

8 Frobenius elements, the Artin map, completions

We continue to work in the *AKLGB* setup: A is a Dedekind domain, K is its fraction field, L is a finite Galois extension of K with Galois group $G := \text{Gal}(L/K)$, and B is the integral closure of A in L (a Dedekind domain with fraction field L). We now add the further assumption that the residue fields A/\mathfrak{p} (and therefore B/\mathfrak{q}) are finite. This holds in the cases we are most interested in, where K is a global field (a finite extension of \mathbb{Q} or $\mathbb{F}_q(t)$). It follows that B/\mathfrak{q} is then a Galois extension of A/\mathfrak{p} : finite fields are perfect, so the extension is separable, and we proved last time that it is always normal (whether the residue fields are finite or not); see Proposition 7.19.

In order to simplify the notation, when working with finite residue fields we may write $\mathbb{F}_q := B/\mathfrak{q}$ and $\mathbb{F}_p := A/\mathfrak{p}$. These are both finite fields of p -power order, where p is the characteristic of \mathbb{F}_p (and of \mathbb{F}_q). There are two distinct possibilities, depending on the characteristic of K . When K has characteristic 0 (for example, if K is a number field), its characteristic differs from the characteristic of the residue fields, which may vary with \mathfrak{p} but is necessarily nonzero because A/\mathfrak{p} is finite; this is referred to as a *mixed characteristic* setting. When K has positive characteristic p , the residue fields will necessarily have the same characteristic (the ring homomorphism $A \rightarrow A/\mathfrak{p}$ must send 1 to 1, and if $p \cdot 1 = 0$ in $A \subseteq K$, then the same holds in A/\mathfrak{p}), this is an *equal characteristic* setting.

8.1 Frobenius elements

Recall that for each nonzero prime \mathfrak{q} of B lying above a prime \mathfrak{p} of A , the decomposition group $D_{\mathfrak{q}} \subseteq G$ is the stabilizer of \mathfrak{q} under the action of G on $\{\mathfrak{q}|\mathfrak{p}\}$, and Corollary 7.21 gives us the exact sequence

$$1 \longrightarrow I_{\mathfrak{q}} \longrightarrow D_{\mathfrak{q}} \xrightarrow{\pi_{\mathfrak{q}}} \text{Gal}(\mathbb{F}_q/\mathbb{F}_p) \longrightarrow 1.$$

The surjective homomorphism $\pi_{\mathfrak{q}}$ sends $\sigma \in D_{\mathfrak{q}} \subseteq G = \text{Gal}(L/K)$ to an induced automorphism $\bar{\sigma} \in \text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ of \mathbb{F}_q ; the automorphism $\bar{\sigma}$ is obtained by restricting σ to B and noting that since $\sigma(B) = B$ and $\sigma(\mathfrak{q}) = \mathfrak{q}$, we obtain an induced automorphism $\bar{\sigma}$ of $B/\mathfrak{q} = \mathbb{F}_q$ that necessarily fixes the subfield $A/\mathfrak{p} = \mathbb{F}_p$ because σ fixes both $A = B \cap K$ and $\mathfrak{p} = \mathfrak{q} \cap K$; thus $\bar{\sigma} \in \text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$. The inertia group $I_{\mathfrak{q}}$ is then defined as the kernel of $\pi_{\mathfrak{q}}$.

If \mathfrak{p} (equivalently, \mathfrak{q}) is unramified, then $e_{\mathfrak{p}} = 1$ and $I_{\mathfrak{q}}$ is trivial. In this case we have an isomorphism

$$\pi_{\mathfrak{q}} : D_{\mathfrak{q}} \xrightarrow{\sim} \text{Gal}(\mathbb{F}_q/\mathbb{F}_p),$$

and since $\mathbb{F}_q/\mathbb{F}_p$ is an extension of finite fields, the Galois group $\text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ (and hence $D_{\mathfrak{q}}$) is very easy to describe. It is the cyclic group of order $f_{\mathfrak{p}} = [\mathbb{F}_q : \mathbb{F}_p]$ generated by the *Frobenius automorphism*

$$x \mapsto x^{\#\mathbb{F}_p}.$$

Note that the cardinality $\#\mathbb{F}_p$ of the finite field \mathbb{F}_p is a power p^n of its characteristic p . If $K = \mathbb{Q}$ and $\mathfrak{p} = (p)$ is a prime of \mathbb{Z} , then $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ is the field with p elements.

Definition 8.1. Assume *AKLGB* with finite residue fields and $\mathfrak{q}|\mathfrak{p}$ unramified. The inverse image of the Frobenius automorphism of $\text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ under $\pi_{\mathfrak{q}} : D_{\mathfrak{q}} \xrightarrow{\sim} \text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ is the *Frobenius element* $\sigma_{\mathfrak{q}} \in D_{\mathfrak{q}} \subseteq G$, also called the *Frobenius substitution* of \mathfrak{q} .

Proposition 8.2. *Assume AKLBG with finite residue fields and $\mathfrak{q}|\mathfrak{p}$ unramified. The Frobenius element $\sigma_{\mathfrak{q}}$ is the unique $\sigma \in G$ such that for all $x \in B$ we have*

$$\sigma(x) \equiv x^{\#\mathbb{F}_{\mathfrak{p}}} \pmod{\mathfrak{q}}.$$

Proof. It is clear that $\sigma_{\mathfrak{q}}$ has this property, we just need to show uniqueness. Suppose $\sigma \in G$ has the desired property. We claim that $\sigma \in D_{\mathfrak{q}}$: if not then for some $x \in \mathfrak{q}$ we have $\sigma(x) \notin \mathfrak{q}$, and then $\sigma(x) \not\equiv 0 \pmod{\mathfrak{q}}$ and $x^{\#\mathbb{F}_{\mathfrak{p}}} \equiv 0 \pmod{\mathfrak{q}}$, which is a contradiction. So $\sigma \in D_{\mathfrak{q}}$, and since both $\pi_{\mathfrak{q}}(\sigma)$ and $\pi_{\mathfrak{q}}(\sigma_{\mathfrak{q}})$ are the Frobenius automorphism $x \mapsto x^{\#\mathbb{F}_{\mathfrak{p}}}$, we must have $\sigma = \sigma_{\mathfrak{q}}$, since $\pi_{\mathfrak{q}}$ is an isomorphism (because \mathfrak{p} is unramified). \square

Proposition 8.3. *Assume AKLBG with finite residue fields and $\mathfrak{q}|\mathfrak{p}$ unramified. For all $\mathfrak{q}'|\mathfrak{p}$ the Frobenius elements $\sigma_{\mathfrak{q}}$ and $\sigma_{\mathfrak{q}'}$ are conjugate in G .*

Proof. G acts transitively on $\{\mathfrak{q}|\mathfrak{p}\}$, so pick $\tau \in G$ such that $\mathfrak{q}' = \tau(\mathfrak{q})$. For any $x \in B$,

$$\sigma_{\mathfrak{q}}(x) \equiv x^{\#\mathbb{F}_{\mathfrak{p}}} \pmod{\mathfrak{q}}.$$

Applying τ to both sides and then replacing x with $\tau^{-1}(x)$ yields

$$\begin{aligned} \tau(\sigma_{\mathfrak{q}}(x)) &\equiv \tau\left(x^{\#\mathbb{F}_{\mathfrak{p}}}\right) \pmod{\tau(\mathfrak{q})} \\ (\tau\sigma_{\mathfrak{q}})(x) &\equiv \tau(x)^{\#\mathbb{F}_{\mathfrak{p}}} \pmod{\mathfrak{q}'} \\ (\tau\sigma_{\mathfrak{q}})(\tau^{-1}(x)) &\equiv \tau(\tau^{-1}(x))^{\#\mathbb{F}_{\mathfrak{p}}} \pmod{\mathfrak{q}'} \\ (\tau\sigma_{\mathfrak{q}}\tau^{-1})(x) &\equiv x^{\#\mathbb{F}_{\mathfrak{p}}} \pmod{\mathfrak{q}'}, \end{aligned}$$

and the uniqueness given by Proposition 8.2 implies $\sigma_{\mathfrak{q}'} = \tau\sigma_{\mathfrak{q}}\tau^{-1}$. \square

Definition 8.4. *Assume AKLBG with finite residue fields and $\mathfrak{q}|\mathfrak{p}$ unramified. The conjugacy class in G of the Frobenius element $\sigma_{\mathfrak{q}}$ is the *Frobenius class* of \mathfrak{p} , denoted $\text{Frob}_{\mathfrak{p}}$.*

It is common to abuse terminology and refer to $\text{Frob}_{\mathfrak{p}}$ as a Frobenius element $\sigma_{\mathfrak{p}} \in G$ representing its conjugacy class; there is no risk of confusion so long as we remember that $\sigma_{\mathfrak{p}}$ is only determined up to conjugacy (which usually governs all the properties we care about). There is, however, one situation where this terminology is entirely correct. If G is abelian then its conjugacy classes all contains just one element, in which case we necessarily have $\sigma_{\mathfrak{q}} = \sigma_{\mathfrak{p}}$ for all $\mathfrak{q}|\mathfrak{p}$.

8.2 Artin symbols

There is another notation commonly used to denote Frobenius elements that includes the field extension in the notation.

Definition 8.5. *Assume AKLBG with finite residue fields. For each unramified prime \mathfrak{q} of L we define the *Artin symbol**

$$\left(\frac{L/K}{\mathfrak{q}}\right) := \sigma_{\mathfrak{q}}.$$

Proposition 8.6. *Assume AKLBG with finite residue fields and $\mathfrak{q}|\mathfrak{p}$ unramified. Then \mathfrak{p} splits completely if and only if $\left(\frac{L/K}{\mathfrak{q}}\right) = 1$.*

Proof. This follows directly from the definitions: if \mathfrak{p} splits completely then $e_{\mathfrak{p}}f_{\mathfrak{p}} = 1$ and $D_{\mathfrak{q}} = \langle \sigma_{\mathfrak{q}} \rangle = \{1\}$. Conversely, if $D_{\mathfrak{q}} = \langle \sigma_{\mathfrak{q}} \rangle = \{1\}$ then $e_{\mathfrak{p}}f_{\mathfrak{p}} = 1$ and \mathfrak{p} splits completely. \square

We will see later in the course that the extension L/K is completely determined by the set of primes \mathfrak{p} that split completely in L . Thus in some sense the Artin symbol captures the essential structure of L/K .

Proposition 8.7. *Assume AKLBG with finite residue fields and let $\mathfrak{q}|\mathfrak{p}$ be unramified. Let E be an intermediate field between K and L , and let $\mathfrak{q}_E = \mathfrak{q} \cap E$. Then*

$$\left(\frac{L/E}{\mathfrak{q}} \right) = \left(\frac{L/K}{\mathfrak{q}} \right)^{[\mathbb{F}_{\mathfrak{q}_E} : \mathbb{F}_{\mathfrak{p}}]}$$

and if E/K is Galois then $\left(\frac{E/K}{\mathfrak{q}_E} \right)$ is the restriction of $\left(\frac{L/K}{\mathfrak{q}} \right)$ to E .

Proof. For the first claim, note that $\#\mathbb{F}_{\mathfrak{q}_E} = (\#\mathbb{F}_{\mathfrak{p}})^{[\mathbb{F}_{\mathfrak{q}_E} : \mathbb{F}_{\mathfrak{p}}]}$. The second claim follows from the commutativity of the lower right square in the commutative diagram of Proposition 7.23: the Frobenius automorphism $x \mapsto x^{\#\mathbb{F}_{\mathfrak{p}}}$ of $\text{Gal}(\mathbb{F}_{\mathfrak{q}_E}/\mathbb{F}_{\mathfrak{p}})$ is just the restriction of the Frobenius automorphism $x \mapsto x^{\#\mathbb{F}_{\mathfrak{p}}}$ of $\text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}})$. \square

When L/K is abelian, the Artin symbol $\left(\frac{L/K}{\mathfrak{q}} \right)$ takes the same value for all $\mathfrak{q}|\mathfrak{p}$ and we may instead write $\left(\frac{L/K}{\mathfrak{p}} \right)$. In this setting we now view the Artin symbol as a function mapping unramified primes \mathfrak{p} to Frobenius elements $\sigma_{\mathfrak{p}} \in G$. We wish to extend this map to a multiplicative homomorphism from the ideal group \mathcal{I}_A to the Galois group $G = \text{Gal}(L/K)$, but ramified primes $\mathfrak{q}|\mathfrak{p}$ cause problems: the homomorphism $\pi_{\mathfrak{q}} : D_{\mathfrak{q}} \rightarrow \text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}})$ is not a bijection when \mathfrak{p} is ramified (it has nontrivial kernel $I_{\mathfrak{q}}$ of order $e_{\mathfrak{q}} = e_{\mathfrak{p}}$).

For any set S of primes of A , let I_A^S denote the subgroup of \mathcal{I}_A generated by the primes of A that do not lie in S .

Definition 8.8. Let A be a Dedekind domain with finite residue fields. Let L be a finite abelian extension of $K = \text{Frac } A$, and let S be the set of primes of A that ramify in L . The *Artin map* is the homomorphism

$$\begin{aligned} \left(\frac{L/K}{\cdot} \right) : \mathcal{I}_A^S &\rightarrow \text{Gal}(L/K) \\ \prod_{i=1}^m \mathfrak{p}_i^{e_i} &\mapsto \prod_{i=1}^m \left(\frac{L/K}{\mathfrak{p}_i} \right)^{e_i}. \end{aligned}$$

Remark 8.9. We will prove shortly that the set S of ramified primes is always finite, but the definition makes sense in any case.

One of the main results of class field theory is that the Artin map is surjective (this is part of what is known as *Artin reciprocity*). This is a deep theorem that we are not yet ready to prove, but we can verify that it holds in some simple examples.

Example 8.10 (Quadratic fields). Let $K = \mathbb{Q}$ and $L = \mathbb{Q}(\sqrt{d})$ for some square-free integer $d \neq 1$. Then $\text{Gal}(L/K)$ has order 2 and is certainly abelian. As you proved on the problem sets, the only ramified primes $\mathfrak{p} = (p)$ of $A = \mathbb{Z}$ are those that divide the *discriminant*

$$D = \text{disc}(L/K) = \begin{cases} d & \text{if } d \equiv 1 \pmod{4}, \\ 4d & \text{if } d \not\equiv 1 \pmod{4}. \end{cases}$$

If we identify $\text{Gal}(L/K)$ with the multiplicative group $\{\pm 1\}$, then

$$\left(\frac{L/K}{\mathfrak{p}}\right) = \left(\frac{\mathbb{Q}(\sqrt{d})/\mathbb{Q}}{(p)}\right) = \left(\frac{D}{p}\right) = \pm 1,$$

where $\left(\frac{D}{p}\right)$ is the *Kronecker symbol*. For odd primes $p \nmid D$ we have

$$\left(\frac{D}{p}\right) = \begin{cases} +1 & \text{if } D \text{ is a nonzero square modulo } p, \\ -1 & \text{if } D \text{ is not a square modulo } p, \end{cases}$$

and for $p = 2$ not dividing D (in which case $D = d \equiv 1 \pmod{4}$) we have

$$\left(\frac{D}{2}\right) = \begin{cases} +1 & \text{if } D \equiv 1 \pmod{8}, \\ -1 & \text{if } D \equiv 5 \pmod{8}. \end{cases}$$

The cyclotomic extensions $\mathbb{Q}(\zeta_n)/\mathbb{Q}$ provide another interesting example that you will have an opportunity to explore on Problem Set 4.

8.3 Localizing a field

In order to make further progress in our understanding of finite extensions L/K of global fields, and in particular, to understand the primes \mathfrak{p} of K that ramify in L , we want to introduce a new tool that allows us to “localize” fields. Fields have no nontrivial prime ideals (they are Dedekind domains of dimension zero), so we can’t localize them in a nontrivial way as we would a Dedekind domain of dimension one. But there is an analogous operation: by taking the *completion* of a global field K with respect to one of its absolute values we can obtain a *local field*, a term we will define shortly.

For the benefit of those who have not seen this construction before, we first review some background material on completions, topological rings, and inverse limits; those familiar with this material may wish to skip ahead to Section 8.4

8.3.1 Completions

Definition 8.11. Let K be a field with absolute value $|\cdot|$. A sequence (x_n) of elements of K *converges* (to the limit x) if there is an $x \in K$ such that for every $\epsilon > 0$ there is an $N \in \mathbb{Z}_{>0}$ such that $|x_n - x| < \epsilon$ for all $n \geq N$; the limit x is necessarily unique. A sequence (x_n) is *Cauchy* if for every $\epsilon > 0$ there is an $N \in \mathbb{Z}_{>0}$ such that $|x_n - x_m| \leq \epsilon$ for all $m, n \geq N$.

Every convergent sequence is also a Cauchy sequence, but the converse need not hold. The field K is *complete* (with respect to $|\cdot|$) if every Cauchy sequence converges. A subring R of K is *complete* if every Cauchy sequence in R converges to an element of R .

Which sequences converge and which are Cauchy depends very much on the absolute value $|\cdot|$ that we use, and as we have seen in the case of \mathbb{Q} , a field may have infinitely many inequivalent absolute values. Equivalent absolute values necessarily agree on which sequences are convergent and which are Cauchy, so if a field is complete with respect to a given absolute value it is also complete with respect to every equivalent absolute value.

Definition 8.12. Let K be a field with absolute value $|\cdot|$. Two Cauchy sequences $(x_n), (y_n)$ are *equivalent* if $|x_n - y_n| \rightarrow 0$ as $n \rightarrow \infty$. The *completion* of K (with respect to $|\cdot|$) is the field \hat{K} whose elements are equivalence classes of Cauchy sequences with the operations

$$[(x_n)] + [(y_n)] = [(x_n + y_n)] \quad \text{and} \quad [(x_n)][(y_n)] = [(x_n y_n)]$$

(one may verify that these satisfy the field axioms with $0 = [(0, 0, \dots)]$ and $1 = [(1, 1, \dots)]$). The field K is canonically embedded in its completion \hat{K} via the map $x \mapsto \hat{x} = [(x, x, \dots)]$ and we thus view \hat{K} as an extension of K . We extend the absolute value $|\cdot|$ to \hat{K} by defining

$$|[(x_n)]| = \lim_{n \rightarrow \infty} |x_n|,$$

and note that if $|\cdot|$ arises from a discrete valuation v on K , then there is a corresponding discrete valuation on \hat{K} defined by $v([(x_n)]) = \lim v(x_n) \in \mathbb{Z}$ that restricts to v on K (so the discrete valuation on \hat{K} extends v with index 1).

We record the following proposition whose proof is a straightforward exercise.

Proposition 8.13. *Let \hat{K} be the completion of a field K with absolute value $|\cdot|$. The field \hat{K} is complete, and it satisfies the following universal property: every embedding of K into a complete field L can be extended to an embedding of \hat{K} into L . Up to a canonical isomorphism, \hat{K} is the unique field with this property.*

Proof. Exercise. □

The proposition implies that the completion of \hat{K} is (isomorphic to) itself (apply the universal property of the completion of \hat{K} to the embedding $K \rightarrow \hat{K}$); in general, completing a field that is already complete has no effect. In particular, the completion of a field with respect to the trivial absolute value is itself: under the trivial absolute value the only sequences that are Cauchy are those that are eventually constant, and these all clearly converge.

8.3.2 Topological fields with an absolute value

If K is a field with an absolute value $|\cdot|$, then it has a natural topology as a metric space under the distance metric $d(x, y) = |x - y|$ induced by the absolute value. If we define an *open ball* to be a set of the form

$$B_{<r}(x) := \{y \in K : |x - y| < r\},$$

(of *radius* $r > 0$ with *center* $x > 0$), then the open balls form a basis for the topology on K . As with all metric spaces, this topology is Hausdorff, and we note that the *closed ball*

$$B_{\leq r}(x) := \{y \in K : |x - y| \leq r\}$$

is a closed set. The trivial absolute value induces the discrete topology, since every open ball of radius $r < 1$ consists of a single point.

When the absolute value $|\cdot|$ is nonarchimedean the topology it induces has some features that may be counterintuitive to the uninitiated. In particular, every open ball is already closed, so the closure of $B_{<r}(x)$ is not $B_{\leq r}(x)$ unless these two sets are already equal. The latter can definitely happen – the map $|\cdot| : K \rightarrow \mathbb{R}_{\geq 0}$ is in general not surjective, and will even be discrete when $|\cdot|$ corresponds to a discrete valuation. This means that is entirely

possible to have $B_{<r}(x) = B_{<s}(x)$ for $r \neq s$; indeed this must occur uncountably often if $|\cdot|$ arises from a discrete valuation.

The reader may wish to verify that the following hold in any metric space defined by a nonarchimedean absolute value:

1. Every point in an open ball is a center, that is, $B_{<r}(y) = B_{<r}(x)$ for all $y \in B_{<r}(x)$.
2. Any pair of open balls are either disjoint or concentric (have a common center).
3. Every open ball is closed and every closed ball is open.
4. The space is *totally disconnected*: every pair of distinct points lie in disjoint open neighborhoods whose union is the whole space.

The field operations of addition and multiplication are both continuous with respect to this topology as functions from $K \times K \rightarrow K$ (where $K \times K$ is endowed with the product topology). For addition this follows from the triangle inequality, and for multiplication and inversion it follows from the fact that absolute values are multiplicative. The continuity of addition and multiplication makes K into a *topological ring*, and the map $K^\times \rightarrow K^\times$ given by $x \mapsto x^{-1}$ is also continuous (where K^\times is given the subspace topology), which makes K a *topological field*. We can also view K and K^\times as *topological groups* under addition and multiplication, respectively.¹

For any topological space X , the continuity of a map $f: X \times X \rightarrow X$ implies that for every fixed $x \in X$ the maps $X \rightarrow X$ defined by $y \mapsto f(x, y)$ and $y \mapsto f(y, x)$ are continuous, since each is the composition $f \circ \phi$ of f with the continuous map $\phi: X \rightarrow X \times X$ defined by $y \mapsto (x, y)$ and $y \mapsto (y, x)$, respectively. For an additive topological group G this means that every translation by h map $g \mapsto g + h$ is a homeomorphism, since it is continuous and has a continuous inverse (translation by $-h$); in particular, translates of open sets are open and translates of closed sets are closed. A consequence of this is that in order to understand the topology of a topological group, it generally suffices to focus on neighborhoods of the identity; any base of open neighborhoods about the identity determines the entire topology. It also means that any topological property of a subgroup (such as being open, closed, or compact) applies to all of its cosets.

Proposition 8.14. *Let K be a field with absolute values $|\cdot|_1$ and $|\cdot|_2$. The induced topologies on K coincide if and only if $|\cdot|_1$ and $|\cdot|_2$ are equivalent.*

Proof. Only the trivial absolute value induces the discrete topology, so we assume $|\cdot|_1$ and $|\cdot|_2$ are nontrivial. The field K must then be infinite, and if $|\cdot|_1$ and $|\cdot|_2$ are inequivalent we can use weak approximation (Theorem 3.26) to construct a sequence that converges in one topology but not the other. On the other hand, if $|\cdot|_1$ and $|\cdot|_2$ are equivalent, say $|\cdot|_2 = |\cdot|_1^\alpha$, then every open ball $B_{<r}(x)$ in the topology induced by $|\cdot|_1$ is also an open ball $B_{<r^\alpha}(x)$ in the topology induced by $|\cdot|_2$; thus the topologies are the same. \square

If \hat{K} is the completion of K with respect to $|\cdot|$, then \hat{K} is also a topological field with the topology induced by $|\cdot|$, and the subspace topology on $K \subseteq \hat{K}$ is the same as the topology on K induced by $|\cdot|$. By construction, K is *dense* in \hat{K} ; indeed, \hat{K} is precisely the set of limit points of K . More generally, every open ball $B_{<r}(x)$ in K is dense in the corresponding open ball $B_{<r}(x)$ in \hat{K} , and these two sets have the same closure in \hat{K} .

¹For a topological group the inversion map $G \rightarrow G$ must be continuous; for a topological ring this automatically holds for its additive group because multiplication by -1 is continuous.

8.3.3 Inverse limits

Inverse limits are a general construction that can be applied in any category with products, although we will only be concerned with inverse limits in familiar concrete categories such as groups, rings, and topological spaces. Recall that a *concrete category* is one whose objects can be defined as sets (more formally, it is a category that admits a faithful functor to the category of sets), which allows us to speak of the elements of an object in the category.

Definition 8.15. A *directed set* is a set I with a relation “ \leq ” that is reflexive ($i \leq i$), anti-symmetric ($i \leq j \leq i \Rightarrow i = j$), and transitive ($i \leq j \leq k \Rightarrow i \leq k$), in which every finite subset has an upper bound (in particular, I is non-empty).


Definition 8.16. An *inverse system* (*projective system*) in a category is a family of objects $\{X_i : i \in I\}$ indexed by a directed set I and a family of morphisms $\{f_{ij} : X_i \leftarrow X_j : i \leq j\}$ (note the direction) such that each f_{ii} is the identity and $f_{ik} = f_{ij} \circ f_{jk}$ for all $i \leq j \leq k$.²

Definition 8.17. Let (X_i, f_{ij}) be an inverse system in a concrete category with products. The *inverse limit* (or *projective limit*) of (X_i, f_{ij}) is the object

$$X = \varprojlim X_i := \left\{ x \in \prod_{i \in I} X_i : x_i = f_{ij}(x_j) \text{ for all } i \leq j \right\} \subseteq \prod_{i \in I} X_i$$

(whenever such an object X exists in the category). The restrictions $\pi_i : X \rightarrow X_i$ of the projections $\prod X_i \rightarrow X_i$ satisfy $\pi_i = f_{ij} \circ \pi_j$ for $i \leq j$.

The object $X = \varprojlim X_i$ has the universal property that if Y is another object with morphisms $\psi_i : Y \rightarrow X_i$ that satisfy $\psi_i = f_{ij} \circ \psi_j$ for $i \leq j$, then there is a unique morphism $Y \rightarrow X$ for which all of the diagrams



commute (this universal property defines an inverse limit in any category with products).

As with other categorical constructions satisfying (or defined by) universal properties, uniqueness is guaranteed, but existence is not. However, in all the categories that we shall consider, inverse limits exist.

8.4 Valuation rings in complete fields

We now want to specialize to absolute values derived from a discrete valuation $v : K^\times \rightarrow \mathbb{Z}$. If we pick a positive real number $c < 1$ and define $|x|_v := c^{v(x)}$ for $x \in K^\times$ and $|0|_v = 0$ then we obtain a nontrivial nonarchimedean absolute value $|\cdot|_v$. Different choices of c yield equivalent absolute values and thus do not change the induced topology or the completion

²Some (but not all) authors reserve the term *projective system* for cases where the f_{ij} are epimorphisms. This distinction is not relevant to us, as our inverse systems will all use epimorphisms (surjections, in fact).

\hat{K} of K with respect to $|\cdot|_v$. We will see later that there is a natural choice for c when the residue field k of the valuation ring of K is finite (one takes $c = 1/\#k$).

It follows from our discussion that the valuation ring

$$\hat{A} := \{x \in \hat{K} : v(x) \geq 0\} = \{x \in \hat{K} : |x|_v \leq 1\}$$

is a closed (and therefore open) ball in \hat{K} , and it is equal to the closure in \hat{K} of the valuation ring A of K . We now give another characterization of \hat{A} as an inverse limit.

Proposition 8.18. *Let K be a field with absolute value $|\cdot|_v$ induced by a discrete valuation v , let A be the valuation ring of K , and let π be a uniformizer. The valuation ring of the completion of K with respect to $|\cdot|_v$ is a complete DVR \hat{A} with uniformizer π , and we have an isomorphism of topological rings*

$$\hat{A} \simeq \varprojlim_{n \rightarrow \infty} \frac{A}{\pi^n A}.$$

It is clear that \hat{A} is a complete DVR with uniformizer π : it is complete because it is closed and therefore contains all its limit points in the complete field \hat{K} , it is a DVR with uniformizer π because v extends to a discrete valuation on \hat{A} with $v(\pi) = 1$.

Before proving the main part of the proposition, let us check that we understand the topology of the inverse limit $\varprojlim_n A/\pi^n A$. The valuation ring A of K is a closed ball $B_{\leq 1}(0)$ (hence an open set) in the nonarchimedean metric space K , and this also applies to each of the sets $\pi^n A$ (they are each closed balls of radius c^n about 0). This implies that each quotient $A/\pi^n A$ has the discrete topology, since the inverse image of any point under the quotient map is a coset of the open subgroup $\pi^n A$. The inverse limit is a subspace of the infinite product $\prod_n A/\pi^n A$, whose open subsets project onto $A/\pi^n A$ for all but finitely many factors (by definition of the product topology). It follows that the open subsets U of the inverse limit are each the full inverse image (under the canonical projection maps given by the inverse limit construction) of a subset of $A/\pi^m A$ for some m . If we consider the case of a singleton set $\{\bar{x}\}$ in $A/\pi^m A$ and choose a representative $x \in A$, taking the image of the coset $x + \pi^m A = B_{\leq c^m}(x)$ in each $A/\pi^n A$ gives the inverse image of $\bar{x} \in A/\pi^m A$ in $\varprojlim A/\pi^n A$ under the projection map.

We can alternatively describe the topology on $\varprojlim_n A/\pi^n A$ in terms of an absolute value: for $x = (x_n) \in \varprojlim_n A/\pi^n A$, let $v(x)$ be the least $n \geq 0$ for which $x_n \neq 0$, and define $|x|_v = c^{v(x)}$; the image of the coset $x + \pi^m A$ in $\varprojlim A/\pi^n A$ is then just a closed ball $B_{\leq c^m}(x)$ as above. If we embed A in $\varprojlim A/\pi^n A$ in the obvious way ($x \mapsto (\bar{x}, \bar{x}, \bar{x}, \dots)$), the absolute values and subspace topology agree.

Proof. The ring \hat{A} is complete and contains A . For each $n > 0$ we define a ring homomorphism $\phi_n: \hat{A} \rightarrow A/(\pi^n)$ as follows: for each $\hat{a} = [(a_i)]$ let $\phi_n(\hat{a})$ be the limit of the eventually constant sequence (\bar{a}_i) of images of a_i in $A/(\pi^n)$. We thus obtain an infinite sequence of surjective maps $\phi_n: \hat{A} \rightarrow A/\pi^n A$ that are compatible in that for all $n \geq m > 0$ and all $a \in \hat{A}$ the image of $\phi_n(a)$ in $A/\pi^m A$ is $\phi_m(a)$. So we have a surjective ring homomorphism $\phi: \hat{A} \rightarrow \varprojlim A/\pi^n A$. Now note that

$$\ker \phi = \bigcap_{n \geq 1} \pi^n \hat{A} = \{0\}, \tag{1}$$

so ϕ is injective and therefore an isomorphism. To show that ϕ is also a homeomorphism, it suffices to note that if $x + \pi^m A$ is a coset of $\pi^m A$ in A and U is the corresponding open

set in $\varprojlim A/\pi^n A$, then $\phi^{-1}(U)$ is the closure of $x + \pi^m A$ in \hat{A} , which is the coset $x + \pi^m \hat{A}$, an open subset in \hat{A} (as explained in the discussion above, every open set in the inverse limit corresponds to a finite union of cosets $x + \pi^m A$ for some m). Conversely ϕ maps open cosets $x + \pi^m \hat{A}$ to open sets in $\varprojlim A/\pi^n A$. \square

Remark 8.19. Given any ring R with an ideal I , one can define the I -adic completion of R as the inverse limit of topological rings $\hat{R} := \varprojlim_n R/I^n$, where each R/I^n is given the discrete topology. Proposition 8.18 shows that when R is a DVR with maximal ideal \mathfrak{m} , taking the completion of R with respect to the absolute value $|\cdot|_{\mathfrak{m}}$ is the same thing as taking the \mathfrak{m} -adic completion. This is not true in general. In particular, the \mathfrak{m} -adic completion of a (not necessarily discrete) valuation ring R with respect to its maximal ideal \mathfrak{m} need not be complete (either in the sense of Definition 8.11 or in the sense of being isomorphic to its \mathfrak{m} -adic completion). The key issue that arises is that the kernel in (1) need not be trivial; indeed, if $\mathfrak{m}^2 = \mathfrak{m}$ (which can happen) it certainly won't be. This problem does not occur for valuation rings that are noetherian, but these are necessarily DVRs.

Example 8.20. Let $K = \mathbb{Q}$ and let $v = v_p$ be the p -adic valuation for some prime p and let $|x|_p := p^{-v_p(x)}$ denote the corresponding absolute value. The completion of \mathbb{Q} with respect to $|\cdot|_p$ is the field \mathbb{Q}_p of p -adic numbers. The valuation ring of \mathbb{Q} corresponding to v_p is the local ring $\mathbb{Z}_{(p)}$. Taking $\pi = p$ as our uniformizer, we get

$$\widehat{\mathbb{Z}_{(p)}} \simeq \varprojlim_{n \rightarrow \infty} \frac{\mathbb{Z}_{(p)}}{p^n \mathbb{Z}_{(p)}} \simeq \varprojlim_{n \rightarrow \infty} \frac{\mathbb{Z}}{p^n \mathbb{Z}} = \mathbb{Z}_p,$$

the ring of p -adic integers.

Example 8.21. Let $K = \mathbb{F}_q(t)$ be the rational function field over a finite field \mathbb{F}_q and let $v = v_t$ be the t -adic valuation and let $|x|_t := q^{-v_t(x)}$ be the corresponding absolute value. with uniformizer $\pi = t$ The completion of $\mathbb{F}_q(t)$ with respect to $|\cdot|_t$ is isomorphic to the field $\mathbb{F}_q((t))$ of Laurent series over \mathbb{F}_q . The valuation ring of $\mathbb{F}_q(t)$ with respect to v_t is the local ring $\mathbb{F}_q[t]_{(t)}$ consisting of rational functions whose denominators have nonzero constant term. Taking $\pi = t$ as our uniformizer, we get

$$\widehat{\mathbb{F}_q[t]_{(t)}} \simeq \varprojlim_{n \rightarrow \infty} \frac{\mathbb{F}_q[t]_{(t)}}{t^n \mathbb{F}_q[t]_{(t)}} \simeq \varprojlim_{n \rightarrow \infty} \frac{F_q[t]}{t^n F_q[t]} \simeq \mathbb{F}_q[[t]],$$

where $\mathbb{F}_q[[t]]$ denotes the power series ring over \mathbb{F}_q .

Example 8.22. The isomorphism $\mathbb{Z}_p \simeq \varprojlim \mathbb{Z}/p^n \mathbb{Z}$ gives us a canonical way to represent elements of \mathbb{Z}_p : we can write $a \in \mathbb{Z}_p$ as a sequence (a_n) with $a_{n+1} \equiv a_n \pmod{p^n}$, where each $a_n \in \mathbb{Z}/p^n \mathbb{Z}$ is uniquely represented by an integer in $[0, p^n - 1]$. In \mathbb{Z}_7 , for example:

$$\begin{aligned} 2 &= (2, 2, 2, 2, 2, \dots) \\ 2002 &= (0, 42, 287, 2002, 2002, \dots) \\ -2 &= (5, 47, 341, 2399, 16805, \dots) \\ 2^{-1} &= (4, 25, 172, 1201, 8404, \dots) \\ \sqrt{2} &= \begin{cases} (3, 10, 108, 2166, 4567 \dots) \\ (4, 39, 235, 235, 12240 \dots) \end{cases} \\ \sqrt[5]{2} &= (4, 46, 95, 1124, 15530, \dots) \end{aligned}$$

While this representation is canonical, it is also redundant. The value of a_n constrains the value of a_{n+1} to just p possible values among the p^{n+1} elements of $\mathbb{Z}/p^{n+1}\mathbb{Z}$, namely, those that are congruent to a_n modulo p^n . We can always write $a_{n+1} = a_n + p^n b_n$ for some $b_n \in [0, p-1]$, namely, $b_n = (a_{n+1} - a_n)/p^n$.

Definition 8.23. Let $a = (a_n)$ be a p -adic integer with each a_n uniquely represented by an integer in $\in [0, p^n - 1]$. The sequence (b_0, b_1, b_2, \dots) with $b_0 = a_1$ and $b_n = (a_{n+1} - a_n)/p^n$ is called the p -adic expansion of a .

Proposition 8.24. Every element of \mathbb{Z}_p has a unique p -adic expansion and every sequence (b_0, b_1, b_2, \dots) of integers in $[0, p-1]$ is the p -adic expansion of an element of \mathbb{Z}_p .

Proof. This follows immediately from the definition: we can recover (a_n) from its p -adic expansion (b_0, b_1, b_2, \dots) via $a_1 = a_0$ and $a_{n+1} = a_n + pb_n$ for all $n \geq 1$. \square

Thus we have a bijection between \mathbb{Z}_p and the set of all sequences of integers in $[0, p-1]$ indexed by the nonnegative integers.

Example 8.25. We have the following p -adic expansion in \mathbb{Z}_7 :

$$\begin{aligned} 2 &= (2, 0, 0, 0, 0, 0, 0, 0, 0, \dots) \\ 2002 &= (0, 6, 5, 5, 0, 0, 0, 0, 0, \dots) \\ -2 &= (5, 6, 6, 6, 6, 6, 6, 6, 6, \dots) \\ 2^{-1} &= (4, 3, 3, 3, 3, 3, 3, 3, 3, \dots) \\ 5^{-1} &= (3, 1, 4, 5, 2, 1, 4, 5, 2, 1, \dots) \\ \sqrt{2} &= \begin{cases} (3, 1, 2, 6, 1, 2, 1, 2, 4, 6, \dots) \\ (4, 5, 4, 0, 5, 4, 5, 4, 2, 0, \dots) \end{cases} \\ \sqrt[5]{2} &= (4, 6, 1, 3, 6, 4, 3, 5, 4, 6, \dots) \end{aligned}$$

You can easily recreate these examples (and many more) in Sage. To create the ring of 7-adic integers, use `Zp(7)`. By default Sage uses 20 digits of p -adic precision, but you can change this to n digits using `Zp(p, n)`.

Performing arithmetic in \mathbb{Z}_p using p -adic expansions is straight-forward. One computes a sum of p -adic expansions $(b_0, b_1, \dots) + (c_0, c_1, \dots)$ by adding digits mod p and carrying to the right (don't forget to carry!). Multiplication corresponds to computing products of formal power series in p , e.g. $(\sum b_n p^n)(\sum c_n p^n)$, and can be performed by hand (or in Sage) using the standard schoolbook algorithm for multiplying integers represented in base 10, except now one works in base p . For more background on p -adic numbers, see [1, 2, 3, 4].

References

- [1] F.Q. Gouvea, *[p-adic numbers](#)*, Springer, 1993.
- [2] N. Koblitz, *[p-adic numbers, p-adic analysis, and zeta functions](#)*, Springer, 1984.
- [3] A.M. Robert, *[A course in p-adic analysis](#)*, Springer, 2000.
- [4] J.-P. Serre, *[A course in arithmetic](#)*, Springer, 1979.