Notes on the Weil conjectures.

Kevin Buzzard

April 26, 2012

Written Feb 2010.

1 The definition.

The idea is that the Riemann zeta function can be thought of as

$$\zeta(s)=\prod_c(1-\#k(c)^{-s})^{-1}$$

where the product is over the closed points c of $\mathrm{Spec}(\mathbf{Z})$, and $k(c)$ is the residue field of c.

The analogue is this: if V is a scheme of finite type over a finite field k of size q, then we can define

$$\zeta_V(s)=\prod_{c\in|V|}(1-\#k(c)^{-s})^{-1}$$

where $|V|$ denotes the closed points of the scheme V.

1.1 Where does this converge?

Well, here's a fundamental but trivial link between closed points and k_n-points, with k_n the extension of k of degree n: a closed point c with $k(c)$ of size q^d (let's call such c a "degree d point") contributes 0 to $V(k_n)$ if $d\nmid n$ (because there are no field maps $k_d\to k_n$) and d points if $d\mid n$ (because in this case there are d field maps). Hence if V is a union of M affine opens, each of which is a closed subspace of some big affine N-space, then a hugely crude bound for the number of degree d closed points is Mq^{Nd}, because this number is M times the size of $\mathbf{A}^N(k_d)$. So now we use the old trick: the product of $(1+x_n)$ converges (to something non-zero) iff $\sigma_n x_n$ converges, so if C_d denotes the number of closed points of degree d then we need to check $\sum_{d\geq 1}C_dq^{-ds}$ converges, and $|q^{-ds}|=|e^{-ds\log(q)}|=e^{-d\sigma\log(q)}=q^{-d\sigma}$, where $\sigma=\mathrm{Re}(s)$. So, because $C_d\leq Mq^{Nd}$, we have absolute convergence if $\sum_{d\geq 1}q^{Nd}q^{-d\sigma}$ converges, so in particular if $|q^{N-\sigma}|<1$, that is $\mathrm{Re}(s)>N$, then the power series converges.

1.2 A more algebraic way of looking at $\zeta_V(s)$.

If C_d is the number of points of degree d, then

$$\zeta_V(s)=\prod_{d\geq 1}(1-q^{-ds})^{-C_d}$$

and if $u=q^{-s}$ then this is

$$Z_V(u):=\prod_{d\geq 1}(1-u^d)^{-C_d}=1+C_1u+\left(\frac{C_1(C_1+1)}{2}+C_2\right)u^2+\dots$$

which is clearly in $\mathbf{Z}[[u]]$ (and, as we've just seen, converges for $|u|$ sufficiently small). Moreover, for u very small, we have $Z_V(u)\simeq 1+C_1u=1+N_1u$, where N_1 is the size of $V(k)$.

1.3 How to relate $Z_V(u)$ to number of k_n-points.

This is just algebra. Here's the trick. If N_n is $\#V(k_n)$ then, because a closed point of degree d contributes d points to N_n if $d \mid n$ (and none otherwise) we have

$$N_n = \sum_{d|n} dC_d.$$

Hence (with $'$ denoting derivative with respect to u)

$$\begin{aligned}
\frac{uZ'_V(u)}{Z_v(u)} &= u(\log(Z_v(u)))' \\
&= u\sum_{d\geq 1}((-C_d \log(1-u^d)))' \\
&= u\sum_{d\geq 1} -C_d \frac{-du^{d-1}}{1-u^d} \\
&= \sum_{d\geq 1} C_d du^d(1-u^d)^{-1} \\
&= \sum_{d,m\geq 1} C_d du^{md} \\
&= \sum_{n\geq 1}(\sum_{d|n} C_d d)u^n \\
&= \sum_{n\geq 1} N_n u^n
\end{aligned}$$

So the generating function for the number of points (which I can sometimes work out, e.g, for projective space or a CM elliptic curve) is related to $Z_V(u)$.

How to go the other way? Say $G_v(u) = \sum_{n\geq 1} N_n u^n = N_1 u + N_2 u^2 + \ldots$ is the generating function. Then $\int (G_V(u)/u)du$ (the formal integral) is $\log(Z_V(u))$, which has no constant term, so if we define the integral to have no constant term then

$$Z_V(u) = \exp\left(\int \frac{G_V(u)}{u} du\right)$$

1.4 How does changing the base field change things?

If k'/k is a finite extension of degree n, and $Z'_{V'}(u')$ is the resulting function, with $u' = u^n$, then there's a simple relation between $Z'_{V'}(u')$ and $Z_V(u)$. Indeed, a closed point of degree d in V gives rise to a $k_d \otimes_k k'$-valued point of V' and hence, if $e = \gcd(d, n)$, e points each of degree dn/e. Now an easy calculation shows that the factor $(1-u^d)^{-1}$ in $Z_V(u)$ turns into $(1-u^{dn/e})^{-e}$ which is just $\prod_\zeta (1-(\zeta u)^d)^{-1}$, with ζ running through the nth roots of unity. We conclude

$$Z'_{V'}(u^n) = \prod_{\zeta^n=1} Z_V(\zeta u).$$

2 Elementary Examples.

2.1 Projective spaces.

We can now compute our first zeta functions. For projective 0-space let's do it the hard way, for a sanity check. We have $G_V(u) = \sum_{n\geq 1} u^n = u/(1-u)$, and the integral of $1/(1-u)$ is $-\log(1-u)$, so $Z_V(u) = (1-u)^{-1}$, which of course can easily be seen from the definition (this is a case where I can compute the zeta function from its definition!) and I get a check on the formula above.

Now let's try projective 1-space. We have $G_V(u) = \sum_{n\geq 1}(1+q^n)u^n$ which is $u/(1-u) + qu/(1-qu)$. Divide by u to get $1/(1-u)+q/(1-qu)$. Integrate to get $-\log(1-u)-\log(1-qu)$. Exponentiate and get $Z_V(u) = (1-u)^{-1}(1-qu)^{-1}$. This to me seems a bit less obvious. It says, for example,

$$\prod_{f\in k[X] \text{ monic irreducible}} (1-u^{\deg(f)}) = 1/(1-qu)$$

as a power series in u (or a complex function of a sufficiently small $u \in \mathbf{C}$).

This idea generalises to projective r-space: it gives $G_V(u) = u/(1-u)+qu/(1-qu)+\ldots+q^ru/(1-q^ru)$, and dividing by u and integrating and exponentiating gives, I guess, $Z_V(u) = (1-u)^{-1}(1-qu)^{-1}\ldots(1-q^ru)^{-1}$. That's pretty awesome.

Now here's a *second* way of computing the zeta function of projective 1-space, which we'll use to great effect for more general curves later on. Let V be projective 1-space. We have $Z_V(u) = \prod_c(1-u^{d(c)})^{-1}$, the product over the closed points, so

$$Z_V(u) = (1-u^{-1})\sum_{D\geq 0} u^{d(D)}$$

where the sum is now over all the positive divisors on $\mathbf{A}^1$. But on $\mathbf{A}^1$ every divisor is principal with a unique monic generator, so

$$Z_V(u) = (1-u)^{-1}\sum_{m\geq 0} q^m u^m = 1/(1-u)(1-qu).$$

We'll use a generalisation of this trick when dealing with more general curves below. While I'm here let me note that if V is projective 1-space then $uZ_V(u)$ is invariant under $u \mapsto 1/(qu)$, or, equivalently, that $\zeta_V(s)q^{-s}$ is invariant under $s \mapsto 1-s$.

2.2 Hyperelliptic curves.

Now we move onto some less trivial examples. Let's take an elliptic curve $E : y^2 = f(x)$ with f monic in $k[x]$ and of degree 3, and with no repeated root, and $p \neq 2$. Here is how Artin dealt with this case (and more generally with the hyperelliptic case when $p \neq 2$). At infinity we have one point (note: if I were dealing with f of smaller degree I might have two points) of degree 1. All the other points correspond to some prime ideal of the ring $R := k[x,y]/(y^2-f(x))$. Now if P is a non-zero prime ideal of $k[x]$, then P factors in R into prime ideals. We know that P remains prime if f is not a square mod P, that P is ramified if $f \in P$, and P splits as $P = Q_1Q_2$ if f is a square mod P. Now if $P = (g)$ with g monic in $k[x]$, then quadratic reciprocity (which Artin has proved in his thesis) says, amongst other things, that if $q = 1$ mod 4 then f is a square mod g iff g is a square mod f! And if $q = 3$ mod 4 then there are easily-controlled sign issues. Anyway let's not get into this. Let χ be the function on prime ideals such that $\chi(P) = 0$ if $f \in P$, $\chi(P) = 1$ if f is a square mod P but not in P, and $\chi(P) = -1$ if f is not a square mod P. Then

$$Z_E(u) = (1-u)^{-1}\prod_{P\subset k[x]}(1-u^{d(P)})^{-1}(1-\chi(P)u^{d(P)})^{-1}$$

where $d(P)$ is the degree of P, and the product is over all prime ideals of $k[x]$. We can compute the product of the first terms by the Weil conjectures for $\mathbf{P}^1$ (it's $1/(1-qu)$), and if we do it in the 2nd way I explained (by expanding the product as a sum) we get a clue as to how to proceed with the other product. Indeed, same technique shows that the other product is $\sum_M \chi(M)u^{d(M)} = \sum_{m\geq 0}\sigma_m u^m$, with $\sigma_m = \sum_{M:d(M)=m}\chi(M)$. But now here's the bombshell: $\sigma_m = 0$ if $m \geq 3$! For by quadratic reciprocity $\chi(M)$ only depends on g mod f, where $M = (g)$ (at least if $q = 1$ mod 4; if it's 3 mod 4 then it also depends on the parity of the degree of g but this can be dealt with) and the sum over all non-zero residues of the character is 0.[1] 1

[1] I never checked this carefully.

The conclusion:

$$Z_E(u) = (1-u)^{-1}(1-qu)^{-1}(\sigma_0 + \sigma_1 u + \sigma_2 u^2)$$

and in particular it's a rational function. And of course for various explicit q and f one can compute these things, and check the functional equation and the Riemann hypothesis!

3 What are these things supposed to satisfy?

The functional equation says that $\sigma_{2g-i} = q^{g-i}\sigma_i$, or equivalently that for E hyperelliptic, $u^{1-g}Z_E(u)$ is invariant under $u \mapsto 1/qu$, or equivalently that $q^{(g-1)s}\zeta_E(s)$ is invariant under $s \mapsto 1-s$. Artin could prove this (although I don't know how he did it) in the hyperelliptic case.

The Riemann hypothesis says that all the roots of the polynomial $L(u)$ (which incidentally is basically $L(\chi, s)$ isn't it) have absolute value $|q|^{1/2}$, which of course can be checked in some cases. Equivalently, $\zeta_E(s)$ has all its zeros on the line $\mathrm{Re}(s) = 1/2$. Artin checked this sort of thing explicitly for certain explicit hyperelliptic curves over certain explicit finite fields.

Note also that (again for any hyperelliptic curve) it now follows easily that the number of points on the affine curve E is $q + \sigma_1$, because

$$\sigma_1 = \sum_{\lambda \in k, f = x - \lambda} \chi(f)$$

and $\chi(f) = 1$ if $f(\lambda)$ is a non-zero square, 0 if it's zero and -1 if it's a non-zero non-square.

Here's a clever observation of Artin though: σ_1 is related to the number of points, but it's also related to the sum of the roots of the polynomial. Now beefing the base (finite) field up by an extension of degree n changes the polynomial on the numerator to one whose roots are nth powers of the old one! We saw this already, when considering what happened to the zeta function under this change. So in fact if one can prove $|\sigma_1| \leq 2g\sqrt{q}$ (or indeed any bound of the form $O(\sqrt{q})$ as long as the implied constant is absolute, that is, independent of the base extension) one has proved the Riemann hypothesis for hyperelliptic curves!

3.1 Work of F. K. Schmidt.

Schmidt was the first person to think of projective curves rather than affine curves (although I wrote all the above for projective curves; Artin's work was in the affine case which makes things a bit messier). Schmidt also realised that Riemann-Roch could be used in place of quadratic reciprocity to make life a lot easier. Schmidt worked with an arbitrary function field. Here for example is a proof of rationality of the zeta function of a mooth projective curve in full generality. Say k is finite and V/k is smooth projective of genus g. We use the usual trick of pulling off a divisor at infinity and expanding over an affine (a Dedekind domain, of course). We have

$$Z_V(u) = R(u) \prod_{v \in |V_0|} (1 - u^{d(v)})^{-1}$$

where $R(x)$ is an explicit rational function (coming from points at infinity) and V_0 is affine and $|V_0|$ is its closed points. Hence we have

$$Z_V(u) = R(u) \sum_{D \geq 0} u^{d(D)} = R(u) \sum_{m \geq 0} \sigma_m u^m.$$

Here σ_m is the number of effective divisors of degree m. The point however is that instead of counting each effective divisor of degree m we may as well fix a random divisor D of degree m (let's assume one exists: we'll need to check this but we'll do it later, the point being that perhaps there are no effective divisors of degree 1) and then sum over C in the class group (degree zero divisors modulo equivalence) of the size of $H^0(C + D)$ (minus 1 because of zero). Now for

m sufficiently large, the dimension of $H^0(C+D)$ is simply $m+1-g$ (independent of C) by Riemann-Roch! Hence

$$\sum_{m\geq 2g} \sigma_m u^m = \sum_{m\geq 2g} h(q^{m+1-g}-1)u^m$$

from which it follows easily that the zeta function is rational.

Finally a trick: if actually there were no zero cycle of degree 1 then the sum above would be only over those m congruent to zero modulo some $r \geq 2$. However one can still prove that the ζ function has a simple pole at $s=1$ (because $R(u)$ is holomorphic there and the sum above is a GP—one just bashes it out[2]. Bashing out it all explicitly one checks that the base change to k_r of 2
the curve now has a pole of order r at $s=1$, a contradiction! Hence $r=1$ after all. We've proved a smooth projective curve over a finite field has a zero cycle of degree 1 using analysis!

Furthermore we only used the soft part of RR, the computation of the dimension for sufficiently large degrees. If we use the full RR then this relates $H^0(C+D)$ to $H^0(K-C-D)$ with K canonical (of degree $2g-2$) and explicitly bashing it out as in Eichler's book gives us the functional equation![3]. 3

That was what I wanted to understand about curves.

4 Higher-dimensional stuff.

If V is smooth and projective of dimension d then the zeta function has a contribution from cohomology in degree i for $0 \leq i \leq 2d$. The odd degrees are on the numerator. There are now possible zeros at $\mathrm{Re}(s) = q^{i/2}$ for $1 \leq i \leq 2d-1$ odd. In particular the Riemann hypothesis needs to be rephrased as some sort of purity statement; one might have zeros at other places in the higher-dimensional case. Compare this with the arithmetical case of projective 1-space over $\mathbf{Z}$. The zeta function of this scheme, defined as a product over closed points, is $\prod_p \zeta_{\mathbf{P}^1/\mathbf{F}_p}(s)$ which is $\prod_p (1-p^{-s})^{-1}(1-p^{1-s})^{-1} = \zeta(s)\zeta(s-1)$ which indeed has non-trivial zeros, conjecturally, only for $\mathrm{Re}(s) = 1/2$ and $\mathrm{Re}(s) = 3/2$.

[2] I didn't do this.

[3] It is in Eichler "Introduction to the theory of algebraic numbers and functions", page 302; I'm too lazy to copy it out. It's completely elementary