Counting Certain Points of Bounded Height in the Function Field Setting

Theorem (Northcott, 1949): Fix a degree d, a dimension n and a positive bound B. There are only finitely many non-zero points $(1, \alpha_1, \cdots, \alpha_{n-1}) \in \overline{\mathbb{Q}}^n$ with $[\mathbb{Q}(\alpha_i) \colon \mathbb{Q}] \leq d$ for all i and with absolute height $H(1, \alpha_1, \ldots, \alpha_{n-1}) \leq B$.

Note that if α is a root of unity, then $H(1, \alpha) = 1$. Thus all three parameters n, d and B must be bounded to get such a result.

Question: Can we estimate the *number* of such points in Northcott's Theorem?

For any field k and a point $P = (\alpha_1 \colon \cdots : \alpha_n)$ in projective space over some algebraic closure, let $k(P)$ denote the field obtained by adjoining to k all possible quotients α_i/α_j.

Notation: For a number field k, degree $d \geq 1$, $n \geq 2$ and positive bound B, let $N_k(n, d, B)$ denote the number of points $P \in \mathbb{P}^{n-1}(\overline{\mathbb{Q}})$ with $[k(P) \colon k] = d$ and $H(P) \leq B$.

Question: What can we say about $N_k(n, d, B)$?

Quote (Weil 1967): Once the presence of the real field, albeit at infinite distance, ceases to be regarded as a necessary ingredient in the arithmetician's brew, it goes without saying that the function-fields over finite fields must be granted a fully simultaneous treatment with number-fields, instead of the segregated status, and at best separate but equal facilities, which hitherto have been their lot.

Question: What can we say about the function field analogs of $N_k(n, d, B)$?

For each place v of $\mathbb{Q}$ (v is either ∞ or a prime) we have a corresponding absolute value:

$$|\cdot|_\infty = \text{the usual absolute value}$$

$$|\cdot|_p = \text{the usual } p\text{-adic absolute value}$$

Set

$$\|\mathbf{x}\|_v = \max_{1\leq i\leq n} \{|x_i|_v\}$$

for $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{Q}^n$.

Product Formula: For all non-zero $x \in \mathbb{Q}$ we have

$$\prod_v |x|_v = 1.$$

Definition: For a non-zero $\mathbf{x} \in \mathbb{Q}^n$, the (absolute) height of $\mathbf{x}$ is

$$H(\mathbf{x}) = \prod_v \|\mathbf{x}\|_v.$$

Note that, thanks to the Product Formula, this is actually a function on projective space $\mathbb{P}^{n-1}(\mathbb{Q})$.

If $P \in \mathbb{P}^{n-1}(\mathbb{Q})$, then P has exactly two representative points of the form $(z_1, \ldots, z_n) \in \mathbb{Z}^n$ where the greatest common divisor of the z_i's is 1.

Thus, $N_{\mathbb{Q}}(n, 1, B)$ is one-half the number of "primitive" lattice points $\mathbf{z} \in \mathbb{Z}^n$ in the cube

$$C(n, B) = \{\mathbf{x} \in \mathbb{R}^n \colon |x_i| \leq B\}.$$

How can one count the number of such lattice points?

For a fixed positive integer a, let us write $N(n, a, B)$ for the number of lattice points $\mathbf{z} \in \mathbb{Z}^n \cap C(n, B)$ satisfying $a | z_i$ for all $i = 1, \ldots, n$.

Then we have the elementary estimate

$$N(n, a, B) = \frac{2^n B^n}{a^n} + O\big((B/a)^{n-1}\big).$$

By Möbius inversion,

$$\begin{aligned} 2N_{\mathbb{Q}}(n, 1, B) &= \sum_{a \geq 1} \mu(a)\big(N(n, a, B) - 1\big) \\ &= \frac{2^n}{\zeta(n)} B^n + O(B^{n-1})^*. \end{aligned}$$

Suppose k is a number field. Each absolute value $|\cdot|_v$ on $\mathbb{Q}$ extends in a well-known way to absolute values $|\cdot|_w$ on k; we write $w|v$ in this situation.

Done in the usual way, we have

$$\prod_{w|v} |x|_w = |x|_v^{[k\colon \mathbb{Q}]}$$

for all $x \in \mathbb{Q}$.

Generalized Product Formula:

$$\prod_w |x|_w = 1 \qquad \text{all } x \in k^\times.$$

As before, we get a height on k^n:

$$H_k(x_1, \dots, x_n) = \prod_w \max_{1 \le i \le n} \{|x_i|_w\}$$

and an *absolute* height

$$H(x_1, \dots, x_n) = H_k^{1/[k\colon \mathbb{Q}]}(x_1, \dots, x_n).$$

As with our original height on $\mathbb{P}^{n-1}(\mathbb{Q})$, the functions H_k and H are actually functions on projective space $\mathbb{P}^{n-1}(k)$ thanks to the Generalized Product Formula.

Moreover, H is "absolute" in the following sense.

Suppose $\mathbf{x} = (x_1, \dots, x_n) \in \overline{\mathbb{Q}}^n$ (non-zero). Then $\mathbf{x} \in k^n$ where $k = \mathbb{Q}(x_1, \dots, x_n)$. Whence our height

$$H(\mathbf{x}) = H_k^{1/[k:\,\mathbb{Q}]}(\mathbf{x}).$$

But of course if K is *any* number field containing k, we have $\mathbf{x} \in K^n$, and thanks to our normalizations above

$$H_K(\mathbf{x}) = H_k(\mathbf{x})^{[K:\,k]}.$$

So

$$H(\mathbf{x}) = H_K^{1/[K:\,\mathbb{Q}]}(\mathbf{x}) = H_k^{1/[k:\,\mathbb{Q}]}(\mathbf{x}).$$

Theorem (Schanuel, 1979): For any number field k we have

$$N_k(n, 1, B) = S_k(n, 1)B^{ne} + O(B^{ne-1})^*,$$

where $e = [k\colon \mathbb{Q}]$ and $S_k(n, 1)$ is the Schanuel Constant.

How is this obtained? You generalize the simple lattice-point argument above, use Möbius inversion on integral ideals, and incorporate methods from the proof of the Dedekind-Weber Theorem to deal with the units.

Note that the number of points P counted in Schanuel's theorem with $\mathbb{Q}(P)$ strictly contained in k is a smaller order of magnitude.

Idea: Is it possible to sum over *all* degree d extensions to estimate $N_k(n, d, B)$:

$$N_k(n, d, B) \sim \sum_{[K:\ k]=d} S_K(n, 1) B^{ned}?$$

Short answer: no!

If $p(X) \in k[X]$ is a defining polynomial for α of degree d over k, then $H^d(1, \alpha)$ is close to the height of the coefficient vector of $p(X)$, which one may view as a point in $\mathbb{P}^d(k)$. Thus, one expects that

$$N_k(2, d, B) \gg\ll B^{(d+1)de},$$

where $e = [k\colon \mathbb{Q}]$ again.

Theorem (Masser & Vaaler, 2009): For any number field k and any degree $d > 1$ we have

$$N_k(2, d, B) = S_k(2, d) B^{(d+1)de} + O(B^{(d+1)de-d} \log B),$$

where $S_k(2, d)$ is the Masser-Vaaler Constant®.

However, we do have

$$N_k(n,d,B) \sim S_k(n,d)B^{ned}$$

with

$$S_k(n,d) = \sum_{[K:\ k]=d} S_K(n,1)$$

in the following cases.

(Schmidt, 1995): $k = \mathbb{Q}$, $d = 2$, $n \geq 4$

(Gao, unpublished thesis): $k = \mathbb{Q}$, $d \geq 3$, $n \geq d+2$

(Widmer, 2009): $[k\colon \mathbb{Q}] = e > 1$, $n > 5d/2+3+2/(ed)$

Now Andre Weil wants to know about function fields!

Fix a prime p, let $\mathbb{F}_p$ be the field with p elements and let T be transcendental over $\mathbb{F}_p$. The places of $\mathbb{F}_p(T)$ correspond to the irreducible polynomials $P(T) \in \mathbb{F}_p[T]$ and the degree function. We have absolute values

$$|z|_{P(T)} = \exp\big(-\operatorname{ord}_{P(T)}(z)\big)$$

$$|z|_{\deg} = \exp\big(\deg(z))\big)$$

for $z \in \mathbb{F}_p[T]$.

To simplify things, just call the negative of the degree of z the "order" at that place (the place corresponding to the degree). We then have an order function for every place v, and these are extended to order functions on the entire field of rational functions $\mathbb{F}_p(T)$ in the obvious manner:

$$\operatorname{ord}_v(z/y) = \operatorname{ord}_v(z) - \operatorname{ord}_v(y), \qquad z, y \in \mathbb{F}_p[T].$$

Observation (Analog to the Product Formula for $\mathbb{Q}$): For all non-zero $x \in \mathbb{F}_p(T)$ we have

$$\sum_v \operatorname{ord}_v(x) \deg(v) = 0,$$

where the degree of a place is the degree of the corresponding irreducible polynomial, or 1 in the case of the place corresponding to the degree function.

Definition: For a non-zero $\mathbf{x} \in \mathbb{F}_p(T)^n$, the (absolute logarithmic) height of $\mathbf{x}$ is

$$h(\mathbf{x}) = -\sum_v \operatorname{ord}_v(\mathbf{x}) \deg(v).$$

As before, the (analog to the) Product Formula shows that this is a function on projective space $\mathbb{P}^{n-1}(\mathbb{F}_p(T))$.

One may exponentiate to get a true analog of the absolute height on $\mathbb{Q}^n$; the traditional choice is to use the prime p for the base:

$$H(\mathbf{x}) = p^{h(\mathbf{x})}.$$

Actually, one isn't really using the *prime* p for the base here so much as the *cardinality* of the field $\mathbb{F}_p$. Why is that? Because ...

What about finite algebraic extensions of $\mathbb{F}_p(T)$, i.e., function fields?

This is somewhat more complicated than the case for number fields, since it's possible to algebraically extend the field $\mathbb{F}_p$.

Fix an algebraic closure $\overline{\mathbb{F}_p}$. Then if k is a finite algebraic extension of $\mathbb{F}_p(T)$, we have

$$k \cap \overline{\mathbb{F}_p} = \mathbb{F}_{q_k}.$$

This field is called the *field of constants* of k.

Definition: The effective degree of the extension k is

$$\frac{[k\colon \mathbb{F}_p(T)]}{[\mathbb{F}_{q_k} : \mathbb{F}_p]}.$$

Suppose k is a function field. Then every order function on $\mathbb{F}_p(T)$ extends in a well-known way to order functions on k, which one may normalize to have image $\mathbb{Z} \cup \{\infty\}$. Moreover, the degree of the places may be extended as well so that the following holds.

Generalized Observation:

$$\sum_v \operatorname{ord}_v(x) \deg(v) = 0, \qquad \text{all } x \in k^\times.$$

This is the well-known statment that the degree of a principal divisor is zero, and it is the analog to our Generalized Product Formula for number fields.

A *divisor* is simply an element of the free abelian group generated by the places:

$$\mathfrak{A} = \sum_v a_v \cdot v, \qquad a_v \in \mathbb{Z} \text{ and } a_v = 0 \text{ a.e.}$$

and the degree of such a divisor is

$$\deg(\mathfrak{A}) = \sum_v a_v \deg(v).$$

Analogous to our supremum norms before, we set

$$\mathrm{ord}_v(x_1, \dots, x_n) = \min_{1 \le i \le n} \{\mathrm{ord}_v(x_i)\}$$

for any $\mathbf{x} \in k^n$.

To each non-zero $\mathbf{x} \in k^n$ we thus get a divisor

$$\mathrm{div}(\mathbf{x}) = \sum_v \mathrm{ord}_v(\mathbf{x}) \cdot v.$$

Analogous to what we had before, we have a height on k^n:

$$h_k(\mathbf{x}) = -\deg\big(\mathrm{div}(\mathbf{x})\big).$$

Again, thanks to our analog to the Generalized Product Formula, this is a function on projective space $\mathbb{P}^{n-1}(k)$.

The *absolute* height is given by

$$h(\mathbf{x}) = \frac{1}{e} h_k(\mathbf{x}),$$

where e is the effective degree of k over $\mathbb{F}_p(T)$.

Thanks to our normalizations, if K is any function field containing k, we get

$$h(\mathbf{x}) = \frac{1}{e'} h_K(\mathbf{x}) = \frac{1}{e} h_k(\mathbf{x}),$$

where e' is the effective degree of K over $\mathbb{F}_p(T)$.

As before, this justifies our use of the term "absolute" since h is genuinely a function on projective space over an algebraic closure of $\mathbb{F}_p(T)$.

If one desires a direct analog of the absolute height on $\overline{\mathbb{Q}}$, just exponentiate:

$$H(\mathbf{x}) = p^{h(\mathbf{x})}.$$

Suppose k is a function field with field of constants $\mathbb{F}_q$ and set $e = [k \colon \mathbb{F}_q(T)]$, the effective degree of k over $\mathbb{F}_p(T)$. If K is an extension field of degree d with the same field of constants, then the effective degree of K over $\mathbb{F}_p(T)$ is ed. Thus, the height of any point P with $k(P) = K$ is necessarily of the form $m/(ed)$ for some non-negative integer m.

Definition: Let k be as above. For integers d, n and m we let $N_k(n, d, m)$ denote the number of points P in projective $n-1$-space with height $h(P) = m/(ed)$ and such that $k(P) = K$ for some function field K of degree d over k with the same field of constants.

Theorem (Thunder & Widmer, 2011): Fix a function field k as above. Then for all integers $n \geq 2d + 4$ and $m \geq 0$ we have

$$N_k(n, d, m) \sim S_k(n, d) q^{mn},$$

where $S_k(n, d)$ is the "Schanuel Constant:"

$$S_k(n, d) = \sum_{[K:\ k]=d} S_K(n, 1)$$

(sum only over those K with field of constants $\mathbb{F}_q$).

Theorem (Kettlestrings, 2011): In the theorem above, when $d = 2$ one may take $n \geq 4$ (at least when $p \neq 2$).

Whither Schanuel's Theorem?

Our lattice point estimate above is directly analogous to counting the number of $\mathbf{x}$ in

$$L(\mathfrak{A}, n) = \{\mathbf{x} \in k^n \colon \operatorname{ord}_v(\mathbf{x}) \geq -\operatorname{ord}_v(\mathfrak{A})\}$$

for a fixed divisor $\mathfrak{A}$.

We note that $L(\mathfrak{A}, n)$ is actually a finite-dimensional vector space over $\mathbb{F}_q$ (the field of constants of k). Denote its dimension by $l(\mathfrak{A}, n)$. Then $q^{l(n,\mathfrak{A})}$ is a direct analog of our lattice-point counting function $N(n, a, B)$.

Whereas one uses geometry of numbers to get estimates for $N(n, a, B)$, in this situation we have much stronger estimates.

Theorem (Riemann-Roch): With the notation above, there is a non-negative integer g (called the *genus* of k) and a class of divisors $\mathfrak{W}$ such that

$$l(\mathfrak{A}, n) = nl(\mathfrak{A}, 1) = n\big(\deg(\mathfrak{A}) + 1 - g + l(\mathfrak{W} - \mathfrak{A}, 1)\big)$$

for all divisors $\mathfrak{A}$. Moreover, we have

$$l(\mathfrak{A}, 1) = \deg(\mathfrak{A}) + 1 - g$$

whenever $\deg(\mathfrak{A}) \geq 2g - 1$ and $l(\mathfrak{A}, 1) = 0$ whenever $\deg(\mathfrak{A}) < 0$.

Theorem (Clifford): In the Riemann-Roch Theorem, if $0 \leq \deg(\mathfrak{A}) < 2g - 1$, we have

$$l(\mathfrak{A}, 1) \leq 1 + \frac{1}{2}\deg(\mathfrak{A}).$$

Recall how before we had

$$2N_{\mathbb{Q}}(n,1,B)=\sum_{a\geq 1}\mu(a)\big(N(n,a,B)-1\big).$$

Now we have

$$(q-1)N_{\mathbb{F}_q(T)}(n,1,m)=\sum_{\mathfrak{A}\geq 0}\mu(\mathfrak{A})\big(q^{l(\mathfrak{A}_0-\mathfrak{A},n)}-1\big),$$

where $\mathfrak{A}_0$ is any divisor of degree m.

Thanks to the Riemann-Roch Theorem

$$(q-1)N_{\mathbb{F}_q(T)}(n,1,m)=\sum_{i=0}^{m}\sum_{\substack{\mathfrak{A}\geq 0\\ \deg(\mathfrak{A})=i}}\mu(\mathfrak{A})\big(q^{n(m-i+1)}-1\big).$$

But our field of rational functions $\mathbb{F}_q(T)$ not only has genus 0, but an exceedingly simple zeta function, so that

$$\sum_{\substack{\mathfrak{A}\geq 0\\ \deg(\mathfrak{A})=i}} \mu(\mathfrak{A}) = \begin{cases} 1 & \text{if } i=0, \\ -(q+1) & \text{if } i=1, \\ q & \text{if } i=2, \\ 0 & \text{if } i>2. \end{cases}$$

We actually get a closed form expression:

$$\begin{aligned}(q-1)N_{\mathbb{F}_q(T)}(n,1,m) = q^{n(m+1)} - 1 \\ -(q+1)(q^{nm}-1) \\ +q(q^{m-1}-1)!\end{aligned}$$

(That's not a factorial!)

The exact same argument (well, you need to sum over divisor classes, too) works for an arbitrary function field. It isn't *quite* as simple, but it's still much cleaner than the case for $\mathbb{Q}$, even, since we have

$$\left| \sum_{\substack{\mathfrak{A} \geq 0 \\ \deg(\mathfrak{A}) = i}} \mu(\mathfrak{A}) \right| \ll q^{i(1+\epsilon)/2}$$

for any $\epsilon > 0$.

(Did I mention that the "Riemann Hypothesis" here isn't a "hypothesis?")

In the end, you get . . .

Theorem (Serre, Di-Pippo, Wan): Let k be a function field with field of constants $\mathbb{F}_q$. Then

$$N_k(n, 1, m) = S_k(n, 1)q^{mn} + O(q^{m(1+\epsilon)/2}),$$

where the "Schanuel Constant" $S_k(n, 1)$ is given by

$$S_k(n, 1) = \frac{J}{(q-1)\zeta_k(n)q^{n(g-1)}}.$$

Here J is the number of divisor classes of degree 0 (i.e., the "class number"), g is the genus and ζ is the zeta function for k.

But this doesn't really help us, though.

Why $n \geq 2d+4$?

When summing over all extensions K of k, one must be careful with the error term in the Schanuel theorem!

Theorem (Thunder & Widmer, 2011): With k as above,

$$N_k(n,1,m) = S_k(n,1)q^{mn} + O\big(q^{m(1+\epsilon)}q^{g(n-2-2\epsilon)}\big)$$

when $m \geq 2g-1$ and $n \geq 4$. When $m < 2g-1$ and $n \geq 2$ we have

$$N_k(n,1,m) \ll q^{m(\epsilon+(n+1)/2)}.$$

The implicit constants above depend only on n, q, ϵ and the effective degree of k over $\mathbb{F}_p(T)$.

This error term is not so bad, except when the genus of k approaches (and exceeds) the bound m. Here one is resorting to Clifford's Theorem instead of Riemann-Roch.

In Kettlestrings' thesis, he gives a more accurate estimate for $l(\mathfrak{A}, n)$ when one assumes there is an $\mathbf{x} \in L(\mathfrak{A}, n)$ generating K over k, assuming K is a quadratic extension.

In fact, this is precisely what we want since we sum over all extensions K of degree d the number of $P \in \mathbb{P}^{n-1}(K)$ with $k(P) = K$ and $h(P) = m/(ed)$.

What is the "Truth?"

The analog of Masser-Vaaler is much simpler here. Since there are no archimedean places, we have $dh(1,\alpha)$ is exactly equal to the height of a defining polynomial.

One readily sees that

$$dN_k(2,d,m) \sim N_k(d+1,1,m) \sim S_k(d+1,1)q^{(d+1)m}.$$

Clearly $N_k(n,d,m) > N_k(2,d,m)$ whenever $n > 2$. Thus, an asymptotic estimate of the kind noted before is only possible when $n > d$. Moreover, another result of Schmidt indicates it is likely the case that this is possible only when $n \geq d+2$.

Perhaps with more work and/or more graduate students, the gap between $n = d + 2$ and $n = 2d + 3$ can be filled in general, so that we would have

$$N_k(n, d, m) \sim q^{mn} \sum_{[K:\ k]=d} S_K(n, 1)$$

whenever $n \geq d + 2$.

That leaves us with the cases $n = 3, \ldots, d + 1$.