The Lamplighter Groups

Anthony Saint-Criq

January 8, 2021

Abstract

We shall first give an informal desciption of the N-th lamplighter group L_N to allow us to construct it algebraically. Identifying the underlying operation on the groups $\mathbb{Z}/N$ and $\mathbb{Z}$, this gives several properties around the presentation of L_N. Next, we focus on the Cayley graph $\Gamma(L_N)$ of this group, which turns out to be a special object on its own : a Diestel-Leader graph. We will present the formal construction of these graphs, and we identify $\Gamma(L_N)$ with $\mathrm{DL}(N)$. Word length will also be of some interest.

1. Decriptions of the groups

The idea behing these groups is the following : consider an infinite street, with lamp posts every so often, and consider a lamplighter, whose role is to go across the street, lighting up or turning off some of the lamp posts, and then ending his journey at at the foot of a light. We may picturally give an example in figure 1, where darkened discs correspond to lamps turned on, where the vertical line represents the origin (assume the street is the real line and a lamp is set at every integer), where the arrow represents the ending position of the lamplighter, and where circles not drawn are lamps that are off.

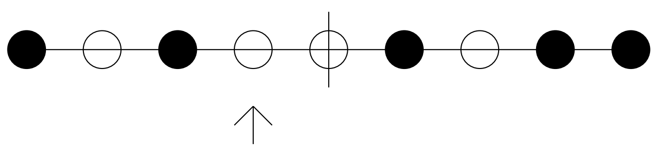

Figure 1. The lamplighter has lit the lamps at positions -4, -2, 1, 3 and 4, and ended his journey at position -1.

Note that the lamplighter may turn on a lamp, perform some actions elsewhere, to then come back and turn back off the lamp. This fact will be important when trying to give a presentation for the lamplighter group.

If this gave an informal description of the elements of the group, we may as well give an as-informal description of the group law. For this, it may be a nice analogy to describe an element as a set of instructions, just like a Turing machine does (although we do not have conditional statements in out case). For the example of figure 1, the group element may be described as the following set of instructions, assuming the starting position of the lamplighter is the origin :

1. Go to the right once.
2. Turn the lamp on.
3. Go to the right twice.
4. Turn the lamp on.
5. Go the right once.
6. Turn the lamp on.
7. Go to the left six times.
8. Turn the lamp on.
9. Go the left twice.
10. Turn the lamp on.
11. Go to the right thrice.

Now, to compose two elements, simply apply the corresponding two sets of instructions one after the order. This corresponds to stacking the diagrams of the two elements, but with the origin of the second element shifted to be aligned with the ending position of the first. Consider the following other element :

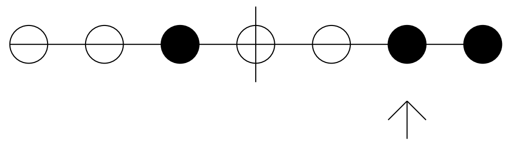

FIGURE 2. The lamplighter has lit the lamps at positions -1, 2 and 3, and ended his journey at position 2.

Stacking this diagram after the one from figure 1, we obtain the following :

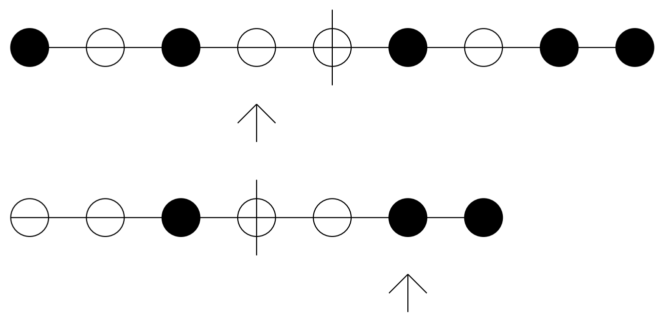

FIGURE 3. How to compute the composition of two elements in the lamplighter group

At last, one checks that concatenation of the two sets of instruction boils down to :

1. Turning on or off the lamps accordingly to the XOR rule (component-wise).
2. The origin of the product is the origin of the first element.
3. The ending position is the ending position of the second element.

For out previous example, it yields :

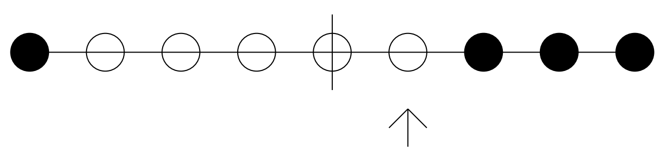

FIGURE 4. The result of the previous composition

Now, to give a more algebraic description of this group, notice that an element is characterized by two things :

1. An integer $k \in \mathbb{Z}$, corresponding to the position of the lamplighter.
2. A finitely-supported sequence $(a_n)_{n\in\mathbb{Z}}$, where $a_n \in \mathbb{Z}/2$ represents the state of the lamp at position n.

Therefore, the underlying set for the lamplighter group L_2 is :

$$L_2 = \mathbb{Z} \times \bigoplus_{n\in\mathbb{Z}} \mathbb{Z}/2.$$

Already, this 2 subscript suggests we shall not restrein ourselves to lamps having only two states. We may therefore define (with $N \geqslant 2$, because $L_1 \cong \mathbb{Z}$) :

$$L_N = \mathbb{Z} \times \bigoplus_{n \in \mathbb{Z}} \mathbb{Z}/N,$$

where in this case, lamps have N possible states. Again, this can be compared to Turing machines, where the construction is similar.
We now need to define composition of elements. Positions are just added together, and the sequences are shifted and then composed component-wise. That is, given $(k, \boldsymbol{a}), (\ell, \boldsymbol{b}) \in L_N$, we may define :

$$(k, \boldsymbol{a}) \star (\ell, \boldsymbol{b}) = (k + \ell, (a_n + b_{n-k})_{n \in \mathbb{Z}})$$

Proposition 1.1. *$(L_N, \star)$ is a group.*

Proof. The neutral element is $(0, \mathbf{0})$. For associativity, we have :

$$\begin{aligned}
[(k, \boldsymbol{a}) \star (\ell, \boldsymbol{b})] \star (m, \boldsymbol{c}) &= (k + \ell, (a_n + b_{n-k})_{n \in \mathbb{Z}}) \star (m, \boldsymbol{c}) \\
&= (k + \ell + m, (a_n + b_{n-k} + c_{n-k-\ell})_{n \in \mathbb{Z}}) \\
&= (k, \boldsymbol{a}) \star (\ell + m, (b_n + c_{n-\ell})_{n \in \mathbb{Z}}) \\
&= (k, \boldsymbol{a}) \star [(\ell, \boldsymbol{b}) \star (m, \boldsymbol{c})].
\end{aligned}$$

At last, one checks that an inverse for $(k, \boldsymbol{a})$ is $(-k, (-a_{n+k})_{n \in \mathbb{Z}})$. ∎

Picturally, in L_2, the inverse of the element of figure 1 would be :

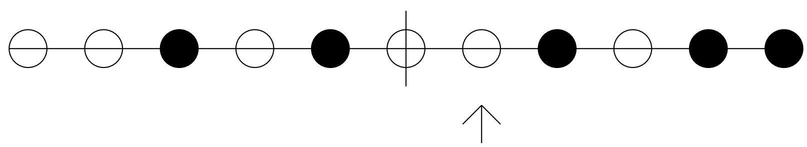

FIGURE 5. The inverse of the element in figure 1. Try stacking the diagrams together to visualize it.

Now, to give a presentation of L_N, we may first derive a generating set for the group, and compute some of the relations. Using the algortihmic point of vue we adopted to define the group law, we can see that two actions are sufficient to generate any element :

1. Moving the lamplighter once to the right.
2. Switching the current lamp to its next state.

Algebraically, those are the following two elements :

$$T = (1, \mathbf{0}) \qquad \text{and} \qquad A = (0, \delta_0),$$

where δ_0 is the sequence $(..., 0, 0, 1, 0, 0, ...)$, with the one being at position zero (recall that 1 is a generator for $\mathbb{Z}/N$). We therefore have :

Proposition 1.2. *$\{T, A\}$ is a generating set for L_N, that is $L_N = \langle T, A \rangle$.*

Proof. Choose an element $(k, \boldsymbol{a}) \in L_N$. We shall proceed by induction over the number r of non-zero elements of the sequence $\boldsymbol{a}$.
If $r = 0$, then $(k, \boldsymbol{a}) = T^{\star k}$.
Suppose that $(k, \boldsymbol{a})$ is an element of $\langle T, A \rangle$ whenever $r = n$, and assume that $r = n + 1$. Let $m \in \mathbb{Z}$ be such that $a_m \neq 0$, and let $b = N - a_m$. Then, define $g = A^{\star b} \star T^{\star(m+k)} = (m, b\delta_{m+k})$, where $b\delta_{m+k}$ is the sequence whose only non-zero term is b at position $m + k$. The hard work is done, since :

$$(k, \boldsymbol{a}) \star g = (k + m, \boldsymbol{a} + b\delta_m),$$

with $(\boldsymbol{a} - k\delta_m)_i = a_i$ if $i \neq m$ and $(\boldsymbol{a} - b\delta_m)_m = 0$. Therefore, $(k, \boldsymbol{a}) \star g$ is an element whom we can apply the induction to, that is : $(k, \boldsymbol{a}) \star g \in \langle T, A \rangle$, and thus, by inverting, since $g^{-1} \in \langle T, A \rangle$, we obtain $(k, \boldsymbol{a}) \in \langle T, A \rangle$. ∎

Heuristically, in the proof, we turned off a lamp to reduce the number of lamps that are not off, to then apply the induction to this new element. The idea behind is exactly this algorithmic interpretation for L_N. However, another description is possible for L_N, that is not involving *sequences* but rather *polynomials*. The core idea is just the same, it's just a rephrasing of the previous description :

Proposition 1.3. *Considering the ring $\mathcal{A}_N = (\mathbb{Z}/N)[X, X^{-1}]$ of polynomials in the formal variables X and X^{-1} whose coefficients are in $\mathbb{Z}/N$, we can identify L_N as a subgroup of $\mathrm{GL}(2, \mathcal{A}_n)$ by :*

$$L_N \cong G_N := \left\{ \begin{pmatrix} X^k & P \\ 0 & 1 \end{pmatrix},\ k \in \mathbb{Z},\ P \in \mathcal{A}_n \right\}.$$

As a remark, recall that $\mathcal{A}_N$ is defined as a quotient ring :

$$\mathcal{A}_N = (\mathbb{Z}/N)[U, V]/\langle UV \rangle,$$

where we identify X to be the class of U and X^{-1} the class of V.

Proof. The isomorphism is explicit :

$$\begin{array}{cccc} \Phi & : & L_N & \to & G_N \\ & & (k, \boldsymbol{a}) & \mapsto & \begin{pmatrix} X^k & \displaystyle\sum_{n\in\mathbb{Z}} a_n X^n \\ 0 & 1 \end{pmatrix}. \end{array}$$

Φ is indeed well-defined, since by hypothesis over $\boldsymbol{a}$, the sum is finitely-supported. Moreover, if

$$P = \sum_{n\in\mathbb{Z}} a_n X^n \qquad \text{and} \qquad Q = \sum_{n\in\mathbb{Z}} b_n X^n,$$

then we have :

$$X^k Q + P = \sum_{n\in\mathbb{Z}} a_n X^n + \sum_{n\in\mathbb{Z}} b_n X^{n+k} = \sum_{n\in\mathbb{Z}} (a_n + b_{n-k}) X^n,$$

and thus, by computing the matrix product, we get that Φ is a morphism. Injectivity is quite evident : if $\Phi(k, \boldsymbol{a}) = \Phi(\ell, \boldsymbol{b})$, then by comparing the first two coefficients of the matrix, we get $k = \ell$ for the first, and $\boldsymbol{a} = \boldsymbol{b}$ for the second, by comparing the coefficients of the two polynomials this time. For surjectivity, it is easy to construct a suitable element. ∎

Now, we shall make use of proposition 1.2 and try to derive relations satisfied by those generators, to then give a presentation of L_N. Evidently, we already see the relation $A^N = (0, \mathbf{0})$. There will be two kinds of relations, with this first being the only one of the sort. The following lemma gives the other relations :

Lemma 1.4. *For all $i, j \in \mathbb{Z}$, we have that $T^i A T^{-i}$ and $T^j A T^{-j}$ commute, that is :*

$$[T^i A T^{-i}, T^j A T^{-j}] = (0, \mathbf{0}).$$

Once again, the idea behind these relations is to see it with the algorithmic point of vue : go switching the state of a lamp, coming back, go switching the state of another lamp and coming back again is the same as doing the same actions but with the opposite order for these lamps. Note that, having the presentation in mind, we do not need to consider powers of A, since we can always take $i = j$.

Proof. We shall in fact prove that these relations hold in G_N, by making use of proposition 1.3. We have :

$$\Phi(T^k) = \begin{pmatrix} X^k & 0 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad \Phi(A) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix},$$

from which we obtain :

$$\Phi(T^k A T^{-k}) = \begin{pmatrix} 1 & X^k \\ 0 & 1 \end{pmatrix}.$$

By computing that

$$\begin{pmatrix} 1 & X^i \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 1 & X^j \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & X^i + X^j \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & X^j \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 1 & X^i \\ 0 & 1 \end{pmatrix},$$

we obtain the result. ■

Finally, this allows us to give a presentation for L_N :

Theorem 1.5. *We have* $L_N \cong \langle t, a | a^N, [t^i a t^{-i}, t^j a t^{-j}],\ i, j \in \mathbb{Z}\rangle$.

However, a direct proof considering the morphism $\Psi : F_2 \to L_N$ is not so easy in that setting, since we are dealing with an *infinite* group. We will therefore use a general construction and the presentation of semi-direct products, that is :

Theorem 1.6. *Let* $G = \langle X|R\rangle$ *and* $H = \langle Y|S\rangle$ *be two groups given by presentation, and let* $\phi : H \to \mathrm{Aut}(G)$*. We have the following presentation for their semi-direct product :*

$$G \rtimes_\phi H \cong \langle X, Y | R, S, yxy^{-1} = \phi(y)(x),\ (x, y) \in X \times Y\rangle.$$

For a proof, see [4].

This is useful, because L_N is a particular semi-direct product, that is a *wreath product* :

Definition 1.7. *Let* G *and* H *be two groups. Their (regular, restricted)* wreath product *is the group* $G \wr H$ *defined as follows : take*

$$K = \bigoplus_{\omega \in H} G,$$

and define an action of H *on* K *by* $h \bullet (g_\omega)_{\omega \in H} = (g_{h^{-1}\omega})_{\omega \in H}$*. This gives a morphism* $\phi : H \to \mathrm{Aut}(K)$*, and this morphism allows us to define* $G \wr H = K \rtimes_\phi H$*.*

Now, it is straight-forward verifications to check that

$$L_N = (\mathbb{Z}/N) \wr \mathbb{Z}.$$

Moreover, a presentation of $\mathbb{Z}/N$ is simply $\langle a|a^N\rangle$, and a presentation of $\mathbb{Z}$ is $\langle t\rangle = \langle t|\varnothing\rangle$. We give the following presentation for the direct sum :

$$\bigoplus_{n \in \mathbb{Z}} \mathbb{Z}/N \cong \langle a_n,\ n \in \mathbb{Z} | a_n^N, a_n a_m a_n^{-1} a_m^{-1},\ m, n \in \mathbb{Z}\rangle.$$

Here, each a_n is the sequence $(x_m)_{m \in \mathbb{Z}}$ whose only non-zero element is at position n and equals the generator 1 of $\mathbb{Z}/N$. At last, the morphism is the following :

$$\begin{array}{cccc} \phi & : & \mathbb{Z} & \to & \mathrm{Aut}\left(\bigoplus_{n \in \mathbb{Z}} \mathbb{Z}/N\right) \\ & & m & \mapsto & [(x_n)_{n \in \mathbb{Z}} \mapsto (x_{n-m})_{n \in \mathbb{Z}}] \end{array}.$$

In particular, we obtain, by keeping in mind that we will use Theorem 1.6 :

$$\phi(t)(a_n) = a_{n+1}.$$

Therefore, we have the following presentation for L_N :

$$L_N \cong \langle t, a_n,\ n \in \mathbb{Z} | a_n^N, t a_n t^{-1} = a_{n+1}, n \in \mathbb{Z},\ a_n a_m a_n^{-1} a_m^{-1},\ m, n \in \mathbb{Z}\rangle.$$

Now, we are almost done. We only need to make use of *Tietze transformations.* Note that the relation $ta_nt^{-1} = a_{n+1}$ implies that $a_n = t^n a_0 t^{-n}$. Therefore, one can replace these generators and re-write the associated relations to obtain :

$$L_N \cong \langle t, a_0 | a_0^N,\ (t^n a_0 t^{-n})^N,\ tt^n a_0 t^{-n} t^{-1} = t^{n+1} a_0 t^{-(n+1)}, [t^n a_0 t^{-n}, t^m a_0 t^{-m}],\ m, n \in \mathbb{Z} \rangle.$$

By noting that the relation $tt^n a_0 t^{-n} t^{-1} = t^{n+1} a_0 t^{-(n+1)}$ is redundant, by noting that $a_0^N = 1$ if and only if $t^n a_0^N t^{-n} = 1$, and by re-labeling a_0 into a, one obtains the claimed presentation for L_N. ■

By looking for litterature about Tietze transformations, one may only find mentions of *elementary* transformations. In our setting, we made use of an *infinite* number of transformations (however, each transformation involved only a *finite* number of generators and relations, but it is possible to generalize). For more on this topic, see [5].

As usual, once we have a presentation, a natural question is about knowing whether we can reduce the number of generators or relations. In this case, we already have the minimal number of generators. Therefore, the question remains whether we can give a *finite* presentation for L_N, that is, with a finite number of *relations*, up to eventually adding some generators. It turns out the answer is no, by making use of a (strong) result from [1] :

Theorem 1.8. (Baumslag) *Let G and H be finitely presented groups. Then their (regular, restricted) wreath product $G \wr H$ is finitely presented if and only if either $G = 1$ or H is finite.*

In our case, it becomes evident that L_N is not satisfying the hypothesis of the theorem, and is therefore not finitely presentable.

2. Diestel-Leader graphs

We will detail the construction presented in [6]. Let T_p be the $(p+1)$-valent tree, $p \geqslant 2$. For instance, see figure 6 for a picture of T_4.

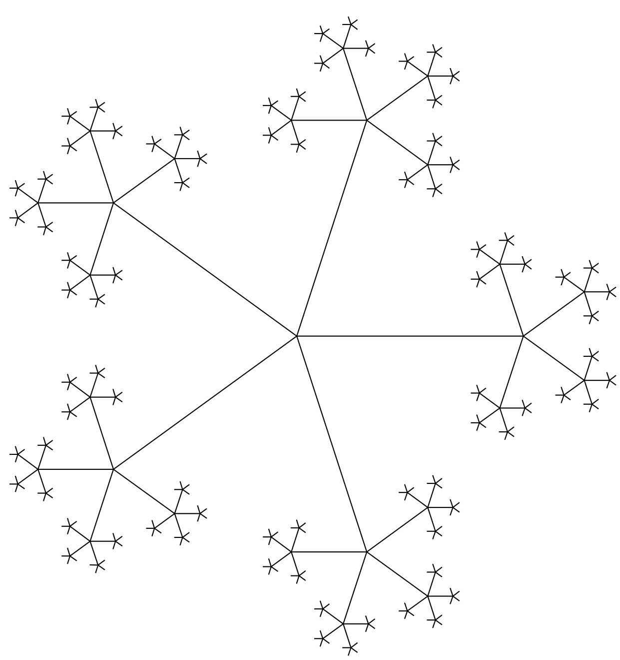

Figure 6. A representation of T_4. There is only a finite number of edges represented, the actual construction is fractal and infinite.

A *ray* is a sequence $(e_i)_{i\in\mathbb{N}}$ of edges of T_p such that $d(e_i, e_j) = |i - j|$ for all $i, j \in \mathbb{N}$, where d desnotes the graph metric on T_p. A *geodesic* is a sequence $(e_i)_{i\in\mathbb{Z}}$ with the same property.

Now, define two rays to be *equivalent* if the symmetric difference of the set of their elements is finite (that is, if, up to shifting one of the sequences, both agree after a certain rank). An *end* is an equivalence class of rays, and we denote the set of ends in T_p by ∂T_p. We also denote $\hat{T_p} = T_p \coprod \partial T_p$. The notation comes from topology, where the hat would denote some kind of compactification of the tree. We also say that a ray R *leads* to an end $\xi \in \partial T_p$ if ξ is represented by R.

By choice of an end $\omega \in \partial T_p$, one could picturally represent the situation as follows :

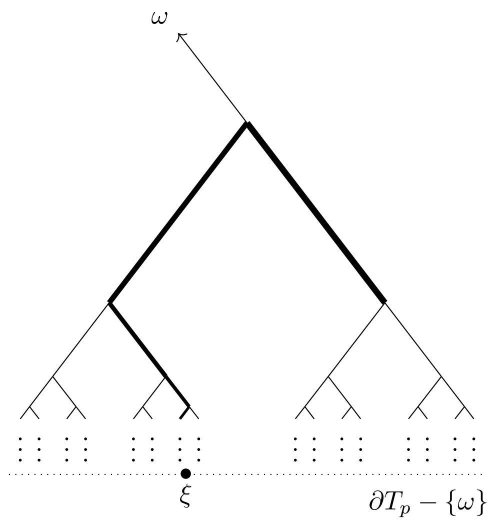

FIGURE 7. Representation of the ends (dotted line) and of a ray (bold) leading to an end $\xi \neq \omega$. Note that we truncated the tree : it goes infinitely far to the left and the right, as well as to the top and the bottom.

To define the Diestel-Leader graphs, we still have some work. First, note that for each $x \in T_p$ and each $\xi \in \partial T_p$, there exists a unique ray starting at x and leading to ξ. Indeed, for existence, if $(e_i)_{i\in\mathbb{N}}$ is a ray leading to ξ, then :

- either $x \in \{e_i,\ i \in \mathbb{N}\}$, and by taking $x = e_m$, we obtain that $(e_{i+m})_{i\in\mathbb{N}}$ is a ray leading to ξ,
- or $x \in \{e_i,\ i \in \mathbb{N}\}$, in which case we may consider the shortest path $x \rightsquigarrow e_0$ (which always exists and is unique in any tree), and write this path as $x = e_{-m} \to e_{-m+1} \to \ldots \to e_{-1} \to e_0$. Then, taking $(e_{i-m})_{i\in\mathbb{N}}$ gives a ray leading to ξ.

For uniqueness, we may prove it by induction. Let $(e_i)_{i\in\mathbb{N}}$ and $(f_i)_{i\in\mathbb{N}}$ both be suitable. We already have $e_0 = x = f_0$. Now, assume that $e_i = f_i$ for all $i \leqslant n$. Suppose that $e_{n+1} \neq f_{n+1}$. We are in this setting :

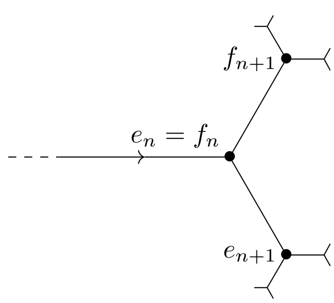

FIGURE 8. The inductive step for proving uniqueness.

By the very definition of a ray, going backwards is not allowed. Therefore, the new rays $(e_{i+n+1})_{i\in\mathbb{N}}$ and $(f_{i+n+1})_{i\in\mathbb{N}}$ are disjoint, yet both leading to ξ since only differing from the original rays by $n+1$ terms. This is contradictory. ■

We could also prove in a similar fashion that given two ends $\xi \neq \chi \in \partial T_p$, there exists a unique geodesic linking ξ and χ (that is, denoting this geodesic as $(e_i)_{i\in\mathbb{Z}}$, the choice of any $m \in \mathbb{Z}$ gives two rays $(e_{m+i})_{i\in\mathbb{N}}$ and $(e_{m-i})_{i\in\mathbb{N}}$ leading respectively to ξ and χ).

We can use this previous property to define the *confluent* $x \curlywedge y$ of two vertices $x, y \in T_p$ with respect to an end $\omega \in \partial T_p$. Denote as $\overline{x\omega}$ the unique ray leading to ω and starting at x. Then both $\overline{x\omega}$ and $\overline{y\omega}$ lead to ω, thus their intersection is also a ray leading to ω. Define $x \curlywedge y$ to be the starting point c of this new ray :

$$\overline{x\omega} \cap \overline{y\omega} = \overline{c\omega}.$$

On a picture, it is simply the point at which the geodesics cross and start coïnciding :

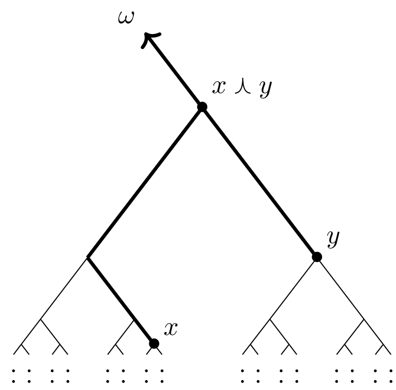

FIGURE 9. Locating the confluent of two vertices.

Now, fix both an end $\omega \in \partial T_p$ and a vertex $o \in T_p$. Define the *Busemann function* (also called *height function*) on T_p by :

$$\mathfrak{h}(x) = d(x, x \curlywedge o) - d(o, x \curlywedge o).$$

It is convenient to define the *horocycles* of this function by :

$$H_k = \{x \in T_p \ / \ \mathfrak{h}(x) = k\}.$$

We immediately see that $o \in H_0$. Horocycles allow us to represent the Busemann function :

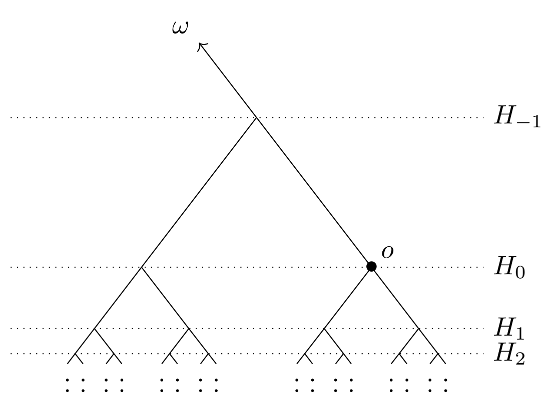

FIGURE 10. The busemann function and its horocycles.

The horocycles satisfies the following :

Proposition 2.1. *Each horocycle is infinite, and $(H_k)_{k\in\mathbb{Z}}$ is a partition of T_p. Moreover, each $x \in H_k$ has one neighbor in H_{k-1} (its parent) and p neighbors in H_{k+1} (its children).*

Proof. The partition is evident. The rest is intuitive, but proofs are no so straight-forward.

Let us first prove that H_k is infinite. We will construct arbitrarily many elements of H_k as follows : denote as $(e_n)_{n\in\mathbb{N}} = \overline{o\omega}$ the ray starting at o and leading to ω. For any $n \in \mathbb{N}$ such that $n+k \geqslant 0$, take $n+k$ steps starting at e_n that are not backtracking (the set of all steps has maximal cardinality) and such that the only such step lying on $\overline{o\omega}$ is e_n itself. For instance, in figure 9, if $y = o$, we have drawn the case $k = 2$ and $n = 1$.

By denoting as γ_n the $n+k$ steps starting at e_n, we end up at a vertex x_n and γ_n is the shortest path $e_n \rightsquigarrow x_n$. We have :

1. $x_n \in H_k$. Indeed, by taking the reversed path $\gamma_n^- : x_n \rightsquigarrow e_n$, and then following the ray $(e_n, e_{n+1}, \ldots)$, we obtain, by uniqueness, the ray $\overline{x_n\omega}$. Similarly, the ray $(e_n, e_{n+1}, \ldots)$ is $\overline{e_n\omega}$. Therefore, we obtain :

$$\overline{x_n\omega} \cap \overline{o\omega} = \overline{e_n\omega},$$

   that is $x_n \curlywedge o = e_n$. This property can be seen on figure 9 as well. We can now compute :

$$\mathfrak{h}(x_n) = d(x_n, e_n) - d(o, e_n) = n + k - n = k,$$

   that is $x_n \in H_k$.

2. For $m \neq n$, we have $x_m \neq x_n$. Indeed, by taking the shortest path $\phi : e_m \rightsquigarrow e_n$, we obtain, by the non-backtracking hypothesis of $\gamma_\bullet$ and by $\gamma_\bullet \cap \overline{o\omega} = \{e_\bullet\}$, that the concatenate $\gamma_m^- \bullet \phi \bullet \gamma_n$ is the shortest path $e_m \rightsquigarrow e_n$. Therefore, we have :

$$d(x_m, x_n) = \ell(\gamma_m^-) + \ell(\phi) + \ell(\gamma_n) \geqslant m + k + n + k > 0$$

   since $m \neq n$.

In particular, we have constructed an injective sequence $(x_n)_{n \geqslant |k|}$ of elements of H_k, so H_k is infinite.

Now, fix $x \in T_p$. Denote as y_0 the neighbor of x along $\overline{x\omega}$ (that is : $\overline{x\omega} = (x, y_0, *, \ldots)$), and let $y_1, \ldots, y_p$ be the remaining p neighboors of x. We shall prove that $\mathfrak{h}(y_0) = \mathfrak{h}(x) - 1$ and that $\mathfrak{h}(y_i) = \mathfrak{h}(x) + 1$ for $i \geqslant 1$. We can construct $\overline{y_i\omega}$ from $\overline{x\omega}$ by either removing x from $\overline{x\omega}$ if $i = 0$, or by appending y_i prior to $\overline{x\omega}$ otherwise. Those observations will allow us to express $y_i \curlywedge o$ conveniently. We need to distinguish two cases :

1. First case : $x \in \overline{o\omega}$. If $x = o$, then $y_0 \curlywedge o = y_0$ and $y_i \curlywedge o = o$ for $i \geqslant 1$, and thus we obtain :

$$\mathfrak{h}(y_0) = d(y_0, y_0) - d(o, y_0) = -1 \qquad \text{and} \qquad \mathfrak{h}(y_i) = d(y_i, o) - d(o, o) = 1 \text{ for } i \geqslant 1.$$

   Now, if $x \neq o$, then $y_0 \curlywedge o = y_0$ and $y_i \curlywedge o = x$ for $i \geqslant 1$ (and $x \curlywedge o = x$). We then have :

$$\mathfrak{h}(y_0) = d(y_0, y_0) - d(o, y_0) = -[d(o, x) + 1] = d(x, x \curlywedge o) - d(o, x \curlywedge o) - 1 = \mathfrak{h}(x) - 1$$

   and

$$\mathfrak{h}(y_i) = \mathfrak{h}(x) + 1$$

   by a similar argument for $i \geqslant 1$.

2. Second case : $x \notin \overline{o\omega}$, in which case we obtain $y_i \curlywedge o = x \curlywedge o$ for all $i \geqslant 0$. In this case too, we can compute directly :

$$\mathfrak{h}(y_0) = \mathfrak{h}(x) - 1 \qquad \text{and} \qquad \mathfrak{h}(y_i) = \mathfrak{h}(x) + 1 \text{ for } i \geqslant 1.$$

■

The previous result will allow us to label the tree T_p. Recall that for any graph $\Gamma = (V, E)$ and any set X, a *labelling* in X of Γ is a function $f : E \to X$.

In T_p, since every vertex has p children, we can choose[1] a labelling of T_p in $\mathbb{Z}/p$ such that given any vertex $x \in T_p$ and its children $y_1, ..., y_p$, the edges $(x, y_1), ..., (x, y_p)$ are all labelled with distinct elements of $\mathbb{Z}/p$. This definition is correct, since all edges are connecting two vertices, with one being the child of the other. Moreover, we can always choose this labelling such that all edges appearing in the ray $\overline{o\omega}$ are labelled 0. We shall denote as $L_{\mathfrak{h}}$ the labelling function of T_p (which is dependant on $\mathfrak{h}$!).

Now, we can give another description of T_p and its horocyclic structure :

Proposition 2.2. *Define Σ_p to be the set of finitely-supported sequences in $\mathbb{Z}/p$. Then, there is a bijection $T_p \cong \Sigma_p \times \mathbb{Z}$, where an element $x \in T_p$ is sent to $((\sigma_n)_{n\in\mathbb{N}}, k)$, with $k = \mathfrak{h}(x)$ and $\sigma_n = L_{\mathfrak{h}}(e_n \to e_{n+1})$, where $(e_n)_{n\in\mathbb{N}} = \overline{x\omega}$. Moreover, if x corresponds to (σ, k), then its parent vertex corresponds to $(T(\sigma), k-1)$, where T is the truncation operator, that is : $T((\sigma_n)_{n\in\mathbb{N}}) = (\sigma_{n+1})_{n\in\mathbb{N}}$.*

Before we give the formal proof of this statement, let us make yet another picture to represent the labelling and the correspondance. When drawing the situation, we can always assume that we order the edges by their label, which provides in particular that the ray $\overline{o\omega}$ is on the right-most part of the picture :

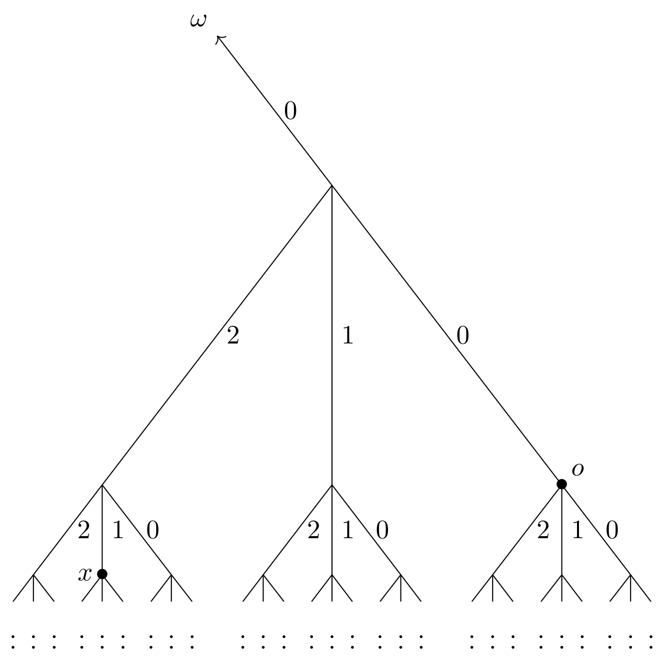

FIGURE 11. Labelling edges on T_p and locating vertices. Here, x corresponds to $(...\overline{0021}, 1)$, where the sequence $(1, 2, 0, 0, ...)$ is represented in number notation. Its parent corresponds to $(...\overline{002}, 0)$.

Proof. First, the application is well-defined, that is, if x is mapped to $((\sigma_n)_{n\in\mathbb{N}}, k)$, we have that $(\sigma_n)_{n\in\mathbb{N}}$ has finite support. Indeed, let $\overline{x\omega} = (e_n)_{n\in\mathbb{N}}$, and let $m \in \mathbb{N}$ be such that $x \curlywedge o = e_m$. Then, we have :

$$n \geqslant m \implies e_n \in \overline{o\omega},$$

and thus : $n \geqslant m \implies \sigma_n = 0$, by using the fact that we chose to label the edges of $\overline{o\omega}$ as 0.

[1]This requires the axiom of choice !

Now, if x has x^- as parent, we have that $\overline{x^-\omega}$ is obtained from $\overline{x\omega}$ by removing the starting x. Therefore, if x is mapped to (σ, k), we indeed have that x^- is mapped to $(T(\sigma), k-1)$.

We shall prove that the application is bijective in two steps :

1. The application is injective : assume both x and y are mapped to $((\sigma_n)_{n\in\mathbb{N}}, k)$. We shall prove by induction over $\max\{n \geqslant \ / \ \sigma_n \neq 0\}$ that $x = y$. Indeed :

   - If $\sigma \equiv 0$, we have that $x, y \in \overline{o\omega}$, and by $\mathfrak{h}(x) = k = \mathfrak{h}(y)$, we obtain $x = y$.
   - Assume the result holds whenever $\max\{n \geqslant 0 \ / \ \sigma_n \neq 0\} = N$, and suppose we are in the case where $\max\{n \geqslant 0 \ / \ \sigma_n \neq 0\} = N + 1$. By the previous observation, we have that both x^- and y^- are represented by $(T(\sigma), k-1)$. Therefore, we can apply the induction hypothesis, and we obtain that $x^- = y^-$, that is, x and y are siblings (children of the same vertex). Now, the labelling $L_\mathfrak{h}$ is one-to-one from the edges linked to the children of x^- to $\mathbb{Z}/p$, and thus $x = y$.

2. The mapping is surjective : choose any $(\sigma, k) \in \Sigma_p \times \mathbb{Z}$. It can also be proven by induction over $\max\{n \geqslant 0 \ / \ \sigma_n \neq 0\}$ that we can find a vertex x mapped to the prescribed element :

   - If $\sigma \equiv 0$, then two cases are to be seperated. At first, if $k \geqslant 0$, then take $x = e_k$, where we wrote $\overline{o\omega} = (e_n)_{n\in\mathbb{N}}$. However, if $k < 0$, then start from o, and take the path $f_0 = e_0 \to f_1 \to \ldots$ from o to its successive descendants, by only choosing the edge labelled 0 at each generation. Taking $x = f_k$ also yields a suitable element.
   - Now, assume we can find a vertex x mapped to (σ, k) (this is taken to be true for all $k \in \mathbb{Z}$) when $\max\{n \geqslant 0 \ / \ \sigma_n \neq 0\} = N$, and suppose $\max\{n \geqslant 0 \ / \ \sigma_n \neq 0\} = N + 1$. We can therefore find a vertex y mapped to $(T(\sigma), k-1)$. Now, choose the child of y that is linked to y by the edge labelled σ_0. This is a suitable element.

■

As a side note, one can notice that if x is represented by (σ, k), then the quantity to which we applied induction is :

$$\max\{n \geqslant 0 \ / \ \sigma_n \neq 0\} = d(x, x \curlywedge o).$$

We are now ready to define the Diestel-Leader graphs :

Definition 2.3. *Let $p, q \geqslant 2$. The* Diestel-Leader graph $\mathrm{DL}(p, q)$ *is the graph whose vertex set is*

$$\{(x, y) \in T_p \times T_q,\ \mathfrak{h}(x) + \mathfrak{h}'(y) = 0\},$$

with $\mathfrak{h}$ and $\mathfrak{h}'$ two Busemann functions on T_p and T_q respectively, and where adjacencies are given by :

$$(x, y) \leftrightarrow (x', y') \iff x \leftrightarrow x' \text{ and } y \leftrightarrow y'.$$

It is not immediate however that this is a well-defined object. Indeed, it may depend on the choices of the Busemann functions on T_p and T_q. However, we will prove that in fact, it yields isomorphic graphs.

Let us first make yet another picture of what the Diestel-Leader graph looks like. For this, plot the two rooted trees next to one another, one being flipped upside down, so that their respective horocycles with opposite heights are on the same level. Then, couples of vertices on the same level are vertices of the Diestel-Leader graph, and adjacencies are given by pairs of edges in the trees. See figure 12 for an example of $\mathrm{DL}(2, 3)$.

Evidently, $\mathrm{DL}(p, q)$ and $\mathrm{DL}(q, p)$ are isomorphic, where the isomorphism is $\varphi(x, y) = (y, x)$. Moreover, we shall denote as $\mathrm{DL}(p) = \mathrm{DL}(p, p)$ the case where $q = p$. Also, from now on, we shall denote a horocycle as $H_k = \{\mathfrak{h} = k\}$, to take the dependancy on the Busemann function into account.

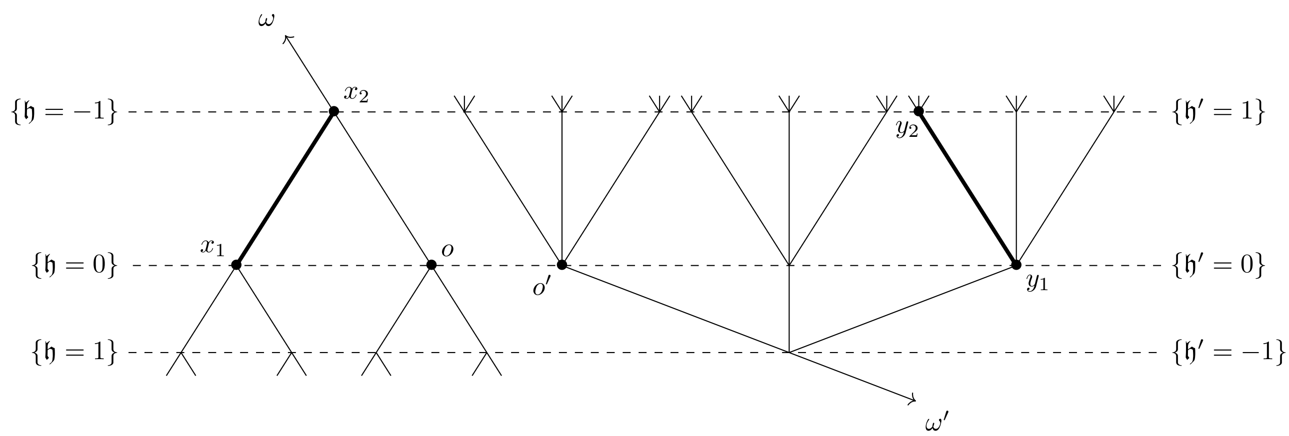

FIGURE 12. An example of a portion of $\mathrm{DL}(2,3)$, where $(T_2, \mathfrak{h})$ and $(T_3, \mathfrak{h}')$ are rooted respectively at $o \in T_2$ and $o' \in T_3$, and with respective ends $\omega \in \partial T_2$ and $\omega' \in \partial T_3$. Here, the vertices (x_1, y_1) and (x_2, y_2) are in $\mathrm{DL}(2,3)$, and are adjacent, as indicated by the bold lines.

We shall now prove that the construction of $\mathrm{DL}(p,q)$ does not depend on the choices of the Busemann functions. Let $\mathfrak{h}_1$ and $\mathfrak{h}_2$ be two Busesmann functions on T_p, with respective roots o_1 and o_2 and fixed and ω_1 and ω_2. Similarly, let $\mathfrak{h}'_1$ and $\mathfrak{h}'_2$ be two Busemann functions on T_q associated to roots and ends o'_1, ω'_1 and o'_2, ω'_2. Define X to be the Diestel-Leader graph whose vertices are $\{(x,y) \in T_p \times T_q \ / \ \mathfrak{h}_1(x) + \mathfrak{h}'_1(y) = 0\}$, and similarly, define Y to be the Diestel-Leader graph whose vertices are $\{(x,y) \in T_p \times T_q \ / \ \mathfrak{h}_2(x) + \mathfrak{h}'_2(y) = 0\}$.

By Proposition 2.2, one has isomorphisms

$$(T_p, \mathfrak{h}_1) \cong \Sigma_p \times \mathbb{Z} \cong (T_p, \mathfrak{h}_2) \qquad \text{and} \qquad (T_q, \mathfrak{h}'_1) \cong \Sigma_q \times \mathbb{Z} \cong (T_q, \mathfrak{h}'_2),$$

which we can compose to get two automorphisms $\varphi : T_p \to T_p$ and $\psi : T_q \to T_q$. One checks that, by definition of the labelling in Proposition 2.2, we have for the horocycles :

$$\varphi(\{\mathfrak{h}_1 = k\}) = \{\mathfrak{h}_2 = k\} \qquad \text{and} \qquad \psi(\{\mathfrak{h}'_1 = k\}) = \{\mathfrak{h}'_2 = k\},$$

as well as $\varphi(o_1) = o_2$ and $\psi(o'_1) = o'_2$. Therefore, defining $f : X \to Y$ by

$$f(x,y) = (\varphi(x), \psi(y))$$

indeed defines an isomorphism from X to Y, by direct computations. ■

As a remark, because of Proposition 2.1, one sees that each vertex of $\mathrm{DL}(p,q)$ has exactly $p+q$ neighbors.

What does $\mathrm{DL}(p,q)$ *look like* in a neighborhood of a vertex ? Up to changing the origins on T_p and T_q, we see that a neighborhood of any vertex will be the same as a neighborhood of (o, o'). In figure 12, denote as $z_1, z_2, z_3 \in \{\mathfrak{h}' = 1\}$ the three children of o' in T_3. Then, the following path is a non-contractible loop in $\mathrm{DL}(2,3)$:

$$(o, o') \to (x_1, z_1) \to (x_2, o') \to (x_1, z_2) \to (o, o').$$

This behaviour generalizes to any $\mathrm{DL}(p,q)$. In particular, even though each vertex has a constant number of neigboors, $\mathrm{DL}(p,q)$ is *not* a tree.

What would balls (centered at (o, o')) of increasing radius look like in $\mathrm{DL}(p,q)$? Something complicated, as in figure 13.

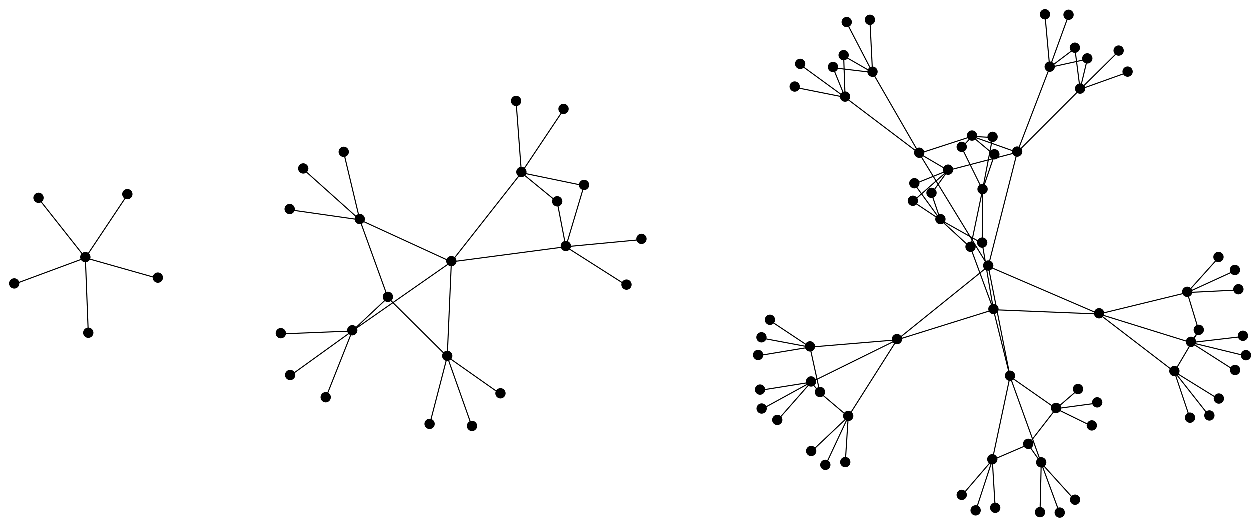

FIGURE 13. Balls in $\mathrm{DL}(2,3)$ centered at (o,o') whose radii increase from 1 to 3.

We shall digress a bit and describe the `Python` algorithm that was made to generate such images. First, the data structure used for the graph is :

```
import networkx as nx
import matplotlib.pyplot as plt

class GraphVisualization:
    def __init__(self):
        self.visual = []
        def addEdge(self,a,b):
            self.visual.append([a,b])
        def visualize(self):
            G = nx.Graph()
            G.add_edges_from(self.visual)
            nx.draw_networkx(G, with_labels=False, node_size=20,node_color="k")
```

Now, we first need to generate the trees T_p and T_q. In fact, we only need the balls of radius R centered at o or o'. This is done by a similar labelling system than in proposition 2.2 :

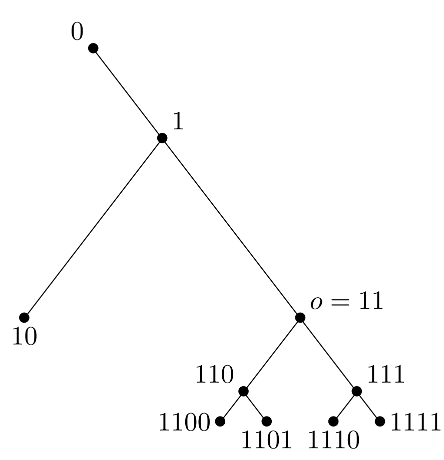

FIGURE 14. Labelling elements in the ball of radius 2 neighborhood in T_2.

Therefore, it is possible to generate the balls in the trees recursively :

```
def T(p,R):
    T = GraphVisualization()
    def addChildren(vertex,gen):
        if(gen>0):
            for k in range(p):
                T.addEdge(vertex, vertex+str(k))
                addChildren(vertex+str(k),gen-1)
    o = R*str(p-1)
    addChildren(o,R)
    for k in range(1,R+1):
        addChildren(o[:(R-k)],R-k)
    T.addEdge("0",str(p-1))
    return T
```

Indeed, we check that o is represented by the string composed with R times the digit $p-1$. Also, notice that this algorithm is far from being efficient : lots of edges are calculated several times.

Now, one sees that the height of a vertex x represented by a string of length ℓ is given by the relation :

$$\mathfrak{h}(x) = \ell - R.$$

Moreover, by simply applying the definition of the Diestel-Leader graph :

$$E(\mathrm{DL}(p,q)) = \{(e,f) \in E(T_p) \times E(T_q) \ / \ \mathfrak{h}(\imath(e)) + \mathfrak{h}'(\imath(f)) = \mathfrak{h}(\jmath(e)) + \mathfrak{h}'(\jmath(f)) = 0\},$$

where $\imath(e)$ denotes the starting point of the edge e, and $\jmath(e)$ its ending point, one obtains the code :

```
def height(v,R):
    if(v=="0"): return -R
    return len(v)-R

def DL(p,q,R):
    Tp = T(p,R)
    Tq = T(q,R)
    G = GraphVisualization()
    for e1 in Tp.visual:
        for e2 in Tq.visual:
            ie1, je1 = e1[1], e1[0]
            ie2, je2 = e2[0], e2[1]
            if(height(ie1,R)+height(ie2,R)==0 and height(je1,R)+height(je2,R)==0):
                G.addEdge(ie1+","+ie2, je1+","+je2)
    G.visualize()
    plt.show()
```

One checks that the three calls `DL(2,3,1)`, `DL(2,3,2)` and `DL(2,3,3)` gives the three pictures in figure 13.

3. The Cayley graph of L_N

Recall that an element of L_N is a couple (σ, k) with $k \in \mathbb{Z}$ and $\sigma = (\sigma_n)_{n\in\mathbb{Z}}$ a doubly-infinite sequence. Now, thanks to proposition 2.2, one can associate to (σ, k) a *unique* element $x \in T_N$, where we denote by $\varphi : \Sigma_N \times \mathbb{Z} \to T_N$ the labelling, by :

$$L(\sigma, k) = \varphi((\sigma_{k-n})_{n\geqslant 0}, k).$$

Similarly, we can also associate a unique element of T_N by :

$$R(\sigma, k) = \varphi((\sigma_{k+n+1})_{n\geqslant 0}, -k).$$

Here, $L(\sigma, k)$ may be called the *left part* of (σ, k), and $R(\sigma, k)$ its *right part*. This allows to define a map $\Phi : L_N \to \mathrm{DL}(N)$ by :

$$\Phi(\sigma, k) = (L(\sigma, k), R(\sigma, k)).$$

Indeed, we have

$$\mathfrak{h}(L(\sigma, k)) = k = -\mathfrak{h}(R(\sigma, k))$$

by definition of φ, so that Φ is well-defined. This is in fact important because of the following :

Proposition 3.1. *The map $\Phi : L_N \to \mathrm{DL}(N)$ is one-to-one.*

Proof. Assume first that $\Phi(\sigma, k) = \Phi(\varsigma, \ell)$. Since φ is one-to-one, we obtain from the left parts :

$$k = \ell \qquad \text{and} \qquad \sigma_{k-n} = \varsigma_{k-n} \text{ for all } n \geqslant 0.$$

Similarly, we obtain $\sigma_{k+n+1} = \varsigma_{k+n+1}$ for all $n \geqslant 0$ from the right parts, so that $(\sigma, k) = (\varsigma, \ell)$. Now, if $(x, y) \in \mathrm{DL}(N)$, then $x = \varphi(\sigma, k)$ and $y = \varphi(\varsigma, -k)$ for some sequences $(\sigma_n)_{n\geqslant 0}$ and $(\varsigma_n)_{n\geqslant 0}$ and some $k \in \mathbb{Z}$, since, once again, φ is one-to-one. Now, consider the sequence

$$\lambda_n = \begin{cases} \sigma_{k-n} & \text{if } n \leqslant k \\ \varsigma_{n-k-1} & \text{if } n > k \end{cases}.$$

Then $(\lambda, k) \in L_N$ is such that $L(\lambda, k) = x$ and $R(\lambda, k) = y$, by construction, that is :

$$\Phi(\lambda, k) = (x, y).$$

■

Now, we will prove that $\mathrm{DL}(N)$ is the Cayley graph of L_N with respect to the following generating set :

$$L_N = \langle T, AT, ..., A^{N-1}T \rangle,$$

where we denoted $T = (\mathbf{0}, 1)$ and $A = (\delta_0, 0)$.

Fix an element $(x_1, x_2) \in \mathrm{DL}(N)$, and choose y_1 to be the child of x_1 downwards the edge labelled w, and denote as x_2^- to be the parent of x_2 :

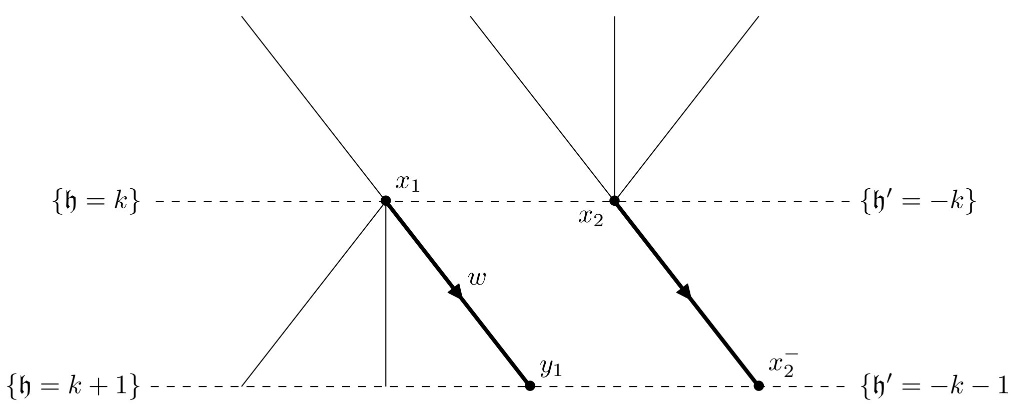

Figure 15. Moving in the Diesteal-Leader graph DL(3).

By the one-to-one correspondance established previously, we can choose a unique $(\sigma, k) \in L_N$ such that $(x_1, x_2) = \Phi(\sigma, k)$. In fact, we even have (and this justifies the terminology) :

$$x_1 = L(\sigma, k) = \varphi((\sigma_{k-n})_{n\geqslant 0}, k) \qquad \text{and} \qquad x_2 = R(\sigma, k) = \varphi((\sigma_{n+k+1})_{n\geqslant 0}, -k).$$

Recall that in proposition 2.2, we denoted as $T(u_n)_{n\geqslant 0} = (u_{n+1})_{n\geqslant 0}$ the truncation operator. Denote as $w \bullet (u_n)_{n\geqslant 0}$ the appending of w to the beginning of the sequence $(u_n)_{n\geqslant 0}$, so that :

$$y_1 = \varphi(w \bullet (\sigma_{k-n})_{n\geqslant 0}, k+1) \qquad \text{and} \qquad x_2^- = \varphi(T(\sigma_{k+n+1})_{n\geqslant 0}, k-1).$$

We therefore have $(y_1, x_2^-) = \Phi(\tilde{\sigma}, k+1)$, with $\tilde{\sigma}$ being the gluing of the two tweaked sequences $x \bullet (\sigma_{k-n})_{n\geqslant 0}$ and $T(\sigma_{k+n+1})_{n\geqslant 0}$. Schematically, we have :

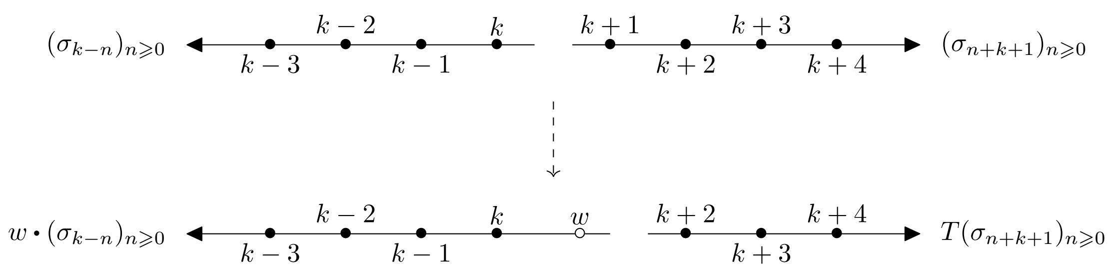

FIGURE 16. Cutting, tweaking and gluing.

In paticular, we see that $\tilde{\sigma}$ differs from σ only at position $k+1$, where we replaced σ_{k+1} by w. This means that

$$(y_1, x_2^-) = \Phi((\sigma, k) \star A^{\ell} T)$$

for some ℓ such that $\sigma_{k+1} + \ell = w$ in $\mathbb{Z}/N$, that is for $\ell = w - \sigma_{k+1}$. We could argue the exact same if taking the parent x_1^- of x_1 and a child y_2 of x_2, and we would instead have

$$(x_1^-, y_2) = \Phi((\sigma, k) \star (A^{\ell} T)^{-1})$$

for some ℓ.

By the labelling, this gives a one-to-one correspondance between edges from (x_1, x_2) and products of (σ, k) with the $2N$ elements $T, AT, ..., A^{N-1}T$ and their inverses. That is, we have proven that the Cayley graph of L_N with respect to the generating set as above is $\mathrm{DL}(N)$. ∎

Here are the balls of radii 2 and 3 in the Cayley graph of L_2 :

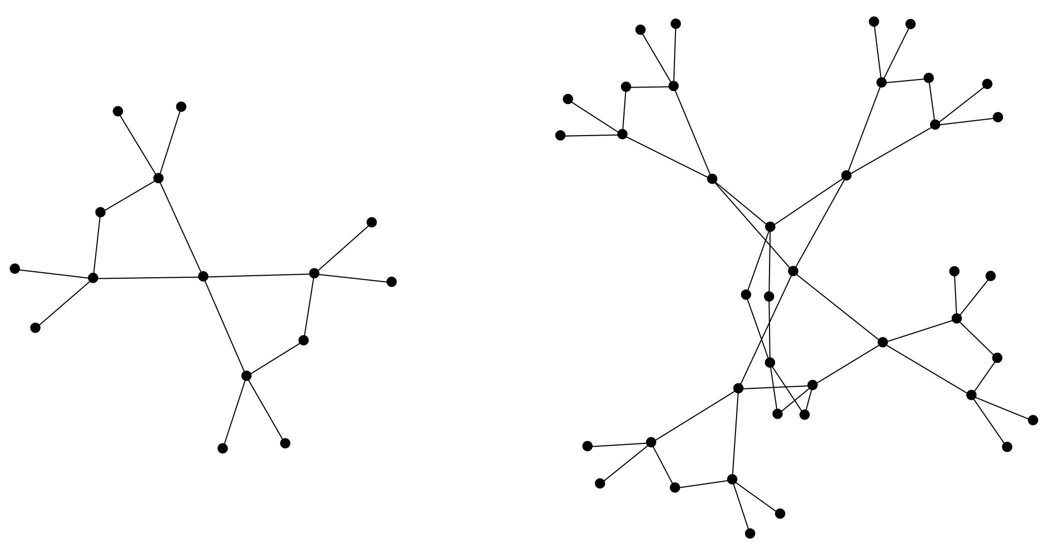

FIGURE 17. Balls of radii 2 and 3 in DL(2), that is, in the Cayley graph of L_2.

4. Word length in L_N

Recall the group presentation for L_N :

$$L_N \cong \langle a, t | a^N,\ [t^i a t^{-i}, t^j a t^{-j}],\ i, j \in \mathbb{Z} \rangle.$$

Defining $\alpha_i = t^i a t^{-i}$, we can re-write this as :

$$L_N \cong \langle a, t | a^N,\ [\alpha_i, \alpha_j],\ i, j \in \mathbb{Z} \rangle.$$

Now, let $g = (\sigma, m) \in L_N$ be an element of the lamplighter group. By definition, $\mathrm{Supp}(\sigma)$ is finite. Define :

$$\mathrm{Supp}(\sigma) \cap \mathbb{N} = \{i_1, ..., i_k\} \qquad \text{and} \qquad \mathrm{Supp}(\sigma) \cap (-\mathbb{N}^\star) = \{-j_1, ..., -j_\ell\},$$

with $0 \leqslant i_1 < ... < i_k$ and $0 < j_1 < ... < j_\ell$. Now, for all $s \in [\![1, k]\!]$, define $e_s = \sigma_{i_s}$, and for $t \in [\![1, \ell]\!]$, define $f_t = \sigma_{j_t}$. This allows us to define :

Definition 4.1. *Using the previous notations, one can check that* $g = \alpha_{i_1}^{e_1} \star ... \star \alpha_{i_k}^{e_k} \star \alpha_{-j_1}^{f_1} \star ... \star \alpha_{-j_\ell}^{f_\ell} \star t^m$. *This is called the* normal form *of* g.

To represent g, this is doing the following :

1. Go to the first positive index where a lamp is not off, and turn it to its state.
2. Move right to the next one, rinse and repeat until the right-most lamp has been lit.
3. Go to the first non-positive index where a lamp is not off, and turn it to its state.
4. Move left to the next one, rinse and repeat again, until done.
5. Move to the ending position.

Indeed, in the normal form, one sees that the $\alpha_i \star \alpha_j$ have cancelling pairs of the form $t^{-i} \star t^j = t^{j-i}$, so this boils down to doing exactly this algorithm. Note that α_i moves to position i, moves the lamp to its next state, and goes back to the origin.

Proposition 4.2. *Let* $g \in L_N$ *have normal form* $g = \alpha_{i_1}^{e_1} \star ... \star \alpha_{i_k}^{e_k} \star \alpha_{-j_1}^{f_1} \star ... \star \alpha_{-j_\ell}^{f_\ell} \star t^m$. *Define :*

$$D(g) = \sum_{s=1}^{k} e_s + \sum_{t=1}^{\ell} f_t + \min\left\{2i_k + j_\ell + |m + j_\ell|, 2j_\ell + i_k + |m - i_k|\right\}.$$

Then $D(g)$ *is the word length of* g *with respect to the generating set* $L_N = \langle A, T \rangle$.

Proof. Denote a $\mathcal{L}(g)$ the word length of g.

- We first have $\mathcal{L}(g) \leqslant D(g)$. Indeed, taking the normal form of g, we see that (dropping the star symbols) :
$$g = t^{i_1} a^{e_1} t^{i_2 - i_1} a^{e_2} ... t^{i_k - i_{k-1}} a^{e_k} t^{-j_1 - i_k} a^{f_1} t^{j_1 - j_2} a^{f_2} ... t^{j_{\ell-1} - j_\ell} a^{f_\ell} t^{m + j_\ell}.$$
  We can count that there are $e_1 + ... + e_k + f_1 + ... + f_\ell$ occurrences of a, and the number of occurrences of t is :
$$\cancel{i_1} + \cancel{i_2} - \cancel{i_1} + ... + i_k - \cancel{i_{k-1}} + \cancel{j_1} + i_k + \cancel{j_2} - \cancel{j_1} + ... + j_\ell - \cancel{j_{\ell-1}} + |m + j_\ell| = 2i_k + j_\ell + |m + j_\ell|.$$
  In particular, we have :
$$\mathcal{L}(g) \leqslant \sum e_s + \sum f_t + 2i_k + j_\ell + |m + j_\ell|.$$
  In a similar fashion, we can take the normal form and re-arrange terms :
$$g = \alpha_{-j_1}^{f_1} \star ... \star \alpha_{-j_\ell}^{f_\ell} \star \alpha i_1{}^{e_1} \star ... \star \alpha_{i_k}^{e_k} \star t^m.$$
  Expanding everything yields this time :
$$\mathcal{L}(g) \leqslant \sum e_s + \sum f_t + 2j_\ell + i_k + |m - i_k|.$$
  Taking minima provides $\mathcal{L}(g) \leqslant D(g)$.

- To get a lower bound, notice that if $g = (\sigma, m)$ with $\#\mathrm{Supp}(\sigma) = n$, then there must be at least n occurrences of *powers* of a to light the corresponding lamps. We see that therefore, there must be at least n' occurences of a with n' being the sum of the powers of a, that is $n' = \sum e_s + \sum f_t$.

  Choosing any minimal representative for g, notice that the sum of the exponents of all occurrences of t adds up to m. Moreover, the partial sums of exponents of all occurrences of t up to being at position p adds up to p. For instance, at the time when the right-most lamp is being lit, the exponent sum of occurrences of t adds up to i_k. Similarly, at the time of lighting up the left-most lamp, the exponent sum is $-j_\ell$.

  There are two cases (still considering a minimal representative for g) :

  (a) If the right-most lamp is lit before the left-most, then the partial sums of the exponents of all occurrences of t take the successive values 0, i_k, $-j_\ell$ and finally m. In particular, in between, there must be at least i_k, $i_k + j_\ell$ and $|m + j_\ell|$ occurrences of t to go from each of these states to the next. This means that there are at least $2i_k + j_\ell + |m + j_\ell|$ occurrences of t.

  (b) Similarly, if the left-most lamp is lit before the right-most, then the partial sums take the values 0, $-j_\ell$, i_k and then m, meaning that at least $j_\ell + i_k + j_\ell + |m - i_k| = 2j_\ell + i_k + |m - i_k|$ steps are needed.

  In either case, we need at least the least of both values occurrences of t, which gives the desired lower-bound $\mathcal{L}(g) \geqslant D(g)$.

■

Notice that using this result, it is possible to use the normal form as in the proof to explicitly give a minimal representative of g :

1. Compute the normal form.
2. Compute the word length.
3. Accordingly to whichever of the two quantities to minimize is the smallest, choose the minimal normal form representing g accordingly.
4. Expand and reduce the cancelling pairs.

This pseudo-code translates directly to Python code as follows :

```
def wordLength(sigma, m):
    indI = []
    indJ = []
    for x in sigma:
        if x[0]>=0: indI.append(x)
        else: indJ.append((-x[0],x[1]))
    indI = sorted(indI,key=lambda x:x[0])
    indJ = sorted(indJ,key=lambda x:x[0])
    EF = sum([ x[1] for x in sigma ])
    ik = indI[-1][0]
    jl = indJ[-1][0]
    length = EF+min(2*ik+jl+abs(m+jl), 2*jl+ik+abs(m-ik))
    w = ""
    if(2*ik+jl+abs(m+jl)<=2*jl+ik+abs(m-ik)):
        if(indI[0][0]!=0): w += "t^{"+str(indI[0][0])+"}"
        w += "a^{"+str(indI[0][1])+"}"
        for i in range(len(indI)-1):
            w += "t^{"+str(indI[i+1][0]-indI[i][0])+"}a^{"+str(indI[i+1][1])+"}"
```

```
        w += "t^{"+str(-indI[-1][0]-indJ[0][0])+"}a^{"+str(indJ[0][1])+"}"
        for j in range(len(indJ)-1):
            w += "t^{"+str(indJ[j][0]-indJ[j+1][0])+"}a^{"+str(indJ[j+1][1])+"}"
        if(m+indJ[-1][0]!=0): w += "t^{"+str(m+indJ[-1][0])+"}"
    else:
        if(indJ[0][0]!=0): w += "t^{"+str(-indJ[0][0])+"}"
        w += "a^{"+str(indJ[0][1])+"}"
        for j in range(len(indJ)-1):
            w += "t^{"+str(indJ[j][0]-indJ[j+1][0])+"}a^{"+str(indJI[j+1][1])+"}"
        w += "t^{"+str(indJ[-1][0]+indI[0][0])+"}a^{"+str(indI[0][1])+"}"
        for i in range(len(indI)-1):
            w += "t^{"+str(indI[i+1][0]-indI[i][0])+"}a^{"+str(indI[i+1][1])+"}"
        if(m-indI[-1][0]!=0): w += "t^{"+str(m-indI[-1][0])+"}"
    return length, w
```

To run it, simply input a list of couples `(i,s)` with `i` being the index of a lit lamp and `s` its state (*i.e.* an integer in $[\![1, N]\!]$), as well as the ending position of the lamplighter. The program returns the word length as well as a minimal representative. For example :

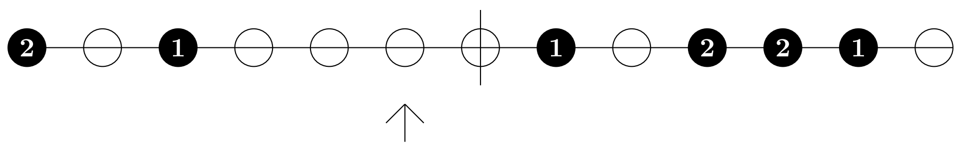

FIGURE 18. This element in L_3 is represented as the list `[(1,1),(3,2),(4,2),(5,1),(-4,1),(-6,2)]` with ending position `-1`. The program returns a world length of 30, and $tat^2a^2ta^2tat^{-9}at^{-2}a^2t^5$ as a minimal representative.

As a final remark, note that this program does *not* depend on the rank N of the lamplighter group L_N.

References

[1] Gilbert Baumslag. "Wreath products and finitely presented groups". In: *Mathematische Zeitschrift* 75.1 (Dec. 1961), pp. 22–28.

[2] Matt Clay and Dan Margalit. *Office Hours with a Geometric Group Theorist*. illustrated. Princeton University Press, 2017. Chap. 15, pp. 310–330.

[3] Sean Cleary and Jennifer Taback. "Dead end words in lamplighter groups and other wreath products". In: *Quarterly Journal of Mathematics* 56 (2003), pp. 165–178.

[4] David L. Johnson. *Presentations of Groups*. Ed. by J. W. Bruce and C. M. Series. illustrated, revised. Vol. 15. London Mathematical Society Student Texts. 1997. Chap. 10, p. 140.

[5] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial Group Theory : Presentations of Groups in Terms of Generators and Relations*. illustrated, unabridged, reprint. Dover books on mathematics. Courier Corporation, 2004. Chap. 1, pp. 48–55.

[6] Wolfgang Woess. "Lamplighters, Diestel–Leader Graphs, Random Walks, and Harmonic Functions". In: *Combinatorics, Probability and Computing* 14.3 (2005), pp. 415–433.

[7] Wolfgang Woess. "What is a horocyclic product, and how is it related to lamplighters?" In: *Internationale Mathematische Nachrichten* 224.1 (2013), pp. 1–27.