The Other Dedekind Domains: Function Fields over $\mathbb{F}_q$

Alison Miller

September 18, 2005

1 Why fields of functions over algebraic curves?

The study of algebraic number fields — finite extensions of $\mathbb{Q}$ — and their rings of integers yields many fascinating results, theories, and open question. These objects form such fertile ground of study in part because the ring of integers $\mathcal{O}_K$ of an algebraic number field has the additional structure of Dedekind domain, allowing one to work in terms of ideal factorization. All such number fields $K/\mathbb{Q}$ must have characteristic 0, so the question arises: are there fields of nonzero characteristic p that give rise to a similar theory? Finite fields $\mathbb{F}_q$ give us nothing to study. Indeed, any subring of a finite field is also a field, with no interesting ideal structure. However, if we add a transcendental element, and consider fields of $\mathbb{F}_q(t)$ and its finite separable extensions, we will find we arrive at a theory very similar to that of algebraic number fields. As a bonus, these fields can also be interpreted as function fields of algebraic curves. Our exposition will first briefly develop this theory of algebraic curves in order to use it as a tool to prove that function fields of algebraic curves over finite fields are Dedekind domains like the algebraic number fields they resemble.

2 Theory of Function Fields over $\mathbb{F}_q$

2.1 Definitions and other Gory Details

Let K be a field: for the purposes of this paper K will always be either a finite field of the form $\mathbb{F}_q$ for some prime power q or K will be the algebraic closure $\bar{F}_q$. To understand algebraic curves, we will introduce a few notions from algebraic geometry. We will consider algebraic curves over K as subsets of either the affine space $\mathbb{A}^n(K) = K^n$ or $\mathbb{P}^n(K) = (K^{n+1} - \{0\})/\text{scalar multiplication by elements of } K$.

Definition. An *(affine) algebraic set* over K is a subset V of $\mathbb{A}^n(K)$ such that there exist polynomials $a_1, a_2, \ldots, a_m \in K[x_1, \ldots, x_n]$ such that $V = \{(x_1, \ldots, x_n \in \mathbb{A}^n(K) \mid a_i(x_1, \ldots x_n) = 0 \text{ for } i = 1, \ldots m\}$. Note that if

$I \subset K[x_1, \ldots, x_n]$ is the ideal generated by $a_1, \ldots a_n$ this condition is equivalent to the stronger condition $a(x_1, \ldots, x_n) = 0$ for any $a \in I$. In such a case we write $V = V(I)$. In the other direction, starting with an algebraic set $V \subset \mathbb{A}^n(K)$, we define $I(V)$ to be the ideal of all polynomials that vanish on V. A *projective algebraic set* is defined similarly as a subset of the projective space $\mathbb{P}^n(K)$ (defined as usual to be $\mathbb{A}^{n+1}(K)/$(multiplication by nonzero elements of K)). In this case $a_1, \ldots, a_m$ are required to be homogeneous polynomials in $K[x_0, \ldots, x_n]$. We can construct $V(I)$ and $I(V)$ exactly as above, where our ideal I will now be a homogeneous ideal (that is, one generated by homogeneous polynomials in $x_0, \ldots, x_n$).

Definition. An *(affine) algebraic curve* is defined as a polynomial $F \in K(x_1, x_2)$ modulo multiplication by nonzero elements of K (that is, polynomials that are scalar multiples of each other correspond to the same curve). Any algebraic curve has an associated algebraic set, $V(F) = \{x_1, x_2 \in \mathbb{A}^2(K) \mid F(x_1, x_2) = 0\}$, which can also be identified as the curve.

Definition. Suppose that V is an algebraic set over K such that $I(V)$ is a prime ideal (it can be shown that this condition is equivalent to saying that V is not the union of two smaller algebraic sets). Then the *coordinate ring* of V is defined as $\Gamma(V) = K(x_1, \ldots, x_n)/I(V)$, which is then an integral domain. Define the *function field* $K(V)$ of V to be the field of fractions of $I(V)$. For an algebraic curve F for which $V(F)$ is irreducible, as a matter of terminology we write $\Gamma(F) = \Gamma(V(F))$ and $K(F) = K(V(F))$.

However, we run into difficulty because our field K will generally be finite, hence not algebraically closed. When working with an algebraic curve F over such a non-algebraically closed field, the ideal $I(V(F))$ need not be generated by F alone. (For example $\mathbb{F}_q(x)$ contains functions such as $x^q - x$ that vanish everywhere and lie in $I(V(F))$ for any F.) However it's a fact of algebraic geometry (which we will not prove) that this is no longer an issue when K is algebraically closed. This means if we look at the field of functions when we consider F as a curve over $\bar{K}$ (we will assume that $V(F)$ remains irreducible when the field of constants is extended), our coordinate ring will be $\Gamma_{\bar{K}} = \bar{K}[x_1, x_2]/(F)$. However, this new ring has $\bar{K}$ as its field of constants, so we restrict back to K by looking at the image of $K[x_1, \ldots, x_n]$ in this quotient. In effect, then, the way we really want to think of $\Gamma(F)$ and of $K(F)$ is simply as the ring $K[x_1, x_2]/(F)$ and its associated field of fractions. However, we will look at them in relationship to the curve over the algebraic closure $\bar{K}$, which we'll call $V_{\bar{K}}(F)$. Although our finite fields K can't be algebraically closed, we will at least require that K be algebraically closed in $K(F)$. As it turns out, this is not a severe restriction, for one can show that we can always replace K with its algebraic closure K^* in $K(F)$ and choose the curve appropriately so as to get a field of functions extending $K^*(x)$.

Definition. A curve F over a field K is nonsingular if at no point p on $V_{\bar{K}}(F)$

do the formal partial derivatives

$$\frac{\partial F}{\partial x}(p), \frac{\partial F}{\partial y}(p)$$

both vanish.

We will restrict our attention to the case when F is nonsingular: this turns out to not be a restriction on the field $V_{\bar{K}}(F)$.

At this point we note that elements of $K(F)$ do in fact act like rational functions on $V_{\bar{K}}(F)$. First, if $p = (p_x, p_y) \in V_{\bar{K}}(F)$, and $f \in \Gamma(F)$, $f(a)$ is well defined (elements of $I(F)$ are zero at a), and this induces a homomorphism from $K(F)$ to $\bar{K}$. Similarly, if $f = \frac{f_1}{f_2} \in K(V)$, we can define $f(V)$ uniquely as $\frac{f_1(a_x,a_y)}{f_2(a_x,a_y)}$, which will be defined as an element of $\bar{K}$ unless $f_2(a) = 0$: if $f_2(a) = 0$ for any choice of $f_1, f_2 \in \Gamma(F)$ with $f = \frac{f_1}{f_2}$, f is said to have a pole at a.

The field $K(F)$ is a finite extension of $K(x) \cong K(x_1)$ generated by x_2, (although this extension need not be separable), and F is an irreducible curve over $\bar{K}$, so we can extend the field of constants to the new field $\bar{K}(F) = K(F) \otimes_K \bar{K}$. We now show a converse: any finite separable extension L of $K(x)$ with base field (that is, field of algebraic elements) K which satisfies the additional condition that the compositum $L \otimes_K \bar{K}$ with the algebraic closure remains a field is isomorphic to some function field of the form $K(F)$. We first invoke the Theorem of the Primitive Element (here we use that $L/K(x)$ is separable). This lets use write $L = K(x)(a)$, where a is algebraic over $K(x)$. Additionally, a cannot be algebraic over K, because if it were, the minimal polynomial of a over K would have more than its full complement of roots in the field $L \otimes_K \bar{K}$, which is impossible. Let $F \in K(x)[y]$ be the minimal polynomial of a over $K(x)$, so that $F(a) = 0$, and clear denominators and remove any common factors in $K[x]$ so that $F \in K[x][y]$ and the coefficents of y are relatively prime in $K[x]$. Then let $\hat{F}$ be the corresponding element of $K[x, y]$, so that $\hat{F}(x, a) = 0$. This ensures that F is irreducible in $K[x, y]$, but we also want the stronger condition that F is irreducible in $\bar{K}[x, y]$. This follows from our condition that $L \otimes_K \bar{K}$ is a field (see [3]Corollary 2.4.8 for details). Hence $K[x, y]/(\hat{F})$ is an integral domain, whose field of fractions is $K(\hat{F})$. We want to show $K(F) \cong L$. For $K(F)$ is an extension of $K(x)$ by an element y satisfying $\hat{F}(x, y) = 0$, so $F(y) = 0$: by irreducibility of F, F is the minimal polynomial of y over $K(x)$. But F is also the minimal polynomial of a over $K(x)$, and $L = K(x)(a)$, so we have an isomorphism $K(F) \cong L$ identifying y with a.

This shows that our construction of finite extensions of $K(x)$ via algebraic curves is fully general. This will yield a general theory of finite extensions of $K(x)$ (satisfying certain conditions given above) similar to the general algebraic number theory of finite extensions of $\mathbb{Q}$. However, there is one important difference. In the number field case, the field $\mathbb{Q}$ embeds uniquely into K. However, the field $K(x)$ can be embedded into the function field $K(F)$ in more than one way. For any element a of $K(F)$ with $a \notin K$, a is transcendental over K, whence we can embed $K(x) \cong K(a)$ into $K(F)$.

It would be nice if, additionally we could always find a $b \in K(F)$ such that $K(F) = K(a, b)$, however this need not be the case. It can happen that $K(F)$ is not a separable extension of $K(a)$, so that the Theorem of the Primitive Element no longer applies. However, there is a simple criterion for showing when this is the case. The proof is technical and not very enlightening, so we will omit it, but we state the criterion below.

Definition. For any $a \in K(F)$, a is a *separating variable* for $K(F)$ if $K(F)$ is a separable extension of $K(a)$.

Proposition 2.1 ([3] Theorem 2.4.6). *If K is a finite field with characteristic p, let $K(F)^p$ be the subfield of $K(F)$ generated by K and the p-th powers of the elements of $K(F)$. Because the finite field K is perfect and the Frobenius map $a \mapsto a^p$ is a homomorphism, this field will consist exactly of the p-th powers in $K(F)$, and so the notation will cause no confusion. Then $[K : K(F)^p] = p$, and for all $a \in K(F)$ which are not contained in $K(F)^p$, $K(F)$ is a separable extension of $K(a)$.*

However, in the general case we can at least say that $K(F)$ is a finite extension of $K(a)$.

Proposition 2.2. *For any $a \in K(F)$, $a \notin K$, $K(F)$ is a finite extension of $K(a)$.*

Proof. We see first that because $K(F)$ is a finite extension of $K(x)$ and hence of the intermediate field $K(x, a)$, we need only show that $K(x, a)$ is a finite extension of $K(a)$. We have that $K(x, a) \subset K(F)$ is a finite extension of $K(x)$, which implies that a is algebraic over $K(x)$ and there exists a nonzero polynomial $p \in K(x)[y]$ such that $p(a) = 0$. By clearing denominators, we can assume that $p \in K[x][y] \cong K[x, y] \cong K[y][x] \cong K[a][x]$. Hence by regrouping there is a corresponding polynomial $p' \in K[a][x]$ with $p'(x) = 0$, and p' is also nonzero because a is transcendental over K. This means that x is algebraic over $K(a)$, giving that $K(x, a)$, so also $K(F)$, is a finite extension of $K(a)$. □

When we know that a is a separating variable for $K(F)$, we can apply the Theorem of the Primitive Element to arrive at the following.

Corollary 2.3. *For any $a \in K(F)$ outside of $K(F)^p$, there exists $b \in K(F)$ such that $K(F) = K(a, b)$.*

Also, we can deal with the projective case similarly, which gives us a projective function field isomorphic to the one from the affine case, so we can identify the two, so that elements of $K(F)$ can have zeroes and poles at projective "points at infinity" as well as the regular points of the affine plane. However, we note that the points at infinity of $V_{\bar{K}}(F)$ will be of the form $(x, y, 0) \in \mathbb{P}^3(K)$, where $F_{\text{hom}}(x, y, 0) = 0$ (where F_{hom} is the homogenized version of F). Because F is not the zero polynomial, the homogenized version F_{hom} will have only finitely many roots of (x, y) on the projective line $\mathbb{P}^2(K)$.

2.2 Zeroes, Poles, and Local Rings

Before moving on to rings of integers, we will prove a couple more facts about the field structure of $K(F)$ and in particular about poles and zeroes. Recall that if $f \in K(F)$, $p \in V_{\bar{K}}(F)$, f is said to have a pole at p if however you express $f = \frac{f_1}{f_2}$, $f_1, f_2 \in \Gamma(F)$, $f_2(p) = 0$. Also, f has a zero at p if $f(p) = 0$. We note the following useful fact:

Proposition 2.4. *Any nonzero $f \in K(F)$ has only finitely many poles and finitely many zeroes.*

Proof. Write $f = \frac{f_1}{f_2}$, $f_1, f_2 \in \Gamma(F)$. Then for $p \in V_{\bar{K}}(F)$, if f has a zero at p, so does f_1, and if f has a pole at p, f_2 has a zero at p. Hence we need only show that any nonzero $f_1 \in \Gamma(F)$ has only finitely many zeroes. Because only finitely many points at infinity lie on $V_{\bar{K}}(F)$, we need only show that f_1 has finitely many zeroes in the affine plane. Now f_1 is the residue of some polynomial $\hat{f}_1 \in K[x, y]$. Because the principal ideal (F) in $K[x, y]$ is prime, (F) is irreducible in $K[x, y]$, so F and f_1 have no common factor $K[x, y] = K[x][y]$. By Gauss's lemma, they then also have no common factor in of $K(x)[y]$, which is a principal ideal domain, so the ideal $(\hat{f}_1, F)$ of $K(x)[y]$ must be the unit ideal. That is, there exists $A, B \in K(x)$ such that $A\hat{f}_1 + BF = 1$. Applying this to any point $p = (p_x, p_y)$ of $V_{\bar{K}}(F)$, we have $1 = A(p_x)\hat{f}_1(p) + B(p_x)F(p) = A(p)f_1(p) + B(p) \cdot 0 = A(p)f_1(p)$. Hence, the only way we can have $f_1(p) = 0$ is if the denominator of $A(p_x)$ is also zero. This denominator is a nonzero polynomial in p_x, so it can be zero for only finitely many values of p_x.

Hence we deduce that p_x can take on only finitely many values at zeroes of f_1. Similarly, p_y can take on only finitely many values at zeroes of f_1, so f_1 has only finitely many affine zeroes, hence only finitely many total zeroes, and the proposition follows. □

Definition. The local ring $\mathcal{O}_p(F)$ is $\{f \in K(F) \mid f \text{ does not have a pole at } p\}$, $p \in V_{\bar{K}}(F)$.

Henceforwards in this section, we assume, without loss of generality, that $p = (p_x, p_y)$ is not a point at infinity (otherwise perform appropriate projective transformations). This condition implies that $\Gamma(F) \subset \mathcal{O}_p(F)$, so the field of fractions of $\mathcal{O}_p(F)$ is the whole field $K(F)$. These rings $\mathcal{O}_p(F)$ are called *local* rings because they have a unique maximal ideal. In this case, the maximal ideal is $M_p(F) = \{f \in \mathcal{O}_p(F) \mid f \text{ has a zero at } p\}$. It is maximal and is unique, because it contains exactly all non-units of $O_p(F)$. Ultimately, we will piece together the structure of the more general rings of integers we define in the following sections from the relatively simple structures of the local rings $\mathcal{O}_p(F)$ for p ranging over all points of the curve. We first introduce another definition:

Definition. An integral domain R is a *discrete valuation ring* if there exists a non-unit $t \in R$ such that any element of R can be represented as $r = ut^n$, where u is a unit of R, and n is a non-negative integer. (In this case, t is called a *uniformizing parameter* for R. Also this representation is unique because otherwise one would arrive at a relation of the form $t^k = u$, where u is a unit, and $k \neq 0$. This implies that t is a unit, which is impossible.)

There is an equivalent definition of discrete valuation rings, which will also be useful, and whose equivalence we will now prove.

Theorem 2.5. *An integral domain R is a discrete valuation ring as defined above if and only it is Noetherian, local, and its maximal ideal is principal.*

Proof. **Proof of $\Rightarrow$:** If R is a discrete valuation ring, the ideal (t) contains exactly the non-units of R, so it is the unique maximal ideal. Hence we need only show that R is Noetherian. Indeed, any ideal I of R is the finitely generated principal ideal (t^n), where n is the least positive integer such that there is a unit u with $ut^n \in I$.

Proof of $\Leftarrow$: Let t generate the maximal ideal of R. We claim that t is a uniformizing parameter for R. Let r be an arbitrary element of R. We claim that there is some $n \geq 0$ such that $t^n \mid r$, $t^{n+1} \nmid r$ in R. For otherwise, we would have that $t^n \mid r$ for all positive integers n, giving us the infinite ascending chain $(r) \subset (rt^{-1}) \subset (rt^{-2}) \subset \cdots$, contradicting the Noetherian condition. Write $r = ut^n$: we need to show that u is a unit. For if not, u is contained in some maximal ideal of R, which must be t, implying the contradiction $t^{n+1} \mid r$. □

We would like to show that our local ring $\mathcal{O}_p(F)$ is a such a ring: we've already shown that it is local, so we need only show that it is Noetherian and that $M_p(F)$ is principal.

Proposition 2.6. *$\mathcal{O}_p(F)$ is Noetherian.*

Proof. The ring $\Gamma(F)$ is Noetherian because it is a quotients of the Noetherian ring $K[x, y]$. Without loss of generality, assume that p is not a point at infinity, so that $\Gamma(F) \subset \mathcal{O}_p(F)$. Then let I be an ideal of $\mathcal{O}_p(F)$, and intersect with $\Gamma(F)$ to get an ideal of $\Gamma(F)$, which is then finitely generated. Then go back to $\mathcal{O}_p(F)$ and you have the same generators. □

Proposition 2.7. *The maximal ideal $M_p(F)$ is principal in $\mathcal{O}_p(F)$.*

Proof. By suitable projective change of coordinates, we can assume that p is the origin, that is, $p_x = p_y = 0$. Because $p = (0, 0)$ is on the curve given by $F(x, y) = 0$, the constant term of F must be nonzero, so write $F = xA + yB$, $A, B \in K[x, y]$. By nondegeneracy of F, one of $A(0, 0) = \frac{\partial F}{\partial x}$, $B(0, 0) = \frac{\partial F}{\partial y}$ must be zero. Without loss of generality assume $B(0, 0) \neq 0$. We now claim that $M_p(F) = (x)$.

Let g be an arbitrary element of $M_p(F)$. We will show $g \in (x)$. We can write $g = \frac{g_1}{g_2}$, where $g_1, g_2 \in \Gamma(F)$ $g_1(0, 0) = 0$, $g_2(0, 0) \neq 0$. Then g_2 is a unit of $O_p(F)$, so we can assume $g \in \Gamma(F)$. Let G be an element of $K[x, y]$ that corresponds to g in the quotient $\Gamma(F) = K[x, y]/(F)$. Then $g(0, 0) = 0$, so we can write $G = xC + yD$, $C, D \in K[x, y]$. We now have two linear equations in the variables x, y, so we can substitute it for y to get

$$G = \frac{x(BC - AD) - F)}{B}.$$

Because F corresponds to $0 in \Gamma(F)$, when we return to $K(F)$, we obtain a corresponding formula of the form

$$g = x\frac{bc - ad}{b},$$

(where the lower case letters are images of their upper case counterparts in the quotient $\Gamma(F)$.) Because $B(0,0) \neq 0$, $\frac{bc-ad}{b}$ does not have a pole at $(0,0)$, so $g \in (x)$.

Hence (x) contains the maximal ideal $M_p(F)$, and so the two ideals are equal and $M_p(F)$ is principal. □

We now deduce our final result:

Proposition 2.8. *The integral domain $\mathcal{O}_a(F)$ is a discrete valuation ring.*

This is a very useful result. We can use it to deduce a simple lemma with two immediate corollaries about zeroes and poles for general elements of $K(F)$. We will return to this result from in Section 3 when we prove that our rings of integers are Dedekind domains.

Lemma 2.9. *Suppose that t is a uniformizing parameter for $\mathcal{O}_p(F)$. Then for any nonzero $f \in K(F)$, we can express $f = ut^n$, where u is a unit of $\mathcal{O}_p(F)$ and n can now be any integer. In fact, this representation is unique, so that we can define $\operatorname{ord}_p(f) = n$: this is the "valuation" of the "discrete valuation ring" $\mathcal{O}_p(F)$ which satisfies $\operatorname{ord}_p(fg) = \operatorname{ord}_p(f) + \operatorname{ord}_p(g)$ and $\operatorname{ord}_p(f+g) \geq \min \operatorname{ord}_p(f), \operatorname{ord}_p(g)$.*

Proof. This is a direct consequence of the definition of a discrete valuation ring and the fact that $K(F)$ is the field of fractions of $\mathcal{O}_p(F)$. Uniqueness holds for the same reason as in $\mathcal{O}_p(F)$ case, and the rest follows automatically. □

Corollary 2.10. *Let $a \in V_{\bar{K}}$ and $f, g \in K(F)$ such that g has a zero at p. Then for sufficently large $n \in \mathbb{N}$, fg^n does not have a pole at p.*

Proof. By the previous lemma, $f = u_f t^{\operatorname{ord}_p(f)}$, $g = u_g, t^{\operatorname{ord}_p(g)}$, where u_f, u_g are units of $\mathcal{O}_p(F)$. Because g has a zero at p, $g \in M_p(F)$, so $\operatorname{ord}_p(g) \geq 1$. Then, for sufficently large $n \in N$, $\operatorname{ord}_p(fg^n) = \operatorname{ord}_p(f) + n \operatorname{ord}_p(g) \geq 0$, so $fg^n = u_f {u_g}^n t^{\operatorname{ord}_p(fg^n)} \in \mathcal{O}_p(F)$ doesn't have a pole at p. □

Corollary 2.11. *If $f \in K(F)$ is nonzero, $p \in V_{\bar{K}}$, f has a zero at p if and only if f^{-1} has a pole at p.*

Proof. Write $f = u_f t^{\operatorname{ord}_p(f)}$, $f^{-1} = {u_f}^{-1} t^{-\operatorname{ord}_p(f)}$. Then f has a zero at p if and only if $\operatorname{ord}_p(f) > 0$ if and only if $-\operatorname{ord}_p(f) < 0$ if and only if f^{-1} has a pole at p. □

2.3 The Ring of Integers $\mathcal{O}_{K(F),S}$

In order to do anything interesting we need to get an analogue of the ring of integers. If we fix a finite set of points S on the projective curve, that is, elements of $V_{\bar{K}}(F)$ (identifying points that are Galois conjugates of each other), we take $\mathcal{O}_{K(F),S} = \{f \in K(F) \mid f \text{ only has poles at points in } S\}$. This is a subring of $\mathcal{O}_{K(F),S}$, hence it is an integral domain. We will neglect to show that in the general case this ring has any elements outside of K. This is a consequence of the Riemann-Roch theorem, which would take us too far afield in our present journey. We content ourselves with the knowledge that we can easily construct such nontrivial elements in many specific cases, and that even if such trivial cases did exist, we would gain nothing from considering them. Hence we can proceed with the assumpion that $\mathcal{O}_{K(F),S}$ is not merely K. In the case of a number field K, the ring of integers $\mathcal{O}_K$ has the property that its field of fractions is all of K, a corollary of Stein 2.3.11, which states that $\mathbb{Q}\mathcal{O}_K = K$. We can show an analogue of this fact for function fields, but the proof is different because we have no canonical analogue of $\mathbb{Q}$. Rather, we must explore more deeply the structure of $V_{\bar{K}}(F)$. Specifically, we will take the route of finding $a, b \in \mathcal{O}_{K(F),S}$ which generate $K(F)$ over K

Proposition 2.12. *For any function field $K(F)$ over K, and any set of points on F, we can find $a, b \in \mathcal{O}_{K(F),S}$ such that $K(F) = K(a, b)$. In fact for any $a \in \mathcal{O}_{K(F),S}$ which is a separating variable for $K(F)$, we can find a b to satisfy the above. As a result, the field of fractions of $\mathcal{O}_{K(F),S}$ is $K(F)$.*

Proof. We have assumed there exists some $a \in \mathcal{O}_{K(F),S}$ outside of K, but we don't know that this a must be a separating variable for $K(F)$. However, we do know that $K(F)$ does contain separating variables by Goldschmidt's criterion. Let s be a separating variable for $K(F)$: then $s \notin K(F)^p$. Without loss of generality, assume that $a \in K(F)^p$ (otherwise we'd be done already). Then we claim that for any pole p of s outside S, there is some nonzero $c_p \in K[a]$ such that c_p has a zero at p. For consider the value $a(p)$: it is well defined because a doesn't have a pole, and it lies in $\bar{K}$, so there is some nonzero polynomial $C_p \in K[x]$ such that $C_p(a(p)) = 0$. So let $c_p = C_p(a)$, which is nonzero because a is transcendental over K. Also each of the c_p are in the field $K(F)^p$. By multiplying by sufficiently high powers of each of the finitely many c_p, we can remove all poles of s to produce an element s' of $\mathcal{O}_{K(F),S}$. Also $\frac{s'}{s} \in K(F)^p$, so since s' is not in the field $K(F)^p$, neither is s. Hence our element s' satisfies all criteria needed, and we can take it for our new a.

Suppose that $a \in \mathcal{O}_{K(F),S}$ is a separating variable for $K(F)$. Write $K(F) = K(a, b')$, for $b' \in K(F)$. Then b' has only finitely many poles outside S: as above, by multiplying by an appropriate element of $K[a]$, we can remove them all. Then $b \in \mathcal{O}_{K(F,S)}$, and $K(a, b) = K(a, b') = K(F)$. This is exactly what we want, and $K(F) = K(a, b)$ is the field of fractions of $K[a, b]$, so also of the larger ring $\mathcal{O}_{K(F),S}$.

□

2.4 Ideals and Prime Ideals

Now we have laid our groundwork and we are ready to do something that looks more like algebraic number theory, that is, study the structure of the ideals of $\mathcal{O}_{K(F),S}$. In particular, we will focus on the prime ideals; unlike in the number field case, where the prime ideals have no special significance, for function fields there is an elegant correspondence between the prime ideals of $\mathcal{O}_{K(F),S}$ and the points of $V_{\bar{K}}(F)$ (up to action of the Galois group) that are not in S. In order to establish this, however, we first need to investigate the structure of the quotient of $\mathcal{O}_{K(F),S}$ by any prime ideal.

Lemma 2.13. *If R is a ring with field of fractions $K(F)$, and $\mathfrak{p}$ a prime ideal of R, the ring $R/\mathfrak{p}$ is algebraic over K.*

Proof. Let a be an arbitrary element of $\mathfrak{p}$. The field $K(F)$ is a finite, hence algebraic, extention of $K(a)$, so that for any $b \in R$, b is algebraic over $K(a)$. Clearing denominators, b is also algebraic over the ring $K[a]$. Then after quotienting by p, the residue $\hat{b}$ of b is algebraic over the image of $K[a]$, which reduces to K. Because $\hat{b}$ is an arbitrary element of $R/\mathfrak{p}$, $R/\mathfrak{p}$ is itself algebraic over K. □

We now think about some concrete examples of prime ideals of $K(F)$. We have previously seen the ideals $M_p(F)$ of $\mathcal{O}_p(F)$, where p is a point on $V_{\bar{K}}(F)$. For any $p \notin S$, these ideals can be restricted to the smaller domain $\mathcal{O}_{K(F),S}$, to produce the ideals

$$\mathfrak{p}_{p,S} = \{f \in \mathcal{O}_{K(F),S} \mid f(p) = 0\}$$

of $\mathcal{O}_{K(F),S}$. The ideal $\mathfrak{p}_{p,S}$ is prime because it is the kernel of the homomorphism $\phi : \mathcal{O}_{K(F),S} \to \bar{K}$ given by $\phi(f) = f(p)$ (which always exists because $p \notin S$, so f can't have a pole at p). We'll now show that any prime ideal of $\mathcal{O}_{K(F),S}$ arises in this manner.

Corollary 2.14. *Any prime ideal of $\mathcal{O}_{K(F),S}$ is of the form $\mathfrak{p}_p$ for some point $p = (p_x, p_y)$ of $V_{\bar{K}}(F)$ outside of S.*

Proof. Let $\mathfrak{p}$ be an arbitrary prime ideal of $\mathcal{O}_{K(F),S}$We know from Lemma 2.13 that $\mathcal{O}_{K(F),S}/\mathfrak{p}$ is algebraic over K, hence it can be embedded into the algebraic closure $\bar{K}$ of K. Let $\phi : \mathcal{O}_{K(F),S}/\mathfrak{p} \to \bar{K}$ be such an embedding. Lift ϕ to $\mathcal{O}_{K(F),S}$ to give a homomorphism $\hat{\phi} : \mathcal{O}_{K(F),S} \to \bar{K}$.

Now let a, b be elements of $\mathcal{O}_{K(F),S}$ that generate $K(F)$ over K. Let $a_p = \phi(\hat{a})$, $b_p = \phi(\hat{b})$. Because $x, y \in K(F)$, there are rational functions $X, Y \in K(A, B)$ such that $X(a, b) = x$, $Y(a, b) = y$. We can think of $K(F)$ as also being parametrized by a, b as opposed to the standard parameters x, y, that is, it is isomorphic to the function field $K(F')$ for some algebraic curve $F' \in K[A, B]$ with $F'(a, b) = 0$. Then if either $X(A, B)$ or $Y(A, B)$ has a pole at the point $A = a_{\mathfrak{p}}$, $B = b_{\mathfrak{p}}$ on $V_{\bar{K}}(F')$, we can apply a appropriate projective transformation to our coordinates x, y to make this not be the case. Then let $X(a_{\mathfrak{p}}, b_{\mathfrak{p}}) = x_{\mathfrak{p}}$, $Y(a_{\mathfrak{p}}, b_{\mathfrak{p}}) = y_{\mathfrak{p}}$.

Now let h be an arbitrary element of $\mathcal{O}_{K(F),S}$, and write $h = H(x, y)$, where $H \in K[X, Y]$. Then we can formally apply our homomorphism $\hat{\phi}$ to h:

$$\begin{aligned}\hat{\phi}(h) &= \hat{\phi}(H(x, y)) \\ &= H(\hat{\phi}(x), \hat{\phi}(y)) \\ &= H(\hat{\phi}(X(a, b)), \hat{\phi}(Y(a, b)) \\ &= H(X(a_\phi, b_\phi), Y(a_\phi, b_\phi)) \\ &= H(x_\mathfrak{p}, y_\mathfrak{p}) \\ &= h(x_\mathfrak{p}, y_\mathfrak{p})\end{aligned}$$

where the last step is true by the formal definition of $h(x_\phi, y_\phi)$. Therefore, the map $\hat{\phi}$ is the map evaluating an element of $\mathcal{O}_{K(F),S}$ at some point $p_\mathfrak{p} = (x_\mathfrak{p}, y_\mathfrak{p})$ on $V_{\bar{K}}(F)$ (which is not on S, becaus ϕ is defined everywhere. Then its kernel $\mathfrak{p}$ must equal the prime ideal $p_{\mathfrak{p}_p}$ corresponding to the point $\mathfrak{p}_p$, which takes the desiredform. □

In this manor, we find there is a correspondence between the prime ideals of $\mathcal{O}_{K(F),S}$ and the points of $V_{\bar{K}}(F)$ outside of S.

3 Dedekind Domains and Unique Ideal Factorization

We have laid our groundwork, and we are ready to prove our big theorem, that $\mathcal{O}_{K(F),S}$ is a Dedekind domain having unique factorization of ideals. In the standard treatment (see [7]) of algebraic number theory, unique factorization of ideals for rings of integers in number fields is deduce as a corollary of the fact that such rings are Dedekind domains. However, due to the difficulty in showing that $\mathcal{O}_{K(F),S}$ is Noetherian, in our case we will effectively prove unique ideal factorization before establishing that $\mathcal{O}_{K(F),S}$ is a Dedekind domain. We recall the definition of a Dedekind domain:

Definition. A ring R is a Dedekind domain if it satisfies:

(i) R is Noetherian.

(ii) R is integrally closed in its field of fractions.

(iii) Prime ideals of R are maximal.

We note that we have already shown the last of these three conditions in the preceding section. We prove that R is Noetherian first, since along the way we will gain a deeper understanding of the structure of ideal factorization in R. In the number field case, the concept of fractional ideals is again useful, so we review it here (for technical reasons, the definition of a fractional ideal used here differs from the one used in Stein [7], but the two definitions are equivalent for Noetherian rings).

Definition. Let R be a ring with field of fractions K. A fractional ideal I of R is an R-submodule of K such that for some nonzero $d \in R$, $dI \subset R$. As with ideals, the product of two fractional ideals I, J is the fractional ideal generated by all elements of the form ij, where $i \in I$, $j \in J$. The inverse of a fractional ideal I is another fractional ideal I^{-1} such that $II^{-1} = R$.

We'll now show that any prime ideal of $\mathcal{O}_{K(F),S}$ has an inverse, so that we will be able to "divide" by prime ideals when factoring ideals later on in our proof.

Proposition 3.1. *For any $p \in V_{\bar{K}}(F)$ outside of S, the prime ideal $\mathfrak{p}_p$ is invertible in $\mathcal{O}_{K(F),S}$, and its inverse is the fractional ideal*

$$\mathfrak{p}_p^{-1} = \{f \in \mathcal{O}_{K(F),S\cup\{p\}} \mid \operatorname{ord}_p(f) \geq -1\}.$$

Proof. We need only show that $\mathfrak{p}_p\mathfrak{p}_p^{-1} = \mathcal{O}_{K(F),S}$. Let t be an element of $\mathcal{O}_{K(F),S}$ with $\operatorname{ord}_p t = 1$. (Such a one exists by Riemann-Roch, but if you would rather not invoke that, just assume that $\operatorname{ord}_p t$ is minimal, and replace 1 by $\operatorname{ord}_p t$ as applicable.) Then $t \in \mathfrak{p}_p$, $t^{-1} \in \mathfrak{p}_p^{-1}$, and so $1 = tt^{-1} \in \mathfrak{p}_p\mathfrak{p}_p^{-1}$, and $\mathcal{O}_{K(F),S} \subset \mathfrak{p}_p\mathfrak{p}_p^{-1}$. For the other inclusion, $\mathfrak{p}_p$ and $\mathfrak{p}_p^{-1}$ are both contained in the larger field $\mathcal{O}_{K,S\cup\{p\}}$. Furthermore, if $i \in \mathfrak{p}_p$, $j \in \mathfrak{p}_p^{-1}$, $\operatorname{ord}_p(ij) = \operatorname{ord}_p(i) + \operatorname{ord}_p(j) \geq 1 + -1 = 0$, so, even after taking linear combinations, a general element of $\mathfrak{p}_p\mathfrak{p}_p^{-1}$ cannot have a pole at p either. We now have that $\mathfrak{p}_p\mathfrak{p}_p^{-1} = \mathcal{O}_{K(F),S}$, and so $\mathfrak{p}_p$ is invertible, as desired. □

Before showing existence of factorizations, we note a couple of facts:

Lemma 3.2. *For a given point $p \in V_{\bar{K}}(F)$,*

$$\bigcap_{k=0}^{\infty} \mathfrak{p}_p^k = \{0\},$$

and for $p_1, p_2, p_3, \ldots \in V_{\bar{K}}(F)$ all distinct,

$$\bigcap_{k=0}^{\infty} \mathfrak{p}_{p_k} = \{0\}$$

Proof. For the first part, we note that for any $f \in \mathfrak{p}_p$, $\operatorname{ord}_p(f) \geq 1$, and $\operatorname{ord}_p(f^k) \geq k$. By the property that $\operatorname{ord}_p(f+g) \geq \min \operatorname{ord}_p(f), \operatorname{ord}_p(g)$, it follows that any nonzero element f of the ideal $\mathfrak{p}_p^k$ must also have $\operatorname{ord}_p(f) \geq k$. But then, if $f \in \bigcap_{k=0}^{\infty} \mathfrak{p}_p^k$, $\operatorname{ord}_p(f)$ must be greater than any positive integer k. Hence this intersection can contain only zero.

For the second part, suppose that $f \in \bigcap_{k=0}^{\infty} \mathfrak{p}_{p_k}$. Then for each $k \geq 0$, f has a zero at p_k. Hence f has infinitely many zeroes, and must be zero.

□

We will now show that an arbitrary ideal I of $\mathcal{O}_{K(F),S}$ can be factored into prime ideals of $\mathcal{O}_{K(F),S}$.

Proposition 3.3. *Let I be a nonzero ideal contained in $\mathcal{O}_{K(F),S}$. Then there exists an $n \in \mathbb{N}$ and (not necessarily distinct) prime ideals $\mathfrak{p}_1, \mathfrak{p}_2, \mathfrak{p}_3 \ldots, \mathfrak{p}_n$ of $\mathcal{O}_{K(F),S}$ such that I can be expressed as the product $\prod_{k=1}^{n} \mathfrak{p}_{p_k}$.*

Proof. We claim that we can inductively construct a (possibly terminating) sequence of ideals $\{I_k\}$ indexed by positive integers, and a sequence of prime ideals $\mathfrak{p}_k$ such that for each k, $I_k \prod_{i=1}^{k} \mathfrak{p}_i = I$. Let $I_0 = I$. Then, for each k, if I_{k-1} is not the unit ideal, let $\mathfrak{p}_k$ be a maximal ideal containing I_k. Additionally, let I_k be the product of fractional ideals $I_{k-1}\mathfrak{p}_k^{-1}$: we claim that this is in fact an integral ideal. After all, $I_k = I_{k-1}\mathfrak{p}_k^{-1} \subset \mathfrak{p}_k\mathfrak{p}_k^{-1} = \mathcal{O}_{K,S}$. Finally, we see inductively that

$$I_k \prod_{i=1}^{k} \mathfrak{p}_i = I_{k-1}\mathfrak{p}_k^{-1} \prod_{i=1}^{k} \mathfrak{p}_i = I_{k-1} \prod_{i=1}^{k-1} \mathfrak{p}_i = I.$$

If for some k, $I_k = \mathcal{O}_{K(F),S}$, we have produced the desired factorization $I = \prod_{i=1}^{k} \mathfrak{p}_i$. Hence we need only show that this process must terminate. For suppose not. We see that $\prod_{i=1}^{k} \mathfrak{p}_i \subset I_k \prod_{i=1}^{k} \mathfrak{p}_i = I$, for any positive integer k. Taking the intersection,

$$\bigcap_{k=1}^{\infty} \prod_{i=1}^{k} \mathfrak{p}_i \supset I.$$

On the other hand, either there are infinitely many distinct $\mathfrak{p}_i$, or some $\mathfrak{p}_i$ appears infinitely many times. In either case, as a result of Lemma 3.2 implies the infinite intersection $\bigcap_{k=1}^{\infty} \prod_{i=1}^{k} \mathfrak{p}_i$ is empty, a contradiction. Hence the process must terminate, and I can be factorized into prime ideals. □

As a consequence, we can prove unique factorization of ideals:

Theorem 3.4 (Unique Ideal Factorization). *If I is an ideal of $\mathcal{O}_{K(F),S}$, there exists a unique integer k unique choice (up to ordering) of prime ideals $\mathfrak{p}_i$ and positive integers n_i for $i = 1, 2, \ldots, k$ such that*

$$I = \prod_{i=1}^{k} \mathfrak{p}_i^{n_i}$$

Proof. We know that such a factorization exists by Theorem 3.3. The proof that these factorizations is unique proceeds exactly as in Stein 3.1.9, using the previously that prime ideals are invertible. □

Although we could stop here now that we have proved the powerful result of unique factorization, we will proceed to show that $\mathcal{O}_{K(F),S}$ additionally shares the other basic properties that Dedekind domains enjoy: namely, the Noetherian property and that of being integrally closed in its field of fractions.

Proposition 3.5. *The integral domain $\mathcal{O}_{K(F),S}$ is Noetherian.*

Proof. We will show that if I is any ideal of $\mathcal{O}_{K(F),S}$, there are only finitely many ideals containing I. By Proposition 3.3, we can factorize I into prime ideals, and can do the same with J. Because J is a product of prime ideals, J has an inverse $J^{-1} = \prod_{j=1}^{m} \mathfrak{p}_j'^{-n_j'}$, and $J^{-1}I \subset J^{-1}J$ is an integral ideal. We can then factor IJ^{-1} into prime ideals as well, and the product of the prime factorizations of J and IJ^{-1} must yield the unique prime ideal factorization of I. Hence only ideals that appear in the prime factorization of I can appear in the prime factorization of J, and no prime ideal can appear in the factorization of J to a higher power that it does for I. This yields only finitely many possibilites for J, proving our claim.

Now the ascending chain condition is clear: if $I_1 \subset I_2 \subset I_3 \ldots$ is an infinite ascending chain of ideals, all ideals in the chain must contain I_1. There are only finitely many such ideals, so the chain must eventually terminate. □

We finally show that $\mathcal{O}_{K(F),S}$ is integrally closed in its field of fractions. For this we take a different tack. First of all, we claim that it suffices to show the result for the fields $\mathcal{O}_a(F)$, where a ranges over all points of $V_{\bar{K}}(F)$ outside of S.

Proposition 3.6. *If for each point $a \in V_{\bar{K}}(F)$ outside S, the field $\mathcal{O}_a(F)$ is integrally closed in its field of fractions $K(F)$, so is $O_{K(F),S}$.*

Proof. First, we note that by definition $\mathcal{O}_{K(F),S} = \bigcap_a \mathcal{O}_a(F)$, where a ranges over all points of $V_{\bar{K}}(F)$ outside S. Then suppose that $\alpha \in K(F)$ satisfies a monic polynomial equation $p(\alpha) = 0$ with coefficients in $\mathcal{O}_{K(F),S}$: we need to show that $\alpha \in \mathcal{O}_{K(F),S}$. For any $a \in V_{\bar{K}}(F) - S$, p also has coefficients in the larger integrally closed ring $\mathcal{O}_a(F)$, so $pin\mathcal{O}_a(F)$. Since this is true for any point a outside of S, we deduce that $\alpha \in \mathcal{O}_{K(F),S}$, which must then also be integrally closed in its field of fractions. □

Now we need only show that $\mathcal{O}_a(F)$ is integrally closed in its field of fractions. This turns out to be a general property of discrete valuation rings.

Proposition 3.7. *Let R be a discrete valuation ring with field of fractions K. Then R is integrally closed in K*

Proof. Suppose not. Then let $\alpha \in K$, $\alpha \notin R$ such that a satisifies a monic polynomial $p \in R[a]$. Then let t be a uniformizing parameter for R. We can write $\alpha = ut^{-n}$, where u is a unit of R and n is a positive integer. Additionally, we can express

$$p(x) = x^N + v_{N-1}t^{a_{N-1}}x^{N-1} + v_{N-2}t^{a_{N-2}}x^{N-2} + \cdots + v_1 t^{a_1} x + v_0 t^{a_0}$$

where all of the t are non-negative. Then

$$t^{(N-1)n}p(\alpha) =$$
$$u^N t^{-n} + v_{N-1}u^{N-1}t^{a_{N-1}} + v_{N-2}u^{N-2}t^{a_{N-2}+n} + \ldots + v_1 u_1 t^{(N-2)n} + v_0 t^{(n-1)N}.$$

The left hand side of this expression is in R, as all all terms of the right hand side save the first. Hence $u^N t^{-n} \in R$, which is impossible. □

Now we are ready to wrap everything up, with our big theorem:

Theorem 3.8. *For an algebraic curve F over a finite field K, and some finite subset S of points of $V_{\bar{K}}(F)$ on the curve, the ring $\mathcal{O}_{K(F),S}$ is a Dedekind domain.*

Proof. We have proven that the maximal ideals of $\mathcal{O}_{K(F),S}$ are maximal, that it is Noetherian, and that $\mathcal{O}_{K(F),S}$ is integrally closed in its field of fractions $K(F)$. This is exactly what we had to show. □

From this, we can rededuce our previous result of unique ideal factorization. Additionally, we know that any theorem about rings of integers in algebraic number fields that uses only the Dedekind domain properties extends immediately to $\mathcal{O}_{K(F),S}$. For example, any ideal of $\mathcal{O}_{K(F),S}$ can be generated by two elements. Additionally, we can define a class group $\mathrm{Cl}(K(F), S)$ in exactly the same way as has been done for algebraic number fields $K/\mathbb{Q}$.

4 Conclusion

Although we have accomplished our major goal of showing that $\mathcal{O}_{K(F),S}$ is, like our previously studied rings of integers in number fields, this is more a beginning than an end. For one thing, our results have allowed us to construct class groups for our rings of integers $\mathcal{O}_{K(F),S}$, which are new objects to study and investigate in the same manner as has been done for number fields. Additionally, well-defined prime ideal factorization allows us to study the question of how prime ideals factor when one moves to a finite extension: whether it ramifies, splits, stays inert, or does some combination of the above. And even without thinking about ideal factorization, we have another object, the group of units, to play with and investigate. There is much room to explore and seek parallels with number fields, although some tools from the number field case will not carry over directly to function fields. This is especially true when exploiting the geometry, analysis, and topology of $\mathbb{Q}$ and its rational extensions as embedded in $\mathbb{C}$. However, there are still many fruitful and far-reaching parallels that we have not have time to delve into.

References

[1] William Fulton, *Algebraic Curves* (New York, 1969)

[2] Robert J. Walker, *Algebraic Curves* (New York, Dover, 1962)

[3] David M. Goldschmidt, *Algebraic Functions and Projective Curves*, online evaluation copy: `http://www.functionfields.org`

[4] J. P. May, *Notes on Dedekind Rings*: `http://www.math.uchicago.edu/~may/Dedekind.pdf`

[5] Felipe Voloch, *Notes on Algebraic Number Theory*:
`http://www.ma.utexas.edu/users/voloch/algnoth.html`

[6] Carlos Moreno, *Algebraic Curves Over Finite Fields* (Cambridge University Press, 1991)

[7] William Stein, *Math 129 textbook*:
`http://modular.fas.harvard.edu/129-05/notes/129.pdf`