Tamagawa numbers via nonabelian Poincaré duality

Jacob Lurie

2014 Young Topologists Meeting (Copenhagen)

Contents

1 Quadratic forms and the Siegel mass formula

1.1 Quadratic forms, the counting problem, and the Siegel mass formula

Definition 1. A *quadratic form* is a homogeneous, degree-2 polynomial in n variables.

Example 2. The polynomials $x^2 + y^2$, $x^2 - y^2$, and $-x^2 - y^2$ are all quadratic forms. Note that these can be defined over *any* ring: all we need is are the coefficients ± 1.

Of course, quadratic forms are defined with respect to a set of coordinates, but we are more interested in the underlying objects (namely, quadric varieties), which may be rendered equivalent by *changes* of coordinates. We are therefore naturally led to ask the following question.

Question 3. Are these quadratic forms related to each other by change of basis?

Answer 4. It depends on the base ring. Over $\mathbb{C}$ the answer is yes, but over $\mathbb{R}$ the answer is no. Notice that this implies that over $\mathbb{Z}$, the answer is no too.

Note that studying $\mathbb{R}$ here tells us things about $\mathbb{Z}$. In fact, quadratic forms over $\mathbb{R}$ (in a given number of variables) are completely classified by their *signature*: we can always change our coordinates so that the quadratic form is given by $\sum a_i \cdot x_i^2$, and then the signature is by definition the integer

$$|\{i : a_i > 0\}| - |\{i : a_i < 0\}|.$$

But over $\mathbb{Z}$, things are much more subtle. For instance, as we just saw, $x^2 + y^2$ and $x^2 + 3y^2$ are equivalent over $\mathbb{R}$, but they can't be equivalent over $\mathbb{Z}$ because they're not even equivalent over $\mathbb{Z}/3$. (In fact, these will be equivalent over a given ring R iff $\sqrt{3} \in R$.) In any case, these tests motivate the following conjecture.

Conjecture 5. *Two quadratic forms over $\mathbb{Z}$ are equivalent iff they're equivalent over $\mathbb{R}$ and over all $\mathbb{Z}/p$.*

For simplicity, let us restrict to the *positive-definite* quadratic forms (still in some fixed number of variables n). It will be convenient to make the following definition.

Definition 6. We say that two (positive-definite) quadratic forms q and q' are *in the same genus* if they are equivalent over $\mathbb{Z}/N$ for all $N > 0$.

Now, it's not true that equivalence classes of quadratic forms are totally classified by their genus, but it turns out that each genus only contains *finitely many* equivalence classes, and moreover we can precisely enumerate these equivalence classes, provided we count with multiplicity in the right way: this formula is called the *Siegel mass formula*. To describe it, we'll need a bit of notation.

Notation 7. Let R be a ring, and let $q(x_1, \ldots, x_n)$ be a quadratic form with coefficients in R. Then, we write

$$O_q(R) = \{A \in GL_n(R) : q = q \circ A\}.$$

Example 8. When $R = \mathbb{R}$ and $q(x_1, \ldots, x_n) = \sum x_i^2$, then $O_q(R)$ is the usual orthogonal group.

Observation 9. If q is positive-definite, then $O_q(\mathbb{Z})$ will be *finite*. This is because its elements will have to preserve a lattice. (There will be some norm for which the lattice is spanned (at least rationally) by all the elements whose norm is bounded by some fixed constant C. So $O_q(\mathbb{Z})$ must act faithfully on these elements, but by positive-definiteness there are only finitely many of them.)

Definition 10. We define the *mass* of the quadratic form q to be

$$\mathrm{mass}(q) = \sum_{[q'] \in \mathrm{genus}(q)} \frac{1}{|O_{q'}(\mathbb{Z})|},$$

where $[q']$ denotes the equivalence class of the quadratic form q'.

We won't give a precise statement of the Siegel mass formula just yet, but suffice it to say that it gives us another expression for this same quantity.

Definition 11. We say that a quadratic form (over $\mathbb{Z}$) is *unimodular* if it's nondegenerate mod p for all primes p.

Example 12. None of the quadratic forms from our original example are unimodular. For instance, mod 2 we have that $x^2 + y^2 \equiv (x + y)^2$, i.e. up to change of coordinates this has just a single variable appearing with a nonzero coefficient.

It turns out that in order for a quadratic form to be unimodular, we must have at least 8 variables. In fact, a positive-definite quadratic form can *only* be unimodular if its number of variables is a multiple of 8.

We can now formulate a special case (or rather a consequence) of the Siegel mass formula:

$$\sum_{q \text{ unimodular}} \frac{1}{|O_q(\mathbb{Z})|} = \frac{\Gamma(1/2) \cdot \Gamma(2/2) \cdots \Gamma(n/2) \cdot \zeta(2) \cdot \zeta(4) \cdots \zeta(n-2) \cdot \zeta(n/2)}{2^{n-1} \pi^{(n^2+n)/4}}.$$

This is totally crazy: a priori the right side isn't even a rational number, let alone related to the left side! One can at least see that it's rational using Bernoulli numbers, but the connection is still mysterious.

Example 13. Let's look at the simplest example of this consequence of the Siegel mass formula. Let's use n to denote the number of variables. So the simplest case is $n = 8$, and then the right side of the equation becomes

$$\frac{1}{2^{14} \cdot 3^5 \cdot 5^2 \cdot 7}.$$

This denominator is actually the order of a certain finite group: if we set q to be the quadratic form corresponding to the E_8-lattice, then $O_q(\mathbb{Z})$ is the Weyl group of E_8, and we have that $|O_q(\mathbb{Z})| = 2^{14} \cdot 3^5 \cdot 5^2 \cdot 7$. But then, since on the left side we're only adding positive quantities and we've already reached our asserted sum, this implies that there are no other summands on the left side: in other words, the E_8 quadratic form is the *unique* unimodular quadratic form in 8 variables.

Remark 14. This example might be misleading, because it might lead one to expect that these numbers are often very small. In fact, at $n = 32$ the right side is about 40,000,000. This tells us that there are *tons* of 32-variable unimodular quadratic forms. (In fact, every quadratic form has at least one automorphism (given by negating all coefficients), so there are at least 80 million unimodular quadratic forms!) As it happens, there are only 2 unimodular quadratic forms in 16 variables and only 24 unimodular quadratic forms in 24 variables (and these have all been classified). But we're *never* going to classify all the 32-variable unimodular quadratic forms by hand, in which sense this result (the Siegel mass formula) is pretty spectacular.

1.2 A reformulation of the counting problem

Now, let's try – and fail – to prove that "same genus" implies "equivalent"; by failing, we'll learn something.

So, suppose that q and q' are in the same genus, i.e. for all $N > 0$ they're equivalent mod N. That is, for each N we have some $A_N \in GL_n(\mathbb{Z}/N\mathbb{Z})$ such that $q = q' \circ A_N$ (as quadratic forms over $\mathbb{Z}/N$).

We claim that without loss of generality, we can assume that if $N|N'$ then $A_{N'}$ reduces to A_N mod N. (This is a simple compactness argument, relying on the fact that there are only finitely many choices at each stage, since $GL_n(\mathbb{Z}/N)$ is finite.) So, we can think about all the A_N together as some single matrix $A \in GL_n(\hat{\mathbb{Z}})$.

Now, recall that

$$\hat{\mathbb{Z}} = \lim \mathbb{Z}/N = \prod_p \mathbb{Z}_p$$

(by the Chinese remainder theorem), so we can actually assume that $q \sim q'$ over $\mathbb{Z}_p$ for all primes p. This implies that $q \sim q'$ over $\mathbb{Q}_p$ for all p (by the factorization

$$\mathbb{Z} \to \mathbb{Z}_p \to \mathbb{Q}_p = \mathbb{Z}_p[p^{-1}]$$

of the unit). Moreover, we have $q \sim q'$ over $\mathbb{R}$, since quadratic forms over $\mathbb{R}$ are classified by their signature. Then, in line with the *Hasse principle* (which in general says that two objects are equivalent over $\mathbb{Q}$ iff they're equivalent over all its completions, i.e. over $\mathbb{R}$ and over all the $\mathbb{Q}_p$), this turns out to imply that $q \sim q'$ over $\mathbb{Q}$, i.e. that there is some $B \in GL_n(\mathbb{Q})$ such that $q = q' \circ B$.

Now, we can use A and B to get an automorphism of q, as

$$q = q' \circ A = (q \circ B^{-1}) \circ A.$$

This is tantamount to saying that $B^{-1}A \in O_q(\mathbb{A}_{fin})$, where $\mathbb{A}_{fin}$ is the ring of *finite adeles*: this can be described as $\mathbb{A}_{fin} = \hat{\mathbb{Z}} \otimes \mathbb{Q}$, i.e.

$$\mathbb{A}_{fin} = \prod_p{}^{res} \mathbb{Q}_p$$

(where the "res" stands for "restricted": we only allow denominators at finitely many places).

Now, to come up with this matrix we made two choices: the choice of A, and the choice of B. hence, to obtain a canonical value we consider

$$B^{-1} \circ A \in O_q(\mathbb{Q}) \backslash O_q(\mathbb{A}_{fin}) / O_q(\hat{\mathbb{Z}}),$$

i.e. we take this as a double-coset. If this were the identity double-coset, then we'd have that $B^{-1} \circ A$ is the identity matrix (so that $A = B$), and moreover since $\hat{\mathbb{Z}} \cap \mathbb{Q} = \mathbb{Z}$ we get that in fact A and B are defined over $\mathbb{Z}$.

In fact, this logic is all reversible, and we come to the following result.

Proposition 15. *The set*

$$O_q(\mathbb{Q})\backslash O_q(\mathbb{A}_{fin})/O_q(\hat{\mathbb{Z}})$$

is in bijection with equivalence classes of quadratic forms in the genus of q.

Before going forward, let's make two modifications to this setup.

1. Rather than work with O, we work with SO: define

   $$SO_q(R) = \{A \in GL_n(R) : q \circ A = q,\ \det(A) = 1\},$$

   and use SO everywhere above where we wrote O. This actually corresponds to a different counting problem (namely, modding out by equivalences via determinant-1 transformations), but this ends up only introducing some determinable power of 2. So, this is no big deal.

2. We'd like to replace $\mathbb{A}_{fin}$ with $\mathbb{A} = \mathbb{A}_{fin} \times \mathbb{R}$. This modification makes our middle terms $SO_q(\mathbb{A}_{fin})$ bigger, by adding in a factor of the compact Lie group $SO_q(\mathbb{R})$. To couteract this, we just divide out by it too.

And so finally, we will want to understand the size of the set

$$SO_q(\mathbb{Q})\backslash SO_q(\mathbb{A})/SO_q(\hat{\mathbb{Z}} \times \mathbb{R}).$$

(Note that $SO_q(\hat{\mathbb{Z}} \times \mathbb{R}) = SO_q(\hat{\mathbb{Z}}) \times SO_q(\mathbb{R})$.)

1.3 Measures

Now, we would also like to take advantage of the topology on $\mathbb{A}$, which is quite nice and makes it into a locally compact topological ring. Hence, there is an induced subspace topology on $SO_q(\mathbb{A}) \subset M_{n\times n}(\mathbb{A})$, which then becomes a locally compact group. Inside of this locally compact group sit the discrete subgroup

$$SO_q(\mathbb{Q}) \subset SO_q(\mathbb{A})$$

and compact open subgroup

$$SO_q(\mathbb{A}) \supset SO_q(\hat{\mathbb{Z}} \times \mathbb{R}).$$

Now, it's always fun to look for invariant measures. Recall that the *Haar measure* is the unique left-invariant measure up to scaling. (In fact, this is a *unimodular* group, meaning left-invariant measures are also right-invariant measure.) Let's write μ for a Haar measure (which, again, is well-defined up to a scalar). This descends to a measure on $SO_q(\mathbb{Q})\backslash SO_q(\mathbb{A})$, which we again denote by μ.

Now, the compact group $SO_q(\hat{\mathbb{Z}} \times \mathbb{R})$ acts on the set $SO_q(\mathbb{Q})\backslash SO_q(\mathbb{A})$ on the right; if this action were *free*, we'd just divide things out to get the resulting measure of the two-sided quotient. But in fact, there's a stabilizer: the stabilizer of the identity left-coset is precisely $SO_q(\mathbb{Z})$. Hence, we obtain that

$$\mu(SO_q(\mathbb{Q})\backslash SO_q(\mathbb{A})) = \mu(SO_q(\hat{\mathbb{Z}} \times \mathbb{R})) \cdot \sum_{q' \in \text{genus}(q)} \frac{1}{|SO_{q'}(\mathbb{Z})|}.$$

And this, finally, is where the mass formula comes from:

$$\text{mass}(q) = \frac{\mu(SO_q(\mathbb{Q})\backslash SO_q(\mathbb{A}))}{\mu(SO_q(\hat{\mathbb{Z}} \times \mathbb{R}))}.$$

Remark 16. This expression is independent of the choice of μ, since μ itself was well-defined up to scaling. But in fact, Tamagawa and Weil saw that it's actually fruitful to consider these numerator and denominator independently (and make a canonical choice of measure). We'll pick up this thread in the next lecture.

2 Weil's conjecture

2.1 The Siegel mass formula and Weil's conjecture

Let q be a positive-definite quadratic form over $\mathbb{Z}$. Recall that the *Siegel mass formula* tells us that

$$\sum_{q' \in \text{genus}(q)} \frac{1}{|SO_{q'}(\mathbb{Z})|} = \frac{\mu(SO_q(\mathbb{Q})\backslash SO_q(\mathbb{A}))}{\mu(SO_q(\hat{\mathbb{Z}} \times \mathbb{R}))},$$

where $\mathbb{A} = \mathbb{R} \times \prod_p^{res} \mathbb{Q}_p$. Observe that this expression on the right is independent of the choice of measure μ, since μ itself is well-defined up to scale. But in fact, there's a canonical choice of measure, the *Tamagawa measure*, which we will now define. Actually, we will define the measure on each factor of the canonical factorization

$$SO_q(\mathbb{A}) = SO_q(\mathbb{R}) \times \prod_p^{res} SO_q(\mathbb{Q}_p).$$

individually.

- First of all, $SO_q(\mathbb{R})$ is a Lie group, and in particular is a manifold. Thus, top-degree differential forms give rise to measures (by integration); if our chosen top-degree form is left-invariant, then so will the resulting measure be left invariant. Left-invariant top-degree forms, in turn, are determined by what they do at the origin. Let's write $V_{\mathbb{R}}$ for the 1-dimensional $\mathbb{R}$-vector space that they span.

  Next, note that $SO_q(\mathbb{R})$ isn't just a compact group, but comes to us as an *algebraic group*: it's a subgroup $SO_q(\mathbb{R}) \subset GL_n(\mathbb{R})$ cut out by certain *polynomial* equations. In general, such an algebraic group is called a *linear algebraic group.*

  Now, this linear algebraic group is not just defined over $\mathbb{R}$, but over $\mathbb{Q}$. we can talk about *algebraic* differential forms defined over $\mathbb{Q}$, too. This gives us a canonical choice of subspace $V_{\mathbb{Q}} \subset V_{\mathbb{R}}$, which is a 1-dimensional $\mathbb{Q}$-vector space: $V_{\mathbb{R}} = V_{\mathbb{Q}} \otimes_{\mathbb{Q}} \mathbb{R}$. This cuts down a huge amount of ambiguity, and we'll leave this thread here for the moment.

- Let's consider the other factors, $SO_q(\mathbb{Q}_p)$ for p prime. This is a *p-adic analytic Lie group*, and so one can say roughly the same words to obtain a function from a 1-dimensional $\mathbb{Q}_p$-vector space $V_{\mathbb{Q}_p}$ to the invariant measures. But once again, this comes to us as a linear algebraic group (this time as a subgroup of $GL_n(\mathbb{Q}_p)$), so again we canonically get $V_{\mathbb{Q}} \subset V_{\mathbb{Q}_p}$. We use the same notation $V_{\mathbb{Q}}$ because this is the same vector space: these conditions are defined by the same equations!

As a result of the previous two analyses, we see that any nonzero $\omega \in V_{\mathbb{Q}}$ determines a left-invariant measure $\mu_{\omega,p}$ on $SO_q(\mathbb{Q}_p)$ as well as a left-invariant measure $\mu_{\omega,\infty}$ on $SO_q(\mathbb{R})$. Using this observation, we define the *Tamagawa measure* on $SO_q(\mathbb{A})$ by taking the product:

$$\mu_{\text{Tam}} = \prod_{\text{places } p} \mu_{\omega,p}.$$

Now, this may look like it depended on the choice of $\omega \in V_{\mathbb{Q}}$, but in fact it does not. If we multiply ω by some rational number r, then the $\mathbb{R}$-factor $\mu_{\omega,\infty}$ would get multiplied by $|r|_\infty$, while the $\mathbb{Q}_p$-factor $\mu_{\omega,p}$ would get multiplied by $|r|_p$. The magic, then, lies in the fact that the product of the absolute values at all the places of *any* nonzero rational number is 1!

Now, let's return to the Siegel mass formula, and use our new canonical measure, the Tamagawa measure. The denominator on the right side is the measure of

$$SO_q(\hat{\mathbb{Z}} \times \mathbb{R}) = SO_q(\mathbb{R}) \times \prod_p SO_q(\mathbb{Z}_p),$$

and by definition we have that

$$\mu_{\text{Tam}}(SO_q(\hat{\mathbb{Z}} \times \mathbb{R})) = \mu_{\omega,\infty}(SO_q(\mathbb{R})) \times \prod_p \mu_{\omega,p}(SO_q(\mathbb{Q}_p)).$$

The first factor is what's contributing all the π's an special values of Γ and whatnot, while the second factor involves a counting problem: the number of $\mathbb{Z}/p$-points of $SO_q(\mathbb{Z}_p)$.

In fact, it turns out that in the statement of the Siegel mass formula, at least in the unimodular case that we saw in the last talk, *all* of the complication came from the denominator. There's another reformulation of the Siegel mass formula, only in terms of the numerator:

$$\mu_{\mathrm{Tam}}(SO_q(\mathbb{Q})\backslash SO_q(\mathbb{A})) = 2.$$

(Going back and forth between these formulations is substantially easier than proving either one by themselves; that's the sense in which this is a reformulation.) But this has another advantage: last time, we only talked about the unimodular version. But this newer statement is good for *any* quadratic form; in fact, it works over $\mathbb{Q}$ and we don't need to restrict to the positive-definite case either.

But one still might not like this formula completely: why 2? In fact, this 2 is related to the fact that SO_q is not simply connected: it has an (algebraic) double cover by an algebraic group, $\mathrm{Spin}_q \to SO_q$. Then, we have the further reformulation

$$\mu_{\mathrm{Tam}}(\mathrm{Spin}_Q(\mathbb{Q})\backslash \mathrm{Spin}_q(\mathbb{A})) = 1$$

of the Siegel mass formula (in the same sense as before).

Now, motivated by this and other examples, Weil conjectured that this mass formula is a *general* phenomenon, not only have to do with quadratic forms but rather with *all* simply-connected algebraic groups. More precisely, he formulated the following.

Conjecture 17 (Weil's conjecture)**.** *Let G be a semisimple, simply-connected algebraic group over $\mathbb{Q}$. Then the Tamagawa measure on $G(\mathbb{A})$ has that*

$$\mu_{\mathrm{Tam}}(G(\mathbb{Q})\backslash G(\mathbb{A})) = 1.$$

Remark 18. In this case, we can take the semisimple and simply-connected assumptions to mean that over $\mathbb{C}$ it is such. As for the Tamagawa measure, one can make all the definitions we've seen above to give a rigorous definition.

Remark 19. If G isn't compact over $\mathbb{R}$, one can reformulate the conjecture as a "density" statement. (The "mass" comes from dividing off $G(\hat{\mathbb{Z}})$ [**this doesn't make sense to me**]).

Weil's conjecture was first checked by Weil himself in a number of cases. Later, Langlands and Lai verified it in a number of other cases. Finally, Kottwitz proved it in general by reducing to those cases.

2.2 Weil's conjecture for function fields

Now, we would like to talk about the (positive-characteristic) *function field* analog of Weil's conjecture. It's a general pattern in number theory that many questions one can ask over $\mathbb{Q}$ also make sense over function fields, and typically these are much easier to answer: we have tools that aren't available in number theory, namely the tools of *algebraic geometry.* This statement seems to be an exception: the number field case was known far earlier than the function field case.

We begin by fixing some notation.

Notation 20. We fix the following.

- Let $\mathbb{F}_q$ denote the finite field with q elements.
- Let X denote a smooth projective algebraic curve over $\mathbb{F}_q$. (One should think of X as something like a compact Riemann surface: this would be the case that we're over $\mathbb{C}$ instead of over $\mathbb{F}_q$.)
- Let K_X be the fraction field of X, i.e. its field of rational functions.

Example 21. The simplest example to keep in mind is $X = \mathbb{P}^1$; in this case, $K_X = \mathbb{F}_q(x)$ (the field obtained by taking $\mathbb{F}_q$ and freely adjoining a transcendental element x – one may actually *define* a function field to be a field which is obtained by this procedure).

Now, function fields are analogous to number fields in many ways, and these relationships organize themselves into a dictionary (two out of three columns of what has been called the *Rosetta stone for the Langlands program*).

number fields	**function fields**
$\mathbb{Q}$ (or any number field)	K_X
places (i.e. prime numbers and ∞)	closed points $x \in X$
$\mathbb{Z}/p$	residue field $\kappa(x)$ (a finite extension of $\mathbb{F}_q$)
$\mathbb{Z}_p$	complete local ring $\mathcal{O}_{X,x}$; noncan. iso. to $\kappa(x)[[t]]$
$\mathbb{Q}_p$ (or $\mathbb{R}$)	fraction field $K_{X,x}$ of $\mathcal{O}_{X,x}$; noncanon. iso. to $\kappa(x)((t))$
$\mathbb{A} = \mathbb{R} \times \prod_p{}^{res} \mathbb{Q}_p$	$\mathbb{A} = \prod^{res} K_{X,x}$

Remark 22. The noncanonical isomorphisms on the right side are determined by a choice of local coordinate.

Remark 23. The last line on each side gives an *exhaustive* list of the completions of the appropriate object, whence the adelic notation. In fact, it turns out that there are again natural topologies on the constituent factors of $\mathbb{A}$ and hence on $\mathbb{A}$ itself, and (on the function field side) this contains K_X as a discrete subring.

Now, let G_0 be a (linear) algebraic group defined over K_X (i.e. it's cut out of $GL_n(K_X)$ by polynomials), assumed to be semisimple and simply-connected (although note that this means something more general than the previous case, since we can't just extend to $\mathbb{C}$ anymore). We can take $G_0(\mathbb{A})$; this inherits the structure of a locally compact group, and then $G_0(K_X) \subset G_0(\mathbb{A})$ becomes a discrete subgroup. We can again discuss invariant measures, and an analogous procedure to the one above determines a measure on each $G_0(K_{X,x})$; once again, as long as one considers all the completions together, the measure will again be independent of the choice. Hence we get a well-defined canonical choice of Haar measure on $G_0(\mathbb{A})$, which we once again call the *Tamagawa measure*.

Conjecture 24 (Weil's conjecture for function fields)**.**

$$\mu_{\mathrm{Tam}}(G_0(K_X) \backslash G_0(\mathbb{A})) = 1.$$

2.3 Counting problems for the function field case

Now, this is a nice clean statement, but it's not as clear what sort of counting problem this should relate to. So, in the remaining time, we give a quick indication of how we can convert this clean statement into a counting problem.

The analogy begins with the fact that having a linear algebraic group G_0 over K_X is parallel to asking about quadratic forms over $\mathbb{Q}$. But to get a mass formula, we need something better, namely an *integral structure*. For this, we observe that we have the inclusion $\mathrm{Spec}(K_X) \to X$ of the generic point, and so we further assume is that we have a pullback diagram

$$\begin{tikzcd} G_0 \arrow[r] \arrow[d] & G \arrow[d] \\ \mathrm{Spec}(K_X) \arrow[r] & X \end{tikzcd}$$

for some smooth affine group scheme G over X.

Now, with this additional structure, we can apply G_0 to not just any ring containing K_X (i.e. things with maps to $\mathrm{Spec}(K_X)$), but we can also apply G to objects living over X. For instance, we have $G(\mathcal{O}_{X,x}) \subset G_0(K_{X,x})$ (a compact open subgroup of a locally compact group), and then we have a map $G(\mathcal{O}_{X,x}) \to G(\kappa(x))$ to a finite group.

Now, we can ask not just about left cosets $G_0(\mathbb{Q})\backslash G_0(\mathbb{A})$, but about double-cosets

$$G_0(K_X)\backslash G_0(\mathbb{A})/ \prod_{x\in X} G(\mathcal{O}_{X,x}),$$

and it turns out that the mass formula counts these double-cosets. In the next lecture, we will describe how we recover this counting problem from the above formulation.

Remark 25. This set of double-cosets can be identified with $H^1(X; G)$, in analogy with our previous discussion of "forms of quadratic forms that locally look the same at all points of $\mathrm{Spec}(\mathbb{Z})$".

3 The geometry of Weil's conjecture for function fields

3.1 More on counting problems

Recall that if X is an algebraic curve over $\mathbb{F}_q$ and $x \in X$, we set the following notation:

- K_X is the function field of X,
- $\kappa(x)$ is the residue field at x (which will be a finite extension of $\mathbb{F}_q$),
- $\mathcal{O}_x = \mathcal{O}_{X,x} \cong \kappa(x)[[t]]$ is the ring of germs of functions at x (where the isomorphism comes only after a choice of coordinate),
- $K_x = K_{X,x} \cong \kappa(x)((t))$ is the fraction field of $\mathcal{O}_x$, and
- $\mathbb{A} = \prod^{res} K_x$ is the restricted product of the fraction fields.

Then, Weil's conjecture for function fields asserts that

$$\mu_{\mathrm{Tam}}(G_0(K_X)\backslash G_0(\mathbb{A})) = 1.$$

As we indicated last time, the first step in the proof is to reduce this to something more concrete, namely to a *counting problem*. In order to do this, we also ask for a *deformation* of G_0: more precisely, we ask for a pullback diagram

$$\begin{tikzcd} G_0 \arrow[r] \arrow[d] & G \arrow[d] \\ \mathrm{Spec}(K_X) \arrow[r] & X. \end{tikzcd}$$

in which G is a smooth group scheme which is affine over X and has connected fibers. Then, whereas we previously already had a left action of $G(K_X) = G_0(K_X)$ on $G(\mathbb{A}) = G_0(\mathbb{A})$, but now we also get a right action on $G(\mathbb{A})$ of $\prod_{x\in X} G(\mathcal{O}_x)$.

Question 26. What do the double-cosets

$$G(K_X)\backslash G(\mathbb{A})/ \prod_{x\in X} G(\mathcal{O}_x)$$

count?

Answer 27. These double-cosets count *isomorphism classes of principal G-bundles* on X.

In order to understand this better, we make two observations.

1. Any principal G-bundle P over X can be trivialized at any (closed) point of X. (This is non-obvious, since the residue fields aren't generally algebraically closed and so a priori there should be obstructions to triviality in its Galois cohomology. But this fact follows from a theorem of Lang, which says that there are no nontrivial G-bundles over a finite field whenever the fibers are connected.) So, P is trivial over $\mathcal{O}_x$ for all $x \in X$.

2. In line with the Hasse principle, this implies that P is also trivial over K_X. (This statement is due to Harder.)

So, we know that we can trivialize P away from a finite set of points. But we can also trivialize it in a formal neighborhood of each of these points, and hence all that we need to keep track of to recover P itself is the gluing data. If the formal disk is associated to the local ring $\mathcal{O}_x$, then the *punctured* formal disk is associated to the field K_x. Hence, if we consider

$$[P] \in \prod_{x \in X}^{res} G(K_x) = G(\mathbb{A}),$$

then the ambiguity in the choice of representative P of $[P]$ is precisely the left-action of $G(K_x)$ and the right action of $\prod_{x \in X} G(\mathcal{O}_x)$. And this recovers for us our answer to the counting problem.

3.2 Stacks

Now, Weil's conjecture has to do with the value $\mu_{\mathrm{Tam}}(G_0(K_x)\backslash G_0(\mathbb{A}))$, and this is associated with a *weighted* count:

$$\frac{\mu_{\mathrm{Tam}}(G(K_x)\backslash G(\mathbb{A}))}{\mu_{\mathrm{Tam}}\left(\prod_{x \in X} G(\mathcal{O}_x)\right)} = \sum_{[P]} \frac{1}{|\mathrm{Aut}(P)|}$$

(where the sum on the right is taken over the double-cosets, i.e. the isomorphism classes of G-bundles P over X), and this is the appropriate Siegel mass formula. Just to totally spell out the analogy, the left side is conjecturally

$$\frac{1}{\mu_{\mathrm{Tam}}\left(\prod_{x \in X} G(\mathcal{O}_x)\right)},$$

and so the right side is counting the number of G-bundles and the left side factors as a product (since the Tamagawa measure is a *product* measure) whose factors are computable.

Remark 28. Actually, it turns out that unless G is trivial, the sum on the right side is infinite: there are infinitely many G-bundles. but luckily, their automorphism groups grow quite quickly, and this sum ends up converging (but only because we were clever enough to count with multiplicity).

Number theory over function fields is supposed to be *easier* than over global fields, and we now take advantage of this. Namely, G-bundles are defined in an algebro-geometric way, and admit an algebro-geometric parametrization: they come in *families.* if Y is an algebraic variety over $\mathbb{F}_q$, then we can think of G-bundles on $X \times Y$ as "families of G-bundles on X which are parametrized by Y". It will be convenient to introduce a bit of notation for encoding this.

Notation 29. Let us write $\mathrm{Bun}_G(X)$ for the *(algebraic) stack of G-bundles of X*: this is an object which is characterized by the property that giving a map $Y \to \mathrm{Bun}_G(X)$ is the same as giving a G-bundle on $X \times Y$.

Remark 30. This object $\mathrm{Bun}_G(X)$ is not so unlike an algebraic variety, except that it's a *stack*: the objects we're trying to classify (namely, G-bundles on X) form a *groupoid* instead of just a set. (One may alternatively set this up by considering the projection map $BG \times X \to X$; then, $\mathrm{Bun}_G(X)$ parametrizes sections of this map.)

Now, in particular we can evaluate $\mathrm{Bun}_G(X)$ on $\mathbb{F}_q$-algebras R, and we have that

$$\mathrm{Bun}_G(X)(R) = \{G\text{-bundles on } X \times \mathrm{Spec}(R)\}.$$

Using $|-|$ to denote the *groupoid Euler characteristic* (determined by $|\mathrm{pt}//G| = 1/|G|$), in particular we have that

$$|\mathrm{Bun}_G(X)(\mathbb{F}_q)| = |\{G\text{-bundles on } X\}| = \sum_{[P]} \frac{1}{|\mathrm{Aut}(P)|},$$

the quantity appearing in Weil's conjecture.

3.3 Trace formulas

Our identification of the quantity as a weighted count of isomorphism classes of G-bundles raises the following related question.

Question 31. Given a variety Y over $\mathbb{F}_q$, how many points does Y have over $\mathbb{F}_q$?

Weil's idea for attacking this question is the following. First of all, we observe that $Y(\mathbb{F}_q) \subset Y(\overline{\mathbb{F}}_q)$; moreover, since Y itself is defined over $\mathbb{F}_q$, then it admits a *(geometric) Frobenius* map $Y \to Y$, given by raising all coordinates to the q^{th} power. (This is because if $f(x_1, \ldots, x_n) = 0$ then $f(x_1^q, \ldots, x_n^q) = 0$ as long as the coefficients of f are all in the ground field $\mathbb{F}_q$.) Using this, we get that this inclusion $Y(\mathbb{F}_q) \subset Y(\overline{\mathbb{F}}_q)$ is precisely the inclusion of the fixedpoints of the Frobenius. Hence, we should have

$$|Y(\mathbb{F}_q)| = \sum_i (-1)^i \mathrm{Tr}(\mathrm{Frob}|H_c^i(Y)),$$

where H_c^i denotes some *conjectural* "compactly-supported cohomology theory for algebraic varieties" with good properties – that is, a *Weil cohomology theory.* In fact, this was one of the great successes of the Grothendieck school of algebraic geometry: they implemented this idea by introducing *ℓ-adic cohomology.* Using this, they obtained the *Grothendieck–Lefschetz trace formula.*

Now, we would like to apply this trace formula not to an algebraic variety Y, but to our algebraic stack $\mathrm{Bun}_G(X)$. But first, we will need to give a reformulation of the trace formula. If Y is a smooth variety of dimension d, then its cohomology satisfies Poincaré duality: $H_c^i(Y, \mathbb{Q}_\ell) \cong (H^{2d-i}(Y, \mathbb{Q}_\ell))^\vee$. We can almost use this to directly rewrite the trace formula, except that this isomorphism isn't equivariant for the action of Frobenius (but rather we always get an extra factor of q), and so the dual form of the trace formula is that

$$\frac{|Y(\mathbb{F}_q)|}{q^d} = \sum_i (-1)^i \mathrm{Tr}(\mathrm{Frob}^{-1}|H^i(Y; \mathbb{Q}_\ell)).$$

Remark 32. This trace formula reflects the heuristic count that a dimension-d variety over $\mathbb{F}_q$ "should" have roughly q^d points, since this is true of affine space, the simplest example. So we should expect this sum to be roughly 1.

Now, we would like to apply this dual formulation of the trace formula to the case where $Y = \mathrm{Bun}_G(X)$: we would like to say that

$$\frac{|\mathrm{Bun}_G(X)(\mathbb{F}_q)}{q^{\dim \mathrm{Bun}_G(X)}} = \sum_i (-1)^i \mathrm{Tr}(\mathrm{Frob}^{-1}|H^i(\mathrm{Bun}_G(X), \mathbb{Q}_\ell)).$$

But the important difference here is that a priori, we expect neither the left side nor the right side to be well-defined. For instance, $|\mathrm{Bun}_G(X)|$ might be given by an infinite sum, and similarly $H^i(\mathrm{Bun}_G(X), \mathbb{Q}_\ell)$ might be nonzero for infinitely many i. So in fact, what we mean to assert is first of all that both of these quantities *are* well-defined, and then moreover that they agree. Luckily, this is a theorem of Behrend, at least if G has good reduction everywhere (and the proof can be generalized with some effort to our case too).

And here we see the entrance of topology. We want to understand the cohomology of $\mathrm{Bun}_G(X)$, including its action of Frobenius. thus, we reformulate Weil's conjecture (after some computations, including computing the "computable" value $\mu_{\mathrm{Tam}}\left(\prod_{x\in X} G(\mathcal{O}_x)\right)$ of the mass formula) as saying that

$$\sum_i (-1)^i \mathrm{Tr}(\mathrm{Frob}^{-1}|H^i(\mathrm{Bun}_G(X), \mathbb{Q}_\ell)) = \prod_{x \in X} \frac{|K(x)|^{\dim(G)}}{|G(K(x))|},$$

where this fraction should be thought of as having numerator "the expected number of points of G over $K(x)$" and as having denominator "the actual number of points of G over $K(x)$", and hence we expect that this fraction should be close to 1. To prove this result, we'll first show that this sum on the left side ought to factor as a product over the points $x \in X$, and then we'll show that each factor should take the form indicated on the right.

4 Nonabelian Poincaré duality

4.1 An overview of the proof

Recall that we have the quantity

$$\frac{\sum_{[P]} \frac{1}{|\mathrm{Aut}(P)|}}{q^{\dim \mathrm{Bun}_G(X)}} = \frac{|\mathrm{Bun}_G(X)(\mathbb{F}_q)|}{q^{\dim \mathrm{Bun}_G(X)}}$$

and the quantity

$$\prod_{x \in X} \frac{|\kappa(x)|^{\dim G}}{|G(\kappa(x))|},$$

and we've restated Weil's conjecture for function fields to be the assertion that these quantities are equal. Recall also Lang's theorem, which says that there's only one G-bundle over $\kappa(x)$ – the trivial bundle – which therefore has automorphism group $G(\kappa(x))$. So, a weighted count of G-bundles on $\mathrm{Spec}(\kappa(x))$ yields $1/|G(\kappa(x))|$. Now, the dimension of BG is the *negative* of the dimension of G, and hence we get the equality

$$\frac{|\{G\text{-bundles over } \kappa(x)\}|}{|\kappa(x)|^{\dim BG}} = \frac{|\kappa(x)|^{\dim G}}{|G(\kappa(x))|}.$$

So, we can now describe our method of attack. On the one hand, we will have a *trace formula* telling us that the first quantity equals

$$\sum_i (-1)^i \mathrm{Tr}(\mathrm{Frob}^{-1} | H^*(\mathrm{Bun}_G(X), \mathbb{Q}_\ell)),$$

and then on the other hand a *local-to-global principle* will give us that this expression agrees with the second quantity as well. Thus, our general goal is to understand the (ℓ-adic) cohomology of $\mathrm{Bun}_G(X)$ over $\mathbb{F}_q$ (or over $\overline{\mathbb{F}}_q$).

4.2 A warm-up case and nonabelian Poincaré duality in topology

As a warm-up, let's first explore these ideas over the field $\mathbb{C}$ and taking G to be a *constant* group scheme, i.e. it's just a semisimple complex connected Lie group (e.g. $G = SL_n$). This lands us in the world of ordinary topology. So, we now think of X as a Riemann surface, and we will think of BG as simply a topological space (rather than as an algebraic stack). Then, $\mathrm{Bun}_G(X)$ is an algebraic stack but has an associated homotopy type, namely $\mathrm{map}(X, BG)$. (Because X is 1-dimensional, there ends up being not much difference between algebraic vector bundles and smooth vector bundles.)

More generally, if X is a d-dimensional manifold and Y is any pointed space (e.g. BG), we might try to describe $\mathrm{map}(X, Y)$ in terms of *local data* – or we'll actually care instead about $\mathrm{map}_c(X, Y)$ in the case that X is noncompact. So, the first fact to realize is that this functor is *covariant* in X with respect to proper maps: if $U \subset X$ is an open embedding, we even get an *embedding* of mapping spaces $\mathrm{map}_c(U, Y) \hookrightarrow \mathrm{map}_c(X, Y)$. As a consequence, we have an abundance of spaces mapping to $\mathrm{map}_c(X, Y)$, namely the spaces $\mathrm{map}_c(U, Y)$ for any $U \subset X$. This naturally leads us to consider the map

$$\mathrm{hocolim}_{U \subset X} \mathrm{map}_c(U, Y) \to \mathrm{map}_c(X, Y).$$

Of course, there's a terminal object of the category of open subsets of X, namely X itself, so this homotopy colimit isn't going to be so interesting – it'll just be constant at $\mathrm{map}_c(X, Y)$, and our map will be an equivalence trivially. But we can still try to get some mileage out of this idea, by restricting to certain subspaces U: if U is an open disk (i.e. an open subset of X that's homeomorphic to $\mathbb{R}^d$), then we will have that

$$\mathrm{map}_c(U, Y) \cong \mathrm{map}((D^n, S^{n-1}), (Y, y)) \simeq \Omega^d Y;$$

more generally, if $U \simeq \mathbb{R}^d \amalg \ldots \amalg \mathbb{R}^d$, then

$$\mathrm{map}_c(U, Y) \simeq \Omega^d Y \times \cdots \times \Omega^d Y.$$

So, on such open subsets we have a very good understanding of the compactly-supported maps into Y.

The miracle, then, is it turns out that this is all we need, at least when Y is sufficiently connected (otherwise taking its d-fold loopspace will entail a loss of information). More precisely we have the following result.

Theorem 33 (Topological nonabelian Poincaré duality)**.** *If Y is $(d-1)$-connected, then*

$$\mathrm{hocolim}_{U\subset X \text{ a disjoint union of } d\text{-disks}}\ \mathrm{map}_c(U,Y) \xrightarrow{\sim} \mathrm{map}_c(X,Y).$$

Remark 34. This theorem is so named for the following reason. If $Y = K(A,n)$ for $n \geq d$ and X is oriented, then the source and target of the duality map are both describable in terms of singular co/chains: the target has homotopy groups

$$\pi_* \mathrm{map}_c(X, K(A,n)) \cong H_c^{n-*}(X;A),$$

while the source has homotopy groups

$$\pi_* \mathrm{hocolim}\ \mathrm{map}_c(U, K(A,n)) \cong H_{*+d-n}(X;A),$$

and indeed the natural map from the latter to the former is the Poincar'e duality map, i.e. the inverse of the map

$$H_c^{n-*}(X;A) \xrightarrow{-\cap[X]} H_{*+d-n}(X;A)$$

(the cap-product).

This statement is good for our purposes, because it gives us a local-to-global principle: we can write $\mathrm{map}_c(X,Y)$ as a (homotopy) colimit, in which

- on the one hand, the category over which we're taking a colimit depends only on X (namely, the category of disjoint unions of d-disks in X), while
- on the other hand, the value of our functor at each object depends only on Y.

Remark 35. In the special case when $Y = BG$, we should think of a map in $\mathrm{map}_c(U, BG)$ as giving a G-bundle on X which is trivialized outside of U (or more precisely, such a datum is given by the image of that map in $\mathrm{map}_c(X, BG)$). Moreover, nonabelian Poincaré duality is telling us that ranging over all open subsets $U \subset X$ which are homeomorphic to a finite disjoint union of disks allows us to describe *all* G-bundles over X in this way, in an essentially unique fashion (up to inclusions $U \subset U'$).

Remark 36. In the further special case that X is a Riemann surface, i.e. a 2-dimensional manifold, our hypotheses dictate that we need BG to be 1-connected; this is why we assume that G itself is connected.

4.3 The Ran space and nonabelian Poincaré duality in algebraic geometry

Let's now return to our actual situation of interest. Recall that we have $G \to X \to \mathrm{Spec}(k)$, and we are interested in the cohomology of the algebraic stack $\mathrm{Bun}_G(X)$ of G-bundles on X.

To mimic the topological side, we define the *Ran space* of X, denoted $\mathrm{Ran}_G(X)$, to be the stack of G-bundles on X which are trivialized away from finitely many points (which is the appropriate analog of our previous notions when we're working in the Zariski topology). Let's be a little more precise: maps $Y \to Ran_G(X)$ are the same thing as G-bundles $P \to X \times Y$ together with a nonempty finite set of maps $f_1, \ldots, f_n : Y \to X$ together with a trivialization of P on $(X \times Y) \backslash (\Gamma_{f_1} \cup \cdots \Gamma_{f_n})$. (In the special case that Y is a point, we get the previous informal description.)

Now, there's a forgetful map $\mathrm{Ran}_G(X) \to \mathrm{Bun}_G(X)$ (which just remembers the G-bundle but forgets the maps $f_i : Y \to X$ and the trivialization), and this leads us to the following result.

Theorem 37 (Algebro-geometric nonabelian Poincaré duality)**.** *If X is an algebraic curve over k, $\ell^{-1} \in k$, and $G \to X$ is a semisimple, generically simply-connected, smooth group scheme which has connected fibers, then the map*

$$\mathrm{Ran}_G(X) \to \mathrm{Bun}_G(X)$$

induces an isomorphism in ℓ-adic cohomology.

Remark 38. In fact, this statement holds under weaker conditions, too – for example, when $G = GL_n$.

Let's first sketch a proof of this for $G = GL_1$. In this case, $\mathrm{Bun}_G(X)$ classifies line bundles on X, while $\mathrm{Ran}_G(X)$ classifies line bundles that are generically trivialized. we're claiming that this induces some sort of equivalence, and this map sort of behaves like a *fibration*, so that we can simply check that the fibers themselves are contractible in the appropriate sense. Now, we have a map pt $\to Bun_G(X)$ selecting the trivial G-bundle, and then the pullback of

$$\begin{tikzcd} \mathrm{Ran}_G(X) \arrow[r] & \mathrm{Bun}_G(X) \\ & \mathrm{pt} \arrow[u] \end{tikzcd}$$

yields the "generic trivializations" of the trivial G-bundle. These correspond to nonzero meromorphic functions, i.e. rational maps $X \dashrightarrow \mathbb{A}^1 \backslash \{0\}$. This set of rational maps is precisely $K_X \backslash \{0\}$, which we should think of as contractible: we think of K_X as an infinite-dimensional vector space (or perhaps slightly better, as a colimit of finite-dimensional vector spaces), and we recall that S^∞ is contractible.

Now, the generalization to $G = GL_n$ goes as follows. now we get $X \dashrightarrow GL_n \subset \mathbb{A}^{n^2}$, and now instead of removing $\{0\} \subset K_X$ we remove some subset $U \subset (K_X)^{n^2}$, whose complement is described informally by

$$\mathrm{map}_{\mathrm{rat'l}}(X, \{A : \det(A) = 0\}).$$

Once again, we're removing something from an infinite-dimensional vector space which has infinite codimension, and so the result should again be thought of as contractible.

Remark 39. One might ask whether this map is an equivalence in some motivic sense. Lurie suspects that this is not true; at the very least, the proof of the statement certainly does not imply that this is so.

5 The computation

5.1 Pushforward along $\mathbf{Ran}_G(X) \to \mathbf{Ran}_{\{e\}}(X)$ and the affine Grassmannian

Let us recall our setup: we're working over a field K, we have a smooth algebraic curve X defined over K, and we have a family of groups $G \to X$ parametrized by X; our goal is to understand $H^*(\mathrm{Bun}_G(X), \mathbb{Q}_\ell)$. In the previous lecture we discussed a tool for studying this, namely the map

$$\mathrm{Ran}_G(X) \to \mathrm{Bun}_G(X);$$

in passing to $\mathrm{Ran}_G(X)$, we're adding the data of a finite set of points of X together with a trivialization of our G-bundle away from those points. Recall that nonabelian Poincaré duality guarantees that under mild hypotheses on G, this map induces an isomorphism on ℓ-adic cohomology. In addition to the map above, we also have a map

$$\mathrm{Ran}_G(X) \xrightarrow{\varphi} \mathrm{Ran}(X) = \mathrm{Ran}_{\{e\}}(X);$$

now, the target simply parametrizes nonempty finite subsets of X. We will see that we can understand the fibers of this map purely in terms of G.

For simplicity, let's assume that G is constant. (One might keep in mind the example $G = SL_n$; then G-bundles are just vector bundles with trivialized determinant.) Then, using a sort of Leray–Serre spectral sequence, we have isomorphisms

$$H^*(\mathrm{Bun}_G(X), \mathbb{Q}_\ell) \cong H^*(\mathrm{Ran}_G(X), \mathbb{Q}_\ell) \cong H^*(\mathrm{Ran}(X), \varphi_* \underline{\mathbb{Q}_\ell})$$

(where of course φ_* really denotes a *derived* pushforward along φ). Let's write $A = \varphi_* \underline{\mathbb{Q}_\ell}$ for this complex of sheaves.

Now, let's consider the fibers of φ. A point of $\mathrm{Ran}(X)$ is a finite subset of X, so the simplest case is a singleton subset $\{x\} \in X \subset \mathrm{Ran}(X)$ (where we consider $X \subset \mathrm{Ran}(X)$ via singletons). In this case, $\varphi^{-1}(\{x\})$ looks like the space of G-bundles on X equipped with a trivialization on $X \backslash \{x\}$. Now, we can fix a trivialization of our G-bundle on a formal neighborhood of the point x, and then all the data is given by the clutching function along the punctured disc: in other words, such G-bundles are parametrized by $G(K_x)$. (Recall that after choosing a coordinate, we have $K_x \cong \kappa(x)((t))$.) But then, given that we actually want to parametrize G-bundles instead of G-bundles equipped with a trivialization, we see that we want to take a quotient to remove the ambiguity introduced when we chose the trivialization at x in the first place: thus, we are actually interested in $G(K_x)/G(\mathcal{O}_x)$. (Recall that after choosing a coordinate, the inclusion $\mathcal{O}_x \subset K_x$ corresponds to the inclusion $\kappa(x)[[t]] \subset \kappa(x)((t))$.)

Now, $G(K_x)/G(\mathcal{O}_x)$ is an algebro-geometric object, a sort of infinite-dimensional algebraic variety (compare with e.g. $\mathbb{P}^\infty = \operatorname{colim} \mathbb{P}^n$). It is called the *affine Grassmannian*, and is denoted $\mathrm{Gr}_G = G(K_x)/G(\mathcal{O}_x)$.

Remark 40. Since "up to homotopy" we're supposed to think of $G(K_x)$ as "maps into G from a circle" and of $G(\mathcal{O}_x)$ as "maps into G from a disc", we should therefore think of Gr_G as being an incarnation of $\Omega G \simeq \Omega^2 BG$. (Perhaps better, we might think of the affine Grassmannian as an *algebraic incarnation* of this two-fold loopspace.)

Now, Gr_G is a colimit of projective varieties, and hence behaves as a sort-of "compact" object. Hence, our map $\mathrm{Ran}_G(X) \xrightarrow{\varphi} \mathrm{Ran}(X)$ behaves like a *proper* map (i.e. it has "compact" fibers). Hence, the *proper base change theorem* tells us that we have a quasi-isomorphism of the stalk $A_x \simeq C^*(\mathrm{Gr}_G, \mathbb{Q}_\ell)$ of A at x with the ℓ-adic cochains on the affine Grassmannian.

More generally, let $S = \{x_1, \ldots, x_n\} \in \mathrm{Ran}(X)$. We can carry out the same analysis, but now we have n disks instead of a single disk, and so we get $\varphi^{-1}(S) = \mathrm{Gr}_G \times \cdots \times \mathrm{Gr}_G$, and thence we obtain an identification of the stalk of A at $S \in \mathrm{Ran}(X)$ as

$$A_S \simeq C^*\left(\prod_n \mathrm{Gr}_G, \mathbb{Q}_\ell\right) \simeq \bigotimes_n C^*(\mathrm{Gr}_G, \mathbb{Q}_\ell) \simeq A_x.$$

5.2 Factorizable sheaves

From this computation, we see that A provides us with the data of a *factorizable sheaf* on X. Namely, if $S = S_1 \coprod S_2$, then we have the "factorization" condition $A_S \simeq A_{S_1} \otimes A_{S_2}$. So, suppose we have two points $x, y \in X$. in the Ran space, what happens if we "move these points together"? More precisely, if we have "paths" from x to z and from y to z, these give us a path $[0,1] \to \mathrm{Ran}(X)$ which changes strata at the endpoint $1 \in [0,1]$. If we pull the sheaf A back to the interval $[0,1]$, we will therefore obtain a sheaf which is the constant sheaf $\underline{C^*(\mathrm{Gr}) \otimes C^*(\mathrm{Gr})}$ over $[0,1)$ but has value $C^*(\mathrm{Gr})$ at 1. To specify such a sheaf, it therefore suffices to define a *cospecialization* map: any stalk at 1 is supposed to extend to a section on some small neighborhood of 1, and so this is asking for a map $C^*(\mathrm{Gr}) \to C^*(\mathrm{Gr}) \otimes C^*(\mathrm{Gr})$.

But this map is actually coming from topology! Recalling the identification $\mathrm{Gr} \approx \Omega^2 BG$, there is a loop multiplication $\Omega^2 BG \times \Omega^2 BG \to \Omega^2 BG$, which is exactly what's inducing the map in the other direction on cochains.

This story is part of a dictionary, at least if we're working in topology instead of algebraic geometry. First of all, we have a containment

$$\{\text{non-unital } E_n\text{-coalgebras}\} \subset \{\text{factorizable sheaves on } \mathrm{Ran}(\mathbb{R}^n)\}.$$

(The non-unitality reflects the fact that we're not considering the empty set as part of the Ran space.) More precisely, $Gr \approx \Omega^2 BG$ is a 2-fold loopspace, and hence $C^*(Gr, \mathbb{Q}_\ell)$ is an E_2-coalgebra.

Now, given any ("nicely symmetric") E_2-coalgebra, for any surface Σ we can construct a factorizable sheaf on $\mathrm{Ran}(\Sigma)$; this construction takes $C^*(\mathrm{Gr})$ to A. We can also consider *reduced* cochains, which will kill off our co-unit, and so we get that $C^*_{red}(\mathrm{Gr})$ gives A_{red}. Now, on both sides these objects are pretty similar, and finally we end up with

$$H^*(\mathrm{Ran}(X), A) \cong H^*(\mathrm{Ran}(X), A_{red}) \oplus \mathbb{Q}_\ell.$$

5.3 Verdier duality and Koszul duality

Now, let's return from working topology to the setting of algebraic geometry over a finite field. Recall that we wanted to compute the number

$$\mathrm{Tr}(\mathrm{Frob}^{-1} | H^*(\mathrm{Ran}(X), A)).$$

This is almost exactly the setting of the Grothendieck–Lefschetz trace formula. Let us take a moment to describe this. Let Y be an algebraic variety over $\mathbb{F}_q$, and let $\mathcal{F}$ be an ℓ-adic sheaf on Y. Then $H^*_c(Y, \mathcal{F})$ carries a Frobenius action, and the trace formula says that

$$\mathrm{Tr}(\mathrm{Frob} | H^*_c(Y, \mathcal{F})) = \sum_{y \in Y(\mathbb{F}_q)} \mathrm{Tr}(\mathrm{Frob} | \mathcal{F}_y).$$

In the case that $\mathcal{F}$ is constant, the summands on the right all equal 1, so this is just counting the set $Y(\mathbb{F}_q)$. So, this looks kind of like what we're trying to do, except for two differences: we're asking about the trace of the *inverse* of Frobenius, and we're also looking at cohomology as opposed to compactly-supported cohomology. So, we'll use *Verdier duality* in our situation in order to apply the trace formula.

First of all, we'd like to say

$$\mathrm{Tr}(\mathrm{Frob}^{-1}|H^*(\mathrm{Ran}(X), A)) = \mathrm{Tr}(\mathrm{Frob}|H^*(\mathrm{Ran}(X), A)^\vee) = \mathrm{Tr}(\mathrm{Frob}|H_c^*(\mathrm{Ran}(X), \mathbb{D}A)),$$

where $\mathbb{D}$ denotes "Verdier duality". Now, this would roughly be true, except that the Ran space is a big, infinite-dimensional object, and it turns out (for this reason) that in fact, the dual $\mathbb{D}A$ is zero, and this equality of traces is just totally false. On the other hand, we can fix this by taking a *weighted* count: in essence, throwing away the basepoint gives us that

$$\mathrm{Tr}(\mathrm{Frob}^{-1}|H^*(\mathrm{Ran}(X), A_{red})) = \mathrm{Tr}(\mathrm{Frob}|H^*(\mathrm{Ran}(X), A_{red})^\vee) = \mathrm{Tr}(\mathrm{Frob}|H_c^*(\mathrm{Ran}(X), \mathbb{D}A_{red}))$$

is true.

Now, Verdier duality takes factorizable sheaves to factorizable sheaves. Moreover, under our dictionary which takes E_2-coalgebras to factorizable sheaves, the Verdier dual of a factorizable sheaf corresponds to the *Koszul dual* of the corresponding E_2-coalgebra. we won't really go into detail here, but we're after the Koszul dual of $C^*_{red}(\Omega^2 BG)$, and this is precisely $C_*^{red}(BG)$. so, we're saying that at a point $x \in X$, the stalk is given by $(\mathbb{D}A_{red})_x \simeq C_*^{red}(BG, \mathbb{Q}_\ell)$, the chain complex computing the reduced homology of BG with $\mathbb{Q}_\ell$-coefficients.

5.4 The computation

Now, let's take a leap of faith and accept that this all works out as it should. Then, we have that

$$1 + \mathrm{Tr}(\mathrm{Frob}|H_c^*(\mathrm{Ran}(X), \mathbb{D}A_{red})) = 1 + \sum_{S \in \mathrm{Ran}(X)(\mathbb{F}_q)} \mathrm{Tr}(\mathrm{Frob}| \bigotimes_{s \in S} C_*^{red}(BG)).$$

What this is saying is that S is a union of orbits for $\mathrm{Gal}(\overline{\mathbb{F}}_q/\mathbb{F}_q)$ on $X(\overline{\mathbb{F}}_q)$, which are exactly the closed points of X, considered as a scheme. So we can rewrite this as

$$1 + \sum_{\emptyset \neq T \subset \text{closed points of } X} \left(\prod_{x \in T} \mathrm{Tr}(\mathrm{Frob}_x | C_*^{red}(BG)) \right).$$

So, this is a sum over all nonempty subsets – plus 1, which corresponds to the empty set – so we can rewrite it as

$$\sum_{T \subset \text{closed points of } X} \left(\prod_{x \in T} \mathrm{Tr}(\mathrm{Frob}_x | C_*^{red}(BG)) \right).$$

Modulo convergence issues, we can rewrite this using the distributive law as

$$\prod_{x \in X} 1 + \mathrm{Tr}(\mathrm{Frob}_x | C_*^{red}(BG)).$$

Once again, the difference between $C_*^{red}(BG)$ and $C_*(BG)$ is just a copy of $\mathbb{Q}_\ell$ on which the Frobenius acts as the identity, so this is

$$\prod_{x \in X} \mathrm{Tr}(\mathrm{Frob}_x | C_*(BG)).$$

Using duality we can rewrite this as

$$\prod_{x \in X} \mathrm{Tr}(\mathrm{Frob}^{-1} | C^*(BG, \mathbb{Q}_\ell)).$$

Finally, applying the trace formula to BG, we see that this equals

$$\prod_{x \in X} \frac{|BG(\kappa(x))|}{|\kappa(x)|^{\dim BG}} = \prod_{x \in X} \frac{|\kappa(x)|^{\dim G}}{|G(\kappa(x))|}.$$

And this is exactly the mass formula that shows up in Weil's conjecture.