Fundamental Limits of Approximate Gradient Coding

SINONG WANG, Department of ECE, The Ohio State University, USA

JIASHANG LIU, Departments of ECE, The Ohio State University, USA

NESS SHROFF, Departments of ECE and CSE, The Ohio State University, USA

In the distributed gradient coding problem, it has been established that, to exactly recover the gradient under s slow machines, the minimum computation load (number of stored data partitions) of each worker is at least linear $(s + 1)$, which incurs a large overhead when s is large [17]. In this paper, we focus on *approximate gradient coding* that aims to recover the gradient with bounded error ϵ . Theoretically, our main contributions are three-fold: (i) we analyze the structure of optimal gradient codes, and derive the information-theoretical lower bound of minimum computation load: $O(\log(n)/\log(n/s))$ for $\epsilon = 0$ and $d \geq O(\log(1/\epsilon)/\log(n/s))$ for $\epsilon > 0$, where d is the computation load, and ϵ is the error in the gradient computation; (ii) we design two approximate gradient coding schemes that exactly match such lower bounds based on random edge removal process; (iii) we implement our schemes and demonstrate the advantage of the approaches over the current fastest gradient coding strategies. The proposed schemes provide order-wise improvement over the state of the art in terms of computation load, and are also optimal in terms of both computation load and latency.

CCS Concepts: • **Mathematics of computing** Information theory; • **Theory of computation** Distributed algorithms; • **Computing methodologies** Parallel computing methodologies; *Machine learning*.

Additional Key Words and Phrases: datasets, neural networks, gaze detection, text tagging

ACM Reference Format:

Sinong Wang, Jiashang Liu, and Ness Shroff. 2019. Fundamental Limits of Approximate Gradient Coding. *Proc. ACM Meas. Anal. Comput. Syst.* 3, 3, Article 52 (December 2019), 22 pages. <https://doi.org/10.1145/3366700>

1 INTRODUCTION

Large-scale machine learning has shown great promise for solving many practical applications [7]. Such applications require massive training datasets and model parameters, and force practitioners to adopt distributed computing frameworks such as Hadoop [4] and Spark [23] to increase the learning speed. However, the speedup gain is far from ideal due to the latency incurred in waiting for a few slow or faulty processors, called “stragglers” to complete their tasks [3]. For example, it was observed in [20] that a straggler may run $8\times$ slower than the average worker performance on Amazon EC2. To alleviate the straggler issue, current frameworks such as Hadoop deploy various straggler detection techniques and usually replicate the straggling tasks on other available nodes.

Recently, *gradient coding* techniques have been proposed to provide an effective way to deal with stragglers for distributed learning applications [17]. The system being considered has n workers, in which the training data is partitioned into n parts. Each worker stores multiple parts of datasets,

Authors' addresses: Sinong Wang, Department of ECE, The Ohio State University, USA, Columbus, Ohio, 43220, wang.7691@osu.edu; Jiashang Liu, Departments of ECE, The Ohio State University, USA, Columbus, Ohio, 43220, liu.3992@osu.edu; Ness Shroff, Departments of ECE and CSE, The Ohio State University, USA, Columbus, Ohio, 43220, shroff.11@osu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

2476-1249/2019/12-ART52 \$15.00

<https://doi.org/10.1145/3366700>

computes a partial gradient over each of its assigned partitions, and returns the linear combination of these partial gradients to the master node. By creating and exploiting coding redundancy in local computation, the master node can reconstruct the full gradient even if only a part of results are collected, and therefore alleviate the impact of straggling workers.

The key performance metric used in gradient coding scheme is the *computation load* $d(s)$, which refers to the number of data partitions that are sent to each node, and characterizes the amount of redundant computations to resist s stragglers. Given the number of workers n and number of stragglers s , the work of [17] establishes a fundamental bound $d(s) \geq s + 1$, and constructs a random code that exactly matches this lower bound. Two subsequent works [5, 14] provide a deterministic construction of the gradient coding scheme. These results imply that, to resist one or two stragglers, the best gradient coding scheme will double or even triple the computation load in each worker, which leads to a large transmission and processing overhead for data-intensive applications.

In practical distributed learning applications, we only need to *approximately* reconstruct the gradients. For example, the gradient descent algorithm is internally robust to the noise of gradient evaluation, and the algorithm still converges when the error of each step is bounded [1]. In other scenarios, adding the noise to the gradient evaluation may even improve the generalization performance of the trained model [13]. These facts motivate the idea to develop an *approximate gradient coding* methodology. More specifically, suppose that s of the n workers are stragglers, the approximate gradient coding allows the master node to reconstruct the full gradient with a multiplicative error ϵ from $n - s$ received results. The computation load in this case is a function $d(s, \epsilon)$ of both number of stragglers s and error ϵ . By introducing the error term, one may expect to further reduce the computation load. Given this formulation, we are interested in the following key questions:

What is the minimum computation load for the approximate gradient coding problem? Can we find an optimal scheme that achieves this lower bound?

There have been two computing schemes proposed earlier for this problem. The first one, introduced in [14], utilizes the expander graph, particularly Ramanujan graphs to provide an approximate construction that achieves a computation load of $O(ns/(n-s)\epsilon)$ given error ϵ . However, expander graphs, especially Ramanujan graphs, are expensive to compute in practice, especially for a large number of workers. Hence, an alternate computing scheme was recently proposed in [2], referred to as *Bernoulli Gradient Code* (BGC). This coding scheme incurs a computation load of $O(\log(n))$ and an error of $O(n/(n-s)\log(n))$ with high probability.

1.1 Main Contribution

In this paper, we provide the **information-theoretic lower bound** of minimum computation load given number of data partitions, straggler tolerance and error of gradient, which implies that the optimum computation load can be far less than what the above two schemes achieve. More specifically, we first show that, if we need to exactly ($\epsilon = 0$) recover the full gradients with high probability, the minimum computation load satisfies

$$d(s, 0) \geq O\left(\frac{\log(n)}{\log(n/s)}\right). \quad (1)$$

We also design a coding scheme, referred to as *d-fractional repetition code* (FRC) that achieves the optimum computation load. This result implies that, if we allow the decoding process to fail with a vanishing probability, the computation load in each worker can be significantly reduced: from $s + 1$ to $O(\log(n)/\log(n/s))$. For example, when $n = 100$ and $s = 10$, each worker in the original gradient coding strategy needs to store $11\times$ data partitions, while in the approximate scheme, we only need to store $2\times$ data partitions.

Table 1. Comparison of Existing Schemes

Scheme	Computation load	Error of gradient
cyclic MDS [17]	$s + 1$	0
expander graph code [14]	$O\left(\frac{ns}{(n-s)\epsilon}\right)$	ϵ
BGC [2] ¹	$O(\log(n))$	$O\left(\frac{n}{(n-s)\log(n)}\right)$
FRC ¹	$O\left(\frac{\log(n)}{\log(n/s)}\right)$	0
BRC ¹	$O\left(\frac{\log(1/\epsilon)}{\log(n/s)}\right)$	ϵ

¹ result holds with high probability, i.e., $1-o(1)$.

Furthermore, we identify the following three-fold fundamental tradeoff among the computation load $d(s, \epsilon)$, recovery error ϵ , and the number of stragglers s in order to approximately recover the full gradients with high probability. The tradeoff reads

$$d(s, \epsilon) \geq O\left(\frac{\log(1/\epsilon)}{\log(n/s)}\right).$$

This result provides a quantitative characterization that the noise of the gradient plays a logarithmic reduction role, i.e., from $O(\log(n))$ to the $O(\log(n)) - O(\log(\epsilon n))$ in the desired computation load. For example, when the error of gradient is $O(1/\log(n))$, the existing BGC scheme in [2] provides a computation load of $O(\log(n))$, while the information-theoretical lower bound is $O(\log(\log(n)))$. We further give an explicit code construction, referred to as *batch raptor code* (BRC), based on random edge removal process that achieves this fundamental tradeoff. We compare our proposed scheme with existing gradient coding schemes in TABLE 1.

We finally implement and benchmark the proposed gradient coding schemes at a super computer center ¹ and empirically demonstrate its performance gain compared with existing strategies.

1.2 Related Literature

The work of Lee et al. [8] initiated the study of using coding technique such as MDS code for mitigating stragglers in the distributed linear transformation problem and the regression problem. Subsequently, one line of studies was centered on designing the coding scheme in distributed linear transformation problem. Dutta et al. [5] constructed a deterministic coding scheme in the product of a matrix and a long vector. Lee et al. [9] designed a type of efficient 2-dimensional MDS code for the high dimensional matrix multiplication problem. Yu et al. [22] proposed the optimal coding scheme, named as polynomial code, in the matrix multiplication problem. Wang et al. [18, 19] further initialized the study of computation load in the distributed transformation problem and design several efficient coding schemes with low density generator matrix.

The second line of researches focus on constructing the coding schemes in the distributed algorithm in machine learning application. The work of [6] first addressed the straggler mitigation in linear regression problem by data encoding. Our results are closely related to designing the code for general distributed gradient descent or the problem of computing sum of functions. The initial study by [17] presented an optimal trade-off between the computation load and straggler tolerance for any loss functions. Two subsequent works in [2, 14] considered the approximate gradient evaluation and proposed the BGC scheme with less computation load compared to the scheme in [17]. Maity et al. [12] applied the existing LDPC to a linear regression model with sparse recovery. Ye et al. [21] further introduced the communication complexity in such problem and constructed an efficient code for reducing both straggler effect and communication overhead. None

¹Omitted due to the double-blind policy

Fig. 1. Gradient coding framework.

of the aforementioned works characterizes the fundamental limits of the approximate gradient coding problem. In the sequel, we will systematically investigate this problem.

2 PRELIMINARIES

2.1 Problem Formulation

The data set is denoted by $D = \{(x_i, y_i)\}_{i=1}^N$ with input feature $x_i \in \mathbb{R}^p$ and label $y_i \in \mathbb{R}$. Most machine learning tasks aim to solve the following optimization problem:

$$\beta^* = \arg \min_{\beta \in \mathbb{R}^p} \sum_{i=1}^N L(x_i, y_i; \beta) + \lambda R(\beta), \quad (2)$$

where $L(\cdot)$ is a task-specific loss function, and $R(\cdot)$ is a regularization function. This problem is usually solved by gradient-based approaches. More specifically, the parameters β are updated according to the iteration $\beta^{(t+1)} = h_R(\beta^{(t)}, g^{(t)})$, where $h_R(\cdot)$ is the proximal mapping of gradient-based iteration, and $g^{(t)}$ is the gradient of the loss function at the current parameter $\beta^{(t)}$, defined as

$$g^{(t)} = \sum_{i=1}^N \nabla L(x_i, y_i; \beta^{(t)}). \quad (3)$$

In practice, the number of data samples N is quite large, i.e., $N \geq 10^9$, the evaluation of the gradient $g^{(t)}$ will become a bottleneck of the above optimization process and should be distributed over multiple workers. Suppose that there are n workers $W_1, W_2, \dots, W_n$, and the original dataset is partitioned into n subsets of equal size $\{D_1, D_2, \dots, D_n\}$. In the traditional distributed gradient descent, each worker i stores the dataset D_i . During iteration t , the master node first broadcasts the current classifier $\beta^{(t)}$ to each worker. Then each worker i computes a partial gradient $g_i^{(t)}$ over data block D_i , and returns it to the master node. The master node collects all the partial gradients to obtain a gradient evaluation $g^{(t)} = \sum_{i=1}^n g_i^{(t)}$ and updates the classifier correspondingly. In the gradient coding framework, as illustrated in Figure. 1, each worker i stores multiple data blocks and computes a linear combination of partial gradients, then the master node receives a subset of results and decodes the full gradient $g^{(t)}$.

More formally, the gradient coding framework can be represented by a coding matrix $A \in \mathbb{R}^{n \times n}$, where the i th worker computes² $\tilde{g}_i = \sum_{j=1}^n A_{ij} g_j$. Let $g \in \mathbb{R}^{n \times p}$ ($\tilde{g} \in \mathbb{R}^{n \times p}$) be a matrix with each row being the (coded) partial gradient

$$\tilde{g} = [\tilde{g}_1; \tilde{g}_2; \dots; \tilde{g}_n] \text{ and } g = [g_1; g_2; \dots; g_n].$$

Therefore, we can represent $\tilde{g} = Ag$. Suppose that there exist s stragglers and the master node receives $n - s$ results indexed by set S . Then the received partial gradients can be represented by $\tilde{g}_S = A_S g$, where $A_S \in \mathbb{R}^{(n-s) \times n}$ is a row submatrix of A containing rows indexed by S . During the decoding process, the master node solves the following problem,

$$u^* = \arg \min_{u \in \mathbb{R}^{n-s}} \|A_S^T u - 1_n\|^2, \quad (4)$$

and recovers the full gradient by $u^* \tilde{g}$, where $1_n \in \mathbb{R}^n$ denotes the all one vector.

DEFINITION 1. (Recovery error) *Given a submatrix $A_S \in \mathbb{R}^{(n-s) \times n}$, the corresponding recovery error is defined as*

$$\text{err}(A_S) = \min_{u \in \mathbb{R}^{n-s}} \|A_S^T u - 1_n\|^2, \quad (5)$$

Instead of directly measuring the error of recovered gradient, i.e., $\min_u \|u A_S g - 1_n g\|$, this metric quantifies how close 1_n is to being in the span of the columns of A_S . It is also worth noting that the overall recovery error is small relative to the magnitude of the gradient, since the minimum decoding error satisfies $\min_u \|u A_S g - 1_n g\| \leq \|g\| \cdot \min_u \|u A_S - 1_n\|$.

DEFINITION 2. (Computation load) *The computation load of a gradient coding scheme A is defined as*

$$\kappa(A) = \max_{1 \leq i \leq n} \|A_i\|_0, \quad (6)$$

where $\|A_i\|_0$ is the number of nonzero coefficients of the i th row A_i .

The existing work [17] shows that the minimum computation load is at least $s + 1$ when we require decoding the full gradient exactly, i.e., $\text{err}(A_S) = 0$, among all $S \subseteq [n]$, $|S| = n - s$. The *approximate gradient coding* relaxes the worst-case scenario to a more realistic setting, the “average and approximate” scenario. Formally, we have the following systematic definition of the approximate gradient codes.

DEFINITION 3. (ϵ -approximate gradient code) *Given number of s stragglers in n workers, the set of ϵ -approximate gradient code is defined as*

$$\mathcal{G}_\epsilon = \{A \in \mathbb{R}^{n \times n} | \mathbb{P}[\text{err}(A_S) > \epsilon n] = o(1)\}, \quad (7)$$

where $A_S \in \mathbb{R}^{(n-s) \times n}$ is a randomly chosen row submatrix of A .

The above definition of gradient code is general and includes most existing works on approximate gradient coding. For example, let $\delta = s/n$, the existing scheme based on Ramanujan graphs is a ϵ -approximate gradient code that achieves computation load of $O(s/(1 - \delta)\epsilon)$; the existing BGC [2] can be regarded as a $O(1/(1 - \delta) \log(n))$ -approximate gradient code that achieves computation load of $O(\log(n)/(1 - \delta))$.

²Here we omit iteration count t .

2.2 Main Results

Our first main result provides the minimum computation load and corresponding optimal code when we want to exactly decode the full gradient with high probability.

THEOREM 1. *Suppose that out of n workers, $s = \delta n$ are stragglers. The minimum computation load of any gradient codes in $\mathcal{G}_0$ satisfies*

$$\kappa_0^*(A) \triangleq \min_{A \in \mathcal{G}_0} \kappa(A) \geq O\left(\max\left\{1, \frac{\log(n)}{\log(1/\delta)}\right\}\right). \quad (8)$$

And we construct a gradient code, we call d -fractional repetition code $A^{FRC} \in \mathcal{G}_0$, such that

$$\lim_{n \rightarrow \infty} \kappa(A^{FRC})/\kappa_0^*(A) = 1. \quad (9)$$

The following main result provides a more general outer bound when we allow the recovered gradient to contain some error.

THEOREM 2. *Suppose that out of n workers, $s = \delta n$ are stragglers. If $0 < \epsilon < O(1/\log^2(n))$, the minimum computation load of any gradient codes in $\mathcal{G}_\epsilon$ satisfies*

$$\kappa_\epsilon^*(A) \triangleq \min_{A \in \mathcal{G}_\epsilon} \kappa(A) \geq O\left(\max\left\{1, \frac{\log(1/\epsilon)}{\log(1/\delta)}\right\}\right).$$

And we construct a gradient code, named as batch raptor code $A^{BRC} \in \mathcal{G}_\epsilon$, such that

$$\lim_{n \rightarrow \infty} \kappa(A^{BRC})/\kappa_\epsilon^*(A) = 1. \quad (10)$$

Theorem 2 provides a fundamental tradeoff among the gradient noise, the straggler tolerance and the computation load. And the gradient noise ϵ provides a factor of logarithmic reduction, i.e., $-\log(\epsilon n)$ of the computation load.

Notation: Suppose that $A_S \in \mathbb{R}^{(n-s) \times n}$ is a row submatrix of A containing $(n-s)$ randomly and uniformly chosen rows. A_i (or $A_{S,i}$) denotes i th column of matrix A (or A_S) and a_i (or $a_{S,i}$) denotes i th row of matrix A (or A_S). The $\text{supp}(x)$ is defined as the support set of vector x . $\|x\|_0$ represents the number of nonzero elements in vector x .

3 0-APPROXIMATE GRADIENT CODE

In this section, we consider a simplified scenario that the error of gradient evaluation is zero. It can be regarded as a probabilistic relaxation of the worst-case scenario in [17]. We first characterize the fundamental limits of the any gradient codes in the set.

$$\mathcal{G}_0 = \{A \in \mathbb{R}^{n \times n} | \mathbb{P}[\text{err}(A_S) > 0] = o(1)\}.$$

Then we design a gradient code to achieve the lower bound.

3.1 Minimum Computation Load

The minimum computation load can be determined by exhaustively searching over all possible coding matrices $A \in \mathcal{G}_0$. However, there exist $\Omega(2^{n^2})$ possible candidates in $\mathcal{G}_0$ and such a procedure is practically intractable. To overcome this challenge, we construct a new theoretical path: (i) we first analyze the structure of the optimal gradient codes, and establish a lower bound of the minimum failure probability $\mathbb{P}(\min_{u \in \mathbb{R}^{n-s}} \|A_S^T u - 1_n\|^2 > 0)$ given computation load d ; (ii) we derive an exact estimation of such lower bound, which is a monotonically non-increasing function of d ; and (iii) we show that this lower bound is non-vanishing when the computation d is less than a specific quantity, which provides the desired lower bound.

The following lemma shows that the minimum probability of decoding failure is lower bounded by the minimum probability that there exists an all-zero column of matrix A_S in a specific set of matrices.

LEMMA 1. Suppose that the computation load $\kappa(A) = d$ and define the set of matrices $\mathcal{A}_n^d = \{A \in \mathbb{R}^{n \times n} | \kappa(A) = d\}$, we have

$$\min_{A \in \mathcal{A}_n^d} \mathbb{P}(\text{err}(A_S) > 0) \geq \min_{A \in \mathcal{U}_n^d} \mathbb{P}\left(\bigcup_{i=1}^n \|A_{S,i}\|_0 = 0\right), \quad (11)$$

where set of matrices $\mathcal{U}_n^d \triangleq \{A \in \mathbb{R}^{n \times n} | \|a_i\|_0 = \|A_i\|_0 = d, \forall i \in [n]\}$.

Based on the inclusion-exclusion principle, we observe that the above lower bound is dependent on the set system $\{\text{supp}(A_i)\}_{i=1}^n$ formed by matrix A . Therefore, one can directly transform the above minimization problem into an integer program. However, due to the non-convexity of the objective function, it is difficult to obtain a closed form expression. To reduce the complexity of our analysis, we have the following lemma to characterize a common structure among all matrices in set $\mathcal{U}_n^d$.

LEMMA 2. For any matrix $A \in \mathcal{U}_n^d$, there exists a set $\mathcal{I}_d \subseteq [n]$ such that $|\mathcal{I}_d| \geq \lfloor n/d^2 \rfloor$ and

$$\text{supp}(A_i) \cap \text{supp}(A_j) = \emptyset, \forall i \neq j, i, j \in \mathcal{I}_d. \quad (12)$$

Based on the results of Lemma 1 and Lemma 2, we can get an estimation of the lower bound (11), and obtain the following theorem.

THEOREM 3. Suppose that out of n workers, $s = \delta n$ are stragglers. If $s = \Omega(1)$, the minimum computation load satisfies

$$d^*(s, 0) \triangleq \min_{A \in \mathcal{G}_0} \kappa(A) \geq \frac{\log(n \log^2(1/\delta)/\log^2(n))}{\log(1/\delta)}; \quad (13)$$

otherwise, $d^*(s, 0) = 1$.

Based on Theorem 5, we can observe the power of probabilistic relaxation in reducing the computation load. For example, if the number of stragglers s is proportional to the number of workers n , the minimum computation load $d^*(s, 0) = O(\log(n))$, while the worst-case lower bound is $\Theta(n)$; if $s = \theta(n^\lambda)$, where $0 < \lambda < 1$ is a constant, the minimum computation load $d^*(s, 0) = 1/(1-\lambda)$ is a constant, while the worst-case one is $\Theta(n^\lambda)$. The Figure 2 provides a quantitative comparison of the proposed ϵ -approximate gradient coding and existing ones.

3.2 d -Fractional Repetition Code

In this subsection, we provide a construction of coding matrix A that asymptotically achieves the minimum computation load. The main idea is based on a generalization of the existing fractional repetition code [17].

DEFINITION 4. (d -Fractional Repetition Code) Divide n workers into d groups of size n/d . In each group, divide all data equally and disjointly, and assign d partitions to each worker. All the groups are replicas of each other. The coding matrix A^{FRC} is defined as

$$A^{\text{FRC}} = \begin{bmatrix} A_b \\ A_b \\ \vdots \\ A_b \end{bmatrix}, A_b = \begin{bmatrix} 1_{1 \times d} & 0_{1 \times d} & \cdots & 0_{1 \times d} \\ 0_{1 \times d} & 1_{1 \times d} & \cdots & 0_{1 \times d} \\ \vdots & \vdots & \ddots & \vdots \\ 0_{1 \times d} & 0_{1 \times d} & \cdots & 1_{1 \times d} \end{bmatrix}_{\frac{n}{d} \times n}.$$

Fig. 2. Information-theoretical lower bound of existing worst-case gradient coding [17] and proposed ϵ -approximate gradient coding when $n = 1000$.

REMARK 1. Compared to the original construction in [17], we do not need the assumption that n is a multiple d . In this case, we can construct the FRC as following: let the size of each group equal to $\lfloor n/d \rfloor$. Then randomly choose $\text{mod}(n, d)$ groups and increase the size of each by one. Besides, we the parameter d is adjustable in our construction, while the original one requires $d = s + 1$.

The decoding algorithm for the above FRC is straightforward: instead of solving the problem (4), the master node sums the partial gradients of any n/d workers that contain disjoint data partitions. The following technical lemma proposed in [10] is useful in our theoretical analysis.

LEMMA 3. (Approximate inclusion-exclusion principle) Let n be integers and $k \geq \Omega(\sqrt{n})$, and let $E_1, E_2, \dots, E_n$ be collections of sets, then we have

$$\mathbb{P}\left(\bigcup_{i \in [n]} A_i\right) = \left(1 + e^{-\frac{2k}{\sqrt{n}}}\right) \sum_{\substack{|I| \leq k \\ I \subseteq [n]}} (-1)^{|I|} \cdot \mathbb{P}\left(\bigcap_{i \in I} A_i\right).$$

The above lemma shows that one can approximately estimate the probability of event $E_1 \cap \dots \cap E_n$ given the probability of events $\bigcap_{i \in I} E_i$, $|I| < \Omega(k^{0.5})$.

THEOREM 4. Suppose that there exist $s = \delta n$ stragglers in n workers. If d satisfies

$$d = \max \left\{ 1, \frac{\log(n \log(1/\delta))}{\log(1/\delta)} \right\}, \quad (14)$$

then we have $\mathbb{P}(\text{err}(A_S^{\text{FRC}}) > 0) = o(1)$.

Combining the results of Theorem 3 and Theorem 4, we can obtain the main argument of Theorem 1. In practical implementation of FRC, once the decoding process fails in k th iteration, a straightforward method is to restart k th iteration. Due to the fact that the decoding failure is less happen during the iteration, such overhead will be amortized. As can be seen in the experimental section, during 100 iterations, only one or two iterations are decoding failure.

4 ϵ -APPROXIMATE GRADIENT CODE

In this section, we consider a more general scenario that the error of gradient evaluation ϵ is larger than zero. We first provide a fundamental three-fold trade-off among the computation load, error

of gradient and the number of stragglers of any codes in the set.

$$\mathcal{G}_\epsilon = \{A \in \mathbb{R}^{n \times n} | \mathbb{P}[\text{err}(A_S) > \epsilon n] = o(1)\}. \quad (15)$$

Then we construct a random code that achieves this lower bound.

4.1 Fundamental Three-fold Tradeoff

Based on the proposed theoretical path in Section 3.1, we can lower bound the probability that the decoding error is larger than ϵn by the one that there exist larger than ϵn all-zero columns of matrix A_S . However, such a lower bound does not admit a close-form expression since the probability event is complicated and contains exponential many partitions. To overcome this challenge, we decompose the above probability event into n dependent events, and analyze its the second-order moment. Then, we use Bienayme-Chebyshev inequality to further lower bound the above probability. The following theorem provides the lower bound of computation load among the feasible gradient codes $\mathcal{G}_\epsilon$.

THEOREM 5. *Suppose that out of n workers, $s = \delta n$ are stragglers, and $\epsilon < O(1/\log^2(n))$, then the minimum computation load satisfies*

$$\kappa^*(A) \triangleq \min_{A \in \mathcal{G}_\epsilon} \kappa(A) \geq \frac{\log(n \log^2(1/\delta)/(2\epsilon n + 4) \log^2(n))}{\log(1/\delta)}.$$

Note that the above result also holds for $\epsilon = 0$, which is slightly lower than the bound in Theorem 3. Based on the result of Theorem 5, we can see that the gradient noise ϵ provides a logarithmic reduction of the computation load. For example, when $n = 1000$ and $s = 100$, the 0-approximate gradient coding requires each worker storing $3\times$ data partitions, while even a 0.01-approximate gradient coding only requires $2\times$ data partitions. Detailed comparison can be seen in Figure 2.

4.2 Random Code Design

Now we present the construction of our random code, we name *batch rapter code* (BRC), which achieves the above lower bound with high probability. The construction of the BRC consists of two layers. In the first layer, the original data set $\{D_i\}_{i=1}^n$ are partitioned into n/b batches $\{B_i\}_{i=1}^{n/b}$ with the size of each batch equal to b . The data in each batch is selected by $B_i = \{D_j\}_{j=1+(i-1)b}^{ib}$, and therefore the intermediate coded partial gradients can be represented by $g_i^b = \sum_{j=1+(i-1)b}^{ib} g_j$, $\forall i \in [n/b]$. In the second step, we construct a type of rapter code taking the coded partial gradients g^b as input block.

DEFINITION 5. ($((b, P)$ -batch rapter code) *Given the degree distribution $P \in \mathbb{R}^{n/b}$ and batches $\{B_i\}_{i=1}^{n/b}$, we define the (b, P) -batch rapter code as: each worker $k \in [n]$, stores the data $\{B_i\}_{i \in I}$ and computes*

$$\tilde{g}_k = \sum_{i \in I} g_i^b = \sum_{i \in I} \sum_{j=1+(i-1)b}^{ib} g_j, \quad (16)$$

where I is a randomly and uniformly subset of $[n/b]$ with $|I| = d$, and d is generated according to distribution P . The coding matrix A^{BRC} is therefore given by

$$A^{BRC} \triangleq \begin{matrix} \text{random } d \text{ nonzero blocks} \\ \begin{bmatrix} 1_{1 \times b} & 0_{1 \times b} & 1_{1 \times b} & \cdots & 0_{1 \times b} \\ 0_{1 \times b} & 0_{1 \times b} & 1_{1 \times b} & \cdots & 0_{1 \times b} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0_{1 \times b} & 1_{1 \times b} & 0_{1 \times b} & \cdots & 1_{1 \times b} \end{bmatrix} \end{matrix}.$$

Note that when n is not a multiple d , we can tackle it using a method similar to the one used in FRC. The decoding algorithm for the BRC goes through a peeling decoding process: it first finds a ripple worker (with only one batch) to recover one batch g_i^b and add it to the sum of gradient g . Then for each collected results, it subtracts this batch if the computed gradients contains this batch. The whole procedure is listed in Algorithm 1.

Example 1. (Batch raptor code, $n = 6, s = 2$) Consider a distributed gradient descent problem with 6 workers. In the batch raptor code, the data is first partitioned into 4 batches with $B_1 = \{D_1\}$, $B_2 = \{D_2\}$, $B_3 = \{D_3, D_4\}$, $B_4 = \{D_5, D_6\}$. After random construction, 6 workers are assigned the tasks: $\tilde{g}_1 = g_1 + g_2$, $\tilde{g}_2 = g_1$, $\tilde{g}_3 = g_2 + (g_5 + g_6)$, $\tilde{g}_4 = (g_3 + g_4) + (g_5 + g_6)$, $\tilde{g}_5 = g_5 + g_6$, $\tilde{g}_6 = g_2 + (g_5 + g_6)$. Suppose that both the 5th and 6th workers are stragglers and the master node collects partial results from workers $\{1, 2, 3, 4\}$. Then we can use the peeling decoding algorithm: first, find a ripple node $\tilde{g}_2$. Then we can use $\tilde{g}_2$ to recover g_2 by $\tilde{g}_1 - \tilde{g}_2$. Further, we can use g_2 to get a new ripple $g_5 + g_6$ by $\tilde{g}_3 - g_2$, and use ripple $g_5 + g_6$ to recover $g_3 + g_4$ by $\tilde{g}_4 - (g_5 + g_6)$. In another case, change the coding scheme of 3th, 4th and 6th worker to $\tilde{g}_3 = g_2$, $\tilde{g}_4 = g_3 + g_4$ and $\tilde{g}_6 = g_1 + g_2$. Suppose that both the 4th and 6th workers are stragglers. We can use a similar decoding algorithm to recover $g_1 + g_2 + g_5 + g_6$ without g_3, g_4 . However, the computation load is decreased from 4 to 2. Actually, the

Algorithm 1 Batch raptor code (master node's protocol)

repeat

The master node assign the data sets according to Definition 5.

until the master node collects results from first finished $n - s$ workers.

repeat

Find a row M_i in received coding matrix M with $\|M_i\|_0 = 1$.

Suppose that the column index of the nonzero element in matrix M_i is k_0 and let $g = g + \tilde{g}_{k_0}$.

for each computation results $\tilde{g}_k$ **do**

if M_{kk_0} is nonzero **then**

$\tilde{g}_k = \tilde{g}_k - M_{kk_0} \tilde{g}_{k_0}$ and set $M_{kk_0} = 0$.

end if

end for

until $n(1 - \epsilon)$ partial gradients is recovered.

above peeling decoding algorithm can be viewed as an edge-removal process in a bipartite graph. We construct a bipartite graph with one partition being the original batch gradients $\{g_i^b\}$ and the other partition being the coded gradients $\{\tilde{g}_i\}$. Two nodes are connected if such a computation task contains that block. As shown in the Figure 3, in each iteration, we find a ripple (degree one node) on the right and remove the adjacent edges of that left node, which might produce some new ripples in the right. Then we iterate this process until we decode all gradients.

Based on the above graphical illustration, the key point of being able to successfully decode for BRC is the existence of the ripple during the edge removal process, which is mainly dependent on the degree distribution P and batch size b . The following theorem shows that, under a specific choice of P and b , we can guarantee the success of decoding process with high probability.

Fig. 3. Example of the batch raptor code and peeling decoding algorithm.

THEOREM 6. Define the degree distribution P_w

$$p_k = \begin{cases} \frac{u}{u+1}, & k = 1 \\ \frac{1}{k(k-1)(u+1)}, & 2 \leq k \leq D \\ \frac{1}{D(u+1)}, & k = D+1 \end{cases}, \quad (17)$$

where $D = \lfloor 1/\epsilon \rfloor$, $u = 2\epsilon(1-2\epsilon)/(1-4\epsilon)^2$ and $b = \lceil 1/\log(1/\delta) \rceil + 1$. Then the $(\lceil 1/\log(1/\delta) \rceil, P_w)$ -batch raptor code with decoding Algorithm 1 satisfies

$$\mathbb{P}(\text{err}(A_S^{BRC}) > c) < e^{-c_0 n},$$

and achieves an average computation load of

$$O\left(\frac{\log(1/\epsilon)}{\log(1/\delta)}\right). \quad (18)$$

The above result is based on applying a martingale argument to the peeling decoding process [11]. In practical implementation, the degree distribution can be further optimized given the n , s and error ϵ [16].

5 SIMULATION RESULTS

In this section, we present the experimental results at a super computing center. We compare our proposed schemes including d -fractional repetition code (FRC) and batch raptor code (BRC) against existing gradient coding schemes: (i) forget- s scheme (stochastic gradient descent): the master node only waits the results of non-straggling workers; (ii) cyclic MDS code [17]: gradient coding scheme that can guarantee the decodability for any s stragglers; (iii) bernoulli gradient code (BGC) [2]: approximate gradient coding scheme that only requires $O(\log(n))$ data copies in each worker. (iv) Asynchronous gradient descent (AsySG) [15]: the master do not need to wait the results of straggling workers and update the parameter every fix number of iterations. To simulate straggler effects in a large-scale system, we randomly pick s workers that are running a background thread.

5.1 Experiment Setup

We implement all methods in python using MPI4py. Each worker stores the data according to the coding matrix A . During the iteration of the distributed gradient descent, the master node broadcasts the current classifier $\beta^{(t)}$ using `Isend()`; then each worker computes the coded partial gradient $\tilde{g}_i^{(t)}$ and returns the results using `Isend()`. Then the master node actively listens to the response from each worker via `Irecv()`, and uses `Waitany()` to keep polling for the earliest finished tasks.

Fig. 4. The generalization AUC versus running time of applying distributed gradient descent in a logistic regression model. The two proposed schemes FRC and BRC are compared against three existing schemes. The learning rate α is fixed for all the experiment.

Upon receiving enough results, the master stops listening and starts decoding the full gradient $g^{(t)}$ and updates the classifier to $\beta^{(t+1)}$.

In our experiment, we ran various schemes to train logistic regression models, a well-understood convex optimization problem that is widely used in practice. We choose the training data from LIBSVM dataset repository. We use $N = 19264097$ samples and a model dimension of $p = 1163024$. We evenly divide the data into n partitions $\{D_k\}_{k=1}^n$. The key step of gradient descent algorithm is

$$\beta^{(t+1)} = \beta^{(t)} + \alpha \sum_{k=1}^n \overbrace{\sum_{i \in D_k} \eta(y_i - h_{\beta^{(t)}}(x_i))x_i}^{g_i^{(t)}},$$

where $h_{\beta^{(t)}}(\cdot)$ is the logistic function, α is the predetermined step size.

5.2 Generalization Error

We first compare the generalization AUC of the above five schemes when number of workers $n = 30$ or 60 and 10% or 20% workers are stragglers. In Figure 4, we plot the generalization AUC versus the running time of all the schemes under different n and s . We can observe that our proposed schemes (FRC and BRC) achieve significantly better generalization error compared to existing ones. Both the forget- s scheme and AsySG scheme converges slowly, since it does not utilize the full gradient and only admits a small step size α compared to other schemes. In particular, when the number of workers increases, our proposed schemes provide even larger speed up over the state of the art.

Fig. 5. The final job completion time of achieving a generalization $AUC = 0.8$ in a logistic regression model. The two proposed schemes FRC and BRC are compared against three existing schemes. The learning rate α is fixed for all the experiments.

5.3 Impact of Straggler Tolerance

We further investigate the impact of straggler tolerance s . We fix the number of workers $n = 30$ or $n = 60$ and increase the fraction of stragglers from 10% to 30%. In Figure 5, we plot the job completion time to achieve a fixed generalization $AUC = 0.8$. The first observation is that our propose schemes reduce the completion time by 50% compared to existing ones. The cyclic MDS code and forget- s (stochastic gradient descent) scheme are sensitive to the number of stragglers. The main reasons are: (i) the computation load of cyclic MDS code is linear in s ; (ii) the available step size of forget- s is reduced when the number of received partial gradients decreases.

The job completion time of the proposed FRC and BRC are not sensitive to the straggler tolerance s , especially when the number of workers n is large. For example, the job completion time of BRC only increases 10% when fraction of straggler increases from 10% to 30%. Besides, we observe that, when the straggler tolerance is small, i.e., $s/n < 0.1$, the job completion time of FRC is slightly lower than that of BRC, because the computation loads are similar for both FRC and BRC in this case, and the FRC utilizes the information of the full gradient.

6 CONCLUSION

In this paper, we formalized the problem of approximate gradient coding and systematically characterized the fundamental three-fold tradeoff among the computation load, noise of gradient ϵ and straggler tolerance s . We further constructed two schemes, FRC and BRC, to achieve this lower bound. Theoretically, the proposed fundamental tradeoff uncovers how the “probabilistic relaxation” and the “noise of gradient” quantitatively influence the computation load: (i) the probabilistic relaxation provides a linear reduction in computation load, i.e., a reduction from $\Theta(n)$ to $O(\log(n))$ redundant computations of each worker; (ii) the gradient noise ϵ introduces another logarithmic reduction in computation load, i.e., a reduction from $O(\log(n))$ to $O(\log(n) - \log(\epsilon n))$ redundant computations. In practice, we have experimented with various gradient coding schemes on super computing center. Our proposed schemes provide 50% speed up over the state of the art.

7 ACKNOWLEDGEMENT

This work has been supported in part by NSF grants CNS-1901057 and CNS- 1717060, ONR grant N00014-17-1-2417), and by the IITP grant (MSIT), (2017-0-00692, Transport-aware Streaming Technique Enabling Ultra Low-Latency AR/VR Services.

REFERENCES

- [1] Léon Bottou. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*. Springer, 177–186.
- [2] Zachary Charles, Dimitris Papailiopoulos, and Jordan Ellenberg. 2017. Approximate gradient coding via sparse random graphs. *arXiv preprint arXiv:1711.06771* (2017).
- [3] Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Andrew Senior, Paul Tucker, Ke Yang, Quoc V Le, et al. 2012. Large scale distributed deep networks. In *Advances in neural information processing systems*. 1223–1231.
- [4] Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: simplified data processing on large clusters. *Commun. ACM* 51, 1 (2008), 107–113.
- [5] Sanghamitra Dutta, Viveck Cadambe, and Pulkit Grover. 2016. Short-dot: Computing large linear transforms distribut- edly using coded short dot products. In *Advances In Neural Information Processing Systems*. 2100–2108.
- [6] Can Karakus, Yifan Sun, Suhas Diggavi, and Wotao Yin. 2017. Straggler mitigation in distributed optimization through data encoding. In *Advances in Neural Information Processing Systems*. 5434–5442.
- [7] Quoc V Le, Jiquan Ngiam, Adam Coates, Abhik Lahiri, Bobby Prochnow, and Andrew Y Ng. 2011. On optimization methods for deep learning. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*. Omnipress, 265–272.
- [8] Kangwook Lee, Maximilian Lam, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran. 2017. Speeding up distributed machine learning using codes. *IEEE Transactions on Information Theory* (2017).
- [9] Kangwook Lee, Changho Suh, and Kannan Ramchandran. 2017. High-dimensional coded matrix multiplication. In *Information Theory (ISIT), 2017 IEEE International Symposium on*. IEEE, 2418–2422.
- [10] Nathan Linial and Noam Nisan. 1990. Approximate inclusion-exclusion. *Combinatorica* 10, 4 (1990), 349–365.
- [11] Michael G Luby, Michael Mitzenmacher, Mohammad Amin Shokrollahi, and Daniel A Spielman. 2001. Efficient erasure correcting codes. *IEEE Transactions on Information Theory* 47, 2 (2001), 569–584.
- [12] Raj Kumar Maity, Ankit Singh Rawat, and Arya Mazumdar. 2018. Robust Gradient Descent via Moment Encoding with LDPC Codes. *SysML* (2018).
- [13] Arvind Neelakantan, Luke Vilnis, Quoc V Le, Ilya Sutskever, Lukasz Kaiser, Karol Kurach, and James Martens. 2015. Adding gradient noise improves learning for very deep networks. *arXiv preprint arXiv:1511.06807* (2015).
- [14] Netanel Raviv, Itzhak Tamo, Rashish Tandon, and Alexandros G Dimakis. 2018. Gradient coding from cyclic MDS codes and expander graphs. (2018).
- [15] Sashank J Reddi, Ahmed Hefny, Suvrit Sra, Barnabas Poczos, and Alexander J Smola. 2015. On variance reduction in stochastic gradient descent and its asynchronous variants. In *Advances in Neural Information Processing Systems*. 2647–2655.
- [16] Amin Shokrollahi. 2006. Raptor codes. *IEEE/ACM Transactions on Networking (TON)* 14, SI (2006), 2551–2567.
- [17] Rashish Tandon, Qi Lei, Alexandros G Dimakis, and Nikos Karampatziakis. 2017. Gradient coding: Avoiding stragglers in distributed learning. In *International Conference on Machine Learning*. 3368–3376.
- [18] Sinong Wang, Jiashang Liu, and Ness Shroff. 2018. Coded Sparse Matrix Multiplication. In *International Conference on Machine Learning*.
- [19] Sinong Wang, Jiashang Liu, Ness Shroff, and Pengyu Yang. 2019. Computation Efficient Coded Linear Transform. In *International Conference on Artificial Intelligence and Statistics*.
- [20] Neeraja J Yadwadkar, Bharath Hariharan, Joseph E Gonzalez, and Randy Katz. 2016. Multi-task learning for straggler avoiding predictive job scheduling. *The Journal of Machine Learning Research* 17, 1 (2016), 3692–3728.
- [21] Min Ye and Emmanuel Abbe. 2018. Communication-computation efficient gradient coding. In *International Conference on Machine Learning*.
- [22] Qian Yu, Mohammad Maddah-Ali, and Salman Avestimehr. 2017. Polynomial codes: an optimal design for high-dimensional coded matrix multiplication. In *Advances in Neural Information Processing Systems*. 4406–4416.
- [23] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. 2010. Spark: Cluster computing with working sets. *HotCloud* 10, 10-10 (2010), 95.

A PROOF OF LEMMA 1

PROOF. Since the event that there exists $i \in [n]$ such that $\|A_{S,i}\|_0 = 0$ implies the event that $\text{err}(A_S) = \|A_S^T u - 1_n\|^2 \geq 1$, we can obtain

$$\mathbb{P}(\text{err}(A_S) > 0) \geq \mathbb{P}\left(\bigcup_{i=1}^n \|A_{S,i}\|_0 = 0\right)$$

Suppose that $A^* = \arg \min_{A \in \mathcal{A}_n^d} \mathbb{P}(\text{err}(A_S) > 0)$, A_S^* is the row submatrix of A^* containing $(n - s)$ randomly and uniformly chosen rows, we have

$$\begin{aligned} \min_{A \in \mathcal{A}_n^d} \mathbb{P}(\text{err}(A_S) > 0) &\geq \mathbb{P}\left(\bigcup_{i=1}^n \|A_{S,i}^*\|_0 = 0\right) \\ &\geq \min_{A \in \mathcal{A}_n^d} \mathbb{P}\left(\bigcup_{i=1}^n \|A_{S,i}\|_0 = 0\right) \end{aligned}$$

We next show that

$$\begin{aligned} &\min_{A \in \mathcal{A}_n^d} \mathbb{P}\left(\bigcup_{i=1}^n \|A_{S,i}\|_0 = 0\right) \\ &= \min_{A \in \mathcal{U}_n^d} \mathbb{P}\left(\bigcup_{i=1}^n \|A_{S,i}\|_0 = 0\right) \end{aligned}$$

We will prove that the above probability is monotonically decreasing with the support size of each row and column of matrix A . Assume that there exists $k \in [n]$ such that $\|a_k\|_0 < \kappa(A)$. We change one zero position of row a_k , i.e., a_{kj} to a nonzero constant. Define the new matrix as A' . For simplicity, define the event E_i as $\|A_{S,i}\|_0 = 0$ and E'_i as $\|A'_{S,i}\|_0 = 0$. Then we can write

$$\begin{aligned} \mathbb{P}\left(\bigcup_{i=1}^n E_i\right) &= \mathbb{P}\left(\bigcup_{i \neq j} E_i\right) + \mathbb{P}\left[E_j \setminus \left(\bigcup_{i \neq j} E_i\right)\right] \\ &\stackrel{(a)}{\geq} \mathbb{P}\left(\bigcup_{i \neq j} E'_i\right) + \mathbb{P}\left[E'_j \setminus \left(\bigcup_{i \neq j} E'_i\right)\right] \\ &= \mathbb{P}\left(\bigcup_{i=1}^n E'_i\right). \end{aligned}$$

The above, step (a) is based on the fact that $E_i = E'_i, i \neq j$ and $E'_j \subset E_j$. Similarly, we can prove the monotonicity for support size of each column. Therefore, based on the monotonicity and the Definition 2 of computation load, the lemma follows. $\square$

B PROOF OF LEMMA 2

PROOF. Given any matrix $A \in \mathcal{U}_n^d$. We construct the set $\mathcal{I}_d$ as follows. First, choose a column A_{i_1} and construct set I_1 as follows,

$$I_1 = \{j \in [n] \mid \text{supp}(A_{i_1}) \cap \text{supp}(A_j) \neq \emptyset, j \neq i_1\} \quad (19)$$

Since $A \in \mathcal{U}_n^d$, suppose that $\text{supp}(A_{i_1}) = \{k_1, k_2, \dots, k_d\}$. We can obtain

$$\begin{aligned} |I_1| &= \left| \bigcup_{l=1}^d \{j \in [n] \mid A_{jk_l}, j \neq i_1\} \right| \\ &\stackrel{(a)}{\leq} \sum_{i=1}^d |\{j \in [n] \mid A_{jk_l}, j \neq i_1\}| \stackrel{(b)}{\leq} d^2. \end{aligned} \quad (20)$$

The above, step (a) utilizes the union bound and step (b) is based on the definition of set $\mathcal{U}_n^d$.

Furthermore, we choose a column A_{i_2} such that $i_2 \in [n] \setminus I_1$. Based on the definition of index set I_1 , we have $\text{supp}(A_{i_2}) \cap \text{supp}(A_{i_1}) = \emptyset$. Similarly, we can construct the index set $I_2 = \{j \in [n] \mid \text{supp}(A_{i_2}) \cap \text{supp}(A_j) \neq \emptyset, j \neq i_2\}$ with $|I_2| \leq d^2$. Continue this process k times, we can construct

a set $\mathcal{I}_d = \{i_1, i_2, \dots, i_k\}$ such that for any $i, j \in \mathcal{I}_d$, $\text{supp}(A_i) \cap \text{supp}(A_j) = \emptyset$, and corresponding $I_1, I_2, \dots, I_k$. Since each $|I_k| \leq d^2$, we have $|\mathcal{I}_d| \geq \lfloor n/d^2 \rfloor$. $\square$

C PROOF OF THEOREM 3

PROOF. Suppose that

$$A^* = \arg \min_{A \in \mathcal{U}_n^d} \mathbb{P} \left(\bigcup_{i=1}^n \|A_{S,i}\|_0 = 0 \right). \quad (21)$$

Based on the results of Lemma 2, we can construct a set $\mathcal{I}_d^*$ such that $|\mathcal{I}_d^*| \geq \lfloor n/d^2 \rfloor$ and

$$\text{supp}(A_i^*) \cap \text{supp}(A_j^*) = \emptyset, \forall i \neq j, i, j \in \mathcal{I}_d^*. \quad (22)$$

Combining the results of Lemma 1, we have

$$\begin{aligned} \min_{A \in \mathcal{A}_n^d} \mathbb{P}(\text{err}(A_S) > 0) &\geq \min_{A \in \mathcal{A}_n^d} \mathbb{P} \left(\bigcup_{i=1}^n \|A_{S,i}\|_0 = 0 \right) \\ &= \min_{A \in \mathcal{U}_n^d} \mathbb{P} \left(\bigcup_{i=1}^n \|A_{S,i}\|_0 = 0 \right) \\ &= \mathbb{P} \left(\bigcup_{i=1}^n \|A_{S,i}^*\|_0 = 0 \right) \\ &\geq \mathbb{P} \left(\bigcup_{i \in \mathcal{I}_d^*} \|A_{S,i}^*\|_0 = 0 \right) \end{aligned} \quad (23)$$

The above, last step is based on the fact that $\mathcal{I}_d^* \subseteq [n]$. Suppose that $|\mathcal{I}_d^*| = t$. Based on the inclusion-exclusion principle, we can write

$$\begin{aligned} &\mathbb{P} \left(\bigcup_{i \in \mathcal{I}_d^*} \|A_{S,i}^*\|_0 = 0 \right) \\ &= \sum_{I \subseteq \mathcal{I}_d^*} (-1)^{|I|+1} \mathbb{P} \left(\bigcap_{i \in I} \|A_{S,i}^*\|_0 = 0 \right) \\ &\stackrel{(a)}{=} \sum_{I \subseteq \mathcal{I}_d^*, |I| \leq \lfloor \frac{s}{d} \rfloor} (-1)^{|I|+1} \binom{n - \sum_{i \in I} |\text{supp}(A_{S,i}^*)|}{s - \sum_{i \in I} |\text{supp}(A_{S,i}^*)|} \bigg/ \binom{n}{s} \\ &= \sum_{k=1}^{\min\{t, \lfloor \frac{s}{d} \rfloor\}} \binom{t}{k} (-1)^{k+1} \binom{n - kd}{s - kd} \bigg/ \binom{n}{s} \\ &\stackrel{(b)}{\geq} \sum_{\substack{k \leq \min\{t, \lfloor \frac{s}{d} \rfloor\} \\ k \text{ is odd}}} \binom{t}{k} \left(\frac{n - kd}{n} \right)^{n - kd + 0.5} \left(\frac{s}{s - kd} \right)^{s - kd + 0.5} \left(\frac{s}{n} \right)^{kd} - \\ &\quad \sum_{\substack{k \leq \min\{t, \lfloor \frac{s}{d} \rfloor\} \\ k \text{ is even}}} \binom{t}{k} \frac{144s(n - kd)}{(12s - 1)(12(n - kd) - 1)} \left(\frac{n - kd}{n} \right)^{n - kd + 0.5} \\ &\quad \left(\frac{s}{s - kd} \right)^{s - kd + 0.5}. \end{aligned} \quad (24)$$

The above, step (a) is based on the property of set I_d^* . Step (b) utilizes the following Sterlin's inequalities

$$\sqrt{2\pi n} \left(\frac{n}{e}\right)^n \leq n! \leq \sqrt{2\pi n} \left(\frac{n}{e}\right)^n \frac{12n}{12n-1}. \quad (25)$$

Case 1: The number of stragglers $s = \delta n$ and $\delta = \Theta(1)$.

Since the event $\|A_{S,i_0}^*\|_0 = 0$ belongs to event $\bigcup_{i \in I_d^*} \|A_{S,i}^*\|_0 = 0$ for some $i_0 \in I_d^*$, we have

$$\begin{aligned} \mathbb{P} \left[\bigcup_{i \in I_d^*} \|A_{S,i}^*\|_0 = 0 \right] &\geq \mathbb{P} [\|A_{S,i_0}^*\|_0 = 0] = \binom{n-d}{s-d} / \binom{n}{s} \\ &\stackrel{(a)}{=} (1+o(1)) \left(1 - \frac{d}{n}\right)^{n-d+0.5} \left(1 + \frac{d}{s-d}\right)^{s-d+0.5} \delta^d \\ &= (1+o(1))\delta^d. \end{aligned} \quad (26)$$

The above, step (a) is based on the Sterlin's approximation. This result implies $d > \Omega(1)$ (otherwise, the failure probability is nonvanishing). Then we have $t = \lfloor n/d^2 \rfloor < \lfloor s/d \rfloor$ and $kd = o(n)$ for any $1 \leq k \leq t$, and obtain the following approximation,

$$\left(\frac{n-kd}{n}\right)^{n-kd+0.5} \left(\frac{s}{s-kd}\right)^{s-kd+0.5} = 1 - \alpha(n, k) \text{ and}$$

$$\lim_{n \rightarrow \infty} \alpha(n, k) = 0, \forall 1 \leq k \leq t.$$

$$\frac{144s(n-kd)}{(12s-1)(12(n-kd)-1)} = 1 + \beta(n, k) \text{ and}$$

$$\lim_{n \rightarrow \infty} \beta(n, k) = 0, \forall 1 \leq k \leq t.$$

Utilizing the above approximation and choosing d such that $\delta^d t \rightarrow 1/e$, we have

$$\begin{aligned} &\mathbb{P} \left(\bigcup_{i \in I_d^*} \|A_{S,i}^*\|_0 = 0 \right) \\ &\geq \sum_{\substack{k \leq t \\ k \text{ is odd}}} \binom{t}{k} (1 - \alpha(n, k)) \left(\frac{s}{n}\right)^{kd} - \sum_{\substack{k \leq t \\ k \text{ is even}}} \binom{t}{k} (1 - \alpha(n, k)) \\ &\quad (1 + \beta(n, k)) \left(\frac{s}{n}\right)^{kd} \\ &= 1 - (1 - \delta^d)^t - \sum_{\substack{k \leq t \\ k \text{ is odd}}} \binom{t}{k} \alpha(n, k) \delta^{kd} + \\ &\quad \sum_{\substack{k \leq t \\ k \text{ is even}}} \binom{t}{k} (\alpha(n, k) - \beta(n, k) + \alpha(n, k)\beta(n, k)) \delta^{kd} \\ &\stackrel{(a)}{=} 1 - (1 - \delta^d)^t + o(1) \\ &\stackrel{(b)}{=} 1 - e^{-e^{-1}} + o(1) > 0.307. \end{aligned} \quad (27)$$

The above, step (a) utilizes the fact that $\binom{t}{k} \leq (et/k)^k$, then the quantity $\binom{t}{k} \delta^{kd} \leq (et\delta^d/k)^k = 1/k^k$ and

$$\sum_{k=1}^t \binom{t}{k} (-1)^{k+1} o(1) \delta^{kd} \leq \sum_{k=1}^t \frac{o(1)}{k^k} = o(1). \quad (28)$$

Step (b) is based on the choice of d such that $\delta^d t \rightarrow 1/e$. It is obvious that the probability $\mathbb{P}(\bigcup_{i \in I_d^*} \|A_{S,i}^*\|_0 = 0)$ is monotonically non-increasing with the computation load d . Therefore, the minimum computation load d^* should satisfy $d^* > d_0$, where $\delta^{d_0} n/d_0^2 \rightarrow 1/e$. It is easy to see that

$$d_0 = \frac{\log(ne \log^2(1/\delta)/\log^2(n))}{\log(1/\delta)}. \quad (29)$$

Case 2: The number of stragglers $s = \delta n$, $\delta = o(1)$ and $\delta = \Omega(1/n)$.

In this case, we can choose $d = d_0$. The conditions $\delta = \Omega(1/n)$ implies that $s = \Omega(1)$.

Then, for $k \in [\min\{t, s/d_0\}]$ and $k = o(s/d_0)$, we have following similar estimation.

$$\left(\frac{n - kd_0}{n}\right)^{n - kd_0 + 0.5} \left(\frac{s}{s - kd_0}\right)^{s - kd_0 + 0.5} = 1 - \alpha(n, k) \quad (30)$$

$$\lim_{n \rightarrow \infty} \alpha(n, k) = 0. \quad (31)$$

For $k \in [\min\{t, s/d_0\}]$ and $k = \Theta(s/d_0) = cs/d_0$, we have

$$\begin{aligned} & \binom{t}{k} \left(\frac{n - kd_0}{n}\right)^{n - kd_0 + 0.5} \left(\frac{s}{s - kd_0}\right)^{s - kd_0 + 0.5} \delta^{kd_0} \\ &= \binom{t}{k} \left(\frac{1}{1 - c}\right)^{(1-c)s + 0.5} \delta^{cs} \\ &\leq \left(\frac{1}{1 - c}\right)^{0.5} \left[\left(\frac{e}{c}\right)^{1/d_0} \left(\frac{1}{1 - c}\right)^{1/c - 1} \delta\right]^{kd_0}. \end{aligned} \quad (32)$$

For all $k \in [\min\{t, s/d_0\}]$, we have

$$\frac{144s(n - kd_0)}{(12s - 1)(12(n - kd_0) - 1)} = 1 + \beta(n, k) \quad (33)$$

$$\lim_{n \rightarrow \infty} \beta(n, k) = 0, \forall 1 \leq k \leq t. \quad (34)$$

Therefore, we can obtain the following estimation.

$$\sum_{k=\Theta(s/d_0)}^t \binom{t}{k} \left(\frac{n - kd_0}{n}\right)^{n - kd_0 + 0.5} \left(\frac{s}{s - kd_0}\right)^{s - kd_0 + 0.5} \delta^{kd_0} = o(1). \quad (35)$$

Utilizing the above approximation, we have

$$\begin{aligned} & \mathbb{P}\left(\bigcup_{i \in I_{d_0}^*} \|A_{S,i}^*\|_0 = 0\right) \\ &\geq \sum_{\substack{k=o(s/d_0) \\ k \text{ is odd}}} \binom{t}{k} (1 - \alpha(n, k)) \left(\frac{s}{n}\right)^{kd_0} - \\ &\quad \sum_{\substack{k=o(s/d_0) \\ k \text{ is even}}} \binom{t}{k} (1 - \alpha(n, k)) (1 + \beta(n, k)) \left(\frac{s}{n}\right)^{kd_0} + o(1) \end{aligned}$$

$$\begin{aligned}
&= 1 - (1 - \delta^{d_0})^t - \sum_{\substack{k=o(s/d_0) \\ k \text{ is odd}}} \binom{t}{k} \alpha(n, k) \delta^{kd_0} + \\
&\quad \sum_{\substack{k=o(s/d_0) \\ k \text{ is even}}} \binom{t}{k} (\alpha(n, k) - \beta(n, k) + \alpha(n, k)\beta(n, k)) \delta^{kd_0} + o(1) \\
&= 1 - (1 - \delta^{d_0})^t + o(1) \\
&= 1 - e^{-e^{-1}} + o(1) > 0.307.
\end{aligned} \tag{36}$$

Therefore, the minimum computation load d^* should satisfy $d^* > d_0$. In the case $s = \Theta(1)$, the lower bound 1 is trivial (otherwise, some gradients are lost). Therefore, the theorem follows. $\square$

D PROOF OF THEOREM 4

PROOF. Based on the structure of coding matrix A^{FRC} and decoding algorithm, we can define the following event

$$E_i \triangleq \bigcap_{j=0}^{d-1} \left(\frac{jn}{d} + i \right) \text{th worker is straggler}, 1 \leq i \leq n/d, \tag{37}$$

and we have

$$\mathbb{P}(\text{err}(A_S^{\text{FRC}}) > 0) = \mathbb{P}\left(\bigcup_{i=1}^{n/d} E_i\right). \tag{38}$$

Utilizing the approximate inclusion-exclusion principle and choose $k = n^{0.6}$ and $d = 1 + \log(n)/\log(1/\delta)$, we have

$$\begin{aligned}
\mathbb{P}(\text{err}(A_S^{\text{FRC}}) > 0) &= (1 + e^{-2n^{0.1}}) \sum_{I \subseteq [n], |I| \leq k} (-1)^{|I|+1} \mathbb{P}\left(\bigcap_{i \in I} E_i\right) \\
&\stackrel{(a)}{=} (1 + o(1)) \sum_{i=1}^k \binom{n/d}{i} (-1)^{i+1} \mathbb{P}\left(\bigcap_{j=1}^i E_j\right) \\
&\stackrel{(b)}{=} (1 + o(1)) \sum_{i=1}^k \binom{n/d}{i} (-1)^{i+1} \binom{n-id}{s-id} \bigg/ \binom{n}{s} \\
&\stackrel{(c)}{\leq} (1 + o(1)) \sum_{\substack{i \leq k \\ i \text{ is odd}}} \binom{n/d}{i} \frac{144s(n-id)}{(12s-1)(12(n-id)-1)} \\
&\quad \left(\frac{n-id}{n} \right)^{n-id+0.5} \left(\frac{s}{s-id} \right)^{s-id+0.5} \left(\frac{s}{n} \right)^{id} - (1 + o(1)) \\
&\quad \sum_{\substack{i \leq k \\ i \text{ is even}}} \binom{n/d}{i} \left(\frac{n-id}{n} \right)^{n-id+0.5} \left(\frac{s}{s-id} \right)^{s-id+0.5} \left(\frac{s}{n} \right)^{id} \\
&\stackrel{(d)}{=} (1 + o(1)) \sum_{\substack{i \leq k \\ i \text{ is odd}}} \binom{n/d}{i} (1 + \beta'(n, k)) (1 - \alpha'(n, k)) \left(\frac{s}{n} \right)^{kd} - \\
&\quad \sum_{\substack{i \leq k \\ i \text{ is even}}} \binom{t}{k} (1 - \alpha'(n, k)) \left(\frac{s}{n} \right)^{kd}
\end{aligned}$$

$$\begin{aligned}
& \stackrel{(e)}{=} 1 - (1 - \delta^d)^{n/d} - \sum_{i=k+1}^{n/d} \binom{n/d}{i} (-1)^{i+1} \delta^{id} + o(1) \\
& \stackrel{(f)}{=} 1 - (1 - \delta^d)^{n/d} + o(1) \\
& \stackrel{(g)}{=} o(1).
\end{aligned} \tag{39}$$

The above, step (a) is based on the symmetry of events E_i . Step (b) utilizes the definition of event E_i and the structure of coding matrix A^{FRC} . Step (c) utilizes Sterlin's inequality (25). In the step (d), since $k = n^{0.6}$, we have $id = o(n)$ for $1 \leq i \leq k$ and

$$\begin{aligned}
& \left(\frac{n-id}{n} \right)^{n-id+0.5} \left(\frac{s}{s-id} \right)^{s-id+0.5} = 1 - \alpha'(n, k) \\
& \lim_{n \rightarrow \infty} \alpha'(n, k) = 0,
\end{aligned}$$

and

$$\begin{aligned}
& \frac{144s(n-id)}{(12s-1)(12(n-id)-1)} = 1 + \beta'(n, k) \\
& \lim_{n \rightarrow \infty} \beta'(n, k) = 0.
\end{aligned}$$

Step (e) is based on the similar argument in the proof of (28). The last step utilizes the fact that, when $i \geq n^{0.6}$ and $d = \log(n \log(1/\delta))/\log(1/\delta)$,

$$\binom{n/d}{i} \delta^{id} \leq \left(\frac{en}{di} \delta^d \right)^i \leq n^{-0.6n^{0.6}}.$$

The last step (g) is based on the choice of d such that $(1 - \delta^d)^{n/d} = e^{-1/\log(n \log(1/\delta))} = 1 - o(1)$. Therefore, the theorem follows. $\square$

E PROOF OF THEOREM 5

PROOF. Define the indicator function

$$X_i = \begin{cases} 1, & \|A_{S,i}\|_0 = 0 \\ 0, & \|A_{S,i}\|_0 > 0 \end{cases} \tag{40}$$

Then we can obtain

$$\mathbb{P}[\text{err}(A_S) > c] \geq \mathbb{P} \left[\sum_{i=1}^n X_i > c \right]. \tag{41}$$

Based on the similar proof of Lemma 1, we have

$$\min_{A \in \mathcal{A}_n^d} \mathbb{P}[\text{err}(A_S) > c] \geq \min_{A \in \mathcal{U}_n^d} \mathbb{P} \left[\sum_{i=1}^n X_i > c \right]. \tag{42}$$

Suppose that

$$A^* = \arg \min_{A \in \mathcal{U}_n^d} \mathbb{P} \left[\sum_{i=1}^n X_i > c \right]. \tag{43}$$

Based on the results in the Lemma 2, we can construct a set I_d^* such that $|I_d^*| \geq \lfloor n/d^2 \rfloor$ and

$$\text{supp}(A_i^*) \cap \text{supp}(A_j^*) = \emptyset, \forall i \neq j, i, j \in I_d^*. \tag{44}$$

Therefore, we have

$$\min_{A \in \mathcal{A}_n^d} \mathbb{P}[\text{err}(A_S) > c] \geq \mathbb{P}[Y(I_d^*) > c], \quad (45)$$

where random variable $Y(I_d^*) = \sum_{i \in I_d^*} X_i$.

Suppose that $d = o(s)$. Each indicator function X_i is a Bernoulli random variable with

$$\begin{aligned} \mathbb{P}(X_i = 1) &= \frac{\binom{n-d}{s-d}}{\binom{n}{s}} \\ &\stackrel{(a)}{=} (1 + o(1)) \left(1 - \frac{d}{n}\right)^{n-d+0.5} \left(1 + \frac{d}{s-d}\right)^{s-d+0.5} \delta^d \\ &\stackrel{(b)}{=} (1 + o(1)) \delta^d. \end{aligned}$$

The above, step (a) utilizes sterlin's approximation; step (b) is based on the fact that $s = \Omega(1)$ and $d = o(s)$. First, the expectation of $Y(I_d^*)$ is given by

$$\mathbb{E}[Y(I_d^*)] = t(1 + o(1)) \delta^d. \quad (46)$$

Furthermore, considering the fact that, for any $i, j \in I_d^*$ with $i \neq j$, the random variable $X_i X_j$ is a also Bernoulli random variable with

$$\begin{aligned} \mathbb{P}(X_i X_j = 1) &= \frac{\binom{n-2d}{s-2d}}{\binom{n}{s}} \\ &= (1 + o(1)) \left(1 - \frac{2d}{n}\right)^{n-2d+0.5} \left(1 + \frac{2d}{s-2d}\right)^{s-2d+0.5} \delta^{2d} \\ &= (1 + o(1)) \delta^{2d}, \end{aligned}$$

the variance of $Y(I_d^*)$ is given by

$$\begin{aligned} \text{Var}[Y(I_d^*)] &= \mathbb{E}[Y^2(I_d^*)] - \mathbb{E}^2[Y(I_d^*)] \\ &= \mathbb{E} \left[\sum_{i \in I_d^*} X_i^2 + 2 \sum_{i, j \in I_d^*, i \neq j} X_i X_j \right] - \mathbb{E}^2[Y(I_d^*)] \\ &= (1 + o(1)) \left[t \delta^d + t(t-1) \delta^{2d} - t^2 \delta^{2d} \right] \\ &= (1 + o(1)) \left[t \delta^d (1 - \delta^d) \right] \end{aligned} \quad (47)$$

Therefore, utilizing the Chebyshev inequality, we have the following upper bound.

$$\begin{aligned} &\mathbb{P} \{ Y(I_d^*) \leq \mathbb{E}[Y(I_d^*)] - 2\mathbb{E}^{0.5}[Y(I_d^*)] \} \\ &\leq \mathbb{P} \{ |Y(I_d^*) - \mathbb{E}[Y(I_d^*)]| \geq 2\mathbb{E}^{0.5}[Y(I_d^*)] \} \\ &\leq \frac{\text{Var}[Y(I_d^*)]}{4\mathbb{E}[Y(I_d^*)]} = (1 + o(1)) \frac{1 - \delta^d}{4}. \end{aligned} \quad (48)$$

Assume that $c \leq \mathbb{E}[Y(I_d^*)] - 2\mathbb{E}^{0.5}[Y(I_d^*)]$, then we have

$$\mathbb{P} \{ Y(I_d^*) \leq c \} \leq \mathbb{P} \{ Y(I_d^*) \leq \mathbb{E}[Y(I_d^*)] - 2\mathbb{E}^{0.5}[Y(I_d^*)] \} \leq 1/4. \quad (49)$$

This result implies that

$$\mathbb{P}[\text{err}(A_S) > c] > \frac{3}{4}, \quad (50)$$

which is a contradiction. Therefore, the parameter c should satisfy $c > \mathbb{E}[Y(I_d^*)] - 2\mathbb{E}^{0.5}[Y(I_d^*)]$, which implies that

$$\mathbb{E}[Y(I_d^*)] < 2c + 4. \quad (51)$$

Since $\mathbb{E}[Y(I_d^*)] = (1 + o(1))t\delta^d$ and $\lfloor n/d^2 \rfloor \delta^d$ is monotonically non-increasing with d , the minimum computation load should satisfy

$$d \geq \frac{\log(n \log^2(1/\delta)/(2c + 4) \log^2(n/(2c + 4)))}{\log(1/\delta)}. \quad (52)$$

Therefore, the theorem follows. $\square$

F PROOF OF THEOREM 6

We use the analysis of the decoding process as described in [11]. Based on the choice of $b = \lceil 1/\log(1/\delta) \rceil + 1$ and $\delta = s/n$, we can obtain that

$$\frac{n(1 - 2\epsilon)}{b(1 - 4\epsilon)} < n - s. \quad (53)$$

Based on the results in [11], to successfully recover $n/b(1 - \epsilon)$ blocks from $\frac{n(1-2\epsilon)}{b(1-4\epsilon)}$ received results with probability $1 - e^{-cn}$, we need to show the following inequality holds.

$$e^{-\frac{1-2\epsilon}{1-4\epsilon}\Omega'(x)} < 1 - x, \forall x \in [0, 1 - \epsilon], \quad (54)$$

where $\Omega'(x)$ is the derivative of the generating function for the degree distribution P_w . Note that

$$\Omega'(x) = \frac{1}{u+1} \left(u - \ln(1-x) + x^D - \sum_{d=D+1}^{\infty} \frac{x^d}{d} \right) \quad (55)$$

Utilizing the fact that $x^D > \sum_{d=D+1}^{\infty} x^d/d$, the theorem follows.

Received August 2019; revised September 2019; accepted October 2019