JOURNAL
de Théorie des Nombres
de BORDEAUX

anciennement Séminaire de Théorie des Nombres de Bordeaux

Michael ROSEN

A Geometric Proof of Hermite's Theorem in Function Fields

Tome 29, n$^{\text{o}}$ 3 (2017), p. 799-813.

<http://jtnb.cedram.org/item?id=JTNB_2017__29_3_799_0>

Journal de Théorie des Nombres de Bordeaux **29** (2017), 799–813

A Geometric Proof of Hermite's Theorem in Function Fields

par Michael ROSEN

Résumé. Un théorème important de C. Hermite énonce que tout ensemble de corps de nombres, dont les discriminants sont bornés en valeur absolue, doit être fini. Correctement formulé, un résultat similaire est valable pour des corps de fonctions d'une variable sur un corps de constantes fini. Cet article donne une nouvelle preuve de ce résultat par l'analogie avec l'approche de la « géométrie des nombres » de H. Minkowski dans le cas des corps de nombres.

Abstract. An important theorem of C. Hermite asserts that any set of algebraic number fields, whose discriminants are bounded in absolute value, must be finite. Properly formulated, a similar theorem holds true for function fields in one variable over a finite constant field. This paper gives a new proof of this result by using an analogue of the geometry of numbers approach due to H. Minkowski in the number field case.

1. Introduction

We will give a new proof of the following theorem originally due to Hermite in the case of number fields. Our proof is along the lines of the classic geometric proof of the number field version (due to H. Minkowski)

Throughout this paper, $\mathbb{F}$ will denote a finite field and $k = \mathbb{F}(T)$ will denote a rational function field over $\mathbb{F}$.

Theorem 1.1. *Let B be a positive constant. There exist only finitely many geometric and separable extensions K/k in a fixed algebraic closure of k for which the discriminant divisor $d_{K/k}$ satisfies $\deg_k d_{K/k} \leq B$.*

This theorem follows from the apparently weaker theorem

Theorem 1.2. *Let B and N be two positive constants. There are only finitely many geometric and separable extensions K/k in a fixed algebraic closure of k for which $[K : k] \leq N$ and the discriminant divisor $d_{K/k}$ satisfies $\deg_k d_{K/k} \leq B$.*

Manuscrit reçu le 17 juin 2016, révisé le 14 septembre 2016, accepté le 12 septembre 2016.

2010 *Mathematics Subject Classification.* 11N56, 14G42.

The author thanks the referee for pointing out numerous misprints and for making several valuable suggestions.

It is important to assume the extensions are geometric since when the constant field is finite there is a constant field extension of every positive integer degree. All these extensions are unramified, so the degree of the discriminant is zero in every case. Thus, Theorem 1.1 is false if we allow constant field extensions. If an extension is inseparable, the trace map is identically zero which implies $d_{K/k} = 0$. However, there are infinitely many inseparable extensions of degree p of k, where p is the characteristic of k. This follows from the fact that k^*/k^{*p} is an infinite dimensional vector space over $\mathbb{Z}/p\mathbb{Z}$. Once again, this is a contradiction to Theorem 1.1.

To show that Theorem 1.2 implies Theorem 1.1 it suffices to show that $\deg_k d_{K/k}$ goes to infinity as $[K : k]$ goes to infinity. To do this one can make use of the Riemann–Hurwitz theorem which asserts

$$2g_K - 2 = n(2g_k - 2) + \deg_K D_{K/k}\,.$$

Here, g_K is the genus of K, g_k the genus of k, $n = [K : k]$, and $D_{K/k}$ is the different divisor of K over k (see [5]). To define $\deg_K$ of a divisor of K we need only define it for prime divisors w and extend by linearity. We define $\deg_K w = [O_w/M_w : \mathbb{F}]$, where M_w is the maximal ideal of O_w the valuation ring at w. We are using the hypothesis that K/k is a geometric extension, i.e. the constant field of K, $\mathbb{F}$, is equal to the constant field of k. Similarly, $\deg_k v = [O_v/M_v : \mathbb{F}]$. Now, a simple calculation using the fact that the norm of the different is equal to the discriminant proves that $\deg_K D_{K/k} = \deg_k d_{K/k}$ (see [5, p. 82]). Thus, for a geometric extension we can restate Riemann–Hurwitz as follows

$$2g_K - 2 = n(2g_k - 2) + \deg_k d_{K/k}\,.$$

We thank Siman Wong for pointing this out to us.

Now, recall that the rational function field has genus zero and we derive

$$\deg_k d_{K/k} = 2(n + g_K - 1) \geq 2n - 2\,.$$

Thus, the degree of the discriminant goes to infinity with $n = [K : k]$ as asserted.

In his 1996 treatise [1], D. Goss outlines a proof of Theorem 1.1 due to Y. Taguchi which uses the theory of moduli spaces of curves in characteristic p. In 2010, E. Widmer proved a theorem about heights of rational points on algebraic curves (see [8]) which can be shown to imply Hermite's theorem in function fields. Finally, S. Wong has written a paper (see [9]) proving Hermite's theorem in the function field case by using the Tchebotarev density theorem and a representation theory style approach.

In December of 1994 I wrote a letter to David Goss outlining a proof of the theorem which uses the classical geometry of numbers approach. The proof was complete if one restricts attention to separable extensions of $k = \mathbb{F}(T)$ which split completely at infinity. A few years ago I was able to give a proof of the theorem in full generality. It has not yet been published. It seems appropriate to include it in this collection of papers in honor of David Goss on his sixtieth birthday.

2. The Number Field Case

Before beginning the proof, we recall some background. Let $E/\mathbb{Q}$ be an algebraic number field of degree n. The classical proof makes use of the usual embedding of a number field $K/\mathbb{Q}$ in $\mathbb{R}^n$. This is done by first observing that $\mathbb{R}$ is the completion of $\mathbb{Q}$ at its unique archimedean prime. The algebraic closure of $\mathbb{R}$ is $\mathbb{C}$, the complex numbers. Let $\psi_1, \psi_2, \dots, \psi_{r_1}$ be the set of real embeddings of K into $\mathbb{C}$ and $\psi_{r_1+1}, \dots \psi_{r_1+r_2}$ a complete set of non-conjugate complex embeddings of K into $\mathbb{C}$. One then embeds K into $R^{r_1}\mathbb{C}^{r_2}$ by the map

$$\Psi(\alpha) = (\psi_1(\alpha), \dots, \psi_{r_1}(\alpha), \psi_{r_1+1}(\alpha), \dots, \psi_{r_1+r_2}(\alpha)) .$$

Finally, by taking a basis $\{1, i\}$ for $\mathbb{C}/\mathbb{R}$ we identify $\mathbb{C}$ with $\mathbb{R}^2$ and thus, $\mathbb{R}^{r_1}\mathbb{C}^{r_2}$ with $\mathbb{R}^n$. The ring of integers $\mathcal{O}_K$ embeds as a lattice in $\mathbb{R}^n$. The classical proof of Hermite's theorem makes use of this embedding and of geometric properties of lattices in Euclidean space. So, to begin with, we must find an analogue of the above construction and also an analogue of one of Minkowski's theorems in the geometry of numbers.

3. Geometry of Numbers in Function Fields

As usual, the analogue of $\mathbb{Z} \subset \mathbb{Q}$ is $\mathbb{A} \subset k$, where $\mathbb{A} = \mathbb{F}[T]$, the ring of polynomials. If K/k is a finite separable extension, we define $\mathcal{O}_K$ to be the integral closure of $\mathbb{A}$ in K. The standard absolute value at the prime at infinity of $k = \mathbb{F}(T)$ is defined on polynomials by $|f(T)|_\infty = q^{\deg f(T)}$ and extended in the obvious way to all of k. The completion of k with respect to this absolute value is the field of formal Laurent series in $1/T$, i.e. $k_\infty = \mathbb{F}((1/T))$. The field k_∞ is our replacement for $\mathbb{R}$. As our replacement for $\mathbb{C}$ we will use the separable closure k_∞^{sep} of k_∞. It is more usual to use the completion of the algebraic closure of k_∞ but for our purposes the separable closure will suffice. Note that $|\alpha|_\infty$ extends uniquely from k to k_∞ and from there to k_∞^{sep}. We will use the same notation $|\alpha|_\infty$ for $\alpha \in k_\infty^{sep}$. Since this will be the only absolute value we consider, from now on we will, for the most part, drop the subscript "∞".

Consider the k_∞ vector space $V = k_\infty^n$. The field k_∞ is locally compact. In fact, the ring $\mathcal{O}_\infty = \mathbb{F}[\![1/T]\!]$ is both open and compact. Moreover, $\mathbb{A} \cap (1/T)\mathcal{O}_\infty = (0)$, so $\mathbb{A}$ is a discrete subring of k_∞. Introducing the product topology on V, we see that V is a locally compact vector space. Let μ denote the unique Haar measure of k_∞ which has the value 1 on $\mathcal{O}_\infty$. It is a standard fact that for $\alpha \in k_\infty^*$ and any measurable set S, one has

$\mu(\alpha S) = |\alpha|\mu(S)$ (see [6, Ch. II, Prop. 2]). Now, let ν be the product measure μ^n on $V = k_\infty^n$. It is a Haar measure on V and has the value one on $\mathcal{O}_\infty^n$.

Definition 3.1. A finitely generated $\mathbb{A}$-submodule L of V is defined to be a lattice if it is discrete and V/L is compact.

An $\mathbb{A}$ lattice must have the form $L = \mathbb{A}v_1 + \mathbb{A}v_2 + \cdots + \mathbb{A}v_n$, where $\{v_1, v_2, \ldots, v_n\}$ is a k_∞ basis for V. Conversely, any such object is an $\mathbb{A}$-lattice.

Definition 3.2. $D(L) = \mathcal{O}_\infty v_1 + \mathcal{O}_\infty v_2 + \cdots + \mathcal{O}_\infty v_n$

This is an "almost" fundamental domain for L. Note that $D(L) \cap L$ is finite. In fact, $D(L) \cap L = \{\sum_i \alpha_i v_i \mid \alpha_i \in \mathbb{F}\}$. Call this set L_o. It is an $\mathbb{F}$ vector space of dimension n. From this it is easily seen that

$$V = \bigcup_{\lambda \in L/L_o} \{\lambda + D(L)\}$$

where the union is taken over a set of coset representatives of L/L_o. This is a disjoint union.

Definition 3.3. If $L \subset V$ is an $\mathbb{A}$-Lattice, define $\mathrm{vol}(L) = \nu(D(L))$.

As we will see, this definition is independent of the $\mathbb{A}$ basis chosen for L. Let $\{e_1, e_2, \ldots, e_n\}$ be the standard basis of k_∞^n, i.e. e_i has all its coefficients 0 except for a 1 at the i-th place. That $\mathrm{vol}(L)$ does not depend on a basis for L follows immediately from the following Proposition.

Proposition 3.4. *Let $L = \mathbb{A}v_1 + \mathbb{A}v_2 + \cdots + \mathbb{A}v_n$ be a lattice in V. Write $v_i = \sum_{j=1}^n c_{ij} e_j$. Then, $|\det(c_{ij})| = \mathrm{vol}(L)$.*

Proof. This is fairly standard fact, so we only sketch a proof. Let P_∞ be the maximal ideal of $\mathcal{O}_\infty$. By a box B about 0 in V we mean an open set of the form

$$B = P_\infty^{m_1} \times P_\infty^{m_2} \times \cdots \times P_\infty^{m_n}\,.$$

For a box of this form we begin by showing $\nu(\mathrm{T}B) = |\det \mathrm{T}|\nu(B)$ for any linear transformation T in $GL_n(k_\infty)$. Every matrix in $GL_n(k_\infty)$ can be written as a product of diagonal matrices, elementary matrices, and permutation matrices, so it is enough to check the result for each class. For elementary matrices, E, EB differs from B only by translation, so the measure of B is unchanged. Also, det $E = 1$. For permutation matrices, P, PB and B have the same measure and $\det P = \pm 1$. For diagonal matrices the result follows from the relation $\mu(\alpha P_\infty^m) = |\alpha|\mu(P_\infty^m)$ for $\alpha \in k_\infty^*$. Since ν is Haar measure, the result holds for the translation of a box. Finally, any open set can be approximated arbitrarily closely by a disjoint union of translations of boxes. Thus, $\nu(\mathrm{T}U) = |\det \mathrm{T}|\nu(U)$ for any open set U in V, etc.

If T denotes the linear transformation given by the matrix (c_{ij}), we have $D(L) = \mathrm{T}(\mathcal{O}_\infty^n)$. Thus, $\mathrm{vol}(L) = \nu(D(L)) = |\det \mathrm{T}|\nu(\mathcal{O}_\infty^n) = |\det \mathrm{T}|$ □

The next Proposition is a substitute for the geometric lemma of Minkowski.

Proposition 3.5. *Let L be a lattice in V and $C \subset V$ a measurable set closed under subtraction, i.e. if $c_1, c_2 \in C$ then $c_1 - c_2 \in C$. If $\nu(C) >$ $\mathrm{vol}(L)$, then $C \cap (L - L_o)$ is not empty.*

Proof. From $C = \bigcup_{\lambda \in L/L_0} (\lambda + D(L) \cap C)$, disjoint union, so

$$\nu(C) = \sum_{\lambda \in L/L_o} \nu\left(\lambda + D(L) \cap C\right) = \sum_{\lambda \in L/L_o} \nu\left(D(L) \cap C - \lambda\right).$$

In the last equality we have used the fact that ν is Haar measure. If the sets $D(L) \cap C - \lambda$ were disjoint, the last sum would be less than or equal to $\nu(D(L)) = \mathrm{vol}(L)$. This contradicts the hypothesis that $\nu(C) > \mathrm{vol}(L)$. Thus, there exist coset representatives $\lambda_1 \neq \lambda_2$ for L/L_o such that $C - \lambda_1 \cap C - \lambda_2$ is non-empty. It follows that there are elements $c_1 \in C$ and $c_2 \in C$ such that $c_1 - \lambda_1 = c_2 - \lambda_2$. Thus, $\lambda_1 - \lambda_2 = c_1 - c_2 \in C$. Thus, $\lambda = \lambda_1 - \lambda_2$ is in $C \cap L$. λ is not in L_o since λ_1 and λ_2 represent different cosets of L/L_o. □

The set C in the Proposition is a subgroup of V. It is non-empty, so if $c \in C$ then $0 = c - c$ is in C and $-c = 0 - c$ is in C, Finally, if $c_1, c_2 \in C$ the $c_1 + c_2 = c_1 - (-c_2)$ is in C. We will use the result when C is a box about the origin of sufficiently large measure.

4. The Geometric Embedding in Function Fields

It turns out that to prove the Main Theorem, it is sufficient to consider the case where K/k is a geometric Galois extension. We will prove this reduction in an appendix. From now on we will make the assumption that K/k is a Galois extension with Galois group G. This considerably simplifies the exposition.

There are n k-embeddings of K into k_∞^{sep}. We will identify k with its image in k_∞. Suppose ψ is a fixed embedding of K into k_∞^{sep}. Then all the embeddings are given by $\{\psi \circ \sigma \mid \sigma \in G\}$. This follows from the fact that $\psi \circ \sigma$ is a k-embedding for all $\sigma \in G$ and that they are all distinct since ψ is one to one. This gives n distinct embeddings and that is all there can be. It follows that for any two embeddings ψ_1 and ψ_2 we have $\psi_1(K) = \psi_2(K)$.

Consider the field $\hat{K} \subset k_\infty^{sep}$ which is the compositum of $\psi(K)$ and k_∞. $\hat{K}$ is a finite extension of k_∞ and so it is complete in the valuation given by $|\alpha|$ on k_∞^{sep}. All the k-embeddings of K into k_∞^{sep} have image in $\hat{K}$ and that image is dense. We also note that since $\psi(K)/k$ is a Galois extension with group $\psi G \psi^{-1}$ we must have $\hat{K}/k_\infty$ is a Galois extension. Denote its Galois group by H. By restriction to $\psi(K)$ we can identify H with a subgroup of $\psi G \psi^{-1}$ and $\psi^{-1} H \psi$ with a subgroup of G.

As usual, we can extend the valuation $|\beta|$ from k to K by choosing an k-embedding ψ_i from K to k_∞^{sep} and defining $|\beta|_{\psi_i} = |\psi_i(\beta)|$ for all $\beta \in K$. It is easily checked that this is an extension of $|\cdot|$ from k to K. All extensions of $|\cdot|$ are obtained this way. As is well known, this subgroup of G given by $\psi^{-1} H \psi$ is the decomposition group of the valuation $|\cdot|_\psi$ on K. For different embeddings ψ_i we get the decomposition groups of the valuations $|\cdot|_{\psi_i}$. For all this, see Serre [6, Ch. 2, §3].

Lets call two embeddings ψ_1 and ψ_2 equivalent if they induce the same valuation on K. From what has been said so far, it is not hard to show that ψ_1 and ψ_2 are equivalent if and only is there is an element h in H such that $\psi_2 = h\psi_1$. Let $m = [\hat{K} : k_\infty] = \#H$, the local degree. If g is defined by $n = gm$, then there are exactly g equivalence classes of embeddings corresponding to g distinct extensions of $|\cdot|$ to K. Let $\{\psi_1, \psi_2, \ldots, \psi_g\}$ be a set of representatives for these classes.

We claim that the set $\{h\psi_i \mid h \in H, 1 \leq i \leq g\}$ is the set of all the k-embeddings of K into k_∞^{sep}. Each element of this set is an embedding and there are $mg = n$ of them, so it is only necessary to show they are all distinct. If $h_1\psi_1 = h_2\psi_2$, then $|h_1\psi_1(\beta)| = |h_2\psi_2(\beta)|$ for all $\beta \in K$, which implies $|\psi_1(\beta)| = |\psi_2(\beta)|$ for all $\beta \in K$, i.e. ψ_1 and ψ_2 are equivalent. It follows that $\psi_1 = \psi_2$. Since $h_2^{-1} h_1 \psi_1 = \psi_1$, and since $\psi_1(K)$ is dense in $\hat{K}$ we can conclude that $h_2^{-1} h_1 = e$, in other words $h_1 = h_2$.

We are now in a position to identify the map we will need to embed K into $V = k_\infty^n$. Let $\{\psi_1, \psi_2, \ldots . \psi_g\}$ be a complete set of inequivalent embeddings and define, for $\beta \in K$,

$$\Psi'(\beta) = (\psi_1(\beta), \psi_2(\beta), \ldots, \psi_g(\beta)) \in \hat{K}^g\,.$$

By choosing a basis for $\hat{K}$ over k_∞ we can identify $\hat{K}$ with k_∞^m and thus $\hat{K}^g$ with $k_\infty^{mg} = k_\infty^n = V$. This is like the classical case where we choose a basis $\{1, i\}$ to identify $\mathbb{C}$ with $\mathbb{R}^2$. The easiest case is when $\hat{K} = k_\infty$, the analogue of K being a totally real number field. In that case Ψ embeds K directly into V. In this situation, the path to the proof of our theorem is very smooth. For the general case there are a number of obstacles to be overcome. One has to show that there are only finitely many possibilities for the field $\hat{K}$. Also, the basis of $\hat{K}$ over k_∞ must be chosen carefully.

5. The local Hermite Theorem

Let's take up the first problem about $\hat{K}$. This problem doesn't arise in the classical case because the completion of $\mathbb{Q}$ at its unique archimedean prime is $\mathbb{R}$. The only extensions of $\mathbb{R}$ are $\mathbb{R}$ itself and $\mathbb{C}$. In our case, however, k_∞ has many extensions. For p-adic fields there are only finitely many extensions of bounded degree. This follows from Krasner's lemma (see Lang's book [4, Ch. II, Prop. 14]). However, in characteristic p we can have infinitely many extensions of bounded degree. As we shall see, this is because of facts about Artin–Schreier extensions in local fields. Proposition 5.2 below can be thought of as a local version of Hermite's theorem.

So as not to introduce too much notation, in this section we will use the letters k and K, for somewhat different objects than in the rest of the paper. We suppose $k = \mathbb{F}((U))$ is a Laurent series field over a finite field $\mathbb{F}$ with $q = p^t$ elements. Let $|\cdot|_k$ be the normalized absolute value, i.e. $|U|_k = q^{-1}$. Set $v_k(\alpha) = -\log_q |\alpha|_k$ for all $\alpha \in k^*$. The function $v_k(\alpha)$ is the normalized additive valuation on k, i.e. $v_k(U) = 1$. If K/k is a finite, separable extension, we define the discriminant of K/k, $d_{K/k}$, to be the discriminant of the valuation ring of K, $\mathcal{O}_K$, as a module over $\mathcal{O}_k$. This is well defined up to the square of a unit in $\mathcal{O}_k$. The absolute value of $d_{K/k}$ is well defined, and therefore so is $v_k(d_{K/k})$. We need to know how this quantity behaves in towers.

Lemma 5.1. *Let L/k be a finite separable extension and K/k an intermediate extension. Then,*

$$v_k(d_{L/k}) = [K:k]v_k(d_{L/K}) + [L:K]v_k(d_{K/k})\,.$$

In particular, if $v_k(d_{L/k})$ is bounded above by B then $v_k(d_{K/k})$ and $v_K(d_{L/K})$ are also bounded above by B.

Proof. We note that v_k extends uniquely to K and the normalized additive valuation on K satisfies $ev_k = v_K$, where $e = e(K/k)$ is the ramification index.

We recall the relation $D_{L/k} = D_{L/K}D_{K/k}$ among the various multiplicative differents involved. Recall also that the discriminant is the norm of the different. Take the norm $N_{L/k}$ of both sides and then take $|\cdot|$ of the result. Since $N_{L/k} = N_{K/k}N_{L/K}$, we find

$$|d_{L/k}|_k = |N_{K/k}(d_{L/K})|_k |d_{K/k}|_k^{[L:K]}\,.$$

In a separable extension the norm of an element is the product of the conjugates of the element. Since we are in the local case, the conjugate of an element has the same absolute value as the element itself. Thus, our relation can be rewritten as

$$|d_{L/k}| = |d_{L/K}|_k^{[K:k]} |d_{K/k}|_k^{[L:K]}\,.$$

The equation in the statement of the Lemma follows by taking $-\log_q$ of both sides. The equation immediately shows that $v_k(d_{K/k}) \leq v_k(d_{L/k})$. Since $e = e(K/k) \leq [K : k]$ the equation also implies

$$v_k(d_{L/k}) \geq [K : k]v_k(d_{L/K})) \geq e\ v_k(d_{L/K}) = v_K(d_{L/K})\,. \qquad \square$$

Proposition 5.2 (The Local Hermite Theorem)**.** *Let k be a Laurent series field in one variable over a finite field $\mathbb{F}$ of characteristic p. Let v_k be its canonical additive valuation. Let L/k be a finite Galois extension of k inside a fixed separable closure of k. Fix positive integers B and N. There are only finitely many L with $[L : k] \leq N$ and $v_k(d_{L/k}) \leq B$.*

Proof. Let G be the Galois group of L/k and $P \subset G$ a p-Sylow subgroup. Consider the fixed field K of P. We must have $[K : k]$ is not divisible by p. In this case the proof of the finiteness of the number of such K with $[K : k]$ bounded can be proved by adopting the proof of the p-adic case given in [4]. There the problem is reduced to the totally ramified case. The theory of Eisenstein polynomials together with a compactness lemma yields the result. This doesn't work in the characteristic p case because an Eisenstein polynomial can be inseparable if p divides the ramification index. However, if p doesn't divide the ramification index, a monic Eisenstein polynomial is separable which is what is needed to invoke Krasner's Lemma. Applying this reasoning to K/k shows that only finitely many such extensions with $[K : k] \leq N$ can exist.

The Galois group of L/K is P, a p-group, so there is a tower

$$K = K_0 \subset K_1 \subset K_2 \subset \cdots \subset K_{t-1} \subset K_t = L\,.$$

where each extension K_{i+1}/K_i is cyclic of degree p. Thus, each of these extensions is generated by the root of an Artin–Schreier equation $x^p - x = a_i$, where $a_i \in K_i$. Under the hypotheses of the Proposition, we will show that for each K_i there are only finitely many possible Artin–Schreier extensions of degree p. This will complete the proof.

By a general structure theorem, each K_i has the form $K_i = \mathbb{E}((\pi_i))$, where π_i is a uniformizing parameter. Since $[\mathbb{E} : \mathbb{F}] \leq [L : k] \leq N$, we see that there are only finitely many possibilities for the finite field $\mathbb{E}$. Let's temporarily drop the subscript "i". Set $E = \mathbb{E}((\pi))$ and let M/E be an Artin- Schreier extension of degree p. Thus, M is generated by a root of $x^p - x - a$ where $a \in E$. If $v_E(a) > 0$, this equation has a root in E. In fact, let $\alpha = \sum_{i=0}^{\infty} a^{p^i}$. This series converges and $\alpha - \alpha^p = a$. It follows that $-\alpha$ is a root of our equation and so $M = E$. Every element in E is the sum of a polynomial over $\mathbb{E}$ in π^{-1} and an element a such that $v(a) > 0$. It therefore suffices to consider Artin–Schreier equations where a is a polynomial over $\mathbb{E}$ in π^{-1}. Let λ be the degree of this polynomial. H. Hasse has shown that if p does not divide λ then $v_K(d_{M/K}) = (p-1)(\lambda+1)$. He also shows that

If p does divide λ then the extension is given with a polynomial in π^{-1} of smaller degree. See H. Hasse [2] or H. Stichtenoth [7, Prop. 3.7.8, p. 127]. Among other things, this shows, in contrast to the p-adic case, there are infinitely many Galois extensions of E of degree p. However, if we also bound v_K of the discriminant, then λ is bounded, and so there are only finitely many such extensions.

It remains to show that if $v_k(d_{L/k})$ is bounded by B then $v_{K_i}(d_{K_{i+1}/K_i})$ is also bounded by B for each i from 0 to $t-1$. This, however, follows by repeated application of Lemma 5.1. □

6. Proof of the Main Theorem

Having explored the local case, we now return to the global setting and recall our original notation. First, $k = \mathbb{F}(T)$ is a rational function field over a finite field $\mathbb{F}$ with characteristic p and K/k is a finite Galois extension with group G and degree n. For a polynomial $f(T)$ let $|\cdot|_\infty$ be the absolute value on k defined by $|f(T)|_\infty = q^{\deg f}$, and k_∞ the completion of k with respect to this absolute value. We have, $k_\infty = \mathbb{F}((T^{-1}))$. Let ψ be an embedding of K into k_∞^{sep} and $\hat{K} = k_\infty \psi(K)$. In our statement of the main theorem, we suppose that both n and the discriminant divisor of K/k is bounded. This easily implies that both $[\hat{K} : k_\infty]$ and $v_{k_\infty}(d_{\hat{K}/k_\infty})$ are bounded. From the local Hermite theorem, Proposition 5.2 above, we see that there are only finitely many possibilities for $\hat{K}$. We recall the mapping Ψ' defined immediately after the proof of Proposition 5.2.

$$\Psi'(\beta) = (\psi_1(\beta), \psi_2(\beta), \ldots, \psi_g(\beta)) \in \hat{K}^g$$

where $\beta \in K$ and the ψ_i vary over a complete set of inequivalent embeddings of K into k_∞^{sep}.

To go further, we now choose an element $\gamma \in \hat{K}$ which is a normal basis element for $\hat{K}/k_\infty$. This means that $\{h_1\gamma, h_2\gamma, \ldots, h_m\gamma\}$ is a basis for $\hat{K}$ over k_∞. Here, h_j vary over the elements in H. For convenience we choose $h_1 = e$. A careful discussion of the existence of a normal basis element is given in section 14.14 of Jacobson's book, Basic Algebra I [3]. It is proven there that γ is a normal basis element if and only if $\det(h_i h_j(\gamma)) \neq 0$. We will use this important fact later. One last restriction; for convenience we assume $|\gamma| \leq 1$. This can always be achieved by multiplying γ with a sufficiently high power of $1/T$ if necessary.

We now, finally, extend the map Ψ' to a map Ψ from K into $k_\infty^n = V$. Let $\beta \in K$ and, for $i = 1, 2, 3, \ldots, g$ write $\psi_i(\beta) = \sum_{j=1}^m a_{ji} h_j(\gamma)$, where $a_{ji} \in k_\infty$. We then define

$$\Psi(\beta) = (a_{11}, a_{21}, \ldots, a_{m1}, a_{12}, a_{22}, \ldots, a_{m2}, \ldots, a_{1g}, a_{2g}, \ldots, a_{mg}) \in k_\infty^n = V.$$

It is easy to see that Ψ is a map of k_∞ vector spaces.

Now, take the equation $\psi_i(\beta) = \sum_{j=1}^m a_{ji}h_j(\gamma)$ and multiply both sides with h_k to get

$$h_k\psi_i(\beta) = \sum_{j=1}^m h_k h_j(\gamma) a_{ji}\,.$$

Set Γ be the $m \times m$ matrix $(\ h_k h_j(\gamma)\ )$. Then,

$$\begin{pmatrix} h_1\psi_i(\beta) \\ h_2\psi_i(\beta) \\ \vdots \\ h_m\psi_i(\beta) \end{pmatrix} = \Gamma \begin{pmatrix} a_{1i} \\ a_{2i} \\ \vdots \\ a_{mi} \end{pmatrix}$$

Let $\Gamma^{(g)}$ be the $n \times n$ matrix with the matrix Γ arranged by blocks along the diagonal and all the other entries zero. As we have seen, the $h_k\psi_i$ run through all the k-embeddings of K into $\hat{K}$. Arrange these lexicographically and rename them $\lambda_1, \lambda_2, \ldots, \lambda_n$. Finally, define $\Lambda(\beta) = (\lambda_1(\beta), \lambda_2(\beta), \ldots, \lambda_n(\beta))$ arranged as a column vector. Then, from the above matrix equation we find

$$\Lambda(\beta) = \Gamma^{(g)}\Psi(\beta) \in \hat{K}^n\,.$$

Let $\{\omega_1, \omega_2, \ldots, \omega_n\}$ be an integral basis for $\mathcal{O}_K$ over $\mathbb{A} = \mathbb{F}[T]$. Finally, we derive the matrix equation which will show that $\Psi(\mathcal{O}_K)$ is an $\mathbb{A}$-lattice in V and enable us to calculate the volume of its fundamental domain. Consider the equation of $n \times n$ matrices

$$(\Lambda(\omega_1), \Lambda(\omega_2), \ldots, \Lambda(\omega_n)) = \Gamma^{(g)}\,(\Psi(\omega_1), \Psi(\omega_2), \ldots, \Psi(\omega_n))\,.$$

Just as in the classical case, the determinant of the left hand side squared is $d_{\mathcal{O}_K/\mathbb{A}}$, the discriminant of $\mathcal{O}_K$ considered as an $\mathbb{A}$-module. Let's just call this discriminant d from now on. It is the finite part of the discriminant of K/k. Since we are assuming K/k is separable, $d \neq 0$. This implies that the $\mathbb{A}$-module generated by the vectors $\Psi(\omega_i)$ is an $\mathbb{A}$ lattice in V and that the volume of its fundamental domain is, by Proposition 3.4,

$$\mathrm{vol}(\Psi(\mathcal{O}_K)) = |\det(\Gamma)|^{-g}\sqrt{|d|}\,.$$

For future reference, we note that the left hand side of this equation is an integral power of q, so this must be true of the right hand side as well. Also note that since we have chosen γ to have absolute value ≤ 1 it follows that $|\det(\Gamma)|^{-g} \geq 1$. Here, and in the rest of the paper, the absolute value $|\cdot|$ refers to the absolute value at infinity, $|\cdot|_\infty$.

We are finally in position to prove the main result, Theorem 1.1. We can assume $n = [K : k] > 1$, since otherwise $K = k$, and there is nothing to prove. The idea is to produce an element $\alpha \in \mathcal{O}_K$ such that the coefficients of $\Psi(\alpha)$ satisfy certain inequalities. One consequence of these inequalities will be that α generates K over k. Secondly, the inequalities put a bound on the coefficients of the minimal polynomial satisfied by α over $\mathbb{A}$. It will follow that there are only finitely many such polynomials and therefore only finitely many such α, and consequently only finitely many fields K satisfying the hypotheses of the Theorem.

Recall that for $1 \le i \le g$ and $\beta \in K$ we set $\psi_i(\beta) = \sum_{j=1}^{m} a_{ji} h_j(\gamma)$ and then the a_{ij} are assembled into the column vector $\Psi(\beta)$. Let's consider the domain in V defined by the inequalities

$$(*) \qquad |a_{11}| \le q^n |\det(\Gamma)|^{-g} \sqrt{|d|} \quad \text{and} \quad |a_{ij}| \le q^{-1} \quad \text{for all } i,j \ne 1,1\,.$$

As we pointed out earlier, the right hand side of the first inequality is an integral power of q. Thus, the sets defined by the inequalities are just powers of P_∞. From this one easily deduces that the volume of the open set defined by all the inequalites is $q \det(\Gamma)^{-g}\sqrt{|d|}$ which is greater than $\mathrm{vol}(\Psi(\mathcal{O}_K)) = |\det(\Gamma)|^{-g}\sqrt{|d|}$. By Proposition 3.5, there is an $\alpha \in \mathcal{O}_K$ such that $\Psi(\alpha)$ satisfies our inequalities.

Assume to begin with that $g > 1$. From $\psi_i(\alpha) = \sum_{j=1}^{m} a_{ji} h_j(\gamma)$ and the above inequalities, we see immediately that $|\psi_i(\alpha)| < 1$ for $i = 2, 3, \ldots, g$. If $|\psi_1(\alpha)| \le 1$, then α would have absolute value ≤ 1 at all infinite places. Since $\alpha \in \mathcal{O}_K$ it has absolute value ≤ 1 at all finite places. This would imply $\prod_v |\alpha|_v < 1$, where the product is over all normalized valuations of K. This contradicts the product formula since $\alpha \ne 0$. Thus, $|\psi_1(\alpha)| > 1$ and also, from the inequalities $(*)$, we also see that we have the inequality $|a_{11}| > 1$.

Now suppose that $g = 1$. In this case, $|H| = |G|$ (equivalently, $m = n$) and so, there is just one place in K above the prime ∞ of k. Thus, all the embeddings $h\psi_1$ are equivalent and, also, $\{h(\gamma) \mid h \in H\}$ is a normal basis for $\hat{K}$ over k_∞. Write $\psi_1(\alpha) = \sum_{j=1}^{n} a_{j1} h_j(\gamma)$. We claim that $|a_{11}| > 1$. If not, then by $(*)$ we have $|\psi_1(\alpha)| \le 1$. Since $\alpha \in \mathcal{O}_K$ the absolute value of α at all primes both finite and infinite is ≤ 1. Since $\alpha \ne 0$ this is only possible if α is a non-zero constant. This will lead to a contradiction. To see this let's change the notation sightly and write $\alpha = \sum_{h \in H} a(h) h(\gamma)$. Let $h' \in H$ and apply h' to both sides of this equation. Since $\alpha \in k \subset k_\infty$, $h'\alpha = \alpha$ which implies

$$\alpha = \sum_{h \in H} a(h) h' h(\gamma)\,.$$

We have two expressions for α. Looking at the coefficient of e in both, we find $a(e) = a(h'^{-1})$. It follows that $a(h)$ is constant on H. Thus, $|a(e)| = |a(h)|$ for all $h \in H$ and $|a(h)| \le q^{-1}$ for $h \ne e$. We deduce $|\psi_1(\alpha)| < 1$. Since α is a non-zero constant it has absolute value 1 at all primes, but this contradicts the fact that we have shown it has absolute value less than 1 at the infinite prime. This concludes the proof that $|a_{11}| = |a(e)| > 1$. Notice that once again we must have $|\psi_1(\alpha)| > 1$.

We have shown that $|a_{11}| > 1$ and $|\psi_1(\alpha)| > 1$ in all cases. From now on we make no restrictions on g, the number of infinite primes.

Now, let us consider once again the equation

$$(**) \qquad h_k\psi_i(\alpha) = \sum_{j=1}^{m} a_{ji}h_kh_j(\gamma)\,.$$

We claim that $\psi_1(\alpha) \neq h_k\psi_i(\alpha)$ except when $i = 1$ and $k = 1$ (we continue to assume, as we can, that $h_1 = e$).

Let us first take the case where $i = 1$. In equation $(**)$, take $i = 1$ and $k \geq 2$. The coefficient of $h_k(\gamma)$ in $h_k\psi_1(\alpha)$ is seen to be a_{11} which has absolute value > 1. The coefficient of $h_k(\gamma)$ in $\psi_1(\alpha)$ has absolute value < 1. Thus, $\psi_1(\alpha) \neq h_k\psi_1(\alpha)$.

Now suppose that $i \geq 2$. If $\psi_1(\alpha) = h_k\psi_i(\alpha)$ take the absolute value of both sides. The left hand side has absolute value > 1 and the right hand side has absolute value < 1, so, once again $\psi_1(\alpha) \neq h_k\psi_i(\alpha)$.

Finally, let's revert to our alternate way of describing the embeddings of K into k_∞^{sep}, namely $\{\psi_1 \circ \sigma \mid \sigma \in G = \mathrm{Gal}(K/k)\}$. The translation of what we have just proved is that $\psi_1(\alpha) \neq \psi_1(\sigma\alpha)$ for all $e \neq \sigma \in G$. Since ψ_1 is one to one, it follows that $\alpha \neq \sigma\alpha$ for all $e \neq \sigma \in G$. Thus, α generates K over k.

It remains to show that if $\alpha \in \mathcal{O}_K$ satisfies the inequalities in $(*)$ there are only finitely many equations of degree n of which it can be a root. To see this, note that $|h_k\psi_1(\alpha)| \leq q^n|\det(\Gamma)|^{-g}\sqrt{|d|}$ for every k from 1 to m and similarly, $|h_k\psi_i(\alpha)| \leq q^{-1}$ for $i \geq 2$ and every k. Now, α is a root of $f(x) = \prod_{\sigma\in G}(x - \sigma\alpha)$. It follows that the coefficients of $f(x)$, which are in $\mathbb{A}$, are bounded by constants which depend only on n and $|d|$ and the choice of $\hat{K}$. There are only finitely many such polynomials for each choice of $\hat{K}$. By the local Hermite theorem, Proposition 5.2, there are only finitely many choices for $\hat{K}$, and it follows that there are only finitely many choices for α. This completes the proof.

Remark 6.1. The proof seems to depend only on d which is the finite component of $d_{K/k}$. However, this is an illusion. We needed a bound on the infinite component of $d_{K/k}$ in order to prove the local Hermite theorem, Proposition 5.2.

Appendix A. Reduction to the Galois Case.

In this appendix it will be shown that to prove Theorem 1.2 it is sufficient to consider Galois extensions of k. To prove this we will need several facts about how the discriminant divisor behaves in towers, These facts are undoubtedly well known, but we have not been able to locate a good reference.

Suppose k is a global function field, i.e. a function field in one variable with a finite field of constants $\mathbb{F}$. Let K_1 and K_2 be two separable finite extensions of k contained in a common algebraic closure. Let $L = K_1K_2$ denote the compositum of K_1 and K_2. For any separable finite extension K of k let $D_{K/k}$ denote the different divisor of K over k.

Proposition A.1. $D_{L/k} \le D_{K_1/k} + D_{K_2/k}$. *(We use the conorm maps to identify D_{K_1} and D_{K_2} with divisors of L.)*

Proof. Since the assertion is local in nature we can reduce to the case where k is complete with respect to a discrete rank one valuation and consequently, so are K_1, K_2, and L.

It is a little more convenient to work with ideals at first. So, we let O, O_1, O_2, and O_3 be the valuation rings in the fields k, K_1, K_2 and L. Let $\mathcal{D}_{K_i/k}$ represents the ideal different from K_i to k. We need to prove that $\mathcal{D}_{K_1/k}\mathcal{D}_{K_2/k} \subseteq \mathcal{D}_{L/k}$, where the left hand side is contained in the right hand side by extension of ideals.

The extension K_2 over k is separable and the residue class fields are finite. Thus, by [6, Ch. III, Prop. 12], there is a $\theta \in O_{K_2}$ such that $O_{K_2} = O_k[\theta]$. Let $f(x) \in O_k[x]$ be the monic irreducible polynomial for θ over k. It is a standard fact that $\mathcal{D}_{K_2/k} = (f'(\theta))$. See [6, Cor. 3 to Prop. 11].

Now, since θ generates K_2 over k it also generates L over K_1. Let $g(x)$ be the monic irreducible for θ over K_1. One has $f(x) = g(x)h(x)$ where $g(x), h(x) \in O_{K_1}[x]$. Differentiate and substitute $x = \theta$ and we find $f'(\theta) = h(\theta)g'(\theta)$. Using [6, Cor. 1 to Prop. 11], one deduces that $g'(\theta) \in \mathcal{D}_{L/K_1}$. Thus

$$\mathcal{D}_{K_2/k} = (f'(\theta)) \subseteq (g'(\theta)) \subseteq \mathcal{D}_{L/K_1}\,.$$

Thus, $\mathcal{D}_{L/k} = \mathcal{D}_{L/K_1}\mathcal{D}_{K_1/k} \supseteq \mathcal{D}_{K_2/k}\mathcal{D}_{K_1/k}$. This establishes the proposition for ideals and it is a simple matter to translate this result back to the language of divisors. □

Corollary A.2. *Suppose $K_1, K_2, \ldots, K_t$ be a set of t finite and separable extensions of a global field k all contained in a common algebraic closure of k. Let $L = K_1K_2 \ldots K_t$. Then,*

$$D_{L/k} \le D_{K_1/k} + D_{K_2/k} + \cdots + D_{K_t/k}\,.$$

Proof. This follows by induction on t using the Proposition. □

We now specialize to the case where $k = \mathbb{F}(T)$ is the rational function field over a finite field. Let K/k satisfy the hypotheses of Theorem 1.2. Let $K = K_1, K_2, \ldots, K_t$ denote the conjugates of K/k inside an algebraic closure of K. Then, $L = K_1K_2 \ldots K_t$ is the Galois closure of K over k. If $n = [K : k]$ it is a fact that $[L : k] \le n!$.

Proposition A.3. *As in Theorem 1.2, let K vary over separable and geometric extensions of k for which $\deg_k d_{K/k} \leq B$ and $[K:k] \leq n$. Then, the Galois closures of these K vary over Galois extensions of L/k for which $\deg_k d_{L/k} \leq n!B$ and $[L:k] \leq n!$.*

Proof. In the above Corollary to Proposition A.1 take the norm, $N_{L/k}$ of both sides of the inequality presented there. The left hand side become $d_{L/k}$. On the right hand side, note that $N_{L/k}D_{K_i/k} = N_{K_i/k}N_{L/K_i}D_{K_i/k} = [L:K_i]d_{K_i}$. Since the K_i are all conjugates of K we have the equality $[L:K_i] = [L:K]$ for all i. Also, $d_{K_i/k} = d_{K/k}$ for all i. Substituting in the above Corollary, we find

$$d_{L/k} \leq t[L:K]d_{K/k} = [L:k]d_{K/k}\,.$$

Thus, $\deg_k d_{L/k} \leq [L:k]\deg_k d_{K/k}$. Since $[K:k] \leq n$ it follows that $[L:k] \leq n!$. Thus, $\deg_k d_{L/k} \leq n!B$ □

Before stating and proving the main result of this Appendix, we need the following Lemma.

Lemma A.4. *Let $\mathbb{E}/\mathbb{F}$ be a finite extension of finite fields. Set $k = \mathbb{F}(T)$ and $k' = \mathbb{E}(T)$. Suppose that A' is a divisor of k'. Then, $[\mathbb{E}:\mathbb{F}]\deg_{k'} A' = \deg_k N_{k'/k}A'$.*

Proof. Both sides of the identity are linear in A' so it suffices to prove the Lemma for a prime divisor w of k'. Let v lie below w. Consider the residue class field at w, $O_w/\mathcal{P}_w$ as an extension of $\mathbb{F}$. There are two intermediate fields, $\mathbb{E}$ and O_v/P_v. Using the multiplicativity of indices in towers we compute

$$[\mathbb{E}:\mathbb{F}]\deg_{k'} w = f(w/v)\deg_k v = \deg_k(N_{k'/k}w)\,.$$

This proves the Lemma for prime divisors and thus in general. □

Theorem A.5. *Suppose Theorem 1.2 can be proven for all geometric, Galois extensions of $k = \mathbb{F}(T)$. Then, Theorem 1.2 is true in general.*

Proof. Suppose K is a separable and geometric extension of k with

$$[K:k] \leq n \text{ and } \deg_k d_{K/k} \leq B\,.$$

Let L be the Galois closure of K over k. Our conclusion seems to follow immediately from Proposition A.3. However, L/k is not, in general, a geometric extension, so there remains some work to be done. Let $\mathbb{E}$ be the constant field of L. Then, L is a geometric and Galois extension of $k' = \mathbb{E}(T)$. Since $[\mathbb{E}:\mathbb{F}] \leq [L:k] \leq n!$, there are only finitely many possibilities for $\mathbb{E}$. Thus, it suffices to show that for each possible $\mathbb{E}$ there are only finitely many possibilities for L as a geometric, Galois extension of k'. We do this by showing there are bounds on $[L:k']$ and $\deg_{k'} d_{L/k'}$ which depend only on n and B. The first requirement is easy, since $[L:k'] \leq [L:k] \leq n!$.

To establish the required bound on $d_{L/k'}$ we begin by noticing that L is the compositum of the fields $K'_i = \mathbb{E}K_i$ which are separable and geometric extensions of k'. Following the proof of Proposition A.3, we derive the following inequality

$$\deg_{k'} d_{L/k'} \leq [L : k'] \deg_{k'} d_{K'/k'} .$$

Since we have assumed that $\deg_k d_{K/k} \leq B$, we will finish the proof by showing $\deg_{k'} d_{K'/k'} = \deg_k d_{K/k}$. We start by considering the extension K'/k with its two intermediate extensions K and k'. Using the formula for the behavior of the different in towers, we find $D_{K'/K} + D_{K/k} = D_{K'/k'} + D_{k'/k}$. Since constant field extensions are unramified, $D_{K'/K}$ and $D_{k'/k}$ are both trivial, so $D_{K/k} = D_{K'/k'}$. Recall that $D_{K/k}$ is identified with a divisor of K' via the conorm map. Now, take the norm $N_{K'/k}$ of both sides using the behavior of the norm in towers, and we derive

$$[\mathbb{E} : \mathbb{F}] d_{K/k} = N_{k'/k} d_{K'/k'} .$$

Finally, apply the lemma to the k'-divisor $d_{K'/k'}$ and we find $\deg_k d_{K/k} = \deg_{k'} d_{K'/k'}$, as asserted. □

References

[1] D. Goss, *Basic Structures of Function Field Arithmetic*, Ergebnisse der Mathematik und ihrer Grenzgebiete. 3, vol. 35, Springer, 1996, xiii+422 pages.

[2] H. Hasse, "Theorie der relativ-zyklischen algebraischen Funktionenkörper, insbesondere bei endlichem Konstantenkörper", *J. Reine Angew. Math.* **172** (1934), p. 27-54.

[3] N. Jacobson, *Basic Algebra I*, 2nd ed., Freeman and Company, 1985, xviii+499 pages.

[4] S. Lang, *Algebraic Number Theory*, corr. 2nd ed., Graduate Texts in Mathematics, vol. 110, Springer, 1986, xiii+354 pages.

[5] M. Rosen, *Number Theory in Function Fields*, Graduate Texts in Mathematics, vol. 210, Springer, 2002, xii+358 pages.

[6] J.-P. Serre, *Local Fields*, Graduate Texts in Mathematics, vol. 67, Springer, 1979, Translated from the French by Marvin Jay Greenberg, vii+241 pages.

[7] H. Stichtenoth, *Algebraic Function Fields and Codes*, 2nd ed., Graduate Texts in Mathematics, vol. 254, Springer, 2009, xiii+355 pages.

[8] M. Widmer, "Small generators of function fields", *J. Théor. Nombres Bordx.* **22** (2010), no. 3, p. 747-753.

[9] S. Wong, "A field theoretic proof of Hermite's theorem for function fields", *Arch. Math.* **105** (2015), no. 4, p. 351-360.

Michael Rosen
Box 1917
151 Thayer Street
Providence, RI 02912, USA
E-mail: michael_rosen@brown.edu
URL: https://www.math.brown.edu/~mrosen