A census of zeta functions of quartic K3 surfaces over $\mathbb{F}_2$

Kiran S. Kedlaya and Andrew V. Sutherland

Department of Mathematics, University of California, San Diego
Department of Mathematics, Massachusetts Institute of Technology
kedlaya@ucsd.edu, drew@math.mit.edu
http://kskedlaya.org/slides/

ANTS-XII: Twelfth Algorithmic Number Theory Symposium
University of Kaiserslautern
August 30, 2016

Kedlaya was supported by NSF grant DMS-1501214, UCSD (Warschawski chair), and a Guggenheim Fellowship.
Sutherland was supported by NSF grants DMS-1115455 and DMS-1522526.

Contents

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

K3 surfaces

Throughout, let K be a field and let X be a *K3 surface* over K, i.e., a geometrically connected, projective variety of dimension 2 such that:

- the canonical bundle $\Omega_{X/K} = \wedge^2 \Omega^1_{X/K}$ is trivial;
- X is not an abelian surface.

Some classes of examples:

- a smooth quartic surface in $\mathbb{P}^3_K$;
- a double cover of $\mathbb{P}^2_K$ branched over a smooth sextic curve;
- a transverse intersection of a smooth quadric and cubic in $\mathbb{P}^4_K$;
- a transverse intersection of three smooth quadrics in $\mathbb{P}^5_K$;
- an elliptic K3 surface.

From the point of view of geometry and arithmetic, K3 surfaces are strongly analogous to elliptic curves. (This analogy extends to *Calabi-Yau threefolds*, but we don't discuss these here.)

Zeta functions of K3 surfaces: initial constraints

Hereafter, assume $K := \mathbb{F}_q$ is finite. By the Weil conjectures and properties of crystalline cohomology[1], the zeta function of X has the form

$$\zeta(X, T) = \frac{1}{(1-T)(1-qT)(1-q^2T)q^{-1}L(qT)}$$

where for some $a_1, \dots, a_{10} \in \mathbb{Z}$ we have

$$L(T) = q + a_1 T + \cdots + a_{10} T^{10} \pm (a_{10} T^{11} + \cdots + a_1 T^{20} + q T^{21})$$

and the roots of L in $\mathbb{C}$ lie on the unit circle. Hence for a given q, these constraints limit $\zeta(X, T)$ to a computable finite set.

In addition to these initial constraints, we also have (for $q \leq 17$) monotonicity constraints like $\#X(\mathbb{F}_{q^2}) \geq \#X(\mathbb{F}_q) \geq 0$, and (for all q) arithmetic constraints derived from Brauer groups as described next.

[1] For $q \neq p$, the p-adic conditions are slightly stronger than stated here.

Zeta functions of K3 surfaces: initial constraints

Hereafter, assume $K := \mathbb{F}_q$ is finite. By the Weil conjectures and properties of crystalline cohomology[1], the zeta function of X has the form

$$\zeta(X, T) = \frac{1}{(1-T)(1-qT)(1-q^2T)q^{-1}L(qT)}$$

where for some $a_1, \ldots, a_{10} \in \mathbb{Z}$ we have

$$L(T) = q + a_1T + \cdots + a_{10}T^{10} \pm (a_{10}T^{11} + \cdots + a_1T^{20} + qT^{21})$$

and the roots of L in $\mathbb{C}$ lie on the unit circle. Hence for a given q, these constraints limit $\zeta(X, T)$ to a computable finite set.

In addition to these initial constraints, we also have (for $q \leq 17$) monotonicity constraints like $\#X(\mathbb{F}_{q^2}) \geq \#X(\mathbb{F}_q) \geq 0$, and (for all q) arithmetic constraints derived from Brauer groups as described next.

[1] For $q \neq p$, the p-adic conditions are slightly stronger than stated here.

Base extension and monotonicity

For $n > 1$, the base extension X_n of X from $\mathbb{F}_q$ to $\mathbb{F}_{q^n}$ has zeta function

$$\zeta(X_n, T) = \frac{1}{(1-T)(1-q^nT)(1-q^{2n}T)q^{-n}L_n(q^nT)}$$

where L_n is the polynomial obtained from L by raising each root to the n-th power. That is, there exist $\alpha_1, \ldots, \alpha_{21} \in \mathbb{C}$ such that

$$L(T) = q\prod_{i=1}^{21}(1-\alpha_i T), \qquad L_n(T) = q^n\prod_{i=1}^{21}(1-\alpha_i^n T).$$

In particular, L_n is uniquely determined by L; for example,

$$L_2(T^2) = L(T)L(-T).$$

Base extension and monotonicity

For $n > 1$, the base extension X_n of X from $\mathbb{F}_q$ to $\mathbb{F}_{q^n}$ has zeta function

$$\zeta(X_n, T) = \frac{1}{(1-T)(1-q^nT)(1-q^{2n}T)q^{-n}L_n(q^nT)}$$

where L_n is the polynomial obtained from L by raising each root to the n-th power. That is, there exist $\alpha_1, \ldots, \alpha_{21} \in \mathbb{C}$ such that

$$L(T) = q\prod_{i=1}^{21}(1-\alpha_i T), \qquad L_n(T) = q^n\prod_{i=1}^{21}(1-\alpha_i^n T).$$

In particular, L_n is uniquely determined by L; for example,

$$L_2(T^2) = L(T)L(-T).$$

Brauer groups and K3 zeta functions

Factor $L(T)$ as $(1-T)^{r-1}L_1(T)$ with $L_1(1) \neq 0$. Under the Tate conjecture[2], r equals the rank of the Néron-Severi lattice $\mathrm{NS}(X)$.

The Artin-Tate formula states that

$$L_1(1) = |\Delta| \, \# \, \mathrm{Br}\, X$$

where Δ is the discriminant of $\mathrm{NS}(X)$ and $\mathrm{Br}\, X$ is the Brauer group of X. The latter is finite and its order is a perfect square.

Even without knowledge of Δ (or even the Tate conjecture), one can compare the Artin-Tate formulas over $\mathbb{F}_q$ and $\mathbb{F}_{q^2}$ to deduce that $L_1(-1)$ is a (possibly zero) perfect square (Elsenhans-Jahnel).

[2] This is apparently now unconditional: the last missing cases in characteristic 2 are handled by Madapusi Pera–Kim, arXiv:1512.02540.

Brauer groups and K3 zeta functions

Factor $L(T)$ as $(1-T)^{r-1}L_1(T)$ with $L_1(1) \neq 0$. Under the Tate conjecture[2], r equals the rank of the Néron-Severi lattice $\mathrm{NS}(X)$.

The Artin-Tate formula states that

$$L_1(1) = |\Delta| \, \# \operatorname{Br} X$$

where Δ is the discriminant of $\mathrm{NS}(X)$ and $\operatorname{Br} X$ is the Brauer group of X. The latter is finite and its order is a perfect square.

Even without knowledge of Δ (or even the Tate conjecture), one can compare the Artin-Tate formulas over $\mathbb{F}_q$ and $\mathbb{F}_{q^2}$ to deduce that $L_1(-1)$ is a (possibly zero) perfect square (Elsenhans-Jahnel).

[2] This is apparently now unconditional: the last missing cases in characteristic 2 are handled by Madapusi Pera–Kim, arXiv:1512.02540.

Brauer groups and K3 zeta functions

Factor $L(T)$ as $(1-T)^{r-1}L_1(T)$ with $L_1(1) \neq 0$. Under the Tate conjecture[2], r equals the rank of the Néron-Severi lattice $\mathrm{NS}(X)$.

The Artin-Tate formula states that

$$L_1(1) = |\Delta| \,\#\, \mathrm{Br}\, X$$

where Δ is the discriminant of $\mathrm{NS}(X)$ and $\mathrm{Br}\, X$ is the Brauer group of X. The latter is finite and its order is a perfect square.

Even without knowledge of Δ (or even the Tate conjecture), one can compare the Artin-Tate formulas over $\mathbb{F}_q$ and $\mathbb{F}_{q^2}$ to deduce that $L_1(-1)$ is a (possibly zero) perfect square (Elsenhans-Jahnel).

[2] This is apparently now unconditional: the last missing cases in characteristic 2 are handled by Madapusi Pera–Kim, arXiv:1512.02540.

The inverse problem for K3 zeta functions

For given q, d, the Honda-Tate theorem specifies which rational functions occur as zeta functions of a d-dimensional abelian variety over $\mathbb{F}_q$.

What about for K3 surfaces? There is a partial analogue of Honda-Tate due to Taelman (to be stated later), but for various reasons it does not tell the whole story.

As a complement, we make a detailed numerical study of the case $q = 2$. For practical reasons, we limit ourselves to smooth quartics; this is a serious limitation from the point of view of zeta functions, but nonetheless we obtain a "reasonable" class in which every eligible candidate actually occurs for some K3 surface.

The inverse problem for K3 zeta functions

For given q, d, the Honda-Tate theorem specifies which rational functions occur as zeta functions of a d-dimensional abelian variety over $\mathbb{F}_q$.

What about for K3 surfaces? There is a partial analogue of Honda-Tate due to Taelman (to be stated later), but for various reasons it does not tell the whole story.

As a complement, we make a detailed numerical study of the case $q = 2$. For practical reasons, we limit ourselves to smooth quartics; this is a serious limitation from the point of view of zeta functions, but nonetheless we obtain a "reasonable" class in which every eligible candidate actually occurs for some K3 surface.

The inverse problem for K3 zeta functions

For given q, d, the Honda-Tate theorem specifies which rational functions occur as zeta functions of a d-dimensional abelian variety over $\mathbb{F}_q$.

What about for K3 surfaces? There is a partial analogue of Honda-Tate due to Taelman (to be stated later), but for various reasons it does not tell the whole story.

As a complement, we make a detailed numerical study of the case $q = 2$. For practical reasons, we limit ourselves to smooth quartics; this is a serious limitation from the point of view of zeta functions, but nonetheless we obtain a "reasonable" class in which every eligible candidate actually occurs for some K3 surface.

Contents

Rolle's theorem as algorithm

We first enumerate the candidates for L consistent with the Weil/crystalline constraints (imposing the others as filters). This resembles finding Pisot/Salem numbers, number fields of fixed signature, etc.

The constraints amount to the existence of a presentation

$$L(T) = (1 \pm T)T^{10}Q(T + T^{-1}),\ Q(T) = qT^{10} + b_1 T^9 + \cdots + b_{10}$$

where all roots of Q in $\mathbb{C}$ are real and in $[-2, 2]$. For a given choice of $\pm$, the transformation between the a_i and b_i is unipotent and integral.

If Q has roots in $[-2, 2]$, then by Rolle's theorem the same is true of

$$\frac{1}{k!}Q^{(k)}(T) = \sum_{i=0}^{10-k} \binom{i+k}{i} b_{10-i-k} T^i \qquad (k = 1, \ldots, 10).$$

It is thus natural to enumerate candidates recursively: given $b_i, \ldots, b_i$, find all b_{i+1} consistent with Rolle's theorem *and other known constraints*.

Rolle's theorem as algorithm

We first enumerate the candidates for L consistent with the Weil/crystalline constraints (imposing the others as filters). This resembles finding Pisot/Salem numbers, number fields of fixed signature, etc.

The constraints amount to the existence of a presentation

$$L(T) = (1 \pm T)T^{10}Q(T + T^{-1}), \; Q(T) = qT^{10} + b_1T^9 + \cdots + b_{10}$$

where all roots of Q in $\mathbb{C}$ are real and in $[-2, 2]$. For a given choice of $\pm$, the transformation between the a_i and b_i is unipotent and integral.

If Q has roots in $[-2, 2]$, then by Rolle's theorem the same is true of

$$\frac{1}{k!}Q^{(k)}(T) = \sum_{i=0}^{10-k} \binom{i+k}{i} b_{10-i-k} T^i \qquad (k = 1, \dots, 10).$$

It is thus natural to enumerate candidates recursively: given $b_i, \dots, b_i$, find all b_{i+1} consistent with Rolle's theorem *and other known constraints*.

Rolle's theorem as algorithm

We first enumerate the candidates for L consistent with the Weil/crystalline constraints (imposing the others as filters). This resembles finding Pisot/Salem numbers, number fields of fixed signature, etc.

The constraints amount to the existence of a presentation

$$L(T) = (1 \pm T)T^{10}Q(T + T^{-1}), \; Q(T) = qT^{10} + b_1 T^9 + \cdots + b_{10}$$

where all roots of Q in $\mathbb{C}$ are real and in $[-2, 2]$. For a given choice of $\pm$, the transformation between the a_i and b_i is unipotent and integral.

If Q has roots in $[-2, 2]$, then by Rolle's theorem the same is true of

$$\frac{1}{k!}Q^{(k)}(T) = \sum_{i=0}^{10-k} \binom{i+k}{i} b_{10-i-k} T^i \qquad (k = 1, \dots, 10).$$

It is thus natural to enumerate candidates recursively: given $b_i, \dots, b_i$, find all b_{i+1} consistent with Rolle's theorem *and other known constraints*.

Rolle's theorem as algorithm

We first enumerate the candidates for L consistent with the Weil/crystalline constraints (imposing the others as filters). This resembles finding Pisot/Salem numbers, number fields of fixed signature, etc.

The constraints amount to the existence of a presentation

$$L(T) = (1 \pm T) T^{10} Q(T + T^{-1}), \ Q(T) = qT^{10} + b_1 T^9 + \cdots + b_{10}$$

where all roots of Q in $\mathbb{C}$ are real and in $[-2, 2]$. For a given choice of $\pm$, the transformation between the a_i and b_i is unipotent and integral.

If Q has roots in $[-2, 2]$, then by Rolle's theorem the same is true of

$$\frac{1}{k!} Q^{(k)}(T) = \sum_{i=0}^{10-k} \binom{i+k}{i} b_{10-i-k} T^i \qquad (k = 1, \ldots, 10).$$

It is thus natural to enumerate candidates recursively: given $b_i, \ldots, b_i$, find all b_{i+1} consistent with Rolle's theorem *and other known constraints*.

Other known constraints

Based on work done by one of us (KK) in 2008, we add some constraints.

- From $b_1, \dots, b_i$, we can compute the first i power sums of either L or Q. Every power sum of L has absolute value at most 21.
- We impose classical symmetric inequalities on the roots of Q.
- We obtain more constraints from Descartes's rule of signs.
- We use Sturm sequences to count roots in intervals; this leads to some additional "lookahead" constraints. (It is unclear whether VCA root isolation would be better here.)

The implementation is also improved from 2008; it uses `Sage` for user-facing code, `FLINT` for low-level operations, and `Cython` in between.

Fewer than 1% of the ends of the search tree lead to solutions. This suggests that there is still significant room for further improvement.

Other known constraints

Based on work done by one of us (KK) in 2008, we add some constraints.

- From $b_1, \ldots, b_i$, we can compute the first i power sums of either L or Q. Every power sum of L has absolute value at most 21.
- We impose classical symmetric inequalities on the roots of Q.
- We obtain more constraints from Descartes's rule of signs.
- We use Sturm sequences to count roots in intervals; this leads to some additional "lookahead" constraints. (It is unclear whether VCA root isolation would be better here.)

The implementation is also improved from 2008; it uses `Sage` for user-facing code, `FLINT` for low-level operations, and `Cython` in between.

Fewer than 1% of the ends of the search tree lead to solutions. This suggests that there is still significant room for further improvement.

Other known constraints

Based on work done by one of us (KK) in 2008, we add some constraints.

- From $b_1, \ldots, b_i$, we can compute the first i power sums of either L or Q. Every power sum of L has absolute value at most 21.
- We impose classical symmetric inequalities on the roots of Q.
- We obtain more constraints from Descartes's rule of signs.
- We use Sturm sequences to count roots in intervals; this leads to some additional "lookahead" constraints. (It is unclear whether VCA root isolation would be better here.)

The implementation is also improved from 2008; it uses `Sage` for user-facing code, `FLINT` for low-level operations, and `Cython` in between.

Fewer than 1% of the ends of the search tree lead to solutions. This suggests that there is still significant room for further improvement.

Other known constraints

Based on work done by one of us (KK) in 2008, we add some constraints.

- From $b_1, \ldots, b_i$, we can compute the first i power sums of either L or Q. Every power sum of L has absolute value at most 21.
- We impose classical symmetric inequalities on the roots of Q.
- We obtain more constraints from Descartes's rule of signs.
- We use Sturm sequences to count roots in intervals; this leads to some additional "lookahead" constraints. (It is unclear whether VCA root isolation would be better here.)

The implementation is also improved from 2008; it uses `Sage` for user-facing code, `FLINT` for low-level operations, and `Cython` in between.

Fewer than 1% of the ends of the search tree lead to solutions. This suggests that there is still significant room for further improvement.

Other known constraints

Based on work done by one of us (KK) in 2008, we add some constraints.

- From $b_1, \ldots, b_i$, we can compute the first i power sums of either L or Q. Every power sum of L has absolute value at most 21.
- We impose classical symmetric inequalities on the roots of Q.
- We obtain more constraints from Descartes's rule of signs.
- We use Sturm sequences to count roots in intervals; this leads to some additional "lookahead" constraints. (It is unclear whether VCA root isolation would be better here.)

The implementation is also improved from 2008; it uses `Sage` for user-facing code, `FLINT` for low-level operations, and `Cython` in between.

Fewer than 1% of the ends of the search tree lead to solutions. This suggests that there is still significant room for further improvement.

Other known constraints

Based on work done by one of us (KK) in 2008, we add some constraints.

- From $b_1, \ldots, b_i$, we can compute the first i power sums of either L or Q. Every power sum of L has absolute value at most 21.
- We impose classical symmetric inequalities on the roots of Q.
- We obtain more constraints from Descartes's rule of signs.
- We use Sturm sequences to count roots in intervals; this leads to some additional "lookahead" constraints. (It is unclear whether VCA root isolation would be better here.)

The implementation is also improved from 2008; it uses `Sage` for user-facing code, `FLINT` for low-level operations, and `Cython` in between.

Fewer than 1% of the ends of the search tree lead to solutions. This suggests that there is still significant room for further improvement.

Other known constraints

Based on work done by one of us (KK) in 2008, we add some constraints.

- From $b_1, \ldots, b_i$, we can compute the first i power sums of either L or Q. Every power sum of L has absolute value at most 21.
- We impose classical symmetric inequalities on the roots of Q.
- We obtain more constraints from Descartes's rule of signs.
- We use Sturm sequences to count roots in intervals; this leads to some additional "lookahead" constraints. (It is unclear whether VCA root isolation would be better here.)

The implementation is also improved from 2008; it uses `Sage` for user-facing code, `FLINT` for low-level operations, and `Cython` in between.

Fewer than 1% of the ends of the search tree lead to solutions. This suggests that there is still significant room for further improvement.

Results for $q = 2$

For $q = 2$, there are 2,971,182 polynomials L satisfying the initial constraints. Of these, 2,195,801 also satisfy the Elsenhans-Jahnel constraint. Of these, 1,672,565 also satisfy monotonicity.

This computation required less than 1 hour on a 24-core machine. We used 512 threads to enumerate the search tree in parallel, using randomized work-stealing.

QA note: we used the same implementation to find monic polynomials of degree 2, 4, . . . , 20 with all roots on the unit circle. Since these must be products of cyclotomic polynomials (Kronecker's theorem), the result can be checked independently.

Results for $q = 2$

For $q = 2$, there are 2,971,182 polynomials L satisfying the initial constraints. Of these, 2,195,801 also satisfy the Elsenhans-Jahnel constraint. Of these, 1,672,565 also satisfy monotonicity.

This computation required less than 1 hour on a 24-core machine. We used 512 threads to enumerate the search tree in parallel, using randomized work-stealing.

QA note: we used the same implementation to find monic polynomials of degree 2, 4, . . . , 20 with all roots on the unit circle. Since these must be products of cyclotomic polynomials (Kronecker's theorem), the result can be checked independently.

Results for $q = 2$

For $q = 2$, there are 2,971,182 polynomials L satisfying the initial constraints. Of these, 2,195,801 also satisfy the Elsenhans-Jahnel constraint. Of these, 1,672,565 also satisfy monotonicity.

This computation required less than 1 hour on a 24-core machine. We used 512 threads to enumerate the search tree in parallel, using randomized work-stealing.

QA note: we used the same implementation to find monic polynomials of degree 2, 4, . . . , 20 with all roots on the unit circle. Since these must be products of cyclotomic polynomials (Kronecker's theorem), the result can be checked independently.

Contents

Enumerating smooth quartic surfaces

We next enumerate quartic surfaces in $\mathbb{P}^3_{\mathbb{F}_q}$ up to PGL_4-equivalence.

- Identify each homogeneous quartic $f \in \mathbb{F}_q[w,x,y,z]$ with a vector $v(f) \in V := \mathbb{F}_q^{35}$ (the $\binom{7}{3} = 35$ monomial f form a basis for V) .
- Identify $\mathrm{PGL}_4(\mathbb{F}_q)$ with $G \subseteq \mathrm{GL}_{35}(\mathbb{F}_q)$. For $q = 2$ we find that $\#G = 20{,}160$ and V has 1,732,564 G-orbits (by Burnside's lemma).
- Using a bitmap M indexed by V we can determine a minimal representative for each G-orbit by simply enumerating orbits.
- Take the index $v(f) \in V$ of the first unmarked bit in M as the representative of its G-orbit, mark all the bits in this G-orbit, repeat.
- Eliminate G-orbit reps $v(f)$ for which the singular locus defined by the Jacobian matrix of f is nonempty.
- For $q = 2$ we get 528,257 PGL_4-inequivalent f defining K3 surfaces.

Enumerating smooth quartic surfaces

We next enumerate quartic surfaces in $\mathbb{P}^3_{\mathbb{F}_q}$ up to PGL_4-equivalence.

- Identify each homogeneous quartic $f \in \mathbb{F}_q[w, x, y, z]$ with a vector $v(f) \in V := \mathbb{F}_q^{35}$ (the $\binom{7}{3} = 35$ monomial f form a basis for V) .
- Identify $\mathrm{PGL}_4(\mathbb{F}_q)$ with $G \subseteq \mathrm{GL}_{35}(\mathbb{F}_q)$. For $q = 2$ we find that $\#G = 20{,}160$ and V has 1,732,564 G-orbits (by Burnside's lemma).
- Using a bitmap M indexed by V we can determine a minimal representative for each G-orbit by simply enumerating orbits.
- Take the index $v(f) \in V$ of the first unmarked bit in M as the representative of its G-orbit, mark all the bits in this G-orbit, repeat.
- Eliminate G-orbit reps $v(f)$ for which the singular locus defined by the Jacobian matrix of f is nonempty.
- For $q = 2$ we get 528,257 PGL_4-inequivalent f defining K3 surfaces.

Enumerating smooth quartic surfaces

We next enumerate quartic surfaces in $\mathbb{P}^3_{\mathbb{F}_q}$ up to PGL_4-equivalence.

- Identify each homogeneous quartic $f \in \mathbb{F}_q[w,x,y,z]$ with a vector $v(f) \in V := \mathbb{F}_q^{35}$ (the $\binom{7}{3} = 35$ monomial f form a basis for V) .
- Identify $\mathrm{PGL}_4(\mathbb{F}_q)$ with $G \subseteq \mathrm{GL}_{35}(\mathbb{F}_q)$. For $q = 2$ we find that $\#G = 20{,}160$ and V has 1,732,564 G-orbits (by Burnside's lemma).
- Using a bitmap M indexed by V we can determine a minimal representative for each G-orbit by simply enumerating orbits.
- Take the index $v(f) \in V$ of the first unmarked bit in M as the representative of its G-orbit, mark all the bits in this G-orbit, repeat.
- Eliminate G-orbit reps $v(f)$ for which the singular locus defined by the Jacobian matrix of f is nonempty.
- For $q = 2$ we get 528,257 PGL_4-inequivalent f defining K3 surfaces.

Enumerating smooth quartic surfaces

We next enumerate quartic surfaces in $\mathbb{P}^3_{\mathbb{F}_q}$ up to PGL_4-equivalence.

- Identify each homogeneous quartic $f \in \mathbb{F}_q[w,x,y,z]$ with a vector $v(f) \in V := \mathbb{F}_q^{35}$ (the $\binom{7}{3} = 35$ monomial f form a basis for V) .
- Identify $\mathrm{PGL}_4(\mathbb{F}_q)$ with $G \subseteq \mathrm{GL}_{35}(\mathbb{F}_q)$. For $q = 2$ we find that $\#G = 20{,}160$ and V has 1,732,564 G-orbits (by Burnside's lemma).
- Using a bitmap M indexed by V we can determine a minimal representative for each G-orbit by simply enumerating orbits.
- Take the index $v(f) \in V$ of the first unmarked bit in M as the representative of its G-orbit, mark all the bits in this G-orbit, repeat.
- Eliminate G-orbit reps $v(f)$ for which the singular locus defined by the Jacobian matrix of f is nonempty.
- For $q = 2$ we get 528,257 PGL_4-inequivalent f defining K3 surfaces.

Enumerating smooth quartic surfaces

We next enumerate quartic surfaces in $\mathbb{P}^3_{\mathbb{F}_q}$ up to PGL_4-equivalence.

- Identify each homogeneous quartic $f \in \mathbb{F}_q[w,x,y,z]$ with a vector $v(f) \in V := \mathbb{F}_q^{35}$ (the $\binom{7}{3} = 35$ monomial f form a basis for V) .
- Identify $\mathrm{PGL}_4(\mathbb{F}_q)$ with $G \subseteq \mathrm{GL}_{35}(\mathbb{F}_q)$. For $q = 2$ we find that $\#G = 20{,}160$ and V has 1,732,564 G-orbits (by Burnside's lemma).
- Using a bitmap M indexed by V we can determine a minimal representative for each G-orbit by simply enumerating orbits.
- Take the index $v(f) \in V$ of the first unmarked bit in M as the representative of its G-orbit, mark all the bits in this G-orbit, repeat.
- Eliminate G-orbit reps $v(f)$ for which the singular locus defined by the Jacobian matrix of f is nonempty.
- For $q = 2$ we get 528,257 PGL_4-inequivalent f defining K3 surfaces.

Enumerating smooth quartic surfaces

We next enumerate quartic surfaces in $\mathbb{P}^3_{\mathbb{F}_q}$ up to PGL_4-equivalence.

- Identify each homogeneous quartic $f \in \mathbb{F}_q[w,x,y,z]$ with a vector $v(f) \in V := \mathbb{F}_q^{35}$ (the $\binom{7}{3} = 35$ monomial f form a basis for V) .
- Identify $\mathrm{PGL}_4(\mathbb{F}_q)$ with $G \subseteq \mathrm{GL}_{35}(\mathbb{F}_q)$. For $q = 2$ we find that $\#G = 20{,}160$ and V has 1,732,564 G-orbits (by Burnside's lemma).
- Using a bitmap M indexed by V we can determine a minimal representative for each G-orbit by simply enumerating orbits.
- Take the index $v(f) \in V$ of the first unmarked bit in M as the representative of its G-orbit, mark all the bits in this G-orbit, repeat.
- Eliminate G-orbit reps $v(f)$ for which the singular locus defined by the Jacobian matrix of f is nonempty.
- For $q = 2$ we get 528,257 PGL_4-inequivalent f defining K3 surfaces.

Enumerating smooth quartic surfaces

We next enumerate quartic surfaces in $\mathbb{P}^3_{\mathbb{F}_q}$ up to PGL_4-equivalence.

- Identify each homogeneous quartic $f \in \mathbb{F}_q[w,x,y,z]$ with a vector $v(f) \in V := \mathbb{F}_q^{35}$ (the $\binom{7}{3} = 35$ monomial f form a basis for V) .
- Identify $\mathrm{PGL}_4(\mathbb{F}_q)$ with $G \subseteq \mathrm{GL}_{35}(\mathbb{F}_q)$. For $q = 2$ we find that $\#G = 20{,}160$ and V has 1,732,564 G-orbits (by Burnside's lemma).
- Using a bitmap M indexed by V we can determine a minimal representative for each G-orbit by simply enumerating orbits.
- Take the index $v(f) \in V$ of the first unmarked bit in M as the representative of its G-orbit, mark all the bits in this G-orbit, repeat.
- Eliminate G-orbit reps $v(f)$ for which the singular locus defined by the Jacobian matrix of f is nonempty.
- For $q = 2$ we get 528,257 PGL_4-inequivalent f defining K3 surfaces.

Computing zeta functions of smooth quartics

Let S be our set of 528,257 K3 surfaces $X\colon f(w,x,y,z)=0$ over $\mathbb{F}_2$. For $X\in S$ we compute $\#X(\mathbb{F}_{2^n})$ for enough n to determine $L_X(T)$.

- Plan: for $x_0,y_0\in\mathbb{F}_{2n}$ and $f\in S$ count roots of $f(w,x_0,y_0,1)$ in $\mathbb{F}_{2^n}$ (also need to count solutions to $f(w,x,y,0)=0$, but this is easy).
- All but 34 $f\in S$ yields cubics we can write as $g(w)=w^3+aw+b$. There are only 2^{2n+1} such g, a lot less than $528{,}257\cdot 2^{2n}$.
- Precompute tables T_n indexed by (a,b) counting roots of w^3+aw+b in $\mathbb{F}_{2^n}$ using Zinoviev's formulas (this actually takes negligible time).
- Iterate over $x_0,y_0\in\mathbb{F}_{2^n}$, instantiate each of 35 quartic monomials at $x=x_0,y=y_0,z=1$, compute $f(w,x_0,y_0,1)$ for $f\in S$ as a linear combination, and look up the number of roots in T_n.
- Compute $\#X(\mathbb{F}_{2^n})$ for $X\in S$ and $n\le 12$. For $n=13,14,\ldots,19$, reduce S to the subset with $L_X(T)$ not yet determined[3] and repeat.

[3] The only ambiguity is $\pm$, which only resolves once we find a nonzero coefficient of L.

Computing zeta functions of smooth quartics

Let S be our set of 528,257 K3 surfaces $X\colon f(w,x,y,z)=0$ over $\mathbb{F}_2$. For $X\in S$ we compute $\#X(\mathbb{F}_{2^n})$ for enough n to determine $L_X(T)$.

- Plan: for $x_0, y_0 \in \mathbb{F}_{2n}$ and $f\in S$ count roots of $f(w,x_0,y_0,1)$ in $\mathbb{F}_{2^n}$ (also need to count solutions to $f(w,x,y,0)=0$, but this is easy).
- All but 34 $f\in S$ yields cubics we can write as $g(w)=w^3+aw+b$. There are only 2^{2n+1} such g, a lot less than $528{,}257\cdot 2^{2n}$.
- Precompute tables T_n indexed by (a,b) counting roots of w^3+aw+b in $\mathbb{F}_{2^n}$ using Zinoviev's formulas (this actually takes negligible time).
- Iterate over $x_0, y_0\in\mathbb{F}_{2^n}$, instantiate each of 35 quartic monomials at $x=x_0, y=y_0, z=1$, compute $f(w,x_0,y_0,1)$ for $f\in S$ as a linear combination, and look up the number of roots in T_n.
- Compute $\#X(\mathbb{F}_{2^n})$ for $X\in S$ and $n\le 12$. For $n=13,14,\ldots,19$, reduce S to the subset with $L_X(T)$ not yet determined[3] and repeat.

[3] The only ambiguity is $\pm$, which only resolves once we find a nonzero coefficient of L.

Computing zeta functions of smooth quartics

Let S be our set of 528,257 K3 surfaces $X\colon f(w,x,y,z)=0$ over $\mathbb{F}_2$. For $X\in S$ we compute $\#X(\mathbb{F}_{2^n})$ for enough n to determine $L_X(T)$.

- Plan: for $x_0,y_0\in\mathbb{F}_{2n}$ and $f\in S$ count roots of $f(w,x_0,y_0,1)$ in $\mathbb{F}_{2^n}$ (also need to count solutions to $f(w,x,y,0)=0$, but this is easy).
- All but 34 $f\in S$ yields cubics we can write as $g(w)=w^3+aw+b$. There are only 2^{2n+1} such g, a lot less than $528{,}257\cdot 2^{2n}$.
- Precompute tables T_n indexed by (a,b) counting roots of w^3+aw+b in $\mathbb{F}_{2^n}$ using Zinoviev's formulas (this actually takes negligible time).
- Iterate over $x_0,y_0\in\mathbb{F}_{2^n}$, instantiate each of 35 quartic monomials at $x=x_0,y=y_0,z=1$, compute $f(w,x_0,y_0,1)$ for $f\in S$ as a linear combination, and look up the number of roots in T_n.
- Compute $\#X(\mathbb{F}_{2^n})$ for $X\in S$ and $n\le 12$. For $n=13,14,\ldots,19$, reduce S to the subset with $L_X(T)$ not yet determined[3] and repeat.

[3] The only ambiguity is $\pm$, which only resolves once we find a nonzero coefficient of L.

Computing zeta functions of smooth quartics

Let S be our set of 528,257 K3 surfaces $X\colon f(w,x,y,z)=0$ over $\mathbb{F}_2$. For $X\in S$ we compute $\#X(\mathbb{F}_{2^n})$ for enough n to determine $L_X(T)$.

- Plan: for $x_0,y_0\in\mathbb{F}_{2n}$ and $f\in S$ count roots of $f(w,x_0,y_0,1)$ in $\mathbb{F}_{2^n}$ (also need to count solutions to $f(w,x,y,0)=0$, but this is easy).
- All but 34 $f\in S$ yields cubics we can write as $g(w)=w^3+aw+b$. There are only 2^{2n+1} such g, a lot less than $528{,}257\cdot 2^{2n}$.
- Precompute tables T_n indexed by (a,b) counting roots of w^3+aw+b in $\mathbb{F}_{2^n}$ using Zinoviev's formulas (this actually takes negligible time).
- Iterate over $x_0,y_0\in\mathbb{F}_{2^n}$, instantiate each of 35 quartic monomials at $x=x_0, y=y_0, z=1$, compute $f(w,x_0,y_0,1)$ for $f\in S$ as a linear combination, and look up the number of roots in T_n.
- Compute $\#X(\mathbb{F}_{2^n})$ for $X\in S$ and $n\le 12$. For $n=13,14,\ldots,19$, reduce S to the subset with $L_X(T)$ not yet determined[3] and repeat.

[3] The only ambiguity is $\pm$, which only resolves once we find a nonzero coefficient of L.

Computing zeta functions of smooth quartics

Let S be our set of 528,257 K3 surfaces $X\colon f(w,x,y,z)=0$ over $\mathbb{F}_2$. For $X \in S$ we compute $\#X(\mathbb{F}_{2^n})$ for enough n to determine $L_X(T)$.

- Plan: for $x_0, y_0 \in \mathbb{F}_{2n}$ and $f \in S$ count roots of $f(w,x_0,y_0,1)$ in $\mathbb{F}_{2^n}$ (also need to count solutions to $f(w,x,y,0)=0$, but this is easy).
- All but 34 $f \in S$ yields cubics we can write as $g(w)=w^3+aw+b$. There are only 2^{2n+1} such g, a lot less than $528{,}257 \cdot 2^{2n}$.
- Precompute tables T_n indexed by (a,b) counting roots of w^3+aw+b in $\mathbb{F}_{2^n}$ using Zinoviev's formulas (this actually takes negligible time).
- Iterate over $x_0, y_0 \in \mathbb{F}_{2^n}$, instantiate each of 35 quartic monomials at $x=x_0, y=y_0, z=1$, compute $f(w,x_0,y_0,1)$ for $f \in S$ as a linear combination, and look up the number of roots in T_n.
- Compute $\#X(\mathbb{F}_{2^n})$ for $X \in S$ and $n \le 12$. For $n=13,14,\ldots,19$, reduce S to the subset with $L_X(T)$ not yet determined[3] and repeat.

[3] The only ambiguity is $\pm$, which only resolves once we find a nonzero coefficient of L.

Computing zeta functions of smooth quartics

Let S be our set of 528,257 K3 surfaces $X\colon f(w,x,y,z)=0$ over $\mathbb{F}_2$. For $X\in S$ we compute $\#X(\mathbb{F}_{2^n})$ for enough n to determine $L_X(T)$.

- Plan: for $x_0,y_0\in\mathbb{F}_{2n}$ and $f\in S$ count roots of $f(w,x_0,y_0,1)$ in $\mathbb{F}_{2^n}$ (also need to count solutions to $f(w,x,y,0)=0$, but this is easy).
- All but 34 $f\in S$ yields cubics we can write as $g(w)=w^3+aw+b$. There are only 2^{2n+1} such g, a lot less than $528{,}257\cdot 2^{2n}$.
- Precompute tables T_n indexed by (a,b) counting roots of w^3+aw+b in $\mathbb{F}_{2^n}$ using Zinoviev's formulas (this actually takes negligible time).
- Iterate over $x_0,y_0\in\mathbb{F}_{2^n}$, instantiate each of 35 quartic monomials at $x=x_0, y=y_0, z=1$, compute $f(w,x_0,y_0,1)$ for $f\in S$ as a linear combination, and look up the number of roots in T_n.
- Compute $\#X(\mathbb{F}_{2^n})$ for $X\in S$ and $n\le 12$. For $n=13,14,\ldots,19$, reduce S to the subset with $L_X(T)$ not yet determined[3] and repeat.

[3] The only ambiguity is $\pm$, which only resolves once we find a nonzero coefficient of L.

Results for $q = 2$

For $q = 2$, on a 32-core machine it took about 2 days to enumerate the set S of PGL_4-inequivalent K3 surfaces X defined by smooth plane quartics, and about 2 weeks to compute $L(T)$ for all $X \in S$.

Most of the time was spent on the roughly 1000 cases in which we computed $\#X(\mathbb{F}_{2^n})$ with $n = 18, 19$.

We actually did more work than necessary (as a sanity check). For example, only 125 cases require $n = 19$, but we computed 283.

Important practical optimization: using Intel's PCLMULQDQ instruction ("carry-less" multiplication) sped up our implementation by a factor of 10.

We find that 52,755 distinct $L(T)$ arise among the 528,257 $X \in S$.

Contents

A theorem of Taelman

For X given, factor $L(T) = L_{\mathrm{alg}}(T)L_{\mathrm{trc}}(T)$ where L_{alg} is the product of all cyclotomic factors of L. (Aside: $\deg(L_{\mathrm{alg}}) + 1 = \operatorname{rank} \mathrm{NS}(X_{\overline{K}})$.)

Theorem (Taelman, 2016)

Assume[a] that K3 surfaces over finite extensions of $\mathbb{Q}_p$ admit potential semistable reduction. For q given, choose any L satisfying the initial constraints. Then for some positive integer n (and hence any multiple thereof), there is a K3 surface over $\mathbb{F}_{q^n}$ whose L_{trc} is the base extension of the one obtained from L.

[a]This hypothesis is made precise by Liedtke–Matsumoto (arxiv.1411.4797). It is is known for K3 surfaces of small degree relative to p.

Question: is it reasonable to hope that one can always take $n = 1$?

A theorem of Taelman

For X given, factor $L(T) = L_{\mathrm{alg}}(T)L_{\mathrm{trc}}(T)$ where L_{alg} is the product of all cyclotomic factors of L. (Aside: $\deg(L_{\mathrm{alg}}) + 1 = \operatorname{rank} \mathrm{NS}(X_{\overline{K}})$.)

Theorem (Taelman, 2016)

Assume[a] that K3 surfaces over finite extensions of $\mathbb{Q}_p$ admit potential semistable reduction. For q given, choose any L satisfying the initial constraints. Then for some positive integer n (and hence any multiple thereof), there is a K3 surface over $\mathbb{F}_{q^n}$ whose L_{trc} is the base extension of the one obtained from L.

[a] This hypothesis is made precise by Liedtke–Matsumoto (arxiv.1411.4797). It is is known for K3 surfaces of small degree relative to p.

Question: is it reasonable to hope that one can always take $n = 1$?

A theorem of Taelman

For X given, factor $L(T) = L_{\mathrm{alg}}(T) L_{\mathrm{trc}}(T)$ where L_{alg} is the product of all cyclotomic factors of L. (Aside: $\deg(L_{\mathrm{alg}}) + 1 = \operatorname{rank} \mathrm{NS}(X_{\overline{K}})$.)

Theorem (Taelman, 2016)

Assume[a] that K3 surfaces over finite extensions of $\mathbb{Q}_p$ admit potential semistable reduction. For q given, choose any L satisfying the initial constraints. Then for some positive integer n (and hence any multiple thereof), there is a K3 surface over $\mathbb{F}_{q^n}$ whose L_{trc} is the base extension of the one obtained from L.

[a] This hypothesis is made precise by Liedtke–Matsumoto (arxiv.1411.4797). It is is known for K3 surfaces of small degree relative to p.

Question: is it reasonable to hope that one can always take $n = 1$?

A positive result for $n = 1$

Take $q = 2$. Since we only have numerical data for smooth quartics, we can only hope to make an affirmative statement towards Taelman's theorem by limiting the set of candidates to those that "probably" come from smooth quartics.

To this end, consider those L satisfying the initial constraints for which

$$L_{\mathrm{alg}}(T) = 1 + T, \quad L_{\mathrm{trc}}(1) = 2, \quad L_{\mathrm{trc}}(-1) > 2.$$

This forces $\operatorname{rank} \mathrm{NS}(X) = \operatorname{rank} \mathrm{NS}(X_{\overline{K}}) = 1$, $|\Delta| = 4$. In particular, X must admit a degree 4 polarization, and so must be either a smooth quartic or a slightly degenerate case thereof.

We find 1995 candidates satisfying these constraints. All of them are realized by smooth quartics!

A positive result for $n = 1$

Take $q = 2$. Since we only have numerical data for smooth quartics, we can only hope to make an affirmative statement towards Taelman's theorem by limiting the set of candidates to those that "probably" come from smooth quartics.

To this end, consider those L satisfying the initial constraints for which

$$L_{\mathrm{alg}}(T) = 1 + T, \quad L_{\mathrm{trc}}(1) = 2, \quad L_{\mathrm{trc}}(-1) > 2.$$

This forces $\operatorname{rank} \mathrm{NS}(X) = \operatorname{rank} \mathrm{NS}(X_{\overline{K}}) = 1$, $|\Delta| = 4$. In particular, X must admit a degree 4 polarization, and so must be either a smooth quartic or a slightly degenerate case thereof.

We find 1995 candidates satisfying these constraints. All of them are realized by smooth quartics!

A positive result for $n = 1$

Take $q = 2$. Since we only have numerical data for smooth quartics, we can only hope to make an affirmative statement towards Taelman's theorem by limiting the set of candidates to those that "probably" come from smooth quartics.

To this end, consider those L satisfying the initial constraints for which

$$L_{\mathrm{alg}}(T) = 1 + T, \quad L_{\mathrm{trc}}(1) = 2, \quad L_{\mathrm{trc}}(-1) > 2.$$

This forces $\operatorname{rank} \mathrm{NS}(X) = \operatorname{rank} \mathrm{NS}(X_{\overline{K}}) = 1$, $|\Delta| = 4$. In particular, X must admit a degree 4 polarization, and so must be either a smooth quartic or a slightly degenerate case thereof.

We find 1995 candidates satisfying these constraints. All of them are realized by smooth quartics!

Contents

Other classes of K3 surfaces

Among our candidates for $q = 2$, we find many which cannot occur for smooth quartics. Namely, some of these have $L_1(1)$ equal to 2 times a large squarefree odd number D, which forces the K3 surface to admit a polarization of degree $2D$ (i.e., every generic hyperplane section has genus $\frac{D}{2} + 1$). Sample values of D include 307, 367, 463.

The moduli space of polarized K3 surfaces consists of one component per polarization degree; for degrees as large as these, these components are of general type. There is thus no hope for an "easy" construction of K3 surfaces matching these zeta functions.

In individual instances, one might be able to make Taelman's method effective: construct a suitable K3 surface over $\mathbb{C}$, descend it to a number field, and find a smooth model over some prime dividing 2.

Other classes of K3 surfaces

Among our candidates for $q = 2$, we find many which cannot occur for smooth quartics. Namely, some of these have $L_1(1)$ equal to 2 times a large squarefree odd number D, which forces the K3 surface to admit a polarization of degree $2D$ (i.e., every generic hyperplane section has genus $\frac{D}{2} + 1$). Sample values of D include 307, 367, 463.

The moduli space of polarized K3 surfaces consists of one component per polarization degree; for degrees as large as these, these components are of general type. There is thus no hope for an "easy" construction of K3 surfaces matching these zeta functions.

In individual instances, one might be able to make Taelman's method effective: construct a suitable K3 surface over $\mathbb{C}$, descend it to a number field, and find a smooth model over some prime dividing 2.

Other classes of K3 surfaces

Among our candidates for $q = 2$, we find many which cannot occur for smooth quartics. Namely, some of these have $L_1(1)$ equal to 2 times a large squarefree odd number D, which forces the K3 surface to admit a polarization of degree $2D$ (i.e., every generic hyperplane section has genus $\frac{D}{2} + 1$). Sample values of D include 307, 367, 463.

The moduli space of polarized K3 surfaces consists of one component per polarization degree; for degrees as large as these, these components are of general type. There is thus no hope for an "easy" construction of K3 surfaces matching these zeta functions.

In individual instances, one might be able to make Taelman's method effective: construct a suitable K3 surface over $\mathbb{C}$, descend it to a number field, and find a smooth model over some prime dividing 2.

Enumerating zeta candidates for $q > 2$

The algorithm for enumerating zeta candidates is sufficiently robust that it can be executed for slightly larger fields. The main difficulty is that the number of candidates over $\mathbb{F}_q$ is $O(q^{10})$, so even enumerating the answers gets tough quickly.

For example, for $q = 3$, in about 2.5 days we find 75,936,610 zeta functions satisfying the initial constraints, of which 49,645,728 satisfy the Elsenhans-Jahnel and monotonicity constraints.

Enumerating zeta candidates for $q > 2$

The algorithm for enumerating zeta candidates is sufficiently robust that it can be executed for slightly larger fields. The main difficulty is that the number of candidates over $\mathbb{F}_q$ is $O(q^{10})$, so even enumerating the answers gets tough quickly.

For example, for $q = 3$, in about 2.5 days we find 75,936,610 zeta functions satisfying the initial constraints, of which 49,645,728 satisfy the Elsenhans-Jahnel and monotonicity constraints.

Enumerating smooth quartics for $q > 2$

For $q > 2$ a bitmap representing all quartic polynomials is too large to work with. Instead, one must find a system of representatives for PGL_4-equivalence by computing some sort of "canonical form" for each quartic polynomial.

For $q = 3$, David Harvey has done this using the intersections with all rational planes, in 215 hours on a 16-core server. As a check, one again uses Burnside's formula to compute the number of PGL_4-equivalence classes; there are 4,127,971,480 of them.

Enumerating smooth quartics for $q > 2$

For $q > 2$ a bitmap representing all quartic polynomials is too large to work with. Instead, one must find a system of representatives for PGL_4-equivalence by computing some sort of “canonical form” for each quartic polynomial.

For $q = 3$, David Harvey has done this using the intersections with all rational planes, in 215 hours on a 16-core server. As a check, one again uses Burnside’s formula to compute the number of PGL_4-equivalence classes; there are 4,127,971,480 of them.

Computing zeta functions for $q > 2$

Our approach for computing zeta functions over $\mathbb{F}_2$ is not feasible for $q > 2$, both in the increased per-instance cost and in light of the number of instances required.

Instead, one should switch to p-adic cohomological methods, particularly the modified version of the Abbott–Kedlaya–Roe method of Costa–Harvey. With this method, a complete census over $\mathbb{F}_3$ is probably doable.

Computing zeta functions for $q > 2$

Our approach for computing zeta functions over $\mathbb{F}_2$ is not feasible for $q > 2$, both in the increased per-instance cost and in light of the number of instances required.

Instead, one should switch to p-adic cohomological methods, particularly the modified version of the Abbott–Kedlaya–Roe method of Costa–Harvey. With this method, a complete census over $\mathbb{F}_3$ is probably doable.