Algebra and Discrete Mathematics
Volume 14 **(2012)**. Number 2. pp. 185 – 216
© Journal "Algebra and Discrete Mathematics"
RESEARCH ARTICLE

Word length in symmetrized presentations of Thompson's group F[1]

Matthew Horak, Alexis Johnson and Amelia Stonesifer

Communicated by V. Nekrashevych

Abstract. Thompson's groups F, T and Z were introduced by Richard Thompson in the 1960's in connection with questions in logic. They have since found applications in many areas of mathematics including algebra, logic and topology, and their metric properties with respect to standard generating sets have been studied heavily. In this paper, we introduce a new family of generating sets for F, which we denote as Z_n, establish a formula for the word metric with respect to Z_1 and prove that F has dead ends of depth at least 2 with respect to Z_1.

1. Introduction

A common goal in geometric group theory is to determine which properties a given group possesses with respect to some, none or all finite generating sets. For Thompson's Group F, few results like this exist, especially for its metric properties such as dead end depth and almost convexity. This is mostly due to the fact that the only generating sets for which a method to calculate the word metric is known are the so-called consecutive generating sets, $X_n = \{x_0, x_1, \ldots, x_n\}$. With respect to every X_n, F is known to contain dead ends [10], leading to the question of whether or not F has dead ends with respect to every finite generating set. In this paper, we introduce a new family of generating sets, Z_n, establish

[1]This work was supported by NSF REU grant DMS 0453421.
2010 MSC: 20F65.
Key words and phrases: Thompson's group F, dead ends, diagram group.

a formula for the word metric with respect to Z_1, and use our formula to show that F has dead end elements of depth at least 2 with respect to Z_1. This result extends the class of generating sets with respect to which F is known to contain dead ends.

Thompson's group F can be viewed as a group of piecewise linear homeomorphisms from the unit interval to itself. In this definition, the generating sets X_n seem slightly asymmetric because each generator x_i is the identity on the subinterval $[0, \frac{2^i-1}{2^i}]$. The generating sets Z_n that we introduce "symmetrize" the standard generating sets by adding elements y_i that are the identity on intervals $[\frac{1}{2^i}, 1]$.

1.1. Word length and dead ends

For a finitely generated group G and fixed finite generating set S, the word length, $l_S(g)$, of g with respect to S is defined to be 0 if g is the identity and to be the minimal length of an expression in elements of $S \cup S^{-1}$ that is equal to g if g is not the identity. Word length depends on the generating set, but if there is a fixed generating set under discussion, we often omit the reference to the generating set and write $l(g)$ for $l_S(g)$. Geometrically, $l(g)$ can be viewed as the distance from the identity to g in the Cayley graph of G with respect to S.

For a group G and fixed generating set S, the element $g \in G$ is called a *dead end* (with respect to S) if $l(g\alpha) \leq l(g)$ for all $\alpha \in S \cup S^{-1}$. Geometrically, no infinite geodesic ray originating at the identity in the Cayley graph can pass through g. More generally, if G is an infinite group, the *dead end depth* of g with respect to S is defined as,

$$\delta(g) = \min\{l(\alpha) \mid l(g\alpha) > l(g)\}.$$

Since G is an infinite group, $\delta(g)$ exists for every $g \in G$. Geometrically, $\delta(g)$ is the distance from g to the complement of the ball of radius $l(g)$ centered at the identity in the Cayley graph of G with respect to S. Note that the element g is a dead end if its dead end depth is at least 2. For a finitely generated infinite group G, the dead end depth of G is,

$$\delta(G) = \max\{\delta(g) \mid g \in G\}$$

if the maximum exists and infinity otherwise. Geometrically, $\delta(G)$ is the least integer N such for any $g \in G$ there is guaranteed to exist a path of length at most N from g to the complement of the ball of radius $l(g)$ centered at the identity. We note that these definitions are consistent with the definition of dead end depth in [7], but is inconsistent with the convention in [10] in which the dead end depth is one less than this dead end depth.

Among other places, dead ends have found application in the proof in [12] demonstrating a random walk that is biased towards the identity on the lamplighter group but that escapes from the identity faster than a simple random walk. Dead ends also played a role in Bogopol′skiĭ's result that infinite commensurable hyperbolic groups are bi-Lipschitz equivalent [4].

1.2. Thompson's group F

We now define Thompson's group F and survey some background results about the metric structure of F. We refer to [6] for a more detailed discussion of the group F.

Definition 1.1. *Thompson's group F* is the set of piecewise linear orientation preserving homeomorphisms from [0,1] to [0,1] with the following properties:

1) Each $f \in F$ has only finitely many points of non-differentiability, each of which occurs at a dyadic rational number.
2) The slope of f at every point of differentiability is a power of 2.

The binary operation is composition of functions.

There are two standard presentations for Thompson's group F: the infinite presentation,

$$F = \langle x_k, k \geq 0 \mid x_i^{-1} x_j x_i = x_{j+1} \text{ if } i < j \rangle \tag{1}$$

and the finite presentation,

$$\langle x_0, x_1 \mid [x_0 x_1^{-1}, x_0^{-1} x_1 x_0], [x_0 x_1^{-1}, x_0^{-2} x_1 x_0^2] \rangle. \tag{2}$$

The generating sets $X_n := \{x_0, x_1, \ldots x_n\}$ are commonly referred to as the *standard* or *standard consecutive* generating sets.

As piecewise linear homeomorphism of the unit interval, the elements x_0 and x_1 are given by,

$$x_0(t) = \begin{cases} \frac{1}{2}t, & 0 \leq t \leq \frac{1}{2} \\ t - \frac{1}{4}, & \frac{1}{2} \leq t \leq \frac{3}{4} \\ 2t - 1, & \frac{3}{4} \leq t \leq 1 \end{cases} \qquad x_1(t) = \begin{cases} t, & 0 \leq t \leq \frac{1}{2} \\ \frac{1}{2}t + \frac{1}{4}, & \frac{1}{2} \leq t \leq \frac{3}{4} \\ t - \frac{1}{8}, & \frac{3}{4} \leq t \leq \frac{7}{8} \\ 2t - 1, & \frac{7}{8} \leq t \leq 1. \end{cases}$$

FIGURE 1. Graphs of x_0 (left) and x_1 (right)

The graphs of x_0 and x_1, are shown in Figure 1.

Note x_1 is the identity on $[0, \frac{1}{2}]$ and the graph of x_1 contains a "copy" of the graph of x_0 scaled down by a factor of $\frac{1}{2}$ and shifted into the upper right corner. The elements x_i are constructed similarly, by defining x_i to be the identity on the interval $[0, \frac{2^i-1}{2^i}]$ and then to act as a scaled version of x_0 on the interval $[\frac{2^i-1}{2^i}, 1]$. This is depicted schematically in Figure 2 for the elements x_1, x_2 and x_3.

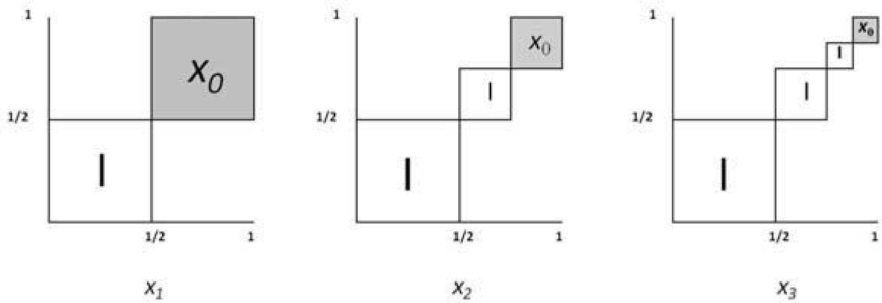

FIGURE 2. Progression of elements x_i

1.3. Combinatorial models for F

A common strategy for studying the geometric and combinatorial properties of F is to find a combinatorial model whose structure faithfully represents the algebraic structure of F but whose elements are easier to manipulate and "compose" or "multiply". This often takes the form of a set of combinatorial objects that are in 1-to-1 correspondence with the elements of F and a description of how the combinatorial representative for an element changes when multiplied by a generator or another element of F.

The most common combinatorial model for F is that of *tree pair diagrams*, for which the reader should consult [6] for a full description. Other models include *forest diagrams* [2], *diagram groups* [9] and *pipe diagrams* [5]. By viewing elements of F geometrically in these ways, researchers have been able to prove that F is not minimally almost convex with respect to any subset of X_n that contains x_1 [3, 11], has dead end depth 3 with respect to X_1 [7], has dead end depth bounded by $\frac{n}{2}$ and $4n-2$ with respect to X_n for $n \geq 3$ [10] and has Hilbert space compression $\frac{1}{2}$ [1]. Except for the Hilbert space compression result, all of the above models and results apply to the standard generating sets, X_n. In this paper, we define and use another model, which we call *wave diagrams* that are better suited to our more symmetrical generating sets, Z_n, which we define in the next subsection.

1.4. Generating set Z_n

For $i \geq 1$, let y_i be the element of Thompson's group F given by,

$$y_i(t) = \begin{cases} 2t, & 0 \leq t \leq \frac{1}{2^{i+2}} \\ t + \frac{1}{2^{i+2}}, & \frac{1}{2^{i+2}} \leq t \leq \frac{1}{2^{i+1}} \\ \frac{t}{2} + \frac{1}{2^{i+1}}, & \frac{1}{2^{i+1}} \leq t \leq \frac{1}{2^{i}} \\ t, & \frac{1}{2^{i}} \leq t \leq 1. \end{cases}$$

Informally, the function y_i is the identity on the interval, $[\frac{1}{2^i}, 1]$ and acts as a scaled down version of x_0^{-1} on the interval $[0, \frac{1}{2^i}]$. The graphs of y_1, y_2 and y_3 are depicted schematically in Figure 3

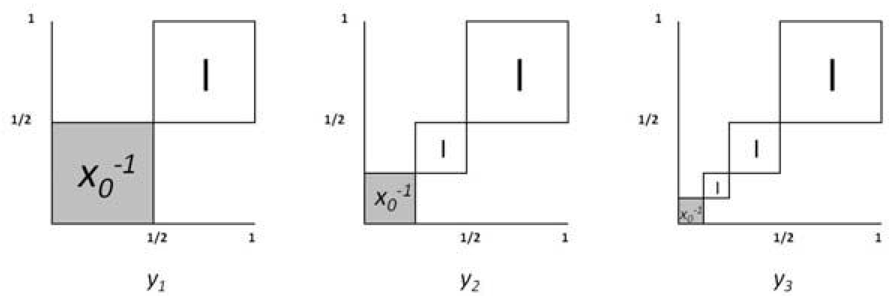

Figure 3. Progression of elements y_i

Definition 1.2. The subset $Z_n \subset F$ is given by

$$Z_n := \{x_0, x_1, y_1, x_2, y_2, ..., x_n, y_n\}.$$

Since the generating set X_n is contained in Z_n, it is clear that Z_n is a generating set for F. For the remainder of the paper, we focus on determining a formula for calculating the word metric with respect to the generating set Z_1 and on using this formula to show that F has dead end depth at least 2 with respect to Z_1. We begin by developing a combinatorial model for elements of F for which it is easy to analyze the effect of multiplication by a generator in Z_n.

2. Wave diagrams

Definition 2.1. A *wave diagram* or simply a *diagram* is a planar graph, D, together with a smooth embedding of D in the plane satisfying the following.

- D contains infinitely many vertices, all of which lie on a distinguished bi-infinite path in G that is embedded in the x-axis in such a way that the vertices have no accumulation point. This path is called the *central line* and its edges are called *central edges*.
- All other edges are embedded either above the x-axis, where they are defined as *upper edges*, or below the x-axis, where they are defined as *lower edges*.
- There are only finitely many upper and lower edges.
- All finite regions have three edges.
- There is a vertex labeled $\frac{1}{2}$ on the top. This vertex cannot be between the two endpoints of an upper edge. There also exists a vertex labeled $\frac{1}{2}$ on the bottom that cannot be between two endpoints of a lower edge. The same vertex may be labeled $\frac{1}{2}$ on both top and bottom.

We will be concerned mainly with the finite portion of a diagram that contains all of the non central edges. This portion of a diagram is referred to as the *essential* portion. Throughout, our convention will be to denote the vertices along the central line by $w_0, w_1, \ldots, w_n$, starting with w_0 as the left-most vertex incident to a non-central edge and ending with w_n, the the right-most vertex incident to a non-central edge. In many arguments, we must distinguish the “left” and “right” endpoints of an

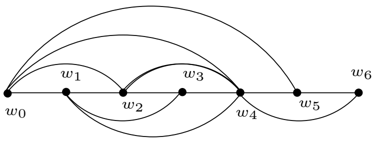

Figure 4. The essential portion of a diagram

edge. If e is an edge incident to vertices w_i and w_j with $i < j$, then we use e^- to refer to w_i and e^+ to refer to w_j.

Figure 4 is an example of the essential portion of a diagram shown with the top and bottom $\frac{1}{2}$ labels omitted. Throughout this paper, we follow the convention of drawing only the essential portion of a diagram unless the top or bottom $\frac{1}{2}$ lies outside of the essential portion, in which case, we draw enough to include both $\frac{1}{2}$ labels.

In order for diagrams to faithfully model elements of F, we must omit diagrams with *dipoles*, defined below, which completely characterize the extent to which two diagrams representing the same element can differ.

Definition 2.2. A *dipole* is a subgraph of a diagram D that consists of three vertices adjacent along the central line, w_i, w_{i+1}, w_{i+2}, an upper edge u, a lower edge l such that $u^- = l^- = w_i, u^+ = l^+ = w_{1+2}$, together with the two central edges connecting w_i, w_{i+1} and w_{i+2}. A dipole is shown in Figure 5.

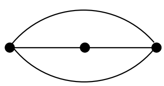

Figure 5. A dipole

Definition 2.3. A diagram D is called *reduced* if it contains no subgraph that is a dipole.

For the remainder of the paper, we deal with only reduced diagrams, but for convenience, we will usually omit the word "reduced". Therefore, all diagrams in the following are assumed to be reduced diagrams unless otherwise stated.

2.1. Labeling and reading a wave diagram

In this section, we describe a 1-to-1 correspondence between (reduced wave) diagrams and elements of Thompson's group F. The first step is to use the upper and lower edges and the $\frac{1}{2}$ labels to label the remaining vertices with "upper" and "lower" labels. Geometrically, we think of "upper" labels as being written above the x-axis and "lower" labels below the x-axis. Diagrams are labeled in such a way that every vertex receives an upper and a lower label. We must first identify the *exposed vertices.*

Definition 2.4. Vertex w_i is *exposed* from the bottom if there exists no lower edge e such that $e^- = w_a$ and $e^+ = w_b$, with $a < i < b$. Vertex w_i is *exposed* from the top if there does not exist an upper, non-central edge, e, such that $e^- = w_a$ and $e^+ = w_b$, where $a < i < b$.

Now we may assign top labels. First work to the right from the vertex labeled $\frac{1}{2}$ on top, labeling the exposed upper vertices in order: "$\frac{3}{4}$, $\frac{7}{8}$, $\frac{15}{16}$,...". Next, work to the left of the vertex labeled $\frac{1}{2}$ on top, label the exposed upper vertices "$\frac{1}{4}$, $\frac{1}{8}$, $\frac{1}{16}$, ...". To assign upper labels to vertices without an upper label, use the rule that if e is an upper edge in which e^- has upper label α and e^+ has upper label β then the non-exposed vertex w contained in the region bounded by e has label $\frac{\alpha+\beta}{2}$. Vertices are assigned lower labels by following the same procedure with lower edges starting with the vertex labeled $\frac{1}{2}$ on the bottom. Results of this process are demonstrated in figure 6.

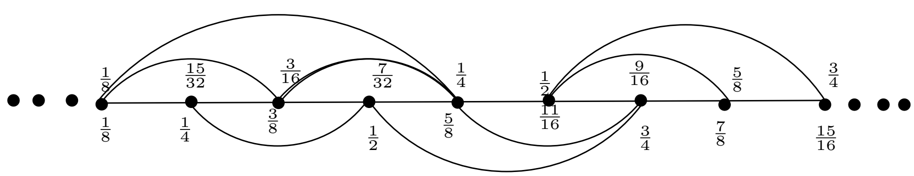

FIGURE 6. Example of labeling a Wave Diagram

Notice that the dots along both ends represent the infinite vertices on the left and right of the non-central edges. In future diagrams we will exclude these ellipsis for simplicity. Even though there is no vertex infinitely far left or right, we sometimes think of an "ideal vertex" lying infinitely far left and one infinitely far to right labeled "0" and "1" respectively.

Two diagrams D_1 and D_2 are isomorphic if there exists an orientation preserving homeomorphism from $\mathbb{R}^2$ to itself that restricts to a label-preserving graph isomorphism between D_1 and D_2 that takes upper edges to upper edges, lower edges to lower edges, and central edges to central edges. As described below, the set of isomorphism classes of reduced diagrams is in in one-to-one correspondence with F. Therefore, it is convenient to abuse terminology and refer to the "diagram D" when one really means the "isomorphism class of diagram D". For the remainder of this paper, we adopt to this convention together with the convention that all diagrams under consideration are reduced. Thus the term "diagram" will always mean "isomorphism class of reduced diagram". Moreover, if D_1 and D_2 are isomorphic diagrams, then there is a unique diagram isomorphism between D_1 and D_2, thus one may unambiguously speak of an edge in an isomorphism class of diagrams.

Each isomorphism class of diagrams corresponds to a unique element of F, as follows. The top and bottom labels on vertices along the central edge of a wave diagram D determine two subdivisions of the interval $[0, 1]$. The piecewise linear homeomorphism f corresponding to D maps the subinterval from the bottom labeling of the endpoints of any edge e linearly to the subinterval for the top labeling of the endpoints of the same edge e. Since all of the subintervals have length which is a power of 2 and since there are only finitely many non-central edges in D, f is a piecewise linear homeomorphism of I to itself with only finitely many points of non-differentiability and slopes that are powers of 2 elsewhere. Therefore, f is an element of D. This correspondence is seen to be a bijection between the set of isomorphism classes of reduced diagrams and F. For $g \in F$, we use the notation $D(g)$ to denote the diagram corresponding to g and for diagram D, we use f_D to denote the element of F corresponding to D.

The most basic of these diagrams, consisting of at most one edge, are the elements of the generating set Z_n as in Figure 7.

2.2. Multiplication of wave diagrams

Though it is not necessary for what follows, we note that one can define a "composition" of these diagrams so that composition of diagrams corresponds exactly to multiplication of the associated group elements. This is done in a way similar to composition in ordinary diagram groups as defined in [9] by identifying the bottom path of $D(f)$ with the top path of $D(g)$ and performing the appropriate "reduction" to arrive at the diagram for $D(fg)$.

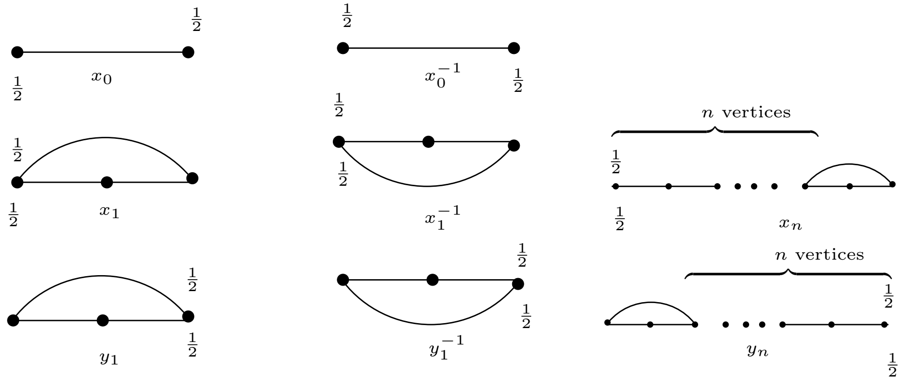

FIGURE 7. Elements of Z_n

While it is not necessary to know how to compose two general diagrams, it is important to establish the rules for constructing $D(fx_i)$ and $D(fy_i)$ from $D(f)$. Before doing so, we classify several different types of edges.

Definition 2.5. Let $f \in F$ and let g and h be edges in $D(f)$. Let $g^- = w_i$, $g^+ = w_j$, $h^- = w_k$, and $h^+ = w_m$. Edge g is said to be *below* edge h if and only if one of the following holds:

1) both g and h are lower edges, $i \leq k$ and $j \geq m$ as shown in Figure 8;
2) g is a lower edge, h is an upper edge and one of the following is true, as shown in Figure 9:
   (a) $k \leq i < m$
   (b) $k < j \leq m$
   (c) $i \leq k$ and $j \geq m$; or
3) both g and h are upper edges, $i \geq k$ and $j \leq m$, as shown in Figure 10.

Conversely, edge h is said to be *above* edge g if edge if g is below h.

We note that the relation $\prec$, defined on the set of edges of a diagram by $g \prec h$ if g is below h, is transitive and antisymmetric, so it defines a partial relation on the set of edges of a diagram

Definition 2.6. An edge, e in $D(f)$, is *bottom-most* if there exists no edge below e.

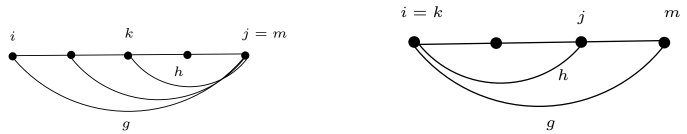

FIGURE 8. Examples of Definition 2.5 (1)

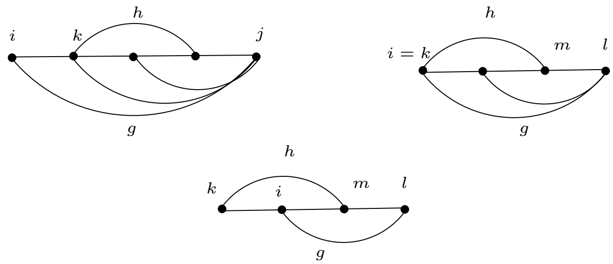

FIGURE 9. Examples of Definition 2.5 (2)

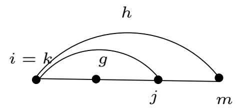

FIGURE 10. Examples of Definition 2.5 (3)

Definition 2.7. An edge e is an *exposed upper edge* if e is an upper edge and there does not exist any non-central edges below e. Edge g in Figure 10 is an exposed upper edge.

We now describe the effect on a diagram that multiplication by a generator in F has. Let f belong to F and and let w_i be the vertex of $D(f)$ labeled $\frac{1}{2}$ on the bottom. For generator $\alpha \in Z_1 \cup Z_1^{-1}$, construct $D(f\alpha)$ from $D(f)$ as follows.

1) If $\alpha = x_0$ then move the bottom $\frac{1}{2}$ label left of w_i to the nearest vertex exposed from the bottom.
2) If $\alpha = x_0^{-1}$ then move the bottom $\frac{1}{2}$ label right of w_i to the nearest vertex exposed from the bottom.
3) If $\alpha = x_1$, proceed as follows.

- If there exists at least one lower edge to the right of w_i, delete the bottom-most of these edges.
- If there does not exist any lower edges to the right of w_i, add a bottom-most upper edge that connects w_i and w_{i+1}. Then, add a vertex subdividing the old central edge from w_i to w_{i+1} into two new central edges, as seen in Figure 11.

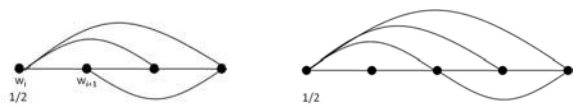

FIGURE 11. Diagram $D(f)$ (left) and $D(f\alpha)$ (right)

4) If $\alpha = x_1^{-1}$, proceed as follows.

- If there exists an exposed upper edge to the right of w_i, delete this edge and the vertex w_{i+1}. Connect w_i and w_{i+2} with a new central edge, as in Figure 12.

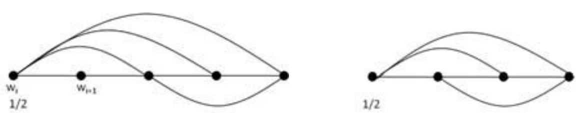

FIGURE 12. Diagram $D(f)$ (left) and $D(f\alpha)$ (right)

- If there does not exist an exposed upper edge to the right of w_i, add a bottom-most lower edge connecting w_i to the second nearest exposed vertex to the right of w_i.

5) If $\alpha = y_1$, then follow Case (3), but consider edges and vertices to the left of w_i.
6) If $\alpha = y_1^{-1}$ then follow Case (4), but consider edges and vertices to the left of w_i.

3. Formula for length with respect to Z_1

Before presenting our length formula, we establish some definitions. For our purposes, we explicitly indicate both the vertices and edges that the path traverses. We use v_i to represent the vertices and q_i to represent the edges in the path.

Definition 3.1. Let $\mathcal{N}$ be the set of non-central edges of the diagram $D(f)$ and let $\lambda = v_0 q_1 v_1 q_2 v_2 \cdots q_l v_l$ be a path in $D(f)$. The function $r_\lambda : \mathcal{N} \to \mathbf{N} \cup \{\infty\}$ is a *reaching function* for λ if the following hold.

1) If $r_\lambda(e) = 0$ then
   (a) Edge e is incident to v_0,
   (b) Every non-central edge below e is incident to v_0,
   (c) No q_j for $j > 0$ is a central edge below e and
   (d) Edge e does not appear as q_j for any $j > 0$.
2) If $0 < r_\lambda(e) < \infty$ then
   (a) Edge e is incident to $v_{r_\lambda(e)}$,
   (b) If h is a non-central edge below e then either h is incident to $v_{r_\lambda(e)}$ or $r_\lambda(h) < r_\lambda(e)$,
   (c) No q_j for $j > r_\lambda(e)$ is a central edge below e and
   (d) Edge e does not appear as q_j for any $j > r_\lambda(e)$.

Frequently we have a path λ together with a fixed reaching function r_λ for λ. In this case, we say that λ *reaches* non-central edge e at vertex v_i if $r_\lambda(e) = i$. We note that this use of the term "reach" depends on the particular reaching function under consideration, so we use this convention only when there is only one such function under consideration. Sometimes, to emphasize r_λ we say that λ reaches e with respect to r_λ at vertex v_i.

Definition 3.2. Let $f \in F$. The path $\lambda = v_0 q_1 v_1 q_2 v_2 \cdots q_l v_l$ is a *legal path* in $D(f)$ if there is a reaching function r_λ for λ such that the following three conditions hold.

1) The vertex v_0 is labeled $\frac{1}{2}$ on the bottom and v_l is labeled $\frac{1}{2}$ on top.
2) With respect to r_λ, the path λ reaches every non-central edge in $D(f)$.
3) For all $0 \leq i \leq l$, if h is a non-central edge below q_i then $r_\lambda(h) < i$.

When we refer to a legal path, we always implicitly assume that we have fixed a reaching function r_λ that reaches every non-central edge in $D(f)$. Thus, the term "legal path" refers to a legal path together with a fixed reaching function.

The next two observations, which we use frequently in the proofs below follow quickly from the definition a legal path.

Observation 3.3. Let e be a non-central edge in the legal path λ. No edge q that occurs in λ before the last occurrence of e in λ can be above e.

Observation 3.3 follows from the fact that if e is below q, it must be reached at a vertex in the path occurring before q, and once an edge is reached by λ, it cannot occur again in λ.

Observation 3.4. If e is a lower and bottom-most edge that is in the legal path λ and if q is reached at a vertex to the left (respectively right) of e^- (respectively e^+) then both endpoints of q lie to the left (respectively right) of e^- (respectively e^+).

Observation 3.4 follows from the fact that if edge q has endpoints on opposite sides of one of the endpoints of e then either e is under q or q is under e, neither of which occurs by Observation 3.3 and the fact that e is a bottom-most edge.

We have the following formula for determining the word length with respect to Z_1 of the element $f \in F$.

Theorem 3.5. For every $f \in F$, the word length of f with respect to the generating set $Z_1 = \{x_0, x_1, y_1\}$ is given by the formula

$$l_{Z_1}(f) = E(f) + P(f)$$

where $E(f)$ is the number of non-central edges in $D(f)$, and $P(f)$ is the length of a minimal length legal path in $D(f)$.

In some of the following proofs, we define a new path λ' that reorders the vertices and edges of an original path λ. Often, we use an interval notation to denote this path. To denote the subinterval of λ between vertices v_i and v_p, we use the notation, $[v_i, v_p]$. Now, when $[v_i, v_p]$ appears in the notation for a new path λ', every vertex and edge between v_i and v_p is to appear in λ', preserving the order from λ. If a bracket is closed, include the corresponding vertex, and if the bracket is open, do not.

Lemma 3.6. Let $f \in F$ and let λ' be a minimal length legal path in $D(f)$. No bottom-most lower edge occurs more than twice in λ.

Proof. Let $\lambda' = v_0 q_1 v_1 q_2 v_2 \cdots q_n v_n$ be a legal path of minimal length in $D(f)$ with reaching function $r_{\lambda'}$. Suppose towards a contradiction that e is a bottom-most lower edge of $D(f)$ such that $e = q_i = q_j = q_k$ with $i < j < k$. Without loss of generality, assume that and q_i,

q_j and q_k are the first three times that e appears in λ'. Then $\lambda' = [v_0, v_{i-1}]e[v_i, v_{j-1}]e[v_j, v_{k-1}]e[v_k, v_n]$. We claim that the shorter path, $\lambda = [v_0, v_{i-1}](v_j, v_{k-1}]e[v_i, v_{j-1})[v_k, v_n]$, is a legal path in f, which contradicts the minimality of λ'. To prove this claim, we construct a reaching function r_λ that reaches every non-central edge in $D(f)$. Rename the vertices and edges appearing in λ by $\lambda = s_0p_1s_1p_2s_2 \cdots p_{n-2}s_{n-2}$. Note that $s_t = v_t$ if $0 \le t \le i-1$, and $s_t = v_{t-2}$ if $t \ge k-1$ since the paths λ and λ' agree at their beginnings and endings. Additionally, the same set of vertices appear in λ and λ' since $v_{i-1} = v_j = v_{k-1}$ and $v_i = v_{j-1} = v_k$. This is illustrated in Figure 13.

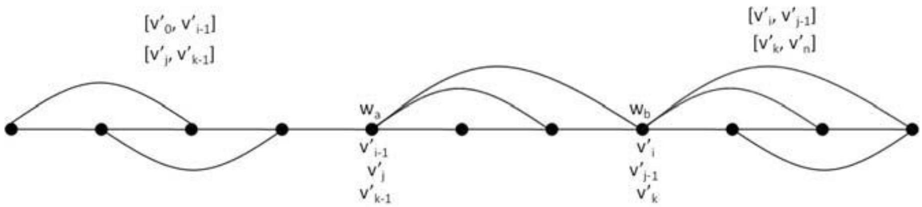

FIGURE 13. Interval Location

Let v be the leftmost vertex of $D(f)$ that is incident to a non-central edge and number the vertices to the right of v in order along the central path as $w_0 = v, w_1, w_2, w_3, \cdots$. Let a and b be such that λ' traverse e at q_i from w_a to w_b. Without loss of generality we may assume that $b > a$. Let h be a non-central edge of $D(f)$. We define r_λ as follows:

1) If $r_{\lambda'}(h) \le i-1$ then $r_\lambda(h) = r_{\lambda'}(h)$.
2) If $i \le r_{\lambda'}(h) \le j-1$ then note that the third segment of λ (together with the vertex v_k) is the same as the second segment of λ'. Identify the vertex s_t in the third segment (including v_k) of λ that is in the same position as the vertex $v_{r_{\lambda'}(h)}$ in the second segment of λ'. Set $r_\lambda(h) = t$.
3) If $j \le r_{\lambda'}(h) \le k-1$ then note that the second segment of λ (together with the vertex v_{i-1}) is the same as the third segment of λ'. Identify the vertex s_t in the second segment (including v_{i-1}) of λ that is in the same position as the vertex $v_{r_{\lambda'}(h)}$ in the third segment of λ'. Set $r_\lambda(h) = t$.
4) If $k \le r_{\lambda'}(h)$ identify the vertex s_t in the fourth segment of λ that is in the same position as vertex $v_{\lambda'(h)}$ in the fourth segment of λ' and set $r_\lambda(h) = t$.

We now prove that λ is a legal path by checking the conditions of Definition 3.2. Definition 3.2 has three parts, (1), (2) and (3). Part (2) requires checking that λ reaches certain edges. This requires checking Definition 3.1, which has many parts, (1a), (1b), (1c), (1d) and (2a), (2b), (2c), (2d). The arguments involved in checking many of the conditions are very similar to each other and involve lengthy case analyses. Therefore, we present the most difficult and representative parts and cases to prove in detail and omit the details for the remaining cases, which are similar to the ones presented in detail. The parts we prove in detail are: Definition 3.2 (1) and the Definition 3.1 (2a), (2b) and (2d) portions of Definition 3.2 (2).

3.2 (1) Since λ' is a legal path, $v_0 = s_0$ is the vertex labeled $\frac{1}{2}$ on the bottom and $v_n = s_{n-2}$ is the vertex labeled $\frac{1}{2}$ on the top. Since λ also begins at s_0 and ends at s_{n-2}, so condition one is satisfied.

3.2 (2) We must show that λ reaches every non-central edge in $D(f)$. Let h be a non-central edge of $D(f)$. The argument proving that λ reaches h in the case that $r_{\lambda'}(h) = 0$ is similar to the argument for the case in which $r_{\lambda'}(h) > 0$, so we present details of the case that $r_{\lambda'}(h) > 0$ and omit the details in the case that $r_{\lambda'}(h) = 0$. So, assume that $r_{\lambda'}(h) = l > 0$. We must verify that the conditions of Definition 3.1 (2a) – (2d) are satisfied. For convenience of notation, let $r_\lambda(h) = N$ so that $s_N = v_l$.

3.1 (2a) Since λ' reaches h at v_l, h is incident to v_l in $D(f)$. By definition, $r_\lambda(h)$ is the s-subscript of the occurrence of the v_l at which h is reached in λ'. Thus, h is incident to $s_{r_\lambda(h)}$.

3.1 (2b) Let m be a non-central edge below h, and let $M = r_\lambda(m)$ and $M' = r_{\lambda'}(m)$, so that $s_M = v_{M'}$. We must show that either $M < N$ or that m is incident to s_N. Assume that m is not incident to s_N. By the construction of r_λ, we see that m is not incident to v_l. Since λ' is a legal path, $M' = r_{\lambda'(m)} < r_{\lambda'}(h) = l$. We separate the argument into cases that depend on the relative positions of v_l and $v_{M'}$.

i. If s_N and s_M are in the same interval of λ then they are in the same interval of λ' and occur in the same order there as in λ. Since λ' is a legal path, $M' < l$, so $M < N$, as required.
ii. If s_N and s_M are in different intervals of λ', there are again several cases.

A. If $k \le l \le n$, then v_l is in the final interval of the path λ' and so s_N is in the final interval of λ. Since λ' is a legal path, m must be reached at $v_{M'}$ in one of the intervals $[v_0, v_{i-1}]$, $[v_i, v_{j-1}]$, or $[v_j, v_{k-1}]$ of λ'. Since these are also the first three of λ and s_N is in the last interval of λ, we have $M < N$.
B. If $i \le l \le (j-1)$, then v_l occurs in the second interval of λ' and s_N occurs in the third interval of λ. Since m is reached before h in λ' and since $v_{M'}$ is in a different interval from v_l, we know that $v_{M'}$ appears in the first interval of λ', which is the same as the first interval of λ. Thus s_M is in the first interval and s_N is in the third interval of λ proving that $M < N$ as required.
C. If $j \le l \le (k-1)$, we show that $v_{M'}$ is in the first interval of λ'. By a similar argument to that of the previous case, we know $v_{M'}$ cannot be in the last interval of λ'. Thus, we need to show that $v_{M'}$ is not in the second interval of λ'. Assume toward contradiction that $v_{M'}$ is in the second interval. By Observation 3.4, we know that m is incident to w_z and w_u with $z, u \ge b$. We also know that h is incident to w_s and w_g with $s, g \le a$. But, we know that $s, g \le a < b \le z, u$, so m is not below h and we have a contradiction.
D. If $0 \le l \le (i-1)$, this case is argued similar to Case (i).

3.1 (2c) Let m be a central edge below h in $D(f)$. We must show that m does not appear in λ after vertex s_N. We let r be an edge in λ occurring after s_N. We must show that $m \ne r$. The argument involves careful case analysis of the relative positions of s_N and r in λ and λ'. The details are similar to the details in 3.1 (2b) and (2d), so we omit them.

3.1 (2d) Since $N = r_\lambda(h)$, we must show that h does not occur in λ after vertex s_N. To do this let $r > N$. We show that $h \ne p_r$ by considering several cases.

i. If s_N and p_r are in the same interval within λ, then they occur in the same interval of λ' and therefore p_r appears after v_l in λ'. Since λ' is a legal path in $D(f)$, it follows from definition 3.1 (2c) that $p_r \ne h$.

ii. If s_N and p_r are not in the same interval within λ, we separate further into cases.

A. If $k \leq l \leq n$, then h is reached by a vertex in the final interval of the path λ'. By the definition of r_λ, we know therefore that $N > l$ and so $r > l$. Therefore, p_r and s_N are in the same interval, and the argument in the previous case shows that $h \neq p_r$.

B. If $i \leq l \leq j-1$, then h is reached by a vertex in the second interval of the path λ'. Since λ' is a legal path, we know that $h \neq q_c$ for any q_c in $[v_j, v_{k-1}]$ or $[v_k, v_n]$. Now, by the definition of r_λ, s_N appears in the third interval of λ. Since p_r appears after s_N in λ and not the same interval, we know that p_r is within the fourth interval of λ, which is contained within the λ' interval, $[v_k, v_n]$. Therefore, we conclude that $h \neq q_p$.

C. If $j \leq l \leq k-1$, then h is reached by a vertex in the third interval of the path λ'. Thus, $h \neq q_c$ for any q_c within $[v_k, v_n]$. We now show that $h \neq q_c$ for q_c within $[v_i, v_{j-1}]$. Assume toward contradiction that $h = q_c$ for some $i < c \leq j-1$. Thus, h is incident to v_c. By Observation 3.3, we conclude that h is incident to some w_u, with $u \geq b$. Since h is reached by λ' at v_l with $j \leq l \leq k-1$, by Observation 3.4 we know that h must be incident to some w_s and w_g with $s, g \leq a$. Therefore, $s, g \leq a < b \leq u$, and we have a contradiction. Thus, $h \neq q_c$ for any edge q_c within the second or fourth intervals of λ'.

Now, by the definition of r_λ, we know that s_N occurs within the second interval of λ, so p_r occurs in the third or fourth intervals of λ, which are contained in the second and fourth intervals of λ'. Since h is not equal to any q_c occurring in these intervals, it follows that $h \neq p_r$.

D. If $0 \leq l \leq i-1$ then by the definition of r_λ, the vertex s_N occurs in the first interval of λ. Therefore, p_r occurs within the second, third or fourth interval of λ, which are contained in the second, third and fourth intervals of λ. Since h is reached by λ' at a vertex in the first interval of λ', definition 3.1 (2c) implies that $h \neq q_c$

for any q_c in the second third or fourth intervals of λ'. Since p_r lies in one of these intervals, $h \neq p_r$.

3.2 (3) Let p_c be an edge in λ, and m a non-central edge below q_c. We must show that λ reaches m before s_c. Since λ' is a legal path, we know that m is reached by λ' at some vertex in λ' that occurs before $v_{r_{\lambda'}(p_c)}$. Again, the argument involves careful case analysis of the relative positions of s_c and m in λ and λ'. The details are similar to the details in 3.1 (2b) and (2d), so we omit them.

This concludes the proof that λ is a legal path in $D(f)$, contradicting the minimality of λ'. Therefore, no bottom-most edge appears more than twice in a legal path. $\square$

In the above lemma, we modified paths in a diagram for a single given element $f \in F$. During the proof that ϕ_{Z_1} computes the length function with respect to Z_1, we occasionally must use paths in one diagram, $D(f)$ to construct paths in the diagram $D(f\alpha)$ for a generator $\alpha \in Z_1 \cup Z_1^{-1}$. As described in Subsection 2.2, if $\alpha \in Z_1 \cup Z_1^{-1}$ then the diagrams $D(f)$ and $D(f\alpha)$ differ by at most the placement of the $\frac{1}{2}$ labels or by at most two edges and one vertex. Although $D(f)$ and $D(f\alpha)$ are formally different diagrams, it will be convenient to abuse notation and refer to a single edge or vertex as “belonging to” both diagrams. We formalize this as follows.

If α is the generator x_0 or x_0^{-1}, then ignoring the labeling of the vertices, $D(f)$ and $D(f\alpha)$ are isomorphic by a unique diagram isomorphism that takes the vertex labeled $\frac{1}{2}$ in $D(f\alpha)$ to the vertex along the bottom path one to the left or right (depending on whether x_0 or x_0^{-1} was used) of the $\frac{1}{2}$ label in $D(f)$ and takes upper edges to upper edges and lower edges to lower edges. If e is an edge of $D(f)$, any reference to “edge e” in $D(f\alpha)$ refers to the image of e under this isomorphism. A similar convention is used for edges of $D(f\alpha)$ and vertices of $D(f)$ and $D(f\alpha)$.

If α is the generator x_1^{-1} and if $D(f\alpha)$ differs from $D(f)$ by the addition of an edge, then there is a unique edge e' in $D(f\alpha)$ such there is an isomorphism between $D(f)$ and $D(f\alpha) \setminus \{e'\}$ that respects the $\frac{1}{2}$ labeling and takes upper edges to upper edges and lower edges to lower edges. Moreover, this isomorphism is unique. If e is an edge of $D(f)$, any reference to “edge e” in $D(f\alpha)$ refers to the image of e under this isomorphism. A similar convention is used for edges of $D(f\alpha) \setminus \{e'\}$ and vertices of $D(f)$ and $D(f\alpha)$.

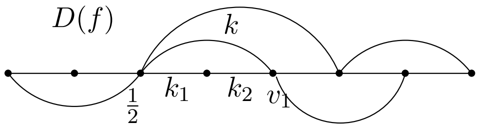

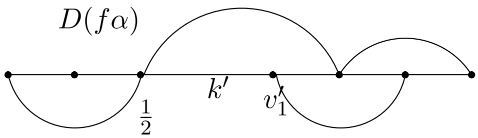

FIGURE 14. Multiplication by x_1^{-1}

If α is the generator x_1^{-1} and if $D(f\alpha)$ differs from $D(f)$ by the deletion of an exposed upper edge in $D(f)$, then the situation is slightly more complicated. In this case, we denote the central edge emanating to the right from the vertex labeled $\frac{1}{2}$ in $D(f\alpha)$ by k'. In the diagram $D(f)$ we denote the central edge emanating to the right from $\frac{1}{2}$ by k_1 and central edge emanating from the right endpoint of k_1 by k_2. We denote the bottommost upper edge emanating from the vertex $\frac{1}{2}$ in $D(f)$ by k. Finally, we denote as v_1 and v_1' the terminal vertices of k and k' respectively, as shown in Figure 14. There is a unique graph isomorphism between $D(f)\setminus\{k,k_1,k_2\}$ and $D(f\alpha)\setminus\{k'\}$ that preserves the $\frac{1}{2}$ labeling, central, upper and lower edges. This isomorphism allows us to refer to edges and vertices as belonging to both $D(f)$ and $D(f\alpha)$ as in the cases for multiplication by x_0 and x_0^{-1}.

The conventions for multiplying by the generator x_1 are obtained by using the conventions for x_1^{-1} and reversing the roles of f and $f\alpha$. The conventions for using this terminology with the generators y_1 and y_1^{-1} are formulated analogously to those for x_1 and x_1^{-1}, with the roles of right and left reversed.

The final tool used in the proof of Theorem 3.5 is the following lemma of Fordham [8], which gives sufficient conditions for a function that guarantee that it computes the length of elements with respect to a given generating set.

Lemma 3.7 ([8], Lemma 3.3.1)**.** Given a group G, a generating set X and a function $\phi : G \to \{0, 1, 2, \ldots\}$, if ϕ has the properties:

1) $\phi(Id_G) = 0$;
2) if $\phi(g) = 0$, then $g = Id_G$;
3) if $g \in G$ and α or α^{-1} is an element of X, then $\phi(g) - 1 \leq \phi(g\alpha)$; and
4) for any non-identity element $g \in G$, there is at least one $\alpha \in G$ with either α or α^{-1} in X such that $\phi(g\alpha) = \phi(g) - 1$,

then $\phi(g) = l(g)$ for all $g \in G$, where $l(g)$ denotes the word length of g with respect to the generating set X.

We now begin the proof of Theorem 3.5.

Proof. Let $f \in F$. We define the function ϕ_{Z_1} by $\phi_{Z_1}(f) = E(f) + P(f)$. To prove that this function computes the word length of f with respect to the generating set Z_1, it suffices show that it satisfies the four properties of Lemma 3.7.

Property 1 of Lemma 3.7: The diagram for the identity element, id in F contains no non-central edges, and the same vertex is labeled $\frac{1}{2}$ on the bottom and the top. Therefore, there exists a legal path with length zero and it follows that $\phi_{Z_n}(id) = 0$.

Property 2 of Lemma 3.7: If $\phi_{Z_1}(f) = 0$, then $E(f) = 0$ and $P(f) = 0$. Thus, f contains no non-central edges and therefore there exists a legal path in f of length zero. It follows from the definition of legal path that the top and bottom $\frac{1}{2}$ labels must be located on the same vertex. Thus, f is the identity.

Property 3 of Lemma 3.7: Let $f \in F$ and $\alpha \in Z_1 \cup Z_1^{-1}$. We must show that $\phi(f) - 1 \leq \phi(f\alpha)$. We do this by analyzing the cases, $\alpha = x_0, \alpha = x_0^{-1}, \alpha = x_1, \alpha = x_1^{-1}, \alpha = y_1, \alpha = y_1^{-1}$ separately. This frequently involves checking the conditions Definition 3.2 through careful and lengthly case analysis, as in the proof of Lemma 3.6. Many of the cases involve arguments very similar to each other or the arguments in Lemma 3.6. Therefore, we present the details for only the most difficult and representative cases and omit the details for the remaining cases.

For ease of notation, let $E(f) = e$, $P(f) = p$, and let v_0 be the vertex labeled $\frac{1}{2}$ on the bottom in $D(f)$.

Case 1: $\alpha = x_0$. Multiplying f on the right by x_0 moves the bottom $\frac{1}{2}$ label left to the nearest exposed vertex along the bottom along, say, edge g. Since we have not added any edges, $E(f\alpha) = e$ and we must show that $P(f) - 1 \leq p$. Assume towards a contradiction that $P(f\alpha) < p - 1$. Thus, there exists a legal path, say $\lambda' = v_0'q_0'v_1'q_1' \cdots q_l'v_l'$, of length l, less than $p-1$. Fix a reaching function $r_{\lambda'}$ for λ'.

Consider the path $\lambda = v_0q_1v_1q_2 \cdots q_{l+1}v_{l+1}$ in $D(f)$ defined by $q_1 = g$, $q_i = q_{i-1}'$ for $i > 1$, and $v_i = v_{i-1}'$ for $i \geq 1$. Define the function r_λ by $r_\lambda(e) = r_{\lambda'}(e) + 1$. We show that r_λ is a reaching function with respect to which λ is a legal path in $D(f)$ by verifying the conditions of Definition 3.1.

Let h be a non-central edge in $D(f)$. Since λ' is a legal path, $r_{\lambda'}(h) = i < \infty$. We now verify the conditions of Definition 3.1 by separately considering the cases $i = 0$ and $i > 0$.

Case $i = 0$: We have $r_\lambda(h) = 1$, so part (2) of Definition 3.1 applies.

3.1 (2a) By definition 3.1 (1a), h is incident to v_0' in $D(f\alpha)$. Since $v_0' = v_1$, h is incident to v_1.

(2b) By definition 3.1 (1b), every non-central edge below h is also incident to v_0' in $D(f\alpha)$, and thus incident to v_1 in $D(f)$.

(2c) By definition 3.1(1c), no central edge below h appears as in λ' q_j' with $j > 0$. Therefore, no central edge below h appears in λ as q_j with $j > 1$.

(2d) By definition 3.1 (1d), h does not appear in λ' as q_j' for any $j > 0$. Thus, h does not appear in λ as q_j for any $j > 1$.

Case $i > 0$: In this case, $r_\lambda(e) = i + 1 \geq 1$ so part (2) of Definition 3.1 applies again.

3.1 (2a) Since $r_{\lambda'}$ is a reaching function for λ', h is incident to $v_i' = v_{i+1}$ in $D(f)$.

(2b) Let m be a non-central edge below h in $D(f)$. Then m is below h in $D(f\alpha)$, and by definition 3.1 (2b) there are two possible cases.

- Edge $r_{\lambda'}(m) \leq i - 1$. In this case, $r_\lambda(m) = j + 1 \leq i < r_\lambda(h)$.
- Edge m is incident to v_i' in $D(f\alpha)$. In this case, m is incident to v_{i+1} in $D(f)$, as required.

(2c) Let m be a central edge below h in $D(f)$. Therefore, m is a central edge below h in $D(f\alpha)$, and by definition 3.1 (2c), we know $m \neq q_j'$ for any $j > i$. Thus $m \neq q_{j+1}$ for any $j > i$. Therefore, $m \neq q_j$ for any $j > i + 1$, as required.

(2d) By definition 3.1 (2d), we know that $q_j' \neq h$ for any $j > i$ since λ' is a legal path in $D(f\alpha)$. Thus, $h \neq q_{j+1}$ in $D(f)$ for any $j + 1 > i + 1$ as required.

We now verify that λ is a legal path with respect to reaching function r_λ.

3.2 (1) Since λ' is a legal path in $D(f\alpha)$, v'_l is the vertex labeled $\frac{1}{2}$ on top. Since λ is the path from v_0 to $v_{l+1} = v'_l$, it is a path from the vertex labeled $\frac{1}{2}$ on bottom to the vertex labeled $\frac{1}{2}$ on top in $D(f)$.

3.2 (2) Let e be a non-central edge of $D(f)$. Since λ' is a legal path, $r_{\lambda'}(e) < \infty$. So, $r_\lambda(e) = r_{\lambda'}(e) + 1 < \infty$, and λ reaches every non-central edge of $D(f)$ with respect to r_λ.

3.2 (3) Let m be an edge below q_i in $D(f)$. We must show that $r_\lambda(m) < i$. Since q_1 is a bottommost edge, i cannot equal 1. Thus, we may assume that $i > 1$. Now m is below edge q'_{i-1} in $D(f\alpha)$. By definition 3.2 (3), m is reached by λ' before v'_{i-1} in $D(f\alpha)$ for all $0 \le i-1 \le l$. Therefore, m is reached by λ before v_i for all $1 \le i \le l+1$.

Thus, λ is a legal path in $D(f)$ of length $l+1$, where $l < p-1$. Therefore, we contradict the assumption that the length of a minimal length legal path in $D(f)$ is p. So, $\phi(f\alpha) \ge \phi(f) - 1$.

Case 2: $\alpha = x_0^{-1}$. Multiplication by α on the right acts similarly to Case 1. The only difference is that v_1 and q_1 will now lie to the right of v_0. Therefore, this case is proved in a similar manner.

Case 3: $\alpha = x_1$.

There must always exist an edge, k, connecting the vertex v_0 to the nearest exposed vertex on the right. Call this vertex v_1. We can see that if there exists a lower edge incident to and to the right of v_0, then k is a lower edge. Otherwise, k is a central edge.

Subcase 1: k is a central edge. Multiplication by α adds an upper edge k' from v_0 to v_1 as shown in Figure 15. Now, $E(f\alpha) = e+1$, so it suffices to show that $P(f\alpha) \ge p-2$. Assume toward contradiction that $P(f\alpha) < p-2$. Then, there exists a legal path, $\lambda' = v'_0 q'_1 v'_1 q'_2 \cdots q'_l v'_l$ with reaching function $r_{\lambda'}$ in $D(f\alpha)$ with $l < p-2$.

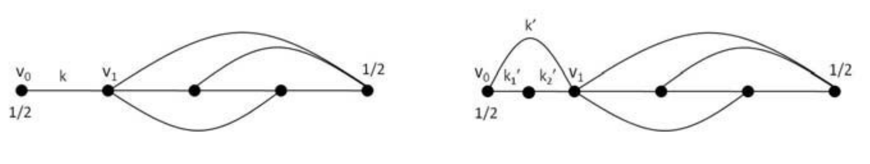

FIGURE 15. Diagrams of $D(f)$ (left) and $D(f\alpha)$ (right)

Since there are no edges below k in $D(f)$, there are not any non-central edges below k' in $D(f\alpha)$. Since λ' is of minimal length, λ' does not use the central edges below k', for otherwise both would have to be used in succession and we could use k' instead and construct a shorter legal path. Therefore we may define path λ in $D(f)$ by $\lambda = v_0q_1v_1q_2v_2\cdots q_lv_l$ in $D(f)$ by $v_i = v_i'$ for all $0 \le x \le l$, q_j=q_j' if $q_j' \ne k'$, and $q_j = k$ if $q_j' = k'$. We define the function r_λ by $r_\lambda(e) = r_{\lambda'}(e)$. We claim that λ is a legal path in $D(f)$ with respect to r_λ. The details in the proof that r_λ is a reaching function for λ and that λ is a legal path with respect to r_λ involve case analysis and arguments similar to the arguments in the case of $\alpha = x_0$, so we omit them here. Since the length of λ is equal to $p-2$, this contradicts the fact that $P(F) = p$, proving that $P(f\alpha) \ge p-2$ as required.

Subcase 2: k is a non-central edge. In this case, k is a lower, bottom-most edge connecting v_0 and v_1. Multiplying by x_1 deletes edge k, so $E(f\alpha) = e-1$, and it suffices to show that $P(f\alpha) \ge p$. We assume toward contradiction that $P(f\alpha) < p$, so there exists a legal path $\lambda' = v_0'q_1'v_1'q_2'\cdots q_l'v_l'$ with reaching function $r_{\lambda'}$ in $D(f\alpha)$ with $l < p$. Consider the path $\lambda = v_0q_1v_1q_2\cdots q_lv_l$ defined by $q_i = q_i'$ and $v_i = v_i'$. We claim that λ is a legal path in $D(f)$ with respect to the reaching function $r_{\lambda'}$. Again, the details in the proof that r_λ is a reaching function for λ and that λ is a legal path with respect to r_λ involve case analysis and arguments similar to the arguments in the case of $\alpha = x_0$, so we omit them here. Since the length of λ is less than p, this contradicts the fact that $P(f) = p$ and proves that $P(f\alpha) \ge p$.

Case 4: $\alpha = y_1$. Multiplying by α on the right acts similarly to Case 3. The only difference is that v_1 is the nearest exposed vertex to the left of v_0. Therefore, this case can be proved in a similar manner.

Case 5: $\alpha = x_1^{-1}$.

Subcase 1:

There does not exist an exposed upper edge emanating to the right of v_0. In this case, multiplication by x_1^{-1} adds a bottom-most edge, k', with $k'^- = v_0$ and k'^+ equal to the second nearest vertex to the right along the bottom path of $D(f)$. There is a unique finite region that is bounded by k' together with two more edges, k_1 and k_2. Let k_1 be the edge incident to k'^- and k_2 be the one incident to k'^+, and let g be the vertex between k_1 and k_2 as shown in Figure 16.

Now, $E(f\alpha) = e+1$, so suffices to show that $P(f\alpha) \ge p-2$. We assume toward contradiction that $P(f\alpha) < p-2$, so there exists a legal path $\lambda' = v_0'q_1'v_1'q_2'v_2'\cdots q_l'v_l'$ with reaching function $r_{\lambda'}$ in $D(f\alpha)$ with length less than $p-2$. From Lemma 3.6, we know that k' appears at most twice in λ'. Consider the following cases:

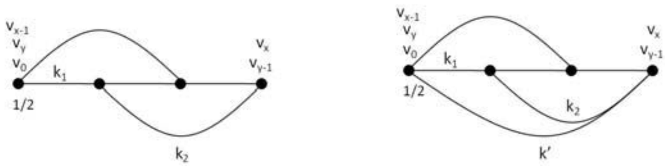

FIGURE 16. Diagrams of $D(f)$ (left) and $D(f\alpha)$ (right)

Case A: Edge k' appears twice in λ'. Let x and y be such that $k' = q'_x$ and $k' = q'_y$ with $1 \le x < y \le l$. Define a path λ in $D(f)$ by following λ', but replacing q'_x by $[k'_1 g k'_2]$, and q'_y by $[k'_2 g k'_1]$. Formally, $\lambda = [v'_0, v'_{x-1}][k_1 g k_2][v'_x, v'_{y-1}][k_2 g k_1][v'_y, v'_l]$. Define the function r_λ by,

$$r_\lambda(e) = \begin{cases} r_{\lambda'}(e) & \text{if } r_{\lambda'}(e) \le x-1 \\ r_{\lambda'}(e)+1 & \text{if } x < r_{\lambda'}(e) \le y-1 \\ r_{\lambda'}(e)+2 & \text{if } r_{\lambda'}(e) \ge y. \end{cases}$$

We now prove that r_λ is a reaching function for λ with respect to which λ is a legal path. For convenience, we rename the vertices and edges in λ by $\lambda = v_0 p_1 v_1 p_2 \cdots v_{l+2} p_{l+2}$.

We begin by verifying that r_λ satisfies the conditions of Definition 3.1. Let h be a non-central edge in $D(f)$.

3.1 (2a) Edge h is also a non-central edge in $D(f\alpha)$, and by the fact that λ' is a legal path with respect to $r_{\lambda'}$ h is incident to $v'_{r_{\lambda'}(h)}$. By the definition of r_λ, we see that h is incident to $v_{r_\lambda(h)}$.

(2b) Let m be a non-central edge below h in $D(f)$. Then, m is a non-central edge below h in $D(f\alpha)$. Suppose that m is not incident to $v_{r_\lambda(h)}$ in $D(f)$. By the definition of r_λ, m is not incident to $v'_{r_{\lambda'}(h)}$. Therefore, since $r_{\lambda'}$ is a reaching function we must have $r_{\lambda'}(m) < r_{\lambda'}(h)$. By the definition of r_λ, this implies that $r_\lambda(m) - r_{\lambda'}(m) \le r_\lambda(h) - r_{\lambda'}(h)$ proving that $r_\lambda(m) < r_\lambda(h)$ as required.

(2c) Let m be a central edge below h in $D(f)$. Let z be an edge in $D(f)$ appearing in λ after $v_{r_\lambda(h)}$. We must show that $m \neq z$.

Subcase 1: $z = q'_a$ for some a.

Since m is a central edge below h in $D(f\alpha)$ and since λ' is a legal path in $D(f\alpha)$, m cannot appear in λ' as q'_a for any $a > r_{\lambda'}(h)$. Thus, z appears in λ' before $r_{\lambda'}(h)$, and by the definitions of λ and

r_λ, z appears in λ before $r_\lambda(h)$. Since z appears in λ after $v_{r_\lambda(h)}$, we have $m \neq z$ as required.

Subcase 2: $z \neq q'_a$ for any a but $z = k_1$.

If k_1 is a non-central edge, since we assumed m is a central edge, $m \neq z$. So, we may assume that k_1 is central. Since $z \neq q'_a$ for any a, k_1 appears in λ because of the fact that k' appeared in λ'. Since z appears in λ after $v_{r_\lambda(h)}$, k' appears in λ' after $v'_{r_{\lambda'}(h)}$. To prove that $m \neq z$, we assume toward contradiction that $m = z = k_1$. Therefore, k_1 is below h. Since k' is below k_1, k' is below h in $D(f\alpha)$. Since λ' is a legal path in $D(f\alpha)$, by definition 3.2 (3), k' must be reached before $v'_{r_{\lambda'}(h)}$. But, k' appears in λ' after $v'_{r_{\lambda'}(h)}$, contradicting condition (2d) of Definition3.1 for $r_{\lambda'}$.

Subcase 3: $z \neq q'_a$ for any a but $z = k_2$. The details in this case are nearly the same as those in the case that $z = k_1$, so we omit the repetition here.

(2d) Let z be an edge in λ after $v_{r_\lambda(h)}$. We must show that $z \neq h$.

Subcase 1: $z = q'_j$ for some j. By the definition of λ and r_λ, edge q'_j appears in λ after $v_{r_\lambda(h)}$ if and only if it appears in λ' after $r_{\lambda'}$. Thus, $z = q'_j$ appears in λ' after after $v_{r_\lambda(h)}$ and so $j > r_{\lambda'}(h)$. Since λ' is a legal path in $D(f\alpha)$, $z = q'_j \neq h$ as required.

Subcase 2: $z \neq q'_j$ for any j but $z = k_1$.

If k_1 is central, then $k_1 \neq h$ since h is non-central. Thus, we may assume that k_1 is non-central. Assume towards a contradiction that $h = z$. Thus, we have $h = z = k_1$. Since $z \neq q'_j$ for any j, k_1 appears in λ because of the fact that k' appeared in λ'. Now, k' cannot appear in λ' after $v_{r_{\lambda'}(h)}$, so k_1 cannot appear in λ after $r_\lambda(h)$. Since $k_1 = z$, this contradicts the assumption that z is an edge of λ occurring after $r_\lambda(h)$, proving that $h \neq z$.

Subcase 3: $z \neq q'_j$ for any j but $z = k_2$. The details in this case are nearly the same as those in the case that $z = k_1$, so we omit the repetition here.

We now check the conditions of Definition 3.2 to verify that λ is a legal path with respect to r_λ.

3.2 (1) This condition is satisfied as usual.

3.2 (2) Let h be a non-central edge in $D(f)$. Since λ' is a legal path with respect to $r_{\lambda'}$, it follows that $r_{\lambda'}(e) < \infty$. By the definition of r_λ, we have therefore that $r_\lambda(e) < \infty$.

3.2 (3) Consider edge p_i in λ and let m be a non-central edge below h in $D(f)$. We must show that $r_\lambda(m) < i$. Since k_1 and k_2 are bottommost edges in $D(f)$, we may assume that $p_i \neq k_1$ and $p_i \neq k_2$. Therefore, p_i appears in λ' as, say, q'_j.

Now, h is below q'_j in λ', so $r_{\lambda'}(h) < j$. Depending on the relative positions of q'_j, q_x and q_y in λ', we may have $i = j, i = j+1$ or $i = j+2$. But, in any of those cases, $r_\lambda(h)$ is increased over $r_{\lambda'}(h)$ by the same amount that i is increased over j. Therefore $r_\lambda(h) < i$, as required.

Case B: The edge k appears only once in λ'. The details of this case are similar to the details of the case in which k appears twice, with the single simplifying difference that the edge k must be replaced only once in the path λ to produce the path λ'. We omit the repeated details of the proof here.

Case C: Edge k does not appear in λ'. In this case, it is straightforward to verify that λ is a legal path in $D(f)$ with respect to the reaching function $r_{\lambda'}$.

Subcase 2: There exists an exposed upper edge, k, emanating to the right from v_0. Denote by k_1 and k_2 the central edges below k. Multiplication by x_1^{-1} deletes edge k. Then $E(f\alpha) = e - 1$, and it is sufficient to show that $P(f\alpha) \geq p$. We assume toward contradiction that $P(f\alpha) < p$, so there exists a legal path $\lambda' = v'_0 q'_1 v'_1 q'_1 \cdots q'_l v'_l$ with respect to reaching function $r_{\lambda'}$ with $l < p$ in $D(f\alpha)$. Define the path λ in $D(f)$ by $\lambda = v_0 q_1 v_1 q_2 v_2 \cdots q_l v_l$ with $v_j = v'_j$, $q_j = q'_j$ if $q'_i \neq k'$, and $q_i = k$ if $q'_i = k'$ for all $0 \leq j \leq l$. We claim that λ is a legal path in $D(f)$ with respect to the reaching function $r_{\lambda'}$. Since the length of λ is less than $P(f)$, we have the desired contradiction, proving that $P(f\alpha) \geq p$. The details involved in this case are similar to those just presented in Subcase 1, so we omit them here.

Case 6: $\alpha = y_1^{-1}$. Multiplication by α in this case acts similarly to Case 5. The only difference is that k emanates to the left of v_0. The proof that $\phi(f) - 1 \leq \phi(f\alpha)$ in this case follows the proof in the case that $\alpha = x_1^{-1}$ with the roles of “left" and “right" exactly reversed, so we omit the repeated details here. This finishes the proof that ϕ satisfies Property 3 of Lemma 3.7.

Property 4 of Lemma 3.7:

Let $f \in F$, $E(f) = e$ and $P(f) = p$. Thus $\phi(f) = e + p$. We must show that there exists an $\alpha \in Z_1 \cup Z_1^{-1}$ such that $\phi(f\alpha) = \phi(f) - 1$. Let $\lambda = v_0q_1v_1q_2v_2...q_pv_p$ be a legal path in $D(f)$ of minimal length, and let r_λ be a reaching function with respect to which λ is a legal path.

Case 1: q_1 is a lower or central edge.

Subcase A: There exists a non-central edge, g with $r_\lambda(g) = 0$.

If g is an upper edge to the right of v_0, take $\alpha = x_1^{-1}$. If g is an upper edge to the left of v_0, take $\alpha = y_1^{-1}$. If g is a lower edge to the right of v_0, take $\alpha = x_1$. If g is a lower edge to the left of v_0, take $\alpha = y_1$. In any of these cases, multiplying f on the right by the corresponding α results in deletion of edge g. Therefore, $E(f\alpha) = e - 1$ We claim that λ is a legal path for $D(f\alpha)$ with respect to r_λ, which shows that $P(f\alpha) \leq p$. This follows from the fact that edge g cannot appear in λ because $r_\lambda(e) = 0$. But by Property 3, which has been proven, $\phi(f\alpha) \geq \phi(f) - 1$. But, in this case, $E(f\alpha) = E(f\alpha) - 1$. So, $P(f\alpha) \geq P(f) = p$. Therefore, $P(f\alpha) = p$, and $\phi(f\alpha) = \phi(f) - 1$.

Subcase B: There does not exist any non-central edge with $r_\lambda(e) = 0$.

If q_1 lies to the left of v_0, take $\alpha = x_0$. If q_1 lies to the right of v_0, take $\alpha = x_0^{-1}$. In either case, multiplication of f on the right by the corresponding α will move the bottom label of $\frac{1}{2}$ from v_0 to v_1. Thus, $E(f\alpha) = e$, and we must show that $P(f\alpha) = p - 1$. We claim that $\lambda' = v_1q_2v_2...q_pv_p$ is legal path for $D(f\alpha)$ with respect to the reaching function $r_{\lambda'}$ defined by $r_{\lambda'}(e) = r_\lambda(e) - 1$. This follows from the fact that since $0 \leq r_{\lambda'}(e) < \infty$ for all non central edges e, we have $1 \leq r_{\lambda'}(e) < \infty$ for all non-central edges. Since all edges are reached by λ' at a vertex numbered one less than the one at which they were reached by λ, all of the conditions of Definitions 3.1 and 3.2 are easily verified. Thus, $P(f\alpha) \leq p - 1$. But by Property 3, and the fact that $E(f\alpha) = E(f)$, we have $P(f\alpha) \geq P(f) - 1$. Therefore, $P(f\alpha) = p - 1$, as required.

Case 2: q_1 is an upper edge.

Subcase A: There exists a non-central edge, g, with $r_\lambda(e) = 0$.

Details of this case are very similar to those of Case 1 Subcase A above, so we omit the repetition.

Subcase B: No non-central edge e satisfied $r_\lambda(e) = 0$.

Since q_1 is the first edge in the path, every non-central edge below q_1 must be reached at v_0. Since no edge is reached by λ at v_0, q_1 must be an exposed upper edge. If q_1 lies to the left of v_0, set $\alpha = y_1^{-1}$. If q_1 lies to the right of v_0, set $\alpha = x_1^{-1}$. In either case, multiplication of f on the right by the corresponding α will delete edge q_1 and the vertex below q_1 that is adjacent to v_0 along to the central line, say w_a and replace the two central edges incident to w_a by a single edge, say q_1'. Thus, $E(f\alpha) = E(f) - 1$ and it suffices to show that $P(f\alpha) = P(f) = p$. Consider the path $\lambda' = v_0 q_1' v_1 q_2 \cdots q_p v_p$. We claim that λ' is a legal path for $D(f\alpha)$ with respect to the reaching function r_λ. The details in the proof of this claim are similar to the details in the previous proofs, so we omit them here. This proves that $P(f\alpha) \leq p$. Now, Property 3, together with the fact that $E(f\alpha) = E(f) - 1$ imply that $P(f\alpha) = p$, as required.

Therefore, there exists an $\alpha \in Z_1 \cup Z_1^{-1}$ such that for all $f \in F$, $\phi(f\alpha) = \phi(f) - 1$.

We have now shown that the formula $\phi_{Z_1}(f) = E(f) + P(f)$ satisfies all four properties of Fordham's Lemma 3.7. Thus, $\phi_{Z_1} = l_{Z_1}$.

□

4. Dead ends

In this section, we conclude by sketching an argument showing F contains dead end elements with respect to Z_1. We claim that the the element $g \in F$ whose wave diagram is pictured in Figure 17 is a dead end with respect to Z_1. The bold edges and arrows indicate a path λ used in determining $l(g)$.

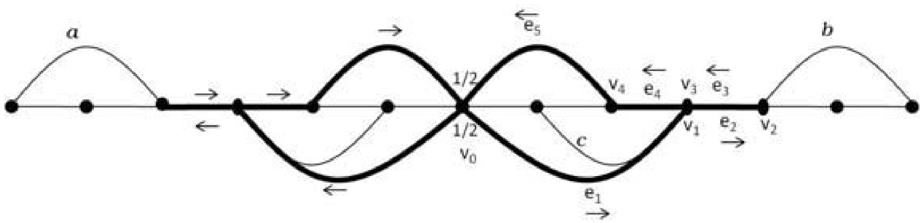

FIGURE 17. Dead end of depth one

It is easy to see that $E(g) = 8$. Consider the length 10 path $\lambda = (v_0 e_1 v_1 e_1 v_3 \ldots v_9 e_{10} v_{10})$ in $D(g)$ starting with e_1 and traversing the bold edges in the directions indicated by the arrows. In order not to clutter the diagram too much, we have omitted the vertex and edge names after e_5. If we define the function r_λ by $r_\lambda(e) = i$ if v_i is the first vertex in λ incident to e, then r_λ is a reaching function for λ with respect to which λ is a legal path. We provide a short explanation as to why λ is the minimal length legal path for g.

Suppose towards a contradiction that $\lambda' = (u_0 f_1 u_2 f_2 \dots f_k u_k)$ is a legal path in $D(g)$ of length $k < 10$. The vertex v_0 divides the diagram into two halves, the "left" half and the "right" half that are isomorphic by a direction-reversing isomorphism. Since λ' has length at most 9, one half of $D(g)$ contains 4 or fewer edges from λ'. Without loss of generality, we assume that it is the right half. Since λ' must reach edge b, we see that the edge e_2 must occur twice, once in each direction, in λ'. This leaves only two more edges of λ' in the right half of $D(g)$. These must be e_1 once in each direction, and the above four edges must form a consecutive subpath $\beta = (v_0 e_1 v_1 e_2 v_2 e_2^{-1} v_2 e_1^{-1} v_0)$ in λ. Now, the edge e_1 cannot be reached by λ' before the end of β since it occurs as the last edge in β. Therefore, the edge c cannot be reached before the end of β since e_1 is below c. However, the rest of λ is to the left of v_0, and thus cannot reach c, contradicting the fact that, as a legal path, λ' reaches every non-central edge of $D(g)$. Therefore, the minimal length of a legal path in $D(g)$ is 10, and so $l(g) = 18$.

To show that g is a dead end element of F with respect to Z_1, we consider the effect of multiplying g by elements of $Z_1 \cup Z_1^{-1}$.

First, we consider multiplication of g by x_0^{-1}. The diagram of gx_0^{-1} is pictured in Figure 18. We see that $E(gx_0^{-1}) = 8$ and the path indicated by the bold edges is a legal path of length 9. Therefore, $l_{Z_1}(gx_0^{-1}) \leq 17$.

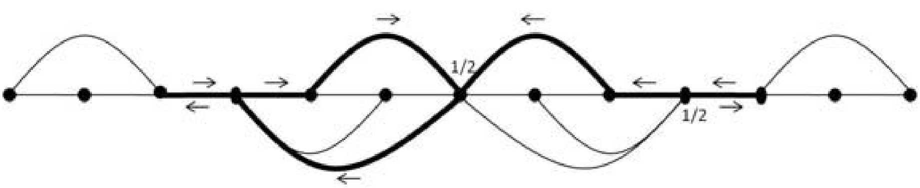

FIGURE 18. Diagram for gx_0^{-1}

Next, we consider multiplication of g by x_1. The diagram of gx_1 is pictured in Figure 19. We see that $E(gx_1) = 7$ and the path indicated by the bold edges is a legal path of length 11. Therefore, $l_{Z_1}(gx_1) \leq 18$.

Next we consider multiplication of g by x_1^{-1}. The diagram of gx_1^{-1} is pictured in Figure 20. We see that $E(gx_1^{-1}) = 9$ and the path indicated by the bold edges is a legal path of length 9. Therefore, $l_{Z_1}(gx_1^{-1}) \leq 18$.

By the symmetry of g and the generators x_1 and y_1, the arguments that y_1 and y_1^{-1} do not increase the length of g are similar to those for

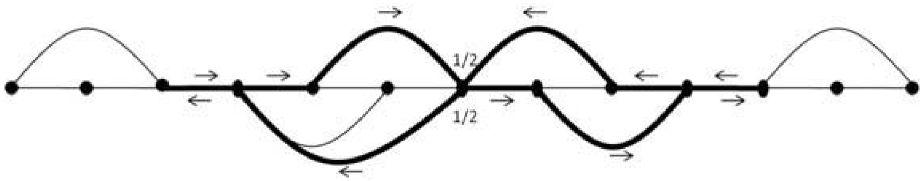

FIGURE 19. Diagram for gx_1

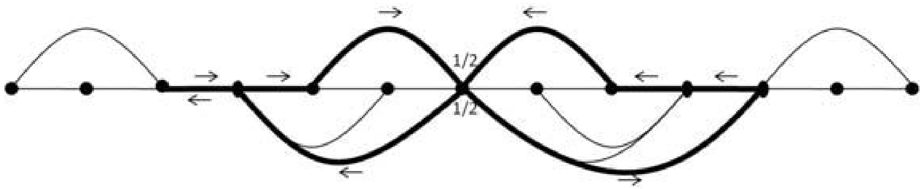

FIGURE 20. Diagram for gx_1^{-1}

x_1 and x_1^{-1}. Therefore, $l_{Z_1}(g\alpha_1) \leq 18$ for all $\alpha_i \in Z_1$, showing that g is a dead end and that F has dead end depth at least 2 with respect to Z_1.

References

[1] G. N. Arzhantseva, V. S. Guba, and M. V. Sapir, *Metrics on diagram groups and uniform embeddings in a Hilbert space*, Comment. Math. Helv. **81** (2006), no. 4, 911–929. MR 2271228 (2007k:20084)

[2] James Belk and Kai-Uwe Bux, *Thompson's group F is maximally nonconvex*, Geometric methods in group theory, Contemp. Math., vol. 372, Amer. Math. Soc., Providence, RI, 2005, pp. 131–146. MR MR2139683 (2006f:20050)

[3] ———, *Thompson's group F is maximally nonconvex*, Geometric methods in group theory, Contemp. Math., vol. 372, Amer. Math. Soc., Providence, RI, 2005, pp. 131–146. MR MR2139683 (2006f:20050)

[4] O. V. Bogopolʹskiĭ, *Infinite commensurable hyperbolic groups are bi-Lipschitz equivalent*, Algebra i Logika **36** (1997), no. 3, 259–272, 357. MR 1485595 (98h:57002)

[5] J. W. Cannon and Floyd, *What is... thompson's group?*, Notices Amer. Math. Soc. **58** (2011), no. 8, 1112–1113.

[6] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Introductory notes on Richard Thompson's groups*, Enseign. Math. (2) **42** (1996), no. 3-4, 215–256. MR 1426438 (98g:20058)

[7] Sean Cleary and Jennifer Taback, *Combinatorial properties of Thompson's group F*, Trans. Amer. Math. Soc. **356** (2004), no. 7, 2825–2849 (electronic). MR 2052598 (2005b:20074)

[8] S. Blake Fordham, *Minimal length elements of Thompson's group F*, Geom. Dedicata **99** (2003), 179–220. MR MR1998934 (2004g:20045)

[9] Victor Guba and Mark Sapir, *Diagram groups*, Mem. Amer. Math. Soc. **130** (1997), no. 620, viii+117. MR MR1396957 (98f:20013)

[10] Matthew Horak, Melanie Stein, and Jennifer Taback, *Computing word length in alternate presentations of Thompson's group F*, Internat. J. Algebra Comput. **19** (2009), no. 8, 963–997. MR 2596882 (2011c:20077)

[11] ———, *A note on convexity properties of Thompson's group F*, J. Group Theory (2011), to appear.

[12] Russell Lyons, Robin Pemantle, and Yuval Peres, *Random walks on the lamplighter group*, Ann. Probab. **24** (1996), no. 4, 1993–2006. MR 1415237 (97j:60014)

Contact information

M. Horak
University of Wisconsin-Stout
Menomonie, WI 54751
USA
E-Mail: horakm@uwstout.edu
URL: www.uwstout.edu/faculty/horakm/index.cfm

A. Johnson
Grand Valley State University
Allendale, MI 49401-9403
USA
E-Mail: johnsoa1@mail.gvsu.edu

A. Stonesifer
St. Olaf College
Northfield, Minnesota 55057
USA
E-Mail: stonesif@stolaf.edu

Received by the editors: 05.12.2011
and in final form 07.09.2012.