

ZETA FUNCTIONS OF NONDEGENERATE HYPERSURFACES IN TORIC VARIETIES VIA CONTROLLED REDUCTION IN p -ADIC COHOMOLOGY

EDGAR COSTA, DAVID HARVEY, AND KIRAN S. KEDLAYA

ABSTRACT. We give an interim report on some improvements and generalizations of the Abbott–Kedlaya–Roe method to compute the zeta function of a nondegenerate ample hypersurface in a projectively normal toric variety over \mathbb{F}_p in linear time in p . These are illustrated with a number of examples including K3 surfaces, Calabi–Yau threefolds, and a cubic fourfold. The latter example is a non-special cubic fourfold appearing in the Ranestad–Voisin coplanar divisor on moduli space; this verifies that the coplanar divisor is not a Noether–Lefschetz divisor in the sense of Hassett.

1. INTRODUCTION

We consider the problem of computing the zeta function $Z(\mathcal{X}, t)$ of an explicitly specified variety \mathcal{X} over a finite field \mathbb{F}_q of characteristic p . For curves and abelian varieties, Schoof’s method and variants [?, ?, ?, ?, ?] can compute $Z(\mathcal{X}, t)$ in time and space polynomial in $\log q$ and exponential in the genus/dimension; these have only been implemented for genus/dimension at most 2. Such methods may be characterized as ℓ -adic, as they access the ℓ -adic cohomology (for $\ell \neq p$ prime) of the variety via torsion points; there also exist p -adic methods which compute approximations of the Frobenius action on p -adic cohomology (Monsky–Washnitzer cohomology), and which have proven to be more viable in practice for large genus. Early examples include Kedlaya’s algorithm [?] for hyperelliptic curves, in which the time/space dependence is polynomial in the genus and quasi-linear in p , and Harvey’s algorithm [?] which improves the dependence on p to $p^{1/2+\epsilon}$. These methods have been subsequently generalized [?, ?, ?, ?], notably by Tuitman’s algorithm [?, ?] which applies to (almost) all curves while keeping the quasi-linear dependence on p . In another direction, Harvey [?] has shown that when computing the zeta functions of reductions of a fixed hyperelliptic curve over a number field, p -adic methods can achieve *average* polynomial time in $\log p$ and the genus; this has been implemented in small genus [?, ?].

One advantage of p -adic methods over ℓ -adic ones is that they scale much better to higher-dimensional varieties. For example, there are several p -adic constructions that apply to *arbitrary* varieties with reasonable asymptotic complexity [?, ?], although we are not aware of any practical implementations. Various algorithms, and some implementations, have been given using Lauder’s *deformation method*

The first author was partially supported by the Simons Collaboration Grant #550029. The second author was supported by the Australian Research Council (grants DP150101689 and FT160100219). The third author was supported by NSF (grants DMS-1101343, DMS-1501214); UC San Diego (Warschawski Professorship); and a Guggenheim Fellowship. All three authors thank ICERM for its hospitality during fall 2015.

of computing the Frobenius action on the Gauss–Manin connection of a pencil [?, ?, ?, ?, ?, ?].

In this paper, we build on an algorithm of Abbott–Kedlaya–Roe [?] which adapts the original approach of [?] to smooth projective hypersurfaces. Here, we add two key improvements.

- We use *controlled reduction* in de Rham cohomology, as described in some lectures of Harvey [?, ?, ?], to preserve sparsity of certain polynomials, thus reducing the time (respectively, space) dependence on p from polynomial to quasi-linear (respectively, $O(\log p)$). The resulting *controlled AKR method* was implemented, with further improvements, in Costa’s Ph.D. thesis [?], with examples of generic surfaces and threefolds over \mathbb{F}_p for $p \sim 10^6$ [?, Section 1.6]; by contrast, the largest p used in [?] is 29. Costa and Harvey are currently preparing a paper on this method; meanwhile, Costa’s GPL-licensed code is available on GITHUB [?], and is slated to be integrated into SAGEMATH [?].
- We also generalize to toric hypersurfaces, subject to a standard genericity condition called *nondegeneracy*. This greatly increases the applicability of the method while preserving much of its efficiency. Some previous attempts have been made to compute zeta functions in this setting, such as work of Castryck–Denef–Vercauteren [?] for curves and Sperber–Voight [?] in general; it is the combination with controlled reduction that makes our approach the most practical to date.

It may be possible to improve the dependence on p to square-root (as in [?]) or average polynomial time (as in [?]), but we do not attempt to do so here.

For reasons of space, we give only a summary of the algorithm, with further details to appear elsewhere. In lieu of these details, we present a number of worked examples in dimensions 2–4 that demonstrate the practicality of this algorithm in a wide range of cases. The results are based on an implementation in C++, using NTL [?] for the underlying arithmetic operations. Our examples in dimensions 2 and 3 were computed on one core of a desktop machine with an Intel(R) Core(TM) i5-4590 CPU @ 3.30GHz; our sole example in dimension 4 was computed on one core of a server with an AMD Opteron Processor 6378 @ 1.6GHz. (We have not yet optimized our vector-matrix multiplications in any way; as a consequence, we observe a serious performance hit whenever the working moduli exceeds 2^{62} .)

In dimensions 2 and 3, our examples are *Calabi–Yau varieties*, i.e., smooth, proper, simply connected varieties with trivial canonical bundle. In dimension 1, these are simply elliptic curves. In dimension 2, they are *K3 surfaces*, whose zeta functions are of computational interest for various reasons. For instance, these zeta functions can (potentially) be used to establish the infinitude of rational curves on a K3 surface (see the introduction to [?] for discussion); there has also been recent work on analogues of the Honda–Tate theorem, establishing conditions under which particular zeta functions are realized by K3 surfaces [?, ?].

As for Calabi–Yau threefolds, much of the interest in their zeta functions can be traced back to *mirror symmetry* in mathematical physics. An early example is the work of Candelas–de la Ossa–Rodriguez Villegas [?] on the Dwork pencil; a more recent example is [?], in which (using p -adic cohomology) certain mirror families of Calabi–Yau threefolds are shown to have related zeta functions.

Our four-dimensional example is a cubic projective fourfold. Such varieties occupy a boundary region between rational and irrational varieties; it is expected that a rational cubic fourfold is *special* in the sense of having a primitive cycle class in codimension 2. The geometry of special cubic fourfolds is in turn closely linked to that of K3 surfaces; in many cases, the Hodge structure of a K3 surface occurs (up to a twist) inside the Hodge structure of a special cubic fourfold, and (modulo standard conjectures) this implies a similar relationship between zeta functions. See [?] for further discussion.

The specific example we consider is related to the geometry of the moduli space of cubic fourfolds over \mathbb{C} . On this space, there exist various divisors consisting entirely of special cubic fourfolds; Hassett calls these *Noether–Lefschetz divisors* (by analogy with the case of surfaces). Recently, Ranestad–Voisin [?] exhibited four divisors which they believed not to be Noether–Lefschetz, but only checked this in one case. Addington–Auel [?] checked two more cases by finding in these divisors some cubic fourfolds over \mathbb{Q} with good reduction at 2, such that the zeta functions over \mathbb{F}_2 show no primitive Tate classes in codimension 2. By replacing the brute-force point counts of Addington–Auel with p -adic methods, we are able to work modulo a larger prime to find an example showing that the fourth Ranestad–Voisin divisor is not Noether–Lefschetz.

To sum up, the overall goal of this project is to vastly enlarge the collection of varieties for which computing the zeta function is practical. It is our hope that doing so will lead to a rash of new insights, conjectures, and theorems of interest to a broad range of number theorists and algebraic geometers.

2. TORIC HYPERSURFACES

We begin by reviewing the construction of a projective toric variety from a lattice polytope. For more details we recommend [?].

Let $n \geq 1$ be an integer. For any commutative ring R , let $R[x^\pm]$ denote the Laurent polynomial ring in n variables x_1, \dots, x_n with coefficients in R . For $\alpha := (\alpha_i) \in \mathbb{Z}^n$, we write x^α for the monomial $x_1^{\alpha_1} \cdots x_n^{\alpha_n}$. We denote the R -torus by $\mathbb{T}_R^n := \text{Spec}(R[x^\pm])$.

Let $\Delta \subset \mathbb{R}^n$ be the convex hull of a finite subset of \mathbb{Z}^n that is not contained in any hyperplane, so that $\dim \Delta = n$. For $r \in \mathbb{R}$, let $r\Delta$ be the r -fold dilation of Δ . For an integer $d \geq 0$, let

$$P_d := \langle x^\alpha : \alpha \in d\Delta \cap \mathbb{Z}^n \rangle_R \quad (\text{resp. } P_d^{\text{Int}} := \langle x^\alpha : \alpha \in \text{Int}(d\Delta) \cap \mathbb{Z}^n \rangle_R)$$

be the free R -module on the set of monomials with exponents in $d\Delta \cap \mathbb{Z}^n$ (resp. $\text{Int}(d\Delta) \cap \mathbb{Z}^n$). Define the R -graded algebras

$$P_\Delta := \bigoplus_{d=0}^{+\infty} P_d \quad \text{and} \quad P_\Delta^{\text{Int}} := \bigoplus_{d=0}^{+\infty} P_d^{\text{Int}}.$$

with the usual multiplication in $R[x^\pm]$. We define the polarized toric variety associated to Δ as the pair $(\mathbb{P}_\Delta, \mathcal{O}_\Delta)$, where $\mathbb{P}_\Delta := \text{Proj } P_\Delta$ and \mathcal{O}_Δ is the ample line bundle on \mathbb{P}_Δ associated to the graded P_Δ -module $P_\Delta(1)$. Note that P_Δ and P_Δ^{Int} admit n commuting degree-preserving differential operators $\partial_i := x_i \frac{\partial}{\partial x_i}$ for $i = 1, \dots, n$.

In order to suppress some expository and algorithmic complexity, we make the simplifying assumption that Δ is a *normal* polytope; that is, the map

$$(\Delta \cap \mathbb{Z}^n)^d \rightarrow d\Delta \cap \mathbb{Z}^n : (x_1, \dots, x_d) \mapsto x_1 + \dots + x_d$$

is surjective for $d \geq 1$. This corresponds to the pair $(\mathbb{P}_\Delta, \mathcal{O}_\Delta)$ being *projectively normal*; this will be the case in our examples. As a consequence, we have that \mathcal{O}_Δ is indeed very ample.

Example 2.1. Let Δ be the regular n -simplex, the convex hull of $0, e_1, \dots, e_n$. We may then identify P_d with the set of homogeneous polynomials of degree d in x_0, \dots, x_n , by identifying $x^\alpha \in P_{\Delta, d}$ with the monomial $x_0^{d-\alpha_1-\dots-\alpha_n} x_1^{\alpha_1} \dots x_n^{\alpha_n}$; then $(\mathbb{P}_\Delta, \mathcal{O}_\Delta)$ is isomorphic to $(\mathbb{P}_R^n, \mathcal{O}(1))$.

We obtain the weighted projective space $\mathbb{P}(w_0, \dots, w_n)$ by taking

$$\Delta = \{(x_0, \dots, x_n) \in \mathbb{R}^{n+1} : \sum_{i=0}^n w_i x_i = w_0 \dots w_n\}, \quad \text{see [?, 1.2.5].}$$

We obtain $\mathbb{P}_R^k \times_R \mathbb{P}_R^r$ by taking Δ to be the Cartesian product of the regular k -simplex by the regular r -simplex [?, §2.4].

We now turn our attention to toric hypersurfaces over $R = \mathbb{F}_q$, the finite field with $q = p^a$ elements and characteristic p . Let \mathcal{Y} be the hypersurface in $\mathbb{T}_{\mathbb{F}_q}^n$ defined by a Laurent polynomial $\bar{f} \in \mathbb{F}_q[x^\pm]$, $\mathcal{Y} := V(\bar{f}) \subset \mathbb{T}_{\mathbb{F}_q}^n$. Let

$$\text{supp } \bar{f} = \{\alpha \in \mathbb{Z}^n : \bar{c}_\alpha \neq 0\}$$

be the support of \bar{f} in \mathbb{R}^n ; the convex hull of $\text{supp } \bar{f}$ is the *Newton polytope* of \bar{f} , which we denote by Δ . We will work under the hypothesis that \bar{f} is (Δ) -*nondegenerate*:¹ for all faces $\tau \subseteq \Delta$ (including Δ itself), the system of equations

$$\bar{f}|_\tau = \partial_1 \bar{f}|_\tau = \dots = \partial_n \bar{f}|_\tau = 0$$

has no solution in $\bar{\mathbb{F}}_q^{\times n}$, where $\bar{\mathbb{F}}_q$ denotes an algebraic closure of \mathbb{F}_q . Furthermore, nondegeneracy implies quasi-smoothness, see [?, Definition 3.1 and Proposition 4.15]. For fixed normal Δ over an infinite field, this condition holds for generic \bar{f} . Others have given point-counting algorithms under this assumption [?, ?].

Let $\mathcal{X} := \text{Proj } P_\Delta / (\bar{f})$ denote the closure of \mathcal{Y} in \mathbb{P}_Δ (placing \bar{f} in degree 1) and set $\mathcal{U} := \mathbb{T}^n \setminus \mathcal{Y}$. Let H_{rig}^i denote the i th rigid cohomology group in the sense of Berthelot [?]. The Lefschetz hyperplane theorem, combined with Poincaré duality, show that the map

$$H_{\text{rig}}^i(\mathbb{P}_\Delta) \rightarrow H_{\text{rig}}^i(\mathcal{X}),$$

induced by the inclusion $\mathcal{X} \hookrightarrow \mathbb{P}_\Delta$ is an isomorphism for $i \neq n-1$ [?, 10.8]. This implies that the “interesting” part of the cohomology of \mathcal{X} occurs in dimension $n-1$ and consists of those classes that do not come from P_Δ . Denote by $PH_{\text{rig}}^{n-1}(\mathcal{X})$ the primitive cohomology group of \mathcal{X} , defined by the (Frobenius-equivariant) exact sequence

$$0 \rightarrow H_{\text{rig}}^{n-1}(\mathbb{P}_\Delta) \rightarrow H_{\text{rig}}^{n-1}(\mathcal{X}) \rightarrow PH_{\text{rig}}^{n-1}(\mathcal{X}) \rightarrow 0$$

With this notation, we may write

$$Z(\mathcal{X}, t) = Z(\mathbb{P}_\Delta, t)Q(t)^{(-1)^n}.$$

¹This condition was introduced by Dwork [?] without a name; the term *nondegenerate* first appears in [?, ?]. Synonyms include Δ -*regular* [?, §4] and *schön* [?].

where

$$Q(t) := \det(1 - t \text{Frob}_q | PH_{\text{rig}}^{n-1}(\mathcal{X})).$$

Thus given \bar{f} , we would like to compute $Q(t)$.

The cohomology group $PH_{\text{rig}}^{n-1}(\mathcal{X})$ is closely related to $H_{\text{rig}}^n(\mathbb{P}_\Delta \setminus \mathcal{X})$. For example, if \mathbb{P}_Δ is a (weighted) projective space, as in [?] and [?], the two cohomology groups are isomorphic; see [?, 10.11].

3. DE RHAM COHOMOLOGY OF TORIC HYPERSURFACES

In preparation for our use of p -adic cohomology to compute $Q(t)$, we give an explicit description of the algebraic de Rham cohomology of a nondegenerate toric hypersurface in characteristic zero. We take R to be the ring \mathbb{Z}_q , the ring of integers of \mathbb{Q}_q , the unramified extension of \mathbb{Q}_p with residue field \mathbb{F}_q .

Let $f \in \mathbb{Z}_q[x^\pm]$ be a lift of \bar{f} to characteristic zero with the same support as f (it will also be nondegenerate). Consider $Y := V(f) \subset \mathbb{T} := \mathbb{T}_{\mathbb{Q}_q}$ and X , the closure of Y in \mathbb{P}_Δ . Write $U := \mathbb{T} \setminus Y$, and $V := \mathbb{P}_\Delta \setminus X \simeq \text{Spec}(A)$, where A is the coordinate ring of V ; explicitly,

$$A \simeq \bigcup_{d=0}^{+\infty} f^{-d} P_d.$$

Let I_f be the ideal in P_Δ generated by $f, \partial_1 f, \dots, \partial_n f$. We call I_f the *toric Jacobian ideal* and the quotient ring $J_f := P_\Delta / I_f$ the *toric Jacobian ring*. Since f is nondegenerate, the ideal I_f is irrelevant in P_Δ and $\text{rank}_{\mathbb{Z}_q} J_f = n! \text{Vol}(\Delta)$; furthermore, $(J_f)_d = 0$ for $d > n$ [?, §4]. If \mathcal{O}_Δ is not very ample, then I_f might not be generated in degree 1 and we might have $(J_f)_d = 0$ only for $d \gg n$.

Let Ω^\bullet denote the logarithmic de Rham complex of V with poles along $\mathbb{P}_\Delta \setminus \mathbb{T}$. Let H^\bullet be the cohomology groups of Ω^\bullet ; these are naturally isomorphic to $H_{\text{dR}}^\bullet(V \cap \mathbb{T} = \mathbb{T} \setminus Y = U)$ and $H_{\text{rig}}^\bullet(\mathbb{T}_{\mathbb{F}_q} \setminus \mathcal{Y} = \mathcal{U})$ [?].

We now provide an explicit description of the group H^n , as in [?, §6 and 7], in which we will compute $Q(t)$. Set

$$\omega := \frac{dx_1}{x_1} \wedge \cdots \wedge \frac{dx_n}{x_n} \in \Omega^n,$$

and define the ascending filtration in Ω^n by

$$\text{Fil}^d \Omega^n := \{g f^{-d} \omega : g \in P_d\}.$$

The associated graded ring

$$\Omega^n := \bigoplus_{d=0}^{\infty} \text{Gr}^d \Omega^n, \quad \text{Gr}^d \Omega^n := \text{Fil}^d \Omega^n / \text{Fil}^{d-1} \Omega^n$$

is then isomorphic to $P_\Delta / (f)$ (again placing f in degree 1).

Equip H^n with the filtration induced from Ω^n , and view H^n as the quotient of Ω^n by the \mathbb{Q}_q -submodule generated by the relations

$$(3.1) \quad \frac{g}{f^d} \omega - \frac{gf}{f^{d+1}} \omega \quad \text{and} \quad \frac{\partial_i(g)}{f^d} \omega - \frac{dg \partial_i(f)}{f^{d+1}} \omega$$

for each $i = 1, \dots, n$, each nonnegative integer d , and each $g \in P_d$. From these relations, we see that

$$\text{Gr}^1 H^n \simeq P_1 / (f) \quad \text{and} \quad \text{Gr}^d H^n \simeq (J_f)_d \quad (d > 1).$$

This gives a way to compute explicitly in H^n : for any $h \in (J_f)_{d+1}$ with $d > n$, we can find a relation of the form

$$(3.2) \quad d \frac{h}{f^{d+1}} \omega = d \frac{g_0 f + \sum_{i=1}^n g_i \partial_i f}{f^{d+1}} \omega \equiv \frac{dg_0 + \sum_{i=1}^n \partial_i g_i}{f^d} \omega.$$

because $P_d \subset (I_f)_d$, so in H^n we can reduce the pole order of any form to at most n . This process was introduced for smooth projective hypersurfaces in [?] and attributed to Dwork; it is commonly known as *Griffiths–Dwork reduction*.

With the above representation of H^n , we may also identify $PH_{\text{dR}}^{n-1}(X)$ with $(P_{\Delta}^{\text{Int}} + I_f)/I_f \subset H^n$, where the filtration by pole order is the Hodge filtration; see [?, ?, §9, §11].

We now introduce a variation of Griffiths–Dwork reduction, called *controlled reduction*. This will be crucial for our application to p -adic cohomology, as careless application of Griffiths–Dwork reduction to a sparse form will easily lead to a dense form. For $d = 1, \dots, n+1$, choose a \mathbb{Z}_q -linear splitting $P_d \approx (J'_f)_d \oplus C_d$, where $(J'_f)_d$ is a lift of $(J_f)_d$ into P_d . Let $\rho_d: P_d \rightarrow (J'_f)_d$ and $\pi_{d,0}, \dots, \pi_{d,n}: P_d \rightarrow P_{d-1}$ be \mathbb{Z}_q -linear maps such that

$$g = \rho_d(g) + \pi_{d,0}(g) \cdot f + \sum_{i=1}^n \pi_{d,i}(g) \cdot \partial_i f; \quad g \in P_d.$$

These maps may be constructed one monomial at a time.

Proposition 3.1 (Controlled reduction). *Let $x^\nu \in P_1$ and $x^\mu \in P_d$ be two monomials and define the following \mathbb{Z}_q -linear maps:*

$$R_{\mu,\nu}(g) := (d+n)\pi_{n+1,0}(x^\nu g) + \sum_{i=1}^n (\partial_i + \mu_i)(\pi_{n+1,i}(x^\nu g))$$

$$S_\nu(g) := \pi_{n+1,0}(x^\nu g) + \sum_{i=1}^n \nu_i \pi_{n+1,i}(x^\nu g)$$

Then for any $g \in P_n$ and any nonnegative integer j , in H^n we have

$$g \frac{x^{(j+1)\nu+\mu}}{f^{d+n+j+1}} \omega \equiv (d+n+j)^{-1} (R_{\mu,\nu}(g) + jS_\nu(g)) \frac{x^{j\nu+\mu}}{f^{d+n+j}} \omega.$$

Proof. This is straightforward from (3.1) and (3.2). \square

Note that Proposition 3.1 enables us to reduce the pole order of a differential form from $d+n+j+1$ to $d+n+j$ without increasing its total number of monomials; we can thus reduce the pole order of a sparse form without making it dense.

Corollary 3.2. *With notation as in Proposition 3.1, let k be a positive integer. Then for any $g \in P_n$,*

$$g \frac{x^{\mu+k\nu}}{f^{d+n+k}} \omega \equiv \frac{\prod_{j=0}^{k-1} (R_{\mu,\nu} + jS_\nu)(g)}{\prod_{j=0}^{k-1} (d+n+j)} \frac{x^\mu}{f^{d+n}} \omega,$$

forming the composition product from left to right.

Using Proposition 3.1 amounts to performing linear algebra on matrices of size $\#(n\Delta \cap \mathbb{Z}^n) \sim n^n \text{Vol}(\Delta)$. One can reduce this by a factor of $n^n/n! \sim e^n$ at the expense of making the expression for the reduction matrix more convoluted; compare [?, Proposition 1.17 and 1.18].

4. MONSKY–WASHNITZER COHOMOLOGY

We now indicate how Monsky–Washnitzer cohomology, as introduced in [?, ?, ?], provides a crucial link between algebraic de Rham cohomology and p -adic rigid cohomology, by transferring to the former the canonical Frobenius action on the latter; see also [?]. To simplify, we assume $p > \max\{n, 2\}$.

Let A^\dagger denote the *weak p -adic completion* of A , the ring consisting of formal sums $\sum_{d=0}^{+\infty} g_d f^{-d}$ such that for some $a, b > 0$, $g_d \in p^{\max\{0, \lfloor ad-b \rfloor\}} P_d$ for all $d \geq 0$. We define the associated logarithmic de Rham complex $\Omega^{\dagger, \bullet}$ by $\Omega^{\dagger, i} := \Omega^i \otimes_A A^\dagger$; denote the cohomology groups of this complex by $H^{\dagger, \bullet}$. We may then obtain p -adic Monsky–Washnitzer cohomology groups $H^{\dagger, \bullet} \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$. The map $\Omega^\bullet \otimes_{\mathbb{Z}_q} \mathbb{Q}_q \rightarrow \Omega^{\dagger, \bullet} \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$ is a quasi-isomorphism [?, ?, ?]; that is, the induced maps $H^i \otimes_{\mathbb{Z}_q} \mathbb{Q}_q \rightarrow H^{\dagger, i} \otimes_{\mathbb{Z}_q} \mathbb{Q}_q$ are isomorphisms. We can thus identify the algebraic de Rham cohomology of U with the Monsky–Washnitzer cohomology of \mathcal{U} .

On the other hand, we also have $H^{\dagger, \bullet} \otimes_{\mathbb{Z}_q} \mathbb{Q}_q \simeq H_{\text{rig}}^\bullet(\mathcal{U})$ and the latter object is functorial with respect to geometry in characteristic p [?]. In this way, $H^{\dagger, i}$ receives an action of the Frobenius automorphism, which we can make explicit by constructing a lift σ of the p -th power Frobenius on \mathbb{F}_q to A^\dagger . To do so, we take the Witt vector Frobenius on \mathbb{Z}_q and set $\sigma(\mu) = \mu^p$ for any monomial $\mu \in P_\Delta$. We then extend σ to A^\dagger by the formula

$$(4.1) \quad \sigma\left(\frac{g}{f^d}\right) := \sigma(g)\sigma(f)^{-d} = \sigma(g) \sum_{i \geq 0} \binom{-d}{i} \frac{(\sigma(f) - f^p)^i}{f^{p(d+i)}}.$$

The above series converges (because p divides $\sigma(f) - f^p$) and the definitions ensure that σ is a semilinear (with respect to the Witt vector Frobenius) endomorphism of A^\dagger . We finally extend σ to $\Omega^{\dagger, \bullet}$ by $\sigma(g dh) := \sigma(g) d(\sigma(h))$.

5. SKETCH OF THE ALGORITHM

We now indicate briefly how to use controlled reduction to compute the Frobenius action on the cohomology of nondegenerate toric hypersurfaces. We start as in [?, Proposition 4.1], by rewriting the Frobenius action in a sparser form.

Lemma 5.1. *For any positive integers d, N and $g \in P_d$, in A^\dagger we have*

$$\sigma\left(\frac{g}{f^d}\right) \equiv \sum_{j=0}^{N-1} \binom{-d}{j} \binom{d+N-1}{d+j} \sigma(g f^j) f^{-p(d+j)} \pmod{p^N}.$$

Proof. This follows from (4.1) by truncating the sum and then rewriting formally; see [?, Lemma 1.10]. \square

In order to compute a p -adic approximation of the Frobenius action on $PH^{n-1}(\mathcal{X})$, we must first fix a basis of the latter; we do this by constructing a monomial basis for $PH_{\text{dR}}^{n-1}(X)$ via explicit linear algebra. We then apply Frobenius to each basis element in the sparse truncated form given by Lemma 5.1; recursively reduce the pole order using Corollary 3.2 (using $k = p$ as much as possible); and project to the chosen monomial basis. The dominant step is controlled reduction, which amounts to $O(pn^N \text{Vol}(\Delta))$ matrix multiplications of size $n! \text{Vol}(\Delta)$ per basis element.

We will not address precision estimates in this report, except to note that the machinery of [?, §3.4] applies. In general, if we want N digits of p -adic accuracy, we must apply Lemma 5.1 with N replaced by $N' = N + O(n + \log N)$ and work

modulo $p^{O(N')}$. Hence, with respect to p alone, we expect our algorithm to run in quasi-linear time in p and use $O(\log p)$ space.

6. K3 SURFACES

We now turn our attention to examples, starting with K3 surfaces. For X a K3 surface, $\dim H^2(X) = 22$ and the Hodge numbers are $(1, 20, 1)$. A common example of a K3 surface is a smooth quartic surface in \mathbb{P}^3 ; however, they also occur in other ways, such as hypersurfaces in weighted projective spaces. Using a criterion of Miles Reid [?], Yonemura [?] found the complete list of (polarized) weighted projective spaces in which a generic hypersurface is a K3 surface; there are 95 of these. For toric varieties, the corresponding classification is that of reflexive 3-dimensional polytopes, of which there are 4,319 in all [?].

In the following examples, we worked modulo p^4 in order to obtain $Q(t)$ with 2 p -adic significant digits. As a result, we observe a performance hit for $p > 2^{16}$.

Example 6.1. Consider the projective quartic surface $\mathcal{X} \subset \mathbb{P}_{\mathbb{F}_p}^3$ defined by

$$x^4 + y^4 + z^4 + w^4 + \lambda xyzw = 0;$$

it is a member of the Dwork pencil. For $p = 2^{20} - 3$ and $\lambda = 1$, using the *controlled AKR algorithm* in 22h7m we compute that


$$Z(\mathcal{X}, t)^{-1} = (1 - t)(1 - pt)^{16}(1 + pt)^3(1 - p^2t)Q(t),$$

where the “interesting” factor is

$$Q(t) = (1 + pt)(1 - 1688538t + p^2t^2).$$

For this family, the remaining factors, apart from $Q(t)$, could have also been deduced by a p -adic formula of de la Ossa-Kadir [?, Chapter 6]. In this context, the Hodge numbers of $PH^2(\mathcal{X})$ are $(1, 19, 1)$.

A similar runtime would be expected if we used our current implementation to compute $Z(\mathcal{X}, t)$ with Δ being the 3-simplex (tetrahedron), as indicated by the outer polytope at right. Instead, we observe that the monomials defining \mathcal{X} generate a sublattice of index 4^2 in \mathbb{Z}^3 ; hence, we can instead run our algorithm with a polytope of significantly smaller volume ($32/3 \approx 10.66$ versus $2/3 \approx 0.66$), as indicated by the inner polytope at right. This leads to a dramatic speedup: with our current implementation, we computed $Q(t)$ in 1m33s.



We present the running times for other p in Table 1; memory usage is about 16MB.

In the new framework, \mathcal{X} is given by the closure (in \mathbb{P}_{Δ}) of the affine surface defined by the Laurent polynomial

$$x^4y^{-1}z^{-1} + \lambda x + y + z + 1,$$

and the Hodge numbers of $PH^2(\mathcal{X})$ are $(1, 1, 1)$, which explains why $\deg Q(t) = 3$.

Since the Dwork pencil is a “small” deformation of the Fermat quartic, we may also use the Pancratz–Tuitman implementation of the *deformation method* [?] to compute $Z(\mathcal{X}, t)$. We did this and verified that our results agree; we compare running times in Table 1. To interpret these fairly, note that Pancratz–Tuitman work in \mathbb{P}^3 and so compute the whole numerator of $Z(\mathcal{X}, t)$ rather than just $Q(t)$. (Note that the algorithm of [?] has a square-root dependence on p , as in [?].)

p	CHK time	PT time	p	CHK time
$2^8 - 5$	0.03s	1.65s	$2^{17} - 1$	11.9s
$2^9 - 3$	0.04s	3.64s	$2^{18} - 5$	23.4s
$2^{10} - 3$	0.04s	7.39s	$2^{19} - 1$	46.9s
$2^{11} - 9$	0.06s	14.65s	$2^{20} - 3$	1m33s
$2^{12} - 3$	0.08s	34.80s	$2^{21} - 9$	3m6s
$2^{13} - 1$	0.13s	34.80s	$2^{22} - 3$	6m15s
$2^{14} - 3$	0.22s	2m33s		
$2^{15} - 19$	0.41s	6m43s		
$2^{16} - 15$	5.72s	14m14s		

TABLE 1. The second and fifth columns use our current implementation to compute $Q(t)$. The third column uses the Pancratz–Tuitman implementation [?] to compute $Z(\mathcal{X}, t)$.

Example 6.2. Consider the projective quartic surface $\mathcal{X} \subset \mathbb{P}_{\mathbb{F}_p}^3$ defined by

$$x^3y + y^4 + z^4 + w^4 - 12xyzw;$$


it contains a hypergeometric motive (see [?, Section 5]). For $p = 2^{15} - 19$, using the *controlled AKR algorithm* in 27m12s we compute that

$$Z(\mathcal{X}, t)^{-1} = (1-t)(1-pt)^2(1+pt)^2(1-pt+p^2t^2)^2(1-p^2t^2+p^4t^4)^2(1-p^2t)Q(t),$$

where the “interesting” factor is (up to rescaling)

$$pQ(t/p) = p + 20508t^1 - 18468t^2 - 26378t^3 - 18468t^4 + 20508t^5 + pt^6.$$

As in the previous example, the Newton polytope has volume 8, but the defining monomials generate a sublattice of index 4 in \mathbb{Z}^3 ; we may thus work instead with a polytope of volume 2 (depicted at right) and observe a significant speedup. In this setting, the Hodge numbers of $PH^2(\mathcal{X})$ are (1, 4, 1). With our current implementation we computed $Q(t)$ in 4s. We present the running times for other p in Table 2, where the memory footprint was about 52MB.



Alternatively, one could try to use MAGMA [?] to confirm $Q(t)$. Unfortunately, MAGMA is only able to confirm the linear coefficient:

```
> C2F2 := HypergeometricData([6,12], [1,1,1,2,3]);
> EulerFactor(C2F2, 2^10 * 3^6, 2^15 - 19: Degree:=1);
1 + 20508*$.1 + 0($.1^2)
```

p	time	p	time	p	time
$2^8 - 5$	0.20s	$2^{13} - 1$	1.12s	$2^{18} - 5$	4m54s
$2^9 - 3$	0.23s	$2^{14} - 3$	2.08s	$2^{19} - 1$	9m46s
$2^{10} - 3$	0.29s	$2^{15} - 19$	4.00s	$2^{20} - 3$	19m32s
$2^{11} - 9$	0.41s	$2^{16} - 15$	1m11s	$2^{21} - 9$	38m58s
$2^{12} - 3$	0.64s	$2^{17} - 1$	2m30s	$2^{22} - 3$	1h18m

TABLE 2. Running times for Example 6.2.

Example 6.3. Consider the closure \mathcal{X} in \mathbb{P}_Δ (which in this case is not a weighted projective space) of the affine surface defined by the Laurent polynomial

$$3x + y + z + x^{-2}y^2z + x^3y^{-6}z^{-2} + 3x^{-2}y^{-1}z^{-2} \\ - 2 - x^{-1}y - y^{-1}z^{-1} - x^2y^{-4}z^{-1} - xy^{-3}z^{-1};$$

it is a K3 surface of geometric Picard rank 6, and the Hodge numbers of $PH^2(\mathcal{X})$ are $(1, 14, 1)$. For $p = 2^{15} - 19$, using our current implementation, in 6m20s we obtain the “interesting” factor of $Z(\mathcal{X}, t)$:


$$pQ(t/p) = (1-t) \cdot (1+t) \cdot (p + 33305t^1 + 1564t^2 - 14296t^3 - 11865t^4 \\ + 5107t^5 + 27955t^6 + 25963t^7 + 27955t^8 + 5107t^9 \\ - 11865t^{10} - 14296t^{11} + 1564t^{12} + 33305t^{13} + pt^{14}).$$

We present the running times for other p in Table 3, where the peak memory usage was about 144MB.

The vertices of the associated polytope correspond to the first six terms displayed; the remaining terms are interior points. We depict this polytope of volume 8 at right.

We know of no previous algorithm that can compute $Z(\mathcal{X}, t)$ for p in this range. The defining polynomial is “dense” from the point of the Sperber–Voight algorithm [?], which is based on Dwork cohomology and scales

with the number of monomials away from the vertices of the Newton polytope.



p	time	p	time	p	time
$2^7 - 1$	6.46s	$2^{10} - 3$	18.93s	$2^{13} - 1$	1m46s
$2^8 - 5$	9.50s	$2^{11} - 9$	31.34s	$2^{14} - 3$	3m24s
$2^9 - 3$	12.64s	$2^{12} - 3$	56.23s	$2^{15} - 19$	6m20s

TABLE 3. Running times for Example 6.3.

Example 6.4. Let \mathcal{X} be the smooth projective surface in \mathbb{P}^3 defined by the fully dense, randomly chosen quartic polynomial

$$-9x^4 - 10x^3y - 9x^2y^2 + 2xy^3 - 7y^4 + 6x^3z + 9x^2yz - 2xy^2z + 3y^3z \\ + 8x^2z^2 + 6y^2z^2 + 2xz^3 + 7yz^3 + 9z^4 + 8x^3w + x^2yw - 8xy^2w \\ - 7y^3w + 9x^2zw - 9xyzw + 3y^2zw - xz^2w - 3yz^2w + z^3w - x^2w^2 \\ - 4xyw^2 - 3xzw^2 + 8yzw^2 - 6z^2w^2 + 4xw^3 + 3yw^3 + 4zw^3 - 5w^4;$$

then Δ is the 3-simplex (tetrahedron) of volume $32/3 \approx 10.66$. For this example, we have $PH^2(\mathcal{X}) \simeq H^3(\mathbb{P}^3 \setminus \mathcal{X})$, the Hodge numbers are $(1, 19, 1)$, and

$$Z(\mathcal{X}, t)^{-1} = (1-t)(1-pt)(1-p^2t)Q(t)$$

where $\deg Q(t) = 21$. For $p = 2^{15} - 19$, we obtain

$$pQ(t/p) = (1+t)(p - 53159t^1 + 10023t^2 - 3204t^3 + 49736t^4 - 56338t^5 + 43086t^6 \\ - 48180t^7 + 44512t^8 - 42681t^9 + 47794t^{10} - 42681t^{11} + 44512t^{12} - 48180t^{13} \\ + 43086t^{14} - 56338t^{15} + 49736t^{16} - 3204t^{17} + 10023t^{18} - 53159t^{19} + pt^{20})$$

using the *controlled AKR algorithm* in 38m27s; our current implementation takes roughly the same time. We present the running times for other p in Table 4. The memory footprint was about 230MB.

Unfortunately, the *deformation method* is not suitable for dense quartics with p in this range. For example, for $p = 31$ the running time was 2h8m and its memory footprint was around 7GB, and both time and space should scale linearly with p .

p	time	p	time	p	time
$2^7 - 1$	25.41s	$2^{10} - 3$	1m30s	$2^{13} - 1$	9m26s
$2^8 - 17$	37.73s	$2^{11} - 9$	2m37s	$2^{14} - 3$	18m42s
$2^9 - 3$	55.82s	$2^{12} - 3$	4m50s	$2^{15} - 19$	36m29s

TABLE 4. Running times for Example 6.4.

7. CALABI–YAU THREEFOLDS

We next consider Calabi–Yau threefolds. Unlike for K3 surfaces, the middle Betti numbers of Calabi–Yau threefolds are not *a priori* bounded; the largest value of which we are aware is 984 (found in [?]).

A common example is a smooth quintic surface in \mathbb{P}^4 . Again, additional constructions arise from generic hypersurfaces in weighted projective spaces, of which there are 7,555 in all, or more generally from toric varieties corresponding to reflexive 4-dimensional polytopes, of which there are 473,800,776 in all [?].

In all of the following examples, we worked modulo p^6 in order to obtain $Q(t)$ and our memory footprint ranged between 100MB and 270MB.

Example 7.1. Consider the projective quintic threefold $\mathcal{X} \subset \mathbb{P}_{\mathbb{F}_p}^3$ defined by

$$x_0^5 + x_1^5 + x_2^5 + x_3^5 + x_4^5 + x_0x_1x_2x_3x_4 = 0;$$

it is a member of the Dwork pencil. We have

$$Z(\mathcal{X}, t) = \frac{R_1(pt)^{20}R_2(pt)^{30}Q(t)}{(1-t)(1-pt)(1-p^2t)(1-p^3t)}$$

where R_1 and R_2 are the numerators of the zeta functions of certain curves given by a formula of Candelas–de la Ossa–Rodríguez Villegas [?].

As it is presented, we would work with $\mathbb{P}_{\Delta} = \mathbb{P}^4$ where Δ is the 4-simplex of volume $625/24$. As in Example 6.1, the monomials of the equation generate a sublattice of index 5^3 in \mathbb{Z}^4 , so we may instead work with a polytope whose volume is smaller by a factor of 5^3 . For $p = 2^{20} - 3$, we compute the “interesting” factor

$$Q(t) = 1 - 1576492860t^1 + 2672053179370pt^2 - 1576492860p^3t^3 + p^6t^4$$

in 11m18s; if we instead had tried to apply the *controlled AKR algorithm* to compute $Q(t)$ (and not the other factors) we extrapolate that it would take us at least 120 days. We present the running times for other p in Table 5.

Since this is a “small” perturbation of the Fermat threefold, we again attempted to confirm these results using the *deformation method*; however, this was again hampered by the fact that the Pancratz–Tuitman implementation works in \mathbb{P}_{Δ} instead of \mathbb{P}^3 . For $p = 7$, it took 5h4m and its memory footprint was around 12GB.

p	time	p	time	p	time
$2^8 - 5$	0.73s	$2^{13} - 1$	6.41s	$2^{18} - 5$	2m50s
$2^9 - 3$	0.77s	$2^{14} - 3$	11.61s	$2^{19} - 1$	5m38s
$2^{10} - 3$	0.80s	$2^{15} - 19$	21.98s	$2^{20} - 3$	11m18s
$2^{11} - 9$	2.54s	$2^{16} - 15$	43.07s	$2^{21} - 9$	22m41s
$2^{12} - 3$	3.80s	$2^{17} - 1$	1m25s	$2^{22} - 3$	52m37s

TABLE 5. Running times for Example 7.1.

Example 7.2. Let \mathcal{X} be the threefold defined by

$$x_0^8 + x_1^5 x_2 + x_0^2 x_1^2 x_2 x_3 + x_1 x_2^3 x_3 + x_1^2 x_3^3 + x_0 x_1 x_2 x_3 x_4 + x_2 x_3 x_4^2$$

in the weighted projective space $\mathbb{P}(1, 14, 18, 20, 25)$. The Newton polytope has volume $11/3 \approx 3.67$; by changing the lattice we may instead work with a polytope of volume $1/3 \approx 0.33$. In this setting, the Hodge numbers of $PH^3(\mathcal{X})$ are $(1, 1, 1, 1)$.

For $p = 2^{20} - 3$, we compute the “interesting” factor of $Z(\mathcal{X}, t)$

$$1 - 618297672t^1 + 390956360946pt^2 - 618297672p^3t^3 + p^6t^4$$

in 32m33s. We present the running times for other p in Table 6.

p	time	p	time	p	time
$2^8 - 5$	1.90s	$2^{13} - 1$	18.2s	$2^{18} - 5$	8m0s
$2^9 - 3$	1.96s	$2^{14} - 3$	32.9s	$2^{19} - 1$	16m8s
$2^{10} - 3$	2.06s	$2^{15} - 19$	1m6s	$2^{20} - 3$	32m33s
$2^{11} - 9$	7.48s	$2^{16} - 15$	2m4s	$2^{21} - 9$	1h5m
$2^{12} - 3$	10.9s	$2^{17} - 1$	4m3s	$2^{22} - 3$	2h23m

TABLE 6. Running times for Example 7.2.

Example 7.3. Let \mathcal{X} be the threefold defined by

$$x_1^7 + x_0^5 x_1 x_2 + x_0^2 x_1^2 x_2 x_3 + x_0^4 x_2 x_4 + x_0 x_2^3 x_3 + x_0^2 x_3^3 + x_0 x_1 x_2 x_3 x_4 + x_2 x_3 x_4^2$$

in the weighted projective space $\mathbb{P}(10, 11, 16, 19, 21)$. Again, by choosing the right lattice, we reduce the volume of the Newton polytope from $55/12 \approx 4.58$ to $11/24 \approx 0.46$, and the Hodge numbers of $PH^3(\mathcal{X})$ are $(1, 2, 2, 1)$. For $p = 2^{20} - 3$, we computed the “interesting” factor of $Z(\mathcal{X}, t)$

$$1 - 2068001468t^1 + 3449674041773pt^2 - 3772715295733197p^2t^3 \\ + 3449674041773p^4t^4 - 2068001468p^6t^5 + p^9t^6$$

in 2h10m. We present the running times for other p in Table 7.

p	time	p	time	p	time
$2^8 - 5$	4.47s	$2^{13} - 1$	1m8s	$2^{18} - 5$	32m25s
$2^9 - 3$	4.60s	$2^{14} - 3$	2m8s	$2^{19} - 1$	1h5m
$2^{10} - 3$	4.96s	$2^{15} - 19$	4m6s	$2^{20} - 3$	2h10m
$2^{11} - 9$	25.8s	$2^{16} - 15$	8m18s	$2^{21} - 9$	4h17m
$2^{12} - 3$	39.1s	$2^{17} - 1$	16m31s	$2^{22} - 3$	9h33m

TABLE 7. Running times for Example 7.3.

Example 7.4. Let \mathcal{X} be the closure in \mathbb{P}_Δ (which is not a weighted projective space) of the threefold defined by the Laurent polynomial

$$xyz^2w^3 + x + y + z - 1 + y^{-1}z^{-1} + x^{-2}y^{-1}z^{-2}w^{-3} = 0.$$

Choosing the correct lattice reduces the volume of the Newton polytope from $9/8 \approx 1.12$ to $3/8 \approx 0.38$, and the Hodge numbers of $PH^3(\mathcal{X})$ are $(1, 2, 2, 1)$. For $p = 2^{20} - 3$, we computed the “interesting” factor of $Z(\mathcal{X}, t)$

$$(1+718pt+p^3t^2) \cdot (1+1188466826t^1+1915150034310pt^2+1188466826p^3t^3+p^6t^4)$$

in 1h15m. We present the running times for other p in Table 8.

p	time	p	time	p	time
$2^8 - 5$	2.74s	$2^{13} - 1$	39.28s	$2^{18} - 5$	18m34s
$2^9 - 3$	2.80s	$2^{14} - 3$	1m13s	$2^{19} - 1$	38m8s
$2^{10} - 3$	3.00s	$2^{15} - 19$	1m21s	$2^{20} - 3$	1h15m
$2^{11} - 9$	14.86s	$2^{16} - 15$	4m45s	$2^{21} - 9$	2h32m
$2^{12} - 3$	22.32s	$2^{17} - 1$	9m12s	$2^{22} - 3$	5h39m

TABLE 8. Running times for Example 7.4.

8. CUBIC FOURFOLDS

For our final example, we consider a cubic fourfold. For X a smooth cubic fourfold in \mathbb{P}^5 , $\dim H^4(X) = 23$ and the Hodge numbers are $(0, 1, 21, 1, 0)$.

In this example, we worked modulo p^6 in order to obtain $Q(t)$.

Example 8.1. Let \mathcal{X} be the smooth projective cubic fourfold in $\mathbb{P}_{\mathbb{F}_p}^5$ defined by

$$x_0^3 + x_1^3 + x_2^3 + (x_0 + x_1 + 2x_2)^3 + x_3^3 + x_4^3 + x_5^3 + 2(x_0 + x_3)^3 + 3(x_1 + x_4)^3 + (x_2 + x_5)^3;$$

it is nondegenerate in \mathbb{P}^5 . For $p = 31$, in 21h31m we computed

$$Z(\mathcal{X}, t)^{-1} = (1 - t)(1 - pt)(1 - p^2t)(1 - p^3t)(1 - p^4t)Q(t)$$

where the “interesting” factor is an irreducible Weil polynomial given by

$$pQ(t/p^2) = p - 7t^1 + 21t^2 - 52t^3 - 8t^4 - 28t^5 + 21t^6 + 35t^7 + 39t^9 + 62t^{10} + 23t^{11} + 62t^{12} + 39t^{13} + 35t^{15} + 21t^{16} - 28t^{17} - 8t^{18} - 52t^{19} + 21t^{20} - 7t^{21} + pt^{22};$$

the coefficient of t^1 may be confirmed independently by counting $\mathcal{X}(\mathbb{F}_p)$ using the Sage function `count_points`. For $p = 127$ the running time was 23h15m and for $p = 499$ it was 24h55m; in both cases, the “interesting” factor is an irreducible Weil polynomial. In these computations, the memory footprint was around 36.5GB.

In dimension 4, the bottleneck seems to be the linear algebra required to set up controlled reduction. For $p = 31$, more than half of the runtime (15h32m) is spent solving a linear problem of size $15,504 \times 37,128$ modulo p^6 . With careful handling of this step (e.g., avoiding Hensel lifts) we would expect a significant speedup.

Note that the defining equation for \mathcal{X} is quite sparse. To assess the effect of this sparsity, as well as to cross-check the answer, we recomputed $Z(\mathcal{X}, t)$ after applying a random linear change of variables to obtain a dense defining equation. For $p = 31$, in 27h55m and using about 41GB we obtained the same value for $Z(\mathcal{X}, t)$ as above.

Recall from the introduction that a cubic fourfold is *coplanar* if it is defined by an expression $\sum_{i=1}^{10} a_i^3$ in which each a_i is a linear form and some four of the a_i are linearly dependent. Ranestad–Voisin [?] show that the Zariski closure D_{copl} of the coplanar locus is a divisor on the moduli space of cubic fourfolds. Example 8.1 is a non-special coplanar cubic fourfold: the existence of a primitive codimension-2 cycle class would imply² that $pQ(t/p^2)$ has a cyclotomic factor. This shows (modulo detailed analysis of the algorithm) that D_{copl} is not a Noether–Lefschetz divisor.

DEPARTMENT OF MATHEMATICS, DARTMOUTH COLLEGE, 6188 KEMENY HALL, HANOVER, NH 03755, USA

Email address: edgarcosta@math.dartmouth.edu

URL: <http://www.math.dartmouth.edu/~edgarcosta/>

SCHOOL OF MATHEMATICS AND STATISTICS, UNIVERSITY OF NEW SOUTH WALES, SYDNEY NSW 2052, AUSTRALIA

Email address: d.harvey@unsw.edu.au

URL: <http://web.maths.unsw.edu.au/~davidharvey/>

UNIV. OF CALIFORNIA, SAN DIEGO, 9500 GILMAN DRIVE #0112, LA JOLLA, CA 92093 USA

Email address: kedlaya@ucsd.edu

URL: <http://kskedlaya.org/>

²While it is not needed here, the Tate conjecture for ordinary cubic fourfolds is known [?].