

# Gradient Coding From Cyclic MDS Codes and Expander Graphs

Netanel Raviv<sup>ID</sup>, *Member, IEEE*, Itzhak Tamo<sup>ID</sup>, *Member, IEEE*, Rashish Tandon,  
and Alexandros G. Dimakis<sup>ID</sup>, *Senior Member, IEEE*

**Abstract**—Gradient coding is a technique for straggler mitigation in distributed learning. In this paper we design novel gradient codes using tools from classical coding theory, namely, cyclic MDS codes, which compare favorably with existing solutions, both in the applicable range of parameters and in the complexity of the involved algorithms. Second, we introduce an approximate variant of the gradient coding problem, in which we settle for approximate gradient computation instead of the exact one. This approach enables graceful degradation, i.e., the  $\ell_2$  error of the approximate gradient is a decreasing function of the number of stragglers. Our main result is that normalized adjacency matrices of expander graphs yield excellent approximate gradient codes, which enable significantly less computation compared to exact gradient coding, and guarantee faster convergence than trivial solutions under standard assumptions. We experimentally test our approach on Amazon EC2, and show that the generalization error of approximate gradient coding is very close to the full gradient while requiring significantly less computation from the workers.

**Index Terms**—Gradient descent, distributed computing, coding theory, expander graphs.

## I. INTRODUCTION

**D**ATA intensive machine learning tasks have become ubiquitous in many real-world applications, and with the increasing size of training data, distributed methods have gained increasing popularity. However, the performance of

Manuscript received September 18, 2018; revised March 16, 2020; accepted September 9, 2020. Date of publication October 7, 2020; date of current version November 20, 2020. The work of Netanel Raviv and Itzhak Tamo was supported in part by ISF under Grant 1030/15 and in part by NSF-BSF under Grant 2015814. The work of Netanel Raviv was supported in part by the postdoctoral fellowship of the Center for the Mathematics of Information (CMI) in the California Institute of Technology and in part by the Lester-Deutsch Postdoctoral Fellowship. This work was supported in part by NSF under Grant 1618689, Grant DMS 1723052, Grant CCF 1763702, Grant AF 1901292 and research gifts by Google, Western Digital, WNCG IAP; and in part by the Fluor Centennial Teaching Fellowship. This article was presented in part at the 2018 International Conference on Machine Learning (ICML). (*Corresponding author: Netanel Raviv.*)

Netanel Raviv is with the Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO 63130 USA (e-mail: netanel.raviv@wustl.edu).

Itzhak Tamo is with the Department of Electrical Engineering–Systems, Tel-Aviv University, Tel-Aviv 39040, Israel.

Rashish Tandon is with Apple, Seattle, WA 98105 USA.

Alexandros G. Dimakis is with the Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX 78712 USA.

Communicated by M. Lentmaier, Associate Editor for Coding Theory.

Color versions of one or more of the figures in this article are available online at <https://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2020.3029396




Fig. 1. Gradient coding for  $n = 3$ ,  $k = 3$ ,  $d = 2$ , and  $s = 1$  [24]. Each worker node  $W_i$  obtains two parts  $S_{i_1}, S_{i_2}$  of the data set  $S = S_1 \cup S_2 \cup S_3$ , computes the partial gradients  $g_{i_1}, g_{i_2}$ , and sends their linear combination back to the master node  $M$ . By choosing the coefficients judiciously, the master node  $M$  can compute the full gradient from any two responses, providing robustness against any one straggler.

distributed methods (in synchronous settings) is strongly dictated by *stragglers*, i.e., nodes that are slow to respond or unavailable. In this paper, we focus on coding theoretic (and graph theoretic) techniques for mitigating stragglers in distributed synchronous gradient descent.

A coding theoretic framework for straggler mitigation called *gradient coding* was first introduced in [24]. It consists of a system with one master and  $n$  worker nodes (or servers), in which the data is partitioned into  $k$  parts, and one or more parts is assigned to each one of the workers. In turn, each worker computes the partial gradient on each of its assigned parts, linearly combines the results according to some predetermined vector of coefficients, and sends this linear combination back to the master node. Choosing the coefficients at each node judiciously, one can guarantee that the master node is capable of reconstructing the full gradient even if *any*  $s$  machines fail to perform their work. The *storage overhead* of the system, which is denoted by  $d$ , refers to the amount of redundant computations, or alternatively, to the number of data parts that are sent to each node (see example in Fig. 1).

The importance of straggler mitigation was demonstrated in a series of recent studies (e.g., [14] and [27]). In particular, it was demonstrated in [24] that stragglers may run up to  $\times 5$  slower than the typical worker ( $\times 8$  in [27]) on Amazon EC2, especially for the cheaper virtual machines; such erratic

behavior is unpredictable and can significantly delay training. One can, of course, use more expensive instances but the goal here is to use coding theoretic methods to provide reliability out of cheap unreliable workers, overall reducing the cost of training.

By and large, the purpose of gradient coding is to enable the master node to compute the exact gradient out of the responses of *any*  $n - s$  non-straggling nodes. The work of [24] established the fundamental bound  $d \geq s + 1$ , which shows that one needs a growing amount of redundancy as the number of stragglers increases. In addition, [24] provided a deterministic construction which achieves it with equality when  $s + 1 | n$ , and a randomized one which applies to all  $s$  and  $n$ . Subsequently, deterministic constructions were also obtained by [7] and [8]. These works have focused on the scenario where  $s$  is known prior to the construction of the system. Furthermore, the exact computation of the full gradient is guaranteed if the number of stragglers is at most  $s$ , but no error bound is guaranteed if this number exceeds  $s$ .

The contribution of this work is twofold. For the computation of the exact gradient we employ tools from classic coding theory, namely, cyclic MDS codes, in order to obtain a deterministic construction which compares favourably with existing solutions; both in the applicable range of parameters, and in the complexity of the involved algorithms. Some of these gains are a direct application of well known properties of these codes.

Second, we introduce an *approximate* variant of the gradient coding problem. In this variant, the requirement for exact computation of the full gradient is traded by an approximate one, where the  $\ell_2$ -deviation of the given solution from the full gradient decreases as the number of stragglers decreases. Note that by this approach, the parameter  $s$  is *not* a part of the system construction, and the system can provide an approximate solution for any  $s < n$ , whose quality deteriorates gracefully as  $s$  increases. In the suggested solution, the coefficients at the worker nodes are based on an important family of graphs called *expanders*. In particular, it is shown that setting these coefficients according to a normalized adjacency matrix of an expander graph, a strong bound on the error term of the resulting solution is obtained. Moreover, this approach enables to break the aforementioned barrier  $d \geq s + 1$ , which is a substantial obstacle in gradient coding, and allows the master node to decode using a very simple algorithm.

This paper is organized as follows. Related work regarding gradient coding (and coded computation in general) is listed in Section II. A framework which encapsulates all the results in this paper is given in Section III. Necessary mathematical notions from coding theory and graph theory are given in Section IV. The former notions are used to obtain an algorithm for exact gradient computation in Section V, and the latter ones are used to obtain an algorithm for the approximate gradient in Section VI. Experimental results are given in Section VII.

## II. RELATED WORK

The work of Lee et al. [13] initiated the use of coding theoretic methods for mitigating stragglers in large-scale learning. This work is focused on linear regression and therefore

can exploit more structure compared to the general gradient coding problem that we study here. The work by Li et al. [17], investigates a generalized view of the coding ideas in [13], showing that their solution is a single operating point in a general scheme of trading off latency of computation to the load of communication.

Further closely related work has shown how coding can be used for distributed MapReduce, as well as a similar communication and computation tradeoff [16], [18]. We also mention the work of [11] which addresses straggler mitigation in linear regression by using a different approach, that is not mutually exclusive with gradient coding. In their work, the data is coded rather than replicated, and the nodes perform their computation on coded data.

The work by [7] generalizes previous work for linear models [13] but can also be applied to general models to yield explicit gradient coding constructions. Our results regarding the exact gradient are closely related to the work by [8], [9] which was obtained independently from our work. In [8], similar coding theoretic tools were employed in a fundamentally different fashion. Both [8] and [7] are comparable in parameters to the randomized construction of [24] and are outperformed by us in a wide range of parameters. A detailed comparison of the theoretical asymptotic behaviour is given in the sequel.

*Remark 1:* None of the aforementioned works studies approximate gradient computations. However, we note that subsequent to this work, two unpublished manuscripts [4], [15] study a similar approximation setting and obtain related results albeit using randomized as opposed to deterministic approaches. Furthermore, the exact setting was also discussed subsequent to this work in [28] and [29]. In [28] it was shown that network communication can be reduced by increasing the replication factor, and respective bounds were given. The work of [29] discussed coded polynomial computation with low overhead, and applies to gradient coding whenever the gradient at hand is a polynomial.

## III. FRAMEWORK

This section provides a unified framework which accommodates straggler mitigation in both the exact and approximate gradient computations which follow. By and large, we use lowercase letters  $a, b, \dots$  to refer to scalars and functions, bold lowercase letter  $\mathbf{a}, \mathbf{b}, \dots$  to refer to vectors, uppercase letters  $A, B, \dots$  to refer to workers (or nodes), uppercase bold letters  $\mathbf{A}, \mathbf{B}$  to refer to matrices, and uppercase calligraphic letters  $\mathcal{A}, \mathcal{B}, \dots$  to refer to sets and codes. E.g., the rows of a matrix  $\mathbf{B}$  are denoted by  $\mathbf{b}_i$ , and its entries are denoted by  $b_{i,j}$ . Unless otherwise stated, all vectors are *row* vectors, and for any field  $\mathbb{F}$  we adopt the common notation  $\mathbb{F}^n$  to refer to  $\mathbb{F}^{1 \times n}$ . We also employ the standard notation  $[n] \triangleq \{1, 2, \dots, n\}$  for an integer  $n$ .

We begin with a brief introduction to machine learning, and the reader is referred to [23] for further reading (in particular, to *gradient descent* [23, Sec. 14.1] and *stochastic gradient descent* [23, Sec. 14.3]). Broadly speaking, the general purpose of machine learning is to find a hypothesis from a given

hypotheses class, which best approximates an unknown target function by observing a *training set*  $\mathcal{S} = \{s_i\}_{i=1}^m$  that is labeled by that target function. Assuming that the hypotheses class is parametrized by real vector  $\mathbf{w} \in \mathbb{R}^p$ , one defines a *loss function*  $\ell(\mathbf{w}, s)$ , which penalizes a given hypothesis  $\mathbf{w}$  for erring on a data point  $s$ . Then, one wishes to find the  $\mathbf{w}$  which minimizes the *empirical risk*  $L_{\mathcal{S}}(\mathbf{w}) \triangleq \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{w}, s_i)$  by using analytic methods (in some parts of this paper, we assume that  $L_{\mathcal{S}}$  is smooth and convex). That is, one starts from an arbitrary point  $\mathbf{w}_0$ , and iteratively computes the *gradient*  $\nabla L_{\mathcal{S}}(\mathbf{w}) = (\frac{\partial}{\partial w_i} L_{\mathcal{S}}(\mathbf{w}))_{i=1}^n$  at the current point, and moves away from the current point in the opposite direction of the gradient, until convergence is achieved. In the *gradient descent* algorithm the exact gradient is computed at every iteration, and in the *stochastic gradient descent* algorithm it is taken from a distribution whose expected value is the exact gradient.

In order to distribute the execution of gradient descent from a master node  $M$  to  $n$  worker nodes  $\{W_j\}_{j=1}^n$  (Algorithm 1), the training set  $\mathcal{S}$  is partitioned to  $n$  disjoint subsets  $\{\mathcal{S}_i\}_{i=1}^n$  of size<sup>1</sup>  $\frac{m}{n}$  each. These subsets are distributed among  $\{W_j\}_{j=1}^n$ , at most  $d$  subsets at each worker for some parameter  $d$ . Then, every node computes the partial gradients  $\nabla L_{\mathcal{S}_i}(\mathbf{w})$  on the  $\mathcal{S}_i$ 's which it obtained, where  $L_{\mathcal{S}_i}(\mathbf{w}) \triangleq \frac{m}{n} \sum_{s \in \mathcal{S}_i} \ell(\mathbf{w}, s)$ . The algorithm operates in iterations, where in iteration  $r \in [t]$  every node evaluates its gradients in the current model  $\mathbf{w}^{(r)}$ , and sends to  $M$  some linear combination of them. After obtaining responses from at least  $n - s_r$  workers, where  $s_r$  is the number of stragglers in iteration  $r$ ,  $M$  aggregates them to form the gradient  $\nabla L_{\mathcal{S}}(\mathbf{w}^{(r)})$  of the overall empirical risk at  $\mathbf{w}^{(r)}$ . In the exact setting the value of  $s_r$  will be some fixed  $s$  for every  $r$ , whereas in the approximate setting this value is at the discretion of the master, in correspondence with the required approximation error.

To support mitigation of stragglers in this setting, the following notions are introduced. Let  $\mathbf{B} \in \mathbb{C}^{n \times n}$  be a matrix whose  $i$ 'th row  $\mathbf{b}_i$  contains the coefficients of the linear combination  $\sum_{j=1}^n b_{i,j} \cdot \nabla L_{\mathcal{S}_j}(\mathbf{w}^{(r)})$  that is sent to  $M$  by  $W_i$ . Note that the support  $\text{supp}(\mathbf{b}_i)$  contains the indices of the sets  $\mathcal{S}_j$  that are to be sent to  $W_i$  by  $M$ . Given a set of non-stragglers  $\mathcal{K} \in \mathcal{P}(n)$ , where  $\mathcal{P}(n)$  is the set of all nonempty subsets of  $[n]$ , a function  $a : \mathcal{P}(n) \rightarrow \mathbb{C}^n$  provides  $M$  with a vector by which the results from  $\{W_i\}_{i \in \mathcal{K}}$  are to be linearly combined to obtain a vector  $\mathbf{v}_r \in \mathbb{R}^p$ . This vector is either the true gradient, in which case Algorithm 1 is gradient descent, or an approximation of it (whose expectation is the true gradient), in which case Algorithm 1 is stochastic gradient descent. We also require that  $\text{supp}(a(\mathcal{K})) \subseteq \mathcal{K}$  for all  $\mathcal{K} \in \mathcal{P}(n)$ . In most of the subsequent constructions, the matrix  $\mathbf{B}$  and the function  $a$  will be defined over  $\mathbb{R}$  rather than over  $\mathbb{C}$ .

Different constructions of the matrix  $\mathbf{B}$  and the function  $a$  in Algorithm 1 enable to compute the gradient either exactly (which requires the storage overhead  $d$  to be at least  $s_r + 1$  for

<sup>1</sup>For simplicity, assume that  $m|n$ . The given scheme could be easily adapted to the case  $m \nmid n$ . Further, the assumption that the number of partitions equals to the number of nodes is a mere convenience, and all subsequent schemes can be adapted to the case where the number of partitions is at most the number of nodes.

---

**Algorithm 1** Gradient Coding

---

**Input:** Data  $\mathcal{S} = \{z_i = (x_i, y_i)\}_{i=1}^m$ , number of iterations  $t > 0$ , learning rates  $\{\eta\}_{r=1}^t > 0$ , straggler tolerance parameters  $\{s_r\}_{r=1}^t$ , a matrix  $\mathbf{B} \in \mathbb{C}^{n \times n}$ , and a function  $a : \mathcal{P}(n) \rightarrow \mathbb{C}^n$ .

Initialize  $\mathbf{w}^{(1)} = (0, \dots, 0)$ .

Partition  $\mathcal{S} = \cup_{i=1}^n \mathcal{S}_i$  and send  $\{\mathcal{S}_j : j \in \text{supp}(\mathbf{b}_i)\}$  to  $W_i$  for every  $i \in [n]$ .

**for**  $r = 1$  **to**  $t$  **do**

$M$  broadcasts  $\mathbf{w}^{(r)}$  to all nodes.

Each  $W_j$  sends  $\frac{1}{n} \sum_{i \in \text{supp}(\mathbf{b}_j)} b_{j,i} \cdot \nabla L_{\mathcal{S}_i}(\mathbf{w}^{(r)})$  to  $M$ .

$M$  waits until at least  $n - s_r$  nodes responded.

$M$  computes  $\mathbf{v}_r = a(\mathcal{K}_r) \cdot \mathbf{C}$ , where the  $i$ 'th row of  $\mathbf{C}$  is  $\frac{1}{n}$  times the response from  $W_i$ , if it responded, and 0 otherwise, and  $\mathcal{K}_r$  is the set of non-stragglers in the current iteration  $r$ .

$M$  updates  $\mathbf{w}^{(r+1)} \triangleq \mathbf{w}^{(r)} - \eta_r \mathbf{v}_r$

**end for**

**Return**  $\frac{1}{t} \sum_{r=1}^t \mathbf{w}^{(r+1)}$ .

---

all  $r \in [t]$  or *approximately*. In what follows, the respective requirements and guarantees from  $a$  and  $\mathbf{B}$  are discussed. In the following definitions, for an integer  $r$  let  $\mathbb{1}_r$  be the vector of  $r$  ones, where the subscript is omitted if clear from context, and for  $\mathcal{K} \subseteq [n]$  let  $\mathcal{K}^c \triangleq [n] \setminus \mathcal{K}$ .

**Definition 2:** A matrix  $\mathbf{B} \in \mathbb{C}^{n \times n}$  and a function  $a : \mathcal{P}(n) \rightarrow \mathbb{C}^n$  satisfy the Exact Computation (EC) condition if for all  $\mathcal{K} \subseteq [n]$  such that  $|\mathcal{K}| \geq \max_{r \in [t]} s_r$ , we have  $a(\mathcal{K}) \cdot \mathbf{B} = \mathbb{1}$ .

**Definition 3:** For a non-decreasing function  $\epsilon : [n-1] \rightarrow \mathbb{R}_{\geq 0}$  such that  $\epsilon(0) = 0$ ,  $a$  and  $\mathbf{B}$  satisfy the  $\epsilon$ -Approximate Computation ( $\epsilon$ -AC) condition, if for all  $\mathcal{K} \in \mathcal{P}(n)$ , we have  $d_2(a(\mathcal{K})\mathbf{B}, \mathbb{1}) \leq \epsilon(|\mathcal{K}^c|)$  (where  $d_2$  is Euclidean distance).

Notice that the error term  $\epsilon$  in Definition 3 is a function of the number of stragglers since it is not expected to decrease if more stragglers are present. The conditions which are given in Definition 2 and Definition 3 guarantee the exact and approximate computation by the following lemmas. In the upcoming proofs, let  $\mathbf{N}(\mathbf{w})$  be the matrix

$$\mathbf{N}(\mathbf{w}) \triangleq \frac{1}{n} \cdot \begin{pmatrix} \nabla L_{\mathcal{S}_1}(\mathbf{w}) \\ \nabla L_{\mathcal{S}_2}(\mathbf{w}) \\ \vdots \\ \nabla L_{\mathcal{S}_n}(\mathbf{w}) \end{pmatrix}. \quad (1)$$

**Lemma 4:** If  $a$  and  $\mathbf{B}$  satisfy the EC condition, then for all  $r \in [t]$  we have  $\mathbf{v}_r = \nabla L_{\mathcal{S}}(\mathbf{w}^{(r)})$ .

**Proof:** For a given  $r \in [t]$ , let  $\mathbf{B}'$  be the matrix whose  $i$ 'th row  $\mathbf{b}'_i$  equals  $\mathbf{b}_i$  if  $i \in \mathcal{K}_r$ , and zero otherwise. By the definition of  $\mathbf{C}$  in Algorithm 1 it follows that  $\mathbf{C} = \mathbf{B}' \cdot \mathbf{N}(\mathbf{w}^{(r)})$ , and since  $\text{supp}(a(\mathcal{K}_r)) \subseteq \mathcal{K}_r$  it follows that  $a(\mathcal{K}_r)\mathbf{B}' = a(\mathcal{K}_r)\mathbf{B}$ . Therefore, we have

$$\begin{aligned} \mathbf{v}_r &= a(\mathcal{K}_r) \cdot \mathbf{C} = a(\mathcal{K}_r) \cdot \mathbf{B}' \cdot \mathbf{N}(\mathbf{w}^{(r)}) \\ &= a(\mathcal{K}_r) \cdot \mathbf{B} \cdot \mathbf{N}(\mathbf{w}^{(r)}) = \mathbb{1} \cdot \mathbf{N}(\mathbf{w}^{(r)}) \end{aligned}$$



$$\begin{aligned}
&= \frac{1}{n} \sum_{i=1}^n \nabla L_{\mathcal{S}_i}(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \frac{1}{m/n} \sum_{\mathbf{z} \in \mathcal{S}_i} \nabla \ell(\mathbf{w}^{(r)}, \mathbf{z}) \\
&= \frac{1}{m} \sum_{\mathbf{z} \in \mathcal{S}} \nabla \ell(\mathbf{w}^{(r)}, \mathbf{z}) = \nabla L_{\mathcal{S}}(\mathbf{w}^{(r)}).
\end{aligned}$$

□

The next lemma bounds the deviance of  $\mathbf{v}_r$  from the gradient of the empirical risk at the current model  $\mathbf{w}^{(r)}$  by using the function  $\epsilon$  and the *spectral norm*  $\|\cdot\|_{\text{spec}}$  of  $\mathbf{N}(\mathbf{w})$ . Recall that for a matrix  $\mathbf{P}$  the spectral norm is defined as  $\|\mathbf{P}\|_{\text{spec}} \triangleq \max_{\mathbf{x}: \|\mathbf{x}\|_2=1} \|\mathbf{P}\mathbf{x}\|_2$ .

**Lemma 5:** For a function  $\epsilon$  as above, if  $a$  and  $\mathbf{B}$  satisfy the  $\epsilon$ -AC condition, then  $d_2(\mathbf{v}_r, \nabla L_{\mathcal{S}}(\mathbf{w}^{(r)})) \leq \epsilon(|\mathcal{K}_r^c|) \cdot \|\mathbf{N}(\mathbf{w}^{(r)})\|_{\text{spec}}$ .

*Proof:* As in the proof of Lemma 4, we have that

$$\begin{aligned}
d_2(\mathbf{v}_r, \nabla L_{\mathcal{S}}(\mathbf{w}^{(r)})) &= \\
&= d_2(a(\mathcal{K}_r) \cdot \mathbf{B}' \cdot \mathbf{N}(\mathbf{w}^{(r)}), \mathbf{1} \cdot \mathbf{N}(\mathbf{w}^{(r)})) \\
&= d_2(a(\mathcal{K}_r) \cdot \mathbf{B} \cdot \mathbf{N}(\mathbf{w}^{(r)}), \mathbf{1} \cdot \mathbf{N}(\mathbf{w}^{(r)})) \\
&= \|(a(\mathcal{K}_r)\mathbf{B} - \mathbf{1}) \cdot \mathbf{N}(\mathbf{w}^{(r)})\|_2 \\
&\leq d_2(a(\mathcal{K}_r)\mathbf{B}, \mathbf{1}) \cdot \|\mathbf{N}(\mathbf{w}^{(r)})\|_{\text{spec}} \\
&\leq \epsilon(|\mathcal{K}_r^c|) \cdot \|\mathbf{N}(\mathbf{w}^{(r)})\|_{\text{spec}}.
\end{aligned}$$

□

Due to Lemma 4 and Lemma 5, in the remainder of this paper we focus on constructing  $a$  and  $\mathbf{B}$  that satisfy either the EC condition (Section V) or the  $\epsilon$ -AC condition (Section VI).

**Remark 6:** In some settings [24], it is convenient to partition the data set  $\mathcal{S}$  to  $\{\mathcal{S}_i\}_{i=1}^k$ , where  $k < n$ . Notice that the definitions of  $a$  and  $\mathbf{B}$  above extend verbatim to this case as well. If  $a$  and  $\mathbf{B}$  satisfy the EC condition, we have that  $a(\mathcal{K})\mathbf{B} = \mathbf{1}_n$  for every large enough  $\mathcal{K} \subseteq [n]$ . Hence, by omitting any  $n - k$  columns of  $\mathbf{B}$  to form a matrix  $\hat{\mathbf{B}}$ , we have that  $a(\mathcal{K})\hat{\mathbf{B}} = \mathbf{1}_k$ , and hence a scheme for any partition of  $\mathcal{S}$  to  $k$  parts emerges instantly. This new scheme  $(a, \hat{\mathbf{B}})$  is resilient to an identical number of stragglers  $s$  and has lesser or equal storage overhead than  $(a, \mathbf{B})$ . Similarly, if  $a$  and  $\mathbf{B}$  satisfy the  $\epsilon$ -AC condition for some  $\epsilon$ , then the scheme  $(a, \hat{\mathbf{B}})$  has lesser or equal storage overhead, and an identical error function  $\epsilon$ , since  $d_2(a(\mathcal{K})\hat{\mathbf{B}}, \mathbf{1}_k) \leq d_2(a(\mathcal{K})\mathbf{B}, \mathbf{1}_n) \leq \epsilon(|\mathcal{K}^c|)$  for any  $\mathcal{K} \in \mathcal{P}(n)$ .

#### IV. MATHEMATICAL NOTIONS

This section provides a brief overview on the mathematical notions that are essential for the suggested schemes. The exact computation (Sec. V) requires notions from coding theory, and the approximate one (Sec. VI) requires notions from graph theory. The coding theoretic material in this section is taken from [22], which focuses on finite fields, and yet the given results extend verbatim to the real or complex case (see also [20], Sec. 8.4).

For a field  $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$  an  $[n, \kappa]$  (linear) code  $\mathcal{C}$  over  $\mathbb{F}$  is a subspace of dimension  $\kappa$  of  $\mathbb{F}^n$ . The minimum distance  $\delta$  of  $\mathcal{C}$  is  $\min\{d_H(\mathbf{x}, \mathbf{y}) : \mathbf{x}, \mathbf{y} \in \mathcal{C}, \mathbf{x} \neq \mathbf{y}\}$ , where  $d_H$  denotes the *Hamming distance*  $d_H(\mathbf{x}, \mathbf{y}) \triangleq |\{i | x_i \neq y_i\}|$ . Since the code is a linear subspace, it follows that the minimum distance of a code is equal to the minimum *Hamming weight*

$w_H(\mathbf{x}) \triangleq \|\mathbf{x}\|_0 = |\text{supp}(\mathbf{x})|$  among the nonzero codewords in  $\mathcal{C}$ . The well-known *Singleton bound* states that  $\delta \leq n - \kappa + 1$ , and codes which attain this bound with equality are called *Maximum Distance Separable* (MDS) codes.

**Definition 7:** [22, Sec. 8] A code  $\mathcal{C}$  is called *cyclic* if the cyclic shift of every codeword is also a codeword, namely,

$$(c_1, c_2, \dots, c_n) \in \mathcal{C} \Rightarrow (c_n, c_1, \dots, c_{n-1}) \in \mathcal{C}.$$

The *dual code* of  $\mathcal{C}$  is the subspace  $\mathcal{C}^\perp \triangleq \{\mathbf{y} \in \mathbb{F}^n | \mathbf{y} \cdot \mathbf{c}^\top = 0 \text{ for all } \mathbf{c} \in \mathcal{C}\}$ . Several well-known and easy to prove properties of MDS codes are used throughout this paper.

**Lemma 8:** If  $\mathcal{C} \subseteq \mathbb{F}^n$  is an  $[n, \kappa]$  MDS code, then

- A1.  $\mathcal{C}^\perp$  is an  $[n, n - \kappa]$  MDS code, and hence its minimum Hamming weight is  $\kappa + 1$ .
- A2. For any subset  $\mathcal{K} \subseteq [n]$  of size  $n - \kappa + 1$  there exists a codeword in  $\mathcal{C}$  whose support (i.e., the set of nonzero indices) is  $\mathcal{K}$ .
- A3. The reverse code  $\mathcal{C}^R \triangleq \{(c_n, \dots, c_1) | (c_1, \dots, c_n) \in \mathcal{C}\}$  is an  $[n, \kappa]$  MDS code.

*Proof:* For A1 see [22, Prob. 4.1]. The proof of A3 is trivial since permuting the coordinates of a code does not alter its minimum distance. For A2, let  $\mathbf{G} \in \mathbb{F}^{\kappa \times n}$  be the *generator matrix* of  $\mathcal{C}$ , i.e., a matrix whose rows are a basis to  $\mathcal{C}$ . By restricting  $\mathbf{G}$  to the columns indexed by  $\mathcal{K}^c$  we get a  $\kappa \times (\kappa - 1)$  matrix, which has a nonzero vector  $\mathbf{y}$  in its left kernel, and hence  $\mathbf{y}\mathbf{G}$  is a codeword in  $\mathcal{C}$  which is zero in the entries that are indexed by  $\mathcal{K}^c$ . Since the minimum distance of  $\mathcal{C}$  is  $n - \kappa + 1$ , it follows that  $\mathbf{y}\mathbf{G}$  has nonzero values in entries that are indexed by  $\mathcal{K}$ , and the claim follows. □

Two common families of codes are used in the sequel—Reed-Solomon (RS) codes and Bose-Chaudhuri-Hocquenghem (BCH) codes. An  $[n, \kappa]$  RS code  $\mathcal{C}$  is defined by a set of  $n$  distinct *evaluation points*  $\alpha_0, \dots, \alpha_{n-1} \in \mathbb{F}$  as

$$\mathcal{C} = \{(f(\alpha_0), f(\alpha_1), \dots, f(\alpha_{n-1})) : f \in \mathbb{F}^{<\kappa}[x]\},$$

where  $\mathbb{F}^{<\kappa}[x]$  is the set of polynomials of degree less than  $\kappa$  and coefficients from  $\mathbb{F}$  in the variable  $x$ . Alternatively, RS codes can be defined as  $\{\mathbf{y}\mathbf{V} | \mathbf{y} \in \mathbb{F}^\kappa\}$ , where  $\mathbf{V} \in \mathbb{F}^{\kappa \times n}$  is a Vandermonde matrix on  $\{\alpha_i\}_{i=0}^{n-1}$ , i.e.,  $v_{i,j} = \alpha_j^{i-1}$  for every  $(i, j) \in [\kappa] \times [n]$ . It is widely known that RS codes are MDS codes, and in some cases, they are also cyclic.

In contrast with RS codes, where every codeword is a vector of *evaluations* of a polynomial, a codeword in a BCH code is a vector of *coefficients* of a polynomial; that is, a codeword  $\mathbf{c} = (c_0, c_1, \dots, c_{n-1})$  is identified by  $c(x) \triangleq c_0 + c_1x + \dots + c_{n-1}x^{n-1}$ . For a field  $\mathbb{K}$  that contains  $\mathbb{F}$  and a set  $\mathcal{R} \subseteq \mathbb{K}$ , a BCH code  $\mathcal{C}$  is defined as  $\mathcal{C} = \{c \in \mathbb{F}^n | c(r) = 0 \text{ for all } r \in \mathcal{R}\}$ . The set  $\mathcal{R}$  is called *the roots of  $\mathcal{C}$* , or alternatively,  $\mathcal{C}$  is said to be a BCH code on  $\mathcal{R}$  over  $\mathbb{F}$ . For example, a set of complex numbers  $\mathcal{R} \subseteq \mathbb{C}$  defines a BCH code on  $\mathcal{R}$  over  $\mathbb{R}$ , which is the set of real vectors whose corresponding polynomials vanish on  $\mathcal{R}$ .

**Lemma 9:** [20], [22] If all elements of  $\mathcal{R} \subseteq \mathbb{C}$  are roots of unity of order  $n$ , then the BCH code  $\mathcal{C}$  on  $\mathcal{R}$  over  $\mathbb{R}$  is cyclic.

*Proof:* If  $c(x)$  is a codeword in  $\mathcal{C}$ , then its cyclic shift is given by  $\tilde{c}(x) \triangleq c(x) \cdot x \bmod (x^n - 1) = x \cdot c(x) - c_{n-1} \cdot (x^n - 1)$ . Since every  $r \in \mathcal{R}$  is a root of unity of order  $n$ ,

it follows that

$$\tilde{c}(r) = r \cdot c(r) - c_{n-1} \cdot (r^n - 1) = r \cdot c(r) = 0,$$

and hence  $\tilde{c}$  is a codeword in  $\mathcal{C}$ .  $\square$

Further, the structure of  $R$  may also imply a lower bound on the distance of  $\mathcal{C}$ .

**Theorem 10:** (The BCH bound) [20], [22] If  $\mathcal{R}$  contains a subset of  $\ell$  consecutive powers of a primitive root of unity (i.e., a subset of the form  $\omega^b, \omega^{b+1}, \dots, \omega^{b+\ell-1}$ , where  $\omega$  is a primitive  $n$ 'th root of unity), then the minimum distance of  $\mathcal{C}$  is at least  $\ell + 1$ .

In the remainder of this section, a brief overview on expander graphs is given. The interested reader is referred to [10] for further details. Let  $G = (\mathcal{V}, \mathcal{E})$  be a  $d$ -regular, undirected, and connected graph on  $n$  nodes. Let  $\mathbf{A}_G \in \mathbb{R}^{n \times n}$  be the *adjacency matrix* of  $G$ , i.e.,  $(\mathbf{A}_G)_{i,j} = 1$  if  $\{i, j\} \in \mathcal{E}$ , and 0 otherwise. Since  $\mathbf{A}_G$  is a real symmetric matrix, it follows that it has  $n$  real eigenvalues  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ , and denote  $\lambda \triangleq \max\{|\lambda_2|, |\lambda_n|\}$ . It is widely known [10] that  $\lambda_1 = d$ , and that  $\lambda_n \geq -d$ , where equality holds if and only if  $G$  is bipartite. Further, it also follows from  $\mathbf{A}_G$  being real and symmetric that it has a basis of orthogonal real eigenvectors  $\mathbf{v}_1 = \mathbf{1}, \mathbf{v}_2, \dots, \mathbf{v}_n$ , and w.l.o.g assume that  $\|\mathbf{v}_i\|_2 = 1$  for every  $i \geq 2$ . The parameters  $\lambda$  and  $d$  are related by the celebrated Alon-Boppana Theorem.

**Theorem 11:** [10] An infinite sequence of  $d$  regular graphs on  $n$  vertices satisfies that  $\lambda \geq 2\sqrt{d-1} - o_n(1)$ , where  $o_n(1)$  is an expression which tends to zero as  $n$  tends to infinity.

Constant degree regular graphs (i.e., families of graphs with fixed degree  $d$  that does not depend on  $n$ ) for which  $\lambda$  is small in comparison with  $d$  are largely referred to as *expanders*. In particular, graphs which attain the above bound asymptotically (i.e.,  $\lambda \leq 2\sqrt{d-1}$ ) are called *Ramanujan graphs*, and several efficient constructions are known [6], [19]. Since explicit constructions of Ramanujan graphs are often rather intricate, one may resort to choosing a random regular graph and verify its expansion by computing the respective eigenvalues. This process that is known to produce a good expander with high probability [10, Theorem 7.10], and is used in our experiments.

## V. EXACT GRADIENT CODING FROM CYCLIC MDS CODES

For a given  $n$  and  $s$ , let  $\mathcal{C}$  be a cyclic  $[n, n-s]$  MDS code over  $\mathbb{F}$  that contains  $\mathbf{1}$  (explicit constructions of such codes are given in the sequel). According to Lemma 8, there exists a codeword  $\mathbf{c}_1 \in \mathcal{C}$  whose support is  $\{1, \dots, s+1\}$ . Let  $\mathbf{c}_2, \dots, \mathbf{c}_n$  be all cyclic shifts of  $\mathbf{c}_1$ , which lie in  $\mathcal{C}$  by its cyclic property. Finally, let  $\mathbf{B}$  be the  $n \times n$  matrix whose *columns* are  $\mathbf{c}_1, \dots, \mathbf{c}_n$ , i.e.,  $\mathbf{B} \triangleq (\mathbf{c}_1^\top, \mathbf{c}_2^\top, \dots, \mathbf{c}_n^\top)$ . The following lemma provides some important properties of  $\mathbf{B}$ .

**Lemma 12:** The matrix  $\mathbf{B}$  satisfies the following properties.

- B1.  $\|\mathbf{b}\|_0 = s+1$  for every row  $\mathbf{b}$  of  $\mathbf{B}$ .
- B2. Every row of  $\mathbf{B}$  is a codeword in  $\mathcal{C}^R$ .
- B3. The column span of  $\mathbf{B}$  is the code  $\mathcal{C}$ .
- B4. Every set of  $n-s$  rows of  $\mathbf{B}$  are linearly independent over  $\mathbb{F}$ .

**Proof:** To prove B1 and B2, observe that  $\mathbf{B}$  is of the following form, where  $\mathbf{c}_1 \triangleq (\beta_1, \dots, \beta_{s+1}, 0, \dots, 0)$ .

$$\begin{pmatrix} \beta_1 & 0 & \cdots & 0 & \beta_{s+1} & \beta_s & \cdots & \beta_2 \\ \beta_2 & \beta_1 & 0 & \cdots & 0 & \beta_{s+1} & \cdots & \beta_3 \\ \vdots & \vdots & \ddots & \ddots & \vdots & \ddots & \ddots & \vdots \\ \beta_s & \beta_{s-1} & \cdots & \beta_1 & 0 & \cdots & 0 & \beta_{s+1} \\ \beta_{s+1} & \beta_s & \cdots & \beta_2 & \beta_1 & 0 & \cdots & 0 \\ 0 & \beta_{s+1} & \cdots & \beta_3 & \beta_2 & \beta_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \beta_{s+1} & \beta_s & \beta_{s-1} & \cdots & \beta_1 \end{pmatrix}.$$

To prove B3, notice that the leftmost  $n-s$  columns of  $\mathbf{B}$  have leading coefficients in different positions, and hence they are linearly independent. Thus, the dimension of the column span of  $\mathbf{B}$  is at least  $n-s$ , and since  $\dim \mathcal{C} = n-s$ , the claim follows.

To prove B4, assume for contradiction that there exist a set of  $n-s$  linearly dependent rows. Hence, there exists a vector  $\mathbf{v} \in \mathbb{F}^n$  of Hamming weight  $n-s$  such that  $\mathbf{v}\mathbf{B} = 0$ . According to B3, the columns of  $\mathbf{B}$  span  $\mathcal{C}$ , and hence the vector  $\mathbf{v}$  lies in the dual code  $\mathcal{C}^\perp$  of  $\mathcal{C}$ . Since  $\mathcal{C}^\perp$  is an  $[n, s]$  MDS code by Lemma 8, it follows that the minimum Hamming weight of a codeword in  $\mathcal{C}^\perp$  is  $n-s+1$ , a contradiction.  $\square$

Since  $\mathcal{C}^R$  is of dimension  $n-s$ , it follows from parts B2 and B4 of Lemma 12 that every set of  $n-s$  rows of  $\mathbf{B}$  are a basis to  $\mathcal{C}^R$ . Furthermore, since  $\mathbf{1} \in \mathcal{C}$  it follows that  $\mathbf{1} \in \mathcal{C}^R$ . Therefore, there exists a function  $a : \mathcal{P}(n) \rightarrow \mathbb{F}^n$  such that for any set  $\mathcal{K} \subseteq [n]$  of size  $n-s$  we have that  $\text{supp}(a(\mathcal{K})) \subseteq \mathcal{K}$  and  $a(\mathcal{K}) \cdot \mathbf{B} = \mathbf{1}$ .

**Theorem 13:** The above  $a$  and  $\mathbf{B}$  satisfy the EC condition (Definition 2).

In the remainder of this section, two cyclic MDS codes over the complex numbers and the real numbers are suggested, from which the construction in Theorem 13 can be obtained. These constructions are taken from [20] (Sec. II.B), and are given with a few adjustments to our case. The contributions of these codes is summarized in the following theorem.

**Theorem 14:** For any given  $n$  and  $s$  there exist explicit complex valued  $a$  and  $\mathbf{B}$  that satisfy the EC-condition with optimal  $d = s+1$ . The respective encoding (i.e., constructing  $\mathbf{B}$ ) and decoding (i.e., constructing  $a(\mathcal{K})$  given  $\mathcal{K}$ ) complexities are  $O(s(n-s))$  and  $O(s \log^2(s) + n \log n)$ , respectively. In addition, for any given  $n$  and  $s$  such that  $n \not\equiv s \pmod{2}$  there exist explicit real valued  $a$  and  $\mathbf{B}$  that satisfy the EC-condition with optimal  $d = s+1$ . The encoding and decoding complexities are  $O(\min\{s \log^2(s), n \log n\})$  and  $O(g_s + s(n-s))$ , where  $g_s$  is the complexity of inverting a generalized Vandermonde matrix.

### A. Cyclic-MDS Codes Over the Complex Numbers

For a given  $n$  and  $s$ , let  $i = \sqrt{-1}$ , and let  $\mathcal{A} \triangleq \{\alpha_j\}_{j=0}^{n-1}$  be the set of  $n$  complex roots of unity of order  $n$ , i.e.,  $\alpha_j \triangleq e^{2\pi i j/n}$ . Let  $\mathbf{G} \in \mathbb{C}^{(n-s) \times n}$  be a complex Vandermonde matrix over  $\mathcal{A}$ , i.e.,  $g_{k,j} = \alpha_{j-1}^{k-1}$  for  $j \in [n]$  and any  $k \in [n-s]$ . Finally, let  $\mathbf{C} \triangleq \{\mathbf{x}\mathbf{G} | \mathbf{x} \in \mathbb{C}^{n-s}\}$ . It is readily verified that  $\mathbf{C}$

is an  $[n, n-s]$  MDS code that contains  $\mathbf{1}$ , whose codewords may be seen as the evaluations of all polynomials in  $\mathbb{C}^{<n-s}[x]$  on the set  $\mathcal{A}$ .

*Lemma 15:* The code  $\mathcal{C}$  is cyclic.

*Proof:* Let  $\mathbf{c} \in \mathcal{C}$  be a codeword, and let  $f_{\mathbf{c}} \in \mathbb{C}^{<n-s}[x]$  be the corresponding polynomial. Consider the polynomial  $f_{\mathbf{c}'}(x) \triangleq f_{\mathbf{c}}(e^{2\pi i/n} \cdot x)$ , and notice that  $\deg f_{\mathbf{c}'} = \deg f_{\mathbf{c}}$ . Further, it is readily verified that any  $j \in \{0, 1, \dots, n-1\}$  satisfies that  $f_{\mathbf{c}'}(\alpha_j) = f_{\mathbf{c}}(\alpha_{j-1})$ , where the indices are taken modulo  $n$ . Hence, the evaluation of the polynomial  $f_{\mathbf{c}'}$  on the set of roots  $\mathcal{A}$  results in the cyclic shift of the codeword  $\mathbf{c}$ , and lies in  $\mathcal{C}$  itself.  $\square$

*Corollary 16:* The code  $\mathcal{C}$  is a cyclic MDS code which contains  $\mathbf{1}$ , and hence it can be used to obtain  $\mathbf{a}$  and  $\mathbf{B}$ , as described in Theorem 13.

Given a set  $\mathcal{K}$  of  $n-s$  non-stragglers, an algorithm for computing the encoding vector  $\mathbf{a}(\mathcal{K})$  in  $O(s \log^2(s) + n \log n)$  operations over  $\mathbb{C}$  (after a one-time initial computation of  $O(s^2 + s(n-s))$ ), is given in Appendix A. The complexity of this algorithm is asymptotically smaller than the corresponding algorithm in [7] and [8] whenever  $s = o(n)$ . Furthermore, the cyclic structure of the matrix  $\mathbf{B}$  enables a very simple algorithm for its construction; this algorithm compares favorably with previous works for any  $s$ , and is given in Appendix A as well.

*Remark 17:* Note that the use of complex rather than real matrix  $\mathbf{B}$  may potentially double the required bandwidth, since every complex number contains two real numbers. A simple manipulation of Algorithm 1 which resolves this issue is given in Appendix C. This optimal bandwidth is also attained by the scheme in the next section, which uses a smaller number of multiplication operations. However, it is applicable only if  $n \not\equiv s \pmod{2}$ .

### B. Cyclic-MDS Codes Over the Real Numbers

If one wishes to abstain from using complex numbers, e.g., in order to reduce bandwidth, we suggest the following construction, which provides a cyclic MDS code over the reals. This construction relies on [20] (Property 3), with an additional specialized property.

*Construction 18:* For a given  $n$  and  $s$  such that  $n \neq s \pmod{2}$ , define the following BCH codes over the reals. In both cases denote  $\omega \triangleq e^{2\pi i/n}$ .

- 1) If  $n$  is even and  $s$  is odd let  $s' \triangleq \lfloor \frac{s}{2} \rfloor$ , and let  $\mathcal{C}_1$  be a BCH code which consists of all polynomials in  $\mathbb{R}^{<n}[x]$  that vanish over the set  $\mathcal{R}_1 \triangleq \{\omega^{n/2-s'}, \omega^{n/2-s'+1}, \dots, \omega^{n/2+s'}\}$ .
- 2) If  $n$  is odd and  $s$  is even let  $n' \triangleq \lfloor \frac{n}{2} \rfloor$ , and let  $\mathcal{C}_2$  be a BCH code which consists of all polynomials in  $\mathbb{R}^{<n}[x]$  that vanish over the set  $\mathcal{R}_2 \triangleq \{\omega^{n'-s/2+1}, \omega^{n'-s/2+2}, \dots, \omega^{n'+s/2}\}$ .

*Lemma 19:* The codes  $\mathcal{C}_1$  and  $\mathcal{C}_2$  from Construction 18 are cyclic  $[n, n-s]$  MDS codes that contain  $\mathbf{1}$ .

*Proof:* According to Lemma 9, it is clear that  $\mathcal{C}_1$  and  $\mathcal{C}_2$  are cyclic. According to the BCH bound (Theorem 10), it is also clear that the minimum distance of  $\mathcal{C}_1$  is at least  $|\mathcal{R}_1|+1 = s+1$ , and the minimum distance of  $\mathcal{C}_2$  is at least  $|\mathcal{R}_2|+1 = s+1$ .

Hence, to prove that  $\mathcal{C}_1$  and  $\mathcal{C}_2$  are MDS codes, it is shown that their code dimensions are  $n-s$ .

Since the sets  $\mathcal{R}_1$  and  $\mathcal{R}_2$  are closed under conjugation (i.e.,  $r$  is in  $\mathcal{R}_i$  if and only if the conjugate of  $r$  is in  $\mathcal{R}_i$ ) it follows that the polynomials  $p_1(x) \triangleq \prod_{r \in \mathcal{R}_1} (x-r)$  and  $p_2(x) \triangleq \prod_{r \in \mathcal{R}_2} (x-r)$  have real coefficients. Hence, by the definition of BCH codes it follows that

$$\begin{aligned} \{p_1(x) \cdot \ell(x) | \ell(x) \in \mathbb{R}^{<n-s}[x]\} &\subseteq \mathcal{C}_1 \\ \{p_2(x) \cdot \ell(x) | \ell(x) \in \mathbb{R}^{<n-s}[x]\} &\subseteq \mathcal{C}_2, \end{aligned} \quad (2)$$

and hence,  $\dim \mathcal{C}_1 \geq n-s$  and  $\dim \mathcal{C}_2 \geq n-s$ . Let  $d(\mathcal{C}_1)$  and  $d(\mathcal{C}_2)$  be the minimum distances of  $\mathcal{C}_1$  and  $\mathcal{C}_2$ , respectively, and notice that by the Singleton bound [22] (Sec. 4.1) it follows that

$$\begin{aligned} \dim \mathcal{C}_1 &\leq n - d(\mathcal{C}_1) + 1 \leq n - (s+1) + 1 = n-s \\ \dim \mathcal{C}_2 &\leq n - d(\mathcal{C}_2) + 1 \leq n - (s+2) + 1 = n-s, \end{aligned}$$

and thus  $\mathcal{C}_1$  and  $\mathcal{C}_2$  satisfy the Singleton bound with equality, or equivalently, they are MDS codes. To prove that  $\mathbf{1}$  is in  $\mathcal{C}_1$  and  $\mathcal{C}_2$ , we ought to show that  $\mathbf{1}(r_1) = \mathbf{1}(r_2) = 0$  (where  $\mathbf{1}$  denotes the all ones polynomial here) for every  $r_1 \in \mathcal{R}_1$  and  $r_2 \in \mathcal{R}_2$ , which amounts to showing that  $\sum_{j=0}^{n-1} r_1^j = \sum_{j=0}^{n-1} r_2^j = 0$ . It is well-known that the sum of the 0'th to  $(n-1)$ 'th power of any root of unity of order  $n$ , other than 1, equals zero. Since

$$\begin{aligned} 1 &\leq \frac{n}{2} - \lfloor \frac{s}{2} \rfloor = \frac{n}{2} - s' < \frac{n}{2} - s' + 1 \\ &< \dots < \frac{n}{2} + s' = \frac{n}{2} + \lfloor \frac{s}{2} \rfloor \leq n-1 \\ &\text{if } n \text{ is even and } s \text{ is odd, and} \\ 1 &\leq \lfloor \frac{n}{2} \rfloor - \frac{s}{2} + 1 = n' - \frac{s}{2} + 1 < n' - \frac{s}{2} + 2 \\ &< \dots < n' + \frac{s}{2} = \lfloor \frac{n}{2} \rfloor + \frac{s}{2} \leq n-1 \\ &\text{otherwise,} \end{aligned}$$

it follows that all powers of  $\omega$  in  $\mathcal{R}_1$  and  $\mathcal{R}_2$  are between 1 and  $n-1$ , and hence  $\mathbf{1} \notin \mathcal{R}_1$  and that  $\mathbf{1} \notin \mathcal{R}_2$ . Hence, we have that  $\mathbf{1}(r_1) = \mathbf{1}(r_2) = 0$  for every  $r_1 \in \mathcal{R}_1$  and  $r_2 \in \mathcal{R}_2$ , which concludes the claim.  $\square$

Algorithms for computing the matrix  $\mathbf{B}$  and the vector  $\mathbf{a}(\mathcal{K})$  for the codes in this subsection are given in Appendix B. The algorithm for construction  $\mathbf{B}$  outperforms previous works whenever  $s = o(n)$ , and the algorithm for computing  $\mathbf{a}(\mathcal{K})$  outperforms previous works for a smaller yet wide range of  $s$  values.

## VI. APPROXIMATE GRADIENT CODING FROM EXPANDER GRAPHS

Recall that in order to retrieve that exact gradient, one must have  $d \geq s+1$ , an undesirable overhead in many cases. To break this barrier, we relax the requirement to retrieve the gradient exactly, and settle for an approximation of it. Note that trading the exact gradient for an approximate one is a necessity in many variants of gradient descent (such as the acclaimed *stochastic gradient descent* [21, Sec. 14.3]), and hence our techniques are aligned with common practices in machine learning.



For a set  $\mathcal{K} \subseteq [n]$  of non-stragglers, let  $\mathbf{1}_{\mathcal{K}}$  be its binary characteristic vector with respect to  $[n]$ . Setting  $\mathbf{B}$  as the identity matrix and  $a$  as the function which maps  $\mathcal{K} \in \mathcal{P}(n)$  to  $\frac{n}{|\mathcal{K}|} \cdot \mathbf{1}_{\mathcal{K}}$  corresponds to executing gradient descent while ignoring the stragglers [5], and averaging the partial gradients from  $\mathcal{K}$  (the factor  $\frac{n}{|\mathcal{K}|}$  corrects the factor  $\frac{1}{n}$  in (1)). In what follows, this is referred to as the *trivial scheme*. The respective  $\mathbf{B}$  and  $a$  clearly satisfy the  $\epsilon$ -AC scheme for  $\epsilon(s) = \sqrt{\frac{ns}{n-s}}$ , where  $s = s(\mathcal{K}) \triangleq n - |\mathcal{K}|$ , since

$$\begin{aligned} d_2(a(\mathcal{K})\mathbf{B}, \mathbf{1}) &= d_2\left(\frac{n}{n-s}\mathbf{1}_{\mathcal{K}}, \mathbf{1}\right) \\ &= \sqrt{\sum_{j \in \mathcal{K}} \left(\frac{n}{n-s} - 1\right)^2 + \sum_{j \notin \mathcal{K}} 1^2} \\ &= \sqrt{\frac{ns}{n-s}}. \end{aligned} \quad (3)$$

We show that this can be outperformed by setting  $\mathbf{B}$  to be a normalized adjacency matrix of a connected regular graph on  $n$  nodes, which is constructed by the master before dispersing the data, and setting  $a$  to be some simple function.

The resulting error function  $\epsilon(s)$  depends on the parameters of the graph, whereas the resulting storage overhead  $d$  is given by its *degree* (i.e., the *fixed* number of neighbors of each node). The error function is given below for a general connected and regular graph, and particular examples with their resulting errors are given in the sequel. In particular, it is shown that taking the graph to be an expander graph provides an error term  $\epsilon$  which is smaller than (3) for any  $s$ .

For a given  $n$  let  $G$  be a connected  $d$ -regular graph on  $n$  nodes, with eigenvalues  $\lambda_1 \geq \dots \geq \lambda_n$ , corresponding (row) eigenvectors  $\mathbf{v}_1 = \mathbf{1}, \mathbf{v}_2, \dots, \mathbf{v}_n$ , and  $\lambda \triangleq \max\{|\lambda_2|, |\lambda_n|\}$  as described in Subsection IV. For a given  $\mathcal{K} \subseteq [n]$  define  $\mathbf{u}_{\mathcal{K}} \in \mathbb{R}^n$  as

$$(\mathbf{u}_{\mathcal{K}})_i = \begin{cases} -1 & i \notin \mathcal{K} \\ \frac{s}{n-s} & i \in \mathcal{K}, \end{cases} \quad (4)$$

and let  $\langle \mathbf{v}_1, \mathbf{v}_2, \dots \rangle$  be the span of  $\mathbf{v}_1, \mathbf{v}_2, \dots$  over  $\mathbb{R}^n$ .

**Lemma 20:** For any  $\mathcal{K} \subseteq [n]$  we have  $\mathbf{u}_{\mathcal{K}} \in \langle \mathbf{v}_2, \dots, \mathbf{v}_n \rangle$ .

*Proof:* First, observe that  $\langle \mathbf{v}_2, \dots, \mathbf{v}_n \rangle$  is exactly the subspace of all vectors whose sum of entries is zero. This follows from the fact that  $\{\mathbf{1}, \mathbf{v}_2, \dots, \mathbf{v}_n\}$  is an orthogonal basis, hence  $\mathbf{v}_i^\top \cdot \mathbf{1} = 0$  for every  $i \geq 2$ , and from the fact that  $\{\mathbf{v}_2, \dots, \mathbf{v}_n\}$  are linearly independent. Since the sum of entries of  $\mathbf{u}_{\mathcal{K}}$  is zero, the result follows.  $\square$

**Corollary 21:** For any  $\mathcal{K} \subseteq [n]$  there exists  $\alpha_2, \dots, \alpha_n \in \mathbb{R}$  such that  $\mathbf{u}_{\mathcal{K}} = \alpha_2 \mathbf{v}_2 + \dots + \alpha_n \mathbf{v}_n$ , and  $\|\mathbf{u}_{\mathcal{K}}\|_2 = \sqrt{\sum_{i=2}^n \alpha_i^2} = \sqrt{\frac{ns}{n-s}}$ .

*Proof:* The first part follows immediately from Lemma 20. The second part follows by computing the  $\ell_2$  norm of  $\mathbf{u}_{\mathcal{K}}$  in two ways, once by its definition (4) and again by using the representation of  $\mathbf{u}_{\mathcal{K}}$  as a linear combination of the orthonormal set  $\{\mathbf{v}_2, \dots, \mathbf{v}_n\}$ .  $\square$

Now, let  $\mathbf{B} \triangleq \frac{1}{d} \cdot \mathbf{A}_G$ , define  $a : \mathcal{P}(n) \rightarrow \mathbb{R}^n$  as  $a(\mathcal{K}) = \mathbf{u}_{\mathcal{K}} + \mathbf{1}$ , and observe that  $\text{supp}(a(\mathcal{K})) = \mathcal{K}$  for all  $\mathcal{K} \in \mathcal{P}(n)$ . Note that computing  $a(\mathcal{K})$  given  $\mathcal{K}$  is done by a straightforward  $O(n)$  algorithm. The error function  $\epsilon$  is given by the following lemma.

**Lemma 22:** For every nonempty<sup>2</sup> set  $\mathcal{K} \subseteq [n]$  of size  $n - s$  we have that  $d_2(a(\mathcal{K})\mathbf{B}, \mathbf{1}) \leq \frac{\lambda}{d} \cdot \sqrt{\frac{ns}{n-s}} \triangleq \epsilon(s)$ .

*Proof:* Notice that the eigenvalues of  $\mathbf{B}$  are  $\mu_i \triangleq \frac{\lambda_i}{d}$ , and hence  $\mu \triangleq \max\{|\mu_2|, |\mu_n|\}$  equals  $\frac{\lambda}{d}$ . Further, the eigenvectors are identical to those of  $\mathbf{A}_G$ . Therefore, it follows from Corollary 21 that

$$\begin{aligned} d_2(a(\mathcal{K})\mathbf{B}, \mathbf{1}) &= d_2((\mathbf{1} + \mathbf{u}_{\mathcal{K}})\mathbf{B}, \mathbf{1}) \\ &= d_2((\mathbf{1} + \alpha_2 \mathbf{v}_2 + \dots + \alpha_n \mathbf{v}_n)\mathbf{B}, \mathbf{1}) \\ &= d_2(\mathbf{1} + \alpha_2 \mu_2 \mathbf{v}_2 + \dots + \alpha_n \mu_n \mathbf{v}_n, \mathbf{1}) \\ &= \|\alpha_2 \mu_2 \mathbf{v}_2 + \dots + \alpha_n \mu_n \mathbf{v}_n\|_2, \end{aligned}$$

and since  $\{\mathbf{v}_2, \dots, \mathbf{v}_n\}$  are orthonormal, it follows that

$$\begin{aligned} &\|\alpha_2 \mu_2 \mathbf{v}_2 + \dots + \alpha_n \mu_n \mathbf{v}_n\|_2 \\ &= \sqrt{\sum_{i=2}^n \mu_i^2 \alpha_i^2} \leq \sqrt{\sum_{i=2}^n \mu^2 \alpha_i^2} \\ &= \mu \sqrt{\sum_{i=2}^n \alpha_i^2} = \frac{\lambda}{d} \sqrt{\frac{ns}{n-s}}. \end{aligned}$$

$\square$

**Corollary 23:** The above  $a$  and  $\mathbf{B}$  satisfy the  $\epsilon$ -AC condition for  $\epsilon(s) = \frac{\lambda}{d} \sqrt{\frac{ns}{n-s}}$ . The storage overhead of this scheme equals the degree  $d$  of the underlying regular graph  $G$ .

**Remark 24:** For the purpose of convergence analysis, that is given in Subsection VI-C which follows, one must multiply the added gradient by some constant  $c$  very close to 1 to formally guarantee convergence. For the sake of simplicity we omit this constant from the following discussion, and provide the full details in Subsection VI-C.

It is evident that in order to obtain small deviation  $\epsilon(s)$ , it is essential to have a small  $\lambda$  and a large  $d$ . However, most constructions of expanders have focused in the case where  $d$  is constant (i.e.,  $d = O(1)$ ). On one hand, constant  $d$  serves our purpose well in terms of storage overhead, since it implies a *constant* storage overhead. On the other hand, a constant  $d$  does not allow  $\lambda/d$  to tend to zero as  $n$  tends to infinity due to Theorem 11.

It is readily verified that our scheme outperforms the trivial one by a multiplicative factor of  $\frac{\lambda}{d}$ , which is less than one for every non-bipartite graph; bipartite graphs are used in a slightly different fashion in the sequel. We conclude the discussion with several examples, the first of which uses *Margulis graphs* ([10], Sec. 8), that are rather easy to construct.

**Example 25:** For any integer  $n$  there exists an 8-regular graph on  $n$  nodes with  $\lambda \leq 5\sqrt{2}$ . For example, by using these graphs with the parameters  $n = 500$ ,  $d = 8$ , we have an improvement factor of  $\frac{\lambda}{d} = \frac{5\sqrt{2}}{8} \approx 0.883$ .

Several additional examples for Ramanujan graphs, which attain an improvement factor  $\leq \frac{2\sqrt{d-1}}{d} \approx \frac{2}{\sqrt{d}}$  but are harder to construct, are as follows.

<sup>2</sup>For  $\mathcal{K} = \emptyset$  we clearly have  $\mathbf{u}_{\mathcal{K}} + \mathbf{1} = 0$ , and hence  $d_2(a(\mathcal{K})\mathbf{B}, \mathbf{1}) = \|\mathbf{1}\|_2 = \sqrt{n}$ .

*Example 26:* [19] Let  $p$  and  $q$  be distinct primes such that  $p \equiv 1 \pmod{4}$ ,  $q \equiv 1 \pmod{4}$ , and such that the Legendre symbol  $\left(\frac{p}{q}\right)$  is 1 (i.e.,  $p$  is a quadratic residue modulo  $q$ ). Then, there exist a non-bipartite Ramanujan graph on  $n = \frac{q(q^2-1)}{2}$  nodes and constant degree  $p+1$ .

- 1) If  $p = 13$  and  $q = 17$  then  $n = 2448$ ,  $d = 14$ , and  $\frac{2\sqrt{d-1}}{d} \approx 0.534$ .
- 2) If  $p = 5$  and  $q = 29$  then  $n = 12180$ ,  $d = 6$ , and  $\frac{2\sqrt{d-1}}{d} \approx 0.816$ .

Restricting  $d$  to be a constant (i.e., not to grow with  $n$ ) is detrimental to the improvement factor  $\frac{\lambda}{d}$  due to Theorem 11, but allows lower storage overhead. If one wishes a smaller (i.e., better) improvement factor at the price of higher overhead, the following is useful.

*Example 27:* [1] There exists a polynomial algorithm (in  $n$ ) to produce a graph  $G$  with the parameters  $(n, d, \lambda) = (2^m, m-1, \sqrt{m \log^3 m})$ . For this family of graphs, the term  $\frac{\lambda}{d}$  goes to zero as  $n$  goes to infinity.

#### A. Bipartite Expanders

The above approximation scheme can be used with a *bipartite* graph  $G$  as well. However, bipartite graphs satisfy that  $\lambda = d$ , and hence the resulting error function  $\epsilon(s) = \sqrt{\frac{ns}{n-s}}$  is identical to the error function of the trivial scheme (3), which requires lower overhead, and hence no gain is attained. However, in what follows it is shown that bipartite graphs on  $2n$  nodes can be employed in a slightly different fashion, and obtain  $\epsilon(s) = \frac{\sigma_2}{d} \sqrt{\frac{ns}{n-s}}$  for some  $\sigma_2 < d$  that is defined shortly.

To this end, we require the notion of *singular value decomposition*, which implies that any matrix  $\mathbf{P} \in \mathbb{R}^{n \times n}$  can be written as  $\mathbf{P} = \mathbf{U}\mathbf{D}\mathbf{V}^\top$ , where  $\mathbf{D} \in \mathbb{R}^{n \times n}$  is a diagonal matrix, and  $\mathbf{U}$  and  $\mathbf{V}$  are orthonormal matrices (i.e.,  $\mathbf{U}\mathbf{U}^\top = \mathbf{V}\mathbf{V}^\top = \mathbf{I}_n$ , where  $\mathbf{I}_n$  is the identity matrix of order  $n$ ). The elements  $\{\sigma_i\}_{i=1}^n$  on the diagonal of  $\mathbf{D}$ , which are nonnegative, are called the *singular values* of  $\mathbf{P}$ , the columns  $\{\mathbf{u}_i^\top\}_{i=1}^n$  of  $\mathbf{U}$  are called *left-singular vectors* of  $\mathbf{P}$ , and the columns  $\{\mathbf{v}_i^\top\}_{i=1}^n$  of  $\mathbf{V}$  are called *right-singular vectors* of  $\mathbf{P}$ . The singular values and singular vectors of  $\mathbf{P}$  can be arranged in triples  $\{(\mathbf{u}_i, \mathbf{v}_i, \sigma_i)\}_{i=1}^n$  such that for all  $i \in [n]$  we have  $\mathbf{P}\mathbf{v}_i^\top = \sigma_i \mathbf{u}_i^\top$  and  $\mathbf{P}^\top \mathbf{u}_i^\top = \sigma_i \mathbf{v}_i^\top$ , which implies that  $\mathbf{P}^\top \mathbf{P} \mathbf{v}_i^\top = \sigma_i^2 \mathbf{v}_i^\top$  and  $\mathbf{P}\mathbf{P}^\top \mathbf{u}_i^\top = \sigma_i^2 \mathbf{u}_i^\top$ .

Let  $G = (\mathcal{L} \cup \mathcal{R}, \mathcal{E})$  be a  $d$ -regular (in both sides) and connected bipartite graph on  $2n$  nodes (and hence  $|\mathcal{L}| = |\mathcal{R}| = n$ ), with an adjacency matrix

$$\mathbf{A}_G = \begin{pmatrix} 0 & \mathbf{C} \\ \mathbf{C}^\top & 0 \end{pmatrix}$$

for some  $n \times n$  real matrix  $\mathbf{C}$ , and eigenvalues  $\lambda_1 = d \geq \lambda_2 \geq \dots \geq \lambda_{2n} = -d$ . Let  $\{(\mathbf{u}_i, \mathbf{v}_i, \sigma_i)\}_{i=1}^n$  be the set of triples of left-singular vectors, right-singular vectors, and singular values of  $\mathbf{C}$ , as explained above, where  $\sigma_1 \geq \dots \geq \sigma_n \geq 0$ . The next well-known lemma presents the connection between the singular values of  $\mathbf{C}$  and the eigenvalues of  $\mathbf{A}_G$ . Its proof is a combination of a few simple exercises, and is given for completeness.

*Lemma 28:*  $\{\lambda_i\}_{i=1}^{2n} = \{\sigma_i\}_{i=1}^n \cup \{-\sigma_i\}_{i=1}^n$ .

*Proof:* For any  $i \in [n]$  we have that

$$(\mathbf{u}_i, \mathbf{v}_i) \begin{pmatrix} 0 & \mathbf{C} \\ \mathbf{C}^\top & 0 \end{pmatrix} = (\mathbf{v}_i \mathbf{C}^\top, \mathbf{u}_i \mathbf{C}) = \sigma_i (\mathbf{u}_i, \mathbf{v}_i),$$

and hence  $\sigma_i \in \{\lambda_i\}_{i=1}^{2n}$ . It is an easy exercise to verify that the eigenvalues of a bipartite graph are symmetric around zero, and hence it follows that  $-\sigma_i \in \{\lambda_i\}_{i=1}^{2n}$  as well.

Conversely, for any  $i \in [2n]$  let  $(\mathbf{s}_i, \mathbf{t}_i)$  be an eigenvector of  $\mathbf{A}_G$  with a corresponding eigenvalue  $\lambda_i$ , where  $\mathbf{s}_i, \mathbf{t}_i \in \mathbb{R}^n$ . Since  $(\mathbf{s}_i, \mathbf{t}_i) \mathbf{A}_G = \lambda_i (\mathbf{s}_i, \mathbf{t}_i)$ , it follows that

$$\begin{cases} \mathbf{t}_i \mathbf{C}^\top = \lambda_i \mathbf{s}_i \\ \mathbf{s}_i \mathbf{C} = \lambda_i \mathbf{t}_i \end{cases}, \text{ and hence } \begin{cases} \mathbf{t}_i \mathbf{C}^\top \mathbf{C} = \lambda_i^2 \mathbf{t}_i \\ \mathbf{s}_i \mathbf{C} \mathbf{C}^\top = \lambda_i^2 \mathbf{s}_i \end{cases}.$$

Therefore, it follows that  $\lambda_i^2 \in \{\sigma_i^2\}_{i=1}^n$ , and thus  $\lambda_i \in \{\sigma_i\}_{i=1}^n \cup \{-\sigma_i\}_{i=1}^n$ .  $\square$

From Lemma 28, the following corollaries are easy to prove.

*Corollary 29:*

- A.  $(1, 1, d) \in \{(\mathbf{u}_i, \mathbf{v}_i, \sigma_i)\}_{i=1}^n$ .
- B.  $\sigma_2 < d$ .

*Proof:*

- A. Repeat the second part of the proof of Lemma 28 with  $i = 1$ , namely, with  $(\mathbf{s}_i, \mathbf{t}_i) = (1, 1)$  and  $\lambda_i = d$ .
- B. Since  $\sigma_1 \geq \dots \geq \sigma_n$ , it follows from Lemma 28 that  $\lambda_2 = \sigma_2$ . Further, since  $G$  is connected, it follows that  $\lambda_2 < d$ , and hence  $\sigma_2 < d$ .  $\square$

Given Corollary 29.A, we may assume without loss of generality that  $\mathbf{u}_1 = \mathbf{v}_1 = 1$ , and that  $\|\mathbf{v}_i\|_2 = \|\mathbf{u}_i\|_2 = 1$  for every  $i \in \{2, \dots, n\}$ . Now, for a given set  $\mathcal{K} \subseteq [n]$  of size  $n-s$ , let  $\mathbf{u}_{\mathcal{K}}$  and  $a(\mathcal{K})$  be as in Section VI. Notice that since  $\{\mathbf{u}_i\}_{i=1}^n$  is an orthonormal basis, and since  $1$  is orthogonal to  $\mathbf{u}_{\mathcal{K}}$ , it follows that  $\mathbf{u}_{\mathcal{K}} \in \langle \mathbf{u}_2, \dots, \mathbf{u}_n \rangle$ . Hence, there exist real  $\alpha_2, \dots, \alpha_n$  such that  $\mathbf{u}_{\mathcal{K}} = \sum_{i=2}^n \alpha_i \mathbf{u}_i$ . Further, it follows that

$$\|\mathbf{u}_{\mathcal{K}}\|_2 = \sqrt{\frac{ns}{n-s}} = \sqrt{\sum_{i=2}^n \alpha_i^2}.$$

By setting  $\mathbf{B} \triangleq \frac{1}{d} \mathbf{C}$ , we have the following lemma, which may be seen as the equivalent of Lemma 22 to the bipartite case.

*Proposition 30:* For every set  $\mathcal{K} \subseteq [n]$  of size  $n-s$ ,  $d_2(a(\mathcal{K})\mathbf{B}, 1) \leq \frac{\sigma_2}{d} \cdot \sqrt{\frac{ns}{n-s}}$ .

*Proof:* By Corollary 29, and since  $\{\mathbf{v}_i\}_{i=2}^n$  is an orthonormal set, it follows that

$$\begin{aligned} d_2(a(\mathcal{K})\mathbf{B}, 1) &= \|(\mathbf{1} + \mathbf{u}_{\mathcal{K}}) \cdot \frac{1}{d} \mathbf{C} - \mathbf{1}\|_2 = \|\frac{1}{d} \mathbf{u}_{\mathcal{K}} \mathbf{C}\|_2 \\ &= \|\frac{1}{d} \left( \sum_{i=2}^n \alpha_i \mathbf{u}_i \right) \mathbf{C}\|_2 = \|\frac{1}{d} \sum_{i=2}^n \alpha_i \sigma_i \mathbf{v}_i\|_2 \\ &= \frac{1}{d} \sqrt{\sum_{i=2}^n \sigma_i^2 \alpha_i^2} \leq \frac{\sigma_2}{d} \sqrt{\frac{ns}{n-s}} \end{aligned}$$

$\square$

Applying the above lemma on several constructions of bipartite expanders provides the following examples.



*Example 31:* Let  $p$  and  $q$  be distinct primes such that  $p \equiv 1 \pmod{4}$ ,  $q \equiv 1 \pmod{4}$ , and such that the Legendre symbol  $\left(\frac{2}{q}\right)$  is  $-1$ . Then, there exists a bipartite Ramanujan graph on  $2n = q(q^2 - 1)$  nodes with  $\sigma_2 \leq 2\sqrt{p}$ .

- 1) If  $p = 5$  and  $q = 13$  then  $n = \frac{2184}{2} = 1092$ ,  $d = 6$ , and  $\frac{\sigma_2}{d} \leq \frac{2\sqrt{5}}{6} \approx 0.745$ .
- 2) If  $p = 13$  and  $q = 5$  then  $n = \frac{120}{2} = 60$ ,  $d = 14$ , and  $\frac{\sigma_2}{d} \leq \frac{2\sqrt{13}}{14} \approx 0.515$ .

### B. Lower Bound

Finally, we have the following lower bound on the approximation error of any *Approximate Computation (AC)* scheme, that establishes asymptotic optimality of our scheme, up to constants, when used with Ramanujan graphs. In what follows, for any set of vertices  $\mathcal{A}$  in a graph  $G$ , let  $\mathcal{N}(\mathcal{A})$  be the set of vertices which has at least one neighbor in  $\mathcal{A}$ , and for a vertex  $V$  let  $d(V)$  denote its degree.

*Lemma 32:* Let  $G = (\mathcal{W} \cup \mathcal{P}, \mathcal{E})$  be a bipartite graph where  $|\mathcal{W}| = |\mathcal{P}| = n$  for some  $n$  and  $d(W) \leq d$  for every  $W \in \mathcal{W}$  and some  $d$ . Then, for every  $r \in \{1, 2, \dots, \lfloor \frac{n}{d} \rfloor\}$  there exists a set  $\mathcal{Q}_r \subseteq \mathcal{P}$  of size  $r$  such that  $|\mathcal{N}(\mathcal{Q}_r)| \leq d|\mathcal{Q}_r|$ .

*Proof:* Let  $\mathcal{P} = \{P_1, \dots, P_n\}$  and without loss of generality assume that  $d(P_1) \leq d(P_2) \leq \dots \leq d(P_n)$ . Due to this ordering of  $P_1, \dots, P_n$ , and due to the bounded degree of vertices in  $\mathcal{W}$ , it follows that  $d \geq \text{avg}_n \geq \text{avg}_{n-1} \geq \text{avg}_1$ , where  $\text{avg}_j \triangleq \frac{1}{j} \sum_{i=1}^j d(P_i)$  for every  $j \in [n]$ . Therefore, the average degree of  $\mathcal{Q}_r \triangleq \{P_1, \dots, P_r\}$  is at most  $d$ , which implies that  $\sum_{i=1}^r d(P_i) \leq rd$ . Since the size of  $\mathcal{N}(\mathcal{Q}_r)$  is at most the sum of degrees in  $\mathcal{Q}_r$ , it follows that  $|\mathcal{N}(\mathcal{Q}_r)| \leq \sum_{i=1}^r d(P_i) \leq rd = d|\mathcal{Q}_r|$ .  $\square$

*Lemma 33:* Consider any  $\mathbf{B} \in \mathbb{R}^{n \times n}$  with each row having at most  $d$  non-zeros. Then, for any  $s$  such that  $d < s < n$  there exists a set  $\mathcal{K} \subseteq [n]$  of size  $n - s$  such that

$$\min_{\substack{\mathbf{a} \in \mathbb{R}^n \\ \text{supp}(\mathbf{a}) \subseteq \mathcal{K}}} d_2(\mathbf{a}\mathbf{B}, \mathbf{1}) \geq \sqrt{\left\lfloor \frac{s}{d} \right\rfloor}. \quad (5)$$

*Proof:* Associate a bipartite graph  $G = (\mathcal{W} \cup \mathcal{P}, \mathcal{E})$  with  $\mathbf{B}$  as follows. Consider left vertices  $\mathcal{W} = \{W_1, W_2, \dots, W_n\}$  corresponding to workers, and right vertices  $\mathcal{P} = \{P_1, \dots, P_n\}$  corresponding to data parts. We draw an edge  $\{W_i, P_j\} \in \mathcal{E}$  between worker  $W_i$  and data part  $P_j$  if  $(\mathbf{B})_{i,j} \neq 0$ . According to Lemma 32, for any  $s$  such that  $d < s < n$  there exists a set  $\mathcal{Q} \subseteq \mathcal{P}$  of size  $\lfloor \frac{s}{d} \rfloor$  such that  $|\mathcal{N}(\mathcal{Q})| \leq d \cdot \lfloor \frac{s}{d} \rfloor \leq s$ .

Now, let  $\mathcal{K}$  be any set of size  $n - s$  which is contained in  $\mathcal{W} \setminus \mathcal{N}(\mathcal{Q})$ ; this set exists since  $|\mathcal{W} \setminus \mathcal{N}(\mathcal{Q})| = n - |\mathcal{N}(\mathcal{Q})| \geq n - s$ . In addition, let  $\mathbf{a} = (a_j)_{j=1}^n \in \mathbb{R}^n$  be any vector with  $\text{supp}(\mathbf{a}) \subseteq \mathcal{K}$ . Notice that  $(\mathbf{B})_{\ell,j} = 0$  whenever  $\ell \notin \mathcal{N}(\mathcal{Q})$  and  $j \in \mathcal{Q}$ , and  $a_\ell = 0$  whenever  $\ell \in \mathcal{N}(\mathcal{Q})$ . Thus, for  $j \in \mathcal{Q}$  we have

$$\begin{aligned} (\mathbf{a}\mathbf{B})_j &= \sum_{\ell=1}^n a_\ell (\mathbf{B})_{\ell,j} \\ &= \sum_{\ell \in \mathcal{N}(\mathcal{Q})} a_\ell (\mathbf{B})_{\ell,j} + \sum_{\ell \notin \mathcal{N}(\mathcal{Q})} a_\ell \mathbf{B}_{\ell,j} = 0, \end{aligned}$$

since in the summation over  $\ell \in \mathcal{N}(\mathcal{Q})$  we have  $a_\ell = 0$ , and in the summation over  $\ell \notin \mathcal{N}(\mathcal{Q})$  we have  $(\mathbf{B})_{\ell,j} = 0$ . Therefore,  $(\mathbf{a}\mathbf{B} - \mathbf{1})_j = -1$  for  $j \in \mathcal{Q}$ , and thus  $d_2(\mathbf{a}\mathbf{B}, \mathbf{1}) \geq \sqrt{|\mathcal{Q}|} \geq \sqrt{\left\lfloor \frac{s}{d} \right\rfloor}$ , which concludes the claim.  $\square$

The above lemma establishes the asymptotic optimality (up to constants) of our scheme, when used with Ramanujan graphs. Recall that for Ramanujan graphs we have  $\lambda \leq 2\sqrt{d}$ . Thus, using our proposed scheme, we get an  $\mathbf{a}$  and  $\mathbf{B}$  that satisfy the  $\epsilon$ -AC condition for  $\epsilon(s) \leq \frac{2}{\sqrt{d}} \sqrt{\frac{s}{1-(s/n)}}$  which tends to  $2\sqrt{\frac{s}{d}}$  as  $s/n \rightarrow 0$ .

### C. A Few Remarks About Convergence

While our results in Section V guarantee the exact computation of the gradient, and hence the convergence of the overall gradient descent algorithm, some care is needed to guarantee convergence when the gradient is approximated, as done in Section VI and Section VI-A. We operate under the standard assumption (e.g. [26]) that the arrival of a computation result from a server, within a predefined time frame, is modeled by a Bernoulli random variable. Then, the convergence of our algorithms can be guaranteed as a special case of the SGD algorithm [23, Sec. 14.3]; this applies for Section VI and Section VI-A, but for brevity we focus on Section VI. Furthermore, the  $\ell_2$ -deviations that were discussed above provide better convergence guarantees than the trivial algorithm, that can also be seen as a special case of SGD under identical assumptions.

Assume that every server in the system responds by some given deadline with probability  $q$ , i.e., it is a straggler with probability  $1 - q$ . Under this premise, it is shown that the expected value of  $\mathbf{v}_r$  in iteration  $r$  of Algorithm 1, when employing the expander scheme from Section VI, is the true gradient  $\nabla L_S(\mathbf{w}^{(r)})$  up to a constant factor  $c$  that is very close to one. Since the analysis is identical for every  $r$ , we omit this superscript from now on.

Let  $X_i$  be a Bernoulli( $q$ ) random variable that equals 1 if the  $i$ 'th server responded, and 0 otherwise. Further, if no server responded we let  $\mathbf{v} = \mathbf{0}$ . Therefore, the averaging of the workers' results by the master can be described using a random variable  $\mathbf{z}$  such that

$$\mathbf{z} = \begin{cases} \mathbf{0} & \text{if } \sum_{i=1}^n X_i = 0 \\ \frac{n}{\sum_{i=1}^n X_i} (X_1, \dots, X_n) & \text{else.} \end{cases}$$

That is, the added gradient is  $\mathbf{v} = \frac{1}{c} \mathbf{z} \mathbf{B} \mathbf{N}(\mathbf{w})$  (see Remark 24, whether  $\mathbf{B} = \frac{1}{d} \mathbf{A}_G$  in our algorithm or  $\mathbf{B} = \mathbf{I}$  in the trivial one), for a constant  $c$  which satisfies  $\mathbb{E} \mathbf{z} = c \mathbf{1}$  and is given by the following lemma.

*Lemma 34:* For  $\mathbf{z}$  and  $q$  defined as above, we have that  $\mathbb{E} \mathbf{z} = (1 - (1 - q)^n) \cdot \mathbf{1}$ .

*Proof:* We have that

$$\mathbb{E} \mathbf{z} = \mathbf{0} \cdot (1 - q)^n + \sum_{\mathcal{F} \subseteq [n], \mathcal{F} \neq \emptyset} q^{|\mathcal{F}|} (1 - q)^{n-|\mathcal{F}|} \cdot \frac{n}{|\mathcal{F}|} \cdot \mathbf{1}_{\mathcal{F}},$$

where  $\mathbb{1}_{\mathcal{F}}$  is the characteristic vector of a set  $\mathcal{F}$ . Therefore, for every  $j \in [n]$ , the  $j$ 'th entry of  $\mathbb{E}z$  equals

$$\sum_{\mathcal{F} \subseteq [n], \mathcal{F} \neq \emptyset} q^{|\mathcal{F}|} (1-q)^{n-|\mathcal{F}|} \cdot \frac{n}{|\mathcal{F}|} \cdot \mathbb{1}(j \in \mathcal{F}),$$

where  $\mathbb{1}(j \in \mathcal{F})$  is a 0-1 indicator for the truth value of “ $j \in \mathcal{F}$ ”. Re-indexing the above sum by all possible sizes of a set  $\mathcal{F}$  that contains  $j$ , we have

$$\begin{aligned} \sum_{i=1}^n \binom{n-1}{i-1} q^i (1-q)^{n-i} \cdot \frac{n}{i} &= \sum_{i=1}^n \binom{n}{i} q^i (1-q)^{n-i} \\ &= \sum_{i=0}^n \binom{n}{i} q^i (1-q)^{n-i} - (1-q)^n \\ &= 1 - (1-q)^n, \end{aligned}$$

where the last inequality follows from the Binomial Theorem. Hence, it follows that  $\mathbb{E}z = (1 - (1-q)^n) \cdot \mathbb{1}$ , which concludes the claim.  $\square$

Therefore, we have that  $c = 1 - (1-q)^n$ . It is widely known that if the gradient update in gradient descent is taken from a probability distribution whose expectation is the true gradient, then convergence is guaranteed [23, Sec. 14.3]. Therefore, according to Lemma 34, and since  $\mathbf{B}$  and  $\mathbf{N}(\mathbf{w})$  do not depend on  $\mathcal{K}$ , it follows that

$$\mathbb{E}v = \frac{1}{c}(\mathbb{E}z) \cdot \mathbf{B}\mathbf{N}(\mathbf{w}) = \mathbb{1}\mathbf{B}\mathbf{N}(\mathbf{w}) \stackrel{(*)}{=} \mathbb{1}\mathbf{N}(\mathbf{w}) = \nabla L_S(\mathbf{w}),$$

where  $(*)$  follows since  $\mathbb{1}$  is an eigenvector for eigenvalue 1, whether  $\mathbf{B} = \frac{1}{d}\mathbf{A}_G$  or  $\mathbf{B} = \mathbf{I}$ . Hence, both the trivial algorithm and the one in Section VI converge. However, in what follows it is shown that the improved  $\ell_2$ -deviation guarantees of our algorithm, with respect to the trivial one, provide better convergence guarantees for an important family of functions.

For  $\beta \in \mathbb{R}$ , a continuously differentiable function  $f$  is  $\beta$ -smooth in a domain  $\mathcal{X}$  if  $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_2 \leq \beta\|\mathbf{x} - \mathbf{y}\|_2$  for every  $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ . The following theorem applies whenever  $L_S$  is  $\beta$ -smooth and convex, and the learning algorithm is constrained to some domain  $\mathcal{X}$ . For the sake of citing the following theorem, notice that the vector  $v$  in Algorithm 1 is a (stochastic) function of  $\mathbf{w}$ , and thus we write  $v(\mathbf{w})$ . In addition, we let  $\mathbf{w}^*$  be the minimizer of  $L_S$  in  $\mathcal{X}$ .

**Theorem 35:** ([3, Theorem 6.3]<sup>3</sup>) For  $R^2 \triangleq \sup_{\mathbf{w} \in \mathcal{X}} \frac{1}{2}\|\mathbf{w}\|_2^2 - \|\mathbf{w}^{(1)}\|_2^2$ , if  $\mathbb{E}\|\nabla L_S(\mathbf{w}) - v(\mathbf{w})\|_2^2 \leq \sigma^2$  for every  $\mathbf{w} \in \mathcal{X}$  and  $\mathbb{E}v(\mathbf{w}) = \nabla L_S(\mathbf{w})$ , then Algorithm 1 with step size  $\frac{1}{(\beta+1/\eta)}$  and  $\eta \triangleq \frac{R}{\sigma} \sqrt{\frac{2}{t}}$  satisfies

$$\mathbb{E}L_S\left(\frac{1}{t} \sum_{s=1}^t \mathbf{w}^{(s+1)}\right) - L_S(\mathbf{w}^*) \leq R\sigma\sqrt{\frac{2}{t}} + \frac{\beta R^2}{t}.$$

Let  $v_{\text{triv}}(\mathbf{w}) = \frac{1}{c}\mathbf{z}\mathbf{N}(\mathbf{w})$  be the gradient update in the trivial scheme, and let  $v_{\text{exp}}(\mathbf{w}) = \frac{1}{c}\mathbf{z}\mathbf{B}\mathbf{N}(\mathbf{w})$  be the one in our algorithm, where  $\mathbf{B} = \frac{1}{d}\mathbf{A}_G$ . Since the convergence guarantees in Theorem 35 are largely determined by the parameter  $\sigma$ ,

we show that a smaller  $\sigma$  can be obtained in our algorithm. Specifically, since  $\nabla L_S(\mathbf{w}) = \mathbb{1}\mathbf{N}(\mathbf{w})$ , we have that

$$\begin{aligned} \mathbb{E}\|\nabla L_S(\mathbf{w}) - v_{\text{triv}}(\mathbf{w})\|_2^2 &\leq \|\mathbf{N}(\mathbf{w})\|_{\text{spec}}^2 \cdot \mathbb{E}\|\mathbb{1} - \frac{1}{c}\mathbf{z}\|_2^2, \text{ and} \\ \mathbb{E}\|\nabla L_S(\mathbf{w}) - v_{\text{exp}}(\mathbf{w})\|_2^2 &\leq \|\mathbf{N}(\mathbf{w})\|_{\text{spec}}^2 \cdot \mathbb{E}\|\mathbb{1} - \frac{1}{c}\mathbf{z}\mathbf{B}\|_2^2. \end{aligned}$$

The following lemma shows that our algorithm provides better convergence guarantees, and is proved in Appendix D.

**Lemma 36:** We have that

$$\mathbb{E}\|\mathbb{1} - \frac{1}{c}\mathbf{z}\mathbf{B}\|_2^2 \leq \mathbb{E}\|\mathbb{1} - \frac{1}{c}\mathbf{z}\|_2^2 - n \cdot \left(1 - \frac{\lambda^2}{d^2}\right) \frac{1-q}{q} (1 - o(1)).$$

## VII. EXPERIMENTAL RESULTS

In this section, we present experimental results of our proposed approximate gradient coding scheme (Sec. VI).

### A. $\ell_2$ -Error

We measured the performance of our approximate coding schemes in terms of the  $\ell_2$ -error for the recovery of  $\mathbb{1}$ . We chose the normalized adjacency matrix of a random  $d$ -regular graph on  $n$  vertices as the matrix  $\mathbf{B}$ . We randomly chose  $n - s$  rows of  $\mathbf{B}$  to be the surviving workers in any particular iteration, where  $s$  is the number of stragglers. For the decoding vector  $a(\mathcal{K})$ , we chose the vector in (4), (called the *linear decoder*), as well as the optimal least squares solution (called the *optimal decoder*):

$$a(\mathcal{K}) = \min_{\mathbf{a} \in \mathbb{R}^{n-s}} \|\mathbf{a}\mathbf{B}(\mathcal{K}, :) - \mathbb{1}\|_2, \quad (6)$$

where  $\mathbf{B}(\mathcal{K}, :)$  is the submatrix of  $\mathbf{B}$  which consists of the rows that are indexed by  $\mathcal{K}$ . Note that even though we have no additional theoretical guarantees for the optimal decoder, it is always possible to compute it in cubic time, e.g., by the singular value decomposition of  $\mathbf{B}(\mathcal{K}, :)$ .

Figure 2 presents the results using graphs on  $n = 30, 50$  vertices, and various values of  $s$  and  $d$ . The results shown are averaged over multiple samples of  $\mathcal{K}$  and multiple draws of the matrix  $\mathbf{B}$ . Figures 2(a), 2(b), 2(d), and 2(e) show the  $\ell_2$ -error versus number of stragglers  $s$ . As the number of stragglers increases, the recovery gets worse for a fixed degree  $d$ .

Figures 2(c) and 2(f) show the  $\ell_2$ -error versus the degree  $d$ . As  $d$  increases, the recovery error gets better for a fixed number of stragglers  $s$ . Also, as expected, in all cases, the optimal decoder does better than the linear decoder in terms of the  $\ell_2$ -error. Interestingly, we can also observe that on average both the linear decoder and the optimal decoder are better than the theoretical upper bound in our paper. One could even think of exploiting this empirically by randomizing the assignment of the rows of  $\mathbf{B}$  to the different workers in every iteration.

### B. Generalization Error

In this section, our approximate gradient coding (AGC) scheme is compared to other approaches. We compare against gradient coding from [24] (GC), as well as the trivial scheme (IS), where the data is divided equally among all workers, but the master only uses the first  $n - s$  gradients.

We measured the performance of our coding schemes in terms of the area under the curve (AUC) on a validation set

<sup>3</sup>This theorem is stated in much greater generality in [3], but cited herein as a special case for simplicity.

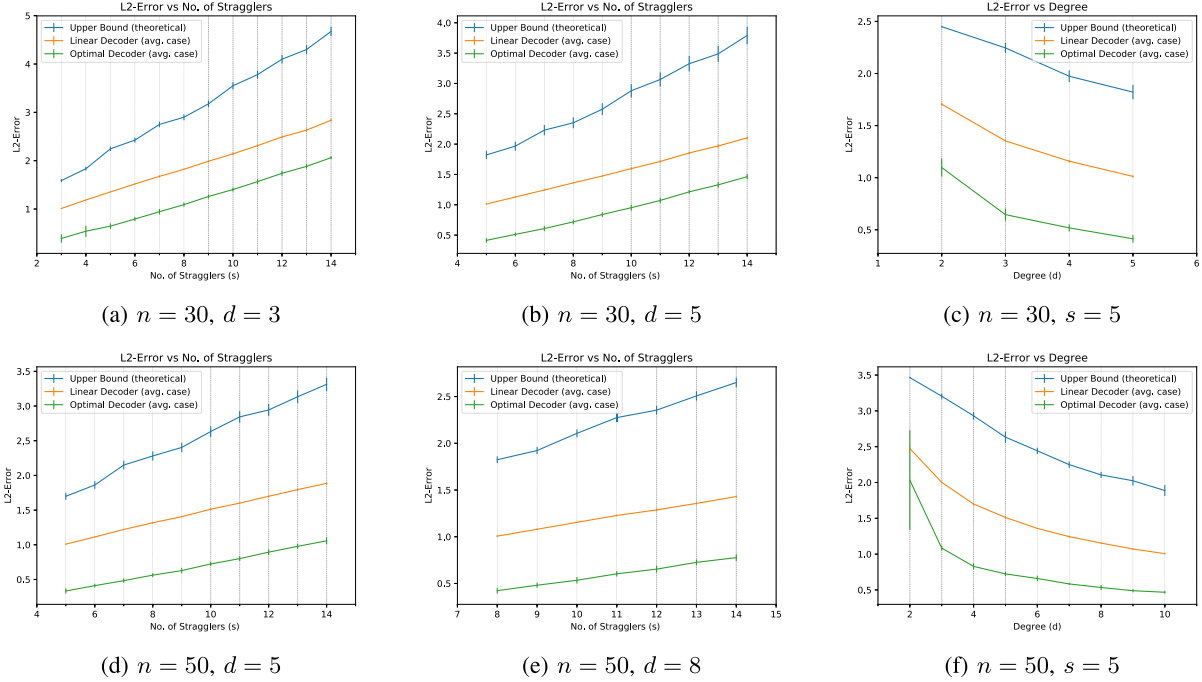


Fig. 2.  $\ell_2$ -error for recovery of 1 using normalized adjacency matrices of random  $d$ -regular graphs.




Fig. 3. Generalization error versus number of iterations using  $n = 30$  t2.micro worker instances on EC2, with  $d = 3$ , and  $s = 5$ . Note that in case of *Gradient Coding* [24], the computational overhead here is  $\times 6$  times (instead of  $\times 3$  in our approach).

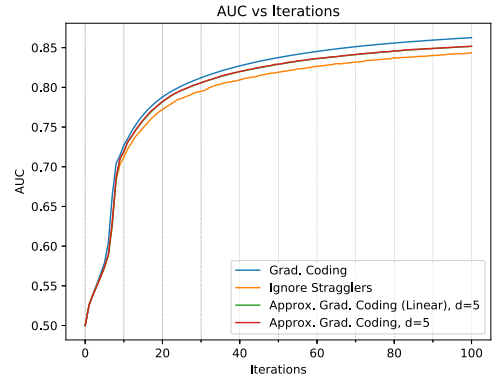


Fig. 4. Generalization error versus number of iterations using  $n = 50$  t2.micro worker instances on EC2, with  $d = 5$ , and  $s = 10$ . Note that in case of *Gradient Coding* [24], the computational overhead here is  $\times 11$  times (instead of  $\times 5$  in our approach).

for a logistic regression problem, on a real dataset. The dataset we used was the Amazon Employee dataset from Kaggle. We used 26,200 training samples, and a model dimension of 241,915 (after one-shot encoding with interaction terms), and used gradient descent to train the logistic regression. For GC we used a constant learning rate, chosen using cross-validation. For AGC and IS we used a learning rate of  $c_1/(r + c_2)$ , which is typical for SGD, where  $c_1$  and  $c_2$  were also chosen via cross-validation.

All our methods were implemented in python using MPI4py (similar to [24]). We ran our experiments using t2.micro worker instance types on Amazon EC2 and a c3.8xlarge master instance type. The results for  $n = 30, 50$  are given in Fig. 3 and Fig. 4, in which AGC corresponds to our approximation schemes with the optimal decoder, whereas

AGC (Linear), termed AGCL is our full proposed approximation scheme.

We observe that both these approaches are only slightly worse than GC, which utilizes the full gradient, and are quite better than the IS approach. Compared to each other, AGC and AGCL seem equivalent, however AGC was marginally better. That being said, AGCL can be faster since computing the linear decoder only requires  $O(n)$  time, in contrast to  $O(n^3)$  time for the optimal decoder.

## APPENDIX A

In this section, efficient algorithms for encoding (i.e., computing the matrix  $\mathbf{B}$ ) and decoding (i.e., computing the vector  $a(\mathcal{K})$  given a set  $\mathcal{K}$  of non-stragglers) are given for the scheme in Section V-A.



Since the matrix  $\mathbf{B}$  is circulant, it suffices to compute only its leftmost column. Further, the leftmost column  $\mathbf{c}_1^\top$  of  $\mathbf{B}$  is a codeword in an  $[n, n-s]$  Reed-Solomon code whose evaluation points are all roots of unity of order  $n$ , denoted  $\{\alpha_i\}_{i=0}^{n-1}$ . Hence, to find  $\mathbf{c}_1$ , one can define the polynomial  $m(x) \triangleq \prod_{j=s+1}^{n-1} (x - \alpha_j)$  and evaluate it over  $\alpha_0, \alpha_1, \dots, \alpha_s$ , which is possible in  $O(s(n-s))$  operations. This compares favorably with the respective  $O(n^2 \log^2(n))$  of [7] (Sec. 5.1.1) for any  $s$ .

As for decoding, for a given set  $\mathcal{K}$  of  $n-s$  non-stragglers we present an algorithm which computes  $a(\mathcal{K})$  in  $O(s \log^2 s + n \log n)$  operations over  $\mathbb{C}$ . This outperforms the respective  $O((n-s) \log^2(n-s))$  in [7] (Sec. 5.2.1) whenever  $s = o(n)$ , and the respective  $O((n-s)^2)$  of [8] for any  $s < \delta n$  with  $\delta < 1$ . The central tool in this section is the well-known *Generalized Reed-Solomon* (GRS) codes. A code  $\mathcal{C} \subseteq \mathbb{C}^n$  is called an  $[n, s]$  GRS code if

$$\mathcal{C} = \{(\ell_1 f(\alpha_1), \ell_2 f(\alpha_2), \dots, \ell_n f(\alpha_n)) : f \in \mathbb{C}[x]^{<s}\},$$

where  $\{\alpha_i\}_{i=1}^n$  are pairwise distinct evaluation points, and  $\{\ell_i\}_{i=1}^n \subseteq \mathbb{C}$  are nonzero *column multipliers*. It is readily verified that any RS code (Section IV) is a GRS code.

**Lemma 37** ([22], Prop. 5.2): If  $\mathcal{C} \subseteq \mathbb{F}^n$  is a GRS code then its dual  $\mathcal{C}^\perp$  is a GRS code with identical evaluation points.

Let  $\mathcal{C}$  be the code from Subsection V-A, and notice that it is a GRS code whose column multipliers are all equal to 1. By Lemma 37, it follows that the generator matrix of  $\mathcal{C}^\perp$  is  $\mathbf{V} \cdot \mathbf{D}$ , where

$$\mathbf{V} = \begin{pmatrix} 1 & 1 & \dots & 1 \\ \alpha_1 & \alpha_2 & \dots & \alpha_n \\ \alpha_1^2 & \alpha_2^2 & \dots & \alpha_n^2 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_1^{s-1} & \alpha_2^{s-1} & \dots & \alpha_n^{s-1} \end{pmatrix}$$

and  $\mathbf{D} = \text{diag}(\ell_1, \dots, \ell_n)$  is a diagonal matrix which contains the nonzero column multipliers of  $\mathcal{C}^\perp$ . In Algorithm 2, for a subset  $\mathcal{K} \subseteq [n]$  let  $\mathbf{V}_{\mathcal{K}}$  be the matrix of *columns* of  $\mathbf{V}$  that are indexed by  $\mathcal{K}$ , and let  $\mathbf{D}_{\mathcal{K}} = \text{diag}(\ell_i)_{i \in \mathcal{K}}$ . In addition, for a vector  $\mathbf{x} \in \mathbb{C}^n$ , let  $\mathbf{x}_{\mathcal{K}}$  be the vector which results from deleting the entries of  $\mathbf{x}$  that are not in  $\mathcal{K}$ .

---

**Algorithm 2** Computing the Decoding Vector  $a(\mathcal{K})$  for the Scheme in Subsection V-A

---

**Data:** Any vector  $\mathbf{x}' \in \mathbb{C}^n$  such that  $\mathbf{x}'\mathbf{B} = \mathbf{1}$ .

**Input:** A set  $\mathcal{K} \subseteq [n]$  of  $n-s$  non-stragglers.

**Output:** A vector  $a(\mathcal{K})$  such that  $\text{supp}(a(\mathcal{K})) \subseteq \mathcal{K}$  and  $a(\mathcal{K})\mathbf{B} = \mathbf{1}$ .

Find  $\mathbf{f} \in \mathbb{C}^s$  such that  $\mathbf{f}\mathbf{V}_{\mathcal{K}^c} = -\mathbf{x}'_{\mathcal{K}^c} \cdot \mathbf{D}_{\mathcal{K}^c}^{-1}$ .

Let  $\mathbf{y} \triangleq \mathbf{f} \cdot \mathbf{V}\mathbf{D}$ .

**return**  $a(\mathcal{K}) = \mathbf{y} + \mathbf{x}'$

---

*Correctness of Algorithm 2:* Since  $\mathbf{y} \in \mathcal{C}^\perp$ , it follows that  $\mathbf{y}\mathbf{B} = \mathbf{0}$ , and hence  $a(\mathcal{K})\mathbf{B} = (\mathbf{y} + \mathbf{x}')\mathbf{B} = \mathbf{1}$ . In addition, since  $\mathbf{y} = \mathbf{f} \cdot \mathbf{V}\mathbf{D}$ , it follows that  $\mathbf{y}_{\mathcal{K}^c} = (\mathbf{f} \cdot \mathbf{V}\mathbf{D})_{\mathcal{K}^c} = \mathbf{f} \cdot \mathbf{V}_{\mathcal{K}^c}\mathbf{D}_{\mathcal{K}^c} = -\mathbf{x}'_{\mathcal{K}^c}$ , and thus  $y_i = -x'_i$  for every  $i \notin \mathcal{K}$ . Therefore, it follows that  $a(\mathcal{K})_i = y_i + x'_i = 0$  for every  $i \notin \mathcal{K}$ , which implies that  $\text{supp}(a(\mathcal{K})) \subseteq \mathcal{K}$ .  $\square$

*Complexity of Algorithm 2:* Notice that –

- 1) Since  $\mathbf{D}_{\mathcal{K}^c}$  is diagonal, computing its inverse requires  $s$  inverse operations in  $\mathbb{C}$ .
- 2) Solving the equation  $\mathbf{f}\mathbf{V}_{\mathcal{K}^c} = -\mathbf{x}'_{\mathcal{K}^c} \cdot \mathbf{D}_{\mathcal{K}^c}^{-1}$  amounts to an interpolation problem, i.e., finding a degree (at most)  $s-1$  polynomial which passes through  $s$  given points. This is possible in  $O(s \log^2 s)$  operations by [12].
- 3) Given  $\mathbf{f}$ , computing the product  $\mathbf{f}\mathbf{V}$  reduces to evaluation of a degree (at most)  $n$  polynomial on all roots of unity of order  $n$ . This is possible in  $O(n \log n)$  operations by utilizing the famous Fast Fourier Transform (FFT) [2].

Hence, the total complexity of Algorithm 2 is  $O(s \log^2 s + n \log n)$ .  $\square$

The pre-computation of  $\mathbf{x}'$  may be done by finding  $\mathbf{x}'' \in \mathbb{C}^{n-s}$  such that  $\mathbf{x}''\mathbf{B}' = \mathbf{1}$ , where  $\mathbf{B}'$  is the upper-left  $(n-s) \times (n-s)$  submatrix of  $\mathbf{B}$ , and padding  $\mathbf{x}''$  with  $s$  zeros. Since  $\mathbf{B}'$  is a lower-triangular matrix in which the support of every column is of size at most  $s+1$ , the equation  $\mathbf{x}''\mathbf{B}' = \mathbf{1}$  can be solved by a simple  $O(s^2 + s(n-s))$  back-substitution algorithm. The computation of  $\mathbf{x}'$  can be done at the encoding phase (described above), and at a comparable complexity.

## APPENDIX B

As in the complex number case (Subsection V-A and Appendix A), an algorithm for computing the matrix  $\mathbf{B}$  and the vector  $a(\mathcal{K})$  for the scheme in Subsection V-B is given. In either of the cases of Construction 18, the resulting code  $\mathcal{C}$  consists of all codewords (seen as coefficients of polynomials in  $\mathbb{R}[x]^{<n}$ ) with  $s$  mutual roots. That is, the code can be described as the right kernel of

$$\mathbf{V} \triangleq \begin{pmatrix} 1 & \alpha_1 & \alpha_1^2 & \dots & \alpha_1^{n-1} \\ 1 & \alpha_2 & \alpha_2^2 & \dots & \alpha_2^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha_s & \alpha_s^2 & \dots & \alpha_s^{n-1} \end{pmatrix},$$

where  $\{\alpha_i\}_{i=1}^s$  are the roots of the code.

As in Appendix A, to compute the matrix  $\mathbf{B}$  it suffices to find the codeword  $\mathbf{c}_1$ , which in this case is a lowest weight codeword in a BCH code. It is readily verified that  $\mathbf{c}_1$  may be given by the coefficients of the *generator polynomial*  $g(x) \triangleq \prod_{i=1}^s (x - \alpha_i)$  (denoted by  $p_1$  and  $p_2$  in Eq. (2)). Finding these coefficients is possible by evaluating  $g(x)$  in  $s+1$  arbitrary and pairwise distinct roots of unity of order  $n$ , and solving an interpolation problem in  $O(s \log^2 s)$  by [12]. However, evaluating  $g(x)$  at  $s+1$  points requires  $O(n \log n)$  operations using the FFT algorithm. Hence, the overall complexity of computing  $\mathbf{B}$  is  $O(\min\{s \log^2 s, n \log n\})$ , an improvement over [7] whenever  $s = o(n)$ .

In Algorithm 3, for a given  $\mathcal{K} \subseteq [n]$ , let  $\mathbf{V}_{\mathcal{K}}$  be the matrix of columns of  $\mathbf{V}$  that are indexed by  $\mathcal{K}$ . The complexity of this algorithm outperforms [7] whenever  $s = o(\log^2 n)$ , and outperforms [8] whenever  $s = o(n^{2/3})$ .

*Correctness of Algorithm 3:* First, note that  $\mathbf{V}_{\mathcal{K}^c}$  is an invertible matrix, since  $\mathcal{C}^\perp$  is an MDS code by Lemma 8 and Lemma 19. Second, since  $\mathbf{y}$  is in the left image of  $\mathbf{V}$  it follows

**Algorithm 3** Computing the Decoding Vector  $a(\mathcal{K})$  for the Scheme in Subsection V-B

**Data:** Any vector  $\mathbf{x}' \in \mathbb{R}^n$  such that  $\mathbf{x}'\mathbf{B} = \mathbf{1}$ .

**Input:** A set  $\mathcal{K} \subseteq [n]$  of  $n - s$  non-stragglers.

**Output:** A vector  $a(\mathcal{K})$  such that  $\text{supp}(a(\mathcal{K})) \subseteq \mathcal{K}$  and  $a(\mathcal{K})\mathbf{B} = \mathbf{1}$ .

Compute  $\mathbf{V}_{\mathcal{K}^c}^{-1}$ .

Let  $\mathbf{f} \triangleq -\mathbf{x}'_{\mathcal{K}^c} \cdot \mathbf{V}_{\mathcal{K}^c}^{-1}$ .

Let  $\mathbf{y} \triangleq \mathbf{f} \cdot \mathbf{V}$ .

**return**  $a(\mathcal{K}) = \mathbf{y} + \mathbf{x}'$ .

that  $\mathbf{y} \in \mathcal{C}^\perp$ , and further,  $y_i = -x'_i$  for every  $i \in \mathcal{K}^c$ . Hence,  $\text{supp}(a(\mathcal{K})) \subseteq \mathcal{K}$  and  $a(\mathcal{K})\mathbf{B} = \mathbf{1}$ .  $\square$

*Complexity of Algorithm 3:* As in the complexity analysis of Algorithm 2, the complexity is clearly  $O(\gamma_s + s(n - s))$ , where  $\gamma_s$  is the complexity of inverting a generalized Vandermonde matrix. While explicit formulas for inverting a generalized Vandermonde matrix were discussed in [25], for simplicity, one may employ the trivial  $O(s^3)$  algorithm.  $\square$

#### APPENDIX C

##### BANDWIDTH REDUCTION IN SUBSECTION V-A

In step  $r \in [t]$  of Algorithm 1, each node  $W_j$  transmits  $\frac{1}{n} \sum_{i \in \text{supp}(\mathbf{b}_j)} b_{j,i} \cdot \nabla L_{\mathcal{S}_i}(\mathbf{w}^{(r)}) = \mathbf{b}_j \cdot \mathbf{N}(\mathbf{w}^{(r)})$  back to the master node. In the scheme which is suggested in Subsection V-A, this step requires transmitting a vector in  $\mathbb{C}^p$ , where  $p$  is the number of entries of  $\mathbf{w}^{(r)}$ . Since the gradient vectors  $\nabla L_{\mathcal{S}_i}(\mathbf{w}^{(r)})$  are real, the resulting bandwidth is larger than the actual amount of information by a factor of 2. In this section it is shown that the optimal bandwidth can be attained by a simple manipulation of the gradient vectors.

For  $p' \triangleq \lceil \frac{p}{2} \rceil$  and any  $r \in [t]$ , denote the columns of the matrix  $\mathbf{N}(\mathbf{w}^{(r)}) \in \mathbb{R}^{n \times p}$  by  $\mathbf{N}_{r,1}, \dots, \mathbf{N}_{r,p}$ , and let  $\mathbf{N}'_r \in \mathbb{C}^{n \times p'}$  be the matrix whose columns are<sup>4</sup>  $\mathbf{N}_{r,1} + i\mathbf{N}_{r,2}, \mathbf{N}_{r,3} + i\mathbf{N}_{r,4}, \dots$ . To obtain optimal bandwidth, replace the transmission  $\mathbf{b}_j \cdot \mathbf{N}(\mathbf{w}^{(r)}) \in \mathbb{C}^p$  of  $W_j$  at iteration  $r$  of Algorithm 1 by  $\mathbf{b}_j \cdot \mathbf{N}'_r \in \mathbb{C}^{p'}$ . In addition, replace the definition of  $\mathbf{w}^{(r+1)}$  by  $\mathbf{w}^{(r)} - \eta e(\mathbf{v}_r)$ , for  $e: \mathbb{C}^{p'} \rightarrow \mathbb{R}^p$  that is defined as

$$e(\mathbf{v}) \triangleq (\Re(v_1), \Im(v_1), \Re(v_2), \Im(v_2), \dots),$$

where  $\Re(y)$  and  $\Im(y)$  are the real and imaginary parts of a given  $y \in \mathbb{C}$ , respectively.

Since  $\mathbf{a}$  and  $\mathbf{B}$  of Subsection V-A satisfy the EC condition, it follows that  $a(\mathcal{K})\mathbf{B}\mathbf{N}'_r = \mathbf{1} \cdot \mathbf{N}'_r$  for every set  $\mathcal{K}$  of  $n - s$  non-stragglers and every  $r \in [t]$ . Further, we have that

$$\begin{aligned} e(\mathbf{1}\mathbf{N}'_r) &= e(\mathbf{1}(\mathbf{N}_{r,1} + i\mathbf{N}_{r,2}), \mathbf{1}(\mathbf{N}_{r,3} + i\mathbf{N}_{r,4}), \dots) \\ &= e(\mathbf{1}\mathbf{N}_{r,1} + i\mathbf{1}\mathbf{N}_{r,2}, \mathbf{1}\mathbf{N}_{r,3} + i\mathbf{1}\mathbf{N}_{r,4}, \dots) \\ &= (\mathbf{1}\mathbf{N}_{r,1}, \mathbf{1}\mathbf{N}_{r,2}, \mathbf{1}\mathbf{N}_{r,3}, \dots) = \mathbf{1}\mathbf{N}(\mathbf{w}^{(r)}), \end{aligned}$$

and hence the correctness of Algorithm 1 under the scheme from Subsection V-A is preserved. Under this framework, the transmission from  $W_j$  to the master node contains  $p'$  complex numbers (i.e.,  $2p'$  real numbers) rather than  $p$

complex numbers (i.e.,  $2p$  real numbers), and it is clearly optimal.

Note that this improvement does not diminish the contribution of the scheme in Subsection V-B. In the current section, to compute the transmission  $\mathbf{b}_j\mathbf{N}'_r$ , any worker node  $W_j$  performs  $dp'$  multiplication operations over  $\mathbb{C}$ , i.e.,  $4dp' \approx 2dp$  multiplication operations over  $\mathbb{R}$ . In contrast, the scheme in Subsection V-B requires any  $W_j$  to perform only  $dp$  multiplications over  $\mathbb{R}$ .

#### APPENDIX D

We operate under the convention that if all servers are stragglers (i.e.,  $\mathcal{K} = \emptyset$ ), then the algorithm outputs the vector  $\mathbf{0}$  as the approximation of the gradient. To analyze the trivial algorithm under this convention, notice that the random variable  $z$  receives the value

$$Y_{\mathcal{K}} = \begin{cases} \frac{n}{|\mathcal{K}|} \cdot \mathbf{1}_{\mathcal{K}} & \text{if } \mathcal{K} \neq \emptyset \\ \mathbf{0} & \text{else,} \end{cases}$$

whenever the set of non-stragglers is  $\mathcal{K}$ . Further, notice that if  $\mathcal{K} = \emptyset$  then  $\|\mathbf{1} - \frac{1}{c}Y_{\mathcal{K}}\|_2^2 = n$ , and otherwise

$$\begin{aligned} &\|\mathbf{1} - \frac{1}{c}Y_{\mathcal{K}}\|_2^2 \\ &= (\mathbf{1} - Y_{\mathcal{K}} + (1 - \frac{1}{c})Y_{\mathcal{K}})(\mathbf{1} - Y_{\mathcal{K}} + (1 - \frac{1}{c})Y_{\mathcal{K}})^\top \\ &= (\mathbf{1} - Y_{\mathcal{K}})(\mathbf{1} - Y_{\mathcal{K}})^\top + 2(1 - \frac{1}{c})(\mathbf{1} - Y_{\mathcal{K}})Y_{\mathcal{K}}^\top \\ &\quad + (1 - \frac{1}{c})^2\|Y_{\mathcal{K}}\|_2^2 \\ &\stackrel{(a)}{=} \frac{ns}{n-s} + 2(1 - \frac{1}{c})\mathbf{1}Y_{\mathcal{K}}^\top + (\frac{1}{c^2} - 1)\|Y_{\mathcal{K}}\|_2^2 \\ &\stackrel{(b)}{=} \frac{ns}{n-s} + 2n(1 - \frac{1}{c}) + (\frac{1}{c^2} - 1)\frac{n^2}{n-s}, \end{aligned} \tag{7}$$

where (a) follows from Corollary 21 since  $\mathbf{1} - Y_{\mathcal{K}} = -\mathbf{u}_{\mathcal{K}}$ , and (b) follows from the definition of  $Y_{\mathcal{K}}$ . To conduct similar analysis on our algorithm, we prove the following technical lemma.

*Lemma 38:* The following statements hold.

- 1) For vectors  $\mathbf{a}, \mathbf{b}$  and a constant  $r > 0$ , if  $\|\mathbf{a} - \mathbf{b}\|_2^2 \leq \alpha$  then  $\|\mathbf{a} - r\mathbf{b}\|_2^2 \leq \alpha + 2(1 - r)\mathbf{a}\mathbf{b}^\top + (r^2 - 1)\|\mathbf{b}\|_2^2$ .
- 2)  $\mathbf{1}((\mathbf{1} + \mathbf{u}_{\mathcal{K}})\mathbf{B})^\top = n$ .
- 3)  $\|(\mathbf{1} + \mathbf{u}_{\mathcal{K}})\mathbf{B}\|_2^2 \leq \frac{n^2}{n-s}$ .

*Proof:*

- 1) We have

$$\begin{aligned} &(\mathbf{a} - r\mathbf{b})(\mathbf{a} - r\mathbf{b})^\top \\ &= (\mathbf{a} - \mathbf{b} + (1 - r)\mathbf{b})(\mathbf{a} - \mathbf{b} + (1 - r)\mathbf{b})^\top \\ &= (\mathbf{a} - \mathbf{b})(\mathbf{a} - \mathbf{b})^\top + 2(1 - r)(\mathbf{a} - \mathbf{b})\mathbf{b}^\top + (1 - r)^2\|\mathbf{b}\|_2^2 \\ &\leq \alpha + 2(1 - r)\mathbf{a}\mathbf{b}^\top + (r^2 - 1)\|\mathbf{b}\|_2^2. \end{aligned}$$

- 2) Since  $d\mathbf{1}$  is an eigenvector of  $\mathbf{A}_G$ , and since  $\mathbf{1}\mathbf{u}_{\mathcal{K}}^\top = 0$ , we have

$$\begin{aligned} \mathbf{1}((\mathbf{1} + \mathbf{u}_{\mathcal{K}})\mathbf{B})^\top &= \mathbf{1}\mathbf{B}^\top(\mathbf{1} + \mathbf{u}_{\mathcal{K}})^\top = \frac{1}{d}\mathbf{1}\mathbf{A}_G^\top(\mathbf{1} + \mathbf{u}_{\mathcal{K}})^\top \\ &= \mathbf{1}(\mathbf{1} + \mathbf{u}_{\mathcal{K}})^\top = n. \end{aligned}$$

- 3) We have

$$\begin{aligned} (\mathbf{1} + \mathbf{u}_{\mathcal{K}})\mathbf{B} \cdot ((\mathbf{1} + \mathbf{u}_{\mathcal{K}})\mathbf{B})^\top &= (\mathbf{1} + \mathbf{u}_{\mathcal{K}})\mathbf{B}\mathbf{B}^\top(\mathbf{1} + \mathbf{u}_{\mathcal{K}})^\top \\ &= \frac{1}{d^2}(\mathbf{1} + \mathbf{u}_{\mathcal{K}})\mathbf{A}_G^2(\mathbf{1} + \mathbf{u}_{\mathcal{K}})^\top \end{aligned}$$

<sup>4</sup>The rightmost column of  $\mathbf{N}'_r$  is  $\mathbf{N}_{r,p-1} + i\mathbf{N}_{r,p}$  if  $p$  is even and  $\mathbf{N}_{r,p}$  otherwise.

$$\begin{aligned}
&= \frac{1}{d^2} (\mathbb{1} \mathbf{A}_G^2 \mathbb{1}^\top + 2 \mathbb{1} \mathbf{A}_G^2 \mathbf{u}_K^\top + \mathbf{u}_K \mathbf{A}_G^2 \mathbf{u}_K^\top) \\
&= \frac{1}{d^2} (d^2 n + \|\mathbf{u}_K \mathbf{A}_G\|_2^2). \tag{8}
\end{aligned}$$

Now, notice that  $\|\mathbf{u}_K \mathbf{A}_G\|_2^2 \leq \|\mathbf{u}_K\|_2^2 \cdot \|\mathbf{A}_G\|_{\text{spec}}^2$ , and recall that the spectral norm of an adjacency matrix is bounded from above by the maximum degree. Therefore, it follows from Corollary 21 that

$$(8) \leq n + \frac{1}{d^2} \cdot \frac{ns}{n-s} \cdot d^2 = n + \frac{ns}{n-s} = \frac{n^2}{n-s}.$$

□

**Corollary 39:** If  $\mathcal{K} = \emptyset$  then  $\|\mathbb{1} - \frac{1}{c}(\mathbb{1} + \mathbf{u}_K)\mathbf{B}\|_2^2 = n$ , and otherwise

$$\|\mathbb{1} - \frac{1}{c}(\mathbb{1} + \mathbf{u}_K)\mathbf{B}\|_2^2 \leq \frac{\lambda^2}{d^2} \cdot \frac{ns}{n-s} + 2n(1 - \frac{1}{c}) + (\frac{1}{c^2} - 1) \frac{n^2}{n-s}$$

*Proof:* The first part is trivial since  $\mathbf{u}_K = -\mathbb{1}$  for  $\mathcal{K} = \emptyset$ . If  $\mathcal{K} \neq \emptyset$ ,

$$\begin{aligned}
&\|\mathbb{1} - \frac{1}{c}(\mathbb{1} + \mathbf{u}_K)\mathbf{B}\|_2^2 \\
&\stackrel{(a)}{\leq} \frac{\lambda^2}{d^2} \cdot \frac{ns}{n-s} + 2(1 - \frac{1}{c})\mathbb{1}((\mathbb{1} + \mathbf{u}_K)\mathbf{B})^\top \\
&\quad + (\frac{1}{c^2} - 1)\|(\mathbb{1} + \mathbf{u}_K)\mathbf{B}\|_2^2 \\
&\stackrel{(b)}{\leq} \frac{\lambda^2}{d^2} \cdot \frac{ns}{n-s} + 2n(1 - \frac{1}{c}) + (\frac{1}{c^2} - 1)\|(\mathbb{1} + \mathbf{u}_K)\mathbf{B}\|_2^2 \\
&\stackrel{(c)}{\leq} \frac{\lambda^2}{d^2} \cdot \frac{ns}{n-s} + 2n(1 - \frac{1}{c}) + (\frac{1}{c^2} - 1) \frac{n^2}{n-s},
\end{aligned}$$

where (a) follows from Lemma 22 and Lemma 38.1, (b) follows from Lemma 38.2, and (c) follows from Lemma 38.3. □

*Proof:* (of Lemma 36) According to (7) and Corollary 39, we have

$$\begin{aligned}
&\mathbb{E}\|\mathbb{1} - \frac{1}{c}\mathbf{z}\|_2^2 - \mathbb{E}\|\mathbb{1} - \frac{1}{c}\mathbf{z}\mathbf{B}\|_2^2 = \\
&\sum_{\mathcal{K} \subseteq [n]} \Pr(\mathcal{K}) \|\mathbb{1} - \frac{1}{c}\mathbf{Y}_K\|_2^2 - \\
&\sum_{\mathcal{K} \subseteq [n]} \Pr(\mathcal{K}) \|\mathbb{1} - \frac{1}{c}(\mathbb{1} + \mathbf{u}_K)\mathbf{B}\|_2^2 \\
&\geq \sum_{\mathcal{K} \neq \emptyset} \Pr(\mathcal{K}) \left( \frac{ns}{n-s} + 2n(1 - \frac{1}{c}) + (\frac{1}{c^2} - 1) \frac{n^2}{n-s} \right) \\
&\quad - \sum_{\mathcal{K} \neq \emptyset} \Pr(\mathcal{K}) \left( \frac{\lambda^2}{d^2} \cdot \frac{ns}{n-s} + 2n(1 - \frac{1}{c}) + (\frac{1}{c^2} - 1) \frac{n^2}{n-s} \right) \\
&= n \cdot \left( 1 - \frac{\lambda^2}{d^2} \right) \sum_{\mathcal{K} \neq \emptyset} \Pr(\mathcal{K}) \cdot \frac{n-|\mathcal{K}|}{|\mathcal{K}|}, \tag{9}
\end{aligned}$$

and since  $|\mathcal{K}|$  is distributed as a sum of independent Bernoulli( $q$ ) variables, we have

$$\begin{aligned}
(9) &= n \cdot \left( 1 - \frac{\lambda^2}{d^2} \right) \sum_{i=1}^{n-1} q^i (1-q)^{n-i} \binom{n}{i} \frac{n-i}{i} \\
&= n \cdot \left( 1 - \frac{\lambda^2}{d^2} \right) \sum_{i=1}^{n-1} q^i (1-q)^{n-i} \binom{n}{i+1} \frac{i+1}{i} \\
&\geq n \cdot \left( 1 - \frac{\lambda^2}{d^2} \right) \frac{1-q}{q} \sum_{i=1}^{n-1} q^{i+1} (1-q)^{n-(i+1)} \binom{n}{i+1} \\
&= n \cdot \left( 1 - \frac{\lambda^2}{d^2} \right) \frac{1-q}{q} \sum_{i=2}^n q^i (1-q)^{n-i} \binom{n}{i}
\end{aligned}$$

$$\begin{aligned}
&= n \cdot \left( 1 - \frac{\lambda^2}{d^2} \right) \frac{1-q}{q} \left( \sum_{i=0}^n q^i (1-q)^{n-i} \binom{n}{i} \right. \\
&\quad \left. - (1-q)^n - q(1-q)^{n-1}n \right) \\
&= n \cdot \left( 1 - \frac{\lambda^2}{d^2} \right) \frac{1-q}{q} (1 - (1-q)^n - q(1-q)^{n-1}n) \\
&= n \cdot \left( 1 - \frac{\lambda^2}{d^2} \right) \frac{1-q}{q} (1 - o(1)).
\end{aligned}$$

□

## ACKNOWLEDGMENT

The work of R. Tandon was done while he was at UT Austin, prior to joining Apple. The authors express their gratitude to Prof. R. Livni for his valuable input. Most of N. Raviv's work was done while he was a postdoctoral researcher at the Department of Electrical Engineering, California Institute of Technology.

## REFERENCES

- [1] Y. Bilu and N. Linial, "Lifts, discrepancy and nearly optimal spectral gap\*," *Combinatorica*, vol. 26, no. 5, pp. 495–519, Oct. 2006.
- [2] E. O. Brigham and E. Brigham, *The Fast Fourier Transform and Its Applications*, vol. 1. Englewood Cliffs, NJ, USA: Prentice-Hall, 1988.
- [3] S. Bubeck, "Convex optimization: Algorithms and complexity," *Found. Trends Mach. Learn.*, vol. 8, nos. 3–4, pp. 231–357, 2015.
- [4] Z. Charles, D. Papailiopoulos, and J. Ellenberg, "Approximate gradient coding via sparse random graphs," 2017, *arXiv:1711.06771*. [Online]. Available: <http://arxiv.org/abs/1711.06771>
- [5] J. Chen, X. Pan, R. Monga, S. Bengio, and R. Jozefowicz, "Revisiting distributed synchronous SGD," 2016, *arXiv:1604.00981*. [Online]. Available: <http://arxiv.org/abs/1604.00981>
- [6] M. B. Cohen, "Ramanujan graphs in polynomial time," in *Proc. IEEE 57th Annu. Symp. Found. Comput. Sci. (FOCS)*, Oct. 2016, pp. 276–281.
- [7] S. Dutta, V. Cadambe, and P. Grover, "Short-dot: Computing large linear transforms distributedly using coded short dot products," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 2100–2108.
- [8] W. Halbawi, "Error-correcting codes for networks, storage and computation," Ph.D. dissertation, Dept. Elect. Eng., California Inst. Technol., Pasadena, CA, USA, 2017.
- [9] W. Halbawi, N. Azizan, F. Salehi, and B. Hassibi, "Improving distributed gradient descent using Reed–Solomon codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2018, pp. 2027–2031.
- [10] S. Hoory, N. Linial, and A. Wigderson, "Expander graphs and their applications," *Bull. Amer. Math. Soc.*, vol. 43, pp. 439–561, Aug. 2006.
- [11] C. Karakus, Y. Sun, S. Diggavi, and W. Yin, "Straggler mitigation in distributed optimization through data encoding," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5440–5448.
- [12] H.-T. Kung, "Fast evaluation and interpolation," Dept. Comput. Sci., Carnegie-Mellon Univ., Pittsburgh, PA, USA, Tech. Rep., 1973.
- [13] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," *IEEE Trans. Inf. Theory*, vol. 64, no. 3, pp. 1514–1529, Mar. 2018.
- [14] M. Li, D. G. Andersen, A. Smola, and K. Yu, "Communication efficient distributed machine learning with the parameter server," in *Proc. 27th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, vol. 1. Cambridge, MA, USA: MIT Press, 2014, pp. 19–27. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2968826.2968829>
- [15] S. Li, S. M. Mousavi Kalan, A. S. Avestimehr, and M. Soltanolkotabi, "Near-optimal straggler mitigation for distributed gradient methods," in *Proc. IEEE Int. Parallel Distrib. Process. Symp. Workshops (IPDPSW)*, May 2018, pp. 857–866.
- [16] S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, "Coded MapReduce," in *Proc. 53rd Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Sep. 2015, pp. 964–971.
- [17] S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, "A unified coding framework for distributed computing with straggling servers," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, Dec. 2016, pp. 1–6.
- [18] S. Li, M. A. Maddah-Ali, Q. Yu, and A. S. Avestimehr, "A fundamental tradeoff between computation and communication in distributed computing," *IEEE Trans. Inf. Theory*, vol. 64, no. 1, pp. 109–128, Jan. 2018.



- [19] A. Lubotzky, R. Phillips, and P. Sarnak, "Ramanujan graphs," *Combinatorica*, vol. 8, no. 3, pp. 261–277, 1988.
- [20] T. Marshall, "Coding of real-number sequences for error correction: A digital signal processing problem," *IEEE J. Sel. Areas Commun.*, vol. 2, no. 2, pp. 381–392, Mar. 1984.
- [21] R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, Eds., *Machine Learning: An Artificial Intelligence Approach*, vol. 1. Palo Alto, CA, USA: Tioga, 1983.
- [22] R. Roth, *Introduction to Coding Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2006.
- [23] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2014.
- [24] R. Tandon, Q. Lei, A. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *Proc. 34th Int. Conf. Mach. Learn.*, Aug. 2017, pp. 3368–3376. [Online]. Available: <http://proceedings.mlr.press/v70/tdandon17a.html>
- [25] L. Verde-Star, "Inverses of generalized Vandermonde matrices," *J. Math. Anal. Appl.*, vol. 131, no. 2, pp. 341–353, May 1988.
- [26] H. Wang, Z. Charles, and D. Papailiopoulos, "ErasureHead: Distributed gradient descent without delays using approximate gradient coding," 2019, *arXiv:1901.09671*. [Online]. Available: <http://arxiv.org/abs/1901.09671>
- [27] N. J. Yadwadkar, B. Hariharan, J. E. Gonzalez, and R. Katz, "Multi-task learning for straggler avoiding predictive job scheduling," *J. Mach. Learn. Res.*, vol. 17, no. 106, pp. 1–37, 2016. [Online]. Available: <http://jmlr.org/papers/v17/15-149.html>
- [28] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 5610–5619.
- [29] Q. Yu, S. Li, N. Raviv, S. M. M. Kalan, M. Soltanolkotabi, and S. Avestimehr, "Lagrange coded computing: Optimal design for resiliency, security and privacy," 2018, *arXiv:1806.00939*. [Online]. Available: <http://arxiv.org/abs/1806.00939>

**Netanel Raviv** (Member, IEEE) received the B.Sc. degree in mathematics and computer science and the M.Sc. and Ph.D. degrees in computer science from Technion, Israel, in 2010, 2013, and 2017, respectively. He is currently an Assistant Professor with the Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO. He was an awardee of the IBM Ph.D. Fellowship for the academic year 2015–2016, the First Prize in the Feder family competition for best student work in communication technology, and the Lester-Deutsche Postdoctoral Fellowship. His research interests include applications of coding theory to privacy, distributed computations, and machine learning.

**Itzhak Tamo** (Member, IEEE) received the B.A. degree in mathematics and the B.Sc. and Ph.D. degrees in electrical engineering from Ben-Gurion University, Israel, in 2008 and 2012, respectively. From 2012 to 2014, he was a Post-Doctoral Researcher with the Institute for Systems Research, University of Maryland, College Park. Since 2015, he has been a Senior Lecturer with the Electrical Engineering Department, Tel-Aviv University, Israel. His research interests include storage systems and devices, coding, information theory, and combinatorics. He was a co-recipient (with Zhiying Wang and Jehoshua Bruck) of the IEEE Communication Society Data Storage Technical Committee 2013 Best Paper Award. He received the 2015 IEEE Information Theory Society Paper Award along with A. Barg. In 2018, he received the Krill Prize.

**Rashish Tandon** received the B.Tech and M.Tech degrees in computer science from the Indian Institute of Technology (IIT), Kanpur, in 2011, and the Ph.D. degree in computer science from the University of Texas at Austin, in 2017. He is currently a Senior Machine Learning Engineer with Apple. His research interests include algorithms for machine learning in distributed and high-dimensional settings.

**Alexandros G. Dimakis** (Senior Member, IEEE) received the Diploma degree from the National Technical University of Athens in 2003 and the Ph.D. degree in electrical engineering and computer sciences from UC Berkeley in 2008. He is currently a Professor with the Electrical and Computer Engineering Department, University of Texas at Austin. He received an ARO Young Investigator Award, the NSF Career, a Google Faculty Award, and the Eli Jury Dissertation Award. He was a co-recipient of several best paper awards including the joint Information Theory and Communications Society Best Paper Award and the James Massey Award in 2018. His research interests include information theory, coding theory, and machine learning.