GALOIS REPRESENTATIONS AND ELLIPTIC CURVES

FALL 2011 JUNIOR PAPER

ABSTRACT. An elliptic curve over a field K is a projective nonsingular genus 1 curve E over K along with a chosen K-rational point O of E, which automatically becomes an algebraic group with identity O. If K has characteristic 0, the n-torsion of E, denoted $E[n]$, is isomorphic to $(\mathbb{Z}/n\mathbb{Z})^2$ over $\overline{K}$. The absolute Galois group G_K acts on these points as a group automorphism, hence it acts on the inverse limit $\varprojlim_n E[n]$, which is isomorphic to $\prod_\ell T_\ell$, where $T_\ell = \varprojlim E[\ell^k]$ is the Tate module. The Tate module is isomorphic to $\mathbb{Z}_\ell^2$ as an abelian group, and since the Galois group acts on finite quotients, it acts continuously, and we get a continuous homomorphism $\rho_{E,\ell} : G_K \to \mathrm{GL}_2(\mathbb{Z}_\ell)$. This paper presents of a proof of Serre's open image theorem, which states that the image is an open subgroup. The bulk of the proof involves facts about ℓ-adic Galois representations constructed from global class field theory.

The paper assumes standard results in algebraic number theory, algebra, and Lie group theory, as well as the basics of the arithmetic of elliptic curves. The paper quotes some results coming from p-adic Hodge theory of étale cohomology in the section on Hodge-Tate decompositions.

Most of the material is from [Ser89], the original source of this theorem.

0. Notation

We summarize some of the notation defined in the text below.

For a number field K, we let Σ_K denote the set of finite places of K, Σ_K^∞ the set of infinite places, and $\overline{\Sigma_K}$ the set of all places. For each place v of K, we let K_v denote the completion at v, $\mathcal{O}_v$ its valuation ring, and U_v the group of units. For a prime p (or ℓ), we thus let U_p denote the p-adic units. We write valuations additively, always setting the valuation of a uniformizer to be 1. When we refer to an *absolute value* $||_v$, we are working multiplicatively, with the standard normalization relative to the given field. We let I_K denote the group of idèles of K, and if S is a finite set of places of K, we let $\mathbf{A}_K^{S,\times}$ denote the idèles supported outside S.

All fields, unless otherwise noted, have characteristic 0. If L/K is a Galois extension of fields, we let $G(L/K)$ denote its Galois group. If both are number fields, with w a place of L lying over a place v of K, we denote by $I(w/v) \subseteq D(w/v) = G(L_w/K_v) \subseteq G(L/K)$ the decomposition and inertia groups, respectively, of w. If K is a field, then $\overline{K}$ denotes its algebraic closure (which is the same as the algebraic closure when K is perfect, the case we mostly consider).

1. Background on p-adic Lie Groups

Just for completeness, we include a quick summary of definitions and facts about p-adic Lie groups, as this material is not standard. Most of the theorems and proofs are similar to those for real Lie groups.

Date: David Corwin.

1.1. Manifolds and Lie Groups.

Definition 1.1. Let U be an open subset of $\mathbb{Q}_p^n$. An *analytic function* from U to $\mathbb{Q}_p$ is a continuous map locally given by convergent power series. We say a function to $\mathbb{Q}_p^m$ is analytic if each of its component functions are.

Definition 1.2. A (analytic) p-adic manifold of dimension n is a topological space with an atlas of open subsets homeomorphic to open subsets of $\mathbb{Q}_p^n$ such that the transition maps on the overlap are analytic.

Definition 1.3. An analytic map between p-adic manifolds is defined as a continuous map whose restriction to any pair of maps in the atlas is analytic.

Example 1.4. Any open subset of a p-adic manifold is a p-adic manifold. Thus $\mathrm{GL}_n(\mathbb{Q}_p)$ and $\mathbb{Z}_p^n$ are p-adic manifolds.

Definition 1.5. A p-adic Lie group is a group and p-adic manifold such that the multiplication and inverse maps are analytic.

Fact 1.6. Any algebraic group over $\mathbb{Q}_p$ is a p-adic Lie group, as polynomial maps are analytic.

1.2. Facts about Lie Groups.

Fact 1.7. A continuous homomorphism between p-adic Lie groups is analytic.

Theorem 1.8. *[P-adic Cartan's Theorem] A closed subgroup of a p-adic Lie group is a Lie subgroup.*

Proof. See [Hoo42]. □

1.3. Representations of Lie Groups.

Definition 1.9. The tangent space $T_P(M)$ to an n-dimensional p-adic manifold M at a point P to be equivalence classes of analytic maps from $\mathbb{Z}_p^n$ to the manifold sending 0 to P, with two maps equivalent if, in some coordinate chart, their linear terms are equal.

If $f : M \to N$ is an analytic map of p-adic manifolds, we have an induced map $d_P f : T_p(M) \to T_{f(P)}(N)$. In particular, if a p-adic Lie group is embedded in $\mathrm{GL}_n(\mathbb{Q}_p)$, then its tangent space at the identity is embedded in the tangent space of $\mathrm{GL}_n(\mathbb{Q}_p)$ at the identity element. As in the real case, this latter vector space is naturally identified with $M_n(\mathbb{Q}_p) = \mathrm{End}(\mathbb{Q}_p^n)$.

The *Lie bracket* of $x, y \in M_n(\mathbb{Q}_p)$ is defined to be $[x, y] = xy - yx$. If $G \subseteq \mathrm{GL}_n(\mathbb{Q}_p)$ is an embedded Lie subgroup, then $T_e(X)$ can be shown to be closed under the Lie bracket operation. Furthermore, this p-adic Lie algebra is an invariant of the Lie group G (i.e., it does not depend on the embedding). We also note that one can define the matrix exponential as in the real case using the same series, which converges for sufficiently small matrices. Then a matrix $x \in \mathfrak{gl}_n(\mathbb{Q}_p)$ is in the Lie algebra of G iff $exp(\lambda x) \in G$ for some $\lambda \in \mathbb{Q}_p$. Thus, the Lie algebra of an open subgroup of a Lie group is the same by continuity of the exponential. We note:

Fact 1.10. The image of the exponential map in G contains an open neighborhood of the identity.

We now discuss representations of p-adic Lie groups.

Definition 1.11. A representation of a p-adic Lie group G is a continuous (hence analytic) homomorphism $G \to \mathrm{GL}(V)$, where V is a finite-dimensional vector space over $\mathbb{Q}_p$.

From a representation $G \to \mathrm{GL}(V)$ of a p-adic Lie group, where V is a vector space over $\mathbb{Q}_p$, we get a representation of the Lie algebra of G, denoted $\mathrm{Lie}(G)$. In distinction to the real case, we have:

Fact 1.12. A morphism of the associated representations of Lie algebras is the same as a morphism of representations of an open subgroup of the Lie group.

A corollary of this fact is:

Corollary 1.13. *A representation of G is abelian, meaning the image of G is abelian, for an open subgroup iff the associated representation of the Lie algebra is abelian, meaning that its image in $\mathfrak{gl}(V)$ is an abelian Lie algebra.*

2. ℓ-adic Galois Representations

We denote by G_K the Galois group of $\overline{K}/K$, where $\overline{K}$ is the algebraic closure of K. We assume that K has characteristic 0.

Definition 2.1. An ℓ-adic ρ representation of a field K is a continuous homomorphism $\rho : G_K \to \mathrm{GL}_n(\mathbb{Q}_l)$ (equivalently, into $\mathrm{Aut}(V)$ for some finite-dimensional vector space over $\mathbb{Q}_l$).

Whenever we refer to a representation of G_K on a vector space V as ℓ-adic, we are assuming V is a vector space over $\mathbb{Q}_\ell$, and we identify $\mathrm{GL}(V)$ with $\mathrm{GL}_n\mathbb{Q}_\ell$ (non-canonically).

Definition 2.2. An ℓ-adic representation ρ of K is said to be *abelian* if $\rho(G_K)$ is abelian.

Proposition 2.3. *The image $G_\ell = \rho(G_K)$ is an ℓ-adic Lie group.*

Proof. As G_K is compact, so its its image, and as $\mathrm{GL}_n(\mathbb{Q}_l)$ is Hausdorff, the image is closed. Thus, by Theorem 1.8, the image is an ℓ-adic Lie subgroup. □

Note that if we restrict the representation to G_L, where L/K is finite Galois, then we restrict to a finite index closed subgroup, hence $\rho(G_L)$ is finite index and compact, hence open in $\rho(G_K)$. This implies that its Lie algebra is the same.

Example 2.4. Assume $\mathrm{char}(K) \neq \ell$. For a prime ℓ, let μ_{ℓ^n} denote the ℓ^n roots of unity in $\overline{K}$. Then G_K acts continuously (through a finite quotient) on μ_{ℓ^n}. We thus have a continuous action of G_K on $T_\ell(\mu) := \varprojlim(\mu_{\ell^n})$, which is isomorphic to $\mathbb{Z}_\ell$. We can tensor with $\mathbb{Q}_\ell$ to get a continuous representation of G_K on $\mathbb{Q}_\ell$, so we have a continuous homomorphism $G_K \to U_\ell = \mathrm{GL}_1(\mathbb{Z}_\ell) = \mathbb{Z}_\ell^\times \subseteq \mathbb{Q}_\ell^\times$. This homomorphism is denoted χ_ℓ and is called the *cyclotomic character of K.*

2.1. **Representations of Number Fields.**

We have reduced the proof of the theorem to the abelian case, but we have a longer way to go. Now that we are considering abelian representations of a number field, we may use the tools of class field theory to analyze ℓ-adic Galois representations in the abelian case. Before we begin, we discuss some general properties of ℓ-adic representations of number fields.

2.1.1. *Decomposition, Inertia, and Frobenius.*

We review some standard facts about decomposition groups, inertia groups, and Frobenius elements, which can be found in any standard algebraic number theory textbook. We note how these extend to infinite extensions.

If K is a number field, we let Σ_K denote the set of finite places (or primes) of K, Σ_K^∞ the infinite ones, and $\overline{\Sigma_K} = \Sigma_K \cup \Sigma_K^\infty$.

If L/K is a Galois extension of number fields, and w is a place of L over v, then we define $I(w/v) \subseteq D(w/v) \subseteq G(L/K)$ to be the inertia, decomposition, and Galois groups. We let $e(w/v)$ and $f(w/v)$ denote the ramification index and residue degree of w over v. If we let K_v denote the completion of K at v, then $D(w/v) = G(L_w/K_v)$. Furthermore, if we have a Galois extension L'/K containing L and a prime w' of L lying over v, then we have the following commutative diagram with exact rows:

$$\begin{tikzcd}
1 \arrow[r] & I(w'/w) \arrow[r] \arrow[d, hook] & I(w'/v) \arrow[r] \arrow[d, hook] & I(w/v) \arrow[r] \arrow[d, hook] & 1 \\
1 \arrow[r] & D(w'/w) \arrow[r] & D(w'/v) \arrow[r] & D(w/v) \arrow[r] & 1
\end{tikzcd}$$

If w'' is a different prime of L' lying over v (or w), then the associated decomposition and inertia groups are conjugate. If $I(w/v) = 1$, i.e. w is unramified over v, then we define the Frobenius element $F_{w/v}$ to be the element of $D(w/v)$ corresponding to the Frobenius element in the extension of residue fields. Note that in the case above, $F_{w'/v}$ maps onto $F_{w/v}$, and $F_{w'/w} = F_{w'/v}^{f(w/v)}$. If the extension is abelian, or if we only care about conjugacy classes, we write simply F_v.

If L'/K is an infinite Galois extension, then we define $\overline{\Sigma}_{L'}$ to the projective limit of the sets $\overline{\Sigma}_L$ for finite Galois extensions L/K, and for $w' \in \overline{\Sigma}_{L'}$, we define $I(w'/v)$ and $D(w'/v)$ to be the projective limit of the groups $I(w/v)$ and $D(w/v)$, respectively, as w ranges over primes lying between w' and v. All of the above diagrams, equalities, and isomorphisms hold for infinite Galois extensions of number fields just as well as for finite ones.

Fact 2.5 (Cheboratev Density Theorem)**.** If L/K is a Galois extension of number fields, then for each conjugacy class σ of $\mathrm{Gal}(L/K)$, the density of primes v of K such that $F_v = \sigma$ is $\dfrac{|\sigma|}{|\mathrm{Gal}(L/K)|}$. In particular, there are infinitely many such primes.

Corollary 2.6. *If $S \subseteq \Sigma_K$ is a finite subset and L/K unramified outside S, then the conjugacy classes $F_v \in G(L/K)$ are dense.*

Proof. Any open subset of $G(L/K)$ contains a coset of an open subset of finite index, which corresponds to an element of $G(M/K)$ for a finite Galois extension M/K. But there exists a prime in $\Sigma_K \setminus S$ whose Frobenius element is equal to that by Chebotarev's Density Theorem, so we are done. □

2.2. **Rational and Compatible Representations.** In all that follows, K denotes a number field.

Definition 2.7. If $\rho : G_K \to \mathrm{GL}(V)$ is an ℓ-adic representation of a number field, we say that the representation is *unramified* at v if $\rho(I(w/v)) = 1$ for all places w of $\overline{K}$ lying over v. If this is the case, we can define $\rho(F_v) \in \mathrm{GL}(V)$ (which is a conjugacy class). We denote this by $F_{v,\rho}$.

Definition 2.8. Let G be an affine algebraic group and $K[G]$ its affine coordinate ring over K. An element of $K[G]$ is said to be *central* if $f(xy) = f(yx)$ for any K-algebra K' and all $x, y \in G(K')$. Equivalently, it lies in the ring of invariants of the action of G on itself by conjugation.

Definition 2.9. The conjugacy class of an element $x \in G(K')$ for a K-algebra K' is said to be *rational* over K if $f(x) \in K$ for any central $f \in K[G]$. It is said to be *integral* if all these values are algebraic integers.

Definition 2.10. A representation $\rho : G_K \to \mathrm{GL}_n(\mathbb{Q}_l)$ whose image lies in an algebraic subgroup H defined over $\mathbb{Q}$ is said to be *rational with values in H* if ρ is unramified at v, and the conjugacy class $\rho(F_v)$ is rational over $\mathbb{Q}$ in H for almost all $v \in \Sigma_K$ (note we only care about the finite primes). The set of v for which this does not hold is called the *support* of ρ.

Lemma 2.11. *The central elements of* GL_n *consist of the ring generated by the coefficients of the characteristic polynomial.*

Proof. The proof in one direction follows from the well-known fact that the characteristic polynomial is invariant under conjugation. Next, note that the coefficients of the characteristic polynomial are ($\pm$) the symmetric functions on the eigenvalues. If an element is central, then it is central on $G(\overline{K})$ (as $\overline{K}$ is certainly a K-algebra), and it is known that the converse is true. Thus the value of a central element on a diagonalizable matrix is determined by its restriction to the set of diagonal matrices, and because conjugation can be used to interchange entries of a diagonal matrix, such an element must be a symmetric function on the eigenvalues of the symmetric polynomial. Thus any central element coincides with a polynomial in the coefficients of the characteristic polynomial. It remains to show that the diagonalizable elements are dense in the Zariski topology. This follows because all matrices whose characteristic polynomial has nonzero discriminant are diagonalizable, and the set of such matrices is clearly dense in the Zariski topology because GL_n is irreducible. □

Example 2.12. As the central elements of $\mathbb{Q}[\mathrm{GL}_n]$ are the coefficients of the characteristic polynomial, a representation is rational if the characteristic polynomial of $F_{v,\rho}$, denoted $P_{v,\rho}(T)$, has rational coefficients for almost all v. We refer to such a representation simply as *rational.*

Definition 2.13. Let G be an algebraic group over $\mathbb{Q}$, and let ρ and ρ' be l- and l'-adic representations in $G(\mathbb{Q}_l)$ and $G(\mathbb{Q}_{l'})$ respectively. Suppose that for $v \in \Sigma_K$ outside a finite set of places of K (called the *support of compatibility*), we have $f(F_{v,\rho}) = f(F_{v,\rho'})$ for all central $f \in K[G]$ (note that f is well-defined on conjugacy classes as it is central). Then we say that the representations ρ and ρ' are compatible as representations with values in G. (As before, when G is not specified, we mean GL_n, and we refer to them simply as compatible.)

For ℓ a prime of $\mathbb{Q}$, let S_ℓ be the set of primes of K lying over ℓ.

Definition 2.14. Suppose, for each ℓ, we are given an ℓ-adic representation ρ_ℓ of K. This collection $(\rho_\ell)_\ell$ is called a *system of ℓ-adic representations.* The system is said to be *compatible* if for any two primes ℓ, ℓ', the representations $\rho_\ell, \rho_{\ell'}$ are compatible. The system is said to be *strictly compatible* with *exceptional set* S if there exists a set S of primes of K such that ρ_ℓ is rational with support contained in $S \cup S_\ell$, and ρ_ℓ and $\rho_{\ell'}$ are compatible with support of compatibility contained in $S \cup S_\ell \cup S_{\ell'}$.

Example 2.15. We recall the ℓ-adic representation of G_K on $T_\ell(\mu)$, previously defined. We note that the extension $K(\mu_{\ell^n})/K$ is ramified only at primes lying above ℓ, meaning the action of inertia is trivial for all primes $v \in S_{\ell'}$ for $\ell' \neq \ell$, hence the representation is unramified outside S_ℓ. Furthermore, if $v \notin S_\ell$, then F_v acts by sending a root of unity to its $\mathrm{N}_{K/\mathbb{Q}}(v)$ power, i.e. $F_{v,\rho} = \mathrm{N}_{K/\mathbb{Q}} \in \mathbb{Q}^\times \subseteq \mathbb{Q}_\ell^\times$. The characteristic polynomial is thus rational, even integral, and it is independent of ℓ. Thus we have a strictly compatible system of integral representations with empty exceptional set.

3. Elliptic Curves

As usual, we suppose E is an elliptic curve defined over a field K, and we let $E[n]$ denote the set of n-torsion points of E over in $E(\overline{K})$. Furthermore, G_K acts as group homomorphisms on $E(\overline{K})$.

We say that E has no complex multiplication over K if $\mathrm{End}_K(E) = \mathbb{Z}$.

3.1. **Siegel's Theorem.** We wish to prove Shafarevich's Theorem, which is a necessary ingredient in Serre's open image theorem. In order to prove that, we must prove Siegel's Theorem. Before that, we begin with a couple definitions and then a lemma.

Definition 3.1. If $S \subseteq \overline{\Sigma_K}$ is a finite set of places containing Σ_K^∞, then $\mathcal{O}_{K,S}$ is the subring of elements of K that are integral outside S. The units of this ring, denoted $\mathcal{O}_{K,S}^\times$, are known as the *S-units.*

Definition 3.2. If $v \in \Sigma_K$ lies over a prime ℓ of $\mathbb{Q}$, we define $n_v = e(v/\ell)f(v/\ell) = [K_v : \mathbb{Q}_\ell]$.

Definition 3.3. If $x \in K^\times$, we define the *height* $H_K(x)$ to be

$$\prod_{v \in \overline{\Sigma_K}} \max(1, |x|_v^{n_v}).$$

Note that for almost all v, the factor is 1, so this is well-defined.

Remark 3.4. If we view x as the point with homogenous coordinates $[x, 1] \in \mathbb{P}^1(K)$, then H_K is the standard height on the projective line.

We first prove the following lemma about H_K:

Lemma 3.5. *For $C > 0$, there are finitely many $x \in \mathcal{O}_{K,S}^\times$ such that $H_K(x) \leq C$.*

Proof. Let $U_{v,C}$ denote the set of $y \in K_v^\times$ for which $|y|_v \leq C$. Then such x lie in $\prod_{v \in \overline{\Sigma}_K} U_{v,C} \subseteq I_K$, where I_K denotes the idèles of K. But this product is compact, and we know that $K^\times$ is discrete in I_K (see, e.g. [Cas67], Chapter II), so there are finitely many such x. □

Before we begin, we quote a theorem from algebraic number theory that will be ingredients in our proof:

Theorem 3.6 (Roth's Theorem)**.** *For a number field K and any $v \in \overline{\Sigma_K}$, $\alpha \in \overline{K}$, and $C > 0$, there are finitely many $x \in K$ such that*

$$|x - \alpha|_v \leq C H_K(x)^{-3}.$$

We are now ready to prove the following lemma:

Lemma 3.7. *If $a, b \in K^\times$, then the equation*

$$ax + by = 1$$

has finitely many solutions $x, y \in \mathcal{O}_{K,S}^\times$.

Proof. Choose $m > 3[K : \mathbb{Q}]\#S$. By Dirichlet's S-Unit Theorem, the group $(\mathcal{O}_{K,S}^\times)/(\mathcal{O}_{K,S}^\times)^m$ is finitely-generated, so we choose a set of coset representatives $c_1, \cdots, c_r$. Suppose the equation has infinitely many solutions in $\mathcal{O}_{K,S}^\times$. Then for some i, j, we find that $ac_iX^m + bc_jY^m = 1$ has infinitely many solutions for $X, Y \in \mathcal{O}_{K,S}^\times$. We thus replace a by ac_i and b by bc_j and continue.

Since Y is supported on a finite set S, we have that, for some v, there are infinitely many solutions for which $|Y|_v \geq |Y|_w$ for all $v \neq w$. It follows that $|Y|_v \geq H_K(Y)^{\frac{1}{[K:\mathbb{Q}]\#S}}$ for all such solutions.

For any $t > 0$, there are infinitely many solutions with $|Y|_v^{[K:\mathbb{Q}]\#S} \geq H_K(Y) > t$ by the finiteness result in the previous lemma.

Suppose we have $g^m = -\dfrac{b}{a}$. .Then we have the factorization $\displaystyle\prod_{\zeta\in\mu(m)}(X-\zeta gY) = \frac{1}{a}$, which gives us

$$\prod_{\zeta\in\mu(m)}\left|\frac{X}{Y}-\zeta g\right|_v = \frac{1}{|aY^m|_v}.$$

We choose $s < |g(1-\zeta)|_v$, we can choose t sufficiently large that $|\dfrac{X}{Y}-g|_v < s$, for some mth root g, and we always pick that root for every (X,Y). Since there are finitely many choices for g, we can assume that we have an infinite number of solutions that all have the same g. Note that, for $\zeta\in\mu(m)$, we have

$$\left|\frac{X}{Y}-\zeta g\right|_v > |g(1-\zeta)|_v - s.$$

It follows that there is a constant

$$c = \frac{1}{|a|_v\prod_{\zeta\in\mu(m)}\min(1,|g(1-\zeta)|_v)}$$

such that

$$\left|\frac{X}{Y}-g\right|_v \le c|Y|_v^{-m}.$$

We also know from

$$\left(\frac{X}{Y}\right)^m = \frac{1}{aY^m} - g^m$$

that for all w, we have

$$\left|\frac{X}{Y}\right|_w^m \quad\le\quad \frac{1}{|a|_w}\left|\frac{1}{Y}\right|_w^m + |g|_w^m$$

It follows that

$$\begin{aligned}
H_K\left(\frac{X}{Y}\right)^m &= \prod_{w\in S}\max\left(1,\left|\frac{X}{Y}\right|_w^{mn_w}\right)\\
&\le \prod_{w\in S}\max\left(1,\frac{1}{|a|_w}\left|\frac{1}{Y}\right|_w^m + |g|_w^m\right)^{n_w}\\
&\le \prod_{w\in S}\max\left(1,\left|\frac{1}{Y^m}\right|_w(|a|_w+|Yg|_w^m)\right)^{n_w}\\
&\le \prod_{w\in S}\max\left(1,\left|\frac{1}{Y^m}\right|_w(|a|_w+t\,|g|_w^m)\right)^{n_w}\\
&\le d^{[K:\mathbb{Q}]}\prod_{w\in S}\max\left(1,\left|\frac{1}{Y^m}\right|_w\right)^{n_w}\\
&\le d^{[K:\mathbb{Q}]}H_K(Y)^m,
\end{aligned}$$

where $d = \max\limits_{w\in S}((|a|_w+t|g|_w^m)\,,1)$.

It follows that

$$\left|\frac{X}{Y}-g\right|_v \le c|Y|_v^{-m} \le cH_K(Y)^{-\frac{m}{[K:\mathbb{Q}]\#S}} \le \frac{c}{d}H_K\left(\frac{X}{Y}\right)^{-\frac{m}{[K:\mathbb{Q}]\#S}}.$$

By Roth's Theorem, there are only finitely many such $\frac{X}{Y}$ as $m > 3[K : \mathbb{Q}]\#S$. But for each $\frac{X}{Y}$, the equation $aX^m + bY^m = 1$ gives only finitely many possible (X, Y). This contradiction proves the lemma. □

Now we can prove Siegel's Theorem:

Theorem 3.8 (Siegel's Theorem)**.** *Suppose $f(x) \in K[x]$ has degree $d \geq 3$. Then $y^2 = f(x)$ has finitely many solutiosn $x, y \in \mathcal{O}_{K,S}^\times$.*

Proof. We enlarge K so that $f(x)$ splits in K as $y^2 = f(x) = \alpha(x - a_1) \cdots (x - a_d)$. We enlarge S so that $\alpha \in \mathcal{O}_{K,S}^\times$, $a_i - a_j \in \mathcal{O}_{K,S}^\times$ for $i \neq j$, and $\mathcal{O}_{K,S}$ is a principal ideal domain (PID). To do the latter, choose a collection T of finite primes that generate the ideal class group, and choose a uniformizer for each prime in T. Then, include all elements of the support of all of these uniformizers in S, and we are done.

If $\mathfrak{p} \notin S$, then $\mathfrak{p}$ divides at most one of $x - a_i$, hence it divides $x - a_i$ with even order. As $\mathcal{O}_{K,S}$ is a PID, we can write $x - a_i = \beta_i z_i^2$ with $\beta_i \in \mathcal{O}_{K,S}^\times$. Now let L be obtained by adjoining the square root of every element of $\mathcal{O}_{K,S}^\times$ (we know that L/K is finite by Dirichlet's S-Unit Theorem). Let T be the set of primes of L lying above S. Then let $\beta_i = b_i^2$, and we have $z_i \in \mathcal{O}_{L,T}$ and $a_i, b_i \in \mathcal{O}_{L,T}^\times$.

Now suppose there are infinitely many solutions. For each x, there are at most two y, so there are infinitely many possible x. As $x = a_1 + (b_1 z_1)^2$, there are infinitely many possible values of

$$x = \frac{1}{4}\left(b_1 z_1 - b_3 z_3 + \frac{a_3 - a_1}{b_1 z_1 - b_3 z_3}\right),$$

hence infinitely many possible values of

$$\frac{a_2 - a_1}{(b_1 z_1 - b_3 z_3)^2} = \left(\frac{b_1 z_1 + b_2 z_2}{b_1 z_1 - b_3 z_3}\right)\left(\frac{b_1 z_1 - b_2 z_2}{b_1 z_1 - b_3 z_3}\right).$$

However, $b_i, z_i \in \mathcal{O}_{L,T}$, and $a_i - a_j = (b_i z_i - b_j z_j)(b_i z_i + b_j z_j) \in \mathcal{O}_{L,T}^\times$, so $b_i z_i \pm b_j z_j \in \mathcal{O}_{L,T}^\times$. By the Lemma, there are finitely many solutions to $w \pm z = 1$ for $w, z \in \mathcal{O}_{L,T}^\times$, but we have

$$\frac{b_1 z_1 \pm b_2 z_2}{b_1 z_1 - b_3 z_3} \mp \frac{b_2 z_2 \pm b_3 z_3}{b_1 z_1 - b_3 z_3} = 1$$

Thus there are finitely many possible possibly values of $\frac{b_1 z_1 \pm b_2 z_2}{b_1 z_1 - b_3 z_3}$, hence of their product, contradicting our assumption. This concludes the proof of the theorem. □

3.2. **Shafarevich's Theorem.** We can now prove the following theorem:

Theorem 3.9 (Shafarevich's Theorem)**.** *Let E be an elliptic curve over K and S a finite set of places. The set of isomorphism classes of elliptic curves over K with bad reduction contained in S is finite.*

Note that this gives us the following important corollary to be used in the proof of the open image theorem:

Corollary 3.10. *If E is an elliptic curve over K, then there are finitely many isomorphism classes of elliptic curves over K isogenous to K.*

Proof. This follows as any two isogenous curves have the same places of good reduction. □

We now prove Shafarevich's Theorem:

Proof. First, enlarge S so that it contains all places lying above 2 and 3, and, as in the proof of Siegel's Theorem, so that $\mathcal{O}_{K,S}$ is a PID. If E has bad reduction contained in S, suppose v is a place outside S. Then because E has good reduction, it has a Weierstrass equation of the form

$$y^2 = 4x^3 - g_{2,v}x - g_{3,v},$$

where $g_{2,v}, g_{3,v} \in \mathcal{O}_{K,v}$ and discriminant $g_{2,v}^3 - 27g_{3,v}^2 \in \mathcal{O}_{K,v}^\times$. Choose a Weierstrass equation $y^2 = 4x^3 - g_2'x - g_3'$ for E. Then we know that, for each $v \notin S$, there is $u_v \in K^\times$ such that $g_{i,v} = u_v^{2i}g_i'$. Now choose $u \in K^\times$ such that $v(u) = v(u_v)$ for all $v \notin S$. Then if we let $g_i = u^{2i}g_i'$, we have $g_i \in \mathcal{O}_{K,S}$ and discriminant in $\mathcal{O}_{K,S}^\times$.

Next, note that by Dirichlet's S-Unit Theorem, the group $(\mathcal{O}_{K,S}^\times)/(\mathcal{O}_{K,S}^\times)^{12}$ is finite. If $X \subseteq \mathcal{O}_{K,S}^\times$ is a set of coset representatives, then every elliptic curve has a Weierstrass equation as above with discriminant in X. But for a given discriminant D, the equation $g_3^2 = \dfrac{g_2^3 - D}{27}$ has finitely many solutions with $g_2, g_3 \in \mathcal{O}_{K,S}$, so there are finitely many such elliptic curves possible, and we are done. □

3.3. Tate Modules.

Definition 3.11. If E is an elliptic curve defined over a field K, the ℓ-adic *Tate module* is the inverse limit $T_\ell(E) = \varprojlim_n E[\ell^n]$.

It is automatically a $\mathbb{Z}_\ell$-module, isomorphic to $\mathbb{Z}_\ell^2$ if K has characteristic 0. Furthermore, since G_K acts continuously on each of the $E[\ell^n]$, we get a continuous action of G_K on T_ℓ, or a continuous homomorphism $G_K \to \mathrm{GL}_2(\mathbb{Z}_\ell)$, which is an ℓ-adic representation. We can take the tensor product with $\mathbb{Q}_\ell$, so that $T_\ell \otimes_{\mathbb{Z}_\ell} \mathbb{Q}_\ell$ becomes a linear representation of G_K. We denote this by V_ℓ. We denote the image of G_K in $\mathrm{GL}(V_\ell)$ by G_ℓ, and its Lie algebra by $\mathfrak{g}_\ell$.

Remark 3.12. Note that $H^1_{\acute{e}t}(E, \mathbb{Z}/\ell^n)$ corresponds to homomorphisms from the pro-ℓ part of $\pi_1^{\acute{e}t}(E)$ to $\mathbb{Z}/\ell^n$, or equivalently separable isogenies from a curve onto E with kernel cyclic of degree dividing ℓ^n along with a chosen element of the kernel, up to isomorphism. By the existence of the dual isogeny and the corresponding Weil pairing (c.f. [Sil09], Exercise 3.15), these correspond dually to ℓ^n torsion points of E. Thus the Tate module $T_\ell(E)$ comes from the dual of the representation of G_K on the étale cohomology $H^1(E, \mathbb{Z}_\ell)$.

We now come to the main topic of this paper, Serre's open image theorem:

Theorem 3.13 (Serre's open image theorem)**.** *If E is an elliptic curve with no complex multiplication defined over a number field K, and $\rho : G_K \to \mathrm{GL}(V_\ell)$ is the associated ℓ-adic representation on the Tate module, then the Lie algebra of $\rho(G_K)$ is $\mathfrak{gl}(V_\ell)$.*

Corollary 3.14. *By Fact 1.10, $\rho(G_K)$ contains an open neighborhood of the identity, hence is open.*

The proof will occupy the rest of this paper.

3.4. Irreducibility of the Tate Module.

First, we show that this representation is irreducible. We begin with a lemma.

Lemma 3.15. *If E is an elliptic curve with no complex multiplication over K, and $E' \to E$ and $E'' \to E$ are isogenies over K with non-isomorphic cyclic kernels, then E' and E'' are not isomorphic over K.*

Proof. Suppose the kernels have orders n' and n''. Suppose $E' \to E''$ is an isomorphism. We take the dual isogeny of $E' \to E$ and compose $E \to E' \to E'' \to E$ to get an isogeny $E \to E$. By the multiplicativity of degree, this isogeny has degree $n'n''$, and it is the preimage in E of a cyclic subgroup of E' of order n''. Because E has no complex multiplication, this isogeny is multiplication by an integer a. It follows that $a^2 = n'n''$, and $n' \mid a$, since the kernel contains a cyclic subgroup of order n'. But the kernel has a quotient isomorphic to $\mathbb{Z}/n''$, so $n'' \mid a$, implying that $a = n' = n''$. But we assumed that $n' \neq n''$, so we have a contradiction. □

Next, assume that E has no complex multiplication over K. We prove the following:

Theorem 3.16. *As G_K-representations V_ℓ is irreducible for all ℓ, and $E[\ell]$ is irreducible for almost all ℓ.*

Proof. Suppose V_ℓ is not irreducible, meaning there is a one-dimensional $\mathbb{Q}_\ell$-subspace Y stable under G_K. Since $T_\ell \subseteq V_\ell$ is stable under G_K, meaning $X = Y \cap T_\ell$ is a free $\mathbb{Z}_\ell$-module of rank 1. Then $X/\ell^n X$ is isomorphic to a subgroup X_n of $E[\ell^n]$ for each n, and because of G_K-invariance, this subgroup is defined over K, hence so is the quotient E/X_n. The isogeny $E/X_n \to E$ dual to the quotient map also has kernel cyclic of order ℓ^n, hence all the curves E/X_n are non-isomorphic by the lemma. But Shafarevich's Theorem implies that there are finitely many (isomorphism classes of) elliptic curves isogenous to a given curve, so we are done.

If $E[\ell]$ is reducible, we have a cyclic subgroup X_ℓ of order ℓ stable under G_K. If this is true for infinitely-many ℓ, the curves E/X_ℓ contradict Shafarevich's Theorem, so we are done. □

3.5. **Reduction to the Abelian Case.** Let $F \subseteq \mathrm{End}(V_\ell)$ be the set of endomorphisms contained in $\mathfrak{g}_\ell$ that commute with the action of $\mathfrak{g}_\ell$. By Schur's lemma, this is a division ring, and because it commutes with $\mathfrak{g}_\ell$, hence itself, it is a field. As V_ℓ has dimension 2, it is either $\mathbb{Q}_\ell$, or a quadratic extension of $\mathbb{Q}_\ell$. By basic Lie theory, $\mathfrak{g}_\ell$ must be either $\mathfrak{gl}(V_\ell)$ or $\mathfrak{sl}(V_\ell)$. The second case cannot happen for then $\wedge^2(V_\ell)$ would be a trivial representation, but this is not the case by the non-degeneracy and Galois invariance of the Weil pairing mapping $\wedge^2(V_\ell)$ to $T_\ell(\mu)$.

If F is quadratic over $\mathbb{Q}_\ell$, then V_ℓ is a one-dimensional vector space over F, and since $\mathfrak{g}_\ell$ commutes with F, it is abelian. Thus, upon moving to a finite open subgroup of G_K (equivalently, a finite extension of K), we have an abelian representation by Corollary 1.13.

3.6. **Reduction to the Existence of a Certain Compatible Representation.** We now note the following about the Tate module:

Theorem 3.17 (Criterion of Néron-Ogg-Shafarevich). *The Tate module $T_\ell(E)$ is unramified at v iff $v \nmid \ell$ and E has good reduction at v.*

Proof. See Chapter VII of [Sil09]. □

Furthermore, it is also proven that the T_ℓ form a system of strictly compatible representations of G_K with exceptional set equal to the set of primes at which E has bad reduction.

Proposition 3.18. *Suppose that for a rational prime ℓ', there exists an ℓ'-adic representation $W_{\ell'}$ compatible with $T_\ell(E)$. Then this representation is irreducible.*

Proof. This representation is compatible with $T_{\ell'}(E)$. Thus the two representations are unramified at almost all primes, hence at almost all primes, the Frobenius elements are defined in G/N, where N is the intersection of the kernels of $W_{\ell'}$ and $T_{\ell'}(E)$. But the Frobenius elements are dense in G/N by the Chebotarev Density Theorem, and because the traces of the Frobenius elements are the same, the traces of the two representations $W_{\ell'}$ and $T_{\ell'}(E)$ are the same. By the following result, the two representations are isomorphic, and $W_{\ell'}$ is irreducible: □

Lemma 3.19. *If M and N are two semi-simple modules over a K-algebra A that are finite-dimensional over K, and every element of A has the same trace when considered as a K-linear endomorphism of M and N, respectively, then M and N are isomorphic.*

Proof. See [Lan71], Chapter XVII, Corollary 3.8. □

We would now like to analyze the representation $T_\ell(E)$ more closely, under the assumption that it is abelian.

4. Hodge-Tate Decompositions

Suppose K is an algebraic extension of $\mathbb{Q}_\ell$ with discrete valuation, i.e. a finite extension of an unramified extension of $\mathbb{Q}_\ell$. Let $\mathbb{C}_\ell$ denote the completion of the algebraic closure of K (which is also the same of $\mathbb{Q}_\ell$). Note that because G_K acts continuously on $\overline{K}$, its action extends continuously and uniquely to $\mathbb{C}_\ell$. Suppose G_K acts G_K-linearly on a finite dimensional $\mathbb{C}_\ell$-vector space, i.e. $\sigma(cw) = \sigma(c)\sigma(w)$ for all $\sigma \in G_K, c \in \mathbb{C}_\ell, w \in W$.

Let $\chi : G_K \to U_\ell$ be the cyclotomic character. For $i \in \mathbb{Z}$, we define

$$W^i = \{w \in W \mid \sigma(w) = \chi(\sigma)^i w \forall \sigma \in G_K\}.$$

It is clear that this is a K-subspace of W, hence we have a natural $\mathbb{C}_\ell$-linear inclusion of $W(i) = \mathbb{C}_\ell \otimes_K W^i$ into W, which commutes with the action of G_K.

Proposition 4.1. *The natural map $\oplus_i W(i) \to W$ is injective.*

Proof. See [Tat67]. □

Definition 4.2. If the map in the preceeding proposition is surjective, we say that W with its G_K-action is of *Hodge-Tate type*, or that it possesses a *Hodge-Tate decomposition.*

If V is a vector space over $\mathbb{Q}_\ell$, and $\rho : G_K \to \mathrm{GL}(V)$ is an ℓ-adic Galois representation, we say that it is of Hodge-Tate type if $\mathbb{C}_\ell \otimes_{\mathbb{Q}_\ell} V$ is of Hodge-Tate type under its action of G_K.

4.1. Tate Modules Have a Hodge-Tate Decomposition.

It can be shown, such as in [Ser66], that such representations possess a Hodge-Tate decomposition. The proof uses Tate's theory of p-divisible groups, which are certain inverse systems of finite flat group schemes (the Tate module of an abelian variety being an example, in particular, even in characteristic $p = \ell$). As we shall see, for the proof of the open image theorem, this is only necessary when E has integral j-invariant.

Remark 4.3. This is the beginning of a much more general theory known as p-adic Hodge theory, which finds such decompositions for representations coming from étale cohomology.

5. Locally Algebraic Representations

Suppose ℓ is a prime number and K a finite extension of $\mathbb{Q}_\ell$. Let T be the torus obtained by Weil restriction of G_m from K to $\mathbb{Q}_\ell$. Then $T(\mathbb{Q}_\ell) = K^\times$, and from the local Artin map $K^\times \to G_K^{ab}$, we get a continuous map $T(\mathbb{Q}_\ell) \to G_K^{ab}$, which we denote by i. If we have a ℓ-adic representation $\rho : G_K^{ab} \to \mathrm{GL}(V)$ where V is a $\mathbb{Q}_\ell$-vector space, we can compose with i to get a continuous map $T(\mathbb{Q}_\ell) \to \mathrm{GL}(V)$.

Definition 5.1. If $\rho : G_K^{ab} \to \mathrm{GL}(V)$ is such a representation, then we say ρ is *locally algebraic* if there is an algebraic homomorphism $r : T(\mathbb{Q}_\ell) \to \mathrm{GL}(V)$ such that $\rho \circ i(x) = r(x^{-1})$ for all x sufficiently close to 1 in the ℓ-adic topology.

Note that the r above is unique because an open subset of $K^\times$ is dense in the Zariski topology. We refer to r as the algebraic morphism *associated with* ρ.

5.1. Hodge-Tate Implies Locally Algebraic.

We merely sketch the proof.

5.1.1. *Invariance Under Change of K.*

Lemma 5.2. *An ℓ-adic Galois representation is of Hodge-Tate type over a field K with discrete valuation iff it holds over a finite extension of an unramified extension of that field.*

The proof involves an application of Hilbert's Theorem 90 for GL_n, and is similar to that of Galois descent for vector spaces (thought slightly more complicated).

5.1.2. *Some Cohomological Results.* We endow $K^\times$ with the trivial G_K-action. A continuous character $G_K \to K^\times$ is the same as an element of $H^1(G_K, K^\times)$. We say for two characters ϕ, ϕ' that $\phi \sim \phi'$ if they map to the same element of $H^1(G_K, \mathbb{C}_\ell^\times)$ under the map induced by the inclusion $K^\times \to \mathbb{C}_\ell^\times$. We refer to ϕ as *admissible* if $\phi \sim 1$.

If ϕ is a character, we can define a new action of G_K on $\mathbb{C}_\ell$ by $(\sigma, c) \mapsto \phi(\sigma)\sigma(c)$, and we denote this G_K-module by $\mathbb{C}_\ell(\phi)$. We have:

Proposition 5.3. *$\mathbb{C}_\ell(\phi)$ and $\mathbb{C}_\ell(\phi')$ are isomorphic as $\mathbb{C}_\ell[G_K]$-modules iff $\phi \sim \phi'$.*

Proof. An isomorphism between the modules is the same as an element c_0 of $\mathbb{C}_\ell$ such that $c_0\phi(\sigma)\sigma(c) = \phi'(\sigma)\sigma(c_0)\sigma(c)$ for all $\sigma \in G_K, c \in \mathbb{C}_\ell$. But this is equivalent to saying that $\phi'\phi^{-1}$ is a coboundary in $H^1(G_K, \mathbb{C}_\ell^\times)$, so we are done. □

Now, suppose that E is of finite degree over $\mathbb{Q}_\ell$ with Galois closure contained in K. Let Γ_E denote the set of embeddings $E \hookrightarrow \overline{\mathbb{Q}_\ell}$.

By local class field theory, we have an exact sequence

$$1 \to U_E \to G_E^{ab} \to \hat{\mathbb{Z}} \to 1$$

Note that a uniformizer π gives splitting of the exact sequence, by local class field theory. Let $\mathrm{pr}_\pi : G_E^{ab} \to U_E$ be the projection associated with this splitting. We then compose $G_K^{ab} \to G_E^{ab} \to U_E$. We get a continuous character known as χ_E.

Now, let V be one-dimensional over E, and suppose G_K acts on V via a continuous homomorphism $\rho : G_K \to U_E$. Then $W = \mathbb{C}_\ell \otimes_{\mathbb{Q}_\ell} V$ is of dimension $d = [E : \mathbb{Q}_\ell]$ over $\mathbb{C}_\ell$. We get a

$\mathbb{C}_\ell$-linear action of E on W that commutes with the action of G_K, and we denote this action by $(w, x) \mapsto a_x(w)$.

We set

$$W_\sigma = \{w \in W \mid a_x(w) = \sigma(x)w \, \forall x \in E\}$$

where $\sigma \in \Gamma_E$.

Then we note the following lemma, which one can prove by an easy calculation:

Lemma 5.4. *Each W_σ is a one-dimensional $\mathbb{C}_\ell$-vector space, stable under G_K, W is the direct sum of them, and the Galois module W_σ is isomorphic to $\mathbb{C}_\ell(\sigma \circ \rho)$.*

Next, we note that, in the case of $\mathbb{Q}_\ell$, the restriction of the cyclotomic character to the inertia group is simply the inverse of the local Artin map. We expand on this idea, and show that for a field E, we actually have $\chi_E \sim \chi$, where χ denotes the cyclotomic character. This part is *the crux of the argument*, for it allows us to relate χ, a key part of the definition of Hodge-Tate modules, to local algebraicity, which is a condition relating to the local Artin map of local class field theory.

Theorem 5.5 (Tate)**.** *Let $\rho : G_K \to \mathrm{GL}(V)$ be an abelian representation on a vector space V_ℓ over $\mathbb{Q}_\ell$. Suppose the representation is irreducible and of Hodge-Tate type. Then ρ is locally algebraic.*

Sketch. By Schur's lemma, the set of endomorphisms of the representation is a finite field extension of $\mathbb{Q}_\ell$, call it E. Then the representation is given by a continuous character $\rho : G_K \to E$. Let L/K contain the Galois conjugates of E. Then the representation over L is also of Hodge-Tate type by Lemma 5.2, and for local algebraicity, it suffices to prove the result for G_L. We then use the results above to show that this implies the extension is locally algebraic. □

The key corollary of this is the following:

Corollary 5.6. *The ℓ-adic representation on the Tate module is locally algebraic.*

Proof. Apply the previous theorem to the fact that the Tate module is of Hodge-Tate type. □

We will now analyze locally algebraic representations more closely and then apply those results to the Tate module.

6. The Group $S_\mathfrak{m}$ and its ℓ-adic Representations

We will now develop the theory of the group algebraic group $S_\mathfrak{m}$ and its associated ℓ-adic representation, which will be a central ingredient in the proof of the open image theorem. Before we begin, we need to discuss the characters and representations of algebraic tori and the idèlic version of class field theory.

6.1. **Algebraic Tori.** We first recall that an algebraic variety can be defined over a field K, which gives a notion of A-rational points for any algebra A over K (in particular, for A an algebraic extension field of K). The Galois group G_K acts on the $\overline{K}$-rational points, and a point is K-rational iff it's fixed under the Galois group (similarly for any intermediate field). Furthermore, G_K acts on rational maps between two varieties defined over K, and a map is defined over K iff it's invariant under G_K (equivalently, if it can be defined by equations with coefficients in K). Applied to maps from a variety to $\mathbb{A}^1$, this corresponds to the fact that G_K acts on the affine coordinate ring of a variety V, and $K[V]$ is the set of fixed points of that action (by Galois descent, it follows that $\overline{K}[V] = \overline{K} \otimes_K K[V]$, and the G_K-action is given by acting on the first component). This is all covered in the first chapter of [Sil09]. If L is an extension of K, then a K-variety is also an L-variety, and the G_L-action comes from restricting to the subgroup $G_L \subseteq G_K$ (scheme-theoretically, this comes from taking the fibre product with $\operatorname{Spec}(L)$).

Definition 6.1. An *algebraic torus* over K is an algebraic group which is isomorphic to a product of copies of G_m over the algebraic closure of K. If it is isomorphic to such a product over a field L, then it is said to be *split* over L.

Definition 6.2. A *character* of an algebraic torus T is an algebraic homomorphism $T \to G_m$ (possibly defined only over an extension field). We denote the set of characters by $X(T)$.

Given any two characters, we can multiply them to get another character of our torus. We have an action of G_K on the characters, any this is a group homomorphism because multiplication in G_m is defined over K. Note that by composing with the map $G_m \to \mathbb{A}^1$, we get a ring homomorphism $\overline{K}[X(T)] \to \overline{K}[T]$ (the former being the group algebra), and this is an isomorphism by examining the affine coordinate ring of a split torus ($\overline{K}[x_1, \cdots, x_n, x_1^{-1}, \cdots, x_n^{-1}]$). Note furthermore that this gives us an action of G_K on $\overline{K}[X(T)]$ and hence on the characters themselves, which are invertible elements of that ring. The set of characters defined over K are those fixed under the action of G_K.

In other words, we have a continuous G_K-action on the free abelian group $X(T)$. It can be shown without much trouble that the category of free abelian groups with G_K action is anti-equivalent to the category of algebraic tori over K and maps defined over K. For an algebraic extension L of K, this equivalence commutes with the forgetful functor from G_K-actions to G_L-actions on free abelian groups.

We note that the torus is *split* over K iff the action of G_K is trivial (if it's split, triviality is clear, and anti-equivalence of the categories proves the other direction). Furthermore, we have the following definition:

Definition 6.3. A torus is said to be *anisotropic* if it has no nontrivial characters defined over K.

Furthermore, if T is a torus, then the intersection of all the kernels of characters defined over K is a subtorus whose character group corresponds to the quotient $X(T)/X(T)^{G_K}$. The quotient torus is split, and its character group corresponds to $X(T)^{G_K}$.

Finally, we have the following important definition:

Definition 6.4. We say that an algebraic is of *multiplicative type* if it is the product of a finite abelian group with a torus over an algebraically closed field.

All of the results extend, if we allow actions on finitely-generated abelian groups, not just free ones. In particular, we still have $\overline{K}[X(T)] = \overline{K}[T]$.

6.1.1. *Weil Restriction of Tori.* For a finite separable field extension L/K and an algebraic group G over L, there is a unique algebraic group $R_{L/K}(G)$ over K such that for any K-algebra A, we have $R_{L/K}(G)(A) = G(A \otimes_K L)$. In particular, $R_{L/K}(G)(K) = G(L)$. We can do this by viewing L as a K-algebra of finite dimension, hence viewing the elements of L as matrices over K. The condition for a matrix to be invertible is an algebraic condition, and $R_{L/K}(G)(A)$ consists of the matrices satisfying those conditions with coefficients in A.

Let T be the algebraic group given by restriction of G_m from L to K. Then $T(\overline{K}) = (\overline{K} \otimes_K L)^\times = \overline{K}^{\times [L:K]}$, where the projection onto the direct factors correspond to embeddings of L/K into $\overline{K}/K$. In particular, the character group is generated by these embeddings, and if E is the Galois closure of L/K, then G_K permutes the characters in the same way it permutes the embeddings, hence its action factors through a finite quotient of G_K corresponding to the Galois closure of L/K.

If E is a subgroup of $T(K) = K^\times$, then the Zariski closure $\overline{E}$ is also a subgroup, and $T/\overline{E}$ is an algebraic group over K. Its character group corresponds to the subset of $X(T)$ that vanishes on E (equivalently, on $\overline{E}$), and this subset is automatically invariant under G_K. If a character is given by a power n_σ for each $\sigma \in \mathrm{Hom}_K(L, \overline{K})$, then $X(T/\overline{E})$ corresponds to those characters for which $\prod \sigma(x)^{n_\sigma} = 1\, \forall x \in E$.

6.1.2. *Representations of Tori.* Suppose T is an algebraic group of multiplicative type defined over a field K, and $\phi : T \to \mathrm{GL}(V)$ is an algebraic representation (defined over some extension L/K) of T on a finite-dimensional vector space V over K. Note that the trace is an algebraic map on $\mathrm{GL}(V)$, hence this pulls back to an algebraic map on T. We denote this by $\theta_\phi \in L[T]$.

Note that $\mathrm{GL}(V \otimes_K L)$ is naturally isomorphic to the set of L-rational points of $\mathrm{GL}(V)$, so T acts on the vector space $V \otimes_K L$. In the particular case that $L = K$, we find that T acts on V. In general, the best we can say is that T acts on $V \otimes_K \overline{K}$.

To recall, a module is *semi-simple* if it is a direct sum of simple modules. By Schur's lemma, any homomorphism between two simple modules is either zero or an isomorphism, and two semi-simple modules are isomorphic iff their direct factors are isomorphic. A representation of G is semi-simple if it is semi-simple as a module over $K[G]$.

Note that all representations of algebraic groupes of multiplicative type are semi-simple, as representations of tori may be diagonalized over the algebraic closure of K, and representations of finite groups are semi-simple. It follows that the trace is a linear combination of elements of $X(T)$ with coefficients in the positive integers.

Conversely, a positive integer linear combination of elements of $X(T)$ clearly is the trace of some representation of T over $\overline{K}$. We wish to show that representations correspond uniquely to such combinations. This follows from Lemma 3.19.

We know that representations defined over K have trace in $K[T] = \overline{K}[X(T)]^{G_K}$. We next show that the converse is true:

Lemma 6.5. *A representation of a group T of multiplicative type over K is defined over K iff its trace is G_K-invariant.*

Proof. Suppose $\chi_1, \cdots, \chi_k$ is the set of conjugates of a character $\chi_1 \in X(T)$. We wish to show that $\sum_{i=1}^k \chi_i \in K[T]$ corresponds to a representation defined over K. Let $G(\chi) \subseteq G_K$ denote the subgroup fixing χ_1, which is closed by the continuity of the action of G_K on $X(T)$. Let L be the fixed field of $G(\chi)$. It follows that $\chi_1 \in L[T]$, hence $\chi_1 : T(L) \to L^\times$ gives a one-dimensional representation of T defined over L. But it is easy to see that if we restrict scalars to K, we get a $[L : K]$-dimensional representation defined over K with trace equal to $\sum \chi_i$, so we are done. □

6.2. **Idèlic CFT.** We recall the basic definitions and facts of idèlic global class field theory.

For a number field K, we let Σ_K denote the set of finite places of K, Σ_K^∞ the set of infinite places, and $\overline{\Sigma_K}$ the set of all places. For each place v of K, we let K_v denote the completion at v, $\mathcal{O}_v$ its valuation ring, and U_v the group of units. We write valuations additively, always setting the valuation of a uniformizer to be 1. We let I_K denote the group of idèles of K, and if S is a finite set of places of K, we let $\mathbf{A}_K^{S,\times}$ denote the idèles supported outside S.

Definition 6.6. A *modulus* $\mathfrak{m}$ of K is a collection of non-negative integers $\{\mathfrak{m}_v\}_{v\in\overline{\Sigma_K}}$ equal to 0 outside a finite set and equal to 0 or 1 for all infinite primes. A modulus is said to be *supported* on a finite set S of places if it is equal to 0 outside S.

For each modulus $\mathfrak{m}$ and $v \in \Sigma_K$, we let $U_{v,\mathfrak{m}}$ be the set of $u \in U_v$ such that $v(u-1) \geq \mathfrak{m}_v$, and if $\mathfrak{m}_v = 1$ and $K_v = \mathbb{R}$ for $v \in \Sigma_K^\infty$, we let $U_{v,\mathfrak{m}} = \mathbb{R}^+$, and $U_{v,\mathfrak{m}} = K_v$ otherwise. If ℓ is a rational prime, we set

$$U_{\ell,\mathfrak{m}} = \prod_{v|\ell} U_{v,\mathfrak{m}}.$$

We set

$$U_\mathfrak{m} = \prod_v U_{\mathfrak{m},v}.$$

Note that this is an open subgroup of I_K, and such subgroups form a basis for open subgroups of I_K (hence, for open sets containing the connected component of the identity).

Let $C_K = I_K/K^\times$ be the idèle class group of K. One can easily show that $C_\mathfrak{m} := I_K/(U_\mathfrak{m}K^\times) = C_K/U_\mathfrak{m}$ is the ray class group associated to $\mathfrak{m}$. By class field theory, we have an isomorphism between G_K^{ab} and the inverse limit of the $C_\mathfrak{m}$, which, by the above, is isomorphic to C_K modulo its connected component D_K. We denote the surjection $I_K \to G_K^{ab}$ by i_K, and this is called the (idèlic) global Artin map.

6.3. **The Group $S_\mathfrak{m}$.** Fix a modulus $\mathfrak{m}$ whose support contains Σ_K^∞. Let $I_\mathfrak{m} = I_K/U_\mathfrak{m}$, and let $E_\mathfrak{m} = K^* \cap U_\mathfrak{m}$. Thus we have an exact sequence

$$1 \to K^\times/E_\mathfrak{m} \to I_\mathfrak{m} \to C_\mathfrak{m} \to 1$$

Let T be $R_{K/\mathbb{Q}}(G_m)$. Let $\overline{E}_\mathfrak{m}$ be the Zariski closure of $E_\mathfrak{m}$ in $K^\times = T(\mathbb{Q})$, and denote the quotient algebraic group over $\mathbb{Q}$ by $T_\mathfrak{m}$. Then we have the following diagram with exact row:

$$\begin{tikzcd}
1 \arrow[r] & K^\times/E_\mathfrak{m} \arrow[r] \arrow[d] & I_\mathfrak{m} \arrow[r] & C_\mathfrak{m} \arrow[r] \arrow[d, "id"'] & 1 \\
1 \arrow[r] & T_\mathfrak{m}(\mathbb{Q}) & & C_\mathfrak{m} \arrow[r] & 1
\end{tikzcd}$$

We would like to complete the bottom row to make it an exact sequence and keep the diagram commutative. To do this, we take the colimit of the diagram:

$$\begin{tikzcd}
K^\times/E_\mathfrak{m} \arrow[r] \arrow[d] & I_\mathfrak{m} \\
T_\mathfrak{m}(\mathbb{Q}) &
\end{tikzcd}$$

which we denote by $S_\mathfrak{m}(\mathbb{Q})$.

Equivalently, we can view $1 \to K^\times/E_\mathfrak{m} \to I_\mathfrak{m} \to C_\mathfrak{m} \to 1$ as an element of $H^2(C_\mathfrak{m}, K^\times/E_\mathfrak{m})$, then map this into $H^2(C_\mathfrak{m}, T_\mathfrak{m}(\mathbb{Q}))$. This shows the quotient is still $C_\mathfrak{m}$. Note that the structure of $S_\mathfrak{m}$ as an algebraic variety is as the disjoint union of $|C_\mathfrak{m}|$ copies of $T_\mathfrak{m}$, and $S_\mathfrak{m}(\mathbb{Q})$ becomes the group of $\mathbb{Q}$-rational points. Furthermore, this shows that $S_{\mathfrak{m}}(A)$ is the pushout of the diagram

$$\begin{tikzcd} K^\times/E_{\mathfrak{m}} \arrow[r] \arrow[d] & I_{\mathfrak{m}} \\ T_{\mathfrak{m}}(A) & \end{tikzcd}$$

for any $\mathbb{Q}$-algebra A. In particular, if $A = \overline{\mathbb{Q}}$, we get an exact sequence

$$1 \to T_{\mathfrak{m}}(\bar{(}\mathbb{Q})) \to S_{\mathfrak{m}}(\overline{\mathbb{Q}}) \to C_{\mathfrak{m}} \to 1$$

As $T_{\mathfrak{m}}(\overline{\mathbb{Q}})$ is divisible, the sequence splits, so $S_{\mathfrak{m}}(\overline{\mathbb{Q}})$ is the product of a finite group with the torus $T_{\mathfrak{m}}$ over $\overline{\mathbb{Q}}$. In particular, $S_{\mathfrak{m}}$ is a group of multiplicative type.

6.4. **The Associated ℓ-adic Representation.** We thus end up with the following commutative diagram with exact rows:

$$\begin{tikzcd}
1 \arrow[r] & E_{\mathfrak{m}} \arrow[r] \arrow[d, hook] & U_{\mathfrak{m}} \arrow[r] \arrow[d, hook] & U_{\mathfrak{m}}/E_{\mathfrak{m}} \arrow[r] \arrow[d, hook] & 1 \\
1 \arrow[r] & T(\mathbb{Q}) = K^\times \arrow[r] \arrow[d, two heads] & I_K \arrow[r] \arrow[d, two heads] & C_K \arrow[r] \arrow[d, two heads] & 1 \\
1 \arrow[r] & K^\times/E_{\mathfrak{m}} \arrow[r] \arrow[d, two heads] & I_{\mathfrak{m}} \arrow[r] \arrow[d, "\epsilon"'] & C_{\mathfrak{m}} \arrow[r] \arrow[d, "id"'] & 1 \\
1 \arrow[r] & T_{\mathfrak{m}}(\mathbb{Q}) \arrow[r] & S_{\mathfrak{m}}(\mathbb{Q}) \arrow[r] & C_{\mathfrak{m}} \arrow[r] & 1
\end{tikzcd}$$

Note that

$$T(\mathbb{Q}_\ell) = (K \otimes_{\mathbb{Q}} \mathbb{Q}_\ell)^\times = (\oplus_{v|\ell} K_v)^\times = \oplus_{v|\ell} K_v^\times.$$

But this is a direct factor of I_K, so we get a projection $\mathrm{pr}_\ell : I_K \to T(\mathbb{Q}_\ell)$. We then get the following commutative diagram:

$$\begin{tikzcd}
I_K \arrow[dr, two heads, "\mathrm{pr}_\ell"] & K^\times = T(\mathbb{Q}) \arrow[l, hook'] \arrow[r, two heads] \arrow[d, hook] & T_{\mathfrak{m}}(\mathbb{Q}) \arrow[r] \arrow[d, hook] & S_{\mathfrak{m}}(\mathbb{Q}) \arrow[d, hook] \\
& T(\mathbb{Q}_\ell) \arrow[r] & T_{\mathfrak{m}}(\mathbb{Q}_\ell) \arrow[r] & S_{\mathfrak{m}}(\mathbb{Q}_\ell)
\end{tikzcd}$$

Consider the map $\epsilon_\ell : I_K \to I_{\mathfrak{m}} \to^\epsilon S_{\mathfrak{m}}(\mathbb{Q}) \hookrightarrow S_{\mathfrak{m}}(\mathbb{Q}_\ell)$. If we compose this with the inclusion $K^\times \hookrightarrow I_K$, we get the map from $K^\times \to S_{\mathfrak{m}}(\mathbb{Q}_\ell)$ coming from $T(\mathbb{Q}) \to T_{\mathfrak{m}}(\mathbb{Q}) \to S_{\mathfrak{m}}(\mathbb{Q}) \to S_{\mathfrak{m}}(\mathbb{Q}_\ell)$ by the commutativity of the first diagram. By the commutativity of the second diagram, this is the same as

$$T(\mathbb{Q}) \hookrightarrow I_K \to^{\mathrm{pr}_\ell} T(\mathbb{Q}_\ell) \to S_{\mathfrak{m}}(\mathbb{Q}_\ell).$$

We let α_ℓ denote the composition $I_K \to^{\mathrm{pr}_\ell} T(\mathbb{Q}_\ell) \to S_{\mathfrak{m}}(\mathbb{Q}_\ell)$. Then we have shown that α_ℓ and ϵ_ℓ coincide on $K^\times$. Thus $\epsilon_\ell \alpha_\ell^{-1} : I_K \to S_{\mathfrak{m}}(\mathbb{Q}_\ell)$ induces a map $C_K \to S_{\mathfrak{m}}(\mathbb{Q}_\ell)$, and since all the maps are continuous in the ℓ-adic topology, and $S_{\mathfrak{m}}(\mathbb{Q}_\ell)$ is totally disconnected, this factors through a map $C_K/D_K \to S_{\mathfrak{m}}(\mathbb{Q}_\ell)$. But $C_K/D_K \cong G_K^{ab}$ by the idèlic Artin map, so we have a continuous map $a_\ell : G_K \twoheadrightarrow G_K^{ab} \to S_{\mathfrak{m}}(\mathbb{Q}_\ell)$.

Definition 6.7. The map a_ℓ defined above is the ℓ-adic representation *associated to K and $\mathfrak{m}$*.

Suppose $v \notin S_\ell \cup \mathrm{Supp}(\mathfrak{m})$. Then $I_v \subseteq G_K^{ab}$ is the image of $U_v \subseteq I_K$ under the Artin map. But U_v is contained in $U_\mathfrak{m}$ and maps to 1 under $\mathrm{p}r_\ell$, so its image in $S_\mathfrak{m}(\mathbb{Q}_\ell)$ is trivial, i.e. the representation is *unramified at v*.

Let x be an idèle which is a uniformizer at v and 1 at all other primes. Then x maps trivially under pr_ℓ, so $a_\ell(x) = \epsilon_\ell(x)$. But x maps onto the Frobenius of v in G_K^{ab}, hence $\epsilon_\ell(x) = F_{v,a_\ell}$. Furthermore, ϵ_ℓ factors through $S_\mathfrak{m}(\mathbb{Q})$, so $\epsilon_\ell(x)$ is *rational and independent of ℓ* (assuming that $v \notin S_\ell$). Thus we have shown:

Theorem 6.8. *The representations $a_\ell : G_K \to S_\mathfrak{m}(\mathbb{Q}_\ell)$ form a strictly compatible system of rational ℓ-adic representations with values in $S_\mathfrak{m}$ and with exceptional set contained in* $\mathrm{Supp}(\mathfrak{m})$.

If we have a representation of $S_\mathfrak{m}$ on a $\mathbb{Q}$-vector space V, we get a corresponding system of rational ℓ-adic representations of G_K.

Furthermore, we have the following:

Proposition 6.9. *Suppose we have a rational algebraic representation $S_\mathfrak{m} \to \mathrm{GL}(V)$, where V is a finite-dimensional vector space over $\mathbb{Q}$. Then for infinitely many ℓ, the representation $S_\mathfrak{m}(\mathbb{Q}_\ell) \to \mathrm{GL}(V \otimes_\mathbb{Q} \mathbb{Q}_\ell)$ is diagonalizable.*

Proof. Note that, because $S_\mathfrak{m}$ is a torus, the representation is diagonalizable over some extension $K/\mathbb{Q}$. For any ℓ that splits completely in K, then E embeds in $\mathbb{Q}_\ell$, so the representation on $V \otimes_\mathbb{Q} \mathbb{Q}_\ell = (V \otimes_\mathbb{Q} E) \otimes_E \mathbb{Q}_\ell$ is also diagonalizable. By the Chebotarev Density Theorem, there are infinitely many such ℓ, so we are done. □

Finally, we note the following lemma and its important corollary:

Lemma 6.10. *The set of Frobenius elements $a_\ell(F_v) \in S_\mathfrak{m}(\mathbb{Q}_\ell)$ are dense in the Zariski topology.*

Proof. As the Zariski topology is coarser than the ℓ-adic topology, a_ℓ is continuous with respect to the Zariski topology, and the Frobenius elements F_v are dense in G_K, the Zariski closure of the Frobenius elements contains $\mathrm{im}(a_\ell)$.

Now note that on $\prod_{v|\ell} U_{v,\mathfrak{m}}$, a_ℓ coincides with α_ℓ^{-1}, hence it maps onto a closed subgroup of $T_\mathfrak{m}(\mathbb{Q}_\ell)$ of finite index (hence open in the ℓ-adic topology). Such a subgroup is Zariski dense in $T_\mathfrak{m}(\mathbb{Q}_\ell)$, and because a_ℓ maps G_K^{ab} onto $C_\mathfrak{m}$ (which is the Galois group of the ray class field of $\mathfrak{m}$ over K), it follows that $\mathrm{im}(a_\ell)$ is Zariski dense in $S_\mathfrak{m}$, so we are done. □

Corollary 6.11. *If a representation defined by a representation of $S_\mathfrak{m}$ over $\mathbb{Q}_\ell$ is rational, then the corresponding representation of $S_\mathfrak{m}$ comes from a rational representation.*

Proof. Choose a basis $x_1, \cdots, x_r$ for the $\mathbb{Q}$-vector space generated by the coefficients of the trace (as an element of $\mathbb{Q}[S_\mathfrak{m}] \otimes_\mathbb{Q} \mathbb{Q}_\ell$) and 1, and suppose $x_1 = 1$. Write the trace as $\sum_{i=1}^{r} \phi_i x_i$, with $\phi_i \in \mathbb{Q}[S_\mathfrak{m}]$. Because the traces of all the Frobenius elements are rational, we have $\phi_i(F_v) = 0$ for $i \neq 1$. Because the F_v are Zariski dense, this is true on all of $S_\mathfrak{m}$, i.e. the trace is in $\mathbb{Q}[S_\mathfrak{m}]$, so by Lemma 6.5, the representation is defined over $\mathbb{Q}$. □

6.5. **Back to the Open Image Theorem.** It follows that, if we show that the representation on the Tate module, if abelian, comes from a representation of $S_{\mathfrak{m}}$ defined over $\mathbb{Q}_\ell$, we will be done by Proposition 3.18, as a diagonalizable two-dimensional representation is clearly not irreducible. As we shall see, this will follow from the local algebraicity of the Tate module.

7. Locally Algebraic Representations and $S_{\mathfrak{m}}$

We now extend the notion of local algebraicity to ℓ-adic representations of number fields.

Definition 7.1. We say that a global abelian ℓ-adic representation $\rho : G_K \to G(\mathbb{Q}_\ell)$, where G is an algebraic group, is *locally algebraic* in the sense of Section 5 if its restriction ρ_v to $G_{K_v} \subseteq G_K$ is locally algebraic for all $v \mid \ell$. If we say a representation is locally algebraic, we assume it is abelian.

Once again, let T be the torus obtained by Weil restriction of G_m from K to $\mathbb{Q}$. Then $T(\mathbb{Q}_\ell) = (K \otimes_{\mathbb{Q}} \mathbb{Q}_\ell)^\times$ is a direct factor of I_K, and it maps onto the decomposition group of v in G_K^{ab}, as the idèlic Artin map i_K is the product of the local Artin maps. Thus a representation on V is locally algebraic iff there is an algebraic morphism $T \to \mathrm{GL}(V)$ defined over $\mathbb{Q}_\ell$ such that the restriction to $T(\mathbb{Q}_\ell)$ is locally equal to $f(x^{-1})$.

7.1. **Modulus of Definition.** In order to relate locally algebraic representations to the groups $S_{\mathfrak{m}}$, we need to know what $\mathfrak{m}$ to choose. We make the following definition:

Definition 7.2. If $\rho : G_K \to \mathrm{GL}(V)$ is a locally algebraic ℓ-adic representation, we say that $\mathfrak{m}$ is a *modulus of definition* for ρ if $\rho \circ i_K$ is trivial on $U_{v,\mathfrak{m}}$ for $v \nmid \ell$ and equal to (the inverse of) the associated algebraic morphism for $v \mid \ell$.

We must first show that every locally algebraic representation has a modulus of definition:

Proposition 7.3. *Every locally algebraic ℓ-adic representation $\rho : G_K \to \mathrm{GL}(V)$ has a modulus of definition.*

Proof. Suppose first that $v \mid p \neq \ell$. Then some open subgroup of $\mathrm{GL}(V)$ is a pro-ℓ group, and its preimage under $\rho \circ i_K$ is open. Furthermore, some open subgroup of $K_v^\times$ is a pro-p group. The intersection of these two groups is therefore open in $K_v^\times$, and because it is a pro-p group, and its image is contained in a pro-ℓ group, its image must be trivial.

Furthermore, by Lie theory, there is an open neighborhood N of the identity in $\mathrm{GL}(V)$ that contains no nontrivial finite subgroup. Its preimage is open in I_K, hence contains U_v for almost all v. But we know that the representation vanishes on an open subgroup of U_v, which is of finite index by the compactness of U_v. Thus for all such $v \nmid \ell$, the image of U_v in $\mathrm{GL}(V)$ is finite, hence trivial as N contains no subgroups of finite index.

Now we know that for almost all v, $\rho \circ i_K$ is trivial on U_v, and for the rest of the $v \nmid \ell$, it vanishes on an open subgroup of U_v, hence we can choose $\mathfrak{m}$ such that it vanishes on $U_{v,\mathfrak{m}}$ for all $v \nmid \ell$. By the definition of local algebraicity, we can choose $\mathfrak{m}$ such that furthermore, it is equal to the algebraic morphism on $U_{v,\mathfrak{m}}$ for $v \mid \ell$. Finally, as $\mathrm{GL}(V)$ is totally disconnected in the ℓ-adic topology, $\rho \circ i_K$ is trivial on the connected components of the infinite places of K, so we are done. □

Finally, we come to the theorem that completes the proof of Serre's open image theorem:

Theorem 7.4. *Let $\rho : G_K^{ab} \to \mathrm{GL}(V)$ be locally algebraic with modulus $\mathfrak{m}$. Then there is a $\mathbb{Q}$-subspace $W \subseteq V$ and an algebraic morphism $\phi : S_{\mathfrak{m}} \to \mathrm{GL}(W)$ defined over $\mathbb{Q}$ such that ρ is equal to $\phi \circ a_\ell$, where $\mathrm{GL}(V)$ is identified with the set of $\mathbb{Q}_\ell$-rational points of $\mathrm{GL}(W)$.*

Proof. Let $f : T \to \mathrm{GL}(V)$ be the algebraic morphism associated with ρ. We have the following commutative diagram:

$$\begin{tikzcd}
U_{\mathfrak{m}} \arrow[r, hook] \arrow[d, two heads, "\mathrm{pr}_\ell"'] & I_K \arrow[d, two heads, "i_K"'] \\
U_{\ell,\mathfrak{m}} \arrow[d, hook] & G_K^{ab} \arrow[d, "\rho"'] \\
T(\mathbb{Q}_\ell) \arrow[r, "f^{-1}"] & \mathrm{GL}(V)
\end{tikzcd}$$

We define a map $\psi : I \to \mathrm{GL}(V)$ by

$$x \mapsto (\rho \circ i_K)(x) f(\mathrm{pr}_\ell(x)).$$

It follows by the commutativity of the diagram that this is trivial on $U_{\mathfrak{m}}$, hence defines a map $\psi : I_{\mathfrak{m}} \to \mathrm{GL}(V)$.

Furthermore, we obtain the commutative diagram:

$$\begin{tikzcd}
E_{\mathfrak{m}} \arrow[r, hook] \arrow[d, hook] & T(\mathbb{Q}) \arrow[r, hook] \arrow[dl, hook] \arrow[d] & T(\mathbb{Q}_\ell) \arrow[d, "f"'] \\
I_K \arrow[r, two heads] & I_{\mathfrak{m}} \arrow[r, "\psi"] & \mathrm{GL}(V)
\end{tikzcd}$$

so f is trivial on $E_{\mathfrak{m}}$.

Thus we have an induced algebraic morphism $f : T_{\mathfrak{m}}(\mathbb{Q}_\ell) \to \mathrm{GL}(V)$. By the universal property of pushout, the morphisms f and ψ induce a homomorphism $\phi : S_{\mathfrak{m}}(\mathbb{Q}_\ell) \to \mathrm{GL}(V)$. We wish to show that $\phi \circ a_\ell$ gives ρ. Note that $\phi \circ \epsilon = \psi = (f \circ \mathrm{pr}_\ell)(\rho \circ i_K)$, and $\phi \circ \alpha_\ell = (f \circ \mathrm{pr}_\ell)$, so $\phi \circ (\epsilon \alpha_\ell^{-1}) = \rho \circ i_K$. Passing to G_K^{ab}, this says that $\phi \circ a_\ell = \rho$, so we are done. □

As we have shown, this completes the proof of Serre's open image theorem.

8. Interpretation in terms of torsion points

The theorem is equivalent to the statement that for sufficiently large n, we have $[K(E[\ell^{n+1}]) : K(E[\ell^n])] = \ell^4$.

References

[Cas67] Cassels, J. W. S., and A. Fröhlich. *Algebraic Number Theory*. London: Academic, 1967.

[Hoo42] Hooke, Robert. Linear p-Adic Groups and Their Lie Algebras. *The Annals of Mathematics*, Second Series, Vol. 43, No. 4 (Oct., 1942), p.641-655.

[Lan71] Lang, Serge. *Algebra*. Reading, Mass. [u.a.: Addison-Wesley, 1971.

[Ser66] Serre, Jean-Pierre. Groupes p-Divisibles. *Séminaire N. Bourbaki*, 1966-1968, exp n○ 318, p.73-86.

[Ser89] Serre, Jean-Pierre. *Abelian l-Adic Representations and Elliptic Curves*. Redwood City, CA: Addison-Wesley, Advanced Book Program, 1989.

[Sil09] Silverman, Joseph H. *The Arithmetic of Elliptic Curves*. Dordrecht: Springer, 2009.

[Tat67] Tate, John T. p-Divisible Groups. *Proceedings of a Conference on Local Fields*, Springer-Verlag, 1967, p.157-183.