PROCEEDINGS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 127, Number 8, Pages 2233–2243
S 0002-9939(99)04874-1
Article electronically published on April 9, 1999

REPRESENTATIONS OF INFINITE PERMUTATIONS BY WORDS (II)

RANDALL DOUGHERTY AND JAN MYCIELSKI

(Communicated by Ronald M. Solomon)

Dedicated to the memory of Roger C. Lyndon

ABSTRACT. We present an argument (due originally to R. C. Lyndon) which completes the proof of the following theorem: Every free group word which is not a proper power can represent any permutation of an infinite set.

1. INTRODUCTION

Let F be a free group, w a member of F, G any group, and g a member of G. We say that w *can represent* g *in* G iff there exists a homomorphism $h\colon F \to G$ such that $h(w) = g$. In other words, treating the generators of F as unknowns, one can solve the equation $w = g$ in G. For any set A, we denote by $S(A)$ the group of all permutations of A.

The purpose of this paper is to give an alternative version, with some additional consequences, of part of the proof of the following theorem of Lyndon and Mycielski.

Theorem. *If w is not a proper power (i.e., $w \neq v^n$ for $v \in F$, $n > 1$) and A is infinite, then w can represent every permutation $\pi \in S(A)$ in $S(A)$.*

Several cases of this theorem were proved earlier. D. Silberger [12] showed it for $w = x^p y^q$, $p, q \neq 0$, and asked whether it is true in general. M. Droste [3] showed it for $w = w_1 w_2$, where w_1 and w_2 have no variables in common, and for $w = x^p y^{-r} x^q y^r$ (and one can even require that the permutations used for x and y lie in certain specific conjugacy classes).

The cases of the theorem when the permutation π to be represented has infinitely many cycles of the same size (and possibly any other cycles) were proved in Mycielski [10]. In particular, if A is uncountable, then π must have $|A|$ cycles of the same size and the theorem holds. R. C. Lyndon [7] proved the rest of the theorem by handling the following cases:

(∗) A is countable and either π has arbitrarily large finite cycles or π has an infinite cycle.

The authors have simplified, somewhat, Lyndon's proof (using the same idea), and we present it here. Unlike the proof of the cases treated in Mycielski [10],

Received by the editors November 1, 1997.
1991 *Mathematics Subject Classification.* Primary 20F05.
The first author was supported by NSF grant number DMS-9158092 and by a fellowship from the Sloan Foundation.

which uses some deep results about the topology of Cayley complexes, the proof assuming (∗) is elementary; the present paper is self-contained. In the third section of this paper, we give modifications of the construction which produce representing permutations with additional properties; we discuss the problem of representing two permutations simultaneously by two given words; and we solve some problems stated in Mycielski [10].

Let us add that groups other than $S(A)$ have also been considered. Ore [11] showed that, if $n \geq 5$, then every element of the alternating group A_n is a commutator, and he proved the same for $S(A)$ for infinite A. Droste [4] showed that, for $p, q \neq 0$, if m is the largest squarefree divisor of pq and $n \geq 4m+1$, then $x^p y^q$ can represent every element of A_n. Brenner, Evans, and Silberger [1] improved this to $n \geq \max(5, \frac{5}{2}\log m)$. For related work on permutations and on other groups, such as $SL(n, \mathbf{R})$, see Ehrenfeucht et al. [5], Mycielski [8], and Mycielski [9].

2. Proof of the theorem under assumption (∗)

We will write actions of permutations on the right: applying permutation π to element a gives element $a\pi$. So a composition $\pi_1\pi_2$ indicates first applying π_1, then π_2.

Let F be a free group on the generators $x_1, \ldots, x_n$. (Examples will be given in terms of two generators, x and y. Actually, the general case can be reduced to the case of two generators [10], but this would not simplify the proof here.) Let $w = w(x_1, \ldots, x_n)$ be a word in the generators $x_1, \ldots, x_n$ which is not a proper power in F, and let π be a permutation of A satisfying (∗). We want to show that w represents π in $S(A)$; that is, there exist $\sigma_1, \ldots, \sigma_n \in S(A)$ such that $w(\sigma_1, \ldots, \sigma_n) = \pi$.

We will use the term *letter*, and the variable r, to denote a generator x_k or an inverse generator x_k^{-1}. So w is a product of letters, say $w = r_1 r_2 \cdots r_l$.

Clearly, we may assume that w is a reduced word. In fact, since words which can be mapped to each other by automorphisms of F (in particular, conjugate words) represent the same permutations, we may assume that w is cyclically reduced. (See Proposition 3 of Silberger [12].) Also, the result is trivial if w has length 1, so assume that w has length greater than 1. (Since w is not a power, it follows that w must use more than one of the generators.)

We will construct the desired solution in the form of a labeled directed graph D_∞, with vertex set V_∞. Each vertex will be labeled with one of the elements of A, and each such element will appear exactly once as a label, so the labeling will give a bijection between A and V_∞. Each edge will be labeled by one of the generators x_k. (There may be multiple edges with different labels joining the same pair of vertices.) For each k, every vertex will have exactly one x_k-edge leading from it and exactly one x_k-edge leading to it. This means that, for each k, the x_k-edges will induce a bijection from V_∞ to V_∞. Hence, we will be able to define $\sigma_k \in S(A)$ by the condition that the x_k-edge starting at the vertex labeled a ends at the vertex labeled $a\sigma_k$. The goal is to construct D_∞ so that the resulting permutations $\sigma_1, \ldots, \sigma_n$ satisfy $w(\sigma_1, \ldots, \sigma_n) = \pi$.

We will construct D_∞ by starting with a larger labeled digraph which includes unlabeled vertices and does not have enough edges to define the permutations fully. Then, repeatedly identifying vertices in this digraph (gluing them together), we end up with the desired configuration D_∞ in the limit.

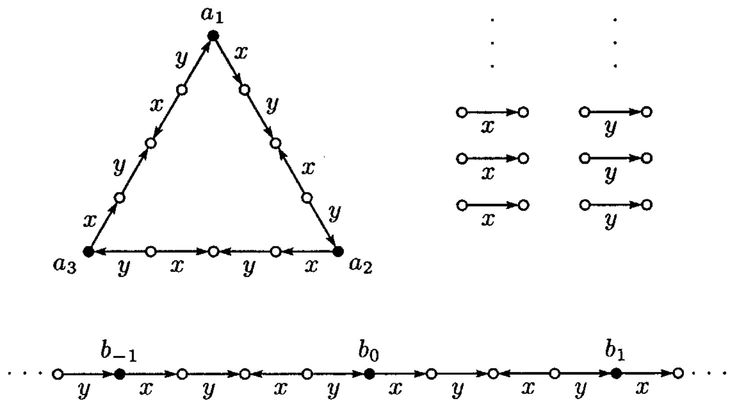

FIGURE 1. The initial labeled digraph for the construction.

The initial labeled digraph D_0, with vertex set V_0, is defined as follows. First, define "x_k^{-1}-edge from v to v'" to mean an x_k-edge from v' to v; this makes it convenient to talk about r-edges without specifying whether r is a generator or an inverse generator. Now, for each $a \in A$, V_0 contains distinct vertices $v_{a,i}$ for $0 \le i < l$ (where $w = r_1 r_2 \cdots r_l$), and vertex $v_{a,0}$ has label a in D_0; the other vertices $v_{a,i}$ are unlabeled. For each i such that $1 \le i < l$, there is an r_i-edge from $v_{a,i-1}$ to $v_{a,i}$; also, there is an r_l-edge from $v_{a,l-1}$ to $v_{a\pi,0}$. Finally, for each $k \le n$, put into D_0 countably infinitely many pairs of (unlabeled) vertices u, u' with no edges except for an x_k-edge from u to u'. (We will call these *free x_k-edges.*) This completes the definition of D_0.

For example, Figure 1 shows what D_0 looks like for the case of two generators x and y, where $w = xyx^{-1}y$ and π consists of a 3-cycle $(a_1\ a_2\ a_3)$ and an infinite cycle $(\dots b_{-1}\ b_0\ b_1\ b_2\ b_3 \dots)$.

The initial digraph D_0 has the following four properties, which will be maintained throughout the construction:

(1) Each $a \in A$ is a label for exactly one vertex.
(2) No vertex has more than one label.
(3) Each edge is labeled by exactly one of the generators x_k.
(4) Each vertex has at most one x_k-edge leading from it and at most one x_k-edge leading to it.

Property (4) holds in D_0 because w is reduced and cyclically reduced.

In D_0 or any other digraph satisfying (4), the following definition makes sense. For a vertex v and a letter r, let vr be the end point of the r-edge that starts at v, if there is one; otherwise, let vr be undefined. Then one can define $vrr' = (vr)r'$, and so on for longer words, so long as the requisite edges exist.

Another property that holds in D_0 and will hold in the later digraphs as well is:

(5) If vertex v has label a, then vw exists and has label $a\pi$.

It follows that vw^{-1} also exists and has label $a\pi^{-1}$, because, if v' is the vertex with label $a\pi^{-1}$, then $v'w$ must be the vertex with label a, namely v. It also follows that, if v is unlabeled, then vw and vw^{-1} are unlabeled if they exist.

One can view (5) as saying that D_0 gives a solution to our problem, except that the resulting 'permutations' are one-to-one maps not on A but on subsets of some set including A. The remaining part of the construction must get rid of the extraneous elements.

We will do this by a sequence of identifications. Suppose D is a digraph with vertex set V and labels of the sort that D_0 has. Then, given a function $g\colon V \to V'$, we can produce a labeled digraph D' on vertex set V': if $v \in V$ has label a in D, then $g(v)$ gets label a in D'; if there is an x_k-edge from u to v in D, then there is an x_k-edge from $g(u)$ to $g(v)$ in D'. (If there is also an x_k-edge from u' to v' where $g(u') = g(u)$ and $g(v') = g(v)$, then put only one x_k-edge from $g(u)$ to $g(v)$ in D'. However, if the edge from u' to v' has a different label, $x_{k'}$, then there will be multiple edges from $g(u)$ to $g(v)$ in D' with different labels.) We will call D' the *g-image of D*. If $g(v) = g(v')$, then v and v' have been identified in D'.

Properties (1) and (3) are preserved under all identifications (g-images), but one has to make extra assumptions about g in order for properties (2) and (4) to be preserved. Note that, if (4) holds in D and D', then, for any $v \in V$ and any word $\bar{w}$, if $v\bar{w}$ exists in D, then $g(v)\bar{w}$ exists in D' and is equal to $g(v\bar{w})$. Hence, in this case, property (5) is also preserved under the identification.

If we also have $g'\colon V' \to V''$, and D'' is the g'-image of D', then D'' is the $(g' \circ g)$-image of D.

Starting with D_0, we will produce a sequence of digraphs $D_0, D_1, D_2, \ldots$; for each $s > 0$, D_s will be the g_s-image of D_{s-1} for some surjection $g_s\colon V_{s-1} \to V_s$. Let $f_s\colon V_0 \to V_s$ be the composition $g_s \circ \cdots \circ g_1$ (and let f_0 be the identity map on V_0); then D_s will be the f_s-image of D_0.

Each identifying map g_s will be produced in the following way. Suppose we have vertices $u_1, \ldots, u_m$ and $v_1, \ldots, v_m$ in D_{s-1} such that:

- The vertices $u_1, \ldots, u_m, v_1, \ldots, v_m$ are all distinct.
- There is no j such that u_j and v_j are both labeled.
- For any j and any letter r, if there are r-edges starting at both u_j and v_j, then there is a j' such that these r-edges end at $u_{j'}$ and $v_{j'}$, respectively.

Then we can let $V_s = V_{s-1} \setminus \{u_1, \ldots, u_m\}$, and define $g_s\colon V_{s-1} \to V_s$ by $g_s(u_j) = v_j$ and $g_s(v) = v$ for all other v. We call this a *basic identification*, and say that D_s is obtained from D_{s-1} by identifying $u_1, \ldots, u_m$ with $v_1, \ldots, v_m$. The assumptions above imply that properties (2) and (4) (and hence all of (1)–(5)) are preserved by basic identifications.

Since we will be producing the digraphs D_s by a sequence of basic identifications, the following property will be preserved automatically:

(6) Each vertex in D_s has only finitely many f_s-preimages in D_0, and only finitely many vertices in D_s have more than one f_s-preimage.

Equivalently, only finitely many original vertices $v_0 \in V_0$ have been identified with another vertex by the map f_s. In particular, for each k, there will be infinitely many free x_k-edges in D_0 for which neither vertex is identified with anything else by f_s; these edges become free x_k-edges in D_s.

Let v_0 be a vertex in a labeled digraph satisfying (1)–(4). Then one can try to construct a path by starting at v_0, taking the r_1-edge to $v_1 = v_0 r_1$, then taking the r_2-edge to $v_2 = v_1 r_2$, and so on. If one reaches $v_l = v_{l-1} r_l = v_0 w$, then one can cycle again through w to define $v_{l+1} = v_l r_1$, $v_{l+2} = v_l r_1 r_2$, and so on. One may be able to continue this forever, or it may terminate when one reaches a vertex v_M at which no $r_{(M \bmod l)+1}$-edge starts. In either case, we call the resulting path $v_0, v_1, \ldots$ the *forward w-path rooted at v_0*.

Similarly, one can define the *backward w-path rooted at v_0* to be the forward w^{-1}-path rooted at v_0, which will be of the form $v_0, v_{-1}, v_{-2}, \ldots$ where $v_{-1} = v_0 r_l^{-1}$, $v_{-2} = v_0 r_l^{-1} r_{l-1}^{-1}$, $v_{-l} = v_0 w^{-1}$, and so on, continuing forever or until the required r_i^{-1}-edge fails to exist. The combined path $\ldots, v_{-2}, v_{-1}, v_0, v_1, v_2, \ldots$ is called the *w-path rooted at v_0*.

Note that a w-path does not double back on itself (use an edge in the forward direction and then immediately use this edge backward, or vice versa), because w is reduced and cyclically reduced.

In D_0, if vertex v has label a, then vw^j is defined for all integers j (and has label $a\pi^j$), so the w-path rooted at v is infinite in both directions. (Note that this infinite path may be periodic, a finite cycle traversed infinitely many times.) The final property we will need for the digraphs D_s is that these are the only infinite w-paths.

(7) For any unlabeled vertex v, the w-path rooted at v is finite (in both directions).

In other words, for any unlabeled v, vw^j exists for only finitely many integers j.

The basic identifications used to produce the digraphs D_s will be carefully chosen so as to preserve property (7) as well as (1)–(6). First, though, we have to see that this property holds in D_0. Clearly an infinite w-path cannot occur within the free x_k-edges (since w-paths cannot double back). The other components of D_0 are themselves w-paths (cyclic or doubly infinite) rooted at labeled vertices. We must see that, if we try to follow one of these paths 'out of phase' (i.e., starting at an unlabeled vertex) as a w-path, then we will only succeed for finitely many steps.

In fact, we will show that the forward w-path rooted at an unlabeled vertex u in one of these components of D_0 must have fewer than l edges. Recalling the notation for the elements of V_0, we must have $u = v_{a,i}$ for some $a \in A$ and some i, $0 < i < l$. If the forward w-path rooted at u lasts as long as l-edges, reaching from $u_0 = u$ to $u_l = uw$, then, since this path cannot double back on itself in D_0, it must follow the paths in D_0 either forward ($u_1 = v_{a,i+1}$, $u_2 = v_{a,i+2}$, and so on) or backward ($u_1 = v_{a,i-1}$, $u_2 = v_{a,i-2}$, and so on). In the former case, we must have $r_1 = r_{i+1}$, $r_2 = r_{i+2}$, $\ldots$, $r_{l-i} = r_l$, $r_{l-i+1} = r_1$, $\ldots$, $r_l = r_i$; this means that the word w is periodic with period $\gcd(i, l)$, which contradicts the assumption that w is not a proper power in F. In the latter case, we get $r_1 = r_i^{-1}$, $r_2 = r_{i-1}^{-1}$, and so on. If i is odd, this gives $r_{(i+1)/2} = r_{(i+1)/2}^{-1}$, which is impossible; if i is even, we get $r_{(i/2)+1} = r_{i/2}^{-1}$, which is impossible because w is reduced. Therefore, u_l cannot exist.

The same argument works for the backward w-path rooted at u, so this must have fewer than l edges as well. Therefore, D_0 satisfies all of the properties (1)–(7).

The sequence of identifications will be set up so as to take care of the following requirements:

(a) Every vertex eventually has a full set of edges attached: for each $v_0 \in V_0$ and each letter r, there is a stage s such that $f_s(v_0)r$ exists in D_s.

(b) Every vertex is eventually identified with a labeled vertex: for each $v_0 \in V_0$, there is a stage s such that $f_s(v_0)$ has a label in D_s.

Each of (a) and (b) can be thought of as a countably infinite list of requirements; (b) gives a requirement for each vertex v_0, and (a) gives a requirement for each pair (v_0, r). Since we have infinitely many stages $s = 1, 2, 3, \ldots$ available, it is easy to arrange a schedule (map requirements to stages) so that just one requirement needs to be taken care of at each stage. Note that, once a requirement has been taken care of at stage s ($f_s(v_0)$ has an r-edge or a label), it will remain true at stages $s' > s$.

Suppose that we have defined D_{s-1} so that properties (1)–(7) hold and, at stage s, we must take care of the requirement that $f_s(v_0)r$ exists. Let $v = f_{s-1}(v_0)$. If vr exists in D_{s-1}, we do not need to do anything, so just let g_s be the identity map on V_{s-1}. If vr does not exist, then find a free r-edge in D_{s-1}, say the edge from u to u', where $u, u' \neq v$ (this is possible because D_{s-1} has more than one free r-edge). Obtain D_s from D_{s-1} by identifying u with v. This is clearly a basic identification, and $f_s(v_0) = g_s(v) = g_s(u)$ is the starting point of an r-edge in D_s. Hence, D_s and f_s defined in this way satisfy (1)–(6). To see that (7) also holds in D_s, note that vertex u' is incident to only one edge in D_s, so it cannot lie on any w-path except possibly at an end of that path. Hence, if $\bar{v}$ is an unlabeled vertex such that the forward w-path rooted at $\bar{v}$ is infinite, then the forward w-path rooted at $\bar{v}w$ is also infinite and avoids u'. But D_s with u' omitted is actually included in D_{s-1}, so there is an infinite w-path in D_{s-1} rooted at $\bar{v}w$, which is an unlabeled vertex. This contradicts (7) for D_{s-1}. A similar contradiction arises if the backward w-path rooted at $\bar{v}$ is infinite, so (7) holds in D_s. Therefore, D_s has all of the desired properties.

Now suppose that the requirement to be taken care of at stage s is that $f_s(v_0)$ has a label. Again, if $f_{s-1}(v_0)$ already has a label in D_{s-1}, we can just let g_s be the identity map, so assume that $f_{s-1}(v_0)$ has no label. Let $v = f_{s-1}(v_0)$. By (7), we can define $\bar{v}_m, \bar{v}_{m+1}, \dots, \bar{v}_M$ to be the w-path rooted at v, where $m \leq 0 \leq M$ and $\bar{v}_0 = v$.

Claim. There is a labeled vertex $u_0 \in V_0$ such that, if $\dots, u_{-2}, u_{-1}, u_0, u_1, u_2, \dots$ is the w-path in D_0 rooted at u_0, then the vertices $u_{m-2l}, u_{m-2l+1}, \dots, u_{M+2l}$ are distinct, and none of them is identified with any other vertex or mapped to one of the vertices $\bar{v}_j$ by f_{s-1}.

Proof. Let N be the number of vertices in D_0 which are identified with another vertex or mapped to some $\bar{v}_j$ by f_{s-1}; by (6), this is finite. By (∗), π has a cycle which is so long that the corresponding w-path in D_0 includes $N+1$ disjoint subpaths of the form $u_{m-2l}, u_{m-2l+1}, \dots, u_{M+2l}$ with u_0 labeled; at least one of these subpaths contains none of the N 'bad' vertices. □

So find such a vertex $u_0 \in V_0$, and let $\bar{u}_j = f_{s-1}(u_j)$ for $m - 2l \leq j \leq M + 2l$. Then $\bar{u}_{m-2l}, \bar{u}_{m-2l+1}, \dots, \bar{u}_{M+2l}$ is part of the w-path rooted at $\bar{u}_0$ in D_{s-1}, and the w-path edges are the only edges incident to these vertices. Obtain D_s from D_{s-1} by identifying $\bar{u}_m, \bar{u}_{m+1}, \dots, \bar{u}_M$ with $\bar{v}_m, \bar{v}_{m+1}, \dots, \bar{v}_M$. Since $\bar{v}_m, \bar{v}_{m+1}, \dots, \bar{v}_M$ is the entire w-path rooted at the unlabeled vertex $v = \bar{v}_0$ (and hence all of the vertices $\bar{v}_{jl} = vw^j$ are unlabeled), this is a basic identification, so D_s satisfies (1)–(6). Now we will show that (7) also holds.

Suppose to the contrary that property (7) fails in D_s; say the forward w-path rooted at the unlabeled vertex $\tilde{v}$ is infinite. Then the forward w-path rooted at $\tilde{v}w^2$ (which is all but the first $2l$ edges of the preceding path) is also infinite. This latter path must contain at least one of the vertices $\bar{u}_{m-1}$ and $\bar{u}_{M+1}$, because D_s without these two vertices is included in D_{s-1}, which satisfies (7). So the forward w-path rooted at $\tilde{v}$ must include as a subpath one of the two finite paths $\bar{u}_{m-2l}, \dots, \bar{u}_{m-1}$

and $\bar{u}_{M+1}, \ldots, \bar{u}_{M+2l}$, either forward or backward. Say, for instance, that it includes $\bar{u}_{M+1}, \ldots, \bar{u}_{M+2l}$ traced forward. Then one of the vertices $\bar{u}_{M+i}$ for $1 \le i \le l$ must be $\tilde{v}w^j$ for some j, and we also have $\tilde{v}w^j r_1 = \bar{u}_{M+i+1}$, $\tilde{v}w^j r_1 r_2 = \bar{u}_{M+i+2}$, and so on to $\tilde{v}w^{j+1} = \bar{u}_{M+i+l}$. But vertices $u_{M+i}, \ldots, u_{M+i+l}$ have not been identified with any other vertices by f_s, so we must have $u_{M+i}w = u_{M+i+l}$ in D_0, and u_{M+i} is unlabeled (because $\tilde{v}w^j$ is unlabeled in D_s). This contradicts the earlier statement that a forward w-path in D_0 rooted at an unlabeled vertex must have length less than l. A similar contradiction arises in the other cases here, so the forward w-path rooted at $\tilde{v}$ must be finite. The same argument works for the backward w-path, so (7) holds for D_s.

This completes the definition of the digraphs D_s and mappings f_s and g_s. We must now see how to obtain a limit digraph D_∞ in which all of the identifications have been made.

Since we have sets V_s and mappings $g_s\colon V_{s-1} \to V_s$, we can define the direct limit in the usual way. (In fact, since the mappings g_s are surjective, the definition is a bit simpler than usual.) So we get a set V_∞, a surjection $f_\infty\colon V_0 \to V_\infty$, and mappings $g_{s,\infty}\colon V_s \to V_\infty$ such that $f_\infty = g_{s,\infty} \circ f_s$. Let D_∞ be the f_∞-image of D_0; this is the same as the $g_{s,\infty}$-image of D_s, for any s.

It remains to show that D_∞ has the properties initially specified for it. Since no two labeled vertices in D_0 were ever identified in D_s, each vertex in D_∞ has at most one label. Since every vertex of D_0 was eventually identified with a labeled vertex, each vertex of D_∞ has exactly one label. For each letter r, property (4) ensures that no vertex in D_s has more than one r-edge leading from it, so the same holds for D_∞; and requirements (a) ensure that each vertex of D_∞ actually does have an r-edge leading from it. Therefore, the permutations $\sigma_k \in S(A)$ are well-defined. Finally, property (5) holds in D_∞ because it holds in D_s, so the permuations σ_k satisfy $w(\sigma_1, \ldots, \sigma_n) = \pi$. This completes the proof of the theorem. □

In the cases treated in Mycielski [10] (where $(*)$ may fail), the representation of π by w is relatively simple. It consists of disjoint connected components, each of which is obtained by starting with the Cayley graph of a group given by a relation of the form $w^n = e$, and taking a certain natural cover followed by a homomorphism. In the present paper, the structure D_∞ is much more irregular. Moreover, there is a great deal of arbitrariness in the structure, because of the choices made at each stage of the construction.

As we will see in the next section, though, these choices can provide useful flexibility, allowing us to produce permutations σ_k satisfying $w(\sigma_1, \ldots, \sigma_n) = \pi$ with additional desirable properties. (Such flexibility is the reason that we introduced free x_k-edges, instead of attaching additional edges to the labeled vertices in the initial configuration D_0 so as to immediately take care of requirements (a) for these vertices.)

3. Extensions, remarks, and problems

1. Using single-cycle permutations. By adding extra requirements to the construction, one can ensure that the permutations in the representation of π by w have additional properties. For instance, one can show: If $(*)$ holds, and $w = w(x_1, \ldots, x_n)$ is a word of length $l > 1$ which is not a proper power in the free group F, then there exist $\sigma_1, \ldots, \sigma_n \in S(A)$ satisfying $w(\sigma_1, \ldots, \sigma_n) = \pi$ such that each σ_k is a single infinite cycle moving all elements of A. Furthermore, we may ensure that, for any $a \in A$, the $2n$ elements $a\sigma_k^{\pm 1}$ for $k = 1, \dots, n$ are distinct, unless $l = 2$ and π has a fixed point.

This can be achieved as follows. We can define the *x_k-path rooted at v* just as we did for w-paths (in fact, more easily). Now we impose two additional properties to be maintained during the construction:

(8) For all k and all vertices v, the x_k-path rooted at v is finite.
(9) For all vertices v and all distinct letters r_1 and r_2, $vr_1 \neq vr_2$.

Property (8) holds for D_0 because w must use more than one of the generators; property (9) also holds for D_0 unless $l = 2$ and π has a fixed point (in which case we do not impose this property). It is not hard to see that the identifications we performed in order to meet requirements (a) and (b) preserve these two new properties.

We also add a third list of requirements to be met during the construction:

(c) All x_k-paths are eventually joined together: for all k and all original vertices $v_0, v_0' \in V_0$, there is a stage s such that $f_s(v_0')$ lies on the x_k-path rooted at $f_s(v_0)$ in D_s.

There is still no problem arranging a schedule so that each stage takes care of one requirement and all requirements are eventually taken care of. Suppose that at stage s we need to take care of requirement (c) for k and v_0, v_0'. If $f_{s-1}(v_0')$ is already on the x_k-path rooted at $f_{s-1}(v_0)$, do nothing at stage s (let g_s be the identity map). Otherwise, let t be the last vertex on the x_k-path rooted at $f_{s-1}(v_0)$, and let t' be the first vertex on the x_k-path rooted at $f_{s-1}(v_0')$. Let (u_i', u_i) for $i = 1, \dots, l$ be distinct free x_k-edges in D_{s-1}. Construct D_s from D_{s-1} by identifying $t, u_1, \dots, u_l$ with $u_1', \dots, u_l', t'$. The effect of this is to link up the formerly separate x_k-paths into one path with l new intermediate edges. This clearly will take care of the specified requirement (c), and it is not hard to show that properties (1)–(9) are preserved. (For (7), note that a w-path cannot contain l consecutive x_k-edges or x_k^{-1}-edges.)

The new requirements (c) ensure that, in D_∞, there is just a single x_k-path containing all of the vertices, so σ_k is a cycle moving all elements of A. And property (9) ensures that $vr_1 \neq vr_2$ for all $v \in V_\infty$ and all letters $r_1 \neq r_2$, so the corresponding property holds for the permutations σ_k.

If $(*)$ fails, then it may not be possible to represent π by w using permutations which are infinite cycles moving all elements of the countable set A. In fact, this can fail even when w is just xy; see below. But one can ask for weaker conditions, and we do not know whether these conditions can be met assuming that A is countable but $(*)$ fails. For instance, can one always represent π by w using permutations such that the corresponding digraph is connected (weakly; that is, without regard to the orientation of the edges)?

2. Products of two cycles. A problem of Gale [6] asks whether every permutation of a countable set can be expressed as a product of two cycles. Gale notes that this is possible for a permutation of a finite set, and hence for a permutation of a countable set which only moves finitely many elements.

It turns out that one can modify the methods used here to solve Gale's problem. In fact, one can show: if π is a permutation of the countable set A which moves infinitely many elements of A, then π can be expressed as a product of two infinite cycles, each of which moves all elements of A. Some changes are needed in the argument, because the hypothesis here is weaker than $(*)$; instead of starting with a larger digraph and identifying vertices, one starts with a smaller digraph and adds edges. The full construction is given in Dougherty [2]. F. Galvin independently found a similar proof of this result.

Note that a permutation of the countably infinite set A which only moves finitely many elements of A need not be expressible as a product of two cycles moving all elements of A; in fact, there can be no such expression if the finite permutation is odd. Galvin has shown that any even permutation of an n-element set can be expressed as a product of two n-cycles, and it follows that any even finite permutation of A can be written as a product of two cycles moving all elements of A.

3. Representing two permutations simultaneously. Suppose that we have two words w_1, w_2 on the generators $x_1, \dots, x_n$, and we want to solve two equations $w_1 = \pi_1$ and $w_2 = \pi_2$ simultaneously. What are the conditions on w_1 and w_2 such that this would be possible for all $\pi_1, \pi_2 \in S(A)$?

First, suppose that π_1 and π_2 satisfy the analogue of $(*)$, which is:

$(**)$ A is countable and, for every finite set W of reduced words in the letters π_j and π_j^{-1}, there exists $a \in A$ such that the elements $a\omega$ for $\omega \in W$ are distinct.

One way of phrasing this is that π_1 and π_2 act locally freely on arbitrarily large neighborhoods in A.

The methods used here can be extended to show that the words w_1 and w_2 can simultaneously represent any pair of permutations (π_1, π_2) satisfying $(**)$, provided that w_1 and w_2 have the following properties:

(i) $w_1 w_2 \neq w_2 w_1$ in F.

(ii) The subgroup $\langle w_1, w_2 \rangle$ of F is *pure*; that is, if $u \in \langle w_1, w_2 \rangle$ and u is a k-th power in F, then u is a k-th power in $\langle w_1, w_2 \rangle$.

(iii) If $u_1, u_2 \in \langle w_1, w_2 \rangle$ and u_1, u_2 are conjugate in F, then they are conjugate in $\langle w_1, w_2 \rangle$.

Properties (ii) and (iii) are equivalent to: If T is the subtree of the Cayley graph of F (using the standard generators) spanned by $\langle w_1, w_2 \rangle$, then any out-of-phase copy of T (i.e., one not rooted at a member of $\langle w_1, w_2 \rangle$) has only a finite overlap with T. They are also equivalent to the statement that the subgroup $H = \langle w_1, w_2 \rangle$ of F has trivial intersection with any conjugate gHg^{-1} for $g \in F \setminus H$; such a subgroup is called (in other contexts) a *Frobenius complement*.

Furthermore, one can show that properties (i)–(iii) are necessary in order for (w_1, w_2) to represent any pair (π_1, π_2) satisfying $(**)$. This is clear for (i), since pairs (π_1, π_2) need not commute. If (ii) fails, let $u = u(w_1, w_2)$ be a counterexample of minimal length. Then u is not a proper power in $\langle w_1, w_2 \rangle$, so one can let $\rho \in S(A)$ be a single infinite cycle, and apply the theorem to find $\pi_1, \pi_2 \in S(A)$ such that $u(\pi_1, \pi_2) = \rho$. In fact, by adding extra requirements to the construction, one can ensure that π_1 and π_2 satisfy $(**)$. But (w_1, w_2) cannot represent (π_1, π_2) if u is a proper power in F, since ρ is not a proper power in $S(A)$. Similarly, if (iii) fails for $u_1(w_1, w_2)$ and $u_2(w_1, w_2)$, then one can find π_1 and π_2 satisfying $(**)$ such that $u_1(\pi_1, \pi_2)$ and $u_2(\pi_1, \pi_2)$ are not conjugate in $S(A)$. (We may assume that the cyclically reduced form of u_1 as a word in w_1, w_2 is no longer than that of u_2. Note that u_1 is not conjugate to u_2^{-1} even in F, because u_1 is conjugate to u_2 in F and no non-identity element of F is conjugate to its own inverse. Therefore, u_1 is not conjugate to any power of u_2 in $\langle w_1, w_2 \rangle$. Now let A be the set of right cosets of $\langle u_2\rangle$ in $\langle w_1, w_2\rangle$, and define $\pi_i \in S(A)$ by $a\pi_i = aw_i$. Then $u_2(\pi_1, \pi_2)$ has a fixed point but $u_1(\pi_1, \pi_2)$ doesn't.) Hence, (w_1, w_2) cannot represent (π_1, π_2).

The same results hold for larger simultaneous systems, when one wants to use words $(w_1, \ldots, w_m)$ to represent permutations $(\pi_1, \ldots, \pi_m)$; this is possible for all $\pi_1, \ldots, \pi_m$ satisfying $(**)$ if and only if $w_1, \ldots, w_m$ satisfy the versions of properties (i)–(iii) for m. In (ii) and (iii), just replace $\langle w_1, w_2\rangle$ with $\langle w_1, \ldots, w_m\rangle$; the new version of (i) is:

(i) $w_1, \ldots, w_m$ freely generate $\langle w_1, \ldots, w_m\rangle$ in F.

However, matters are more difficult if one wants to use w_1 and w_2 to represent *all* pairs (π_1, π_2), not just those satisfying $(**)$. Here are three examples where w_1 and w_2 satisfy (i)–(iii) but do not represent (π_1, π_2).

$$w_1 = xy, \quad w_2 = y^2x^2; \qquad \pi_1 = \text{identity} \neq \pi_2.$$
$$w_1 = x, \quad w_2 = y^2xy; \qquad \pi_1 = \text{identity}, \quad \pi_2 = (a\ b\ c).$$
$$w_1 = x, \quad w_2 = y^2xy^{-1}; \qquad \pi_1 = (a\ b), \quad \pi_2 = (b\ c).$$

One can use these to obtain additional necessary conditions on w_1 and w_2 in order for them to represent all pairs (π_1, π_2). The only sufficient condition the authors know of at present is $\langle w_1, w_2\rangle = \langle x, y\rangle$.

4. Intersections of Cayley cycles. We take this opportunity to note that the problem stated on page 239 of Mycielski [10] has a negative solution. Let G be the Cayley graph of a group presented by one relation, and assume that the relator word is reduced and cyclically reduced. By a *basic cycle*, we mean any cycle in G obtained by spelling the relator word of G, beginning at any vertex of G; if the starting vertex is a and the relator word is $w = r_1r_2\ldots r_l$, then the basic cycle is the closed curve

$$IL(a) = a, ar_1, ar_1r_2, \ldots, ar_1r_2\cdots r_{l-1}, a.$$

(By a theorem of Weinbaum [13], the basic cycles of one-relator groups are always simple closed curves.) Then *the intersection of two basic cycles of a one-relator group presentation may be disconnected.* For example, the presentation $\langle a, b; ababa = e\rangle$ has basic cycles $IL(e) = e, a, ab, aba, abab, e$ and $IL(a) = a, a^2, a^2b, a^2ba, a^2bab, a$. Since $a^2bab = e$ and $aba = a^2b$ (this follows from $ab = abababa = ba$), the intersection $IL(e) \cap IL(a)$ has two connected components, the edge $\{e, a\}$ and the singleton $\{a^2b\}$.

By a *membrane* we mean a closed two-dimensional cell whose boundary is a basic cycle. *There are one-relator group presentations in which the union of some pair of membranes is an orientable (cylindrical) strip or a Möbius strip.* Examples are $\langle a, b; abaabaa = e\rangle$ and $\langle a, b, c; (abcb^{-1})^2a = e\rangle$, respectively.

Acknowledgments

We would like to thank D. Silberger and R. Solomon for helpful comments.

References

1. J. L. Brenner, R. J. Evans, and D. M. Silberger, *The universality of words x^ry^s in alternating groups*, Proc. Amer. Math. Soc. **96** (1986), 23–28. MR **88a:**20047
2. R. Dougherty, *Products of two cycles (91-6(b)) — solution*, in Mathematical Entertainments column, D. Gale, ed., Math. Intelligencer **16** (2) (1994), 30–31.

3. M. Droste, *Classes of universal words for the infinite symmetric group*, Algebra Universalis **20** (1985), 205–216. MR **87d:**20049
4. ———, *On the universality of words for the alternating groups*, Proc. Amer. Math. Soc. **96** (1986), 18–22. MR **87c:**20063
5. A. Ehrenfeucht, S. Fajtlowicz, J. Malitz, and J. Mycielski, *Some problems on the universality of words in groups*, Algebra Universalis **11** (1980), 261–263. MR **81k:**20045
6. D. Gale, *Mathematical entertainments*, Math. Intelligencer **15** (3) (1993), 56–61.
7. R. C. Lyndon, *Words and infinite permutations*, Mots, Lang. Raison Calc., Hermès, Paris, 1990, pp. 143–152. MR **95c:**20006
8. J. Mycielski, *Can one solve equations in groups?*, Amer. Math. Monthly **84** (1977), 723–726.
9. ———, *Equations unsolvable in $GL_2(C)$ and related problems*, Amer. Math. Monthly **85** (1978), 263–265. MR **57:**9867
10. ———, *Representations of infinite permutations by words*, Proc. Amer. Math. Soc. **100** (1987), 237–241. MR **88c:**20044
11. O. Ore, *Some remarks on commutators*, Proc. Amer. Math. Soc. **2** (1951), 307–314. MR **12:**671e
12. D. M. Silberger, *Are primitive words universal for infinite symmetric groups?*, Trans. Amer. Math. Soc. **276** (1983), 841–852. MR **84c:**20011
13. C. M. Weinbaum, *On relators and diagrams for groups with one defining relation*, Illinois J. Math. **16** (1972), 308–322. MR **45:**6901

Department of Mathematics, Ohio State University, Columbus, Ohio 43210
E-mail address: rld@math.ohio-state.edu

Department of Mathematics, University of Colorado, Boulder, Colorado 80309
E-mail address: jmyciel@euclid.colorado.edu