GROUP CANCELLATION AND RESOLUTION

A. CARBONE

Institut des Hautes Études Scientifiques
35, route de Chartres
91440 – Bures-sur-Yvette (France)

Février 2002

IHES/M/02/07

Group cancellation and resolution

A. Carbone

February 20, 2002

Abstract

This paper contains the fourth chapter of my upcoming book *Combinatorial Geometry of Formal Proofs*. This explains the basic material in Sections 1 and 2 and Theorem 3. It was written during my stay at Courant Institute in the spring 2001.

We establish a connection between the geometric methods developed in the combinatorial theory of small cancellation and the propositional *resolution* calculus. We define a precise correspondence between *resolution proofs* in logic and *diagrams* in small cancellation theory, and as a consequence, we derive that a resolution proof is a 2-dimensional process. The isoperimetric function defined on diagrams corresponds to the length of resolution proofs.

1 The resolution calculus

The language: literals and clauses. Let $p_1, \dots, p_n$ be propositional variables. A *literal* is either a propositional variable p_i or a negation of a propositional variable $\neg p_i$. A *clause* is a disjunction of literals, i.e. a formula of the form $q_1 \vee q_2 \vee \dots \vee q_l$, where the q_i's are literals. The disjunction can contain one single literal, or may be none. In the latter case we say that the clause is *empty* and we denote it with the symbol $\perp$.

Normal forms. Any propositional formula built out of propositional variables and logical symbols $\wedge, \vee, \neg$ can be put in a *conjunctive normal form*, i.e. it can be rewritten as a conjunction of disjunction of literals. Once a formula is written in its normal form, we can think of it, with no ambiguity, as being a *set of clauses*. For instance, the formula $(\neg a \wedge b) \vee (c \vee a)$ is equivalent to its normal form $(\neg a \vee c \vee a) \wedge (b \vee c \vee a)$ and can be seen as the set of two clauses $\{\neg a \vee c \vee a, b \vee c \vee a\}$. In what follows, a formula is intended to be its associated set of clauses.

Tacit assumptions. When we speak of a clause, we usually have in mind any clause equivalent to a given one which can be obtained by permutating the literals. For instance, if $b \vee c \vee a$ is the given clause then, in practice, one thinks of the clause itself and of any of its permutations $c \vee a \vee b$, $a \vee b \vee c$ and $a \vee c \vee b$ as being the same clause. (Notice that the first two clauses are circular permutations and the third is not.)

Another implicit assumption is that any multiple occurrence of a literal in a clause can be identified. For instance, consider the clause $a \vee c \vee a \vee b$. It is tacitly thought to be the clause $a \vee c \vee b$.

Resolution rule, resolution calculus and resolution proofs. The rule of *resolution* takes two clauses containing a literal and its negation respectively, and combines them into a new clause which merges all the other literals belonging to the clauses into a larger clause. It is schematized as follows

$$\frac{b_1 \vee \ldots \vee b_n \vee a \quad \neg a \vee c_1 \vee \ldots \vee c_k}{b_1 \vee \ldots \vee b_n \vee c_1 \vee \ldots \vee c_k} \tag{1}$$

where we say that the rule *resolves* the literal a. We call *resolution calculus* the calculus defined by the resolution rule. A *resolution proof* is a binary tree of clauses, where the root of the theorem is labelled by the empty clause, i.e. $\perp$, the leaves are labelled by starting clauses, and the internal nodes are labelled by clauses derived by applying the resolution rule to the clauses labeling the antecedents of the node in question. A *derivation* in the resolution calculus is a tree of clauses as above, which does not necessarily end with the empty clause. It can be shown that the resolution calculus is *complete* and *valid*: any true formula can be derived from the calculus and any formula derived from the calculus is true. A formula A is *proved by resolution* if the empty clause $\perp$ is derived from the set of clauses associated to $\neg A$.

Example. We want to derive b from the set of clauses $b \vee a$, $\neg a \vee c \vee d$, $\neg d \vee b$ and $\neg c \vee b$. To do this, we add the clause $\neg b$ to the set of original clauses, and we try to derive the contradiction, i.e. $\perp$. This is a resolution proof

$$\cfrac{\cfrac{\cfrac{\cfrac{b \vee a \qquad \neg a \vee c \vee d}{b \vee c \vee d} \qquad \neg d \vee b}{b \vee c} \qquad \neg c \vee b}{b} \qquad \neg b}{\bot} \tag{2}$$

that combines the first two clauses by merging the literals $a, \neg a$, then combines the resulting clause with a third one by merging the literals $d, \neg d$, then a fourth one by merging $c, \neg c$ and finally merges the result with the clause $\neg b$ to obtain the empty clause. Implicitly, some other identification of literals have been performed by the applications of the rule. Namely, at the second application of the resolution rule the double occurrence of the literal b in the resulting clause has been reduced to one, and the same happened at the third application of the rule.

The size of a derivation in the resolution calculus. It is the number of clauses in the derivation tree.

2 Diagrams

Let G be a group and $< X, R >$ its presentation. Let F be the free group on X and let N be the normal closure of R in F. Clearly, $G = F/N$ and also an element $w \in G$ represents the *identity* iff $w \in N$. In particular, $w \in N$ iff, in the free group F, w is a product of conjugates of elements of $R^{\pm 1}$

$$w = \prod_{i=1}^{n} u_i r_i^{\pm 1} u_i^{-1} \tag{3}$$

with $u_i \in F$ and r_i in R, for all i.

Reduced words. A word w in G (and in F) is said to be *reduced* if no subword of the form ss^{-1} or $s^{-1}s$, with $s \in X$ occurs in w. We say that w is *cyclically reduced* if all cyclic permutations of w are reduced.

Cyclically reduced relators and symmetrization. We think of relators of G as *finite words* over an alphabet $X \cup X^{-1}$. Also, we shall think of a relator $r \in R$ as being *cyclically reduced.* With the symbol R^* we denote the set of all distinct cyclic permutations of the defining relators $r \in R$ and of their inverses r^{-1}. This set is called *symmetrization.*

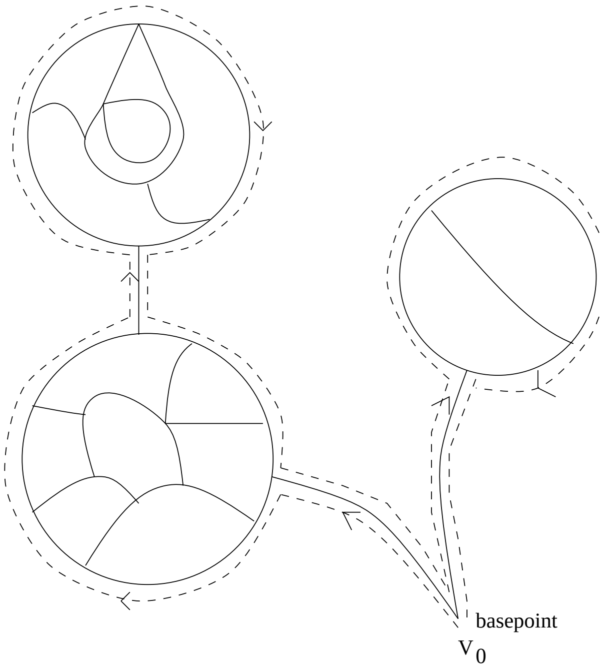

Figure 1: A 2-dimensional complex M, i.e. a tassellated multidisc.

To a product of conjugates, as in (3), we associate a *diagram* in the Euclidean plane which contains all the essential information about the product itself. Diagrams are used as a tool to study membership in N of F and equality in G. We start with some terminology and some basic concept.

Let $\mathbb{E}^2$ denote the Euclidean plane. If $S \subset \mathbb{E}^2$, then δS will denote the boundary of S, and $\bar{S}$ will denote the topological closure of S. A *vertex* is a point of $\mathbb{E}^2$. An *edge* is a bounded subset of $\mathbb{E}^2$ homeomorphic to the open unit interval. A *region* is a bounded set homeomorphic to the open unit disc.

2-*dimensional complexes.* A 2-*dimensional complex* M is a finite collection of vertices, edges and regions which are pairwise disjoint and satisfy the following properties:

1. if e is an edge of M, there are vertices a and b (not necessarily distinct) in M such that $\bar{e} = e \cup \{a\} \cup \{b\}$, and

2. the boundary δD of each region D of M is connected and there is a set of edges $e_1, \ldots, e_n$ in M such that $\delta D = \bar{e}_1 \cup \ldots \cup \bar{e}_n$.

We consider 2-dimensional complexes with oriented edges. The boundary of M is denoted δM. If M is constituted by several regions, then M is called a *multidisc*. See Fig. 1.

Oriented edges. If e is an edge with $\bar{e} = e \cup \{a\} \cup \{b\}$, then a and b are called *endpoints* of e. A *closed edge* is an edge e together with its endpoints. An edge might be traversed in either of the two directions. If e is an oriented edge running from a point v to a point w, the vertex v is the *initial vertex* and the vertex w is the final *vertex*. The oppositely oriented edge, or *inverse* of e, is denoted by e^{-1} and it runs from w to v.

Paths. A *path* is a sequence of oriented closed edges $e_1, \dots, e_n$ such that the initial vertex of e_{i+1} is the initial vertex of e_i, for $1 \leq i \leq n-1$. A *closed path* or a *cycle* is a path such that the initial vertex of e_1 is the final vertex of e_n. A path is *reduced* if it does not contain a subsequent pair of edges of the form ee^{-1}. A path is *simple* if all edges have distinct endpoints.

Diagrams. A *diagram over* $< X, R >$ *with boundary* P is a triple (M, f, P), where M is a 2-dimensional complex, f is a labelling map and P is a boundary path of M which starts and ends in a given basepoint, such that the following conditions are satisfied:

1. the space underlying M is homeomorphic to a simply connected closed subset of the plane;

2. f associates to each edge x of M a letter from $X \cup X^{-1}$; moreover $f(x^{-1}) = f(x)^{-1}$, for all oriented edges of M;

3. the label of every simple boundary path of a 2-cell of M is an element of R^*;

4. the boundary path P starts and ends at the basepoint.

A diagram is *reduced* if all its paths are reduced, i.e. there are no successive pairs of edges labelled xx^{-1} or $x^{-1}x$, where $x \in X$.

Example of a diagram. The 2-dimensional complex in Fig. 1, where for simplicity the edges have not being labelled, is a diagram. The boundary path P is read by going from left to right along the edges of the multidisc, starting from the basepoint v_0 and following the dotted line indicated in the figure until the basepoint is reached again and no more edges have to be read.

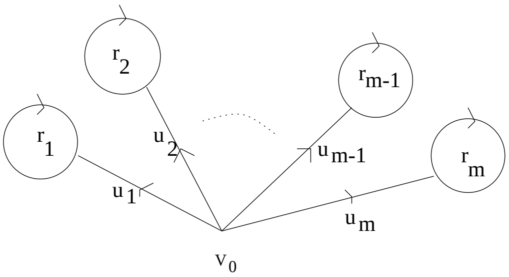

Figure 2: The representation of a trivial word as a product of conjugates.

2.1 Cancellation is a 2-dimensional process

van Kampen introduced an operation on diagrams that brings to light the fact that cancellation is a 2-dimensional process. He noticed that to any word w in the free group F one can associate a diagram M whose boundary P is w. Then, he showed that to any trivial word w in a group G, one can associate a diagram where the space underlying M is homeomorphic to a simply connected closed subset of the plane.

Diagrams and words over a free group: van Kampen procedure. To associate a diagram to a word in the free group F is simple. Any word w in F can be written in the form $w = u^{-1}su$, where s is a cyclically reduced word. In particular, any product $w = w_1 \dots w_n$ can be written as a product of conjugates of the form (3). To build a diagram for w, one starts with representing the product w as a 2-dimensional complex as illustrated in Fig. 2: each conjugate w_i is represented by a path labelled u_i followed by a disc associated to the reduced word r_i. A basepoint v_0 is common to all 2-dimensional complexes associated to the conjugates and the *order* of the conjugates of the product is respected.

The boundary of the 2-dimensional complex (as illustrated in Fig 1) reads as w. To see this, one starts at the basepoint v_0, reads first u_1 then goes around the disc and reads r_1, and finally goes back to v_0 by reading u_1^{-1}. Once one arrives to v_0 again, then will continue by reading the following path u_2, and so on until no more paths are to be read.

If w is reduced, we have obtained the desired diagram. If not, we reduce the label w of P by *sewing-up* subpaths xy of P which are products of two consecutive oriented edges whose labels are inverses of each other. This process needs to be iterated until no more sewing can be performed. At

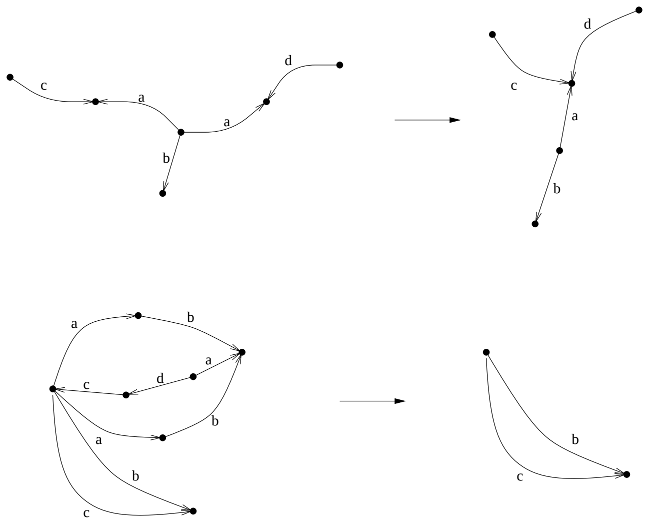

Figure 3: Top: step of sewing-up; bottom: cancellation of a sphere in a diagram.

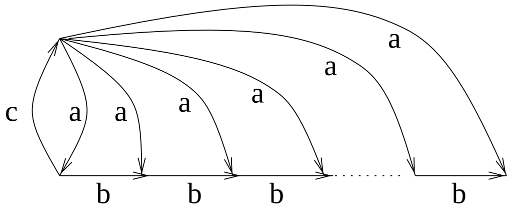

Figure 4: Diagram showing that the words $b^n = ca$ are trivial, for $n \geq 1$, in the group $\{a, b, c \mid ac = 1, ab = a\}$.

some step of this process it might happen that a 2-sphere could be formed, i.e. a disc whose boundary is of the form ss^{-1} or $s^{-1}s$, for some word s. In this situation, the 2-sphere should be discarded, together with the superfluous tail that might connect the sphere and the rest of the diagram. The outcome of the process is a diagram whose boundary path P_0 has label w_0. An example of a sewing-up and of the cancellation of a sphere are given in Fig. 3.

Before we continue, let us mention that a 2-sphere is an higher dimensional object that does not contribute any 1-dimensional information (i.e. it does not carry any 1-boundary) and for this reason it can be discharged.

Example. Consider the finitely presented group $\{a, b, c \mid ac = 1, ab = a\}$. The disc in Fig.4 shows that the words $b^n = ca$ are trivial, for $n \geq 1$.

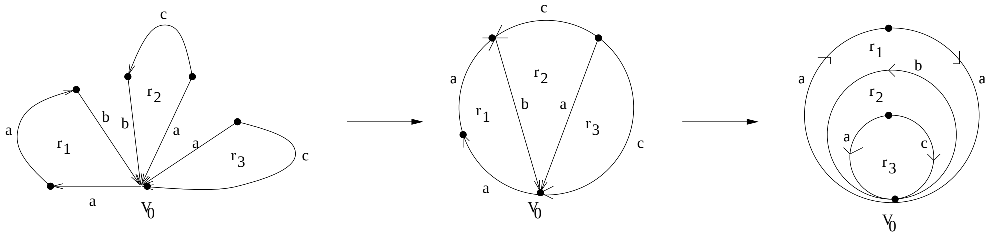

Figure 5: The van Kampen's procedure is applied to the three relations r_1, r_2, r_3 and results in a disc with boundary aa.

Diagrams and words over a finitely presented group G. The process explained for the free group F can be used for finitely presented groups $G =< X, R >$. In this case, one considers a word w in G that can be written as in (3). Then, one builds the diagram of the product as in Fig. 2 by considering discs with boundaries r_i's together with their tails as before (notice that r_i is a relator in R and that for a finitely presented group, simple discs are associated only to relators r_i). This multidisc can be reduced by following the steps of the construction above. van Kampen shows that if the word is *trivial*, then one ends-up with a disc which is homeomorphic to a simply connected closed subset of the plane, i.e. a disc whose boundary is w. On the other hand, he also shows that any diagram which is a disc with boundary w, implies that w is *trivial.*

Example. Consider the group $G =< X, R >$ where $X = \{a, b, c\}$ and $R = \{r_1, r_2, r_3\}$ with $r_1 = a^2b, r_2 = b^{-1}c^{-1}a$ and $r_3 = a^{-1}c$. In Fig. 5 we show that the word a^2 is trivial in G. The first step illustrated in the figure represents two steps of the procedure: one cancels two edges labelled b which belong to two distinct discs, and the other cancels two edges labelled a in a similar manner. The second step in the picture, identifies consecutive edges labelled c lying along the boundary. The resulting diagram is a disc with boundary aa, and by van Kampen's Theorem it follows that the word aa is trivial in G.

Operations of identification. van Kampen's procedure does not allow identifications between non-consecutive edges which belong to the boundary of the multidisc, or between adjacent vertices which are oriented in the same way. In particular, the *sewing-up* and the *discard of a sphere*, are very special kinds of cancellation. These restrictions ensure that the *only* constructible 3-dimensional object is the sphere. Arbitrary combinatorial operations of

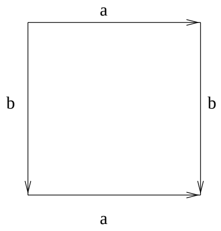

Figure 6: The basic pattern that forms a torus after identification of the edges labelled a and of the edges labelled b.

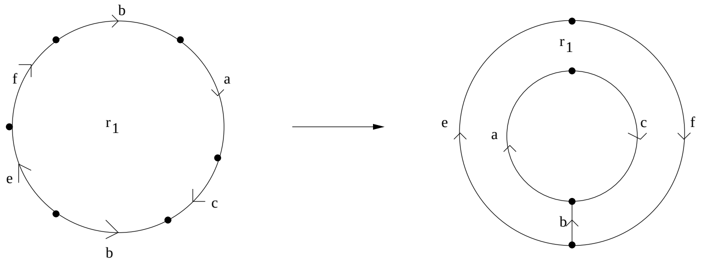

Figure 7: After identification of the edges labelled b and $\neg b$ in the boundary of a disc, a disconnection of the boundary takes place.

cancellation would bring into the picture more complicated surfaces, some embedded in higher dimensional spaces. A classical example is the torus which results from the identification of opposite sides of the square illustrated in Fig. 6. In Fig. 7, cancellation implies the disconnection of the boundary.

3 Resolution is a 2-dimensional process

The cancellation process for groups resembles the *resolution* process for propositional formulas described in Section 1. We modify the logical hypothesis of van Kampen's construction to capture the geometry of a derivation in resolution. The purpose is to show that logical proofs are *high dimensional* objects, even in the simple case of the resolution calculus. We define a combinatorial object, called *resolution diagram*, associated to a clause $p_1 \vee \ldots \vee p_k$. It is a 2-dimensional complex whose boundary is a sequence of oriented edges which starts and ends in a basepoint. Each edge of the boundary is labelled with a literal p_i, and the order of the edges (starting at a basepoint and following a clockwise direction) follows the presentation of the clause from right to left. The boundary reads $p_1 \cdot p_2 \cdot \ldots \cdot p_k$.

Notation. To be coherent with the notation used in Section 2, we denote a negative literal $\neg p$ with the symbol p^{-1}. Similarly, we talk about *composition* of literals instead of disjunction of literals. The interpretation remains unchanged.

Reduced words and reduced clauses. Similarly to groups, where diagrams are defined from *cyclically reduced* relators, we consider resolution proofs where clauses do *not* contain the literals p, p^{-1}, for some p. We shall call such clauses *reduced.* From a logical point of view, this assumption is not restrictive since derivations based on resolution attempt to show that a set of clauses is contradictory, and reduced clauses are not obviously true. Also, if a resolution proof Π contains some true clause, then one can transform it into a resolution proof of smaller size which is free of true clauses. To do this is easy. Given Π, there is a true clause $p, p^{-1}, a_1, \ldots, a_n$ and a clause $p^{-1}, b_1, \ldots, b_m$ that resolves the literal p (this is because Π is a *proof* and the last clause is empty, therefore all literals have to be resolved) by producing the clause $p^{-1}, a_1, \ldots, a_n, b_1, \ldots, b_m$. One can eliminate this application of the resolution rule by directly considering the clause $p^{-1}, b_1, \ldots, b_m$ instead of $p^{-1}, a_1, \ldots, a_n, b_1, \ldots, b_m$. Eventually, one should eliminate from the proof also those steps that resolve the literals $a_1, \ldots, a_n$. By performing this transformation on all true clauses in Π, we end-up with the desired proof. Based on these considerations, we assume that *derivations* also contain reduced clauses only.

Basic regions and structural regions. To a *starting* clause $p_1 \vee \ldots \vee p_k$ we associate a relation $p_1 \cdot \ldots \cdot p_k = 1$ called *basic relation*, and a region with boundary $p_1 \cdot \ldots \cdot p_k$, called *basic region.* Besides basic relations, we allow

$$pwpw^{-1}p^{-1} = 1 \tag{4}$$

with regions of boundary $pwpw^{-1}p^{-1}$, where p is a literal and w is a composition of literals. These relations come from the tacit assumption that, in resolution, any disjunction of the form $p \vee w \vee p$ is equivalent to $p \vee w$, where p is a literal and w is any disjunction (maybe an empty disjunction) of literals. There is a second implicit relation that is considered in resolution proofs

$$pwp^{-1}w^{-1} = 1 \tag{5}$$

and it corresponds to the fact that any disjunction $w \vee p$ is equivalent to $p \vee w$. We allow regions associated to this relation as well. The regions for (4) and (5) are called *structural regions.*

Symmetrization. Let R be a set of reduced basic relations and structural relations of the form (4) and (5). The symbol R^* denotes the set of all distinct cyclic permutations of relations $r \in R$ and of their inverses. The set R^* is called *symmetrization* of R.

Resolution diagrams. A *resolution diagram* with boundary w is a triple (M, f, w) where M is a 2-dimensional complex, f is a labelling map, and w is the boundary path of M which starts and ends in a given basepoint, such that the following properties are satisfied

1. the space underlying M is homeomorphic to a simply connected closed subset of the plane;
2. f associates to each edge x in M a positive literal; moreover, $f(x^{-1}) = f(x)^{-1}$, for all oriented edges of M;
3. the label of every simple boundary path of a region of M is an element of R^*,
4. the boundary path w starts and ends at the basepoint.

Resolution diagrams and resolution proofs. The construction of a resolution diagram associated to a resolution proof goes as follows. We start with a sequence of basic and structural regions which are all connected to a basepoint. Basic regions correspond to the starting clauses of the proof. At each intermediate stage of the construction, two adjacent discs (possibly constituted by several regions, and built in some previous step of the construction) are merged by identifying (parts of) their boundaries. The boundaries of the discs correspond either to the two clauses which have to be resolved at the current stage, or to structural rearrangements of the literals in the clause. The last step of the procedure is applied to two discs with opposite boundaries, and the result is a 3-dimensional sphere. As for group cancellation, we allow the *discard of the sphere*, and as a result, the resolution diagram vanishes.

As for group cancellation, we identify two edges in the boundary of the 2-dimensional complex when they are adjacent and directed towards opposite directions.

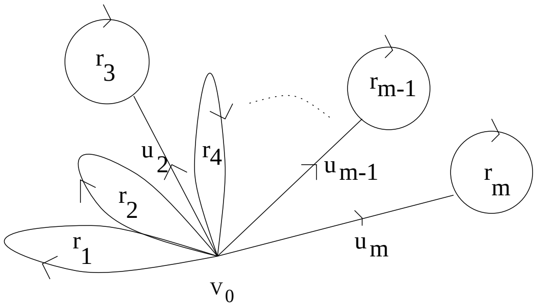

Figure 8: The structure of a resolution diagram for a proof. The stems with a disc on their top describe structural relations (4) and (5) which appear between words s, s^{-1}, for some s. These relations are tacitly applied along the derivation. The symbol $*$ labels the basepoint.

Theorem 1 *Let* Π *be a resolution proof derived from the set of clauses* S. *There is a* 2*-dimensional complex* M*, constituted by basic regions from* S *and by structural regions connected through a basepoint, which vanishes.*

Proof. By induction on the height of the resolution proof, we construct a 2-dimensional complex M, and a disc D (corresponding to the derivation) obtained from M by identifying its edges. The boundary of D corresponds to the clause resulting from the derivation. At the last step of the procedure the disc vanishes.

For each starting clause $p_1 \vee \ldots \vee p_n$ in the resolution proof, we let M and D be the region with boundary $p_1 \cdot \ldots \cdot p_n$.

Suppose that two clauses composed through the resolution rule have the form $p_1 \vee \ldots \vee p_n \vee p$ and $\neg p \vee q_1 \vee \ldots \vee q_s$. By induction they are associated to two complexes M_1, M_2 and two discs D_1, D_2. We identify the basepoints of M_1, M_2 and call M' the resulting complex. The boundaries of D_1, D_2 are $p_1 \cdot \ldots \cdot p_n \cdot p$ and $p^{-1} \cdot q_1 \cdot \ldots \cdot q_s$, where the edges labelled p and p^{-1} have the basepoint in common. We identify p and p^{-1} and we obtain a disc D' with boundary $p_1 \cdot \ldots \cdot p_n \cdot q_1 \cdot \ldots \cdot q_s$. If $C = p_1 \vee \ldots \vee p_n \vee q_1 \vee \ldots \vee q_s$ is the resulting clause C' of the resolution proof then we let M to be M' and D' to be D (notice that D is obtained from M by identification of edges). If C in not C' (this possibility is discussed in Section 1), then $C' = r_1 \vee \ldots \vee r_m$ must be identical to C up to commutation of literals and cancellation of multiple copies of the p_i's and q_j's. Hence, to obtain a disc D with boundary $r_1 \cdot \ldots \cdot r_m$ we need to use structural regions in the obvious way. That is, for any implicit application of the relation $pwp = pw$ ($pw = wp$) for instance, the structural region R (associated to $pwp = pw$) should interact with D'. For this, we define M to be the complex constructed by identifying the basepoints of M' and R. We identify D' and R along the sequence of edges pwp (pw) in the boundary. If the boundary of D' is $spwps'$ ($spws'$), where s, s' are words, then we consider the 2-complex R with boundary $s'^{-1}p^{-1}w^{-1}p^{-1}pws'$ instead, and after the identification of the boundaries, we end-up with a disc D of boundary $spws'$ ($swps'$). Several structural relations might have to be applied to D', and we define accordingly the complex M associated to them. See Fig. 8.

By repeatedly performing the operation of identification of pairs of literals p, p^{-1} with the help of structural regions, we construct a complex made out of two discs with opposite boundaries, i.e. one of them has a boundary p and the other p^{-1}, for some literal p. We identify the two discs and form a sphere with no boundary. By discarding the sphere, the diagram vanishes and this corresponds to the fact that the empty clause is proved.

□

Corollary 2 *Let* Π *be a derivation involving at most* k *distinct literals. There is a resolution diagram* M *associated to* Π *which is reducible to a disc* D *such that* $|\delta D| \leq k$*. If* n *is the number of resolution steps in* Π *then* D *has at most* $2 \cdot k \cdot n$ *regions.*

Proof. The bounds follow directly from the construction in the proof of Theorem 1.

The clauses in the derivation have length $\leq k$. This means that after a step of resolution, we obtain a disc with boundary length $\leq 2 \cdot (k-1)$. This means that resolution performs *implicitly* a reduction of at most $k-1$ literals through tacit identifications. After such identifications, reducing the length of the clause down to k, one might need to apply commutative operations to exchange the order of the edges along the boundary. The number of commutative operations might be reduced to 1 (since identification in Π is done on exactly one literal, which needs to reside either on the left or on the right of a resolving clause), but in general one might want to apply commutativity at most k times.

To conclude, the construction requires at most $2 \cdot k$ implicit operations for each step of resolution, and each operation corresponds to the presence of a structural region in the resolution diagram associated to the resolution

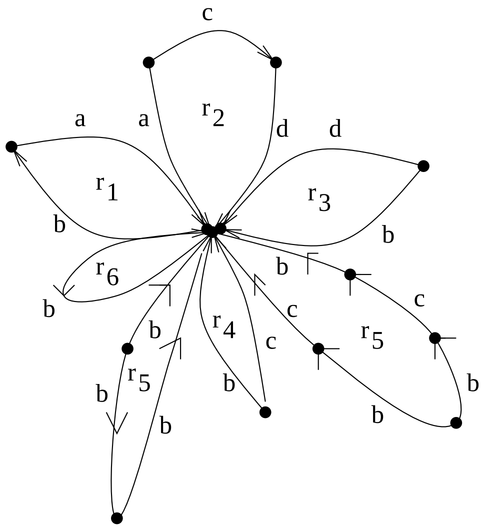

Figure 9: The resolution diagram for the proof displayed in (2). The symbol $*$ labels the basepoint.

proof. This means that D has at most $2 \cdot k \cdot n$ regions.

□

Example of a resolution diagram. Fig. 9 illustrates the resolution diagram associated to the resolution proof (2). In Fig. 10, we reduce the diagram to a point. In the first step, the discs corresponding to the basic relations ba and $a^{-1}cd$ are glued together through the cancellation of the literals $a, \neg a$. In the second step we identify d and $\neg d$, and apply relation r_5 to identify two occurrences of b along the boundary. The last step illustrated in the figure represents the identification of $c, \neg c$ followed by the identification of two occurrences of b in the boundary. The resulting multidisc is composed by two discs with opposite boundary b, b^{-1}. The very last step of the procedure cancels the literals $b, \neg b$ along the boundaries of the discs by forcing the creation of a 3-sphere with no 1-boundary. The vanishing of the diagram corresponds to the derivation of the empty clause.

3.1 Resolution proofs as products of words

Given a word $p_1 \cdot \ldots \cdot p_n$ labeling the boundary of some resolution diagram, we can write down $p_1 \cdot \ldots \cdot p_n$ as a product of conjugates $u_i r^{\pm 1} u_i^{-1}$, for $i = 1 \ldots n$, as in (3), where the r_i's are either basic relations or structural relations, and

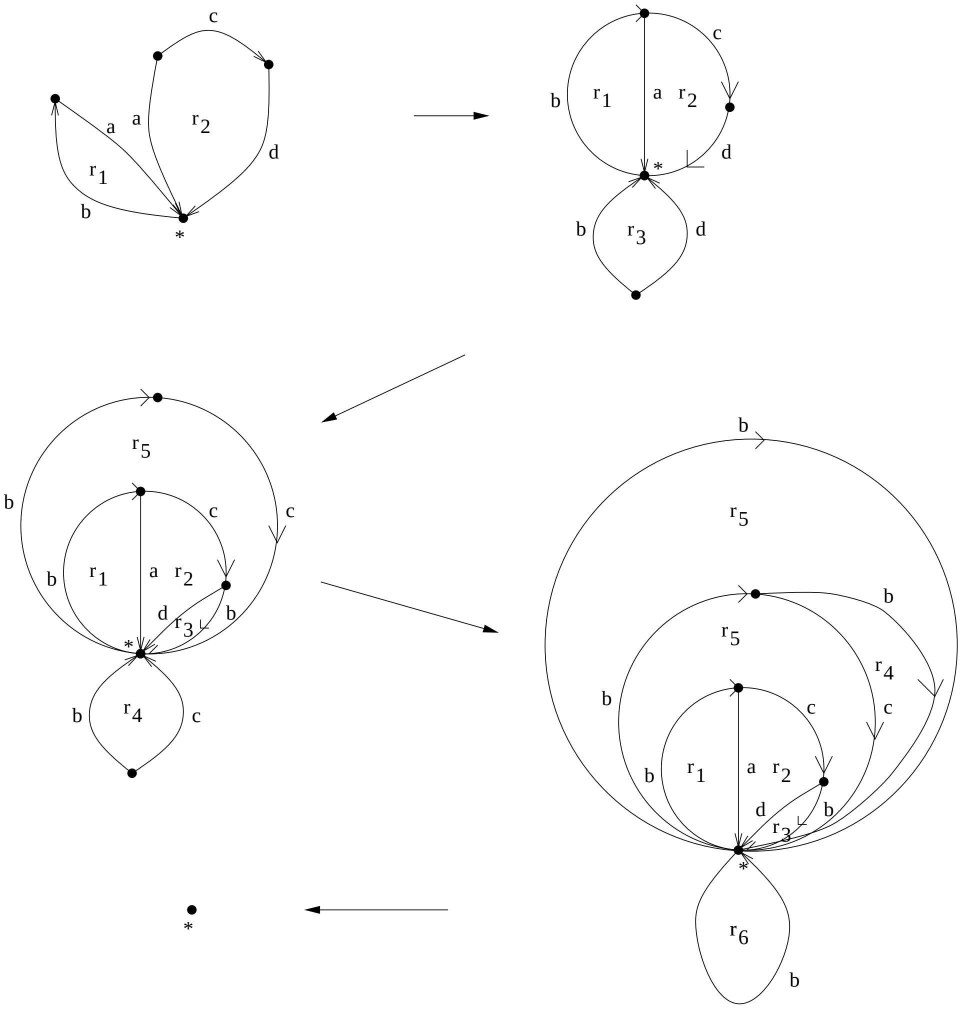

Figure 10: The sequence of interactions between discs during the reduction of the resolution diagram in Fig. 9. For convenience, at each step of identification, only the relevant discs are illustrated.

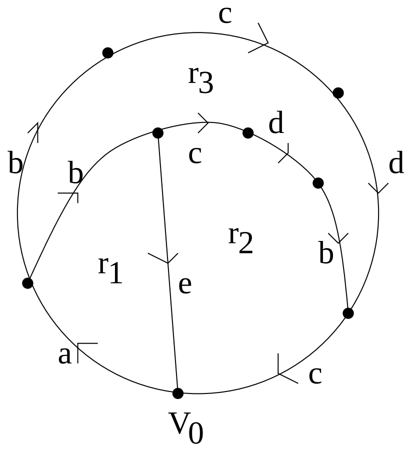

Figure 11: A disc with boundary *abcdc* and combinatorial area 3.

where the u_i's might be *non-trivial* words only if the corresponding r_i is a structural relation. (See legend of Fig. 8 and proof of Theorem 1.) To do this is rather straightforward: one reads the boundary of the regions of the disc, from left to right, starting from the basepoint V_0, and makes the free product of these words. (Notice that these words are conjugates.) For instance, the diagram in Fig 11 gives the words abe, $abcdb^{-1}d^{-1}c^{-1}b^{-1}a^{-1}$ and $e^{-1}cdbc$ associated to its three regions. The free product of the three words (read from left to right) is the word $abcdb^{-1}d^{-1}c^{-1}b^{-1}a^{-1}abee^{-1}cdbc$, which after simplification is reduced to $abcdc$, i.e. the boundary of the disc.

Theorem 3 (Bounding length) *Let D be a resolution diagram homeomorphic to a disc with boundary w. The word w is expressed as a free product of n conjugates of basic and structural relations as in 3, and we can rewrite the product so that each u_i has length at most $(|w| + 2k)2^n$, where k is the maximum length of the relators.*

Proof. Suppose that w is expressed as a product of conjugates $u_i r_i^{\pm 1} u_i^{-1}$, for $i = 1 \ldots n$. We have seen above how to do it. To show the statement, we make use of the *dual* of a van Kampen diagram whose construction is described as follows (see top left of Fig. 12). We start by drawing the conjugates of relators in clockwise order around a point, each $u_i r_i^{\pm 1} u_i^{-1}$ appearing as a path (representing u_i) terminated by a loop (representing $r_i^{\pm 1}$). We call this a *bouquet*. Reducing the product of conjugates in the free group means pairing off adjacent edges with inverse labels. We mark by pairing off with

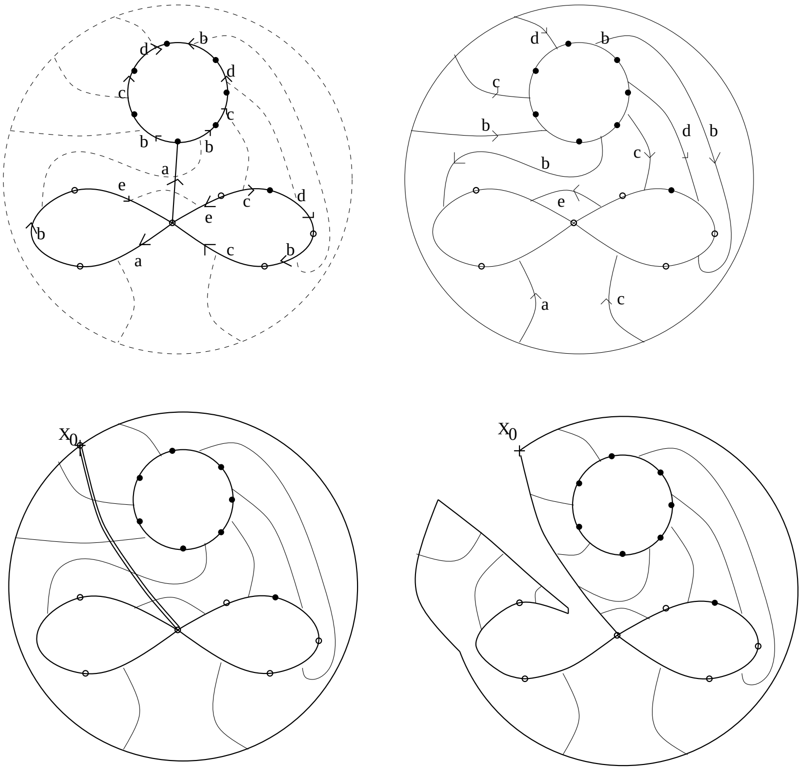

Figure 12: Top left: the construction of the dual of the van Kampen diagram in Fig. 11. The bouquet of conjugates is drawn with thicker lines. Each arrow between two nodes represents a literal in the relator. Top right: the outer circle, the inner circles and the connecting curves represent the dual of the van Kampen diagram. Bottom: the cutting of the diagram. The point X_0 represents the basepoint of the dual diagram.

edges connecting the midpoints of the edges (dotted lines in the figure). After edges can no longer be paired off, the remaining edges spell out w. We draw w as a circle surrounding the diagram, and pair off the remaining edges of the bouquet with the edges of w.

We interpret the region between the outer and the inner relator loops as a disk with holes, and the identifications as edges running between the various boundary components (top right of Fig. 12). Edges might also join to make a closed loop, but this corresponds to a total cancellation of the relators contained in the interior of the loop, so we could have omitted these terms from the product in the first place.

Observe that, given a dual diagram for w we can reconstruct a product of conjugates by drawing non-intersecting paths from the basepoint to each of the holes, and reading off the path labels from the edges intersected, keeping track of orientation. This concludes the construction of the dual of a van Kampen diagram.

Let us go back to the proof of our statement (see bottom of Fig. 12). The basepoint X_0 of w can be connected to a point of one of the loops in the dual of D by a path u that crosses at most $\frac{1}{2}(|w| + k)$ edges. To see this, cut the disc with holes along all edges that begin or end at a loop, and look at the connected component P containing X_0. The only remaining edges are those running from w to itself, and there are at most $\frac{1}{2}|w|$ of them. To connect X_0 with some hole, we need to cross at most these many edges. Also, we might need to cross another $\frac{1}{2}k$ edges to get to a suitable point of the relator.

We now cut the diagram open along the path u, getting a disc with one fewer holes. The boundary of the new disc is $w' = wurw^{-1}$, where r is some relator or its inverse and u has length at most $\frac{1}{2}(|w|+k)$. Thus w' has length at most $2(|w|+k)$. Since w' can be written as the product of $n-1$ conjugates of relators and $w = w'uru^{-1}$, the result follows by induction on n.

□

There are many ways to write down a word w as a product of conjugates based on basic and structural relations. This corresponds to the fact that there are many discs with boundary w and to the fact that there are many proofs of the same theorem.

Combinatorial area of a diagram, a word and a sphere. Following the terminology used in combinatorial group theory, we say that a *resolution diagram* D has combinatorial area, denoted $area(D)$, n if n is the number of regions that compose the disc. The disc in Fig 11 has combinatorial area 3, with two regions associated to basic relations and one to structural relations. The combinatorial area of a *word* w, denoted $area(w)$, is the minimal combinatorial area of a disc spanning w. If a set of basic relations S allows for the construction of a resolution diagram that vanishes after the formation of a *sphere*, then we say that the combinatorial area of S, denoted $area(S)$, is the minimal combinatorial area of the spheres that can be built from S (and possibly from structural relations), where the combinatorial area of a sphere is the sum of the combinatorial areas of the two discs forming the sphere.

Isoperimetric function. Suppose that $\mathcal{S} = \{S_i\}_{i=1}^{\infty}$ is a family of sets of basic relations S_i defined on the literals X_i, where $|X_i| = i^{\mathcal{O}(1)}$, for all $i \geq 1$, such that each S_i allows for the construction of a resolution diagram that vanishes. The *isoperimetric function* of $\mathcal{S}$ is defined by

$$\phi(i) = area(S_i) \tag{6}$$

Thinking the sets S_i as being inconsistent sets of clauses, we have that the isoperimetric function induces a complexity measure on resolution proofs.

Theorem 4 *There is a family $\mathcal{S} = \{S_i\}_{i=1}^{\infty}$ of relations S_i, defined on the literals X_i, where $|X_i| = i^{\mathcal{O}(1)}$, for all $i \geq 1$, such that each S_i allows for a vanishing resolution diagram. The isoperimetric function of $\mathcal{S}$ grows exponentially.*

Proof. From Haken's exponential lower bound for resolution, we know that the sets of clauses representing the negations of the pigeon-hole principle $\neg PHP_n$, for all $n \geq 1$, are inconsistent and that the proof of inconsistency must be of *exponential* size in n. This means that the resolution discs associated to the resolution proofs of PHP_n, for all n, should have exponential combinatorial area.

□

4 Bibliographical remarks and others

Resolution calculus. The resolution calculus was introduced by Blake [Bla37]. Building on work of Herbrand [Her71], there was much activity in theorem proving in the early '60 by Prawitz [Pra60], Davis and Putnam [DaPu60], Gilmore [Gil60], Robinson [Rob65]. A proof of completeness for resolution can be found in [DaPu60]. The introduction of logic programming, which uses resolution as an inference rule, is mainly due to Kowalski [Kow74] and Colmerauer [Col73].

Lower bounds for resolution. In [Tse68], Tseitin showed a lower bound for *regular resolution.* A sub-exponential lower bound for resolution was shown by Haken in [Hak85], for the pigeon-hole principle. An exponential lower bound was found by Urquhart in [Ur87] for Tseitin's tautologies (see below). Haken's lower bound was improved and generalized by Buss and Turán in [BuT88] for PHP_n^m, i.e. the pigeon-hole principle for m pigeons and n holes. They show that any resolution proof of PHP_n has at least $2^{\Omega(\frac{n^2}{m})}$ clauses. A proof of this result can be also found in [Kra95].

Theory of small cancellation. The exposition follows closely the presentation in [LS77] and [S90]. The reader can find there more information. Diagrams have been introduced by van Kampen in 1933 [vKa33] even though he did not make himself much use of them, and other authors did not consider them until 1966, when they have been rediscovered by Lyndon who used them to start a geometric study of cancellation in groups [Lyn66]. In these same years, Weinbaum discovered van Kampen's paper and used diagrams to prove results in small cancellation theory [Wei66]. van Kampen's diagrams are, at times, called *Dehn's diagrams.*

van Kampen's Theorem. The intuitive description of the construction of diagrams from products of conjugates seems to be the only type of proof of van Kampen's Theorem present in the literature. A formalized proof would need to involve too many subcases: diagrams would need to be dismantled, simplified and reassembled in the course of the construction.

The finitely presented group $\{a, b, c \ : \ a^2b = 1, b^{-1}c^{-1}a = 1, a^{-1}c = 1\}$ has been considered in [S90]. The proof of Theorem 3 is the same as in Lemma 2.2.4 pp 42 of [EetAl].

Towards diagram groups? The role of van Kampen diagrams for groups is similar to the role of semi-group diagrams in the study of semi-groups [LS77]. Recently Guba and Sapir looked at the structure governing operations applied to semi-group diagrams, monoid pictures, annular diagrams, braided pictures and cylindric pictures, and developed the theory of *diagram groups* [GSa96]. Their approach already inspired the work on proof structures for LK in [Car99a], and is likely to be relevant in the comprehension of the 2-dimensional processes underlying proof structures of logical systems which are structurally more complicated than resolution.

Number of regions in a diagram and size of resolution proofs. Given a word w, what is the number of relations that one might need to apply to get a resolution diagram with boundary w? From the proof of Proposition 2, at any given stage of the procedure one obtains a disc whose boundary has length at most k. There are 2^k many such words, and therefore the worse estimate is *exponential* in the number of literals. Similarly, given any tautology, the number of steps needed to prove it is, in the worse case, exponential in the size of the tautology.

Linking proofs to groups, and other proof systems. This question has been formulated in [CS97b] for finitely presented groups based on a finite set of generators:

Question. Let $R_1, \ldots, R_k$ be a finite set of relations. Suppose that all words w_i of length i which are equivalent to the empty word, have combinatorial area $\leq C^i$, for some constant C. Are there quantifier-free proofs of the triviality of the words w_i, which are of polynomial-size in i?

Here, we think of quantifier-free proofs as being proofs in a first order language where no quantifiers are used. No assumptions on the form of the (finite set of) axioms are made.

References

[Bla37] A. Blake. *Canonical Expressions in Boolean Algebra.* PhD thesis, University of Chicago, 1937.

[BuT88] S. R. Buss and G. Turán. Resolution proofs of generalized pigeonhole principles. *Theoretical Computer Science*, 62:311–317, 1988.

[Car99a] A. Carbone. Streams and strings of formal proofs. *Theoretical Computer Science.* To appear.

[CS97b] A. Carbone and S. Semmes. *A Graphic Apology for Symmetry and Implicitness.* Mathematical Monographs, Oxford University Press, 2000.

[Col73] A. Colmerauer, H. Kanoui, P. Roussel and R. Pasero. *Un système de communication homme-machine en français.* Groupe de Recherche en Intelligence Artificielle, Université d'Aix-Marseille, 1973.

[DaPu60] M. Davis, H. Putnam. A computing procedure for quantification theory. *Journal of ACM*, 7:201–215, 1960.

[EetAl] D.B.A. Epstein, J.W. Cannon, D.F. Holt, S.V.F. Levy, M.S. Paterson and W.P. Thurston. *Word Processing in Groups.* Jones and Bartlett Publishers, Boston and London, 1992.

[Gil60] P.C. Gilmore. A proof method for quantification theory. *IBM Journal Res. Develop.*, 4:28–35, 1960.

[GSa96] V. Guba and M. Sapir. *Diagram Groups.* Memoirs of the American Mathematical Society, number 620, volume 130, 1996.

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

[Her71] J. Herbrand. *Logical Writings.* Edited by W.D. Goldfarb, Harvard University Press, 1971.

[vKa33] E.R. van Kampen. On some lemmas in the theory of groups. *American Journal of Mathematics*, 55:268–273, 1933.

[Kow74] R.A. Kowalski. Predicate logic as a programming language. *IFIP 74*, 569–574, 1974.

[Kra95] J. Krajíček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory.* Encyclopedia of Mathematics and its Applications 60, Cambridge University Press, Cambridge, 1995.

[Lyn66] R.C. Lyndon. On Dehn's algorithm. *Math. Annalen*, 166:208–228, 1966.

[LS77] R.C. Lyndon and P.E. Schupp. *Combinatorial Group Theory.* Ergebnisse der Mathematik und iher Grenzgebiete 89, *A Series of Modern Surveys in Mathematics*, Springer-Verlag, 1977.

[Pra60] D. Prawitz. An improved proof procedure. *Theoria*, 26:102–139, 1960.

[Rob65] J.A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of ACM*, 1:23–41, 1965.

[S90] R. Strebel. Small cancellation groups. In *Sur les groupes hyperbolique d'apres Mikhael Gromov*, E. Ghys and P. de la Harpe (editors), Progress in Mathematics, volume 83, Birkhäuser, Boston, 1990.

[Tse68] G.S. Tseitin. Complexity of a derivation in the propositional calculus. *Zap. Nauchn. Sem. Leningrad Otd. Mat. Inst. Akad. Nauk SSSR*, 8:234–259, 1968. Also appeared in *Studies in Mathematics and mathematical Logic*, Part II, ed. A.O. Slissenko, 115–125.

[Ur87] A. Urquhart. Hard examples for resolution *Journal of the Association for Computing Machinery*, 34(1):209–219, 1987.

[Wei66] C.M. Weinbaum. Visualizing the word problem with an application to sixth group. *Pacific Journal of Mathematics*, 16:557–578, 1966.

A. Carbone
Institut des Hautes Études Scientifiques
35, route de Chartres, 91440 Bures-sur-Yvette, France
and
Laboratoire d'Algorithmique, Complexité et Logique
Université de Paris 12
e-mail: `carbone@ihes.fr`