Algebraic number theory, Michaelmas 2015

January 22, 2016

Contents

1 Discrete valuation rings

Definition 1.1. *Let A be a ring. We say that A is a discrete valuation ring (DVR) if A is a principal ideal domain with a unique non-zero prime ideal $\mathfrak{m}_A \subset A$.*

If A is a DVR, then $k_A = A/\mathfrak{m}_A$ is a field, its residue field. If $\pi \in \mathfrak{m}_A$ is a generator, then π is an irreducible element of A, and every other irreducible element of A has the form $\pi' = u\pi$, $u \in A^\times$. Such an element is called a uniformizer of A.

If A is a DVR with $\operatorname{Frac} A = K$, then any element $x = a/b \in K^\times$ admits a unique expression as $x = \pi^n u$, with $n \in \mathbb{Z}$ and $u \in A^\times$. We call $v(x) = n$ the valuation of the element x. The function $v : K^\times \to \mathbb{Z}$ satisfies the following conditions:

1. $v : K^\times \to \mathbb{Z}$ is a surjective homomorphism.
2. If $x, y \in K$ then we have $v(x+y) \geq \min(v(x), v(y))$, with equality if $v(x) \neq v(y)$. (By convention, we set $v(0) = \infty$ and extend v to a map $v : K \to \mathbb{Z} \cup \{\infty\}$.)

If K is any field and $v : K^\times \to \mathbb{Z}$ is a function satisfying these conditions, then v is called a valuation. We can then define $A = \{x \in K \mid v(x) \geq 0\}$; then A is a DVR with non-zero prime ideal $\mathfrak{m}_A = \{x \in K \mid v(x) > 0\}$.

Lemma 1.2. *Let K be a field. Then this process defines a bijection between the set of subrings $A \subset K$ which are discrete valuation rings and the set of valuations $v : K^\times \to \mathbb{Z}$.*

Definition 1.3. *A ring A is called a local ring if it has a unique maximal ideal.*

Lemma 1.4 (Nakayama's lemma)**.** *Let A be a local ring with maximal ideal $\mathfrak{m}$, and let M be a finitely generated A-module.*

1. *If $\mathfrak{m}M = M$, then $M = 0$.*
2. *If $N \subset M$ is a submodule such that $N \to M/\mathfrak{m}M$ is surjective, then $N = M$.*

Proof. We observe that $A - \mathfrak{m} = A^\times$. For the first part, suppose that M is generated by elements $m_1, \dots, m_n$, and that $n \geq 1$ is minimal with respect to this property. We can write $m_n = a_1 m_1 + \cdots + a_n m_n$ with $a_i \in \mathfrak{m}$, hence $(1 - a_n) m_n = a_1 m_1 + \dots a_{n-1} m_{n-1}$. Since $1 - a_n$ is a unit, we conclude that A is generated by $m_1, \dots, m_{n-1}$, a contradiction to the minimality of n.

For the second part, the hypothesis implies $M = N + \mathfrak{m}M$. We apply the first part to the quotient M/N to conclude $M/N = 0$, hence $M = N$. □

Proposition 1.5. *Let A be a Noetherian domain. Then the following are equivalent:*

1. *A is a DVR.*
2. *A is integrally closed and has exactly one non-zero prime ideal.*

Proof. We recall that if $R \subset S$ are rings, then an element $s \in S$ is said to be integral over R if it satisfies a monic equation $s^n + r_1 s^{n-1} + \cdots + r_n = 0$ with $r_i \in R$; equivalently, $R[s]$ is a finitely generated R-module. The integral closure of R in S consists of all elements of S integral over R, and is a subring of S which contains R. The domain A is said to be integrally closed if A equals the integral closure of A in $K = \mathrm{Frac}(A)$.

If A is a DVR, then A is integrally closed: if an element $a = u\pi^r$ satisfies a relation $a^n + a_1 a^{n-1} + \cdots + a_n = 0$, then after multiplying through by u^{-n}, we can assume $a = \pi^r$. If $r < 0$ then we get $1 = a_1 \pi^{-r} + \cdots + a_n \pi^{-nr}$. The left-hand side lies in $\mathfrak{m}_A$, a contradiction. It clearly has a unique non-zero prime ideal.

Suppose conversely that A is a Noetherian domain which is integrally closed, and has exactly one non-zero prime ideal $\mathfrak{p}$. Note that if $\mathfrak{a} \subset A$ is a proper non-zero ideal, then we can find $n \geq 1$ such that $\mathfrak{p}^n \subset \mathfrak{a} \subset \mathfrak{p}$. Indeed, if not then we can find a proper ideal $\mathfrak{a} \subset A$ maximal with respect to the property that it contains no $\mathfrak{p}^n$. Then $\mathfrak{a} \neq \mathfrak{p}$, so we can find $a, b \in A$ such that $ab \in \mathfrak{a}$ but $a \notin \mathfrak{a}$ and $b \notin \mathfrak{a}$. The inclusions $\mathfrak{a} \subset \mathfrak{a} + (a)$ and $\mathfrak{a} \subset \mathfrak{a} + (b)$ are then proper, so we can find $m \geq 1$ such that $\mathfrak{p}^m \subset \mathfrak{a} + (a)$ and $\mathfrak{p}^m \subset \mathfrak{a} + (b)$, hence $\mathfrak{p}^{2m} \subset (\mathfrak{a} + (a))(\mathfrak{a} + (b)) \subset \mathfrak{a}$, a contradiction.

Let us therefore choose a non-zero element $y \in \mathfrak{p}$, and $n \geq 1$ such that $\mathfrak{p}^n \subset (y) \subset \mathfrak{p}$. We can choose $n \geq 1$ minimal with respect to this property. We are going to show that $\mathfrak{p}$ is principal. If $n = 1$, then we're done. Otherwise $n \geq 2$, and we have $\mathfrak{p}^{n-1} \not\subset (y)$, and we can choose $x \in \mathfrak{p}^{n-1} - (y)$. Let $z = x/y \in K$. We have $x\mathfrak{p} \subset \mathfrak{p}^n \subset (y)$, hence $z\mathfrak{p} \subset A$ is an ideal. If $z\mathfrak{p} \subset \mathfrak{p}$, then $A[z]$ injects into $\mathrm{End}_A(\mathfrak{p})$, a finitely generated A-module (since A is Noetherian). Thus z is integral over A, hence $z \in A$, hence $x = yz \in (y)$, a contradiction. Thus $z\mathfrak{p} = A$, and $\mathfrak{p} = z^{-1}A$ and $\mathfrak{p}$ is principal.

Let $\pi = z^{-1}$. We claim that every element $a \in A$ admits a unique expression $a = u\pi^n$ with $u \in A^\times$, which will show that A is a PID. The uniqueness is clear. To show existence, it is enough to show that $\mathfrak{q} = \cap_{n \geq 1}(\pi^n A) = 0$ (as then $A = \cup_{n \geq 0}(\pi^n A - \pi^{n+1} A)$). However, we have $\pi\mathfrak{q} = \mathfrak{q}$, hence $\mathfrak{q} = 0$ (by Nakayama's lemma). This completes the proof. □

We are going to study rings which are 'locally DVRs'. We first need to define localization. Let A be a ring. A multiplicative subset $S \subset A$ is by definition a subset containing 1, not containing 0, and which is closed under multiplication.

Definition 1.6. *We define $S^{-1}A$ to be the set of equivalence classes of pairs (a, s) with $a \in A$, $s \in S$, with respect to the equivalence relation $(a, s) \sim (b, t)$ if there exists $u \in S$ such that $u(at - bs) = 0$. We represent the equivalence class of (a, s) using the symbol a/s.*

Lemma 1.7.
1. *This is an equivalence relation, and $S^{-1}A$ becomes a ring with the operations $a/s + b/t = (at + sb)/(st)$ and $a/s \cdot b/t = (ab)/(st)$.*
2. *The map $A \to S^{-1}A$, $a \mapsto a/1$ is a ring homomorphism with kernel $\{a \in A \mid \exists s \in S, sa = 0\}$.*
3. *If A is an integral domain, then $\mathrm{Frac}\, A = (A - \{0\})^{-1}A$ and $S^{-1}A$ is identified with a subring of $\mathrm{Frac}\, A$.*

Proof. The symmetry and reflexivity of the relation is clear. To show transitivity, suppose given (a, s), (a', s') and (a'', s''), together with $u, u' \in S$ such that $u(as' - a's) = 0$ and $u'(a's'' - a''s') = 0$. We then have $uu's'as'' = u's''ua's = usu'a''s'$, showing that $a/s = a''/s''$ in $S^{-1}A$. The ring axioms are easily verified.

It is clear that the map $A \to S^{-1}A$ is a ring homomorphism. We have $a/1 = 0/1$ if and only if there exists $s \in S$ such that $as = 0$, as in the statement of the lemma.

If A is an integral domain, then $a/s = a'/s'$ if and only if $as' = a's$ (because A has no zero divisors). This shows that $(A - \{0\})^{-1}A = \operatorname{Frac} A$, by definition, and that the map $S^{-1}A \to \operatorname{Frac} A$, $a/s \mapsto a/s$, is a well-defined injection. □

A common choice of S is when $S = \{1, f, f^2, \dots\}$ for some $f \in A$ which is not nilpotent; then we write $S^{-1}A = A[1/f]$. Another common choice is $S = A - \mathfrak{p}$, where $\mathfrak{p}$ is a prime ideal of A. In fact, an ideal $I \subset A$ is prime if and only if $A - I$ is a multiplicative subset, as follows from the definitions.

Definition 1.8. *Let A be a ring with multiplicative subset S, and let M be an A-module. We define the localized module $S^{-1}M$ to be the set of equivalence classes of pairs (m, s) with $m \in M$ and $s \in S$. Two pairs (m, s) and (m', s') are said to be equivalent if there exists $u \in S$ such that $u(s'm - sm') = 0$. The equivalence class of the pair (m, s) is denoted using the symbol m/s.*

Again, one shows that $S^{-1}M$ is an $S^{-1}A$-module, with multiplication $a/s \cdot m/t = (am)/(st)$. Localization is a functor: if M, N are A-modules and $f : M \to N$ is a homomorphism of A-modules, then there is a natural map $S^{-1}f : M \to N$ given by $S^{-1}f(m/s) = f(m)/s$. Similarly, if there is a ring homomorphism $g : A \to B$, then $S^{-1}B$ is naturally a ring and the map $S^{-1}g : S^{-1}A \to S^{-1}B$ is a ring homomorphism. An important feature of localization is its exactness:

Lemma 1.9. *Let $M' \xrightarrow{f} M \xrightarrow{g} M''$ be an exact sequence of A-modules (i.e. $\ker g =$ im f). Then the sequence $S^{-1}M' \xrightarrow{S^{-1}f} S^{-1}M \xrightarrow{S^{-1}g} S^{-1}M''$ is exact.*

Proof. Since localization is functorial, we have $S^{-1}g \circ S^{-1}f = S^{-1}(gf) = 0$. This shows that im $S^{-1}f \subset \ker S^{-1}g$. For the other inclusion, suppose $S^{-1}g(m/s) = g(m)/s = 0$. Then there exists $u \in S$ such that $ug(m) = g(um) = 0$, hence $m' \in M'$ such that $f(m') = um$, hence $S^{-1}f(m'/us) = um/us = m/s$. □

This shows, for example, that if $\mathfrak{a} \subset A$ is an ideal, then $S^{-1}\mathfrak{a}$ can be naturally identified with an ideal of $S^{-1}A$ (which is the ideal $\mathfrak{a}S^{-1}A$ generated by $\mathfrak{a}$).

Lemma 1.10. *Let A be a ring, and let $f : A \to S^{-1}A$ be the natural map. The maps $\mathfrak{p} \mapsto \mathfrak{p} \cdot S^{-1}A$, $\mathfrak{q} \mapsto f^{-1}(\mathfrak{p})$ define inclusion-preserving bijections between the set of prime ideals of A which do not intersect S and the set of prime ideals of $S^{-1}A$.*

Proof. We first check that $\mathfrak{p} \cdot S^{-1}A$ is indeed a prime ideal. The map $A \to A/\mathfrak{p}$ gives a map $S^{-1}A \to S^{-1}(A/\mathfrak{p})$, and the latter is non-zero since $A/\mathfrak{p}$ is non-zero and no element is killed by an element of S. An element a/s is in the kernel if and only if there exists $u \in S$ such that $ua = 0$ in $A/\mathfrak{p}$, if and only if $a \in \mathfrak{p}$. Thus $\mathfrak{p} \cdot S^{-1}A$ is the kernel of this map, and is a prime ideal of $S^{-1}A$.

We now show that these maps are mutually inverse. If $\mathfrak{p} \subset S^{-1}A$ is a prime ideal, then $ff^{-1}(\mathfrak{p}) \subset \mathfrak{p}$, hence $f^{-1}(\mathfrak{p}) \cdot S^{-1}A \subset \mathfrak{p}$. On the other hand, if $a/s \in \mathfrak{p}$ then $a/1 \in \mathfrak{p}$, hence $a \in f^{-1}(\mathfrak{p})$ and $a/s \in f^{-1}(\mathfrak{p}) \cdot S^{-1}A$. This shows one direction. In the other direction, suppose that $\mathfrak{q} \subset A$ is a prime ideal disjoint from S. Then $\mathfrak{q} \subset f^{-1}(\mathfrak{q} \cdot S^{-1}A)$. On the other hand, if $a \in A$ and $a/1 \in \mathfrak{q} \cdot S^{-1}A$, then $a/1 = qb/s$ for some $q \in \mathfrak{q}$, $b \in A$, $s \in S$, and hence there exists $u \in S$ such that $u(as - qb) = 0$, hence $uas = qbu \in \mathfrak{q}$. Since $\mathfrak{q}$ is prime and disjoint from S, we conclude $a \in \mathfrak{q}$, showing that $f^{-1}(\mathfrak{q} \cdot S^{-1}A) \subset \mathfrak{q}$, as required. □

This shows that if $\mathfrak{p} \subset A$ is a prime ideal and $S = A - \mathfrak{p}$, then $S^{-1}A$ has a unique maximal ideal, namely $S^{-1}\mathfrak{p}$, and is therefore a local ring. We will write $A_\mathfrak{p} = (A - \mathfrak{p})^{-1}A$ for this localization; its residue field is naturally identified with $\mathrm{Frac}(A/\mathfrak{p})$.

Proposition 1.11. *Let A be a Noetherian domain with field of fractions K. Then the following are equivalent:*

1. *For all non-zero prime ideals $\mathfrak{p} \subset A$, $A_\mathfrak{p}$ is a discrete valuation ring.*
2. *A is integrally closed and every non-zero prime ideal $\mathfrak{p} \subset A$ is a maximal ideal.*

A ring A satisfying these equivalent conditions is called a Dedekind domain*.*

Proof. First suppose that A satisfies the first condition. If $\mathfrak{q} \subset \mathfrak{p}$ are prime ideals of A, then let $S = A - \mathfrak{p}$. We have $S^{-1}\mathfrak{q} \subset S^{-1}\mathfrak{p}$, hence either $\mathfrak{q} = 0$ or $\mathfrak{q} = \mathfrak{p}$. This shows that if $\mathfrak{p} \neq 0$ then $\mathfrak{p}$ is maximal. If $x \in K$ is integral over A, then it lies in $A_\mathfrak{p}$ for every non-zero prime $\mathfrak{p} \subset A$, hence we can write $x = a_\mathfrak{p}/s_\mathfrak{p}$, where $a_\mathfrak{p} \in A$ and $s_\mathfrak{p} \in A - \mathfrak{p}$ depend on $\mathfrak{p}$. The ideal $\mathfrak{a}$ generated by all $s_\mathfrak{p}$ is the unit ideal, because it is contained in no non-zero prime ideal of A; hence we can write $1 = \sum_\mathfrak{p} s_\mathfrak{p} t_\mathfrak{p}$, where only finitely many $t_\mathfrak{p}$ are non-zero. We then have $x = \sum x s_\mathfrak{p} t_\mathfrak{p} = \sum a_\mathfrak{p} t_\mathfrak{p} \in A$, showing that A is integrally closed in K.

Now suppose instead that A satisfies the second condition. If $\mathfrak{p} \subset A$ is a non-zero prime ideal, then the ring $A_\mathfrak{p}$ is a local ring with a unique non-zero prime ideal, and we must show (by the previous proposition) that it is integrally closed. If $a/s \in A_\mathfrak{p}$ satisfies an equation $(a/s)^n + a_1/s_1 \cdot (a/s)^{n-1} + \cdots + a_n/s_n = 0$, then the element $a(s_1 \dots s_n)$ is integral over A (by clearing denominators), so lies in A. Then $a/s = (as_1 \dots s_n)/(ss_1 \dots s_n)$ lies in $A_\mathfrak{p}$, as required. □

If A is an integral domain with field of fractions K, then a fractional ideal of A is by definition a finitely generated A-submodule of K.

Lemma 1.12. *Let A be a Noetherian integral domain with field of fractions K and let $S \subset A$ be a multiplicative subset. Let $\mathfrak{a}, \mathfrak{b} \subset K$ be fractional ideals. Then:*

1. *$(S^{-1}\mathfrak{a})(S^{-1}\mathfrak{b}) = S^{-1}(\mathfrak{a}\mathfrak{b})$ and $S^{-1}\mathfrak{a} + S^{-1}\mathfrak{b} = S^{-1}(\mathfrak{a} + \mathfrak{b})$.*
2. *Define $(\mathfrak{a} : \mathfrak{b}) = \{x \in K \mid x\mathfrak{b} \subset \mathfrak{a}\}$. Then $(\mathfrak{a} : \mathfrak{b})$ is a fractional ideal of A and $S^{-1}(\mathfrak{a} : \mathfrak{b}) = (S^{-1}\mathfrak{a} : S^{-1}\mathfrak{b})$.*

Proof. Suppose that $\mathfrak{b} = (b_1, \dots, b_n)$, where the b_i are non-zero. Then $(\mathfrak{a} : \mathfrak{b}) = \cap_{i=1}^n b_i^{-1}\mathfrak{a}$, which shows that $(\mathfrak{a} : \mathfrak{b})$ is finitely generated (as A is Noetherian), hence a fractional ideal of A. The other assertions follow easily from the fact that localization commutes with finite sums and intersections of submodules. □

Proposition 1.13. *Let A be a Dedekind domain. Then the set I of non-zero fractional ideals of A forms a group under multiplication.*

Proof. Multiplication of fractional ideals gives an associative composition law in I, with identity $(1) = A$. It remains to show the existence of inverses. We claim that for any non-zero fractional ideal $\mathfrak{a} \subset A$, we have $\mathfrak{a} \cdot (A : \mathfrak{a}) = A$. Since $(A : \mathfrak{a})$ is a fractional ideal, this will be enough. If A is a DVR then we have $\mathfrak{a} = \pi^n A$ for some $n \in \mathbb{Z}$, and $(A : \mathfrak{a}) = \pi^{-n}A$; so this is clear.

In general, we have for any non-zero prime $\mathfrak{p} \subset A$,

$$(\mathfrak{a}(A : \mathfrak{a}))_{\mathfrak{p}} = \mathfrak{a}_{\mathfrak{p}}(A : \mathfrak{a})_{\mathfrak{p}} = A_{\mathfrak{p}},$$

by the lemma. It therefore suffices to show that for any fractional ideal $\mathfrak{b} \subset K$, if $\mathfrak{b}_{\mathfrak{p}} = A_{\mathfrak{p}}$ for all non-zero prime ideals $\mathfrak{p}$ of A, then $\mathfrak{b} = A$. Let us show the more general statement that if $\mathfrak{b}, \mathfrak{c}$ are fractional ideals such that $\mathfrak{b}_{\mathfrak{p}} \subset \mathfrak{c}_{\mathfrak{p}}$ for all $\mathfrak{p}$, then $\mathfrak{b} \subset \mathfrak{c}$.

It is enough to show $\mathfrak{b} \subset \mathfrak{c}$. Let $b \in \mathfrak{b}$. For all non-zero prime ideals $\mathfrak{p} \subset A$, we can find an expression $b = c_{\mathfrak{p}}/s_{\mathfrak{p}}$ with $c_{\mathfrak{p}} \in \mathfrak{c}$ and $s_{\mathfrak{p}} \in A - \mathfrak{p}$. The ideal $(s_{\mathfrak{p}})$ is the unit ideal, since it can be contained in no maximal ideal of A; we can therefore find non-zero prime ideals $\mathfrak{p}_1, \dots, \mathfrak{p}_n$ and $a_i \in A$ such that $1 = \sum_{i=1}^n a_i s_{\mathfrak{p}_i}$. Multiplying through, we obtain

$$b = \sum_{i=1}^n a_i b s_{\mathfrak{p}_i} = \sum_{i=1} a_i c_{\mathfrak{p}_i} \in \mathfrak{c},$$

as desired. □

Lemma 1.14. *Let A be a Dedekind domain. Then for each non-zero ideal $\mathfrak{a} \subset A$, there are only finitely many non-zero prime ideals $\mathfrak{p} \subset A$ such that $\mathfrak{a} \subset \mathfrak{p}$.*

Proof. We can assume that $\mathfrak{a} = (x)$ is principal. The set of ideals $\mathfrak{b}$ of A containing x satisfies the descending chain condition. Indeed, if $\mathfrak{b}_1 \supset \mathfrak{b}_2 \supset \cdots \supset (x)$ then the chain $\mathfrak{b}_1^{-1} \subset \mathfrak{b}_2^{-1} \subset \cdots \subset x^{-1}A$ is eventually stationary, since A is Noetherian. This implies the same for the original chain.

Suppose therefore that there are infinitely many distinct primes $\mathfrak{p}_1, \mathfrak{p}_2, \dots$ containing x. The sequence $\mathfrak{p}_1, \mathfrak{p}_1 \cap \mathfrak{p}_2, \dots$ is eventually stationary; suppose that it becomes stationary at step k. Then we have

$$\mathfrak{p}_{k+1} \supset \mathfrak{p}_1 \cap \cdots \cap \mathfrak{p}_{k+1} = \mathfrak{p}_1 \cap \cdots \cap \mathfrak{p}_k \supset \mathfrak{p}_1 \dots \mathfrak{p}_k.$$

We can find for each $i = 1, \dots, k$ an element $x_i \in \mathfrak{p}_i - \mathfrak{p}_{k+1}$. Then the element $x_1 \dots x_k \in \mathfrak{p}_1 \dots \mathfrak{p}_k$ must lie in $\mathfrak{p}_{k+1}$. Since the ideal $\mathfrak{p}_{k+1}$ is prime, this forces some x_i to lie in $\mathfrak{p}_{k+1}$, a contradiction. □

If A is a Dedekind domain, then for each non-zero prime $\mathfrak{p} \subset A$, we have the valuation $v_\mathfrak{p} : K^\times \to \mathbb{Z}$ which is associated to the DVR $A_\mathfrak{p}$. If $\mathfrak{a} \subset K$ is a non-zero fractional ideal, then we have $\mathfrak{a}_\mathfrak{p} = aA_\mathfrak{p}$ for some $a \in K^\times$, and we define $v_\mathfrak{p}(\mathfrak{a}) = v_\mathfrak{p}(a)$. This defines a surjective homomorphism $v_\mathfrak{p} : I \to \mathbb{Z}$.

Proposition 1.15. *Let A be a Dedekind domain, and let $\mathfrak{a} \subset K$ be a non-zero fractional ideal.*

1. *We have $v_\mathfrak{p}(\mathfrak{a}) = 0$ for all but finitely many $\mathfrak{p}$.*

2. *We have $\mathfrak{a} = \prod_\mathfrak{p} \mathfrak{p}^{v_\mathfrak{p}(\mathfrak{a})}$, where the product has only finitely many terms.*

Proof. We recall that if $\mathfrak{a}, \mathfrak{b} \subset K$ are fractional ideals, then $\mathfrak{a} \subset \mathfrak{b}$ if and only if $\mathfrak{a}_\mathfrak{p} \subset \mathfrak{b}_\mathfrak{p}$ for all $\mathfrak{p}$. If $\mathfrak{p} \neq \mathfrak{q}$ are non-zero prime ideals of A, then $\mathfrak{p}_\mathfrak{q} = A_\mathfrak{q}$. It follows that if $\mathfrak{a} \subset A$ is a non-zero ideal, then $\mathfrak{a} \subset \mathfrak{p}$ if and only if $\mathfrak{a}_\mathfrak{p} \subset \mathfrak{p}_\mathfrak{p}$, if and only if $v_\mathfrak{p}(\mathfrak{a}) > 0$. By the previous lemma, this can be true for only finitely many primes $\mathfrak{p}$. If $\mathfrak{a} \subset K$ is a non-zero fractional ideal, then we can find $x \in A - \{0\}$ such that $x\mathfrak{a} \subset A$ is an ideal. Then $v_\mathfrak{p}(x) = 0$ for all but finitely many $\mathfrak{p}$, and $v_\mathfrak{p}(x\mathfrak{a}) = 0$ for all but finitely many $\mathfrak{p}$. Since $v_\mathfrak{p}(x\mathfrak{a}) = v_\mathfrak{p}(x) + v_\mathfrak{p}(\mathfrak{a})$ for all non-zero prime ideals of A, this shows that $v_\mathfrak{p}(\mathfrak{a}) = 0$ for all but finitely many $\mathfrak{p}$.

The two fractional ideals $\mathfrak{a}$ and $\prod_\mathfrak{p} \mathfrak{p}^{v_\mathfrak{p}(\mathfrak{a})}$ have the same localizations, so are the same. □

Corollary 1.16. *Let A be a Dedekind domain. Then every non-zero ideal $\mathfrak{a} \subset A$ admits a unique expression $\mathfrak{a} = \prod_{i=1}^n \mathfrak{p}_i^{a_i}$, where the $\mathfrak{p}_i$ are pairwise distinct prime ideals of A. This expression is uniquely determined up to re-ordering of terms.*

2 Extensions of Dedekind domains

Lemma 2.1. *Let E/K be a finite separable extension of fields. Then the K-bilinear form $S : E \times E \to K$ given by the formula $S(x, y) = \mathrm{tr}_{E/K}(xy)$ is non-degenerate.*

Proof. We must show that for any non-zero $x \in E$, there exists $y \in E$ such that $\mathrm{tr}_{E/K}(xy) \neq 0$. It is clearly enough to show that $\mathrm{tr}_{E/K} \neq 0$ as a homomorphism $E \to K$. If L/K denotes the Galois closure of E, then there are exactly n distinct K-embeddings $\sigma_1, \dots, \sigma_n : E \hookrightarrow L$; and we have $\mathrm{tr}_{E/K}(x) = \sigma_1(x) + \cdots + \sigma_n(x)$. The lemma therefore follows from the fact that the embeddings $\sigma_1, \dots, \sigma_n$ are linearly independent over L (even as homomorphisms $E^\times \to L^\times$). □

Let A be a Dedekind domain with field of fractions K, and let E/K be a finite separable extension. Let B denote the integral closure of A in E.

Proposition 2.2. *B is a finitely generated A-module, which spans E as a K-vector space. It is a Dedekind domain.*

Proof. Let $e_1, \dots, e_n$ be a K-basis for E. After multiplying through by elements of A, we can suppose that each $e_i \in B$. The map $S(x, y) = \operatorname{tr}_{E/K}(xy)$ defines a non-degenerate K-bilinear pairing $S : E \times E \to K$. Let us note that if $z \in B$, then $\operatorname{tr}_{E/K} z$ is integral over A and lies in K, hence lies in A.

Let $f_1, \dots, f_n \in E$ be the dual K-basis with respect to S, so that $S(e_i, f_j) = \delta_{ij}$. Choose $c \in A - \{0\}$ such that each $cf_j \in B$. We claim that $B \subset A \cdot c^{-1}e_1 \oplus \cdots \oplus A \cdot c^{-1}e_n$. If $z \in B$, then each zcf_j is integral over A, so $\operatorname{tr}_{E/K}(zcf_j) \in A$. Writing $z = \sum_i r_i e_i$ with $r_i \in K$, we therefore have $S(z, cf_j) = cr_j \in A$, hence $r_j \in c^{-1}A$. This shows the claim. It follows that B is contained in a finitely generated A-module, hence is a finitely generated A-module as A is Noetherian, hence is a finitely generated A-algebra, hence is a Noetherian ring, by the Hilbert basis theorem.

Let $\mathfrak{q} \subset B$ be a non-zero prime ideal, and let $\mathfrak{p} = \mathfrak{q} \cap A$. Then $\mathfrak{p} \subset A$ is prime. It is also non-zero: if $b \in \mathfrak{q}$ is a non-zero element, then we can find an equation of minimal degree with $a_i \in A$:

$$b^n + a_1 b^{n-1} + \cdots + a_n = 0.$$

Then $a_n \in \mathfrak{q} \cap A = \mathfrak{p}$ and $a_n \neq 0$, by minimality. There is an injective homomorphism $A/\mathfrak{p} \hookrightarrow B/\mathfrak{q}$, $A/\mathfrak{p}$ is a field, and $B/\mathfrak{q}$ is a finite $A/\mathfrak{p}$-algebra and a domain, hence a field. This shows that $\mathfrak{q}$ is maximal, hence B is a Dedekind domain. □

The ring $\mathbb{Z}$ of rational integers is a PID, so is a Dedekind domain. We deduce:

Corollary 2.3. *Let K be a number field. Then the ring of integers $\mathcal{O}_K$ is a Dedekind domain. In particular, any non-zero ideal $\mathfrak{a} \subset \mathcal{O}_K$ admits a factorization $\mathfrak{a} = \prod_{i=1}^m \mathfrak{p}_i^{a_i}$, where the $\mathfrak{p}_i$ are pairwise distinct prime ideals and a_i non-negative integers, and this factorization is unique up to re-ordering.*

3 Complete discrete valuation rings

Definition 3.1. *Suppose given for each $i = 1, 2, \dots$ a group A_i and a homomorphism $f_i : A_{i+1} \to A_i$. The inverse limit is by definition*

$$\varprojlim_i A_i = \{(a_i) \in \prod_{i=1}^{\infty} A_i \mid \forall i \geq 1, f_i(a_{i+1}) = a_i\}.$$

It is a group. If the A_i are all abelian groups (resp. ring) and the f_i are homomorphisms of abelian groups (resp. rings) then it is naturally an abelian group (resp. ring).

Definition 3.2. *Let A be a discrete valuation ring. We say that A is complete if the natural map $A \to \varprojlim_i A/\mathfrak{m}_A^i$ is an isomorphism.*

The terminology is justified by the following lemma.

Lemma 3.3. *Let A be a discrete valuation ring with fraction field K and valuation $v : K^\times \to \mathbb{Z}$. Then the following are equivalent:*

1. *A is complete.*

2. *K is a complete metric space with respect to the metric $d(x, y) = 2^{-v(x-y)}$.*

Proof. Suppose first that A is complete, and let $(x_i)_{i=0}^{\infty}$ be a Cauchy sequence in K. This is equivalent to the following condition: for all $M > 0$, there exists $N > 0$ such that $v(x_i - x_j) \geq M$ for all $i, j \geq M$. In particular, we can assume after discarding finitely many elements that $v(x_i - x_j) \geq 0$ for all $i, j \geq 0$. Replacing x_i by $x_i - x_1$, we can therefore assume that $v(x_i) \geq 0$ for all i, i.e. $x_i \in A$ for all i.

After passing to a subsequence, we can assume that $v(x_i - x_{i+1}) \geq i$ for all i, or in other words $x_{i+1} \equiv x_i \bmod \mathfrak{m}_A^i$. It follows that $(x_i)_i \in \varprojlim_i A/\mathfrak{m}_A^i$, hence there exists $x \in A$ such that $v(x - x_i) \geq i$ for all $i \geq 0$. This is the desired limit.

Now suppose instead that K is complete. The map $A \to \varprojlim_i A/\mathfrak{m}_A^i$ is injective, so we must show it is surjective. Let $(x_i)_{i=0}^{\infty} \in \varprojlim_i A/\mathfrak{m}_A^i$. Then for all $i, j \geq N$ we have $v(x_i - x_j) \geq N$, hence (x_i) is a Cauchy sequence, hence there exists $x \in K$ such that $v(x - x_i) \to \infty$ as $i \to \infty$. In particular, we have $v(x) \geq 0$, hence $x \in A$ and x maps to $(x_i)_{i=0}^{\infty}$. This completes the proof. □

Proposition 3.4. *Let A be a DVR with uniformizer π, and define $\widehat{A} = \varprojlim_i A/\mathfrak{m}_A^i$. Then;*

1. *$\widehat{A}$ is a complete DVR with uniformizer π.*

2. *For each $i \geq 0$, the natural map $A/\mathfrak{m}_A^i \to \widehat{A}/\mathfrak{m}_{\widehat{A}}^i$ is an isomorphism.*

3. *Let $X \subset A$ be a set of representatives for the residue field k_A containing 0. Then every element $x \in \widehat{A}$ admits a unique π-adic expansion*

$$x = \sum_{i=0}^{\infty} a_i \pi^i$$

with $a_i \in X$ for each $i \geq 0$, and conversely every such expansion defines an element of $\widehat{A}$.

Proof. We prove the last part first. We observe that the map $A/\pi \to \pi^i A/\pi^{i+1}A$ given by multiplication by π^i is an isomorphism. It follows that each element of the ring A/π^{i+1} admits a unique expression of the form $\sum_{j=0}^{i} a_j \pi^j$ with $a_j \in X$. The map $A/\pi^{i+1} \to A/\pi^i$ corresponds to throwing away the term $a_i \pi^i$. Since an element of $\widehat{A}$ is a compatible system of elements of the quotients A/π^{i+1}, we see that the elements of $\widehat{A}$ are in bijection with the expressions $\sum_{j=0}^{\infty} a_j \pi^j$ with $a_j \in X$.

By definition, we have for each $i \geq 0$ a surjection $\widehat{A} \to A/\pi^i A$. An element $x = \sum_j a_j \pi^j$ lies in the kernel if and only if all of its digits $a_0, \ldots, a_{i-1}$ are 0. In this case, we have $x = \pi^i \sum_{j=0}^{\infty} a_{j+i} \pi^j$, showing that the kernel of this map is $\pi^i \widehat{A}$, and hence we have an isomorphism $\widehat{A}/\pi^i \widehat{A} \cong A/\pi^i A$.

The natural map $A \to \widehat{A}$ is injective. Every element $x = \sum_j a_j \pi^j$ with $a_0 \neq 0$ is a unit. Indeed, we write $x = a_0(1 - \pi y)$ for some $y \in \widehat{A}$. The element a_0 is a unit in A, hence in $\widehat{A}$, and $(1 - \pi y)$ has inverse given by

$$1 + \pi y + (\pi y)^2 + (\pi y)^3 + \dots,$$

this series converging in $\widehat{A}$. This shows that every non-zero element of $\widehat{A}$ has a unique expression as $u\pi^n$ with $u \in \widehat{A}^\times$ and $n \geq 0$, and hence that $\widehat{A}$ is a DVR.

To complete the proof, we show that $\widehat{A}$ is complete. However, we have

$$\widehat{A} = \varprojlim_i A/\pi^i A \cong \varprojlim_i \widehat{A}/\pi^i \widehat{A}$$

by the second part of the proposition, so we're done. □

We refer to the ring $\widehat{A}$ as the completion of A; the proof of the proposition shows that if A is complete, then $A \cong \widehat{A}$.

4 The p-adic numbers

We can now study our first new example of a complete discrete valuation ring.

Definition 4.1. *Let p be a prime. The ring of p-adic integers $\mathbb{Z}_p$ is by definition the completion of $\mathbb{Z}_{(p)}$. The field of p-adic numbers $\mathbb{Q}_p$ is its fraction field.*

Thus $\mathbb{Z}_p$ is a complete DVR with residue field $\mathbb{Z}_p/(p) \cong \mathbb{Z}/p\mathbb{Z} = \mathbb{F}_p$. Every element of $\mathbb{Z}_p$ admits a unique expression

$$a_0 + a_1 p + a_2 p^2 + \dots$$

with $a_i \in \{0, \dots, p-1\}$. Multiplication and addition is done in the same way as for formal power series, except we now need to 'carry' digits.

The following lemma is one example of a family of results which are referred to by the name 'Hensel's lemma'. The key idea is that in complete DVRs, one can efficiently solve equations by successive approximation.

Lemma 4.2. *Let A be a complete DVR with valuation v, and let $f(X) \in A[X]$ be a monic polynomial. Suppose there exists $x \in A$ such that $v(f(x)) > 2v(f'(x))$. Then there exists a unique $\alpha \in A$ such that $f(\alpha) = 0$ and $v(\alpha - x) > v(f'(x))$.*

Proof. We use Newton's approximation. Let $a_1 = x$, and define a sequence inductively by the formula

$$a_{n+1} = a_n - \frac{f(a_n)}{f'(a_n)}.$$

Let $t = v(f(a_1)/f'(a_1)^2) > 0$. We will show by induction on $n \geq 1$ that the following conditions are satisfied:

- $v(a_n) \geq 0$.
- $v(f'(a_n)) = v(f'(a_1))$.
- $v(f(a_n)) \geq 2v(f'(a_n)) + 2^{n-1}t$.

These conditions hold for $n = 1$, by hypothesis. We now treat the induction step. We have $v(a_{n+1}) \geq 0$ if and only if $f(a_n)/f'(a_n) \in A$. We have

$$v(f(a_n)) - v(f'(a_n)) \geq v(f'(a_n)) + 2^{n-1}t = v(f'(a_1)) + 2^{n-1}t \geq 0.$$

We have $f'(a_{n+1}) - f'(a_n) = (a_{n+1} - a_n)y$ for some $y \in A$, hence

$$v(f'(a_{n+1}) - f'(a_n)) \geq v(a_{n+1} - a_n) = v(f(a_n)) - v(f'(a_n)) > v(f'(a_n)),$$

hence $v(f'(a_{n+1})) = v(f'(a_n))$. We can write (using the Taylor expansion of the polynomial $f(X)$)

$$f(a_{n+1}) = f(a_n) - \frac{f(a_n)}{f'(a_n)} f'(a_n) + \left(\frac{f(a_n)}{f'(a_n)} \right)^2 z$$

for some $z \in A$. We thus have

$$v(f(a_{n+1})) \geq 2(v(f(a_n)) - v(f'(a_n))) \geq 2v(f'(a_n)) + 2^n t.$$

This completes the induction. Incidentally we have shown $v(a_{n+1} - a_n) \geq v(f'(a_1)) + 2^{n-1}t$, which shows that (a_n) is a Cauchy sequence which has a limit $\alpha \in A$, and $f(\alpha) = 0$.

To establish uniqueness, we suppose $\beta \in A$ is another root with $f(\beta) = 0$ and $v(\beta - x) > v(f'(x))$. We write $\beta = \alpha + h$ and calculate

$$f(\beta) = 0 = f(\alpha) + hf'(\alpha) + h^2 w$$

for some $w \in A$, hence $hf'(\alpha) = -h^2 w$, hence (assuming $h \neq 0$) $v(f'(\alpha)) \geq v(h)$. On the other hand, we have

$$v(h) = v(\beta - \alpha) \geq \min(v(\alpha - x), v(\beta - x)) > v(f'(x)) = v(f'(\alpha)).$$

This contradiction shows that the root α is unique, as desired. □

Corollary 4.3. *Let A be a complete DVR with valuation v, and let $f(X) \in A[X]$ be a monic polynomial. Let $\overline{f}(X) \in k_A[X]$ be the reduction of $f(X)$ modulo $\mathfrak{m}_A$. Suppose there exists $x \in k_A$ such that $\overline{f}(x) = 0$ and $\overline{f}'(x) \neq 0$. Then there exists a unique $y \in A$ such that $f(y) = 0$ and $y \equiv x \bmod \mathfrak{m}_A$.*

Example 4.4. We can use this to understand which elements of $\mathbb{Q}_p$ are squares. Any non-zero element has a unique expression $p^n u$ with $u \in \mathbb{Z}_p^\times$, so it is equivalent to understand which elements of $\mathbb{Z}_p^\times$ are squares. A necessary condition is that $u \bmod p \in \mathbb{F}_p^\times$ is a square, so let us suppose this condition holds.

We apply Hensel's lemma to the polynomial $f(X) = X^2 - u$. Let v be an element of $\mathbb{Z}_p^\times$ such that $v^2 \equiv u \bmod p$. If p is odd, then $f'(v) = 2v$ is a p-adic unit, so Hensel's lemma shows that there is a unique $w \in \mathbb{Z}_p^\times$ such that $w^2 = u$.

If p is even, then a necessary condition is that $u \bmod 8 \in (\mathbb{Z}/8\mathbb{Z})^\times$ is a square, i.e. that $u \equiv 1 \bmod 8$. This being the case, we have $f(1) \equiv 0 \bmod 8$ and $f'(1) = 2$, hence $v(f(1)) \geq 3$ and $v(f'(1)) = 1$. Hensel's lemma shows in this case that there is a unique $w \in \mathbb{Z}_p^\times$ such that $w^2 = u$.

Example 4.5. The group homomorphism $\mathbb{Z}_p^\times \to \mathbb{F}_p^\times$ is surjective. We can use Hensel's lemma to construct a section to this map, i.e. a homomorphism $\tau : \mathbb{F}_p^\times \to \mathbb{Z}_p^\times$ such that $\tau(a) \bmod p = a$. Indeed, let $f(X) = X^p - X$. Then $f'(X) \equiv -1 \bmod p$, so the simple version of Hensel's lemma shows that for each $a \in \mathbb{F}_p$, there is a unique $b = \tau(a) \in \mathbb{Z}_p$ such that $b^p = b$ and $b \equiv a \bmod p$. The uniqueness implies that $\tau(aa') = \tau(a)\tau(a')$, hence τ is a group homomorphism. The map $a \mapsto \tau(a)$ is called the Teichmüller lift, and exists for any complete DVR with finite residue field.

5 Extensions of Dedekind domains, II

Let A be a Dedekind domain, let $K = \mathrm{Frac}(A)$, and let E/K be a finite separable extension, and let B denote the integral closure of A in E. We have seen that B is a Dedekind domain with field of fractions E.

Let $\mathfrak{q} \subset B$ be a non-zero prime ideal. Then $\mathfrak{q} \cap A$ is a non-zero prime ideal of A: if $b \in \mathfrak{q} - \{0\}$, then we can find an equation $b^n + a_1 b^{n-1} + \cdots + a_n = 0$ of minimal degree. This forces $a_n \neq 0$, and then $a_n \in \mathfrak{q} \cap A$.

Definition 5.1. *If $\mathfrak{q} \subset B$ is a non-zero prime ideal and $\mathfrak{p} = \mathfrak{q} \cap A$, then we say that $\mathfrak{q}$ lies above $\mathfrak{p}$.*

Lemma 5.2. *Let $\mathfrak{q} \subset B$, $\mathfrak{p} \subset A$ be non-zero prime ideals. Then the following are equivalent:*

1. *$\mathfrak{q}$ lies above $\mathfrak{p}$.*

2. *$\mathfrak{q}$ appears in the prime factorization of $\mathfrak{p}B \subset B$ (in other notation, $v_\mathfrak{q}(\mathfrak{p}B) > 0$).*

Proof. We recall that if $\mathfrak{a}, \mathfrak{b} \subset B$ are non-zero ideals, then $\mathfrak{b} = \mathfrak{a}\mathfrak{c}$ for some ideal $\mathfrak{c} \subset B$ if and only if $\mathfrak{b} \subset \mathfrak{a}$. If $\mathfrak{q}$ lies above $\mathfrak{p}$, then $\mathfrak{p}B \subset \mathfrak{q}B = \mathfrak{q}$, so $\mathfrak{q}$ divides $\mathfrak{p}B$. Conversely, if $\mathfrak{q}$ divides $\mathfrak{p}B$, then $\mathfrak{p} \subset \mathfrak{p}B \subset \mathfrak{q}$, hence $\mathfrak{p} \subset \mathfrak{q} \cap A$. Since $\mathfrak{p}$ is a maximal ideal, this forces $\mathfrak{p} = \mathfrak{q} \cap A$. □

Definition 5.3. *Suppose that $\mathfrak{q} \subset B$ is a non-zero prime ideal, and let $\mathfrak{p} = \mathfrak{q} \cap A$. Then there is a natural embedding of fields $A/\mathfrak{p} \hookrightarrow B/\mathfrak{q}$. We define the ramification index $e_{\mathfrak{q}/\mathfrak{p}} = v_\mathfrak{q}(\mathfrak{p}B)$ and the residue degree $f_{\mathfrak{q}/\mathfrak{p}} = [B/\mathfrak{q} : A/\mathfrak{p}]$.*

We say that the prime $\mathfrak{q}$ is unramified over A if $e_{\mathfrak{q}/\mathfrak{p}} = 1$ and the extension of residue fields is separable. If every prime lying above $\mathfrak{p}$ is unramified over A, then we say that $\mathfrak{p}$ is unramified in B.

Proposition 5.4. *Let $\mathfrak{p} \subset A$ be a non-zero prime ideal, and let $n = [E : K]$. Then we have $n = \sum_{\mathfrak{q}|\mathfrak{p}} e_{\mathfrak{q}/\mathfrak{p}} f_{\mathfrak{q}/\mathfrak{p}}$, the sum running over all primes of B lying over $\mathfrak{p}$.*

Proof. Let $S = (A - \mathfrak{p})$. After replacing A by $S^{-1}A$ and B by $S^{-1}B$, we can assume that A is a DVR. (Note that $S^{-1}B$ is the integral closure of $S^{-1}A$, and if $\mathfrak{q} \cap S = \emptyset$ then replacing $\mathfrak{q}$ by $S^{-1}\mathfrak{q}$ does not change e or f.) Then B is a torsion-free finitely generated A-module, so is free of some rank. We have $(A - \{0\})^{-1}B = E$, so this rank is $\dim_K E = n$.

To prove the proposition, we will calculate this rank in another way. Reducing modulo $\mathfrak{p}$, we also have $n = \dim_{A/\mathfrak{p}} B/\mathfrak{q}B$. By the Chinese remainder theorem, there is an isomorphism

$$B/\mathfrak{p}B \cong \prod_{\mathfrak{q}|\mathfrak{p}} B/\mathfrak{q}^{e_{\mathfrak{q}/\mathfrak{p}}},$$

so we get $n = \sum_{\mathfrak{q}/\mathfrak{p}} \dim_{A/\mathfrak{p}} B/\mathfrak{q}^{e_{\mathfrak{q}/\mathfrak{p}}}$. There is a filtration $B \supset \mathfrak{q} \supset \mathfrak{q}^2 \supset \cdots \supset \mathfrak{q}^{e_{\mathfrak{q}/\mathfrak{p}}}$, and each $\mathfrak{q}^i/\mathfrak{q}^{i+1}$ is a $B/\mathfrak{q}$-vector space of dimension 1, hence an $A/\mathfrak{p}$-vector space of dimension $f_{\mathfrak{q}/\mathfrak{p}}$. We obtain

$$\dim_{A/\mathfrak{p}} B/\mathfrak{q}^{e_{\mathfrak{q}/\mathfrak{p}}} = e_{\mathfrak{q}/\mathfrak{p}} f_{\mathfrak{q}/\mathfrak{p}},$$

which completes the proof. □

Now suppose that the extension E/K is Galois. In the case the group $G = \mathrm{Gal}(E/K)$ fixes A, so acts on B. In particular, if $\mathfrak{p} \subset A$ is a non-zero prime, then G acts on the set of primes $\mathfrak{q} \subset B$ lying over A by the formula $\mathfrak{q} \mapsto \sigma(\mathfrak{q})$.

Proposition 5.5. *Suppose that the extension E/K is Galois, and let $\mathfrak{p} \subset A$ be a non-zero prime ideal. Then:*

1. *Let $\mathfrak{q}$ be a prime lying over $\mathfrak{p}$, and let $\sigma \in G$. Then $f_{\sigma(\mathfrak{q})/\mathfrak{p}} = f_{\mathfrak{q}/\mathfrak{p}}$ and $e_{\sigma(\mathfrak{q})/\mathfrak{p}} = e_{\mathfrak{q}/\mathfrak{p}}$.*

2. *The group G acts transitively on the set of primes of B lying above $\mathfrak{p}$.*

3. *For any prime $\mathfrak{q}$ of B lying over $\mathfrak{p}$, we have $[E : K] = e_{\mathfrak{q}/\mathfrak{p}} f_{\mathfrak{q}/\mathfrak{p}} g_{\mathfrak{q}/\mathfrak{p}}$, where $g_{\mathfrak{q}/\mathfrak{p}}$ is the number of distinct primes of B lying over $\mathfrak{p}$.*

Proof. The first part is clear. The third part follows from combining the first and second with the previous proposition.

For the second, we can assume again that A is a DVR. Let $\mathfrak{q}, \mathfrak{q}'$ be distinct primes of B lying above $\mathfrak{p}$, and let $\pi \in B$ be a generator of $\mathfrak{q}$. Then $N_{E/K}(\pi) = \prod_{\sigma \in G} \sigma(\pi) \in \mathfrak{q} \cap A = \mathfrak{p} \subset \mathfrak{q}'$. Since $\mathfrak{q}'$ is prime there exists σinG such that $\sigma(\pi) \in \mathfrak{q}'$, hence $\sigma(\mathfrak{q}) \subset \mathfrak{q}'$, hence $\sigma(\mathfrak{q}) = \mathfrak{q}'$, as desired. □

If the extension E/K is Galois and $\mathfrak{q}$ is a non-zero prime of B lying above the prime $\mathfrak{p}$ of A, then we define $D_{\mathfrak{q}/\mathfrak{p}} = \mathrm{Stab}_G(\mathfrak{q})$, and call it the decomposition group at the prime $\mathfrak{q}$.

Proposition 5.6. *Suppose that the extension E/K is Galois, and let $\mathfrak{q}$ be a prime of B lying above the prime $\mathfrak{p}$ of A such that the corresponding extension $k_{\mathfrak{q}}/k_{\mathfrak{p}}$ of residue fields is separable. Then it is Galois, and the induced map $D_{\mathfrak{q}/\mathfrak{p}} \to \mathrm{Gal}(k_{\mathfrak{q}}/k_{\mathfrak{p}})$ is surjective.*

Proof. Let $\overline{x} \in k_{\mathfrak{q}}$, and let $x \in B$ be any element with $\overline{x} = x \bmod \mathfrak{q}$. Let $f(X) = \prod_{\sigma \in G}(X - \sigma(x)) \in A[X]$. Let $\overline{f}(X) = f(X) \bmod \mathfrak{p} \in k_{\mathfrak{p}}[X]$. By construction, the polynomial $\overline{f}(X)$ has $\overline{x}$ as a root and splits into linear factors in $k_{\mathfrak{q}}[X]$. Since $\overline{x}$ was arbitrary, this shows that $k_{\mathfrak{q}}$ is normal, hence Galois over $k_{\mathfrak{p}}$.

To show that the map $D_{\mathfrak{q}/\mathfrak{p}} \to \operatorname{Gal}(k_{\mathfrak{q}}/k_{\mathfrak{p}})$ is surjective, let $\overline{x} \in k_{\mathfrak{q}}$ be a primitive element; this exists since the extension $k_{\mathfrak{q}}/k_{\mathfrak{p}}$ is separable. We can assume that $\overline{x}$ is not 0 (since otherwise $k_{\mathfrak{q}} = k_{\mathfrak{p}}$ and the proposition is trivial). By the Chinese remainder theorem, we can choose an element $x \in B$ lifting $\overline{x}$ such that $x \in \mathfrak{q}'$ if $\mathfrak{q}' \neq \mathfrak{q}$ is any other prime of B lying over the prime $\mathfrak{p}$. Form the polynomial $f(X)$ as before; we now have $\overline{f}(X) = X^d \prod_{\sigma \in D_{\mathfrak{q}/\mathfrak{p}}}(X - \sigma(\overline{x}))$ for some $d \geq 0$.

If $\tau \in \operatorname{Gal}(k_{\mathfrak{q}}/k_{\mathfrak{p}})$, then $\tau(\overline{x})$ is a non-zero root of $\overline{f}(X)$, hence there exists $\sigma \in D_{\mathfrak{q}/\mathfrak{p}}$ such that $\sigma(\overline{x}) = \tau(\overline{x})$. Since $\overline{x}$ is a primitive element, this shows that the image of σ in $\operatorname{Gal}(k_{\mathfrak{q}}/k_{\mathfrak{p}})$ equals τ, as desired. □

Definition 5.7. *In the situation of the proposition, we call* $I_{\mathfrak{q}/\mathfrak{p}} = \ker(D_{\mathfrak{q}/\mathfrak{p}} \to \operatorname{Gal}(k_{\mathfrak{q}}/k_{\mathfrak{p}}))$ *the inertia group at the prime* $\mathfrak{q}$.

Observe that we have $\#I_{\mathfrak{q}/\mathfrak{p}} = e_{\mathfrak{q}/\mathfrak{p}}$; in particular, the inertia group is trivial if and only if the prime $\mathfrak{q}$ is unramified over $\mathfrak{p}$.

6 Extensions of complete DVRs

Theorem 6.1. *Let A be a complete DVR and let $K = \operatorname{Frac}(A)$. Let L/K be a finite separable extension, and let B denote the integral closure of A in L. Then B is a complete DVR.*

Proof. Let π be a uniformizer of A. We know that B is a Dedekind domain, and a finite free A-module. It follows that the natural map $B \to \varprojlim_i B/\pi^i B$ is an isomorphism (since this holds for A^d). On the other hand, we have $\pi B = \mathfrak{q}_1^{e_1} \dots \mathfrak{q}_r^{e_r}$, where the $\mathfrak{q}_i$ are the pairwise distinct non-zero prime ideals of B. Being a Dedekind domain with finitely many prime ideals, B is a PID. By the Chinese remainder theorem, we have for each $j \geq 0$

$$B/\pi^j B = B/\mathfrak{q}_1^{je_1} \dots \mathfrak{q}_r^{je_r} \cong \prod_{i=1}^{r} B/\mathfrak{q}_i^{je_i},$$

hence

$$B \cong \varprojlim_j B/\pi^j B \cong \prod_{i=1}^{r} \varprojlim_j B/\mathfrak{q}_i^j.$$

Since B is a subring of L, it is a domain. It follows that we must have $j = 1$, hence B has a unique non-zero prime ideal. It follows that B is a DVR. If we write ϖ for a uniformizer of B, then we have $\pi B = \varpi^e B$, hence $\varprojlim_j B/\pi^j B \cong \varprojlim_j B/\varpi^j B$. We find that B is a complete DVR. □

In particular, it makes sense to define $e_{L/K} = e_{\mathfrak{q}/\mathfrak{p}}$ and $f_{L/K} = f_{\mathfrak{q}/\mathfrak{p}}$, where $\mathfrak{p}, \mathfrak{q}$ are the unique non-zero prime ideals of A, B.

Corollary 6.2. *Suppose that the extension L/K is Galois with valuation v_L. Then for any $x \in L$ and $\sigma \in \mathrm{Gal}(L/K)$, we have $v_L(\sigma(x)) = v_L(x)$.*

Corollary 6.3. *For any $x \in L^\times$, we have $v_L(x) = \frac{1}{f_{L/K}} v_K(N_{L/K}(x))$.*

Proof. Let E/K be the Galois closure of L, and let $\sigma_1, \dots, \sigma_n : L \hookrightarrow E$ be the distinct K-embeddings. Then we have $N_{L/K}(x) = \prod_{i=1}^n \sigma_i(x)$, and

$$[L : K]v_L(x) = v_L(\prod_{i=1}^n \sigma(x)) = v_L(N_{L/K}(x)) = e_{L/K} v_K(N_{L/K}(x)).$$

The result now follows from the formula $[L : K] = e_{L/K} f_{L/K}$. □

Working in a complete DVR is very pleasant. We now give several examples of this.

Definition 6.4. *Let A be a DVR and let $v : K^\times \to \mathbb{Z}$ be the corresponding valuation. If $f(X) = X^d + a_1 X^{d-1} + \cdots + a_d \in K[X]$ is a polynomial with $a_d \neq 0$, then we define the Newton polygon of f to be the graph of the largest continuous piecewise linear function $N : [0, d] \to \mathbb{R}$ satisfying the following conditions:*

- *$N(0) = 0$ and $N(d) = v(a_d)$.*
- *For all $j = 1, \dots, d$, $N(j) \leq v(a_j)$.*
- *The derivative of N is non-decreasing away from its points of discontinuity.*

In other words, the Newton polygon is the lower convex hull of the set of points $(j, v(a_j))$.

Proposition 6.5. *Suppose that $f(X)$ factors as $f(X) = \prod_{i=1}^d (X - \alpha_i)$ with $\alpha_i \in K$. Suppose that the Newton polygon of $f(X)$ has slopes $\gamma_1 \leq \gamma_2 \leq \cdots \leq \gamma_d$, counted with multiplicity. Then after re-ordering we have $\gamma_i = v(\alpha_i)$ for each i.*

Proof. Let $\lambda_i = v(\alpha_i)$; we can assume after re-ordering that $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_d$. Define the λ-polygon to be the graph of the continuous piecewise linear function $L : [0, d] \to \mathbb{R}$ with $L(0) = 0$ and $L'(x) = \lambda_i$ for $x \in (i-1, i)$. Then $L(d) = \lambda_1 + \cdots + \lambda_d = v(\alpha_1 \dots \alpha_d) = v(a_d)$, so the Newton polygon and the λ-polygon have the same endpoints. We must show that they are in fact the same.

We first show that the Newton polygon lies above the λ-polygon. For this, it suffices to show that for all $j = 1, \dots, d$, we have $v(a_j) \geq \lambda_1 + \cdots + \lambda_j$. But we can write $a_j = \pm \sum \alpha_{i_1} \dots \alpha_{i_j}$, and each term in the sum has valuation $\lambda_{i_1} + \cdots + \lambda_{i_j} \geq \lambda_1 + \cdots + \lambda_j$. We therefore have $v(a_j) \geq \lambda_1 + \cdots + \lambda_j$, by the ultrametric property.

We now show that the two polygons are in fact equal. We will use the following general fact: if $x_1, \dots, x_r \in K$ and $v(x_1) < v(x_j)$ for all $j > 1$, then $v(x_1 + \dots + x_r) = v(x_1)$. Suppose that $\lambda_1 = \dots = \lambda_{k_1} < \lambda_{k_1+1} = \dots = \lambda_{k_2} < \dots$. We can write

$$a_{k_j} = \alpha_1 \dots \alpha_{k_j} + \sum \alpha_{i_1} \dots \alpha_{i_{k_j}}.$$

The terms in the sum all involve α_r for some $r > k_j$, hence have valuation strictly greater than $\lambda_1 + \dots + \lambda_{k_j}$. We conclude that $v(a_{k_j}) = v(\alpha_1 \dots \alpha_{k_j}) = \lambda_1 + \dots + \lambda_{k_j}$.

This implies that the two polygons must in fact coincide. Indeed, we have shown that the Newton polygon lies above the λ-polygon, and that the Newton polygon shares a vertex with each of the vertices of the λ-polygon. This completes the proof. □

Corollary 6.6. *Let A be a complete DVR with $\mathrm{Frac}(A) = K$ of characteristic 0, and let $f(X) = X^d + a_1 X^{d-1} + \dots + a_d \in K[X]$ be a polynomial with $a_d \neq 0$. Suppose that the slopes of the Newton polygon of $f(X)$ are $\lambda_1 \leq \dots \leq \lambda_k$, each appearing with width $w_1, \dots, w_k$. Then there is a unique factorization $f(X) = \prod_{i=1}^k g_i(X)$ in $K[X]$ such that $\deg g_i(X) = w_i$ and the Newton polygon of $g_i(X)$ has a single segment of slope λ_i.*

Proof. Let L/K denote the splitting field of $f(X)$, and let v_L denote the valuation of L, so that $v_L(x) = e_{L/K} v_K(x)$ for $x \in K^\times$. Let $\alpha_1, \dots, \alpha_n$ be the roots of $f(X)$ in L. We know that the numbers λ_i are exactly the $e_{L/K}^{-1} v_L(\alpha_j)$, with the width w_i being the number of j with $v_L(\alpha_j) = e_{L/K}\lambda_i$. We therefore define $g_i(X) = \prod(X - \alpha_j)$, the product being over the roots with $v_L(\alpha_j) = e_{L/K}\lambda_i$.

We clearly have $f(X) = \prod_{i=1}^k g_i(X)$ in $L[X]$. We claim that each $g_i(X)$ in fact lies in $K[X]$. The group $\mathrm{Gal}(L/K)$ permutes the α_i leaving the values $v_L(\alpha_i)$ invariant, so fixes the coefficients of each $g_i(X)$, so we have $g_i(X) \in K[X]$ by Galois theory. The uniqueness of the given factorization is clear. □

Example 6.7. The polynomial $X^3 + X^2 - 2X + 8$ has 3 distinct roots in $\mathbb{Q}_2$, because its Newton polygon has 3 distinct slopes.

We introduce some language to go with the above theorem.

Definition 6.8. *A pair (K, v) where K is a field and $v : K^\times \to \mathbb{Z}$ is a valuation is called a discrete valuation field (DVF). If v is clear from the context (for example, if $K = \mathbb{Q}_p$), then we will refer to K itself as a DVF. We then write A_K for its valuation ring, $\mathfrak{m}_K \subset A_K$ for the maximal ideal, and $k_K = A_K/\mathfrak{m}_K$ for the residue field.*

If A_K is complete, then we call (K, v) a complete discrete valuation field (CDVF). In this case, we say that L/K is an extension of compete discrete valuation fields if L/K is a separable field extension and L is endowed with its canonical structure of CDVF (i.e. with A_L the integral closure of A_K in L).

Let L/K be an extension of complete CDVFs. It is called unramified if $e_{L/K} = 1$ and the extension k_L/k_K of residue fields is separable. If L/K is unramified and Galois, then so is k_L/k_K and the map $\mathrm{Gal}(L/K) \to \mathrm{Gal}(k_L/k_K)$ is an isomorphism. This can be viewed as a special case of the following proposition:

Proposition 6.9. *Let A be a complete DVR with fraction field K and residue field $k_K = A/\mathfrak{m}_A$. Let l/k_K be a finite separable extension. Then there exists a finite separable extension L/K with the following property:*

1. *There is a k_K-isomorphism $k_L \cong l$, and L/K is unramified.*

2. *For any finite separable extension E/K equipped with a k-embedding $l \hookrightarrow k_E$, there is a unique K-embedding $L \hookrightarrow E$ which induces this map on residue fields.*

In particular, the field L with this property is unique up to unique isomorphism.

Proof. We first address existence. Let π be a uniformizer of A. Let $\overline{x} \in l$ be a primitive element, and let $\overline{f}(X) \in k_K[X]$ be its minimal polynomial. Thus there is an isomorphism $l \cong k_K[X]/(\overline{f}(X))$. Let $f(X) \in A[X]$ be an arbitrary monic lift of $\overline{f}(X)$, and let $B = A[X]/(f(X))$. The polynomial $f(X)$ is irreducible, because a factorization could be reduced modulo $\mathfrak{m}_A$ to give a factorization of $\overline{f}(X)$. It is separable for the same reason (look at common divisors of $f(X)$ and $f'(X)$). In particular, B embeds in the field $L = K[X]/(f(X))$, showing that B is a domain. For any $y \in L$ we have $\pi^n y \in B$ for some n, showing that $\operatorname{Frac} B = L$. Finally, we have $B/(\pi) = k_K[X]/(\overline{f}(X)) \cong l$.

We claim that B is integrally closed in L. Since B is clearly integral over A, this will imply that B is the integral closure of A in L, hence that B is a complete DVR with residue field l and that the extension L/K is unramified. To establish the claim, choose $y \in L$ and suppose that y is integral over B. We can write $y = z/\pi^n$ for some $z \in B$ and integer $n \geq 0$; let us assume that n is minimal with respect to this property.

Take an equation $y^d + b_1 y^{d-1} + \cdots + b_d = 0$ with $b_i \in B$. Substituting, we obtain the equation $z^d + \pi^n b_1 z^{d-1} + \cdots + \pi^{nd} b_d = 0$ in B. If $n > 0$, this shows that the element $\overline{z} = z \bmod \pi B$ of $B/\pi B \cong l$ satisfies $\overline{z}^d = 0$. Since l is a field, this forces $\overline{z} = 0$, implying that z is divisible by π in B, a contradiction. Thus $n = 0$ and we in fact have $y \in B$, showing that B is integrally closed.

We now establish the universal property of L/K. Let E/K be a finite separable extension equipped with a k_K-embedding $l \hookrightarrow k_E$, and let C denote the integral closure of A in E. The polynomial $\overline{f}(X)$ splits into distinct linear factors in $l[X]$, so by Hensel's lemma there is a unique element $x \in C$ lifting the image of $\overline{x} \in k_E$ and satisfying $f(x) = 0$. This determines a homomorphism $B = A[X]/(f(X)) \to C$ by the formula $X \mapsto x$. Passing to fraction fields gives the desired K-embedding $L \hookrightarrow E$.

The same argument establishes uniqueness: if $L \hookrightarrow E$ is any K-embedding, then we get an induced map $\phi : B \hookrightarrow C$, hence an element $x' = \phi(X \bmod f(X)) \in C$ such that $f(x') = 0$. The compatibility with the embedding $l \hookrightarrow k_E$ means that $x' \bmod \mathfrak{m}_C = \overline{x}$. The uniqueness part of Hensel's lemma then forces $x' = x$, showing that there is exactly one embedding $L \hookrightarrow E$ with the desired properties. □

Corollary 6.10. *Let A be a complete DVR with field of fractions K, and let E/K be a finite separable extension such that the corresponding extension k_E/k_K of residue fields is separable. Then there is a unique intermediate subfield $E/E_0/K$ with the following property: E_0/E is unramified, and if $E' \subset E$ is any other intermediate field which is unramified, then $E' \subset E_0$.*

We call E_0 the maximal unramified subextension of E.

Proof. We take L/K to be the unramified extension associated by the proposition to k_E/k_K, and E_0 to be the image of L under the corresponding K-embedding $L \hookrightarrow E$. The corollary then follows from the universal property of L. □

Corollary 6.11. *Let A be a complete DVR with field of fractions K, and let E/K be a finite separable extension such that the corresponding extension k_E/k_K of residue fields is separable. Then the following sets are in canonical bijection:*

1. *The set of intermediate extensions $E/L/K$ such that L/K is unramified.*
2. *The set of intermediate extensions $k_E/l/k_K$.*

If L/K is an extension of CDVFs with k_L/k_K separable, then the extension L/L_0 satisfies $f_{L/L_0} = 1$ and $e_{L/L_0} = [L : L_0]$. Such extensions are said to be totally ramified. We now characterize these extensions:

Proposition 6.12. *Let K be a CDVF of characteristic 0.*

1. *Let $f(X) = X^d + a_1 X^{d_1} + \cdots + a_d \in A_K[X]$ be a polynomial which is Eisenstein, i.e. such that $v(a_d) = 1$ and $v(a_i) \geq 1$ for each $i = 1, \ldots, d-1$. Then $f(X)$ is irreducible and the extension $L = K[X]/(f(X))$ is totally ramified.*
2. *Suppose conversely that L/K is a finite extension which is totally ramified, and let $\pi_L \in L$ be a uniformizer with minimal polynomial $f(X) = X^d + a_1 X^{d_1} + \cdots + a_d \in A_K[X]$. Then $f(X)$ is Eisenstein and $A_L = A_K[\pi_L]$.*

Proof. The condition that $f(X)$ is Eisenstein is equivalent to the condition that the Newton polygon $N_K(f)$ has a single segment of slope $1/d$. Let E denote the splitting field of $f(X)$, and let $\pi \in E$ be a root, $L = K(\pi)$. The Newton polygon of $f(X)$ over L has slopes $e_{L/K}$ times the slopes of the Newton polygon over K, hence equal to $e_{L/K}/d$. On the other hand these equal $v_L(\pi)$. Since $e_{L/K} \leq d$, it follows that $e_{L/K} = d$ and $v_L(\pi) = 1$, showing that L/K is totally ramified of degree d and that $f(X)$ is irreducible.

Suppose instead that L/K is a totally ramified extension, and let B denote the integral closure of A in L, $\pi_L \in B$ a uniformizer with minimal polynomial $f(X) \in A[X]$. Let $\pi_K \in A$ be a uniformizer of A. Then the Newton polygon of $f(X)$ has a single segment of slope $1/d$, so is Eisenstein. Any element $x \in B$ admits an expression $x = \sum_{i=0}^{\infty} a_i \pi_L^i$ with $a_i \in A$; it follows that the map $A[\pi_L] \to B/\pi_K B = B/\pi_L^e B$ is surjective. Applying Nakayama's lemma to the finitely generated A-module $B/A[\pi_L]$, we find that $B = A[\pi_L]$. □

We state a related result:

Proposition 6.13. *Let L/K be an extension of CDVFs of characteristic 0 with k_L/k_K separable. Then there exists $x \in A_L$ such that $A_L = A_K[x]$.*

Proof. Let $L/L_0/K$ be the maximal unramified subextension. We first choose $\overline{y} \in k_L$ which is a primitive element for $k_L = k_{L_0}/k_K$, and let $y \in A_{L_0}$ be any lift of $\overline{y}$. Let $\pi \in A_L$ be any uniformizer. Then the previous proposition shows that $A_L = A_{L_0}[\pi] = A_K[y, \pi]$.

Let $f(X) \in A_K[X]$ be the minimal polynomial of y over K. Then $\overline{f}(X) \in k_K[X]$ is the minimal polynomial of $\overline{y}$ over k_K, and we can write

$$f(y + \pi) = f(y) + \pi f'(y) + \pi^2 z = \pi f'(y) + \pi^2 z,$$

where $z \in A_L$. The element $f'(y)$ is a unit, because it reduces modulo π to $\overline{f}'(\overline{y}) \neq 0$. We deduce that $v_L(f(y + \pi)) = 1$, and hence $f(y + \pi)$ is a uniformizer of A_L. We set $x = y + \pi$.

To show that $A_L = A_K[x]$, it is enough to (by Nakayama's lemma) to show that the map $A_K[x] \to A_L/\mathfrak{m}_K A_L$ is surjective. We observe that the map $A_K[x] \to k_L$ is surjective, so every element of k_L can be represented by an element of $A_K[x]$. Since $f(x)$ is a uniformizer of A_L, every element of $A_L/\mathfrak{m}_K A_L$ admits a representative of the form $\sum_{i=0}^{e_{L/K}-1} a_i f(x)^i$ with $a_i \in A_K[x]$. But the polynomial $f(X)$ has coefficients in K, so any such element lies in $A_K[x]$. This completes the proof. □

Finally, we discuss passage to completion.

Proposition 6.14. *Let A be a Dedekind domain with field of fractions K, let L/K be a finite separable extension, let B be the integral closure of A in L, and let $\mathfrak{q}$ be a non-zero prime ideal of B lying above the prime $\mathfrak{p}$ of A. Then:*

1. *There is a canonical embedding of DVRs $\widehat{A}_\mathfrak{p} \to \widehat{B}_\mathfrak{q}$ extending the embedding $A \to B$.*
2. *Let $L_\mathfrak{q}, K_\mathfrak{p}$ denote the fields of fractions of $\widehat{B}_\mathfrak{q}$ and $\widehat{A}_\mathfrak{p}$, respectively. Then $L_\mathfrak{q}/K_\mathfrak{p}$ is a finite separable extension, $\widehat{B}_\mathfrak{q}$ is the integral closure of $K_\mathfrak{p}$ in $L_\mathfrak{q}$, and we have $e_{L_\mathfrak{q}/K_\mathfrak{p}} = e_{\mathfrak{q}/\mathfrak{p}}$, $f_{L_\mathfrak{q}/K_\mathfrak{p}} = f_{\mathfrak{q}/\mathfrak{p}}$.*
3. *Suppose further that the extension L/K is Galois. Then $L_\mathfrak{q}/K_\mathfrak{p}$ is Galois and the natural map $D_{\mathfrak{q}/\mathfrak{p}} \to \mathrm{Gal}(L_\mathfrak{q}/K_\mathfrak{p})$ given by passage to completion is an isomorphism.*

Proof. We first note that we have $\widehat{A}_\mathfrak{p} \cong \varprojlim_i A/\mathfrak{p}^i$, and similarly for $\widehat{B}_\mathfrak{q}$. Indeed, it suffices to note that the natural map $A/\mathfrak{p}^i \to A_\mathfrak{p}/\mathfrak{p}_\mathfrak{p}^i$ is an isomorphism for any $i \geq 1$. It is surjective by the existence of π-adic expansions. It is injective because $\mathfrak{p}_\mathfrak{p}^i \cap A = \mathfrak{p}^i$: if $a/s \in \mathfrak{p}_\mathfrak{p}^i \cap A$, then we have $a = sb$ for some $b \in A$, $a \in \mathfrak{p}^i$, and $b \in A - \mathfrak{p}$. We then get $v_\mathfrak{p}(b) = v_\mathfrak{p}(a) \geq i$, showing that $b \in \mathfrak{p}^i$.

We have natural maps $A/\mathfrak{p}^i \to B/\mathfrak{q}^i$ for each $i \geq 1$, and passage to the inverse limit gives a map $\widehat{A}_\mathfrak{p} \to \widehat{B}_\mathfrak{q}$. It is injective because the map $A \to B$ is injective. If we write $\mathfrak{q} = \mathfrak{q}_1, \ldots, \mathfrak{q}_r$ for the distinct primes of B above $\mathfrak{p}$, then there is an isomorphism

$$\varprojlim_i B/\mathfrak{p}^i B \cong \prod_{i=1}^r \widehat{B}_{\mathfrak{q}_i}.$$

Since B is a finite free A-module, $\varprojlim_i B/\mathfrak{p}^i B$ is a finite free $\widehat{A}_\mathfrak{p}$-module. We conclude that the quotient $\widehat{B}_\mathfrak{q}$ is a finite free $\widehat{A}_\mathfrak{p}$-module. In particular, $\widehat{B}_\mathfrak{q}$ is integral over $\widehat{A}_\mathfrak{p}$, hence is the integral closure of $\widehat{A}_\mathfrak{p}$ in $L_\mathfrak{q}$.

We also see (using π-adic expansions and Nakayama's lemma) that $\widehat{B}_\mathfrak{q} = \widehat{A}_\mathfrak{p} \cdot B$ (i.e. $\widehat{B}_\mathfrak{q}$ is generated as a ring by these two subrings). This implies that $L_\mathfrak{q} = K_\mathfrak{p} \cdot L$ (compositum of subfields of $L_\mathfrak{q}$), hence that $L_\mathfrak{q}/K_\mathfrak{p}$ is separable. The isomorphisms $\widehat{B}_\mathfrak{q}/\mathfrak{q}\widehat{B}_\mathfrak{q} \cong B/\mathfrak{q}$ and $\widehat{B}_\mathfrak{q}/\mathfrak{p}\widehat{B}_\mathfrak{q} \cong B/\mathfrak{q}^e_{\mathfrak{q}/\mathfrak{p}}B$ show that $f_{L_\mathfrak{q}/K_\mathfrak{p}} = f_{\mathfrak{q}/\mathfrak{p}}$ and $e_{L_\mathfrak{q}/K_\mathfrak{p}} = e_{\mathfrak{q}/\mathfrak{p}}$.

Suppose finally that L/K is Galois. Any element of $D_{\mathfrak{q}/\mathfrak{p}}$ acts on $\widehat{B}_\mathfrak{q}$, by passage to completion, so we obtain a map $D_{\mathfrak{q}/\mathfrak{p}} \hookrightarrow \mathrm{Aut}(L_\mathfrak{q}/K_\mathfrak{p})$. Since the source has order $e_{\mathfrak{q}/\mathfrak{p}}f_{\mathfrak{q}/\mathfrak{p}}$ and the target has dimension at most $[L_\mathfrak{q} : K_\mathfrak{p}] = e_{L_\mathfrak{q}/K_\mathfrak{p}}f_{L_\mathfrak{q}/K_\mathfrak{p}}$, we find that this map is an isomorphism, that $\mathrm{Aut}(L_\mathfrak{q}/K_\mathfrak{p}) = [L_\mathfrak{q} : K_\mathfrak{p}]$, and hence that the extension $L_\mathfrak{q}/K_\mathfrak{p}$ is Galois. □

7 Number fields

Definition 7.1. *A number field is a finite extension $K/\mathbb{Q}$. We write $\mathcal{O}_K$ for the integral closure of $\mathbb{Z}$ in K, and call it the ring of integers of K.*

We observe that $\mathcal{O}_K$ is a Dedekind domain.

Lemma 7.2. *Let K be a number field and let $\mathfrak{p} \subset \mathcal{O}_K$ be a non-zero prime ideal. Then $\mathcal{O}_K/\mathfrak{p}$ is a finite field.*

Proof. Let $\mathfrak{p} \cap \mathbb{Z} = (p)$, for a prime number p. Then $\mathcal{O}_K/\mathfrak{p}$ is a finite extension of $\mathbb{F}_p$, so is a finite field. □

If L/K is a Galois extension of number fields, and $\mathfrak{q} \subset \mathcal{O}_L$ is a non-zero prime ideal lying above the prime $\mathfrak{p}$ of $\mathcal{O}_K$, then the map $D_{\mathfrak{q}/\mathfrak{p}} \to \mathrm{Gal}(k_\mathfrak{q}/k_\mathfrak{p})$ is surjective. The group $\mathrm{Gal}(k_\mathfrak{q}/k_\mathfrak{p})$ has a canonical generator, the Frobenius automorphism $x \mapsto x^{\#k_\mathfrak{p}}$. If $e_{\mathfrak{q}/\mathfrak{p}} = 1$, then the map $D_{\mathfrak{q}/\mathfrak{p}} \to \mathrm{Gal}(k_\mathfrak{q}/k_\mathfrak{p})$ is an isomorphism and this automorphism therefore lifts to a canonical element, the Frobenius element $\mathrm{Frob}_{\mathfrak{q}/\mathfrak{p}} \in \mathrm{Gal}(L/K)$, which depends only on the prime $\mathfrak{q}$. If $\sigma \in \mathrm{Gal}(L/K)$ then $\mathrm{Frob}_{\sigma(\mathfrak{q})/\mathfrak{p}} = \sigma \, \mathrm{Frob}_{\mathfrak{q}/\mathfrak{p}} \, \sigma^{-1}$, which shows that the Frobenius conjugacy class $\mathrm{Frob}_\mathfrak{p}$ depends only on the underlying prime $\mathfrak{p}$ (recall that $\mathrm{Gal}(L/K)$ acts transitively on the set of primes of $\mathcal{O}_L$ above $\mathfrak{p}$).

In general, we can study the decomposition group by passage to completion: we write $\mathcal{O}_{K_\mathfrak{p}}$ for the localization and completion of $\mathcal{O}_K$ at the prime $\mathfrak{p}$, and $K_\mathfrak{p} = \mathrm{Frac}\, \mathcal{O}_{K_\mathfrak{p}}$. Then there is an isomorphism $D_{\mathfrak{q}/\mathfrak{p}} \cong \mathrm{Gal}(L_\mathfrak{q}/K_\mathfrak{p})$, given by "passage to completion".

Proposition 7.3. *Let K be a number field, and let $L = K(\alpha)$ be a finite extension, $f(X) \in K[X]$ the minimal polynomial of α. Let $\mathfrak{p}$ be a non-zero prime ideal of $\mathcal{O}_K$. Then the following two sets are in canonical bijection:*

1. *The irreducible factors of $f(X)$ in $K_\mathfrak{p}[X]$.*

2. The primes of $\mathcal{O}_L$ lying above the prime $\mathfrak{p}$ of $\mathcal{O}_K$.

Proof. Let E/K be the splitting field of $f(X)$, and let $G = \mathrm{Gal}(E/K)$, $H = \mathrm{Gal}(E/L)$. Let $\mathfrak{q}$ be a fixed prime of $\mathcal{O}_E$ above $\mathfrak{p}$, and let $D = D_{\mathfrak{q}/\mathfrak{p}} = \mathrm{Gal}(E_\mathfrak{q}/K_\mathfrak{p})$. We first note that there is an isomorphism of G-sets (i.e. sets with left G-action) between G/H and the set of roots of $f(X)$ in E, hence an isomorphism of $D_{\mathfrak{q}/\mathfrak{p}}$-sets between G/H and the set of roots of $f(X)$ in $E_\mathfrak{q}$. It follows that the irreducible factors of $f(X)$ are in $K_\mathfrak{p}[X]$ are in bijection with the set of orbits of $D_{\mathfrak{q}/\mathfrak{p}}$ on G/H. This set of orbits is identified with the double quotient $D\backslash G/H$.

On the other hand, there is an isomorphism of G-sets between G/D and the set of prime ideals of $\mathcal{O}_E$ lying above $\mathfrak{p}$, because G acts transitively with stabilizer D. The set of primes ideals of $\mathcal{O}_L$ lying above $\mathfrak{p}$ is in bijection with the set of H-orbits, because H acts transitively on the set of prime ideals of $\mathcal{O}_E$ above a given prime of $\mathcal{O}_L$. We conclude that the set of primes of $\mathcal{O}_L$ is in bijection with the double quotient $H\backslash G/D$.

These two double quotients are in bijection via the map $D\sigma H \mapsto H\sigma^{-1}D$. This is the bijection of the proposition. It can be interpreted directly as follows: if $g(X)$ is an irreducible factor of $f(X)$ in $K_\mathfrak{p}[X]$, we choose $\sigma \in G$ such that $g(\sigma(\alpha)) = 0$ in $E_\mathfrak{q}$. Then we take the prime ideal $\sigma^{-1}(\mathfrak{q}) \cap \mathcal{O}_L$ of $\mathcal{O}_L$.

To show that this bijection is canonical, we must show that it is independent of the choice of $\mathfrak{q}$. Let $\tau \in G$, and let $\mathfrak{q}' = \tau(\mathfrak{q})$. Let $g(X)$ be an irreducible factor of $f(X)$ in $K_\mathfrak{q}[X]$. Then τ induces an isomorphism $E_\mathfrak{q} \to E_{\mathfrak{q}'}$ that respects $K_\mathfrak{p}$, so we get $\tau(g(\sigma(\alpha)) = g(\tau(\sigma(\alpha)) = 0$. In the bijection defined using $\mathfrak{q}'$, we then take the prime ideal $(\tau\sigma)^{-1}(\mathfrak{q}') \cap \mathcal{O}_L = \sigma^{-1}\tau^{-1}\tau(\mathfrak{q}) \cap \mathcal{O}_L = \sigma^{-1}(\mathfrak{q}) \cap \mathcal{O}_L$. This shows the independence of the choice of $\mathfrak{q}$. □

These kinds of methods are very useful for studying the Galois theory of number fields. Here is a simple example: let $K = \mathbb{Q}(\sqrt{d})$, where $d \neq 0, 1$ is a square-free integer. Then K is a quadratic field, and we can use the proposition to factor the ideal $p\mathcal{O}_K$ for any prime number p by factoring the polynomial $f(X) = X^2 - d$ in $\mathbb{Q}_p[X]$. We split into 3 cases:

- If p is odd and $p \nmid d$, then the polynomial $\overline{f}(X) = f(X) \bmod p \in \mathbb{F}_p[X]$ has distinct roots modulo p. If d is a quadratic residue modulo p then Hensel's lemma shows that $f(X)$ splits into linear factors in $\mathbb{Q}_p[X]$. Otherwise $\overline{f}(X)$, and hence $f(X)$, is irreducible and $p\mathcal{O}_K$ is prime.

- If $p|d$, then $f(X)$ is Eisenstein, hence irreducible, and p is ramified in K.

- If $p = 2$ and $p \nmid d$, then the behaviour depends on the image of d in $\mathbb{Q}_2^\times/(\mathbb{Q}_2^\times)^2$. If $d \equiv 1 \bmod 8$, then $f(X)$ splits into linear factors in $\mathbb{Q}_p[X]$. If $d \equiv 5 \bmod 8$, then $f(X)$ is irreducible and p is unramified in $f(X)$ (because $\mathbb{Q}_p(\sqrt{5})/\mathbb{Q}_p$ is unramified). If $d \equiv 3, 7 \bmod 8$, then $f(X)$ is irreducible and p is ramified in K.

The behaviour of Frobenius elements at unramified primes can be analyzed in general as follows. Let $f(X) \in \mathbb{Z}[X]$ be a monic irreducible polynomial, let $L/\mathbb{Q}$ be its splitting field, and let $\alpha_1, \ldots, \alpha_d \in \mathcal{O}_L$ be its roots; then $\mathrm{Gal}(L/\mathbb{Q})$ is identified with a transitive subgroup of S_d, the symmetric group on d letters.

Let p be a prime not dividing $\operatorname{disc} f$ (which means all but finitely many primes p), and let $\overline{f}(X) = f(X) \bmod p$. We claim that p is unramified in $\mathcal{O}_L$, and that the cycle type of the Frobenius element Frob_p as an element of S_d depends only on the factorization of $\overline{f}(X)$ into irreducibles in $\mathbb{F}_p[X]$. Note that the assumption that p does not divide $\operatorname{disc} f$ means exactly that $\overline{f}(X)$ has no repeated roots.

To see this, let $\mathfrak{q}$ be a prime of $\mathcal{O}_L$ lying above p. Then $L_\mathfrak{q} = \mathbb{Q}(\alpha_1, \dots, \alpha_d)$. Indeed, we know that $L_\mathfrak{q} = \mathbb{Q}_p(\alpha)$, where $\alpha \in L$ is a primitive element. But α is a polynomial in the $\alpha_1, \dots, \alpha_d$ with $\mathbb{Q}$-coefficients, and conversely each α_i can be expressed as a polynomial in α with $\mathbb{Q}$-coefficients. Let $K \subset L_\mathfrak{q}$ denote the maximal unramified subextension. We know that $\overline{f}(X)$ factors into linear factors over $k_\mathfrak{q}$, so Hensel's lemma shows that $f(X)$ factors into linear factors in $\mathcal{O}_K[X]$. This shows that $\alpha_1, \dots, \alpha_d \in \mathcal{O}_K$, and hence $K = L_\mathfrak{q}$ and p is indeed unramified.

We now show how to calculate $\mathrm{Frob}_{\mathfrak{q}/p}$. Let $\overline{\alpha}_i = \alpha_i \bmod \mathfrak{q}$; then the map $\{\alpha_1, \dots, \alpha_d\} \to \{\overline{\alpha}_1, \dots, \overline{\alpha}_d\}$ is a bijection. The Frobenius automorphism of $k_\mathfrak{q}$ acts on the elements $\overline{\alpha}_1, \dots, \overline{\alpha}_d$ by cyclic permutations, with cycle types $(d_1)(d_2)\dots(d_k)$, say. The irreducible factors of $\overline{f}(X)$ are the products $\prod(X - \overline{\alpha})$, the product running over the roots of $\overline{\alpha}$ in a given orbit of Frobenius. We see therefore that $\overline{f}(X)$ factors as a product of k distinct irreducible polynomials, with degrees $d_1, \dots, d_k$.

Now let's consider a different example. Let l be a prime, and consider a polynomial $f(X) = X^l - aX - b \in \mathbb{Z}[X]$, where the integers $(l-1)a$ and lb are coprime. Suppose that $f(X)$ is irreducible, and let $K/\mathbb{Q}$ be the splitting field of $f(X)$. We claim that if p is a prime of $\mathbb{Z}$ ramified in $\mathcal{O}_K$, and $\mathfrak{p}$ is a prime of $\mathcal{O}_K$ above p, then $e_{\mathfrak{p}/p} = 2$ and the inertia group $I_{\mathfrak{p}/p} \subset \mathrm{Gal}(K/\mathbb{Q}) \subset S_l$ is generated by a transposition.

To see this, we first note that $\overline{f}(X) = f(X) \bmod p$ has a repeated root; otherwise, our previous argument shows that p is in fact unramified in $\mathcal{O}_K$. The equation $Xf'(X) - lf(X) = alX + lb - aX = (l-1)aX + lb$ shows that the GCD of $\overline{f}(X)$ and $\overline{f}'(X)$ divides $(l-1)aX + lb$. Our assumption that $(l-1)a$ and lb are coprime shows that this polynomial is non-zero modulo p. Since it is linear, we see that if $\overline{f}(X)$ has a repeated root, then we must have $\overline{f}(X) = \overline{g}(X)^2\overline{h}(X)$, where $\overline{g}(X) \in \mathbb{F}_p[X]$ is linear, $\overline{h}(X) \in \mathbb{F}_p[X]$ has distinct roots, and $\overline{g}$, $\overline{h}$ have no roots in common.

Let K_0 be the maximal unramified subextension of $K_\mathfrak{p}/\mathbb{Q}_p$. Hensel's lemma shows that we can factorize $f(X) = r(X)h(X)$ in $\mathcal{O}_{K_0}[X]$, where $h(X) \in \mathcal{O}_{K_0}[X]$ lifts $\overline{h}(X)$ and splits into linear factors in $\mathcal{O}_{K_0}[X]$, and $r(X)$ lifts $\overline{g}(X)^2$. Since $K_\mathfrak{p}$ is generated over $\mathbb{Q}_p$ by the roots of $f(X)$, we conclude that $K_\mathfrak{p}$ is generated over K_0 by adjoining the roots of $r(X)$, hence is an extension of degree at most 2. Since $e_{\mathfrak{p}/p} = [K_\mathfrak{p} : K_0]$, and we have assumed $e_{\mathfrak{p}/p} > 1$, we get $e_{\mathfrak{p}/p} = [K_\mathfrak{p} : K_0] = 2$. We also see that the inertia group at such a prime permutes the roots of $r(X)$ and fixes the roots of $h(X)$, so is generated by a transposition.

This completes the proof of the claim, and shows that $\mathrm{Gal}(K/\mathbb{Q}) = S_l$: the group S_l acts transitively on l letters, and l is a prime, so the Galois group contains an l-cycle. Every number field is ramified above some prime p of $\mathbb{Z}$, so we find that the Galois group contains a transposition as well, hence equals the whole of S_l. Writing $L = K^{A_l}$ for the quadratic extension of $\mathbb{Q}$ fixed by A_l, we see that the extension K/L is an everywhere unramified extension of number fields with Galois group A_l.

As a concrete example, consider the polynomial $f(X) = X^5 - X + 1$. Then $f(X)$ is irreducible, and its discriminant equals $2869 = 19 \cdot 151$. These primes are both ramified in K, the splitting field of $f(X)$, and we find that the extension $K/\mathbb{Q}(\sqrt{2869})$ is an everywhere unramified extension with Galois group A_5.

8 Lower ramification groups

Let L/K be a Galois extension of CDVFs such that k_L/k_K is separable.

Definition 8.1. *Let $G = \mathrm{Gal}(L/K)$. For each $i \geq 0$ we define the i^{th} ramification group $G_i = \ker(\mathrm{Gal}(L/K) \to \mathrm{Aut}(A_L/\mathfrak{m}_L^{i+1}))$.*

Thus $G_0 = I_{L/K} = \ker(\mathrm{Gal}(L/K) \to \mathrm{Gal}(k_L/k_K))$ is the usual inertia group.

Lemma 8.2. *1. Each G_i is a normal subgroup of G. We have $G \supset G_0 \supset G_1 \supset \dots$ and $\cap_{i \geq 0} G_i = \{1\}$.*

2. Suppose that $A_L = A_K[x]$ for some $x \in A_L$, and for $s \in G$ define $i_G(s) = v_L(s(x) - x)$. Then i_G is independent of the choice of x and we have $G_i = \{s \in G \mid i_G(s) \geq i+1\}$. (By convention we set $i_G(1) = \infty$.)

Proof. The first part is clear from the definition. For the second, we have $s \in G_i$ if and only if $s(x) - x \in \mathfrak{m}_L^{i+1}$, if and only if $v_L(s(x) - x) = i_G(s) \geq i + 1$. This shows that i_G is independent of the choice of x. □

Example 8.3. Let $K = \mathbb{Q}_2(\sqrt{2}, i)$, a compositum of two ramified quadratic extensions. The quadratic extensions are distinct, and we have $G = \mathrm{Gal}(K/\mathbb{Q}_2) \cong (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})$. Indeed, $\zeta = (1+i)/\sqrt{2}$ is an 8^{th} root of unity, and $\zeta - 1$ satisfies the polynomial $(X+1)^4 + 1 = X^4 + 4X^3 + 6X^2 + 4X + 2$, which is Eisenstein over $\mathbb{Q}_2$. We find that $K/\mathbb{Q}_2$ is totally ramified, so $G = G_0$.

Let B denote the integral closure of $\mathbb{Z}_2$ in K. We have $B = \mathbb{Z}_2[\zeta - 1] = \mathbb{Z}_2[\zeta]$, and we calculate $s(\zeta) - \zeta$ for the various elements of G:

$$s_1 : \frac{1-i}{\sqrt{2}} - \frac{1+i}{\sqrt{2}} = -\sqrt{2}i, \; i_G(s_1) = 2,$$

$$s_2 : -\frac{1+i}{\sqrt{2}} - \frac{1+i}{\sqrt{2}} = -\sqrt{2}(1+i), \; i_G(s_2) = 4,$$

$$s_3 : -\frac{1-i}{\sqrt{2}} - \frac{1+i}{\sqrt{2}} = -\sqrt{2}, \; i_G(s_3) = 2.$$

We find $G_0 = G_1 = \{1, s_1, s_2, s_3\}$, $G_2 = G_3 = \{1, s_2\}$, and $G_4 = \{1\}$.

Lemma 8.4. *Let $\pi \in A_L$ be a uniformizer, and let $s \in G_0$, $i \geq 0$. Then $s \in G_i$ if and only if $s(\pi)/\pi \equiv 1$ mod (π^i).*

Proof. We can assume, after replacing K by the maximal unramified subextension of L/K, that L/K is totally ramified and $A_L = A_K[\pi]$. If $s \in G_0$, then we have $i_G(s) = v_L(s(\pi) - \pi) = v_L(s(\pi)/\pi - 1) + 1$, hence $i_G(s) \geq i + 1$ if and only if $s(\pi)/\pi \equiv 1 \bmod \pi^i B$. □

If $s \in G_0$, then the element $s(\pi)/\pi$ lies in $A_L^\times$ (since it has valuation 0). We define a filtration of the group $U_L = A_L^\times$ by the formula $U_L^i = \ker(A_L^\times \to (A_L/\pi^i)^\times)$. We thus have $U_L/U_L^1 \cong (A_L/\mathfrak{m}_L)^\times = k_L^\times$. On the other hand for $i \geq 1$ we have $U_L^i = 1 + \pi^i A_L$, and we have an isomorphism $\pi^i A_L/\pi^{i+1} A_L \cong U_L^i/U_L^{i+1}$ given by $x \mapsto 1 + x$. It is a homomorphism because $(1 + x)(1 + y) = 1 + x + y + xy$, and $xy \in \pi^{i+1} A_L$.

Proposition 8.5. *Let $\pi \in A_L$ be a uniformizer.*

1. *There is an injection $G_0/G_1 \to k_L^\times$ given by the formula $s \mapsto s(\pi)/\pi$ mod $\mathfrak{m}_L$.*

2. *For each $i \geq 1$, there is an injection $G_i/G_{i+1} \to \pi^i A_L/\pi^{i+1} A_L$ given by the formula $s \mapsto s(\pi)/\pi - 1$.*

3. *The quotient G_0/G_1 is cyclic. If k_L has characteristic 0, then the group G_1 is trivial. If k_L has characteristic $p > 0$, then the group G_1 is the unique p-Sylow subgroup of G_0.*

Proof. We have already proved the first two parts. We prove the third. The group G_0 is finite, and G_0/G_1 is a finite subgroup of the multiplicative group of a field, which is therefore cyclic. If k_L has characteristic 0, then each of the quotients G_i/G_{i+1}, $i > 0$, is a finite group which injects into a $\mathbb{Q}$-vector space, which is therefore trivial. Since $\cap_i G_i = \{1\}$, we see that the group G_1 must be trivial.

If k_L has characteristic p, then the same argument shows that G_1 has order a power of p: each G_i/G_{i+1} is a subgroup of the additive group of a 1-dimensional k_L-vector space. On the other hand, the quotient G_0/G_1 injects into the torsion subgroup of $k_L^\times$, which has order prime to p. The result follows. □

Corollary 8.6. *The group $I_{L/K} = G_0$ is soluble. If the residue field k_K is finite, then the group* Gal(L/K) *is soluble. There is no Galois extension $E/\mathbb{Q}_p$ with Galois group A_5.*

Definition 8.7. *If L/K is a Galois extension of complete discrete non-archimedean valued fields, we say that L/K is tamely ramified if the group G_1 is trivial. We say that it is wildly ramified if G_1 is non-trivial.*

Corollary 8.8. *Suppose that the extension L/K is Galois and totally tamely ramified of degree n. Then K contains the n^{th} roots of unity, and there exists a uniformizer $\pi_K \in A_K$ such that $L = K(\sqrt[n]{\pi_K})$.*

Proof. There is an embedding $G_0 \hookrightarrow k_K^\times$. Any finite subgroup of $k_K^\times$ is cyclic, so we find that G_0 is cyclic of degree n. By the existence of Teichmüller representatives, we find that K itself contains the n^{th} roots of unity. Moreover, if $\pi_L \in A_L$ is a uniformizer and $\sigma \in \text{Gal}(L/K)$ is a generator, then there is a primitive n^{th} root of unity $\zeta \in A_K^\times$ such that $\sigma(\pi_L) = \zeta \pi_L \bmod \pi_L^2$.

Let $\alpha = \pi_L + \zeta^{-1}\sigma(\pi_L) + \zeta^{-2}\sigma^2(\pi_L) + \cdots + \zeta^{-(n-1)}\sigma^{n-1}(\pi_L)$. Then $\sigma(\alpha) = \sigma(\pi_L) + \cdots + \zeta\pi_L = \zeta\alpha$, hence $\alpha^n \in A_K$. We claim that α is a uniformizer of A_L. We calculate modulo π_L^2: here, we have

$$\alpha \bmod \pi_L^2 \equiv \pi_L + \zeta^{-1}\zeta\pi_L + \zeta^{-2}\zeta^2\pi_L + \cdots = n\pi_L \bmod \pi_L^2.$$

Since L/K is tamely ramified, $n \in A_K^\times$ and hence $v(\alpha) = v(\pi_L) = 1$. It follows that $A_L = A_K[\alpha]$ and $\alpha^n = \pi_K$ is a uniformizer of A_K. □

9 Upper ramification groups

Suppose again that L/K is a Galois extension of CDVFs with k_L/k_K separable. Consider an intermediate extension $L/E/K$ with E/K Galois, and let $H = \mathrm{Gal}(L/E)$, a normal subgroup of G. It is immediate from the definition that $i_H = i_G|_H$, and hence that the ramification filtrations are compatible: $H_i = G_i \cap H$. However, this is not true for passage to quotient, as the following example shows:

Example 9.1. Let $L = \mathbb{Q}_2(\sqrt{2}, i)$. We have seen that $G = \mathrm{Gal}(L/\mathbb{Q}_2) \cong (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})$. We have $G = G_0 = G_1$, and $G_2 = G_3 = \mathrm{Gal}(L/\mathbb{Q}_2(i))$, $G_4 = \{1\}$. We calculate the ramification groups for some quadratic subextensions of L.

If $K = \mathbb{Q}_2(i)$, $H = \mathrm{Gal}(L/K)$, $G/H = \{1, s\}$, then $i_{G/H}(s) = v_K(-2i) = 2$, so $(G/H)_0 = (G/H)_1$, $(G/H)_2 = \{1\}$. In this case $(G/H)_i$ equals the image of G_i in G/H for each $i \geq 0$.

If $E = \mathbb{Q}_2(\sqrt{2})$, $N = \mathrm{Gal}(L/E)$, $G/N = \{1, t\}$, then $i_{G/N}(t) = v_E(-\sqrt{2} - \sqrt{2}) = 3$, so $(G/N)_0 = (G/N)_1 = (G/N)_2$, $(G/N)_3 = \{1\}$. In particular, $(G/N)_3$ is not equal to the image of G_3 in G/N.

We therefore now show how to modify the numbering of ramification subgroups so that they become compatible with passage to quotient. This work will later play an essential role in our formulation of global class field theory.

Definition 9.2. *If $u \geq 0$ is a real number, then we define $G_u = G_{\lceil u \rceil} = \{s \in G \mid i_G(s) \geq u + 1\}$. The ramification function is $\varphi(u) = \varphi_{L/K}(u) = \int_{t=0}^{u} [G_0 : G_u]^{-1} dt$.*

Lemma 9.3.

1. *The function φ is a continous piecewise linear, increasing homeomorphism of $[0, \infty)$ to itself.*
2. *We have the formula $\varphi(u) + 1 = \frac{1}{\#G_0} \sum_{s \in G} \min(i_G(s), u + 1)$.*

Proof. We can write an explicit formula: let $g_i = \#G_i$, and suppose that $m \leq u \leq m + 1$, where m is an integer. Then we have

$$\varphi(u) = \sum_{i=1}^{m} g_i/g_0 + (u - m)(g_{m+1}/g_0).$$

This shows that φ is increasing and piecewise linear, with discontinuities in the derivative only possible at integer values of u.

To show the second part of the lemma we first observe that the right-hand side is a continuous, increasing, piecewise linear function, with discontinuities in the derivative only possible at integer values of u. It takes the value 1 at $u = 0$. It therefore suffices to show that for non-integer values of u, both the left-hand side and the right-hand side have the same derivative.

At a point $m < u < m + 1$, the right-hand side has derivative equal to g_0^{-1} times the number of $s \in G$ such that $i_G(s) < u + 1$. This is $g_0^{-1} \# G_u = [G_0 : G_u]^{-1}$, by definition, and this equals the derivative of the left-hand side. □

Definition 9.4. *For $v \in [0, \infty)$, we define $\psi(v) = \psi_{L/K}(v) = \varphi^{-1}(v)$. Thus ψ is a piecewise linear, increasing homeomorphism of $[0, \infty)$ to itself. The ramification groups G^v are then defined for all real numbers $v \in [0, \infty)$ by the formula $G^v = G_{\psi(v)}$.*

It is clear from the definition that $\cap_{v>0} G^v = \{1\}$ and that $G^0 = G_0 = I_{L/K}$. The filtration is left-continuous, in the sense that $G^v = \cap_{\epsilon>0} G^{v-\epsilon}$ for all $v \geq 0$. We say that a real number v is a jump in the upper-numbering filtration if the inclusion $G^{v+\epsilon} \subset G^v$ is strict for all $\epsilon > 0$. In contrast to the lower numbering, jumps can occur at rational numbers which are not integers!

Lemma 9.5. *Let $\sigma \in G/H$. Then we have the formula*

$$i_{G/H}(\sigma) = \frac{1}{e_{L/E}} \sum_{s \in \sigma} i_G(s).$$

Proof. Choose elements $x \in A_E$, $y \in A_L$ such that $A_E = A_K[x]$ and $A_L = A_K[y]$. Let $\sigma \in G/H$ be non-trivial, and choose $s \in \sigma$. Then we have

$$i_{G/H}(\sigma) = v_E(\sigma(x) - x) = e_{L/E}^{-1} v_L(\sigma(x) - x).$$

Let $f(X) \in A_E[X]$ denote the minimal polynomial of y over E. Then we have $f(X) = \prod_{t \in H}(X - t(y))$, hence $s(f)(X) = \prod_{t \in H}(X - st(y))$, and the right-hand side in the statement of the lemma equals $\frac{1}{e_{L/E}} v_L(s(f)(y))$. Writing $a = \sigma(x) - x$, $b = s(f)(y)$, we must therefore show that a, b generate the same ideal of A_L.

We show the divisibility in each direction. The polynomial $s(f)(X) - f(X) \in A_E[X]$ has all coefficients divisible by $s(x) - x = \sigma(x) - x$, which shows that $s(f)(y) - f(y) = s(f)(y) = b$ is divisible by $s(x) - x = a$. On the other hand, we can write $x = g(y)$ for some polynomial $g(X) \in A_K[X]$ (because $x \in A_L = A_K[y]$). Then $g(X) - x \in A_E[X]$ has y as a root, and we can therefore write $g(X) - x = f(X)h(X)$ for some polynomial $h(X) \in A_E[X]$. Then $g(X) - s(x) = s(f)(X)s(h)(X)$, and hence (evaluating at $X = y$)

$$g(y) - s(x) = x - s(x) = a = s(f)(y)s(h)(y) = bs(h)(y),$$

showing that b divides a. This concludes the proof. □

Lemma 9.6. *1. Let $\sigma \in G/H$, and let $j(\sigma) = \sup_{s \in \sigma} i_G(s)$. Then $i_{G/H}(\sigma) - 1 = \varphi_{L/E}(j(\sigma) - 1)$.*

2. *Suppose that $v = \varphi_{L/E}(u)$. Then $(G/H)_v = im(G_u \to G/H)$.*

Proof. Choose $s \in G$ such that $i_G(s) = j(\sigma) = m$, say. If $t \in H$ then $i_G(st) \leq m$. If t lies in H_{m-1} then $i_G(t) \geq m$, hence $i_G(st) \geq m$, hence $i_G(st) = m$. If $t \notin H_{m-1}$ then $i_G(t) < m$, and so $i_G(st) = i_G(t)$. In either case we have $i_G(st) = \min(i_G(t), m)$. We then have the formula

$$i_{G/H}(\sigma) = \frac{1}{e_{L/E}} \sum_{s \in \sigma} i_G(s) = \frac{1}{h_0} \sum_{t \in H} \min(i_G(t), m) = 1 + \varphi_{L/E}(m-1),$$

which shows the first part. For the second, we have $\mathrm{im}(G_u \to G/H) = G_u H/H$, and then for $\sigma \in G/H$,

$$\sigma \in (G/H)_v \Leftrightarrow \varphi_{L/E}(j(\sigma) - 1) \geq \varphi_{L/E}(u) \Leftrightarrow j(\sigma) \geq u + 1 \Leftrightarrow \sigma \in G_u H/H,$$

as desired. □

Lemma 9.7. *We have the formula $\varphi_{L/K} = \varphi_{E/K} \circ \varphi_{L/E}$.*

Proof. Again both sides are continuous, piecewise linear, increasing homeomorphisms of $[0, \infty)$ into itself. It suffices therefore to show that the derivatives are equal, wherever they are defined. The derivative of the left-hand side at u equals $[G_0 : G_u]^{-1}$, while the derivative of the right-hand side equals $\varphi'_{E/K}(\varphi_{L/E}(u))\varphi'_{L/E}(u) = [G/H : (G/H)_{\varphi_{L/E}(u)}]^{-1}[H_0 : H_u]^{-1}$. By the previous lemma, this is equal to $[G/H : G_u H/H]^{-1}[H_0 : H_u]^{-1} = [G_0 : G_u]^{-1}$, as required. □

Theorem 9.8. *For all $v \geq 0$, we have $(G/H)^v = G^v H/H$.*

Proof. We have $(G/H)^v = (G/H)_{\psi_{E/K}(v)} = G_u H/H$, where $u = \psi_{L/E}(\psi_{E/K}(v))$, or equivalently $v = \varphi_{E/K}(\varphi_{L/E}(u)) = \varphi_{L/K}(u)$, by the lemma. This in turn implies that $G_u = G^v$, as desired. □

Example 9.9. Let $L = \mathbb{Q}_2(\sqrt{2}, i)$, $K = \mathbb{Q}_2(i)$, and $E = \mathbb{Q}_2(\sqrt{2})$. Let $G = \mathrm{Gal}(L/\mathbb{Q}_2)$, $H = \mathrm{Gal}(L/K)$, $N = \mathrm{Gal}(L/E)$. We calculate $G^1 = G$, $G^2 = H$ (and the jumps are at $v = 1, 2$).

We have $(G/H)^1 = G/H$, and the jump is at $v = 1$. We have $(G/N)^2 = G/N$, and the jump is at $v = 2$. This is in accordance with the theorem.

The existence of the upper numbering allows us to generalize the notion of maximal unramified subextension. For simplicity, we will restrict here to the case of Galois extensions only.

Theorem 9.10. *Let L/K be as above, and let $a \geq 0$ be a real number.*

1. *If $a = 0$, then L^a is the maximal unramified subextension of L/K.*

2. *Let $L/E/K$ be an intermediate extension, Galois over K. Then $E^a = L^a \cap E$.*

3. *Let E_1, E_2 be intermediate extensions, Galois over K. Then $E_1^a \cdot E_2^a \subset (E_1 \cdot E_2)^a$. In particular, if $E_1^a = E_1$ and $E_2^a = E_2$ then $(E_1 \cdot E_2)^a = E_1 \cdot E_2$.*

Proof. We have $G^a = G^0 = G_0 = I_{L/K}$, so the first part is clear. For the second, the subgroup of G fixing E^a is the pre-image of $(G/H)^a$ in G; by Herbrand's theorem, this is exactly $G^a H$, which is the subgroup of G fixing $L^a \cap E$.

For the third part, let $H_1 = \mathrm{Gal}(L/E_1)$, $H_2 = \mathrm{Gal}(L/E_2)$. Then $E_1 \cdot E_2$ is the fixed field of $H_1 \cap H_2$, and we must show that the pre-image of $(G/H_1 \cap H_2)^a$ in G, namely $G^a(H_1 \cap H_2)$, is contained inside $(G^a H_1) \cap (G^a H_2)$. This is clear. □

This notion becomes particularly useful in the case of an abelian extension, thanks to the Hasse–Arf theorem:

Theorem 9.11. *Let L/K be an abelian extension of complete fields with k_L/k_K separable, and let $v \geq 0$ be a jump in the ramification filtration (i.e. a real number such that $\mathrm{Gal}(L/K)^v \neq \mathrm{Gal}(L/K)^{v+\epsilon}$ for all $\epsilon > 0$). Then $v \in \mathbb{Z}$.*

Definition 9.12. *Let L/K be a Galois extension of CDVFs with k_L/k_K separable. If $\mathrm{Gal}(L/K)$ is abelian, then we define the conductor of L/K to be the ideal $\mathfrak{f}_{L/K} = \mathfrak{m}_K^a$, where $a \geq 0$ is the smallest non-negative integer such that $\mathrm{Gal}(L/K)^a = \{1\}$ (or equivalently $L^a = L$).*

Proposition 9.13. *Let E/K be a Galois extension of CDVFs with k_E/k_K separable, and let $L_1, L_2/K$ be intermediate extensions, abelian over K. Then $L_1 \dot{L}_2$ is abelian over K, and $\mathfrak{f}_{L_1 \cdot L_2/K} = \mathrm{lcm}(\mathfrak{f}_{L_1/K}, \mathfrak{f}_{L_2/K})$.*

Proof. If $a \in [0, \infty)$, then we have $(L_1 \cdot L_2)^a = (L_1 \cdot L_2) \cap E^a$, hence $\mathrm{Gal}(L_1 \cdot L_2/K)^a = \{1\}$ if and only if $L_1 \cdot L_2 \subset E^a$, if and only if $L_1 \subset E^a$ and $L_2 \subset E^a$, if and only if $\mathrm{Gal}(L_1/K)^a = \{1\}$ and $\mathrm{Gal}(L_2/K)^a = \{1\}$. The result follows. □

10 The different

Let A be a Dedekind domain with field of fractions K, and let V be a finite-dimensional K-vector space.

Definition 10.1. *An A-lattice in K is a finitely generated A-submodule $M \subset V$ which spans V.*

For example, if $e_1, \dots, e_n$ is a K-basis of V, then $\oplus_{i=1}^n Ae_i$ is an A-lattice. If A is a PID (so that every finitely generated torsion-free module is free), then every A-lattice of V has this form.

Now suppose that V is endowed with a symmetric bilinear form $S : V \times V \to K$. In this case, if M is an A-lattice of V we define $M^\vee = \{v \in V \mid S(v, M) \subset A\}$.

Lemma 10.2. *Let M, N be an A-lattice of V.*

1. *$M^\vee$ is an A-lattice of V.*

2. *If $M \subset N$, then $N^\vee \subset M^\vee$.*

3. *If $S \subset A$ is a multiplicative subset, then $S^{-1}M$ is an $S^{-1}A$-lattice of V, and $(S^{-1}M)^\vee = S^{-1}(M^\vee)$.*

4. *We have $M = (M^\vee)^\vee$.*

Proof. We first note that if $M \subset N$, then $N^\vee \subset M^\vee$, from the definition. If $M = \oplus_{i=1}^n Ae_i$, where $e_1, \dots, e_n$ is a K-basis of V, then $M^\vee = \oplus_{j=1}^n Af_j$, where $f_1, \dots, f_j$ is the dual K-basis (with respect to S). In particular, $M^\vee$ is a lattice. In general, we can find a sandwich $M_1 \subset M \subset M_2$, where M_1, M_2 are free A-modules; then we get $M_2^\vee \subset M \subset M_1^\vee$, which shows that M spans V (it contains a lattice) and is finitely generated (it is contained inside a lattice), hence M is itself a lattice.

If M is a lattice and S is a multiplicative subset of A, then $S^{-1}M$ is finitely generated and spans V, so is an $S^{-1}A$-lattice. We have $M^\vee \subset (S^{-1}M)^\vee$, hence $S^{-1}(M^\vee) \subset (S^{-1}M)^\vee$. On the other hand if $b_1, \dots, b_m$ are A-module generators for M and $v \in (S^{-1}M)^\vee$, then we can write $S(v, b_i) = a_i/s_i$ for $a_i \in A, s_i \in S$, hence $S(s_iv, b_i) \in A$, hence $s_1 \dots s_m v \in M^\vee$, hence $v \in S^{-1}(M^\vee)$.

Finally we show $M = (M^\vee)^\vee$. If $m \in M$, then $S(m, M^\vee) = S(M^\vee, m) \subset A$ (because S is symmetric and by the definition of $M^\vee$). This shows $M \subset (M^\vee)^\vee$. If M is A-free, then $(M^\vee)^\vee = M$ (because the dual basis of the dual basis is the original basis). Using that passage to the dual lattice commutes with localization, we see that the inclusion $M \subset (M^\vee)^\vee$ becomes an isomorphism after localization at any non-zero prime ideal $P \subset A$.

We therefore need to show that if $M \subset N$ are A-lattices of V such that $M_P = N_P$ for all such primes P, then $M = N$. Let $n \in N$; then for any non-zero prime ideal P, we can write $n = m_P/s_P$ with $m_P \in M$, $s_p \in A - P$. The ideal generated by all s_P equals the unit ideal, so we can find primes $P_1, \dots, P_r$ and $t_1, \dots, t_r \in A$ such that $\sum_{i=1}^r t_i s_{P_i} = 1$. We get $n = \sum_{i=1}^r t_i s_{P_i} n = \sum_{i=1}^r t_i m_{P_i} \in M$, showing that $M = N$, as desired. □

We apply this formalism in the following situation. Let E/K be a finite separable extension, and let B denote the integral closure of A in E. We have a non-degenerate symmetric K-bilinear form $S : E \times E \to K$ given by the formula $S(x, y) = \mathrm{tr}_{E/K}(xy)$.

Definition 10.3. *The codifferent is $\mathfrak{c}_{B/A} = B^\vee = \{x \in E \mid \mathrm{tr}_{E/K}\, xB \subset A\}$. Note that it is stable under multiplication by B, so is even a non-zero fractional ideal of B (not just a lattice). We define the different as the inverse ideal $\mathfrak{d}_{B/A} = \mathfrak{c}_{B/A}^{-1}$.*

Note that $B \subset B^\vee$, since $\mathrm{tr}_{E/K}(B) \subset A$, so $\mathfrak{d}_{B/A} \subset B$ is in fact an ideal (not just a fractional ideal).

Lemma 10.4. *Let $S \subset A$ be a multiplicative subset. Then $\mathfrak{d}_{S^{-1}B/S^{-1}A} = S^{-1}\mathfrak{d}_{B/A}$.*

Proof. This follows from the two facts that localization commutes with taking the dual lattice, and with taking the inverse of a non-zero fractional ideal. □

Lemma 10.5. *Let $\mathfrak{q} \subset B$ be a non-zero prime lying above the prime $\mathfrak{p}$ of A. Then $\mathfrak{d}_{\widehat{B}_\mathfrak{q}/\widehat{A}_\mathfrak{p}} = \mathfrak{d}_{B/A}\widehat{B}_\mathfrak{q}$.*

Proof. By the previous lemma, we can assume after replacing A by $A_{\mathfrak{p}}$ that A is a DVR. Let $\mathfrak{q} = \mathfrak{q}_1, \dots, \mathfrak{q}_g$ be the primes of B above $\mathfrak{p} = (\pi)$. Then we have an isomorphism

$$\widehat{B} = \varprojlim_i B/\mathfrak{p}^i B \cong \prod_{i=1}^{g} \widehat{B}_{\mathfrak{q}_i},$$

and we define $\widehat{E} = \widehat{B}[1/\pi] = \prod_{i=1}^{g} E_{\mathfrak{q}_i}$. We observe that $\widehat{E}$ is a $\widehat{K}$-vector space of dimension $[E : K]$; and if $x_1, \dots, x_n$ is a K-basis of E, then it is also a $\widehat{K}$-basis of $\widehat{E}$. We define a map $\mathrm{tr}_{\widehat{E}/\widehat{K}} : \widehat{E} \to \widehat{K}$ in the usual way; this restricts to $\mathrm{tr}_{E/K}$ on the subring $E \subset \widehat{E}$.

Define $\mathfrak{c}_{\widehat{B}/\widehat{A}} = \{b \in \widehat{E} \mid \mathrm{tr}_{\widehat{E}/\widehat{K}}\, b\widehat{B} \subset \widehat{A}\}$. Then $\mathfrak{c}_{\widehat{B}/\widehat{A}} = \mathfrak{c}_{B/A} \cdot \widehat{B}$. Indeed, if $x_1, \dots, x_n$ is an A-basis of B, then it is also a $\widehat{A}$-basis of $\widehat{B}$. If $y_1, \dots, y_n$ is the dual K-basis of E with respect to the trace pairing, then we get

$$\mathfrak{c}_{\widehat{B}/\widehat{A}} = \oplus_{i=1}^{n} \widehat{A} \cdot y_i = \widehat{B}\mathfrak{c}_{B/A}.$$

On the other hand, we can decompose $\mathrm{tr}_{\widehat{E}/\widehat{K}} = \prod_{i=1}^{g} \mathrm{tr}_{E_{\mathfrak{q}_i}/K_{\mathfrak{p}}}$, which shows that

$$\mathfrak{c}_{\widehat{B}/\widehat{A}} = \prod_{i=1}^{g} \mathfrak{c}_{\widehat{B}_{\mathfrak{q}_i}/\widehat{A}}.$$

We finally obtain

$$\mathfrak{c}_{\widehat{B}_{\mathfrak{q}_i}/\widehat{A}_{\mathfrak{p}}} = \mathfrak{c}_{\widehat{B}/\widehat{A}} \widehat{B}_{\mathfrak{q}_i} = \mathfrak{c}_{B/A} \widehat{B}_{\mathfrak{q}_i},$$

as desired. □

Lemma 10.6. *Suppose that L/E is a finite separable extension, and let C denote the integral closure of B in L. Then we have the equality $\mathfrak{d}_{C/A} = \mathfrak{d}_{B/A} \cdot \mathfrak{d}_{C/B}$ as ideals of C.*

Proof. We use the formula $\mathrm{tr}_{L/K} = \mathrm{tr}_{E/K} \circ \mathrm{tr}_{L/E}$. Note that if $\mathfrak{a}$ is a non-zero fractional ideal of A, and $\mathfrak{b}$ is a fractional ideal of B, then we have $\mathrm{tr}_{E/K}\, \mathfrak{b} \subset \mathfrak{a}$ if and only if $\mathfrak{b} \subset \mathfrak{a}\mathfrak{c}_{B/A}$. Indeed, we calculate

$$\mathrm{tr}\, \mathfrak{b} \subset \mathfrak{a} \Leftrightarrow \mathfrak{a}^{-1} \mathrm{tr}\, \mathfrak{b} = \mathrm{tr}\, \mathfrak{a}^{-1}\mathfrak{b} \subset A \Leftrightarrow \mathfrak{a}^{-1}\mathfrak{b} \subset \mathfrak{c}_{B/A} \Leftrightarrow \mathfrak{b} \subset \mathfrak{a}\mathfrak{c}_{B/A}.$$

We can then calculate for an arbitrary fractional ideal $\mathfrak{c} \subset E$:

$$\mathfrak{c} \subset \mathfrak{c}_{C/A} \Leftrightarrow \mathrm{tr}_{L/K}\, \mathfrak{c} = \mathrm{tr}_{E/K}\, \mathrm{tr}_{L/E}\, \mathfrak{c} \subset A \Leftrightarrow \mathrm{tr}_{L/E}\, \mathfrak{c} \subset \mathfrak{c}_{B/A} \Leftrightarrow \mathfrak{c} \subset \mathfrak{c}_{B/A}\mathfrak{c}_{C/B}.$$

This shows that $\mathfrak{c}_{C/A} = \mathfrak{c}_{B/A}\mathfrak{c}_{C/B}$. Multiplying by $\mathfrak{d}_{C/A}\mathfrak{d}_{B/A}\mathfrak{d}_{C/B}$ on either side now gives the result. □

Lemma 10.7. *Let $\alpha \in B$ be a primitive element for the extension E/K, and let $f(X) \in K[X]$ be its minimal polynomial. Then $A[\alpha] \subset E$ is a lattice and $A[\alpha]^\vee = f'(\alpha)^{-1} A[\alpha]$.*

Proof. We claim that if $i \geq 0$, then $\operatorname{tr}_{E/K} \alpha^i / f'(\alpha)$ equals 0 if $0 \leq i \leq n-2$; equals 1 if $i = n-1$; and lies in A if $i \geq n$. If $i \geq n$, then α^i is an A-linear combination of the elements $1, \alpha, \dots, \alpha^{n-1}$, so it suffices to establish the claim in the cases $i = 0, \dots, n-1$. Before establishing the claim, we show why it implies the lemma. Let $M = f'(\alpha)^{-1} A[\alpha]$. The claim shows that $\operatorname{tr}_{E/K} MA[\alpha] \subset A$, so $M \subset A[\alpha]^\vee$. Let $f_1, \dots, f_n$ denote the dual basis to $1, \alpha, \dots, \alpha^{n-1}$. We just need to show $f_1, \dots, f_n \in M$. But the claim implies that $f_n = f'(\alpha)^{-1}$ and also for each $i = 1, \dots, n-1$, $f_{n-i} - \alpha^i / f'(\alpha) \in \oplus_{j=0}^{i-1} A f_{n-j}$. By induction, we see that each $f_n, \dots, f_1$ lies in M, as required.

We now establish the claim. Let L/K be the Galois closure of E, and let $\alpha_1, \dots, \alpha_n$ be the Galois conjugates of α. For each $j = 0, \dots, n-1$ we have an identity

$$\sum_{k=1}^{n} \frac{f(X)}{X - \alpha_k} \frac{\alpha_k^j}{f'(\alpha_k)} = X^j.$$

Indeed, both sides are polynomials of degree at most $n-1$ which agree at the n points $X = \alpha_1, \dots, \alpha_n$. The coefficient of X^{n-1} in the left-hand side equals $\sum_{k=1}^n \alpha_k^j / f'(\alpha_k) = \operatorname{tr}_{E/K} \alpha^j / f'(\alpha)$. The coefficient of X^{n-1} in the right-hand side equals 0 if $j < n-1$, and 1 if $j = n-1$. This concludes the proof. □

Proposition 10.8. *Let $\alpha \in B$ be a primitive element for the extension E/K, and let $f(X) \in K[X]$ be its minimal polynomial. Then $(f'(\alpha)) \subset \mathfrak{d}_{B/A}$, with equality if and only if $B = A[\alpha]$.*

Proof. Let $C = A[\alpha]$, and let $\mathfrak{a} = \{b \in B \mid bB \subset C\}$, a non-zero ideal of B. We have for any non-zero $b \in B$:

$$b \in \mathfrak{a} \Leftrightarrow bB \subset C \Leftrightarrow C^\vee \subset (bB)^\vee \Leftrightarrow f'(\alpha)^{-1} C \subset b^{-1} \mathfrak{c}_{B/A} \Leftrightarrow b \in f'(\alpha) \mathfrak{c}_{B/A}.$$

This shows that $\mathfrak{a} = f'(\alpha) \mathfrak{c}_{B/A}$, hence $f'(\alpha) = \mathfrak{a} \mathfrak{d}_{B/A}$. We see that $(f'(\alpha)) \subset \mathfrak{d}_{B/A}$, with equality if and only if $\mathfrak{a} = B$, i.e. $B = C$. □

Proposition 10.9. *Let $\mathfrak{q} \subset B$ be a non-zero prime ideal lying above the prime $\mathfrak{p}$ of A, and let $v_\mathfrak{q} : E^\times \to \mathbb{Z}$ be the corresponding valuation. Suppose that the corresponding extension $k_\mathfrak{q}/k_\mathfrak{p}$ of residue fields is separable. Then $v_\mathfrak{q} \mathfrak{d}_{B/A} \geq e_{\mathfrak{q}/\mathfrak{p}} - 1$, with equality if and only if $e_{\mathfrak{q}/\mathfrak{p}}$ is coprime to the residue characteristic of $\mathfrak{q}$, i.e. the ramification is tame. In particular, $v_\mathfrak{q} \mathfrak{d}_{B/A} = 0$ if and only if $e_{\mathfrak{q}/\mathfrak{p}} = 1$, i.e. $\mathfrak{q}$ is unramified over $\mathfrak{p}$.*

Proof. After localization and completion, we can assume that both A and B are complete DVRs. By the transitivity of the different, and the existence of the maximal unramified subextension of E/K, we can assume either that $e = 1$ or that $e > 1$ and $f = 1$. In the first case, we choose a primitive element $\overline{\alpha}$ for k_E/k_K, and let $f(X) \in A[X]$ be a monic polynomial lifting the minimal polynomial of $\overline{\alpha}$, and $\alpha \in B$ the unique root of $f(X)$ in B lifting $\overline{\alpha}$. We have seen that $B = A[\alpha]$ and $\overline{f}'(\overline{\alpha}) \neq 0$, showing that $\mathfrak{d}_{B/A} = B$ is the unit ideal.

In the second case, we let $\pi \in B$ be a uniformizer and let $f(X) \in A[X]$ be its minimal polynomial. In this case $f(X)$ is an Eisenstein polynomial and $B = A[\pi]$, so we get $\mathfrak{d}_{B/A} = f'(\pi)B$. Using that $f(X)$ is Eisenstein we have

$$f'(\pi) = e\pi^{e-1} \bmod \pi^e B.$$

Thus $v_{\mathfrak{q}}(\mathfrak{d}_{B/A}) \geq e - 1$, with equality if and only if $e \in B^\times$. This completes the proof. □

Corollary 10.10. *Only finitely many primes of A can ramify in B.*

Example 10.11. We calculate the ring of integers of $K = \mathbb{Q}(\alpha)$, where α is a root of $f(X) = X^3 - 2$. Let $C = \mathbb{Z}[\alpha] \subset \mathcal{O}_K$. We claim that in fact $C = \mathcal{O}_K$. This will be the case if and only if $\mathfrak{d}_{\mathcal{O}_K/\mathbb{Z}} = f'(\alpha)\mathcal{O}_K = 3\alpha^2\mathcal{O}_K$. The polynomial $X^3 - 2$ is Eisenstein, so the prime 2 is totally and tamely ramified in $\mathcal{O}_K$. It follows that there is a unique prime $\mathfrak{q}_2 = (\alpha)$ of $\mathcal{O}_K$ above 2, and $v_{\mathfrak{q}_2}(\mathfrak{d}_{\mathcal{O}_K/\mathbb{Z}}) = 2$. We have $(X-1)^3 - 2 = X^3 - 3X^2 + 3X - 3$. This polynomial is Eisenstein, showing that 3 is also totally wildly ramified in K, with $3 = \mathfrak{q}_3^3 = (\alpha+1)^3$. This shows that $v_{\mathfrak{q}_3}(\mathfrak{d}_{\mathcal{O}_K/\mathbb{Z}}) \geq 3$. We find that $3\alpha^2\mathcal{O}_K = \mathfrak{q}_2^2\mathfrak{q}_3^3$ divides $\mathfrak{d}_{\mathcal{O}_K/\mathbb{Z}}$, which divides $f'(\alpha)\mathcal{O}_K = 3\alpha^2\mathcal{O}_K$. Thus equality holds and $C = \mathcal{O}_K$.

11 Cyclotomic fields

If $m \geq 1$ is an integer, then we typically use the notation ζ_m for a primitive m^{th} root of unity, and let $\mathbb{Q}(\zeta_m)$ denote the splitting field of $X^m - 1$. A field of this form is called a cyclotomic field.

We will study the case where $m = p^r$ is a prime power. We first study the local case. Let $\Phi(X) = (X^{p^r} - 1)/(X^{p^{r-1}} - 1)$; then the roots of Φ are exactly the primitive $(p^r)^{\text{th}}$ roots of unity. We claim that $\Phi(X)$ is irreducible over $\mathbb{Q}_p$. Evaluating at $X = 1$, we have

$$\prod_{a \in (\mathbb{Z}/p^r\mathbb{Z})^\times} (\zeta^a - 1) = \Phi(1) = p.$$

Let $\zeta = \zeta_{p^r}$ and $K = \mathbb{Q}_p(\zeta)$. Then $\mathbb{Z}_p[\zeta] \subset \mathcal{O}_K$. We observe that for any $a \in (\mathbb{Z}/p^r\mathbb{Z})^\times$ with $ab \equiv 1 \bmod p^r$, $\zeta - 1$ divides $\zeta^a - 1$ in $\mathcal{O}_K$, as we can write $(\zeta^a - 1)/(\zeta - 1) = 1 + \zeta + \cdots + \zeta^{a-1}$. Similarly, $\zeta^a - 1$ divides $\zeta^{ab} - 1 = \zeta - 1$, as we can write $(\zeta^{ab} - 1)/(\zeta^a - 1) = 1 + \zeta^b + \cdots + \zeta^{(a-1)b}$. It follows that as ideals $(1 - \zeta) = (1 - \zeta^a)$ for any $a \in (\mathbb{Z}/p^r\mathbb{Z})^\times$, and hence (using the above identity) we obtain the equality of ideals

$$(1 - \zeta)^{(p-1)p^{r-1}} = p\mathcal{O}_K.$$

This shows that the ramification index $e_{K/\mathbb{Q}_p}$ is at least $(p-1)p^{r-1}$. Since $e_{K/\mathbb{Q}_p} \leq [K : \mathbb{Q}_p] \leq \deg \Phi(X) = (p-1)p^{r-1}$, it follows that equality holds, $e_{K/\mathbb{Q}_p} = [K : \mathbb{Q}_p] = (p-1)p^{r-1}$, and $1 - \zeta$ is a uniformizer of $\mathcal{O}_K$. Moreover, the injective homomorphism $\mathrm{Gal}(K/\mathbb{Q}_p) \to (\mathbb{Z}/p^r\mathbb{Z})^\times$, which sends $\sigma \in \mathrm{Gal}(K/\mathbb{Q}_p)$ to the unique a such that $\sigma(\zeta) = \zeta^a$, is an isomorphism.

Now let $E = \mathbb{Q}(\zeta)$. Since the polynomial $\Phi(X)$ is irreducible over $\mathbb{Q}_p$, it is irreducible over $\mathbb{Q}$, so we see that $[E : \mathbb{Q}] = [K : \mathbb{Q}_p]$, and there is a unique prime ideal $\mathfrak{p}$ of $\mathcal{O}_E$ above p which satisfies $E_{\mathfrak{p}} = K$. We claim that we even have $\mathcal{O}_E = \mathbb{Z}[\zeta]$. We know that this is true if and only if $\mathfrak{d}_{E/\mathbb{Q}} = (\Phi'(\zeta)) = p^r(1 - \zeta_p)^{-1}\mathcal{O}_E$. This equality holds after localization at p, since $\mathfrak{d}_{E/\mathbb{Q}}$ respects localization and completion and we have $\mathcal{O}_K = \mathbb{Z}_p[\zeta]$. On the other hand, we always have $(\Phi'(\zeta)) \subset \mathfrak{d}_{E/\mathbb{Q}}$, and the left-hand side divides $p^r\mathcal{O}_E$, hence has zero valuation at all primes of $\mathcal{O}_E$ not lying above p. We conclude that equality holds and $\mathcal{O}_E = \mathbb{Z}[\zeta]$. Moreover, no prime other than p ramifies in the field E.

Let $\mathfrak{l}$ be a prime of $E = \mathbb{Q}(\zeta_{p^r})$ lying above a rational prime $l \neq p$; then l is unramified in E. Recall that the decomposition group $D_{\mathfrak{l}/l} \cong \mathrm{Gal}(\mathbb{F}_{\mathfrak{l}}/\mathbb{F}_l)$ is cyclic, generated by the Frobenius element $\mathrm{Frob}_{\mathfrak{l}/l} : x \mapsto x^l$. In fact, $\mathrm{Frob}_{\mathfrak{l}/l} \in \mathrm{Gal}(E/\mathbb{Q}) \cong (\mathbb{Z}/p^r\mathbb{Z})^\times$ is identified with the residue class of l modulo p^r: we have an injection

$$\mathcal{O}_E^\times[p^r] \hookrightarrow \mathbb{F}_{\mathfrak{l}}^\times.$$

By definition, this is compatible with the action of $\mathrm{Frob}_{\mathfrak{l}/l}$ on the left-hand side and the map $x \mapsto x^l$ on the right hand side. It follows that $\mathrm{Frob}_{\mathfrak{l}/l}$ agrees with multiplication by l on $\mathcal{O}_E^\times[p^r]$, which is equivalent to saying that $\mathrm{Frob}_{\mathfrak{l}/l} = l$ in $\mathrm{Gal}(E/\mathbb{Q})$. We observe that $\mathrm{Frob}_{\mathfrak{l}/l} = \mathrm{Frob}_l$ depends only on the prime l, and not on the choice of prime $\mathfrak{l}$ of $\mathcal{O}_E$ above it; this is a general feature of abelian extensions (because of the formula $\sigma\, \mathrm{Frob}_{\mathfrak{l}/l}\, \sigma^{-1} = \mathrm{Frob}_{\sigma(\mathfrak{l})/l}$).

We can use this to give a quick proof of the quadratic reciprocity law. Let p be an odd prime. Since the group $(\mathbb{Z}/p\mathbb{Z})^\times$ is cyclic of order $p-1$, it has a unique index 2 subgroup; by Galois theory, since means that $\mathbb{Q}(\zeta_p)$ contains a unique subfield which is quadratic over $\mathbb{Q}$. Such a subfield can be ramified only at the prime p, which implies (by an earlier calculation) that it must be $\mathbb{Q}(\sqrt{p^*})$, where $p^* = \left(\frac{-1}{p}\right) p$. We can characterize $\mathrm{Gal}(\mathbb{Q}(\zeta_p)/\mathbb{Q}(\sqrt{p^*})) \subset (\mathbb{Z}/p\mathbb{Z})^\times$ as the set of quadratic residues modulo p.

We show that $\left(\frac{p^*}{q}\right) = \left(\frac{q}{p}\right)$ by looking at how q splits in $\mathbb{Q}(\sqrt{p^*})$. We have $\mathrm{Frob}_q = q \bmod p$ in $\mathrm{Gal}(\mathbb{Q}(\zeta_p)/\mathbb{Q})$. The prime q splits in $\mathbb{Q}(\sqrt{p^*})$ if and only $D_{\mathfrak{q}/q} \subset \mathrm{Gal}(\mathbb{Q}(\zeta_p)/\mathbb{Q}(\sqrt{p^*}))$, if and only if q is a square mod p. On the other hand the prime q splits in $\mathbb{Q}(\sqrt{p^*})$ if and only if the equation $X^2 - p^*$ has a solution in $\mathbb{F}_q$, if and only if p^* is a square mod q.

12 Class field theory

We are now going to discuss global class field theory. Let K be a number field. Recall that an extension E/K is said to be abelian if it is Galois with abelian Galois group. The main goal of class field theory is to describe all abelian extensions of K in terms of the 'internal arithmetic' of K. Let I denote the group of fractional ideals of $\mathcal{O}_K$, and $P \subset I$ the subgroup of principal fractional ideals $\alpha\mathcal{O}_K$, $\alpha \in K^\times$. The ideal class group of $\mathcal{O}_K$ is the quotient $H(\mathcal{O}_K) = I/P$. Class field theory gives a description of the abelian extensions of K in terms of so-called ray class groups, which are generalizations of the ideal class group. We begin by describing these.

Definition 12.1. *A divisor of K is a formal product $\mathfrak{c} = \mathfrak{c}_0 \cdot \mathfrak{c}_\infty$, where $\mathfrak{c}_0 \subset \mathcal{O}_K$ is a non-zero ideal and $\mathfrak{c}_\infty$ is a (possibly empty) set of embeddings $\tau : K \hookrightarrow \mathbb{R}$.*

If $\mathfrak{c}, \mathfrak{d}$ are divisors we write $\mathfrak{c} \leq \mathfrak{d}$ if $\mathfrak{c}_0 | \mathfrak{d}_0$ and $\mathfrak{c}_\infty \subset \mathfrak{d}_\infty$. Thus the set of divisors of K is partially ordered, with minimal element $\mathcal{O}_K$ (the unit ideal of $\mathcal{O}_K$ with empty set of infinite places).

(The terminology of divisors is supposed to remind you of algebraic curves, where divisors are formal sums of points.) If $\mathfrak{c}$ is a divisor, then we write $I(\mathfrak{c}) \subset I$ for the group of non-zero fractional ideals which are prime to $\mathfrak{c}_0$, i.e. satisfying $v_\mathfrak{p}(\mathfrak{a}) = 0$ for all $\mathfrak{p}|\mathfrak{c}_0$. We write $K_\mathfrak{c} \subset K^\times$ for the subgroup of $\alpha \in K^\times$ satisfying the following conditions:

- If $\mathfrak{p}$ is a prime of $\mathcal{O}_K$ dividing $\mathfrak{c}_0$, then $v_\mathfrak{p}(\alpha) \geq 0$ and $v_\mathfrak{p}(\alpha - 1) \geq v_\mathfrak{p}(\mathfrak{c}_0)$.
- If $\tau : K \hookrightarrow \mathbb{R}$ is a real embedding which lies in $\mathfrak{c}_\infty$, then $\tau(\alpha) > 0$.

We write $P_\mathfrak{c} \subset P$ for the subgroup of principal fractional ideals of the form (α) for some $\alpha \in K_\mathfrak{c}$. It is clear from the definition that $P_\mathfrak{c} \subset I(\mathfrak{c})$, and we accordingly call the quotient $H(\mathfrak{c}) = I(\mathfrak{c})/P_\mathfrak{c}$ the ray class group of level $\mathfrak{c}$.

We now discuss the relation of the groups $H(\mathfrak{c})$ with the usual ideal class group. Let $P(\mathfrak{c}) = P \cap I(\mathfrak{c})$ denote the subgroup of principal fractional ideals which are prime to $\mathfrak{c}_0$. Then there is are obvious maps $I/P_\mathfrak{c} \to I(\mathfrak{c})/P(\mathfrak{c}) \hookrightarrow I/P$, hence $H(\mathfrak{c}) \to H(\mathcal{O}_K)$.

Proposition 12.2. *1. The natural map $I(\mathfrak{c})/P(\mathfrak{c}) \to I/P = H(\mathcal{O}_K)$ is an isomorphism.*

2. Let $U = \mathcal{O}_K^\times$ and $U_\mathfrak{c} = U \cap K_\mathfrak{c}$. Then there are natural short exact sequences of abelian groups

$$\begin{tikzcd} 0 \arrow[r] & P(\mathfrak{c})/P_\mathfrak{c} \arrow[r] & H(\mathfrak{c}) \arrow[r] & H(\mathcal{O}_K) \arrow[r] & 0 \end{tikzcd}$$

$$\begin{tikzcd} 0 \arrow[r] & U/U_\mathfrak{c} \arrow[r] & (\mathcal{O}_K/\mathfrak{c}_0)^\times \times \prod_{\tau \in \mathfrak{c}_\infty} \{\pm 1\} \arrow[r] & P(\mathfrak{c})/P_\mathfrak{c} \arrow[r] & 0. \end{tikzcd}$$

In particular, $H(\mathfrak{c})$ is a finite group, of order $h_\mathfrak{c} = \frac{h_K \phi(\mathfrak{c})}{[U:U_\mathfrak{c}]}$, where $h_K = \#H(\mathcal{O}_K)$ and

$$\phi(\mathfrak{c}) = 2^{\#\mathfrak{c}_0}(\mathcal{O}_K/\mathfrak{c}_0)^\times.$$

Proof. We begin by recalling the basic finiteness results proved in Part II Number Fields. First, the group $H(\mathcal{O}_K)$ is finite. Second, let $\tau_1, \ldots, \tau_n : K \hookrightarrow \mathbb{C}$ denote the $n = [K : \mathbb{Q}]$ distinct complex embeddings of K. We define r_1 to be the number of embeddings that actually take values in $\mathbb{R}$, and $r_2 = (n - r_1)/2$. The Dirichlet unit theorem says that there is an isomorphism $\mathcal{O}_K^\times \cong \Delta \times \mathbb{Z}^{r_1 + r_2 - 1}$, where Δ is the finite group of roots of unity in K.

To show that the map $I(\mathfrak{c})/P(\mathfrak{c}) \to I/P$ is an isomorphism, it remains to show that it is surjective; in other words, we must show that for any non-zero fractional ideal $\mathfrak{a} \subset K$, we can find $\alpha \in K^\times$ such that $\alpha\mathfrak{a} \in I(\mathfrak{c})$. We can assume without loss of generality that $\mathfrak{a} \subset \mathcal{O}_K$ is a non-zero ideal. By the Chinese remainder theorem, we can find $\alpha \in K^\times$ such that $v_\mathfrak{p}(\alpha) = v_\mathfrak{p}(\mathfrak{a})$ if $\mathfrak{p}|\mathfrak{c}_0$. Then $\alpha^{-1}\mathfrak{a}$ is prime to $\mathfrak{c}_0$, as desired.

We now come to the second part of the proposition. The existence of the first exact sequence follows immediately from the first part of the lemma. To get the second exact sequence, let $K(\mathfrak{c}) \subset K^\times$ denote the set of elements α such that $(\alpha) \in P(\mathfrak{c})$. Then $K_\mathfrak{c} \subset K(\mathfrak{c})$, and we have a commutative diagram with exact rows:

$$\begin{tikzcd} 0 \arrow[r] & U \arrow[r] & K(\mathfrak{c}) \arrow[r] & P(\mathfrak{c}) \arrow[r] & 0 \\ 0 \arrow[r] & U_\mathfrak{c} \arrow[r] \arrow[u] & K_\mathfrak{c} \arrow[r] \arrow[u] & P_\mathfrak{c} \arrow[r] \arrow[u] & 0. \end{tikzcd}$$

By the snake lemma, there is an exact sequence

$$\begin{tikzcd} 0 \arrow[r] & U/U_\mathfrak{c} \arrow[r] & K(\mathfrak{c})/K_\mathfrak{c} \arrow[r] & P(\mathfrak{c})/P_\mathfrak{c} \arrow[r] & 0. \end{tikzcd}$$

To complete the proof, we need to show that the natural map

$$K(\mathfrak{c})/K_{\mathfrak{c}} \to (\mathcal{O}_K/\mathfrak{c}_0)^\times \times \prod_{\tau \in \mathfrak{c}_\infty} \{\pm 1\}$$

is an isomorphism. (On each factor $\tau \in \mathfrak{c}_\infty$, the map is $\alpha \mapsto \operatorname{sign} \tau(\alpha)$.) This map is injective, by definition. To show it is surjective, we observe that the Chinese remainder theorem shows that for any $x \in (\mathcal{O}_K/\mathfrak{c}_0)^\times$, we can find $\alpha \in \mathcal{O}_K$ such that $\alpha \equiv x \bmod \mathfrak{c}_0$ and $\tau(x) > 0$ for all $\tau \in \mathfrak{c}_\infty$. Indeed, we can find an element x satisfying the first condition. Let $N \geq 1$ be an integer in $\mathfrak{c}_0$ (e.g. the integer $[\mathcal{O}_K : \mathfrak{c}_0]$). Then for $a \geq 1$ sufficiently large, $\tau(x + aN) > 0$ for all $\tau \in \mathfrak{c}_\infty$.

We therefore just need to show that for any subset $S \subset \mathfrak{c}_\infty$, we can find $\alpha \in K(\mathfrak{c})$ such that $\tau(\alpha) > 0$ if $\tau \in S$ and $\tau(\alpha) < 0$ if $\tau \notin S$. To accomplish this, choose $\beta \in \mathcal{O}_K$ such that $K = \mathbb{Q}(\beta)$, let $\tau_1, \ldots, \tau_r$ be the elements of $\mathfrak{c}_\infty$, and let $\beta_j = \tau_j(\beta)$. Identify S with a subset of $\{1, \ldots, r\}$. Then there is (e.g. by the Chinese remainder theorem) an isomorphism $\mathbb{R}[X]/(\prod_{j=1}^r (X - \beta_j)) \cong \mathbb{R}^r$, which sends X to $(\beta_1, \ldots, \beta_r)$. We can therefore find a polynomial $f(X) \in \mathbb{R}[X]$ such that $f(\beta_j) > 0$ if $j \in S$ and $f(\beta_j) < 0$ if $j \notin S$. Since $\mathbb{Q}$ is dense in $\mathbb{R}$, we can even assume $f(X) \in \mathbb{Q}[X]$. Clearing denominators (which does not affect signs), we can even assume $f(X) \in \mathbb{Z}[X]$. Then we can take $\alpha = f(\beta)$; for we have $\tau_j \alpha = \tau_j f(\beta) = f(\beta_j)$, which has the required sign. If α is not prime to $\mathfrak{c}_0$, we replace it by $1 + N\alpha$, where $N \in \mathfrak{c}_0$ is a sufficiently large integer. □

Example 12.3. Suppose that $K = \mathbb{Q}$ and $\mathfrak{c} = (N) \cdot \infty$, for some integer $N \geq 1$. Then $U = \{\pm 1\}$ and $U_{\mathfrak{c}}$ is trivial; $I/P = H(\mathbb{Z})$ is trivial. We see that $H(\mathfrak{c}) \cong (\mathbb{Z}/N\mathbb{Z})^\times$ in this case. If $\mathfrak{a} \subset \mathbb{Z}$ is an ideal which is prime to N, then this isomorphism sends the class $[\mathfrak{a}]$ of $\mathfrak{a}$ in $H(\mathfrak{c})$ to the integer m, where $m \geq 1$ is the unique *positive* generator of $\mathfrak{a}$.

Now suppose instead that K is a real quadratic field and $\mathfrak{c} = \mathfrak{c}_\infty$ is the the set of real places. Suppose further that $h_K = 1$, i.e. the ideal class group of $\mathcal{O}_K$ is trivial. By the unit theorem, there is a 'fundamental unit' $\epsilon \in \mathcal{O}_K^\times$ such that every element of $\mathcal{O}_K^\times$ has the form $\pm \epsilon^n$ for a unique $n \in \mathbb{Z}$. Then $H(\mathfrak{c}) = (\{\pm 1\} \times \{\pm 1\})/\{\pm \epsilon^{\mathbb{Z}}\}$. The group $H(\mathfrak{c})$ is trivial if and only if ϵ has opposite signs at the infinite place, if and only if $\mathbf{N}_{K/\mathbb{Q}} \epsilon = -1$. Otherwise, $H(\mathfrak{c})$ has 2 elements.

Either case can occur: examples are given by $\mathbb{Q}(\sqrt{2})$ (which has class number 1 and fundamental unit $1 + \sqrt{2}$) and $\mathbb{Q}(\sqrt{3})$ (which has class number 1 and fundamental unit $2 + \sqrt{3}$).

With these preliminaries out of the way, we can discuss class field theory. Let L/K be an abelian extension of number fields. If $\mathfrak{p}$ is a non-zero prime of $\mathcal{O}_K$ which is unramified in $\mathcal{O}_L$, then we define the Artin symbol $(\mathfrak{p}, L/K) = \mathrm{Frob}_{\mathfrak{q}/\mathfrak{p}}$, where $\mathfrak{q}$ is any prime of $\mathcal{O}_L$ lying above $\mathfrak{p}$.

If $\mathfrak{c}$ is a divisor such that $\mathfrak{c}_0$ is divisible by all primes which ramify in $\mathcal{O}_L$, then we can extend the Artin symbol to a homomorphism $\psi_{L/K} : I(\mathfrak{c}) \to \mathrm{Gal}(L/K)$ by specifying its values on prime ideals: we set $\psi_{L/K}(\mathfrak{p}) = (\mathfrak{p}, L/K)$.

We introduce a divisor $\mathfrak{f}_{L/K}$, which we call the support of L, as follows. We define $\mathfrak{f}_{L/K,\infty}$ to be the set of embeddings $\tau : K \hookrightarrow \mathbb{R}$ which do not extend to an embedding $L \hookrightarrow \mathbb{R}$

(of course, it will extend to an embedding $L \hookrightarrow \mathbb{C}$). We define $\mathfrak{f}_{L/K} = \prod_{\mathfrak{p}} \mathfrak{p}^{a_{L/K}}$, where $a_{L/K}$ is the least integer $a \geq 0$ such that $D^a_{\mathfrak{q}/\mathfrak{p}} = 1$, and $\mathfrak{q}$ is any choice of prime of $\mathcal{O}_L$ above $\mathfrak{p}$. Note that $a_{L/K} > 0$ if and only if $\mathfrak{p}$ is ramified in L; thus this product has only finitely many terms that are not 1.

Let $\mathfrak{c}$ be a divisor of K such that $\mathfrak{c} \geq \mathfrak{f}_{L/K}$. It follows from the definition that $\mathfrak{c}_0$ is divisible by all prime ideals of $\mathcal{O}_K$ which are ramified in $\mathcal{O}_L$, so the map $\psi_{L/K} : I(\mathfrak{c}) \to \mathrm{Gal}(L/K)$ is defined. The first main theorem of class field theory is then as follows:

Theorem 12.4. *Let L/K be an abelian extension of number fields, and let $\mathfrak{c} \geq \mathfrak{f}_{L/K}$ be a divisor of K. Then the map $\psi_{L/K}$ is surjective and its kernel contains $P_{\mathfrak{c}}$. It particular, it factors through $I(\mathfrak{c}) \to I(\mathfrak{c})/P_{\mathfrak{c}} = H(\mathfrak{c})$, giving a surjective homomorphism $\psi_{L/K} : H(\mathfrak{c}) \to \mathrm{Gal}(L/K)$.*

The second main theorem of class field theory is as follows:

Theorem 12.5. *Let $\mathfrak{c}$ be a divisor of K. Then there is a canonical bijection between the following two sets:*

1. *The set of abelian extensions L/K such that $\mathfrak{f}_{L/K} \leq \mathfrak{c}$.*

2. *The set of subgroups of the finite group $H(\mathfrak{c})$.*

The bijection is given by the map $L/K \mapsto \ker \psi_{L/K}$. In particular, the maximal abelian extension $L_{\mathfrak{c}}/K$ of support at most $\mathfrak{c}$, which we call the ray class field of level $\mathfrak{c}$, satisfies $\mathrm{Gal}(L_{\mathfrak{c}}/K) \cong H(\mathfrak{c})$.

An important point is that if $L_1, L_2/K$ are abelian extensions of K (say inside a fixed algebraic closure $\overline{K}/K$), and $\mathfrak{f}_{L_1/K} \leq \mathfrak{c}$, $\mathfrak{f}_{L_2/K} \leq \mathfrak{c}$, then $\mathfrak{f}_{L_1 \cdot L_2/K} \leq \mathfrak{c}$. Indeed, they are subfields of the maximal extension $L_{\mathfrak{c}}/K$ of support $\mathfrak{c}$. This explains why we need to use the upper ramification subgroups in defining these objects.

We now discuss examples. The basic example is when $\mathfrak{c} = \mathcal{O}_K$ is the trivial divisor, and $H(\mathfrak{c}) = H(\mathcal{O}_K)$ is the usual ideal class group. Then class field theory says that there is an abelian extension L/K which is everywhere unramified (and in which the real embeddings remain real) for which the Artin map gives an isomorphism $H(\mathcal{O}_K) \cong \mathrm{Gal}(L/K)$. The field L is called the Hilbert class field of K, and is contained inside every ray class field of K.

As an example, consider the polynomial $f(X) = X^3 - X + 1$. As $X^3 + aX + b$ has discriminant $-4a^3 - 27b^2$, the polynomial $f(X)$ has discriminant -23, and is irreducible (even over $\mathbb{F}_3$). Thus the splitting field L of $f(X)$ has Galois group S_3 over $\mathbb{Q}$, and contains the quadratic extension $K = \mathbb{Q}(\sqrt{-23})$. We claim that L is the Hilbert class field of L/K. For this we need to know two things:

- The extension L/K is everywhere unramified. (Since K has no real places, these do not play a role.) This will show that L is contained inside the Hilbert class field of K.

- The class number of K is 3. This will show that L actually equals the Hilbert class field of K.

The class number can be calculated using either the Minkowski bound (as in Part II Number Fields) or using the theory of binary quadratic forms (see the next section). To show that L/K is everywhere unramified, we can refer to an earlier calculation (we showed that for a polynomial like $f(X)$, all of the inertia groups in its splitting field are generated by transpositions).

The isomorphism $H(\mathcal{O}_K) \cong \mathrm{Gal}(L/K)$ sends a prime $\mathfrak{p}$ to $(\mathfrak{p}, L/K)$. In particular, $\mathfrak{p}$ splits in L/K if and only if $\mathrm{Frob}_{\mathfrak{p}} = (\mathfrak{p}, L/K)$ is trivial, if and only if $\mathfrak{p}$ is principal. For example, $f(X)$ is irreducible over $\mathbb{F}_3$, and $3\mathcal{O}_K = \mathfrak{p}_3\overline{\mathfrak{p}}_3$ splits in $K = \mathbb{Q}(\sqrt{-23})$. We see that $\mathfrak{p}_3$ is not a principal ideal.

For an example where the infinite primes play a role, let us consider the field $K = \mathbb{Q}(\sqrt{3})$, with $h_K = 1$ and fundamental unit $2 + \sqrt{3}$ of norm 1. Let $\mathfrak{c} = \mathfrak{c}_\infty$ denote the divisor consisting of all infinite places. We observed that $H(\mathfrak{c})$ is cyclic of order 2. Class field theory therefore tells us that $\mathbb{Q}(\sqrt{3})$ has no everywhere unramified extensions which embed in $\mathbb{R}$, but does have an everywhere unramified quadratic extension if we allow embeddings in $\mathbb{C}$. It is given by $L = K(i) = K(\sqrt{-3})$.

The class field theory of $\mathbb{Q}$ is particularly explicit. Consider the divisor $\mathfrak{c} = (N) \cdot \infty$, for an integer $N \geq 1$; every divisor is dominated by one of this form. We have already computed that the ray class group $H(\mathfrak{c})$ is isomorphic to $(\mathbb{Z}/N\mathbb{Z})^\times$. On the other hand, we have computed (see example sheet 3) that the support of the cyclotomic field $\mathbb{Q}(\zeta_{p^r})/\mathbb{Q}$ is exactly $(p^r) \cdot \infty$ (except if $p^r = 2$, in which case it is trivial). If $N = \prod_{i=1}^k p_i^{r_i}$, then $\mathbb{Q}(\zeta_N)$ is the composite of the fields $\mathbb{Q}(\zeta_{p_i}^{r_i})$, hence has support $\leq \mathfrak{c}$. By the first main theorem, we find that the Artin map gives a surjective homomorphism $(\mathbb{Z}/N\mathbb{Z})^\times \to \mathrm{Gal}(\mathbb{Q}(\zeta_N)/\mathbb{Q})$, which sends a prime ideal (p) (p a prime number not dividing N) to Frob_p. We know from Part II Galois theory (or can prove directly) that the group $\mathrm{Gal}(\mathbb{Q}(\zeta_N)/\mathbb{Q})$ has cardinality $\#(\mathbb{Z}/N\mathbb{Z})^\times$, so this surjective map must in fact be an isomorphism, showing that $\mathbb{Q}(\zeta_N)$ is in fact the ray class field of level $(N) \cdot \infty$. Applying the second main theorem, we deduce the Kronecker–Weber theorem:

Theorem 12.6. *Let $K/\mathbb{Q}$ be an abelian extension. Then there exists an integer $N \geq 1$ such that $K \subset \mathbb{Q}(\zeta_N)$.*

13 Binary quadratic forms

In the remainder of the course, we will analyze the following problem: when is a given prime p represented by a given positive definite binary quadratic form $f(x, y)$? We have for an odd prime p, $p = x^2 + y^2$ if and only if $p \equiv 1 \bmod 4$ (Fermat). We also have $p = x^2 + 2y^2$ if and only if $p \equiv 1, 3 \bmod 8$ (Euler), and $p = x^2 + 5y^2$ if and only if $p \equiv 1, 9 \bmod 20$ (Gauss).

On the other hand, one can show that $p = x^2 + 14y^2$ if and only if the equations $x^2 = -14$ and $(y^2 + 1)^2 = 8$ have a solution in $\mathbb{F}_p$. This shows that the problem cannot always be described in terms of congruence conditions on p. We'll see that the solution to this problem is intimately tied up with the class field theory of imaginary quadratic fields.

Definition 13.1. *A binary quadratic form is a function $f(x, y) = ax^2 + bxy + cy^2$ where a, b, c are integers. We say that the form f represents a given integer m if there exist values*

$x_0, y_0 \in \mathbb{Z}$ such that $f(x_0, y_0) = m$.

Two forms $f(x, y)$, $g(u, v)$ are said to be properly equivalent (or just equivalent) if there exists a matrix $\begin{pmatrix} A & B \\ C & D \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z})$ such that $g(u, v) = f(Au + Bv, Cu + Dv)$. The discriminant of a binary quadratic form is the integer $\Delta(f) = b^2 - 4ac$. The form $f(x, y)$ is primitive if its coefficients a, b, c are coprime.

Lemma 13.2. *If the forms $f(x, y)$ and $g(x, y)$ are equivalent, then $\Delta(f) = \Delta(g)$ and f, g represent the same integers. The form $f(x, y)$ is positive definite (i.e. represents only positive integers) if and only if $\Delta(f) < 0$ and $a > 0$.*

Proof. It is clear that equivalent forms represent the same integers. A calculation shows that the discriminant is invariant. If $f(x, y)$ is positive definite, then $a \neq 0$, and we can write $f(x, y) = \frac{1}{a}((ax + \frac{b}{2}y)^2 + (ac - \frac{b^2}{4})y^2)$. Since $f(x, y)$ is definite, we must have $\Delta(f) < 0$. Since $f(x, y)$ is positive definite, we must have $a > 0$. Conversely, if $a > 0$ then we can write f in this form, and if $\Delta(f) < 0$ then this expression shows that f is positive definite. □

If $D \leq 1$ is an integer which equals $\Delta(f)$ for some binary quartic form f, then $D \equiv 0, 1 \bmod 4$. Conversely, every such integer D appears as the discriminant of a primitive positive definite binary quartic form: we can take $x^2 - \frac{D}{4}y^2$ or $x^2 + xy + \frac{1-D}{4}y^2$. We call these forms the principal forms of discriminant D. In this case, we write $C(D)$ for the set of equivalence classes of primitive positive definite binary quadratic forms of discriminant D.

We will see that the arithmetic of the binary quadratic forms of discriminant D is intimately tied up with the arithmetic of the quadratic field $K = \mathbb{Q}(\sqrt{D})$. To this end, we recall that a $\mathbb{Z}$-submodule $M \subset K$ is called a $\mathbb{Z}$-lattice if it spans K as a $\mathbb{Q}$-vector space and is finitely generated as a $\mathbb{Z}$-module.

If M is a $\mathbb{Z}$-lattice of K, then we define $\operatorname{disc} M = \det \begin{pmatrix} \alpha & \beta \\ \overline{\alpha} & \overline{\beta} \end{pmatrix}^2$, where $\{\alpha, \beta\}$ is a $\mathbb{Z}$-basis of M. This is clearly independent of the choice of basis.

Definition 13.3. *An order in K is a $\mathbb{Z}$-lattice $\mathcal{O} \subset K$ that is also a subring.*

Proposition 13.4. *1. Let $\mathcal{O} \subset K$. Then $\mathcal{O} = \mathbb{Z} + c\mathcal{O}_K$ for a uniquely determined integer $c \geq 1$. In particular, $\mathcal{O} \subset \mathcal{O}_K$ and $\operatorname{disc} \mathcal{O} = c^2 \operatorname{disc} \mathcal{O}_K$.*

2. Let $M \subset K$ be a $\mathbb{Z}$-lattice, and let $\mathcal{O}_M = \{x \in K \mid xM \subset M\}$. Then $\mathcal{O}_M$ is an order.

Proof. For the first part, let $\alpha \in \mathcal{O}_K$ be such that $\mathcal{O}_K = \mathbb{Z} \oplus \mathbb{Z}\alpha$. Then $\mathcal{O}_c = \mathbb{Z} + c\mathcal{O}_K = \mathbb{Z} \oplus \mathbb{Z}c\alpha$ is a subring, hence an order, and $[\mathcal{O}_K : \mathcal{O}_c] = c$. Conversely, suppose that $\mathcal{O} \subset K$ is an order. Then $\mathcal{O}$ is a finitely generated $\mathbb{Z}$-module, so is integral over $\mathbb{Z}$, so contained in $\mathcal{O}_K$. Let $c = [\mathcal{O}_K : \mathcal{O}]$. We have $\mathbb{Z} \subset \mathcal{O}$ (since $1 \in \mathcal{O}$), which means that $c = [\mathcal{O}_K/\mathbb{Z} : \mathcal{O}/\mathbb{Z}]$. The quotient $\mathcal{O}_K/\mathbb{Z}$ is a cyclic group of infinite order, generated by the element α; so $\mathcal{O}/\mathbb{Z}$ is the subgroup generated by $c\alpha$, hence $\mathcal{O} = \mathbb{Z} \oplus \mathbb{Z}c\alpha = \mathcal{O}_c$.

For the second part, we note that $\mathcal{O}_M$ is a ring and clearly spans K (since for any $x \in K$, we have $Nx \in \mathcal{O}_M$ for some integer $N \geq 1$). It will be an order if it is finitely generated over $\mathbb{Z}$. But we get, by definition, an injection $\mathcal{O}_M \hookrightarrow \mathrm{End}_{\mathbb{Z}}(M)$ into a finitely generated $\mathbb{Z}$-module, so $\mathcal{O}_M$ is itself finitely generated. □

Corollary 13.5. *For each integer $D \leq -1$ such that $D \equiv 0, 1 \bmod 4$, there is a unique imaginary quadratic field K and order $\mathcal{O}_D \subset K$ such that* $\operatorname{disc} \mathcal{O}_D = D$.

Proof. If $\mathcal{O} \subset K$ is an order of discriminant D, then $K = \mathbb{Q}(\sqrt{D})$, $\mathcal{O} \subset \mathcal{O}_K$, and hence $\mathcal{O}$ is the unique order $\mathbb{Q}(\sqrt{D})$ such that $\operatorname{disc} \mathcal{O} = [\mathcal{O}_K : \mathcal{O}]^2 \operatorname{disc} \mathcal{O}_K$. This shows uniqueness.

To show existence, we set $\mathcal{O} = \mathbb{Z}[\sqrt{D}]$ or $\mathcal{O} = \mathbb{Z}[(1+\sqrt{D})/2]$, depending on the value of D modulo 4. A calculation shows that this order has the correct discriminant. □

From now on, we will write $\mathcal{O}_D$ for the unique order of discriminant D. We will always view the field $\mathbb{Q}(\sqrt{D})$ as being a subfield of $\mathbb{C}$, with $\sqrt{D}$ the square root of D which has positive imaginary part.

This gives us a way to classify lattices $M \subset K$: we have a discrete invariant, namely the order $\mathcal{O}_M \subset K$. We say that two lattices M, M' are equivalent if there exists $\lambda \in K^\times$ such that $M' = \lambda M$; then clearly we have $\mathcal{O}_M = \mathcal{O}_{M'}$. If $M \subset K$ is a lattice, we define its norm $\mathbf{N}M$ to be the index $[\mathcal{O}_M : M]$, if $M \subset \mathcal{O}_M$; otherwise, we choose $a \geq 1$ such that $aM \subset \mathcal{O}_M$, and define the norm $\mathbf{N}M = a^{-2}[\mathcal{O}_M : aM]$. This is clearly independent of the choice of a. In all cases we have $\operatorname{disc} M = (\mathbf{N}M)^2 \operatorname{disc} \mathcal{O}_M$, and if $\lambda \in K^\times$ then $\mathbf{N}(\lambda M) = \mathbf{N}_{K/\mathbb{Q}}(\lambda)\mathbf{N}M$.

We can compute $\mathcal{O}_M$ as follows. Choose a $\mathbb{Z}$-basis α, β of M. After multiplying through by α^{-1}, we can assume that $\alpha = 1$. We then apply the following lemma:

Lemma 13.6. *Let $M = \mathbb{Z} \oplus \mathbb{Z}\beta$, where $\beta \in K - \mathbb{Q}$ lies in an imaginary quadratic field. Let $f(X) = aX^2 + bX + c \in \mathbb{Z}[X]$ be the unique polynomial such that $f(\beta) = 0$, $a > 0$, and a, b, c are coprime. Then $\mathcal{O}_M = \mathbb{Z} \oplus \mathbb{Z}a\beta$,* $\operatorname{disc} \mathcal{O}_M = b^2 - 4ac$, *and* $\mathbf{N}M = a^{-1}$.

Proof. The proof is by direct calculation. Let $\gamma = A + B\beta$ with $A, B \in \mathbb{Q}$, and suppose that $\gamma \in \mathcal{O}_M$; equivalently, $A + B\beta \in M$ and $(A + B\beta)\beta \in M$. This happens if and only if the rational numbers $A, B, (A - Bb/a)$ and Bc/a are all in fact integers; equivalently, if and only if $A, B, Bb/a$ and Bc/a are all integers. Since $(a, b, c) = 1$, this is equivalent to asking that $A \in \mathbb{Z}$ and $B \in a\mathbb{Z}$, i.e. $\gamma \in \mathbb{Z} \oplus a\beta\mathbb{Z}$. We then have $\operatorname{disc} \mathcal{O}_M = (a\beta - a\overline{\beta})^2 = b^2 - 4ac$, $\mathbf{N}M = a^{-1}$, as claimed. □

The connection with binary quadratic forms is made as follows. Fix an identification $K = \mathbb{Q}(\sqrt{D}) \subset \mathbb{C}$, where $\sqrt{D}$ is the square-root with positive imaginary part. Choose a $\mathbb{Z}$-basis α, β of M such that β/α has positive imaginary part. We define a binary quadratic form $f = f_M$ by the formula

$$f(x, y) = \mathbf{N}_{K/\mathbb{Q}}(\alpha x + \beta y)/\mathbf{N}M = \frac{(\alpha x + \beta y)(\overline{\alpha}x + \overline{\beta}y)}{\mathbf{N}M}.$$

This form has discriminant $\operatorname{disc} M/\mathbf{N}M^2 = \operatorname{disc} \mathcal{O}_M$. It depends on the choice of basis, but any other basis differs from the chosen one by the action of $\mathrm{SL}_2(\mathbb{Z})$. (This is because a change of basis preserves the sign of the imaginary part if and only if it lies in $\mathrm{SL}_2(\mathbb{Z}) \subset \mathrm{GL}_2(\mathbb{Z})$.) It follows that the equivalence class of f_M depends only on M and not on the choice of basis.

Theorem 13.7. *Let $D \leq -1$ be an integer congruent to $0, 1 \bmod 4$, and let $K = \mathbb{Q}(\sqrt{D})$. The map $M \mapsto f_M$ induces a bijection between the following two sets:*

1. *The set of equivalence classes of lattices $M \subset K$ such that $\mathcal{O}_M = \mathcal{O}_D$.*

2. *The set of equivalence classes of primitive positive definite binary quadratic forms of discriminant D.*

Proof. A change of basis changes by a linear substitution in $\mathrm{SL}_2(\mathbb{Z})$, so the map is well-defined. We first show that this map is injective. If M_1, M_2 are lattices such that f_{M_1} and f_{M_2} are equivalent, then we can find bases α, β of M_1 and α', β' of M_2 such that $f_{M_1} = f_{M_2}$ (i.e. the forms are actually equal). Since scaling M and its basis does not change f_M, we can further assume that $\alpha = \alpha' = 1$, hence we have $\mathbf{N}_{K/\mathbb{Q}}(x+\beta y)/\mathbf{N}M_1 = \mathbf{N}_{K/\mathbb{Q}}(x+\beta' y)/\mathbf{N}M_2$ for all $x, y \in \mathbb{Z}$. Comparing coefficients of x^2, xy, y^2, we see that $\mathbf{N}M_1 = \mathbf{N}M_2$, and β, β' have the same characteristic minimal polynomial over $\mathbb{Q}$. Since β, β' are non-real with positive imaginary part, this implies $\beta = \beta'$, hence $M_1 = M_2$.

We now show that the map is surjective. Let $f(x,y) = ax^2 + bxy + cy^2$ be a primitive form of discriminant D, and let $\beta \in K$ be the unique root of $f(x,1) = ax^2 + bx + c$ in K with positive imaginary part. Let $M = \mathbb{Z} \oplus \mathbb{Z}\beta$. The lemma shows that $\mathcal{O}_M = \mathbb{Z} \oplus \mathbb{Z}a\beta$, $\operatorname{disc} \mathcal{O}_M = D$, $\mathbf{N}M = a^{-1}$; and then the form f_M is $\mathbf{N}_{K/\mathbb{Q}}(x+\beta y)/\mathbf{N}M = a(x+\beta y)(x+\overline{\beta}y) = a(x^2 + bxy/a + cy^2a) = f(x,y)$, as desired. □

The set $C(D)$ of equivalence classes of lattices with $\mathcal{O}_M = \mathcal{O}_D$ in fact forms a group under multiplication. If $D = \operatorname{disc} \mathcal{O}_K$ then this is just the usual ideal class group:

Theorem 13.8. *Let $D \equiv 0, 1$ mod 4 be a negative integer. Then the set $C(D)$ becomes a group under the law $[M] \cdot [M'] = [MM']$. If $m \geq 1$ is an integer, then a primitive form $f(x,y) = f_M(x,y)$ of discriminant D represents the integer m if and only if the inverse class $[M]^{-1}$ contains an proper ideal $\mathfrak{a} \subset \mathcal{O}_D$ such that $\mathbf{N}\mathfrak{a} = m$.*

Proof. If M is any $\mathbb{Z}$-lattice of K, we write $\overline{M}$ for its complex conjugate. We first check that for any choice of M, we have $M\overline{M} = \mathbf{N}M\mathcal{O}_M$. We can assume after rescaling that $M = \mathbb{Z} \oplus \mathbb{Z}\beta$, where β satisfies the polynomial $aX^2 + bX + c$, $a, b, c \in \mathbb{Z}$ coprime integers. Then $M\overline{M} = \langle 1, \beta, \overline{\beta}, \beta\overline{\beta} \rangle = \langle 1, \beta, b/a, c/a \rangle$. Since a, b, c are coprime, this equals $\langle 1/a, \beta \rangle = 1/a\mathcal{O}_M = \mathbf{N}M\mathcal{O}_M$.

If $\mathcal{O}, \mathcal{O}' \subset K$ are orders and $z\mathcal{O} = \mathcal{O}'$ for some rational number z, then $\mathcal{O} = \mathcal{O}'$ and $z = \pm 1$. Therefore if M, M' are lattices such that $\mathcal{O}_M = \mathcal{O}'_M = \mathcal{O}$, then we have $(MM')(\overline{MM'}) = (M\overline{M})(M'\overline{M}') = \mathbf{N}M\mathbf{N}M'\mathcal{O}_D = \mathbf{N}(MM')\mathcal{O}_{MM'}$. We conclude that $\mathcal{O}_{MM'} = \mathcal{O}_D$, showing that the set $C(D)$ is preserved by multiplication of lattices. We also see that $\mathbf{N}(MM') = \mathbf{N}M\mathbf{N}M'$. The lattice $\mathcal{O}_D$ is clearly a multiplicative identity, so to show that $H(D)$ is a group we just need to show the existence of inverses. But we have $\mathcal{O}_{\overline{M}} = \overline{\mathcal{O}}_M = \mathcal{O}_D$, so this follows from the identity $M\overline{M} = \mathbf{N}M\mathcal{O}_D$.

It remains to show that f_M represents m if and only if there exists $M' \in [M]^{-1}$ such that $\mathbf{N}M' = m$ and $M' \subset \mathcal{O}_D$. The form f_M represents the integer m if and only if there exists $\gamma \in M$ such that $\mathbf{N}_{K/\mathbb{Q}}(\gamma) = m\mathbf{N}M$. If there exists such a γ, let $M' = \gamma M^{-1}$, where $MM^{-1} = \mathcal{O}_D$. Then $\mathbf{N}M' = \mathbf{N}_{K/\mathbb{Q}}(\gamma)\mathbf{N}M^{-1} = m$, and if $a \in M^{-1}$ then $a\gamma \in \mathcal{O}_D$, hence $M' \subset \mathcal{O}_D$. Conversely, if there exists $M' \in [M]^{-1}$ such that $\mathbf{N}M' = m$ and $M' \subset \mathcal{O}_D$, we write $M' = \gamma M^{-1}$ for $\gamma \in K$. Then $\mathbf{N}M' = \mathbf{N}_{K/\mathbb{Q}}(\gamma)\mathbf{N}(M)^{-1} = m$, hence $\mathbf{N}_{K/\mathbb{Q}}(\gamma) =$

$m\mathbf{N}M$. Moreover, we have $MM' = \gamma\mathcal{O}_D \subset M\mathcal{O}_D = M$, hence $\gamma \in M$. This completes the proof. □

We thus have a group $C(D)$ that tells us a lot about representation of integers by primitive binary quadratic forms of discriminant D. In order to relate this to class field theory, we need to relate $C(D)$ to generalized ideal class groups. If $D = \operatorname{disc}\mathcal{O}_K$, then $C(D) = H(\mathcal{O}_K)$ is the usual ideal class group and the relation is immediate. We thus obtain the following corollary:

Corollary 13.9. *Suppose that $D = \operatorname{disc}\mathcal{O}_K$, and let p be a prime number not dividing D. Then the binary quadratic form $f(x, y) = x^2 - Dy^2/4$ (resp. $x^2 + xy + (1 - D)y^2/4$) represents p if and only if p splits in the Hilbert class field of K.*

Proof. We first observe that if $D \leq -1$ is any negative integer congruent to 0 mod 4, then the identity element of $C(D)$ corresponds to the form $x^2 - Dy^2/4$. Indeed, we have $\mathcal{O}_D = \mathbb{Z} \oplus \mathbb{Z}\sqrt{D}/2$, so the corresponding form is $\mathbf{N}_{K/\mathbb{Q}}(x + \sqrt{D}y/2) = x^2 - Dy^2/4$. Similarly, if $D \equiv 1 \bmod 4$ then the identity element corresponds to the form $x^2 + xy + (1 - D)y^2/4$.

We see that when $D = \operatorname{disc}\mathcal{O}_K$, so that $C(D) = H(\mathcal{O}_K)$, the prime p is represented by this form if and only if there exists an ideal $\mathfrak{a} \subset \mathcal{O}_K$ such that $\mathbf{N}\mathfrak{a} = p$ and $\mathfrak{a}$ is principal. The first condition forces $\mathfrak{a}$ to be prime (as norm is multiplicative) and p either to be ramified or split in $\mathcal{O}_K$. Since p does not divide D, it splits in K as $\mathfrak{p}\overline{\mathfrak{p}}$, and the condition that $\mathfrak{p}$ is principal is then equivalent, by class field theory, to the condition that $\mathfrak{p}$ splits in the Hilbert class field of K. □

We can use this to understand our first examples. If $D = -4 = \operatorname{disc}\mathbb{Z}[i]$, then the principal form is $x^2 + y^2$. The field $\mathbb{Q}(i)$ has class number 1, hence trivial Hilbert class field, so we see that an odd prime p is represented by the form $x^2 + y^2$ if and only if it splits $\mathbb{Q}(i)$, if and only if $x^2 = -1$ has a solution modulo p, if and only if $p \equiv 1 \bmod 4$.

Similarly the field $\mathbb{Q}(\sqrt{-2})$ ($\operatorname{disc}\mathcal{O}_K = -8$) has class number 1, and p is represented by $x^2 + 2y^2$ if and only if the equation $x^2 = -2$ has a solution modulo p, if and only if $p \equiv 1, 3 \bmod 8$.

The field $\mathbb{Q}(\sqrt{-5})$ ($\operatorname{disc}\mathcal{O}_K = -20$) has class number 2, and its Hilbert class field is $\mathbb{Q}(\sqrt{-5}, \sqrt{5}) = \mathbb{Q}(\sqrt{-5}, i)$. The prime p is represented by $x^2 + 5y^2$ if and only if p splits in $\mathbb{Q}(\sqrt{-5})$, if and only if the equations $x^2 = -1$ and $y^2 = 5$ have solutions in $\mathbb{F}_p$, if and only if $p \equiv 1, 9 \bmod 20$.

Finally, the field $\mathbb{Q}(\sqrt{-14})$ ($\operatorname{disc}\mathcal{O}_K = -4\times 14$) has class number 4. You'll show on the example sheet that its Hilbert class field is $\mathbb{Q}(\sqrt{-14}, \sqrt{2\sqrt{2} - 1})$. The minimal polynomial of the element $\sqrt{2\sqrt{2} - 1}$ is $(X^2 + 1)^2 - 8$. It follows that for a prime not dividing the discriminant of this polynomial, p is represented by $x^2 + 14y^2$ if and only if p splits in the Hilbert class field, if and only if the equations $x^2 = -14$ and $(y^2 + 1)^2 = 8$ have solutions in $\mathbb{F}_p$.

This is the story for D of the form $\operatorname{disc}\mathcal{O}_K$. To understand what happens for general discriminants D, we need to work a bit harder. We start with a lemma.

Lemma 13.10. *Let $\mathcal{O}$ be an order in an imaginary quadratic field, and let $c \geq 1$ be an integer. Then the following two sets are in canonical bijection:*

1. *The set of non-zero ideals $\mathfrak{a} \subset \mathcal{O}$ such that $\mathfrak{a} + c\mathcal{O} = \mathcal{O}$ (i.e. $\mathfrak{a}$ is coprime to $c\mathcal{O}$).*

2. *The set of non-zero ideals $\mathfrak{b}$ of $\mathcal{O}[1/c]$.*

The bijection is given by $\mathfrak{a} \mapsto \mathfrak{a}[1/c]$ and $\mathfrak{b} \mapsto \mathfrak{b} \cap \mathcal{O}$. In particular, it preserves multiplication of ideals. Moreover, we have $\mathcal{O}/\mathfrak{a} \cong \mathcal{O}[1/c]/\mathfrak{b}$.

Proof. We first note that for a non-zero ideal $\mathfrak{a} \subset \mathcal{O}$, $\mathfrak{a} + c\mathcal{O} = \mathcal{O}$ if and only if $\mathfrak{a} = \mathfrak{a}[1/c] \cap \mathcal{O}$, if and only if $\mathcal{O}/\mathfrak{a}$ has order prime to c. Indeed, the first condition says that multiplication by c on $\mathcal{O}/\mathfrak{a}$ is surjective. The second condition says that multiplication by c is injective. Since $\mathcal{O}/\mathfrak{a}$ is a finite group, these conditions are equivalent.

If $\mathfrak{a} \subset \mathcal{O}$ is a non-zero ideal prime to c, then $\mathfrak{a} \subset \mathfrak{a}[1/c] \cap \mathcal{O}$. The condition that $\mathfrak{a}$ is prime to c implies that equality holds.

If $\mathfrak{b} \subset \mathcal{O}[1/c]$ is a non-zero ideal, then $\mathfrak{a} = \mathfrak{b} \cap \mathcal{O}$ is prime to c, because if $x \in \mathcal{O}$ and $cx \in \mathfrak{a}$, then $x = c^{-1}cx \in \mathfrak{a}$. We have $\mathfrak{a}[1/c] \subset \mathfrak{b}$. To show equality, note that if $x \in \mathfrak{b}$ then $c^n x \in \mathcal{O}$ for some $n \geq 1$, hence $c^n x \in \mathfrak{a}$, hence $x \in \mathfrak{a}[1/c]$.

It is clear that the map $\mathfrak{a} \mapsto \mathfrak{a}[1/c]$ preserves multiplication of ideals. Since it is bijective, its inverse also preserves multiplication of ideals. Since localization is an exact functor, and c is invertible in the quotient $\mathcal{O}/\mathfrak{a}$, we have $\mathcal{O}/\mathfrak{a} \cong \mathcal{O}/\mathfrak{a}[1/c] \cong \mathcal{O}[1/c]/\mathfrak{b}$. □

Corollary 13.11. *Let $D \leq -1$, $D \equiv 0, 1$ mod 4 be a discriminant, and let $K = \mathbb{Q}(\sqrt{D})$. Let $c = [\mathcal{O}_K : \mathcal{O}_D]$. Then there is a multiplication-preserving bijection between the following two sets of ideals:*

1. *The set of non-zero ideals of $\mathcal{O}_K$, prime to c.*

2. *The set of non-zero ideals of $\mathcal{O}_D$, prime to c.*

The map is given by $\mathfrak{a} \mapsto \mathfrak{a} \cap \mathcal{O}_D$, with inverse $\mathfrak{a} \mapsto \mathfrak{a}[1/c] \cap \mathcal{O}_K$.

Proof. We just need to observe that $\mathcal{O}_D[1/c] = \mathcal{O}_K[1/c]$. □

This will allow us to relate the group $C(D)$ to a ray class group of K, using the following lemma:

Lemma 13.12. *Let $D \leq -1$, $D \equiv 0, 1$ mod 4 be a discriminant, and let $K = \mathbb{Q}(\sqrt{D})$. Let $c = [\mathcal{O}_K : \mathcal{O}_D]$. Then:*

1. *If $\mathfrak{a} \subset \mathcal{O}_D$ is an ideal prime to c, then $\mathfrak{a}$ is a lattice and $\mathcal{O}_\mathfrak{a} = \mathcal{O}_D$, and $\mathbf{N}\mathfrak{a} = [\mathcal{O}_D : \mathfrak{a}]$ is prime to c.*

2. *Every lattice M such that $\mathcal{O}_M = \mathcal{O}_D$ is equivalent to an ideal $\mathfrak{a} \subset \mathcal{O}_D$ prime to c.*

Proof. For the first part, we note that if $\mathfrak{a} + c\mathcal{O}_D = \mathcal{O}_D$, and $\beta \in K$ satisfies $\beta\mathfrak{a} \subset \mathfrak{a}$, then β is integral over $\mathbb{Z}$, so satisfies $\beta \in \mathcal{O}_K$. We get $\beta\mathcal{O}_D = \beta\mathfrak{a} + c\beta\mathcal{O}_D \subset \mathfrak{a} + c\mathcal{O}_K \subset \mathcal{O}_D$. Since $\mathcal{O}_D$ is its own order, this shows that $\beta \in \mathcal{O}_D$.

For the second part, we recall that such an ideal $\mathfrak{a}$ exists if and only if the binary quadratic form $f_{M^{-1}}(x, y)$ represents an integer prime to c. It therefore suffices to show that any primitive binary quadratic form represents integer prime to c. By the Chinese remainder theorem, it suffices to show that for any prime p, a primitive binary quadratic form $f(x,y)$ represents integers not divisible by p. But one of $f(1,0)$, $f(1,1)$ and $f(0,1)$ will be prime to p, because $f(x,y)$ is primitive. □

Corollary 13.13. *There is a surjective homomorphism $H(c\mathcal{O}_K) \to C(D)$.*

Proof. The map will be induced by the map which sends an ideal $\mathfrak{a} \subset \mathcal{O}_K$ which is prime to c to the ideal $\mathfrak{a}[1/c] \cap \mathcal{O}_D = \mathfrak{a} \cap \mathcal{O}_D$. This is a bijection on ideals prime to c. We know that this map preserves multiplication of ideals, so we get a homomorphism $I(c\mathcal{O}_K) \to C(D)$, which is surjective by the previous lemma. To show that this descends to the quotient, we must show that every ideal $\mathfrak{a} \subset \mathcal{O}_K$ of the form $\alpha\mathcal{O}_K$, where $\alpha \in \mathcal{O}_K$ satisfies $\alpha \equiv 1 \bmod c\mathcal{O}_K$, is send to the trivial class in $C(D)$.

However, we have $\alpha\mathcal{O}_K \cap \mathcal{O}_D = \alpha\mathcal{O}_D$, which is indeed in the trivial class. This completes the proof. □

By extending the arguments in the proof of the corollary, it is possible to calculate explicitly the kernel of the map $H(c\mathcal{O}_K) \to C(D)$ and to give a formula for the order of the finite group $C(D)$, as we have already done for the group $H(c\mathcal{O}_K)$.

Theorem 13.14. *Let $D \leq -1$, $D \equiv 0, 1$ mod 4 be a discriminant, and let $K = \mathbb{Q}(\sqrt{D})$. Let $c = [\mathcal{O}_K : \mathcal{O}_D]$. Then there exists an abelian extension K_D/K, called the ring class field of K of discriminant D, which satisfies the following properties:*

1. *K_D is contained inside the ray class field of level $c\mathcal{O}_K$, and there is an isomorphism $\phi_{K_D/K} : C(D) \cong \mathrm{Gal}(K_D/K)$, uniquely characterized as follows: for every prime ideal $\mathfrak{p} \subset \mathcal{O}_K$ not dividing $c\mathcal{O}_K$, the isomorphism sends the class of the lattice $\mathfrak{p} \cap \mathcal{O}_D$ to $(\mathfrak{p}, K_D/K)$.*

2. *Let p be a prime not dividing D. Then p splits in K_D if and only if p is represented by the principal form of discriminant D.*

Proof. This is now a matter of assembling the ingredients. By the second main theorem of class field theory, the quotient $H(c\mathcal{O}_K) \to C(D)$ corresponds to an abelian extension K_D/K, contained inside the ray class field of level $c\mathcal{O}_K$, for which the Artin map gives a surjection $\phi_{K_D/K} : H(c\mathcal{O}_K) \to \mathrm{Gal}(K_D/K)$ which factors through an isomorphism $C(D) \cong \mathrm{Gal}(K_D/K)$. The group $H(c\mathcal{O}_K)$ is generated by the classes of non-zero prime ideals $\mathfrak{p} \subset \mathcal{O}_K$ prime to c, so the group $C(D)$ is generated by the classes $\mathfrak{p} \cap \mathcal{O}_D$, and $\phi_{L/K}(\mathfrak{p}) = (\mathfrak{p}, K_D/K)$.

On the other hand, we know that we can identify the group $C(D)$ with the set of primitive, positive definite binary quadratic forms of discriminant D; and that a given such form $f(x,y)$ represents a prime p if and only if the class $[f]^{-1}$ contains a lattice $\mathfrak{a} \subset \mathcal{O}_D$ such that $\mathbf{N}\mathfrak{a} = p$, if and only if the class $[f]$ contains a lattice $\mathfrak{a} \subset \mathcal{O}_D$ such that $\mathbf{N}\mathfrak{a} = p$ (because complex conjugation preserves norms and acts by inversion on the group $C(D)$).

If furthermore p is prime to D, then the ideal $\mathfrak{a}$ must be prime to the conductor, and this is equivalent to asking for an ideal $\mathfrak{a} \subset \mathcal{O}_K$ such that $\mathbf{N}\mathfrak{a} = p$. Norm of ideals is multiplicative (by the Chinese remainder theorem), so this happens if and only if $\mathfrak{a}$ is prime, and in this case p must be split in K. We find that the form $f(x,y)$ represents the prime p if and only if there exists a non-zero prime ideal $\mathfrak{p} \subset \mathcal{O}_K$ such that $\mathfrak{p}$ divides p and $[\mathfrak{p} \cap \mathcal{O}_D] = [f]$ in $\mathcal{O}_D$.

Now suppose that $f(x,y)$ is the principal form, corresponding to the trivial element $[\mathcal{O}_D]$ in $C(D)$. Then we see that for a prime p not dividing D, p is represented by $f(x,y)$ if and only if p splits in $\mathcal{O}_K$ and there is a prime $\mathfrak{p} \subset \mathcal{O}_K$ lying above p such that $[\mathfrak{p} \cap \mathcal{O}_D] = [\mathcal{O}_D]$, if and only if $\phi_{K_D/K}([\mathfrak{p} \cap \mathcal{O}_D]) = (\mathfrak{p}, K_D/K) = 1$, if and only if $\mathfrak{p}$ splits in K_D/K, if and only if p splits in $K/\mathbb{Q}$. □

Example 13.15. Let $D = -4 \times 27$, $\mathcal{O}_D = \mathbb{Z}[\sqrt{-27}]$. Then $K = \mathbb{Q}(\sqrt{D}) = \mathbb{Q}(\sqrt{-3})$, and one can show that $K_D = K(\sqrt[3]{2})$. We find that for a prime $p > 3$, p is represented by the principal form $f(x,y) = x^2 + 27y^2$ of discriminant D if and only if p splits in K_D, if and only if $p \equiv 1 \bmod 3$ and the equation $x^3 = 2$ has a solution in $\mathbb{F}_p$.

We can now analyze concrete examples by doing explicit calculations. To this end, it is helpful to recall the following result in reduction theory from Part II Number Theory:

Theorem 13.16. *Each primitive form of discriminant D is properly equivalent to a unique primitive form $ax^2 + bxy + cy^2$ which is reduced, i.e. satisfies $|b| \leq a \leq c$, and $b \geq 0$ if either $|b| = a$ or $a = c$.*

Since a reduced form has $-D = 4ac - b^2 \geq 4a^2 - a^2 = 3a^2$, hence $a \leq \sqrt{-D/3}$, we can always calculate $C(D)$, at least as a set, by enumerating all reduced forms of discriminant D. It is also possible to describe the group law on $C(D)$ explicitly at the level of binary quadratic forms, without passage to ideal classes. This was done by Gauss, and gives rise to the famous composition law of binary quadratic forms.

A useful observation is that if $f(x,y) = ax^2 + bxy + cy^2$ is a primitive, positive definite binary quadratic form of discriminant D, then the inverse class of $[f]$ is represented by $ax^2 - bxy + cy^2$. Using this one can show that, if $f(x,y)$ is reduced, then it corresponds to a class of order dividing 2 in $C(D)$ if and only if either $b = 0$, $a = b$, or $a = c$.

Let's use this to calculate the ideal class group of the field $K = \mathbb{Q}(\sqrt{-5})$; this was used in an earlier calculation. We have $\operatorname{disc} \mathcal{O}_K = -20$, so any reduced form has $|a| \leq 2$. Enumerating all possibilities, we get $x^2 + 5y^2$ and $2x^2 + 2xy + 3y^2$. The class group has order 2, and the Hilbert class field is $K(\sqrt{5})$. We find that for a prime p not dividing 20, the prime p is represented by $2x^2 + 2xy + 3y^2$ if and only if p splits in K but does not split in $\mathbb{Q}(\sqrt{5})$, if and only if $p \equiv 3, 7 \bmod 20$.

To end the course, we prove a result about the number of primes which are represented by a given positive definite, primitive binary quadratic form.

Definition 13.17. *Let K be a number field, and let S be a set of prime ideals of $\mathcal{O}_K$, $\delta \in [0,1]$. We say that the set S has density δ if the limit*

$$\lim_{X \to \infty} \frac{\{\mathfrak{p} \in S \mid \mathbf{N}\mathfrak{p} \leq X\}}{\{\mathfrak{p} \subset \mathcal{O}_K \textit{ prime } \mid \mathbf{N}\mathfrak{p} \leq X\}}$$

exists and equals δ. We write $\mathbf{N}\mathfrak{p}$ for the index $[\mathcal{O}_K : \mathfrak{p}]$ of additive groups; it is easy to see that the numerator and denominator are finite for any $X < \infty$.

Theorem 13.18. *Let L/K be a Galois extension of number fields, let $G = \mathrm{Gal}(L/K)$, and let $C \subset G$ be a conjugacy class. Let S denote the set of prime ideals of $\mathcal{O}_K$ unramified in $\mathcal{O}_L$ such that $\mathrm{Frob}_{\mathfrak{q}/\mathfrak{p}} \in C$ for some (equivalently, every) prime ideal $\mathfrak{q} \subset \mathcal{O}_L$ lying above $\mathfrak{p}$.*

Then the set S has density equal to $\#C/\#G$.

This is the Chebotarev density theorem. Its proof uses class field theory and L-functions. Note that as a particular consequence, we see that if L/K is a Galois extension of number fields, then the set of prime ideals $\mathfrak{p} \subset \mathcal{O}_K$ which split in L has density equal to $1/[L : K]$.

Theorem 13.19. *Let $f(x, y)$ be a primitive, positive definite binary quadratic form of discriminant D, and let S denote the set of primes not dividing D which are represented by $f(x, y)$. Then the set S has density equal to either $1/2\#C(D)$ or $1/\#C(D)$, depending on whether the class of $f(x, y)$ in the group $C(D)$ divides 2 or not.*

Proof. Let K_D/K be the ring class field of discriminant D. Then $K_D/\mathbb{Q}$ is Galois, and its Galois group $\mathrm{Gal}(K_D/K) \cong C(D) \rtimes \mathrm{Gal}(K/\mathbb{Q})$ is a semi-direct product, with $\mathrm{Gal}(K/\mathbb{Q})$ acting on $C(D)$ by inversion (which in this case, agrees with complex conjugation).

Let $\sigma_f = \phi_{K_D/K}([f]) \in \mathrm{Gal}(K_D/K)$. We know that for a given prime p, not dividing D, the form $f(x, y)$ represents p if and only if $p = \mathfrak{p}\overline{\mathfrak{p}}$ splits in K and $\sigma_f \in \{(\mathfrak{p}, K_D/K), (\overline{\mathfrak{p}}, K_D/K)\}$. We can calculate the density of the set of such primes using the Chebotarev density theorem.

First suppose that σ_f has order dividing 2. Then the conjugacy class of σ_f is just $\{\sigma_f\}$, so we see that the set has density $1/\#\,\mathrm{Gal}(K_D/\mathbb{Q}) = 1/2\#C(D)$.

Now suppose that σ_f has order not dividing 2, so that $\sigma_f \neq \sigma_f^{-1}$. Then the conjugacy class of σ_f is $\{\sigma_f, \sigma_f^{-1}\}$, and p is represented by $f(x, y)$ if and only if $\mathrm{Frob}_{\mathfrak{q}/p} \in \{\sigma_f, \sigma_f^{-1}\}$ for some (equivalently every) prime $\mathfrak{p} \subset \mathcal{O}_{K_D}$ lying above p. Calculating again using the Chebotarev density theorem, we find that the set of such primes has density $2/\#\,\mathrm{Gal}(K_D/\mathbb{Q}) = 1/\#C(D)$. □

Corollary 13.20. *Let $n \geq 1$ be an integer. Then there are infinitely many primes of the form $p = x^2 + ny^2$.*

Proof. Apply the theorem to the principal form of discriminant $D = \operatorname{disc} \mathbb{Z}[\sqrt{-n}]$. □