Mathematical Structures in Computer Science (2026), **36**, e33, pp. 1–42
doi:10.1017/S0960129525000143

CAMBRIDGE UNIVERSITY PRESS

PAPER

Elimination and cut-elimination in multiplicative linear logic

Daniel Murfet and William Troiani

Mathematics and Statistics, University of Melbourne Institute, Melbourne, Australia
Corresponding author: Daniel Murfet; Email: d.murfet@unimelb.edu.au

(Received 14 June 2023; revised 25 October 2024; accepted 7 April 2025)

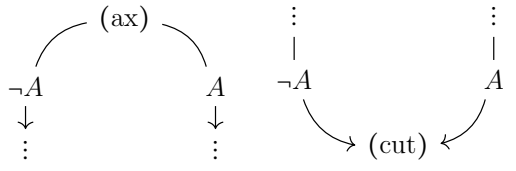
Abstract
We associate with every proof structure in multiplicative linear logic an ideal that represents the logical content of the proof as polynomial equations. We show how cut-elimination in multiplicative proof nets corresponds to instances of the Buchberger algorithm for computing Gröbner bases in elimination theory.

Keywords: Multiplicative linear logic; proof nets; cut-elimination; Gröbner bases; elimination theory

1. Introduction

The correspondences between logic and computation that can be found in the work of Curry-Howard Howard (1980) and Gentzen-Mints-Zucker Murfet and Troiani (2020) have as their deep content the observation that intuitionistic proofs can be understood as *patterns of equality* that relate differently named occurrences of the 'same' variable. Such patterns can be enacted on other patterns in a dynamical process that we call computation. By now, we have many useful ways of thinking about these patterns as static objects, but the mathematical theory of these patterns in motion, what Girard calls the *geometry of interaction* Girard (1989), is not yet as well-developed.

This paper is the first of a series in which we bring a fresh perspective to this geometry of interaction, using algebraic geometry. Here, we only consider the simplest kind of proof, multiplicative proof nets in linear logic, and their cut-elimination process. We assign to every multiplicative proof net π (from here on out just *proof net*) a polynomial ring P_π, in which the variables are occurrences of atomic propositions in the formulas labelling edges of π, and an ideal I_π in this ring (Definition 3.21). The generating set G_π for this ideal contains, in particular, for every (ax) or (cut) link in π on an atomic proposition A



a polynomial $A' - A$ where A, A' are two distinct copies of the same atomic proposition, one for each edge incident at the link (informally, both links say $A = A$ and this equation is in G_π). For any reduction γ of π yielding another proof net π', there is a natural inclusion $P_{\pi'} \subseteq P_\pi$ where the variables in P_π not in $P_{\pi'}$ are atomic propositions on edges eliminated by γ. The relationship

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

between I_π and $I_{\pi'}$ is straightforward (Corollary 4.7)

$$I_{\pi'} = I_\pi \cap P_{\pi'} . \tag{1}$$

If we think of I_π as all consequences of the generating equations in G_π, then this says that the pattern of equalities between the occurrences of atomic propositions in π' *derived from* the equations in G_π is the same as the pattern we would write down directly from the structure of π'. There is, however, an important subtlety here: the generating set $G_{\pi'}$ for $I_{\pi'}$ *cannot* be derived from G_π by simple intersection, as

$$G_{\pi'} \neq G_\pi \cap P_{\pi'} . \tag{2}$$

The difference between G_π and I_π is analogous to the difference between a set of axioms and the set of all propositions that may be derived from those axioms. The reason for the lack of equality in (2) is easy to understand: if γ eliminates a (cut) link as above then $A' - A \in G_\pi$, and if $A_1 - A', A - A_2 \in G_\pi$ are generators associated to other links in π then

$$A_1 - A_2 = (A_1 - A') + (A' - A) + (A - A_2) \in I_\pi$$

and hence $A_1 - A_2 \in I_{\pi'} = I_\pi \cap P_\pi$. As we will see, this is one of the generators in $G_{\pi'}$ (since in the reduced proof net A_1, A_2 are incident at a single link). However, if we simply intersect G_π with P_π, then we delete all three of the equations $A_1 - A', A' - A, A - A_2$ since each contains some eliminated variable, and the connection between A_1 and A_2 is lost.

In algebraic geometry, *Elimination Theory* (Cox et al. 2015, Ch. 3) provides a general framework for eliminating variables from systems of polynomial equations, in the sense that A', A have been eliminated from $A_1 - A', A' - A, A - A_2$ to obtain $A_1 - A_2$. In our present circumstances, the theory says the following:

(1) Choose a lexicographic monomial order $<_\gamma$ for P_π in which the variables eliminated by the reduction γ (e.g. A', A) are large.
(2) Compute a Gröbner basis for I_π from G_π using this monomial order.
(3) Intersect the resulting Gröbner basis with $P_{\pi'}$.

The main theorem of Elimination Theory (recalled as Theorem 5.7 here) says that the intersection in (3) is a Gröbner basis for $I_{\pi'}$. We prove in Theorem 7.2 that, if the monomial order is chosen carefully and a particular form of the Buchberger algorithm for computing Gröbner bases is used in (2), then (3) computes $G_{\pi'}$ (which is a Gröbner basis for $I_{\pi'}$ with respect to a monomial order determined by the ordering of atomic propositions along persistent paths). In this sense, the Buchberger algorithm can be viewed as an algebraic interpretation of cut-elimination for multiplicative proof nets.

The goals of this paper are to set up a basic dictionary between

- multiplicative proof nets and ideals
- reduction sequences and monomial orders
- cut-elimination and elimination

While there is some small novelty in the form of the Buchberger algorithm that makes this connection clearest (we call it Buchberger with Early Stopping), there is nothing deep here logically, algebraically or geometrically. Indeed, the ideals I_π do not have a very interesting geometry: the associated affine variety is just an intersection of pairwise diagonals. In subsequent papers in this series, we introduce, on top of the foundations laid here, more interesting algebra and geometry (see Section 8.1 for some hints).

Previous authors have also developed Girard's Geometry of Interaction (GoI) programme; we mention a few here. In Abramsky et al. (1999), the connection between GoI and linear combinatory algebras is explored. There is also Abramsky (1994), which introduces traced monoidal

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

categories as a framework for capturing feedback loops and iterative processes, making it particularly relevant to modelling GoI and cut-elimination. This work is built upon in Haghverdi and Scott (2006), which considers the cut-elimination process within the framework of Abramsky (1994). We believe that the work considered in the present paper can be organised into the language of GoI situations and traced monoidal categories, a connection that we plan to explore in future work.

2. Motivation: The Canonical Detour

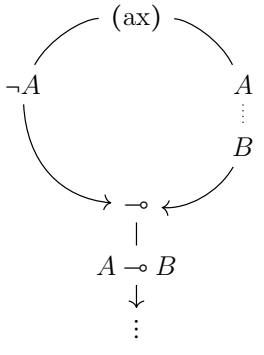
The reader wishing to dive straight into the details should skip to Section 3. In this section, we recall some of the concepts behind proof nets, following (Girard 2011, §11.2.2). The deduction rules for intuitionistic linear natural deduction include the introduction and elimination rules for linear implication:

$$\frac{\begin{array}{c}[A]^{\mathbf{1}}\\ \vdots\\ B\end{array}}{A \multimap B}\,(\multimap I)^{\mathbf{1}} \qquad\qquad \frac{\begin{array}{cc}\vdots & \vdots\\ A & A \multimap B\end{array}}{B}\,(\multimap E) \tag{3}$$

Among other criticisms of this system of natural deduction, Girard points to the need for a non-local 'gimmick' (the de Bruijn index **1**) to 'physically' link the hypothesis A to the introduction rule (Girard 2011, §11.2.1). This non-locality is resolved in proof nets, in two steps: firstly, by giving up active hypotheses and having only conclusions, and secondly by bringing these hypotheses-as-conclusions down to meet the introduction rule that employs them. The first step involves introducing the (ax) and (cut) links



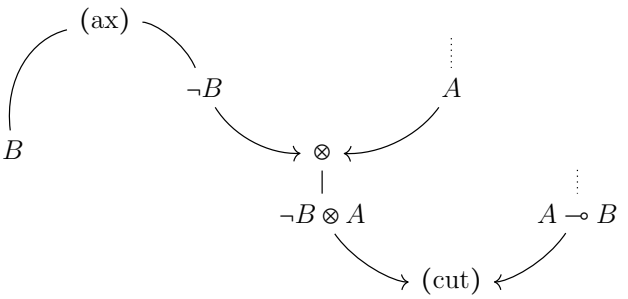
which 'replace a hypothesis by a conclusion and 'vice versa'. The second step, in which $\neg A$ is brought down to meet the introduction rule, is made clear by the proof net interpreting the introduction rule in (3)



(4)

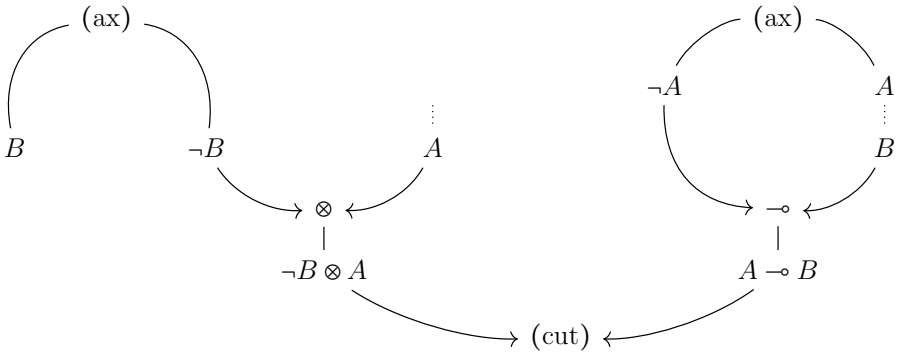
To explain Girard's 'vice versa' note that we can read the elimination rule of (3) as converting a proof of the *conclusion* $A \multimap B$ into a proof of B on *hypothesis* A, and this is achieved in proof nets via (cut), as shown in the following diagram:

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143



(5)

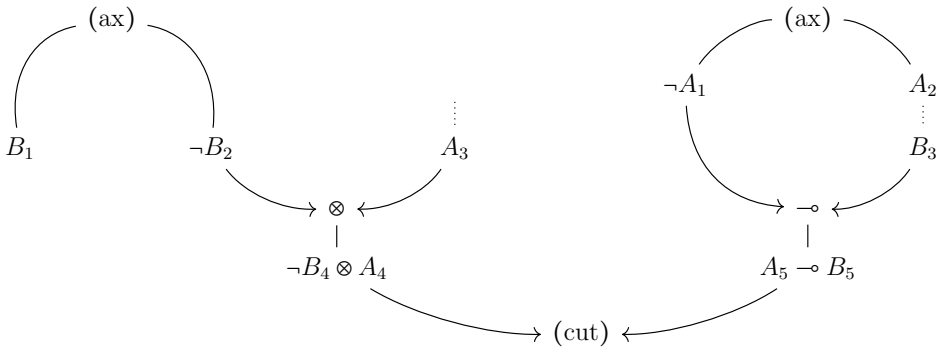
The reduction step in natural deduction is the elimination of the 'detour' created when the introduction rule in (3) appears as the proof of $A \multimap B$ in the elimination rule; this detour is eliminated to obtain a more straightforward proof of B by combining the given proof of A and the proof of B on hypothesis A. The canonical example of a detour is shown below:



(6)

In the language of (linear) lambda terms, the proof net (4) is $\lambda x\,.\,M$ for some term M determined by the proof of B from A in ellipsis. If the proof of A in ellipsis in (5) has term N, then the cut (6) corresponds to $((\lambda x\,.\,M)N)$ reducing to $M[x := N]$ with the 'output' routed through the (cut)and $\otimes$ links to the conclusion B on the left. We can imagine a path (a more refined version of which we will later call a *persistent path*) from the proof of A down through the $\otimes$ link and (cut) links, up through the $\multimap$ link and around to the left through $\neg A$ to the (ax) link, then back down through the proof of B from A (the one with term M) through the $\multimap$ and (cut) links, then up through $\otimes$ and the (ax) for B until it exits the conclusion B. Along this path, 'nothing happens' except when we pass through the proof of B from A, and we will get the same result if we 'straighten out' the rest of the proof net, leaving just this interesting part *with the proof of A feeding into it*. This is the result of cut-elimination for this proof net.

Let us sketch the algebro-geometric point of view on this situation. We first distinguish every occurrence of the formulas A, B in the proof net above by giving it an index[1]



Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Suppose the atomic propositions occurring in A are $X_1, \ldots, X_n$ and label the occurrence of X_j in A_i by $X_j^{(i)}$. Similarly let $Y_1, \ldots, Y_m$ be the atomic propositions in B and write $Y_j^{(i)}$ for the occurrence of Y_j in B_i. We let $\mathbf{X}^{(i)}$ (resp. $\mathbf{Y}^{(i)}$) stand for all the X-variables (resp. Y-variables) with superscript i. Then, associated with the path described above through the proof net are the identifications

$$\mathbf{X}^{(3)} = \mathbf{X}^{(4)} = \mathbf{X}^{(5)} = \mathbf{X}^{(1)} = \mathbf{X}^{(2)} \tag{7}$$

leading to the 'input' of the proof B of A, and the identifications

$$\mathbf{Y}^{(3)} = \mathbf{Y}^{(5)} = \mathbf{Y}^{(4)} = \mathbf{Y}^{(2)} = \mathbf{Y}^{(1)} \tag{8}$$

leading from the 'output' back through the proof net to the conclusion B. Here, $\mathbf{X}^{(i)} = \mathbf{X}^{(i')}$ stands for a sequence of equations $X_j^{(i)} = X_j^{(i')}$ for $1 \leq j \leq n$. We have grouped these equations by how they appear on the path from A to B through the proof net, but we can also group them by link

$$(\mathrm{ax})_B \quad \mathbf{Y}^{(1)} = \mathbf{Y}^{(2)} \tag{9}$$

$$(\mathrm{ax})_A \quad \mathbf{X}^{(1)} = \mathbf{X}^{(2)} \tag{10}$$

$$(\otimes) \quad \mathbf{X}^{(3)} = \mathbf{X}^{(4)}, \mathbf{Y}^{(2)} = \mathbf{Y}^{(4)} \tag{11}$$

$$(\multimap) \quad \mathbf{Y}^{(3)} = \mathbf{Y}^{(5)}, \mathbf{X}^{(1)} = \mathbf{X}^{(5)} \tag{12}$$

$$(\mathrm{cut}) \quad \mathbf{X}^{(4)} = \mathbf{X}^{(5)}, \mathbf{Y}^{(4)} = \mathbf{Y}^{(5)}\,. \tag{13}$$

For some base field k, we consider the polynomial ring P generated by all the atomic propositions

$$P = k[\mathbf{X}^{(1)}, \ldots, \mathbf{X}^{(5)}, \mathbf{Y}^{(1)}, \ldots, \mathbf{Y}^{(5)}] \tag{14}$$

and to the sets of equations (9)–(13) we associate the ideals (one for each link)

$$I_{(\mathrm{ax})_B} = \langle \mathbf{Y}^{(1)} - \mathbf{Y}^{(2)} \rangle = \left\langle Y_1^{(1)} - Y_1^{(2)}, \ldots, Y_m^{(1)} - Y_m^{(2)} \right\rangle \tag{15}$$

$$I_{(\mathrm{ax})_A} = \langle \mathbf{X}^{(1)} - \mathbf{X}^{(2)} \rangle \tag{16}$$

$$I_{\otimes} = \langle \mathbf{X}^{(3)} - \mathbf{X}^{(4)}, \mathbf{Y}^{(2)} - \mathbf{Y}^{(4)} \rangle \tag{17}$$

$$I_{\multimap} = \langle \mathbf{Y}^{(3)} - \mathbf{Y}^{(5)}, \mathbf{X}^{(1)} - \mathbf{X}^{(5)} \rangle \tag{18}$$

$$I_{(\mathrm{cut})} = \langle \mathbf{X}^{(4)} - \mathbf{X}^{(5)}, \mathbf{Y}^{(4)} - \mathbf{Y}^{(5)} \rangle\,. \tag{19}$$

Assume that somehow ideals

$$\mathfrak{n} \subseteq k[\mathbf{X}^{(3)}]\,, \qquad \mathfrak{m} \subseteq k[\mathbf{X}^{(2)}, \mathbf{Y}^{(3)}] \tag{20}$$

have been assigned to the proofs of A, and of B from A, in ellipsis. The *ideal of the proof net* (6) is (see Section 3 for the general definition)

$$I = I_{(\mathrm{ax})_B} + I_{(\mathrm{ax})_A} + I_{\otimes} + I_{\multimap} + I_{(\mathrm{cut})} + \mathfrak{m} + \mathfrak{n} \subseteq P\,. \tag{21}$$

This ideal represents the pattern of equality of the proof net in terms of the variables occurring on every edge. When it comes to the interaction of this proof net with other nets under cut-elimination, this pattern in the 'interior' only matters insofar as it implicitly determines a pattern of equality between the variables in $\mathbf{X}^{(3)}$ and $\mathbf{Y}^{(1)}$, which live on the 'boundary' of the proof net. Computation is the process of making this *implicit* pattern in the interior *explicit* on the boundary (Section 7).

The pattern of equalities on the boundary is the set of equations between the variables $\mathbf{X}^{(3)}, \mathbf{Y}^{(1)}$ which are implied by the equations between the variables in the interior. These equations form an ideal $J \subseteq k[\mathbf{X}^{(3)}, \mathbf{Y}^{(1)}]$ and to say that the equations in J follow from those in I is to say that under the inclusion

$$\iota : k[\mathbf{X}^{(3)}, \mathbf{Y}^{(1)}] \longrightarrow P$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

we have $\iota(J) \subseteq I$. To say that *every* equation between the variables $\mathbf{X}^{(3)}, \mathbf{Y}^{(1)}$ implied by the equations in I is in J is to say further that

$$J = I \cap k[\mathbf{X}^{(3)}, \mathbf{Y}^{(1)}]\,. \tag{22}$$

So, what is a generating set for this ideal? With a little computation using the above equations, it is easy to see that $J = \mathfrak{n} + \mathfrak{m}'$ is the sum of $\mathfrak{n}$ with the ideal $\mathfrak{m}'$ obtained from $\mathfrak{m}$ by replacing every generator $f(\mathbf{X}^{(2)}, \mathbf{Y}^{(3)})$ of this ideal with the polynomial $f(\mathbf{X}^{(3)}, \mathbf{Y}^{(1)})$ obtained by substituting $X_i^{(3)}$ for $X_i^{(2)}$ in f for all $1 \leq i \leq n$ and $Y_j^{(1)}$ for $Y_j^{(3)}$ for all $1 \leq j \leq m$. This parallels the substitution $M[x := N]$.

In this paper, we treat such substitution processes systematically from the point of view of Elimination Theory, as instances of the Buchberger algorithm for computing a Gröbner basis of J from I.

3. The Ideal of a Proof

To any proof structure π, we will associate a polynomial ring P_π. This is the universal algebraic object in which we can add and multiply the atomic propositions in the formulas adorning the edges of the proof structure (noting that each occurrence gets its own copy in the polynomial ring) and multiply them by scalars from some fixed base field k. The idea of making propositions into objects of arithmetic or algebra is due to Boole Boole (1854). The ring P_π only reflects the set of occurrences of atomic propositions in π not the structure of the proof; this structure is encoded in an ideal $I_\pi \subseteq P_\pi$ (Definition 3.21).

3.1 Background

We assume familiarity with linear logic, for an introduction see Girard et al. (1989); Laurent (2013); Murfet (2017); Troiani (2020).

Definition 3.1. *There is an infinite set of unoriented atoms $X, Y, Z, \ldots$ and an oriented atom (or atomic proposition) is a pair $(X, +)$ or $(X, -)$ where X is an unoriented atom. Let $\mathcal{A}$ denote the set of oriented atoms.*

For $x \in \{+, -\}$, we write $\bar{x}$ for the negation, so $\bar{+} = -$, $\bar{-} = +$.

Definition 3.2. *The set of pre-formulas is defined as follows:*

- *Any atomic proposition is a pre-formula.*
- *If A, B are pre-formulas then so are $A \otimes B$, $A \parr B$.*
- *If A is a pre-formula then so is $\neg A$.*

The set of formulas is the quotient of the set of pre-formulas by the equivalence relation $=$ generated, for arbitrary formulas A, B and unoriented atom X, by

$$\neg(A \otimes B) = \neg B \parr \neg A, \qquad \neg(A \parr B) = \neg B \otimes \neg A, \qquad \neg(X, x) = (X, \bar{x})$$

Let $\mathcal{F}$ denote the set of formulas.

The following easy lemma is left to the reader:

Lemma 3.3. *For all formulas A we have $\neg\neg A = A$.*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Definition 3.4. *Let $\mathcal{A}^* = \bigcup_{n\geq 0} \mathcal{A}^n$ be the set of sequences of oriented atoms of length ≥ 0. We define an involution r on $\mathcal{A}^*$ as follows:*

$$r : \mathcal{A}^* \longrightarrow \mathcal{A}^*$$
$$\big((X_1, x_1), \ldots, (X_n, x_n)\big) \longmapsto \big((X_n, \bar{x}_n), \ldots, (X_1, \bar{x}_1)\big)$$

For the empty string $\emptyset \in \mathcal{A}^*$, we define $r(\emptyset) = \emptyset$.

The set $\mathcal{A}^*$ is a monoid under concatenation $c : \mathcal{A}^* \times \mathcal{A}^* \longrightarrow \mathcal{A}^*$ with identity $\emptyset$.

Definition 3.5. *We denote by $\otimes : \mathcal{F} \times \mathcal{F} \longrightarrow \mathcal{F}$ the function which maps a pair of formulas (A, B) to the formula $A \otimes B$. Similarly, $\parr : \mathcal{F} \times \mathcal{F} \longrightarrow \mathcal{F}$ denotes the function such that $\parr(A, B) = A \parr B$ and $\neg : \mathcal{F} \longrightarrow \mathcal{F}$ denotes the function such that $\neg(A) = \neg A$. We denote by $\mathrm{inc} : \mathcal{A} \longrightarrow \mathcal{F}$ the map which sends an oriented atom (X, x) to itself (X, x), and lastly we denote by $\iota : \mathcal{A} \longrightarrow \mathcal{A}^*$ the function which maps an oriented atom (X, x) to the sequence consisting only of (X, x).*

Lemma 3.6. *There is a unique map $a : \mathcal{F} \longrightarrow \mathcal{A}^*$ making the following diagrams commute*

$$\begin{array}{ccc} \mathcal{F}\times\mathcal{F} & \xrightarrow{a\times a} & \mathcal{A}^*\times\mathcal{A}^* \\ {\scriptstyle\otimes}\downarrow & & \downarrow{\scriptstyle c} \\ \mathcal{F} & \xrightarrow{a} & \mathcal{A}^* \end{array} \qquad \begin{array}{ccc} \mathcal{F}\times\mathcal{F} & \xrightarrow{a\times a} & \mathcal{A}^*\times\mathcal{A}^* \\ {\scriptstyle\parr}\downarrow & & \downarrow{\scriptstyle c} \\ \mathcal{F} & \xrightarrow{a} & \mathcal{A}^* \end{array} \tag{23}$$

$$\begin{array}{ccc} \mathcal{F} & \xrightarrow{a} & \mathcal{A}^* \\ {\scriptstyle\neg}\downarrow & & \downarrow{\scriptstyle r} \\ \mathcal{F} & \xrightarrow{a} & \mathcal{A}^* \end{array} \qquad \begin{array}{ccc} \mathcal{A} & \xrightarrow{\mathrm{inc}} & \mathcal{F} \\ & {\scriptstyle\iota}\searrow & \downarrow{\scriptstyle a} \\ & & \mathcal{A}^* \end{array} \tag{24}$$

Proof. Left to the reader. □

Definition 3.7. *Let A be a formula. The sequence of oriented atoms of A is*

$$a(A) = (X_1, x_1), \ldots, (X_n, x_n)$$

as defined by the previous lemma. The sequence of unoriented atoms of A is $X_1, \ldots, X_n$ and the set of unoriented atoms of A is the disjoint union $U_A = \{X_1\} \coprod \ldots \coprod \{X_n\}$.

Our definition of *multiplicative proof structures* follows that of Laurent (2013) but with edges labelled by formulas in the sense of Definition 3.2. We remind the reader of the notion of a multigraph (a graph with multiple edges between vertices).

Definition 3.8. *A directed multigraph is a triple (V, E, φ) where:*

- *V is a set of vertices or nodes.*
- *E is a set of edges.*
- *$\Gamma : E \longrightarrow \{(x, y) \mid x, y \in V\}$ is a function from the set of edges to the set of ordered pairs of vertices.*

For $e \in E$, the first element of $\Gamma(e)$ is the source and the second element is the target.

Definition 3.9. *A proof structure is a directed multigraph with edges labelled by formulas and with vertices labelled by* (ax), (cut), $\otimes$, $\parr$ *or* c. *The incoming edges of a vertex are called its premises, the outgoing edges are its conclusions. Proof structures are required to adhere to the following conditions:*

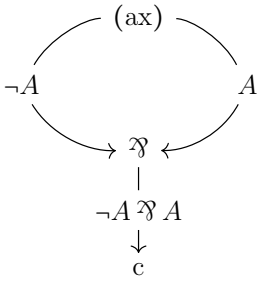
Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

- *Each vertex labelled* (ax) *has exactly two conclusions and no premise, the conclusions are labelled A and* $\neg A$ *for some A. We call this an axiom link.*
- *Each vertex labelled* (cut) *has exactly two premises and no conclusion, where the premises are labelled A and* $\neg A$ *for some A. We call this a cut link.*
- *Each vertex labelled* $\otimes$ *has exactly two premises and one conclusion. The premises are ordered, the smallest one is called the left premise of the vertex, the biggest one is called the right premise. The left premise is labelled A, the right premise is labelled B and the conclusion is labelled* $A \otimes B$*, for some A, B. We call this a tensor link.*
- *Each vertex labelled* $\parr$ *has exactly two ordered premises and one conclusion. The left premise is labelled A, the right premise is labelled B and the conclusion is labelled* $A \parr B$*, for some A, B. We call this a par link.*
- *Each vertex labelled* c *has exactly one premise and no conclusion. Such a premise of a vertex labelled* c *is called a conclusion of the proof structure.*

A link is any of an axiom link, tensor link, par link, or cut link. Conclusions are not links.

Definition 3.10. *A proof net is a proof structure which is the image under the translation map of a sequent calculus proof in multiplicative linear logic, see (Troiani 2020, Definition 1.0.12).*

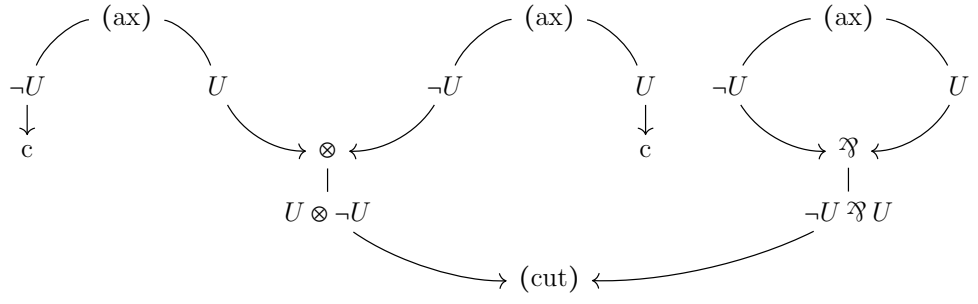
Example 3.11. *The smallest proof net with a single conclusion is*



This is the translation of the following sequent calculus proof.

$$\dfrac{\dfrac{}{\vdash \neg A, A}\,(\mathrm{ax})}{\vdash \neg A \parr A}\,(\parr)$$

Example 3.12. *As in Section 2, the following proof net is the translation of the linear lambda term* $(\lambda x.x)x$*.*

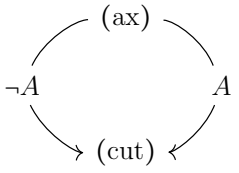


Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

This is the translation of the following sequent style proof.

$$\dfrac{\dfrac{\dfrac{}{\vdash \neg U, U}\,(\mathrm{ax}) \quad \dfrac{}{\vdash \neg U, U}\,(\mathrm{ax})}{\vdash \neg U, U \otimes \neg U, U}\,(\otimes) \quad \dfrac{\dfrac{}{\vdash \neg U, U}\,(\mathrm{ax})}{\vdash \neg U \parr U}\,\parr}{\vdash \neg U, U}\,(\mathrm{cut})$$

Example 3.13. *The following is a proof structure which is not a proof net.*



Suppose the above proof structure was equal to $T(\pi)$ where π is a sequent calculus proof. Then there exist sequent calculus proofs π_1, π_2 such that π is given by a cut rule with π_1 on the left branch and π_2 on the right branch. Since sequent calculus proofs are binary trees with leaves labelled by instances of the Axiom rule, it follows that π admits at least two instances of an Axiom rule, a contradiction.

3.2 The polynomial ring

Let k be a field.

Definition 3.14. *The polynomial ring of a formula A is the free commutative k-algebra P_A on the set U_A of unoriented atoms of A.*

Remark 3.15. *Some technical remarks:*

- *The algebra P_A does not depend on an ordering of the unoriented atoms in A. For example, we can take as our concrete realisation of P_A the symmetric algebra $\mathrm{Sym}(V_A)$ on the free vector space V_A on the set U_A.*
- *If we choose an ordering on U_A so that $U_A = \{X_1, \ldots, X_n\}$ then $P_A \cong k[X_1, \ldots, X_n]$. It will be convenient to present P_A in this way in the following, but unless we explicitly introduce an ordering on U_A no meaning should be associated to the indices i in X_i.*
- *The oriented atom $(X, +)$ appears twice in $A = (X, +) \otimes (X, +)$, so the sequence of unoriented atoms of A is X, X and the set of unoriented atoms is a disjoint union $\{X, X'\} = \{X\} \coprod \{X\}$ inside which we can distinguish the two copies of X (the precise nature of these copies is not important here). Hence $P_A \cong k[X, X']$.*

Definition 3.16. *The set of unoriented atoms of a proof structure π is the disjoint union $U_\pi = \coprod_{e \in E} U_{A_e}$ where E is the set of edges in π and for each $e \in E$ we write U_{A_e} for the set of unoriented atoms of the formula A_e labelling e.*

Definition 3.17. *The polynomial ring of a proof structure π is the free commutative k-algebra on the set U_π of unoriented atoms of π.*

The polynomial ring of a proof structure π is the polynomial ring with a generator for each pair consisting of an edge e in π and an unoriented atom of the formula labelling e. If we choose

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

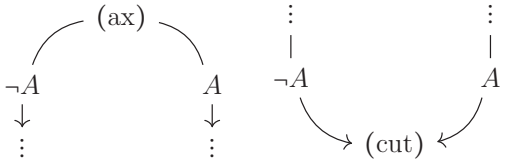
an ordering on the edges e of π, then we have an isomorphism

$$P_\pi = k\left[\coprod_{e\in E} U_{A_e}\right] \cong \bigotimes_{e\in E} P_{A_e}\,. \tag{25}$$

3.3 The ideal of a link

The ideal generated by a set S is denoted $\langle S\rangle$. For each type of link l, we define the *generating set* G_l and *link ideal* $I_l = \langle G_l\rangle$ in the polynomial ring generated by the set of unoriented atoms of formulas labelling edges incident at the link. Below $(X_1, x_1), \ldots, (X_n, x_n)$ is the sequence of oriented atoms of a formula A, and $(Y_1, y_1), \ldots, (Y_m, y_m)$ is the sequence of oriented atoms of B.

Definition 3.18 (Cut and Axiom). *If l is one of the links*

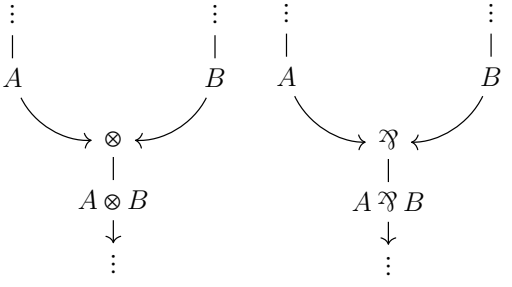


then in the polynomial ring generated by the disjoint union of unoriented atoms $U_A \coprod U_{\neg A}$

$$G_l = \left\{X_i - X_i'\right\}_{i=1}^n \tag{26}$$

where we write X_i' for the elements of $U_{\neg A}$.

Definition 3.19 (Tensor and Par links). *If l is one of the links*



then in polynomial ring generated by the disjoint union $U_A \coprod U_B \coprod U_{A\boxtimes B}$ (where $\boxtimes$ is either $\otimes$ or $\parr$)

$$G_l = \left\{X_i - X_i'\right\}_{i=1}^n \cup \left\{Y_j - Y_j'\right\}_{j=1}^m$$

where X_i', Y_j' denote elements of $U_{A\boxtimes B}$.

In Definition 3.17, we defined the polynomial ring P_π of a proof structure π. If l is a link of π, then in the obvious way we can identify G_l with a subset of P_π.

Definition 3.20. *The ideal of a link l in a proof structure π is the ideal $I_l = \langle G_l\rangle \subseteq P_\pi$ generated by G_l in P_π.*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Definition 3.21. *Let π be a proof structure with set of links $\mathcal{L}$. Then,*

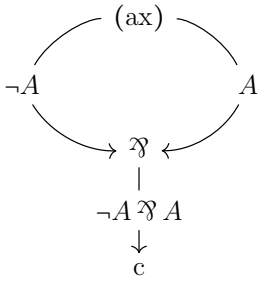
$$G_\pi := \bigcup_{l \in \mathcal{L}} G_l \,. \tag{27}$$

The defining ideal I_π is the ideal in P_π generated by G_π, or equivalently

$$I_\pi = \sum_{l \in \mathcal{L}} I_l \,. \tag{28}$$

The coordinate ring of π is the quotient $R_\pi = P_\pi / I_\pi$.

Example 3.22. *Consider the proof net π of Example 3.11 in which we set A to be the atomic proposition $A = (X, +)$*



(29)

The set of unoriented atoms of π is $\{X, X', X'', X'''\}$ where we associate X to A, X' to $\neg A$ and we let X'', X''' be the sequence of unoriented atoms of $\neg A \parr A$. Then

$$G_{(\mathrm{ax})} = \{X - X'\} \qquad G_{\parr} = \{X - X''', X' - X''\} \,.$$

Hence,

$$R_\pi = P_\pi / I_\pi = k[X, X', X'', X'''] / (I_{(\mathrm{ax})} + I_{\parr}) \cong k[X, X'] / I_{(\mathrm{ax})} \cong k[X]$$

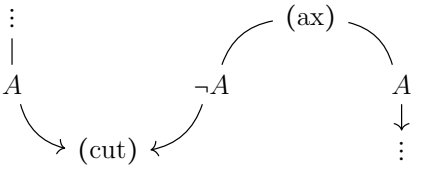
in which $X = X' = X'' = X'''$.

Example 3.23. *Consider the proof structure ρ of Example 3.13 (which is not a proof net) with $A = (X, +)$. Then $G_{(\mathrm{ax})} = \{X - X'\} = G_{(\mathrm{cut})}$ so $R_\rho = k[X, X'] / (X - X') \cong k[X]$.*

4. Reduction

We have now associated to each proof structure π a pair (I_π, P_π) consisting of a polynomial ring P_π generated by the unoriented atoms of π and an ideal I_π in this ring generated by differences $U - V$ where U, V are 'the same' variable on two different edges incident at a link in π. We now turn to the cut-elimination process for multiplicative proof nets and study how to think about a reduction $\pi \longrightarrow \pi'$ in terms of $(I_\pi, P_\pi), (I_{\pi'}, P_{\pi'})$.

Definition 4.1. *An a-redex in a proof structure π is a subgraph of the form*



(30)

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

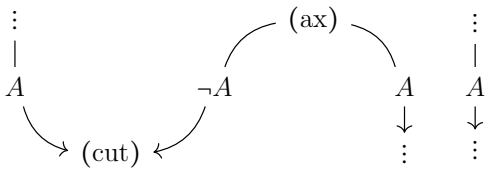
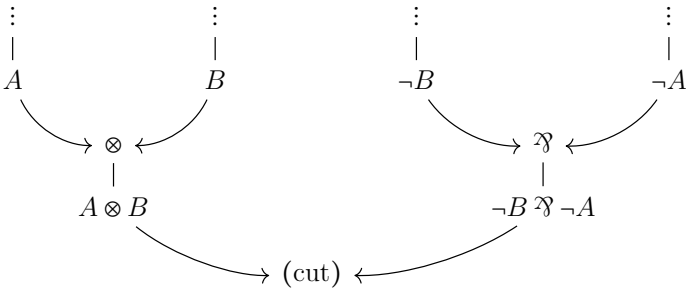


Figure 1. Reduction of an *a*-redex.

Definition 4.2. *An m-redex in a proof structure π is a subgraph of the form*

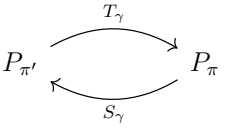


(31)

Definition 4.3. *A reduction $\gamma : \pi \longrightarrow \pi'$ between proof structures π, π' is an a-redex or m-redex γ in π such that reducing γ in π yields π', where*

- *if γ is an a-redex (30), the reduction is the proof structure obtained by replacing this subgraph by what is displayed on the right in* Figure 1.
- *if γ is an m-redex (31), the reduction is the proof structure obtained by replacing this subgraph by what is displayed on the bottom in* Figure 2.

Definition 4.4. *Let $\gamma : \pi \longrightarrow \pi'$ be a reduction of proof structures. We define morphisms S_γ, T_γ of k-algebras*



(32)

as follows: an unoriented atom in π' (resp. π) not associated to an edge in the subgraph γ is sent to itself in π (resp. π') by T_γ (resp. S_γ), other unoriented atoms are mapped according to the schematic in Figure 3 *(resp. Figure 4).*

It is easily checked that

Lemma 4.5. $S_\gamma T_\gamma = 1$.

The maps T_γ, S_γ are morphisms of the pairs $(I_\pi, P_\pi), (I_{\pi'}, P_{\pi'})$ in the following sense:

Proposition 4.6. *If $\gamma : \pi \longrightarrow \pi'$ is a reduction of proof structures then*

$$T_\gamma(I_{\pi'}) \subseteq I_\pi, \qquad S_\gamma(I_\pi) \subseteq I_{\pi'} \tag{33}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

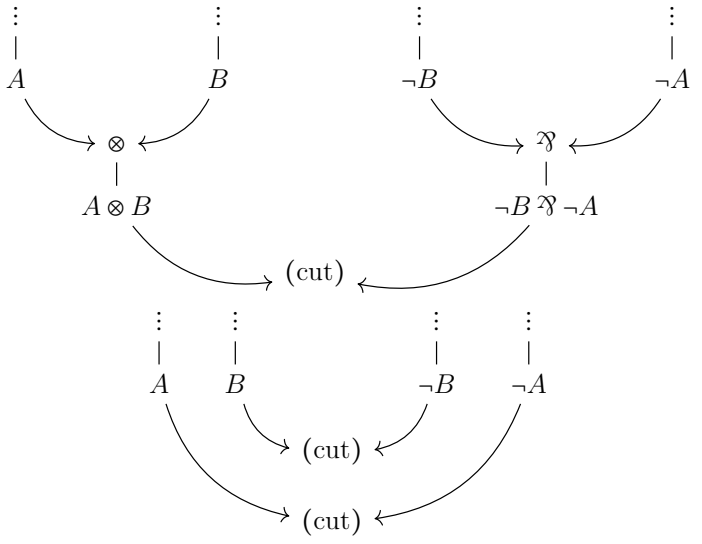


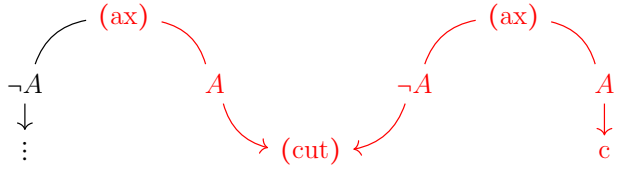
Figure 2. Reduction of a *m*-redex.

and the induced morphisms of k-algebras $\overline{T}_\gamma, \overline{S}_\gamma$ *making the squares in the following diagram commute*

$$\begin{array}{ccccc} I_\pi & \longrightarrow & P_\pi & \xrightarrow{p} & R_\pi \\ & & {\scriptstyle S_\gamma}\downarrow \uparrow{\scriptstyle T_\gamma} & & {\scriptstyle \overline{S}_\gamma}\downarrow \uparrow{\scriptstyle \overline{T}_\gamma} \\ I_{\pi'} & \longrightarrow & P_{\pi'} & \xrightarrow{p'} & R_{\pi'} \end{array} \tag{34}$$

are mutually inverse isomorphisms where p, p' *are the quotient maps.*

Proof. To prove $S_\gamma(I_\pi) \subseteq I_{\pi'}$, we first choose a link l in π. If l does not appear in the reduction (i.e. it is not the particular (ax) or (cut) involved in an a-redex nor the particular $\otimes, \parr$, (cut) in an m-redex) and no edge of l is in the redex, then every generator of I_l is in $I_{\pi'}$. If l does not appear in γ but l, γ share an edge, then $S_\gamma(I_l) \subseteq I_{\pi'}$ clearly holds in the case where γ is an m-redex. In the case of (ax)/(cut) (as e.g., the red (ax) link in Diagram (35) below)



(35)

the equations generating I_l are of the form $X - X'$, where X is an unoriented atom of the label of an edge in the redex and X' is an unoriented atom of the label of an edge occuring in l. By inspection of Figure 1, we see $S_\gamma(X - X') \in I_{\pi'}$.

Now, we consider the case where l is in the redex and γ reduces an a-redex. For all generators $X - X'$ of I_l, we have $S_\gamma(X) = S_\gamma(X')$ and so $S_\gamma(X - X') = 0 \in I_{\pi'}$. Now, consider the case where l is an m-redex. If l is the displayed tensor or par link then any generator $X - X'$ of $I_{\pi'}$ associated to l satisfies $S_\gamma(X) = S_\gamma(X')$ and so just as we had in the previous case we have $S_\gamma(X - X') = 0 \in I_{\pi'}$. The final case to consider is when l is the displayed cut link. Let $X - X'$ be a generator for I_l.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

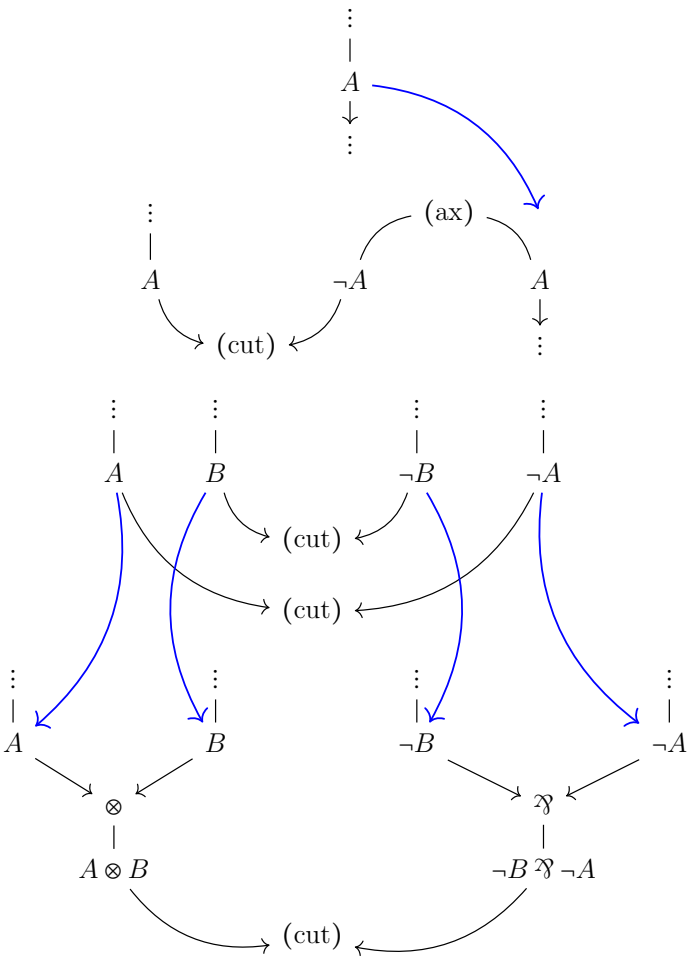


Figure 3. Schematic for T_γ.

Let l_1 denote the cut link in π with premises labelled $B, \neg B$ and l_2 denote the cut link with premises $A, \neg A$. We have that $S_\gamma(X) - S_\gamma(X') \in I_{l_1}$ if X is an undirected atom of A and $S_\gamma(X) - S_\gamma(X') \in I_{l_2}$ if X is an undirected atom of B. In either case, we have that $S_\gamma(X) - S_\gamma(X') \in I_{\pi'}$.

Now, we prove $T_\gamma(I_{\pi'}) \subseteq I_\pi$. A generator $X - X'$ of $I_{\pi'}$ is associated to a link l of π'. If γ is a reduction of an a-redex then l is also present in π and it is clear $T_\gamma(X - X') \in I_\pi$. If γ is a reduction of an m-redex and l is not one of the cut links shown in the bottom diagram of Figure 2 then the same reasoning applies. If l is one of these cut links the generator is either $X - X'$ for some unoriented atom X in A or $Y - Y'$ for some unoriented atom Y in B, where we are using the notation of Figure 2. There exists unoriented atoms X'', X''' of $A \otimes B$ and $\neg A \parr \neg B$, respectively, which satisfy $X - X'', X'' - X''', X''' - X' \in I_\pi$. Hence, $T_\gamma(X - X') = (X - X'') + (X'' - X''') + (X''' - X') \in I_\pi$, and similarly for $Y - Y'$. Hence, $T_\gamma(I_{\pi'}) \subseteq I_\pi$.

Hence, $\overline{S}_\gamma, \overline{T}_\gamma$ exist. To prove they are mutually inverse, it suffices to prove:

$$\overline{T}_\gamma \overline{S}_\gamma p = p \tag{36}$$

$$\overline{S}_\gamma \overline{T}_\gamma p' = p' \tag{37}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

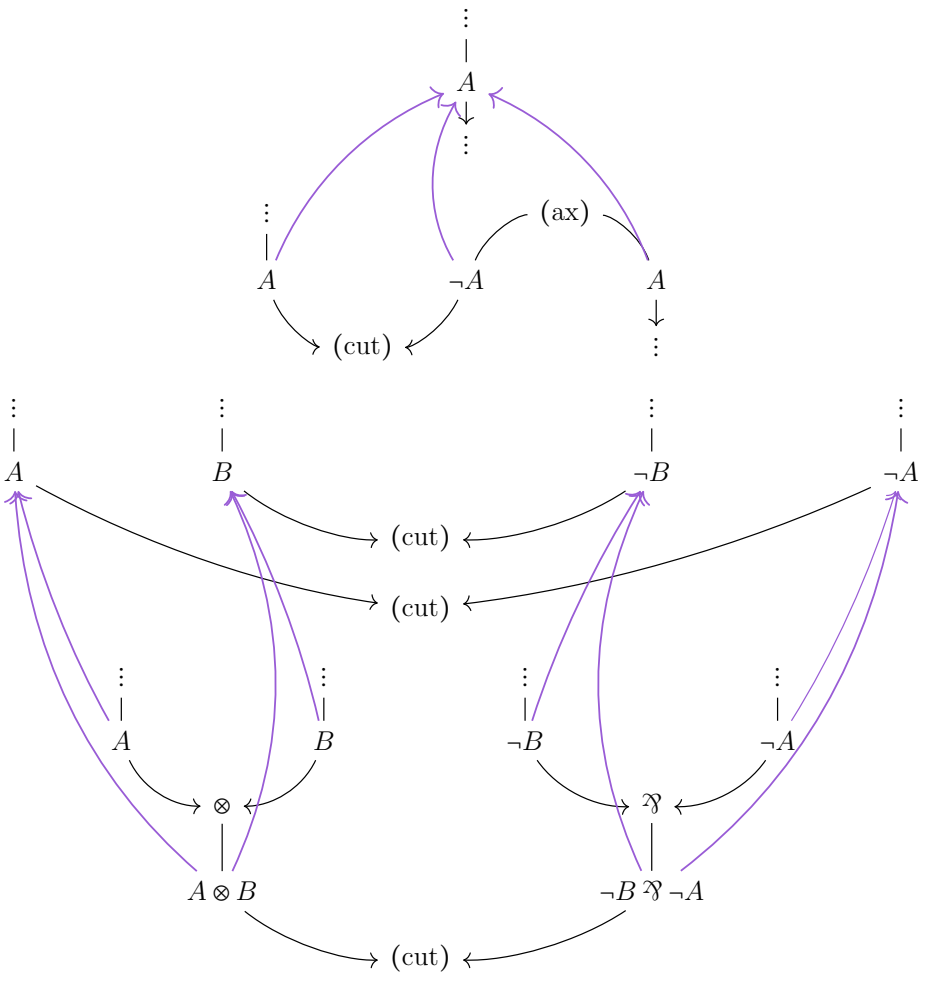


Figure 4. Schematic for S_γ.

as p, p' are surjective. By commutativity of (34), this is equivalent to showing $p'S_\gamma T_\gamma = p'$, $pT_\gamma S_\gamma = p$, or $p'(S_\gamma T_\gamma - 1) = 0, p(T_\gamma S_\gamma - 1) = 0$. It suffices to check this on generators, that is, on unoriented atoms. It is clear that $S_\gamma T_\gamma = 1$, however, we have $T_\gamma S_\gamma \neq 1$. The circumstances where this is the case is indicated schematically in Figure 5.

First, we consider the case in the first diagram of Figure 5. Let l be the displayed (ax) link and let X be an unoriented atom of A. Then, by inspection of the figure, we have $TS(X) - X \in I_l$. Now, let l' denote the displayed cut link. There exists an unoriented atom X' such that $X - X' \in I_{l'}$ and $TS(X) - X' \in I_l$. Hence, $TS(X) - X = TS(X) - X' + X' - X \in I_\pi$. This shows that $p(T_\gamma S_\gamma - 1)(X) = 0$.

Now, we consider the case in the second diagram of Figure 5. Let X be an unoriented atom of $A \otimes B$ and denote the displayed tensor link by l. There exists an unoriented atom X' either of A or of B such that $X - X' \in I_l$. In either case, we have that $X - TS(X) = X - X' \in I_l$ and so $p(T_\gamma S_\gamma - 1)(X) = 0$. The remaining cases are similar. □

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

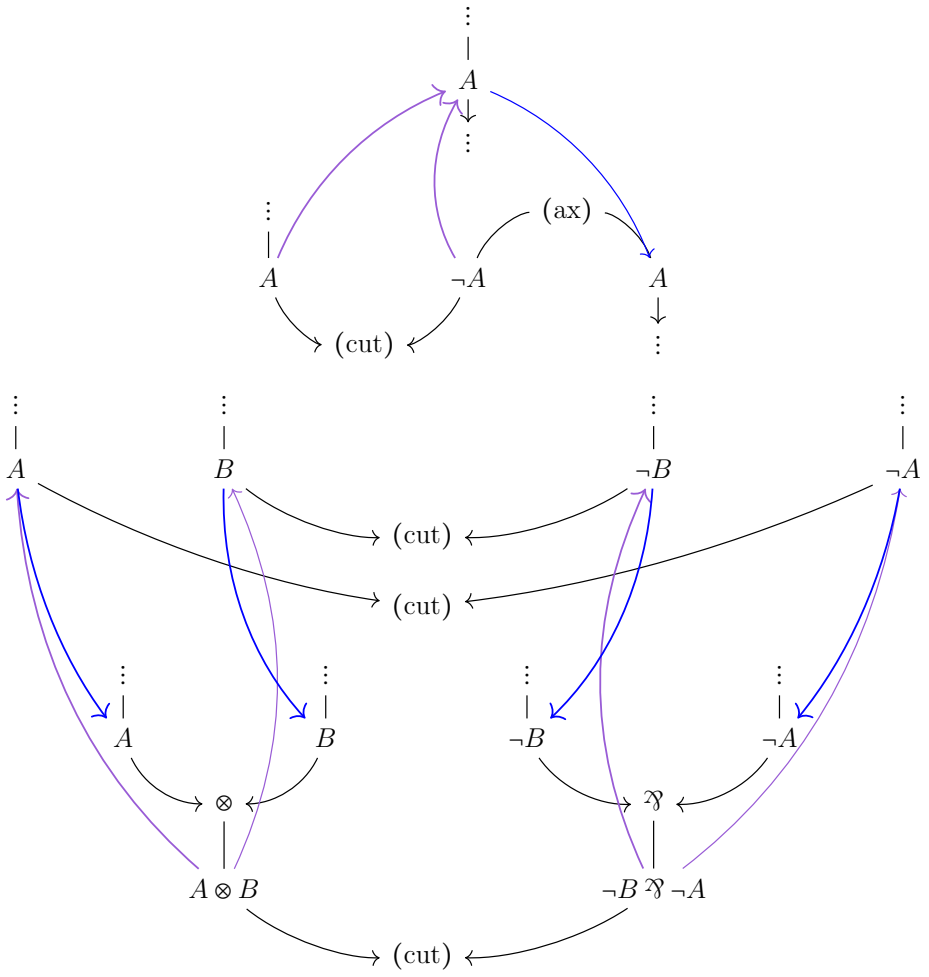


Figure 5. Understanding the composite *TS*.

Corollary 4.7. *Let $\gamma : \pi \longrightarrow \pi'$ be a reduction of proof structures. Then, $I_{\pi'} = T_\gamma^{-1}(I_\pi)$ or, identifying $P_{\pi'}$ as a subring of P_π with inclusion T_γ,*

$$I_{\pi'} = I_\pi \cap P_{\pi'}\,. \tag{38}$$

Proof. Proposition 4.6 shows that $T_\gamma(I_{\pi'}) \subseteq I_\pi$, so we have $I_{\pi'} \subseteq T_\gamma^{-1}(I_\pi)$. If $T_\gamma(f) \in I_\pi$ then by Lemma 4.5 $f = S_\gamma T_\gamma(f) \in S_\gamma(I_\pi) \subseteq I_{\pi'}$ giving the reverse inclusion. □

It is not difficult to understand the effect of taking the quotient of P_π by the ideal I_π. All those unoriented atoms connected across links (in the sense that $U - V$ appears in G_π) are identified in R_π with each other. This common image can be thought of as the *path* which connects all these differently named occurrences of the 'same' variable; such paths are sometimes called persistent paths (see Definition 4.12).

Definition 4.8. *Let π be a proof structure and $U, V \in U_\pi$. We write $U \sim V$ if $U - V \in G_\pi$ or $V - U \in G_\pi$. Let $\approx$ be the equivalence relation on U_π generated by $\sim$.*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Lemma 4.9. *If $U, V \in U_\pi$, then we have $U \approx V$ if and only if $U - V \in I_\pi$.*

Proof. If $U \approx V$, then clearly $U - V \in I_\pi$. If $U - V \in I_\pi$, then we can write $U - V = \sum_j q_j(U_j - V_j)$ with $U_j - V_j \in G_\pi$. Comparing the degrees of the left and right hand side, q_j may be assumed to be in k, and in fact in $\{1, -1\}$. From this sum, we deduce a chain of $U_j \sim V_j$ relations leading from V to U as claimed. □

Lemma 4.10. *Let $\gamma : \pi \longrightarrow \pi'$ be a reduction of proof structures. Then, if $U, V \in U_{\pi'}$ have $T_\gamma(U) \approx T_\gamma(V)$ in U_π if and only if $U \approx V$ in $U_{\pi'}$.*

Proof. Follows from Lemma 4.9 and Corollary 4.7. □

Recall that we call a formula A labelling an edge incident at a conclusion of π a *conclusion* of π. If the sequence of oriented atoms of A is $(U_1, u_1), \ldots, (U_n, u_n)$ then we say that U_i is *positive* in A if $u_i = +$ and otherwise it is *negative*.

Proposition 4.11. *Let π be a proof net with single conclusion A, and let*

$$(Z_1, z_1), \ldots, (Z_n, z_n) \tag{39}$$

be the sequence of oriented atoms of A. Then $n = 2m$ is even, there are an equal number of positive and negative atoms, and if $\mathbf{U} = U_1, \ldots, U_m$ denotes the subsequence of positive unoriented atoms and $\mathbf{V} = V_1, \ldots, V_m$ the subsequence of negative unoriented atoms then

(i) *The inclusions $k[\mathbf{U}] \longrightarrow P_\pi$ and $k[\mathbf{V}] \longrightarrow P_\pi$ followed by the quotient $P_\pi \longrightarrow R_\pi$ are isomorphisms β_+, β_- as in the diagram*

$$\begin{array}{ccc} k[\mathbf{U}] & & \\ \downarrow & \searrow^{\beta_+} & \\ P_\pi & \twoheadrightarrow & R_\pi \\ \uparrow & \nearrow_{\beta_-} & \\ k[\mathbf{V}] & & \end{array} \tag{40}$$

(ii) *The composite $\beta_-^{-1}\beta_+ : k[\mathbf{U}] \longrightarrow k[\mathbf{V}]$ is*

$$\beta_-^{-1}\beta_+(U_i) = V_{\sigma(i)}, \quad 1 \leq i \leq m \tag{41}$$

for some permutation σ_π of $\{1, \ldots, m\}$.

(iii) *Each equivalence class of the relation $\approx$ is the underlying set of a sequence*

$$\mathscr{P} = (Z_1, \ldots, Z_r) \tag{42}$$

where for some $1 \leq i \leq m$, we have $Z_1 = V_{\sigma(i)}, Z_r = U_i$ and $Z_i \sim Z_{i+1}$ for $1 \leq i < r$.

Proof. We define the *weight* of a proof net π to be the sum, over all (cut) links on some formula $A, \neg A$, of the weight $|A|$ of A which is defined by recursively as follows: if A is an atomic proposition, then $|A| = 1$, if $A = B \otimes C$ with $\otimes \in \{\otimes, \parr\}$ then $|A| = |B| + |C| + 1$, and $|\neg A| = |A|$. Note that the weight of a proof net strictly decreases under cut-elimination.

We prove the proposition by induction on the weight. All statements can be treated one persistent path at a time by Lemma 4.10. In the base case, a persistent path proceeds from a negative unoriented atom V in the conclusion up through $\otimes, \parr$ links to an (ax) link and then downwards

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

through $\parr, \otimes$ links to the positive unoriented atom U in the conclusion. If the unoriented atoms involved are $V = Z_1, \ldots, Z_r = U$, then (40) is

$$\begin{array}{ccc} k[U] & & \\ \downarrow & \searrow^{\beta_+} & \\ k[Z_1, \ldots, Z_r] & \twoheadrightarrow & k[\mathbf{Z}]/I \\ \uparrow & \nearrow_{\beta_-} & \\ k[V] & & \end{array} \tag{43}$$

where $I = (Z_2 - Z_1, \ldots, Z_r - Z_{r-1})$ so it is clear that β_+, β_- are isomorphisms, and that $\beta_-^{-1}\beta_+(U) = V$ is the map sending the end of each persistent path to its beginning, a permutation of the set of unoriented atoms in the conclusion.

For the inductive hypothesis, suppose the proposition is true of π' and that there is a reduction $\gamma : \pi \longrightarrow \pi'$. Note π' has the same conclusion as π. The claims follow from commutativity of the diagram

$$\begin{array}{ccccc} P_\pi & & \longrightarrow & & R_\pi \\ \uparrow & \nwarrow & & \nearrow^{\beta_+} & \uparrow \\ T_\gamma & & k[\mathbf{U}] & & \cong\ \overline{T}_\gamma \\ \downarrow & \swarrow & & \searrow_{\beta_+} & \downarrow \\ P_{\pi'} & & \longrightarrow & & R_{\pi'} \end{array} \tag{44}$$

and a similar one for $\mathbf{V}$. □

Definition 4.12. *Let π be a proof net with single conclusion A. The sequences $\mathscr{P}$ of (42) whose underlying sets are the equivalence classes of $\approx$ are called persistent paths.*

Example 4.13. *The proof net π in Figure 6 is the canonical detour (6) of the introduction with $B = A, \multimap$ written in terms of $\parr$ and the multiple conclusions combined with $\parr$ to a single conclusion. With $A_i = A = (X, +)$ for $1 \leq i \leq 12$ the set of unoriented atoms of π is a disjoint union of twelve copies of $\{X\}$ and we name the atoms $X_1, \ldots, X_{12}$.*

In the notation of the proposition, $\mathbf{U} = \{X_{12}\}, \mathbf{V} = \{X_1\}$ and the single persistent path of π is $X_1, X_2, \ldots, X_{12}$. Note that the persistent path goes 'through' the (cut) *link twice, once in each direction.*

Remark 4.14. *By the definition of persistent paths, it follows from Proposition 4.6 that given any reduction $\gamma : \pi \longrightarrow \pi'$ and persistent path of a proof structure π, there is an associated persistent path $\mathscr{P}'$ in π' given by removing the variables which occur in π but not in π' from $\mathscr{P}$. This will be used in the proofs of Theorems 7.2, 7.11.*

Definition 4.15. *A reduction sequence $\Gamma : \pi \longrightarrow \pi'$ between proof structures π, π' is a nonempty sequence of reductions (Definition 4.3)*

$$\pi = \pi_1 \xrightarrow{\gamma_1} \cdots \xrightarrow{\gamma_{n-1}} \pi_n = \pi'. \tag{45}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

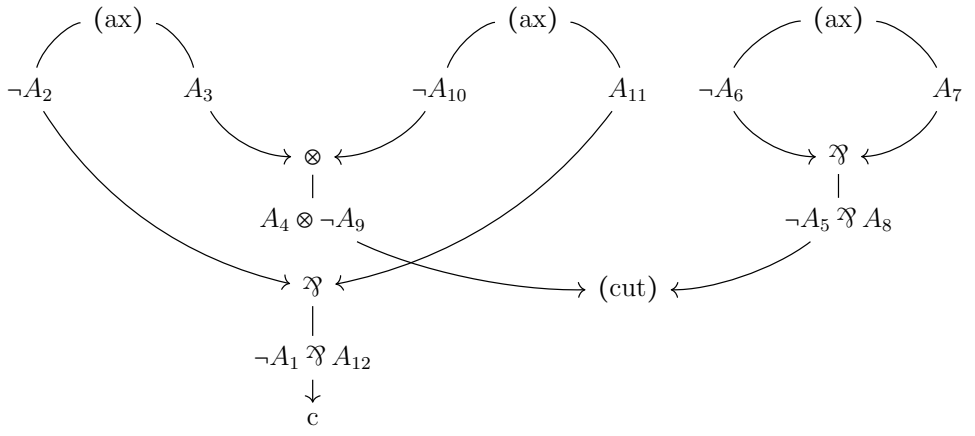


Figure 6. The canonical detour as a proof net, with $A_i = A = (X, +)$ for all i.

This induces a sequence of k-algebra morphisms

$$P_{\pi'} \xrightarrow{T_{\gamma_{n-1}}} \cdots \xrightarrow{T_{\gamma_1}} P_{\pi} \tag{46}$$

the composite of which we denote by $T_\Gamma : P_{\pi'} \longrightarrow P_\pi$.

5. Elimination Theory

5.1 Buchberger's algorithm

The Buchberger algorithm computes, given a generating sequence for an ideal I, a Gröbner basis for that ideal. There are a number of variations on this algorithm in the literature. We present here an optimisation of the algorithm of (Cox et al. 2015, §2.10 Theorem 9) for the purposes of Elimination Theory, using a form of Euclidean Division with 'early stopping' (Algorithm 1).

In this section, k is a field and $k[X_1, \ldots, X_n]$ a polynomial ring with a monomial order. We use the terminology of leading terms LT, leading monomials LM, leading coefficient LC, and multidegree (Cox et al. 2015, §2).

Lemma 5.1. *Let* $G = (f_1, \ldots, f_s)$ *be a sequence of elements of* $k[X_1, \ldots, X_n]$. *Given* $f \in k[X_1, \ldots, X_n]$ *let* $(q_1, \ldots, q_s, r)$ *be the output of Algorithm 1. Then,*

$$f = q_1 f_1 + \cdots + q_s f_s + r \tag{47}$$

and if $r \neq 0$ *the leading term of* r *is not divisible by any* $\mathrm{LT}(f_i)$ *with* $1 \leq i \leq s$. *Furthermore, if* $q_i f_i \neq 0$, *then* $\mathrm{multideg}(f) \geq \mathrm{multideg}(q_i f_i)$.

Proof. The expression $f = \sum_{i=1}^{s} q_i f_i + p + r$ holds at the end of every iteration of the outer while loop, and it terminates with $p = 0$ so (47) holds. If the algorithm returns $r \neq 0$, then it came from $r \leftarrow p$ with DivOcc = `false`, so the claim about divisibility of LT (r) follows. Termination and the other property follow as in the proof of the standard division algorithm (Cox et al. 2015, §2.3). □

Definition 5.2. *Let* $G = (f_1, \ldots, f_s)$ *be a sequence of polynomials and* f *a polynomial. We denote by* $\mathrm{div}(f, G)$ *(resp.* $\mathrm{div}_{es}(f, G)$*) the remainder* r *produced by the Euclidean division (resp. Algorithm 1).*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Algorithm 1 Euclidean Division with Early Stopping.

Require: $(f_1, \ldots, f_s), f$
  $p \leftarrow f$
  $q_1, \ldots, q_s \leftarrow 0, \ldots 0$
  $r \leftarrow 0$
  while $p \neq 0$ **do**
    DivOcc $\leftarrow$ `False`
    $i \leftarrow 1$
    while $i \leq s$ and DivOcc $=$ `false` **do**
      if LT $f_i |$ LT p **then**
        $q_i \leftarrow q_i +$ LT $p /$ LT f_i
        $p \leftarrow p -$ (LT $p /$ LT $f_i) f_i$
        DivOcc $\leftarrow$ `True`
      else
        $i \leftarrow i + 1$
      end if
    end while
    if DivOcc $=$ `false` **then**
      $r \leftarrow p$
      $p \leftarrow 0$
    end if
  end while
return $(q_1, \ldots, q_s, r)$

To state our variation of Buchberger's algorithm, we need some notation. For $i \neq j$ set

$$[i, j] = \begin{cases} (i, j) & \text{if } i \leq j \\ (j, i) & \text{if } i \geq j \end{cases}$$

Recall that the *S-polynomial* of $g, h \in k[X_1, \ldots, X_n]$ is

$$S(g, h) := \frac{X^\gamma}{\mathrm{LT}\,(g)} g - \frac{X^\gamma}{\mathrm{LT}\,(h)} h\,. \tag{48}$$

where $X^\gamma = \mathrm{LCM}\,(\mathrm{LM}\,(g), \mathrm{LM}(h))$. The only significant difference between Algorithm 2 below and the form of Buchberger's algorithm in (Cox et al. 2015, §2.10 Theorem 9) is that we replace the line $S \leftarrow \mathrm{div}\,(S(f_i, f_j), G)$ using normal Euclidean division with the line $S \leftarrow \mathrm{div}_{es}\,(S(f_i, f_j), G)$ using the division algorithm with early stopping. We also divide S by its leading coefficient $\mathrm{LC}\,(S)$ before adding it to G, but this is insignificant.

Definition 5.3. *Given a set $G = \{f_1, \ldots, f_s\}$ of polynomials we say that f reduces to zero modulo G, written $f \longrightarrow_G 0$ if f can be written in the form $f = a_1 f_1 + \cdots + a_s f_s$ such that whenever $a_i f_i \neq 0$ we have* $\mathrm{multideg}\,(f) \geq \mathrm{multideg}\,(a_i f_i)$ *(see [Cox et al. 2015, §2.9 Definition 1]).*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Algorithm 2 Buchberger with Early Stopping.

Require: $F = (f_1, \ldots, f_s)$, returns a Gröbner basis for $\langle f_1, \ldots, f_s \rangle$.

$B \leftarrow \{(i,j) \mid 1 \le i < j \le s\}$

$G \leftarrow F$

$t \leftarrow s$

while $B \neq \emptyset$ **do**

Let $(i,j) \in B$ be first in the lexicographic order.

if LCM (LM (f_i), LM (f_j)) $\neq$ LM (f_i) LM (f_j) **and** Criterion(f_i, f_j, B) **is false then**

$S \leftarrow \mathrm{div}_{es}\,(S(f_i, f_j), G)$

if $S \neq 0$ **then**

$t \leftarrow t + 1$

$f_t \leftarrow \frac{1}{\mathrm{LC}\,(S)} S$

$G \leftarrow G \cup \{f_t\}$

$B \leftarrow B \cup \{(i,t) \mid 1 \le i \le t-1\}$

end if

end if

$B \leftarrow B \setminus \{(i,j)\}$

end while

return G

where Criterion(f_i, f_j, B) is true provided that there is some $k \notin \{i,j\}$ for which the pairs $[i,k]$ and $[j,k]$ are *not* in B and LM (f_k) divides LCM (LM (f_i), LM (f_j)).

Proposition 5.4. *Algorithm 2 terminates and the output is a Gröbner basis.*

Proof. Algorithm 2 is a variation on (Cox et al. 2015, §2.10 Theorem 9) and the proof here is similar. Suppose that in some pass through the while loop, G is enlarged to $G \cup \{f_t\}$. By Lemma 5.1, we have that LT (f_t) is not divisible by LT (f) for any $f \in G$, and therefore, $\langle \mathrm{LT}\,(G) \rangle$ is a proper subset of $\langle \mathrm{LT}\,(G \cup \{f_t\}) \rangle$. Since $k[X_1, \ldots, X_n]$ is Noetherian this means that at some point G stops growing, but then in each subsequent step the size of B decreases by one. Hence, the Algorithm terminates in finitely many steps.

To prove the output is a Gröbner basis, we first establish the following claim: at the end of every pass through the while loop, B has the property that if $1 \le i < j \le t$ and $(i,j) \notin B$, then

$$S(f_i, f_j) \longrightarrow_G 0 \quad \text{or} \quad \text{Criterion}(f_i, f_j, B) \text{ holds}. \tag{49}$$

This is initially true since B is the set of all possible pairs. To prove the claim, we need to argue that if it holds for B, then it continues to hold when B changes to B'.

To prove this, assume $(i,j) \notin B'$. If $(i,j) \in B$ then $B' = B \setminus \{(i,j)\}$. In the step, before (i,j) is removed if either $\mathrm{LCM}(\mathrm{LT}\,(f_i), \mathrm{LT}\,(f_j)) = \mathrm{LT}\,(f_i)\,\mathrm{LT}\,(f_j)$ or Criterion(f_i, f_j, B) is true, then one argues as in the original proof from Cox, Little, and O'Shea (2015). If not, then the algorithm computes $S \leftarrow \mathrm{div}_{es}\,(S(f_i, f_j), G,)$. If $S = 0$, then $S(f_i, f_j) \longrightarrow_G 0$ by Lemma 5.1. If $S \neq 0$ then we enlarge G to $G' = G \cup \{\frac{1}{\mathrm{LC}\,(S)} S\}$. Examining the division algorithm with early stopping it clear that $\mathrm{div}_{es}\,(S(f_i, f_j), G') = 0$ so $S(f_i, f_j) \longrightarrow_{G'} 0$ and (49) holds for B', G'.

It remains to consider the case $(i,j) \notin B$. Then, (49) holds for B. Let G, G' be the state of the sequence of polynomials in the two steps, so $G' = G$ or G' has one additional element. If $S(f_i, f_j) \longrightarrow_G 0$ then $S(f_i, f_j) \longrightarrow_{G'} 0$. If Criterion($f_i, f_j, B$) holds then there is some $k \in \{i,j\}$ with

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

$k \leq |G|$ such that $[i,k]$, $[j,k]$ are not in B and LT (f_k) divides LCM (LT (f_i), LT (f_j)). Any new pairs in B' involve $|G|+1$, so $[i,k], [j,k] \notin B'$ and hence Criterion(f_i, f_j, B') holds.

The argument that G is a Gröbner basis when $B = \emptyset$ goes through as before. □

Definition 5.5. *Given a sequence $F = (f_1, \ldots, f_s)$ of polynomials and a monomial order $<$ on $k[X_1, \ldots, X_n]$, we denote by $\mathbb{B}_{es}(F, <)$ the output of Algorithm 2.*

We state the main theorem of Elimination Theory. Let $\mathbf{X} = \{X_1, \ldots, X_n\}$ and $\mathbf{Y} = \{Y_1, \ldots, Y_m\}$ be variables. We suppose given an ideal $I \subseteq k[\mathbf{X}, \mathbf{Y}]$ and we are interested to know the equations between the $\mathbf{X}$-variables that are implied by the equations in I. We call this set of equations the elimination ideal:

Definition 5.6. *The elimination ideal of I is the ideal $I \cap k[\mathbf{X}]$ in $k[\mathbf{X}]$.*

In the situation of Section 2, we can think of the $\mathbf{X}$-variables as living at the 'boundary' of the proof net, and $\mathbf{Y}$ as being the 'interior' variables.

Theorem 5.7 (The Elimination Theorem). *Let $I \subseteq k[\mathbf{X}, \mathbf{Y}]$ be an ideal and G a Gröbner basis of I with respect to a lexicographic order where $X_i < Y_j$ for all $1 \leq i \leq n, 1 \leq j \leq m$. Then, $G \cap k[\mathbf{X}]$ is a Gröbner basis for $I \cap k[\mathbf{X}]$.*

Proof. See (Cox et al. 2015, §3.1 Theorem 2). □

5.2 Buchberger and falling roofs

We introduce graphical notation for understanding the operation of Algorithm 2 on a simple class of ideals. Fix a polynomial ring $k[X_1, \ldots, X_n]$, let $<$ be a total order on the set $\{X_1, \ldots, X_n\}$ and take the lexicographic monomial order determined by $<$. Let σ be the permutation uniquely defined by

$$X_{\sigma^{-1}1} < X_{\sigma^{-1}2} < \cdots < X_{\sigma^{-1}n}\,.$$

The position of X_i in this sequence is $\sigma(i)$. We view σ as assigning a *height* to variables:

Definition 5.8. *The realisation of $<$ is the oriented graph $\mathscr{R}_<$ with vertices*

$$\big\{(i, \sigma i) \mid 1 \leq i \leq n\big\} \subseteq \mathbb{R}^2 \tag{50}$$

with an edge between $(i, \sigma i)$, $(i+1, \sigma(i+1))$ for $i < n$, and $(i, \sigma i)$ decorated with X_i. The orientation of the edge is from $(i, \sigma i)$ to $(j, \sigma j)$ if $X_i < X_j$.

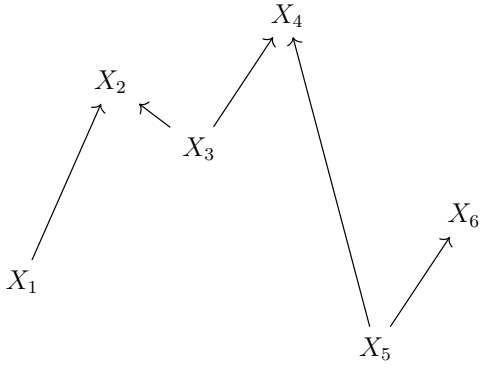
Definition 5.9. *A $<$-graph is an oriented graph on the vertex set (50) with the property that if there is an edge from a vertex $(i, \sigma i)$ decorated with X_i to a vertex decorated with X_j then $X_i < X_j$.*

Definition 5.10. *A $<$-graph is linear if every vertex has valence at most two.*

Clearly, $\mathscr{R}_<$ is a linear $<$-graph. In this section, we simply write *graph* for $<$-graph, since the ordering $<$ on the variables is fixed. We write $e : X_i \longrightarrow X_j$ for an edge from a vertex decorated with X_i to a vertex decorated with X_j.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Example 5.11. *Let $X_1, \ldots, X_6$ be ordered by $X_5 < X_1 < X_6 < X_3 < X_2 < X_4$. Then $\mathscr{R}_<$ is*



(51)

Remark 5.12. *The realisation involves two orderings on the variables: the order reading along the x-axis (the $X_1, \ldots, X_n$ order) and the order reading along the y-axis (the $<$-order). The realisation is a graphical presentation of the relation between these orders.*

Definition 5.13. *A roof in a graph $\mathscr{S}$ is an ordered pair (e, e') of edges $e : X_i \longrightarrow X_l, e' : X_k \longrightarrow X_l$ with the same endpoint X_l and $X_i < X_k$. We call X_l the tip of the roof.*

Definition 5.14. *Given a graph $\mathscr{S}$, we define*

$$G_{\mathscr{S}} = \left\{ X_j - X_i \mid e : X_i \longrightarrow X_j \text{ is an edge in } \mathscr{S} \right\}.$$

Definition 5.15. *Let $\mathscr{S}$ be a graph with edges $e : X_i \longrightarrow X_j, e' : X_k \longrightarrow X_l$. Then, $e < e'$ if the pair (X_j, X_i) precedes (X_l, X_k) lexicographically, that is, either $X_j < X_l$ or $X_j = X_l$ and $X_i < X_k$. This is a total order on the set of edges of $\mathscr{S}$.*

The map sending $e : X_i \longrightarrow X_j$ to $X_j - X_i$ is a bijection between the edge set of $\mathscr{S}$ and $G_{\mathscr{S}}$ and we give the latter set the total order induced by this bijection and the total order on the former set just defined.

Definition 5.16. *A roof (e, e') precedes a roof (d, d') if $e < d$ or $e = d$ and $e' < d'$.*

We will only consider roofs in linear graphs: since in a linear graph distinct roofs have distinct tips, $(e, e') < (d, d')$ simply means that $e < d$, or what is the same, the tip of the first roof is lower in the realisation than the tip of the second.

We now explain the operation of Algorithm 2 on input $(G_{\mathscr{S}}, <)$ for a linear graph $\mathscr{S}$ by giving an 'isomorphic' algorithm (Algorithm 3 below), which is organised around roofs and is easier to understand. The algorithm works by marking a subset of edges in a graph $\mathscr{N}$ derived from $\mathscr{S}$ as *dead* (an edge i.e., not dead is *live*). A roof is *live* if both the edges in it are live. We indicate dead edges diagrammatically by dotted lines.

Proposition 5.17. *For any linear graph $\mathscr{S}$:*

(i) *Algorithm 3 terminates.*

(ii) *At every step, $\mathscr{N}$ is a graph and the subgraph of live edges in $\mathscr{N}$ is linear.*

(iii) *If the output of of Algorithm 2 on input $(G_{\mathscr{S}}, <)$ is G and the output of Algorithm 3 on input $\mathscr{S}$ is $\mathscr{N}$ then $G = G_{\mathscr{N}}$.*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Algorithm 3 Falling Roofs.

Require: Linear graph $\mathscr{S}$

$\mathscr{N} \leftarrow \mathscr{S}$

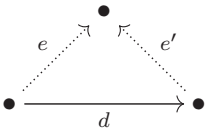
Mark all edges in $\mathscr{N}$ as live

while $\mathscr{N}$ contains a live roof **do**

$(e, e') \leftarrow$ the first live roof in $\mathscr{N}$

Mark e, e' as dead

If it does not exist, add to $\mathscr{N}$ a live edge d as shown below:



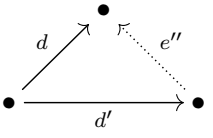
(52)

while d is part of a live roof in $\mathscr{N}$ **do**

if (d, e'') is a live roof in $\mathscr{N}$, **then**

Mark e'' as dead

if it does not exist, add to $\mathscr{N}$ a live edge d' as shown below:



(53)

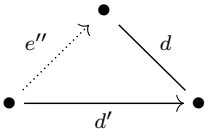
Remove d from $\mathscr{N}$

$d \leftarrow d'$

else if (e'', d) is a live roof in $\mathscr{N}$ **then**

Mark e'' as dead

if it does not exist, add to $\mathscr{N}$ a live edge d' as shown below:



(54)

Remove d from $\mathscr{N}$

$d \leftarrow d'$

end if

end while

end while

return $\mathscr{N}$

Proof. Given an edge $e : X_i \longrightarrow X_j$, we write f_e for $X_j - X_i$. We prove (iii) and along the way (i), (ii) by stepping through Algorithm 2 on input $G_{\mathscr{S}}$ (with intermediate variables G, B, S) and Algorithm 3 on input $\mathscr{S}$ (with intermediate variable $\mathscr{N}$). At each step, let $B_{rel} \subseteq B$ denote the set of pairs (e, e') which are *relevant* in the sense that

$$\mathrm{LCM}(\mathrm{LM}(f_e), \mathrm{LM}(f_{e'})) \neq \mathrm{LM}(f_e)\,\mathrm{LM}(f_{e'}) \textbf{ and } \mathrm{Criterion}(f_e, f_{e'}, B) \textbf{ is false}.$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

We will prove by induction on the number of steps through Algorithm 2 that at each step the following proposition **IP** holds:

IP: $\mathcal{N}$ is a graph, the live edges in $\mathcal{N}$ form a linear graph, no dead edge in $\mathcal{N}$ ends where a live edge begins, $G = G_{\mathcal{N}}$, and B_{rel} is the set of live roofs in $\mathcal{N}$.

It is clear that these conditions are all initially true, except perhaps the last claim about B_{rel}. To examine this condition at the beginning of the algorithm, let $e : X_i \longrightarrow X_j$ and $e' : X_k \longrightarrow X_l$ be edges in $\mathcal{N} = \mathcal{S}$. Then $\mathrm{LT}(f_e) = X_j$, $\mathrm{LT}(f_{e'}) = X_l$ and

$$\mathrm{LCM}(\mathrm{LM}(f_e), \mathrm{LM}(f_{e'})) = \mathrm{LCM}(X_j, X_l) = \begin{cases} X_j X_l & X_j \neq X_l \\ X_j & X_j = X_l \end{cases}$$

so, $\mathrm{LCM}\,(\mathrm{LM}\,(f_e), \mathrm{LM}\,(f_{e'})) \neq \mathrm{LM}\,(f_e)\,\mathrm{LM}\,(f_{e'})$ if and only if $X_j = X_l$, that is, if and only if e, e' have a common endpoint in $\mathcal{N}$. Suppose this is the case. Then, Criterion$(f_e, f_{e'}, B)$ is true if there exists some edge $e'' \notin \{e, e'\}$ for which $[e, e'']$, $[e', e'']$ are *not* in B and the edges e, e', e'' all have the same endpoint. But since $\mathcal{N}$ is linear this is impossible, so the criterion is false. Hence, a pair (e, e') is in B_{rel} if and only if it is a roof in $\mathcal{S}$, and since initially all roofs are live this proves the claim.

Suppose that **IP** is true up to some step of the outer **while** loops of the two algorithms. Let $\mathcal{N}_0$ denote the state of $\mathcal{N}$ at the beginning of this step of Algorithm 3. Algorithm 2 looks in order through pairs (e, e') for one that is relevant, or what is the same by the inductive hypothesis, for a live roof in $\mathcal{N}$. Let (e, e') be a live roof with $e : X_i \longrightarrow X_j$, $e' : X_k \longrightarrow X_j$. Algorithm 2 computes the S-polynomial

$$S := S(f_e, f_{e'}) = f_e - f_{e'} = X_k - X_i \tag{55}$$

which is f_d for $d : X_i \longrightarrow X_k$ in Algorithm 3. The next step of Algorithm 2 is to compute $\mathrm{div}_{es}\,(S, G)$. By **IP,** we have $G = G_{\mathcal{N}}$. A polynomial in G with leading term X_k is an edge in $\mathcal{N}$ ending at X_k, which must be live by **IP** as e' was live. If there is no polynomial in G with leading term X_k, then the division algorithm (with early stopping) returns S, $\frac{1}{\mathrm{LC}\,(S)} S$ is added to G, d is added to $\mathcal{N}$, e, e' are marked as dead and the step is complete. It is straightforward to check that **IP** still holds.

Otherwise there is a polynomial f in G with leading term X_k. If $f = S$, then $\mathrm{div}_{es}\,(S, G) = 0$ so nothing is added to G by Algorithm 2, nothing is added to $\mathcal{N}$ by Algorithm 3 and the step is complete; again **IP** still holds. We are left with the case where $f \neq S$ exists in G with leading term X_k. For this, we analyse the inner **while** loop in Algorithm 3. The key observations are that at the beginning of each step of this inner loop:

(a) d is part of at most one live roof in $\mathcal{N}$.

(b) there is no dead edge in $\mathcal{N}$ ending where d ends.

Both claims are true for the initial d by **IP**. We call d Type 1 if it comes from d' in (53) and Type 2 if it comes from d' in (54). If d is Type 1 both claims (a),(b) follow from **IP**. In Type 2, we have to analyse the previous steps; if the previous step was Type 2 (or there was no previous step) then there is a live edge in $\mathcal{N}_0$ beginning where d ends and from this we deduce (a),(b). Otherwise there are some Type 1 steps intervening between the current step and the most recent Type 2 step (or the original roof (52) if there are no previous Type 2 steps). However, in Type 1, steps the d, d' edges start from the same vertex, so the most recent Type 2 step (or original roof) provides a live edge beginning where d ends, from which we deduce (a),(b) again from **IP**.

We now prove that the inner **while** loop of Algorithm 3 aligns with the division step of Algorithm 2. By (a), there is a unique polynomial $f = f_{e''}$ in G with the same leading term as f_d. The division algorithm replaces $S = f_d$ by the dividend

$$S \leftarrow S - (\mathrm{LT}(S)/\,\mathrm{LT}(f_{e''}))f_{e''} = f_d - f_{e''} = \pm f_{d'} \tag{56}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

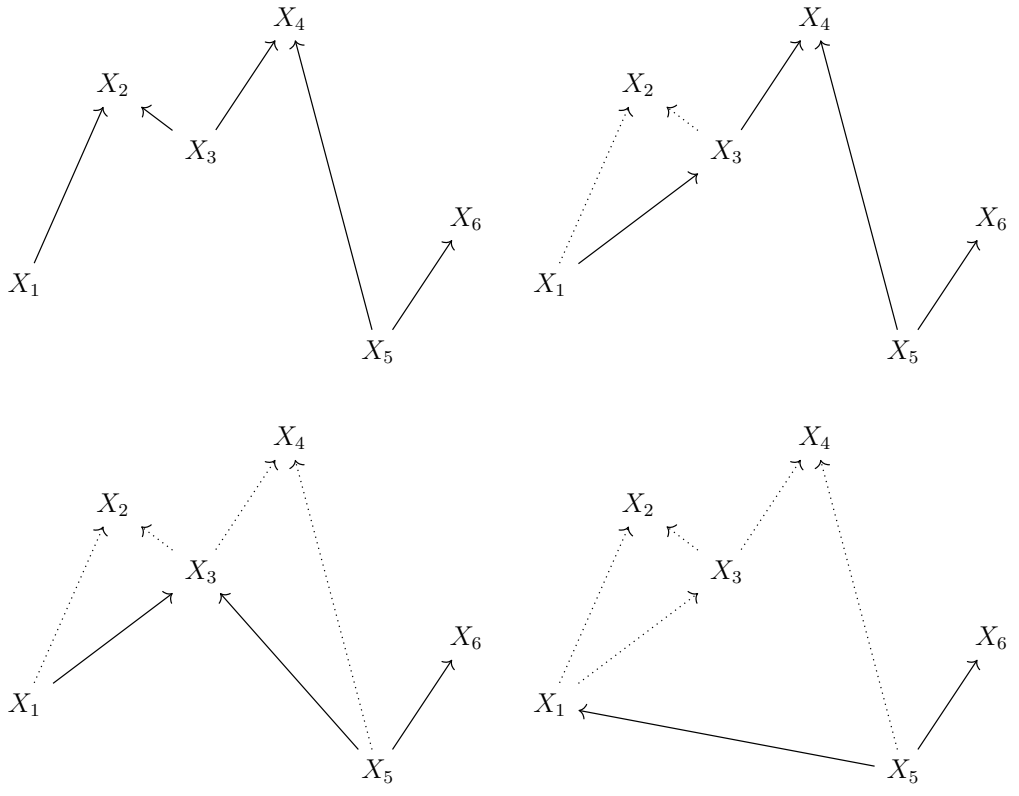


Figure 7. Algorithm 3 applied to the graph of Example 5.11, reading from left to right and top to bottom.

There are two cases, corresponding to the two **if** statements in Algorithm 3. In the first case, $f_{d'} = S$, and in the second, $f_{d'} = -S$.

If $f_{d'} \in G$, then by linearity of $\mathscr{N}_0$ and (b) there is no other polynomial in G with the same leading term as $f_{d'}$ and so the division algorithm returns zero. Likewise, the inner **while** loop of Algorithm 3 terminates because there is no live roof involving d'. To prove **IP** note that the only changes to $\mathscr{N}$ during the inner **while** loop are to mark some edges as dead. Marking the e'' as dead does not change the set of live roofs by (a).

If $f_{d'} \notin G$ then by (b) the leading term of $f_{d'}$ is the same as the leading term of some element of G (i.e.„ division continues) if and only if d' is part of a live roof in $\mathscr{N}$ (i.e.„ the inner **while** loop continues). If division and the **while** loop both stop, then $f_{d'} = \frac{1}{\mathrm{LC}\, S} S$ is added to G and d' is added to $\mathscr{N}$. As before one can check that **IP** is still true. Otherwise the division and **while** loop continue, and we repeat the above analysis. This completes the proof of the inductive step, showing that **IP** holds in the next iteration of the outer **while** loop and so (iii) follows by induction.

(i) Algorithm 2 terminates when B is empty, at which step B_{rel} must be empty and so there are no live roofs in $\mathscr{N}$ so Algorithm 3 also terminates. □

Figure 7 shows the Falling Roofs algorithm applied to the graph of Example 5.11.

Corollary 5.18. *If $\mathscr{S}$ is a linear graph and $\mathscr{N}$ is the output of Algorithm 3 on $\mathscr{S}$, then $G_{\mathscr{N}}$ is a Gröbner basis for the ideal generated by $G_{\mathscr{S}}$.*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Examining the algorithm, it is easy to verify the following:

Lemma 5.19. *Let $\mathscr{S}$ be a linear graph and $\mathscr{N}$ the output of Algorithm 3 on $\mathscr{S}$. If a vertex has valence one in $\mathscr{S}$ then it is incident with precisely one live edge in $\mathscr{N}$.*

5.3 Falling roofs and elimination

We consider the Falling Roofs algorithm (Algorithm 3) in the context of Elimination Theory. Let $\mathbf{Z} = \{Z_1, \ldots, Z_r\}$ be a set of variables and $<$ a total order on $\mathbf{Z}$. We denote the associated lexicographic monomial order also by $<$. We assume that given a partition $\mathbf{Z} = \mathbf{X} \cup \mathbf{Y}$ such that if $Z_i \in \mathbf{X}$ and $Z_j \in \mathbf{Y}$ then $Z_i < Z_j$. That is, the $\mathbf{X}$-variables are 'low' and the $\mathbf{Y}$-variables are 'high'. Let $\mathscr{S}$ be a linear $<$-graph with associated ideal

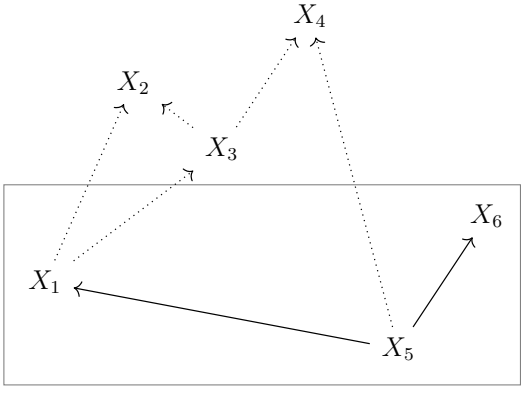
$$I = \langle G_{\mathscr{S}} \rangle \subseteq k[\mathbf{X}, \mathbf{Y}] .$$

Let $\mathscr{N}$ be the output of the Falling Roofs algorithm on $\mathscr{S}$. Then, by Corollary 5.18, $G_{\mathscr{N}}$ is a Gröbner basis for I and hence by the Elimination Theorem (Theorem 5.7)

$$G_{\mathscr{N}} \cap k[\mathbf{X}] \text{ is a Gröbner basis for } I \cap k[\mathbf{X}] .$$

If we write $X_1, \ldots, X_n$ for the variables in $\mathbf{X}$ and similarly $Y_1, \ldots, Y_m$ for the variables in $\mathbf{Y}$ then $G_{\mathscr{N}} \cap k[\mathbf{X}]$ is simply all the polynomials $f_e = X_j - X_i$ where $e : X_i \longrightarrow X_j$ is an edge in $\mathscr{N}$ between $\mathbf{X}$-variables, that is, we delete any edges involving $\mathbf{Y}$-variables. This has a natural interpretation as 'cutting off the top' of the realisation: If we draw a line through the realisation at the height of the highest $\mathbf{X}$-variable, intersecting with $k[\mathbf{X}]$ means deleting any vertices above the line and all edges incident with such vertices.

Example 5.20. *In Example 5.11, take $\mathbf{X} = \{X_1, X_5, X_6\}$ and $\mathbf{Y} = \{X_2, X_3, X_4\}$. Then, the set $G_{\mathscr{N}} \cap k[\mathbf{X}]$ is a Gröbner basis for $I \cap k[\mathbf{X}]$ where $G_{\mathscr{N}} \cap k[\mathbf{X}] = \{X_6 - X_5, X_1 - X_5\}$ is determined by the subgraph in the box below:*



Note that we follow the notational convention introduced just prior to Proposition 5.17, of denoting some edges of $\mathscr{N}$ with solid lines (the live edges) and some with dotted lines (the dead) edges. Shown is the state of these labels at the end of the algorithm. These are all edges of $\mathscr{N}$ and so for example $X_3 - X_1 \in G_{\mathscr{N}}$.

We assume $X_1, \ldots, X_n$ is the order these variables appear in the list $Z_1, \ldots, Z_r$, that is, the order the $\mathbf{X}$-variables are enumerated is their order of appearance from left to right in the realisation (which is *not* necessarily the $<$ order).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Proposition 5.21. *Suppose that $\mathscr{S}$ is a subgraph of the realisation $\mathscr{R}_<$ satisfying the following conditions*

(i) Every vertex in $\mathscr{S}$ of valence one is labelled with an element from $\mathbf{X}$.

(ii) We have $X_1 < X_2 < \cdots < X_n$.

Let $\mathscr{N}^+$ denote the subgraph of live edges in $\mathscr{N}$. Then, $\mathscr{N}^+$ is a linear graph, every vertex of nonzero valence in $\mathscr{N}^+$ is labelled with an element of $\mathbf{X}$, *and $G_{\mathscr{N}^+} = G_{\mathscr{N}} \cap k[\mathbf{X}]$ is a Gröbner basis for $I \cap k[\mathbf{X}]$.*

Proof. By Proposition 5.17(ii), the graph $\mathscr{N}^+$ is linear and by construction it contains no roofs. Suppose there were a vertex with nonzero valence in $\mathscr{N}^+$ labelled by $Y \in \mathbf{Y}$. By examination of the Falling Roofs algorithm, it is easy to see that Y must have had nonzero valence in $\mathscr{S}$. It belongs to a connected component of $\mathscr{S}$ of the form

$$Z_{i_1} \text{ —— } \cdots \text{ —— } Y \text{ —— } \cdots \text{ —— } Z_{i_l} \tag{57}$$

where $Z_{i_1}, Z_{i_l} \in \mathbf{X}$ by (i). By Lemma 5.19, the vertices Z_{i_1}, Z_{i_l} have valence one in $\mathscr{N}^+$ and again by examining the algorithm, we can see that Y is still connected to these vertices in $\mathscr{N}^+$ (one proves by induction that if two vertices connected by an unoriented path in $\mathscr{S}$ have nonzero valence in $\mathscr{N}^+$ then they are still connected by an unoriented path in $\mathscr{N}^+$). That is, we have also in $\mathscr{N}^+$ a connected component that looks like (57).

Now, we invoke a discrete version of Rolle's theorem: since (57) starts low (at Z_{i_1}) ends low (at Z_{i_l}) and is high somewhere in between (at Y) it must have a local maxima (i.e.„ a roof) as is easily proven by induction. But this contradicts the non-existence of live roofs in $\mathscr{N}$. Hence, all vertices with nonzero valence in $\mathscr{N}^+$ are labelled by $\mathbf{X}$.

For the last statement, since $G_{\mathscr{N}} \cap k[\mathbf{X}]$ is a Gröbner basis for $I \cap k[\mathbf{X}]$, it suffices for the last claim to prove that

$$G_{\mathscr{N}^+} = G_{\mathscr{N}} \cap k[\mathbf{X}]\,. \tag{58}$$

Suppose $f \in G_{\mathscr{N}} \cap k[\mathbf{X}]$ is $f = f_e$ for an edge $e : X_i \longrightarrow X_j$ which is not live. In the Falling Roofs algorithm, the only way e can be marked as dead is if at some point during the algorithm there are two live edges with target X_j. Since by hypothesis every $\mathbf{Y}$ variable is greater in the $<$-order than every $\mathbf{X}$ variable, the source of both of these edges with target X_j must be in $\mathbf{X}$. But this is impossible by (ii). This contradiction proves (58). □

6. Monomial Orders

Let π be a proof net with single conclusion A. We learned in Proposition 4.11 that the variables of P_π may be organised into sequences called persistent paths. It is natural to use these sequences to define a monomial order on P_π. Let

$$\mathscr{P}_1, \ldots, \mathscr{P}_m \tag{59}$$

be the persistent paths of π ordered so that if U_i is the last unoriented atom in $\mathscr{P}_i$, then $U_1, \ldots, U_m$ is the order that these atoms appear in A. Let us name the variables X_i so that (59) is $X_1, \ldots, X_n$ and $P_\pi = k[X_1, \ldots, X_n]$.

Definition 6.1. *We write $U <_0 V$ if U is before V in (59). The monomial order $<_0$ on P_π is the lexicographic order determined by $<_0$ on the variables.*

In this section, the monomial order on P_π is $<_0$.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Lemma 6.2. *If f, g are distinct elements of G_π then $\mathrm{LM}(f) \neq \mathrm{LM}(g)$, where leading monomials are defined with respect to $<_0$.*

Recall the relation $\sim$ on the set U_π of unoriented atoms of π, which generates the equivalence relation $\approx$ whose equivalence classes are persistent paths. We have $U \sim V$ if U, V appear next to each other on a persistent path.

Definition 6.3. *Let $\mathscr{S}_0$ be the oriented graph with vertex set U_π where two variables $U, V \in U_\pi$ are connected by an edge $e : U \longrightarrow V$ if $U \sim V$ and $U <_0 V$.*

In the terminology of Section 5.2, this is a linear $<_0$-graph. This is a subgraph of the realisation $\mathscr{R}_{<_0}$.

Definition 6.4. *We denote by $G_\pi^{(0)}$ the ordered set of polynomials $G_{\mathscr{S}_0}$ of Definition 5.14,*

$$G_\pi^{(0)} = \{V - U \mid e : U \longrightarrow V \text{ is an edge in } \mathscr{S}_0\}\,. \tag{60}$$

Up to signs, $G_\pi^{(0)}$ is just G_π.

Lemma 6.5. *If f, g are distinct elements of $G = G_\pi^{(0)}$, then $S(f,g) \longrightarrow_G 0$.*

Proof. Suppose $f = X - X', g = Y - Y'$. By Lemma 6.2, we may assume $X >_0 Y$. Then,

$$\begin{aligned} S(f,g) &= \frac{XY}{X}(X - X') - \frac{XY}{Y}(Y - Y') \\ &= Y(X - X') - X(Y - Y') \\ &= XY' - X'Y \end{aligned}$$

Hence, $\mathrm{LT}\,(S(f,g)) = XY'$. Rearranging

$$\begin{aligned} &= Y'(X - X') - X'(Y - Y') \\ &= Y'f - X'g \end{aligned}$$

and $\mathrm{LT}(Y'f) = XY'$, $\mathrm{LT}(X'g) = X'Y < XY'$ so $S(f,g) \longrightarrow_G 0$. □

Proposition 6.6. *$G_\pi^{(0)}$ is a minimal Gröbner basis for I_π with respect to $<_0$.*

Proof. To prove $G = G_\pi^{(0)}$ is a Gröbner basis, we need to (by (Cox et al. 2015, §2.9 Theorem 3)) prove that $S(f,g) \longrightarrow_G 0$ for distinct elements $f, g \in G$. But this is Lemma 6.5. To prove G is a minimal Gröbner basis, we can treat each persistent path $\mathscr{P} = X_1, \ldots, X_r$ separately. The part of G associated to $\mathscr{P}$ is $\{X_{i+1} - X_i\}_{i=1}^{r-1}$. With $f = X_{i+1} - X_i$, it is clear that $\mathrm{LC}\,(f) = 1$ and $\langle \mathrm{LT}\,(G \setminus \{f\})\rangle = \langle X_2, \ldots, X_i, X_{i+2}, \ldots, X_r\rangle$, which does not contain $\mathrm{LT}\,(f)$. □

Corollary 6.7. *G_π is a Gröbner basis for I_π.*

Remark 6.8. *While $G_\pi^{(0)}$ is minimal, it is not the reduced Gröbner basis of I_π as soon as there is a persistent path with more than two variables in π. The reduced Gröbner basis for I_π is the union, over all persistent paths $\mathscr{P}$ in the above notation, of $X_2 - X_1, \ldots, X_i - X_1, \ldots, X_r - X_1$.*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

6.1 Reduction and monomial orders

In this section, $\Gamma : \pi \longrightarrow \pi'$ is a reduction sequence between proof nets with single conclusion A. Let $\mathscr{P}_1, \ldots, \mathscr{P}_m$ denote the persistent paths of π ordered as in (59) and let us name the variables X_i so that $\mathscr{P}_1, \ldots, \mathscr{P}_m$ is $X_1, \ldots, X_n$ and $P_\pi = k[X_1, \ldots, X_n]$. That is, we name the variables according to their $<_0$-order.

Let $\mathscr{Q}_i$ be the subsequence of $\mathscr{P}_i$ consisting just of those unoriented atoms in π' (those in the image of T_Γ) so that $\mathscr{Q}_1, \ldots, \mathscr{Q}_m$ is a partition of the unoriented atoms of π'. Let $\mathscr{P}_i \setminus \mathscr{Q}_i$ denote the complement of the subsequence $\mathscr{Q}_i$. Then,

$$\mathscr{Q}_1, \ldots, \mathscr{Q}_m, \mathscr{P}_1 \setminus \mathscr{Q}_1, \ldots, \mathscr{P}_m \setminus \mathscr{Q}_m \tag{61}$$

is the set of unoriented atoms of π arranged in an order that depends on the reduction Γ. Note that $\mathscr{P}_i \setminus \mathscr{Q}_i$ are the variables in $\mathscr{P}_i$ eliminated during the reduction sequence.

Definition 6.9. *We write $U <_\Gamma V$ if U is before V in (61), reading from left to right. The monomial order $<_\Gamma$ on P_π is the lexicographic order determined by $<_\Gamma$ on the variables.*

From the point of view of the realisation $\mathscr{R}_{<_\Gamma}$ of Section 5.2, the order $<_\Gamma$ 'lifts' the variables in P_π that are to be eliminated above those that persist in $P_{\pi'}$. Observe that this realisation refers to the relation between $<_0$ and $<_\Gamma$ (Remark 5.12).

Definition 6.10. *Let $\mathscr{S}_\Gamma$ be the oriented graph with vertex set U_π, where two variables $U, V \in U_\pi$ are connected by an edge $e : U \longrightarrow V$ if $U \sim V$ and $U <_\Gamma V$.*

In the terminology of Section 5.2, this is a linear $<_\Gamma$-graph, which is a subgraph of the realisation $\mathscr{R}_{<_\Gamma}$. Note that the connected components of $\mathscr{S}_\Gamma$ correspond to persistent paths, since unoriented atoms in different persistent paths are not related by $\sim$.

Definition 6.11. *Let $G_\pi^{(\Gamma)}$ be the ordered set of polynomials $G_{\mathscr{S}_\Gamma}$ of Definition 5.14,*

$$G_\pi^{(\Gamma)} = \left\{ V - U \mid e : U \longrightarrow V \text{ is an edge in } \mathscr{S}_\Gamma \right\}. \tag{62}$$

Example 6.12. *Let γ be the reduction of the m-redex of the canonical detour π in Example 4.13. The sequence (61) is*

$$\mathscr{Q}, \underline{\mathscr{P} \setminus \mathscr{Q}} = X_1, X_2, X_3, X_6, X_7, X_{10}, X_{11}, X_{12}, \underline{X_4, X_5, X_8, X_9}$$

where we underline the atoms in $\mathscr{P} \setminus \mathscr{Q}$ for the reader's convenience. The graph $\mathscr{S}_\gamma$ is shown in Figure 8.

The sequence $G_\pi^{(\gamma)}$ is

$$\begin{aligned} &X_2 - X_1, X_3 - X_2, X_7 - X_6, X_{11} - X_{10}, X_{12} - X_{11}, \\ &\quad X_4 - X_3, X_5 - X_6, X_5 - X_4, X_8 - X_7, X_9 - X_{10}, X_9 - X_8 . \end{aligned}$$

Note that the persistent path $\mathscr{P}$ passes through the (cut) *link twice, which we can see clearly if we underline the atoms in $\mathscr{P} \setminus \mathscr{Q}$:*

$$X_1, X_2, X_3, \underline{X_4}, \underline{X_5}, X_6, X_7, \underline{X_8}, \underline{X_9}, X_{10}, X_{11}, X_{12} . \tag{63}$$

Each pass contributes two consecutive unoriented atoms in $\mathscr{P} \setminus \mathscr{Q}$. Note that the sequence (63), in which the atoms are ordered by appearance in the persistent path, is $<_\gamma$ increasing except at two points of 'defect' where the path exits an edge incident at a (cut) *link:*

$$X_1 <_\gamma X_2 <_\gamma X_3 <_\gamma \underline{X_4} <_\gamma \underline{X_5} >_\gamma X_6 <_\gamma X_7 <_\gamma \underline{X_8} <_\gamma \underline{X_9} >_\gamma X_{10} <_\gamma X_{11} <_\gamma X_{12} . \tag{64}$$

These are of course the roofs in $\mathscr{S}_\gamma$.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

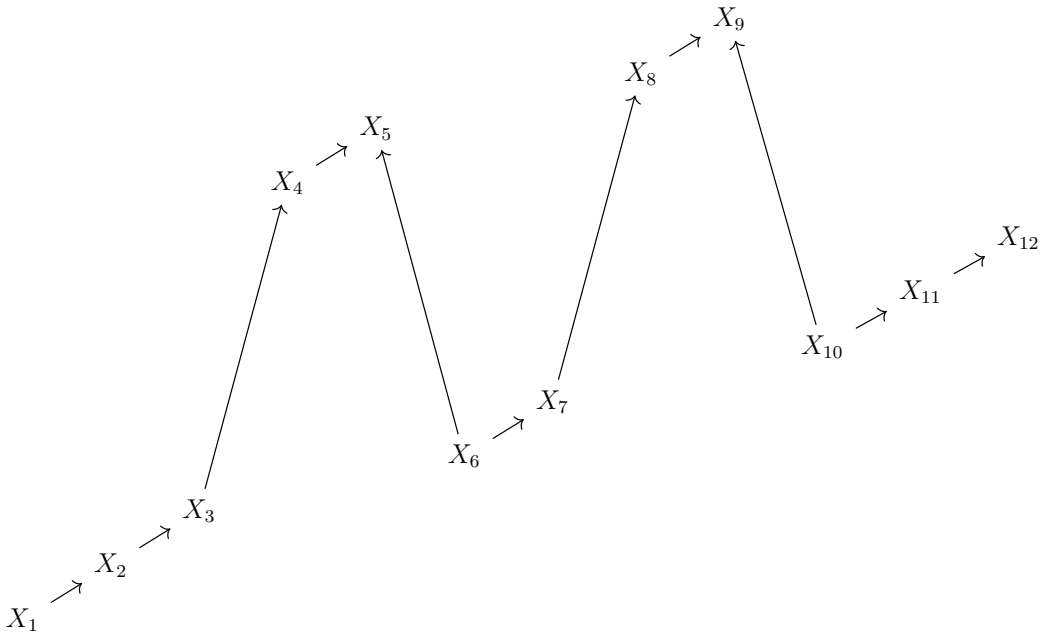


Figure 8. The graph $\mathscr{S}_\gamma$ for the canonical detour of Example 4.13.

Lemma 6.13. *$G_\pi^{(\Gamma)}$ is not a Gröbner basis for I_π with respect to $<_\Gamma$.*

From this, it is also clear that G_π is not a Gröbner basis with respect to $<_\Gamma$.

Remark 6.14. *The definition of the monomial orders $<_0, <_\Gamma$ involves ordering the persistent paths by the order of appearance of their last unoriented atom in A. This order could be chosen arbitrarily without changing our results; some choice must be made and this one appears relatively natural. The ordering on the persistent paths matters insofar as the Falling Roofs algorithm chooses roofs in earlier persistent paths before later ones.*

7. Main Theorems

Let π be a proof net with single conclusion A. We wrote informally in Section 2 about how the structure of π implicitly defines a relationship between unoriented atoms on the 'boundary' of the proof net, and how computation is the process of making this implicit relationship explicit. Let us now make these comments more precise.

Let $\mathbf{U}$ be the subsequence of positive unoriented atoms in A, and $\mathbf{V}$ the subsequence of negative unoriented atoms in A. The composite

$$k[\mathbf{U}, \mathbf{V}] \longrightarrow P_\pi \longrightarrow P_\pi / I_\pi \tag{65}$$

is by Proposition 4.11 a surjective morphism of k-algebras, and we denote by $I_{\partial\pi}$ the ideal which is its kernel. There is a commutative diagram

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

$$\begin{array}{ccc} k[\mathbf{U},\mathbf{V}] & \xrightarrow{\text{inc}} & P_\pi \\ \downarrow & & \downarrow \\ k[\mathbf{U},\mathbf{V}]/I_{\partial\pi} & \xrightarrow[\cong]{} & P_\pi/I_\pi \end{array} \tag{66}$$

By Proposition 4.11(iii)

$$I_{\partial\pi} = \langle \{U_i - V_{\sigma i}\}_i \rangle \tag{67}$$

for a permutation σ, which contains the structural information of the normal form of π (see Appendix A). The structure of π, as represented say by the generating set G_π for the ideal I_π, contains the same information but in an implicit form; the computational process of reduction of π to normal form may thus be identified with a process of discovering the generating set $\{U_i - V_{\sigma i}\}_i$ for $I_{\partial\pi}$ from the generating set G_π for I_π.

There are many such processes: indeed, we can now sketch a machine for generating them. Let $<$ be a total order on the set of unoriented atoms in π

$$U_\pi = \mathbf{U} \sqcup \mathbf{V} \sqcup \mathbf{W} \text{ with } \mathbf{V} < \mathbf{U} < \mathbf{W}$$

in the sense that for any i, j, k, we have $V_i < U_j < W_k$. Let $\mathscr{S} = \mathscr{S}_<$ denote the graph (see Section 6) with an edge $X \longrightarrow X'$ if $X \sim X'$ and $X < X'$. Then, $G_{\mathscr{S}}$ is, up to signs, the generating set G_π of I_π. As our computational process we take the Buchberger Algorithm with Early Stopping applied to $G_{\mathscr{S}}$, or what is the same, the Falling Roofs algorithm applied to $\mathscr{S}$. These algorithms return $G_{\mathscr{N}}$ and $\mathscr{N}$, respectively, (Proposition 5.17). Let $\mathscr{N}^+$ be the subgraph of live edges in $\mathscr{N}$.

Theorem 7.1. *$G_{\mathscr{N}^+} = \{U_i - V_{\sigma i}\}_i$ is a Gröbner basis for $I_{\partial\pi}$.*

Proof. The graph $\mathscr{S}$ satisfies the hypotheses of Proposition 5.21 with $\mathbf{X} = \mathbf{V} \cup \mathbf{U}$, so the subgraph $\mathscr{N}^+$ of live edges in $\mathscr{N}$ satisfies that $G_{\mathscr{N}^+} = G_{\mathscr{N}} \cap k[\mathbf{U},\mathbf{V}]$ is a Gröbner basis of $I_\pi \cap k[\mathbf{U},\mathbf{V}] = I_{\partial\pi}$. For each i, there is a persistent path beginning at $V_{\sigma i}$ and ending at U_i, and these endpoints maintain valence one in the graph of live edges throughout the Falling Roofs algorithm; since only vertices within the same persistent path may be connected by an edge in $\mathscr{N}$ this means that only edges in $\mathscr{N}^+$ are of the form $V_{\sigma i} \longrightarrow U_i$, as claimed. □

In this section, we answer two further questions:

- Is there a choice of total order on U_π under which the Gröbner basis algorithms match the familiar process of cut-elimination as represented by the reduction rules of Definition 4.3? Theorem 7.2 answers this in the affirmative.
- The normal form of π contains more unoriented atoms than just $\mathbf{U}, \mathbf{V}$, so is there a choice of total order on U_π for which the Gröbner basis algorithms produce the normal proof net $\widetilde{\pi}$ with its ideal $I_{\widetilde{\pi}}$?

7.1 Elimination versus cut-elimination

Let $\Gamma : \pi \longrightarrow \pi'$ be a reduction sequence between proof nets with single conclusion A. We view $P_{\pi'}$ as a subring of P_π via the algebra homomorphism T_Γ of Definition 4.15. Recall the following notation:

- $<_\Gamma$ is a monomial order; variables eliminated by Γ are 'large' (Definition 6.9).
- $G_\pi^{(\Gamma)}$ is a generating set for I_π (Definition 6.11). Up to signs, it is G_π.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

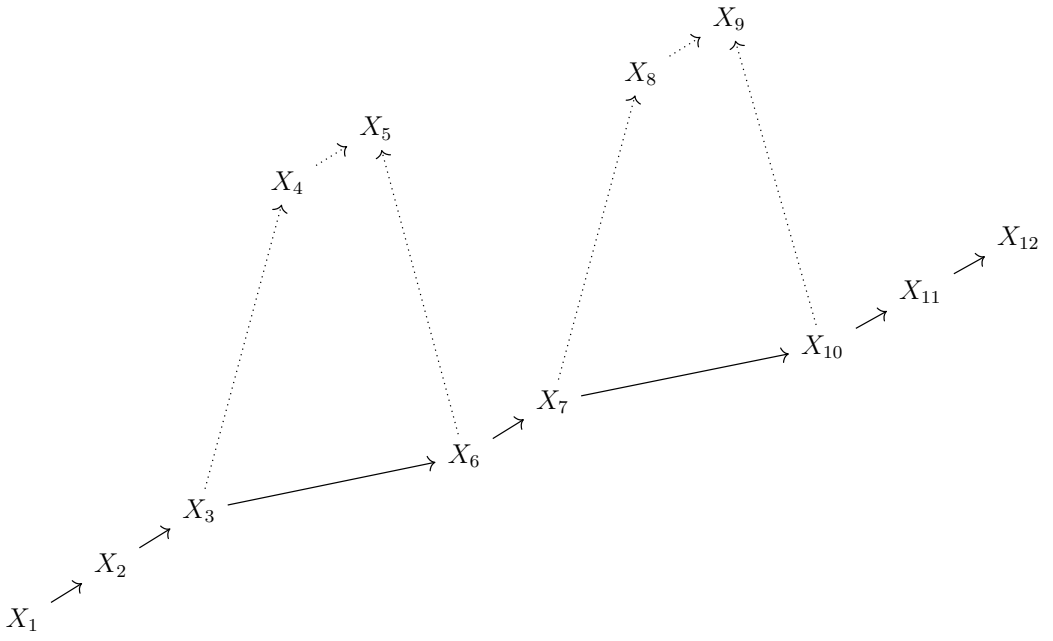


Figure 9. The Falling Roofs algorithm on the canonical detour.

- $\mathbb{B}_{es}(G_\pi^{(\Gamma)}, <_\Gamma)$ is a Gröbner basis for I_π obtained by applying the Buchberger with Early Stopping algorithm (Algorithm 2) to $G_\pi^{(\Gamma)}$.
- $G_{\pi'}^{(0)}$ is a (minimal) Gröbner basis for $I_{\pi'}$ (Definition 6.4).

Theorem 7.2. *There is an equality of sets*

$$G_{\pi'}^{(0)} = \mathbb{B}_{es}\big(G_\pi^{(\Gamma)}, <_\Gamma\big) \cap P_{\pi'}\,. \tag{68}$$

Proof. By definition $G_\pi^{(\Gamma)} = G_{\mathscr{S}}$, where $\mathscr{S}$ is the graph of Definition 6.10. If $\mathscr{N}$ is the output of the Falling Roofs algorithm applied to $\mathscr{S}$, then by Proposition 5.17

$$\mathbb{B}_{es}\big(G_\pi^{(\Gamma)}, <_\Gamma\big) = G_{\mathscr{N}}. \tag{69}$$

We are in the situation of Section 5.3, where the **X**-variables are those in $P_{\pi'}$. The monomial order $<_\Gamma$ satisfies the conditions of Proposition 5.21, from which we deduce that $G_{\mathscr{N}^+} = G_{\mathscr{N}} \cap P_{\pi'}$ is a Gröbner basis for $I_{\pi'}$. Since the **X** variables are ordered by $<_\Gamma$ in P_π as they are by $<_0$ in $P_{\pi'}$, it is clear that $\mathscr{N}^+$ is just the graph $\mathscr{S}_0$ of Definition 6.3 and hence $G_{\mathscr{N}^+} = G_{\pi'}^{(0)}$. □

The core idea of the theorem can be seen in the following example.

Example 7.3. *Let π denote the canonical detour as in Example 4.13. We continue Example 6.12 by presenting the Falling Roofs algorithm on the graph of Figure 8. The result of the algorithm is the graph $\mathscr{N}$ in Figure 9. Note that in the context of Section 5.3, if we partition the variables of our polynomial ring*

$$P_\pi = k[X_1, X_2, X_3, \underline{X_4}, \underline{X_5}, X_6, X_7, \underline{X_8}, \underline{X_9}, X_{10}, X_{11}, X_{12}]$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

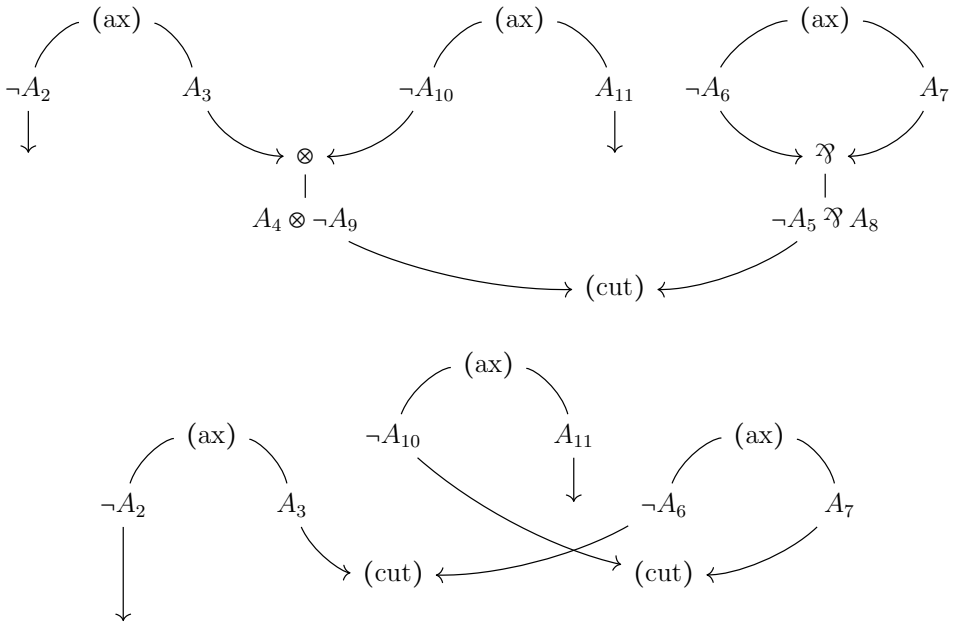


Figure 10. Cut reduction on the canonical detour.

into the set $\mathbf{X}$ *of those variables without underlines and the set* $\mathbf{Y}$ *of those variables with underlines, then 'cutting off the top' of the graph* $\mathscr{N}$ *means cutting it at a height with all the underlined variables above and all the other variables below. Recall from Proposition 5.17 that the Buchberger algorithm with Early Stopping on* $G_\pi^{(\gamma)}$ *computes* $\mathscr{N}$:

$$\mathbb{B}_{es}(G_\pi^{(\gamma)}, <_\gamma) = G_{\mathscr{N}}.$$

The Elimination Theorem (Theorem 5.7) says that

$$G_{\mathscr{N}} \cap k[\mathbf{X}] = \big\{X_2 - X_1, X_3 - X_2, X_6 - X_3, X_7 - X_6, \\ X_{10} - X_7, X_{11} - X_{10}, X_{12} - X_{11}\big\}$$

is a Gröbner basis for $I_\pi \cap k[\mathbf{X}]$. *Note that eliminating the cut in the canonical detour leads to the proof net* π', *the main part of which is shown in Figure 10. By construction,* $T_\gamma : P_{\pi'} \longrightarrow P_\pi$ *is the inclusion of*

$$P_{\pi'} = k[\mathbf{X}] = k[X_1, X_2, X_3, X_6, X_7, X_{10}, X_{11}, X_{12}]$$

into P_π *(note that* X_1, X_{12} *are not shown in the Figure for compactness).*

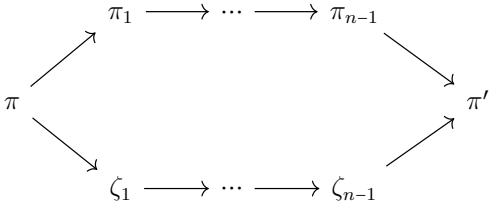
7.2 Normal Buchberger and normal forms

In Theorem 7.2, we made a connection between our modification to the Buchberger algorithm (Algorithm 2) and single step cut-reduction. We now examine the 'standard' form of Buchberger's Algorithm (Cox et al. 2015, §2.10 Theorem 9) and show that it relates to *normalisation*, that is, a maximal sequence of cut reductions yielding a normal form. In this section, π denotes a proof net with single conclusion A.

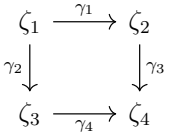
Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Lemma 7.4. *If* $\Gamma, \Omega : \pi \longrightarrow \pi'$ *are reduction sequences between proof nets with* π' *cut-free then* $T_\Gamma = T_\Omega$.

Proof. Let the reduction sequences be

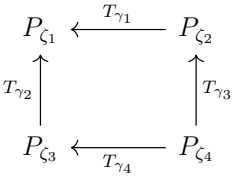


with Γ on top and Ω on the bottom. Recall that any two reduction paths have the same length (Troiani 2020, Corollary 3.0.7). Cut reduction is confluent (Troiani 2020, Proposition 3.0.3), so it suffices to show that any commutative diagram of reductions



(70)

induces a commutative diagram



(71)

For any reduction $\gamma : \pi \longrightarrow \pi'$, the schematics used to define T_γ indicate a mapping from the edges of π to the edges of π'. Commutativity of (71) amounts to confluence of this induced map. This follows from inspection of Figure 3. □

Let Γ denote any reduction path from π to its normal form $\widetilde{\pi}$ (the unique cut-free multiplicative proof net equivalent under cut-reduction to π). By the previous lemma, the k-algebra morphism

$$T = T_\Gamma : T_{\widetilde{\pi}} \longrightarrow T_\pi$$

is independent of the choice of Γ, and in this section, we simply denote it by T. There is a simple characterisation of the variables of P_π in the image of T. Recall that a proof net is an oriented multigraph, and a *directed path* in π means a sequence of edges traversed in the same direction as the orientation.

Lemma 7.5. *An unoriented atom* U *of* π *is in the image of* T *if and only if it is above a conclusion in the sense that the unique maximal directed path in* π *starting at the edge associated to* U *ends at a conclusion.*

Proof. Let ρ be the maximal directed path whose first edge is the one associated to U. First, we prove that if ρ ends at a conclusion, then every unoriented atom of every formula labelling an edge in ρ is in the image of T.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Write $\rho = (e_1, \ldots, e_n)$ and notice first that none of the edges $e_2, \ldots, e_n$ can appear in any redex. So, we only need to consider e_1, which is labelled by A. If e_1 is part of a redex, it is necessarily part of an a-redex. Inspection of (30) shows that reducing this a-redex results in a proof with the edge labelled A still present. This resulting edge now either may be part of an a-redex or not. Reducing any other a-redex not involving the edge labelled A does not remove this edge labelled A. Reducing any m-redex also does not remove this edge labelled A. It follows that the edge labelled A appears in the normal form π'. That is, U is in the image of T.

Now, we prove that if ρ ends at a cut link then none of the unoriented atoms of A are in the image of T. We proceed by induction on the length n of ρ. Say $n = 1$, if the edge labelled A is premise to a cut link, then it is part of an a-redex which when reduced, removes the edge labelled A. Thus, U cannot be in the image of T.

Now, assume that $n > 1$ and the result has been proved in the setting where ρ is of length $n - 1$. If A is above a cut link, then reducing this cut link results in a proof net where A is still above a cut link, the result follows directly from the inductive hypothesis. □

Note that an unoriented atom of π, which is not above a conclusion must be *above a cut* in the sense that the unique maximal directed path in π starting at the edge associated to the atom ends at a cut link. These are the variables eliminated by Γ, or what is the same, the variables *not* in the image of T.

Lemma 7.6. *Let $\mathscr{P} = (Z_1, \ldots, Z_r)$ be a persistent path in π which traverses a cut link. Then, there exists a (necessarily unique) pair of indices $1 < t_1, t_2 < r$ such that Z_i is above a cut if and only if $t_1 \leq i \leq t_2$.*

Proof. Let Z be a variable above a conclusion. By Lemma 7.5, every unoriented atom of every formula labelling an edge in the directed path starting at the edge of Z and ending at a conclusion is also above a conclusion. Thus, either $(Z_1, \ldots, Z)$ or $(Z, \ldots, Z_r)$ is a subsequence of $\mathscr{P}$ consisting only of variables above a conclusion. Thus, t_1 can be taken to be the maximal integer such that $(Z_1, \ldots, Z_{r_1-1})$ is a subsequence of $\mathscr{P}$ consisting only of variables above a conclusion, and t_2 can be taken to be the maximal integer such that $(Z_{t_2+1}, \ldots, Z_r)$ is a subsequence of $\mathscr{P}$ consisting only of variables above a conclusion. □

Let $\mathscr{P}_1, \ldots, \mathscr{P}_m$ denote the persistent paths of π, ordered in the way described at the beginning of Section 6.1. Let $\mathscr{Q}_i$ denote the subsequence of $\mathscr{P}_i$ consisting of those unoriented atoms in the normalisation $\widetilde{\pi}$, that is, those atoms in the image of T. In light of Lemma 7.6, we can divide the sequence $\mathscr{P}_i$ into three parts:

$$\mathscr{P}_i = \mathscr{Q}_i^1, \mathscr{P}_i \setminus \mathscr{Q}_i, \mathscr{Q}_i^2 \tag{72}$$

where $\mathscr{Q}_i^1$ is the subsequence of $\mathscr{P}_i$ consisting of those unoriented atoms in the persistent path coming *before* all the atoms above a cut, and $\mathscr{Q}_i^2$ is the subsequence of atoms coming *after* all those above a cut. By definition, $\mathscr{Q}_i = \mathscr{Q}_i^1 \sqcup \mathscr{Q}_i^2$.

Consider the following permutation of $\mathscr{P}_i$

$$\widetilde{\mathscr{P}}_i = (\mathscr{Q}_i^1)^{\mathrm{op}}, \mathscr{Q}_i^2, (\mathscr{P}_i \setminus \mathscr{Q}_i)^{\mathrm{op}} \tag{73}$$

where $\mathscr{R}^{\mathrm{op}}$ means the reversal of the sequence $\mathscr{R}$. Combining these re-ordered persistent paths we obtain a total order on the set of unoriented atoms of π:

$$\widetilde{\mathscr{P}}_1, \ldots, \widetilde{\mathscr{P}}_m\,. \tag{74}$$

Definition 7.7. *We write $U <_n V$ if U is before V in (74), reading from left to right. The monomial order $<_n$ on P_π is the lexicographic order determined by $<_n$ on the variables.*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

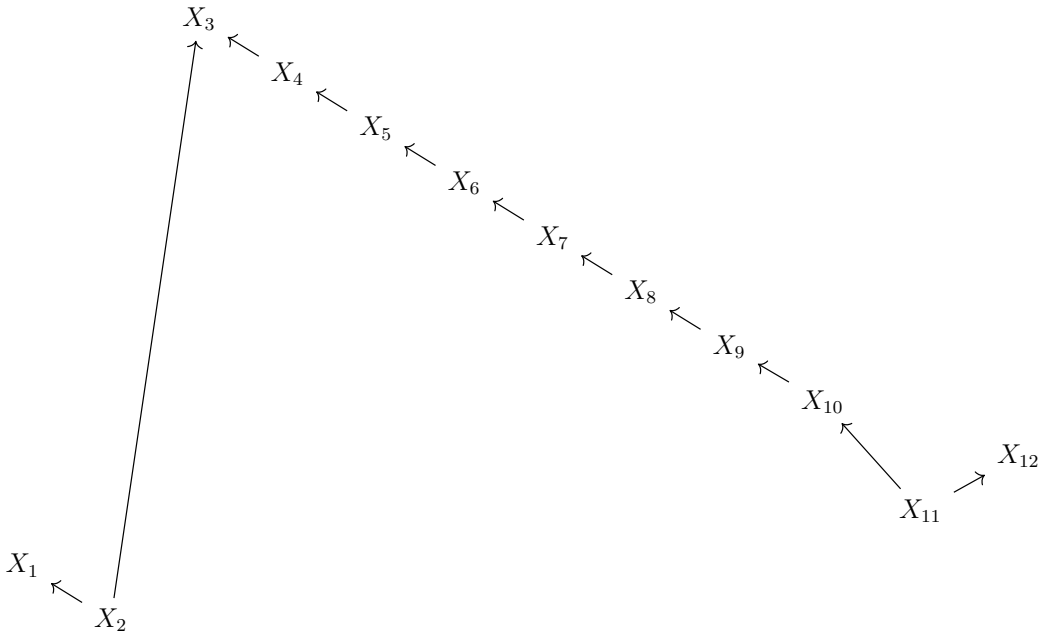


Figure 11. The graph $\mathscr{S}_n$ for the canonical detour of Example 7.10.

Definition 7.8. *Let $\mathscr{S}_n$ be the oriented graph with vertex set U_π, where two variables $U, V \in U_\pi$ are connected by an edge $e : U \longrightarrow V$ if $U \sim V$ and $U <_n V$.*

Definition 7.9. *Let $G_\pi^{(n)}$ be the ordered set of polynomials $G_{\mathscr{S}_n}$ of Definition 5.14,*

$$G_\pi^{(n)} = \left\{ V - U \mid e : U \longrightarrow V \text{ is an edge in } \mathscr{S}_n \right\}. \tag{75}$$

Example 7.10. *Let Γ be the normalisation of the canonical detour π in Example 4.13. Then, X_1, X_2, X_{11}, X_{12} are above a conclusion and the rest are above a cut, so $\mathscr{P}$ is*

$$X_1, X_2, \underline{X_3, X_4, X_5, X_6, X_7, X_8, X_9, X_{10}}, X_{11}, X_{12}$$

where we underline the atoms in $\mathscr{P} \setminus \mathscr{Q}$. Hence $\mathscr{Q}^1$ is X_1, X_2, $\mathscr{Q}^2$ is X_{11}, X_{12} and (73) is

$$(\mathscr{Q}^1)^{\mathrm{op}}, \mathscr{Q}^2, (\mathscr{P} \setminus \mathscr{Q})^{\mathrm{op}} = X_2, X_1, X_{11}, X_{12}, \underline{X_{10}, X_9, X_8, X_7, X_6, X_5, X_4, X_3}\,.$$

The graph $\mathscr{S}_n$ is shown in Figure 11

We identify $P_{\widetilde{\pi}}$ as a subring of P_π using T.

Theorem 7.11. *Let π be a proof net and $\widetilde{\pi}$ the normal form. Let $\mathbb{B}(G_\pi^{(n)}, <_n)$ denote the result of applying the Buchberger algorithm. Then there is an equality of sets*

$$G_{\widetilde{\pi}}^{(n)} = \mathbb{B}(G_\pi^{(n)}, <_n) \cap P_{\widetilde{\pi}}\,. \tag{76}$$

Proof. We write $G_\pi^{(n)} = \{f_1, \ldots, f_s\}$. Buchberger's Algorithm (Cox et al. 2015, §2.10 Theorem 9) begins by choosing a pair (i, j) with $i < j$ and considering the S-polynomial $S(f_i, f_j)$. The only interesting case is when f_i, f_j share a leading term, so we assume this below.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

The polynomials f_i, f_j must lie on a common persistent path $\mathscr{P}$, which traverses a cut. We have that $f_i = Z_{r_1+1} - Z_{r_1+2}$ and $f_j = Z_{r_1+1} - Z_{r_1}$ for some unoriented atoms in the persistent path $\mathscr{P} = Z_1, \ldots, Z_l$ and

$$\begin{aligned} S(f_i, f_j) &= (Z_{r_1+1} - Z_{r_1+2}) - (Z_{r_1+1} - Z_{r_1}) \\ &= Z_{r_1} - Z_{r_1+2}\,. \end{aligned}$$

We now divide by the sequence $(f_1, \ldots, f_s)$.

$$\begin{array}{l} (f_1, \ldots, f_s) \quad \overline{)Z_{r_1} - Z_{r_1+2}} \\ \qquad\qquad\quad \underline{Z_{r_1+3} - Z_{r_1+2}} \\ \qquad\qquad\qquad\quad Z_{r_1+3} - Z_{r_1} \\ \qquad\qquad\qquad\quad \underline{Z_{r_1+3} - Z_{r_1+4}} \\ \qquad\qquad\qquad\qquad\quad Z_{r_1+4} - Z_{r_1} \\ \qquad\qquad\qquad\qquad\quad \vdots \\ \qquad\qquad\qquad\qquad\quad \underline{Z_{r_2-1} - Z_{r_1}} \\ \qquad\qquad\qquad\qquad\qquad\quad Z_{r_2-1} - Z_{r_2} \\ \qquad\qquad\qquad\qquad\qquad\quad \underline{\qquad\qquad} \\ \qquad\qquad\qquad\qquad\qquad\qquad\quad Z_{r_2} - Z_{r_1} \end{array} \tag{77}$$

The division algorithm necessarily terminates here as there is no polynomial in $G_\pi^{(n)}$ with leading term either Z_{r_2} or Z_{r_1}. It remains to show that $Z_{r_2} - Z_{r_1}$ appears in $G_{\pi'}^{(n)}$.

The sequence $(Z_{r_1}, Z_{r_1+1}, \ldots, Z_{r_2-1}, Z_{r_2})$ is a subsequence of $\mathscr{P}$ and there is an associated persistent path $\mathscr{P}'$ in π' for which (Z_{r_1}, Z_{r_2}) is a subsequence. By the way that $<_n$ is defined, we have that $Z_{r_2} - Z_{r_1} \in G_{\pi'}^{(n)}$. This proves the containment $\supseteq$ in (76).

Next, we show the reverse inclusion. Let $f \in G_{\pi'}^{(n)}$ be given. If $f \in G_\pi^{(n)}$, then there is nothing to show, as $G_\pi^{(n)} \subseteq \mathbb{B}(G_\pi^{(n)}, <_n)$. So, suppose $f \notin G_\pi^{(n)}$. Then, $f = X' - X$ for some persistent path in π with subsequence $(X, W_1, \ldots, W_t, X')$ where $W_1, \ldots, W_t$ are variables not in $P_{\pi'}$. Let $g = W_1 - W_2$ and $g' = W_1 - X$. Then, $g, g' \in G_\pi^{(n)}$ and the first half of this proof shows that division of $S(g, g')$ by $G_\pi^{(n)}$ yields f, so $f \in \mathbb{B}(G_\pi^{(n)}, <_n)$. □

8. Conclusion

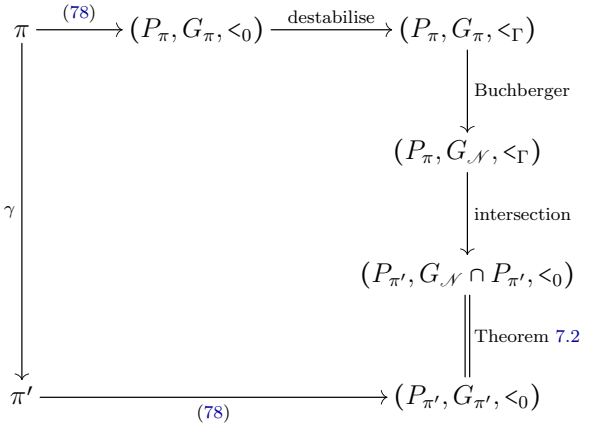
Let us summarise the dictionary between proofs and ideals developed in this paper. We have shown how to associate to a proof net π in multiplicative linear logic a triple

$$\pi \longmapsto (P_\pi, G_\pi, <_0) \tag{78}$$

consisting of a polynomial ring P_π, a generating set G_π for an ideal I_π in this ring, and a monomial order $<_0$ on P_π. We have shown that G_π is a Gröbner basis for I_π with respect to this monomial order (Proposition 6.6).

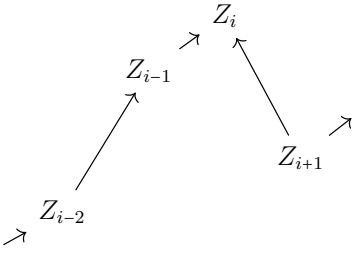
There are many other monomial orders on P_π. With respect to some monomial orders, G_π will be 'stable' in the sense that it is a Gröbner basis, while it will be 'unstable' (not a Gröbner basis) with respect to others and running the Buchberger algorithm makes nontrivial changes. The correspondence in (78) maps proof nets to stable triples. However, if π contains (cut) links, then we can destabilise it by using a monomial order $<_\Gamma$ derived from some reduction sequence. The 'restabilisation' of $(P_\pi, G_\pi, <_\Gamma)$ is the process of running the Buchberger algorithm and intersecting with the subring $P_{\pi'}$ of 'low' variables with respect to $<_\Gamma$. Our main Theorem says that this process ends with the image of the reduction π' under the map (78), or in other words the following diagram commutes:

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143



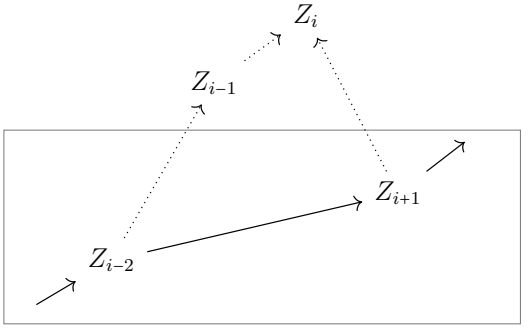
(79)

The language of stability is best understood in the context of the realisations: the combinatorial 1-manifolds together with a discrete Morse function ('height on the page') that we have associated to any pair $(G_\pi, <)$ where $<$ is a lexicographic monomial order. If the realisation contains no local maximum (that is, a roof) then by Corollary 5.18, G_π is a Gröbner basis with respect to $<$. Otherwise, the Buchberger algorithm (or equivalently Falling Roofs) proceeds to 'straighten out' the pair consisting of the 1-manifold with its Morse function, until the subgraph of live edges contains no such local maxima. That is, the Falling Roofs algorithm transforms a segment looking like



(80)

into a segment looking like



(81)

The algebro-geometric procedure in the right column of (79) and the topological procedure in (80), (81) represent our attempt to find mathematical shadows of the process of cut-elimination process for multiplicative proof nets.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

8.1 Towards geometry

Let π be a proof net with ideal $I_\pi \subseteq P_\pi$. The algebraic set

$$\mathbb{V}(I_\pi) = \{a \in \mathbb{A}^\pi \mid f(a) = 0 \text{ for all } f \in I_\pi\}$$

is an intersection of pairwise diagonals, which is not geometrically very interesting. More geometry can be introduced into the picture by considering pairs $\mathbb{V}(I_\pi) \subseteq V$ where V is an algebraic set. For example, given any $N \geq 2$, the polynomial $W = \sum_{f \in G_\pi} f^N$ is contained in I_π and so we have an inclusion

$$\mathbb{V}(I_\pi) \subseteq \mathbb{V}(W)\,.$$

The subscheme $\mathbb{V}(I_\pi)$ naturally determines an object in the bounded derived category of coherent sheaves of the singular hypersurface $\mathbb{V}(W)$, and in closely related categories like the category of matrix factorisations of W. In sequels to the present paper, we will explore this idea as a way of associating objects in triangulated categories to proofs; the role of the Buchberger algorithm in the present paper is played by constructions in homological algebra such as the perturbation lemma.

Note

1 These indices are not a form of de Brujin index because they only matter locally, as we will explain.

References

Abramsky, S. (1994). Retracing some paths in process algebra. In: *Proceedings of CONCUR 1994: Concurrency Theory*, Berlin, Heidelberg: Springer, 1–17.

Abramsky., S., Haghverdi., E. and Scott, P. J. (1999). Geometry of interaction and linear combinatory algebras. *Mathematical Structures in Computer Science* **9** (3) 481–520.

Abramsky., S. and Jagadeesan, R. (1994). New foundations for the geometry of interaction. *Information and Computation* **111** (1) 53–119.

Boole, G. (1854). *An Investigation of the Laws of Thought: On Which Are Founded the Mathematical Theories of Logic and Probabilities*, Walton and Maberly, London.

Cox, D., Little, J. and O'Shea, D. (2015). *Ideals, Varieties, and Algorithms*, Fourth Edition, Springer.

Girard, J.-Y. (1987). *Multiplicatives, Logic and Computer Science: New Trends and Applications*, Rosenberg & Sellier, 11–34.

Girard, J.-Y. (1989). Towards a geometry of interaction. In: Gray, J. W. and Scedrov, A. (eds.) *Categories in Computer Science and Logic*, Vol. 92 of Contemporary Mathematics, AMS, 69–108.

Girard, J.-Y. (2011). *The Blind Spot: Lectures on Logic*, European Mathematical Society.

Girard, J.-Y., Lafont, Y. and Taylor, P. (1989). *Proofs and Types*, Cambridge Tracts in Theoretical Computer Science 7, Cambridge University Press.

Haghverdi., E. and Scott, P. J. (2006). A categorical model for the geometry of interaction. *Theoretical Computer Science* **350** (2-3) 252–274.

Howard, W. A. (1980). The formulae-as-types notion of construction. In: *Seldin and Hindley To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, Academic press.

Laurent, O. (2013). An Introduction to Proof Nets, Available at http://perso.ens-lyon.fr/olivier.laurent/pn.pdf.

Murfet, D. (2017). Logic and Linear Algebra: An Introduction, preprint. Available at https://arxiv.org/abs/1407.2650v3

Murfet, D. and Troiani, W. (2020). Gentzen-Mints-Zucker Duality, preprint. Available at https://arxiv.org/abs/2008.10131

Troiani, W. (2020). Linear Logic, lecture notes. Available at https://williamtroiani.github.io/MathNotes/LinearLogic.pdf

A. Geometry of Interaction: Multiplicatives

In Troiani (2020), a permutation is associated to each proof net, as was originally done by Girard Girard (1987). In short, these permutations act on the disjoint union of all the unoriented atoms of all the conclusions to axiom links of π. The permutation of interest here is α_π (Definition [Troiani 2020, 4.1.3]) which permutes two unoriented atoms X, Y if and only if X, Y are respectively unoriented atoms of formulas $A, \neg A$ where $A, \neg A$ are the conclusions of some axiom link in π, and X, Y

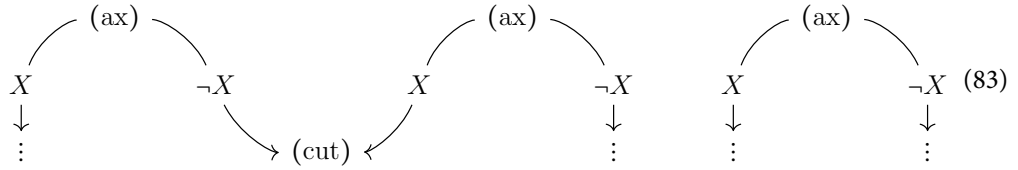
Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

appear at the same index of the *sequences* of unoriented atoms of A, $\neg A$. This permutation is equal to σ, which is introduced in Proposition 4.11.

Proposition A.1. *Let π be a proof net with single conclusion A with sequence of oriented atoms given by: $\big((U_1, u_1), \ldots, (U_n, u_n)\big)$ and let σ be as given in Proposition 4.11. Let δ_π be the permutation defined in (Troiani 2020, Definition 4.2.10). Then for all $i = 1, \ldots, n$ we have:*

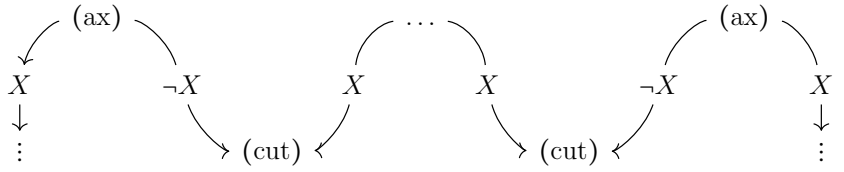
$$\delta_\pi(U_i) = U_{\sigma(i)} \tag{82}$$

Proof. First, we consider the special case when there are no m-redexes in π, and all conclusions of all axiom links are atomic. We will use some structure of such proof nets proved in Troiani (2020). First, we observe that all reductions are of the following form (see (Troiani 2020, Lemma 4.2.5)).



(83)

Hence, all the cut links of π appear inside 'chains' of axiom and cut links, as in the following Diagram.



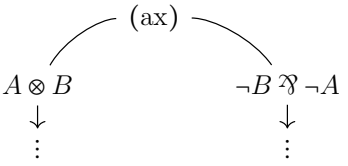
Also, by Proposition (Troiani 2020, 4.1.7), a cut-free proof net with sole conclusion B satisfying the property that every conclusion of every axiom link is atomic, is unique up to axiom links. Since each chain of axiom and cut links can be replaced by a single axiom link, it follows from this Proposition that π is uniquely determined up to these chains. The proof net π is thus a root labelled c with an incoming edge labelled A, the body of π is a finite binary tree uniquely determined by A, with some choice of chains connecting the leaves of the body of π. See Example (Troiani 2020, 4.1.9) for an example. Consider the subgraph of π given by removing the chains, we notice that the order of the atomic propositions at the leaves (read from left to right) is the same as the order of the sequence of oriented atoms of A. Hence, $U_{\sigma(i)}$ is the unoriented atom corresponding to an atomic proposition connected via a chain to the formula whose corresponding unoriented atom is U_i.

On the other hand, as described in the proof of Theorem (Troiani 2020, 4.2.11), the permutation δ_π is a product of transpositions, where two unoriented formulas X, Y are swapped if and only if they are at the extreme ends of a common chain. Hence, $\delta_\pi(U_i) = U_{\sigma(i)}$.

Now, we consider the case where π may admit m-redexes, we still assume that every conclusion of every axiom link is atomic though. Let π' denote the proof net obtained by reducing all m-redexes, by confluence of cut-reduction (Proposition [Troiani 2020, 3.0.3]), π' is independent of the order in which these m-redexes are reduced. We claim that $\delta_\pi = \delta_{\pi'}$ and also that the permutation σ obtained from π is equal to the permutation σ' obtained from π'. Both of these facts follow from Lemma (Troiani 2020, 4.2.4).

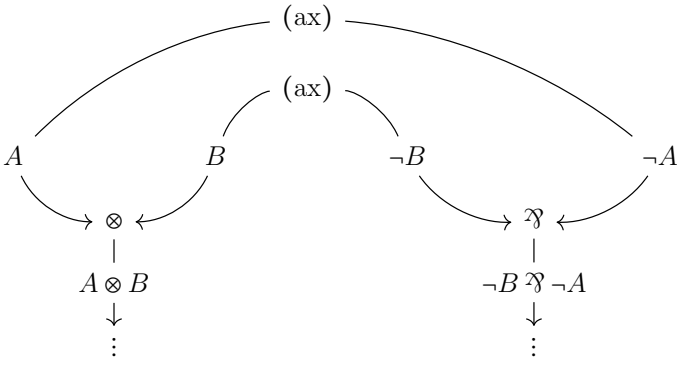
Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143

Finally, we consider the general case. In (Troiani 2020, Definition 3.0.9) we defined an *η-increx* to be a subgraph of a proof net of the following form.



(84)

We can replace this by the subgraph (85) below, and then repeat that process to eventually obtain a proof net π' such that every every conclusion of every axiom link of π' is atomic.



(85)

One must check that π' is independent of the choice of η-expansion procedure, and also that this process terminates. This is done just after Lemma 3.0.10 in Troiani (2020). It now remains to show that $\delta_\pi = \delta_{\pi'}$ and that the permutation σ obtained from π is equal to the permutation σ' obtained from π'. Both of these facts follow from inspection of (85). □

Cite this article: Murfet D and Troiani W (2026). Elimination and cut-elimination in multiplicative linear logic. *Mathematical Structures in Computer Science* **36**(e33), 1–42. https://doi.org/10.1017/S0960129525000143

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 16 Nov 2026 at 10:36:19, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0960129525000143