Topology and its Applications 24 (1986) 243-252
North-Holland

FREE SUBGROUPS OF THE HOMEOMORPHISM GROUP OF THE REALS

Andreas BLASS and James M. KISTER
Department of Mathematics, University of Michigan, Ann Arbor, MI 48109-1003, USA

Received 22 April 1985
Revised 24 September 1985

After proving the known result that the homeomorphism group $\mathscr{H}(\mathbb{R})$ of the reals has a free subgroup of rank equal to the cardinality of the continuum, we apply similar techniques to give criteria for the existence of many (a comeager set in a natural complete metric topology) homeomorphisms independent of a given subgroup of $\mathscr{H}(\mathbb{R})$.

AMS(MOS) Subj. Class.: Primary 57S25, 54E52; Secondary 26A48

homeomorphism group of the reals
compact metric topology comeager set

In this note we give a short self-contained proof of a result attributed to Ehrenfeucht by Mycielski [3] that the homeomorphism group of the real line, $\mathscr{H}(\mathbb{R})$, contains a subgroup which is free of rank $2^{\aleph_0}$ (Theorem 1). Our methods allow us to show that, in many cases, a subgroup G of $\mathscr{H}(\mathbb{R})$ can be enlarged by the choice of an independent generator h to get a subgroup of $\mathscr{H}(\mathbb{R})$ isomorphic to the free product $G * \mathbb{Z}$ (Theorem 2). In fact 'most' (in the sense of Baire category) choices of h will work.

Our interest in the problem began with trying (and failing) to exhibit just two explicit elements of $\mathscr{H}(\mathbb{R})$ which generate a free subgroup. H. Friedman has asked whether the homeomorphisms $g(x)=x+1$ and $h(x)=x^3$ are such generators, and we suggest (for the intrepid) the following conjecture.

Conjecture. *If $h_\alpha:(0,1)\to(0,1)$ is the homeomorphism given by the quadratic function $h_\alpha(x)=x+\alpha(x^2-x)$, for each $\alpha\in(0,1)$, then $\{h_\alpha \mid \alpha\in(0,1)\}$ is a free set of generators of a subgroup of $\mathscr{H}((0,1))$.*

Here we have substituted the interval $(0,1)$ as a homeomorphic copy of $\mathbb{R}$ in order to easily describe a set of homeomorphisms. We shall use $(0,1)$ again in the rest of this paper in order to describe a complete metric, which goes back to [4], on the space of homeomorphisms. Then we use a Baire category argument and an idea of Mycielski to obtain Theorem 1. We have learned that D. Mauldin and R.

0166-8641/86/$3.50

Kallman also obtained the same result by a different method. M. Freedman has pointed out to us that, by lifting to $\mathbb{R}$ two Schottky homeomorphisms of the circle, one gets two explicit free generators, but this method does not produce more than a countable number of generators.[1]

Let $\mathcal{G}$ be the group of all homeomorphisms of $I=[0,1]$ leaving each endpoint fixed. This is easily seen to be isomorphic to the subgroup of $\mathcal{H}(\mathbb{R})$ consisting of orientation-preserving homeomorphisms. Define a metric d on $\mathcal{G}$ by:

$$d(g,h)=\sup_{x\in I}|g(x)-h(x)|+\sup_{x\in I}|g^{-1}(x)-h^{-1}(x)|,$$

where g and h are in $\mathcal{G}$. It is an easy exercise in uniform convergence to check that this is a complete metric for $\mathcal{G}$. The metric topology makes $\mathcal{G}$ a topological group.

Let F be the free group generated by the countable set $\{\gamma_1,\gamma_2,\ldots\}$. Then any reduced word, w, in F has a finite number of distinct γ_i appearing in it, and we write $w=w(\gamma_{i_1},\ldots,\gamma_{i_k})$ to indicate that the γ's shown are precisely the ones that occur in w. If $g_1,\ldots,g_k\in\mathcal{G}$, we write $w(g_1,\ldots,g_k)$ for the homeomorphism in $\mathcal{G}$ obtained by substituting each g_j for the corresponding γ_{i_j} in the word w. For example, if $w=w(\gamma_2,\gamma_7)=\gamma_2^{-1}\gamma_7^3\gamma_2$ and $g\in\mathcal{G}$, then $w(g,\mathrm{id})=g^{-1}\,\mathrm{id}^3\,g=\mathrm{id}$, but $w(\mathrm{id},g)=g^3$. Denote the k-fold Cartesian product of $\mathcal{G}$ by $\mathcal{G}^k$ and, for any $w\in F$, let $\mathcal{I}_w\subseteq\mathcal{G}^k$ be the subset defined by

$$\mathcal{I}_w=\{(g_1,g_2,\ldots,g_k)\in\mathcal{G}^k\mid w(g_1,g_2,\ldots,g_k)=\mathrm{id}\}.$$

Lemma 1. *For any reduced word w involving $k>0$ of the generators γ_i, the set $\mathcal{I}_w$ is closed and nowhere dense in $\mathcal{G}^k$.*

Proof. To see that I_w is closed, simply observe that the function from $\mathcal{G}^k$ to $\mathcal{G}$ sending $(g_1,g_2,\ldots,g_k)$ to $w(g_1,g_2,\ldots,g_k)$ is continuous, since $\mathcal{G}$ is a topological group. $\mathcal{I}_w$ is the preimage of the closed set $\{\mathrm{id}\}$ under this function. To show that I_w is nowhere dense, we prove the following stronger result; a similar but non-local construction was given, for other purposes, in [1].

Lemma 2. *For any reduced word $w=w(\gamma_{i_1},\ldots,\gamma_{i_k})\in F$ with $k>0$, any $(g_1,\ldots,g_k)\in\mathcal{G}^k$, any $x\in(0,1)$, and any $\varepsilon>0$, there exists $(g_1',\ldots,g_k')\in\mathcal{G}^k$ such that $d(g_i,g_i')<\varepsilon$ for $i=1,2,\ldots,k$, and $w(g_1',\ldots,g_k')(x)\neq x$.*

Proof. We first introduce some more notation. If $(g_1',\ldots,g_k')\in\mathcal{G}^k$ where the g_i' are distinct and if $w(g_1',\ldots,g_k')=h_p\circ h_{p-1}\circ\cdots\circ h_1$, where each h_i is g_j' or $(g_j')^{-1}$ for some j and adjacent h_i's are not inverses of each other, then we let $x_0=x$ and

[1] A. Ehrenfeucht asked (cf. [3, p. 47]) whether the automorphism group of a linear order must have a free subgroup of rank $2^{\aleph_0}$ provided it has one of rank 2. This question has recently been answered affirmatively by W. Charles Holland (Varieties of automorphism groups of orders, Trans. Amer. Math. Soc. 288 (1985) 755-763). Combined with Freedman's observation, this result provides another proof that the homeomorphism group of the real line has a free subgroup of rank $2^{\aleph_0}$.

$x_i = h_i(x_{i-1})$ for $i = 1, 2, \ldots, p$. To prove Lemma 2, we shall show, by induction on p, that the g'_j can be chosen close to the g_j and so that the x_i will all be distinct.

The case $p = 0$ is vacuous. Now assume that $p > 0$ and that we have found distinct $g'_1, \ldots, g'_k$ such that $d(g_j, g'_j) < \frac{1}{2}\varepsilon$ for $j = 1, \ldots, k$ and $x_0, \ldots, x_{p-1}$ are all distinct. We shall show that, if $h_p = g'_m$ or $h = g'^{-1}_m$, then we can alter g'_m slightly so that, when we replace g'_m by the alteration and recompute the x_i, then $x_0, \ldots, x_p$ are all distinct. Of course the previously computed x_i (for $i < p$) might be altered by this process if g'_m or its inverse occurs as h_i for some $i < p$. We shall, however, take precautions to ensure that our alteration of g'_m does not affect the values of x_i for $i < p$.

We assume that $h_p = g'_m$; the case of g'^{-1}_m can be reduced to the case we consider by interchanging γ_{i_m} and its inverse wherever they occur in w and replacing g'_m with its inverse.

We wish to find $g'' \in \mathscr{G}$ near g'_m such that

(a) if $h_i = g'_m$ for some $i < p$, then $g''(x_{i-1}) = g'_m(x_{i-1}) = x_i$,

(b) if $h_i = g'^{-1}_m$ for some $i < p$, then $g''^{-1}(x_{i-1}) = g'^{-1}_m(x_{i-1}) = x_i$, and

(c) $g''(x_{p-1})$ is distinct from $x_0, \ldots, x_{p-1}$.

First note that the image of x_{p-1} under g'' is not specified by (a) or (b), for this could happen only if (b) applied to $i = p-1$, which means that $h_{p-1} = g'^{-1}_m$. This cannot occur since $h_p = g'_m$ and w is a reduced word.

Now let $P = \{0 = z_0 < z_1 < \cdots < z_n = 1\}$ be a fine partition of I containing $\{x_0, \ldots, x_{p-1}\}$. Choose y close to $g'_m(x_{p-1})$ but not in $\{x_0, \ldots, x_{p-1}\}$, and define g'' to be the unique map which is linear on each subinterval $[z_{i-1}, z_i]$, agrees with g'_m on $P - \{x_{p-1}\}$, and sends x_{p-1} to y. It is not hard to see that if P is sufficiently fine and y is sufficiently close to $g'_m(x_{p-1})$ then g'' is a homeomorphism satisfying $d(g'_m, g'') < \frac{1}{2}\varepsilon$ and conditions (a), (b), and (c). We complete the induction by replacing g'_m with g''. This finishes the proof of Lemma 2 and hence the proof of Lemma 1. □

Before turning to our main application of Lemma 1, we mention, at the referee's suggestion, an immediate consequence of this lemma. Most, in the sense of Baire category, k-tuples in $\mathscr{G}^k$ freely generate free subgroups of $\mathscr{G}$ of rank k. Indeed, those that do not constitute the union of the countably many nowhere dense sets $\mathscr{I}_w$, where w ranges over words involving k generators.

Theorem 1. *Let $\mathscr{G}$ have the complete metric topology defined above. There exists a Cantor set $C \subseteq \mathscr{G}$ whose elements satisfy no non-trivial group relations. Hence the subgroup of $\mathscr{G}$ generated by C is a free group of rank $2^{\aleph_0}$.*

Proof. Let $w_1, w_2, \ldots$ be a listing of all the non-trivial reduced words in the free group F generated by $\{\gamma_1, \gamma_2, \ldots\}$. We use an idea of Mycielski [2] and recursively define open subsets $V(\sigma)$ of $\mathscr{G}$, indexed by all finite sequences of zeros and ones, and satisfying:

(a) if σ and τ are distinct sequences of the same length, then $V(\sigma)$ and $V(\tau)$ are disjoint,
(b) if σ is a proper initial segment of τ, then $\mathrm{cl}(V(\tau)) \subseteq V(\sigma)$,
(c) if σ has length $n > 0$, then $0 < \mathrm{diam}(V(\sigma)) < 1/n$, and
(d) if $g_1, \ldots, g_k$ are taken from $k > 0$ distinct sets of the form $V(\sigma)$ with σ's of length n, and if, for some $j < n$, w_j has exactly k variables, then $w_j(g_1, \ldots, g_k) \neq \mathrm{id}$ in $\mathcal{G}$.

To start the recursion, let $V(\emptyset)$, where $\emptyset$ is the empty sequence, be any open subset of $\mathcal{G}$ not containing the identity. Condition (d) is easily verified, and the other three conditions are vacuous in this case.

Now assume that $V(\sigma)$ has been defined for every σ of length n in such a way that conditions (a) through (d) hold. In each such $V(\sigma)$ choose two disjoint open balls U_0 and U_1 whose closures are contained in $V(\sigma)$ and whose diameters are smaller than $1/(n+1)$. This is possible because $\mathcal{G}$ has no isolated points. View these U_i as preliminary values of $V(\tau)$ for the two one-term extensions τ of the sequence σ. Then (a), (b) and (c) continue to hold, but we must shrink the U_i further to achieve (d). This takes a finite number of steps, one step for each tuple $(w, \tau_1, \ldots, \tau_k)$ where $w = w_j$ for some $j < n+1$, k is the number of variables in w, and the τ_i are distinct dyadic sequences of length $n+1$. Such a step is taken to ensure that (d) holds for this w_j when each g_i is chosen from the corresponding $V(\tau_i)$.

At the step corresponding to $(w, \tau_1, \ldots, \tau_k)$, suppose the balls which served as the preliminary values of the $V(\tau_i)$ have already been shrunk to balls T_i. Since $\mathcal{I}_w$ is closed and nowhere dense, by Lemma 1, we can find open balls $S_1, \ldots, S_k$ such that

$$S_1 \times S_2 \times \cdots \times S_k \subseteq (T_1 \times T_2 \times \cdots \times T_k) - \mathcal{I}_w.$$

Replace each of the balls T_i by the corresponding S_i; clearly this makes (d) true for the cases being treated at this step. Thus, when all the steps have been completed, we have all four of the induction hypotheses satisfied for all sequences of length at most $n+1$. This completes the inductive construction of the $V(\sigma)$'s.

If we set $C = \bigcap_n (\bigcup_{\sigma \text{ of length } n} V(\sigma))$, then requirements (a), (b) and (c) together with the completeness of $\mathcal{G}$ imply that C is a Cantor set. Every k-tuple of distinct elements of C satisfies the hypothesis of (d) for all sufficiently large n and therefore, by virtue of (d), is not sent to id by any non-trivial word w. This completes the proof of Theorem 1. $\square$

We now take up the problem of enlarging a prescribed subgroup G of homeomorphisms by finding an independent generator h, i.e., a homeomorphism such that the subgroup generated by G and h is isomorphic to the free product $G * \mathbb{Z}$, the isomorphism being the identity on G and sending h to the generator 1 of $\mathbb{Z}$. An equivalent condition on h is that it satisfy no non-trivial group-theoretic equations with coefficients in G. We begin with some examples showing that, even when G

is cyclic, finding an independent generator is not as easy as one might think; in particular, the set of independent generators need not be dense.

Example 1. Let g be a non-trivial homeomorphism in $\mathscr{G}$ that fixes pointwise a neighborhood of 0 and a neighborhood of 1, and let G be the subgroup of $\mathscr{G}$ that it generates. Let $[a, b]$, with $0 < a < b < 1$, contain all the points moved by g. Choose $h_0 \in \mathscr{G}$ with $h_0[a, b]$ disjoint from $[a, b]$. (A piecewise linear h_0 will do.) Since there is a positive distance between $[a, b]$ and $h_0[a, b]$, every $h \in \mathscr{G}$ sufficiently near h_0 will have the property that $h[a, b]$ is disjoint from $[a, b]$. Since all points moved by g are in $[a, b]$, all points moved by hgh^{-1} must be in $h[a, b]$, and it follows from the disjointness of these intervals that g commutes with hgh^{-1}. Thus, we have a neighborhood of h_0 none of whose members h is independent of G.

Example 2. Again let g be a non-trivial homeomorphism in $\mathscr{G}$, but assume only that it fixes pointwise a neighborhood of 0, so all the points it moves lie in $[a, 1]$ for some $a > 0$. Let h_0 be an orientation-reversing element of $\mathscr{H} = \mathscr{H}([0, 1])$ whose unique fixed point is $< a$. Then the same is true of every $h \in \mathscr{H}$ sufficiently close to h_0. For such h, all the points moved by hgh^{-1} are in $h[a, 1] = [0, h(a)]$ which is disjoint from $[a, 1]$. So, as in Example 1, g commutes with hgh^{-1}, and therefore no h near h_0 is independent of the subgroup G generated by g.

Remark. If the g in Example 2 does not fix any neighborhood of 1 pointwise (so it is not covered by Example 1) then Theorem 2(a) below shows that the homeomorphisms h independent of g are dense in $\mathscr{G}$, so the use of an orientation-reversing h in Example 2 is essential.

A subset of $\mathscr{H} = \mathscr{H}([0, 1])$ or $\mathscr{G}$ (or indeed any complete metric space) is called *comeager* if it includes the intersection of some countable collection of dense open sets; equivalently, a set is comeager if its complement is of first category. In particular, the Baire category theorem asserts that comeager sets are dense.

Theorem 2. (a) *If $g \in \mathscr{G}$ then $\{h \in \mathscr{G} \mid g$ and h are free generators$\}$ is comeager in $\mathscr{G}$ if and only if g has non-fixed points arbitrarily close to at least one endpoint of I.*

(b) *If $g \in \mathscr{H}$ then $\{h \in \mathscr{H} \mid g$ and h are free generators$\}$ is comeager in $\mathscr{H}$ if and only if g^2 has non-fixed points arbitrarily close to each endpoint of I.*

(c) *If G is a countable subgroup of $\mathscr{G}$ and if the only element of G that pointwise fixes a non-degenerate interval is the identity, then $\{h \in \mathscr{G} \mid h$ is independent of $G\}$ is comeager in $\mathscr{G}$. The same is true with $\mathscr{H}$ in place of $\mathscr{G}$.*

Proof. Example 1 gives the 'only if' part of (a), and Example 2, with g replaced by g^2, gives the 'only if' part of (b). Indeed, the examples show that $\{h \mid g$ and h

are free generators} will not even be dense, let alone comeager, in $\mathscr{G}$ or $\mathscr{H}$ unless the condition in (a) or (b) (respectively) is met.

To prove the 'if' part of (a), let $g \in \mathscr{G}$ have non-fixed points arbitrarily near 1. Let $w = w(\gamma_1, \gamma_2)$ be a non-trivial reduced word in the free group generated by γ_1 and γ_2. It suffices to show that

$$\mathscr{N}_w = \{h \in \mathscr{G} \mid w(g, h) \neq \mathrm{id}\}$$

is open and dense in $\mathscr{G}$, since

$$\{h \in \mathscr{G} \mid g \text{ and } h \text{ are free generators}\} = \bigcap_w \mathscr{N}_w.$$

$\mathscr{N}_w$ is open because its complement is a section of the closed set $\mathscr{I}_w$ of Lemma 1. To see that $\mathscr{N}_w$ is dense, let $h \in \mathscr{G}$ and $\varepsilon > 0$. We must find $h'' \in \mathscr{N}_w$ such that $d(h, h'') < \varepsilon$. We use an argument similar to that in Lemma 2, except that this time we are not allowed to alter g, so we must choose x judiciously.

As a preliminary step, we modify the given h slightly (by $<\frac{1}{2}\varepsilon$) to an h' that is the identity on a neighborhood of 1. To do this, first find, since h preserves orientation, a point $(p, h(p))$ on the graph of h such that both p and $h(p)$ are within $\frac{1}{2}\varepsilon$ of 1. Then choose c smaller than 1 but larger than both p and $h(p)$. Replace the part of the graph of h from $(p, h(p))$ to $(1, 1)$ with the broken line segment from $(p, h(p))$ to (c, c) to $(1, 1)$. The result is the graph of the desired h'. We note, for use in the proof of (b) later, that if h had been orientation-reversing, we could have modified it slightly to agree with the function $1 - x$ near both endpoints of I. Furthermore, there is nothing special about the identity or $1 - x$ here; any orientation-preserving (resp. orientation-reversing) homeomorphism f could be used instead. Indeed, if we modify $f^{-1}h$ slightly (how slightly depends on the modulus of continuity of f) to get a k that is the identity near the endpoints, then fk is a slight modification of h that agrees with f near the endpoints.

Returning to the proof of (a), we seek an h'', within $\frac{1}{2}\varepsilon$ of our h' (that fixes $[c, 1)$ pointwise), such that $w(g, h'') \neq \mathrm{id}$ for a certain specified word w. We consider two cases.

Case 1. 1 is a limit of fixed points of g.

Since 1 is also, by hypothesis, a limit of non-fixed points of g, we can find, successively, points a, x_0, b such that $c < a < x_0 < b < 1$, a and b are fixed by g and x_0 is not. As the fixed point set is closed, we can increase a and decrease b, if necessary, to arrange in addition that g has no fixed points in the open interval (a, b). Note that (a, b) is invariant under both g and h'. We shall alter h' only on (a, b) to obtain an h'' with $d(h', h'') < \frac{1}{2}\varepsilon$ and $w(g, h'') \neq \mathrm{id}$ and in fact $w(g, h'')(x_0) \neq x_0$.

To construct h'' we proceed as in the proof of Lemma 2 with minor modifications. To make the notation agree with that lemma, let g_1 and g_2 be g and h' and let $(0, 1)$ there be replaced by (a, b) here. For g_2' near g_2, let $w(g_1, g_2') = h_p \circ h_{p-1} \circ \cdots \circ h_1$ and $x_i = h_i(x_{i-1})$ as before. We intend to modify g_2' slightly so as to make all the x_i

distinct. As before, we use induction on p, but we need to carry along a stronger induction hypothesis to compensate for our inability to modify g_1. In addition to requiring the x_i to be distinct, we require that, whenever $h_i = g_2'$ or $g_2'^{-1}$, then $x_i = h_i(x_{i-1})$ is in a different g_1-orbit from $x_0, x_1, \ldots, x_{i-1}$. Thus, if the next few functions $h_{i+1}, \ldots, h_{i+k}$ in the composition defining $w(g_1, g_2')$ are g_1 (or g_1^{-1}), then the resulting $x_{i+1}, \ldots, x_{i+k}$ will be distinct from $x_0, \ldots, x_{i-1}$. They will also be distinct from x_i and from each other because g_1 acts freely on the interval (a, b) in which we are working. Thus, there will be no need to modify g_1 as in the proof of Lemma 2.

To show that this stronger induction hypothesis can be maintained, consider, as in Lemma 2, the case that $h_p = g_2'$ and that g_2' satisfies the induction hypothesis that $x_0, \ldots, x_{p-1}$ are distinct and that each $h_i = g_2'$ or $g_2'^{-1}$ among $h_1, \ldots, h_{p-1}$ leads to a new g_1-orbit. We modify g_2' slightly so that x_i is unaltered for $i \leq p-1$ but $g_2'(x_{p-1})$ is in a different g_1-orbit from $x_0, \ldots, x_{p-1}$. This is done exactly as in the proof Lemma 2, except that the point y in that proof must be chosen outside the g_1-orbits of $x_0, \ldots, x_{p-1}$. This new requirement excludes only countably many possible values of y, so there is no difficulty in finding an appropriate y.

Case 2. 1 is not a limit of fixed points of g.

In this case, we have an interval $(d, 1)$ on which g acts freely; d is either the largest fixed point of g or 0 if g has no fixed point. We assume, without loss of generality, that $c > d$ (where c is still the left end of the interval $[c, 1)$ pointwise fixed by h'). We would like to proceed as in Case 1, using the fixed-point-free invariant interval $(d, 1)$ for g as we used (a, b) there. Unfortunately, if some x_i is too close to the left end, d, of this interval, and $h_{i+1} = g_2'$, we may find that x_{i+1}, which is near $h'(x_i)$, is outside $(d, 1)$. (This problem did not arise in Case 1, since the interval (a, b) was pointwise fixed by h'.) To be safe from this difficulty, we work in $(c, 1)$, where h' is the identity. But this interval is not invariant under g, so again the x_i sequence may escape from it. To avoid this, we choose x_0 (which in Case 1 could be any element of (a, b)) to be so close to 1 that there is no danger of the x_i's getting smaller than c. Specifically, if there are N occurrences of g or g^{-1} in $w(g, h')$, we require x_0 to be so close to 1 that all of $g^{-N}(x_0)$, $g^{-N+1}(x_0), \ldots, g^{N-1}(x_0)$, $g^N(x_0)$ are larger than c. (This is achievable since g is continuous and order-preserving.) Now proceed as in Case 1, making the modifications of h' so slight that all x_i remain in $(c, 1)$ at all stages of the induction. This completes the proof of (a).

Before turning to (b), we point out that the proof of (a) differs from Lemma 2 in essentially two respects. One is the strengthened induction hypothesis, which, as we saw, is not difficult to satisfy since it excludes only countably many potential values of y at each stage of the induction. The second difference is the need to work in a region where g has no fixed points; this is essential for keeping x_i distinct from $x_{i+1}, \ldots, x_{i+k}$ when $h_{i+1} = \cdots = h_{i+k} = g_1$, since we cannot modify g_1. Most of the work in the proof of (a) was needed in order to keep the x_i-sequence within such a fixed-point-free region. The same will be true in the proof of (b), and a bit more work will be needed. Instead of presenting the proof of (b) in detail, we indicate only the additional arguments needed to avoid fixed points of g; the rest of the proof is as in (a).

For orientation-preserving g, a point x not fixed by g will not be fixed by any power of g (for if $x < g(x)$, say, then as g preserves order $g(x) < g^2(x) < g^3(x) < \cdots$), but this is not so for orientation-reversing g. Indeed, such a g has a unique fixed point, which is also the unique fixed point of each of its odd powers, but the even powers may have more fixed points, and we must avoid these points as well. Of course, as g^2 preserves orientation, it suffices to avoid the fixed points of g^2.

We turn now to the proof of the 'if' part of (b). We are given $g \in \mathscr{H}$ such that g^2 has non-fixed points arbitrarily close to each endpoint of I, we are given a non-trivial reduced word w, and we are given $h \in \mathscr{H}$. We seek an h'' near h such that $w(g, h'')$ moves some x_0. If g and h are both orientation-preserving, we use the proof of part (a). We consider the remaining cases, subdividing them according to how many endpoints of I are limits of fixed points of g^2.

Case 1. g reverses orientation, and g^2 has fixed points arbitrarily near both 0 and 1.

As in the preliminary step in part (a), modify h slightly to get an h' that agrees, on little intervals $(0, c_0)$ and $(c_1, 1)$ with id if h preserves orientation and with g if h reverses orientation. Then, as in Case 1 of part (a), find an invariant fixed-point-free interval I_1 of g^2 so close to 1 that $I_1 \subseteq (c_1, 1)$ and $I_0 = g(I_1) \subseteq (0, c_0)$. Then $I_0 \cup I_1$ is invariant under both g and h' and contains no fixed points of g^2, hence no fixed points of any power of g. So we can proceed as in Case 1 of part (a), modifying h' only within this set.

Case 2. g reverses orientation and neither 0 nor 1 is a limit of fixed points of g^2.

Modify h slightly to agree near 0 and 1 with id or with $1-x$ (according as h preserves or reverses orientation), and let $(0, 1-d)$ and $(d, 1)$ contain no fixed points of g^2. Then $(0, 1-d) \cup (d, 1)$ contains no fixed points of g, and we keep the x_i sequence within this set by choosing x_0 close enough to 1 (or to 0), as in Case 2 of part (a).

The two cases just discussed cover all the possibilities for orientation-reversing g, since if one endpoint is a limit of fixed points, a_i of g^2, then the other endpoint is the limit of the fixed points $g(a_i)$ of g^2. It remains, therefore, to consider the cases where g preserves orientation and h reverses it. Note that in these cases, the fixed points of g^2 are the same as those of g, so we omit the squaring.

Case 3. g preserves orientation and has fixed points arbitrarily near each endpoint of I; h reverses orientation.

As in Case 1 of part (a), there are invariant fixed-point free intervals I_0 (resp. I_1) arbitrarily near 0 (resp. 1). Modify h slightly (as before) to get an h' that interchanges one of the intervals I_0 with one of the intervals I_1. Then proceed, as in Case 1 of part (a), to modify h' within the set $I_0 \cup I_1$, which is invariant under both g and h' and contains no fixed points of g.

Case 4. g preserves orientation and neither 0 nor 1 is a limit of fixed points; h reverses orientation.

Arrange for $h'(x)$ to be $1-x$ for x near the endpoints, and proceed as in Case 2 of part (a).

Case 5. g preserves orientation and has 0 but not 1 as a limit of fixed points; h reverses orientation.

We assume that $g(x)>x$ for x near 1; otherwise work with g^{-1}. Choose a and b near 1 (nearer than any fixed point) so that $a<b<g(a)$. Then the interval $I_0=[a, b]$ and its images $I_1=g[a, b]$, $I_2=g^2[a, b], \ldots$ are disjoint and approach 1 monotonically. (They approach 1 because otherwise the supremum of their union would be a fixed point closer to 1 than a and b are.) At the other end of the interval, where g has fixed points, choose a sequence $J_0, J_1, \ldots$ of disjoint invariant intervals whose interiors are fixed-point-free for g, approaching 0 monotonically. Then modify h to get an h' that interchanges each I_n with J_n; this is a slight modification provided all these intervals are close enough to the endpoints of I. Now work in the union of all the I_n's and J_n's. This union contains no fixed points of g, is invariant under h', and is mapped into itself by g, though not by g^{-1}. To avoid having the x_i sequence leave the union, start with $x_0 \in I_N$ (or J_N) for N larger than the length of the word w, essentially as in Case 2 of part (a).

These cases, and the symmetric one to Case 5 with 0 and 1 interchanged, exhaust the possibilities, so (b) is established.

Finally, we prove (c). Let G satisfy the hypothesis of (c). For each element $w \in G * \mathbb{Z}$ and each $h \in \mathscr{G}$, let $w(h)$ be the result of replacing in w each $n \in \mathbb{Z}$ with h^n, i.e. the image of w under the unique homomorphism $G * \mathbb{Z} \to \mathscr{G}$ that is the identity on G and sends 1 to h. Let $\mathcal{N}_w = \{h \in \mathscr{G} \mid w(h) \neq \mathrm{id}\}$. We must show that the set

$$\{h \in \mathscr{G} \mid \text{for all } w \neq 1,\ w(h) \neq \mathrm{id}\} = \bigcap_{w \neq 1} \mathcal{N}_w$$

is comeager, so we show that each $\mathcal{N}_w$ is open and dense. Openess is immediate, as before, since $w(h)$ is a continuous function of h. The proof of density is exactly as in the preceding parts of the proof. We view $w(h)$ as a composite $h_p \circ h_{p-1} \circ \cdots \circ h_1$ where each h_i is h or h^{-1} or an element of G. We wish to choose x_0 and slightly modify h so that the points x_i defined by $x_i = h_i(x_{i-1})$ are all distinct and, whenever $h_i = h$ or $h_i = h^{-1}$, x_i lies in a different G-orbit from $x_0, \ldots, x_{i-1}$ and is not fixed by any non-identity element of G. As before, this is done by induction on p. At each stage, the new value (called y in previous proofs) of h' must be chosen to avoid finitely many G-orbits and the fixed point set of every $g \in G - \{\mathrm{id}\}$. But G is countable, so finitely many G-orbits contain only countably many points. And each $g \in G - \{\mathrm{id}\}$ has a fixed-point-set that is closed and nowhere dense, so the union of the fixed-point-sets is of first category. That leaves a comeager set of y's that can be used, so the induction can proceed and the proof is complete. □

The referee has pointed out that the group generated by the lifts to $\mathbb{R}$ of two Schottky homeomorphisms of the circle satisfies the hypotheses of Theorem 2(c).

Acknowledgment

We thank J. Mycielski for providing information about the history of our original problem and for suggesting conjectures that led to Theorem 2.

References

[1] C.J. Everett and S. Ulam, On ordered groups, Trans. Amer. Math. Soc. 57 (1945) 208–214.
[2] J. Mycielski, Independent sets in topological algebras, Fund. Math. 55 (1964) 139–147.
[3] J. Mycielski, Almost every function is independent, Fund. Math. 81 (1973) 43–48.
[4] J. Schreier and S. Ulam, Sur les transformations continues des sphères euclidiennes, C.R. Acad. Sci. Paris 197 (1933) 967–968.