ON THE CONDUCTOR OF COHOMOLOGICAL TRANSFORMS

ÉTIENNE FOUVRY, EMMANUEL KOWALSKI, AND PHILIPPE MICHEL

Abstract. In the analytic study of trace functions of ℓ-adic sheaves over finite fields, a crucial issue is to control the conductor of sheaves constructed in various ways. We consider cohomological transforms on the affine line over a finite field which have trace functions given by linear operators with an additive character of a rational function in two variables as a kernel. We prove that the conductor of such transforms is bounded in terms of the complexity of the input sheaf and of the rational function defining the kernel, and discuss applications of this result, including motivating examples arising from the Polymath8 project.

Contents

1. Introduction

This paper considers a problem which appeared in special cases in [4, 6] in our study of analytic applications of trace functions over finite fields. We are given a constructible ℓ-adic

---

Date: January 17, 2019, 7:18.

2010 *Mathematics Subject Classification.* 11T23,14F20,11G20.

Key words and phrases. Étale cohomology, conductor, ℓ-adic sheaves, Riemann Hypothesis over finite fields, exponential sums.

Ph. M. was partially supported by the SNF (grant 200021-137488) and the ERC (Advanced Research Grant 228304). É. F. thanks ETH Zürich, EPF Lausanne and the Institut Universitaire de France for financial support. Ph.M. and E.K. were supported by the SNF-DFG Grant 200020L_175755.

sheaf $\mathcal{K}$ on $\mathbf{A}^1 \times \mathbf{A}^1$ (or, potentially, on another algebraic surface) over a finite field $\mathbf{F}_q$, and we use it to define a "cohomological transform" with "kernel" $\mathcal{K}$, that maps a constructible ℓ-adic sheaf $\mathcal{F}$ on $\mathbf{A}^1_{\mathbf{F}_q}$ to

$$T_{\mathcal{K}}(\mathcal{F}) = R^1 p_{1,!}(p_2^*\mathcal{F} \otimes \mathcal{K})(1/2),$$

where p_1, p_2 are the two projections $p_i : \mathbf{A}^1 \times \mathbf{A}^1 \to \mathbf{A}^1$. The problem is then to *estimate the conductor of* $T_{\mathcal{K}}(\mathcal{F})$, as defined in [4], in terms of that of $\mathcal{F}$.

The arithmetic interpretation of this problem, and our motivation, is that for suitable input sheaves $\mathcal{F}$ (as described later in more detail), the trace function $t_{\mathcal{G}}$ of $\mathcal{G} = T_{\mathcal{K}}(\mathcal{F})$ is related to the trace functions $t_{\mathcal{K}}$ and $t_{\mathcal{F}}$ by

$$t_{\mathcal{G}}(x) = -\frac{1}{\sqrt{q}} \sum_{y \in \mathbf{F}_q} t_{\mathcal{F}}(y) t_{\mathcal{K}}(x, y),$$

for all $x \in \mathbf{F}_q$. In other words, for all $\mathcal{F}$, we have $t_{\mathcal{G}} = \mathrm{T}_K(t_{\mathcal{F}})$, where

$$K(x, y) = t_{\mathcal{K}}(x, y)$$

and T_K is the (normalized) linear map defined on the space $C(\mathbf{F}_q)$ of complex-valued functions on $\mathbf{F}_q$ by the kernel K, i.e.,

$$\mathrm{T}_K(\varphi)(x) = -\frac{1}{\sqrt{q}} \sum_{y \in \mathbf{F}_q} K(x, y)\varphi(y).$$

The most important example of such transforms arises for $K(x, y) = \psi(xy)$, where ψ is a non-trivial additive character, which corresponds to $\mathcal{K} = \mathcal{L}_{\psi(XY)}$ (where X, Y are the coordinates on $\mathbf{A}^1 \times \mathbf{A}^1$): the corresponding linear operator T_K on trace functions is (minus the) normalized Fourier transform (which we denote also FT_ψ) on $C(\mathbf{F}_q)$, namely

$$\mathrm{FT}_\psi(\varphi)(x) = -\frac{1}{\sqrt{p}} \sum_{y \in \mathbf{F}_p} \varphi(y)\psi(xy).$$

The sheaf-theoretic construction, in that case, is due to Deligne, and it was extensively studied by Laumon [16].

This special case is crucial in [4] (and the following papers). In particular, it is essential for our applications that we have an estimate for the conductor of the Fourier transform $\mathcal{G}$ in terms only of the conductor of $\mathcal{F}$, which follows from the estimate

$$\mathbf{c}(\mathcal{G}) \leqslant 10\, \mathbf{c}(\mathcal{F})^2, \tag{1.1}$$

proved in [4, Prop. 8.2]. In order to establish this result, which we view as a form of "continuity" of the sheaf-theoretic Fourier transform, we used the deep theory of the local Fourier transform of Laumon [16, 11].

The general case of these transforms is a natural approach to estimates for two-variable character sums (and more complicated algebraic sums) based on Deligne's work, and an estimate for the conductor leads for instance easily to strong statements of "control of the diagonal" (see Proposition 5.6 for a precise statement.)

It is not at all clear if a local theory like Laumon's applies to the general transforms we consider (the same applies to the theory of "affine cohomological transforms" of Katz [13], which has diophantine applications to stratification results for sums of trace functions.) Thus, our present goal is to prove estimates for more general cohomological transforms. These will be weaker than (1.1), but more accessible. We will be able to do so when the kernel $\mathcal{K}$ is a rather general Artin-Schreier sheaf, or in other words (in the case when $q = p$ is prime) when

$$K(x, y) = e\Big(\frac{f(x, y)}{p}\Big)$$

for some rational function $f \in \mathbf{F}_p(X, Y)$.

The precise statement is given in Theorem 2.3 in the next section. In the case of the Fourier transform, this gives a form of the important property (1.1) which is more accessible than Laumon's theory. Section 8 treats this case fully, in order to motivate and clarify the algebraic tools used in the general case. In Section 5, we discuss some first applications of these bounds; for instance, we show how the ideas lead to an account of the character sums considered by Conrey and Iwaniec in [1]. This section can, to a large extent, be read independently of the part of the paper where the main results are proved.

Remark 1.1. (1) In work in progress, W. Sawin has developed a much more general and powerful theory of complexity measures of ℓ-adic sheaves on schemes, including all so-called 6 operations and derived category objects. His work subsumes ours entirely, but also involves much more difficult algebraic geometry. We (finally) submit the present paper for publication as an illustration of fairly simple manipulations of the formalism of étale cohomology, in the hope that it will be helpful to readers with a more analytic background.

(2) In recent work, I. Petrow and M. Young [18] generalized the estimate of Conrey and Iwaniec to more general characters. They need to estimate slightly different sums than those in [1], and the first draft of their preprint refered to this paper for this purpose. W. Sawin has also observed that their sums (as those of Conrey and Iwaniec) and special cases of hypergeometric sums, and can be directly estimated by an simple appeal to Katz's book [11].

(3) A slightly different definition of the conductor suggested by W. Sawin leads to better estimates (e.g., a linear bound for the conductor of the Fourier transform instead of (1.1)). Since our work is in any case very restricted (see Remark (1)), and the most important qualitative feature is not affected for applications, we have not incoroporated all the changes required by this adjustment.

Acknowledgments. Part of the original motivation for this paper in 2013/2104 arose in online discussions related to the Polymath8 project.

Notation. By "sheaf", or "ℓ-adic sheaf", we will always mean "constructible $\bar{\mathbf{Q}}_\ell$-sheaf", where ℓ will be a prime number different from the characteristic of the base field.

For a power $q \neq 1$ of a prime p and any integer $w \in \mathbf{Z}$, a *q-Weil number of weight w* is an algebraic number $\alpha \in \mathbf{C}$ such that all Galois-conjugates β of α satisfy $|\beta| = q^{w/2}$.

An algebraic variety over a field k is a finite type, separated, reduced scheme over k. If X/k is an algebraic variety over a field k, and $\bar{k}$ is an algebraic closure of k, we denote by $\bar{X}$ or $X_{\bar{k}}$ the base change $X \times \bar{k}$.

If X_k is an algebraic variety over a field k and $x \in X(k)$, we denote by $\bar{x}$ a geometric point above x. If k is algebraically closed, we take $\bar{x} = x$. If $\mathcal{F}$ is an étale sheaf on X, then $\mathcal{F}_{\bar{x}}$ denotes the stalk of $\mathcal{F}$ at $\bar{x}$.

Whenever a prime ℓ is given, we assume *fixed* an isomorphism $\iota : \bar{\mathbf{Q}}_\ell \to \mathbf{C}$, and we use it as an implicit identification.

For any ℓ-adic sheaf $\mathcal{F}$ on an algebraic variety $X_{\mathbf{F}_q}$, we write $t_{\mathcal{F}}(x)$ for the value at x of the trace function of $\mathcal{F}$, i.e., we have

$$t_{\mathcal{F}}(x) = \iota((\operatorname{tr}\mathcal{F})(\mathbf{F}_q, x)),$$

the trace of the Frobenius of $\mathbf{F}_q$ acting on the stalk of $\mathcal{F}$ at x.

If $k/\mathbf{F}_q$ is a finite extension, we write

$$t_{\mathcal{F}}(x, |k|) = t_{\mathcal{F}}(x, k) = \iota((\operatorname{tr}\mathcal{F})(k, x)).$$

2. Statement of the main result

We first recall the definition of the conductor of a constructible ℓ-adic sheaf $\mathcal{F}$ on the affine line over a finite field $\mathbf{F}_q$. Indeed, since in this work it will be important to work with general constructible sheaves, and not only the middle-extension sheaves considered in our previous works, we need to adapt the definition slightly.

Let $\mathcal{F}$ be a constructible ℓ-adic sheaf over $\mathbf{A}^1_{\mathbf{F}_q}$. Let $U \subset \mathbf{P}^1$ be the maximal dense open subset where $\mathcal{F}$ is lisse. Let $j : U \hookrightarrow \mathbf{P}^1$ be the corresponding open immersion. Recall that there is a canonical (adjunction) map

$$\mathcal{F} \longrightarrow j_*j^*\mathcal{F},$$

and that $\mathcal{F}$ is said to be a *middle-extension sheaf* if this is an isomorphism. In general, if we let

$$\mathcal{F}_0 = j_*j^*\mathcal{F},$$

then one shows that $\mathcal{F}_0$ is a middle-extension sheaf on $\mathbf{P}^1_{\mathbf{F}_q}$, which is isomorphic to $\mathcal{F}$ when restricted to U. We define

$$\mathbf{c}(\mathcal{F}) = \operatorname{rank}(\mathcal{F}_0) + \sum_x \operatorname{Swan}_x(\mathcal{F}_0) + n(\mathcal{F}) + \operatorname{pct}(\mathcal{F}),$$

where:

– $n(\mathcal{F}) = |(\mathbf{P}^1 - U)(\bar{\mathbf{F}}_q)|$ is the number of singularities of $\mathcal{F}$ in $\mathbf{P}^1(\bar{\mathbf{F}}_q)$;

– the sum is over $\mathbf{P}^1(\bar{\mathbf{F}}_q)$, with all but finitely many terms vanishing;

– we define

$$\operatorname{pct}(\mathcal{F}) = \dim H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F}).$$

Remark 2.1. (1) If $\mathcal{F}$ is a middle-extension sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$, we have $\mathcal{F} = \mathcal{F}_0$ (on $\mathbf{A}^1$) and

$$\mathbf{c}(\mathcal{F}) = \operatorname{rank}(\mathcal{F}) + \sum_x \operatorname{Swan}_x(\mathcal{F}) + n(\mathcal{F}),$$

as in our previous works.

(2) Let $\mathcal{P}$ be the kernel of the map

$$\mathcal{F} \longrightarrow j_*j^*\mathcal{F}.$$

Then $\mathcal{P}$ has finite support; if this support is $S \subset \mathbf{A}^1(\bar{\mathbf{F}}_q)$, then

$$|S| \leqslant \operatorname{pct}(\mathcal{F}) \leqslant \sum_{s\in S} \dim \mathcal{F}_{\bar{s}}$$

(see [9, §4.4, 4.5] for a discussion).

(3) Note that $n(\mathcal{F})$ takes into account the fact that a general constructible sheaf might have "artificial" singularities, which are not singularities of the associated middle-extension

sheaf. These may also be seen as the contribution to the conductor of the cokernel $\mathcal{F}$ of the map

$$\mathcal{F} \longrightarrow j_* j^* \mathcal{F},$$

which is also a sheaf with finite support.

For instance, let $U = \mathbf{P}^1 - \mathbf{A}^1(\mathbf{F}_p)$ over $\mathbf{F}_p$, and let

$$j \,:\, U \longrightarrow \mathbf{P}^1$$

be the open immersion. Consider

$$\mathcal{F} = j_! \bar{\mathbf{Q}}_\ell,$$

the extension by zero to $\mathbf{P}^1$ of the trivial sheaf on U. Then $\mathcal{F}_0$ is the trivial sheaf on $\mathbf{P}^1$, with $n(\mathcal{F}_0) = 0$, and $\mathbf{c}(\mathcal{F}_0) = 1$, while $\mathbf{c}(\mathcal{F}) = 1 + n(\mathcal{F}) = 1 + |\mathbf{A}^1(\mathbf{F}_p)| = p + 1$ because of the artificial singularities created at the points in $\mathbf{A}^1(\mathbf{F}_p)$. It is necessary here to have a big conductor if we want some basic qualitative features of the Riemann Hypothesis to hold.

We note the following useful property:

$$\mathbf{c}(\mathcal{F}_1 \oplus \mathcal{F}_2) \leqslant \mathbf{c}(\mathcal{F}_1) + \mathbf{c}(\mathcal{F}_2) \tag{2.1}$$

for two constructible sheaves on $\mathbf{A}^1$ (more generally, if

$$0 \longrightarrow \mathcal{F}_1 \longrightarrow \mathcal{F}_3 \longrightarrow \mathcal{F}_2 \longrightarrow 0$$

is a short exact sequence of constructible sheaves on $\mathbf{A}^1$, then we have

$$\mathbf{c}(\mathcal{F}_3) \leqslant \mathbf{c}(\mathcal{F}_1) + \mathbf{c}(\mathcal{F}_2)$$

as one can check.)

We also recall the definition of the drop of a constructible sheaf $\mathcal{F}$ on $\mathbf{A}^1_{\bar{\mathbf{F}}_q}$ at a point $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$: we have

$$\mathrm{drop}_x(\mathcal{F}) = \mathrm{rank}(\mathcal{F}_0) - \dim \mathcal{F}_x, \tag{2.2}$$

where $\mathcal{F}_x$ is the stalk of $\mathcal{F}$ at x. Note that the rank of $\mathcal{F}_0$ is also the "generic" rank of $\mathcal{F}$, i.e., the dimension of the fiber at a geometric generic point.

As mentioned in the introduction, we consider in this paper a kernel $\mathcal{K}$ which is an Artin-Schreier sheaf, with trace function given by additive characters of rational function. We give a formal definition to avoid any ambiguity concerning the behavior at the poles or points of indeterminacy of a rational function in two variables.

Definition 2.2 (Artin-Schreier sheaf on $\mathbf{A}^n$)**.** Let $\mathbf{F}_q$ be a finite field of characteristic p, and let $\ell \neq p$ be a prime number, and ψ a non-trivial additive ℓ-adic character of $\mathbf{F}_q$. Let $\mathcal{L}_\psi$ denote the associated Artin-Schreier sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$.

Let $f \in \mathbf{F}_q(X_1, \ldots, X_n)$ be a rational function for some $n \geqslant 1$. Write $f = f_1/f_2$ where $f_i \in \mathbf{F}_q[X_1, \ldots, X_n]$ and where f_1 is coprime with f_2. Let $U \subset \mathbf{A}^n$ be the open set where f_2 is invertible, $j \,:\, U \hookrightarrow \mathbf{A}^n$ the corresponding open immersion, and let

$$f_U \,:\, U \longrightarrow \mathbf{A}^1$$

be the morphism associated to the restriction of f to U.

The *Artin-Schreier sheaf on* $\mathbf{A}^n$ *associated to* f is the constructible ℓ-adic sheaf on $\mathbf{A}^n_{\mathbf{F}_q}$ given by

$$\mathcal{L}_{\psi(f)} = j_! f_U^* \mathcal{L}_\psi.$$

We also write $\mathcal{L}_{\psi(f(X_1,\ldots,X_n))}$ for this sheaf. We define its conductor to be

$$\mathbf{c}(\mathcal{L}_{\psi(f)}) = 1 + \deg(f_1) + \deg(f_2),$$

and we will also sometimes just speak of the conductor $\mathbf{c}(f)$ of f.

We will find a satisfactory generalization of (1.1) for transforms associated to a kernel which is an Artin-Schreier sheaf.

Theorem 2.3 (Conductor of Artin-Schreier transforms)**.** *Let $\mathbf{F}_q$ be a finite field of order q and characteristic p, ℓ a prime distinct from p. Let $\mathcal{K}$ be an ℓ-adic sheaf on $\mathbf{A}^1 \times \mathbf{A}^1$ over $\mathbf{F}_q$ of the form*

$$\mathcal{K} = \mathcal{L}_{\psi(f(X,Y))},$$

where ψ is a non-trivial additive ℓ-adic character and $f \in \mathbf{F}_q(X, Y)$ is a rational function with conductor $< p$.

For constructible sheaves $\mathcal{F}$ on $\mathbf{A}^1_{\mathbf{F}_q}$, and $0 \leqslant i \leqslant 2$, let

$$T^i_{\mathcal{K}}(\mathcal{F}) = R^i p_{1,!}(p_2^*\mathcal{F} \otimes \mathcal{K}).$$

There exists an integer $A \geqslant 1$ such that for any middle-extension sheaf $\mathcal{F}$ on $\mathbf{A}^1_{\mathbf{F}_q}$, and $0 \leqslant i \leqslant 2$, we have

$$\mathbf{c}(T^i_{\mathcal{K}}(\mathcal{F})) \leqslant (2\,\mathbf{c}(\mathcal{K})\,\mathbf{c}(\mathcal{F}))^A.$$

In particular, if f is obtained by reduction modulo p of a fixed non-constant rational function f_1/f_2, where $f_i \in \mathbf{Z}[X, Y]$, and if we have some integer $M \geqslant 1$ and, for each p, we consider a sheaf $\mathcal{F}_p$ modulo p with conductor $\leqslant M$, then we have

$$\mathbf{c}(T^1_{\mathcal{K}}\mathcal{F}_p) \ll 1$$

for all primes. This allows us to apply all our estimates for trace functions to the trace functions of these sheaves; we give some examples in Section 5.

Remark 2.4. (1) The statement makes sense because it is known that $T^i_{\mathcal{K}}(\mathcal{F})$ is itself, for all constructible sheaves $\mathcal{F}$ and all i, a constructible sheaf (see, e.g., [2, Arcata, IV, Th. 6.2] or [8, Th. 7.8.1].)

(2) Note that we omitted the Tate twist in this statement, since it concerns purely geometric and algebraic facts.

We consider all transforms $T^i_{\mathcal{K}}$, and not only $T^1_{\mathcal{K}}$ because this will turn out to be useful in the proof, which will be interleaved with the proof of the following other useful fact:

Theorem 2.5 (Bounds for Betti numbers)**.** *Let $\mathbf{F}_q$ be a finite field of order q and characteristic p, ℓ a prime distinct from p. Let $\mathcal{K}$ be an ℓ-adic sheaf on $\mathbf{A}^1 \times \mathbf{A}^1$ over $\mathbf{F}_q$ of the form*

$$\mathcal{K} = \mathcal{L}_{\psi(f(X,Y))},$$

where ψ is a non-trivial additive ℓ-adic character and $f \in \mathbf{F}_q(X, Y)$ is a rational function with conductor $< p$.

There exists an integer $B \geqslant 1$ such that for any middle-extension sheaf $\mathcal{F}$ on $\mathbf{A}^1_{\mathbf{F}_q}$ and for $0 \leqslant i \leqslant 4$, we have

$$\dim H^i_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, p_2^*\mathcal{F} \otimes \mathcal{K}) \leqslant (2\,\mathbf{c}(f)\,\mathbf{c}(\mathcal{F}))^B.$$

Roughly speaking, we will proceed as follows: (1) we prove Theorem 2.3 for $\mathcal{F}$ the trivial sheaf, and observe that Theorem 2.5 is a known fact in that case, from bounds on Betti numbers due to Bombieri, Adolphson-Sperber and Katz; (2) using Theorem 2.3 for the trivial sheaf, we first prove Theorem 2.5 for all input sheaves $\mathcal{F}$ and $i = 2$; (3) finally, we prove Theorem 2.3 in general and deduce Theorem 2.5 for all i.

3. Diophantine motivation of the proof

The arguments of the proof of Theorem 2.3 are purely algebraic and geometric, and exercise much of the basic formalism of étale cohomology, as well as a simple use of spectral sequences. However, there is a concrete analytic motivation from (expected) properties of sums of trace functions, and we will first present it. This is based on the Riemann Hypothesis over finite fields, and is similar in principle to the discussion [15, Lecture IV, Interlude] by Katz that motivates the crucial step in his paper.

The first ingredient is a lemma that, essentially, allows one to estimate, in terms of accessible global invariants, the conductor of a middle-extension sheaf, satisfying some conditions, assuming one already knows estimates for the rank and the number of singularities. In other words, it provides a bound for the sum of Swan conductors in global terms, assuming that the rank and number of singularities are under control.

To be slightly more precise, assume that $\mathcal{F}$ is a middle-extension sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$ which is pointwise pure of weight 0, and assume in addition the following conditions:

(1) $\mathcal{F}$ has no geometrically trivial Jordan-Hölder factor;

(2) the Frobenius action on $H^1_c(\bar{U}, \mathcal{F})$ is *pure* of weight 1, for the maximal dense open set U on which $\mathcal{F}$ is lisse.

We then define the invariant

$$\tilde{\sigma}(\mathcal{F}) = \limsup_{\nu \to +\infty} \frac{|S_\nu(\mathcal{F})|}{q^{\nu/2}},$$

where

$$S_\nu(\mathcal{F}) = \sum_{x \in U(\mathbf{F}_{q^\nu})} t_{\mathcal{F}}(x, q^\nu),$$

for $\nu \geqslant 1$ (in other words, these are the sums of trace functions over extension fields). Then we have

$$\mathbf{c}(\mathcal{F}) \leqslant 3 \operatorname{rank}(\mathcal{F}) + n(\mathcal{F}) + \tilde{\sigma}(\mathcal{F}).$$

Indeed, using (1), the Lefschetz trace formula applied to U over $\mathbf{F}_{q^\nu}$ gives

$$S_\nu(\mathcal{F}) = -\operatorname{tr}(\mathrm{Fr}^\nu \mid H^1_c(\bar{U}, \mathcal{F})),$$

so that the purity assumption implies

$$\tilde{\sigma}(\mathcal{F}) = \dim H^1_c(\bar{U}, \mathcal{F})$$

and then the stated bound follows from Lemma 4.11 below (which is an elementary application of the Euler-Poincaré formula.)

We now consider the situation of Theorem 2.3. We will assume (and this is where the argument is not easy to make rigorous in a decent generality) that the sheaves $\mathcal{G} = T_{\mathcal{K}}(\mathcal{F})$ whose conductor we wish to control always satisfy the conditions above (i.e., that they are middle-extensions, pointwise of weight 0, and (1), (2) hold.) We first assume that we can find suitable estimates of the rank, of the number of singularities, and of the punctual part of

$\mathcal{G}$ (intuitively, this is possible because these amounts to fiber-by-fiber considerations, which boil down to properties of one-variable sheaves, which are relatively well-understood; the case of the trivial sheaf $\mathcal{F}$ is quite elementary, but the details will turn out to be a bit involved in the general case). We then need to estimate $\tilde{\sigma}(\mathcal{G})$. For this purpose, we proceed in two steps.

In Step 1, we consider only the trivial input sheaf $\mathcal{F} = \bar{\mathbf{Q}}_\ell$. We then have

$$\frac{S_\nu(\mathcal{G})}{q^{\nu/2}} = -\frac{1}{q^\nu}\sum_{x\in\mathbf{F}_{q^\nu}}\Bigl(\sum_{y\in\mathbf{F}_{q^\nu}}\psi_\nu(f(x,y))\Bigr) = -\frac{1}{q^\nu}\sum_{(x,y)\in\mathbf{F}_{q^\nu}\times\mathbf{F}_{q^\nu}}\psi_\nu(f(x,y)) \tag{3.1}$$

and the two-variable character sum (under Assumption (2) for $\mathcal{G}$) has square-root cancellation, so that the bounds on Betti numbers of [14] (or often their predecessors, due to Bombieri and Adolphson-Sperber) give

$$\limsup_{\nu\to+\infty}\frac{|S_\nu(\mathcal{G})|}{q^{\nu/2}} \leqslant C$$

where $C \geqslant 1$ depends only on the conductor of f.

In Step 2, we handle the case of a general sheaf $\mathcal{F}$. We then have

$$\frac{S_\nu(\mathcal{G})}{q^{\nu/2}} = -\frac{1}{q^\nu}\sum_{x\in\mathbf{F}_{q^\nu}}\sum_{y\in\mathbf{F}_{q^\nu}}t_{\mathcal{F}}(y,q^\nu)\psi_\nu(f(x,y)) = -\frac{1}{q^\nu}\sum_{y\in\mathbf{F}_{q^\nu}}t_{\mathcal{F}}(y,q^\nu)\sum_{x\in\mathbf{F}_{q^\nu}}\psi_\nu(f(x,y)). \tag{3.2}$$

The basic point is that this is the inner-product of the trace functions of the dual sheaf of $\mathcal{F}$ and of the sheaf $R^1p_{2,!}\mathcal{L}_{\psi(f(X,Y))}$. This last sheaf, *by the first step* (applied to $\mathcal{L}_{\psi(f(Y,X))}$), has conductor bounded by a constant depending only on the conductor of f. By assumption again, we have square-root cancellation in this sum as $\nu \to +\infty$, and by the quasi-orthogonality formulation of Deligne's proof of the Riemann Hypothesis over finite fields [3], we obtain

$$\tilde{\sigma}(\mathcal{G}) = \limsup_{\nu\to+\infty}\frac{|S_\nu(\mathcal{G})|}{q^{\nu/2}} \leqslant C',$$

where C' depends only on the conductors of $\mathcal{F}$ and of f.

Remark 3.1. In terms of linear operators and of the standard (unnormalized) inner-product on functions on $\mathbf{F}_q$, we exploit the obvious identity

$$\sum_{x\in\mathbf{F}_q}(\mathrm{T}_K\,\varphi)(x) = \langle \mathrm{T}_K\,\varphi, 1\rangle = \langle\varphi, \mathrm{T}_K^*\,1\rangle,$$

where the adjoint operator T_K^* has kernel $K^*(x,y) = \overline{K(y,x)}$; the first step in our sketch amounts to bounding (the complexity of) $\mathrm{T}_K^*\,1$, and the second applies standard inequalities to deduce a bound for the sum over x.

In contrast with this sketch, the proof of Theorem 2.3 below is entirely algebraic and does not require the Riemann Hypothesis over finite fields. It also applies in greater generality, so that the assumptions (1) and (2) are not needed. Roughly speaking, instead of sums of trace functions, we control directly the dimension $\tilde{\sigma}(\mathcal{G})$ of $H^1_c(\bar{U},\mathcal{G})$ for the transformed sheaf $\mathcal{G}$. The "combination of sums" in (3.1) and the "exchange of order of summation" in (3.2) are replaced by arguments based on spectral sequences (compare again with [15, Lecture IV, Interlude], and the dictionary [2, Sommes Trig., §2].) The proof is however complicated by the fact that we must also control the possible punctual part of the transformed sheaf.

Before giving the proof, we will present some algebraic preliminaries and then discuss first the motivating applications in Section 5 (Section 4 may be skipped in a first reading, since Section 5 will only refer to it incidentally). We then set up the proof in Section 6, and follow by presenting an (almost) self-contained account of the Fourier transform and of the special case which is relevant to the Polymath8 project (see Section 8). Finally, we give the full proof of Theorems 2.3 and 2.5.

4. Preliminaries

We begin by some preliminary results.

4.1. Étale cohomology.

We first state formally some properties of étale cohomology that we will often use.

Proposition 4.1. (1) *Let $f : Y_k \longrightarrow X_k$ be a morphism of algebraic varieties over an algebraically closed field k, with fibers of dimension $\leqslant n$. Let $\mathcal{F}$ be a constructible ℓ-adic sheaf on Y. We have $R^i f_! \mathcal{F} = 0$ for $i < 0$ and for $i > 2n$. In particular, if $\mathcal{F}$ is a sheaf on X and X has dimension $\leqslant n$, we have $H^i_c(X, \mathcal{F}) = 0$ for $i < 0$ and for $i > 2n$.*

(2) *Let X_k be an algebraic variety over an algebraically closed field k, let $U \subset X$ be an open subset and $C = X - U$ its complement. Let $\mathcal{F}$ be a constructible ℓ-adic sheaf on X. We have a long sequence sequence*

$$\cdots \longrightarrow H^i_c(U, \mathcal{F}) \longrightarrow H^i_c(X, \mathcal{F}) \longrightarrow H^i_c(C, \mathcal{F}) \longrightarrow H^{i+1}_c(U, \mathcal{F}) \longrightarrow \cdots, \tag{4.1}$$

and in particular, for all $i \geqslant 0$, we have

$$\dim H^i_c(X, \mathcal{F}) \leqslant \dim H^i_c(U, \mathcal{F}) + \dim H^i_c(C, \mathcal{F}) \tag{4.2}$$

$$\dim H^i_c(U, \mathcal{F}) \leqslant \dim H^i_c(X, \mathcal{F}) + \dim H^{i-1}_c(C, \mathcal{F}). \tag{4.3}$$

(3) *Let X_k be a smooth algebraic variety over an algebraically closed field k, pure of dimension $n \geqslant 0$, and let $\mathcal{F}$ be a lisse ℓ-adic sheaf on X. We have*

$$H^i_c(X, \mathcal{F}) = 0 \tag{4.4}$$

for $0 \leqslant i < n$.

(4) *Let $f : X_k \longrightarrow Y_k$ be a morphism of algebraic varieties over an algebraically closed field k, and let $\mathcal{F}$ be an ℓ-adic constructible sheaf on X. Then, for $y \in Y$ and $i \geqslant 0$, the stalk of $R^i f_! \mathcal{F}$ at y is naturally isomorphic to $H^i_c(f^{-1}X, \mathcal{F})$.*

Proof. (1) is the cohomological dimension property; the vanishing of $R^i f_! \mathcal{F}$ for $i < 0$ is immediate by definition, while the vanishing for $i > 2n$ can be found, e.g., in [2, Arcata, IV, Th. 6.1] or [8, th. 7.4.5]; the case of H^i_c follows by considering $f : X \longrightarrow \mathrm{Spec}(k)$, the structure morphism.

(2) is the so-called "excision" long-exact sequence, see for instance [2, Sommes Trig., (2.5.1)*]); the inequality (4.2) for a given $i \geqslant 0$ is an immediate consequence of the fragment

$$H^i_c(U, \mathcal{F}) \longrightarrow H^i_c(X, \mathcal{F}) \longrightarrow H^i_c(C, \mathcal{F}),$$

and (4.3) is a consequence of

$$H^{i-1}_c(C, \mathcal{F}) \longrightarrow H^i_c(U, \mathcal{F}) \longrightarrow H^i_c(X, \mathcal{F}).$$

(3) is the property of affine cohomological dimension for lisse sheaves; it follows for instance from the Poincaré duality

$$H^i_c(X,\mathcal{F}) \simeq H^{n-i}(X,\mathcal{F}^*)$$

where $\mathcal{F}^*$ is the dual of $\mathcal{F}$ (see for instance [2, Sommes Trig., Remarque 1.18 (c)]; note that the right-hand side is a cohomology group with no restriction of compact support) and the vanishing property

$$H^i(X,\mathcal{F}) = 0$$

for an affine scheme X and $i < \dim(X)$ (see, e.g., [2, Arcata, IV, Th. 6.4]).

(4) is a special case of the proper base change theorem, (see, e.g., [2, Arcata, IV, Th. 5.4] or [8, Th. 7.4.4 (i)]). ☐

The following lemma will also be used frequently:

Lemma 4.2. *Let $\mathcal{F}$ and $\mathcal{G}$ be middle-extension ℓ-adic sheaves on $\mathbf{A}^1_{\mathbf{F}_q}$. Then we have*

$$H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F}\otimes\mathcal{G}) = 0,$$

i.e., the tensor product has no punctual part.

Proof. In general, for a constructible sheaf $\mathcal{H}$ lisse on a dense open set $U \subset \mathbf{A}^1$, the condition

$$H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{H}) = 0$$

amounts to saying that, for all $x \in (\mathbf{A}^1 - U))(\bar{\mathbf{F}}_q)$, the specialization map

$$\mathcal{H}_x \longrightarrow \mathcal{H}_{\bar{\eta}}^{I_x}$$

is injective (see [9, §4.4] for instance), where $I_{\bar{x}}$ is the inertia group at x. We now have

$$(\mathcal{F}\otimes\mathcal{G})_x = \mathcal{F}_x \otimes \mathcal{G}_x \hookrightarrow \mathcal{F}_{\bar{\eta}}^{I_x} \otimes \mathcal{G}_{\bar{\eta}}^{I_x} \subset (\mathcal{F}_{\bar{\eta}} \otimes \mathcal{G}_{\bar{\eta}})^{I_x} = (\mathcal{F}\otimes\mathcal{G})_{\bar{\eta}}^{I_x}.$$

☐

4.2. **Properties of Artin-Schreier sheaves.** We next recall the crucial link between Swan conductor and order of a pole for a rational function in one variable.

Lemma 4.3. *Let $\mathbf{F}_q$ be a finite field of order q and characteristic p, $\ell \neq p$ a prime number. Let $\mathcal{L} = \mathcal{L}_{\psi(g(X))}$ be an ℓ-adic Artin-Schreier sheaf on $\mathbf{A}^1$ over $\mathbf{F}_q$, where $g \in \mathbf{F}_q(X)$ is a non-constant rational function. For $x \in \mathbf{P}^1(\bar{\mathbf{F}}_q)$, the Swan conductor of $\mathcal{L}$ at x is at most equal to the order of the pole of g at x, and there is equality if the numerator and denominator of g have degree $< p$.*

Proof. This is a standard property (see, for instance, [2, Sommes Trig., (3.5.4)]). ☐

We next discuss relations between two-variable Artin-Schreier sheaves and specializations of one variable. We need first some notation.

Definition 4.4 (Specializations)**.** Let $\mathbf{F}_q$ be a finite field of order q and characteristic p, ℓ a prime distinct from p. Let $f \in \mathbf{F}_q(X,Y)$ be a non-constant rational function.

(1) If $x \in \bar{\mathbf{F}}_q$ is such that $X - x$ does not divide the denominator of f, we denote by $f_x \in \mathbf{F}_q(Y)$ the specialization $f(x,Y)$ of f.

(2) Let $\mathcal{L} = \mathcal{L}_{\psi(f(X,Y))}$ be the Artin-Schreier sheaf on $\mathbf{A}^2_{\mathbf{F}_q}$ associated to f. For every finite extension $k/\mathbf{F}_q$ and every $x \in k$, the *specialization of* $\mathcal{L}$ *at* x is the ℓ-adic constructible sheaf on $\mathbf{A}^1_k$ given by

$$\mathcal{L}_x = j_x^* \mathcal{L}, \tag{4.5}$$

where $j_x : \{x\} \times \mathbf{A}^1 \hookrightarrow \mathbf{A}^2$ is the closed immersion.

These two definitions are related as follows:

Lemma 4.5. *Let $\mathbf{F}_q$ be a finite field of order q and characteristic p, $\ell \neq p$ a prime number. Let $\mathcal{L} = \mathcal{L}_{\psi(f(X,Y))}$ be an ℓ-adic Artin-Schreier sheaf on $\mathbf{A}^2$ over $\mathbf{F}_q$, where $f \in \mathbf{F}_q(X,Y)$ is a rational function.*

(1) *For any finite extension $k/\mathbf{F}_q$ and $x \in k$, we have*

$$\mathcal{L}_x = 0$$

if $X - x$ divides the denominator of f, and otherwise

$$\mathcal{L}_x = j_! \mathcal{L}_{\psi(f_x)}$$

where $j : U_x \longrightarrow \{x\} \times \mathbf{A}^1$ is the open immersion of the open subset of $\{x\} \times \mathbf{A}^1$ which is the intersection of $\{x\} \times \mathbf{A}^1$ and the open set of $\mathbf{A}^2$ where the denominator of f is invertible.

If $\{x\} \times \mathbf{A}^1$ does not intersect the zero set of the numerator of f, then $\mathcal{L}_x$ is isomorphic to the Artin-Schreier sheaf $\mathcal{L}_{\psi(f_x)}$ associated to f_x.

(2) *For every finite extension $k/\mathbf{F}_q$ and all $x \in k$, we have*

$$\mathbf{c}(\mathcal{L}_x) \leqslant 2\,\mathbf{c}(f).$$

Proof. (1) If $X - x$ divides the denominator of f, then by definition the sheaf $\mathcal{L}$ is zero on $\{x\} \times \mathbf{A}^1$, and hence $\mathcal{L}_x = 0$.

If $X - x$ does not divide the denominator of f, then there are only finitely many points where $\{x\} \times \mathbf{A}^1$ intersects the open set U where the denominator is invertible. The sheaf $\mathcal{L}_x$ has zero stalk at these points, and is isomorphic to the one-variable sheaf $\mathcal{L}_{\psi(f(x,Y))}$ on the complementary open set, which is the result we claim.

If $\{x\} \times \mathbf{A}^1$ does not intersect the zero set of the numerator of f, then the points in $\{x\} \times \mathbf{A}^1$ where $\mathcal{L}_x$ has zero stalk are precisely the poles of f_x, which means that $j_! \mathcal{L}_{\psi(f_x)} = \mathcal{L}_{\psi(f_x)}$ as Artin-Schreier sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$.

(2) If $\mathcal{L}_x = 0$, then the conductor bound is trivial, and otherwise we obtain from (1) the bound

$$\mathbf{c}(\mathcal{L}_x) \leqslant 1 + \deg_Y f(x,Y) + \sum_{y \in \mathbf{P}^1(\bar{\mathbf{F}}_q)} \mathrm{ord}_y(f(x,Y)) \leqslant 1 + 2\deg_Y f(x,Y) \leqslant 2\,\mathbf{c}(f),$$

as claimed. ☐

Remark 4.6. Note that $\mathcal{L}_x$ is not always isomorphic to the Artin-Schreier sheaf $\mathcal{L}_{\psi(f_x))}$ on $\mathbf{A}^1$: for instance, if $f = X/Y$ and $x = 0$, we have $\mathcal{L}_{\psi(f(x,Y))} = \bar{\mathbf{Q}}_\ell$, but $\mathcal{L}_0 = j_! \bar{\mathbf{Q}}_\ell$, where $j : \mathbf{A}^1 - \{0\} \hookrightarrow \mathbf{A}^1$ is the open immersion. Thus $\mathcal{L}_x$ has zero stalk at 0. However, this subtlety will not be a problem for us, in particular because the set of x for which this behavior happens (and the set of y such that the stalk of $\mathcal{L}_x$ at y is not the same as that of $\mathcal{L}_{\psi(f(x,Y))}$) is finite and – since these points must be common zeros of the numerator f_1 and the denominator f_2 of f – has size bounded by $\deg(f_1)\deg(f_2)$, e.g. by Bezout's theorem.

4.3. **Basic estimates.** Another frequently-used fact, which is implicit in our previous work in the case of middle-extension sheaves, is the control of Betti numbers of constructible sheaves on $\mathbf{A}^1$ in terms of the conductor:

Lemma 4.7. *Let $\mathbf{F}_q$ be a finite field of characteristic p, let $\ell \neq p$ be a prime number and $\mathcal{F}$ an ℓ-adic constructible sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$. For $i = 0, 2$, we have*

$$\dim H^i_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F}) \leqslant \mathbf{c}(\mathcal{F})$$

and

$$\dim H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F}) \leqslant 2\,\mathbf{c}(\mathcal{F}) + \mathbf{c}(\mathcal{F})^2.$$

Proof. For $i = 0$, this is obvious from the definition of $\mathrm{pct}(\mathcal{F}) \leqslant \mathbf{c}(\mathcal{F})$. For $i = 2$, we use the fact that if $\mathcal{F}$ is lisse on a dense open subset $U \subset \mathbf{A}^1$, we have

$$H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F}) = H^2_c(\bar{U}, \mathcal{F}) \simeq (\mathcal{F}_{\bar{\eta}})_{\pi_1(\bar{U},\bar{\eta})},$$

the coinvariant space for the action of the geometric fundamental group on the geometric generic fiber (see, e.g., [2, Sommes Trig., Rem. 1.18 (d)]; the first equality is also a consequence of excision) and hence

$$\dim H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F}) \leqslant \mathrm{rank}(\mathcal{F}) \leqslant \mathrm{rank}(\mathcal{F}) \leqslant \mathbf{c}(\mathcal{F}).$$

For $i = 1$, we use the Euler-Poincaré formula (see [10, 8.5.2, 8.5.3]) to get

$$\begin{aligned}\dim H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F}) = -\,\mathrm{rank}(\mathcal{F}) + \dim H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F}) + \dim H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F})\\ + \sum_x (\mathrm{drop}_x(\mathcal{F}) + \mathrm{Swan}_x(\mathcal{F})) + \mathrm{Swan}_\infty(\mathcal{F})\end{aligned} \tag{4.6}$$

where the sum is over $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$, and all but finitely many terms are zero, and the result follows from the definition of the conductor since $\mathrm{drop}_x(\mathcal{F}) = \mathrm{rank}(\mathcal{F}) - \dim \mathcal{F}_x \leqslant \mathrm{rank}(\mathcal{F})$. ∎

The following was also proved for middle-extensions in our previous works.

Lemma 4.8. *Let $\mathbf{F}_q$ be a finite field of characteristic p, let $\ell \neq p$ be a prime number and $\mathcal{F}_1$ and $\mathcal{F}_2$ be ℓ-adic constructible sheaves on $\mathbf{A}^1_{\mathbf{F}_q}$. We have*

$$\mathbf{c}(\mathcal{F}_1 \otimes \mathcal{F}_2) \leqslant 8\,\mathbf{c}(\mathcal{F}_1)^2\,\mathbf{c}(\mathcal{F}_2)^2$$

Proof. One checks easily (as in [4, Prop. 8.2 (2)]) that for the middle-extension part $(\mathcal{F}_1 \otimes \mathcal{F}_2)_0$ we have

$$\mathbf{c}((\mathcal{F}_1 \otimes \mathcal{F}_2)_0) \leqslant 6\,\mathbf{c}(\mathcal{F}_1)^2\,\mathbf{c}(\mathcal{F}_2)^2,$$

and as for the punctual part, we have

$$\mathrm{pct}(\mathcal{F}_1 \otimes \mathcal{F}_2) \leqslant (n_1 + n_2) m_1 m_2$$

where n_i is the number of points where there are punctual sections of $\mathcal{F}_i$, while m_i is the maximal dimension of the space of sections supported at a single point. Since

$$(n_1 + n_2) m_1 m_2 \leqslant (\mathbf{c}(\mathcal{F}_1) + \mathbf{c}(\mathcal{F}_2)) m_1 m_2 \leqslant 2\,\mathbf{c}(\mathcal{F}_1)\,\mathbf{c}(\mathcal{F}_2) m_1 m_2 \leqslant 2\,\mathbf{c}(\mathcal{F}_1)^2\,\mathbf{c}(\mathcal{F}_2)^2,$$

we get the result. ∎

In particular, we get the following corollary from the previous three lemmas. The statement uses the notation (4.5).

Corollary 4.9. *Let $\mathbf{F}_q$ be a finite field of order q and characteristic p, ℓ a prime distinct from p. Let $\mathcal{L} = \mathcal{L}_{\psi(f(X,Y))}$ be an ℓ-adic Artin-Schreier sheaf on $\mathbf{A}^2$ over $\mathbf{F}_q$, where $f \in \mathbf{F}_q(X,Y)$ is a non-constant rational function. Let $\mathcal{F}$ be a middle-extension ℓ-adic sheaf on $\mathbf{A}^1$ over $\mathbf{F}_q$.*

For every $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$, we have

$$\dim H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x) \leqslant 3 \cdot 2^{10}\, \mathbf{c}(f)^4\, \mathbf{c}(\mathcal{F})^4.$$

Proof. Combine Lemma 4.7, Lemma 4.8 and Lemma 4.5. □

4.4. **Number of singularities.** We will also use a criterion to bound the number of singularities in terms of estimates for the punctual part.

Lemma 4.10. *Let $\mathbf{F}_q$ be a finite field of characteristic p, $\ell \neq p$ a prime number and $\mathcal{F}$ an ℓ-adic constructible sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$. Let $U \subset \mathbf{A}^1$ be a dense open set such that the dimension of the stalks $\mathcal{F}_x$ is constant, equal to some integer $d \geqslant 0$, for all $x \in U(\bar{\mathbf{F}}_q)$. We then have*

$$n(\mathcal{F}) \leqslant |(\mathbf{P}^1 - U)(\bar{\mathbf{F}}_q)| + \mathrm{pct}(\mathcal{F}).$$

Proof. Since U contains the generic point η of $\mathbf{A}^1$, we have

$$\mathrm{rank}(\mathcal{F}) = \dim \mathcal{F}_{\bar{\eta}} = d.$$

Let $U_1 \subset U$ be the open dense subset where $\mathcal{F}$ is lisse, and let $x \in (U - U_1)(\bar{\mathbf{F}}_q)$, i.e., a point of U where $\mathcal{F}$ is not lisse. Let $\varphi : \mathcal{F}_x \longrightarrow \mathcal{F}_{\bar{\eta}}^{I_x}$ be the canonical map. The image has dimension $< d$ (since otherwise, for dimension reasons, I_x would act trivially on the geometric generic fiber $\mathcal{F}_{\bar{\eta}}$, and $\mathcal{F}$ would be lisse at x), and since $\dim \mathcal{F}_x = d$, it follows that

$$\dim \ker \varphi \geqslant 1,$$

which means that x is in the support of the punctual part of $\mathcal{F}$. Thus the number of such x is at most the size of this support, which is bounded by $\mathrm{pct}(\mathcal{F})$. Adding the points of $(\mathbf{P}^1 - U)(\bar{\mathbf{F}}_q)$ leads to the result. □

4.5. **Global conductor bound.** We now come to the lemma which contains the first idea in the proof of Theorem 2.3: it allows us to replace the sum of Swan conductors, in the definition of the conductor of a sheaf, by a global invariant that is more accessible to algebraic manipulations.

Lemma 4.11 (Global conductor bound)**.** *Let $\mathcal{F}$ be a middle-extension sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$ which is lisse on some dense open set U. We have*

$$\mathbf{c}(\mathcal{F}) \leqslant 3\,\mathrm{rank}(\mathcal{F}) + n(\mathcal{F}) - \chi(\mathcal{F}) \leqslant 3\,\mathrm{rank}(\mathcal{F}) + n(\mathcal{F}) + \sigma(\mathcal{F}),$$

where

$$\chi(\mathcal{F}) = \chi_c(\bar{U}, \mathcal{F}) = \dim H^0_c(\bar{U}, \mathcal{F}) \;- \dim H^1_c(\bar{U}, \mathcal{F}) + \dim H^2_c(\bar{U}, \mathcal{F})$$

is the compactly supported Euler-Poincaré characteristic of $\mathcal{F}$ over $\bar{U}$ and

$$\sigma(\mathcal{F}) = \dim H^1_c(\bar{U}, \mathcal{F}).$$

Proof. By the Euler-Poincaré formula (see, e.g., [10, 2.3.1]), we have

$$-\chi_c(\bar{U}, \mathcal{F}) = -\mathrm{rank}(\mathcal{F})\chi_c(\bar{U}, \bar{\mathbf{Q}}_\ell) + \sum_x \mathrm{Swan}_x(\mathcal{F}),$$

where the sum is over the points in $(\mathbf{P}^1 - U)(\bar{\mathbf{F}}_q)$, and hence

$$\sum_x \mathrm{Swan}_x(\mathcal{F}) = -\chi_c(\bar{U}, \mathcal{F}) + (2 - 2|(\mathbf{P}^1 - U)(\bar{\mathbf{F}}_q)|)\,\mathrm{rank}(\mathcal{F}) \leqslant 2\,\mathrm{rank}(\mathcal{F}) - \chi(\mathcal{F})$$

so that the first result follows from the definition of the conductor (we do not insist here that U be the maximal open set where $\mathcal{F}$ is lisse, but use the fact that the sum of Swan conductors is independent of the choice of a dense open set where $\mathcal{F}$ is lisse).

The second estimate is obtained by noting that

$$-\chi_c(\bar{U}, \mathcal{F}) = \dim H^1_c(\bar{U}, \mathcal{F}) - \dim H^2_c(\bar{U}, \mathcal{F}) \leqslant \dim H^1_c(\bar{U}, \mathcal{F}) = \sigma(\mathcal{F}).$$

□

5. Examples and applications

5.1. **Preliminaries.** The simplest applications of our results arise by using the trace functions of transform sheaves $T_{\mathcal{K}}(\mathcal{F})$ in any result involving these functions. One must be slightly careful since many results are stated for middle-extension sheaves which are pointwise pure of weight 0, and $T_{\mathcal{K}}(\mathcal{F})$ may not have these properties (and neither is it usually irreducible even if $\mathcal{F}$ is.)

There is a potential notational subtlety (which did not arise in our previous works) involving the definition of weights. For an integer $n \in \mathbf{Z}$, recall (see [3, Def. 1.2.2]) that an ℓ-adic sheaf $\mathcal{F}$ on $X_{\mathbf{F}_q}$ is *pointwise pure of weight* n if, for all finite extensions $k/\mathbf{F}_q$ and for all $x \in X(k)$, the eigenvalues of Frobenius acting on $\mathcal{F}_{\bar{x}}$ are $|k|$-Weil numbers of some weight $w = n$. A sheaf $\mathcal{F}$ on X is *mixed of weights* $\leqslant n$ if it has a finite filtration

$$0 = \mathcal{F}_0 \subset \mathcal{F}_1 \subset \cdots \subset \mathcal{F}_m = \mathcal{F}$$

where the successive quotients $\mathcal{F}_i/\mathcal{F}_{i-1}$ are pointwise pure with weight $n_i \leqslant n$.

On the other hand (see [11, (7.3.7)]), a middle-extension sheaf $\mathcal{F}$ on a curve $U_{\mathbf{F}_q}$ is *pure of weight* n if, given a dense open set $V \subset U$ where $\mathcal{F}$ is lisse, for all $k/\mathbf{F}_q$ and all $x \in V(k)$, the eigenvalues of Frobenius on $\mathcal{F}_{\bar{x}}$ are $|k|$-Weil numbers of weight n. It follows from results of Deligne (in particular [3, Lemme 1.8.1], and the Riemann Hypothesis) that such a sheaf is also mixed of weights $\leqslant n$, i.e., the eigenvalues of Frobenius at the "missing points" $U - V$ are also Weil numbers with weight $\leqslant n$. However, these weights may be $< n$. In other words, a middle-extension sheaf may be pure of weight n without being pointwise pure of weight n.

The following lemma encapsulates a reduction of trace functions of constructible sheaves to middle-extension sheaves:

Lemma 5.1 (Trace function of constructible sheaf)**.** *Let $\mathbf{F}_q$ be a finite field of characteristic p, let $\ell \neq p$ be a prime and let $\mathcal{F}$ be an ℓ-adic constructible sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$ which is mixed of weights $\leqslant 0$.*

There exists a decomposition of the trace function $t_{\mathcal{F}}$ of $\mathcal{F}$ of the form

$$t_{\mathcal{F}} = t_{\mathcal{F}^{mid}} + t_1 + t_2,$$

where $\mathcal{F}^{mid}$ is a middle-extension sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$ which is pure of weight 0, and where:

(1) *The function t_1 is zero except for a set of values of $x \in \mathbf{F}_q$ of size at most $2\,\mathbf{c}(\mathcal{F})$, and it satisfies*

$$|t_1(x)| \leqslant 2\,\mathbf{c}(\mathcal{F})$$

for all $x \in \mathbf{F}_q$.

(2) *The function* t_2 *satisfies*
$$|t_2(x)| \leqslant \mathbf{c}(\mathcal{F})q^{-1/2}$$
for all $x \in \mathbf{F}_q$.

Proof. This is a classical "dévissage". We begin by observing that
$$|t_{\mathcal{F}}(x)| \leqslant \mathbf{c}(\mathcal{F})$$
for all $x \in \mathbf{F}_q$: indeed, by assumption, all eigenvalues of Frobenius on the stalk $\mathcal{F}_{\bar{x}}$ are of modulus at most 1, and the maximal dimension of a stalk is bounded by the conductor (including where there is a punctual part of the sheaf.)

Let $\mathcal{F}^0$ be the direct sum of quotients which are pointwise pure of weight 0 in a filtration of $\mathcal{F}$ with successive quotients which are pointwise pure of some weight $\leqslant 0$, and let $\mathcal{F}^1$ be the direct sum of the remaining quotients. We have
$$t_{\mathcal{F}}(x) = t_{\mathcal{F}^0}(x) + t_{\mathcal{F}^1}(x),$$
and trivially
$$|t_{\mathcal{F}^1}(x)| \leqslant p^{-1/2}\,\mathbf{c}(\mathcal{F})$$
for all $x \in \mathbf{F}_q$. We put $t_2 = t_{\mathcal{F}^1}$.

Next, let
$$0 \longrightarrow \mathcal{F}^{pct} \to \mathcal{F}^0 \to \mathcal{F}^{npct} \to 0$$
be the short exact sequence associated to the inclusion of the punctual part $\mathcal{F}^{pct}$ of $\mathcal{F}^0$. We have
$$t_{\mathcal{F}^0}(x) = t_{\mathcal{F}^{pct}}(x) + t_{\mathcal{F}^{npct}}(x),$$
and $t_{\mathcal{F}^{pct}}$ is zero except for $\leqslant \mathbf{c}(\mathcal{F})$ values of x for which we have
$$|t_{\mathcal{F}^{pct}}(x)| \leqslant \dim H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_p, \mathcal{F}^0) \leqslant \mathbf{c}(\mathcal{F}).$$

Finally, let $j : U \hookrightarrow \mathbf{A}^1$ be the open immersion of the maximal dense open subset where $\mathcal{F}^{npct}$ is lisse, and let
$$\mathcal{F}^{mid} = j_*j^*\mathcal{F}^{npct}.$$

This is a middle-extension sheaf, pointwise pure of weight 0, with trace function equal to that of $\mathcal{F}$ for $x \in U(\mathbf{F}_q)$. Thus the difference
$$t_{\mathcal{F}} - t_{\mathcal{F}^{mid}}$$
is zero except for at most $\mathbf{c}(\mathcal{F})$ values of $x \in \mathbf{F}_q$, and has modulus $\leqslant 2\,\mathbf{c}(\mathcal{F})$ for all x. We obtain the desired decomposition by taking
$$t_1 = t_{\mathcal{F}^{pct}} + t_{\mathcal{F}} - t_{\mathcal{F}^{mid}}.$$
$\square$

We can apply this to the trace functions of the transform sheaves $T^1_{\mathcal{K}}(\mathcal{F})$ considered in this paper. We introduce a definition for convenience.

Definition 5.2 (f-disjoint sheaf)**.** Let $\mathbf{F}_q$ be a finite field of characteristic p, let $\ell \neq p$ be a prime. Let $f \in \mathbf{F}_q(X, Y)$ be a rational function and let $\mathcal{K} = \mathcal{L}_{\psi(f)}$ be the Artin-Schreier sheaf on $\mathbf{A}^2_{\mathbf{F}_q}$ associated to f.

A middle-extension sheaf $\mathcal{F}$ on $\mathbf{A}^1_{\mathbf{F}_q}$ is called *f-disjoint* or *$\mathcal{K}$-disjoint* if
$$H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x) = 0$$

for all $x \in \bar{\mathbf{F}}_q$.

Corollary 5.3 (Artin-Schreier transforms as trace functions)**.** *Let $\mathbf{F}_q$ be a finite field of characteristic p, let $\ell \neq p$ be a prime. Let $f \in \mathbf{F}_q(X, Y)$ be a rational function given by $f = f_1/f_2$ with $f_i \in F_q[X, Y]$ coprime polynomials, and assume that $\mathbf{c}(f) < p$.*

Let $\mathcal{F}$ be a middle-extension sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$ which is pointwise pure of weight 0 and f-disjoint.

There exists an absolute constant $A \geqslant 1$, independent of f and $\mathcal{F}$, such that for all $x \in \mathbf{F}_q$, we have

$$\frac{1}{\sqrt{q}} \sum_{\substack{y \in \mathbf{F}_q \\ f_2(x,y) \neq 0}} t_{\mathcal{F}}(y) \psi(f(x, y)) = -t_0 + t_1 + t_2,$$

where t_0 is the trace function of a middle-extension sheaf $\mathcal{G}_0$ of weight 0 on $\mathbf{A}^1_{\mathbf{F}_q}$ with

$$\mathbf{c}(\mathcal{G}_0) \leqslant (2\,\mathbf{c}(f)\,\mathbf{c}(\mathcal{F}))^A.$$

the function t_1 is zero for a set of values of $x \in \mathbf{F}_q$ of size at most $(2\,\mathbf{c}(f)\,\mathbf{c}(\mathcal{F}))^A$, and it satisfies

$$|t_1(x)| \leqslant (2\,\mathbf{c}(f)\,\mathbf{c}(\mathcal{F}))^A,$$

for all $x \in \mathbf{F}_q$, while the function t_2 satisfies

$$|t_2(x)| \leqslant (2\,\mathbf{c}(f)\,\mathbf{c}(\mathcal{F}))^A q^{-1/2}$$

for all $x \in \mathbf{F}_q$.

Proof. Let

$$\mathcal{G} = T^1_{\mathcal{K}}(\mathcal{F})(1/2), \qquad \mathcal{G}_i = T^i_{\mathcal{K}}(\mathcal{F})(1/2)$$

for $0 \leqslant i \leqslant 2$.

By the Riemann Hypothesis [3] (taking into account the Tate twist) the sheaf $\mathcal{G}$ is mixed of weight $\leqslant 0$, while $\mathcal{G}_i$ is mixed of weight $\leqslant i - 1$. The trace function of $\mathcal{G}$, by the proper base change theorem (see Proposition 4.1, (4)) and the Grothendieck-Lefschetz trace formula, is

$$t_{\mathcal{G}}(x) = -\frac{1}{\sqrt{q}} \sum_{y \in \mathbf{F}_q} t_{\mathcal{F}}(y) t_{\mathcal{K}}(x, y) + t_{\mathcal{G}_0}(x) + t_{\mathcal{G}_2}(x)$$

for $x \in \mathbf{F}_q$.

The stalk of $\mathcal{G}_0$ over x is

$$H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x)$$

by Lemma 4.2) and that of $\mathcal{G}_2$ is

$$H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x) = 0$$

since $\mathcal{F}$ is f-disjoint. Hence we obtain

$$t_{\mathcal{G}}(x) = -\frac{1}{\sqrt{q}} \sum_{y \in \mathbf{F}_q} t_{\mathcal{F}}(y) t_{\mathcal{K}}(x, y)$$

for all $x \in \mathbf{F}_q$. By definition (Definition 2.2) we have

$$t_{\mathcal{K}}(x, y) = \begin{cases} \psi(f(x, y)) & \text{if } f_2(x, y) \neq 0 \\ 0 & \text{otherwise,} \end{cases}$$

and by Theorem 2.3, there exists $A \geqslant 1$ such that the constructible sheaf $\mathcal{G}$ satisfies

$$\mathbf{c}(\mathcal{G}) \leqslant (2\,\mathbf{c}(f)\,\mathbf{c}(\mathcal{F}))^A.$$

Thus the result follows by applying Lemma 5.1 to $\mathcal{G}$. $\square$

Remark 5.4. For $\mathcal{K}$ as in this corollary, the condition that $\mathcal{F}$ is f-disjoint is valid in many cases. We list some of them for convenience. The assumption of Corollary 5.3 holds:

(1) If $\mathcal{F}$ is irreducible of rank at least 2 (e.g., Kloosterman sheaves in one or more variables), or more generally if $\mathcal{F}$ is irreducible and no isomorphic to an Artin-Schreier sheaf;

(2) If $\mathcal{F}$ is tamely ramified and there is no specialization f_x of f which is constant as an element in $\mathbf{F}_q(Y)$ (e.g., any Kummer sheaf with trace function $\chi(g(x))$ for a multiplicative character χ, provided no f_x is constant); in particular, if $\mathcal{F}$ is the trivial sheaf with constant trace function 1, it is enough that no specialization f_x be constant.

(3) If $\mathcal{F}$ is an Artin-Schreier sheaf $\mathcal{L}_{\psi(g)}$ with trace function $\psi(g(x))$ and there is no $x \in \bar{\mathbf{F}}_q$ such that $g + f_x$ is constant.

5.2. **Automorphic twists.** We begin by explaining one setting where the application of our result is very easy:

Proposition 5.5. *Let f be a Hecke cusp form of level $N \geqslant 1$ with Fourier coefficients $\varrho_f(n)$ at ∞. Let g_1, $g_2 \in \mathbf{Z}[X, Y]$ be two non-constant coprime polynomials, and let $g = g_1/g_2 \in \mathbf{Q}(X, Y)$.*

Let V be a smooth function on $]0, +\infty[$ with compact support. Let p be a prime number, let K be an irreducible trace function modulo p associated to a middle-extension sheaf $\mathcal{F}$ which is $(g \,(\mathrm{mod}\, p))$-disjoint. For $\varepsilon > 0$, we have

$$\sum_{n\geqslant 1} \varrho_f(n) \frac{1}{\sqrt{p}} \Bigl(\sum_{\substack{x\in\mathbf{F}_p \\ g_2(n,x)\neq 0 \,(\mathrm{mod}\, p)}} K(x) e\Bigl(\frac{g_1(n,x)\overline{g_2(n,x)}}{p}\Bigr)\Bigr) V(n/p) \ll p^{1-1/8+\varepsilon}$$

where the implied constant depends on $(f, V, \varepsilon, \mathbf{c}(K), \mathbf{c}(g))$.

Proof. The main result of [4] shows that

$$\sum_{n\geqslant 1} \varrho_f(n) K(n) V(n/p) \ll p^{1-1/8+\varepsilon}$$

if K is the trace function of a geometrically isotypic middle-extension sheaf which is pointwise pure of weight 0. We will show how to deduce the result from this.

By Corollary 5.3 (applied with ψ chosen so that $t_{\mathcal{L}_\psi}(x) = e(x/p)$ for $x \in \mathbf{F}_p$), we have a decomposition

$$\frac{1}{\sqrt{p}} \Bigl(\sum_{\substack{x\in\mathbf{F}_p \\ g_2(n,x)\neq 0}} K(x) e\Bigl(\frac{g_1(n,x)\overline{g_2(n,x)}}{p}\Bigr) = -t_0 + t_1 + t_2$$

where $-t_0$ is the trace function of a middle-extension sheaf which is pure of weight 0 and has conductor $\leqslant C = (2\,\mathbf{c}(f)\,\mathbf{c}(K))^A$, while t_1 is zero except for $\leqslant C$ values of $x \in \mathbf{F}_p$, where it has modulus at most C, while $|t_2| \leqslant Cp^{-1/2}$. We have then

$$\sum_{n\geqslant 1} \varrho_f(n) t_i(n)(n) V(n/p) \ll p^{1-1/8+\varepsilon}$$

for $i = 1$, 2, and we are reduced to the case of t_0. Decomposing t_0 in trace functions of its geometrically isotypic components, we conclude by applying [4]. $\square$

5.3. **Two-variable sums and the example of Conrey-Iwaniec.** A basic application of bounds on conductors like those of Theorem 2.3 concerns two-variable exponential sums of quite general type. We present the very general principle before giving a concrete example.

Given a trace function $K(x, y)$ in two variables, e.g. $K(x, y) = \chi(f_1(x, y))e(f_2(x, y)/p)$ for rational functions f_1 and $f_2 \in \mathbf{F}_p(X, Y)$ and for a multiplicative character χ modulo p, one wishes to obtain square-root cancellation (when possible) for

$$\sum_{x,y} K(x, y).$$

This may be written as

$$\sum_x \sum_y K(x, y),$$

i.e., as the inner product of the constant function 1 (i.e., the trace function of the trivial sheaf) and (essentially) the trace function of

$$T^1_{\mathcal{K}}(\bar{\mathbf{Q}}_\ell)$$

where $\mathcal{K}$ is the sheaf with trace function K. It may happen that K is given naturally as a product

$$K(x, y) = K_1(x)K_2(y)K_3(x, y)$$

for trace functions K_1 and K_2 modulo p and another trace function K_3 in two variables; in such a case, it may be better to write the sum as

$$\sum_x K_1(x) \sum_y K_2(y)K_3(x, y),$$

which is the inner-product of K_1 with the trace function of $T^1_{\mathcal{K}_3}(\mathcal{K}_2)$, with obvious notation.

From a direct application of the Riemann Hypothesis, we obtain the following qualitative information concerning these types of sums:

Proposition 5.6 (Small diagonal principle)**.** *Let $\mathbf{F}_q$ be a finite field of characteristic p, let $\ell \neq p$ be a prime number. Let $\mathcal{K}$ be a constructible ℓ-adic sheaf mixed of weight $\leqslant 0$ on $\mathbf{A}^2_{\mathbf{F}_q}$.*

Let $\mathcal{F}_2$ be a middle-extension sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$, pointwise pure of weight 0 such that $T^2_{\mathcal{K}}(\mathcal{F}_2)$ is generically 0.

There exists a finite set $X(\mathcal{K}, \mathcal{F}_2)$ of geometrically irreducible middle-extension sheaves which are pointwise pure of weight 0, of cardinality bounded in terms of the conductor of $T^1_{\mathcal{K}}(\mathcal{F}_2)$, such that if $\mathcal{F}_1$ is a middle-extension sheaf of weight 0, geometrically irreducible, and not geometrically isomorphic to any of the sheaves in $X(\mathcal{K}, \mathcal{F}_2)$, then

$$\sum_{x,y \in \mathbf{F}_q} t_{\mathcal{F}_1}(x)t_{\mathcal{F}_2}(y)t_{\mathcal{K}}(x, y) \ll q,$$

where the implied constant depends only on the conductor of $\mathcal{F}_1$ and of $T^1_{\mathcal{K}}(\mathcal{F}_2)$.

Proof. Let $X(\mathcal{K}, \mathcal{F}_2)$ be the set of geometric isomorphism classes of geometrically irreducible components of the weight 1 part of $T^1_{\mathcal{K}}(\mathcal{F}_2)$. This is finite with cardinality bounded by the rank of $T^1_{\mathcal{K}}(\mathcal{F}_2)$, hence bounded in terms of the conductor of $\mathcal{F}_2$.

Under the assumptions of the proposition, for $\mathcal{F}_1$ geometrically irreducible and not in $X(\mathcal{K},\mathcal{F}_2)$, we have

$$\sum_{x,y\in\mathbf{F}_q} t_{\mathcal{F}_1}(x)t_{\mathcal{F}_2}(y)t_{\mathcal{K}}(x,y) = -\operatorname{tr}(\mathrm{Fr}\mid H^1_c(\mathbf{A}^1\times\bar{\mathbf{F}}_q,\mathcal{F}_1\otimes T^1_{\mathcal{K}}(\mathcal{F}_2)))$$

and the cohomology space is mixed of weights $\leqslant 2$ (since the H^2_c vanishes by definition.) Using conductor bounds (Lemmas 4.7 and 4.8), we obtain the result. ☐

Although this proposition does not, by itself, give square-root cancellation in any individual case, it implies for instance that

$$\sum_x e\Bigl(\frac{ax^2}{p}\Bigr)\sum_y t_{\mathcal{F}_2}(y)t_{\mathcal{K}}(x,y) \ll p$$

(working over $\mathbf{F}_p$) for all $a\in\mathbf{F}_p$ except for a number of exceptions bounded in terms of the conductors of $\mathcal{F}_2$ and $\mathcal{K}$ only. In quite a few applications, this type of qualitative "control of the diagonal" is sufficient (for instance, similar ideas are crucial in [4].) However, this is not always the case, and one needs to attempt some further analysis if a more precise result is needed.

We present a concrete example, taken from the important work of Conrey and Iwaniec on the third moment of special values of automorphic L-functions [1]. Given a prime p and two multiplicative characters χ_1 and χ_2 modulo p, Conrey and Iwaniec consider the sum

$$S(\chi_1,\chi_2)=\sum_{x,y\in\mathbf{F}_p}\chi_1(xy(x+1)(y+1))\chi_2(xy-1),$$

and prove:

Theorem 5.7 (Conrey-Iwaniec)**.** *Let χ_1 be a non-trivial multiplicative character modulo p, and let χ_2 be any multiplicative character modulo p. Then*

$$S(\chi_1,\chi_2)\ll p$$

where the implied constant is absolute.

This is [1, Lemma 13.1], slightly generalized, since we do not assume that χ_1 is a real character. Conrey and Iwaniec remark [1, Remarks, p. 1208] that their main result concerning L-functions would be considerably weakened if (for χ_1 a real character modulo p, for many primes p) there existed a single character χ_2 for which the size of the sum would be $p^{3/2}$.

We now explain how to prove Theorem 5.7 using the ideas of cohomological transforms. The sums $S(\chi_1,\chi_2)$ are naturally presented in the form discussed above, namely as the inner product of the trace function of the dual of the Kummer sheaf

$$\mathcal{F}_1=\mathcal{L}_{\chi_1(X(X+1))}$$

with that of the transform sheaf

$$\mathcal{G}=T^1_{\mathcal{K}}(\mathcal{L}_{\chi_1(Y(Y+1))}),\qquad \mathcal{K}=\mathcal{L}_{\chi_2(XY-1)}$$

(the latter is defined as the extension by 0 of the Kummer sheaf $\mathcal{L}_{\chi_2(XY-1)}$ on the open set complement of the curve $XY-1$, see below for the general definition.)

More precisely, the trace function of $\mathcal{G}$ is

$$t_{\mathcal{G}}(x) = -\sum_{y\in\mathbf{F}_p} \chi_1(y(y+1))\chi_2(xy-1)$$

for all $x \in \mathbf{F}_p$, provided $\chi_1 \neq 1$: indeed, by the trace formula and the proper base change theorem, it is enough to show that $T^0_{\mathcal{K}}(\mathcal{L}_{\chi_1(Y(Y+1))}) = T^2_{\mathcal{K}}(\mathcal{L}_{\chi_1(Y(Y+1))}) = 0$ in that case. The former is true by Lemma 4.2, and the latter because the fiber above $x \in \bar{\mathbf{F}}_p$ is

$$H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_p, \mathcal{L}_{\chi_1(Y(Y+1))} \otimes \mathcal{L}_{\chi_2(xY-1)}) = 0$$

(since $\chi_1 \neq 1$, this can only be non-zero if the second tensor factor is ramified at 0 and -1, but it is in fact always unramified at 0.)

The kernel $\mathcal{K}$ is not of the type considered in Theorem 2.3. However, it is easy to adapt the proof of this result to derive an analogue for Kummer sheaves. These we define in general in analogy with Definition 2.2: for a multiplicative ℓ-adic character of $\mathbf{F}_q^\times$, if $f \in \mathbf{F}_q(X,Y)$ is a rational function, U the open set where the numerator and denominator are both non-zero, with $j : U \hookrightarrow \mathbf{A}^2$ the open immersion, and if $f_U : U \longrightarrow \mathbf{G}_m$ is the associated morphism, then we define

$$\mathcal{L}_{\chi(f)} = j_! f_U^* \mathcal{L}_\chi.$$

Theorem 5.8 (Conductor of Kummer transforms)**.** *Let $\mathbf{F}_q$ be a finite field of order q and characteristic p, ℓ a prime distinct from p. Let $\mathcal{K}$ be an ℓ-adic sheaf on $\mathbf{A}^1 \times \mathbf{A}^1$ over $\mathbf{F}_q$ of the form $\mathcal{K} = \mathcal{L}_{\chi(f)}$.*

For constructible sheaves $\mathcal{F}$ on $\mathbf{A}^1_{\mathbf{F}_q}$, and $0 \leqslant i \leqslant 2$, let

$$T^i_{\mathcal{K}}(\mathcal{F}) = R^i p_{1,!}(p_2^*\mathcal{F} \otimes \mathcal{K}).$$

There exists an absolute constant $A \geqslant 1$ such that

$$\mathbf{c}(T^i_{\mathcal{K}}(\mathcal{F})) \leqslant (2\,\mathbf{c}(\mathcal{K})\,\mathbf{c}(\mathcal{F}))^A$$

and moreover

$$\dim H^i_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, p_2^*\mathcal{F} \otimes \mathcal{K}) \leqslant (2\,\mathbf{c}(f)\,\mathbf{c}(\mathcal{F}))^A.$$

Sketch of proof. One can follow line by line the proof of Theorems 2.3 and 2.5 as we will explain it below. The only differences are:

(1) to bound the Betti numbers

$$\dim H^i_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{K})$$

(i.e., when the input sheaf is trivial), one uses the results of Adolphson-Sperber or Katz [14, Th. 12] instead of those of Bombieri (which are only proved for additive characters), although an alternative is to lift the tame sheaves to characteristic 0.

(2) the specialized Kummer sheaves $\mathcal{L}_{\chi(f_x)}$ are ramified, with drop 1, at zeros and poles of f_x, and not only at poles; on the other hand, the Swan conductors always vanish since these sheaves are tamely ramified. Thus one must use a slightly different version of Lemma 12.1, (1) and (3), and use it to prove the analogue of Lemma 12.3 (1) for Kummer kernels. ☐

In particular, in our case, $\mathcal{G} = T^1_{\mathcal{K}}(\mathcal{L}_{\chi_1(Y(Y+1))})$ has conductor absolutely bounded as χ_1, χ_2 and p vary. By the Riemann Hypothesis, the sheaf $\mathcal{G}$ is also mixed of weights $\leqslant 1$, and therefore the principle above shows that, for all primes p, and for all characters χ_2, we have

$$S(\chi_1, \chi_2) \ll p$$

with an absolute implied constant, for all but a bounded number of multiplicative characters χ_1 modulo p (since $\mathcal{L}_{\chi_1(X(X+1))} \simeq \mathcal{L}_{\chi_1'(X(X+1))}$ if and only if $\chi_1 = \chi_1'$.)

In order to go deeper and show that, in fact, these exceptions do not exist, we must look a bit more carefully at $\mathcal{G}$.

Proposition 5.9. (1) *For all non-trivial multiplicative character χ_2, the constructible sheaf $T^1_{\mathcal{K}}(\mathcal{L}_{\chi_2(Y(Y+1))})$ has generic rank* 2.

(2) *Let $\mathcal{F}$ be a geometrically irreducible middle-extension sheaf on $\mathbf{A}^1_{\mathbf{F}_p}$, pointwise pure of weight* 0, *and assume $\mathcal{F}$ is not geometrically isomorphic to $\mathcal{L}_{\chi_1}$. Then the part of weight* 1 *of the sheaf $T^1_{\mathcal{K}}(\mathcal{F})$ is geometrically irreducible, in the sense that the associated middle-extension sheaf is geometrically irreducible.*

If we grant this proposition, then the weight one part of our particular sheaf

$$\mathcal{G} = T^1_{\mathcal{K}}(\mathcal{L}_{\chi_1(Y(Y+1))})$$

is geometrically irreducible (since $\mathcal{L}_{\chi_2(Y(Y+1))}$, which is ramified at -1, is certainly not geometrically isomorphic to $\mathcal{L}_{\chi_1}$). Since $\mathcal{G}$ is mixed of weights $\leqslant 1$, and of rank 2, and hence is not geometrically isomorphic to the geometrically irreducible Kummer sheaf $\mathcal{L}_{\chi_1(x(x+1))}$, we obtain

$$S(\chi_1, \chi_2) = \sum_x \chi_1(x(x+1)) t_{\mathcal{G}}(x) \ll p$$

where the implied constant is absolute. This is the result of Conrey and Iwaniec and finishes our proof of Theorem 5.7.

For the proof of (2), we recall a very useful diophantine criterion for irreducibility of Katz (see [12, Lemma 7.0.3]).

Lemma 5.10 (Irreducibility criterion)**.** *Let $\mathbf{F}_q$ be a finite field of characteristic p, let $\ell \neq p$ be a prime number and let $\mathcal{F}$ be an ℓ-adic constructible sheaf on $\mathbf{A}^1_{\mathbf{F}_p}$ which is mixed of weights $\leqslant 0$. Then we have*

$$\frac{1}{q^\nu} \sum_{x \in \mathbf{F}_{q^\nu}} |t_{\mathcal{F}}(x, q^\nu)|^2 = 1 + O(q^{-\nu/2}) \tag{5.1}$$

for $\nu \geqslant 1$, if and only if the middle-extension part of weight 0 *of $\mathcal{F}$ is geometrically irreducible, i.e., if and only if, for any dense open subset U where $\mathcal{F}$ is lisse, the restriction of the weight* 0 *part of $\mathcal{F}$ to $U \times \bar{\mathbf{F}}_q$ corresponds to an irreducible representation of the geometric fundamental group of U.*

Proof. For $\nu \geqslant 1$ fixed, let

$$t_{\mathcal{F}}(x, q^\nu) = t_{\mathcal{F}^{mid}}(x, q^\nu) + t_1(x) + t_2(x)$$

for $x \in \mathbf{F}_{q^\nu}$ be the decomposition of Lemma 5.1 (applied to $\mathbf{F}_{q^\nu}$). We wish to prove that $\mathcal{F}^{mid}$ is geometrically irreducible. From the properties of t_1 and t_2, we see that

$$\frac{1}{q^\nu} \sum_{x \in \mathbf{F}_{q^\nu}} |t_{\mathcal{F}}(x, q^\nu)|^2 = \frac{1}{q^\nu} \sum_{x \in \mathbf{F}_{q^\nu}} |t_{\mathcal{F}^{mid}}(x, q^\nu)|^2 + O(q^{-\nu})$$

for $\nu \geqslant 1$. Now let U be a dense open subset of $\mathbf{A}^1$ where $\mathcal{F}^{mid}$ is lisse. Then we have

$$\frac{1}{q^\nu} \sum_{x \in U(\mathbf{F}_{q^\nu})} |t_{\mathcal{F}^{mid}}(x, q^\nu)|^2 = \frac{1}{q^\nu} \sum_{x \in \mathbf{F}_{q^\nu}} |t_{\mathcal{F}^{mid}}(x, q^\nu)|^2 + O(q^{-\nu}),$$

for $\nu \geqslant 1$, since the complement is finite. Therefore, we have (5.1) if and only if

$$\frac{1}{q^\nu}\sum_{x\in U(\mathbf{F}_{q^\nu})}|t_{\mathcal{F}^{mid}}(x,q^\nu)|^2 = 1+O(q^{-\nu/2})$$

for $\nu \geqslant 1$. But by [12, Lemma 7.0.3] applied to the lisse sheaf $\mathcal{F}^{mid}$ on U, which is pure of weight 0, this last condition holds if and only if $\mathcal{F}^{mid}$ is geometrically irreducible on U. $\square$

Proof of the proposition. (1) The fiber of $\mathcal{G}$ over $x \in \mathbf{F}_q$ is

$$H^1_c(\mathbf{A}^1\times\bar{\mathbf{F}}_q, \mathcal{L}_{\chi_1(Y(Y+1))}\otimes\mathcal{L}_{\chi_2(xY-1)}).$$

By the Euler-Poincaré formula (see (4.6)), its dimension is

$$\dim H^1_c(\mathbf{A}^1\times\bar{\mathbf{F}}_q, \mathcal{L}_{\chi_1(Y(Y+1))}\otimes\mathcal{L}_{\chi_2(xY-1)}) = -1+3=2$$

if $x\neq -1$ (so that the sheaf is ramified at the three points $y=0$, -1 and $1/x$). Hence the generic rank is 2.

(2) We apply the irreducibility criterion to the Tate twist $\mathcal{G}(1)$, which is mixed of weights $\leqslant 0$.

For $\nu\geqslant 1$, denoting by $\chi_{1,\nu}$ the extension $\chi_1\circ N_{\mathbf{F}_{q^\nu}/\mathbf{F}_q}$ of χ_1 to $\mathbf{F}_{q^\nu}$, we have

$$\begin{aligned}\sum_{x\in\mathbf{F}_{q^\nu}}|t_{\mathcal{G}(1)}(x,q^\nu)|^2 &= \frac{1}{q^{2\nu}}\sum_{x\in\mathbf{F}_{q^\nu}}\Big|\sum_{y\in\mathbf{F}_{q^\nu}}t_{\mathcal{F}}(x,q^\nu)\chi_{1,\nu}(xy-1)\Big|^2\\ &= \frac{1}{q^{2\nu}}\sum_{y_1,y_2\in\mathbf{F}_{q^\nu}}t_{\mathcal{F}}(y_1,q^\nu)\overline{t_{\mathcal{F}}(y_2,q^\nu)}\sum_{x\in\mathbf{F}_{q^\nu}}\chi_{1,\nu}(xy_1-1)\overline{\chi_{1,\nu}(xy_2-1)}.\end{aligned}$$

The contribution of the diagonal terms $y_1=y_2$ to this sum is

$$\frac{q^\nu-1}{q^{2\nu}}\sum_{y\in\mathbf{F}_{q^\nu}}|t_{\mathcal{F}}(y,q^\nu)|^2=\frac{1}{q^\nu}\sum_{y\in\mathbf{F}_{q^\nu}}|t_{\mathcal{F}}(y,q^\nu)|^2+O(q^{-\nu})=1+O(q^{-\nu})$$

by the irreducibility criterion applied to $\mathcal{F}$, which is assumed to be a geometrically irreducible middle-extension sheaf which is pure of weight 0.

If $y_1\neq y_2$, the map

$$x\mapsto\frac{xy_1-1}{xy_2-1}$$

is a bijection on $\mathbf{P}^1(\mathbf{F}_{q^\nu})$. Hence, in that case, we have

$$\sum_{x\in\mathbf{F}_{q^\nu}}\chi_{1,\nu}(xy_1-1)\overline{\chi_{1,\nu}(xy_2-1)}=-\chi_{1,\nu}(y_1)\overline{\chi_{1,\nu}(y_2)}$$

(we write it in this way to incorporate the case $y_2=0$, in which case the map is a bijection of $\mathbf{F}_{q^\nu}$, while otherwise the sum over $x\in\mathbf{F}_{q^\nu}$ misses the point y_1/y_2.)

Thus we get an off-diagonal contribution equal to

$$-\frac{1}{q^{2\nu}}\sum_{\substack{y_1,y_2\in\mathbf{F}_{q^\nu}\\ y_1\neq y_2}}t_{\mathcal{F}}(y_1,q^\nu)\overline{t_{\mathcal{F}}(y_2,q^\nu)}\chi_{1,\nu}(y_1)\overline{\chi_{1,\nu}(y_2)}.$$

Inserting the diagonal in this sum, we find that it is equal to

$$-\frac{1}{q^{2\nu}}\Big(\Big|\sum_{y\in\mathbf{F}_{q^\nu}} t_{\mathcal{F}}(y,q^\nu)\chi_{1,\nu}(y)\Big|^2 - \sum_{y\in\mathbf{F}_{q^\nu}^\times} |t_{\mathcal{F}}(y,q^\nu)|^2\Big).$$

By the Riemann Hypothesis, since $\mathcal{F}$ is geometrically irreducible but not geometrically isomorphic to $\mathcal{L}_{\chi_1}$, we have

$$\Big|\sum_{y\in\mathbf{F}_{q^\nu}} t_{\mathcal{F}}(y,q^\nu)\chi_{1,\nu}(y)\Big|^2 = O(q^\nu),$$

while the bound

$$\sum_{y\in\mathbf{F}_{q^\nu}^\times} |t_{\mathcal{F}}(y,q^\nu)|^2 = O(q^\nu)$$

is immediate. Hence the off-diagonal contribution is $O(q^{-\nu})$, and the irreducibility criterion does apply. ☐

5.4. **A generalization of a theorem of Fouvry and Michel.** As another application, we show that our bounds for cohomological transforms lead to a generalization of (a slightly weakened form) of a theorem of Fouvry and Michel [7, Th. 3.1].

Let p be a prime number, $f\in\mathbf{F}_p(X)^\times$ a non-zero rational function and ψ a non-trivial additive character of $\mathbf{F}_p$. We consider the two-variable family of exponential sums given by

$$S(x,y;p) = -\frac{1}{\sqrt{p}}\sum_{t\in\mathbf{F}_p}\psi(xf(t)+yt)$$

for $(x,y)\in\mathbf{F}_p\times\mathbf{F}_p$ (where t runs implicitly over $t\in\mathbf{F}_p$ which is not a pole of f). We then consider another non-constant rational function $g\in\mathbf{F}_p(X)$, and the sums

$$R(x,y;p) = \begin{cases} S(g(x),y;p) & \text{if } x \text{ is not a pole of } g\\ 0 & \text{otherwise.}\end{cases}$$

obtained by change of variable. We finally define

$$C(y,z) = \frac{1}{\sqrt{p}}\sum_{x\in\mathbf{F}_p} R(x,y;p)\overline{R(x,z;p)}$$

for $(y,z)\in\mathbf{F}_p\times\mathbf{F}_p$.

If $f=X^k$ for some $k\in\mathbf{Z}$, these are the sums that occur in [7, Th. 3.1]. The goal is to have square-root cancellation in this sum, except in "diagonal situations" where y and z are related. Fouvry and Michel prove such a statement for $f=X^k$, where the diagonal is defined by the condition $y^k=z^k$.

Our general statement is qualitatively similar but less precise with respect to the diagonal:

Theorem 5.11 (Fouvry-Michel sums)**.** *With notation as above, and under the assumption that f is not a polynomial of degree $\leqslant 1$, there exists a constant $B\geqslant 0$ depending only on $\mathbf{c}(f)$ and $\mathbf{c}(g)$ and a subset $\Delta=\Delta_{f,g}(\mathbf{F}_p)\subset\mathbf{F}_p^2$ such that:*

(1) *For each $y\in\mathbf{F}_p$, there are $\leqslant B$ values of $z\in\mathbf{F}_p$ such that $(y,z)\in\Delta$, where B depends only on the degrees of numerator and denominator of f and g, and conversely, for each z there are at most B values of y.*

(2) *For* $(y,z)\notin\Delta$, *we have*

$$|C(y,z)|\leqslant B.$$

(3) *For* $(y,z)\in\Delta$, *we have*

$$|C(y,z)|\leqslant B\sqrt{p}.$$

We begin the proof by expressing the sums in terms of linear operators applied to trace functions. We recall that FT_ψ denotes the Fourier transform on functions modulo p.

Lemma 5.12. (1) *For any fixed* $z\in\mathbf{F}_p$, *let* $D(y)=C(y,z)$. *Then we have*

$$D=\mathrm{FT}_\psi(\mathrm{T}_K(\varphi_z))$$

where

$$K(x,y)=\begin{cases}\psi(g(y)f(x)) & \text{if } y \text{ is not a pole of } g \text{ and } x \text{ is not a pole of } f\\ 0 & \text{otherwise,}\end{cases}$$

and

$$\varphi_z(x)=\overline{R(x,z;p)}.$$

(2) *For any* $z\in\mathbf{F}_p$, *we have*

$$\varphi_z=\mathrm{T}_{\tilde K}(\psi_{-z})$$

where $\psi_a(x)=\psi(ax)$ *and* $\tilde K(x,y)=\overline{K(y,x)}$.

Proof. We have

$$\begin{aligned}D(y)&=\frac{1}{\sqrt{p}}\sum_x R(x,y;p)\overline{R(x,z;p)}\\ &=-\frac{1}{\sqrt{p}}\sum_t\psi(yt)\frac{1}{\sqrt{p}}\sum_x\varphi_z(x)\psi(g(x)f(t))=\mathrm{FT}_\psi(\mathrm{T}_K(\varphi_z))(y),\end{aligned}$$

where the sum over x (resp. t) is over points in $\mathbf{F}_p$ which are not poles of g (resp. f), which gives (1).

Next, we have by definition

$$\varphi_z(x)=-\frac{1}{\sqrt{p}}\sum_t\psi(-g(x)f(t)-zt)=\mathrm{T}_{\tilde K}(\psi_{-z}),$$

if x is not a pole of g, with t ranging over $t\in\mathbf{F}_p$ not a pole of f. This gives (2) for such x, and the other case is immediate. ☐

Let $\mathcal{K}=\mathcal{L}_{\psi(g(Y)f(X))}$ (resp. $\tilde{\mathcal{K}}=\mathcal{L}_{\psi(-g(X)f(Y))}$); this Artin-Schreier ℓ-adic sheaf on $\mathbf{A}^2_{\mathbf{F}_p}$ has trace function equal to $K(x,y)$ (resp. to $\tilde K(x,y)$) by definition. Note that $\mathbf{c}(\mathcal{K})$ and $\mathbf{c}(\tilde{\mathcal{K}})$ are bounded in terms of $\mathbf{c}(f)$ and $\mathbf{c}(g)$ only.

It is convenient (in this section only) to normalize $T^i_{\mathcal{K}}$ and $T^i_{\tilde{\mathcal{K}}}$ to have weights $\leqslant 0$ for $i=1$, i.e., we denote

$$T^i_{\mathcal{K}}(\mathcal{F})=R^ip_{1,!}(p_2^*\mathcal{F}\otimes\mathcal{K})(1/2),\qquad T^i_{\tilde{\mathcal{K}}}(\mathcal{F})=R^ip_{1,!}(p_2^*\mathcal{F}\otimes\tilde{\mathcal{K}})(1/2).$$

We also denote by FT^i_ψ the corresponding cohomological operation for the Fourier transform

$$\mathrm{FT}^i_\psi(\mathcal{F})=R^ip_{1,!}(p_2^*\mathcal{F}\otimes\mathcal{L}_{\psi(XY)})(1/2).$$

We need the following simple lemma concerning the Fourier transform:

Lemma 5.13. (1) *We have*

$$\mathrm{FT}^1_\psi(\boldsymbol{\delta}_a) = \mathrm{FT}^2_\psi(\boldsymbol{\delta}_a) = 0,$$
$$\mathrm{FT}^2_\psi(\mathcal{L}_{\psi(aX)}) = \boldsymbol{\delta}_a(-1),$$
$$\mathrm{FT}^2_\psi(\mathcal{F}) = 0$$

where $\boldsymbol{\delta}_a$ denotes the punctual constructible sheaf supported on $\{a\}$ of rank 1 with trivial Frobenius action at $\{a\}$, and where $\mathcal{F}$ is any middle-extension sheaf which has no Artin-Schreier sheaf $\mathcal{L}_{\psi(aX)}$ as Jordan-Hölder factor.

(2) *Let $\mathcal{F}$ be a constructible sheaf modulo p. Let*

$$0 \longrightarrow \mathcal{F}_1 \longrightarrow \mathcal{F} \longrightarrow \mathcal{F}^{mid} \longrightarrow \mathcal{F}_2 \longrightarrow 0$$

be the canonical exact decomposition where $\mathcal{F}_1$ and $\mathcal{F}_2$ are punctual and $\mathcal{F}^{mid}$ is a middle-extension. We then have canonical isomorphisms

$$\mathrm{FT}^2_\psi(\mathcal{F}) \simeq \mathrm{FT}^2_\psi(\mathcal{F}^{mid}).$$

(3) *Let $\mathcal{F}$ be a middle-extension sheaf modulo p. Then $\mathrm{FT}^2_\psi(\mathcal{F})$ is punctual and its support has size bounded in terms of the conductor of $\mathcal{F}$ only.*

Proof. (1) The fiber over $x \in \bar{\mathbf{F}}_p$ of $\mathrm{FT}^i_\psi(\mathcal{F})$ is

$$H^i_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_p, \mathcal{F} \otimes \mathcal{L}_{\psi(xX)})$$

for a constructible sheaf $\mathcal{F}$ by the proper base change theorem. If $\mathcal{F} = \boldsymbol{\delta}_a$, then $\mathcal{F} \otimes \mathcal{L}_{\psi(xX)}$ is punctual and supported on $P = \{a\}$. By excision, we find that for $i = 1$, 2, we have

$$H^i_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_p, \boldsymbol{\delta}_a \otimes \mathcal{L}_{\psi(xX)}) \simeq H^i_c(P \times \bar{\mathbf{F}}_p, \boldsymbol{\delta}_a \otimes \mathcal{L}_{\psi(xX)}) = 0.$$

Similarly, the fiber of $\mathrm{FT}^2_\psi(\mathcal{L}_{\psi(aX)})$ at x is

$$H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_p, \mathcal{L}_{\psi((a+x)X)})$$

which vanishes unless $x = -a$. For $x = -a$, it is of dimension 1, and the local Frobenius acts trivially.

Finally, the last property of middle-extension sheaves is one of the defining properties of Fourier sheaves (as in [10, §8]).

(2) The functors $\mathcal{F} \mapsto \mathrm{FT}^i_\psi(\mathcal{F})$ acting on constructible sheaves form a cohomological functor, i.e., a short exact sequence

$$0 \longrightarrow \mathcal{F}_1 \longrightarrow \mathcal{F}_2 \longrightarrow \mathcal{F}_3 \longrightarrow 0$$

of constructible sheaves gives a long(ish) exact sequence

$$\begin{aligned}(5.2)\quad 0 \longrightarrow \mathrm{FT}^0_\psi(\mathcal{F}_1) \longrightarrow \mathrm{FT}^0_\psi(\mathcal{F}_2) \longrightarrow \mathrm{FT}^0_\psi(\mathcal{F}_3)&\\ \longrightarrow \mathrm{FT}^1_\psi(\mathcal{F}_1) \longrightarrow \mathrm{FT}^1_\psi(\mathcal{F}_2) \longrightarrow \mathrm{FT}^1_\psi(\mathcal{F}_3) \longrightarrow&\\ \mathrm{FT}^2_\psi(\mathcal{F}_1) \longrightarrow \mathrm{FT}^2_\psi(\mathcal{F}_2) \longrightarrow \mathrm{FT}^2_\psi(\mathcal{F}_3) \longrightarrow 0.&\end{aligned}$$

Consider the canonical adjunction map

$$\alpha : \mathcal{F} \longrightarrow \mathcal{F}^{mid} = j_* j^* \mathcal{F},$$

where $j : U \hookrightarrow \mathbf{A}^1$ is the open immersion of a dense open set where $\mathcal{F}$ is lisse. Applying the long exact sequence to the short exact sequence

$$0 \longrightarrow \ker(\alpha) \longrightarrow \mathcal{F} \longrightarrow \operatorname{im}(\alpha) \longrightarrow 0$$

we obtain an isomorphism

$$\mathrm{FT}^2_\psi(\mathcal{F}) \simeq \mathrm{FT}^2_\psi(\operatorname{im}(\alpha))$$

since $\ker(\alpha)$ is punctual and thus $\mathrm{FT}^1_\psi(\ker(\alpha)) = 0$ by (1) (extended by induction to punctual support with possibly more than one point). Then, applying the long exact sequence to the short exact sequence

$$0 \longrightarrow \operatorname{im}(\alpha) \longrightarrow \mathcal{F}^{mid} \longrightarrow \operatorname{coker}(\alpha) \longrightarrow 0,$$

we get

$$\mathrm{FT}^2_\psi(\operatorname{im}(\alpha)) \simeq \mathrm{FT}^2_\psi(\mathcal{F}^{mid})$$

since $\operatorname{coker}(\alpha)$ is also punctual. Together, these isomorphisms give (2).

(3) This is clear from (1), by induction on the length of a composition series of $\mathcal{F}$. □

Proof of the theorem. We will sheafify the formulas of Lemma 5.12. For clarity of notation, we denote by $t(\cdot;\mathcal{F})$ the trace function of a sheaf in this proof. The Grothendieck trace formula gives

$$\mathrm{T}_K(t(\cdot;\mathcal{F})) = \sum_{i=0}^{2} (-1)^i t(\cdot; T^i_{\mathcal{K}}(\mathcal{F}))$$

for any constructible sheaf $\mathcal{F}$, resp.

$$\mathrm{FT}_\psi(t(\cdot;\mathcal{F})) = \sum_{i=0}^{2} (-1)^i t(\cdot; \mathrm{FT}^i_\psi(\mathcal{F})).$$

By the Riemann Hypothesis over finite fields, if $\mathcal{F}$ is mixed of weights $\leqslant 0$, then $T^i_{\mathcal{K}}(\mathcal{F})$ and $\mathrm{FT}^i_\psi(\mathcal{F})$ are mixed of weights $\leqslant i-1$ (the -1 comes from the Tate twist in the definitions).

Inserting these formulas in Lemma 5.12 we find, for given z, that

$$D = \sum_{i=0}^{2}\sum_{j=0}^{2}\sum_{k=0}^{2}(-1)^{i+j+k} t(\cdot; \mathrm{FT}^i_\psi(T^j_{\mathcal{K}}(T^k_{\tilde{\mathcal{K}}}(\mathcal{L}))))$$

where $\mathcal{L} = \mathcal{L}_{\psi(zX)}$. For given (i,j,k), the sheaf

$$\mathrm{FT}^i_\psi(T^j_{\mathcal{K}}(T^k_{\tilde{\mathcal{K}}}(\mathcal{L})))$$

is mixed of weights $\leqslant i+j+k-3$. In fact, the fiber over x of $T^k_{\tilde{\mathcal{K}}}(\mathcal{L})$ is

$$H^i_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_p, \mathcal{L}_{\psi(-zY)} \otimes \mathcal{L}_{\psi(-g(x)f(Y))})$$

so that $T^k_{\tilde{\mathcal{K}}}(\mathcal{L}) = 0$ for $k \neq 1$ if, as assumed, f is not a polynomial of degree $\leqslant 1$. The remaining sheaves

$$\mathrm{FT}^i_\psi(T^j_{\mathcal{K}}(T^1_{\tilde{\mathcal{K}}}(\mathcal{L})))$$

are mixed of weights $\leqslant i+j-2$. In particular, they are mixed of weight $\leqslant 1$, except for the sheaf

$$\mathcal{G} = \mathrm{FT}^2_\psi(T^2_{\mathcal{K}}(T^1_{\tilde{\mathcal{K}}}(\mathcal{L}))).$$

But, by Lemma 5.13 (2) and (3), we see that $\mathcal{G}$ is punctual with support (say Δ_z) of size bounded by the conductor of $T^2_{\mathcal{K}}(T^1_{\tilde{\mathcal{K}}}(\mathcal{L}))$. We have

$$\mathbf{c}(\mathcal{L}) \leqslant 3$$

for all z, and hence Theorem 2.3 implies that there exists a constant $B_1 \geqslant 0$, depending only on the conductors of g and f, such that

$$\mathbf{c}(T^2_{\mathcal{K}}(T^1_{\tilde{\mathcal{K}}}(\mathcal{L}))) \leqslant B_1$$

for all z. Hence there exists $B_2 \geqslant 0$, depending only on $\mathbf{c}(g)$ and $\mathbf{c}(f)$ such that $|D(y)| \leqslant B_2$ for $z \notin \Delta_z$ (by the Riemann Hypothesis as in the proof of Proposition 5.6).

Finally, by definition of the punctual part of the conductor, for each point $a \in \Delta_z$, the dimension of the fiber of $\mathcal{G}$ at a is bounded by B, and this gives

$$|D(a)| \leqslant B_3\sqrt{p}$$

again for a constant B_3 that depends only on $\mathbf{c}(f)$ and $\mathbf{c}(g)$.

To finish the proof of Theorem 5.11, we need to consider the behavior of $z \mapsto C(y,z)$ for fixed y. But since $C(y,z) = \overline{C(z,y)}$, this reduces to the previous case. ☐

In their special case, Fouvry and Michel use more delicate information concerning the sheaves with trace function $x \mapsto R(x,y;p)$ for $f = X^k$, and in particular use results of Katz to determine that these sheaves are, for y varying, geometrically irreducible and geometrically isomorphic only for $y^k = z^k$. This leads to their more precise computation of the diagonal set.

The motivation of Fouvry and Michel was the study of certain short character sums (going below the Polya-Vinogradov range). We expect that Theorem 5.11 will have similar applications, and we hope to come back to this soon.

5.5. **Final remarks.** Although the previous examples (and other applications) show that Theorem 2.3 is extremely useful, there remains the natural question of extending the result to more general kernel sheaves $\mathcal{K}$ on $\mathbf{A}^2$. The difficulty in this case is partly conceptual: it is not clear what should be the definition of the conductor of $\mathcal{K}$, and without such a notion, the problem can not even be stated properly. However, one may consider other "ad-hoc" cases. Indeed, in our study of the ternary divisor function [6], we studied (in effect) the case of the kernel

$$\mathcal{K} = m^*\mathcal{K}\ell_2$$

(on $\mathbf{G}_m \times \mathbf{G}_m$) where $m : \mathbf{G}_m \times \mathbf{G}_m \to \mathbf{G}_m$ is the multiplication map and $\mathcal{K}\ell_2$ is Deligne's one-variable *Kloosterman sheaf*, with trace function (over $\mathbf{F}_p$) given by

$$t_{\mathcal{K}\ell_2}(x) = -\frac{1}{\sqrt{p}} \sum_{y \in \mathbf{F}_p^\times} e\Bigl(\frac{xy + y^{-1}}{p}\Bigr).$$

In this case, the associated transform can be easily studied, because by opening the Kloosterman sum, one obtains the expression

$$t_{T_{\mathcal{K}}(\mathcal{F})}(x) = -\frac{1}{\sqrt{p}} \sum_{y \in \mathbf{F}_p^\times} \widehat{t_{\mathcal{F}}}(y) e\Bigl(\frac{xy^{-1}}{p}\Bigr)$$

for $x \in \mathbf{F}_p$, which shows that the transform has a character kernel when expressed in terms of the Fourier transform.

6. Setting up the proof

To clarify the proof of Theorems 2.3 and 2.5, and in view of further generalizations, we introduce the following definition:

Definition 6.1 (Continuity)**.** (1) Let

$$i\,:\,(f,\mathcal{F})\mapsto i(f,\mathcal{F})$$

be any real-valued map taking a pair $(f,\mathcal{F})$ as input, where f is a non-constant rational function in $\mathbf{F}_q(X,Y)$ for some finite field $\mathbf{F}_q$ and $\mathcal{F}$ is a middle-extension ℓ-adic sheaf on the affine line over $\mathbf{F}_q$. Then we say that i is *continuous* if and only if there exists an integer $C\geqslant 1$ such that

$$|i(f,\mathcal{F})|\leqslant (2\,\mathbf{c}(f)\,\mathbf{c}(\mathcal{F}))^C$$

for all pairs $(f,\mathcal{F})$ as above such that $\mathbf{c}(f)<p$.[1]

(2) Similarly, if

$$j\,:\,f\mapsto j(f),\ \text{resp. }\ k\,:\,\mathcal{F}\mapsto k(\mathcal{F})$$

are real-valued maps taking as input a non-constant rational function $f\in\mathbf{F}_q(X,Y)$ for some finite field $\mathbf{F}_q$ (resp. a middle-extension ℓ-adic sheaf $\mathcal{F}$ on the affine line over $\mathbf{F}_q$), then we say that j (resp. k) is *continuous* if and only if there exists an integer $C\geqslant 1$ such that

$$|j(f)|\leqslant (2\,\mathbf{c}(f))^C,\ \text{resp. }\ |k(\mathcal{F})|\leqslant (2\,\mathbf{c}(\mathcal{F}))^C,$$

for all f with $\mathbf{c}(f)<p$ (resp. all middle-extension sheaves $\mathcal{F}$).

Remark 6.2. Some of our arguments are easier to follow and check if one uses a weaker definition of continuity, where one only asks that

$$|i(f,\mathcal{F})|\leqslant \Psi(\mathbf{c}(f),\mathbf{c}(\mathcal{F}))$$

for some function Ψ taking positive integral values. For some basic applications, such a statement is also sufficient, and the reader might wish to consider this as the notion of continuity in a first reading.

Example 6.3. For instance, Theorem 2.3 asserts that the maps

$$(f,\mathcal{F})\mapsto \mathbf{c}(T^i_{\mathcal{K}}(\mathcal{F}))$$

are continuous, and Theorem 2.5 that the maps

$$(f,\mathcal{F})\mapsto \dim H^i_c(\mathbf{A}^2\times\bar{\mathbf{F}}_q, p_2^*\mathcal{F}\otimes\mathcal{K})$$

are continuous. Lemma 4.7 proves that the functions

$$\mathcal{F}\mapsto \dim H^i_c(\mathbf{A}^1\times\bar{\mathbf{F}}_q,\mathcal{F})$$

are continuous.

[1] This restriction may seem artificial, and it is possible that it would not be needed for our results. But it has no influence on the applications.

Clearly, if we fix one argument of a continuous map $i(f, \mathcal{F})$ and let the other vary, this gives a continuous map of this second argument. Also, a sum $i_1 + i_2$ of continuous functions is also continuous, as well as a product $i_1 i_2$.

For simplicity, we denote

$$\begin{aligned} c_i(f, \mathcal{F}) &= \mathbf{c}(T^i_{\mathcal{K}}(\mathcal{F})), \qquad 0 \leqslant i \leqslant 2 \\ h^j(f, \mathcal{F}) &= \dim H^j_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, p_2^*\mathcal{F} \otimes \mathcal{L}_{\psi(f)}), \qquad 0 \leqslant j \leqslant 2 \\ m(f, \mathcal{F}) &= \mathrm{rank}(T^1_{\mathcal{K}}(\mathcal{F})) + n(T^1_{\mathcal{K}}(\mathcal{F})) + \mathrm{pct}(T^1_{\mathcal{K}}(\mathcal{F})). \end{aligned}$$

The proof of Theorems 2.3 and 2.5 will be based on the following steps:

Proposition 6.4. *The following assertions are true:*

(1) *The map c_0 is continuous;*

(2) *For $0 \leqslant j \leqslant 2$, the map*

$$f \mapsto h^j(f, \bar{\mathbf{Q}}_\ell) = \dim H^j_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(f)})$$

is continuous;

(3) *If $h^2(f, \mathcal{F})$ is continuous, then $c_2(f, \mathcal{F})$ is continuous;*

(3bis) *If $h^2(f, \bar{\mathbf{Q}}_\ell)$ is continuous, then $c_2(f, \bar{\mathbf{Q}}_\ell)$ is continuous;*

(4) *If $f \mapsto c_1(f, \bar{\mathbf{Q}}_\ell)$ and $f \mapsto c_2(f, \bar{\mathbf{Q}}_\ell)$ are continuous, then $h^2(f, \mathcal{F})$ is continuous;*

(5) *If $c_2(f, \mathcal{F})$ is continuous, then $m(f, \mathcal{F})$ is continuous;*

(5bis) *If $c_2(f, \bar{\mathbf{Q}}_\ell)$ is continuous, then $m(f, \bar{\mathbf{Q}}_\ell)$ is continuous;*

(6) *If $m(f, \mathcal{F})$ and $h^2(f, \mathcal{F})$ are both continuous, then $c_1(f, \mathcal{F})$ is continuous;*

(6bis) *If $m(f, \bar{\mathbf{Q}}_\ell)$ and $h^2(f, \bar{\mathbf{Q}}_\ell)$ are both continuous, then $c_1(f, \bar{\mathbf{Q}}_\ell)$ is continuous.*

(7) *If $c_i(f, \mathcal{F})$ is continuous for $0 \leqslant i \leqslant 2$, then $h^j(f, \mathcal{F})$ is continuous for all $0 \leqslant j \leqslant 2$.*

We now explain how to deduce Theorems 2.3 and 2.5 from this proposition. Since this may also look like spaghetti-mathematics, the reader may also wish to go straight to Sections 7 and 8 (possibly in the opposite order) which together give an account of the proof for the special case of the Fourier transform (and discuss another example arising in the Polymath8 project), in which case the flow of the proof is much easier to follow.

First of all, c_0 is continuous by (1), so we must show that c_1, c_2 and the h^j are continuous.

Step 1. Using (2), we can apply (3bis) and deduce that $f \mapsto c_2(f, \bar{\mathbf{Q}}_\ell)$ is continuous. By (5bis), it follows that $m(f, \bar{\mathbf{Q}}_\ell)$ is continuous. Combining this with (6bis) and (2) again, we deduce that $c_1(f, \bar{\mathbf{Q}}_\ell)$ is continuous.

At this point, we have proved both theorems in the special case when $\mathcal{F} = \bar{\mathbf{Q}}_\ell$ is the trivial sheaf.

Step 2. From (4) and Step 1, we see that $h^2(f, \mathcal{F})$ is continuous. This fact combined with (3) shows that $c_2(f, \mathcal{F})$ is continuous. In turn, (5) then proves that $m(f, \mathcal{F})$ is continuous, and finally (6) allows us to conclude that $c_1(f, \mathcal{F})$ is continuous.

At this point we have proved Theorem 2.3 (and the continuity of $h^2(f, \mathcal{F})$); by (7), we deduce that all h^j are continuous.

Remark 6.5. (1) We will in fact prove (3) and (3bis) simultaneously by proving a direct relation between $c_2(f, \mathcal{F})$ and $h^2(f, \mathcal{F})$, and similarly for (5) and (5bis), (6) and (6bis).

(2) The most crucial points in Proposition 6.4 are (4), which allows us to pass from properties known for the trivial sheaf only, to properties of all sheaves, and (2), which gives the starting point of the argument for the trivial sheaf, and which comes from the bounds for

Betti numbers of Bombieri, Adolphson-Sperber and Katz. On the other hand, what turns out to be most involved (though relatively elementary) is the proof of (5) and (5bis), which amounts to controlling all invariants defining the conductor of $T^1_{\mathcal{K}}(\mathcal{F})$ except for the sum of Swan conductors. This part was essentially swept under the rug in Section 3, which explains partly why the proof of Theorem 2.3 is quite a bit longer than that motivating sketch might suggest.

(3) It is only in the proof of (5) and (5bis) that we will use the restriction that continuity applies to f with $\mathbf{c}(f) < p$.

7. Spectral sequence argument

We state here the few simple facts about spectral sequences that we require. We first recall the basic formalism, referring to [17, Appendix. B] for a survey and [21, Ch. 10] for details.

Let k be a fixed field. A converging (first quadrant) spectral sequence

$$E_2^{p,q} \Rightarrow E^n,$$

of k-vector spaces involves (1) vector spaces $E_2^{p,q}$ defined for p, $q \geqslant 0$; (2) vector spaces E^n defined for $n \geqslant 0$; (3) linear maps

$$d_2^{p,q} \;:\; E_2^{p,q} \longrightarrow E_2^{p+2,q-1}, \tag{7.1}$$

(called differentials)[2] for all p and q (with the convention $E_2^{p,q} = 0$ if p or q is negative), such that

$$d_2^{p,q} \circ d_2^{p-2,q+1} = 0.$$

Remark 7.1. The use of the indices p and q for the spectral sequence is almost universal, although it clashes with the usual convention that p is a prime and q a power of p. We will use i and j instead of p and q when both notation are involved, although the difference in context should avoid confusion.

One defines

$$E_3^{p,q} = \ker d_2^{p,q} / \operatorname{im} d_2^{p+2,q-1}, \tag{7.2}$$

and one shows that there are linear maps

$$d_3^{p,q} \;:\; E_3^{p,q} \longrightarrow E_3^{p+3,q-2}, \tag{7.3}$$

such that $d_3^{p,q} \circ d_3^{p-3,q+2} = 0$. This process is then suitably iterated to obtain $E_j^{p,q}$ for all $j \geqslant 2$, and differentials

$$d_j^{p,q} \;:\; E_j^{p,q} \longrightarrow E_j^{p+j,q-j+1}$$

(with composites vanishing).

One says that *the spectral sequence degenerates at the E_j-level* (where $j = 2$ or 3) if $d_i^{p,q} = 0$ for all p, $q \geqslant 0$ and $i \geqslant j$. When this is the case, the formalism gives (among other things) the following relation between the $E_j^{p,q}$ and the spaces E^n: we have for all $n \geqslant 0$, an isomorphism

$$E^n \simeq \bigoplus_{p=0}^{n} E_j^{p,n-p}, \tag{7.4}$$

[2] Note that these differentials show that p and q do not play symmetric roles.

of k-vector spaces. (There is often more structure involved, but this will suffice for us.)

Furthermore, whether the spectral sequence degenerates at the E_2 or E_3 level or not, there is an exact sequence

$$0 \to E_2^{1,0} \longrightarrow E^1 \longrightarrow E_2^{0,1} \longrightarrow E_2^{2,0}. \tag{7.5}$$

All these facts are stated in [17, p. 307–309]. The next proposition then summarizes all results we will need from spectral sequences:

Proposition 7.2. *Let k be a field and let*

$$E_2^{p,q} \Rightarrow E^n$$

be a converging spectral sequence as above. Assume that $E_2^{p,q} = 0$ unless $0 \leqslant p \leqslant 2$ and $0 \leqslant q \leqslant 2$.

(1) *The spectral sequence degenerates at the E_3-level and we have*

$$E^2 \simeq E_3^{0,2} \oplus E_2^{1,1} \oplus E_3^{2,0}. \tag{7.6}$$

(2) *We have*

$$\dim E^n \leqslant \sum_{p=0}^{n} \dim E_2^{p,n-p},$$

and

$$\dim E_2^{0,2} \leqslant \dim E^2 + \dim E_2^{2,1}.$$

(3) *Assume in addition that $E_2^{p,q} = 0$ if $q = 0$. We have then $E_2^{0,1} \simeq E^1$.*

Proof. (1) From (7.2), we see that $E_3^{p,q} = 0$ unless $0 \leqslant p, q \leqslant 2$ since it is a quotient of a subspace of $E_2^{p,q}$. But then (7.3) shows that, for any p, q, either the source of the target of $d_3^{p,q}$ is zero. In fact, for all $i \geqslant 3$, either the target or the source of $d_i^{p,q}$ vanishes, and therefore the spectral sequence degenerates at that level.

By (7.4) we deduce that

$$E^2 \simeq E_3^{0,2} \oplus E_3^{1,1} \oplus E_3^{0,2},$$

but

$$E_3^{1,1} = \ker d_2^{1,1} / \operatorname{im} d_2^{-1,2},$$

and since $d_2^{1,1}$ and $d_2^{-1,2}$ are both zero (the target of the first and the source of the second are zero), we have $E_3^{1,1} = E_2^{1,1}$, hence (7.6).

(2) By (1) and (7.4), we have

$$E^n \simeq E_3^{n,0} \oplus E_3^{n-1,1} \oplus \cdots \oplus E_3^{0,n}.$$

Since

$$\dim E_3^{p,q} \leqslant \dim E_2^{p,q},$$

for all p and q, by (7.2), we obtain

$$\dim E^n = \sum_{p=0}^{n} \dim E_3^{p,q} \leqslant \sum_{p=0}^{n} \dim E_2^{p,q}.$$

Similarly, we note that

$$E_3^{0,2} = \ker d_2^{0,2} / \operatorname{im} d_2^{-2,1} = \ker d_2^{0,2},$$

and hence we have a short exact sequence

$$0 \longrightarrow E_3^{0,2} \longrightarrow E_2^{0,2} \xrightarrow{d_2^{0,2}} E_2^{2,1}$$

which implies that

$$\dim E_2^{0,2} \leqslant \dim E_3^{0,2} + \dim E_2^{2,1}.$$

From the degeneracy at the E_3-level, we then get

$$\dim E_3^{0,2} \leqslant \dim E^2,$$

hence the bound for $E_2^{0,2}$.

(3) The exact sequence (7.5), under the assumptions that $E_2^{p,0} = 0$, becomes

$$0 \longrightarrow E^1 \longrightarrow E_2^{0,1} \longrightarrow 0,$$

hence the result. □

The spectral sequences we use are given by the following lemma:

Lemma 7.3. *Let $\mathbf{F}_q$ be a finite field of characteristic p, $\ell \neq p$ a prime number. Let $f \in \mathbf{F}_q(X,Y)$ be a rational function, and denote*

$$\mathcal{K} = \mathcal{L}_{\psi(f(X,Y))},$$

where ψ is a non-trivial additive ℓ-adic character. Denote $f^*(X,Y) = f(Y,X) \in \mathbf{F}_q$, and

$$\mathcal{K}^* = \mathcal{L}_{\psi(f^*)}.$$

Let $\mathcal{F}$ be a constructible ℓ-adic sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$.

(1) For any dense open subsets U, V of $\mathbf{A}^1_{\mathbf{F}_q}$, with p_1, p_2 denoting the projection maps $U \times V \longrightarrow U$ and $U \times V \longrightarrow V$, respectively, there exist converging spectral sequences

$$E_2^{i,j} = H^i_c(\bar{U}, T^j_{\mathcal{K}}(\mathcal{F})) \Rightarrow H^{i+j}_c(\bar{U} \times \mathbf{A}^1, p_2^*\mathcal{F} \otimes \mathcal{K}),$$

$$E_2^{i,j} = H^i_c(\bar{V}, \mathcal{F} \otimes T^j_{\mathcal{K}^*}(\bar{\mathbf{Q}}_\ell)) \Rightarrow H^{i+j}_c(\mathbf{A}^1 \times \bar{V}, p_2^*\mathcal{F} \otimes \mathcal{K})$$

of $\bar{\mathbf{Q}}_\ell$-vector spaces

(2) These two spectral sequences satisfy $E_2^{i,j} = 0$ unless $0 \leqslant i \leqslant 2$ and $1 \leqslant j \leqslant 2$.

Proof. (1) The first spectral sequence is the Leray spectral sequence of the first projection map $p_1 : U \times \mathbf{A}^1 \to U$ and of the sheaf $p_2^*\mathcal{F} \otimes \mathcal{K}$ (see, e.g., [8, Th. 7.4.4 (ii)] or [17, Th. VI.3.2 (c)].)

The second spectral sequence arises from the Leray spectral sequence of the second projection $p_2 : \mathbf{A}^1 \times V \to V$ and of the sheaf $p_2^*\mathcal{F} \otimes \mathcal{K}$, namely

$$E_2^{i,j} = H^i_c(\bar{V}, R^j p_{2,!}(p_2^*\mathcal{F} \otimes \mathcal{K})) \Rightarrow H^{i+j}_c(\mathbf{A}^1 \times \bar{V}, p_2^*\mathcal{F} \otimes \mathcal{K})$$

together with the facts that

$$R^j p_{2,!}(p_2^*\mathcal{F} \otimes \mathcal{K}) = \mathcal{F} \otimes R^j p_{2,!}(\mathcal{L}_{\psi(f(X,Y))})$$

by the projection formula (see, e.g., [8, Th. 7.4.7]), and that we can identify $R^j p_{2,!}(\mathcal{L}_{\psi(f(X,Y))})$ with $T^j_{\mathcal{K}^*}(\bar{\mathbf{Q}}_\ell)$ (restricted to V).

(2) The fact that $E_2^{i,j} = 0$ unless $0 \leqslant i, j \leqslant 2$ is immediate from (1) and from the vanishing of cohomology of curves (resp. of higher-direct image sheaves for maps with curves as fibers) in Proposition 4.1, (1): the former constrains i to be between 0 and 2, and the second constrains similarly j.

For the vanishing when $j = 0$, we note that the stalk at $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$ of $R^0 p_{1,!}(p_2^* \mathcal{F} \otimes \mathcal{K})$ is, by the proper base change theorem, equal to

$$H_c^0(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_{\psi(f(x,Y))}) = 0$$

by Lemma 4.2. Similarly, the stalk of $R^0 p_{2,!}(\mathcal{K})$ at y is

$$H_c^0(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(f(X,y))}) = 0,$$

and these facts show that $E_2^{i,0} = 0$ for all i in both spectral sequences. □

8. Working out an example

This section is independent of the remainder of the proof of Proposition 6.4. What we do here is prove Theorem 2.3 in the special case of the Fourier transform, and of the special case which is important in the Polymath8 project (which turns out to be related). This should help understand the arguments in the coming sections (and the tools of the previous one). We will not strictly keep track of the fact that the conductor bounds for the Fourier transform are of polynomial size in terms of $\mathbf{c}(\mathcal{F})$, but this is easily checked to follow from the argument.

Thus we consider first

$$f = XY \in \mathbf{F}_q[X, Y],$$

and we write $\mathrm{FT}_\psi(\mathcal{F})$ for the corresponding transform

$$\mathrm{FT}_\psi(\mathcal{F}) = R^1 p_{1,!}(p_2^* \mathcal{F} \otimes \mathcal{L}_{\psi(XY)}),$$

which is the "naive" Fourier transform (see [10, Chap. 8]). Note that $\mathbf{c}(f) = 2$, independently of q. We will not need, however to restrict to primes $p > 2$.

Let $\mathcal{F}$ be a middle-extension sheaf. We start by attempting to evaluate $m(\mathrm{FT}_\psi(\mathcal{F}))$. By the proper base change theorem, the fiber of $\mathrm{FT}_\psi(\mathcal{F})$ at $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$ is

$$H_c^1(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_{\psi(xY)}).$$

From Lemmas 4.8 and 4.7, we already see that the maximum over x of the dimension of these spaces, hence the rank of $\mathrm{FT}_\psi(\mathcal{F})$, is bounded in terms of $\mathbf{c}(\mathcal{F})$.

We now wish to apply the Euler-Poincaré formula to compute more precisely the dimension of this cohomology space. Since $H_c^0(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_{\psi(xY)}) = 0$ (by Lemma 4.2), it is easiest to do so if the second cohomology group happens to vanish. But if the group

$$H_c^2(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_{\psi(xY)})$$

is non-zero, this means that $\mathcal{L}_{\psi(-xY)}$ is a quotient of $\mathcal{F}$ (by the irreducibility of $\mathcal{L}_{\psi(xY)}$ and the coinvariant formula for H_c^2, which we already used in the proof of Lemma 4.7.) Since the sheaves $\mathcal{L}_{\psi(xY)}$ are pairwise non-isomorphic as x varies in $\bar{\mathbf{F}}_q$, this may happen, at most, for $\leqslant \mathrm{rank}(\mathcal{F}) \leqslant \mathbf{c}(\mathcal{F})$ values of x (and it never happens if $\mathcal{F}$ is a Fourier sheaf in the sense of [10, Def. 8.2.1.2].) Let $U \subset \mathbf{A}^1$ be the open dense subset where $H_c^2(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_{\psi(xY)}) = 0$ for $x \in U(\bar{\mathbf{F}}_q)$.

Denote $\mathcal{G} = \mathrm{FT}_\psi(\mathcal{F})$. For $x \in U(\bar{\mathbf{F}}_q)$, we have

$$\dim \mathcal{G}_x = -\mathrm{rank}(\mathcal{F}) + \sum_{y \in \mathbf{A}^1(\bar{\mathbf{F}}_q)} (\mathrm{Swan}_y(\mathcal{F}) + \mathrm{drop}_y(\mathcal{F})) + \mathrm{Swan}_\infty(\mathcal{F} \otimes \mathcal{L}_{\psi(xY)}) \tag{8.1}$$

(see [10, 8.5.3]) where we used the fact that $\mathcal{L}_{\psi(xY)}$ is lisse of rank 1 on $\mathbf{A}^1$, so that

$$\mathrm{drop}_y(\mathcal{F} \otimes \mathcal{L}_{\psi(xY)}) = \mathrm{drop}_y(\mathcal{F}), \qquad \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_{\psi(xY)}) = \mathrm{Swan}_y(\mathcal{F}),$$

for all $y \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$.

In particular, we see that the rank of the stalks is constant on the set of $x \in U(\bar{\mathbf{F}}_q)$ where the Swan conductor $\mathrm{Swan}_\infty(\mathcal{F} \otimes \mathcal{L}_{\psi(xY)})$ is constant. A relatively simple argument (see [10, Cor. 8.5.5], which we will generalize straightforwardly in Lemma 12.3 (3), below) shows that this Swan conductor is constant except for at most $\mathrm{rank}\,\mathcal{F} \leqslant \mathbf{c}(\mathcal{F})$ values of x. Let U_1 be the dense open subset of U where the Swan conductor is constant, say equal to $d \geqslant 0$. We then have

$$\dim \mathcal{G}_x = -\,\mathrm{rank}(\mathcal{F}) + \sum_{y \in \mathbf{A}^1(\bar{\mathbf{F}}_q)} (\mathrm{Swan}_y(\mathcal{F}) + \mathrm{drop}_y(\mathcal{F})) + d$$

for all $x \in U_1(\bar{\mathbf{F}}_q)$.

We next claim that $\mathrm{pct}(\mathcal{G}) = 0$; it follows then that $\mathcal{G}$ is lisse on U_1 (this is similar, but slightly more precise, than Lemma 4.10), and consequently that $m(\mathcal{G})$ is bounded in terms of $\mathbf{c}(\mathcal{F})$.

We use the first spectral sequence of Lemma 7.3 (taking there $U = \mathbf{A}^1$) and apply Proposition 7.2, (3) to deduce

$$\mathrm{pct}(\mathcal{G}) = \dim H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{G}) = \dim H^1_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, p_2^*\mathcal{F} \otimes \mathcal{L}_{\psi(XY)}).$$

Let $S \subset \mathbf{A}^1$ be the finite set of singularities of $\mathcal{F}$ in $\mathbf{A}^1$ and $T = \mathbf{A}^1 \times S \subset \mathbf{A}^2$. The sheaf $\mathcal{M} = p_2^*\mathcal{F} \otimes \mathcal{L}_{\psi(XY)}$ is lisse on the dense open set $W = \mathbf{A}^2 - T$. Applying excision (4.1), we get an exact sequence

$$\cdots \longrightarrow H^1_c(\bar{W}, \mathcal{M}) \longrightarrow H^1_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{M}) \longrightarrow H^1_c(\bar{T}, \mathcal{M}) \longrightarrow \cdots.$$

We have

$$H^1_c(\bar{W}, \mathcal{M}) = 0$$

by (4.4), because W is an affine surface and $\mathcal{M}$ is lisse on W. Also, since T is a disjoint union of "horizontal" lines, we have

$$H^1_c(\bar{T}, \mathcal{M}) = \bigoplus_{y \in S} H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(yX)}) \otimes \mathcal{F}_y = 0$$

because $H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(yX)}) = 0$ for all $y \in S$ (including $y = 0$). The excision exact sequence then gives $H^1_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{M}) = 0$, as claimed.

By Lemma 4.11, we deduce that the conductor of $\mathrm{FT}_\psi(\mathcal{F})$ is bounded in terms of the conductor of $\mathcal{F}$, and of the invariant

$$\sigma(\mathcal{G}) = \dim H^1_c(\bar{U}_1, \mathcal{G}).$$

By the first spectral sequence of Lemma 7.3 (taking U to be U_1) and Proposition 7.2, (1), we get

$$\dim H^1_c(\bar{U}_1, \mathcal{G}) \leqslant \dim H^2_c(\bar{U}_1 \times \mathbf{A}^1, \mathcal{M}).$$

To compute this last group, we first use the second spectral sequence, which shows that

$$H^2_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{M}) \simeq E_3^{0,2} \oplus E_3^{1,1} \oplus E_3^{2,0} \simeq E_3^{0,2} \oplus E_3^{1,1}.$$

We have

$$\begin{aligned}E_2^{1,1} &= H_c^1(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, R^1 p_{2,!}(p_2^*\mathcal{F} \otimes \mathcal{L}_{\psi(XY)}))\\ &= H_c^1(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes R^1 p_{2,!}\mathcal{L}_{\psi(XY)})\end{aligned}$$

by the projection formula. But the sheaf $R^1 p_{2,!}\mathcal{L}_{\psi(XY)}$ is zero, since the fiber at any $y \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$ is

$$H_c^1(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(yX)}) = 0.$$

As for $E_3^{0,2}$, it is a subspace of

$$E_2^{0,2} = H_c^0(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, R^2 p_{2,!}(p_2^*\mathcal{F} \otimes \mathcal{L}_{\psi(XY)})) = H_c^0(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes R^2 p_{2,!}\mathcal{L}_{\psi(XY)}).$$

The stalk of $R^2 p_{2,!}\mathcal{L}_{\psi(XY)}$ at $y \in \bar{\mathbf{F}}_q$ is

$$H_c^2(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(yX)}) = \begin{cases} \bar{\mathbf{Q}}_\ell & \text{if } y = 0\\ 0 & \text{otherwise,}\end{cases}$$

so the sheaf $R^2 p_{2,!}\mathcal{L}_{\psi(XY)}$ is punctual and supported at 0 with stalk $\mathcal{F}_0$. Hence the dimension of $E_2^{0,2}$ is at most the rank $\mathrm{rank}(\mathcal{F}) \leqslant \mathbf{c}(\mathcal{F})$. Thus we obtain

$$\dim H_c^2(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{M}) \leqslant \mathbf{c}(\mathcal{F}).$$

We finally apply excision (4.1) again to compare $H_c^2(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{M})$ and $H_c^2(\bar{U}_1 \times \mathbf{A}^1, \mathcal{M})$. Let $C = \mathbf{A}^1 - U_1$. We have an exact sequence

$$\cdots \longrightarrow H_c^1(\bar{C}, \mathcal{M}) \longrightarrow H_c^2(\bar{U}_1 \times \mathbf{A}^1, \mathcal{M}) \longrightarrow H_c^2(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{M}) \longrightarrow \cdots$$

so that (by the above) we have

$$\dim H_c^2(\bar{U}_1 \times \mathbf{A}^1, \mathcal{M}) \leqslant \dim H_c^1(\bar{C}, \mathcal{M}) + \mathbf{c}(\mathcal{F})$$

(we just spelled-out the inequality (4.3) explicitly.)

From the definition of C, we get

$$H_c^1(\bar{C}, \mathcal{M}) \simeq \bigoplus_{x \in (\mathbf{A}^1 - U_1)(\bar{\mathbf{F}}_q)} H_c^1(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_{\psi(xY)}).$$

From Corollary 4.9, we know that the maximum dimension of the

$$H_c^1(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_{\psi(xY)})$$

is bounded in terms of $\mathbf{c}(\mathcal{F})$, and since we saw already that $|\mathbf{A}^1 - U_1|$ is also bounded, we conclude that $\sigma(\mathrm{FT}_\psi(\mathcal{F}))$ is bounded in terms of the conductor of $\mathcal{F}$. This concludes the proof of Theorem 2.3 for the Fourier transform. We state it formally for convenience:

Corollary 8.1. *Let $\mathbf{F}_q$ be a finite field of characteristic p, $\ell \neq p$ a prime number. Let ψ be a non-trivial additive ℓ-adic character of $\mathbf{F}_q$. There exists a function $n \mapsto C(n)$ with positive integral values such that, for any middle-extension sheaf $\mathcal{F}$ on $\mathbf{A}^1_{\mathbf{F}_q}$, the naive Fourier transform $\mathrm{FT}_\psi(\mathcal{F})$ satisfies $\mathrm{pct}(\mathrm{FT}_\psi(\mathcal{F})) = 0$ and we have*

$$\mathbf{c}(\mathrm{FT}_\psi(\mathcal{F})) \leqslant C(\mathbf{c}(\mathcal{F})).$$

As we already mentioned in the introduction, we obtain in [4, Prop. 8.2] the estimate

$$\mathbf{c}(\mathrm{FT}_{\psi}(\mathcal{F})) \leqslant 10\,\mathbf{c}(\mathcal{F})^2$$

for $\mathcal{F}$ a Fourier sheaf on $\mathbf{A}^1_{\mathbf{F}_p}$, using the local study of the Fourier transform, due to Laumon [16]. It is clear from the arguments above that they can also be used to give a completely effective upper bound.

Remark 8.2. (1) A Fourier sheaf is defined to be a middle-extension sheaf which has no subsheaf or quotient sheaf geometrically isomorphic to an Artin-Schreier sheaf $\mathcal{L}_{\psi(aX)}$. For a sheaf which is not of this type, the naive Fourier transform is not the right object to consider, but this is of course not due to a failure of continuity.

For instance, if $\mathcal{F} = \mathcal{L}_{\psi(Y)}$ (a typical non-Fourier sheaf!) we have

$$R^1 p_{1,!}(p_2^*\mathcal{F} \otimes \mathcal{L}_{\psi(XY)}) = 0$$

since the stalk of this sheaf at $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$ is

$$H^1_c(\mathbf{A}^1 \otimes \bar{\mathbf{F}}_q, \mathcal{L}_{\psi((1+x)Y)}) = 0$$

for all values of x. This certainly has bounded conductor!

(2) For Fourier sheaves, other properties of the Fourier transform are established, relatively elementarily, in [10, 8.2.5, 8.4.1]: the Fourier transform is again a Fourier sheaf, and the Fourier transform of a geometrically irreducible Fourier sheaf is again geometrically irreducible.

We can deduce a version of the irreducibility property (which suffices in many applications) from the diophantine irreducibility criterion of Lemma 5.10. Indeed, if $\mathcal{F}$ is a middle-extension Fourier sheaves which is pointwise pure of weight 0, we have the discrete Plancherel formula

$$\frac{1}{q^{\nu}} \sum_{x \in \mathbf{F}_{q^{\nu}}} |t_{\mathcal{F}}(x, q^{\nu})|^2 = \frac{1}{q^{2\nu}} \sum_{t \in \mathbf{F}_{q^{\nu}}} |t_{\mathrm{FT}_{\psi}(\mathcal{F})}(t, q^{\nu})|^2$$

for $\nu \geqslant 1$. The Fourier transform $\mathrm{FT}_{\psi}(\mathcal{F})(1/2)$ is mixed of weight $\leqslant 0$ by the Riemann Hypothesis (in fact, it is known to be pure of weight 0, but this is again a deeper fact), hence Lemma 5.10 implies that $\mathcal{F}$ is geometrically irreducible if and only if the weight 0 part of $\mathrm{FT}_{\psi}(\mathcal{F})$ is geometrically irreducible.

We next consider quickly an example from the Polymath8 project, which we will reduce to a Fourier transform. We let

$$f = \frac{1}{X(X+Y)} + hY$$

where $h \in \mathbf{F}_q^{\times}$ is a parameter, and we wish to bound the conductor of

$$R^1 p_{1,!}\mathcal{L}_{\psi(f)},$$

i.e., the corresponding transform of the trivial sheaf, by a constant (independent of q).

We outline the steps that prove such a bound, leaving some details to the reader.

– It is equivalent to bound the conductor of

$$R^1 p_{1,!}\mathcal{L}_{\psi(g)}$$

where $g = (XY)^{-1} + hY - hX$ (applying the automorphism $(X, Y) \mapsto (X, X + Y)$). By the projection formula, we have

$$R^1 p_{1,!}\mathcal{L}_{\psi(g)} = \mathcal{L}_{\psi(-hX)} \otimes \mathcal{G},$$

where

$$\mathcal{G} = R^1 p_{1,!}\mathcal{L}_{\psi(h)}, \qquad h = \frac{1}{XY} + hY.$$

By Lemma 4.8, it is enough to estimate the conductor of $\mathcal{G}$.

– Note that the trace function of $\mathcal{G}$ is

$$\sum_{y \neq 0} \psi\Bigl(\frac{1}{xy} + hy\Bigr) = \sum_{v \neq 0} \psi\Bigl(\frac{1}{v} + \frac{h}{x}v\Bigr)$$

for $x \neq 0$, which is visibly a Kloosterman sum with parameter h/x. Let

$$\pi : \begin{cases} \mathbf{G}_m \times \mathbf{G}_m \longrightarrow \mathbf{G}_m \times \mathbf{G}_m \\ (x, y) \mapsto (hx^{-1}, xy) \end{cases}$$

and

$$\nu : \begin{cases} \mathbf{G}_m \longrightarrow \mathbf{G}_m \\ x \mapsto hx^{-1} \end{cases}$$

Then the analogue of the change of variable $(u, v) = \pi(x, y) = (h/x, xy)$ that establishes this identity is the isomorphism

$$\nu^* R^1 p_{1,!}\mathcal{L}_{\psi(V^{-1}+UV)} \simeq R^1 p_{1,!}\mathcal{L}_{\psi((XY)^{-1}+hY)}$$

of sheaves over the multiplicative group $\mathbf{G}_m = \mathbf{A}^1 - \{0\}$ over $\mathbf{F}_q$, which is a consequence of the isomorphism $\pi^*\mathcal{L}_{\psi(V^{-1}+UV)} \simeq \mathcal{L}_{\psi((XY)^{-1}+Y)}$.

Note that

$$R^1 p_{1,!}\mathcal{L}_{\psi(V^{-1}+UV)} = \mathrm{FT}_\psi(\mathcal{L}_{\psi(V^{-1})}),$$

which has bounded conductor independently of q. Since ν is an automorphism, and since the dimensions of the stalk of $\mathcal{G}$ at 0 is bounded, it follows from the fact that $\mathcal{G}$ coincides with $(\nu^{-1})^* \mathrm{FT}_\psi(\mathcal{L}_{\psi(V^{-1})})$ on $\mathbf{G}_m$ that the conductor of $\mathcal{G}$ is bounded for all q, as desired.

Remark 8.3. The Fourier transform of $\mathcal{L}_{\psi(X^{-1})}$ is the *Kloosterman sheaf* (in one variable), that was originally defined by Deligne. See [10] for its properties, and generalizations to more than one variable.

9. Beginning of the proof

We will now begin the general proof of Proposition 6.4. The complications which account for the length of the proof, compared with the case of the Fourier transform, are that the cohomology of the specializations $\mathcal{L}_{\psi(f(x,Y))}$ are not as simple as that of $\mathcal{L}_{\psi(xY)}$ (so that one can not compute $\mathrm{pct}(T_{\mathcal{K}}(\mathcal{F}))$ so easily as we showed that $\mathrm{pct}(\mathrm{FT}_\psi(\mathcal{F})) = 0$) and that, in general, the analogue of the formula (8.1) for the dimension of the fibers contains more terms depending on x. Nevertheless, the reader will clearly see that many concrete cases could be analyzed in much the same spirit as in Section 8.

We first deal with Points (1) and (2) in Proposition 6.4.

(1) We claim that $T^0_{\mathcal{K}}(\mathcal{F}) = 0$ for all f and $\mathcal{F}$. Indeed, by the proper base change theorem (Proposition 4.1, (4)), the stalk of $R^0 p_{1,!}(p_2^*\mathcal{F} \otimes \mathcal{K})$ over $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$ is

$$H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_{\psi(f(x,Y))}) = 0$$

by Lemma 4.2.

(2) By the bounds of Bombieri, Adolphson-Sperber and Katz (see, e.g., [14, Th. 12]), the sum of Betti numbers

$$\sum_{i=0}^{4} \dim H^i_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(f)})$$

is bounded by $(1 + \mathbf{c}(f))^B$ for some absolute constant $B \geqslant 1$, which proves the continuity of $h^i(f, \bar{\mathbf{Q}}_\ell)$. Precisely, in order to apply the result of Katz, one writes $f = f_1/f_2$ with $f_i \in \mathbf{F}_q[X,Y]$ and f_1 coprime to f_2, then one notes that if $U_2 \subset \mathbf{A}^2$ is the open subset where the denominator f_2 is invertible, we have

$$H^i_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(f)}) = H^i_c(U_2 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(f)})$$

by definition of cohomology with compact support. Define $Z \subset \mathbf{A}^3$, where $\mathbf{A}^3$ has coordinates (U, X, Y), to be the zero set of the polynomial $Uf_2(X,Y) - 1$. Then the morphism

$$\alpha \begin{cases} Z \longrightarrow U_2 \\ (u,x,y) \mapsto (x,y) \end{cases}$$

is an isomorphism such that $\alpha^*\mathcal{L}_{\psi(f)}$ is isomorphic to the lisse sheaf $\mathcal{L}_\psi(\tilde{f})$ for the polynomial $\tilde{f} = Uf_1(X,Y) \in \mathbf{F}_q[U,X,Y]$. Katz's theorem gives precisely the upper-bound

$$\sum_{i=0}^{4} \dim H^i_c(\bar{Z}, \mathcal{L}_{\psi(\tilde{f})}) \leqslant 3\Big(1 + 1 + \max(1 + \deg f_1, 1 + \deg f_2)\Big)^{3+1},$$

and hence the result.

The other parts of the proof are more involved, and require the tools of Section 7. However, before going further we will deal directly with the special case when $f \in \mathbf{F}_q(X) + \mathbf{F}_q(Y)$ (the reader is invited to figure out the analogue of Section 3 in this case).

So assume that

$$f = f_1 + f_2,$$

with $f_1 \in \mathbf{F}_q(X)$ and $f_2 \in \mathbf{F}_q(Y)$. We have $\mathcal{K} = p_1^*\mathcal{L}_1 \otimes p_2^*\mathcal{L}_2$, where $\mathcal{L}_i = \mathcal{L}_{\psi(f_i)}$, hence

$$R^i p_{1,!}(p_2^*\mathcal{F} \otimes p_2^*\mathcal{L}_2 \otimes p_1^*\mathcal{L}_1) \simeq \mathcal{L}_1 \otimes R^i p_{1,!}(p_2^*(\mathcal{F} \otimes \mathcal{L}_2)),$$

for $0 \leqslant i \leqslant 2$, by the projection formula (see, e.g., [8, Th. 7.4.7]).

But the sheaf $R^i p_{1,!}(p_2^*(\mathcal{F} \otimes \mathcal{L}_2))$ is the constant sheaf associated to $H^i_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_2)$: indeed, applying [2, Arcata, IV, Th. 5.4] to the cartesian diagram

$$\begin{array}{ccc} \mathbf{A}^1 & \xleftarrow{p_2} & \mathbf{A}^2 \\ s_2 \downarrow & & \downarrow p_1 \\ \operatorname{Spec} \mathbf{F}_q & \xleftarrow{s_1} & \mathbf{A}^1 \end{array}$$

and the sheaf $\mathcal{F} \otimes \mathcal{L}_2$ on $\mathbf{A}^1$, we obtain

$$s_1^* R^i s_{2,!}(\mathcal{F} \otimes \mathcal{L}_2) \simeq R^i p_{1,!}(p_2^*(\mathcal{F} \otimes \mathcal{L}_2)),$$

and the left-hand side is a constant sheaf (since it is pulled-back from $\mathbf{F}_q$) and has fiber $R^i s_{2,!}(\mathcal{F}\otimes\mathcal{L}_2) = H^i_c(\mathbf{A}^1\times\bar{\mathbf{F}}_q, \mathcal{F}\otimes\mathcal{L}_2)$, by the definition of cohomology with compact support and higher-direct images.

Hence we have (see (2.1))

$$c_i(f_1+f_2,\mathcal{F}) \leqslant (\dim H^i_c(\mathbf{A}^1\times\bar{\mathbf{F}}_q,\mathcal{F}\otimes\mathcal{L}_2))\times \mathbf{c}(\mathcal{L}_1),$$

which is continuous as a function of $\mathbf{c}(\mathcal{F})$ and $\mathbf{c}(f)$ by Lemmas 4.7 and 4.8.

This proves Theorem 2.3 in this special case, and Theorem 2.5 follows either from the argument in Section 15 (which is general) or from an application of the Künneth formula (see, e.g., [2, Sommes Trig., (2.4)*] and of Lemma 4.7.

10. Proof of (3) and (3bis)

We now prove (3) and (3bis) in Proposition 6.4. We first assume that the function $(f,\mathcal{F})\mapsto h^2(f,\mathcal{F})$ is continuous, and we consider $c_2(f,\mathcal{F}) = c(T^2_{\mathcal{K}}(\mathcal{F}))$. We can assume that $f\notin \mathbf{F}_q(X)+\mathbf{F}_q(Y)$, as we already treated the case when $f\in\mathbf{F}_q(X)+\mathbf{F}_q(Y)$ at the end of Section 6.

Lemma 10.1. *Assume $f\notin\mathbf{F}_q(X)+\mathbf{F}_q(Y)$. Then $T^2_{\mathcal{K}}(\mathcal{F})$ vanishes generically.*

Proof. Denote

$$\mathcal{G} = T^2_{\mathcal{K}}(\mathcal{F}) = R^2 p_{1,!}(p_2^*\mathcal{F}\otimes\mathcal{K}).$$

Let $(\mathcal{F}_i)$ be the (geometric) Jordan-Hölder factors of $\mathcal{F}$; then the geometric Jordan-Hölder factors of $p_2^*\mathcal{F}\otimes\mathcal{K}$ are the $p_2^*\mathcal{F}_i\otimes\mathcal{K}$, so that we may assume that $\mathcal{F}$ if geometrically irreducible.

Let $\eta = \mathrm{Spec}(\mathbf{F}_q(X))$ be the generic point of the affine line $\mathbf{A}^1_{\mathbf{F}_q}$ (with coordinate X), let $\bar\eta = \mathrm{Spec}(\overline{\mathbf{F}_q(X)})$ be a geometric point above η. By constructibility, the stalks of $\mathcal{G}$ vanish for all x in a dense open subset if and only if the stalk $\mathcal{G}_{\bar\eta}$ is zero.

By the proper base change theorem, we have

$$\mathcal{G}_{\bar\eta} = H^2_c(\mathbf{A}^1\times\overline{\mathbf{F}_q(X)},\mathcal{F}\otimes\mathcal{L}_{\psi(f_X(Y))}),$$

where $f_X(Y)=f(X,Y)$. Assume this stalk is non-zero. Then, using the coinvariant formula for the second cohomology group on a curve, it follows that there exists an open subset U of the affine line (with coordinate Y) over $\overline{\mathbf{F}_q(X)}$ such that

$$\mathcal{F}\simeq\mathcal{L}_{\psi(-f_X(Y))}$$

as sheaves on $U\times\overline{\mathbf{F}_q(X)}$. Since they are middle-extension sheaves, they are isomorphic as sheaves on the affine line over $\overline{\mathbf{F}_q(X)}$.

Note that $\mathcal{F}$ is pulled back from the affine line $\mathbf{A}^1$ over $\mathbf{F}_q$ (still with coordinate Y), and so the classification of Artin-Schreier sheaves shows that f is, up to an additive "constant" in $\overline{\mathbf{F}_q(X)}$, an element in $\mathbf{F}_q(Y)$, i.e., we have

$$f=f_1+f_2,$$

with $f_1\in\mathbf{F}_q(X)$ and $f_2\in\mathbf{F}_q(Y)$. ∎

Remark 10.2. One can also prove this lemma using more elementary arguments on rational functions, by looking at the vanishing at individual stalks and the classification of Artin-Schreier sheaves on $\mathbf{A}^1_{\mathbf{F}_q}$.

Because of this lemma, the conductor of $\mathcal{G} = T^1_{\mathcal{K}}(\mathcal{F})$ is equal to $\mathrm{pct}(\mathcal{G})$ (the generic rank is 0, and thus the action of all inertia groups on the generic fiber is trivial, which implies that $n(\mathcal{G}) = 0$ and hence the Swan conductors also vanish.) Hence

$$\mathbf{c}(\mathcal{G}) = \dim H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, T^2_{\mathcal{K}}(\mathcal{F})).$$

In the first spectral sequence of Lemma 7.3, with $U = \mathbf{A}^1$, we must therefore bound $\dim E_2^{0,2}$. By the last part of Proposition 7.2 (2), we have

$$\begin{aligned}(10.1)\quad \dim E_2^{0,2} \leqslant \dim E^2 + \dim E_2^{2,1} &= \dim H^2_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, p_2^*\mathcal{F} \otimes \mathcal{K}) + \\ \dim H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, T^1_{\mathcal{K}}(\mathcal{F})) &= h^2(f, \mathcal{F}) + \dim H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, T^1_{\mathcal{K}}(\mathcal{F})).\end{aligned}$$

We have already recalled in the proof of Lemma 4.7 that

$$\dim H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, T^1_{\mathcal{K}}(\mathcal{F})) \leqslant \mathrm{rank}(T^1_{\mathcal{K}}(\mathcal{F})).$$

Using the notation of Definition 4.4 and the proper base change theorem, we get

$$\dim H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, T^1_{\mathcal{K}}(\mathcal{F})) \leqslant \max_{x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)} \dim T^1_{\mathcal{K}}(\mathcal{F})_x \leqslant \max_{x \in \bar{\mathbf{F}}_q} \dim H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x),$$

and by Corollary 4.9, this shows that $(f, \mathcal{F}) \mapsto \dim H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, T^1_{\mathcal{K}}(\mathcal{F}))$ is continuous. The inequality (10.1) then finishes the proof of (3).

In the case of (3bis), we observe that if $\mathcal{F} = \bar{\mathbf{Q}}_\ell$ is trivial, then the above still shows that $c_2(f, \bar{\mathbf{Q}}_\ell)$ is bounded polynomially in terms of $h^2(f, \bar{\mathbf{Q}}_\ell)$, proving (3bis).

11. Proof of (4)

We now prove (4) in Proposition 6.4. Thus we assume that the function $f \mapsto c_i(f, \bar{\mathbf{Q}}_\ell)$ are continuous for $i = 1$ and $i = 2$, and we consider $h^2(f, \mathcal{F})$. Again we may assume that $f \notin \mathbf{F}_q(X) + \mathbf{F}_q(Y)$.

We apply the second spectral sequence of Lemma 7.3, with $V = \mathbf{A}^1$, and the first part of Proposition 7.2, (2) with $n = 2$: this gives

$$h^2(f, \mathcal{F}) = \dim E^2 \leqslant \dim E_2^{2,0} + \dim E_2^{1,1} + \dim E_2^{0,2},$$

where

$$E_2^{i,j} = H^i_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes T^j_{\mathcal{K}^*}(\bar{\mathbf{Q}}_\ell)).$$

We note that $\mathbf{c}(\mathcal{K}^*) = \mathbf{c}(\mathcal{K})$. We have $E_2^{2,0} = 0$, and

$$\dim E_2^{1,1} = \dim H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes T^1_{\mathcal{K}^*}(\bar{\mathbf{Q}}_\ell))$$

is continuous by Lemma 4.7 and 4.8, since the conductor of $T^1_{\mathcal{K}^*}(\bar{\mathbf{Q}}_\ell)$ is bounded polynomially in terms of the conductor of f by assumption.

Finally, we have

$$\dim E_2^{0,2} = \dim H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes T^2_{\mathcal{K}^*}(\bar{\mathbf{Q}}_\ell)) \leqslant \mathbf{c}(\mathcal{F} \otimes T^2_{\mathcal{K}^*}(\bar{\mathbf{Q}}_\ell)).$$

By assumption, $f \mapsto c_2(f^*, \bar{\mathbf{Q}}_\ell)$ is continuous, and therefore the function $\dim E_2^{0,2}$ is continuous (Lemma 4.7). Thus $h^2(f, \mathcal{F})$ is also continuous.

12. Further lemmas

We now come to some (mostly) elementary lemmas about rational functions in $\mathbf{F}_q(X,Y)$. These will be the main ingredients in the next section when estimating $m(f,\mathcal{F})$. Although the proofs are a bit lengthy, it should be noted that the statements can often be checked very quickly in concrete special cases.

Below, for a rational function $f \in \mathbf{F}_q(X,Y)$, we will write

$$f = g_1/g_2$$

where $g_i \in \mathbf{F}_q[X,Y]$ are polynomials and g_1 and g_2 are coprime. We will call g_1 (resp. g_2) the numerator (resp. denominator) of f, although they are only unique up to multiplication by a non-zero element $\alpha \in \mathbf{F}_q^\times$. We will sometimes write

$$d_i = \deg_Y g_i, \qquad e_i = \deg_X g_i.$$

We will consider open subsets U associated to f. We will say that such an open set $U \subset \mathbf{A}^1$ has *bounded complement* if $(\mathbf{A}^1 - U)(\bar{\mathbf{F}}_q)$ is of finite order such that

$$|(\mathbf{A}^1 - U)(\bar{\mathbf{F}}_q)| \leqslant (2\,\mathbf{c}(f))^C$$

for some absolute constant $C \geqslant 1$, i.e., if the map $f \mapsto |(\mathbf{A}^1 - U)(\bar{\mathbf{F}}_q)|$ is continuous. Thus, for instance, the set of x where f_x is defined has bounded complement. Below, when the data also involves a sheaf $\mathcal{F}$, an open set will be said to have bounded complement if the size of the complement is a continuous function of $(f,\mathcal{F})$.

Lemma 12.1. *Let $\mathbf{F}_q$ be a finite field of order q and characteristic p, and let $f \in \mathbf{F}_q(X,Y)$ be a non-constant rational function.*

(1) *There exists a finite subset $C \subset \mathbf{P}^1$ such that $y \in C$ if and only if y is a pole of all but finitely many specializations f_x; the size of C is bounded in terms of $\mathbf{c}(f)$ only, and for $y \in C$, the set of $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$ for which y is a pole of f_x has bounded complement.*

(2) *There exists a dense open subset $U \subset \mathbf{A}^1$ with bounded complement such that, for $x \in U(\bar{\mathbf{F}}_q)$, the sum*

$$\sum_{y \in (\mathbf{A}^1 - C)(\bar{\mathbf{F}}_q)} \operatorname{ord}_y f_x$$

of the order of the poles of f at $y \notin C$ is constant. This open set has bounded complement, and the value of this constant is a continuous function of f.

(3) *Similarly, there exists a dense open subset $V \subset \mathbf{A}^1$ with bounded complement such that, for $x \in V(\bar{\mathbf{F}}_q)$, the number of poles in $\mathbf{A}^1$ of f_x is constant. This number is a continuous function of f.*

Proof. (1) The set C is the set of those y such that $Y - y$ divides g_2; in particular $|C| \leqslant d_2$. If $y \in C$, then it is a pole of f_x unless $g_1(x,Y)$ has a zero at y of order at least equal to that of $g_2(x,Y)$ at y. The number of x where this last property holds is bounded by $\mathbf{c}(f)$, because $g_1(X,y) \neq 0$ (since $Y - y$ does not divide g_1).

(2) If $g_1(x,Y)$ and $g_2(x,Y)$ are coprime, we have

$$\sum_{y \in \bar{\mathbf{F}}_q - C} \operatorname{ord}_y f_x + \sum_{y \in C} \operatorname{ord}_y f_x = \deg g_2(x,Y),$$

the number of poles of f_x in $\mathbf{A}^1$, counted with multiplicity. For x outside of the set Z of zeros of the leading term $a(X)$ of g_2, as a polynomial in Y, the degree $\deg g_2(x,Y)$ is equal to d_2. We have $|Z| \leqslant \deg(a) \leqslant e_2$. Furthermore, $g_1(x, Y)$ and $g_2(x, Y)$ are coprime unless x is a zero of the resultant

$$r = \mathrm{Res}(g_1, g_2) \in \mathbf{F}_q[X]$$

(where g_1, g_2 are viewed as elements of $\mathbf{F}_q[X][Y]$). Since g_1 and g_2 are coprime, the polynomial r is non-zero. Furthermore, its degree is bounded in terms of $\mathbf{c}(f)$. Thus, for x outside a set of size bounded in terms of $\mathbf{c}(f)$, we have

$$\sum_{y \in \bar{\mathbf{F}}_q - C} \mathrm{ord}_y f_x + \sum_{y \in C} \mathrm{ord}_y f_x = d_2.$$

From (the proof of) (1), we have

$$\sum_{y \in C} \mathrm{ord}_y f_x = \sum_{y \in C} v_{Y-y}(g_2),$$

for all $x \in \bar{\mathbf{F}}_q$ with a number of exceptions depending only on $\mathbf{c}(f)$, where $v_{Y-y}(\cdot)$ is the valuation corresponding to the irreducible polynomial $Y - y \in \mathbf{F}_q[X, Y]$. So for all x in an open set with bounded complement, we have

$$\sum_{y \in \bar{\mathbf{F}}_q - C} \mathrm{ord}_y f_x = d_2 - \sum_{y \in C} v_{Y-y}(g_2),$$

a constant which is $\leqslant d_2 \leqslant \mathbf{c}(f)$.

(3) The proof is similar to that of (2), except that we must count poles without multiplicity. From the open set U in (2), we start by removing the $x \in \bar{\mathbf{F}}_q$ where two different irreducible factors of g_2, neither of which is of the form $Y - y$, have a common root; since such factors are coprime, arguing in terms of resultant again shows that we obtain a dense open subset U_1 with bounded complement. For x in $U_1(\bar{\mathbf{F}}_q)$, the number of poles of f_x is the sum of the number of poles coming from the factors of g_2. We then show that each factor has a constant contribution.

Consider then a fixed factor $g^k \in \mathbf{F}_q[X, Y]$ of g_2, where $k \geqslant 1$ and g is irreducible and different from $Y - y$ (up to a scalar factor). For $x \in U_1(\bar{\mathbf{F}}_q)$ (so that, in particular, $g(x, Y)$ is coprime with $g_1(x, Y)$) the sum of the multiplicities of the poles arising from g is equal to $k \deg_Y g$. If g is separable as a polynomial in $\mathbf{F}_q[X][Y]$, so that g is coprime with g' (where the derivative is taken with respect to Y), then the number of poles coming from g is $\deg_Y g$ for all $x \in U_1(\bar{\mathbf{F}}_q)$, except the roots of the resultant of g and g', which form a set of size bounded in terms of $\mathbf{c}(f)$. If g is not separable,[3] so that $g = \tilde{g}(Y^{p^\nu})$ for some $\nu \geqslant 1$ and some separable polynomial $\tilde{g} \in \mathbf{F}_q[X](Y)$, then outside of the roots of the resultant of $\tilde{g}$ and $\tilde{g}'$, the number of poles coming from g is $\deg \tilde{g}$. □

Example 12.2. We illustrate the last two points.

(1) Consider

$$f = \frac{1}{(XY^2 + 4Y + X)(X^2Y - X - 6)}$$

for $p \neq 2, 3$.

Then the sum of the multiplicities of the poles of f_x is equal to 3, except for $x = 0$. On the other hand, the number of poles is 3 except for $x \in \{0, 2, -2\}$; it is equal to 1 for $x = 0$,

[3] Note that this case can only happen if $\mathbf{c}(f) \geqslant p$, which we will exclude for the purpose of Theorem 2.3 anyway.

to 2 for $x = 2$ (where the first factor $XY^2 + 4Y + X$ has a double zero), and to 1 for $x = -2$ (when the double zero of $-2Y^2 + 4Y - 2$ at $Y = 1$ coincides with the zero of $4Y - 4$).

(2) Consider

$$f = \frac{1}{Y^p - X} \in \mathbf{F}_q(X, Y).$$

In that case, there is a unique pole of f_x, which always has multiplicity p, for every $x \in \bar{\mathbf{F}}_q$.

In the next lemma, we show that various invariants of $\mathcal{L}_{\psi(f)}$ are constant on open sets with bounded complement. Some steps may be compared with those in [10, §8.2–8.5], where the case of the Fourier transform is dealt with (among deeper properties). It is the only step where we need to assume that $\mathbf{c}(f) < p$.

Lemma 12.3. *Let $\mathbf{F}_q$ be a finite field of order q and characteristic p, ℓ a prime distinct from p. Let $f \in \mathbf{F}_q(X, Y)$ and let $\mathcal{K} = \mathcal{L}_{\psi(f(X,Y))}$, where ψ is a non-trivial additive ℓ-adic character. Assume $\mathbf{c}(f) < p$. Further, let $\mathcal{F}$ be a middle-extension sheaf on $\mathbf{A}^1_{\bar{\mathbf{F}}_q}$.*

(1) *There exists a dense open set $V_1 \subset \mathbf{A}^1$ with bounded complement such that*

$$\sum_{y \in \bar{\mathbf{F}}_q} \mathrm{drop}_y(\mathcal{F} \otimes \mathcal{L}_x),$$

is constant for $x \in V_1(\bar{\mathbf{F}}_q)$.

(2) *There exists a dense open set $V_2 \subset \mathbf{A}^1$ with bounded complement such that*

$$\sum_{y \in \bar{\mathbf{F}}_q} \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x)$$

is constant for $x \in V_2(\bar{\mathbf{F}}_q)$.

(3) *There exists a dense open set $V_3 \subset \mathbf{A}^1$ with bounded complement such that*

$$\mathrm{Swan}_\infty(\mathcal{F} \otimes \mathcal{L}_x)$$

is constant for $x \in V_3(\bar{\mathbf{F}}_q)$.

In the proof, since we now have input data $(f, \mathcal{F})$, recall an open set $U \subset \mathbf{A}^1$ associated to $(f, \mathcal{F})$ has bounded complement if $|(\mathbf{A}^1 - U)(\bar{\mathbf{F}}_q)|$ is a continuous function of $(f, \mathcal{F})$.

Proof. From the description of $\mathcal{L}_x$ in Lemma 4.5 and the remark that follows, we may restrict our attention to the open set of those $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$ where $\mathcal{L}_x$ is isomorphic to the Artin-Schreier sheaf $\mathcal{L}_{\psi(f_x)}$ on $\mathbf{A}^1_{\bar{\mathbf{F}}_q}$ (since the set of $x \in \bar{\mathbf{F}}_q$ where this does not hold is contained in the finite set of size $\leqslant \mathbf{c}(f)$ where either $X - x$ divides the denominator of f or for which there exists y such that (x, y) is a common zero of the numerator and denominator.) We assume that all x in the arguments below satisfy this property.

(1) The stalk of $\mathcal{F} \otimes \mathcal{L}_x$ at $y \in \bar{\mathbf{F}}_q$ is $\mathcal{F}_y \otimes \mathcal{L}_{x,y}$, and hence it is either of dimension $\dim \mathcal{F}_y$ if y is not a singularity of $\mathcal{L}_x$, or it vanishes. Denoting by L_x (resp. by C_1) the set of singularities of $\mathcal{L}_x$ (resp. of $\mathcal{F}$) in $\mathbf{A}^1$, the definition (2.2) gives

$$\sum_{y \in \bar{\mathbf{F}}_q} \mathrm{drop}_y(\mathcal{F} \otimes \mathcal{L}_x) = \sum_{y \in L_x} \mathrm{rank}(\mathcal{F}) + \sum_{y \in C_1 - L_x} (\mathrm{rank}(\mathcal{F}) - \dim \mathcal{F}_y).$$

Let $C \subset \mathbf{A}^1$ be the set given by Lemma 12.1, (1). For x in an open set with bounded complement, the set $C \cap C_1$ is the set of common singularities of $\mathcal{L}_x$ and $\mathcal{F}$, so that $C_1 - L_x =$

$C_1 - (C \cap C_1)$ is then independent of x, and therefore so is the sum

$$\sum_{y \in C_1 - L_x} (\mathrm{rank}(\mathcal{F}) - \dim \mathcal{F}_y).$$

Finally, the size of L_x is also constant outside of an open set with bounded complement by Lemma 12.1, (3), and hence we get (1).

(2) Let again $C_1 \subset \mathbf{A}^1$ be the set of singularities of $\mathcal{F}$, and let $C \subset \mathbf{A}^1$ be the set given by Lemma 12.1, (1). Thus $C \cap C_1$ is, for all x in an open set with bounded complement, a set of order bounded in terms of $\mathbf{c}(f)$ where $\mathcal{F}$ and $\mathcal{L}_x$ have a common singularity. Let $S_x \subset \mathbf{A}^1$ be the set of singularities in $\mathbf{A}^1$ of $\mathcal{F} \otimes \mathcal{L}_x$. We then have a disjoint union

$$S_x = (C_1 - C) \cup \tilde{L}_x \cup T_x,$$

for all x in an open set with bounded complement, where $\tilde{L}_x$ is the set of singularities of $\mathcal{L}_x$ outside of C (which are not singularities of $\mathcal{F}$) and T_x is contained in $C \cap C_1$ (some points in $C \cap C_1$ might not be actual singularities.) In particular

$$\sum_{y \in \bar{\mathbf{F}}_q} \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x) = \sum_{y \in C_1 - C} \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x) + \sum_{y \in \tilde{L}_x} \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x) + \sum_{y \in C \cap C_1} \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x),$$

(where we are allowed to replace T_x by the larger set $C \cap C_1$ since the swan conductors are zero by definition in the complement outside of T_x.) It is enough then to show that each of the three terms in this sum is constant in some open set with bounded complement.

We have

$$\mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x) = \mathrm{Swan}_y(\mathcal{F}) \text{ if } y \in C_1 - C$$
$$\mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x) = \mathrm{rank}(\mathcal{F})\, \mathrm{Swan}_y(\mathcal{L}_x) \text{ if } y \in \tilde{L}_x.$$

The first formula immediately implies that

$$\sum_{y \in C_1 - C} \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x)$$

is constant on an open set with bounded complement. The second gives

$$\sum_{y \in \tilde{L}_x} \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x) = \mathrm{rank}(\mathcal{F}) \sum_{y \in \tilde{L}_x} \mathrm{Swan}_y(\mathcal{L}_x) = \mathrm{rank}(\mathcal{F}) \sum_{y \in \tilde{L}_x} \mathrm{ord}_y f_x$$

by Lemma 4.3, and then Lemma 12.1, (2) shows that this expression is constant in an open set with bounded complement.

To handle the last component, we claim that, for each $y \in C_1 \cap C$, there exists an open set U_y with bounded complement such that

$$\mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x)$$

is constant for $x \in U_y(\bar{\mathbf{F}}_q)$. Since $C \cap C_1$ has bounded order, the intersection of the U_y for $y \in C \cap C_1$ is an open set with bounded complement where the last sum is constant. And the claim is, up to changing y to ∞, identical with part (3) which we will now prove.

(3) We consider now the issue of the variation of $\mathrm{Swan}_\infty(\mathcal{F}\otimes\mathcal{L}_x)$, arguing as in [10, Lemma 8.5.4, Cor. 8.5.5].[4]

The Swan conductor $\mathrm{Swan}_\infty \mathcal{L}_x$ is constant outside of an open set U with bounded complement (see Lemma 4.3. Let $d \geqslant 0$ be this constant value, which is also the unique break of $\mathcal{L}_x$ as a representation of the inertia group I_∞ (since $\mathcal{L}_x$ has rank 1). We may assume $d \geqslant 1$, since otherwise $\mathcal{L}_x$ is unramified at ∞ for all $x \in U(\bar{\mathbf{F}}_q)$, and $\mathrm{Swan}_x(\mathcal{F}\otimes\mathcal{L}_x) = \mathrm{Swan}_x(\mathcal{F})$ is constant on U.

Now let M be an irreducible component of $\mathcal{F}$, seen as a representation of the wild inertia group at ∞. This has a unique break $\lambda \geqslant 0$. If $\lambda \neq d$, then

$$\mathrm{Swan}_\infty(M\otimes\mathcal{L}_x) = \max(d,\lambda)\dim M,$$

which is independent of $x \in U(\bar{\mathbf{F}}_q)$, and we are also done.

Now assume $\lambda = d$. We have $\mathrm{Swan}_\infty(M\otimes\mathcal{L}_x) \leqslant d\dim M$. Assume there exists one $x_0 \in U(\bar{\mathbf{F}}_q)$ such that $\mathrm{Swan}_\infty(M\otimes\mathcal{L}_{x_0}) < d\dim M$, or equivalently such that the unique (because of irreducibility) break of $M\otimes\mathcal{L}_{x_0}$ is $< d$. We will show that this strict inequality happens *at most* for a bounded number of other values of $x \in U(\bar{\mathbf{F}}_q)$.

Indeed, we can write

$$f(X,Y) = a_d(X)Y^d + a_{d-1}(X)Y^{d-1} + \cdots + a_0(X) + g(X,Y) = a_d(X)Y^d + f_1(X,Y),$$

where $a_j \in \mathbf{F}_q[X]$, $a_d(x) \neq 0$ for $x \in U$ and $g(X,Y) \in \mathbf{F}_q(X,Y)$ is defined at infinity for $x \in U$. In fact, we may assume that a_d is not a constant polynomial, since if it were, we could replace $\mathcal{F}$ by $\mathcal{F}\otimes\mathcal{L}_{\psi(a_dY^d)}$ and correspondingly d by $d-1$, thus concluding by induction.

Now, for any $x \in U(\bar{\mathbf{F}}_q)$, we write

$$M\otimes\mathcal{L}_x = (M\otimes\mathcal{L}_{x_0})\otimes\mathcal{L}_{\psi(f_x-f_{x_0})}.$$

By assumption, the first factor has unique break $< d$, while on the other hand, the second has unique break equal to d for all x such that

$$a_d(x) \neq a_d(x_0).$$

Since a_d is not constant, this fails to happen for at most $\deg(a_d) \leqslant \mathbf{c}(f)$ values of x. Now applying these bounds to all irreducible components M (their number is at most the rank of $\mathcal{F}$, hence $\leqslant \mathbf{c}(\mathcal{F})$) we finish the proof of (3). ☐

And finally we will use a cohomological computation:

Lemma 12.4. *Let $\mathbf{F}_q$ be a finite field of characteristic p and $\ell \neq p$ a prime number. Let $f \in \mathbf{F}_q(X,Y)$ be a rational function and*

$$\mathcal{K} = \mathcal{L}_{\psi(f)}$$

where ψ is a non-trivial ℓ-adic additive character. Let $\mathcal{F}$ be a constructible ℓ-adic sheaf on $\mathbf{A}^1_{\mathbf{F}_q}$. Write $f = g_1/g_2$ with $g_i \in \mathbf{F}_q[X,Y]$ coprime.

Let C be the union of the zero set of g_2, seen as a reduced subscheme of $\mathbf{A}^2$, and of the lines

$$\mathbf{A}^1 \times \{y\} \subset \mathbf{A}^2,$$

where $y \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$ is a singularity of $\mathcal{F}$. Let $W \subset \mathbf{A}^2$ be the open subset complement of C.

[4] Note that if $\mathcal{F}$ is at most tamely ramified at ∞, or if $\mathcal{L}_x$ is generically unramified at ∞ (e.g., if $f = g_1/g_2$ with $\deg_X g_2 > \deg_X g_1$), then $\mathrm{Swan}_\infty(\mathcal{F}\otimes\mathcal{L}_x)$ is clearly constant on a dense open set with bounded complement, so that we are already done in such a case.

(1) *We have*
$$H^1_c(\bar{W}, p_2^*\mathcal{F} \otimes \mathcal{K}) = 0.$$
(2) *The map*
$$(f, \mathcal{F}) \mapsto \dim H^1_c(\bar{C}, p_2^*\mathcal{F} \otimes \mathcal{K})$$
is continuous.

Proof. (1) The open subset W is a smooth affine surface, and $p_2^*\mathcal{F} \otimes \mathcal{K}$ is lisse on W, so (4.4) gives vanishing of the first cohomology group.

(2) Write C_1 for the zero set of g_2 (as a reduced scheme) and
$$C_2 = \bigcup_{y \in \tilde{S}} \mathbf{A}^1 \times \{y\},$$
where y ranges over those singularities of $\mathcal{F}$ in $\mathbf{A}^1(\bar{\mathbf{F}}_q)$ such that $\mathbf{A}^1 \times \{y\}$ is *not* contained in C_1.

Let $S = C_1 \cap C_2$ be the intersection of these two sets; because of the last restriction, this is a finite set, and its order is bounded polynomially in terms of $\mathbf{c}(\mathcal{F})$ and $\mathbf{c}(f)$ (e.g., by Bezout's Theorem for plane curves). Applying the excision exact sequence (4.1) to C and the complement U (in C) of the closed set C_1, we get by (4.2) the bound
$$\dim H^1_c(\bar{C}, p_2^*\mathcal{F} \otimes \mathcal{K}) \leqslant \dim H^1_c(\bar{U}, p_2^*\mathcal{F} \otimes \mathcal{K}) + \dim H^1_c(\bar{C}_1, p_2^*\mathcal{F} \otimes \mathcal{K}) = \dim H^1_c(\bar{U}, p_2^*\mathcal{F} \otimes \mathcal{K})$$
since $\mathcal{K}$, by definition, is zero on C_1.

We have $U = C_2 - S$, and we apply again the excision exact sequence to C_2 and its open set U, obtaining by (4.3) the bound
$$\dim H^1_c(\bar{C}, p_2^*\mathcal{F} \otimes \mathcal{K}) \leqslant \dim H^1_c(\bar{U}, p_2^*\mathcal{F} \otimes \mathcal{K}) \leqslant \dim H^1_c(\bar{C}_2, p_2^*\mathcal{F} \otimes \mathcal{K})$$
(because $H^0_c(\bar{S}, p_2^*\mathcal{F} \otimes \mathcal{K}) = 0$ since $S \subset C_1$, so that $\mathcal{K}$ is zero on S). Finally, we have
$$H^1_c(\bar{C}_2, p_2^*\mathcal{F} \otimes \mathcal{K}) = \bigoplus_{y \in \tilde{S}} \mathcal{F}_y \otimes H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{L}_{\psi(f(x,Y))}),$$
and moreover both $\tilde{S}$ and S have order bounded in terms of $\mathbf{c}(\mathcal{F})$ and $\mathbf{c}(f)$, so that we obtain the result. ☐

13. Proof of (5) and (5bis)

We now prove (5) and (5bis) in Proposition 6.4. Thus we assume that the function $(f, \mathcal{F}) \mapsto c_2(f, \mathcal{F})$ is continuous, and we consider $m(f, \mathcal{F})$. We still assume that $f \notin \mathbf{F}_q(X) + \mathbf{F}_q(Y)$.

We estimate first the rank of $T^1_{\mathcal{K}}(\mathcal{F})$, then the punctual part $\mathrm{pct}(T^1_{\mathcal{K}}(\mathcal{F}))$, and finally the number of singularities.

We will use some of the notation introduced in the previous section. We write
$$\mathcal{G} = T^1_{\mathcal{K}}(\mathcal{F}) = R^1 p_{1,!}(p_2^*\mathcal{F} \otimes \mathcal{K}).$$
For $x \in \mathbf{A}^1(\bar{\mathbf{F}}_q)$, the stalk of $\mathcal{G}$ at x is
$$\mathcal{G}_x = H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x)$$
by the proper base change theorem.

The generic rank of $\mathcal{G}$ is at most the maximal value of the dimension of this stalk. Hence, by Corollary 4.9, it is a continuous function of $(f, \mathcal{F})$.

Next, we have

$$\mathrm{pct}(\mathcal{G}) = \dim H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{G})$$

Applying Proposition 7.2, (3) to the first spectral sequence of Lemma 7.3 (with $U = \mathbf{A}^1$), we obtain

$$\mathrm{pct}(\mathcal{G}) = \dim H^1_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, p_2^*\mathcal{F} \otimes \mathcal{K}).$$

Let $W \subset \mathbf{A}^2 \times \bar{\mathbf{F}}_q$ be the (dense) open set defined in Lemma 12.4 and $C = \mathbf{A}^2 - W$ its complement. From the excision inequality (4.2) and $H^1_c(\bar{W}, p_2^*\mathcal{F} \otimes \mathcal{K}) = 0$ (by Lemma 12.4), we get

$$\mathrm{pct}(\mathcal{G}) = \dim H^1_c(\mathbf{A}^2 \times \bar{\mathbf{F}}_q, p_2^*\mathcal{F} \otimes \mathcal{K}) \leqslant \dim H^1_c(\bar{C}, p_2^*\mathcal{F} \otimes \mathcal{K}),$$

and the second part of Lemma 12.4 shows that $(f, \mathcal{F}) \mapsto \mathrm{pct}(\mathcal{G})$ is continuous.

There only remains to estimate the number of singularities of $\mathcal{G}$. We use the criterion of Lemma 4.10. Since we already have an estimate for $\mathrm{pct}(\mathcal{G})$, this shows that it is enough to prove that the dimension of the stalk of $\mathcal{G}$ is constant on an open set with bounded complement, by which we mean as before an open set with complement of order bounded in terms of the conductor of f and of the conductor of $\mathcal{F}$ only.

Since $\mathcal{F}$ is a middle-extension sheaf and since $\mathcal{L}_x \simeq \mathcal{L}_{\psi(f(x,Y))}$ is also one for all x in an open set with bounded complement (Lemma 4.5), we have

$$H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x) = 0$$

for all x in an open set with bounded complement by Lemma 4.2. From Lemma 10.1, and from the fact that $\mathrm{pct}(T^2_{\mathcal{K}}(\mathcal{F})) = c_2(f, \mathcal{F})$ is continuous, there exists therefore an open set U_1 with bounded complement such that

$$H^0_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x) = H^2_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x) = 0$$

for $x \in U_1(\bar{\mathbf{F}}_q)$.

By the Euler-Poincaré formula, we have

$$\dim \mathcal{G}_x = \sum_{y \in \bar{\mathbf{F}}_q} (\mathrm{drop}_y(\mathcal{F} \otimes \mathcal{L}_x) + \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x)) + \mathrm{Swan}_\infty(\mathcal{F} \otimes \mathcal{L}_x) - \mathrm{rank}(\mathcal{F})$$

for $x \in U_1(\bar{\mathbf{F}}_q)$, where all but finitely many terms are zero (see [10, 8.5.2, 8.5.3]).

We write

$$\dim \mathcal{G}_x = N_1(x) + N_2(x) + N_3(x) - \mathrm{rank}(\mathcal{F})$$

where

$$N_1(x) = \sum_y \mathrm{drop}_y(\mathcal{F} \otimes \mathcal{L}_x), \qquad N_2(x) = \sum_y \mathrm{Swan}_y(\mathcal{F} \otimes \mathcal{L}_x),$$
$$N_3(x) = \mathrm{Swan}_\infty(\mathcal{F} \otimes \mathcal{L}_x).$$

By Lemma 12.3, for each $i = 1, 2, 3$, there exists an open set $V_i \subset U_1$ with bounded complement such that $N_i(x)$ is constant on V_i. It follows that $\dim \mathcal{G}_x$ is constant on the open set with bounded complement $U_1 \cap V_1 \cap V_2 \cap V_3$, and by the criterion of Lemma 4.10, that the number of singularities of $\mathcal{G}$ is bounded polynomially in terms of $\mathbf{c}(\mathcal{F})$ and $\mathbf{c}(\mathcal{K})$.

This argument proves (5), but the reader can now check that (5bis) is also a consequence.

14. Proof of (6) and (6bis)

We now prove (6) and (6bis) in Proposition 6.4. Thus we assume that the functions $(f, \mathcal{F}) \mapsto m(f, \mathcal{F})$ and $(f, \mathcal{F}) \mapsto h^2(f, \mathcal{F})$ are continuous, and we consider $c_1(f, \mathcal{F})$. We continue assuming that $f \notin \mathbf{F}_q(X) + \mathbf{F}_q(Y)$.

Let U be the maximal dense open set on which $T^1_{\mathcal{K}}(\mathcal{F})$ is lisse. From Lemma 4.11, and the assumption that $m(f, \mathcal{F})$ is continuous, we see that it is enough to prove that

$$\dim H^1_c(\bar{U}, T^1_{\mathcal{K}}(\mathcal{G}))$$

is continuous.

We use the first spectral sequence of Lemma 7.3 (with the open set U). By Proposition 7.2 (1) and (7.4), we have

$$E^2 \simeq E_3^{0,2} \oplus E_2^{1,1} \oplus E_3^{2,0}$$

and in particular

$$\dim H^1_c(\bar{U}, T^1_{\mathcal{K}}(\mathcal{F})) = \dim E_2^{1,1} \leqslant \dim E^2 = \dim H^2_c(\bar{U} \times \mathbf{A}^1, p_2^*\mathcal{F} \otimes \mathcal{K}).$$

To conclude, we relate the cohomology of $\bar{U} \times \mathbf{A}^1$ with that of $\mathbf{A}^2$, which will allow us the use the continuity of $h^2(f, \mathcal{F})$.

Let $W = U \times \mathbf{A}^1$ and $C = \mathbf{A}^2 - (U \times \mathbf{A}^1)$. By the excision inequality (4.3), we get

$$\dim H^2_c(\bar{W}, p_2^*\mathcal{F} \otimes \mathcal{K}) \leqslant h^2(f, \mathcal{F}) + \dim H^1_c(\bar{C}, p_2^*\mathcal{F} \otimes \mathcal{K})$$

We have

$$C = \bigcup_{x \in (\mathbf{A}^1 - U)(\bar{\mathbf{F}}_q)} \{x\} \times \mathbf{A}^1,$$

where the union is disjoint, and hence

$$H^1_c(\bar{C}, p_2^*\mathcal{F} \otimes \mathcal{K}) = \bigoplus_{x \in (\mathbf{A}^1 - U)(\bar{\mathbf{F}}_q)} H^1_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, \mathcal{F} \otimes \mathcal{L}_x).$$

The number of summands is at most $m(f, \mathcal{F})$, and the dimension of each summand is continuous, again by Corollary 4.9. Since $h^2(f, \mathcal{F})$ is also assumed to be continuous, this proves (6). The case of (6bis) is then clear by inspection.

15. Proof of (7)

We now prove (7) in Proposition 6.4. Thus we assume that the functions $(f, \mathcal{F}) \mapsto c_i(f, \mathcal{F})$ are continuous for $0 \leqslant i \leqslant 2$, and we consider $h^j(f, \mathcal{F})$. (There is no need here to assume that $f \notin \mathbf{F}_q(X) + \mathbf{F}_q(Y)$.)

We use the first spectral sequence of Lemma 7.3 with $U = \mathbf{A}^1$. For any j, it implies

$$h^j(f, \mathcal{F}) = \dim E^j \leqslant \sum_{p=0}^{j} \dim H^p_c(\mathbf{A}^1 \times \bar{\mathbf{F}}_q, T^{j-p}_{\mathcal{K}}(\mathcal{F})).$$

By Lemma 4.7, and the continuity of $c_{j-p}(f, \mathcal{F})$, each term in the sum is a continuous function, and hence so is $h^j(f, \mathcal{F})$.

References

[1] J.B. Conrey and H. Iwaniec: *The cubic moment of central values of automorphic L-functions*, Ann. of Math. (2) 151 (2000), no. 3, 1175–1216.

[2] P. Deligne: *Cohomologie étale*, S.G.A $4\frac{1}{2}$, L.N.M 569, Springer Verlag (1977).

[3] P. Deligne: *La conjecture de Weil, II*, Publ. Math. IHÉS 52 (1980), 137–252.

[4] É. Fouvry, E. Kowalski, Ph. Michel: *Algebraic twists of modular forms and Hecke orbits* Preprint (2012).

[5] É. Fouvry, E. Kowalski, Ph. Michel: *Algebraic trace weights over the primes* Preprint (2012).

[6] É. Fouvry, E. Kowalski, Ph. Michel: *On the exponent of distribution of the ternary divisor function* Preprint (2013).

[7] É. Fouvry and Ph. Michel: *Sur certaines sommes d'exponentielles sur les nombres premiers*, Ann. Sci. École Norm. Sup. (4) 31 (1998), 93–130.

[8] L. Fu: *Etale cohomology theory*, Nankai Tracts in Math. 13, World Scientific 2011.

[9] N.M. Katz: *Sommes exponentielles*, Astérisque 79 (1980).

[10] N.M. Katz: *Gauss sums, Kloosterman sums and monodromy groups*, Annals of Math. Studies 116, Princeton Univ. Press (1988).

[11] N.M. Katz: *Exponential sums and differential equations*, Annals of Math. Studies 124, Princeton Univ. Press (1990).

[12] N.M. Katz: *Rigid local systems*, Annals of Math. Studies 139, Princeton Univ. Press (1993).

[13] N.M. Katz: *Affine cohomological transforms*, Journal of the A.M.S 6 (1993), 149–222.

[14] N.M Katz: *Sums of Betti numbers in arbitrary characteristic*, Finite Fields Appl. 7 (2001), no. 1, 29–44.

[15] N.M. Katz: *L-functions and monodromy: four lectures on Weil II*, Advanced in Math. 160 (2001), 81–132.

[16] G. Laumon: *Transformation de Fourier, constantes d'équations fonctionnelles et conjecture de Weil*, Publ. Math. IHÉS, 65 (1987), 131–210.

[17] J. Milne: *Étale cohomology*, Princeton Univ. Press, 1990.

[18] I. Petrow and M. P. Young: *The Weyl bound for Dirichlet L-functions of cube-free conductor*, preprint arXiv:1811.02452v1

[19] Polymath8 project, http://terrytao.wordpress.com/tag/polymath8/

[20] Polymath8 project, "An improved Type I estimate", http://terrytao.wordpress.com/2013/07/27/an-improved-type-i-estimate/

[21] J.J. Rotman: *An introduction to homological algebra*, Second Edition, Springer Universitext, 2009.

Université Paris Sud, Laboratoire de Mathématique, Campus d'Orsay, 91405 Orsay Cedex, France

Email address: etienne.fouvry@math.u-psud.fr

ETH Zürich – D-MATH, Rämistrasse 101, CH-8092 Zürich, Switzerland

Email address: kowalski@math.ethz.ch

EPFL/SB/IMB/TAN, Station 8, CH-1015 Lausanne, Switzerland

Email address: philippe.michel@epfl.ch