THEORY OF COMPUTING, Volume 9 (5), 2013, pp. 253–272
www.theoryofcomputing.org

Constructing Small-Bias Sets from Algebraic-Geometric Codes*

Avraham Ben-Aroya† Amnon Ta-Shma‡

Received February 21, 2011; Revised October 29, 2012; Published February 20, 2013

Abstract: We give an explicit construction of an ε-biased set over k bits of size

$$O\left(\frac{k}{\varepsilon^2 \log(1/\varepsilon)}\right)^{5/4}.$$

This improves upon previous explicit constructions when ε is roughly (ignoring logarithmic factors) in the range $[k^{-1.5}, k^{-0.5}]$. The construction builds on an algebraic-geometric code. However, unlike previous constructions, we use low-degree divisors whose degree is significantly smaller than the genus.

ACM Classification: F.2.2, G.2

AMS Classification: 94B27, 12Y05

Key words and phrases: small-bias sets, algebraic geometry, AG codes, Goppa codes

1 Introduction

Explicitly constructing combinatorial objects with certain properties (such as expander graphs, extractors, error correcting codes and others) is an intriguing challenge in computer science. Often, it is easy to verify that a random object satisfies the required property with high probability, while it is difficult to pin down such an explicit object.

*A preliminary version of this paper appeared in FOCS 2009 [3].

†Supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities, and by the European Commission under the Integrated Project QAP funded by the IST directorate as Contract Number 015848.

‡Supported by Israel Science Foundation grant 217/05 and by USA Israel BSF grant 2004390.

DOI: 10.4086/toc.2013.v009a005

In most cases it is believed (and sometimes proven) that a random object is nearly optimal. Therefore, giving an optimal explicit construction becomes a derandomization problem. There are, however, rare cases in which explicit constructions outperform naive random constructions. Perhaps the most remarkable example of this type is that of algebraic-geometric codes (AG codes). In the seminal work of Tsfasman et al. [10] it was shown that there are algebraic-geometric codes over constant size alphabets that lie above the Gilbert-Varshamov bound, a bound that random codes achieve and that was believed to be optimal at the time.

The important case of *binary* error correcting codes is still open. In the binary case, the Gilbert-Varshamov bound gives the best known (explicit or non-explicit) codes to date. For codes with distance close to half, the bound shows that random linear codes of length $n = O(k/\varepsilon^2)$ are $[n,k,(1/2-\varepsilon)n]_2$ codes. Finding an explicit construction that attains this bound is an open problem.

Another closely related question is that of finding an $[n,k,(1/2-\varepsilon)n]_2$ binary code, in which the relative weight of every non-zero codeword is in the range $[1/2-\varepsilon, 1/2+\varepsilon]$. Such codes are called *ε-balanced* and they are related to another kind of combinatorial objects called *ε-biased sets*. An ε-biased set is a set $S \subseteq \{0,1\}^k$ such that for every non-empty subset $T \subseteq [k]$, the binary random variable $\bigoplus_{i\in T} s_i$, where s is sampled uniformly from S, has bias at most ε. It turns out that ε-biased sets are just ε-balanced codes in a different guise: the rows of a matrix whose columns generate an ε-balanced code form an ε-biased set, and vice versa. In terms of parameters, an $[n,k]_2$ ε-balanced code is equivalent to an ε-biased set $S \subseteq \{0,1\}^k$ of size n.

The status of ε-balanced codes is similar to that of $[n,k,(1/2-\varepsilon)n]_2$ codes. In both cases the probabilistic method gives non-explicit $[n,k]_2$ ε-balanced codes with $n = O(k/\varepsilon^2)$, whereas the best lower bound is

$$n = \Omega\left(\frac{k}{\varepsilon^2 \log(\frac{1}{\varepsilon})}\right).$$

For a discussion of these bounds see [2, Section 7].

There are several *explicit* constructions of such codes. Naor and Naor [7] give a construction with $n = k \cdot \mathrm{poly}(\varepsilon^{-1})$, which was later improved in [1] to $n = O(k/\varepsilon^3)$. Alon et al. [2] establish the incomparable bound

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

Concatenating algebraic-geometric codes with the Hadamard code gives

$$n = O\left(\frac{k}{\varepsilon^3 \log(\frac{1}{\varepsilon})}\right).$$

In this paper we show an explicit construction of an $[n,k]_2$ ε-balanced code with

$$n = O\left(\frac{k}{\varepsilon^2 \log(\frac{1}{\varepsilon})}\right)^{5/4},$$

which improves upon previous explicit constructions when ε is roughly (ignoring logarithmic factors) in the range of $k^{-1.5} \le \varepsilon \le k^{-0.5}$ (see Figure 1).

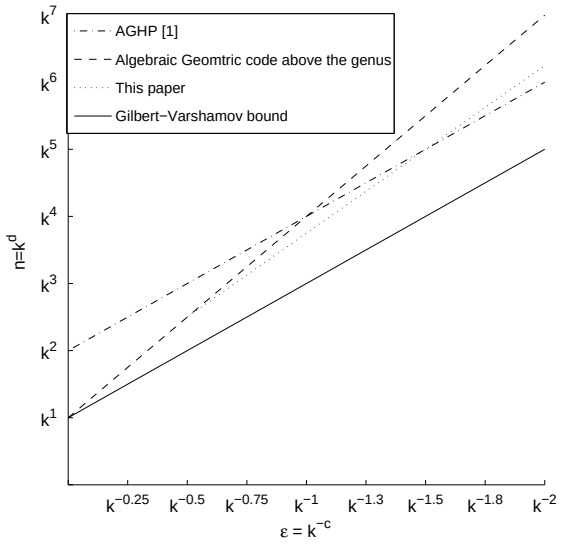

Figure 1: Constructions of ε-biased sets for $\varepsilon = k^{-c}$.

The construction is simple and can be described by elementary means. We first take a finite field $\mathbb{F}_q$ of the appropriate size. We then carefully choose a subset A of $\mathbb{F}_q \times \mathbb{F}_q$. The elements in the ε-biased set are indexed by pairs $((a,b),c) \in A \times \mathbb{F}_q$. For each $((a,b),c) \in A \times \mathbb{F}_q$ the corresponding element is the bit vector $\left(\langle (a^i b^j), c\rangle_2\right)_{i,j}$, where (i,j) range over all integers i,j whose sum is bounded by an appropriately chosen parameter and the inner product is of the binary representation of the elements in $\mathbb{F}_q$. The analysis of the construction relies on Bézout's Theorem.

To put the construction in context, we need to move to the terminology of algebraic function fields. AG codes are *evaluation* codes where a certain set of *evaluation functions* is evaluated at a chosen set of *evaluation points*. The space of evaluation functions used is a vector space (this is the reason we get a linear error correcting code) and is determined by a *divisor* G of an algebraic function field F. We explain these notions in Section 3, and for the time being continue with an intuitive discussion. We denote the code associated with a divisor G by $C(G)$.

The code $C(G)$ has the following parameters.

- The *length* of the code is the number of evaluation points and is denoted by $N = N(F)$.
- The *distance* of the code is $N - \deg(G)$, where $\deg(G)$ is the degree of G (formally defined in Section 3).
- The *dimension* of the code, $\dim(G)$, is the dimension of the vector space of evaluation functions.

When the "degree" of G is larger than the *genus* g of the function field F (again, defined in Section 3) the Riemann-Roch Theorem [8, Thm I.5.17] tells us exactly what the dimension $\dim(G)$ is, and it turns out that

$$\dim(G) = \deg(G) - g + 1\,.$$

This almost matches the Singleton bound, $\dim(G) \leq \deg(G) + 1$, except for a loss of g. Thus, our goal is to get as many evaluation points as possible while keeping the genus small. Indeed, a lot of research has been done on the best possible ratio between the length of the code $N(F)$ and the genus g. The bottom line of this research, roughly speaking, is that $N(F)$ can be larger than the genus by at most a multiplicative $\sqrt{q} - 1$ factor and this is essentially optimal.

A simple check shows that when $\deg(G)$ is larger than the genus, an AG code concatenated with the Hadamard code cannot give ε-balanced codes with n better than

$$O\left(\frac{k}{\varepsilon^3 \log(\frac{1}{\varepsilon})}\right)$$

(see Section 3.2). In contrast, our construction takes as an outer code an AG code $C(G)$ where $\deg(G)$ is much smaller than the genus, and we show that this leads to a better code.

A natural question is whether the ε-balanced codes we achieve are the best binary codes one can achieve using this approach. We do not know the answer to this question. When $\deg(G)$ is smaller than the genus, one cannot use the Riemann-Roch Theorem, and estimating $\dim(G)$ is often a challenging task. Furthermore, $\dim(G)$ now depends on G itself, and not just on its degree as before. However, we can formulate the question as follows. The important thing for us is not the best possible ratio between the number of evaluation points $N(F)$ and the genus. Instead, we are interested in the best possible ratio between $N(F)$ and $\deg(G)$, where G is a *low-degree* divisor having a *large dimension*.

Following the preliminary version of our work [3], Felipe Voloch [11] used a variant of Castelnuovo's bound to show our approach cannot lead to error correcting codes approaching the Gilbert-Varshamov bound. We show that a careful analysis of Voloch's argument implies that all k-dimensional ε-balanced codes built using our approach must have length

$$n = \Omega\left(\frac{k}{\varepsilon^{2.5} \log^2(\frac{1}{\varepsilon})}\right).$$

The rest of the paper is organized as follows. In Section 2 we describe the construction and its analysis using Bézout's Theorem. Section 3 contains a description of the same construction in the terminology of algebraic function fields. In Subsection 3.1 we give the necessary background on algebraic function fields and geometric Goppa codes. Finally, in Section 4 we analyze the limits of our approach based on Voloch's work.

2 A self-contained elementary description of the construction

We first recall the definition of an ε-biased set.

Definition 2.1. A set $S \subseteq \{0,1\}^k$ is *ε-biased* if for every nonempty $T \subseteq [k]$,

$$\frac{1}{|S|}\left|\sum_{s\in S}(-1)^{\sum_{i\in T} s_i}\right| \leq \varepsilon\,.$$

The construction Given k and ε, let $p = 2^\ell$ be a power of 2 in the range

$$\left[\frac{1}{2}\left(\frac{k}{\varepsilon^2}\right)^{1/4}, \left(\frac{k}{\varepsilon^2}\right)^{1/4}\right].$$

That is,

$$\frac{1}{16}\frac{k}{\varepsilon^2} \leq p^4 \leq \frac{k}{\varepsilon^2}\,.$$

Define $q = p^2$ and $r = \varepsilon p^3$. Let $\mathbb{F}_q$ denote the finite field with q elements and $\mathbb{F}_p$ its subfield with p elements. Consider the vector space of bivariate polynomials over $\mathbb{F}_q$ with total degree at most $r/(p+1)$.

$$V = \left\{\phi \in \mathbb{F}_q[x,y] \;:\; \deg(\phi) \leq \frac{r}{p+1}\right\} \;=\; \mathrm{Span}\left\{x^i y^j \;:\; i+j \leq \frac{r}{p+1}\right\}.$$

We denote the dimension of this space (over $\mathbb{F}_q$) by k'. It follows that

$$k' = \Theta\left(\frac{r^2}{p^2}\right) = \Theta\left(\frac{\varepsilon^2 p^6}{p^2}\right) = \Theta\left(\varepsilon^2 p^4\right) = \Theta(k)\,.$$

Let $A \subseteq \mathbb{F}_q \times \mathbb{F}_q$ be the set of roots of the polynomial $y^p + y - x^{p+1}$. The ε-biased set over k' bits that we construct is

$$S = \left\{\left(\left\langle \mathrm{bin}(a^i b^j), \mathrm{bin}(c)\right\rangle_2\right)_{i+j\leq \frac{r}{p+1}} \;:\; (a,b) \in A \text{ and } c \in \mathbb{F}_q\right\},$$

where $\mathrm{bin} : \mathbb{F}_q \to \mathbb{Z}_2^{2\ell}$ is any isomorphism between the additive group of $\mathbb{F}_q$ and the vector space $\mathbb{Z}_2^{2\ell}$ and $\langle\cdot,\cdot\rangle_2$ denotes inner product over $\mathbb{Z}_2^{2\ell}$.

The analysis Clearly, $|S| = q|A|$. We further claim:

Claim 2.2. *The cardinality of A is p^3.*

Proof. The trace function $\mathrm{Tr}(y) = y^p + y$ maps $\mathbb{F}_q$ to $\mathbb{F}_p$. We claim that for every $\alpha \in \mathbb{F}_p$, the number of solutions in $\mathbb{F}_q$ to $\mathrm{Tr}(y) = \alpha$ is p. To see this, observe that Tr is a linear function. Hence, the set of solutions to $\mathrm{Tr}(y) = 0$ is a subgroup of $\mathbb{F}_q$ that has at most p elements. For every $\alpha \in \mathbb{F}_p$, the set of solutions to $\mathrm{Tr}(y) = \alpha$ is either empty or a coset of this subgroup. As every element of $\mathbb{F}_q$ is in one of these cosets, it must be the case that for every $\alpha \in \mathbb{F}_p$ there are exactly p solutions.

The norm function $\mathbb{N}(x) = x^{p+1}$ also maps $\mathbb{F}_q$ to $\mathbb{F}_p$. Thus, for every $\alpha \in \mathbb{F}_q$ there are exactly p values $\beta \in \mathbb{F}_q$ such that $\mathrm{Tr}(\beta) = \mathbb{N}(\alpha)$. Therefore, $|A| = p^3$. □

In order to bound the bias ε, we need to invoke Bézout's theorem, reviewed below. First, we need to show that the bivariate polynomial $y^p + y - x^{p+1}$ is irreducible. We do this using Eisenstein's Criterion:

Theorem 2.3 (Eisenstein's Criterion [6, Thm 3.1])**.** *Let U be a unique factorization ring and let K be its field of fractions. Let $f(x) = \sum_{i=0}^{n} a_i x^i$ be a polynomial of degree $n \geq 1$ in $U[x]$. Let ρ be a prime of U, and assume:*

- $a_n \neq 0 \pmod{\rho}$,
- *for every $i < n$, $a_i = 0 \pmod{\rho}$,*
- $a_0 \neq 0 \pmod{\rho^2}$.

Then $f(x)$ is irreducible in $K[x]$.

With that we conclude:

Claim 2.4. *The polynomial $y^p + y - x^{p+1}$ is irreducible over $\mathbb{F}_q$.*

Proof. This follows from Eisenstein's Criterion. The unique factorization ring we consider is $U = \mathbb{F}_q[y]$. The prime element we use is $\rho = y$. The leading coefficient is $a_{p+1} = -1$ and $-1 \neq 0 \pmod{y}$. Every other coefficient except the last is 0, hence it is $0 \pmod{y}$. The last coefficient is $a_0 = y^p + y = 0 \pmod{y}$. Finally, since $p \geq 2$, $y^p = 0 \pmod{y^2}$ but $y \neq 0 \pmod{y^2}$, hence $a_0 = y^p + y \neq 0 \pmod{y^2}$. Therefore the univariate polynomial (in x) is irreducible over the field of fractions. As one of the coefficients is -1, it follows that the bivariate polynomial is irreducible over the field $\mathbb{F}_q$ (see [6, Thm 2.3]). □

We are now ready to recall Bézout's Theorem and apply it to prove S is indeed ε-biased.

Theorem 2.5 (Bézout's Theorem [4, Section 5.3])**.** *Suppose ϕ and ψ are two bivariate polynomials over some field. If ϕ and ψ have more than $\deg(\phi) \cdot \deg(\psi)$ common roots then they have a common factor.*

Theorem 2.6. *For every k and ε such that $\varepsilon < 1/\sqrt{k}$, S is an ε-biased set over $k' = \Omega(k)$ bits of size*

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

Proof. By Claim 2.2,

$$|S| = |A| \cdot q = p^5 = O\left(\frac{k}{\varepsilon^2}\right)^{5/4}.$$

We now show that S is ε-biased. Let $T \subseteq [k']$ be some non-empty set. We identify $[k']$ with the set

$$\left\{(i,j) \,:\, i + j \leq \frac{r}{p+1}\right\}$$

and T with the corresponding subset.

Let $s \in S$ be an element specified by the pair $((a,b),c) \in A \times \mathbb{F}_q$. Then,

$$\sum_{(i,j)\in T} s_{(i,j)} = \sum_{(i,j)\in T} \langle \mathrm{bin}(a^i b^j), \mathrm{bin}(c) \rangle_2 = \left\langle \mathrm{bin}\left(\sum_{(i,j)\in T} a^i b^j \right), \mathrm{bin}(c) \right\rangle_2 .$$

The polynomial $\phi_T = \sum_{(i,j)\in T} x^i y^j$ is a non-zero polynomial. Clearly, for any (a,b) which is not a root of ϕ_T, the inner product will be unbiased when ranging over c (i. e., exactly half of the values for c will make the inner product 0). From the assumption $\varepsilon < 1/\sqrt{k}$ it follows that $\deg(\phi_T) < p+1$, since

$$\frac{\deg(\phi_T)}{p+1} \leq \frac{r}{(p+1)^2} < \varepsilon p \leq k^{1/4}\sqrt{\varepsilon} < 1 .$$

Hence, by Claim 2.4 it follows that ϕ_T and $y^p + y - x^{p+1}$ have no common factors. Therefore, by Bézout's theorem we conclude that the number of roots of ϕ_T that are in A is at most $(\frac{r}{p+1}) \cdot (p+1) = r$ and

$$\frac{1}{|S|} \left| \sum_{s\in S} (-1)^{\sum_{i\in T} s_i} \right| \leq \frac{r}{|A|} = \varepsilon . \qquad \square$$

Remark 2.7. The above construction can be improved to an ε-biased set of size

$$O\left(\frac{k}{\varepsilon^2 \log(\frac{1}{\varepsilon})} \right)^{5/4} \quad \text{for every } k \text{ and } \varepsilon \text{ such that } \quad \frac{\varepsilon}{\sqrt{\log(\frac{1}{\varepsilon})}} < \frac{1}{\sqrt{k}} .$$

To achieve this we choose

$$p = \Theta\left(\frac{k}{\varepsilon^2 \log(\frac{1}{\varepsilon})} \right)^{1/4} .$$

We then observe that instead of taking a basis for V over $\mathbb{F}_q$, we can actually afford to take a basis over $\mathbb{F}_2$. Finally, we need to use the fact that by the constraints we have on ε, it follows that $\log(1/\varepsilon) = \Theta(\log(p))$. When we restate the construction in the terminology of algebraic function fields, we also include this improvement.

3 Restating the construction in the terminology of algebraic function fields

Without putting the above construction in the proper context, it may appear coincidental. We now describe the general framework of algebraic-geometric codes and explain why the above construction fits into this framework.

3.1 Algebraic geometry

We recall a few notions from the theory of algebraic function fields. A detailed exposition of the subject can be found, e. g., in [8].

Let $\mathbb{F}_q$ denote the finite field with q elements. The polynomial ring $\mathbb{F}_q[x]$, where x is transcendental over $\mathbb{F}_q$, is the set of all *polynomials* in x with coefficients in $\mathbb{F}_q$. The *rational* function field $\mathbb{F}_q(x)$, where x is transcendental over $\mathbb{F}_q$, contains all *rational* functions in x with coefficients in $\mathbb{F}_q$. A field F is an algebraic function field over $\mathbb{F}_q$, denoted $F/\mathbb{F}_q$, if F is a *finite* algebraic extension of $\mathbb{F}_q(x)$.

A *place* P of $F/\mathbb{F}_q$ is a maximal ideal of some valuation ring O of the function field. We denote by O_P the valuation ring that corresponds to the place P. We denote by v_P the *discrete valuation* that corresponds to the valuation ring O_P. Therefore, we can write P and O_P as

$$P = \{y \in F \,:\, v_P(y) > 0\} \quad \text{and} \quad O_P = \{y \in F \,:\, v_P(y) \geq 0\}.$$

Since P is a maximal ideal, $F_P = O_P/P$ is a field. In fact, it is a finite field [8, Proposition I.1.14]. For every $y \in O_P$, $y(P)$ denotes $y \pmod P$ and is an element of F_P. It can be thought of as the evaluation of the function y at the "evaluation point" P. The degree of a place P is defined to be $\deg(P) = [F_P : \mathbb{F}_q]$, i. e., the dimension of F_P as a vector space over $\mathbb{F}_q$. In particular, if a place P is of degree 1 then F_P is isomorphic to $\mathbb{F}_q$ and the evaluation of y at the place P is an element of $\mathbb{F}_q$.

We proceed with a simple example that illustrates the above notions.

Example 3.1. The rational function field $\mathbb{F}_q(x)/\mathbb{F}_q$ is the simplest algebraic function field. For every irreducible polynomial $p(x)$ let O_p denote the set of rational functions $r(x) = u(x)/w(x)$ whose denominator $w(x)$ is not divisible by p. The set O_p forms a ring. Furthermore, for every rational function r, either r or r^{-1} belongs to O_p, making O_p a valuation ring. The place P_p is the set of rational functions $r(x) = u(x)/w(x)$ for which p divides u but does not divide w.

There is exactly one more valuation ring in $\mathbb{F}_q(x)$. Let O_∞ denote the set of rational functions $r(x) = u(x)/w(x)$ for which $\deg(w) \geq \deg(u)$. Again, O_∞ is a ring and furthermore, for every rational function $r(x) = u(x)/w(x)$, either r or r^{-1} belongs to O_∞, making O_∞ a valuation ring. The place P_∞ is the set of rational functions $r(x) = u(x)/w(x)$ for which $\deg(w) > \deg(u)$.

For an irreducible polynomial p, the discrete valuation v_p corresponding to O_p is defined as follows. For a polynomial $u(x) \in \mathbb{F}_q[x]$, $v_p(u)$ is the largest integer k such that p^k divides u. For a rational function $r(x) = u(x)/w(x)$, $v_p(r) = v_p(u) - v_p(w)$. Thus, $v_p(r)$ counts the number of zeroes (or poles) that r has when the irreducible polynomial $p(x)$ is zero. If p is linear, i. e., $p(x) = x - \alpha$, $v_p(r)$ counts the number of zeroes (or poles) the polynomial $p(x)$ has when setting $x = \alpha$. If in addition $r = u/w$ belongs to O_p, i. e., p does not divide w, then $r(P_{x-\alpha}) = r \bmod (x-\alpha)$ is well defined, and, in fact, is the element $u(\alpha)/w(\alpha)$ in $\mathbb{F}_q$.

The discrete valuation v_∞ corresponding to O_∞ is defined as follows. For a polynomial $u(x) \in \mathbb{F}_q[x]$, $v_\infty(u)$ is $-\deg(u)$. For a rational function $r(x) = u(x)/w(x)$, $v_\infty(r) = v_\infty(u) - v_\infty(w) = \deg(w) - \deg(u)$. Thus, $v_\infty(r)$ counts the number of zeroes (or poles) that r has when $x = \infty$. If $r(x) = u(x)/w(x)$ belongs to O_∞, i. e., $\deg(w) \geq \deg(u)$, then $r(P_\infty)$ is well defined, and is either zero (if $\deg(w) > \deg(u)$) or the element u_k/w_k in $\mathbb{F}_q$, where u_k and w_k are the leading coefficients of the polynomials u and w respectively.

We let $\mathcal{P}_F$ denote the set of places of F, and $N(F)$ the number of places of *degree* 1 in $F/\mathbb{F}_q$. $N(F)$ is always finite. $\mathcal{D}_F$ is the free abelian group over the places of F. A *divisor* is an element in this group, i. e., it is a formal sum $G = \sum_{P \in \mathcal{P}_F} n_P P$ with $n_P \in \mathbb{Z}$ and where $n_P \neq 0$ for only a finite number of places. We also denote $v_P(G) = n_P$. The *degree* of the divisor $\sum_P n_P P$ is defined to be $\sum_P n_P \cdot \deg(P)$, and is always

finite. We say $G_1 \geq G_2$ if G_1 is component-wise larger than G_2, i. e., $v_P(G_1) \geq v_P(G_2)$ for every place P. The *support* of a divisor G is $\mathrm{Supp}(G) = \{P \in \mathcal{P}_F \ : \ v_P(G) \neq 0\}$.

Each element $0 \neq x \in F$ is associated with two divisors. The first is called the *principal divisor* of x and it is defined by

$$(x) = \sum_P v_P(x)P\,.$$

The degree of a principal devisor is always 0. The second is the *pole divisor* of x and it is defined by

$$(x)_\infty = \sum_{P:v_P(x)<0} -v_P(x)P\,.$$

If $x \in F \setminus \mathbb{F}_q$ then $\deg((x)_\infty) = [F : \mathbb{F}_q(x)]$.

Example 3.2 (Continued from Example 3.1)**.** Let $u \in \mathbb{F}_q[x]$ be an arbitrary nonconstant polynomial. Then, u has a pole at P_∞ (since $v_\infty(u) < 0$) and zero at P_p, for every irreducible polynomial p that divides u (since $v_p(u) > 0$). Thus, $(u)_\infty = \deg(u)P_\infty$ and $\deg((u)_\infty) = \deg(u)$. Also, it turns out that $deg(P_p) = \deg(p)$. Thus, $\sum_p v_p(u)$ is the number of irreducible factors u has, and $\sum_p v_p(u)\deg(P_p)$ is the degree of u. In total, $\deg((u)) = 0$.

For a divisor G, we define the *Riemann-Roch space* $\mathcal{L}(G)$ to be:

$$\mathcal{L}(G) = \{x \in F \ : \ (x) \geq -G\} \cup \{0\}\,.$$

Example 3.3 (Continued from Examples 3.1 and 3.2)**.** For the divisor $G = kP_\infty$ the Riemann-Roch space $\mathcal{L}(G)$ is the set of all polynomials of degree at most k over $\mathbb{F}_q$.

The set $\mathcal{L}(G)$ is a vector space and furthermore has finite dimension. We define the dimension of G by $\dim(G) = \dim \mathcal{L}(G)$ and we use the two notations interchangeably. The fact that the degree of each principal divisor is 0 implies that if $\deg(G) < 0$ then $\dim(\mathcal{L}(G)) = 0$.

3.1.1 Geometric Goppa codes

A Goppa code is specified by a triplet (F, Y, G), where $F/\mathbb{F}_q$ is a function field, $Y = \{P_1, \ldots, P_n\}$ is a set of places of degree 1 and G is an arbitrary divisor with no support over any place in Y. Notice that for any $x \in \mathcal{L}(G)$, $v_{P_i}(x) \geq 0$ and therefore $x \in O_{P_i}$ and $x(P_i) \in \mathbb{F}_q$. The triplet (F, Y, G) specifies the code:

$$C(Y;G) = \{(x(P_1), \ldots, x(P_n)) \ : \ x \in \mathcal{L}(G)\} \subseteq {\mathbb{F}_q}^n\,.$$

Claim 3.4 ([8, Cor II.2.3])**.** *If* $\deg(G) < n$ *then* $C(Y;G)$ *is an* $[n, \dim(G), n - \deg(G)]$ *linear code over* $\mathbb{F}_q$.

We call $n - \deg(G)$ the *designated* distance of the code. Given the designated distance d, we would like to find a code G with that designated distance and maximal dimension, i. e., maximize $\dim(G)$. It turns out that for every G, $\dim(G) \leq \deg(G) + 1$ (which is the Singleton bound in coding theory). Our goal is to minimize the gap between $\dim(G)$ and $\deg(G)$. It turns out that for any function field $F/\mathbb{F}_q$ there exists a constant $g \in \mathbb{N}$, such that for any divisor $G \in \mathcal{D}_F$, $\deg(G) - \dim(G) \leq g - 1$. The minimal integer with this property is called the *genus* of $F/\mathbb{F}_q$. The Riemann-Roch Theorem states that:

Theorem 3.5 ([8, Thm I.5.17]). *If* $\deg(G) \geq 2g-1$ *then* $\dim(G) = \deg(G) - g + 1$.

This, in particular, allows one to easily compute the dimension of the code when $\deg(G) \geq 2g-1$. The only remaining question is whether there are function fields $F/\mathbb{F}_q$ with a large number $N = N(F)$ of evaluation points (i. e., degree 1 places) and a small genus g. A *negative* answer to that question is given by the Hasse-Weil bound:

Theorem 3.6 (Hasse-Weil bound [8, Thm V.2.3]). *Let* $F/\mathbb{F}_q$ *be a function field of genus* g. *Then, the number* N *of places of degree one satisfies* $N \leq (q+1) + 2\sqrt{q}g$.

The Drinfeld-Vlăduţ bound tells us that when g tends to infinity, the bound can be strengthened by about a factor of 2, and roughly speaking, $N \leq g(\sqrt{q}-1)$. This is tight for prime power squares q.

On the positive side, the good news is that following much research, there are several beautiful explicit constructions that meet the Drinfeld-Vlăduţ bound, and we refer the interested reader to the beautiful survey paper [5, Chapter 1].

In this paper we look at divisors G whose degree is smaller than the genus. Much less is known about such small-degree divisors. In this regime, $\dim(G)$ depends on the divisor G itself, and not only on its degree, as is the case when $\deg(G) \geq 2g-1$. For some special algebraic function fields the vector space $\mathcal{L}(G)$ (and therefore also its dimension) is known in full. We discuss this below.

3.2 Concatenating AG codes with Hadamard

In this section we consider the concatenation of an outer code with the Hadamard code. If the outer code is an $[n_1, k_1, d]_q$ code and q is a power of two, then concatenating it with the $[q, \log q, q/2]_2$ Hadamard code gives an $[n = n_1 q, k = k_1 \log q]_2$ code that is $\varepsilon = (n_1 - d)/n_1$ balanced, because non-zero symbols in the outer code expand by the concatenation to perfectly balanced blocks.[1]

Using a $[q, k_1, q-k_1+1]_q$ Reed-Solomon code as the outer code, one gets an $[n = q^2, k = k_1 \log q]_2$ code that is $\varepsilon < k_1/q$ balanced. Rearranging the parameters, this gives an $[n,k]_2$ ε-balanced code with

$$n = O\left(\frac{k}{\varepsilon \log(\frac{k}{\varepsilon})}\right)^2. \tag{3.1}$$

This is one of the constructions in [2].

Taking the outer code to be an $[N, \dim(G), N - \deg(G)]_q$ AG code $C(Y;G)$ over $\mathbb{F}_q$ and concatenating it with Hadamard, we get an $[n = Nq, k = \dim(G)\log q]_2$ code that is $\varepsilon = \deg(G)/N$ balanced. We can choose an AG code which uses a curve of genus g with $N = \Theta(g\sqrt{q})$ degree 1 places (the asymptotic is over g going to infinity). Picking the divisor G to be of degree $\deg(G) \geq 2g$ and setting $q = 1/\varepsilon^2$ results in

$$N = \frac{\deg(G)}{\varepsilon} = \frac{\dim(G) + g - 1}{\varepsilon} = \Theta\left(\frac{k}{\varepsilon \log(\frac{1}{\varepsilon})}\right),$$

[1] If q is a power of 2, then the resulting concatenated code is linear. Concatenation is well defined even when q is not a power of 2. In such a case we embed $\mathbb{F}_q$ into $\mathbb{F}_2^{\lceil \log q \rceil}$ using any one-to-one mapping. The resulting (non-linear) code has essentially the same dimension and distance as in the previous case – the only difference is a small loss due to the fact that $2^{\lceil \log q \rceil}$ is slightly larger than q. From now on we will discuss the simpler case where q is a power of two, keeping in mind that everything also holds for arbitrary q.

where the second equality follows from the Riemann-Roch Theorem. Thus, we get an ε-balanced code of length

$$n = Nq = O\left(\frac{k}{\varepsilon^3 \log(\frac{1}{\varepsilon})}\right). \tag{3.2}$$

In fact, if one takes an AG code over $\mathbb{F}_q$ with large genus $g \geq \sqrt{q}$ then

$$N \geq \frac{\dim(G)}{\varepsilon} = \frac{k}{\varepsilon \log q} \quad \text{and} \quad q = \Omega\left(\frac{1}{\varepsilon^2}\right)$$

and equation (3.2) is tight. Taking an AG code with a small genus $g \leq \sqrt{q}$ is essentially equivalent to taking a Reed-Solomon outer code and cannot be better (up to constant factors) than equation (3.1). In what follows, we show one can improve on both bounds when the AG code has degree much smaller than the genus.

So we now turn our attention to the case where $\deg(G) \leq 2g-1$. In this case $\dim(G)$ depends on the divisor G and not just its degree. One special case is the case where $G = rQ$, $r \in \mathbb{N}$ and Q is a place of degree 1. For any such r, $\dim(rQ)$ is either equal to $\dim((r-1)Q)$ or to $\dim((r-1)Q)+1$. In the former case r is said to be a *gap number* of Q. The Weierstrass Gap Theorem [8, Thm I.6.7] says that for any place Q there are exactly $g = \text{genus}(F/\mathbb{F}_q)$ gap numbers, and they are all in the range $[1, 2g-1]$.

The non-gap numbers (also called *pole numbers*) form a semigroup of $\mathbb{N}$ (i. e., a set that is closed under addition). This semigroup is sometimes referred to as the *Weierstrass semigroup* of Q. We say that a semi-group S is *generated* by a set of elements $\{g_i\}$, if each $g_i \in S$ and, furthermore, every element $s \in S$ can be expressed as $s = \sum a_i g_i$ with $a_i \in \mathbb{N}$.

The structure of the Weierstrass semigroup is crucial to our construction. We know that there are exactly g non-gap elements of this semigroup in the range $[1, 2g]$. If these elements are too concentrated on the upper side of the range then the behavior of $\dim(rQ)$ will be very similar to the case where $r > 2g-1$. Thus, our goal is to find a function field F that has many places of degree 1, say, $N(F) \geq \Omega(g\sqrt{q})$, while at the same time F has a degree 1 place Q with a "good" Weierstrass semigroup.

3.3 The construction

Let p be a prime power and $q = p^2$. The Hermitian function field over $\mathbb{F}_q$ (see [8, Lemma VI.4.4]) can be represented as the extension field $\mathbb{F}_q(x,y)$ of the rational function field $\mathbb{F}_q(x)$ with $y^p + y = x^{p+1}$. This function field has $1 + p^3$ places of degree one. First, there is the common pole Q_∞ of x and y. Moreover, for each pair $(\alpha, \beta) \in \mathbb{F}_q$ with $\beta^p + \beta = \alpha^{p+1}$ there is a unique place $P_{\alpha,\beta}$ of degree one such that $x(P_{\alpha,\beta}) = \alpha$ and $y(P_{\alpha,\beta}) = \beta$ and we already saw there are p^3 such pairs. The genus of the Hermitian function field is $g = p(p-1)/2$.

For the outer code we take the Goppa code $C_r = C(Y, G = rQ_\infty)$, where Y is the set of all degree 1 places $P_{\alpha,\beta}$ mentioned above and $r = \varepsilon p^3$. The Weierstrass semigroup of G is generated by p and $p+1$, and a basis for $\mathcal{L}(G) = \mathcal{L}(rQ_\infty)$ is

$$\left\{x^i y^j \ : \ j \leq p-1 \text{ and } ip + j(p+1) \leq r\right\}.$$

The dimension of the code is

$$\left|\{(i,j) \ : \ j \leq p-1 \text{ and } ip + j(p+1) \leq r\}\right|.$$

We can now see the similarity between this construction and the one in Section 2. The parameter r is chosen such that the constraint $ip+j(p+1)\leq r$ forces $j\leq p-1$. Therefore, both use evaluations of low degree bivariate polynomials over the same set of p^3 points.[2]

Theorem 3.7. *For every k and every ε such that*

$$\frac{\varepsilon}{\sqrt{\log(1/\varepsilon)}}\leq\frac{1}{\sqrt{k}},$$

there exists an explicit $[n,\Omega(k)]_2$ code that is ε-balanced, with

$$n=O\left(\frac{k}{\varepsilon^2\log(1/\varepsilon)}\right)^{5/4}.$$

Proof. For a given k and ε, let

$$p\in\left[\frac{1}{2}\left(\frac{k}{\varepsilon^2\log(1/\varepsilon)}\right)^{1/4},\left(\frac{k}{\varepsilon^2\log(1/\varepsilon)}\right)^{1/4}\right]$$

be a power of two. It can verified that

$$\frac{1}{16p^4}\;\leq\;\varepsilon\;\leq\;\frac{1}{p}$$

as

$$\frac{1}{p}\;\geq\left(\frac{\varepsilon^2\log(1/\varepsilon)}{k}\right)^{1/4}\geq\left(\varepsilon^2\log(1/\varepsilon)\cdot\frac{\varepsilon^2}{\log(1/\varepsilon)}\right)^{1/4}=\;\varepsilon$$

and

$$\varepsilon\;=\;\varepsilon^2\cdot\frac{1}{\varepsilon}\;\geq\;\varepsilon^2\cdot\log\left(\frac{1}{\varepsilon}\right)\;\geq\;\frac{k}{16p^4}\;\geq\;\frac{1}{16p^4},$$

and so $\log(1/\varepsilon)=\Theta(\log(p))$.

Let $r=\varepsilon p^3$ and let $\mathbb{F}_q$ be the field with $q=p^2$ elements. Let F denote the Hermitian function field over $\mathbb{F}_q$ and let Y denote its set of places of degree 1, excluding Q_∞. This implies that $|Y|=p^3$. Define the divisor G to be $G=rQ_\infty$. Since $r\leq p^2$,

$$\dim(rQ_\infty)\geq\left(\frac{r}{2(p+1)}\right)^2=\Omega\left(\varepsilon^2p^4\right)=\Omega\left(\frac{k}{\log(p)}\right).$$

By Claim 3.4, the Goppa code that is obtained from the triplet (F,Y,G) is a

$$\left[p^3,\Omega\left(\frac{k}{\log(p)}\right),p^3-r\right]_{p^2}$$

[2]The only slight difference is that in this construction we take all bivariate polynomials with bounded *weighted* total degree. However, the weight is nearly identical for both variables and so this does not affect much the parameters of the construction.

code. Concatenating this code with Hadamard gives a $[p^5, \Omega(k)]_2$ code that is ε-balanced (since $r/p^3 = \varepsilon$). Now, by our choice of p, it follows that

$$\frac{k}{\varepsilon^2 \log(\frac{1}{\varepsilon})} = \Theta(p^4)$$

and therefore

$$n = p^5 = O\left(\left(\frac{k}{\varepsilon^2 \log(\frac{1}{\varepsilon})}\right)^{5/4}\right)$$

as desired. □

4 Limits of the approach

As explained in Section 3.1.1, the genus measures the maximal loss in dimension compared to the degree. The Drinfeld-Vlăduţ bound implies that the number of evaluation points (which is bounded by the number of degree one places $N(F)$) is at most $O(g\sqrt{q})$ when $N(F) \gg q$. In Section 3.2 we saw this implies that when $\deg(G) > 2g$, concatenating the best AG code $C(Y;G)$ with the Hadamard code cannot give ε-balanced codes of dimension k and length

$$n = o\left(\frac{k}{\varepsilon^3 \log(1/\varepsilon)}\right).$$

Our construction shows that substantially better results are possible when $\deg(G) \ll g$. Namely, we show that there exists a code $C(Y;G)$ with $\deg(G) \ll g$ such that when this code is concatenated with Hadamard, it gives a k-dimensional ε-balanced code of length

$$n = O\left(\left(\frac{k}{\varepsilon^2 \log(1/\varepsilon)}\right)^{5/4}\right).$$

It is therefore natural to ask what are the limits of our approach. More concretely we ask what are the best codes one can construct by concatenating an AG code with a Hadamard code? Let us state the question precisely. We look at constructions of the following structure:

- An outer AG code $C = C(Y;G)$, defined by an algebraic function field $F/\mathbb{F}_q$, a set of degree 1 places Y and a divisor $G \in \mathcal{D}_F$ with no support over any place in Y.
- An inner Hadamard code.

In the analysis we view C as a $[|Y|, \dim(G), |Y| - \deg(G)]_q$ code, and then the concatenated code has parameters

$$\left[|Y|q,\ \dim(G)\log(q),\ \frac{1}{2} - \frac{\deg(G)}{|Y|}\right]_2.$$

Notice that it may be the case that C has better distance than the so-called *designated* distance, but as far as we are concerned the analysis does not take advantage of that, and we take the distance to be $|Y| - \deg(G)$.

In this section we prove:

Theorem 4.1. *Any ε-balanced $[n,k]_2$ code that is constructed and analyzed as above, must have*

$$n \geq \Omega\left(\frac{k}{\varepsilon^2} \cdot \min\left\{\frac{k}{\log^2(\frac{k}{\varepsilon})}, \frac{1}{\sqrt{\varepsilon}\log(\frac{k}{\varepsilon})}\right\}\right).$$

For the proof we need definitions and theorems about finite extensions of algebraic functions fields. Specifically, for an extension F of a function field F', we use the following notation:

- A place $P \in \mathcal{P}_F$ *lying over* a place $P' \in \mathcal{P}_{F'}$, denoted by $P|P'$, see [8, Def III.1.3],
- The *ramification index* of P over P', denoted by $e(P|P')$, see [8, Def III.1.5],
- The *conorm* of a divisor $G' \in \mathcal{D}_{F'}$, denoted by $\mathrm{Con}_{F/F'}(G')$, see [8, Def III.1.8].

For more details we refer the reader to [8, Chapter III].

4.1 AG theorems about degree vs. dimension

It turns out that the above question boils down to the question of whether there are function fields with many degree 1 places (compared to the genus) and with low-degree divisors (of degree much smaller than the genus) of high dimension. We start by presenting two AG theorems relating degree to dimension in the small degree regime (when the degree is smaller than the genus).

The first argument we present shows that any divisor with non-trivial dimension must have degree at least $N(F)/(q+1)$. The argument was shown to us by Henning Stichtenoth [9].

Lemma 4.2. *Let $F/\mathbb{F}_q$ be a function field and $G \in \mathcal{D}_F$ a divisor with* $\dim(G) > 1$. *Then $N(F) \leq \deg(G) \cdot (q+1)$.*

Proof. As $\dim(G) > 1$, there exists some $x \in F \setminus \mathbb{F}_q$ such that $(x) \geq -G$. Fix any such x. In particular, $\deg((x)_\infty) \leq \deg(G)$. Also, by [8, Thm I.4.11], $\deg(x)_\infty = [F : \mathbb{F}_q(x)]$. We may view F as a finite extension over the rational function field $\mathbb{F}_q(x)$. Every place of degree 1 of F lies above some place of degree 1 of $\mathbb{F}_q(x)$. There are exactly $q+1$ places of degree 1 of $\mathbb{F}_q(x)$, and each one of them may split to at most $[F : \mathbb{F}_q(x)]$ places of degree 1 of F (by the fundamental equality, [8, Thm III.1.11]). Altogether, $N(F) \leq (q+1)[F : \mathbb{F}_q(x)] = (q+1)\deg(x)_\infty \leq (q+1)\deg(G)$. □

Remark 4.3. Lemma 4.2 only uses the fact that G is non-trivial. We wonder if one can strengthen the lemma for divisors G of high dimension. In particular, is it true that if $\dim(G) > \ell$ then

$$N(F) \leq \frac{\deg(G) \cdot (q+1)}{f(\ell)}$$

for some function f that goes to infinity with ℓ?

We now move to the second theorem. For a set $S \subseteq F$, where F is a field, let $\mathrm{Closure}(S)$ denote the minimal subfield of F that contains S. Following our work, Voloch [11] showed, based on the Castelnuovo bound, that:

Theorem 4.4 ([11, based on the Castelnuovo bound])**.** *Let K be an arbitrary field. Let F/K be a function field of genus g. Let $G \in \mathcal{D}_F$ be a divisor with degree $d+1$ and dimension $\ell+2$ such that* $\text{Closure}(\mathcal{L}(G)) = F$*. Let $m = d \text{ div } \ell$ and $r = d \bmod \ell$. Then*

$$g \leq m(m-1)\ell + m(2r+1)\,,$$

and, in particular, $g \leq m(m+1)\ell$.

Using Theorem 4.4 requires an assumption on the AG code, namely, that the closure of the Riemann space of the divisor used to define the code is the entire function field F. The following lemma, based on private communication with Voloch, shows that this assumption is inessential when analyzing the rate versus distance problem.

Lemma 4.5. *Let K be a finite field, F/K be a function field, $G \in \mathcal{D}_F$ is a divisor. Let C be the Goppa code of length n, dimension k, and designated relative distance δ, specified by some triplet (F, Y, G). Define a function field $F' = \text{Closure}(\mathcal{L}(G))$. Then there exists a Goppa code C' defined by a triplet (F', Y', G'), of length $n' \leq n$, dimension k and designated relative distance $\delta' \geq \delta$, such that* $\text{Closure}(\mathcal{L}(G')) = F'$*.*

Proof. We first define the new triplet (F', Y', G').

- We already have that $F' = \text{Closure}(\mathcal{L}(G))$.
- Next let $B = \{s_1, \ldots, s_k\}$ be a basis for $\mathcal{L}(G)$. Define,

  $$G'' = \sum_{P' \in \mathcal{P}_{F'}} \max_i \{-v_{P'}(s_i)\} \cdot P'\,.$$

  We would like to exchange G'' with an equivalent divisor that has no support over places of degree 1. By the Weak Approximation Theorem [8, Thm I.3.1] there exists $z \in F/K$ such that for every place P of degree 1, $v_P(z) = -v_P(G)$ and we let

  $$G' = G'' + (z)\,.$$

- Define a set $Y' \subset \mathcal{P}_{F'}$ by

  $$Y' = \left\{P' \in \mathcal{P}_{F'} \,:\, \exists P \in Y \text{ such that } P \mid P'\right\}.$$

  Observe that since Y consists only of places of degree 1 this is also true for Y' (see [8, Proposition III.1.6]).

Consider the Goppa code C' defined by the triplet (F', Y', G'). Notice that Y' does not intersect G' because Y' contains only degree 1 places and G' has no support over degree 1 places. We will prove:

- The dimension of C' is the same as C, i. e., $\dim(\mathcal{L}'_F(G')) = k$.
- The length of C' is at most the length of C, i. e., $n' = |Y'| \leq |Y| = n$.

- The designated relative distance of C' is at least as good as in C, i. e.,

$$\delta' = 1 - \frac{\deg(G')}{|Y'|} \geq \delta\,.$$

For the proof we will show:

Claim 4.6.

$$G \geq \mathrm{Con}_{F/F'}(G'')\,.$$

With that we can prove the three assertions above about C':

Dimension: Since $\mathcal{L}_{F'}(G'') \subseteq \mathcal{L}_F(\mathrm{Con}_{F/F'}(G''))$, it follows that

$$\dim(\mathcal{L}_{F'}(G'')) \leq \dim(\mathcal{L}_F(\mathrm{Con}_{F/F'}(G')))\,.$$

Thus, by Claim 4.6,

$$\dim(\mathcal{L}_F(G)) \geq \dim(\mathcal{L}_F(\mathrm{Con}_{F/F'}(G''))) \geq \dim(\mathcal{L}_{F'}(G'')) \geq |B| = \dim(\mathcal{L}_F(G))\,,$$

and therefore

$$\dim(\mathcal{L}_{F'}(G'')) = \dim(\mathcal{L}_F(G)) = k\,.$$

The claim follows since $\dim(\mathcal{L}_{F'}(G')) = \dim(\mathcal{L}_{F'}(G''))$ by [8, Lemma I.4.6].

Length: $|Y| \geq |Y'|$, since every place of Y lies over exactly one place of Y', see [8, Proposition III.1.7].

Designated distance: By Claim 4.6, $\deg(G) \geq \deg(\mathrm{Con}_{F/F'}(G''))$. By [8, Cor III.1.13],

$$\deg(\mathrm{Con}_{F/F'}(G'')) = [F : F'] \cdot \deg(G'')\,.$$

Since $\deg(G'') = \deg(G')$ it follows that

$$\deg(G) \geq [F : F'] \cdot \deg(G')\,.$$

Also, since every place P' can split to at most $[F : F']$ places in F/K we have

$$|Y| \leq |Y'| \cdot [F : F']\,.$$

Altogether,

$$\delta' = 1 - \frac{\deg(G')}{|Y'|} \geq 1 - \frac{\deg(G)}{[F : F'] \cdot |Y'|} \geq 1 - \frac{\deg(G)}{|Y|} = \delta\,.$$

□

We are left with proving Claim 4.6.

Proof of Claim 4.6. By definition, $B \subseteq \mathcal{L}(G'')$, and therefore

$$F' = \text{Closure}(\mathcal{L}(G)) = \text{Closure}(B) \subseteq \text{Closure}(\mathcal{L}(G'')),$$

and $\text{Closure}(\mathcal{L}(G'')) = F'$.

Also, for any $P|P'$ (where $P' \in \mathcal{P}_{F'}$ and $P \in \mathcal{P}_F$) and for any i,

$$e(P|P') \cdot v_{P'}(s_i) = v_P(s_i) \geq -v_P(G),$$

where the last inequality is simply because $s_i \in \mathcal{L}(G)$. Therefore

$$v_{P'}(G'') = \max\{-v_{P'}(s_i)\} = \max\left\{-\frac{v_P(s_i)}{e(P|P')}\right\} \leq \frac{v_P(G)}{e(P|P')},$$

and the claim follows from the definition of the conorm. □

4.2 The bound

We are now ready to prove Theorem 4.1.

Proof of Theorem 4.1. Assume a code is obtained by concatenating the AG code specified by the triplet $(F/\mathbb{F}_q, Y, G)$ with the Hadamard code. Let $\ell = \dim(G)$ and $d = \deg(G)$. The AG code $C(Y;G)$ is a $[|Y|, \ell]_q$ code, with designated distance $|Y| - d$. The concatenated code is therefore a

$$[n = |Y| \cdot q, k = \ell \log(q)]_2$$

code which is ε-balanced for

$$\varepsilon = \frac{d}{|Y|}.$$

By Lemma 4.5 we can assume without loss of generality that $\text{Closure}(\mathcal{L}(G)) = F$.

There are two extreme cases that we handle separately:

Large base field: If the base field size q is too large the theorem is trivially true. Namely, If $q > \frac{k}{\varepsilon^3}$ we are done because $n \geq q > k/\varepsilon^3$. We can therefore assume without loss of generality that

$$q \leq \frac{k}{\varepsilon^3}$$

and

$$\log(q) = O\left(\log\left(\frac{k}{\varepsilon}\right)\right). \tag{4.1}$$

Few evaluation points: If the number of evaluation points is about the field size, we are essentially in the Reed-Solomon case and we are done. Specifically, if $|Y| \leq 4q$ then

$$4n = |Y| \cdot 4q \geq |Y|^2 = \frac{d^2}{\varepsilon^2} \geq \frac{\ell^2}{\varepsilon^2} = \frac{k^2}{\varepsilon^2 \log^2(q)} = \Omega\left(\frac{k^2}{\varepsilon^2 \log^2(\frac{k}{\varepsilon})}\right),$$

and we are done. We can therefore assume without loss of generality that

$$|Y| > 4q.$$

This also implies that $\sqrt{q} < g$ since by Theorem 3.6, $|Y| \leq N(F) \leq q+1+2g\sqrt{q}$. We can therefore conclude (again, by Theorem 3.6) that

$$N(F) \leq 4g\sqrt{q}. \tag{4.2}$$

Let $m = d \operatorname{div} \ell \geq 1$. By Theorem 4.4, $g \leq 2m^2\ell$, and by equation (4.2), $N(F) \leq 8m^2\ell\sqrt{q}$. Thus,

$$n = |Y| \cdot q = \frac{N(F) \cdot |Y| \cdot q}{N(F)} \geq \frac{N(F) \cdot |Y| \cdot q}{8m^2\ell\sqrt{q}}.$$

Substituting $m \leq d/\ell$ and $d = \varepsilon|Y|$ we see that

$$n \geq \frac{N(F) \cdot \sqrt{q} \cdot \ell}{8\varepsilon^2|Y|}.$$

Substituting $\ell = k/\log(q)$ and using $N(F) > |Y|$,

$$n \geq \frac{\sqrt{q} \cdot k}{8\varepsilon^2 \log(q)} = \Omega\left(\frac{\sqrt{q} \cdot k}{\varepsilon^2 \log(\frac{k}{\varepsilon})}\right),$$

where the last equality follows from equation (4.1). To finish the argument notice that by Lemma 4.2, $N(F) \leq d(q+1)$. This implies

$$\frac{d}{\varepsilon} = |Y| \leq d(q+1) \quad \text{and} \quad \varepsilon \geq \frac{1}{q+1},$$

hence,

$$n = \Omega\left(\frac{k}{\varepsilon^{2.5} \log(\frac{k}{\varepsilon})}\right). \qquad \square$$

4.3 An open problem

Can one strengthen the above lower bound to match the parameters given in Section 3? More specifically we ask whether it is possible to get a concatenated code with $n = \tilde{O}(k/\varepsilon^{2.5})$, where the $\tilde{O}$ notation is used to hide poly-logarithmic factors in q (or equivalently in k and ε). We know the following:

- $n = \tilde{O}(k^2/\varepsilon^2)$ implies $N(F) = \tilde{\Omega}(qm^2)$. (We already saw that in the proof of Theorem 4.1.)
- A similar calculation shows $n = \tilde{O}\left(k/\varepsilon^{2.5}\right)$ implies $N(F) = \tilde{\Omega}(q^{2/3}m^{5/3}\ell)$.

We also know two upper bounds on $N(F)$, namely:

- $N(F) = \tilde{O}(qm\ell)$ (follows from $N(F) \le d(q+1)$), and,
- $N(F) = \tilde{O}(q^{1/2}m^2\ell)$ (since we can assume $N(F) \ge 2(q+1)$, as explained in the proof of Theorem 4.1).

Solving the constraints we get $m = \tilde{\Theta}(\sqrt{q})$. We thus see that the approach can lead to codes with $n = \tilde{O}(k/\varepsilon^{2.5})$ if and only if the following question has a positive answer:

Open Problem 4.7. *Given a prime power q and an integer $d = \tilde{O}(q)$ is there an algebraic function field $F/\mathbb{F}_q$ with $\tilde{\Omega}(q^2)$ places of degree one, and a divisor G such that* $\deg(G) = d$ *and* $\dim(G) \ge \tilde{O}(d/\sqrt{q})$.

One might suspect such a high dimension, low-degree divisor does not exist. However, Theorem 4.4 and Lemma 4.2 are not strong enough to disprove it. We remark that the lower bound could be improved, if Lemma 4.2 could be strengthened to use the high-dimension of G, as suggested in Remark 4.3.

References

[1] NOGA ALON, JEHOSHUA BRUCK, JOSEPH NAOR, MONI NAOR, AND RON M. ROTH: Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. *IEEE Trans. Inform. Theory*, 38(2):509–516, 1992. Preliminary version in ISIT'91. [doi:10.1109/18.119713] 254

[2] 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. Preliminary version in FOCS'90. [doi:10.1002/rsa.3240030308] 254, 262

[3] AVRAHAM BEN-AROYA AND AMNON TA-SHMA: Constructing small-bias sets from algebraic-geometric codes. In *Proc. 50th FOCS*, pp. 191–197. IEEE Comp. Soc. Press, 2009. [doi:10.1109/FOCS.2009.44] 253, 256

[4] WILLIAM FULTON: *Algebraic Curves: An Introduction to Algebraic Geometry*. Third edition, 2008. Author's version. 258

[5] ARNALDO GARCIA AND HENNING STICHTENOTH, editors. *Topics in Geometry, Coding Theory and Cryptography*. Volume 6. Springer, 2007. Available from Springer. 262

[6] SERGE LANG: *Algebra*. Springer, revised third edition, 2002. Available from Springer. 258

[7] JOSEPH NAOR AND MONI NAOR: Small-bias probability spaces: Efficient constructions and applications. *SIAM J. Comput.*, 22(4):838–856, 1993. Preliminary version in STOC'90. [doi:10.1137/0222053] 254

[8] HENNING STICHTENOTH: *Algebraic Function Fields and Codes*. Springer, 1993. Available from Springer. 256, 259, 260, 261, 262, 263, 266, 267, 268

[9] HENNING STICHTENOTH: Private communication, 2009. 266

[10] MICHAEL A. TSFASMAN, SERGE G. VLĂDUŢ, AND THOMAS ZINK: Modular curves, Shimura curves, and Goppa codes, better than Varshamov-Gilbert bound. *Mathematische Nachrichten*, 109(1):21–28, 1982. [doi:10.1002/mana.19821090103] 254

[11] JOSÉ FELIPE VOLOCH: Special divisors of large dimension on curves with many points over finite fields. *Portugaliae Mathematica*, 68(1):103–107, 2011. [doi:10.4171/PM/1882] 256, 266, 267

AUTHORS

Avraham Ben-Aroya
postdoctoral researcher
The Weizmann Institute of Science
Rehovot, Israel
avraham.ben-aroya@weizmann.ac.il
http://www.wisdom.weizmann.ac.il/~benaroya/

Amnon Ta-Shma
professor
Tel-Aviv University
Tel-Aviv, Israel
amnon@tau.ac.il
http://www.cs.tau.ac.il/~amnon

ABOUT THE AUTHORS

AVRAHAM BEN-AROYA received his Ph. D. from Tel-Aviv University. His advisors were Oded Regev and Amnon Ta-Shma. He is currently a postdoc at the Weizmann Institute of Science.

AMNON TA-SHMA is a theoretical computer scientist at Tel-Aviv University.