Visualizing the p-adic Upper Half Plane and the Bruhat-Tits Tree

Fahad Hossaini[1] **Mustafa Motiwala**[1]

[1]University of Toronto

September 18, 2025

Abstract

In this exposé, we introduce the p-adic upper half plane, the Bruhat-Tits tree, the connections between them and finally, a foray into rigid cocycles all with motivating and explanatory examples.

Keywords. p-adic upper half plane · Bruhat-Tits tree · reduction map

1 Annuli and Affinoids

To work with the p-adic upper half plane, we first need to get situated with $\mathbb{C}_p$ and some useful sets defined in the half plane. These sets are incredibly important in defining important p-adic analogues of meromorphic functions from complex analysis.

1.1 $\mathbb{C}_p$

To begin, we need a brief idea of $\mathbb{C}_p$'s structure.

To this extent, we will assume the reader is acquainted with $\mathbb{Q}_p$, the absolute value on it where $|x|_p = p^{-v_p(x)}$, and the p-adic valuation.

We will extend $|x|_p$ so that it is well defined on any finite field extension of $\mathbb{Q}_p$. In particular, we need to know a few things, which we will assume without proof, but can be find in [1].

(i) If $|x|_p$ were to exist on a finite field extension of $\mathbb{Q}_p$, say K, then K would be complete with respect to $|x|_p$.
(ii) If $|x|_p$ exists, it will be unique. Furthermore, $|x|_p$ on any finite extensions of K will have the same value for elements in K.
(iii) Automorphisms of K, denoted by σ, that fix the base field $\mathbb{Q}_p$ in our field extension, have the following property: $|x|_p = |\sigma(x)|_p$ (this fact secretly uses the point above).
(iv) The following equation holds: $\prod_\sigma |\sigma(x)|_p = |x|_p^n$. Then, recalling that the field norm, denoted $\mathcal{N}_{K/\mathbb{Q}_p}$ is the same as $\prod_\sigma |\sigma|$, we can rearrange to get that:

$$|x| = \sqrt[n]{|\mathcal{N}_{K/\mathbb{Q}_p}(x)|_p}$$

We will assume this is an absolute value without proving it. Now, we can show that any finite field extension of $\mathbb{Q}_p$ is not algebraically closed by finding a transcendental root in $\mathbb{Q}_p$.

Proof.

Consider a sequence $\zeta_1, \zeta_2, ...$ to be roots of unity not in $\mathbb{Q}_p$ of growing degree, so that $\zeta_i \mid \zeta_{i+1}$ and $\left[\mathbb{Q}_p(\zeta_i) : \mathbb{Q}_p\right] < \left[\mathbb{Q}_p(\zeta_{i+1}) : \mathbb{Q}_p\right]$. Then, consider the field extensions $\mathbb{Q}_p(\zeta_1), \mathbb{Q}_p(\zeta_2), ...$

Now, consider the function $x^{m_i} - 1$, where m_i is the degree of the $\mathbb{Q}_p(\zeta_i)$ extension. There exists a unique root of this function that lives strictly in $\mathbb{Q}_p(\zeta_i)$ and not in $\mathbb{Q}_p(\zeta_{i-1})$. Create a sequence of these roots α_i.

Finally, consider the series $\sum_{i=0}^{\infty} \zeta_i p^i$. This limit converges, but it can not converge to an element in a finite extension of $\mathbb{Q}_p$. To see this, let us assume for the sake of contradiction that the series converges to c, which lives in a finite extension of $\mathbb{Q}_p$. So, $\left[\mathbb{Q}_p(c) : \mathbb{Q}_p\right] = n$ for some finite n. But choose i so that $n < m_i$. Thus, α_i is not a root in $\mathbb{Q}_p(c)$, and thus, our series can not converge to c in a finite field extension. Indeed, we found a transcendental root, and thus, $\mathbb{Q}_p$ is incomplete with respect to the extended absolute value.

■

Since finite field extensions of $\mathbb{Q}_p$ are not algebraically closed, we get that the algebraic closure is a field extension of infinite degree. This is called $\overline{\mathbb{Q}}_p$, which we'll denote by $\mathbb{C}_p$ throughout the rest of the paper. So, $|x|_p = p^y$ for some $y \in \mathbb{Q}$ when $x \in \overline{\mathbb{Q}}_p$.

1.2 Projective Space

Now that we have $\mathbb{C}_p$ constructed, we will introduce the p-adic upper half plane. Formally, we say that the p-adic upper half plane is:

$$\mathcal{H}_p = \mathbb{P}_1(\mathbb{C}_p) - \mathbb{P}_1(\mathbb{Q}_p)$$

Let us unravel this definition using $\mathbb{C}$ and $\mathbb{R}$ as an example. Elements of $\mathbb{P}_1(\mathbb{R})$ are equivalence classes of elements in $\mathbb{R}^2$, where the relation is that $(x_1, y_1) \sim (x_2, y_2) \Leftrightarrow \exists c \neq 0, x_1 = cx_2$ and $y_1 = cy_2$. So, for example, the point $(3, 2)$ is equivalent to $(6, 4)$. We denote elements as follows: $[x : y]$.

Note that we would like a 'nice' representative of the equivalence class, and choose the element so that either x or y is a unit or the multiplicative identity over our field. These are called unimodular coordinates.

$\mathbb{P}_1(\mathbb{R}) = \{[0 : 0], [1 : 0]\}$. One way we can think about this is that each element represents a line (excluding the origin as we need to scale), and any point on that line is in the same equivalence class.

So, elements in $\mathbb{P}_1(\mathbb{C})$ are: $\{[x : 1] \mid x \in \mathbb{C}\} \cup \{[0 : 0], [1 : 0]\}$. These two points we add are of importance. The first element is the origin while the second element is called the point at infinity. By scaling, we can see that $\mathbb{P}_1(\mathbb{C}) - \mathbb{P}_1(\mathbb{R})$ is the unit circle in $\mathbb{C}$ except the points $(1, 0)$ and $(-1, 0)$. We can find any element in $\mathbb{C} - \mathbb{R}$ by multiplying by a scalar value as necessary. So we really can get every number in $\mathbb{C} - \mathbb{R}$ by this construction.

Notice that this construction splits the unit circle in two distinct halves. This is not true in the p-adic case, but the definition is still the same. Mainly, $\mathbb{P}_1(\mathbb{C}_p) = \left\{[x : 1] \mid x \in \overline{\mathbb{C}}_p\right\} \cup \{[0 : 0], [1 : 0]\}$. Then, we remove all elements from $\mathbb{P}_1(\mathbb{Q}_p)$ and we get the p-adic upper half plane.

1.3 p-adic Annuli

The next two sections relies heavily on [2] and [3].

An annulus in geometry is a region with a hole in it. The most classic annulus is a ring. In traditional Euclidean space, it is a circle with a smaller circle cut out. And we can visualize $\mathbb{R}^2$

as a plane. But visualizing $\mathcal{H}_p$ will not be an easy task, so we will have to define some things before we can do so.

First, we have to define the p-adic absolute value in two dimensions since elements in $\mathcal{H}_p$ are of two dimensions. Given (x_1, y_1) and (x_2, y_2) as points in $\mathcal{H}_p$, we define $|(x_1, y_1) - (x_2, y_2)| = |x_1 y_2 - x_2 y_1|$.

We define a set called an annulus: $\mathcal{W}_s = \left\{x \in \mathcal{H}_p : \frac{1}{p} < |x - s| < 1, s = 0, 1, ..., p-1\right\}$. This creates a circle centered at point s that includes all elements of absolute value less than 1 but greater than $\frac{1}{p}$. Note that these elements must strictly be in $\mathcal{H}_p$. $\mathcal{W}_0$ is called the standard annulus.

We also define the annulus at infinity: $\mathcal{W}_\infty = \left\{x \in \mathcal{H}_p : 1 < |x - s| < p, s = 0, 1, ..., p-1\right\}$

We provide an illustration for $p = 2$ and $p = 3$:

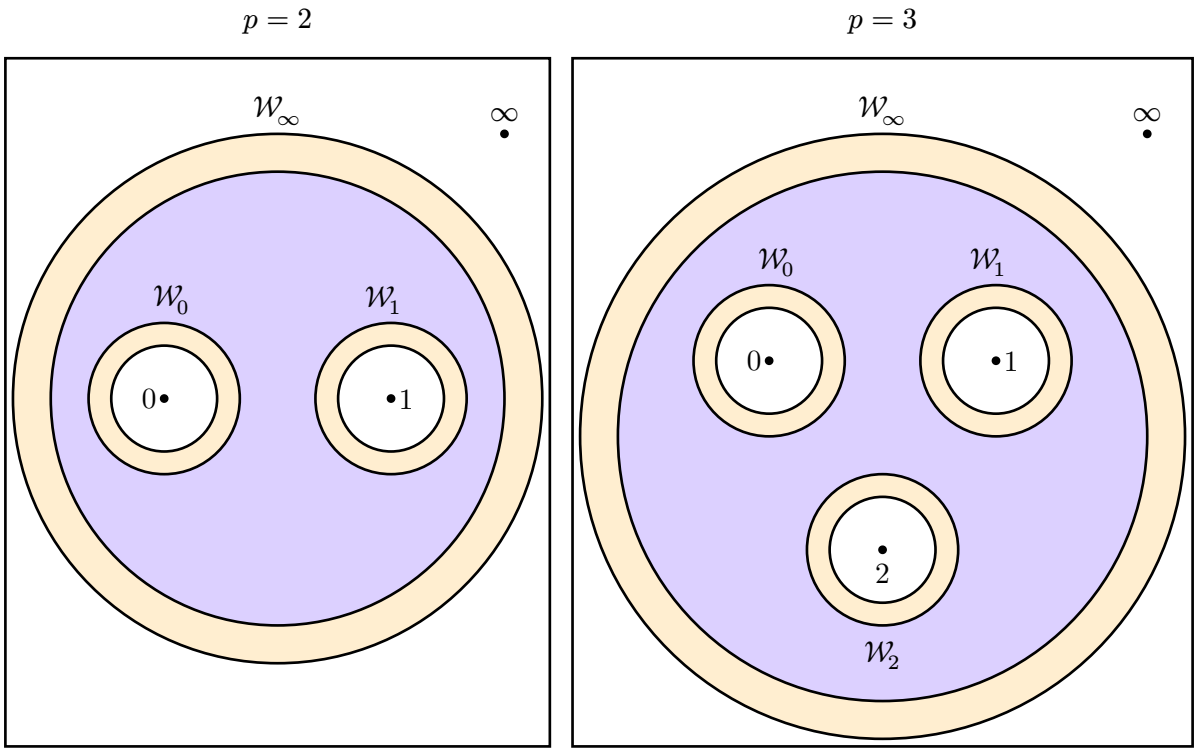

Keep in mind that this illustration is flawed as $\mathbb{C}_p$ is totally disconnected so we do not really have circles. Despite that, the points illustrated in the drawings really are points. We also decided to put the drawing in a box, but there should not be a border of the drawing, like how we visualize $\mathbb{C} \cup \{\infty\}$ as a circle rather than a plane.

This drawing also avoids the fact that the annulus at infinity really is an annulus with infinity in the middle, but these pictures make more sense in terms of definitions. Also, the points are in $\mathbb{P}_1(\mathbb{Q}_p)$, so they are not in $\mathcal{H}_p$. Finally, we note that everything in the drawing (each annulus and the white space inside the box) is in $\mathbb{P}_1(\mathbb{C}_p) \cup \mathbb{P}_1(\mathbb{Q}_p)$.

Let us take a look at an example. We know that $\sqrt{2} \notin \mathbb{Q}_3$, so $\sqrt{2} \in \mathbb{C}_3 - \mathbb{Q}_3$. In particular, we see that $1 + \sqrt{2}$ will have field norm 3 ($a^2 + 2b^2 = 1^2 + 2(1)^2 = 3$). Thus, $|1 + \sqrt{2}|_3 = \sqrt{3}$. Thus, $1 + \sqrt{2} \in \mathcal{W}_0$.

What we would like to do (we will explain why later) is define even more precise annuli, or smaller annuli. For example, say $\mathcal{W}_3$ with $p = 2$. How would we do this? A finer definition:

$$\mathcal{W}_{s,n} = \left\{ x \in \mathcal{H}_p : \frac{1}{p^n} < |x - s| < \frac{1}{p^{n-1}}, s = 0, 1, ..., p^n - 1 \right\}$$

Note: In an attempt to clarify notation and not confuse different sized annuli, we will put a second subscript to denote which power of p we will consider. This is not traditional notation.

By repeating this process, we would get smaller annuli inside $\mathcal{W}_{0,1}$ or $\mathcal{W}_{1,1}$ in some sort of recursive picture. But we also need to deal with $\mathcal{W}_{\infty,n}$. we will define it as follows:

$\mathcal{W}_{\frac{1}{s},n} = \left\{ x \in \mathcal{H}_p : 1 < |x - \frac{1}{s}| < p^n, s = \infty, p, 2p, ..., p^{n-1} - p \right\}$, and by convention, when $s = \infty$, this is the annuli at infinity.

Before we add more annuli, we will make one small change to our diagram. Instead of drawing $\mathcal{W}_\infty$ as a ring around the diagram, we will draw it as a separate annuli around infinity. we will also add the second subscript. So, we get:

$p = 2$ $\qquad\qquad$ $p = 3$

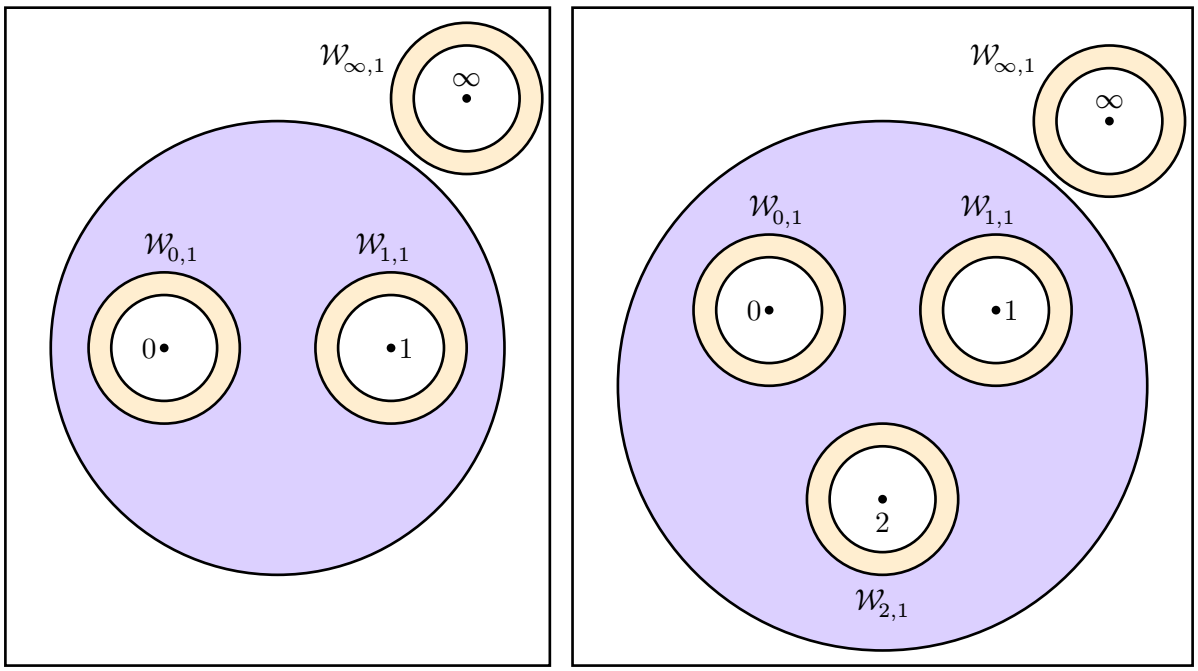

Now, we can draw the diagrams for $\mathcal{W}_{s,2}$. So, here goes:

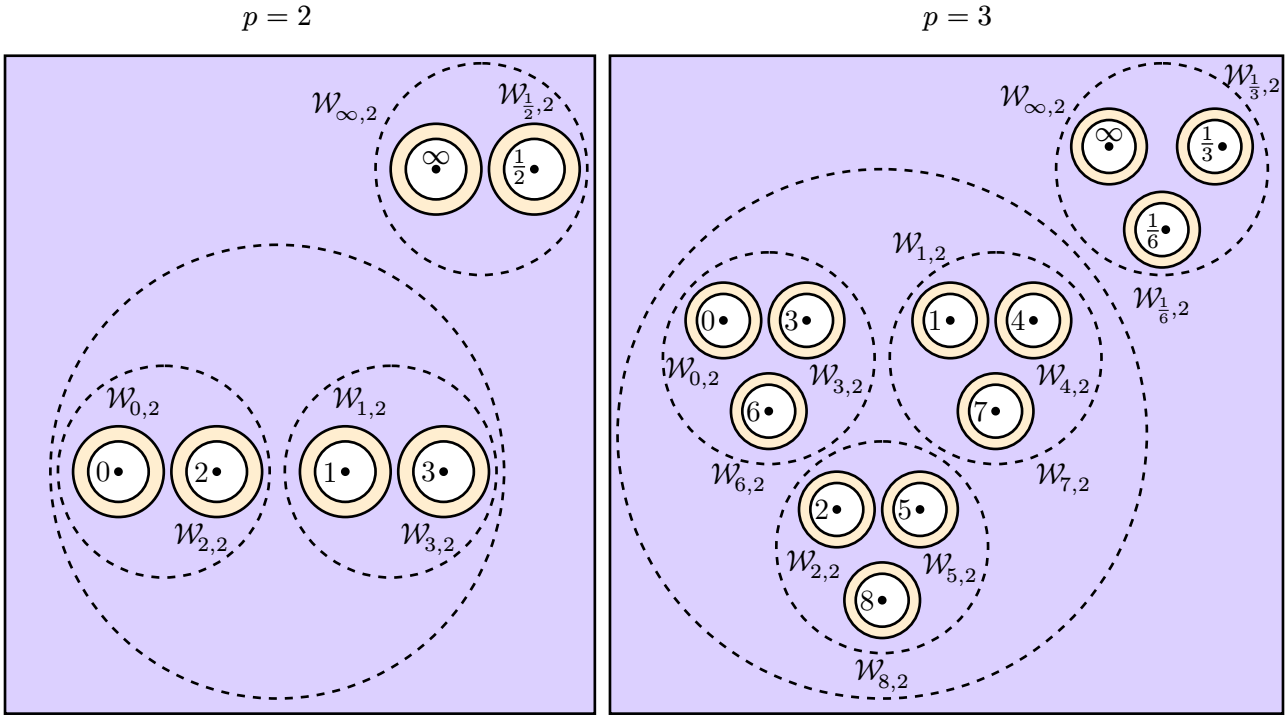

1.4 p-adic Affinoids

As we can see, the purple area really is growing and gaining more points in $\mathcal{H}_p$ as we refine the annuli. We call this purple area the affinoid, denoted Ω. In particular, we correlate the radius of the annuli with the size of the affinoid. i.e. Ω_1 is the purple area in the diagram where we only have annuli of radius 1 and Ω_2 is the purple area in the diagram where we only have annuli of radius 2 and so on. We call Ω_1 the standard affinoid, and denote it by $\mathcal{A}$.

In our diagram of Ω_2, the white area in the box is now in the affinoid. The color is lighter for visualization purposes only. If we had no square enclosing the diagram, this area would be everything outside the annuli from Ω_1. This is the only time something so strange happens.

There are two very nice properties of the affinoid.

(i) $\Omega_n \subset \Omega_{n+1}$. We will not formally define Ω, so we can not prove this, but the diagram makes it clear. Look at [2] for a more careful construction.
(ii) $\cup_{n\in\mathbb{N}}\, \Omega_n = \mathcal{H}_p$. Or, the limit as n goes to infinity makes the affinoid all of $\mathcal{H}_p$.

The big idea is that we can work with the affinoids, take limits (which work nicely), and end up working with $\mathcal{H}_p$. There are two big reasons (other than the obvious 'small things are nicer to work with') we like to work with affinoids over $\mathcal{H}_p$ directly. One of them is that affinoids have a nice algebraic expression. We will cite a theorem from without proof from [2]:

Theorem: Let $[a, 1]$ be a representative of elements in $\mathbb{Z}_p/p^n\mathbb{Z}_p$ and $[1, b]$ be a representative of elements in $p\mathbb{Z}_p/p^n\mathbb{Z}_p$. Let $x = [a, 1]$ or $[1, b]$ depending on the circumstance. Let $z \in \mathcal{H}_p$. Then, $|x - z| \in \Omega_n$ when:

(i) $v_p(z - a) \leq n - 1$
(ii) $v_p\left(z - \frac{1}{b}\right) \leq n - 1 - 2v_p(b)$
(iii) $v_p(z) \geq 1 - n$

We can find out which elements are in which affinoid by checking valuations in $\mathbb{C}_p$.

The second reason is that there is a nice way to visualize $\mathcal{H}_p$ in a very different, but useful way that helps illuminate many of the theorems that require knowledge of $\mathcal{H}_p$. The rest of the paper is dedicated to this goal.

2 Lattices in $\mathbb{Q}_p$

Lattices in $\mathbb{R}^n$ are easy to visualize. For example, we think of $\mathbb{Z}^2$. But we can't visualize $\mathbb{Q}_p$ lattices. $\mathbb{Z}_p$ is dense in $\mathbb{Q}_p$ and $\mathbb{Z}_p$ is uncountable. So, we can't visualize lattices in $\mathbb{Q}_p$ Yet shockingly, the traditional visualization should be sufficient, since we will be talking about linear transformations of the lattice, and this extra complexity does not pose a problem.

We will be looking at lattices in $\mathbb{Q}_p^2$. We say that two lattices, say L with (e_1, e_2) as a basis and L' with (e_1', e_2') are homothetic if the basis elements in L can be rewritten as basis elements in L' scaled by a prime, or, $(e_1, e_2) = (p^n e_1', p^n e_2')$ for $n \in \mathbb{Z}$. This is an equivalence relation, where our relation is called homothety. Intuitively, our lattices are the same up to some prime multiplication in both basis elements.

2.1 Distance Between Lattices

Let L_1, L_2 be two lattices. We aim to quantify the difference between L_1, L_2. To this end, fix a basis $\{v, w\}$ for L_1. We claim that there is a unique way of obtaining, from $\{v, w\}$, a basis for L_2.

Indeed, by the invariant factor theorem, we can find integers $a, b \in \mathbb{Z}$ such that $\{p^a v, p^b w\}$ is a basis for L_2. We define the distance between L_1, L_2 to be the distance between a, b; that is,

$$d(L_1, L_2) = |a - b|$$

That is, L_2 can be obtained from L_1 by scaling the two basis vectors independently and we measure the distance from L_2 to L_1 by measuring how different the two scaling factors are. The definition of homothety is that the two scaling factors actually coincide, so homothetic lattices are actually distance zero apart. Thus, the measurement lifts in a well-defined way to homothety equivalence classes: given L_1, L_2, we can define in the same way as above:

$$d([L_1], [L_2]) = |a - b|$$

Even though our choice of a, b ostensibly depends on L_1, L_2, their distance is actually independent of the representatives L_1, L_2; if $L_1 \sim L_1'$ and $L_2 \sim L_2'$, then there exist $z_1, z_2 \in \mathbb{Q}_p$ such that

$$L_1' = z_1 L_1 \qquad \text{and} \qquad L_2' = z_2 L_2$$

so that if $\{v, w\}$ is a basis for L_1 and $\{p^a v, p^b\}$ is a basis for L_2, then $\{z_1 v, z_1 w\}$ is a basis for L_1 and

$$\{p^a z_2 v, p^b z_2 w\} = \left\{ p^{a + v_p\left(\frac{z_2}{z_1}\right)} z_1 v, p^{b + v_p\left(\frac{z_2}{z_1}\right)} z_2 w \right\}$$

is a basis for L_2', giving

$$d([L_1'], [L_2']) = \left| \left(a + v_p\left(\frac{z_2}{z_1}\right) \right) - \left(b + v_p\left(\frac{z_2}{z_1}\right) \right) \right| = |a - b|$$

By placing the basis vectors of a lattice as the columns of a matrix, we can represent lattices as 2×2 matrices over $\mathbb{Q}_p$. We can describe lattices via a 2×2 matrix. For example, if L is the lattice generated by $\{2e_1, e_1 + 3e_2\}$, then the corresponding matrix would be: $\begin{pmatrix} 2 & 1 \\ 0 & 3 \end{pmatrix}$.

In this representation, lattices of distance one from $\mathbb{Z}_p^2$ are either $\begin{pmatrix} 1 & 0 \\ 0 & p \end{pmatrix}$ or of the form $\begin{pmatrix} p & a \\ 0 & 1 \end{pmatrix}$ for some $a \in \{0, 1, ..., p-1\}$.

There are two very nice reasons to use matrix representations. First, any lattice in the same homothety class will be conjugate by matrices in $\mathrm{SL}_2(\mathbb{Z})$. So, L' is in the same homothety class as L if: $L' = \gamma L \gamma^{-1}$ where $\gamma \in \mathrm{SL}_2(\mathbb{Z})$.

Secondly, finding distances between a lattice and the lattice represented by the identity matrix is very easy. If L has associated matrix $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, and L' has associated matrix γ, then:

$$d(L, L') = v_p(\det\left(\gamma \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right) = v_p(\det(\gamma))$$

2.2 Trees

In this section, we discuss what kind of structure arises from the distance metric we've developed on (homothety classes of) lattices. Since the metric is discrete, a natural representation is as a graph. That is, we take the set of homothety classes as vertices and let "distance 1 apart" define an adjacency relation. This gives us a graph, which we denote by $\mathcal{T}$ and, following [2], call "the Bruhat-Tits tree". Of course, the designation of "tree" requires some justification. We will soon show that the graph $\mathcal{T}$ described here is indeed a tree, but first, let us take a moment to visualize portions of $\mathcal{T}$ in the cases $p = 2$ and $p = 3$.

2.2.1 Examples

Let $L = \mathbb{Z}_p^2 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$. Then, we will explicitly list out the distinct homothety classes of L of distance 1 and draw the associated graph.

For $p = 2$, we get: $L_1 = \begin{pmatrix} 1 & 0 \\ 0 & p \end{pmatrix}, L_2 = \begin{pmatrix} p & 0 \\ 0 & 1 \end{pmatrix}, L_3 = \begin{pmatrix} p & 1 \\ 0 & 1 \end{pmatrix}$. For $p = 3$, we add $L_4 = \begin{pmatrix} p & 2 \\ 0 & 1 \end{pmatrix}$.

$p = 2$ $\qquad$ $p = 3$

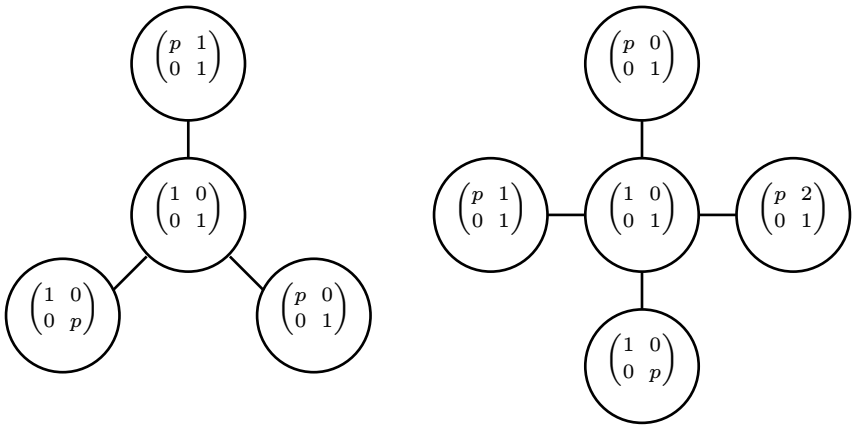

Since distance can be arbitrarily large, this graph has countably many nodes and vertices. We will draw one more complicated example for $p = 2$ and include vertices of up to distance 2 away from the center of the graph. Since we can label any vertex with the matrix $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ via a change of basis, each vertex will have $p + 1$ neighbors, which we'll prove later in this paper. This gives us a graph with 17 nodes.

Indeed, the number of nodes of distance d away, say $V(d)$, is given by the recurrence relation: $V(d) = \begin{cases} 1 & \text{if } d=1 \\ pV(d-1)+2 & \text{otherwise} \end{cases}$ We'll omit the calculation and display the final graph. We have:

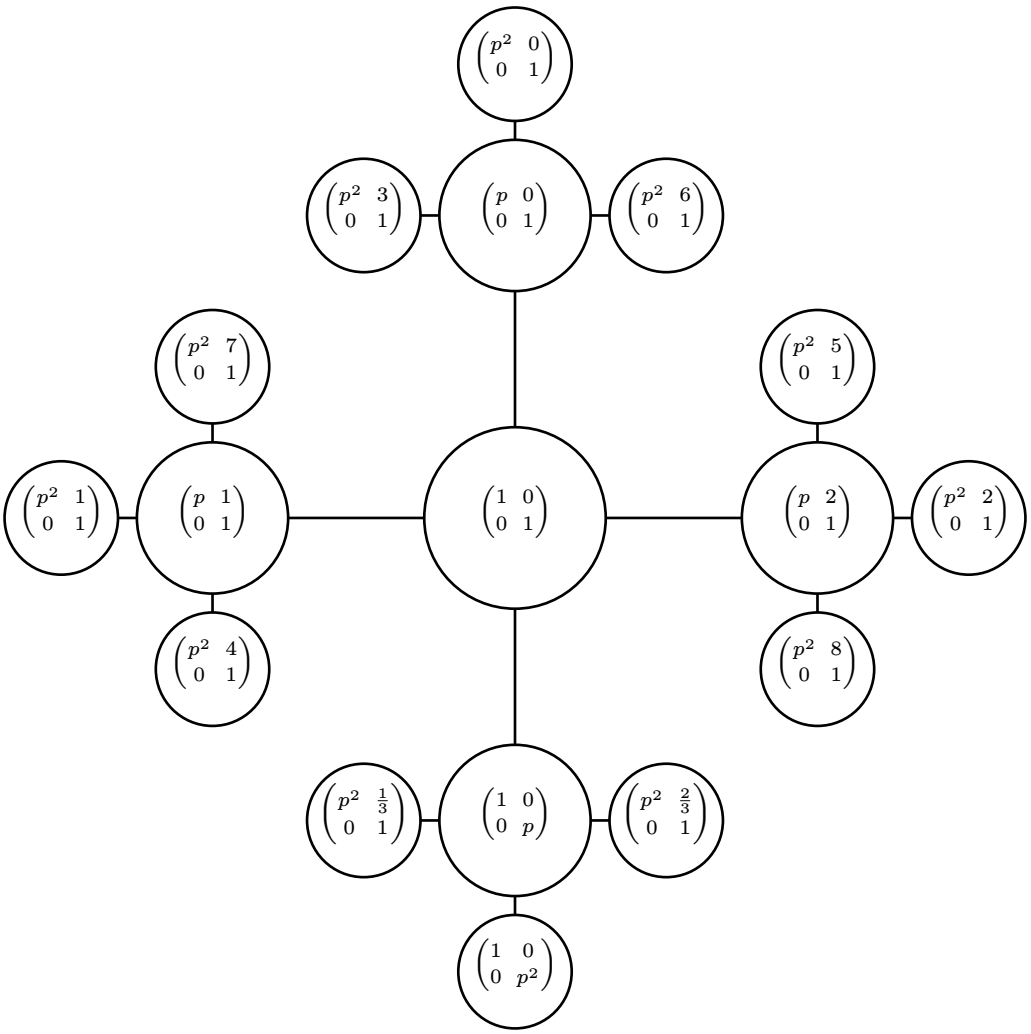

I hope you can see some connection between the holes in the affinoid drawing and the top right entry of each lattice. As we'll explore in chapter 3, this connection goes deeper.

2.2.2 Why the Bruhat-Tits tree is really a tree

First, we offer a different perspective on the adjacency relation.

Claim. *Two vertices $[L_1], [L_2]$ are adjacent if and only if $pL_2 \subsetneq L_1 \subsetneq L_2$.*

Proof.

($\Longrightarrow$) If $d([L_1],[L_2]) = 1$, then (by choosing appropriate representatives), we can assume $L_1 = \langle e_1, e_2 \rangle$ and $L_2 = \langle pe_1, e_2 \rangle$. The inclusions $pL_2 \subsetneq L_1 \subsetneq L_2$ are then clear.

($\Longleftarrow$) Write $L_2 = \langle e_1, e_2 \rangle$ and $L_1 = \langle p^a e_1, p^b e_2 \rangle$ for some integers a, b. From $L_1 \subsetneq L_2$, we obtain $a, b \geq 0$ and from $pL_2 \subsetneq L_1$, we obtain $a, b \leq 1$. The fact that both inclusions are proper ensure that $a \neq b$, which implies $|a - b| = 1$, as needed. ∎

Claim. *The Bruhat-Tits tree $\mathcal{T}$ is indeed a tree.*

Proof. We need to show $\mathcal{T}$ is both connected and acyclic. We focus on the latter first. We borrow this proof from [2].

Suppose for the sake of contradiction that $\mathcal{T}$ contains a cycle; let us take a minimal cycle represented by a chain of lattices

$$L' \subsetneq L_d \subsetneq L_{d-1} \subsetneq \cdots \subsetneq L_1 \subsetneq L$$

where L', L are homothetic and none of the other intermediate lattices are. Since L/L is a cyclic module but L/L' is not, there must be some smallest index i for which L/L_i is cyclic but L/L_{i+1} is not. Then L_{i-1}/L_{i+1} is a length 2 module which is not cyclic, so $L_{i+1} = pL_{i-1}$; that is, L_{i+1} and L_{i-1} are homothetic, contradicting minimality of the cycle.

Now, as for connectedness, let $[L], [L']$ be any vertices of the tree and consider a Jordan-Hölder sequence for L/L':

$$L' = L_n \subset L_{n-1} \subset \cdots \subset L_0 = L$$

Then for each $0 \leq i < n$, we have L_i/L_{i+1} a simple module of length 1 which implies $d(L_{i+1}, L_i) = 1$. In other words, adjacent lattices in the composition series are adjacent in the graph, and hence $[L'], [L]$ are connected, as needed. ■

2.3 Structural Properties of $\mathcal{T}$

An important property of $\mathcal{T}$ is that it is *regular* with degree $p+1$; that is, every vertex has exactly $p+1$ neighbors. This gives $\mathcal{T}$ a fractal-like, self-similar structure. In particular, no particular vertex of $\mathcal{T}$ is inherently more "special" than any other vertex. So, we might as well designate a vertex of our choice to be "privileged". We choose

$$v_0 = [\mathbb{Z}_p^2] \qquad \text{and} \qquad e_0 = \left\{ v_0, \begin{pmatrix} p & 1 \\ 0 & 1 \end{pmatrix} v_0 \right\}$$

to be the privileged vertex and edge respectively. Note that the point $v_1 := \begin{pmatrix} p & 1 \\ 0 & 1 \end{pmatrix} v_0$ is obtained from the natural action of linear transformations on lattices (simply apply the linear transformation to every point on the lattice). It is easy to check that this multiplication is compatible with the homothety relation.

2.4 Ends and $\mathbb{Q}_p$

We define $\text{Ends}(\mathcal{T})$ to be the set of infinite, simple (as in, no repeated vertices) paths in $\mathcal{T}$ quotiented by the relation of being "eventually equal". In other words, two infinite paths $P = ([L_0], [L_1], ...), Q = ([L_0'], [L_1'], ...)$ are equal if they differ in only finitely many places. Intuitively, each path is heading off towards a point "at infinity" and we consider two paths to be the same if they end up at the same point.

We aim to construct a topology on $\text{Ends}(\mathcal{T})$. To that end, we define some basic neighbourhoods $U(e)$ for each edge $e \in \mathcal{T}$ as follows. Fix some edge $e = [L_0] \leftrightarrow [L_1]$. We let

$$U(e) = \{[P] : P = ([L_0], [L_1], ...)\}$$

To visualize $U(e)$, orient the edges of $\mathcal{T}$ so that they radiate out from $[L_0]$ and consider the subtree S of the resulting graph rooted at $[L_1]$: see Figure 1

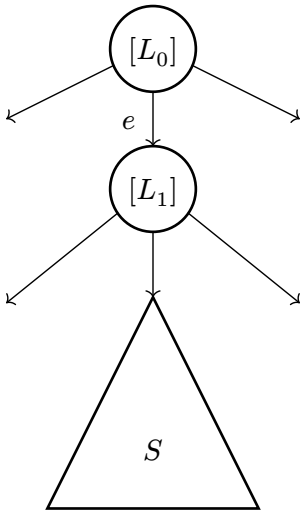

Figure 1: $U(e)$ is the set of paths eventually residing in $\mathcal{S}$

The set $U(e)$ is the set of all paths that eventually reside in S.

Claim. *The set*

$$\mathcal{B} = \{U(e) : e \text{ is an oriented edge in } \mathcal{T}\}$$

forms a basis for a topology on $\text{Ends}(\mathcal{T})$.

Proof. Given any path P starting with the edge e, we clearly have $[P] \in U(e)$, so $\mathcal{B}$ covers $\text{Ends}(\mathcal{T})$. Now, suppose $P \in U(e_1) \cap U(e_2)$. We want to find some e_3 such that $P \in U(e_3)$ and $U(e_3) \subseteq U(e_1) \cap U(e_2)$. Write $P = ([L_0], [L_1], ...)$ and let S_1, S_2 be the subtrees corresponding to e_1, e_2 as in the picture above. Let n be the smallest integer such that $[L_k]$ is in both S_1, S_2 for all $k \geq n$. The existence of this n is guaranteed by the existence of P. Taking e_3 to be the edge from $[L_n]$ to $[L_{n+1}]$, we see that $P \in U(e_3)$ and $U(e_3) \subseteq U(e_1) \cap U(e_2)$.

Thus, $\mathcal{B}$ forms a basis for a topology. ∎

So, we have these "ends" which are like points at infinity, and a topology on these ends; what is this topology? Surprisingly, it is the projective line on $\mathbb{Q}_p$!

The idea behind the correspondence is simple; orient $\mathcal{T}$ as a directed graph moving away from the privileged vertex $v_0 = [\mathbb{Z}_p^2]$. As we have seen, $\mathcal{T}$ is $(p+1)$-regular, so we associate the $(p+1)$-edges exiting each vertex with the digits $0, 1, ..., p-1$ and ∞. The sequence of edges in a path then correspond to digits of a number in $\mathbb{P}^1(\mathbb{Q}_p)$, with edges corresponding to ∞ resulting in a division by p (or moving the decimal point to the right!).

Making this idea precise requires the notion of an inverse limit. First, because we are quotienting out differences in the initial segment of a path, we can identify every element of $\text{Ends}(\mathcal{T})$ with an infinite path starting at the privileged vertex $v_0 = [\mathbb{Z}_p^2]$. On any such path, a vertex w of distance n away from v_0 (i.e, points at the nth position on the path) are in bijection with points of $\mathbb{P}^1(\mathbb{Z}_p/p^n\mathbb{Z}_p)$. Reduction modulo p gives a point in $\mathbb{P}^1(\mathbb{Z}_p/p^{n-1}\mathbb{Z}_p)$, which in turn corresponds to a point distance $n-1$ away (see [2]). In particular, we can do this in a way that we get back a neighbor of w. These reduction maps alongside the points of $\mathcal{T}$ thus form an inverse system which gives

$$\text{Ends}(\mathcal{T}) = \varprojlim \mathbb{P}^1(\mathbb{Z}_p/p^n\mathbb{Z}_p) = \mathbb{P}^1(\mathbb{Z}_p) = \mathbb{P}^1(\mathbb{Q}_p)$$

3 The Connection

All of this effort has been for one big payoff. We can create an equivariant homeomorphism between $\mathcal{T}$ and the drawing in $\mathcal{H}_p$.

3.1 The Reduction Map

We first define the set $PGL_2(\mathbb{Q}_p) = GL_2(\mathbb{Q}_p)/Z_2(\mathbb{Q}_p)$ where $Z_2(\mathbb{Q}_p)$ is the set of scalar transformations. Intuitively, this is the set of matrices with nonzero determinant where scaling keeps you in the same equivalence class. This is similar to how homothety classes are unaffected by scaling by p. In this case, it's any constant.

We define a map $\text{red} : \mathcal{H}_p \to \mathcal{T}$ that is equivariant under $PGL_2(\mathbb{Q}_p)$. Meaning, that for any subset Ω in $\mathcal{H}_p$ and any matrix $\gamma \in PGL_2(\mathbb{Q}_p)$, we have that: $\text{red}(\gamma\Omega) = \gamma \cdot \text{red}(\Omega)$

This is called the reduction map, and its definition is on the sets in $\mathcal{H}_p$. We have:

$$\text{red}(\Omega_1) = \text{red}(\mathcal{A}) = v_0$$
$$\text{red}(\mathcal{W}_{0,1}) = e_0$$

There's one more key thing about this reduction map. Let's denote $\mathcal{T}_n$ to be the subtree of $\mathcal{T}$ with vertices at most $n-1$ away from v_0. Then we have:

$$\text{red}(\Omega_n) = \text{red}(\mathcal{T}_n)$$

Let's visualize these three properties, using color coding to (hopefully) make everything clear:

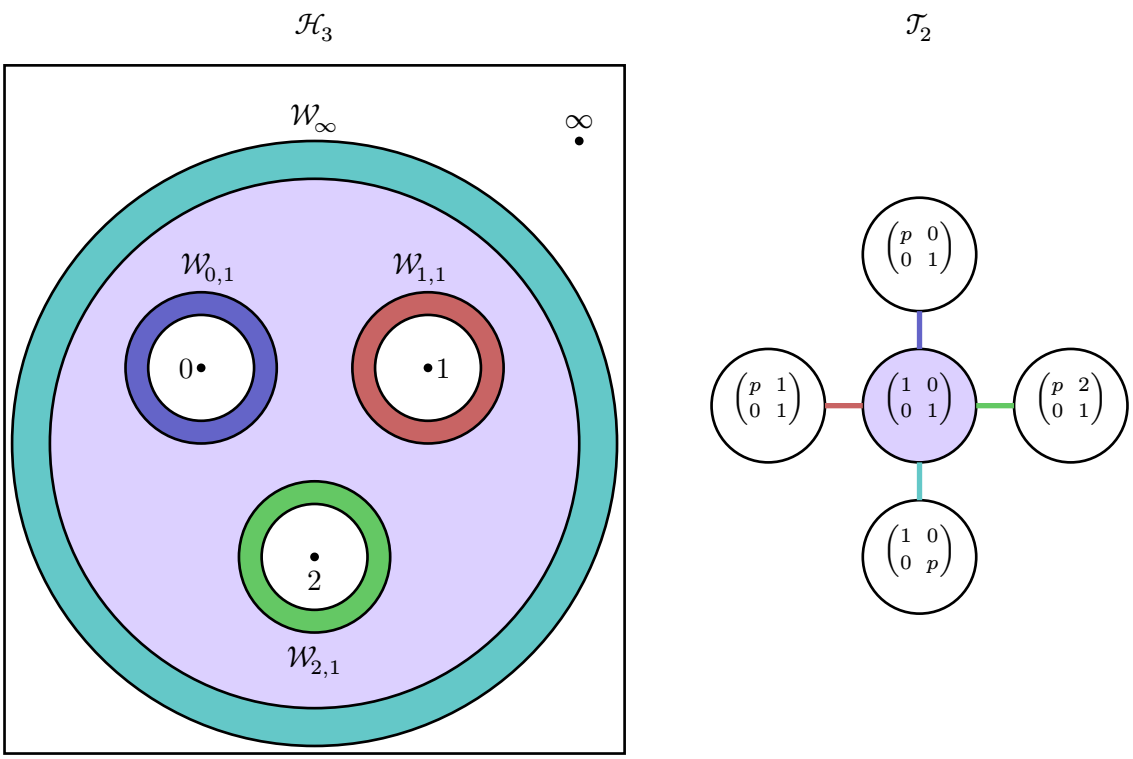

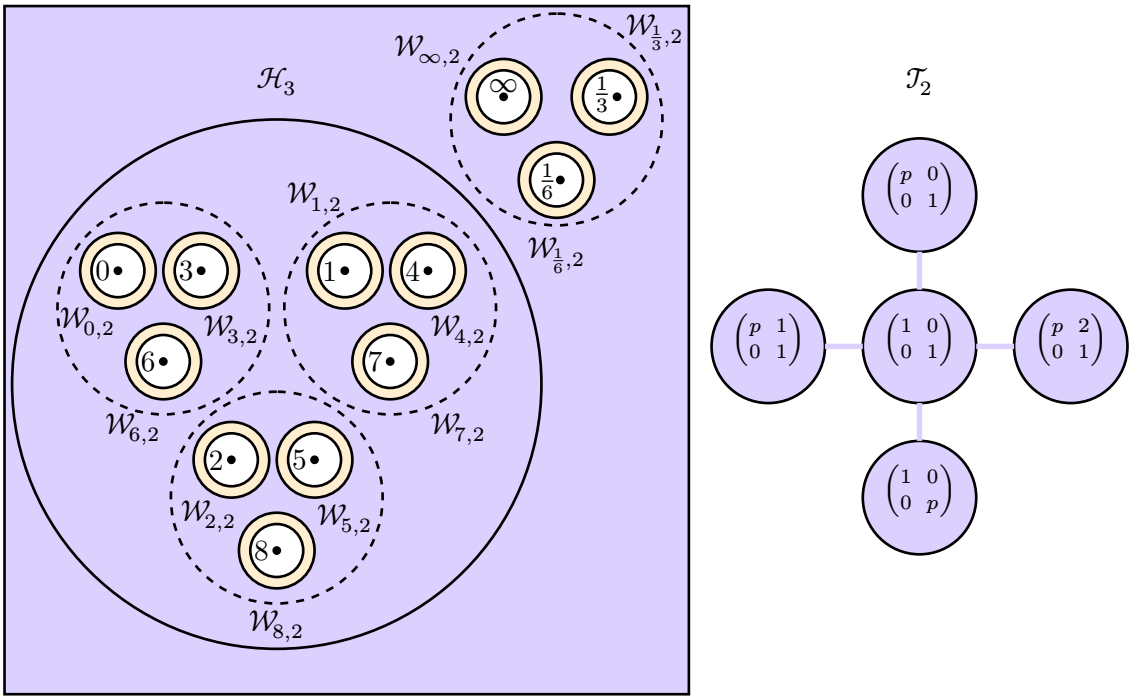

One question we have to deal with is what is the inverse image of vertices under the reduction map. While we will not formally define subaffinoids, the following picture should give a good idea of what is happening:

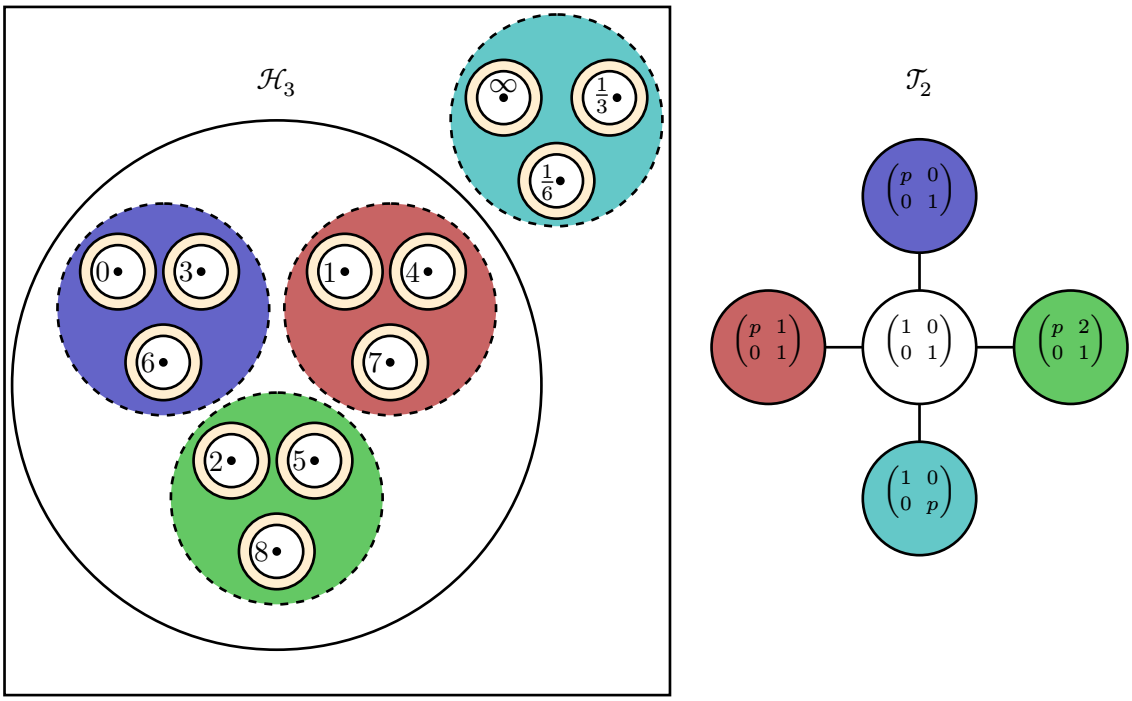

3.2 Ends and $\mathcal{H}_p$

As stated in Chapter 2, we can identify an infinite path in $\mathcal{T}$ with a number in $\mathbb{P}_1(\mathbb{Q}_p)$, or, $\mathbb{Q}_p \cup \{\infty\}$. When taking a look at the diagram and the reduction map, we see that the infinite path is related to the holes in our affinoid.

In particular, assume a path approaches a number $a \in \mathbb{P}_1(\mathbb{Q}_p)$. Let e_n be an edge of distance n away from v_0 in that path. Then, $\lim_{n\to\infty} \mathrm{red}^{-1}(e_n) = \lim_{n\to\infty} \mathcal{W}_{a,n}$ (note that limits are not formally defined).

This infinite path approaches the hole at a.

3.3 Meromorphic Functions

We can now define rigid analytic functions and meromorphic functions on $\mathcal{H}_p$.

Let e be an edge of $\mathcal{T}$ connecting v_1 and v_2. Define $[e] = \{e, v_1, v_2\}$, and call this the closed edge. We define the standard affinoid at $[e]$ as follows:

$$\mathcal{A}_{[e]} := \mathrm{red}^{-1}([e])$$

We say a $\mathbb{C}_p$ valued function f on $\mathcal{H}_p$ is rigid analytic if, for all edges $e \in \mathcal{T}$, the restriction of f to $\mathcal{A}_{[e]}$ is a uniform limit with respect to the sup norm of rational functions on $\mathbb{P}_1(\mathbb{C}_p)$ having poles outside of $\mathcal{A}_{[e]}$.

Finally, we say a $\mathbb{C}_p$ valued function f on $\mathcal{H}_p$ is meromorphic if $f = \frac{g}{h}$ where g, h are two rigid analytic functions and h is non-trivial.

This way of defining meromorphic functions recovers a key property we want from a p-adic analogue of a complex meromorphic function: connectedness. If we were to define a p-adic meromorphic function without $\mathcal{H}_p$ and affinoids, we don't have connectedness since $\mathbb{C}_p$ is totally disconnected.

Bibliography

[1] F. Q. Gouvêa, *P-adic numbers: An introduction*. 2020.

[2] S. Dasgupta and J. Teitelbaum, "The p-adic upper-half plane," vol. 45. American Mathematical Society, pp. 65–121, 2008.

[3] I. Negrini, "On the computation of p-adic theta functions arising from the Hurwitz quaternions," Master's thesis, Università degli Studi di Milano, Concordia University, 2017.