5 Dedekind extensions

In this lecture we prove that the integral closure of a Dedekind domain in a finite extension of its fraction field is also a Dedekind domain; this implies, in particular, that the ring of integers of a number field is a Dedekind domain. Having accomplished this task we consider the factorization of prime ideals in Dedekind extensions.

5.1 Dual modules, pairings, and lattices

In this section we work in a more general setting, where A is any commutative (unital) ring.

Definition 5.1. Let A be a commutative ring and M an A-module. The *dual module* $M^\vee$ is the A-module $\operatorname{Hom}_A(M,A)$ with scalar multiplication $(af)(m) = af(m)$, where $a \in A$, $f \in \operatorname{Hom}_A(M,A)$, and $m \in M$. If $\varphi\colon M \to N$ is an A-module homomorphism, the dual homomorphism $\varphi^\vee\colon N^\vee \to M^\vee$ is defined by $\varphi^\vee(g)(m) = g(\varphi(m))$, for $g \in N^\vee$ and $m \in M$.

It is easy to check that taking duals preserves identity maps and is compatible with composition: if $\varphi_1 : M \to N$ and $\varphi_2 : N \to P$ are A-module homomorphisms, then $(\varphi_2\varphi_1)^\vee = \varphi_1^\vee\varphi_2^\vee$. We thus have a contravariant functor from the category of A-modules to itself. We have already seen some examples of dual modules: invertible fractional ideals are dual to their inverses.

Proposition 5.2. *Let A be an integral domain with fraction field K and let M be a nonzero A-submodule of K. Then $M^\vee \simeq (A:M) := \{x \in K : xM \subseteq A\}$; in particular, if M is an invertible fractional ideal then $M^\vee \simeq M^{-1}$ and $M^{\vee\vee} \simeq M$.*

Proof. For any $x \in (A:M)$ the map $m \mapsto xm$ is an A-linear map from M to A, hence an element of $M^\vee$, and this defines an A-module homomorphism $\varphi\colon (A:M) \to M^\vee$, since the map $x \mapsto (m \mapsto xm)$ is itself A-linear. Since $M \subseteq K$ is a nonzero A-module, it contains some nonzero $a \in A$ (if $a/b \in M$, so is $ba/b = a$). If $f \in M^\vee$ and $m = b/c \in M$ then

$$f(m) = f\left(\frac{b}{c}\right) = \frac{ac}{ac}f\left(\frac{b}{c}\right) = \frac{b}{ac}f\left(\frac{ac}{c}\right) = \frac{b}{ac}f(a) = f(a)a^{-1}m,$$

where we have used the fact that $a_1f(a_2/a_3) = a_2f(a_1/a_3)$ for any $a_1,a_2,a_3 \in A$ with $a_1/a_3, a_2/a_3 \in M$, by the A-linearity of f. It follows that f corresponds to multiplication by $x = f(a)a^{-1}$, which lies in $(A:M)$ since $xm = f(m) \in A$ for all $m \in M$. The map $f \mapsto f(a)a^{-1}$ defines an A-module homomorphism $M^\vee \to (A:M)$ inverse to φ, so φ is an isomorphism. When M is an invertible fractional ideal we have $M^\vee \simeq (A:M) = M^{-1}$, by Lemma 3.8, and $M^{\vee\vee} \simeq (M^{-1})^{-1} = M$ follows. □

If A is a field, then M is a vector space and $M^\vee$ is its dual space and we have $M^{\vee\vee} \simeq M$, as in the proposition above. But in general not every A-module is isomorphic to its double dual; those that are are said to be *reflexive*.

Example 5.3. As a $\mathbb{Z}$-module, we have $\mathbb{Q}^\vee = \{0\}$ because there are no non-trivial $\mathbb{Z}$-linear homomorphisms from $\mathbb{Q}$ to $\mathbb{Z}$; indeed, $\mathbb{Q}$ is a divisible group and $\mathbb{Z}$ contains no non-trivial divisible subgroups. It follows that $\mathbb{Q}^{\vee\vee} = \{0\}$ (but as $\mathbb{Q}$-modules we have $\mathbb{Q} \simeq \mathbb{Q}^\vee \simeq \mathbb{Q}^{\vee\vee}$). Similarly, the dual of any finite $\mathbb{Z}$-module (any finite abelian group) is the zero module, as is the double dual. More generally, if A is an integral domain every dual (and double dual) A-module must be torsion free, but not all A-modules are torsion free.

One situation where we can recover many of the standard results that hold for vector spaces of finite dimension (with essentially the same proofs), is when M is a free module of finite rank. In particular, not only is M reflexive, we have $M \simeq M^\vee$ (non-canonically) and may explicitly construct a dual basis.

Theorem 5.4. *Let A be a commutative ring and let M be a free A-module of rank n. Then $M^\vee$ is also a free A-module of rank n, and each basis $(e_1, \ldots, e_n)$ of M uniquely determines a dual basis $(e_1^\vee, \ldots e_n^\vee)$ of $M^\vee$ with the property*

$$e_i^\vee(e_j) = \delta_{ij} := \begin{cases} 1 & i = j, \\ 0 & i \neq j. \end{cases}$$

Proof. If $n = 0$ then $M = M^\vee = \{0\}$ and the theorem holds. Now assume $n \geq 1$ and fix an A-basis $\mathbf{e} := (e_1, \ldots, e_n)$ for M. For each $\mathbf{a} := (a_1, \ldots, a_n) \in A^n$, define $f_{\mathbf{a}} \in M^\vee$ by setting $f_{\mathbf{a}}(e_i) = a_i$ and extending A-linearly. The map $\mathbf{a} \mapsto f_{\mathbf{a}}$ gives an A-module homomorphism $A^n \to M^\vee$ with inverse $f \mapsto (f(e_1), \ldots, f(e_n))$ and is therefore an isomorphism. It follows that $M^\vee \simeq A^n$ is a free A-module of rank n.

Now let $e_i^\vee := f_{\mathbf{a}_i}$, where $\mathfrak{a}_i = (0, \ldots, 0, 1, 0, \ldots, 0) \in A^n$ has a 1 in the ith position. Then $(e_1^\vee, \ldots, e_n^\vee)$ is a basis for $M^\vee$, since $(\mathfrak{a}_1, \ldots, \mathfrak{a}_n)$ is a basis for A^n, and we have $e_i^\vee(e_j) = \delta_{ij}$. This choice of $e_i^\vee$ is unique because it must be the image of $(\mathbf{a}_1, \ldots, \mathbf{a}_n)$ under the isomorphism $\mathbf{a} \mapsto f_{\mathbf{a}}$. □

Definition 5.5. Let A be a commutative ring and M an A-module. A (bilinear) *pairing* on M is an A-linear map $\langle \cdot, \cdot \rangle \colon M \times M \to A$. Explicitly, this means that for all $u, v, w \in M$ and $\lambda \in A$ we have

$$\begin{aligned} \langle u + v, w \rangle &= \langle u, w \rangle + \langle v, w \rangle, \\ \langle u, v + w \rangle &= \langle u, v \rangle + \langle u, w \rangle, \\ \langle \lambda u, v \rangle &= \langle u, \lambda v \rangle = \lambda \langle u, v \rangle. \end{aligned}$$

If $\langle v, w \rangle = \langle w, v \rangle$ then $\langle \cdot, \cdot \rangle$ is *symmetric*, if $\langle v, w \rangle = -\langle w, v \rangle$ then $\langle \cdot, \cdot \rangle$ is *skew-symmetric*, and if $\langle v, v \rangle = 0$ then $\langle \cdot, \cdot \rangle$ is *alternating* (the last two are equivalent provided $\operatorname{char}(A) \neq 2$). The pairing $\langle \cdot, \cdot \rangle$ induces an A-module homomorphism

$$\begin{aligned} \varphi \colon M &\to M^\vee \\ m &\mapsto (n \mapsto \langle m, n \rangle) \end{aligned}$$

If $\ker \varphi = \{0\}$ then $\langle \cdot, \cdot \rangle$ is *nondegenerate*, and if φ is an isomorphism then $\langle \cdot, \cdot \rangle$ is *perfect*.

Every perfect pairing is necessarily nondegenerate. If M is a vector space of finite dimension the converse holds, but this is not true in general, not even for free modules of finite rank: consider the pairing $\langle x, y \rangle := 2xy$ on $\mathbb{Z}$, which is non-degenerate but not perfect.

If M is a free A-module with basis $(e_1, \ldots, e_n)$ and $\langle \cdot, \cdot \rangle$ is a perfect pairing, we can apply the inverse of the isomorphism $\varphi \colon M \xrightarrow{\sim} M^\vee$ induced by the pairing to the dual basis $(e_1^\vee, \ldots, e_n^\vee)$ given by Theorem 5.4 to obtain a basis $(e_1', \ldots, e_n')$ for M that satisfies

$$\langle e_i', e_j \rangle = \delta_{ij}.$$

When $\langle \cdot, \cdot \rangle$ is symmetric we can similarly recover $(e_1, \ldots, e_n)$ from $(e_1', \ldots, e_n')$ in the same way. We record this fact in the following proposition.

Proposition 5.6. *Let A be a commutative ring and let M be a free A-module of rank n with a perfect pairing $\langle\cdot,\cdot\rangle$. For each A-basis $(e_1,\ldots,e_n)$ of M there is a unique basis $(e_1',\ldots,e_n')$ for M such that $\langle e_i',e_j\rangle=\delta_{ij}$.*

Proof. Existence follows from the discussion above: apply the inverse of the isomorphism $\varphi\colon V\to V^\vee$ induced by $\langle\cdot,\cdot\rangle$ to the dual basis $(e_1^\vee,\ldots,e_n^\vee)$ given by Theorem 5.4 to obtain a basis $(e_1',\ldots,e_n')$ for M with $e_i'=\varphi^{-1}(e_i^\vee)$. We then have $e_i^\vee=\varphi(e_i')=m\mapsto\langle e_i',m\rangle$ and

$$\langle e_i',e_j\rangle=\varphi(e_i')(e_j)=e_i^\vee(e_j)=\delta_{ij}$$

for $1\le i,j\le n$. If $(f_1',\ldots,f_n')$ is another basis for M with the same property then for each i we have $\langle e_i'-f_i',e_j\rangle=\delta_{ij}-\delta_{ij}=0$ for every e_j, and therefore $\langle e_i'-f_i',m\rangle=0$ for all $m\in M$, but then $e_i'-f_i'\in\ker\varphi=\{0\}$, since the perfect pairing $\langle\cdot,\cdot\rangle$ is nondegenerate, and therefore $f_i'=e_i'$ for each i; uniqueness follows. □

Remark 5.7. In what follows the commutative ring A in Proposition 5.6 will typically be a field K and the free A-module M will be a K-vector space that we will denote V. We may then use A to denote a subring of K and M to denote an A-submodule of V. A perfect paring $\langle\cdot,\cdot\rangle$ on the K-vector space V will typically not restrict to a perfect pairing on the A-module M. For example, the perfect pairing $\langle x,y\rangle=xy$ on $\mathbb{Q}$ does not restrict to a perfect pairing on the $\mathbb{Z}$-module $2\mathbb{Z}$ because the induced map $\varphi\colon 2\mathbb{Z}\to 2\mathbb{Z}^\vee$ defined by $\varphi(m)=(n\mapsto mn)$ is not surjective: the map $x\mapsto x/2$ lies in $2\mathbb{Z}^\vee=\operatorname{Hom}_{\mathbb{Z}}(2\mathbb{Z},\mathbb{Z})$ but it is not in the image of φ.

We now introduce the notion of a lattice in a vector space.

Definition 5.8. Let A be an integral domain with fraction field K and let V be a K-vector space of finite dimension. A (full) *A-lattice in V* is a finitely generated A-submodule M of V that spans V as a K-vector space.

Remark 5.9. Some authors require A-lattices to be free A-modules. When $A=\mathbb{Z}$ (or any PID) this is not a restriction because M is necessarily torsion-free (it lies in a vector space) and any finitely generated torsion-free module over a PID is free (by the structure theorem for finitely generated modules over a PID). But when A is not a PID, finitely generated torsion-free A-modules will typically *not be free*. We do not want to exclude this case! In particular if L/K is an extension of number fields the ring of integers $\mathcal{O}_L$ will typically not be a free $\mathcal{O}_K$-module (even though it is a free $\mathbb{Z}$-module, as we shall shortly prove), but we still want to treat $\mathcal{O}_L$ as an $\mathcal{O}_K$-lattice in L (this will be important in later lectures when we define the *different ideal* $\mathcal{D}_{L/K}$).

Definition 5.10. Let A be a noetherian domain with fraction field K, and let V be a K-vector space of finite dimension with a perfect pairing $\langle\cdot,\cdot\rangle$. If M is an A-lattice in V, its *dual lattice* (with respect to the pairing $\langle\cdot,\cdot\rangle$) is the A-module

$$M^*:=\{x\in V:\langle x,m\rangle\in A\text{ for all }m\in M\}.$$

It is clear that M^* is an A-submodule of V, but it is not clear that it is an A-lattice in V (it must be finitely generated and to span V), nor is it obvious that it is isomorphic to the dual module $M^\vee$. In order to justify the term *dual lattice*, let us now prove both facts. We will need to use the hypothesis that A is noetherian, since in general the dual of a finitely generated A-module need not be finitely generated. Notice that $\langle\cdot,\cdot\rangle$ is a perfect pairing on the K-module V; it typically does not restrict to a perfect pairing on the A-module M (so Proposition 5.6 applies with $A = K$ and $M = V$, but typically not with $A = A$ and $M = M$).

Theorem 5.11. *Let A be a noetherian domain with fraction field K, let V be a K-vector space with a perfect pairing $\langle\cdot,\cdot\rangle$, and let M be an A-lattice in V. The dual lattice M^* is an A-lattice in V isomorphic to $M^\vee$.*

Proof. Let $\mathbf{e} := (e_1, \ldots, e_n)$ be a K-basis for V that lies in M, and let $(e'_1, \ldots, e'_n)$ be the unique K-basis for V given by Proposition 5.6 that satisfies $\langle e'_i, e_j\rangle = \delta_{ij}$.

To see that M^* spans V, if we write a finite set of generators for M in terms of the basis $\mathbf{e}$ with coefficients in K and let d be the product of all denominators that appear, then $d\mathbf{e}'$ lies in M^*, and it is a K-basis for V; here we are using the bilinearity of $\langle\cdot,\cdot\rangle$ and (crucially) the fact that $\langle e'_i, e_j\rangle = \delta_{ij} \in A$ (since the ring A contains 0 and 1).

We now show M^* is finitely generated. Let

$$N := \{a_1e_1 + \cdots + a_ne_n : a_1, \ldots, a_n \in A\} \simeq A^n$$

be the free A-submodule of M spanned by $\mathbf{e}$. The A-module N spans V and is finitely generated, so it is an A-lattice in V. The K-basis $\mathbf{e}'$ for V lies in N^*, since $\langle e'_i, e_j\rangle = \delta_{ij} \in A$, and we claim it is an A-basis for N^*. Given $x \in N^*$, write $x = \sum_{i=1}^n x_ie'_i$ with $x_i \in K$; then $\langle x, e_i\rangle = x_i\langle e'_i, e_i\rangle = x_i \in A$, so x lies in the A-span of $\mathbf{e}'$. It follows that N^* is a free A-module of rank n, and in particular, a finitely generated over a noetherian ring and therefore a noetherian module (all submodules finitely generated); see [1, Thm. 16.19]. From the definition of the dual lattice we have $N \subseteq M \Rightarrow M^* \subseteq N^*$, so M^* is a submodule of a noetherian module, hence finitely generated.

We now show $M^* \simeq M^\vee$. We have an obvious A-module homomorphism $\varphi\colon M^* \to M^\vee$ given by $x \mapsto (m \mapsto \langle x, m\rangle)$. We now proceed as in the proof of Theorem 5.2. Given $f \in M^\vee$, we can write any $m \in M$ as $\frac{1}{d}\sum_i a_ie_i$ for some $d, a_1, \ldots, a_n \in A$. We then have

$$f(m) = \frac{d}{d}f\left(\frac{1}{d}\sum_{i=1}^n a_ie_i\right) = \sum_{i=1}^n \frac{a_i}{d}f(e_i) = \sum_{i=1}^n \frac{a_i}{d}\sum_{j=1}^n \langle f(e_j)e'_j, e_i\rangle = \langle x, m\rangle \in A,$$

where $x = \sum_{i=1}^n f(e_i)e'_i \in M^*$ does not depend on m. The map $f \mapsto \sum_{i=1}^n f(e_i)e'_i$ is thus an A-module homomorphism $M^\vee \to M^*$ inverse to φ, so φ is an isomorphism. □

Corollary 5.12. *Let A be a noetherian domain with fraction field K, let V be a K-vector space with a perfect pairing $\langle\cdot,\cdot\rangle$, and let M be a free A-lattice in V with A-basis $(e_1, \ldots, e_n)$. The dual lattice M^* is a free A-lattice in V that has a unique A-basis $(e_1^*, \ldots, e_n^*)$ that satisfies $\langle e_i^*, e_j\rangle = \delta_{ij}$.*

Proof. This follows from the proof of Theorem 5.11 with $N = M$ and $e_i^* = e'_i$. □

You might wonder whether $M^{**} = M$ for an A-lattice M in a vector space V. This is false in general, but it is true when A is a Dedekind domain and we have a symmetric perfect pairing on V. To prove this we first show that the dual lattice respects localization.

Lemma 5.13. *Let A be a noetherian domain with fraction field K, let V be a K-vector space of finite dimension with a perfect pairing $\langle\cdot,\cdot\rangle$, let M be an A-lattice in V, and let S be a multiplicative subset of A. Then $S^{-1}M$ and $S^{-1}M^*$ are $S^{-1}A$-lattices with $(S^{-1}M)^* = S^{-1}M^*$ as $S^{-1}A$.*

Proof. It is clear that $S^{-1}M$ is an $S^{-1}M^*$ are both $S^{-1}A$-lattices: each contains a basis for V (since M and M^* do), and both are finitely generated as $S^{-1}A$-modules (since M and M^* are finitely generated as A-modules).

Let $m_1, \ldots m_n$ be A-module generators for M (and therefore $S^{-1}A$-module generators for $S^{-1}M$). If x is an element of $(S^{-1}M)^*$ then for each m_i we have $\langle x, m_i\rangle = a_i/s_i$ for some $a_i \in A$ and $s_i \in S$, and if we put $s = s_1 \cdots s_n$ then $\langle sx, m_i\rangle \in A$ for every m_i, hence for all $m \in M$; thus $sx \in M^*$ and $x \in S^{-1}M^*$. Conversely, if $x = y/s$ is an element of $S^{-1}M^*$ with $y \in M^*$ and $s \in S$, then $\langle y, m_i\rangle \in A$ for every m_i and $\langle x, m_i\rangle = \langle y, m_i\rangle/s \in S^{-1}A$ for every m_i, hence for all $m \in S^{-1}M$, and it follows that $x \in (S^{-1}M)^*$. □

Proposition 5.14. *Let A be a Dedekind domain with fraction field K, let V be a K-vector space of finite dimension with a symmetric perfect pairing $\langle\cdot,\cdot\rangle$, and let M be an A-lattice in V. Then $M^{**} = M$.*

Proof. By Proposition 2.7, it suffices to show $(M^{**})_{\mathfrak{p}} = M_{\mathfrak{p}}$ for each maximal ideal $\mathfrak{p}$ of A. By Lemma 5.13 we have $(M^{**})_{\mathfrak{p}} = M_{\mathfrak{p}}^{**}$, so it is enough to show that the proposition holds when A is replaced by one of its localizations $A_{\mathfrak{p}}$ (a DVR, since A is a Dedekind domain).

So let us assume A that is a DVR. Then A is a PID and M and M^* are both torsion-free modules over a PID, hence free A-modules. So let us choose an A-basis $(e_1, \ldots, e_n)$ for M, and let $(e_1^*, \ldots, e_n^*)$ be the unique dual A-basis for M^* that satisfies $\langle e_i^*, e_j\rangle = \delta_{ij}$ (given by Corollary 5.12). If we now let $(e_1^{**}, \ldots, e_n^{**})$ be the unique A-basis for M^{**} that satisfies $\langle e_i^{**}, e_j^*\rangle = \delta_{ij}$ and note that $\langle e_i, e_j^*\rangle = \delta_{ij}$ (since $\langle\cdot,\cdot\rangle$ is symmetric), by uniqueness, we must have $e_i^{**} = e_i$ for all i, and therefore $M^{**} = M$. □

5.2 Extensions of Dedekind domains

Let A be a Dedekind domain with fraction field K, let L/K be a finite extension, and let B be the integral closure of A in L. We wish to prove that B is a Dedekind domain, which we will do by showing that it is an A-lattice in L; this will imply, in particular, that B is finitely generated, which is really the only difficult thing to show. Let us first show that B spans L as a vector space (and in fact L is its fraction field).

Proposition 5.15. *Let A be a Dedekind domain with fraction field K, let L/K be a finite extension, and let B be the integral closure of A in L. Every element of L can be written as b/a with $a \in A$ and $b \in B$. In particular, B spans L as a K-vector space and L is the fraction field of B.*

Proof. Let $\alpha \in L$. By multiplying the minimal polynomial of α in $K[x]$ by the product of the denominators of its coefficients, we obtain an irreducible polynomial in $A[x]$:

$$g(x) = a_n x^n + a_{n-1}x^{n-1} + \cdots + a_1 x + a_0,$$

with $a_n \neq 0$, that has α as a root. We can make this polynomial monic by replacing x with x/a_n and multiplying through by a_n^{n-1} to obtain

$$a_n^{n-1} g(x/a_n) = x^n + a_{n-1}x^{n-1} + a_n a_{n-2} x^{n-2} \cdots + a_n^{n-2} a_1 x + a_n^{n-1} a_0.$$

This is a monic polynomial with coefficients in A that has $a_n\alpha \in L$ as a root. Therefore $a_n\alpha \in B$, since B is the integral closure of A in L, and $\alpha = b/a_n$ for some $b \in B$ and $a_n \in A$ as claimed. It follows that B generates L as a K-vector space (we have $\alpha = b \cdot \frac{1}{a_n}$ with $\frac{1}{a_n} \in K$), and $B \subseteq L \subseteq \operatorname{Frac} B$ implies $L = \operatorname{Frac} B$ (no smaller field can contain B). □

Proposition 5.16. *Let A be a Dedekind domain with fraction field K, let L/K be a finite extension, and let B be the integral closure of A in L. Then $\mathrm{N}_{L/K}(b) \in A$ and $\mathrm{T}_{L/K}(b) \in A$ for all $b \in B$.*

Proof. The minimal polynomial $f = \sum_{i=0}^{d} a_i x^i \in K[x]$ of b has coefficients in A, by Proposition 1.25, and it then follows from Proposition 4.44 that $N_{L/K}(b) = (-1)^{de} a_0^e \in A$ and $T_{L/K}(b) = -e a_{d-1} \in A$ (where $e = [L : K(b)] \in \mathbb{Z}$). □

Definition 5.17. Let L/K be a finite extension of fields. The *trace pairing* on L is the map $L \times L \to K$ defined by

$$\langle x, y \rangle_{L/K} := \mathrm{T}_{L/K}(xy).$$

Proposition 5.18. *Let L/K be a finite extension of fields. The trace pairing $\langle \cdot, \cdot \rangle_{L/K}$ is a symmetric bilinear pairing; it is perfect if and only if L/K is separable.*

Proof. Bilinearity follows from the K-linearity of the trace map $\mathrm{T}_{L/K}$, and symmetry is immediate. As proved on Problem Set 2, $T_{L/K}$ is the zero map if and only if L/K is inseparable. If $T_{L/K}$ is the zero map then the trace pairing is clearly not perfect, and otherwise we may pick $z \in L$ for which $\mathrm{T}_{L/K}(z) \neq 0$. Then for every $x \in L^\times$ we have $\langle x, z/x \rangle_{L/K} = \mathrm{T}_{L/K}(z) \neq 0$, which means that map $x \mapsto \langle x, y \rangle_{L/K}$ is not the zero map, so the kernel of the map $L \to L^\vee$ induced by $\langle \cdot, \cdot \rangle_{L/K}$ is trivial. Thus the pairing $\langle \cdot, \cdot \rangle_{L/K}$ is non-degenerate and therefore perfect, since K is a field. □

We now assume that L/K is separable. For the next several lectures we will be working in the following setting: A is a Dedekind domain with fraction field K, the extension L/K is finite separable, and B is the integral closure of A in L (which we will shortly prove is a Dedekind domain). As a convenient shorthand, we will write "assume $AKLB$" to indicate that we are using this setup.

Proposition 5.19. *Assume $AKLB$. Then B is an A-lattice in L under the trace pairing $\langle \cdot, \cdot \rangle_{L/K}$. In particular, B is finitely generated as an A-module.*

Proof. By Proposition 5.15, B spans L as a K-vector space, so it contains a basis $(e_1, \ldots, e_n)$ for L as a K-vector space. Let $M \subseteq B$ be the A-span of $(e_1, \ldots, e_n)$. Then M is an A-lattice in L contained in B, and it has a dual lattice M^* that contains the A-module

$$B^* := \{x \in L : \langle x, b \rangle \in A \text{ for all } b \in B\}.$$

Proposition 5.16 implies that $B \subseteq B^*$, and we thus have inclusions

$$M \subseteq B \subseteq B^* \subseteq M^*.$$

By Theorem 5.11, M^* is an A-lattice in L, hence finitely generated, hence noetherian. It follows that its A-submodule B is finitely generated and thus an A-lattice in L. □

Remark 5.20. When L/K is inseparable, B need not be finitely generated as an A-module, not even when A is a PID; see [2, Ex. 11, p. 205]. We used the separability hypothesis in order to get a perfect pairing, which plays a crucial role in the proof of Theorem 5.11.

Lemma 5.21. *Let B/A be an extension of domains with B integral over A, and let $\mathfrak{q}_0 \subsetneq \mathfrak{q}_1$ be primes of B. Then $\mathfrak{q}_0 \cap A \subsetneq \mathfrak{q}_1 \cap A$ and $\dim A \geq \dim B$.*

Proof. We first replace B with $B/\mathfrak{q}_0$ and replace A, $\mathfrak{q}_0$, and $\mathfrak{q}_1$ with their images in $B/\mathfrak{q}_0$ (the new B is integral over the new A, since the image of a monic polynomial in $A[x]$ is a monic polynomial in $(A/(\mathfrak{q}_0 \cap A))[x]$). Then $\mathfrak{q}_0 = (0)$ and $\mathfrak{q}_1$ is a nonzero prime ideal. Let $\alpha \in \mathfrak{q}_1$ be nonzero. Its minimal polynomial $x^n + a_{n-1}x^{n-1} + \cdots + a_0 = 0$ has coefficients in A and nonzero constant coefficient a_0 (otherwise divide by x). We have $a_0 = -a_1\alpha - \cdots - \alpha^n \in \mathfrak{q}_1$, thus $0 \neq a_0 \in \mathfrak{q}_1 \cap A$. So $\mathfrak{q}_1 \cap A$ is not the zero ideal and therefore properly contains $\mathfrak{q}_0 \cap A = \{0\}$. We can apply this result repeatedly to any chain of distinct prime ideals in B to get a corresponding chain of distinct prime ideals in A, thus $\dim A \geq \dim B$. □

Theorem 5.22. *Let A be a Dedekind domain with fraction field K, let L/K be a finite separable extension, and let B be the integral closure of A in L. Then B is a Dedekind domain.*

Proof. Recall that we defined a Dedekind domain as an integrally closed noetherian domain of dimension at most one. Let us verify that each of these conditions holds:

- B is an integrally closed domain (by definition);
- B is finitely generated over the noetherian ring A (by Prop. 5.19), hence noetherian;
- B has dimension at most 1, since $\dim B \leq \dim A \leq 1$, by Lemma 5.21.

Thus B is a Dedekind domain. □

Remark 5.23. Theorem 5.22 holds without the assumption that L/K is separable. This follows from the Krull-Akizuki Theorem, see [4, Thm. 11.7] or [3, §VII.2.5], which is used to prove that B is noetherian even when it is not finitely generated as an A-module.

Corollary 5.24. *The ring of integers of a number field is a Dedekind domain.*

5.3 Splitting primes in Dedekind extensions

We continue in the $AKLB$ setup, in which A is a Dedekind domain, K is its fraction field, L/K is a finite separable[1] extension, and B is the integral closure of A, which we now know is a Dedekind domain with fraction field L. As we proved in earlier lectures, every nonzero ideal in a Dedekind domain can be uniquely factored into prime ideals. Understanding the ideal structure of a Dedekind domain thus boils down to understanding its prime ideals.

In order to simplify the language, whenever we have a Dedekind domain A, by a *prime* of A (or of its fraction field K), we always mean a **nonzero** prime ideal; the prime elements of A are precisely those that generate nonzero principal prime ideals, so this generalizes the usual terminology. Note that 0 is (by definition) not prime, even though (0) is a prime ideal; thus when we refer to *a prime of A* we are specifically excluding the zero ideal.

If A has dimension zero then so does B there are no primes to consider, so we may as well assume $\dim A = 1$, in which case $\dim B = 1$ as well (if B is a field than so is $B \cap K = A$). Henceforth our $AKLB$ setup will include the assumption that $A \neq K$.

Given a prime $\mathfrak{p}$ of A, we can consider the ideal $\mathfrak{p}B$ it generates in B (its extension to B under the inclusion map). The ideal $\mathfrak{p}B$ need not be prime, but it can be uniquely factored

[1] Most of our proofs will not actually use the separability hypothesis (and even when they do, there may be another way to prove the same result, as with Theorem 5.22). In order to simplify the presentation we will use the separability assumption whenever it would be awkward not to. The cases we care most about (extensions of local and global fields) are going to be separable in any case.

into nonzero prime ideals in the Dedekind domain B. We thus have

$$\mathfrak{p}B = \prod_{\mathfrak{q}} \mathfrak{q}^{e_{\mathfrak{q}}},$$

where $\mathfrak{q}$ ranges over primes of B and the exponents $e_{\mathfrak{q}} \geq 0$ are zero for all but finitely many primes $\mathfrak{q}$. The primes $\mathfrak{q}$ for which $e_{\mathfrak{q}} > 0$ are said to *lie over* or *above* the prime ideal $\mathfrak{p}$. As an abuse of notation, we will often write $\mathfrak{q}|\mathfrak{p}$ to indicate this relationship (there is little risk of confusion, the prime ideal $\mathfrak{p}$ is maximal hence not divisible by any prime ideals of A other than itself).

Lemma 5.25. *Let A be a ring of dimension one contained in a Dedekind domain B. Let $\mathfrak{p}$ be a prime of A and let $\mathfrak{q}$ be a prime of B. Then $\mathfrak{q}|\mathfrak{p}$ if and only if $\mathfrak{q} \cap A = \mathfrak{p}$.*

Proof. If $\mathfrak{q}$ divides $\mathfrak{p}B$ then it contains $\mathfrak{p}B$ (to divide is to contain), and therefore $\mathfrak{q} \cap A$ contains $\mathfrak{p}B \cap A$ which contains $\mathfrak{p}$; the ideal $\mathfrak{p}$ is maximal and $\mathfrak{q} \cap A \neq A$ (since $1 \notin \mathfrak{q}$), so $\mathfrak{q} \cap A = \mathfrak{p}$. Conversely, if $\mathfrak{q} \cap A = \mathfrak{p}$ then $\mathfrak{q} = \mathfrak{q}B$ certainly contains $(\mathfrak{q} \cap A)B = \mathfrak{p}B$, and B is a Dedekind domain, so $\mathfrak{q}$ divides $\mathfrak{p}B$ (in a Dedekind domain to contain is to divide). □

Lemma 5.25 implies that contraction gives us a surjective map $\operatorname{Spec} B \to \operatorname{Spec} A$ defined by $\mathfrak{q} \mapsto \mathfrak{q} \cap A$; to see why it is surjective, note that $(0) \cap A = (0)$, and if $\mathfrak{p}$ is a nonzero element of $\operatorname{Spec} A$ then $\mathfrak{p}B$ is nonzero and not the unit ideal, and therefore divisible by at least one $\mathfrak{q} \in \operatorname{Spec} B$. The fibers of this map are finite; we use $\{\mathfrak{q}|\mathfrak{p}\}$ to denote the fiber above a prime $\mathfrak{p}$ of A.

The primes $\mathfrak{p}$ of A are all maximal ideals (since $\dim A = 1$), so each has an associated residue field $A/\mathfrak{p}$, and similarly for primes $\mathfrak{q}$ of B. If $\mathfrak{q}$ lies above $\mathfrak{p}$ then we may regard the residue field $B/\mathfrak{q}$ as a field extension of $\mathfrak{q}$: the kernel of the map $A \hookrightarrow B \to B/\mathfrak{q}$ is $\mathfrak{p} = A \cap \mathfrak{q}$, and the induced map $A/\mathfrak{p} = A/(\mathfrak{q} \cap A) \to B/\mathfrak{q}$ is a ring homomorphism of fields, hence injective.

Definition 5.26. Assume $AKLB$, and let $\mathfrak{p}$ be a prime of A. The exponent $e_{\mathfrak{q}}$ in the factorization $\mathfrak{p}B = \prod_{\mathfrak{q}|\mathfrak{p}} \mathfrak{q}^{e_{\mathfrak{q}}}$ is the *ramification index of* $\mathfrak{q}$, and the degree $f_{\mathfrak{q}} = [B/\mathfrak{q} : A/\mathfrak{p}]$ of the corresponding residue field extension is the *residue degree* (or *inertia degree*) of $\mathfrak{q}$. In situations where more than one extension of Dedekind domains is under consideration, we may write $e_{\mathfrak{q}/\mathfrak{p}}$ for $e_{\mathfrak{q}}$ and $f_{\mathfrak{q}/\mathfrak{p}}$ for $f_{\mathfrak{q}}$.

Lemma 5.27. *Let A be a Dedekind domain with fraction field K, let $M/L/K$ be a tower of finite separable extension, and let B and C be the integral closures of A in L and M respectively. Then C is the integral closure of B in M, and if $\mathfrak{r}$ is a prime of M lying above a prime $\mathfrak{q}$ of L lying above a prime $\mathfrak{p}$ of K then $e_{\mathfrak{r}/\mathfrak{p}} = e_{\mathfrak{r}/\mathfrak{q}}e_{\mathfrak{q}/\mathfrak{p}}$ and $f_{\mathfrak{r}/\mathfrak{p}} = f_{\mathfrak{r}/\mathfrak{q}}f_{\mathfrak{q}/\mathfrak{p}}$.*

Proof. It follows from Proposition 1.18 that the integral closure of B in M lies in C, and it contains C, since $A \subseteq B$. We thus have a tower of Dedekind extensions $C/B/A$. If $\mathfrak{r}|\mathfrak{q}|\mathfrak{p}$ then the factorization of $\mathfrak{p}C$ in C refines the factorization of $\mathfrak{p}B$ in B, so $e_{\mathfrak{r}/\mathfrak{p}} = e_{\mathfrak{r}/\mathfrak{q}}e_{\mathfrak{q}/\mathfrak{p}}$, and the residue field embedding $A/\mathfrak{p} \hookrightarrow C/\mathfrak{r}$ factors as $A/\mathfrak{p} \hookrightarrow B/\mathfrak{q} \hookrightarrow C/\mathfrak{r}$, so $f_{\mathfrak{r}/\mathfrak{p}} = f_{\mathfrak{r}/\mathfrak{q}}f_{\mathfrak{q}/\mathfrak{p}}$. □

Example 5.28. Let $A = \mathbb{Z}$, with $K = \operatorname{Frac} A = \mathbb{Q}$, and let $L = \mathbb{Q}(i)$ with $[L : K] = 2$. The prime $\mathfrak{p} = (5)$ factors in $B = \mathbb{Z}[i]$ into two distinct prime ideals:

$$5\mathbb{Z}[i] = (2 + i)(2 - i).$$

The prime $(2+i)$ has ramification index $e_{(2+i)} = 1$, and $e_{(2-i)} = 1$ as well. The residue field $\mathbb{Z}/(5)$ is isomorphic to the finite field $\mathbb{F}_5$, and we also have $\mathbb{Z}[i]/(2+i) \simeq \mathbb{F}_5$ (this can be determined by counting the $\mathbb{Z}[i]$-lattice points in a fundamental parallelogram of the sublattice $(2+i)$ in $\mathbb{Z}[i]$), so $f_{(2+i)} = 1$; we similarly have $f_{(2-i)} = 1$.

The $\mathfrak{p} = (7)$ remains prime in $B = \mathbb{Z}[i]$; its prime factorization is simply

$$7\mathbb{Z}[i] = (7),$$

where the (7) on the RHS denotes a principal ideal in B (this is clear from context). The ramification index of (7) is thus $e_{(7)} = 1$, but its residue field degree is $f_{(7)} = 2$, because $\mathbb{Z}/(7) \simeq \mathbb{F}_7$, but $\mathbb{Z}[i]/(7) \simeq \mathbb{F}_{49}$ has dimension 2 has an $\mathbb{F}_7$-vector space.

The prime $\mathfrak{p} = (2)$ factors as

$$(2) = (1+i)^2,$$

since $(1+i)^2 = (1+2i-1) = (2i) = (2)$ (note that i is a unit). You might be thinking that $(2) = (1+i)(1-i)$ factors into distinct primes, but note that $(1+i) = -i(1+i) = (1-i)$. Thus $e_{(1+i)} = 2$, and $f_{(1+i)} = 1$ because $\mathbb{Z}/(2) \simeq \mathbb{F}_2 \simeq \mathbb{Z}[i]/(1+i)$.

Let us now compute the sum $\sum_{\mathfrak{q}|\mathfrak{p}} e_{\mathfrak{q}} f_{\mathfrak{q}}$ for each of the primes $\mathfrak{p}$ we factored above:

$$\begin{aligned}
\sum_{q|(2)} e_{\mathfrak{q}} f_{\mathfrak{q}} &= e_{(1+i)} f_{(1+i)} = 2 \cdot 1 = 2,\\
\sum_{q|(5)} e_{\mathfrak{q}} f_{\mathfrak{q}} &= e_{(2+i)} f_{(2+i)} + e_{(2-i)} f_{(2-i)} = 1 \cdot 1 + 1 \cdot 1 = 2,\\
\sum_{q|(7)} e_{\mathfrak{q}} f_{\mathfrak{q}} &= e_{(7)} f_{(7)} = 2 \cdot 1 = 2.
\end{aligned}$$

In all three cases we obtain $2 = [\mathbb{Q}(i) : \mathbb{Q}]$; as we shall shortly prove, this is not an accident.

Example 5.29. Let $A = \mathbb{C}[x]$, with $K = \operatorname{Frac} A = \mathbb{C}(x)$, and let $L = \mathbb{C}(\sqrt{x}) = \operatorname{Frac} B$, where $B = \mathbb{C}[x,y]/(y^2 - x)$. Then $[L : K] = 2$. The prime $\mathfrak{p} = (x-4)$ factors in B into two distinct prime ideals:

$$(x-4) = (y^2-4) = (y+2)(y-2).$$

We thus have $e_{(y+2)} = 1$, and $f_{(y+2)} = [B/(y+2) : A/(x-4)] = [\mathbb{C} : \mathbb{C}] = 1$. Similarly, $e_{(y-2)} = 1$ and $f_{(y-2)} = 1$. The prime $\mathfrak{p} = x$ factors in B as

$$(x) = (y^2) = (y)^2,$$

and $e_{(y)} = 2$ and $f_{(y)} = 1$. As in the previous example, $\sum_{\mathfrak{q}|\mathfrak{p}} e_{\mathfrak{q}} f_{\mathfrak{q}} = [L : K]$ in both cases:

$$\begin{aligned}
\sum_{\mathfrak{q}|(x-4)} e_{\mathfrak{q}} f_{\mathfrak{q}} &= e_{(y+2)} f_{(y+2)} + e_{(y-2)} f_{(y-2)} = 1 \cdot 1 + 1 \cdot 1 = 2,\\
\sum_{\mathfrak{q}|(x)} e_{\mathfrak{q}} f_{\mathfrak{q}} &= e_{(y)} f_{(y)} = 2 \cdot 1 = 2.
\end{aligned}$$

Before proving that $\sum_{\mathfrak{q}|\mathfrak{p}} e_{\mathfrak{q}} f_{\mathfrak{q}} = [L : K]$ always holds, let us consider the quotient ring $B/\mathfrak{p}B$. The ring $B/\mathfrak{p}B$ is in general not a field extension of $A/\mathfrak{p}$ (because it is not necessarily a field), but it is an $A/\mathfrak{p}$-algebra, and in particular, an $A/\mathfrak{p}$-vector space. Let us determine its dimension.

Lemma 5.30. *Assume $AKLB$ and let $\mathfrak{p}$ be a prime of A. The dimension of $B/\mathfrak{p}B$ as an $A/\mathfrak{p}$-vector space is equal to the dimension of L as a K-vector space.*

Proof. Let $S = A - \mathfrak{p}$, let $A' = S^{-1}A = A_\mathfrak{p}$ and let $B' = S^{-1}B$ (note that S is closed under finite products, both as a subset of A and as a subset of B, so this makes sense). Then

$$A'/\mathfrak{p}A' = (S^{-1}A)/(\mathfrak{p}S^{-1}A) = A_\mathfrak{p}/(\mathfrak{p}A_\mathfrak{p}) \simeq A/\mathfrak{p},$$

and

$$B'/\mathfrak{p}B' = S^{-1}B/\mathfrak{p}S^{-1}B \simeq B/\mathfrak{p}B,$$

Thus if the lemma holds when $A = A_\mathfrak{p}$ is a DVR then it also holds for A; we now assume without loss of generality that A is a DVR.

By Proposition 5.19, B is finitely generated as an A module, and as an integral domain containing A, it must be torsion free. It follows from the structure theorem for finitely generated modules over a PID that B is free of finite rank over A. By Proposition 5.15, B spans L as a K-vector space, so any A-basis for B is a K-basis for L. Conversely, after clearing denominators any K-basis for L yields an A-basis for B. It follows that B has rank $[L : K]$ as a free A-module, and it follows that $B/\mathfrak{p}B$ has rank $[L : K]$ as a a free $A/\mathfrak{p}$-module (in fact an $A/\mathfrak{p}$ vector space).

To see this explicitly, let $\mathbf{e} = (e_1, \ldots, e_n)$ be an A-basis for B. Its image $\bar{\mathbf{e}}$ in $B/\mathfrak{p}B$ spans $B/\mathfrak{p}B$ as an $A/\mathfrak{p}$-module, since the reduction map is surjective. If $\sum_i a_i e_i \in \mathfrak{p}B$ for some $a_1, \ldots, a_n \in A$, then $\sum_i a_i e_i = \sum_j b_j p_j$ for some $b_j = \sum_i a_{ij} e_i \in B$ with $a_{ij} \in A$ and $p_j \in \mathfrak{p}$. Each $a_i = \sum a_{ij} p_j$ therefore lies in $\mathfrak{p}$, which shows that $\bar{\mathbf{e}}$ is $A/\mathfrak{p}$-linearly independent. It follows that $\bar{\mathbf{e}}$ is an $A/\mathfrak{p}$-basis for $B/\mathfrak{p}B$, which therefore has rank $n = [L : K]$. □

Theorem 5.31. *Assume $AKLB$. For each prime $\mathfrak{p}$ of A we have*

$$\sum_{\mathfrak{q}|\mathfrak{p}} e_\mathfrak{q} f_\mathfrak{q} = [L : K].$$

Proof. We have

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

Applying the previous proposition gives

$$\begin{aligned}[L : K] &= [B/\mathfrak{p}B : A/\mathfrak{p}]\\ &= \sum_{\mathfrak{q}|\mathfrak{p}} [B/\mathfrak{q}^{e_q} : A/\mathfrak{p}]\\ &= \sum_{\mathfrak{q}|\mathfrak{p}} e_\mathfrak{q}[B/\mathfrak{q} : A/\mathfrak{p}]\\ &= \sum_{\mathfrak{q}|\mathfrak{p}} e_\mathfrak{q} f_\mathfrak{q}.\end{aligned}$$

The second equality comes from the Chinese Remainder Theorem, and the third uses the fact that $B/\mathfrak{q}^{e_\mathfrak{q}}$ has dimension $e_\mathfrak{q}$ as a $B/\mathfrak{q}$-vector space. Indeed, we have

$$\mathfrak{q}^{e_\mathfrak{q}} = \{x \in B : v_\mathfrak{q}(x) \le e_\mathfrak{q}\},$$

and if $\pi \in \mathfrak{q}$ is a uniformizer for $B_\mathfrak{q}$ (a generator $\mathfrak{q}B_\mathfrak{q}$ that we can force to lie in $\mathfrak{q}$ by clearing denominators), the images of $(\pi^0, \pi^1, \ldots, \pi^{e_\mathfrak{q}-1})$ in $B/\mathfrak{q}^{e_\mathfrak{q}}$ are a $B/\mathfrak{q}$-basis for $B/\mathfrak{q}^{e_\mathfrak{q}}$. □

For each prime $\mathfrak{p}$ of A, let $g_{\mathfrak{p}} := \#\{\mathfrak{q}|\mathfrak{p}\}$ denote the cardinality of the fiber above $\mathfrak{p}$.

Corollary 5.32. *Assume $AKLB$ and let $\mathfrak{p}$ be a prime of A. Then $g_{\mathfrak{p}}$ is an integer in the interval $[1, n]$, where $n = [L : K]$, as are $e_{\mathfrak{q}}$ and $f_{\mathfrak{q}}$ for each $\mathfrak{q}|\mathfrak{p}$.*

We now define some standard terminology that we may use in the $AKLB$ setting to describe how a prime $\mathfrak{p}$ of K splits in L (that is, for a nonzero prime ideal $\mathfrak{p}$ of A, how the ideal $\mathfrak{p}B$ factors into nonzero prime ideals $\mathfrak{q}$ of B).

Definition 5.33. Assume $AKLB$, let $\mathfrak{p}$ be a prime of A.

- L/K is *totally ramified at* $\mathfrak{q}$ if $e_{\mathfrak{q}} = [L : K]$ (equivalently, $f_{\mathfrak{q}} = 1 = g_{\mathfrak{p}} = 1$).
- L/K is *unramified at* $\mathfrak{q}$ if $e_{\mathfrak{q}} = 1$ **and** $B/\mathfrak{q}$ is a separable extension of $A/\mathfrak{p}$.
- L/K is *unramified above* $\mathfrak{p}$ if it is unramified at all $\mathfrak{q}|\mathfrak{p}$, equivalently, if $B/\mathfrak{p}B$ is a finite étale algebra over $A/\mathfrak{p}$.

When L/K is unramified above $\mathfrak{p}$ we say that

- $\mathfrak{p}$ *remains inert in* L if $\mathfrak{p}B$ is prime (equivalently, $e_{\mathfrak{q}} = g_{\mathfrak{p}} = 1$, and $f_{\mathfrak{q}} = [L : k]$).
- $\mathfrak{p}$ *splits completely in* L if $g_{\mathfrak{p}} = [L : K]$ (equivalently, $e_{\mathfrak{q}} = f_{\mathfrak{q}} = 1$ for all $\mathfrak{q}|\mathfrak{p}$).

References

[1] Allen Altman and Steven Kleiman, *A term of commutative algebra*, Worldwide Center of Mathematics, 2013.

[2] Z.I. Borevich and I.R. Shafarevich, *Number theory*, Academic Press, 1966.

[3] N. Bourbaki, *Commutative Algebra: Chapters 1–7*, Springer, 1989.

[4] H. Matsumura, *Commutative ring theory*, Cambridge University Press, 1986.

18.785 Number Theory I
Fall 2016