12 The different and the discriminant

12.1 The different

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. Recall that a prime $\mathfrak{q}|\mathfrak{p}$ of L is unramified if and only if the residue field $B/\mathfrak{q}$ is a finite étale $A/\mathfrak{p}$-algebra; this is equivalent to requiring $v_\mathfrak{q}(\mathfrak{p}B) = 1$ with $B/\mathfrak{q}$ a separable extension of $A/\mathfrak{p}$. A prime $\mathfrak{p}$ of K is unramified if and only if all the primes $\mathfrak{q}|\mathfrak{p}$ lying above it are unramified.[1]

Our main tools for doing are the *different ideal* $\mathcal{D}_{B/A}$ and the *discriminant ideal* $D_{B/A}$. The different ideal is an ideal of B and the discriminant ideal is an ideal of A (the norm of the different ideal, in fact). We will show that the primes of B that ramify are exactly those that divide $\mathcal{D}_{B/A}$, the primes of A that ramify are exactly those that divide $D_{B/A}$. Moreover, the valuation $v_\mathfrak{q}(\mathcal{D}_{B/A})$ will give us information about the ramification index $e_\mathfrak{q}$ (its exact value in the tamely ramified case). We could just define $\mathcal{D}_{B/A}$ and $D_{B/A}$ to have the properties we want, but the key is to define them in an intrinsic way that makes it possible to compute them without knowing which primes ramify; indeed, there main purpose is to allow us to determine these primes.

Recall from Lecture 5 the trace pairing $L\times L \to K$ defined by $(x, y) \mapsto \mathrm{T}_{L/K}(xy)$; under our assumption that L/K is separable, it is a perfect pairing (see Proposition 5.18). An A-lattice M in L is a finitely generated A-module that spans L as a K-vector space (see Definition 5.8). Associated to any A-lattice M is its *dual lattice* (with respect to the trace pairing), which is defined by

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

(see Definition 5.10); it is an A-lattice isomorphic to the dual module $M^\vee := \mathrm{Hom}_A(M, A)$ (see Theorem 5.11), and in our $AKLB$ setting we have $M^{**} = M$ (see Proposition 5.14).

Every fractional ideal I of B is finitely generated as a B-module, and therefore finitely generated as an A module (since B is finite over A). If I is nonzero, it spans L (if $e_1, \ldots, e_n$ is a K-basis for L in B and $a \in I$ is nonzero then $ae_1, \ldots, ae_n$ is a K-basis for L in I). It follows that every element of the group $\mathcal{I}_B$ of nonzero fractional ideals of B is an A-lattice in L. We now show that $\mathcal{I}_B$ is closed under the operation of taking duals.

Lemma 12.1. *Assume $AKLB$ and let $I \in \mathcal{I}_B$. Then $I^* \in \mathcal{I}_B$.*

Proof. As noted above, I is an A-lattice in L, as is its dual lattice I^* which is a nonzero finitely generated A-module; if I^* is a B-module then it is certainly finitely generated, hence a fractional ideal of B. Thus to show $I^* \in \mathcal{I}_B$ we just need to show that I^* is a B-module. For any $b \in B$ and $x \in I^* \subseteq L$ the product bx lies in L, we just need to check that it lies in I^*. For any $m \in I$ we have $bm \in I$, since I is a B-module, and $T_{L/K}(x(bm)) \in A$, by the definition of I^*. Thus $\mathrm{T}_{L/K}((bx)m) = \mathrm{T}_{L/K}(x(bm)) \in A$ so $bx \in I^*$. $\square$

[1] As usual, by a *prime* of A or K (resp., B or L) we mean a nonzero prime ideal of A (resp., B). In our $AKLB$ setting the notation $\mathfrak{q}|\mathfrak{p}$ means that $\mathfrak{q}$ is a prime of B lying above $\mathfrak{p}$ (so $\mathfrak{p} = \mathfrak{q} \cap A$ and $\mathfrak{q}$ divides $\mathfrak{p}B$).

Lecture by Andrew Sutherland

Definition 12.2. Assume $AKLB$. The *different ideal* is the inverse of B^* in $\mathcal{I}_B$. That is,

$$B^* := \{x \in L : \mathrm{T}_{L/K}(xb) \in A \text{ for all } b \in B\},$$
$$\mathcal{D}_{B/A} := (B^*)^{-1} = (B : B^*) = \{x \in L : xB^* \subseteq B\}.$$

Note that $B \subseteq B^*$, since $\mathrm{T}_{L/K}(ab) = T_{L/K}(b) \in A$ for all $a, b \in B$, and this implies $(B^*)^{-1} \subseteq B^{-1} = B$; so $\mathcal{D}_{B/A}$ is an ideal, not just a fractional ideal.

We now show that the different respects localization and completion.

Proposition 12.3. *Assume $AKLB$ and let S be a multiplicative subset of A. Then*

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

Proof. This follows the fact that inverses and duals are both compatible with localization; see Lemmas 3.13 and 5.13. Note that a multiplicative subset of A is also a multiplicative subset of B and the localization of a B-module with respect to S is the same as its localization as an A-module with respect to S. $\square$

Proposition 12.4. *Assume $AKLB$ and let $\mathfrak{q}|\mathfrak{p}$ be a prime of B. Then*

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

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 11.20, 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 11.23, $\hat{B} = \prod_{\mathfrak{q}|\mathfrak{p}} \hat{B}_{\mathfrak{q}}$, 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 $(\hat{B}^*_{\mathfrak{q}})^{-1} = \mathcal{D}_{\hat{B}_{\mathfrak{q}}/\hat{A}}$. $\square$

12.2 The discriminant

Definition 12.5. Let S/R be a ring extension with S free as an R-module. For any $x_1, \ldots, x_n \in S$ we define the *discriminant*

$$\mathrm{disc}(x_1, \ldots, x_n) := \det[\mathrm{T}_{S/R}(x_i x_j)]_{i,j} \in R.$$

(note that the $e_1, \ldots, e_n$ may be any elements of S, they need not be an R-basis).

In our $AKLB$ setup, we have in mind the case where $e_1, \ldots, e_n \in B$ is a basis for L as a K-vector space, in which case $\mathrm{disc}(e_1, \ldots, e_n) = \det[\mathrm{T}_{L/K}(e_i e_j)]_{ij} \in A$. Note that we are not assuming B is a free A-module, but L is certainly a free K-module, so we can compute the discriminant of any set of elements of L (including elements of B).

Proposition 12.6. *Let L/K be a finite separable extension of degree n, and let Ω/K be a field extension for which there are distinct $\sigma_1, \ldots, \sigma_n \in \mathrm{Hom}_K(L, \Omega)$. For any $e_1, \ldots, e_n \in L$*

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

and for any $x \in L$ we have

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

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

Proof. For $1 \le i, j \le n$ we have $\mathrm{T}_{L/K}(e_ie_j) = \sum_{k=1}^n \sigma_k(e_ie_j)$, by Theorem 4.46. Therefore

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

since the determinant is multiplicative and and $\det M = \det M^{\mathrm{t}}$ for any matrix M.

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

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

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

Definition 12.7. 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 12.8. $\operatorname{disc}(x^2 + bx + c) = b^2 - 4c$ and $\operatorname{disc}(x^3 + ax + b) = -4a^3 - 27b^2$.

Now assume $AKLB$ and let M be an A-lattice in L. 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 $\mathrm{T}(e_ie_j)$ to $\mathrm{T}(e_i'e_j')$. Explicitly, writing bases as row-vectors, let $e = (e_1, \ldots, e_n)$, $e' = (e_1', \ldots, e_n')$ with $e' = eP$. Then

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

where we have used the fact that $\mathrm{T}_{L/K}$ is K-linear, the determinant is multiplicative, and $\det P^{\mathrm{t}} = \det P$.

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 discriminant of every basis is the same. 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 (multiplying a fractional ideal by a unit does nothing).

Definition 12.9. Assume $AKLB$ and let M be an A-lattice in L. The *discriminant* $D(M)$ of M is the A-submodule of K generated by $\{\operatorname{disc}(e_1, \ldots, e_n) : e_1, \ldots, e_n \in M\}$.

When M is free, $D(M)$ is the principal fractional ideal generated by $\operatorname{disc}(e_1, \ldots, e_n)$, where $e \coloneqq e_1, \ldots, e_n$ is any A-basis for M. Given any n-tuple $e' = (e'_1, \ldots, e'_n)$ of elements in M, if we view e and e' as row vectors we can write $e' = eP$ for some (not necessarily invertible) matrix $P \in A^{n\times n}$, and we always have $\operatorname{disc}(e') = (\det P)^2 \operatorname{disc}(e) \in (\operatorname{disc}(e))$.

Lemma 12.10. *Assume $AKLB$ and let $M' \subseteq M$ be free A-lattices in L. If $D(M') = D(M)$ then $M' = M$.*

Proof. Fix A-bases e and e' for M and M'. Then $e' = eP$ for some $P \in A^{n\times n}$, and we have

$$D(M') = (\operatorname{disc}(e')) = (\operatorname{disc}(eP)) = ((\det P)^2) \operatorname{disc}(e)) = (\det P)^2 D(M),$$

as fractional ideals of A. The fact that e is a basis for L and the trace pairing is nondegenerate guarantees that $\operatorname{disc}(e) \neq 0$. Now A is a Dedekind domain, so if $D(M') = D(M)$ then $(\det P)$ must be the unit ideal (multiply both sides by $D(M)^{-1}$), and $\det P$ must be a unit, which implies P is invertible. We then have $e = e'P^{-1}$, thus $M \subseteq M'$ and $M' = M$. $\square$

Proposition 12.11. *Assume $AKLB$ and let M be an A-lattice in L. Then $D(M) \in \mathcal{I}_A$.*

Proof. The A-module $D(M) \subseteq K$ is nonzero because M contains a K-basis $e = (e_1, \ldots, e_n)$ for L and $\operatorname{disc}(e) \neq 0$ because the trace pairing is nondegenerate. To show that $D(M)$ is a 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 in L generated by e. The A-lattice M is finitely generated, so we can pick a nonzero $a \in A$ such that $M \subseteq a^{-1}N$: write each generator for M in terms of the K-basis e and let a be the product of all the denominators that appear. We then have $D(M) \subseteq D(a^{-1}N)$, and since $a^{-1}N$ is a free A-lattice, $D(a^{-1}N)$ is a principal fractional ideal of A, hence a noetherian A-module (since A is noetherian), and this implies that the A-submodule $D(M)$ is finitely generated. $\square$

Definition 12.12. Assume $AKLB$. The *discriminant* of L/K is the discriminant of B as an A-module:

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

Example 12.13. 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 12.14. 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 12.15. *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. $\square$

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 6.1 and 6.4). 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 12.16. *Assume AKLB. Then $D_{B/A} = N_{B/A}(\mathcal{D}_{B/A})$.*

Proof. The different and discriminant are both compatible with localization, by Propositions 12.3 and 12.15, and the fractional ideals $D_{B/A}$ and $N_{B/A}(\mathcal{D}_{B/A})$ of A are both determined by the intersections of their localizations at prime ideals (Proposition 2.7), so it suffices to prove that the theorem holds when $A = A_\mathfrak{p}$ is a DVR, and in particular a PID (here we are using the fact that A is a Dedekind domain). In this case B is a free A-lattice in L, and we can choose a basis $(e_1, \ldots, e_n)$ for B as an A-module. The dual A-lattice

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

is also a free A-module, with basis $(e_1^*, \ldots, e_n^*)$ uniquely determined by $\mathrm{T}_{L/K}(e_i^* e_j) = \delta_{ij}$. If we write $e_i = \sum a_{ij} e_j^*$ in terms of the K-basis $(e_1^*, \ldots, e_n^*)$ for L then

$$T_{L/K}(e_i e_j) = T_{L/K}\left(\sum_k a_{ik} e_k^* e_j\right) = \sum_k a_{ik} T_{L/K}(e_k^* e_j) = \sum_k a_{ik}\delta_{kj} = a_{ij},$$

so $P := [\mathrm{T}_{L/K}(e_i e_j)]_{ij}$ is the change-of-basis matrix from $e^* := (e_1^*, \ldots, e_n^*)$ to $e := (e_1, \ldots, e_n)$ (as row vectors we have $e = e^* P$). If we let $\phi\colon B^* \to B$ denote the linear transformation with matrix P, then ϕ is an isomorphism of free A-modules and

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

where $[B^* : B]_A$ is the module index (see Definition 6.1). Applying Corollary 6.7 yields

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

Corollary 12.17. *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. □

12.3 Ramification

Having defined the different and discriminant ideals we now consider what they can tell us about ramification. Recall that in our $AKLB$ setup, if $\mathfrak{p}$ is a prime of A with

$$\mathfrak{p}B = \prod \mathfrak{q}_1^{e_1} \cdots \mathfrak{q}_r^{e_r},$$

each prime $\mathfrak{q}_i$ is unramified if and only if $e_i = 1$ and the residue field $B/\mathfrak{q}_i$ is a separable extension of $A/\mathfrak{p}$, and $\mathfrak{p}$ is unramified if and only if all the $\mathfrak{q}_i$ are unramified. As noted in Definition 5.33, an equivalent definition is that $B/\mathfrak{p}B$ is a finite étale $A/\mathfrak{p}$-algebra (a finite product of finite separable extensions of $A/\mathfrak{p}$). To see this, note that the Chinese remainder theorem implies

$$B/\mathfrak{p}B \simeq B/\mathfrak{q}_1^{e_1} \times \cdots \times B/\mathfrak{q}_r^{e_r},$$

and if any $e_i > 1$ then $B/\mathfrak{p}B$ contains a nonzero nilpotent element (take a uniformizer for $\mathfrak{q}_i$). In this case $B/\mathfrak{p}B$ cannot be étale, since a product of fields has no nonzero nilpotents. If every $e_i = 1$, then $B/\mathfrak{p}B$ is isomorphic to the product of the residue fields $B/\mathfrak{q}_i$, each of which is a finite extension of $A/\mathfrak{p}$. In this case $B/\mathfrak{p}B$ is étale if and only if these extensions are all separable, equivalently, if and only if all the $\mathfrak{q}_i$ are unramified.

We now relate the property of being finite étale to the discriminant.

Lemma 12.18. *Let k be a field and let R be a commutative k-algebra that is a finite dimensional k-vector space with basis $r_1, \ldots, r_n$. Then R is a finite étale k-algebra if and only if* $\operatorname{disc}(r_1, \ldots, r_n) = \det[\mathrm{T}_{R/k}(r_i r_j)]_{ij} \neq 0$.

Proof. We first note that the choice of basis is immaterial, changing the basis will not change whether the discriminant is zero or nonzero.

Suppose R contains a nonzero nilpotent r (so $r^m = 0$ for some $m > 1$). In this case R cannot be finite étale, and we can extend $\{r\}$ to a basis, so we may assume $r_1 = r$ is nilpotent. Every multiple of r_1 is also nilpotent, and it follows that the first row of the matrix $[\mathrm{T}_{R/k}(r_i r_j)]_{ij}$ is zero, since the trace of any nilpotent element s is zero (the eigenvalues of the multiplication-by-s map must all be zero). Therefore $\operatorname{disc}(r_1, \ldots, r_n) \det[\mathrm{T}_{R/k}(r_i r_j)]_{ij} = 0$.

Suppose R contains no nonzero nilpotents. Then R is isomorphic to a product of fields, each of which is a finite extension of k (this is a standard result of commutative algebra which follows, for example, from Lemmas 10.52.2-5 of [3]). Without loss of generality we can assume our basis contains k-bases for each of these field extensions, grouped together so that the matrix $[\mathrm{T}_{R/k}(r_i r_j)]_{ij}$ is block diagonal. The determinant is then nonzero if and only if the determinant of each block is nonzero, so we can reduce to the case where R/k is a field extension. The proof then follows from the fact that the trace pairing $\mathrm{T}_{R/k}$ is nondegenerate if and only if R/k is separable (see Proposition 5.18). □

Theorem 12.19. *Assume AKLB, let $\mathfrak{q}$ be a prime of B lying above a prime $\mathfrak{p}$ of A. The extension L/K is unramified at $\mathfrak{q}$ if and only if $\mathfrak{q}$ does not divide $\mathcal{D}_{B/A}$, and it is unramified at $\mathfrak{p}$ if and only if $\mathfrak{p}$ does not divide $D_{B/A}$.*

Proof. We first consider the different ideal $\mathcal{D}_{B/A}$. By Proposition 12.4, the different is compatible with completion, so it suffices to consider the case that A and B are complete DVRs (complete K at $\mathfrak{p}$ and L at $\mathfrak{q}$ and apply Theorem 11.20). We then have $[L:K] = e_{\mathfrak{q}} f_{\mathfrak{q}}$, where $e_{\mathfrak{q}}$ is the ramification index and $f_{\mathfrak{q}}$ is the residue field degree, and $\mathfrak{p}B = \mathfrak{q}^{e_{\mathfrak{q}}}$.

Since B is a DVR with maximal ideal $\mathfrak{q}$, we must have $\mathcal{D}_{B/A} = \mathfrak{q}^m$ for some $m \geq 0$. By Theorem 12.16 we have

$$D_{B/A} = N_{B/A}(\mathcal{D}_{B/A}) = N_{B/A}(\mathfrak{q}^m) = \mathfrak{p}^{f_{\mathfrak{q}} m}.$$

Thus $\mathfrak{q} | \mathcal{D}_{B/A}$ if and only if $\mathfrak{p} | D_{B/A}$. Since A is a PID, B is a free A-module and we may choose an A-module basis $e_1, \ldots, e_n$ for B that is also a K-vector space for L. Let $k \coloneqq A/\mathfrak{p}$, and let $\bar{e}_i$ be the reduction of e_i to the k-algebra $R \coloneqq B/\mathfrak{p}B$. Then $(\bar{e}_1, \ldots, \bar{e}_n)$ is a k-basis for R: it clearly spans, and we have $[R:k] = [B/\mathfrak{q}^{e_{\mathfrak{q}}} : A_{\mathfrak{p}}] = e_{\mathfrak{q}} f_{\mathfrak{q}} = [L:K] = n$.

Since B has an A-module basis, we may compute its discriminant as

$$D_{B/A} = (\operatorname{disc}(e_1, \ldots, e_n)).$$

Thus $\mathfrak{p} | D_{B/a}$ if and only if $\operatorname{disc}(e_1, \ldots, e_n) \in \mathfrak{p}$, equivalently, $\operatorname{disc}(\bar{e}_1, \ldots, \bar{e}_n) = 0$ (note that $\operatorname{disc}(e_1, \ldots, e_n)$ is a polynomial in the $\mathrm{T}_{L/K}(e_i e_j)$ and $T_{R/k}(\bar{e}_i \bar{e}_j)$ is the trace of the multiplication-by-$\bar{e}_i \bar{e}_j$ map, which is the same as the reduction to $k = A/\mathfrak{p}$ of the trace of the multiplication-by-$e_i e_j$ map $T_{L/K}(e_i e_j) \in A$). By Lemma 12.18, $\operatorname{disc}(\bar{e}_1, \ldots, \bar{e}_n) = 0$ if and only if the k-algebra $B/\mathfrak{p}B$ is not finite étale, equivalently, if and only if $\mathfrak{p}$ is ramified. Thus $\mathfrak{p} | D_{B/A}$ if and only if $\mathfrak{p}$ is ramified. There is only one prime $\mathfrak{q}$ above $\mathfrak{p}$, so we also have $\mathfrak{q} | \mathcal{D}_{B/A}$ if and only if $\mathfrak{q}$ is ramified. □

We now note an important corollary of Theorem 12.19.

Corollary 12.20. *Assume AKLB. Only finitely many primes of A (or B) ramify.*

Proof. Both A and B are Dedekind domains, so the ideals $D_{B/A}$ and $\mathcal{D}_{B/A}$ both have unique factorizations into prime ideals in which only finitely many primes appear. □

12.4 The discriminant of an order

Recall from Lecture 6 that an order $\mathcal{O}$ is a noetherian domain of dimension one whose conductor is nonzero (see Definitions 6.15 and 6.18), and the integral closure of an order is always a Dedekind domain. In our $AKLB$ setup, the orders with integral closure B are precisely the A-lattices in L that are rings (see Proposition 6.21); if $L = K(\alpha)$ with $\alpha \in B$ then $A[\alpha]$ is an example. The discriminant $D_{\mathcal{O}/A}$ of such an order $\mathcal{O}$ is its discriminant $D(\mathcal{O})$ as an A-module. The fact that $\mathcal{O} \subseteq B$ implies that $D(\mathcal{O}) \subseteq D_{B/A}$ is an A-ideal.

If $\mathcal{O}$ is an order of the form $A[\alpha]$, where $\alpha \in B$ generates $L = K(\alpha)$ with minimal polynomial $f \in A[x]$, then $\mathcal{O}$ is a free A-lattice with basis $1, \alpha, \ldots, \alpha^{n-1}$, where $n = \deg f$, and we may compute its discriminant as

$$D_{\mathcal{O}/A} = (\operatorname{disc}(1, \alpha, \ldots, \alpha^{n-1})) = (\operatorname{disc}(f)),$$

which is a principal A-ideal contained in $D_{B/A}$. If B is also a free A-lattice, then as in the proof of Lemma 12.10 we have

$$D_{\mathcal{O}/A} = (\det P)^2 D_{B/A} = [B : \mathcal{O}]_A^2 D_{B/A},$$

where P is the matrix of the A-linear map $\phi\colon B \to \mathcal{O}$ that sends an A-basis for B to an A-basis for $\mathcal{O}$ and $[B:\mathcal{O}]_A$ is the module index (a principal A-ideal).

In the important special case where $A = \mathbb{Z}$ and L is a number field, the integer $(\det P)^2$ is uniquely determined and it necessarily divides $\operatorname{disc}(f)$, the generator of the principal ideal $D(\mathcal{O}) = D(A[\alpha])$. It follows that if $\operatorname{disc}(f)$ is squarefree then we must have $B = \mathcal{O} = A[\alpha]$. More generally, any prime p for which $v_p(\operatorname{disc}(f))$ is odd must be ramified, and any prime that does not divide $\operatorname{disc}(f)$ must be unramified.

Another useful observation that applies when $A = \mathbb{Z}$ is that in this case the module index $[B : \mathcal{O}]_{\mathbb{Z}} = ([B : \mathcal{O}])$ is the principal ideal generated by the index of $\mathcal{O}$ in B (as $\mathbb{Z}$-lattices), and we have

$$D_{\mathcal{O}/A} = [B:\mathcal{O}]^2 D_{B/A}.$$

Example 12.21. Consider $A = \mathbb{Z}$, $K = \mathbb{Q}$ with $L = \mathbb{Q}(\alpha)$, where $\alpha^3 - \alpha - 1 = 0$. We would like to determine the primes that ramify in L and describe its ring of integers $B = \mathcal{O}_L$. We can compute the absolute discriminant of $\mathbb{Z}[\alpha]$ as

$$\operatorname{disc}(1, \alpha, \alpha^2) = \operatorname{disc}(x^3 - x - 1) = -4(-1)^3 - 27(-1)^2 = -23.$$

This immediately implies that 23 is the only prime of that ramifies. The $\mathbb{Z}$-ideal $D(\mathbb{Z}[\alpha])$ is principal (because $\mathbb{Z}$ is a PID) and therefore must be generated by the integer $-23/m^2$, where $m = [\mathcal{O}_L : \mathbb{Z}[\alpha]]$; this implies $m = 1$, so $\mathcal{O}_L = \mathbb{Z}[\alpha]$.

More generally, we have the following theorem.

Theorem 12.22. *Assume AKLB and let $\mathcal{O}$ be an order with integral closure B and conductor $\mathfrak{c}$. Then $D_{\mathcal{O}/A} = \mathrm{N}_{B/A}(\mathfrak{c}) D_{B/A}$.*

Proof. See Problem Set 6. □

12.5 Computing the discriminant and different

We conclude with a number of results that allow one to explicitly compute the discriminant and different in many cases.

Proposition 12.23. *Assume AKLB. If $B = A[\alpha]$ for some $\alpha \in L$ and $f \in A[x]$ is the minimal polynomial of α, then*

$$\mathcal{D}_{B/A} = (f'(\alpha))$$

is the B-ideal generated by $f'(\alpha)$.

Proof. See Problem Set 6. □

The assumption $B = A[\alpha]$ in Proposition 12.23 does not always hold, but if we want to compute the power of $\mathfrak{q}$ that divides $\mathcal{D}_{B/A}$ we can complete L at $\mathfrak{q}$ and K at $\mathfrak{p} = \mathfrak{q} \cap A$ so that A and B become complete DVRs, in which case $B = A[\alpha]$ does hold (by Lemma 10.15), so long as the residue field extension is separable (always true if K and L are global fields, since the residue fields are then finite, hence perfect). The following definition and proposition give an alternative approach.

Definition 12.24. Assume $AKLB$ and let $\alpha \in B$ have minimal polynomial $f \in A[x]$. The *different of α* is defined by

$$\delta_{B/A}(\alpha) = \begin{cases} f'(\alpha) & \text{if } L = K(\alpha), \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 12.25. *Assume AKLB. Then* $\mathcal{D}_{B/A} = \big(\delta_{B/A}(\alpha) : \alpha \in B\big)$.

Proof. See [1, Thm. III.2.5]. $\square$

We can now more precisely characterize the ramification information given by the different ideal.

Theorem 12.26. *Assume AKLB and let* $\mathfrak{q}$ *be a prime of* L *lying above* $\mathfrak{p} = \mathfrak{q} \cap A$ *for which the residue field extension* $(B/\mathfrak{q})/(A/\mathfrak{p})$ *is separable. Let* $s = v_\mathfrak{q}(\mathcal{D}_{B/A})$*, let* $e = e_\mathfrak{q}$ *be the ramification index of* $\mathfrak{q}$ *over* $\mathfrak{p}$*, and let* p *be the characteristic of* $A/\mathfrak{p}$*. If* $p \nmid e$ *then*

$$s = e - 1$$

and if $p|e$ *then*

$$e \le s \le e - 1 + e v_\mathfrak{p}(e)$$

Proof. See Problem Set 6. $\square$

We also note the following proposition, which shows how the discriminant and different behave in a tower of extensions.

Proposition 12.27. *Assume AKLB and let* M/L *be a finite separable extension and let* C *be the integral closure of* A *in* M*. Then*

$$\mathcal{D}_{C/A} = \mathcal{D}_{C/B} \cdot \mathcal{D}_{B/A}$$

(where the product on the right is taken in C*), and*

$$D_{C/A} = (D_{B/A})^{[M:L]} N_{B/A}(D_{C/B}).$$

Proof. See [2, Prop. III.8]. $\square$

If $M/L/K$ is a tower of finite separable extensions, we note that the primes $\mathfrak{p}$ of K that ramify are precisely those that divide either $D_{L/K}$ or $N_{L/K}(D_{M/L})$.

References

[1] J. Neukirch, *Algebraic number theory*, Springer, 1999.

[2] J.-P. Serre, *Local fields*, Springer, 1979.

[3] Stacks Project Authors, *Stacks Project*, `http://stacks.math.columbia.edu`.

MIT OpenCourseWare
https://ocw.mit.edu

18.785 Number Theory I
Fall 2016

For information about citing these materials or our Terms of Use, visit: https://ocw.mit.edu/terms.