Improved Nearly-MDS Expander Codes

Ron M. Roth and Vitaly Skachek

***Abstract*— A construction of expander codes is presented with the following three properties: (i) the codes lie close to the Singleton bound, (ii) they can be encoded in time complexity that is linear in their code length, and (iii) they have a linear-time bounded-distance decoder. By using a version of the decoder that corrects also erasures, the codes can replace MDS outer codes in concatenated constructions, thus resulting in linear-time encodable and decodable codes that approach the Zyablov bound or the capacity of memoryless channels. The presented construction improves on an earlier result by Guruswami and Indyk in that any rate and relative minimum distance that lies below the Singleton bound is attainable for a significantly smaller alphabet size.**

Keywords: Concatenated codes, Expander codes, Graph codes, Iterative decoding, Linear-time decoding, Linear-time encoding, MDS codes.

I. Introduction

In this work, we consider a family of codes that are based on expander graphs. The notion of graph codes was introduced by Tanner in [19]. Later, the explicit constructions of Ramanujan expander graphs due to Lubotsky, Philips, and Sarnak [8, Chapter 4], [13] and Margulis [15], were used by Alon *et al.* in [1] as building blocks to obtain new polynomial-time constructions of asymptotically good codes in the low-rate range (by "asymptotically good codes" we mean codes whose rate and relative minimum distance are both bounded away from zero). Expander graphs were used then by Sipser and Spielman in [16] to present polynomial-time constructions of asymptotically good codes that can be decoded in time complexity which is linear in the code length. By combining ideas from [1] and [16], Spielman provided in [18] an asymptotically good construction where both the decoding and encoding time complexities were linear in the code length.

While the linear-time decoder of the Sipser-Spielman construction was guaranteed to correct a number of errors that is a positive fraction of the code length, that fraction was significantly smaller than what one could attain by bounded-distance decoding—namely, decoding up to half the minimum distance of the code. The guaranteed fraction of linear-time correctable errors was substantially improved by Zémor in [20]. In his analysis, Zémor considered the special (yet abundant) case of the Sipser-Spielman construction where the underlying Ramanujan graph is bipartite, and presented a linear-time iterative decoder where the correctable fraction was $1/4$ of the relative minimum distance of the code. An additional improvement by a factor of two, which brought the (linear-time correctable) fraction to be essentially equal to that of bounded-distance decoding, was then achieved by the authors of this paper in [17], where the iterative decoder of Zémor was enhanced through a technique akin to generalized minimum distance (GMD) decoding [10], [11].

In [12], Guruswami and Indyk used Zémor's construction as a building block and combined it with methods from [1], [3], and [4] to suggest a code construction with the following three properties:

(P1) The construction is nearly-MDS: it yields for every designed rate $R \in (0,1]$ and sufficiently small $\epsilon > 0$ an infinite family of codes of rate at least R over an alphabet of size

$$2^{O\left((\log(1/\epsilon))/(R\epsilon^4)\right)} , \tag{1}$$

and the relative minimum distance of the codes is greater than

$$1 - R - \epsilon .$$

(P2) The construction is linear-time encodable, and the time complexity per symbol is $\text{POLY}(1/\epsilon)$ (i.e., this complexity grows polynomially with $1/\epsilon$).

(P3) The construction has a linear-time decoder which is essentially a bounded-distance decoder: the correctable number of errors is at least a fraction $(1-R-\epsilon)/2$ of the code length. The time complexity per symbol of the decoder is also $\text{POLY}(1/\epsilon)$.

In fact, the decoder described by Guruswami and Indyk in [12] is more general in that it can handle a combination of errors and erasures. Thus, by using their codes as an outer code in a concatenated construction, one obtains a linear-time encodable code that attains the Zyablov bound [9, p. 1949], with a linear-time bounded-distance decoder. Alternatively, such a concatenated construction approaches the capacity of any given memoryless channel: if the inner code is taken to have the smallest decoding error exponent, then the overall decoding error probability behaves like Forney's error exponent [10], [11] (the time complexity of searching for the inner code, in turn, depends on ϵ, yet not on the overall length of the concatenated code).

Codes with similar attributes, both with respect to the Zyablov bound and to the capacity of memoryless channels, were presented also by Barg and Zémor in a sequence of papers [5], [6], [7] (yet in their constructions, only the decoding is guaranteed to be linear-time).

In this work, we present a family of codes which improves on the Guruswami-Indyk construction. Specifically, our codes will satisfy properties (P1)–(P3), except that the alphabet size in property (P1) will now be only

$$2^{O\left((\log(1/\epsilon))/\epsilon^3\right)} . \tag{2}$$

The authors are with the Computer Science Department, Technion, Haifa 32000, Israel, e-mail: {ronny, vitalys}@cs.technion.ac.il. This work was supported by the Israel Science Foundation (Grant No. 746/04). Part of this work was presented at the *2004 IEEE Int'l Symposium on Information Theory (ISIT'2004),* Chicago, Illinois (June 2004).

The basic ingredients of our construction are similar to those used in [12] (and also in [3] and [4]), yet their layout (in particular, the order of application of the various building blocks), and the choice of parameters will be different. Our presentation will be split into two parts. We first describe in Section II a construction that satisfies only the two properties (P1) and (P3) over an alphabet of size (2). These two properties will be proved in Sections III and IV. We also show that the codes studied by Barg and Zémor in [5] and [7] can be seen as concatenated codes, with our codes serving as the outer codes.

The second part of our presentation consists of Section V, where we modify the construction of Section II and use the resulting code as a building block in a second construction, which satisfies property (P2) as well.

II. Construction of linear-time decodable codes

Let $\mathcal{G} = (V' : V'', E)$ be a bipartite Δ-regular undirected connected graph with a vertex set $V = V' \cup V''$ such that $V' \cap V'' = \emptyset$, and an edge set E such that every edge in E has one endpoint in V' and one endpoint in V''. We denote the size of V' by n (clearly, n is also the size of V'') and we will assume hereafter without any practical loss of generality that $n > 1$. For every vertex $u \in V$, we denote by $E(u)$ the set of edges that are incident with u. We assume an ordering on V, thereby inducing an ordering on the edges of $E(u)$ for every $u \in V$. For an alphabet F and a word $\boldsymbol{z} = (z_e)_{e \in E}$ (whose entries are indexed by E) in $F^{|E|}$, we denote by $(\boldsymbol{z})_{E(u)}$ the sub-block of $\boldsymbol{z}$ that is indexed by $E(u)$.

Let F be the field $\mathrm{GF}(q)$ and let $\mathcal{C}'$ and $\mathcal{C}''$ be linear $[\Delta, r\Delta, \theta\Delta]$ and $[\Delta, R\Delta, \delta\Delta]$ codes over F, respectively. We define the code $\mathsf{C} = (\mathcal{G}, \mathcal{C}' : \mathcal{C}'')$ as the following linear code of length $|E|$ over F:

$$\mathsf{C} = \Big\{ \boldsymbol{c} \in F^{|E|} \,:\, (\boldsymbol{c})_{E(u)} \in \mathcal{C}' \text{ for every } u \in V' \text{ and } (\boldsymbol{c})_{E(v)} \in \mathcal{C}'' \text{ for every } v \in V'' \Big\}$$

(C is the primary code considered by Barg and Zémor in [5]).

Let Φ be the alphabet $F^{r\Delta}$. Fix some linear one-to-one mapping $\mathcal{E} : \Phi \to \mathcal{C}'$ over F, and let the mapping $\psi_{\mathcal{E}} : \mathsf{C} \to \Phi^n$ be given by

$$\psi_{\mathcal{E}}(\boldsymbol{c}) = \left(\mathcal{E}^{-1}\left((\boldsymbol{c})_{E(u)} \right) \right)_{u \in V'} , \quad \boldsymbol{c} \in \mathsf{C} . \tag{3}$$

That is, the entries of $\psi_{\mathcal{E}}(\boldsymbol{c})$ are indexed by V', and the entry that is indexed by $u \in V'$ equals $\mathcal{E}^{-1}\left((\boldsymbol{c})_{E(u)} \right)$. We now define the code $(\mathsf{C})_\Phi$ of length n over Φ by

$$(\mathsf{C})_\Phi = \{ \psi_{\mathcal{E}}(\boldsymbol{c}) \,:\, \boldsymbol{c} \in \mathsf{C} \} .$$

Every codeword $\boldsymbol{x} = (\boldsymbol{x}_u)_{u \in V'}$ of $(\mathsf{C})_\Phi$ (with entries $\boldsymbol{x}_u$ in Φ) is associated with a unique codeword $\boldsymbol{c} \in \mathsf{C}$ such that

$$\mathcal{E}(\boldsymbol{x}_u) = (\boldsymbol{c})_{E(u)} , \quad u \in V' .$$

Based on the definition of $(\mathsf{C})_\Phi$, the code C can be represented as a concatenated code with an inner code $\mathcal{C}'$ over F and an outer code $(\mathsf{C})_\Phi$ over Φ. It is possible, however, to use $(\mathsf{C})_\Phi$ as an outer code with inner codes other than $\mathcal{C}'$. Along these lines, the codes studied in [5] and [7] can be represented as concatenated codes with $(\mathsf{C})_\Phi$ as an outer code, whereas the inner codes are taken over a sub-field of F.

III. Bounds on the code parameters

Let $\mathsf{C} = (\mathcal{G}, \mathcal{C}' : \mathcal{C}'')$, Φ, and $(\mathsf{C})_\Phi$ be as defined in Section II. It was shown in [5] that the rate of C is at least $r + R - 1$. From the fact that C is a concatenated code with an inner code $\mathcal{C}'$ and an outer code $(\mathsf{C})_\Phi$, it follows that the rate of $(\mathsf{C})_\Phi$ is bounded from below by

$$\frac{r + R - 1}{r} = 1 - \frac{1}{r} + \frac{R}{r} . \tag{4}$$

In particular, the rate approaches R when $r \to 1$.

We next turn to computing a lower bound on the relative minimum distance of $(\mathsf{C})_\Phi$. By applying this lower bound, we will then verify that $(\mathsf{C})_\Phi$ satisfies property (P1). Our analysis is based on that in [7], and we obtain here an improvement over a bound that can be inferred from [7]; we will need that improvement to get the reduction of the alphabet size from (1) to (2). We first introduce several notations.

Denote by $A_{\mathcal{G}}$ the adjacency matrix of $\mathcal{G}$; namely, $A_{\mathcal{G}}$, is a $|V| \times |V|$ real symmetric matrix whose rows and columns are indexed by the set V, and for every $u, v \in V$, the entry in $A_{\mathcal{G}}$ that is indexed by (u, v) is given by

$$(A_{\mathcal{G}})_{u,v} = \begin{cases} 1 & \text{if } \{u, v\} \in E \\ 0 & \text{otherwise} \end{cases} .$$

It is known that Δ is the largest eigenvalue of $A_{\mathcal{G}}$. We denote by $\gamma_{\mathcal{G}}$ the ratio between the second largest eigenvalue of $A_{\mathcal{G}}$ and Δ (this ratio is less than 1 when $\mathcal{G}$ is connected and is nonnegative when $n > 1$; see [8, Propositions 1.1.2 and 1.1.4]).

When $\mathcal{G}$ is taken from a sequence of Ramanujan expander graphs with constant degree Δ, such as the LPS graphs in [13], we have

$$\gamma_{\mathcal{G}} \le \frac{2\sqrt{\Delta - 1}}{\Delta} .$$

For a nonempty subset S of the vertex set V of $\mathcal{G}$, we will use the notation $\mathcal{G}_S$ to stand for the subgraph of $\mathcal{G}$ that is induced by S: the vertex set of $\mathcal{G}_S$ is given by S, and its edge set consists of all the edges in $\mathcal{G}$ that have each of their endpoints in S. The degree of u in $\mathcal{G}_S$, which is the number of adjacent vertices to u in $\mathcal{G}_S$, will be denoted by $\deg_S(u)$.

Theorem 3.1: The relative minimum distance of the code $(\mathsf{C})_\Phi$ is bounded from below by

$$\frac{\delta - \gamma_{\mathcal{G}} \sqrt{\delta / \theta}}{1 - \gamma_{\mathcal{G}}} .$$

In particular, this lower bound approaches δ when $\gamma_{\mathcal{G}} \to 0$.

The proof of the theorem will make use of Proposition 3.3 below, which is an improvement on Corollary 9.2.5 in Alon and Spencer [2] for bipartite graphs, and is also an improvement on Lemma 4 in Zémor [20]. We will need the following technical lemma for that proposition. The proof of this lemma can be found in Appendix A.

Denote by $\mathcal{N}(u)$ the set of vertices that are adjacent to vertex u in $\mathcal{G}$.

Lemma 3.2: Let χ be a real function on the vertices of $\mathcal{G}$ where the images of χ are restricted to the interval $[0, 1]$. Write

$$\sigma = \frac{1}{n} \sum_{u \in V'} \chi(u) \qquad \text{and} \qquad \tau = \frac{1}{n} \sum_{v \in V''} \chi(v) .$$

Then

$$\begin{aligned}\frac{1}{\Delta n}\sum_{u\in V'}\sum_{v\in\mathcal{N}(u)}\chi(u)\chi(v) &\le \sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma(1-\sigma)\tau(1-\tau)}\\ &\le (1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}\ .\end{aligned}$$

(Comparing to the results in [20], Lemma 4 therein is stated for the special case where the images of χ are either 0 or 1. Our first inequality in Lemma 3.2 yields a bound which is always at least as tight as Lemma 4 in [20].)

Proposition 3.3: Let $S \subseteq V'$ and $T \subseteq V''$ be subsets of sizes $|S| = \sigma n$ and $|T| = \tau n$, respectively, such that $\sigma+\tau > 0$. Then the sum of the degrees in the graph $\mathcal{G}_{S\cup T}$ is bounded from above by

$$\sum_{u\in S\cup T}\deg_{S\cup T}(u) \le 2\left((1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}\right)\Delta n\ .$$

Proof: We select $\chi(u)$ in Lemma 3.2 to be

$$\chi(u) = \begin{cases}1 & \text{if } u\in S\cup T\\ 0 & \text{otherwise}\end{cases}\ .$$

On the one hand, by Lemma 3.2,

$$\sum_{u\in V'}\sum_{v\in\mathcal{N}(u)}\chi(u)\chi(v) \le \left((1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}\right)\Delta n\ .$$

On the other hand,

$$2\sum_{u\in V'}\sum_{v\in\mathcal{N}(u)}\chi(u)\chi(v) = \sum_{u\in S\cup T}\deg_{S\cup T}(u)\ .$$

These two equations yield the desired result. □

Proof of Theorem 3.1: First, it is easy to see that $(\mathsf{C})_\Phi$ is a linear subspace over F and, as such, it is an Abelian subgroup of Φ^n. Thus, the minimum distance of $(\mathsf{C})_\Phi$ equals the minimum weight (over Φ) of any nonzero codeword of $(\mathsf{C})_\Phi$.

Pick any nonzero codeword $\boldsymbol{x} \in (\mathsf{C})_\Phi$, and let $\boldsymbol{c} = (c_e)_{e\in E}$ be the unique codeword in C such that $\boldsymbol{x} = \psi_{\mathcal{E}}(\boldsymbol{c})$. Denote by $Y \subseteq E$ the support of $\boldsymbol{c}$ (over F), i.e.,

$$Y = \{e \in E \,:\, c_e \ne 0\}\ .$$

Let S (respectively, T) be the set of all vertices in V' (respectively, V'') that are endpoints of edges in Y. In particular, S is the support of the codeword $\boldsymbol{x}$. Let σ and τ denote the ratios $|S|/n$ and $|T|/n$, respectively, and consider the subgraph $\mathcal{G}(Y) = (S : T, Y)$ of $\mathcal{G}$. Since the minimum distance of $\mathcal{C}'$ is $\theta\Delta$, the degree in $\mathcal{G}(Y)$ of every vertex in V' is at least $\theta\Delta$. Therefore, the number of edges in $\mathcal{G}(Y)$ satisfies

$$|Y| \ge \theta\Delta\cdot\sigma n\ .$$

Similarly, the degree in $\mathcal{G}(Y)$ of every vertex in V'' is at least $\delta\Delta$ and, thus,

$$|Y| \ge \delta\Delta\cdot\tau n\ .$$

Therefore,

$$|Y| \ge \max\{\theta\sigma,\delta\tau\}\cdot\Delta n\ .$$

On the other hand, $\mathcal{G}(Y)$ is a subgraph of $\mathcal{G}_{S\cup T}$; hence, by Proposition 3.3,

$$|Y| \le \frac{1}{2}\sum_{u\in S\cup T}\deg_{S\cup T}(u) \le \left((1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}\right)\Delta n\ .$$

Combining the last two equations yields

$$\max\{\theta\sigma,\delta\tau\} \le (1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}\ . \tag{5}$$

We now distinguish between two cases.

Case 1: $\sigma/\tau \le \delta/\theta$. Here (5) becomes

$$\delta\tau \le (1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}$$

and, so,

$$\sigma \ge \frac{\delta-\gamma_{\mathcal{G}}\sqrt{\sigma/\tau}}{1-\gamma_{\mathcal{G}}} \ge \frac{\delta-\gamma_{\mathcal{G}}\sqrt{\delta/\theta}}{1-\gamma_{\mathcal{G}}}\ . \tag{6}$$

Case 2: $\sigma/\tau > \delta/\theta$. By exchanging between σ and τ and between θ and δ in (6), we get

$$\tau \ge \frac{\theta-\gamma_{\mathcal{G}}\sqrt{\theta/\delta}}{1-\gamma_{\mathcal{G}}}\ .$$

Therefore,

$$\sigma > \frac{\delta}{\theta}\cdot\tau \ge \frac{\delta}{\theta}\cdot\frac{\theta-\gamma_{\mathcal{G}}\sqrt{\theta/\delta}}{1-\gamma_{\mathcal{G}}} = \frac{\delta-\gamma_{\mathcal{G}}\sqrt{\delta/\theta}}{1-\gamma_{\mathcal{G}}}\ .$$

Either case yields the desired lower bound on the size, σn, of the support S of $\boldsymbol{x}$. □

The next example demonstrates how the parameters of $(\mathsf{C})_\Phi$ can be tuned so that the improvement (2) of property (P1) holds.

Example 3.1: Fix $\theta = \epsilon$ for some small $\epsilon \in (0, 1]$ (in which case $r > 1-\epsilon$), and then select q and Δ so that $q > \Delta \ge 4/\epsilon^3$. For such parameters, we can take $\mathcal{C}'$ and $\mathcal{C}''$ to be generalized Reed-Solomon (GRS) codes over F. We also assume that $\mathcal{G}$ is a Ramanujan bipartite graph, in which case

$$\gamma_{\mathcal{G}} \le \frac{2\sqrt{\Delta-1}}{\Delta} < \epsilon^{3/2}\ .$$

By (4), the rate of $(\mathsf{C})_\Phi$ is bounded from below by

$$1-\frac{1}{1-\epsilon}+\frac{R}{1-\epsilon} > R-\epsilon\ ,$$

and by Theorem 3.1, the relative minimum distance is at least

$$\begin{aligned}\frac{\delta-\gamma_{\mathcal{G}}\sqrt{\delta/\theta}}{1-\gamma_{\mathcal{G}}} &\ge \delta-\gamma_{\mathcal{G}}\sqrt{\delta/\theta} > \delta-\epsilon^{3/2}\cdot\frac{1}{\sqrt{\epsilon}}\\ &= \delta-\epsilon > 1-R-\epsilon\ .\end{aligned}$$

Thus, the code $(\mathsf{C})_\Phi$ approaches the Singleton bound when $\epsilon \to 0$. In addition, if q and Δ are selected to be (no larger than) $O(1/\epsilon^3)$, then the alphabet Φ has size

$$|\Phi| = q^{r\Delta} = 2^{O\left((\log(1/\epsilon))/\epsilon^3\right)}\ .$$

□

From Example 3.1 we can state the following corollary.

Corollary 3.4: For any designed rate $R \in (0, 1]$ and sufficiently small $\epsilon > 0$ there is an infinite family of codes $(\mathsf{C})_\Phi$ of rate at least R and relative minimum distance greater than $1-R-\epsilon$, over an alphabet of size as in (2).

IV. Decoding algorithm

Let $\mathsf{C} = (\mathcal{G}, \mathcal{C}' : \mathcal{C}'')$ be defined over $F = \mathrm{GF}(q)$ as in Section II. Figure 1 presents an adaptation of the iterative decoder of Sipser and Spielman [16] and Zémor [20] to the code $(\mathsf{C})_\Phi$, with the additional feature of handling erasures (as well as errors over Φ): as we show in Theorem 4.1 below, the algorithm corrects any pattern of t errors and ρ erasures, provided that $t + (\rho/2) < \beta n$, where

$$\beta = \frac{(\delta/2) - \gamma_{\mathcal{G}}\sqrt{\delta/\theta}}{1 - \gamma_{\mathcal{G}}} .$$

Note that β equals approximately half the lower bound in Theorem 3.1. The value of ν in the algorithm, which is specified in Theorem 4.1 below, grows logarithmically with n.

We use the notation "?" to stand for an erasure. The algorithm in Figure 1 makes use of a word $\boldsymbol{z} = (z_e)_{e \in E}$ over $F \cup \{?\}$ that is initialized according to the contents of the received word $\boldsymbol{y}$ as follows. Each sub-block $(\boldsymbol{z})_{E(u)}$ that corresponds to a non-erased entry $\boldsymbol{y}_u$ of $\boldsymbol{y}$ is initialized to the codeword $\mathcal{E}(\boldsymbol{y}_u)$ of $\mathcal{C}'$. The remaining sub-blocks $(\boldsymbol{z})_{E(u)}$ are initialized as erased words of length Δ. Iterations $i = 3, 5, 7, \ldots$ use an error-correcting decoder $\mathcal{D}' : F^\Delta \to \mathcal{C}'$ that recovers correctly any pattern of less than $\theta\Delta/2$ errors (over F), and iterations $i = 2, 4, 6, \ldots$ use a combined error-erasure decoder $\mathcal{D}'' : (F \cup \{?\})^\Delta \to \mathcal{C}''$ that recovers correctly any pattern of a errors and b erasures, provided that $2a + b < \delta\Delta$ (b will be positive only when $i = 2$).

Theorem 4.1: Suppose that

$$\sqrt{\theta\delta} > 2\gamma_{\mathcal{G}} > 0 , \tag{7}$$

and fix σ to be a positive real number such that

$$\sigma < \beta = \frac{(\delta/2) - \gamma_{\mathcal{G}}\sqrt{\delta/\theta}}{1 - \gamma_{\mathcal{G}}} . \tag{8}$$

If

$$\nu = 2\left\lfloor \log\left(\frac{\beta\sqrt{\sigma n} - \sigma}{\beta - \sigma}\right)\right\rfloor + 3$$

then the decoder in Figure 1 recovers correctly any pattern of t errors (over Φ) and ρ erasures, provided that

$$t + \frac{\rho}{2} \le \sigma n . \tag{9}$$

The proof of the theorem makes use of the following lemma.

Lemma 4.2: Let χ, σ, and τ be as in Lemma 3.2, and suppose that the restriction of χ to V'' is not identically zero and that $\gamma_{\mathcal{G}} > 0$. Let δ be a real number for which the following condition is satisfied for every $v \in V''$:

$$\chi(v) > 0 \implies \sum_{u \in \mathcal{N}(v)} \chi(u) \ge \frac{\delta\Delta}{2}.$$

Then

$$\sqrt{\frac{\sigma}{\tau}} \ge \frac{(\delta/2) - (1-\gamma_{\mathcal{G}})\sigma}{\gamma_{\mathcal{G}}} .$$

The proof of Lemma 4.2 can be found in Appendix A. This lemma implies an upper bound on τ, in terms of σ; it can be verified that this bound is always at least as tight as Lemma 5 in [20].

Proof of Theorem 4.1: For $i \ge 2$, let U_i be the value of the set U at the end of iteration i in Figure 1, and let S_i be the set of all vertices $u \in U_i$ such that $(\boldsymbol{z})_{E(u)}$ is in error at the end of that iteration. Let $\chi_1 : (V' \cup V'') \to \{0, \frac{1}{2}, 1\}$ be the function

$$\chi_1(u) = \begin{cases} 1 & \text{if } u \in V' \text{ and } \boldsymbol{y}_u \text{ is in error} \\ \frac{1}{2} & \text{if } u \in V' \text{ and } \boldsymbol{y}_u \text{ is an erasure} \\ 0 & \text{otherwise} \end{cases} ,$$

and, for $i \ge 2$ define the function $\chi_i : (V' \cup V'') \to \{0, \frac{1}{2}, 1\}$ recursively by

$$\chi_i(u) = \begin{cases} 1 & \text{if } u \in S_i \\ 0 & \text{if } u \in U_i \setminus S_i \\ \chi_{i-1}(u) & \text{if } u \in U_{i-1} \end{cases} ,$$

where $U_1 = V'$.

Denote

$$\sigma_i = \frac{1}{n}\sum_{u \in U_i} \chi_i(u) .$$

Obviously, $\sigma_1 n = t + (\rho/2)$ and, so, by (9) we have $\sigma_1 \le \sigma$.

Let ℓ be the smallest positive integer (possibly ∞) such that $\sigma_\ell = 0$. Since both $\mathcal{D}'$ and $\mathcal{D}''$ are bounded-distance decoders, a vertex $v \in U_i$ can belong to S_i for even $i \ge 2$, only if the sum $\sum_{u \in \mathcal{N}(v)} \chi_i(u)$ (which equals the sum $\sum_{u \in \mathcal{N}(v)} \chi_{i-1}(u)$) is at least $\delta\Delta/2$. Similarly, a vertex $v \in U_i$ belongs to S_i for odd $i > 1$, only if $\sum_{u \in \mathcal{N}(v)} \chi_i(u) \ge \theta\Delta/2$. It follows that the function χ_i satisfies the conditions of Lemma 4.2 (with θ taken instead of δ for odd i) and, so,

$$\sqrt{\frac{\sigma_{i-1}}{\sigma_i}} \ge \begin{cases} \dfrac{\delta}{2\gamma_{\mathcal{G}}} - \dfrac{1-\gamma_{\mathcal{G}}}{\gamma_{\mathcal{G}}}\sigma_{i-1} & \text{for even } 0 < i < \ell \\ \dfrac{\theta}{2\gamma_{\mathcal{G}}} - \dfrac{1-\gamma_{\mathcal{G}}}{\gamma_{\mathcal{G}}}\sigma_{i-1} & \text{for odd } 1 < i < \ell \end{cases} . \tag{10}$$

Using the condition $\sigma_1 \le \sigma < \beta$, it can be verified by induction on $i \ge 2$ that

$$\frac{\sigma_{i-1}}{\sigma_i} \ge \begin{cases} \delta/\theta & \text{for even } 0 < i < \ell \\ \theta/\delta & \text{for odd } 1 < i < \ell \end{cases} . \tag{11}$$

Hence, for every $i > 2$,

$$\frac{\sigma_{i-2}}{\sigma_i} = \frac{\sigma_{i-2}}{\sigma_{i-1}} \cdot \frac{\sigma_{i-1}}{\sigma_i} \ge \frac{\delta}{\theta} \cdot \frac{\theta}{\delta} = 1 ;$$

in particular, $\sigma_i \le \sigma$ for odd i and $\sigma_i \le \sigma_2$ for even i. Incorporating these inequalities into (10) yields

$$\frac{1}{\sqrt{\sigma_i}} \ge \frac{\delta}{2\gamma_{\mathcal{G}}\sqrt{\sigma_{i-1}}} - \frac{1-\gamma_{\mathcal{G}}}{\gamma_{\mathcal{G}}}\sqrt{\sigma} \qquad \text{for even } 0 < i < \ell \tag{12}$$

and

$$\frac{1}{\sqrt{\sigma_i}} \ge \frac{\theta}{2\gamma_{\mathcal{G}}\sqrt{\sigma_{i-1}}} - \frac{1-\gamma_{\mathcal{G}}}{\gamma_{\mathcal{G}}}\sqrt{\sigma_2} \qquad \text{for odd } 1 < i < \ell . \tag{13}$$

By combining (12) and (13) we get that for even $i > 0$,

$$\begin{aligned} \frac{2\gamma_{\mathcal{G}}}{\theta\sqrt{\sigma_{i+1}}} + \frac{2(1-\gamma_{\mathcal{G}})}{\theta}\sqrt{\sigma_2} \;&\ge\; \frac{1}{\sqrt{\sigma_i}} \\ &\ge\; \frac{\delta}{2\gamma_{\mathcal{G}}\sqrt{\sigma_{i-1}}} - \frac{1-\gamma_{\mathcal{G}}}{\gamma_{\mathcal{G}}}\sqrt{\sigma} , \end{aligned}$$

Input: Received word $\boldsymbol{y} = (\boldsymbol{y}_u)_{u \in V'}$ in $(\Phi \cup \{?\})^n$.

Initialize: For $u \in V'$ do: $\quad (\boldsymbol{z})_{E(u)} \leftarrow \begin{cases} \mathcal{E}(\boldsymbol{y}_u) & \text{if } \boldsymbol{y}_u \in \Phi \\ ??\ldots? & \text{if } \boldsymbol{y}_u = ? \end{cases}$.

Iterate: For $i = 2, 3, \ldots, \nu$ do:

(a) If i is odd then $U \equiv V'$ and $\mathcal{D} \equiv \mathcal{D}'$, else $U \equiv V''$ and $\mathcal{D} \equiv \mathcal{D}''$.

(b) For every $u \in U$ do: $(\boldsymbol{z})_{E(u)} \leftarrow \mathcal{D}\left((\boldsymbol{z})_{E(u)}\right)$.

Output: $\psi_{\mathcal{E}}(\boldsymbol{z})$ if $\boldsymbol{z} \in \mathsf{C}$ (and declare 'error' otherwise).

Fig. 1. Decoder for $(\mathsf{C})_\Phi$.

or

$$\begin{aligned} \frac{1}{\sqrt{\sigma_{i+1}}} &\geq \frac{\theta\delta}{4\gamma_{\mathcal{G}}^2\sqrt{\sigma_{i-1}}} - \frac{1-\gamma_{\mathcal{G}}}{\gamma_{\mathcal{G}}}\left(\frac{\theta\sqrt{\sigma}}{2\gamma_{\mathcal{G}}} + \sqrt{\sigma_2}\right) \\ &\geq \frac{\theta\delta}{4\gamma_{\mathcal{G}}^2\sqrt{\sigma_{i-1}}} - \frac{1-\gamma_{\mathcal{G}}}{\gamma_{\mathcal{G}}}\left(\frac{\theta}{2\gamma_{\mathcal{G}}} + \sqrt{\frac{\theta}{\delta}}\right)\sqrt{\sigma} \\ &= \frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\left(\frac{1}{\sqrt{\sigma_{i-1}}} - \frac{\sqrt{\sigma}}{\beta}\right) + \frac{\sqrt{\sigma}}{\beta}\ , \qquad (14) \end{aligned}$$

where the second inequality follows from $\sigma_2 \leq \sigma \cdot \theta/\delta$ (see (11)), and the (last) equality follows from the next chain of equalities:

$$\begin{aligned} &\frac{1-\gamma_{\mathcal{G}}}{\gamma_{\mathcal{G}}}\left(\frac{\theta}{2\gamma_{\mathcal{G}}} + \sqrt{\frac{\theta}{\delta}}\right)\sqrt{\sigma} \\ &\qquad = \frac{1-\gamma_{\mathcal{G}}}{2\gamma_{\mathcal{G}}^2}\left(2\gamma_{\mathcal{G}} + \sqrt{\theta\delta}\right)\sqrt{\frac{\sigma\theta}{\delta}} \\ &\qquad = -\frac{1-\gamma_{\mathcal{G}}}{2\gamma_{\mathcal{G}}^2} \cdot \frac{4\gamma_{\mathcal{G}}^2 - \theta\delta}{\sqrt{\theta\delta} - 2\gamma_{\mathcal{G}}}\sqrt{\frac{\sigma\theta}{\delta}} \\ &\qquad = -\left(1 - \frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)\frac{(1-\gamma_{\mathcal{G}})\sqrt{\sigma}}{(\delta/2) - \gamma_{\mathcal{G}}\sqrt{\delta/\theta}} \\ &\qquad = -\left(1 - \frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)\frac{\sqrt{\sigma}}{\beta}\ . \end{aligned}$$

Consider the following first-order linear recurring sequence $(\Lambda_j)_{j \geq 0}$ that satisfies

$$\Lambda_{j+1} = \frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\left(\Lambda_j - \frac{\sqrt{\sigma}}{\beta}\right) + \frac{\sqrt{\sigma}}{\beta}\ , \quad j \geq 0\ ,$$

where $\Lambda_0 = 1/\sqrt{\sigma}$. From (14) we have $1/\sqrt{\sigma_{i+1}} \geq \Lambda_{i/2}$ for even $i \geq 0$. By solving the recurrence for (Λ_j), we obtain

$$\frac{1}{\sqrt{\sigma_{i+1}}} \geq \Lambda_{i/2} = \left(\left(\frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)^{i/2}\left(1 - \frac{\sigma}{\beta}\right) + \frac{\sigma}{\beta}\right)\frac{1}{\sqrt{\sigma}}\ . \quad (15)$$

From the condition (7) we thus get that σ_{i+1} decreases exponentially with (even) i. A sufficient condition for ending the decoding correctly after ν iterations is having $\sigma_\nu < 1/n$, or

$$\frac{1}{\sqrt{\sigma_\nu}} > \sqrt{n}\ .$$

We require therefore that ν be such that

$$\frac{1}{\sqrt{\sigma_\nu}} \geq \left(\left(\frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)^{(\nu-1)/2}\left(1 - \frac{\sigma}{\beta}\right) + \frac{\sigma}{\beta}\right)\frac{1}{\sqrt{\sigma}} > \sqrt{n}\ .$$

The latter inequality can be rewritten as

$$\left(\frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)^{(\nu-1)/2} > \frac{\sqrt{n\sigma} - (\sigma/\beta)}{1 - (\sigma/\beta)} = \frac{\beta\sqrt{n\sigma} - \sigma}{\beta - \sigma}\ ,$$

thus yielding

$$\nu > 2\log\left(\frac{\beta\sqrt{n\sigma} - \sigma}{\beta - \sigma}\right) + 1\ ,$$

where the base of the logarithm equals $(\theta\delta)/(4\gamma_{\mathcal{G}}^2)$. In summary, the decoding will end with the correct codeword after

$$\nu = 2\left\lfloor \log\left(\frac{\beta\sqrt{n\sigma} - \sigma}{\beta - \sigma}\right)\right\rfloor + 3$$

iterations (where the base of the logarithm again equals $(\theta\delta)/(4\gamma_{\mathcal{G}}^2)$.) $\square$

In Lemma B.1, which appears in Appendix B, it is shown that the number of actual applications of the decoders $\mathcal{D}'$ and $\mathcal{D}''$ in the algorithm in Figure 1 can be bounded from above by $\omega \cdot n$, where

$$\omega = 2 \cdot \left\lceil \frac{\log\left(\frac{\Delta\beta\sqrt{\sigma}}{\beta - \sigma}\right)}{\log\left(\frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)} \right\rceil + \frac{1 + \frac{\theta}{\delta}}{1 - \left(\frac{4\gamma_{\mathcal{G}}^2}{\theta\delta}\right)^2}\ .$$

Thus, if θ and δ are fixed and the ratio σ/β is bounded away from 1 and $\mathcal{G}$ is a Ramanujan graph, then the value of ω is bounded from above by an absolute constant (independent of Δ).

The algorithm in Figure 1 allows us to use GMD decoding in cases where $(\mathsf{C})_\Phi$ is used as an outer code in a concatenated code. In such a concatenated code, the size of the inner code is $|\Phi|$ and, thus, it does not grow with the length n of $(\mathsf{C})_\Phi$. A GMD decoder will apply the algorithm in Figure 1 a number of times that is proportional to the minimum distance of the inner code. Thus, if the inner code has rate that is bounded away from zero, then the GMD decoder will have time complexity that grows linearly with the overall code length. Furthermore, if $\mathcal{C}'$, $\mathcal{C}''$, and the inner code are codes that have a polynomial-time bounded-distance decoder—e.g., if they are GRS codes—then the multiplying constant in the linear expression of the time complexity (when measured in operations in F) is $\text{POLY}(\Delta)$. For the choice of parameters in Example 3.1, this constant is $\text{POLY}(1/\epsilon)$ and, since F is chosen in that example to have size $O(1/\epsilon^3)$, each operation in F can in turn be implemented by $\text{POLY}(\log(1/\epsilon))$ bit operations. (We remark that in all our complexity estimates, we assume that the graph $\mathcal{G}$ is "hard-wired" so that we can ignore the complexity of figuring out the set of incident edges of a given vertex in $\mathcal{G}$. Along these lines, we assume that each access to an entry takes constant time, even though the length of the index of that entry may grow logarithmically with the code length. See the discussion in [16, Section II].)

When the inner code is taken as $\mathcal{C}'$, the concatenation results in the code $\mathsf{C} = (\mathcal{G}, \mathcal{C}' : \mathcal{C}'')$ (of length Δn) over F, and the (linear-time) correctable fraction of errors is then the product $\theta \cdot \sigma$, for any positive real σ that satisfies (8). A special case of this result, for $F = \mathrm{GF}(2)$ and $\mathcal{C}' = \mathcal{C}''$, was presented in our earlier work [17], yet the analysis therein was different. A linear-time decoder for C was also presented by Barg and Zémor in [7], except that their decoder requires finding a codeword that minimizes some weighted distance function, and we are unaware of a method that performs this task in time complexity that is $\mathrm{POLY}(\Delta)$—even when $\mathcal{C}'$ and $\mathcal{C}''$ have a polynomial-time bounded-distance decoder.

V. CONSTRUCTION WHICH IS ALSO LINEAR-TIME ENCODABLE

In this section, we use the construction $(\mathsf{C})_\Phi$ of Section II as a building block in obtaining a second construction, which satisfies all properties (P1)–(P3) over an alphabet whose size is given by (2).

A. *Outline of the construction*

Let $\mathsf{C} = (\mathcal{G}, \mathcal{C}' : \mathcal{C}'')$ be defined over $F = \mathrm{GF}(q)$ as in Section II. The first simple observation that provides the intuition behind the upcoming construction is that the encoding of C, and hence of $(\mathsf{C})_\Phi$, can be easily implemented in linear time if the code $\mathcal{C}'$ has rate $r = 1$, in which case $\Phi = F^\Delta$. The definition of C then reduces to

$$\mathsf{C} = \left\{ \boldsymbol{c} \in F^{|E|} \,:\, (\boldsymbol{c})_{E(v)} \in \mathcal{C}'' \text{ for every } v \in V'' \right\} .$$

We can implement an encoder of C as follows. Let $\mathcal{E}'' : F^{R\Delta} \to \mathcal{C}''$ be some one-to-one encoding mapping of $\mathcal{C}''$. Given an information word $\boldsymbol{\eta}$ in $F^{R\Delta n}$, it is first recast into a word of length n over $F^{R\Delta}$ by sub-dividing it into sub-blocks $\boldsymbol{\eta}_v \in F^{R\Delta}$ that are indexed by $v \in V''$; then a codeword $\boldsymbol{c} \in \mathsf{C}$ is computed by

$$(\boldsymbol{c})_{E(v)} = \mathcal{E}''(\boldsymbol{\eta}_v) \ , \quad v \in V'' \ .$$

By selecting $\mathcal{E}$ in (3) as the identity mapping, we get that the respective codeword $\boldsymbol{x} = (\boldsymbol{x}_u)_{u \in V'} = \psi_{\mathcal{E}}(\boldsymbol{c})$ in $(\mathsf{C})_\Phi$ is

$$\boldsymbol{x}_u = (\boldsymbol{c})_{E(u)} \ , \quad u \in V' \ .$$

Thus, each of the Δ entries (over F) of the sub-block $\boldsymbol{x}_u$ can be associated with a vertex $v \in \mathcal{N}(u)$, and the value assigned to that entry is equal to one of the entries in $\mathcal{E}''(\boldsymbol{\eta}_v)$.

While having $\mathcal{C}' = \Phi$ $(= F^\Delta)$ allows easy encoding, the minimum distance of the resulting code $(\mathsf{C})_\Phi$ is obviously poor. To resolve this problem, we insert into the construction another linear $[\Delta, r_0\Delta, \theta_0\Delta]$ code $\mathcal{C}_0$ over F. Let H_0 be some $((1-r_0)\Delta) \times \Delta$ parity-check matrix of $\mathcal{C}_0$ and for a vector $\boldsymbol{h} \in F^{(1-r_0)\Delta}$, denote by $\mathcal{C}_0(\boldsymbol{h})$ the following coset of $\mathcal{C}_0$ within Φ:

$$\mathcal{C}_0(\boldsymbol{h}) = \{ \boldsymbol{v} \in \Phi \,:\, H_0 \boldsymbol{v} = \boldsymbol{h} \} \ .$$

Fix now a list of vectors $\boldsymbol{s} = (\boldsymbol{h}_u)_{u \in V'}$ where $\boldsymbol{h}_u \in F^{(1-r_0)\Delta}$, and define the subset $\mathsf{C}(\boldsymbol{s})$ of C by

$$\mathsf{C}(\boldsymbol{s}) = \{ \boldsymbol{c} \in \mathsf{C} \,:\, (\boldsymbol{c})_{E(u)} \in \mathcal{C}_0(\boldsymbol{h}_u) \text{ for every } u \in V' \} \ ;$$

accordingly, define the subset $(\mathsf{C}(\boldsymbol{s}))_\Phi$ of $(\mathsf{C})_\Phi$ by

$$(\mathsf{C}(\boldsymbol{s}))_\Phi = \left\{ \psi_{\mathcal{E}}(\boldsymbol{c}) = ((\boldsymbol{c})_{E(u)})_{u \in V'} \,:\, \boldsymbol{c} \in \mathsf{C}(\boldsymbol{s}) \right\} .$$

Now, if $\boldsymbol{s}$ is all-zero, then $\mathsf{C}(\boldsymbol{s})$ coincides with the code $\mathsf{C}(\boldsymbol{0}) = (\mathcal{G}, \mathcal{C}_0 : \mathcal{C}'')$; otherwise, $\mathsf{C}(\boldsymbol{s})$ is either empty or is a coset of $\mathsf{C}(\boldsymbol{0})$, where $\mathsf{C}(\boldsymbol{0})$ is regarded as a linear subspace of C over F. From this observation we conclude that the lower bound in Theorem 3.1 applies to any nonempty subset $(\mathsf{C}(\boldsymbol{s}))_\Phi$, except that we need to replace θ by θ_0.

In addition, a simple modification in the algorithm in Figure 1 adapts it to decode $(\mathsf{C}(\boldsymbol{s}))_\Phi$ so that Theorem 4.1 holds (again under the change $\theta \leftrightarrow \theta_0$): during odd iterations i, we apply to each sub-block $(\boldsymbol{z})_{E(u)}$ a bounded-distance decoder of $\mathcal{C}_0(\boldsymbol{h}_u)$, instead of the decoder $\mathcal{D}'$.

Therefore, our strategy in designing the linear-time encodable codes will be as follows. The raw data will first be encoded into a codeword $\boldsymbol{c}$ of C (where $\mathcal{C}' = \Phi$). Then we compute the n vectors

$$\boldsymbol{h}_u = H_0 \cdot (\boldsymbol{c})_{E(u)} \ , \quad u \in V' \ ,$$

and produce the list $\boldsymbol{s} = (\boldsymbol{h}_u)_{u \in V'}$; clearly, $\boldsymbol{c}$ belongs to $\mathsf{C}(\boldsymbol{s})$. The list $\boldsymbol{s}$ will then undergo additional encoding stages, and the result will be merged with $\psi_{\mathcal{E}}(\boldsymbol{c})$ to produce the final codeword. The parameters of $\mathcal{C}_0$, which determine the size of $\boldsymbol{s}$, will be chosen so that the overhead due to $\boldsymbol{s}$ will be negligible.

During decoding, $\boldsymbol{s}$ will be recovered first, and then we will apply the aforementioned adaptation to $(\mathsf{C}(\boldsymbol{s}))_\Phi$ of the decoder in Figure 1, to reconstruct the information word $\boldsymbol{\eta}$.

B. *Details of the construction*

We now describe the construction in more detail. We let F be the field $\mathrm{GF}(q)$ and Δ_1 and Δ_2 be positive integers. The construction makes use of two bipartite regular graphs,

$$\mathcal{G}_1 = (V' : V'', E_1) \qquad \text{and} \qquad \mathcal{G}_2 = (V' : V'', E_2) \ ,$$

of degrees Δ_1 and Δ_2, respectively. Both graphs have the same number of vertices; in fact, we are making a stronger assumption whereby both graphs are defined over the same set of vertices. We denote by n the size of V' (or V'') and by Φ_1 and Φ_2 the alphabets F^{Δ_1} and F^{Δ_2}, respectively. The notations $E_1(u)$ and $E_2(u)$ will stand for the sets of edges that are incident with a vertex u in $\mathcal{G}_1$ and $\mathcal{G}_2$, respectively.

We also assume that we have at our disposal the following four codes:

- a linear $[\Delta_1, r_0\Delta_1, \theta_0\Delta_1]$ code $\mathcal{C}_0$ over F;
- a linear $[\Delta_1, R_1\Delta_1, \delta_1\Delta_1]$ code $\mathcal{C}_1$ over F;
- a linear $[\Delta_2, R_2\Delta_2, \delta_2\Delta_2]$ code $\mathcal{C}_2$ over F;

- a code $\mathcal{C}_{\rm m}$ of length n and rate $r_{\rm m}$ over the alphabet $\Phi_{\rm m} = F^{R_2 \Delta_2}$.

The rates of these codes need to satisfy the relation

$$(1-r_0)\Delta_1 = r_{\rm m} R_2 \Delta_2 \ ,$$

and the code $\mathcal{C}_{\rm m}$ is assumed to have the following properties:

1) Its rate is bounded away from zero: there is a universal positive constant κ such that $r_{\rm m} \ge \kappa$.
2) $\mathcal{C}_{\rm m}$ is linear-time encodable, and the encoding time per symbol is $\text{POLY}(\log |\Phi_{\rm m}|)$.
3) $\mathcal{C}_{\rm m}$ has a decoder that recovers in linear-time any pattern of up to μn errors (over the alphabet $\Phi_{\rm m}$), where μ is a universal positive constant. The time complexity per symbol of the decoder is $\text{POLY}(\log |\Phi_{\rm m}|)$.

(By a universal constant we mean a value that does not depend on any other parameter, not even on the size of $\Phi_{\rm m}$.) For example, we can select as $\mathcal{C}_{\rm m}$ the code of Spielman in [18], in which case κ can be taken as $1/4$.

Based on these ingredients, we introduce the codes

$$\mathsf{C}_1 = (\mathcal{G}_1, \Phi_1 : \mathcal{C}_1) \qquad \text{and} \qquad \mathsf{C}_2 = (\mathcal{G}_2, \Phi_2 : \mathcal{C}_2)$$

over F. The code C_1 will play the role of the code C as outlined in Section V-A, whereas the codes $\mathcal{C}_{\rm m}$ and C_2 will be utilized for the encoding of the list $\boldsymbol{s}$ that was described there.

The overall construction, which we denote by $\mathbb{C}$, is now defined as the set of all words of length n over the alphabet

$$\Phi = \Phi_1 \times \Phi_2$$

that are obtained by applying the encoding algorithm in Figure 2 to information words $\boldsymbol{\eta}$ of length n over $F^{R_1 \Delta_1}$. A schematic diagram of the algorithm is shown in Figure 3. (In this algorithm, we use a notational convention whereby entries of information words $\boldsymbol{\eta}$ are indexed by V'', and so are codewords of $\mathcal{C}_{\rm m}$.)

From the discussion in Section V-A and from the assumption on the code $\mathcal{C}_{\rm m}$ it readily follows that the encoder in Figure 2 can be implemented in linear time, where the encoding complexity per symbol (when measured in operations in F) is $\text{POLY}(\Delta_1, \Delta_2)$. The rate of $\mathbb{C}$ is also easy to compute: the encoder in Figure 2 maps, in a one-to-one manner, an information word of length n over an alphabet of size $q^{R_1 \Delta_1}$, into a codeword of length n over an alphabet Φ of size $q^{\Delta_1 + \Delta_2}$. Thus, the rate of $\mathbb{C}$ is

$$\frac{R_1 \Delta_1 n}{(\Delta_1 + \Delta_2) n} = \frac{R_1}{1 + (\Delta_2/\Delta_1)} \ . \tag{16}$$

In the next section, we show how the parameters of $\mathbb{C}$ can be selected so that it becomes nearly-MDS and also linear-time decodable.

C. Design, decoding, and analysis

We will select the parameters of $\mathbb{C}$ quite similarly to Example 3.1. We assume that the rates R_1 and R_2 of $\mathcal{C}_1$ and $\mathcal{C}_2$ are the same and are equal to some prescribed value R, and define

$$\alpha_R = 8 \cdot (1-R) \cdot \max\{R/\mu, 2/\kappa\}$$

(notice that α_R can be bounded from above by a universal constant that does not depend on R, e.g., by $16/\min\{2\mu, \kappa\}$). We set $\theta_0 = \kappa \cdot \epsilon$ for some positive $\epsilon < R$ (in which case $1-r_0 < \kappa \cdot \epsilon$), and then select q, Δ_1, and Δ_2 so that $q > \Delta_1 \ge \alpha_R/\epsilon^3$ and

$$\Delta_2 = \frac{(1-r_0)\Delta_1}{r_{\rm m} R} \quad (< \Delta_1) \ ; \tag{17}$$

yet we also assume that q is (no larger than) $O(1/\epsilon^3)$. The graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ are taken as Ramanujan graphs and $\mathcal{C}_0$, $\mathcal{C}_1$, and $\mathcal{C}_2$ are taken as GRS codes over F. (Requiring that both Δ_1 and Δ_2 be valid degrees of Ramanujan graphs imposes some restrictions on the value $(1-r_0)/(r_{\rm m} R)$. These restrictions can be satisfied by tuning the precise rate of $\mathcal{C}_{\rm m}$ last.)

Given this choice of parameters, we obtain from (17) that $\Delta_2/\Delta_1 < \epsilon/R$ and, so, the rate (16) of $\mathbb{C}$ is greater than

$$\frac{R}{1 + (\epsilon/R)} > R - \epsilon \ . \tag{18}$$

The alphabet size of $\mathbb{C}$ is

$$|\Phi| = |\Phi_1| \cdot |\Phi_2| = q^{\Delta_1 + \Delta_2} = 2^{O\left((\log(1/\epsilon))/\epsilon^3\right)} \ ,$$

as in (2), where we have absorbed into the $O(\cdot)$ term the constants κ and μ.

Our next step in the analysis of the code $\mathbb{C}$ consists of showing that there exists a linear-time decoder which recovers correctly any pattern of t errors and ρ erasures, provided that

$$2t + \rho \le (1-R-\epsilon) n \ . \tag{19}$$

This, in turn, will also imply that the relative minimum distance of $\mathbb{C}$ is greater than $1-R-\epsilon$, thus establishing with (18) the fact that $\mathbb{C}$ is nearly-MDS.

Let $\boldsymbol{x} = (\boldsymbol{x}_u)_{u \in V'}$ be the transmitted codeword of $\mathbb{C}$, where

$$\boldsymbol{x}_u = \left((\boldsymbol{c})_{E_1(u)}, (\boldsymbol{d})_{E_2(u)} \right) \ ,$$

and let $\boldsymbol{y} = (\boldsymbol{y}_u)_{u \in V'}$ be the received word; each entry $\boldsymbol{y}_u$ takes the form $(\boldsymbol{y}_{u,1}, \boldsymbol{y}_{u,2})$, where $\boldsymbol{y}_{u,1} \in \Phi_1 \cup \{?\}$ and $\boldsymbol{y}_{u,2} \in \Phi_2 \cup \{?\}$. Consider the application of the algorithm in Figure 4 to $\boldsymbol{y}$, assuming that $\boldsymbol{y}$ contains t errors and ρ erasures, where $2t + \rho \le (1-R-\epsilon)n$.

Step (D1) is the counterpart of the initialization step in Figure 1 (the entries of $\boldsymbol{z}$ here are indexed by the edges of $\mathcal{G}_2$).

The role of Step (D2) is to compute a word $\tilde{\boldsymbol{w}} \in \Phi_{\rm m}^n$ that is close to the codeword $\boldsymbol{w}$ of $\mathcal{C}_{\rm m}$, which was generated in Step (E3) of Figure 2. Step (D2) uses the inverse of the encoder $\mathcal{E}_2$ (which was used in Step (E4)) and also a combined error-erasure decoder $\mathcal{D}_2 : (F \cup \{?\})^{\Delta_2} \to \mathcal{C}_2$ that recovers correctly any pattern of a errors (over F) and b erasures, provided that $2a + b < \delta_2 \Delta_2$. The next lemma provides an upper bound on the Hamming distance between $\boldsymbol{w}$ and $\tilde{\boldsymbol{w}}$ (as words of length n over $\Phi_{\rm m}$).

Lemma 5.1: Under the assumption (19), the Hamming distance between $\boldsymbol{w}$ and $\tilde{\boldsymbol{w}}$ (as words over $\Phi_{\rm m}$) is less than μn.

Input: Information word $\boldsymbol{\eta} = (\boldsymbol{\eta}_v)_{v \in V''}$ of length n over $F^{R_1 \Delta_1}$.

(E1) Using an encoder $\mathcal{E}_1 : F^{R_1\Delta_1} \to \mathcal{C}_1$, map $\boldsymbol{\eta}$ into a codeword $\boldsymbol{c}$ of C_1 by

$$(\boldsymbol{c})_{E_1(v)} \leftarrow \mathcal{E}_1(\boldsymbol{\eta}_v) \ , \quad v \in V'' \ .$$

(E2) Fix some $((1-r_0)\Delta_1) \times \Delta_1$ parity-check matrix H_0 of $\mathcal{C}_0$ over F, and compute the n vectors

$$\boldsymbol{h}_u \leftarrow H_0 \cdot (\boldsymbol{c})_{E_1(u)} \ , \quad u \in V' \ ,$$

to produce the list $\boldsymbol{s} = (\boldsymbol{h}_u)_{u \in V'}$.

(E3) Regard $\boldsymbol{s}$ as a word of length $(1-r_0)\Delta_1 n$ $(= r_{\mathrm{m}} R_2 \Delta_2 n)$ over F, and map it by an encoder of $\mathcal{C}_{\mathrm{m}}$ into a codeword $\boldsymbol{w} = (\boldsymbol{w}_v)_{v \in V''}$ of $\mathcal{C}_{\mathrm{m}}$.

(E4) Using an encoder $\mathcal{E}_2 : F^{R_2\Delta_2} \to \mathcal{C}_2$, map $\boldsymbol{w}$ into a codeword $\boldsymbol{d}$ of C_2 by

$$(\boldsymbol{d})_{E_2(v)} \leftarrow \mathcal{E}_2(\boldsymbol{w}_v) \ , \quad v \in V'' \ .$$

Output: Word $\boldsymbol{x} = (\boldsymbol{x}_u)_{u \in V'}$ in $(\Phi_1 \times \Phi_2)^n$ whose components are given by the pairs

$$\boldsymbol{x}_u = ((\boldsymbol{c})_{E_1(u)}, (\boldsymbol{d})_{E_2(u)}) \ , \quad u \in V' \ .$$

Fig. 2. Encoder for $\mathbb{C}$.

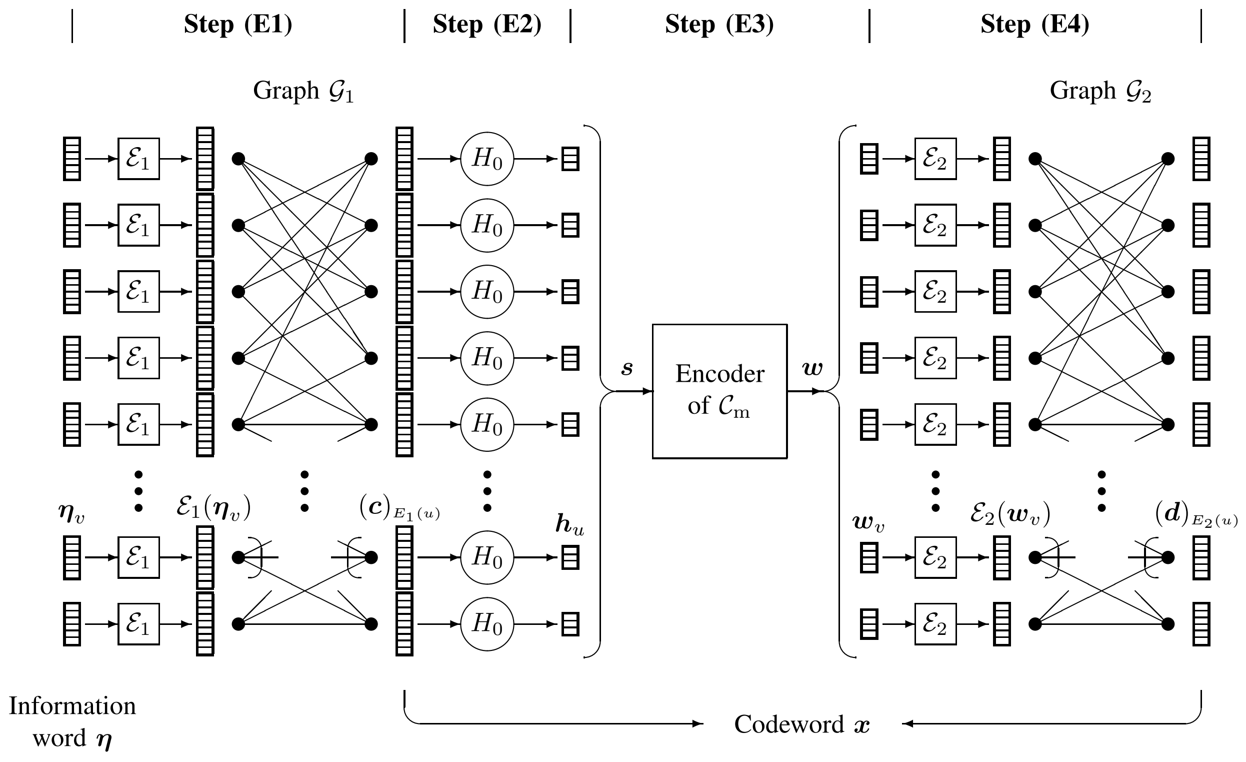

Fig. 3. Schematic diagram of the encoder for $\mathbb{C}$.

Input:	Received word $\boldsymbol{y} = (\boldsymbol{y}_u)_{u \in V'}$ in $(\Phi \cup \{?\})^n$.
(D1)	For $u \in V'$ do: $\quad (\boldsymbol{z})_{E_2(u)} \leftarrow \begin{cases} \boldsymbol{y}_{u,2} & \text{if } \boldsymbol{y}_{u,2} \in \Phi_1 \\ ??\ldots? & \text{if } \boldsymbol{y}_{u,2} = ? \end{cases}$.
(D2)	For $v \in V''$ do: $\quad \tilde{\boldsymbol{w}}_v \leftarrow \mathcal{E}_2^{-1}\left(\mathcal{D}_2\left((\boldsymbol{z})_{E_2(v)}\right)\right)$.
(D3)	Apply a decoder of $\mathcal{C}_{\mathrm{m}}$ to $\tilde{\boldsymbol{w}} = (\tilde{\boldsymbol{w}}_v)_{v \in V''}$ to produce an information word $\hat{\boldsymbol{s}} \in F^{(1-r_0)\Delta_1 n}$.
(D4)	Apply a decoder for $(\mathsf{C}_1(\hat{\boldsymbol{s}}))_{\Phi_1}$ to $(\boldsymbol{y}_{u,1})_{u \in V'}$, as described in Section V-A, to produce an information word $\hat{\boldsymbol{\eta}} = (\hat{\boldsymbol{\eta}}_v)_{v \in V''}$.
Output:	Information word $\hat{\boldsymbol{\eta}} = (\hat{\boldsymbol{\eta}}_v)_{v \in V''}$ of length n over $F^{R\Delta_1}$.

Fig. 4. Decoder for $(\mathsf{C})_\Phi$.

Proof: Define the function $\chi : (V' \cup V'') \to \{0, \frac{1}{2}, 1\}$ by

$$\chi(u) = \begin{cases} 1 & \text{if } u \in V' \text{ and } \boldsymbol{y}_{u,2} \text{ is in error} \\ \frac{1}{2} & \text{if } u \in V' \text{ and } \boldsymbol{y}_{u,2} \text{ is an erasure} \\ 1 & \text{if } u \in V'' \text{ and } \tilde{\boldsymbol{w}}_u \neq \boldsymbol{w}_u \\ 0 & \text{otherwise} \end{cases} .$$

Assuming that $\tilde{\boldsymbol{w}} \neq \boldsymbol{w}$, this function satisfies the conditions of Lemma 4.2 with respect to the graph $\mathcal{G}_2$, where σn equals $t + (\rho/2)$ and τn equals the number of vertices $v \in V''$ such that $\tilde{\boldsymbol{w}}_v \neq \boldsymbol{w}_v$. By that lemma we get

$$\begin{aligned} \sqrt{\frac{\sigma}{\tau}} \quad &\geq \quad \frac{(\delta_2/2) - (1-\gamma_2)\sigma}{\gamma_2} \geq \frac{(\delta_2/2) - \sigma}{\gamma_2} \\ &> \quad \frac{1-R-2\sigma}{2\gamma_2} \geq \frac{\epsilon}{2\gamma_2} , \end{aligned} \tag{20}$$

where γ_2 stands for $\gamma_{\mathcal{G}_2}$ and the last inequality follows from (19). Now, by (17) we have

$$\Delta_2 = \frac{(1-r_0)\Delta_1}{r_{\mathrm{m}} R} > \frac{\epsilon \Delta_1}{R} \geq \frac{\alpha_R}{R \cdot \epsilon^2} \geq \frac{8(1-R)}{\mu \cdot \epsilon^2} ,$$

from which we get the following upper bound on the square of γ_2:

$$\gamma_2^2 \leq \frac{4(\Delta_2 - 1)}{\Delta_2^2} < \frac{4}{\Delta_2} \leq \frac{\mu \cdot \epsilon^2}{2(1-R)} .$$

Combining this bound with (20) yields

$$\frac{\sigma}{\tau} > \frac{1-R}{2\mu} ,$$

namely, $\tau < 2\mu\sigma/(1-R) < \mu$. □

It follows from Lemma 5.1 that Step (D2) reduces the number of errors in $\tilde{\boldsymbol{w}}$ to the extent that allows a linear-time decoder of $\mathcal{C}_{\mathrm{m}}$ to fully recover the errors in $\tilde{\boldsymbol{w}}$ in Step (D3). Hence, the list $\hat{\boldsymbol{s}}$, which is computed in Step (D3), is identical with the list $\boldsymbol{s}$ that was originally encoded in Step (E2).

Finally, to show that Step (D4) yields complete recovery from errors, we apply Theorem 4.1 to the parameters of the code $(\mathcal{G}_1, \mathcal{C}_0 : \mathcal{C}_1)$. Here $\theta_0 = \kappa \cdot \epsilon$ and

$$\gamma_1 = \gamma_{\mathcal{G}_1} < \frac{2}{\sqrt{\Delta_1}} \leq \frac{2\epsilon^{3/2}}{\sqrt{\alpha_R}} \leq \frac{\epsilon^{3/2}}{2\sqrt{(1-R)/\kappa}} ;$$

therefore,

$$\beta = \frac{(\delta_1/2) - \gamma_1\sqrt{\delta_1/\theta_0}}{1-\gamma_1} > \frac{1-R}{2} - \gamma_1\sqrt{\frac{1-R}{\theta_0}} > \frac{1-R-\epsilon}{2}$$

and, so, by (19), the conditions of Theorem 4.1 hold for $\sigma = (1-R-\epsilon)/2$ (note that $\beta > 0$ yields $\sqrt{\theta_0\delta_1} > 2\gamma_1$, thus (7) holds).

Appendix A

We provide here the proofs of Lemmas 3.2 and 4.2.

Given a bipartite graph $\mathcal{G} = (V' : V'', E)$, we associate with $\mathcal{G}$ a $|V'| \times |V''|$ real matrix $X_{\mathcal{G}}$ whose rows and columns are indexed by V' and V'', respectively, and $(X_{\mathcal{G}})_{u,v} = 1$ if and only if $\{u, v\} \in E$. With a proper ordering on $V' \cup V''$, the matrix $X_{\mathcal{G}}$ is related to the adjacency matrix of $\mathcal{G}$ by

$$A_{\mathcal{G}} = \left(\begin{array}{c|c} 0 & X_{\mathcal{G}} \\ \hline X_{\mathcal{G}}^T & 0 \end{array} \right) . \tag{21}$$

Lemma A.1: Let $\mathcal{G} = (V' : V'', E)$ be a bipartite Δ-regular graph where $|V'| > 1$. Then Δ^2 is the largest eigenvalue of the (symmetric) matrix $X_{\mathcal{G}}^T X_{\mathcal{G}}$ and the all-one vector $\mathbf{1}$ is a corresponding eigenvector. The second largest eigenvalue of $X_{\mathcal{G}}^T X_{\mathcal{G}}$ is $\gamma_{\mathcal{G}}^2 \Delta^2$.

Proof: We compute the square of $A_{\mathcal{G}}$,

$$A_{\mathcal{G}}^2 = \left(\begin{array}{c|c} X_{\mathcal{G}} X_{\mathcal{G}}^T & 0 \\ \hline 0 & X_{\mathcal{G}}^T X_{\mathcal{G}} \end{array} \right) ,$$

and recall the following two known facts:

(i) $X_{\mathcal{G}} X_{\mathcal{G}}^T$ and $X_{\mathcal{G}}^T X_{\mathcal{G}}$ have the same set of eigenvalues, each with the same multiplicity [14, Theorem 16.2].

(ii) If λ is an eigenvalue of $A_{\mathcal{G}}$, then so is $-\lambda$, with the same multiplicity [8, Proposition 1.1.4].

We conclude that λ is an eigenvalue of $A_{\mathcal{G}}$ if and only if λ^2 is an eigenvalue $X_{\mathcal{G}}^T X_{\mathcal{G}}$; furthermore, when $\lambda \neq 0$, both these eigenvalues have the same multiplicities in their respective matrices. The result readily follows. □

For real column vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^m$, let $\langle \boldsymbol{x}, \boldsymbol{y} \rangle$ be the scalar product $\boldsymbol{x}^T \boldsymbol{y}$ and $\|\boldsymbol{x}\|$ be the norm $\sqrt{\langle \boldsymbol{x}, \boldsymbol{x} \rangle}$.

Lemma A.2: Let $\mathcal{G} = (V' : V'', E)$ be a bipartite Δ-regular graph where $|V'| = n > 1$ and let $\boldsymbol{s} = (s_u)_{u \in V'}$ and $\boldsymbol{t} = (t_u)_{u \in V''}$ be two column vectors in $\mathbb{R}^n$. Denote by σ and τ the averages

$$\sigma = \frac{1}{n} \sum_{u \in V'} s_u \qquad \text{and} \qquad \tau = \frac{1}{n} \sum_{u \in V''} t_u ,$$

and let the column vectors $\boldsymbol{y}$ and $\boldsymbol{z}$ in $\mathbb{R}^n$ be given by

$$\boldsymbol{y} = \boldsymbol{s} - \sigma \cdot \mathbf{1} \qquad \text{and} \qquad \boldsymbol{z} = \boldsymbol{t} - \tau \cdot \mathbf{1} \; .$$

Define the vector $\boldsymbol{x} \in \mathbb{R}^{2n}$ by

$$\boldsymbol{x} = \begin{pmatrix} \boldsymbol{s} \\ \boldsymbol{t} \end{pmatrix} .$$

Then,

$$|\langle \boldsymbol{x}, A_{\mathcal{G}} \boldsymbol{x} \rangle - 2\sigma\tau\Delta n| \le 2\gamma_{\mathcal{G}} \Delta \|\boldsymbol{y}\| \cdot \|\boldsymbol{z}\| \; .$$

Proof: First, it is easy to see that $X_{\mathcal{G}} \mathbf{1} = X_{\mathcal{G}}^T \mathbf{1} = \Delta \cdot \mathbf{1}$ and that $\langle \boldsymbol{y}, \mathbf{1} \rangle = \langle \boldsymbol{z}, \mathbf{1} \rangle = 0$; these equalities, in turn, yield the relationship:

$$\langle \boldsymbol{y}, X_{\mathcal{G}} \boldsymbol{z} \rangle = \langle \boldsymbol{s}, X_{\mathcal{G}} \boldsymbol{t} \rangle - \sigma\tau\Delta n \; .$$

Secondly, from (21) we get that

$$\langle \boldsymbol{x}, A_{\mathcal{G}} \boldsymbol{x} \rangle = 2 \langle \boldsymbol{s}, X_{\mathcal{G}} \boldsymbol{t} \rangle \; .$$

Hence, the lemma will be proved once we show that

$$|\langle \boldsymbol{y}, X_{\mathcal{G}} \boldsymbol{z} \rangle| \le \gamma_{\mathcal{G}} \Delta \|\boldsymbol{y}\| \cdot \|\boldsymbol{z}\| \; . \tag{22}$$

Let

$$\lambda_1 \ge \lambda_2 \ge \ldots \ge \lambda_n$$

be the eigenvalues of $X_{\mathcal{G}}^T X_{\mathcal{G}}$ and let $\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_n$ be corresponding orthonormal eigenvectors where, by Lemma A.1,

$$\lambda_1 = \Delta^2 \; , \qquad \lambda_2 = \gamma_{\mathcal{G}}^2 \Delta^2 \; , \qquad \text{and} \qquad \boldsymbol{v}_1 = (1/\sqrt{n}) \cdot \mathbf{1} \; .$$

Write

$$\boldsymbol{z} = \sum_{i=1}^{n} \beta_i \boldsymbol{v}_i \; ,$$

where $\beta_i = \langle \boldsymbol{z}, \boldsymbol{v}_i \rangle$. Recall, however, that $\beta_1 = (1/\sqrt{n}) \cdot \langle \boldsymbol{z}, \mathbf{1} \rangle = 0$; so,

$$\begin{aligned} \|X_{\mathcal{G}} \boldsymbol{z}\|^2 &= \langle \boldsymbol{z}, X_{\mathcal{G}}^T X_{\mathcal{G}} \boldsymbol{z} \rangle \\ &= \Big\langle \sum_{i=2}^{n} \beta_i \boldsymbol{v}_i, \sum_{i=2}^{n} \lambda_i \beta_i \boldsymbol{v}_i \Big\rangle = \sum_{i=2}^{n} \lambda_i \beta_i^2 \|\boldsymbol{v}_i\|^2 \\ &\le \lambda_2 \sum_{i=2}^{n} \beta_i^2 = \lambda_2 \|\boldsymbol{z}\|^2 = \gamma_{\mathcal{G}}^2 \Delta^2 \|\boldsymbol{z}\|^2 \; . \end{aligned}$$

The desired result (22) is now obtained from the Cauchy-Schwartz inequality. $\square$

Lemma A.3: Let $\mathcal{G} = (V' : V'', E)$ be a bipartite Δ-regular graph where $|V'| = n > 1$ and let $\chi : (V' \cup V'') \to \mathbb{R}$ be a function on the vertices of $\mathcal{G}$. Define the function $w : E \to \mathbb{R}$ and the average $\mathsf{E}_{\mathcal{G}}\{w\}$ by

$$w(e) = \chi(u)\chi(v) \quad \text{for every edge } e = \{u, v\} \text{ in } \mathcal{G}$$

and

$$\mathsf{E}_{\mathcal{G}}\{w\} = \frac{1}{\Delta n} \sum_{e \in E} w(e) \; .$$

Then

$$\left| \mathsf{E}_{\mathcal{G}}\{w\} - \mathsf{E}'_{\mathcal{G}}\{\chi\} \cdot \mathsf{E}''_{\mathcal{G}}\{\chi\} \right| \le \gamma_{\mathcal{G}} \sqrt{\mathsf{Var}'_{\mathcal{G}}\{\chi\} \cdot \mathsf{Var}''_{\mathcal{G}}\{\chi\}} \; ,$$

where

$$\mathsf{E}'_{\mathcal{G}}\{\chi^i\} = \frac{1}{n} \sum_{u \in V'} (\chi(u)^i) \; ,$$

$$\mathsf{E}''_{\mathcal{G}}\{\chi^i\} = \frac{1}{n} \sum_{u \in V''} (\chi(u)^i) \; ,$$

$$\mathsf{Var}'_{\mathcal{G}}\{\chi\} = \mathsf{E}'_{\mathcal{G}}\{\chi^2\} - (\mathsf{E}'_{\mathcal{G}}\{\chi\})^2 \; ,$$

and

$$\mathsf{Var}''_{\mathcal{G}}\{\chi\} = \mathsf{E}''_{\mathcal{G}}\{\chi^2\} - (\mathsf{E}''_{\mathcal{G}}\{\chi\})^2 \; .$$

Proof: Define the column vectors

$$\boldsymbol{s} = (\chi(u))_{u \in V'} \; , \quad \boldsymbol{t} = (\chi(u))_{u \in V''} \; ,$$

and

$$\boldsymbol{x} = \begin{pmatrix} \boldsymbol{s} \\ \boldsymbol{t} \end{pmatrix} ,$$

and denote by σ and τ the averages

$$\sigma = \frac{1}{n} \sum_{u \in V'} s_u \qquad \text{and} \qquad \tau = \frac{1}{n} \sum_{u \in V''} t_u \; .$$

The following equalities are easily verified:

$$\mathsf{E}_{\mathcal{G}}\{w\} = \frac{\langle \boldsymbol{x}, A_{\mathcal{G}} \boldsymbol{x} \rangle}{2\Delta n} \; ,$$

$$\mathsf{E}'_{\mathcal{G}}\{\chi\} = \sigma \; , \qquad \mathsf{E}''_{\mathcal{G}}\{\chi\} = \tau \; ,$$

$$\mathsf{Var}'_{\mathcal{G}}\{\chi\} = \frac{1}{n} \cdot \|\boldsymbol{s} - \sigma \cdot \mathbf{1}\|^2 \; ,$$

and

$$\mathsf{Var}''_{\mathcal{G}}\{\chi\} = \frac{1}{n} \cdot \|\boldsymbol{t} - \tau \cdot \mathbf{1}\|^2 \; .$$

The result now follows from Lemma A.2. $\square$

Proof of Lemma 3.2: Using the notation of Lemma A.3, write

$$\mathsf{E}_{\mathcal{G}}\{w\} = \frac{1}{\Delta n} \sum_{u \in V'} \sum_{v \in \mathcal{N}(u)} \chi(u)\chi(v) \; , \tag{23}$$

$$\mathsf{E}'_{\mathcal{G}}\{\chi\} = \frac{1}{n} \sum_{u \in V'} \chi(u) = \sigma \; , \tag{24}$$

and

$$\mathsf{E}''_{\mathcal{G}}\{\chi\} = \frac{1}{n} \sum_{u \in V''} \chi(u) = \tau \; . \tag{25}$$

Since the range of χ is restricted to the interval $[0, 1]$, we have

$$\mathsf{E}'_{\mathcal{G}}\{\chi^2\} \le \mathsf{E}'_{\mathcal{G}}\{\chi\} \qquad \text{and} \qquad \mathsf{E}''_{\mathcal{G}}\{\chi^2\} \le \mathsf{E}''_{\mathcal{G}}\{\chi\} \; ;$$

hence, the values $\mathsf{Var}'_{\mathcal{G}}\{\chi\}$ and $\mathsf{Var}''_{\mathcal{G}}\{\chi\}$ can be bounded from above by

$$\mathsf{Var}'_{\mathcal{G}}\{\chi\} \le \sigma - \sigma^2 \qquad \text{and} \qquad \mathsf{Var}''_{\mathcal{G}}\{\chi\} \le \tau - \tau^2 \; . \tag{26}$$

Substituting (23)–(26) into Lemma A.3 yields

$$\left| \frac{1}{\Delta n} \Big(\sum_{u \in V'} \sum_{v \in \mathcal{N}(u)} \chi(u)\chi(v) \Big) - \sigma\tau \right| \le \gamma_{\mathcal{G}} \sqrt{\sigma(1-\sigma)\tau(1-\tau)} \; ;$$

so,

$$\frac{1}{\Delta n}\sum_{u\in V'}\sum_{v\in\mathcal{N}(u)}\chi(u)\chi(v)$$
$$\begin{aligned}&\le \sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma(1-\sigma)\tau(1-\tau)}\\&= (1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}\left(\sqrt{\sigma\tau}+\sqrt{(1-\sigma)(1-\tau)}\right)\\&\le (1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}\ ,\end{aligned}$$

as claimed. □

Proof of Lemma 4.2: We compute lower and upper bounds on the average

$$\frac{1}{\Delta n}\sum_{v\in V''}\sum_{u\in\mathcal{N}(v)}\chi(u)\chi(v)\ .$$

On the one hand, this average equals

$$\frac{1}{\Delta n}\sum_{\substack{v\in V'':\\ \chi(v)>0}}\chi(v)\underbrace{\sum_{u\in\mathcal{N}(v)}\chi(u)}_{\ge\delta\Delta/2}\ge\frac{1}{\Delta n}\cdot\frac{\delta\Delta}{2}\underbrace{\sum_{v\in V''}\chi(v)}_{\tau n}=\frac{\delta\tau}{2}\ ,$$

where the inequality follows from the assumed conditions on χ. On the other hand, this average also equals

$$\frac{1}{\Delta n}\sum_{u\in V'}\sum_{v\in\mathcal{N}(u)}\chi(u)\chi(v)\le(1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}\ ,$$

where the inequality follows from Lemma 3.2. Combining these two bounds we get

$$\frac{\delta\tau}{2}\le(1-\gamma_{\mathcal{G}})\sigma\tau+\gamma_{\mathcal{G}}\sqrt{\sigma\tau}\ ,$$

and the result is now obtained by dividing by $\gamma_{\mathcal{G}}\tau$ and rearranging terms. □

Appendix B

When analyzing the complexity of the algorithm in Figure 1, one can notice that the decoder $\mathcal{D}\in\{\mathcal{D}',\mathcal{D}''\}$ needs to be applied at vertex u, only if $(\boldsymbol{z})_{E(u)}$ has been modified since the last application of $\mathcal{D}$ at that vertex. Based on this observation, we prove the following lemma.

Lemma B.1: The number of (actual) applications of the decoders $\mathcal{D}'$ and $\mathcal{D}''$ in the algorithm in Figure 1 can be bounded from above by $\omega\cdot n$, where

$$\omega=2\cdot\left\lceil\frac{\log\left(\dfrac{\Delta\beta\sqrt{\sigma}}{\beta-\sigma}\right)}{\log\left(\dfrac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)}\right\rceil+\frac{1+\dfrac{\theta}{\delta}}{1-\left(\dfrac{4\gamma_{\mathcal{G}}^2}{\theta\delta}\right)^2}\ .$$

Proof: Define i_T by

$$i_T\ =\ 2\cdot\left\lceil\frac{\log\left(\dfrac{\Delta\beta\sqrt{\sigma}}{\beta-\sigma}\right)}{\log\left(\dfrac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)}\right\rceil\ .$$

It is easy to verify that

$$\left(\frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)^{i_T/2}\left(\frac{1}{\sqrt{\sigma}}-\frac{\sqrt{\sigma}}{\beta}\right)\ \ge\ \Delta\ . \qquad (27)$$

In the first i_T iterations in Figure 1, we apply the decoder $\mathcal{D}$ (which is either $\mathcal{D}'$ or $\mathcal{D}''$) at most $i_T\cdot n$ times.

Next, we evaluate the total number of applications of the decoder $\mathcal{D}$ in iterations $i=i_T+1,i_T+2,\cdots,\nu$. We hereafter use the notations U_i and S_i as in the proof of Theorem 4.1. Recall that we need to apply the decoder $\mathcal{D}$ to $(\boldsymbol{z})_{E(u)}$ for a vertex $u\in U_{i+2}$, only if at least one entry in $(\boldsymbol{z})_{E(u)}$ — say, the one that is indexed by the edge $\{u,v\}\in E(u)$ — has been altered during iteration $i+1$. Such an alteration may occur only if v is a vertex in U_{i+1} with an adjacent vertex in S_i. We conclude that $\mathcal{D}$ needs to be applied at vertex u during iteration $i+2$ only if $u\in\mathcal{N}(\mathcal{N}(S_i))$. The number of such vertices u, in turn, is at most $\Delta^2|S_i|=\Delta^2\cdot\sigma_i n$.

We now sum the values of $\Delta^2\sigma_i n$ over iterations $i=i_T+1,i_T+2,\cdots,\nu$:

$$\begin{aligned}&\Delta^2 n\cdot\sum_{i=i_T+1}^{\nu}\sigma_i\\&\quad=\ \Delta^2 n\left(\sum_{j=i_T/2}^{\lfloor(\nu-1)/2\rfloor}\sigma_{2j+1}+\sum_{j=i_T/2}^{\lfloor(\nu-2)/2\rfloor}\sigma_{2j+2}\right)\\&\quad\le\ \Delta^2 n\cdot\sum_{j=i_T/2}^{\lfloor(\nu-1)/2\rfloor}\sigma_{2j+1}\left(1+\frac{\theta}{\delta}\right)\ ,\end{aligned}\qquad (28)$$

where the last inequality is due to (11).

From (15) (and by neglecting a positive term), we obtain

$$\frac{1}{\sqrt{\sigma_{i+1}}}\ \ge\ \left(\frac{\theta\delta}{4\gamma_{\mathcal{G}}^2}\right)^{i/2}\left(\frac{1}{\sqrt{\sigma}}-\frac{\sqrt{\sigma}}{\beta}\right)$$

for even $i\ge i_T$. Therefore, the expression in (28) is bounded from above by

$$\begin{aligned}&\frac{\Delta^2 n\left(1+\dfrac{\theta}{\delta}\right)\cdot\left(\dfrac{4\gamma_{\mathcal{G}}^2}{\theta\delta}\right)^{i_T}}{\left(1-\left(\dfrac{4\gamma_{\mathcal{G}}^2}{\theta\delta}\right)^2\right)\left(\dfrac{1}{\sqrt{\sigma}}-\dfrac{\sqrt{\sigma}}{\beta}\right)^2}\\&\qquad\le\ \frac{\Delta^2 n\left(1+\dfrac{\theta}{\delta}\right)\cdot\dfrac{1}{\Delta^2}}{1-\left(\dfrac{4\gamma_{\mathcal{G}}^2}{\theta\delta}\right)^2}\\&\qquad=\ \frac{n\left(1+\dfrac{\theta}{\delta}\right)}{1-\left(\dfrac{4\gamma_{\mathcal{G}}^2}{\theta\delta}\right)^2}\ ,\end{aligned}$$

where the inequality follows from (27).

Adding now the number of applications of the decoder $\mathcal{D}$ during the first i_T iterations, we conclude that the total number of applications of the decoder $\mathcal{D}$ is at most $\omega\cdot n$, where

$$\omega=i_T+\frac{1+\dfrac{\theta}{\delta}}{1-\left(\dfrac{4\gamma_{\mathcal{G}}^2}{\theta\delta}\right)^2}\ .$$

□

REFERENCES

[1] N. Alon, J. Bruck, J. Naor, M. Naor, and R.M. Roth, "Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs," *IEEE Trans. Inf. Theory,* vol. 38, no. 2, pp. 509–516, Mar. 1992.

[2] N. Alon and J.H. Spencer, *The Probabilistic Method,* 2nd ed. New York: Wiley, 2000.

[3] N. Alon, J. Edmonds, and M. Luby, "Linear time erasure codes with nearly optimal recovery," in *Proc. 36th Annual IEEE Symp. on Foundations of Computer Science (FOCS),* Milwaukee, Wisconsin, Oct. 1995, pp. 512–519.

[4] N. Alon and M. Luby, "A linear time erasure-resilient code with nearly optimal recovery," *IEEE Trans. Inf. Theory,* vol. 42, no. 6, pp. 1732–1736, Nov. 1996.

[5] A. Barg and G. Zémor, "Error exponents of expander codes," *IEEE Trans. Inf. Theory,* vol. 48, no. 6, pp. 1725–1729, June 2002.

[6] A. Barg and G. Zémor, "Error exponents of expander codes under linear-complexity decoding," *SIAM J. Discrete Math.,* vol. 17, no. 3, pp. 426–445, 2004.

[7] A. Barg and G. Zémor, "Concatenated codes: serial and parallel," *IEEE Trans. Inf. Theory.*, vol. 51, no. 5, pp. 1625–1634, May 2005.

[8] G. Davidoff, P. Sarnak, and A. Valette, *Elementary Number Theory, Group Theory, and Ramanujan Graphs.* Cambridge, UK: Cambridge University Press, 2003.

[9] I. Dumer, "Concatenated codes and their multilevel generalizations," in *Handbook of Coding Theory, Volume II,* V.S. Pless and W.C. Huffman (Editors), Amsterdam, Netherlands: North-Holland, pp. 1911–1988, 1998.

[10] G.D. Forney, Jr., *Concatenated Codes.* Cambridge, Massachusetts: MIT Press, 1966.

[11] G.D. Forney, Jr., "Generalized minimum distance decoding," *IEEE Trans. Inf. Theory,* vol. 12, no. 2, pp. 125–131, Apr. 1966.

[12] V. Guruswami and P. Indyk, "Near-optimal linear-time codes for unique decoding and new list-decodable codes over smaller alphabets," in *Proc. 34th Annual ACM Symposium on Theory of Computing (STOC),* Montréal, Quebec, Canada, pp. 812–821, May 2002.

[13] A. Lubotsky, R. Philips, and P. Sarnak, "Ramanujan graphs," *Combinatorica,* vol. 8, pp. 261–277, 1988.

[14] C.C. MacDuffee, *The Theory of Matrices.* New York: Chelsea, 1946.

[15] G.A. Margulis, "Explicit group theoretical constructions of combinatorial schemes and their applications to the design of expanders and concentrators," *Probl. Inf. Transm.,* vol. 24, no. 1, pp. 39–46, July 1988.

[16] M. Sipser and D.A. Spielman, "Expander codes," *IEEE Trans. Inf. Theory,* vol. 42, no. 6, pp. 1710–1722, Nov. 1996.

[17] V. Skachek and R.M. Roth, "Generalized minimum distance iterative decoding of expander codes," in *Proc. IEEE Inf. Theory Workshop (ITW),* Paris, France, pp. 245–248, Mar.-Apr. 2003.

[18] D.A. Spielman, "Linear-time encodable and decodable error-correcting codes," *IEEE Trans. Inf. Theory,* vol. 42, no. 6, pp. 1723–1731, Nov. 1996.

[19] R.M. Tanner, "A recursive approach to low-complexity codes," *IEEE Trans. Inf. Theory,* vol. 27, no. 5, pp. 533–547, Sep. 1981.

[20] G. Zémor, "On expander codes," *IEEE Trans. Inf. Theory,* vol. 47, no. 2, pp. 835–837, Feb. 2001.