11 Completing extensions, different and discriminant ideals

11.1 Local extensions come from global extensions

Let $\hat{L}$ be a local field. From our classification of local fields (Theorem 9.7), we know $\hat{L}$ is a finite extension of $\hat{K} = \mathbb{Q}_p$ (some prime $p \le \infty$) or $\hat{K} = \mathbb{F}_q((t))$ (some prime power q). We also know that the completion of a global field at any of its nontrivial absolute values is such a local field (Corollary 9.5). It thus reasonable to ask whether $\hat{L}$ is the completion of a corresponding global field L that is a finite extension of $K = \mathbb{Q}$ or $K = \mathbb{F}_q(t)$.

More generally, for any fixed global field K and local field $\hat{K}$ that is the completion of K with respect to one of its nontrivial absolute values $|\ |$, we may ask whether every finite extension of local fields $\hat{L}/\hat{K}$ necessarily corresponds to an extension of global fields L/K, where $\hat{L}$ is the completion of L with respect to one of its absolute values (whose restriction to K must be equivalent to $|\ |$). The answer is yes. In order to simplify matters we restrict our attention to the case where $\hat{L}/\hat{K}$ is separable, but this is true in general.

Theorem 11.1. *Let K be a global field with a nontrivial absolute value $|\ |$, and let $\hat{K}$ be the completion of K with respect to $|\ |$. Every finite separable extension $\hat{L}$ of $\hat{K}$ is the completion of a finite separable extension L of K with respect to an absolute value that restricts to $|\ |$. Moreover, one can choose L so that $\hat{L}$ is the compositum of L and $\hat{K}$ and $[\hat{L} : \hat{K}] = [L : K]$.*

Proof. Let $\hat{L}/\hat{K}$ be a separable extension of degree n. Let us first suppose that $|\ |$ is archimedean. Then K is a number field and $\hat{K}$ is either $\mathbb{R}$ or $\mathbb{C}$; the only nontrivial case is when $\hat{K} = \mathbb{R}$ and $n = 2$, and we may then assume that $\hat{L} \simeq \mathbb{C}$ is $\hat{K}(\sqrt{-d})$ where $-d \in \mathbb{Z}_{<0}$ is a nonsquare in K (such a $-d$ exists because $K/\mathbb{Q}$ is finite). We may assume without loss of generality that $|\ |$ is the Euclidean absolute value on $\hat{K} \simeq \mathbb{R}$ (it must be equivalent to it), and uniquely extend $|\ |$ to $L = K(\sqrt{-d})$ by requiring $|\sqrt{-d}| = \sqrt{d}$. Then $\hat{L}$ is the completion of L with respect to $|\ |$, and clearly $[\hat{L} : \hat{K}] = [L : K] = 2$, and $\hat{L}$ is the compositum of L and $\hat{K}$.

We now suppose that $|\ |$ is nonarchimedean, in which case the valuation ring of $\hat{K}$ is a complete DVR and $|\ |$ is induced by the corresponding discrete valuation. By the primitive element theorem (Theorem 4.33), we may assume $\hat{L} = \hat{K}[x]/(f)$ where $f \in \hat{K}[x]$ is monic, irreducible, and separable. The field K is dense in its completion $\hat{K}$, so we can find a monic $g \in K[x] \subseteq \hat{K}[x]$ that is arbitrarily close to f: such that $\|g - f\|_1 < \delta$ for any $\delta > 0$. It then follows from Proposition 10.30 that $\hat{L} = \hat{K}[x]/(g)$ (and that g is separable). The field $\hat{L}$ is a finite separable extension of the fraction field of a complete DVR, so by Theorem 9.25 it is itself the fraction field of a complete DVR and has a unique absolute value that extends the absolute value $|\ |$ on $\hat{K}$.

Now let $L = K[x]/(g)$. The polynomial g is irreducible in $\hat{K}[x]$, hence in $K[x]$, so $[L : K] = \deg g = [\hat{L} : \hat{K}]$. The field $\hat{L}$ contains both $\hat{K}$ and L, and it is clearly the smallest field that does (since g is irreducible in $\hat{K}[x]$), so $\hat{L}$ is the compositum of $\hat{K}$ and L. The absolute value on $\hat{L}$ restricts to an absolute value on L extending the absolute value $|\ |$ on K, and $\hat{L}$ is complete, so $\hat{L}$ contains the completion of L with respect to $|\ |$. On the other hand, the completion of L with respect $|\ |$ contains both L and $\hat{K}$, so it must be $\hat{L}$. $\square$

In the preceding theorem, when the local extension $\hat{L}/\hat{K}$ is Galois one might ask whether the corresponding global extension L/K is also Galois, and whether $\mathrm{Gal}(\hat{L}/\hat{K}) \simeq \mathrm{Gal}(L/K)$. As shown by the following example, this need not be the case.

Example 11.2. Let $K = \mathbb{Q}$, $\hat{K} = \mathbb{Q}_7$ and $\hat{L} = \hat{K}[x]/(x^3 - 2)$. The extension $\hat{L}/\hat{K}$ is Galois because $\hat{K} = \mathbb{Q}_7$ contains ζ_3 (we can lift the root 2 of $x^2 + x + 1 \in \mathbb{F}_7[x]$ to a root of $x^2 + x + 1 \in \mathbb{Q}_7[x]$ via Hensel's lemma), and this implies that $x^3 - 2$ splits completely in $L_w = \mathbb{Q}_7(\sqrt[3]{2})$. But $L = K[x]/(x^3 - 2)$ is not a Galois extension of K because it contains only one root of $x^3 - 2$. However, we can replace K with $\mathbb{Q}(\zeta_3)$ without changing $\hat{K}$ (take the completion of K with respect to the absolute value induced by a prime above 7) or $\hat{L}$, but now $L = K[x]/(x^3 - 2)$ is a Galois extension of K.

In the example we were able to adjust our choice of the global field K without changing the local fields extension $\hat{L}/\hat{K}$ in a way that ensures that $\hat{L}/\hat{K}$ and L/K have the same automorphism group. Indeed, this is always possible.

Corollary 11.3. *For every finite Galois extension $\hat{L}/\hat{K}$ of local fields there is a corresponding Galois extension of global fields L/K and an absolute value $|\ |$ on L such that $\hat{L}$ is the completion of L with respect to $|\ |$, $\hat{K}$ is the completion of K with respect to the restriction of $|\ |$ to K, and $\mathrm{Gal}(\hat{L}/\hat{K}) \simeq \mathrm{Gal}(L/K)$.*

Proof. The archimedean case is already covered by Theorem 11.1 (take $K = \mathbb{Q}$), so we assume $\hat{L}$ is nonarchimedean and note that we may take $|\ |$ to be the absolute value on both $\hat{K}$ and on $\hat{L}$ (by Theorem 9.25). The field $\hat{K}$ is an extension of either $\mathbb{Q}_p$ or $\mathbb{F}_q((t))$, and by applying Theorem 11.1 to this extension we may assume $\hat{K}$ is the completion of a global field K with respect to the restriction of $|\ |$. As in the proof of the theorem, let $g \in K[x]$ be a monic separable polynomial irreducible in $\hat{K}[x]$ such that $\hat{L} = \hat{K}[x]/(g)$ and define $L := K[x]/(g)$ so that $\hat{L}$ is the compositum of $\hat{K}$ and L.

Now let M be the splitting field of g over K, the minimal extension of K that contains all the roots of g (which are distinct because g is separable). The field $\hat{L}$ also contains these roots (since $\hat{L}/\hat{K}$ is Galois) and $\hat{L}$ contains K, so $\hat{L}$ contains a subextension of K isomorphic to M (by the universal property of a splitting field), which we now identify with M; note that $\hat{L}$ is also the completion of M with respect to the restriction of $|\ |$ to M.

We have a group homomorphism $\varphi\colon \mathrm{Gal}(\hat{L}/\hat{K}) \to \mathrm{Gal}(M/K)$ induced by restriction, and φ is injective (each $\sigma \in \mathrm{Gal}(\hat{L}/\hat{K})$ is determined by its action on any root of g in M). If we now replace K by the fixed field of the image of φ and replace L with M, the completion of K with respect to the restriction of $|\ |$ is still equal to $\hat{K}$, and similarly for L and $\hat{L}$, and now $\mathrm{Gal}(L/K) = \mathrm{Gal}(\hat{L}/\hat{K})$ as desired. □

11.2 Completing a separable extension of Dedekind domains

We now return to our general $AKLB$ setup: A is a Dedekind domain with fraction field K with a finite separable extension L/K, and B is the integral closure of A in L, which is also a Dedekind domain. Recall from Theorem 5.11 that if $\mathfrak{p}$ is a nonzero prime of A, each prime $\mathfrak{q}|\mathfrak{p}$ gives a valuation $v_\mathfrak{q}$ of L that extends the valuation $v_\mathfrak{p}$ of K with index $e_\mathfrak{q}$, meaning that $v_\mathfrak{q}|_K = e_\mathfrak{q} v_\mathfrak{p}$. Moreover, every valuation of L that extends $v_\mathfrak{p}$ arises in this way. We now want to look at what happens when we complete K with respect to the absolute value $|\ |_\mathfrak{p}$ induced by $v_\mathfrak{p}$, and similarly complete L with respect to $|\ |_\mathfrak{q}$ for some $\mathfrak{q}|\mathfrak{p}$. This includes the case where L/K is an extension of global fields, in which case we get a corresponding extension $L_\mathfrak{q}/K_\mathfrak{p}$ of local fields for each $\mathfrak{q}|\mathfrak{p}$, but note that $L_\mathfrak{q}/K_\mathfrak{p}$ may have strictly smaller degree than L/K because if we write $L \simeq K[x]/(f)$, the irreducible polynomial $f \in K[x]$ need not be irreducible over $K_\mathfrak{p}$. Indeed, this will necessarily be the case if there is more than one prime $\mathfrak{q}$ lying above $\mathfrak{p}$; there is a one-to-one correspondence between factors of f

in $K_{\mathfrak{p}}[x]$ and primes $\mathfrak{q}|\mathfrak{p}$. If L/K is Galois, so is $L_{\mathfrak{q}}/K_{\mathfrak{p}}$ and each $\mathrm{Gal}(L_{\mathfrak{q}}/K_{\mathfrak{p}})$ is isomorphic to the decomposition group $D_{\mathfrak{q}}$ (which perhaps helps to explain the terminology).

The following theorem gives a complete description of the situation.

Theorem 11.4. *Assume AKLB, let $\mathfrak{p}$ be a prime of A, and let $\mathfrak{p}B = \prod_{\mathfrak{q}|\mathfrak{p}} \mathfrak{q}^{e_{\mathfrak{q}}}$ be the factorization of $\mathfrak{p}B$ in B. Let $K_{\mathfrak{p}}$ denote the completion of K with respect to $|\ |_{\mathfrak{p}}$, and let $\hat{\mathfrak{p}}$ denote the maximal ideal of its valuation ring. For each $\mathfrak{q}|\mathfrak{p}$, let $L_{\mathfrak{q}}$ denote the completion of L with respect to $|\ |_{\mathfrak{q}}$, and let $\hat{\mathfrak{q}}$ denote the maximal ideal of its valuation ring. The following hold:*

1. *Each $L_{\mathfrak{q}}$ is a finite separable extension of $K_{\mathfrak{p}}$;*
2. *Each $\hat{\mathfrak{q}}$ is the unique prime of $L_{\mathfrak{q}}$ lying over $\hat{\mathfrak{p}}$.*
3. *Each $\hat{\mathfrak{q}}$ has ramification index $e_{\hat{\mathfrak{q}}} = e_{\mathfrak{q}}$ and residue field degree $f_{\hat{q}} = f_{\mathfrak{q}}$.*
4. *$[L_{\mathfrak{q}} : K_{\mathfrak{p}}] = e_{\mathfrak{q}} f_{\mathfrak{q}}$;*
5. *The map $L \otimes_K K_{\mathfrak{p}} \to \prod_{\mathfrak{q}|\mathfrak{p}} L_{\mathfrak{q}}$ defined by $\ell \otimes x \mapsto (\ell x, \ldots, \ell x)$ is an isomorphism of finite étale $K_{\mathfrak{p}}$-algebras.*
6. *If L/K is Galois then each $L_{\mathfrak{q}}/K_{\mathfrak{p}}$ is Galois and we have isomorphisms of decomposition groups $D_{\mathfrak{q}} \simeq D_{\hat{\mathfrak{q}}} = \mathrm{Gal}(L_{\mathfrak{q}}/K_{\mathfrak{p}})$ and inertia groups $I_{\mathfrak{q}} \simeq I_{\hat{\mathfrak{q}}}$.*

Proof. We first note that the $K_{\mathfrak{p}}$ and the $L_{\mathfrak{q}}$ are all fraction fields of complete DVRs; this follows from Proposition 8.18 (note: we are not assuming they are local fields, in particular, their residue fields need not be finite).

(1) For each $\mathfrak{q}|\mathfrak{p}$ the embedding $K \hookrightarrow L$ induces an embedding $K_{\mathfrak{p}} \hookrightarrow L_{\mathfrak{q}}$ via the map $[(a_n)] \mapsto [(a_n)]$ on equivalence classes of Cauchy sequences; a sequence (a_n) that is Cauchy in K with respect to $|\ |_{\mathfrak{p}}$, is also Cauchy in L with respect to $|\ |_{\mathfrak{q}}$ because $v_{\mathfrak{q}}$ extends $v_{\mathfrak{p}}$. We thus view $K_{\mathfrak{p}}$ as a subfield of $L_{\mathfrak{q}}$, which also contains L. There is thus a K-algebra homomorphism $\phi_{\mathfrak{q}} \colon L \otimes_K K_{\mathfrak{p}} \to L_{\mathfrak{q}}$ defined by $\ell \otimes x \mapsto \ell x$, which we may view as a linear map of $K_{\mathfrak{p}}$ vector spaces. We claim that $\phi_{\mathfrak{q}}$ is surjective.

If $\alpha_1, \ldots, \alpha_m$ is any basis for $L_{\mathfrak{q}}$ then its determinant with respect to $\mathcal{B}$, i.e., the $m \times m$ matrix whose jth row contains the coefficients of α_j when written as a linear combination of elements of $\mathcal{B}$, must be nonzero. The determinant is a polynomial in the entries of this matrix, hence a continuous function with respect to the topology on $L_{\mathfrak{q}}$ induced by the absolute value $|\cdot|_{\mathfrak{q}}$. It follows that if we replace $\alpha_1, \ldots, \alpha_m$ with $\ell_1, \ldots, \ell_m$ chosen so that $|\alpha_j - \ell_j|_{\mathfrak{q}}$ is sufficiently small, the matrix of $\ell_1, \ldots, \ell_m$ with respect to $\mathcal{B}$ must also be nonzero, and therefore $\ell_1, \ldots, \ell_m$ is also a basis for $L_{\mathfrak{q}}$. We can thus choose a basis $\ell_1, \ldots, \ell_m \in L$, since L is dense in its completion $L_{\mathfrak{q}}$. But then $\{\ell_j\} = \{\phi_{\mathfrak{q}}(\ell_j \otimes 1)\} \subseteq \mathrm{im}\, \phi_{\mathfrak{q}}$ spans $L_{\mathfrak{q}}$, so $\phi_{\mathfrak{q}}$ is surjective as claimed.

The $K_{\mathfrak{p}}$-algebra $L \otimes_K K_{\mathfrak{p}}$ is the base change of a finite étale algebra, hence finite étale, by Proposition 4.34. It follows that $L_{\mathfrak{q}}$ is a finite separable extension of $K_{\mathfrak{p}}$: it certainly has finite dimension as a $K_{\mathfrak{p}}$-vector space, since $\phi_{\mathfrak{q}}$ is surjective, and it is separable because every $\alpha \in L_{\mathfrak{q}}$ is the image $\phi_{\mathfrak{q}}(\beta)$ of an element $\beta \in L \otimes_K K_{\mathfrak{p}}$ that is a root of a separable (but not necessarily irreducible) polynomial $f \in K_{\mathfrak{p}}[x]$, as explained after Definition 4.29; we then have $0 = \phi_{\mathfrak{q}}(0) = \phi_{\mathfrak{q}}(f(\beta)) = f(\alpha)$, so α is a root of f, hence separable.

(2) The valuation rings of $K_{\mathfrak{p}}$ and $L_{\mathfrak{q}}$ are complete DVRs, so this follows immediately from Theorem 9.20.

(3) The valuation $v_{\hat{\mathfrak{q}}}$ extends $v_{\mathfrak{q}}$ with index 1, which in turn extends $v_{\mathfrak{p}}$ with index $e_{\mathfrak{q}}$. The valuation $v_{\hat{\mathfrak{p}}}$ extends $v_{\mathfrak{p}}$ with index 1, and it follows that $v_{\hat{q}}$ extends $v_{\hat{p}}$ with index $e_{\mathfrak{q}}$ and therefore $e_{\hat{\mathfrak{q}}} = e_{\mathfrak{q}}$. The residue field of $\hat{\mathfrak{p}}$ is the same as that of $\mathfrak{p}$: for any Cauchy

sequence (a_n) over K the a_n will eventually all have the same image in the residue field at $\mathfrak{p}$ (since $v_\mathfrak{p}(a_n - a_m) > 0$ for all sufficiently large m and n). Similar comments apply to each $\hat{\mathfrak{q}}$ and $\mathfrak{q}$, and it follows that $f_{\hat{\mathfrak{q}}} = f_\mathfrak{q}$.

(4) It follows from (2) that $[L_\mathfrak{q} : K_\mathfrak{p}] = e_{\hat{\mathfrak{q}}} f_{\hat{\mathfrak{q}}}$, since $\hat{\mathfrak{q}}$ is the only prime above $\hat{\mathfrak{p}}$, and (3) then implies $[L_\mathfrak{q} : K_\mathfrak{p}] = e_\mathfrak{q} f_\mathfrak{q}$.

(5) Let $\phi = \prod_{\mathfrak{q}|\mathfrak{p}} \phi_\mathfrak{q}$, where $\phi_\mathfrak{q}$ are the surjective $K_\mathfrak{p}$-algebra homomorphisms defined in the proof of (1). Then $\phi\colon L \otimes_K K_\mathfrak{p} \to \prod_{\mathfrak{q}|\mathfrak{p}} L_\mathfrak{q}$ is a $K_\mathfrak{p}$-algebra homomorphism. Applying (4) and the fact that base change preserves dimension (see Proposition 4.34):

$$\dim_{K_\mathfrak{p}} (L \otimes_K K_\mathfrak{p}) = \dim_K L = [L : K] = \sum_{\mathfrak{q}|\mathfrak{p}} e_\mathfrak{q} f_\mathfrak{q} = \sum_{\mathfrak{q}|\mathfrak{p}} [L_\mathfrak{q} : K_\mathfrak{p}] = \dim_{K_\mathfrak{p}} \left(\prod_{\mathfrak{q}|\mathfrak{p}} L_\mathfrak{q} \right).$$

The domain and range of ϕ thus have the same dimension, and ϕ is surjective (since the $\phi_\mathfrak{q}$ are), so it is an isomorphism.

(6) We now assume L/K is Galois. Each $\sigma \in D_\mathfrak{q}$ acts on L and respects the valuation $v_\mathfrak{q}$, since it fixes $\mathfrak{q}$ (if $x \in \mathfrak{q}^n$ then $\sigma(x) \in \sigma(\mathfrak{q}^n) = \sigma(\mathfrak{q})^n = \mathfrak{q}^n$). It follows that if (x_n) is a Cauchy sequence in L, then so is $(\sigma(x_n))$, thus σ is an automorphism of $L_\mathfrak{q}$, and it fixes $K_\mathfrak{p}$. We thus have a group homomorphism $\varphi\colon D_\mathfrak{q} \to \operatorname{Aut}_{K_\mathfrak{p}}(L_\mathfrak{q})$.

If $\sigma \in D_\mathfrak{q}$ acts trivially on $L_\mathfrak{q}$ then it acts trivially on $L \subseteq L_\mathfrak{q}$, so $\ker \varphi$ is trivial. Also,

$$e_\mathfrak{q} f_\mathfrak{q} = |D_\mathfrak{q}| \le \#\operatorname{Aut}_{K_\mathfrak{p}}(L_\mathfrak{q}) \le [L_\mathfrak{q} : K_\mathfrak{p}] = e_\mathfrak{q} f_\mathfrak{q},$$

by Theorem 11.4, so $\#\operatorname{Aut}_{K_\mathfrak{p}}(L_\mathfrak{q}) = [L_\mathfrak{q} : K_\mathfrak{p}]$ and $L_\mathfrak{q}/K_\mathfrak{p}$ is Galois, and this also shows that φ is surjective and therefore an isomorphism. There is only one prime $\hat{q}$ of $L_\mathfrak{q}$, and it is necessarily fixed by every $\sigma \in \operatorname{Gal}(L_\mathfrak{q}/K_\mathfrak{p})$, so $\operatorname{Gal}(L_\mathfrak{q}/K_\mathfrak{p}) \simeq D_{\hat{\mathfrak{q}}}$. The inertia groups $I_\mathfrak{q}$ and $I_{\hat{\mathfrak{q}}}$ both have order $e_\mathfrak{q} = e_{\hat{\mathfrak{q}}}$, and φ restricts to a homomorphism $I_\mathfrak{q} \to I_{\hat{\mathfrak{q}}}$, so the inertia groups are also isomorphic. □

Corollary 11.5. *Assume AKLB and let $\mathfrak{p}$ be a prime of A. For every $\alpha \in L$ we have*

$$\mathrm{N}_{L/K}(\alpha) = \prod_{\mathfrak{q}|\mathfrak{p}} \mathrm{N}_{L_\mathfrak{q}/K_\mathfrak{p}}(\alpha) \qquad \textit{and} \qquad \mathrm{T}_{L/K}(\alpha) = \prod_{\mathfrak{q}|\mathfrak{p}} \mathrm{T}_{L_\mathfrak{q}/K_\mathfrak{q}}(\alpha).$$

where we view α as an element of $L_\mathfrak{q}$ via the canonical embedding $L \hookrightarrow L_\mathfrak{q}$.

Proof. The norm and trace are defined as the determinant and trace of K-linear maps $L \xrightarrow{\times\alpha} L$ that are unchanged upon tensoring with $K_\mathfrak{p}$; the corollary then follows from the isomorphism in part (5) of Theorem 11.4, which commutes with the norm and trace. □

Remark 11.6. Theorem 11.4 can be stated more generally in terms of (equivalence classes of) absolute values (or *places*). Rather than working with a prime $\mathfrak{p}$ of K and primes $\mathfrak{q}$ of L above $\mathfrak{p}$, one works with an absolute value $|\ |_v$ of K (for example, $|\ |_\mathfrak{p}$) and inequivalent absolute values $|\ |_w$ of L that extend $|\ |_v$. Places will be discussed further in the next lecture.

Corollary 11.7. *Assume AKLB with A a DVR with maximal ideal $\mathfrak{p}$. Let $\mathfrak{p}B = \prod \mathfrak{q}^{e_\mathfrak{q}}$ be the factorization of $\mathfrak{p}B$ in B. Let $\hat{A}$ denote the completion of A, and for each $\mathfrak{q}|\mathfrak{p}$, let $\hat{B}_\mathfrak{q}$ denote the completion of $B_\mathfrak{q}$. Then $B \otimes_A \hat{A} \simeq \prod_{\mathfrak{q}|\mathfrak{p}} \hat{B}_\mathfrak{q}$.*

Proof. Since A is a DVR (and therefore a torsion-free PID), the ring extension B/A is a free A module of rank $n := [L:K]$, and therefore $B \otimes_A \hat{A}$ is a free $\hat{A}$-module of rank n. And $\prod \hat{B}_{\mathfrak{q}}$ is a free $\hat{A}$-module of rank $\sum_{\mathfrak{q}|\mathfrak{p}} e_{\mathfrak{q}} f_{\mathfrak{q}} = n$. These two $\hat{A}$-modules lie in isomorphic $K_{\mathfrak{p}}$-vector spaces, $L \otimes_K K_{\mathfrak{p}} \simeq \prod L_{\mathfrak{q}}$, by part (5) of Theorem 11.4. To show that they are isomorphic it suffices to check that they are isomorphic after reducing modulo $\hat{\mathfrak{p}}$, the maximal ideal of $\hat{A}$.

For the LHS, note that $\hat{A}/\hat{\mathfrak{p}} \simeq A/\mathfrak{p}$, so

$$B \otimes_A \hat{A}/\hat{\mathfrak{p}} \simeq B \otimes_A A/\mathfrak{p} \simeq B/\mathfrak{p}B.$$

On the RHS we have

$$\prod_{\mathfrak{q}|\mathfrak{p}} \hat{B}_{\mathfrak{q}}/\hat{\mathfrak{p}}\hat{B}_{\mathfrak{q}} \simeq \prod_{\mathfrak{q}|\mathfrak{p}} \hat{B}_{\mathfrak{q}}/\mathfrak{p}\hat{B}_{\mathfrak{q}} \simeq \prod_{\mathfrak{q}|\mathfrak{p}} B_{\mathfrak{q}}/\mathfrak{p}B_{\mathfrak{q}} = \prod_{\mathfrak{q}|\mathfrak{p}} B_{\mathfrak{q}}/\mathfrak{q}^{e_{\mathfrak{q}}}B_{\mathfrak{q}}$$

which is isomorphic to $B/\mathfrak{p}B$ on the LHS because $\mathfrak{p}B = \prod_{\mathfrak{q}|\mathfrak{p}} \mathfrak{q}^{e_{\mathfrak{q}}}$. □

11.3 The different ideal

We continue in our usual $AKLB$ setup: A is a Dedekind domain, K is its fraction field, L/K is a finite separable extension, and B is the integral closure of A in L (a Dedekind domain with fraction field L). We would like to understand the primes that ramify in L/K, that is, the primes $\mathfrak{q}$ of B for which $e_{\mathfrak{q}} > 1$, or, at a coarser level, primes $\mathfrak{p}$ of A that have a ramified prime $\mathfrak{q}$ lying above them. Our main tool for doing so is the *different ideal* $\mathcal{D}_{B/A}$, a fractional ideal of B that will give us an exact answer to this question: the primes of B that ramify are exactly those that divide the different ideal, and $v_{\mathfrak{q}}(\mathcal{D}_{B/A})$ will give us information about the ramification index $e_{\mathfrak{q}}$ (its exact value in the tamely ramified case). Of course we could just define $\mathcal{D}_{B/A}$ to have the properties we want, but the key is to define it in a way that makes it independently computable, allowing us to determine the primes $\mathfrak{q}$ that ramify in B, which we typically do not know *a priori*.

Recall from Lecture 4 the trace pairing $L \times L \to K$ defined by $(x,y) \mapsto \mathrm{T}_{L/K}(xy)$. Since L/K is separable, this pairing is nondegenerate, by Proposition 4.58. For any A-module $M \subseteq L$, we defined the *dual A-module*

$$M^* := \{x \in L : \mathrm{T}_{L/K}(xm) \in A \ \forall m \in M\}$$

(see Definition 4.59). Note that if $M \subseteq N$ are two A-modules in L, then it is clear from the definition that $N^* \subseteq M^*$ (taking duals reverses inclusions).

If M is a free A-*lattice* (see Definition 6.1) then it has an A-module basis $e_1, \ldots, e_n$ that is also a K-basis for L. The dual A-module M^* is then also a free A-lattice, and it has the *dual basis* $e_1^*, \ldots, e_n^*$, which is the unique K-basis for L that satisfies

$$\mathrm{T}_{L/K}(e_i^* e_j) = \delta_{ij} := \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise} \end{cases}$$

(see Proposition 4.54) and also an A-module basis for M^*.

Every B-module $M \subseteq L$ (including all fractional ideals of B) is also a (not necessarily free) A-module in L, and in this case the dual A-module M^* is also a B-module: for any $x \in M^*$, $b \in B$, and $m \in M$ we have $\mathrm{T}((bx)m) = \mathrm{T}(x(bm)) \in A$, since $bm \in M$ and $x \in M^*$, so $bx \in M^*$. If M is a finitely generated as a B-module, then it is a fractional ideal of B (by definition), and provided it is nonzero, it is invertible, since B is a Dedekind domain, and therefore an element of the ideal group $\mathcal{I}_B$. We now show that $M^* \in \mathcal{I}_B$.

Lemma 11.8. *Assume AKLB and suppose $M \in \mathcal{I}_B$. Then $M^* \in \mathcal{I}_B$.*

Proof. Since M is a B-module, so is M^* (as noted above), and M^* is clearly nonzero: if $M = \frac{1}{b}I$ with $b \in B$ nonzero and I a B-ideal, then $bm \in B$ and $T_{L/K}(bm) \in A$ for all $m \in M$ so $b \in M^*$. We just need to check that M^* is finitely generated. Here we use the use the standard trick: find a free submodule of M, take its dual to get a free module that contains M^*, and then note that M^* is a submodule of a noetherian module.

Let $e_1, \ldots, e_n$ be a K-basis for L. By clearing denominators, we may assume the e_i lie in B (since $L = \operatorname{Frac} B$). If m is any element of M, then $me_1, \ldots, me_n$ is a K-basis for L that lies in M. Let C be the free A-submodule of M generated by $me_1, \ldots, me_n$; this is a free A-lattice, and it follows that $M^* \subseteq C^*$ is contained in the free A-lattice C^*, which is obviously finitely generated. As a finitely generated module over a noetherian ring, the A-module C^* is a noetherian module, which means that every A-submodule of C^* is finitely generated, including M^*. We have $A \subseteq B$, so if M^* is finitely generated as an A-module, it is certainly finitely generated as a B-module. $\square$

Definition 11.9. Assume $AKLB$. The *inverse different ideal* (or *codifferent*) of B is the dual of B as an A-module:

$$B^* := \{x \in L : \mathrm{T}_{L/K}(xb) \in A \ \forall b \in B\} \in \mathcal{I}_B.$$

The *different ideal* (or *different*) $\mathcal{D}_{B/A}$ is the inverse of B^* as a fractional B-ideal.

To justify the name, we should check that $\mathcal{D}_{B/A}$ is actually an ideal, not just a fractional ideal. The dual module B^* clearly contains 1, since $\mathrm{T}_{L/K}(1 \cdot b) = \mathrm{T}_{L/K}(b) \in A$ for all $b \in B$. It follows that

$$\mathcal{D}_{B/A} = (B^*)^{-1} = (B : B^*) = \{x \in L : xB^* \subseteq B\} \subseteq B,$$

so $\mathcal{D}_{B/A}$ is indeed a B-ideal.

We now show that the different respects localization and completion.

Proposition 11.10. *Assume AKLB, let S be a multiplicative subset of A. Then*

$$S^{-1}\mathcal{D}_{B/A} = \mathcal{D}_{S^{-1}B/S^{-1}A}.$$

Proof. Since taking inverses respects localization, it suffices to show that $S^{-1}B^* = (S^{-1}B)^*$, where $(S^{-1}B)^*$ denotes the dual of $S^{-1}B$ as an $S^{-1}A$-module in L. If $x = s^{-1}y \in S^{-1}B^*$ with $s \in S$ and $y \in B^*$, and $m = t^{-1}b \in S^{-1}B$ with $t \in S$ and $b \in B$ then

$$\mathrm{T}_{L/K}(xm) = (st)^{-1}\mathrm{T}_{L/K}(yb) \in S^{-1}A,$$

since the trace is K-linear and $S \subseteq A \subseteq K$; this shows that $S^{-1}B^* \subseteq (S^{-1}B)^*$, For the reverse inclusion, let $\{b_i\}$ be a finite set of generators for B as an A-module and let $x \in (S^{-1}B)^*$. For each b_i we have $\mathrm{T}_{L/K}(xb_i) \in S^{-1}A$, since $(S^{-1}B)^*$ is an $S^{-1}B$-module and therefore a B-module. So each $\mathrm{T}_{L/K}(xb_i) = s_i^{-1}a_i$ for some $s_i \in S$ and $a_i \in A$. If we now put $s = \prod s_i$ (a finite product), then $\mathrm{T}_{L/K}(sxb_i) \in A$ for all b_i (here we are again using the K-linearity of $T_{L/K}$). So $sx \in B^*$, and therefore $x \in S^{-1}B^*$ as desired. $\square$

Proposition 11.11. *Assume AKLB and let $\mathfrak{q}|\mathfrak{p}$ be a prime of B. Then*

$$\mathcal{D}_{\hat{B}_{\mathfrak{q}}/\hat{A}_{\mathfrak{p}}} = \mathcal{D}_{B/A} \cdot \hat{B}_{\mathfrak{q}}.$$

Proof. We can assume without loss of generality that A is a DVR by localizing at $\mathfrak{p}$. Let $\hat{L} := L \otimes \hat{K}$. By (5) of Theorem 11.4, we have $\hat{L} = \prod_{\mathfrak{q}|\mathfrak{p}} \hat{L}_{\mathfrak{q}}$. This is not a field, in general, but $\mathrm{T}_{\hat{L}/\hat{K}}$ is defined as for any ring extension, and we have $\mathrm{T}_{\hat{L}/\hat{K}}(x) = \sum_{\mathfrak{q}|\mathfrak{p}} \mathrm{T}_{\hat{L}_{\mathfrak{q}}/\hat{K}}(x)$.

Now let $\hat{B} := B \otimes \hat{A}$. By Corollary 11.7, $\hat{B} = \prod_{\mathfrak{q}|\mathfrak{p}} \hat{B}_{\mathfrak{q}}$, and therefore $\hat{B}^* \simeq \prod_{\mathfrak{q}|\mathfrak{p}} \hat{B}^*_{\mathfrak{q}}$ (since the trace is just a sum of traces). It follows that $\hat{B}^* \simeq B^* \otimes_A \hat{A}$. Thus B^* generates the fractional ideal $\hat{B}^*_{\mathfrak{q}} \in \mathcal{I}_{\hat{B}_{\mathfrak{q}}}$. Taking inverses, $\mathcal{D}_{B/A} = (B^*)^{-1}$ generates $(\hat{B}^*_{\mathfrak{q}})^{-1} = \mathcal{D}_{\hat{B}_{\mathfrak{q}}/\hat{A}}$. ☐

11.4 The discriminant

Definition 11.12. Let B/A be a ring extension with B free as an A-module. For any $e_1, \ldots, e_n \in B$ we define the *discriminant*

$$\operatorname{disc}(e_1, \ldots, e_n) = \det[\mathrm{T}_{B/A}(e_i e_j)]_{i,j} \in A,$$

where $\mathrm{T}_{B/A}(b)$ is the trace from B to A (see Definition 4.40).[1]

We have in mind the case where $e_1, \ldots, e_n$ is a basis for L as a K-vector space. In our usual $AKLB$ setup, if $e_1, \ldots, e_n \in B$ then $\operatorname{disc}(e_1, \ldots, e_n) \in A$.

Proposition 11.13. *Let L/K be a finite separable extension of degree n, and let Ω/K be a field extension for which there are distinct $\sigma_1, \ldots, \sigma_n \in \hom_K(L, \Omega)$. For any $e_1, \ldots, e_n \in L$*

$$\operatorname{disc}(e_1, \ldots, e_n) = (\det[\sigma_i(e_j)]_{ij})^2 .$$

Furthermore, for any $x \in L$

$$\operatorname{disc}(1, x, x^2, \ldots, x^{n-1}) = \prod_{i<j} (\sigma_i(x) - \sigma_j(x))^2 .$$

Note that such an Ω exists, since L/K is separable (just take a normal closure).

Proof. For $1 \le i, j \le n$ we have $\mathrm{T}_{L/K}(e_i e_j) = \sum_{k=1}^n \sigma_k(e_i e_j)$, by Theorem 4.44. Therefore

$$\begin{aligned} \operatorname{disc}(e_1, \ldots, e_n) &= \det[\mathrm{T}_{L/K}(e_i e_j)]_{ij} \\ &= \det\left([\sigma_k(e_i)]_{ik}[\sigma_k(e_j)]_{kj}\right) \\ &= \det\left([\sigma_k(e_i)]_{ik}[\sigma_k(e_j)]^{\mathrm{t}}_{jk}\right) \\ &= (\det[\sigma_i(e_j)]_{ij})^2 \end{aligned}$$

since the determinant is multiplicative and invariant under taking transposes.

Now let $x \in L$ and define $e_i := x^{i-1}$ for $1 \le i \le n$. Then

$$\operatorname{disc}(1, x, x^2, \ldots, x^{n-1}) = \left(\det[\sigma_i(x^{j-1})]_{ij}\right)^2 = \prod_{i<j} (\sigma_i(x) - \sigma_j(x))^2 ,$$

since $[\sigma_i(x)^{j-1})]_{ij}$ is a Vandermonde matrix. ☐

[1] This definition is consistent with Definition 4.49 where we defined the discriminant of a bilinear pairing.

Definition 11.14. For a polynomial $f(x) = \prod_i (x - \alpha_i)$, the *discriminant* of f is

$$\operatorname{disc}(f) := \prod_{i<j} (\alpha_i - \alpha_j)^2.$$

Equivalently, if A is a Dedekind domain, $f \in A[x]$ is a monic separable polynomial, and α is the image of x in $A[x]/(f(x))$, then

$$\operatorname{disc}(f) = \operatorname{disc}(1, \alpha, \alpha^2, \ldots, \alpha^{n-1}) \in A.$$

Example 11.15. $\operatorname{disc}(x^2 + bx + c) = b^2 - 4c$ and $\operatorname{disc}(x^3 + ax + b) = -4a^3 - 27b^2$.

Now assume $AKLB$ and let M be an A-lattice in L (Definition 6.1). Then M is a finitely generated A-module that contains a basis for L as a K-vector space, but we would like to define the discriminant of M in a way that does not require us to choose a basis.

Let us first consider the case where M is a free A-lattice. If $e_1, \ldots, e_n \in M \subseteq L$ and $e_1', \ldots, e_n' \in M \subseteq L$ are two bases for M, then

$$\operatorname{disc}(e_1', \ldots, e_n') = u^2 \operatorname{disc}(e_1, \ldots, e_n)$$

for some unit $u \in A^\times$; this follows from the fact that the change of basis matrix $P \in A^{n\times n}$ is invertible and its determinant is therefore a unit u. This unit gets squared because we need to apply the change of basis twice in order to change $\mathrm{T}(e_i e_j)$ to $\mathrm{T}(e_i' e_j')$. Explicitly, writing bases as row-vectors, let $e = (e_1, \ldots, e_n)$, $e' = (e_1', \ldots, e_n')$ and suppose $e' = eP$. We then have

$$\begin{aligned}
\operatorname{disc}(e') &= \det[\mathrm{T}_{L/K}(e_i' e_j')]_{ij} \\
&= \det[\mathrm{T}_{L/K}((eP)_i (eP)_j)]_{ij} \\
&= \det[P^{\mathrm{t}} \mathrm{T}_{L/K}(e_i e_j) P]_{ij} \\
&= (\det P^{\mathrm{t}}) \operatorname{disc}(e) (\det P) \\
&= (\det P)^2 \operatorname{disc}(e),
\end{aligned}$$

where we have (repeatedly) used the fact that $\mathrm{T}_{L/K}$ is A-linear.

This actually gives us an unambiguous definition when $A = \mathbb{Z}$: the only units in $\mathbb{Z}$ are $u = \pm 1$, so we always have $u^2 = 1$ and get the same discriminant no matter which basis we choose. In general we want to take the principal fractional ideal of A generated by $\operatorname{disc}(e_1, \ldots, e_n)$, which does not depend on the choice of basis. This suggests how we should define the discriminant of M in the general case, where M is not necessarily free.

Definition 11.16. Assume $AKLB$ and let M be an A-lattice in L. The *discriminant* $D(M)$ of M is the A-module generated by the set $\{\operatorname{disc}(e_1, \ldots, e_n) : e_1, \ldots, e_n \in M\}$.

In the case that M is free, $D(M)$ is equal to the principal fractional ideal generated by $\operatorname{disc}(e_1, \ldots, e_n)$, for any fixed basis $e = (e_1, \ldots, e_n)$. For any n-tuple $e' = (e_1', \ldots, e_n')$ of elements in L, we can write $e' = eP$ for some (not necessarily invertible) matrix P; we will have $\operatorname{disc}(e') = 0$ whenever e' is not a basis.

Lemma 11.17. *Assume $AKLB$ and let $M \subseteq M'$ be free A-lattices in L. If $D(M) = D(M')$ then $M = M'$.*

Proof. Fix bases e and e' for M and M'. If $D(M) = (\operatorname{disc}(e)) = (\operatorname{disc}(e')) = D(M')$ as fractional ideals of A, then the change of basis matrix from M' to M is invertible over A, which implies $M' \subseteq M$ and therefore $M = M'$. $\square$

In general, $D(M)$ is a fractional ideal of A, but it need not be principal.

Proposition 11.18. *Assume $AKLB$ and let M be an A-lattice in L. Then $D(M) \in \mathcal{I}_A$.*

Proof. The A-module $D(M)$ is nonzero because M contains a K-basis $e_1, \ldots, e_n$ for L and $\operatorname{disc}(e_1, \ldots, e_n) \neq 0$ because the trace pairing is nondegenerate, and it is clearly a submodule of the fraction field K of A (it is generated by determinants of matrices with entries in K). To show that $D(M)$ is finitely generated as an A-module we use the usual trick: show that it is a submodule of a noetherian module. Let N be the free A-lattice generated by a K-basis of L in M. Since N is finitely generated, we can pick a nonzero $a \in A$ such that $M \subseteq a^{-1}N$. Then $D(M) \subset D(a^{-1}N)$, and since $a^{-1}N$ is a free A-lattice, $D(a^{-1}N)$ is finitely generated and therefore a noetherian module, since A is noetherian. Every submodule of a noetherian module is finitely generated, so $D(M)$ is finitely generated. $\square$

Definition 11.19. Assume $AKLB$. The *discriminant* of L/K is the discriminant of B as an A-module:

$$D_{L/K} := D_{B/A} := D(B) \in \mathcal{I}_A.$$

Note that $D_{L/K}$ is a fractional ideal (in fact an ideal, by Corollary 11.24 below), not an element of A (but see Remark 11.21 below).

Example 11.20. Consider the case $A = \mathbb{Z}$, $K = \mathbb{Q}$, $L = \mathbb{Q}(i)$, $B = \mathbb{Z}[i]$. Then B is a free A-lattice with basis $(1, i)$ and we can compute $D_{L/K}$ in three ways:

- $\operatorname{disc}(1, i) = \det \begin{bmatrix} \mathrm{T}_{L/K}(1 \cdot 1) & \mathrm{T}_{L/K}(1 \cdot i) \\ \mathrm{T}_{L/K}(i \cdot 1) & \mathrm{T}_{L/K}(i \cdot i) \end{bmatrix} = \det \begin{bmatrix} 2 & 0 \\ 0 & -2 \end{bmatrix} = -4.$
- The non-trivial automorphism of L/K fixes 1 and sends i to $-i$, so we could instead compute $\operatorname{disc}(1, i) = \left(\det \begin{bmatrix} 1 & 1 \\ i & -i \end{bmatrix} \right)^2 = (-2i)^2 = -4.$
- We have $B = \mathbb{Z}[i] = \mathbb{Z}[x]/(x^2 + 1)$ and can compute $\operatorname{disc}(x^2 + 1) = -4$.

In every case the discriminant ideal $D_{L/K}$ is $(-4) = (4)$.

Remark 11.21. If $A = \mathbb{Z}$ then B is the ring of integers of the number field L, and B is a free A-lattice, because it is a torsion-free module over a PID and therefore a free module. In this situation it is customary to define the *absolute discriminant* D_L of the number field L to be the *integer* $\operatorname{disc}(e_1, \ldots, e_n) \in \mathbb{Z}$, for any basis $(e_1, \ldots, e_n)$ of B, rather than the ideal it generates. As noted above, this integer is independent of the choice of basis because $u^2 = 1$ for any $u \in \mathbb{Z}^\times$; in particular, the sign of D_L is well defined. In the example above, the absolute discriminant is $D_L = -4$ (not 4).

We now show that the discriminant respects localization.

Proposition 11.22. *Assume $AKLB$ and let S be a multiplicative subset of A. Then $S^{-1}D_{B/A} = D_{S^{-1}B/S^{-1}A}$.*

Proof. Let $x = s^{-1}\operatorname{disc}(e_1,\ldots,e_n) \in S^{-1}D_{B/A}$ for some $s \in S$ and $e_1,\ldots,e_n \in B$. Then $x = s^{2n-1}\operatorname{disc}(s^{-1}e_1,\ldots,s^{-1}e_n)$ lies in $D_{S^{-1}B/S^{-1}A}$. This proves the forward inclusion.

Conversely, for any $e_1,\ldots,e_n \in S^{-1}B$ we can choose a single $s \in S \subseteq A$ so that each se_i lies in B. We then have $\operatorname{disc}(e_1,\ldots,e_n) = s^{-2n}\operatorname{disc}(se_1,\ldots,se_n) \in S^{-1}D_{B/A}$, which proves the reverse inclusion. □

We have now defined two different ideals associated to a finite separable extension of Dedekind domains B/A in the $AKLB$ setup. We have the different $\mathcal{D}_{B/A}$, which is a fractional ideal of B, and the discriminant $D_{B/A}$, which is a fractional ideal of A. We now relate these two ideals in terms of the ideal norm $N_{B/A}\colon \mathcal{I}_B \to \mathcal{I}_A$, which for $I \in \mathcal{I}_B$ is defined as $N_{B/A}(I) := (B : I)_A$, where $(B : I)_A$ is the module index (see Definitions 6.2 and 6.5). We recall that $N_{B/A}(I)$ is also equal to the ideal generated by the image of I under the field norm $\mathrm{N}_{L/K}$; see Corollary 6.8.

Theorem 11.23. *Assume $AKLB$. Then $D_{B/A} = N_{B/A}(\mathcal{D}_{B/A})$.*

Proof. The different respects localization at any prime $\mathfrak{p}$ of A (see Proposition 11.10), and we just proved that this is also true of the discriminant. Since A is a Dedekind domain, the A-modules on both sides of the equality are determined by the intersections of their localization, so it suffices to consider the case that $A = A_{\mathfrak{p}}$ is a DVR, and in particular a PID. In this case B is a free A-lattice in L (torsion-free over a PID implies free), and we can choose a basis $e_1,\ldots,e_n$ for B as an A-module. The dual A-module

$$B^* = \{x \in L : \mathrm{T}_{L/K}(xb) \in A \ \forall b \in B\} \in \mathcal{I}_B$$

is also a free A-lattice in L, with basis $e_1^*,\ldots,e_n^*$ uniquely determined by $\mathrm{T}_{L/K}(e_i^*e_j) = \delta_{ij}$.

If M is any free A-lattice with basis $m_1,\ldots,m_n$, then $[\mathrm{T}_{L/K}(m_ie_j)]_{ij}$ is precisely the change of basis matrix from $e_1^*,\ldots,e_n^*$ to $m_1,\ldots,m_n$. Applying this to the free A-lattice B, we then have

$$D_{B/A} = \left(\det[\mathrm{T}_{L/K}(e_ie_j)]_{ij}\right) = (B^* : B)_A,$$

by the definition of the module index for free A-modules (see Definition 6.2).

For any $I \in \mathcal{I}_B$ we have $(B : I) = I^{-1} = (I^{-1} : B)$ as B-modules, and it follows that $(B : I)_A = (I^{-1} : B)_A$. Applying this with $I^{-1} = B^*$ gives

$$D_{B/A} = (B^* : B)_A = (B : (B^*)^{-1})_A = (B : \mathcal{D}_{B/A})_A = N_{B/A}(\mathcal{D}_{B/A})$$

as claimed. □

Corollary 11.24. *Assume $AKLB$. The discriminant $D_{B/A}$ is an A-ideal.*

Proof. The different $\mathcal{D}_{B/A}$ is a B-ideal, and the field norm $N_{L/K}$ sends elements of B to A; it follows that $D_{B/A} = N_{B/A}(\mathcal{D}_{B/A}) = (\{N_{L/K}(x) : x \in \mathcal{D}_{B/A}\})$ is an A-ideal. □

18.785 Number Theory I
Fall 2015