

Zeta functions of algebraic varieties over finite fields

Kiran S. Kedlaya

Department of Mathematics, University of California San Diego

kedlaya@ucsd.edu

These slides can be downloaded from <https://kskedlaya.org/slides/>.

Quantum computer science seminar

Google (virtual)

August 6, 2024

Supported by  (grant DMS-2053473) and  (Warschawski Professorship).

I acknowledge that my workplace occupies unceded ancestral land of the [Kumeyaay Nation](#).



Contents

- 1 Zeta functions of algebraic varieties
- 2 The (classical) zeta function problem and some classical results
- 3 The (apparent) quantum advantage

Finite fields

For every prime power q , there is a unique (up to isomorphism) finite field of q elements, denoted \mathbb{F}_q . For example, if $q = p$ is prime, these are just the integers modulo p .

For every positive integer n , we can view \mathbb{F}_{q^n} as an extension of \mathbb{F}_q . Consequently, if we consider a system of polynomial equations

$$P_1(x_1, \dots, x_m) = \dots = P_k(x_1, \dots, x_m) = 0$$

with coefficients in \mathbb{F}_q , we can form the set of solutions over \mathbb{F}_{q^n} for each n .

These solutions can formally be reinterpreted as the \mathbb{F}_{q^n} -valued points of the **affine algebraic variety**

$$X := \text{Spec } \mathbb{F}_q[x_1, \dots, x_m] / (P_1, \dots, P_k).$$

The **dimension** of X is (roughly) $m - k$. The complexity of X is also controlled by the **degrees** of the P_i .

Finite fields

For every prime power q , there is a unique (up to isomorphism) finite field of q elements, denoted \mathbb{F}_q . For example, if $q = p$ is prime, these are just the integers modulo p .

For every positive integer n , we can view \mathbb{F}_{q^n} as an extension of \mathbb{F}_q . Consequently, if we consider a system of polynomial equations

$$P_1(x_1, \dots, x_m) = \dots = P_k(x_1, \dots, x_m) = 0$$

with coefficients in \mathbb{F}_q , we can form the set of solutions over \mathbb{F}_{q^n} for each n .

These solutions can formally be reinterpreted as the \mathbb{F}_{q^n} -valued points of the **affine algebraic variety**

$$X := \text{Spec } \mathbb{F}_q[x_1, \dots, x_m] / (P_1, \dots, P_k).$$

The **dimension** of X is (roughly) $m - k$. The complexity of X is also controlled by the **degrees** of the P_i .

Finite fields

For every prime power q , there is a unique (up to isomorphism) finite field of q elements, denoted \mathbb{F}_q . For example, if $q = p$ is prime, these are just the integers modulo p .

For every positive integer n , we can view \mathbb{F}_{q^n} as an extension of \mathbb{F}_q . Consequently, if we consider a system of polynomial equations

$$P_1(x_1, \dots, x_m) = \dots = P_k(x_1, \dots, x_m) = 0$$

with coefficients in \mathbb{F}_q , we can form the set of solutions over \mathbb{F}_{q^n} for each n .

These solutions can formally be reinterpreted as the \mathbb{F}_{q^n} -valued points of the **affine algebraic variety**

$$X := \text{Spec } \mathbb{F}_q[x_1, \dots, x_m] / (P_1, \dots, P_k).$$

The **dimension** of X is (roughly) $m - k$. The complexity of X is also controlled by the **degrees** of the P_i .

Finite fields

For every prime power q , there is a unique (up to isomorphism) finite field of q elements, denoted \mathbb{F}_q . For example, if $q = p$ is prime, these are just the integers modulo p .

For every positive integer n , we can view \mathbb{F}_{q^n} as an extension of \mathbb{F}_q . Consequently, if we consider a system of polynomial equations

$$P_1(x_1, \dots, x_m) = \dots = P_k(x_1, \dots, x_m) = 0$$

with coefficients in \mathbb{F}_q , we can form the set of solutions over \mathbb{F}_{q^n} for each n .

These solutions can formally be reinterpreted as the \mathbb{F}_{q^n} -valued points of the **affine algebraic variety**

$$X := \text{Spec } \mathbb{F}_q[x_1, \dots, x_m] / (P_1, \dots, P_k).$$

The **dimension** of X is (roughly) $m - k$. The complexity of X is also controlled by the **degrees** of the P_i .

Zeta functions

The infinite sequence

$$\#X(\mathbb{F}_q), \#X(\mathbb{F}_{q^2}), \#X(\mathbb{F}_{q^3}), \dots$$

can be represented in a finitistic way: the formal power series

$$Z(X, T) = \exp \left(\sum_{n=1}^{\infty} \frac{T^n}{n} \#X(\mathbb{F}_{q^n}) \right)$$

represents a **rational function** of T . Concretely, this means that there exist some $\alpha_1, \dots, \alpha_r, \beta_1, \dots, \beta_s \in \mathbb{C}$ such that

$$\#X(\mathbb{F}_{q^n}) = \alpha_1^n + \dots + \alpha_r^n - \beta_1^n - \dots - \beta_s^n.$$

Moreover, the degree of the rational function (equivalently, the values of r and s) can be bounded in terms of n and the degrees of the P_i .

Zeta functions

The infinite sequence

$$\#X(\mathbb{F}_q), \#X(\mathbb{F}_{q^2}), \#X(\mathbb{F}_{q^3}), \dots$$

can be represented in a finitistic way: the formal power series

$$Z(X, T) = \exp \left(\sum_{n=1}^{\infty} \frac{T^n}{n} \#X(\mathbb{F}_{q^n}) \right)$$

represents a **rational function** of T . Concretely, this means that there exist some $\alpha_1, \dots, \alpha_r, \beta_1, \dots, \beta_s \in \mathbb{C}$ such that

$$\#X(\mathbb{F}_{q^n}) = \alpha_1^n + \dots + \alpha_r^n - \beta_1^n - \dots - \beta_s^n.$$

Moreover, the degree of the rational function (equivalently, the values of r and s) can be bounded in terms of n and the degrees of the P_i .

Examples

For X equal to the affine space (i.e., no equations!), we have $\#X(\mathbb{F}_{q^n}) = q^{mn}$, so

$$Z(X, T) = \frac{1}{1 - q^m T}.$$

Let $Q(x) \in \mathbb{F}_q[x]$ be a squarefree cubic polynomial. For X the associated affine elliptic curve:

$$X = \text{Spec } \mathbb{F}_q[x, y]/(y^2 - Q(x))$$

we have

$$Z(X, T) = \frac{1 - aT + qT^2}{1 - qT}$$

for some integer a with $|a| \leq 2\sqrt{q}$. There is no “easy” formula for a in terms of Q .

Examples

For X equal to the affine space (i.e., no equations!), we have $\#X(\mathbb{F}_{q^n}) = q^{mn}$, so

$$Z(X, T) = \frac{1}{1 - q^m T}.$$

Let $Q(x) \in \mathbb{F}_q[x]$ be a squarefree cubic polynomial. For X the associated affine elliptic curve:

$$X = \text{Spec } \mathbb{F}_q[x, y]/(y^2 - Q(x))$$

we have

$$Z(X, T) = \frac{1 - aT + qT^2}{1 - qT}$$

for some integer a with $|a| \leq 2\sqrt{q}$. There is no “easy” formula for a in terms of Q .

Contents

- 1 Zeta functions of algebraic varieties
- 2 The (classical) zeta function problem and some classical results
- 3 The (apparent) quantum advantage

Formulation of the problem

Problem

Given on input a prime power q and a sequence of polynomials $P_1, \dots, P_k \in \mathbb{F}_q[x_1, \dots, x_m]$, return $Z(X, T)$ in the form of a rational function in T .

As formulated this might appear to be NP-hard, since the data of $Z(X, T)$ includes $\#X(\mathbb{F}_q)$, and 3-SAT can be expressed as the question of whether a certain sequence of polynomials has a common zero.

However, one must measure polynomiality in the length of the input **and the output**, and the degree of $Z(X, T)$ grows exponentially with m .

In any case, we will generally treat m as fixed, in which case the input (in a dense representation) and the output both have length polynomial in $\log q \deg P_1 \cdots \deg P_k$. It is an **open question** to give a randomized polynomial time algorithm.

Formulation of the problem

Problem

Given on input a prime power q and a sequence of polynomials $P_1, \dots, P_k \in \mathbb{F}_q[x_1, \dots, x_m]$, return $Z(X, T)$ in the form of a rational function in T .

As formulated this might appear to be NP-hard, since the data of $Z(X, T)$ includes $\#X(\mathbb{F}_q)$, and 3-SAT can be expressed as the question of whether a certain sequence of polynomials has a common zero.

However, one must measure polynomiality in the length of the input **and the output**, and the degree of $Z(X, T)$ grows exponentially with m .

In any case, we will generally treat m as fixed, in which case the input (in a dense representation) and the output both have length polynomial in $\log q \deg P_1 \cdots \deg P_k$. It is an **open question** to give a randomized polynomial time algorithm.

Formulation of the problem

Problem

Given on input a prime power q and a sequence of polynomials $P_1, \dots, P_k \in \mathbb{F}_q[x_1, \dots, x_m]$, return $Z(X, T)$ in the form of a rational function in T .

As formulated this might appear to be NP-hard, since the data of $Z(X, T)$ includes $\#X(\mathbb{F}_q)$, and 3-SAT can be expressed as the question of whether a certain sequence of polynomials has a common zero.

However, one must measure polynomiality in the length of the input **and the output**, and the degree of $Z(X, T)$ grows exponentially with m .

In any case, we will generally treat m as fixed, in which case the input (in a dense representation) and the output both have length polynomial in $\log q \deg P_1 \cdots \deg P_k$. It is an **open question** to give a randomized polynomial time algorithm.

Formulation of the problem

Problem

Given on input a prime power q and a sequence of polynomials $P_1, \dots, P_k \in \mathbb{F}_q[x_1, \dots, x_m]$, return $Z(X, T)$ in the form of a rational function in T .

As formulated this might appear to be NP-hard, since the data of $Z(X, T)$ includes $\#X(\mathbb{F}_q)$, and 3-SAT can be expressed as the question of whether a certain sequence of polynomials has a common zero.

However, one must measure polynomiality in the length of the input **and the output**, and the degree of $Z(X, T)$ grows exponentially with m .

In any case, we will generally treat m as fixed, in which case the input (in a dense representation) and the output both have length polynomial in $\log q \deg P_1 \cdots \deg P_k$. It is an **open question** to give a randomized polynomial time algorithm.

Zeta functions and cohomology

Most approaches to studying or computing $Z(X, T)$ use some sort of **Lefschetz trace formula for Frobenius**, i.e., the data of some finite-dimensional vector spaces V_i over some field K of characteristic 0, plus invertible linear transformations $F: V_i \rightarrow V_i$ for which

$$\#X(\mathbb{F}_{q^n}) = \sum_i (-1)^i \text{Trace}(F^n, V_i).$$

The existence of any such data already implies that $Z(X, T)$ is a rational function: its poles (resp. zeroes) are the reciprocals of the eigenvalues of F^n acting on the V_i for i even (resp. for i odd).

Zeta functions and cohomology

Most approaches to studying or computing $Z(X, T)$ use some sort of **Lefschetz trace formula for Frobenius**, i.e., the data of some finite-dimensional vector spaces V_i over some field K of characteristic 0, plus invertible linear transformations $F: V_i \rightarrow V_i$ for which

$$\#X(\mathbb{F}_{q^n}) = \sum_i (-1)^i \text{Trace}(F^n, V_i).$$

The existence of any such data already implies that $Z(X, T)$ is a rational function: its poles (resp. zeroes) are the reciprocals of the eigenvalues of F^n acting on the V_i for i even (resp. for i odd).

Schoof's algorithm

Theorem (Schoof, 1985)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(y^2 - P(x))$ with $P \in \mathbb{F}_q[x]$ a squarefree cubic polynomial, there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q$.

This follows from the corresponding statement for X an arbitrary **elliptic curve** over \mathbb{F}_q . The latter has the structure of an algebraic group; for every prime ℓ not dividing q , the ℓ -torsion points (over an algebraic closure of \mathbb{F}_q) form a two-dimensional \mathbb{F}_ℓ -vector space on which Frobenius acts.

This is not quite a trace formula because the characteristic of \mathbb{F}_ℓ is not 0, so we only get $Z(X, T)$ modulo ℓ . To fix this, we use **several** primes ℓ plus the Chinese remainder theorem. This “multimodular” strategy also occurs in my work with Umans on polynomial factorization (FOCS 2008), and work of Lin–Mook–Wichs on private information retrieval (STOC 2023).

Aside: there has been some speculation that this could be used to derandomize polynomial factorization over \mathbb{F}_q (Agrawal, Poonen, etc.) but as yet no results.

Schoof's algorithm

Theorem (Schoof, 1985)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(y^2 - P(x))$ with $P \in \mathbb{F}_q[x]$ a squarefree cubic polynomial, there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q$.

This follows from the corresponding statement for X an arbitrary **elliptic curve** over \mathbb{F}_q . The latter has the structure of an algebraic group; for every prime ℓ not dividing q , the ℓ -torsion points (over an algebraic closure of \mathbb{F}_q) form a two-dimensional \mathbb{F}_ℓ -vector space on which Frobenius acts.

This is not quite a trace formula because the characteristic of \mathbb{F}_ℓ is not 0, so we only get $Z(X, T)$ modulo ℓ . To fix this, we use **several** primes ℓ plus the Chinese remainder theorem. This “multimodular” strategy also occurs in my work with Umans on polynomial factorization (FOCS 2008), and work of Lin–Mook–Wichs on private information retrieval (STOC 2023).

Aside: there has been some speculation that this could be used to derandomize polynomial factorization over \mathbb{F}_q (Agrawal, Poonen, etc.) but as yet no results.

Schoof's algorithm

Theorem (Schoof, 1985)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(y^2 - P(x))$ with $P \in \mathbb{F}_q[x]$ a squarefree cubic polynomial, there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q$.

This follows from the corresponding statement for X an arbitrary **elliptic curve** over \mathbb{F}_q . The latter has the structure of an algebraic group; for every prime ℓ not dividing q , the ℓ -torsion points (over an algebraic closure of \mathbb{F}_q) form a two-dimensional \mathbb{F}_ℓ -vector space on which Frobenius acts.

This is not quite a trace formula because the characteristic of \mathbb{F}_ℓ is not 0, so we only get $Z(X, T)$ modulo ℓ . To fix this, we use **several** primes ℓ plus the Chinese remainder theorem. This “multimodular” strategy also occurs in my work with Umans on polynomial factorization (FOCS 2008), and work of Lin–Mook–Wichs on private information retrieval (STOC 2023).

Aside: there has been some speculation that this could be used to derandomize polynomial factorization over \mathbb{F}_q (Agrawal, Poonen, etc.) but as yet no results.

The one-dimensional case

Theorem (Pila, 1990; Adleman–Huang, 2001)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P)$ with $P \in \mathbb{F}_q[x, y]$ of **fixed degree**, there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q$.

This uses a similar (not quite a) trace formula for the **Jacobian variety** associated to a projective algebraic curve. This can also be interpreted in terms of the **étale cohomology** of the curve with coefficients in \mathbb{F}_ℓ .

Unfortunately, there seems to be way to reduce the complexity in $\deg P$ below exponential. More on this later.

It is expected that one can do something similar in dimension > 1 if we again fix the characteristic and the polynomial degrees, but this depends on ongoing work to make étale cohomology computationally explicit; more on this later.

The one-dimensional case

Theorem (Pila, 1990; Adleman–Huang, 2001)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P)$ with $P \in \mathbb{F}_q[x, y]$ of **fixed degree**, there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q$.

This uses a similar (not quite a) trace formula for the **Jacobian variety** associated to a projective algebraic curve. This can also be interpreted in terms of the **étale cohomology** of the curve with coefficients in \mathbb{F}_ℓ .

Unfortunately, there seems to be way to reduce the complexity in $\deg P$ below exponential. More on this later.

It is expected that one can do something similar in dimension > 1 if we again fix the characteristic and the polynomial degrees, but this depends on ongoing work to make étale cohomology computationally explicit; more on this later.

The one-dimensional case

Theorem (Pila, 1990; Adleman–Huang, 2001)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P)$ with $P \in \mathbb{F}_q[x, y]$ of **fixed degree**, there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q$.

This uses a similar (not quite a) trace formula for the **Jacobian variety** associated to a projective algebraic curve. This can also be interpreted in terms of the **étale cohomology** of the curve with coefficients in \mathbb{F}_ℓ .

Unfortunately, there seems to be way to reduce the complexity in $\deg P$ below exponential. More on this later.

It is expected that one can do something similar in dimension > 1 if we again fix the characteristic and the polynomial degrees, but this depends on ongoing work to make étale cohomology computationally explicit; more on this later.

The one-dimensional case

Theorem (Pila, 1990; Adleman–Huang, 2001)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P)$ with $P \in \mathbb{F}_q[x, y]$ of **fixed degree**, there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q$.

This uses a similar (not quite a) trace formula for the **Jacobian variety** associated to a projective algebraic curve. This can also be interpreted in terms of the **étale cohomology** of the curve with coefficients in \mathbb{F}_ℓ .

Unfortunately, there seems to be way to reduce the complexity in $\deg P$ below exponential. More on this later.

It is expected that one can do something similar in dimension > 1 if we again fix the characteristic and the polynomial degrees, but this depends on ongoing work to make étale cohomology computationally explicit; more on this later.

Fixed characteristic

Theorem (Lauder–Wan, 2008)

Fix a positive integer m and a prime p . For $X = \text{Spec } \mathbb{F}_q[x_1, \dots, x_m]/(P_1, \dots, P_k)$ where q is a power of p , there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P_1 \cdots \deg P_k$.

This uses a different trace formula arising from p -**adic cohomology**, which can be computed efficiently in terms of differential forms. Variants of this method for special classes of varieties even work well in practice!

Unfortunately, there seems to be no way to reduce the complexity in p below square-root (Harvey).

Aside: if one starts with polynomials over \mathbb{Z} and reduces modulo various primes p , there are ways to amortize that cost over p (Harvey–Sutherland etc.). This use case comes up in number theory when studying L -**functions**, as in the Langlands program.

Fixed characteristic

Theorem (Lauder–Wan, 2008)

Fix a positive integer m and a prime p . For $X = \text{Spec } \mathbb{F}_q[x_1, \dots, x_m]/(P_1, \dots, P_k)$ where q is a power of p , there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P_1 \cdots \deg P_k$.

This uses a different trace formula arising from **p -adic cohomology**, which can be computed efficiently in terms of differential forms. Variants of this method for special classes of varieties even work well in practice!

Unfortunately, there seems to be no way to reduce the complexity in p below square-root (Harvey).

Aside: if one starts with polynomials over \mathbb{Z} and reduces modulo various primes p , there are ways to amortize that cost over p (Harvey–Sutherland etc.). This use case comes up in number theory when studying **L -functions**, as in the Langlands program.

Fixed characteristic

Theorem (Lauder–Wan, 2008)

Fix a positive integer m and a prime p . For $X = \text{Spec } \mathbb{F}_q[x_1, \dots, x_m]/(P_1, \dots, P_k)$ where q is a power of p , there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P_1 \cdots \deg P_k$.

This uses a different trace formula arising from **p -adic cohomology**, which can be computed efficiently in terms of differential forms. Variants of this method for special classes of varieties even work well in practice!

Unfortunately, there seems to be no way to reduce the complexity in p below square-root (Harvey).

Aside: if one starts with polynomials over \mathbb{Z} and reduces modulo various primes p , there are ways to amortize that cost over p (Harvey–Sutherland etc.). This use case comes up in number theory when studying **L -functions**, as in the Langlands program.

Fixed characteristic

Theorem (Lauder–Wan, 2008)

Fix a positive integer m and a prime p . For $X = \text{Spec } \mathbb{F}_q[x_1, \dots, x_m]/(P_1, \dots, P_k)$ where q is a power of p , there is a deterministic algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P_1 \cdots \deg P_k$.

This uses a different trace formula arising from **p -adic cohomology**, which can be computed efficiently in terms of differential forms. Variants of this method for special classes of varieties even work well in practice!

Unfortunately, there seems to be no way to reduce the complexity in p below square-root (Harvey).

Aside: if one starts with polynomials over \mathbb{Z} and reduces modulo various primes p , there are ways to amortize that cost over p (Harvey–Sutherland etc.). This use case comes up in number theory when studying **L -functions**, as in the Langlands program.

Contents

- 1 Zeta functions of algebraic varieties
- 2 The (classical) zeta function problem and some classical results
- 3 The (apparent) quantum advantage

Shor's algorithm, version 1

Theorem (Shor, 1994)

For G a “black box abelian group”, there is a quantum polynomial time algorithm to compute the order of the subgroup generated by a list of elements.

Clarification: a “black box abelian group” is a group whose underlying elements are a set of bitstrings, together with an oracle that can:

- test a bitstring for membership;
- compute the inverse of a bitstring;
- compute the group operation on two bitstrings.

By generating random elements of G , this promotes to an algorithm for computing the order of G itself.

Shor's algorithm, version 1

Theorem (Shor, 1994)

For G a “black box abelian group”, there is a quantum polynomial time algorithm to compute the order of the subgroup generated by a list of elements.

Clarification: a “black box abelian group” is a group whose underlying elements are a set of bitstrings, together with an oracle that can:

- test a bitstring for membership;
- compute the inverse of a bitstring;
- compute the group operation on two bitstrings.

By generating random elements of G , this promotes to an algorithm for computing the order of G itself.

Shor's algorithm, version 1

Theorem (Shor, 1994)

For G a “black box abelian group”, there is a quantum polynomial time algorithm to compute the order of the subgroup generated by a list of elements.

Clarification: a “black box abelian group” is a group whose underlying elements are a set of bitstrings, together with an oracle that can:

- test a bitstring for membership;
- compute the inverse of a bitstring;
- compute the group operation on two bitstrings.

By generating random elements of G , this promotes to an algorithm for computing the order of G itself.

A quantum advantage for zeta functions: dimension 1 case

Theorem (K, 2006)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P(x, y))$, there is a quantum algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P$.

It suffices to compute $Z(C, T)$ where C is the associated smooth projective curve. This has the form

$$Z(C, T) = \frac{Q(T)}{(1-T)(1-qT)}$$

where Q is a polynomial of degree $2g$ where g is the **genus** of C (roughly the degree of P).

The Jacobian variety J has the property that for every n ,

$$Q(1)Q(e^{2\pi i/n}) \cdots Q(e^{2\pi i(n-1)/n}) = \#J(\mathbb{F}_{q^n}).$$

Treating $J(\mathbb{F}_{q^n})$ as a black box abelian group and using Shor to compute its order for $O(g)$ values of n , we can (classically) recover $Q(T)$.

A quantum advantage for zeta functions: dimension 1 case

Theorem (K, 2006)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P(x, y))$, there is a quantum algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P$.

It suffices to compute $Z(C, T)$ where C is the associated smooth projective curve. This has the form

$$Z(C, T) = \frac{Q(T)}{(1-T)(1-qT)}$$

where Q is a polynomial of degree $2g$ where g is the **genus** of C (roughly the degree of P).

The Jacobian variety J has the property that for every n ,

$$Q(1)Q(e^{2\pi i/n}) \cdots Q(e^{2\pi i(n-1)/n}) = \#J(\mathbb{F}_{q^n}).$$

Treating $J(\mathbb{F}_{q^n})$ as a black box abelian group and using Shor to compute its order for $O(g)$ values of n , we can (classically) recover $Q(T)$.

A quantum advantage for zeta functions: dimension 1 case

Theorem (K, 2006)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P(x, y))$, there is a quantum algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P$.

It suffices to compute $Z(C, T)$ where C is the associated smooth projective curve. This has the form

$$Z(C, T) = \frac{Q(T)}{(1-T)(1-qT)}$$

where Q is a polynomial of degree $2g$ where g is the **genus** of C (roughly the degree of P).

The Jacobian variety J has the property that for every n ,

$$Q(1)Q(e^{2\pi i/n}) \cdots Q(e^{2\pi i(n-1)/n}) = \#J(\mathbb{F}_{q^n}).$$

Treating $J(\mathbb{F}_{q^n})$ as a black box abelian group and using Shor to compute its order for $O(g)$ values of n , we can (classically) recover $Q(T)$.

A quantum advantage for zeta functions: dimension 1 case

Theorem (K, 2006)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P(x, y))$, there is a quantum algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P$.

It suffices to compute $Z(C, T)$ where C is the associated smooth projective curve. This has the form

$$Z(C, T) = \frac{Q(T)}{(1-T)(1-qT)}$$

where Q is a polynomial of degree $2g$ where g is the **genus** of C (roughly the degree of P).

The Jacobian variety J has the property that for every n ,

$$Q(1)Q(e^{2\pi i/n}) \cdots Q(e^{2\pi i(n-1)/n}) = \#J(\mathbb{F}_{q^n}).$$

Treating $J(\mathbb{F}_{q^n})$ as a black box abelian group and using Shor to compute its order for $O(g)$ values of n , we can (classically) recover $Q(T)$.

Shor's algorithm, version 2

For ℓ a prime, by a “black box \mathbb{F}_ℓ -vector space” I mean a black box abelian group in which every element is killed by ℓ .

Let V be a black box \mathbb{F}_ℓ -vector space with a known basis v_1, \dots, v_m . Given another element w , we can find the coefficients c_1, \dots, c_m for which $c_1 v_1 + \dots + c_m v_m = w$, by recovering the kernel of the linear transformation $\mathbb{F}_\ell^m \times V \rightarrow V$ taking (a_1, \dots, a_m, v) to $a_1 v_1 + \dots + a_m v_m - v$.

Now let $F: V \rightarrow V$ be a “black box endomorphism” of V . We can then recover the matrix expressing F in terms of v_1, \dots, v_m , and from that the characteristic polynomial of V . (If we don't start with a basis of V , pick enough elements to generate with high probability, then make a similar kernel computation to reduce to a basis.)

Shor's algorithm, version 2

For ℓ a prime, by a “black box \mathbb{F}_ℓ -vector space” I mean a black box abelian group in which every element is killed by ℓ .

Let V be a black box \mathbb{F}_ℓ -vector space with a known basis v_1, \dots, v_m . Given another element w , we can find the coefficients c_1, \dots, c_m for which $c_1 v_1 + \dots + c_m v_m = w$, by recovering the kernel of the linear transformation $\mathbb{F}_\ell^m \times V \rightarrow V$ taking (a_1, \dots, a_m, v) to $a_1 v_1 + \dots + a_m v_m - v$.

Now let $F: V \rightarrow V$ be a “black box endomorphism” of V . We can then recover the matrix expressing F in terms of v_1, \dots, v_m , and from that the characteristic polynomial of V . (If we don't start with a basis of V , pick enough elements to generate with high probability, then make a similar kernel computation to reduce to a basis.)

Shor's algorithm, version 2

For ℓ a prime, by a “black box \mathbb{F}_ℓ -vector space” I mean a black box abelian group in which every element is killed by ℓ .

Let V be a black box \mathbb{F}_ℓ -vector space with a known basis v_1, \dots, v_m . Given another element w , we can find the coefficients c_1, \dots, c_m for which $c_1 v_1 + \dots + c_m v_m = w$, by recovering the kernel of the linear transformation $\mathbb{F}_\ell^m \times V \rightarrow V$ taking (a_1, \dots, a_m, v) to $a_1 v_1 + \dots + a_m v_m - v$.

Now let $F: V \rightarrow V$ be a “black box endomorphism” of V . We can then recover the matrix expressing F in terms of v_1, \dots, v_m , and from that the characteristic polynomial of V . (If we don't start with a basis of V , pick enough elements to generate with high probability, then make a similar kernel computation to reduce to a basis.)

A quantum advantage for zeta functions: dimension 1 case (again)

Theorem (K, 2006 but not with this proof)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P(x, y))$, there is a quantum algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P$.

For each prime ℓ not divisible by p , we can represent the ℓ -torsion of the Jacobian variety of C as a black box \mathbb{F}_ℓ -vector space of dimension $2g$, and Frobenius as a black box endomorphism. In the expression

$$Z(C, T) = \frac{Q(T)}{(1 - T)(1 - qT)},$$

$Q(T)$ reduces modulo ℓ to the characteristic polynomial of Frobenius; we can thus recover $Q(T)$ by another multimodular calculation.

A quantum advantage for zeta functions: dimension 1 case (again)

Theorem (K, 2006 but not with this proof)

For $X = \text{Spec } \mathbb{F}_q[x, y]/(P(x, y))$, there is a quantum algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P$.

For each prime ℓ not divisible by p , we can represent the ℓ -torsion of the Jacobian variety of C as a black box \mathbb{F}_ℓ -vector space of dimension $2g$, and Frobenius as a black box endomorphism. In the expression

$$Z(C, T) = \frac{Q(T)}{(1 - T)(1 - qT)},$$

$Q(T)$ reduces modulo ℓ to the characteristic polynomial of Frobenius; we can thus recover $Q(T)$ by another multimodular calculation.

A quantum advantage for zeta functions: higher dimensional case?

Conjecture

For fixed m , for $X = \text{Spec } \mathbb{F}_q[x_1, \dots, x_m]/(P_1, \dots, P_k)$, there is a quantum algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P_1 \cdots \deg P_k$.

This would follow if one had an efficient black box representation of étale cohomology with coefficients in \mathbb{F}_ℓ . So far this is only known for the first cohomology group (Roy–Saxena–Venkatesh, 2024).

A quantum advantage for zeta functions: higher dimensional case?

Conjecture

For fixed m , for $X = \text{Spec } \mathbb{F}_q[x_1, \dots, x_m]/(P_1, \dots, P_k)$, there is a quantum algorithm to compute $Z(X, T)$ in time polynomial in $\log q \deg P_1 \cdots \deg P_k$.

This would follow if one had an efficient black box representation of étale cohomology with coefficients in \mathbb{F}_ℓ . So far this is only known for the first cohomology group (Roy–Saxena–Venkatesh, 2024).

Some food for thought

Question for the audience: The use of Shor's algorithm and its variants means that we are talking about quantum circuits of (currently) impractically large depth. Is there a "time-depth tradeoff"?

This would of course have **much more significant** consequences, including to the security of classical (not quantum-resistant) public key cryptography.

Some food for thought

Question for the audience: The use of Shor's algorithm and its variants means that we are talking about quantum circuits of (currently) impractically large depth. Is there a "time-depth tradeoff"?

This would of course have **much more significant** consequences, including to the security of classical (not quantum-resistant) public key cryptography.

Some food for thought

Question for the audience: The use of Shor's algorithm and its variants means that we are talking about quantum circuits of (currently) impractically large depth. Is there a “time-depth tradeoff”?

This would of course have **much more significant** consequences, including to the security of classical (not quantum-resistant) public key cryptography.

Thank you!

Questions?