Drinfeld modules

Jared Weinstein

November 8, 2017

1 Motivation

Drinfeld introduced his so-called *elliptic modules* in an important 1973 paper[1] of the same title. He introduces this paper by observing the unity of three phenomena that appear in number theory:

1. The theory of cyclotomic extensions of $\mathbb{Q}$, and class field theory over $\mathbb{Q}$.
2. The theory of elliptic curves with complex multiplication, relative to an imaginary quadratic field.
3. The theory of elliptic curves in the large, over $\mathbb{Q}$.

To these, Drinfeld added a fourth:

4. The theory of Drinfeld modules over a function field.

Thus, Drinfeld modules are some kind of simultaneous generalization of groups of roots of unity (which are rank 1), but also of elliptic curves (which are rank 2, in the appropriate sense). Furthermore, Drinfeld modules can be any rank whatsoever; there is no structure that we know of which is an analogue of a rank 3 Drinfeld module over $\mathbb{Q}$.

In later work, Drinfeld extended his notion to a rather more general gadget called a *shtuka*[2]. Later, Laurent Lafforgue used the cohomology of moduli spaces of shtukas to prove the Langlands conjectures for $\mathrm{GL}(n)$, generalizing

[1] Try not to think too hard about the fact that Drinfeld was 20 years old that year.

[2] Russian slang for "thingy", from German *Stück*, "piece".

what Drinfeld had done for $n = 2$, and receiving a Fields medal in 2002 for those efforts. Whereas the Langlands conjectures are still wide open for number fields, even for GL(2) and even for $\mathbb{Q}$!

Before stating the definition of a Drinfeld module, it will be helpful to review items (1)-(3) above and highlight the common thread.

1.1 Cyclotomic fields

Let $n \geq 1$. We begin with the observation that the algebraic number $\zeta_n = e^{2\pi i/n}$ happens to be an algebraic number, in fact a root of $x^n - 1$. This is a small miracle and you should think deeply about why this is so. One interpretation is that the *multiplicative group* $\mathbf{G}_m$, a priori just an algebraic group over $\mathbb{Z}$, admits a *complex uniformization* $\mathbb{C} \to \mathbf{G}_m(\mathbb{C})$ given by $z \mapsto e^z$, whose kernel is the discrete subgroup $2\pi i\mathbb{Z} \subset \mathbb{C}$. Multiplication by n on $\mathbb{C}$ translates over to the map $z \mapsto z^n$ on $\mathbf{G}_m(\mathbb{C})$, which just happens to come from an endomorphism of the algebraic groups $\mathbf{G}_m$.

All of this is very basic, so let's move on to the arithmetic of cyclotomic extensions. The number field $\mathbb{Q}(\zeta_n)$ is abelian over $\mathbb{Q}$, and there is an isomorphism

$$r\colon \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \to (\mathbb{Z}/n\mathbb{Z})^\times$$

characterized by the equation $\sigma(\zeta_n) = \zeta_n^{r(n)}$, for all $\sigma \in \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$. This relation makes plain the following *reciprocity law*: For a prime p not dividing n, let Frob_p be the Frobenius element of $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$; then

$$r(\mathrm{Frob}_p) = p \pmod{n}.$$

This law tells you how primes decompose in $\mathbb{Q}(\zeta_n)/\mathbb{Q}$. For instance, p splits completely if and only if $\mathrm{Frob}_p = 1$, which is true if and only if $p \equiv 1 \pmod{n}$.

The reciprocity law can be phrased in terms of L-series. If χ is a Dirichlet character modulo n, then $\chi \circ r$ is a continuous homomorphism $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathbb{C}^\times$; that is, $\chi \circ r$ is an Artin character of dimension 1. The reciprocity law states that $L(\chi \circ r, s) = L(\chi, s)$. Furthermore, the *Kronecker-Weber theorem* states that every abelian extension of $\mathbb{Q}$ is contained in $\mathbb{Q}(\zeta_n)$ for some n. From this one can derive the fact that there is a bijection between Artin characters of dimension 1 and primitive Dirichlet characters, and this bijection preserves L-series. What we have here is a connection between a class of *Galois representations* (in this case, 1-dimensional Artin characters) and *automorphic representations* (in this case, Dirichlet characters).

1.2 Elliptic curves with complex multiplication

Hilbert's 12th problem asks whether special values of transcendental functions can be used to systematically construct abelian extensions of a given number field K, as the exponential function $\exp(2\pi iz)$ does for $\mathbb{Q}$. When K is an *imaginary quadratic field*, Kronecker's *Jugendtraum* provides an affirmative answer, in the form of modular functions. In modern language, we say that abelian extensions of K come from adjoining the torsion of elliptic curves with CM by (an order in) K.

Recall that a *lattice* in $\mathbb{C}$ is a discrete subgroup $\Lambda \subset \mathbb{C}$ which is free of rank 2 as an abelian group. Another way of saying this is that $\Lambda = \mathbb{Z}\alpha \oplus \mathbb{Z}\beta$, where $\{\alpha, \beta\}$ is a basis for $\mathbb{C}/\mathbb{R}$. Given a lattice Λ, there exists a corresponding elliptic curve $E_\Lambda/\mathbb{C}$ which is uniformized by $\mathbb{C}$, in the sense that there is an isomorphism of complex tori

$$\begin{aligned} \mathbb{C}/\Lambda &\to E_\Lambda(\mathbb{C}) \\ z &\mapsto (\wp(z), \wp'(z)). \end{aligned}$$

Explicitly, E_Λ is the projective cubic curve with affine equation

$$y^2 = 4x^3 - g_2(\Lambda)x - g_3(\Lambda),$$

where

$$\begin{aligned} g_2(\Lambda) &= 60 \sum_{\lambda \in \Lambda \setminus \{0\}} \lambda^{-4} \\ g_4(\Lambda) &= 140 \sum_{\lambda \in \Lambda \setminus \{0\}} \lambda^{-6} \end{aligned}$$

and

$$\wp(z) = \frac{1}{z^2} + \sum_{\lambda \in \Lambda \setminus \{0\}} \left[\frac{1}{(z-\lambda)^2} - \frac{1}{\lambda^2} \right].$$

It is something of a miracle that the abstractly defined complex torus $\mathbb{C}/\Lambda$ should admit an algebraic description, and it is exactly this miracle that is exploited for the Jugendtraum.

The association $\Lambda \mapsto E_\Lambda$ induces an equivalence between the following two $\mathbb{Z}$-linear categories:

1. Lattices $\Lambda \subset \mathbb{C}$, where $\mathrm{Hom}(\Lambda, \Lambda')$ is defined as $\left\{ \alpha \in \mathbb{C} \;\middle|\; \alpha\Lambda \subset \Lambda' \right\}$, and the composition law is multiplication in $\mathbb{C}$,

2. Elliptic curves over $\mathbb{C}$.

It is a good exercise to remind yourself why $\Lambda \mapsto E_\Lambda$ really is a functor, and why it is fully faithful. The functoriality means that if $\alpha \in \mathbb{C}$ satisfies $\alpha\Lambda \subset \Lambda'$, then the induced map of complex analytic groups $E_\Lambda(\mathbb{C}) \to E_{\Lambda'}(\mathbb{C})$ actually arises from a morphism $E_\Lambda \to E_{\Lambda'}$ of elliptic curves; the full faithfulness means that every such morphism arises this way.

Now let K be an imaginary quadratic field, considered as a subfield of $\mathbb{C}$, and let $\mathcal{O}_K$ be its ring of integers. The above equivalence of categories restricts to a bijection between:

1. Homothety class of lattices $\Lambda \subset \mathbb{C}$ which are stable under multiplication by $\mathcal{O}_K$; i.e. $\mathcal{O}_K\Lambda = \Lambda$,

2. Elliptic curves over $\mathbb{C}$ with complex multiplication (CM) by $\mathcal{O}_K$ (meaning that there exists $\mathcal{O}_K \to \operatorname{End} E$ whose derivative is the inclusion $\mathcal{O}_K \to \mathbb{C}$).

But the first set is finite: it is in bijection with the ideal class group of $\mathcal{O}_K$. Therefore so is the second set. Now we apply the miracle: having CM by $\mathcal{O}_K$ is a purely algebraic property, so that if E has CM by $\mathcal{O}_K$, and σ is an automorphism of $\mathbb{C}/K$, then $E^\sigma = E \times_{\mathbb{C},\sigma} \mathbb{C}$ also has CM by $\mathcal{O}_K$. Therefore E^σ only runs through finitely many isomorphism classes of elliptic curves, and thus E can be defined over a finite extension of K.

In fact E can be defined over the Hilbert class group H/K, and $H = K(j(E))$. Now, just as the torsion in $\mathbf{G}_m$ generated abelian extensions of $\mathbb{Q}$, the torsion in E generates abelian extensions of K. This is essentially because $E[n]$ is a free $\mathcal{O}_K/n$-module of rank 1. The analogue of the Kronecker-Weber theorem states that every abelian extension of K is contained in $H(E[n])$ for some n.

As with cyclotomic fields, there is also a reciprocity law. One version of this is that if $\mathfrak{p}$ is a nonzero prime ideal of $\mathcal{O}_K$, and Λ is a lattice with $\mathcal{O}_K\Lambda = \Lambda$, then $\operatorname{Frob}_p(j(E_\Lambda)) = j(E_{\mathfrak{p}^{-1}\Lambda})$. There is a consequence for L-functions, too: we have an equality

$$L(E/H, s) = L(\psi, s)L(\overline{\psi}, s)$$

where ψ is a certain Hecke character of H, and $L(E/H, s)$ is the Hasse-Weil L-series.

1.3 Elliptic curves in the large

When we drop the CM assumption, there are infinitely many isomorphism classes of elliptic curves. So we consider the *moduli space* of elliptic curves, as a variety (actually a stack, but don't worry about this now) $Y/\mathbb{Q}$. Y by itself isn't all that interesting: the j-invariant gives an isomorphism of it onto the affine line. We add *level structures* to make it interesting. For instance, let $Y(N)$ be the moduli space of pairs (E, α), where E is an elliptic curve and $\alpha\colon E[N] \cong \mathbb{Z}/N\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z}$ is an isomorphism. Y and $Y(N)$ are called *modular curves*. Note that we have a morphism $Y(M) \to Y(N)$ whenever N divides M. The curve $Y(N)$ is not complete; let $X(N)$ be the completion of $Y(N)$.

There is once again a reciprocity law, and also an equality between L-series, but it takes some time to set up, and the details are outside the scope of this lecture. What follows is the ultra-brief version.

Modular curves are closely linked with modular forms. For instance, a cusp form of weight 2 for $\Gamma(N)$ is the same thing as a holomorphic differential on $X(N)$. Spaces of modular forms admit actions by a commutative ring of Hecke operators, so we may talk of eigenforms. Let $f(z) = \sum_{n\geq 1} a_n q^n$ be a (new) cuspidal eigenform. Then the a_n generate a number field E. A theorem of Eichler-Shimura and Deligne associates to f a family of Galois representations $\rho_f\colon \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \overline{\mathbb{Q}}_\ell$ (one for each embedding $E \to \overline{\mathbb{Q}}_\ell$, where ℓ is a prime), satisfying the equality

$$L(f, s) = L(\rho_f, s),$$

where $L(f, s) = \sum_{n\geq 1} a_n n^{-s}$ and $L(\rho_f, s)$ is a product of Euler factors, almost all of which are of the form $\det(1 - \rho_f(\mathrm{Frob}_p)p^{-s})^{-1}$ (this last expression lies in E and is independent of the choice of $E \to \overline{\mathbb{Q}}_\ell$).

The theorem of Eichler-Shimura and Deligne above may be interpreted as a "non-abelian reciprocity law". Its proof involves a study of the étale cohomology $H^1(X(N)_{\overline{\mathbb{Q}}}, \overline{\mathbb{Q}}_\ell)$, which is a finite-dimensional $\overline{\mathbb{Q}}_\ell$-vector space admitting an action of $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$.

1.4 Elliptic modules?

Admittedly there are few data points to work with, but the above three examples point to a general phenomenon, which goes like this: Start with a

number field $K \subset \mathbb{C}$ and an integer $d \geq 1$. Consider discrete rank d $\mathcal{O}_K$-modules $\Lambda \subset \mathbb{C}$, and for each one, try to consider $\mathbb{C}/\Lambda$ as an algebraic group. Call such an algebraic group an *elliptic $\mathcal{O}_K$-module* of rank d. Thus $\mathbf{G}_m$ is an elliptic $\mathbb{Z}$-module of rank 1, an elliptic curve is an elliptic $\mathbb{Z}$-module of rank 2, and an elliptic curve with CM by $\mathcal{O}_K$ is an elliptic $\mathcal{O}_K$-module of rank 1.

It is really not clear how to give any other examples of elliptic modules than these, essentially for the reason that $\dim_{\mathbb{R}} \mathbb{C} = 2$. It is impossible to have a discrete rank d $\mathcal{O}_K$-submodule of $\mathbb{C}$ unless $d[K : \mathbb{Q}] \leq 2$; and if $K \neq \mathbb{Q}$ it must be imaginary quadratic. We are therefore in one of the cases (1)-(3) as above.

2 Drinfeld modules: definition and first examples

2.1 Global fields

The situation is quite different when K is replaced with a function field. Recall that a *function field* is a finite extension $K/\mathbb{F}_p(T)$, where p is prime. More intrinsically, $\operatorname{Spec} K$ is the generic point of a curve X (curve = nonsingular projective integral scheme of dimension 1) over a finite field.

Also recall the following basic definitions: a *global field* K is a number field or a function field. A *place* of K is a nontrivial metric, up to equivalence; these are either nonarchimedean (ultrametric) or else they are archimedean, in which case they are either real or complex. Let $|K|$ be the set of places of K. If K is the function field of a curve X, then $|K|$ consists only of nonarchimedean places; it may be identified with the set of closed points of X. If $v \in |K|$, there is a usual choice of norm $|f|_v$ on K. If v is nonarchimedean, it corresponds to a discrete valuation $\nu \colon K \to \mathbb{Z} \cup \{\infty\}$, and then $|f|_v = q^{-\nu(f)}$, where q is the cardinality of the residue field of v. An important special case is $K = \mathbb{F}_q(T)$, $v = \infty$, in which case $|f(x)| = q^{\deg f}$ for $f(x) \in \mathbb{F}_q[T]$. If v is archimedean, it corresponds to an embedding $\iota \colon K \to \mathbb{C}$ (up to complex conjugation), and then $|x|_v = |\iota(x)|^a$, where $a = 1$ or 2 depending on whether ι is real or complex.

Under these normalizations, we have the global product formula

$$\prod_{v \in |K|} |f|_v = 1,$$

valid for all $f \in K^\times$. It might be worthwhile to review the proof of this fact: for the fields $\mathbb{Q}$ and $\mathbb{F}_p(T)$, this can be worked out rather amusingly by hand, and then the result can be extended to general K once you know how norms interact with finite extensions.

Let $S \subset |K|$ be a nonempty finite set of places containing the archimedean places. We define

$$\mathcal{O}_{K,S} = \left\{ x \in K \;\middle|\; |x|_v \leq 1,\ v \notin S \right\}.$$

Because we included all archimedean places in S, $\mathcal{O}_{K,S}$ is closed under addition. In fact $\mathcal{O}_{K,S}$ is a Dedekind ring. Note that if K is a number field and S is the set of archimedean places, then $\mathcal{O}_K = \mathcal{O}_{K,S}$. Also note that if K is the function field of a curve X, then $\mathcal{O}_K, S = H^0(X \backslash S, \mathcal{O}_X)$ is the ring of functions which are regular outside S. The Dirichlet unit theorem describes the unit group $\mathcal{O}_{K,S}^\times$; up to torsion it is free of rank $\#S - 1$.

For a place $v \in |K|$, we have the completion K_v, a locally compact field. If v is archimedean, then K_v is either $\mathbb{R}$ or $\mathbb{C}$; if v is nonarchimedean then K_v it is a finite extension of either $\mathbb{Q}_p$ or $\mathbb{F}_p(\!(T)\!)$. Furthermore, if K is characteristic p, then $K_v \cong \mathbb{F}_q(\!(T)\!)$, where $\mathbb{F}_q$ is the residue field of v.

We would like to carry the notion of an elliptic module over to the case of a general global field. For an analogue of $\mathbb{C}$, we might choose a place $v \in |K|$, and let C be the completion of an algebraic closure $\overline{K}_v$ of K_v. Recall that if v is nonarchimedean, $\overline{K}_v$ is not complete. For instance, if $K_v = \mathbb{Q}_p$, the completion C of $\overline{\mathbb{Q}}_p$ (this is often called $\mathbb{C}_p$) has infinite transcendence degree over $\mathbb{Q}_p$. This is in stark contrast to the case $K_v = \mathbb{R}$, $C = \mathbb{C}$.

The next step is to consider a discrete finite-rank $\mathcal{O}_{X,S}$-submodule of C to serve as our lattice Λ. This can only exist if $\mathcal{O}_{X,S}$ is itself a discrete subring of C.

Exercise. Show that $\mathcal{O}_{X,S} \subset C$ is discrete if and only if $S = \{v\}$.

Thus if a theory of elliptic modules is to get off the ground, we need $S = \{v\}$. In the case that K is a number field, this forces $S = \{\infty\}$, which means that K has only one archimedean place; i.e. $K = \mathbb{Q}$ or else K is an imaginary quadratic field. In that case $C = \mathbb{C}$ and $\mathcal{O}_{K,S} = \mathcal{O}_K$, and then the rank of a discrete $\mathcal{O}_K$-submodule of C can only be 1 or 2 (if $K = \mathbb{Q}$) or 1 (if K is imaginary quadratic). I call this *the tyranny of the archimedean place*[3].

[3]For his part, Drinfeld defines an *admissible triple* as a pair (K, ∞, d), where $\infty \in |K|$ is a place such that all places except possibly ∞ are nonarchimedean, and $d \leq [\overline{K} : K]$; the number field examples are those enumerated above.

No such tyranny exists in the world of function fields: we are free to choose any place $v \in |K|$, and then $C = \widehat{K}_v$ contains discrete $\mathcal{O}_{K,\{v\}}$-submodules of arbitrary finite rank.

2.2 Drinfeld modules: Analytic point of view

Now we are in the function field setting only: let X be a curve over $\mathbb{F}_q$, and let K be its function field. We choose a closed point of X, corresponding to a place of K; we call this point ∞, and let $A = \mathcal{O}_{K,\{\infty\}} = H^0(X \setminus \{\infty\}, \mathcal{O}_X)$. We have the completion K_∞, and $C = \widehat{K}_\infty$.

Warm-up exercise. Show that the map $x \mapsto x^q - x = \prod_{a \in \mathbb{F}_q}(x - a)$ induces an isomorphism of topological groups $C/\mathbb{F}_q \cong C$.

We are interested in quotients C/Λ, where $\Lambda \subset C$ is a discrete finite-rank A-module. If A happens to be a PID, this means that $\Lambda = \bigoplus_{i=1}^{d} A\alpha_i$, where $\{\alpha_1, \ldots, \alpha_d\}$ are K_v-linearly independent elements of C.

Theorem 2.2.1 (Drinfeld)**.** *For $z \in C$, let*

$$\wp_\Lambda(z) = z \prod_{\lambda \in \Lambda \setminus \{0\}} \left(1 - \frac{z}{\lambda}\right).$$

1. *The product defining $\wp_\Lambda(z)$ converges in C.*
2. *$\wp_\Lambda : C \to C$ is a surjective continuous homomorphism of topological $\mathbb{F}_q$-vectors spaces, whose kernel is exactly Λ. Thus $C/\Lambda \cong C$ as topological $\mathbb{F}_q$-vector spaces.*

Some remarks before the proof: note that $\wp_\Lambda$ has been designed to have a simple root at each element of Λ. This is in analogy with the Weierstrass $\wp$-function, which is designed to have a simple pole at every point of a lattice.

Generally in complex analysis, one has a Weierstrass product for an entire function; the product will match the function up to a nonvanishing entire function. The situation is a bit simpler in the nonarchimedean world: no exponential factors are needed to make a Weierstrass product converge. What's more, we have the following analogue of Picard's theorem from complex analysis: There are no nonconstant nonvanishing entire functions. (This can be proven by analyzing Newton polygons [cite, probably BGR].)

Proof. (1) will follow from the discreteness of Λ: for every $N > 0$, there exist only finitely many $\lambda \in \Lambda$ with $|\lambda| < N$. Thus for every $\epsilon > 0$, there exist only finitely many λ with $|z/\lambda| < \epsilon$. In the nonarchimedean world, this is enough to ensure convergence of the product (exercise!).

For (2), let $\Lambda_0 \subset \Lambda$ be a finite $\mathbb{F}_q$-submodule, and let

$$\wp_{\Lambda_0} = z \prod_{\lambda \in \Lambda_0 \setminus \{0\}} \left(1 - \frac{z}{\lambda}\right) \in C[z],$$

a polynomial in z. We claim that $\wp_{\Lambda_0}(x+y) = \wp_{\Lambda_0}(x) + \wp_{\Lambda_0(y)}$ in $C[x,y]$. Let $f(x,y) = \wp_{\Lambda_0}(x+y) - \wp_{\Lambda_0}(x) - \wp_{\Lambda_0}(y)$. Evidently $f(\lambda, \lambda') = 0$ for all $\lambda, \lambda' \in \Lambda_0$. For all $\lambda \in \Lambda_0$, $f(x,\lambda) \in C[x]$ has at least $\#\Lambda_0$ roots, but $\deg f(x,\lambda) < \#\Lambda_0$, therefore $f(x,\lambda) = 0$ identically. Now consider $f(x,y) \in C(x)[y]$; as a polynomial in y this has at least $\#\Lambda_0$ roots (namely $y \in \Lambda_0$) but once again its degree in y is $< \#\Lambda_0$. Therefore $f(x,y) = 0$ indentically. A similar argument shows that $f(az) = af(z)$ for $a \in \mathbb{F}_q$. □

Since C/Λ is an A-module, the isomorphism $\wp_\Lambda \colon C/\Lambda \to C$ gives C an exotic A-module structure, one quite distinct from the structure arising from the inclusion $A \to C$. *The set C endowed with this exotic A-module structure is the Drinfeld module associated to Λ.* Now, the new $\mathbb{F}_q$-module structure on C is the same as the usual one, since $\wp_\Lambda$ is $\mathbb{F}_q$-linear. What is really new is the endomorphism $\phi_\alpha \colon C \to C$ carried over from the action of multiplication by α on C/Λ.

Theorem 2.2.2. *Let $\alpha \in A$. There exists a* polynomial $\phi_\alpha(x) \in C[x]$ *making the following diagram commute:*

$$\begin{array}{ccc} C/\Lambda & \xrightarrow{\wp_\Lambda} & C \\ {\scriptstyle \alpha}\downarrow & & \downarrow{\scriptstyle \phi_\alpha} \\ C/\Lambda & \xrightarrow[\wp_\Lambda]{} & C. \end{array}$$

This polynomial is $\mathbb{F}_q$-linear; that is, it takes the form

$$\phi_\alpha(x) = a_0 + a_1 x^q + a_2 x^{q^2} + \ldots, \ a_i \in C$$

Finally, $\phi_{\alpha\beta} = \phi_\alpha \circ \phi_\beta$ (composition of polynomials in $C[x]$ for all $\alpha, \beta \in A$.

Proof. We may assume $\alpha \neq 0$. A priori we have a $\mathbb{F}_q$-linear entire map $\phi_\alpha \colon C \to C$ which makes the diagram commute, but we need to know that it is a polynomial. Its kernel is the image under $\wp_\Lambda$ of $\ker(\alpha | C/\Lambda) = \alpha^{-1}\Lambda/\Lambda$, which is finite (because Λ is a finite-rank A-module, and because A/a is finite). We claim that

$$\phi_\alpha(z) = az \prod_{\lambda \in \alpha^{-1}\Lambda/\Lambda \setminus \{0\}} \left(1 - \frac{z}{\wp_\Lambda(\lambda)}\right).$$

Indeed, both sides are entire functons with the same set of simple zeros; thus by the nonarchimedean Picard theorem the quotient is constant. Examining coefficients of z shows that this constant is 1.

The rest of the claims about ϕ_α follow formally from the corresponding properties of multiplication by α on C/Λ. □

In light of Theorem 2.2.2, the association

$$\begin{array}{ccc} A & \to & C[z] \\ \alpha & \mapsto & \phi_\alpha \end{array}$$

has the following properties:

1. $\phi'_\alpha(0) = \alpha$.
2. ϕ_α is an $\mathbb{F}_q$-linear polynomial; that is,
$$\phi_\alpha(z) = \alpha z + a_1 z^q + a_2 z^{q^2} + \dots.$$
3. $\phi_{\alpha+\beta} = \phi_\alpha + \phi_\beta$.
4. $\phi_{\alpha\beta} = \phi_\alpha \circ \phi_\beta$.

Properties (2)-(4) mean that $\alpha \mapsto \phi_\alpha$ defines an $\mathbb{F}_q$-algebra homomorphism $A \to \mathrm{End}_{\mathbb{F}_q} \mathbb{G}_{a,C}$, where $\mathbb{G}_{a,C} = \operatorname{Spec} C[z]$ is the additive $\mathbb{F}_q$-vector space scheme over C, and $\mathrm{End}_{\mathbb{F}_q}$ means endomorphisms in the category of $\mathbb{F}_q$-vector space schemes over C. Property (1) means that the derivative of this action agrees with the inclusion $A \to C$.

Yet another interpretation: Let $C[z]^{\mathbb{F}_q\text{-lin}}$ be the set of polynomials in $C[z]$ which are $\mathbb{F}_q$-linear. This set becomes a noncommutative C-algebra, if we define the C-vector space structure as usual, but the multiplication operation is

interpreted as *composition*. (Check that these satisfy the ring axioms.) With this $\mathbb{F}_q$-algebra structure, $C[z]^{\mathbb{F}_q\text{-lin}}$ is isomorphic to the twisted polynomial ring $C\{\tau\}$, whose underlying C-vector space is $C[\tau]$, but which is subject to the rule $\tau a = a^q \tau$, for $a \in C$. Then ϕ may be recast as a ring homomorphism $A \to C\{\tau\}$.

3 The Carlitz module

It is enlightening to work out some details of the following special case, which was worked out by Carlitz in the 1930s (long before Drinfeld). Let

$$\begin{aligned} K &= \mathbb{F}_q(T) \\ A &= \mathbb{F}_q[T] \\ K_\infty &= \mathbb{F}_q(\!(1/T)\!) \\ \Lambda &= A. \end{aligned}$$

Thus ∞ is the point at infinity in the projective line over $\mathbb{F}_q$. Recall that if $f \in A$ then $|f|_\infty = q^{\deg f}$.

We would like to investigate the A-module structure on C induced by the isomorphism $\wp_A \colon C/A \to C$, where

$$\wp_A(z) = z \prod_{a \in A \setminus \{0\}} \left(1 - \frac{z}{a}\right).$$

For $\alpha \in \mathbb{F}_q[x]$, let $\phi_{A,\alpha}$ for the polynomial appearing in Theorem 2.2.2. To know what $\phi_{A,\alpha}$ is for general α, it is enough to know $\phi_{A,T}$. Now $\ker \phi_{A,T}$ is isomorphic to the kernel of multiplication by T on C/A, which is in turn isomorphic to $A/T = \mathbb{F}_q$. Thus $\deg \phi_{A,T} = q$, which means that

$$\phi_{A,T}(z) = Tz + \beta z^q$$

for some nonzero $\beta \in C$.

What we will now do is *rescale the lattice* $\Lambda = A \subset C$ to get rid of β above. For a nonzero element $\xi \in C$, consider the lattice $\xi A \subset C$; by inspection we have $\phi_{\xi A}(\xi z) = \xi \phi_A(z)$ and $\phi_{\xi A,\alpha}(\xi z) = \xi \phi_{A,\alpha}(z)$, so that $\phi_{\xi A,T}(z) = Tz + \xi^{1-q}\beta$. Since C is algebraically closed, we may choose ξ so that $\xi^{q-1} = \beta$, and then

$$\phi_{\xi A,T}(z) = Tz + z^q.$$

The *Carlitz module* $\mathcal{C}$ is the A-module whose underlying $\mathbb{F}_q$-vector space is C, for which the action of $\alpha \in A$ is through $\phi_{\xi A,\alpha}$. Note that $\phi_{\xi A,\alpha}(z) \in A[T]$ for all $\alpha \in A$.

The power series $e_{\mathcal{C}}(z) = \wp_{\xi A}(z)$ is called the *Carlitz exponential*. It is an $\mathbb{F}_q$-linear power series with the properties $e'_{\mathcal{C}}(z) = 1$ and $e_{\mathcal{C}}(T) = Te_{\mathcal{C}}(T) + e_{\mathcal{C}}(T)^q$. These properties determine $e_{\mathcal{C}}(z)$ completely:

$$e_{\mathcal{C}}(z) = \sum_{n=0}^{\infty} \frac{x^{q^n}}{D_n}, \tag{1}$$

where the D_n are determined by the recursion $D_0 = 1$, $D_n = (T^{q^n} - T)D_{n-1}^q$. It is a fun exercise to verify that D_n is the product of all monic polynomials in $\mathbb{F}_q[T]$ of degree n. Therefore

$$z \prod_{\alpha \in \mathbb{F}_q[T] \setminus \{0\}} \left(1 - \frac{z}{\xi\alpha}\right) = \sum_{n=0}^{\infty} \frac{z^{q^n}}{D_n} \in K[\![z]\!] \tag{2}$$

This picture lines up beautifully with the archimedean story. The quotient $\mathbb{C}/\mathbb{Z}$ may be identified with $\mathbb{C}^\times$ by means of a power series $\exp(2\pi z) \in \mathbb{C}[\![z]\!]$. After renormalizing $\mathbb{Z}$ to $2\pi i\mathbb{Z}$, the power series has rational coefficients $\exp(z) \in \mathbb{Q}[\![z]\!]$. After examining the Weierstrass product form for the entire function $\exp(z) - 1$, we arrive at the familiar product formula

$$\sin(z) = z \prod_{n \geq 1} \left(1 - \frac{z^2}{\pi^2 n^2}\right),$$

from which follows Euler's calculation that $\zeta(2k) \in \pi^{2k}\mathbb{Q}$ for $k = 1, 2, \ldots$.

Evidently the element $\xi \in C$ is a characteristic p analogue for $2\pi i$. With a little more analysis, one can even show that $\xi \in K_\infty \cdot \overline{K}$, just as $2\pi i \in \mathbb{R} \cdot \overline{\mathbb{Q}}$.

4 Drinfeld modules: algebraic approach

As usual, K is a function field, $\infty \in |K|$, $A = \mathcal{O}_{K,\{\infty\}}$. For $\alpha \in A$, write $\deg \alpha = -v_\infty(\alpha)$, where v_∞ is the ($\mathbb{Z}$-valued) valuation on K corresponding to ∞. By the product formula, $\deg \alpha \geq 0$ for all nonzero $\alpha \in A$.

We are ready to define Drinfeld modules over a field (the generalization to arbitrary scheme bases will come later). Let L be an *A-field*; that is, a

field L equipped with a ring homomorphism $\iota\colon A \to L$. We do not require that ι be injective; the prime ideal $\mathfrak{p} = \ker \iota$ is called the *characteristic* of the A-field L. Define the twisted polynomial ring $L\{\tau\} = \mathrm{End}_{\mathbb{F}_q} \mathbb{G}_{a,L}$.

Definition 4.0.3. A *Drinfeld A-module* ϕ over L is a homomorphism of $\mathbb{F}_q$-algebras $A \to L\{\tau\}$, $\alpha \mapsto \phi_\alpha$, such that for all $\alpha \in A$, the constant term of ϕ_α is $\iota(\alpha)$. It is required that ϕ be something other than the map $\alpha \mapsto \iota(\alpha)$.

A morphism $\phi \to \phi'$ between Drinfeld modules is an element $f \in L\{\tau\}$ such that for all $\alpha \in A$, $f\phi_\alpha = \phi'_\alpha f$.

We could have also defined a Drinfeld A-module over L this way: it is an A-module scheme over L, whose underlying $\mathbb{F}_q$-vector space scheme is $\mathbb{G}_{a,L}$, such that the derivative of $A \to \mathrm{End}_{\mathbb{F}_q} \mathbb{G}_{a,L}$ at the origin is $\iota\colon A \to L$.

In important observation is that Drinfeld modules form an A-linear category: that is, all Hom sets $\mathrm{Hom}(\phi, \phi')$ are A-modules (and composition is A-bilinear). Indeed, for $f\colon \phi \to \phi'$ and $\beta \in A$, one defines $\beta \cdot f = \phi'_\beta f$; this also lies in $\mathrm{Hom}(\phi, \phi')$ because for all $\alpha \in A$,

$$(\beta \cdot f)\phi_\alpha = \phi'_\beta f \phi_\alpha = \phi'_\beta \phi'_\alpha f = \phi'_\alpha \phi'_\beta f = \phi'_\alpha (\beta \cdot f).$$

This is in line with the philosophy of elliptic A-modules for general A: these shouls always constitute an A-linear category.

4.1 Rank and height

Our first order of business is to define the rank of a Drinfeld module ϕ. For every $\alpha \in A$, we have the degree $\deg \phi_\alpha$, where ϕ_α is considered as a polynomial in τ. This function $A \mapsto \mathbb{Z}_{\geq 0} \cup \{-\infty\}$ satisfies the properties

1. $\deg \phi_{\alpha\beta} = \deg \phi_\alpha + \deg \phi_\beta$,
2. $\deg \phi_{\alpha+\beta} \leq \max\{\deg \phi_\alpha, \deg \phi_\beta\}$,
3. There exists α with $\deg \phi_\alpha > 0$.

This means that $\alpha \mapsto \exp \deg \phi_\alpha$ extends to a place of K, which is ≤ 1 on A; this place must be equivalent to ∞. Therefore there exists $d \in \mathbb{Q}_{>0}$ such that for all $\alpha \in A$,

$$\deg \phi_\alpha = d \deg \alpha.$$

In fact d is an integer, known as the *rank* of ϕ.

We can play a similar game with the number of times τ divides ϕ_α (rather than its degree). For each α, let $m(\alpha) = m$ be the least integer for which τ^m appears in ϕ_α or ∞ if $\phi_\alpha = 0$. Now, if the characteristic of L is 0, then ϕ_α has constant term $\iota(\alpha)$, so that $m(\alpha) = 0$ for all nonzero α. Now suppose $\mathfrak{p} = \ker \iota$ is nonzero. Then $\alpha \mapsto \exp(-m(\alpha))$ extends to a nontrivial real-valued function $K \to \mathbb{R}_{\geq 0}$ which is multiplicative and ultrametric; *i.e.* it is a place of K. Since $m(\alpha) > 0$ for $\alpha \in \mathfrak{p}$, this place must be equivalent to the one associated to $\mathfrak{p}$. Thus there exists $h \in \mathbb{Q}_{\geq 0}$ such that $m(\alpha) = hv_{\mathfrak{p}}(\alpha)$ for all $\alpha \in A$. In fact h is an integer, known as the *height* of ϕ. If L has A-characteristic 0 then the height of ϕ is defined to be 0.

The significance of the rank and height becomes apparent when we consider the A-torsion in the Drinfeld module ϕ. Let $\overline{L}$ be an algebraic closure of L. For $\alpha \in A$, let

$$\phi[\alpha] = \left\{ x \in \overline{L} \;\middle|\; \phi_\alpha(x) = 0 \right\},$$

so that $\phi[\alpha]$ is an A/α-module. Now, note that if $\alpha \notin \mathfrak{p}$ then $\phi_\alpha(z) \in \mathbb{F}_q[z]$ is a separable polynomial (its derivative is $\iota(\alpha) \neq 0$), and so $\#\phi[\alpha] = q^{\deg \phi_\alpha} = q^{d \deg \alpha}$. Otherwise, $\#\phi[\alpha]$ will be strictly less than $q^{\deg \phi_\alpha}$.

For an ideal $I \subset A$ we may define $\phi[I]$ as the intersection of $\phi[\alpha]$ for all $\alpha \in I$.

Theorem 4.1.1. *Suppose I is relatively prime to $\mathfrak{p}$. Then $\phi[I]$ is a free A/I-module of rank d. Furthermore, for all $e \geq 1$, $\phi[\mathfrak{p}^e]$ is a free $A/\mathfrak{p}^e$-module of rank $d - h$. (In particular d and h are integers.)*

Proof. Let $P \subset A$ be a nonzero prime ideal, and consider $\phi[P^\infty] = \cup_{n \geq 1} \phi[P^n]$. Also let A_P be the localization of A at P, a DVR. Since $\phi[P^n]$ is a module over $A/P^n = A_P/P^n$, $\phi[P^\infty]$ is a module over A_P. In fact it is a *divisible* A_P-module, essentially because $z \mapsto \phi_\alpha(z)$ is surjective on C for each nonzero $\alpha \in A$. By the structure theorem for divisible modules over a DVR[4], A_P is isomorphic to a direct sum of r copies of K_P/A_P, where r is some cardinal. Now suppose $P^e = (\alpha)$ is principal. Then $\phi_\alpha(z) \in C[z]^{\mathbb{F}_q\text{-lin}}$ is a polynomial of separable degree $q^{d \deg \alpha}$ if $P \neq \mathfrak{p}$ (resp., $q^{(d-h) \deg \alpha}$ if $P = \mathfrak{p}$), so that $\#\phi[\alpha]$ is $q^{d \deg \alpha}$ (resp., $q^{(d-h) \deg \alpha}$). This implies that $r = d$ (resp., $r = d - h$).

[4]If you are reading this – do you know a reference for this structure theorem? Sketch: every element of such a module is contained in a copy of K_P/A_P, which (being divisible, hence injective) must be a direct summand. Apply Zorn's lemma.

This proves the theorem in the case that I is a prime power; the general case reduces to this case by a suitable Chinese remainder theorem. $\square$

5 Moduli of Drinfeld modules over C

5.1 Drinfeld modules and lattices

As usual, K is a function field with residue field $\mathbb{F}_q$, $\infty \in |K|$ is a place, and $A = \mathcal{O}_{K,\{\infty\}}$. In time we will define a moduli stack M^d of Drinfeld A-modules of rank d. We cannot do this quite yet (we need to define Drinfeld modules over general scheme bases first), but we can make sense of $M^d(C)$, where $C = \hat{\overline{K}}_\infty$, as a rigid-analytic space over C. Happily, this situation bears a strong analogy to the classical situation of elliptic curves over $\mathbb{C}$, where one can identify the moduli space of complex elliptic curves with the quotient $\mathcal{H}/\operatorname{SL}_2(\mathbb{Z})$, where $\mathcal{H}$ is the complex upper half-plane.

Theorem 5.1.1. *Let $d \geq 1$. The functor $\Lambda \mapsto \psi_\Lambda$ is an equivalence between the following A-linear categories:*

1. *A-lattices $\Lambda \subset C$ of rank d (meaning discrete A-submodules which are locally free of rank d), where* $\operatorname{Hom}(\Lambda, \Lambda') = \left\{ \alpha \in C \;\middle|\; \alpha\Lambda \subset \Lambda' \right\}$.
2. *Drinfeld A-modules over C.*

Proof. Let us at least sketch the proof that this functor is essentially surjective. Suppose $\psi\colon A \to C\{\tau\}$ is a Drinfeld A-module of rank d. The idea is to find an entire $\mathbb{F}_q$-linear function $\wp\colon C \to C$ which is periodic in Λ and which satisfies

$$\wp(\alpha z) = \psi_\alpha(\wp(z)), \text{ all } \alpha \in A \text{ and } \wp'(0) = 0. \tag{3}$$

We may interpret (3) as a functional equation to be solved in the ring $C[\![\tau]\!]$ of noncommutative formal power series (satisfying the usual relation $\tau a = a^q \tau$. To wit, we are looking for an element $\wp$ in $C[\![\tau]\!]$ for which satisfies

$$\wp^{-1}\psi_\alpha\wp = \alpha \tag{4}$$

for all $\alpha \in A$, and also the constant term of $\wp$‘ is 1.

Let $\alpha \in A$. By completely formal means, one may find a $\wp_\alpha \in C[\![\tau]\!]$ with constant term 1 satisfying (4); its coefficients can be defined recursively. If $\wp'_\alpha$ is another such element, then $\wp'_\alpha \wp_\alpha^{-1}$ commutes with α. If we assume that α is transcendental over $\mathbb{F}_q$, then it is easy to see that the centralizer of α is C itself, so that in fact $\wp_\alpha = \wp'_\alpha$.

Let $\beta \in A$ be another transcendental element. Since ϕ_α commutes with ϕ_β, $\wp_\beta^{-1}\phi_\alpha\wp_\beta$ commutes with $\wp_\beta^{-1}\phi_\beta\wp_\beta = \beta$, and thus must be a constant; examining constant terms gives $\wp_\beta^{-1}\phi_\alpha\beta = \alpha$, and so in fact $\wp_\alpha = \wp_\beta$.

Letting $\wp$ be the common value of the $\wp_\alpha$ for the transcendental $\alpha \in A$, we have an $\mathbb{F}_q$-linear power series which at least formally satisfies (3). Now one has to do a calculation to verify that the recursion giving the coefficients of $\wp$ decay rapidly enough to ensure that $\wp$ defines an entire function $C \to C$.

Let $\Lambda = \wp^{-1}(0)$. Then (3) ensures that Λ is an A-module. By a general fact about zero sets of analytic functions (analogous with the complex version), Λ is discrete. Also one finds that for all $\alpha \in A$ nonzero, $\wp$ descends to an isomorphism $\alpha^{-1}\Lambda/\Lambda \cong \phi[\alpha]$. Since ϕ has rank d, $\alpha^{-1}\Lambda/\Lambda$ is a free A/α-module of rank d. This forces the rank of Λ to be d. Finally, since $\wp\colon C \to C$ is an entire $\mathbb{F}_q$-linear function vanishing simply on Λ, the uniqueness of the Weierstrass product shows that $\wp = \wp_\Lambda$, and so $\phi = \phi_\Lambda$. □

Let us define $M^d(C)$ to be the set of isomorphism classes of Drinfeld A-modules of rank d, leaving M^d undefined for the moment. By the above theorem, $M^d(C)$ is in bijection with the set of homothety classes of A-lattices $\Lambda \subset C$ of rank d. From there it is possible to give $M^d(C)$ the structure of a rigid analytic space over C. Before doing this, let us identify some invariants of lattices and examine some special cases.

5.2 Eisenstein series

Let $\Lambda \subset C$ be an A-lattice of rank d. For $k \geq 1$ we define the *Eisenstein series*

$$E_k(\Lambda) = \sum_{\lambda \in \Lambda}{}' \lambda^{-k},$$

where $\sum'$ means that 0 has been omitted. Then $E_k(\Lambda) = 0$ unless $k \equiv 0 \pmod{q-1}$, owing to the scalars $\mathbb{F}_q^\times \subset A^\times$. Furthermore we have the properties $E_k(\lambda\Lambda) = \lambda^{-k}E_k(\Lambda)$ for $\lambda \in C^\times$ and $E_{qk}(\Lambda) = E_k(\Lambda)^q$.

As with the classical Weierstrass $\wp$-function, the Eisenstein series appear as Taylor coefficients of the entire function

$$\wp_\Lambda(z) = z \prod_{\lambda \in \Lambda}{}' \left(1 - \frac{z}{\lambda}\right).$$

Lemma 5.2.1. *We have*

$$\frac{z}{\wp_\Lambda(z)} = 1 - \sum_{k \geq 1} E_k(\Lambda) z^k.$$

Proof. Formally taking the logarithmic derivative of $\wp_\Lambda(z)$ and noting that $\wp'_\Lambda(z) = 1$ gives

$$\begin{aligned} z/\wp_\Lambda(z)^{-1} &= z \sum_{\lambda \in \Lambda} \frac{1}{z - \lambda} \\ &= \sum_{\lambda \in \Lambda} \frac{1}{1 - \lambda/z} \\ &= 1 - \sum_{\lambda \in \Lambda \setminus \{0\}} \frac{z}{\lambda} \frac{1}{1 - z/\lambda} \\ &= 1 - \sum_{k \geq 1} E_k(\Lambda) z^k. \end{aligned}$$

□

5.3 The case $d = 1$

Recall that A is a Dedekind ring, and as such the following groups are naturally isomorphic:

1. Isomorphism classes of locally free A-modules of rank 1, under tensor product.

2. The class group of A, defined as fractional ideals modulo principal ideals.

3. The quotient $K^\times \backslash \mathbf{A}_K^{\infty,\times} / \prod_v \mathcal{O}_v^\times$, where $\mathbf{A}^\infty$ is the adele ring (away from ∞) and v runs over places of K other than ∞.

Theorem 5.3.1. *There is a bijection $M^1(C) \to \operatorname{Pic} A$ which sends $\Lambda \subset C$ to the isomorphism class of Λ as a locally free A-module of rank 1.*

Proof. For injectivity: Suppose we have A-lattices $\Lambda, \Lambda' \subset C$ and an isomorphism of (abstract) A-modules $f\colon \Lambda \to \Lambda'$. Let $i\colon K \to \Lambda \otimes_A K$ be an isomorphism of K-vector spaces. Then we have two embeddings of K-vector spaces $K \hookrightarrow C$, namely i and $f \circ i$. These must differ by a nonzero constant α, which then satisfies $\alpha\Lambda = \Lambda'$.

For surjectivity: every abstract Λ is isomorphic to a fractional ideal of K, which is naturally a lattice in C. □

As a variation on this idea, there is a functor $\Lambda \mapsto \bigwedge^d \Lambda$ from locally free A-modules of rank d to locally free A-modules of rank 1, which induces a map $M^d(C) \to M^1(C) \cong \operatorname{Pic} A$. This shows that we cannot expect $M^d(C)$ to be connected in general.

5.4 The Carlitz module revisited: $d = 1$, $A = \mathbb{F}_q[T]$

Recall the transcendental element $\xi \in C$, which is the A-analogue of $2\pi i$. Recall that $e_C(z) = \wp_{\xi A}(z)$ is the Carlitz exponential. Lemma 5.2.1 shows that

$$z/e_{\mathcal{C}}(z) = 1 + E_{q-1}(\xi A) z^{q-1} + \dots,$$

so that

$$e_{\mathcal{C}}(z) = z - E_{q-1}(\xi A) z^q + \dots.$$

Comparing with (1) shows that

$$E_{q-1}(A) = \sum_{\alpha \in \mathbb{F}_q[T] \setminus \{0\}} \frac{1}{\alpha^{q-1}} = -\frac{\xi^{q-1}}{T^q - T},$$

which is something of an analogue of $\zeta(2) = \pi^2/6$.

5.5 The case $d = 2$, $A = \mathbb{F}_q[T]$

In the case $d = 2$, $A = \mathbb{F}_q[T]$ there are many beautiful parallels to the theory of classical (elliptic) modular forms and modular functions. Note that $\operatorname{Pic} A = 0$, so that every A-lattice is a free A-module.

Using (5.2.1), we find that the first few terms of $\wp_\Lambda$ are

$$\wp_\Lambda(z) = z + E_{q-1}(\Lambda) z^q + [E_{q^2-1}(\Lambda) + E_{q-1}(\Lambda)^{q+1}] z^{q^2} + \dots. \tag{5}$$

We have the Drinfeld module $\phi = \phi_\Lambda$, which is characterized by $\wp(Tz) = \phi_T(\wp(z))$. Let us write

$$\phi_T = T + g(\Lambda)\tau + \Delta(\Lambda)\tau^2,$$

for elements $g(\Lambda), \Delta(\Lambda) \in C$. Note that $\Delta(\Lambda) \neq 0$, since ϕ must be rank 2. Comparing (5) with the functional equation $\wp_\Lambda(Tz) = \phi_T(\wp_\Lambda(z))$ we find the following relations:

$$\begin{aligned} g(\Lambda) &= (T^q - T)E_{q-1}(\Lambda) \\ \Delta(\Lambda) &= (T^q - T)^q E_{q-1}(\Lambda)^{q+1} + (T^{q^2} - T)E_{q^2-1}(\Lambda) \end{aligned}$$

Note that $g(\Lambda)$ and $\Delta(\Lambda)$ have weights $q-1$ and q^2-1, in the sense that $g(\lambda\Lambda) = \lambda^{-(q-1)}g(\Lambda)$, etc. Therefore

$$j(\Lambda) := g(\Lambda)^{q+1}/\Delta(\Lambda)$$

is a homothety invariant.

These functions of lattices (E_k, g, Δ, j) can be turned into functions of a single variable z by setting $E_k(z) = E_k(A + Az)$ (and similarly for g, Δ, j). Since $A + Az$ is a lattice if and only if $z \notin K_\infty$, these functions have as their domain the *Drinfeld half-plane*

$$\Omega = C \backslash K_\infty,$$

considered as a rigid-analytic variety over C. One can show that $E_k(z)$ is analytic on Ω, and (since $\Delta(z)$ is nowhere vanishing) so is $j(z)$.

Two elements in Ω determine homothetic A-lattices if and only if they are in the same $\mathrm{GL}_2(A)$-orbit, so that the set of homothety classes of A-lattices of rank 2 is in bijection with $\Omega/\mathrm{GL}_2(A)$. The function j, being a homothety invariant, descends to an analytic function on this quotient.

Theorem 5.5.1. *The function $j\colon \Omega/\mathrm{GL}_2(A) \to C$ is a bijection.*

Proof. To show that an element $a \in C$ lies in the image of j, one simply has to produce elements $g, \Delta \in C$ with $\Delta \neq 0$ and $g^{q+1}/\Delta = a$; then the lattice Λ whose Drinfeld module is $\phi_T = T + g\tau + \Delta\tau^2$ has j-invariant a.

For injectivity, we observe that two Drinfeld A-modules ϕ and ϕ' of rank 2 are isomorphic if and only if there exists $u \in C^\times$ such that $\phi' = u^{-1}\phi u$. If $\phi_T = T + g\tau + \Delta\tau^2$, and similarly for ϕ', this condition means that $g' = u^{q-1}g$, $\Delta' = u^{q^2-1}\Delta$. It is a simple matter to check that such a u exists if ϕ_T and ϕ'_T have the same j-invariant. □

Definition 5.5.2. A *Drinfeld modular form* of weight $k \geq 0$ for $\mathrm{GL}_2(A)$ is an analytic function $f\colon \Omega \to C$ which satisfies

$$f\left(\frac{az+b}{cz+d}\right) = (cz+d)^k f(z)$$

for all $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{GL}_2(A)$, and which is "analytic at ∞".

Recall that for a classical modular form f for the group $\mathrm{SL}_2(\mathbb{Z})$, the condition of being analytic at ∞ is the condition that f be holomorphic as a function of the parameter $q = e^{2\pi i z}$ at the point $q = 0$. The analogue of q in the Drinfeld case is $Q = e_C(\xi z)^{-1}$, where e_C is the Carlitz exponential; note that Q is invariant under translation by A. A Drinfeld modular form admits a Q-expansion[5] $\sum_{n\geq 0} a_n Q^n$, with $a_n \in C$, which is convergent on an open subset of Ω.

Let M_k be the space of Drinfeld modular forms of weight k. Note that $M_k = 0$ unless $k \equiv 0 \pmod{q-1}$, owing to the presence of matrices $\begin{pmatrix} a & \\ & 1 \end{pmatrix} \in \mathrm{GL}_2(A)$ for $a \in \mathbb{F}_q$. Let $M = \bigoplus_{k\geq 0} M_k$, a graded C-algebra.

Theorem 5.5.3 (Gekeler)**.** *$M = C[g, \Delta]$. Furthermore, up to a scalar, Δ admits an integral Q-expansion:*

$$\Delta(z) = -\xi^{q^2-1} Q^{q-1} \prod_{\alpha \in A monic} f_\alpha(Q),$$

where $f_\alpha(X) = X^{q^{\deg \alpha}} \phi_\alpha(X^{-1}) \in A[X]$, and ϕ is the Carlitz module. Therefore M admits an A-model M_0 consiting of Drinfeld modular forms whose Q-expansions lie in A, namely $M_0 = A[\tilde{g}, \tilde{\Delta}]$, with $\tilde{g} = \xi^{1-q} g$ and $\tilde{\Delta} = \xi^{1-q^2}\Delta$.

The result about Δ is analogous to Ramanujan's expansion of the classical Δ function,

$$\Delta = (2\pi i)^{12} \prod_{n\geq 1} (1-q^n)^{24}.$$

Practically any technique, result or conjecture about classical modular forms can be carried over into the setting of Drinfeld modular forms. There are Drinfeld modular forms of higher level. Spaces of Drinfeld modular forms

[5]The letter q being reserved for the size of the residue field of K.

have Hecke operators acting upon them, and there is a notion of a Hecke eigenform; these latter have Galois representations attached. Note though that the coefficients of these Galois representations lie in a completion of A, rather than $\mathbb{Z}$.

5.6 Adelic description of Drinfeld modular varieties over C

We return to the general case of a function field K, a place ∞ and $d \geq 1$ arbitrary. We will give a description of the set of homothety classes of A-lattices $\Lambda \subset C$ of rank d, which we have called $M^d(C)$. This will make it clear that $M^d(C)$ has the structure of a $(d-1)$-dimensional rigid-analytic space over C.

A *based A-lattice* of rank d is an A-lattice $\Lambda \subset C$ equipped with a specified A-basis: $\Lambda = \sum_{1 \leq i \leq d} Ax_i$, $x_i \in C$. A d-tuple of elements $x_1, \ldots, x_d \in C$ spans an A-lattice of rank d if and only if it is linearly independent over K_∞. Let Ω^d (*Drinfeld's half-space*) be the set of points $[x_1 : \cdots : x_d] \in \mathbf{P}^d(C)$ for which the x_i are linearly independent over K_∞. That is,

$$\Omega^d = \mathbf{P}^{d-1}(C) \backslash \bigcup_H H$$

where H runs over K_∞-rational hyperplanes in $\mathbf{P}^{d-1}$. Note that $\Omega^2 = \Omega = \mathbf{P}^1(C) \backslash \mathbf{P}^1(K_\infty)$ is Drinfeld's half-plane.

One shows that $\Omega^d \subset \mathbf{P}^{d-1}(C)$ is an admissible open subset, and thus that Ω^d is a rigid-analytic space. It admits an action by the group $\mathrm{GL}_d(K_\infty)$.

Homothety classes of based A-lattices of rank d are classified by Ω^d, so that the quotient $\Omega^d / \mathrm{GL}_d(A)$ classifies homothety classes of A-lattices of rank d which admit an A-basis; ie, those that are free as A-modules. Thus if $\operatorname{Pic} A = 0$ then $M^d(C) = \Omega^d / \mathrm{GL}_d(A)$.

Of course $\operatorname{Pic} A$ is nontrivial in general. To proceed, we need the adelic group $\mathrm{GL}_d(\mathbf{A}_K^\infty)$, a locally compact toplogical group. A basis of neighborhoods of the origin is given by the *congruence subgroups*

$$U_N = \left\{ (g_v)_{v \neq \infty} \in \mathrm{GL}_d(A_v) \;\middle|\; g_v \equiv 1 \pmod{N} \right\},$$

where N ranges over nonzero ideals of A.

Proposition 5.6.1. *We have a bijection*

$$M^d(C) \cong \mathrm{GL}_d(A)\backslash(\mathrm{GL}_d(\mathbf{A}_K^\infty) \times \Omega^d)/U_1,$$

where in the double coset space above, $\mathrm{GL}_d(A)$ *acts on* $\mathrm{GL}_d(\mathbf{A}_K^\infty)$ *(via the inclusion* $K \subset \mathbf{A}_K^\infty$*) and* Ω^d *(via the inclusion* $K \subset K_\infty$*, and* U_1 *acts on* $\mathrm{GL}_d(\mathbf{A}_K^\infty)$ *acts via right multiplication and on* Ω^d *trivially.*

Proof. Given an element $(g_v) \in \mathrm{GL}_d(\mathbf{A}_K^\infty)$, and an element $[x_1 : \cdots : x_d] \in \Omega^d$, we define a K-subspace $\Lambda_K = \sum_i Kx_i$ (this is well-defined up to $C^\times$). For all places $v \neq \infty$, the K_v-vector space $\Lambda_K \times_K K_v \cong K_v^n$ contains a distinguished A_v-lattice Λ_{A_v}, namely the span of the x_i. Finally, we let

$$\Lambda = \Lambda_K \cap \bigcap_{v \neq \infty} g_v \Lambda_{A_v},$$

an A-lattice in C of rank d which is well-defined up to homothety. One can now check that the lattice Λ only depends on the class of $((g_v), x)$ in the required double coset space.

For the reverse direction, we are given an A-lattice $\Lambda \subset C$ of rank d. Let $x_1, \ldots, x_d \in C$ be a basis for $\Lambda \otimes_K K^n$, and let $g_v \in \mathrm{GL}_d(K_v)$ be a matrix carrying $\sum_i A_v x_i$ onto $\Lambda \otimes_A A_v$. Then $((g_v), [x_1 : \cdots : x_d])$ corresponds to Λ. □

Recall the map $M^d(C) \to M^1(C) = \mathrm{Pic}\, A$, which sends Λ to the invertible A-module $\wedge^d \Lambda$. In light of the above description of $M^d(C)$ this map fits in a commutative diagram

$$\begin{array}{ccc} M^d(C) & \longrightarrow & \mathrm{GL}_d(K)\backslash(\mathrm{GL}_d(\mathbf{A}_K^\infty) \times \Omega^d)/\prod_{v\neq\infty} \mathrm{GL}_d(A_v) \\ \downarrow & & \downarrow \det \\ M^1(C) & \longrightarrow & K^\times\backslash\mathbf{A}_K^{\infty,*}/\prod_{v\neq\infty} A_v^*, \end{array}$$

where the horizontal maps are bijections. On the right-hand side, one checks that the fibers of the map labeled "det" are of the form $\Omega^d/\backslash\Gamma$, where $\Gamma \subset \mathrm{GL}_d(K_\infty)$ is a discrete subgroup.

For a nonzero ideal $N \subset A$, let $M_N^d(C)$ denote the set of homothety classes of A-lattices $\Lambda \subset C$ of rank d equipped with an isomorphism of A/N-modules $(A/N)^{\oplus d} \to N^{-1}\Lambda/\Lambda$. Then

$$M_N^d(C) = \mathrm{GL}_d(K)\backslash(\mathrm{GL}_d(\mathbf{A}_K^\infty) \times \Omega^d)/U_N.$$

For large enough N, the group $\mathrm{GL}_d(K) \times U_N$ acts strictly discontinuously on $\mathrm{GL}_d(\mathbf{A}_K^\infty) \times \Omega^d$, so that $M_N^d(C)$ really has the structure of a rigid-analytic space over C.

Those who have studied Shimura varieties in some level of generality will recognize the analogy. For a reductive group $G/\mathbb{Q}$, one defines a tower of Shimura varieties by

$$\mathrm{Sh}_U = G(\mathbb{Q})\backslash(G(\mathbf{A}_\mathbb{Q}^\infty) \times X)/U,$$

where $U \subset G(\mathbf{A}_\mathbb{Q}^\infty)$ is a compact open subgroup, and X is a hermitian symmetric space for $G(\mathbb{R})$. For instance, if $G = \mathrm{GSp}_{2n}$ is the general symplectic group and X is Siegel's upper half-space, then Sh_U is a *Siegel modular variety*, which classifies abelian surfaces of dimension n. The trouble is that for general G, such an X often does not exist (example: $G = \mathrm{GL}_n$ for $n \geq 3$), and so the theory of Shimura varieties for such groups remains elusive. Once again, we suffer under the tyranny of the archimedean place.

6 Drinfeld modular forms over general schemes, and their moduli

6.1 Line bundles

Let S be a scheme. If G is a group scheme over S, the *Lie algebra* $\mathrm{Lie}\, G$ is an $\mathcal{O}_S$-module. If $S = \mathrm{Spec}\, R$ is affine, then $\mathrm{Lie}\, G$ is the kernel of $G(R[\varepsilon]) \to G(R)$; one shows this is an R-module and that it "glues" to give a functor from group schemes over S to $\mathcal{O}_S$-modules.

A *line bundle* over S is a group scheme $L \to S$ such that locally on S we have $L \cong \mathbf{G}_{a,S}$. (Of course there is a similar definition for an n-dimensional vector bundle; these are locally isomorphic to $\mathbf{G}_{a,S}^n$.) Since $\mathrm{Lie}\, \mathbf{G}_{a,S} = \mathcal{O}_S$, the Lie algebra $\mathrm{Lie}\, L$ is a locally free $\mathcal{O}_S$-module of rank 1; that is, $\mathrm{Lie}\, L$ is an *invertible* $\mathcal{O}_S$-module.

In literature the terms "line bundle" and "invertible $\mathcal{O}_S$-module" (or "vector bundle" and "locally free $\mathcal{O}_S$-module") are sometimes confused[6]. There are functors in both directions: We have $L \mapsto \mathrm{Lie}\, L$ in one direction. In

[6] Adding to the confusion is the fact that invertible modules can be confused with Weil divisors. There is a homomorphism $\mathrm{Pic}\, S \to \mathrm{Cl} S$ from classes of invertible modules to classes of Weil divisors, which is an isomorphism when S is sufficiently nice.

the other, we may associate to an invertible $\mathcal{O}_S$-module $\mathcal{L}$ the line bundle $\operatorname{Spec}\operatorname{Sym}\mathcal{L}$, where $\operatorname{Sym}\mathcal{L} = \bigoplus_{n\geq 0}\operatorname{Sym}^n\mathcal{L}$, an $\mathcal{O}_S$-algebra.

Nonetheless, the functor line bundles to invertible $\mathcal{O}_S$-modules is not necessarily an equivalence of categories! It is essentially surjective and full (owing to the functor in the opposite direction), but it is not faithful. The issue is that in characteristic p, a morphism between line bundles can have derivative 0. To wit, the Frobenius morphism $x \mapsto x^p$ on $\mathbf{G}_{a,S}$ (S of characteristic p) induces the zero morphism on $\operatorname{Lie}\mathbf{G}_{a,S} = \mathcal{O}_S$.

Indeed, for an affine scheme $S = \operatorname{Spec} R$ of characteristic p, we have $\operatorname{End}\mathbf{G}_{a,S} = R\{\tau\}$, the twisted polynomial ring with $\tau a = a^p\tau$, $a \in R$. Explanation: Such an endomorphism corresponds to an R-algebra homomorphism $R[z] \to R[z]$, which much be the substitution $z \mapsto f(z)$ for some polynomial $f(z) \in R[z]$. The condition that the endomorphism preserves the group structure on $\mathbf{G}_{a,S}$ corresponds to the condition that $f(z)$ is *additive*, namely that $f(x+y) = f(x) + f(y)$. Thus $\operatorname{End}\mathbf{G}_{a,S} \cong R[z]^{\mathbb{F}_p\text{-lin}} \cong R\{\tau\}$.

The derivative map $\operatorname{End}\mathbf{G}_{a,S} \to \operatorname{End}\operatorname{Lie}\mathbf{G}_{a,S}$ corresponds to the ring homomorphism $R\{\tau\} \to R$ sending a polynomial to its constant term. (Or, if you like, it sends an $\mathbb{F}_p$-linear polynomial $f(z) \in R[z]$ to its linear term $f'(0)$.)

6.2 Drinfeld modules over schemes

Definition 6.2.1. Let S be a scheme equipped with a morphism $S \to \operatorname{Spec} A$. A *Drinfeld A-module* of rank d over S is a line bundle $L \to S$ together with a ring homomorphism $\phi\colon A \to \operatorname{End} L$. It is required that:

1. The derivative $\phi\colon A \to \operatorname{End}\operatorname{Lie} L$ agrees with the map $A \to H^0(S, \mathcal{O}_S) \to \operatorname{End}\operatorname{Lie} L$,
2. For every point $x = \operatorname{Spec} F$ of S, the composite $A \to \operatorname{End} L \to \operatorname{End} L_x \cong F\{\tau\}$ is a Drinfeld A-module of rank d over F.

Locally we have $S = \operatorname{Spec} R$ and $L = \mathbf{G}_{a,\operatorname{Spec} R}$, so that $\operatorname{End} L \cong R\{\tau\}$ (where τ is the pth power Frobenius map). Thus our Drinfeld module is a ring homomorphism $\phi\colon A \to R\{\tau\}$, $\alpha \mapsto \phi_\alpha$.

Some easy observations: An element $a_0 + a_1\tau_q + \dots$ of $R\{\tau\}$ is a unit if and only if $a_0 \in R^\times$ and if $a_1, a_2, \dots$ are nilpotent. If a unit in $R\{\tau\}$ has prime-to-p order, it must lie in $R^\times$.

Let us write $q = p^f$ for the cardinality of the residue field of K. Condition (1) states that ϕ_α has constant term α for all $\alpha \in A$. By the observation of the previous paragraph, we have $\phi_a = a$ for $a \in \mathbb{F}_q \subset A$. Therefore ϕ_α commutes with $\mathbb{F}_q^\times$ for all $\alpha \in A$, from which it follows that each ϕ_α is actually a polynomial in $\tau_q = \tau^f$, so that ϕ is a ring homomorphism $A \to R\{\tau_q\}$.

Condition (2) states that the coefficient of $\tau_q^{d \deg \alpha}$ of $\phi_\alpha \in R\{\tau_q\}$ is invertible (it lies in no prime ideal), and all coefficients beyond $\tau_q^{d \deg \alpha}$ are nilpotent (they lie in every prime ideal). We now apply the following lemma:

Lemma 6.2.2 (Drinfeld, Prop. 5.2)**.** *Let $f = \sum_n a_n \tau_q^n \in R\{\tau_q\}$. Assume that there exists $d \geq 1$ such that a_d is invertible and such that a_n is nilpotent for all $n > d$. Then there exists a unique $g \in R\{\tau_q\}$ with $g \equiv 1 + g_1\tau_q + g_2\tau_q^2 + \cdots \in R\{\tau_q\}$ such that g_i is nilpotent for all $i \geq 1$ and such that $g^{-1}fg$ has degree d.*

Therefore there exists a unique $g_\alpha \in R\{\tau_q\}^\times$ with constant term 1 such that $g_\alpha^{-1}\phi_\alpha g_\alpha$ has degree $d \deg \alpha$. Since ϕ_α and ϕ_β commute for $\alpha, \beta \in A$, a now-familiar argument shows that all g_α are equal (to g, say), so that $\phi' = g^{-1}\phi g$ has the property that $\deg \phi'_\alpha = d \deg \alpha$ for all $\alpha \in A$. Such a homomorphism is called a *standard* Drinfeld A-module of rank d over R; we have just shown that every Drinfeld A-module of rank d is locally isomorphic to a standard one, and that every automorphism of a standard Drinfeld A-module is conjugation by an element of $R^\times$.

6.3 Torsion subgroups and level structures

Let $S \to \operatorname{Spec} A$ be a scheme, let (L, ϕ) be a Drinfeld A-module of rank d over S. For a nonzero $\alpha \in A$, we have the morphism $\phi_\alpha \colon L \to L$ of group schemes over S. The *torsion subgroup* $\phi[\alpha]$ is defined as the scheme-theoretic kernel of ϕ_α. It is a finite group scheme over S, which admits an action of A/α.

Proposition 6.3.1. *The following are equivalent.*

1. *The image of S in $\operatorname{Spec} A$ is disjoint from $V(\alpha)$.*
2. *For all points $\operatorname{Spec} F$ of S, the pullback of (L, ϕ) to F has height 0.*
3. *The morphism $\phi_\alpha \colon L \to L$ is étale.*
4. *The torsion subgroup $\phi[\alpha]$ is a finite étale group scheme over S.*

Proof. Since these properties are local on S, it suffices to treat the case of a standard Drinfeld A-module over an A-algebra R, with structure morphism $i\colon A \to R$. There, the morphism $\phi_\alpha\colon \mathbf{G}_{a,R} \to \mathbf{G}_{a,R}$ is given by $\phi_\alpha(z) = i(\alpha)z + a_1 z^q + a_2 z^{q^2} + \cdots + a_d z^{q^d}$, which is étale if and only if $\phi'_\alpha(z) = i(\alpha)$ is invertible in R, which is the case if and only if the image of $\operatorname{Spec} R \to \operatorname{Spec} A$ is disjoint from $V(\alpha)$. Similarly, the torsion subgroup $\phi[\alpha]$ is $\operatorname{Spec} R[z]/\phi_\alpha(z)$, which is étale if and only if the linear term of $\phi_\alpha(z)$ is invertible in R. □

For a nonzero ideal $N \subset A$, we define the torsion subgroup $\phi[N]$ as the scheme-theoretic intersection of the $\phi[\alpha]$ for the nonzero $\alpha \in N$. Then $\phi[N]$ is a finite group scheme with an action of A/N. It is étale if and only if the image of $S \to \operatorname{Spec} A$ is disjoint from $V(N)$.

Naively, a level N structure on (L, ϕ) is an isomorphism of A/N-modules $\psi\colon (N^{-1}/A)^d \to \phi[\alpha](S)$. This definition works fine if $S = \operatorname{Spec} F$ for an A-field F whose characteristic is prime to N. But we run into problems if S is disconnected (in which case $\phi[\alpha](S)$ can have rank higher than d) or if its image in $\operatorname{Spec} A$ meets $V(N)$. In the latter case, one can have $S = \operatorname{Spec} F$ for an algebraically closed A-field F, but $\phi[\alpha](S)$ will fail to be a free A/N-module of rank d. These problems are remedied by the following definition.

Definition 6.3.2. A *Drinfeld level N structure* on (L, ϕ) is a homomorphism of A/N-modules

$$\psi\colon (N^{-1}/A)^d_S \to L(S),$$

such that, as divisors of L, $\phi[N]$ coincides with $\sum_{\beta \in (N^{-1}/A^d}[\psi(\alpha)]$.

(We remark that it is equivalent to impose the above condition for all $\mathfrak{p} \in V(N)$, which is the way Drinfeld defines it.)

In the case of a standard Drinfeld module $\phi\colon A \to R\{\tau\}$, a Drinfeld level N structure is an A/N-module homomorphism $\psi\colon (N^{-1}/A)^d \to R$, such that the ideal generated by $\phi_\alpha(z) \in R[z]$ for $\alpha \in N$ coincides with the principal ideal generated by $(\prod_{\beta \in (N^{-1}/A)^d}(z - \psi(\beta))$.

Exercise. If $S \to \operatorname{Spec} A$ is disjoint from $V(N)$, show that a Drinfeld level N structure is the same as an isomorphism of A/N-module schemes $(N^{-1}/A)^d_S \to \phi[N]$. If N is not the unit ideal, then the existence of a Drinfeld level N structure shows that the line bundle $L \to S$ is trivial.

Exercise. If $S = \operatorname{Spec} F$ for an A-field F of nonzero characteristic $\mathfrak{p}$, show that for all $e \geq 1$, an A/gp^e-module homomorphism $\psi\colon (\mathfrak{p}^{-e}/A^d \to \phi[\mathfrak{p}^e](F)$ is a Drinfeld level $\mathfrak{p}^e$-structure if and only if ψ is surjective. In

particular, if the height of ϕ is d, then there is exactly one Drinfeld level $\mathfrak{p}^e$-structure, namely the zero map.

Lemma 6.3.3. *Let $N \subset A$ be a nonzero ideal. Given a Drinfeld A-module (L, ϕ) over S, there exists an étale surjection $S' \to S$ such that the pullback of (L, ϕ) to S' admits a Drinfeld level N structure.*

6.4 Moduli of Drinfeld modules

It is define to define moduli spaces of Drinfeld modules over general schemes.

Definition 6.4.1. Let $N \subset A$ be a nonzero ideal, and let $d \geq 1$. For a scheme $S \to \operatorname{Spec} A$, we let $M_N^d(S)$ denote the set of isomorphism classes of triples (L, ϕ, ψ), where (L, ϕ) is a Drinfeld A-modules of rank d over S and ψ is a Drinfeld level N structure.

Proposition 6.4.2. *Assume that N is divisible by at least two primes. Then M_N^d is representable by a scheme of finite type over* $\operatorname{Spec} A$.

Proof. Let P be a prime dividing N. It suffices to show that the restriction of M_N^d to the category of schemes over $\operatorname{Spec} A \setminus \{P\}$ is a representable by a scheme of finite type. Indeed, one can then repeat the process for another prime P' dividing N, and then glue the resulting schemes together to obtain the desired scheme over $\operatorname{Spec} A$.

Let $v \in (P^{-1}/A)^d$ be nonzero. Let $S = \operatorname{Spec} R \to \operatorname{Spec} A \setminus \{P\}$ be an scheme, and let $(L, \phi, \psi) \in M_N^d(S)$. The restriction of ψ to $(P^{-1}/A)^d$ induces an isomorphism onto $\phi[P](S)$ Then $\psi(v) \in L(S)$ is nonzero everywhere on S, so that we have a distinguished trivialization $L \cong \mathbf{G}_{a,S}$. Therefore $M_N^d(S)$ is in bijection with the set of isomorphism classes of pairs (ϕ, ψ), where $\phi \colon A \to R\{\tau\}$ is a standard Drinfeld module and $\psi \colon (N^{-1}/A)^d \to R$ is a level N structure.

Now we make an important observaton: an isomorphism between two standard Drinfeld A-modules over R is given by an element of $R^\times$. Furthermore, $\psi(v) \in R$ is invertible. Therefore, every isomorphism class of standard Drinfeld A-modules over R contains a *unique* pair (ϕ, ψ) such that $\psi(v) = 1$.

We can now write down a scheme which represents M_N^d over $\operatorname{Spec} A \setminus \{P\}$. Let $\alpha_1, \ldots, \alpha_r$ generate A as a k-algebra. We can create an affine space over $\operatorname{Spec} A \setminus \{P\}$ of dimension $d \sum_i \deg \alpha_i$ whose coordinates represent the coefficients of each $\phi_{\alpha_i} \in R\{\tau\}$. Atop this, we can create a larger affine space

which parametrizes group homomorphisms $\psi\colon (N^{-1}/A)^d \to R$. Within this large affine space, we can impose the conditions which ensure that ϕ is a Drinfeld A-module of rank d (namely, that the leading coefficients of ϕ_{α_i} are invertible, and that they commute with one another) and that ψ is a Drinfeld level N structure satisfying $\psi(v) = 1$ (this can be done by using the condition that the ideal generated by the $\phi_{\alpha_i}(z) \in R[z]$ for $\alpha \in N$ coincides with the principal ideal generated by $(\prod_{\beta \in (N^{-1}/A)^d}(z - \psi(\beta)))$. We conclude that the restriction of M_N^d to $\operatorname{Spec} A \setminus \{P\}$ is a locally closed subset of an affine space. □

6.5 A digression on stacks

7 Formal A_v-modules

We have constructed a finite type scheme $M_N^d \to \operatorname{Spec} A$ which parametrizes Drinfeld A-modules with level structure, but the construction of M_N^d tells little about its geometry. It is not even clear what the dimensions of the fibers of this map are! The purpose of this section is to prove the following theorem:

Theorem 7.0.1. *Assume that $N \subset A$ is a nonzero ideal divisible by at least two primes. Then M_N^d is a smooth k-scheme of dimension d. Furthermore, $M_N^d \to \operatorname{Spec} A$ is smooth of relative dimension $d-1$ away from $V(N)$. Finally, if $N' \subset N$, then $M_{N'}^d \to M_N^d$ is finite and flat.*

The statements above are of a local nature, so that we can verify them all in a neighborhood of a point $x \in M_N^d$. If x lies over the generic point of A, we can appeal to our analytic description of $M_N^d(C)$ as a smooth rigid-analytic space of dimension $d-1$.

Therefore let $v \in \operatorname{Spec} A$ be a nonzero prime, with residue field k_v, and suppose that $x \in M_N^d(k_v)$. Let $\mathcal{O}_{M_N^d,x}$ be the local ring at this point. To prove the above theorem we can focus on the completion $\hat{\mathcal{O}}_{M_N^d,x}$, which is an algebra over the completion $A_v = \hat{\mathcal{O}}_{\operatorname{Spec} A,v}$. We claim that $\hat{\mathcal{O}}_{M_N^d,x}$ is a regular local ring, and that if v does not divide N, then $\hat{\mathcal{O}}_{M_N^d,x}$ is formally smooth of dimension $d-1$ over A_v. The key to proving these claims is to prove $\hat{\mathcal{O}}_{M_N^d,x}$ can be interpreted as a deformation ring in its own right.

7.1 Formal $\mathcal{O}$-modules: definition

Let K be a nonarchimedean local field, with ring of integers $\mathcal{O}$, uniformizer π, and residue field $k = \mathbb{F}_q$.

Definition 7.1.1. Let R be an $\mathcal{O}$-algebra. A (one-dimensional) *formal $\mathcal{O}$-module* $\mathcal{F}$ over R consists of the following data:

1. A power series $X +_{\mathcal{F}} Y = X + Y + \cdots \in R[X, Y]$, which is commutative and associative,

2. For each $\alpha \in \mathcal{O}$, a power series $[\alpha]_{\mathcal{F}}(X) = \alpha X + \cdots \in R[\![X]\!]$; the $[\alpha]_{\mathcal{F}}$ must commute with one another and distribute over $X +_{\mathcal{F}} Y$.

A morphism $f\colon \mathcal{F} \to \mathcal{G}$ of formal $\mathcal{O}$-modules is a power series $f(X) \in R[\![X]\!]$ for which $f(X +_{\mathcal{F}} Y) = X +_{\mathcal{G}} Y$ and $f([\alpha]_{\mathcal{F}}(X)) = [\alpha]_{\mathcal{G}}(X)$.

Note the close connection with Drinfeld A-modules. A Drinfeld A-module is an A-module structure on the scheme $\mathbf{A}^1$, whereas a formal $\mathcal{O}$-module is an $\mathcal{O}$-module struture on the formal scheme $\hat{\mathbf{A}}^1$.

Examples:

1. The formal additive group $\hat{\mathbf{G}}_a$ over $\mathcal{O}$ has addition law $X + Y$ and multiplication law αX.

2. The formal multiplicative group $\hat{\mathbf{G}}_m$ over $\mathbb{Z}_p$ becomes a formal $\mathbb{Z}_p$-module, via $X +_{\hat{\mathbf{G}}_m} Y = X + Y + XY$ and $[\alpha]_{\mathbf{G}_m}(X) = (1 + X)^\alpha - 1$ (the latter considered as a power series in $\mathbb{Z}_p[\![X]\!]$).

3. Suppose E is an elliptic curve over a $\mathbb{Z}_p$-algebra R. Choose a local coordinate X around the origin, so that the completion $\hat{E}$ may be identified with the formal scheme $\operatorname{Spf} R[\![X]\!]$. Then $\hat{E}$ has the structure of a formal $\mathbb{Z}_p$-module.

4. Suppose $\phi\colon A \to R\{\tau\}$ is a Drinfeld A-module over R, where R is a complete A_v-algebra. After completing v-adically we obtain a ring homomorphism $\phi_v\colon A_v \to R\{\{\tau\}\}$. This defines a formal $\mathcal{O}_v$-module structure over R, where the addition law is $X + Y$ and where multiplication by $\alpha \in A_v$ is the k-linear power series $\phi_{v,\alpha}(X)$.

7.2 Formal $\mathcal{O}$-modules over a k-algebra

Definition 7.2.1. Let R be a local k-algebra, and let $\mathcal{F}$ be a formal $\mathcal{O}$-module over R. We say $\mathcal{F}$ is *π-divisible* if the substitution map $R[\![X]\!] \to R[\![X]\!]$ sending $X \to [\pi]_{\mathcal{F}}(X)$ is finite. If so, the *height* of $\mathcal{F}$ is the rank of $R[\![X]\!]$ over itself via this map.

In fact the height h is the largest integer for which $[\pi]_{\mathcal{F}}(X) = g(X^{q^h})$ for a powerseries $g(X) \in R[\![X]\!]$. The v-adic completion of a rank d Drinfeld A-module is a formal A_v-module of height d. The formal additive group $\hat{\mathbf{G}}_a$ is not π-divisible.

The category of π-divisible $\mathcal{O}$-modules over $\overline{k}$ is well understood.

Proposition 7.2.2. *For each $h \geq 1$, there exists a formal $\mathcal{O}$-module over $\overline{k}$ of height h, which is unique up to isomorphism.*

7.3 Deformations of formal $\mathcal{O}$-modules

7.4 Connection to moduli of Drinfeld modules

8 Drinfeld modules of rank 1: "complex multiplication"

8.1 A review of class field theory

Let K be a global field. We have the topological group of ideles, denoted $\mathbf{A}_K^\times$, and its discrete subgroup $K^\times \subset \mathbf{A}_K^\times$. Let $\mathbf{J}_K = K^\times \backslash \mathbf{A}_K^\times$, the *idele class group*. When K is a number field, J_K is related to the usual class group by the exact sequence

$$1 \to \mathcal{O}_K^\times \to \prod_{v \nmid \infty} \mathcal{O}_{K_v}^\times \times \prod_{v \mid \infty} K_v^\times \to J_K \to \mathrm{Cl}_K \to 1.$$

When K is the function field of a curve X with field of constants k, there is an analogous sequence:

$$1 \to k^\times \to \prod_{v \in |X|} \mathcal{O}_{K_v}^\times \to J_K \to \mathrm{Pic}\, X \to 1$$

The thrust of class field theory is that there exists a continuous homomorphism

$$\mathrm{rec}_K \colon \mathbf{J}_K \to \mathrm{Gal}(K^{\mathrm{ab}}/K),$$

the *reciprocity map*, which satisfies a few key properties, including:

- rec_K has dense image.
- The kernel of rec_K is the neutral component of J_K.
- For each $v \in |K|$, rec_K is compatible with the local reciprocity map for K_v. In particular, if v is finite and $\pi_v \in K_v$ is a uniformizer, then $\mathrm{rec}_K(\pi_v)$ is a Frobenius element for v.
- For a finite abelian extension L/K, rec_K induces an isomorphism

$$\mathbf{J}_K/N_{L/K}(\mathbf{J}_L) \cong \mathrm{Gal}(L/K).$$

- The map $L \mapsto N_{L/K}(\mathbf{J}_L)$ is a one-to-one correspondence between finite abelian extensions L/K and open subgroups of $\mathbf{J}_K$ of finite index.

In no case is rec_K an isomorphism of topological groups; indeed $\mathbf{J}_K$ is never a profinite group. If K is a number field, then rec_K is surjective but not injective; the kernel of rec_K is the manifold $(\mathbb{R}^\times_{>0})^{r_1} \times (\mathbb{C}^\times)^{r_2}$. If K is a function field, the reverse is true: rec_K is injective but not surjective. In this case, the failure of rec_K to be surjective is explained by the field of constants k: we have a commutative diagram

$$\begin{array}{ccc} J_K & \xrightarrow{\mathrm{rec}_K} & \mathrm{Gal}(K^{\mathrm{ab}}/K) \\ \downarrow & & \downarrow \\ \mathbb{Z} & \longrightarrow & \hat{\mathbb{Z}} \end{array}$$

where the bottom right $\hat{\mathbb{Z}}$ is $\mathrm{Gal}(K\overline{k}/K) \cong \mathrm{Gal}(\overline{k}/k)$.

There arises the natural problem of recovering a finite abelian extension L/K from its norm subgroup $N_{L/K}(\mathbf{J}_L)$. This is the essence of *Hilbert's 12th problem*. It suffices to accomplish this for a system of open finite index subgroups of $\mathbf{J}_K$. There is a convenient such system indexed by what are classically called moduli. A *modulus* is a formal product $\mathfrak{m} = prod_v v^{n_v}$ for nonnegative integers n_v, almost all of which are 0. We demand $n_v \in \{0, 1\}$

if v is real, and $n_v = 0$ if v is complex. For such an $\mathfrak{m}$ we let $U_{\mathfrak{m}} \subset \mathbf{A}_K^\times$ be the subgroup defined by imposing the following conditions on an element $(a_v) \in \mathbf{A}_K^\times$:

- $v(a_v - 1) \geq n_v$ for v nonarchimedean with $a_v > 0$,
- $a_v \in \mathcal{O}_v^\times$ for v nonarchimedean with $a_v = 0$,
- $a_v > 0$ for v real with $a_v = 1$.

Observe that any open subgroup of $\mathbf{A}_K^\times$ is contained in $U_{\mathfrak{m}}$ for some modulus $\mathfrak{m}$. Define the *ray class group*

$$C_{\mathfrak{m}} = \mathbf{J}_K / U_{\mathfrak{m}},$$

a (discrete) abelian group. To specify an open subgroup of finite index in $\mathbf{J}_K$, it suffices to give a modulus $\mathfrak{m}$ and a finite index subgroup $H \subset C_{\mathfrak{m}}$. The corresponding abelian extension K_H/K should have these properties:

1. K_H/K is unramified outside $\mathfrak{m}$.
2. There is an isomorphism $C_{\mathfrak{m}}/H \to \mathrm{Gal}(K_H/K)$, which carries (the image of) a uniformizer at an nonarchimedean place $v \nmid \mathfrak{m}$ (respectively, a negative element at an archimedean place $v \nmid \mathfrak{m}$) to the Frobenius element $\mathrm{Frob}_v \in \mathrm{Gal}(K_H/K)$.

Let $K_{\mathfrak{m}}$ be the compositum of all finite extensions obtained this way; then $\mathrm{Gal}(K_{\mathfrak{m}}/K)$ is isomorphic to the profinite completion of $C_{\mathfrak{m}}$. Also note that K^{ab} is the compositum of the $K_{\mathfrak{m}}$ as $\mathfrak{m}$ runs through all moduli of K.

8.2 Hilbert's 12th problem for number fields

If K is a number field, $C_{\mathfrak{m}}$ is finite. There is an interpretation of $C_{\mathfrak{m}}$ in terms of ideals: it is the group of fractional ideals which are prime to $\mathfrak{m}$ modulo the subgroup of principal ideals generated by elements of $K^\times \cap U_{\mathfrak{m}}$. To solve Hilbert's 12th problem for K, it suffices to construct, for all moduli $\mathfrak{m}$, the finite abelian extension $K_{\mathfrak{m}}/K$. When $\mathfrak{m} = 1$ is the trivial modulus, $C_{\mathfrak{m}} = \mathrm{Cl}_K$ is the class group of K, and K_1 is the Hilbert class field.

When $K = \mathbb{Q}$, Hilbert's 12th problem has a complete solution: if $\mathfrak{m} = m(\infty)$ for an integer $m \geq 1$, then $\mathbb{Q}_{\mathfrak{m}} = \mathbb{Q}(\mu_m)$ is a cyclotomic field. (If $\mathfrak{m} = m$, then $\mathbb{Q}_{\mathfrak{m}}$ is the maximal totally real subfield of $\mathbb{Q}(\mu_m)$.)

When K is an imaginary quadratic field, Hilbert's 12th problem has a complete solution as well. This time, a modulus may be identified with a nonzero ideal $\mathfrak{m} \subset \mathcal{O}_K$. The Hilbert class field K_1 is obtained by adjoining to K the j-invariant of one elliptic curve (or equivalently of all elliptic curves) with complex multiplication by $\mathcal{O}_K$. Then $K_\mathfrak{m}$ is obtained by adjoining to K_1 the $\mathfrak{m}$-torsion of one (equivalently, all) of these curves.

8.3 Hilbert's 12th problem for function fields

If $K = k(X)$ is a function field, then a modulus $\mathfrak{m}$ is one and the same as an effective divisor of X. The ray class group C_1 may be identified with $\operatorname{Pic} X$, the quotient of the divisor group by the subgroup of principal divisors. There is an exact sequence

$$0 \to \operatorname{Pic}^0 X \to \operatorname{Pic} X \to \mathbb{Z} \to 0$$

given by the degree map. Note that $\operatorname{Pic}^0 X$ can be "geometrized", in the sense that there is an abelian variety J/k, the *Jacobian* of X, with $J(k) = \operatorname{Pic}^0 X$. In particular $\operatorname{Pic}^0 X$ is finite. The Jacobian represents the functor which assigns to a k-scheme S the quotient of $\operatorname{Pic}^0(X \times_k S)$ by the image of $\operatorname{Pic}^0 S$.

For a general modulus $\mathfrak{m}$, $C_\mathfrak{m}$ is isomorphic to the group $\operatorname{Div}^\mathfrak{m} X$ of divisors away from $\mathfrak{m}$ by the subgroup generated by principal divisors attached to functions$f \equiv 1 \pmod{\mathfrak{m}}$. (This latter expression is shorthand for $f \in K^\times \cap U_\mathfrak{m}$.) Once again there is an exact sequence

$$0 \to C_\mathfrak{m}^0 \to C_\mathfrak{m} \to \mathbb{Z} \to 0.$$

The group $C_\mathfrak{m}^0$ can also be geometrized. There is a commutative group scheme $J_\mathfrak{m}/k$, the *generalized Jacobian*, which fits into an exact sequence of group schemes

$$0 \to H_\mathfrak{m} \to J_\mathfrak{m} \to J \to 0,$$

where $H_\mathfrak{m}$ is a product of a torus and a unipotent group. We refer the reader to Serre's book *Algebraic Groups and Class Fields* for the construction and properties of $J_\mathfrak{m}$. For now we just note that $J_\mathfrak{m}(k) = C_\mathfrak{m}^0$.

Recall that to solve Hilbert's 12th problem for K, we want to associate to each finite index subgroup $H \subset C_\mathfrak{m}$ its corresponding finite abelian extension K_H/K. There's one easy special case: If H is the preimage of $n\mathbb{Z} \subset \mathbb{Z}$ under the degree map $C_\mathfrak{m} \to \mathbb{Z}$, then $K_H = Kk_n$, where k_n/k is the extension of degree n.

Let $\infty \in |X|$ be a point not dividing $\mathfrak{m}$, and let $A = H^0(X \setminus \{\infty\}, \mathcal{O}_X)$ as usual. We may consider $\mathfrak{m}$ as an ideal of A. Let $\pi_\infty \in K_\infty$ be a uniformizer. Then $\pi_\infty^{\mathbb{Z}} \subset C_{\mathfrak{m}}$ has finite index, and the quotient, call it $\mathrm{Cl}_{\mathfrak{m}} A$, can be computed as follows:

$$\begin{aligned} \mathrm{Cl}_{\mathfrak{m}} A &= K^\times \backslash \mathbf{A}_K^\times / U_{\mathfrak{m}} \pi^{\mathbb{Z}} \\ &= K^\times \backslash \mathbf{A}_K^{\infty,\times} / U_{\mathfrak{m}}^\infty \\ &\cong I^{(\mathfrak{m})} / P^{(\mathfrak{m})}, \end{aligned}$$

where $I^{(\mathfrak{m})}$ is the group of fractional ideals of A which are prime to $\mathfrak{m}$, and $P^{(\mathfrak{m})}$ is the subgroup of principal ideals of the form (f), where $f \equiv 1 \pmod{\mathfrak{m}}$. In particular $\mathrm{Cl}_1 A$ is the usual class group of A.

Let $K_{(\infty),\mathfrak{m}}$ be the extension corresponding to the finite quotient $\mathrm{Cl}_{\mathfrak{m}} A$ of $\mathbf{J}_K$. Since the image of π_∞ in $\mathrm{Cl}_{\mathfrak{m}} A$ is trivial, we find that $\mathrm{Frob}_\infty = 1$ in this extension, which is to say that $K_{(\infty),\mathfrak{m}}/K$ is split at ∞. Let $n = \deg \infty$; then $K_{(\infty),\mathfrak{m}} \cap \overline{k} = k_n$. We have an exact sequence

$$0 \to C_m^\circ \to \mathrm{Cl}_{\mathfrak{m}} A \to \mathbb{Z}/n\mathbb{Z} \to 0.$$

Here is the relevant diagram of fields:

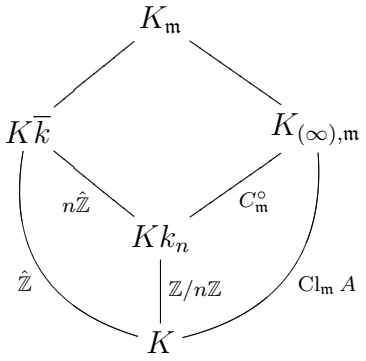

We have the following interpretation for $K_{(\infty)} = \cup_{\mathfrak{m}} K_{\infty,m}$: it is the maximal abelian extension of K in which ∞ splits.

We have now reduced the problem to constructing the field $K_{(\infty),\mathfrak{m}}$. After replacing K with $K\overline{k}_n$, we may assume that $n = \deg \infty = 1$. Since $K_{(\infty),\mathfrak{m}} \cap \overline{k} = k$, $K_{(\infty),\mathfrak{m}}$ is the function field of an absolutely irreducible curve $X_{(\infty),\mathfrak{m}}$, which is a generically étale cover of X with group $C_{\mathfrak{m}}^\circ$. This cover is unramified outside $\mathfrak{m}$.

We have a rational map $X \dashrightarrow J_{\mathfrak{m}}$ defined by $P \mapsto [P]-[\infty]$. (It is defined away from the support of $\mathfrak{m}$.) We also define the *Lang isogeny* $L\colon J_{\mathfrak{m}} \to J_{\mathfrak{m}}$, defined by $x \mapsto \mathrm{Frob}_q(x) - x$. Its kernel is $J_{\mathfrak{m}}(\mathbb{F}_q) = C^0_{\mathfrak{m}}$.

Theorem 8.3.1. *$X_{(\infty),\mathfrak{m}}$ fits into a Cartesian diagram*

$$\begin{array}{ccc} X_{(\infty),\mathfrak{m}} & \dashrightarrow & J_{\mathfrak{m}} \\ \downarrow & & \downarrow L \\ X & \dashrightarrow & J_{\mathfrak{m}} \end{array}$$

8.4 Explicit construction of abelian extensions of function fields: Examples

Example 8.4.1. Let $X = \mathbf{P}^1_k$, with coordinate T, so that $K = k(T)$. Let $t_1, \ldots, t_n \in k$ be distinct, and let $\mathfrak{m} = \prod_i (T - t_i)$. Then $J_{\mathfrak{m}}$ is the quotient of the torus $\mathbf{G}^n_{m,k}$ by the image of the diagonal map $\mathbf{G}_{m,k} \to \mathbf{G}^n_{m,k}$. Let us identify $J_{\mathfrak{m}} \cong \mathbf{G}^{n-1}_{m,k}$ by projecting onto the first $n-1$ factors. The Lang isogeny on $J_{\mathfrak{m}}$ becomes $(x_1, \ldots, x_{n-1}) \mapsto (x_1^{q-1}, \ldots, x_{n-1}^{q-1})$. Furthermore, the map $X \backslash V(\mathfrak{m}) \to J_{\mathfrak{m}}$ is $T \mapsto ((T-t_1)/(T-t_n), \ldots, (T-t_{n-1})/(T-t_n))$. Therefore

$$K_{(\infty),\mathfrak{m}} = K\left(\sqrt[q-1]{\frac{T-t_i}{T-t_n}} \right)_{1 \le i \le n-1}$$

Example 8.4.2. Once again let $X = \mathbf{P}^1_k$, but this time let $\mathfrak{m} = (T^2)$. Then $J_{\mathfrak{m}}$ is the group scheme over k for which $J_{\mathfrak{m}}(R) = (R[T]/T^2)^*/R^*$ for an R algebra. Then $J_{\mathfrak{m}} \cong \mathbf{G}_{a,k}$ via $a + bT \mapsto a^{-1}b$. The morphism $X \backslash \{0\} \to J_{\mathfrak{m}}$ can be computed as follows: given a nonzero $t \in X$, the divisor $(t) - (\infty)$ is principal with generator $T - t$, whose image in $J_{\mathfrak{m}} \cong \mathbf{G}_{a,k}$ is $-t^{-1}$. The Lang map on $\mathbf{G}_{a,k}$ is $x \mapsto x^q - x$, and so $K_{(\infty),\mathfrak{m}} = K(u)$, where $u^q - u = -T^{-1}$.

Example 8.4.3. Let X/k be a curve of genus 1, and let $\infty \in X(k)$ be a rational point. Then we can give X the structure of an elliptic curve with origin ∞. Let $\mathfrak{m} = 1$, so that $J_{\mathfrak{m}} = J$ is the Jacobian of X. Then $X \to J$ is an isomorphism, and $X_{(\infty),1} \to X$ is the Lang isogeny $P \mapsto \mathrm{Frob}(P) - P$ on X. Thus in this case $K_{(\infty),1}$ happens to be isomorphic to K.

8.5 The shtuka correspondence for $d = 1$

There is another approach to constructing abelian extensions of function fields, which leverages the analogy between Drinfeld modules and elliptic curves with CM. Let K be a function field, let L/K be a finite extension, and let $\phi\colon A \to L\{\tau\}$ be a Drinfeld module of rank 1. Let $N \subset A$ be a nonzero ideal. Then $L(\phi[N])/L$ is an abelian extension: indeed, its Galois group is a subgroup of $(A/N)^\times$.

Example 8.5.1. Let $K = k(T)$, $A = k[T]$, and let $\phi\colon A \to A\{\tau\}$ be the Carlitz module, so that $\phi_T = T + \tau$. We can construct abelian extensions of K by adjoining to K the torsion in ϕ. Let $N \subset A$ be a nonzero ideal, with generator $\prod_i (T - t_i)$, where the $t_i \in k$ are distinct. Then

$$K(\phi[N]) = K(\sqrt[q-1]{-(T - t_i)})_{1 \le i \le n}.$$

This is a Galois extension of K with group $(A/N)^\times$. Note that the fixed field of the subgroup $k^\times \subset (A/N)^\times$ is the field $K_{(\infty),N}$ of Example 8.4.1. In fact, one can check that if $n \ge 2$, then the moduli space M^1_N is exactly the spectrum of the integral closure of A in $K_{(\infty),N}$. (One must quotient by $k^\times$ in the formation of M^1_N because this is the automorphism group of ϕ.)

Example 8.5.2. Once again let $K = k(T)$ and $A = k[T]$, but this time we consider the case $N = (T^2)$. Let $t = \sqrt[q-1]{-T}$, so that $K(\phi[T]) = K(t)$, and then

$$K(\phi[T^2]) = K(t, v),\ v^q + Tv = t,$$

a Galois extension of K with group $(k[T]/T^2)^\times$. The fixed field of $k^\times$ in this extension is $K(u)$, where $u = t^{-1}v$. Then $u^q - u = -T^{-1}$, so that $K(u)$ is in fact the ray class field $K_{(\infty),T^2}$ from 8.4.2.

The examples above suggest a close relationship between rank 1 Drinfeld modules and class field theory for function fields. The formal relationship can be formalized by the following theorem.

Theorem 8.5.3 (The shtuka correspondence for $d = 1$)**.** *Assume that $N \subset A$ is a nonzero ideal which is divisible by more than one prime. Let A_N be the integral closure of A in $K_{(\infty),N}$. There is an isomorphism $M^1_N \cong \operatorname{Spec} A_N$ of schemes over A.*

We will prove a generalization of this theorem in the next section, but first let's examine a consequence of it. We have $\mathrm{Gal}(K_{(\infty),N}/K) \cong (A/N)^\times/k^\times$. After taking quotients by $(A/N)^\times$, the isomorphism of the theorem becomes an isomorphism of stacks $M_1^1 \cong [(\operatorname{Spec} A_{(\infty),1})/k^\times]$ over $\operatorname{Spec} A$. In particular, *there should exist a rank 1 Drinfeld module defined over (the ring of integers of) $K_{(\infty),1}$*, just as there exists a CM elliptic curve defined over (the ring of integers of) the Hilbert class field of an imaginary quadratic field.

Example 8.5.4. In particular if $\operatorname{Pic} A = 0$, so that $K_{(\infty),1} = K$, then Theorem 8.5.3 predicts that there should exist a Drinfeld A-module of rank 1 over A itself. We know this to be the case when $A = k[T]$, because of the Carlitz module. There are a few other examples where $\operatorname{Pic} A = 0$ (in fact there are only finitely many, even if one varies the residue field k). One such is

$$A = \mathbb{F}_2[x,y]/(y^2 + y + x^3 + x + 1).$$

Then $A = H^0(X \setminus \{\infty\}, \mathcal{O}_X)$ for an elliptic curve $X/\mathbb{F}_2$ with identity $\infty \in X(\mathbb{F}_2)$. One checks that $\operatorname{Pic} A \cong X(\mathbb{F}_2)$ is the trivial group. We exhibit a Drinfeld A-module over A of degree 1:

$$\begin{aligned} \phi_x &= x + (x^2 + x)\tau + \tau^2 \\ \phi_y &= y + (y^2 + y)\tau + x(y^2 + y)\tau^2 + \tau^3 \end{aligned}$$

Remarkably, ϕ_x and ϕ_y commute with each other, and satisfy the same polynomial relation satisfied by x and y. [This example is due to Hayes.]

9 The shtuka correspondence

9.1 The category of Drinfeld shtukas

As usual, let X be a nonsingular projective curve over a finite field k. For a k-scheme S, let $X_S = X \times_k S$ be the base change. Given a morphism $\iota\colon S \to X$, we have the graph Γ_ι, which is the image of $\iota \times \mathrm{id}\colon S \to X_S$. Because X is a curve, Γ_ι has codimension 1; i.e., it is a Weil divisor of X_S.

We have the qth power Frobenius endomorphism $1 \times \mathrm{Frob}_S\colon X_S \to X_S$, which we will abbreviate simply as Frob_S.

Definition 9.1.1. Let X be an A-scheme, with structure morphism $\iota\colon X \to \operatorname{Spec} A$. A *Drinfeld A-shtuka* of rank d over S is a diagram

$$\begin{array}{ccccc} \operatorname{Frob}_S^* \mathcal{F} & & \xrightarrow{\alpha} & & \\ & & & & \mathcal{F}' \\ & & \xrightarrow{\beta} & & \\ \mathcal{F} & & & & \end{array} \tag{6}$$

in which $\mathcal{F}$ and $\mathcal{F}'$ are locally free $\mathcal{O}_X$-modules of rank d. It is required that:

1. $\operatorname{cok}\alpha$ is supported on Γ_ι, and the restriction of $\operatorname{cok}\alpha$ to Γ_ι is invertible.
2. $\operatorname{cok}\beta$ is supported on $\infty_S = \{\infty\} \times S \subset X_S$.

We write the above data as $\operatorname{Frob}_S^* \mathcal{F} \to \mathcal{F}' \leftarrow \mathcal{F}$.

Remark 9.1.2. For an arbitrary scheme X, a divisor $Z \subset X$, and an $\mathcal{O}_X$-module $\mathcal{M}$, the condition that $\mathcal{M}|_Z$ is invertible is equivalent to the condition that $\mathcal{M} \cong i_*\mathcal{I}$, where $i\colon Z \to X$ is the inclusion map and $\mathcal{I}$ is an invertible $\mathcal{O}_Z$-module.

Remark 9.1.3. Given an A-shtuka as in the definition, the composite $\beta \circ \alpha^{-1}$ defines a rational map $\phi\colon \operatorname{Frob}_S^* \mathcal{F} \dashrightarrow \mathcal{F}$, which is regular away from ∞_S. In fact, $\mathcal{F}'$ can be reconstructed from the rational map ϕ.

There is a slight generalization of this notion which is slightly more symmetric. Let us suppose the k-curve X is given, but that there is no special point $\infty \in X(k)$. Instead, let S be a k-scheme, and let $\iota_0\colon S \to X$ and $\iota_\infty\colon S \to X$ be two morphisms, corresponding to two points $O, \infty \in X(S)$. A *Drinfeld X-shtuka with zero O and pole ∞* is a diagram as in (6), such that $\operatorname{cok}\alpha \cong \Gamma_{\iota_0,*}\mathcal{I}_0$ for an invertible sheaf $\mathcal{I}_0$ on S, and similarly for ∞.

9.2 The general shtuka correspondence

Let $\infty \in X(k)$ be a k-rational point, and let $A = H^0(X \setminus \{\infty\}, \mathcal{O}_X)$.

Theorem 9.2.1. *Let $S \to \operatorname{Spec} A$ be a scheme. The following categories are equivalent:*

A. Drinfeld A-modules of rank d over S.

B. Drinfeld A-shtukas of rank d over S, which satisfy $\chi(\mathcal{F}) = 0$.

10 Automorphic forms and moduli of Drinfeld modules