GROWTH OF NILPOTENT GROUPS

— SEMINAR NOTES —

HEINRICH-GREGOR ZIRNSTEIN

Contents

1. Introduction

These notes present some material from Gábor Pete's notes "Probability and Geometry on Groups" [Pet13]. My main goal was to spell out some details in order to make the text even more accessible to undergraduate students.

2. Cayley graphs

Let Γ be a group. We say that a subset $S \subset \Gamma$ *generates* the group if every element of Γ can we written as a product of elements from S and their inverses. In this case, we write $\Gamma = \langle S \rangle$. The elements of S are called *generators*.

A group is called *finitely generated* if we can find a generating set of finite size.

Definition 2.1. *Let Γ be a group and S a finite generating set. The* Cayley graph $Cay(\Gamma, S)$ *is the graph with*

$$\textit{vertex set } V = G \textit{ and edge set } E = \{(g, sg) : g \in \Gamma, s \in S\}.$$

In other words, each group element is a vertex in the Cayley graph, and two group elements are connected by an edge if we can reach one from the other by multiplying a generator on the left.

We usually only consider Cayley graphs where the generating set S is *symmetric*, that is for every element $a \in S$ it contains, it already contains its inverse $a^{-1} \in S$. This way, the Cayley graph is an undirected graph.

Example. The group $\mathbb{Z}^2$ is generated by two elements $\mathbb{Z}^2 = \langle (1,0), (0,1) \rangle$. Its Cayley graph $Cay(\mathbb{Z}^2, S)$ is the planar grid.

Date: January 9, 2014.

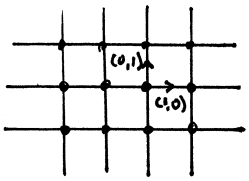

Cayley graph of the group $\mathbb{Z}^2$.

We would like to understand groups by studying their Cayley graphs. Of course, a single group may have many different Cayley graphs, depending on the choice of generators. However, they are all related by being *quasi-isometric*, a notion which we now define.

But first, note that every connected graph $G = (V, E)$ can also be interpreted as a metric space $X = (V, d)$ whose points are the vertices of the graph, and where the distance $d(x, y)$ between two vertices $x, y \in V$ is the minimum number of edges connecting these two vertices. Two vertices are linked by single edge iff their distance equals 1.

Definition 2.2. *Two metric spaces (X, d_X) and (Y, d_Y) are said to be* quasi-isometric *if there is a map $f : X \to Y$ such that the following two conditions hold*

(1) *There are $A, B > 0$ such that for all points $x, x' \in X$, we have*
$$\frac{1}{A} d_X(x, x') - \frac{B}{A} \leq d_Y(f(x), f(x')) \leq A d_X(x, y) + B.$$

(2) *For every $y \in Y$, there exists a point $x \in X$ such that*
$$d_Y(y, f(x)) < B.$$

The first condition says that the map f stretches or shortens distances only by a factor A and maybe adds some constant distance B. The second condition says that the image $f(X)$ is sufficiently dense in the space Y.

Example. $\mathbb{Z}^2$ is quasi-isometric to $\mathbb{R}^2$.

Lemma 2.3. *Let G be a group and S_1, S_2 two finite generating sets. Then, the Cayley graphs $Cay(G, S_1)$ and $Cay(G, S_1)$ are quasi-isometric.*

Proof. Exercise. □

As said, we want to understand groups by studying their Cayley graphs. We can now refine this plan: we want to study groups up to quasi-isomorphism (of their Cayley graphs).

For instance, we can consider random walks on Cayley graphs and ask whether they are recurrent or transient. It turns out that the answer only depends on the quasi-isometry class of the group.

Note that this plan only makes sense for infinite groups, as any two finite graphs are quasi-isomorphic. More generally, we don't distinguish between groups that differ only by some "finite" data, as the following lemma shows:

Lemma 2.4. *Let Γ be a finitely generated group and $H \leq \Gamma$ be a* finite index subgroup*, that is $[\Gamma : H] < \infty$. Then, the group H is also finitely generated, and any two Cayley graphs of the groups H and Γ are quasi-isometric.*

Remember that the index of a subgroup is the number of different cosets

$$[\Gamma : H] = \#\{gH : g \in \Gamma\}.$$

Keep in mind that an arbitrary subgroup of a finitely generated group does not need to be finitely generated itself, so it is not obvious that the subgroup H should be finitely generated.

Actually, let us prove a generalization of the lemma, which also makes the proof easier. Remember one of the most important notions form group theory: that of a group action. Groups never come alone, they always act on something.

Definition 2.5. *Let Γ be a group and X be some set. A* group action *of Γ on X is a homomorphism $\phi : \Gamma \to \operatorname{Aut} X$ from the group to the bijections of X. For simplicity, we often suppress the name of the map and write*

$$gx := g(x) := \phi(g)x$$

We also use the shorthand notation $\Gamma \curvearrowright X$ to mean that the group acts on the set.

Definition 2.6. *Let $\Gamma \curvearrowright X$ be a group action. The* orbit *of a point $x \in X$ is the set of points*

$$\Gamma x := \{gx : g \in \Gamma\} \subseteq X.$$

Note that the orbits partition the set X. The action is called transitive *if there is only a single orbit, that is every point can be reached from a single starting point by acting with some group element.*

Definition 2.7. *Let $\Gamma \curvearrowright X$ be a group action. The* stabilizer *of a point $x \in X$ is the set of group elements that leave the point invariant*

$$\operatorname{Stab}(x) = \{g \in \Gamma : gx = x\}.$$

Note that this is actually a subgroup of Γ. The action is called free *if the stabilizer group of every point is the trivial subgroup.*

Proposition 2.8. *Let G be a connected graph so that each vertex has finite degree. Let Γ be a group that acts on the graph by isometries. Moreover, assume that the action has finite stabilizers and finitely many orbits. Then, the group Γ is finitely generated and its Cayley graph is quasi-isometric to the graph G.*

Proof. Pick a reference point x_0. Since there are only finitely many orbits, we can find a ball $B = B_R(x_0)$ of radius R which contains a point from every orbit. This implies that the graph is a union of all translates of the ball, $V = \bigcup_{g\in\Gamma} g(B)$.

Consider the translates $g(B)$ of this ball. Some of them will intersect the slightly enlarged ball $B' = B_{R+1}(x_0)$, while others will be disjoint from it. Let

$$S = \{g \in \Gamma : g(B) \cap B' \neq \emptyset\} \setminus \{1\}$$

be the set of group elements where the translate intersects. Note that if $g \notin S$, $g \neq 1$ is any other element, then we have $d(g(B), B) \geq 2$ because $d(g(B), B') \geq 1$ and we must cross one additional edge to get from the larger ball to the original ball.

Claim: The set S is finite. Otherwise, by the pigeonhole principle, we could find two points $x \in B$, $y \in B'$ such that an infinite number of group elements $g_1, g_2, \dots$ would map one point to the other, $g_i x = y$. However, this would imply that infinitely many group elements $g_2^{-1} g_1, g_3^{-1} g_1, \dots$ stabilize the point x, in contradiction to our assumption that the stabilizers are finite.

Claim: The set S is a generating set. To prove this, let $h \in \Gamma$ be any element. Choose a path $x_0, x_1, \dots, x_{n+1} = h(x_0)$ of minimum length that connects the points x_0 and $h(x_0)$. In particular, $d(x_j, x_{j+1}) = 1$. Since the translates of the ball B cover the whole graph, there are group elements g_j such that $x_j \in g_j(B)$, where we choose $g_0 = 1$ and $g_{n+1} = h$. Moreover, we have

$$1 = d(x_j, x_{j+1}) \geq d(g_j(B), g_{j+1}(B)) = d(g_{j+1}^{-1} g_j(B), B).$$

By definition, this means that all the elements $h_j = g_{j+1}^{-1} g_j$ are either trivial or lie in S, and we have written $h = h_n \dots h_1 h_0$ as a product of generators.

It remains to be shown that the map $h \to h(x_0)$ is a quasi-isometry of graphs. In one direction, we have just given a representation in terms of generators such that

$$d_S(1, h) \leq d(x_0, h(x_0)).$$

In the other direction, the generators $h_j \in S$ cannot map the reference point too far away. For an element h with shortest representation $h = h_n \dots h_1 h_0$, we have

$$d(x_0, h(x_0)) = \sum_{j=0}^{n} d(x_0, h_j(x_0)) \leq 2R(n+1) = 2R d_S(1, h).$$

To conclude the proof, note that the metrics are invariant under the group action

$$d(x_0, h(x_0)) = d(g(x), gh(x_0)) \text{ and } d_S(1, h) = d_S(g, gh).$$

□

Proof of Lemma 2.4. The subgroup $H \leq \Gamma$ acts on the Cayley graph $Cay(\Gamma, S)$ by multiplication on the *right*, $h(g) := gh$. The orbits are precisely the right cosets $\{gH, g \in \Gamma\}$, of which there are only finitely many. Moreover, the stabilizers are trivial and the result follows. □

Exercise. Let Γ be a group and $F \lhd \Gamma$ be a finite normal subgroup such that the quotient Γ/F is finitely generated. Prove that the group Γ is also finitely generated and that the Cayley graphs of the groups Γ/F and Γ are quasi-isometric.

These two observations lead to the following definition:

Definition 2.9. *Two groups Γ_1, Γ_2 are called* virtually isomorphic *if there are finite index subgroups $H_i \leq \Gamma_i$ and finite normal subgroups $F_i \lhd H_i$ such that the quotients are isomorphic, $H_1/F_1 \cong H_2/F_2$.*

We have just shown that virtually isomorphic groups are quasi-isometric, i.e. they have essentially the same "geometry". Unfortunately, the reverse is not quite true. See [DK09] for a counterexample.

3. The Free Group

We now want to introduce the so-called free group, which is the grandmother of all groups.

The elements of the free group are defined as (equivalence classes of) words over an alphabet. In particular, consider an alphabet of exactly four letters a,b,a^{-1} and b^{-1}. It may be weird to think of the symbols "a^{-1}" and "b^{-1}" as "letters", but look at them as similar to the english letter "i" which consists of several disconnected components.

We can string these letters together to form words, for instance ab, $a^{-1}aa$ and $bab^{-1}a^{-1}$. There is also the empty word, which is kind of tricky to write down, so we denote it by the symbol ε.

Definition 3.1 (Free group on two generators.)**.** *Let $\Sigma = \{a, b, a^{-1}, b^{-1}\}$ be an alphabet of four letters. The set of all words on this alphabet is denoted by Σ^*. The* free group on two generators *$F_2 \subset \Sigma^*$ is defined as the set of those words which do not contain the following pairs of letters in sequence: aa^{-1}, $a^{-1}a$, bb^{-1} and $b^{-1}b$. The group operations are defined as*

1. *The identity element $1 = \varepsilon$ is the empty word.*
2. *The inverse u^{-1} is obtained by exchanging all letters $a \leftrightarrow a^{-1}$, $b \leftrightarrow b^{-1}$ in the word u.*
3. *The product $u \cdot v$ of two elements is obtained from the concatenated word uv by repeatedly removing all adjacent pairs aa^{-1}, $a^{-1}a$, bb^{-1} and $b^{-1}b$.*

Example. $(aab^{-1}) \cdot (ba^{-1}a) = aab^{-1}ba^{-1}a = aaa^{-1}a = aa$.

In other words, multiplication is defined such that the letters a and a^{-1} are treated as inverse of each other, ditto for b and b^{-1}.

This definition comes with a proof obligation: we have to check the group laws. But a more pressing issue is this: why is the group multiplication well-defined? It is not obvious that no matter in which *order* we cancel adjacent letters, we still get the same word at the end.

To show this, define two relation $\to, \to^*$ on words. We write $w \to w'$ if the word w' is obtained from w by cancelling exactly *one* pair of letters aa^{-1}, $a^{-1}a$, bb^{-1} or $b^{-1}b$. We say that one word *reduces* to the other, $w \to^* w'$, if the word w' is either equal to w, or if it is obtained by repeatedly cancelling letter pairs, no matter how often. (The relation $\to^*$ is the transitive closure of the relation $\to$.)

Lemma 3.2. *If we can cancel letters in two ways, $w \to w_1$ and $w \to w_2$, then there exists a word w' such that the reductions converge again, $w_1 \to^* w'$ and $w_2 \to^* w'$.*

Proof. First, consider the case where the cancelled pairs don't overlap. Example:

$$w = x \dots aa^{-1} \dots y \dots bb^{-1} \dots z$$

and we can cancel the first, $w_1 = x \dots y \dots bb^{-1} \dots z$, or the second pair, $w_2 = x \dots aa^{-1} \dots y \dots z$. Clearly, $w' = x \dots y \dots z$ will do the trick.

The interesting case is where the cancelled pairs do overlap. Example: If $w = aa^{-1}a$, then we can cancel either the first pair $(aa^{-1})a$ or the second pair $a(a^{-1}a)$. Fortunately, the reductions are actually equal $w_1 = w_2$. □

Lemma 3.3 (Confluence)**.** *Let w be a word that we can reduce in two ways $w \to^* w_1$ and $w \to^* w_2$. Then, there exists a word w' such that the reductions converge again, $w_1 \to^* w'$ and $w_2 \to^* w'$.*

Proof. Induction. □

This proves that the group multiplication is well-defined — we will always end up at the same word w'. Since word concatenation is associative, this also proves that group multiplication is associative.

Let us gain some more experience with the free group. Apparently, it is generated by two elements $F_2 = \langle a, b \rangle$. What does the Cayley graph look like?

Remark. The Cayley graph of the free group with respect to set of generators $S = \{a, b, a^{-1}, b^{-1}\}$ is the infinite 4-regular tree.

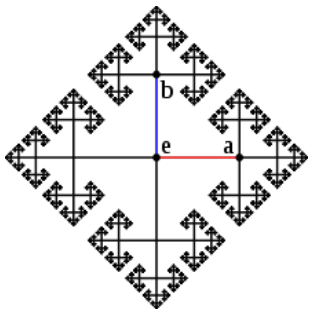

Cayley graph of the free group on two generators. (image from Wikipedia)

Remark. The free group has *exponential growth*, that is the set of vertices in the Cayley graph with distance $\leq R$ from a given point has $O(\alpha^R)$ elements. In this case, $\alpha = 3$.

Definition 3.4. *The* free group on k generators, F_k, *is defined in the same way, but with k letters $a_1, a_2, \dots a_k$ and their inverses $a_1^{-1}, a_2^{-1}, \dots, a_k^{-1}$. Similarly, we can define the free group on infinite sets of letters.*

The free group is the grandmother of all groups in the following sense

Lemma 3.5 (Freeness)**.** *Let S be a set, Γ be a group. Any map $f : S \to \Gamma$ can be extended to a homomorphism $\tilde{f} : F_S \to \Gamma$.*

Proof. Define $\tilde{f}(s_1^{k_1} s_2^{k_2} \cdots s_n^{k_n}) := f(s_1)^{k_1} f(s_2)^{k_2} \cdots f(s_n)^{k_n}$ and check that this is a homomorphism. □

Lemma 3.6. *Every group Γ is the quotient of a free group.*

Proof. Pick a generating set S for the group Γ. If the group is finitely generated, then the set S can be chosen to be finite; in the worst case, we have to choose $\Gamma = S$. The inclusion $S \to \Gamma$ gives rise to a homomorphism $\tilde{f} : F_S \to \Gamma$. Since the set is generating, this homomorphism is surjective and we have $\Gamma \cong F_S / \ker \tilde{f}$. □

A more concrete way to understand this is the following: any group Γ can be described in terms of *generators* and *relations*. For instance, consider the group $\mathbb{Z}^2$. It is generated by two elements $a = (1, 0)$ and $b = (0, 1)$ that fulfill the relation $ab = ba$.

$$\mathbb{Z}^2 \cong \langle a, b \mid ab = ba \rangle.$$

The notation on the right hand side is to be understood as the quotient of the free group generated by two elements a, b by the normal subgroup generated by the element $aba^{-1}b^{-1}$:

$$\langle a, b \mid ab = ba \rangle := F_2 / \langle\langle aba^{-1}b^{-1} \rangle\rangle$$

$$\langle\langle R \rangle\rangle = \text{smallest normal subgroup that contains the set } R$$

This is because an equation of the form $ab = ba$ is equivalent to an equation of the form $1 = aba^{-1}b^{-1}$, so every relation can, in fact, be represented by a group element that should become equal to the identity element. Taking the quotient will then "force" this element to be equal to the identity element.

Definition 3.7. *A group $\Gamma = \langle S|R \rangle$ given by generators S and relations R is said to be* finitely presented *if both S and R are finite.*

While a finitely presented group may seem very concrete at first, it can actually be fairly difficult to understand. The reason is that the relations can imply more complicated relations that are not obvious. In fact, the situation is worse:

Theorem 3.8 (Unsolvability of the word problem)**.** *There exists a group presentation $\Gamma = \langle S|R \rangle$ such that there is no algorithm which, given any word $g = \langle S \rangle$, can decide whether the corresponding group element $g \in \Gamma$ is equal to the identity element or not.*

We will not prove this theorem here.

Exercise. Prove that indeed $\mathbb{Z}^2 \cong \langle a, b \mid ab = ba \rangle$. In particular, solve the word problem for this group.

We will see a more elaborate example of generators and relations in the next section.

4. The Heisenberg Group

We have seen that the commutative groups $\mathbb{Z}^d$ have polynomial group $O(R^d)$, while the highly non-commutative free groups F_k have exponential growth. We now want to start with commutative groups and inject a little non-commutativity into them, which will lead us to the notion of *nilpotent groups*.

Before exploring the general notion, let first us consider a prototypical example, the *Heisenberg group*.

Definition 4.1. *The* Heisenberg group *H is the following group of upper triangular matrices*

$$H = \left\{ \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} : a, b, c \in \mathbb{Z} \right\}$$

where group multiplication is given by matrix multiplication.

We want to understand the Cayley graph of the Heisenberg group. To that end, we need a system of generators. consider the three matrices

$$x = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} = \mathbf{1} + e_{23}, \quad y = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} = \mathbf{1} + e_{12}, \quad z = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} = \mathbf{1} + e_{13}$$

where $\mathbf{1}$ is the identity matrix and e_{ij} is the matrix which has an entry 1 at row i and j and vanishing entries everywhere else. We have

$$e_{ij}e_{kl} = \begin{cases} e_{il} & \text{if } j = k \\ 0 & \text{otherwise} \end{cases}$$

and hence

$$\begin{aligned} x^a y^b z^c &= (\mathbf{1} + e_{23})^a (\mathbf{1} + e_{12})^b (\mathbf{1} + e_{13})^c = (\mathbf{1} + ae_{23})(\mathbf{1} + be_{12})(\mathbf{1} + ce_{13}) \\ &= \mathbf{1} + ae_{23} + be_{12} + ce_{13} = \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix}. \end{aligned}$$

In other words, the elements x, y and z generate the Heisenberg group, and every element can be written uniquely as a product $x^a y^b z^c$.

Moreoever, we have the following identities for multiplication from on the left

$$\begin{aligned} yx = (\mathbf{1} + e_{12})(\mathbf{1} + e_{23}) = \mathbf{1} + e_{12} + e_{23} + e_{13} = xyz, \quad zx = xz, \quad zy = yz \\ x(x^a y^b z^c) = x^{a+1} y^b z^c, \quad y(x^a y^b z^c) = x^a y^{b+1} z^{a+c}, \quad z(x^a y^b z^c) = x^a y^b z^{c+1}. \end{aligned}$$

We can now draw the Cayley graph of the Heisenberg group. The vertices of the graph are points in a three-dimensional grid and the edges are drawn according to the equations above.

Exercise. Draw the Cayley graph of the Heisenberg group.

Exercise. Show that the Heisenberg group is isomorphic to the group $\langle x, y, z \mid xy = yxz, xz = zx, yz = zy \rangle$ defined by generators and relations.

Exercise. Show that the Heisenberg group has quartic growth $O(R^4)$. This is in contrast to the group $\mathbb{Z}^3$, which has only cubic growth.

5. Nilpotent Groups

The *commutator* of two group elements a, b is defined as the group element

$$[a, b] := aba^{-1}b^{-1}.$$

In a non-commutative group, we generally cannot interchange two group elements. However, we can do so at the price of introducing a commutator

$$ab = [a, b]ba.$$

Of course, this only helps if commutators are easier to understand than the group elements themselves. For instance, in the Heisenberg group, we have $[x, y] = z$, $[x, z] = 1$ and $[y, z] = 1$, so the commutators are generated by the element z and we can commute elements at the price of introducing extra zs.

For subsets $A, B \subseteq \Gamma$ of a group, we write

$$[A, B] = \langle \{[a, b] : a \in A, b \in B\} \rangle$$

for the subgroup generated by commutators of their elements. The special subgroup $[\Gamma, \Gamma]$ is also a normal subgroup and the quotient

$$\Gamma^{Ab} := \Gamma / [\Gamma, \Gamma]$$

is an abelian group, called the *abelianization* of Γ. It is the group you obtain when all elements of the original group were to commute with each other.

We can define a general notion of groups whose commutators decrease in complexity.

Definition 5.1 (Nilpotent group)**.** *Let Γ be a group. Its* lower central series *is the sequence of groups*

$$\Gamma_0 \trianglerighteq \Gamma_1 \trianglerighteq \ldots \trianglerighteq \Gamma_s \trianglerighteq \ldots$$

where

$$\Gamma_0 = \Gamma \text{ and } \Gamma_{n+1} = [\Gamma_n, \Gamma].$$

The group is called s-step nilpotent *if the series reaches the trivial group after s steps, $\Gamma_s = \{1\}$.*

We have to show that the groups are indeed normal subgroups, as indicated. In fact, we have $\Gamma \trianglerighteq \Gamma_n$. This can proven inductively. If $\Gamma \trianglerighteq H$, then first of all $H \geq [H, \Gamma]$ because $[h, x] = h(xh^{-1}x^{-1}) \in H$. Moreover, we have

$$g[h, x]g^{-1} = (ghg^{-1})(gxg^{-1})(gh^{-1}g^{-1})(gx^{-1}g^{-1}) = [ghg^{-1}, gxg^{-1}] \in [H, \Gamma]$$

for any $g \in \Gamma$.

Example. The 1-step nilpotent groups are precisely the abelian groups.

Example. The Heisenberg group is 2-step nilpotent.

Note that when the group $\Gamma = \langle S \rangle$ is finitely generated, it is not necessarily true that the commutator subgroup $[\Gamma, \Gamma]$ is also finitely generated. In particular, it is generally not true that the commutators of the generators already generate the commutator subgroup, we often have $[S, S] \subsetneq [\Gamma, \Gamma]$. The trouble is that we also have elements like $[a, bc] \in [\Gamma, \Gamma]$ with $a, b, c \in S$.

However, for the case of a nilpotent group, we can give a finite system of generators:

Proposition 5.2. *Let $\Gamma = \langle S \rangle$ be a group. Then, the commutator subgroup $[\Gamma, \Gamma]$ is generated by all nested commutators of the form $[\ldots[[s_1, s_2], s_2]\ldots, s_n]$ with $s_i \in S$. This set is finite if the group is finitely generated and nilpotent.*

Proof. Any commutator can be represented by a word with an equal number of letters s and s^{-1}. Example:

$$[ab, c^{-1}] = abc^{-1}b^{-1}a^{-1}c.$$

Consider the rightmost letter c and use the commutator identity $xc = cx[x^{-1}, c^{-1}]$ to move it to left until it cancels with its corresponding inverse c^{-1}:

$$abc^{-1}b^{-1}a^{-1}\underline{c} = abc^{-1}b^{-1}\underline{c}a^{-1}[a, c^{-1}] = abc^{-1}\underline{c}b^{-1}[b, c^{-1}]a^{-1}[a, c^{-1}] = abb^{-1}[b, c^{-1}]a^{-1}[a, c^{-1}].$$

Proceed with the second leftmost letter from the original word in the same fashion. This time, we will obtain nested commutators:

$$\begin{aligned} abb^{-1}[b,c^{-1}]\underline{a^{-1}}[a,c^{-1}] &= abb^{-1}\underline{a^{-1}}[[b,c^{-1}]^{-1},a][a,c^{-1}] = abb^{-1}\underline{a^{-1}}[[c^{-1},b],a][a,c^{-1}] \\ &= [[c^{-1},b],a][a,c^{-1}]. \end{aligned}$$

And so on, until all letters have been cancelled and only nested commutators remain. □

Exercise. Let Γ be an s-step nilpotent group. Show that the group Γ_n of the lower central series is $(s-n)$-step nilpotent.

Exercise. Every subgroup of a nilpotent group is nilpotent.

One of the most celebrated results of geometric group theory is that nilpotent groups are characterized by their growth function. A group Γ is called "virtually nilpotent" if it is virtually isomorphic to a nilpotent group.

Theorem 5.3. *Every finitely generated nilpotent group Γ has polynomial growth.*

Theorem 5.4 (Gromov)**.** *Any group Γ of polynomial growth is virtually nilpotent.*

The growth of nilpotent groups can be calculated precisely. Let Γ be a nilpotent group with lower central series $\Gamma \trianglerighteq \Gamma_1 \trianglerighteq \ldots \trianglerighteq \Gamma_s = \{1\}$. The factor groups Γ_k/Γ_{k+1} are finitely generated abelian groups, so let $\mathrm{rk}(\Gamma_k/\Gamma_{k+1})$ denote their free rank.

Theorem 5.5 (Bass)**.** *Every nilpotent group Γ has growth $V(n) \sim n^{d(\Gamma)}$ where the degree of the polynomial is given by*

$$d(\Gamma) = \sum_{k=1}^{s} k \cdot \mathrm{rk}(\Gamma_{k-1}/\Gamma_k).$$

Together with Gromov's classification theorem, this calculation be used to classify the recurrence of random walks on groups:

Theorem 5.6. *Let Γ be a finitely generated, infinite group. Then, the random walk on Γ is recurrent if and only if the group contains $\mathbb{Z}$ or $\mathbb{Z}^2$ as a finite index subgroup.*

For more details on this, see [Woe00], section 3.B.

We won't prove Gromov's theorem here, but we can at least show that every nilpotent group has polynomial growth.

Write $l_S(g) = d_S(1,g)$ for the distance of a group element g from the unit element with respect to some generating set S. Enlarging the set of generators will only make it easier to reach group elements, hence $l_S(g) \geq l_{S'}(g)$ for any two generating sets $S \subset S'$.

Definition 5.7. *Let Γ be a finitely generated group, $H \leq \Gamma$ a subgroup, and $S_H \subset S_\Gamma$ corresponding finite generating sets. This pair is said to have* polynomial distortion *if we can bound $l_H(h) \leq p(l_\Gamma(h))$ by a polynomial $p(n) = a_d n^d + \ldots$. In other words, the adding generators from Γ will shorten distances only by a polynomial amount.*

Lemma 5.8. *Let Γ be a finitely generated nilpotent group. Then, the normal subgroup $[\Gamma,\Gamma]$ has polynomial distortion.*

Proof. Polynomial distortion is a property that is invariant under quasi-isometry, so we can restrict our attention to a particular choice of generating systems.

Write $H = [\Gamma, \Gamma]$ and consider the quotient map

$$\pi : \Gamma \longrightarrow \Gamma/H \cong \Gamma^{Ab}.$$

As a quotient of the finitely generated group Γ, the abelian group Γ^{Ab} is finitely generated. By the structure theorem for abelian groups, it has the form

$$\Gamma^{Ab} \cong \mathbb{Z}^r \oplus \mathbb{Z}/d_1\mathbb{Z} \oplus \ldots \oplus \mathbb{Z}/d_n\mathbb{Z}.$$

Let $e_1, \ldots, e_r$ be the free generators in Γ^{Ab} and $t_1, \ldots, t_r$ corresponding preimages in the group Γ. To prove polynomial distortion, it is enough to consider the finite index subgroup $\Gamma' = \langle t_1, \ldots, t_r, H\rangle$ generated by the t_i and the kernel.

Let $h \in H$ be an element of the subgroup and let w be a corresponding word of minimal length in the generators of Γ'. For example,

$$w = t_1 h_1 t_1^{-1} h_2 t_2 t_2 \ldots.$$

Projecting to Γ^{Ab}, we see that this word must have an equal number of letters t_i and their inverses t_i^{-1}. We can remove these letters in favor of generators from H by apply the same procedure that we used to show that the commutator subgroup is finitely generated (Proposition 5.2). This time, however, we have to keep track of the number of nested commutators, and show that only polynomially many are added. This would show $l_H(h) \leq p(l_{\Gamma'}(h))$ as desired.

It is convenient to enlarge the generating system of the subgroup H so that it also contains all the nested commutators $[\ldots[[s, t_i], t_j]\ldots, t_k] \in S_H$ for each generator $s \in S_H$. Since the group is nilpotent, only finitely many elements of this form need to be added.

Now, we need to keep track of the number of nested commutators that arise when moving a letter t_i to the left to cancel its inverse t_i^{-1}. For a word w, let

$$l_k(w) = \#\{\text{letters of } w \text{ that are elements of } \Gamma_k \setminus \Gamma_{k-1}\}.$$

Clearly, $l_\Gamma(w) = \sum_{k=0}^{s} l_k(w)$. Whenever we move the generator one step to the left,

$$x t_i = t_i x [x^{-1}, t_i^{-1}],$$

we incur one additional letter, but it is a commutator and hence "one level lower". Thus, if w' denotes the word obtained from w by moving the rightmost generator t_i to the left, we have

$$l_k(w') \leq \begin{cases} l_0(w) - 2 & \text{if } k = 0 \\ l_k(w) + l_{k-1}(w) & \text{if } 1 \leq k \leq s \\ 0 & \text{else} \end{cases}.$$

Repeating this process several times, we get a sequence of words $w = w_0 \to w_1 \to w_2 \to \ldots \to w_n$ which starts at the original word w and ends after n steps at a word in the commutator subgroup $w_n \in [\Gamma, \Gamma]$, where $n = l_0(w)$ is the number of generators t_i occuring in the original word. We can establish the estimates

$$l_k(w_j) \leq j n^k \text{ for } 1 \leq k \leq s$$
$$l_k(w_j) \leq n^{k+1} \text{ for all } k$$

by induction, using the induction step

$$l_k(w_{j+1}) \leq l_k(w_j) + l_{k-1}(w_j) \leq jn^k + n^k \leq (j+1)n^k \text{ for } 1 \leq k.$$

In other words, the number of letters for each kind is bounded by a polynomial in n, and the sum over the s different kinds will be bounded by a polynomial in n as well. □

Proof of Theorem 5.3. We use induction on the number of steps s of the nilpotent group. The 0-step nilpotent group is the trivial group, which clearly has polynomial growth.

Given an s-step nilpotent group Γ, we proceed similarly to the previous proof. As before, let $H = [\Gamma, \Gamma]$ be the commutator subgroup, let $t_1, \ldots, t_r$ be generators that are free abelian in the quotient Γ/H, and focus on the group $\Gamma' = \langle t_1, \ldots, t_r, H \rangle$ that is quasi-isometric to the group Γ.

Consider a group element $g \in \Gamma'$ from a ball of radius R, i.e. $l_{\Gamma'}(g) \leq R$. Similar to the previous proof, we can represent the element by a word of minimal length and move all the generators t_j to the left, obtaining a representation

$$g = t_1^{k_1} t_2^{k_2} \ldots t_r^{k_r} h, \quad \text{with } h \in H.$$

Since the number of letters t_i never increases while moving, we have

$$|k_1| + |k_2| + \cdots + |k_r| \leq l_{\Gamma'}(g) = R.$$

We don't keep track of the length of the remainder h in all detail, we just estimate it using the triangle inequality as $l_{\Gamma'}(h) \leq 2R$.

Now, the commutator subgroup $H = [\Gamma, \Gamma]$ is $(s-1)$-step nilpotent, so we can apply the induction hypothesis. Together with polyomial distortion, we obtain that there are only $q(2R)$ many elements h that have length $l_{\Gamma'}(h) \leq 2R$ where q is a polynomial. This gives a polynomial count of at most $R^r q(2R)$ distinct possible elements g inside the ball of radius R. □

References

[DK09] Cornelia Drutu and Michael Kapovich. *Lectures on Geometric Group Theory*. Dec. 2009. URL: https://www.math.ucdavis.edu/~kapovich/EPR/ggt.pdf.

[Pet13] Gábor Pete. *Probability and Geometry on Groups*. Oct. 2013. URL: http://www.math.bme.hu/~gabor/PGG.html.

[Woe00] Wolfgang Woess. *Random Walks on Infinite Graphs and Groups*. Cambridge University Press, 2000. ISBN: 0 521 55292 3.