Intuitionistic proofs without syntax

Willem B. Heijltjes
Dept. Computer Science
University of Bath
UK

Dominic J. D. Hughes
Logic Group
U.C. Berkeley
USA

Lutz Straßburger
Parsifal
Inria
France

***Abstract*—We present Intuitionistic Combinatorial Proofs (ICPs), a concrete geometric semantics of intuitionistic logic based on the principles of the second author's classical Combinatorial Proofs.**

An ICP naturally factorizes into a linear fragment, a graphical abstraction of an IMLL proof net (an *arena net*), and a parallel contraction-weakening fragment (a *skew fibration*). ICPs relate to game semantics, and can be seen as a strategy in a Hyland-Ong arena, generalized from a tree-like to a dag-like strategy.

Our first main result, *Polynomial Full Completeness*, is that ICPs as a semantics are complexity-aware: the translations to and from sequent calculus are size-preserving (up to a polynomial). By contrast, lambda-calculus and game semantics incur an exponential blowup. Our second main result, *Local Canonicity*, is that ICPs abstract fully and faithfully over the non-duplicating permutations of the sequent calculus, analogously to the first and second authors' recent result for MALL.

I. Introduction

The research we present here started out with the goal of understanding the computational content behind classical combinatorial proofs [1], [2] by attempting to elicit an intuitionistic counterpart. While our primary result, a definition of ***intuitionistic combinatorial proof*** (***ICP***), achieves this goal, perhaps even more interesting are some unanticipated connections to prior art. An ICP turned out to be a chimera which can be viewed variously as:

(1) a canonical representation of an intuitionistic sequent calculus proof modulo non-duplicating rule permutations;
(2) an ***abstract IMLL***[1] ***proof net*** [3] skew-fibred over an intuitionistic formula;
(3) an ***abstract innocent strategy*** [4] skew-fibred over an arena, admitting disciplined sub-strategy sharing;
(4) a λ-calculus term admitting disciplined sub-term sharing.

While (1) and (2) are fleshed out in this paper (Sections II–VII), (3) and (4) represent new and unexplored avenues for future work, prompting more questions than answers; we report initial results on (3) in Section VIII. We refer to our main result (1) as ***polynomial full completeness***, since it is a full completeness (or full abstraction) result in which the surjection from intuitionistic sequent proofs to ICPs is a polynomial-time function. (We also provide a polynomial-time right inverse to the surjection.) This contrasts with other abstract representations of intuitionistic sequent proofs, such as game semantics or λ-calculus terms, whose surjections incur an exponential blow-up. Thus ICPs provide a semantics of intuitionistic proofs which not only *respects structure*, but also *respects complexity*.

ICP full completeness is also *faithful*: two sequent proofs are equivalent modulo non-duplicating rule transpositions if and only if they have the same ICP. Due to the intimate relationship between duplicating rule permutations and exponential-blowup, our polynomial full and faithful completeness theorem can be recast as:

- ICPs provide the most abstract representation of intuitionistic sequent proofs possible, subject to the constraint that the translation from sequent calculus is polynomial time.

Combinatorial proofs. Combinatorial proofs were introduced by the second author as a syntax-free reformulation of classical propositional logic [1], [2]. For example, here is a combinatorial proof of Peirce's Law $((P \Rightarrow Q) \Rightarrow P) \Rightarrow P$:

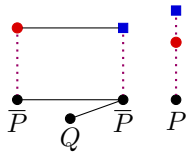

The lower graph abstracts the formula (one vertex per propositional variable, edges encoding conjunctive relationships); the upper graph has two colour classes, ● ● and ■ ■, each expressing an axiom $P \Rightarrow P$; the dotted lines define a *skew fibration* from the upper graph to the lower graph, a lax notion of graph fibration. The upper graph captures the axioms and logical rules in a proof, the lower graph captures the formula proved, and the skew fibration captures all contraction and weakening, simultaneously and in parallel [2], [5].

The intuitionistic setting required reformulating combinatorial proofs with directed edges for implicative relationships. We model a formula or sequent as a partitioned dag (directed acyclic graph) called an ***arena***, akin to an arena in game semantics [4], [6]. Here is a sequent and its arena:

$$(P \Rightarrow P) \Rightarrow Q \vdash Q \wedge Q$$

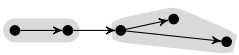

The five atom-occurrences become five vertices, and vertices are in the same partition (shown as a shaded region) if they are associated with the same atom (P or Q). Here are two

[1]Intuitionistic multiplicative linear logic.

978-1-7281-3608-0/19/$31.00

ICPs on $(P \Rightarrow P) \Rightarrow Q \vdash Q \wedge Q$:

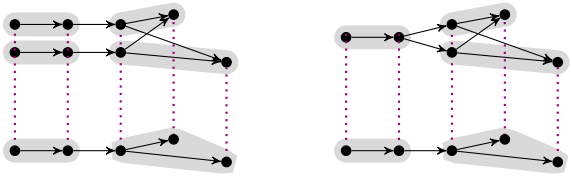

The lower graph, the *base*, is the arena of the sequent being proved. The upper graph, the *cover*, is also an arena, an abstraction of an IMLL proof net which we call an ***arena net***: each partition has two elements, and can be viewed as an axiom link. The dotted lines, defining the skew fibration, can be thought of as indicating where the two ends of each axiom link end up in the intuitionistic sequent.

It can be convenient to represent partitions using colours or labels, for example, rendering the two ICPs above as

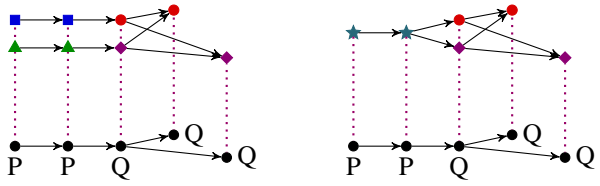

Here we have chosen vertex colour (and shape) on the cover, and the original atom labels on the base. Leaving base graphs implicit, we can render the combinatorial proofs concisely:

Using this notation, Figure 1 shows step-by-step translations of intuitionistic sequent calculus proofs into the respective ICPs above. Figure 2 shows the corresponding lambda calculus translations. The resulting lambda terms are identical (modulo bound variable renaming), and the right translation duplicates $\lambda z.z$. Because of iterated duplications, translation to a lambda term is exponential-time in the size of the proof. In contrast, translating the right proof to an ICP involves no duplication.

Abstract innocent strategies, shared sub-strategies, and implicit time. We illustrate a surjection from ICPs (Section VIII) to innocent strategies with an example. Below we have taken the ICP above-left and drawn the top row of its cover below-left, and the bottom row below-right.

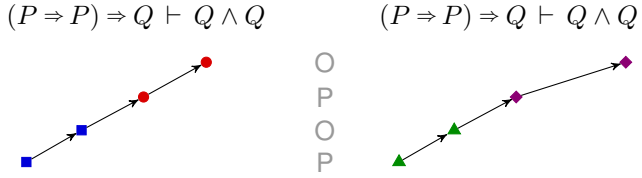

Each is a maximal 4-move play in an arena dialogue game alternating between O and P with time flowing downwards and justification pointers back up to previous moves. Each play is a P-view (O-shortsighted), and taken together the two P-views constitute an innocent strategy. Thus we can view an ICP as an *abstract innocent strategy* (the cover), skew-fibred over an arena. The sibling ICP drawn earlier above-right induces the same two P-views, hence the same innocent strategy. However, since the "top row" ★→★→●→● and "bottom row" ★→★→◆→◆ share ★→★, the cover is an abstract innocent strategy with a *shared sub-strategy* (the ★→★ part). As abstract innocent strategies, ICPs leave time implicit. For example, the ICPs

each unfold into a single play (maximal P-view), respectively

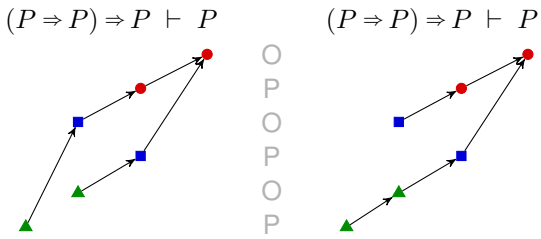

There is no explicit linear ordering in the covers: time appears spontaneously during the unfolding of the play.

Abstract IMLL nets. As remarked earlier, every ICP cover is an *arena net*, an abstraction of an IMLL proof net in the *connective-free* spirit of Retoré [7]. To illustrate the relationship to IMLL proof nets, observe that the arena net below-left is the arena of the formula below-right.

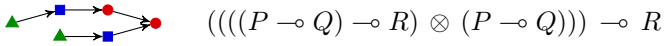

We have suggestively used linear implication $\multimap$ and conjunction $\otimes$ in place of $\Rightarrow$ and $\wedge$: the resulting binary IMLL formula constitutes an IMLL net in the sense of Lamarche [3].

Acknowledgments. The first author was supported by EPSRC Project EP/R029121/1 *Typed Lambda-Calculi with Sharing and Unsharing*. The second author is grateful to his hosts at the Berkeley Logic Group, Wes Holliday and Dana Scott. The third author was supported by the ANR-FWF international grant ANR-5-CE25-0014 *The Fine Structure of Formal Proof Systems and their Computational Interpretations (FISP)*. We thank the referees for their diligence.

II. Arenas

Our first objective is to give a graphical representation of intuitionistic formulas, as a counterpart to the ***cographs*** for classical formulas in [1]. We will call our graphs ***arenas***, for the similarities with game semantics [4], [6], [8].

In implication-only intuitionistic logic, formulas are naturally interpreted as a tree: $A_1 \Rightarrow \ldots \Rightarrow A_n \Rightarrow P$ generates the tree with P as the root and (inductively) the trees for each A_i as subtrees. For example:

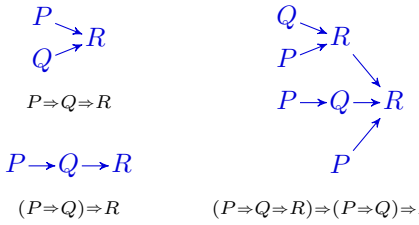

The *currying* isomorphism $A \Rightarrow B \Rightarrow C \cong (A \wedge B) \Rightarrow C$ then naturally gives the interpretation of conjunction, as the *sibling*

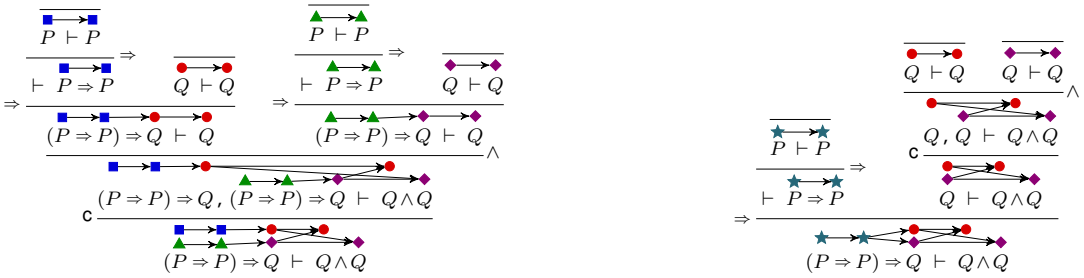

Fig. 1. Translating two intuitionistic sequent calculus proofs to intuitionistic combinatorial proofs (ICPs). The translation is relatively simple to describe: (1) place a pair of tokens atop the propositional variables in each axiom, with a rightward directed edge; (2) trace the tokens down through the proof; (3) each left implication rule and right conjunction rule inserts edges between tokens above the active formulas (details in Section IV).

$$\dfrac{\dfrac{\dfrac{x:P\vdash x:P}{\vdash \lambda x.x : P\Rightarrow P}\quad w:Q\vdash w:Q}{f_1 : (P\Rightarrow P)\Rightarrow Q \;\vdash\; f_1(\lambda x.x) : Q}\quad \dfrac{\dfrac{y:P\vdash y:P}{\vdash \lambda y.y : P\Rightarrow P}\quad v:Q\vdash v:Q}{f_2 : (P\Rightarrow P)\Rightarrow Q \;\vdash\; f_2(\lambda y.y) : Q}}{f_1 : (P\Rightarrow P)\Rightarrow Q\,,\; f_2 : (P\Rightarrow P)\Rightarrow Q \;\vdash\; \langle f_1(\lambda x.x), f_2(\lambda y.y)\rangle \;:\; Q\wedge Q}}{f : (P\Rightarrow P)\Rightarrow Q \;\vdash\; \langle f(\lambda x.x), f(\lambda y.y)\rangle \;:\; Q\wedge Q}$$

$$\dfrac{\dfrac{\dfrac{z:P\vdash z:P}{\vdash \lambda z.z : P\Rightarrow P}}{f : (P\Rightarrow P)\Rightarrow Q \;\vdash\; \langle f(\lambda z.z), f(\lambda z.z)\rangle \;:\; Q\wedge Q}}{}\quad \dfrac{\dfrac{v_1:Q\vdash v_1:Q\quad v_2:Q\vdash v_2:Q}{v_1:Q\,,\; v_2:Q \;\vdash\; \langle v_1,v_2\rangle \;:\; Q\wedge Q}}{v : Q \;\vdash\; \langle v, v\rangle \;:\; Q\wedge Q}$$

Fig. 2. Translating the same two intuitionistic sequent calculus proofs into λ-calculus terms. Note that the two terms are the same (up to bound variable renaming). On the right, the subterm $\lambda z.z$ from the left sub-proof is duplicated, because of extensionality. In contrast, the translation to a combinatorial proof does not require such a duplication: on the right of Figure 1, the final rule keeps only one pair of tokens over $P \Rightarrow P$, from the left sub-proof.

relation. However, a choice must be made for how to interpret a conjunction as the *consequent* of an implication, $P{\Rightarrow}(Q{\wedge}R)$: via *distribution* $A \Rightarrow (B \wedge C) \cong (A \Rightarrow B) \wedge (A \Rightarrow C)$ as below left, relaxing trees to forests; or as below right, relaxing trees to directed acyclic graphs, or *dags*.

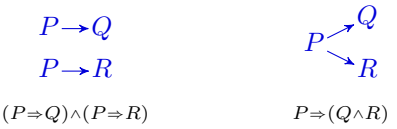

We make a fundamental choice to reject distribution:

Complexity awareness. We wish to be complexity aware, but distribution may incur exponential growth.

IMLL proof nets. We aim to interpret IMLL proof nets, but distribution is not valid in IMLL.

Our arenas will be dags. While every tree corresponds to a formula, not every dag does. We need two conditions, akin to the P_4-*free* [9], [7] (or *N-free* [10]) condition of cographs [10]: the ***L-free*** and ***Σ-free*** conditions. We define arenas after some preliminaries.

We generate ***formulas*** $A, B, C, \ldots$ from a countably infinite set of ***atoms*** $\{P, Q, R, \ldots\}$ by the binary connectives $\wedge$ (***conjunction***) and $\Rightarrow$ (***implication***). By convention $\wedge$ binds more strongly than $\Rightarrow$ and parentheses associate to the right:

$$A \wedge B \Rightarrow C = (A \wedge B) \Rightarrow C$$
$$A \wedge B \wedge C = A \wedge (B \wedge C)$$
$$A \Rightarrow B \Rightarrow C = A \Rightarrow (B \Rightarrow C).$$

A ***graph*** $\mathcal{G}$ is a finite set $V(\mathcal{G})$ of ***vertices*** and a binary ***edge*** relation $\to_{\mathcal{G}}$, abbreviated to $\to$ if $\mathcal{G}$ is clear. A ***path*** of length n from v to w is a vertex sequence $u_0 \ldots u_n$ with $u_0 = v$, $u_n = w$ and $u_{i-1} \to u_i$ for $1 \leq i < n$. Write $v \to^* w$ if there is a path from v to w, and $v \to^n w$ if there is a path of length n.

A graph $\mathcal{G}$ is ***acyclic***, or a ***dag***, if $v \to^n v$ implies $n = 0$. A ***root*** is a vertex r with no outgoing edges, $r \not\to$.

A graph $\mathcal{G}$ is ***partitioned*** if it carries an equivalence $\sim_{\mathcal{G}}$ on vertices (written $\sim$ if $\mathcal{G}$ is clear). An equivalence class is a ***partition***. An L***-labelled*** graph $\mathcal{G}$ is one equipped with a ***labelling*** $\lambda_{\mathcal{G}} \colon V(\mathcal{G}) \to L$ (written λ if $\mathcal{G}$ is clear), which induces partitioning via $v \sim w \iff \lambda(v) = \lambda(w)$.

Definition 1 (Arena)**.** A ***pre-arena*** is a non-empty dag that is
L-free : $v \leftarrow u \to x \to w$ implies $v \to w$; and
Σ-free : $u \leftarrow x \to v \leftarrow y \to w$ implies $x \to w$ or $y \to u$.
An ***arena*** is a pre-arena with a partitioning.

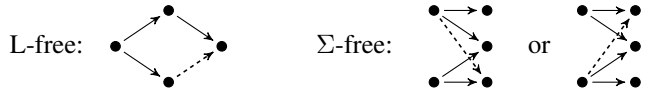

Example 2. The dag below-left is both L-free and Σ-free; that below-centre, Σ-free only; that below-right, neither.

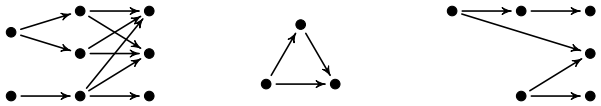

Example 3. The arena with vertices $\{u, v, w, x\}$, edge relation $u \to v \to x \leftarrow w \leftarrow u$, and equivalence classes $\{u, v\}$ and $\{w, x\}$ is shown below, in three ways: (i) with explicit vertex names, and with equivalence (ii) as shading, and (iii) as vertex shape.

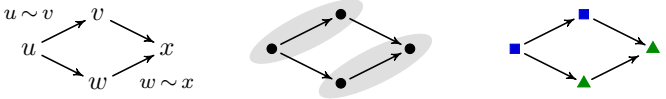

We shall model a formula A first as an atom-labelled arena $\lfloor A \rfloor$, i.e. one whose partitioning is given by labelling vertices with propositional atoms. By forgetting the labels (but not the partitioning), we then obtain the arena $[\![A]\!]$.

Definition 4 (Arena of a formula)**.** For (L-labelled) dags $\mathcal{G}$ and $\mathcal{H}$, write $\mathcal{G}+\mathcal{H}$ for their ***disjoint union***, and define the

subjunction $\mathcal{G} \triangleright \mathcal{H}$ of $\mathcal{G}$ onto $\mathcal{H}$ as $\mathcal{G}+\mathcal{H}$ plus an edge from every root of $\mathcal{G}$ to every root of $\mathcal{H}$.

$\mathcal{G}+\mathcal{H}$: 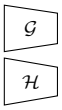

$\mathcal{G} \triangleright \mathcal{H}$: 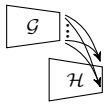

The ***(atom-)labelled arena*** $\lfloor A \rfloor$ of a formula A is given by: $\lfloor P \rfloor$ is a single vertex labelled P, and

$$\lfloor A \wedge B \rfloor = \lfloor A \rfloor + \lfloor B \rfloor \qquad \lfloor A \Rightarrow B \rfloor = \lfloor A \rfloor \triangleright \lfloor B \rfloor \ .$$

The ***arena*** $[\![A]\!]$ of A is the ***underlying arena*** of $\lfloor A \rfloor$, which forgets the labelling but retains the partitioning.

Proposition 5. *For every formula A: $\lfloor A \rfloor$ is a well-defined labelled arena, and $[\![A]\!]$ is a well-defined arena.*

Proof. A routine induction on the size of the formula. □

Observe that a formula with n atom-occurrences and d distinct atoms generates an arena with n vertices and d partitions.

Example 6. Consider the formulas

$$\begin{array}{rcl} A & = & (((P \Rightarrow P) \Rightarrow Q) \wedge Q) \ \Rightarrow \ (R \wedge ((S \Rightarrow S) \Rightarrow S)) \\ B & = & (((P \Rightarrow P) \Rightarrow Q) \Rightarrow Q) \ \Rightarrow \ (((S \Rightarrow S) \Rightarrow S) \wedge R) \\ C & = & (((Q \Rightarrow Q) \Rightarrow R) \Rightarrow R) \ \Rightarrow \ (((T \Rightarrow T) \Rightarrow T) \wedge S) \ . \end{array}$$

The labelled arena $\lfloor A \rfloor = \lfloor B \rfloor$ is below-left; $\lfloor C \rfloor$ is below-right.

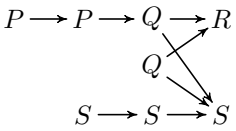

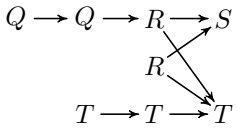

The arena $[\![A]\!] = [\![B]\!] = [\![C]\!]$ of all three is below, with partitioning through shading (left) and vertex shape (right).

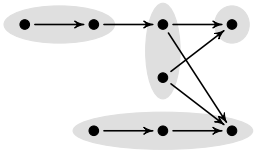

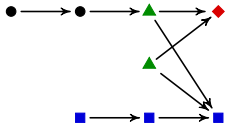

Our main theorem for arenas is that every arena arises as the interpretation of some formula. We state it here, but leave the proof to the end of this section.

Theorem 7 (Arena surjectivity)**.** *The translation $[\![-]\!]$ from formulas to arenas is surjective.*

We establish that the representation of a formula A by an arena $[\![A]\!]$ factors out the standard isomorphism ($\cong$) on formulas generated by currying, associativity, and commutativity (below), together with renaming of atoms (as in Example 6).

$$\begin{array}{rcl} (A \wedge B) \Rightarrow C & \cong & A \Rightarrow (B \Rightarrow C) \\ A \wedge (B \wedge C) & \cong & (A \wedge B) \wedge C \\ A \wedge B & \cong & B \wedge A \end{array}$$

Theorem 8 (Arena canonicity)**.** $[\![A]\!] = [\![B]\!]$ *if and only if* $A \cong B$.

Proof. A straightforward induction, via labelled arenas. □

We provide intuition for the L-free and Σ-free conditions. In a labelled arena $\lfloor A \rfloor$, an edge $P \rightarrow Q$ (identifying vertices with their labels for convenience) means A has a subformula of the form $..P.. \Rightarrow ..Q..$; call this a ***subformula constraint***.

The L-pattern of the L-free condition is the overlap of the configurations below. Each gives a subformula constraint:

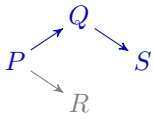

$(..P..\Rightarrow..Q..)\Rightarrow..S..$

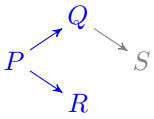

$..P..\Rightarrow(..Q..\wedge..R..)$

The formula below reconciles these constraints, and it implies the edge $R \rightarrow S$. This explains the L-free condition.

$(..P..\Rightarrow(..Q..\wedge..R..))\Rightarrow..S..$

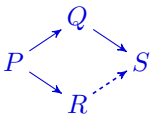

The Σ-pattern is the overlap of these two configurations:

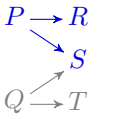

$..P..\Rightarrow(..R..\wedge..S..)$

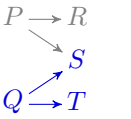

$..Q..\Rightarrow(..S..\wedge..T..)$

The constraints are resolved in the following formula, which replaces $..S..$ above left by the formula above right.

$..P..\Rightarrow(..R..\wedge(..Q..\Rightarrow(..S..\wedge..T..)))$

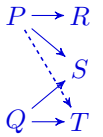

The formula implies the edge $P \rightarrow T$. In addition to the symmetric case, which implies $Q \rightarrow R$, a third solution is

$..P..\Rightarrow..Q..\Rightarrow(..R..\wedge..S..\wedge..T..)$

which implies both edges, and is thus subsumed by the other solutions. This explains the Σ-free condition.

We continue exploring the formal consequences of the L-free and Σ-free conditions, to then prove that every arena is the translation of an intuitionistic formula (Theorem 7). A reader interested mainly in the definition of combinatorial proof may skip ahead to the next section (Section III).

The ***depth*** $\mathsf{d}(v)$ of a vertex v is the length of the longest path from v. A dag is ***stratified*** if $v \rightarrow w$ implies $\mathsf{d}(v) = \mathsf{d}(w) + 1$, i.e., every path from v to a root has length $\mathsf{d}(v)$.

Lemma 9. *An L-free dag is stratified.*

Proof. For a contradiction assume $\mathcal{G}$ is not stratified. Then it contains a configuration $u \leftarrow v \rightarrow w$ with $\mathsf{d}(u) < \mathsf{d}(w)$. Hence $\mathsf{d}(w) \geq 1$. Let x be the first node on a longest path from w, i.e. $w \rightarrow x$ with $\mathsf{d}(x) = \mathsf{d}(w) - 1$. An edge from u to x would imply $\mathsf{d}(u) \geq \mathsf{d}(x) + 1 = \mathsf{d}(w)$, a contradiction. Hence,

$u \leftarrow v \rightarrow w \rightarrow x$ but not $u \rightarrow x$, and $\mathcal{G}$ is not L-free. □

For an edge $v \rightarrow w$, we call v an ***antecedent*** of w and w a ***consequent*** of v. Write $(v{\rightarrow}^n)$ for the set $\{x \mid v \rightarrow^n x\}$, the n-fold consequents of v; and $(v{\rightarrow}^*) = \bigcup_{n\geq 0}(v{\rightarrow}^n)$. Define $({\rightarrow}^n v)$ analogously. The smallest subgraph of $\mathcal{G}$ containing the vertex v closed under $\rightarrow$, i.e., induced by the vertex set $(v{\rightarrow}^*)$, is the ***cone*** of v. A stratified dag is ***full*** if from one level to the next, it contains all edges, i.e., $x \rightarrow y$ whenever x is at depth $n+1$ and y is at depth n.

Lemma 10. *In an L-free dag every cone is full.*

Proof. We will show that $x \rightarrow y$ whenever $v \rightarrow^i x$ and $v \rightarrow^{i+1} y$ by induction on i. The case $i = 0$, where $x = v$, is immediate. For the inductive step, let $v \rightarrow^{i+2} y$ via $v \rightarrow^i w \rightarrow x \rightarrow y$. For any x' such that $v \rightarrow^{i+1} x'$, by the inductive hypothesis $w \rightarrow x'$. Then w, x, x', y form an L, forcing $x' \rightarrow y$, as required. □

The Σ-freeness condition of (pre-)arenas governs how cones may or may not overlap. Two vertices v and w ***meet*** at a vertex u if u is a vertex of greatest depth belonging to both cones, thus $v \rightarrow^n u$ and $w \rightarrow^m u$ for least n and m. We say u is a ***meeting vertex*** or ***meet*** of v and w. Σ-freeness implies that where two cones meet, one becomes included in the other; that is, their meeting vertices belong entirely to one or the other.

Lemma 11. *Let v, w, and y be vertices in a pre-arena. If $v \rightarrow^n y \leftarrow^m w$ then $(v{\rightarrow}^n) \subseteq (w{\rightarrow}^m)$ or $(w{\rightarrow}^m) \subseteq (v{\rightarrow}^n)$.*

Proof. We proceed by contradiction. Let

$$v \rightarrow^{n-1} a \rightarrow y \leftarrow b \leftarrow^{m-1} w$$

and assume $v \rightarrow^n x \not\leftarrow^m w$ and $v \not\rightarrow^n z \leftarrow^m w$. Since cones are full (Lemma 10), we have $a \rightarrow x$ and $b \rightarrow z$. Then a, b, x, y, z form a Σ, a contradiction. □

A consequence is that at any depth, but in particular at depth zero (the roots), a pre-arena can be partitioned into sections of largest cones, where nodes in separate sections do not share an antecedent. Different sections correspond to composition with $+$, while a largest cone can be decomposed along $\triangleright$.

Theorem 12. *A dag is a pre-arena if and only if it is constructed from single vertices by the operations $+$ and $\triangleright$.*

Proof. First, observe that that a single vertex is a pre-arena, and both properties (L-free and Σ-free) are preserved by $+$ and $\triangleright$. Conversely, we must show that a pre-arena $\mathcal{F}$ is either a single vertex or is composed from two pre-arenas $\mathcal{G}$ and $\mathcal{H}$ as $\mathcal{G}+\mathcal{H}$ or as $\mathcal{G} \triangleright \mathcal{H}$. We proceed by induction on the size of $\mathcal{F}$, the base case being trivial. If the roots of $\mathcal{F}$ can be partitioned into two sets with no common antecedent, then $\mathcal{F}$ has the form $\mathcal{G}+\mathcal{H}$. Otherwise, by Lemma 11 there is a largest cone $(v{\rightarrow}^*)$ which includes all roots. Let n be the depth of v, so that $(v{\rightarrow}^n)$ are the roots, and $Y = (v{\rightarrow}^{n-1})$ is a set of nodes connecting to all roots. Then $\mathcal{F}$ has the form $\mathcal{G} \triangleright \mathcal{H}$ for $\mathcal{G}$ the subgraph over the nodes $\{x \mid x \rightarrow^* y, y \in Y\}$. □

Proof of Theorem 7. Immediate by Theorem 12. □

Remark 13 (Empty arenas)**.** We are very close to modelling the conjunctive unit $\top$ by the empty dag. We have unit cancellation $A \wedge \top \cong A$ and 0-ary currying $\top \Rightarrow A \cong A$ via the equations $\mathcal{G} + \varnothing = \mathcal{G}$ respectively $\varnothing \triangleright \mathcal{G} = \mathcal{G}$, but a formula $A \Rightarrow \top$ is problematic. Subjunction currently gives $\mathcal{G} \triangleright \varnothing = \mathcal{G}$, incorrectly reducing it to A. One solution is to admit 0-ary distribution, $A \Rightarrow \top \cong \top$, by defining the special case $\mathcal{G} \triangleright \varnothing = \varnothing$ instead. While complexity is not an objection, others remain: the isomorphism is not valid in IMLL; it would mean rejecting binary but admitting 0-ary distribution; and it creates a rather ad-hoc case distinction for subjunction ($\triangleright$). We opt to simply not model $\top$, and reject empty arenas, even though this makes interpreting the *weakening* rule of sequent calculus slightly inconvenient; we use empty dags for this.

III. Arena nets

The *cover* part of an intuitionistic combinatorial proof will be an ***arena net***: an IMLL proof net [3] over an *arena* instead of a formula, in the same way that classical combinatorial proofs [1] use MLL proof nets over a *cograph* [7] instead of a formula. We will first give our definition.

The ***parity*** of a vertex v is the parity of its depth, ***even*** or ***odd***, which we indicate by v° respectively $v^\bullet$. The ***parity*** of an edge is that of its target, ***even*** $(v^\bullet \dashrightarrow w^\circ)$, indicated by a dashed arrow, or ***odd*** $(v^\circ \rightarrow w^\bullet)$.

Definition 14 (Arena Net)**.** An arena $\mathcal{G}$ is ***linked*** if every partition is of the form $\{v^\bullet, w^\circ\}$, inducing a ***link*** $v^\bullet \overset{\frown}{} w^\circ$. Its ***link graph*** $\overset{\frown}{\mathcal{G}}$ is the graph of its vertices, odd edges, and links. An ***arena net*** is a linked arena $\mathcal{G}$ satisfying:

- ***Acyclicity:*** The link graph $\overset{\frown}{\mathcal{G}}$ is acyclic.
- ***Functionality:*** Every path in $\overset{\frown}{\mathcal{G}}$ from an odd vertex $v^\bullet$ to a root includes a vertex w° such that $v^\bullet \dashrightarrow w^\circ$.

For simplicity, we will take the vertices of a linked arena $[\![A]\!]$ to be the atoms $P, Q, R, \ldots$ of A themselves, annotated for even and odd; then each link is of the form $P^\bullet \overset{\frown}{} P^\circ$.

Example 15. Below are three arena nets.

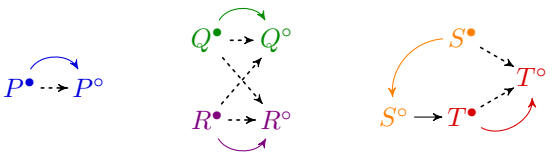

We will show that arena nets are isomorphic to IMLL proof nets, up to the formula equivalence ($\cong$) factored out by arenas. First, we give a concise treatment of IMLL. Formulas are

$$A, B, C \ ::= \ P \mid A \otimes B \mid A \multimap B$$

and sequent proofs are generated by the rules in Figure 3. An ***IMLL proof net*** [3] forgets the rule structure of a proof, retaining just the conclusion formula, plus which atoms are linked by an axiom. We capture this as follows:

Definition 16 (IMLL proof nets)**.** An ***IMLL proof net*** is the conclusion formula of a proof with distinct propositional atoms for each axiom.

$$\frac{}{P \vdash P} \qquad \frac{\Gamma, A, B \vdash C}{\Gamma, A \otimes B \vdash C} \qquad \frac{\Gamma, A \vdash B}{\Gamma \vdash A \multimap B}$$

$$\frac{\Gamma \vdash A \qquad \Delta \vdash B}{\Gamma, \Delta \vdash A \otimes B} \qquad \frac{\Gamma \vdash A \qquad B, \Delta \vdash C}{\Gamma, A \multimap B, \Delta \vdash C}$$

Fig. 3. A sequent calculus for IMLL

$$\frac{}{\bullet \vdash \bullet} \qquad \frac{\Phi, \mathcal{F}, \mathcal{G} \vdash \mathcal{H}}{\Phi, \mathcal{F} + \mathcal{G} \vdash \mathcal{H}} \qquad \frac{\Phi, \mathcal{F} \vdash \mathcal{G}}{\Phi \vdash \mathcal{F} \triangleright \mathcal{G}}$$

$$\frac{\Phi \vdash \mathcal{F} \qquad \Psi \vdash \mathcal{G}}{\Phi, \Psi \vdash \mathcal{F} + \mathcal{G}} \qquad \frac{\Phi \vdash \mathcal{F} \qquad \mathcal{G}, \Psi \vdash \mathcal{H}}{\Phi, \mathcal{F} \triangleright \mathcal{G}, \Psi \vdash \mathcal{H}}$$

Fig. 4. A sequent calculus for arena nets

Describing a proof net as a special formula streamlines the translation to arena nets. First, we interpret IMLL formulas as arenas in the same way as intuitionistic ones:

$$\lfloor A \otimes B \rfloor = \lfloor A \rfloor + \lfloor B \rfloor \qquad \lfloor A \multimap B \rfloor = \lfloor A \rfloor \triangleright \lfloor B \rfloor \ .$$

Definition 17. (IMLL-constructed arenas) An arena $[\![A]\!]$ is ***IMLL-constructed*** if A is an IMLL proof net.

Example 18. The arena nets in Example 15 are the interpretations of the following proof nets:

$$P \multimap P \qquad Q \multimap R \multimap (Q \otimes R) \qquad (S \otimes (S \multimap T)) \multimap T$$

By applying the translation $[\![-]\!]$ to the formulas of the sequent rules in Figure 3, we obtain a sequent calculus for arenas, in Figure 4. Here, Φ and Ψ are multisets of arenas, and each axiom creates a distinct partition, as indicated by shading. The following propositions are straightforward.

Proposition 19. *An arena is IMLL-constructed if and only if it is generated by a proof in the calculus of Figure 4.*

Proposition 20. *An IMLL-constructed arena is an arena net.*

The remainder of this section will be devoted to demonstrating that every arena net is IMLL-constructed (Theorem 25). The standard correctness criterion for IMLL is Lamarche's essential net condition [3]. Our ***functionality*** condition (see Definition 14) is a direct adaptation to arenas. To prove the correspondence with IMLL nets, we will relate it to the original essential net condition.

The condition relies on a special notion of path over a formula tree. We will use ***step*** and ***walk*** when traversing a formula, to distinguish from ***edge*** and ***path*** as used on dags. We refer to a position in the tree by its subformula. In line with the previous notion of parity, a subformula is ***even*** respectively ***odd*** if it occurs within the antecedent of an even respectively odd number of linear implications. Steps in a formula tree are then as follows:

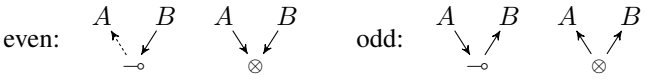

Observe that these steps generate three kinds of walk: ***neutral*** walks up or down the formula tree, ***even*** walks crossing an even implication (using the dashed step above), and ***odd*** walks crossing an odd implication. Note that at most one implication can be crossed (no walk can first go up and then down a tree).

We will assume that in a formula, atoms occur in pairs of an even and an odd occurrence $P^\bullet, P^\circ$. Each pair generates an ***axiom link***, a step from $P^\bullet$ to P°. An ***essential walk*** is a composite of odd walks, neutral walks, and axiom links.

Theorem 21 (Lamarche)**.** *A formula A is an IMLL proof net if and only if it satisfies the following correctness condition:*

- *No essential walk through the proof net is a cycle;*
- *For even $B \multimap C$ all essential walks from B to A visit C.*

(The latter is the essential net condition [3].) The key observation in connecting arena nets to IMLL proof nets is that the even and odd walks through a formula A correspond to the even and odd edges of the arena $[\![A]\!]$, where even arena edges have the opposite direction of even formula walks. The following lemmas formalize this correspondence.

Lemma 22. *The roots of $[\![A]\!]$ are those atoms with a neutral walk to A.*

Proof. By induction on A. The case $A = P$ is immediate; if $A = B \otimes C$ the roots are those of B and C, which both have a step to A; if $A = B \multimap C$ the roots of A are those of C, which has a step to A. □

Lemma 23. *There is an even edge $P \dashrightarrow Q$ (respectively an odd edge $P \rightarrow Q$) in $[\![A]\!]$ if and only if there is an even walk from Q to P (respectively an odd walk from P to Q) in A.*

Proof. By induction on A. The case $[\![P]\!]$ is immediate. For $[\![B \otimes C]\!] = [\![B]\!] + [\![C]\!]$, any edge $P \rightarrow Q$ must be created in $[\![B]\!]$ or $[\![C]\!]$, while there are no walks between B and C. For $[\![B \multimap C]\!] = [\![B]\!] \triangleright [\![C]\!]$, the parity of an edge created in $[\![B]\!]$ is inverted, as are the parity and direction of a walk in B (note that a walk between atoms P and Q must first go down then up a tree, and so traverse exactly one implication). Next, the even edge $P \dashrightarrow Q$ is created if P is a root of $[\![B]\!]$ and Q a root of $[\![C]\!]$. By Lemma 22, this is exactly when there are walks from P to B and from Q to C; since the direction of walks is inverted in B, this gives an even walk from Q to P across $B \multimap C$, as required. □

Example 24. Below is an IMLL net:

$$((P \multimap P) \multimap Q) \otimes R \ \multimap \ Q \otimes ((R \multimap S) \multimap S)$$

We make parity, steps, and axiom links explicit below left, and give its interpretation as an arena net below right.

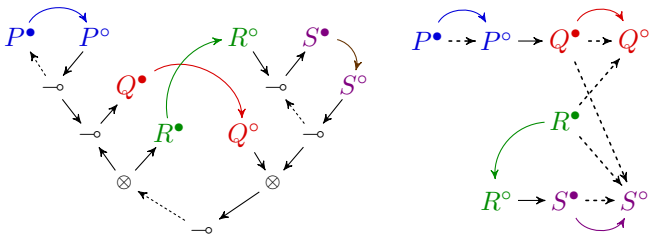

$$\frac{}{P \vdash P}\,\mathsf{ax} \qquad \frac{\Gamma,\, A,\, A \vdash B}{\Gamma,\, A \vdash B}\,\mathsf{c} \qquad \frac{\Gamma \vdash B}{\Gamma,\, A \vdash B}\,\mathsf{w}$$

$$\frac{\Gamma,\, A,\, B \vdash C}{\Gamma,\, A \wedge B \vdash C}\,\wedge\mathsf{L} \qquad \frac{\Gamma \vdash A \qquad \Delta \vdash B}{\Gamma,\, \Delta \vdash A \wedge B}\,\wedge\mathsf{R}$$

$$\frac{\Gamma \vdash A \qquad B,\, \Delta \vdash C}{\Gamma,\, A \Rightarrow B,\, \Delta \vdash C}\,\Rightarrow\mathsf{L} \qquad \frac{\Gamma,\, A \vdash B}{\Gamma \vdash A \Rightarrow B}\,\Rightarrow\mathsf{R}$$

Fig. 5. The cut-free sequent calculus LIp

Theorem 25. *For a linked arena $[\![A]\!]$ these are equivalent:*

1) *$[\![A]\!]$ is an arena net;*
2) *A is an IMLL proof net;*
3) *$[\![A]\!]$ is IMLL-constructed.*

Proof. We show that 1) implies 2) below; 2) implies 3) by Theorem 21; and 3) implies 1) by Proposition 20.

We show that the functionality condition on arenas corresponds to the essential net condition for IMLL proof nets. Consider an even subformula $B \multimap C$ in A and a walk from B to A; we will show it visits C. From B, this walk must first reach an odd atom $P_1^\bullet$ in B, then traverse a sequence of alternating axiom links (on atoms P_1 through P_n) and odd walks in A, and finally go from P_n° to A straight down the syntax tree. In $[\![A]\!]$, by Lemma 22, $P_1^\bullet$ is a root of $[\![B]\!]$, and by Lemma 23, every walk from P_i° to $P_{i+1}^\bullet$ has a corresponding odd edge in $[\![A]\!]$. This gives a path in $[\![A]\!]$, consisting of an odd edge $P_i^\circ \to P_{i+1}^\bullet$ for each i, plus the links $P_i^\bullet \rightsquigarrow P_i^\circ$. By the functionality condition, this path must visit a consequent of $P_1^\bullet$; let this be P_i°. The consequents of $P_1^\bullet$ are the roots of $[\![C]\!]$, so that by Lemma 22, there is a neutral walk in A from P_i° to C. Observe that this neutral walk allows no possibility to deviate: it does not cross an odd implication, and the only branching occurs at odd conjunctions. It then follows that the walk in A from B must pass from P_i° to C, as required. That there is no cyclic essential walk follows by similar (but simpler) reasoning from the acyclicity condition in arena nets. □

IV. Intuitionistic Combinatorial Proofs

In this section we define intuitionistic combinatorial proofs and show how to translate sequent calculus proofs into ICPs. We work with conjunction-implication intuitionistic sequent calculus, given in Figure 5. A ***sequent*** $\Gamma \vdash C$ is a pair of a formula multiset $\Gamma = A_1, \ldots, A_n$ (written as list, separated by comma) and a formula C. We extend the mapping $[\![\cdot]\!]$ to sequents by defining the labelled arena of a sequent by

$$\lfloor A_1, \ldots, A_n \vdash C \rfloor = (\lfloor A_1 \rfloor + \cdots + \lfloor A_1 \rfloor) \triangleright \lfloor C \rfloor$$

and by letting $[\![\Gamma \vdash C]\!]$ be the underlying arena of $\lfloor \Gamma \vdash C \rfloor$.

The ***meeting depth*** $\mathsf{m}(v, w)$ of distinct vertices v and w in an arena is the maximal depth of a vertex u with $v \to^* u \leftarrow^* w$ when such u exists, and -1 otherwise. Distinct vertices v and w are ***conjuncts***, denoted $v \curlywedge w$, if $\mathsf{m}(v, w)$ is odd.

Definition 26 (Skew fibration)**.** A ***skew fibration*** $f \colon \mathcal{G} \to \mathcal{H}$, of an arena $\mathcal{G}$ over an arena $\mathcal{H}$, is a function $f \colon V(\mathcal{G}) \to V(\mathcal{H})$ which preserves

edges:	$v \to_{\mathcal{G}} w$	implies	$f(v) \to_{\mathcal{H}} f(w)$
equivalence:	$v \sim_{\mathcal{G}} w$	implies	$f(v) \sim_{\mathcal{H}} f(w)$
roots:	$r \not\to_{\mathcal{G}}$	implies	$f(r) \not\to_{\mathcal{H}}$
conjuncts:	$v \curlywedge_{\mathcal{G}} w$	implies	$f(v) \curlywedge_{\mathcal{H}} f(w)$

and satisfies the following ***skew lifting*** condition [1]:

- if $f(v) \curlywedge_{\mathcal{H}} w$ there exists u with $v \curlywedge_{\mathcal{G}} u$ and $f(u) \not\curlywedge_{\mathcal{H}} w$.

Remark 27. In the next section we show that instead of demanding the preservation of roots, we could also demand the preservation of parity. Even though preservation of parity is the weaker condition, we added preservation of roots to Definition 26 because it is conceptually simpler and easier to check.

Definition 28. An ***intuitionistic combinatorial proof*** or ***ICP*** is a skew fibration f of an arena net $\mathcal{G}$ over an arena $\mathcal{H}$. An ***ICP*** of a formula A (resp. sequent $\Gamma \vdash A$) is an ICP over its arena $[\![A]\!]$ (resp. $[\![\Gamma \vdash A]\!]$). We denote this by $f \colon \mathcal{G} \to [\![A]\!]$ or $f \colon \mathcal{G} \to [\![\Gamma \vdash A]\!]$.

For translating sequent proofs to combinatorial proofs, we need the following operations on mappings between dags:

$$\begin{aligned}
1 &= \{(\bullet, \bullet)\} : \bullet \to \bullet \\
f_1 \triangleright f_2 &= f_1 \cup f_2 : \mathcal{G}_1 \triangleright \mathcal{G}_2 \to \mathcal{H}_1 \triangleright \mathcal{H}_2 \qquad (\mathcal{G}_2, \mathcal{H}_2 \neq \varnothing) \\
f_1 + f_2 &= f_1 \cup f_2 : \mathcal{G}_1 + \mathcal{G}_2 \to \mathcal{H}_1 + \mathcal{H}_2 \\
[f_1, f_2] &= f_1 \cup f_2 : \mathcal{G}_1 + \mathcal{G}_2 \to \mathcal{H} \qquad (\mathcal{H} = \mathcal{H}_1 = \mathcal{H}_2) \\
\varnothing_{\mathcal{H}} &= \varnothing : \varnothing \to \mathcal{H}
\end{aligned}$$

(illustrated below), to construct ICPs inductively.

$$\begin{array}{cccccc}
\vdots & \mathcal{G}_1 \triangleright \mathcal{G}_2 & \mathcal{G}_1 + \mathcal{G}_2 & \mathcal{G}_1 + \mathcal{G}_2 & \varnothing & \\
 & f_1 \big| \quad \big| f_2 & f_1 \big| \quad \big| f_2 & f_1 \diagdown \diagup f_2 & & \\
 & \mathcal{H}_1 \triangleright \mathcal{H}_2 & \mathcal{H}_1 + \mathcal{H}_2 & \mathcal{H} & \mathcal{H} & \\
1 & f_1 \triangleright f_2 & f_1 + f_2 & [f_1, f_2] & \varnothing_{\mathcal{H}} &
\end{array} \tag{1}$$

We use the notation $k_1, \ldots, k_n \vdash f$ for $(k_1 + \cdots + k_n) \triangleright f$. More generally, we let

$$k_1 \colon \mathcal{K}_1 \to [\![B_1]\!], \ldots, k_n \colon \mathcal{K}_n \to [\![B_n]\!] \vdash f \colon \mathcal{F} \to [\![C]\!] \tag{2}$$

stand for

$$(k_1 + \cdots + k_n) \triangleright f \colon (\mathcal{K}_1 + \cdots + \mathcal{K}_n) \triangleright \mathcal{F} \to [\![\Gamma \vdash C]\!] \tag{3}$$

where Γ is $B_1, \ldots, B_n$. This notation allows us to use lower case Greek letters $\phi, \psi, \ldots$ for multisets of expressions $k_i \colon \mathcal{K}_i \to [\![B_i]\!]$, separated by comma, in the same way as we use upper case Greek letters for multisets of formulas.

Remark 29. Observe that even though $(k_1 + \cdots + k_n) \triangleright f$ is an ICP, i.e. it is a skew fibration and $(\mathcal{K}_1 + \cdots + \mathcal{K}_n) \triangleright \mathcal{F}$ is an arena net, the components $k_i \colon \mathcal{K}_i \to [\![B_i]\!]$ and $f \colon \mathcal{F} \to [\![C]\!]$ do in general not have this property.

$$\frac{}{1\colon \bullet \to [\![P]\!] \;\vdash\; 1\colon \bullet \to [\![P]\!]}\ [\![\mathrm{ax}]\!]^1$$

$$\frac{\varphi \;\vdash\; f\colon \mathcal{F} \to [\![A]\!] \qquad \psi \;\vdash\; g\colon \mathcal{G} \to [\![B]\!]}{\varphi,\, \psi \;\vdash\; f+g\colon \mathcal{F}+\mathcal{G} \to [\![A \wedge B]\!]}\ [\![\wedge \mathsf{R}]\!]$$

$$\frac{\varphi,\, k\colon \mathcal{F} \to [\![A]\!],\, l\colon \mathcal{G} \to [\![B]\!] \;\vdash\; h\colon \mathcal{H} \to [\![C]\!]}{\varphi,\, k+l\colon \mathcal{F}+\mathcal{G} \to [\![A \wedge B]\!] \;\vdash\; h\colon \mathcal{H} \to [\![C]\!]}\ [\![\wedge \mathsf{L}]\!]$$

$$\frac{\varphi,\, k\colon \mathcal{G} \to [\![A]\!] \;\vdash\; f\colon \mathcal{H} \to [\![B]\!]}{\varphi \;\vdash\; k \triangleright f\colon \mathcal{G} \triangleright \mathcal{H} \to [\![A \Rightarrow B]\!]}\ [\![\Rightarrow \mathsf{R}]\!]$$

$$\frac{\varphi \;\vdash\; f\colon \mathcal{F} \to [\![A]\!] \quad k\colon \mathcal{G} \to [\![B]\!],\, \psi \;\vdash\; h\colon \mathcal{H} \to [\![C]\!]}{\varphi,\, f \triangleright k\colon \mathcal{F} \triangleright \mathcal{G} \to [\![A \Rightarrow B]\!],\, \psi \;\vdash\; h\colon \mathcal{H} \to [\![C]\!]}\ [\![\Rightarrow \mathsf{L}]\!]^2$$

$$\frac{\varphi \vdash f\colon \mathcal{F} \to [\![A]\!] \quad \varnothing\colon \varnothing \to [\![B]\!],\, \psi \vdash h\colon \mathcal{H} \to [\![C]\!]}{\varphi_\varnothing,\, \varnothing\colon \varnothing \to [\![A \Rightarrow B]\!],\, \psi \vdash h\colon \mathcal{H} \to [\![C]\!]}\ [\![\Rightarrow \mathsf{L}]\!]^3$$

$$\frac{\varphi,\, k\colon \mathcal{F} \to [\![A]\!],\, l\colon \mathcal{G} \to [\![A]\!] \;\vdash\; f\colon \mathcal{H} \to [\![B]\!]}{\varphi,\, [k,l]\colon \mathcal{F}+\mathcal{G} \to [\![A]\!] \;\vdash\; f\colon \mathcal{H} \to [\![B]\!]}\ [\![\mathsf{c}]\!]$$

$$\frac{\varphi \;\vdash\; f\colon \mathcal{H} \to [\![B]\!]}{\varphi,\, \varnothing_{[\![A]\!]}\colon \varnothing \to [\![A]\!] \;\vdash\; f\colon \mathcal{H} \to [\![B]\!]}\ [\![\mathsf{w}]\!]$$

1) The two vertices in the net form a link
2) $k \neq \varnothing_{[\![B]\!]}$
3) $\varphi \;=\; g_1\colon \mathcal{G}_1 \to [\![D_1]\!], \ldots, g_n\colon \mathcal{G}_n \to [\![D_n]\!]$ and
   $\varphi_\varnothing \;=\; \varnothing\colon \varnothing \to [\![D_1]\!], \ldots, \varnothing\colon \varnothing \to [\![D_n]\!]$

Fig. 6. Interpretation of sequent proofs as ICPs

Definition 30 (Interpretation)**.** Figure 6 defines inductively the ***interpretation*** $[\![\pi]\!]$ of a sequent proof π in LIp (shown in Figure 5) as a sequent of mappings between dags $[\![\pi]\!] = k_1, \ldots, k_n \vdash f$. The component maps $k_1, \ldots, k_n, f$ are called the ***doors*** of $[\![\pi]\!]$.

In Section VI we will show that (i) $[\![\pi]\!]$ is always an ICP, (ii) the size of $[\![\pi]\!]$ is polynomial in the size of π, and (iii) for each ICP g there is a LIp sequent proof π with $[\![\pi]\!] = g$ whose size is polynomial in the size of g. For obtaining these results, we need to first investigate the properties of the doors of $[\![\pi]\!]$, which is the purpose of the next section.

V. Skew fibrations

The purpose of this section is to show for skew fibrations a similar result as for arenas in Theorems 7 and 12. More precisely, we will show that all skew fibrations are generated from 1 and $\varnothing$ via $\triangleright$, $+$, and $[\cdot,\cdot]$, shown in (1), and that each map that is constructed in a certain way from these primitives is indeed a skew fibration. This result is central for all applications of ICPs, it justifies the notation we introduced in the previous section, and it allows us to prove sequentialization.

Recall that a pre-arena is a non-empty L-free and Σ-free dag. A map $f\colon \mathcal{G} \to \mathcal{H}$ is a ***pre-arena morphism*** if it preserves edges and the parity of vertices. Note that $\mathcal{G}$ and $\mathcal{H}$ do not need to be pre-arenas; it suffices that they are stratified. In particular, they can be empty.

For two *distinct* vertices v and w of a pre-arena $\mathcal{G}$ recall that v and w are ***conjuncts***, denoted $v \curlywedge w$, if $\mathsf{m}(v,w)$ is odd. Analogously, define v and w as ***disjuncts***, denoted $v \curlyvee w$, if $\mathsf{m}(v,w)$ is even. We write $v \not\curlywedge w$ (resp. $v \not\curlyvee w$) if v and w are not conjuncts (resp. disjuncts).[2] Note that $v \not\curlyvee w$ iff $v \curlywedge w$ or $v = w$, and $v \not\curlywedge w$ iff $v \curlyvee w$ or $v = w$.

Example 31. In the pre-arena below, since u, v, and w pairwise meet at w, we have $u \curlywedge v \curlywedge w \curlywedge u$, and since y has no meet with x or z, we have $x \curlywedge y \curlywedge z$. Other vertices are disjuncts: for example, u, x and z pairwise meet at z, with even depth, so $u \curlyvee x \curlyvee z$.

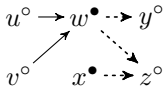

$(u \wedge v \Rightarrow w) \Rightarrow y \wedge (x \Rightarrow z)$

Definition 32. Let $f\colon \mathcal{G} \to \mathcal{H}$ be an pre-arena morphism where $\mathcal{G}$ and $\mathcal{H}$ are both pre-arenas. We say that f is an ***even skew fibration*** if the following two conditions are fulfilled:

- f preserves $\curlywedge$, i.e., $v \curlywedge w$ implies $f(v) \curlywedge f(w)$, and
- if $f(v) \curlywedge w'$ then there exists u with $v \curlywedge u$ and $f(u) \not\curlywedge w'$.

Now let $k\colon \mathcal{G} \to \mathcal{H}$ be a pre-arena morphism where $\mathcal{H}$ is a pre-arena and $\mathcal{G}$ is either a pre-arena or empty. Then k is an ***odd skew fibration*** if we have:

- k preserves $\curlyvee$, i.e., $v \curlyvee w$ implies $k(v) \curlyvee k(w)$, and
- if $k(v) \curlyvee w'$ then there exists u with $v \curlyvee u$ and $k(u) \not\curlyvee w'$.

Example 33. Below are six examples. The left column shows even skew fibrations, the middle column odd skew fibrations, and the examples in the right column are neither even nor odd skew fibrations. A map can be an even and an odd skew fibration at the same time only if it is an identity map.

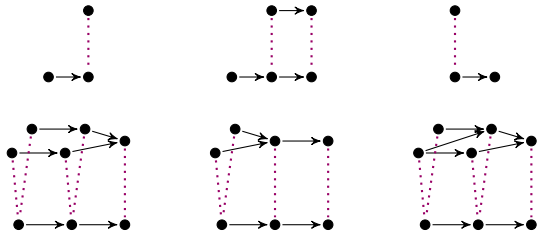

Below we use f and g to denote even skew fibrations, and k and l for odd skew fibrations. We immediately have:

Proposition 34. *Let $f\colon \mathcal{F} \to \mathcal{F}'$ and $g\colon \mathcal{G} \to \mathcal{G}'$ be even skew fibrations, and let $k\colon \mathcal{K} \to \mathcal{K}'$ and $l\colon \mathcal{L} \to \mathcal{L}'$ be odd skew fibrations.*

1) *$1_\mathcal{H}$ is an even and odd skew fibrations for all $\mathcal{H}$.*
2) *$\varnothing_\mathcal{H}$ is an odd skew fibrations for all $\mathcal{H}$.*
3) *$f+g\colon \mathcal{F}+\mathcal{G} \to \mathcal{F}'+\mathcal{G}'$ is an even skew fibration.*
4) *$k+l\colon \mathcal{K}+\mathcal{L} \to \mathcal{K}'+\mathcal{L}'$ is an odd skew fibration.*

[2] The $\curlywedge$-relation corresponds to the R-edges in Rétoré's proof nets [7].

5) *If* $\mathcal{K}' = \mathcal{L}' = \mathcal{H}$ *then* $[k,l]\colon \mathcal{K}+\mathcal{L} \to \mathcal{H}$ *is an odd skew fibration.*
6) $k \triangleright f\colon \mathcal{K} \triangleright \mathcal{F} \to \mathcal{K}' \triangleright \mathcal{F}'$ *is an even skew fibration.*
7) $f \triangleright k\colon \mathcal{F} \triangleright \mathcal{K} \to \mathcal{F}' \triangleright \mathcal{K}'$ *is an odd skew fibration.*

Proof. This follows straighforwardly from the definitions. □

The converse is the main result of this section:

Theorem 35. *Every even skew fibration is of one of the forms:*

$$1 \qquad f+g \qquad k \triangleright f \tag{4}$$

And every odd skew fibration is of one of the following forms:

$$1 \qquad k+l \qquad f \triangleright k \qquad [k,l] \qquad \varnothing_{\mathcal{H}} \tag{5}$$

where f *and* g *are even skew fibrations and* k *and* l *are odd skew fibrations.*

The remainder of this section is dedicated to the proof.

Two vertices v and w in a dag are ***connected*** if they are connected in the underlying undirected graph, i.e., there is a path from v to w if we forget the direction of the edges.

Lemma 36. *Let* $h\colon \mathcal{G} \to \mathcal{H}$ *be a graph morphism. If* v *and* w *are connected in* $\mathcal{G}$*, then* $h(v)$ *and* $h(w)$ *are connected in* $\mathcal{H}$*.*

Lemma 37. *Let* $h\colon \mathcal{G} \to \mathcal{H}$ *be a pre-arena morphism where* $\mathcal{G}$ *is a pre-arena and* $\mathcal{H} = \mathcal{H}'+\mathcal{H}''$ *such that* $\mathcal{H}'$ *and* $\mathcal{H}''$ *are both pre-arenas. Then* $h = h'+h''$ *with* $h'\colon \mathcal{G}' \to \mathcal{H}'$ *and* $h''\colon \mathcal{G}'' \to \mathcal{H}''$ *for some* $\mathcal{G}'$ *and* $\mathcal{G}''$ *with* $\mathcal{G} = \mathcal{G}'+\mathcal{G}''$*. (Note that one of* $\mathcal{G}'$ *and* $\mathcal{G}''$ *can be empty.)*

Proof. Straightforward, making use of Lemma 37 and the fact that $+$ produces disconnected parts of the pre-arena. □

We say that a pre-arena morphism $h\colon \mathcal{G} \to \mathcal{H}$ ***preserves depth*** if for all vertices v in $\mathcal{G}$ we have $\mathsf{d}(v) = \mathsf{d}(h(v))$. We say that h ***preserves roots*** if $\mathsf{d}(v) = 0$ implies $\mathsf{d}(h(v)) = 0$.

Lemma 38. *Even skew fibrations preserve roots.*

Proof. Let $f\colon \mathcal{G} \to \mathcal{H}$ be an even skew fibration and let v be a root of $\mathcal{G}$. For a contradiction, assume $f(v)$ is not a root of $\mathcal{H}$. However, $f(v)$ is even, since v is even. Hence, there is an odd vertex y in $\mathcal{H}$ such that there is a path from $f(v)$ to y and from y to a root z of $\mathcal{H}$. We have $f(v) \curlywedge y$ in $\mathcal{H}$. Hence, there is a vertex u in $\mathcal{G}$, such that $v \curlywedge u$ (and therefore also $f(v) \curlywedge f(u)$) and $f(u) \not\curlywedge y$. If $f(u)$ and $f(v)$ are not connected, then $f(u)$ and y are not connected, and hence $f(u) \curlywedge y$. Therefore $f(u)$ and $f(v)$ must be connected. This means that there must be a path from $f(u)$ to the root z (because of Lemmas 10 and 11). Then the two paths from $f(u)$ to z and from $f(v)$ to z must meet at a vertex x at odd depth. Consequently $f(u)$ and y either meet at x or at y. Since both are at odd depth we have $f(u) \curlywedge y$, a contradiction. □

Lemma 39. *Odd skew fibrations preserve roots.*

Proof. Similar to the proof of Lemma 38 but simpler. □

Lemma 40. *Even and odd skew fibrations preserve depth.*

Proof. Via Lemma 38 and 39, and induction on the depth. □

Lemma 41. *Let* $\mathcal{G} = \mathcal{G}_1 \triangleright \mathcal{G}_2$ *and* $\mathcal{H} = \mathcal{H}' \triangleright \mathcal{H}''$ *where* $\mathcal{G}_1$*,* $\mathcal{G}_2$*,* $\mathcal{H}'$*, and* $\mathcal{H}''$ *are arenas, and let* $h\colon \mathcal{G} \to \mathcal{H}$ *be an even or odd skew fibration. Then* $h = h' \triangleright h''$ *with* $h'\colon \mathcal{G}' \to \mathcal{H}'$ *and* $h''\colon \mathcal{G}'' \to \mathcal{H}''$ *for some* $\mathcal{G}'$ *and* $\mathcal{G}''$ *with* $\mathcal{G} = \mathcal{G}' \triangleright \mathcal{G}''$*. (Note that* $\mathcal{G}'$ *but not* $\mathcal{G}''$ *can be empty.)*

Proof. This is similar to Lemma 37, but makes additional crucial use of Lemma 40. □

Proof of Theorem 35. The proof proceeds by case analysis. We show here only the case where $h\colon \mathcal{G} \to \mathcal{H}$ is an even skew fibration and $\mathcal{H} = \mathcal{H}' \triangleright \mathcal{H}''$ for some $\mathcal{H}'$ and $\mathcal{H}''$. We define two vertex sets $V' = V(\mathcal{G}) \cap V(\mathcal{H}')$ and $V'' = V(\mathcal{G}) \cap V(\mathcal{H}'')$. Then we have that $V'' \neq \varnothing$ because h preserves roots. If $V' = \varnothing$ then $h''\colon \mathcal{G} \to \mathcal{H}''$ is an even skew fibration and we have $h = \varnothing_{\mathcal{H}'} \triangleright h''$. Now assume $V' \neq \varnothing$. Then $\mathcal{G}$ cannot be a single vertex and can also not be of shape $\mathcal{G}_1+\mathcal{G}_2$ because then it would not preserve $\curlywedge$. Hence, $\mathcal{G} = \mathcal{G}_1 \triangleright \mathcal{G}_2$, such that we can apply Lemma 41 and get $h = h' \triangleright h''$ with $h'\colon \mathcal{G}' \to \mathcal{H}'$ and $h''\colon \mathcal{G}'' \to \mathcal{H}''$ for some $\mathcal{G}'$ and $\mathcal{G}''$ with $\mathcal{G} = \mathcal{G}' \triangleright \mathcal{G}''$. Since h preserves depth and $\curlywedge$, so do h' and h''. Now let $v \in V(\mathcal{G}'')$ and $w \in V(\mathcal{H}'')$ such that $h(v) \curlywedge w$. Then there is a $u \in V(\mathcal{G})$ such that $u \curlywedge v$ and $h(u) \not\curlywedge w$. Hence $u \in V(\mathcal{G}'')$ because otherwise $u \curlyvee v$. Hence, h'' is an even skew fibration. The same argument applies to h', but we conclude that h' is an odd skew fibration since the parity of all vertices is inverted. All other cases are similar. □

Observe that a map on arenas is a skew fibration (Definition 26) if and only if it is an even skew fibration (Definition 32) that preserves the equivalence relation. It follows from Theorem 35 that if the expression in (2) is an ICP, then $k_1, \ldots, k_n$ are odd skew fibrations, and f is an even skew fibration. This is crucial for sequentialization.

VI. Polynomial Full Completeness

In this section we will define ***sequentialization***, the translation from ICPs to sequent proofs. We establish that both sequentialization and de-sequentialization are size-preserving (up to a polynomial): our ***polynomial full completeness*** result.

For convenience, we use $f :: A$ to abbreviate $f\colon \mathcal{G} \to [\![A]\!]$, and similarly $\varphi :: \Gamma$ for a sequence φ of such maps.

We describe sequentialization as a rewrite relation, where intermediate stages of the translation process are captured as sequent proofs containing (still-to-be-translated) ICPs as subproofs. We formalize this by adding a special axiom (icp) to the sequent calculus, given below, which incorporates an ICP $\varphi :: \Gamma \vdash f :: A$ (depicted in a box for easy recognition) as a sub-proof of $\Gamma \vdash A$. Such a proof will be a ***hybrid proof***.

$$\frac{\boxed{\varphi :: \Gamma \vdash f :: A}}{\Gamma \vdash A}\ (\mathsf{icp})$$

Definition 42 (Hybrid proofs)**.** The calculus of ***hybrid proofs*** is the sequent calculus LIp extended with the icp axiom.

Note that embedding the ICP $f :: A$ into a hybrid proof forces an interpretation of the arena $[\![A]\!]$ as a fixed formula A, which may be chosen up to isomorphism. We define sequentialization

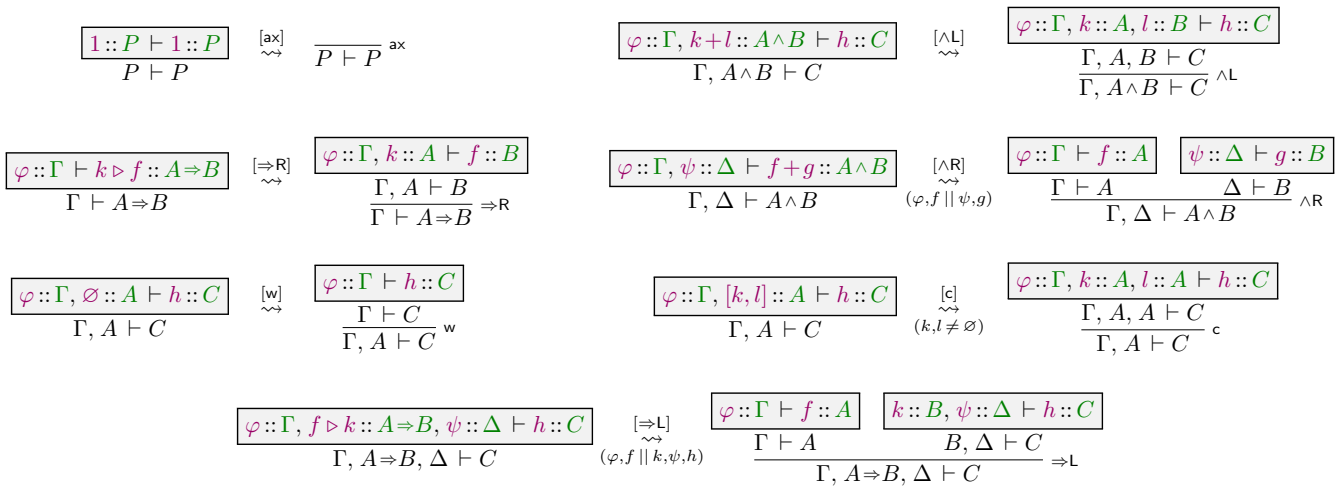

Fig. 7. Sequentialization

via a rewrite relation $(\rightsquigarrow)$ on hybrid proofs, given in Figure 7, which interprets constructions on ICPs as proof rules. The side conditions for $[\Rightarrow\mathsf{L}]$ and $[\wedge\mathsf{R}]$, which split an ICP in two, ensure that both parts are ICPs. They use the following notion:

Definition 43. Two sequences θ, ψ of doors of an ICP $\varphi \vdash f$ are ***separate***, $\theta \parallel \psi$, if their covers do not share a partition.

Definition 44 (Sequentialization)**.** The ***sequentialization*** rewrite relation $(\rightsquigarrow)$ on hybrid proofs is given by the rules in Figure 7. An ICP ***sequentializes*** $(\rightsquigarrow\!\!\!\!\shortmid)$ by:

$$f :: A \rightsquigarrow\!\!\!\!\shortmid \pi \quad \iff \quad \frac{\vdash\ f :: A}{\vdash A} \ \rightsquigarrow^*\ \pi\ .$$

Sequentialization and de-sequentialization (Figure 6, Definition 30) are near-inverses. In one direction they cancel out:

Theorem 45. *If $f \rightsquigarrow\!\!\!\!\shortmid \pi$ then $[\![\pi]\!] = f$.*

Proof. By inspection of both translations. □

In the other direction, from sequent proof to ICP to sequent proof, note:

- The rule $[\Rightarrow\mathsf{L}]$ is a counterpart only to the first version of $[\![\Rightarrow\mathsf{L}]\!]$ in Figure 6, where the consequent k of the introduced implication $f \triangleright k :: A \Rightarrow B$ is non-empty.
- The rule $[\mathsf{c}]$ requires that both components k, l of $[k, l]$ are non-empty, to avoid non-termination via $k = [k, \varnothing]$.

To prove that ICPs sequentialize, we show that $(\rightsquigarrow)$ reduces any ICP in a hybrid proof, so that its normal forms are standard LIp proofs. We use a result that corresponds to the "splitting tensor lemma" in the proof net literature [11], [12].

For an ICP $\varphi :: \Gamma \vdash f :: A$, call the doors in φ ***input*** doors, and f an ***output*** door. An input door is ***open*** if it is:

- a ***weakening*** $\varnothing :: A$,
- a ***contraction*** $[k, l] :: A$ where $k, l \neq \varnothing$,
- a ***conjunction*** $k + l :: A \wedge B$, or
- an ***implication*** $f \triangleright k :: A \Rightarrow B$ that is ***splitting***: it occurs as

$$\varphi :: \Gamma,\ f \triangleright k :: A \Rightarrow B,\ \psi :: \Delta \vdash h :: C \quad \text{where} \quad \varphi, f \ \parallel\ k, \psi, h\ .$$

An output door is ***open*** if it is:

- an ***implication*** $k \triangleright f :: A \Rightarrow B$, or
- a ***conjunction*** $f + g :: A \wedge B$ that is ***splitting***: it occurs as

$$\varphi :: \Gamma,\ \psi :: \Delta \vdash f + g :: A \wedge B \quad \text{where} \quad \varphi, f \ \parallel\ \psi, g\ .$$

By the correspondence with MLL, we call an input implication and output conjunction a ***tensor***, and an input conjunction and output implication a ***par***. A door that is not open is ***closed***. By the inductive characterization of even and odd skew fibrations (Theorem 35), a closed door is either ***atomic*** $1 :: P$, or a tensor that is not splitting. Observe that a sequentialization step applies to every ICP with an open door.

Lemma 46 (Splitting tensor)**.** *An ICP is either an axiom or has an open door.*

Proof. An ICP that contains a weakening, contraction, or par, has an open door. Otherwise, it consists of atomic doors $1 : \bullet \to [\![P]\!]$ and tensors: input $f \triangleright k : \mathcal{F} \triangleright \mathcal{K} \to [\![A \Rightarrow B]\!]$ and output $f + g : \mathcal{F} + \mathcal{G} \to [\![A \wedge B]\!]$. The (top-level) structure of the ICP is then mirrored in its arena net. Since the latter is IMLL-constructed (by Theorem 25), the ICP is either an axiom, or has a splitting tensor. □

Theorem 47 (Sequentialization)**.** *For every ICP $f :: A$ there is a sequent proof π such that $f :: A \rightsquigarrow\!\!\!\!\shortmid \pi$.*

Proof. By the splitting tensor lemma (Lemma 46), to any ICP in a hybrid proof a sequentialization step applies. By arena net sequentialization (Theorem 25) and the inductive characterization of even and odd skew fibrations (Theorem 35), the result of a sequentialization step is again a hybrid proof. □

With Theorem 45, de-sequentialization is surjective:

Corollary 48 (Surjectivity)**.** *Every ICP is the image $[\![\pi]\!]$ of a sequent proof π.*

Our main result is that, in addition to the above, both translations $[\![-]\!]$ and $(\rightsquigarrow\!\!\!\!\shortmid)$ are polynomial. Intuitively, a sequent

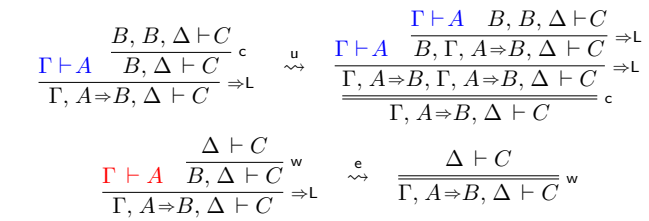

Fig. 8. Unfolding and excising

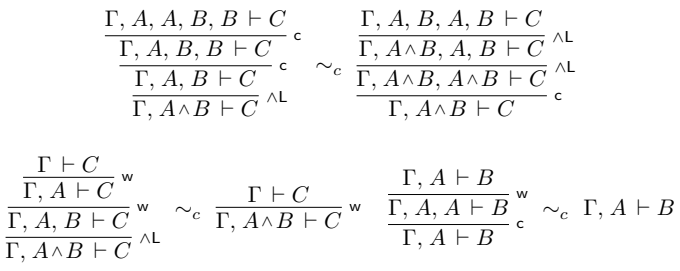

Fig. 9. Comonoid transformations

proof and an ICP have comparable size because the *axioms* of a sequent proof π make up the *partitions* of the arena net of $[\![\pi]\!]$ (unless they are deleted by a $[\![\Rightarrow\mathsf{L}]\!]$ translation step), while the conclusion of π directly gives the base arena of $[\![\pi]\!]$.

Formally, we define the size of sequent proofs and of ICPs in the standard way. The size $|\Gamma \vdash A|$ of a sequent $\Gamma \vdash A$ is the number of symbols in it, and the size $|\pi|$ of a proof π is the sum of the sizes of the sequents in it. The size $|\mathcal{G}|$ of an arena $\mathcal{G}$ is the number of its vertices. Finally, the size $|g|$ of an ICP $g\colon \mathcal{G} \to [\![A]\!]$ is $|\mathcal{G}| + |[\![A]\!]|$. By inspecting Figures 6 and 7, we immediately have the following:

Theorem 49 (Polynomial full completeness)**.** *There are polynomials p and q such that:*

- $|[\![\pi]\!]| \leq p(|\pi|)$ *for every sequent proof π;*
- *if* $f \rightsquigarrow\!\!\!\!\circ\, \pi$ *then* $|\pi| \leq q(|f|)$.

VII. Local canonicity

The sequent calculus admits many proof transformations between proofs that, morally or semantically, denote the same object. We will consider these from the perspective of complexity, and *reject* those transformations that incur exponential growth. We factor out the remaining, computationally feasible transformations: our ***local canonicity*** result.

We will distinguish several kinds of transformations, and identify whether they are ***local***, i.e. do not affect the rest of the proof. First, the interaction of $\Rightarrow\mathsf{L}$ with contraction and weakening gives rise to two non-local (but semantically sound) transformations we call ***unfolding*** (u) and ***excising*** (e), which *duplicate* respectively *delete* subproofs, in Figure 8.

Iterated unfolding induces exponential growth, despite the fact that we are working with cut-free proofs, and we reject it on complexity grounds. Figure 1 in the Introduction illustrates how two sequent proofs that differ by this rewrite step give rise to distinct ICPs. Note also that dag-arenas are crucial here. An interpretation of the left-hand side would be of the form:

$$\dfrac{\varphi \vdash f\colon \mathcal{F} \to [\![A]\!] \qquad \dfrac{k\colon \mathcal{G} \to [\![B]\!],\, l\colon \mathcal{H} \to [\![B]\!],\, \psi \vdash h}{[k,l]\colon \mathcal{G}+\mathcal{H} \to [\![B]\!],\, \psi \vdash h}\,[\![\mathsf{c}]\!]}{\varphi,\, f \triangleright [k,l]\colon \mathcal{F} \triangleright (\mathcal{G}+\mathcal{H}) \to [\![A \Rightarrow B]\!],\, \psi \vdash h}\,[\![\Rightarrow\mathsf{L}]\!]$$

Tree-arenas must express $\mathcal{F} \triangleright (\mathcal{G}+\mathcal{H})$ as $(\mathcal{F} \triangleright \mathcal{G}) + (\mathcal{F} \triangleright \mathcal{H})$, however, forcing the application of unfolding.

For excising, the weakened formula B would be interpreted as $\varnothing\colon \varnothing \to [\![B]\!]$. Since we cannot admit the construction $\mathcal{F} \triangleright \varnothing$ (as it equates to $\mathcal{F}$), we cannot interpret the left-hand side. Instead, the interpretation of a sequent proof as an ICP, in Figure 6, applies excising by way of the second $[\![\Rightarrow\mathsf{L}]\!]$ rule: reducing ϕ to the empty ϕ_0 in that rule corresponds to deleting the subproof of $\Gamma \vdash A$ in the excising rule of Figure 8.

Define a ***permutation*** $\sim_p$ as the exchange of two consecutive inferences on *distinct* formulas in a sequent; e.g.:

$$\dfrac{\Gamma \vdash A \quad \dfrac{B, \Delta, C \vdash D}{B, \Delta \vdash C \Rightarrow D}\Rightarrow\mathsf{R}}{\Gamma, A \Rightarrow B, \Delta \vdash C \Rightarrow D}\Rightarrow\mathsf{L} \quad \sim_p \quad \dfrac{\dfrac{\Gamma \vdash A \quad B, \Delta, C \vdash D}{\Gamma, A \Rightarrow B, \Delta, C \vdash D}\Rightarrow\mathsf{L}}{\Gamma, A \Rightarrow B, \Delta \vdash C \Rightarrow D}\Rightarrow\mathsf{R}$$

Each permutation is a simple, local exchange of independent inferences, since we use a *context-splitting* (*multiplicative*) and not *context-sharing* (*additive*) formulation of the sequent rules.

Transformations of two rules applied to the *same* formula in a sequent involve a contraction or weakening (since no two logical rules apply to the same formula) and another left-rule: $\Rightarrow\mathsf{L}$, $\wedge\mathsf{L}$, c, or w. The interaction with $\Rightarrow\mathsf{L}$ gave rise to *unfolding* and *excising* above. The remaining ***comonoid transformations*** ($\sim_c$) are given in Figure 9. These are *local*, and correspond to the following equalities:

$$[k,l] + [m,n] = [k+m, l+n] \qquad \varnothing + \varnothing = \varnothing \qquad [k, \varnothing] = k$$

Finally, the formula isomorphisms ($\cong$), applied to the conclusion of a proof, induce a further set of proof transformations, the ***formula-isomorphism transformations***. For example:

$$\dfrac{\dfrac{\Gamma \vdash A \quad \Delta \vdash B}{\Gamma, \Delta \vdash A \wedge B}\wedge\mathsf{R} \quad C, \Lambda \vdash D}{\Gamma, \Delta, (A \wedge B) \Rightarrow C, \Lambda \vdash D}\Rightarrow\mathsf{L} \quad \cong \quad \dfrac{\Gamma \vdash A \quad \dfrac{\Delta \vdash B \quad C, \Lambda \vdash D}{\Delta, B \Rightarrow C, \Lambda \vdash D}\Rightarrow\mathsf{L}}{\Gamma, \Delta, A \Rightarrow B \Rightarrow C, \Lambda \vdash D}\Rightarrow\mathsf{L}$$

Definition 50 (Proof equivalence)**.** Proof transformations generate the following equivalence relations on proofs:

Local rule-equivalence ($\sim$) is generated by permutations ($\sim_p$) and comonoid transformations ($\sim_c$).

Local equivalence ($\cong$) is generated by local rule-equivalence ($\sim$) and formula-isomorphism transformations.

The following theorems (i) complete the relation between sequentialization and de-sequentialization; (ii) observe that any two sequentializations of the same ICP are locally equivalent; and (iii) formalize local canonicity. Proofs are omitted.

Theorem 51. *If* $[\![\pi]\!] \rightsquigarrow\!\!\!\!\circ\, \pi'$ *then* $\pi\ (\sim \cup \overset{\mathsf{e}}{\rightsquigarrow})^*\ \pi'$.

Theorem 52. *If* $f \rightsquigarrow\!\!\!\!\circ\, \pi$ *and* $f \rightsquigarrow\!\!\!\!\circ\, \rho$ *then* $\pi \cong \rho$.

Theorem 53 (Local canonicity)**.** *ICPs canonically model sequent proofs modulo local equivalence and excising:*

$$[\![\pi]\!] = [\![\rho]\!] \iff \pi\ (\overset{\mathsf{e}}{\rightsquigarrow} \cup \cong \cup \overset{\mathsf{e}}{\leftsquigarrow})^*\ \rho\ .$$

VIII. Relation with innocent strategies

We present a very natural surjection from ICPs to the Hyland-Ong/Nickau arena innocent strategies [4], [6] corresponding to λ-calculus terms. We have the following relationships:

$$\text{ICPs} \xrightarrow[\text{surjection}]{} \begin{matrix}\text{winning}\\ \text{innocent}\\ \text{strategies}\end{matrix} \xleftrightarrow[\text{bijection}]{} \begin{matrix}\eta\text{-expanded}\\ \beta\text{-normal}\\ \lambda\text{ terms}\end{matrix}$$

The latter bijection is game semantics folklore; one formalization can be found in [13] (which extends it to second-order). We reformulate the bijection over our partitioned arenas below.

A sequence is ***justified*** [4] if every element but the first has a pointer to an earlier element, its ***justifier***, skipping an even number of elements. An element preceded by an even (resp. odd) number of others is an O-***move*** (resp. P-***move***). A ***play*** in an arena $\mathcal{G}$ is a justified sequence over $V(\mathcal{G})$ such that the first move (if present) is a root in $\mathcal{G}$, and a v-occurrence justifies a w-occurrence only if $v \leftarrow_{\mathcal{G}} w$. A play is O-***shortsighted*** if every pointer from an O-move skips no moves, and P-***uniform***[3] if whenever v is the vertex of a P-move and w is the vertex of the preceding O-move, then $w \sim_{\mathcal{G}} v$. A (P-)***view*** is a play which is O-shortsighted and P-uniform.

The ***predecessor*** of a non-empty play is the result of deleting the final move (and its pointer); the converse is the ***successor*** relation. A ***winning innocent strategy*** in an arena $\mathcal{G}$ is a finite predecessor-closed set σ of views in $\mathcal{G}$ such that for every view p in σ, if p has even (resp. odd) length, then every (resp. exactly one) successor of p is in σ.[4]

Theorem 54. *Winning innocent strategies in a formula A are in bijection with η-expanded, β-normal λ-terms of type A.*

Proof. A special case of the main theorem of [13, §7] (omitting quantifiers). That theorem was in terms of atom-labelled arenas which were forests; the inductive proof (itself a variant of the definability proof in [4]) goes through with partitions instead of atom labels, and with dags instead of forests. □

An ***abstract view*** in a arena net $\mathcal{G}$ is a (possibly-empty) reverse path in its link graph $\widehat{\mathcal{G}}$, starting from the root.

Proposition 55. *Let $f : \mathcal{G} \to \mathcal{H}$ be an ICP. The image under f of an abstract view in $\mathcal{G}$ is a view in $\mathcal{H}$.*

Proof. The *Functionality* condition of arena nets (Def. 14) provides a target vertex for each justification pointer, which is unique by the *Acyclicity* condition. P-uniformity holds because every equivalence class in an arena net comprises an even and an odd vertex, and skew fibrations preserve equivalence. □

Theorem 56. *Let $f : \mathcal{G} \to \mathcal{H}$ be an ICP. The image under f of all abstract views in $\mathcal{G}$ is a winning innocent strategy in $\mathcal{H}$.*

Proof. Let σ be the image under f of the abstract views in $\mathcal{G}$. Let p be a view in σ ending with the vertex $v = f(\widehat{v})$. If p has even length, by the ICP skew lifting condition, the neighbourhood $\{w : v \leftarrow_{\mathcal{H}} w\}$ in $\mathcal{H}$ is in bijection with the neighbourhood $\{u : \widehat{v} \leftarrow_{\mathcal{G}} u\}$ in $\mathcal{G}$. Thus every possible successor of p is in σ. If p has odd length, $\widehat{v}$ is an even vertex in $\mathcal{G}$, hence an abstract view ending in $\widehat{v}$ has only one possible extending vertex, the mate $\widehat{v}'$ of $\widehat{v}$ (the opposite vertex in the link of $\widehat{v}$). Thus there is a unique successor of p in σ. □

Corollary 57. *For every formula A, there is a surjection from ICPs of A to winning innocent strategies in A, hence (via Theorem 54) to η-expanded, β-normal λ-calculus terms on A.*

IX. Conclusion and related work

We have constructed intuitionistic combinatorial proofs, as a concrete, geometric semantics of intuitionistic logic with good properties: *local canonicity* (factoring out non-duplicating conversions), as pioneered for MALL proof nets in [15], and *polynomial full completeness*, a complexity-aware version of full completeness. As a (locally) canonical semantics, ICPs naturally link to many fundamental ideas in the literature.

ICPs are a close relative to *proof nets* [11], [3], and *arena nets* can be seen as the IMLL version of the cograph-based proof nets of Rétoré [7].

The stratification of a proof into separate *logical* and *structural* (contraction/weakening) phases is analogous to Herbrand's Theorem [16] (as formulated for prenex forms in Gentzen's sharpened Hauptsatz or midsequent theorem [17]). While impossible in sequent calculus [18], *deep-inference* proof theory does allow such stratification [19], [20], [5], [21], [22], as well as the embedding of classical skew fibrations [1], [2] in deep-inference proofs.

In *game semantics* [4], [14] skew fibrations are mostly implicit as the repeating of moves in an arena, though Boudes' construction of *thick subtree* [23] is effectively a skew fibration on trees, defined inductively. Many games models use tree-arenas, with the notable exceptions of McCusker's for FPC [8] and *event structures* (see e.g. [24], [25]). Yet our graph-based *arena nets* are a direct and natural interpretation of IMLL proof nets, and the sharing they enable is the crucial ingredient for our *polynomial full completeness* result.

Current work extends intuitionistic combinatorial proofs with normalization (cut elimination), and extends them to first order (see [26] for the classical case).

[3] This is the ***copycat condition*** of [13], reformulated in a partioned arena. The condition was called *token-reflecting* in [14].

[4] For technical convenience and brevity we have formulated innocent strategies directly as sets of views, and assume P-uniformity in all views.

References

[1] D. J. D. Hughes, "Proofs Without Syntax," *Annals of Mathematics*, vol. 164, no. 3, pp. 1065–1076, 2006.

[2] ——, "Towards Hilbert's 24th problem: Combinatorial proof invariants: (preliminary version)," *ENTCS*, vol. 165, pp. 37–63, 2006.

[3] F. Lamarche, "Proof Nets for Intuitionistic Linear Logic: Essential Nets," Inria, Research Report, 2008. [Online]. Available: https://hal.inria.fr/inria-00347336

[4] J. M. E. Hyland and C. L. Ong, "On full abstraction for PCF: I, II, and III," *Information and Computation*, vol. 163, no. 2, pp. 285–408, 2000.

[5] L. Straßburger, "A characterisation of medial as rewriting rule," in *RTA'07*, ser. LNCS, F. Baader, Ed., vol. 4533. Springer, 2007, pp. 344–358.

[6] H. Nickau, "Hereditarily sequential functionals," in *International Symposium on Logical Foundations of Computer Science*, 1994, pp. 253–264.

[7] C. Retoré, "Handsome proof-nets: perfect matchings and cographs," *Theoretical Computer Science*, vol. 294, no. 3, pp. 473–488, 2003.

[8] G. McCusker, "Games and full abstraction for FPC," *Information and Computation*, vol. 160, no. 1–2, pp. 1–61, 2000.

[9] D. G. Corneil, H. Lerchs, and L. K. Stewart-Burlingham, "Complement reducible graphs," *Disc. Appl. Math.*, 1981.

[10] A. Das and L. Straßburger, "On linear rewriting systems for Boolean logic and some applications to proof theory," *Logical Methods in Computer Science*, vol. Volume 12, Issue 4, 2017.

[11] J.-Y. Girard, "Linear logic," *Theor. Comp. Science*, vol. 50, pp. 1–102, 1987.

[12] V. Danos and L. Regnier, "The structure of multiplicatives," *Arch. Math. Log.*, vol. 28, no. 3, pp. 181–203, 1989.

[13] D. J. D. Hughes, "Games and definability for system F," in *12th Annual IEEE Symposium on Logic in Computer Science (LICS)*, 1997, pp. 76–86.

[14] A. S. Murawski and C.-H. L. Ong, "Evolving games and essential nets for affine polymorphism," in *International Conference on Typed Lambda Calculi and Applications*. Springer, 2001, pp. 360–375.

[15] D. J. D. Hughes and W. Heijltjes, "Conflict nets: Efficient locally canonical MALL proof nets," in *LICS'16*, M. Grohe, E. Koskinen, and N. Shankar, Eds. ACM, 2016, pp. 437–446.

[16] J. Herbrand, "Investigations in proof theory: The properties of true propositions," in *From Frege to Gödel: A source book in mathematical logic, 1879–1931*, J. van Heijenoort, Ed. Harvard University Press, 1967, pp. 525–581.

[17] G. Gentzen, "Untersuchungen über das logische Schließen I & II," *Mathematische Zeitschrift*, vol. 39, pp. 176–210 405–431, 1935.

[18] K. Brünnler, "Two restrictions on contraction," *Logic Journal of the Interest Group in Pure and Applied Logics*, vol. 11, no. 5, pp. 525—529, 2003.

[19] K. Brünnler and A. Tiu, "A local system for classical logic," in *8th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, ser. LNCS, vol. 2250, 2001, pp. 347–361.

[20] L. Straßburger, "MELL in the calculus of structures," *Theor. Comput. Sci.*, vol. 309, pp. 213–285, 2003.

[21] A. Aler Tubella, A. Guglielmi, and B. Ralph, "Removing cycles from proofs," in *26th EACSL Annual Conference on Computer Science Logic (CSL)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 82, 2017, pp. 9:1–17.

[22] L. Straßburger, "Combinatorial flows and their normalisation," in *2nd Int. Conf. Formal Structures for Computation and Deduction (FSCD)*, 2017.

[23] P. Boudes, "Thick subtrees, games and experiments," in *Typed Lambda Calculi and Applications (TLCA)*, 2009.

[24] P.-A. Melliès, "Asynchronous games 2: The true concurrency of innocence," in *Selected Papers of CONCUR 2004*, ser. Theoretical Computer Science, vol. 358, 2006, pp. 200–228.

[25] S. Castellan, P. Clairambault, and G. Winskel, "Symmetry in concurrent games," in *CSL-LICS'14*, 2014, p. 28.

[26] D. J. D. Hughes, "First-order proofs without syntax," 2019, available at arxiv.org.