2 Localization and Dedekind domains

After a brief review of some commutative algebra background on localizations, in this lecture we begin our study of Dedekind domains, which are commutative rings that play a key role in algebraic number theory and arithmetic geometry (named after Richard Dedekind).

2.1 Localization of rings

Let A be a commutative ring (unital, as always), and let S be a multiplicative subset of A; this means that S is closed under finite products (including the empty product, so $1 \in S$), and S does not contain zero. The *localization* of A with respect to S is a ring $S^{-1}A$ equipped with a ring homomorphism $\iota\colon A \to S^{-1}A$ that maps S to $A^\times$ and satisfies the following universal property: if $\varphi\colon A \to B$ is a ring homomorphism with $\varphi(S) \subseteq B^\times$ then there is a unique ring homomorphism $S^{-1}A \to B$ that makes the following diagram commute:

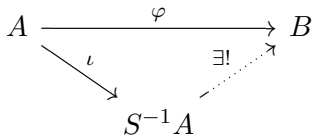

One says that φ *factors uniquely through* $S^{-1}A$ (or more precisely, through ι). As usual with universal properties, this guarantees that $S^{-1}A$ is unique (hence well-defined) if it exists. To prove existence we construct $S^{-1}A$ as the quotient of $A \times S$ modulo the equivalence relation

$$(a, s) \sim (b, t) \Leftrightarrow \exists u \in S \text{ such that } (at - bs)u = 0. \tag{1}$$

We then use a/s to denote the equivalence class of (a, s) and define $\iota(a) := a/1$; one can easily verify that $S^{-1}A$ is a ring with additive identity $0/1$ and multiplicative identity $1/1$ and that $\iota : A \to S^{-1}A$ is a ring homomorphism. If s is invertible in a we can view a/s either as the element as^{-1} of A or the equivalence class of (a, s) in $S^{-1}A$; we have $(a, s) \sim (a/s, 1)$, since $(a \cdot 1 - a/s \cdot s) \cdot 1 = 0$, so there is no real risk of confusion. For $s \in S$ we have $\iota(s)^{-1} = 1/s$, since $(s/1)(1/s) = s/s = 1/1 = 1$, thus $\iota(S) \subseteq (S^{-1}A)^\times$ as required.

If $\varphi\colon A \to B$ is a ring homomorphism with $\varphi(S) \subseteq B^{-1}$, then $\varphi = \pi \circ \iota$, where π is defined by $\pi(a/s) := \varphi(a)\varphi(s)^{-1}$. If $\pi\colon S^{-1}A \to B$ is any ring homomorphism that satisfies $\varphi = \pi \circ \iota$, then $\varphi(a)\varphi(s)^{-1} = \pi(\iota(a))\pi(\iota(s))^{-1} = \pi(\iota(a)\iota(s)^{-1}) = \pi((a/1)(1/s)) = \pi(a/s)$, so π is unique.

In the case of interest to us, A is actually an integral domain, in which case $(a, s) \sim (b, t)$ if and only if $at - bs = 0$ (we can always take $u = 1$ in the equivalence relation (1) above), and we can then identify $S^{-1}A$ with a subring of the fraction field of A (which we note is the localization $A^{-1}A$ of A with respect to itself), and if T is a multiplicative set of the integral domain A that contains S, then $S^{-1}A \subseteq T^{-1}A$.

Moreover, when A is an integral domain, the map $\iota\colon S \to S^{-1}A$ is injective and we may identify A with its image $\iota(A) \subseteq S^{-1}A$ (in general, ι is injective if and only if S contains no zero divisors). Thus when A is an integral domain we may (and will) view $S^{-1}A$ as an intermediate ring that lies between A and its fraction field.

2.2 Ideals in localizations of rings

If $\varphi\colon A \to B$ is a ring homomorphism and $\mathfrak{b}$ is a B-ideal, then $\varphi^{-1}(\mathfrak{b})$ is an A-ideal called the *contraction* of $\mathfrak{b}$ (to A) and sometimes denoted $\mathfrak{b}^c$; when A is a subring of B and φ is

the inclusion map we simply have $\mathfrak{b}^c = \mathfrak{b} \cap A$. If $\mathfrak{a}$ is an A-ideal then $\varphi(\mathfrak{a})$ is in general not a B-ideal; but the B-ideal generated by $\varphi(\mathfrak{a})$ is called the *extension* of $\mathfrak{a}$ (to B) and sometimes denoted $\mathfrak{a}^e$.

In the case of interest to us, A is an integral domain, $B = S^{-1}A$ is the location of A with respect to some multiplicative set S, and $\varphi = \iota$ is injective, so we view A as a subring of B. We then have

$$\mathfrak{a}^e = \mathfrak{a}B := \{ab : a \in \mathfrak{a}, b \in B\}. \tag{2}$$

We clearly have $\mathfrak{a} \subseteq \varphi^{-1}((\varphi(\mathfrak{a}))) = \mathfrak{a}^{ec}$ and $\mathfrak{b}^{ce} = (\varphi(\varphi^{-1}(\mathfrak{b}))) \subseteq \mathfrak{b}$; one might ask whether these inclusions are equalities. In general the first is not: if $B = S^{-1}A$ and $\mathfrak{a} \cap S \neq \emptyset$ then $\mathfrak{a}^e = \mathfrak{a}B = B$ and $\mathfrak{a}^{ec} = B \cap A$ are both unit ideals, but we may still have $\mathfrak{a} \subsetneq A$. However when $B = S^{-1}A$ the second inclusion is always an equality; see [1, Prop. 11.19] or [2, Prop. 3.11] for a short proof. We also note the following theorem.

Theorem 2.1. *The map $\mathfrak{q} \mapsto \mathfrak{q} \cap A$ defines a bijection from the set of prime ideals of $S^{-1}A$ and the set of prime ideals of A that do not intersect S. The inverse map is $\mathfrak{p} \mapsto \mathfrak{p}S^{-1}A$.*

Proof. See [1, Cor. 11.20] or [2, Prop. 3.11.iv]. □

Remark 2.2. An immediate consequence of (2) is that if $a_1, \dots, a_n \in A$ generate $\mathfrak{a}$ as an A-ideal, then they also generate $\mathfrak{a}^e = \mathfrak{a}B$ as a B-ideal. As noted above, when $B = S^{-1}A$ we have $\mathfrak{b} = \mathfrak{b}^{ce}$, so every B-ideal is of the form $\mathfrak{a}^e$ (take $\mathfrak{a} = \mathfrak{b}^c$). It follows that if A is noetherian then so are all its localizations, and if A is a PID then so are all of its localizations.

An important special case of localization occurs when $\mathfrak{p}$ is a prime ideal in an integral domain A, and $S = A - \mathfrak{p}$ (the complement of the set $\mathfrak{p}$ in the set A). In this case it is customary to denote $S^{-1}A$ by

$$A_{\mathfrak{p}} := \{a/b : a \in A, b \notin \mathfrak{p}\}/\sim, \tag{3}$$

and call it the *localization of A at $\mathfrak{p}$*. The prime ideals of $A_{\mathfrak{p}}$ are then in bijection with the prime ideals of A that lie in $\mathfrak{p}$. It follows that $\mathfrak{p}A_{\mathfrak{p}}$ is the unique maximal ideal of $A_{\mathfrak{p}}$ and $A_{\mathfrak{p}}$ is therefore a local ring (whence the term *localization*).

Warning 2.3. The notation in (3) makes it tempting to assume that if a/b is an element of $\operatorname{Frac} A$, then $a/b \in A_{\mathfrak{p}}$ if and only if $b \notin \mathfrak{p}$. This is not necessarily true! As an element of $\operatorname{Frac} A$, the notation "a/b" represents an equivalence class; if $a/b = a'/b'$ with $b' \notin A_{\mathfrak{p}}$, then in fact $a/b = a'/b' \in A_{\mathfrak{p}}$. As a trivial example, take $A = \mathbb{Z}$, $\mathfrak{p} = (3)$, $a/b = 9/3$ and $a'/b' = 3/1$. You may object that we should write a/b in lowest terms, but when A is not a unique factorization domain it is not clear what this means.

Example 2.4. For a field k, let $A = k[x]$ and $\mathfrak{p} = (x - 2)$. Then

$$A_{\mathfrak{p}} = \{f \in k(x) : f \text{ is defined at } 2\}.$$

The ring A is a PID, so $A_{\mathfrak{p}}$ is a PID with a unique nonzero maximal ideal (the ideal $\mathfrak{p}A_{\mathfrak{p}}$), hence a DVR. Its maximal ideal is

$$\mathfrak{p}A_{\mathfrak{p}} = \{f \in k(x) : f(2) = 0\}.$$

The valuation on the field $k(x) = \operatorname{Frac} A$ corresponding to the valuation ring $A_{\mathfrak{p}}$ measures the order of vanishing of functions $f \in k(x)$ at 2. The residue field is $A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}} \simeq k$, and the quotient map $A_{\mathfrak{p}} \twoheadrightarrow A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}}$ sends f to $f(2)$.

Example 2.5. Let $p \in \mathbb{Z}$ be a prime. Then $\mathbb{Z}_{(p)} = \{a/b : a, b \in \mathbb{Z}, p \nmid b\}$. As in the previous example, $\mathbb{Z}$ is a PID and $\mathbb{Z}_{(p)}$ is a DVR; the valuation on $\mathbb{Q}$ is the p-adic valuation.

2.3 Localization of modules

The concept of localization generalizes immediately to modules. As above, let A be a commutative ring, let S a multiplicative subset of A, and let M be an A-module. The localization $S^{-1}M$ of M with respect to S is an $S^{-1}A$-module equipped with an A-module homomorphism $\iota\colon M \to S^{-1}M$ with the universal property that if N is an $S^{-1}A$-module and $\varphi\colon M \to N$ is an A-module homomorphism, then φ factors uniquely through $S^{-1}M$ (via ι). Note that in this definition we are viewing $S^{-1}A$-modules as A-modules via the canonical homomorphism $A \to S^{-1}A$ that is part of the definition of $S^{-1}A$. Our definition of $S^{-1}M$ reduces to the definition of $S^{-1}A$ in the case $M = A$.

The explicit construction of $S^{-1}M$ is exactly the same as $S^{-1}A$, one takes the quotient of $M \times S$ module the same equivalence relation as in (1):

$$(a, s) \sim (b, t) \Leftrightarrow \exists u \in S \text{ such that } (at - bs)u = 0,$$

where a and b now denote elements of M, and $\iota(a) \coloneqq a/1$ as before. Alternatively, one can define $S^{-1}M := M \otimes_A S^{-1}A$ (see [2, Prop. 3.5] for a proof that this is equivalent). In other words, $S^{-1}M$ is the *base change* of M from A to $S^{-1}A$; we will discuss base change more generally in later lectures.

The map $\iota\colon M \to S^{-1}M$ is injective if and only if the map $M \xrightarrow{\times s} M$ is injective for every $s \in S$. This is a strong condition that does not hold in general, even when A is an integral domain (the annihilator of M may be non-trivial), but it applies to all the cases we care about. In particular, if A lies in a field K (in which case A must be an integral domain whose fraction field lies in K) and M is an A-module that is contained in a K-vector space. In this setting multiplication by any nonzero $s \in A$ is injective and we can view M as an A-submodule of any of its localizations $S^{-1}M$.

We will mostly be interested in the case $S = A - \mathfrak{p}$, where $\mathfrak{p}$ is a prime ideal of A, in which case we write $M_\mathfrak{p}$ for $S^{-1}M$, just as we write $A_\mathfrak{p}$ for $S^{-1}A$.

Proposition 2.6. *Let A be a subring of a field K, and let M be an A-module contained in a K-vector space V (equivalently, for which the map $M \to M \otimes_A K$ is injective).*[1] *Then*

$$M = \bigcap_\mathfrak{m} M_\mathfrak{m} = \bigcap_\mathfrak{p} M_\mathfrak{p},$$

where $\mathfrak{m}$ ranges over the maximal ideals of A, $\mathfrak{p}$ ranges over the prime ideals of A, and the intersections take place in V.

Proof. The fact that $M \subseteq \bigcap_\mathfrak{m} M_\mathfrak{m}$ is immediate. Now suppose $x \in \bigcap_\mathfrak{m} M_\mathfrak{m}$ and consider the A-ideal $\mathfrak{a} \coloneqq \{a \in A : ax \in M\}$. For each maximal ideal $\mathfrak{m}$ we can write $x = m/s$ for some $m \in M$ and $s \in A - \mathfrak{m}$; we then have $sx \in M$ and $s \in \mathfrak{a}$, but $s \notin \mathfrak{m}$ to $\mathfrak{a} \not\subseteq \mathfrak{m}$. It follows that $\mathfrak{a}$ and must be the unit ideal, so $1 \in \mathfrak{a}$ and $x = 1 \cdot x \in M$; thus $\cap_\mathfrak{m} M_\mathfrak{m} \subseteq M$.

We now note that each $M_\mathfrak{p}$ contains some $M_\mathfrak{m}$ (since each $\mathfrak{p}$ is contained in some $\mathfrak{m}$), and every maximal ideal is prime, so $\cap_\mathfrak{m} M_\mathfrak{m} = \cap_\mathfrak{p} M_\mathfrak{p}$. □

An important special case of this proposition occurs when $K = \operatorname{Frac} A$ and $V = K$, in which case M is an A-submodule of K. Every I of A is an A-submodule of K, and each can be localized as above. The localization of I (as an A-module) at a prime ideal $\mathfrak{p}$ of A is

[1] The image is a tensor product of A-modules that is also a K-vector space. We need the natural map to be injective in order to embed M in it. Note that V necessarily contains a subspace isomorphic to $M \otimes_A K$.

the same thing as the extension of I (as an A-ideal) to the localization of A at $\mathfrak{p}$. In other words,

$$I_{\mathfrak{p}} = IA_{\mathfrak{p}}.$$

We also have the following corollary of Proposition 2.6.

Corollary 2.7. *Let A be an integral domain. Every ideal I of A (including $I = A$) is equal to the intersection of its localizations at the maximal ideals of A (and also to the intersection of its localizations at the prime ideals of A).*

Example 2.8. If $A = \mathbb{Z}$ then $\mathbb{Z} = \bigcap_p \mathbb{Z}_{(p)}$ in $\mathbb{Q}$.

Proposition 2.6 and Corollary 2.7 are powerful tools, because they allow us work in local rings (rings with just one maximal ideal), which often simplifies matters considerably. For example, to prove that an ideal I in an integral domain A satisfies a certain property, it is enough to show that this property holds for all its localizations $I_{\mathfrak{p}}$ at prime ideals $\mathfrak{p}$ and is preserved under intersections. We now want to show that if when A satisfies some further assumptions its localizations become even easier to work with.

2.4 Dedekind domains

Proposition 2.9. *Let A be a noetherian domain. The following are equivalent:*

(i) *For every nonzero prime ideal $\mathfrak{p} \subset A$ the local ring $A_{\mathfrak{p}}$ is a DVR.*

(ii) *The ring A is integrally closed and* $\dim A \leq 1$.

Proof. If A is a field then (i) and (ii) both hold, so let us assume that A is not a field, and put $K := \operatorname{Frac} A$. We first show that (i) implies (ii). Recall that $\dim A$ is the supremum of the length of all chains of prime ideals. It follows from Theorem 2.1 that every chain of prime ideals $(0) \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_n$ extends to a corresponding chain in $A_{\mathfrak{p}_n}$ of the same length; conversely, every chain in $A_{\mathfrak{p}}$ contracts to a chain in A of the same length. Thus

$$\dim A = \sup\{\dim A_{\mathfrak{p}} : \mathfrak{p} \in \operatorname{Spec} A\} = 1,$$

since every $A_{\mathfrak{p}}$ is either a DVR ($\mathfrak{p} \neq (0)$), in which case $\dim A_{\mathfrak{p}} = 1$, or a field ($\mathfrak{p} = (0)$), in which case $\dim A_{\mathfrak{p}} = 0$. Any $a \in K$ that is integral over A is integral over every $A_{\mathfrak{p}}$ (since they all contain A), and the $A_{\mathfrak{p}}$ are integrally closed, since they are DVRs. So $a \in \bigcap_{\mathfrak{p}} A_{\mathfrak{p}} = A$, and therefore A is integrally closed, which shows (ii).

To show that (ii) implies (i), we first show that the following properties are all inherited by localizations of a ring: (1) no zero divisors, (2) noetherian, (3) dimension at most one, (4) integrally closed. (1) is obvious, (2) was noted in Remark 2.2, and (3) follows from the fact that every chain of prime ideals in $A_{\mathfrak{p}}$ extends to a chain of prime ideals in A of the same length, so $\dim A_{\mathfrak{p}} \leq \dim A$. To show (4), suppose $x \in K$ is integral over $A_{\mathfrak{p}}$. Then

$$x^n + \frac{a_{n-1}}{s_{n-1}}x^{n-1} + \cdots + \frac{a_1}{s_1}x + \frac{a_0}{s_0} = 0$$

for some $a_0, \ldots, a_{n-1} \in A$ and $s_0, \ldots, s_{n-1} \in A - \mathfrak{p}$. Multiplying both sides by s^n, where $s = s_0 \cdots s_{n-1} \in S$, shows that sx is integral over A, hence an element of A, since A is integrally closed. But then $sx/s = x$ is an element of $A_{\mathfrak{p}}$, so $A_{\mathfrak{p}}$ is integrally closed as claimed.

Thus (ii) implies that every $A_\mathfrak{p}$ is an integrally closed noetherian local domain of dimension at most 1, and for $\mathfrak{p} \neq (0)$ we must have $\dim A_\mathfrak{p} = 1$. Thus for every nonzero prime ideal $\mathfrak{p}$, the ring $A_\mathfrak{p}$ is an integrally closed noetherian local domain of dimension 1, and therefore a DVR, by Theorem 1.14. □

Definition 2.10. A noetherian domain satisfying either of the equivalent properties of Proposition 2.9 is called a *Dedekind domain*.

Corollary 2.11. *Every PID is a Dedekind domain. In particular, $\mathbb{Z}$ is a Dedekind domain, as is $k[x]$ for any field k.*

Remark 2.12. Every PID is both a UFD and a Dedekind domain. Not every UFD is a Dedekind domain (consider $k[x, y]$, for any field k), and not every Dedekind domain is a UFD (consider $\mathbb{Z}[\sqrt{-13}]$, in which $(1 + \sqrt{-13})(1 - \sqrt{-13}) = 2 \cdot 7 = 14$). However (as we shall see), every ring that is both a UFD and a Dedekind domain is a PID.

One of our first goals in this course is to prove that the ring of integers of a number field is a Dedekind domain. More generally, we will prove that if A is a Dedekind domain and L is a finite separable extension of its fraction field, then the integral closure of A in L is a Dedekind domain. The two main cases of interest to us are when $A = \mathbb{Z}$, in which case L is a number field, and $A = \mathbb{F}_q[t]$ for some finite field $\mathbb{F}_q$, in which case L is a global function field. The finite extension of $\mathbb{Q}$ and $\mathbb{F}_q(t)$ (number fields and global function fields) are collectively known as *global fields*, for reasons that will become clear in later lectures.

Remark 2.13. Unlike $\mathbb{Q}$, not every finite extension of $\mathbb{F}_q(t)$ is separable. But every finite extension K of $\mathbb{F}_q(t)$ contains a subfield isomorphic to $\mathbb{F}_q(t)$ over which it is separable; one can always pick a *separating element* $s \in K$ that is transcendental over $\mathbb{F}_q$ such that $K/\mathbb{F}_q(s)$ is separable. More generally, by a theorem of Schmidt, every finitely generated extension of a perfect field k is *separably generated*, meaning that it is a separable algebraic extension of a purely transcendental extension of k; see [3, Thm. 7.20] for a proof.

2.5 Fractional ideals

Throughout this subsection, A is a noetherian domain (not necessarily a Dedekind domain) and K is its fraction field.

Definition 2.14. A *fractional ideal* of a noetherian domain A is a finitely generated A-submodule of its fraction field.

Fractional ideals generalize the notion of an ideal: when A is noetherian the ideals of A are precisely the finitely generated A-submodules of A, and when A is also a domain we can extend this notion to its fraction field. Every ideal of A is also a fractional ideal of A, but fractional ideals are typically not ideals because they need not be contained in A. Some authors use the term *integral ideal* to distinguish the fractional ideals that lie in A (and are thus ideals) but we will not use this terminology.

Lemma 2.15. *Let A be a noetherian domain with fraction field K, and let $I \subseteq K$ be an A-module. Then I is finitely generated if and only if $aI \subseteq A$ for some nonzero $a \in A$.*

Proof. For the forward implication, if $r_1/s_1, \dots, r_n/s_n$ generate I as an A-module, then $aI \subseteq A$ for $a = s_1 \cdots s_n$. Conversely, if $aI \subseteq A$, then aI is an ideal, hence finitely generated (since A is noetherian), and if $a_1, \dots, a_n$ generate aI then $a_1/a, \dots, a_n/a$ generate I. □

Remark 2.16. Lemma 2.15 gives an alternative definition of fractional ideals that can be extended to domains that are not necessarily noetherian; they are A-submodules I of K for which there exists a nonzero $r \in A$ such that $rI \subseteq A$. When A is noetherian this coincides with our definition above.

Corollary 2.17. *Every fractional ideal of A can be written in the form $\frac{1}{a}I$, for some nonzero $a \in A$ and ideal I.*

Example 2.18. The set $I = \frac{1}{2}\mathbb{Z} = \{\frac{n}{2} : n \in \mathbb{Z}\}$ is a fractional ideal of $\mathbb{Z}$. As a $\mathbb{Z}$-module it is generated by $1/2 \in \mathbb{Q}$, and we have $2I \subseteq \mathbb{Z}$.

Definition 2.19. A fractional ideal of A is *principal* if it is generated by one element, that is, it has the form xA for some $x \in K$. We will also use the notation $(x) := xA$ to denote the principal fractional ideal generated by $x \in K$.

As with ideals, we can add and multiply fractional ideals:

$$I + J := (i + j : i \in I, j \in J), \qquad IJ := (ij : i \in I, j \in J).$$

Here the notation (S) means the A-module generated by $S \subseteq K$. As with ideals, we actually have $I + J = \{i + j : i \in I, j \in J\}$, but the ideal IJ is typically not the same as set $\{ij : i \in I, j \in J\}$, it consists of all finite sums of elements in this set. We also have a new operation, corresponding to division. For any fractional ideals I, J with J nonzero, the set

$$(I : J) := \{x \in K : xJ \subseteq I\}$$

is called a *colon ideal*. Some texts refer to $(I : J)$ as the *ideal quotient* of I by J, but note that it is **not** a quotient of A-modules; for example, $(\mathbb{Z} : \mathbb{Z}) = \mathbb{Z}$ but $\mathbb{Z}/\mathbb{Z} = \{0\}$.

We do not assume $I \subseteq J$ (or $J \subseteq I$), the definition makes sense for any fractional ideals I and J (including $J = \{0\}$, in which case $(I : K) = K$). If $I = (x)$ and $J = (y)$ are principal fractional ideals then $(I : J) = (x/y)$, so colon ideals can be viewed as a generalization of division in $K^\times$.

The colon ideal $(I : J)$ is an A-submodule of K, and it is finitely generated, hence a fractional ideal. This is easy to see when $I, J \subseteq A$: let j be any nonzero element of $J \subseteq A$ and note that $j(I : J) \subseteq I \subseteq A$, so $(I : J)$ is finitely generated, by Lemma 2.15. More generally, choose a and b so that $aI \subseteq A$ and $bJ \subseteq A$. Then $(I : J) = (abI : abJ)$ with $abI, abJ \subseteq A$ and we may apply the same argument.

References

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

[2] Michael Atiyah and Ian MacDonald, *Introduction to commutative algebra*, Addison–Wesley, 1969.

[3] Anthony W. Knapp, *Advanced Algebra*, Digital Second Edition, 2016.

18.785 Number Theory I
Fall 2017