http://topology.auburn.edu/tp/

TOPOLOGY PROCEEDINGS

Volume 46, 2015

Pages 181-204

http://topology.nipissingu.ca/tp/

Visual Boundaries of Diestel-Leader Graphs

by

Keith Jones and Gregory A. Kelsey

Electronically published on June 6, 2014

Topology Proceedings

Web: http://topology.auburn.edu/tp/
Mail: Topology Proceedings
Department of Mathematics & Statistics
Auburn University, Alabama 36849, USA
E-mail: topolog@auburn.edu
ISSN: 0146-4124

http://topology.auburn.edu/tp/
http://topology.nipissingu.ca/tp/

TOPOLOGY PROCEEDINGS

Volume 46 (2015)
Pages 181-204

E-Published on June 6, 2014

VISUAL BOUNDARIES OF DIESTEL-LEADER GRAPHS

KEITH JONES AND GREGORY A. KELSEY

Abstract. Diestel-Leader graphs are neither hyperbolic nor CAT(0), so their visual boundaries may be pathological. Indeed, we show that for $d > 2$, $\partial\mathrm{DL}_d(q)$ carries the indiscrete topology. On the other hand, $\partial\mathrm{DL}_2(q)$, while not Hausdorff, is T_1, totally disconnected, and compact. Since $\mathrm{DL}_2(q)$ is a Cayley graph of the lamplighter group L_q, we also obtain a nice description of $\partial\mathrm{DL}_2(q)$ in terms of the lamp stand model of L_q and discuss the dynamics of the action.

1. Introduction

The *visual boundary* ∂M of a complete CAT(0) metric space M is the topological space obtained by giving the set of asymptotic equivalence classes of geodesic rays in M the compact-open topology [2, Ch. II.8]. For any base point $p \in M$, one can simply take $\partial_p M$ to be the set of geodesic rays emanating from p, and $\partial_p M$ and ∂M are homeomorphic. In this setting, the visual boundary has nice properties; for instance, $M \cup \partial M$ is contractible, and if M is *proper* (a metric space in which every closed metric ball is compact), then ∂M provides a compactification of M under the "Cone topology." An action of a group G by isometries on M can be extended to an action by homeomorphisms on ∂M, and studying the dynamics of this action can prove quite fruitful. One can define the visual boundary more generally (i.e., outside the context of CAT(0) spaces) and

2010 *Mathematics Subject Classification.* Primary 20F65, 20F69; Secondary 20E22, 05C25.

Key words and phrases. visual boundary, Diestel-Leader graphs, lamplighter groups.

Support for the first author from a summer research grant from SUNY Oneonta is gratefully acknowledged.

ask whether these nice properties still arise or whether the study of the action on the boundary is still fruitful.

When a group G acts geometrically on a space, one may take the boundary of the space as a boundary of the group. For word hyperbolic groups, the visual boundary is unique and has proven very useful [8]. Outside this class of groups, the situation is not so nice. Christopher B. Croke and Bruce Kleiner have shown that even CAT(0) groups may not have unique visual boundaries [4]. Even worse, outside this context one may run into pathological situations. In [7], it is shown that the visual boundary of the Cayley graph of $\mathbb{Z}^2$ with respect to the standard generating set is uncountable, yet it has the indiscrete (a.k.a. trivial) topology. In short, this occurs because one is able to play the asymptotic classes (two rays are equivalent if they are close in the long term) against the compact-open topology (two rays are close if they agree in the short term) to obtain a sequence of asymptotic rays representing an arbitrary point of the boundary and whose limit is another arbitrary point of the boundary.

This paper investigates whether visual boundaries for non-hyperbolic, non-CAT(0) groups may carry interesting topologies, and if so what this might tell us about those groups. In particular, we study the family of lamplighter groups $L_q = \mathbb{Z}_q \wr \mathbb{Z}$, $q \geq 2$ an integer. Using the appropriate generating set, one obtains a particularly nice Cayley graph for L_q, called the Diestel-Leader graph $\mathrm{DL}_2(q)$ [1], [13, §2]. The boundary $\partial\mathrm{DL}_2(q)$ is not a canonical boundary for L_q, since using a different Cayley graph might give rise to a different boundary. However, $\partial\mathrm{DL}_2(q)$ is appealing in that it can be well understood using the standard "lamp stand" model for the lamplighter group. This model and the (well-known) geometry of Diestel-Leader graphs provide ample tools for studying the visual boundary.

More generally, the Diestel-Leader graph $\mathrm{DL}(q_1, q_2, \ldots, q_d)$ can be realized as a subspace of the product of d trees, having respective valence $q_1 + 1$, $q_2 + 1$, $\ldots$, $q_d + 1$ [1]. The notation $\mathrm{DL}_d(q)$ is used when each tree has valence $q + 1$. While we will state our results only for the cases when the degrees are equal, allowing the degrees to vary between trees has no effect on our analysis. Recently, Melanie Stein and Jennifer Taback have described the metric on these graphs [11], and Moon Duchin, Samuel Lelièvre, and Christopher Mooney have discussed geodesics in $\mathrm{DL}_d(q)$ in their work on sprawl [5]. While not all Diestel-Leader graphs are Cayley graphs [6, Theorem 1.4], [1, Corollary 2.15], this paper discusses the relation to L_2 when applicable and also has results that apply in the non-Cayley graph case. Because Diestel-Leader graphs inherit much of their structure from trees, which are prototypical CAT(0) and hyperbolic spaces, it seems natural to ask whether the boundaries of such graphs inherit any nice properties from the boundaries of trees.

In §2, we provide some background on visual boundaries, lamplighter groups, and Diestel-Leader graphs.

In §3, we collect some basic facts about geodesic rays in Diestel-Leader graphs and prove the following.

Theorem A (Corollary 3.8)**.** *As a set, $\partial DL_2(q)$ is a disjoint union of two punctured Cantor sets.*

In §4, we prove the following theorem.

Theorem B. *$\partial DL_2(q)$ is not Hausdorff, but it is T_1, compact, and totally disconnected.*

The proof of this theorem is collected in Observation 4.10, Observation 4.12, Proposition 4.13, and Observation 4.14. Additionally, we discuss the dynamics of the action by L_2 on $\partial \mathrm{DL}_2(q)$ in Theorem 4.17 and Corollary 4.18.

In §5 we discuss the geometry of geodesics in $\mathrm{DL}_d(q)$ and prove the following.

Theorem C (Theorem 5.13)**.** *For $d > 2$, the topology of $\partial DL_d(q)$ is indiscrete.*

Roughly speaking, this is a consequence of the additional degree of freedom a third tree provides. Along the way, we establish in Theorem 5.4 a strong restriction on the kinds of paths in $\mathrm{DL}_d(q)$ which may be geodesics.

One feature of CAT(0) and hyperbolic spaces is that the horofunction[1] boundary $\partial_h X$ is naturally homeomorphic to ∂X [2, II.8.13] since all horofunctions are Busemann functions (i.e., they come from geodesic rays). However, outside this setting, one may find horofunctions which are not Busemann functions [12]. An investigation of the horofunction boundary will appear in a forthcoming work, where we will show that while $\partial \mathrm{DL}_2(q)$ embeds in $\partial_h \mathrm{DL}_2(q)$, there are many horofunctions which are not Busemann functions.

[1] In short, for a metric space X, a *horofunction* is a point in $C(X)$ (the space of continuous functions on X with the topology of compact convergence on bounded subsets) which is a limit of a sequence of functions $d_y(x) = d_X(x, y)$, $y \in X$. Horofunctions represent points of the horofunction boundary $\partial_h(X)$. Busemann functions are those horofunctions obtained as a limit of points along a geodesic ray in X. See [2, II.8.12-14] for details.

2. Background

2.1. The visual boundary.

Let X be a geodesic space with base point x_0. Two geodesic rays $\gamma, \gamma' : [0, \infty) \to X$ are said to be *asymptotic* if there is a $\lambda \geq 0$ such that for all $t \geq 0$, $d(\gamma(t), \gamma'(t)) \leq \lambda$. The *visual boundary* of X is the space ∂X consisting of all asymptotic equivalence classes of geodesic rays in X, endowed with the quotient topology from the topology of uniform convergence on compact sets. The *based visual boundary* of X with base point x_0, denoted $\partial(X, x_0)$, is the same topology restricted to the subset of geodesic rays emanating from x_0. In general, the based and unbased visual boundaries need not agree.

In this paper we consider (based) Diestel-Leader graphs, which are known to be vertex transitive [1, Proposition 2.4], so the based visual boundary is independent of base point. In Proposition 3.11, we show that in $\mathrm{DL}_2(q)$, the based and unbased boundaries are the same, so throughout sections 3 and 4, we abuse notation and use $\partial\mathrm{DL}_2(q)$ to refer to the based visual boundary of $\mathrm{DL}_2(q)$. In §5, we still abuse notation and use $\partial\mathrm{DL}_d(q)$ to refer to the based visual boundary, even though we do not consider whether it is the same as the unbased visual boundary.

2.2. The Diestel-Leader graph $\mathrm{DL}_d(q)$.

For an integer q, let T be the regular $q+1$-valent simplicial tree. Following [13, §2] and [11, §2], we orient the edges of T so that each vertex v has exactly one *predecessor* v^- and q *successors*. This induces a partial ordering on the set of vertices of T, under which any two vertices v and w have a greatest common ancestor $v \curlywedge w$. Choosing a base vertex o in T allows us to define a height function $h(v) = d_T(v, v \curlywedge o) - d_T(o, v \curlywedge o)$, where the function d_T measures distance in T when each edge is given length 1. The partial ordering provides a chosen endpoint ω of T, obtained by any geodesic ray that always follows predecessors, and this height function is the Busemann function for ω corresponding to the ray emanating from o. For a vertex v in the horocycle $H_k = \{v \in T \mid h(v) = k\}$, its unique predecessor v^- is in H_{k-1}, and each of its q successors is in H_{k+1} (see Figure 1). In particular, for a given initial vertex v, there is a unique "downward" path of length k, for each k, and a unique downward ray: that which leads to ω.

We now define $\mathrm{DL}_d(q)$. Let $T_0, T_1, \ldots, T_{d-1}$ be copies of T, with base points $o_i \in T_i$. The Diestel-Leader graph $\mathrm{DL}_d(q)$ is the graph whose vertex set consists of d-tuples $v = (x_0, x_1, \ldots, x_{d-1})$, $x_i \in T_i$ a vertex, such that $\sum_{i=0}^{d-1} h(x_i) = 0$. Let $h_i(v)$ denote the height function $h(x_i)$ on T_i. There is a natural base point $(o_0, o_1, \ldots, o_{d-1})$ for $\mathrm{DL}_d(q)$.

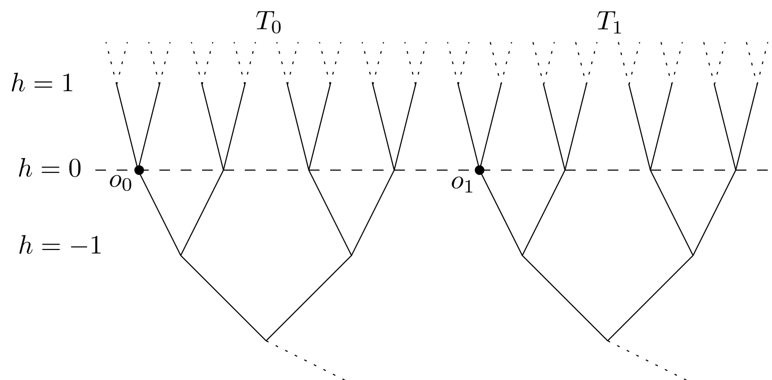

FIGURE 1. A neighborhood of the origin in $\mathrm{DL}_2(2)$

The edges of $\mathrm{DL}_d(q)$ correspond to pairs $(v, w) = ((x_0, ..., x_{d-1}), (y_0, ..., y_{d-1}))$ such that there are i and j, $i \neq j$, with an edge joining x_i to y_i in T_i, an edge joining x_j to y_j in T_j, and $x_k = y_k$ for all $k \neq i, j$. The relation

$$h(y_i) - h(x_i) = \pm 1 = h(x_j) - h(y_j)$$

follows from the definition of vertices of $\mathrm{DL}_d(q)$. Thus, moving along an edge in $\mathrm{DL}_d(q)$ means simultaneously choosing one tree in which to increase height and choosing another tree in which to decrease height, while holding the position constant in every other tree.

Convention. We adopt the convention that all geodesic rays in $\mathrm{DL}_d(q)$ under discussion emanate from o, unless otherwise stated.

2.3. Lamplighter groups.

The Diestel-Leader graph $\mathrm{DL}_2(q)$ is the Cayley graph of the lamplighter group $L_q = \mathbb{Z}_q \wr \mathbb{Z}$ with generating set $\{t, at, a^2t, ..., a^{q-1}t\}$ (t is the generator of $\mathbb{Z}$ in the wreath product and a is the generator of $\mathbb{Z}_q$). Each element of the lamplighter group (and thus each vertex of $\mathrm{DL}_2(q)$) is associated with a "lamp stand." In the case $q = 2$, the lamp stand consists of a row of lamps in bijective correspondence with $\mathbb{Z}$, a finite number of which are lit, and a lamplighter positioned at one of the lamps. If $q > 2$, then the lamps have q settings; these can be interpreted as off and $q - 1$ levels of brightness while lit or $q - 1$ different colors [10]. The Diestel-Leader graph $\mathrm{DL}_3(q)$ has a similar interpretation, except that there is a rhombic grid of lamps [3].

We should also note that for $d > 3$, if q has a prime factor p such that $p < d - 1$, it is open whether $\mathrm{DL}_d(q)$ is a Cayley graph of some group. If

q has no such prime factor, then $\mathrm{DL}_d(q)$ is a Cayley graph [1, Corollary 3.17]. If q does have such a prime factor, it is only known that $\mathrm{DL}_d(q)$ is quasi-isometric to a Cayley graph [1, Corollary 3.21].

In the $d = 2$ case, the base vertex (o_0, o_1) corresponds to the lamp stand with no lit lamps and the lamplighter at position 0. In general, the lamp stand corresponding to vertex (x_0, x_1) has the lamplighter at position $h(x_0)$. An edge in $\mathrm{DL}_2(q)$ corresponds to the lamplighter stepping between adjacent lamps. If the edge is associated with generator t, then the lamplighter moves without toggling any bulbs. If the edge is associated with generator at, then the lamplighter toggles the bulb before he leaves (if he is moving in the positive direction) or after he arrives (if he is moving in the negative direction) [13].

So, for example, the word $t^3(at)t^{-2}(at)^{-2}t^{-1}$ corresponds with the lamplighter starting at position 0, moving three to the right, lighting lamp 3, stepping to position 4, stepping back to position 2, then stepping back to light lamps 1 and 0, and finally stepping back to lamp -1. The end result is the lamp stand pictured in Figure 2. Notice that this lamp stand is also obtainable by the word $(at)^2t(at)t^{-5}$, and so these words represent the same element of L_2.

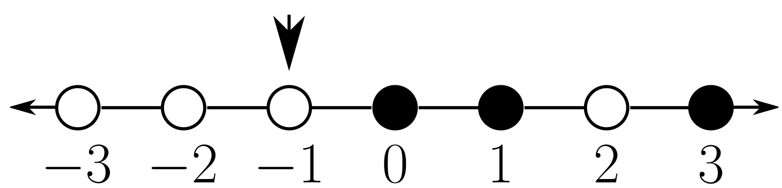

FIGURE 2. Lamp stand example for $t^3(at)t^{-2}(at)^{-2}t^{-1}$

Multiplication of elements corresponds to "composition" of lamp stands. To compute the lamp stand for the element $g \cdot h$, take the lamp stand for g, then have the lamplighter perform the same toggling of lamps as in h, but starting from the lamplighter's end position in g instead of at 0. So, for example, for g the lamp stand in Figure 2, $g \cdot t$ would have the same set of lit lamps, but the lamplighter would be at 0 instead of -1. The lamp stand for $g \cdot tat$ would have the lamplighter at position 1 and lit lamps only at positions 1 and 3.

Note that L_q has presentation

$$\langle a, t \mid [t^i a t^{-i}, t^j a t^{-j}] = a^q = 1 \text{ for all } i, j \in \mathbb{Z} \rangle,$$

from which we obtain the epimorphism $exp_t : L_2 \to \mathbb{Z}$ mapping $a \mapsto 0$ and $t \mapsto 1$, recording the exponent sum of t for a given element of L_2.

3. The Boundary of $\mathbf{DL}_2(q)$

3.1. Using projections to bound distance.

We begin with two observations that apply in the general case. Since a path in $\mathrm{DL}_d(q)$ projects to a path in a tree T_i, we have a lower bound on the distance between two vertices.

Observation 3.1. *Let $v = (v_0, v_1, \ldots, v_{d-1})$ and $w = (w_0, w_1, \ldots, w_{d-1})$ be two vertices in $DL_d(q)$. Then $d(v, w) \geq \max\{d_{T_i}(v_i, w_i) \,|\, 0 \leq i \leq d-1\}$.*

Moreover, there is a simple upper bound on the distance as well.

Lemma 3.2. *Let $v = (v_0, v_1, \ldots, v_{d-1})$ and $w = (w_0, w_1, \ldots, w_{d-1})$ be two vertices in $DL_d(q)$. Then*

$$d(v, w) \leq \sum_{i=0}^{d-1} d_{T_i}(v_i, w_i).$$

Proof. We may assume v is the origin o, since any path between vertices may be translated via isometry to a path from the origin. We will now construct a path in $\mathrm{DL}_d(q)$ from o to w that has length at most $\sum_{i=0}^{d-1} d_{T_i}(o_i, w_i)$.

Let k be the index of a tree such that $h_k(w)$ is minimal. Then $h_k(w) \leq 0$. For each tree T_i other than T_k, in turn, follow the path from o_i to w_i, always compensating in tree T_k (i.e., moving up in T_k when moving down in T_i, and vice versa), following the rule that when the current vertex in T_k has negative height and we must move up in T_k, we choose to stay on the ray from o_k to the distinguished end ω_k. After all trees other than T_k have been so traversed, a total distance of $\sum_{i=0, i\neq k}^{d-1} d_{T_i}(o_i, w_i)$ has been traveled. At this point, the current vertex $x \in T_k$ lies at height $h_k(w)$ and along the ray from o_k to ω_k, implying that x lies on the geodesic from o_k to w_k, and so $d_{T_k}(x, w_k) \leq d_{T_k}(o_k, w_k)$. Thus, we can move from o to w taking no more than $\sum_{i=0}^{d-1} d_{T_i}(o_i, w_i)$ steps. We can compensate for the steps that move T_k into position using any one other tree. Those compensating steps will undo themselves, and the vertex in that tree at the end will be the same as it was before. □

3.2. Asymptotic equivalence classes.

Definition 3.3. We can project a path $\gamma = (v_0, v_1, \ldots, v_n)$ through $\mathrm{DL}_d(q)$ to a tree T_j. Let $v_i = (x_{i,0}, x_{i,1}, \ldots, x_{i,d-1})$ be the i^{th} vertex along γ. Then $\gamma^{(j)}$ is the corresponding path $(x_{0,j}, x_{1,j}, \ldots, x_{n,j})$ through T_j. We will refer to $\gamma^{(j)}$ as the *projection* of γ to tree T_j.

For the rest of this section, we restrict our attention to $\mathrm{DL}_2(q)$, though the notation and terminology we use will extend to the general case.

Let p be a path between vertices v and w in $\mathrm{DL}_2(q)$. Thinking of p as a sequence $(v = v_0, v_1, \dots, v_n = w)$ of vertices of $\mathrm{DL}_2(q)$, we refer to a subsequence v_i, v_{i+1}, v_{i+2} such that $h_j(v_i) = h_j(v_{i+1}) + 1 = h_j(v_{i+2})$ as a *turn* in p, a *bottoming out* in the j^{th} tree. In the lamp stand interpretation for $d = 2$, bottoming out corresponds with the lamplighter turning around and moving in the opposite direction along the row of lamps; if the path bottoms out in T_0, then the lamplighter stops moving to the left (towards $-\infty$) along the lamp stand and begins moving to the right, and vice versa for T_1.

In [5, Figure 1] it is shown that a geodesic in $\mathrm{DL}_2(q)$ has at most two turns. However, the case of geodesic rays is simpler.

Lemma 3.4. *There are no 2-turn geodesic rays in $DL_2(q)$.*

Proof. Let γ be a ray in $\mathrm{DL}_2(q)$ emanating from o with two turns and no backtracking. Then γ "bottoms out" once in T_i, $i \in \{0, 1\}$, and then once in T_{1-i}. Let $k \in \mathbb{Z}$, $k > 0$, be the distance traveled before the first turn, so that γ bottoms out in T_i at $\gamma(k)$ with heights $h_i(\gamma(k)) = -k$ and $h_{1-i}(\gamma(k)) = k$. Let $l \in \mathbb{Z}$, $l > 0$, be the distance traveled before the second turn, so that γ bottoms out in T_{1-i} at $\gamma(k + l)$ with heights $h_i(\gamma(k + l)) = l - k$ and $h_{1-i}(\gamma(k + l)) = k - l$. Because γ has exactly two turns, γ then proceeds to descend eternally in T_i while ascending in T_{1-i}. Consider the vertex $\gamma(k + 2l)$. In T_i, γ descends to $\gamma^{(i)}(k)$, then ascends for l edges followed by a descent of the same distance, so that $\gamma^{(i)}(k) = \gamma^{(i)}(k + 2l)$.

We now show there is a path through $\mathrm{DL}_2(q)$ from o that arrives at $\gamma(k + 2l)$ in shorter time. It may be helpful to refer to Figure 3, which provides an example. Let d be the distance in T_{1-i} from o_{1-i} to $\gamma^{(1-i)}(k+2l)$. Then either $d = k$ if $k \geq l$ or $d = 2l - k$ if $l > k$. Choose a path γ' from o to $\gamma(k + 2l)$ so that $\gamma'^{(1-i)}$ is the geodesic in T_{1-i} from o_{1-i} to $\gamma^{(1-i)}(k + 2l)$, and so that $\gamma^{(i)}$ changes height appropriately with each edge in $\gamma^{(1-i)}$. The actual choice of $\gamma^{(i)}$ is irrelevant; regardless of how it ascends, it must then descend to $\gamma^{(i)}(k + 2l) = \gamma^{(i)}(k)$. Thus, $\gamma'(d) = \gamma(k + 2l)$ and $d < k + 2l$, so γ cannot be a geodesic ray. $\square$

It is worth noting that because there are 2-turn geodesic paths in $\mathrm{DL}_2(q)$, Lemma 3.4 implies that $\mathrm{DL}_2(q)$ is not *geodesically complete*; i.e., there are geodesics which cannot be extended to geodesic rays. This fact is proved in [11, Theorem 12], where the authors demonstrate that $\mathrm{DL}_d(q)$ has *dead-end elements*.

Let γ be a geodesic ray based at the origin vertex $o \in \mathrm{DL}_2(q)$, and suppose γ first descends to height $-h \leq 0$ in T_i before ascending eternally

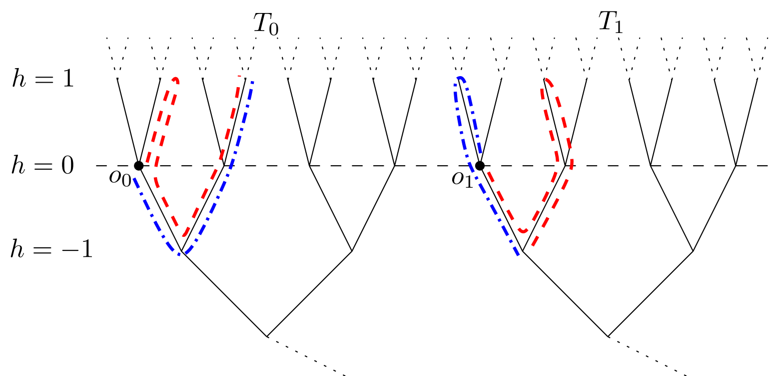

FIGURE 3. Two paths in DL(2, 2) having the same endpoints: a 1-turn path which is geodesic and a 2-turn path which is minimally non-geodesic.

in T_i. Then $\gamma^{(i)}|_{[0,h]}$ is fixed, while $\gamma^{(1-i)}|_{[0,h]}$ follows one of q^h paths. On the other hand, $\gamma^{(i)}|_{[h,\infty]}$ chooses some endpoint other than ω_i of T_i, while $\gamma^{(1-i)}|_{[h,\infty]}$ must approach ω_{1-i}.

Lemma 3.5. *If γ and γ' are geodesic rays in $DL_2(q)$ emanating from o which both descend to height $-h \leq 0$ in T_i before turning, then they are asymptotic if and only if $\gamma^{(i)} = \gamma'^{(i)}$. In this case, they eventually merge and upon merging, never split.*

Proof. First, assume that $\gamma^{(i)} = \gamma'^{(i)}$. Because $h_{1-i}(\gamma^{(1-i)}(h)) = h_{1-i}(\gamma'^{(1-i)}(h)) = h$, and both of these vertices have o_{1-i} as an ancestor, we have $\gamma^{(1-i)}(2h) = \gamma'^{(1-i)}(2h) = o_{1-i}$. From this point on, since both $\gamma^{(1-i)}$ and $\gamma'^{(1-i)}$ are descending in T_{1-i}, they are the same. Hence, $\gamma|_{[2h,\infty)} = \gamma'|_{[2h,\infty)}$. So γ and γ' merge at or before distance $2h$, and they are asymptotic.

Now, assume that there exists t such that $\gamma^{(i)}(t) \neq \gamma'^{(i)}(t)$. Since both rays descend to height $-h$, we must have $t > h$. Since γ and γ' are geodesic rays bottoming out at height h in T_i, it follows that $\gamma^{(i)}|_{[h,\infty)}$ and $\gamma'^{(i)}|_{[h,\infty)}$ are geodesic rays in T_i. By assumption, they are not equal, so since T_i is a tree, they are not asymptotic in T_i. From Observation 3.1, γ and γ' are not asymptotic. □

Furthermore, Observation 3.1 also ensures the following.

Observation 3.6. *Let γ and γ' be geodesic rays in $DL_2(q)$.*

(1) *If γ begins by descending in T_i, while γ' begins by ascending in T_i, or vice-versa, then γ and γ' are not asymptotic.*

(2) *If γ and γ' are geodesic rays descending to heights h and h', respectively, before turning, with $h \neq h'$, then γ and γ' are not asymptotic.*

Combining lemmas 3.4 and 3.5 and Observation 3.6, we obtain the following description of asymptotic equivalence classes of geodesic rays in $\partial \mathrm{DL}_2(q)$.

Theorem 3.7. *Two geodesic rays γ and γ' in $DL_2(q)$ are asymptotic if and only if their projections $\gamma^{(0)}$ and $\gamma'^{(0)}$ approach the same end of T_0 and their projections $\gamma^{(1)}$ and $\gamma'^{(1)}$ approach the same end of T_1.*

Corollary 3.8. *The family of geodesic rays whose projections do not approach $\omega_i \in \partial T_i$, $i \in \{0,1\}$, is in one-to-one correspondence with a Cantor set minus the point corresponding to ω_i. Hence, as a set, $\partial DL_2(q)$ is a disjoint union of two deleted Cantor sets:*

$$\partial DL_2(q) = ((\partial T_0 - \omega_0) \times \{\omega_1\}) \coprod ((\partial T_1 - \omega_1) \times \{\omega_0\}).$$

It is perhaps not surprising that $\partial \mathrm{DL}_2(q)$ should be so closely related to a Cantor set, given that $\mathrm{DL}_2(q)$ is a one-dimensional subset of a product of trees. Lemma 3.5 leads to the picture in Figure 4 of a typical element of $\partial \mathrm{DL}_2(2)$.

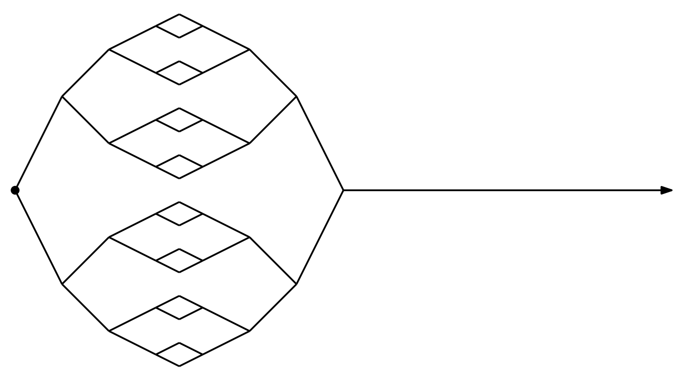

FIGURE 4. A typical asymptotic equivalence class in $\mathrm{DL}_2(2)$.

3.3. LAMP STAND INTERPRETATION OF $\partial \mathrm{DL}_2(2)$.

We can understand the visual boundary using the lamp stand. For each geodesic ray starting at the base point, the lamplighter starts at position 0 on an unlit row of lamps. He starts moving in a direction (always the same direction as the projection of the ray to T_0), perhaps lighting lamps along the way. If the ray "bottoms out" in one tree, then the lamplighter will turn around and proceed in the other direction, again possibly toggling lamps along the way. So each element of the visual boundary corresponds with a lamp stand with the lamplighter standing at either $+\infty$ or $-\infty$. If the lamplighter is at $+\infty$, then the set of lit lamps (if it is non-empty) has a minimum. If the lamplighter is at $-\infty$, then the set of lit lamps (if it is non-empty) has a maximum.

Notice that if the lamplighter turns, he can reset the lamps he has already passed to undo any lighting that he has done or to light any lamps that he missed the first time. In this way, we can see how the "pre-turn" segment of the ray does not affect the asymptotic equivalence class.

Since the height of the associated vertex in tree T_0 is the position of the lamplighter, this means that the points in $(\partial T_0 - \omega_0) \times \omega_1$ have the lamplighter at $+\infty$ and the points in $(\partial T_1 - \omega_1) \times \omega_0$ have the lamplighter at $-\infty$.

The lamp stand interpretation for $\partial \mathrm{DL}_2(q)$ is essentially the same, except that the lamps can take on q different states, instead of simply on and off.

3.4. Action of L_2 on $\partial \mathbf{DL}_2(2)$.

We can compute the action of the lamplighter group L_2 on the visual boundary $\partial \mathrm{DL}_2(q)$ by using the lamp stand interpretation in §3.3. For γ a geodesic ray in $\mathrm{DL}_2(q)$, we write $[\gamma]$ for its asymptotic equivalence class in $\partial \mathrm{DL}_2(q)$. For $g \in L_2$ and $[\gamma] \in \partial \mathrm{DL}_2(2)$, to compute the lamp stand for $g \cdot [\gamma]$, start with the lamp stand for g. Then have the lamplighter perform the lighting prescribed by $[\gamma]$, but starting from the lamplighter's end position in g instead of at position 0. See Figure 5 for an example.

Notice that for any $[\gamma] \in (\partial T_0 - \omega_0) \times \omega_1$ and any $g \in L_2$, we will have $g \cdot [\gamma] \in (\partial T_0 - \omega_0) \times \omega_1$. Similarly, $(\partial T_1 - \omega_1) \times \omega_0$ is also invariant under the action of L_2.

Observation 3.9. *The action of the generators t and at on the lamp stand model for $\partial DL_2(2)$ is as follows:*

- *t shifts the lit lamps one spot to the right (i.e., towards $+\infty$);*
- *t^{-1} shifts the lit lamps one spot to the left (i.e., towards $-\infty$);*
- *at shifts the lamps one spot to the right and then toggles the lamp located at 0;*
- *$(at)^{-1}$ toggles the lamp located at 0 and then shifts the lamps one spot to the left.*

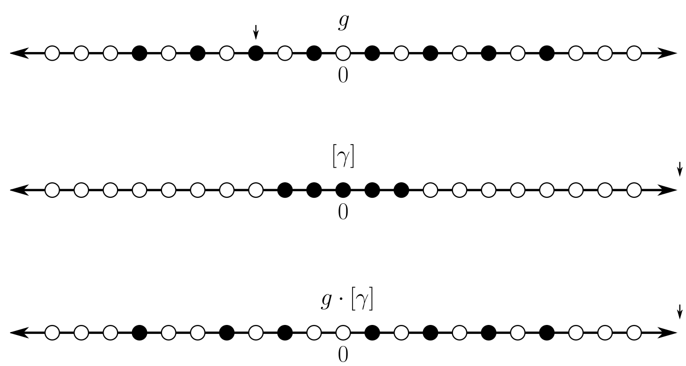

FIGURE 5. The action of an element $g \in L_2$ on an asymptotic equivalence class $[\gamma] \in \partial \mathrm{DL}_2(2)$

For $k \in \mathbb{Z}$, let a_k represent the element $t^k a t^{-k} \in L_2$. Notice that in the lamp stand model for L_2, this is the element associated with only lamp k lit and the lamplighter at position 0.

Observation 3.10. *The action of a_k on the lamp stand model of $\partial DL_2(2)$ is to switch the lamp at position k.*

In §4.6, we use the lamp stand interpretation of $\partial \mathrm{DL}_2(2)$ to compute the dynamics of this action.

3.5. $\partial\mathbf{DL}_2(q)$ WITHOUT A BASE POINT.

In §2.1 we introduced the based and unbased visual boundaries. When X is CAT(0) or δ-hyperbolic, these agree. The following shows that the same is true for $\mathrm{DL}_2(q)$.

Proposition 3.11. *Let γ be a geodesic ray in $DL_2(q)$. Then there exists a geodesic ray τ emanating from the origin which is asymptotic to γ.*

Proof. In one tree T_i, $i \in \{0,1\}$, γ chooses a non-distinguished end $e \neq \omega_i$, and in T_{1-i}, γ approaches ω_{1-i}. Let τ be any geodesic ray emanating from o that approaches $e \in T_i$ and $\omega_{1-i} \in T_{1-i}$. Because the projections $\gamma^{(i)}$ and $\tau^{(i)}$ approach the same end of T_i, they must merge since T_i is a tree; i.e., there exist $r_1, r_2 \in \mathbb{Z}$ such that $\gamma^{(i)}(r_1) = \tau^{(i)}(r_2)$. Similarly, there are $s_1, s_2 \in \mathbb{Z}$ such that $\gamma^{(1-i)}(s_1) = \tau^{(1-i)}(s_2)$. Setting $n = \max\{r_1, r_2, s_1, s_2\}$, one of $\gamma^{(i)}(k)$ and $\tau^{(i)}(k)$ is an ancestor of the other for all $k \geq n$, and the opposite relation holds for $\gamma^{(1-i)}(k)$ and $\tau^{(1-i)}(k)$. The distance from $\gamma(k)$ to $\tau(k)$ is constant, regardless of k, and so the rays are asymptotic. □

4. Topology of $\partial\mathbf{DL}_2(q)$

4.1. Some important sets.

The natural topology on the visual boundary of a space is the topology of uniform convergence on compact sets. Informally, this means that two asymptotic equivalence classes are close if there are representatives of those classes that share a long initial segment. More formally, given a ray γ, a compact subset $[0,k]$ of $[0,\infty)$, and $0<\epsilon<1$, define the set

$$B_{[0,k]}(\gamma,\epsilon)=\{\gamma' \mid \sup\{d(\gamma(x),\gamma'(x)) \mid x\in[0,k]\}<\epsilon\}.$$

The sets $B_{[0,k]}(\gamma,\epsilon)$ form a basis for the topology on the set of geodesic rays. Often in our proofs, we will work with representatives in the space of rays, rather than the equivalence classes themselves. We will denote the equivalence class of a ray γ by $[\gamma]$. Abusing notation, we will write $B_{[0,k]}([\gamma],\epsilon)$ for the image of $B_{[0,k]}(\gamma,\epsilon)$ in the quotient space.

Observation 4.1. *The sets $B_{[0,k]}([\gamma],\epsilon)$ form a basis for the topology on the visual boundary (the set of* equivalence classes *of rays).*

Definition 4.2. For $i\in\{0,1\}$ and $n\in\mathbb{N}$, we define C_n^i to be the set of equivalence classes of geodesic rays that "bottom out" in T_i after descending for exactly n edges. We define C_0^i to be the set of equivalence classes of rays that ascend forever in T_i without ever turning.

Notice that when equipped with the subspace topology, the sets C_n^i are homeomorphic to the Cantor set.

In terms of the lamp stand, elements of C_n^0 (for $n>0$) have a lit lamp at position $-n$, no lit lamps below that position, and the lamplighter at $+\infty$. Similarly, elements of C_n^1 (for $n>0$) have a lit lamp at position $n-1$, no lit lamps above that position, and the lamp lighter at $-\infty$. The lamp stand for an element of C_0^0 has the lamplighter at $+\infty$ and no lamps lit below 0. The lamp stand for an element of C_0^1 has the lamplighter at $-\infty$ and no lit lamps above -1.

Definition 4.3. For $k\in\mathbb{N}$, we define the set $C_{k,\infty}^i=\cup_{n=k}^{\infty}C_n^i$, which is the set of equivalence classes of geodesic rays that descend at least k edges in T_i before turning and ascending in T_i forever.

When equipped with the subspace topology, the sets $C_{k,\infty}^i$ are homeomorphic to the punctured Cantor set.

We can use these sets to better understand the topology on $\partial\mathrm{DL}_2(q)$.

Observation 4.4. *If $[\gamma]\in C_n^i$ for $n>0$ and $k\le n$, then*

$$B_{[0,k]}([\gamma],\epsilon)=C_{k,\infty}^i\cup C_0^{1-i}.$$

If $[\gamma] \in C^i_n$ *for* $n > 0$ *and* $k > n$*, then*

$$B_{[0,k]}([\gamma], \epsilon) = \{[\tau] \in C^i_n \mid \tau^{(i)} \text{ agrees with } \gamma^{(i)} \text{ on } [0, k]\}.$$

If $[\gamma] \in C^i_0$*, then*

$$B_{[0,k]}([\gamma], \epsilon) = C^{1-i}_{k,\infty} \cup \{[\tau] \in C^i_0 \mid \tau^{(i)} \text{ agrees with } \gamma^{(i)} \text{ on } [0, k]\}.$$

Proof. These statements are easily verified; recall that $0 < \epsilon < 1$. □

See Figure 6 for some examples of nested basis elements.

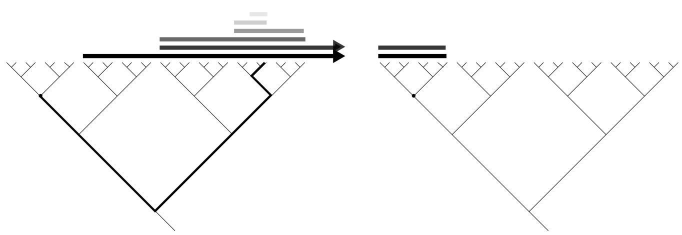

FIGURE 6. Some nested basis elements of $\partial\mathrm{DL}_2(2)$.

We now prove some of the important properties of these sets.

Lemma 4.5. *For* $n > 0$*, the set* C^i_n *is open in* $\partial DL_2(q)$*.*

Proof. Fix $n > 0$. For each $j \in \{1, 2, ..., q-1\}$, let $[\gamma_j] \in C^i_n$ such that if $j \neq j'$, then $\gamma^{(i)}_j(n+1) \neq \gamma^{(i)}_{j'}(n+1)$ (note that $\gamma^{(i)}_j|_{[0,n]} = \gamma^{(i)}_{j'}|_{[0,n]}$).
For $0 < \epsilon < 1$, notice that

$$C^i_n = \bigcup_{j=1}^{q-1} B_{[0,n+1]}([\gamma_j], \epsilon).$$

Thus, C^i_n is open. □

Lemma 4.5 does not apply when $n = 0$ because, in this case, the open sets $B_{[0,1]}([\gamma_j], \epsilon)$ include all elements of $C^{1-i}_{1,\infty}$ (i.e., every class that bottoms out in the opposite tree). Hence, C^i_0 cannot be formed as a union in the same way.

Lemma 4.6. *The set* C^i_0 *is not open.*

Proof. Fix $[\gamma] \in C^i_0$. For each $n > 0$, let γ_n be a ray that agrees with γ on the first n edges, but then bottoms out in T_{1-i} and ascends in T_{1-i} forever. In other words, $\gamma(x) = \gamma_n(x)$ for all $x \in [0, n]$ and $[\gamma_n] \in C^{1-i}_n$. Notice that $[\gamma_n] \notin C^i_0$.

Consider a basis element $B_{[0,k]}(\gamma,\epsilon)$ of the pre-quotient topology. If $n > k$, then $\gamma_n \in B_{[0,k]}(\gamma,\epsilon)$, and thus $[\gamma_n] \in B_{[0,k]}([\gamma],\epsilon)$. Thus, $[\gamma]$ is a limit point of $\{[\gamma_n]\}$, and so the complement of C_0^i is not closed. Hence, C_0^i is not open. □

Observation 4.7. *For any $k \in \mathbb{N}$, the set $C_0^i \cup C_{k,\infty}^{1-i}$ is open.*

Proof. This follows directly from Observation 4.4. □

Observation 4.8. *For $n \geq 0$, the set C_n^i is closed.*

Proof. The complement is open by the previous observations. □

4.2. **Separability.**

The boundary $\partial \mathrm{DL}_2(q)$ has some interesting separability properties that distinguish it from visual boundaries of hyperbolic or CAT(0) spaces.

Definition 4.9 ([9, §2.6])**.** A topological space X is T_1 if, for every pair of points $x, y \in X$, there exist open sets O_x and O_y such that $x \in O_x, y \notin O_x$ and $y \in O_y, x \notin O_y$.

This is a weaker form of separability than the Hausdorff condition (also known as T_2), which requires that the open sets O_x and O_y be disjoint.

Observation 4.10. *The visual boundary $\partial DL_2(q)$ is not Hausdorff.*

Proof. Let γ and γ' be distinct geodesic rays that ascend forever in T_i with no turns (i.e., $[\gamma], [\gamma'] \in C_0^i$). Notice that $[\gamma] \neq [\gamma']$. For each $n > 0$, let $[\gamma_n] \in C_n^{1-i}$ be as in the proof of Lemma 4.6; that is, γ_n agrees with γ on the first n edges before bottoming out in T_{1-i}. Notice that in the asymptotic equivalence class of γ_n, there is an element γ_n' that agrees with γ' on the first n edges before bottoming out in tree T_{1-i}.

Thus, $[\gamma]$ and $[\gamma']$ are distinct limit points of the sequence $\{[\gamma_n]\} = \{[\gamma_n']\}$, and so the topology is not Hausdorff. □

We could also prove that the topology is not Hausdorff using the following observation.

Observation 4.11. *Any open set containing an element of C_0^i necessarily contains $C_{k,\infty}^{1-i}$ for some k.*

Proof. This follows directly from Observation 4.4. □

Observation 4.12. *The visual boundary $\partial DL_2(q)$ is T_1.*

Proof. Let γ and γ' be geodesic rays that are not asymptotic to each other. So there exists some $k \in \mathbb{N}$ such that $\gamma(n) \neq \gamma'(n)$ for all $n \geq k$. Consider the basis elements $B_{[0,k]}([\gamma],\epsilon)$ and $B_{[0,k]}([\gamma'],\epsilon)$. For any ray $\tilde{\gamma} \in [\gamma]$, notice that $\tilde{\gamma}(k) \neq \gamma'(k)$, so $d(\gamma(k), \gamma'(k)) \geq 1 > \epsilon$, and $[\gamma] \notin B_{[0,k]}([\gamma'],\epsilon)$. By symmetry, the reverse holds as well. □

4.3. Compactness.

For X non-positively curved, ∂X is homeomorphic to the horofunction boundary of X and is also an inverse limit of compact sets [2, §II.8], both of which imply compactness. Since $\mathrm{DL}_2(q)$ is not CAT(0) or unique geodesic, we have to prove compactness directly.

Proposition 4.13. *$\partial DL_2(q)$ is compact.*

Proof. Let $\mathcal{A} = \{A_i\}_{i\in I}$ for some index set I be an open cover of $\partial\mathrm{DL}_2(q)$. Without loss of generality, we may assume that these open sets are basis elements.

As sets, $\partial T_0 \sqcup \partial T_1 = \partial\mathrm{DL}_2(q) \sqcup \{\omega_0, \omega_1\}$. We extend $\mathcal{A}$ to a cover $\bar{\mathcal{A}}$ of $\partial T_0 \sqcup \partial T_1$ by defining

$$\bar{A}_i = A_i \cup \{x \mid x = \omega_j \text{ for } j = 0, 1 \text{ and } C^j_{k,\infty} \subseteq A_i \text{ for } k > 0\}.$$

Since $\mathcal{A}$ covers $\partial\mathrm{DL}_2(q)$, Observation 4.11 ensures $\bar{\mathcal{A}}$ covers $\partial T_0 \sqcup \partial T_1$ (one can also see this from Observation 4.4 since our sets A_i are assumed to be basis elements).

We now define covers $\bar{\mathcal{A}}^0$ and $\bar{\mathcal{A}}^1$ of ∂T_0 and ∂T_1, respectively, by $\bar{A}^0_i = \bar{A}_i \cap \partial T_0$ and $\bar{A}^1_i = \bar{A}_i \cap \partial T_1$. Since the A_i are basis elements, it is easy to verify from Observation 4.4 that these projections $\bar{A}^0_i$ and $\bar{A}^1_i$ are open sets in the boundaries of the trees.

Thus, the covers $\bar{\mathcal{A}}^0$ and $\bar{\mathcal{A}}^1$ are open. Since ∂T_0 and ∂T_1 are compact, there exists a finite $F \subset I$ such that $\{\bar{A}^0_i\}_{i\in F}$ covers ∂T_0 and $\{\bar{A}^1_i\}_{i\in F}$ covers ∂T_1. Then $\{A_i\}_{i\in F}$ is a finite subcover of $\partial\mathrm{DL}_2(q)$. □

4.4. Connectedness.

We have been considering the visual boundary through the Cantor sets C^i_n and punctured Cantor sets $C^i_{k,\infty}$, so it is reasonable to expect that the visual boundary is disconnected in a similar manner to a Cantor set.

Observation 4.14. *$\partial DL_2(q)$ is totally disconnected.*

Proof. Let S be a subset of $\partial\mathrm{DL}_2(q)$ containing at least two elements.

Suppose that $S \cap C^i_n \neq \varnothing$ for some $n > 0$ and some $i \in \{0, 1\}$. If $S \subseteq C^i_n$, then S is disconnected since C^i_n is a Cantor set. Else, since C^i_n is both open and closed, it and its complement form a separation of S.

If $S \cap C^i_n = \varnothing$ for all $n > 0$ and $i \in \{0, 1\}$, then $S \subseteq \left(C^0_0 \cup C^1_0\right)$. So S is a subset of a Cantor set and thus is disconnected. □

4.5. Intuitive picture of the topology of $\partial\mathbf{DL}_2(q)$.

Intuitively, the visual boundary $\partial\mathrm{DL}_2(q)$ can be viewed as a pair of punctured Cantor sets in which the punctures are "filled" by a portion of the other Cantor set. Specifically, every open neighborhood of ω_i becomes an open neighborhood of C_0^{1-i}. Figure 7 illustrates this notion. Clearly, $\partial\mathrm{DL}_2(q)$ is not homogeneous.

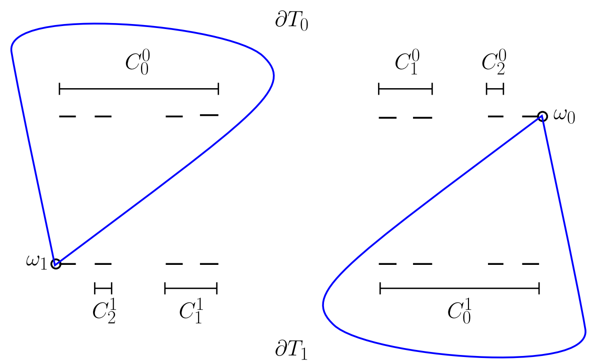

Figure 7. An informal visualization of $\partial\mathrm{DL}_2(2)$

4.6. Dynamics of the action of L_2 on $\partial\mathbf{DL}_2(2)$.

Notice that the exponent sum of t in a word representing an element g of L_2 is equal to the position of the lamplighter in the lamp stand representation of g. Thus, the exponent sum is an invariant of the group element. In §2.3, we defined the function $exp_t(g)$ to denote the exponent sum of t for g.

Notice that for an element $g \in L_2$, if $exp_t(g) = 0$, then g^2 is trivial (since the second application of g will switch off all the lights that the first application of g switched on). Notice also that if $exp_t(g) \neq 0$, then g will have infinite order since the lamplighter for g^n with $n \neq 0$ is at position $n \cdot exp_t(g) \neq 0$. In other words, for $g \in L_2$ non-trivial, the order of g is either 2 (when $exp_t(g) = 0$) or infinite (when $exp_t(g) \neq 0$).

Definition 4.15. Let $g \in L_2$ with $exp_t(g) > 0$. Its lamp stand has no lit lamps below some position m. Consider the lamp stand for g^n for $n \in \mathbb{N}$. The lamplighter for g^n is at position $n \cdot exp_t(g)$ and no matter how many more times we multiply by g, the lamps below position $n \cdot exp_t(g) + m$ will not be switched again. Thus, since $n \cdot exp_t(g) + m \to \infty$ as $n \to \infty$, we have a well-defined lamp stand for g^∞. This lamp stand can be realized by a geodesic ray in $\mathrm{DL}_2(2)$ (since it is the Cayley graph of L_2) by starting with $t^m(at)$ and then multiplying by t or at for each successive lamp, depending on whether the lamp is lit or unlit in g^∞. Thus, g^∞ is an element of $\partial\mathrm{DL}_2(2)$.

We can similarly define g^∞ for g with $exp_t(g) < 0$ (except that it will have no lit lamps *above* m).

Intuitively, g^∞ is the "lamp stand limit" of g^n. For example, t^∞ is the lamp stand with no lit lamps and the lamplighter at $+\infty$.

Definition 4.16. For $g \in L_2$ with $exp_t(g) \neq 0$, we define $g^{-\infty}$ to be $(g^{-1})^\infty$.

Theorem 4.17. *If a non-trivial element g of L_2 has $exp_t(g) = 0$, then its action on $\partial DL_2(2)$ will be periodic of order 2. Otherwise, g will act with north-south dynamics on the boundary, with the attractor in $(\partial T_0 - \omega_0) \times \omega_1$ and the repeller in $(\partial T_1 - \omega_1) \times \omega_0$ if $exp_t(g) > 0$ and vice versa if $exp_t(g) < 0$.*

Proof. If $exp_t(g) = 0$, then the action of g on an element of $\partial\mathrm{DL}_2(q)$ will simply be to switch a finite set of lamps (the ones that are lit in the lamp stand interpretation of g).

For $g \in L_2$ with $exp_t(g) > 0$ and $[\gamma] \in \partial\mathrm{DL}_2(q)$, define $[\gamma_n] \in \partial\mathrm{DL}_2(q)$ to be $g^n \cdot [\gamma]$. Assume $[\gamma]$ (and thus $[\gamma_n]$ for all n) has the lamplighter at $+\infty$ (i.e., $[\gamma] \in (\partial T_0 - \omega_0) \times \omega_1$). Thus, there is a minimum lit lamp, say at position m, in the lamp stand for $[\gamma]$. Then, in the representation for $g^n \cdot [\gamma]$, all lamps at positions below $n \cdot exp_t(g) + m$ will be lit or unlit according to g^n's lamp stand. For any $[0, k] \subseteq [0, \infty)$ compact and any $0 < \epsilon < 1$, let n be large enough so that $n \cdot exp_t(g) + m > k$. Then notice that $[\gamma_n] \in B_{[0,k]}(g^\infty, \epsilon)$. Thus, $[\gamma_n] \to g^\infty$.

Similar arguments show the rest of the result. $\square$

Corollary 4.18. *The action on $\partial DL_2(2)$ of a non-torsion element of L_2 is hyperbolic.*

5. $\partial\mathbf{DL}_d(q)$ for $d > 2$

5.1. Geodesics in $\mathbf{DL}_d(q)$.

Label each edge of each tree T_i by an $\alpha \in \{0, \ldots, q-1\}$, so that, for each vertex $v \in T_i$, the edges moving up from v correspond to $\{0, \ldots, q-1\}$. Modifying a concept from [11], say an edge of $\mathrm{DL}_d(q)$ has *type* $(i(\alpha)-j)$, $0 \le i \neq j < d$, $0 \le \alpha < q$, if it ascends in T_i along an edge labeled α and descends in T_j, or $(i(\alpha) - j(\beta))$ if we wish to keep track of the descending label as well. If i, j, k, l are pairwise distinct, any edge of type $(i(\alpha) - j)$ "commutes" with any edge of type $(k(\beta) - l)$, in the sense that, given an initial vertex in $\mathrm{DL}_d(q)$, the two (uniquely determined) paths of type

$(i(\alpha)-j)(k(\beta)-l)$ and $(k(\beta)-l)(i(\alpha)-j)$ have the same terminal vertex. Moreover, $(i(\alpha)-j)$ commutes with $(k(\beta)-j)$. In addition, an adjacent pair of edges having type $(i(\alpha)-j)(k(\beta)-i)$ can be replaced by the single edge of type $(k(\beta)-j)$, since this pair creates an unnecessary backtrack in T_i. Finally, $(j(\alpha)-i(\beta))(i(\beta')-k)$ can be replaced with $(j(\alpha)-k)$ if and only if $\beta=\beta'$, as that is the only case with backtracking.

This notation gives us a way to define turns in the $d>2$ case.

Definition 5.1. A *turn* in T_i in a path in $\mathrm{DL}_d(q)$ is a subpath that begins with an edge of type $(j(\alpha)-i)$ for some j and α, ends with an edge of type $(i(\beta)-k)$ for some k and β, and no other edge type in the subpath involves i.

In the case $d=2$, this is equivalent to our definition in §3.2.

We now show that any path whose projection to T_i turns back up the same edge is not geodesic.

Lemma 5.2. *Let p be a path in $DL_d(q)$ following edges $e_1,e_2,\dots,e_n$ in order. Suppose for some $0\le t<s\le n$, e_t is of type $(j(\alpha)-i(\beta))$, and e_s is of type $(i(\beta)-k)$, and all edges between e_t and e_s do not involve T_i. Then p is not geodesic.*

Proof. We have a sequence of edge types

$$(j(\alpha)-i(\beta))(a_0(\delta_0)-b_0)\dots(a_{s-t-1}(\delta_{s-t-1})-b_{s-t-1})(i(\beta)-k),$$

a_x, b_x, i pairwise distinct, $0\le x<s-t$. Using the commuting relations discussed above, we may replace this subsequence with either

$$(j(\alpha)-i(\beta))(i(\beta)-k)(a_0(\delta_0)-b_0)\dots(a(\delta_{s-t-1})_{s-t-1}-b_{s-t-1})$$

if there is no x with $a_x=k$, or if such an x exists, by

$$(j(\alpha)-i(\beta))\dots(a_y(\delta_y)-b_y)(i(\beta)-k)\dots(a_{s-t-1}(\delta_{s-t-1})-b_{s-t-1})$$

(where $y=\max\{x\mid a_x=k\}$). In either case, again by the preceding discussion, we may replace a two edge subsequence with a single edge without affecting the endpoints of the subsequence. Hence, a shorter path is found and p is not geodesic. □

The following observation is trivial, but will be key in the proof of Theorem 5.4.

Observation 5.3. *Let π be a path in $DL_d(q)$ containing a subpath ρ of length l. Consider the family P_ρ of paths of length l that begin at $\rho(0)$ and end at $\rho(l)$. For any $\rho'\in P_\rho$, the path π' constructed from π by replacing ρ with ρ' has the same initial and terminal vertices and the same length as π.*

Theorem 5.4. *A geodesic in $DL_d(q)$ has no more than one turn in each tree.*

Proof. Suppose a path π of length n has more than one turn in a tree T_i. Let $v_1, v_2 \in T_i$ be the vertices where the first two turns in T_i bottom out, in order. If $h_i(v_1) \geq h_i(v_2)$, then π descends to v_1, turns, and must descend through v_1 again to reach v_2. If $h_i(v_2) \geq h_i(v_1)$, then π must ascend to v_2 (having passed through v_1) and, in fact, ascend above v_2, before it can turn at v_2. Either way, there are $k, l \in \mathbb{Z}$ satisfying $0 < k < k+1 < l < n$, and $z \in \{1, 2\}$ such that $\pi^{(i)}(k) = v_z = \pi^{(i)}(l)$. Moreover, by assumption, the subpath $\rho = \pi(k), \pi(k+1), \dots, \pi(l)$ has no turns in T_i and has at least one ascent in T_i followed by at least one descent. So, applying Observation 5.3, ρ can be replaced with a subpath ρ', making the resultant path π' satisfy Lemma 5.2. (To see this, when $h_i(v_1) \geq h_i(v_2)$, we can choose ρ' so that $\pi^{(i)}(k-1) = \rho'^{(i)}(k+1) = \rho'^{(i)}(l-1)$, and if $h_i(v_2) \geq h_i(v_1)$, then we can choose ρ' so that $\rho'^{(i)}(k+1) = \rho'^{(i)}(l-1) = \pi^{(i)}(l+1)$.) Hence, π' is not geodesic, and since π and π' have the same initial and terminal vertices and the same length, neither is π. □

Corollary 5.5. *For any geodesic ray γ in $DL_d(q)$ and any $0 \leq i < d$, $\gamma^{(i)}$ approaches at most one end point of T_i, and if $\gamma^{(i)}$ consists of finitely many edges, then $\gamma^{(i)}$ is eventually constant in T_i.* □

5.2. Asymptotic geodesic rays in $\mathrm{DL}_d(q)$.

We will show that $\partial \mathrm{DL}_d(q)$, $d > 2$, has the indiscrete topology.

Definition 5.6. We say that two geodesic rays *have the same ends* if whenever one of them has a projection to a tree that has infinitely many edges, so does the other, and the two projections go to the same end of that tree.

The visual boundary of $\mathrm{DL}_d(q)$, $d > 2$, will be significantly larger than that of $\mathrm{DL}_2(q)$, as sets, not just because additional punctured Cantor sets will be added for the trees, but also because it is no longer guaranteed that two geodesic rays having the same ends will be asymptotic due to the additional degree of freedom offered by a third tree. However, since we aim to show that when $d > 2$ the boundary has the indiscrete topology, we will not delve into this. We will show that any point of $\partial \mathrm{DL}_d(q)$ is topologically indistinguishable[2] from a point that approaches a distinguished end ω_i in some T_i and a non-distinguished end e_j, $j \neq i$, in some T_j, and is trivial in every other tree. Thus, the above issue can be avoided in proving that $\partial \mathrm{DL}_d(q)$ has the indiscrete topology.

[2]Two points are *topologically indistinguishable* from each other if every open set that contains one of these points contains the other as well.

Observation 5.7. *If τ_n are geodesic rays in $DL_d(q)$ that are asymptotic to another geodesic ray γ for all $n \in \mathbb{N}$, and τ is a geodesic ray that is a limit point of $\{\tau_n\}$ in the compact-open topology on geodesic rays before we quotient by the asymptotic equivalence classes, then $[\gamma]$ and $[\tau]$ are topologically indistinguishable elements of $\partial DL_d(q)$.*

Proof. Clearly, every neighborhood of $[\tau]$ contains $[\gamma]$ and, since the basis definition of the topology is symmetric, every neighborhood of $[\gamma]$ contains $[\tau]$. □

Lemma 5.8. *Let γ be a geodesic ray in $DL_d(q)$ for $d > 2$. Partition the set $\{T_0, T_1, T_2, \ldots, T_{d-1}\}$ into sets $\mathcal{I}$ and $\mathcal{F}$, where the projection of γ to any tree in $\mathcal{F}$ is eventually constant and the height of the projection to any tree in $\mathcal{I}$ approaches $\pm\infty$. Then we can construct a geodesic ray τ that is asymptotic to γ and such that the projection of τ to any tree in $\mathcal{F}$ is trivial. (Here trivial means the image is constant at the origin.)*

Proof. Let M be large enough that for each $T_i \in \mathcal{F}$, all edges of γ that project onto T_i come before M, and for each tree in $\mathcal{I}$ in which γ bottoms out, γ does so before M.

Let ρ be a geodesic ray such that for each tree T_i in $\mathcal{I}$, the projection $\rho^{(i)}$ approaches the same end as the projection $\gamma^{(i)}$, chosen so that the projection of ρ to any tree in $\mathcal{F}$ is trivial, and all of the turns in ρ come before $N \geq M$.

Let τ be defined by $\tau|_{[0,N]} = \rho|_{[0,N]}$, and for $n > N$, the nth edge of τ simply "tracks" the nth edge of γ. That is, when the nth edge of γ moves upward in some T_{i_n} and downward in some T_{j_n}, τ does the same, choosing the upward branch that takes it toward the same point of ∂T_{i_n} that $\gamma^{(i_n)}$ approaches. Since $\rho|_{[0,N]}$ is a geodesic and all turns in ρ occur before N, τ is a geodesic ray.

The ray τ has been chosen so that for each tree T_i, for $n > N$, $d_{T_i}(\tau^{(i)}(n), \gamma^{(i)}(n))$ is constant. Lemma 3.2 then ensures τ and γ are asymptotic. □

Lemma 5.9. *For $d > 2$, let γ be a geodesic ray in $DL_d(q)$ with empty projection to T_i. Let τ be another geodesic ray whose projections in trees other than T_i have the same ends as γ and whose projection to T_i is infinite. Then, $[\gamma]$ and $[\tau]$ are topologically indistinguishable.*

Proof. Let N be large enough so that all turns and finite projections of γ and τ come before N. For $n > N$, define τ_n to be the ray that matches τ up through n edges in T_i and then "tracks" γ by going up and down in the same trees for each edge as in the proof of Lemma 5.8. By the same argument as in the proof of Lemma 5.8, each τ_n is asymptotic to γ. But clearly $\tau_n \to \tau$, so by Observation 5.7 we are done. □

Corollary 5.10. *For $d > 2$, an element of $\partial DL_d(q)$ is topologically indistinguishable from at least one other element that only has non-empty projections in two trees: one that eventually ascends in height without bound and the other which eventually descends in height without bound.*

Proof. This follows immediately from Lemma 5.8 and Lemma 5.9. □

Lemma 5.11. *Suppose that γ is a geodesic ray in $DL_d(q)$ for $d > 2$ that has no projection to T_i and infinite projection to T_j. Then we can construct a geodesic ray τ such that τ has no projection to T_j, $[\tau]$ is in every open set that contains $[\gamma]$, and γ's and τ's edges are exactly the same except that whenever γ has an edge that would project to T_j, τ projects to the same exact edge in T_i.*

In other words, γ and τ are the same ray, just swapping the projections in T_i and T_j (one of which is empty), and the asymptotic equivalence classes of γ and τ are topologically indistinguishable in $\partial DL_d(q)$.

Proof. We begin by assuming that $\gamma^{(j)}$ eventually increases without bound in height. The descending case is analogous. Let ℓ be the number of edges that $\gamma^{(j)}$ descends before increasing forever.

In the obvious way, we can construct a geodesic ray τ that exactly matches γ, except that the infinite projection to T_j and the empty projection to T_i are swapped. We will now construct a sequence of geodesic rays τ_n such that $\tau_n \in [\gamma]$ and $\tau_n \to \tau$, which will show topological indistinguishability.

Let N be large enough so that every turn and finite projection of γ (and thus of τ also) occurs before N. For $n > N$, we construct τ_n as follows.

The first n edges of τ_n exactly match the first n edges of τ. By choice of N, we have partitioned the trees into "up," "down," and "empty" for projections of τ. That is, any edge after N has its up projection in one of the "up" trees and its down projection in one of the "down" trees (after N there are no edge projections in the "empty" trees). Notice that T_j is "empty" for τ, but is "up" for γ.

So for the next ℓ edges of τ_n, continue to copy τ, except that the "down" projections of the edges should all be in T_j instead. By the definition of ℓ, these down edges will exactly reach the point where $\gamma^{(j)}$ turns.

For all subsequent edges, τ_n "mimics" γ by going up and down in the same trees as γ. As a result, for $x \geq n + \ell$, the distance between $\tau_n(x)$ and $\gamma(x)$ will be equal to the distance between $\tau_n(n+\ell)$ and $\gamma(n+\ell)$, so the two rays are in the same asymptotic equivalence class. Since $\tau_n \to \tau$, by Observation 5.7, we are done. □

Corollary 5.12. *For $d > 2$, an element of $\partial DL_d(q)$ is topologically indistinguishable from at least one other element that only has non-empty projections in trees T_0 and T_1 such that the projection to T_0 eventually ascends in height without bound, and the projection to T_1 eventually descends in height without bound.*

Theorem 5.13. *For $d > 2$, $\partial DL_d(q)$ has the indiscrete topology.*

Proof. Let γ and γ' be geodesic rays in $\mathrm{DL}_d(q)$. By Corollary 5.12, we may assume that the only non-empty projections of γ and γ' are in trees T_0 and T_1, that both $\gamma^{(0)}$ and $\gamma'^{(0)}$ eventually ascend in height without bound, and that both $\gamma^{(1)}$ and $\gamma'^{(1)}$ eventually descend in height without bound. Let ℓ (and ℓ') be the number of down edges in $\gamma^{(0)}$ ($\gamma'^{(0)}$, respectively) before turning.

For N sufficiently large so that all the turns in γ and γ' occur after N and for $n > N$, define τ_n so that the first n edges go up always taking the leftmost edge in T_2 (recall $d > 2$) and down in T_1. For the next ℓ edges, τ_n goes down in T_0 and up in T_2 (again, always taking the leftmost edge). After that, τ_n goes up in T_0 and down in T_1 towards the ends of γ. We define τ'_n similarly, but using ℓ' and γ'. Notice that $\tau_n|_{[0,n]} = \tau'_n|_{[0,n]}$.

The ray τ_n is asymptotic to γ, since the two rays are never farther apart than their distance at $\tau_n(n+\ell)$ and $\gamma(n+\ell)$. Similarly, τ'_n is asymptotic to γ'.

But notice that for any $[0,k] \subseteq [0,\infty)$ and any $0 < \epsilon < 1$, we have $\tau_n \in B_{[0,k]}(\tau'_n, \epsilon)$ for any $n \geq l$. Thus, $[\gamma]$ and $[\gamma']$ are topologically indistinguishable. □

Acknowledgments. The authors thank Moon Duchin and Melanie Stein for their helpful conversations. The authors appreciate the helpful comments from the anonymous referee, particularly those which improved the approach in §4. Both authors also thank the Institute for Advanced Study for its hospitality at the Park City Math Institute Summer Session 2012.

References

[1] Laurent Bartholdi, Markus Neuhauser, and Wolfgang Woess, *Horocyclic products of trees*, J. Eur. Math. Soc. (JEMS) **10** (2008), no. 3, 771–816.

[2] Martin R. Bridson, and André Haefliger, *Metric Spaces of Non-Positive Curvature*. Grundlehren der Mathematischen Wissenschaften, 319. Berlin: Springer-Verlag, 1999.

[3] Sean Cleary and Tim R. Riley, *Erratum to: "A finitely presented group with unbounded dead-end depth"* [Proc. Amer. Math. Soc. **134** (2006), no. 2, 343–349; MR2176000], Proc. Amer. Math. Soc. **136** (2008), no. 7, 2641–2645.

[4] Christopher B. Croke and Bruce Kleiner, *Spaces with nonpositive curvature and their ideal boundaries*, Topology **39** (2000), no. 3, 549–556.

[5] Moon Duchin, Samuel Lelièvre, and Christopher Mooney, *Statistical hyperbolicity in groups*, Algebr. Geom. Topol. **12** (2012), no. 1, 1–18.

[6] Alex Eskin, David Fisher, and Kevin Whyte, *Coarse differentiation of quasi-isometries I: Spaces not quasi-isometric to Cayley graphs*, Ann. of Math. (2) **176** (2012), no. 1, 221–260.

[7] Kyle Kitzmiller and Matt Rathbun, *The visual boundary of* $\mathbb{Z}^2$, Involve **4** (2011), no. 2, 103–116.

[8] Ilya Kapovich and Nadia Benakli, *Boundaries of hyperbolic groups* in Combinatorial and Geometric Group Theory. Ed. Sean Cleary, Robert Gilman, Alexei G. Myasnikov, and Vladimir Shpilrain. Contemporary Mathematics, 296. Providence, RI: Amer. Math. Soc., 2002. 39–93.

[9] James R. Munkres, *Topology: A First Course*. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1975.

[10] Walter Parry, *Growth series of some wreath products*, Trans. Amer. Math. Soc. **331** (1992), no. 2, 751–759.

[11] Melanie Stein and Jennifer Taback, *Metric properties of Diestel-Leader groups*, Michigan Math. J. **62** (2013), no. 2, 365–386.

[12] Corran Webster and Adam Winchester, *Busemann points of infinite graphs*, Trans. Amer. Math. Soc. **358** (2006), no. 9, 4209–4224 (electronic).

[13] Wolfgang Woess, *Lamplighters, Diestel-Leader graphs, random walks, and harmonic functions*, Combin. Probab. Comput. **14** (2005), no. 3, 415–433.

(Jones) Department of Mathematics, Computer Science & Statistics; State University of New York at Oneonta; 108 Ravine Parkway; Oneonta, NY 13820

E-mail address: keith.jones@oneonta.edu

(Kelsey) Department of Mathematics; Trinity College; 300 Summit St.; Hartford, CT 06106

E-mail address: gregory.kelsey@trincoll.edu