Trace Codes Arising from Algebraic Curves

Gil Cohen Dean Doron Noam Goldgraber Tomer Manket

Taking a Swing at Gilbert-Varshamov

Gil Cohen Dean Doron Noam Goldgraber Tomer Manket

Rate-distance tradeoff: AG vs. GV

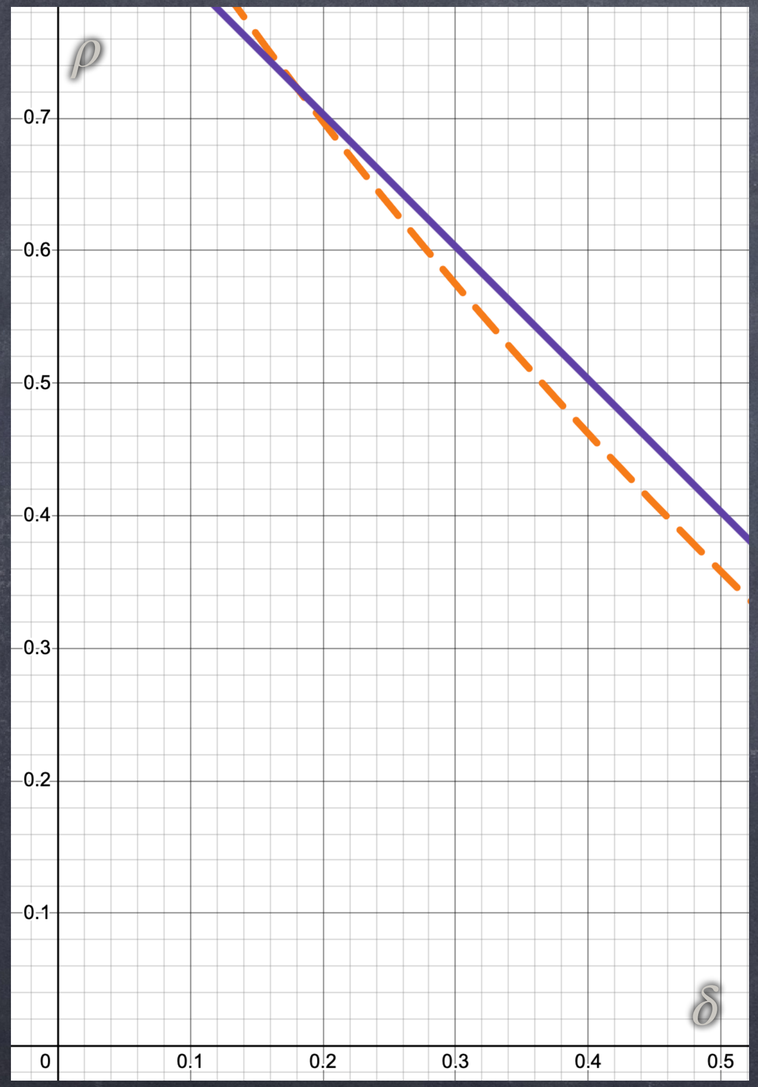

GF(q)

GV $\rho + \delta = 1 - \dfrac{1}{\Omega_\delta(\ln q)}$

AG $\rho + \delta = 1 - \dfrac{1}{\sqrt{q} - 1}$

RS $\rho + \delta = 1$

Our Approach

* Start with AG
* Reduce alphabet
* Keep tradeoff strong

Code Concatenation

Pro: Plug-and-play.

Con: Blind to structure, cares only for distance & rate.

Tracing the Code

Pro: Goes beyond black-box combinatorics

Con: More challenging — requires deeper math.

Tracing the History of AG Codes

Trace codes of AG codes were studied since the 90s, mainly in connection with subfield subcodes [Delsarte '75, var der Vlugt '91, ...].

A closely related recent result to ours is on character sums over AG codes by [Kopparty, Ta-Shma, Yakirevitch '24].

Warm-Up:
Tracing Reed-Solomon

Tracing Reed-Solomon

Take $q = p^m$ and recall $Tr : \mathrm{GF}(q) \to \mathrm{GF}(p)$ is given by

$$Tr(x) = x^p + x^{p^2} + \cdots + x^{p^m}$$

Messages correspond to degree d polynomials over GF(q)

The codeword corresponding to f is

$$C(f) = \left(\mathrm{Tr}(f(P_1)), \ldots, \mathrm{Tr}(f(P_q))\right)$$

Analysis

$$C(f) = \left(\mathrm{Tr}(f(P_1)), \ldots, \mathrm{Tr}(f(P_q))\right)$$

Consider the curve given by

$$y^p - y = f(x)$$

Let N_f = # GF(q)-points on the curve.

Claim

$$\frac{N_f}{p} = \#\text{zeros in } C(f)$$

The proof follows from Hilbert's Theorem 90.

In our Artin-Schreier extension $y^p - y = f(x)$, we have

(Corollary of) The Hasse-Weil Theorem

$$|N_f - (q+1)| \leq d\sqrt{q}$$

⇒ For p=2, Tr∘RS has distance $\delta = \frac{1}{2} - \varepsilon$ with $n = \frac{k^4}{\varepsilon^2}$

In comparison, RS∘Had satisfies $n = \frac{k^2}{\varepsilon^2}$ [AGHP '92], whereas GV achieves $n = \frac{k}{\varepsilon^2}$. The state-of-the-art is $\frac{k}{\varepsilon^{2+o(1)}}$ [Ta-Shma '17].

What have we learned?

- Tr∘RS inevitably leads to algebraic curves.
- Trace degrades the rate–distance tradeoff.

Our approach

Study Tr∘AG codes — shorter traces may reduce the damage and still keep Tr∘AG ahead of GV.

Problem! Hasse–Weil breaks down.

Tracing AG Codes

Tracing RS

Messages correspond to degree d polynomials over GF(q)

The codeword corresponding to f is

$$C(f) = \left(\mathrm{Tr}(f(P_1)), \ldots, \mathrm{Tr}(f(P_q))\right)$$

Tracing AG

Messages correspond to "degree" d functions over the curve F

The codeword corresponding to f is

$$C(f) = \left(\mathrm{Tr}(f(P_1)), \ldots, \mathrm{Tr}(f(P_{N_F}))\right)$$

Analysis

$$C(f) = \left(\mathrm{Tr}(f(P_1)), \ldots, \mathrm{Tr}(f(P_{N_F}))\right)$$

Consider the curve given by

$$z^p - z = f(x, y)$$

Let N_F = #points on the original curve F

N_L = #points on the extended curve L

Claim

$$\delta = 1 - \frac{N_L}{pN_F}$$

Recall the Hasse-Weil Theorem

$$|N - (q+1)| \leq 2g\sqrt{q}$$

Grothendieck's Trace Formula gives

$$|N_L - N_F| \leq 2(g_L - g_F)\sqrt{q}$$

Doesn't help in our setting.

$$|N_L - N_F| \leq 2(g_L - g_F)\sqrt{q}$$

By Hurwitz Genus Formula

$g_L = p g_F + \Delta \geq 2g_F$ where $\Delta \approx pd$ is the Different degree

Ideally, we seek a relative Hasse–Weil bound—one that doesn't "pay" g_F and depends only on Δ.

Our Contribution

The core mathematical result of this work is

Theorem

For "nice enough" F and f,

$$|N_L - N_F| = O\left(\sqrt{\Delta \cdot g_F}\right) \cdot \sqrt{q}$$

- Applies well beyond Artin–Schreier (e.g., Kummer extensions).
- New exponential-sum bounds as a direct corollary.

Tracing Herm

Instantiating with the Hermitian function field, Tr∘Herm can have distance $\delta = 1/2 - \varepsilon$ with

$$n = \left(\frac{k}{\varepsilon^4}\right)^{3/2}$$

Not as good as Herm∘Had [Ben-Aroya Ta-Shma '13] which gives

$$n = \left(\frac{k}{\varepsilon^2}\right)^{5/4}$$

Broader Perspective

Full Spectrum Analysis

Code Concatenation

Pro: Plug-and-play.

Con: Blind to structure, cares only for distance & rate.

Tracing the Code

Pro: Goes beyond black-box combinatorics

Con: More challenging — requires deeper math.

Zig-Zag Product and Friends

Pro: Plug-and-play.

Con: Blind to structure, cares only for spectral expansion.

Marcus-Spielman-Srivastava and Friends

Pro: Goes beyond black-box combinatorics

Con: More challenging — requires free probability theory.

L-functions

Each curve has an associated L-function—a degree-$2g$ polynomial

$$L(t) = \prod_{i=1}^{2g} (1 - \alpha_i t)$$

such that

$$-\sum_{i=1}^{2g} \alpha_i = N - (q+1)$$

The Hasse-Weil Theorem states that $|\alpha_i| = \sqrt{q}$.

Full-Spectrum Analogy

Shared setup: Each object carries an associated polynomial satisfying a functional equation.

Objective: Study the trace associated with the object.

Approach: Incorporating the object's entire spectrum.

Payoff: Capture the global spectral structure, hopefully yielding tighter results.

Thanks!

Proof idea & related work

Our proof is based on a variant of Bombieri's proof for Hasse-Weil.

Kopparty, Ta-Shma, and Yakirevitch obtained related results for Kummer extensions (multiplicative characters).

Our proof, when restricted to Kummer extensions, is simpler (no derivatives) and has additional technical advantages.

Backup slides

Theorem (RVW)

For all graphs G,H

$$\omega_{G \circ H} \leq \omega_G + \omega_H$$

By exploiting the entire spectrum of the graphs, one can prove

Theorem (CCM)

For every H there exists G s.t.

$$\omega_Z = min \frac{1}{x} \sqrt{1 - \frac{h(x)}{x h'(x)}}$$

The core mathematical result of this work is

Theorem

$$|N_L - N_F| = O\left(\sqrt{(g_L - pg_F)g_F}\right)$$

Put differently, if $t = \tau g_F$ then

$$\frac{N_L}{N_F} = 1 + O(\sqrt{\tau})$$

Thus, we are forced to work below the genus.

With our bound, one cannot obtain anything stronger than

$$n = \frac{k}{\varepsilon^4}$$

Still, this work is the first to provide a non-trivial guarantee for Tr∘AG, marking the starting point for future works.

Our Approach

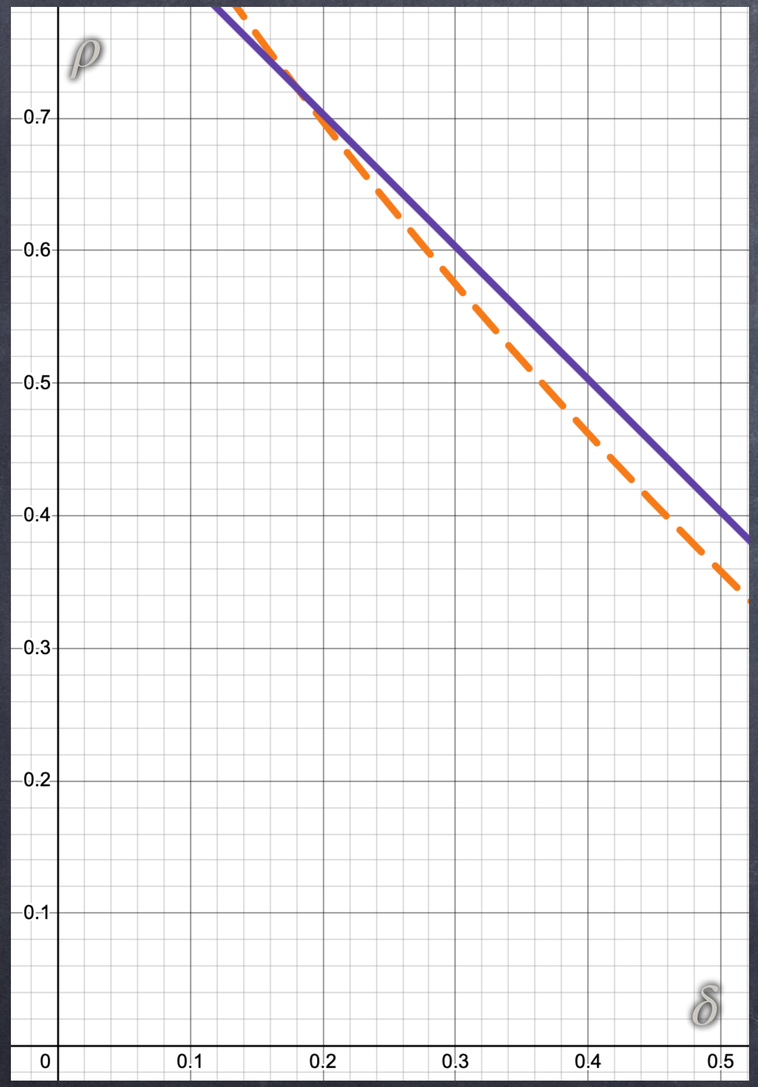

GF(q)

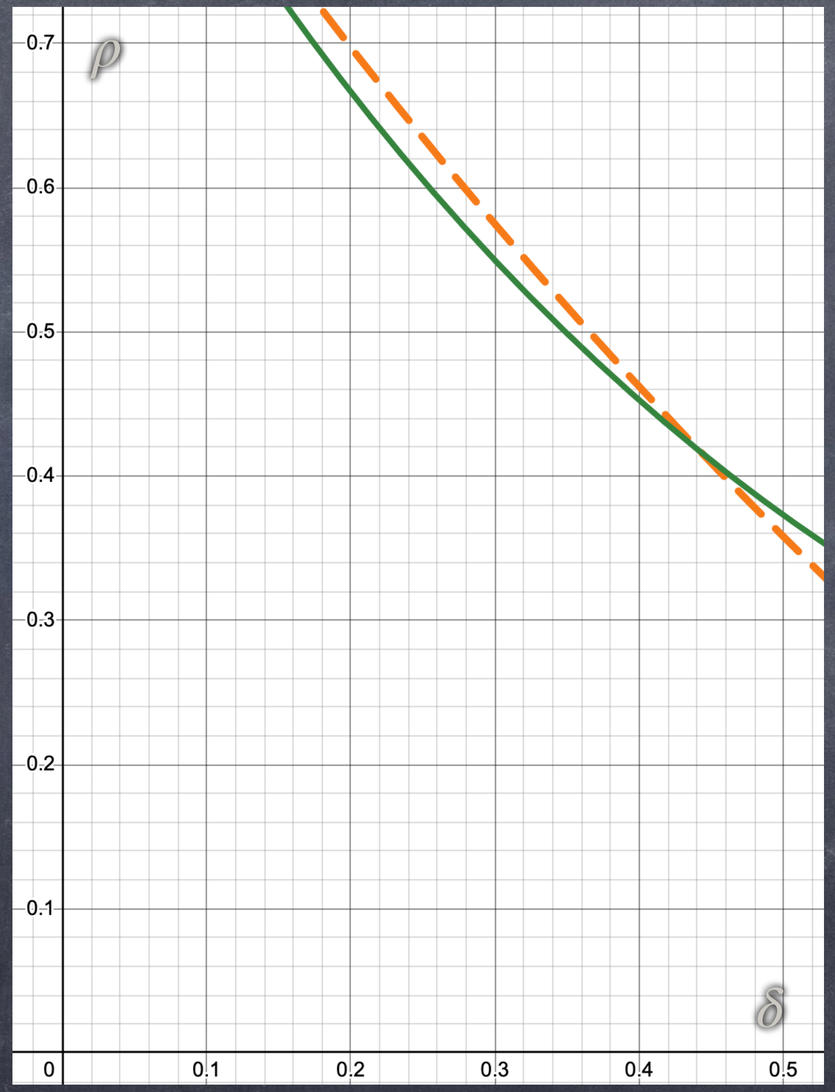

GF(2) - a wishful sketch

* Start with AG
* Reduce alphabet
* Keep tradeoff strong

Weierstrass Poles & Gaps

On a curve, every function has a "degree".

Not all degrees are realized by functions, but:

* From $2g$ upwards all are.
* Below $2g$ only half are attainable.

Rate-distance tradeoff: AG vs. GV

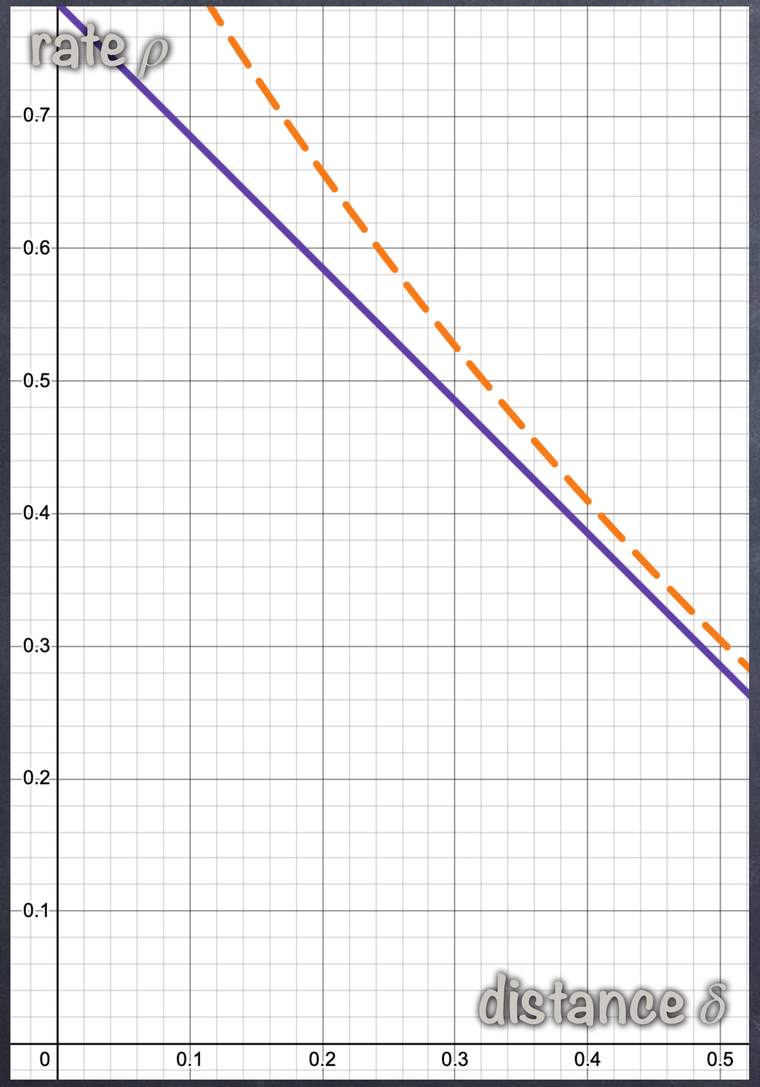

GF(32)

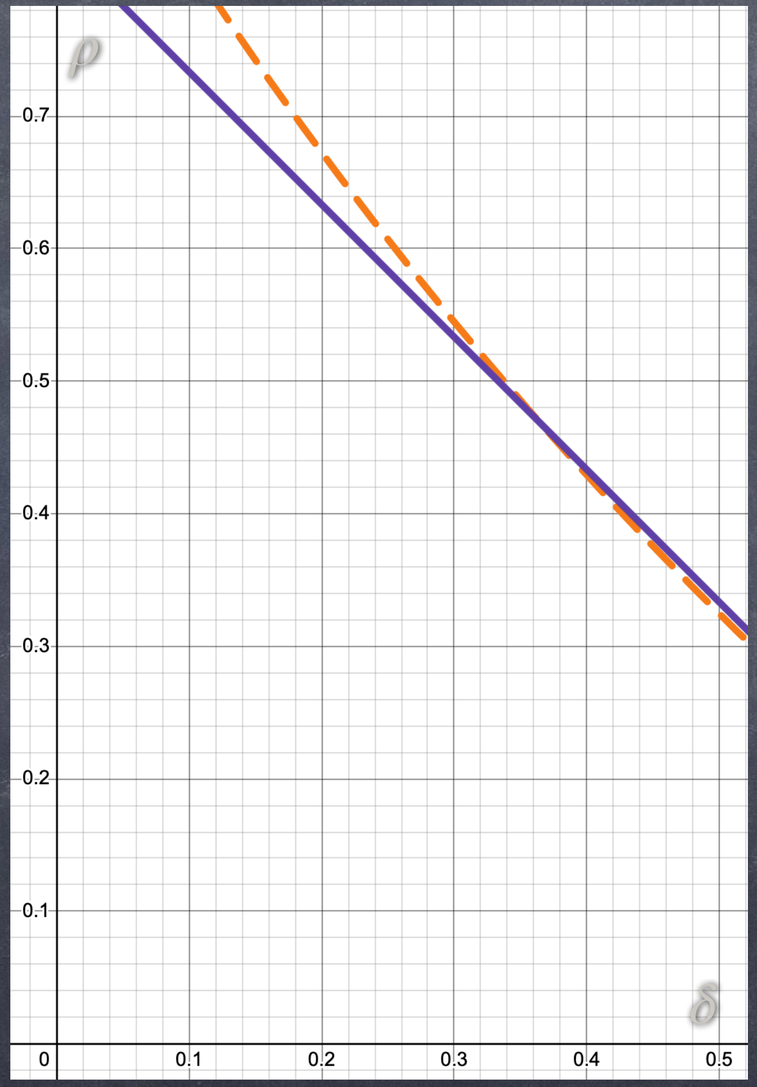

GF(49)

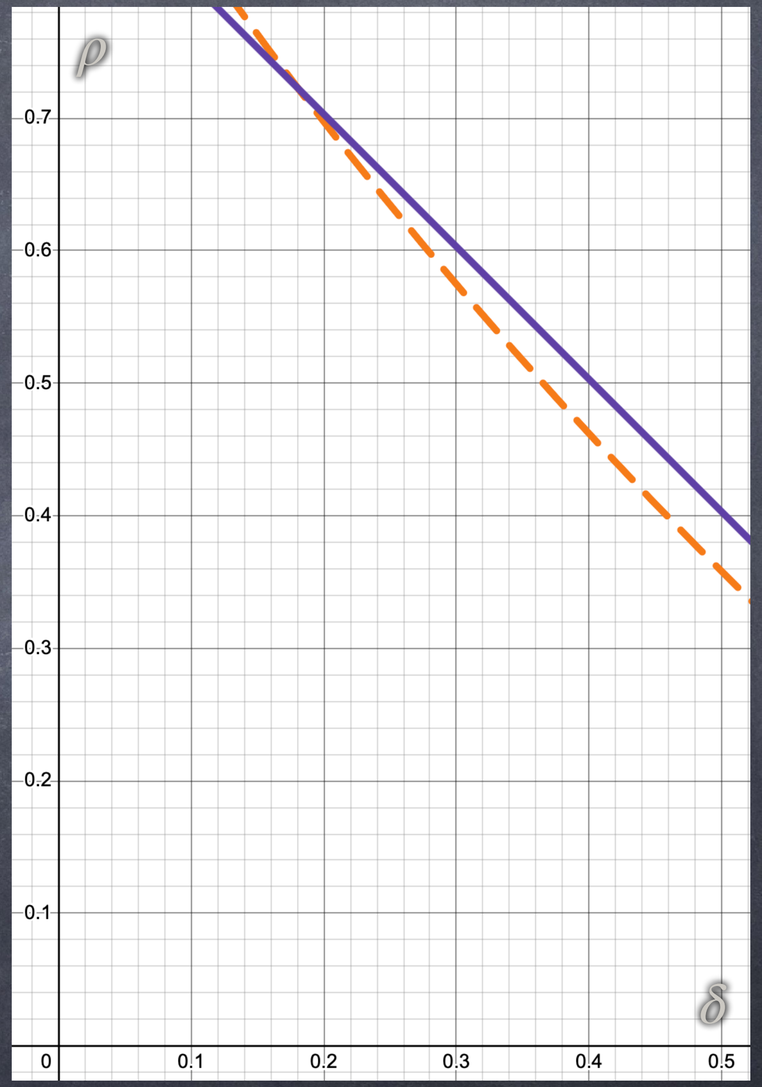

GF(256)

An algebraic curve is defined by polynomial equations: 'one fewer' equation than variables.

Example: The Hermitian Curve

Assume $q = r^2$

$$y^r - y = x^{r+1}$$

$$z^r - z = y^{r+1}$$

Algebraic Curves 101

An algebraic curve is defined by polynomial equations: 'one fewer' equation than variables.

Two key parameters of a curve: number of points N and genus g.

Genus? We'll use an informal, operative point of view.

The Hasse–Weil Theorem

$$|N - (q+1)| \leq 2g\sqrt{q}$$