdash \exists x.A, \bar{A}$: the sequent has one unification net, but infinitely many witness nets (one for each term t). If existential quantification is indeed analogous to a sum, Lambek’s notion of generality is more faithfully captured by witness nets than unification nets.

For future work, there are a few natural questions. First is that of a geometric criterion for unification nets. For this, it seems essential to reconstruct dependencies between products and existential quantifiers globally, as the *expansion* condition provides for witness nets. These interact in highly intricate ways, which the unifying coalescence algorithm resolves incrementally and locally. To do so globally, as would be necessary for a geometric criterion, is a major combinatorial challenge.

A second is whether coalescence (or *contractibility* [4]) applies to MLL1 unification nets [18]. We believe it would, straightforwardly (without requiring a dependency or leap edges).

A final question is whether the current approach can be extended to obtain proof nets for MALL1. We see two ways forward that could succeed: (i) combine witness nets with MALL slice nets [21] using a slicing criterion; (ii) combine witness nets or unification nets with MALL conflict nets [20] using a coalescence criterion. For both, there are still significant combinatorial challenges, for example because coalescence for MALL is not a straightforward extension of that for ALL. Other approaches are much less certain: MALL conflict nets do not yet have a slicing criterion (though this looks feasible), and MALL slice nets do not support coalescence (which seems fundamentally problematic).

References

- 1 Samson Abramsky. Sequentiality vs. Concurrency in Games and Logic. *Mathematical Structures in Computer Science*, 13(4):531–565, 2003.
- 2 Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In *CONCUR*, volume 6269 of *LNCS*, pages 222–236, 2010.
- 3 Robin Cockett and Luigi Santocanale. On the Word Problem for $\Sigma\Pi$ -Categories, and the Properties of Two-Way Communication. In *Computer Science Logic (CSL), 18th Annual Conference of the EACSL*, pages 194–208, 2009.
- 4 Vincent Danos. *La Logique Linéaire appliquée à l’étude de divers processus de normalisation (principalement du Lambda-calcul)*. PhD thesis, Université Paris 7, 1990.
- 5 Vincent Danos and Laurent Regnier. The structure of multiplicatives. *Archive for Mathematical Logic*, 28:181–203, 1989.
- 6 Didier Galmiche and Jean-Yves Marion. Semantic Proof Search Methods for ALL – a first approach –. Short paper in Theorem Proving with Analytic Tableaux, 4th International Workshop (TABLEAUX’95). Available from the first author’s webpage, 1995.
- 7 Jean-Yves Girard. Linear Logic. *Theoretical Computer Science*, 50(1):1–102, 1987.
- 8 Jean-Yves Girard. Proof-nets: the parallel syntax for proof-theory. *Logic and Algebra*, pages 97–124, 1996.
- 9 Stefano Guerrini and Andrea Masini. Parsing MELL proof nets. *Theoretical Computer Science*, 254(1-2):317–335, 2001.
- 10 Willem Heijltjes. Classical proof forestry. *Ann. Pure Appl. Logic*, 161(11):1346–1366, 2010.
- 11 Willem Heijltjes. Proof nets for additive linear logic with units. In *26th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 207–216, 2011.
- 12 Willem Heijltjes and Dominic J. D. Hughes. Complexity bounds for sum–product logic via additive proof nets and Petri nets. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 80–91, 2015.
- 13 Willem Heijltjes, Dominic J. D. Hughes, and Lutz Straßburger. Proof nets for first-order additive linear logic. Technical Report RR-9201, INRIA, 2018. URL: hal.inria.fr/hal-01867625/.

- 14 Jacques Herbrand. Investigations in proof theory: The properties of true propositions. In Jean van Heijenoort, editor, *From Frege to Gödel: A source book in mathematical logic, 1879–1931*, pages 525–581. Harvard University Press, 1967.
- 15 Stefan Hetzl. A sequent calculus with implicit term representation. In *Computer Science Logic (CSL)*, volume 6247 of *LNCS*, pages 351–365, 2010.
- 16 Kohei Honda, Vasco Vasconcelos, and Makoto Kubo. Language primitives and type disciplines for structured communication-based programming. In *European Symposium on Programming*, pages 122–138, 1998.
- 17 Dominic J. D. Hughes. Proofs Without Syntax. *Annals of Mathematics*, 164(3):1065–1076, 2006.
- 18 Dominic J. D. Hughes. Unification nets: canonical proof net quantifiers. In *33rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2018.
- 19 Dominic J. D. Hughes. First-order proofs without syntax. Available at arXiv.org, 2019.
- 20 Dominic J. D. Hughes and Willem Heijltjes. Conflict nets: efficient locally canonical MALL proof nets. In *31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2016.
- 21 Dominic J. D. Hughes and Rob van Glabbeek. Proof nets for unit-free multiplicative-additive linear logic. *Transactions on Computational Logic*, 6(4):784–842, 2005.
- 22 Andre Joyal. Free Lattices, Communication and Money Games. *Proc. 10th Int. Cong. of Logic, Methodology and Philosophy of Science*, 1995.
- 23 Joachim Lambek. Deductive Systems and Categories I, II, III. *Theory of Computing Systems (I), Lecture Notes in Mathematics (II, III)*, 1968–1972.
- 24 Roberto Maieli. Retractable Proof Nets of the Purely Multiplicative and Additive Fragment of Linear Logic. In *14th Int. Conf. Logic for Programming Artificial Intelligence and Reasoning (LPAR)*, pages 363–377, 2007.
- 25 Dale Miller. A Compact Representation of Proofs. *Studia Logica*, 46(4):347–370, 1987.
- 26 Samuel Mimram. The Structure of First-Order Causality. *Mathematical Structures in Computer Science*, 21(1):65–110, 2011.
- 27 Lutz Straßburger. A Characterisation of Medial as Rewriting Rule. In Franz Baader, editor, *Term Rewriting and Applications, RTA'07*, volume 4533 of *LNCS*, pages 344–358, 2007.
- 28 Philip Wadler. Propositions as Sessions. *Journal of Functional Programming*, 24(2-3):384–418, 2014.

A Proofs

In this appendix we restate and prove all theorems, and add two supporting lemmata.

► **Theorem 6.** *For any ALL1 proof π , the witness net $[\pi]$ sequentializes to π .*

Proof. It follows by induction on π that if $\lambda_\Sigma = [\pi \vdash A, B]_\sigma^{A', B'}$ where $A'\sigma = A$ and $B'\sigma = B$, then $\lambda_\Sigma^* \triangleright A, B$ reduces in (\rightarrow) to $\{(A', B')_\sigma^\pi\} \triangleright A, B$. The statement is the case $\sigma = \emptyset$. ◀

► **Theorem 7.** *If $\lambda_\Sigma \triangleright A, B$ sequentializes to π , then $[\pi]$ is $\lambda_\Sigma \triangleright A, B$.*

Proof. By induction on the sequentialization path $\lambda_\Sigma^* \triangleright A, B \rightarrow^* \{(A, B)_\emptyset^\pi\} \triangleright A, B$ it follows that in every pre-net $\kappa_\emptyset^\Phi \triangleright A, B$ on this path, λ_Σ is equal to the union over the de-sequentialization of all proof labels ϕ in Φ :

$$\lambda_\Sigma = \bigcup \{ [\phi]_\sigma^{C, D} \mid (C, D)_\sigma^\phi \in \kappa_\emptyset^\Phi \}.$$

The statement is then the case $\kappa_\emptyset^\Phi = \{(A, B)_\emptyset^\pi\}$. ◀

► **Theorem 8.** *Witness nets are canonical: $[\pi] = [\phi]$ if and only if $\pi \sim \phi$.*

Proof. From left to right is by inspection of the critical pairs of sequentialization (\rightarrow). From right to left is by inspection of the rule permutations in Figure 2. \blacktriangleleft

► **Lemma 23.** *Strict coalescence preserves and reflects correctness.*

Proof. For a strict coalescence step $L \rightarrow R$, we will show that the witness pre-net L is correct if and only if R is. Let $L = \lambda_\Sigma \triangleright A, B$ and $R = \kappa_\Theta \triangleright A, B$. In each case, exact coverage and local eigenvariables are immediately preserved and reflected. For slice-correctness, we will demonstrate that the left-hand side and right-hand side of each rule belong to the same slice of $\vdash A, B$, or in the case of $\exists S$, naturally corresponding slices. For dependency-correctness, we will briefly show how acyclicity of the columns of the involved links is preserved.

- $(C, D_i)_\sigma \rightarrow (C, D_1 + D_2)_\sigma$
A slice S_B of B and \emptyset containing one of (D_1, τ) , (D_2, τ) , and $(D_1 + D_2, \tau)$ must also contain the other two. A slice S of $\vdash A, B$ then contains all three of $(C, D_1)_\sigma$, $(C, D_2)_\sigma$, and $(C, D_1 + D_2)_\sigma$, or none. It follows that $S \cap \lambda_\Sigma$ is a singleton if and only if $S \cap \kappa_\Theta$ is. Since other slices are unaffected, L is slice-correct if and only if R is.
For dependency-correctness, the column of $(C, D_i)_\sigma$ is that of $(C, D_1 + D_2)_\sigma$ plus the pair $(D_i, \sigma|_{\text{EV}_B(D_i)})$ itself, which is minimal in the order \preceq .
- $(C, D_1)_\sigma, (C, D_2)_\sigma \rightarrow (C, D_1 \times D_2)_\sigma$
A slice S of $\vdash A, B$ contains $(C, D_1 \times D_2)_\sigma$ if and only if it contains either of $(C, D_1)_\sigma$ or $(C, D_2)_\sigma$, and cannot contain both. Then $S \cap \lambda_\Sigma$ is a singleton if and only if $S \cap \kappa_\Theta$ is. Dependency-correctness is immediate, as above.
- $(C, D)_\sigma \rightarrow (C, \exists x.D)_{\sigma-x}$
A slice S of $\vdash A, B$ contains $(C, \exists x.D)_\tau$ if and only if it contains all links $(C, D)_{\tau[t/x]}$ for any term t . Letting $\tau[t/x] = \sigma$, then $S \cap \lambda_\Sigma$ is the singleton $\{(C, D)_\sigma\}$ if and only if $S \cap \kappa_\Theta$ is $\{(C, \exists x.D)_{\sigma-x}\}$.
For dependency-correctness, the column of $(C, D)_\sigma$ is that of $(C, \exists x.D)_{\sigma-x}$ plus a pair (D, τ) , which is minimal in (\preceq) .
- $(C, D)_\sigma \rightarrow (C, \forall x.D)_\sigma$
A slice S of $\vdash A, B$ contains $(C, D)_\sigma$ if and only if it contains also $(C, \forall x.D)_\sigma$, and hence $S \cap \lambda_\Sigma$ is a singleton if and only if $S \cap \kappa_\Theta$ is.
For dependency-correctness, the column of $(C, D)_\sigma$ is that of $(C, \forall x.D)_\sigma$ plus a pair (D, τ) . The side-condition of the coalescence step is that $x \notin \sigma$; then x does not occur free in any (X, ρ) , and (D, τ) is minimal in (\preceq) . \blacktriangleleft

► **Lemma 24.** *To a correct witness pre-net $\lambda_\Sigma \triangleright A, B$ a coalescence step applies, unless it is fully coalesced already, $\lambda_\Sigma = \{(A, B)_\emptyset\}$.*

Proof. Let the **depth** of a link $(C, D)_\sigma$ be a pair of integers (n, m) , where n is the distance from C to the root of A , and m that from D to B . We order link depth in the product order: $(i, j) \leq (n, m)$ if and only if $i \leq n$ and $j \leq m$. We will demonstrate that a link at maximal depth may always be coalesced, unless it is the unique link $(A, B)_\emptyset$ at $(0, 0)$.

To see that a maximally deep link coalesces, first note that a link $(C, D_i)_\sigma$ where D_i occurs in $D_0 + D_1$ may always coalesce, as may a link $(C, D)_\sigma$ where D occurs in $\exists x.D$. This leaves the following cases:

- $(A, D_i)_\sigma$ with D_i occurring in $D = D_1 \times D_2$.
Without loss of generality, let $i = 1$. A slice S_1 of $\vdash A, B$ containing $(A, D_1)_\sigma$ has a counterpart S_2 containing $(A, D_2)_\sigma$. The depth of $(A, D_2)_\sigma$ is the same as that of $(A, D_1)_\sigma$. By correctness $S_2 \cap \lambda_\Sigma$ is a singleton; by the assumption of maximality it

22:20 Proof Nets for First-Order Additive Linear Logic

may not contain a deeper link than $(A, D_2)_\sigma$; and it may not contain a shallower one since that would be shared with $S_1 \cap \lambda_\Sigma$. Then $\lambda_\Sigma \triangleright A, B$ contains both $(A, D_1)_\sigma$ and $(A, D_2)_\sigma$, and these contract to $(A, D)_\sigma$.

- $(A, D)_\sigma$ with D in $\forall x.D$.

The step $(A, D)_\sigma \rightarrow (A, \forall x.D)_\sigma$ applies if $x \notin \sigma$. By way of contradiction, assume $x \in \sigma$. The column of $(A, D)_\sigma$ contains (D, σ_D) and $(\forall x.D, \tau)$ where $\tau = \sigma|_{\text{EV}_B(\forall x.D)}$. By the exact coverage condition, $\sigma = \sigma_A \cup \sigma_D$, and since the free existential variables in D and $\forall x.D$ are the same, $\text{EV}_B(D) = \text{EV}_B(\forall x.D)$, so that $\tau = \sigma_D$. (Note that since $\sigma_A = \emptyset$, we get $\sigma = \sigma_D = \tau$, but this is not essential to the argument.) Since $x \in \sigma$ we have $x \in \tau$, and in the column of $(A, D)_\sigma$ we have $(\forall x.D, \tau) \preceq (D, \tau)$ since D occurs as $\forall x.D$. But we already have $(D, \tau) \preceq (\forall x.D, \tau)$ because $D \leq \forall x.D$, contradicting antisymmetry of (\preceq) . Then $x \notin \sigma$, and the step $(A, D)_\sigma \rightarrow (A, \forall x.D)_\sigma$ applies.

- $(C_i, D_j)_\sigma$ in $C = C_1 \times C_2$ and $D = D_1 \times D_2$.

Without loss of generality, let $i = j = 1$. By minimal depth and using similar reasoning to the first case above, the pre-net must contain one of the following three configurations.

1. $(C_1, D_1)_\sigma, (C_1, D_2)_\sigma, (C_2, D_1)_\sigma, (C_2, D_2)_\sigma$
 2. $(C_1, D_1)_\sigma, (C_1, D_2)_\sigma, (C_2, D)_\sigma$
 3. $(C_1, D_1)_\sigma, (C_2, D_1)_\sigma, (C, D_2)_\sigma$
- $C_1 \times C_2$

 $D_1 \times D_2$

$C_1 \times C_2$

 $D_1 \times D_2$

$C_1 \times C_2$

 $D_1 \times D_2$

In the second case, the step $(C_1, D_1)_\sigma, (C_1, D_2)_\sigma \rightarrow (C_1, D)_\sigma$ applies; in the third case, $(C_1, D_1)_\sigma, (C_2, D_1)_\sigma \rightarrow (C, D_1)_\sigma$; and in the first case, both.

- $(C_i, D)_\sigma$ in $C = C_1 \times C_2$ and $\forall x.D$.

Without loss of generality let $i = 1$. If $x \notin \sigma$ the rewrite step $(C_1, D)_\sigma \rightarrow (C_1, \forall x.D)_\sigma$ applies. Otherwise, let $x \in \sigma$. The slice S_1 of $\vdash A, B$ containing $(C_1, D)_\sigma$ has a counterpart S_2 containing $(C_2, D)_\sigma$, which must include exactly one link of λ_Σ . By the assumption of minimal depth, it cannot have greater depth than $(C_2, D)_\sigma$. It cannot be $(C, D)_\sigma$ or any shallower link, since that would be shared with the slice S_1 which already contains $(C_1, D)_\sigma$. It cannot be $(C_2, \forall x.D)_\sigma$ or any shallower link $(C_2, X)_\tau$ (i.e. with $\forall x.D \leq X$) because $x \in \sigma$. This would mean either $x \in \tau$ which contradicts the *eigenvariables not free* convention, or $x \in \text{FV}(\sigma(y))$ where $\forall x.D < \exists y.Y \leq X$ which creates a cyclic column, as in the second case above. It follows that $S_2 \cap \lambda_\Sigma = \{(C_2, D)_\sigma\}$, so that the rewrite step $(C_1, D)_\sigma, (C_2, D)_\sigma \rightarrow (C, D)_\sigma$ applies.

- $(C, D)_\sigma$ in $\forall x.C$ and $\forall y.D$.

A rewrite step $(C, D)_\sigma \rightarrow (\forall x.C, D)_\sigma$ or $(C, D)_\sigma \rightarrow (C, \forall y.D)_\sigma$ applies unless $x, y \in \sigma$. But that would generate a cycle in the column of $(C, D)_\sigma$, in one of three ways. If $x \in \sigma_C$ or $y \in \sigma_D$ then, since $\sigma_C = \sigma_{\forall x.C}$ and $\sigma_D = \sigma_{\forall y.D}$, respectively:

$$(C, \sigma_C) \preceq (\forall x.C, \sigma_C) \preceq (C, \sigma_C) \quad (D, \sigma_D) \preceq (\forall y.D, \sigma_D) \preceq (D, \sigma_D) .$$

Otherwise, if $x \in \sigma_D$ and $y \in \sigma_C$ then

$$(C, \sigma_C) \preceq (\forall x.C, \sigma_C) \preceq (D, \sigma_D) \preceq (\forall x.D, \sigma_D) \preceq (C, \sigma_C) . \quad \blacktriangleleft$$

► **Theorem 12.** *A witness pre-net that strict-coalesces is correct, and a correct witness pre-net strongly strict-coalesces.*

Proof. For the first statement, we proceed by induction on the coalescence path from $\lambda_\Sigma \triangleright A, B$ to $\{(A, B)_\emptyset\} \triangleright A, B$, with the end result as the base case. It is slice-correct: every slice of $\vdash A, B$ contains $(A, B)_\emptyset$, so every slice of $\{(A, B)_\emptyset\} \triangleright A, B$ is the singleton $\{(A, B)_\emptyset\}$. It is also dependency-correct: the column of $(A, B)_\emptyset$ is the set $\{(A, \emptyset), (B, \emptyset)\}$, where A and B are unrelated in (\preceq) . For the inductive step, by Lemma 23 coalescence

reflects correctness, so that any pre-net along the coalescence path is correct, in particular $\lambda_\Sigma \triangleright A, B$.

For the second statement, let $\lambda_\Sigma \triangleright A, B$ be correct. By Lemma 24 either the net has coalesced, or a coalescence step applies. By Lemma 23 the result of any coalescence step is again correct. Since links strictly move towards the roots of both formula trees, it follows that this process terminates, and the pre-net $\lambda_\Sigma \triangleright A, B$ strongly strict-coalesces. ◀

► **Theorem 15.** *If proof nets $\lambda_\Sigma \triangleright A, B$ and $\kappa_\Theta \triangleright \overline{B}, C$ sequentialize to π and ϕ respectively, then their composition $(\lambda_\Sigma \triangleright A, B); (\kappa_\Theta \triangleright \overline{B}, C)$ is well-defined (i.e. all fixed points are finite) and sequentializes to a normal form ψ of $\pi; \phi$.*

Proof. By Theorem 12 the proof nets $L = \lambda_\Sigma \triangleright A, B$ and $R = \kappa_\Theta \triangleright \overline{B}, C$ strongly coalesce. We may then interleave their coalescence sequences as follows: if a synchronized step in L and R on the interface B and \overline{B} is available, apply it; otherwise perform steps in L on A and in R on C until it is. This gives the following combined sequence.

$$\begin{array}{ccccccc} L & = & L_1 & \xrightarrow{?} & L_2 & \xrightarrow{?} & \dots & \xrightarrow{?} & L_n \\ R & = & R_1 & \xrightarrow{?} & R_2 & \xrightarrow{?} & \dots & \xrightarrow{?} & R_n \\ \Downarrow & & \Downarrow & & \Downarrow & & & & \Downarrow \\ L; R & = & L_1; R_1 & \rightarrow & L_2; R_2 & \rightarrow & \dots & \rightarrow & L_n; R_n \end{array}$$

(Here, $(\xrightarrow{?})$ is the relation $(\rightarrow) \cup (=)$, but we assume that at least $L_i \rightarrow L_{i+1}$ or $R_i \rightarrow R_{i+1}$.) The path along the top and right of this diagram sequentializes L to π' and R to ϕ' (equivalent to π and ϕ respectively), and then composes to $L_n; R_n = \{(A, C)_{\emptyset}^{\pi'; \phi'}\} \triangleright A, C$.

Each square of the diagram converges as one of the critical pairs of sequentialization and composition discussed above. Then each path along the diagram from top left (L and R) to bottom right ($L_n; R_n$) gives a sequentialization, with cuts, of $L_n; R_n$. Let the path taking the vertical step from L_i and R_i to $L_i; R_i$ sequentialize to ψ_i , so that $\psi_n = \psi'$. By the way each square converges, we have that ψ_i is reached from ψ_{i+1} by a cut-elimination or permutation step.

Finally, in L and R every link is an axiom link. Any link in $L; R$ is composed from two links $(a, b)_\sigma$ in L and $(\overline{b}, c)_\tau$ in R , which yields $(a, c)_\rho$ where $\rho = \sigma_a \tau_c \cdot \overline{\sigma_b \tau_b}$. This sequentializes to the axiom $\vdash a\rho, c\rho$, which is in normal form. Then $L; R$ is a proof net (it has an axiom linking and it coalesces), and it sequentializes to a normal form of ψ . ◀

► **Lemma 18.** *In (\rightsquigarrow) , if $\lambda_\Sigma \triangleright A, B$ sequentializes to π then $\lambda_\star \triangleright A, B$ sequentializes to $\pi' \leq \pi$.*

Proof. The sequentialization path $\lambda_\Sigma^\star \triangleright A, B = L_1 \rightsquigarrow L_2 \rightsquigarrow \dots \rightsquigarrow L_n = (A, B)_{\emptyset}^\pi \triangleright A, B$ has a corresponding path $\lambda_\star^\star \triangleright A, B = R_1 \rightsquigarrow R_2 \rightsquigarrow \dots \rightsquigarrow R_n = (A, B)_{\emptyset}^{\pi'} \triangleright A, B$ where the same links (but with potentially different witness maps) are coalesced. It follows by induction on this path (where the base case is L_1 and R_1) that for every corresponding pair of links $(C, D)_\sigma^\phi$ in L_i and $(C, D)_\tau^\psi$ in R_i we have $\tau \leq \sigma$ and $\psi \leq \phi$. ◀

► **Lemma 19.** *If $\lambda_\star \triangleright A, B$ unifying-sequentializes to π then there exists a witness assignment Σ and substitution ρ such that $\lambda_\Sigma \triangleright A, B$ strict-sequentializes to π and $\lambda_\Sigma = \lambda_\star \rho$.*

Proof. By induction on the sequentialization path $\lambda_\star \triangleright A, B \rightsquigarrow^\star (A, B)_{\emptyset}^\pi \triangleright A, B$. For the end result, the statement holds with $\rho = \emptyset$. For the inductive step, consider a step $L \rightsquigarrow R$. We show the case $(\times U)$; the other cases are immediate.

22:22 Proof Nets for First-Order Additive Linear Logic

$$\blacksquare (C, D_1)_\sigma, (C, D_2)_\tau \rightsquigarrow (C, D_1 \times D_2)_{\sigma \vee \tau}$$

By the inductive hypothesis, $R\rho'$ strict-sequentializes to π . Let $\sigma \vee \tau = \sigma\rho'' = \tau\rho''$ and let $\rho = \rho''\rho'$. Then $L\rho$ strict-sequentializes to π by

$$(C, D_1)_{\sigma\rho}, (C, D_2)_{\tau\rho} \rightarrow (C, D_1 \times D_2)_{(\sigma \vee \tau)\rho'}. \quad \blacktriangleleft$$

► **Theorem 20.** *If $[\pi \vdash A, B]$ is $\lambda_\Sigma \triangleright A, B$ then $\lambda \triangleright A, B$ unifying-sequentializes to $\pi' \leq \pi$.*

Proof. By Theorem 6, $\lambda_\Sigma \triangleright A, B$ sequentializes to π in (\rightarrow) , and hence also in (\rightsquigarrow) . Then by Lemma 18 $\lambda_* \triangleright A, B$ sequentializes to $\pi' \leq \pi$. \blacktriangleleft

► **Theorem 21.** *If $\lambda \triangleright A, B$ sequentializes to π , then $[\pi] = \lambda_\Sigma \triangleright A, B$ for some Σ .*

Proof. By Lemma 19, since $\lambda \triangleright A, B$ sequentializes to π there is a net $\lambda_\Sigma \triangleright A, B$ that sequentializes to π . By Theorem 7, $[\pi] = \lambda_\Sigma \triangleright A, B$. \blacktriangleleft

► **Theorem 22.** *If $\lambda \triangleright A, B$ sequentializes to π and $\kappa \triangleright \overline{B}, C$ to ϕ then their composition $\lambda; \kappa \triangleright A, C$ sequentializes to a proof $\psi' \leq \psi$ where ψ is a normal form of $\pi; \phi$.*

Proof. By Lemma 19 there are witness labellings Σ and Θ such that $\lambda_\Sigma \triangleright A, B$ strict-sequentializes to π and $\kappa_\Theta \triangleright \overline{B}, C$ to ϕ . By Theorem 15 their composition $(\lambda_\Sigma; \kappa_\Theta) \triangleright A, C$ strict-sequentializes to a normal form ψ of $\pi; \phi$. By Lemma 18 the net $(\lambda; \kappa)_* \triangleright A, C$ unifying-sequentializes to $\psi' \leq \psi$. \blacktriangleleft