The Cost of a Cycle Is a Square

A. Carbone

The Journal of Symbolic Logic, Vol. 67, No. 1. (Mar., 2002), pp. 35-60.

Stable URL:
http://links.jstor.org/sici?sici=0022-4812%28200203%2967%3A1%3C35%3ATCOACI%3E2.0.CO%3B2-3

The Journal of Symbolic Logic is currently published by Association for Symbolic Logic.

THE JOURNAL OF SYMBOLIC LOGIC
Volume 67, Number 1, March 2002

THE COST OF A CYCLE IS A SQUARE

A. CARBONE

Abstract. The logical flow graphs of sequent calculus proofs might contain oriented cycles. For the predicate calculus the elimination of cycles might be non-elementary and this was shown in [Car96]. For the propositional calculus, we prove that if a proof of k lines contains n cycles then there exists an acyclic proof with $\mathcal{O}(k^{n+1})$ lines. In particular, there is a polynomial time algorithm which eliminates cycles from a proof. These results are motivated by the search for general methods on proving lower bounds on proof size and by the design of more efficient heuristic algorithms for proof search.

§1. Introduction. A formal proof presents an underlying graph structure which is induced by the flow of its formula occurrences. By applying formal rules one after the other, one implicitly defines logical relations between the formula occurrences in a deduction, and by tracing the graph of these relations one usually obtains rather complicated graphs.

Different formal systems correspond to different properties of such graphs. The logical graphs of resolution proofs and cut-free proofs for instance, are essentially *trees* (or better say, collections of trees), while proofs with cuts might present instead a quite intricate 'topology'. In particular, they might contain cycles and at times nested loops. A study of the combinatorial properties of graphs of proofs with cuts is done in [Car97b].

In [Bus91], Buss introduces the idea of tracing formula occurrences in a proof and he calls *logical flow graph* the underlying graph structure of a proof. In [Bus93], Buss observes that logical flow graphs might contain oriented cycles. In [Car97] it was proved that cut-free proofs are cycle-free and it was sketched there that proofs containing cuts but no contractions are also cycle-free.

In this paper we will be concerned with *oriented* cycles but we should notice however that *non-oriented* cycles might appear in proofs and even in cut-free ones. Take for instance any cut-free proof of the pigeon hole principle PHP_n (i.e. the pigeon hole principle relative to n holes: if $n+1$ pigeons are placed in one of n holes, then there is a hole which contains at least two pigeons) written as a sequent $\Gamma_n \rightarrow \Delta_n$ where Γ_n is the collection of formulas

$$\Gamma_n = \{\bigvee_{j=1}^{n} p_{i,j} \ : \ i = 1, \dots, n+1\}$$

Received December 4, 1997; revised March 16, 2000.

and Δ_n is the collection of formulas

$$\Delta_n = \{p_{i,j} \wedge p_{m,j} \ : \ 1 \leq i < m \leq n+1, 1 \leq j \leq n\}$$

The $p_{i,j}$'s are propositional variables representing the fact that the i-th pigeon is placed in the j-th hole. The number of sequents in any cut-free proof of PHP_n is exponential in n (this was first proved by Haken in [Hak85] for resolution proofs and the result can be translated in the formalism of the cut-free sequent calculus). Since the size of the sequent $\Gamma_n \rightarrow \Delta_n$ is $\mathcal{O}(n^3)$, the number of axioms in the proof should be exponential in n and furthermore one can see that there are two atomic occurrences in the sequent PHP_n which have to be linked by an exponential number of logical paths passing through an exponential number of axioms. (A formal proof of this fact is given in Proposition 13.) By Haken's lower bound these logical paths cannot be eliminated and indeed they form an exponential number of unoriented cycles.

Oriented cycles can always be eliminated in propositional logic by applying the cut-elimination procedure to the proof but the price to pay for this reduction is (at times) an exponential expansion of the number of lines of the proof (this is a direct consequence of [Sta78]). In [Car96] it is shown that the elimination of cycles in predicate logic requires a non-elementary expansion. In fact, through cycles one can codify a repeated substitution of terms and find formulas containing large terms whose proofs need to contain cycles to be 'short' (i.e. to have a polynomial number of lines). If one is ready to slightly change the rules of the calculus for predicate logic, [Car97a] shows that proofs with cycles can be turned into cycle-free proofs with only an elementary expansion. Geometrically, cycles are turned into spirals.

For the propositional case the following question is open and furnishes a motivation for our analysis of cycles in proofs

> *Does the presence of cycles in proofs of propositional logic help to shorten the length of a proof?*

An equivalent formulation of the question is whether or not any cycle in a propositional proof (with axioms defined on atomic formulas) can be eliminated without substantially increasing the number of lines of the original proof. It is important to point out here that there are fundamental questions in complexity theory which are related to the length of proofs. In [CR79] it is shown that the existence of a proof system for propositional logic where classical tautologies can be proved by polynomial size proofs in the size of the tautology would imply $NP = coNP$. Since this equality looks unlikely to hold, one believes in the existence of tautologies with only exponential size proofs. For the sequent calculus with cuts the search for these tautologies seems a particularly hard task and any result on the structure of proofs might be of help to better understand its combinatorics and related complexity questions. (The reader can refer to [CS97] for a survey on this topics and references.)

Another motivation for our study is related to the desire of finding a common combinatorial language that expresses complexity phenomena present in both acyclic boolean circuits *and* proofs. Several analogies between proofs and circuits have been observed and for a survey we refer the reader to [Pud96]. One idea discussed there (pp. 612-618) and which was used to prove lower bounds for propositional proof systems, concerns *effective interpolation*. It is due to Krajíček [Kra97] and can be stated as follows: suppose we can show that the propositional implication $A \rightarrow B$ does not have simple interpolant I (i.e. there is no formula I that is written with symbols used both in A and B and such that $A \rightarrow I$ and $I \rightarrow B$ are provable), then it cannot have a simple proof. Several lower bounds results for propositional proof systems, which provide a *circuit* of polynomial size computing the interpolant, have been obtained. (See [Pud96] for references.)

Another result is stated in [Car97b]. It is shown that given any acyclic boolean circuit there is a propositional proof which simulates the computation of the same boolean function, and where the underlying graph of the proof *is* the graph of the circuit. The propositional proof might be required to contain cuts and its underlying graph is acyclic, as the circuit is.

A basic doubt regarding an effective correspondence between proofs and acyclic boolean circuits comes from the possibility of having *cyclic* underlying graphs for proofs. Therefore one can ask whether or not cyclic proofs are polynomially equivalent to acyclic ones

> *Given a proof Π (possibly containing cycles), is there an acyclic proof for the same sequent only polynomially larger than Π?*

This question was asked in [Bus95] and we show that we can transform a proof with n cycles into an acyclic one with an expansion bounded by a polynomial of degree $n + 1$. In particular, we can always eliminate a single cycle with a *quadratic* expansion of the size of the proof.

THEOREM 1. *Let $\Pi : S$ be a proof that contains n cycles. If Π has k lines then there is an acyclic proof $\Pi' : S$ of $\mathcal{O}(k^{n+1})$ lines.*

The plan of the paper goes as follows. In Section 2 we recall the sequent calculus *LK* and the notion of logical flow graph. In Section 3 we prove some basic facts on cycles in proofs. In Section 4 we introduce two notions which we will use to eliminate redundancies from the structure of proofs and to ensure some regularity conditions which will be of help in the proof of Theorem 1. In Section 5 we state two structural results on proofs and we use them to define a procedure that reduces the complexity of cut-formulas in a proof. This procedure is a subtle version of cut-elimination and based on it we show Theorem 1.

We thank an anonymous referee for his/her suggestions.

§2. Basic definitions and notations.

In this section we quickly recall known concepts. We present the rules of the *sequent calculus LK* [Gir87a, Tak87] for propositional logic and define what is a logical flow graph of a proof [Bus91].

2.1. The sequent calculus.

The system *LK* was introduced by Gentzen in 1934 [Gen34]. Here we present a propositional formulation of it restricted to the connectives $\wedge, \vee, \neg$. The connective $\supset$ will not be considered since it is easily definable from the $\neg, \vee$ symbols and the restriction will not affect the generality of our results as observed in Remark 8.

The *axioms* are sequents of the form

$$A, \Gamma \rightarrow \Delta, A$$

where A is an *atomic* formula and Γ, Δ are any collections of formulas. (In a collection we allow multiple occurrences of the same formula.) The formulas A are called *distinguished* occurrences and the formulas in Γ, Δ are called *weak* formulas. (Notice the notion of *weak occurrence* of a formula *in a proof* given in Section 2.2.)

Logical rules are used to introduce connectives, and they are given as follows:

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

$$\wedge : \textit{right} \quad \frac{\Gamma_1 \rightarrow \Delta_1, A \quad \Gamma_2 \rightarrow \Delta_2, B}{\Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2, A \wedge B}$$

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

$$\vee : \textit{left} \quad \frac{A, \Gamma_1 \rightarrow \Delta_1 \quad B, \Gamma_2 \rightarrow \Delta_2}{A \vee B, \Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2}$$

$$\vee : \textit{right} \quad \frac{\Gamma \rightarrow \Delta, A, B}{\Gamma \rightarrow \Delta, A \vee B}$$

The *structural rules* do not involve connectives and are the following:

$$\textit{Cut} \quad \frac{\Gamma_1 \rightarrow \Delta_1, A \quad A, \Gamma_2 \rightarrow \Delta_2}{\Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2}$$

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

The pair of formulas A in the upper sequent of the contraction rule are called *contraction formulas*. In axioms and rules, the formulas in Γ, Δ are called *side* formulas. In logical and structural rules, the formulas A (B) in the upper sequent(s) are referred to as *auxiliary* formulas and the one obtained by the application of a rule is called *main* formula. Notice that the cut-rule has no main formula.

In 1934 Gentzen ([Gen34]; see also [Gir87a, Tak87]) proved the following

THEOREM 2 (Cut Elimination Theorem). *Any proof in LK can be effectively transformed into a proof which never uses the cut rule. This works for both propositional and predicate logic.*

This is a fundamental result in logic and we will refer to it several times in the paper. Let us just mention here that the 'price' of cut elimination is that the cut-free proof may have to be much larger than proofs without cuts. There are propositional tautologies for which cut-free proofs must be exponentially larger than proofs with cuts, and in predicate logic the expansion can be non-elementary. See [Ore79, Ore93, Sta78, Tse68]. For an introduction to the combinatorics and complexity of cut elimination the reader can consult [CS97].

In the paper we will say that a *weak connective* is either an $\wedge$ occurring negatively or an $\vee$ occurring positively in a sequent. A *strong connective* is either an $\wedge$ occurring positively or an $\vee$ occurring negatively in a sequent.

The *size* $|A|$ of a formula A is its number of symbols. The size $|S|$ of a sequent S is the sum of the sizes of its formulas. The size $|\Pi|$ of a proof Π is the sum of the sizes of its sequents. The number of axioms and rules of inference in Π is denoted #lines(Π). The symbol $N(\Pi)$ denotes the number of axioms, cuts and logical rules (i.e. contractions are not counted) in Π. Clearly $N(\Pi) \leq$ #lines(Π), for all proofs Π.

The *logical complexity* $d(A)$ of a formula is its logical depth. Formally we say that $d(A) = 0$ when A is an atomic formula, $d(A \wedge B) = d(A \vee B) = max\{d(A), d(B)\} + 1$, and $d(\neg A) = d(A) + 1$.

2.2. The logical flow graph. To each proof in the sequent calculus we can associate a *logical flow graph* which is an oriented graph that traces the flow of occurrences within the proof. This concept was introduced by Buss [Bus91]. Properties of logical flow graphs are discussed in [Car97] where the notion is used to analyze the structure of classical proofs. A different but related graph was introduced earlier by Girard [Gir87] for linear logic proofs: *proof nets* represent the first tool used to analyze a proof as a "global" object.

We illustrate the notion of logical flow graph through an example. A formal definition together with an informal discussion of some of its properties follow (see also [Bus91] and [Car97]). Take the following proof

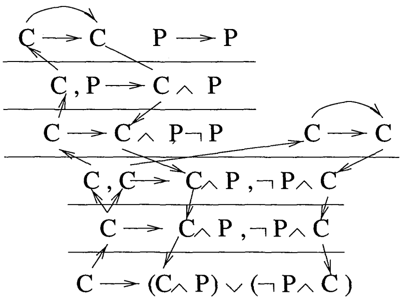

where the occurrences of the atomic formula C in the end-sequent are logically linked in the proof through oriented paths. For each application of a rule, we trace an edge between an occurrence of C in the upper sequent(s) and the corresponding occurrence of C in the lower sequent of the rule. (The correspondence between formula occurrences is induced by the rule in the obvious way.) For the axioms $C \rightarrow C$, we trace an edge from the antecedent to the consequent of the axiom. Similar links are present for the occurrences P and also for formulas which are logically more complicated, as $C \wedge P$, $\neg P$, $\neg P \wedge C$ and $(C \wedge P) \vee (\neg P \wedge C)$. For instance, we will say that there is a path from the occurrence $C \wedge P$ in the second line of the proof going down until the end-sequent, and similarly for the other subformulas. Notice that since our axioms are on atomic formulas, the only paths that will turn over axioms will be paths along atomic formulas. Notice also that paths can split apart at contractions and that this rule is the only one that allows the branching.

Let us now detail the formal definition of *logical flow graph*.

For each axiom, we trace an edge, called *axiom-edge*, from the distinguished occurrence in the antecedent of the axiom to the distinguished occurrence in its consequent. (Notice that axioms are defined over *atomic* formulas.)

For each rule, we trace an edge between any positive occurrence of a formula B in the upper sequent(s) of the rule and the corresponding occurrence of B in the lower sequent of the rule. Similarly, we trace an edge between any negative occurrence of a formula B in the lower sequent of the rule and the corresponding occurrence of B in the upper sequent(s) of the rule.

If the rule is a contraction, then there are two edges going to (coming from) the negative (positive) occurrences of B in the contraction formulas and coming from (going to) the relative occurrence of B in the main formula.

In the case of a cut rule applied to sequents $\Gamma_1 \rightarrow \Delta_1, A$ and $A, \Gamma_2 \rightarrow \Delta_2$, edges, called *cut-edges*, are traced between the two cut-formulas A as follows: for any subformula B of A occurring positively (negatively) in $\Gamma_1 \rightarrow \Delta_1, A$, there is an edge going from (coming to) it to (from) the correspondent occurrence of B in A of $A, \Gamma_2 \rightarrow \Delta_2$.

This concludes the definition of the edges of the graph. We say that the *logical flow graph* of a proof Π is the directed graph which we can read off the proof, whose nodes are labelled by the occurrences of formulas in Π, and whose edges are the links induced by the rules of Π, as defined above.

The logical flow graph carries a natural orientation, as discussed in [Car97, Car97a]. Negative occurrences in a sequent have edges going up in the proof and positive occurrences have edges going down. Negative occurrences are linked to positive ones through axioms, and positive occurrences are linked to negative through cuts, as in the example below

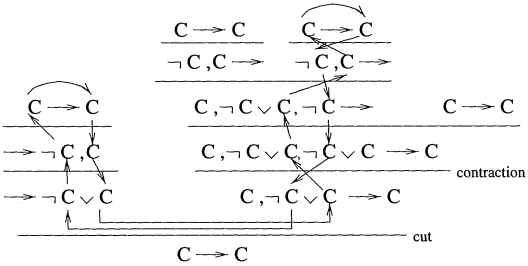

where the presence of a cut permits the path to turn up again. We linked only some of the occurrences of C to emphasize the presence of a closed sequence of edges in the logical flow graph. Any closed sequence of edges in a logical flow graph is called a *cycle*. When the nodes of the sequence have only one incoming edge and only one outcoming edge, the cycle is called *simple*. The path illustrated in the picture is a simple cycle. Notice that axioms in our proofs are defined on atomic formulas, and therefore, cyclic paths can only pass through *atomic* formulas.

Before concluding this section let us add some more terminology which will be useful later. A *bridge* is a path starting from a positive occurrence in the end-sequent S of a proof Π and ending in a negative occurrence in S. The starting and ending point of a bridge are called *extremes* of the bridge. Notice that the proof Π is not assumed to be cut-free and this allows a bridge to pass through cuts in Π.

We say that a logical path is *direct* when it links two formula occurrences without passing through cut-edges or axiom-edges. Direct paths in a logical flow graph go up towards axiom sequents, or go down towards either cut-formulas or formulas lying in the end-sequent. Because of the presence of contractions in proofs, there might be *several* direct paths linking a given formula occurrence to axioms in Π.

We say that a *weak occurrence* of a formula in a proof Π is a formula whose direct paths all go to or come from weak formulas in axioms of Π.

§3. **Cycles.** In this section we present a number of easy facts about *oriented* cycles in proofs. Since we are concerned only with oriented cycles, from now on we do not refer to the orientation anymore and we simply call them *cycles*.

We mentioned already that cycles can always be eliminated by applying the cut elimination procedure to a proof. Let us start by illustrating here how cyclic paths are split by the process of cut elimination. Suppose to have the following diagram

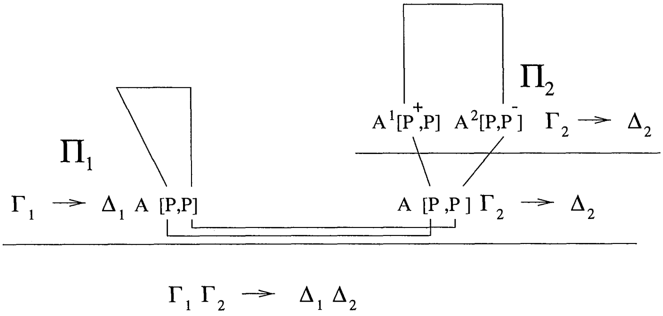

where P^+ and P^- are atomic occurrences of P and they appear positively and negatively inside the occurrences A^1 and A^2 respectively. For instance, A^1 and A^2 might be of the form $S \vee \neg S$, where P^+ and P^- correspond to the two different occurrences of S.

Imagine first that there is a path from P^+ to P^- passing through Π_2. Imagine also that there is a path starting at P^+ inside A^1 and going down through the contraction and the cut, up into Π_1, and then down again and back through the cut and contraction a second time to arrive at P^- inside A^2, as shown in the picture. After the transformation this cannot happen anymore.

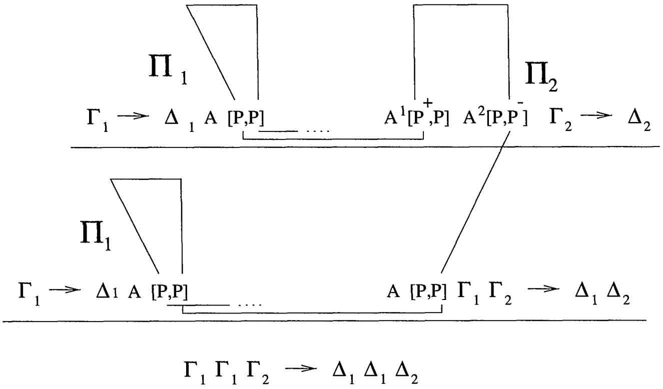

The path starting at P^+ inside A^1 and going into the (upper) copy of Π_1 will not have the opportunity to go back through A^2, it can only go back into A_1 through Π_2, as in the figure. In this way a simple cycle can be broken.

REMARK 3. The phenomenon we described explains how paths (not necessarily cyclic) split during the process of elimination of cuts. If the path from P^- to P^+ in Π_2 did not exist, we would not have a cycle in the proof but we would simply have a connection from P^+ to P^- through Π_1 that would break after duplication.

(See [Car97, Car96, Car97a] for more information about these phenomena, the formation of patterns in proofs and their role in the structure of formal proofs.)

The proof displayed in Section 2 contains one cycle and the experts certainly recognized already its correspondence with the construction of the Church numeral 2 in Lambda Calculus. If we look at the proofs corresponding to the Church numeral n, for any natural number n, we can observe that such proofs contain $n - 1$ cycles. For instance consider the proof corresponding to the Church numeral 3

$$
\begin{prooftree}
\AxiomC{$A \rightarrow A$}
\UnaryInfC{$\rightarrow \neg A, A$}
\UnaryInfC{$\rightarrow \neg A \vee A$}
\AxiomC{$A \rightarrow A$}
\UnaryInfC{$A, \neg A \rightarrow$}
\AxiomC{$A \rightarrow^1 A$}
\UnaryInfC{$A, \neg A \rightarrow$}
\BinaryInfC{$\neg A \vee A, \neg A, A \rightarrow$}
\AxiomC{$A \rightarrow^2 A$}
\UnaryInfC{$A, \neg A \rightarrow$}
\BinaryInfC{$\neg A \vee A, \neg A \vee A, \neg A, A \rightarrow$}
\RightLabel{\textit{contraction}}
\UnaryInfC{$\neg A \vee A, \neg A, A \rightarrow$}
\AxiomC{$A \rightarrow A$}
\BinaryInfC{$\neg A \vee A, \neg A \vee A, A \rightarrow A$}
\RightLabel{\textit{contraction}}
\UnaryInfC{$\neg A \vee A, A \rightarrow A$}
\RightLabel{\textit{cut}}
\BinaryInfC{$A \rightarrow A$}
\end{prooftree}
$$

This proof contains two simple cycles passing through the axioms labelled by 1 and 2.

Proposition 4. *For any natural number $n \geq 1$, there is a proof with only 1 cut rule and n simple cycles.*

Proof. The family of proofs corresponding to Church numerals n (for all $n > 0$) satisfies the condition. ⊣

The number of cycles that a proof can contain is bounded only by the number of its axioms and it does not depend on the number of cuts contained in the proof. In fact there is a known procedure (a folklore) which transforms a proof with many cuts into a proof with only one cut, with essentially the same number of lines. This procedure leaves unaltered the 'topological structure' of the logical flow graph of the proof. Namely, no new contraction rules are introduced in the proof by the transformation and no old ones are eliminated. (Insights and precise results on the topological structures of proofs can be found in [CS97a, Car97b].) For completeness we state here this fact.

Proposition 5. *Let $\Pi : S$ be a proof containing cuts. Then there is a proof $\Pi' : S$ containing only one cut and the same number of contractions as Π. Moreover, the logical flow graph of Π has the same topological structure of the logical flow graph of Π'.*

Proof. Suppose Π to contain m cuts on formulas $A_1, \dots, A_m$. To each pair of cut-formulas A_i in Π instead of applying a cut to it we apply a $\neg$:*left* followed by a $\vee$:*left* rule to derive $\neg A_i \vee A_i$. This will give a cut-free proof of the sequent $\neg A_1 \vee A_1, \dots, \neg A_m \vee A_m, \Gamma \rightarrow \Delta$ which has the same number of lines as Π. In m applications of the $\wedge$:*left* rule we can derive from this sequent

$$\neg A_1 \vee A_1 \wedge \dots \wedge \neg A_m \vee A_m, \Gamma \rightarrow \Delta$$

and in $\mathcal{O}(m)$ steps (without contractions) we can derive

$$\rightarrow \neg A_1 \vee A_1 \wedge \dots \wedge \neg A_m \vee A_m.$$

By combining the two sequents with a cut we derive $\Gamma \rightarrow \Delta$. The new proof has only one cut as required, it is of linear size in the length of Π and has the same number of contractions as the original proof.

It is clear from the construction that the logical flow graph of Π' has the same 'topological structure' of the logical flow graph of Π. The paths have been stretched but nothing dramatic has been done to the underlying graph (in particular no paths have been split). No new contractions have been introduced and all the old ones have been preserved. ⊣

In [Car97] it was proved that cut-free proofs do not contain cycles. It was also sketched there that the presence of cycles depends on the presence of contractions in the proof. Here we prove this fact as a consequence of the acyclicity of cut-free proofs. We say that a contraction *lies above* a cut when the contracted formula has a direct path to or from a cut-formula of the cut.

PROPOSITION 6. *Let* Π *be a proof with no contractions lying above its cuts. Then the logical flow graph of* Π *is acyclic.*

PROOF. We want to eliminate all cuts in Π except those where at least one of the cut-formulas is weak. These cuts are called *innocuous* since no cyclic path can pass through them. We shall prove this fact and derive the statement afterwards.

To show that cyclic paths do not pass through innocuous cuts, observe that paths passing through weak cut-formulas have to end-up, or have to come from, weak occurrences in axioms. But weak occurrences in axioms correspond to starting or ending points of paths, since axiom-edges are defined only on distinguished formulas. Therefore these paths cannot be cyclic.

To show the statement, we apply Gentzen's procedure of cut elimination. Since we do not have contractions and we want to eliminate only non-innocuous cuts, the procedure will simply consist on the reduction of the logical complexity of the cut-formulas, on the permutation of cuts and on the elimination of atomic cuts from axioms. In particular, we do not remove (since only innocuous cuts require the removal of subproofs) or duplicate any subproof, and we do not introduce any contraction formula. In this way, the shape of the proof will be altered but the 'topological structure' of the logical flow graph will remain unchanged. If Π contains a cyclic path, by the application of the procedure the cyclic path might be deformed in his shape but there would be no way it would disappear from the proof. This means that after cut-elimination we should obtain a proof which contains the cycle, but this is absurd since cyclic paths cannot pass through innocuous cuts and any cyclic path needs to pass through at least one cut (by the acyclicity of cut-free proofs). ⊣

The family of proofs that corresponds to Church numerals shows that there is no relation between the number of cycles in a proof and the size of the end-sequent. In particular, the formulas occurring in these cycles are logically linked with the formulas occurring in the end-sequent. Usually, this is not required. Consider, for instance, the following proof

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{A \to A}{\to \neg A, A} \quad B \to B, A}{B \to B, A \wedge \neg A, A} \quad B \to B, \neg A}{B, B \to B, B, A \wedge \neg A, A \wedge \neg A}}{B \to B, A \wedge \neg A} \qquad \dfrac{\dfrac{A \to A}{A, \neg A \to}}{A \wedge \neg A \to}}{B \to B}$$

where a cycle passes through occurrences of A, and where no A appears in the end-sequent. One could ask whether or not such cycles help to shorten a proof.

Remark 7. Proofs with *monotone* cut-formulas (i.e. they do not contain negations) cannot contain cycles. To see this is easy since a cycle will always link at least two distinguished occurrences lying in the same cut-formula (by Lemma 23). One of the occurrences will be positive and the other negative. But a monotone formula will contain only positive or only negative occurrences by definition.

Remark 8. Any proof Π formalized with the connective $\supset$ can be easily transformed into a proof Π' where occurrences of the form $A \supset B$ are replaced by formulas $\neg A \vee B$, and the rules $\supset$:*left* and $\supset$:*right* are substituted by the obvious combination of $\neg$:*left*, $\vee$:*left* and $\neg$:*right*, $\vee$:*right*, respectively. Since the transformation does not introduce new contractions and does not eliminate any axiom, cut, nor contraction, the topological structure of Π' remains the same as the topological structure of Π.

§4. Two notions to reduce a structure. All the examples of proofs in the previous section present a similar structure in the formation of cycles and they might seem not too convincing since one can easily imagine a proof of their end-sequent which is smaller in size and cycle-free. Indeed, these examples can be reduced easily to acyclic proofs by uniform procedures which uniformly act on the structure of proofs and eliminate trivial "redundancies". We will present here two structural notions and will describe how to transform a proof in such a way to avoid redundancies. These two notions furnish some regularity conditions for proofs that will be used to derive the bounds in Theorem 1.

The first notion has its origin in [Car97].

Definition 9. A proof Π is *reduced* when the following two requirements are satisfied:

1. (*local condition*) no binary rule is applied to a weak auxiliary formula in Π, no unary logical rule is applied to two weak auxiliary formulas and no contraction rule is applied to a weak auxiliary formula;
2. (*global condition*) no occurrence in a cut-formula of Π is weak.

We say that a proof is *locally reduced* when condition 1 is satisfied. Similarly we speak about *global reducibility* when condition 2 is required to hold.

In the example above Remark 7, the proof is *not* reduced because of the $\wedge$:*right* rule applied to the sequents $\to \neg A, A$ and $B \to B, A$ to introduce the formula $\neg A \wedge A$. The auxiliary formula A in $B \to B, A$ is weak and the local condition fails to hold. The following proof provides another example

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

where A in the conjunction of $A \wedge B \to A$ is a weak occurrence. Here, the global condition fails to hold.

Proposition 10. *Let $\Pi : S$ be a proof of k lines. Then there is an effective way to transform Π into a reduced proof $\Pi' : S$ which has at most k lines. If Π contains no cuts then the same is true for Π'.*

Proof. It is enough to give a procedure that first transforms a proof into a new one satisfying the local condition, and then transforms it into a new proof satisfying the global condition.

The first step was achieved in Lemma 3.2 (pp. 261) of [Car97] (see also Proposition 6.28 pp. 142 in [CS97a]), and the second step was shown in Lemma 3.11 (pp. 265) of [Car97].

Since the second step of transformation can give a proof which falsifies the local requirement, one needs to iterate the first and second step until both conditions are simultaneously satisfied. Such a moment will exist since both steps of transformation always reduce either the size (i.e. the number of symbols) or the number of lines of the proofs. ⊣

The proposition above says that a reduced proof eliminates insignificant structure from a proof. Some aspects of this are given in the next three results.

Proposition 11. *Let $\Pi : S$ be a reduced proof, let c be the number of contractions and a be the number of axioms in Π. Then $a \geq c/2$.*

Proof. In [Car96]. ⊣

Proposition 12. *Let $\Pi : S$ be a proof with k lines. Then there exists a reduced proof $\Pi' : S$ with at most k lines and whose cut-formulas have size $\leq 2 \cdot k$.*

Proof. In [Car97] (Lemma 3.12, pp. 267). There the statement is slightly different but its proof gives the bound we look for. (The proof is a direct consequence of the global condition in Definition 9.) ⊣

Proposition 13. *Let S be a sequent of size n and Π be a reduced cut-free proof of S of $2^{\mathcal{O}(n)}$ lines. Then there is a pair of atomic occurrences B^1, B^2 in S such that the logical flow graph of Π contains $2^{\mathcal{O}(n)}$ paths starting at B^1 and ending in B^2.*

Proof. We shall observe first a few properties that follow directly from the assumptions. Since S is of size n, then the number of distinct subformulas of S is $\leq n$. Also, by the subformula property each formula occurring in Π should appear in S, and since the number of sequents of Π is $2^{\mathcal{O}(n)}$, then there should be exponentially many main formulas in Π which are subformulas of S. This implies that these formulas are identified in Π by an exponential number of contractions, say c. By Proposition 11 the number of axioms a should be $a \geq c/2$, since the proof is reduced, and therefore the number of axioms in Π is exponentially large. Now, notice that there is a logical path which passes through each axiom, it comes from a negative occurrence in S and it goes to a positive occurrence in S. Since there are $\mathcal{O}(n^2)$ such pairs of occurrences in S, at least one of these pairs should be associated to an exponential number of paths. This proves the claim. ⊣

The second notion that we want to define has been introduced in [Car97]. It is formulated as follows

Definition 14. 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 $\Pi_1 : C, \Gamma_1 \rightarrow \Delta_1$ and $\Pi_2 : C, \Gamma_2 \rightarrow \Delta_2$ (respectively $\Pi_1 : \Gamma_1 \rightarrow \Delta_1, C$ and $\Pi_2 : \Gamma_2 \rightarrow \Delta_2, C$) 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.

The definition says 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. The proof displaying a cycle in Section 2.2 is *not* separated.

In [Car97] (Proposition 4.18, pp. 278) it is shown that any *propositional* proof can be transformed into a separated proof

Proposition 15. *Let Π be a proof in propositional logic. Suppose that Π has k lines. There is an effective method to transform Π into a proof Π' : S which is separated and has at most k lines.*

Roughly speaking, the transformation claimed in the proposition replaces a contraction on a formula C in Π with contractions on subformulas of C. This step turns out to be possible only for *propositional* proofs since the presence of quantifiers could induce to force contractions on pairs of subformulas of C which differ in their eigenvariables.

Other properties of separated proofs can be found in [Car97]. Here we only want to point out one more fact. Let S be a sequent and C be a formula in S. We say that two occurrences in C are *strongly linked* when the minimal subformula of C which contains both occurrences is either positive in S with $\wedge$ as main connective, or it is negative in S with $\vee$ as main connective. We speak of *weak link* when the occurrence of the minimal subformula is either positive in S with $\vee$ as main connective, or it is negative in S with $\wedge$ as main connective.

Lemma 16. *Let Π : S be a separated proof and C be any formula in S. Then there is no pair of occurrences B^1, B^2 in C which are strongly linked and connected by a logical path departing from B^1, going up to an axiom, turning over it and descending to B^2 through a direct path.*

Proof. The statement follows from the separation property since all pairs of paths from any two occurrences which are strongly linked should depart into different subproofs. If this is not the case, then the two occurrences should lie in some weak formula and therefore no logical path would connect them. ⊣

A reduced proof might be non-separated. Take for instance the proofs corresponding to Church numerals. They are clearly *not* separated. On the other hand it is trivial to build examples of proofs which are separated but not reduced. Take for instance any separated proof and apply a binary rule to a weak occurrence in it.

Looking back at the examples of proofs with cycles given in Sections 2.2 and 3, we observe that all these proofs are non-separated and that Proposition 15 is a *linear* procedure for the elimination of their cycles. There are proofs which are reduced, separated and contain cycles nevertheless. We provide here an example.

Let Π be a proof of the sequent $(C \vee D) \vee \neg A \to \neg A, C, D$, whose last rule is a cut over formulas $(\neg(A \vee \neg A) \vee C) \vee (\neg(A \vee \neg A) \vee D) \to C, D$ obtained by combining the subproofs Π_1, Π_2 defined as follows. Let the proof Π_1 be

$$\dfrac{\dfrac{\dfrac{A \to^* A}{\to A, \neg A}}{\to A \vee \neg A} \qquad \dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{A \to^* A \quad \dfrac{A \to A}{\neg A \to \neg A}}{A \vee \neg A \to A, \neg A}}{A \vee \neg A \to A \vee \neg A}}{A \vee \neg A, \neg(A \vee \neg A) \to} \quad C \to C}{A \vee \neg A, \neg(A \vee \neg A) \vee C \to C} \quad \dfrac{\dfrac{\dfrac{\dfrac{\dfrac{A \to A \quad \dfrac{A \to^* A}{\neg A \to \neg A}}{A \vee \neg A \to A, \neg A}}{A \vee \neg A \to A \vee \neg A}}{A \vee \neg A, \neg(A \vee \neg A) \to} \quad D \to D}{A \vee \neg A, \neg(A \vee \neg A) \vee D \to D}}{A \vee \neg A, A \vee \neg A, (\neg(A \vee \neg A) \vee C) \vee (\neg(A \vee \neg A) \vee D) \to C, D}}{A \vee \neg A, (\neg(A \vee \neg A) \vee C) \vee (\neg(A \vee \neg A) \vee D) \to C, D}\,\textit{contraction}}{(\neg(A \vee \neg A) \vee C) \vee (\neg(A \vee \neg A) \vee D) \to C, D}\,\textit{cut}$$

and the proof Π_2 be

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{A \to^* A}{\neg A, A \to} \quad \neg A \to \neg A}{\neg A, A \vee \neg A \to \neg A} \quad A \to A}{A \vee \neg A, A \vee \neg A \to A, \neg A}}{\neg A, A \vee \neg A, A \vee \neg A \to \neg A} \quad \dfrac{C \to C \quad D \to D}{C \vee D \to C, D}}{(C \vee D) \vee \neg A, A \vee \neg A, A \vee \neg A \to \neg A, C, D}}{(C \vee D) \vee \neg A \to \neg(A \vee \neg A), \neg(A \vee \neg A), \neg A, C, D}}{(C \vee D) \vee \neg A \to \neg(A \vee \neg A), \neg(A \vee \neg A) \vee C, \neg A, D}}{(C \vee D) \vee \neg A \to \neg A, \neg(A \vee \neg A) \vee C, \neg(A \vee \neg A) \vee D}}{(C \vee D) \vee \neg A \to \neg A, (\neg(A \vee \neg A) \vee C) \vee (\neg(A \vee \neg A) \vee D)}$$

The cycle goes along the path on A occurrences, passes through the axioms whose sequent arrow is marked by the symbol * and also through both cuts in the proof.

REMARK 17. The class of separated and reduced propositional proof structures is algorithmically very poor since it does not codify Church numerals inside it, but indeed general enough to formulate any complexity question. This is because the set of tautologies proved by reduced and separated proofs is the complete set of classical tautologies and for each proof there is a reduced and separated one with the same complexity (i.e. the number of lines does not augment).

The class of reduced proofs will be important in the next section, to show that a cycle can be eliminated from proofs in *quadratic* time. It is this regularity of the structure of proofs that allows to compute a quadratic bound. The condition of *separation* is not essential for Theorem 1, though to consider separated proofs will slim the details of the transformations.

§5. The elimination of cycles. We will prove here our main result (i.e. Theorem 1). We state first two simple lemmas which will be crucial to the proof of the theorem. Let us denote $A[B]$ for a formula A containing B as a subformula (the formula B might be A itself). By this notation we intend B to be a specific occurrence of the subformula B in A.

Lemma 18. *Let* $\Pi : \Gamma \to \Delta, B[A_1 \vee A_2]$ *be a proof where* $A_1 \vee A_2$ *occurs positively in the sequent (i.e. the disjunction is a weak connective). Then there is a separated proof* $\Pi' : \Gamma \to \Delta, B[A_1], B[A_2]$ *such that* $N(\Pi') \leq 2 \cdot \#\text{lines}(\Pi)$. *Moreover, in the logical flow graph of* Π' *there is a path in* Π' *linking two occurrences in* Γ, Δ *only if there is already a path linking them in* Π.

Proof. By Proposition 15 we can assume the proof Π to be separated and we want to transform Π into a proof Π' by steps. (The hypothesis of separation is not crucial but simplifies the argument.) We look first at all logical paths in Π which end-up into the occurrence $A_1 \vee A_2$ in the end-sequent (notice that these paths do not pass through cut-formulas since they should first pass through axioms and our axioms are on atomic formulas only).

These paths might start from some weak occurrence $B'[A_1 \vee A_2]$ in some axiom of Π which is logically linked to the subformula $B'[A_1 \vee A_2]$ in the end-sequent. In this case, we replace $B'[A_1 \vee A_2]$ in the axioms of Π with two weak occurrences $B'[A_1], B'[A_2]$ and we substitute the pair of formulas along the paths until we find a rule with auxiliary formula $B'[A_1 \vee A_2]$ (in Π). Before explaining what to do next, let us consider first those paths that start in some main formula $A_1 \vee A_2$ of a $\vee$:*right* rule. In this case we remove the application of the rule in Π' and substitute the pair of formulas A_1, A_2 along the paths until we find a rule with auxiliary formula $A_1 \vee A_2$ (in Π).

What to do next? Remember that we want to construct two formulas $B[A_1], B[A_2]$ instead of $B[A_1 \vee A_2]$.

Suppose that for any subproof $\Pi_* : \Gamma_* \to \Delta_*, B'[A_1 \vee A_2]$ of Π, where $B'[A_1 \vee A_2]$ is an auxiliary formula in Π_* which has a path to $B[A_1 \vee A_2]$, we have built a proof $\Pi'_* : \Gamma_* \to \Delta_*, B'[A_1], B'[A_2]$. Then for all *binary* rules applied to $\Pi_\circ : \Gamma_\circ \to \Delta_\circ, C$ and Π_* in Π, we have a pair of binary rules in Π' that are applied to two copies of $\Pi_\circ : \Gamma_\circ \to \Delta_\circ, C$ and the auxiliary formulas $B'[A_1], B'[A_2]$ in Π'_*, respectively. We obtain a subproof whose end-sequent contains *two* copies of $\Gamma_\circ, \Delta_\circ$. We contract these formulas pairwise and we do this just after the application of both the binary rules. (Notice that the side formulas Γ_*, Δ_* and $\Gamma_\circ, \Delta_\circ$ in Π'_* and $\Pi_\circ$ do not contain any occurrence directly linked to $A_1 \vee A_2$ in the end-sequent because of the hypothesis of separation. This implies that no subproof of $\Pi_\circ$ has been duplicated by a previous step of the construction.)

The *unary* rules in Π applied to the auxiliary formulas $B'[A_1 \vee A_2]$ and C in $\Pi_* : \Gamma_* \to \Delta_*, B'[A_1 \vee A_2]$ are handled similarly. Suppose that we have built a proof $\Pi'_* : \Gamma'_* \to \Delta'_*, B'[A_1], B'[A_2]$, then we apply two unary rules, one to $B'[A_1]$ and the other to $B'[A_2]$. Since unary rules have two premises, we need to add one occurrence C as a weak occurrence, to ensure the auxiliary formula for one of the unary rules. This addition does not enlarge the size of the proof (as proved in Lemma 4.8 pp. 274 in [Car97]). We combine the weak occurrence with $B'[A_2]$.

Any other rule in Π is applied in Π' with no manipulation. The treatment of subproofs $\Pi_* : B'[A_1 \vee A_2], \Gamma_* \to \Delta_*$ is analogous. This concludes the construction of Π'.

The bound on the number of lines of Π' follows from the hypothesis of separation. By construction the treatment of binary rules implies a duplication of some subproof $\Pi_\circ$ of Π. As observed above, subproofs $\Pi_\circ$ which are duplicated will not be duplicated again in further steps of the construction, because Π is separated. This implies that Π' contains at most $2 \cdot \#\text{lines}(\Pi)$ logical rules, axioms and cuts.

The property of the logical flow graph of Π' is easily proved by induction on the height of the subproofs of Π. (Notice that paths linking two occurrences in the end-sequent of a proof should pass through *atomic* occurrences.) $\dashv$

Notice that Lemma 18 holds only for *propositional* proofs. As a counterexample take the provable sequent $\rightarrow \forall x.A(x) \vee \neg A(x)$. Clearly, the sequent $\rightarrow \forall x.A(x), \forall x.\neg A(x)$ is not provable.

REMARK 19. The transformation described in the proof of Lemma 18 introduces *new* weak formulas and *new* contractions. This implies that in Π' we might have new paths linking occurrences in the end-sequent to weak occurrences in axioms, and also new bridges. New bridges can be classified in *four* different groups:
1. a bridge in Π whose extremes do *not* lie in $B[A_1 \vee A_2]$, duplicates in Π': both bridges pass through subproofs $\Pi_\circ$ of Π and their extremes are identified in Π' through contractions applied to pairs of formulas in $\Gamma_\circ, \Delta_\circ$ of Π'_*.
2. a bridge in Π with *one* of the extremes lying in $B[A_1 \vee A_2]$ (but not in $A_1 \vee A_2$) and the other extreme lying in Γ, Δ duplicates into a pair of bridges in Π' having a common extreme lying in Γ, Δ (the common extreme is obtained through contractions applied to pairs of formulas in $\Gamma_\circ, \Delta_\circ$ for some subproof $\Pi_\circ$ which is duplicated in Π') and the other extremes lying in $B[A_1], B[A_2]$, respectively.
3. a bridge in Π with *both* extremes lying in $B[A_1 \vee A_2]$ (but not in $A_1 \vee A_2$) duplicates into a pair of bridges in Π': the extremes of one of the bridges lie in $B[A_1]$ and the extremes of the other lie in $B[A_2]$.
4. a bridge in Π with *both* extremes lying in $B[A_1 \wedge A_2]$ (but not in $A_1 \wedge A_2$) duplicates into a pair of bridges in Π': one of the extremes is in common and lies in $B[A_1]$ (let us call D the atomic occurrence in $B[A_1]$ corresponding to this extreme), and the other pair of extremes lie in $B[A_1], B[A_2]$, respectively. Also, the atomic occurrence in $B[A_2]$ which corresponds to D, is weak.

It is routine to check that this *classification* covers all possible cases of duplication induced by Lemma 18, and that no bridge with an extreme in A_i is duplicated. We shall make use of these facts in the proof of Theorem 1.

LEMMA 20. *Let* $\Pi : \Gamma \rightarrow \Delta, B[A_1 \vee A_2]$ *be a proof where* $A_1 \vee A_2$ *occurs negatively in the sequent* (*i.e. the disjunction is a strong connective*). *Then there are separated proofs* $\Pi_1 : \Gamma \rightarrow \Delta, B[A_1]$ *and* $\Pi_2 : \Gamma \rightarrow \Delta, B[A_2]$ *such that* $\#\text{lines}(\Pi_1), \#\text{lines}(\Pi_2) \leq \#\text{lines}(\Pi)$. *Moreover, there is a path between two occurrences in* Π_1 (Π_2) *only if there is already a path linking them in* Π.

PROOF. We transform the proof Π into a separated proof Π^* whose number of lines is at most $\#\text{lines}(\Pi)$ (by Proposition 15). We want to build from Π^* a proof $\Pi_1 : \Gamma \rightarrow \Delta, B[A_1]$. A proof $\Pi_2 : \Gamma \rightarrow \Delta, B[A_2]$ can be built similarly.

We look first at all logical paths in Π^* which start from the occurrence $A_1 \vee A_2$ in the end-sequent (notice that these paths do not pass through cut-formulas as remarked in the proof of Lemma 18 also). These paths might end into some weak occurrence $B'[A_1 \vee A_2]$ in some axiom of Π. In this case, we substitute $B'[A_1 \vee A_2]$ with the occurrence $B'[A_1]$ all along the path.

There might be also paths that end in some main formula $A_1 \vee A_2$ of a $\vee$:*left* rule and in this case we do not apply the rule in Π_1 but we consider as part of Π_1 only the subproof supporting the auxiliary formula A_1. More precisely, suppose that the premises of the $\vee$:*left* rule are $\Pi_\circ : A_1, \Gamma_\circ \rightarrow \Delta_\circ$ and $\Pi_* : A_2, \Gamma_* \rightarrow \Delta_*$. Then we will consider as part of Π_1, the proof $\Pi'_\circ : A_1, \Gamma_\circ, \Gamma_* \rightarrow \Delta_\circ, \Delta_*$ obtained from $\Pi_\circ$ by adding the formulas Γ_*, Δ_* as weak occurrences. (This addition does not increase the number of lines of $\Pi_\circ$ as proved in Lemma 4.8 pp. 274 in [Car97]). We need to do this for all occurrences of $A_1 \vee A_2$ as main formulas. At the end, the proof Π_1 will roughly look as Π^* except for the subproofs supporting the occurrences A_2, which will be removed. This is because Π^* is separated. (Notice that the hypothesis of separation ensures also that the side formulas Γ_*, Δ_* and $\Gamma_\circ, \Delta_\circ$ do not contain subformulas linked to $A_1 \vee A_2$ in the end-sequent.)

The last part of the statement is easily proved by induction on the height of the subproofs of Π. ⊣

Remark 21. If Lemma 20 is applied to a *reduced* proof Π, notice that we might obtain proofs Π_1, Π_2 which are not anymore reduced because of the addition of new weak occurrences (due to the cancellation of some of the subproofs of Π). In fact, non-weak formulas in Π might become weak in Π_1, Π_2. The same can be noticed for Lemma 18: unary rules might be applied to pairs of weak formulas (where one of the formulas has been added along the construction).

Remark 22. In Lemma 20, the situation is simpler than the one discussed in Remark 19 for Lemma 18. The transformation described in the proof of Lemma 20 introduces *new* weak formulas but *no* new contractions. This immediately implies that no duplication of bridges can arise. On the other hand, it might be that given a pair of occurrences in the end-sequents of Π_1 and Π_2, both proofs contain a bridge between these occurrences. This can happen when the occurrences do not lie in A_1, A_2. This fact will play a role in the duplication of bridges (and cycles) in Theorem 1.

An important point is that, if a bridge in Π links two occurrences lying in A_1 *and* A_2, the transformation in Lemma 20 will obviously break this path. We will use this simple observation in the proof of Theorem 1 where we will combine Lemma 18 and Lemma 20 to reduce the complexity of specific cut-formulas in a proof and therefore disrupt cyclic paths in a finite number of steps of reduction. The procedure we present in the proof of Theorem 1 allows to focus on *arbitrary* subformulas of a cut-formula and reduce their complexity. We will be interested only on those subformulas where a cyclic path passes through and we will reduce the complexity only of such subformulas. In this sense the procedure can be seen as a subtle version of cut elimination, where instead one considers *all* subformulas of a cut-formula.

Statements similar to Lemma 18 and Lemma 20 hold when the disjunction is replaced by a conjunction. In the sequel we will refer to Lemmas 18 and 20 even if we talk about conjunctions. The following statement is a consequence of Lemma 16 and tells us something about the structure of cyclic paths in proofs.

Lemma 23. *Let* Π *be a proof that contains a simple cycle. Then the last rule in* Π *where the cycle passes through is a cut, which we call* special. *Moreover, each cut-formula in the special cut contains a pair of distinguished occurrences that are connected by a logical path in* Π *and this path is part of the cycle.*

Proof. We consider the minimal subproof Π' of Π that contains the cycle. By minimality, the end-sequent of Π' does not contain the cyclic path, the last rule of Π' is a cut and that the cycle passes through its cut-formulas.

When a cyclic path arrives to a premise of the cut-rule, it passes through the cut-formulas and turns upward towards the axioms. If it does not do that, it will go down to the end-sequent of Π' by definition of logical flow graph, and this is in contradiction with the minimality. (Notice that if Π' is Π then the path will end-up into the end-sequent of Π and this is in contradiction with the cyclicity of the path.) In particular, the cyclic path needs to pass *twice* through the cut-formulas since any path coming to the cut-formula and passing through it needs to be rejoined when cyclic. But this can happen only if once going towards the axioms, the path comes back (maybe passing through several other cuts) to the cut-formula, passes through it and joins the other ending. (Note that there might be several crossing of the cyclic path through the same cut-formula.)

We proved that there is at least a pair of occurrences which lie in one of the cut-formulas of the last inference of Π', and that the occurrences are connected by a logical path in Π' which is part of the cyclic path. ⊣

A simple cycle might pass several times through its special cut, back and forth.

Definition 24. Let $\mathscr{C}$ be a simple cycle in the logical flow graph of Π. The set of cut-edges of $\mathscr{C}$ which belong to the special cut of $\mathscr{C}$ is called *special set* for $\mathscr{C}$.

The simple cycle lying in the formal proof illustrated in Section 2.2 has a special set defined by the two cut-edges indicated in the picture.

Notice that the cardinality of a special set is at least 2. This is an immediate corollary of Lemma 23.

Definition 25. Let $\mathscr{C}_1, \mathscr{C}_2$ be two simple cycles in the logical flow graph of Π. If $\mathscr{C}_1, \mathscr{C}_2$ have the same special set, then we say that $\mathscr{C}_1, \mathscr{C}_2$ are *nested* cycles.

The two cycles that occur in the proof corresponding to the Church numeral 3 in Section 3, are nested.

By Definition 25, if the special set of edges for $\mathscr{C}_1$ is *strictly* contained in the special set of edges for $\mathscr{C}_2$, then the cycles $\mathscr{C}_1$ and $\mathscr{C}_2$ are *not* nested. For short, we shall call *nested cycle* the collection of *all* those simple cycles of a logical flow graph of a proof which have the same special set. We shall refer to this latter as *special set* for the nested cycle.

We are now ready to show our main theorem.

Proof of Theorem 1. Without loss of generality we assume that Π is reduced *and* separated. If this is not the case then we apply propositions 10 and 15 repeatedly to Π and obtain a separated and reduced proof which is bounded in size by the size of Π. (Notice that after a finite number of iterations we always end-up with a reduced and separated proof since each transformation induced by propositions 10 and 15 reduces either the size, i.e. the number of symbols, or the number of lines of the proof. This is easy to check.)

We will start showing how to eliminate one simple cycle from a proof of k lines with a quadratic expansion of its number of lines (the new proof will have at most $\leq \mathscr{O}(k^2)$ lines).

We consider the special cut associated to the simple cycle and we assume it to be applied to the subproofs $\Pi_1 : \Gamma_1 \rightarrow \Delta_1, C$ and $\Pi_2 : C, \Gamma_2 \rightarrow \Delta_2$.

The size of the cut-formula C is bounded by $2 \cdot k$ by Lemma 12 and the fact that the proof is reduced. But one can do better. In fact, the size of C can be bounded by $2 \cdot min(\#\text{lines}(\Pi_1), \#\text{lines}(\Pi_2)) \leq k$. To see this, note that both cut-formulas C in Π_1, Π_2 do not contain occurrences which are weak in Π_1, Π_2, because Π is reduced. This means that all their atomic occurrences have a direct path to a distinguished occurrence in some axiom of Π_1, Π_2 respectively. But there are $2 \cdot \#\text{lines}(\Pi_1), 2 \cdot \#\text{lines}(\Pi_2)$ many such occurrences, respectively. Hence, the bound.

We look at the cut-formula C. If its main connective is a negation we look at its immediate subformula and so on until we find the larger subformula B which is not a negation. Let B_1, B_2 be the immediate subformulas of B. If there are two occurrences in B where the cycle passes through which lie one in B_1 and the other in B_2, we will proceed as described below. Otherwise, we choose the B_i where the cycle passes through and repeat the testing over B_i until we find a subformula B which is either a conjunction or a disjunction and such that there are two occurrences where the cycle passes through, lying in distinct immediate subformulas of B. Notice that such a formula B should exist because of Lemma 23.

Suppose that the main connective of B is strong in Π_1 and weak in Π_2. (The treatment is similar if the main connective of B is weak in Π_1 and strong in Π_2.) Suppose also that B_1, B_2 are the immediate subformulas of B. Then we apply Lemma 20 to Π_1 and Lemma 18 to Π_2 and combine the resulting subproofs $\Pi_{1,1} : \Gamma_1 \rightarrow \Delta_1, C[B_1]$, $\Pi_{1,2} : \Gamma_1 \rightarrow \Delta_1, C[B_2]$ and $\Pi_2' : C[B_1], C[B_2]\Gamma_2 \rightarrow \Delta_2$ through two new cuts on $C[B_1]$ and $C[B_2]$. The order of the cuts is important here. One should be careful in cutting first the formula $C[B_i]$ with B_i of larger logical complexity, and then cut the other. This choice will optimize the number of duplications induced on the proof by the algorithm. In fact, by the splitting of the subproofs, it might be that the cyclic path was split and if this is the case then we are done. Otherwise, the algorithm will proceed to analyze the current special cut, which is the *lowest* cut through which the cyclic path passes. By cutting last the formula $C[B_i]$, where B_i has smaller logical complexity, we ensure that the lowest cut has $|B_i| \leq |B|/2 \leq |C|/2$.

The procedure will be iterated at most $\log k$ times because the size of C is at most k and also because at the next step of the procedure we analyze a subformula of B_i (the subformula might be B_i itself) in $C[B_i]$ of size $\leq |B|/2$ and find a subformula B_i' of B_i such that $|B_i'| \leq |B_i|/2$. (Note that we do *not* need to look at the whole $C[B_i]$.)

By Lemma 23, a cyclic path passes through at least two occurrences in B. On the other hand, after $\log k$ steps of iteration we are guaranteed to end-up with an *atomic* formula B and this implies a disruption of the cyclic path. Of course, the procedure can stop in less than $\log k$ steps, i.e. as soon as the cycle is disrupted.

Since the number of steps in the algorithm is at most $\log k$, the expansion during the transformation of Π into a resulting proof Π' is

$$N(\Pi') \leq 2^{\log k} \cdot k = k^2$$

where the inequality is a consequence of Lemma 18 and Lemma 20 (in each step of the procedure we start from a proof Π^*, we split and recombine its subproofs and obtain a proof with at most $2 \cdot \#\text{lines}(\Pi^*)$ axioms, cuts and logical rules).

To derive the bound on the *number of lines*, we need to transform the proof Π' into a *reduced* proof Π'' (notice that after the transformation it might be that the proof Π' is not anymore reduced as explained in Remark 21). By Proposition 11, the number of contractions in Π'' is bounded by $2 \cdot \#\text{axioms}(\Pi'')$. Therefore

$$\begin{aligned}\#\text{lines}(\Pi'') &\le N(\Pi'') + 2 \cdot \#\text{axioms}(\Pi'') \\ &\le 3 \cdot k^2\end{aligned}$$

Before we proceed in explaining how to transform a proof with n simple cycles into an acyclic one, we would like to clarify a point concerning the creation of new simple cycles that we left undiscussed until now. There are *several* ways in which new simple cycles might be built through the transformation. We shall group them in three families and discuss the way they arise.

Let us start with a trivial situation that arises with the application of Lemma 20 to Π_1; it will constitute the first family. Suppose that there is a simple cycle $\mathscr{C}$ whose special cut lies in Π_1. The "duplication" of Π_1 into $\Pi_{1,1}, \Pi_{1,2}$ might preserve a copy of $\mathscr{C}$ in the logical flow graph of $\Pi_{1,1}$ and another copy in the logical flow graph of $\Pi_{1,2}$. Clearly, these simple cycles have different special cuts: one lies in $\Pi_{1,1}$ and the other in $\Pi_{1,2}$.

The second family arises when there is a simple cycle $\mathscr{C}$ which passes through the part of the proof under transformation (i.e. Π_1, Π_2) and whose special cut lies *below* this part.

The first situation that we present occurs when Lemma 20 is applied to Π_1. Suppose that there are bridges in Π, whose extremes lie in some side formula in Γ_1, Δ_1, that appear in both $\Pi_{1,1}$ and $\Pi_{1,2}$ (as explained in Remark 22). Since the side formulas Γ_1, Δ_1 of $\Pi_{1,1}, \Pi_{1,2}$ are contracted by construction (just below $\Pi_{1,1}, \Pi_{1,2}$), the bridges will end-up to have common extremes. If the bridge in Π_1 with extremes in Γ_1, Δ_1 was part of a simple cycle $\mathscr{C}$ in Π, then the two new bridges will now be part of two nested simple cycles. In fact, with the exception of the two bridges which identify after the contractions, the simple cycles coincide. This is visualized in the following diagram

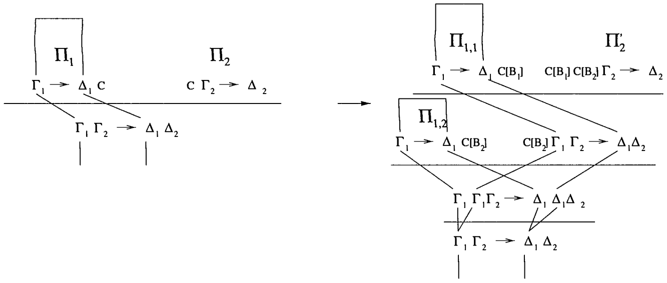
Π_1
Π_2
$\Gamma_1 \to \Delta_1\, C$
$C\, \Gamma_2 \to \Delta_2$
$\Gamma_1\, \Gamma_2 \to \Delta_1\, \Delta_2$
$\Pi_{1,1}$
Π_2'
$\Gamma_1 \to \Delta_1\, C[B_1]$
$C[B_1]\, C[B_2]\, \Gamma_2 \to \Delta_2$
$\Pi_{1,2}$
$\Gamma_1 \to \Delta_1\, C[B_2]$
$C[B_2]\, \Gamma_1\, \Gamma_2 \to \Delta_1 \Delta_2$
$\Gamma_1\, \Gamma_1 \Gamma_2 \to \Delta_1\, \Delta_1 \Delta_2$
$\Gamma_1\, \Gamma_2 \to \Delta_1\, \Delta_2$

Clearly, the two simple cycles in Π' have the same special cut and the same special set. (Notice that the special cut lies below the contractions through which the bridges identify.)

The symmetric situation occurs with the application of Lemma 18 to Π_2. Suppose that there is a bridge in Π_2 whose extremes lie in Γ_2, Δ_2. Lemma 18 might duplicate this bridge in two copies whose extremes coincide pairwise and lie in Γ_2, Δ_2 of Π_2' (see the first group in the classification described in Remark 19). If the bridge in Π_2

was part of a simple cycle $\mathscr{C}$ in Π, then the duplication would imply the formation of a new *simple* cycle in Π' which is nested to $\mathscr{C}$. This cycle will have the same special cut and the same special set of $\mathscr{C}$. (Notice that the special cut for the two cycles lies below the subproof of Π which is affected by the transformation.)

Another situation is illustrated by the picture below

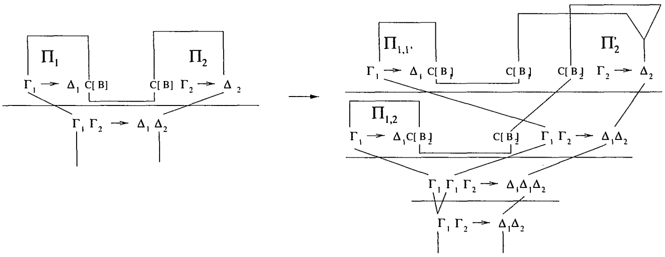

where Lemma 18 transforms a bridge in Π_2 into two bridges passing through $C[B_1], C[B_2]$ and identifies one of their extremes lying in Γ_2, Δ_2 (see the second group in the classification described in Remark 19). On the other hand, Lemma 20 induces two bridges in $\Pi_{1,1}, \Pi_{1,2}$ where one of their extremes is identified by the contractions on the collections Γ_1, Δ_1 in Π'. If the bridge in Π was part of a simple cycle $\mathscr{C}$, then after the transformation a new simple cycle would be formed. Again, the two cycles are nested.

A variation of the previous case is illustrated in the following figure

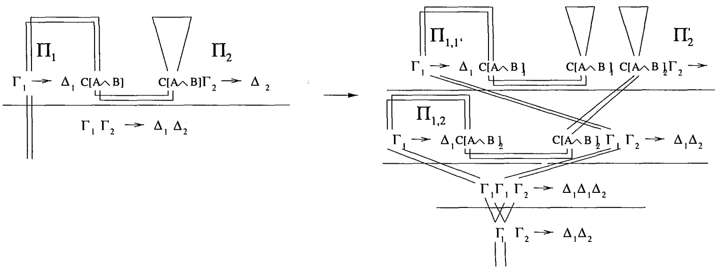

Suppose that there is a simple cycle $\mathscr{C}$ passing through A *and* Γ_1, Δ_1 in Π. If the bridge in Π_2 is duplicated as in the third group in the classification described in Remark 19, then a new simple cycle will be nested to $\mathscr{C}$.

Another variation is illustrated in the following figure

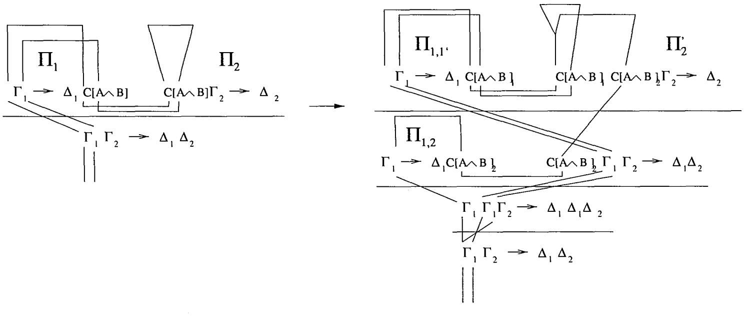

where the bridge in Π_2 duplicates as in the fourth group in the classification described in Remark 19 and recombines, after the "duplication" of Π_1, through contractions.

Suppose that the path in Π is part of a simple cycle. Then, after the transformation, this cyclic path is duplicated into two simple cycles whose special cut coincides.

The symmetric counterpart to these last two cases is illustrated in the following diagram

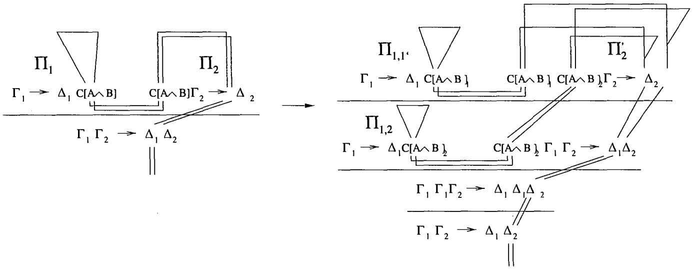

where the bridges in Π_2 duplicate as in the second group in the classification described in Remark 19. Suppose that there is a simple cycle $\mathscr{C}$ passing through A *and* Γ_2, Δ_2 in Π. The procedure induces a duplication of the parts of the cyclic path passing through Π_1, Π_2 and a new simple cycle nested to $\mathscr{C}$ is constructed. (Again, the special cut of the nested cycles lies below the part of the proof where the transformation takes place.)

The third family arises when a cycle $\mathscr{C}$ passes through the part of the proof under transformation (i.e. Π_1, Π_2) and its special cut *is* the cut on C applied to Π_1, Π_2. This means that the cut on C is *special* for both $\mathscr{C}$ and the simple cycle under consideration (in the procedure). Let the two simple cycles pass through "disjoint" subformulas of C, i.e. C contains a subformula $A \wedge B$ (or $A \vee B$) where one simple cycle passes through A and the other one passes through B. Suppose, for simplicity, that B is of the form $B_1 \wedge B_2$ and that the simple cycle passing through B, passes through both B_1 and B_2. (In general, B can have a more complicated logical form; for instance, it can be of the form $\neg \ldots \neg(B_1 \wedge B_2)$ with arbitrarily many $\neg$ preceding the binary connective, or of the form $D \vee (B_1 \wedge B_2)$. Our analysis applies to these cases as well.) If we were interested in eliminating the cycle passing through B, then our procedure would apply Lemmas 18 and 20 and furnish proofs of three sequents, say $\Pi_{1,1} : \Gamma_1 \rightarrow \Delta_1, C[A \wedge B_1]$, $\Pi_{1,2} : \Gamma_1 \rightarrow \Delta_1, C[A \wedge B_2]$ and $\Pi_2' : C[A \wedge B_1], C[A \wedge B_2]\Gamma_2 \rightarrow \Delta_2$, where one of the cycles passes through B_1, B_2 and the other through the A's.

There are several situations that might satisfy this setting. We start by illustrating the simpler one

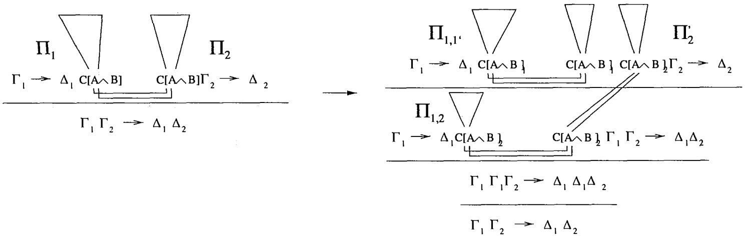

where the bridge in Π_2 duplicates as in the third group in the classification described in Remark 19. The cycle $\mathscr{C}$ passing through A duplicates into two simple cycles, and both the new cuts are special cuts for the pair of simple cycles. (Notice that, if the bridge in Π_2 duplicates as in the fourth group in the classification described in Remark 19, then there would be no duplication of the simple cycle since the bridge in $\Pi_{1,2}$ would end-up into a *weak* occurrence in $C[A \wedge B_2]$. This is easy to check.)

A variation of this situation is illustrated as follows

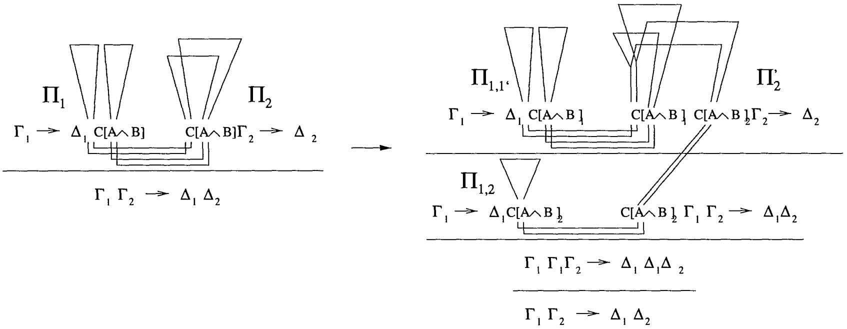
Π_1
Π_2
$\Gamma_1 \rightarrow \Delta_1\, C[A\wedge B]$ $C[A\wedge B]\Gamma_2 \rightarrow \Delta_2$
$\Gamma_1\Gamma_2 \rightarrow \Delta_1\Delta_2$
$\Pi_{1,1'}$
Π_2'
$\Gamma_1 \rightarrow \Delta_1\, C[A\wedge B]_1$ $C[A\wedge B]_1\, C[A\wedge B]_2\Gamma_2 \rightarrow \Delta_2$
$\Pi_{1,2}$
$\Gamma_1 \rightarrow \Delta_1 C[A\wedge B]_2$ $C[A\wedge B]_2\,\Gamma_1\,\Gamma_2 \rightarrow \Delta_1\Delta_2$
$\Gamma_1\,\Gamma_1\Gamma_2 \rightarrow \Delta_1\,\Delta_1\Delta_2$
$\Gamma_1\,\Gamma_2 \rightarrow \Delta_1\,\Delta_2$

Here, we have a simple cycle in Π that passes four times through the cut-formula C. The two bridges in Π_2 have been duplicated in Π_2' as in the fourth group in the classification described in Remark 19, and they are combined with the bridges in $\Pi_{1,1}, \Pi_{1,2}$ to form two simple cycles. One of these cycles passes through $C[A \wedge B_1]$ only, and the other one passes through both formulas $C[A \wedge B_1], C[A \wedge B_2]$. As one can check from the figure, the special cuts of the two simple cycles are the two cuts introduced by the transformation.

Let us assume now that $\mathscr{C}$ passes through both B *and* a subformula A which is "disjoint" from B in C. Again, we have that cycles might be duplicated. In the following picture the simple cycle $\mathscr{C}$ in Π passes through both A and B_1. We illustrate a transformation that generates two copies of the cycle

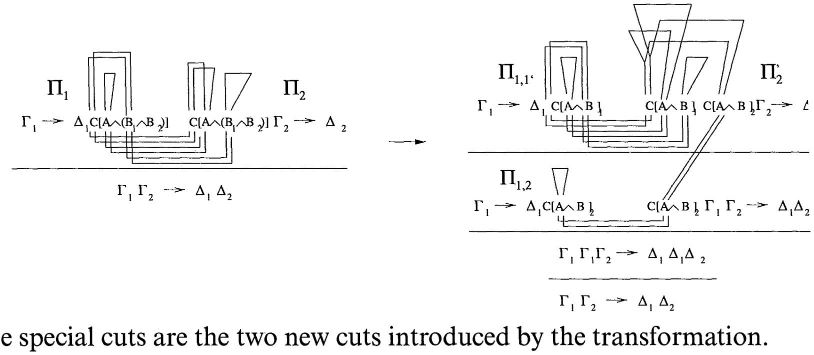
Π_1
Π_2
$\Gamma_1 \rightarrow \Delta_1 C[A\wedge(B_1\wedge B_2)]$ $C[A\wedge(B_1\wedge B_2)]\,\Gamma_2 \rightarrow \Delta_2$
$\Gamma_1\,\Gamma_2 \rightarrow \Delta_1\,\Delta_2$
$\Pi_{1,1'}$
Π_2'
$\Gamma_1 \rightarrow \Delta_1\, C[A\wedge B]_1$ $C[A\wedge B]_1\, C[A\wedge B]_2\Gamma_2 \rightarrow \Delta$
$\Pi_{1,2}$
$\Gamma_1 \rightarrow \Delta_1 C[A\wedge B]_2$ $C[A\wedge B]_2\,\Gamma_1\,\Gamma_2 \rightarrow \Delta_1\Delta_2$
$\Gamma_1\,\Gamma_1\Gamma_2 \rightarrow \Delta_1\,\Delta_1\Delta_2$
$\Gamma_1\,\Gamma_2 \rightarrow \Delta_1\,\Delta_2$

whose special cuts are the two new cuts introduced by the transformation.

Before concluding the analysis of the third family of situations, let us notice that a simple cycle $\mathscr{C}$, which passes through $B_1 \wedge B_2$ but not through the disjoint formula A, is *not* duplicated by the transformation. This is because the cut considered by the procedure is special for $\mathscr{C}$, and because Lemmas 18 and 20 do not induce any duplication of bridges whose extremes lie in B_1, B_2. (In particular, notice that the cycle that the procedure means to break is not duplicated either.)

The situations treated in the three families discussed above can be used as templates to describe *all* other cases where new simple cycles arise. In fact, to be exhaustive, we have to consider also the cases where $\mathscr{C}$ passes back and forth through the cut-formulas C several times. Our analysis can be applied, in a straightforward way, to capture these remaining cases as well. Take for instance the following transformation

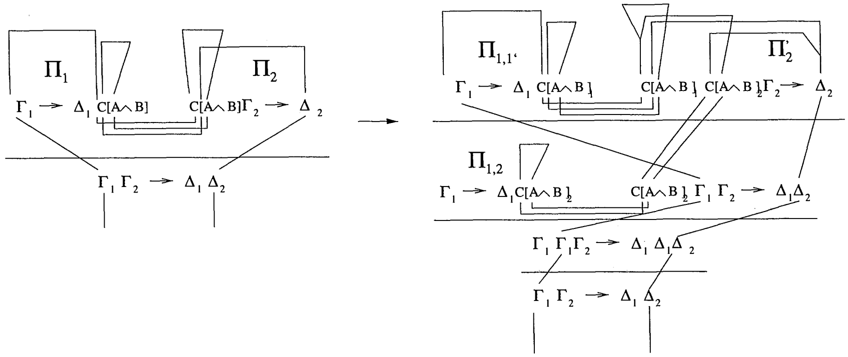

which can be classified in the second family, and where the paths in Π_2 duplicate as in the third and fourth group of the classification described in Remark 19. If the path in Π is part of $\mathscr{C}$, then the simple cycle is duplicated into two simple cycles whose special cut coincides.

It is important to notice that by Lemmas 18 and 20, any pair of occurrences in the cut-formulas C which are *unlinked* will remain unlinked after the procedure has been applied to C. Hence, no extra cycle is formed through the transformation, except the ones which are formed by bridges and cut-edges that already exist.

Notice that, if n different parts of the same simple cycle are involved in a step of transformation, the duplication might induce the formation of 2^n new nested simple cycles in Π'. Take for instance the following transformation

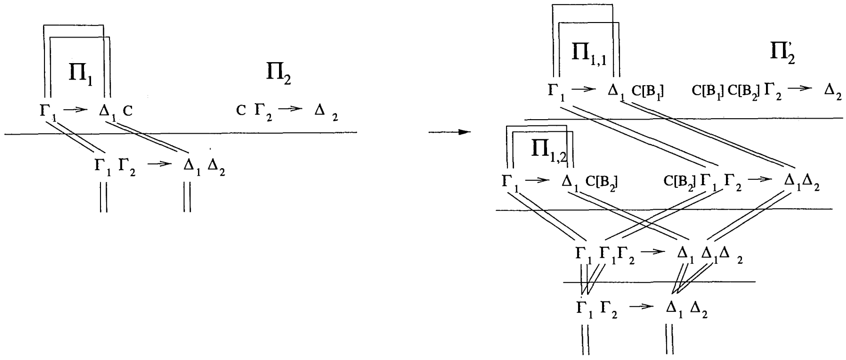

and suppose that the two bridges in Π are both parts of the same simple cycle. Then, after the transformation, we shall have four different simple cycles, one nested into the other.

All situations classified in the second family induce the number of *simple* cycles to increase with the construction but the number of *nested* cycles to remain the same. On the other hand, the first and third family induce the number of nested cycles to *augment*. The procedure we shall design for the elimination of all cycles will avoid *all* situations classified in the first and third families.

We are now ready to present the procedure and discuss its complexity. Let Π be a proof with n simple cycles and k_0 special cuts. Clearly $k_0 \leq n$ since there is exactly one special cut for each simple cycle. Let $n_{0,i}$ be the number of simple cycles associated to the i-th special cut of Π, for $i = 1, \dots, k_0$. Clearly $n_{0,1} + \dots + n_{0,k_0} = n$.

A simple cycle is based on a special set. Let $m_{0,i}$ be the number of distinct special sets on which the $n_{0,i}$ simple cycles are based. Notice that $m_{0,i}$ is the number of distinct nested cycles in Π passing through the i-th special cut. In particular, $m_{0,i} \leq n_{0,i}$ and therefore $m_{0,1} + \dots + m_{0,k_0} \leq n$.

We want to transform Π into an acyclic proof Π'. The idea is to repeat the procedure (that we shall describe in a moment) until the $m_{0,1} + \dots + m_{0,k_0}$ nested cycles are disrupted. Our aim is to eliminate *nested* cycles and not just simple cycles as before. The approach is similar though: the procedure considers, one after the other, the nested cycles in Π and it reduces the number of cut-edges in their special sets. (Put it in another way, it reduces the number of formula occurrences of the special cut where the nested cycle passes through.) By doing this, we shall guarantee that the number of special cuts k_l is $\leq n$ and that $m_{l,1} + \dots + m_{l,k_l} \leq n$, where $m_{l,j}$ is the number of nested cycles associated, at stage l, to the j-th special cut. In particular, we shall guarantee that the number of nested cycles does not increase. In other words, we need to avoid the first and third family of cases in the analysis above.

By Lemma 23, a special set has cardinality at least 2 and by the procedure, after a finite number of steps, the cardinality of any special set associated to some nested cycle is decreased to 1. This ensures the disruption of the nested cycles and the *termination* of the procedure.

The procedure is defined as follows. We start from a proof Π which is reduced and separated. As argued for the case of a simple cycle, this implies that the size of cut-formulas in Π is bounded by k, where k is the number of lines in Π. We want to eliminate first those nested cycles whose special cut is higher up in the proof. Namely, we want to treat first those nested cycles lying in subproofs that do not contain special cuts except for their last rule of inference. Let C be the cut-formula. Similarly to the procedure dealing with one simple cycle, we look for the larger conjunction or disjunction B in C which satisfies one of the following two conditions:

1. either B contains two atomic occurrences through which *a* nested cycle passes (remember that the special cut that we are considering might be associated to several nested cycles based on different sets of cut-edges, and that we look for any one of them here), and which lie in distinct immediate subformulas B_1, B_2 of B,
2. or C is a special cut for a pair of nested cycles, one that passes through B_1 and the other that passes through B_2, where B_1, B_2 are the immediate subformulas of B. (Notice that if the cycle passing through B_1 passes also through B_2 then condition 1 is satisfied.)

The procedure proceeds, as for the case of a simple cycle, by splitting and recombining subproofs Π_1, Π_2 through two new cuts on the formulas $C[B_1]$ and $C[B_2]$. It will not induce any trivial duplication of nested cycles since the duplicated subproofs have logical flow graphs which do not contain special cuts. This means that the first family of cases is ruled out. The third family of cases is also ruled out. In fact, it cannot be the case that a cycle passes through both B and a disjoint subformula in C because the formula B is the largest subformula of C satisfying condition 1, and also, it cannot be the case that a cycle passes through a subformula of C which is disjoint from B because of condition 2. On the other hand, it might happen that several new bridges are "glued" together by the *new* contractions introduced in the proof (i.e. the cases classified in the second family can take place). This might imply the number of simple cycles to increase but *not* the number of nested cycles.

Also, notice that the number of special cuts might increase (since a cut is transformed into two cuts by Lemmas 18 and 20). This happens when at least two nested cycles pass through a special cut on the cut-formula C and, after splitting, one of them passes through one of the new cut-formulas, say $C[B_1]$, and not through the other. If the cut on $C[B_1]$ precedes the cut on $C[B_2]$ in the new proof, then it is certainly special for the nested cycle passing through it. The cut on $C[B_2]$, being lower, will remain a special cut for the other nested cycles associated to C which pass through $C[B_2]$ and have not being disrupted by the splitting.

Since all nested cycles in the new proof were already present in the proof (this is guaranteed because no case classified in the first and third family can take place), the number of special cuts remains bounded by n and the inequality $m_{l,1} + \ldots + m_{l,k_l} \leq n$ will continue to hold. Notice that, if both cuts on $C[B_1]$ and $C[B_2]$ are special, then $k_l = k_{l-1} + 1$. Also, if at the l-th step of iteration, r nested cycles have been disrupted (it might be that several nested cycles are disrupted in a single step), then $m_{l,1} + \ldots + m_{l,k_l} = m_{l-1,1} + \ldots + m_{l-1,k_{l-1}} - r \leq n$.

The procedure reduces a special cut associated to a nested cycle in at most $\log k$ steps, where k is the number of lines of Π. (The argument is the same as for the case of a simple cycle.) Since the number of special cuts remains bounded by n, it is clear that the special cuts in Π will be reduced in at most $n \cdot \log k$ steps. Moreover, notice that the procedure might be activated at most n times by condition 2, since we can have at most n nested cycles which occur in "disjoint" subformulas in the special cuts.

If Π' is the acyclic proof, we have $N(\Pi') \leq 2^{n \cdot \log k + n} \cdot k = 2^n \cdot k^{n+1}$, since the number of steps of the procedure is at most $n \cdot \log k + n$.

As for the case of a simple cycle

$$\#\text{lines}(\Pi'') \leq 3 \cdot 2^n \cdot k^{n+1}$$

where Π'' is a reduced proof, obtained from Π' using Proposition 11. ⊣

REFERENCES

[Bus91] S. Buss, *The undecidability of k-provability*, ***Annals of Pure and Applied Logic***, vol. 53 (1991), pp. 72–102.

[Bus95] ———, *Some remarks on lengths of propositional proofs*, ***Archive for Mathematical Logic***, vol. 34 (1995), pp. 377–394.

[Bus93] ———, ***Personal communication***, June 1993.

[Car97] A. Carbone, *Interpolants, cut elimination and flow graphs for the propositional calculus*, ***Annals of Pure and Applied Logic***, vol. 83 (1997), pp. 249–299.

[Car97b] ———, *Duplication of directed graphs and exponential blow up of proofs*, ***Annals of Pure and Applied Logic***, vol. 100 (1999), pp. 1–76.

[Car97a] ———, *Turning cycles into spirals*, ***Annals of Pure and Applied Logic***, vol. 96 (1999), pp. 57–73.

[Car96] ———, *Cycling in proofs and feasibility*, ***Transactions of the American Mathematical Society***, vol. 352 (2000), pp. 2049–2075.

[CS97] A. Carbone and S. Semmes, *Making proofs without modus ponens: an introduction to the combinatorics and complexity of cut elimination*, ***Bulletin of the American Mathematical Society***, vol. 34 (1997), pp. 131–159.

[CS97a] ———, *A Graphic Apology for Symmetry and Implicitness — Combinatorial Complexity of Proofs, Languages and Geometric Constructions*, ***Mathematical monographs***, Oxford University Press, 2000.

[CR79] S. Cook and R. Reckhow, *The relative efficiency of propositional proof systems*, this Journal, vol. 44 (1979), pp. 36–50.

[Gen34] Gerhard Gentzen, *Untersuchungen über das logische Schließen I–II*, ***Mathematische Zeitschrift***, vol. 39 (1934), pp. 176–210, 405–431.

[Gir87] Jean-Yves Girard, *Linear logic*, ***Theoretical Computer Science***, vol. 50 (1987), pp. 101–102.

[Gir87a] ———, ***Proof theory and logical complexity***, Volume 1 of *Studies in Proof Theory, Monographs*, Bibliopolis, Naples, Italy, 1987.

[Hak85] A. Haken, *The intractability of resolution*, ***Theoretical Computer Science***, vol. 39 (1985), pp. 297–308.

[Kra97] J. Krajíček, *Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic*, this Journal, vol. 62 (1997), no. 2, pp. 457–486.

[Ore79] Vladimir P. Orevkov, *Lower bounds for increasing complexity of derivations after cut elimination*, ***Journal of Soviet Mathematics***, vol. 20 (1982), no. 4.

[Ore93] ———, *Complexity of proofs and their transformations in axiomatic theories*, ***Translations of Mathematical Monographs 128***, American Mathematical Society, Providence, RI, 1993.

[Pud96] P. Pudlák, *The lengths of proofs*, ***Handbook of proof theory*** (S. Buss, editor), North-Holland, Amsterdam, 1996.

[Pud97] ———, *Lower bounds for resolution and cutting plane proofs and monotone computations*, this Journal, vol. 62 (1997), no. 3, pp. 981–998.

[Sta78] R. Statman, *Bounds for proof-search and speed-up in the predicate calculus*, ***Annals of Mathematical Logic***, vol. 15 (1978), pp. 225–287.

[Tak87] G. Takeuti, ***Proof theory***, 2nd ed., Studies in Logic 81, North-Holland, Amsterdam, 1987.

[Tse68] G.S. Tseitin, *Complexity of a derivation in the propositional calculus*, ***Zapiski Nauchnykh Seminarov, Leningrad Otdelenie Matematicheskiĭ Institut, Akademiya Nauka SSSR***, vol. 8 (1968), pp. 234–259.

MATHÉMATIQUES/INFORMATIQUE
UNIVERSITÉ DE PARIS XII
61 AVENUE DU GÉNÉRAL DE GAULLE
94010 CRÉTEIL CEDEX, FRANCE
E-mail: carbone@ihes.fr

References

[Car96] **Cycling in Proofs and Feasibility**
A. Carbone
Transactions of the American Mathematical Society, Vol. 352, No. 5. (May, 2000), pp. 2049-2075.
Stable URL:
http://links.jstor.org/sici?sici=0002-9947%28200005%29352%3A5%3C2049%3ACIPAF%3E2.0.CO%3B2-9

[CR79] **The Relative Efficiency of Propositional Proof Systems**
Stephen A. Cook; Robert A. Reckhow
The Journal of Symbolic Logic, Vol. 44, No. 1. (Mar., 1979), pp. 36-50.
Stable URL:
http://links.jstor.org/sici?sici=0022-4812%28197903%2944%3A1%3C36%3ATREOPP%3E2.0.CO%3B2-G