ELSEVIER

Annals of Pure and Applied Logic 83 (1997) 249–299

ANNALS OF PURE AND APPLIED LOGIC

Interpolants, cut elimination and flow graphs for the propositional calculus

A. Carbone*,[1]

Equipe de Logique, UFR de Mathématiques - Tour 45-55, Université Paris 7 - CNRS (URA 753), 2, Place Jussieu, 75251 Paris Cedex 05, France

Received 28 February 1995; revised 22 April 1996
Communicated by van Dalen

Dedicated to the memory of Spiro-Giorgio Mantzivis

Abstract

We analyse the structure of propositional proofs in the sequent calculus focusing on the well-known procedures of Interpolation and Cut Elimination. We are motivated in part by the desire to understand why a tautology might be 'hard to prove'. Given a proof we associate to it a logical graph tracing the flow of formulas in it (Buss, 1991). We show some general facts about logical graphs such as acyclicity of cut-free proofs and acyclicity of contraction-free proofs (possibly containing cuts), and we give a proof of a strengthened version of the Craig Interpolation Theorem based on flows of formulas. We show that tautologies having minimal interpolants of *non-linear* size (i.e. number of symbols) must have proofs with certain precise structural properties. We then show that given a proof Π and a cut-free form Π' associated to it (obtained by a particular cut elimination procedure), certain subgraphs of Π' which are logical graphs (i.e. graphs of proofs) correspond to subgraphs of Π which are logical graphs for the same sequent. This locality property of cut elimination leads to new results on the complexity of interpolants, which *cannot* follow from the known constructions proving the Craig Interpolation Theorem.

Keywords: Cut elimination; Interpolation; Graphs of proofs

0. Introduction

Fundamental questions in complexity theory are formulated in terms of the co-NP-complete set of propositional tautologies. In [5], Cook and Reckhow pointed out that there exists a propositional proof system in which all tautologies have proofs of polynomial size (i.e. number of symbols) if and only if NP = co-NP. A number of proof systems have been proposed and a hierarchy of them has been defined with respect

* E-mail: ale@logique.jussieu.fr.

[1] Partially supported by a grant from the EC (HCM) and by a grant from the Consiglio Nazionale delle Ricerche (CNR).

PII S0168-0072(96)00019-X

to their polynomial simulation (for a survey see [18]). To exhibit 'concrete' examples (i.e. families of tautologies) to prove exponential lower bounds for the proof systems of the hierarchy seems to be a difficult task. Certain combinatorial principles provide good candidates for this task but it is equally difficult to show whether or not such principles have polynomial size proofs.

We claim that this difficulty has its origin in the fact that *all* short proofs for such tautologies satisfy some non-trivial structural properties. With this idea in mind, we analyse the structure of propositional proofs in the *sequent calculus* (which is located high up in the hierarchy) and attempt to understand *why* a tautology might be 'hard to prove'. Our analysis will focus on the well-known procedures of Interpolation and Cut Elimination [7].

The Craig Interpolation Theorem says that given a sequent $A \to B$ there is a formula C, called an *interpolant*, that is made up of subformulas 'common' to A and B and such that $A \to C$ and $C \to B$ are provable sequents. We give another proof of this theorem and we strengthen its statement by describing the precise logical relations between occurrences of formulas in A, B and C. The proof is obtained as a corollary of the Cut Elimination theorem and as a consequence the size (i.e. the number of symbols) of the interpolant induced by the construction turns out to be in the worse case exponential in the size of $A \to B$. Under certain conditions (on the cut-formulas in the proof) though, we show how to avoid cut elimination by building the interpolant directly from the original proof and we obtain the interpolant's size to be at most quadratic.

It is an open question whether or not the size of the smallest interpolant can be polynomially bounded by the size of the tautology. A positive answer would imply an important consequence in complexity theory, namely that $\mathrm{NP} \cap \text{co-NP} \subseteq \mathrm{P/poly}$ [1]. Fixing a k, to find families of tautologies $A_n \to B_n$ of size $\mathcal{O}(n)$ with interpolants C_n of minimal size bounded from below by some polynomial in n of degree k, is difficult. In [16], Mundici proves a lower bound for $k = 2$. We will show that indeed to have a *non-linear* lower bound for the size of C_n the structure of all proofs for $A_n \to B_n$ should satisfy some non-trivial requirement on the logical relations between formula occurrences in $A_n \to B_n$ induced by the proof.

This analysis is made by studying the properties of the graph defined by tracing the flow of formula occurrences in a proof. The formulation of 'graph of a proof' (called *logical flow graph*) we will be using in the paper was introduced by Buss in [2] to prove that the k-provability problem (i.e. given a formula A and an integer k, to determine if A has a proof with k or fewer lines) for first order logic is undecidable. The idea of using the flow of occurrences to study the structure of proofs is fundamental in the work of Girard [8] where graphs called *proof nets* are associated to linear logic proofs. [1]

Our second task is the analysis of the transformation of logical flow graphs during the procedure of cut elimination (i.e. an effective way to transform any proof in the

[1] The notions of *logical flow graph* and *proof net* are strongly related. We should emphasize that the concept of *flow* which will be fundamental for most of the results presented here, has a natural formulation in terms of proof nets as well.

sequent calculus into a proof with simpler structural properties). When we look at the procedure of cut-elimination geometrically, in terms of its effect on the logical flow graph, we see that it enjoys a certain similarity with the application of 'local rules' for cellular automata [21] or 'division laws' for coloured simplicial complexes [9], where nodes ('cells') of a graph are transformed into more complicated objects ('complexes') by a finite number of rules whose applicability in general depends on the neighbourhood relations of the node in the graph. In our case we want to transform a logical flow graph into another logical flow graph. Each rule of transformation is applied to a certain subgraph and transforms it into another subgraph, at times doubling the number of nodes.

As for cellular automata or cell divisions it is interesting to study the relations between local and global behaviour of the transformation steps. We will show that given a proof Π (possibly containing cuts) and its cut-free form Π', for *any* subgraph of Π' that is a logical flow graph and that satisfies some global conditions, there is a subgraph of Π that is a logical graph with the same logical properties, namely the proofs associated to both subgraphs are proofs for the same sequent. The idea is illustrated by the following diagram:

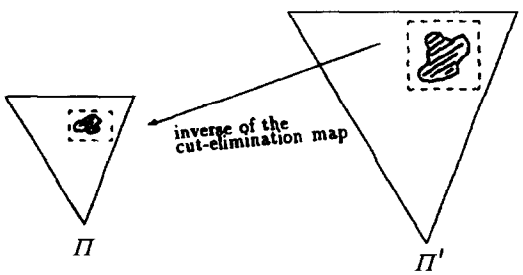

In order to have this inversion property one has to be careful about the method of cut elimination. It is not satisfied for whatever procedure of cut elimination we might apply (for instance it does not hold for the well-known cut elimination procedure introduced by Gentzen) because some of the information content of the original proof might be lost with the transformation. We shall give here a method of cut elimination for propositional proofs which does preserve this information. Our procedure does not work in predicate logic however.

An additional motivation for this property lies in the observation that the complexity of the subgraph in Π is in general much smaller than the complexity of the subgraph in Π'. In fact the complexity of any procedure of cut elimination induces, in the worse case, at least an *exponential* expansion of the number of lines of the original proof (this fact is a consequence of the doubling of the number of lines we mentioned before).

This inversion property suggests conditions under which given a proof (possibly with cuts) for $A \to B$ with no *logical* links (determined by the proof) between A and B, there is a proof either of $A \to$ or of $\to B$ with no greater complexity than the complexity of the original proof for $A \to B$. (While this problem may appear absurd in

terms of reasoning, it is less trivial combinatorially.) This result cannot be derived as a consequence of the usual construction used to prove the Craig Interpolation Theorem because, as we observed above, the latter is based on the Cut Elimination Theorem which increases the length of proofs very substantially.

The plan of the paper is as follows. In Section 1 we review the propositional formulation of Gentzen's Sequent Calculus; in Section 2 we present the definition of logical flow graph, we prove acyclicity of cut-free proofs and of contraction-free proofs (possibly containing cuts), and we justify the intuition of *inner* proof (i.e. a proof associated to some subgraph of a logical flow graph); in Section 3 we give a proof based on logical flow graphs of a strengthened version of the Craig Interpolation Theorem; the result is obtained as a consequence of the Cut Elimination Theorem but we show how, under certain conditions, an interpolant can be constructed directly over proofs containing cuts; in Section 4 the formal definition of inner proof is introduced together with properties of the structure of propositional proofs which are of general interest; the transformation of a logical flow induced by the well-known Gentzen's Cut Elimination procedure is analysed, a new procedure for the elimination of cuts is introduced and the Inversion Theorem on subgraphs is shown. In Section 5 we obtain some new bounds on the complexity of interpolants which are suggested by the Inversion Theorem.

Although the language of this paper is largely that of logic and complexity, in order to understand what is actually happening it is frequently better to think in pictures. Induction is convenient for making arguments precise, but it is also conducive to losing ideas in a wash of syntax. The main points are typically clearer if one imagines a formula moving towards an axiom or across a cut. One of our intentions here is to try to bring proofs closer to a geometric understanding.

Part of the results presented in this paper appear in [3]. The author thanks Melvin Fitting for suggesting the analysis of the Craig Interpolation Theorem with logical graphs, Rohit Parikh for helpful discussions, Sam Buss and Stephen Semmes for comments in an earlier version of the work, and Alexander Razborov for pointing out Example 4.22 which lead to the concept of *compact flow*, fundamental for the Inversion Theorem.

1. The sequent calculus

This section contains a brief review of the *sequent calculus* for propositional logic formulated by Gentzen. For a detailed exposition the reader can refer to [7, 19].

The sequent calculus is formulated in a language with logical symbols $\wedge, \vee, \neg$[2] and propositional variables $p, q, p_1, p_2, p_3, \ldots, q_1, q_2, q_3, \ldots$. *Formulas* are defined as usual; we will denote them with capital letters $A, B, C, \ldots, A_1, A_2, \ldots, B_1, \ldots$.

A *sequent* is a line of the form

$$A_1, \ldots, A_k \rightarrow B_1, \ldots, B_l$$

[2] The implication symbol $\supset$ is sometimes included in the set of symbols formulating to a propositional language. The results presented in the paper can be extended in a natural way to the extended language.

where the A_i's and B_j's are formulas; its intended meaning is $\wedge_i A_i \supset \vee_j B_j$. We permit k and l to be zero. A sequence of formulas separated by commas is a *cedent*; in the sequent above, $A_1,\ldots,A_k$ is the *antecedent* and $B_1,\ldots,B_l$ is the *succedent*. We will often refer to antecedents and succedents in a sequent using capital letters of the Greek alphabet. For instance $\Gamma \to \Delta$ denotes a sequent. In the following a sequence $A_1,\ldots,A_k$ will be a *multiset* of formulas, i.e. finite (possibly empty) set of formulas, in which repetitions of some formulas are admitted; the order of formulas in a multiset is not essential but for every member of the multiset the number of its occurrences is important. By the symbol Γ,Δ we denote the sum of the multisets Γ and Δ (i.e. the multiset containing all formulas in Γ and Δ so that if n_1 and n_2 are the number of occurrences of a formula A in Γ and Δ respectively, then $n_1 + n_2$ is the number of occurrences of A in the union Γ,Δ). The multiset A,Γ is obtained from Γ by adjoining the formula A. We will write Γ_1,Γ_2 as $\Gamma_{1,2}$, for short.

It should be pointed out that a proof in the sequent calculus is intended to be a *tree* of sequents; each sequent must either be an axiom (in this case the sequent is labeling a leaf of the tree) or be derived by one of the rules of inference we will give below (the sequent is a label for an internal node of the tree). Every occurrence of a sequent in a proof other than the end-sequent is used exactly once as a premise of an inference. Notice that a proof could be defined as sequence of sequents; but obviously any proof defined as such can be transformed into a tree-like proof by duplicating subproofs to derive intermediate results multiple times.

The axioms for our sequent calculus are of the form $A,\Gamma \to \Delta,A$ where A is atomic, and we will call them *logical axioms*. Formulas occurring in the cedents Γ,Δ of the above axioms are referred to as *weak* formulas of the sequent. Both occurrences of A are called *distinguished* occurrences.

The rules of inference are divided into two groups: the *logical rules* that introduce logical connectives to the left and to the right side of a sequent, and the *structural rules* (note that cut rule and contraction rule will be the only structural rules of the calculus).

The *logical rules* are the following:

$$\neg : \textit{left} \quad \frac{\Gamma \to \Delta,A}{\neg A,\Gamma \to \Delta} \qquad \neg : \textit{right} \quad \frac{A,\Gamma \to \Delta}{\Gamma \to \Delta,\neg A}$$

$$\wedge : \textit{right} \quad \frac{\Gamma_1 \to \Delta_1,A \quad \Gamma_2 \to \Delta_2,B}{\Gamma_{1,2} \to \Delta_{1,2},A \wedge B}$$

$$\wedge : \textit{left} \quad \frac{A,\Gamma \to \Delta}{A \wedge B,\Gamma \to \Delta} \qquad \frac{A,\Gamma \to \Delta}{B \wedge A,\Gamma \to \Delta}$$

$$\vee : \textit{left} \qquad \frac{A, \Gamma_1 \to \Delta_1 \quad B, \Gamma_2 \to \Delta_2}{A \vee B, \Gamma_{1,2} \to \Delta_{1,2}}$$

$$\vee : \textit{right} \qquad \frac{\Gamma \to \Delta, A}{\Gamma \to \Delta, A \vee B} \qquad \frac{\Gamma \to \Delta, A}{\Gamma \to \Delta, B \vee A}$$

The *structural rules* are

$$\textit{Cut} \qquad \frac{\Gamma_1 \to \Delta_1, A \quad A, \Gamma_2 \to \Delta_2}{\Gamma_{1,2} \to \Delta_{1,2}}$$

$$\textit{Contraction} \qquad \frac{\Gamma \to \Delta, A, A}{\Gamma \to \Delta, A} \qquad \frac{A, A, \Gamma \to \Delta}{A, \Gamma \to \Delta}$$

The formula occurrence, which is introduced explicitly into the lower sequent of a given rule of inference is called the *main* formula of the inference, and the formula(s), which are distinguished in the upper sequent(s) are the *auxiliary* formula(s). For example in the $\vee : \textit{left}$ rule the formula $A \vee B$ is the main formula and the formulas A and B are the auxiliary formulas. The other formulas of a sequent (i.e. the formulas in Γ, Δ) are called *side formulas.*

The cut rule has no main formulas; its auxiliary formulas are also called *cut-formulas.* An application of the cut rule will simply be called a *cut.* The auxiliary formulas of a contraction rule are also called *contraction formulas.*

We do not need the usual Weakening and Exchange rules in our calculus (as they appear in Gentzen's original formalization), because it uses multisets of formulas instead of sequences. The reader familiar with Gentzen's original formulation may like to notice that weak formulas appearing in logical axioms play in fact the role of formulas introduced by Weakening. This idea will be made explicit in Lemmas 4.8 and 4.9.

Our formalization will be referred to as *PLK*, ignoring the slight differences with Gentzen's classical formalization.

In the sequel, we will be concerned with different measures of complexity for proofs and formulas. The *number of lines* in a proof is defined as the number of nodes in its tree-like structure. The *size* of a formula A (denoted $|A|$) is its number of symbols. Namely, $|A|$ is 1 if A is atomic; $|A \wedge B| = |A \vee B| = |A| + |B| + 1$; $|\neg A| = |A| + 1$. The size of a sequent is the sum of the sizes of its formulas. The size of a proof is the sum of the sizes of its sequents.

2. The logical flow graph of a proof

In [2], Buss introduces the notion of *logical flow graph* to study how the influence of a formula spreads through a proof in LK_e (i.e. the first order sequent calculus *LK* together with axioms for equality). Following his introduction, we present the notion of logical flow graph formulated for our propositional sequent calculus *PLK*.

Let Π be a proof. An *s-formula* is an *occurrence* of a subformula of a formula in Π (here, '*s*-' stands for '*semi*-' or '*sub*-'). It has to be emphasized that an *s*-formula is an *occurrence* of a subformula in the proof as opposed to the subformula itself which may occur many times in the proof. Two distinct occurrences of the same formula A are called *variants* (we will use superscripts A^i, A^j to distinguish variants of A). The *logical flow graph* (formally defined below) is a directed graph whose nodes are *s*-formulas in Π; two *s*-formulas will be connected by an edge only if they are variants of each other; any two *s*-formulas connected by an edge will be in (distinct) sequents of some inference or will both be in an axiom on opposite sides of the sequent arrow.

We define the logical flow graph by specifying the edges.

First, in an axiom $A, \Gamma \to \Delta, A$ there is an edge directed from the left-hand A to the right-hand A.

Second, in any logical and structural inferences listed above, there is an edge directed from each side formula in the antecedent Γ in the *lower* sequent to the corresponding side formula of Γ in the *upper* sequent(s). There is an edge directed from each formula in the succedent Δ in the *upper* sequent to the corresponding formula of Δ in the *lower* sequent.

Third, in any logical inference or in a contraction rule if A (or B) is an auxiliary formula which appears in the succedent of an upper sequent of an inference then there is an edge directed from that A (or B) to the corresponding *s*-formula in the lower sequent. If A (or B) is an auxiliary formula which appears in the antecedent of an upper sequent of an inference then there is an edge directed towards that A (or B) from the corresponding *s*-formula in the lower sequent.

Fourth, in a cut inference there is an edge directed from the cut-formula A in the succedent of the left-hand upper sequent to the occurrence of A in the antecedent of the right-hand upper sequent.

Fifth, suppose there is a directed edge from an *s*-formula A^1 to A^2 and suppose B^1 is a subformula of A^1. Since A^1 and A^2 are variants there is a subformula B^2 of A^2 which corresponds to the subformula B^1 of A^1; the *s*-formulas B^1 and B^2 are, of course, variants. If B^1 occurs positively in A^1 then there is an edge from B^1 to B^2. If B^1 occurs negatively in A^1 then there is an edge from B^2 to B^1. Recall that B occurs positively (negatively) in A if B occurs an even (odd) number of times in the scope of a negation or in the left-hand operand of an implication. Clearly B^1 occurs positively in A^1 if and only if B^2 occurs positively in A^2. This concludes the definition of logical flow graph.

As an example consider the following proof:

$$\dfrac{\dfrac{\dfrac{\dfrac{A \to A \quad B \to B}{A, B \to A \wedge B} \quad C \to C}{A \vee C, B \to A \wedge B, C}}{(A \vee C) \wedge B \to A \wedge B, C}}{(A \vee C) \wedge B \to (A \wedge B) \vee C} \quad \vee : \textit{right and Contraction}$$

with logical flow graph restricted to the formula A (edges for $B, C, A \vee C, (A \vee C) \wedge B, A \wedge B$ and $(A \wedge B) \vee C$ are not indicated)

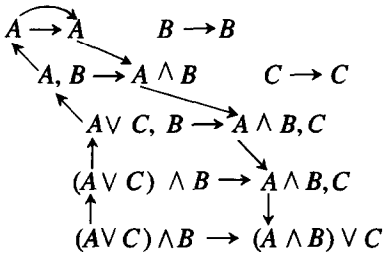

Notice that all pairs of variants A (respectively B and C) in Π are connected by a sequence of edges. In the following, we will call *connected variants* any pair of variants linked by *some* sequence of edges. We do not require the edges to respect the orientation.

A formula B occurs positively (negatively) in a sequent $\Gamma \to \Delta$ if B occurs negatively (positively) in a formula of Γ or positively (negatively) in a formula of Δ. If not otherwise indicated, in the following we intend that a positive or negative occurrence of a formula be defined relative to sequents. Notice that in a logical flow graph there are four kinds of edges: edges connecting positive occurrences, that are directed downwards; edges connecting negative occurrences, that are directed upwards; edges defined on axioms, that are directed from negative occurrences towards positive occurrences; edges defined on cut-formulas, that are directed from positive occurrences towards negative occurrences.

If a proof is cut-free, its logical flow graph will contain only three kinds of edges. Moreover, from a positive (negative) occurrence at most 1 edge goes out (comes in). If a *contraction* rule is applied to a positive (negative) occurrence, there are 2 edges directed downwards (upwards) to positive (negative) occurrences in the main formula of the rule application. From an easy checking of the rules of the calculus, it follows that a contraction rule is the only way in which two distinct edges can be directed to/from the same node of a logical flow graph. Hence, each node of a logical flow graph, whenever labeled by a positive (negative) occurrence can have at most 2 incoming (outgoing) edges and at most 1 outgoing (incoming) edge. From this, it follows that each node of a logical flow graph is at most of *degree* 3, where the degree of a node is the number of edges which depart or arrive at the node.

A connected subgraph of a logical flow graph in general is *not* a tree. Take for instance the following proof:

$$\dfrac{\dfrac{A \to A}{\to \neg A, A} \qquad \dfrac{\dfrac{A \to A}{A \to A} \quad \dfrac{\dfrac{A \to A}{\neg A, A \to}}{A, \neg A \vee A \to A} \quad \dfrac{A \to A}{\neg A, A \to}}{A, \neg A \vee A, \neg A \vee A \to A}\ \textit{Contraction}}{\dfrac{\to \neg A \vee A \qquad A, \neg A \vee A \to A}{A \to A}\ \textit{Cut}}$$

and notice the connected subgraph obtained by tracing some of the logical relations between the occurrences of the formula A in it

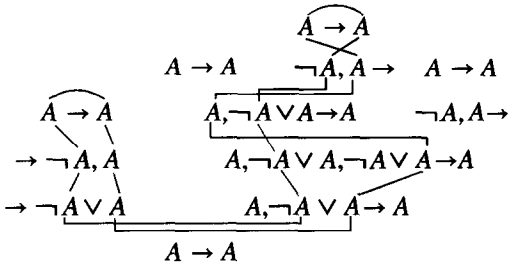

We call any sequence of consecutive edges in the logical flow graph $\mathscr{L}$ of $\Pi : S$ (i.e. the proof Π with end-sequent S) a *logical path* (two consecutive edges in this sequence should meet in a vertex which is a source for one and a sink for the other) or simply a *path*. We call any logical path starting and ending with two (distinct) s-formulas occurring in S a *bridge*. We call *direct path* a logical path passing through either positive or negative occurrences *only*. Notice that a direct path cannot cross an axiom or a cut, since this would force the path to pass through both positive and negative formula occurrences. Moreover, notice that the number of variants in a direct path is bounded by the height of Π.

Even though we will not use this fact in the following, it might interest the reader to see that cut-free proofs do not contain *cycles* (i.e. paths starting with an occurrence of a formula and going back to it). We call *acyclic* those logical graphs that do not contain cycles.

2.1. Proposition. *Let $\Pi : S$ be a cut-free proof. The logical flow graph of Π is acyclic.*

Proof. Let Π be a cut-free proof. Then there are only three kinds of edges connecting positive and negative occurrences of formulas in it, as remarked above. Notice that if all occurrences of formulas in a path are positive (negative) then the edges of the path should always go upwards (downwards), and therefore the path cannot be a cycle. Suppose there is a cycle in Π. By the latter observation, both negative and positive occurrences of formulas must occur in it. In particular we should have a way to connect positive occurrences to negative occurrences but there are no such edges in a logical flow graph of a cut-free proof. Therefore a cut-free proof cannot have cycles. □

It is also provable that *contraction-free* proofs do not contain cycles.

2.2. Proposition. *Let $\Pi : S$ be a contraction-free proof (possibly containing cuts). The logical flow graph of Π is acyclic.*

Proof. (*sketch*). Suppose Π to be a contraction free proof with logical flow graph containing a cycle. Let $\Pi_1, \ldots, \Pi_k$ be the intermediate proofs obtained by applying

the Gentzen procedure of cut elimination [7] to all pairs of cut-formulas that are non-weak. Let Π_1 be Π and Π_k be a proof possibly with cuts on weak formulas only. By inspection of the logical flow graphs induced by the transformation of the procedure of cut elimination it can be shown (here we use the assumption that the proof is contraction free) that if Π_i contains a cycle then Π_{i+1} must contain a cycle. Therefore Π_k contains a cycle. This gives a contradiction. In fact, it is enough to observe that a cycle cannot contain as a (sub)path any direct logical path passing through a cut-formula that is weak. This implies that a cycle cannot pass through any of the cut-formulas in Π_k. But as a corollary of Proposition 2.1 we know that a cycle must pass through a cut-formula. □

As a consequence of Propositions 2.1 and 2.2 it is clear that 'complicated' logical flow graphs arise from the interaction of both *cut* and *contraction* rules in the proof. Other facts of general interest are the following:

2.3. Proposition. *Let $\Pi : S$ be a proof and $\Pi' : S$ be a cut-free proof obtained by applying to Π the Gentzen procedure of cut elimination. If there is a logical path between two variants in the end-sequent S of Π' then there is a logical path between the same pair of variants of S in Π.*

Proof. (*sketch*). The procedure of cut elimination either duplicates or rearranges or eliminates subproofs of Π. It is a routine to check that for each one of these steps no new logical paths between pairs of variants in S are introduced, but simply eliminated or duplicated. □

The following is the formulation in terms of logical paths of the well-known *sub-formula property* (i.e. each formula in a cut-free proof $\Pi : S$ is a subformula of some formula in S). The property is formulated for arbitrary occurrences of formulas in Π.

2.4. Proposition. *Let $\Pi : S$ be a cut-free proof. For all positive (negative) occurrences of a s-formula D in Π, there is exactly one positive (negative) variant D' in S with a logical path from (to) D.*

Proof. By induction on the height of the proof Π. The result follows because all logical and structural rules have the property that the incoming and outgoing logical paths to the premises are naturally extended to the conclusion. For instance consider $\wedge$: *right* to be the last rule of inference in Π and the proof Π be of the form

$$\frac{\Gamma_1 \overset{\Pi_1}{\rightarrow} \Delta_1, A \quad \Gamma_2 \overset{\Pi_2}{\rightarrow} \Delta_2, B}{\Gamma_{1,2} \rightarrow \Delta_{1,2}, A \wedge B}$$

Suppose D is a positive (negative) occurrence in Π. If D is in Π_1, then by induction hypothesis, there is exactly one positive (negative) occurrence of a s-formula D'' in

$\Gamma_1 \to \Delta_1, A$ with logical path from D. The natural extension of such path activates exactly one formula D' in $\Gamma_{1,2} \to \Delta_{1,2}, A \wedge B$, variant of D''. If D is in $\Gamma_{1,2} \to \Delta_{1,2}, A \wedge B$, the claim is obviously satisfied for D' being D. □

Notice that a path links a pair of positive and negative variants if and only if it passes through an axiom of Π. As a consequence all weak occurrences in a cut-free proof are linked to only one variant in the end-sequent and such a variant must be of the same sign.

The *orientation* of the edges in a logical flow graph can be justified by the usual notion of soundness for axioms and logical rules. For instance, consider the edge connecting the occurrences of a formula A in the upper and lower sequents of a $\neg$: *left* rule

$$\begin{array}{c} \Gamma \to \Delta, A \\ \swarrow \\ \neg A, \Gamma \to \Delta \end{array}$$

and suppose that $A \to A'$ is a valid sequent. The rule obtained by substituting A with A' in the lower sequent (as the direction of the logical edge suggests), i.e.

$$\frac{\Gamma \to \Delta, A}{\neg A', \Gamma \to \Delta}$$

is a sound rule. On the other hand

$$\frac{\Gamma \to \Delta, A'}{\neg A, \Gamma \to \Delta}$$

is not sound.

In the following, we focus on the *logical relations* between *atomic* formulas; namely, we will work with the subgraph of a logical flow graph whose nodes are atomic formulas. Whenever we refer to s-formulas we will intend them to be atomic.

If one decides not to look at the whole logical flow graph for a proof Π but at subgraphs of it, one may observe that there are subgraphs tracing *proofs* involving only certain occurrences of formulas in Π. We will refer to them as *inner* proofs in Π. Let us illustrate the idea with an example. Consider the following proof:

$$\dfrac{\dfrac{A \to A \quad D \to D}{A \vee D \to D, A}}{A \vee D \to A, C \vee D}$$

and the portion of the logical flow graph not relying on A

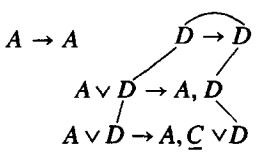

It is easy to see that it describes the inner proof

$$\frac{D \to D}{D \to C \vee D}$$

where we will refer to $D \to C \vee D$ as *inner* sequent of $A \vee D \to A, C \vee D$.

On the other hand, not all subgraphs of Π induce a proof. For instance, consider the subgraph

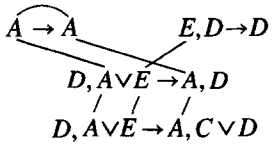

Notice that a given proof may contain more than one inner proof. Consider the following very simple proof:

$$\dfrac{A \to A \quad \dfrac{B \to B \quad C \to C}{B \vee C \to B, C}}{A, B \vee C \to B, A \wedge C}$$

and the two logical flows for it

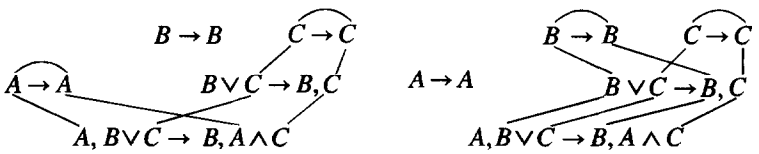

describing respectively the inner proofs

$$\frac{A \to A \quad C \to C}{A, C \to A \wedge C} \quad \text{and} \quad \frac{B \to B \quad C \to C}{B \vee C \to C, B}.$$

The idea of *inner* proof is a basic intuition for both the proof of the Craig Interpolation Theorem and the result on the Cut Elimination Theorem we want to present. Given a proof of $A \to B$ we will 'extract' from its cut-free form a proof of $A \to C$ ($C \to B$) by keeping the logical paths linked to A (B) and forgetting the ones only linked to B (A). Second, we show that if we can 'extract' (by keeping and forgetting paths properly) a proof from the cut-free form of Π of a certain inner sequent, then we can 'extract' a proof directly from Π of the same inner sequent.

3. Craig's interpolation theorem

We give a proof of the Craig's Interpolation Theorem based on logical flow graphs. The statement describes in a precise way the logical relations between a tautology $A \to B$ and its interpolant C.

3.1. Theorem (Craig's interpolation theorem). *Let $\Pi : A \to B$ be a proof. If A and B have at least one propositional variable in common, then there exists a formula C (called the* interpolant *of $A \to B$) containing only those propositional variables that occur in both A and B, such that*

(1) *there exist proofs $\Pi^A : A \to C$ and $\Pi^B : C \to B$;*

(2) *if p occurs in C then there is a p^A in A and a p^B in B such that Π^A (Π^B) contains a logical path between p and p^A (p^B) and Π contains a logical path between p^A and p^B;*

(3) *if Π^A (Π^B) contains a logical path between two occurrences p^1 and p^2 in A (B) then Π contains a logical path between p^1 and p^2.*

If A and B contain no common propositional variable then either $A \to$ or $\to B$ is provable.

The interpolant will be built by steps on the height of the subproofs using some easy considerations on the logical flows of formulas in the proof. By building the interpolant C, proofs of $A \to C$ and $C \to B$ are explicitly constructed from Π and essentially the idea of the construction is to 'forget' all those paths in Π that are linked only to B and A respectively. Each axiom and rule in Π will turn out to belong either to the proof of $A \to C$ or to the proof of $C \to B$ (Corollaries 3.3 and 3.6).

The same technique can be used to prove the intuitionistic version of the theorem and the theorem for the full predicate logic ([3]); in [4] the technique is used to deduce the Interpolation property for schematic systems (i.e. a generalization of the notion of proof system introduced in [17]).

3.1. Proof of the interpolation theorem

We need first to prove an easy lemma on the use of weak formulas in proofs. We say that a *weak occurrence* of a formula A in a proof Π is a formula whose direct paths all go to weak formulas A in axioms of Π. (The proof needs two simple lemmas on the elimination and introduction of weak occurrences to be established in Section 4.)

3.2. Lemma. *Let $\Pi : S$ be a proof of k lines. Then there is a proof $\Pi' : S$ of at most k lines and such that no weak formula in Π' is an auxiliary formula for some rule in Π'. Furthermore, if Π is cut-free then Π' is cut-free.*

Proof. This is proved by induction on the height of the proof Π. The only interesting cases arise when the last rule of Π is either a binary rule or a contraction

rule. In the first case suppose that a $\wedge$:*right* rule (similarly for $\vee$:*left* and *cut*) is applied to the subproofs $\Pi_1 : \Gamma_1 \to \Delta_1, A$ and $\Pi_2 : \Gamma_2 \to \Delta_2, B$ where A is a weak occurrence in Π_1. By eliminating A from Π_1 (Lemma 4.9) and adding to the resulting proof weak occurrences $\Gamma_2, \Delta_2, A \wedge B$ (Lemma 4.8) we obtain a proof with the same number of lines as Π_1, of $\Gamma_1, \Gamma_2 \to \Delta_1, \Delta_2, A \wedge B$. Let Π' be such proof.

In case the last rule of inference of Π is a contraction applied to $\Gamma \to \Delta, A, A$, and one of the A's is weak then applying Lemma 4.9 to it we obtain the desired proof of $\Gamma \to \Delta, A$. □

We are now ready to prove the theorem.

Proof of Theorem 3.1. Let Π' be the cut-free proof of $A \to B$ obtained as a result of a cut elimination procedure applied to a given proof of $A \to B$. By Lemma 3.2 there is a proof Π_* with no rules applied to weak auxiliary formulas. Without loss of generality we can assume both A, B to be non-weak and therefore having a propositional variable in common. (If $L(A) \cap L(B) = \emptyset$ then either A or B is weak in Π_*. This is not too difficult to check.) If B (A) is weak, the sequent $A \to (\to B)$ is provable and if P is a predicate in $L(A) \cap L(B)$ then $A \to P \wedge \neg P$ and $P \wedge \neg P \to B$ are provable sequents. Notice that Π_* does not contain weak occurrences, therefore all axioms are of the form $D \to D$.

For each subproof Π of Π_* with end-sequent $\Gamma \to \Delta$ we show that

(1) there exist subcollections Γ^A, Δ^A and Γ^B, Δ^B of Γ, Δ such that
   (a) each formula occurrence in Γ^A, Δ^A (Γ^B, Δ^B, respectively) has a direct path to A (B, respectively), and
   (b) Γ is Γ^A, Γ^B and Δ is Δ^A, Δ^B,
(2) either there exists a formula C_Π such that
   (a) $L(C_\Pi) \subset L(A) \cap L(B)$, and
   (b) $\Gamma^A \to \Delta^A, C_\Pi$ and $C_\Pi, \Gamma^B \to \Delta^B$ are provable sequents,
   or C_Π is not defined and, either $\Gamma^A \to \Delta^A$ or $\Gamma^B \to \Delta^B$ is a provable sequent.

Notice that if Π is Π_* then Γ^A is A, Δ^A is $\emptyset$, Γ^B is $\emptyset$, Δ^B is B and $A \to C_{\Pi_*}$, $C_{\Pi_*} \to B$ turn out to be provable sequents.

Let us start first by considering subproofs Π of Π_*, of height 1 of the form $D \to D$. If the l.h.s. (r.h.s.) distinguished occurrence of D has a direct logical path to A in Π_*, while the r.h.s. (l.h.s.) distinguished occurrence of D has a direct logical path to B, let C_Π be D ($\neg D$), the proof Π^A be $D \to D$ ($\Pi^A : \to \neg D, D$), the proof Π^B be $D \to D$ ($\Pi^B : \neg D, D \to$).

If both l.h.s. and r.h.s. occurrences of D have a direct logical path to A (B), let C_Π and Π^B (Π^A) be undefined, define Π^A (Π^B) to be $D \to D$.

For subproofs Π of height $k > 1$ we will examine the last rule of inference R. Suppose R is a $\vee$: *left* rule of the form

$$\frac{\overset{\Pi_1}{D, \Theta_1 \to \Lambda_1} \quad \overset{\Pi_2}{E, \Theta_2 \to \Lambda_2}}{D \vee E, \Theta_{1,2} \to \Lambda_{1,2}}$$

By construction there are (at most) two (defined) pairs of proofs Π_1^A, Π_2^A and Π_1^B, Π_2^B, and (at most) two (defined) formulas C_{Π_1}, C_{Π_2} such that $L(C_{\Pi_1}), L(C_{\Pi_2}) \subset L(A) \cap L(B)$. We will use this information to build the formula C_Π (whenever definable) and, Π^A : $\Gamma_\Pi^A \to \Delta_\Pi^A, C_\Pi$ and $\Pi^B : C_\Pi, \Gamma_\Pi^B \to \Delta_\Pi^B$. There are two cases:

(1) the main formula $D \vee E$ has a direct logical path to A. In case C_{Π_1} and C_{Π_2} are both defined then let C_Π be $C_{\Pi_1} \vee C_{\Pi_2}$, and define Π^A to be

$$\frac{\dfrac{\dfrac{\overset{\Pi_1^A}{D, \Theta_1^A \to \Lambda_1^A, C_{\Pi_1}}}{D, \Theta_1^A \to \Lambda_1^A, C_{\Pi_1} \vee C_{\Pi_2}} \quad \dfrac{\overset{\Pi_2^A}{E, \Theta_2^A \to \Lambda_2^A, C_{\Pi_2}}}{E, \Theta_2^A \to \Lambda_2^A, C_{\Pi_1} \vee C_{\Pi_2}}}{D \vee E, \Theta_{1,2}^A \to \Lambda_{1,2}^A, C_{\Pi_1} \vee C_{\Pi_2}, C_{\Pi_1} \vee C_{\Pi_2}}}{D \vee E, \Theta_{1,2}^A \to \Lambda_{1,2}^A, C_{\Pi_1} \vee C_{\Pi_2}}$$

and Π^B to be

$$\frac{\overset{\Pi_1^B}{C_{\Pi_1}, \Theta_1^B \to \Lambda_1^B} \quad \overset{\Pi_2^B}{C_{\Pi_2}, \Theta_2^B \to \Lambda_2^B}}{C_{\Pi_1} \vee C_{\Pi_2}, \Theta_{1,2}^B \to \Lambda_{1,2}^B}$$

where Θ_Π^A is $\Theta_{1,2}^A$, Λ_Π^A is $\Lambda_{1,2}^A$, Θ_Π^B is $\Theta_{1,2}^B$, Λ_Π^B is $\Lambda_{1,2}^B$. Clearly, $L(C_{\Pi_1} \vee C_{\Pi_2}) \subset L(A) \cap L(B)$. If C_{Π_1} (C_{Π_2}) is defined and C_{Π_2} (C_{Π_1}) is not, then let C_Π be C_{Π_1} (C_{Π_2}); if both formulas are undefined, also C_Π will remain undefined. The proofs Π^A and Π^B are built as expected.

(2) the main formula $D \vee E$ has a direct logical path to B. Let C_Π be $C_{\Pi_1} \wedge C_{\Pi_2}$ (whenever defined) and define Π^A to be

$$\frac{\overset{\Pi_1^A}{\Theta_1^A \to \Lambda_1^A, C_{\Pi_1}} \quad \overset{\Pi_2^A}{\Theta_2^A \to \Lambda_2^A, C_{\Pi_2}}}{\Theta_{1,2}^A \to \Lambda_{1,2}^A, C_{\Pi_1} \wedge C_{\Pi_2}}$$

and Π^B to be

$$\frac{\dfrac{\dfrac{\overset{\Pi_1^B}{D, C_{\Pi_1}, \Theta_1^B \to \Lambda_1^B}}{D, C_{\Pi_1} \wedge C_{\Pi_2}, \Theta_1^B \to \Lambda_1^B} \quad \dfrac{\overset{\Pi_2^B}{E, C_{\Pi_2}, \Theta_2^B \to \Lambda_2^B}}{E, C_{\Pi_1} \wedge C_{\Pi_2}, \Theta_2^B \to \Lambda_2^B}}{D \vee E, C_{\Pi_1} \wedge C_{\Pi_2}, C_{\Pi_1} \wedge C_{\Pi_2}, \Theta_{1,2}^B \to \Lambda_{1,2}^B}}{D \vee E, C_{\Pi_1} \wedge C_{\Pi_2}, \Theta_{1,2}^B \to \Lambda_{1,2}^B}$$

The case where C_{Π_1} and C_{Π_2} are not both defined, is handled as expected.

If R is a $\wedge$: *right* rule the way to build the formula C_Π and the proofs associated to it is symmetric to the case we just discussed. If R is a unary rule the formula C_Π is C_{Π_1} (whenever defined) and Π^A, Π^B are built with the same ideas already used in the above case. This concludes the construction.

Call C, Π^A, Π^B respectively, the formula $C_{\Pi *}$ and the derivations of $A \to C_{\Pi *}$ and $C_{\Pi *} \to B$. Conditions (2), (3) of the statement are direct consequences of the construction above and Proposition 2.3. This concludes the proof. □

The following statements are direct consequences of the proof of Theorem 3.1. Let $\Pi * : A \to B$, C, Π^A, Π^B be defined as in Theorem 3.1.

3.3. Corollary (Inclusion-exclusion formula). *The number of axioms in $\Pi *$ is exactly the sum of the number of axioms in Π^A and the number of axioms in Π^B, minus the number of axioms where bridges connecting A and B pass through.*

Proof. It is enough to observe that the number of axioms in Π^A (Π^B) is the number of bridges from A to C (C to B) plus the number of bridges from A to A (B to B), and that the number of bridges from A to C (C to B) is indeed the number of bridges from A to B in $\Pi *$. □

3.4. Corollary. *If $\Pi *$ has k lines then Π^A and Π^B have at most $3k$ lines.*

3.5. Corollary. *There are no contraction rules in Π^A, Π^B where the direct paths from C pass through.*

This last corollary follows from the fact that along the construction we never applied any contraction to the subformulas of the interpolant. Furthermore, the size of the interpolant is controlled by the number of binary rules in the proof. This is part of our next corollary.

3.6. Corollary. *The size of the interpolant C is exactly the number of axioms in $\Pi *$ where bridges connecting A to B pass through.*

Proof. This is a consequence of Theorem 3.1 (condition (2)) and Corollary 3.5. □

This last corollary might suggest that the complexity of an interpolant for a proof Π (possibly containing cuts) does *not* depend on the number of bridges which start and end in either A or B. In Section 5.2 we will prove this intuition to be false by showing that a tautology of size n has a *non-linear* lower bound for its minimal interpolant (i.e. the interpolant has size at least n^2) only when all its proofs do contain bridges from A back to A and from B back to B.

3.7. Remark. Because of Lemma 3.2 we do not need to consider an extension of the language L to include dummy symbols $\top, \bot$ as in the construction of the interpolant

proposed by Maehara (see [19, 7]). The use of flows in our construction makes explicit the relations between the interpolant and the formulas A, B in the end-sequent $A \to B$ (this fact is expressed in conditions (2), (3) of the theorem) and justifies in a natural way the sets of formulas used in Maehara's construction.

3.2. *Interpolation and proofs containing cuts*

There are cases where the interpolant for $\Pi : A \to B$ can be built directly from a proof *containing* cuts.

3.8. Definition. Let $\Pi : A \to B$ be a proof. A cut-formula in Π is $A(B)$-*monochromatic* if none of the paths passing through it connect to B (A).

3.9. Definition. A proof $\Pi : A \to B$ is *monochromatic* if each of its cut-formulas is either A-monochromatic or B-monochromatic, and if any path that does not start or end in the end-sequent either passes through A-monochromatic cut-formulas only or B-monochromatic cut-formulas only.

Any monochromatic proof might contain paths starting in A (B) and ending in B (A). These paths do not pass through cut-formulas.

3.10. Theorem.[3] *Let $\Pi : A \to B$ be a monochromatic proof (possibly with cuts) with k lines. Suppose A and B have at least one propositional variable in common. There is an interpolant C for $A \to B$ such that $|C| \leqslant \mathcal{O}(k)$ and there are two proofs $\Pi^A : A \to C$ and $\Pi^B : C \to B$ such that $\#lines(\Pi^A), \#lines(\Pi^B) \leqslant \mathcal{O}(k)$ and their cut-formulas have size $\mathcal{O}(k)$.*

To show the bound on the interpolant C we need to prove two general facts for propositional proofs. The proof of the first fact needs a lemma on the introduction of weak occurrences to be established in Section 4.2.

3.11. Lemma. *Let $\Pi : A \to B$ be a proof (possibly with cuts) of k lines. Then there exists a proof $\Pi' : A \to B$ with at most k lines such that each positive (negative) s-formula in a cut-formula of Π' has a direct path from (to) a distinguished occurrence in some axiom.*

Proof. Apply Lemma 3.2 to Π in order to obtain a proof Π^* where weak occurrences are not used as auxiliary formulas in rules. If Π^* satisfies the condition of the statement we are done. Otherwise consider a subformula R of some cut-formula C in Π^* which

[3] A similar result holds for LK-proofs as well. For the full predicate calculus the statement says that given a proof $\Pi : A \to B$ with k lines and monochromatic cuts only, there is an interpolant C for $A \to B$ such that $|C| \leqslant 2^{|A|+|B|+\mathcal{O}(k)}$ and $\#lines(\Pi^A), \#lines(\Pi^B) \leqslant \mathcal{O}(k)$. The proof we presented holds for the full predicate calculus, but instead of Lemma 3.12 one should use Lemma 3.1 in [11].

admits no direct paths to axioms of Π^* from its s-formulas and which is maximal with respect to this property. The formula R must be a *proper* subformula of C because C cannot be weak; moreover R has to arise from $\wedge$: *left* rules or $\vee$: *right* rules. More precisely, either there is a formula U such that $U \wedge R$ is a subformula of C or there is a formula V such that $V \vee R$ is a subformula of C. (Negations could always be absorbed into R, contradicting maximality.) In the first case there is at least one direct path from R which ends up in the main formula $U \wedge R$ of a $\wedge$: *left* rule, with U as the auxiliary formula, and in the second case there is at least one direct path from R which ends up in the main formula $V \vee R$ of a $\vee$: *right* rule, with V as the auxiliary formula. All other paths (if any) should end up in formulas of the form $U' \wedge U''$ introduced by a $\wedge$: *left* applied to an auxiliary formula U'', or $U' \vee U''$ introduced by a $\vee$: *right* applied to an auxiliary formula U'', where R is a proper subformula of U'. (The choice of connective $\wedge$ and $\vee$ here is independent of the preceding choice.) Binary rules could not be used to insert the connective ($\wedge$ or $\vee$) attached to R, because R would then have to have a direct path to an axiom.

We would like to show that we can modify Π^* in such a way as to remove R from the cut formula C and its dual. Assume for the sake of argument that R arises in C as $U \wedge R$. We can replace all $U \wedge R$ occurrences lying in the direct paths from C by U occurrences. By doing this we should also remove from the proof the $\wedge$: *left* rules originally applied to U. We have to look now at what happens in the cut formula dual to C. There might be direct paths from $U \wedge R$ in the dual of C which end up inside the main formula of a $\wedge$: *left* or a $\vee$: *right* rule, that is, in the part of the main formula that is added by the rule, and not the auxiliary formula. In this case we can replace all $U \wedge R$ occurrences in these direct paths by U occurrences. What else can occur? Since we have ruled out weak occurrences the connective in $U \wedge R$ has to be introduced by the binary rule $\wedge$: *right* (because on the opposite side of the cut it was introduced by unary rules $\wedge$: *left*). In this case the $\wedge$: *right* rule will be applied to subproofs $\Pi^R : \Gamma_1 \to \Delta_1, R$ and $\Pi^U : \Gamma_2 \to \Delta_2, U$. The proof we look for will be defined by replacing in Π^* the $\wedge$: *right* rule applied to Π^R, Π^U with the proof $\Pi'^U : \Gamma_{1,2} \to \Delta_{1,2}, U$ obtained from Π^U by adding Γ_1, Δ_1 as weak occurrences (Lemma 4.8). We should then replace all $U \wedge R$ occurrences in the direct path going to the cut formula dual of C with U.

One treats $V \vee R$ similarly since it either ends up in the easy part of a unary rule or in a binary rule.

In all situations the size (i.e. number of symbols) of the resulting proof is strictly smaller than the size of Π^*. Call Π the new proof and apply the procedure again to some formula R (if any) occurring in a cut-formula and having no direct paths to axioms. Since the proof has a finite number of formula occurrences and the size of the proof is decreasing by applying the procedure, after a finite number of steps we will obtain the desired proof Π'. □

The following lemma can be seen as the propositional counterpart of Lemma 3.1 in [11] which is formulated for the system *LK* (it appeared also in earlier work of Parikh

[17] and in the paper of Farmer [6]). Notice that we claim a bound on the *size* of the formulas in a proof though and not simply on the logical depth. Our statement is indeed stronger.

3.12. Lemma. *Let $\Pi : S$ be a proof (possibly with cuts) of k lines. Then there exists a proof $\Pi' : S$ with at most k lines and in which formulas have size at most $|S|+\mathcal{O}(k)$. In particular cut-formulas in it have size $\mathcal{O}(k)$.*

Proof. Apply Lemma 3.11 to Π and call Π^* the resulting proof. Clearly all formulas in Π^* with direct path to S have size at most $|S|$. So we need to look only at cut-formulas in Π^*. Since for all s-formulas B (occurring in cut-formulas C in Π^*) there is a *direct* path from B to an axiom in Π^*, it follows that $|C| \leqslant \mathcal{O}(k)$ (since there are at most k axioms, all axioms are defined on atomics and distinct s-formulas in a cut-formula should be directly linked to distinct distinguished occurrences in axioms). Thus the proof Π^* is indeed the proof Π' we are looking for. □

Proof of Theorem 3.10. The proof is based on two observations. For the first we assume that we are given a proof with only monochromatic cuts. For any subproofs Π' of it, having a cut as last rule of inference applied to $\Pi_1 : \Theta_1 \to \Lambda_1, D$ and $\Pi_2 : D, \Theta_2 \to \Lambda_2$ where D is a monochromatic cut-formula (say, having all paths not going to (coming from) B), we can build the interpolant by defining Π'^A as

$$\dfrac{\dfrac{\overset{\Pi_1^A}{\Theta_1^A \to \Lambda_1^A, D^A, C_{\Pi_1}} \quad \overset{\Pi_2^A}{D^A, \Theta_2^A \to \Lambda_2^A, C_{\Pi_2}}}{\Theta_{1,2}^A \to \Lambda_{1,2}^A, C_{\Pi_1}, C_{\Pi_2}}\,Cut}{\Theta_{1,2}^A \to \Lambda_{1,2}^A, C_{\Pi_1} \vee C_{\Pi_2}}\,\vee : left \text{ and } Contraction$$

and Π'^B as

$$\dfrac{\overset{\Pi_1^B}{C_{\Pi_1}, \Theta_1^B \to \Lambda_1^B} \quad \overset{\Pi_2^B}{C_{\Pi_2}, \Theta_2^B \to \Lambda_2^B}}{C_{\Pi_1} \vee C_{\Pi_2}, \Theta_{1,2}^B \to \Lambda_{1,2}^B}$$

(All paths not linked to A but passing through an A-monochromatic cut formula, do pass only through A-monochromatic cut-formulas in Π because Π is monochromatic. Those paths will then belong to Π'^A.)

Second, notice that by Lemma 3.12, for each $\Pi : A \to B$ with k lines there is a proof $\Pi' : A \to B$ with k lines and cut-formulas in Π' of size at most $\mathcal{O}(k)$. Furthermore, if Π has monochromatic cut-formulas *only* then Π' has monochromatic cut-formulas *only*. This property is a direct consequence of the proof of Lemma 3.12.

To show that there is an interpolant C for $A \to B$ such that $|C| \leqslant \mathcal{O}(k)$ it is enough to consider $\Pi' : A \to B$ and use an induction on the construction of the interpolant (defined by combining the proof of Theorem 3.1 and the first observation). The bounds on the complexity of the proofs for $A \to C$ and $C \to B$ follow directly from the construction. □

In Theorem 3.10, cut-formulas in Π do not need to be in $L(A) \cap L(B)$ but be linked either to A or to B. In Theorem 3.13 we will consider proofs with cut-formulas in $L(A) \cap L(B)$, without assuming any hypothesis on logical links.

3.13. Theorem. *Let $\Pi : A \to B$ be a proof (possibly with cuts) of k lines. Suppose A and B have at least one propositional variable in common. Suppose also that all cut-formulas in Π are in the language $L(A) \cap L(B)$. Then there is an interpolant C for $A \to B$ such that $|C| \leqslant \mathcal{O}(k^2)$ and there are two proofs $\Pi' : A \to C$ and $\Pi'' : C \to B$ such that $\#lines(\Pi'), \#lines(\Pi') \leqslant \mathcal{O}(k^2)$.*

Proof. From Π we build a proof of $\bigwedge_{i=1}^{n}(\neg A_i \vee A_i), A \to B$ where the A_i's are all the cut-formulas in Π. To do that we substitute in the obvious way all cut rules on A_i by a pair of rules of $\neg : \mathit{left}$ and $\vee : \mathit{left}$. We will obtain a proof of

$$\neg A_1 \vee A_1, \ldots, \neg A_n \vee A_n, A \to B$$

By $n-1$ applications of $\wedge : \mathit{right}$ we finally have a cut-free proof of $\bigwedge_{i=1}^{n}(\neg A_i \vee A_i), A \to B$ with $\mathcal{O}(k)$ lines. By Lemma 3.12 the size of this sequent is $|A| + |B| + \mathcal{O}(k^2)$ (since $n \leqslant k$). We then apply Theorem 3.1 to Π' and build an interpolant I for this sequent and two proofs $\Pi'_1 : \bigwedge_{i=1}^{n}(\neg A_i \vee A_i), A \to I$ and $\Pi'_2 : I \to B$ of $\mathcal{O}(k)$ lines. Since $\bigwedge_{i=1}^{n}(\neg A_i \vee A_i)$ is in $L(A) \cap L(B)$, I will be in $L(A) \cap L(B)$ as well. The formula $I \vee \neg(\bigwedge_{i=1}^{n}(\neg A_i \vee A_i))$ is clearly an interpolant for $A \to B$ of size $\mathcal{O}(k^2)$. In particular Π' and Π'' are built from Π'_1 and Π'_2 in the obvious way (notice that the sequent $\neg \bigwedge_{i=1}^{n}(\neg A_i \vee A_i) \to$ is provable with $\mathcal{O}(k^2)$ lines and Π'_1, Π'_2 are in size $\mathcal{O}(k)$ by Corollary 3.4). □

4. Cut elimination and flows

The aim of this section is to show that given a proof $\Pi : S$ and a cut-free proof Π' obtained by eliminating cuts from Π, an inner proof for Π' (say $\Pi'_* : S_*$) of an inner sequent S_* of S will induce an inner proof of S_* (say $\Pi_* : S_*$) in Π preserving the logical relations between the occurrences of formulas in the end-sequent (i.e. if there is a sequence of edges between a pair of variants of S_* in Π' then there is a sequence of edges between the corresponding pair of variants in S for Π).

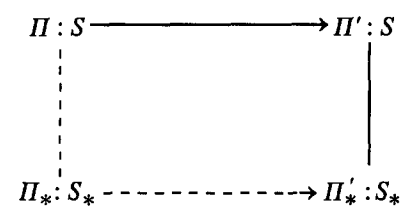

We will show that the picture above is *roughly* what one can prove; namely one should allow Π' to contain cuts on weak formulas (by eliminating these cuts one would indeed eliminate logical links between variants that we want to preserve). We call any such cut *trivial*.

We will show that the property is *not* true if Π' is obtained by the usual Gentzen's procedure of Cut Elimination. We present a counterexample to this fact together with an alternative procedure for the elimination of cuts that will satisfy the statement. Notice that this procedure works only for the *propositional* fragment of Gentzen sequent calculus. In purely combinatorial terms *a* cut elimination procedure *deforms* paths of a logical flow graph by adding new edges and cancelling old ones, and in the process it *disconnects* paths. It is this last feature that makes the analysis of the dynamics of the procedure particularly hard. The procedure we present allows one to control in some precise way the disconnecting of paths.

The intuition of *inner* proof is formally captured by the notion of *compact flow*. In this section we introduce the notion of flow and the fact that any flow (as subgraph of a logical flow graph) is a graph of a proof. Compact flows are a natural subclass of flows (see Section 4.1). Then we will show some general facts on the structure of propositional proofs. Finally we present the procedure of cut elimination and prove our main theorem, i.e. Theorem 4.27 (Section 4.3). A natural generalization of this result (Theorem 4.34) to compact subgraphs of *sub*proofs is the Inversion Property of cut elimination we described and illustrated in the introduction.

There are different possible choices that can be made to define a subgraph as a graph for a proof. They depend on the structure of the proofs we consider. Since the notion of *flow* is *new* we want to introduce it in its most straightforward form. Because of this we will prove the main result of this section (Theorem 4.27) for a class of proofs Π called *separated* (separated proofs satisfy a certain regularity condition on contraction formulas). For arbitrary proofs Π the idea for the proof of the theorem remains the same but the notion of flow should be refined (so that an analogue of Proposition 4.24 can be proved). In Proposition 4.18 we will show that for any proof of k lines always there is a *separated* proof of the same sequent with at most k lines (and the same number of symbols). Therefore from the perspective of complexity, to consider *separated* proofs is *not* restrictive (all tautologies are proved by some separated proof). Moreover, *separated* proofs better represent which logical relations between formula occurrences are *essential* for the provability of a sequent. But we will comment on this later.

To avoid annoying syntactical details let us replace unary rules for binary connectives in *PLK* with the following pair of rules:

$$\wedge : \textit{left} \quad \frac{A, B, \Gamma \to \Delta}{A \wedge B, \Gamma \to \Delta} \qquad \vee : \textit{right} \quad \frac{\Gamma \to \Delta, A, B}{\Gamma \to \Delta, A \vee B}$$

We will still call this calculus *PLK* since it is obviously equivalent. With this change Lemma 3.2 should be interpreted as saying that for all proofs Π of k lines there is a proof Π' of at most k lines such that no weak formula is used as auxiliary

formula for *binary* rules in Π', or for *unary* rules at one premise, or for contraction rules.

4.1. Flows of PLK

In the following we give a formal definition of *inner* proof by means of the notion of *flow* of formulas. As we have already observed in Section 2, we will focus on the logical relations between *atomic* formulas occurring in a proof, therefore whenever we refer to s-formulas we will intend them to be atomic.

We now define the notion of *full path.*

4.1. Definition. Let Π be a proof, $\mathscr{L}$ its logical flow graph, and f be a path in $\mathscr{L}$ between s-formulas. We call f a *full path* if there is a formula C_1 of f with no incoming edges (belonging to $\mathscr{L}$) and a formula C_2 of f with no outgoing edges (belonging to $\mathscr{L}$); f may have no edges in which case C_1 is C_2. Moreover, for each $i = 1,2$ either C_i occurs in the end-sequent of Π or C_i is an s-formula occurring in some non-distinguished formula of the axioms of Π (i.e. in the cedents Γ, Δ).

Notice that whenever C_i is a non-distinguished occurrence in some axiom, it does not give any logical contribution to the proof, in the sense that the validity of the sequent in which C_i appears, resides on the logical relations concerning formulas other than C_i.

Atomic formulas that are nodes of a path are called *active* formulas in the path. More generally we say that a formula (not necessarily atomic) is active in a set of paths whenever some (possibly all) of its s-formulas are active. Thus if A is not active in $\neg A$ then $\neg A$ is not active, and if A is active in $A \wedge B$ then $A \wedge B$ is active even if B is not active.

The active part of a formula occurrence in a proof is defined as image of the map introduced below.

4.2. Definition. Let $\Pi : S$ be a proof and $\mathscr{F}$ be a subgraph of $\mathscr{L}$ in Π. The 'forgetful' map $H_{\mathscr{F}}$ from occurrences of formulas in Π to formulas (possibly none) is defined as follows :

(1) if A is an atomic formula then $H_{\mathscr{F}}(A) = \begin{cases} A & \text{if } A \text{ is active in } \mathscr{F} \\ \emptyset & \text{otherwise} \end{cases}$

(2) $H_{\mathscr{F}}(\neg A) = \begin{cases} \neg H_{\mathscr{F}}(A) & \text{if } H_{\mathscr{F}}(A) \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases}$

(3) $H_{\mathscr{F}}(A \vee B) = \begin{cases} H_{\mathscr{F}}(A) \vee H_{\mathscr{F}}(B) & \text{if } H_{\mathscr{F}}(A), H_{\mathscr{F}}(B) \neq \emptyset \\ H_{\mathscr{F}}(A) & \text{if } H_{\mathscr{F}}(B) = \emptyset \\ H_{\mathscr{F}}(B) & \text{if } H_{\mathscr{F}}(A) = \emptyset \\ \emptyset & \text{if } H_{\mathscr{F}}(A), H_{\mathscr{F}}(B) = \emptyset \end{cases}$

(4) $H_{\mathscr{F}}(A \wedge B)$ is defined as in (3).

Let A be a formula. The formula $H_{\mathscr{F}}(A)$ when defined, is not necessarily a subformula of A. On the other hand, not all subformulas of A can be images of some forgetful map $H_{\mathscr{F}}$. Take for instance A to be $\neg(C \wedge (D \vee E))$; the subformula D cannot be the image of any forgetful map and the formula $\neg(C \wedge E)$ is not a subformula of A but it can be obtained from A by forgetting D. Because of these properties, we will call $H_{\mathscr{F}}(A)$ an *inner formula of* A.

Let Γ be a multiset of formulas $A_1 \cdots A_n$; with the symbol $H_{\mathscr{F}}(\Gamma)$ we denote the multiset $H_{\mathscr{F}}(A_1) \cdots H_{\mathscr{F}}(A_n)$. If Γ, Δ are two multisets of formulas then $H_{\mathscr{F}}(\Gamma \to \Delta)$ denotes the sequent $H_{\mathscr{F}}(\Gamma) \to H_{\mathscr{F}}(\Delta)$ where $H_{\mathscr{F}}(\Gamma), H_{\mathscr{F}}(\Delta)$ are multisets. We will call $H_{\mathscr{F}}(\Gamma \to \Delta)$ an *inner sequent of* $H_{\mathscr{F}}(\Gamma \to \Delta)$.

Notice that the map $H_{\mathscr{F}}$ is $\mathscr{F}$-dependent.

4.3. Definition (*Natural extension of a path*). Let $\Pi : S$ be a proof of height greater than 1 (i.e. Π is not an axiom) with last rule R. Let R be either a binary rule applied to $\Pi_1 : S_1$, $\Pi_2 : S_2$ or a unary rule applied to $\Pi_1 : S_1$.

(1) Let f be a path in Π_i passing through an s-formula C occurring positively (negatively) in S_i (where i is either 1 or 2 in case R is binary, and 1 in case R is unary). If R is a binary logical rule or R is a unary rule or C occurs in some side formula of a cut rule, then a path f' in Π is the *natural extension* of f if it is defined by extending f with the outgoing (incoming) edge from (to) C to (from) the variant C' in S (in the sense of the logical flow graph);

(2) let R be a cut rule, f_1 be a path in Π_1 passing through a subformula C occurring positively (negatively) in the cut-formula A of S_1, f_2 be a path in Π_2 passing through the corresponding subformula C of A occurring negatively (positively) in the cut-formula of S_2. Then, a path f' in Π is the *natural extension* of f_1, f_2 if it is defined by connecting f_1, f_2 with the outgoing (incoming) edge from (to) C in Π_1 to (from) C in Π_2 (in the sense of the logical flow graph).

We are now ready to define the notion of *flow*, that as we anticipated at the end of Section 2, will be used to bring out the intuition of *inner* proof. The definition is split into two cases. First we consider a proof Π to be an axiom, second we examine the last rule of inference in Π and define the flow with respect to flows existing for immediate subproof(s) of it.

4.4. Definition (*flow*). Let $\Pi : S$ be a proof with last rule R (if any) and logical flow graph $\mathscr{L}$.

(1) Suppose Π is an axiom of the form $A, \Gamma \to \Delta, A$; the subgraph $\mathscr{F}$ for Π of $\mathscr{L}$, is called a *flow* if it contains the edge linking the atomic distinguished occurrences of A (remember that by assumption, axioms are defined on atomic formulas);

(2.a) Suppose R is a $\vee$:*left*, or a $\wedge$:*right* applied to the subproofs $\Pi_1 : S_1$ and $\Pi_2 : S_2$ of Π. A subgraph $\mathscr{F}$ for Π is called a *flow* if one of the following conditions is satisfied:

1. there exist flows $\mathscr{F}_1, \mathscr{F}_2$ for Π_1, Π_2 (where the auxiliary formulas in S_1, S_2 might both be non-active, and) where $\mathscr{F}$ consists of exactly the natural extensions of paths in $\mathscr{F}_1$, $\mathscr{F}_2$, or
2. there exists a flow $\mathscr{F}_1$ for Π_1 (respectively $\mathscr{F}_2$ for Π_2) with active auxiliary formula in S_1 such that all natural extensions of paths in $\mathscr{F}_1$ (respectively $\mathscr{F}_2$) are in $\mathscr{F}$; any other (possibly none) path in $\mathscr{F}$ is an s-formula of S for which there exists a logical edge to/from some side formula in S_2 (respectively $\mathscr{F}_1$), or
3. there exists a flow $\mathscr{F}_1$ for Π_1 (respectively $\mathscr{F}_2$ for Π_2) with non-active auxiliary formula in S_1 such that all natural extensions of paths in $\mathscr{F}_1$ (respectively $\mathscr{F}_2$) are in $\mathscr{F}$; any other (possibly none) path in $\mathscr{F}$ is an s-formula of S for which there exists a logical edge to/from some side formula in S_2 (respectively $\mathscr{F}_1$), or is an s-formula of the main formula in S.

(Note that in the last parts of cases 2 and 3, only the s-formula in S is included in the flow, and not its counterpart in S_2 (S_1).)

(2.b) Suppose R is a *cut* rule applied to the subproofs $\Pi_1 : S_1$ and $\Pi_2 : S_2$ of Π, and let A be the cut-formula of R. A subgraph $\mathscr{F}$ for Π is called a *flow* if

1. there exist flows $\mathscr{F}_1, \mathscr{F}_2$ for Π_1, Π_2 such that $H_{\mathscr{F}_1}(A)$ and $H_{\mathscr{F}_2}(A)$ are the same inner formulas and $\mathscr{F}$ consists of exactly the natural extensions of paths in $\mathscr{F}_1, \mathscr{F}_2$, or
2. there exists a flow $\mathscr{F}_1$ for Π_1 (respectively $\mathscr{F}_2$ for Π_2) where $H_{\mathscr{F}_1}(A)$ (respectively $H_{\mathscr{F}_2}$) is empty, the flow $\mathscr{F}$ is the natural extension of $\mathscr{F}_1$ (respectively $\mathscr{F}_2$); any other (possibly none) path in $\mathscr{F}$ is an s-formula of S for which there exists a logical edge to/from some side formula in S_2 (respectively S_1). (Again only the s-formula in S is included in the flow in this last part.)

(2.c) Suppose R is a *unary* rule applied to the subproof $\Pi_1 : S_1$ of Π. A subgraph $\mathscr{F}$ for Π is called a *flow* if there exists a flow $\mathscr{F}_1$ for Π_1 such that all the natural extensions of paths in $\mathscr{F}_1$ are in $\mathscr{F}$.

(2.d) Suppose R is a *contraction* rule applied to $\Pi_1 : S_1$. A subgraph $\mathscr{F}$ is called a *flow* if there exists a flow $\mathscr{F}_1$ for Π_1 such that both images of contraction formulas in Π_1 are the same under $H_{\mathscr{F}_1}$ and where $\mathscr{F}$ consists of exactly the natural extensions of paths in $\mathscr{F}_1$.

This concludes the definition of *flow*. The sequent $H_{\mathscr{F}}(S)$ induced by $\mathscr{F}$ is called a *base for* $\mathscr{F}$. We say that an occurrence of the formula A in S does not belong to the base whenever it is not active (i.e. when $H_{\mathscr{F}}(A) = \emptyset$).

Notice that a flow $\mathscr{F}$ on Π with a binary rule R as last rule of inference, is sometimes passing only through one of the immediate subproofs of Π; more formally, the flow $\mathscr{F}$ does not contain subgraphs of one of the immediate subproofs of Π. Moreover, notice that the images under $H_{\mathscr{F}}$ of the auxiliary formulas for R being a cut rule or a contraction rule, must coincide. The flow $\mathscr{F}$ obtained from $\mathscr{F}_1, \mathscr{F}_2$ (possibly only $\mathscr{F}_1$) is also referred to as *natural extension* of $\mathscr{F}_1, \mathscr{F}_2$.

4.5. Remark. As a consequence of Definition 4.4, for all subproofs Π' of Π the restriction of a flow $\mathscr{F}$ to Π' (i.e. an edge belongs to the restriction if belongs both to $\mathscr{F}$ and to the logical flow graph of Π') whenever non-empty is a flow for Π'.

The crucial property of flows is emphasized in the following proposition. Its proof shows how, given a flow for a proof, one can 'extract' an inner proof from it.

4.6. Proposition. *Let $\Pi : S$ be a proof of k lines and $\mathscr{F}$ a flow for Π. Then there exists a proof $\Pi' : H_{\mathscr{F}}(S)$ with at most k lines.*

Proof. By induction on the height of Π.

$h(\Pi) = 1$: the proof Π must be an axiom of the form $A, \Gamma \to \Delta, A$; if $\mathscr{F}$ exists then $H_{\mathscr{F}}(A) = A$. Since $A, H_{\mathscr{F}}(\Gamma) \to H_{\mathscr{F}}(\Delta), A$ is an axiom then let Π' be $A, H_{\mathscr{F}}(\Gamma) \to H_{\mathscr{F}}(\Delta), A$.

$h(\Pi) = k + 1$: assume that $\mathscr{F}$ is a flow for Π; we can proceed by inspecting the form of the last rule of inference R.

Let R be a $\neg$: *left* rule applied to $\Pi_1 : \Gamma_1 \to \Delta_1, A$. Then apply the induction hypothesis to Π_1 and obtain a proof Π_1' such that $\#\ lines(\Pi_1') \leqslant \#\ lines(\Pi_1)$. If A is active in $\mathscr{F}$ then apply to Π_1' a $\neg$: *left*; if not, notice that $H_{\mathscr{F}}(S_1)$ is $H_{\mathscr{F}}(S)$ and therefore Π' can be taken to be Π_1'. Clearly $\#\ lines(\Pi') \leqslant \#\ lines(\Pi)$.

The $\neg$: *right* case is handled similarly.

Let R be a $\vee$: *left* rule applied to $\Pi_1 : A, \Gamma_1 \to \Delta_1$ and $\Pi_2 : B, \Gamma_2 \to \Delta_2$. If $\mathscr{F}$ is defined in both Π_1 and Π_2 then A and B might be both active and the induction hypothesis can be applied to Π_1, Π_2 to find proofs Π_1', Π_2' to combine with a $\vee$: *left* and obtain Π'. Clearly $\#\ lines(\Pi') = \#\ lines(\Pi_1') + \#\ lines(\Pi_2') + 1 \leqslant \#\ lines(\Pi_1) + \#\ lines(\Pi_2) + 1 = \#\ lines(\Pi)$. If at most one of the formulas A, B is active, say A, then apply the induction hypothesis to Π_1 to find the proof Π_1'. Then apply Lemma 4.8 to Π_1' to add all weak occurrences $H_{\mathscr{F}}(C)$ for C in Γ_2, Δ_2. Let Π' be the resulting proof. If both A, B are non-active, then let Π' be either Π_1' or Π_2', and add weak occurrences as before. If $\mathscr{F}$ is defined only in Π_1 (respectively in Π_2) then apply the induction hypothesis to it to obtain Π_1' and Lemma 4.8 to add all weak occurrences $H_{\mathscr{F}}(C)$ for C in Γ_2, Δ_2 (respectively Γ_1, Δ_1). If $H_{\mathscr{F}}(A) = \emptyset$ then apply Lemma 4.8 to add $H_{\mathscr{F}}(A \vee B)$ as well. Call the resulting proof Π'. Clearly $\#\ lines(\Pi_1') \leqslant \#\ lines(\Pi_1)$ and by Lemma 4.8 $\#\ lines(\Pi') \leqslant \#\ lines(\Pi_1') \leqslant \#\ lines(\Pi)$. The $\wedge$: *right* rule is handled similarly.

Let R be a $\vee$: *right* rule applied to $\Pi_1 : \Gamma_1 \to \Delta_1, A, B$ and Π with end-sequent $\Gamma_1 \to \Delta_1, A \vee B$. If A, B are active in $\mathscr{F}$ then apply the induction hypothesis to Π_1 to obtain Π_1' and the $\vee$: *right* rule. In case either A or B is not active, then apply the induction hypothesis to Π_1. Clearly $\#\ lines(\Pi_1') \leqslant \#\ lines(\Pi_1)$ and therefore $\#\ lines(\Pi') \leqslant \#\ lines(\Pi)$.

The *contraction* rule and the $\wedge$: *left* rule are handled similarly. Note that condition (2.d) in the definition of flow for a contraction rule, plays a crucial role here.

Let R be a *cut* rule applied to $\Pi_1 : \Gamma_1 \to \Delta_1, A$ and $\Pi_2 : A, \Gamma_2 \to \Delta_2$. If the cut-formulas are active in $\mathscr{F}$ then apply the induction hypothesis to Π_1, Π_2 to obtain

the sequents $H_{\mathscr{F}}(\Gamma_1) \to H_{\mathscr{F}}(\Delta_1), H_{\mathscr{F}}(A)$ and $H_{\mathscr{F}}(A), H_{\mathscr{F}}(\Gamma_2) \to H_{\mathscr{F}}(\Delta_2)$, where both distinguished occurrences of $H_{\mathscr{F}}(A)$ are the same inner formula of A. Then apply a cut rule to these sequents. If $\mathscr{F}$ is active in Π_1 (or similarly Π_2) but not in the cut formula A of $\Gamma_1 \to \Delta_1, A$ (respectively, $A, \Gamma_1 \to \Delta_1$), we can apply the induction hypothesis to obtain Π_1' and Lemma 4.8 to add the weak occurrences $H_{\mathscr{F}}(B)$ for B in Γ_2, Δ_2 (respectively Γ_1, Δ_1). □

4.7. Corollary. *Let $\Pi : S$ be a proof and $\mathscr{F}$ a flow for it. Then $H_{\mathscr{F}}(S) \neq \emptyset$.*

Proof. By induction on the height of Π. Notice that if the last rule of Π is a cut rule applied to $\Pi_1 : \Gamma_1 \to \Delta_1, A^1$ and $\Pi_2 : A^2, \Gamma_2 \to \Delta_2$ such that $H_{\mathscr{F}}(A^1) = H_{\mathscr{F}}(A^2) \neq \emptyset$ and $H_{\mathscr{F}}(\Gamma_1) = H_{\mathscr{F}}(\Delta_1) = H_{\mathscr{F}}(\Gamma_2) = H_{\mathscr{F}}(\Delta_2) = \emptyset$, then by Proposition 4.6, the sequents $\to A$ and $A \to$ are both provable. This is in contradiction with the consistency of the calculus. □

4.2. Structural properties of propositional proofs

In this section we present some properties of the structure of proofs in the propositional calculus. We will present a number of transformations of proofs preserving the number of lines (as well as the number of symbols) and satisfying certain natural properties over logical flow graphs. Their interest is independent of the particular use we make of them in this paper.

The first result essentially says that to each sequent we can always add new weak formulas without augmenting the complexity of the proof. In other words, the usual weakening rule can be simulated by the system *PLK*. Remember that a *weak occurrence* in a proof Π was defined to be a formula A whose direct paths all go to weak formulas A in axioms of Π.

4.8. Lemma (Addition of weak occurrences). *Let $\Pi : \Gamma \to \Delta$ be a proof of k lines and Λ, Θ be multisets. A proof $\Pi' : \Gamma, \Lambda \to \Delta, \Theta$ with k lines can be constructed such that if Π' has a flow $\mathscr{F}'$ then there exists a flow $\mathscr{F}$ for Π with base $H_{\mathscr{F}}(\Gamma \to \Delta) = H_{\mathscr{F}'}(\Gamma \to \Delta)$.*

Proof. By induction on the height of Π. If Π is an axiom then Π' is an axiom as well. If Π ends with a rule of inference R then apply the induction hypothesis to one of the premises whenever R is binary, or to the only one premise whenever R is unary; apply again the rule R to the result.

Since the structure of Π' is essentially the same as the structure of Π (the only difference consists on the presence of weak occurrences Λ, Θ in Π') then $\#\ lines(\Pi') = \#\ lines(\Pi)$ and any flow of Π' is a flow of Π as well whenever it is restricted to formulas of Π (notice that the formulas in Λ, Θ are weak occurrences, hence they do not interfere with the logical structure of Π). □

We can always eliminate weak occurrences from an end-sequent without augmenting the complexity of the proof.

4.9. Lemma (Elimination of weak occurrences). *Let* $\Pi : A, \Gamma \to \Delta$ ($\Pi : \Gamma \to \Delta, A$) *be a proof with k lines (possibly with cuts) and A be a weak occurrence. Then there is a proof* $\Pi' : \Gamma \to \Delta$ *with k lines such that* Π' *has a flow* $\mathscr{F}'$ *if and only if* Π *has a flow* $\mathscr{F}$. *Moreover, for each flow* $\mathscr{F}'$ *and any inner formula* A' *of* A *there is a flow* $\mathscr{F}$ *for* Π *such that* $H_{\mathscr{F}}(\Gamma \to \Delta) = H_{\mathscr{F}'}(\Gamma \to \Delta)$ *and* $H_{\mathscr{F}}(A) = A'$.

Proof. By induction on the height of Π. If $A, \Gamma \to \Delta$ is an axiom then $\Gamma \to \Delta$ is an axiom as well (since A is a non-distinguished occurrence). If R is the last rule of inference in Π we apply the induction hypothesis to the premise where a connected variant of A occurs and the rule R to the resulting proof. This can be done because a weak occurrence can be a main formula in Π only for a contraction rule. If R is a contraction rule on two variants of A, we should apply the induction hypothesis twice. Otherwise we apply it only once. We call Π' the resulting proof. Clearly Π' is a proof of k lines.

If Π' has a flow $\mathscr{F}'$ then define $\mathscr{F}$ in Π initially from $\mathscr{F}'$ in the obvious manner. Clearly the direct paths from (to) A in Π are non-active. In particular they will end up in some weak occurrences of axioms of Π. If the axiom is active, then force A in it to be active in $\mathscr{F}$ with image A' and activate the paths from the active s-formulas of A till the end-sequent. In case the axiom is not active, consider the first active sequent where the path is passing through and activate the variant A in it with image A' and the paths from its active s-formulas till the end-sequent. □

The idea behind the following propositions is that for any non-atomic formula in the end-sequent of a proof we can always find a proof of its 'premises' without augmenting the complexity of the proof, sometimes reducing it indeed.

4.10. Proposition (Elimination of negative conjunctions). *Let* $\Pi : A \wedge B, \Gamma \to \Delta$ *be a proof (possibly with cuts) of k lines. Then there is a proof* $\Pi' : A, B, \Gamma \to \Delta$ *of at most* $k - 1$ *lines whenever* $A \wedge B$ *is non-weak, and of at most k lines if* $A \wedge B$ *is weak. Moreover, if* $\mathscr{F}'$ *is a flow for* Π', *then there is a flow* $\mathscr{F}$ *for* Π *such that* $H_{\mathscr{F}'}(\Gamma \to \Delta) = H_{\mathscr{F}}(\Gamma \to \Delta)$, $H_{\mathscr{F}'}(A) = H_{\mathscr{F}}(A)$ *and* $H_{\mathscr{F}'}(B) = H_{\mathscr{F}}(B)$.

Proof. By induction on the number of lines of the proof.

If # *lines*$(\Pi) = 1$ then Π is of the form $A \wedge B, \Gamma \to \Delta$ where $A \wedge B$ is a weak occurrence (since axioms are defined on atomic formulas only). Define Π' to be the axiom $A, B, \Gamma \to \Delta$ with A, B weak occurrences.

If # *lines*$(\Pi) = k > 1$ then let R be the last rule of inference of Π. If $A \wedge B$ is not the main formula for R then we simply apply the induction hypothesis to the premise(s) and again R to define Π'. If it is the main formula for R then there are two cases we should consider.

Suppose R is a $\wedge$: *left* rule applied to $\Pi_1 : A, B, \Gamma \to \Delta$. Then define Π' to be Π_1.

Suppose R is a *contraction* applied to $\Pi_1 : A \wedge B^1, A \wedge B^2, \Gamma \to \Delta$. If either $A \wedge B^1$ or $A \wedge B^2$ is weak then apply Lemma 4.9 to Π_1, to obtain a proof $\Pi_1' : A \wedge B, \Gamma \to \Delta$ of $k - 1$ lines. Then apply the induction hypothesis to Π_1' to prove the assertion. If neither one of the contraction formulas is weak, then apply the induction hypothesis twice to obtain the proof $\Pi_1' : A, A, B, B, \Gamma \to \Delta$ of $k - 3$ lines. By applying suitable contraction rules, we obtain a proof $\Pi' : A, B, \Gamma \to \Delta$ of $k - 1$ lines.

This concludes the construction of Π'.

Given a flow $\mathscr{F}'$ for Π', we will define $\mathscr{F}$ essentially as $\mathscr{F}'$. If # *lines*$(\Pi) = 1$ then we force all s-formulas in $A \wedge B, \Gamma \to \Delta$ to be active whenever their corresponding variant is active in $A, B, \Gamma \to \Delta$. If # *lines*$(\Pi) > 1$ then we define $\mathscr{F}$ essentially as natural extension of flows in the premise(s) of the last rule of inference R. In the *contraction* case we might need to use Lemma 4.9. □

There is a dual version of Proposition 4.10 regarding formulas $A \vee B$ occurring positively in the end-sequent.

4.11. Proposition (Elimination of positive disjunctions). *Let $\Pi : \Gamma \to \Delta, A \vee B$, be a proof (possibly with cuts) of k lines. Then there is a proof $\Pi' : \Gamma \to \Delta, A, B$ of at most $k - 1$ lines whenever $A \vee B$ is non-weak, and of at most k lines if $A \vee B$ is weak. Moreover, if $\mathscr{F}'$ is a flow for Π', then there is a flow $\mathscr{F}$ for Π such that $H_{\mathscr{F}'}(\Gamma \to \Delta) = H_{\mathscr{F}}(\Gamma \to \Delta)$, $H_{\mathscr{F}'}(A) = H_{\mathscr{F}}(A)$ and $H_{\mathscr{F}'}(B) = H_{\mathscr{F}}(B)$.*

Propositions 4.10 and 4.11 speak about unary rules for binary connectives in a proof. Similar properties can be proved for binary rules.

4.12. Proposition (Elimination of negative disjunctions). *Let $\Pi : A \vee B, \Gamma \to \Delta$ be a proof (possibly with cuts) of k lines. Then there are proofs $\Pi^A : A, \Gamma \to \Delta$ and $\Pi^B : B, \Gamma \to \Delta$ of at most k lines. Moreover any pair of variants in Γ, Δ which are connected in Π are also connected in Π^A and Π^B.*

Proof. By induction on # *lines*(Π). The case of # *lines*$(\Pi) = 1$ and the case where $A \vee B$ is not a main formula for the last rule of inference R, are treated as in Proposition 4.10. In case R is a $\vee$: *left* rule applied to $\Pi_1 : A, \Gamma_1 \to \Delta_1$ and $\Pi_2 : B, \Gamma_2 \to \Delta_2$, let Π^A be the proof defined by applying a cut to $\Pi_1' : A, \Gamma_1 \to \Delta_1, B$ and $\Pi_2 : B, \Gamma_2 \to \Delta_2$ (where Π_1' is obtained from Π_1 by adding B as a weak occurrence). (It is important to use the cut rule here instead of adding weak occurrences to maintain the logical connections in the proof.) Define Π^B similarly. Clearly # *lines*(Π^A), # *lines*(Π^B) = # *lines*(Π). In case R is a *contraction* rule applied to $\Pi_1 : A \vee B^1, A \vee B^2, \Gamma \to \Delta$ then we proceed as follows. If $A \vee B^1$ (analogously $A \vee B^2$) is weak then apply Lemma 4.9 to obtain $\Pi_1' : A \vee B^2, \Gamma \to \Delta$ of $k - 1$ lines and then the induction hypothesis to Π_1' to prove the assertion. If neither one of the contraction formulas is weak, then apply the induction hypothesis to $A \vee B^1$ to obtain $\Pi_1^{A'} : A, A \vee B^2, \Gamma \to \Delta$ and $\Pi_1^{B'} : B, A \vee B^2, \Gamma \to \Delta$

(of $k-1$ lines) then we apply it a second time to $A \vee B^2$ in $\Pi_1^{A'}$ and $\Pi_1^{B'}$ to obtain $\Pi_1^{A''} : A, A, \Gamma \to \Delta$ and $\Pi_1^{B''} : B, B, \Gamma \to \Delta$ (of $k-1$ lines). We define Π_1^A, Π_1^B (of k lines) from $\Pi_1^{A''}, \Pi_1^{B''}$ by applying a contraction on A and on B respectively. □

4.13. Remark. If we do not ask that logical relations between formulas Π be preserved, then we can find proofs Π^A, Π^B of at most $k-1$ lines. The proof is the same as for Proposition 4.12 except when R is a $\vee : left$ rule. In this case we simply add weak occurrences Γ_2, Δ_2 to Π_1 and we save at least one sequent.

The dual version of Proposition 4.12 is stated as follows:

4.14. Proposition (Elimination of positive conjunctions). *Let $\Pi : \Gamma \to \Delta, A \wedge B$ be a proof (possibly with cuts) of k lines. Then there are proofs $\Pi^A : \Gamma \to \Delta, A$ and $\Pi^B : \Gamma \to \Delta, B$ of at most k lines. Moreover any pair of variants in Γ, Δ which are connected in Π are also connected in Π^A and Π^B.*

We will call Π an *amalgam* of Π^A and Π^B if the latter are produced in the manner of Propositions 4.12 and 4.14. Note that in Propositions 4.12 and 4.14 we did not state how flows behave with respect to amalgams. We will do this in Proposition 4.24.

4.15. Remark. It should be noticed here that Proposition 4.12 (and Proposition 4.14, respectively) does not define Π^A, Π^B from Π just by *reversing* the order of the rules of inference. This operation would indeed remove contractions on $A \vee B$ but it would not imply certain properties of flows we desire to be satisfied. We illustrate this point with an example. Take the proof Π with flow $\mathscr{F}$

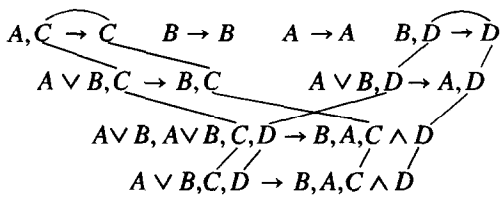

and the proof Π' (obtained by reversing in Π the order of the rules) with flow $\mathscr{F}'$

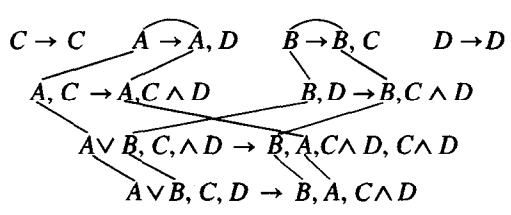

It is easy to observe that the *validity* of the end-sequent $A \vee B, D, C \to A, B, C \wedge D$ is 'justified' by different inner sequents in Π and Π' (i.e. $C, D \to C \wedge D$ and $A \vee B \to A, B$ respectively). In particular Π and Π' do not have flows of base $A \vee B \to A, B$ and

$C, D \to C \wedge D$ respectively. Proposition 4.24 will ensure that by transforming Π into Π' (as described in Proposition 4.12) 'justifications' are preserved.

4.16. Proposition (Elimination of negations). *Let $\Pi : \neg A, \Gamma \to \Delta$ (respectively $\Pi : \Gamma \to \Delta, \neg A$) be a proof (possibly with cuts) of k lines. Then there is a proof $\Pi' : \Gamma \to \Delta, A$ (respectively $\Pi' : A, \Gamma \to \Delta$) of at most $k-1$ lines whenever $\neg A$ is non-weak, and of at most k lines if $\neg A$ is weak. Moreover, if $\mathcal{F}'$ is a flow for Π', then there is a flow $\mathcal{F}$ for Π such that $H_{\mathcal{F}'}(\Gamma \to \Delta) = H_{\mathcal{F}}(\Gamma \to \Delta)$, $H_{\mathcal{F}'}(A) = H_{\mathcal{F}}(A)$.*

Proof. By induction on the # *lines*(Π), as in the proof of Proposition 4.10. □

Now we present a structural result concerning *contractions* in proofs. We say that a proof Π is *separated* when all pairs of directed paths from its non-weak contraction formulas, parallel each other until some binary rule separates them. We show that given an arbitrary proof of k lines, one can always find a separated proof of the same sequent of at most k lines. More formally.

4.17. Definition. A proof Π is *separated* if for all pairs of non-weak contraction formulas C^1, C^2 in Π and all pairs of direct paths f^1, f^2 from (to) C^1, C^2 respectively, there is a binary rule applied to subproofs $\Pi_1 : C, \Gamma_1 \to \Delta_1$ and $\Pi_2 : C, \Gamma_2 \to \Delta_2$ (respectively $\Pi_1 : \Gamma_1 \to \Delta_1, C$ and $\Pi_2 : \Gamma_2 \to \Delta_2, C$) such that the occurrence C in Π_i is a variant in f^i, for $i = 1, 2$. The rule R is called *separation rule* and the proofs Π_1, Π_2 are called *separation subproofs* for C.

Note that the proof of the sequent $A \to A$ presented in Section 2 and showing the presence of a cycle in its logical flow graph, is *not* separated.[4]

4.18. Proposition. *Let $\Pi : \Gamma \to \Delta$ be a proof (possibly with cuts) of k lines. Then there is a separated proof $\Pi' : \Gamma \to \Delta$ of at most k lines.*

Proof. Without loss of generality we can assume the last rules of inference of Π to be n contraction rules applied to $\Pi^* : C^0, C^1, \ldots, C^n, \Gamma \to \Delta$ where $C^0, C^1, \ldots, C^n$ are non-weak and the last rule of Π^* is different from a contraction over some C^i. If one of the C^i's is weak, notice that by applying Lemma 4.9 we can always eliminate it and define Π' to be the resulting proof followed by $n-1$ contractions. Moreover, if # *lines*(Π^*) $= 1$ then $n-1$ occurrences of C must be weak and in this case again the contraction rules can be eliminated. So, let R be the last rule of inference of Π^* and let h be the number of lines in Π^*.

We show the claim by induction on h.

If none of the C^i's is a main formula for R then we might have two situations. If R is binary, the formulas $C^0, C^1, \ldots, C^n$ might belong to distinct premises, say $\Pi_1^* : C^0, C^1, \ldots, C^m, \Gamma_1 \to \Delta_1$ and $\Pi_2^* : C^{m+1}, \ldots, C^n, \Gamma_2 \to \Delta_2$ with number of lines h_1, h_2

[4] It can be shown that there are separated proofs *with* cycles though.

respectively (where $h_1 + h_2 + 1 = h$). In this case we consider the proofs Π_1^*, Π_2^* respectively followed by m and $n-m-1$ contractions over the C^i's, and we apply the induction hypothesis to the resulting proofs to obtain separated proofs $\Pi_1^{**} : C, \Gamma_1 \to \Delta_1$ and $\Pi_2^{**} : C, \Gamma_2 \to \Delta_2$ of $h_1 + m$ and $h_2 + n - m - 1$ lines respectively. By applying first the binary rule R and then a contraction rule to the C occurrences, we obtain a separated proof of $h_1 + h_2 + n + 1 = h + n$ lines of $C, \Gamma \to \Delta$.

In case $C^0, C^1, \ldots, C^n$ belong to the same premise then we proceed similarly (by applying the induction hypothesis to a proof of at most $h - 1$ lines).

Suppose now that one of the C^i's is a main formula for R, say $i = 0$. Notice that R is *not* a contraction on formulas C by hypothesis, so we need to handle only two cases.

If R is a $\wedge$: *left* rule applied to $\Pi_1 : A^0, B^0, A \wedge B^1, \ldots, A \wedge B^n, \Gamma \to \Delta$ then we can apply Proposition 4.10 to all $A \wedge B^i$'s (for $i = 1 \cdots n$) to obtain a proof $\Pi_1' : A^0, B^0, A^1, B^1, \ldots, A^n, B^n, \Gamma \to \Delta$ of $h - 1 - n$ lines (remember that the $A \wedge B$'s are non-weak). Applying twice the induction hypothesis over the A's and B's to Π_1' followed by n contractions over A and n contractions over B, we obtain a separated proof of $A, B, \Gamma \to \Delta$ with $h + n - 1$ lines (note that the first induction hypothesis is applied to $h - n - 1$ and the second to $h - 1$). By applying the rule R we prove the statement.

If R is a $\vee$: *left* rule applied to $\Pi_1 : A^0, A \vee B^1, \ldots, A \vee B^m, \Gamma_1 \to \Delta_1$ and $\Pi_2 : B^0, A \vee B^{m+1}, \ldots, A \vee B^n, \Gamma_2 \to \Delta_2$ (with $m \geqslant 0$ and h_1, h_2 lines respectively, where $h = h_1 + h_2 + 1$) then we can apply Proposition 4.12 to the $A \vee B$'s in Π_1 and Π_2 to obtain proofs $\Pi_1' : A^0, A^1, \ldots, A^m, \Gamma_1 \to \Delta_1$ and $\Pi_2' : B^0, B^{m+1}, \ldots, B^n, \Gamma_2 \to \Delta_2$ of at most h_1, h_2 lines. Applying the induction hypothesis to Π_1' and Π_2' followed by m and $n - m$ contractions respectively, we find two separated proofs of $A, \Gamma_1 \to \Delta_1$ and $B, \Gamma_2 \to \Delta_2$ respectively. To prove the statement we need only to apply R to the pair of separated proofs.

The $\wedge$: *right* and $\vee$: *right* cases are handled similarly. For the $\neg$: *left* and $\neg$: *right* cases we should use Proposition 4.16. □

Since we are interested in the complexity (i.e. the number of lines[5]) of proofs for a given sequent, the last proposition says that we can always restrict ourselves to separated proofs. Moreover from its proof it directly follows that for any path linking two occurrences in $\Gamma \to \Delta$ of Π', there is a path linking the corresponding occurrences of $\Gamma \to \Delta$ in Π (but not the other way around). In this sense, separated proofs turn out to be a better approximation for what is *essential* to the provability of a sequent.

Let us observe now that all constructions used in Propositions 4.10–4.16, transform separated proofs Π into separated proofs Π', Π^A, Π^B. We show this fact only for Proposition 4.10 since the other cases can be seen similarly.

[5] It can be easily checked that the size of the proof Π' in Proposition 4.18 is also bounded by the size of Π.

4.19. Proposition. *Let Π and Π' as in Proposition* 4.10. *If Π is separated then Π' is separated.*

Proof. By induction on the steps of construction of Π' from Π. Notice that Π' is essentially obtained from Π by *deleting* $\wedge$*:left* rules over pairs A,B wherever they are applied, and by substituting weak formulas $A \wedge B$ in axioms of Π with pairs A,B. Hence, the construction does not essentially change the tree-structure of the proof (i.e. for each subproof of Π there is a corresponding subproof of Π'). Since the construction does not touch any formula other than $A \wedge B$, this implies that all pairs of contraction formulas that do *not* have direct link to $A \wedge B$ in Π', satisfy the condition of *separation* in Π' (in fact, they were satisfying the condition of separation already in Π).

In case the *contraction* rule R is applied to $\Pi_1 : A \wedge B^1, A \wedge B^2, \Gamma \to \Delta$, notice that the construction furnishes a proof for $A^1,B^1,A^2,B^2,\Gamma \to \Delta$ where A^1,B^1 and A^2,B^2 correspond pairwise to the occurrences A,B in $A \wedge B^1, A \wedge B^2$ respectively. Since Π is separated, for any pair of direct paths f^1, f^2 from $A \wedge B^1, A \wedge B^2$ respectively, there are in Π_1 two subproofs, say $\Pi_1^1 : A \wedge B, \Gamma_1 \to \Delta_1$ and $\Pi_1^2 : A \wedge B, \Gamma_2 \to \Delta_2$ such that the $A \wedge B$'s in Π_1^1 and Π_1^2 are variants of the direct paths f^1, f^2 respectively. The construction transforms Π_1^1 and Π_1^2 into two proofs $\Pi_1^{1'} : A,B,\Gamma_1 \to \Delta_1$ and $\Pi_1^{2'} : A,B,\Gamma_2 \to \Delta_2$ where A,B in $\Pi_1^{1'}$ and $\Pi_1^{2'}$ are linked respectively to A^1,B^1 and A^2,B^2. Therefore A^1,A^2 and B^1,B^2 satisfy the condition of separation in Π'. □

4.20. Remark. If Π is an amalgam of Π^A, Π^B and Π is separated, notice that the tree-like structures of Π^A, Π^B, Π are essentially the same (this is a direct consequence of the proof of Proposition 4.12). In particular to each pair of separation subproofs for $A \vee B$ $(A \wedge B)$ in Π, correspond a pair of separation subproofs for A in Π^A and a pair of separation subproofs for B in Π^B.

We finally state the last property. It is a property of flows for amalgams of *separated* proofs and it is fundamental for the proof of Theorem 4.27. In general given two flows for Π^A, Π^B with the same base (with the obvious exception for A and B), there is not a priori a flow for the amalgam Π preserving such base. The next example illustrates this point.

4.21. Example. Let Π be a *non*-separated proof of the form

$$\dfrac{\dfrac{\dfrac{A \to A \quad \dfrac{A \to A \quad B \to B}{A,B \to A \wedge B}}{A \vee B, A \to A, A \wedge B} \quad \dfrac{B \to B,C}{B \to B \vee C}}{A \vee B, A \vee B, \to A, A \wedge B, B \vee C}}{A \vee B \to A, A \wedge B, B \vee C}\, \textit{Contraction}$$

then by applying the procedure described in Proposition 4.12 we obtain the proofs Π^A of the form $A \to A, A \wedge B, B \vee C$ and Π^B of the form

$$\dfrac{B \to B, C, A, A \wedge B}{B \to B \vee C, A, A \wedge B}$$

Let $\mathscr{F}^A, \mathscr{F}^B$ be flows for Π^A, Π^B with base $H_{\mathscr{F}^A}(A \to A, A \wedge B, B \vee C) = A \to A, B \vee C$ and $H_{\mathscr{F}^B}(B \to A, A \wedge B, B \vee C) = B \to A, B \vee C$ respectively. Notice that there is no flow for Π with base $A \vee B \to A, B \vee C$. In fact any flow forcing the occurrence B on the left-hand side of the end-sequent in Π, must force both occurrences B in the right-hand side as well.

To restrict ourselves to *separated* proofs though, it is still not enough. Indeed given two arbitrary flows for Π^A and Π^B, we might not be able to combine them in a flow for their amalgam. We illustrate this point with the following example.

4.22. Example. Consider the proofs Π^A, Π^B

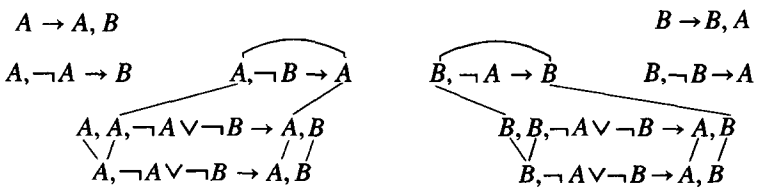

with flows $\mathscr{F}^A, \mathscr{F}^B$ of base $A \to A, B$ and $B \to A, B$ respectively, and amalgam Π

$$\dfrac{\dfrac{\dfrac{\dfrac{A \to A}{A, \neg A \to B \to B}}{A \vee B, \neg A \to B} \quad \dfrac{A \to A \quad \dfrac{B \to B}{B, \neg B \to}}{A \vee B, \neg B \to A}}{A \vee B, A \vee B, \neg A \vee \neg B \to B, A}}{A \vee B, \neg A \vee \neg B \to B, A}\ \textit{Contraction}$$

Both $\mathscr{F}^A, \mathscr{F}^B$ are flows with the same base but there is no flow of base $A \vee B \to A, B$ that can be defined over Π.

Given a proof Π with logical flow graph $\mathscr{L}$, we say that a subgraph $\mathscr{L}'$ of $\mathscr{L}$ is *compact* if $\mathscr{L}'$ is a union of connected components of $\mathscr{L}$ (i.e. any variant in $\mathscr{L}$ connected to a variant in $\mathscr{L}'$ belongs already to $\mathscr{L}'$, and any edge of $\mathscr{L}$ connecting two variants in $\mathscr{L}'$ is an edge of $\mathscr{L}'$), and if moreover any variant in $\mathscr{L}'$ has a connected variant in the end-sequent of Π. A compact subgraph which is a flow is called a *compact flow*.

Notice that a compact subgraph is uniquely determined by its base (derived from the end-sequent). This is easy to check, from the definitions. One can try to go backwards, take a collection of points in the logical flow graph of a proof that come from the end-sequent, and ask whether it is the base of a compact subgraph. One can always take the compact subgraph that these points generate, i.e. the smallest compact subgraph that contains them, but that might lead to additional points in the base. In order to know that the original collection of points is the base for a compact subgraph one needs to know that all variants in the end-sequent which are connected through the proof are already contained in the original collection.

Notice that a compact subgraph need not be a flow.

4.23. Proposition. *Let Π be a proof and $\mathscr{L}$ its logical flow graph. If $\mathscr{L}'$ is a compact subgraph of $\mathscr{L}$, then $\mathscr{L}'$ is a compact flow if and only if all weak occurrences in $\mathscr{L}'$ are supported by axioms whose distinguished formulas lie in $\mathscr{L}'$.*

Proof. By induction on the height of the proof; the claim follows readily from the definition of a flow. □

We show next that any pair of compact flows for Π^A, Π^B having the same base can be combined to obtain a compact flow for the amalgam Π which preserves that base.

4.24. Proposition. *Let $\Pi : A \vee B, \Gamma \to \Delta$ ($\Pi : \Gamma \to \Delta, A \wedge B$) be a separated cut-free proof (possibly containing trivial cuts, i.e. cuts on weak occurrences) and an amalgam for Π^A, Π^B. Let $\mathscr{F}^A$ and $\mathscr{F}^B$ be compact flows for Π^A and Π^B such that $H_{\mathscr{F}^A}(\Gamma \to \Delta) = H_{\mathscr{F}^B}(\Gamma \to \Delta)$. Then there is a compact flow $\mathscr{F}$ for Π such that $H_{\mathscr{F}}(\Gamma \to \Delta) = H_{\mathscr{F}^A}(\Gamma \to \Delta) = H_{\mathscr{F}^B}(\Gamma \to \Delta)$ and $H_{\mathscr{F}}(A) = H_{\mathscr{F}^A}(A)$ and $H_{\mathscr{F}}(B) = H_{\mathscr{F}^B}(B)$.*

Proof. We show the statement for $\Pi : A \vee B, \Gamma \to \Delta$. The proof works exactly the same way for $\Pi : \Gamma \to \Delta, A \wedge B$.

For all proofs $\Pi : A \vee B, \Gamma \to \Delta$, *all* direct paths from $A \vee B$ either end into an axiom of Π where the variant $A \vee B$ is weak or they end into a main formula of a $\vee$: *left* rule applied to some subproofs $\Pi_1^* : A, \Gamma_1 \to \Delta_1$ and $\Pi_2^* : B, \Gamma_2 \to \Delta_2$. If the proof Π is *separated*, in the second case we can say that there is no non-weak formula in $\Gamma_{1,2} \to \Delta_{1,2}$ directly linked to some subformula of $A \vee B$ in the end-sequent of Π. This is a direct consequence of the definition of *separated* proof. Therefore Π^A, Π^B are in *essence* the same as Π except that all variants in direct paths from $A \vee B$ to axioms in Π are substituted by variants A and B, and all subproofs Π^* (which combine proofs Π_1^* and Π_2^* with a $\vee$: *left* rule) are replaced with $\Pi_1^{*\prime}$ of the form

$$\frac{A, \Gamma_1 \overset{\Pi_1^{**}}{\to} \Delta_1, B \quad B, \Gamma_2 \overset{\Pi_2^*}{\to} \Delta_2}{A, \Gamma_{1,2} \to \Delta_{1,2}}$$

and $\Pi_2^{*\prime}$ of the form

$$\frac{B, \Gamma_2 \overset{\Pi_2^{**}}{\to} \Delta_2, A \quad A, \Gamma_1 \overset{\Pi_1^*}{\to} \Delta_1}{A, \Gamma_{1,2} \to \Delta_{1,2}}$$

(where Π_1^{**} and Π_2^{**} are obtained from Π_1^* and Π_2^* by adding B and A as weak occurrences, respectively). Moreover, all variants occurring in the path between $A \vee B$ in Π^* and $A \vee B$ in the end-sequent of Π are respectively substituted by A and B.

With this construction in mind (which we described in Proposition 4.12 for proofs possibly containing non-trivial cuts), we can define the flow $\mathscr{F}$ as follows.

If Π does *not* contain separation subproofs for $A \vee B$ (i.e. $A \vee B$ is not contracted over non-weak occurrences) then either $A \vee B$ is weak or there is at most one pair of subproofs Π_1^*, Π_2^* in Π whose main formula is a variant of $A \vee B$. Even though the construction is easy let us say precisely how $\mathscr{F}$ is defined.

In case $A \vee B$ is a weak occurrence with direct paths to some weak formulas in axioms of Π^*, then as in (the proof of) Proposition 4.12, the proofs Π_1^* and Π_2^* are essentially the same proof Π^* where A and B are added as weak occurrences in it. We choose to define $\mathscr{F}$ either forcing $\mathscr{F}^A$ or forcing $\mathscr{F}^B$ over Π^* in Π and imposing the image $H_{\mathscr{F}}(A \vee B) = H_{\mathscr{F}^A}(A) \vee H_{\mathscr{F}^B}(B)$ for all variants $A \vee B$ (belonging to the direct paths) occurring in active sequents of Π. In case $H_{\mathscr{F}^A}(A)$ ($H_{\mathscr{F}^B}(B)$) is $\emptyset$, the image $H_{\mathscr{F}}(A \vee B)$ will be $H_{\mathscr{F}^B}(B)$ ($H_{\mathscr{F}^A}(A)$) (notice that there are no connected variants of A and B in S other than A, B themselves). Since $\mathscr{F}^A, \mathscr{F}^B$ are compact, $\mathscr{F}$ will be compact.

In case $A \vee B$ is non-weak, then we claim that the compact subgraph with base $H_{\mathscr{F}^A}(A) \vee H_{\mathscr{F}^B}(B), H_{\mathscr{F}^A}(\Gamma) \to H_{\mathscr{F}^A}(\Delta)$ is a compact flow. To show this we will use the fact that $\mathscr{F}^A, \mathscr{F}^B$ are compact flows (with the same image on Γ, Δ) and the following property: if C^1, C^2 are connected variants of Π occurring in A, B respectively, then there is a variant C^3 in Γ, Δ such that C^1, C^3 are connected variants in Π^A and C^2, C^3 are connected variants in Π^B. (Notice that all paths in Π, Π^A, Π^B starting or ending in A (and similarly in B) can pass through exactly one axiom because the proofs have no cuts except on weak occurrences. The property follows readily from this. In particular, the property holds for all those proofs Π which contain cuts only on weak formulas.) The argument for proving our claim goes as follows. All connected subgraphs in the logical flow graph of Π^A and Π^B which are linked to some active occurrence in Γ, Δ, are active for both $\mathscr{F}^A$ and $\mathscr{F}^B$ (this is because $\mathscr{F}^A$ and $\mathscr{F}^B$ are compact). Therefore we can define $\mathscr{F}$ as in $\mathscr{F}^A, \mathscr{F}^B$ over all formulas of Π except for those variants which are connected to some occurrence in A, B; over those variants we should define $\mathscr{F}$ as in Π^A, Π^B respectively. Notice that any pair of connected variants C^1, C^2 in A, B of Π will be either both active or non-active since there is a variant C^3 connected to them (by the property mentioned above) and the image of C^3 under $\mathscr{F}^A, \mathscr{F}^B$ coincides by hypothesis. This shows that $\mathscr{F}$ is a compact subgraph. To show that $\mathscr{F}$ is a flow we use Proposition 4.23.

This concludes the construction of $\mathscr{F}$ in case Π does *not* contain separation subproofs. The flow $\mathscr{F}$ is defined in a similar way in case contractions over non-weak occurrences of $A \vee B$ exist in Π. To do this we argue as in the previous case. The compactness condition on $\mathscr{F}^A, \mathscr{F}^B$ guarantees that contraction formulas $A \vee B$ can be forced to have the same image under $\mathscr{F}$. □

4.25. Remark. If Π contains for instance cuts on atomic formulas (which are non-trivial) the claim in Proposition 4.24 can fail to hold. To see this consider a proof Π

with the following structure:

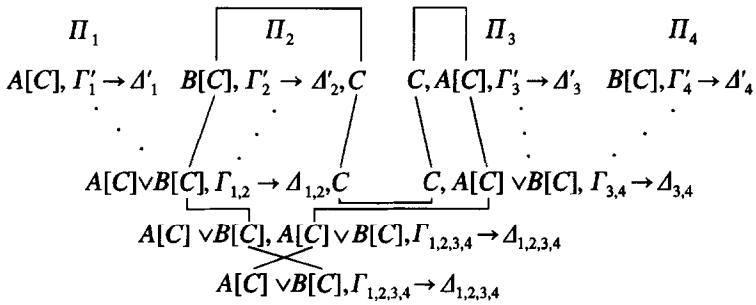

where the separation rule is a cut on a pair of non-weak atomic formulas C, and formula occurrences A, B, C are linked as indicated in the picture. As in Proposition 4.12 the proofs Π^A, Π^B will be of the form

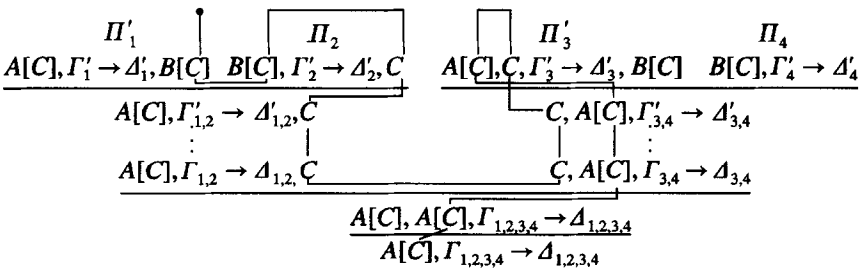

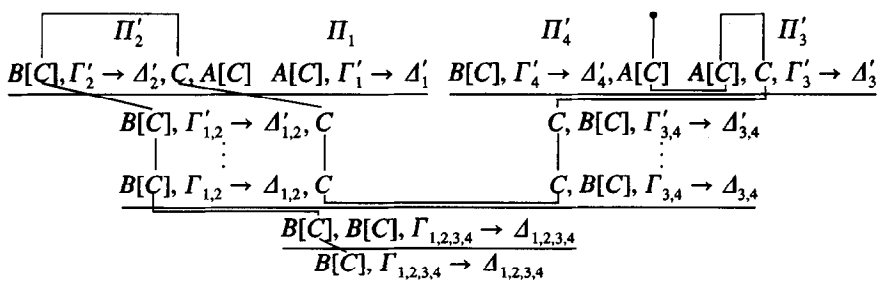

Notice that in Π^A (Π^B) the cut-formula C occurring on the right-hand (left-hand) side of the sequent ends up into a weak occurrence. Clearly given any flow $\mathscr{F}^A$ where C^A is active and any flow $\mathscr{F}^B$ where C^B is not active, we will not be able to define a flow for Π which is compact (i.e. the splitting of Π induces a *partition* of the set of connected variants).

4.26. Remark. The statement of Proposition 4.24 says that a compact flow $\mathscr{F}$ can be defined out of two compact flows $\mathscr{F}^A$ and $\mathscr{F}^B$. We can also try to produce $\mathscr{F}$ given

a compact flow $\mathscr{F}^A$ ($\mathscr{F}^B$) for Π^A (Π^B) without having a compact flow $\mathscr{F}^B$ ($\mathscr{F}^A$) for Π^B (Π^A). This can happen when $H_{\mathscr{F}^A}(A) \neq \emptyset$ and $H_{\mathscr{F}^A}(\Gamma \to \Delta) = \emptyset$ ($H_{\mathscr{F}^B}(B) \neq \emptyset$ and $H_{\mathscr{F}^B}(\Gamma \to \Delta) = \emptyset$). If this is the case we define $\mathscr{F}$ as $\mathscr{F}^A$ ($\mathscr{F}^B$).

4.3. The statement of the Theorem, an example and a procedure of cut elimination

We are now ready to state the main result of this section. Its local version (i.e. its relativization to *sub*proofs) described in the introduction, is a natural generalization of this statement and it is formulated in Theorem 4.34.

4.27. Theorem (Inversion property). *Let $\Pi : \Gamma \to \Delta$ (possibly with cuts) be a separated proof. Then there is an effective procedure that transforms Π into a proof $\Pi' : A_1 \vee \neg A_1, \ldots, A_n \vee \neg A_n, \Gamma \to \Delta$ with only trivial cuts if any, where the A_i's are atomic formulas. Moreover if $\mathscr{F}'$ is a compact flow for Π' with $H_{\mathscr{F}'}(A_i \vee \neg A_i)$ being either $A_i \vee \neg A_i$ or $\emptyset$ for each $i = 1, \ldots, n$, then Π has a compact flow $\mathscr{F}$ with base $H_{\mathscr{F}}(\Gamma \to \Delta) = H_{\mathscr{F}'}(\Gamma \to \Delta)$.*

The idea of the proof consists in showing that a compact flow for a proof can be transformed into a compact flow for another proof obtained by applying a certain procedure of cut elimination. We present a procedure different from the well known one proposed by Gentzen (see for instance [7]). The main reason for which this change is essential concerns the behaviour of flows through contraction formulas. More precisely, suppose Π to be of the form

$$\dfrac{\dfrac{\Pi_1}{\Gamma_1 \to \Delta_1, C} \quad \dfrac{\Pi_2}{C, C, \Gamma_2 \to \Delta_2}}{C, \Gamma_2 \to \Delta_2}}{\Gamma_{1,2} \to \Delta_{1,2}}$$
$$\vdots \lambda$$

Gentzen's procedure defines the transformation Π' as follows:

$$\dfrac{\dfrac{\Pi_1}{\Gamma_1 \to \Delta_1, C} \quad \dfrac{\dfrac{\Pi_1}{\Gamma_1 \to \Delta_1, C} \quad \dfrac{\Pi_2}{C, C, \Gamma_2 \to \Delta_2}}{C, \Gamma_{1,2} \to \Delta_{1,2}}}{\Gamma_{1,1,2} \to \Delta_{1,1,2}}$$
$$\vdots \textit{ contractions}$$
$$\Gamma_{1,2} \to \Delta_{1,2}$$
$$\vdots \lambda$$

If we have a flow for Π' passing through the occurrences of C in Π_2, it is not necessarily true that this flow would have the same image over both C occurrences (remember that this condition is required by the definition of flow), therefore we would not be a priori able to define from such a flow, a flow for Π. The following example illustrates the point.

4.28. Example. Let Π be the proof

$$\dfrac{\dfrac{\dfrac{A \to A}{\to A, \neg A} \quad C \to C}{C \to A \wedge C, \neg A}}{C \to (A \wedge C) \vee \neg A} \qquad \dfrac{\dfrac{\dfrac{A \to A}{A \wedge C \to A} \quad \dfrac{A \to A}{\neg A, A \to}}{(A \wedge C) \vee \neg A, A \to A} \qquad \dfrac{\dfrac{C \to C}{A \wedge C \to C} \quad \dfrac{\dfrac{A \to A}{\to \neg A, A}}{\neg A \to \neg A}}{(A \wedge C) \vee \neg A, \to \neg A, C}}{(A \wedge C) \vee \neg A, \to \neg A \vee C}}{(A \wedge C) \vee \neg A, (A \wedge C) \vee \neg A, A \to A \wedge (\neg A \vee C)}}{(A \wedge C) \vee \neg A, A \to A \wedge (\neg A \vee C)}$$

$$C, A \to A \wedge (\neg A \vee C)$$

and let $\mathscr{F}'$ be the flow for Π' defined by duplicating Π_1 over the contraction formulas $(A \wedge C) \vee \neg A$

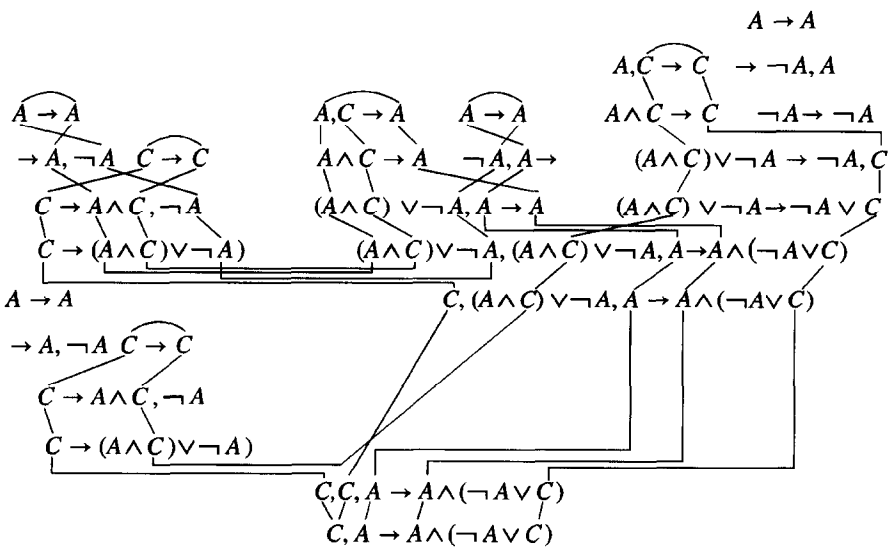

Clearly the contraction formulas in Π' have different images over the flow indicated in the picture (one of them is $(A \wedge C) \vee \neg A$ and the other C). On the other hand there is no flow $\mathscr{F}$ in Π having the same base as $\mathscr{F}'$. In fact any possible flow with some active occurrence A would force all occurrences of A in the end-sequent to be active. In particular, if the Gentzen's procedure is applied to eliminate all cuts from Π, it gives a proof of the form

$$\dfrac{\dfrac{}{A \to A} \quad \dfrac{C \to C}{C \to \neg A \vee C}}{C, A \to A \wedge (\neg A \vee C)}$$

having a flow with base $C, A \to A \wedge C$ (notice that this is the only cut-free form for our example which comes from Gentzen's procedure since Π' does not contain contractions on cut-formulas and the procedure in this case is essentially deterministic). Since the

original proof Π cannot have any flow with such a base, the claim fails to hold for the usual Cut Elimination procedure.[6]

Therefore, to prove the claim we modify the Gentzen's procedure to avoid the duplication of the proof Π_1 over the pair of contraction formulas C in Π'. The new procedure will depend on the logical form of C. Namely given any cut

$$\dfrac{\begin{array}{cc}\Pi_1 & \Pi_2\\ \Gamma_1 \to \Delta_1, C & C, \Gamma_2 \to \Delta_2\end{array}}{\Gamma_{1,2} \to \Delta_{1,2}}$$

in Π, if C is (non-weak and) of the form $A \wedge B$ then we define Π' to be of the form

$$\dfrac{\begin{array}{cc}\Pi_1^B & \dfrac{\begin{array}{cc}\Pi_1^A & \Pi_2'\\ \Gamma_1 \to \Delta_1, A & A, B, \Gamma_2 \to \Delta_2\end{array}}{B, \Gamma_{1,2} \to \Delta_{1,2}}\\ \Gamma_1 \to \Delta_1, B & \end{array}}{\begin{array}{c}\Gamma_{1,1,2} \to \Delta_{1,1,2}\\ \vdots \textit{contractions}\\ \Gamma_{1,2} \to \Delta_{1,2}\\ \vdots \lambda\end{array}}$$

where Π_2' is obtained by applying Proposition 4.10 to Π_2 (and Lemma 4.8 whenever needed) and Π_1^A, Π_1^B are obtained by Proposition 4.14 applied to Π_1. The case of $A \vee B$ can be treated in the same way.

If C is (non-weak and) of the form $\neg A$ then we apply Proposition 4.16 to Π_1 and Π_2 to obtain proofs $\Pi_1' : A, \Gamma_1 \to \Delta_1$ and $\Pi_2' : \Gamma_2 \to \Delta_2, A$ respectively. We define Π' as

$$\dfrac{\begin{array}{cc}\Pi_2' & \Pi_1'\\ \Gamma_2 \to \Delta_2, A & A, \Gamma_1 \to \Delta_1\end{array}}{\begin{array}{c}\Gamma_{1,2} \to \Delta_{1,2}\\ \vdots \lambda\end{array}}$$

If C is (non-weak and) *atomic* we will convert the cut into a disjunction $C \vee \neg C$ introduced on the left-hand side of the sequent. One can also think to eliminate cuts on atomics by following the usual procedure suggested by Gentzen. Since our purpose is to analyse the transformations of flows though, we will be content with eliminating cuts except for trivial and atomic ones (see Remark 4.31).

This concludes the description of the cut elimination procedure. Notice that our procedure can be applied to arbitrary proofs Π (i.e. the proof Π might not be *separated*). In fact Propositions 4.10, 4.12 and 4.16 do not require any hypothesis on the structure of the proof Π. The notion of separation is needed though to prove the claim because the property of flows described in Proposition 4.24 requires that the proofs be separated.

[6] Similar counterexamples can be constructed for proofs containing no weak formulas. They are a bit more complicated though.

4.29. Proposition. *Let Π be a separated proof. Any step of the procedure of cut elimination described above, transforms Π into a separated proof Π'.*

Proof. Suppose the cut-formula C has the form $A \wedge B$ and that the cut is applied to subproofs Π_1, Π_2 of Π. Then by Proposition 4.19 (and Remark 4.20) the proofs Π_1^A, Π_1^B, Π_2' are separated proofs because Π_1, Π_2 are separated.

Moreover notice that the procedure introduces new contractions on side formulas. These new contraction formulas come from different subproofs and therefore the resulting proof Π' is separated.

If C has the form $A \vee B$ or $\neg A$, the claim is shown similarly. □

4.30. Remark. Gentzen's procedure does *not* transform separated proofs into separated proofs. This is because of the step of permutation of cuts upwards (see [19]).

4.31. Remark. Since Proposition 4.24 should be applied to *separated* cut-free proofs, and Gentzen's procedure (even when restricted to cuts on atomics) does not preserve separation by Remark 4.30, we are forced to keep track of atomic cuts by transforming them into appropriate disjunctions.

4.32. Remark. Our procedure does not handle explicitly the case of weak cut-formulas. This is because we want to *preserve* the structure of the proof Π during the transformation (in order to preserve compactness) and the elimination of weak cut-formulas (as suggested in Gentzen's procedure) would lead to the pruning of some subproofs of it. By maintaining links between variants in Π, we essentially reduce our cut elimination procedure to two simple combinatorial operations; namely, we either stretch and relax paths of a flow graph, or disconnect those pairs of paths which start and end into the atomic occurrences A_i of those formulas $A_i \vee \neg A_i$ in the end-sequent of Π'.

4.33. Remark. To check whether or not a compact subgraph of a proof Π is a flow can be accomplished by checking just the axioms of Π. This is proved in Proposition 4.23 and holds for proofs Π containing cuts. One would like to have a way to decide if there is a flow for Π with a certain given base, where the flow might not have an underlying graph which is compact. One might hope to use cut elimination to analyze this question, but Example 4.22 shows that the property of being a flow is not preserved by the procedure. Theorem 4.27 tells us what is preserved in the process of eliminating cuts, however.

4.4. Proof of the Theorem

We are now ready to prove Theorem 4.27. We remind the reader that the *degree of a formula* is inductively defined as follows: $d(A) = 1$ if A is atomic, $d(A \wedge B) =$

$d(A \vee B) = \sup\{d(A)+1, d(B)+1\}$ and $d(\neg A) = d(A)+1$. The *degree of a cut* is the degree of its cut-formulas.

Proof of Theorem 4.27. If Π is cut-free then the claim holds.

In the following we will consider the procedure of cut elimination described in Section 4.3 and we will study how compact flows are transformed by it. The procedure transforms any cut of degree $\geqslant 1$ in Π into a cut(s) of lower degree, possibly no cuts. This operation must be repeated until every cut of degree $\geqslant 1$ has systematically been eliminated (except for trivial cuts). We require that each time a cut is eliminated the supporting subproofs are already free from non-trivial cuts (that is, we eliminate cuts from the top down). The proof Π' will be obtained from Π by a finite number ($k > 0$) of steps of reduction transforming Π_i into Π_{i+1} and where Π_1 is Π and Π_k is Π'. We say that Π_{i+1} has been obtained from Π_i if a step of the reduction procedure is applied to a cut on a formula in the proof Π_i.

Assume that Π_{i+1} has a compact flow with base $H_{i+1}(\Gamma \to \Delta)$; we want to show that Π_i has a compact flow where $H_i(\Gamma \to \Delta) = H_{i+1}(\Gamma \to \Delta)$ and with either full or empty images for those disjunctions $A_j \vee \neg A_j$ which are introduced to eliminate cuts on atomic formulas A_j, in some step $j \leqslant i$. Notice that this is sufficient to show the claim.

There are two kinds of reduction we must discuss:

(1) the reducibility of the complexity of a cut when it is of degree > 1;

(2) elimination of cuts of degree $= 1$.

Let us begin with reductions of the first kind. Suppose C is $A \wedge B$ and let Π_i, Π_{i+1} contain subproofs Π, Π' defined as in the description of the cut elimination procedure given in Section 4.3. We follow the notation used there.

By construction the subproof Π_1 contained as a subproof in Π_i is separated, it is cut-free (possibly it contains trivial cuts) and it is an amalgam of Π_1^A, Π_1^B. Moreover the flow $\mathscr{F}_{i+1}$ (defined over Π_{i+1}) when restricted to the subproofs Π_1^A, Π_1^B gives two compact flows $\mathscr{F}_1^A, \mathscr{F}_1^B$ where $H_{\mathscr{F}_1^A}(\Gamma_1 \to \Delta_1) = H_{\mathscr{F}_1^B}(\Gamma_1 \to \Delta_1)$ (this is because $\mathscr{F}_{i+1}$ is a flow and contractions on Γ_1, Δ_1 must have the same image by definition of a flow). Hence by Proposition 4.24, from $\mathscr{F}_1^A, \mathscr{F}_1^B$ we can define a compact flow $\mathscr{F}_1$ for Π_1 such that $H_{\mathscr{F}_1}(\Gamma_1 \to \Delta_1) = H_{\mathscr{F}_1^A}(\Gamma_1 \to \Delta_1) = H_{\mathscr{F}_1^B}(\Gamma_1 \to \Delta_1)$, $H_{\mathscr{F}_1}(A) = H_{\mathscr{F}_1^A}(A)$ and $H_{\mathscr{F}_1}(B) = H_{\mathscr{F}_1^B}(B)$. If only one of the flows is defined, say $\mathscr{F}_1^A$ (resp. $\mathscr{F}_1^B$) for instance, then apply Remark 4.26 instead of Proposition 4.24. Define $\mathscr{F}_i$ over Π_1 as $\mathscr{F}_1$. By Proposition 4.10 there is a compact flow $\mathscr{F}_2$ for Π_2 with the same base as the compact flow $\mathscr{F}_{i+1}$ on Π_2'. Define $\mathscr{F}_i$ over Π_2 as $\mathscr{F}_2$. Finally, define $\mathscr{F}_i$ on λ as for $\mathscr{F}_{i+1}$.

In a similar way, we treat the cases where C has the form $A \wedge B$ or $\neg A$, by applying Proposition 4.19 and Proposition 4.16 respectively. Define the flow on λ in Π_i as for λ in Π_{i+1}.

Reductions of the second kind simply consist in re-establishing a cut from a disjunction. Notice that the flow $\mathscr{F}_i$ is defined as $\mathscr{F}_{i+1}$ where the path through the cut-formula A is deformed in the obvious way. Notice that compactness holds for $\mathscr{F}_i$ because we asked the flow $\mathscr{F}_{i+1}$ to be either full of empty over the disjunction. This means that

the variants connected to one of the disjuncts in Π_{i+1} are active if and only if the variants connected to the other disjunct are active.

This concludes the proof. □

Next we want to discuss a more *local* version of Theorem 4.27. The inversion property given there does have the nice feature of going *inside* proofs, through the notion of compact flows, but the inner proofs that it treats are still *global* in a certain way. To have a more local result we need to be able to work with compact subgraphs with respect to any subproof. This is the content of Theorem 4.34. It shows that for *any* compact subgraph of the logical flow graph in Π' (lying in some subproof of Π') that is a graph for some proof there is a subgraph of Π that is a graph for a proof of the same sequent. In particular, this means that for any subproof of the cut-free proof Π' (possibly containing trivial cuts) associated to Π, we can always find an inner proof of the same sequent in Π.

4.34. Theorem (Locality of the inversion property). *Let the proofs* $\Pi : \Gamma \to \Delta$ *and* $\Pi' : A_1 \vee \neg A_1, \ldots, A_n \vee \neg A_n, \Gamma \to \Delta$ *be as in Theorem* 4.27. *For all subproofs* $\Pi'_* : A_{i_1} \vee \neg A_{i_1}, \ldots, A_{i_k} \vee \neg A_{i_k}, \Gamma'_* \to \Delta'_*$ *of* Π' *there is a subproof* $\Pi_* : \Gamma_* \to \Delta_*$ *of* Π *such that for all compact flows* $\mathscr{F}'_*$ *of* Π'_* *with* $H_{\mathscr{F}'_*}(A_{i_j} \vee \neg A_{i_j})$ *being either* $A_{i_1} \vee \neg A_{i_1}$ *or* $\emptyset$ *for all* $j = 1, \ldots, k$, *there is a compact flow* $\mathscr{F}_*$ *of* Π_* *such that* $H_{\mathscr{F}'_*}(\Gamma'_* \to \Delta'_*) = H_{\mathscr{F}_*}(\Gamma_* \to \Delta_*)$.

Proof. (*sketch*). For any subproof $\Pi'_* : A_{i_1} \vee \neg A_{i_1}, \ldots, A_{i_k} \vee \neg A_{i_k}, \Gamma'_* \to \Delta'_*$ of Π' we find a subproof $\Pi_* : \Gamma_* \to \Delta_*$ of Π and a flow $\mathscr{F}_*$ for Π_* such that $H_{\mathscr{F}_*}(\Gamma_* \to \Delta_*) = \Gamma'_* \to \Delta'_*$. Intuitively the flow $\mathscr{F}_*$ (induced by the procedure) can be thought as a 'maximal' flow in Π satisfying $H_{\mathscr{F}_*}(\Gamma_* \to \Delta_*) = \Gamma'_* \to \Delta'_*$, i.e. any compact flow for Π'_* of some *inner* sequent of $\Gamma'_* \to \Delta'_*$ will evolve (through the steps of the procedure) into a subgraph of $\mathscr{F}_*$ that is a compact flow for the same inner sequent of $H_{\mathscr{F}_*}(\Gamma_* \to \Delta_*)$.

To find the subproof Π_*, consider a subproof $\Pi'_* : \Gamma'_* \to \Delta'_*$ of Π' and let $\mathscr{F}'_*$ be the logical flow graph of it restricted to atomic formulas only. By definition $\mathscr{F}'_*$ is a compact flow (note that all formulas are active). The procedure described for Theorem 4.27 assures that all steps of cut elimination (from a proof Π_{i+1} to a proof Π_i) preserve the image $H_{\mathscr{F}'_*}(\Gamma'_* \to \Delta'_*)$ of the end-sequent of the subproofs in Π_{i+1} and Π_i where the flow passes through. A detailed proof would be a replica of the proof of Theorem 4.27. □

An immediate consequence of Theorem 4.34 is the following:

4.35. Corollary (Inverse image of subproofs). *Let the proofs* $\Pi : \Gamma \to \Delta$ *and* $\Pi' : A_1 \vee \neg A_1, \ldots, A_n \vee \neg A_n, \Gamma \to \Delta$ *be as in Theorem* 4.27. *For all subproofs* $\Pi'_* : A_{i_1} \vee \neg A_{i_1}, \ldots, A_{i_k} \vee \neg A_{i_k}, \Gamma'_* \to \Delta'_*$ *of* Π' *there is a subproof* $\Pi_* : \Gamma_* \to \Delta_*$ *of* Π *for which there is a compact flow* $\mathscr{F}_*$ *of* Π_* *such that* $H_{\mathscr{F}_*}(\Gamma_* \to \Delta_*) = \Gamma'_* \to \Delta'_*$.

5. On the complexity of interpolants

The construction of the interpolant proposed in the proof of Craig's Interpolation Theorem (Theorem 3.1) points out a correspondence between the size of the interpolant for a tautology $A \to B$ and the complexity of a cut-free proof of it (Corollary 3.6). For the first order predicate calculus with equality, it has been proved by Meyer [14] that there is no general recursive bound on the length of the smallest interpolant. It is an open question whether or not the size of the smallest interpolant in the propositional case can be polynomially bounded by the size of the tautology. A positive answer to this open question would have a consequence in complexity theory, namely that NP $\cap$ co-NP $\subset$ P/poly. To see this consider a language $L \subset \{0,1\}^*$ and $L \in$ NP $\cap$ co-NP (therefore the complement $\bar{L}$ of L is in NP). It is well-known that there are sequences of propositional formulas $A_n(p_1,\ldots,p_n,q_1,\ldots,q_m), B_n(p_1,\ldots,p_n,r_1,\ldots,r_s)$ (where $p_1,\ldots,p_n,q_1,\ldots,q_m,r_1,\ldots,r_s$ are the only propositional variables occurring within A_n, B_n) such that $|A_n|, |B_n| \leqslant n^{\mathcal{O}(1)}$ and

$$L = \{(p_1,\ldots,p_n) \in \{0,1\}^n \mid \exists y_1 \ldots y_m A_n(p_1,\ldots,p_n,y_1,\ldots,y_m) \text{ holds}\}$$

$$\bar{L} = \{(p_1,\ldots,p_n) \in \{0,1\}^n \mid \exists z_1 \ldots z_s B_n(p_1,\ldots,p_n,z_1,\ldots,z_s) \text{ holds}\}$$

Clearly $A_n \to \neg B_n$ is a tautology (for all n). By hypothesis there is an interpolant $I_n(p_1,\ldots,p_n)$ for $A_n \to \neg B_n$ such that $|I_n| \leqslant f(|A_n| + |B_n|)$ (for all n) where f is some function computable in polynomial time. Therefore $L \in$ P/poly. This observation has been made also in [16].

The results we will present, suggest that the complexity of an interpolant for $A \to B$ depends on the *logical relations* of the s-formulas in A and B. Certain tautologies will turn out to be 'easier' to prove than others and, they will have an interpolant of low complexity (for a discussion concerning related topics and motivations see [20]).

5.1. Splitting a proof

As suggested by Theorem 4.27, to find new results on the complexity of interpolation we will look at *compact subgraphs* of the logical flow graph of a proof Π. Without loss of generality we assume the logical flow graph of Π to be compact, i.e. all s-formulas in Π are connected to the end-sequent. In fact we have the following (recall the notion of *bridge* defined just before Proposition 2.1):

5.1. Lemma. *Let $\Pi : A \to B$ be a proof. There is a proof $\Pi' : A \to B$ whose logical flow graph is compact and # lines(Π') $\leqslant$ # lines(Π). If Π does not have any bridge between A and B, Π' does not have any bridge between A and B either.*

Proof. For the argument that follows it will be convenient to use the extra propositional symbol $\bot$ and the corresponding extension of *PLK* (where the axiom $\bot \to$ is added), but there will not be any occurrence of $\bot$ in the proof Π' that we produce. We

substitute all occurrences of atomic formulas in Π which are not lying in the compact subgraph of base $A \to B$ with $\bot$. Call this proof Π^*. We have that axioms $C, \Gamma \to \Delta, C$ will be replaced by the sequent $\bot, \Gamma^* \to \Delta^*, \bot$ (with Γ^*, Δ^* obtained by substituting $\bot$ to those occurrences in Γ, Δ which are not connected to the end-sequent) which we can think to be the axiom $\bot, \Gamma^* \to \Delta^*$ where the right occurrence $\bot$ is added as a weak formula. It is clear then, that all positive occurrences of $\bot$ which appear in cut-formulas of Π^* are directly linked to a weak occurrence in some axiom. Applying Lemma 3.11 we transform Π^* into a proof Π' with no $\bot$ occurrences in the cut-formulas of the proof (notice that the proof of Lemma 3.11 does not require axioms in Π^* to be logical axioms, i.e. axioms of the form $C, \Gamma \to \Delta, C$). Lemma 3.11 is indirectly removing all the distinguished axiomatic occurrences of $\bot$ from the proof and since there are no $\bot$ occurrences in the end-sequent of Π^* the proof Π' satisfies the claim. □

Call *A-axioms* (*B-axioms*) those axioms in $\Pi : A \to B$ whose distinguished formulas are connected to variants in A (B) but not to variants in B (A). Notice that certain axioms might be neither A-axioms nor B-axioms since they might be connected to both A and B; if so, we refer to them as *AB-axioms*. We might also have axioms whose distinguished formulas are not connected to the end-sequent, but they can be removed as in the lemma.

5.2. Corollary. *Let $\Pi : A \to B$ be a proof. There is a proof $\Pi' : A \to B$ whose axioms are either A-axioms or B-axioms or AB-axioms, and # lines(Π') $\leqslant$ # lines(Π). If there are no bridges between A and B in Π, then axioms in Π' are either A-axioms or B-axioms.*

5.3. Theorem. *Let $\Pi : A \to B$ be a proof (possibly with cuts) with no bridges between A and B. Suppose that all s-formulas occurring in weak formulas of Π and lying in the compact subgraph of base $A \to$ ($\to B$), occur in A-axioms (B-axioms). Then there exists a proof $\Pi' : A \to (\Pi' : \to B)$ such that # lines(Π') $\leqslant$ # lines(Π).*

A compact subgraph is uniquely determined by its base (as discussed just before Proposition 4.23). If there are no bridges between A and B the compact subgraph of base $A \to$ and the compact subgraph of base $\to B$ always exist. The theorem says that under certain conditions, these compact subgraphs are graphs of proofs. In other words, whenever the languages of A and B are disjoint and weak formulas are properly distributed in Π, it says that we can find either a proof of $A \to$ or a proof of $\to B$ with complexity bounded by the complexity of the original proof. This bound cannot be obtained from the construction used to prove the Craig Interpolation Theorem because the proof Π may contain cuts and the bound induced by the theorem would be in general exponential in the number of lines of the original proof (since it uses the Cut Elimination Theorem).

Proof of Theorem 5.3. By Proposition 4.23 the compact subgraph of base $A\rightarrow\ (\rightarrow B)$ is a compact flow. Apply Proposition 4.6 and find a proof $\Pi' : A\rightarrow\ (\Pi' : \rightarrow B)$ with the desired bounds on the complexity of the proof Π'. □

5.4. Corollary. *Let $\Pi : A\rightarrow B$ (possibly with cuts) be a proof with no bridges between A and B. If all weak occurrences in Π have either direct paths to A or to B and they are used as auxiliary formulas only by unary rules, then there exists a proof $\Pi' : A\rightarrow$ or a proof $\Pi' : \rightarrow B$ such that # lines$(\Pi')\leqslant$# lines(Π).*

Proof. (*sketch*). Given any proof Π where all weak formulas have direct paths to A or to B, one can always rearrange weak formulas properly in order to find either a proof whose weak occurrences directly linked to A lie in A-axioms or a proof whose weak occurrences directly linked to B lie in B-axioms. Such a proof has the same number of lines as Π. (It is crucial here that weak occurrences are not used as auxiliary formulas of binary rules.) By applying Theorem 5.3 we derive the claim. □

A weak formula C in $\Pi : A\rightarrow B$ is called *A-weak* (*B-weak*) if each s-formula in C has at least one connected variant in A (B).

5.5. Theorem. *Let*

$$A(p_1,\ldots,p_n,q_1,\ldots,q_m)\rightarrow B(p_1,\ldots,p_n,r_1,\ldots,r_s)$$

be a tautology where $p_1,\ldots,p_n$ are the only variables common to A and B. Suppose there is a proof (possibly containing cuts) of k lines for the tautology and suppose that A-weak (B-weak) formulas occur only in A-axioms (B-axioms). For any assignment σ of $\top,\bot$ to the variables p_i's, the sequent $A(\sigma(p_1),\ldots,\sigma(p_n),q_1,\ldots,q_m)\rightarrow$ $(\rightarrow B(\sigma(p_1),\ldots,\sigma(p_n),r_1,\ldots,r_s))$ is provable with a proof of at most $k+2|A|+1$ $(k+2|B|+1)$ lines.

In essence the theorem considers those proofs where A-weak (B-weak) formulas can only lie in axioms of the form $q_i,\Gamma\rightarrow\Delta,q_i$ $(r_j,\Theta\rightarrow\Lambda,r_j)$.

Proof of Theorem 5.5. Let Π be a proof of k lines for the tautology. Substitute all occurrences p_i in Π with $\sigma(p_i)\in\{\top,\bot\}$ and call the new proof Π'. We have that axioms $p_i,\Gamma\rightarrow\Delta,p_i$ will be replaced either by $\top,\Gamma'\rightarrow\Delta',\top$ or by $\bot,\Gamma'\rightarrow\Delta',\bot$ where we may think of $\top,\Gamma'\rightarrow\Delta',\top$ $(\bot,\Gamma'\rightarrow\Delta',\bot)$ as the axiom $\Gamma'\rightarrow\Delta',\top$ $(\bot,\Gamma'\rightarrow\Delta')$ where the left (right) occurrence $\top$ ($\bot$) is added as weak occurrence. (The multisets Γ',Δ' are obtained from Γ,Δ by substituting variables p_i with $\sigma(p_i)$ in them.) We will think of axioms $\top,\Gamma'\rightarrow\Delta',\top$ and $\bot,\Gamma'\rightarrow\Delta',\bot$ in this way *only* when their negative occurrence $\top$ and their positive occurrence $\bot$ are *directly* linked to the end-sequent.

We want to 'forget' all direct paths from positive occurrences of $\neg^k\bot$ (if $k=0$ or k is even), negative occurrences of $\neg^k\bot$ (if k is odd), negative occurrences of $\neg^k\top$ (if $k=0$ or k is even) and positive occurrences of $\neg^k\top$ (if k is odd) in $A\rightarrow B$. (We assume here that k is taken as large as possible for the given occurrence.) They will

clearly be *weak* occurrences because of our interpretation of the axioms. We describe the transformation of the proof Π' for the case of $\neg^k\bot$ occurrences only. For $\neg^k\top$ the treatment is similar.

If a positive occurrence $\neg^k\bot$ in $A \to B$ has paths to auxiliary formulas of some $\vee$:*right* (if $k = 0$ or k is even) or $\wedge$:*left* (if k is odd), then we erase all variants connected to it which lie in its direct paths. (The proof structure should be reshaped in order to eliminate extra $\neg$:*left*, $\neg$:*right*, $\vee$:*right*, $\wedge$:*left* rules but the details are obvious.) If a positive occurrence $\neg^k\bot$ in $A \to B$ has paths to auxiliary formulas of some $\vee$:*left* (if k is odd) or $\wedge$:*right* (if $k = 0$ or k is even), we should be a bit more careful. If the rule is an $\wedge$:*right* for instance, and it is applied to $\Pi_1 : \Gamma_1 \to \Delta_1, \neg^k\bot$ and $\Pi_2 : \Gamma_2 \to \Delta_2, C$ then we substitute it with a cut on $\Pi_1' : \Gamma_1 \to \Delta_1, \bot$ and $\Pi_2' : \bot, \Gamma_2 \to \Delta_2, C$, where Π_1' is obtained from Π_1 by erasing the rules of $\neg$:*right*, $\neg$:*left* applied to $\bot$, and Π_2' is obtained by adding $\bot$ as weak occurrence to Π_2. (Note that we are not removing the paths from $\neg^k\bot$ in this case, but we simply 'redirect' them into weak occurrences. Again we should modify the earlier proof below the place where the binary rule was used to account for this redirecting.) It is important to notice here that the only weak occurrence we add is a $\bot$ occurrence. For a $\vee : left$ rule with k odd we also add only a $\bot$ on the right side of the sequent as a weak occurrence.

Call $A^\circ \to B^\circ$ the resulting sequent, and Π° the proof. It is easy to check that we can always prove $A \to A^\circ$ and $B^\circ \to B$ in at most $2|A|$ and $2|B|$ lines. (The argument is similar to Lemma 5.8 below.) Therefore to derive our statement, we only need to show that we can 'split' Π° and find a proof of $A^\circ \to\ (\to B^\circ)$ with complexity bounded by k.

We would like to apply Theorem 5.3 to $\Pi^\circ : A^\circ \to B^\circ$, or rather an obvious extension of Theorem 5.3 to the language containing $\top, \bot$. The first point to observe is that Π° contains no bridges between A° and B°. This is an immediate consequence of our construction of Π° (which replaces every bridge with a path that ends in some weak occurrence). For Theorem 5.3 we would also like to know that any s-formula which occurs as a weak formula in an axiom and which lies in the compact subgraph with base $A^\circ \to$ (respectively $\to B^\circ$) actually lies in an A°-axiom (B°-axiom). Our hypotheses ensure that this is the case for the weak formulas from the original proof Π, but in the preceding construction we added some new weak occurrences. One can check that all of these additional weak occurrences are either positive occurrences of $\top$ or negative occurrences of $\bot$. These occurrences may not be linked correctly to A°-axioms (B°-axioms), but when they are not correctly linked it is easy to adapt to the situation using the fact that $\to \top$ and $\bot \to$ can be used as axioms. With this observation it is easy to extend Theorem 5.3 to the present situation, following the same argument as before. □

Theorems 5.3 and 5.5 consider proofs Π with cuts. If the proof Π is cut-free we can show stronger statements. In fact if there are no bridges between A and B, we can always find (by simply applying Lemma 3.2) either a proof $\Pi' : A \to$ or a proof $\Pi' : \to B$ such that $\#\ lines(\Pi') \leqslant \#\ lines(\Pi)$.

5.2. *Interpolants of linear size*

In this section we show that if A (respectively B) has a 'relatively simple construction' in $\Pi : A \to B$, then the size of the interpolant for $A \to B$ is linearly bounded by the size of A. More precisely, we show that whenever A (respectively, B) does not contain pairs of subformulas linked by a bridge one to the other, then there is an interpolant with size *linearly* bounded by the size of A (respectively, B).

Our claim is precisely stated as follows:

5.6. Theorem.[7] *Let $\Pi : A \to B$ be a proof (possibly with cuts); suppose there are no bridges from A back to A in Π (from B back to B, respectively). If there is a bridge between A and B, then there is an interpolant C for $A \to B$ such that*

(i) $|C| \leqslant 4|A|$ ($|C| \leqslant 4|B|$, *respectively*), *and*

(ii) $\#lines(\Pi^A) \leqslant 5|A|$ *and* $\#lines(\Pi^B) \leqslant 3 \cdot \#lines(\Pi)$, *where* $\Pi^A : A \to C$ *and* $\Pi^B : C \to B$ (*and symmetrically in the other case*).

If there is no bridge between A and B, then either there is a proof $\Pi' : A \to$ or a proof $\Pi' : \to B$ such that $\#lines(\Pi') \leqslant \#lines(\Pi)$.

To prove Theorem 5.6 we need some more definitions and considerations about logical flow graphs. We will extend the first order language L with atomic formulas $\top$, $\bot$ and the system *PLK* with new axioms $\Theta \to \Lambda, \top$ and $\bot, \Theta \to \Lambda$. We call the new language and the new system L^+ and PLK^+, respectively.

5.7. Definition. Let B be a formula in L. A *variation* B^* of B is a formula in L^+ obtained by substituting in B none, one or more of its s-formulas with $\top$ or $\bot$ as follows:

(1) if C is an s-formula that occurs *positively* in B then C is replaced by $\top$;

(2) if C is an s-formula that occurs *negatively* in B then C is replaced by $\bot$.

Clearly, there are more than one variation associated to a given formula B. Given a variation B^*, we denote $^*B$ the formula obtained from B^* replacing *all* occurrences of $\top, \bot$ with $\bot, \top$ respectively. For instance, if $\top \wedge B$ is the variation $(A \wedge B)^*$ then $\bot \wedge B$ will be denoted $^*(A \wedge B)$.

Notice that a variation $(\neg B)^*$ of $\neg B$ can also be written as $\neg\ ^*B$ and the formula $^*(\neg B)$ as $\neg B^*$, for some variation B^*.

5.8. Lemma. *Let A be a formula in L. For all variations A^* of A we have that the sequents $A \to A^*$ and $^*A \to A$ are PLK^+-provable with a proof Π such that $\#lines(\Pi) \leqslant 2|A|$.*

[7] The result holds for the full predicate calculus *LK* as well (as proved in [3]). The proof is the same but notice that the complexity of the interpolant C in (i) should be modified to be $|C| \leqslant (2k+5)|A|$ ($|C| \leqslant (2k+5)|B|$, respectively), where k is the arity of some predicate constant common to A and B.

Proof. By induction on the size of A.

$|A| = 1$: If A^* is A then the claim is obviously satisfied since $A \to A$ is an axiom of PLK; if A^* is $\top$ then $^*A$ is $\bot$ and the claim holds because $A \to \top$ and $\bot \to A$ are axioms for PLK^+;

$|A| > 1$: We will discuss in detail only the case in which A is of the form $\neg B$. The other cases can be handled similarly.

If A is of the form $\neg B$ and $(\neg B)^*$ is an arbitrary variation of it, we want to show that $\neg B \to (\neg B)^*$ and $^*(\neg B) \to \neg B$ are PLK^+-provable sequents.

By induction we know that $B \to B^*$ and $^*B \to B$ are PLK^+-provable sequents, for all variations B^* of B.

Then, for all pair of variations B^* it is easy to derive the sequent $\neg B \to \neg {}^*B$ from $^*B \to B$, and the sequent $\neg B^* \to \neg B$ from $B \to B^*$. Hence, the sequents $\neg B \to (\neg B)^*$ and $^*(\neg B) \to \neg B$ are PLK^+-provable for all variations $(\neg B)^*$.

The bound on the number of lines of a proof of $A \to A^*$ or $^*A \to A$ is clear. □

5.9. Lemma. *Let A be a formula in L and A^* the variation of it made up of $\top, \bot$ and logical symbols only. Then the sequents $\to A^*$ and $^*A \to$ are PLK^+-provable in $|A|$ steps.*

Proof. By an easy induction on the complexity of A. □

Observe that in a sequent $A_1 \ldots A_k \to B_1, \ldots, B_l$ all formulas A_i (for $i = 1, \ldots, k$) appear negatively in the sequent. This justifies the statement of the following lemma (with respect to Definition 5.7):

5.10. Lemma. *Let $\Pi : S$ be a PLK^+-proof and let A be an s-formula in L (i.e. A is neither $\top$ or $\bot$). Then either there is a bridge between A and a variant of it in S or there is a PLK^+-proof $\Pi' : S'$ where the sequent S' is obtained by replacing A with $\top$ ($\bot$) and $\#lines(\Pi') \leqslant \#lines(\Pi)$.*

Proof. Let A be an s-formula occurring in S and suppose there is no bridge to S starting from A.

Transform $\Pi : S$ into $\Pi' : S'$ by replacing A and all its variants (linked to A by the logical flow graph) by $\top$ ($\bot$). The logical flow graph of Π' is essentially the same as the logical flow graph of Π; the only difference concerns the sequence of edges associated to A (in Π), that in Π' is associated to $\top$ ($\bot$) (because of the renaming).

We have to show that Π' *is* a proof. We discuss only a few crucial cases; the others are treated similarly.

First, where Π has an axiom $A, \Gamma \to \Delta, A$, Π' might contain $\top, \Gamma' \to \Delta', \top$ ($\bot$, $\Gamma' \to \Delta', \bot$), where both A's are replaced by $\top$ ($\bot$) and Γ', Δ' are obtained from Γ, Δ by replacing some variants of A by $\top$ ($\bot$). The new sequent is clearly an axiom.

Second, where Π has a contraction, Π' might contain

$$\frac{\Gamma \to \Delta, C', C''}{\Gamma \to \Delta, C'''}$$

where C', C'', C''' are obtained from the formula C replacing some variant of A by $\top$ ($\bot$). By definition of logical flow graph, each s-formula positively (negatively) occurring in C''' has incoming (outgoing) edges from (to) the corresponding s-formulas occurring in C' and C''. This implies that C', C'', C''' are actually the same formula and the inference in Π' is a valid inference. □

5.11. Remark. The preceding argument is similar to the one used by Buss to show Proposition 6 in [2].

Proof of Theorem 5.6. Suppose that there is at least one bridge between A and B. This implies that there is at least one propositional variable common to A and B. Let R be such a propositional variable.

Apply Lemma 5.10 repeatedly to all s-formulas in A that do not have a bridge to some variant in $A \to B$ in Π so that the s-formulas occurring positively in A are replaced with $\top$ and the s-formulas occurring negatively in A are replaced with $\bot$. In this way one obtains the proof $\Pi' : A^* \to B$. Notice that by hypothesis A does not have bridges to itself, so all s-formulas in A^* are either linked to some variant in B or their form is $\top, \bot$. This means that all predicates occurring in A^* are common to A and B. By Lemma 5.8 we also know that $A \to A^*$ is provable; and by definition of variation, the size of A^* is the same as the size of A. Now we need just to transform A^* into a formula C of the original language L. By replacing $\top$ with $\neg R \vee R$ and $\bot$ with $\neg R \wedge R$, we can transform A^* into the formula C, i.e. the desired interpolant. Since $|A^*| = |A|$ then by the transformation we have that $|C| \leqslant 4|A|$. Moreover, $\#lines(\Pi^A) \leqslant 5|A|$ (note that axioms of the form $\Gamma \to \Delta, \top, \bot, \Gamma \to \Delta$ should be replaced by 3 line proofs and the bound follows from Lemma 5.8) and $\#lines(\Pi^B) \leqslant 3 \cdot \#lines(\Pi)$.

If there are no bridges between A and B then we notice that A^* is a formula made up of $\top, \bot$ and logical symbols only. In particular, all s-formulas in A^* of $\Pi' : A^* \to B$ are weak in Π'. We now erase all direct paths which start and end in A^* (which amounts to deleting weak occurrences, as in Lemma 4.9). This gives us a proof of $\to B$ in PLK^+. We want a proof in PLK, so that we need to delete the remaining occurrences of $\bot$ and $\top$. This can be accomplished using the obvious extension of Lemma 3.11 to PLK^+, as in Lemma 5.1. That is, the remaining occurrences of $\bot$ and $\top$ in the proof have to go through cuts, since they have been removed from the end-sequent. Each occurrence will be weak on one side of the cut, and so we can eliminate them using Lemma 3.11. This gives a proof Π^B with the correct bounds. □

5.12. Remark. Since a bridge links a negative to a positive occurrence, there are no possible bridges from a monotone (i.e. negation-free) formula back to itself. Therefore if either A or B is monotone then the tautology $A \rightarrow B$ has an interpolant C such that $|C| \leqslant \mathcal{O}(|A| + |B|)$.

5.13. Corollary. *Let $A \rightarrow B$ be a tautology of size $m > 5$. If the minimal interpolant of $A \rightarrow B$ has size at least m^2, then all proofs of $A \rightarrow B$ must contain bridges from A back to A, and bridges from B back to B.*

The converse of Corollary 5.13 does not hold, i.e. there are tautologies with small interpolants but for which bridges are necessary. Take for instance the tautology

$$y \wedge (y \leqslant \leftrightarrow (x \oplus y)) \rightarrow \neg z \vee (z \leqslant \leftrightarrow (x \oplus z))$$

where $\oplus$ denotes addition modulo 2 (i.e. $x \oplus y$ abbreviates $((x \wedge \neg y) \vee (\neg x \wedge y))$) and the symbol $x \leftrightarrow y$ abbreviates $\neg((x \wedge \neg y) \vee (\neg x \wedge y))$. The tautology has interpolant $\neg x$ and all proofs of it have bridges between the second and third occurrence of y, and the second and third occurrences of z. (If there were a proof in which the second and third occurrences of y were not logically linked, then we would still have a tautology after renaming the occurrences of y so that they were different. This would clearly not be a tautology.)

Nevertheless the corollary shows the necessity of bridges from A to A and from B to B to have complex interpolants, no matter the number of lines of the proof or the size of the tautology. This suggests that when investigating the complexity of the interpolant we should not consider standard quantities (like the number of lines or the size), but instead look at measurements of structure, such as the number of bridges from A to A and from B to B. This point is illustrated by the following result.

5.14. Proposition. *For any n, arbitrarily large, there exists a tautology $A_n \rightarrow B_n$ of size $\mathcal{O}(n)$ with smallest interpolant of size at least $n^2/400$. Any proof of $A_n \rightarrow B_n$ should contain at least n bridges from A_n back to A_n and at least n bridges from B_n back to B_n.*

The lower bound for the size of the smallest interpolant stated in Proposition 5.14 has been proved in [16] (as a consequence of Krapchenko's lower bound; see [12, 13]) by considering a sequent $A_n \rightarrow B_n$ with A_n of the form

$$y_n \wedge (y_2 \leftrightarrow (x_2 \oplus x_1)) \wedge (y_3 \leftrightarrow (x_3 \oplus y_2)) \wedge \ldots \wedge (y_n \leftrightarrow (x_n \oplus y_{n-1}))$$

and B_n of the form

$$\neg(\neg z_n \wedge (z_2 \leftrightarrow (x_2 \oplus x_1)) \wedge (z_3 \leftrightarrow (x_3 \oplus z_2)) \wedge \ldots \wedge (z_n \leftrightarrow (x_n \oplus z_{n-1})))$$

for n a natural number. For any proof of the tautology $A_n \rightarrow B_n$, the pairs of occurrences of the variables y_i and z_i in A_n and B_n, must be linked by bridges. For if this were not the case, we could rename those pairs of occurrences by different propositional letters, and the sentence so obtained would no longer be a tautology.

References

[1] R.B. Boppana and M. Sipser, The complexity of finite functions, in: Jan Van Leeuwen, ed., The Handbook of Theoretical Computer Science, Vol. A: Algorithms and Complexity (Elsevier/MIT Press, Amsterdam/Cambridge, 1990).

[2] S. Buss, The undecidability of k-provability, Ann. Pure Appl. Logic 53 (1991) 72–102.

[3] A. Carbone, On logical flow graphs, Ph.D. Dissertation, Graduate School, The City Univ. of New York, May 1993.

[4] A. Carbone, The Craig interpolation theorem for schematic systems, in: Collegium Logicum, Annals of the Kurt Gödel Society, Vol. 2 (Springer, Berlin, 1996).

[5] S.A. Cook and R. Reckhow, The relative efficiency of propositional proof systems, J. Symbolic Logic 44 (1979) 36–50.

[6] W.M. Farmer, A unification-theoretic method for investigating the k-provability problem, Ann. Pure Appl. Logic 51 (1991) 173–214.

[7] J.Y. Girard, Proof theory and logical complexity, in: Studies in Proof Theory, Monographs, Vol. 1 (Bibliopolis, Naples, 1987).

[8] J.Y. Girard, Linear logic, Theoret. Comput. Sci. 50 (1987) 1–102.

[9] M. Gromov, Cell division and hyperbolic geometry, Publication of IHES/M/90/54, Institut des Hautes Etudes Scientifiques, Bures-sur-Yvette, June 1990.

[10] S.C. Kleene, Introduction to Metamathematics, Bibliotheca Mathematica, Monographs on Pure and Applied Mathematics, Vol. 1 (Groningen, Walters-Noordhoff, 1988).

[11] J. Krajíček and P. Pudlák, The number of proof lines and the size of proofs in first order logic, Arch. Math. Logic 27 (1988) 69–84.

[12] V.M. Krapchenko, Complexity of the realization of a linear function in the class of Π-circuits, Mat. Zamet. 9 (1971) 35–40. English translation in Math. Notes Acad. Sci. USSR 10 (1971) 21–23.

[13] V.M. Krapchenko, A method of obtaining lower bounds for the complexity of Π-schemes, Mat. Zametki, 10 (1971) 83–92. English translation in Math. Notes Acad. Sci. USSR 11 (1972) 474–479.

[14] A.R. Meyer, A note on the length of Craig's interpolants, Technical report, Laboratory for Computer Science, M.I.T., October 1980.

[15] D. Mundici, NP and Craig's interpolation theorem. in: G. Longo, G. Lolli and A. Marcja, eds., Logic Colloquium '82 (Elsevier, Amsterdam, 1984).

[16] D. Mundici, A quadratic lower bound for the complexity of boolean interpolants. Technical report, Dept. of Computer Science, Univ. of Milan, 1993.

[17] R. Parikh, Some results on the length of proofs, Trans. Amer. Math. Soc. 177 (1973) 29–36.

[18] P. Pudlák, The length of proofs, in: S. Buss, ed., Handbook of Proof Theory (North-Holland, Amsterdam, 1995).

[19] G. Takeuti, Proof theory, in: Studies in Logic and the Foundations of Mathematics, Vol. 81 (North-Holland, Amsterdam, 1975).

[20] A. Urquhart, Complexity of proofs in classical propositional logic, in: Y.N. Moschovakis, ed., Logic from Computer Science (Springer, Berlin, 1992) 597–608.

[21] S. Wolfram, Cellular Automata and Complexity, collected papers (Addison-Wesley, Reading, MA, 1994).