6.897 Algorithmic Introduction to Coding Theory October 3, 2001

Lecture 8

Lecturer: Madhu Sudan *Scribe: Shien Jin Ong*

Today we will discuss upper bounds on the rate of any family code, given a lower bound on its relative distance. Specifically we will present the Plotkin bound and the Elias-Bassalygo bounds. En route we will also encounter a different bound, called the Johnson bound. The unifying theme for the lecture is that of finding upper bounds on the rate of codes by geometric arguments. In particular, we will embed Hamming space into Euclidean space and use the embeddings, in combination with geometric facts, to derive our proofs.

These notes include extensions of various proofs to q-ary cases - the lecture only covered the binary case. Throughout this lecture, we will use R to denote the rate of some (unspecified) family of codes and δ to denote the relative distance of the same family.

1 Embedding Hamming spaces in Euclidean spaces

To motivate our first bound, let us recall our current state of knowledge, for binary codes. On the one hand we have the Singleton and Hamming upper bounds on codes, with the latter dominating the former and showing $R \leq 1 - H(\delta/2)$. The best existence result, the Gibert-Varshamov (GV) bound, shows there exists a family with $R \geq 1 - H(\delta)$. For any $\delta > 0$, the bounds are far away from each other. However to get a qualitative sense of the gap, consider the largest distance that these bounds suggest are feasible for codes of positive rate. The Hamming bound rules out a relative distance of 1 for codes of positive rate. One the other hand, the GV bound only finds codes of positive rate with relative distance close to $\frac{1}{2}$. Clearly there is a qualitative gap here — and we address this gap first.

It is reasonably easy to guess which of these bounds is closer to the truth. Over a binary alphabet, random words have a relative distance of $\frac{1}{2}$ from each other and it seems quite impossible to construct codes with better distance. The Hamming bound on the other seems quite weak around these parts. We just need a way to formalize our intuition, and we will do so geometrically, by embedding the binary Hamming space into Euclidean space. We develop the embedding below.

Definition 1 (Embedding) *The embedding function* Embed : $\{0,1\} \to \mathbb{R}$*, mapping bits to the reals is given by* $\text{Embed}(0) = +1$ *and* $\text{Embed}(1) = -1$*. For $n \geq 1$, the n-dimensional embedding function extends the embedding above, with* Embed : $\{0,1\}^n \to \mathbb{R}^n$ *being given by*

$$\text{Embed}(\langle b_1, \ldots, b_n \rangle) = \langle \text{Embed}(b_1), \ldots, \text{Embed}(b_n) \rangle.$$

The property of this embedding is that Hamming distances are preserved as Euclidean distances, or in inner products. We recall some familiar definitions for vector spaces over the reals.

Definition 2 *For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, the inner product between $\mathbf{x}$ and $\mathbf{y}$, denoted $\langle \mathbf{x}, \mathbf{y} \rangle$, equals $\sum_{i=1}^n x_i y_i$. The norm of a vector $\mathbf{x}$, denoted $\|\mathbf{x}\|$, is $\sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$, The Euclidean distance between $\mathbf{x}$ and $\mathbf{y}$, is simply the norm of $\mathbf{x} - \mathbf{y}$.*

The following proposition lists some of the elementary properties of our embedding from Hamming to Euclidean spaces. The proof is easily verified from the definitions and hence omitted.

Proposition 3 *For $\mathbf{b}, \mathbf{c} \in \{0,1\}^n$, the function* Embed *satisfies* $\langle \text{Embed}(\mathbf{b}), \text{Embed}(\mathbf{c}) \rangle = n - 2\Delta(\mathbf{b}, \mathbf{c})$*, where $\Delta(\cdot,\cdot)$ is the Hamming distance function. Hence, we have*

$$\|\text{Embed}(\mathbf{b})\|^2 = n, \text{ and } \|\text{Embed}(\mathbf{b}) - \text{Embed}(\mathbf{c})\|^2 = 4\Delta(\mathbf{b}, \mathbf{c}).$$

The embedding above thus allows us to transform questions about Hamming space into questions about Euclidean space. We will then appeal to our geometric intuition backed by linear algebra for proofs of coding-theoretic statements.

2 The Plotkin bound

Theorem 4 (Plotkin bound [6])

1. *An $(n, k, d)_2$ code with $d \geq \frac{n}{2}$ has at most $2n$ codewords. In other words, $k \leq \log 2n$.*
2. *If an $(n, k, d)_2$ code exists, then $k \leq n - 2d + \log(4d)$.*

Proof The first part is the harder part and the second part follows easily by using restrictions. We start with the first part.

Let $\mathbf{c}_1, \ldots, \mathbf{c}_m$ be all the codewords of the $(n, k, d)_2$ code. Let $\mathbf{x}_1, \ldots, \mathbf{x}_m$ be their embeddings, i.e., $\mathbf{x}_i = \text{Embed}(\mathbf{c}_i)$. By Proposition 3, we have that the inner product between $\mathbf{x}_i$ and $\mathbf{x}_j$, for $i \neq j$, is equal to $n - 2\Delta(\mathbf{c}_i, \mathbf{c}_j) \leq 0$ from the fact that the code has distance $d \geq n/2$. In Lemma 6 we show that in $\mathbb{R}^n$ there can be at most $2n$ vectors such that every pair has an non-positive inner product. The first part of the theorem follows.

To see the second part, let us write $n = 2d + \ell$ and suppose C is an $(n, k, d)_2$ code. Then by restricting C to the most commonly occuring pattern in the first ℓ coordinates and deleting these coordinates, we get a $(2d, k - \ell, d)_2$ code. By the first part of the theorem, we have $k - \ell \leq \log(4d)$. ■

Before stating or proving the critical Lemma 6, we state the asymptotic version of Plotkin's bound.

Corollary 5 *For any family of binary codes $\mathcal{C}$ with rate R and relative distance δ, it is the case that $R \leq 1 - 2\delta$.*

We now move on to proving Lemma 6,

3 Geometric assertions, Linear-algebraic proofs

Our first goal here is to prove a geometric fact: In n dimensions there exist at most $2n$ vectors that pairwise subtend an angle of at least $\frac{\pi}{2}$ at the origin. We start with an intuitive, inductive proof. However the proof actually uses a fair bit of intuition about Euclidean spaces that we haven't (or won't) prove. We will then give an alternate, linear-algebraic proof that only uses the fact that the norm of a vector is non-negative, and that any $n + 1$ vectors in n dimensions are linearly dependent.

Before proving the lemma, let us see why it is proving the right fact. We already know that the Hadamard code matches the Plotkin bound (or the first part of it) and so its embedding should match the lemma below tightly. But we can come up with a simpler example (geometrically the same, actually!) which shows that the lemma is tight. Take $\mathbf{x}_1, \ldots, \mathbf{x}_n$ to be the unit vectors along the coordinate axes, and let $\mathbf{x}_{n+i} = -\mathbf{x}_i$ for $i \in [n]$. This gives $2n$ non-zero vectors that are mutually at an angle of at least $\pi/2$.

Lemma 6 *If $\mathbf{x}_1, \ldots, \mathbf{x}_m \in \mathbb{R}^n$ are non-zero and satisfy $\langle \mathbf{x}_i, \mathbf{x}_j \rangle \leq 0$ for every $i \neq j \in [m]$, then $m \leq 2n$.*

Proof We prove the lemma by induction on n. Without loss of generality we may assume that the vector $\mathbf{x}_m$ is the unit vector $\langle 1, 0, \ldots, 0 \rangle$. (The fact that this assumption follows without loss of generality is intuitively obvious, but would require some work if we decided to prove it!) Write $\mathbf{x}_i = \langle \alpha_i, \mathbf{y}_i \rangle$, where $\mathbf{y}_i \in \mathbb{R}^{n-1}$. Since we know that all other vectors have a non-positive inner product with $\mathbf{x}_m$, we find that $\alpha_i = \langle \mathbf{x}_i, \mathbf{x}_m \rangle \leq 0$. It follows that for distinct $i, j \in [m-1]$, we have

$$\langle \mathbf{y}_i, \mathbf{y}_j \rangle = \langle \mathbf{x}_i, \mathbf{x}_j \rangle - \alpha_i \alpha_j \leq \langle \mathbf{x}_i, \mathbf{x}_j \rangle \leq 0.$$

So we have $2m - 1$ vectors in $n - 1$ dimensions, as so we should be able to apply the inductive hypothesis — right? Well, that would be too strong and would yield $m \leq n$ — a bit better than some of the examples we have. Why does this happen. Well, the inductive hypothesis assumes $\mathbf{y}_i$'s are non-zero, and we didn't prove this yet. So to complete the lemma, we note that at most one of the $\mathbf{y}_i$'s may be zero. (If two, say $\mathbf{y}_1$ and $\mathbf{y}_2$ are zero, then their inner product would be positive!) We delete the zero vector and then we are left with $m-2$ non-zero vectors in $n-1$ dimensions with a pairwise non-positive inner product. This allows us to apply induction and the lemma is proved. ■

What if we were actually given that pairwise the vectors have a stricly negative inner product of say $-\alpha$? Could we improve the bound? The reader may try modifying the proof above to show that it this case the number of vectors is at most $1+\frac{1}{\alpha}$, a bound independent of the number of dimensions. But the proof also starts to get more tedious. Motivated by such tasks, we now state a stronger lemma and give a self-contained proof. In particular, the proof is easier to verify (though possibly harder to conceive).

Lemma 7

1. *If α is a positive number and $\mathbf{x}_1,\ldots,\mathbf{x}_m \in \mathbb{R}^n$ are unit vectors i.e., $\|\mathbf{x}_i\| = 1$, that satisfy $\langle \mathbf{x}_i, \mathbf{x}_j\rangle \leq -\alpha$, for every distinct pair $i,j \in [m]$, then $m \leq 1+\frac{1}{\alpha}$.*
2. *If $\mathbf{y},\mathbf{x}_1,\ldots,\mathbf{x}_m \in \mathbb{R}^n$ satisfy $\langle \mathbf{x}_i, \mathbf{x}_j\rangle \leq 0$ for distinct i,j, while $\langle \mathbf{y}, \mathbf{x}_i\rangle > 0$, then $m \leq n$.*
3. *If $\mathbf{x}_1,\ldots,\mathbf{x}_m \in \mathbb{R}^n$ are non-zero and satisfy $\langle \mathbf{x}_i, \mathbf{x}_j\rangle \leq 0$ for every $i \neq j \in [m]$, then $m \leq 2n$.*

Proof We prove the three parts in order.

1. Let $\mathbf{z} = \mathbf{x}_1 + \cdots + \mathbf{x}_m$. On the one hand, we have $\langle \mathbf{z}, \mathbf{z}\rangle \geq 0$. On the other, we have

$$\begin{aligned}\langle \mathbf{z}, \mathbf{z}\rangle &= \sum_{i=1}^{m} \langle \mathbf{x}_i, \mathbf{x}_i\rangle + \sum_{i\neq j\in[m]} \langle \mathbf{x}_i, \mathbf{x}_j\rangle \\ &\leq m\cdot 1 + m(m-1)\cdot(-\alpha) \\ &= m\cdot(1-(m-1)\alpha).\end{aligned}$$

   Putting the two together, we have $1-(m-1)\alpha \geq 0$, implying $m \leq 1+\frac{1}{\alpha}$.

2. For this part, assume for the sake of contradiction that $m \geq n+1$. Then there must exist a linearly dependent set of vectors among the $\mathbf{x}_i$'s. Specifically there exist disjoint sets $S,T \subseteq [m]$ and positive λ_i, for $i \in S\cup T$, such that $\sum_{i\in S}\lambda_i\mathbf{x}_i = \sum_{j\in T}\lambda_j\mathbf{x}_j$. It is not necessary that both S and T be non-empty, but at least one is non-empty. Assume without loss of generality that S is non-empty. Let $\mathbf{z} = \sum_{i\in S}\lambda_i\mathbf{x}_i = \sum_{j\in T}\lambda_j\mathbf{x}_j$. Our analysis divides into two cases depending on whether $\mathbf{z} = \mathbf{0}$ or not.

   Case: $\mathbf{z} \neq \mathbf{0}$: Here we obtain the following contradication:

$$\begin{aligned}0 &< \langle \mathbf{z}, \mathbf{z}\rangle \\ &= \langle \sum_{i\in S}\lambda_i\mathbf{x}_i, \sum_{j\in T}\lambda_j\mathbf{x}_j\rangle \\ &= \sum_{i\in S}\sum_{j\in T}\lambda_i\lambda_j\langle \mathbf{x}_i, \mathbf{x}_j\rangle \\ &\leq 0,\end{aligned}$$

   where the last inequality uses the fact that S and T are disjoint and so $\langle \mathbf{x}_i, \mathbf{x}_j\rangle \leq 0$ for every $i \in S$ and $j \in T$.

Case: $\mathbf{z} = \mathbf{0}$: Here we use the existence of the vector $\mathbf{y}$ and obtain a contradiction as follows:

$$\begin{aligned} 0 &= \langle \mathbf{y}, \mathbf{0} \rangle \\ &= \langle \mathbf{y}, \mathbf{z} \rangle \\ &= \langle \mathbf{y}, \sum_{i \in S} \lambda_i \mathbf{x}_i \rangle \\ &= \sum_{i \in S} \lambda_i \langle \mathbf{y}, \mathbf{x}_i \rangle \\ &> 0. \end{aligned}$$

The last inequality is strict since $S \neq \emptyset$, $\lambda_i > 0$ and $\langle \mathbf{y}, \mathbf{x}_i \rangle > 0$.

3. Finally we move to Part (3) which is exactly the same statement as that of Lemma 6. Our new proof follows easily from Part (2) of the current lemma. Pick a vector $\mathbf{y}$ in general position, i.e., so that $\langle \mathbf{y}, \mathbf{x}_i \rangle \neq 0$ for any $i \in [m]$. At least half the vectors $\mathbf{x}_i$ must have a positive inner product with either $\mathbf{y}$ or $-y$. Assume without loss of generality that $\mathbf{x}_1, \ldots, \mathbf{x}_{\lceil m/2 \rceil}$ have a positive inner product with $\mathbf{y}$. Applying Part(2) to these vectors and $\mathbf{y}$, we get $\lceil m/2 \rceil \leq n$.

■

Exercise: (A) Give an example showing the result in Part (1) of Lemma 7 is tight. (B) Interpret this part as a coding-theoretic bound. (C) Give codes that make this interpretation tight.

The exercise above gives the natural motivation for Part (1) of the lemma. Part (3) was already motivated by the Plotkin bound. What motivates us to study Part (2). On the one hand it provides a simple proof of Part (3). But it actually turns out to be even more important on its own and we will use it several times in the rest of these notes.

4 The Elias-Bassalygo bound

We now return to the task of bounding the rate of a code, given its relative distance. The current picture of the upper bounds involves two incomparable bounds: the Hamming bound is stronger for smaller δ and the Plotkin bound is stronger for larger δ. Our next bound unifies the two techniques and thus gets a bound which is always stronger, though the bound is very close to the Hamming bound for small δ.

To motivate this bound, we introduce a new notion of error-correction. Later we will refer to this notion as that of "list-decoding". Currently, we will use a terminology that is more reminiscent of the notion of "t-error-correcting codes" of Hamming.

Definition 8 ((t, ℓ)-error-correcting code) *A code $C \subseteq \Sigma^n$ is a (t, ℓ) error correcting code if for every received word $\mathbf{y} \in \Sigma^n$, the ball of radius t around $\mathbf{y}$, $B(\mathbf{y}, t)$, contains at most ℓ codewords of C.*

For a code C and integer ℓ, we refer to the largest t for which C is a (t, ℓ)-error-correcting code to be the list of ℓ error-correcting radius of C.

Recall that Hamming's notion of a t-error-correcting code becomes a $(t, 1)$-error-correcting code in this new definition. Let us take a peek back at Hamming's proof of the Hamming bound for binary codes. The crux of the proof was that the balls $B(\mathbf{c}, t)$ around the codewords of a t-error-correcting code are disjoint. Thus if the code has 2^k codewords we get $2^k \mathrm{Vol}(t, n) \leq 2^n$, where $\mathrm{Vol}(t, n) = |B(\mathbf{c}, t)|$. Note that we didn't say exactly this when we proved the Hamming bound. Instead we considered balls of radius $(d-1)/2$ around codewords, where d was the minimum distance, and implicitly used the fact that such a code is a $(d-1)/2$-error-correcting code. But when we generalize the Hamming bound it will be better to explicit with notion of t-error-correcting codes.

Proposition 9 *Suppose a $(n,k,d)_2$ code C is a (t,ℓ)-error-correcting code. Then $2^k \mathrm{Vol}(t,n) \le \ell \cdot 2^n$.*

Proof The proposition follows easily. If we consider the balls of radius t around the codewords, then any word in $\{0,1\}^n$ is considered at most ℓ times. Thus the sum of the volumes of these balls is at most $\ell \cdot 2^n$. ■

Of course, the proposition does not immediately translate into new asymptotic relationships between rate and relative minimum distance. To get such relationships we have to relate the minimum distance of a code to its list of ℓ-error-correcting radius for non-trivial values of ℓ. Any $\ell > 2$, but less than $2^{\epsilon n}$ would be of interest. We will study such a bound next. Such bounds are closely related to bounds studied by S. Johnson [4, 5] and are termed the Johnson bounds.

Theorem 10 (Johnson bound [4]) *Every $(n,k,\delta n)_2$ code is also a $(\tau n - 1, n)$-error-correcting code for $\tau = \frac{1}{2} \cdot (1 - \sqrt{1-2\delta})$.*

Proof As usual we turn the problem into a geometric one by using the embedding function Embed. Let $\mathbf{c}_1, \ldots, \mathbf{c}_m$ be codewords of a code of minimum distance $d = \delta n$ that are within a Hamming ball of radius $t = \tau n - 1$ from a received vector $\mathbf{b}$. We wish to show $m \le n$.

We will embed the vectors into Euclidean space, scaling them by a factor of $1/\sqrt{n}$ to get vectors of unit norm. For $i \in [m]$, let $\mathbf{x}_i = \frac{1}{\sqrt{n}}\mathrm{Embed}(\mathbf{c}_i)$ and let $\mathbf{y} = \frac{1}{\sqrt{n}}\mathrm{Embed}(\mathbf{b})$. By the properties of the embedding function (Proposition 3), we get: (1) $\|\mathbf{x}_i\| = \|\mathbf{y}\| = 1$. (2) $\langle \mathbf{x}_i, \mathbf{x}_j \rangle \le 1 - 2\delta$, if $i \ne j$. (3) $\langle \mathbf{y}, \mathbf{x}_i \rangle > 1 - 2\tau$. In other words, we have a collection of unit vectors $\mathbf{x}_i$ whose pairwise inner product is small, but which have a common, large inner product with $\mathbf{y}$. Notice the syntactic similarity to Part (2) of Lemma 7. In fact we will reduce our problem to exactly this case. How? We will just shift our origin to a new vector $\mathbf{v}$ so that from this vector, the vectors $\mathbf{x}_i$ mutually subtend an angle of at least $\pi/2$. And how do we find such a vector $\mathbf{v}$? Well the most natural idea is to shift the origin closer to the scene of action — namely, towards $\mathbf{y}$. Specifically, we will move to some point $\alpha\mathbf{y}$ and inspect our world from there. The following claim asserts that we will see what we hope to see.

Claim 11 *There exists an α such that for every $i \ne j \in [m]$, $\langle \mathbf{x}_i - \alpha\mathbf{y}, \mathbf{x}_j - \alpha\mathbf{y} \rangle \ge 0$, while for every i, $\langle \mathbf{x}_i - \alpha\mathbf{y}, \mathbf{y} - \alpha\mathbf{y} \rangle > 0$.*

Proof We will not specify α yet only that it will lie in the interval $0 \le \alpha < 1$. For such α, note that

$$\langle \mathbf{x}_i - \alpha\mathbf{y}, \mathbf{x}_j - \alpha\mathbf{y} \rangle \le 1 - 2\delta - 2\alpha(1-2\tau) + \alpha^2 = (1-\alpha)^2 + 4\alpha\tau - 2\delta.$$

The right-hand side is minimized at $\alpha = 1 - 2\tau$. For this setting, the RHS above equals $4\tau - 4\tau^2 - 2\delta$. Recall we set $\tau = \frac{1}{2}(1 - \sqrt{1-2\delta})$, and so we have $(1-2\tau)^2 = 1 - 2\delta$, which in turn implies $4\tau - 4\tau^2 - 2\delta = 0$. We conclude that for this setting $\langle \mathbf{x}_i - \alpha\mathbf{y}, \mathbf{x}_j - \alpha\mathbf{y} \rangle \ge 0$, as desired.

To conclude, we note that for the same setting $\alpha = 1 - 2\tau$, we have

$$\langle \mathbf{x}_i - \alpha\mathbf{y}, (1-\alpha)\mathbf{y} \rangle > (1-\alpha)(1-2\tau) - (\alpha)(1-\alpha) = 0,$$

which yields the other part of the claim. ■

We are now in a position to apply Lemma 7, Part (2), to the vectors $\{\mathbf{x}_i - \alpha\mathbf{y}\}_i$ and $\mathbf{y} - \alpha\mathbf{y}$ to conclude that $m \le n$. This concludes the proof of the theorem. ■

Combining Proposition 9 with Theorem 10 gives us the Elias-Bassalygo upper bound on the rate of a family of codes with relative distance δ.

Theorem 12 (Elias-Bassalygo bound [2, 7]) *If $\mathcal{C}$ is an infinite family of binary codes with rate R and relative distance δ, then $R \le 1 - H(\frac{1}{2} \cdot (1 - \sqrt{1-2\delta}))$.*

Proof The theorem follows essentially from Proposition 9 and Theorem 10. The missing ingredients are dry exercises showing that one can pick n large enough so as to overcome all problems posed by identities which hold only asymptotically. (The volume of Hamming balls is not exactly related to the binary entropy function; the Johnson bound only lower bounds the (t, ℓ)-error-correcting radius of codes when ℓ is a growing function of n. And the bound only allows a list-decoding radius of $\tau n - 1$ and not τn.) We'll spare the reader the details. ■

Before concluding the section, let us digress briefly to understand when and why the Elias-Bassalygo bound is better. To do so, let us recall two of the previous bounds that we have worked with. The Hamming upper bound says $R \leq 1 - H(\delta/2)$, while the GV lower bound says $R \geq 1 - H(\delta)$. The Elias-Bassalygo bound shows $R \leq 1 - H(\tau(\delta))$, for $\tau(\delta) = \frac{1}{2} \cdot (1 - \sqrt{1 - 2\delta})$. First note that $\delta/2 \leq \tau(\delta) \leq \delta$ and H_2 is monotone decreasing, and so the Elias-Bassalygo bound is always between the Hamming bound and the GV bound. Further if $\delta > 0$, then $\tau(\delta) > \delta/2$ and so the Elias-Bassalygo bound is strictly better than the Hamming bound. However if δ is close to zero then $\delta/2$ is a very good approximation to $\tau(\delta)$ and so for small values of δ the Elias-Bassalygo bound is not much better than the Hamming bound. However for large values of δ, $\tau(\delta)$ starts to get closer to δ. In particular, if $\delta = \frac{1}{2} - \epsilon$, then $\tau(\delta) = \frac{1}{2} - \sqrt{\epsilon/2}$ and so τ approaches $\frac{1}{2}$ as δ approaches $\frac{1}{2}$. So as $\delta \to \frac{1}{2}$, the Elias-Bassalygo bound really starts to get better and approaches the GV bound!

What else could we hope for? While the Elias-Bassalygo bound gives us the right bound for $\delta = \frac{1}{2}$, it does not quite have the right growth around this point. In particular, the GV bound shows that one can find codes of rate $O(\epsilon^2)$ and relative distance $\frac{1}{2} - \epsilon$, as $\epsilon \to 0$. The Elias-Bassalygo bound only rules out codes of rate $\Omega(\epsilon)$ at this distance. Which bound is closer to the truth? Turns out the GV bound is correct here, and the E-B bound is too weak. In the next lecture we will describe a different upper bound, called the Linear Programming (LP) bound, which ends up showing the tightness of the GV bound.

5 q-ary bounds

We now extend the results of earlier section to codes over general alphabets. We start with q-ary embeddings. The definition does not extend the previous definition, but in fact, gives an alternate embedding which works as well.

Fix an arbitrary bijection $\mathrm{ind} : \mathbb{F}_q \to [q]$ be any bijection between $\mathbb{F}_q$ and $[q]$. For $1 \leq i \leq n$, Let $\mathbf{e}_{i,n}$ be the unit vector along the ith coordinate direction in $\mathbb{R}^n$. We now define our q-ary embeddings.

Definition 13 (Embedding q-ary space in Euclidean space) *The embedding function q-Embed : $\mathbb{F}_q \to \mathbb{R}^q$ is defined as follows:*

$$q\text{-}Embed(\alpha) = \mathbf{e}_{\mathrm{ind}(\alpha),q}.$$

For $n \geq 1$, the n-dimensional embedding function extends the embedding above, with q-Embed : $\mathbb{F}_q^n \to \mathbb{R}^{qn}$ being given by

$$q\text{-}Embed(\langle \alpha_1, \ldots, \alpha_n \rangle) = \langle q\text{-}Embed(\alpha_1), \ldots, q\text{-}Embed(\alpha_n) \rangle.$$

Proposition 14 *For vectors $\alpha, \beta \in \mathbb{F}_q^n$, the embedding q-Embed satisfies:*

$$\|q\text{-}Embed(\alpha)\|^2 = n, \ \langle q\text{-}Embed(\alpha), q\text{-}Embed(\beta) \rangle = n - \Delta(\alpha, \beta).$$

It will be preferrable to index our qn-dimensional space by two indices i, j with $i \in [n]$ and $j \in [q]$. Let H_i denote the hyperlane in $\mathbb{R}^{qn}$ given by $\sum_{j=1}^{q} x_{ij} = 1$. Note that the q-ary embedding lies in the affine subspace $\mathcal{H}$ given by the intersection of the hyperplanes H_i, i.e., $\mathcal{H} = \cap_{i=1}^{n} H_i$. Let $\mathbf{Q}_n \in \mathbb{R}^{qn}$ be the vector $\langle \frac{1}{q}, \ldots, \frac{1}{q} \rangle$. Then Q_n also lies in $\mathcal{H}$ and will play the role of the origin in $\mathcal{H}$. We will use the following proposition to tighten our results.

Proposition 15 *The vectors* $\{q\text{-}Embed(\mathbf{x}) - \mathbf{Q}_n | \mathbf{x} \in \mathbb{F}_q^n\}$ *lie in an* $(q-1)n$*-dimensional vector space over* $\mathbb{R}$.

We are now ready to prove some bounds. We start with the q-ary Plotkin bound.

Theorem 16 (q-ary Plotkin bound) *If C is an $(n,k,d)_q$ code, then* $k \le n - \frac{q}{q-1} \cdot d + \log_q \left(\frac{q^2}{q-1} d\right)$.

Proof It suffices to prove the theorem for $d \ge \frac{q-1}{q} n$. The remaining cases follow by the restriction argument. For the case $d \ge \frac{q}{q-1} n$, we need to show that the number of codewords is at most qn. (This bound is met by Reed-Muller codes of degree 1.) It would be slightly easier to prove a bound of $2(q-1)n$. This may satisfy mere mortals, but since we're superhuman, we'll prove the correct result.

Let C be an $(n,k,d)_q$ code with $d \ge \frac{q-1}{q} n$. Let α be the *least* commonly occurring symbol in the first coordinate among codewords of C. Let C' be the code obtained by throwing away from C all codewords that have an α in the first coordinate position. Note that $|C'| \ge \frac{q-1}{q} |C|$. Thus it will suffice to prove that $|C'| \le (q-1)n$ (to get $|C| \le qn$). We will do so below.

Let $\mathbf{c}_1, \ldots, \mathbf{c}_m$ be the codewords of C'. Let $\mathbf{x}_1, \ldots, \mathbf{x}_m \in \mathbb{R}^{qn}$ be the vectors given $\mathbf{x}_i = q\text{-Embed}(\mathbf{c}_i) - Q_n$. Let $\mathbf{y} = -(\langle q\text{-Embed}(\alpha), \mathbf{Q}_{n-1}\rangle - \mathbf{Q}_n)$. We will show below that the following are true: (1) $\langle \mathbf{x}_i, \mathbf{x}_j \rangle \le 0$, (2) $\langle \mathbf{x}_i, \mathbf{y} \rangle > 0$, and (3) $\mathbf{x}_i, \mathbf{y}$'s are contained in a $(q-1)n$ dimensional real vector space. Applying Lemma 7 to these vectors, we get that $m \le (q-1)n$, as desired. Thus it suffices to verify the three conditions above to prove the theorem. We do so below.

1. Note that

$$\begin{aligned} \langle \mathbf{x}_i, \mathbf{x}_j \rangle &= \langle q\text{-Embed}(\mathbf{c}_i) - \mathbf{Q}_n, q\text{-Embed}(\mathbf{c}_j) - \mathbf{Q}_n \rangle \\ &\le n - \Delta(\mathbf{c}_i, \mathbf{c}_j) - \frac{n}{q} \\ &\le 0. \end{aligned}$$

2. We need to show $\langle \mathbf{x}_i, \mathbf{y} \rangle > 0$. Since $\mathbf{y}$ is zero on all but the first q coordinates, it suffices to consider the contribution to the inner product from the first q terms. Let the first coordinate of $\mathbf{c}_i = \beta$. Then $\langle \mathbf{x}_i, \mathbf{y} \rangle = \langle q\text{-Embed}(\beta) - \mathbf{Q}_1, -q\text{-Embed}(\alpha) + \mathbf{Q}_1 \rangle$. Since $\alpha \ne \beta$, the first inner product is zero, while the others are $\frac{1}{q}$. Summing them all up, we get $\langle \mathbf{x}_i, \mathbf{y} \rangle = \frac{1}{q} > 0$.

3. The final part follows from Proposition 14 and the fact that $\mathbf{y}$ also lies on the intersection of hyperplanes in which each of the n blocks of q coordinates sum to zero.

■

We move on to the Johnson bound for q-ary codes.

Theorem 17 (q-ary Johnson bound) *Every $(n,k,\delta n)_q$ code is also a $(\tau n - 1, (q-1)n)$-error-correcting code for* $\tau = \frac{q-1}{q} \cdot (1 - \sqrt{1 - \frac{q}{q-1}\delta})$.

Proof Let $d = \delta n$ and $t = \tau n - 1$. Let C be an $(n,k,d) - q$ code. For a received vector $\mathbf{b}$, let $\mathbf{c}_1, \ldots, \mathbf{c}_m$ be codewords of C within a Hamming distance of t from $\mathbf{b}$. Define $\mathbf{x}_i = q\text{-Embed}(\mathbf{c}_i) - \mathbf{Q}_n$, for $i \in [m]$. Define $\mathbf{y} = q\text{-Embed}(\mathbf{b}) - \mathbf{Q}_n$. The vectors $\mathbf{x}_i$, for $i \in [m]$, are contained in a $(q-1)n$ dimensional subspace of $\mathbb{R}^{qn}$ with the property that their pairwise inner product is small, while each has a large inner product with some fixed vector $\mathbf{v}$. As in the proof of Theorem 10, we can conclude that under the setting $\tau = \frac{q-1}{q} \cdot (1 - \sqrt{1 - \frac{q}{q-1}\delta})$, we can find an α such that the vectors $\{\mathbf{x}_i - \alpha \mathbf{v}\}_i$, have a pairwise non-positive inner product, while their inner product with $\mathbf{v}$ is positive. Applying Part (2) of Lemma 7 (our favorite workhorse) we get $m \le (q-1)n$. ■

The q-ary Elias-Bassalygo bound is now straightforward. We state it for completeness.

Theorem 18 (q-ary Elias-Bassalygo bound) *If $\mathcal{C}$ is a family of q-ary codes with rate R and relative distance δ, then*

$$R \leq 1 - H_q\left(\frac{q-1}{q}\left(1 - \sqrt{1 - \frac{q}{q-1}\delta}\right)\right).$$

Bibliographic Notes

The Plotkin bound was shown in [6], and the Johnson bound in [4, 5]. The Elias-Bassalygo bound was discovered independently by P. Elias and L. Bassalygo. Elias seemingly discovered the bound in the 1950s but never published his result — it just got integrated into the folklore of coding theory in the US. The first journal paper to mention Elias's proof seems to be a paper by Shannon, Gallager, and Berlekamp [7] in 1967. In the meanwhile, L.A. Bassalygo discovered the same bound in 1965 [2]. The Johnson bounds are from some intermediate period (between Elias's observation, and Bassalygo's publication). The proofs of the Johnson bound in this notes are not from the original papers, but rather from more recent work. The proofs over the binary alphabet are from Agrell, Vardy, and Zeger [1]. The q-ary version is from [3].

References

[1] Erik Agrell, Alexander Vardy, and Kenneth Zeger. Upper bounds for constant-weight codes. *IEEE Transactions on Information Theory*, 46:2373–2395, 2000.

[2] L.A. Bassalygo. New upper boundes for error-correcting codes. *Problems of Information Transmission*, 1(1):32–35, 1965.

[3] Venkatesan Guruswami and Madhu Sudan. Extensions to the Johnson bound. *Manuscript*, February 2001.

[4] Selmer M. Johnson. A new upper bound for error-correcting codes. *IEEE Transactions on Information Theory*, 8:203–207, 1962.

[5] Selmer M. Johnson. Improved asymptotic bounds for error-correcting codes. *IEEE Transactions on Information Theory*, 9:198–205, 1963.

[6] M. Plotkin. Binary codes with specified minimum distance. *IRE Transactions on Information Theory*, 6:445–450, 1960.

[7] Claude E. Shannon, Robert G. Gallager, and Elwyn R. Berlekamp. Lower bounds to error probability for coding on discrete memoryless channels. *Information and Control*, 10:65–103 (Part I), 522–552 (Part II), 1967.