# Tracing AG Codes: Toward Meeting the Gilbert–Varshamov Bound

**Gil Cohen**
Tel Aviv University, Israel

**Dean Doron**
Ben-Gurion University, Be'er-Sheva, Israel

**Noam Goldgraber**
Ben-Gurion University, Be'er-Sheva, Israel
Tel Aviv University, Israel

**Tomer Manket**
Tel Aviv University, Israel

## Abstract

One of the oldest problems in coding theory is to match the Gilbert–Varshamov bound with explicit binary codes. Over larger – yet still constant-sized – fields, algebraic-geometry codes are known to beat the GV bound. In this work, we leverage this phenomenon by taking traces of AG codes. Our hope is that the margin by which AG codes exceed the GV bound will withstand the parameter loss incurred by taking the trace from a constant field extension to the binary field. In contrast to concatenation, the usual alphabet-reduction method, our analysis of trace-of-AG (TAG) codes uses the AG codes' algebraic structure throughout – including in the alphabet-reduction step.

Our main technical contribution is a Hasse–Weil–type theorem that is well-suited for the analysis of TAG codes. The classical theorem (and its Grothendieck trace-formula extension) are inadequate in this setting. Although we do not obtain improved constructions, we show that a constant-factor strengthening of our bound would suffice. We also analyze the limitations of TAG codes under our bound and prove that, in the high-distance regime, they are inferior to code concatenation. Our Hasse–Weil–type theorem holds in far greater generality than is needed for analyzing TAG codes. In particular, we derive new estimates for exponential sums.

**2012 ACM Subject Classification** Mathematics of computing → Coding theory; Theory of computation → Pseudorandomness and derandomization

**Keywords and phrases** Coding Theory, Algebraic Geometry, Finite Fields, Exponential Sums

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2026.68

**Category** Track A: Algorithms, Complexity and Games

**Related Version** *Full Version*: `https://arxiv.org/abs/2511.08788` [5]

**Funding** *Gil Cohen*: Supported by ERC starting grant 949499 and by the Israel Science Foundation grant 2989/24.
*Dean Doron*: Supported in part by NSF-BSF grant 2022644.
*Noam Goldgraber*: Supported by NSF-BSF grant 2022644.
*Tomer Manket*: Supported by ERC starting grant 949499.

**Acknowledgements** We are grateful to Amnon Ta-Shma for illuminating discussions and for sharing his insights on trace codes of algebraic geometric codes.

## 1 Introduction

The study of the tradeoff between rate and distance is as old as coding theory itself, dating back to the late 1940s and early 1950s. Gilbert [10] proved the existence of codes with distance $\delta$ and rate $\rho \geq 1 - H_2(\delta)$, and Varshamov [27] subsequently proved that such codes





can be taken to be linear. Naturally, the ultimate goal is to achieve *explicit* constructions, and indeed much effort has focused on obtaining efficiently encodable codes with a good rate-vs.-distance tradeoff.

For binary codes, notable early constructions include Justesen's code [13], the first explicit construction to get constant rate and constant relative distance, and expander-based codes [24, 19], which further enjoy very efficient decoding. The large distance regime, $\delta = 1/2 - \varepsilon$, has been the subject of extensive and fruitful research over the past decades, in part due to its importance in complexity theory via its relation to small-bias sets. The GV bound guarantees the existence of codes with such distance $\delta$ and rate $\Omega(\varepsilon^2)$. Earlier explicit constructions achieved suboptimal rates [1, 18, 2], but a breakthrough construction of Ta-Shma attained nearly optimal rate $\varepsilon^{2+o(1)}$ via a sophisticated expander-based bias-reduction technique [22]. Progress toward the GV bound in the regime where $\delta$ is bounded away from $\frac{1}{2}$ has been comparatively limited.

Constructing codes over large alphabets has proved easier, particularly with respect to the rate–distance tradeoff. Reed–Solomon codes are the prototypical example: they attain the optimal tradeoff by meeting the Singleton bound with alphabet size equal to the block length. Consequently, a particularly useful tool for constructing codes over small alphabets is *code concatenation*, which reduces the alphabet size by combining large-alphabet outer codes with short, small-alphabet inner codes, yielding arbitrarily long codes over a small alphabet. Importantly, in the $\delta = 1/2 - \varepsilon$ regime, several constructions are concatenation-based (say, [1, 2]), and in particular, one can show that concatenating optimal algebraic-geometric codes with the Hadamard code yield codes of rate $\Omega(\varepsilon^3)$.

Remarkably, the GV bound is suboptimal for sufficiently large constant alphabets. A major line of research, initiated by Goppa's 1981 introduction of algebraic-geometry (AG) codes [11], constructs codes from algebraic curves over finite fields. This approach laid the foundation for the breakthrough of Tsfasman, Vlăduţ, and Zink [25], who proved that for sufficiently large fields (e.g., already over $\mathbb{F}_{49}$), AG codes can asymptotically achieve a rate–distance tradeoff exceeding the GV bound. Their result relies on families of curves with many rational points relative to their genus. Subsequently, Garcia and Stichtenoth [8, 9] constructed explicit recursive towers of function fields attaining the Drinfel'd-Vlăduţ [29] upper bound on the number of rational points, thereby providing explicit families of AG codes that meet the Tsfasman–Vlăduţ-Zink bound in a fully constructive way. Moreover, as the field size increases, the quantitative improvement becomes more pronounced, and the range of parameters for which the GV bound can be exceeded broadens accordingly.

## 1.1 Our Approach: Trace-Based Alphabet Reduction for AG Codes

Inspired by the work of Kopparty, Ta-Shma, and Yakirevitch [14] and by a talk of Ta-Shma at the Simons Institute (Berkeley) [23], this paper studies a family of explicit constructions over constant-size fields that aim to meet – or even surpass – the GV bound, with a particular focus on binary codes. Our primary technical contribution is to develop tools for analyzing these candidate constructions and for understanding their limitations. The resulting statements are general and extend beyond the original motivation, with potential further applications.

Our idea is simple: we aim to leverage the underlying algebraic structure of AG codes – which enables them to beat the GV bound – in the alphabet-reduction step as well. The hope is that, by exploiting this structure, the reduction will not substantially compromise the excellent rate–distance tradeoff of the original AG code. In its most basic form, the alphabet-reduction method we propose applies the field trace to each coordinate of every

codeword. While this is the variant we focus on in this work, we view it as a special case within a broader family of constructions whose common feature is the aforementioned strategy of leveraging structure for alphabet reduction.

This approach is, in a sense, an antithesis of the off-the-shelf technique of the black-box analysis of code concatenation. The latter ignores the code's internal structure: the rate and distance of the resulting code are simply the products of the corresponding parameters of the outer and inner codes (see Section 3). One can hope to get better guarantees by exploiting additional, more complicated, structure.[1]

Concatenation is reminiscent of other composition-type primitives, such as the zig–zag product in expander-graph constructions, where the roles of rate and distance are played by the graph's degree and spectral expansion. In contrast, exploiting more of the underlying structure – in the graph setting, the entire spectrum – yields stronger, and in fact optimal, analyses [4].

Of course, this is not the first work to consider trace codes: they have been studied for decades, most notably since Delsarte [6], who established their connection to subfield subcodes and, in particular, to dual BCH codes. Trace codes of AG codes have likewise been investigated since the 1990s (e.g., [26, 20, 16] and subsequent works). However, the aspects examined in that literature, largely motivated by the analysis of BCH and cyclic codes, are of limited relevance to our goals here. The papers by Kopparty, Ta-Shma, and Yakirevitch [14, 15] as well as an earlier paper by Vlăduţ [28] are the most relevant to our work. We discuss Vlăduţ's work in Section 2.2 and [14, 15] in Section 2.3.

To clarify the challenges in analyzing trace codes of AG codes and to place our technical results in context, we must first turn to describe the simplest case: the trace of Reed–Solomon codes. This (by now standard) analysis exposes the intimate relationship between taking traces of codes – AG codes or otherwise – and the geometry of algebraic curves. It also shows why existing fundamental results from the theory of algebraic curves, while effective for Reed–Solomon codes, fall short for trace codes of general AG codes. With this perspective, we present our main result, which constitutes a first step toward overcoming these obstacles.

### Concurrent work

In concurrent, independent work, Kopparty, Ta-Shma, and Yakirevitch [15] also study trace codes of AG codes, focusing on the Hermitian function field; this extends their earlier paper [14], which inspired the present work. We give a technical comparison with [14, 15] after presenting our results (see Section 2.3).

## 1.2 Trace Codes of Reed–Solomon: Analysis via the Hasse–Weil Theorem

We recall that a Reed–Solomon (RS) code over $\mathbb{F}_q$ is defined by identifying messages with polynomials $f \in \mathbb{F}_q[x]$ of degree $< k$ and evaluating them at distinct field elements. Specifically, for evaluation points $\mathfrak{p}_1, \ldots, \mathfrak{p}_n \in \mathbb{F}_q$ (with $n \leq q$), the codeword corresponding to $f$ is $(f(\mathfrak{p}_1), \ldots, f(\mathfrak{p}_n))$. Let $p$ denote the characteristic of $\mathbb{F}_q$, and write $q = p^m$. The trace code of the Reed–Solomon code is obtained by applying the (absolute) field trace $\mathsf{Tr}$ to each

[1] We note that while our trace code approach is inherently "non black-box", there have been few attempts at exploiting structure in concatenation-based construction, most recently in [7], with a similar goal of attaining the GV bound.

coordinate. Thus the encoder maps $\mathbb{F}_q^k \to \mathbb{F}_p^n$ via

$$f \longmapsto \big(\mathsf{Tr}(f(\mathfrak{p}_1)), \ldots, \mathsf{Tr}(f(\mathfrak{p}_n))\big), \tag{1.1}$$

where $\mathsf{Tr}(x) = x + x^p + \cdots + x^{p^{m-1}}$. Note that the resulted code is $\mathbb{F}_p$-linear. As it stands, as long as the encoder is injective, the rate of the code is

$$\rho = \frac{k \log q}{n \log p} = \frac{mk}{n}, \tag{1.2}$$

which is a factor-$m$ improvement over the rate of the underlying RS code. We now turn to the distance analysis, which is more challenging.[2]

Fix a polynomial $f \in \mathbb{F}_q[x]$ of degree $t < k$. Using Hilbert's Theorem 90, it can be shown that the number of zeros $z_f$ in the codeword corresponding to $f$ in Equation (1.1) is related to the number of pairs $(x, y) \in \mathbb{F}_q^2$ satisfying

$$y^p - y = f(x), \tag{1.3}$$

that is, to the number of $\mathbb{F}_q$-rational points $n_f$ on this curve. More precisely, we have that

$$z_f \leq \frac{n_f}{p}, \tag{1.4}$$

and equality holds when $n = q$. Therefore, analyzing the distance of the trace code is *equivalent* to counting points on curves, one curve for each message $f$.

Counting points on curves over finite fields is a difficult task. A deep and powerful result, the Hasse–Weil theorem (see [21, Chapter 5]), provides a bound on the number of points. We will discuss the Hasse–Weil theorem in more depth later on; for now it suffices to note that in our current setting, where the curve has the form of Equation (1.3), the theorem implies that the number of $\mathbb{F}_q$-rational points $n_f$ satisfies

$$|n_f - (q+1)| \leq (t-1)(p-1)\sqrt{q}.\ ^3 \tag{1.5}$$

Combining this with Equation (1.4), we readily obtain a bound on the distance. Assuming, for simplicity, that we evaluate over all field elements (i.e., $n = q$), we get

$$\delta \geq 1 - \frac{1}{p} - \frac{1}{np} - \frac{p-1}{p} \cdot \frac{k-2}{\sqrt{n}} \geq 1 - \frac{1}{p} - \frac{k}{\sqrt{n}}.$$

Since one cannot expect a distance better than $1 - \frac{1}{p}$, we see that the "loss" term is $\frac{k}{\sqrt{n}}$. This, in particular, means that to obtain a nontrivial bound on the distance we must take $k = O(\sqrt{n})$, which forces the rate to vanish at an inverse–square-root rate, $\rho = O\Big(\frac{1}{\sqrt{n}}\Big)$.

The above construction extends naturally to trace codes of AG codes, and this is what we undertake in the following sections. However, the distance analysis based on the Hasse–Weil theorem does not carry over. This motivates our results, and in particular our main theorem: a Hasse–Weil–type bound suitable to the analysis of trace codes of AG codes, which yields a meaningful lower bound on their distance.

To keep this introductory section accessible, we continue to assume no prior knowledge of algebraic function fields. As a result, this section has an expository flavor, while the formal treatment appears later in the technical sections. In particular, in Section 1.3 we provide a brief, informal introduction to algebraic curves.

---

[2] As it turns out, analyzing the distance will require a slight tweak to the construction, incurring a small loss in the rate computed in Equation (1.2).

## 1.3 A Brief Introduction to Algebraic Curves

Informally, an algebraic curve over a finite field $\mathbb{F}_q$ is the set of points in $\mathbb{F}_q^m$ satisfying $m-1$ independent polynomial equations. A good example to have in mind is the Hermitian plane curve over a field of size $q = r^2$ for some prime power $r$, consisting of all points $(x, y) \in \mathbb{F}_q^2$ satisfying $y^r + y = x^{r+1}$. It is not hard to show that this curve has $r^3 = q^{3/2}$ points. In general, the number of points $n$ on a curve is a key parameter; in the coding-theoretic context, it governs the block length of the associated code, as we shall see. Generally, these points are denoted by $\mathfrak{p}_1, \ldots, \mathfrak{p}_n$.

AG codes are obtained by evaluating functions on a fixed curve. The notion of a function on a curve is somewhat delicate. For example, on the Hermitian curve the two polynomials $y^r$ and $x^{r+1} - y$, although different as formal polynomials, are identical as functions, since they agree on all points of the curve. Nonetheless, there is a notion analogous to the *degree* of a function irrespective of its representation as a polynomial. As with polynomials, this notion of degree bounds the number of zeros a function may have on the curve. Moreover, the set of all functions of degree at most $k$ forms an $\mathbb{F}_q$-vector space.

Given a curve, we associate the corresponding AG code in a manner analogous to the Reed–Solomon code. We fix a degree $k$ and identify the messages with the subspace of functions of degree at most $k$. The codeword corresponding to such a function $f$ is obtained by evaluating $f$ at all points on the curve, $(f(\mathfrak{p}_1), \ldots, f(\mathfrak{p}_n)) \in \mathbb{F}_q^n$.

A second important parameter is the curve's *genus*. This natural number, denoted by $g$, is a measure of the curve's complexity. We will not give a formal definition of the genus here, but rather adopt the following operative perspective, which is based on the Hasse-Weil Theorem. The latter is fundamental in the study of algebraic curves over finite fields; it is, in fact, the proof of the Riemann Hypothesis over finite fields. We will discuss the theorem itself later; however, an important corollary – referred to here as the Hasse–Weil bound – states that the number of points $n$ on a curve of genus $g$ satisfies

$$|n - (q+1)| \leq 2\sqrt{q} \cdot g. \quad [4]$$

From this result we deduced Equation (1.5). Thus, the smaller the genus, the better the bound on the number of points on the curve. That is, we view the genus as the parameter that controls $|n - (q+1)|$ (up to the factor $2\sqrt{q}$).

## 1.4 TAG Codes

With the concepts in Section 1.3 in place, we are ready to give an informal definition of trace codes of AG codes (TAG codes, for short). Fix an algebraic curve over a finite field $\mathbb{F}_q$ of characteristic $p$, choose $n$ points $\mathfrak{p}_1, \ldots, \mathfrak{p}_n$, and identify the message space with functions of degree less than $k$. As in (1.1), we map a message $f$ by

$$f \longmapsto \big(\mathsf{Tr}(f(\mathfrak{p}_1)), \ldots, \mathsf{Tr}(f(\mathfrak{p}_n))\big) \in \mathbb{F}_p^n. \tag{1.6}$$

As in the analysis of the trace of RS codes in Section 1.2, here too there is a precise relation between the number of zeros $z_f$ in the codeword corresponding to $f$ and the number of points on a certain curve. This time, the latter curve depends on both the original curve and the function $f$. In particular, analogously to Equation (1.3), if the original curve lies in an ambient space of dimension $m$, then the new curve we consider lies in an ambient space of dimension $m+1$, and in addition to the $m-1$ polynomial relations among $x_1, \ldots, x_m$ dictated by the original curve, we have the relation

$$x_{m+1}^p - x_{m+1} = f(x_1, \ldots, x_m). \tag{1.7}$$

Analogously to Equation (1.4), there is an exact relation between the number of points $n_f$ on the new curve and $n$, the number of points on the original curve. It takes a slightly less simple form, which is quantitatively almost identical; for simplicity, in this informal section we will use the same relation as before, namely,

$$z_f = \frac{n_f}{p}. \tag{1.8}$$

Recall that, to analyze the distance of the trace of RS codes, we relied on the Hasse–Weil bound as stated in Equation (1.5). Looking more closely at that equation, the term $q$ counts the number of points on the base curve underlying the Reed–Solomon code, namely the line over $\mathbb{F}_q$. The Hasse–Weil bound controls the difference between the number of points on the base curve and the number of points on the new curve. The bound in Equation (1.5) is in terms of the genus of the new curve. What should the bound be in the case of TAG codes? There are now two genera involved: the genus of the curve defining the code, denoted by $g$, and the genus of the new, extended curve, denoted by $g_f$. In the original case, since the genus of the line is 0, it made no appearance in the Hasse–Weil bound. There is a sense in which the $\sqrt{q}$ appearing in the Hasse–Weil bound provides the natural "units" for measuring the number of points, independent of whether we extend the line or another curve (see Section 2.2). Thus, it is conceivable that the difference of genera should be used in the bound. Indeed, the general result in this context is Grothendieck's trace formula [12] which states exactly this; namely, it yields the bound

$$|n_f - n| \leq 2\sqrt{q} \cdot (g_f - g). \tag{1.9}$$

Unfortunately, this bound, which is tight in general, doesn't give any nontrivial bounds on the distance of TAG codes. To see why, we rely on two results from the theory of algebraic curves. First, there is a formula relating the genus $g_f$ of the extended curve and the genus $g$ of the base curve, known as the Hurwitz genus formula. For our introductory purposes it suffices to say that

$$g_f = pg + \Delta, \tag{1.10}$$

where $\Delta \geq 0$ is an integer. [5] Hence, the bound given by Equation (1.9) is at least $2\sqrt{q}\,(p-1)g$.

The second result is the Drinfel'd–Vlăduţ bound [29], which, informally, states that asymptotically, as the genus of the curve tends to infinity, we have

$$\frac{n}{g} \leq \sqrt{q} - 1. \tag{1.11}$$

Combining this with the calculation above shows that the bound in Equation (1.9) is worse than $2(p-1)n$. Consequently, the resulting upper bound on $n_f$ is no better than $(2p-1)n$. Together with Equation (1.8), this yields an upper bound of $\frac{(2p-1)n}{p}$ on the number of zeros in a codeword, which is trivial since this quantity exceeds $n$.[6]

As discussed above, a distance analysis for TAG codes demands estimates stronger than those yielded by Grothendieck's trace formula. This is the main technical contribution of the present work. In the next section we state our main result and its applications to the analysis of TAG codes, as well as its broader consequences.

---

[5] For readers familiar with algebraic function fields, $\Delta$ is the degree of the Different divisor of the corresponding function field extension.

[6] We remark that even a bound of the form $\sqrt{q}\,(g_f - g)$ in Equation (1.9) – which may be obtainable asymptotically – still does not yield a nontrivial distance bound.

## 2 Our Results

As discussed in Section 1.4, the issue with Grothendieck's trace formula for analyzing the distance of TAG codes is that the upper bound it yields for $|n_f - n|$ depends too heavily on the genus $g$ of the underlying curve. In other words, regardless of how simple the extension is (equivalently, regardless of the degree $t$ of $f$), the bound is too loose when the underlying function field is of large genus. Indeed, the degree of $f$ is encoded in $\Delta$, and the previous section showed that even if the contribution of $\Delta$ to the bound is ignored, one still does not obtain a meaningful bound on the distance. Thus, we seek a bound that depends on the complexity of the *extension* as captured by $\Delta$, rather than on how complex is the base curve, as encoded by the genus $g$. More precisely, one can show that $\Delta = \Theta(pt)$, and in particular it can be taken quite significantly smaller than the genus $g$.

A bound of the form

$$|n_f - n| \;=\; \Phi(\Delta)\,\sqrt{q}, \tag{2.1}$$

for some function $\Phi$ (e.g., linear in $\Delta$) is most desirable, as it is completely independent of the genus $g$. Our main technical contribution is a step toward such a result, in which the bound we obtain is the *geometric mean* of $\Delta$ (the quantity we wish to appear in the bound) and the genus $g$, which by itself is too large.

For the analysis of TAG codes we only need to consider specific types of extensions, as in Equation (1.7). Such extensions are called Artin–Schreier extensions. Our result, however, holds in much greater generality. In this introductory section we choose to be somewhat informal and consider only the special case required for analyzing TAG codes.

▶ **Theorem 1** (main result; informal). *Let $\mathbb{F}_q$ be a finite field of characteristic $p$. Let $C$ be a "nice" curve over $\mathbb{F}_q$ with genus $g$ and $n$ points. Let $f$ be a "nice" function on $C$ of degree $t$, and define the curve $C_f$ by the additional polynomial constraint $y^p - y = f$, where $y$ is a new formal variable. Then the number of points $n_f$ on the curve $C_f$ satisfies*

$$|n_f - n| = O\big(p^2\sqrt{tg}\sqrt{q}\big)\,. \tag{2.2}$$

In fact, the result also extends to Kummer extensions and, more generally, to extensions of the form $\phi(y) = f$ for an arbitrary polynomial $\phi$, as formalized in Theorem 17. The theorem further applies to even more general extensions under technical conditions – hidden in the two "nice" instances highlighted in the informal Theorem 1. For the complete, formal statement, see Proposition 14.

### 2.1 Implications for Error Correcting Codes

As discussed in Section 1.4, an upper bound on $n_f$ directly translates into a lower bound on the weight of the codeword corresponds to $f$ in the TAG construction. In this section, we examine the resulting codes parameters, also recognizing a possible limitation of this approach. We start by giving an example.

#### 2.1.1 The Hermitian TAG code

In this section, we illustrate Theorem 1 by giving an explicit instantiation of TAG codes based on the Hermitian curve introduced in Section 1.3. Let $r = p^\ell$ be a power of a prime $p$, and let $q = r^2$. Let $A \subseteq \mathbb{F}_q \times \mathbb{F}_q$ be the set of roots of the polynomial $y^r + y - x^{r+1}$. For $T \geq 0$, let

$$B = \{(i,j) \in \mathbb{N} \times \mathbb{N} : ir + j(r+1) \le T, \quad i \text{ is odd and } j \text{ is even}, \quad j < r/6\}. \tag{2.3}$$

The Hermitian TAG code is defined by

$$\mathcal{C} = \left\{ \left( \mathsf{Tr}\left( \sum_{(i,j)\in B} c_{i,j}\alpha^i\beta^j \right) \right)_{(\alpha,\beta)\in A} : \; c_{i,j} \in \mathbb{F}_q \right\}.$$

▶ **Theorem 2** (Hermitian TAG codes; informal)**.** *The Hermitian TAG code with message length $k$ and relative distance $1/2 - \varepsilon$ has block length*

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

In fact, the result shows that the Hermitian TAG code not only has relative distance $\frac{1}{2} - \varepsilon$ but is actually $\varepsilon$-*balanced*: every nonzero codeword has relative Hamming weight in $[\,\frac{1}{2} - \varepsilon, \frac{1}{2} + \varepsilon\,]$. In Section 7, we further instantiate our results for additional curves – specifically the Hermitian tower of function fields and the norm–trace function field – obtaining $\varepsilon$-balanced codes with varying parameter trade-offs.

### 2.1.2 The high distance regime: TAG codes vs. concatenation

As our case studies in Section 7 show, in the regime $\delta = \frac{1}{2} - \varepsilon$ all the TAG codes we consider are still far from the GV bound, requiring $n = \omega\left(\frac{k}{\varepsilon^2}\right)$. This prompts the question: does there exist a curve for which the corresponding TAG code matches the GV bound?

In Section 8 we present strong evidence for a negative answer. We show that, when analyzed using our bound from Theorem 1, in the high-distance regime $\delta = \frac{1}{2} - \varepsilon$, TAG codes instantiated from a given AG code are outperformed by concatenating the same AG code with Hadamard. This remains true even under a bound of the form Equation (2.1) with $\Phi(\Delta)$ that is linear in $\Delta$. In particular, we establish the following.

▶ **Theorem 3.** *Assume the bound given by Equation* (2.1) *is tight up to a constant. Then, any TAG code with message length $k$ and relative distance $1/2 - \varepsilon$ has block length*

$$n = \Omega\left(\frac{k}{\varepsilon^3}\right).$$

*Moreover, assuming the bound given in Theorem 1 is tight up to a constant, $n = \Omega(k/\varepsilon^6)$.*

### 2.1.3 The constant distance regime

As implied by Section 2.1.2, under our analysis, attaining the GV bound requires operating in the regime where the distance $\delta$ is bounded away from $\frac{1}{2}$. To approach the GV bound in this regime, we must have a nonvanishing rate. In this short section, we examine the implications for the parameters of the underlying AG code.

For a curve $C$, let $\ell(T)$ denote the dimension of the vector space of functions on $C$ of degree less than $T$. Assume that Theorem 1 applies to every function in this space, and let $\mathcal{C}$ be the resulting TAG code. By Equation (2.2) together with Equation (1.8), the encoding of a function $f$ of degree $t$ is nonzero provided the right-hand side is $< (p-1)\,n$.

For simplicity, assume the implicit constant hidden in the big-$O$ of Equation (2.2) is 1, and that $n/g = \sqrt{q} - 1$ (the optimal AG-code guarantee, meeting the Drinfel'd-Vlăduţ bound). To obtain a nontrivial bound on the number of zeros of $f$, we require

$$p^2\sqrt{tg}\,\sqrt{q} < pn,$$

which is equivalent (since $n = g(\sqrt{q} - 1)$) to

$$t < \frac{g}{p^2}\left(1 - \frac{1}{\sqrt{q}}\right)^2.$$

Hence, to apply our result we must work in the regime $T < g/p^2$. The resulting TAG code $\mathcal{C}$ has rate

$$\rho \;=\; \frac{\ell(T)\,\log_p q}{n}.$$

Since $\ell(T) \le T \le g/p^2$ and $n = (\sqrt{q} - 1)g$, we obtain

$$\rho \;=\; \frac{\log_p q}{\sqrt{q} - 1} \cdot \frac{\ell(T)}{g} \;\le\; \frac{\log_p q}{(\sqrt{q} - 1)\,p^2}.$$

Thus, to obtain a family with constant rate, we must fix $q$ to a constant and choose curves (and $T \le g/p^2$) so that $\ell(T) = \Omega_q(g)$.

This is a somewhat nonstandard regime for AG codes. Typically one works with functions of degree at least $2g$, where Riemann–Roch guarantees linear growth of $\ell(T)$ (indeed, $\ell(T) = T + 1 - g$ for $T \ge 2g - 1$). In particular, among degrees up to $(1 + \alpha)g$ one obtains at least $\alpha g$ attainable degrees. By contrast, working below $g$ is subtler: the attainable sub-$g$ degrees – the *Weierstrass non-gaps* – depend on the curve. In our TAG setting, one must therefore understand this sub-$g$ regime.

A property closely related to the latter was examined in [2, Section 4.1]. However, we are not aware of any optimal family of function fields exhibiting this behavior. For instance, in [30], the authors show that certain Riemann–Roch spaces on the Garcia–Stichtenoth tower do not satisfy this property. We formalize this as an open problem.

▶ **Open Problem 4.** *Does there exist a prime power $q = p^\ell$ and a family of function fields $F_i/\mathbb{F}_q$ with $\frac{n_i}{g_i} = \Theta_q(1)$, such that each $F_i$ of genus $g_i$ contains a Riemann–Roch space $\mathcal{L}_i$ of degree $T_i \le \frac{g_i}{p^2}$ satisfying $\ell_i(T_i) = \Omega_q(g_i)$?*

### 2.1.4 Limitations and consequences of improving the bound

In Section 2.1.3, we assumed that the constant appearing in the error term in Equation (2.2) is equal to 1. Suppose instead that we could obtain a small constant $0 < c < 1$ such that, for every "nice" function $f$ of degree $t$, the bound

$$n_f \le n + cp^2\sqrt{tg}\sqrt{q}$$

holds. If a $\beta$-fraction of all functions of degree up to $T$ would be "nice", then the resulting TAG code would have rate and relative distance, correspondingly,

$$\rho = \beta \cdot \frac{\log_p q}{\sqrt{q} - 1} \cdot \frac{\ell(T)}{g},$$
$$\delta \ge 1 - \frac{1}{p} - c \cdot \frac{p\sqrt{gT}\sqrt{q}}{n}.$$

Taking $T = 2g$, the Riemann–Roch theorem gives $\ell(T) = g + 1$. We then have

$$\rho \geq \beta \cdot \frac{\log_p q}{\sqrt{q} - 1},$$
$$\delta \geq 1 - \frac{1}{p} - cp\frac{\sqrt{2q}}{\sqrt{q} - 1}.$$

From this we see that if our bound were sharpened by reducing the leading constant, TAG codes would work in the usual AG regime (e.g., for degrees $\geq 2g - 1$) and could lead to better code parameters. Note that the constant $c$ cannot be taken arbitrarily small; doing so would violate the MRRW linear-programming upper bounds [17]. This offers an unusual direction of inference: using coding-theoretic limits to deduce statements about algebraic curves.

## 2.2 Exponential Sums over Curves

Exponential sums over finite fields play a central role in number theory and algebraic geometry. A. Weil was the first to observe a deep connection between exponential sums of polynomials and Artin–Schreier covers. Specifically, for a polynomial $f(x)$ of degree $d$ over $\mathbb{F}_q$ such that $f \neq u^p - u$ for all $u \in \overline{\mathbb{F}_q}(x)$, the exponential sum

$$S(f) = \sum_{\alpha \in \mathbb{F}_q} e^{\frac{2\pi i}{p} \mathsf{Tr}(f(\alpha))} \tag{2.4}$$

can be expressed as $S(f) = \omega_1 + \cdots + \omega_D$, where the $\omega_i$ are a subset of the reciprocals of the roots of the zeta function of the Artin–Schreier curve

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

Weil also showed that if $\deg f$ is relatively prime to $p$, then $D = d - 1$. Combined with his proof of the Riemann Hypothesis for curves over finite fields, which asserts that each $\omega_i$ satisfies $|\omega_i| = \sqrt{q}$, this yields the bound

$$|S(f)| \leq (d - 1)\sqrt{q}.$$

In 1966, Bombieri [3, Theorem 5] extended Weil's results from polynomials to general rational functions defined over algebraic curves.

▶ **Theorem 5** ([3], Theorem 5). *Let $C$ be a complete, irreducible, non-singular curve of genus $g$ over $\mathbb{F}_q$, which contains $n$ points. Let $f \in \mathbb{F}_q(C)$ be a function on $C$ such that $f \neq u^p - u$ for all $u \in \overline{\mathbb{F}_q}(C)$. Let $\mathcal{P}$ be the set of poles of $f$. Then the exponential sum*

$$S(f) = \sum_{\mathfrak{p} \in C \setminus \mathcal{P}} e^{\frac{2\pi i}{p} \mathsf{Tr}(f(\mathfrak{p}))}, \tag{2.5}$$

*can be expressed as*

$$S(f) = \omega_1 + \cdots + \omega_D, \tag{2.6}$$

*where $|\omega_i| = \sqrt{q}$ for all $1 \leq i \leq D$, and*

$$D \leq 2g - 2 + |\mathcal{P}| + \deg(f).$$

*Moreover, if $f$ has a single pole of degree relatively prime to $p$, then $D = 2g - 1 + \deg(f)$.*

▶ **Corollary 6** ([3]). *In the setting of Theorem 5, we have*

$$|S(f)| \leq (2g - 1 + \deg(f))\sqrt{q}.$$

Note that a trivial bound $|S(f)| \leq n - |\mathcal{P}|$ follows directly from Equation (2.5). Hence, the bound in Corollary 6 is non-trivial only when

$$(2g - 1 + \deg(f))\sqrt{q} < n - |\mathcal{P}|. \tag{2.7}$$

Corollary 6 was used in [16] to compute the dimension of certain TAG codes defined over curves for which Equation (2.7) holds.

There is a major limitation of Bombieri's bound in our setting: If the curve $C$ satisfies $2g\sqrt{q} \geq n$, then the bound becomes trivial for any function on the curve. By the Drinfeld–Vlăduţ bound, Equation (1.11), this inequality holds for all curves of sufficiently large genus. Since such curves are precisely those used in AG codes constructions, Bombieri's bound is not applicable in this regime.

In [28], Vlăduţ proves nontrivial bounds for the case of Hermitian and Hansen–Stichtenoth curves, showing that for certain functions $f$ of bounded degree one has $S(f) \leq cn$ for some constant $0 < c < 1$. Our results provide new bounds on exponential sums of functions over curves of large genus, yielding in particular that $S(f) = o(n)$ for functions $f$ of degree $o(g)$.

▶ **Theorem 7** (exponential sums; informal). *Let $\mathbb{F}_q$ be a finite field of characteristic $p$. Let $C$ be a "nice" curve over $\mathbb{F}_q$ with genus $g$ and $n$ rational points, and let $f$ be a "nice" function on $C$ of degree $t$. Then,*

$$S(f) = O\left(p^3\sqrt{tg}\sqrt{q}\right).$$

The reader is referred to Theorem 21 for the formal statement. Interestingly, in conjunction with Theorem 5, this implies that the complex roots $\omega_i$, each of absolute value $\sqrt{q}$, cannot be all in the same direction.

Note that this bound remains meaningful even for families of curves with $g \to \infty$. However, since $g\sqrt{q} = \Theta(n)$, as in Section 2.1.3, the function $f$ must satisfy $t \ll g$ for the bound to be non-trivial.

## 2.3 Comparison with [14, 15]

In this section, we compare our results with those obtained in an earlier work by Kopparty, Ta-Shma, and Yakirevitch [14], and with a concurrent work by the same authors [15]. These works focus primarily on the Hermitian function field. For a technical comparison, we instantiate Proposition 22 with this function field.

▶ **Theorem 8.** *Let $d > 1$ be such that $d \mid q - 1$. Let $f \in \mathcal{L}(q\mathfrak{P}_\infty)$ of degree $t = ir + j(r + 1)$, such that $(d, t) = 1$, $(i + j, d) = 1$ and $j < \frac{r}{3d}$. Then,*

$$\frac{1}{N}\left|\left\{\mathfrak{q} \in \mathbb{P}_F^1 \setminus \{\mathfrak{p}\} : \exists\, y \in \mathbb{F}_q^\times,\ y^d = f(\mathfrak{q})\right\}\right| \leq \frac{1}{d} + O\left(p\sqrt{\frac{t}{q}} + pd\,\frac{t}{q}\right).$$

The corresponding theorem from [14] goes as follows.

▶ **Theorem 9** ([14]). *Assume that $r > 500$ is a prime number, $q = r^2$ and let $d$ be a prime number that divides $q - 1$. Let $f \in \mathcal{L}(q\mathfrak{P}_\infty)$ of degree $t = ir + j(r + 1)$, such that $(d, t) = 1$, $(i + j, d) = 1$, and $d < O(\sqrt{t} + \frac{q^{3/2}}{t})$. Then,*

$$\frac{1}{N}\left|\left\{\mathfrak{q} \in \mathbb{P}_F^1 \setminus \{\mathfrak{p}\} : \exists\, y \in \mathbb{F}_q^\times,\ y^d = f(\mathfrak{q})\right\}\right| \leq \frac{1}{d} + O\left(\sqrt{\frac{t}{q}}\right).$$

There are several distinctions between these results. First, Theorem 9 is restricted to the case where $r$ is a prime number, reflecting the limitations of the derivatives method as used in [14]. In contrast, our result applies to general prime powers $r = p^e$; but the error term in our bound depends on $p$. For the purpose of TAG codes, one can take $p$ to be a constant, particularly $p = 2$ (and we are free to vary $r$) so it has no effect on our bound. Importantly, our result is meaningful for all $e > 2$ whereas Theorem 9 requires $e = 1$.

In [15], the authors extend the bound for the case $d = 2$ to all functions $f \in \mathcal{L}\left(q^{3/4}\,\mathfrak{P}_\infty\right)$ for which the polynomial $T^2 - f$ is absolutely irreducible, without imposing any restriction on $\deg(f)$. However, this comes at the cost of replacing the error term $O\left(\sqrt{\frac{t}{q}}\right)$ with $O\left(\frac{t}{q^{3/4}}\right)$. Moreover, in [15], the authors provide a bound on the exponential sum Equation (2.5) in the setting of the Hermitian curve, when $p = 2$ and $q$ is an arbitrary power of 2, for any function $f \in \mathcal{L}\left(q^{3/4}\,\mathfrak{P}_\infty\right)$ of odd degree; the resulting error term is $O\left(\frac{t}{q^{3/4}}\right)$, whereas our result Theorem 21 yields the sharper bound $O\left(\sqrt{\frac{t}{q}}\right)$ as $t = \Omega(\sqrt{q})$.

## 3 Preliminaries

We assume familiarity with basic background on algebraic function fields, such as places, valuations, extensions, decomposition and ramification of places, etc. A detailed exposition of the subject can be found in [21]. In the preliminaries section in the full version of the paper [5], we recall some basic notions concerning algebraic function fields, the trace map over finite fields, $\varepsilon$-balanced codes, and code concatenation.

## 4 TAG Codes and Function Field Extensions

In this short section, we formally define the construction of the trace code associated with an algebraic geometric code. We then review the relation of its minimum distance to the number of rational points in a certain elementary abelian $p$-extension of the underlying function field. This connection is a known result, which we include here for completeness.

Let $p$ be a prime number, and let $p \le \ell \le q$ be powers of $p$ such that $\mathbb{F}_p \subseteq \mathbb{F}_\ell \subseteq \mathbb{F}_q$. Throughout, $\mathsf{Tr}$ is the trace function from $\mathbb{F}_q$ down to $\mathbb{F}_\ell$. Let $F/\mathbb{F}_q$ be a function field of genus $g$ with $N = n + 1$ rational points, one of which is denoted by $\mathfrak{p}$, and the remaining by $\mathfrak{p}_1, \ldots, \mathfrak{p}_n$. For an integer $r$, we denote the set of pole numbers up to $r$ at $\mathfrak{p}$ by

$$\mathsf{WS}_r = \{i \in [r] : \mathcal{L}(i\mathfrak{p}) \neq \mathcal{L}((i-1)\mathfrak{p})\},$$

and for each $i \in \mathsf{WS}_r$, let $b_i \in \mathcal{L}(i\mathfrak{p}) \setminus \mathcal{L}((i-1)\mathfrak{p})$. We now focus on those functions whose pole order is coprime to the characteristic $p$: let $i_1 < i_2 < \cdots < i_k$ be the elements of $\mathsf{WS}_r$ such that $p \nmid i$, and define

$$B_r = \{b_i : i \in \mathsf{WS}_r \text{ such that } p \nmid i\}$$

to be the corresponding set of functions.

With these notations, we define the $\mathbb{F}_\ell$-linear code $\mathsf{TC}\colon \mathbb{F}_q^k \to \mathbb{F}_\ell^n$ as follows. Given $m = (m_1, \ldots, m_k) \in \mathbb{F}_q^k$, define the function $f_m = \sum_{j=1}^k m_j b_{i_j}$. Then, set

$$\mathsf{TC}(m) = \left(\mathsf{Tr}(f_m(\mathfrak{p}_1)), \ldots, \mathsf{Tr}(f_m(\mathfrak{p}_n))\right).$$

Put differently, we identify the domain of $\mathsf{TC}$ with $\mathsf{Span}_{\mathbb{F}_q}(B_r)$, and for a given input function $f \in \mathsf{Span}_{\mathbb{F}_q}(B_r)$, we evaluate it at the $n$ rational points $\mathfrak{p}_1, \ldots, \mathfrak{p}_n$, followed by applying the trace function from $\mathbb{F}_q$ to $\mathbb{F}_\ell$ coordinate-wise.

We now turn to analyze the distance of the code $\mathsf{TC}$. As hinted above, the key fact used is that the distance is closely related to the number of rational points in a certain extension of the function field $F$. To make this connection precise, fix a function $0 \neq f \in \mathsf{Span}_{\mathbb{F}_q}(B_r)$ and consider the field extension $L = F(z)$ defined by the equation

$$z^\ell - z = f.$$

We first note that $L/F$ is an elementary abelian $p$-extension, which is a generalization of Artin-Schreier extensions (which are the case $\ell = p$). Indeed, notice that the polynomial $a(T) = T^\ell - T$ is additive, and its set of roots is $\mathbb{F}_\ell \subseteq \mathbb{F}_q$. In addition, $v_{\mathfrak{p}}(f) < 0$ with $p \nmid v_{\mathfrak{p}}(f)$, and $v_{\mathfrak{q}}(f) \geq 0$ for all $\mathfrak{q} \in \mathbb{P}_F \setminus \{\mathfrak{p}\}$. Thus, all the conditions of [21, Proposition 3.7.10] are satisfied, and the extension $L/F$ is an elementary abelian $p$-extension of degree $\ell$.

▷ Claim 10. For every $i \in [n]$, we have $\mathsf{Tr}(f(\mathfrak{p}_i)) = 0$ if and only if the place $\mathfrak{p}_i$ splits completely in the extension $L/F$.

Proof. We refer the reader to the full version of the paper [5]. ◁

As a consequence of Claim 10, we obtain the following result. We denote by $N_L$ the number of rational points of $L$.

▶ **Corollary 11.** *The relative-distance of the code* $\mathsf{TC}$ *defined above is*

$$\delta = 1 - \frac{1}{\ell} \cdot \frac{N_L - 1}{n}.$$

By Corollary 11, to lower bound the distance of $\mathsf{TC}$, it suffices to upper bound the ratio $\frac{N_L-1}{n} = \frac{N_L-1}{N_F-1}$, or, essentially equivalently, the more natural ratio $\frac{N_L}{N_F}$. Note that the trivial upper bound $N_L \leq \ell n + 1$ yields only the trivial lower bound $\delta \geq 0$. Obtaining a nontrivial upper bound on $N_L$ is the content of Section 5.

## 5 Our Bound on the Number of Rational Points

In this section, we prove the following fairly general - though somewhat technical to state - proposition, which serves as the foundation for our results on elementary abelian $p$-extensions and Kummer extensions. The statement of Proposition 12 is formulated in a broad setting and thus involves several unspecified parameters. In Corollary 13, we present new normalized parameters, and rewrite the bound with these notations. In Proposition 14 we instantiate these parameters, and obtain the main general result, although still quite technical. In Section 5.3 we show that a quite general family of extensions satisfy the conditions of Proposition 14, and conclude Theorem 17, which will be used in all our applications.

From here onwards, $p$ is a prime number and $q = p^u$ for some integer $u \geq 1$. Moreover, $F/\mathbb{F}_q$ is a function field with $N$ rational places and genus $g$. We focus on curves in the regime

$$g = \Omega\left(\frac{N}{\sqrt{q}}\right). \tag{5.1}$$

For the complementary regime, one can apply the Grothendieck's trace formula to bound the number of rational places over extensions of $F$; see the discussion in Section 1.4. Let $\mathfrak{p}$ be a rational place of $F$, and let $x \in \mathcal{L}(s\mathfrak{p}) \setminus \mathcal{L}((s-1)\mathfrak{p})$ for some integer $s > 1$. Hence $\deg(x)_\infty = s$. We consider field extensions of the form $L/F$ where $L = F(z)$. We denote by

$$t \triangleq \deg(z)_\infty.$$

The reader may think of $t$ as given as input, and the resulted bound depends on $t$. The parameter $s$ on the other hand is "internal" to the proof and it will be beneficiary to choose it as small as possible.[7]

We will need to introduce a few more parameters. Let $1 \leq m \leq q$ be a power of $p$. Denote

$$A \triangleq \left\lfloor \frac{Nm}{2q} \right\rfloor$$

and let $B \in \mathbb{N}$ be a parameter.

▶ **Proposition 12.** *Let $L/F$ be a function field extension of the form $L = F(z)$, and let $d \triangleq [L : F]$. Assume that $\mathfrak{p}$ is totally ramified in $L$, and denote by $\mathfrak{P}$ the unique place of $L$ lying above $\mathfrak{p}$. Let $\{b_i\}_{i\in\mathcal{I}}$ be a basis of $\mathcal{L}(A\mathfrak{P})$ such that $v_{\mathfrak{P}}(b_i) = -i$. Assume that the elements*

$$S = \left\{ b_i x^{jm} z^{km} : i \leq A,\, j \leq B,\, k < d \right\} \tag{5.2}$$

*are linearly independent over $\mathbb{F}_q$. Further assume that*

$$(A - g_L)Bd > \frac{N}{2} + dsB + (d-1)t, \tag{5.3}$$

*where $g_L$ denotes the genus of $L$. Then, the number of rational points $N_L$ of $L$ satisfies*

$$N_L \leq \left( sdB + \frac{N}{2q} + td \right) m.$$

**Proof of Proposition 12.** First, notice that as $\mathfrak{p}$ is totally ramified, $\mathbb{F}_q$ is the full constant field of $L$. Let $U$ be the $\mathbb{F}_q$ vector space spanned by $S$ as given in Equation (5.2). Consider the $\mathbb{F}_p$ linear map $\Psi : U \to L$ defined by

$$\Psi\left( \sum_{\substack{i\leq A \\ j\leq B, k<d}} c_{i,j,k} b_i x^{jm} z^{km} \right) = \sum_{i,j,k} c_{i,j,k}^{q/m} b_i^{q/m} x^j z^k,$$

where $c_{i,j,k} \in \mathbb{F}_q$. Note that

$$\mathrm{Im}\Psi \subseteq \mathcal{L}\left( (N/2 + dsB)\,\mathfrak{P} + (d-1)(z)_\infty \right),$$

where we used that $v_{\mathfrak{P}}(x) = e(\mathfrak{P}|\mathfrak{p}) v_{\mathfrak{p}}(x) = ds$. This, together with Claim 3.2 in [5], shows that

$$\dim_{\mathbb{F}_p} \mathrm{Im}\Psi \leq (N/2 + dsB + (d-1)t + 1) \cdot [\mathbb{F}_q : \mathbb{F}_p].$$

Per our assumption on $S$, as an $\mathbb{F}_q$-vector space,

$$\dim_{\mathbb{F}_p} U = \dim_{\mathbb{F}_q} \mathcal{L}(A\mathfrak{P}) \cdot (B+1)\, d \cdot [\mathbb{F}_q : \mathbb{F}_p] \geq (A - g_L)\,(B+1)\, d \cdot [\mathbb{F}_q : \mathbb{F}_p],$$

where the last inequality follows by Riemann's Theorem.

Therefore our assumption Equation (5.3) implies that $\dim_{\mathbb{F}_p} U > \dim_{\mathbb{F}_p} \mathrm{Im}\Psi$ and so there exists a nonzero element $h \in U$ such that $\Psi(h) = 0$. Indeed, this follows since $\Psi$ is additive. Note however that for every degree-1 place $\mathfrak{Q} \neq \mathfrak{P}$ of $L$,

$$0 = \Psi(h)^m(\mathfrak{Q}) = h(\mathfrak{Q}).$$

[7] Note that by Claim 3.1 in the full version [5], $s \geq \frac{N}{q+1}$.

That is, $h$ vanishes on all rational places $\mathfrak{Q} \neq \mathfrak{P}$ of $L$, and so $N_L \leq \deg(h)_\infty + 1$. But

$$h \in U \subseteq \mathcal{L}\left((A + sdBm)\,\mathfrak{P} + (d-1)m(z)_\infty\right),$$

hence $\deg(h)_\infty \leq A + sdBm + t(d-1)m$, which concludes the proof. ◀

We define the parameters $\tau, \beta, \gamma, \mu$ and $\sigma$ which satisfies the following relations with regards to the previously defined parameters in this section so far:

$$\begin{aligned} t &= \tau\frac{N}{\sqrt{q}} \\ B &= \beta\frac{\sqrt{q}}{d} \\ g &= \gamma\frac{N}{\sqrt{q}} \\ m &= \mu\sqrt{q} \\ s &= \sigma\frac{N}{q} \end{aligned} \tag{5.4}$$

Note that per Equation (5.1) we assume that $\gamma = \Omega(1)$.

With this we have the following corollary, retaining the notation from Proposition 12.

▶ **Corollary 13.** *Assume that*

$$g_L \leq d(g+t) \tag{5.5}$$

*and that*

$$\beta\mu \geq 1 + 2(\tau+\gamma)\beta d + \frac{2}{\sqrt{q}}\left(\tau(d-1) + \sigma\beta\right) \tag{5.6}$$

*hold. Assume further, as in Proposition 12, that the elements in $S$ are linearly independent over $\mathbb{F}_q$. Then,*

$$N_L \leq \mu N\left(\sigma\beta + \tau d + \frac{1}{2\sqrt{q}}\right). \tag{5.7}$$

**Proof.** We refer the reader to the full version of the paper [5]. ◀

## 5.1 Before the Instantiation: A Parameter Walkthrough

Before proceeding to instantiate the parameters $\beta, \mu$ in Corollary 13, we provide the reader with some informal insight into the magnitudes of the various parameters. While this discussion is not rigorous, it is intended to offer intuition about the bounds one should expect.

The five parameters appearing in Equation (5.4) can be grouped into three distinct categories. As previously noted, the reader should view $\tau$ as the input parameter – the bound we derive for $N_L$ depends on $\tau$, and it is the nature of this dependence that we aim to understand. The parameters $\beta$ and $\mu$ are under our control; by selecting them appropriately, we can optimize the resulting bound. In contrast, the parameters $\gamma$ and $\sigma$ are intrinsic to the function field under consideration, and for the purposes of this discussion, the reader may assume $\gamma = \sigma = 1$.

As a starting point, we ignore terms that vanish as $q \to \infty$. This simplification allows us to gain an initial understanding of how the bound depends on the dominant quantities, namely, $p$, $d$ and $\tau$. It also offers guidance on how to optimize the bound by appropriately choosing the parameters $\beta$ and $\mu$. The assumption in Equation (5.6) then simplifies to

$$\beta\mu \geq 1 + 2(\tau+1)\beta d.$$

We take this to hold with equality, yielding the following relation between $\beta$ and $\mu$:

$$\mu = \frac{1}{\beta} + 2(\tau+1)d. \tag{5.8}$$

Now, the bound guaranteed by Corollary 13 simplifies, as discussed above, to

$$\frac{N_L}{N} \leq \mu\left(\beta + \tau d\right).$$

Plugging in the relation between $\mu$ and $\beta$, we obtain

$$\frac{N_L}{N} \leq \left(\frac{1}{\beta} + 2(1+\tau)d\right)\left(\beta + \tau d\right).$$

Generally, for $a, b > 0$, the minimum of the function $f(x) = (x+a)\left(\frac{1}{x} + b\right)$ in $(0, \infty)$ is $(1+\sqrt{ab})^2$, attained at $x = \sqrt{\frac{a}{b}}$. Applying this, we get that setting $\beta = \sqrt{\frac{\tau}{2(1+\tau)}}$ yields the bound

$$\frac{N_L}{N} \leq \left(1 + \sqrt{2\tau(1+\tau)}\, d\right)^2.$$

Analogously to the Hasse–Weil theorem, where the number of rational points on a curve lying over the projective line is expressed as $q + 1 + \mathrm{Err}$, the presence of the constant term 1 is natural and expected. The remaining contribution is the "error term". Note that a bound of $d$ on the ratio is trivial, since $L/F$ is an extension of degree $d$. Thus, the bound is meaningful only when $\tau \ll 1$, in which case we can simplify the informal (and inaccurate) bound to

$$\frac{N_L}{N} \leq \left(1 + \sqrt{2\tau}\, d\right)^2 = 1 + 2\sqrt{2\tau}\, d + 2\tau d^2.$$

Note that the contributions of the two summands on the right-hand side to the error term are incomparable; which one dominates depends on whether $2\tau < \frac{1}{d^2}$ or not.

### The dependence on $q$

Of course, this bound is not accurate – in particular, it ignores the dependence on $q$, which is actually quite important in applications, as it determines how the alphabet reduction (i.e., the minimal value of $q$ for which we start) affects the bound. Moreover, the parameter $\beta$ should be chosen such that $B$ is an integer, and $\mu$ should be chosen such that $m$ is a power of $p$, which we will take care of later.

However, the above discussion provides useful insight into how to choose $\beta$ and $\mu$. In particular, we set $\beta = \sqrt{\frac{\tau}{2(1+\tau)}}$, which is well-approximated by $\sqrt{\tau/2}$ under the assumption that $\tau \ll 1$. Thus, we proceed with the choice $\beta = \sqrt{\tau/2}$. Substituting this into Equation (5.8), and using again that $\tau \ll 1$, a suitable choice for $\mu$ is

$$\mu = \frac{1}{\sqrt{\tau/2}} + 2d. \tag{5.9}$$

Plugging this into Equation (5.7), we obtain

$$\frac{N_L}{N} \leq \left(1+\sqrt{2\tau}d\right)^2 + \frac{1}{\sqrt{q}}\left(d+\frac{1}{\sqrt{2\tau}}\right). \tag{5.10}$$

Notice that by Claim 3.1 in [5], we have $\tau = t\frac{\sqrt{q}}{N} \geq \frac{1}{\sqrt{q}} \cdot \frac{q}{q+1}$. Thus,

$$\frac{d}{\sqrt{q}} \leq \tau d \cdot \left(1+q^{-1}\right),$$
$$\frac{1}{\sqrt{2q\tau}} \leq \sqrt{\tau}.$$

Hence, the bound in Equation (5.10) can be rewritten as

$$\frac{N_L}{N} \leq 1 + O(\sqrt{\tau}d + \tau d^2). \tag{5.11}$$

## 5.2 Instantiation of Parameters

In this subsection, we formalize the above discussion and establish the following general result.

▶ **Proposition 14.** *Let $L/F$ be a function field extension of the form $L = F(z)$, and let $d \triangleq [L:F]$. Assume that $\mathfrak{p}$ is totally ramified in $L$, and denote by $\mathfrak{P}$ the unique place of $L$ lying above $\mathfrak{p}$. Let $\{b_i\}_{i\in\mathcal{I}}$ be a basis of $\mathcal{L}(A\mathfrak{P})$ such that $v_{\mathfrak{P}}(b_i) = -i$. Assume that the elements*

$$S = \left\{b_i x^{jm} z^{km} : i \leq A,\ j \leq B,\ k < d\right\} \tag{5.12}$$

*are linearly independent over $\mathbb{F}_q$. Further, assume that $g_L \leq d(g+t)$. Then,*

$$\frac{N_L}{N} \leq \sigma + O_\sigma\left(\sqrt{\tau}d\gamma p + \tau d^2\gamma p + \frac{d}{\sqrt{q}}\left(\frac{1}{\sqrt{\tau}}+d\right)\gamma p\right). \tag{5.13}$$

▶ Remark 15. If $\sigma \sim 1, \gamma \sim 1$ as in Section 5.1, this is the same error term as in Equation (5.10) up to a factor of $p$.

**Proof.** We refer the reader to the full version of the paper [5]. ◀

## 5.3 Extensions that Satisfy the Conditions of Proposition 14

▶ **Proposition 16.** *Let $\mathfrak{p} \in \mathbb{P}_F$ be a rational place, and let $f \in F$ have poles only at $\mathfrak{p}$. Let $L/F$ be a function field extension defined by*

$$L = F(z) \quad \text{such that} \quad \phi(z) = f,$$

*where $\phi(T) \in \mathbb{F}_q[T]$ is a polynomial of degree $d$. Assume that the polynomial $\phi(T) - f \in F[T]$ is irreducible. Assume that $\mathfrak{p}$ is totally ramified in $L$, and denote by $\mathfrak{P}$ the unique place of $L$ lying above $\mathfrak{p}$. Let $\{b_i\}_{i\in\mathcal{I}}$ be a basis of $\mathcal{L}(A\mathfrak{P})$ such that $v_{\mathfrak{P}}(b_i) = -i$. Then,*

1. $t \triangleq \deg_L(z)_\infty = -v_{\mathfrak{P}}(z) = \deg_F(f)_\infty$.

2. *If we can write* $\deg_F(f)_\infty = t = \ell s + \varepsilon \frac{N}{q}$ *for some* $\ell$ *relatively prime to* $d$*, and* $|\varepsilon| < \frac{\sigma - 1/2}{d-1}$*, then the set*

$$S = \left\{ b_i x^{jm} z^{km} : i \le A,\ j \le B,\ k < d \right\} \tag{5.14}$$

*is linearly independent over* $\mathbb{F}_q$*.*

**Proof.** We refer the reader to the full version of the paper [5]. ◀

As a corollary, we state the main theorem which will be used in all our applications.

▶ **Theorem 17.** *Let* $\mathfrak{p} \in \mathbb{P}_F$ *be a rational place, and let* $f \in F$ *have poles only at* $\mathfrak{p}$*. Assume that* $\deg_F(f)_\infty$ *is relatively prime to* $d$*, and that we can write* $\deg_F(f)_\infty = \ell s + \varepsilon \frac{N}{q}$ *for some* $\ell$ *relatively prime to* $d$*, and* $0 \le \varepsilon < \frac{\sigma - 1/2}{d-1}$*. Let* $L/F$ *be a function field extension defined by*

$$L = F(z) \quad \text{such that} \quad \phi(z) = f,$$

*where* $\phi(T) \in \mathbb{F}_q[T]$ *is a polynomial of degree* $d$*. Assume that the polynomial* $\phi(T) - f \in F[T]$ *is irreducible. Assume that* $\mathfrak{p}$ *is totally ramified in* $L$*. Further, assume that* $g_L \le d(g + t)$*. Then* $\deg_F(f)_\infty = t$*, and*

$$\frac{N_L}{N} \le \sigma + O_\sigma\left(\sqrt{\tau} d \gamma p + \tau d^2 \gamma p\right). \tag{5.15}$$

**Proof.** We refer the reader to the full version of the paper [5]. ◀

# 6 Abelian Extensions and Character Sums

In this section, we instantiate the bound from Theorem 17 for two families of abelian extensions. The first family consists of elementary abelian $p$-extensions of the form $L = F(z)$, where $z^\ell - z = f$ and $d$ is a prime power. We apply the bound to this setting to derive a lower bound on the minimum distance of TAG codes, as discussed in Section 4. Moreover, in the Artin–Schreier case (i.e., when $\ell = p$ is prime), the bound yields an upper bound on exponential sums of functions over curves. The second family we consider is Kummer extensions, namely those of the form $L = F(z)$ with $z^d = f$, where $d$ is not divisible by the characteristic $p$ of $F$.

## 6.1 Elementary Abelian $p$-Extensions and Exponential Sums

In this subsection we consider elementary abelian $p$-extensions of the form $L = F(z)$, where $z^\ell - z = f$ and $\ell$ is a prime power. These extensions satisfy the conditions of Theorem 17. As discussed in Section 4, any non-trivial bound for the number of rational points on this type of extensions immediately yields a lower bound on the minimum distance of the corresponding TAG code. In Proposition 18, we use this connection to obtain a lower bound for the distance of TAG codes; in Corollary 20 we give an argument that also yields an *upper* bound. We use this in Theorem 21 to get an upper bound for exponential sums of function on curves.

Recall the setting of Section 5. Let $\ell > 1$ be a power of $p$. Let $\phi(T) = T^\ell - T$, let $f \in \mathcal{L}(r\mathfrak{p})$ and assume that $t = \deg(f)_\infty = -v_\mathfrak{p}(f)$ is not divisible by $p$. Let $L = F(z)$ with $\phi(z) = f$. By [21, Proposition 3.7.10(a)], $L/F$ is an elementary abelian extension of exponent $p$. By part (d) of this proposition, the prime $\mathfrak{p}$ is totally ramified in $L$, and moreover by part (e) we have

$$g_L = \ell g + \frac{\ell - 1}{2}(-2 + (t+1)) \le \ell(g + t).$$

Thus we are in a position to apply Theorem 17 to $L/F$.

▶ **Proposition 18.** *In the above setting, assume further that $f$ satisfies the condition in Item 2 of Proposition 16. Then,*

$$\frac{N_L}{N} \leq \sigma + O_\sigma\left(\ell\sqrt{\tau}\gamma p + \ell^2\tau\gamma p\right). \tag{6.1}$$

Let $\mathsf{Tr}$ be the trace function from $\mathbb{F}_q$ to $\mathbb{F}_\ell$. Claim 10 relates splitting places in $L/F$ to the vanishing of the trace of $f$ at these places. We use this relation together with Proposition 18 to get an upper bound for the number of vanishing traces of evaluations of $f$.

▶ **Corollary 19.** *For every function $f \in \mathcal{L}(r\mathfrak{p})$ such that $t = \deg(f)_\infty$ is not divisible by $p$, we have*

$$\left|\left\{\mathfrak{q} \in \mathbb{P}_F^1 \setminus \{\mathfrak{p}\} : \mathsf{Tr}(f(\mathfrak{q})) = 0\right\}\right| = \frac{N_L - 1}{\ell}.$$

*If moreover $f$ satisfies the condition in Item 2 of Proposition 16, then*

$$\frac{1}{N-1}\left|\left\{\mathfrak{q} \in \mathbb{P}_F^1 \setminus \{\mathfrak{p}\} : \mathsf{Tr}(f(\mathfrak{q})) = 0\right\}\right| \leq \frac{\sigma}{\ell} + O_\sigma\left(p\sqrt{\tau}\gamma + p\ell\tau\gamma\right). \tag{6.2}$$

As we will see in Section 7, for all function fields in which we instantiate our results, it will be possible to choose $x \in F$ such that $\sigma \leq 1$. Moreover, Proposition 16 provides a condition for the set $S = S(f)$ to be linearly independent over $\mathbb{F}_q$, depending only on $t$. Consequently, if the result applies to $f$, it also applies to $f - \alpha$ for all $\alpha \in \mathbb{F}_q$.
In this case, we can also obtain a lower bound.

▶ **Corollary 20.** *Assume that $\sigma \leq 1$. For every function $f \in \mathcal{L}(r\mathfrak{p})$ of degree not divisible by $p$, and that satisfies the condition in Item 2 of Proposition 16, we have*

$$\frac{1}{N-1}\left|\left\{\mathfrak{q} \in \mathbb{P}_F^1 \setminus \{\mathfrak{p}\} : \mathsf{Tr}(f(\mathfrak{q})) = \beta\right\}\right| = \frac{1}{\ell} + O\left(\ell\sqrt{\tau}\gamma p\right).$$

**Proof.** We refer the reader to the full version of the paper [5]. ◀

In Section 7 we use these results to construct and analyze TAG codes on some curves.

Lastly, from the special case of Artin-Schreier extensions, i.e. when $\ell = p$, we conclude an upper bound for the exponential sums arising from these functions.

▶ **Theorem 21.** *Let $\ell = p$. Assume that $\sigma \leq 1$. For every function $f \in \mathcal{L}(r\mathfrak{p})$ of degree not divisible by $p$, and that satisfies the condition in Item 2 of Proposition 16, we have*

$$\frac{1}{N}\left|\sum_{\mathfrak{q} \in \mathbb{P}_F^1 \setminus \{\mathfrak{p}\}} e^{\frac{2\pi i}{p}\mathsf{Tr}(f(\mathfrak{q}))}\right| = O\left(\sqrt{\tau}\gamma p^3\right).$$

## 6.2 Kummer Extensions and Multiplicative Character Sums

In this section, following a similar outline to Section 6.1, we instantiate the bound from Theorem 17 in the setting of Kummer extensions to get an upper bound on the number of rational places in such extensions. Similarly to Corollary 20, in Proposition 23 we give an argument that also yields a *lower* bound. We then apply this bound to derive upper bounds for multiplicative character sums of functions on curves, in direct analogy with the upper bounds for exponential sums obtained for the Artin–Schreier case.

Recall the setting of Section 5. Let $\phi(T) = T^d$ for $d > 1$ such that $d \mid q-1$. Let $f \in \mathcal{L}(r\mathfrak{p})$ and assume that $t = \deg(f)_\infty = -v_\mathfrak{p}(f)$ is relatively prime to $d$. Let $L = F(z)$ with $\phi(z) = f$. Since $d \mid q-1$, the field $\mathbb{F}_q$ contains a primitive $d$-th root of unity. By Corollary 3.7.4

in [21], $L/F$ is a cyclic Kummer extension of degree $d$, and hence $\phi(T) - f$ is irreducible. By Proposition 3.7.3(b) in [21], the place $\mathfrak{p}$ is totally ramified in $L$, and moreover by Corollary 3.7.4 in [21] we have

$$g_L = 1 + d(g-1) + \frac{1}{2}\sum_{\mathfrak{q}\in\mathbb{P}_F} (d - (d, v_{\mathfrak{q}}(f))) \deg \mathfrak{q}.$$

Note that for $\mathfrak{q} \in \mathbb{P}_F$, if $v_{\mathfrak{q}}(f) = 0$ then $d - (d, v_{\mathfrak{q}}(f)) = 0$. Otherwise, $d - (d, v_{\mathfrak{q}}(f)) \leq d$. Therefore,

$$\begin{aligned} g_L &\leq 1 + d(g-1) + \frac{1}{2} \cdot \sum_{\mathfrak{q}\in\mathbb{P}_F \,:\, v_{\mathfrak{q}}(f)\neq 0} d \deg \mathfrak{q} \\ &\leq dg + \frac{d}{2}\left(\deg(f)_0 + \deg(f)_\infty\right) \leq d(g+t) \end{aligned}$$

as $\deg(f)_0 = \deg(f)_\infty = t$. Thus we are in a position to apply Theorem 17 to $L/F$.

▶ **Proposition 22.** *In the above setting, assume further that* $\deg_F(f)_\infty$ *is relatively prime to* $d$*, and satisfies the condition in Item 2 of Proposition 16. Then,*

$$\frac{N_L}{N} \leq \sigma + O_\sigma\left(\sqrt{\tau} d\gamma p + \tau d^2 \gamma p\right). \tag{6.3}$$

Similarly to Section 6.1, we can also achieve a lower bound under further assumptions.

▶ **Proposition 23.** *Assume that* $\sigma \leq 1$*. For every function* $f \in \mathcal{L}(r\mathfrak{p})$ *such that* $\deg_F(f)_\infty$ *is relatively prime to* $d$*, and satisfies the condition in Item 2 of Proposition 16, we have*

$$\frac{N_L}{N} = 1 + O\left(\sqrt{\tau} d^2 \gamma p\right).$$

*Furthermore, we have*

$$\frac{1}{N}\left|\left\{\mathfrak{q} \in \mathbb{P}_F^1 \setminus \{\mathfrak{p}\} : \exists\, y \in \mathbb{F}_q^\times,\ y^d = f(\mathfrak{q})\right\}\right| = \frac{1}{d} + O\left(\sqrt{\tau} d\gamma p\right).$$

**Proof.** We refer the reader to the full version of the paper [5]. ◀

Lastly, we conclude the bound for multiplicative character sums arising from these functions.

▶ **Theorem 24.** *For all multiplicative characters* $\chi \in \widehat{\mathbb{F}_q^\times}$ *of order* $d$*, and for all functions* $f \in F$ *which satisfy the conditions of Proposition 23, we have*

$$\frac{1}{N}\left|\sum_{\mathfrak{q}\in\mathbb{P}_F^1\setminus\{\mathfrak{p}\}} \chi(f(\mathfrak{q}))\right| = O\left(\sqrt{\tau} d^2 \gamma p\right).$$

## 7 TAG Codes Instantiations

In this section, we use the results obtained in the previous sections to construct and analyze TAG codes of specific function fields $F$. As discussed in Section 2.1.3, the functions captured by our result are roughly those of degree $t < \frac{g}{p^2}$. Therefore, in order to construct TAG codes with high rate, it is necessary to consider curves that admit a place for which the associated Riemann-Roch spaces of degree $< \frac{g}{p^2}$ have sufficiently large dimensions. In the full version of the paper [5] we instantiate our results for curves satisfying these conditions, such as the Hermitian curve, Norm-Trace curves and the Hermitian tower.

## 8 TAG Codes vs. Concatenation in the High Distance Regime

In this section we compare binary TAG codes with the familiar approach of concatenating with the Hadamard code, thereby proving Theorem 3. We consider here the case where the error term guaranteed in Corollary 19 is $O(\tau)$, which would lead to the $\Omega(\varepsilon^3)$ lower bound on the rate. The "Moreover" part of Theorem 3, corresponds to $O(\sqrt{\tau})$, is similar.

Let $C$ be a one-point evaluation AG code, defined by an algebraic function field $F/\mathbb{F}_q$ with genus $g$, and a set of $N$ rational places $Y$ with respect to a divisor $G = T\mathfrak{P}$ such that $\mathfrak{P} \notin Y$, $T < g$. Assume that $F$ contains an element $x \in F$ with $\frac{q}{N}\deg(x)_\infty = 1 + O(q^{-1/2})$, and $N/g = \Omega(\sqrt{q})$. Let $\ell = \ell(T\mathfrak{P})$, the dimension of the corresponding Riemann-Roch space. This AG code has parameters $n_1 = N$, designated distance $d_1 = N - T$, and dimension $\ell$, and it is defined over $\mathbb{F}_q$. We compare the parameters of the two methods described above:

1. **TAG codes.** Without loss of generality, assume that every function $f \in \mathcal{L}(T\mathfrak{P})$ satisfies the conditions of Corollary 19. The error term guaranteed in Corollary 19 is $O(\tau) = O\left(\frac{t\sqrt{q}}{N}\right)$. Then the resulting TAG code has parameters
$$\begin{aligned} n_T &= N, \\ k_T &= \ell \log(q), \\ \varepsilon_T &= \Omega\left(\frac{T\sqrt{q}}{N}\right). \end{aligned}$$
2. **Concatenation with Hadamard.** The resulting code has parameters
$$\begin{aligned} n_H &= Nq, \\ k_H &= \ell \log(q), \\ \varepsilon_H &\leq \frac{T}{N}. \end{aligned}$$

We let $k = k_H = k_T$. Ta-Shma and Ben-Aroya [2] proved that in the regime of $g = \Omega(\sqrt{q})$ we have

$$n_H = \Omega\left(\frac{k}{\varepsilon_H^{2.5}\log(k/\varepsilon_H)}\right). \tag{8.1}$$

Assume that $\alpha, \beta > 0$ are such that

$$n_H = \Omega\left(\frac{k^\alpha}{\varepsilon_H^\beta}\right), \tag{8.2}$$

and note that per Equation (8.1), $\beta \geq 2.5$. Observe that $\varepsilon_T = \Omega(\varepsilon_H\sqrt{q})$, $n_T = n_H/q$. Hence,

$$n_T = \frac{n_H}{q} = \Omega\left(\frac{k^\alpha \varepsilon_H^{-\beta}}{q}\right) = \Omega\left(k^\alpha \varepsilon_T^{-\beta} q^{\beta/2-1}\right).$$

Notice that by Claim 3.1 in [5], in order to have $\mathcal{L}(T\mathfrak{P}) \neq \mathbb{F}_q$, we must have $T \geq \frac{N}{q+1}$, hence as $\varepsilon_T = \Omega(\frac{T\sqrt{q}}{N})$ we have $\sqrt{q} = \Omega(\frac{1}{\varepsilon_T})$. Finally, we obtain

$$n_T = \Omega\left(k^\alpha \varepsilon_T^{-\beta-(\beta-2)}\right) = \Omega\left(\frac{k^\alpha}{\varepsilon_T^{2\beta-2}}\right).$$

Comparing this with Equation (8.2), recalling that $\beta \geq 2.5$, wee see that the exponent in the dependence in $\varepsilon$ in the TAG code, $2\beta - 2$, is larger by at least 0.5 from the corresponding exponent for the concaentated code. Indeed, the difference is $(2\beta - 2) - \beta = \beta - 2 \geq 0.5$.

## References


1 Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost $k$-wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992. doi:10.1002/RSA.3240030308.

2 Avraham Ben-Aroya and Amnon Ta-Shma. Constructing small-bias sets from Algebraic-Geometric codes. *Theory of Computing*, 9(5):253–272, 2013. doi:10.4086/TOC.2013.V009A005.

3 Enrico Bombieri. On exponential sums in finite fields. *American Journal of Mathematics*, 88(1):71–105, 1966.

4 Gil Cohen, Itay Cohen, and Gal Maor. Tight bounds for the zig-zag product. In *Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1470–1499. IEEE, 2024. doi:10.1109/FOCS61266.2024.00094.

5 Gil Cohen, Dean Doron, Noam Goldgraber, and Tomer Manket. Tracing AG codes: Toward meeting the Gilbert-Varshamov bound. *arXiv*, 2025. (manuscript). doi:10.48550/arXiv.2511.08788.

6 P. Delsarte. On subfield subcodes of modified Reed–Solomon codes. *IEEE Transactions on Information Theory*, 21(5):575–576, 1975.

7 Dean Doron, Jonathan Mosheiff, and Mary Wootters. When do low-rate concatenated codes approach the gilbert-varshamov bound? In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, volume 317, pages 53:1–53:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.APPROX/RANDOM.2024.53.

8 Arnaldo García and Henning Stichtenoth. A tower of Artin–Schreier extensions of function fields attaining the Drinfeld–Vlăduţ bound. *Inventiones mathematicae*, 121:211–222, 1995.

9 Arnaldo García and Henning Stichtenoth. On the asymptotic behaviour of some towers of function fields over finite fields. *Journal of Number Theory*, 61(2):248–273, 1996.

10 Edgar N. Gilbert. A comparison of signalling alphabets. *Bell System Technical Journal*, 31(3):504–522, 1952.

11 V. D. Goppa. Codes on algebraic curves. *Doklady Akademii Nauk SSSR*, 259(6):1289–1290, 1981.

12 Alexander Grothendieck. Formule de lefschetz. In *Cohomologie $\ell$-adique et Fonctions L (SGA 5)*, volume 589 of *Lecture Notes in Mathematics*, pages 73–137. Springer, 1977.

13 J. Justesen. A class of constructive asymptotically good algebraic codes. *IEEE Transactions on Information Theory*, 18(5):652–656, 1972. doi:10.1109/TIT.1972.1054893.

14 Swastik Kopparty, Amnon Ta-Shma, and Kedem Yakirevitch. Character sums over AG codes. *ECCC*, 2024. (manuscript). URL: https://eccc.weizmann.ac.il/report/2024/069/.

15 Swastik Kopparty, Amnon Ta-Shma, and Kedem Yakirevitch. Trace Hermitian codes have vanishing bias, 2025. Personal communication.

16 Phong Le and Sunil Chetty. On the dimension of AG trace codes. *arXiv*, 2016. (manuscript). arXiv:0902.1729.

17 Robert McEliece, Eugene Rodemich, Howard Rumsey, and Lloyd Welch. New upper bounds on the rate of a code via the Delsarte–MacWilliams inequalities. *IEEE transactions on Information Theory*, 23(2):157–166, 1977. doi:10.1109/TIT.1977.1055688.

18 Joseph Naor and Moni Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993. doi:10.1137/0222053.

19 Michael Sipser and Daniel A. Spielman. Expander codes. *IEEE Transactions on Information Theory*, 42(6):1710–1722, 2002.

20 Alexei N. Skorobogatov. The parameters of subcodes of algebraic-geometric codes over prime subfields. *Discrete Applied Mathematics*, 33(1-3):205–214, 1991. doi:10.1016/0166-218X(91)90116-E.

21 Henning Stichtenoth. *Algebraic function fields and codes*, volume 254 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, second edition, 2009.

22 Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In *Annual Symposium on Theory of Computing (STOC)*, pages 238–251. ACM, 2017. doi:10.1145/3055399.3055408.

23 Amnon Ta-Shma. The hermitian trace code. https://simons.berkeley.edu/talks/amnon-ta-shma-tel-aviv-university-2024-04-08, 2024. Invited talk.

24 R. Tanner. A recursive approach to low complexity codes. *IEEE Transactions on Information Theory*, 27:533–547, 1982.

25 M. A. Tsfasman, S. G. Vlăduţ, and Th. Zink. Modular curves, Shimura curves, and Goppa codes, better than the Varshamov–Gilbert bound. *Mathematische Nachrichten*, 109(1):21–28, 1982.

26 Marcel van der Vlugt. On the dimension of trace codes. *IEEE Transactions on Information Theory*, 37(1):196–199, 1991. doi:10.1109/18.61140.

27 R. R. Varshamov. Estimate of the number of signals in error-correcting codes. *Doklady Akademii Nauk SSSR*, 117:739–741, 1957. English translation reprinted in I. F. Blake (ed.), *Algebraic Coding Theory: History and Development*, Dowden, Hutchinson & Ross, 1973, pp. 68–71.

28 S. G. Vladut. Two applications of Weil-Serre's explicit formula. *Applicable Algebra in Engineering, Communication and Computing*, 7(4):279–288, 1996. doi:10.1007/BF01195533.

29 S. G. Vlăduţ and V. G. Drinfel'd. Number of points of an algebraic curve. *Functional Analysis and Its Applications*, 17(1):53–54, 1983.

30 Shudi Yang and Chuangqiang Hu. Weierstrass semigroups from a tower of function fields attaining the Drinfeld-Vladut bound. *arXiv*, 2019. (manuscript). arXiv:1911.04269.