Constructing Small-Bias Sets from Algebraic-Geometric Codes

Avraham Ben-Aroya
The Blavatnik School of Computer Science
Tel-Aviv University
Israel 69978. *

Amnon Ta-Shma
The Blavatnik School of Computer Science
Tel-Aviv University
Israel 69978. †

Abstract

We give an explicit construction of an ϵ-biased set over k bits of size $O\left(\frac{k}{\epsilon^2 \log(1/\epsilon)}\right)^{5/4}$. 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.

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.

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. [8] it was shown that there are Algebraic-Geometric codes over constant size alphabets that lie above the Gilbert-Varshamov bound, a bound that was believed to be optimal at the time.

*Email: abrhambe@tau.ac.il. 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.

†Email: amnon@tau.ac.il. Supported by Israel Science Foundation grant 217/05 and by USA Israel BSF grant 2004390.

The important case of *binary* error correcting codes is still open. The Gilbert-Varshamov bound gives the best known (explicit or non-explicit) codes to date. Finding an explicit construction that attains this bound is an open problem as well. The above statements also apply if we restrict ourselves to codes with distance close to half, which is a case of special interest.

Another closely related question is that of finding an $[n, k, \frac{1}{2} - \epsilon]_2$ binary code, in which the relative weight of every non-zero codeword is in the range $[\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon]$. 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 columns of a matrix whose rows generate an ϵ-balanced code form an ϵ-biased set, and vise 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 $[n, k, , \frac{1}{2} - \epsilon]_2$ codes. In both cases the probabilistic method gives non-explicit $[n, k]_2$ ϵ-balanced codes with $n = O(\frac{k}{\epsilon^2})$, whereas the best lower bound is $n = \Omega(\frac{k}{\epsilon^2 \log(\frac{1}{\epsilon})})$. For a discussion of these bounds see [1, Section 7].

There are several *explicit* constructions of such codes. Naor and Naor [5] give a construction with $n = k \cdot \mathrm{poly}(\epsilon^{-1})$. Alon et al. [1] have the incomparable bound $n = O(\frac{k^2}{\epsilon^2 \log^2(k/\epsilon)})$. Concatenating Algebraic-Geometric codes with the Hadamard code gives $n = O(\frac{k}{\epsilon^3 \log(\frac{1}{\epsilon})})$. In this paper we show an explicit construction of an $[n, k]_2$ ϵ-balanced code with $n = O(\frac{k}{\epsilon^2 \log(\frac{1}{\epsilon})})^{5/4}$, which improves upon previous explicit constructions when ϵ is roughly (ignoring logarithmic factors) in the range of $k^{-1.5} \leq \epsilon \leq k^{-0.5}$ (see Figure 1).

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 algebraic function fields terminology. 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. We explain what a divisor is and other terminology 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 $N = N(F)$ (F is the algebraic function field). The distance of the code is $N - \deg(G)$ ($\deg(G)$ is the degree of G, we explain what it is 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* (we explain what the genus is in Section 3), the Riemann-Roch Theorem [6, Thm I.5.17] tells us exactly what the dimension $\dim(G)$ is, and it turns out to be $\deg(G) - g + 1$. This almost matches the Singleton bound, except for a loss of g in the dimension. Thus, our goal is to get as many evaluation points while keeping the genus small. Indeed, a lot of research was done on the best possible ratio between the length of

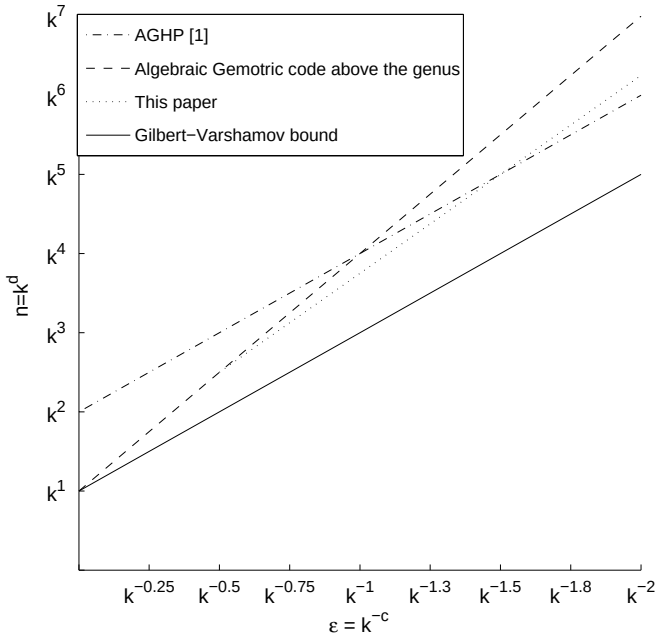

Figure 1: Constructions of ϵ-biased sets for $\epsilon = k^{-c}$

the code $N(F)$ and the genus. 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 Hadamard cannot give ϵ-balanced codes with n better than $O(\frac{k}{\epsilon^3 \log(\frac{1}{\epsilon})})$. 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 $\deg(G)$ 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 to us is not the best possible ratio between the number of rational 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 our work Felipe Voloch [9] 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 imply that all dimension k, ϵ-balanced codes built using our approach must have length $n = \Omega(\frac{k}{\epsilon^{2.5} \log^2(\epsilon)})$.

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 algebraic function fields terminology. In Subsections 3.1 and 3.1.1 we give the necessary background on algebraic function fields and geometric Goppa codes. Finally, in Section 4 we analyze the limits of our approach bases on Voloch's work.

2 A self-contained elementary description of the construction

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

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

$$\frac{1}{|S|}\Big|\sum_{s\in S}(-1)^{\sum_{i\in T} s_i}\Big| \leq \epsilon.$$

The construction: Given k and ϵ, let $p = 2^\ell$ be a power of 2 in the range $\left[\left(\frac{k}{\epsilon^2}\right)^{1/4}, 2\left(\frac{k}{\epsilon^2}\right)^{1/4}\right]$. Define $q = p^2$ and $r = \epsilon p^3$. 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 $r/(p+1)$:

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

The dimension of this space (over $\mathbb{F}_q$) is $k' = \Omega(\frac{r^2}{p^2}) = \Omega(k)$.

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

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

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: The following claim will be used to bound the size of S.

Claim 1. *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 $\mathrm{N}(x) = x^{p+1}$ 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 $\mathrm{Tr}(\beta) = \mathrm{N}(\alpha)$. Therefore, $|A| = p^3$. ∎

We want to apply Bézout's Theorem on the bivariate polynomial $y^p + y - x^{p+1}$. However, we first need to show it is irreducible. We need Eisenstein's Criterion for irreducibility:

Theorem 2 (Eisenstein's Criterion [4, 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 3. *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 -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 also $0 \pmod{y}$. Finally, since $p \geq 2$, $y^p = 0 \pmod{y^2}$ but $y \neq 0 \pmod{y^2}$, hence $y^p + y \neq 0 \pmod{y^2}$. Therefore the univariate polynomial (in x) is irreducible over the field of fractions and in particular over $\mathbb{F}_q[y]$. This implies the bivariate polynomial is irreducible over the field $\mathbb{F}_q$. ■

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

Theorem 4 (Bézout's Theorem [2, Section 5.3])**.** *Suppose* ϕ *and* ψ *are two bivariate polynomials over some field. If* ϕ *and* ψ *have more then* $\deg(\phi) \cdot \deg(\psi)$ *common roots than they have a common factor.*

Theorem 5. *For every* k *and* ϵ *such that* $\epsilon < \frac{1}{\sqrt{k}}$, S *is an* ϵ*-biased set over* $k' = \Omega(k)$ *bits of size* $O\left(\frac{k}{\epsilon^2 \log(1/\epsilon)}\right)^{5/4}$.

Proof: By Claim 1, $|S| = |A| \cdot q = p^5 = O\left(\frac{k}{\epsilon^2}\right)^{5/4}$.

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,

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

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 $\epsilon < \frac{1}{\sqrt{k}}$ it follows that $\deg(\phi_T) < p + 1$. Hence, by Claim 3 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,

$$\frac{1}{|S|} \Big| \sum_{s\in S} (-1)^{\sum_{i\in T} s_i} \Big| \leq \frac{r}{|A|} = \epsilon.$$

■

Remark 6. *The above construction can be improved to an ϵ-biased set of size $O\left(\frac{k}{\epsilon^2 \log(1/\epsilon)}\right)^{5/4}$ for every k and ϵ such that $\frac{\epsilon}{\sqrt{\log(1/\epsilon)}} < \frac{1}{\sqrt{k}}$. To achieve this we choose $p = \Theta\left(\frac{k}{\epsilon^2 \log(1/eps)}\right)^{1/4}$. 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 $\log(1/\epsilon) = \Theta(\log(p))$. When we restate the construction in algebraic function fields terminology, we also include this improvement.*

3 Restating the construction in algebraic function fields terminology

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 [6].

$\mathbb{F}_q$ denotes the finite field with q elements. $\mathbb{F}_q(x)$, where x is transcendental over $\mathbb{F}_q$, is the *rational* function field, and it contains all rational functions in x with coefficients in $\mathbb{F}_q$. $F/\mathbb{F}_q$ is an algebraic function field, 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

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

Since P is a maximal ideal, $F_P = O_P/P$ is a field. For every $x \in O_P$, $x(P)$ denotes $x(\bmod P)$ and is an element of F_P. The degree of a place P is defined to be $\deg(P) = [F_P : \mathbb{F}_q]$. In particular, if a place is of degree 1 then F_P is isomorphic to $\mathbb{F}_q$. $\mathcal{P}_F$ is the set of places of F. $N(F)$ is the number places of *degree* 1 (also called *rational points*) in $F/\mathbb{F}_q$ and 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 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 it 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 any place P.

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

For a divisor G, we define the *Riemann-Roch space* is

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

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 7 ([6, Cor II.2.3]). *If* $\deg(G) < n$ *then* $C(Y; G)$ *is an* $[n, \dim(\mathcal{L}(G)), n - \deg(G)]$ *linear code over* $\mathbb{F}_q$.

We want the gap between $\dim(\mathcal{L}(G))$ and $\deg(G)$ to be small. 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 $G \in \mathcal{D}_F$, $\deg(G) - \dim(\mathcal{L}(G)) \leq g - 1$. The minimal integer with this property is called the *genus* of $F/\mathbb{F}_q$. The Riemann-Roch Theorem says that:

Theorem 8 ([6, Thm I.5.17]). *If* $\deg(G) \geq 2g - 1$ *then* $\dim(\mathcal{L}(G)) = \deg(G) - g + 1$.

This, in particular, allows one to easily compute the dimension of the code when $\deg(G) > 2g$. The only remaining question is whether there are function fields with a large number $N = N(F)$ of rational points, and a small genus g. This is addressed in:

Theorem 9 (Hasse-Weil bound [6, 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, and several explicit constructions meet the bound (see [3, 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, $\dim(\mathcal{L}(G))$ depends on the divisor G itself, and not only on its degree, as is the case when $\deg(G) > 2g$. For some special algebraic function fields the vector space $\mathcal{L}(G)$ (and therefore also its dimension) is known in full. We talk more about this below.

3.2 Concatenating AG codes with Hadamard

We concatenate 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 $[2^{k_2}, k_2 = \log(q), \frac{1}{2}]_2$ Hadamard code

gives an $[n = 2n_1q, k = k_1(1 + \log q)]_2$ code that is $\epsilon = \frac{n_1-d}{n_1}$ balanced, because non-zero symbols in the outer code expand by the concatenation to perfectly balanced blocks.[1]

Using Reed-Solomon code as the outer code, one gets an $[n = 2q^2, k = k_1(\log q + 1)]_2$ code that is $\epsilon = \frac{k_1}{q}$ balanced. Rearranging parameters, this gives an $[n, k]_2$ ϵ-balanced code with $n = O((\frac{k}{\epsilon \log(\frac{k}{\epsilon})})^2)$. This is one of the constructions in [1].

Taking the outer code to be an AG code $C(Y; G)$ over $\mathbb{F}_q$, with $\deg(G) > 2g$ and optimal length $g\sqrt{q}$, one gets an $[n = g\sqrt{q}, k = \deg(G) + 1 - g]_2$ code that is $\epsilon = \frac{\deg(G)}{g\sqrt{q}}$-biased. Doing the calculation one sees that $n = O(\frac{k}{\epsilon^3 \log(1/\epsilon)})$. As these are the best AG codes possible for the case $\deg(G) > 2g$, no improvement is possible here unless we consider low-degree divisors G.

So we now turn our attention to the case where $\deg(G) \leq 2g - 1$. In this case $\dim \mathcal{L}(G)$ 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, $\dim \mathcal{L}(rQ)$ is either equal to $\dim \mathcal{L}((r-1)Q)$ or to $\dim \mathcal{L}((r-1)Q) + 1$. In the former case r is said to be a *gap number* of Q. Weierstrass Gap Theorem [6, 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 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 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 $\dim \mathcal{L}(rQ)$ 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 $\mathbb{F}_q$ 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 points. 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 $r = \epsilon p^3$. The Weierstrass semigroup of G is generated by p and $p + 1$, and a basis for $\mathcal{L}(G) = \mathcal{L}(rQ_\infty)$ is

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

[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.

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 will be chosen such that the constraint $j \leq p-1$ will be nullified. Therefore, both use evaluations of low degree bivariate polynomials over the same set of p^3 points.[2]

Theorem 10. *For every k and every ϵ such that $\frac{\epsilon}{\sqrt{\log(1/\epsilon)}} \leq \frac{1}{\sqrt{k}}$, there exists an explicit $[n, \Omega(k)]_2$ code that is ϵ-balanced, with $n = O\left(\frac{k}{\epsilon^2 \log(1/\epsilon)}\right)^{5/4}$.*

Proof: For a given k and ϵ, let

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

be a power of two. It can verified that $\frac{1}{p^3} \leq \epsilon \leq \frac{1}{p}$ and so $\log(1/\epsilon) = \Theta(\log(p))$.

Let $r = \epsilon p^3$ 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 $r \leq p^2$, $\dim \mathcal{L}(rQ_\infty) \geq (\frac{r}{2(p+1)})^2 = \Omega(\epsilon^2 p^4) = \Omega(\frac{k}{\log(p)})$. By Claim 7, the Goppa code that is obtained from the triplet (F, Y, G) is a

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

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

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

and therefore $n = p^5 = O((\frac{k}{\epsilon^2 \log(\frac{1}{\epsilon})})^{5/4})$ as desired. ∎

4 The approach limits

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 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 Hadamard cannot give ϵ-balanced codes of dimension k and length $n = O(\frac{k}{\epsilon^3 \log(1/\epsilon)})$.

Our construction shows 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 dimension k, ϵ-balanced code of length $n = O((\frac{k}{\epsilon^2 \log(1/\epsilon)})^{5/4})$.

[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.

It is therefore natural to ask what are the limits of our approach. It turns out this boils down to the question whether there are function fields with many rational points (compared to the genus) and with low-degree divisors (of degree much smaller than the genus) of high dimension.

The first argument we present shows any divisor with non-trivial dimension must have degree at least $N(F)/(q+1)$. The argument was shown to us by Henning Stichtenoth [7].

Lemma 11. *Let $F/\mathbb{F}_q$ be a function field and $G \in \mathcal{D}_F$ a divisor with $\dim(\mathcal{L}(G)) > 1$. Then $N(F) \leq \deg(G) \cdot (q+1)$.*

Proof: As $\dim(\mathcal{L}(G)) > 1$, 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 [6, Thm I.4.11], $\deg(x)_\infty = [F : \mathbb{F}_q(x)]$. On the other hand, 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, [6, 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 12. *Lemma 11 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 $\dim(\mathcal{L}(G)) > \ell$ then $N(F) \leq \frac{\deg(G)\cdot(q+1)}{f(\ell)}$ for some function f that goes to infinity with ℓ ?*

Lemma 11 shows the best we can hope for is an algebraic function field F with $N(F) \sim g\sqrt{q}$, and a divisor G of low degree $\frac{g}{\sqrt{q}} \leq \deg(G) \ll g$ and high dimension ℓ. But do such divisors exist with parameters better than those of Section 3?

Following our work, Voloch [9] showed, based on Castelnuovo bound, that:

Theorem 13 ([9, based on 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 $\ell+2$. Let $m = d \operatorname{div} \ell$ and $r = d \bmod \ell$. Then $g \leq m(m-1)\ell + m(2r+1)$, and, in particular, $g \leq m(m+1)\ell$.*

Voloch used Theorem 13 to prove our codes cannot get close to the Gilbert-Varshamov bound. Doing the analysis more carefully leads to:

Theorem 14. *Any ϵ-balanced $[n,k]_2$ code that is constructed by concatenating an AG code with the Hadamard code, must have*

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

Proof: Assume the code Z is obtained by concatenating the AG code $C(Y;G)$ specified by the triplet $(F/\mathbb{F}_q, Y, G)$ with the Hadamard code. Thus, $Y \subseteq \mathcal{P}_F$ is some subset of degree 1 places and $G \in \mathcal{D}_F$ with no support over any place in Y. The AG code $C(Y;G)$ is a $[|Y|, \ell, |Y| - d]_q$ code, where $\ell = \dim(G)$ and $d = \deg(G)$. The code Z is therefore a $[n = |Y| \cdot q, k = \ell\log(q)]_2$ code which is $\epsilon = \frac{d}{|Y|}$-balanced. Let $m = d \operatorname{div} \ell \geq 1$.

Assume $n \leq \frac{k^2}{3\epsilon^2\log^2(q)}$. A simple calculation shows this implies $|Y| \geq 3m^2q$, and, in particular, $N(F) \geq |Y| \geq 2(q+1)$. We already know by Theorem 9 that $N(F) \leq (q+$

$1) + 2g_F\sqrt{q}$, where g_F is the genus of F. Thus, $N(F) \leq 4g_F\sqrt{q}$. By Theorem 13, $N(F) \leq 8m^2\ell\sqrt{q}$. A simple calculation shows this implies $n \geq \frac{k\sqrt{q}}{8\epsilon^2 \log q}$.

If $q > \frac{k}{\epsilon^3}$ we are trivially done (because $n \geq q$), and so we can assume $\log(q) = O(\log(\frac{k}{\epsilon}))$. Thus, $n = O(\frac{k^2}{\epsilon^2 \log^2(\frac{k}{\epsilon})})$ implies $n = \Omega(\frac{k\sqrt{q}}{\epsilon^2 \log(\frac{k}{\epsilon})})$. To finish the argument notice that by Lemma 11, $N(F) \leq d(q+1)$. This implies $\frac{d}{\epsilon} = |Y| \leq d(q+1)$ and $\epsilon \geq \frac{1}{q+1}$, hence, $n = \Omega(\frac{k}{\epsilon^{2.5} \log(\frac{k}{\epsilon})})$. ∎

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 $n = \tilde{O}(\frac{k}{\epsilon^{2.5}})$, where the $\tilde{O}$ notation is used to hide poly-logarithmic factors in q (or equivalently in k and ϵ). We know the following:

- $n = \tilde{O}(\frac{k^2}{\epsilon^2})$ implies $N(F) = \tilde{\Omega}(qm^2)$. (We already saw that in the proof of Theorem 14.)
- A similar calculation shows $n = \tilde{O}(\frac{k}{\epsilon^{2.5}})$ 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 $N(F) \leq d(q+1)$), and,
- $N(F) = \tilde{O}(q^{1/2}m^2\ell)$ (since we can assume $N(F) \geq 2(q+1)$, as explained in the proof of Theorem 14).

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

Open Problem 15. *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 $\dim(G) \geq \tilde{O}(\frac{d}{\sqrt{q}})$.*

One might suspect such a high dimension, low-degree divisor does not exist. However, Theorem 13 and Lemma 11 are not strong enough to disprove it. We remark that the lower bound could be improved, if Lemma 11 could be strengthened to use the high-dimension of G, as suggested in Remark 12.

References

[1] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple constructions of almost k–wise independent random variables. *Random Structures and Algorithms*, 3(3):289–303, 1992.

[2] W. Fulton. *Algebraic Curves*. Third edition, 2008.

[3] A. Garcia and Eds. Stichtenoth, H. *Topics in Geometry, Coding Theory and Cryptography (Algebra and Applications)*. Springer-Verlag, 2006.

[4] S. Lang. *Algebra*. Springer, revised third edition, 2002.

[5] J. Naor and M. Naor. Small–bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993.

[6] H. Stichtenoth. *Algebraic function fields and codes*. Springer Verlag, 1993.

[7] H. Stichtenoth. Private communication, 2009.

[8] M.A. Tsfasman, S.G. Vladutx, and T. Zink. Modular curves, Shimura curves, and Goppa codes, better than Varshamov-Gilbert bound. *Mathematische Nachrichten*, 109(1), 1982.

[9] F. Voloch. Special divisors of large dimension on curves with many points over finite fields. To appear in Portugaliae Mathematica, 2009.