THE SHIMURA-TANIYAMA FORMULA AND p-DIVISIBLE GROUPS

DANIEL LITT

1. Notation and Introduction

Let us fix the following notation:

- K is a number field;
- L is a CM field with totally real subfield L^+;
- (A, i) is an Abelian variety/K with CM $i : L \to \mathrm{End}^0_K(A)$ (remark: by increasing K we may assume A has good reduction at all places). By passing to an isogenous Abelian variety, we may assume $\mathcal{O}_L$ acts on A, rather than just an order.
- Φ is the CM type of (A, i), realized as the collection of embeddings $\phi_i : L \to \mathbb{C}$ appearing via the action of $L \otimes_{\mathbb{Q}} \mathbb{C}$ on $T_0(A(\mathbb{C})^{an})$. Recall that this is a set of representatives for the orbits of $\mathrm{Gal}(L/L^+)$ on the set of all embeddings $L \to \mathbb{C}$.
- $\mathcal{A}/\mathcal{O}_K$ is the proper Néron model of A; this is an Abelian scheme over $\mathrm{Spec}(\mathcal{O}_K)$. By functoriality of the Néron model construction, we have a map $i : L \hookrightarrow \mathrm{End}^0_S(\mathcal{A})$.
- $\mathfrak{p}$ is a prime of K with residue field k of size q; the place of K corresponding to $\mathfrak{p}$ will be denoted w;
- By v, I will denote a place of L with $v|q$.
- $\overline{A}$ is the reduction of $\mathcal{A}$ over k.

If you'd like, you can simply think of A as being defined by equations with coefficients in $\mathcal{O}_K$; then the reduction mod $\mathfrak{p}$ is simply the Abelian variety over k whose equations are given by reducing mod $\mathfrak{p}$ (which is smooth by assumption). Of course we've already used the Néron model repeatedly (for example, to get to the situation I've just described), and we will continue to do so, but if you'd like you can just imagine you've been given an Abelian scheme $\mathcal{A}$ over $\mathcal{O}_{K,\mathfrak{p}}$ with good reduction mod $\mathfrak{p}$ and with an embedding $i : L \hookrightarrow \mathrm{End}^0_{\mathcal{O}_K}(\mathcal{A})$, and the rest of the talk will go through without reference to Néron models.

The goal of this talk is to try to better understand the map

$$L \xhookrightarrow{i} \mathrm{End}^0_K(A) \hookrightarrow \mathrm{End}^0_{\mathcal{O}_K}(\mathcal{A}) \hookrightarrow \mathrm{End}^0_k(\mathcal{A}_k).$$

where the last map is injective by e.g. considering the ℓ-adic Tate module with $(\ell, q) = 1$.

There is a distinguished element of $\mathrm{End}_k(\mathcal{A}_k) \subset \mathrm{End}^0_k(\mathcal{A}_k)$, namely Frob_q, the so-called "absolute Frobenius" if I recall the terminology correctly. On an affine scheme $\mathrm{Spec}(R)/\mathbb{F}_q$, this is the map induced by the map

$$\mathrm{Frob}_q : R \to R$$

$$x \mapsto x^q.$$

By the "universal commutativity" of the Frobenius, Frob_q lies in the center of $\mathrm{End}^0_k(\mathcal{A}_k)$, which is simply L for reasons of size (that is, L is its own centralizer, by results Arnav proved—one may see this also by looking at $V_\ell(\mathcal{A}_k)$). Indeed, this map is an isogeny (it suffices to note it's non-zero) and thus it lies in the image of $L^\times$. Let $\pi \in L^\times$ be the unique element mapping to Frob_q. We'd like to try to pin down π as best we can, using only the data of the CM type Φ of A.

Note that $\pi \in \mathcal{O}_L$ as $\mathrm{Frob}_q \in \mathrm{End}_k(\mathcal{A}_k) \subset \mathrm{End}^0_k(\mathcal{A}_k)$, which is finitely generated as a $\mathbb{Z}$-module, so Frob_q lies in a $\mathbb{Z}$-finite subring of L, which is thus contained in $\mathcal{O}_L$. Let us try to pin it down a bit further.

Fix ℓ with $(\ell, q) = 1$. As $V_\ell(A) \otimes \overline{\mathbb{Q}_\ell}$ is a free rank one $L \otimes_{\mathbb{Q}} \overline{\mathbb{Q}_\ell} \simeq \prod_{\phi : L \to \overline{\mathbb{Q}_\ell}} \overline{\mathbb{Q}_\ell}_\phi$-module, we may identify the eigenspaces of Frobenius with the multiplicands corresponding to the embeddings $\phi : L \to \overline{\mathbb{Q}_\ell}$. As the determinant of Frobenius acting on the ℓ-adic Tate module is $q^{\dim(A)}$, the ℓ-Newton polygon for the characteristic polynomial of Frobenius is flat, and so all the Frobenius eigenvalues have ℓ-adic valuations 1.

These eigenvalues correspond to the image of π under the embeddings $L \to \overline{\mathbb{Q}_\ell}$, so we have that for any place $w|\ell$, $(\ell, q) = 1$, that $v_w(\pi) = 0$.

Furthermore, the Riemann hypothesis for Abelian varieties implies that the Archimedean valuations of π are all $q^{1/2}$. So all that is left is to pin down the valuations for places w with $w|\mathfrak{p}$; then we will have determined π up to a root of unity. And this is of course the best one can hope for if we are only using the CM type, since taking Galois twists of Abelian varieties will change the Frobenius lifts by a root of unity without changing the CM type.

So let us fix $v|q$. Let $\overline{\mathbb{Q}_p}$ be an algebraic closure of $\mathbb{Q}_p$. Then it is clear that any embedding $L \to \overline{\mathbb{Q}_p}$ induces a place on L over $\mathbb{Q}_p$; let H_v be the set of embeddings inducing v. $\#H_v = [L_v : \mathbb{Q}_p]$, as we have

$$\mathrm{Hom}_{\mathbb{Q}}(L, \overline{\mathbb{Q}_p}) = \mathrm{Hom}_{\mathbb{Q}_p}(\mathbb{Q}_p \otimes_{\mathbb{Q}} L, \overline{\mathbb{Q}_p}) = \mathrm{Hom}_{\mathbb{Q}_p}(\prod_{v|p} L_v, \overline{\mathbb{Q}_p}) = \bigsqcup_{v|p} \mathrm{Hom}_{\mathbb{Q}_p}(L_v, \overline{\mathbb{Q}_p})$$

whence we may identify both H_v with the set of embeddings $L_v \to \overline{\mathbb{Q}_p}$, which has cardinality $[L_v : \mathbb{Q}_p]$. Let $\Phi_v = \Phi \cap H_v$. The theorem we'd like to prove is:

Theorem 1. *With the notation we've just described*

$$\frac{\mathrm{ord}_v(\pi)}{\mathrm{ord}_v(q)} = \frac{\#\Phi_v}{\#H_v} \left(= \frac{\#\Phi_v}{[L_v : \mathbb{Q}_p]} \right).$$

Remark 1. *We've stated this theorem in the global setting—we'll prove the slightly more general result where $\mathcal{A}$ is an Abelian scheme over $\mathcal{O}_{K_w}$ with CM by L, whence the "global case" will follow immediately. Apparently the original proof of Shimura and Taniyama was global; the proof we give, due to Tate, is purely local.*

To prove this theorem, we'll have to make contact between the CM-type, which we've defined via the action of L on the tangent space of the generic fiber of $\mathcal{A}$, and the Frobenius action, which live over the special fiber. To do this, we'll study the p-divisible group associated to $\mathcal{A}$.

2. p-divisible Groups

Good references include Tate's original paper and notes by Brandon and Mike from last year.

Let A/k be an Abelian variety. If ℓ is prime to the characteristic of k, the scheme $A[\ell]$ is finite étale, and thus the Galois representation $A[\ell](\bar{k})$ contains essentially all of the information of $A[\ell]$ (in the sense that $A[\ell]$ can be recovered from $A[\ell](\bar{k})$ with the Galois action). Thus we may view the Tate module $T_\ell(A)$ as either the "pro-object"

$$\cdots \to A[\ell^n] \xrightarrow{\ell} A[\ell^{n-1}] \to \cdots$$

or as the inverse limit

$$\varprojlim A[\ell^n](\bar{k})$$

in the category of group objects in $\mathrm{Gal}(\bar{k}/k)$-sets. Indeed, there is an equivalence of categories between finite étale group schemes over $\mathrm{Spec}(k)$ and group objects in finite $\mathrm{Gal}(\bar{k}/k)$ sets, and similarly with inverse systems (though I'll state this formally in a second).

If p is not prime to the characteristic of k, the group schemes $A[p]$ need not be étale, though they are flat by the fibral flatness criterion. So we cannot forget the scheme structure on $A[p]$ when we study it. This motivates the following definition.

Definition 1. *A p-divisible group of height h over S is an inductive system $G = (\{G_v\}_{v \in \mathbb{N}}, G_v \to G_{v+1})$ satisfying:*

- *Each G_v is a finite flat commutative group scheme over S of order p^{hv}.*
- $$0 \to G_v \xrightarrow{i_v} G_{v+1} \xrightarrow{p^v} G_{v+1}$$
  is exact for each v.

A morphism of p-divisible groups is simply a map of inductive systems in the category of group schemes; kernels, cokernels etc. are defined at finite level.

Of course, the prototypical example we are interested is the case where A/S is an Abelian scheme of relative dimension g; then setting $G_v = A[p^v]$ with the obvious inclusion maps gives a p-divisible group of height $2g$ over S (finiteness can be checked on geometric fibers, and flatness follows by fibral flatness). Other examples include the constant p-divisible group of height h

$$(\mathbb{Z}/p^v\mathbb{Z})^h \to (\mathbb{Z}/p^{v+1}\mathbb{Z})^h$$

or the p^k torsion groups of a torus.

In the case where each G_v is étale, the theory is somewhat boring. In particular, after choosing a geometric point of S, there is an equivalence of categories between étale p-divisible groups over a connected S and free $\mathbb{Z}_p$-modules with a $\pi_1(S)$-action, carrying height to $\mathbb{Z}_p$ rank. We've already remarked on this in the case of the Tate module for ℓ prime to the characteristic. We'll call the $\mathbb{Z}_p$-module obtained in this case the p-adic Tate module of a p-divible group G, and denote it $T_p(G)$.

Now let $S = \mathrm{Spec}(R)$, where R is Henselian. Recall that in this case, there is a short exact sequence

$$0 \to G^0 \to G \to G^{et} \to 0$$

where G^0 is terminal among connected finite flat group schemes over S mapping to G and G^{et} is initial for maps $G \to H$ with H finite étale over S. (We need Henselian-ness to construct G^0.) The functors $G \mapsto G^0, G \mapsto G^{et}$ are exact functors on the category of finite flat group schemes over S, so applying them levelwise gives for each p-divisible group G a short exact sequence

$$0 \to G^0 \to G \to G^{et} \to 0,$$

allowing us to separate the "interesting" connected stuff from the étale stuff.

Let's discuss the p-divisible group associated to a CM Abelian scheme $\mathcal{A}$ over $\mathcal{O}_{K_w}$, where the CM is by L (with $\mathcal{O}_L$ acting honestly). $\mathcal{A}[p^n]$ is a finite flat group scheme over $\mathcal{O}_{K_w}$, with an action by $\mathcal{O}_L$ which factors through

$$\mathcal{O}_L/p^n \simeq \prod_{v|p} \mathcal{O}_{L_v}/p^n.$$

Decomposing by idempotents gives a functorial decomposition

$$\mathcal{A}[p^n] = \prod_{v|p} G_{v,n}$$

where each $G_{w,n}$ has an action by $\mathcal{O}_{L_v}/p^n$. We claim first that each $G_{v,n}$ is flat. Indeed, the following lemma, due to Serre, gives the proof:

Lemma 1. *Let A and B be C-algebras with $A \otimes_C B$ flat over C, and such that the map $C \to A$ has a section. Then B is flat over C.*

Proof. By assumption $A = C \oplus J$, so $A \otimes_C B = B \oplus (J \otimes_C B)$, whence each summand is flat. □

In particular, each $G_{v,n}$ admits a section and the product $\mathcal{A}[p^n]$ is flat, so each $G_{w,n}$ is flat. As the generic fiber of $G_{v,n}$ is étale, it is free of rank 1 over $\mathcal{O}_{L_v}/p^n$ on geometric points. Thus $G_{v,n}$ is a subscheme of $G_{v,n+1}$.

A priori the sequence

$$0 \to G_{v,1} \to G_{v,n+1} \xrightarrow{p} G_{v,n}$$

is only left exact, but examining orders on the fiber over the closed point gives right exactness. Thus $G_v := \{G_{v,n}\}_{n\geq 1}$ is a p-divisible group with height $\#H_v = [L_v : \mathbb{Q}_p]$. Indeed, note that the generic fiber is an étale p-divisible group, whose "Tate module" is free of rank one over $\mathcal{O}_{L_v}$.

3. The Serre-Tate Theorem and the Proof

The main idea we'll use to study these G_v is the following beautiful theorem of Serre and Tate. Strictly speaking we won't use this theorem (though I might say something about the proof if I have time), but the ideas in it will strongly influence the proof of the Shimura-Taniyama formula.

Theorem 2 (Serre-Tate)**.** *Let $(R, \mathfrak{m})$ be a complete local Noetherian ring with residue characteristic p.*

(1) *Let G be a connected p-divisible group over R with height h. Then the augmented topological R-algebra*

$$\mathcal{O}_G := \varprojlim \mathcal{O}_{G_n}$$

is topologically isomorphic to the R-algebra $R[[X_1, ..., X_d]]$ equipped with the $(\mathfrak{m}, \underline{X})$-adic topology, and so is a commutative formal Lie group for which $[p]$ is finite flat of degree p^h. Furthermore, $\mathcal{O}_{G_n}$ is the quotient ring corrsesponding to the p^n-torsion, and these quotients are a cofinal system of $(\mathfrak{m}, \underline{X})$-adic quotients of $R[[\underline{X}]]$.

(2) *If $\mathcal{R} = R[[\underline{X}]]$ is a commutative formal group law so that the maps $[p]^* : \mathcal{R} \to \mathcal{R}$ are finite flat of degree p^h, then the kernels $\mathcal{G}_n = \ker([p^n]_{\mathrm{Spf}(\mathcal{R})})$ form a p-divisible group over R that recovers $\mathcal{R}$ via (1).*

The number d is called the dimension of G. A key point is the following easy result:

Proposition 1. *If $G = \{\mathcal{A}[p^n]^0\}$, then the recipe in (1) recovers the formal group law on the complete ring $\hat{\mathcal{O}_{A,0}}$ along the identity section, with the formal group law inherited from A, whence $d = \dim(A)$ in this case.*

Proof. The question is supported at the identity section, so we must simply show that the ideals of $\ker[p^k]^* \subset R[[X_1, ..., X_d]] = \hat{\mathcal{O}_{A,0}}$ are $(\mathfrak{m}, \underline{X})$-adically cofinal. But

$$[p]^*(X_j) = pX_j + \text{higher order terms}$$

so this is obvious. □

Thus we will think of the p-divisible group $\{\mathcal{A}[p^n]^0\}$ as "the same" as the formal lie group $\hat{\mathcal{A}_0}$ over $\mathcal{O}_{K_w}$. Let $\hat{G_v}$ be the formal group constructed from $\{G^0_{v,n}\}$ by the same recipe; we have that

$$\prod_{v|q} \hat{G_v} = \hat{\mathcal{A}_0}.$$

One can see from this (without Serre-Tate) that the same construction gives formal Lie groups associated to the G_v's; the content is in showing that the rings are power series rings over $\mathcal{O}_{K_w}$, but this follows by e.g. the same trick as in Serre's trick.

In particular, as $\dim(\mathcal{A}) = \#\Phi$, and $\prod_v G_v = \{\mathcal{A}[p^n]\}$,

$$\sum_{v|q} \dim \hat{G_v} = \#\Phi.$$

Theorem 3.

$$\dim \hat{G_v} = \#\Phi_v.$$

Proof. As

$$\prod_{v|q} T(\hat{G_v}) = T(\hat{\mathcal{A}_0})$$

we see that $T(\hat{G_v})$ is the $\mathcal{O}_{L_v} \otimes_{\mathbb{Z}_p} \mathcal{O}_{K_w}$ component of $T(\hat{\mathcal{A}_0})$ as an $\mathcal{O}_L \otimes \mathcal{O}_{K_w}$ module (e.g. by the functorial description of tangent spaces). We can compute the dimension at the generic fiber of $\mathrm{Spec}(\mathcal{O}_{K_w})$, and indeed after base-changing the generic fiber to e.g. $\overline{\mathbb{Q}_p}$. Thus we wish to compute the dimension of the subspace of $T_0(A_{\overline{\mathbb{Q}_p}})$ where L acts via L_v. But $T_0(A_{\overline{\mathbb{Q}_p}})$ is a module over $L \otimes_{\mathbb{Q}} A_{\overline{\mathbb{Q}_p}}$, where the action is according to the CM type Φ by the definition of Φ; that is, it is precisely

$$T_0(A_{\overline{\mathbb{Q}_p}}) \simeq \prod_{\phi \in \Phi} (\overline{\mathbb{Q}_\ell})_\phi$$

with the action on the ϕ component via the embedding $\phi : L \to \overline{\mathbb{Q}_\ell}$. Thus the dimension of the w component is precisely $\#\Phi \cap H_v = \#\Phi_v$ as desired. □

Let us summarize what we've gotten so far. We've defined p-divisble groups G_v, whose product is the p-divisible group $\{\mathcal{A}[p^n]\}$. The G_w have height $\#H_v$ and dimension $\#\Phi_v$. Finally, $\mathcal{O}_{L_v}$ acts on G_v, with $\pi \in \mathcal{O}_{L_v}$ acting via Frobenius on the closed fiber.

In this situation, we may using Serre-Tate prove a general Shimura-Taniyama formula for p-divisible groups with CM. I will only prove this in the case of the p^n torsion for an Abelian variety (so we will not

need Serre-Tate), but an essentially identical proof works in general after the initial input from Serre-Tate. Here is the statement:

Theorem 4. *Let Γ be a p-divisible group over $\mathcal{O}_{K_w}$ of dimension d and height $h > 0$. Assume that $\mathcal{O}_{L_v}$ acts on Γ with $[L_v : \mathbb{Q}_p] = h$, so that the p-adic Tate module of Γ over the generic fiber (defined as the inverse limit of the level-wise $\overline{K_w}$-points) is a rank one $\mathcal{O}_{L_v}$-module.*

If there is an element $\pi \in \mathcal{O}_{L_v} - \{0\}$ lifting the Frobenius action on the closed fiber, then $d/h = \mathrm{ord}_v(\pi)/\mathrm{ord}_v(q)$, where q is the cardinality of the residue field k of $\mathcal{O}_F$.

Proof. (Proof in the case $\Gamma = G_v$) An equivalent statement of the theorem is the statement that π^h/q^d is a unit.

First, consider the case where G_v is étale. Then $d = 0$ by the definition of dimension, and Frob acts as an automorphism. In particular, π^r cannot be a unit multiple of p, as the action of p has non-trivial kernel. Thus π is thus a unit in $\mathcal{O}_{L_v}$. Thus in this case the ratio $\pi^h/q^d = \pi^h$ is a unit and we win.

Now consider the connected étale sequence

$$0 \to G_v^0 \to G_v \to G_v^{\text{ét}} \to 0.$$

Passing to the generic fiber and taking the inverse limit over geometric points, we get a short exact sequence

$$0 \to T_p((G_v)^0_{K_w}) \to T_p((G_v)_{K_w}) \to T_p((G_v)^{\text{ét}}_{K_w}) \to 0.$$

As remarked earlier, the middle term is a rank 1 module over $\mathcal{O}_{L_v}$, so either $G_v^0 = 0$ or $G_v^{\text{ét}} = 0$. In the former case we're done by the previous remarks, so only the connected case remains.

The (in my opinion awesome) idea is to identify the normalized absolute value on $\mathcal{O}_{L_v}$ with the *degree* of the induced map on G_v (make sure you know what the degree of this map means—it's a little subtle in the case of a general p-divisible group, but less so when the p-divisible group comes from an abelian variety). Indeed, the map $[p]$ on G_v is finite flat of order $p^h = ||p||_{L_v}$. The map $x \mapsto \deg([x]_{G_v})$ is multiplicative and sends $\mathcal{O}_{L_v} - \{0\}$ to $p^{\mathbb{Z}}$. It sends units to 1 and p to $||p||_{L_v}$, so it agrees with the normalized absolute value as desired.

So all that remains is to show that π^h and q^d have the same degree. Of course $\deg[p]_{G_v} = p^h$ by the definition of the height of a p-divisible group, so $\deg[q^d]_{G_v} = q^{dh}$. So we want $\deg[\pi]_{G_v} = q^d$. But again by flatness we may compute the degree on the closed fiber, whence π acts via the Frobenius. But this manifestly has degree q^d as desired (because this is the degree of the q-th power map on $k[[Y_1, ..., Y_d]]$). □