Generalized Minimum Distance Iterative Decoding of Expander Codes[1]

Vitaly Skachek Ron M. Roth
Computer Science Department
Technion
Haifa 32000, Israel
{vitalys, ronny}@cs.technion.ac.il.

Abstract — **Recently, Zémor proposed an improvement on the Sipser-Spielman analysis of expander codes and presented a linear-time iterative decoder that can correct a number of errors up to approximately $1/4$ the known lower bound on the minimum distance of the code. In this work, we propose an improvement on Zémor's decoder for $F = \mathrm{GF}(2)$, with the number of correctable errors becoming close to half the lower bound on the minimum distance. The improvement is obtained by inserting into the decoding algorithm features akin to generalized minimum distance decoding of concatenated codes.**

I. Introduction

In [11], Zémor considered a special class of the Sipser-Spielman construction of expander codes, where the underlying graph is bipartite [8], [9]. We summarize the construction next.

Let $\mathcal{G} = (V, E)$ be a Δ-regular bipartite undirected graph with a vertex set $V = A \cup B$ such that $A \cap B = \emptyset$ and $|A| = |B| = n$, and an edge set E of size $N = n\Delta$ such that every edge in E has one endpoint in A and one endpoint in B. For every vertex $v \in V$, denote by $E(v)$ the set of edges incident with v. We assume an ordering on V, thereby inducing an ordering on the edges of $E(v)$ for every $v \in V$. Let $F = \mathrm{GF}(2)$, and for a word $\boldsymbol{x} = (x_e)_{e \in E}$ (whose entries are indexed by E) in F^N, denote by $(\boldsymbol{x})_v$ the sub-word of $\boldsymbol{x}$ that is indexed by $E(v)$. Observe that the subset A (respectively, B) induces on every word $\boldsymbol{x} \in F^N$ a partition into n non-overlapping sub-words $(\boldsymbol{x})_v \in F^\Delta$, where v ranges over the elements of A (respectively, B).

Fix $\mathcal{C}$ to be a linear $[\Delta, k{=}r\Delta, d{=}\delta\Delta]$ code over F. The code $\mathbb{C} = (\mathcal{G}, \mathcal{C})$ is defined in [11] to be the following linear $[N, K, D]$ code over F:

$$\mathbb{C} = \left\{ \boldsymbol{c} \in F^N \;:\; (\boldsymbol{c})_v \in \mathcal{C} \text{ for every } v \in V \right\} .$$

It can be readily verified that $K/N \geq 2r{-}1$ [9], and it has been shown in [8] that

$$D \geq N \left(\frac{\delta - (\lambda/\Delta)}{1 - (\lambda/\Delta)} \right)^2 = N \left(\delta^2 - O(\lambda/\Delta) \right) , \qquad (1)$$

where $\lambda = \lambda(\mathcal{G})$ is the second largest eigenvalue of the adjacency matrix of $\mathcal{G}$ (the largest eigenvalue is Δ, and when writing the $O(\cdot)$ expression we have assumed that λ/Δ is bounded away from 1). For any Δ such that $\Delta{-}1$ is an odd prime, there are known explicit constructions of sequences of Δ-regular bipartite graphs with arbitrarily large number of vertices such that each graph $\mathcal{G}$ in the sequence is a *Ramanujan graph*, namely, it satisfies the inequality $\lambda(\mathcal{G}) \leq 2\sqrt{\Delta{-}1}$ [6]. (In particular, $\mathcal{G}$ is an *expander graph*, as the separation between the eigenvalues Δ and λ implies certain expansion properties on $\mathcal{G}$ [10]. Constructions of bipartite Ramanujan graphs other than those in [6] are also known.) If the graph $\mathcal{G}$ in the construction of $\mathbb{C}$ is taken to be a Ramanujan graph, we obtain from (1) that the $O(\cdot)$ expression goes to zero as Δ becomes large.

The iterative decoding algorithm of Zémor is shown in Figure 1, where $\mathcal{D} : F^\Delta \to \mathcal{C}$ stands for a decoder for $\mathcal{C}$ that recovers correctly any pattern of less than $d/2$ errors. It is shown in [11] that the algorithm in Figure 1 can correct any error word whose Hamming weight does not exceed

$$\frac{1}{2} \cdot \alpha N \delta \left(\frac{\delta}{2} - \frac{\lambda}{\Delta} \right) = \frac{1}{4} \cdot \alpha N \left(\delta^2 - O(\lambda/\Delta) \right) ,$$

for any fixed positive constant $\alpha < 1$. The number of iterations m in Figure 1 can be taken as $\lfloor (\log n)/\log(2{-}\alpha) \rfloor$. Using similar arguments as in [8], it can be shown that Zémor's algorithm can be implemented in time complexity $O(N)$, assuming that α and the code $\mathcal{C}$ are fixed (in particular, this assumption implies that $\mathcal{D}$ can be implemented in constant time).

Input: Received word $\boldsymbol{y} = (y_e)_{e \in E}$.
Let $\boldsymbol{z} \leftarrow \boldsymbol{y}$.
For $t \leftarrow 1$ **to** m **do** {
  Let X stand for A if t is odd, and for B otherwise.
  Iteration t**:** For every $v \in X$ let $(\boldsymbol{z})_v \leftarrow \mathcal{D}((\boldsymbol{z})_v)$.
}
Output: $\boldsymbol{z}$.

Figure 1: Zémor's decoder.

In this work, we modify Zémor's algorithm so that it recovers correctly any error pattern with at most

$$\begin{aligned} J_{\mathbb{C}}(\alpha) &= \frac{1}{2} \cdot \alpha N \delta \left(\delta - \frac{2\lambda}{\Delta} \right) \\ &= \frac{1}{2} \cdot \alpha N \left(\delta^2 - O(\lambda/\Delta) \right) \qquad (2) \end{aligned}$$

[1] This work was supported by Grant No. 94/99 from the Israel Science Foundation.

errors, for any fixed constant $\alpha < 1$, while still maintaining a linear decoding time. Note that (2) reaches approximately half the lower bound (1) on the minimum distance of $\mathbb{C}$. Our improvement is achieved by borrowing ideas from generalized minimum distance decoding of concatenated codes [2]. Specifically, if the number of errors found in a sub-block $(\boldsymbol{z})_v$ during the first iteration exceeds a given threshold, then the entries of that sub-block are regarded as erasures during the next iteration. Up to $\lceil d/2 \rceil$ $(= O(1))$ different threshold values are checked, until the right one is found.

More recently, Barg and Zémor have generalized in [1] the construction $\mathbb{C}$ by replacing the code $\mathcal{C}$ with two different codes—one for each side of the bipartite graph; in addition, they allow the graph to contain parallel edges. This generalization results in a code construction with a larger lower bound on the minimum distance. Whether our method of GMD-like decoding can be applied to the Barg-Zémor generalization is still under investigation.

It is worth comparing our results with the recent papers by Guruswami and Indyk [3]–[5], where they show a construction of a sequence of linear-time decodable (and encodable) binary codes that approach the Zyablov bound, with the number of correctable errors approaching half the designed minimum distance of the code. Their construction uses a (conventional) concatenation of an outer expander code with an inner binary code that attains the Gilbert-Varshamov bound. For a given rate $R = K/N$, their lower bound, J_{GI}, on the number of correctable errors satisfies

$$\frac{J_{\mathrm{GI}}}{N} \simeq \frac{1}{2} \max_{R \le r \le 1} \left(1 - \frac{R}{r}\right) \cdot \mathsf{H}^{-1}(1-r) ,$$

where $\mathsf{H}^{-1} : [0,1] \to [0, \frac{1}{2}]$ is the inverse of the binary entropy function $\mathsf{H}(x) = -x \log_2 x - (1-x) \log_2 (1-x)$. This bound is higher than the lower bound

$$\frac{J_{\mathbb{C}}(1)}{N} \simeq \frac{1}{2} \left(\mathsf{H}^{-1}\left(\tfrac{1}{2}(1-R)\right)\right)^2$$

obtained from (2) (the difference between the two bounds becomes significant in the low-rate range). Yet, the use of concatenation in the Guruswami-Indyk construction makes the multiplying constants in the $O(N)$ complexity expression much larger.

II. Decoding algorithm

Our improved decoding algorithm is presented in Figure 2. We next describe the changes made compared to the algorithm in Figure 1 and explain the new notations used in Figure 2. Hereafter, $\mathsf{d}(\cdot, \cdot)$ stands for Hamming distance over F.

1. Iteration 0 in Figure 2 is the same as Iteration 1 in Figure 1, except that the result of the decoding is stored in an intermediate word $\boldsymbol{x} \in F^N$.
2. The remaining iterations are performed within a loop that enumerates over a threshold ϑ that takes values on $\{1, 2, \ldots, \lceil d/2 \rceil\}$.
3. Iteration 1 in Figure 2 modifies the result of Iteration 0, depending on the threshold ϑ. Specifically, for every vertex $v \in A$, the sub-vector $(\boldsymbol{z})_v$ is set to what Iteration 1 in Figure 1 would have set, whenever the latter iteration produces a sub-block whose Hamming distance from $(\boldsymbol{y})_v$ is less than ϑ. Otherwise, $(\boldsymbol{z})_v$ is set to a word $\boldsymbol{\varepsilon}$ of length Δ whose entries are all equal to a special 'erasure' symbol (not in F), denoted '?'.
4. Iteration 2 differs from its counterpart in Figure 1 in that it uses a combined error-erasure decoder $\widetilde{\mathcal{D}} : (F \cup \{?\})^\Delta \to \mathcal{C}$, instead of the decoder $\mathcal{D}$. It is assumed that $\widetilde{\mathcal{D}}$ recovers correctly any pattern of τ errors and ρ erasures, provided that $2\tau + \rho < d$ (such a decoder always exists).
5. Iterations 3 through m in Figure 2 are identical to those in Figure 1.
6. The decoding terminates when the result of Iteration m (in any of the enumeration steps over ϑ) is a word $\boldsymbol{z}$ whose Hamming distance from $\boldsymbol{y}$ is at most the value $J_{\mathbb{C}}(\alpha)$ given by (2); here α is any fixed positive constant smaller than 1, and the value m can be taken as $\lfloor (\log n)/\log(2-\alpha) \rfloor + 1$.

Input: Received word $\boldsymbol{y} = (y_e)_{e \in E}$.

Iteration 0: For every $v \in A$ let $(\boldsymbol{x})_v \leftarrow \mathcal{D}((\boldsymbol{y})_v)$.

For $\vartheta = 1$ **to** $\lceil d/2 \rceil$ **do** {

    Iteration 1: For every $v \in A$ let
$$(\boldsymbol{z})_v = \begin{cases} (\boldsymbol{x})_v & \text{if } \mathsf{d}((\boldsymbol{y})_v, (\boldsymbol{x})_v) < \vartheta \\ \boldsymbol{\varepsilon} & \text{otherwise} \end{cases} .$$

    Iteration 2: For every $v \in B$ let $(\boldsymbol{z})_v \leftarrow \widetilde{\mathcal{D}}((\boldsymbol{z})_v)$.

    For $t \leftarrow 3$ **to** m **do** {

        Let X stand for A if t is odd, and for B otherwise.

        Iteration t: For every $v \in X$ let $(\boldsymbol{z})_v \leftarrow \mathcal{D}((\boldsymbol{z})_v)$.

    }

    If $\mathsf{d}(\boldsymbol{y}, \boldsymbol{z}) \le J_{\mathbb{C}}(\alpha)$ break.

}

Output: $\boldsymbol{z}$.

Figure 2: Improved decoder.

We next state our result.

Theorem 1 *Suppose that $d \ge 3\lambda$ and let $J_{\mathbb{C}}(\alpha)$ be given by (2) for some positive $\alpha < 1$. Then the algorithm in Figure 2 recovers correctly any pattern of up to $\lfloor J_{\mathbb{C}}(\alpha) \rfloor$ errors.*

As is the case with the decoding algorithms of Sipser-Spielman and Zémor, the algorithm in Figure 2 can be implemented in time complexity that is linear in N, where the multiplying constant depends on α and on the code $\mathcal{C}$. Note that the enumeration over the threshold ϑ does introduce a slow-down by a constant factor of $\lceil d/2 \rceil$, compared to the algorithm in Figure 1.

An analysis of the algorithm, which leads to the proof of Theorem 1, is given in Section III.

III. Analysis of the algorithm

Throughout this section, we let the code $\mathbb{C} = (\mathcal{G}, \mathcal{C})$ be as in Section I and fix the transmitted codeword $\boldsymbol{c} = (c_e)_{e \in E}$ in $\mathbb{C}$ and the received word $\boldsymbol{y} = (y_e)_{e \in E}$ in F^N.

We say that an edge $e \in E$ is *corrupted* at a given time during the execution of the algorithm (in Figure 1 or 2) if the respective entry in $\boldsymbol{z}$ is in F yet differs from c_e. A vertex $v \in V$ is corrupted if $E(v)$ contains a corrupted edge. The core of the analysis in [11] is showing an upper bound on the number of corrupted vertices that decreases by a constant factor with each iteration. We will apply a similar technique here, treating separately Iterations 1 and 2 where the two algorithms differ.

For $\vartheta = 1, 2, \ldots, \lceil d/2 \rceil$, denote by $\boldsymbol{z}(\vartheta) = (z_e)_{e \in E}$ (respectively, $\boldsymbol{w}(\vartheta) = (w_e)_{e \in E}$) the result $\boldsymbol{z}$ of Iteration 1 (respectively, 2) while the external loop assumes a threshold ϑ. We let $Z(\vartheta)$ and $W(\vartheta)$ be the respective subsets of corrupted edges, i.e.,

$$\begin{aligned} Z(\vartheta) &= \{ e \in E \,:\, z_e \in F \text{ and } z_e \neq c_e \} \;, \\ W(\vartheta) &= \{ e \in E \,:\, w_e \neq c_e \} \;, \end{aligned}$$

and denote by $S(\vartheta)$ and $T(\vartheta)$ the respective subsets of corrupted vertices:

$$\begin{aligned} S(\vartheta) &= \{ v \in A \,:\, E(v) \cap Z(\vartheta) \neq \emptyset \} \;, \\ T(\vartheta) &= \{ v \in B \,:\, E(v) \cap W(\vartheta) \neq \emptyset \} \;. \end{aligned}$$

We also define the subset of *erased* vertices in Iteration 1, namely,

$$R(\vartheta) = \{ v \in A \,:\, z_e = ? \;\text{ for } e \in E(v) \} \;.$$

A. Iteration 1

An analog of the next proposition appears as a key ingredient in establishing the correctness of generalized minimum distance (GMD) decoding of concatenated codes [2]. The proposition provides a relation between the number of corrupted vertices and the number of erased vertices in A after Iteration 1.

Proposition 2 *There exists a threshold*

$$\vartheta \in \{1, 2, \ldots, \lceil d/2 \rceil\}$$

for which $2|S(\vartheta)| + |R(\vartheta)| \leq 2\,\mathsf{d}(\boldsymbol{y}, \boldsymbol{c})/d$.

The proof of Proposition 2 is similar to its GMD counterpart. The reader can refer to a proof of Theorem 2 in [2].

B. Iteration 2

Our method for bounding the size of $T(\vartheta)$ makes use of the following graph construction. The *split graph* of $\mathcal{G}$ is the graph $\widehat{\mathcal{G}} = (\widehat{V}, \widehat{E})$, where for each vertex $v \in V$ we let $\widehat{V}$ contain two *descendant* vertices v_1 and v_2, and for each edge $e = a$—b in E with endpoints $a \in A$ and $b \in B$, we let $\widehat{E}$ contain the following four descendant edges:

$$\begin{aligned} e_1 = a_1 \text{—} b_1 \;, &\quad e_1' = a_1 \text{—} b_2 \;, \\ e_2 = a_2 \text{—} b_1 \;, &\quad e_2' = a_2 \text{—} b_2 \;. \end{aligned}$$

An ordering on $\widehat{V}$ is assumed to be inherited from V, where v_i precedes v_j' if either $i < j$ or both $i = j$ and the parent vertex v precedes v' in V. This ordering induces an ordering on the set $\widehat{E}(v_i)$ of incident edges with v_i in $\widehat{\mathcal{G}}$.

It is easy to see that $\widehat{\mathcal{G}}$ is a 2Δ-regular bipartite graph whose vertex set is the union of the two sets $\widehat{A} = \bigcup_{v \in A} \{v_1, v_2\}$ and $\widehat{B} = \bigcup_{v \in B} \{v_1, v_2\}$ (each of size $2n$), and every edge has one endpoint in $\widehat{A}$ and one endpoint in $\widehat{B}$. The second largest eigenvalue of $\widehat{\mathcal{G}}$ equals 2λ [7, p. 84].

Define the function $\varphi : (F \cup \{?\})^N \to F^{4N}$ that maps a word $(x_e)_{e \in E}$ into a word $(\widehat{x}_f)_{f \in \widehat{E}}$ according to the following rule:

- If $x_e \neq ?$ we let $\widehat{x}_f = x_e$ for all $f \in \{e_1, e_1', e_2, e_2'\}$.
- Otherwise, we let $\widehat{x}_f = c_e$ for $f \in \{e_1, e_1'\}$ and $\widehat{x}_f = c_e \oplus 1$ for $f \in \{e_2, e_2'\}$.

Write $\varphi(\boldsymbol{c}) = \widehat{\boldsymbol{c}} = (\widehat{c}_f)_{f \in \widehat{E}}$ and $\varphi(\boldsymbol{z}(\vartheta)) = \widehat{\boldsymbol{z}}(\vartheta) = (\widehat{z}_f)_{f \in \widehat{E}}$ and define the following subsets of corrupted edges and vertices by

$$\widehat{Z}(\vartheta) = \left\{ f \in \widehat{E} \,:\, \widehat{z}_f \neq \widehat{c}_f \right\}$$

and

$$\widehat{S}(\vartheta) = \left\{ u \in \widehat{A} \,:\, \widehat{E}(u) \cap \widehat{Z}(\vartheta) \neq \emptyset \right\} \;.$$

It follows from the definition of $\widehat{\mathcal{G}}$ and φ that

$$|\widehat{S}(\vartheta)| = 2|S(\vartheta)| + |R(\vartheta)| \;,$$

and, so, from Proposition 2 we obtain the following.

Lemma 3 *There exists a threshold*

$$\vartheta \in \{1, 2, \ldots, \lceil d/2 \rceil\}$$

for which $|\widehat{S}(\vartheta)| \leq 2\,\mathsf{d}(\boldsymbol{y}, \boldsymbol{c})/d$.

Let $\widehat{\mathcal{C}}$ denote the linear $[2\Delta, k, 2d]$ code over F that is defined by

$$\widehat{\mathcal{C}} = \{ (\boldsymbol{u}\; \boldsymbol{u}) \,:\, \boldsymbol{u} \in \mathcal{C} \} \;,$$

and let $\widehat{\mathcal{D}} : F^{2\Delta} \to F^{2\Delta}$ be a decoder for $\widehat{\mathcal{C}}$ that recovers correctly any pattern of less than d errors and decodes *incorrectly* any pattern of d errors or more (e.g., when the pre-image is at Hamming distance d or more from every codeword in $\widehat{\mathcal{C}}$, we set the image under $\widehat{\mathcal{D}}$ to be a non-codeword).

Let $\widehat{\boldsymbol{w}}(\vartheta) = (\widehat{w}_f)_{f \in \widehat{E}}$ be the word in F^{4N} that is defined by

$$(\widehat{\boldsymbol{w}}(\vartheta))_u = \widehat{\mathcal{D}}\big((\widehat{\boldsymbol{z}}(\vartheta))_u\big) \quad \text{for every} \quad u \in \widehat{B} \;.$$

Observe that $\widehat{\boldsymbol{w}}(\vartheta)$ is the result of the application of Iteration 2 in Figure 1 to the code $(\widehat{\mathcal{G}}, \widehat{\mathcal{C}})$, with the word $\boldsymbol{z}$ in the algorithm taking the value $\widehat{\boldsymbol{z}}(\vartheta)$. We let $\widehat{T}(\vartheta)$ be the set of corrupted vertices in $\widehat{B}$, namely,

$$\widehat{T}(\vartheta) = \left\{ u \in \widehat{B} \,:\, \widehat{E}(u) \cap \widehat{W}(\vartheta) \neq \emptyset \right\} \;,$$

where

$$\widehat{W}(\vartheta) = \left\{ f \in \widehat{E} \,:\, \widehat{w}_f \neq \widehat{c}_f \right\} \;.$$

The next lemma relates between the sizes of $T(\vartheta)$ and $\widehat{T}(\vartheta)$.

Lemma 4 *For every $\vartheta \in \{1, 2, \ldots, \lceil d/2 \rceil\}$,*

$$2|T(\vartheta)| \le |\widehat{T}(\vartheta)| \ .$$

Proof. Suppose that $v \in T(\vartheta)$. We show that the two descendants of v in $\widehat{\mathcal{G}}$ are in $\widehat{T}(\vartheta)$. Denote by τ the size of $E(v) \cap Z(\vartheta)$ and let ρ be the number of vertices in $R(\vartheta)$ that are adjacent to v in $\mathcal{G}$. Note that τ (respectively, ρ) is the number of errors (respectively, erasures) in $(\boldsymbol{z}(\vartheta))_v$. By the construction of $\widehat{\mathcal{G}}$ and the definition of $\widehat{\boldsymbol{z}}(\vartheta)$ we have for each descendant u of v,

$$\mathsf{d}\left((\widehat{\boldsymbol{z}}(\vartheta))_u, (\widehat{\boldsymbol{c}})_u\right) = 2\tau + \rho \ . \tag{3}$$

Since v is in $T(\vartheta)$, the decoder $\tilde{\mathcal{D}}$ necessarily errs on $(\boldsymbol{z}(\vartheta))_v$; yet this is possible only if $2\tau + \rho \ge d$. Hence, from (3) we deduce that the decoder $\widehat{\mathcal{D}}$ errs on $(\widehat{\boldsymbol{z}}(\vartheta))_u$, i.e., $(\widehat{\boldsymbol{w}}(\vartheta))_u \ne (\widehat{\boldsymbol{c}})_u$; equivalently, $\widehat{E}(u) \cap \widehat{W}(\vartheta) \ne \emptyset$. Therefore, $u \in \widehat{T}(\vartheta)$. $\square$

We will also need a relation between the sizes of $\widehat{S}(\vartheta)$ and $\widehat{T}(\vartheta)$. To this end, we will make use of the next lemma, which is an application—with some modifications—of Lemmas 4 and 5 in [11] to the graph $\widehat{\mathcal{G}}$.

Lemma 5 *Suppose that $d \ge 2\lambda$ and that there exist subsets $\mathcal{S} \subseteq \widehat{A}$, $\mathcal{T} \subseteq \widehat{B}$, and $\mathcal{Z} \subseteq \widehat{E}$ that satisfy the following three conditions:*

1. *$|\mathcal{S}| \le n\left(\delta - (2\lambda/\Delta)\right)$;*
2. *every edge of $\mathcal{Z}$ has one of its endpoints in $\mathcal{S}$;*
3. *every vertex of $\mathcal{T}$ is incident with at least d edges of $\mathcal{Z}$.*

Then $|\mathcal{T}| \le |\mathcal{S}|$.

Proof. Let $\mathcal{H} = (\mathcal{S} \cup \mathcal{T}, \mathcal{E})$ be the subgraph of $\widehat{\mathcal{G}}$ that is induced by the vertex set $\mathcal{S} \cup \mathcal{T}$, and denote by γ the average degree in $\mathcal{H}$, namely, $\gamma = 2|\mathcal{E}|/|\mathcal{S} \cup \mathcal{T}|$. Applying Lemma 4 in [11] to $\mathcal{H}$ we get

$$\gamma \le \frac{2|\mathcal{S}||\mathcal{T}|}{|\mathcal{S}| + |\mathcal{T}|} \, \frac{\Delta}{n} + 2\lambda \ .$$

On the other hand,

$$\gamma = \frac{2|\mathcal{E}|}{|\mathcal{S}| + |\mathcal{T}|} \ge \frac{2|\mathcal{E} \cap \mathcal{Z}|}{|\mathcal{S}| + |\mathcal{T}|} \ge \frac{2d|\mathcal{T}|}{|\mathcal{S}| + |\mathcal{T}|} \ ,$$

where the last inequality follows from conditions 2 and 3. Combining the last two equations we obtain

$$2d|\mathcal{T}| \le 2|\mathcal{S}||\mathcal{T}|\frac{\Delta}{n} + 2\lambda(|\mathcal{S}| + |\mathcal{T}|) \ .$$

Finally, we make use of condition 1 and replace the first instance of $|\mathcal{S}|$ in the last equation by the upper bound $(n/\Delta)(d - 2\lambda)$. The result follows. $\square$

The following proposition yields an upper bound on the size of $T(\vartheta)$.

Proposition 6 *Suppose that $d \ge 2\lambda$ and that $\mathsf{d}(\boldsymbol{y}, \boldsymbol{c}) \le J_{\mathbb{C}}(\alpha)$ for some $\alpha < 1$. Then there exists a threshold $\vartheta \in \{1, 2, \ldots, \lceil d/2 \rceil\}$ for which*

$$|T(\vartheta)| \le \alpha n \left(\frac{\delta}{2} - \frac{\lambda}{\Delta}\right) \ .$$

Proof. Let ϑ be the threshold guaranteed by Lemma 3. We verify that the sets $\mathcal{S} = \widehat{S}(\vartheta)$, $\mathcal{T} = \widehat{T}(\vartheta)$, and $\mathcal{Z} = \widehat{Z}(\vartheta)$ satisfy conditions 1–3 in Lemma 5. Indeed, by Lemma 3 and the assumption on $\mathsf{d}(\boldsymbol{y}, \boldsymbol{c})$ we have

$$\begin{aligned} |\widehat{S}(\vartheta)| &\le 2\,\mathsf{d}(\boldsymbol{y}, \boldsymbol{c})/d \le 2J_{\mathbb{C}}(\alpha)/d \\ &= \alpha n \left(\delta - (2\lambda/\Delta)\right) < n\left(\delta - (2\lambda/\Delta)\right) \ , \end{aligned}$$

thereby satisfying condition 1. Condition 2 follows from the definition of the set $\widehat{S}(\vartheta)$, and condition 3 holds since the decoder $\widehat{\mathcal{D}} : F^{2\Delta} \to F^{2\Delta}$ errs only when it attempts to correct d errors or more. An application of Lemma 5 yields

$$|\widehat{T}(\vartheta)| \le |\widehat{S}(\vartheta)| \le \alpha n \left(\delta - (2\lambda/\Delta)\right) \ ,$$

and the result follows from Lemma 4. $\square$

C. Subsequent iterations

The upper bound on $|T(\vartheta)|$ in Proposition 6 coincides with the upper bound in [11] on the number of corrupted vertices after Iteration 1 in Figure 1 (see Equation (2) in [11]). Hence, any upper bound found in [11] on the number of corrupted vertices at Iteration t in Figure 1 is also an upper bound on that number at Iteration $t+1$ in Figure 2, provided that the threshold ϑ in the outer loop is the one guaranteed by Proposition 2. This, in turn, yields Theorem 1.

REFERENCES

[1] A. BARG, G. ZÉMOR, *Error exponents of expander codes, IEEE Trans. Inform. Theory,* 48 (2002), 1725–1729.

[2] G.D. FORNEY, *Generalized minimum distance decoding, IEEE Trans. Inform. Theory,* 12 (1966), 125–131.

[3] V. GURUSWAMI, P. INDYK, *Expander-based constructions of efficiently decodable codes, Proc. 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS'2001),* Las Vegas, Nevada (2001), 658–667.

[4] V. GURUSWAMI, P. INDYK, *Linear-time codes to correct a maximum possible fraction of errors, 39th Annual Allerton Conference on Communication, Control, and Computing,* Monticello, Illinois (2001).

[5] V. GURUSWAMI, P. INDYK, *Near-optimal linear-time codes for unique decoding and new list-decodable codes over smaller alphabets, Proc. 34th Annual ACM Symposium on Theory of Computing (STOC'2002),* Montréal, Quebec, Canada (2002), 812–821.

[6] A. LUBOTSKY, R. PHILIPS, P. SARNAK, *Ramanujan graphs, Combinatorica,* 8 (1988), 261–277.

[7] C.C. MACDUFFEE, *The Theory of Matrices,* Chelsea, New York, 1946.

[8] M. SIPSER, D.A. SPIELMAN, *Expander codes, IEEE Trans. Inform. Theory,* 42 (1996), 1710–1722.

[9] M. TANNER, *A recursive approach to low-complexity codes, IEEE Trans. Inform. Theory,* 27 (1981), 533–547.

[10] M. TANNER, *Explicit construction of concentrators from generalized N-gons, SIAM J. Alg. Disc. Meth.,* 5 (1984), 287–293.

[11] G. ZÉMOR, *On expander codes, IEEE Trans. Inform. Theory,* 47 (2001), 835–837.