The Grothendieck-Lefschetz trace formula

Tony Feng

April 19, 2017

Contents

1 Introduction to the trace formula

1.1 The Lefschetz fixed point formula

Weil's motivation for the trace formula was by analogy to the classical theorem of Lefschetz:

Theorem 1.1 (Lefschetz fixed point theorem)**.** *Let M be a compact, oriented manifold. If $\psi \colon M \to M$ is a continuous map with isolated fixed points, then*

$$\# \operatorname{Fix}(\psi) = \sum_i (-1)^i \operatorname{Tr}(\psi^*, H^i(M, \mathbf{R})).$$

We caution that $\# \operatorname{Fix}(\psi)$ has to be interpreted with some care: it really denotes a fixed point count *with multiplicities.* A better way to phrase it as follows. Inside $M \times M$ we have two submanifolds: the diagonal Δ and the graph Γ_ψ. We can then try to take

$$\# \operatorname{Fix}(\psi) := \Delta \cdot \Gamma_\psi = \text{“intersection number of } \Delta \text{ and } \Gamma_\psi\text{”}.$$

When Δ and Γ_ψ intersect transversely this intersection number coincides with the naïve point count, but it can be defined more generally. The “correct” version of the theorem is then

$$\Delta \cdot \Gamma_\psi = \sum_i (-1)^i \operatorname{Tr}(\psi, H^i(M, \mathbf{R})).$$

Weil envisioned that an algebraic analogue of this story would yield an interesting trace formula for algebraic varieties over a finite field. If X_0 is a variety over $\mathbf{F}_q$, and $X = X_0 \times_{\mathbf{F}_q} \overline{\mathbf{F}}_q$, then we could take $\psi \colon X \to X$ to be the Frobenius morphism $\psi = (\mathrm{Frob}_{X_0} \times \mathrm{Id})$. Then $\# \operatorname{Fix}(\psi)$ should count the $\overline{\mathbf{F}}_q$-points of X whose coordinates are *fixed by Frobenius*, which is the same as saying that they lie in $\mathbf{F}_q$.

To summarize, a version of the Lefschetz trace formula in étale cohomology would say: for a smooth proper variety $X/\mathbf{F}_q$,

$$\# X(\mathbf{F}_q) = \sum_i (-1)^i \operatorname{Tr}(\psi, H^i(X, \mathbf{Q}_\ell)).$$

The purpose of these notes is to prove a slight generalization of this statement. (We can remove the assumptions on smoothness or properness, trading cohomology for compactly-supported cohomology, and allow general étale sheaves.)

1.2 The Frobenius endomorphism

We recall the Frobenius action on sheaves, which factor into the statement of the Lefschetz trace formula.

For any scheme $X_0/\mathbf{F}_q$, let F_{X_0} be its q-Frobenius, inducing $f \mapsto f^p$ on rings. For any étale sheaf $\mathcal{F}_0$ on X_0, we have an isomorphism

$$\mathrm{Fr}_{\mathcal{F}_0} \colon F_{X_0}^* \mathcal{F}_0 \to \mathcal{F}_0.$$

Letting $X = X_0 \times_{\mathbf{F}_q} \overline{\mathbf{F}}_q$ and $\mathcal{F}$ be the pullback of $\mathcal{F}_0$ to X, we have an endomorphism on cohomology:

$$F \colon H^i(X, \mathcal{F}) \xrightarrow{(F_{X_0} \times \mathrm{Id})^*} H^i(X, (F_{X_0} \times \mathrm{Id})^* \mathcal{F}) \xrightarrow{\mathrm{Fr}_{\mathcal{F}_0}} H^i(X, \mathcal{F}).$$

We remind you that F coincides with the action of *geometric Frobenius* on $H^i(X, \mathcal{F})$ induced by $(\mathrm{Id}_{X_0} \times \mathrm{Frob}_{\overline{\mathbf{F}}_q}^{-1})$.

More generally, let $\mathcal{G}$ be an étale sheaf on $\operatorname{Spec} k = \mathbf{F}_{q^d}$. Then we have a morphism

$$\mathrm{Fr}_{\mathcal{G}}^d \colon F_{\operatorname{Spec} k}^{d*} \mathcal{G} = \mathcal{G} \to \mathcal{G}.$$

Upon choosing a separable closure of k to identify $\mathcal{G}$ with a $\operatorname{Gal}(\overline{k}/k)$-modules M (under which $M = \mathcal{G}_{k^s}$), the map endomorphism $\mathrm{Fr}_{\mathcal{G}}^d$ corresponds to the "arithmetic Frobenius" in $\operatorname{Gal}(\overline{k}/k)$ acting on M.

- If we apply this to $\mathcal{G} = H^i(X_0, \mathcal{F}_0)$ then $d = 1$ and

  $$\mathrm{Fr}_{\mathcal{G}} \colon H^i(X, \mathcal{F}) \to H^i(X, \mathcal{F})$$

  coincides with the geometric Frobenius action on the $\operatorname{Gal}(\overline{k}/k)$-module corresponding to $\mathcal{G}$.

- We will also apply this discussion to $k = \kappa(x)$ for a closed point $x \in |X|$, and denote by
$$F_x \colon \mathcal{F}_{\overline{x}} \to \mathcal{F}_{\overline{x}}$$
the endomorphism corresponding to the *geometric Frobenius* in $\mathrm{Gal}(\overline{k}/k)$.

1.3 Statement of the formula

Theorem 1.2. *Let X_0 be a variety over $\mathbf{F}_q$ and $\mathcal{F}_0$ a constructible $\mathbf{Q}_\ell$-sheaf on X_0. Let $X = X_0 \times_{\mathbf{F}_q} \overline{\mathbf{F}}_q$ and $\mathcal{F}$ be the pullback of $\mathcal{F}_0$ to X. Then we have*

$$\sum_{x \in X(\mathbf{F}_q)} \mathrm{Tr}(F_x, \mathcal{F}_x) = \sum_i (-1)^i \,\mathrm{Tr}(F, H^i_c(X, \mathcal{F})).$$

Example 1.3. Let $X_0 = \mathbf{P}^1$. Then we know that

$$H^i_c(X, \mathbf{Q}_\ell) = H^i(X, \mathbf{Q}_\ell) = \begin{cases} \mathbf{Q}_\ell & i = 0 \\ 0 & i = 1 \\ \mathbf{Q}_\ell(-1) & i = 2 \end{cases}.$$

(The only interesting aspect of this is the identification of the Galois structure on H^2, which was explained in my previous talk on Poincaré duality.) Therefore, we have

$$\mathrm{Tr}(F, H^i_c(X, \mathcal{F})) = \begin{cases} 1 & i = 0 \\ 0 & i = 1 \\ q & i = 2 \end{cases}$$

so the trace formula tells us that

$$\#\mathbf{P}^1(\mathbf{F}_q) = 1 + q$$

as we know!

Example 1.4. Let $X_0 = \mathbf{A}^1$. Then

$$H^i_c(X, \mathbf{Q}_\ell) = H^i(X, \mathbf{Q}_\ell) = \begin{cases} 0 & i = 0 \\ 0 & i = 1 \\ \mathbf{Q}_\ell(-1) & i = 2 \end{cases}$$

The trace formula tells us that

$$\#\mathbf{A}^1(\mathbf{F}_q) = q$$

as expected.

2 Weil's trace formula

The proof of Theorem 1.2 will proceed by reduction to the case of curves, where eventually something geometric happens. We will go through this geometric step first. The goal is to establish a special case due to Weil:

Theorem 2.1 (Weil)**.** *Let C be a smooth projective curve over $k = \overline{\mathbf{F}}_q$ and $\psi \colon C \to C$ an endomorphism over k. Then*

$$\Delta \cdot \Gamma_\psi = \sum_{i=0}^{2} (-1)^i \operatorname{Tr}(\psi^*, H^i(C, \mathbf{Q}_\ell)).$$

Once we have a fairly robust apparatus of étale cohomology at our disposal, the proof becomes quite formal (analogous to the proof of Leschetz's trace formula in topology).

Remark 2.2. Weil didn't have access to étale cohomology, so he proved Theorem 2.1 (or really, a cohomology-free reformulation of it) by using the theory of the Jacobian. The point is that he essentially knew what étale cohomology of smooth projective curves should be, in terms of Jacobians.

2.1 Cohomology classes of algebraic cycles

Let X be a smooth variety of dimension d over an algebraically closed field $k = \overline{k}$. Our first mini-goal is to explain how to associated to a closed subvariety $i \colon Y \hookrightarrow X$ a cycle class $\mathrm{cl}_X(Y) \in H^*(X; \mathbf{Q}_\ell)$.

We may assume that Y is irreducible (the cycle class of a union will be the sum of the cycle classes of the components), of dimension e. When Y is also smooth, this can be defined as follows. We have a pullback map

$$i^* \colon H^*_c(X; \mathbf{Q}_\ell) \to H^*_c(Y; \mathbf{Q}_\ell).$$

We can dualize this to obtain

$$(i^*)^\vee \colon H^*_c(Y; \mathbf{Q}_\ell)^\vee \to H^*_c(X; \mathbf{Q}_\ell)^\vee.$$

Then we can identify these duals with compactly-supported cohomology using Poincaré duality (which is where we use the smoothness assumption):

$$\begin{array}{ccc} H^i_c(Y; \mathbf{Q}_\ell(e))^\vee & \xrightarrow{(i^*)^\vee} & H^i_c(X; \mathbf{Q}_\ell(e))^\vee \\ \| & & \| \\ H^{2e-i}(Y; \mathbf{Q}_\ell) & \longrightarrow & H^{2d-i}(X; \mathbf{Q}_\ell(d-e)). \end{array} \tag{2.1}$$

Definition 2.3. If Y is smooth and irreducible, the *cycle class*

$$\mathrm{cl}_X(Y) \in H^{2d-2e}(X; \mathbf{Q}_\ell(d-e))$$

is the image of $1 \in H^0(Y; \mathbf{Q}_\ell)$ under $(i^*)^\vee$, using the identification (2.1).

If Y is not smooth, then unfortunately we don't have Poincaré duality available. However, there is a dense open subset $j \colon U \subset Y$ which is smooth which induces an isomorphism

$$H_c^{2e}(U) \cong H_c^{2e}(Y)$$

(as one can see by looking at the the long exact sequence of the pair). Therefore, the fundamental class in $H_c^{2e}(U; \mathbf{Q}_\ell(e))$ transfers to a fundamental class in $H_c^{2e}(Y, \mathbf{Q}_\ell(e))$, which we can then stuff into $H^{2d-2e}(X; \mathbf{Q}_\ell(d-e))$.

Definition 2.4. In general, we define the *cycle class*

$$\mathrm{cl}_X(Y) \in H^{2d-2e}(X; \mathbf{Q}_\ell(d-e))$$

to be the image of the fundamental class in $H_c^{2e}(Y; \mathbf{Q}_\ell)$ under $(i^*)^\vee$, using the identification

$$\begin{array}{ccc} H_c^i(Y; \mathbf{Q}_\ell(e))^\vee & \xrightarrow{(i^*)^\vee} & H_c^i(X; \mathbf{Q}_\ell(e))^\vee \\ & & \| \\ & & H^{2d-i}(X; \mathbf{Q}_\ell(d-e)) \end{array}$$

That was a little abstract, so here is a concrete interpretation. Thanks to Poincaré duality, the class $\mathrm{cl}_X(Y)$ is uniquely determined by the data of the pairings

$$\langle \mathrm{cl}_X(Y), \beta \rangle_X \text{ for all } \beta \in H_c^*(X; \mathbf{Q}_\ell)$$

and we have

$$\langle \mathrm{cl}_X(Y), \beta \rangle_X = \int_Y i^* \beta.$$

I like to write this in the following form, which makes sense when Y is smooth:

$$\langle i_* \alpha, \beta \rangle_X = \langle \alpha, i^* \beta \rangle_Y \text{ for all } \alpha \in H^*(Y; \mathbf{Q}_\ell), \beta \in H_c^*(X; \mathbf{Q}_\ell).$$

Remark 2.5. There is another interpretation of the cycle class for a *divisor* in terms of Chern classes. Let $D \subset X$ be a divisor. Then we get a line bundle $\mathcal{O}(D)$, to which we can assign a *first Chern class* $c_1(\mathcal{O}(D)) \in H^2(X; \mathbf{Q}_\ell(1))$ assembled from the boundary homomorphisms

$$H^1(X; \mathbf{G}_m) \to H^2(X; \mu_{\ell^n}).$$

2.2 Proof of Weil's Theorem

By the preceding section, we can associated cohomology classes to the embeddings $\Delta \hookrightarrow C \times C$ and $\Gamma_\psi \hookrightarrow C \times C$, obtaining $\mathrm{cl}_{C\times C}(\Delta)$ and $\mathrm{cl}_{C\times C}(\Gamma_\psi)$. (Note that in this case Δ and Γ are both smooth, both being isomorphic to C.)

The first order of business is to "identify" the class $\mathrm{cl}_{C\times C}(\Delta)$. Since Δ is so canonical, it must have a canonical description. To express this more conveniently, we *fix* a trivialization $\mu_{\ell^\infty} \cong \mathbf{Z}_\ell$, thus identifying $\mathbf{Q}_\ell(j) \cong \mathbf{Q}_\ell$ for all j. The point is that this makes $H^*(C \times C; \mathbf{Q}_\ell)$ self-dual.

Lemma 2.6. *Let $\{e_i\}$ be a basis for $H^*(C; \mathbf{Q}_\ell)$ and $\{f_i\}$ its dual basis, so $\langle e_i, f_j \rangle_X = \delta_{ij}$. Under the Künneth isomorphism*

$$H^*(C \times C; \mathbf{Q}_\ell) \cong H^*(C; \mathbf{Q}_\ell) \otimes H^*(C; \mathbf{Q}_\ell)$$

we have

$$\Delta \leftrightarrow \sum e_i \otimes f_i$$

The proof is a straightforward exercise in unwinding the definition of the push-forward

$$H^*(\Delta; \mathbf{Q}_\ell) \to H^{*+2}(C \times C; \mathbf{Q}_\ell).$$

Instead of going through these motions, I would rather tell you about a picture I have, which "explains" the formula. A cohomology class $\alpha \in H^*(C \times C; \mathbf{Q}_\ell)$ defines a correspondence on $H^*(C)$, as follows:

$$\alpha \mapsto [u \in H^*(C; \mathbf{Q}_\ell) \mapsto p_{2*}(\alpha \smile p_1^*(u))] \,.$$

This defines a map

$$H^*(C \times C; \mathbf{Q}_\ell) \to \mathrm{End}(C).$$

In fact, if the cohomology class α has a more refined interpretation, e.g. as a cycle class, then we can realize this action at the correspondingly more refined level. Obviously the class of Δ should act as the identity, so this map should take

$$\mathrm{cl}_{C\times C}(\Delta) \mapsto \mathrm{Id}\,.$$

We can unwind this map using Künneth and Poincaré duality:

$$\begin{array}{ccc} H^*(C; \mathbf{Q}_\ell) \otimes H^*(C; \mathbf{Q}_\ell) & \overset{\text{PD}}{=\!=\!=} & H^*(C; \mathbf{Q}_\ell)^\vee \otimes H^*(C; \mathbf{Q}_\ell) \\ \text{Künneth}\Big\| & & \Big\| \\ H^*(C \times C; \mathbf{Q}_\ell) & \longrightarrow & \mathrm{End}(H^*(C; \mathbf{Q}_\ell)). \end{array}$$

Now the formula of Lemma 2.6 is the familiar expression of $\mathrm{Id} \in \mathrm{End}(C)$ as the sum of a basis tensored with its dual basis.

Exercise 2.7. Check the claims made in the preceding paragraphs.

Now suppose

$$\mathrm{Cl}_{C\times C}(\Gamma_\psi) = \sum_{ij} c_{ji}\cdot f_j \otimes e_i,$$

By the picture we just explained, the action on cohomology is given by

$$\psi^*(e_k) = p_{2*}((e_k\otimes 1)\smile \sum_{i,j} c_{ji}\cdot f_j\otimes e_i) = \sum_i c_{ki}e_i.$$

Therefore, the graded trace is

$$\mathrm{Tr}(\psi^*, \sum (-1)^\nu H^\nu(C;\mathbf{Q}_\ell)) = \sum c_{ii}(-1)^{\deg e_i}.$$

while

$$\begin{aligned}\langle \Delta, \Gamma_\psi\rangle &= \sum_{i,j} c_{ij}\sum_k \langle e_k\otimes f_k, f_j\otimes e_i\rangle \\ &= \sum_{i,j} c_{ij}\sum_{j,k}\delta_{kj}\delta_{ki}(-1)^{\deg e_i} \\ &= \sum_i c_{ii}(-1)^{\deg e_i}\end{aligned}$$

2.3 Intersection number vs cup product

One issue we glossed over in the preceding section is in what sense the cohomological pairing $\langle\Delta,\Gamma_\psi\rangle$ is an accurate reflection of "$\Delta\cdot\Gamma_\psi$". Another interpretation of the latter quantity is in terms of the intersection product on Chow groups. This is itself subtle to define in general, but in some cases it is clear what the answer should be. For instance, we might ask the following questions:

- When Δ and Γ_ψ intersect transversely, does the cohomological pairing coincide with the naïve $\#\Delta\cap\Gamma_\psi$? (We know that the intersection pairing does have this property.)
- If $\psi = \mathrm{Frob}_{X_0}\times\mathrm{Id}$ for some $X_0/\mathbf{F}_q$ descending X, do any of these coincide with $\#X_0(\mathbf{F}_q)$?

Let's address the second question first. Frobenius enjoys the happy property that its graph *always* intersects Δ transversely. This is because $d(f^p)=0$ in characteristic p, so $d\,\mathrm{Frob}=0$. So $T_{(x,x)}\Delta$ is a diagonally embedded copy of T_xC, while $T_{(x,y)}\Gamma_\psi$ is the copy of T_xC embedded into the first factor of $T_xC\times T_yC$. These tangent spaces are always transverse.

We know that the (Chow-theoretic) intersection product always coincides with the naïve point count in the happy case of a transverse intersection, so it suffices to see that the cohomology pairing of curves on a surface coincides with the intersection product in general. It will be useful to recall the definition of the intersection product on a surface.

Let S be a smooth projective surface over an algebraically closed field k. We can define an intersection product on $\mathrm{CH}^1(S) = \mathrm{Pic}(S)$ as follows. By linearity, it suffices to define the intersection $[C] \cdot [D]$ where C, D are (effective) curves. When C and D intersect transversally, we have the short exact sequence

$$0 \to \mathcal{O}_C(-D) \to \mathcal{O}_C \to \mathcal{O}_{C\cap D} \to 0$$

and $\#(C \cap D) = \deg \mathcal{O}_{C\cap D} = -\deg \mathcal{O}_C(-D)$. Motivated, by this, we define in general:

$$[C] \cdot [D] := \deg \mathcal{O}(D)|_C.$$

Let's compare this to the cohomological definition. The inclusion $i\colon C \hookrightarrow S$ induces

$$i_*\colon H^*(C; \mathbf{Q}_\ell) \xrightarrow{\sim} H^{*+2}(S, \mathbf{Q}_\ell(1))$$

with the property that

$$\langle i_*\alpha, \beta\rangle_S = \langle \alpha, i^*\beta\rangle_C.$$

In particular, for $\beta = \mathrm{Cl}_S(D)$ and $\alpha = 1$ (so $i_*\alpha = \mathrm{Cl}_S(C)$), we have

$$\langle \mathrm{Cl}_S(C), \mathrm{Cl}_S(D)\rangle = \int i^*\beta.$$

By Remark 2.5, β is the first Chern class of the divisor $\mathcal{O}(D)$ so $i^*\beta$ is the first Chern class of its restriction to C, by functoriality of Chern classes (which is very easy in this case).

3 Perfect complexes

3.1 Reformulation of the trace formula

Consider the statement of Theorem 1.2. We know that the constructible $\mathbf{Q}_\ell$-sheaf $\mathcal{F}$, and its cohomology, are really defined from pro-systems of constructible torsion sheaves $(\mathcal{F}_n)$. We want to deduce it from the version for $\mathcal{F}_n$ for each n:

$$\sum_{x\in X(\mathbf{F}_q)} \mathrm{Tr}(F_x, (\mathcal{F}_n)_x) = \sum_i (-1)^i \,\mathrm{Tr}(F, H^i_c(X, \mathcal{F}_n)).$$

Now what does this mean? Each $\mathcal{F}_n$ is a constructible $\mathbf{Z}/\ell^{n+1}\mathbf{Z}$-module, so at best this is an equality of numbers in $\mathbf{Z}/\ell^{n+1}\mathbf{Z}$.

Once we start working with modules over $\mathbf{Z}/\ell^{n+1}\mathbf{Z}$ another question arises: what is the trace? We should probably take $\mathcal{F}_n$ to have stalks that are flat over $\mathbf{Z}/\ell^{n+1}\mathbf{Z}$ in order to talk about the trace on the left side. However, we then run into the problem that even if $\mathcal{F}_n$ has stalks flat over $\mathbf{Z}/\ell^{n+1}\mathbf{Z}$, its *cohomology* groups $H^i_c(X, \mathcal{F}_n)$ don't necessarily retain this property.

The way out of this conundrum is based on the following idea. If we have projective module over a ring Λ, then we have a good notion of trace of an endomorphism. If we have a module that is not projective, then we try to replace it by a *complex* of projectives, and extend the endomorphism to an endomorphism of complexes.

3.2 Perfect complexes

Let Λ be a ring, which we'll often take to be $\mathbf{Z}/\ell^n\mathbf{Z}$. First we review how to the trace of an endomorphism of a finite projective module M over Λ. Perhaps we should first review how to take the trace if M is *free*. In this case we can choose an isomorphism $M \cong \Lambda^N$, so that any endomorphism $u \colon M \to M$ can be represented by a matrix (λ_{ij}). We then define the trace of u to be the sum of the diagonal entries λ_{ii}. That this is well-defined amounts to

$$\mathrm{Tr}(ABA^{-1}) = \mathrm{Tr}(B)$$

for any A, which is a consequence of the cyclicity property $\mathrm{Tr}(AB) = \mathrm{Tr}(BA)$.

If M is projective, then it is a summand of a free module Λ^N so we can choose a decomposition

$$\Lambda^N \cong M \oplus M'.$$

We can then extend an endomorphism u of M to Λ^N by 0 on M', and define the trace of u to be the trace of this extension.

If $P_\cdot$ is a finite complex of finite projective modules over Λ, and $u_\cdot \colon P_\cdot \to P_\cdot$ is an endomorphism, meaning a commutative diagram

$$\begin{array}{ccccccccc}
\dots & \longrightarrow & P_{-1} & \longrightarrow & P_0 & \longrightarrow & P_1 & \longrightarrow & \dots \\
 & & \downarrow{\scriptstyle u_{-1}} & & \downarrow{\scriptstyle u_0} & & \downarrow{\scriptstyle u_1} & & \\
\dots & \longrightarrow & P_{-1} & \longrightarrow & P_0 & \longrightarrow & P_1 & \longrightarrow & \dots
\end{array}$$

then we define

$$\mathrm{Tr}(u_\cdot) := \sum_i (-1)^i \,\mathrm{Tr}\, u_i.$$

Definition 3.1. We say $K_\cdot \in D(X)$ is *perfect* if it is isomorphic to a finite complex of finitely generated (stalk-)projective étale sheaves of Λ-modules. We denote by $D_{\mathrm{perf}}(X)$ the full subcategory of perfect complexes.

If we have a map $f \colon K_\cdot \to K_\cdot$ between perfect complexes in $D_c(X)$, we can try to define $\mathrm{Tr}\, f$ by picking an isomorphism α to a finite complex of finite projective modules $P_\cdot$, and setting

$$\mathrm{Tr}\, f := \mathrm{Tr}(\alpha \circ f \circ \alpha^{-1}, P_\cdot).$$

This depended on many choices, so it requires a little bit of work to show that this is well-defined.

The goal is to prove a more general formula

$$\sum_{x \in X(\mathbf{F}_q)} \operatorname{Tr}(F_x, K_x) = \sum_i (-1)^i \operatorname{Tr}(F, H^i_c(X, K))$$

for perfect complexes K. Implicit in this goal is assertion that $H^i_c(X, K)$ is a perfect complex (in $D(\operatorname{Spec} \mathbf{F}_q)$) if K was a perfect complex in $D(X)$. Our current definition isn't suited for checking something like this, so we need a more robust characterization.

Proposition 3.2. *$K \in D^b(X)$ is perfect if and only if the following two conditions are satisfied:*

1. *the homology sheaves $H^i(K)$ are constructible,*
2. *K has finite Tor-dimension.*

Proof sketch. The conditions are clearly necessary. For the converse, start building a quasi-isomorphism $P_. \to K_.$ with each P_i projective and finitely generated over Λ (possible by the finiteness assumption on the homology groups). This may continue indefinitely, but truncating below the Tor-dimension of K yields a quasi-isomorphic complex of flat and finitely generated, hence projective, Λ-modules. □

Theorem 3.3. *Let $f \colon X \to S$ be a map of schemes. If $K \in D_{\text{perf}}(X)$, then $Rf_! K \in D_{\text{perf}}(S)$.*

Proof. First we handle the constructibility. We have a spectral sequence

$$R^q f_!(\underline{H}^p(K)) \implies R^{p+q} f_!(K)$$

and by assumption $\underline{H}^p(K)$ is constuctible on X. Since $f_!$ preserves constructibility, $R^q f_!(\underline{H}^p(K))$ is also constructible, hence so is $R^{p+q} f_!(K)$.

Next let's examine the Tor-dimension. Let N be a sheaf on S. Then

$$(Rf_! K) \overset{L}{\otimes} N = Rf_!(K \overset{L}{\otimes} f^* N).$$

Since K has finite Tor-dimension, $K \overset{L}{\otimes} f^* N$ has bounded homology, so $Rf_!(K \overset{L}{\otimes} f^* N)$ also has bounded homology.

□

4 Reduction to the case of curves

We now put forth the general version of the trace formula (for Frobenius):

Theorem 4.1. *Let X_0 be a variety over $\mathbf{F}_q$ and $K_0 \in D_{\text{perf}}(X_0)$. Let $X = X_0 \times_{\mathbf{F}_q} \overline{\mathbf{F}}_q$ and K be the pullback of K_0 to X. Then we have*

$$\sum_{x \in X(\mathbf{F}_q)} \operatorname{Tr}(F_x, K_x) = \sum_i (-1)^i \operatorname{Tr}(F, H^i_c(X, K)) \tag{4.1}$$

with both sides understood in the sense of §3.2.

Repeating the sort of reasoning that we have used many times, we're going to reduce this to the case of curves, which we'll handle in the next section. In the reduction there will be one step which is rather subtle. We'll gloss over it at the start and then point it out later; the reader may amuse himself or herself by trying to spot it.

The basic mechanism for the reduction is as follows:

1. If $X_0 = U_0 \sqcup Y_0$ is a partition of X_0 into a closed subset Y_0 and its closed complement U_0, and (4.1) holds for U_0 and Y_0, then it holds for X_0.

2. If $f : X_0 \to S_0$ is a map and (4.1) holds for S_0, and all fibers of f, then it holds for X_0.

If we establish these two claims, and grant the case where X_0 is a curve, then we are done by induction on $\dim X$: we can chop up $X_0 = U_0 \sqcup Y_0$ such that U_0 admits a curve fibration to a variety of smaller dimension, and Y_0 is also "smaller" than X_0. So it suffices to prove the two assertions above. Let $T_l(X_0, K)$ be the LHS of (4.1) and $T_g(X_0, K)$ be the RHS (l for "local" and g for "global").

Proof of 1. Let $j : U_0 \hookrightarrow X_0$ and $i : Y_0 \hookrightarrow X_0$ be the two inclusion maps. It is easy to see that if $K_0 \in D_{\text{perf}}(X_0)$, then $j^* K_0 \in D_{\text{perf}}(U_0)$and $i^* K_0 \in D_{\text{perf}}(Y_0)$. Therefore, the assumption that (4.1) holds for U_0 and V_0 implies that

$$T_l(U_0, j^*K) = T_g(U_0, j^*K)$$
$$T_l(Y_0, i^*K) = T_g(Y_0, i^*K)$$

Since T_l is purely local, we manifestly have

$$T_l(U_0, j^*K) + T_l(Y_0, i^*K) = T_l(X_0, K).$$

By the additivity of the trace in long exact sequences, applied to the cohomology LES induced by the exact triangle

$$j_! j^* K \to K \to i_* i^* K$$

and the fact that $H_c^*(U_0, j^*K) = H_c^*(X_0, j_! j^*K)$ and $H_c^*(Y_0, i^*K) = H_c^*(X_0, i_* i^*K)$, we have also

$$T_g(U_0, j^*K) + T_g(Y_0, i^*K) = T_g(X_0, K), \tag{4.2}$$

thus establishing the claim in light of the preceding equations. □

Proof of 2. Let $f\colon X_0 \to S_0$ be a map. By Theorem 3.3, $Rf_!K \in D_{\text{perf}}(S)$. Then (4.1) applied to $Rf_!K$ tells us that

$$T_l(S_0, Rf_!K) = T_g(S_0, Rf_!K).$$

Since $\mathbf{R}\Gamma_c(X, Rf_!K) \cong R\Gamma_c(X, K)$ we have

$$T_g(S_0, Rf_!K) = T_g(X_0, K).$$

On the other hand,

$$T_l(S_0, Rf_!K) = \sum_{s \in S_0(\mathbf{F}_q)} \text{Tr}(F_s, (Rf_!K)_s)$$

and (4.1) applied to $f^{-1}(s)$ plus proper base change says that $\text{Tr}(F_s, (Rf_!K)_s) = T_l(f^{-1}(s), K_{f^{-1}(s)})$. Thus, we find that

$$T_l(S_0, Rf_!K) = \sum_{s \in S_0(\mathbf{F}_q)} T_l(f^{-1}(s), K_{f^{-1}(s)}) = T_l(X, K).$$

□

Now we should explain where we have lied. The subtlety is in (4.2). In fact it is *not* true in general that for an endomorphism of an exact triangle in $D_{\text{perf}}(X)$, the trace of the middle term is the sum of the traces of the outer terms! This is true for an endomorphism of complexes, but the point is that you can't always lift a "commutative diagram" in the derived category to a commutative diagram in chain complexes. The derived category is like the homotopy category, so you can lift it to a diagram which is "commutative up to homotopy", but it turns out that this isn't good enough.

However, the additivity (4.2) *does* hold in our situation. The point is that it should hold whenever the endomorphism is produced in a "canonical" way, which tends to always be the case in practice. There are at least two ways of articulating this.

- One is the original method of Deligne. The idea is to define an enhanced version of the derived category, where one starts with an abelian category consisting out of modules *plus* endomorphism. Then one takes the derived category of this abelian category, and any exact triangle in such induces an endomorphism of an exact triangle in the usual derived category. The point is that by baking the endomorphism into the objects before passing to homotopy, we ensure that the endomorphisms are "sufficiently canonical".

- Another is the "filtered derived category" approach of Illusie. Here one also enhances the abelian category, by carrying around the data of a filtration on the module. A two-step filtration then induces an exact triangle, and again an endomorphism produced in a filtered way is "sufficiently canonical" to enjoy the additivity property. The point is that the triangle $j_! j^* \to \mathrm{Id} \to i_* i^*$ can be lifted to the filtered derived category, essentially because it can be applied at the sheaf level.

5 The case of curves

5.1 Reductions and setup

Finally we are reduced to showing (4.1) for a curve. Weil's Theorem 2.1 tells us that (4.1) is true for a proper curve and constant sheaf (and arbitrary morphism). The idea is to reduce to this case, by chopping up the curve into pieces where the sheaves becomes locally constant, and then passing to a cover that trivializes them.

We begin with some simplifications. We can easily prove directly the 0-dimensional case of Theorem 4.1 (it is not totally content-free!), which by excision lets us cut out a finite set of closed points any time we wish.

- We may assume that K is a finite complex of finitely generated projective modules. By filtering K by its "naïve filtration", we reduce to the case where $K = \mathcal{F}$ is a constructible Λ-sheaf with finitely generated projective stalks.

- By restricting to a stratum, we may assume that $\mathcal{F}$ is locally constant.

- By replacing X by an open subset U, we may assume that X is smooth and irreducible.

- By partitioning X into the union of its rational points and their complement, we may assume that $X(\mathbf{F}_q) = \emptyset$. The goal is then to show that

$$\mathrm{Tr}(F, R\Gamma_c(X, \mathcal{F})) = 0.$$

Since $\mathcal{F}$ is locally constant, we can find a finite étale Galois cover $f \colon Y_0 \to X_0$ such that $f^*\mathcal{F}$ is constant. We can find smooth compactifications $X_0 \hookrightarrow \overline{X_0}$ and $Y_0 \hookrightarrow \overline{Y_0}$ extending to a ramified covering:

$$\begin{array}{ccc} Y & \hookrightarrow & \overline{Y} \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle \overline{f}} \\ X & \hookrightarrow & \overline{X} \end{array}$$

Roughly speaking, the idea is to try to compute $\mathrm{Tr}(F, R\Gamma_c(X, \mathcal{F}))$ in terms of $\mathrm{Tr}(F, R\Gamma_c(Y, f^*\mathcal{F})))$. The latter is something that we can sort of handle using Weil's Theorem 2.1.

There is a counit map

$$f_* f^* \mathcal{F} = f_! f^! \mathcal{F} \to \mathcal{F}.$$

Let $G = \mathrm{Gal}(Y_0/X_0)$. The sheaf $f_* f^* \mathcal{F}$ has an action of G coming from the fact that it is pushed forward from Y_0, while $\mathcal{F}$ doesn't, so this map factors through

$$(f_* f^* \mathcal{F})_G \to \mathcal{F}.$$

We can check on stalks that this is an isomorphism. Indeed, for $\overline{x} \in X$

$$(f_* f^* \mathcal{F})_{\overline{x}} = \bigoplus_{\overline{y} \in f^{-1}(\overline{x})} \mathcal{F}_{\overline{x}}$$

with the G-action induced by its permutation action on the G-torsor $f^{-1}(\overline{x})$.

Say $f^* M$ has value group $M := H^0(Y, f^* \mathcal{F})$. Then M is a $\Lambda[G]$-module, and

$$f^* f_* (\mathcal{F}) = f^* (\Lambda \otimes \underline{M}) = (f_* \Lambda) \otimes_\Lambda M$$

with the G-action being the diagonal one. Therefore,

$$\boxed{\mathcal{F} = ((f_* \Lambda) \otimes_\Lambda M) \otimes_{\Lambda[G]} \Lambda} \tag{5.1}$$

We claim that these tensor products really coincide with the derived tensor products. The only nonobvious point is why $(f_* \Lambda) \otimes_\Lambda M$ is projective over $\Lambda[G]$. Since $f_* \Lambda = \Lambda[G]$ is the regular representation, it is certainly projective over $\Lambda[G]$. Since M is flat over Λ, the tensor product $(f_* \Lambda) \otimes_\Lambda M$ is obviously projective if G-action is through the left factor. Unfortunately, this is not the case: the G-action is diagonal. But in fact, these are isomorphic as G-representations! In other words, there is an isomorphism of G-representations

$$\Lambda[G] \otimes M \to \Lambda[G] \otimes \underline{M}$$

where $\underline{M}$ has the trivial G-action, given by $g \otimes m \mapsto g \otimes g^{-1} m$.

So we have

$$\begin{aligned} R\Gamma_c(X, \mathcal{F}) &= R\Gamma_c(X, ((f_* \Lambda) \otimes_\Lambda M) \otimes_{\Lambda[G]} \Lambda) \\ &= R\Gamma_c(X, ((f_* \Lambda) \overset{L}{\otimes}_\Lambda M) \overset{L}{\otimes}_{\Lambda[G]} \Lambda) \\ &= (R\Gamma_c(X, f_* \Lambda) \otimes_\Lambda M) \otimes_{\Lambda[G]} \Lambda \\ &= (R\Gamma_c(Y, \Lambda) \otimes_\Lambda M) \otimes_{\Lambda[G]} \Lambda \end{aligned}$$

5.2 Computing the trace

We have just seen that

$$R\Gamma_c(X, \mathcal{F}) = (R\Gamma_c(Y, \Lambda) \otimes_\Lambda M) \otimes_{\Lambda[G]} \Lambda \tag{5.2}$$

We want to understand the action of Frobenius on the LHS in terms of the action of Frobenius on $\Gamma_c(X, f_*\Lambda)$. We begin by understanding the effect of the G-coinvariants.

More generally, suppose we have a projective $\Lambda[G]$-module P, for a finite group G, equipped with an additional endomorphism $F\colon P \to P$ commuting with the G-action, hence descending to $F\colon P_G \to P_G$. This is the same as saying that P is a module over $\Lambda[G \times \mathbf{N}]$.

Lemma 5.1. *We have*

$$|G| \operatorname{Tr}(F, P_G) = \sum_{g \in G} \operatorname{Tr}((g, F), P).$$

Remark 5.2. It is possible to prove a more refined equality

$$\operatorname{Tr}(F, P_G) = \frac{1}{|G|} \sum_{g \in G} \operatorname{Tr}((g, F), P).$$

First we have to give meaning to the right hand side. We can view P as a module over the group ring $\Lambda[G]$, and take its trace over $\Lambda[G]$, which takes values in the quotient by the subgroup generatored by commutators. The resulting module has a basis indexed by conjugacy classes, and the coefficient of the identity is a canonical "$|G|$th root" of $\sum_{g \in G} \operatorname{Tr}((g, F), P)$. For the details, see [1].

Proof. We want to prove the equality

$$|G| \operatorname{Tr}(F, P_G) = \sum_{g \in G} \operatorname{Tr}((g, F), P_G).$$

Now comes a trick: the endomorphism $\sum_{g \in G} g$ factors through $P \to P_G$:

$$\begin{array}{ccc} P & \xrightarrow{\pi} & P_G \\ \downarrow & \swarrow^{\Sigma} & \downarrow {\scriptstyle \sum(g,F)} \\ P & \xrightarrow[\pi]{} & P_G \end{array}$$

where Σ is defined by

$$\sum_{g \in G} (g, F) = \pi \circ \Sigma.$$

By the properties of the trace,

$$\operatorname{Tr}(\pi \circ \Sigma\colon P_G \to P_G) = \operatorname{Tr}(\Sigma \circ \pi\colon P \to P).$$

Therefore we deduce

$$|G| \operatorname{Tr}(F, P_G) = \sum_{g \in G} \operatorname{Tr}((g, F), P).$$

□

Using this in (5.2), we deduce that

$$|G|\operatorname{Tr}(F, R\Gamma_c(X,\mathcal{F})) = \sum_{g\in G} \operatorname{Tr}((g,F), R\Gamma_c(Y,\Lambda)\otimes_\Lambda M) \tag{5.3}$$

Here we recall that $R\Gamma_c(Y,\Lambda)\otimes_\Lambda M$ has the diagonal G-action, while the action of F is trivial on the M factor.

At this point we'd like to factorize the trace into the piece on $R\Gamma_c(Y,\Lambda)$ and the one on M, using $\operatorname{Tr}(A\otimes B) = \operatorname{Tr}(A)\operatorname{Tr}(B)$:

$$\operatorname{Tr}((g,F), R\Gamma_c(Y,\Lambda)\otimes_\Lambda M) = \operatorname{Tr}((g,F), R\Gamma_c(Y,\Lambda))\cdot \operatorname{Tr}((g,F), M).$$

Recall that we want to show that

$$\operatorname{Tr}(F, R\Gamma_c(X,\mathcal{F})) = 0. \tag{5.4}$$

If we show that

$$|G|\operatorname{Tr}(F, R\Gamma_c(X,\mathcal{F})) = 0 \tag{5.5}$$

for all coefficient rings $\Lambda = \mathbf{Z}/m\mathbf{Z}$ with m coprime to the characteristic of k, then (5.4) will follow for all such Λ. (The equation (5.5) for $\mathbf{Z}/m|G|\mathbf{Z}$ implies (5.4) for $\mathbf{Z}/m\mathbf{Z}$.) For this purpose, by (5.3) it suffices to show that $\operatorname{Tr}((g,F), R\Gamma_c(Y,\Lambda))$ vanishes for every g. In this case the cohomology endomorphism (g,F) is induced by a map of schemes $g^{-1}\circ F_{Y_0}\colon Y\to Y$:

$$\operatorname{Tr}((g,F), R\Gamma_c(Y,\Lambda)) = \operatorname{Tr}((g^{-1}\circ F_{Y_0})^*, R\Gamma_c(Y,\Lambda)).$$

5.3 Coup de grâce

We have reduced to showing that

$$\operatorname{Tr}((g^{-1}\circ F_{Y_0})^*, R\Gamma_c(Y,\Lambda)) = 0.$$

Consider the compactified ramified covering map

$$\begin{array}{ccc} Y & \hookrightarrow & \overline{Y} \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle \overline{f}} \\ X & \hookrightarrow & \overline{X} \end{array}$$

By Weil's Theorem 2.1, we know that

$$\operatorname{Tr}((g^{-1}\circ F_{Y_0})^*, R\Gamma_c(\overline{Y},\Lambda)) = \sum_{y\in \operatorname{Fix}(Fg)} \operatorname{Tr}(Fg, \Lambda_y)$$

Now by excision,

$$\operatorname{Tr}((g^{-1}\circ F_{Y_0})^*, R\Gamma_c(\overline{Y},\Lambda)) = \operatorname{Tr}((g^{-1}\circ F_{Y_0})^*, R\Gamma_c(Y,\Lambda)) + \operatorname{Tr}((g^{-1}\circ F_{Y_0})^*, R\Gamma_c(\overline{Y}-Y,\Lambda))$$

Let $i\colon \overline{Y} - Y \to \overline{Y}$ denote the inclusion (this is just the inclusion of finitely many closed points). By the 0-dimensional case of Theorem 4.1, we have

$$\mathrm{Tr}((g^{-1} \circ F_{Y_0})^*, R\Gamma_c(\overline{Y} - Y, \Lambda)) = \sum_{y \in \mathrm{Fix}(g^{-1} \circ F_{Y_0}) \setminus Y} \mathrm{Tr}(Fg, \Lambda_y).$$

Therefore, it suffices to see that there are no points $y \in Y$ fixed by $g^{-1} \circ F_{Y_0}$. Indeed, if there were then $f(y) \in X$ would be fixed by F_{X_0}, but by assumption X has no rational points. We are done.

References

[1] Deligne, Pierre. "Cohomologie étale", Séminaire de Géométrie Algébrique du Bois-Marie SGA $4\frac{1}{2}$. Lecture Notes in Mathematics, Vol. 569. Springer-Verlag, Berlin-New York, 1977. iv+312pp

[2] Freitag, Eberhard; Kiehl, Reinhardt. "Étale cohomology and the Weil conjecture". Translated from the German by Betty S. Waterhouse and William C. Waterhouse. With an historical introduction by J. A. Dieudonné. Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], 13. Springer-Verlag, Berlin, 1988. xviii+317 pp. ISBN: 3-540-12175-7