First-order proofs without syntax

Dominic J. D. Hughes

Stanford University & U.C. Berkeley*

4 June 2019

Proofs are traditionally syntactic, inductively generated objects. This paper reformulates first-order logic (predicate calculus) with proofs which are graph-theoretic rather than syntactic. It defines a *combinatorial proof* of a formula φ as a lax fibration over a graph associated with φ. The main theorem is soundness and completeness: a formula is a valid if and only if it has a combinatorial proof.

1 Introduction

Proofs are traditionally syntactic, inductively generated objects. For example, Fig. 1 shows a syntactic proof of $\exists x\,(px \Rightarrow \forall y\,py)$. This paper reformulates first-order logic (predicate calculus) [Fre79] with proofs which are graph-theoretic rather than syntactic. It defines a *combinatorial proof* of a formula φ as a lax graph fibration $f : K \to \mathcal{G}(\varphi)$ over a graph $\mathcal{G}(\varphi)$ associated with φ, where K is a partially coloured graph. For example, if $\varphi = \exists x\,(px \Rightarrow \forall y\,py)$ then $\mathcal{G}(\varphi)$ is

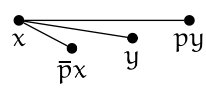

and a combinatorial proof $f : K \to \mathcal{G}(\varphi)$ of φ is

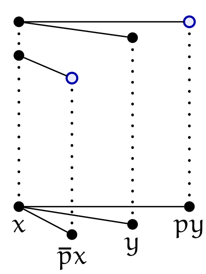

The upper graph is K (two coloured vertices ∘∘ and three uncoloured vertices), the lower graph is $\mathcal{G}(\varphi)$, and the dotted lines define f. Additional combinatorial proofs are depicted in Fig. 2. The combinatorial proof $f : K \to \mathcal{G}(\varphi)$ above can be condensed by leaving $\mathcal{G}(\varphi)$ implicit and drawing K over the formula φ:

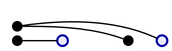

$$\exists x\,(\,px \Rightarrow \forall y\,py\,)$$

The reader may contrast this with the syntactic proof of the same formula in Fig. 1. The four combinatorial proofs of Fig. 2 are rendered in condensed form in Fig. 3.

The main theorem of this paper is soundness and completeness: a formula is valid if and only if it has a combinatorial proof (Theorem 6.4). The propositional fragment was presented in [Hug06a].

*I conducted this research as a Visiting Scholar at Stanford then Berkeley. Many thanks to my hosts, Vaughan Pratt (Stanford Computer Science), Sol Feferman (Mathematics) and Wes Holliday (Berkeley Logic Group). I am extremely grateful for very helpful feedback from Willem Heijltjes, Lutz Straßburger, Grisha Mints, Sam Buss, Martin Hyland, Marc Bagnol and Nil Demirçubuk. In memoriam Sol and Grisha.

$$
\begin{array}{cl}
\hline
py \vdash py & \\
\hline
py \vdash py,\ \forall y\, py & \text{weaken} \\
\hline
\vdash py,\ py \Rightarrow \forall y\, py & \text{implies} \\
\hline
\vdash py,\ \exists x(px \Rightarrow \forall y\, py) & \text{there exists} \\
\hline
\vdash \forall y\, py,\ \exists x(px \Rightarrow \forall y\, py) & \text{for all} \\
\hline
py \vdash \forall y\, py,\ \exists x(px \Rightarrow \forall y\, py) & \text{weaken} \\
\hline
\vdash py \Rightarrow \forall y\, py,\ \exists x(px \Rightarrow \forall y\, py) & \text{implies} \\
\hline
\vdash \exists x(px \Rightarrow \forall y\, py),\ \exists x(px \Rightarrow \forall y\, py) & \text{there exists} \\
\hline
\vdash \exists x(px \Rightarrow \forall y\, py) & \text{contract}
\end{array}
$$

Figure 1. A syntactic proof of $\exists x(px \Rightarrow \forall y\, py)$, in Gentzen's sequent calculus LK [Gen35].

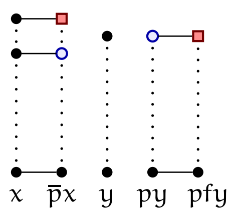

$(\forall x\, px) \Rightarrow \forall y\,(py \wedge pfy)$

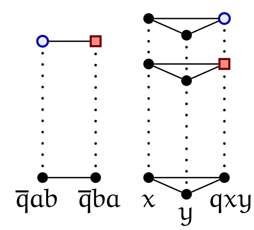

$qab \vee qba \Rightarrow \exists x\, \exists y\, qxy$

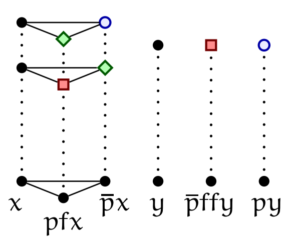

$\big(\forall x(pfx \Rightarrow px)\big) \Rightarrow \forall y\,(pffy \Rightarrow py)$

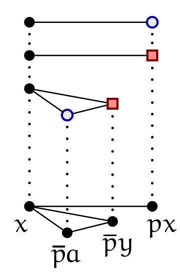

$\exists x\,(pa \vee py \Rightarrow px)$

Figure 2. Four combinatorial proofs, each shown above the formula proved. Here x and y are variables, f is a unary function symbol, a and b are constants (nullary function symbols), p is a unary predicate symbol, and q is a binary predicate symbol.

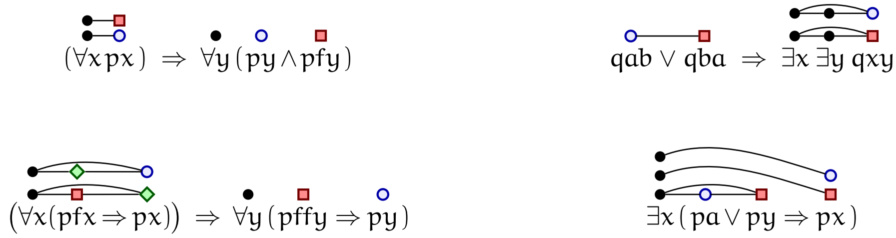

Figure 3. Condensed forms of the four combinatorial proofs in Fig. 2.

2 Notation and terminology

First-order logic. We mostly follow the notation and terminology of [Joh87] for first-order logic without equality [Fre79]. Terms and atoms (atomic formulas) are generated inductively from variables $x, y, z, \ldots$ by: if γ is an n-ary function (resp. predicate) symbol and $t_1, \ldots, t_n$ are terms then $\gamma t_1 \ldots t_n$ is a term (resp. atom). For technical convenience we assume every predicate symbol p is assigned a ***dual*** predicate symbol $\bar{p}$ with $\bar{p} \neq p$ and $\bar{\bar{p}} = p$, and extend duality to atoms with $\overline{pt_1 \ldots t_n} = \bar{p}t_1 \ldots t_n$. ***Formulas*** are generated from atoms by binary $\wedge$ and $\vee$ and quantifiers $\forall x$ and $\exists x$ per variable x. Define $\neg$ and $\Rightarrow$ as abbreviations: $\neg(\alpha) = \bar{\alpha}$ on atoms α, $\neg(\varphi \wedge \theta) = (\neg\varphi) \vee (\neg\varphi)$, $\neg(\varphi \vee \theta) = (\neg\varphi) \wedge (\neg\theta)$, $\neg\forall x\, \varphi = \exists x \neg\varphi$, $\neg\exists x\, \varphi = \forall x \neg\varphi$, and $\varphi \Rightarrow \theta = (\neg\varphi) \vee \theta$. A formula is ***rectified*** if all bound variables are distinct from one another and from all free variables, *e.g.* $(px \vee \exists y\, qy) \wedge \exists z\, rz$ but not $(px \vee \exists x\, qx) \wedge \exists x\, rx$. We assume all formulas are rectified (losing no generality since every unrectified formula has a logically equivalent rectified form).

Graphs. An ***edge*** on a set V is a two-element subset of V. A ***graph*** (V, E) is a finite set V of ***vertices*** and a set E of edges on V. Write V_G and E_G for the vertex and edge sets of a graph G, and vw for $\{v,w\}$. The ***complement*** of (V, E) is the graph (V, E^c) with $vw \in E^c$ if and only if $vw \notin E$. A graph G is (partially) ***coloured*** it carries a partial equivalence relation $\sim$ on V_G such that $v \sim w$ only if $vw \notin E_G$; each equivalence class is a ***colour***. A graph is ***labelled*** in a set L if each vertex has an element of L associated with it, its ***label***. A ***vertex renaming*** of (V, E) along a bijection $(\hat{\ }) : V \to V'$ is the graph $(V', \{\hat{v}\hat{w} : vw \in E\})$, with colouring or labelling inherited (*i.e.*, $\hat{v} \sim \hat{w}$ if $v \sim w$, and the label of $\hat{v}$ that of v). Following standard graph theory, we identify graphs modulo vertex renaming. Let $G = (V, E)$ and $G' = (V', E')$ be graphs. A ***homomorphism*** $h : G \to G'$ is a function $h : V \to V'$ such that if $vw \in E$ then $h(v)h(w) \in E'$. Without loss of generality, assume $V \cap V' = \emptyset$ (by renaming vertices if needed). The ***union*** $G + G'$ is $(V \cup V', E \cup E')$ and ***join*** $G \times G'$ is $(V \cup V', E \cup E' \cup \{vv' : v \in V, v' \in V'\})$; any colourings or labellings are inherited. G is ***disconnected*** if $G = G_1 + G_2$ for graphs G_i, else ***connected***, and ***coconnected*** if its complement is connected. The subgraph of (V, E) ***induced*** by $W \subseteq V$ is $(W, E{\restriction_W})$ for $E{\restriction_W}$ the restriction of E to edges on W. A graph is G-***free*** if G is not an induced subgraph. A ***cograph*** is a P_4-free graph, where $P_4 =$ •–•–•–• $= (\{v_1,v_2,v_3,v_4\}, \{v_1v_2, v_2v_3, v_3v_4\})$. In (V, E) the ***neighbourhood*** $N(v)$ of $v \in V$ is $\{w : vw \in E\}$, a ***module*** is a set $M \subseteq V$ such that $N(v) \setminus M = N(w) \setminus M$ for all $v,w \in M$, and M is ***strong*** if every module M' satisfies $M' \cap M = \emptyset$, $M' \subseteq M$ or $M' \supseteq M$. A ***directed graph*** (V, E) is a set V of vertices and a set $E \subseteq V \times V$ of ***directed edges***. A directed graph ***homomorphism*** $h : (V, E) \to (V', E')$ is a function $h : V \to V'$ such that $\langle v,w \rangle \in E$ implies $\langle h(v), h(w) \rangle \in E'$.

3 Fographs (first-order cographs)

A cograph is ***logical*** if every vertex is labelled by a variable or atom, and it has at least one atom-labelled vertex. Write $\bullet\lambda$ for a λ-labelled vertex.

Definition 3.1. The ***graph*** $\mathcal{G}(\varphi)$ of a formula φ is the logical cograph defined inductively by:

$$\mathcal{G}(\alpha) \;=\; \bullet\alpha \quad \text{for every atom } \alpha$$

$$\begin{aligned} \mathcal{G}(\,\varphi \vee \theta\,) &= \mathcal{G}(\varphi) + \mathcal{G}(\theta) & \qquad \mathcal{G}(\,\forall x\, \varphi\,) &= \bullet x \,+\, \mathcal{G}(\varphi) \\ \mathcal{G}(\,\varphi \wedge \theta\,) &= \mathcal{G}(\varphi) \times \mathcal{G}(\theta) & \qquad \mathcal{G}(\,\exists x\, \varphi\,) &= \bullet x \,\times\, \mathcal{G}(\varphi) \end{aligned}$$

For example, $\exists x(\overline{p}x \vee \forall y\, py)$ and $\exists x \forall y\,(py \vee \overline{p}x)$ have the same graph D:

$$\begin{aligned} D \;&=\; \mathcal{G}\;\big(\,\exists x(\overline{p}x \vee \forall y\, py)\,\big) \\ &=\; \mathcal{G}\;\big(\,\exists x \forall y\,(py \vee \overline{p}x)\,\big) \;=\; \end{aligned}$$

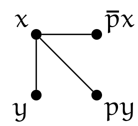

Vertices of $\mathcal{G}(\varphi)$ correspond to occurrences of atoms and quantifiers in φ: each occurence of an atom α in φ becomes an α-labelled vertex, and each occurrence of a quantifier $\forall x$ or $\exists x$ becomes an x-labelled vertex. A ***literal*** is an atom-labelled vertex and a ***binder*** is a variable-labelled vertex. Thus D has two literals, $\bullet\overline{p}x$ and $\bullet py$, and two binders, $\bullet x$ and $\bullet y$ (obtained from $\exists x$ and $\forall y$).

A module is ***proper*** if it has two or more vertices. The ***scope*** of a binder b is the smallest proper strong module strictly containing b (*i.e.*, containing b and at least one other vertex).[1,2] For example, in D, the scope of $\bullet y$ is $\{\,\bullet y, \bullet\overline{p}x, \bullet py\,\}$, and the scope of $\bullet x$ is $\{\,\bullet x, \bullet y, \bullet\overline{p}x, \bullet py\,\}$, illustrated below by shading.

scope of $\bullet y$ $=$

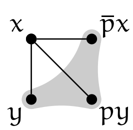

scope of $\bullet x$ $=$

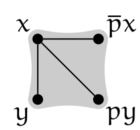

A binder is ***universal*** if its scope contains no edge, otherwise ***existential***. In D, for example, $\bullet y$ is universal and $\bullet x$ is existential (corresponding to $\forall y$ and $\exists x$ in the formula(s) generating D). An x-***binder*** is a binder with variable x, which is ***legal*** if its scope contains at least one literal and no other x-binder.

Definition 3.2. A ***fograph*** or ***first-order cograph*** is a logical cograph whose binders are legal.

For example, D above is a fograph, but $x\bullet\!\!-\!\!\bullet y\;\;\bullet p$ is not (since neither binder scope contains a literal), nor is $\bullet x\;\;\bullet x\;\;\bullet px$ (since each x-binder is in the other's scope).

Lemma 3.3. The graph $\mathcal{G}(\varphi)$ of every formula φ is a fograph.[3]

Proof. By structural induction on φ. The base case with φ an atom is immediate. For the induction step, note that all four operations defined in Def. 3.1 preserve the property of being a fograph, since all formulas are rectified.[4] $\square$

An x-***literal*** is one whose atom contains the variable x. An x-binder ***binds*** every x-literal in its scope. In D above, for example, $\bullet x$ binds $\bullet\overline{p}x$ and $\bullet y$ binds $\bullet py$. An x-binder is ***rectified*** if it is the only x-binder and its scope contains every x-literal. A fograph is ***rectified*** if its binders are rectified. For example, D is rectified but $x\bullet\!\!-\!\!\bullet px\;\;x\bullet\!\!-\!\!\bullet qx$ is not (since it has two x-binders), nor is $x\bullet\!\!-\!\!\bullet px\;\;\bullet qx$ (since $\bullet x$ does not bind $\bullet qx$). To ***rectify*** an unrectified x-binder b in a fograph G is to change its label to a variable x' which is fresh (not in any label of G) and substitute x' for x in the label of every literal bound by b. A ***rectified form*** is any result of rectifying binders until reaching a rectified form. For example, $\bullet px\;\;x\bullet\!\!-\!\!\bullet qx\;\;x\bullet\!\!-\!\!\bullet rxx$ has the rectified form

[1] Since, by definition, every logical cograph has a literal, the requisite strong module in the scope definition exists.

[2] To discern scope it is helpful to draw the modular decomposition tree [Gal67], *i.e.*, cotree [CLS81]. See Lemma 10.60.

[3] In §10 we will observe that $\mathcal{G}$ is a surjection onto fographs (Lemma 10.1).

[4] Naively applying $\mathcal{G}$ to an unrectified formula such as $(\forall x\, px) \vee (\forall x(qx \vee rx))$ yields $\bullet x\;\;\bullet px\;\;\bullet x\;\;\bullet qx\;\;\bullet rx$ with all three literals bound ambiguously by both binders. Whence our assumption that every formula is in rectified form.

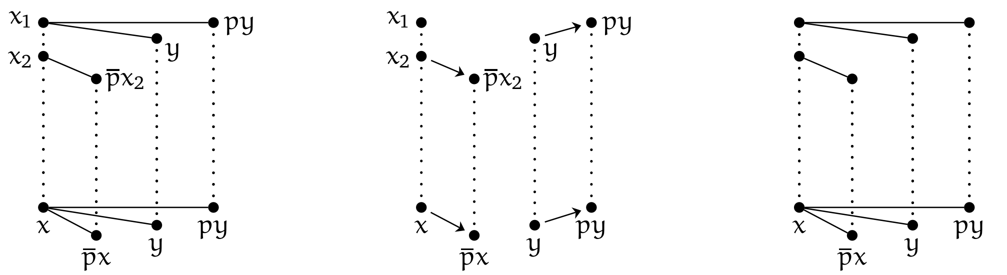

Figure 4. A skew bifibration (left), its binding fibration (centre), and its skeleton (right).

$\bullet px \quad y\bullet\!\!-\!\!\bullet qy \quad z\bullet\!\!-\!\!\bullet rzz$. This is analogous to the unrectified formula $(px \vee \exists x\, qx) \vee \exists x\, rxx$ having the rectified form $(px \vee \exists y\, qy) \vee \exists z\, rzz$.[5]

The ***binding graph*** $\vec{G}$ of a fograph G is the directed graph $(V_G, \{\langle b, l\rangle : b \text{ binds } l\})$. For example, the binding graph of D above is

$$\vec{D} \quad = \quad \begin{array}{c} x\bullet\longrightarrow\bullet\,\bar{p}x \\[1em] y\bullet\longrightarrow\bullet\,py \end{array}$$

4 Skew bifibrations

A directed graph homomorphism $f : (V, E) \to (V', E')$ is a ***fibration*** [Gro60, Gra66] if for all $v \in V$ and $\langle w, f(v)\rangle \in E'$ there exists a unique $\widetilde{w} \in V$ with $\langle \widetilde{w}, v\rangle \in E$ and $f(\widetilde{w}) = w$. This definition is illustrated below-left.

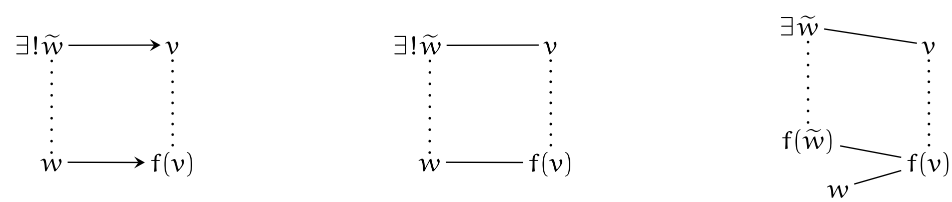

Similarly, an undirected graph homomorphism $f : (V, E) \to (V', E')$ is a ***fibration*** if for all $v \in V$ and $w\, f(v) \in E'$ there exists a unique $\widetilde{w} \in V$ with $\widetilde{w}\, v \in E$ and $f(\widetilde{w}) = w$. This definition is illustrated above-centre.[6] An undirected graph homomorphism $f : (V, E) \to (V', E')$ is a ***skew fibration*** [Hug06a] if for all $v \in V$ and $w\, f(v) \in E'$ there exists $\widetilde{w} \in V$ with $\widetilde{w}\, v \in E$ and $f(\widetilde{w})w \notin E'$. This definition is illustrated above-right. Since $f(\widetilde{w}) = w$ implies $f(\widetilde{w})w \notin E'$, skew fibrations generalize fibrations.

A graph homomorphism $f : K \to G$ between fographs ***preserves labels*** if for every vertex $v \in V_K$ the label of v in K equals the label of $f(v)$ in G, and ***preserves existentials*** if for every existential binder b in K the vertex $f(b)$ is an existential binder in G.

DEFINITION 4.1. A ***skew bifibration*** $f : K \to G$ between fographs is a label- and existential-preserving graph homomorphism such that

- $f : K \to G$ is a skew fibration
- $f : \vec{K} \to \vec{G}$ is a fibration.

We refer to $f : \vec{K} \to \vec{G}$ as the ***binding fibration***. For example, a skew bifibration is shown in Fig. 4, with its binding fibration. The ***skeleton*** of a skew bifibration is the result of dropping labels from its

[5] In §10 we will observe that $\mathcal{G}$ is a surjection onto rectified fographs (Lemma 10.1).

[6] An undirected graph fibration is a special case of a topological fibration [Whi78], by viewing every edge as a copy of the unit interval.

Figure 5. A fonet N (left) with unique dualizer $\{x{\mapsto}z, y{\mapsto}fz\}$ and its leap graph $\mathcal{L}_N$ (right).

source. Fig. 4 shows an example. We identify a skew bifibration with its skeleton. No information is lost since the source labels can be lifted from the target (because skew bifibrations preserve labels, by definition).

5 Fonets (first-order nets)

Definition 5.1. A coloured fograph is ***linked*** if

- every colour, called a ***link***, comprises two literals, and
- every literal is in a link.

Fig. 5 shows a linked fograph N with two links, $\{\blacksquare\bar{p}x, \blacksquare pz\}$ and $\{\circ\bar{q}y, \circ qfz\}$.

Let K be a linked fograph. Without loss of generality, assume K is rectified (by rectifying binders as needed). A ***dualizer*** for K is a function δ assigning to each existential binder variable x a term such that, for every link $\{\bullet\alpha_1, \bullet\alpha_2\}$, the atoms $\alpha_1\delta$ and $\alpha_2\delta$ are dual, where $\alpha\delta$ denotes the result of substituting $\delta(x)$ for x throughout α (simultaneously for each x). For example, $\{x{\mapsto}z, y{\mapsto}fz\}$ is a dualizer[7] for N (Fig. 5) since $\bar{p}x\{x{\mapsto}z, y{\mapsto}fz\} = \bar{p}z$ is dual to pz, and $\bar{q}y\{x{\mapsto}z, y{\mapsto}fz\} = \bar{q}fz$ is dual to qfz; this is the unique dualizer for N.

A ***dependency*** $\{\bullet x, \bullet y\}$ of K is an existential binder $\bullet x$ and a universal binder $\bullet y$ such that every dualizer for K assigns to x a term containing y.[8] For example, $\{\bullet y, \bullet z\}$ is a dependency of N (Fig. 5) since the unique dualizer $\{x{\mapsto}z, y{\mapsto}fz\}$ assigns fz to y. A ***leap*** is a dependency or link. The ***leap graph*** $\mathcal{L}_K$ is the graph (V_K, L_K) where L_K comprises all leaps of K. See Fig. 5 for an example.

A graph (V, E) is a ***matching*** if V is non-empty and for all $v \in V$ there is a unique $v' \in V$ with $vv' \in E$. A set W ***induces a bimatching*** in a linked fograph K if W induces matchings in K and $\mathcal{L}_K$.

Definition 5.2. A ***fonet*** or ***first-order net*** is a linked fograph which has a dualizer but no induced bimatching.

See Fig. 5 for an example of a fonet N.

6 Combinatorial proofs

Definition 6.1. A ***combinatorial proof*** of a fograph G is a skew bifibration $f : N \to G$ from a fonet N. A combinatorial proof of a formula φ is a combinatorial proof of its graph $\mathcal{G}(\varphi)$.

For examples, see §1.

Theorem 6.2 (Soundness). *A formula is valid if it has a combinatorial proof.*

Proof. Section 10. □

Theorem 6.3 (Completeness). *Every valid formula has a combinatorial proof.*

Proof. Section 11. □

[7] In the context of a function we write $a{\mapsto}b$ for the ordered pair $\langle a, b\rangle$.

[8] In §14 we show that all dependencies can be constructed in polynomial time, despite quantification over *every dualizer*.

Figure 6. A standard combinatorial proof (left) and a homogeneous combinatorial proof (right) of Peirce's law $\big((p \Rightarrow q) \Rightarrow p\big) \Rightarrow p \;=\; \big((\overline{p} \vee q) \wedge \overline{p}\big) \vee p$.

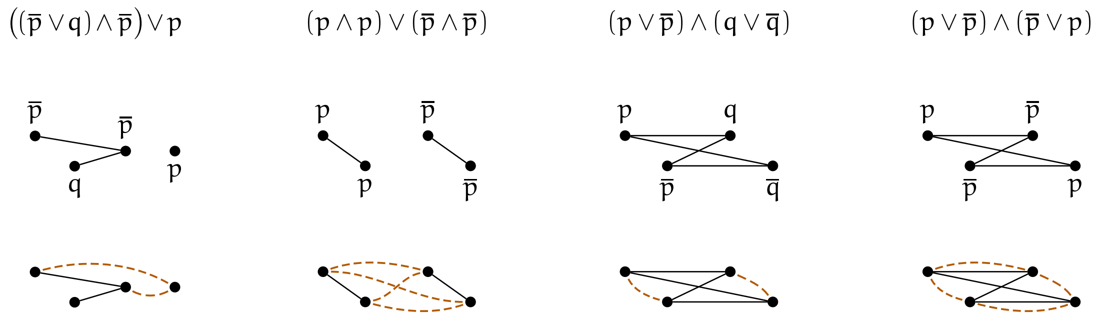

Figure 7. Four propositions φ (top row), their fographs $\mathcal{G}(\varphi)$ (middle row), and their dualizing graphs $\mathcal{D}(\varphi)$. Each vertex in $\mathcal{G}(\varphi)$ and $\mathcal{D}(\varphi)$ is aligned vertically with the corresponding atom occurence in φ. Dualities are shown dashed and curved.

Combining the two theorems above, we obtain the main theorem of this paper:

Theorem 6.4 (Soundness & Completeness). *A formula of first-order logic is valid if and only if it has a combinatorial proof.*

7 Propositional combinatorial proofs without labels

A ***proposition*** is a formula with no quantifiers or terms, *e.g.* $\big((\overline{p} \vee q) \wedge \overline{p}\big) \vee p$. This section provides an alternative representation of fographs and combinatorial proofs in the propositional case, without labels (variables and atoms). An illustrative example is shown in Fig. 6. The left side shows a standard combinatorial proof (Def. 6.1) of Peirce's law $\big((p \Rightarrow q) \Rightarrow p\big) \Rightarrow p = \big((\overline{p} \vee q) \wedge \overline{p}\big) \vee p$. The right side shows the label-free form, called a *homogeneous combinatorial proof*, defined below. The source colouring and target labels ($\overline{p}$, p and q) have disappeared, and both are replaced by *duality* edges, shown dashed and curved. The adjective *homogeneous* reflects the common type of the source and target (both cographs with additional duality edges), in contrast to a standard combinatorial proof skeleton which is *heterogeneous* (the source is coloured, while the target is labelled).

7.1 Dualizing graphs

A graph is ***triangle-free*** if it is C_3-free, where $C_3 = \triangle = (\{v_1, v_2, v_3\}, \{v_1v_2, v_2v_3, v_3v_1\})$.

Definition 7.1. A ***dualizing graph*** is a non-empty cograph D equipped with a second set $\perp_D$ of undirected edges on V_D, called ***dualities***, such that $(V_D, \perp_D)$ is a triangle-free cograph.

Four examples of dualizing graphs are shown in the bottom row of Fig. 7.

Definition 7.2. The dualizing graph $\mathcal{D}(\varphi)$ of a proposition φ is the dualizing graph D with

- $V_D = \{\,$occurrences of predicate symbols in $\varphi\,\}$,
- $vw \in E_D$ ies f and only if $v \neq w$ and the smallest subformula of φ containing both v and w is a conjunction (*i.e.*, of the form $\theta \wedge \psi$)
- $vw \in \perp_D$ if and only if v and w have dual predicate symbols (*e.g.*, p and $\overline{p}$).

For example, for each proposition φ in the top row of Fig. 7, the bottom row shows the corresponding dualizing graph $\mathcal{D}(\varphi)$. For comparison, the fograph $\mathcal{G}(\varphi)$ is shown in the middle row.

LEMMA 7.3. $\mathcal{D}(\varphi)$ is a well-defined dualizing graph for every proposition φ.[9]

Proof. Let $D = \mathcal{D}(\varphi)$. We must show (V_D, E_D) and $(V_D, \perp_D)$ are P_4-free, and $(V_D, \perp_D)$ is C_3-free.

Suppose $(\{v_1,v_2,v_3,v_4\},\{v_1v_2,v_2v_3,v_3v_4\})$ is an induced subgraph of (V_G, E_D). Since $v_1v_2 \in E_D$ there exist subformulas φ_1 and φ_2 of φ containing v_1 and v_2, respectively, with $\varphi_1 \wedge \varphi_2$ a subformula of φ. Necessarily v_3 is in φ_1, otherwise (since φ is a syntactic tree) $v_1v_3 \in E_D$ (a contradiction), and similarly v_4 is in φ_2, otherwise $v_2v_4 \in E_D$ (a contradiction). But then $v_1v_4 \in E_D$, a contradiction.

Suppose $(\{v_1,v_2,v_3,v_4\},\{v_1v_2,v_2v_3,v_3v_4\})$ is an induced subgraph of $(V_G, \perp_D)$, where v_i is an occurrence of the nullary predicate symbol p_i. By definition of $\perp_D$, we have $\overline{p}_1 = p_2$, $\overline{p}_2 = p_3$ and $\overline{p}_3 = p_4$. Thus $p_3 = p_1$, hence $p_4 = \overline{p}_1$, so $v_1v_4 \in \perp_D$, a contradiction.

Suppose $(\{v_1,v_2,v_3\},\{v_1v_2,v_2v_3,v_3v_1\})$ is an induced subgraph of $(V_G, \perp_D)$, where v_i is an occurrence of the nullary predicate symbol p_i. By definition of $\perp_D$, we have $\overline{p}_1 = p_2$, $\overline{p}_2 = p_3$ and $\overline{p}_3 = p_1$. Thus $p_3 = \overline{p}_2 = \overline{\overline{p}}_1 = p_1$, contradicting $\overline{p}_3 = p_1$. $\square$

7.2 Dualizing nets

A set $W \subseteq V_D$ ***induces a bimatching*** in a dualizing graph D if W induces matchings in both (V_D, E_D) and $(V_D, \perp_D)$.

DEFINITION 7.4. A ***dualizing net*** N is a dualizing graph with no induced bimatching, such that $(V_N, \perp_N)$ is a matching.

For example,  is a dualizing net, while  and  are not. The third dualizing graph in the bottom row of Fig 7 is a dualizing net, while the other three in the bottom row are not.

7.3 Propositional homogeneous combinatorial proofs

A ***skew fibration*** $f : C \to D$ of dualizing graphs is a skew fibration $f : (V_C, E_C) \to (V_D, E_D)$ such that $f : (V_C, \perp_C) \to (V_D, \perp_D)$ is a homomorphism.

DEFINITION 7.5. A ***homogeneous combinatorial proof*** of a dualizing graph D is a skew fibration $f : N \to D$ from a dualizing net N. A ***homogeneous combinatorial proof*** of a proposition φ is a homogeneous combinatorial proof of its dualizing graph $\mathcal{D}(\varphi)$.

For example, a homogeneous combinatorial proof of Peirce's law $\big((p \Rightarrow q) \Rightarrow p\big) \Rightarrow p = \big((\overline{p} \vee q) \wedge \overline{p}\big) \vee p$ is shown on the right of Fig 6.

7.4 Propositional homogeneous soundness and completeness

THEOREM 7.6 (Propositional homogeneous soundness and completeness). *A proposition is valid if and only if it has a homogeneous combinatorial proof.*

Proof. A corollary of Theorem 6.4, detailed in § 12.1. $\square$

Figure 8. A combinatorial proof $f : N \to G$ of the monadic formula $\exists x(px \Rightarrow \forall y\, py)$ (left), copied from the Introduction, and its homogeneous combinatorial proof $f' : N' \to G'$ (right). The directed edges of N' and G' are those of the binding graphs $\vec{N}$ and $\vec{G}$, the dashed edge of N' captures the colour of N, and the dashed edge of G' captures the duality between the two predicate symbols p and $\bar{p}$ in G.

$\forall x\,\big((px \wedge px) \vee \exists y\, \bar{p}y\big)$

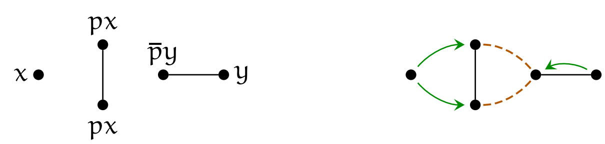

Figure 9. A monadic formula φ, its fograph $\mathcal{G}(\varphi)$, and its mograph $\mathcal{M}(\varphi)$, respectively.

8 Monadic combinatorial proofs without labels

A formula is ***monadic*** if its predicate symbols are unary and it has no function symbols, *e.g.*, $\exists x(px \Rightarrow \forall y\, py)$. This section extends homogeneous combinatorial proofs to the monadic case. Fig. 8 shows an illustrative example: on the left is the combinatorial proof of $\exists x(px \Rightarrow \forall y\, py)$ (copied from the Introduction), and on the right is the corresponding homogeneous combinatorial proof, to be defined below.

For technical convenience throughout this section we assume every monadic formula is closed, *i.e.*, has no free variables. This loses no generality because a formula φ with free variables $x_1, \dots, x_n$ is valid if and only if its closure $\forall x_1 \dots \forall x_n\, \varphi$ is valid.

Given a directed edge $e = \langle v, w\rangle$, v is the ***source*** of e, w is the ***target*** of e, and v and w are ***in*** e.

Definition 8.1. A ***pre-monadic graph*** or ***pre-mograph*** is a dualizing graph M equipped with a non-empty set B_M of directed edges on V_M, called ***bindings***, such that if a vertex v is the target of a binding then v is in no other binding.[10]

An example of a pre-mograph is shown on the right of Fig. 9, with two dualities (dashed and curved) and three bindings (directed and curved). A vertex in a pre-mograph M is a ***literal*** if it is the target of a binding, otherwise a ***binder***. If $\langle b, l\rangle \in B_M$ we say that b ***binds*** l.[11] The ***scope*** of a binder b in M is the smallest proper strong module of (V_M, E_M) strictly containing b (*i.e.*, containing b and at least one other vertex).

Definition 8.2. A ***mograph*** M is a pre-mograph such that no binder is in a duality, every binder has non-empty scope, and $\langle b, l\rangle \in B_M$ only if l is in the scope of b.

For example, the pre-mograph on the right of Fig. 9 is a mograph.

[9] We will observe in §13 that $\mathcal{D}$ is a surjection from propositions onto dualizing graphs (Lemma 13.2).

[10] In other words, if $\langle w, v\rangle \in B_M$, then (1) $\langle v, u\rangle \notin B_M$ for all vertices u, and (2) $\langle w', v\rangle \in B_M$ implies $w' = w$.

[11] Note that, by the condition in the definition of pre-mograph, necessarily b is a binder.

Definition 8.3. The ***mograph*** $\mathcal{M}(\varphi)$ of a closed monadic formula φ is the mograph defined by:

- $V_M = \{$ occurrences of atoms and quantifiers in $\varphi\,\}$,
- $vw \in E_M$ if and only if $v \neq w$ and either
  - the smallest subformula containing both v and w is a conjunction (*i.e.*, of the form $\varphi\wedge\theta$)
  - v is an existential quantifier and w is in its scope,
- $vw \in \perp_M$ if and only if v and w are atoms with dual predicate symbols (*e.g.*, px and $\overline{p}y$), and
- $\langle v, w\rangle \in B_M$ if and only if v is a quantifier, w is an atom, and v binds w.

For example, in Figure 9, the closed rectified monadic formula $\varphi = \forall x\,\big((px \wedge px) \vee \exists y\,\overline{p}y\big)$ on the left has the mograph $\mathcal{M}(\varphi)$ on the right.

Lemma 8.4. $\mathcal{M}(\varphi)$ is a well-defined mograph for every closed monadic formula φ.[12]

Proof. Let $M = \mathcal{M}(\varphi)$. Since every atom-occurrence in φ has a single variable, each literal is the target of at most one binding in M, and since no atom-occurrence binds another atom-occurrence, M satisfies the condition on bindings in the definition of pre-mograph (Def. 8.1). By reasoning as in the proof of Lemma 7.3, $(V_M, \perp_M)$ is P_4-free and C_3-free. By definition of $\mathcal{M}$, no binder is in a duality. It remains to show that (V_M, E_M) is a cograph, every binder has non-empty scope, and $\langle b, l\rangle \in B_M$ only if l is in the scope of b. We proceed by induction on the structure of φ.

Base case: $\varphi = px$ for some p and x, so M is a single vertex, hence a mograph.

Induction case: $\varphi = \varphi_1 * \varphi_2$ for $* \in \{\wedge, \vee\}$. By induction hypothesis $M_i = \mathcal{M}(\varphi_i)$ is a mograph $(i = 1, 2)$. By definition of E_M, we have $(V_M, E_M) = (V_{M_1}, E_{M_1}) \times (V_{M_2}, E_{M_2})$ or $(V_M, E_M) = (V_{M_1}, E_{M_1}) + (V_{M_2}, E_{M_2})$, thus (V_M, E_M) is a cograph since each (V_{M_i}, E_{M_i}) is a cograph. The scope of a binder b in M is at least the scope of b in the M_i containing b, thus the scope of b in M is non-empty and contains every literal bound by b, since M_i is a mograph.

Induction case: $\varphi = \nabla x\,\varphi'$ for $\nabla \in \{\forall, \exists\}$. By induction hypothesis, $M' = \mathcal{M}(\varphi')$ is a mograph. By definition of E_M we have $(V_M, E_M) = b + M'$ or $(V_M, E_M) = b \times M'$ for a vertex b (the initial occurrence of ∇x in φ), thus (V_M, E_M) is a cograph since $(V_{M'}, E_{M'})$ is a cograph. The scope of b in M comprises every literal, and is therefore non-empty and contains every literal bound by b. The scope of any other binder b' in M is equal the scope of b' in M', so is non-empty and contains every literal bound by b', since M' is a mograph. $\square$

8.1 Monets

A mograph is ***linked*** if every literal is in a unique duality. An example of a linked mograph is shown in Fig. 10 (left).

Definition 8.5. Let M be a linked mograph. Its ***binder equivalence*** $\simeq_M$ is the equivalence relation on binders generated by $b_1 \simeq_M b_2$ if there exist literals l_1 and l_2 with $\langle b_1, l_1\rangle, \langle b_2, l_2\rangle \in B_M$ and $l_1 l_2 \in \perp_M$.

Thus $b_1 \simeq_M b_2$ if and only if there exists a binding/duality pattern of the form

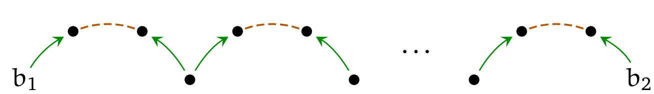

Let M be a linked mograph. A binder in M is ***universal*** if its scope contains no edge, otherwise ***existential***. A ***conflict*** in M is a pair $\{b, c\}$ of distinct universal binders b and c such that $b \simeq_K c$.

Definition 8.6. A mograph is ***consistent*** if it has no conflict.

A ***dependency*** of M is a pair $\{b, c\}$ of binders with $b \simeq_K c$, b existential, and c universal. A ***leap*** is a duality or dependency.

[12]We will observe in §13 that $\mathcal{M}$ is a surjection from closed monadic formulas onto mographs (Lemma 13.3).

Figure 10. A monet N (left) and its leap graph $\mathcal{L}_N$ (right).

Definition 8.7. The ***leap graph*** $\mathcal{L}_M$ of a linked mograph M is (V_M, L_M) for L_M the set of leaps of M.

An example of a leap graph is shown in Fig. 10 (right). A set of vertices $W \subseteq V_M$ ***induces a bimatching*** in a linked mograph M if W induces matchings in both (V_M, E_M) and $\mathcal{L}_M$.

Definition 8.8. A ***monet*** (*monadic net*) is a consistent linked mograph with no induced bimatching.

An example of a monet is shown in Fig. 10.

8.2 Monadic homogeneous combinatorial proofs

A function $f : V_N \to V_M$ between mographs ***preserves existentials*** if for every existential binder b in N the vertex $f(b)$ is an existential binder in M.

Definition 8.9. A ***skew bifibration*** $f : N \to M$ between mographs is an existential-preserving skew fibration

- $f : (V_N, E_M) \to (V_N, E_M)$ such that
- $f : (V_N, \bot_M) \to (V_N, \bot_M)$ is a homomorphism and
- $f : (V_N, B_M) \to (V_N, B_M)$ is a fibration.

An example of a skew bifibration between mographs is shown on the right of Fig. 8.

Definition 8.10. A ***homogeneous combinatorial proof*** of a mograph M is a skew bifibration $f : N \to M$ from a monet N. A ***homogeneous combinatorial proof*** of a closed monadic formula φ is a homogeneous combinatorial proof of its mograph $\mathcal{M}(\varphi)$.

A homogenous combinatorial proof of $\exists x(px \Rightarrow \forall y\, py)$ is shown in Fig 8 (right).

8.3 Monadic homogeneous soundness and completeness

Theorem 8.11 (Monadic homogeneous soundness and completeness). *A closed monadic formula is valid if and only if it has a homogeneous combinatorial proof.*

Proof. A corollary of Theorem 6.4, detailed in § 12.2. ☐

9 Modal combinatorial proofs

A ***modal*** formula is generated from the ***modal operators*** $\Box$ (necessity) and $\Diamond$ (possibility) instead of quantifiers and has nullary predicate symbols, *e.g.* $\Diamond(p \Rightarrow \Box p)$. Every modal formula abbreviates a standard first-order one [Min92, §3.3]: replace every $\Box$ by $\forall x$, $\Diamond$ by $\exists x$, and predicate symbol p by px. For example, $\Diamond(p \Rightarrow \Box p)$ abbreviates $\exists x(px \Rightarrow \forall x\, px)$, or $\exists x(px \Rightarrow \forall y\, py)$ in rectified form.

Definition 9.1. A ***modal combinatorial proof*** of a modal formula μ is a standard combinatorial proof (Definition 6.1) of the first-order formula abbreviated by μ.

For example, a modal combinatorial proof of $\Diamond(p \Rightarrow \Box p)$ is shown below-left, in condensed form.

$\Diamond\,(\,p \Rightarrow \Box\,p\,)$

$\exists x\,(\,px \Rightarrow \forall y\,py\,)$

It abbreviates the first-order combinatorial proof above-right (copied from the Introduction).

Theorem 9.2 (S5 Modal Soundness & Completeness). *A modal formula is valid in S5 modal logic if and only if it has a modal combinatorial proof.*

Proof. By Theorem 3.2 of [Min92, p. 42], a modal formula is valid in S5 if and only if the first-order formula it abbreviates is valid in first-order logic. Thus the result follows from Theorem 6.4. ☐

9.1 Modal combinatorial proofs without labels

A modal formula μ is ***closed*** if every predicate symbol occurrence is bound by a modal operator, *e.g.* $\Diamond(p \Rightarrow \Box p)$ but not $p \Rightarrow \Box p$.

Definition 9.3. The mograph $\mathscr{M}(\mu)$ of a closed modal formula μ is the mograph M defined by

- $V_M = \{$ occurrences of predicate symbols and modal operators in $\mu\}$,
- $vw \in E_M$ if and only if $v \neq w$ and either
  - the smallest subformula containing both v and w is a conjunction (*i.e.*, of the form $\varphi \wedge \theta$)
  - v is a $\Diamond$ and w is in its scope,
- $vw \in \perp_M$ if and only if v and w are dual predicate symbols, and
- $\langle v, w\rangle \in B_M$ if and only if v is a modal operator, w is a predicate symbol, and v binds w.

For example,

$$\mathscr{M}(\Diamond(p \Rightarrow \Box p)) \quad = $$

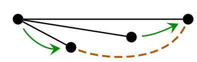

Definition 9.4. A ***homogeneous combinatorial proof*** of a closed modal formula μ is a homogeneous combinatorial proof of its mograph $\mathscr{M}(\mu)$.

Theorem 9.5 (Modal homogeneous soundness and completeness). *A closed modal formula is valid in S5 modal logic if and only if it has a homogeneous combinatorial proof.*

Proof. Since $\mathscr{M}(\mu) = \mathcal{M}(\mu')$ for μ' the first-order formula encoded by μ, the result is a corollary of Theorem 9.2. ☐

10 Proof of the Soundness Theorem

In this section we prove the Soundness Theorem, Theorem 6.2.

Lemma 10.1. The function $\mathcal{G}$ (Def. 3.1) is a surjection from rectified formulas onto rectified fographs. Two rectified formulas have the same graph if and only if they are equal modulo

$$\varphi \wedge \theta = \theta \wedge \varphi \qquad \varphi \wedge (\theta \wedge \psi) = (\varphi \wedge \theta) \wedge \psi \qquad \exists x \exists y\, \varphi = \exists y \exists x\, \varphi \qquad \varphi \wedge \exists x\, \theta = \exists x\, (\varphi \wedge \theta)$$

$$\varphi \vee \theta = \theta \vee \varphi \qquad \varphi \vee (\theta \vee \psi) = (\varphi \vee \theta) \vee \psi \qquad \forall x \forall y\, \varphi = \forall y \forall x\, \varphi \qquad \varphi \vee \forall x\, \theta = \forall x\, (\varphi \vee \theta)$$

Proof. A routine induction. ☐

Let G be a rectified fograph. Using the above Lemma, choose a formula φ such that $\mathcal{G}(\varphi) = G$. Define G as ***valid*** if φ is valid. This is well-defined with respect to choice of φ since every equality in Lemma 10.1 is a logical equivalence. Define a coloured fograph as valid if its underlying uncoloured fograph is valid.

Write $\models \chi$ to assert that a formula or fograph χ is valid, and $\varphi\{x \mapsto t\}$ for the result of substituting a term t for all occurrences of the variable x in a formula φ, but only if no variable in t is a bound variable of φ [TS96, §1.1.2].

Lemma 10.2. Let φ, θ and ψ be formulas.

1. $\models \varphi \wedge \theta$ if and only if ($\models \varphi$ and $\models \theta$).
2. $\models \varphi \vee \theta$ if ($\models \varphi$ or $\models \theta$).
3. $\models (\varphi \vee \theta) \wedge \psi$ implies $\models (\varphi \wedge \psi) \vee (\theta \wedge \psi)$.
4. $\models \forall x\, \varphi$ if and only if $\models \varphi$.
5. $\models \varphi\{x \mapsto t\}$ implies $\models \exists x\, \varphi$.
6. $\models \varphi \vee \theta\{x \mapsto t\}$ implies $\models \varphi \vee \exists x\, \theta$.
7. $\models (\varphi \vee \theta) \wedge \psi$ implies $\models \varphi \vee (\theta \wedge \psi)$.

Proof. 1–6 are standard inferences and properties of validity in first-order classical logic. See [TS96] and [Joh87], for example. Property 7 follows from 1 and 3. $\square$

10.1 Soundness of fonets

In this section we prove that fonets are sound, *i.e.*, every fonet is valid (Lemma 10.23 below).

Let G be a fograph. A set $P \subseteq V_G$ is ***well-founded*** if P contains a binder only if P contains a literal.

Definition 10.3. A ***portion*** of a rectified fograph G is a set $P \subseteq V_G$ such that P and $V_G \setminus P$ are well-founded, and P is closed under adjacency and binding: if $vw \in E_G$ or $\langle v,w \rangle \in E_{\vec{G}}$, then $v \in P$ if and only if $w \in P$.

A variable x in a fograph G is ***bound*** if G contains an x-binder, and ***free*** if G contains an x-literal but no x-binder. Two fographs are ***independent*** if any variable in both is free in both.

10.1.1 Fusion

Definition 10.4. Let G and G' be independent rectified fographs with respective portions P and P'. The ***fusion*** of G and G' at P and P' is the union $G+G'$ together with edges between every vertex in P and every vertex in P'.

For example, if $G = x\bullet\!\!-\!\!\bullet\overline{p}x \;\; py\bullet$, $G' = \bullet q \;\bullet\overline{q} \;\bullet z$, $P = \{\, py\bullet \,\}$ and $P' = \{\, \bullet q,\ \bullet\overline{q} \,\}$, then the fusion of G and G' at P and P' is $x\bullet\!\!-\!\!\bullet\overline{p}x \;\; py\bullet\!\!-\!\!\bullet q \;\bullet\overline{q} \;\bullet z$. Colourings are inherited during fusion, since they are inherited during graph union $+$. For example, if $K = x\bullet\!\!-\!\!\square\overline{p}x \;\; py\square$, $K' = \diamond q \;\diamond\overline{q} \;\bullet z$, $P = \{\, py\square \,\}$ and $P' = \{\, \diamond q,\ \diamond\overline{q} \,\}$, then the fusion of K and K' at P and P' is $x\bullet\!\!-\!\!\square\overline{p}x \;\; py\square\!\!-\!\!\diamond q \;\diamond\overline{q} \;\bullet z$.

The subgraph of a graph (V, E) ***induced*** by $W \subseteq V$ is $(W, E{\restriction_W})$ for $E{\restriction_W}$ the restriction of E to edges on W.

Lemma 10.5. Every fusion of valid rectified fographs is valid.

Proof. Let F be the fusion of valid rectified fographs G and G' at portions P and P'. We consider four cases.

1. P or P' is empty. Without loss generality, we may assume both are empty, since with one portion empty the fusion operation no longer depends on the other. Thus $F = G + G'$ for rectified fographs G and G', so by Lemma 10.1 there exist formulas φ and φ' with $\mathcal{G}(\varphi) = G$ and $\mathcal{G}(\varphi') = G'$. Since G and G' are independent, $\mathcal{G}(\varphi \vee \varphi') = F$. Since $\models G$ and $\models G'$ we have $\models \varphi$ and $\models \varphi'$, hence $\models \varphi \vee \varphi'$ by Lemma 10.2.2. Thus $\models F$.
2. $P = V_G$ and $P' = V_{G'}$. Thus $F = G \times G'$. As in the previous case we have valid formulas φ and φ' with $\mathcal{G}(\varphi) = G$ and $\mathcal{G}(\varphi') = G'$. Thus $\models F$ since $\models \varphi \wedge \varphi'$ by Lemma 10.2.1 and $F = \mathcal{G}(\varphi \wedge \varphi')$.

3. $P = V_G$ or $P' = V_{G'}$, and the previous two cases do not hold. Without loss of generality assume $P' = V_{G'}$, so $\emptyset \neq P \neq V_G$. Let $P^* = V_G \setminus P \neq \emptyset$. Thus $F = G[P^*] + (G[P] \times G')$. By Lemma 10.1 there exist formulas φ^*, φ and φ' with $\mathcal{G}(\varphi^*) = G[P^*]$, $\mathcal{G}(\varphi) = G[P]$ and $\mathcal{G}(\varphi') = G'$. Since $\models G'$ we have $\models \varphi'$, and since $\models G$ and $\mathcal{G}(\varphi^* \vee \varphi) = G$, we have $\models \varphi^* \vee \varphi$. Thus $\models (\varphi^* \vee \varphi) \wedge \varphi'$ by Lemma 10.2.1, so $\models \varphi^* \vee (\varphi \wedge \varphi')$ by Lemma 10.2.7, hence $\models F$ since $F = \mathcal{G}(\varphi^* \vee (\varphi \wedge \varphi'))$.

4. Otherwise $\emptyset \neq P \neq V_G$ and $\emptyset \neq P' \neq V_{G'}$. Let $P^* = V_G \setminus P \neq \emptyset$ and $P'^* = V_{G'} \setminus P' \neq \emptyset$. Thus the rectified fograph F is $G[P^*] + G'[P'^*] + (G[P] \times G'[P'])$. By Lemma 10.1 there exist formulas φ^*, φ'^*, φ and φ' with $\mathcal{G}(\varphi^*) = G[P^*]$, $\mathcal{G}(\varphi'^*) = G'[P'^*]$, $\mathcal{G}(\varphi) = G[P]$, and $\mathcal{G}(\varphi') = G'[P']$. Since $\models G$ and $\mathcal{G}(\varphi^* \vee \varphi) = G$ we have $\models \varphi^* \vee \varphi$, and since $\models G'$ and $\mathcal{G}(\varphi'^* \vee \varphi') = G'$ we have $\models \varphi'^* \vee \varphi'$. Thus $\models (\varphi^* \vee \varphi'^*) \vee (\varphi \wedge \varphi')$, hence $\models F$ since $F = \mathcal{G}((\varphi^* \vee \varphi'^*) \vee (\varphi \wedge \varphi'))$. $\square$

LEMMA 10.6. Every fusion of two rectified fonets is a rectified fonet.

Proof. Let F be a fusion of rectified fonets K and K'. Since each portion is closed under adjacency, F is a union of cographs, hence is a cograph. Every binder scope contains a literal, by inheritance from K and K'. Since K and K' are rectified, and no links traverse between the two in F, every union of dualizers for K and K' is a dualizer for F, and vice versa. Thus the set of dependencies of F is the union of those of K and K', so any $W \subseteq V_F$ inducing a bimatching in F would induce a bimatching in K or K'. $\square$

10.1.2 Universal quantification

DEFINITION 10.7. Let G be a rectified fograph with no x-binder. The ***universal quantification*** of G by x is $\bullet x + G$.

LEMMA 10.8. Every universal quantification of a valid rectified fograph is valid.

Proof. Let $G = \bullet x + H$ be the universal quantification of a valid rectified fograph H by x. By Lemma 10.1 there exists a formula φ such that $\mathcal{G}(\varphi) = H$, and $\models \varphi$ since $\models G$. Thus $\mathcal{G}(\forall x\, \varphi) = G$, hence $\models G$ since $\models \forall x\, \varphi$ if and only if $\models \varphi$, by Lemma 10.2.4. $\square$

If K is a coloured rectified fograph, in the universal quantification $\bullet x + K$ we assume that the colouring of K is inherited, while $\bullet x$ remains uncoloured.

LEMMA 10.9. Every universal quantification of a rectified fonet is a rectified fonet.

Proof. Let K' be the universal quantification $\bullet x + K$. Dualizers for K are dualizers for K', and vice versa, since if x occurs in K, it has merely transitioned from free to bound. The leap graph of K' is that of K together with additional dependencies involving $\bullet x$. Since $\bullet x$ is in no edge, any $W \subseteq V_{K'}$ inducing a bimatching in K' would induce a bimatching in K. $\square$

10.1.3 Existential quantification

DEFINITION 10.10. Let G be a rectified fograph without the variable x, let P be a non-empty portion of G, and let ω be a set of occurrences of a term t in labels of literals in P, such that t contains no bound variable of G. The ***existential quantification*** of G by x at ω in P is $\bullet x + G\{t \mapsto_\omega x\}$ together with an edge between $\bullet x$ and each vertex in P, where $G\{t \mapsto_\omega x\}$ is the result of substituting x for every occurrence of t in ω.

For example, if $G = \bullet pfgy \;\; \bullet \bar{p}fgy$, $P = \{\bullet pfgy\}$ and ω is the occurrence of the term gy in $\bullet pfgy$, the existential quantification of G by x at ω is $x\bullet\!\!-\!\!\bullet pfx \;\; \bullet\bar{p}fgy$, while if ω is empty the existential quantification becomes $x\bullet\!\!-\!\!\bullet pfgy \;\; \bullet\bar{p}fgy$. If $P = \{\bullet pfgy, \bullet\bar{p}fgy\}$ and ω comprises both occurrences of the term fgy in P, then the existential quantification is $x\bullet\!\!-\!\!\bullet px \;\; \bullet\bar{p}x$

LEMMA 10.11. Every existential quantification of a valid rectified fograph is valid.

Proof. Let H be the existential quantification of a valid rectified fograph G by x at a set ω of occurrences of the term t in the non-empty portion P. Thus $H = \bullet x + G\{t \mapsto_\omega x\}$ plus edges from $\bullet x$ to every vertex in P. We consider two cases.

1. Suppose $P = V_G$. Thus $H = \bullet x \times G\{t \mapsto_\omega x\} = \bullet x \times \mathcal{G}(\varphi) = \mathcal{G}(\exists x\, \varphi)$. By Lemma 10.1 there exists a formula φ such that $\mathcal{G}(\varphi) = G\{t \mapsto_\omega x\}$. Since x does not occur in G we have $\mathcal{G}(\varphi\{x \mapsto t\}) = G$, and $\models \varphi\{x \mapsto t\}$ since $\models G$. By Lemma 10.2.5 we have $\models \exists x\, \varphi$ since $\models \varphi\{x \mapsto t\}$, thus $\models H$.

2. Otherwise $\emptyset \neq P \neq V_G$. Let $P^* = V_G \setminus P \neq \emptyset$. Since P is a portion, it is well-founded and closed under adjacency and binding, $G\{t \mapsto_\omega x\} = G\{t \mapsto_\omega x\}[P^*] + G\{t \mapsto_\omega x\}[P]$ with $G\{t \mapsto_\omega x\}[P^*]$ and $G\{t \mapsto_\omega x\}[P]$ both rectified fographs, and $G\{t \mapsto_\omega x\}[P^*] = G[P^*]$ since ω does not intersect P^*. Thus $G\{t \mapsto_\omega x\} = G[P^*] + G\{t \mapsto_\omega x\}[P]$. By Lemma 10.1 there exist formulas θ^* and θ with $\mathcal{G}(\theta^*) = G[P^*]$ and $\mathcal{G}(\theta) = G\{t \mapsto_\omega x\}[P]$. Thus

$$H \;=\; G[P^*] + \bullet x \times G\{t \mapsto_\omega x\}[P] \;=\; \mathcal{G}(\theta^*) + \mathcal{G}(\exists x\, \theta) \;=\; \mathcal{G}(\theta^* \vee \exists x\, \theta)$$

   Since $\mathcal{G}(\theta) = G\{t \mapsto_\omega x\}[P]$ and x does not occur in G we have $\mathcal{G}(\theta\{x \mapsto t\}) = G[P]$. Thus

$$G \;=\; G[P^*] + G[P] \;=\; \mathcal{G}(\theta^*) + \mathcal{G}(\theta\{x \mapsto t\}) \;=\; \mathcal{G}(\theta^* \vee \theta\{x \mapsto t\})$$

   Since $\models G$ we have $\models \theta^* \vee \varphi\{x \mapsto t\}$, so by Lemma 10.2.6 we have $\models \theta^* \vee \exists x\, \theta$, hence $\models H$. $\square$

When quantifying a coloured rectified fograph existentially, the colouring is inherited, while the added binder remains uncoloured. For example, if $K = \square pfgy\;\; \square pfgy$, $P = \{\, \square pfgy \,\}$ and ω is the occurrence of the term y in $\square pfgy$, the existential quantification of K by x at ω in P is $x\bullet\!\!-\!\!\square pfgx\;\; \square \bar{p}fgy$. In the remainder of this section (§10.1.3) we prove that every existential quantification of a rectified fonet is a rectified fonet (Lemma 10.14).

Let K be a linked rectified fograph. An ***existential*** (resp. ***universal***) variable of K is one labelling an existential (resp. universal) binder in K. An ***output*** of a function is any element of its image. A ***stem*** of a dualizer δ for K is a variable in an output of δ but not in K. For example, if $K = x\bullet\!\!-\!\!\square px\;\; y\bullet\!\!-\!\!\square \bar{p}y\;\; \bullet z$ and z_1 and z_2 are variables, the dualizer $\{x \mapsto z_1, y \mapsto z_1\}$ has one stem z_1, $\{x \mapsto fz_1z_2, y \mapsto fz_1z_2\}$ has two stems z_1 and z_2, $\{x \mapsto fz_1z, y \mapsto fz_1z\}$ has one stem z_1, and $\{x \mapsto z, y \mapsto z\}$ has no stem. A dualizer δ ***generalizes*** a dualizer δ' if δ yields δ' by substituting terms for stems: there exists a function σ from the stems of δ to terms such that $\delta'(x) = \delta(x)\sigma$ for every existential variable x of K, where $e\sigma$ denotes the result of substituting $\sigma(z_1)$ for z_1 in e, simultaneously for each stem z_1 of δ. For example, if $K = x\bullet\!\!-\!\!\square px\;\; y\bullet\!\!-\!\!\square \bar{p}y\;\; \bullet z$ and z_1 is a variable, the dualizer $\delta = \{x \mapsto z_1, y \mapsto z_1\}$ generalizes $\delta' = \{x \mapsto fza, y \mapsto fza\}$ via $\{z_1 \mapsto fza\}$ since $\delta'(x) = \delta(y) = z_1\{z_1 \mapsto fza\} = fza$. A dualizer δ is ***most general*** if it generalizes every other dualizer. For example, $\{x \mapsto z_1, y \mapsto z_1\}$ is a most general dualizer for $x\bullet\!\!-\!\!\square px\;\; y\bullet\!\!-\!\!\square \bar{p}y\;\; \bullet z$ but $\{x \mapsto z, y \mapsto z\}$ is not. A linked rectified cograph is ***dualizable*** if it has a dualizer.

LEMMA 10.12. Every dualizable linked rectified fograph has a most general dualizer.

Proof. Let K be the dualizable linked rectified fograph. Every dualizer for K is, by definition, a unifier for the unification problem $\approx_K$ (binary relation on terms) [TS96, §7.2] defined by $t_i \approx_K t'_i$ for each link $\{\, \bullet pt_1 \ldots t_n, \bullet \bar{p}t'_1 \ldots t'_n \,\}$ and $1 \leqslant i \leqslant n$, solved for the existential variables. Let δ be a most general unifier of $\approx_K$ [TS96, §7.2]. By renaming variables as needed, we may assume that no output of δ contains an existential variable. Define δ' as the restriction of δ to existential variables. Since δ is a most general unifier, δ' is a most general dualizer. $\square$

Let K be a linked rectified fograph with dualizer δ. A pair $\{\bullet x, \bullet y\}$, with $\bullet x$ an existential binder and $\bullet y$ a universal binder, is a ***dependency*** of δ if $\delta(x)$ contains y.

LEMMA 10.13. Let K be a linked rectified fograph with a most general dualizer δ. A pair $\{\bullet x, \bullet y\}$ is a dependency of K if and only if $\{\bullet x, \bullet y\}$ is a dependency of δ.

Proof. Since δ is most general, for any dualizer δ' every dependency of δ is a dependency of δ'. By definition, $\{\bullet x, \bullet y\}$ is a dependency of K if and only if it is a dependency of every dualizer for K. Thus $\{\bullet x, \bullet y\}$ is a dependency of K if and only if it is a dependency of δ. $\square$

A variable in a rectified fonet is ***existential*** (resp. ***universal***) if it is the label of an existential (resp. universal) binder.

LEMMA 10.14. Every existential quantification of a rectified fonet is a rectified fonet.

Proof. Let K' be the existential quantification of K by x at ω in P, where ω is a set of occurrences of the term t. Since P is closed under adjacency, K' is a cograph and every binder scope in K' contains a literal.

For any dualizer δ for K, the function $\delta' = \delta' \cup \{x \mapsto t\}$ is a dualizer for K', since the links of K' are those of K but for some occurrences of t becoming x. The dependencies of δ in K are the same as those of δ' in K', since t contains no binder variable of K. Every dependency of K' is a dependency of K: a dependency of K' is a dependency of every dualizer of K', hence a dependency of δ' for every dualizer δ for K, thus a dependency of K.

Conversely, given a most general dualizer γ for K' we construct a dualizer $\hat{\gamma}$ for K. Let δ be a most general dualizer for K. Since γ is most general for K', there exists a function σ from the stems of γ to terms such that $t = \delta'(x) = \gamma(x)\sigma$. Let $\tilde{\sigma}$ be the restriction of σ to stems appearing in $\gamma(x)$. Define $\tilde{\gamma}$ by $\tilde{\gamma}(y) = \gamma(y)\tilde{\sigma}$, for every existential variable of K'. Thus, in particular, $\tilde{\gamma}(x) = t$. The function $\tilde{\gamma}$ is a dualizer for K' (since it is γ with terms substituted for stems), and has the same dependencies as γ because $\gamma(x)\tilde{\sigma} = t$ so $\tilde{\sigma}(z)$ is a sub-term of t for every stem z of γ in $\gamma(x)$, and t contains no bound variable of K, hence no bound variable of K'. Define $\hat{\gamma}$ as the restriction of $\tilde{\gamma}$ to the existential variables of K. (Thus $\tilde{\gamma} = \hat{\gamma} \cup \{x \mapsto t\}$.) The function $\hat{\gamma}$ is a dualizer for K since for every link $\{\bullet p t_1 \ldots t_n, \bullet \bar{p} u_1 \ldots u_n\}$ in K we have $t_i\hat{\gamma} = u_i\hat{\gamma}$, because for the corresponding link $\{\bullet p t'_1 \ldots t'_n, \bullet \bar{p} u'_1 \ldots u'_n\}$ in K' we have $t'_i\tilde{\gamma} = u'_i\tilde{\gamma}$ with $t_i = t'_i\{x \mapsto t\}$ and $u_i = u'_i\{x \mapsto t\}$, and by construction $\tilde{\gamma}(x) = t$. The dualizer $\hat{\gamma}$ is a restriction of $\tilde{\gamma}$, which has the same dependencies as γ, thus $\hat{\delta}$ has the same dependencies as γ. Since γ is most general for K', its dependencies are those of K', hence every dependency of K' is a dependency of K.

Since the dependencies of K and K' coincide, the leap graphs $\mathcal{L}_K$ and $\mathcal{L}_{K'}$ are identical but for an extra vertex $\bullet x$ in the latter which is not in any leap. Thus induced bimatchings of K and K' coincide, so K' is a fonet because K is a fonet. $\square$

10.1.4 Soundness of fonets

An ***axiom*** is a coloured rectified fograph comprising two dual literals of the same colour, for example, $\square\, pxfy \quad \square\, \bar{p}xfy$.

Lemma 10.15. Every coloured rectified fograph constructed from axioms by fusion and quantification is a rectified fonet.

Proof. Every axiom is a rectified fonet, and fusion and quantification preserve the property of being a rectified fonet, by Lemmas 10.6, 10.9, and 10.14. $\square$

A fonet is ***universal*** if it has a binder in no edge (necessarily a universal binder).

Lemma 10.16. Every universal rectified fonet is a universal quantification of a rectified fonet.

Proof. Let N be a universal rectified fonet, with (universal) binder $\bullet x$ in no edge. The result N^- of deleting $\bullet x$ from N is a fonet, since N^- inherits all dualizers from N (because x goes from being universal to being free) and if W induces a bimatching in N^- then W induces a bimatching in N. Since N^- is an induced subgraph of a rectified fograph, N^- is rectified. Since $N = \bullet x + N^-$, the rectified fonet N is the universal quantification of the rectified fonet N^- by x. $\square$

Lemma 10.17. Every fonet with no edge and no binder is a union $\lambda_1 + \ldots + \lambda_n$ of axioms λ_i $(n \geqslant 1)$.

Proof. Since N has no edges, it has no existential binders, hence the empty dualizer. Thus every link in N has literals with dual atoms, so the subgraph of N induced by each link is an axiom. Since N has no edges, it is the union of these axioms. $\square$

Lemma 10.18. Let N be a rectified fonet with underlying uncoloured fograph $G = K_1 + (H_1 \times H_2) + K_2$ for each H_i a fograph and each K_j empty or a fograph. Suppose no leap of N is between $V_{K_1} \cup V_{H_1}$ and $V_{H_2} \cup V_{K_2}$. Then N is a fusion of rectified fonets.

Proof. Since N is a fograph and no leap goes between $V_{K_1} \cup V_{H_1}$ and $V_{H_2} \cup V_{K_2}$, the graphs $K_1 + H_1$ and $H_2 + K_2$ are well-defined fonets. Thus N is a fusion of rectified fonets. $K_1 + H_1$ and $H_2 + K_2$ at V_{H_1} and V_{H_2}. $\square$

LEMMA 10.19. Let N be a rectified fonet with underlying uncoloured fograph $G = K_1 + (\bullet x \times H) + K_2$ for H a fograph and each K_i empty or a fograph. Suppose no leap of N is between $V_{K_1} \cup \{\bullet x\}$ and $V_H \cup V_{K_2}$. Then the binder $\bullet x$ is in no leap of N.

Proof. In this proof *leap supposition* refers to the supposition on leaps in the Lemma statement. Suppose for a contradiction that $\{\bullet x, \bullet y\}$ is a leap, hence dependency, of N. By the leap supposition, the universal binder $\bullet y$ is in K_1. Let δ be a most general dualizer for N, which exists by Lemma 10.12. Since $\{\bullet x, \bullet y\}$ is a dependency, the term $\delta(\bullet x)$ contains y, by Lemma 10.13. There must be a link $\{v, w\}$ such that the atom label of the literal v contains x, otherwise $\delta(\bullet x) = z$ for a stem variable z not occurring in N, so $\delta(\bullet)$ would not contain y. Since N is rectified, the literal v must be in the scope of $\bullet x$, thus v is in H. The atom label of w cannot contain y, since w would then be in K_1 (because N is rectified so w must be in the scope of $\bullet y$, which is in K_1), and $\{v, w\}$ would be a link (hence leap) between H and K_1, contradicting the leap supposition. Thus, for $\delta(x)$ to be a term containing y, there must be a link $\{v, w\}$ with the label of v containing x and the label of w containing an existential variable x' such that the term $\delta(x')$ contains y. Therefore N has a leap $\{\bullet x', \bullet y\}$. Since v is in H and $\{v, w\}$ is a link, hence a leap, by the leap supposition w must be in H or K_2. Because N is rectified, the literal w must be in the scope of the existential binder $\bullet x'$, so $\bullet x'$ is in H or K_2. Since $\bullet y$ is in K_1, the leap $\{\bullet x', \bullet y\}$ is between K_1 and H or K_2, contradicting the leap supposition. □

LEMMA 10.20. Let N be a rectified fonet with underlying uncoloured fograph $G = K_1 + (\bullet x \times H) + K_2$ for H a fograph and each K_i empty or a fograph. Suppose no leap of N is between $V_{K_1} \cup \{\bullet x\}$ and $V_H \cup V_{K_2}$. Then N is an existential quantification of a rectified fonet by x.

Proof. By Lemma 10.19 the existential binder $\bullet x$ is in no leap of N. Let δ be a most general dualizer for N and let $t = \delta(x)$. Define N' as the result of deleting $\bullet x$ from N and substituting t for x in the atom label of every literal. Since N is a rectified fograph and $\bullet x$ is in no leap of N, N' is a rectified fograph. Thus N is an existential quantification of N' by x at ω in the portion $V_{G_{m+1}}$ for ω the set of occurrences of t in N' which replaced occurrences of x in N during the construction of N. □

The ***mate*** of a literal l in a linked fograph is the other literal in the unique link containing l.

LEMMA 10.21. Every non-universal rectified fonet with at least one edge is a fusion of rectified fonets or an existential quantification of a rectified fonet.

Proof. Let N be a non-universal fonet with an edge, and let G be its underlying uncoloured fograph. Since G is a (labelled) cograph, it has the form $G = (G_1 \times G_2) + (G_3 \times G_4) + \ldots + (G_{n-1} \times G_n) + L$ for fographs G_i and L, where L has no binder or edge, and $n \geqslant 1$ since N (hence G) has an edge. Let Ω be the graph whose vertices are the G_i with $G_i G_j \in E(\Omega)$ if and only if N has an edge or leap $\{v, w\}$ with $v \in V(G_i)$ and $w \in V(G_j)$. A *1-factor* is a set of pairwise disjoint edges whose union contains all vertices. Since N is a fonet, $Z = \{G_1G_2, G_3G_4, \ldots, G_{n-1}G_n\}$ is the only 1-factor of Ω. For if Z' is another 1-factor, then $Z' \setminus Z$ determines a set of leaps in N whose union induces a bimatching in N: for each $G_iG_j \in Z' \setminus Z$ pick a leap $\{v, w\}$ with $v \in V(G_i)$ and $w \in V(G_j)$. Since Ω has a unique 1-factor, some $G_m G_{m+1} \in Z$ is a bridge [Kot59, LP86], *i.e.*, $(V_\Omega, E_\Omega \setminus G_m G_{m+1}) = X + Y$ with $G_m \in V_X$ and $G_{m+1} \in V_Y$. Without loss of generality assume $G_i \in V_X$ for $i \leqslant m$ and $G_j \in V_Y$ for $i \geqslant m+1$. Let L_X (resp. L_Y) be the restriction of L to literals with mate in a vertex of X (resp. Y), thus $L = L_X + L_Y$ since L contains only literals and no binders. Define $K_1 = L_X + (G_1 \times G_2) + \ldots + (G_{m-2} \times G_{m-1})$ and $K_2 = L_Y + (G_{m+2} \times G_{m+3}) + \ldots + (G_{n-2} \times G_n)$. Thus $G = K_1 + (G_m \times G_{m+1}) + K_2$.

Since L comprises literals only, each of K_1 and K_2 is either empty or a fograph. If G_m and G_{m+1} both contain a literal, they are fographs, so we can appeal to Lemma 10.18 with $H_1 = G_m$ and $H_2 = G_{m+1}$ to conclude that N is a fusion of rectified fonets. Otherwise one of G_m or G_{m+1}, say G_m, has no literal, thus $G_m = \bullet x$. Then G_{m+1} must contain a literal, since G hence $G_m \times G_{m+1}$ is a fograph, therefore G_{m+1} is a fograph. Applying Lemma 10.20 with $H = G_{m+1}$, we conclude that N is an existential quantification of a rectified fonet. □

LEMMA 10.22. Every rectified fonet can be constructed from axioms by fusion and quantification.

Proof. Let N be a rectified fonet. We proceed by induction on the number of binders and edges in N. In the base case with no edge or binder, by Lemma 10.17, N is a union of axioms, hence a fusion of axioms since union is a special case of fusion (with empty portions). If N is universal, apply Lemma 10.16 then appeal to induction with one less binder. Thus we may assume N is non-universal with a binder or edge. Had N no edge, it would have no binder (since every existential binder must be in an edge, and a universal binder would make N universal), thus N has at least one edge. Apply Lemma 10.21 then appeal to induction with fewer edges. $\square$

LEMMA 10.23 (Fonet soundness). Every fonet is valid.

Proof. By Lemma 10.22 every fonet can be constructed from axioms by fusion and quantification. Since every axiom is valid, and fusion and quantification preserve validity by Lemmas 10.5, 10.8, and 10.11, every fonet is valid. $\square$

10.2 Soundness of skew bifibrations

Throughout this section we no longer assume implicitly that every formula is rectified.

An ***intrusion*** is a formula of the form $\varphi \vee \forall x\,\theta$, $(\forall x\,\theta) \vee \varphi$, $\varphi \wedge \exists x\,\theta$, or $(\exists x\,\theta) \wedge \varphi$, for any variable x and formulas φ and θ. A formula is ***extruded*** if no subformula is an intrusion. For any variable x, an x***-quantifier*** is a quantifier of the form $\forall x$ or $\exists x$. A formula is ***unambiguous*** if no x-quantifier is in the scope of another x-quantifier, for every variable x. A formula is ***clear*** if it is extruded and unambiguous.

DEFINITION 10.24. The ***graph*** $\mathbb{G}(\varphi)$ of a clear formula φ is the logical cograph defined inductively by:

$$\mathbb{G}(\alpha) \;=\; \bullet\alpha \quad \text{for every atom } \alpha$$

$$\begin{aligned} \mathbb{G}(\,\varphi \vee \theta\,) &= \mathbb{G}(\varphi) + \mathbb{G}(\theta) & \qquad \mathbb{G}(\,\forall x\,\varphi\,) &= \bullet x + \mathbb{G}(\varphi) \\ \mathbb{G}(\,\varphi \wedge \theta\,) &= \mathbb{G}(\varphi) \times \mathbb{G}(\theta) & \qquad \mathbb{G}(\,\exists x\,\varphi\,) &= \bullet x \times \mathbb{G}(\varphi) \end{aligned}$$

Note that $\mathbb{G}$ coincides with $\mathcal{G}$ (Def. 3.1) on extruded rectified formulas.

LEMMA 10.25. The function $\mathbb{G}$ is a surjection from clear formulas onto fographs. Two clear formulas have the same graph if and only if they are equal modulo

$$\begin{aligned} \varphi \wedge \theta &= \theta \wedge \varphi & \qquad \varphi \wedge (\theta \wedge \psi) &= (\varphi \wedge \theta) \wedge \psi & \qquad \exists x\, \exists y\, \varphi &= \exists y\, \exists x\, \varphi \\ \varphi \vee \theta &= \theta \vee \varphi & \qquad \varphi \vee (\theta \vee \psi) &= (\varphi \vee \theta) \vee \psi & \qquad \forall x\, \forall y\, \varphi &= \forall y\, \forall x\, \varphi \end{aligned}$$

Proof. A routine induction, akin to the proof of Lemma 10.1. $\square$

Let G be a fograph. Using the above Lemma, choose a clear formula φ such that $\mathbb{G}(\varphi) = G$. Define G as ***valid*** if φ is valid. This is well-defined with respect to choice of φ since every equality in Lemma 10.25 is a logical equivalence.

Fographs G and H are $\wedge$***-compatible*** if $G \times H$ is a well-defined fograph and $\overrightarrow{G \times H} = \vec{G} + \vec{H}$, and $\vee$***-compatible*** if $G + H$ is a well-defined fograph and $\overrightarrow{G + H} = \vec{G} + \vec{H}$. Thus $\vee$- and $\wedge$-compatibility ensure that no new bindings are created during graph union and join. For any variable x, a fograph G is x***-compatible*** if G does not contain an x-binder $\bullet x$.

DEFINITION 10.26. Let G and H be fographs. Define the ***fograph connectives*** $\wedge$, $\vee$, $\forall$ and $\exists$ by:

- if G and H are $\wedge$-compatible, define $G \wedge H \;=\; G \times H$
- if G and H are $\vee$-compatible, define $G \vee H \;=\; G + H$
- for any variable x, if G is x-compatible, define $\forall x\, G \;=\; \bullet x + G$
- for any variable x, if G is x-compatible, define $\exists x\, G \;=\; \bullet x \times G$.

LEMMA 10.27. The fograph connectives $\wedge$, $\vee$, $\forall$ and $\exists$ are well-defined on fographs. In other words, given fographs as input(s), each connective produces a fograph as output.

Proof. By the compatibility constraints, no x-binder of $G \wedge H$, $G \vee H$, $\forall x\, G$, or $\exists x\, G$ can be in the scope of another x-binder. ☐

LEMMA 10.28. The following equalities hold for clear formulas:

$$\begin{array}{rclcrcl} \mathbb{G}(\,\varphi \vee \theta\,) & = & \mathbb{G}(\varphi) \vee \mathbb{G}(\theta) & \qquad & \mathbb{G}(\,\forall x\, \varphi\,) & = & \forall x\ \mathbb{G}(\varphi) \\ \mathbb{G}(\,\varphi \wedge \theta\,) & = & \mathbb{G}(\varphi) \wedge \mathbb{G}(\theta) & & \mathbb{G}(\,\exists x\, \varphi\,) & = & \exists x\ \mathbb{G}(\varphi) \end{array}$$

Proof. Since $\varphi \vee \theta$ and $\varphi \wedge \theta$ are clear, $\mathbb{G}(\varphi)$ and $\mathbb{G}(\theta)$ are $\wedge$- and $\vee$-compatible, thus $\mathbb{G}(\varphi) \vee \mathbb{G}(\theta)$ and $\mathbb{G}(\varphi) \wedge \mathbb{G}(\theta)$ are well-defined. Because $\forall x\, \varphi$ and $\exists x\, \varphi$ are clear, no x-quantifier occurs in φ, so $\mathbb{G}(\varphi)$ contains no binder $\bullet x$, thus $\forall x\,(\mathbb{G}(\varphi))$ and $\exists x\,(\mathbb{G}(\varphi))$ are well-defined. ☐

LEMMA 10.29. A labelled graph is a fograph if and only if it can be constructed from literals by the fograph connectives $\wedge$, $\vee$, $\forall$ and $\exists$.

Proof. Let G be a fograph. By Lemma 10.25 there exists a clear formula φ such that $\mathbb{G}(\varphi) = G$. By Lemma 10.28 the $\times$ and $+$ operations in the inductive translation $\mathbb{G}$ of φ are well-defined $\wedge$, $\vee$, $\forall$ and $\exists$ operations on fographs. Thus G can be constructed from literals by fograph connectives. Conversely, any labelled graph constructed from literals is a fograph, by repeated application of Lemma 10.27, starting from the fact that any literal vertex is a fograph. ☐

A ***map*** is a label-preserving graph homomorphism between fographs.

DEFINITION 10.30. Extend the fograph connectives to maps $f : G \to H$ and $f' : G' \to H'$ as follows:

- if $G \wedge G'$ and $H \wedge H'$ are well-defined, define $f \wedge f' : G \wedge G' \to H \wedge H'$ as $f \cup f'$
- if $G \vee G'$ and $H \vee H'$ are well-defined, define $f \vee f' : G \vee G' \to H \vee H'$ as $f \cup f'$
- if $\forall x\, G$ and if $\forall x\, H$ are well-defined, define $\forall x\, f : \forall x\, G \to \forall x\, H$ as $f \cup \{\bullet x \mapsto \bullet x\}$
- if $\exists x\, G$ and if $\exists x\, H$ are well-defined, define $\exists x\, f : \exists x\, G \to \exists x\, H$ as $f \cup \{\bullet x \mapsto \bullet x\}$.

LEMMA 10.31. The fograph connectives are well-defined on skew bifibrations: if f and f' are skew bifibrations, then, when defined, each of the maps $f \wedge f'$, $f \vee f'$, $\forall x\, f$ and $\exists x\, f$ is a skew bifibration, where x is any variable.

Proof. Due to the compatibility constraint in the definitions of the fograph connectives, the skew fibration condition is preserved and the directed graph homomorphisms between binding graphs are fibrations. In the $\wedge$ and $\exists$ connectives, additional requisite skew liftings are created across the corresponding graph join. ☐

LEMMA 10.32. Skew bifibrations between fographs compose: if $f : G \to H$ and $f' : H \to K$ are skew bifibrations between fographs, their composition $f;f' : G \to K$ is a skew bifibration.

Proof. Skew fibrations between cographs compose [Hug06b, Cor. 3.5], and directed graph fibrations compose [Gro60]. Existential preservation is transitive. ☐

DEFINITION 10.33. If G is a fograph and $G \vee G$ is well-defined, define ***pure contraction*** C_G as the canonical map $G \vee G \to G$. If G and H are fographs and $G \vee H$ is well-defined, define ***pure weakening*** W_G^H as the canonical map $G \to G \vee H$.

LEMMA 10.34. Every pure contraction and pure weakening is a skew bifibration.

Proof. Immediate from the definitions of pure contraction and pure weakening. ☐

DEFINITION 10.35. A ***contraction*** is any map generated from a pure contraction by fograph connectives, and a ***weakening*** is any map generated from a pure weakening by fograph connectives.

LEMMA 10.36. Every contraction and weakening is a skew bifibration.

Proof. Pure contraction and pure weakening are skew bifibrations by Lemma 10.34, and fograph connectives are well-defined on skew bifibrations by Lemma 10.31. ☐

DEFINITION 10.37. A ***structural map*** is any map constructed from isomorphisms, contractions, and weakenings by composition.

LEMMA 10.38. Every structural map is a skew bifibration.

Proof. Every isomorphism is a skew bifibration, and every contraction and weakening is a skew bifibration by Lemma 10.36. Skew bifibrations compose by Lemma 10.32. □

LEMMA 10.39. Structural maps are sound: if G is a valid fograph and $f : G \to H$ is a structural map, then H is valid.

Proof. Isomorphisms, pure contraction, pure weakening, composition and fograph connectives are sound. □

10.2.1 The image of a skew bifibration is a fograph

We recall the *modular decomposition* [Gal67] of a cograph, called its ***cotree*** [CLS81].

A directed graph $(N, \prec)$ is ***acyclic*** if the transitive closure of $\prec$ (viewed as a binary relation on N) is irreflexive. A ***forest*** is an acyclic directed graph $(N, \prec)$ such that for every $n \in N$ there exists at most one $m \in N$ with $\langle n, m \rangle \in \prec$. We refer to the vertices of a forest as ***nodes***. Write $m \prec n$ or $n \succ m$ for $\langle n, m \rangle \in \prec$, and say that m is a ***child*** of n and n is the ***parent*** of m. A ***leaf*** (resp. ***root***) is a node with no child (resp. parent). A ***tree*** is a forest with a unique root. A $+\times$ ***tree*** is a tree in which a node is labelled $+$ or $\times$ if and only if it is not a leaf. Each node labelled $+$ or $\times$ is a $+\times$ ***node***. An isomorphism $\iota : (N, \prec) \to (N', \prec')$ of $+\times$ trees is a bijection $\iota : N \to N'$ such that $m \prec n$ if and only if $\iota(m) \prec' \iota(n)$ and $\iota(n)$ is a $+$ (resp. $\times$) node if and only if n is a $+$ (resp. $\times$) node. We identify $+\times$ trees up to isomorphism.

Given $+\times$ trees $T_1, \ldots, T_n$ for $n \geqslant 1$ define $+T_1 \ldots T_n$ (resp. $\times T_1 \ldots T_n$) as the disjoint union of the T_i together with a $+$ (resp. $\times$) root node r and an edge to r from the root of each T_i $(1 \leqslant i \leqslant n)$. Write $\bullet$ for the $+\times$ tree with a unique node. For example, the $+\times$ tree $+(\times\bullet\bullet)\bullet\big(\times\bullet\bullet(+\bullet\bullet)\big)$ is below-left and $+(\times\bullet\bullet)(+\bullet)\big(\times\bullet(\times\bullet(+\bullet\bullet))\big)$ is below-right.

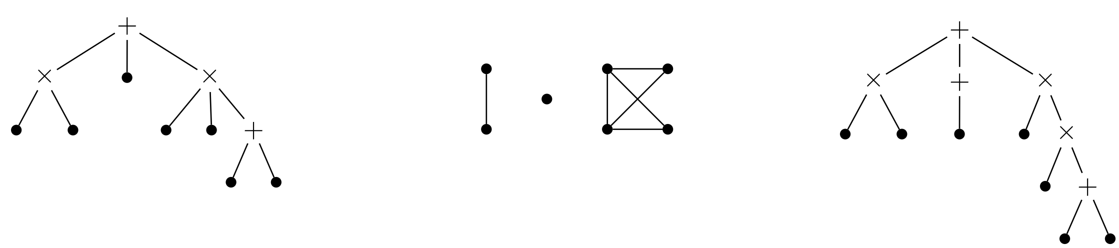

DEFINITION 10.40. The ***cograph*** $G(T)$ of a $+\times$ tree T is the cograph defined inductively by

$$G(\bullet) = \bullet \qquad G(+T_1 \ldots T_n) = G(T_1) + \ldots + G(T_n) \qquad G(\times T_1 \ldots T_n) = G(T_1) \times \ldots \times G(T_n)$$

For example, the cograph of the $+\times$ tree above-left is shown above-center; this cograph is also the cograph of the $+\times$ tree above-right.

LEMMA 10.41. The leaves of a $+\times$ tree T are in bijection with the vertices of its cograph $G(T)$.

Proof. Induction on the number of vertices in G, pattern-matching the three cases in Def. 10.40. □

A $+\times$ node ***repeats*** if it has a parent with the same label, and is ***unary*** if it has a unique child. A $+\times$ tree ***alternates*** if it has no repeating $+\times$ node and ***branches*** if it has no unary $+\times$ node.

DEFINITION 10.42. A ***cotree*** is a branching and alternating $+\times$ tree.

For example, the $+\times$ tree above-left is a cotree, while the $+\times$ tree above-right is not (since it has a repeating $\times$ node and a unary and repeating $+$ node). We recall the following definition from [CLS81].

DEFINITION 10.43. The ***cotree*** $T(G)$ of a cograph G is the cotree defined inductively by

$$T(\bullet) \;=\; \bullet \qquad \begin{aligned} T(G_1 + \ldots + G_n) &= +T(G_1)\ldots T(G_n) \quad \text{if } G_i \text{ is connected for } 1 \leqslant i \leqslant n \geqslant 2 \\ T(G_1 \times \ldots \times G_n) &= \times T(G_1)\ldots T(G_n) \quad \text{if } G_i \text{ is coconnected for } 1 \leqslant i \leqslant n \geqslant 2. \end{aligned}$$

The following Lemma articulates a standard property of cotrees. Recall from §2 that a module in a graph is *proper* if it has two or more vertices. A module M of a cograph G is ***connected*** (resp. ***coconnected***) if the induced subgraph $G[M]$ is connected (resp. coconnected).

LEMMA 10.44. The nodes of the cotree $T(G)$ of a cograph G correspond to the strong modules of G, and the $\times$ (resp. $+$) nodes correspond to proper connected (resp. coconnected) strong modules.

Proof. Induction on the number of vertices in G [CLS81]. ☐

The following Lemma is also a standard cotree property.

LEMMA 10.45. The function $T(-)$ is a bijection from cographs to cotrees.

Proof. Induction on the number of vertices in the cograph [CLS81]. ☐

LEMMA 10.46. The cotree $T(G)$ of a cograph G is the unique branching and alternating $+\times$ tree T such that $G(T) = G$.

Proof. A routine induction on the number of vertices in G. ☐

LEMMA 10.47. The vertices of a cograph G are in bijection with the leaves of its cotree $T(G)$.

Proof. Lemmas 10.41 and 10.46. ☐

Let n be a node in a tree $T = (N, \prec)$. Define the ***absorption*** $T{\uparrow}n$ of n in T as the result of deleting n (and incident edges) from T and, if n has a parent $\hat{n}$, adding an edge from each child of n to $\hat{n}$. Thus $N_{T\uparrow n} = N_T \setminus \{n\}$ and $m \prec_{T\uparrow n} m'$ if and only if $m \prec_T m'$ or $m \prec_T n \prec_T m'$.

DEFINITION 10.48. Given a $+\times$ tree T define its ***cotree*** $|T|$ as the cotree obtained by iteratively and exhaustively absorbing unary $+\times$ nodes and repeating $+\times$ nodes in T.

For example, if T is the $+\times$ tree above-right of Def. 10.40 then its cotree $|T|$ is above-left of Def. 10.40.

LEMMA 10.49. $G(T) = G(|T|)$ for every $+\times$ tree T.

Proof. By induction on the number of nodes in T, pattern-matching the three cases in Def. 10.40, combined with the associativity and commutativity of the graph union $+$ and join $\times$ operations. ☐

Recall that $G[U]$ is the subgraph of a graph G induced by a set of vertices U. Define the $+\times$ tree $T[U]$ ***induced*** by a non-empty set of leaves U in a $+\times$ tree T by deleting from T every leaf not in U, and then iteratively and exhaustively deleting any resulting childless $+\times$ nodes. For example, if T is the cotree below-left and U comprises the left-most four leaves of T, then the $+\times$ tree $T[U]$ is below-center, and the cotree $|T[U]|$ is below-right.

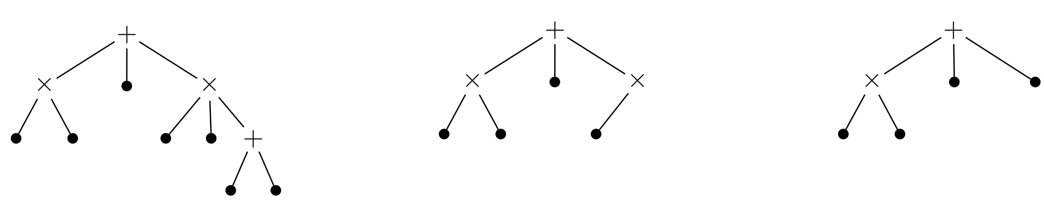

LEMMA 10.50. If U is a non-empty set of leaves in a cotree T, then $G(T[U]) \;=\; G(T)[U]$.

Proof. Induction on the number of nodes in T. ☐

LEMMA 10.51. If U is a non-empty set of vertices in a cograph G, then $G(|T(G)[U]|) = G[U]$.

Proof. By Lemma 10.49, $G(|T(G)[U]|) = G(T(G)[U])$, which is $G(T(G))[U]$ by Lemma 10.50, hence $G[U]$ by Lemma 10.46. ☐

LEMMA 10.52. If U is a non-empty set of vertices in a cograph G, then $\mathsf{T}(G[U]) = |\mathsf{T}(G)[U]|$.

Proof. By Lemma 10.46 it suffices to show that $\mathsf{G}(|\mathsf{T}(G)[U]|) = G[U]$, which is Lemma 10.51. $\square$

Write N_T for the set of nodes of a tree T, $\prec_T$ for its set of directed edges, $<_T$ for the transitive closure of $\prec_T$, and $\leqslant_T$ for the reflexive closure of $<_T$ (a partial order on N_T). Define $m >_T n$ as $n <_T m$, and say that m is ***above*** n or n is ***below*** m; define $m \geqslant_T n$ as $n \leqslant_T m$, and say that m is ***at or above*** n or n is ***at or below*** m. Define the ***meet*** $m \odot n$ of nodes m and n in a tree T as the $\leqslant_T$-least node o with $m \leqslant_T o$ and $n \leqslant_T o$.

LEMMA 10.53. Let G be a cograph and $v, w \in V_G$. Then $vw \in E_G$ if and only if $v \odot w$ in the cotree $\mathsf{T}(G)$ is a $\times$ node and $vw \notin E_G$ if and only if $v \odot w$ is a $+$ node or $v = w$.

Proof. This follows directly from Lemma 10.44. $\square$

Write $v \otimes w$ (resp. $v \oplus w$) for $v \odot w$ if it is a $\times$ (resp. $+$) node. For a cograph G write $\prec_G$, $<_G$, and $\leqslant_G$ for $\prec_{\mathsf{T}(G)}$, $<_{\mathsf{T}(G)}$, and $\leqslant_{\mathsf{T}(G)}$, respectively.

LEMMA 10.54. If G is a cograph with $vw, vu \in E_G$, $wu \notin E_G$, $w \neq u$, then $v <_G v \otimes w >_G w \oplus u >_G w, u$.

Proof. By Lemma 10.53 $v \odot w$ is a $\times$ node and $w \odot u$ is a $+$ node. Since $v \otimes w >_G w <_G w \oplus u$ and $\mathsf{T}(G)$ is a tree, either $v \otimes w <_G w \oplus u$ or $v \otimes w >_G w \oplus u$. If $v \otimes w <_G w \oplus u$ then $v \odot u$ is a $+$ node, contradicting $vu \in E_G$ (by Lemma 10.53), so $v \otimes w >_G w \oplus u$. $\square$

LEMMA 10.55. If G is a cograph with $vw, vu \notin E_G$ and $wu \in E_G$ then $v <_G v \oplus w >_G w \otimes u >_G w, u$.

Proof. Necessarily $w \neq u$ since $wu \in E_G$, hence $v \neq w$ and $v \neq u$. Thus we can apply Lemma 10.54 to the complement of G. $\square$

LEMMA 10.56. If $f : G \to H$ is a skew fibration between cographs and $v \prec_G v \oplus w >_G w$ for $v, w \in V_G$ with $f(v) \neq f(w)$, then $f(v) <_H f(v) \oplus f(w) >_H f(w)$.

Proof. Since $f(v) \neq f(w)$ the meet $f(v) \odot f(w)$ in $\mathsf{T}(H)$ is a $+$ or $\times$ node. If the former, then we have $f(v) <_H f(v) \oplus f(w) >_H f(w)$ as desired. Otherwise $f(v) <_H f(v) \otimes f(w) >_H f(w)$. By Lemma 10.53 $vw \notin E_G$ since $v \prec_G v \oplus w >_G w$, and $f(v)f(w) \in E_H$ since $f(v) <_H f(v) \otimes f(w) >_H f(w)$. Because f is a skew fibration and $f(v)f(w) \in E_H$, there exists $u \in V_G$ with $vu \in E_G$ and $f(w)f(u) \notin E_H$. Since f is a graph homomorphism, $f(v)f(u) \in E_H$ and $wu \notin E_G$, and $w \neq u$ (otherwise $vw \in E_G$ since $vu \in E_G$, contradicting $vw \notin E_G$). Since $wv, wu \notin E_G$ and $vu \in E_G$, by Lemma 10.55 we have $w <_G w \oplus v >_G v \otimes u >_G v$, hence $v <_G v \otimes u <_G v \oplus w$, contradicting $v \prec_G v \oplus w$. $\square$

The following Lemma refines $f(v) <_H f(v) \oplus f(w)$ in the above Lemma to $f(v) \prec_H f(v) \oplus f(w)$.

LEMMA 10.57. If $f : G \to H$ is a skew fibration between cographs and $v \prec_G v \oplus w >_G w$ for $v, w \in V_G$ with $f(v) \neq f(w)$, then $f(v) \prec_H f(v) \oplus f(w) >_H f(w)$.

Proof. By Lemma 10.56 we have $f(v) <_H f(v) \oplus f(w) >_H f(w)$. Suppose not $f(v) \prec_H f(v) \oplus f(w)$. Then $f(v) <_H f(v) \otimes u \prec_H f(v) \oplus f(w) >_H f(w)$ for some $u \in V_H$. Since $f(v) <_H f(v) \otimes u >_H u$ we have $f(v)u \in E_H$ by Lemma 10.53. Because f is a skew fibration and $f(v)u \in E_H$, there exists $\tilde{u} \in V_G$ with $v\tilde{u} \in E_G$ and $f(u)f(\tilde{u}) \notin E_H$. Necessarily $\tilde{u}w \in E_G$, otherwise Lemma 10.55 applied to $vw, \tilde{u}w \notin E_G$ and $v\tilde{u} \in E_G$ yields $w <_G w \oplus v >_G v \otimes \tilde{u} >_G v$ so $v <_G v \otimes \tilde{u} <_G w \oplus v$ contradicting $v \prec_G v \oplus w$. $\square$

LEMMA 10.58. If $f : G \to H$ is a skew fibration of cographs and $m \prec_H m \otimes n \succ_H n$, then $f(v) \leqslant_H m$ for some $v \in V_G$ if and only if $f(w) \leqslant_H n$ for some $w \in V_G$.

Proof. Assume $m \neq n$, otherwise the result is immediate. Suppose $f(v) \leqslant_H m$ for $v \in V_G$. Choose $u \in V_H$ with $u \leqslant_H n$. Thus $f(v) \leqslant_H m \prec_H m \otimes n \succ_H n \geqslant_H u$. Since $m \otimes n = f(v) \otimes u$ is a $\times$ node, we have $f(v)u \in E_H$ by Lemma 10.53. Because f is a skew fibration and $f(v)u \in E_H$, there exists $\hat{u} \in V_G$ with $v\hat{u} \in E_G$ and $f(\hat{u})u \notin E_H$. Since $v\hat{u} \in E_G$ and f is a graph homomorphism, we have $f(v)f(\hat{u}) \in E_G$. If $u = f(\hat{u})$ then since $n \geqslant_H u$ we have $f(\hat{u}) \leqslant_H n$ as desired. Otherwise $u \neq f(\hat{u})$, so applying Lemma 10.54 to $f(v)u, f(v)f(\hat{u}) \in E_H$, $uf(\hat{u}) \notin E_H$, $u \neq f(\hat{u})$ yields $f(v) <_H f(v) \otimes u >_H u \oplus f(\hat{u}) >_H u, f(\hat{u})$. Thus since $f(v) \otimes u = m \otimes n$ is the parent of n, both $f(v) \otimes u \succ_H n \geqslant_H u$ and $f(v) \otimes u >_H w \oplus f(\hat{u}) >_H u$, so because $\mathsf{T}(H)$ is a tree, we have $n \geqslant_H w \oplus f(\hat{u})$, hence $f(\hat{u}) \leqslant_H n$. $\square$

Given a function $f : V \to W$ write $f(V)$ for $\{ f(v) : v \in V \} \subseteq W$.

DEFINITION 10.59. Let $f : G \to H$ be a graph homomorphism between cographs. Define the ***image*** $\mathsf{Im}\, f$ as the subgraph $H[f(V_G)]$ of H induced by $f(V_G)$.

Define $v \rhd_T w$ if v and w are leaves and $\widehat{v} >_T w$ for the parent $\widehat{v}$ of v. Recall that a cograph is *logical* if every vertex is a binder or literal (*i.e.*, is labelled by a variable or atom), and at least one vertex is a literal. Write $S_G(b)$ for the scope of a binder b in a logical cograph G. The following Lemma shows that the scope of b is the set of leaves below the parent of b in the cotree $T(G)$.

LEMMA 10.60. For any vertex v and binder b in a logical cograph G, $v \in S_G(b)$ if and only if $b \rhd_{T(G)} v$.

Proof. By definition the scope of b is the smallest proper strong module containing b, which corresponds to the parent of b in the cotree $T(G)$ by Lemma 10.44. ☐

LEMMA 10.61. For any $+\times$ tree T and non-root $+\times$ node n in T, if $v \rhd_T w$ then $v \rhd_{T\uparrow n} w$.

Proof. Immediate from the definition of $T{\uparrow}n$. ☐

LEMMA 10.62. For any $+\times$ tree T, if $v \rhd_T w$ then $v \rhd_{|T|} w$.

Proof. Iterate Lemma 10.61 for every absorption step in the construction of $|T|$ in Def. 10.48. ☐

LEMMA 10.63. For any $+\times$ tree T and non-empty set U of leaves in T, if $v \rhd_T w$ and $v, w \in U$ then $v \rhd_{T[U]} w$.

Proof. The edge relation $\prec_{T[U]}$ is a subset of $\prec_T$. ☐

LEMMA 10.64. Let H logical cograph which is an induced subgraph of a logical cograph G. For every vertex v and binder b in H, if $v \in S_G(b)$ then $v \in S_H(b)$.

Proof. Let $v \in S_G(b)$. By Lemma 10.60, $b \rhd_{T(G)} v$. By Lemma 10.63, $b \rhd_{T(G)[V_H]} v$. By Lemma 10.62, $b \rhd_{|T(G)[V_H]|} v$. By Lemma 10.52, $|T(G)[V_H]| = T(G[V_H])$, and $T(G[V_H]) = T(H)$, thus $b \rhd_{T(H)} v$. Therefore $v \in S_H(b)$ by Lemma 10.60. ☐

A fograph map $f : G \to H$ ***preserves universals*** if every universal binder b in G maps to a universal binder $f(b)$ in H.

LEMMA 10.65. Every skew fibration between fographs preserves universals.

Proof. By Lemma 10.44, a binder is universal if and only if its parent in the cotree is a $+$ node. Thus the result follows from Lemma 10.57. ☐

A fograph map ***preserves binders*** if every universal (resp. existential) binder b in G maps to a universal (resp. existential) binder $f(b)$ in H.

LEMMA 10.66. Every skew bifibration between fographs preserves binders.

Proof. Skew bifibrations preserve existentials by definition, and universals by Lemma 10.65. ☐

LEMMA 10.67. For every skew bifibration $f : G \to H$ of fographs, the image $\mathsf{Im}\, f$ is a logical cograph.

Proof. Every vertex of $\mathsf{Im}\, f$ is inherited from H, and is therefore a binder or literal. Since $\mathsf{Im}\, f$ is an induced subgraph of a cograph H, it is a cograph. Because G is a logical cograph, it contains a literal l, thus $\mathsf{Im}\, f$ contains the literal $f(l)$ (a literal since f preserves labels). ☐

LEMMA 10.68. For every skew bifibration $f : G \to H$ between fographs and universal binder b in $\mathsf{Im}\, f$, the scope $S_H(b)$ contains a literal in $\mathsf{Im}\, f$.

Proof. Choose $\tilde{b}$ in G with $f(\tilde{b}) = b$. By Lemma 10.66, $\tilde{b}$ is universal. Since G is a fograph there exists a literal $l \in S_G(\tilde{b})$. Thus $\tilde{b} \prec_G \tilde{b} \oplus l >_G l$ by Lemma 10.60, and since $f(\tilde{b}) \neq f(l)$ by label preservation, by Lemma 10.57 we have $b \prec_H b \oplus f(l) >_H f(l)$, hence $f(l) \in S_H(b)$, by Lemma 10.60. ☐

LEMMA 10.69. For every skew bifibration $f : G \to H$ between fographs and universal binder b in $\mathsf{Im}\, f$, the scope $S_{\mathsf{Im}\, f}(b)$ contains a literal.

Proof. By Lemma 10.68, the scope $S_H(b)$ contains a literal $f(l)$. Since $\mathsf{Im}\, f$ is an induced subgaph of H, by Lemma 10.64 we have $f(l) \in S_{\mathsf{Im}\, f}(b)$. $\square$

LEMMA 10.70. For every skew bifibration $f : G \to H$ between fographs and existential binder b in $\mathsf{Im}\, f$, the scope $S_H(b)$ contains a literal in $\mathsf{Im}\, f$.

Proof. Since H is a fograph there exists a literal k in H in the scope of b. Thus $b \prec_H b \otimes k >_H k$ by Lemma 10.60. Therefore $b \prec_H b \otimes k \succ_H n \geqslant_H k$ for some child n of $b \otimes k$. Since b is in $\mathsf{Im}\, f$ there exists $\tilde{b} \in V_G$ with $f(\tilde{b}) = b$, so we may apply Lemma 10.58 with $m = b$ and $v = \tilde{b}$ to obtain $w \in V_G$ with $f(w) \leqslant_H n$. If w is a literal, then the literal $f(w)$ is in $S_H(b)$, and the Lemma holds.

Otherwise w is a binder, hence $f(w)$ is a binder. We proceed by induction on the number of vertices in the scope $S_H(b)$. Since $f(w) \leqslant_H n \geqslant_H k$ for $f(w)$ a binder and k a literal, n must be a $+$ or $\times$ node, and since $+$ and $\times$ alternate in a cotree, n is a $+$ node because its parent $b \otimes k$ is a $\times$ node. Let n' be the parent of $f(w)$. Thus $b \prec_H b \otimes k \succ_H n \geqslant_H n' \succ_H f(w)$. If n' is a $+$ node, then $f(w)$ is universal so by Lemma 10.69 the scope of $f(w)$ contains a literal in $\mathsf{Im}\, f$, *i.e.*, a literal $f(l)$ for some literal l in G. Therefore $b \prec_H b \otimes k \succ_H n \geqslant_H n' \succ_H f(l)$, so $f(l)$ is also in $S_H(b)$. Otherwise n' is a $\times$ node, so $f(w)$ is existential. Since $S_H(f(w))$ is strictly contained in $S_H(b)$, by induction there exists a literal l in G such that $f(l)$ is in $S_H(f(w))$, thus the literal $f(l)$ is in $S_H(b)$. $\square$

LEMMA 10.71. For every skew bifibration $f : G \to H$ between fographs and binder b in $\mathsf{Im}\, f$, the scope $S_H(b)$ contains a literal in $\mathsf{Im}\, f$.

Proof. If b is universal (resp. existential) apply Lemma 10.68 (resp. 10.70). $\square$

LEMMA 10.72. For every skew bifibration $f : G \to H$ between fographs and existential binder b in $\mathsf{Im}\, f$, the scope $S_{\mathsf{Im}\, f}(b)$ contains a literal.

Proof. By Lemma 10.70 there exists a literal l with $f(l)$ in $S_H(b)$. Since $\mathsf{Im}\, f$ is an induced subgraph of H, we have $f(l)$ in $S_{\mathsf{Im}\, f}(b)$ by Lemma 10.64. $\square$

DEFINITION 10.73. A logical cograph G is ***fair*** if binders b and b' have the same variable only if $bb' \notin E_G$.

Note that every rectified fograph is fair.

LEMMA 10.74. For every skew bifibration $f : G \to H$ between fographs with H fair, $\mathsf{Im}\, f$ is a fair fograph.

Proof. By Lemma 10.67 $\mathsf{Im}\, f$ is a logical cograph, and $\mathsf{Im}\, f$ is fair since if $bb' \in E_{\mathsf{Im}\, f}$ for binders b and b' with the same variable, then $bb' \in E_H$ since $\mathsf{Im}\, f$ is an induced subgraph, contradicting the fairness of H. It remains to show that (1) for every binder b in $\mathsf{Im}\, f$ the scope $S_{\mathsf{Im}\, f}(b)$ contains a literal, and (2) for every variable x and every x-binder b in $\mathsf{Im}\, f$, the scope $S_{\mathsf{Im}\, f}(b)$ contains no other x-binder.

(1) If b is universal (resp. existential), then by Lemma 10.69 (resp. 10.72), the scope $S_{\mathsf{Im}\, f}(b)$ contains a literal.

(2) Suppose b' were another x-binder with $b' \in S_{\mathsf{Im}\, f}(b)$, *i.e.*, $b \prec_{\mathsf{Im}\, f} b \odot b' >_{\mathsf{Im}\, f} b'$. If b is universal, then $b \odot b' = b \oplus b'$, so $b \prec_H b \oplus b' >_H b'$, whence $b' \in S_H(b)$, contradicting the fact that H is a fograph. Otherwise, b is existential, and $b \odot b' = b \otimes b'$, so $bb' \in E_{\mathsf{Im}\, f}$. Since $\mathsf{Im}\, f$ is an induced subgraph of H, we have $bb' \in E_H$, contradicting the fairness of H. $\square$

The following example illustrates why fairness of H is required to ensure no x-binder is in the scope of another in Lemma 10.74. Let $G = \mathbb{G}((\exists x\, px) \vee (\exists x\, px)) = x\bullet\!\!-\!\!\bullet px \;\; x\bullet\!\!-\!\!\bullet px$, let $H = \mathbb{G}(q \vee \exists x\, px) = \bullet q \;\; x\bullet\!\!-\!\!\bullet px$, and let f be the unique label-preserving graph homomorphism $G \to H$, which is a skew bifibration between fographs. Then $f \wedge f : G \wedge G \to H \wedge H$ is a skew bifibration between fographs, but its image $\mathsf{Im}\, f$ is $(x\bullet\!\!-\!\!\bullet px) \wedge (x\bullet\!\!-\!\!\bullet px)$, which is not well-defined fograph since each $\bullet x$ is in the scope of the other.

10.2.2 Marking and pruning

Let G be a cograph and let $\mathsf{U} \subseteq \mathsf{V_G}$. A node n in the cotree $\mathsf{T(G)}$ is ***over*** U if $\mathsf{n} \geqslant_\mathsf{G} \mathsf{u}$ for some vertex $\mathsf{u} \in \mathsf{U}$. Define the ***support*** $\mathsf{U}^* \subseteq \mathsf{N}_{\mathsf{T(G)}}$ as the set of nodes over U, and say that U is ***balanced*** for G if, for every $\times$ node n in $\mathsf{T(G)}$ and child m of n, we have $\mathsf{m} \in \mathsf{U}^*$ if $\mathsf{n} \in \mathsf{U}^*$.

LEMMA 10.75. If $\mathsf{f} : \mathsf{G} \to \mathsf{H}$ is a skew fibration between cographs then $\mathsf{f}(\mathsf{V_G}) \subseteq \mathsf{V_H}$ is balanced for H.

Proof. A corollary of Lemma 10.58. □

Let G be a cograph and let $\mathsf{U} \subseteq \mathsf{V_G}$. A $+\times$ node n in U^* is ***literal-supported*** if there exists a literal $\mathsf{l} \in \mathsf{U}$ with $\mathsf{n} \geqslant_\mathsf{G} \mathsf{l}$. We say that U is ***binding-closed*** if for every literal $\mathsf{l} \in \mathsf{U}$ and binder b in G, if b binds l then $\mathsf{b} \in \mathsf{U}$.

DEFINITION 10.76. Let G be a fograph. A set $\mathsf{U} \subseteq \mathsf{V_G}$ is a ***marking*** for G if it is balanced, every $+\times$ node of U^* is literal-supported, and U is binding-closed.

LEMMA 10.77. If $\mathsf{f} : \mathsf{G} \to \mathsf{H}$ is a skew bifibration between fographs then $\mathsf{f}(\mathsf{V_G})$ is a marking for H.

Proof. Let $\mathsf{U} = \mathsf{f}(\mathsf{V_G})$. By Lemma 10.58, U is balanced, by Lemma 10.71, every node n in U^* is literal-supported, and U is binding-closed since $\mathsf{f} : \vec{\mathsf{G}} \to \vec{\mathsf{H}}$ is a directed graph fibration. □

Let n be the child of a $+$ node m in a $+\times$ tree T. The node n is ***critical*** to m if n is the only child of m which is at or above a literal. If n is an x-binder for some variable x, then n is ***vacuous*** if it is the unique node in the subtree rooted at m whose label contains x.[13]

DEFINITION 10.78. A node n in a $+\times$ tree T is ***pareable*** if:

1. n has a parent m, a $+$ node,
2. n is not critical to m, and
3. if n is a binder (necessarily universal) then it is vacuous.

To ***pare*** a pareable node n in a $+\times$ tree T is to delete the subtree rooted at n.

DEFINITION 10.79. A ***pruning*** is any result of iteratively paring zero or more $+$ nodes.

LEMMA 10.80. Let G be a fograph with marking U, and let T be a $+\times$ tree such that $\mathsf{G(T)} = \mathsf{G}$. There exists a pruning T' of T with $\mathsf{G}(\mathsf{T}') = \mathsf{G}[\mathsf{U}]$.

Proof. A routine induction on the number of nodes in T. □

10.2.3 Decomposition of skew bifibrations

DEFINITION 10.81. If G is a connected fograph without the variable x, define ***slackening*** $\mathsf{S}^\mathsf{x}_\mathsf{G}$ as the canonical inclusion map $\mathsf{G} \to \forall\mathsf{x}\,\mathsf{G}$.

LEMMA 10.82. Every slackening is a structural map.

Proof. Weaken G to $\mathsf{G} \vee \forall\mathsf{x}\,\mathsf{G}$, which is $\forall\mathsf{x}(\mathsf{G} \vee \mathsf{G})$, then contract under $\forall\mathsf{x}$ to $\forall\mathsf{x}\,\mathsf{G}$. (Note that $\mathsf{G} \vee \mathsf{G}$ is well-defined because G is connected.) □

DEFINITION 10.83. A WS-***map*** is any map constructed from isomorphisms, weakenings and slackenings by composition and fograph connectives.

LEMMA 10.84. Every WS-map is a structural map.

Proof. Iterate Lemma 10.82. □

LEMMA 10.85. Let G be a fograph, let T be a $+\times$ tree such that $\mathsf{G(T)} = \mathsf{G}$, let T' be the result of paring a pareable node in T, and let $\mathsf{G}' = \mathsf{G}(\mathsf{T}')$. There exists a WS-map $\mathsf{G}' \to \mathsf{G}$.

[13]Thus in the cograph $\mathsf{G(T)}$, the universal binder n binds no literal.

Proof. If the paring is of a vacuous binder (condition 3 in Def. 10.78), then we obtain a slackening in the context of a fograph connective, otherwise (condition 2 in Def. 10.78) we obtain a weakening in the context of a fograph connective. □

LEMMA 10.86. Let G be a fograph, let T be a $+\times$ tree such that $\mathsf{G}(T) = G$, let T' be a pruning of T, and let $G' = \mathsf{G}(T')$. There exists a WS-map $G' \to G$.

Proof. Apply Lemma 10.85 to each paring in the pruning. □

LEMMA 10.87. Let $f : G \to H$ be a skew bifibration with H fair. The inclusion $\mathsf{Im}\, f \to H$ is a WS-map.

Proof. By Lemma 10.74, $\mathsf{Im}\, f$ is a fograph. Let $U = f(V_G)$, thus $\mathsf{Im}\, f$ is the induced subgraph $H[U]$. By Lemma 10.77, U is a marking. Let T be the cotree $\mathsf{T}(H)$. By Lemma 10.77, there exists a pruning T' of T with $\mathsf{G}(T') = H[U] = \mathsf{Im}\, f$. By Lemma 10.86, there exists a WS-map $\mathsf{G}(T') \to \mathsf{G}(T)$, *i.e.*, $\mathsf{Im}\, f \to H$. □

Let $f : G \to H$ be a skew fibration and let K be a connected component of H. The ***multiplicity*** of K is the number of connected components of $f^{-1}(K)$, and the ***weight*** of K is one more than its multiplicity. The ***weight*** of f is the sum of the weights of the connected components of H. A skew bifibration is ***shallow*** if the multiplicity of every connected component of H is at most one.

LEMMA 10.88. Every skew bifibration into a fair fograph is a structural map.

Proof. By induction on the weight of the skew bifibration $f : G \to H$ and its multiplicity. By Lemma 10.87 (and the fact that every WS-map is a structural map by Lemma 10.84) we may assume f is a surjection, and by pre-composing with contractions we may assume f is shallow. If $H = \bullet x + H'$ then $G = \bullet x + G'$ since f is a shallow surjection, hence $f = \forall x\, f'$, and by induction f' is a structural map. Otherwise if $H = H_1 + H_2$ then $H = H_1 \vee H_2$ (since H is not of the form $\bullet x + H'$). Since f is a shallow surjection, $G = G_1 \vee G_2$, so $f = f_1 \vee f_2$ for $f_i : G_i \to H_i$. By induction each f_i is a structural map, hence f is structural. Otherwise H is connected. If H has no edge then f is an isomorphism from a literal to a literal, hence is a structural map. Thus we may assume H has an edge. If $H = \bullet x \times H'$ then $G = \bullet x \times G'$ since f is a shallow surjection, hence $f = \exists x\, f'$, and by induction f' is a structural map. Otherwise $H = H_1 \times H_2$ for fographs H_i, with H_i not of the form $\bullet x \times H'_i$. Thus $H = G_1 \wedge G_2$, hence $G = G_1 \wedge G_2$ with $f(V_{G_i}) \subseteq V_{H_i}$. Therefore $f = f_1 \vee f_2$ for skew bifibrations $f_i : G_i \to H_i$, and by induction each f_i is a structural map, so f is a structural map. □

LEMMA 10.89 (Soundness of skew bifibrations). If G is a valid fograph and $f : G \to H$ is a skew bifibration with H fair, then H is valid.

Proof. By Lemma 10.88, f is a structural map, which is sound by Lemma 10.39. □

10.3 Proof of the Soundness Theorem

Proof of Soundness Theorem (Theorem 6.2). Let $f : N \to \mathcal{G}(\varphi)$ be a combinatorial proof of a formula φ. By Lemma 10.23 N is valid, thus by Lemma 10.89 $\mathcal{G}(\varphi)$ is valid (applicable since $\mathcal{G}(\varphi)$ is rectified, hence fair), therefore φ is valid. □

11 Proof of the Completeness Theorem

In this section we prove the Completeness Theorem, Theorem 6.3.

We shall employ variants $\mathsf{F1}$ and $\mathsf{F2}$ of the syntactic proof system $\mathsf{GS1}$ [TS96, §3.5.2] for first-order logic, which is a reformulation of Gentzen's LK [Gen35]. The system $\mathsf{F2}$ allows us to factorize a proof of any valid formula φ into two phases, *logical* then *structural*. The logical phase yields a fonet N, and the structural phase provides a skew bifibration $N \to \mathcal{G}(\varphi)$, hence a combinatorial proof of φ. This completeness strategy generalizes that of the propositional case in [Hug06b].

Throughout this section we no longer assume formulas are implicitly in rectified form. A ***sequent*** is a finite multiset of formulas, *i.e.*, a finite sequence of formulas modulo reordering [TS96, §1.1.4]. We write comma for disjoint union on sequents, and identify a singleton sequent with its formula.

Let $\Gamma = \varphi_1, \ldots, \varphi_n$ be a sequent, and without loss of generality assume $\varphi_1, \ldots, \varphi_n$ are ordered according to some fixed linear order on the set of all formulas. Define the ***formula*** $\Phi(\Gamma)$ of Γ as Γ, φ. Define Γ as ***valid*** if its formula is valid. Recall (from §10) that $\varphi\{x{\mapsto}t\}$ denotes the result of substituting a term t for all occurrences of the variable x in φ, but only if no variable in t becomes bound [TS96, §1.1.2]. The inference rules of F1, from which F1 proofs are generated, are as follows, in which Γ and Δ are arbitrary sequents, φ and θ are arbitrary formulas, and α is any atom.

$$\frac{}{\alpha, \overline{\alpha}} \qquad \frac{\Gamma,\, \varphi,\, \theta}{\Gamma,\, \varphi \vee \theta}\,\vee \qquad \frac{\Gamma}{\Gamma,\, \varphi}\,\mathsf{w} \qquad \frac{\Gamma,\, \varphi, \varphi}{\Gamma, \varphi}\,\mathsf{c}$$

$$\frac{\Gamma,\, \varphi \qquad \theta,\, \Delta}{\Gamma,\, \varphi \wedge \theta,\, \Delta}\,\wedge \qquad \frac{\Gamma,\, \varphi\{x{\mapsto}t\}}{\Gamma,\, \exists x\, \varphi}\,\exists \qquad \frac{\Gamma,\, \varphi}{\Gamma,\, \forall x\, \varphi}\,\forall \quad (x \text{ not free in } \Gamma)$$

We refer to the rule c as ***sequent-weakening*** and w as ***sequent-contraction***. Each sequent above a rule is a ***hypothesis*** of the rule, and the sequent below a rule is the ***conclusion*** of the rule.

LEMMA 11.1 (F1 soundness & completeness). A sequent is valid if and only if it has an F1 proof.

Proof. System F1 is equivalent to GS1, which is sound and complete [TS96, §3.5.2]. It differs only in the choice of $\wedge$ and $\vee$ rules, which are equivalent in the presence of c and w. ☐

A ***context*** [TS96, §1.1.3] is the variant of a formula with the symbol $\text{\textvisiblespace}$ as a generator in addition to all atoms, but with the restriction that $\text{\textvisiblespace}$ occurs exactly once in a context. Given a context χ we write $\chi[\varphi]$ for the result of substituting φ for $\text{\textvisiblespace}$ in χ. For example $\chi = \exists x(\text{\textvisiblespace} \vee \forall y\, py)$ is a context, and $\chi[\overline{p}x] = \exists x(\overline{p}x \vee \forall y\, py)$. A formula φ is ***connected*** if it is not of the form $\forall x\, \varphi'$ or $\varphi_1 \vee \varphi_2$ (thus a rectified formula φ is connected if and only if its graph $\mathcal{G}(\varphi)$ is connected). Define system F2 by:

$$\frac{}{\alpha, \overline{\alpha}} \qquad \frac{\Gamma,\, \varphi,\, \theta}{\Gamma,\, \varphi \vee \theta}\,\vee \qquad \frac{\Gamma,\, \chi[\varphi_i]}{\Gamma,\, \chi[\varphi_1 \vee \varphi_2]}\,\mathsf{W}_i \qquad \frac{\Gamma,\, \chi[\varphi \vee \varphi]}{\Gamma,\, \chi[\varphi]}\,\mathsf{C} \quad (\varphi \text{ connected})$$

$$\frac{\Gamma,\, \varphi \qquad \theta,\, \Delta}{\Gamma,\, \varphi \wedge \theta,\, \Delta}\,\wedge \qquad \frac{\Gamma,\, \varphi\{x{\mapsto}t\}}{\Gamma,\, \exists x\, \varphi}\,\exists \qquad \frac{\Gamma,\, \varphi}{\Gamma,\, \forall x\, \varphi}\,\forall \quad (x \text{ not free in } \Gamma)$$

11.1 Soundness and completeness of F2

In this section we prove that F2 is sound and complete. The completeness of F2 is a key step towards showing that combinatorial proofs are complete in §11. The ***size*** $\#\Pi$ of a proof Π is the number of occurrences of rules in Π.

LEMMA 11.2. Every F1 proof Π of $\Gamma, \forall x\, \varphi$ can be transformed into an F1 proof Π' of $\Gamma, \forall x\, \varphi$ such that $\#\Pi' \leqslant \#\Pi$ and the last rule of Π' takes one of the following two forms:

$$\frac{\Gamma,\, \varphi}{\Gamma,\, \forall x\, \varphi}\,\forall \qquad\qquad \frac{\Gamma}{\Gamma,\, \forall x\, \varphi}\,\mathsf{w}$$

Proof. By induction on $\#\Pi$, transposing $\forall$ rules down through non-$\forall$ rules. The only non-trivial case is when Π ends with a c rule, so Π has the form

$$\begin{array}{c}\vdots\;\Pi_1\\ \dfrac{\Gamma, \forall x\,\varphi, \forall x\,\varphi}{\Gamma, \forall x\,\varphi}\,\mathsf{c}\end{array}$$

for Π_1 the $\mathsf{F1}$ proof of $\Gamma, \forall x\,\varphi, \forall x\,\varphi$ obtained by deleting the last rule from Π. Since Π_1 has one less rule than Π we can apply induction to obtain either $\Pi_\forall$ ending in a $\forall$ rule or Π_{w} ending in a w rule:

$$\begin{array}{c}\vdots\;\Pi_\forall\\ \dfrac{\dfrac{\Gamma, \forall x\,\varphi, \varphi}{\Gamma, \forall x\,\varphi, \forall x\,\varphi}\,\forall}{\Gamma, \forall x\,\varphi}\,\mathsf{c}\end{array} \qquad\qquad \begin{array}{c}\vdots\;\Pi_{\mathsf{w}}\\ \dfrac{\dfrac{\Gamma, \forall x\,\varphi}{\Gamma, \forall x\,\varphi, \forall x\,\varphi}\,\mathsf{w}}{\Gamma, \forall x\,\varphi}\,\mathsf{c}\end{array}$$

In the latter case we take Π' to be Π_{w}. In the former case, we once again apply induction, and obtain one of the following two proofs:

$$\begin{array}{c}\vdots\;\Pi_{\forall\forall}\\ \dfrac{\dfrac{\dfrac{\Gamma, \varphi, \varphi}{\Gamma, \forall x\,\varphi, \varphi}\,\forall}{\Gamma, \forall x\,\varphi, \forall x\,\varphi}\,\forall}{\Gamma, \forall x\,\varphi}\,\mathsf{c}\end{array} \qquad\qquad \begin{array}{c}\vdots\;\Pi_{\forall\mathsf{w}}\\ \dfrac{\dfrac{\dfrac{\Gamma, \varphi}{\Gamma, \forall x\,\varphi, \varphi}\,\mathsf{w}}{\Gamma, \forall x\,\varphi, \forall x\,\varphi}\,\forall}{\Gamma, \forall x\,\varphi}\,\mathsf{c}\end{array}$$

Define Π', respectively, as:

$$\begin{array}{c}\vdots\;\Pi_{\forall\forall}\\ \dfrac{\dfrac{\Gamma, \varphi, \varphi}{\Gamma, \varphi}\,\mathsf{c}}{\Gamma, \forall x\,\varphi}\,\forall\end{array} \qquad\qquad \begin{array}{c}\vdots\;\Pi_{\forall\mathsf{w}}\\ \dfrac{\Gamma, \varphi}{\Gamma, \forall x\,\varphi}\,\forall\end{array}$$

□

LEMMA 11.3. Every $\mathsf{F1}$ proof Π of $\Gamma, \varphi \vee \theta$ can be transformed into an $\mathsf{F1}$ proof Π' of $\Gamma, \varphi \vee \theta$ such that $\#\Pi' \leqslant \#\Pi$ and the last rule of Π' takes one of the following two forms:

$$\dfrac{\Gamma, \varphi, \theta}{\Gamma, \varphi \vee \theta}\,\vee \qquad\qquad \dfrac{\Gamma}{\Gamma, \varphi \vee \theta}\,\mathsf{w}$$

Proof. Reason as in the proof of Lemma 11.2, but with a $\vee$-rule instead of a $\forall$-rule. □

A c rule $\frac{\Gamma, \chi[\varphi \vee \varphi]}{\Gamma, \chi[\varphi]}\mathsf{c}$ is ***connected*** if the formula φ is connected.

LEMMA 11.4. Every $\mathsf{F1}$ proof Π of a sequent Γ can be transformed into an $\mathsf{F1}$ proof Π' of Γ in which every c rule is connected.

Proof. By induction on $\#\Pi$. In the base case $\#\Pi = 1$, the proof Π is just an axiom $\overline{\alpha, \overline{\alpha}}$, so we take $\Pi' = \Pi$. Otherwise, if the last rule ρ of Π is not a c rule, or is a connected c rule, we appeal to induction with the result Π_1 of deleting ρ from Π to obtain Π_1', then append ρ to Π_1' to construct Π'. The key case is when ρ is a non-connected c rule. Thus ρ takes one of the following two forms:

$$\dfrac{\Gamma, \forall x\,\varphi, \forall x\,\varphi}{\Gamma, \forall x\,\varphi}\,\mathsf{c} \qquad\qquad \dfrac{\Gamma, \varphi \vee \theta, \varphi \vee \theta}{\Gamma, \varphi \vee \theta}\,\mathsf{c}$$

In the former case we apply Lemma 11.2 to Π to obtain Π_0 ending with a $\forall$ or w rule, and in the latter case we apply Lemma 11.3 to Π to obtain Π_0 ending with a $\vee$ or w rule, and $\#\Pi_0 \leqslant \#\Pi$. We can now apply our earlier induction step to Π_0 since its last rule is a $\forall$, $\vee$ or w rule (not a c rule). □

Define $\mathsf{F2w}$ as the extension of $\mathsf{F2}$ with the w rule of $\mathsf{F1}$.

LEMMA 11.5. Every $\mathsf{F1}$ proof Π of a sequent Γ can be transformed into an $\mathsf{F2w}$ proof Π' of Γ.

Proof. Apply Lemma 11.4 to Π to obtain an $\mathsf{F1}$ proof Π_0 in which every c rule is connected. Construct Π' from Π_0 by replacing every c rule of the form below-left in Π by the pair of rules below-right.

$$\frac{\Gamma,\,\varphi,\,\varphi}{\Gamma,\,\varphi}\,\mathsf{c} \qquad\qquad \frac{\dfrac{\Gamma,\,\varphi,\,\varphi}{\Gamma,\,\varphi\vee\varphi}\,\vee}{\Gamma,\,\varphi}\,\mathsf{C}$$

The C rule is well-defined since every c rule in Π_0 is connected. (Note that this instance of a C rule is the special case with trivial context $\chi = \_$.) $\square$

LEMMA 11.6. Every $\mathsf{F2w}$ proof Π of a sequent Γ can be transformed into an $\mathsf{F2w}$ proof Π' of Γ in which every w rule is below every non-w rule.

Proof. We exhaustively transpose every w rule ρ of Π down through any immediately-following non-w rule ρ'. For example, here are the two cases where ρ' is a $\wedge$ rule,

$$\frac{\dfrac{\Gamma,\,\theta}{\Gamma,\,\varphi,\,\theta}\,\mathsf{w} \qquad \psi,\,\Delta}{\Gamma,\,\varphi,\,\theta\wedge\psi,\,\Delta}\,\wedge \quad\to\quad \frac{\dfrac{\Gamma,\,\theta \qquad \psi,\,\Delta}{\Gamma,\,\theta\wedge\psi,\,\Delta}\,\wedge}{\Gamma,\,\varphi,\,\theta\wedge\psi,\,\Delta}\,\mathsf{w} \qquad\qquad \frac{\dfrac{\Gamma,\,\theta}{\Gamma,\,\theta,\,\varphi}\,\mathsf{w} \qquad \psi,\,\Delta}{\Gamma,\,\theta,\,\varphi\wedge\psi,\,\Delta}\,\wedge \quad\to\quad \frac{\Gamma,\,\theta}{\Gamma,\,\theta,\,\varphi\wedge\psi,\,\Delta}\,\mathsf{w}^+$$

where w^+ denotes a sequence of $n+1$ instances of the w rule and n is the number of formulas occurring in Δ. All other transpositions are similar. $\square$

LEMMA 11.7. Every $\mathsf{F2w}$ proof Π of a formula φ can be transformed into an $\mathsf{F2}$ proof Π' of φ.

Proof. By Lemma 11.6 we obtain from Π an $\mathsf{F2w}$ proof Π' of φ in which every w rule is below every non-w rule. This proof has zero w rules, since if the last rule were a w rule then φ would be a sequent with at least two formulas, rather than a single formula φ. Thus Π' is a $\mathsf{F2}$ proof. $\square$

LEMMA 11.8. Every $\mathsf{F1}$ proof Π of a formula φ can be transformed into an $\mathsf{F2}$ proof Π' of φ.

Proof. By Lemma 11.5 we obtain an $\mathsf{F2w}$ proof of φ, whence an $\mathsf{F2}$ proof by Lemma 11.7. $\square$

LEMMA 11.9 ($\mathsf{F2}$ soundness & completeness). A formula is valid if and only if it has an $\mathsf{F2}$ proof.

Proof. $\mathsf{F2}$ is sound since every rule is sound: the conclusion is valid whenever the hypothesis/hypotheses is/are valid. Completeness of $\mathsf{F2}$ follows from the completeness of $\mathsf{F1}$, by Lemma 11.8. $\square$

Note that $\mathsf{F2}$ is complete for formulas, but not for sequents in general. For example, $p\vee\overline{p}$ has an $\mathsf{F2}$ proof, but $p\vee\overline{p}, q$ does not; however $(p\vee\overline{p})\vee q$ does.[14]

11.2 Preliminaries to the proof of the Completeness Theorem

Recall from §10.2 that a formula is *extruded* if no subformula is an intrusion, where an *intrusion* is a formula of the form $\varphi\vee\forall x\,\theta$, $(\forall x\,\theta)\vee\varphi$, $\varphi\wedge\exists x\,\theta$, or $(\exists x\,\theta)\wedge\varphi$, for any variable x and formulas φ and θ. Define a formula φ as ***extrudable*** if no variable occurs both free and bound in φ. The ***extruded form*** of an extrudable formula is the extruded formula which results from exhaustively applying the following ***extrusion*** subformula rewrites:

$$\varphi\wedge\exists x\,\theta \;\to\; \exists x\,(\varphi\wedge\theta) \qquad\qquad (\exists x\,\theta)\wedge\varphi \;\to\; \exists x\,(\theta\wedge\varphi)$$

$$\varphi\vee\forall x\,\theta \;\to\; \forall x\,(\varphi\vee\theta) \qquad\qquad (\forall x\,\theta)\vee\varphi \;\to\; \forall x\,(\theta\vee\varphi)$$

[14] See [Hug10] for an analysis of sequent calculi which are complete for formulas but not for sequents.

Note that these rewrites are not applicable if x occurs free in φ, due to the unintended capture of x by the quantifier in the output of the rewrite; hence our restriction of these rewrites to extrudable formulas.

Recall from §10.2 that a formula is *unambiguous* if no x-quantifier is in the scope of another x-quantifier, for every variable x. Define a formula as ***pristine*** if it is extrudable and its extruded form is unambiguous. The ***graph*** $\mathbb{G}(\varphi)$ of a pristine formula φ is the graph $\mathbb{G}(\varphi')$ of its extruded form φ', as defined in Def. 10.24. Define a sequent $\Gamma = \varphi_1, \ldots, \varphi_n$ as pristine if its formula $\Phi(\Gamma) = \varphi_1 \vee (\varphi_2 \vee \ldots (\varphi_{n-1} \vee \varphi_n) \ldots)$ is pristine, and define the graph of a pristine sequent Γ as the graph $\mathbb{G}(\Phi(\Gamma))$ of its (pristine) formula.

LEMMA 11.10. If the conclusion sequent below a W or C rule is pristine, the hypothesis sequent above the rule is pristine.

Proof. Going upwards from the pristine conclusion $\Gamma, \chi[\varphi]$ to the hypothesis $\Gamma, \chi[\varphi \vee \varphi]$, the C rule replaces a connected subformula φ by $\varphi \vee \varphi$. Since the free and bound variables of $\varphi \vee \varphi$ are those of φ, the hypothesis $\Gamma, \chi[\varphi \vee \varphi]$ is extrudable. Since φ is connected, the extruded form φ' of φ cannot be a $\forall$ or $\vee$ formula, thus the extruded form of $\varphi \vee \varphi$ is $\varphi' \vee \varphi'$, so the extruded form of $\Gamma, \chi[\varphi \vee \varphi]$ is unambiguous, and $\Gamma, \chi[\varphi \vee \varphi]$ is pristine.

Going upwards from the pristine conclusion $\Gamma, \chi[\varphi \vee \theta]$ to the hypothesis $\Gamma, \chi[\varphi]$, the W rule replaces a subformula $\varphi \vee \theta$ by φ. Since this amounts to the deletion of the subformula θ, the sequent $\Gamma, \chi[\varphi]$ is pristine: every variable that occurs free and bound in $\Gamma, \chi[\varphi]$ also occurs free and bound in $\Gamma, \chi[\varphi \vee \theta]$, and any ambiguity in the extruded form of $\Gamma, \chi[\varphi]$ would yield an ambiguity in the extruded form of $\Gamma, \chi[\varphi \vee \theta]$. ☐

The rules W and C of F2 are ***structural***, and the remaining rules are ***logical***.

DEFINITION 11.11. A proof of F2 is ***phased*** if every logical rule is above every structural rule.

LEMMA 11.12. Every F2 proof of a formula φ can be transformed into a phased proof of φ.

Proof. A routine induction on the size of the proof, by exhaustively commuting structural rules downwards through logical rules. ☐

A proof is ***logical*** if every rule is logical. The ***heart*** of a phased F2 proof is the formula(-occurrence) which is the hypothesis of the first structural rule, the ***logical phase*** is the logical subproof ending at the heart, and the rules below the heart constitute the ***structural phase***.

A sequent is rectified if all bound variables are distinct from one another and from all free variables. (Thus a sequent Γ is rectified if and only if its formula $\Phi(\Gamma)$ is rectified.)

LEMMA 11.13. Every F2 proof Π of a logical proof of a sequent Γ with rectified form Γ' can be transformed into a logical proof Π' of Γ', such that $\#\Pi' = \#\Pi$ and every sequent in Π' is rectified.

Proof. By induction on $\#\Pi$. The base case with Π just an axiom $\overline{\alpha, \overline{\alpha}}$ is immediate. For the induction step, let σ be the sequence of renamings of occurrences of bound variables in Γ which yields the rectified form Γ'. Since the last rule ρ of Π is not a C rule, applying σ to each hypothesis sequent Γ_i above ρ yields a rectified sequent Γ'_i. Apply the induction hypothesis to the subproofs above each Γ_i. ☐

Define the graph $\mathcal{G}(\Gamma)$ of a rectified sequent $\Gamma = \varphi_1, \ldots, \varphi_n$ as the rectified fograph $\mathcal{G}(\Phi(\Gamma)) = \mathcal{G}(\varphi_1 \vee (\varphi_2 \vee \ldots (\varphi_{n-1} \vee \varphi_n) \ldots))$.

LEMMA 11.14. Each non-axiom logical rule of F2 with rectified hypothesis sequent(s) Γ_i and rectified conclusion sequent Γ determines a fograph operation or identity which constructs a rectified fograph $\mathcal{G}(\Gamma)$ from the rectified fograph(s) $\mathcal{G}(\Gamma_i)$.

Proof. The $\wedge$ rule determines the fusion (Def. 10.4) of $\mathcal{G}(\Gamma, \varphi)$ and $\mathcal{G}(\theta, \Delta)$ to produce $\mathcal{G}(\Gamma, \varphi \wedge \theta, \Delta)$. The $\vee$ rule acts as the identity on $\mathcal{G}(\Gamma, \varphi, \theta) = \mathcal{G}(\Gamma, \varphi \vee \theta)$, The $\forall$ rule applies universal quantification (Def. 10.7) to $\mathcal{G}(\Gamma, \varphi)$ to form $\mathcal{G}(\Gamma, \forall x\, \varphi) = \forall x\, \mathcal{G}(\Gamma)$. The $\exists$ rule applies existential quantification (Def. 10.10) to $\Gamma, \varphi\{x \mapsto t\}$ to form $\mathcal{G}(\Gamma, \forall x\, \varphi) = \forall x\, \mathcal{G}(\Gamma)$, ☐

LEMMA 11.15. Every logical F2 proof of a rectified formula φ determines a fonet with underlying fograph $\mathcal{G}(\varphi)$.

Proof. Each axiom rule determines a fonet axiom (as defined in §10.1). By Lemma 11.14 every rule determines a fograph operation, which by Lemma 10.15 construct a fonet from the fonet axioms. □

LEMMA 11.16. Every structural rule whose hypothesis sequent Γ and conclusion sequent Δ are pristine determines a structural map $\mathbb{G}(\Gamma) \to \mathbb{G}(\Delta)$.

Proof. A C rule determines a contraction fograph map and a W determines a weakening fograph map (Def. 10.35). □

11.3 Proof of the Completeness Theorem

Proof of Completeness Theorem, Theorem 6.3. Let φ be a valid rectified formula. We will construct a combinatorial proof $f : N \to \mathcal{G}(\varphi)$.

By Lemma 11.9 there exists an F2 proof Π of φ. By Lemma 11.12 we may assume Π is a phased. Let Π_L be the logical phase of Π, culminating in the heart μ of Π. By Lemma 11.13 there exists a logical proof Π'_L of a rectified form μ' of μ. By Lemma 11.2, Π'_L determines a fonet N' whose underlying rectified fograph is $\mathcal{G}(\mu')$.

The structural phase Π_S of Π is a sequence of C and W rules beginning with μ and ending with φ. Since φ is rectified, it is pristine. Thus by iterating Lemma 11.10 upwards from φ, every formula occurring in Π_S is pristine. By Lemma 11.16 we obtain a sequence of contraction and weakening maps, one per rule of Π_S, whose composite is a structural map $f : \mathbb{G}(\mu) \to \mathbb{G}(\varphi)$, since structural maps are (by definition) closed under composition. By Lemma 10.38 f is a skew bifibration. Since φ is rectified, $\mathbb{G}(\varphi) = \mathcal{G}(\varphi)$, and since μ' is a rectified form of μ, the underlying fograph of the fonet N' is a rectified form of $\mathbb{G}(\mu)$. Thus applying the colouring of N' to $\mathbb{G}(\mu)$ yields a fonet N, so f is a skew bifibration $N \to \mathcal{G}(\varphi)$. □

12 Homogeneous soundness and completeness proofs

12.1 Propositional homogeneous soundness and completeness proof

In this section we prove the propositional homogeneous soundness and completeness theorem, Theorem 7.6. We begin by observing that the function $\mathcal{D}$ from propositions to dualizing graphs (Def. 7.2) factorizes through propositional fographs. A fograph is ***propositional*** if every predicate symbol is nullary. For example, the middle row of Fig. 7 (p. 7) shows four propositional fographs.

DEFINITION 12.1. The ***dualizing graph*** $\mathbb{D}(G)$ of a propositional fograph is the dualizing graph D with $V_D = V_G$, $E_D = E_G$, and $vw \in \perp_D$ if and only if v and w have dual predicate symbols.[15]

For example, for each propositional fograph G in the middle row of Fig. 7 (p. 7), the corresponding dualizing graph $\mathbb{D}(G)$ is shown below G.

LEMMA 12.2. $\mathbb{D}(G)$ is a well-defined dualizing graph for every propositional fograph G.

Proof. Let $D = \mathbb{D}(G)$. Since $V_D = V_G$, $E_D = E_G$ and G is a fograph, D is a cograph. By reasoning as in the proof of Lemma 7.3, $(V_G, \perp_D)$ is P_4- and C_3-free. □

LEMMA 12.3. The function $\mathcal{D}$ from propositions to dualizing graphs (Def. 7.2) factorizes through propositional fographs: $\mathcal{D}(\varphi) = \mathbb{D}(\mathcal{G}(\varphi))$ for every proposition φ.

Proof. A routine induction on the structure of φ. □

LEMMA 12.4 (Propositional homogeneous soundness). A proposition is valid if it has a homogeneous combinatorial proof.

[15] In §13 we will show that $\mathbb{D}$ is a surjection from propositional fographs onto dualizing graphs (Lemma 13.1).

Proof. Suppose $f : N \to \mathcal{D}(\varphi) = D$ is a homogeneous combinatorial proof of the proposition φ. By Lemma 12.3, $D = \mathbb{D}(\mathcal{G}(\varphi))$. Define N' as the cograph (V_N, E_N) with a link $\{v, w\}$ for each $vw \in \perp_N$ and the label of a vertex v in N' defined as the label of $f(v)$ in $\mathcal{G}(\varphi)$, where $f(v) \in V_D$ can be viewed as a vertex of $\mathcal{G}(\varphi)$ since $V_D = V_{\mathcal{G}(\varphi)}$ es by definition of $\mathbb{D}$. Since N is a dualizing net, N' is a fonet: (a) the predicate symbols on the atoms of every link are dual, since $f : (V_N, \perp_N) \to (V_D, \perp_D)$ is an undirected graph homomorphism, (b) N' trivially has a dualizer, the empty assignment, since it is propositional, with no existential variables, and (c) N' has no induced bimatching, since the leap graphs $\mathcal{L}_{N'}$ and $\mathcal{L}_N$ are equal and N has no induced bimatching. We claim that $f : N' \to \mathcal{G}(\varphi)$ is a skew bifibration. Since $f : N \to D$ is a skew fibration, $(V_{N'}, E_{N'}) = (V_N, E_N)$, and $(V_{\mathcal{G}(\varphi)}, E_{\mathcal{G}(\varphi)}) = (V_D, E_D)$, we know $f : N' \to \mathcal{G}(\varphi)$ is a skew fibration. Because the label of v in N' is that of $f(v)$ in $\mathcal{G}(\varphi)$, $f : N' \to \mathcal{G}(\varphi)$ preserves labels. Since there are no binders, existentials are preserved trivially and $f : \vec{N}' \to \vec{\mathcal{G}}(\varphi)$ is trivially a directed graph fibration. Thus $f : N' \to \mathcal{G}(\varphi)$ is a skew bifibration, hence a combinatorial proof (since N' is a fonet). By Theorem 6.2, φ is valid. ☐

LEMMA 12.5 (Propositional homogeneous completeness). Every valid proposition has a homogeneous combinatorial proof.

Proof. Let φ be a valid proposition. By Theorem 6.3 there exists a (standard) combinatorial proof $f : N \to \mathcal{G}(\varphi)$. Let N' be the dualizing graph obtained from N by replacing each link (colour) $\{v, w\}$ by a duality $vw \in \perp_{N'}$. Since, by definition of a linked fograph, every literal is in exactly one link, $(V_{N'}, \perp_{N'})$ is a matching, and since N is a propositional fonet, N' has no induced bimatching; thus N' is a dualizing net. Let $D = \mathcal{D}(\varphi)$. By Lemma 12.3, $\mathcal{D}(\varphi) = \mathbb{D}(\mathcal{G}(\varphi))$, thus $(V_D, E_D) = (V_{\mathcal{G}(\varphi)}, E_{\mathcal{G}(\varphi)})$. We claim that $f : N' \to D$ is a homogeneous combinatorial proof, *i.e.*, (1) $f : N' \to D$ is a skew fibration and (2) $f : (V_{N'}, \perp_{N'}) \to (V_D, \perp_D)$ is a graph homomorphism. By definition, (1) holds if $f : (V_{N'}, E_{N'}) \to (V_D, E_D)$ is a skew fibration, which is true because f is a skew bifibration, $(V_{N'}, E_{N'}) = (V_N, E_N)$, and $(V_D, E_D) = (V_{\mathcal{G}(\varphi)}, E_{\mathcal{G}(\varphi)})$. For (2), suppose $vw \in \perp_{N'}$. Since N is a fonet, it has a dualizer, so the labels of v and w are dual, say, p and $\overline{p}$, respectively. Because f preserves labels, $f(v)$ and $f(w)$ are labelled p and $\overline{p}$, thus $f(v)\,f(w) \in \perp_D$, and (2) holds. ☐

Proof of Theorem 7.6 (Propositional homogeneous soundness and completeness). Lemmas 12.4 and 12.5. ☐

12.2 Monadic homogeneous soundness and completeness proof

In this section we prove the monadic homogeneous soundness and completeness theorem, Theorem 8.11. The proof of completeness is similar to that of the propositional case, Lemma 12.5: transform a standard first-order combinatorial proof of a monadic formula into a homogeneous combinatorial proof. The proof of soundness is more subtle. In the propositional case, Lemma 12.4, we transformed a homogeneous combinatorial proof directly into a standard one, with the same vertices in both source and target. The monadic case involves quotienting indistinguishable vertices in the source net.

12.2.1 Factorization through closed monadic fographs

A fograph is ***closed*** if it contains no free variables, and ***monadic*** if its predicate symbols are unary and it has no function symbols.

DEFINITION 12.6. The ***mograph*** $\mathbb{M}(G)$ of a closed monadic fograph G is the mograph M with

- $V_M = V_G$,
- $E_M = E_G$,
- $vw \in \perp_M$ if and only if v and w are literals whose predicate symbols are dual, and
- $\langle v, w \rangle \in B_M$ if and only if v binds w.

For example, the closed monadic fograph G in the centre of Fig. 9 (p. 9)has the mograph $\mathbb{M}(G)$ to its right.

LEMMA 12.7. $\mathbb{M}(G)$ is a well-defined mograph for every closed monadic fograph G.

Proof. The underlying cograph (V_M, E_M) is inherited directly from G. By reasoning as in the proof of Lemma 7.3, $(V_G, \perp_D)$ is P_4- and C_3-free. It remains to show (a) every target of a binding in B_M is in no other binding, (b) no binder is in a duality, (c) the scope of every binder b is non-empty, and (d) $\langle b, l\rangle \in B_M$ only if l is in the scope of b.

(a) Since G is monadic, every literal label contains exactly one variable, hence is bound by at most one binder in G. By definition of $\vec{G}$, no literal binds any other vertex, thus every literal target of a binding is in no other binding.

(b) Dualities are defined as pairs of literals in G, which become literals in M since G is closed. (Every literal in G is bound by a binder in G, so becomes a literal in M.)

(d) By definition of fograph binding, l is bound by a binder b only if l is in the scope of b. $\square$

LEMMA 12.8. The function $\mathcal{M}$ from closed monadic formulas to mographs (Def. 8.3) factorizes through closed monadic fographs: $\mathcal{M}(\varphi) = \mathbb{M}(\mathcal{G}(\varphi))$ for every closed monadic formula φ.

Proof. A routine induction on the structure of φ. $\square$

12.2.2 Collapsing indistinguishable vacuous universal binders

Given an equivalence relation $\sim$ on a set V write $[v]_\sim$ for the $\sim$-equivalence class $\{\, w \in V : w \sim v \,\}$ and $V/\!\sim$ for the set of $\sim$-equivalence classes $\{\, [v]_\sim : v \in V \,\}$. For a set E of edges on V define $E/\!\sim$ as the set $\{\, [v]_\sim [w]_\sim : v, w \in E \,\}$ of edges on $V/\!\sim$. Given a mograph M and an equivalence relation $\sim$ on V_M define the ***quotient*** mograph $M/\!\sim$ by $V_{M/\sim} = V_M/\!\sim$, $E_{M/\sim} = E_M/\!\sim$, $\perp_{M/\sim} = \perp_M/\!\sim$, and $B_{M/\sim} = B_M/\!\sim$.

A binder in a mograph is ***vacuous*** if it binds no literal. Let $f : N \to M$ be a skew bifibration of mographs. Vacuous universal binders b, c in N are ***indistinguishable*** if their images and neighbourhoods are equal, *i.e.*, $f(b) = f(c)$ and $N(b) = N(c)$. Define $\asymp$ as the equivalence relation on V_N generated by indistinguishability, and the ***collapse*** $f_\asymp : N/\!\asymp \to M$ as the canonical function on the quotient, *i.e.*, $f_\asymp([b]_\asymp) = f(v)$, a well-defined function since $b \asymp c$ implies $f(b) = f(c)$.

LEMMA 12.9. Let M be a mograph and N a monet. If $f : N \to M$ is a homogeneous combinatorial proof then its collapse $f_\asymp : N/\!\asymp \to H$ is a homogeneous combinatorial proof.

Proof. $N/\!\asymp$ is a monet because if $W \subseteq V_{N/\asymp}$ induces a bimatching in $N/\!\asymp$ then it induces a bimatching in N: since indistinguishable vertices are vacuous binders, they cannot be in both a leap and an edge of $E_N/\!\asymp$, so cannot occur in W. $f_\asymp : (V_N/\!\asymp, E_N/\!\asymp) \to (V_M, E_M)$ is a skew fibration because $f : (V_N, E_N) \to (V_M, E_M)$ is a skew fibration and indistinguishable vertices have the same image and neighbourhood. $f_\asymp : (V_N/\!\asymp, \perp N/\!\asymp) \to (V_M, \perp_M)$ is a homomorphism because $f : (V_N, E_N) \to (V_M, \perp_M)$ is a homomorphism and no binder is in a duality edge. $f_\asymp : (V_N/\!\asymp, B_N/\!\asymp) \to (V_M, B_M)$ is a fibration because $f : (V_N, E_N) \to (V_M, E_M)$ is a fibration and indistinguishable binders are vacuous, hence are absent from bindings. $\square$

12.2.3 Monadic fonets without dualizers

Monets were defined (§8.1) without need for dualizers, in terms of the binder equivalence relation $\simeq_M$. In this section we take an analogous approach with monadic fonets (§5).

Let ***rmf*** abbreviate *rectified monadic fograph*. Two atoms are ***pre-dual*** if their predicate symbols are dual (*e.g.* px and $\bar{p}y$), and two literals are pre-dual if their atom labels are pre-dual.

DEFINITION 12.10. Let K be a linked rmf whose links are pre-dual. ***Variable equivalence*** $\simeq_K$ is the equivalence relation on binders generated by $x \simeq_K y$ for each link $\{\, \bullet px, \bullet \bar{p}y \,\}$ in K.

In the above definition p is any predicate symbol (necessarily unary, since K is monadic).

A ***conflict*** in K is a pair $\{\, x, y \,\}$ of distinct non-existential variables x and y such that $x \simeq_K y$.

DEFINITION 12.11. A linked rmf is ***consistent*** if its links are pre-dual and it has no conflict.

LEMMA 12.12. A linked rmf has a dualizer if and only if it is consistent.

Proof. Let K be the linked rmf.

Suppose K has a dualizer. By Lemma 10.12 K has a most general dualizer δ. Thus for every link $\{\circ px, \circ qy\}$ we have $(px)\delta$ dual to $(qy)\delta$. (Recall that $\alpha\delta$ denotes the result of substituting $\delta(x)$ for x throughout α, simultaneously for each x.) Therefore $q = \overline{p}$ so $\{\circ px, \circ qy\}$ is pre-dual. For a contradiction, suppose $\{z_1, z_2\}$ were a conflict in K, *i.e.*, $z_1 \simeq_K z_2$ for non-existential variables $z_1 \neq z_2$. Since $z_1 \simeq_K z_2$ we have variables $x_1, \ldots, x_n$ for $n \geqslant 1$ with $x_1 = z_1$, $x_n = z_2$, and for $1 \leqslant i < n$ there exists a link $\{\bullet p_i x_i, \bullet \overline{p}_i x_{i+1}\}$ for some predicate symbol p_i. Since δ is a dualizer we have $(p_i x_i)\delta$ dual to $(\overline{p}_i x_{i+1})\delta$, so $x_i\delta = x_{i+1}\delta$. Thus $x_1\delta = x_n\delta$ so $z_1\delta = z_2\delta$. Since z_1 and z_2 are non-existential, we have $z_1\delta = z_1$ and $z_2\delta = z_2$, hence $z_1 = z_2$, contradicting $z_1 \neq z_2$.

Conversely, suppose K is consistent. Let $e_1, \ldots, e_n$ be the equivalence classes of $\simeq_K$. Define y_i as the unique non-existential variable in e_i, if it exists (necessarily unique since K is consistent), and otherwise define y_i as a fresh variable, where *fresh* means not in K and distinct from y_j for $1 \leqslant j < i$. Given an existential variable x, define $\delta(x) = y_i$ if e_i is the equivalence class containing x_i.

We must show that for every link $\{\circ px, \circ qy\}$ in K we have $(px)\delta$ dual to $(\overline{p}y)\delta$. Since K is consistent, its links are pre-dual, hence $q = \overline{p}$. Thus it remains to show that $x\delta = y\delta$. Since x and y are in the same link, they are in the same equivalence class e_i (for some i). We consider three cases.

1. Both x and y are existential. Since x and y are in e_i, we have $\delta(x) = \delta(y) = y_i$.

2. Both x and y are non-existential. Therefore $x\delta = x$ and $y\delta = y$, so we require $x = y$. This holds because $x \neq y$ would imply that $\{x, y\}$ is a conflict, contradicting the consistency of K.

3. Exactly one of x and y is existential, say x. Since y is non-existential, $y\delta = y$, and y is the unique y_i non-existential variable in e_i. Since x is also in e_i, we have $\delta(x) = y_i$. □

LEMMA 12.13. Let $\bullet x$ and $\bullet y$ be binders in a consistent linked rmf K, with $\bullet x$ existential and $\bullet y$ universal. The pair $\{\bullet x, \bullet y\}$ is a dependency of K if and only if $x \simeq_K y$.

Proof. By Lemma 10.13, the dependencies of K are those of a most general dualizer δ, so it suffices to show that $x \simeq_K y$ if and only if $\delta(x) = y$. Since every predicate symbol in K is unary, $\simeq_K$ is the transitive closure of the unification problem $\approx_K$ (see the proof of Lemma 10.12). Thus the dualizer δ defined in the proof of Lemma 12.12 is most general, and by construction $x \simeq_K y$ if and only if $\delta(x) = y$. □

Note that the above lemmas simplify the definition of (standard, non-homogeneous) monadic combinatorial proof $f : K \to G$:

- Instead of checking for the existence of a dualizer for (the rectified form of) K, we merely check that K is consistent, via the variable relation $\simeq_K$, using Lemma 12.12.

- Instead of building the leap graph $\mathcal{L}_K$ with dependencies via a dualizer, we read dependencies directly from $\simeq_K$, using Lemma 12.13.

12.2.4 The linked mograph of a linked closed monadic fograph

DEFINITION 12.14. The ***linked mograph*** $\wedge(K)$ of a linked closed monadic fograph K is the linked mograph M with

- $V_M = V_K$,

- $E_M = E_K$,

- $vw \in \perp_M$ if and only if $\{v, w\}$ is a link

- $\langle v, w\rangle \in B_M$ if and only if v binds w.

LEMMA 12.15. $\wedge(K)$ is a well-defined linked mograph for every linked closed monadic fograph K.

Proof. Let $\lambda_1, \ldots, \lambda_n$ be the links of K for $\lambda_i = \{ l_i, k_i \}$. Choose distinct predicate symbols $p_1, \ldots, p_n$, and define K' by replacing the predicate symbols in the labels of l_i and k_i by p_i and $\overline{p}_i$, respectively. By construction, $\Lambda(K') = \Lambda(K)$, and since two literals in K' are pre-dual in K' if and only if they constitute a link, we have $\Lambda(K') = \mathbb{M}(K')$. Thus $\Lambda(K) = \mathbb{M}(K')$ which is a well-defined mograph by Lemma 12.7. Since the p_i are distinct, every literal of K' is in a unique duality, so K' is linked. ☐

LEMMA 12.16. A linked closed monadic fograph K is a fonet if and only if its linked mograph $\Lambda(K)$ is a monet.

Proof. Without loss of generality we may assume K is rectified. By Lemma 12.12, K has a dualizer if and only if it is consistent in the sense of Def. 12.11, and consistency of K coincides with consistency of $\Lambda(K)$ (Def. 8.6). By Lemma 12.13 the dependencies of K are those pairs $\{ x, y \}$ of variables with x existential, y universal and $x \simeq_K y$, which, by definition of Λ, correspond to pairs $\{ b_x, b_y \}$ of binders in $\Lambda(K)$ with b_x and b_y the unique binders corresponding to the variables x and y, and $b_x \simeq_{\Lambda(K)} b_y$. Thus the leap graphs of K and $\Lambda(K)$ are the same, so K has an induced bimatching if and only if $\Lambda(K)$ has an induced bimatching. ☐

12.2.5 Proof of monadic homogeneous combinatorial soundness

Recall that, by definition of $\mathbb{M}$, $V_{\mathbb{M}(G)} = V_G$ for every closed monadic fograph G.

LEMMA 12.17. Let G be a closed monadic fograph. A vertex is a literal in G if and only if it is a literal in the mograph $\mathbb{M}(G)$.

Proof. Immediate from the definition of the binding set $B_{\mathbb{M}(G)}$ (Def. 12.6) and that, by definition, a vertex is a literal in a mograph if and only if it is the target of a binding. ☐

LEMMA 12.18. Let G be a closed monadic fograph. A binder is universal in G if and only if it is universal in the mograph $\mathbb{M}(G)$.

Proof. By definition $(V_G, E_G) = (V_{\mathbb{M}(G)}, E_{\mathbb{M}(G)})$, and in both cases, a binder is universal if and only if its scope contains no edge. ☐

Define the ***type*** $\mathsf{typ}_G(v) \in \{ *, \forall, \exists \}$ of a vertex v in a mograph or fograph G as $*$ if v is a literal, $\forall$ if v is a universal binder, and $\exists$ if v is an existential binder.

LEMMA 12.19. For every closed monadic fograph G, $\mathsf{typ}_G(v) = \mathsf{typ}_{\mathbb{M}(G)}(v)$ for every vertex v.

Proof. Lemmas 12.17 and 12.18. ☐

LEMMA 12.20. Every mograph skew bifibration $f : N \to M$ preserves vertex type, *i.e.*, $\mathsf{typ}_N(v) = \mathsf{typ}_M(f(v))$ for every vertex v in N.

Proof. A vertex is a literal if and only if it is the target of a binding, and since $f : (V_N, B_N) \to (V_M, B_M)$ is a fibration, a vertex v in V_N is the target of a binding if and only if $f(v)$ is the target of a binding. Thus f maps literals to literals and binders to binders. By definition (Def. 8.9) a skew bifibration maps existential binders to existential binders, so it remains to show that universal binders map to universal binders. This follows from the proof of Lemma 10.65, which applies in the homogeneous setting because it does not depend on labels. ☐

LEMMA 12.21 (Monadic homogeneous soundness). A closed monadic formula is valid if it has a homogeneous combinatorial proof.

Proof. Suppose $f : N \to \mathcal{M}(\varphi) = M$ is a homogeneous combinatorial proof of the proposition φ. Without loss of generality, we may assume f is collapsed, by Lemma 12.9. Define N' as the coloured labelled cograph with $V_{N'} = V_N$, $E_{N'} = E_N$, a colour $\{ v, w \}$ for each $vw \in \perp_N$, and the label of v in N' defined as the label of $f(v)$ in $\mathcal{G}(\varphi)$, where $f(v) \in V_{\mathcal{M}(\varphi)}$ can be viewed as a vertex in $V_{\mathcal{G}(\varphi)}$ since $\mathcal{M}(\varphi) = \mathbb{M}(\mathcal{G}(\varphi))$ by Lemma 12.3 and, by definition of $\mathbb{M}$ (Def. 12.6), $V_{\mathbb{M}(G)} = V_G$ for any closed monadic fograph G.

We claim that N' is a well-defined fograph (Def. 3.2, p. 4). By Lemma 12.19, $\mathsf{typ}_{\mathbb{M}(\varphi)} = \mathsf{typ}_{\mathcal{G}(\varphi)}$ so by Lemma 12.20, $\mathsf{typ}_{N'} = \mathsf{typ}_N$ (since the label of v in N' is that of $f(v)$ in $\mathcal{G}(\varphi)$). Thus N' has a literal since N has one (because it is a mograph), so N' is a logical cograph. We must show, for all variables x, that every x-binder b is legal, *i.e.*, the scope of b contains (a) at least one literal and (b) no other x-binder. For (a), the scope $S_G(b)$ of b in a fograph or mograph G depends only on the underlying cograph (V_G, E_G), thus $S_{N'}(b) = S_N(b)$. Thus $S_{N'}(b)$ has a literal because $S_N(b)$ does (since N is a mograph). For a contradiction to (b), Suppose $c \neq b$ were an x-binder in $S_{N'}(b)$. Let T' be the cotree $T(N')$ and let $\widehat{b}$ be the parent of b in T'. If b is existential, then $bc \in E_{N'}$ (since, by Lemma 10.60, all distinct vertices in the scope of an existential binder are in an edge, since $\widehat{b}$ is a $\times$ node in T'), contradicting $f(b) = f(c)$ (which holds because, without loss of generality, φ is rectified, so there is a unique x-binder in $\mathcal{G}(\varphi)$). Otherwise b is universal. Since $c \in S_{N'}(b)$, by Lemma 10.60 we have $b \rhd_{T'} c$, *i.e.*, $b \prec_{T'} \widehat{b} >_{T'} c$, with $\widehat{b}$ a $+$-node, since b is universal. Because f is collapsed and $f(b) = f(c)$, we cannot have $\widehat{b} \succ_{T'} c$ (otherwise b and c would be indistinguishable, contradicting f being collapsed), thus $b \prec_{T'} \widehat{b} \succ_{T'} c \otimes v$ for some vertex v. Since $cv \in E_{N'}$ and f is a graph homomorphism we have $f(c)f(v) \in E_{\mathcal{G}(\varphi)}$, so $f(b)f(v) \in E_{\mathcal{G}(\varphi)}$ (because $f(b) = f(c)$). Since f is a skew fibration, there exists $w \in V_{N'}$ such that $wb \in E_{N'}$ and $f(w)f(v) \notin E_{\mathcal{G}(\varphi)}$. Because $wb \in E_{N'}$, the meet $b \odot w$ is a $\times$-node, *i.e.*, $b \odot w = b \otimes w$, and since the parent $\widehat{b}$ of b is a $+$-node, we must have $b \otimes w >_{T'} \widehat{b}$, hence $w >_{T'} b \otimes w >_{T'} v$. Therefore $wv \in E_{N'}$, so $f(w)f(v) \in E_{\mathcal{G}(\varphi)}$ a contradiction. Thus we have proved that N' is a well-defined fograph. Since every literal label in N' comes from $\mathcal{G}(\varphi)$, N' is monadic, and since f is a directed graph fibration $(V_N, B_N) \to (V_M, B_M)$, N' is closed.

By construction, $N' = \Lambda(N)$ (Def. 12.14), so by Lemma 12.16, N' is a fonet. Since $f : N \to M$ is a skew bifibration of mographs, $f : N' \to \mathcal{G}(\varphi)$ is a skew bifibration of fographs, hence a (standard) combinatorial proof, so φ is valid by Theorem 6.2. □

The crux of the soundness proof above is to transform a collapsed monadic homogeneous combinatorial proof into a standard combinatorial proof. The following example shows why collapse occurs before this transformation. A monadic homogeneous combinatorial proof of the closed monadic formula $\forall x \exists y (py \vee \overline{p}y)$ is shown below-left.

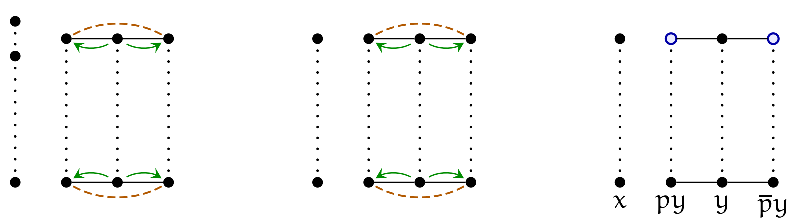

Its collapse, also a monadic homogeneous combinatorial proof (by Lemma 12.9), is shown above-centre. Above-right is the standard combinatorial proof constructed from the collapse in the soundness proof above. Observe that, were we to attempt to construct a standard combinatorial proof from the uncollapsed form, it would have two source vertices above $\bullet x$ in the target, each implicitly labelled x (implicit since we are drawing the skeleton), so the source would have a (universal) x-binder in the scope of another x-binder and therefore fail to be a well-defined fograph.

LEMMA 12.22 (Monadic homogeneous completeness). Every valid closed monadic formula has a homogeneous combinatorial proof.

Proof. Let φ be a valid closed monadic formula. By Theorem 6.3 there exists a (standard) combinatorial proof $f : N \to \mathcal{G}(\varphi)$. Let N' be the linked mograph obtained from N with $V_{N'} = V_N$, $E_{N'} = E_N$, $vw \in \perp_{N'}$ if and only if and $\{v, w\}$ is a link (colour) in N, and $B_{N'} = E_{\overrightarrow{N}}$ (*i.e.*, $vw \in B_{N'}$ if and only if v binds w in N). Since, by definition of a linked fograph, every literal is in exactly one link, every literal of N' is in a unique duality, no binder of N' is in a duality, and N' has no induced bimatching because N is a fonet; thus N' is a monet. Let $M = \mathbb{M}(\mathcal{G}(\varphi)) = \mathcal{M}(\varphi)$.

We claim that $f : N' \to M$ is a homogeneous combinatorial proof, *i.e.*, (1) f preserves existential binders, (2) $f : N' \to M$ is a skew fibration, (3) $f : (V_{N'}, \perp_{N'}) \to (V_M, \perp_M)$ is an undirected graph homomorphism, and (4) $f : (V_{N'}, B_{N'}) \to (V_M, B_M)$ is a directed graph fibration. (1) holds because $f : N \to \mathcal{G}(\varphi)$ preserves existential binders, and by construction the existential binders of N' and

N coincide, as do those of $\mathcal{G}(\varphi)$ and M. By definition (2) holds if $f : (V_{N'}, E_{N'}) \to (V_M, E_M)$ is a skew fibration, which is true because f is a skew bifibration, $(V_{N'}, E_{N'}) = (V_N, E_N)$, and $(V_M, E_M) = (V_{\mathcal{G}(\varphi)}, E_{\mathcal{G}(\varphi)})$. For (3), suppose $vw \in \perp_{N'}$. Since N is a fonet, it has a dualizer, so the labels of v and w are dual, say, p and $\overline{p}$, respectively. Because f preserves labels, $f(v)$ and $f(w)$ are labelled p and $\overline{p}$, thus $f(v)\,f(w) \in \perp_M$, and (3) holds. (4) holds because $f : N \to \mathcal{G}(\varphi)$ is a skew bifibration, thus $f : \vec{N} \to \vec{\mathcal{G}}(\varphi)$ is a directed graph fibration, and by construction $(V_{N'}, B_{N'}) = \vec{N}$ and $(V_M, B_{\mathcal{G}(\varphi)}) = \vec{\mathcal{G}}(\varphi)$. □

Proof of Theorem 8.11 (Monadic homogeneous soundness and completeness).
Lemmas 12.21 and 12.22. □

13 Homogeneous surjections

We observed in §10 that $\mathcal{G}$ is a surjection from rectified formulas to rectified fographs (Lemma 10.1), and that $\mathbb{G}$ is a surjection from clear formulas onto fographs (Lemma 10.25). In this section we exhibit similar surjections onto duality graphs and mographs.

LEMMA 13.1. $\mathbb{D}$ is a surjection from propositional fographs onto dualizing graphs.

Proof. Let D be a dualizing graph. We construct a fograph G such that $\mathbb{D}(G) = D$. Define $V_G = V_D$ and $E_G = E_D$, with a nullary predicate symbol label on each vertex defined as follows. Since $(V_D, \perp_D)$ is P_4-free and C_3-free, it is a disjoint union of complete bipartite graphs[16] $K_1, \ldots, K_n$. Choose distinct nullary predicate symbols $p_1, \ldots, p_n$ such that $\overline{p}_i \neq p_j$ $(1 \leqslant i, j \leqslant n)$. If K_i has no edges, it has a single vertex v_i; assign p_i as the label of v_i. Otherwise, $K_i = K_i' \times K_i''$ for K_i without edges. Assign the label p_i to every vertex in K_i' and the label $\overline{p}_i$ to every vertex in K_i''. The graph G is a non-empty cograph with vertices labelled by nullary predicate symbols, hence G is a propositional fograph. By construction, $\mathbb{D}(G) = D$. □

LEMMA 13.2. The function $\mathcal{D}$ from propositions to dualizing graphs (Def. 7.2) is a surjection.

Proof. By Lemma 10.1 $\mathcal{G}$ is a surjection from (rectified) formulas onto fographs. The restriction of $\mathcal{G}$ to propositions is a surjection onto propositional fographs. Since $\mathcal{D} = \mathbb{D} \circ \mathcal{G}$ by Lemma 12.3, and $\mathbb{D}$ Lemma 13.1, $\mathcal{D}$ is a surjection. □

LEMMA 13.3. $\mathbb{M}$ is a surjection from closed monadic fographs onto mographs.

Proof. Let M be a mograph. We will construct a closed monadic fograph G with $\mathbb{M}(G) = M$. Define $V_G = V_M$ and $E_G = E_M$, and define the predicate symbol in the label of each vertex of V_G exactly as in the proof of Lemma 13.1, only this time we shall make each such predicate symbol p unary rather than nullary by adding a variable after p. For each binder b in M, choose a distinct variable x_b, set the label of b to x_b, and for every literal l with $\langle b, l \rangle \in B_M$, add the variable x_b to the label of l as the argument of the predicate symbol already assigned to l. Since every binder in M has non-empty scope, every binder in G has non-empty scope. By construction every literal label is a unary predicate symbol followed by a variable, so G is monadic. Because every variable x_b is distinct for each binder b, no literal in G can be bound by two binders in G. Thus G is a rectified monadic fograph. Since, by definition of a literal in a mograph, every literal in M is the target of a binding in B_M, every literal in G is bound, so G is closed. By construction, $\mathbb{M}(G) = M$. □

LEMMA 13.4. The function $\mathcal{M}$ from closed monadic formulas to mographs graphs (Def. 8.3) is a surjection.

Proof. By Lemma 10.1 $\mathcal{G}$ is a surjection from (rectified) formulas onto fographs. The restriction of $\mathcal{G}$ to closed monadic formulas is a surjection onto closed monadic fographs. Since $\mathcal{M} = \mathbb{M} \circ \mathcal{G}$ by Lemma 12.8, and $\mathbb{M}$ is a surjection by Lemma 13.3, $\mathcal{M}$ is a surjection. □

[16] Recall that a complete bipartite graph is one of the form $G \times H$ for edgeless graphs G and H.

14 Polynomial-time verification

In this section we show that a combinatorial proof can be verified in polynomial time. Thus combinatorial proofs constitute a formal *proof system* [CR79].

The ***size*** of a graph G is the sum of the number of vertices in G and the number of edges in G.

LEMMA 14.1. The dependencies of a linked rectified fograph K can be constructed in time polynomial in the size of K.

Proof. Let $x_1, \ldots, x_n$ be the existential variables in K. The primary unification algorithm of [MM76] provides in linear time an assignment $\{x_1 \mapsto u_1, \ldots, x_n \mapsto u_n\}$ with x_i not in u_j for $i \leqslant j$, such that the most general unifier σ is $\{x_1 \mapsto t_1, \ldots, x_n \mapsto t_n\}$ for $t_i = u_i\{x_{i+1} \mapsto u_{i+1}\} \ldots \{x_n \mapsto u_n\}$ (the sequential composition of $n-i$ one-variable substitutions applied to u_i). Let $\{y_{i1}, \ldots, y_{im_i}\}$ be the set of variables occurring in u_i, and define u'_i as $f_i y_{i1} \ldots y_{im_i}$ for a fresh m_i-ary function symbol f_i. The assignment $\sigma' = \{x_1 \mapsto t'_1, \ldots, x_n \mapsto t'_n\}$ for $t'_i = u'_i\{x_{i+1} \mapsto u'_{i+1}\} \ldots \{x_n \mapsto u'_n\}$ has the same dependencies as σ but can be constructed in polynomial time since each x_j appears at most once in each u'_i. □

The above proof is essentially the first part of the proof of Theorem 3 in [Hug18].

LEMMA 14.2. The correctness of fonet can be verified in time polynomial in its size.

Proof. Let N be a fonet of size n. By Lemma 14.1 we can construct all dependencies of N in polynomial time, hence the leap graph $\mathcal{L}_N$ in polynomial time. By Lemma 10.22 every fonet is constructible from axioms by fusion and quantification. Since there can be at most n fusions and/or quantifications, it suffices to show that each step in the inductive decomposition of a fonet in the proof of Lemma 10.22 can be performed in polynomial time. In the first case the proof of Lemma 10.22, N has no edges (which can be determined in polynomial time), and to confirm that N is a union of axioms takes polynomial time. In the second case, N is universal, and the universal binder can be found and deleted in polynomial time, by inspecting each vertex of N in succession.

In the final case, N is not universal and has at least one edge, and we seek to decompose N as a fusion or existential quantification via Lemma 10.21. Henceforth we follow the proof of Lemma 10.21 closely. The graph Ω in the proof of Lemma 10.21 can be constructed in polynomial time from the cotree, which can built in polynomial time [CLS81]. The bridge $G_m G_{m+1}$ can be located in polynomial time (by iterating through the edges of Ω), and K_1 and K_2 can be determined in polynomial time by traversing edges. The underlying fograph G of N is $K_1 + (G_m \times G_{m+1}) + K_2$. Depending on whether both G_m and G_{m+1} both contain literals, the proof of Lemma 10.21 now provides either N as a fusion of $K_1 + G_m$ and $G_{m+1} + K_2$, and we recurse with each half of the fusion, or $N = \bullet x + N'$, and we delete the existential binder $\bullet x$ and recurse with N'. □

Define the ***size*** of a combinatorial proof $f : N \to G$ as the sum of the size of N and the size of G.

THEOREM 14.3. *The correctness of a combinatorial proof can be verified in time polynomial in its size.*

Proof. Let $f : N \to G$ be a combinatorial proof. By Lemma 14.2 the fonet N can be verified in polynomial time. Verifying that f is a skew bifibration is polynomial time because the skew fibration and directed graph fibration conditions apply to pairs of vertices, one in N and one in G, seeking the existence of a vertex in N, which can be found be iterating through each vertex of N in turn. □

15 Cut combinatorial proofs

Just as sequent calculus proofs may include cuts [Gen35], combinatorial proofs can be extended with cuts. An ***n-cut combinatorial proof*** of a formula φ as a combinatorial proof of a formula $\varphi \vee (\theta_1 \wedge \neg\theta_1) \vee \ldots \vee (\theta_n \wedge \neg\theta_n)$ for arbitrary formulas $\theta_1, \ldots, \theta_n$. Each formula $\theta_i \wedge \neg\theta_i$ is a ***cut***. A ***cut combinatorial proof*** is an n-cut combinatorial proof for some $n \geqslant 0$; if $n = 0$ the combinatorial proof is ***cut-free***.

THEOREM 15.1. *A formula is valid if and only if it has a cut combinatorial proof.*

Proof. Since $\varphi \vee (\theta_1 \wedge \neg\theta_1) \vee \ldots \vee (\theta_n \wedge \neg\theta_n)$ is valid if and only if φ is valid, the result follows from Theorem 6.4. □

16 Conclusion and related work

This paper reformulated classical first-order logic with combinatorial rather than syntactic proofs (§3–§6), extending the propositional case of [Hug06a] to quantifiers. The proofs of soundness (§10) and completeness (§11) were significantly more intricate than those of the propositional case [Hug06a, §5].[17] In the propositional, monadic and S5-modal special cases, labels can be removed from a combinatorial proof, and colouring from the source, for a homogeneous form (§7–§9).

Propositional combinatorial proofs are related to sequent calculus [Gen35] in [Hug06b] and [Car10], and to other syntactic systems (including resolution and analytic tableaux) in [Str17] and [AS18]. Skew fibrations are decomposed as propositional structural maps (composites of contraction and weakening maps) in [Hug06b] and [Str07]. Combinatorial proofs may provide an avenue to tackle Hilbert's 24th problem [TW02, Thi03, Hug06b, Str19].

Combinatorial proofs for non-classical logics are being actively pursued. For example, combinatorial proofs for propositional intuitionistic logic are presented in [HHS19a]. A potential topic of future research would be first-order intuitionistic combinatorial proofs. Cut elimination procedures for propositional cut combinatorial proofs are presented in [Hug06b] and [Str17]. Natural open questions include the extension of propositional intuitionistic combinatorial proofs to first-order, and cut elimination procedures for first-order combinatorial proofs (classical and intuitionistic).

Links between literals in fonets, which become dual only after applying a dualizer/unifier, are akin to the first-order *connections* or *matings* employed in automated theorem proving [Bib81, And81]. Bibel in [Bib81, p. 4] proposed *link* as an alternative name for a connection, and we have adapted that terminology in the present paper. Since a combinatorial proof can be verified in polynomial time (§14), combinatorial proofs constitute a formal *proof system* [CR79], in contrast to the connection and mating methods. The roots of first-order connections/matings with unification can be found in Prawitz [Pra70] and [Qui55]. Unification in the context of first-order logic can be traced back even further, to Robinson's resolution [Rob65] and Herbrand's theorem [Her30]. Propositional links between dual literals can be found in predecessors to the first-order connections/matrix method [Dav71, Bib74, And76], and sets of such propositional links form a category [LS05]. The pairing of propositional dual occurrences can be found in the study of other forms of syntax, such as closed categories [KM71] (see also [EK66]), contraction-free predicate calculus [KW84] and linear logic [Gir87].

A precursor to a fonet, called a *unification net*, was presented in [Hug18], building on proof nets for first-order multiplicative linear logic [Gir87, Gir91]. Unification nets are also available for first-order additive linear logic [HHS19b]. Propositional fonets correspond to the *nicely coloured cographs* of [Hug06a], which in turn correspond to the *alternate elementary acyclic R&B cographs* of [Ret03]. For background on cographs (complement-reducible graphs) see [Ler81, Sum73, CLS81]. That cographs are exactly the P_4-free graphs is proved in [Sum73].

Abstract representations of first-order quantifiers with explicit witnesses are in [Hei10] (extending expansion trees [Mil84]) and [McK10a] (for classical logic) and [HHS19b] (for additive linear logic). Composition of witnesses is analysed in [Mim11] and [ACHW18].

Proof nets [Gir87] were extended to propositional classical logic in [Gir91] (developed in detail in [Rob03]). The paper [McK13] fixes issues of redundancy due to contraction and weakening nodes and relates classical propositional proof nets to propositional combinatorial proofs [Hug06a, Hug06b].

Peirce [Pei33, vol. 4:2] provides an early graphical representation of propositional formulas.

A Fograph examples

Figs. 11–12 show a progression of instructive examples of formula graphs.

[17] This may relate to the fact that first-order logic is undecidable, while propositional logic is decidable.

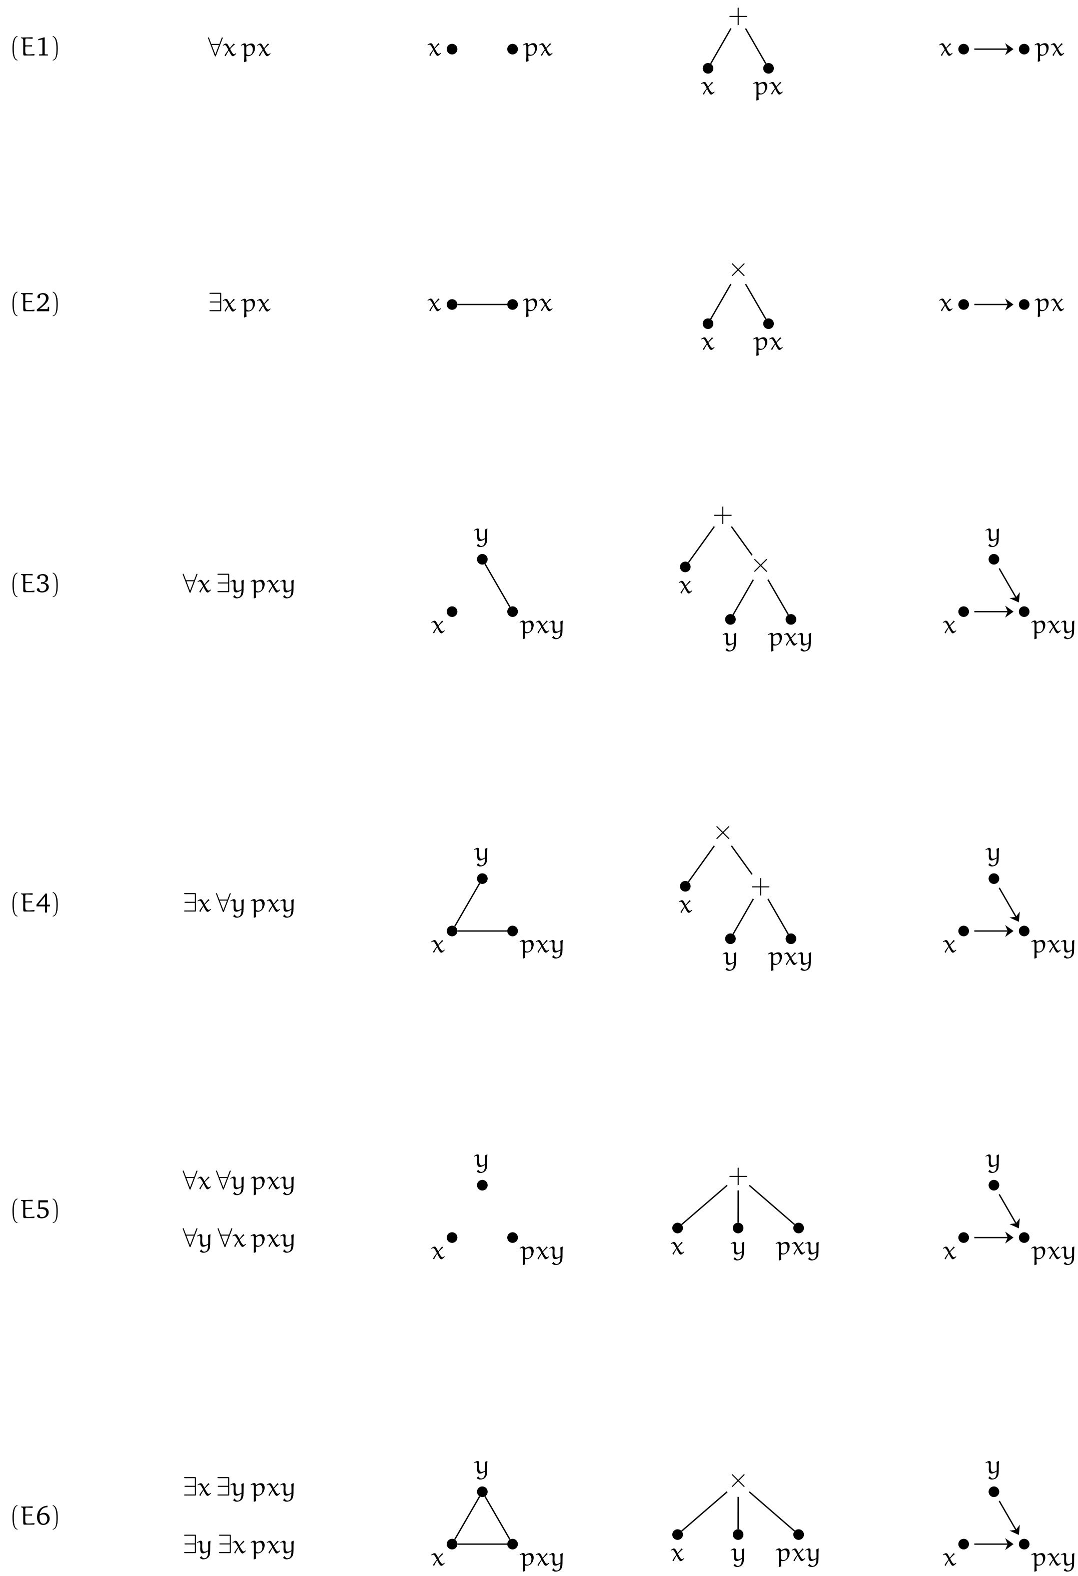

Figure 11. Examples E1–6. Each syntactic formula φ is followed by its graph $\mathsf{G} = \mathcal{G}(\varphi)$, cotree $\mathsf{T}(\mathsf{G})$, and binding graph $\vec{\mathsf{G}}$. Examples E5 and E6 show two syntactic formulas with the same combinatorial formulas.

(E7)

$\forall x\,(px \wedge qx)$

$\forall x\,(qx \wedge px)$

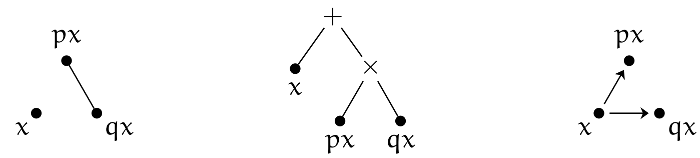

(E8)

$\exists x\,(px \vee qx)$

$\exists x\,(qx \vee px)$

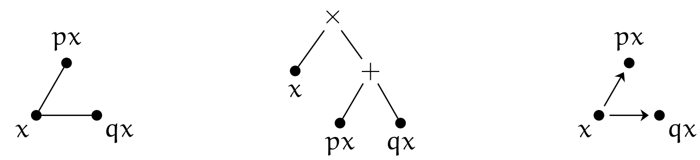

(E9)

$\exists x\,(\,px \wedge q\,)$

$\exists x\,(\,q \wedge px\,)$

$(\,\exists x\, px\,) \wedge q$

$q \wedge (\,\exists x\, px\,)$

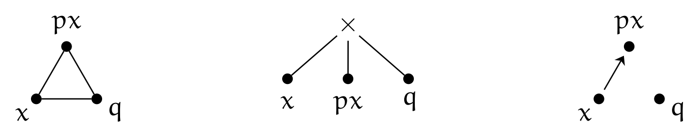

(E10)

$(\exists x\, px) \vee (\exists y\, qy)$

$(\exists y\, qy) \vee (\exists x\, px)$

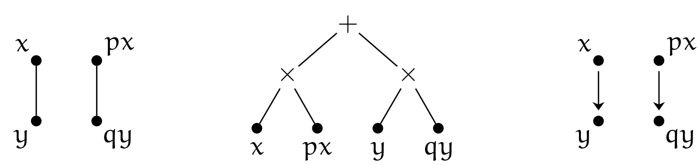

(E11)

$(\forall x\, px) \wedge (\forall y\, qy)$

$(\forall y\, qy) \wedge (\forall x\, px)$

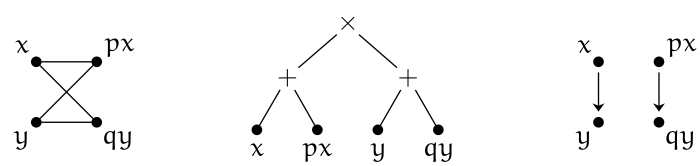

(E12)

$\exists x\,\big(\,px \vee (\forall y\, qy)\,\big)$

$\exists x\,\big(\,(\forall y\, qy) \vee px\,\big)$

$\exists x\,\big(\,\forall y\,(px \vee qy)\,\big)$

$\exists x\,\big(\,\forall y\,(qy \vee px)\,\big)$

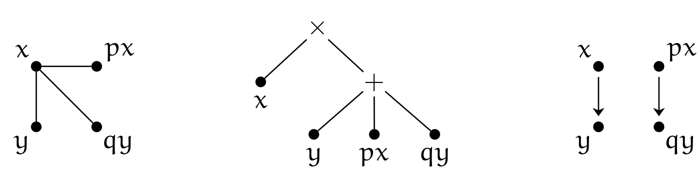

Figure 12. Examples E7–12. Each syntactic formula φ is followed by its graph $G = \mathcal{G}(\varphi)$, cotree $T(G)$, and binding graph $\vec{G}$. Each example shows multiple syntactic formulas with the same combinatorial formula.

References

[ACHW18] A. Alcolei, P. Clairambault, J. M. E. Hyland, and G. Winskel. The true concurrency of herbrand's theorem. In *CSL'18)*, 2018.

[And76] P. B. Andrews. Refutations by matings. *IEEE Transactions on Computers*, 25(8), 1976.

[And81] P. B. Andrews. Theorem proving via general matings. *J. ACM*, 28, 1981.

[AS18] M. Acclavio and L. Straßburger. From syntactic proofs to combinatorial proofs. In *Int. Joint Conf. on Automated Reasoning*. Springer, 2018.

[Bib74] W. Bibel. An approach to a systematic theorem proving procedure in first-order logic. *Computing*, 12(1), 1974.

[Bib81] W. Bibel. On matrices with connections. *J. ACM*, 28(4), 1981.

[BK92] G. Bellin and J. Ketonen. A decision procedure revisited: Notes on direct logic, linear logic and its implementation. *Theor. Comp. Sci.*, 95:115–142, 1992.

[Car10] A. Carbone. A new mapping between combinatorial proofs and sequent calculus proofs read out from logical flow graphs. *Inf. Comput.*, 208(5), 2010.

[CLS81] D. G. Corneil, H. Lerchs, and L. K. Stewart-Burlingham. Complement reducible graphs. *Disc. Appl. Math.*, 1981.

[CR79] S. A. Cook and R. A. Reckhow. The relative efficiency of propositional proof systems. *J. Symb. Logic*, 44:36–50, 1979.

[Dav71] G. V. Davydov. The synthesis of the resolution method with the inverse method. *Zapiski Nauchnykh Seminarov Lomi 20 1971 24–35.*, 20:24–35, 1971. Translation in J. Sov. Math. 1(1) 1973 12-–18.

[EK66] S. Eilenberg and G. M. Kelly. A generalization of the functorial calculus. *J. Algebra*, 3:366–375, 1966.

[Fre79] G. Frege. *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Halle, 1879.

[Gal67] T. Gallai. Transitiv orientierbare graphen. *Acta Math. Acad. Sci. Hungar*, 18:25–66, 1967.

[Gen35] G. Gentzen. Untersuchungen über das logische Schließen. II. *Mathematische Zeitschrift*, 39:405–431, 1935.

[Gir87] J.-Y. Girard. Linear logic. *Theor. Comp. Sci.*, 50:1–102, 1987.

[Gir89] J.-Y. Girard. Towards a geometry of interaction. In *Categories in Computer Science and Logic*, volume 92 of *Contemporary Mathematics*, pages 69–108, 1989. Proc. of June '87 meeting in Boulder, Colorado.

[Gir91] J.-Y. Girard. A new constructive logic: Classical logic. *Mathematical Structures in Comp. Sci.*, 1:255–296, 1991.

[Gra66] J. W. Gray. Fibred and cofibred categories. In *Proc. Conf. on Categorical Algebra '65*, pages 21–83. Springer, 1966.

[Gro60] A. Grothendieck. Technique de descente et théorèmes d'existence en géométrie algébrique. I. Généralités. Descente par morphismes fidèlement plats. In *Séminaire Bourbaki: années 1958/59–1959/60, exposés 169-204*. Société mathématique de France, 1960.

[Hei10] W. Heijltjes. Classical proof forestry. *Annals of Pure and Applied Logic*, 161(11), 2010.

[Her30] J. Herbrand. *Recherches sur la théorie de la démonstration*. PhD thesis, Sorbonne, Paris, 1930.

[HHS19a] W. Heijltjes, D. J. D. Hughes, and L. Straßburger. Intuitionistic proofs without syntax. In *Proc. LICS'19*, 2019.

[HHS19b] W. Heijltjes, D. J. D. Hughes, and L. Straßburger. Proof nets for first-order additive linear logic. In *Proc. FSCD'19*, 2019.

[Hug06a] D. J. D. Hughes. Proofs without syntax. *Annals of Mathematics*, 143:1065–1076, 2006.

[Hug06b] D. J. D. Hughes. Towards Hilbert's 24^{th} Problem: Combinatorial Proof Invariants. In *Proc. WOLLiC'06*, volume 165 of *Lec. Notes in Comp. Sci.*, 2006.

[Hug10] D. J. D. Hughes. A minimal classical sequent calculus free of structural rules. *Annals of Pure and Applied Logic*, 161(10), 2010.

[Hug18] D. J. D. Hughes. Unification nets: canonical proof net quantifiers. In *Proc. LICS '18*, 2018.

[Joh87] P. T. Johnstone. *Note on Logic and Set Theory*. Cambridge University Press, 1987.

[KM71] G. M. Kelly and S. Mac Lane. Coherence in closed categories. *J. Pure Appl. Algebra*, 1:97–140, 1971.

[Kot59] A. Kotzig. On the theory of finite graphs with a linear factor ii. *Mat.-Fyz. Casopis. Slovensk. Akad. Vied*, 9(3), 1959.

[KW84] J. Ketonen and R. Wehyrauch. A decidable fragment of predicate calculus. *Theor. Comp. Sci.*, 32:297–307, 1984. Error in proof of main theorem corrected in [BK92].

[Ler81] H. Lerchs. On cliques and kernels. Tech. report, U. Toronto, 1981.

[LP86] L. Lovász and M. D. Plummer. *Matching Theory*. North-Holland, 1986.

[LS05] F. Lamarche and L. Straßburger. Naming proofs in classical propositional logic. In *Proc. Typed Lambda Calculus '05*, Lec. Notes in Comp. Sci., 2005.

[McK10a] R. McKinley. Expansion nets: Proof nets for for propositional classical logic. In *LPAR 17*, volume 6397 of *LNCS*, 2010.

[McK10b] R. McKinley. Proof nets for herbrand's theorem. *ACM Trans. Comp. Logic*, 14, 2010.

[McK13] R. McKinley. Canonical proof nets for classical logic. *Annals of Pure and Applied Logic*, 164(6):702–732, 2013.

[Mil84] D. A. Miller. Expansion tree proofs and their conversion to natural deduction proofs. *Lec. Notes in Comp. Sci.*, 170, 1984.

[Mim11] S. Mimram. The structure of first-order causality. *Mathematical Structures in Comp. Sci.*, 21, 2011.

[Min92] G. Mints. *A short introduction to modal logic*. CSLI, 1992.

[MM76] A. Martelli and U. Montanari. Unification in linear time and space: a structured presentation. Technical report, U. Pisa, 1976.

[Pei33] C. S. Peirce. Description of a notation for the logic of relatives, resulting from an amplification of the conceptions of Boole's calculus of logic. In *Collected Papers of Charles Sanders Peirce. III. Exact Logic*. Harvard University Press, 1933.

[Pra70] D. Prawitz. A proof procedure with matrix reduction. *Lecture Notes in Mathematics*, 125, 1970.

[Qui55] W. V. Quine. A proof procedure for quantification theory. *J. Symbolic Logic*, 20(2), 1955.

[Ret03] C. Retoré. Handsome proof-nets: perfect matchings and cographs. *Theor. Comp. Sci.*, 2003.

[Rob65] J. A. Robinson. A machine-oriented logic based on the resolution principle. *J. ACM*, 12(1), 1965.

[Rob03] E. Robinson. Proof nets for classical logic. *J. Logic and Computation*, 13(5):777–797, 2003.

[Str07] L. Straßburger. A characterisation of medial as rewriting rule. In *Term Rewriting and Applications, RTA07*, volume 4533. Lec. Notes in Comp. Sci., 2007.

[Str17] L. Straßburger. Combinatorial Flows and Their Normalisation. In *Proc. FSCD 2017*, volume 84 of *Leibniz Int. Proc. Informat.*, 2017.

[Str19] L. Straßburger. The problem of proof identity, and why computer scientists should care about hilbert's 24th problem. *Phil. Trans. Royal Soc. A*, 377(2140), 2019.

[Sum73] D. P. Sumner. Graphs indecomposable with respect to the x-join. *Discr. Math.*, 6, 1973.

[Thi03] R. Thiele. Hilbert's twenty-fourth problem. *The American Mathematical Monthly*, 110(1), 2003.

[TS96] A. S. Troelstra and H. Schwichtenberg. *Basic Proof Theory*. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, Cambrige, U.K., 1996.

[TW02] R. Thiele and L. Wos. Hilbert's twenty-fourth problem. *Journal of Automated Reasoning*, 29(1):67–89, 2002.

[Whi78] G.W Whitehead. *Elements of Homotopy Theory*. Springer-Verlag, 1978.