Notes on P-adic numbers and Serre Trees

Rich Schwartz

December 12, 2019

1 Algebraic Background

Rings and Integral Domains: Previously in the class, we have defined what groups are. A *ring* is an abelian group $(R, +)$ with a second operation $\times$ having the following properties:

- $a \times (b \times c) = (a \times b) \times c$ for all $a, b, c \in \boldsymbol{R}$.
- $a \times (b + c) = a \times b + a \times c$ for all $a, b, c \in \boldsymbol{R}$.
- $(a + b) \times c = a \times c + b \times c$ for all $a, b, c \in \boldsymbol{R}$.

R is called *commutative* if $a \times b = b \times a$ for all $a, b \in R$. The ring R *has a* 1 if there is some element, called 1, such that $1 \times a = a \times 1 = a$ for all $a \in R$. A commutative ring R with 1 is said to be an *integral domain* if the condition $a \times b = 0$ implies that either $a = 0$ or $b = 0$. All the rings we consider will be integral domains. The integers are a classic example of an integral domain. Another example is the ring of polynomials with integer coefficients.

Field of Fractions: Any field is an integral domain. Conversely, any integral domain is contained in a field, called its *field of fractions*. The field of fractions of an integral domain R is the set F of equivalence classes a/b with $b \neq 0$, such that $ra/rb \sim a/b$ for all nonzero $r \in R$. Addition and multiplication in F work just like you think:

$$[a/b] + [c/d] = [(ad + bc)/cd], \qquad [a/b] \times [c/d] = [(a \times c)/(b \times d].$$

It is important to note that the "division sign" is purely formal. It has no meaning in R. However, division makes sense in F. When b, c are nonzero, we have

$$\frac{[a/b]}{[c/d]} = [(a \times d)/(b \times c)].$$

It is a routine exercise to check that all the operations above are well-defined and make F into a ring. F is a field because every nonzero element $[a/b]$ has its multiplicative inverse, $[b/a]$. There is a natural inclusion of R in F. The element $a \in R$ corresponds to $[a/1] \in F$.

Modules: Let R be a ring. An abelian group M is called an *R-module* if there is an operation $\times$ such that $r \times m \in M$ for all $r \in R$ and $m \in M$, subject to the following rules.

- $r \times (s \times m) = (r \times s) \times m$ for all $r, s \in R$ and $m \in M$.
- $(r + s) \times m = (r \times m) + (s \times m)$ for all $r, s \in R$ and $m \in M$.
- $r \times (m + n) = (r \times m) + (r \times n)$ for all $r \in R$ and $m, n \in M$.

In the first rule above, $r \times s$ is defined relative to the multiplication on R. This is almost the same definition as for a ring, but with a subtle difference: The new $\times$ operation has input an element in a ring and an element in a module. rather than 2 elements in a ring.

Exercise 1: let R be an integral domain and F be the field of fractions. We choose 2 vectors $V_j = (x_j, y_j) \in \boldsymbol{F}^2$ and define M to be the set of all linear combinations $a_1V_1 + a_2V_2$ with $a_1, a_2 \in R$. Prove that M is an R-module. This example shows how modules are generalizations of vector spaces.

Ring Homomorphisms: Let R_1 and R_2 be rings. A map $\phi : R_1 \to R_2$ is a *homomorphism* if $\phi(a + b) = \phi(a) + \phi(b)$ and $\phi(a \times b) = \phi(a) \times \phi(b)$ for all $a, b \in R_1$. The addition and multiplication laws are meant to take place in the respective rings. The key example for us is the homomorphism $\boldsymbol{Z}/mn \to \boldsymbol{Z}/n$ which just amounts to reducing mod n. Here $\boldsymbol{Z}/m$ is the cyclic group of residues mod m and likewise $\boldsymbol{Z}/mn$. We call ϕ a *reduction homomorphism*.

2 The P-adic Numbers

Let p be any prime. We have the reduction homomorphism

$$\phi : \boldsymbol{Z}/p^{n+1} \to \boldsymbol{Z}/p^n.$$

For example, when $p = 2$ and $n = 4$ the map $\phi : \boldsymbol{Z}/8 \to \boldsymbol{Z}/4$ has the action $(0, 1, 2, 3, 4, 5, 6, 7) \to (0, 1, 2, 3, 0, 1, 2, 3)$.

A *p-adic integer* is an infinite sequence $\{a_n\}_{n=1}^{\infty}$ such that

- $a_n \in \boldsymbol{Z}/p^n$.
- $\phi(a_{n+1}) = a_n$ for all n.

The set of these sequences is denoted $\boldsymbol{Z}_p$. Componentwise addition and multiplication make $\boldsymbol{Z}_p$ into a ring:

- $\{a_n\} + \{b_n\} = \{a_n + b_n\}$.
- $\{a_n\} \times \{b_n\} = \{a_n \times b_n\}$.

This works because all the reduction maps are ring homomorphisms. This means that the sums and products of p-adic integers are again p-adic integers.

Lemma 2.1 *$\boldsymbol{Z}_p$ is an integral domain.*

Proof: $\boldsymbol{Z}_p$ is clearly commutative, and the 1 element is just $1, 1, 1, ...$ We just have to see that there are no 0-divisors. Suppose that $\{a_n b_n\} = 0$. Then p^n divides $a_n b_n$ for all n. That means that there are infinitely many even indices such that either p^n divides a_{2n} or p^n divides b_{2n}. After switching a and b if necessary, we can assume that the former occurs. So, p^n divides a_{2n} infinitely often. But if p^n divides a_{2n} then p^n also divides a_n because $a_{2n} \equiv a_n \bmod p^n$. But then $a_n = 0$ in $\boldsymbol{Z}/p^n$. So, $\{a_n\}$ is 0 infinitely often. But this means that $a_n = 0$ for all n. ♠

Since $\boldsymbol{Z}_p$ is an integral domain, it makes sense to take its field of fractions. This is called $\boldsymbol{Q}_p$. The space $\boldsymbol{Q}_p$ is called the *field of p-adic numbers*.

Exercise 2: A *unit* in a ring R is an element $a \in R$ such that $ab = 1$ for some $b \in R$. Prove that $\{a_n\}$ in $\boldsymbol{Z}_p$ is a unit provided that $a_1 \neq 0$.

Exercise 3: Use exercise 2 to prove that every p-adic number is equivalent to one the form a/p^k where $a \in \boldsymbol{Z}_p$ and k is some integer. In other words, a p-adic number is almost a p-adic integer; the only difference is that you are allowed to have a power of p in the denominator.

Building on Exercise 3, we see that every p-adic number α can be written uniquely in the form $\alpha = p^j u$ where $u \in \boldsymbol{Z}_p$ is a unit. The integer j is known as the p-adic *valuation* of α.

3 Geometry of the P-adic Integers

There is a rooted tree T_p associated to $\boldsymbol{Z}_p$. For this purpose, it is convenent to define

$$\boldsymbol{Z}/p^0 = \boldsymbol{Z}/1 = \{0\}.$$

The vertices of our tree are all the elements of $\boldsymbol{Z}/p^n$ for every $n = 0, 1, 2, ...$ We join an element $a \in \boldsymbol{Z}/p_{n+1}$ to the element $\phi(a) \in \boldsymbol{Z}/p_n$. We do this for every n and every n. Figure shows the beginning of this tree for $p = 2$.

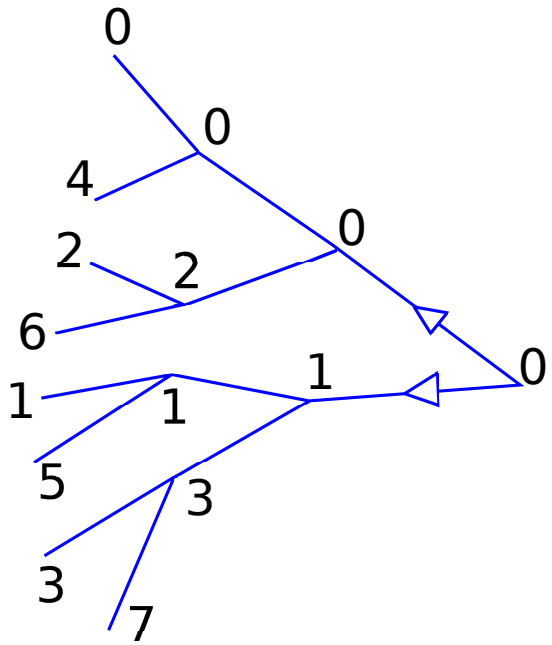

Figure 1. The beginning of the rooted tree for $p = 2$.

In general, T_p is a rooted infinite tree, in which the initial node has degree p and the remaining notes have degree $p+1$. We direct the edges of T_p outward from the root. One can think of $\boldsymbol{Z}_p$ as the set of infinite directed paths in T_p.

There is a natural metric on $\boldsymbol{Z}_p$. The distance between $\{a_n\}$ and $\{b_n\}$ is 2^{-N} where N is the smallest index such that $a_N \neq b_N$. Geometrically, N is the length that the two paths representing $\{a_n\}$ and $\{b_n\}$ agree.

Exercise 4: Prove that $+$ and $\times$ are continuous maps from $\boldsymbol{Z}_p \times \boldsymbol{Z}_p$ into $\boldsymbol{Z}_p$. A ring with this property is called a *topological ring*. So, the p-adic integers form a topological ring.

Lemma 3.1 *Equipped with this topology $\boldsymbol{Z}_p$ is homeomorphic to a Cantor set!*

Proof: I'll explain this for the case $p = 2$. The general case is similar. We can associate a binary sequence $\{\beta_i\}$ to each 2-adic integer $\{a_i\}$. Here $\beta_i = 0$ if $a_{i+1} = a_i$ and $\beta_i = 1$ otherwise. We map $\{a_i\}$ to the point

$$\sum_{i=1}^{\infty} 3^{-i} \times (2\beta_i).$$

Call this map Φ. By construction $\Phi(\{a_i\})$ is a real number whose base 3 decimal expansion has only 0s and 2s. This is a point in the middle third Cantor set. Given the definition of $\boldsymbol{Z}_p$, and the definition of our metric, Φ is a continuous bijection. Moreover, if two 2-adic integers $\{a_i\}$ and $\{b_i\}$ are not that close together, then the base-3 expansions of $\Phi(\{a_n\})$ and $\Phi(\{b_n\})$ differ at an early stage. Hence their images are far apart. In other words, Φ^{-1} is also continuous. Hence Φ is a homeomorphism from $\boldsymbol{Z}_2$ to the middle third Cantor set. ♠

It is a beautiful fact that the Cantor set is a topological ring in many different ways – one way for each prime. You might enjoy trying to visualize what the map $x \to x+1$ looks like e.g. on $\boldsymbol{Z}_2$ or what the map $x \to 2x$ looks like on $\boldsymbol{Z}_3$.

Exercise 5: Extend the metric on $\boldsymbol{Z}_p$ to a metric on $\boldsymbol{Q}_p$ in such a way that $\boldsymbol{Q}_p$ becomes a topological field.

4 Serre Trees: A First Pass

The problem with the tree T_p constructed in the previous section is that it is not quite a homogeneous object. It has a special vertex of degree p whereas all the other vertices have degree $p+1$. It turns out that there is a way to associate the regular $(p+1)$ valent infinite tree to $\boldsymbol{Z}_p$. This object, which we call $\widehat{T}_p$, is known as the *Serre p-adic tree*. In this section I will give the definition that Serre gives in his famous book, *Trees*. The problem with this definition is that it is then hard to see that it gives a tree. In the next section I will explain a much more concrete construction of $\widehat{T}_p$.

The description of $\widehat{T}_p$ builds on Exercise 1. The vertices of $\widehat{T}_p$ are equivalences classes of rank-two $\boldsymbol{Z}_p$ submodules of $\boldsymbol{Q}_p^2$. These are modules of the form

$$M = \{a_1 V_1 + a_2 V_2 |\ a_1, a_2 \in \boldsymbol{Z}_p\}.$$

Here $V_1, V_2 \in (\boldsymbol{Q}_p)^2$ are vectors. We insist that these vectors are linearly independent over $\boldsymbol{Q}_p$. We sometimes write

$$M = \begin{bmatrix} x_1 & y_1 \\ x_2 & y_2 \end{bmatrix}.$$

Here $V_j = (x_j, y_j)$. Here are the two important auxiliary definitions:

- Two such modules M_1 and M_2 are *equivalent* if there is some $\lambda \in \boldsymbol{Q}_p$ such that $\lambda M_1 = M_2$.
- Serre calls two such modules $[M_1]$ and $[M_2]$ *adjacent* if there are representatives $M_j' \in [M_j]$ such that $M_1' \subset M_2'$ and the quotient module M_2'/M_1' is isomorphic to $\boldsymbol{Z}/p$.

The Serre tree $\widehat{T}_p$ is defined to be the graph whose vertices are equivalence classes of rank two $\boldsymbol{Z}_p$ modules and whose edges join adjacent equivalence classes. Serre's fundamental theorem is that $\widehat{T}_p$ is isomorphic as a graph to the regular infinite $(p+1)$ valent tree.

One view of the hyperbolic plane is that space of lattices in $\boldsymbol{R}^2$ modulo similarity. From this perspective, the hyperbolic plane and the Serre tree $\widehat{T}_p$ are pretty close relatives.

5 Serre Trees: A Second Pass

One way to understand the Serre tree is just to start building it. We will take $p = 2$ as usual. First of all, we can view T_p as a subtree of $\widehat{T}_p$. Recall that the vertices of T_p are pairs $(k, 2^n)$ where k is an element of $\boldsymbol{Z}/2^n$. We can associate to $(k, 2^n)$ the lattice

$$M(k, 2^n) = \begin{bmatrix} 2^n & 0 \\ k & 1 \end{bmatrix}$$

Notice that $M(k_1, 2^{n_1})$ and $M(k_2, 2^{n_2})$ really are adjacent when $(k_1, 2^{n_1})$ and $(k_2, 2^{n_2})$ are adjacent vertices of T_p. An example will illustrate this. Consider $(1, 4)$ and $(5, 8)$. The corresponding lattices are

$$M(k, n) = \begin{bmatrix} 4 & 0 \\ 1 & 1 \end{bmatrix}, \qquad \begin{bmatrix} 8 & 0 \\ 5 & 1 \end{bmatrix}.$$

The second one is a sub-module of the first, and the quotient is just $\boldsymbol{Z}/2$. Any element of $M(1, 4)$ either belongs to $M(3, 8)$ or has the form $W + (1, 1)$ where $W \in M(5, 8)$. For instance,

$$(4, 0) = (8, 0) - (5, 1) + (1, 1).$$

The argument above, suitably generalized, identifies T_p as a subtree of $\widehat{T}_p$. There is one fine point, however. We would need to see that two distinct members of T_p give rise to inequivalent lattices.

Lemma 5.1 *$M(k_1, 2^{n_1})$ and $M(k_2, 2^{n_2})$ are equivalent lttices only if we have $(k_1, n_1) = (k_2, n_2)$.*

Proof: We will suppose that M_1 and M_2 are equivalent and we will show that $M_1 = M_2$. Suppose there is some $\lambda \in Q_p$ such that $M_1 = \lambda M_2$. Then the vector $(k_1, 1)$ is the $\boldsymbol{Z}_p$ span of the vectors $(\lambda 2^{n_2}, 0)$ and $(\lambda k_2, \lambda)$. By looking at the second coordinate, we see that this is possible only if λ has non-positive p-adic valuation. That is, $\lambda = p^j u$ for some $j \leq 0$. On the other hand, we have $M_2 = \lambda^{-1} M_1$. Therefore λ^{-1} has non-positive p-adic valuation. But the p-adic valuation of λ^{-1} is the negative of the p-adic valuation of λ. Hence λ has p-adic valuation 0. That is, λ is a unit in $\boldsymbol{Z}_p$. But the lattices uM_2 and M_2 coincide for any unit u. In particular, $M_1 = \lambda M_2 = M_2$. ♠

Now we explain the difference between $\widehat{T}_p$ and T_p. Again, we consider the case $p = 2$. The root vertex in T_2 is $(0, 1)$. The corresponding lattice is

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

This lattice is adjacent to the three lattices:

$$\begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}, \qquad \begin{bmatrix} 2 & 0 \\ 2 & 1 \end{bmatrix}, \qquad \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}.$$

The third lattice is new: It belongs to $\widehat{T}_2$ and not T_2. We can identify an entire new ray in $\widehat{T}_2$. This ray corresponds to the lattices

$$\begin{bmatrix} 1 & 0 \\ 0 & 2^n \end{bmatrix}, \qquad n = 1, 2, 3, ...$$

Figure 2 shows how this ray sits with respect to T_2.

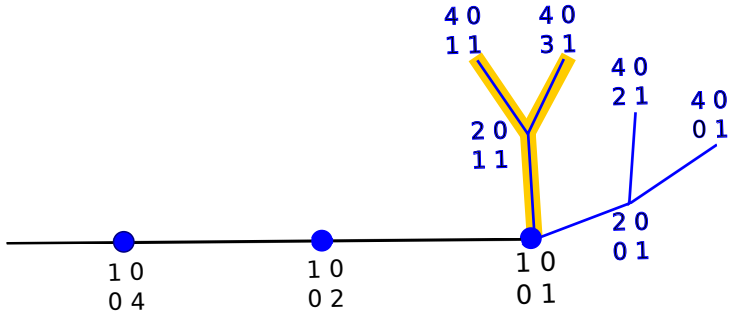

Figure 2. The tree T_2 and its tail in $\widehat{T}_2$.

Figure 2 gives us a hint about how to fill in the rest of the Serre tree. We have highlighted "the odd part" of T_2 in orange. The idea is simply to copy the odd part of T_2 downwards, appropriately changing the lower right entry of the matrix. Figure 3 shows how this works.

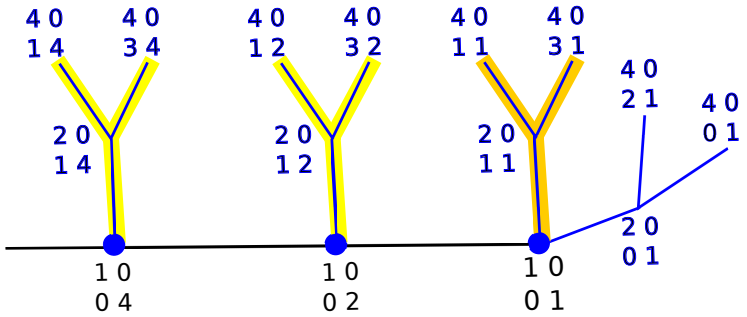

Figure 3. The tree $\widehat{T}_2$.

Figure 3 suggests an infinite 3-valent tree labeled by adjacent and inequivalent lattices. This tree is a sub-graph of $\widehat{T}_2$. At this point, all we have done is shown that $\widehat{T}_2$ contains the regular 3-valent tree as a subgraph. If you believe Serre's theorem, then the subgraph above must be precisely $\widehat{T}_2$.

I am not going to actually prove Serre's theorem in these notes, but let me discuss how far away what we've done is from the proof. A proof would involve 2 more ingredients:

- The result that $\widehat{T}_2$ is connected.
- The result that the vertices of $\widehat{T}_2$ have degree at most 3.

These two facts, together with knowledge of the tree from Figure 3 would complete the proof of Serre's theorem.

Exercise 6: Work out a similar description for the Serre tree $\widehat{T}_3$.

6 The Group Action

You might wonder what is the point of defining the Serre tree. After all, the construction of regular infinite trees is well known. At first glance, Serre's construction just looks like the world's most complicated definition of a regular infinite tree. What makes Serre's construction come alive is that the group $GL_2(\boldsymbol{Q}_p)$, meaning the group of 2×2 invertible matrices with p-adic entries, acts in a natural way on $\widehat{T}_p$.

Given an element $\Lambda \in GL_2(\boldsymbol{Q}_p)$ and a vertex $[M]$ of the Serre tree, we choose a representive M of $[M]$ and let V_1 and V_2 be two vectors which span M. We then define $\Lambda([M])$ to be the equivalence class of the $\boldsymbol{Z}_p$ span of the vectors $\Lambda(V_1)$ and $\Lambda(V_2)$. This definition is independent of all choices.

Lemma 6.1 *The action of Λ is an automorphism of $\widehat{T}_p$.*

Proof: Since Λ is invertible, Λ is a bijection from the vertices of $\widehat{T}_p$ to the vertices of $\widehat{T}_p$. We just have to check that the action of Λ preserves adjacency. If $[M_1]$ and $[M_2]$ are adjacent lattices, then we may choose representatives M_1 and M_2 so that $M_1 \subset M_2$ and $M_2/M_1 = \boldsymbol{Z}/p$. But $\Lambda(M_1) \subset \Lambda(M_2)$ and Λ induces a surjective homomorphism from M_2/M_1 to $\Lambda(M_2)/\Lambda(M_1)$. Since $\Lambda(M_1) \neq \Lambda(M_2)$, this forces $\Lambda(M_1)/\Lambda(M_2) = \boldsymbol{Z}/p$. This means that $\Lambda(M_1)$

and $\Lambda(M_2)$ are adjacent. Hence Λ preserves adjacency in the Serre tree. ♠

Here is a simple example. Suppose that $\Lambda(x, y) = (x, 2y)$. Then Λ is essentially downward translation of the tree shown in Figure 3. Indeed, one can use the action of this Λ to figure how $\widehat{T}_2$ extends T_2. Just take the orbit of T_2 under the action and see how the new vertices are created.

Let's concentrate on the case $p = 2$. Even with the explicit parametrization suggested by Figure 3, figuring out the action of Λ on $\widehat{T}_2$ can be tricky. The problem is that the two vectors $\Lambda(V_1)$ and $\Lambda(V_2)$ are not necessarily in a form which makes it easy to see how they are $\boldsymbol{Z}_2$ linear combinations of the vectors generating one of the special lattices shown in Figure 3. Here are 2 tricks to help do this:

1. It might turn out that Λ can be decomposed into simpler matrices, say $\Lambda = \Lambda_1 ... \Lambda_n$. Then it suffices to figure out the action of Λ_j for each $j = 1, ..., n$.

2. One can figure out the action of Λ in an iterative way. Let M_0 be the equivalence class of the identity matrix in $\widehat{T}_2$. Assuming we want to compute the action of Λ on some M_k, we consider the path in $\widehat{T}_2$ connecting M_0 to M_k. Call this path $M_0, ..., M_k$. Once we locate $\Lambda(M_i)$, we have narrowed down the possibility for $\Lambda(M_{i+1})$ to just 3 lattices. This presents us with a much easier identification problem.

We have stated these tricks for $p = 2$, but the general case is similar.

7 Sketch of an Algorithm

My computer program computes the action of certain matrices on the Serre trees $\widehat{T}_2$ and $\widehat{T}_3$. In this section I will explain how this works. Certain of the routines will look a bit idiosyncratic, due to the way I set up the program. (For instance, I have to take transposes of matrices sometimes; a better implementation would straighten all this out.)

My basic routine works for the matrices

$$\begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}, \qquad \begin{bmatrix} 3 & 0 \\ 0 & 1 \end{bmatrix}, \qquad \begin{bmatrix} 2 & 1 \\ 0 & 2 \end{bmatrix}$$

and their transposes and their adjoints. Many other matrices, such as the generators of the Long-Reid group, can be written as words in these. The algorithm also works directly for some other matrices, for not for all of them.

Here are some routines:

Adjoint:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \to \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}.$$

Clear Factors:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \to \frac{1}{K} \begin{bmatrix} a & b \\ c & d \end{bmatrix}, \qquad K = \gcd(a, b, c, d).$$

Act:

$$\mathrm{Act}(M, A) = \mathrm{clearFactors}(M^t.A)^t.$$

The transposes are just present because of the way I have set up my program.

Brittle Equivalance We compute

$$C = \mathrm{Act}(\mathrm{Adjoint}(A), B)$$

and then check if $\det(C) = \pm 1$. If so we declare that A and B are equivalent. This is a quick way to tell if two lattices are equivalent, but it does not always work.

Dividing by Units: Let $V \in \boldsymbol{Z}^2$ be a vector. We define

$$\begin{bmatrix} x \\ y \end{bmatrix} \to \frac{1}{K} \begin{bmatrix} x \\ y \end{bmatrix}, \qquad K = \gcd_p(x, y).$$

Call the new vector V'. Here we obtain K by taking the GCD and then removing all factors of the prime p of interest – either 2 or 3 in this case. To V and the new vector are $\boldsymbol{Z}_p$ multiples of each other.

Tweaking a Lattice: Starting with the lattice $L = (a, b)$, with $a, b \in \boldsymbol{Z}^2$, there are 4 ways we replace it with an equivalent representation.

1. $(a, b) \to (a', b')$.
2. $(a, b) \to (b, (a + b)')$.

3. $(a, b) \to (a, (a - b)')$..

We let L^k be the kth tweak of L.

Robust Equivalence: Given lattices A and B_0, we do the following:

- Let $B = (B_0)^1$, the first tweak of B_0.
- If easyMatch(A^1, B) is true, return true.
- If easyMatch(A^{12}, B) is true, return true.
- If easyMatch(A^{13}, B) is true, return true.

Otherwise return false. Here, for instance, A^{12} is the second tweak of the first tweak of A. The robust evaluation seems to work without fail for the several matrices listed above.

Guided action: We precompute where our given matrix Λ sends the base lattice

$$A_0 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

We then take a path $A_0, ..., A_n$ to the lattice A_n of interest to us. Assuming that we have determined $\Lambda(A_k)$ we compute $\Lambda(A_{k+1})$ and compare it to each of the neighbors of $\Lambda(A_k)$ using the robust comparison. We then pick the winner and move on. That's it.