



Universiteit
Leiden
The Netherlands

On class field theory for curves

Stolk, A.P.

Citation

Stolk, A. P. (2006). *On class field theory for curves*.

Version: Not Applicable (or Unknown)

License: [License to inclusion and publication of a Bachelor or Master thesis in the Leiden University Student Repository](#)

Downloaded from: <https://hdl.handle.net/1887/3597550>

Note: To cite this publication please use the final published version (if applicable).

On Class Field Theory for Curves

Arjen Stolk



Doctoraalscriptie Wiskunde, verdedigd op 28 juni 2006

Scriptiebegeleider - prof. dr. S. J. Edixhoven

Mathematisch Instituut, Universiteit Leiden

Contents

1	Introduction	1
2	Curves	3
3	Endomorphisms of the Additive Group	7
4	Drinfeld Modules	13
5	Constructions with Drinfeld Modules	20
6	Abelian Extensions	34
7	References	43

1 Introduction

Many theorems of mathematics assert that given certain data, there exists an object with some desired properties. It may at first seem surprising, but frequently these theorems, or indeed their proofs, do not actually tell you how to find said object! Class field theory is one such area where the theorems aren't particularly helpful in actually finding the objects which they predict.

This thesis belongs within the study of *explicit* class field theory, which seeks to give constructions for the Abelian extensions predicted by the main theorems of class field theory. Surprisingly, the techniques used in this theory are not at all the ones used for proving the abstract class field theory.

Explicit CFT is by no means a complete theory. Indeed, when the base field is a number field, not very much is known. The theorem of Kronecker and Weber gives a perfect description of the class fields of the rational numbers, the simplest number field. For imaginary quadratic fields we can also do explicit CFT. This is the theory of elliptic curves with complex multiplication. Beyond these cases there is not much general theory for number fields.

But the theorems of CFT don't just work in the context of number fields. There is also *local* CFT, where the original theorems do actually provide explicit constructions. The last situation in which CFT works is that of global function fields. These are the function fields of curves over finite fields. It is this theory we will be focusing on in this thesis.

What makes these fields interesting is their geometric interpretation. Within algebraic geometry there is a lot of general theory and machinery that can be applied to the curves associated to these fields. These methods largely focus on considering the whole curve, which is a projective variety and therefore very nice to work with.

We do not take this approach. Instead, we break the symmetry of the curve, by throwing out one of its points. The resulting object, although perhaps not as pretty as the original curve, has a much simpler structure, closer to the setup we encounter in the study of number fields.

In the late 1930's, R. Carlitz developed an explicit class field theory for the function field of the projective line. His theorem resembles the that of Kronecker and Weber for the rational numbers. The price Carlitz pays for breaking the symmetry is that not all class fields are obtained using his theorem. In some sense, his theory misses the part that comes from the point that has been left out.

In the 1970's, V.G. Drinfeld, a Russian mathematician, developed a theory of what he called elliptic modules. To some extent they resemble the elliptic curves that we encounter in algebraic number theory. Again, his constructions start with a curve with one point removed. Drinfeld is not directly interested in explicit class field theory. His focus is on the Langlands conjectures. These are in some sense a very strong theoretical generalisation of class field theory. The Langlands conjectures for number fields are still open and much work is being done trying to at least understand the rank 2 theory. The rank 1 theory corresponds roughly to class field theory.

Drinfeld also has a theory of shtukas. These are objects that generalise his elliptic modules. Moreover, they restore the symmetry of the curve. Indeed there is a very nice way of formulating part of class field theory geometrically using shtukas. L. Lafforgue has in recent years succeeded in *proving* the Langlands conjectures for global function fields using generalisations of Drinfeld's shtukas and the advanced machinery of modern algebraic geometry.

What we are concerned with in this thesis is using Drinfeld's elliptic modules, or Drinfeld modules as we call them today, to do explicit class field theory for global function fields. This theory was developed in detail by D. Hayes. Just like Carlitz' theory, the theory of Hayes does not produce all the class fields, but misses what happens at the point that we have taken out of the curve. Hayes' development of the theory mirrors the theory of elliptic curves with complex multiplication, however the scope of Hayes' theory is much larger. CM-theory works exclusively for imaginary quadratic fields, while Hayes' theory applies equally to all global function fields.

We begin by showing the relation between the geometry of curves over finite fields and their function fields. This theory, collected in chapter two, is not used explicitly within the next chapters, but it provides the theoretical framework within which we can place these fields. Chapter three provides more technical prerequisites.

Chapters four and five introduce Drinfeld modules and derive some of the basic machinery one uses to understand them. In particular we look at their morphisms and several invariants associated to a Drinfeld module. We also give an analytical description of Drinfeld modules over the global function field equivalent of the complex numbers. Readers with some knowledge of the theory of elliptic curves will be struck by the similarities that exist between that theory and this one.

In the last chapter we apply the theory of Drinfeld modules that we have created in order to do explicit class field theory. The theorems that come out actually come in two flavours. What we get from Hayes' construction isn't quite what we want for class field theory. The fields are a little too big as his construction does allow for slightly larger extensions at the point that we have removed. However, there isn't much that we can do with this extra bit of information rather than control it so much that we can remove it, producing a theorem that is truly an explicit construction of fields predicted by class field theory.

To conclude, a few words on the sources I have used to prepare this text. For chapter two, I have mostly used R. Hartshorne's standard work on algebraic geometry and knowledge gained from lectures by prof. H.W. Lenstra. My exposition of the theory of Drinfeld modules and doing class field theory with them largely follows Hayes' overview article [Hayes1]. I have also found D. Goss' book on the subject, [Goss], pleasantly accessible and have used it extensively for chapter three.

Occasionally I have referred to Drinfeld's original article [Dri] in which he introduced the notion of elliptic modules. On several occasions I have used two other articles of Hayes, [Hayes2] and [Hayes3] and an article by Deligne and Husemöller [D-H] to provide an alternative view on the same material. Mumford's book on abelian varieties [Mum] and Silverman's work on elliptic curves provided me with additional insight into the parallel world within the theory of number fields. Lastly, the little book of Atiyah and MacDonal [A-M] is an indispensable reference for all things related to commutative algebra.

2 Curves

Definition 2.1. Let X be an integral scheme. Then the function field $k(X)$ of X is the residue field at the generic point ζ_X of X .

It is easy to see that in fact $k(X)$ is the local ring at ζ_X .

Let X be an integral scheme, then there is a natural map $\text{Spec}(k(X)) \longrightarrow X$, which is just the inclusion of the generic point. This means that there are maps $\phi_U : \Gamma(U, \mathcal{O}_X) \longrightarrow k(X)$ for all non-empty opens U of X . These maps are compatible with restrictions.

Lemma 2.2. For every non-empty open U of X , the map ϕ_U is injective, so we can indentify the ring $\Gamma(U, \mathcal{O}_X)$ with its image inside $k(X)$. Then we have that $\Gamma(U, \mathcal{O}_X)$ is $\bigcap_{\text{Spec}(A)} A$, where the intersection runs over all non-empty affine opens $\text{Spec}(A)$ inside U .

Proof. For non-empty affine opens $\text{Spec}(A)$ we can easily identify $\phi_{\text{Spec}(A)}$. The generic point of X is in $\text{Spec}(A)$ and corresponds to the ideal (0) there. Thus $k(X) = A_{(0)}$ is the fraction field $Q(A)$ of A . We conclude that the map is injective for all non-empty affine opens.

Let $f \in \Gamma(U, \mathcal{O}_X)$. For every non-empty affine open $\text{Spec}(A)$ inside U we see that $f|_{\text{Spec}(A)} = \phi_U(f)$, so $\phi_U(f) \in A$. Suppose that $\phi_U(f) = 0$ then we see that $f|_{\text{Spec}(A)} = 0$ for every non-empty affine open $\text{Spec}(A) \subset U$. But then $f = 0$. We conclude that ϕ_U is injective and that $\Gamma(U, \mathcal{O}_X) \subset A$ for all $\text{Spec}(A) \subset U$ non-empty.

Conversely, let $g \in k(X)$ and suppose that $g \in A$ for all non-empty affine opens $\text{Spec}(A)$ inside U . Then we have a section on all these $\text{Spec}(A)$'s and these sections are compatible, so we get a section on their union, which is U . \square

Corollary 2.3. If X is an integral scheme of finite type over a field k , then $k(X)$ is a finitely generated field extension of k whose transcendence degree is the dimension of X .

Proof. Take $\text{Spec}(A) \subset X$ a non-empty affine open. Then A is a finitely generated k -algebra that is a domain. It follows from dimension theory for such algebras that the dimension is the transcendence degree of the fraction field. \square

Definition 2.4. A curve over a field k is a separated integral scheme of dimension 1 that is of finite type over k .

Proposition 2.5. Let X be a curve. The following are equivalent

1. X is normal;
2. X is regular;
3. for any non-empty affine open subset $\text{Spec}(A)$ of X , A is a Dedekind domain.

Proof. Let $\text{Spec}(A) \subset X$ be a non-empty affine open. Since X is of finite type over a field, X is Noetherian. So A is a Noetherian domain of dimension 1. The theory of Dedekind

domains now states that A is Dedekind if and only if it is integrally closed in $Q(A)$. Also A is Dedekind if and only if all the local rings are discrete valuation rings, that is, regular local rings of dimension 1. This shows the conditions are indeed equivalent. \square

Definition 2.6. A curve is called *complete* if it is proper over k .

Lemma 2.7. *Let X be a curve over k . Then there is a map from the set of points of X to the set of valuation rings of $k(X)$ containing k , sending $x \in X$ to $R_x = \mathcal{O}_{X,x}$. If X is regular, this map is injective and we have $\Gamma(U, \mathcal{O}_X) = \bigcap_{x \in U} R_x$ for all nonempty opens U of X . If X is complete, this map is surjective.*

Proof. This is a direct application of the valuative criteria for separatedness and properness. The observation about regular functions follows from the fact that for any domain R with field of fractions K the intersection of all valuation rings of K containing R is the integral closure of R in K . \square

Let K and L be fields containing k . Recall that a *place* over k from K to L is a valuation ring R of K containing k and a ring homomorphism $f : R \rightarrow L$ such that $f(x) = 0$ for all x in the maximal ideal of R .

If X and Y are curves over k and $f : Y \rightarrow X$ is a morphism, then we get a place f^* from $k(X)$ to $k(Y)$ by considering the map $\mathcal{O}_{X,f(\xi_Y)} \rightarrow \mathcal{O}_{Y,\xi_Y} = k(Y)$. It is a place because it is a local homomorphism of local rings.

Lemma 2.8. *Let X and Y be curves over k and $f : Y \rightarrow X$ a morphism, then the following are equivalent*

1. f sends the generic point of Y to the generic point of X ;
2. f^* is a morphism of fields;
3. f is a dominant morphism, i.e., the image of f is dense in X ;
4. f is non-constant.

Proof. It is clear that 1 and 2 are equivalent, by the definition of f^* and the fact that the valuation ring corresponding to f^* uniquely determines the point $f(\xi_X)$.

The generic point of X is dense in X , so 1 implies 3. Suppose that $f(\xi_Y) = a$ is not the generic point of X . Then $f^{-1}(a)$ is a closed subset of Y containing ξ_Y , so it is all of Y and f is constant. This shows that 4 implies 1.

If f is dominant and constant, then it must send every point of Y to the generic point of X . In particular, we see that 1 holds in this case, so 2 holds. So f^* is a morphism of fields. We view $k(X)$ as a subfield of $k(Y)$ via f^* . Then $k(Y)$ is a finite (hence algebraic) extension of $k(X)$, as they both have transcendence degree 1 over k and $k(Y)$ is finitely generated over k . We conclude that the integral closure of $k(X)$ in $k(Y)$ is all of $k(Y)$, but this is the intersection of all the valuation rings of points that map to ξ_X , that is, all of Y . Contradiction, so 3 implies 4. \square

Lemma 2.9. *Let X and Y be curves over k and let f and g be morphisms from Y to X such that $f^* = g^*$. Then $f = g$.*

Proof. We begin by showing that f and g induce the same map on the underlying topological spaces. If f^* (and therefore g^*) is not a morphism of fields, then f and g are constant. The point that they map to is determined by the valuation ring on which f^* and g^* are defined. So f and g are the same in this case.

If $j = f^* = g^*$ is a morphism of fields, then the image of a point $y \in Y$ can be determined from j as follows. $j^{-1}R_y$ is a valuation ring of $k(X)$, so there is at most one point $x \in X$ such that $R_x = j^{-1}R_y$, or equivalently, such that $jR_x \subset R_y$. Since this must in particular hold for $x = f(y)$ and for $x = g(y)$, we conclude that f and g are the same.

Lastly we observe that the map $f_y^\# : \mathcal{O}_{X,f(y)} \rightarrow \mathcal{O}_{Y,y}$ can be obtained from f^* by restriction. This means that these maps are the same for f and g also. It now follows that f and g are the same as morphisms of schemes. \square

Lemma 2.10. *Let X and Y be curves over k with X complete and Y regular. Let $j : k(X) \rightarrow k(Y)$ be a morphism of k -algebras. Then there is a unique morphism $f : Y \rightarrow X$ such that $f^* = j$.*

Proof. From the previous lemma we see that f is unique if it exists. We now consider pairs (U, f_U) where U is an open of Y and f_U is a morphism of curves $U \rightarrow X$ such that $f_U^* = j$. If we have two such pairs (U, f_U) and (V, f_V) , then we also have a pair $(U \cup V, f_{U \cup V})$. This is because on the intersection of U and V , f_U and f_V coincide by the previous lemma.

So we are done if we can find for every $y \in Y$ an open neighbourhood where we can define the map f . Let $y \in Y$ and let R_y be the valuation ring corresponding to y . Let $x \in X$ be the point corresponding to the valuation ring $j^{-1}R_y$. Here we need that X is complete. Let $\text{Spec}(A)$ be an affine open of X such that $x \in \text{Spec}(A)$. Pick generators x_1, \dots, x_t of A as an algebra over k .

Let $a \in k(Y)$ nonzero and let $\text{Spec}(B) \subset Y$ be a non-empty affine open. We can write $a = s/t$ with s, t in B . Now there are only finitely many prime ideals of B such that t lies inside that prime ideal. As we can cover Y by finitely many such $\text{Spec}(B)$, we see that there are only finitely many points z of Y such that $a \notin R_z$. So the set $U_a = \{z \in Y \mid a \in R_z\}$ is open and a is in $\Gamma(U_a, \mathcal{O}_Y)$ by lemma 2.7.

We apply this to $j(x_1), \dots, j(x_t)$. Let U be the intersection of the finitely many opens we find. Note that $y \in U$, as all the x_i are in $j^{-1}R_y$. Let $\text{Spec}(B)$ be an open affine subset of U containing y . Then $j(x_1), \dots, j(x_t)$ are in B , so j maps A into B . We conclude that we have a morphism $f_{\text{Spec}(B)} : \text{Spec}(B) \rightarrow \text{Spec}(A)$ such that $f_{\text{Spec}(B)}^* = j$ and $y \in \text{Spec}(B)$. \square

Corollary 2.11. *If two complete, regular curves over k have isomorphic function fields, they are isomorphic.*

Proof. Apply the previous lemma a few times to the isomorphism, its inverse and the identity on either side. \square

Lemma 2.12. *Let X be a complete curve over k and ℓ a finite extension of $k(X)$. Then there is a unique complete and regular curve \tilde{X} over k such that the function field of \tilde{X} is ℓ . Moreover, the morphism $\tilde{X} \rightarrow X$ corresponds to the inclusion $k(X) \subset \ell$ and is finite. We call \tilde{X} the normalisation of X in ℓ .*

Proof. Let $\text{Spec}(A) \subset X$ be a non-empty affine open. Put \tilde{A} the integral closure of A in ℓ . Now $\text{Spec}(\tilde{A})$ is regular and the map to $\text{Spec}(A)$ is finite. Moreover, the maps are compatible by lemma 2.9. Therefore they glue to a scheme \tilde{X} which has the required properties. \square

We consider the category \mathcal{C}_k of complete regular curves over k with dominant morphisms.

Theorem 2.13. *Let $X \in \mathcal{C}_k$ and define \mathcal{C}_X as the category of curves $Y \in \mathcal{C}_k$ together with a morphism $Y \rightarrow X$. Let $\mathcal{F}_{k(X)}$ be the category of finitely generated fields ℓ of transcendence degree 1 over k with a k -algebra morphism $k(X) \rightarrow \ell$. Then there is an anti-equivalence of categories between \mathcal{C}_X and $\mathcal{F}_{k(X)}$. The functors are taking the function field in one direction and taking the normalisation of X in the other direction.*

Proof. This is now more or less immediate. Taking the function field is a functor from the curves to the fields and normalising X gives a functor in the other direction. \square

Corollary 2.14. *All maps in \mathcal{C}_k are finite. All curves in \mathcal{C}_k are projective over k .*

Proof. For the first part, let $f : Y \rightarrow X$ be such a morphism and apply the previous theorem to X . We see that Y is isomorphic to the normalisation \tilde{X} of X in $k(Y)$ and that f corresponds to the natural map from \tilde{X} to X , which is finite. For the second part, note that all finite maps are projective. Now let α be a transcendental element of $k(X)$. Then there is a morphism of fields $k(\alpha) \subset k(X)$ and $k(\alpha)$ is the function field of the projective line, which is projective over k . The composition of projective morphisms is again projective, so X is projective over k . \square

3 Endomorphisms of the Additive Group

- **Group Schemes**

Let ℓ be a field and S a scheme over ℓ . Then S represents a contravariant functor from commutative ℓ -algebras to sets, sending an algebra R to the set $S(R)$ of R -points of S . From the Yoneda lemma and the fact that schemes are everywhere locally affine, the schemes form a full subcategory of the category of contravariant functors from commutative ℓ -algebras to sets.

A ℓ -rational point p in $S(\ell)$ gives rise to a compatible choice of a point in $S(R)$ for any commutative ℓ -algebra R . This means that the pair (S, p) represents a functor to pointed sets.

We want to have a type of object that represents a functor from commutative ℓ -algebras to *groups*. A group, as is well-known, is a set G with a distinguished *unit* element e , a *multiplication* map $m : G \times G \rightarrow G$ and an *inverse* map $i : G \rightarrow G$ satisfying a few conditions. This leads us to the following definition.

Definition 3.1. Let ℓ be a field. A *group scheme* over ℓ consists of a scheme G over ℓ , a ℓ -rational point e and two maps of ℓ -schemes $m : G \times_{\ell} G \rightarrow G$ and $i : G \rightarrow G$, such that together they represent a functor from commutative ℓ -algebras to groups. A morphism of group schemes is a morphism of functors.

Example 3.1. The additive group \mathbf{G}_a represents the forgetful functor sending a ℓ -algebra to its underlying additive group. As a scheme, \mathbf{G}_a is the affine line over ℓ .

Example 3.2. The multiplicative group \mathbf{G}_m represents the functor that sends a ℓ -algebra to its unit group. As a scheme, \mathbf{G}_m is the spectrum of $\ell[x, y]/(xy - 1)$.

Example 3.3. Let E be an elliptic curve over ℓ . Then the formulas obtained from the chord-and-tangent process allow us to specify for every commutative ℓ -algebra R a group structure on $E(R)$. Thus an elliptic curve is a group scheme.

Let G be a group scheme over ℓ . Let $\ell[\epsilon]$ be the commutative ℓ -algebra $\ell[x]/(x^2)$. Denote by ρ the map $G(\ell[\epsilon]) \rightarrow G(\ell)$ which corresponds to the ℓ -algebra map $\ell[\epsilon] \rightarrow \ell$ sending ϵ to 0.

Definition 3.2. The *tangent space* at e , $T_G(e)$ of G is the set of ϕ in $G(\ell[\epsilon])$ such that $\rho\phi$ is the unit element e in $G(\ell)$.

Note that $T_G(e)$ is a ℓ -vector space. A morphism $f : G \rightarrow H$ of group schemes gives rise to a ℓ -linear map $T_f : T_G(e) \rightarrow T_H(e)$. This makes T into a functor.

Example 3.4. The tangent space of \mathbf{G}_a is a one-dimensional ℓ -vector space. The same is true for \mathbf{G}_m .

Note that the set $\text{End}_\ell(G)$ of endomorphism of a commutative group scheme G is in a natural way a ring. Composition of endomorphisms is the multiplication and addition is done pointwise. The tangent functor we have just seen gives rise to a ring homomorphism $D : \text{End}_\ell(G) \longrightarrow \text{End}_\ell(T_G)$.

Let A be a commutative ring. Then an A -module is an Abelian group M together with a morphism $\phi : A \longrightarrow \text{End}(M)$. This motivates the following definition.

Definition 3.3. Let A be a commutative ring. An A -module scheme over ℓ is a commutative group scheme M over ℓ together with a ring homomorphism $\phi : A \longrightarrow \text{End}_\ell(M)$.

Note that such an object represents a functor from commutative ℓ -algebras to A -modules.

Example 3.5. Any commutative group scheme is in a natural way a \mathbf{Z} -module scheme.

Example 3.6. If E is an elliptic curve over \mathbf{C} with complex multiplication by some order \mathcal{O} in an imaginary quadratic number field, then E is an \mathcal{O} -module scheme over \mathbf{C} .

Note that an A -module scheme M gives rise to a ring homomorphism $D\phi$ from A to $\text{End}_\ell(T_M)$. This makes T_M into an A -module in a way that respects the existing ℓ -module structure.

• Additive Polynomials

Let ℓ be a field. We consider the additive group scheme \mathbf{G}_a over ℓ . As a scheme \mathbf{G}_a is isomorphic to $\text{Spec} \ell[x]$. The endomorphisms of this scheme are ℓ -algebra morphisms from $\ell[x]$ to itself. These are uniquely determined by the image of x and therefore correspond bijectively to the elements of $\ell[x]$. Such a polynomial $f \in \ell[x]$ induces for every commutative ℓ -algebra m an evaluation map $\mathbf{G}_a(m) \longrightarrow \mathbf{G}_a(m)$ sending λ to $f(\lambda)$.

In order for such a map to be an endomorphism of algebraic groups, we must have that $f(\lambda + \mu) = f(\lambda) + f(\mu)$ for all λ and μ in every commutative ℓ -algebra m . This is equivalent to demanding that $f(x + y) = f(x) + f(y)$ holds in $\ell[x, y]$. We call a polynomial that satisfies this equality an *additive polynomial*.

Let f be a polynomial in $\ell[x]$. Then its formal derivative f' is defined as follows. Let $\ell[\epsilon]$ be the ℓ -algebra $\ell[y]/(y^2)$. It is an ℓ -vector space of dimension two with basis $1, \epsilon$. We define f' to be the unique polynomial in $\ell[x]$ such that $f(x + \epsilon) = f(x) + f'(x)\epsilon$ holds in $\ell[\epsilon][x]$.

Lemma 3.4. *Let f be an additive polynomial in $\ell[x]$. Then f' constant polynomial.*

Proof. If f is additive, then in particular we must have in $\ell[\epsilon][x]$ that

$$f(x) + f(\epsilon) = f(x + \epsilon) = f(x) + f'(x)\epsilon.$$

If we write f as $\sum_{i=0}^n a_i x^i$ with the a_i in ℓ , then we see that $f(\epsilon) = a_0 + a_1 \epsilon$. We conclude that $a_0 = 0$ and $f'(x) = a_1$, that is, f' is a constant polynomial. \square

Corollary 3.5. *If ℓ has characteristic 0, the only additive polynomials are the scalar multiplications: λx with λ in ℓ .*

Theorem 3.6. *Let ℓ be a field of characteristic p . Then the ring of additive polynomials is naturally isomorphic to the skew polynomial ring $\ell\{\tau_p\}$ generated by ℓ and τ_p and satisfying the relation $\tau_p \lambda = \lambda^p \tau_p$ for all λ in ℓ . The isomorphism identifies τ_p with the additive polynomial x^p .*

Proof. Note that the polynomial x^p is indeed additive. We identify this polynomial with τ_p and λ in ℓ with the scalar multiplication λx , which is also an additive polynomial. The required relation is satisfied, so $\ell\{\tau_p\}$ is in a natural way a subring of the additive polynomials.

What remains to be shown is that all additive polynomials are actually elements of $\ell\{\tau_p\}$. That is, they are of the form $\sum_{i=0}^n \lambda_i x^{p^i}$ with the λ_i in ℓ .

The requirement that must be satisfied for a polynomial f to be additive,

$$f(x + y) = f(x) + f(y)$$

in $\ell[x, y]$, is homogeneous. The term of degree n in f gives rise to the terms with total degree n in the above equation. So it suffices to check which monomials satisfy this relation. Let n be a positive integer and write $n = p^r m$ with m non divisible by p . Then we see that

$$(x + y)^n = (x^{p^r} + y^{p^r})^m = (x^{p^r})^m + m (x^{p^r})^{m-1} (y^{p^r})^1 + \text{lower order terms in } x$$

holds in $\mathbf{F}_p[x, y]$. It is therefore necessary that $m = 1$ if this monomial is to satisfy the required relation. We conclude that all additive polynomials are indeed in $\ell\{\tau_p\}$. \square

Let ℓ be a field that contains \mathbf{F}_q . Then we have a natural way to turn \mathbf{G}_a over ℓ into an \mathbf{F}_q -module scheme. We call a polynomial in $\ell[x]$ an \mathbf{F}_q -linear polynomial if it induces an endomorphism of \mathbf{G}_a as an \mathbf{F}_q -module scheme.

Lemma 3.7. *Let ℓ be a field containing \mathbf{F}_q . The ring of \mathbf{F}_q -linear polynomials over ℓ is the subring of $\ell\{\tau_p\}$ generated by ℓ and $\tau_q = x^q$. We write $\ell\{\tau_q\}$ for this ring.*

Proof. An \mathbf{F}_q -linear polynomial f satisfies for all ζ in \mathbf{F}_q the relation $f(\zeta x) = \zeta f(x)$ in $\ell[x]$. Note that f is also an additive polynomial, so that only monomials of degree x^{p^i} occur. Comparing the monomial of degree p^i on either side we conclude that $\zeta a_i = a_i \zeta^{p^i}$ must hold for all ζ in \mathbf{F}_q . So, either a_i is 0, or p^i is a power of q , which is what we wanted to prove. \square

- **Basic Properties of Additive Polynomials**

There is another way to characterise the separable polynomials that are additive, using their roots.

Lemma 3.8. *Let f be a non-zero separable polynomial in $\ell[x]$ and denote by Z the set of roots of f in some fixed separable closure $\bar{\ell}$ of ℓ . Then f is additive if and only if Z is an additive subgroup of $\bar{\ell}$. If ℓ contains \mathbf{F}_q then f is \mathbf{F}_q -linear if and only if Z is a vector space over \mathbf{F}_q .*

Proof. The implications in one direction are clear. If f is additive (\mathbf{F}_q -linear) then Z is a subgroup (vector space over \mathbf{F}_q .)

Let f be a separable polynomial whose zero-set Z is an additive subgroup. Let c be the leading coefficient of f . Then we know that $f = c \prod_{z \in Z} (x - z)$ holds in $\bar{\ell}[x]$. Let λ be in $\bar{\ell}$. Note that for all μ in Z we have

$$f(\lambda + \mu) = c \prod_{z \in Z} (\lambda + (\mu - z)) = c \prod_{z \in Z} (\lambda - z) = f(\lambda) = f(\lambda) + f(\mu).$$

Therefore, the polynomials $f(\lambda + x)$ and $f(\lambda) + f(x)$ coincide on Z . Moreover, their degree is $\#Z$ and they have the same leading coefficient, so in fact $f(\lambda + x) = f(\lambda) + f(x)$. We have this for all λ in $\bar{\ell}$. We conclude that f is additive.

Suppose that Z is also a vector space over \mathbf{F}_q . Then for all ζ in \mathbf{F}_q we have

$$f(\zeta x) = c \prod_{z \in Z} (\zeta x - z) = \zeta c \prod_{z \in Z} (x - \zeta^{-1}z) = \zeta c \prod_{z \in Z} (x - z) = \zeta f(x).$$

So f is \mathbf{F}_q -linear. □

By the degree of an element in $\ell\{\tau_q\}$ we mean its degree as a polynomial in τ_q , not in x . That is, if we can write our polynomial as $\sum_{i=0}^n \lambda_i \tau_q^i$ with λ_n non-zero, then the degree is n . The degree of this polynomial in x is q^n . Note that the degree satisfies

$$\deg(fg) = \deg(f) + \deg(g)$$

and the inequality

$$\deg(f + g) \leq \max\{\deg(f), \deg(g)\}$$

for any f and g non-zero, with equality if $\deg(f)$ and $\deg(g)$ are not the same.

Lemma 3.9. *The ring $\ell\{\tau_q\}$ has a right Euclidean algorithm. That is, for any two additive polynomials f and g with g non-zero there are unique additive polynomials d and r such that $f = dg + r$ and the degree of r is less than that of g .*

Proof. The proof works in exactly the same way as the normal proof of this fact for polynomials. We proceed by induction on the degree of f . If this degree is less than that of g , we see that the only possibility for d is 0 and thus $r = f$.

Now suppose $\deg(f) = m$ and the claim holds for all polynomials of degree less than m . Let n be the degree of g . Let λ and μ be the leading coefficients of f and g respectively. Then the polynomial

$$\tilde{f} = f - (\lambda\mu^{-p^{m-n}}\tau_q^{m-n})g$$

has degree less than m , as $\lambda - \lambda\mu^{-q^{m-n}}\mu^{q^{m-n}} = 0$. Now let \tilde{d} and \tilde{r} be the unique polynomials such that $\tilde{f} = \tilde{d}g + \tilde{r}$. Putting $d = \lambda\mu^{-q^{m-n}}\tau_p^{m-n} + \tilde{d}$ and $r = \tilde{r}$ we see that $f = dg + r$. Comparing leading coefficients, one sees that d and r are in unique. \square

Corollary 3.10. *All left ideals of $\ell\{\tau_p\}$ are principal.*

There is a natural map $i : \ell \longrightarrow \ell\{\tau_p\}$ sending λ to the polynomial $\lambda\tau_p^0$. The derivative of any additive polynomial is constant, which gives us a map $D : \ell\{\tau_p\} \longrightarrow \ell$ in the opposite direction. Note that Di is the identity on ℓ . The invertible elements of $\ell\{\tau_p\}$ must be polynomials of degree 0, so we have $\ell\{\tau_p\}^\times = i(\ell^\times)$.

• Additive Power Series

Definition 3.11. A formal power series $e(z)$ in $\ell[[z]]$ is called *additive* if we have $e(z+w) = e(z) + e(w)$ in $\ell[[z, w]]$. It is called \mathbf{F}_q -linear if in addition we have $e(\zeta z) = \zeta e(z)$ in $\ell[[z]]$ for all ζ in \mathbf{F}_q .

Most things we said about additive and \mathbf{F}_q -linear polynomials carry over to power series with only the obvious modifications. In fact, the proof of lemma 3.4 and theorem 3.6 are written in such a way that we can just replace polynomials by power series everywhere and they still hold.

Proposition 3.12.

- *The formal derivative of an additive power series is constant.*
- *A power series is additive if and only if the only monomials that appear have degree a power of p . It is \mathbf{F}_q -linear if and only if all the monomials that appear have degree a power of q .*
- *The composition ring of additive (resp. \mathbf{F}_q -linear) power series is naturally isomorphic to the skew power series ring $\ell\{\{\tau\}\}$ generated by ℓ and τ , satisfying the relation $\tau\lambda = \lambda^p\tau$ (resp. $\tau\lambda = \lambda^q\tau$) for all λ in ℓ .*

Theorem 3.13. *Let σ be an \mathbf{F}_q -linear power series, whose derivative $s = D(\sigma)$ is transcendental over \mathbf{F}_q . Then there is a unique \mathbf{F}_q -linear power series e_σ such that $D(e_\sigma) = 1$ and $e_\sigma s = \sigma e_\sigma$.*

Proof. Write $\sigma = \sum s_i \tau^i$, so that $s = s_0$. We first show uniqueness. Suppose $e_\sigma = \sum e_i \tau^i$ works. Then comparing the coefficients of τ^i in the relation $e_\sigma s = \sigma e_\sigma$ we see

$$s^{q^i} e_i = \sum_{j=0}^i s_j e_{i-j}^{q^j}.$$

This can be rewritten as

$$(s^{q^i} - s) e_i = \sum_{j=1}^i s_j e_{i-j}^{q^j},$$

which expresses a non-zero (as s is transcendental of \mathbf{F}_q) multiple of e_i in terms of the e_j with j less than i and the (known) s_j . As $e_0 = 1$ is fixed in the requirements for e_σ , we see that the coefficients of e are uniquely determined.

For the existence let $e_\sigma = \sum e_i \tau^i$ with the e_i given by the relation we have just determined. This \mathbf{F}_q -linear power series works. \square

Corollary 3.14. *Let σ be an \mathbf{F}_q -linear power series whose derivative $s = D(\sigma)$ is transcendental over \mathbf{F}_q and let r be an element of ℓ . Then $\rho = e_\sigma r e_\sigma^{-1}$ is the unique \mathbf{F}_q -linear power series with $D(\rho) = r$ that commutes with σ .*

4 Drinfeld Modules

Let X be a complete, regular curve over \mathbf{F}_p . Let \mathbf{F}_q be its field of constants, where $q = p^m$, and let k be its function field. We consider X as a curve over \mathbf{F}_q . Fix a closed point ∞ of X . Now $X - \infty$ is affine, say $\text{Spec} A$. We know that A is a Dedekind domain, with field of fractions k . Its class group is finite and we have $A^\times = \mathbf{F}_q^\times$.

- **Definition**

Definition 4.1. Let ℓ be any field containing \mathbf{F}_q and δ be an \mathbf{F}_q -algebra map $A \rightarrow \ell$. A *Drinfeld A -module* over ℓ is a ring homomorphism $\phi : A \rightarrow \ell\{\tau_p\}$ such that $D\phi = \delta$. We exclude the trivial case $\phi = i\delta$. A morphism of Drinfeld modules from ϕ to ψ is an additive polynomial f such that $f\phi(x) = \psi(x)f$ for all x in A .

As we have seen before, the ring $\ell\{\tau_p\}$ is the ring of endomorphisms of the additive group over ℓ . Giving a ring homomorphism $\phi : A \rightarrow \ell\{\tau_p\}$ therefore is the same as giving an A -module scheme structure on \mathbf{G}_a over ℓ . In other words, we give an A -module structure on the additive group of all commutative ℓ -algebras in a functorial manner.

Since A is a commutative ring, the image of ϕ is a commutative subring. As ϕ maps A^\times into $(\ell\{\tau_p\})^\times = \ell^\times$, all the elements of \mathbf{F}_q map to scalar multiplications and therefore, as δ is \mathbf{F}_q -linear, to themselves. Combining these two facts, we note that ϕ actually lands in the subring $\ell\{\tau_q\}$ of \mathbf{F}_q -linear polynomials. Also, any morphism of Drinfeld modules must be \mathbf{F}_q -linear.

From now on we consider all Drinfeld modules as ring homomorphism to $\ell\{\tau\} = \ell\{\tau_q\}$ and also consider all morphisms as elements of $\ell\{\tau\}$. The degree of an element of $\ell\{\tau\}$ is its degree as a polynomial in τ .

Lemma 4.2. *Let ϕ be a Drinfeld module over ℓ . Then ϕ is injective.*

Proof. The ring $\ell\{\tau\}$ has no zero divisors other than 0 itself, so ϕA is a domain. Therefore, the kernel of ϕ is a prime ideal of A . If it is a maximal prime ideal, then ϕA is a field. The largest field contained in $\ell\{\tau\}$ is ℓ . But this implies that $\phi = i\delta$, which we have excluded. As A is a Dedekind domain, the only non-maximal prime ideal is (0) , so ϕ is injective. \square

- **The Rank and Height of a Drinfeld Module**

Let ϕ be a Drinfeld module over ℓ . Then the map v sending a non-zero element a of A to $-\deg(\phi(a))$ satisfies $v(ab) = v(a) + v(b)$ and $v(a + b) \geq \min\{v(a), v(b)\}$. Therefore, v is a valuation of ℓ . It is non-trivial since $\phi \neq i\delta$. Moreover, it is non-positive for all elements of A . Therefore it corresponds to the point ∞ of X . The theory of valuations now tells us that there is a non-negative rational number r such that $-\deg(\phi(a)) = r \cdot \deg(\infty) \text{ord}_\infty(a)$ holds for all $a \in A$ non-zero. We call this number the rank of the module ϕ .

Let L be an algebraically closed field containing ℓ . Any Drinfeld module ϕ over ℓ induces an A module structure on L , which we shall denote by Φ . Note that Φ is divisible as an

A -module, i.e., the action of every non-zero element of A is surjective. This follows from lemma 4.2 and the fact that any non-zero polynomial with coefficients in ℓ has a root in L .

Let \mathfrak{c} be the kernel of δ . It is a prime ideal of A , possibly (0) . Any element a of A that is not divisible by \mathfrak{c} acts on Φ by a polynomial whose derivative is a non-zero constant, so in particular the polynomial is separable. Therefore, $\phi(a)$ has $q^{\deg(\phi(a))}$ distinct roots in L .

Let \mathfrak{a} be a non-zero ideal of A . Then the \mathfrak{a} -torsion submodule $\Phi[\mathfrak{a}]$ of Φ is the set of elements λ of L such that $a\lambda = 0$ for all a in \mathfrak{a} . These are precisely the elements of L that are roots of all the polynomials in the left ideal generated by the $\phi(a)$ with a in \mathfrak{a} . Note that this ideal is principal and write $\phi(\mathfrak{a})$ for its monic generator. Then $\Phi[\mathfrak{a}]$ is the zero set of the polynomial $\phi(\mathfrak{a})$.

Lemma 4.3. *Let \mathfrak{p} be a prime ideal of A and π a local uniformiser at \mathfrak{p} . Then there is a positive integer t , depending only on \mathfrak{p} and for every positive integer e an isomorphism f_e from $(A/\mathfrak{p}^e)^t$ to $\Phi[\mathfrak{p}^e]$ such that the diagram*

$$\begin{array}{ccc} \Phi[\mathfrak{p}^{e+1}] & \xrightarrow{\pi} & \Phi[\mathfrak{p}^e] \\ \uparrow f_{e+1} & & \uparrow f_e \\ (A/\mathfrak{p}^{e+1})^t & \longrightarrow & (A/\mathfrak{p}^e)^t \end{array}$$

commutes.

Proof. We prove these claims by induction on e . For $e = 1$, note that $\Phi[\mathfrak{p}]$ is a finite torsion A -module which is annihilated by \mathfrak{p} . By the structure theory of finitely generated A -modules there is t such that $\Phi[\mathfrak{p}]$ is isomorphic to $(A/\mathfrak{p})^t$. Fix such an isomorphism f_1 from $(A/\mathfrak{p})^t$ to $\Phi[\mathfrak{p}]$ and let $f_1^{(1)}$ up to $f_1^{(t)}$ be the the images of the standard generators of $(A/\mathfrak{p})^t$.

Suppose now that the claim holds for a certain e . Now pick elements $f_{e+1}^{(1)}$ up to $f_{e+1}^{(t)}$ in Φ such that we have $f_{e+1}^{(i)} = \pi f_e^{(i)}$ for all i . This can be done since Φ is a divisible A -module. Note that for all i , $f_{e+1}^{(i)}$ is in $\Phi[\mathfrak{p}^{e+1}]$ and that we have $\pi^e f_{e+1}^{(i)} = f_1^{(i)}$. Let f_{e+1} be the A -module morphism sending the standard generators of $(A/\mathfrak{p}^{e+1})^t$ to $f_{e+1}^{(1)}$ up to $f_{e+1}^{(t)}$. By construction, this map fits into the commutative diagram from the lemma. Moreover, we have the following large commutative diagram with exact rows.

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \Phi[\mathfrak{p}^e] & \longrightarrow & \Phi[\mathfrak{p}^{e+1}] & \xrightarrow{\pi^e} & \Phi[\mathfrak{p}] & \longrightarrow & 0 \\ & & \uparrow f_e & & \uparrow f_{e+1} & & \uparrow f_1 & & \\ 0 & \longrightarrow & (A/\mathfrak{p}^e)^t & \xrightarrow{\pi} & (A/\mathfrak{p}^{e+1})^t & \longrightarrow & (A/\mathfrak{p})^t & \longrightarrow & 0 \end{array}$$

Since f_e and f_1 are isomorphisms, we conclude that f_{e+1} is also an isomorphism. \square

Theorem 4.4. *Let ϕ be a Drinfeld module over ℓ of rank r . Then r is a positive integer and for any non-zero ideal \mathfrak{a} of A not divisible by \mathfrak{c} we have $\Phi[\mathfrak{a}] \cong (A/\mathfrak{a})^r$.*

Proof. Let \mathfrak{p} be a maximal ideal of A different from \mathfrak{c} . Let $t = t_{\mathfrak{p}}$ be the positive integer such that we have $\Phi[\mathfrak{p}^e] \cong (A/\mathfrak{p}^e)^t$ for all positive integers e .

In particular, if we let h be the class number of A and a a generator of \mathfrak{p}^h then we see that $\Phi[a] \cong (A/a)^t$. Comparing the number of elements on either side we see that

$$q^{\deg(\phi(a))} = \#\ker(\phi(a)) = \#\Phi[a] = (\#(A/a))^t.$$

We use here that $\phi(a)$ is a separable polynomial, as \mathfrak{p} is not \mathfrak{c} . From the product formula we have that $\deg(\mathfrak{p})\text{ord}_{\mathfrak{p}}(a) = -\deg(\infty)\text{ord}_{\infty}(a)$. From this we see

$$\#(A/a) = q^{\deg(\mathfrak{p})\text{ord}_{\mathfrak{p}}(a)} = q^{-\deg(\infty)\text{ord}_{\infty}(a)}$$

We conclude that $\deg(\phi(a)) = -t \deg(\infty)\text{ord}_{\infty}(a)$, so $r = t$. We conclude that t is independent of \mathfrak{p} and that r is a positive integer.

The fact that we have $\Phi[\mathfrak{a}] \cong (A/\mathfrak{a})^t$ for every non-zero ideal now follows from the fact we know this for all powers of prime ideals, using the Chinese remainder theorem. \square

Corollary 4.5. *Let ϕ be a Drinfeld module over ℓ of rank r and let \mathfrak{a} be a non-zero ideal of A not divisible by \mathfrak{c} . Then we have*

$$\deg(\phi(\mathfrak{a})) = -r \deg(\infty)\text{ord}_{\infty}(\mathfrak{a}) = r \sum_{\mathfrak{p}} \deg(\mathfrak{p})\text{ord}_{\mathfrak{p}}(\mathfrak{a}),$$

where the sum runs over the non-zero primes \mathfrak{p} of A .

Let x be a non-zero element of A . Put $j(x)$ the smallest integer k such that the coefficient of τ^k in $\phi(x)$ is non-zero. Note that we have $j(xy) = j(x)j(y)$ and $j(x+y) \geq \min(j(x), j(y))$ for all x and y in A non-zero. We see that j is a valuation on k . If δ is injective, j is the trivial valuation, otherwise it corresponds to the prime $\mathfrak{c} = \ker(\delta)$ of A . In this case there is a positive rational number h such that $j(x) = h \deg(\mathfrak{c})\text{ord}_{\mathfrak{c}}(x)$ for all non-zero x in A . We call this number the height of ϕ . If δ is injective, we say that ϕ has height 0.

Lemma 4.6. *Let ϕ be a Drinfeld module over ℓ and suppose that $\mathfrak{c} = \ker(\delta)$ is non-zero. Then the height h of ϕ is a positive integer and we have $\Phi[\mathfrak{c}^e] \cong (A/\mathfrak{c}^e)^{r-h}$ for all positive integers e .*

Proof. By lemma 4.3 we see that $\Phi[\mathfrak{c}^e] \cong (A/\mathfrak{c}^e)^t$ for some positive integer t that does not depend on e .

When we take e equal to the class number of A , \mathfrak{c}^e is a principal ideal. Let c be a generator of this ideal. Then $\Phi[\mathfrak{c}^e]$ is equal to the set of roots of $\phi(c)$ in $\bar{\ell}$. Therefore we have

$$\#\Phi[\mathfrak{c}^e] = q^{\deg(\phi(c)) - j(c)}.$$

We know already that

$$\deg(\phi(c)) = -r \deg(\infty)\text{ord}_{\infty}(c) = r \deg(\mathfrak{c})\text{ord}_{\mathfrak{c}}(c) = re \deg(\mathfrak{c})$$

holds. Also, by definition of the height, we have

$$j(c) = h \deg(\mathfrak{c})\text{ord}_{\mathfrak{c}}(c) = he \deg(\mathfrak{c}).$$

Now when we compare the number of elements on either side in $\Phi[\mathfrak{c}^e] = (A/\mathfrak{c}^e)^t$, we conclude that $q^{re \deg(\mathfrak{c}) - he \deg(\mathfrak{c})} = q^{te \deg(\mathfrak{c})}$, so $r - h = t$ as required. \square

- **Morphisms of Drinfeld Modules**

Let ϕ and ψ be two Drinfeld modules over ℓ . Recall that a morphism from ϕ to ψ is an element f of $\ell\{\tau\}$ such that $\phi(x)f = f\psi(x)$ holds for all x in A . Comparing degrees we see that the only morphism between Drinfeld modules of different ranks is the zero morphism. The same also holds for Drinfeld modules of different height.

An isomorphism of Drinfeld modules must be an invertible element of $\ell\{\tau\}$, so it is a scalar multiplication by a non-zero element of ℓ .

Lemma 4.7. *Let ϕ and ψ be two Drinfeld modules over ℓ . Then $\text{Hom}(\phi, \psi)$ is a torsion-free A -module.*

Proof. Note that for all a in A , $\phi(a)$ is an endomorphism of ϕ , as $\phi(a)$ commutes with $\phi(x)$ for every x in A . Write $[a]$ for this endomorphism. We know that ϕ is injective, so $[a]$ is non-zero for all non-zero a in A .

We define the A -module structure on $\text{Hom}(\phi, \psi)$ as follows. Let a be in A and f be in $\text{Hom}(\phi, \psi)$. Then we put af equal to $[a] \circ f$. Suppose that we have $af = 0$ for some a in A and f in $\text{Hom}(\phi, \psi)$. Then either $[a] = 0$ or $f = 0$ must hold, as the ring $\ell\{\tau\}$ has no zero-divisors. Note that $[a] = 0$ implies $a = 0$. We conclude that $\text{Hom}(\phi, \psi)$ is a torsion free A -module. \square

Definition 4.8. Let \mathfrak{p} be a prime ideal of A . Then we define the *Tate module* of ϕ at \mathfrak{p} to be the inverse limit

$$T_{\mathfrak{p}}(\phi) = \varprojlim \Phi[\mathfrak{p}^e].$$

Lemma 4.9. *$T_{\mathfrak{p}}$ is a covariant functor from the category of Drinfeld modules over ℓ to the category of free $A_{\mathfrak{p}}$ modules. It is injective on Hom 's. If ϕ is a Drinfeld module of rank r and height h then $T_{\mathfrak{p}}(\phi)$ has rank r if \mathfrak{p} is different from $\ker(\delta)$ and $r - h$ if \mathfrak{p} is equal to $\ker(\delta)$.*

Proof. It follows at once from lemma 4.3 that the Tate module is a free $A_{\mathfrak{p}}$ module of rank t . In theorem 4.4 and lemma 4.6, we computed that we have $t = r$ if \mathfrak{p} is not $\ker(\delta)$ and $t = r - h$ if \mathfrak{p} is $\ker(\delta)$.

Let f be a morphism from ϕ to ψ . Then f gives rise to compatible morphisms from $\Phi[\mathfrak{p}^e]$ to $\Psi[\mathfrak{p}^e]$ for every positive integer e . Thus we get a map $T_{\mathfrak{p}}(f)$ from $T_{\mathfrak{p}}(\phi)$ to $T_{\mathfrak{p}}(\psi)$. It is clear that this construction makes $T_{\mathfrak{p}}$ into a covariant functor. Moreover, the map

$$T_{\mathfrak{p}} : \text{Hom}(\phi, \psi) \longrightarrow \text{Hom}_{A_{\mathfrak{p}}}(T_{\mathfrak{p}}(\phi), T_{\mathfrak{p}}(\psi))$$

is a morphism of A -modules. To check it is injective, we note that $T_{\mathfrak{p}}(f) = 0$ implies f induces the zero map on $\Phi[\mathfrak{p}^e]$ for every positive integer e . So the zero set of the polynomial f contains arbitrarily large sets $\Phi[\mathfrak{p}^e]$. We conclude that f is zero. \square

Theorem 4.10. *Let \mathfrak{p} be a prime ideal different from $\ker(\delta)$. Then the natural map*

$$\mathrm{Hom}(\phi, \psi) \otimes_A A_{\mathfrak{p}} \longrightarrow \mathrm{Hom}_{A_{\mathfrak{p}}}(T_{\mathfrak{p}}(\phi), T_{\mathfrak{p}}(\psi))$$

is injective.

Proof. We follow Silverman's proof in [The Arith. of E.C., III thm 7.4]. Note that we may assume ϕ and ψ have the same rank, r . Let M be a finitely generated sub- A -module of $\mathrm{Hom}(\phi, \psi)$. We show that

$$M_d = \{f \in \mathrm{Hom}(\phi, \psi) : af \in M \text{ for some } a \text{ in } A\}$$

is again finitely generated. Note that M_d injects naturally into the finite dimensional k -vector space $V = M \otimes_A k$, as $\mathrm{Hom}(\phi, \psi)$ is torsion-free. Moreover as k is a localisation of A (at the zero-ideal) all elements of V are pure tensors.

Let $N : k \longrightarrow \mathbf{Q}$ be the map sending x to $q^{-r \deg(\infty) \mathrm{ord}_{\infty}(x)}$. Note that it is a non-archimedean norm on k corresponding to the valuation at ∞ . We can extend N to V by putting

$$\begin{aligned} N : M \otimes_A k &\longrightarrow \mathbf{Q} \\ f \otimes x &\longmapsto q^{\deg(f)} N(x) \end{aligned}$$

Note that this map is well-defined as we have

$$q^{\deg(af)} N(x) = q^{\deg(f)} q^{\deg(\psi(a))} N(x) = q^{\deg(f)} q^{-r \deg(\infty) \mathrm{ord}_{\infty}(a)} N(x) = q^{\deg(f)} N(ax).$$

Clearly, N is a norm on V as a k -vector space. Let U be the open ball $U = \{f \in V : N(f) < 1\}$. Then we see that $M_d \cap U$ contains only 0, as any other element of M_d is an f in $\mathrm{Hom}(\phi, \psi)$ and $N(f) = q^{\deg(f)}$ is at least 1. We conclude that M_d is a discrete A -module inside the finite dimensional k -vector space V and therefore is finitely generated. Moreover, M_d is torsion free, so it is a finitely generated projective A -module.

Now suppose f be in $\mathrm{Hom}(\phi, \psi) \otimes A_{\mathfrak{p}}$ maps to 0. Pick M in $\mathrm{Hom}(\phi, \psi)$ finitely generated such that f is in $M \otimes A_{\mathfrak{p}}$. By the above, M_d is finitely generated projective. So $M_d \otimes A_{\mathfrak{p}}$ is free. Pick f_1, \dots, f_t in M_d such that they are a basis for $M_d \otimes A_{\mathfrak{p}}$. Write f as $\sum_i \alpha_i f_i$ with the α_i in $A_{\mathfrak{p}}$. As the kernel is torsion-free, we may assume that at least one of the α_i is invertible in $A_{\mathfrak{p}}$.

Recall that $\mathrm{Cl}(A)$ is finite, say of order h , so that \mathfrak{p}^h is a principal ideal. Let m be a generator of \mathfrak{p}^h . Pick a_1, \dots, a_t in A such that we have $\alpha_i \equiv a_i$ modulo m for all i . Put g equal to $\sum a_i f_i$ in M_d . By construction $T_{\mathfrak{p}}(g)$ is 0 modulo m , meaning that $\Phi[m]$ is contained in the kernel of g .

The \mathbf{F}_q -linear polynomial g need not be separable. Put V equal to its zero set and let g_0 be the monic separable \mathbf{F}_q -linear polynomial with this zero set. Let e be the minimal exponent for which the coefficient of g at τ^e is non-zero. Then g has a root of multiplicity q^e at 0. Moreover, as g is additive, all roots of g have the same multiplicity. The \mathbf{F}_q -linear polynomial $\tau^e g_0$ has the same roots as g with the same multiplicities. Thus they differ by a scalar multiplication.

Note that the zero set W of the separable \mathbf{F}_q -linear polynomial $\phi(m)$ is contained in V . They are both finite dimensional \mathbf{F}_q -vector spaces. As $\phi(m)$ is \mathbf{F}_q -linear, we see that $\phi(m)V$

is an \mathbf{F}_q -vector space of dimension $\dim V - \dim W$. Let h_0 be the separable \mathbf{F}_q -linear polynomial with this set as zero set. Then $h_0\phi(m)$ has zero set V and degree $q^{\dim V}$. We conclude that $h_0\phi(m)$ and g_0 differ by a scalar multiplication.

It follows that g can be written as $h\phi(m)$ for some \mathbf{F}_q -linear polynomial h . For any x in A we have $g\phi(x) = \psi(x)g$ and therefore

$$h\phi(x)\phi(m) = h\phi(m)\phi(x) = g\phi(x) = \psi(x)g = \phi(x)h\phi(m)$$

holds, so we have $h\phi(x) = \psi(x)h$ for all x in A and therefore h is in fact in $\text{Hom}(\phi, \psi)$.

By the definition of the A -action on this Hom-set, we conclude that g is mh for some h in $\text{Hom}(\phi, \psi)$. By construction of M_d we therefore have $h = \sum_i b_i f_i$ for some b_i in A . As the f_i are independent over A , it follows that we have $a_i = mb_i$ for all i . But the all the a_i are in mA and therefore, all the α_i are in mA_p . This contradicts the assumption that at least one of the α_i is invertible. \square

Corollary 4.11. *Let ϕ be a Drinfeld module of rank r . Then the endomorphism ring $\text{End}(\phi)$ of ϕ is a projective A -module of rank at most r^2 .*

Proof. As $\text{End}(\phi)$ is a torsion-free A -module, it has finite rank over A if and only if $\text{End}(\phi) \otimes_A A_p$ has finite rank over A_p . If this is the case then these ranks are equal. By the previous theorem, $\text{End}(\phi) \otimes_A A_p$ can be embedded into a free A_p module of rank r^2 . So $\text{End}(\phi)$ has finite rank at most r^2 and being torsion-free this implies it is projective. \square

Corollary 4.12. *Let ϕ be a Drinfeld module of rank 1. Then $\text{End}(\phi)$ is isomorphic to A and therefore the natural map $\mathbf{F}_q^\times \rightarrow \text{Aut}(\phi)$ is an isomorphism.*

• Formal Drinfeld Modules

Suppose in this section that $\delta : A \rightarrow \ell$ is injective. It then extends to an inclusion of fields $\delta : k \rightarrow \ell$.

Definition 4.13. A formal Drinfeld module over ℓ is a ring homomorphism $\phi : k \rightarrow \ell\{\{\tau\}\}$ such that $D\phi = \delta$ and ϕ is a non-constant power series for some element of k .

Lemma 4.14. *Every Drinfeld module over ℓ can be extended uniquely to a formal Drinfeld module over ℓ .*

Proof. As $D\phi(a)$ is non-zero for all a in A , the image of $A - \{0\}$ under ϕ in $\ell\{\{\tau\}\}$ lands inside the multiplicative group. We can therefore extend ϕ to a ring homomorphism $k \rightarrow \ell\{\{\tau\}\}$ by putting $\phi(ba^{-1}) = \phi(b)\phi(a)^{-1}$ for all a and b in A with a non-zero. We then have $D\phi = \delta$ by construction of ϕ and δ . All non-constant elements of A give rise to non-constant power series. \square

Lemma 4.15. *Let ϕ be a formal Drinfeld module over ℓ . Then there is a unique power series e_ϕ such that we have $D(e_\phi) = 1$ and $\phi(x) = e_\phi \delta(x) e_\phi^{-1}$ holds for all x in k .*

Proof. Let x in k be non-constant. Then $\delta(x)$ is transcendental over \mathbb{F}_q in ℓ . By theorem 3.13 there is a unique power series e_ϕ such that $\phi(x) = e_\phi \delta(x) e_\phi^{-1}$. Moreover as for every y in k we have $\phi(x)\phi(y) = \phi(y)\phi(x)$, $\phi(y)$ is the unique power series with $D(\phi(y)) = \delta(y)$ that commutes with $\phi(x)$, so we have $\phi(y) = e_\phi \delta(y) e_\phi^{-1}$. \square

A morphism of formal Drinfeld modules from ϕ to ψ is an f in $\ell\{\{\tau\}\}$ such that $f\phi(x) = \psi(x)f$ for all x in k . Any morphism of Drinfeld modules is also a morphism between the formal Drinfeld modules they extend to.

Theorem 4.16. *Let ϕ be a formal Drinfeld module over ℓ . Then $D : \text{End}(\phi) \longrightarrow \ell$ is injective.*

Proof. Fix a non-constant element x of k . Then by theorem 3.13 and its corollary, a power series f such that $f\phi(x) = \phi(x)f$ is uniquely determined by $D(f)$. \square

Corollary 4.17. *Let $\delta : A \longrightarrow \ell$ be injective. Then for every Drinfeld module ϕ over ℓ , we have that $\text{End}(\phi)$ is commutative.*

5 Constructions with Drinfeld Modules

We continue with the notations for the previous section. In addition, we put K the completion of k at ∞ and C the completion of an algebraic closure of K at ∞ . C is an algebraically closed, complete field containing k .

• Action of the Ideals of A

Let \mathfrak{a} be any non-zero ideal of A . For any a in \mathfrak{a} and x in A we have that $\phi(a)$ and $\phi(x)$ commute. Therefore, right multiplication by $\phi(x)$ maps the left ideal generated by the $\phi(a)$ into itself. We conclude that there is a unique $\phi'(x)$ such that $\phi(a)\phi(x) = \phi'(x)\phi(a)$.

Lemma 5.1. *The map ϕ' defined above is a Drinfeld module of the same rank and height as ϕ . We write $\mathfrak{a} * \phi$ for ϕ' . The \mathbf{F}_q -linear polynomial $\phi(\mathfrak{a})$ is a non-zero morphism from ϕ to $\mathfrak{a} * \phi$.*

Proof. We verify that ϕ' is a ring homomorphism $A \longrightarrow \ell\{\tau\}$. First note that $\phi'(1) = 1$. Let x and y be elements of A . Note that we have

$$\begin{aligned} (\phi'(x) + \phi'(y))\phi(\mathfrak{a}) &= \phi'(x)\phi(\mathfrak{a}) + \phi'(y)\phi(\mathfrak{a}) = \phi(\mathfrak{a})\phi(x) + \phi(\mathfrak{a})\phi(y) \\ &= \phi(\mathfrak{a})(\phi(x) + \phi(y)) = \phi(\mathfrak{a})(\phi(x + y)) \\ &= \phi'(x + y)\phi(\mathfrak{a}) \end{aligned}$$

and

$$\begin{aligned} (\phi'(x)\phi'(y))\phi(\mathfrak{a}) &= \phi'(x)(\phi'(y)\phi(\mathfrak{a})) = \phi'(x)\phi(\mathfrak{a})\phi(y) \\ &= \phi(\mathfrak{a})\phi(x)\phi(y) = \phi(\mathfrak{a})\phi(xy) \\ &= \phi'(xy)\phi(\mathfrak{a}), \end{aligned}$$

so ϕ' is a ring homomorphism as required.

Let j be the smallest positive integer such that the coefficient at τ^j of $\phi(\mathfrak{a})$ is non-zero. Let λ be this coefficient. Comparing coefficients at τ^j in the equation $\phi'(x)\phi(\mathfrak{a}) = \phi(\mathfrak{a})\phi(x)$ gives us

$$D(\phi'(x))\lambda = D(\phi(x))^{q^j}\lambda.$$

From lemma 4.6 we see that $j = h \deg(\mathfrak{c}) \text{ord}_{\mathfrak{c}}(\mathfrak{a})$, where $\mathfrak{c} = \ker(\delta)$. If δ is injective, we conclude that $j = 0$, so $D\phi' = \delta$ as required. Otherwise, the image of δ is a field isomorphic to A/\mathfrak{c} . This is a finite field of order $q^{\deg(\mathfrak{c})}$. As $\deg(\mathfrak{c})$ is a divisor of j , q^j -th powering is the identity on this field, so $D(\phi'(x)) = \delta(x)^{q^j} = \delta(x)$ holds for all x in A , as required.

Note that $\phi(\mathfrak{a})$ is a morphism from ϕ to ϕ' by construction of ϕ' . It is non-zero, so it preserves height and rank. \square

Remark 5.2. If $\mathfrak{a} = (a)$ is a principal ideal then $\phi(\mathfrak{a}) = \mu^{-1}\phi(a)$, where μ is the leading coefficient of $\phi(a)$ and we have $(\mathfrak{a} * \phi)(x) = \mu^{-1}\phi(x)\mu$ for all x in A .

Lemma 5.3. *For any two non-zero ideals \mathfrak{a} and \mathfrak{b} of A we have $\phi(\mathfrak{a}\mathfrak{b}) = (\mathfrak{b} * \phi)(\mathfrak{a})\phi(\mathfrak{b})$ and $\mathfrak{a} * \mathfrak{b} * \phi = (\mathfrak{a}\mathfrak{b}) * \phi$.*

Proof. For all a in \mathfrak{a} , we observe $(\mathfrak{b} * \phi)(a)\phi(\mathfrak{b}) = \phi(\mathfrak{b})\phi(a)$. Note that for any b in \mathfrak{b} , there is an x such that $\phi(b) = x\phi(\mathfrak{b})$. We conclude that $\phi(b)\phi(a) = x(\mathfrak{b} * \phi)(a)\phi(\mathfrak{b})$ is an element of $(\mathfrak{b} * \phi)(\mathfrak{a})\phi(\mathfrak{b})$. Since the $\phi(a)\phi(\mathfrak{b})$ generate the left ideal of $\phi(\mathfrak{a}\mathfrak{b})$, we conclude that the left ideal generated by $\phi(\mathfrak{a}\mathfrak{b})$ is contained in that generated by $(\mathfrak{b} * \phi)(\mathfrak{a})\phi(\mathfrak{b})$.

Conversely, since the $\phi(b)$ with b in \mathfrak{b} generate the left ideal generated by $\phi(\mathfrak{b})$ we know that for every a in \mathfrak{a} , $\phi(\mathfrak{b})\phi(a)$ is in the left ideal generated by the $\phi(b)\phi(a)$ with b in \mathfrak{b} . Therefore we see that all elements of the form $(\mathfrak{b} * \phi)(a)\phi(\mathfrak{b})$ are in the left ideal of $\phi(\mathfrak{a}\mathfrak{b})$. This proves the other inclusion. Since $\phi(\mathfrak{a}\mathfrak{b})$ and $(\mathfrak{b} * \phi)(\mathfrak{a})\phi(\mathfrak{b})$ are now two monic generators of the same left ideal, they are equal.

The second equality is an immediate consequence of the first. \square

• The Leading Coefficient Map

Let ϕ be a Drinfeld module over ℓ of rank r . We consider in this section the map μ from $A - \{0\}$ to ℓ^\times that sends a non-zero a in A to the leading coefficient of $\phi(a)$.

This map satisfies the following multiplicative relation for all x and y in A non-zero:

$$\mu(xy) = \mu(x)\mu(y)^{q^{\deg(\phi(x))}} = \mu(x)\mu(y)^{q^{-r \deg(\infty)\text{ord}_\infty(x)}}.$$

Unfortunately, μ does not exhibit good behaviour with respect to the addition in A . Let x and y be two non-zero elements of A with $x + y$ non-zero. Then we have $\mu(x + y) = \mu(x) + \mu(y)$ in case the ord_∞ of x , y and $x + y$ are all the same. However, if $\text{ord}_\infty(x)$ is strictly larger than $\text{ord}_\infty(y)$, we have $\mu(x + y) = \mu(x)$. If x and y have the same ord_∞ , but $x + y$ has lower ord_∞ , all bets are off.

In this section we define an alternative leading coefficient map on a different ring, which allows us to retain the nice multiplicative behaviour of μ but also exhibits good additive behaviour. The approach is, regrettably, through formulas. However, the appealing nature of some of the results suggest there may be something more intrinsic behind it. Just what is not clear to me at present.

Recall that K is the completion of k at ∞ . Extend ord_∞ to this field. Write \mathcal{O} for the ring of integers of K , which consists of the elements with non-negative ord_∞ . This ring is local and its maximal ideal \mathfrak{m} consists of the elements with positive ord_∞ . Let \mathbf{F}_Q be \mathcal{O}/\mathfrak{m} . It is the residue field at ∞ , so $Q = q^{\deg(\infty)}$.

Fix an element π such that $\text{ord}_\infty(\pi)$ is 1, that is, a generator of the maximal ideal \mathfrak{m} . Then in fact K is isomorphic to the Laurent series ring $\mathbf{F}_Q((\pi))$, that is, expressions of the form $x = \sum_{i \geq \text{ord}_\infty(x)} \alpha_i \pi^i$ with all the α_i in \mathbf{F}_Q and $\alpha_{\text{ord}_\infty(x)}$ non-zero. The local ring \mathcal{O} is the power series ring $\mathbf{F}_Q[[\pi]]$.

Note that K comes with a natural filtration coming from ord_∞ . We put

$$\text{Fil}_i(K) = \{x \in K : \text{ord}_\infty(x) \geq -i\} = \mathfrak{m}^{-i}.$$

The graded ring of K is the ring

$$\text{Gr}(K) = \bigoplus_{i \in \mathbf{Z}} \text{Fil}_i(K) / \text{Fil}_{i-1}(K) = \bigoplus_{i \in \mathbf{Z}} \mathfrak{m}^{-i} / \mathfrak{m}^{-i+1}.$$

In addition, we write $R(K)$ for the subring of $\text{Gr}(K)$ consisting of the parts of non-negative degree,

$$R(K) = \bigoplus_{i \geq 0} \mathfrak{m}^{-i} / \mathfrak{m}^{-i+1}.$$

For convenience of notation we write

$$N(x) = q^{-r \deg(\infty) \text{ord}_\infty(x)}$$

whenever this makes sense.

There is a natural map ρ from A to $R(K)$ sending 0 to 0 and x in A non-zero to the class of x in the degree $-\text{ord}_\infty(x)$ part of $R(K)$. This map is well-behaved with respect to multiplication, but not with respect to addition, as one can readily see. In fact, it has the same problems as the original μ . Our aim is to define μ on $R(K)$ in such a way that we recover the original map by composing with ρ .

Lemma 5.4. *For every sufficiently large positive integer n , the image of ρ in the degree n part of $R(K)$ is all of $\mathfrak{m}^{-n} / \mathfrak{m}^{-n+1}$.*

Proof. Note that $\mathcal{O} \cap k$ is the valuation ring of k corresponding to ∞ . Therefore $\mathfrak{m}^{-n} \cap k$ is the stalk at ∞ of the line bundle $\mathcal{O}_X(n\infty)$ on X . Therefore $\mathfrak{m}^{-n} \cap A$ is the set of global sections of this line bundle.

Now there is an exact sequence of vector spaces over \mathbf{F}_q

$$0 \longrightarrow \mathfrak{m}^{-n+1} \cap A \longrightarrow \mathfrak{m}^{-n} \cap A \longrightarrow \mathfrak{m}^{-n} / \mathfrak{m}^{-n+1}.$$

From the Riemann Roch theorem, we know that the dimensions of the first two differ by $\deg(\infty)$ for n sufficiently large. This is precisely the dimension of $\mathfrak{m}^{-n} / \mathfrak{m}^{-n+1}$ as this is isomorphic to the residue field \mathbf{F}_Q of K . We conclude that for n sufficiently large all elements of $\mathfrak{m}^{-n} / \mathfrak{m}^{-n+1}$ are in the class of some x from $\mathfrak{m}^{-n} \cap A$. \square

Proposition 5.5. *There is a unique additive group homomorphism*

$$\mu : R(K) \longrightarrow \ell$$

satisfying

$$\mu(xy) = \mu(x)\mu(y)^{N(x)}$$

for all homogeneous x and y in $R(K)$.

Proof. Suppose x and y are different representatives in A of the same class in $\mathfrak{m}^{-n} / \mathfrak{m}^{-n+1}$. Then $\text{ord}_\infty(x) = \text{ord}_\infty(y)$ is equal to n and $\text{ord}_\infty(x - y)$ is less than n . Therefore the degree of $\phi(x - y)$ is less than that of $\phi(x)$, so $\mu(x)$ and $\mu(y)$ are the same. We conclude that we can extend μ to $\mathfrak{m}^{-n} / \mathfrak{m}^{-n+1}$ for n sufficiently large and that it still satisfies the relation above.

Now let x be in $\mathfrak{m}^k / \mathfrak{m}^{k+1}$ and let y be in $\mathfrak{m}^{-n} / \mathfrak{m}^{-n+1}$ for n sufficiently large. Then $\mu(xy)$ and $\mu(y)$ are both defined. Now we define $\mu(x)$ by requiring

$$\mu(xy) = \mu(x)\mu(y)^{N(x)}.$$

We must check that this is independent of y . Let z be an element of $\mathfrak{m}^{-n'}/\mathfrak{m}^{-n'+1}$ with n' sufficiently large. Then we know that

$$\mu(y)\mu(z)^{N(y)} = \mu(yz) = \mu(z)\mu(y)^{N(z)}$$

and

$$\mu(xz)\mu(y)^{N(xz)} = \mu(xyz) = \mu(xy)\mu(z)^{N(xy)}.$$

Combining these we see that

$$\left(\frac{\mu(y)}{\mu(z)}\right)^{N(x)} = \left(\frac{\mu(y)^{N(z)}}{\mu(z)^{N(y)}}\right)^{N(x)} = \frac{\mu(y)^{N(xz)}}{\mu(z)^{N(xy)}} = \frac{\mu(xy)}{\mu(xz)}$$

and so

$$\mu(xy)\mu(y)^{-N(x)} = \mu(xz)\mu(z)^{-N(x)}.$$

In other words, the two possible definitions of $\mu(x)$ coincide.

We have now defined μ on $\mathfrak{m}^{-n}/\mathfrak{m}^{-n+1}$ for all non-negative n . For n sufficiently large, it is clear that $\mu(x+y) = \mu(x) + \mu(y)$ holds, as x and y are both represented by elements of A . From the way we extended μ to all $\mathfrak{m}^{-n}/\mathfrak{m}^{-n+1}$ it is clear that $\mu(x+y) = \mu(x) + \mu(y)$ still holds for all x and y in these sets. We can now extend μ additively to the entire direct sum. \square

Lemma 5.6. *The restriction ι_ϕ of μ_ϕ to $\mathbf{F}_Q = \mathfrak{m}^0/\mathfrak{m}^1$, the residue field at ∞ , depends only on the isomorphism class of ϕ .*

Proof. Recall that all isomorphisms of Drinfeld modules are elements of ℓ^\times . Let λ be an element of ℓ^\times and put $\psi = \lambda\phi\lambda^{-1}$. For every x in A , we have

$$\mu_\psi(x) = \lambda^{1-N(x)}\mu_\phi(x).$$

From the construction of the leading coefficient map it follows that this relation holds for all x . In particular, if x is in \mathfrak{m}^0 , then $N(x) = 1$, so $\mu_\psi(x) = \mu_\phi(x)$. \square

Remark 5.7. If ℓ is perfect, then we can also take q -th roots in ℓ , in a unique way. In this case, we can extend μ uniquely to all of $\text{Gr}(K)$. Moreover, we then get a map $\mu : K^\times \rightarrow \ell^\times$ satisfying $\mu(xy) = \mu(x)\mu(y)^{N(x)}$ for all x and y .

• Drinfeld Modules over C

We begin this section with some results concerning analysis in C . Recall that C is an algebraically closed, complete field. Let v be a non-trivial valuation on C , \mathcal{O} for the valuation ring, \mathfrak{p} for its maximal ideal and κ the residue class field.

Definition 5.8. An entire function on C is a formal power series $e = \sum_{i \geq 0} e_i z^i$ in $C[[z]]$ which is everywhere convergent, i.e. for all λ in C , the sum $\sum_{i \geq 0} e_i \lambda^i$ converges.

Lemma 5.9. *Let $e = \sum_{i \geq 0} e_i z^i$ be a formal power series in $C[[z]]$. For any π in \mathfrak{p} put*

$$e_\pi = \sum_{i \geq 0, v(e_i) < v(\pi)} e_i z^i$$

Then the following are equivalent:

1. e is entire;
2. $\lim_{n \rightarrow \infty} v(e_n)/n = \infty$;
3. For any π in \mathfrak{p} , e_π is a polynomial and moreover $\lim_{t \rightarrow \infty} \deg(e_{\pi^t})/t = 0$.

Proof. We begin by showing 1 and 2 are equivalent. Recall that v is a non-archimedean valuation. Therefore an infinite sum converges if and only if the terms go to zero. Let λ be in C . Then the n -th term of $e(\lambda)$ is $e_n \lambda^n$, which has valuation $v(e_n) + v(\lambda)n$. We see that if $v(e_n)/n$ does not eventually become arbitrarily large, we can pick a λ such that the terms eventually get negative valuation, so they become big, rather than going to 0.

Next, we show that 2 implies 3. Clearly, if $v(e_n)$ tends to infinity, only a finite number of terms have valuation less than a given bound. So all the e_π are polynomials. Let π in C have positive valuation and let $\lambda > 0$. We must show that for all t sufficiently large, $\deg(e_{\pi^t}) < \lambda t$. Note that $\deg(e_{\pi^t}) = n$ if and only if $v(e_n) < v(\pi)t$ and $v(e_m) \geq v(\pi)t$ for all m larger than n .

Suppose now that $n = \deg(e_{\pi^t})$ is larger than λt . Then we have

$$v(e_n) < v(\pi)t = \frac{v(\pi)}{\lambda} \lambda t < \frac{v(\pi)}{\lambda} n,$$

which cannot happen for infinitely many n , as we know that $v(e_n)$ tends to infinity faster than any multiple of n . Hence 2 implies 3.

To show that 3 implies 2 we have to be a little more careful. Again we fix a π in C with positive valuation. Let λ be positive and let ϵ be strictly between 0 and λ . Suppose that $v(e_n)$ is less than $(\lambda - \epsilon)n$ for infinitely many n . We rewrite this as

$$v(e_n) < v(\pi) \left(\frac{\lambda n}{v(\pi)} - \frac{\epsilon n}{v(\pi)} \right).$$

Note that for all n sufficiently large, $\frac{\epsilon n}{v(\pi)}$ is at least 1, so if we put t the largest integer below $\frac{\lambda n}{v(\pi)}$, we have $v(e_n) < v(\pi)t$ and $t < \frac{\lambda n}{v(\pi)}$. From this we conclude that we have $\deg(e_{\pi^t}) > n > \frac{v(\pi)}{\lambda} t$ infinitely often. Assuming 3 this cannot happen, so we must have for every ϵ and every λ that $v(e_n) > (\lambda - \epsilon)n$ for all n sufficiently large, which implies 2. \square

Next we prove a version of Hensel's lemma for entire functions.

Lemma 5.10. *Let f be an entire function that lies in $\mathcal{O}[[z]]$ such that the constant term of f is in \mathcal{O}^\times . Suppose that \bar{g} and \bar{h} are coprime polynomials in $\kappa[z]$ such that \bar{f} , the reduction of f to $\kappa[z]$ factors as $\bar{g}\bar{h}$. Then there is a polynomial g in $\mathcal{O}[z]$ of degree $\deg(\bar{g})$ and an entire function h in $\mathcal{O}[[z]]$ such that g reduces to \bar{g} , h reduces to \bar{h} and f factors as gh .*

Proof. Note that the only non-zero coefficients of \bar{f} are at those places where f has a coefficient of valuation 0. As the remaining coefficients have valuations that eventually become large, there is a minimal positive value for a valuation of a coefficient of f . Let π in C be an element with valuation positive valuation less than this minimum.

For all $n \geq 0$ we let f_n be the polynomial in $\mathcal{O}[z]$ of the smallest degree, such that $f - f_n$ is in $\pi^{2^n} \mathcal{O}[[z]]$.

We will prove by induction that there are polynomials a_n, b_n, g_n and h_n in $\mathcal{O}[z]$ such that the relations

$$\begin{aligned} f_n &\equiv g_n h_n \pmod{\pi^{2^n} \mathcal{O}[z]} \\ a_n g_n + b_n h_n &\equiv 1 \pmod{\pi^{2^n} \mathcal{O}[z]} \end{aligned}$$

hold and we have $\deg(g_n) = \deg(g_0)$ and $\deg(h_n) = \deg(f_n) - \deg(g_n)$.

For $n = 0$, note that there are polynomials a and b in $\kappa[z]$ such that $a\bar{g} + b\bar{h} = 1$ holds in $\kappa[z]$. Now we just pick any lifts of the appropriate degree of \bar{g}, \bar{h}, a and b and use them for g_0, h_0, a_0 and b_0 .

Suppose now that we have g_n, h_n, a_n and b_n satisfying the required conditions. Write f_{n+1} as $g_n h_n + \pi^{2^n} r_n$, where r_n is a polynomial in $\mathcal{O}[z]$ of degree at most $\deg(f_{n+1})$. Let u_n be the unique polynomial in $\mathcal{O}[z]$ of degree at most $\deg(g_n)$ that satisfies

$$u_n \equiv b_n r_n \pmod{g_n \mathcal{O}[z]}.$$

Note that we now have

$$u_n h_n \equiv r_n \pmod{g_n \mathcal{O}[z] + \pi^{2^n} \mathcal{O}[z]},$$

so there is a polynomial v_n of degree at most $\deg(r_n) - \deg(g_n)$ such that we have

$$u_n h_n \equiv r_n - v_n g_n \pmod{\pi^{2^n} \mathcal{O}[z]}.$$

Now we put g_{n+1} and h_{n+1} equal to $g_n + u_n \pi^{2^n}$ and $h_n + v_n \pi^{2^n}$ respectively. Note that we have

$$\begin{aligned} g_{n+1} h_{n+1} &= g_n h_n + (u_n h_n + v_n g_n) \pi^{2^n} + \pi^{2^{n+1}} u_n v_n \\ &\equiv g_n h_n + r_n \pi^{2^n} \pmod{\pi^{2^{n+1}} \mathcal{O}[z]} \\ &\equiv f_{n+1} \pmod{\pi^{2^{n+1}} \mathcal{O}[z]}. \end{aligned}$$

Moreover, by construction $\deg(g_{n+1})$ is the same as $\deg(g_n)$ and $\deg(h_{n+1})$ is bounded above by $\deg(r_n) - \deg(g_n) \leq \deg(f_{n+1}) - \deg(g_{n+1})$. As we have

$$g_{n+1} h_{n+1} \equiv f \pmod{\pi^{2^{n+1}} \mathcal{O}[[z]]},$$

we see that the degree of h_{n+1} is also at least $\deg(f_{n+1}) - \deg(g_{n+1})$. So the required degree relations hold.

We also need to produce the new auxillary polynomials a_{n+1} and b_{n+1} . Let t_n in $\mathcal{O}[z]$ satisfy

$$a_n g_{n+1} + b_n h_{n+1} = 1 + t_n \pi^{2^n} \pmod{\pi^{2^{n+1}}}$$

and put $a_{n+1} = (1 - t_n \pi^{2^n}) a_n$ and $b_{n+1} = (1 - t_n \pi^{2^n}) b_n$. Now we compute that

$$\begin{aligned} a_{n+1} g_{n+1} + b_{n+1} h_{n+1} &= (1 - t_n \pi^{2^n}) (a_n g_{n+1} + b_n h_{n+1}) \\ &\equiv (1 - t_n \pi^{2^n}) (1 + t_n \pi^{2^n}) \pmod{\pi^{2^{n+1}} \mathcal{O}[z]} \\ &\equiv 1 \pmod{\pi^{2^{n+1}} \mathcal{O}[z]} \end{aligned}$$

holds, as required. This concludes the induction step.

Note that for all $n \geq 0$, the differences $g_n - g_{n+1}$ and $h_n - h_{n+1}$ are both in $\pi^{2^n} \mathcal{O}[z]$. It follows that both the g_n and the h_n converge coefficient-wise. In the case of the g_n , the degree is bounded, so the limit g is again a polynomial of degree $\deg(\bar{g})$. The h_n converge to a power series h in $C[[z]]$. By construction, it is clear that we will have $f = gh$, so we are done if we can show that h is entire.

To see this, we note that, using the notation from the previous lemma, we have $\deg(f_n)$ is $\deg(f_{\pi^{2^t}})$ and $\deg(h_n)$ is $\deg(h_{\pi^{2^n}})$. Since f is entire, it therefore satisfies $\lim_{n \rightarrow \infty} \frac{\deg(f_n)}{2^n} = 0$ and as $\deg(f_n)$ and $\deg(h_n)$ differ only by a constant, the same holds for the h_n , so h is also entire. \square

Corollary 5.11. *Any non-constant entire function on C is onto.*

Proof. Let e be a non-constant entire function. It suffices to show that we can find a zero of e . If the constant term of e is 0, we are done, as then 0 is a zero of e . Multiplying with the inverse of the constant term, we may now assume that e has constant term 1. By applying a substitution $z \mapsto \lambda z$, we can make sure that all the coefficients of e are in \mathcal{O} and that all of them except two (the constant term and one other) are in \mathfrak{p} . The resulting power series in $\mathcal{O}[[z]]$ reduces to a polynomial \bar{e} of positive degree in $\kappa[z]$. The lemma then show that we can lift this to a polynomial factor of the entire function. As C is algebraically closed, this polynomial factor has roots. \square

Next we give a way of constructing entire functions using their roots.

Lemma 5.12. *Let $\{\lambda_i\}_{i \geq 1}$ be a sequence in C^\times such that $\{v(\lambda_i)\}_{i \geq 1}$ is a non-decreasing sequence that tends to ∞ . Then the infinite product*

$$e(z) = \prod_{i \geq 1} (1 - z\lambda_i)$$

converges to an entire function. The zeroes of e , with counted multiplicities are given by the sequence $\{\lambda_i^{-1}\}_{i \geq 0}$.

Proof. Any finite number of factors will not affect the convergence of the product. As the $v(\lambda_i)$ are positive for all i except possibly a finite number, we may just throw out this finite number of λ 's and assume that all the $v(\lambda_i)$ are positive.

If we expand the infinite product formally, we obtain the following infinite sum for the coefficient e_n of z^n in this expansion:

$$e_n = (-1)^n \sum_{1 \leq i_1 < \dots < i_n} \lambda_{i_1} \cdots \lambda_{i_n}$$

In order to show that the infinite product converges to an entire function we must show that all these e_n exist and that $v(e_n)$ goes to ∞ faster than any constant multiple of n as n gets large.

Note that $v(\lambda_{i_1} \cdots \lambda_{i_n})$ is equal to $v(\lambda_{i_1}) + \cdots + v(\lambda_{i_n})$, which will tend to infinity as any one of the i 's gets large. Therefore, the sum defining e_n converges. Moreover, this valuation is non-decreasing in each of the i 's, so we can estimate

$$v(e_n) \geq \inf_{1 \leq i_1 < \cdots < i_n} v(\lambda_{i_1} \cdots \lambda_{i_n}) \geq v(\lambda_1 \cdots \lambda_n).$$

Let $r > 0$ and let N be such that for all $i \geq N$ we have $v(\lambda_i) \geq r + 1$. Then for all $n > (r + 1)N$ we have

$$\begin{aligned} v(e_n) &\geq (v(\lambda_1) + \cdots + v(\lambda_N)) + (v(\lambda_{N+1}) + \cdots + v(\lambda_n)) \\ &\geq (n - N)v(\lambda_{N+1}) \\ &\geq (n - N)(r + 1) = nr + n - N(r + 1) > nr \end{aligned}$$

and therefore, e is an entire function. \square

Remark 5.13. In fact, an entire function is determined up to scalar multiplication by its set of roots with counted multiplicities. We do not prove this result here, but it follows with a bit more work from the things we have shown.

Lemma 5.14. *Let ϕ be a Drinfeld module over C . Let e_ϕ be the formal power series associated to ϕ as per lemma 4.15. Then e_ϕ is entire.*

Proof. Write e_ϕ as $\sum_{i \geq 0} e_i z^{q^i}$. We must show that $v(e_i)$ goes to ∞ faster than any fixed multiple of q^i . It therefore suffices to show that there is a $\delta > 0$ and a real number c such that $v(e_n) > \delta(n - c)q^n$ holds for all n . Note that given δ we can always choose c such that this relation holds for all n below any given n_0 .

Let y be in A be non-constant and write $\phi(y)$ as $\sum_{i=0}^d s_i \tau^i$. Put e_{-1} up to e_{-d} equal to 0. Then we know from the construction of theorem 3.13 that we have the relation

$$(s_0^{q^n} - s_0)e_n = s_1 e_{n-1}^q + \cdots + s_d e_{n-d}^{q^d}.$$

Note the $s_0 = y$ has a negative valuation, so $v(s_0^{q^n})$ is smaller than $v(s_0)$ and therefore $v(s_0^{q^n} - s_0) = v(s_0^{q^n}) = q^n v(s_0)$. Considering the valuation on either side of the equation defining e_n we get

$$\begin{aligned} q^n v(s_0) + v(e_n) &\geq \min(v(s_1) + qv(e_{n-1}), \dots, v(s_d) + q^d v(e_{n-d})) \\ &\geq \min(v_1, \dots, v_d) + \min(qv(e_{n-1}), \dots, q^d v(e_{n-d})). \end{aligned}$$

Suppose that for all $m < n$ we have $v(e_m) > \delta(m - c)q^m$. Then we derive from the previous estimate that

$$\begin{aligned} v(e_n) &\geq -q^n v(s_0) + \min(v_1, \dots, v_d) + \min(\delta(n - (c + 1))q^n, \dots, \delta(n - (c + d))q^n) \\ &\geq -q^n v(s_0) + \min(v_1, \dots, v_d) + \delta(n - (c + d))q^n \\ &= \delta(n - c)q^n + (-v(s_0) - \delta d)q^n + \min(v_1, \dots, v_d). \end{aligned}$$

Recall that $v(s_0)$ is negative, so if we take δ a sufficiently small positive number, $-v(s_0) - \delta d$ will be positive. We then pick n_0 be such that $(-v(s_0) - \delta d)q^{n_0} > -\min(v_1, \dots, v_d)$. Lastly, we fix c such that $v(e_m) > \delta(m - c)q^m$ holds for all $m \leq n_0$. Then the inequality we have just derived shows that we will have $v(e_n) \geq \delta(n - c)q^n$ for all n , as required. \square

Definition 5.15. An lattice Λ in C is a discrete sub- A -module of C such that $K\Lambda$ is a finite dimensional K -vector space.

Note that a lattice Λ is a discrete A -module inside a finite dimensional K -vector space and is therefore finitely generated and torsion free, so it is projective.

Lemma 5.16. Let M be a finitely generated A -submodule of K . Then M is discrete if and only if M is of rank 1.

Proof. Clearly, if M has rank 1, then M is discrete. To prove the other implication, suppose that a and b are two A -linearly independent elements of M .

Let π be a generator of the maximal ideal of K . Recall that K is isomorphic to the Laurent series ring $\mathbb{F}_Q((\pi))$. So we can write a as $\sum_{i \geq -n_a} \alpha_i \pi^i$ and b as $\sum_{i \geq -n_b} \beta_i \pi^i$. Moreover, by lemma 5.4 we know that there is a positive integer N such that for all n greater than N , we have elements of A whose ord_∞ is $-n$ with any leading coefficient. So let x be any non-zero of A , then, through what is essentially long division, we can find a y in A such that $xa - yb$ has ord_∞ at least $-N$. But this gives us infinitely many points in M inside a finite ball. So M is not discrete. \square

Corollary 5.17. A finitely generated projective A -submodule Λ of C is a lattice if and only if the rank of Λ is the same as the K -dimension of $K\Lambda$.

Definition 5.18. Let Λ and Λ' be two lattices. If they do not have the same rank, there is only one morphism between them, the zero morphism. If the rank of Λ is the same as that of Λ' , then a morphism from Λ to Λ' is an element α of C such that $\alpha\Lambda$ is contained in Λ' .

Lemma 5.19. Let Λ be a lattice. Then the exponential function e_Λ given by

$$e_\Lambda(z) = z \prod_{0 \neq \lambda \in \Lambda} (1 - z\lambda^{-1}).$$

is an \mathbb{F}_q -linear entire function.

Proof. Let V be the K -span of Λ . It is a finite-dimensional K -vector space. Note that Λ is a discrete subset of V . It follows that for every real r the ball $\{x \in C | v(x) \leq r\}$ contains at most a finite number of points of Λ . This implies that the λ^{-1} form a sequence that satisfies the conditions of lemma 5.12. We see that e_Λ is entire.

For positive integer n the set

$$\Lambda_n = \{\lambda \in \Lambda : v(\lambda) > -n\}$$

is finite. By the non-archimedean triangle equality, it is an additive subgroup of Λ and as all elements of \mathbb{F}_q^\times have valuation 0, it is even an \mathbb{F}_q -vector space.

If we put

$$e_n(z) = z \prod_{0 \neq \lambda \in \Lambda_n} (1 - z\lambda)$$

then clearly e_Λ is the limit of the e_n . Each of the e_n is additive by lemma 3.8. So e_Λ is an additive power series. \square

Lemma 5.20. *Let Λ be a projective A -module of rank r . Then $\Lambda/a\Lambda$ is isomorphic to $(A/aA)^r$ for all a in A .*

Proof. If Λ is a free A -module of rank r , this is obvious. For projective Λ we know by the structure theory for modules over Dedekind domains that Λ is the direct sum of a free module of rank $r - 1$ and an ideal I of A .

We are therefore reduced to showing that I/aI is isomorphic to A/aA . We do this by localisation. Write $a = \prod_i \mathfrak{p}_i^{e_i}$ and consider the localisation of A/aA at \mathfrak{p}_i . We see that the factors corresponding to \mathfrak{p}_j 's with $j \neq i$ all vanish and what we are left with is $(A/aA)_{\mathfrak{p}_i} = A/\mathfrak{p}_i^{e_i}A$. Similarly, $(I/aI)_{\mathfrak{p}_i}$ is equal to $I/\mathfrak{p}_i^{e_i}I$, which is the same as $I_{\mathfrak{p}_i}/\mathfrak{p}_i^{e_i}I_{\mathfrak{p}_i}$. But now, $I_{\mathfrak{p}_i}$ is a free $A_{\mathfrak{p}_i}$ module of rank 1, so equality clearly holds. \square

Theorem 5.21. *Let Λ be a lattice. Then there is a unique Drinfeld module ϕ_Λ over C such that we have a short exact sequence*

$$0 \longrightarrow \Lambda \longrightarrow C \xrightarrow{e_\Lambda} \Phi \longrightarrow 0$$

of A -modules, where Φ is the A -module on the additive group of C given by ϕ_Λ . The rank of ϕ_Λ is the same as the rank of Λ .

Proof. First of all, it is clear that e_Λ has kernel Λ and is onto (as it is a non-constant entire function). The tricky bit is showing that the A -action induced on C via e_Λ comes from a Drinfeld module.

Let a be a non-zero element of A . Note that the lattice Λ is of finite index in the lattice $a^{-1}\Lambda$. Moreover, we have that an x in C is in $a^{-1}\Lambda$ if and only if ax is in Λ , so $e_{a^{-1}\Lambda}(z)$ and $e_\Lambda(az)$ are separable entire functions with the same set of roots and therefore differ by a scalar multiplication. Comparing the derivatives we see that we have

$$ae_{a^{-1}\Lambda}(z) = e_\Lambda(az)$$

in $C[[z]]$.

The set $W_a = e_\Lambda(a^{-1}\Lambda)$ is a finite \mathbf{F}_q -vector space isomorphic to $(a^{-1}\Lambda)/\Lambda$. Let $\phi(a)$ be the unique separable \mathbf{F}_q -linear polynomial whose set of roots is W_a and whose derivative is a . Then $[\phi(a)](e_\Lambda(z))$ is a separable entire function whose set of roots is $a^{-1}\Lambda$. Looking at the derivatives we conclude that

$$[\phi(a)](e_\Lambda(z)) = ae_{a^{-1}\Lambda}(z) = e_\Lambda(az)$$

holds. Transporting this equality to $C\{\{\tau\}\}$ we conclude that we have

$$\phi(a)e_\Lambda = e_\Lambda a$$

for all a in A nonzero. Note that if a is in \mathbf{F}_q , then $a^{-1}\Lambda$ is the same as Λ so that we find $\phi(a) = a$ in $C\{\tau\}$.

From corollary 3.14 we conclude that the relation we have just derived defines the $\phi(a)$ uniquely for all non-constant a in A . From this we conclude that the map ϕ is a ring

homomorphism from A to $C\{\{\tau\}\}$. By construction we have $D(\phi(a)) = a$ for all a in A . Note that $a^{-1}\Lambda/\Lambda$ is isomorphic to $\Lambda/a\Lambda$, so by the previous lemma, the zero set of $\phi(a)$ is equal to $(A/aA)^r$, as required. \square

Theorem 5.22. *There is an equivalence of categories between the lattices in C and the Drinfeld modules over C . The rank r lattices correspond to the rank r Drinfeld modules.*

Proof. Let ϕ be a Drinfeld module over C of rank r . Write e for the exponential function of the Drinfeld module. We write Λ for its set of roots. We want to show that Λ is a lattice. Note that Λ is discrete in C as it is the set of roots of an entire function.

As e is an additive function, Λ is an additive subgroup of C . To show that Λ is an A -module, note that for any a in A we have

$$e(az) = [\phi(a)](e(z))$$

for all z in Z , so that if z is in Λ , az is also in Λ .

Lastly we must show that $K\Lambda$ has finite dimension. Suppose that $\lambda_1, \dots, \lambda_s$ in Λ are K -linearly independent. Write E for the K -span of these vectors. Then $\Lambda' = \Lambda \cap E$ is a discrete A -submodule of a finite-dimensional K -vector space, whose K -span is the entire space. We conclude that Λ' is finitely generated projective of rank s .

From lemma 5.20 we conclude that $\Lambda'/a\Lambda'$ is isomorphic to $(A/a)^s$ for all a in A . Note that $\Lambda' \cap a\Lambda$ is just $a\Lambda'$, so that $\Lambda'/a\Lambda'$ is a submodule of $\Lambda/a\Lambda$. We know that $\Lambda/a\Lambda$ is isomorphic to the a -torsion submodule of ϕ and therefore to $(A/a)^r$, where r is the rank of ϕ . We conclude that $s < r$, which shows that $K\Lambda$ is finite dimensional.

From the constructions it is clear that the maps from Drinfeld modules to lattices and the other way are eachothers inverse. So we are done if we can show that morphisms on the one side correspond to morphisms on the other side.

Suppose that $\alpha\Lambda$ is of finite index in Λ' . Note that the elements of $\alpha\Lambda$ are precisely the roots of $e_{\Lambda}\alpha^{-1}$ so that the set $[e_{\Lambda}\alpha^{-1}](\Lambda')$ is a finite \mathbf{F}_q -vector space. Let f be the separable \mathbf{F}_q -linear polynomial with precisely this set of roots whose derivative is α . Then we see that $e_{\Lambda'}$ and $f e_{\Lambda}\alpha^{-1}$ are separable entire functions with the same roots and the same derivative, so they are equal. Note that for every x in A we have

$$\phi_{\Lambda'}(x)f = e_{\Lambda'}x e_{\Lambda'}^{-1}f = f e_{\Lambda}\alpha^{-1}x \alpha e_{\Lambda}^{-1}f^{-1}f = f e_{\Lambda}x e_{\Lambda}^{-1} = f\phi_{\Lambda}(x),$$

so f is a morphism of Drinfeld modules from ϕ_{Λ} to $\phi_{\Lambda'}$.

Conversely, suppose that f is a morphism of Drinfeld modules from ϕ to ϕ' . Let α be the derivative of f . Then the computation above shows that we have $e'_{\phi} = f e_{\phi}\alpha^{-1}$, from which we conclude that $\alpha\Lambda_{\phi}$ is a sublattice of finite index in $\Lambda_{\phi'}$. \square

Corollary 5.23. *The set of isomorphism classes of rank 1 lattices in C is isomorphic as a set with $\text{Cl}(A)$ -action to the set of rank 1 Drinfeld modules over C . The former is clearly a torsor for $\text{Cl}(A)$, so the latter is as well.*

Proof. A non-zero ideal \mathfrak{a} of A acts on the lattices in C by sending a lattice Λ to $\mathfrak{a}^{-1}\Lambda$. Note that Λ is a sublattice of finite index in $\mathfrak{a}^{-1}\Lambda$. As in the proof of the previous theorem, it follows that there is a monic, separable \mathbb{F}_q -linear polynomial f such that $e_{\mathfrak{a}^{-1}\Lambda} = fe_{\Lambda}$. One readily computes that f is equal to $\phi_{\Lambda}(\mathfrak{a})$ and therefore that $\phi_{\mathfrak{a}^{-1}\Lambda} = \mathfrak{a} * \phi_{\Lambda}$. This shows that the actions of the ideals on either side are compatible. The result now follows. \square

• Fields of Definition

Let ϕ be a Drinfeld module over C . We say that ϕ is defined over a subfield L of C containing k , if $\phi(x)$ has coefficients in L for all x in A .

We say that ϕ can be defined over a subfield L of C containing k , if ϕ is isomorphic over C to some ψ which is defined over L . In this case we call L a field of definition for ϕ .

Lemma 5.24. *Let ϕ be a Drinfeld module over C of rank 1. Then ϕ can be defined over K , the completion of k at ∞ .*

Proof. By the uniformiation theorem, ϕ is isomorphic to $\phi_{\mathfrak{a}}$ for some ideal \mathfrak{a} of A viewed as a one-dimensional lattice inside C . Note that the construction of $\phi_{\mathfrak{a}}$ can then be carried out inside K , as all the elements of the lattice are in K and the coefficients that come up in the construction are infinite sums of these lattice elements. \square

Lemma 5.25. *Let ϕ be a Drinfeld module over C . Let x be a non-constant element of A . Then the subfield L of C generated by k and the coefficients of $\phi(x)$ is independent of the chosen x . It is the smallest field over which ϕ is defined.*

Proof. Let e be the exponential function of the Drinfeld module ϕ , following lemma 4.15. From the proof of theorem 3.13 we note that the coefficients of e are in L . \square

Lemma 5.26. *Let ϕ be a Drinfeld module over C and \mathfrak{a} an ideal of A . If ϕ is defined over L , then $\mathfrak{a} * \phi$ is also defined over L .*

Proof. Note that $\phi(\mathfrak{a})$ has coefficients in L . It follows that $\mathfrak{a} * \phi$ has coefficients in L . \square

Theorem 5.27. *Let ϕ be a Drinfeld module over C . Then there is a minimal field of definition L_{ϕ} which is contained in every field of definition of ϕ .*

Proof. Let a be a non-constant element of A . By lemma 5.25 a Drinfeld module ϕ' over C is defined over a field L if and only if the coefficients of $\phi'(a)$ are in L .

Let ζ be any element of C^{\times} and let $\phi' = \zeta\phi\zeta^{-1}$. Write $\phi(a)$ as $\sum_{i=1}^s \lambda_i \tau^{n_i}$ with all the λ_i non-zero. Then $\phi'(a)$ is equal to $\sum_{i=1}^s \lambda'_i \tau^{n_i}$, with $\lambda'_i = \zeta^{1-q^{n_i}} \lambda_i$ for all i .

Let g be the greatest common divisor of the $q^{n_i} - 1$ and pick integers e_i such that g can be written as $\sum_{i=1}^s e_i (q^{n_i} - 1)$. Now let ζ and ζ' be solutions in C of $\zeta^g = \prod_{i=1}^s \lambda_i^{e_i}$ and $(\zeta')^g = \prod_{i=1}^s \lambda_i^{e_i}$ respectively. Note that we have

$$(\zeta')^g = \prod_{i=1}^s (\lambda_i')^{e_i} = \zeta^{\sum_{i=1}^s e_i (1 - q^{n_i})} \prod_{i=1}^s \lambda_i^{e_i} = (\zeta^{-1} \zeta)^g.$$

As all the $1 - q^{n_i}$ are multiples of g , we have $(\zeta')^{1 - q^{n_i}} = (\zeta^{-1} \zeta)^{1 - q^{n_i}}$ for all i . This implies that the elements $\mu_i = \zeta^{1 - q^{n_i}} \lambda_i$ and $\mu'_i = (\zeta')^{1 - q^{n_i}} \lambda_i'$ are equal and therefore depend only on the isomorphism class of ϕ .

Thus the subfield L_ϕ of C generated by k and the elements μ_1 up to μ_s is contained in every field of definition for ϕ . Moreover, the Drinfeld module $\phi_{\min} = \zeta \phi \zeta^{-1}$ has $\phi_{\min}(a)$ equal to $\sum_{i=1}^s \mu_i \tau^{n_i}$ and therefore is defined over L_ϕ . \square

• Sign functions and Normalisation

Recall that K is the completion of k at ∞ .

Definition 5.28. A sign function ϵ on K is a group homomorphism

$$\epsilon : K^\times \longrightarrow \mathbf{F}_Q^\times$$

which is the identity on $\mathbf{F}_Q^\times \subset K$.

Lemma 5.29. Let \mathcal{O} be the ring of integers of the local field K and \mathfrak{p} its maximal ideal. Let ϵ be a sign function. Then $\epsilon(x) = 1$ for all x in \mathcal{O} congruent to 1 modulo \mathfrak{p} .

Proof. Suppose that α in \mathcal{O} is congruent to 1 mod \mathfrak{p} . Then the equation $x^{Q-1} - \alpha$ in $\mathcal{O}[x]$ reduces to $x^{Q-1} - 1$ in $\mathbf{F}_Q[x]$, the residue field. This is a separable polynomial and 1 is one of its roots, so by Hensel's lemma there is a β in \mathcal{O} with $\beta \equiv 1 \pmod{\mathfrak{p}}$ that is a root of $x^{Q-1} - \alpha$. This implies that

$$\epsilon(\alpha) = \epsilon(\beta^{Q-1}) = \epsilon(\beta)^{Q-1} = 1$$

as \mathbf{F}_Q^\times is annihilated by $Q - 1$. \square

Lemma 5.30. The set of all sign functions is a torsor for \mathbf{F}_Q^\times .

Proof. Let ϵ be a sign function and ζ in \mathbf{F}_Q^\times . Then the map

$$(\zeta\epsilon) : K^\times \longrightarrow \mathbf{F}_Q^\times \\ x \longmapsto \epsilon(x) \zeta^{\text{ord}_\infty(x)}$$

is again a sign function, as it is obviously a group homomorphism and $\text{ord}_\infty(x) = 0$ for all x in \mathbf{F}_Q^\times . It is clear that this gives an action of \mathbf{F}_Q^\times on the set of sign functions. If ϵ and ϵ' are two sign functions, then the function $x \mapsto \epsilon(x)/\epsilon'(x)$ is a group homomorphism from K^\times to \mathbf{F}_Q^\times which is 1 on all of \mathcal{O}^\times . Therefore, it factors via $\text{ord}_\infty : K^\times \longrightarrow \mathbf{Z}$. So there is a unique ζ in \mathbf{F}_Q^\times such that $\epsilon(x)/\epsilon'(x) = \zeta^{\text{ord}_\infty(x)}$ for all x in K^\times . \square

Let ϕ be a Drinfeld module over C of rank 1. As C is perfect, we can extend the leading coefficient map μ of ϕ to a map $\mu : K^\times \rightarrow C^\times$. Suppose that the image of this map lands inside \mathbf{F}_Q^\times . Then it is a group homomorphism, as we then have

$$\mu(xy) = \mu(x)\mu(y)^{q^{-\deg(\infty)\text{ord}_\infty(x)}} = \mu(x)\mu(y)^{Q^{-\text{ord}_\infty(x)}} = \mu(x)\mu(y).$$

As before, write ι for the restriction of μ to \mathbf{F}_Q^\times . Note that the map $\iota^{-1} \circ \mu$ is a sign function. It is clearly a group homomorphism and composition with ι^{-1} makes it the identity on \mathbf{F}_Q^\times . If $\iota^{-1} \circ \mu$ is equal to ϵ , then we say that ϕ is ϵ -normalised.

Lemma 5.31. *Let ϕ be a Drinfeld module over C of rank 1 and ϵ a sign function. Then the set of ϵ normalised Drinfeld modules isomorphic to ϕ is a torsor for $\mathbf{F}_Q^\times/\mathbf{F}_q^\times$.*

Proof. We begin by showing that the set is non-empty. Let π be a generator of \mathfrak{m} , the maximal ideal of \mathcal{O} , such that $\epsilon(\pi) = 1$. Fix a $(Q-1)$ -th root ζ of $1/\mu_\phi(\pi^{-1})$ in C . Let ϕ' be the Drinfeld module $\zeta\phi\zeta^{-1}$. Then $\mu_{\phi'}(\pi^{-1})$ is 1. Let x in K be non-zero. Then we can write x as $\pi^{-n}x'(1-\pi u)$ for some non-negative integer n , some x' in \mathbf{F}_Q^\times and some u in \mathcal{O} . We see that $\epsilon(x) = \epsilon(x') = x'$ as the other factors have sign 1. Also $\mu_{\phi'}(x) = \mu'_\phi(\pi^{-n}x')$ as the other term has a smaller ord_∞ . It follows that

$$\mu_{\phi'}(x) = \mu_{\phi'}(\pi^{-n}x') = \mu_{\phi'}(x') = \iota(x') = \iota(\epsilon(x))$$

holds for all x in K , so ϕ' is ϵ -normalised.

Suppose that ψ is a ϵ -normalised Drinfeld module. Suppose that ζ is an element of C^\times such that $\psi' = \zeta\psi\zeta^{-1}$ is also ϵ -normalised. Then we have

$$1 = \mu_{\psi'}(\pi^{-1}) = \zeta^{1-Q}\mu_\psi(\pi^{-1}) = \zeta^{1-Q},$$

so ζ is in \mathbf{F}_Q^\times . Conversely, if ϕ is ϵ -normalised, then for every ζ in \mathbf{F}_Q^\times the module $\zeta\phi\zeta^{-1}$ is as well. Lastly, recall that the automorphisms of ψ are the elements of \mathbf{F}_q^\times , so the result of the lemma follows. \square

Theorem 5.32. *Let ϵ be a sign function. Then the set X of ϵ -normalised Drinfeld modules is a torsor for the group $\text{Cl}^+(A)$, the quotient of the ideal group of A by the principal ideals (a) with $\epsilon(a) = 1$.*

Proof. Let ϕ be a ϵ -normalised Drinfeld module. For every non-zero ideal \mathfrak{a} of A and every x in A we have

$$(\mathfrak{a} * \phi)(x)\phi(\mathfrak{a}) = \phi(\mathfrak{a})\phi(x),$$

by the definition of $\mathfrak{a} * \phi$. Comparing the leading coefficients, we conclude that $\mu_{\mathfrak{a} * \phi}(x) = \mu_\phi(x)^{q^{\deg(\phi(\mathfrak{a}))}}$ holds for all x in A . From corollary 4.5 we know that the degree of $\phi(\mathfrak{a})$ is divisible by $\deg(\infty)$. As μ_ϕ lands in \mathbf{F}_Q^\times , we conclude that $\mu_{\mathfrak{a} * \phi}(x) = \mu_\phi(x)$ holds for all x in A . It follows that $\mathfrak{a} * \phi$ is also ϵ -normalised.

It remains to show that an ideal \mathfrak{a} acts trivially if and only if it has a principal generator a with $\epsilon(a) = 1$. From corollary 5.23 we know that an ideal fixes the isomorphism class of the Drinfeld module if and only if it is principal. Remark 5.2 states that (a) acts on ϕ via conjugation with $\mu_\phi(a)^{-1}$. If ϕ is ϵ -normalised, $\mu_\phi(a)$ is $\iota(\epsilon(a))$, so (a) acts trivially if and only if $\epsilon(a) = 1$. \square

6 Abelian Extensions

In this chapter we use the ϵ -normalised Drinfeld modules, which we have constructed in the previous chapter, to explicitly write down certain class fields of the field k .

- **The normalising field**

Proposition 6.1. *Let ϕ be an ϵ -normalised Drinfeld module. Then the smallest field over which ϕ is defined is independent of ϕ . We call this field the normalising field for ϵ -normalised Drinfeld modules and we write H_A^+ for it.*

Proof. We already know that this field is generated over k by the coefficients of $\phi(a)$ for one such Drinfeld module ϕ and some a in A non-constant. By lemma 5.26 and the fact that the action of the ideals is transitive, this field is independent of the choice of ϕ . \square

Lemma 6.2. *Let σ in $\text{Gal}(C/k)$ and ϕ a Drinfeld module over C . Then $(\sigma\phi)(x) = \sigma(\phi(x))$ is also a Drinfeld module over C . If ϕ is ϵ -normalised, so is $\sigma\phi$. The action of $\text{Gal}(C/k)$ on X , the set of ϵ -normalised Drinfeld modules, is compatible with the action of $\text{Cl}^+(A)$.*

Proof. It is clear that $\sigma\phi$ is again a Drinfeld module over C . Moreover, we see $\mu_{\sigma\phi} = \sigma\mu_\phi$ and therefore we also have $\iota_{\sigma\phi} = \sigma\iota_\phi$. If ϕ is ϵ -normalised, we conclude that

$$\mu_{\sigma\phi} = \sigma\mu_\phi = \sigma\iota_\phi\epsilon = \iota_{\sigma\phi}\epsilon$$

holds and therefore $\sigma\phi$ is also ϵ -normalised. To check that the action is compatible with that of the ideals, one simply looks at the formulas. \square

Corollary 6.3. *The extension H_A^+/k is Abelian with Galois group isomorphic to a subgroup of $\text{Cl}^+(A)$.*

Proof. It follows from the previous lemma that H_A^+ is finite and normal over k . Let ϕ be an ϵ -normalised Drinfeld module. Note that H_A^+ contains the minimal field of definition L_ϕ of ϕ . This field is contained in K and finite, so it is algebraic and therefore separable over k . Pick ξ in C such that $\phi' = \xi\phi\xi^{-1}$ is defined over L_ϕ . Let x be a non-constant element of A with $\epsilon(x) = 1$. Then we have

$$\mu_{\phi'}(x) = \xi^{1-q^{\deg(\infty)\text{ord}_\infty(x)}}\mu(x) = \xi^{1-q^n}$$

for some positive integer n . As $\mu_{\phi'}(x)$ is in L_ϕ , this implies that $L_\phi(\xi)$ is a separable extension of L_ϕ . Since H_A^+ is contained in $L_\phi(\xi)$ it is also separable over L_ϕ and therefore over k .

As the action of the Galois group $\text{Gal}(H_A^+/k)$ on X commutes with the action of $\text{Cl}^+(A)$ —for which this set is a torsor—we obtain an injective group homomorphism

$$\text{Gal}(H_A^+/k) \longrightarrow \text{Cl}^+(A).$$

It follows that the extension is Abelian. \square

- **Reduction of Drinfeld Modules**

Let L be a subfield of C which contains k and fix a non-zero prime ideal of the integral closure of A in L . Let \mathcal{O} be the corresponding valuation ring of L and \mathfrak{p} its maximal ideal. Write κ for the residue field \mathcal{O}/\mathfrak{p} . Let ρ be the reduction map $\mathcal{O} \rightarrow \kappa$. Note that it induces a reduction map $\mathcal{O}\{\tau\} \rightarrow \kappa\{\tau\}$ by reducing coefficient-wise, which we shall again call ρ . Let δ_r be the composed map $A \subset \mathcal{O} \rightarrow \kappa$ and D_r be the derivative map $\kappa\{\tau\} \rightarrow \kappa$.

Suppose that ϕ is a Drinfeld module defined over L . Then we say that ϕ has coefficients in \mathcal{O} if for all x in A , $\phi(x)$ is a polynomial whose coefficients lie in \mathcal{O} . In this case, we put ϕ_r equal to $\rho \circ \phi$, the reduction of ϕ at \mathfrak{p} . Note that for all f in $\mathcal{O}\{\tau\}$ we have $D_r(\rho(f)) = \rho(D(f))$. We conclude that $D\phi_r = \delta_r$ holds. In general, ϕ_r need not be a Drinfeld module, as $\phi_r(x)$ could be a constant polynomial for all x .

Definition 6.4. Let ϕ be a Drinfeld module defined over L . Then we say ϕ has stable reduction at \mathfrak{p} if there is a ϕ' isomorphic to ϕ over L which has coefficients in \mathcal{O} , such that ϕ'_r is a Drinfeld module.

Lemma 6.5. *Let ϕ be a Drinfeld module with stable reduction at \mathfrak{p} . Then the reduced module is unique up to isomorphism.*

Proof. Suppose ϕ' and ϕ'' are both isomorphic to ϕ , that they both have coefficients in \mathcal{O} and that both ϕ'_r and ϕ''_r are Drinfeld modules. As ϕ' and ϕ'' are isomorphic, there is a λ in L^\times such that we have $\phi'' = \lambda\phi'\lambda^{-1}$. Our aim is to show that λ is in \mathcal{O}^\times , so that it reduces to an isomorphism $\rho(\lambda)$ between ϕ'_r and ϕ''_r .

As ϕ'_r is a Drinfeld module, there is an x in A such that $\phi'(x)$ is a polynomial $\sum_i x'_i \tau^i$ with x'_j in \mathcal{O}^\times for some $j > 0$. Note that

$$\phi''(x) = \lambda\phi'(x)\lambda^{-1} = \sum_i \lambda^{1-q^i} x'_i \tau^i$$

is in $\mathcal{O}\{\tau\}$, so $\lambda^{1-q^j} x'_j$ is in \mathcal{O} and therefore λ^{1-q^j} is in \mathcal{O} . We conclude λ^{-1} is a root of the polynomial $x^{q^j-1} - \lambda^{1-q^j}$ with coefficients in \mathcal{O} , so λ^{-1} is integral over \mathcal{O} . Being a valuation ring, \mathcal{O} is integrally closed, so λ^{-1} is in \mathcal{O} .

The same holds if we reverse the roles of ϕ' and ϕ'' , in this case the isomorphism is λ^{-1} , so λ is also in \mathcal{O} and therefore in \mathcal{O}^\times . □

Lemma 6.6. *Every Drinfeld module ϕ defined over L has potentially stable reduction at \mathfrak{p} , that is, there is a finite extension M of L and a prime \mathfrak{q} that lies above \mathfrak{p} , such that ϕ has stable reduction at \mathfrak{q} .*

Proof. Let M be a finite extension of L , \mathfrak{q} a prime that lies above \mathfrak{p} and \mathcal{O}_M the corresponding valuation ring of L . Let v be the normalised valuation on M corresponding to \mathfrak{q} , so that \mathcal{O}_M consists of the x in M with $v(x) \geq 0$ and \mathfrak{q} consists of the x in M with $v(x) > 0$.

Let ϕ be a Drinfeld module defined over L and let λ be a non-zero element of M^\times . Write ϕ' for the Drinfeld module $\lambda\phi\lambda^{-1}$, which is defined over M . Write $\phi'_n(x)$ for the coefficient

at τ^n of $\phi'(x)$. We have to find an M , a \mathfrak{q} and a λ such that $v(\phi'_n(x))$ is non-negative for all $n \geq 0$ and x in A and $v(\phi'_n(x)) = 0$ for some x in A and some $n \geq 1$.

Let e be the ramification index of \mathfrak{q} over \mathfrak{p} . Note that for all x in A and all $n \geq 0$ we have

$$v(\phi'_n(x)) = v(\lambda^{1-q^n} \phi_n(x)) = ev_{\mathfrak{p}}(\phi_n(x)) - v(\lambda)(q^n - 1),$$

where $v_{\mathfrak{p}}$ is the normalised valuation of L at \mathfrak{p} and $\phi_n(x)$ is the coefficient of $\phi(x)$ at τ^n .

Recall that A is a finitely generated \mathbf{F}_q -algebra. Let x_1, \dots, x_s be a set of generators of A as an \mathbf{F}_q -algebra. It suffices to show that we can pick M , \mathfrak{p} and λ such that $v(\phi'_n(x_i)) \geq 0$ for all $i = 1, \dots, s$ and all $n \geq 0$ and $v(\phi'_n(x_i)) = 0$ for some $1 \leq i \leq s$ and $n \geq 1$.

Let $j \geq 1$ and $1 \leq m \leq s$ be such that

$$w = \frac{v_{\mathfrak{p}}(\phi_j(x_m))}{q^m - 1}$$

is minimal. Pick M and \mathfrak{q} in such a way that e is a multiple of the denominator of w and pick λ in M such that $v(\lambda) = ew$. For all $n \geq 1$ and $1 \leq i \leq s$ we now have

$$v(\phi'_n(x_i)) = ev_{\mathfrak{p}}(\phi_n(x_i)) - v(\lambda)(q^n - 1) = e(v_{\mathfrak{p}}(\phi_n(x_i)) - w(q^n - 1))$$

and

$$\frac{v_{\mathfrak{p}}(\phi_n(x_i))}{q^n - 1} \geq w$$

so $v(\phi'_n(x))$ is non-negative. Moreover, we have equality for $i = j$ and $n = m$, so $v(\phi'_m(x_j)) = 0$. \square

Definition 6.7. Let ϕ be a Drinfeld module defined over L . Then we say that ϕ has good reduction at \mathfrak{p} if it has stable reduction at \mathfrak{p} and the rank of the reduced module is the same as that of ϕ .

Note that for Drinfeld modules of rank 1, stable and good reduction are the same thing, as the rank of the reduced module is clearly at most the rank of the original module and every Drinfeld module has positive rank.

Lemma 6.8. Let ϕ be a rank 1 Drinfeld module defined over L and suppose that there is a non-zero x in A such that $\mu_{\phi}(x)$ is in \mathcal{O}^{\times} . Then ϕ has coefficients in \mathcal{O} and the reduction is good.

Proof. From lemma 6.6 we know that there is a finite extension M of L and a prime \mathfrak{q} above \mathfrak{p} such that ϕ has good reduction at \mathfrak{q} . So there is a λ in M such that $\phi' = \lambda\phi\lambda^{-1}$ has coefficients in \mathcal{O}_M . Moreover, as ϕ' also has rank 1, the leading coefficient $\mu'_{\phi}(x)$ is in \mathcal{O}_M^{\times} . As we have $\mu'_{\phi}(x) = \lambda^{1-N(x)}\mu_{\phi}(x)$ we see that $\lambda^{1-N(x)}$ is in \mathcal{O}_M^{\times} . It follows that λ itself is also in \mathcal{O}_M^{\times} , as \mathcal{O}_M is integrally closed.

Now we have for all y in A that $\phi(a) = \lambda^{-1}\phi'(a)\lambda$ has coefficients in \mathcal{O}_M . But it also has coefficients in L , so it has coefficients in $\mathcal{O} = \mathcal{O}_M \cap L$. Moreover, $\mu'_{\phi}(y)$ is in \mathcal{O}_M^{\times} for all y (as the reduction of ϕ' is good), so $\mu_{\phi}(y) = \lambda^{N(y)-1}\mu'_{\phi}(y)$ is in \mathcal{O}_M^{\times} . It is also in L , so it is in $\mathcal{O}^{\times} = \mathcal{O}_M^{\times} \cap L$. This is what we wanted to show. \square

Corollary 6.9. *Let ϕ be an ϵ -normalised Drinfeld module. Then, for all x in A , $\phi(x)$ has coefficients in B and the leading coefficient of $\phi(x)$ is in B^\times .*

Proof. Recall that B , the integral closure of A in H_A^+ , is the intersection of all the valuation rings of H_A^+ at primes that lie above a prime of A . The result now follows at once from the previous lemma. \square

Lemma 6.10. *Let ϕ and ψ be two ϵ -normalised Drinfeld modules and let \mathfrak{q} be a non-zero prime of B . Let ϕ_r and ψ_r be the reductions of ϕ and ψ at \mathfrak{q} . Then ϕ_r is equal to ψ_r if and only if ϕ is equal to ψ .*

Proof. Let \mathfrak{p} be the prime ideal of A that lies below \mathfrak{q} . By theorem 5.32, there is an ideal \mathfrak{a} of A such that ψ is $\mathfrak{a} * \phi$. We show that we can choose \mathfrak{a} such that it is coprime to \mathfrak{p} .

Let x be an element of k such that x is congruent to 1 mod ∞ and $v_{\mathfrak{p}}(x) = -v_{\mathfrak{p}}(\mathfrak{a})$. Then the ideal $x\mathfrak{a}$ is a fractional ideal of A coprime to \mathfrak{p} . So there are ideals \mathfrak{d} and \mathfrak{n} , coprime to \mathfrak{p} , such that $x\mathfrak{a} = \mathfrak{n}\mathfrak{d}^{-1}$. In $\text{Cl}^+(A)$ we have $(x) = 1$, so

$$\mathfrak{a} = x\mathfrak{a} = \mathfrak{n}\mathfrak{d}^{-1} = \mathfrak{n}\mathfrak{d}^{\#X-1}$$

holds and we may assume that \mathfrak{a} is coprime to \mathfrak{p} .

By definition of $\phi(\mathfrak{a})$ we have that $\psi(x)\phi(\mathfrak{a}) = \phi(\mathfrak{a})\phi(x)$ holds for all x in A . Reducing this equation mod \mathfrak{q} , we see that

$$\phi_r(x)\phi(\mathfrak{a}) \equiv \phi(\mathfrak{a})\phi_r(x) \pmod{\mathfrak{q}}.$$

As ϕ_r has rank 1, we know that $\text{End}(\phi_r)$ is isomorphic to A , so there is an y in A such that

$$\phi(\mathfrak{a}) \equiv \phi(y) \pmod{\mathfrak{q}}.$$

As $\phi(\mathfrak{a})$ is monic, we conclude that $\mu(y)$ is 1 mod \mathfrak{q} and as $\mu(y)$ is also in \mathbf{F}_Q^\times , we know that in fact $\mu(y) = 1$, so $\epsilon(y) = 1$. Therefore, we are done if we can show that $\mathfrak{a} = (y)$.

We do this by comparing the torsion modules of ϕ_r . Put $\mathfrak{b} = \mathfrak{a} + (y)$. Note that $\phi_r(\mathfrak{a}) = \phi_r(y) = \phi_r(\mathfrak{b})$. Choose an algebraic closure of B/\mathfrak{q} and let Φ_r be the A -module structure on this algebraic closure induced by ϕ_r . As \mathfrak{a} , \mathfrak{b} and y are coprime to \mathfrak{p} and ϕ_r has rank 1, we have $\Phi_r[\mathfrak{a}] \cong A/\mathfrak{a}$, $\Phi_r[\mathfrak{b}] \cong A/\mathfrak{b}$ and $\Phi_r[y] \cong A/(y)$. But $\Phi_r[\mathfrak{a}]$, $\Phi_r[\mathfrak{b}]$ and $\Phi_r[y]$ are all equal to the kernel of $\phi_r(y)$, so \mathfrak{a} , \mathfrak{b} and (y) all have the same norm, so $\mathfrak{a} = \mathfrak{b} = (y)$, as required. \square

Theorem 6.11. *The normalising field H_A^+ is unramified at every non-zero prime of A . For a non-zero ideal \mathfrak{a} of A denote by $\sigma_{\mathfrak{a}}$ the image of \mathfrak{a} in $\text{Gal}(H_A^+/k)$ under the Artin map. Let ϕ be an ϵ -normalised Drinfeld module. Then we have $\sigma_{\mathfrak{a}}\phi = \mathfrak{a} * \phi$ for all non-zero ideals \mathfrak{a} of A . Therefore, $\text{Gal}(H_A^+/k)$ is naturally isomorphic with $\text{Cl}^+(A)$.*

Proof. Let \mathfrak{p} be a non-zero prime of A . Let σ be in the inertia group of \mathfrak{p} . Let \mathfrak{q} be a prime of B that lies above \mathfrak{p} . Then, by definition, σ induces the identity mod \mathfrak{q} . So for any ϵ -normalised Drinfeld module ϕ , we have

$$\sigma\phi \equiv \phi \pmod{\mathfrak{q}}$$

and therefore $\sigma\phi = \phi$, by the previous lemma. But H_A^+ is generated by the coefficients of any $\phi(y)$ with y in A non-constant, so this implies that σ is the identity on H_A^+ . We conclude that the inertia group of \mathfrak{p} is trivial, that is, \mathfrak{p} is unramified.

To check that $\sigma_{\mathfrak{a}}\phi = \mathfrak{a} * \phi$ for every non-zero ideal \mathfrak{a} of A , it suffices to check this for non-zero prime ideals, by lemma 5.3. So let \mathfrak{p} be a non-zero prime ideal of A and \mathfrak{q} a prime of B that lies above \mathfrak{p} . By the previous lemma, it is enough to check that

$$\sigma_{\mathfrak{p}}\phi \equiv \mathfrak{p} * \phi \pmod{\mathfrak{q}}.$$

Note that modulo \mathfrak{q} , $\sigma_{\mathfrak{p}}$ is just the Frobenius-endomorphism, i.e. the $q^{\deg(\mathfrak{p})}$ -th powering map.

Write ϕ_r for the reduction of ϕ at \mathfrak{q} . Note that the kernel of the map $A \rightarrow B/\mathfrak{q}$ is \mathfrak{p} and that ϕ_r must have non-zero height. Since ϕ has rank 1, the height must be 1. So $\phi_r(\mathfrak{p})$ is a monic polynomial in τ of degree

$$v(p) = -r \deg(\infty) \text{ord}_{\infty}(\mathfrak{p}) = \deg(\mathfrak{p}) \text{ord}_{\mathfrak{p}}(\mathfrak{p}) = \deg(\mathfrak{p})$$

and the lowest degree of a non-zero term is

$$j(p) = h \deg(\mathfrak{p}) \text{ord}_{\mathfrak{p}}(\mathfrak{p}) = \deg(p).$$

We conclude that $\phi_r(\mathfrak{p}) = \tau^{\deg(p)}$. It follows that

$$(\mathfrak{p} * \phi_r)(x) \tau^{\deg(p)} = \tau^{\deg(p)} \phi_r(x)$$

holds for all x in A . But we know that pulling $\tau^{\deg(p)}$ through a polynomial results in all the coefficients being raised to the $q^{\deg(p)}$ -th power. So we have

$$\sigma_{\mathfrak{p}}\phi \equiv \mathfrak{p} * \phi \pmod{\mathfrak{q}}.$$

as required.

From corollary 6.3 we know that there is an injective map $\text{Gal}(H_A^+/k) \rightarrow \text{Cl}^+(A)$. The result we have just derived implies that this map is also onto. \square

- **The Hilbert Class Field**

In this section we derive a variant of the main theorem we proved in the previous section. This variant is more natural from the point of view of class field theory.

Let ϕ be an ϵ -normalised Drinfeld module. Let L be the minimal field of definition of ϕ . Following the proof of theorem 5.27 we let ζ in C be such that $\phi_{\min} = \zeta\phi\zeta^{-1}$ is defined over L . The greatest common divisor g from that proof must be $q - 1$, as any g -th root of unity in C will be an automorphism of ϕ and $\text{Aut}(\phi)$ is isomorphic to \mathbf{F}_q^\times . We conclude that ζ^g is in H_A^+ , the smallest field over which ϕ is defined. Moreover, we see that H_A^+ is equal to $L(\zeta^g)$, as, for every y in A , the coefficients of $\phi(y)$ can be expressed in terms of those of ϕ_{\min} and powers of ζ^g . Let $\pi \in A$ have $\text{ord}_\infty(\pi) = -1$ and $\epsilon(\pi) = 1$. Then we have

$$1 = \mu_\phi(\pi) = \zeta^{1-q^{-\deg(\infty)\text{ord}_\infty(\pi)}} \mu_{\phi_{\min}}(\pi) = \zeta^{1-Q} \mu_{\phi_{\min}}(\pi)$$

and therefore, ζ^{Q-1} is in L . We conclude that $[H_A^+ : L]$ is at most $(Q - 1)/(q - 1)$.

By lemma 5.31, the group $\mathbf{F}_Q^\times/\mathbf{F}_q^\times$ acts on the set X of ϵ -normalised Drinfeld modules and the orbits consists precisely of all the ϵ -normalise Drinfeld modules in a single class. From theorem 5.32 and corollary 5.23 we conclude that there is a short exact sequence of Abelian groups

$$0 \longrightarrow \mathbf{F}_Q^\times/\mathbf{F}_q^\times \longrightarrow \text{Cl}^+(A) \longrightarrow \text{Cl}(A) \longrightarrow 0.$$

Moreover, we know from theorem 6.11 that $\text{Cl}^+(A)$ is isomorphic to $\text{Gal}(H_A^+/k)$.

Let y be a non-constant element of A . An element λ of \mathbf{F}_Q^\times will map an ϵ -normalised Drinfeld module ϕ to an isomorphic module $\phi' = \lambda\phi\lambda^{-1}$. This element gives an element σ_λ of the Galois group by saying that it sends the coefficients of $\phi(y)$, which generate H_A^+ , to the coefficients of $\phi'(y)$. From the proof of theorem 5.27 we see that this implies that σ_λ fixes the generators of the minimal field of definition of ϕ . We conclude that the minimal field of definition of ϕ is fixed by the subgroup $\mathbf{F}_Q^\times/\mathbf{F}_q^\times$ of $\text{Gal}(H_A^+/k)$. By the degree estimate we made earlier, it cannot be fixed by any larger subgroup.

We see that there is a single minimal field of definition for all rank 1 Drinfeld modules over C . We call this field H_A . It is an Abelian extension of k and $\text{Gal}(H_A/k)$ is isomorphic to $\text{Cl}(A)$. As H_A is a subfield of H_A^+ it is unramified at all non-zero primes of A . As it is a subfield of K as per lemma 5.24, H_A is completely split at ∞ . By class field theory, we therefore know that H_A is the maximal Abelian extension of k with these properties, which we call the Hilbert class field of A .

The extension H_A^+/H_A is generated by the element ζ^{q-1} from before. Recall that ζ^{Q-1} is in H_A and that the degree of the extension is $(Q - 1)/(q - 1)$. As \mathbf{F}_Q is contained H_A , the $(Q - 1)/(q - 1)$ -th roots of unity are in H_A . We see that H_A^+/H_A is a Kummer extension.

- **Ramified extensions**

Let \mathfrak{m} be a non-zero proper ideal of A . In this section we describe how to obtain class fields with conductor \mathfrak{m} using Drinfeld modules over C . Again we first fix a sign function ϵ and work with ϵ -normalised modules. Afterwards we ‘fix’ the behaviour at ∞ .

Let $I_{\mathfrak{m}}$ be the group of ideals of A that are coprime to \mathfrak{m} . Let $P_{\mathfrak{m}}$ be the subgroup of principal ideals (x) with $x \equiv 1 \pmod{\mathfrak{m}}$ and $P_{\mathfrak{m}}^+$ the subgroup of $P_{\mathfrak{m}}$ of those principal ideals (x) that have $\epsilon(x) = 1$. We call the quotient $\text{Cl}_{\mathfrak{m}}(A) = I_{\mathfrak{m}}/P_{\mathfrak{m}}$ the ray class group modulo \mathfrak{m} and $\text{Cl}_{\mathfrak{m}}^+(A) = I_{\mathfrak{m}}/P_{\mathfrak{m}}^+$ the narrow ray class group modulo \mathfrak{m} relative to ϵ .

Let ϕ be an ϵ -normalised Drinfeld module. As before, we write $\Phi[\mathfrak{m}]$ for the set of \mathfrak{m} -torsion points of ϕ in C . Note that $\Phi[\mathfrak{m}]$ is a cyclic A -module, isomorphic to A/\mathfrak{m} . This module has $\#(A/\mathfrak{m})^\times$ elements that generate it as an A -module. Let $X_{\mathfrak{m}}$ be the set of pairs (ϕ, λ) where ϕ is an ϵ -normalised Drinfeld module and λ is a generator of $\Phi[\mathfrak{m}]$ in C .

Lemma 6.12. *Let an ideal \mathfrak{a} of A coprime to \mathfrak{m} , an ϵ -normalised Drinfeld module ϕ and a generator λ of the \mathfrak{m} -torsion of ϕ be given. Then $[\phi(\mathfrak{a})](\lambda)$ is a generator of the \mathfrak{m} -torsion of $\mathfrak{a} * \phi$.*

Proof. Let \mathfrak{b} be an ideal coprime to \mathfrak{m} such that $\mathfrak{a}\mathfrak{b}$ is a principal ideal generated by x , where $x \in A$ is congruent to 1 modulo \mathfrak{m} and has $\epsilon(x) = 1$. Write ψ for the ϵ -normalised Drinfeld module $\mathfrak{a} * \phi$. Note that $\mathfrak{b} * \psi$ is ϕ .

It is clear that $\phi(\mathfrak{a})$ induces a map from the \mathfrak{m} -torsion of ϕ to that of ψ . Also, $\psi(\mathfrak{b})$ induces a map from the \mathfrak{m} -torsion of ψ to that of ϕ . The compositions satisfy

$$\psi(\mathfrak{b})\phi(\mathfrak{a}) = (\mathfrak{a} * \phi)(\mathfrak{b})\phi(\mathfrak{a}) = \phi(\mathfrak{a}\mathfrak{b}) = \mu_{\phi}(x)^{-1}\phi(x) = \phi(x)$$

and

$$\phi(\mathfrak{a})\psi(\mathfrak{b}) = (\mathfrak{b} * \psi)(\mathfrak{a})\psi(\mathfrak{b}) = \psi(\mathfrak{a}\mathfrak{b}) = \mu_{\psi}(x)^{-1}\psi(x) = \psi(x)$$

and are therefore both the identity. We conclude that $\phi(\mathfrak{a})$ is an isomorphism from the \mathfrak{m} -torsion of ϕ to that of ψ and therefore maps a generator of the one to a generator of the other. \square

Theorem 6.13. *Let \mathfrak{a} be an ideal of A coprime to \mathfrak{m} and (ϕ, λ) in $X_{\mathfrak{m}}$. Then we define $\mathfrak{a} * (\phi, \lambda)$ as $(\mathfrak{a} * \phi, [\phi(\mathfrak{a})](\lambda))$. This operation induces an action of $\text{Cl}_{\mathfrak{m}}^+(A)$ on $X_{\mathfrak{m}}$. Under this action, $X_{\mathfrak{m}}$ is a torsor for $\text{Cl}_{\mathfrak{m}}^+(A)$.*

Proof. We first identify precisely when an ideal \mathfrak{a} of A acts trivially on $X_{\mathfrak{m}}$. From theorem 5.32 we know that \mathfrak{a} fixes an ϵ -normalised Drinfeld module ϕ if and only if \mathfrak{a} is a principal ideal with a generator x such that $\epsilon(x) = 1$. On the \mathfrak{m} -torsion, such an x acts, by definition, through multiplication by x . The only way for x to leave a generator λ of this torsion module fixed, is if x is congruent to 1 modulo \mathfrak{m} .

We conclude that the operation descends to an action of $\text{Cl}_{\mathfrak{m}}^+(A)$ on $X_{\mathfrak{m}}$ which is faithful, as we have precisely divided out the ideals that do nothing. To show that it is in fact a torsor, we count elements on either side. The number of elements of $X_{\mathfrak{m}}$ is the number of

ϵ -normalised Drinfeld modules times the number of generators for the \mathfrak{m} -torsion of one such module, Similarly, we see that there is an exact sequence

$$0 \longrightarrow (A/\mathfrak{m})^\times \longrightarrow \text{Cl}_\mathfrak{m}^+(A) \longrightarrow \text{Cl}^+(A) \longrightarrow 0$$

so the number of elements of $\text{Cl}_\mathfrak{m}^+(A)$ is the number of elements of $\text{Cl}^+(A)$ times the number of invertible elements in A/\mathfrak{m} . We conclude that the two sets have the same size. \square

Proposition 6.14. *Let ϕ be an ϵ -normalised Drinfeld module. Then the field $K_\mathfrak{m}^+ = H_A^+(\Phi[\mathfrak{m}])$ is independent of the choice of ϕ . It is finite Galois over k and the Galois group is a subgroup of $\text{Cl}_\mathfrak{m}^+(A)$ and therefore Abelian.*

Proof. Let \mathfrak{a} be an ideal of A coprime to \mathfrak{m} . Then we know from lemma 6.12 that $\phi(\mathfrak{a})$ is an isomorphism from the \mathfrak{m} -torsion of ϕ to that of $\mathfrak{a} * \phi$. As $\phi(\mathfrak{a})$ is a polynomial with coefficients in H_A^+ , we see that the field generated by the \mathfrak{m} -torsion points of $\mathfrak{a} * \phi$ is contained in $H_A^+(\Phi[\mathfrak{m}])$. By the transitivity of the action of $\text{Cl}_\mathfrak{m}^+(A)$ it follows that the field does not depend on the chosen ϕ .

Note that $K_\mathfrak{m}^+$ is obtained by adjoining all the roots of the polynomial $\phi(\mathfrak{m})$ with coefficients in H_A^+ to H_A^+ and is therefore finite Galois over H_A^+ . As this field is itself finite Galois over k , we conclude that the $K_\mathfrak{m}^+$ is finite Galois over k .

Let (ϕ, λ) be in $X_\mathfrak{m}$ and σ in $\text{Gal}(K_\mathfrak{m}^+/k)$. Then $\sigma\phi$ is again an ϵ -normalised Drinfeld module. Moreover σ sends the \mathfrak{m} -torsion of ϕ to that of $\sigma\phi$, so it sends λ to a generator of the \mathfrak{m} -torsion of $\sigma\phi$. We conclude that $\sigma(\phi, \lambda) := (\sigma\phi, \sigma(\lambda))$ gives an action of $\text{Gal}(K_\mathfrak{m}^+/k)$ on $X_\mathfrak{m}$. From the defining formulas we see that this action commutes with the action of $\text{Cl}_\mathfrak{m}^+(A)$. Therefore there is an injective map from $\text{Gal}(K_\mathfrak{m}^+/k)$ as required. \square

Lemma 6.15. *Let \mathcal{O} be the integral closure of A in $K_\mathfrak{m}^+$, let ϕ be an ϵ -normalised Drinfeld module and let \mathfrak{p} be a non-zero prime ideal of \mathcal{O} that does not divide \mathfrak{m} . Then the \mathfrak{m} -torsion of ϕ is contained in \mathcal{O} and the reduction map to the \mathfrak{m} -torsion of the reduction of ϕ at \mathfrak{p} is bijective.*

Proof. Recall that B is the integral closure of A in H_A^+ . The coefficients of $\phi(x)$ are in B for every x in A . Therefore, the coefficients of $\phi(\mathfrak{m})$ are in B . Moreover, $\phi(\mathfrak{m})$ is monic, so the roots of $\phi(\mathfrak{m})$ are integral over B and therefore over A . These roots lie in $K_\mathfrak{m}^+$, so they are in \mathcal{O} .

Write $\phi_\mathfrak{p}$ for the reduction of ϕ modulo \mathfrak{p} . The polynomial $\phi(\mathfrak{m})$, which is separable, factors into distinct linear factors in \mathcal{O} . So $\phi_\mathfrak{p}(\mathfrak{m})$ also factors into linear factors mod \mathfrak{p} , however, they need no longer be distinct. However, as \mathfrak{m} is coprime to \mathfrak{p} , we know that $\phi_\mathfrak{p}(\mathfrak{m})$ is a separable polynomial, so there are no double factors. We conclude that the reduction map is bijective as required. \square

Theorem 6.16. *The field K_m^+ is unramified at all the non-zero primes of A that are coprime to m . For every ideal \mathfrak{a} of A coprime to m and every (ϕ, λ) in X_m we have*

$$\sigma_{\mathfrak{a}}(\phi, \lambda) = \mathfrak{a} * (\phi, \lambda),$$

where $\sigma_{\mathfrak{a}}$ is the image of \mathfrak{a} in $\text{Gal}(K_m^+/k)$ under the Artin map. It follows that $\text{Gal}(K_m^+/k)$ is isomorphic to $\text{Cl}_m^+(A)$.

Proof. Let \mathfrak{p} be a non-zero prime of A that is coprime to m . Let \mathfrak{q} be an extension prime of \mathfrak{p} in \mathcal{O} . Suppose that σ is in the inertia group of \mathfrak{p} . Let (ϕ, λ) be in X . Write ϕ_r and λ_r for the reductions of ϕ and λ modulo \mathfrak{q} .

By assumption σ acts as the identity modulo \mathfrak{q} , so we have $\sigma\phi_r = \phi_r$ and $\sigma(\lambda_r) = \lambda_r$. The former implies, by lemma 6.10 that $\sigma\phi = \phi$ and then the latter implies that $\sigma(\lambda) = \lambda$ by the previous lemma. We conclude that σ acts trivially on X and is therefore 1. Thus the inertia group at \mathfrak{p} is trivial, i.e., K_m^+ is unramified at \mathfrak{p} .

From the proof of theorem 6.11, we know that $\phi_r(\mathfrak{p})$ acts as $\sigma_{\mathfrak{p}}$ modulo \mathfrak{p} , so we see that the required relation holds for prime ideals coprime to m . This implies that it is true for all ideals of A that are coprime to m . \square

Just as in the case of unramified extensions we have a natural copy of $\mathbf{F}_Q^\times/\mathbf{F}_q^\times$ inside $\text{Gal}(K_m^+/k)$ and $\text{Cl}_m^+(A)$. In fact, cf. [Hayes2], this subgroup is both the decomposition and inertia group at ∞ . It follows that K_m , the fixed field, is unramified at all the primes of A coprime to m and totally split at ∞ . Its Galois group is isomorphic to $\text{Cl}_m(A)$ under the Artin map. The field K_m is the maximal Abelian extension of k which is completely split at ∞ , unramified outside m and such that a prime splits completely if and only if it is congruent to 1 modulo m , i.e., it is a ray class field of A .

7 References

[A-M]

Atiyah, MacDonald, *Introduction to Commutative Algebra*
Addison-Wesley, 1969

[CM67]

AMS, *Current Trends in Arithmetical Algebraic Geometry*
Series: Contemporary Mathematics, 67
American Mathematical Society, 1987

[D-H]

Deligne, Husemöller, *Survey of Drinfel'd modules*
in [CM67], pp. 25–92

[Dri]

Drinfeld, *Elliptic Modules*
in Math. USSR Sbornic, 23 (1974), pp. 561–592

[GHR]

Goss, Hayes, Rosen (eds.), *The Arithmetic of Function Fields*
Series: Ohio State University Mathematical Research Institute Publications, 2
de Gruyter, 1992

[Goss]

Goss, *Basic Structures of Function Field Arithmetic*
Series: Ergebnisse der Mathematik und ihrer Grenzgebiete (3. Folge), 35
Springer-Verlag, 1996

[Har]

Hartshorne, *Algebraic Geometry*
Series: Graduate Texts in Mathematics, 52
Springer-Verlag, 1997

[Hayes1]

Hayes, *A Brief Introduction to Drinfeld Modules*
in [GHR], pp. 1–32

[Hayes2]

Hayes, *Stickelberger elements in function fields*
in Composito Math, 55 (1985), pp. 209–239

[Hayes3]

Hayes, *Explicit Class Field Theory in Global Function Fields*
in Advances in Mathematics, Supplementary Studies, 6 (1979), pp. 173–218

[Mum]

Mumford, *Abelian Varieties*
Oxford Univ. Press, 1970

[Si]

Silverman, *The Arithmetic of Elliptic Curves*
Series: Graduate Texts in Mathematics, 106
Springer-Verlag, 1986