

Infinite Groups

Notes from Miklos Abert's course in the
University of Chicago 2009 REU
June 22 – July 17

Monday, June 22

We are going to work towards solving world hunger by proving the Banach–Tarski paradox: it is possible to cut an orange into finitely many pieces and then move these pieces around with rigid motions of 3-space so that the result is two oranges each identical to the original. However, this doesn't work for pancakes: you can't double 1- or 2-dimensional food; it only works in 3-dimensions.

Some group theory

It turns out this has to do with group theory, so we must first discuss “What is a group?”

Definition 1. A *group* is a set X with a binary operation, denoted (X, \cdot) , satisfying the axioms:

- $a \cdot (b \cdot c) = (a \cdot b) \cdot c$ for all $a, b, c \in X$.
- There exists an identity element $e \in X$ so that $e \cdot a = a = a \cdot e$ for all $a \in X$.
- For every $a \in X$ there is an inverse, denoted a^{-1} , such that $a \cdot a^{-1} = e$.

Most interesting groups arise as a group of transformations. For example the set of “rigid motions” of 2-dimensional space \mathbb{R}^2 forms a group denoted $\text{Isom}(\mathbb{R}^2)$. There are three kinds of transformations: reflection about a line, rotation about a point, and translation in a direction. These form a group under composition (i.e., do one transformation and then do the other). For example, the composition of two reflections is either a rotation (if the two lines intersect) or a translation (if the two lines are parallel). If the two lines intersect, notice that the direction of rotation depends on which reflection you do first. Thus this group is *non-commutative*, i.e., $a \cdot b \neq b \cdot a$.

When are two groups the same? To address this, we must first introduce the

Definition 2. A *group homomorphism* between two groups G and H is a map (or function) $f: G \rightarrow H$ such that for all $a, b \in G$ one has

$$f(a \cdot b^{-1}) = (f(a)) \cdot ((f(b))^{-1}).$$

That is, a homomorphism is simply a map which respects the group operations.

Example 3. Some examples of group homomorphisms are:

- The identity homomorphism $\text{Id}: G \rightarrow G$.
- The trivial homomorphism $G \rightarrow \{1\}$.
- The “mod n ” homomorphism $\mathbb{Z} \rightarrow \mathbb{Z}/n\mathbb{Z} = \{0, 1, \dots, n-1\}$.
- The map $\text{Isom}(\mathbb{R}^2) \rightarrow \{0, 1\}$ which sends a transformation of \mathbb{R}^2 to either 0 or 1 depending on whether it is orientation preserving or not. (To see that this is a group homomorphism one needs to check that the composition of two orientation-reversing transformations is orientation-preserving, that the composition of an orientation-reversing transformation with an orientation-preserving one is orientation-reversing, etc ...)
- Similarly, the “sign” homomorphism $\sigma: \text{Sym}_n \rightarrow \mathbb{Z}/2\mathbb{Z}$, where Sym_n is the *symmetric group* (that is, the set of all bijections of the set $\{1, 2, \dots, n\}$ with group operation given by function composition) and σ is defined as

$$\sigma(f) = \left| \{(i, j) \mid i < j \text{ and } f(i) > f(j)\} \right| \bmod 2.$$

Exercise 1. Prove that $\sigma: \text{Sym}_n \rightarrow \mathbb{Z}/2\mathbb{Z}$ is a group homomorphism.

We may now say what it means for two groups to be the same:

Definition 4. Two groups G and H are *isomorphic* if there exists a bijective homomorphism $f: G \rightarrow H$.

Note that it is insufficient to merely require two homomorphisms $G \rightarrow H$ and $H \rightarrow G$ (or two injective or two surjective homomorphisms). However, this would suffice for sets:

Theorem 5 (Bernstein–Schröder). *If A and B are sets and there exist injective maps $A \rightarrow B$ and $B \rightarrow A$, then there exists a bijection $A \rightarrow B$.*

Some measure theory

Here are some more exercises. Let $S^1 = \{z \in \mathbb{C} : |z| = 1\}$ be the unit circle and let α be an irrational rotation (that is, $\alpha(z) = e^{2\pi i\theta}z$ for some fixed irrational number $\theta \in \mathbb{R}$). Show the following:

Exercise 2. Every orbit is dense, that is, for any $z \in S^1$, the set $\{z, \alpha(z), \alpha^2(z), \dots\}$ is dense in S^1 .

Exercise 3. If $A \subseteq S^1$ is measurable and invariant under α (i.e. $\alpha(A) \subseteq A$), then either $\lambda(A) = 0$ or $\lambda(S^1 \setminus A) = 0$ (here λ denotes Lebesgue measure). This means that the action of α is *ergodic*.

Exercise 4. For all $\varepsilon > 0$ there exists a measurable set $A \subseteq S^1$ such that $\lambda(A) > \frac{1}{2} - \varepsilon$ and $A \cap A^\alpha = \emptyset$ (here $A^\alpha = \alpha(A) = \{\alpha(z) : z \in A\}$).

Exercise 5. There does not exist any measurable subset $A \subseteq S^1$ for which $\lambda(A) = 1/2$ and $A \cap A^\alpha = \emptyset$.

We must be careful to only talk about *measurable* subsets in the above exercises because non-measurable sets do exist. For example, one may define an equivalence relation on S^1 by setting $x \sim y$ if $y = \alpha^i(x)$ for some $i \in \mathbb{Z}$. If $X \subseteq S^1$ is a set consisting of exactly one point from each equivalence class, then X cannot be measurable (note that it requires the axiom of choice to construct such an X).

There are two standard ways to avoid such pathologies: One may restrict ones attention to *Borel sets* (those sets which can be obtained from the open sets by repeatedly taking countable unions, intersections, and compliments) and only consider *Borel measures* (i.e., measures which are well-defined on the Borel sets). Or one may look at *means* instead of measures (a mean is like a measure, except it is only finitely additive).

Theorem 6. *There exists a mean on (all subsets of) \mathbb{R}^2 which is invariant under Euclidean transformations.*

Here are two fun problems:

1. Take two ellipses, one inside the other. Pick a starting point on the outside ellipse, go along a straight line which is tangent to the inner ellipse and stop when you hit the outer ellipse again. Repeat this procedure and assume that you get back where you started after n steps. If this is true, then no matter where you start you will return there after n steps.
2. Consider a disk of diameter 9 (e.g., a wheel of cheese). Your goal is to cover the disk by finitely many infinite strips of width 1 (e.g., long logs) so as to hide it from some pesky birds. Show that this cannot be done with less than 9 strips.

Exercise 6. Fact: There is an invariant mean on (all subsets of) \mathbb{Z} . Try to find one.

The free group

Let $\mathcal{W}(a, b)$ be the set of all finite-length words in the letters $\{a, a^{-1}, b, b^{-1}\}$. A word in $\mathcal{W}(a, b)$ is reduced if no letter is followed by its inverse (i.e., you're not allowed to write aa^{-1} or bb^{-1} etc ...). Two words are equivalent if you can get from one to the other by inserting or deleting strings such as aa^{-1} , $b^{-1}b$, etc.

Exercise 7. There is a unique reduced word in each equivalence class.

We define the *free group* $F(a, b)$ to be the set of equivalence classes in $\mathcal{W}(a, b)$ with operation given by concatenation.

Tuesday, June 23

Let's address some of the exercises from last time:

Solution to Exercise 2. The set $B = \{z, \alpha(z), \alpha^2(z), \dots\}$ is an infinite subset of S^1 . Thus there exist points $p, q \in B$ that are within ε of each other. It follows that B contains an ε -net for any $\varepsilon > 0$. This is only possible if B is dense. \square

More measure theory

Before tackling more exercises, we should first discuss some basics about measures. Let X be a set. A σ -algebra on X is a collection of subsets of X (i.e., a subset $P \subset \mathcal{P}(X)$) that contains the empty set \emptyset and is closed under countable unions and complements. Those sets in the σ -algebra are called *measurable*.

Definition 7. A *measure* is a function $\mu : P \rightarrow [0, \infty]$ which satisfies:

- σ -additivity: if $A_1, A_2, \dots \in P$ are disjoint, then $\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i)$.
- $\mu(\emptyset) = 0$.
- If $A \subseteq B$ are measurable, then $\mu(A) \leq \mu(B)$.

There are many examples of measures. The most natural/geometric is the *Lebesgue measure* on \mathbb{R}^n , this is built by declaring that the unit cube $[0, 1]^n \subset \mathbb{R}^n$ has measure 1 and extending this to other sets in the most natural way — basically, Lebesgue measure agrees with how big you think a set should be.

Another good example of a measure is the Dirac measure δ_{x_0} concentrated at a point $x_0 \in X$. This is defined by setting $\delta_{x_0}(A) = 1$ if $x_0 \in A$ and $\delta_{x_0}(A) = 0$ otherwise.

Definition 8. A *mean* is a similar thing as a measure, but one dispenses with the σ -algebra and only requires finite additivity (rather than σ -additivity).

An important theorem about measures is

Theorem 9. Let μ be a nice measure on a space X (e.g., Lebesgue measure on \mathbb{R}^n or S^1) and let $A \subseteq X$ be any measurable set. Then for any $\varepsilon > 0$ there is an open set $\mathcal{O} \supseteq A$ such that $\mu(A) \leq \mu(\mathcal{O}) < \mu(A) + \varepsilon$.

Example 10. The rational points in $(0, 1)$ have measure 0 (countable sets have measure 0). Think about how to construct an open neighborhood of this set which has measure $< \varepsilon$.

We can use this theorem to prove the ergodicity exercise from last time

Solution to Exercise 3. For an interval $I \subseteq S^1$ define the *relative density* of a subset $A \subseteq S^1$ to be $d(A, I) = \frac{\lambda(A \cap I)}{\lambda(I)}$. For an invariant subset A , the relative density is a constant c independent of the interval I . The measure of A may be approximated by an open set \mathcal{O} which is a countable union of intervals. By σ -additivity we know that $\lambda(\mathcal{O})c = \lambda(A)$ which, by varying \mathcal{O} , shows that $\lambda(A)$ is either 0 or 1. \square

For a measurable set $A \subseteq S^1$ and a point $x \in S^1$ we may alternately define the density of A at x by

$$d_\lambda(A, x) = \lim_{\varepsilon \rightarrow 0} \frac{\lambda(A \cap (x - \varepsilon, x + \varepsilon))}{2\varepsilon}.$$

Since any invariant subset must have the same density at each point, the above exercise is also a consequence of the important

Theorem 11 (Lebesgue density theorem). *If A is measurable, then $d_\lambda(A, x) = 1$ for almost every point $x \in A$.*

Now we'll construct an invariant mean on \mathbb{Z} (c.f. Exercise 6). The natural candidate is the density

$$D(A) = \lim_{n \rightarrow \infty} \frac{|[-n, n] \cap A|}{2n + 1},$$

but the problem is that many sets don't have a density (i.e., the above limit needn't exist). To get around this difficulty, we need to introduce ultralimits.

Definition 12. An *ultrafilter* on \mathbb{N} is a subset $\mathcal{U} \subset \mathbb{N}$ such that

- $\emptyset \notin \mathcal{U}$.
- If $A \subseteq B$ and $A \in \mathcal{U}$, then $B \in \mathcal{U}$.
- If $A, B \in \mathcal{U}$, then $A \cap B \in \mathcal{U}$.
- For all $A \subseteq X$, either $A \in \mathcal{U}$ or $X \setminus A \in \mathcal{U}$ (this is the crazy axiom which cannot be satisfied without the axiom of choice).

Said another way, an ultrafilter is a mean on \mathbb{N} such that every subset of \mathbb{N} gets assigned either the value 0 (the small sets $A \notin \mathcal{U}$) or the value 1 (the large sets $B \in \mathcal{U}$).

People have only ever seen the trivial (principle) ultrafilters $\mathcal{U}_n = \{A \subseteq \mathbb{N} : n \in A\}$, which only depend on a point $n \in \mathbb{N}$. However, it is easy to construct nontrivial *filters* (collections which satisfy only the first three axioms above). For example, the collection of co-finite sets (those sets with finite compliment) is a perfectly good filter but fails to be an ultrafilter. Moreover, any filter can be extended to an ultrafilter by using Zorn's lemma (every filter is contained in a maximal filter which must necessarily be an ultrafilter).

Theorem 13 (Ax). *If $f: \mathbb{C}^n \rightarrow \mathbb{C}^n$ is a polynomial map which is injective, then it is also surjective.*

Using ultrafilters we can do crazy things, such as take the limit of any bounded sequence:

Definition 14. Let \mathcal{U} be an ultrafilter on \mathbb{N} , and let $\{a_n\}$ be a bounded sequence in \mathbb{R} . We say that the *ultralimit* of $\{a_n\}$ (with respect to \mathcal{U}) is the number a , which we denote by

$$\lim_{\mathcal{U}} a_n = a,$$

if each neighborhood V of a contains a \mathcal{U} -majority of the a_n (i.e., if $\{i \in \mathbb{N} \mid a_i \in V\} \in \mathcal{U}$).

Notice that if the ultra limit exists it is unique (\mathbb{R} is Hausdorff!). Moreover, the ultralimit always exists: if not we could cover the compact interval with finitely many open sets all of which contain a \mathcal{U} -minority of the a_n but whose union contains all of the a_n . Alternately one may zero-in on the ultralimit by repeatedly cutting the compact interval in half (exactly one of the halves must contain a \mathcal{U} -majority of the a_n).

Standard things work with ultralimits: the ultralimit of a sum (product) is the sum (product) of the ultralimits. One can change finitely many values of a sequence without changing its ultralimit. Also, the ultralimit of a convergent sequence is same as the normal limit (provided the ultrafilter is non-principle). However, removing finitely many terms from the sequence CAN change the ultralimit.

Now it is easy to construct an invariant mean on \mathbb{Z} : Let \mathcal{U} be a non-principle ultrafilter on \mathbb{N} . Then for any subset $A \subseteq \mathbb{Z}$ let

$$s_n(A) = \frac{|[-n, n] \cap A|}{2n + 1} \quad \text{and define} \quad \mu(A) = \lim_{\mathcal{U}} s_n(A).$$

One easily checks that this gives an mean. To see that it is invariant it suffices to just translate with 1 and observe that

$$|s_n(A) - s_n(A + 1)| \leq \frac{2}{2n + 1},$$

which implies that these sequences have the same ultralimit. One could also construct an invariant mean by using the Hahn–Banach theorem from functional analysis; this essentially amounts to the same thing.

Definition 15. A countable group Γ is *amenable* if there exists a Γ -invariant mean on Γ . Here Γ -invariant means that $\mu(\gamma A) = \mu(A)$ for all $\gamma \in \Gamma$ and $A \subseteq \Gamma$.

There is a nice characterization of amenable groups in terms of random walks, which we will discuss later. Amenable groups are relevant to our primary goal of doubling food: The reason you can double oranges is because $\text{Isom}(S^2)$ is not amenable, and the reason you cannot double pancakes is because $\text{Isom}(R^2)$ is amenable (cf. Theorem 6). Some basic examples of amenable groups are: \mathbb{Z} (by Exercise 6), any finite group, any product of amenable groups, or any quotient of an amenable group. In fact, all abelian groups are amenable. However, free groups are not amenable (meta-argument: if they were, then any quotient of a free group (i.e., any group) would be amenable).

Solution to Exercise 7. Consider the regular 4-valent tree with directed edges labeled by a and b . (this is the Cayley graph for $F(a, b)$). An arbitrary word in $\mathcal{W}(a, b)$ is the same thing as a path in the tree, and, since cancellation in $\mathcal{W}(a, b)$ corresponds to backtracking in the tree, we see that two equivalent words lead to the same point. Thus, the unique reduced word is the unique shortest path to that point (which clearly exists in the tree). \square

Using the perspective of this tree, one can see why $F(a, b)$ is not amenable. Indeed, for $x \in \{a, b, a^{-1}, b^{-1}\}$ let F_x denote the set of words in $F(a, b)$ which start with the letter x . Then we have $F_a \cup F_b \cup F_{a^{-1}} \cup F_{b^{-1}} = F(a, b) \setminus \{e\}$, but we may also write $F_a \cup a \cdot F_{a^{-1}} = F(a, b) = F_b \cup b \cdot F_{b^{-1}}$. Thus, any invariant mean must satisfy $\mu(F(a, b)) = 2\mu(F(a, b))$, which is absurd.

Our treatment of free groups has been rather abstract, but free groups exist in real life:

Theorem 16. *The matrices $A = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$ generate a free group.*

Proof. We need to prove that any freely reduced word in the letters A and B is not equal to the identity matrix $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$. Think about what the matrices A and B do to a vector $\begin{pmatrix} x \\ y \end{pmatrix}$. Notice that if the reduced word ends with A , then it cannot send the vector $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ to itself and therefore cannot be the identity. \square

A similar (but much uglier) proof shows that one can also find two rotations of S^2 which generate a free group.

Exercise 8 (Rokhlin's Lemma). Let $f: S^1 \rightarrow S^1$ be a measure preserving (i.e., $\mu(f(A)) = \mu(A)$) bijection such that f^n fixes no point on S^1 for $n \neq 0$. Then for all $\varepsilon > 0$ there exists some measurable $A \subseteq S^1$ such that $\mu(A) > \frac{1}{2} - \varepsilon$ and $A \cap f(A) = \emptyset$.

Thursday, June 25

The Banach–Tarski paradox

Today we will double the ball and prove the Banach–Tarski paradox. Interestingly, although pancakes can't be doubled, our approach is essentially 2-dimensional in that we will first double the sphere S^2 .

The proof relies on the free group F_2 . In Theorem 16 we said that the two matrices $A = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$ “generate a free group,” what we mean by this is the content of the following

Exercise 9. Explain why the group generated by A and B is isomorphic to the free group $F(a, b)$. (The group generated by A and B is a subgroup of the group of all invertible 2×2 matrices; see Definition 22 for more explanation.)

Definition 17. Two sets $A, B \subseteq S^1$ (or \mathbb{R}^n etc ...) are *congruent* if there exists a Euclidean transformation φ for which $\varphi(A) = B$. Furthermore, A and B are *equidecomposable*, denoted $A \sim B$, if we can write A and B as finite disjoint unions $A = \sqcup_{i=1}^n A_i$ and $B = \sqcup_{i=1}^n B_i$ such that A_i congruent to B_i for each $1 \leq i \leq n$ (here \sqcup is used to denote disjoint union).

With these definitions we can now state the

Theorem 18 (Banach–Tarski paradox). *If B^3 denotes the unit ball in \mathbb{R}^3 , then $B^3 \sim B^3 \sqcup B^3$.*

Definition 19. Let’s say that A fits into B , which we will denote by $A \subseteq B$, if there exists a subset $C \subseteq B$ for which $A \sim C$.

Exercise 10. Show that the circle S^1 fits into the circle minus a point $S^1 \setminus \{s_0\}$.

Hint: It is illuminating to first answer the question: Does the set $\{0, 1, 2, 3, \dots\}$ fit into $\{1, 2, 3, \dots\}$? (Here in the universe of \mathbb{Z} , the allowed, “Euclidean transformations” are just translations.) \square

As indicated by the notation, “equidecomposable” is an equivalence relation (one can check transitivity by ‘intersecting’ the decompositions). We will also make use of the fact that $A \subseteq B$ and $B \subseteq A$ implies $A \sim B$; however, before proving this we should first understand the proof of the Bernstein–Schröder theorem.

Proof of Theorem 5. Consider two injective maps $f: A \rightarrow B$ and $g: B \rightarrow A$ between disjoint sets A and B . Define a graph whose vertex set is $A \sqcup B$ and whose directed edges are the pairs $(a, f(a))$ and $(b, g(b))$. Since each vertex has one edge going out and at most one edge coming in, a connected component of this graph can either be: an even cycle, a bi-infinite line, or a ray (that is, a half-infinite line) — there are no other possibilities. It is clear how to build a bijection between the A - and B -vertices of a given connected component, and taking all of these bijections together proves the theorem. \square

This in fact proves something slightly stronger: there exist decompositions $A = A_1 \sqcup A_2$ and $B = B_1 \sqcup B_2$ such that f restricts to a bijection $A_1 \rightarrow B_1$ and g restricts to a bijection $B_2 \rightarrow A_2$. This formulation easily implies the following

Theorem 20. *If A fits into B and B fits into A , then A and B are equidecomposable.*

Proof. The hypotheses imply the existence of injective maps $f: A \rightarrow B$ and $g: B \rightarrow A$ which are built from finitely many Euclidean transformations. Therefore the bijection produced by Bernstein–Schröder actually consists of finitely many Euclidean transformations (a subset of those which comprise f and g), that is, the bijection is an equidecomposition. Alternately, one could use the partitions of A and B to build a graph as above. This graph then leads to an equidecomposition in a similar manner. \square

Recall the free group $F = F(a, b)$ and that F_x denotes the set elements in the free group which start with the letter x . We saw that $F_a \sqcup a \cdot F_{a^{-1}} = F = F_b \sqcup b \cdot F_{b^{-1}}$ and also that $F \setminus \{e\} = F_a \sqcup F_b \sqcup F_{a^{-1}} \sqcup F_{b^{-1}}$. Hence we have $F \setminus \{e\} \sim F \sqcup F$ (where in this context the allowed “Euclidean transformations” are left-multiplication by elements of F .)

Exercise 11. Show that $F \sim F \sqcup F$.

Exercise 12. If $C \subseteq S^2$ is a countable subset, then $S^2 \sim S^2 \setminus C$.

Solution. Pick an axis that doesn't hit C . The set of rotations α such that $C \cap C^\alpha \neq \emptyset$ is countable. Thus we can find a rotation such that $C \cap C^\alpha = \emptyset$ and use it to send the orbit of C into itself. \square

Consider the 2-sphere S^2 and choose rotations a and b of S^2 that generate a free group: $\langle a, b \rangle \cong F_2$. Since a product of rotations is a rotation, for any word $w \in F_2$ in the free group the transformation $w(a, b)$ has two fixed points. If $C = \{x \in S^2 : w(a, b) \cdot x = x \text{ for some } w \in F_2\}$ is the set of all fixed points, then F_2 acts *freely* on the set $Y = S^2 \setminus C$ (that is, $w_1 \neq w_2 \in F_2 \implies w_1(a, b)y \neq w_2(a, b)y$ for every $y \in Y$).

Theorem 21. $S^2 \setminus C = \left((S^2 \setminus C) \sqcup (S^2 \setminus C) \right)$

Proof. The *orbit* of a point $y \in Y = S^2 \setminus C$ is the set $\mathcal{O}_y = \{w(a, b) \cdot y : w \in F_2\}$. Choose exactly one point from each orbit and call this set X . Then $F_2 \cdot X = Y$ and we can use the equidecomposition $F_2 \sim F_2 \sqcup F_2$ to produce the equidecomposition between Y and $Y \sqcup Y$ (so, if A_i is a piece of the equidecomposition of F_2 , then use $A_i \cdot X$ as a piece of the equidecomposition of Y). \square

Exercise 13. Use this to double the sphere: $S^2 \sim S^2 \sqcup S^2$.

Since $B^3 \setminus \{0\} = S^2 \times (0, 1]$ we can now double $B^3 \setminus \{0\}$ by doing the *same* thing at each radius $r > 0$. With this it is easy to double B^3 , which completes the proof of the Banach–Tarski paradox.

Theorem 18 also lets you fit a big ball into a small ball: Cover the big ball by finitely many copies of the small ball and translate them all away from each other. Now the big ball is equidecomposable with a subset of these finitely many small balls which are collectively equidecomposable with the small ball. Since the small ball clearly fits into the big ball they are in fact equidecomposable.

This whole thing works because we can double the free group F_2 ; this is essentially like implementing a Ponzi scheme on F_2 — in fact, all non-amenable groups have Ponzi schemes.

Amenability

By definition, amenability means the existence of an invariant mean, but we will prove this is equivalent to the existence of finite subsets $A_n \subseteq \Gamma$ such that for all $g \in \Gamma$ we have $\lim_{n \rightarrow \infty} \frac{|(gA_n) \cap A_n|}{|A_n|} = 0$. We will also prove that for finitely-generated groups amenability is equivalent to random walks returning with large probability. Before we can discuss this, we need to say a few more things about groups.

Definition 22. Let Γ be a group. A subset $S \subseteq \Gamma$ *generates* Γ if no proper subgroup of Γ contains S , that is, if every element Γ can be obtained as a product of elements of S and their inverses; in this case we write $\langle S \rangle = \Gamma$.

If $\langle S \rangle = \Gamma$, then Γ is a homomorphic image of $F(S)$ (the free group in the letters of S). This is called a *presentation* of Γ . For example, every group which is generated by one element is a homomorphic image of $F(\{1\}) \cong \mathbb{Z}$ and is therefore isomorphic to either \mathbb{Z} or $\mathbb{Z}/n\mathbb{Z}$.

Claim 23. *The matrices $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ generate $\mathrm{SL}_2 \mathbb{Z} = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} : ad - bc = 1, a, b, c, d \in \mathbb{Z} \right\}$.*

To see this, use multiplication by $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ together with the Euclidean algorithm to get any matrix into the form $\begin{pmatrix} 1 & 0 \\ * & * \end{pmatrix}$, which is clearly in the group $\langle \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \rangle$.

Take a finitely-generated group Γ with the finite and symmetric (i.e., $S = S^{-1}$) generating set S . To this we can associate the *Cayley graph* $\mathrm{Cay}(\Gamma, S)$. The vertex set of $\mathrm{Cay}(\Gamma, S)$ is just the group Γ , and for every $x \in \Gamma$ and every $s \in S$ there is an edge (x, xs) which is labeled by s . Notice that the Cayley graph is connected because S generates Γ .

For example, the bi-infinite line is the Cayley graph of \mathbb{Z} with generating set $\{1\}$, the infinite 4-regular tree is the Cayley graph of the free group $F_2 = \langle a, b \rangle$, and the infinite grid is the Cayley graph of $\mathbb{Z}^2 = \langle \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \rangle$. Notice that all Cayley graphs are *homogeneous*, that is, they look the same at each vertex.

A *random walk* on a Cayley graph is obtained by starting at a vertex and crossing a neighboring edge at random at each stage. The cardinal question about random walks is: what is the probability that a random walk returns to its starting point? If a random walk will return with probability 1, then the graph is called *recurrent*.

Fact 24. The Cayley graph of \mathbb{Z}^2 is recurrent (i.e., every random walk on \mathbb{Z}^2 returns with probability 1), whereas the Cayley graph of \mathbb{Z}^3 is not recurrent. In fact, a group is recurrent if and only if it has a finite index subgroup which is either trivial, \mathbb{Z} , or \mathbb{Z}^2 .

Let P_{2n} be the probability that a random walk returns to the identity after $2n$ steps. We will see that amenable groups are precisely those groups for which

$$\lim_{n \rightarrow \infty} \sqrt[2n]{P_{2n}} = 1.$$