We prove the analogue of the Riemann hypothesis for ζ-functions associated to curves over finite fields. This result was claimed by André Weil in 1940. Weil published his proof in 1948. Our proof follows E. Bombieri's 1973 Séminaire Bourbaki exposé of Stepanov's proof.

Let $\mathbf{F}_q$ be a finite field of characteristic p. Let X be a smooth projective absolutely irreducible curve over $\mathbf{F}_q$ of genus g. Stepanov's proof proceeds by constructing a rational function f on X that vanishes to high order m at the $\mathbf{F}_q$-rational points of X, but whose number of poles is bounded by B. Since the number of poles of f is equal to its number of zeroes, we find that

$$m\#X(\mathbf{F}_q) \leq B.$$

This implies an upper bound for $\#X(\mathbf{F}_q)$. A modification of this idea leads to a lower bound and to an estimate for the absolute values of the reciprocal zeroes ϕ of the zeta functions $Z_X(T)$ of X. The functional equation implies then that $|\phi| = \sqrt{q}$ which is the analogue of the Riemann hypothesis.

Let ∞ be a point in $X(\mathbf{F}_q)$. We study the spaces

$$H^0(m(\infty)) = \{f \in \mathbf{F}_q(X) : \mathrm{ord}_\infty(f) \geq -m\}$$

for $m \geq 0$. By the Riemann-Roch Theorem, the dimension of $H^0(m(\infty))$ is equal to $m - g + 1$ whenever $m > 2g - 2$. We have that

$$\ldots \quad \subset \quad H^0(m(\infty)) \quad \subset \quad H^0((m+1)(\infty)) \quad \subset \quad \ldots$$

Let π denote a uniformizer at ∞. The sequences

$$0 \longrightarrow \quad H^0(m(\infty)) \quad \longrightarrow \quad H^0((m+1)(\infty)) \quad \xrightarrow{e} \quad \mathbf{F}_q$$

given by $e(g) = (g\pi^{m+1})(\infty)$ are exact. Therefore the codimension of $H^0(m(\infty))$ inside $H^0((m+1)(\infty))$ is at most 1. It follows that we can choose an $\mathbf{F}_q$-basis $e_1, \ldots, e_t$ of $H^0(m(\infty))$ for which

$$\mathrm{ord}_\infty(e_1) > \mathrm{ord}_\infty(e_2) > \ldots > \mathrm{ord}_\infty(e_t). \qquad (*)$$

For a positive integer μ we let $H^0(m(\infty))^{p^\mu}$ denote the $\mathbf{F}_q$-vector space of functions of the form f^{p^μ} where $f \in H^0(m(\infty))$. By $H^0(a(\infty))^{p^\mu} H^0(b(\infty))^q$ we denote the $\mathbf{F}_q$-subspace of $H^0((ap^\mu + bq)(\infty))$ generated by products fg of functions $f \in H^0(a(\infty))^{p^\mu}$ and $g \in H^0(b(\infty))^q$.

Lemma 1. *Let $a, b > 2g - 2$ and let p^μ be a power of p for which $ap^\mu < q$. Let $e_1, \ldots, e_{a-g+1}$ and $f_1, \ldots, f_{b-g+1}$ denote bases as in (*) of the vector spaces $H^0(a(\infty))$ and $H^0(b(\infty))$ respectively. Then the products $e_i^{p^\mu} f_j^q$ for $1 \leq i \leq a-g+1$ and $1 \leq j \leq b-g+1$ form an $\mathbf{F}_q$-basis for $H^0(a(\infty))^{p^\mu} H^0(b(\infty))^q$.*

Proof. Clearly the products $e_i^{p^\mu} f_j^q$ generate the vector space $H^0(a(\infty))^{p^\mu} H^0(b(\infty))^q$. The point is to show that they are independent. To this end we observe that the functions $e_i^{p^\mu} f_j^q$ have poles of distinct orders at ∞. Indeed, if $\mathrm{ord}_\infty(e_i^{p^\mu} f_j^q) = \mathrm{ord}_\infty(e_{i'}^{p^\mu} f_{j'}^q)$ we have that

$$p^\mu(\mathrm{ord}_\infty(e_i) - \mathrm{ord}_\infty(e_{i'})) = q(\mathrm{ord}_\infty(f_{j'}) - \mathrm{ord}_\infty(f_j)).$$

Since the absolute value of $p^\mu(\mathrm{ord}_\infty(e_i)-\mathrm{ord}_\infty(e_{i'}))$ does not exceed $p^\mu a < q$, we conclude that $i = i'$ and hence that $j = j'$.

The lemma now follows from the fact that any functions $h_i \in H^0(m(\infty))$ with poles of distinct orders at ∞, are necessarily linearly independent. This proves the lemma.

Proposition 2. *Let X be a curve of genus g over $\mathbf{F}_q$. If q is a square and $q > (2g+2)^2$ then we have that*

$$\#X(\mathbf{F}_q) < q + (3g+1)\sqrt{q}.$$

Proof. We may assume that $X(\mathbf{F}_q) \neq \emptyset$ and choose a point $\infty \in X(\mathbf{F}_q)$. We consider the space $H^0(a(\infty))^{p^\mu} H^0(b(\infty))^q$ of Lemma 1 with $p^\mu = \sqrt{q}$ and $a = \sqrt{q} - 1$. Then $a > 2g-2$. We choose $b > 2g-2$ later. These choices imply that the dimension of the space $H^0(a(\infty))^{p^\mu} H^0(b(\infty))^q$ is equal to

$$\dim H^0(a(\infty)) \times \dim H^0(b(\infty)) = (a-g+1)(b-g+1).$$

We define an $\mathbf{F}_q$-linear homomorphism

$$H^0(a(\infty))^{p^\mu} H^0(b(\infty))^q \longrightarrow H^0((ap^\mu + b)(\infty))$$

by mapping the basis vectors $e_i^{p^\mu} f_j^q$ to $e_i^{p^\mu} f_j$. This homomorphism is *well defined* by Lemma 1. The dimension of $H^0((ap^\mu + b)(\infty))$ is equal to $ap^\mu + b - g + 1$. Therefore, if

$$(a-g+1)(b-g+1) > ap^\mu + b - g + 1,$$

there is a non-zero function $f = \sum_{ij} c_{ij} e_i^{p^\mu} f_j^q$ in the kernel of this homomorphism. This function f has the property that

$$f(P) = \sum_{ij} c_{ij} e_i(P)^{p^\mu} f_j(P)^q = \sum_{ij} c_{ij} e_i(P)^{p^\mu} f_j(P) = 0$$

whenever $P \in X(\mathbf{F}_q)$. Since $p^\mu < q$, this shows that f has a zero of multiplicity at least p^μ in every $P \in X(\mathbf{F}_q) - \{\infty\}$. On the other hand the order of its pole at ∞ is at most $ap^\mu + bq$. It follows that

$$p^\mu(\#X(\mathbf{F}_q) - 1) \leq \#\{\text{zeroes of } f\} = \#\{\text{poles of } f\} \leq ap^\mu + bq.$$

Substituting $p^\mu = \sqrt{q}$ and $a = \sqrt{q} - 1$ this means that

$$\#X(\mathbf{F}_q) \leq (b+1)\sqrt{q}.$$

Now we choose b. This number should be at least $2g-2$ and satisfy $(a-g+1)(b-g+1) > ap^\mu - g + 1$. This means that any choice of b must satisfy

$$b > g - 1 + \frac{(q-1)\sqrt{q}}{\sqrt{q} - 1 - g}.$$

Note that the right hand side of this inequality is larger than $2g-2$. We choose b as small as possible. This b satisfies $b \le g + \frac{(q-1)\sqrt{q}}{\sqrt{q}-1-g}$ so that

$$\#X(\mathbf{F}_q) \le (1+g+\frac{(q-1)\sqrt{q}}{\sqrt{q}-1-g})\sqrt{q} \le \sqrt{q}(g+1)+q+\frac{g}{\sqrt{q}-1-g}$$

which is at most $q+(3g+1)\sqrt{q}$ since $\sqrt{q} > 2g+2$. This proves the Proposition.

Theorem 3. *Let X be a curve of genus g over a finite field k. Then there is an extension field $k \subset \mathbf{F}_q$ for which q is a square larger than $(2g+2)^2$ and*

$$\#X(\mathbf{F}_{q^k}) = q^k + O(q^{k/2})$$

for all $k > 0$. Here the O-symbol depends only on the curve X.

Proof. Let $f \in \mathbf{F}_q(X)$ be a function that is not a p-th power. Then the field extension $\mathbf{F}_q(f) \subset \mathbf{F}_q(X)$ is separable and corresponds to a non-constant separable morphism $\pi_X : X \longrightarrow \mathbf{P}^1$. Let K be a finite extension of $\mathbf{F}_q(X)$ that is also a Galois extension of $\mathbf{F}_q(f)$. Then K is the function field of an absolutely irreducible projective curve Y over $\mathbf{F}_{q^k}$ for some $k \ge 1$. It may happen that the field of constants $\mathbf{F}_{q^k}$ of K is strictly larger than $\mathbf{F}_q$. Extending the field of constants of $\mathbf{F}_q(f)$ and $\mathbf{F}_q(X)$ if necessary, we may assume that all three curves $\mathbf{P}^1$, X and Y have the same field of constants, which we denote by $\mathbf{F}_q$ again.

Corresponding to the field inclusions $\mathbf{F}_q(f) \subset \mathbf{F}_q(X) \subset \mathbf{F}_q(Y) = K$ we have the $\mathbf{F}_q$-morphisms

$$Y \quad \xrightarrow{\pi} \quad X \quad \xrightarrow{\pi_X} \quad \mathbf{P}^1.$$

Let $\pi_Y = \pi_X \cdot \pi$ be the composite morphism $Y \longrightarrow \mathbf{P}^1$. Let $G = \mathrm{Gal}(K/\mathbf{F}_q(f)) = \mathrm{Gal}(Y/\mathbf{P}^1)$ and let H denote its subgroup $\mathrm{Gal}(K/\mathbf{F}_q(X)) = \mathrm{Gal}(Y/X)$.

Fix $k \ge 1$. Let A denote the set of unramified points $P \in Y(\overline{\mathbf{F}}_q)$ whose image $\pi_Y(P)$ is in $\mathbf{P}^1(\mathbf{F}_{q^k})$. Since $\mathbf{P}^1$ has q^k+1 rational points over $\mathbf{F}_{q^k}$, it is immediate that

$$\#A = (q^k+1)\#G + O(1)$$

where $O(1)$ indicates a number that is at most the number of ramification points of the covering Y of $\mathbf{P}^1$ over $\overline{\mathbf{F}}_q$. In particular, it is bounded independently of k.

Let φ denote the *Frobenius morphism.* It raises the coordinates of a point to the power q. The $\mathbf{F}_q$-morphisms π, π_X and π_Y introduced above all commute with φ. It follows that f or every unramified point $P \in A$ the conjugate point $\varphi(P)$ maps to the same point in $\mathbf{P}^1$ as P does. Therefore there is a unique $\sigma \in G$ such that $\varphi(P) = \sigma(P)$. For every $\sigma \in G$ put

$$A_\sigma = \{P \in A\colon \varphi(P) = \sigma(P)\}.$$

The set A is a disjoint union of the A_σ's.

By taking q so large that $q > (2g+2)^2$ we may assume that there is a point $\infty \in Y(\mathbf{F}_q)$ and that the conditions of Proposition 2 are satisfied. Let $\sigma \in G$. If $\sigma = \mathrm{id}_Y$, Proposion 2 implies that $\#A_\sigma \le q^k + (3g+2)\sqrt{q^k}$. We modify the proof of Prop. 2 and show that the same inequality holds for every σ. We let $a = \sqrt{q} - 1$ and $p^\mu = \sqrt{q}$ and consider the vector space $H^0(a(\infty))^{p^\mu} H^0(b(\infty))^q$ of functions in $\mathbf{F}_{q^k}(Y)$. We map it to $H^0(p^\mu a(\infty) + b(\sigma^{-1}(\infty))$ by mapping, in the notation of the proof of Prop. 2, the basis vector $e_i^{p^\mu} f_j^q$ to $e_i^{p^\mu} \cdot (f_j \circ \sigma)$. Any function $f = \sum_{ij} c_{ij} e_i^{p^\mu} f_j^q$ in the kernel of this map has the property that

$$f(P) = \sum_{ij} c_{ij} e_i(P)^{p^\mu} f_j(P)^q = \sum_{ij} c_{ij} e_i(P)^{p^\mu} f_j(\varphi(P)) = \sum_{ij} c_{ij} e_i(P)^{p^\mu} f_j(\sigma(P)) = 0.$$

for every point $P \in A_\sigma$. The same arguments as in Prop. 2 now show that

$$\#A_\sigma \leq q^k + (2g+3)\sqrt{q^k}$$

for *every* $\sigma \in G$. Since

$$\sum_{\sigma \in G} \#A_\sigma = \#A = (q+1)\#G + O(1),$$

there is a constant $C > 0$ so that $\#A_\sigma \geq q^k + 1 - Cq^{k/2}$. Since $\cup_{\sigma \in H} A_\sigma$ is precisely the set of unramified points $P \in Y(\overline{\mathbf{F}}_q)$ for which $\pi(P) \in X(\mathbf{F}_q)$, we see that

$$\#H \cdot \#X(\mathbf{F}_q) = \sum_{\sigma \in H} \#A_\sigma = \#H(q^k + 1) + O(q^{k/2})$$

as $k \to \infty$. This proves the Theorem.

Corrolary 3. *(A. Weil, 1948)* Let X be a curve of genus g over a finite field $\mathbf{F}_q$. Then the reciprocal roots $\phi \in \mathbf{C}$ of the function $Z_X(T)$ satisfy

$$|\phi| = \sqrt{q}.$$

Proof. Since the reciprocal zeroes of the zeta function of X over $\mathbf{F}_{q^m}$ are the m-th powers of the zeroes of $Z_X(T)$, it suffices to give the proof for a power of q. We will call this power q again and choose it so large that the condition of Proposition 2 is satisfied: q is a square exceeding $(2g+2)^2$ and the curves X, Y and the morphisms π_Y, π_X and π of Proposition 2 are all defined over $\mathbf{F}_q$. We deduce that for large enough k we have that

$$\#X(\mathbf{F}_{q^k}) = q^k + 1 + O(q^{k/2}).$$

and therefore, with the usual notation, that

$$\sum_\phi \phi^k = O(q^{k/2}).$$

This implies that the function $f(z) = \sum_\phi \frac{1}{1-\phi z}$ has a radius of convergence at least as large as $q^{-1/2}$. Therefore $|\phi| \leq \sqrt{q}$ for each ϕ. The theorem now follows from the functional equation satisfied by $Z_X(T)$: when ϕ is a reciprocal root, so is q/ϕ and it follows that $|\phi| = \sqrt{q}$ for all ϕ as required.