PROCEEDINGS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 119, Number 1, September 1993

THE LENGTH OF ELEMENTS IN FREE SOLVABLE GROUPS

CARL DROMS, JACQUES LEWIN, AND HERMAN SERVATIUS

(Communicated by Ronald Solomon)

ABSTRACT. We examine the relationship between the complexity of the word problem for a presentation and the complexity of the problem of determining the length of a shortest word equivalent to a given word. Our main result is that the length of the element represented by a word in a free solvable group can be determined in polynomial time.

Let $\mathcal{P} = \langle X \mid R \rangle$ be a group presentation. Let F be free on X, let N be the normal closure of R in F, and let $G = F/N$ be the group presented by $\mathcal{P}$. A *word over* X is a finite string of symbols x and x^{-1} ($x \in X$). The *length* of a word w is the length of w regarded as a string. Clearly, each word over X represents a unique element of G. The *length* of an element $g \in G$ is defined to be the number of symbols in any *shortest* word representing g.

Let $\mathcal{C}$ be the Cayley graph of G with respect to the presentation $\mathcal{P}$. That is, $\mathcal{C}$ is the graph with vertex set G and (directed) edge set $E = G \times X$, where the edge $e = (g, x)$ goes from g to gx. We call g the *initial vertex* of e, denoted $\iota(e)$, and gx its *terminal vertex,* denoted $\tau(e)$. The generator x is called the *label* of e. The length of a group element g is then the distance (under the graph metric) from 1 to g in $\mathcal{C}$.

The *Word Problem* $\mathcal{W}(\mathcal{P})$ for a presentation $\mathcal{P}$ asks for an algorithm which decides, given any word w over X, whether w represents the identity of G, and its *Length Problem* $\mathcal{L}(\mathcal{P})$ asks for an algorithm which determines, given such a word w, what the length of the corresponding group element is.

It is clear that $\mathcal{W}$=$\mathcal{W}(\mathcal{P})$ is solvable if and only if $\mathcal{L}$=$\mathcal{L}(\mathcal{P})$ is; if $\mathcal{L}$ is solvable, then $\mathcal{W}$ is, as well, since the identity is the only element of length 0. In fact, $\mathcal{W}$ is no more complex than $\mathcal{L}$. On the other hand, if $\mathcal{W}$ is solvable, then so is $\mathcal{L}$; given a word w, one can construct, one after another, the spheres of the Cayley graph of G. Each time a new element e is constructed, one can solve $\mathcal{W}$ for $e^{-1}w$. Since, in general, the number of elements of a given length in a group is an exponential function of the length, this solution of $\mathcal{L}$ may be exponentially more complex than the solution of $\mathcal{W}$.

In fact, the following (metabelian) example of Parry [4] shows that it is possible for a presentation to have a word problem which is solvable in polynomial time and a length problem which is $\mathcal{NP}$–complete:

EXAMPLE: Let G be the wreath product $Z_2 \wr (Z \times Z)$, given by the standard infinite presentation

$$\langle t, a, b \mid [a, b] = 1, t^2 = 1, t^x t^y = t^y t^x \text{ for each } x, y \in \text{gp}\{a, b\}\rangle$$

Each element of G can be expressed in the form

$$t^{y_1} t^{y_2} \cdots t^{y_k} x$$

where $x \in \text{gp}\{a, b\}$ and $y_1, \ldots, y_n$ are distinct elements of $\text{gp}\{a, b\}$, and this expression is unique up to the order of the conjugating elements y_i. Now, it is clear that any word of length n in a, b and t can be rewritten in this form in at most $O(n^2)$ steps by commuting all the a's and b's to the right of all the t's. Since two such words represent the same element of G if and only if the respective lists of conjugating elements are permutations of one another and the values of "x" are the same, the word problem is solvable in time $O(n^2)$.

On the other hand, the length problem for G can be reduced to the *Euclidean Travelling Salesman Problem* in the integer lattice in the plane, the Cayley graph of the group $\langle a, b \mid [a, b] = 1\rangle$. An element

$$t^{y_1} t^{y_2} \cdots t^{y_k} x$$

can be represented by a path in the integer lattice which starts at 1, stops at each vertex y_i to "light a lamp," and then proceeds to x [4]. (It is for this reason that this group has been called a "Lamplighter Group" by J. Cannon.)

Finding a *shortest* word representing an element of the form

$$t^{y_1} t^{y_2} \cdots t^{y_m}$$

is thus equivalent to finding a shortest circuit in the integer lattice in the plane which begins at the origin and passes through each of the vertices corresponding to the y_i, a problem which is known to be $\mathcal{NP}$–complete [3]. □

However, if a group has polynomial growth with respect to the given generators and if the word problem can be solved in polynomial time, then the length problem is solvable in polynomial time, as well: suppose that the number of elements of length n is $O(n^k)$ for some positive integer k. Then having constructed the $(n-1)$–sphere of $\mathcal{C}$, one can construct the n–sphere in time $O(n^{2k} w(n))$, where $w(n)$ is the time needed to solve the word problem for two words of length $\leq n$; given a "standard" word W_g of length $n-1$ representing each element $g \in G$ of length $n-1$, form the words $W_g a$ and $W_g a^{-1}$ for each generator a. Each of the $O(n^k)$ new words must be compared to the $O(n^k)$ "standard" words of length $n-1$ and $n-2$, as well as to each of the other new words, resulting in $O(n^{2k})$ solutions of the word problem for words of length $\leq n$. Hence, the n–ball of $\mathcal{C}$ can be constructed in polynomial time, and since the word problem is solvable in polynomial time, so is the length problem.

In certain cases—free groups, for example—the length problem can be solved in polynomial time, even though the group grows exponentially. In fact, this is true of any hyperbolic group with exponential growth, since in this case there is a quadratic algorithm for computing geodesic representatives (see [2].)

The purpose of this note is to prove

Theorem 1 *Let $\mathcal{P} = \langle X \mid R \rangle$ be a presentation with F free on X and N the normal closure of R. If the length of the element of $G = F/N$ represented by a word w of length n can be determined in time $O(t(n))$, then the length of the element of $G_1 = F/N'$ represented by w can be determined in time $O(n^2 t(n))$.*

For example, if the word problem for F/N can be solved in polynomial time and F/N is infinite, then F/N' has exponential growth, but its length problem can be solved in polynomial time. Similarly, if $N^{(k)}$ denotes the *k–th* derived group of N, then the groups $F/N^{(k)}$ $(k \geq 2)$ all have exponential growth and polynomial length problem. In particular, the word problem for the free abelian group F/F' has a word problem solvable in linear time, and so we have

Corollary 1 *If F is freely generated by the finite set X, and if $F^{(k)}$ denotes the k–*th *derived group of F, then the length of the element of $F/F^{(k)}$ represented by a word over X of length n can be determined in time $O(n^{2k-1})$*

Let $\mathcal{C}$ be the Cayley graph of F/N. Then the fundamental group of $\mathcal{C}$ is naturally isomorphic to N—the label of any edge path in $\mathcal{C}$ can be interpreted as an element of F, and the path is closed precisely when this element belongs to N. Thus, the abelianized group N/N' is the first homology of $\mathcal{C}$, so a path P in $\mathcal{C}$ represents an element of N' if and only if it is homologically trivial.

Given a word w over X, let $\overline{w}$ denote the element of $G = F/N$ represented by w, and let P_w denote the edge–path in $\mathcal{C}$ beginning at the vertex 1 and labelled "w." Note that $\overline{w}$ is the ending vertex of P_w. We will define two functions $\epsilon_w : E \to \mathbb{Z}$ and $\nu_w : V \to \mathbb{Z}$, respectively the "edge numbering" and the "vertex numbering" associated to the path P_w.

For each edge $(g, x) \in E = G \times X$, we set $\epsilon_w(g, x)$ equal to the *net* number of times the edge (g, x) is traversed by P_w—each traversal from g to gx counting $+1$ and each traversal from gx to g counting -1. Since $\mathcal{C}$ is a one–dimensional complex, P_w represents the trivial homology class in $H_1(\mathcal{C})$ if and only if ϵ_w is identically 0.

Given two words w and u, a straightforward computation shows that for any edge (g, x) of $\mathcal{C}$,

$$\epsilon_{wu^{-1}}(g, x) \quad = \quad \epsilon_w(g, x) - \epsilon_u(\overline{u}\,\overline{w}^{-1} g, x) \tag{1}$$

For each vertex $g \in V = G$, define $\nu_w(g)$ equal to the net number of "entrances into" the vertex g of $\mathcal{C}$ along the path P_w. That is,

$$\nu_w(g) = \sum_{g=\tau(e)} \epsilon_w(e) - \sum_{g=\iota(e)} \epsilon_w(e)$$

(Note in particular that ν_w is completely determined by ϵ_w.)

It is clear that if w is any word, then ν_w satisfies "Kirchhoff's Law"; that is, either ν_w is identically 0 (if $\overline{w} = 1$), or $\nu_w(1) = -1$, $\nu_w(\overline{w}) = +1$ and $\nu_w(h) = 0$ for each other element $h \in G$ (if $\overline{w} \neq 1$.)

Lemma 1 *Two words w and u over X represent the same element of $G_1 = F/N'$ if and only if the functions ϵ_w and ϵ_u are identically equal.*

PROOF: If w and u are equivalent mod N', then $\overline{w} = \overline{u}$ and $\epsilon_{wu^{-1}} \equiv 0$. Thus, by equation (1), $\epsilon_w \equiv \epsilon_u$.

Conversely, if $\epsilon_w \equiv \epsilon_u$, then $\overline{w} = \overline{u}$, since the edge paths P_w and P_u must end at the same point. Therefore, $\epsilon_{wu^{-1}} \equiv 0$, by equation (1), so $P_{wu^{-1}}$ represents $0 \in H_1(\mathcal{C})$. That is to say, $w \equiv u \,(\mathrm{mod}\, N')$. □

Let $n : E \to \mathbb{Z}$ be an arbitrary integer–valued function on the edge set E of $\mathcal{C}$. Then we can define a function $\overline{n} : G \to \mathbb{Z}$ by setting

$$\overline{n}(g) = \sum_{g=\tau(e)} n(e) - \sum_{g=\iota(e)} n(e)$$

We will say that the function n is *geometric* if $\sum_{e\in E} |n(e)| < \infty$ and the associated function $\overline{n}$ on G satisfies Kirchhoff's Law. If n is geometric and $\overline{n}$ is zero, we will say n is *closed,* otherwise n is *open.*

It is clear that if w is any word over X, then ϵ_w is geometric, $\overline{\epsilon_w} = \nu_w$, and ϵ_w is closed if and only if $\overline{w} = 1$. Conversely, given an arbitrary "edge numbering" $n : E \to \mathbb{Z}$, we shall see that there is a word w over X with $\epsilon_w = n$ if and only if n is geometric, and that the corresponding edge path P_w is closed if and only if n is closed. Any word w with $\epsilon_w = n$ will be said to *realize* n (and we remark that w is *not* unique.)

Given $n : E \to \mathbb{Z}$, let $\mathrm{supp}(n)$ denote the subgraph of $\mathcal{C}$ consisting of all edges e with $n(e) \neq 0$, together with their endpoints, and let $\mathrm{supp}^+(n)$ denote the graph $\mathrm{supp}(n) \cup \{1\}$.

Lemma 2 *Let n be a geometric numbering on $\mathcal{C}$ with supp(n)* connected.

1. *If n is closed, then for each $v \in$ supp(n), there is a loop in $\mathcal{C}$ based at v which traverses each edge e exactly $n(e)$ times.*

2. *If n is open, then $1 \in$ supp(n), and there is a path in $\mathcal{C}$ beginning at 1 which traverses each edge e exactly $n(e)$ times.*

In particular, no edge is traversed in both *directions by such a path.*

PROOF: The proof of the Lemma is essentially the same as the proof that any connected graph in which the vertex degrees are all even possesses an Euler circuit. □

Note that the length of any such path is $\sum_{e\in E} |n(e)|$.

Let n be a geometric (closed or open) numbering on E. Let $s_1, s_2, \ldots, s_r$ be the connected components of $\mathrm{supp}^+(n)$, where s_1 is the component containing 1. It is clear that if n is open, then the unique vertex v with $\overline{n}(v) = +1$ lies in s_1.

For each $i, j \leq r$, let $d_{ij} = d_{ji}$ be the distance in $\mathcal{C}$ from s_i to s_j, and let p_{ij} be a path in $\mathcal{C}$ from s_i to s_j of length d_{ij}, where, for convenience, we take p_{ji} to be the inverse of p_{ij}. Let v_{ij} be the initial point of p_{ij}, so that the terminal point of p_{ij} is v_{ji}. (See Figure 1.)

Let K_r be a complete *undirected* graph with vertices $k_1, \ldots, k_r$ in 1–1 correspondence with the s_i, and assign the weight d_{ij} to the edge joining k_i and k_j. Let $T(n)$ be a *minimal–weight* spanning tree for K_r, and let $W(n)$ be the sum of the weights of the edges of $T(n)$—note that, while $T(n)$ is not uniquely determined by n, $W(n)$ is. It is clear that any path in $\mathcal{C}$ which realizes n must have length at least $\sum_{e\in E} |n(e)| + 2W(n)$. Any such path having this length will be called a *minimal realization* of n.

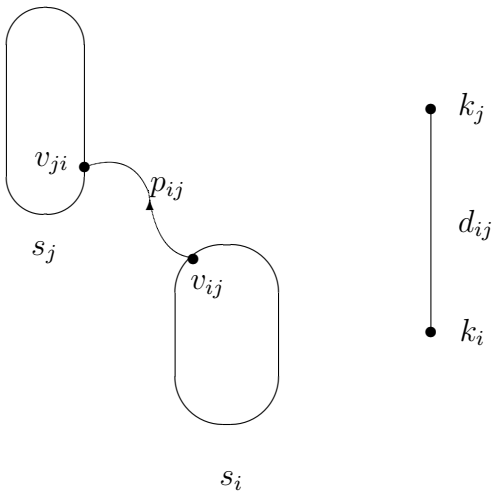

Figure 1:

Lemma 3 *If $n : E \to \mathbb{Z}$ is geometric, then there is a minimal realization of n beginning at the vertex 1 in $\mathcal{C}$.*

PROOF: The proof is by induction on r, the number of components of $\text{supp}^+(n)$. If $r = 1$, the result follows from Lemma 2.

If $r > 1$, choose a pendent vertex k_j of $T(n)$ such that $1 \notin s_j$, and let k_i be the vertex of $T(n)$ to which k_j is adjacent. Define new integer–valued functions n' and n'' on E as follows:

$$n'(e) = \begin{cases} n(e) & \text{if } e \notin s_j \\ 0 & \text{otherwise} \end{cases}$$

$$n''(e) = \begin{cases} n(e) & \text{if } e \in s_j \\ 0 & \text{otherwise} \end{cases}$$

It is clear that n' and n'' are geometric, n'' is closed, $\text{supp}(n'') = s_j$, and $\text{supp}^+(n')$ has $r-1$ components. Furthermore, $W(n') = W(n) - d_{ij}$. Thus, there is a minimal realization p' of n' beginning at 1, whose length is

$$\sum_{e \in E} |n'(e)| + 2(W(n) - d_{ij})$$

The vertex v_{ij} lies on the path p'. To get a path which realizes n, follow p' until v_{ij}, follow p_{ij} to $v_{ji} \in s_j$, follow a minimal realization of n'' beginning and ending at v_{ji} (which exists by Lemma 2, since n'' is closed), follow p_{ij}^{-1} back to v_{ij}, then proceed along p'. Clearly, the length of this path is $\sum_{e \in E} |n(e)| + 2W(n)$, so it is a minimal realization of n. □

Theorem 2 *Let w be a word over X. Then the length of the element of $G_1 = F/N'$ spelled "w" is $\sum_{e \in E} |\epsilon_w(e)| + 2W(\epsilon_w)$.*

The Algorithm.

In this section we will prove Theorem 1, which we restate here for convenience:

Theorem 1 *Let $\mathcal{P} = \langle X \mid R \rangle$ be a presentation with F free on X and N the normal closure of R. If the length of the element of $G = F/N$ represented by a word w of length n can be determined in time $O(t(n))$, then the length of the element of $G_1 = F/N'$ represented by w can be determined in time $O(n^2 t(n))$.*

PROOF: Let $\mathcal{C}$ denote the Cayley graph of G corresponding to the presentation $\mathcal{P}$. We will describe an algorithm to solve the length problem for G_1.

Let w be a word over X of length n and let p be the path at 1 labelled "w." For $i = 1, 2, \ldots, n$, let w_i be the *i–th* letter of w, and let w^i be the *i–prefix* of w—that is, $w^i = w_1 w_2 \cdots w_i$. Let $v_1, v_2, \ldots, v_k$ be the vertices visited by p. (Note that k will in general be $< n$, since p may visit a single vertex more than once.)

In the first part of the algorithm, we will scan w from left to right, and construct two matrices: D, whose ij–entry will be the distance in $\mathcal{C}$ from the vertex v_i to the vertex v_j, and S, whose ij–entry is the net number of traversals of the edge from v_i to v_j (0 if v_i and v_j are not adjacent.) Note that D is symmetric, and that for $i \neq j$, the ij– and ji–entries of S are additive inverses.

Suppose that we have read the prefix w^{i-1}, and that the matrices S and D have been constructed for the word w^{i-1}. Let $1 = v_1, v_2, \ldots, v_\ell$ be the vertices visited by the path at 1 labelled "w^{i-1}" in $\mathcal{C}$, and let v be the end vertex of the path at 1 labelled w^i. At the *i–th* step, we read w_i, and

1. solve $\mathcal{W}$ for $\mathcal{P}$ ℓ times to see if the current vertex v is new.
2. if v is new, set $v_{\ell+1} = v$, add row $\ell + 1$ and column $\ell + 1$ to D, and compute the entries in the new row and column by repeatedly solving the length problem for $\mathcal{P}$.
3. if v is new, add row $\ell + 1$ and column $\ell + 1$ to S, and set their entries to 0.
4. add or subtract 1 to the appropriate entries in S to reflect the fact that the edge from v_j to v has been traversed ± 1 times more.

Note that for each of the n iterations, the first two steps require at most $O(nt(n))$ steps each, the third requires at most $O(n)$ and the fourth a constant number of steps. Thus, the first part of the algorithm can be accomplished in time at most $O(n^2 t(n))$.

In the second part of the algorithm, we

1. identify the connected components of the graph with incidence matrix M obtained by setting each non–zero entry of S to 1.
2. compute the distances in $\mathcal{C}$ between these connected components, by examining D.
3. identify a minimal spanning tree for the complete graph whose edge weights are the numbers computed in the previous step.

Here, each step requires at most $O(n^2)$ steps (see, for example, [1].) Thus, the whole algorithm requires time at most $O(n^2 t(n))$. $\square$

Corollary 1 now follows from the fact that the length problem can be solved for $F/F' = F/F^{(1)}$ in time $O(n)$.

References

[1] A. Aho, J. Hopcroft, J. Ullman, "The Design and Analysis of Computer Algorithms," Addison–Wesley, Reading, 1974.

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

[3] C. H. Papadimitriou, The Euclidean traveling salesman problem is $\mathcal{NP}$–complete, Theoretical Computer Science **4** (1977), 237—244.

[4] W. Parry, Growth series of some wreath products, preprint.