Graphs and Groups

Rich Schwartz

April 19, 2020

1 Definition of a Group

A *group* is a set G together with a binary operation, $*$, which satisfies the following axioms:

- For all $g, h \in G$ the element $g * h$ is defined and lies in G.
- $(a * b) * c = a * (b * c)$ for all $a, b, c \in G$.
- There exists a (unique) $e \in G$ such that $g * e = e * g = g$ for all $g \in G$. The element e is called *the identity*.
- For all $g \in G$ there exists a (unique) $h \in G$ such that $g * h = h * g = e$. The element h is called *the inverse of* g and is often written $h = g^{-1}$.

Note: It is not always true in a group that $g*h = h*g$. When this is always true, the group is called *Abelian*.

Here are some examples of groups.

- The group with one element, e. The rule is $e * e = e$.
- G is any of $\boldsymbol{Z}$, $\boldsymbol{Q}$, $\boldsymbol{R}$, $\boldsymbol{C}$ and $*$ is addition. Here $e = 0$ and g^{-1} is usually written $-g$.
- G is any of $\boldsymbol{Q} - \{0\}$, $\boldsymbol{R} - \{0\}$ or $\boldsymbol{C} - \{0\}$ and $*$ is multiplication. Here $e = 1$ and g^{-1} is usually written $1/g$.
- $\boldsymbol{Z}/n$, the group of residue classes of integers mod n. Here $*$ is addition mod n.

- Let X be any set and let S_X denote the set of bijections from X to itself. In this case $*$ is composition, and e is the identity map, and g^{-1} is the inverse of g in the sense of mappings. When $X = \{1, ..., n\}$ the group S_X is often denoted S_n and called *the permutation group*.

- Whenever G_1 and G_2 are groups, the product $G_1 \times G_2$ is a group in a natural way. The two group laws are done coordinatewise.

- The set of rigid motions of Euclidean space forms a group. The group law is composition. This group is known as the *Euclidean group*.

- The same goes for the group of rotations of the sphere. This group is called the *rotation group*, at least when the dimension of the sphere is specified in advance.

A *subgroup* H of a group G is a subset which is also a group, with respect to the same operation. Here are some examples of subgroups.

- The trivial group is a subgroup of every group.

- The set of powers of a single element $g \in G$, namely

  $$g, g * g, g * g * g, ...$$

  is a subgroup of any finite group. Such a group is known as a *cyclic subgroup*. It is not unlike a cycle in a graph.

- The set $2\boldsymbol{Z}$ of even integers is a subgroup of $\boldsymbol{Z}$.

- When $X \subset Y$, there is a natural way to think of S_X as a subgroup of S_Y.

- Each element g of S_n can be considered as a linear transformation $T_g : \boldsymbol{R}^n \to \boldsymbol{R}^n$. The map T_g sends the standard basis element e_j to the basis element $e_{g(j)}$. That is,

  $$T_g\Big(\sum a_j e_j\Big) = \sum a_j e_{g(j)}.$$

  It turns out that $\det(T_g) = \pm 1$. The element g is called *even* if $\det(T_g) = 1$. The set of even permutations of S_n is a subgroup. It is denoted by A_n. It has $n!/2$ elements.

2 Group Isomorphisms

Here is a little more information about group theory. An *isomorphism* between groups G_1 and G_2 is a bijection $f : G_1 \to G_2$ which respects the group laws. That is,

$$f(a * b) = f(a) * f(b), \qquad \forall\, a, b \in G_1.$$

Here the first instance of $*$ is the group law in G_1 and the second instance is the group law in G_2. When two groups are isomorphic, they are essentially the same group, except that the elements have different names.

Here is a nice example. Let G_1 be the group of nonzero real numbers, with respect to multiplication. Let G_2 be the real numbers with respect to addition. Let $f(x) = \log(x)$. The famous equation

$$\log(ab) = \log(a) + \log(b)$$

really says that the log function is an isomorphism between the two groups.

Theorem 2.1 (Cayley) *Let G be a group. G is isomorphic to a subgroup of S_G.*

Proof: Given an element $g \in G$ we want to cook up a permutation $T = f(g)$ of G (as a set). Here is the formula

$$T(h) = g * h.$$

The map T is a bijection because the inverse map is given by

$$T^{-1}(h) = g^{-1} * h.$$

The associative law shows that

$$f(g) \circ f(h) = f(g * h).$$

Hence, our map $f : G \to S_G$ is an isomorphism onto the image $f(G)$, which is a subgroup. ♠

So, if you don't like abstract groups, you can think of a group as a subgroup of the group of bijections of some set. This is often not the best way to think of a group, but sometimes it is useful.

3 Automorphisms of Graphs

Let Γ be a graph. An *automorphism* of Γ is a graph isomorphism from Γ to itself. In other words, an automorphism of Γ is a permutation of the vertices of Γ which maps incident vertices to incident vertices. If you take two automorphisms of Γ, you can compose them and get a third. Thus, the set of automorphisms of Γ forms a group, denoted $A(\Gamma)$. The identity element is the automorphism which maps every vertex to itself. The inverse of an automorphism is just the inverse permutation.

Note that $A(\Gamma)$ automatically is defined as a subgroup of a permutation group – though if you were to run Cayley's Theorem you'd get a different permutaton group! To be even more concrete, suppose Γ is a finite graph. If you label the vertices of Γ as $1, ..., n$ then every element of $A(\Gamma)$ is naturally an element of S_n. That is $A(\Gamma) \subset S_n$. (Note that Cayley's theorem would give $A(\Gamma) \subset S_N$ where N is the order of $A(\Gamma)$. Often $N \neq n$.) Here are some examples.

- The automorphism group of the graph with one vertex is the trivial group.
- The automorphism group of K_n is S_n.
- The automorphism group of any nontrivial path is $\boldsymbol{Z}/2$. All you can do is reverse the direction of the path, so to speak.
- The automorphism group of the cycle C_n has $2n$ elements. Half the elements are cyclic permutations of the vertices and the other half are "reflections". This group is often denoted D_n, and it is an example of a *dihedral group*.
- The automorphism group of the Peterson graph is S_5. Here is a nice way to think about it. You can make a graph whose vertices are 2 element subsets of $\{1, 2, 3, 4, 5\}$ and whose edges join disjoint subsets. For instance $\{1, 2\}$ is joined to each of $\{3, 4\}$ and $\{4, 5\}$ and $\{3, 5\}$. This is a graph with 10 vertices in which every vertex has degree 3. If you draw it, you will see that it is the Peterson graph. Any permutation of $\{1, 2, 3, 4, 5\}$ naturally gives a permutation of the graph.
- Make the graph whose edges are the seams on a soccer ball. The group of automorphisms of the graph which can be achieved by rotations of the soccer ball is A_5.

4 Cayley Graphs

Let G be a group. A *generating set* for G is a subset $S \subset G$ such that every nontrivial element of G is a product of elements of S. That is

$$g = s_1 * ... * s_n.$$

Here n depends on g and so to $s_1, ..., s_n \in S$. Here are some examples.

- $G = \boldsymbol{Z}/n$ and $S = \{1\}$.
- $G = \boldsymbol{Z}$ and $S = \{\pm 1\}$.
- $G = \boldsymbol{Z}$ and $S = \{\pm 2, \pm 3\}$. This also works when S is any pair of relatively prime integers (and their negatives).
- $G = S_n$ and S is the set of transpositions. A *transposition* is a permutation which switches two elements and fixes the others.
- $G = A_5$ and S is the two permutations

  $$A = (12345), \qquad B = (12)(34)(5).$$

  This is the cycle notation for permutations. The permutation A cycles the indices by adding 1 mod 5. The permutation B swaps 1 and 2, swaps 3 and 4, and fixes 5. Note that these are both even permutations. It is a bit of work to show that $\{A, B\}$ generates A_5.

The *Cayley graph* $\Gamma(G, S)$ is a directed graph whose vertices are the elements of G and whose directed edges have the form

$$g \to g * s, \qquad \forall g \in G, \quad \forall s \in S.$$

When G is a finite group and $|S| = k$, every vertex of G has detree $2k$. However, there is a convention that when $s \in S$ has order 2, the two directed edges

$$g \to g * s, \qquad g * s \to g * s * s = g,$$

are collapsed into a single undirected edge. So, in the last example given above, the Cayley graph $\Gamma(A_5, \{A, B\})$ would have degree 3 with this convention.

Let's work out the Cayley graphs of the examples given above.

- $\Gamma(Z/n, \{1\})$ is C_n.
- $\Gamma(\boldsymbol{Z}, \{\pm 1\})$ is the bi-infinite path.
- It is most convenient to draw $\Gamma(\boldsymbol{Z}, \{\pm 2, \pm 3\})$ on an infinite grid, as suggested by Figure 1.

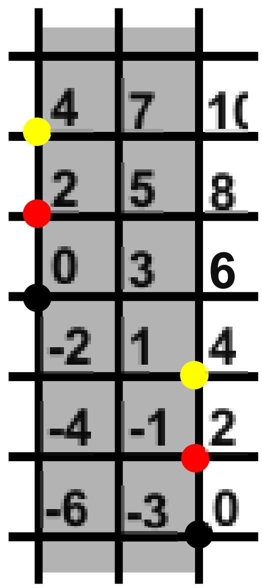

Figure 1: Numbers on a grid

To get the graph, cut out the infinite grey strip and glue the sides according to the rule suggested by the colored vertices. This gives you a kind of grid on an infinite cylinder. You can check that each integer appears exactly once on this cylindrical grid. Were you to use other relatively prime pairs of integers, you would get a similar kind of cylindrical grid, though typically the cylinder would be fatter.

- This next example is hard to draw in general, but Figure 2 shows what you get for the case $n = 3$. In this case S_3 has 6 elements. Using the convention above, Γ has degree 3. The edges are color-coded according to which element in S they represent.

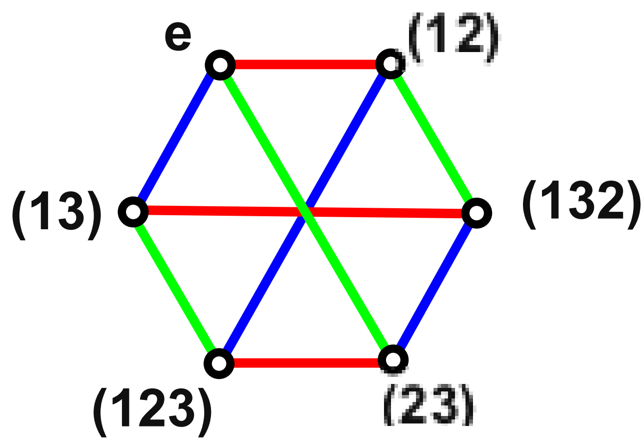

Figure 2: Numbers on a grid

- For the last example, the graph has 60 vertices and (with the convention) each vertex has degree 3. Note

$$A * A * A * A * A = e, \qquad A * B * A * B * A * B = e.$$

So, every vertex of the Cayley graph belongs to two 6-cycles and one 5-cycle, as indicated in Figure 3.

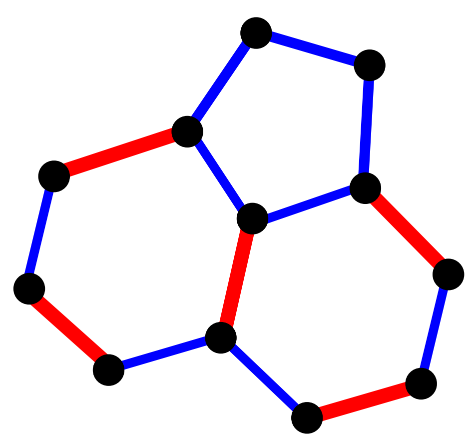

Figure 3: Part of a soccer ball

Each vertex is contained in a unique 5-cycle. Since there are 60 vertices, there are 12 5-cycles. Each of these 5-cycles is surrounded by 6-cycles. At the same time, each 6-cycle shares an edge with 3 5-cycles and 3 6-cycles. The only possible pattern is that of the soccer ball graph.

5 Quasi-Isometries

The Cayley graph $\Gamma(G,S)$ depends a lot on the choice of S. For instance, $\Gamma(G,G)$ is always the complete graph on $|G|$ vertices, at least if we follow the convention of collapsing directed 2-cycles into single edges. (This is a generalization of the convention mentioned above.) It turns out that the situation is a bit different for $\Gamma(G,S)$ when G is an infinite group and S is a finite set. In a sense, the Cayley graph does not depend much on S. This section will explain the sense in which this is meant.

Let X and Y be *metric spaces*. Given some $K \geq 1$, a K-*net* in X is a subset $S \subset X$ such that every point of X is within K units of a point of S. a K-*bi-lipchitz* map from if a map $f: X \to Y$ such that

$$(1/K)d_X(a,b) \leq d_Y(f(a),f(b)) \leq K d_X(a,b), \qquad \forall a,b \in X.$$

A K-*quasi-isometry* from X to Y is a K-bilipshitz bijection between a K-net X' of X and a K-net Y' of Y. Two spaces X and Y are said to be *quasi-isometric* if there is a K-quasi-isometry between them for some K. It turns out that this is an equivalence relation. So, is natural to consider metric spaces up to the equivalence of quasi-isometry. Here are some examples.

- Every compact metric space is quasi-isometric to a point. So, this equivalence relation is pretty boring for compact spaces.

- $\boldsymbol{Z}^n$ and $\boldsymbol{R}^n$ are quasi-isometric. You can take $K = \sqrt{n}+1$. Allowing the additive slop $\pm K$ in the definition allows two spaces to be quasi-isometric even when they look very different on small scales. When we look at tables and chairs and walls we usually think of them as continuous objects, but they are mostly empty space and very granular on small scales. So, our brains naturally do quasi-isometries.

- If K is any compact metric space and X is an arbitrary metric space, then X and $X \times K$ are quasi-isometric. For instance, the line and the infinite cylinder are quasi-isometric. This is also something our brain does. When we look at a telephone wire we might think of a line, but really it is a kind of tube. We are willing to forget a bounded amount of thickness.

- It turns out that the line and the ray are not quasi-isometric.

- It turns out that $\boldsymbol{R}^m$ and $\boldsymbol{R}^n$ are quasi-isometric if and only if $m = n$.
- When $k, \ell \geq 3$, the regular infinite tree of degree k is quasi-isometric to the regular infinite tree of degree ℓ. This is one of your HW problems in the case $k = 3$ and $\ell = 4$.

Any graph is naturally a metric space. We think of each edge as being a copy of the unit line segment, and then the distance between any two points in the graph is the length of the shortest path. For instance, the distance between adjacent vertices is 1, and the distance between the midpoint of an edge and either incident vertex is $1/2$. The graph is naturally a union of unit length segments.

Theorem 5.1 *Let G be an infinite group and let S, T be two finite generating sets. Then $\Gamma(G, S)$ and $\Gamma(G, T)$ are quasi-isometric.*

Proof: Let $\Gamma_S = \Gamma(G, S)$ and likewise define Γ_T. Let $f : \Gamma_S \to \Gamma_T$ be the identity on vertices, and arrange that f maps each edge to one of the two vertices incident to it. The choice doesn't matter. It suffices to prove that f is a quasi-isometry when restricted to the vertices, because every edge is at most 1 unit away from its endpoints.

Let $s_1, ..., s_m$ be the elements of S and let $t_1, ..., t_n$ be the elements of T. The distance between vertices a, b in Γ_S is the minimum number k such that $b^{-1}a = s_1 * ... * s_k$, for $s_1, ..., s_k \in S$. A similar statement applies to Γ_T. There is some K such that

- Each element of S can be written as the product of at most K elements of T.
- Each element of T can be written as the product of at most K elements of S.

This means that distances between vertices, relative to the two graphs, are the same up to a factor of K. Hence, on vertices, the map f is a K-quasi-isometry. As mentioned above, this is all we need to know. ♠

6 The Coarse Hyperbolic Plane

This infinite graph has the following description. Start with a vertical ray divided into segments of lengths $...4, 2, 1, 1/2, 1/4, ...$ Attach an infinite horizontal row of squares having side length 2^s to the edge of the vertical ray which has length 2^s. Figure 1 shows part of the picture. The picture is meant to extend in all directions. The union of squares fills the upper half plane.

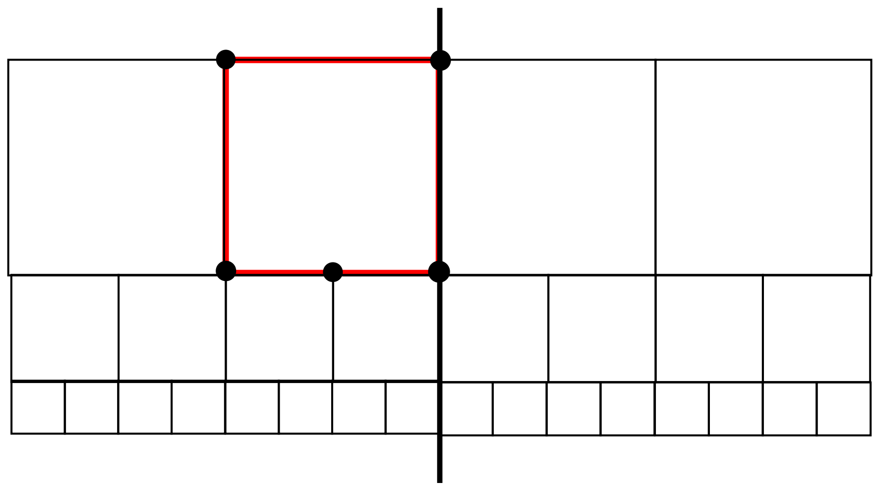

Figure 1: The coarse hyperbolic plane

This graph is an infinite union of 5-cycles. One of the 5 cycles is highlighted in the figure. This graph is sometimes called the *coarse hyperbolic plane*. We put a metric on the coarse hyperbolic plane by declaring that all its edges have length 1. The distance between any two points is defined to be the length of the shortest path which joins them. Let's compare the coarse hyperbolic plane to the hyperbolic plane.

The *hyperbolic plane* is the upper half plane equipped with a funny way of measuring distances. At the point (x, y) the inner product of two vectors V and W is defined by the formula

$$\langle V, W\rangle_{(x,y)} = \frac{1}{y^2}(V \cdot W). \tag{1}$$

In particular the langth of the vector V at (x, y) is

$$\|V\|_{(x,y)} = \frac{1}{y}\sqrt{V \cdot V}. \tag{2}$$

This way of measuring lengths of vectors is called the *hyperbolic metric*. The hyperbolic plane is denoted $\boldsymbol{H}^2$. The length of a parametrized curve $\alpha : [0,1] \to \boldsymbol{H}^2$ is defined to be

$$\text{length}(\alpha) = \int_0^1 \|\alpha'(t)\|_{\alpha(t)}\, dt. \tag{3}$$

The distance between two points in the hyperbolic plane is defined to be the infimum of the lengths of curves joining these points. In general, this kind of definition might not work to define a metric space, but in this case it does. The shortest curves joining two points in $\boldsymbol{H}^2$ are either vertical line segments or else arcs of circles which meet the real axis at right angles.

A map: $f : \boldsymbol{H}^2 \to \boldsymbol{H}^2$ is an *isometry* if it preserves the distances between points. It is easy to check that $f(x,y) = (x+t, y)$ and $g(x,y) = (rx, ry)$ are isometries for all choices of t and r. If you draw the coarse hyperbolic plane inside $\boldsymbol{H}^2$ then all the 5-cycles have the same size. Using the maps just described you can find an isometry from one 5-cycle to any other. This means that the edges of the graph, when drawn in $\boldsymbol{H}^2$, are all the same length to within some factor of K. Probably you can take $K = 2$.

This is the beginning of the proof that the coarse hyperbolic plane is quasi-isometric to $\boldsymbol{H}^2$. This is why it is called the coarse hyperbolic plane. There are about C^r distinct points inside the disk of radius r in the coarse hyperbolic plane. This is the beginning of the proof that the coarse hyperbolic plane is not quasi-isometric to the Euclidean plane.

7 The Heisenberg Graph

We introduce 3 symbols A, B, C and consider the group of words in these symbols, subject to the relations that

$$AC = CA, \quad BC = CB, \quad ABA^{-1}B^{-1} = C.$$

Here A^{-1} is such that $AA^{-1} = A^{-1}A = e$, the empty word. There are other relations implied by these. For instance

$$BAB^{-1}A^{-1} = C^{-1}, \qquad ABA^{-1}B^{-1}C^{-1} = e.$$

Two finite words are declared equivalent if there is a finite number of substitions of the relations above which brings the one word into the other. For instance

$$CBA \sim CBA(A^{-1}B^{-1}ABC^{-1}) \sim CABC^{-1} \sim CC^{-1}AB \sim AB.$$

In general, any word in the Heisenberg group is equivalent to

$$A^aB^bC^c, \qquad a, b, c \in \boldsymbol{Z}.$$

The *Heisenberg graph* is the Cayley graph $\Gamma(G, S)$ where G is the Heisenberg group and S is the generating set $\{A, B, C, A^{-1}, B^{-1}, C^{-1}\}$. When we draw $\Gamma(G, S)$ we use the convention of omitting the edges labeled by A^{-1}, B^{-1}, C^{-1}. The idea is that if ew go backwards along the edge labeled X we are doing X^{-1}.

Let Γ be the Heisenberg graph. Let Γ_0 be the $\boldsymbol{Z}^2$ grid. There is a map $\pi : \Gamma \to \Gamma_0$ which just collapses the C edges. More precisely,

$$\pi(A^aB^bC^c) = (a, b).$$

Figure 2 shows the kind of corkscrew picture which maps to a single square in $\boldsymbol{Z}^2$.

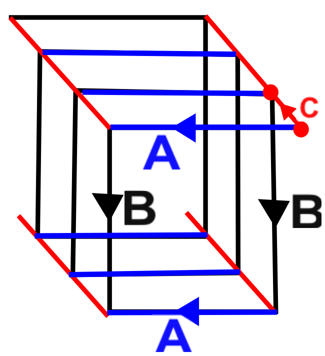

Figure 2: One corkscrew in the Heisenberg graph

The way you get the Heisenberg graph is that you build one corkecrew per unit square in the $\boldsymbol{Z}^2$ graph and glue these corkscrews together along their edges.

The Heisenberg graph has some obvious symmetries. The map which sends the vertex $A^aB^cC^c$ to $A^aB^aC^{c+1}$ respects the edge relations and gives a graph automorphism. More generally, the map $T_n : \Gamma \to \Gamma$ given by

$$T^n(A^aB^bC^c) = A^aB^bC^{c+n} \tag{4}$$

gives a graph automorphism for any $n \in \boldsymbol{Z}$. We call these automorphisms *vertical translations*.

Let γ be a walk in Γ which only uses A and B edges. We say that γ is a *lift* of the path $\gamma_0 = \pi(\gamma)$. Each path γ_0 in the Γ_0 has infinitely many lifts, but they all differ by vertical translations. Suppose γ_0 is a closed loop and γ is some lift. Then π maps the two endpoints of γ to the same point. So, these endpoints differ just by c^k. In other words, one endpoint is $A^aB^bC^c$ and the other is $A^aB^bC^{c+k}$. We call k the *vertical displacement* of the lift.

Lemma 7.1 *The vertical displacement of a lift of a closed loop equals the signed area of the loop.*

Proof: Let A be the signed area enclosed by γ_0. Let $v(\gamma_0)$ denote the vertical displacement of a lift of γ. We want to prove that $v(\gamma_0) = A$. The proof goes by induction on the absolute value of the number of squares enclosed by γ_0. We can write $\gamma_0 = \alpha_0\beta_0$, where α_0 encloses one fewer square and β_0 winds once around the square. The direction which β_0 winds around depends on the loop γ_0. Let's suppose that it winds counterclockwise, so that (by convention) $v(\beta_0) = 1$. Then $v(\alpha_0) = A - 1$ Using the fact that Γ is a group, we have $v(\gamma) = v(\alpha) + v(\beta)$, where α and β are lifts of α_0 and β_0. Hence $V(\gamma_0) = 1 + (A - 1) = A$. ♠

Using this lemma you can see that a path from $A^0B^0C^0$ of length $8n$ can reach any point of the form $A^aB^cC^c$ where $|a| < n$ and $|b| < n$ and $|c| < n^2$. Hence, the number of points of Γ in the ball of radius n is on the order of n^4. By symmetry – i.e. using the fact that the Heisenberg group acts as a group of automorphisms of Γ – we see that the same result holds for balls around any point of Γ. This is the beginning of the proof that Γ is not quasi-isometric to $\boldsymbol{R}^3$.

8 SOL

The graph called SOL is based on the choice of a 2×2 matrix in $SL_2(\boldsymbol{Z})$ which has one eigenvalue greater than 1 and one less than 1. Any two choices lead to quasi-isometric graphs. So, we'll work with the matrix

$$M = \begin{bmatrix} 2 & 1 \\ 3 & 2 \end{bmatrix} \tag{5}$$

We form a graph whose vertices are $\boldsymbol{Z}^3$. First, we add the edges

$$(a,b,k) \leftrightarrow (a \pm 1, b, k), \qquad (a,b,k) \leftrightarrow (a, b \pm 1, k).$$

This creates an infinite stack of copies of the square grid. Next, we add the edges

$$(a,b,k) \to (M(a,b), k+1). \tag{6}$$

For instance, there is an edge joining $(1,1,0)$ to $(2,3,1)$ because we have $M(1,1) = (2,3)$. One can realize SOL as the Cayley graph of a group. Rather than take this point of view, we'll just exhibit lots of symmetries of SOL.

Lemma 8.1 *One can map any vertex of SOL to any other vertex by an automorphism.*

Proof: Let's change notation so that a vertex of SOL is denoted (V,n) where $V \in \boldsymbol{Z}^2$ is a vector. Note that $(V,n) \leftrightarrow (W,n+1)$ if and only if $(V,m) \leftrightarrow (W,m+1)$. For this reason, the map $T(V,m) = (V,n)$ is an automorphism for any integers m and n. In particular, you can use an automorphism to move any point in SOL to a point of the form $(V,0)$. Now, choose some vector W and define

$$T(V,n) = (V + M^n(W), n).$$

This is also an automorphism. In each copy of $\boldsymbol{Z}^2$ the map is just a translation. Also, the vector

$$T(V,n) = (V + M^n(W), n)$$

is connected to

$$(M(V) + M^{n+1}(W), n+1) = T(M(V), n+1)$$

and the vector (V,n) is connected to $(M(V), n+1)$. This shows that T is a graph automorphism. Setting $W = V$, we can map $(V,0)$ to $(0,0)$ by an automorphism. Since we can map any point to $(0,0)$ using an automorphism, we can map any point to any other using an automorphism. ♠

The interesting thing about SOL is that the balls of radius n have about C^n points in them for some $C > 1$. I'll sketch how this is proved. By symmetry, it suffices to check this for balls centered at the origin. The proof is going to be slightly unusual in that we're going to slice SOL by a plane that does not quite contain vertices of SOL. We think of the vertices of SOL as being $\boldsymbol{Z}^3$ and the edges as being straight line segments. When drawn in space, these segments are generally very long, but in the natural metric they all have unit length.

Let E be the eigenvector of M corresponding to the eigenvalue greater than 1. Let Π be the plane spanned by the Z axis and the line through the origin parallel to E. Let $\widehat{\Pi}$ denote the slab of thickness 1 centered on Π. That is, $\widehat{\Pi}$ is the set of all points which are at most one unit from Π in the Euclidean metric. The way to picture $\widehat{\Pi}$ is that is intersect each horizontal plane $R \times \{m\}$ in an infinite strip of width 2. Any portion of this strip having length L intersects about L points of $\boldsymbol{Z}^2$.

Given that $M(L) = L$ and M expands distances by a factor of $\lambda > 1$. We see that all the points in

$$(\boldsymbol{Z}^2 \times \{m\}) \cap \widehat{\Pi} \tag{7}$$

which are within λ^m of the origin can be reached by a path of length about m. But this already gives about λ^m points in the ball of radius m about the origin.

It turns out that $\widehat{\Pi}$ intersects SOL in a graph that is quasi-isometric to the hyperbolic plane. Fattened up planes parallel to the other eigenvector also intersect SOL in graphs which are quasi-isometric to the hyperbolic plane. So, SOL has these two different directions in which you can slice it and get something quasi-isometric to the hyperbolic plane.

There is quite a bit more to say about the geometry of SOL, but these notes are just an introduction.