Straggler Mitigation With Tiered Gradient Codes

by

Shanuja Sasi, Lalitha Vadlamani, Vaneet Aggarwal, Sundar Rajan B

Report No: IIIT/TR/2020/-1



Centre for Communications International Institute of Information Technology Hyderabad - 500 032, INDIA August 2020

Straggler Mitigation With Tiered Gradient Codes

Shanuja Sasi, *Student Member, IEEE*, V. Lalitha^D, *Member, IEEE*, Vaneet Aggarwal^D, *Senior Member, IEEE*, and B. Sundar Rajan^D, *Fellow, IEEE*

Abstract-Coding theoretic techniques have been proposed for synchronous Gradient Descent (GD) on multiple servers to mitigate stragglers. These techniques provide the flexibility that the job is complete when any k out of n servers finish their assigned tasks. The task size on each server is found based on the values of k and n. However, it is assumed that all the n jobs are started when the job is requested. In contrast, we assume a tiered system, where we start with $n_1 \geq k$ tasks, and on completion of c tasks, we start $n_2 - n_1$ more tasks. The aim is that as long as k servers can execute their tasks, the job gets completed. This paper exploits the flexibility that not all servers are started at the request time to obtain the achievable task sizes on each server. The task sizes are in general lower than starting all n_2 tasks at the request times thus helping achieve lower task sizes which helps to reduce both the job completion time and the total server utilization.

Index Terms—Tiered gradient codes, distributed computing, gradient descent, stragglers, server utilization, completion time.

I. INTRODUCTION

ANY distributed machine learning applications require multiple servers to perform distributed computation of gradient descent. Distributed gradient descent involves division of gradient descent tasks across multiple servers, and the job is finished when all the tasks are completed. The slowest tasks that determine the job execution time are called stragglers. Coding theoretic techniques have been proposed to achieve high-quality algorithmic results in the face of uncertainty, including mitigation of stragglers. [1]–[7]. These approaches have been shown to be essential to manage stragglers in distributed computation of gradient descent [1], [2], [8]. However,

Manuscript received August 2, 2019; revised December 9, 2019 and March 1, 2020; accepted April 26, 2020. Date of publication May 6, 2020; date of current version August 14, 2020. This work was supported in part by the India Science and Engineering Research Board (SERB) through J. C. Bose National Fellowship to B. S. Rajan, Early Career Research Award (ECR/2016/000954) to V. Lalitha, and Overseas Visiting Doctoral Fellowship (OVDF) to S. Sasi, V. Aggarwal support of U.S. National Science Foundation through CSR-1618335 and Cisco. The associate editor coordinating the review of this article and approving it for publication was Q. Huang. (*Corresponding author: Vaneet Aggarwal.*)

Shanuja Sasi was with Purdue University, West Lafayette, IN 47907 USA. She is now with the Department of Electrical Communication Engineering, Indian Institute of Science, Bengaluru 560012, India (e-mail: shanuja@iisc.ac.in).

V. Lalitha is with the Signal Processing and Communications Research Center, IIIT Hyderabad, Hyderabad 500032, India (e-mail: lalitha.v@iiit.ac.in).

Vaneet Aggarwal is with the School of Industrial Engineering, Purdue University, West Lafayette, IN 47907 USA, also with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA, and also with the Department of Mathematics, King Abdulaziz University, Jeddah 21589, Saudi Arabia (e-mail: vaneet@purdue.edu).

B. Sundar Rajan is with the Department of Electrical Communication Engineering, Indian Institute of Science, Bengaluru 560012, India (e-mail: bsrajan@iisc.ac.in).

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

Digital Object Identifier 10.1109/TCOMM.2020.2992721

 W_4 W_1 W_2 W_3 D_1 D_2 D_3 D_4 D_2 D_3 D_4 D_1 D_1 D_3 D_A D_2 $a'_2g_2 + b'_2g_3$ $c'_2 g_4$ $b'_3g_4 + c'_3g_5$ $a'_3q_3 \downarrow$ $a_1'g_1 + b_1'g_2 + c_1'g_3$ $a'_4g_4 + b'_4g_1 + c'_4g_2$ Master

Fig. 1. Gradient coding: Each server transmits a scalar and master calculates the sum from the result of any two servers, with the total number of servers being four.

these approaches assume that all the distributed tasks are started at the same time, which can be shown to have a large server utilization cost. To alleviate that, this paper aims to provide a tiered framework for efficient gradient code designs that allow for starting certain tasks at the completion of some tasks with an aim to have an efficient tradeoff between the completion time of the job and the server utilization cost to complete the tasks.

In this paper, we propose a coding-theoretic approach for gradient coding, called Tiered Gradient Coding. Initially at the service request time, tasks are launched on n_1 servers. On the completion of tasks from c of the servers, tasks are launched on $n_2 - n_1 > 0$ more servers, where n_2 is the total number of servers. We note that the earlier studied gradient codes [1], [2], [8] do not have two tiers and the tasks for n_2 servers are decided at the same time. By having the flexibility of obtaining the results from c servers leads to reduction of per-server workload as compared to deciding tasks for n_2 servers at the same time. Consider as an example of gradient coding scheme in Fig. 1, where the data is split into 4 partitions D_1, D_2, D_3 and D_4 . Server W_1 computes the gradients g_1, g_2 and g_3 of the partitions D_1, D_2 and D_3 respectively. Similarly, server W_2 computes the gradients g_2, g_3 and g_4 , server W_3 computes the gradients g_3, g_4 and g_1 and server W_4 computes the gradients g_4, g_1 and g_2 . Each server sends a linear combination of the gradients they have computed. It is enough to get the results from any two servers to calculate the overall sum of gradients. The techniques to calculate the linear combination are provided in [1]. The computation cost per server is proportional to $\frac{3}{4}$. In Fig. 2, we describe the proposed tiered gradient coding framework, where the data is split into 3 partitions D'_1, D'_2 and D'_3 . Initially, only three servers $(W_1, W_2 \text{ and } W_3)$ are launched. Server W_1 computes the gradients g'_1 and g'_2 of the partitions

0090-6778 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See https://www.ieee.org/publications/rights/index.html for more information.



Fig. 2. Tiered gradient coding: Initially 3 servers, i.e. W_1, W_2 and W_3 , are launched. Without loss of generality, assume W_1 completes the task first. Then the fourth server W_4 is launched. Master waits for one more server to finish the task to calculate the sum.

 D'_1 and D'_2 respectively. Similarly, server W_2 computes the gradients g'_2 and g'_3 and server W_3 computes the gradients g'_3 and g'_1 . Without loss of generality, assume that the server W_1 finishes its task first, *i.e.*, W_1 sends a linear combination of the gradients g'_1 and g'_2 . Then, server W_4 is launched, which computes the gradients g'_3 and g'_1 of the partitions D'_3 and D'_1 respectively. The partitions assigned to W_4 depends on the server which had completed the task initially. Master waits for one of the servers - W_2 , W_3 and W_4 to complete the task. Master can calculate the sum from the result from W_1 and any one of the servers - W_2 , W_3 or W_4 . For example, if W_4 completes the task first, $g'_1 + g'_2 + g'_3 = (\frac{g'_1}{2} + g'_2) + (\frac{g'_1}{2} + g'_3)$. Thus, we see that the task per server reduces from $\frac{3}{4}$ to $\frac{2}{3}$ for the same number of four servers, and both schemes guarantee that any two servers completion can provide the required computation result.

We note that waiting to launch tasks $n_2 - n_1$ servers can affect the job completion time negatively, while the decreased task size per server affects the job completion time positively, thus making it apriori unclear whether the completion time increases or decreases. Similarly, server utilization cost (total time any server is being used for computation) may also increase if the completion time is increased, and decrease since $n_2 - n_1$ servers are not used till c complete the computation. A tradeoff between the two metrics of completion time and server utilization cost have been considered for coded tasks [9], while we show that efficient coding can also decrease the task size when such multi-forking capability (starting n_2 after c have finished) can lead to decreased task size further reducing both the metrics. As an example, we consider $n_2 =$ 15, c = 1, and k = 5, where k is the number of servers that must complete execution for the job to be completed. For three models of the time taken for each server to complete the task (with the models described in Section VI), Fig. 3 shows for different values of n_1 , the tradeoff (between service completion time and server utilization cost) points that can be achieved with tiered-gradient codes as compared to gradient coding where all n_2 tasks must be decided apriori. Since the gradient codes are independent of n_1 , they are only a point, while we achieve different tradeoff points for different n_1 . We see that the proposed codes perform significantly better

in both the completion time and server utilization costs and efficient parameters can be decided based on the application requirements.

The key contribution of the paper is a new framework for tiered gradient codes which allows for a delayed start of the tasks at the servers. A novel code construction is provided that exploits this flexibility, and reduces the amount of computation that each server has to perform.

Notation: Throughout this paper, we let d denote the number of samples, n_2 denote the total number of servers, Q denote the number of data partitions, and s denote the number of stragglers/failures. Let k denote the minimum number of servers required to finish their task such that the overall gradient can be computed. Let $n_1 < n_2$ denote the number of servers launched in the first phase. We wait for c < k servers to finish their job first when n_1 servers are launched. In the second phase the rest of $n_2 - n_1$ servers are launched. The n_2 servers are denoted as $\{W_1, W_2, \ldots, W_{n_2}\}$. The partial gradients over Q data partitions are denoted as $\{g_1, g_2, \ldots, g_Q\}$. All matrices under consideration are over real numbers. Let [z] denote the set $\{1, 2, \ldots, z\}$ and $[z_1, z_2]$ denote the set $\{z_1, z_1 + 1, \dots, z_2\}$. The support of a vector $\mathbf{f} = [f_1 \ f_2 \ f_3 \dots f_m], \ supp(\mathbf{f}), \ gives a set of all positions of$ **f** which are non zeros, i.e., $supp(\mathbf{f}) = \{i : f_i \neq 0\}$. Given a set S, 'c subset' of S is defined as a subset of S which is of length c.

II. RELATED WORK

Coding theoretic techniques have been proposed to achieve efficient algorithmic results in the face of uncertainty, including mitigation of stragglers. The different distributed computing applications include coded matrix multiplication [10], [11], MapReduce models [10], [12], [13], and computation of gradients [14]–[18].

In this paper, we focus on gradient computation, which was introduced in [1]. In [2], coding techniques to reduce the running time of distributed learning tasks have been provided. A stochastic block code and an efficient decoding method for approximate gradient recovery are provided in [19]. A distributed computing scheme called Batched Coupon's Collector to mitigate the effect of stragglers in gradient methods is proposed in [20]. In [21], a straggler mitigation scheme that facilitates the implementation of distributed gradient descent in a computing cluster is presented. They also proposed a theoretical delay model which allows to minimize the expected running time. In [22], an approximate variant of the gradient computation is introduced, in which approximate gradient computation.

A cost vs. latency analysis of using simple replication or erasure coding for straggler mitigation in executing jobs with many tasks is studied in [9], [23], [24]. In these works, the authors have shown that the delayed relaunch of stragglers yields significant reduction in cost and latency. The authors of [9], [24] considered that the computation could be complete from the result of any k out of n servers, and evaluated the notions of service completion time and server utilization cost with delayed relaunch. The work in [9], [24] does not propose



(a) Task completion time distributed as SE1 (b) Task completion time distributed as SE2 (c) Task completion time distributed as Pa

Fig. 3. Server utilization cost as a function of service completion time when we vary $n_1 \in [k, n_2]$ for $n_2 = 15$, c = 1 and k = 5.

new coding mechanisms and MDS codes cannot be directly applied for gradient coding. We show that having a delayed relaunch based on when c servers finish execution rather than a constant time as in [9], significant computation time is saved on each server. Delayed relaunch based on c servers finishing execution has been considered in [24], and such stochastic analysis can complement the study to come up with the metrics of service completion time and server utilization cost, while is not the focus of this work. The focus of our work is efficient code designs for distributed gradient computation that exploits the flexibility of delayed relaunch of servers. In this paper, we show that efficient coding further allows reduction of task size per server with a delayed execution of tasks, a flexibility which had not been studied earlier. Thus, our coding-theoretic techniques can further help reduce the job completion time by exploiting a better choice of the parameters since starting more servers need larger task size at each server in general. Thus, this paper aims at finding efficient code constructions that minimizes per-server task sizes with the flexibility of tiered launching of tasks.

III. REVIEW OF GRADIENT CODES

A. Distributed Gradient Descent Computation

Given a dataset D with d examples, $D = \{(X_i, Y_i)\}_{i=1}^d$, where $X_i \in \mathbb{R}^p$ and $Y_i \in \mathbb{R}$, we want to learn parameters $\beta \in \mathbb{R}^p$ by minimizing a generic loss function $L(D;\beta) =$ $\sum_{i=1}^{d} L(X_i, Y_i; \beta)$. We update the parameter β according to the following rule: $\beta^{(t+1)} = h(\beta^{(t)}; g^{(t)})$, where $g^{(t)} = \nabla L(D; \beta^{(t)}) = \sum_{i=1}^{d} \nabla L(X_i, Y_i; \beta^{(t)})$ is the gradient of the loss at the current estimate of the parameters and his a gradient-based optimizer. We consider the problem of distributed synchronized gradient descent where the d data samples are divided into Q partitions, D_1, D_2, \ldots, D_Q . The partial gradient computed on the j^{th} partition is given by $g_{i}^{(t)} =$ $\sum_{(X,Y)\in D_j} \nabla L(X,Y;\beta^{(t)})$. The overall gradient required to compute the update on $\beta^{(t)}$ is given by $g^{(t)} = \sum_{j=1}^{Q} g_j^{(t)}$. We will omit the superscript t in this paper to simplify the notation. Next, we provide a review of two classes of conventional gradient codes known as fractional repetition gradient codes and cyclic repetition gradient codes [1].

B. Gradient Coding Framework

For n_2 workers and Q data partitions, we have a set of linear equations: $AF = 1_{f \times Q}$, where f denotes the number

of combinations of surviving servers/non-stragglers, $1_{f \times Q}$ is the all 1 matrix of dimension $f \times Q$ and we have matrices $A \in \mathbb{R}^{f \times n_2}, F \in \mathbb{R}^{n_2 \times Q}$. The *i*th row of F, \mathbf{f}_i , is associated with the i^{th} server W_i . The support of \mathbf{f}_i , $supp(\mathbf{f}_i)$, represents the data partitions corresponding to the server W_i and the entries of f_i encode a linear combination over their gradients that server W_i transmits. Let $\mathbf{g} \in \mathbb{R}^{Q \times p}$ be a matrix with each row being the partial gradient of a data partition i.e. $\mathbf{g} = [g_1, g_2, \dots, g_Q]^T$. Then, server W_i transmits $\mathbf{f}_i \mathbf{g}$. Each row of A, denoted by \mathbf{a}_i , is associated with a specific straggler scenario, to which tolerance is desired. In particular, any row \mathbf{a}_i , with support $supp(\mathbf{a}_i)$, corresponds to the scenario where the server indices in $supp(\mathbf{a}_i)$ are non-stragglers. The entries of a_i encode a linear combination which, when taken over the transmitted gradients of the non-straggler servers, $\{\mathbf{f}_{u}\mathbf{g}\}_{u \in supp(\mathbf{a}_{i})}$, would yield the full gradient. We refer to this system as (n_2, k) gradient code where k is the number of non stragglers.

C. Fractional Repetition Gradient Codes [1]

Consider the case when $n_2 - k + 1$ divides n_2 . Let $Q = n_2$. Consider the following matrix $F_j = 1_{(n_2-k+1)\times(n_2-k+1)}$, $1 \le j \le \frac{n_2}{n_2-k+1}$. The matrix F of the fractional repetition gradient code is constructed as follows:

$$F = \begin{bmatrix} F_1 & 0 & \dots & 0 \\ 0 & F_2 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & F_{\frac{n_2}{n_2 - k + 1}} \end{bmatrix}$$

D. Cyclic Repetition Gradient Codes [1]

This class of gradient codes exist for all values of k and n_2 . Let $Q = n_2$ and let the columns of F be indexed by $[0, n_2-1]$. The support structure of the matrix F is as follows:

$$supp(\mathbf{f}_i) = [i - 1, i + (n_2 - k - 1)] \mod n_2.$$
 (1)

Now, we will present a randomized construction of the matrix F. Consider a matrix H of size $(n_2 - k) \times n_2$ whose first $(n_2 - 1)$ columns are picked at random i.i.d. from a Gaussian distribution $\mathcal{N}(0, 1)$. The last column of H is obtained as follows: $H(:, n_2 - 1) = -\sum_{i=0}^{n_2-2} H(:, i)$. Each vector \mathbf{f}_i is calculated by solving the following equation $\mathbf{f}_i|_{L_i}H(:, L_i)^T = 0$, where L_i is the support of \mathbf{f}_i as given by Equation (1).

IV. TIERED GRADIENT CODE FRAMEWORK

In this section, we give the framework of tiered gradient coding and present a span condition which is necessary for a code to be tiered gradient code. Based on this framework, the various parameters achieved by tiered gradient codes are provided in Section V.

A. Tiered Gradient Codes

In the conventional gradient code framework, we assume that there are n_2 servers which start computing the partial gradients assigned to them. We want to be able to compute the overall gradient whenever $k < n_2$ servers finish. Each server sends a linear combination of the partial gradients which it has computed and sends it back to the master node. The master node aggregates all the linear combinations of the partial gradients and performs a linear combination in turn to obtain the overall gradient $\sum_{j=1}^{Q} g_i$.

In the tiered gradient code framework, we consider two phases. Initially, before the first phase where the first n_1 servers start computing the partial gradients from the data partitions assigned to them, the master node transfers the assigned data partitions to the first n_1 servers. For the rest of the $n_2 - n_1$ servers, the master node transfers all the data partitions. Hence, the first n_1 servers store only the data partitions assigned to them while all other servers store all the data partitions.

In the first phase, n_1 servers start computing the partial gradients from the data partitions assigned to them and c out of the n_1 servers complete their gradient computation by the end of the first phase. In the second phase, $n_1 - c$ servers continue their tasks which were started in the first phase and $n_2 - n_1$ servers start computing the partial gradients assigned to them in the second phase depending upon which servers have completed their tasks in the first phase.

The assignment of the data partitions to $n_2 - n_1$ servers for computation in the second phase is decided based on which c servers out of the n_1 servers have finished. The number of data partitions assigned to a server launched in the first phase to do computation is the same as that assigned to a server launched in the second phase. Hence, even though the entire data is available with the later initialized nodes, we do not take that into account in this paper. We wish to reduce the computation load. The assignment of data partitions for computation is the same in both the phases. We want to be able to compute the overall gradient whenever k servers out of the n_2 servers finish. This condition is equivalent to saying that we would need the results from any k-c out of the n_2-c servers to complete in the second phase, so that we can compute the overall gradient. We call this set up as (n_1, n_2, k, c) tiered gradient coding.

B. Span Condition of Tiered Gradient Codes

Consider the Q partial gradients arranged in a column vector as $\mathbf{g} = [g_1, g_2, \dots, g_Q]^T$. Let F denote a matrix of size $n_1 \times Q$ over \mathbb{R} . The *i*th row of the F matrix is denoted by \mathbf{f}_i , $1 \leq i \leq n_1$. The support of \mathbf{f}_i indicate the partial gradients which are computed on the i^{th} server. The quantity $f_i g$ is the linear combination sent by the i^{th} server to the master node.

Let $M \subset [n_1]$ denote the set of c servers which have finished their tasks at the end of the first phase and \mathcal{M} denote the set of all possible c subsets of $[n_1]$. Let $\{B_M, M \in \mathcal{M}\}$ denote a set of matrices, each of size $(n_2 - n_1) \times Q$. The *i*th row of B_M is denoted by $\mathbf{b}_i, 1 \leq i \leq n_2 - n_1$. The support of \mathbf{b}_i indicate the partial gradients which are computed on the *i*th server among the $n_2 - n_1$ servers started in the second phase. The quantity \mathbf{b}_i g is the linear combination sent by the *i*th server to the master node. We also assume that for all the $\binom{n_1}{c}$ possibilities of c servers completing the tasks from a total of n_1 servers (set of all possible M), the B_M matrices are stored by the master node initially. After the first phase, the master node knows the set M. Hence the corresponding B_M matrix is taken and the master node assigns the data partitions to the rest of the servers from the respective rows of the B_M matrix.

Let $\{A_M, M \in \mathcal{M}\}\$ denote a set of matrices, each of size $N \times n_2$, where N represents the number of possible straggler configurations which can be tolerated by a tiered gradient code. Columns of the A_M matrix are indexed by the servers. The rows of the A_M matrix are denoted by \mathbf{a}_i . Each \mathbf{a}_i has non-zeros in the c positions corresponding to the subset M. The rows of the A_M matrix are such that each row will have non-zeros in a distinct subset of k-c out of the n_2-c positions. Hence, the number of possible straggler configurations which can be tolerated by a tiered gradient code described above, for a given set of c servers, is $N = \binom{n_2-c}{k-c}$.

Definition 1: Consider a set of matrices $(F, \{(A_M, B_M), M \in \mathcal{M}\})$, where F is a $n_1 \times Q$ matrix, A_M is a $N \times n_2$ matrix and B_M is a $(n_2 - n_1) \times Q$ matrix, for all $M \in \mathcal{M}$. If this set of matrices obeys Equation (2), then we say that this set of matrices generates a (n_1, n_2, k, c) tiered gradient code, where $1_{N \times Q}$ denotes a matrix with all its entries being 1.

$$A_M \begin{bmatrix} F\\ B_M \end{bmatrix} = 1_{N \times Q}, \quad \forall M \in \mathcal{M}.$$
 (2)

Lemma 1 provides the necessary condition for a code to be a tiered gradient code.

Claim 1: The partial gradients which are computed on the first n_1 servers have to constitute a (n_1, k) gradient code.

Proof: We need to be able to compute the overall gradient whenever k servers finish. This includes the c servers which have computed the gradient in the first phase. Now, since $n_1 \ge k$, all the k servers can be potentially from the first n_1 servers and since the property has to held for all possibilities of c servers, the claim follows.

Lemma 1 (Span Condition): Consider a tiered gradient code generated by the set of matrices $(F, \{(A_M, B_M), M \in \mathcal{M}\})$. For each $M \in \mathcal{M}$, $I_1 \subseteq [n_1] \setminus M$ and $I_2 \subseteq [n_2 - n_1]$, such that $|I_1 \cup I_2| = k - c$, it holds that $1_{1 \times Q} \in span\{\mathbf{f}_i, \mathbf{b}_j | i \in M \cup I_1, j \in I_2\}$.

Proof: We consider one row of Equation (2) for a fixed M and is given by

$$\mathbf{a}_i \begin{bmatrix} F\\ B_M \end{bmatrix} = \mathbf{1}_{1 \times Q}.\tag{3}$$

4635

Let $L_i = supp(\mathbf{a}_i) = M \cup I_1 \cup \{I_2\}$. The above equation can be rewritten as $\mathbf{a}_i|_{L_i} \begin{bmatrix} F|_{M \cup I_1} \\ B_M|_{I_2} \end{bmatrix} = \mathbf{1}_{1 \times Q}$, where $\mathbf{a}_i|_{L_i}$ represents the vector of length $|L_i|$ obtained by taking the non zero elements in the vector \mathbf{a}_i . The matrix $F|_{M \cup I_1}$ is obtained from F by taking only the rows in F corresponding to the indices in $M \cup I_1$ and $B_M|_{I_2}$ is obtained by taking the rows in the B_M matrix corresponding to the indices in I_2 . From Equation (3), it is clear that there exists a non-zero vector $\mathbf{a}_i|_{L_i}$ such that the above equation is true. Hence, we can solve for $\mathbf{a}_i|_{L_i}$ and hence \mathbf{a}_i can be solved using the above equation.

The span condition for the conventional gradient code framework in [1] follows from Lemma 1 by considering $n_2 = n_1$ and $A_M = A$.

In order to show that the span condition in Lemma 1 is satisfied by the tiered gradient codes with probability 1, it is enough to show that the following support condition holds for the code under consideration and rest of the arguments follow exactly as in the proof of Lemma 1 in [1] (and hence omitted). We will now present the support condition for tiered gradient codes which is a sufficient condition to show that the randomized construction of the F and B_M matrices satisfy the span condition in Lemma 1.

Condition 1: Consider a set of matrices $(F, \{(A_M, B_M), M \in \mathcal{M}\})$. For every $M \in \mathcal{M}$, $I_1 \subseteq [n_1] \setminus M, I_2 \subseteq [n_2 - n_1]$, such that $|I_1 \cup I_2| = k - c$ and for every $\mathcal{T}_1 \subset M \cup I_1$ and $\mathcal{T}_2 \subset I_2$ of size $|\mathcal{T}_1| + |\mathcal{T}_2| = \ell$ $(1 \leq \ell \leq k)$, it needs to satisfy the following inequality for the above set of matrices to represent a tiered gradient code: $|\bigcup_{i \in \mathcal{T}_1} L_i \cup_{j \in \mathcal{T}_2} Z_j| \geq (n_1 - k) + \ell$.

V. MAIN RESULTS

In this section, we provide our results for (n_1, n_2, k, c) tiered gradient codes for the entire range of n_1 and n_2 . We define the amount of computation per server as the fraction of data that is used by a server to perform computation. More the data, more is the computation time.

Theorem 1: The amount of computation per server of $(n_1, n_2, k, c = 1)$ tiered gradient code is as follows

- 1) For $n_1 = k, n_2 n_1 = 1$ and even k, the amount of computation per server is $\frac{2(k-1)}{k^2}$.
- 2) For $k \le n_1 \le 2(k-1)$ and $n_2 > 2(k-1)$, the amount of computation per server is $\frac{1}{2}$.
- 3) For $n_1, n_2 \ge 3(k-1)$, the amount of computation per server is $\frac{n_2-k+1-G_1}{n_2-G_1}$, where $G_1 = \max\{\min\{n_2 - (n_1 + p^*), \lfloor \frac{n_1+p^*-k+1}{k-1} \rfloor\}, \min\{n_2 - n^+, C_{n+}^*\}, n_2 - n_{min}\}, p^* = \lceil n_2 - n_1 - \frac{n_2-k+1}{k} \rceil, n^+ = \min\{n' \in \lceil \max\{n_1, n_2 - 6\}, n_2 - 1 \rceil\}$ such that

 $C_{n^+}^* = \max_{n' \in [\max\{n_1, n_2 - 6\}, n_2 - 1]} C_{n'}^*, n_{min} = \min\{n'' \in [\max\{n_1, n_2 - 6\}, n_2 - 1]\}$ such that $n_2 \leq n'' + C_{n''}^*$. The values of $C_n^* \forall n \in [n_2]$ are provided in Table II.

4) For $2(k-1) < n_1 < 3(k-1)$ and $n_2 \ge 3(k-1)$, the amount of computation per server is $\frac{n_2-k+1-G_2}{n_2-G_2}$, where $G_2 = \max\{\min\{n_2 - (n_1 + p^*), \lfloor \frac{n_1 + p^* - k + 1}{k-1} \rfloor\},\$



Fig. 4. The amount of computation per server required as a function of n_1 when we vary $n_1 \in [k, n_2]$ for $n_2 = 19, c = 1$ and k = 5.

 $\min \{n_2 - n^+, C_{n^+}^*\}, n_2 - n_{\min}\}, p^* = \\ \max \{3(k-1) - n_1, \lceil n_2 - n_1 - \frac{n_2 - k + 1}{k} \rceil\}, n^+ = \\ \min \{n' \in [\max \{3(k-1), n_2 - 6\}, n_2 - 1]\} \quad \text{such} \\ \text{that} \quad C_{n^+}^* = \max_{n' \in [\max \{n_1, n_2 - 6\}, n_2 - 1]} C_{n'}^*, \\ n_{\min} = \min \{n'' \in [\max \{3(k-1), n_2 - 6\}, n_2 - 1]\} \\ \text{such that} \quad n_2 \leq n'' + C_{n''}^*. \text{ The values of } C_n^* \; \forall n \in [n_2] \\ \text{are provided in Table II.}$

Theorem 2: The amount of computation per server of $(n_1, n_2, k, c > 1)$ tiered gradient code is as follows

- 1) For $n_1 \leq 2(k-1)$ and $n_2 > 2(k-1)$, the amount of computation per server is $\frac{1}{2}$.
- 2) For $n_1 \ge 2(k-1) + (k-c)$, the amount of computation per server is $\frac{n_2-k+1-G_3}{n_2-G_3}$, where $G_3 = \min\{n_2 n_1, \lfloor \frac{n_1+p^*-k+c}{k-1} \rfloor\}, p^* = \lceil n_2 n_1 \frac{n_2-k+c}{k} \rceil$.
- 3) For $2(k-1) < n_1 < 2(k-1) + (k-c)$ and $n_2 \ge 2(k-1) + (k-c)$, the amount of computation per server is $\frac{n_2 k + 1 G_4}{n_2 G_4}$, where

$$G_{4} = \min\{n_{2} - (2(k-1) + (k-c)), \\ \lfloor \frac{n_{1} + p^{*} - k + c}{k-1} \rfloor\}, \\ p^{*} = \max\{(2(k-1) + (k-c)) - n_{1}, \\ \lceil n_{2} - n_{1} - \frac{n_{2} - k + c}{k} \rceil\}.$$

We note that the amount of computation per server is $\frac{n_2-k+1}{n_2}$ for the (n_2, k) gradient code, and the proposed tiered gradient codes reduce this amount to that in the statement of Theorem 1 due to the flexibility of delayed start of $n_2 - n_1$ servers when c have finished computation. Fig. 4 illustrates the savings in the amount of computation per server for tiered gradient codes as compared to the gradient codes for c = 1. We note that as the value of n_1 increase, the amount of computation per server for smaller n_1 , one possibility is to use the code construction of larger n_1 and only start the required number of servers initially and rest can be started when c smaller n_1 providing more savings as n_1 decreases.

For all the cases discussed, the tiered gradient coding performs better than the gradient coding in terms of the amount



distributed as SE2.

Fig. 5. Server utilization cost as a function of service completion time when we vary $n_1 \in [k, n_2]$ for $n_2 = 12$, c = 1, and k = 3.

of computations per server required. We provide the code constructions for Theorem 1 in Section VII. Further, Section VIII provides the code construction for c > 1 (Theorem 2).

VI. NUMERICAL EVALUATIONS

In this section, we compare two metrics for the proposed tiered gradient codes to that for the gradient codes in [1]. The first metric is the service completion time, defined as the time taken for the k tasks to complete. The second metric is the server utilization cost, which is the sum over all n_2 servers, the time during which each of the server is used till the job completes. Since both the metrics are random variables, dependent on the execution times of the tasks, we average the metrics over 10^4 random trials to get mean results.

Two distribution models are typically used to model the task execution times at the servers, both these distributions model the effect of stragglers in the job computation. The first is the shifted exponential distribution. For a shifted exponential distribution with probability distribution given by $\Pr(T > x) = e^{-\mu(x-d)^+}$ for all x > 0, the shift parameter d, the mean is given by $d + \frac{1}{u}$. The distribution is motivated by Tahoe experiments [25], Ceph Experiments [26], and Amazon S3 experiments [27]. The shift is in part due to the read/write times and the minimum transfer time, and the randomness is due to the background processes. In our simulation, we have taken the total number of data points as D = 10,000. The shift parameter happens from a combination of disk I/O and computation, and thus we consider two models for this. The first, called SE1, is where d = 0.5 times the computation per server required as in Theorem 1 whereas $\frac{1}{\mu}$ is taken as 2 times the the computation per server required. The second, called SE2, is where d = 100 and is independent of the computation per server required, modeling more of the disk I/O rather than computation. For gradient coding case considered in the second model, $\frac{1}{\mu}$ is taken as 2 times the the computation per server required which is around 8,000 (for $n_2 = 12, c = 1$ and k = 3).

The second distribution model that is considered is the Pareto distribution, which has probability distribution of task execution at each server as $Pr(T > x) = (\min(x, x_m)/x)^{\alpha}$ for all x, where x_m is the scale parameter and α is the shape parameter. For our evaluations, we let $\alpha = 1.5$, and have



function of c as a function of c Fig. 6. Server utilization cost and service completion time as a function of c for $n_2 = 15$, $n_1 = 8$, and k = 5. Task completion time is assumed to be

 $x_m = 1$ times the computation per server requires as given in Theorem 1. This distribution is label Pa.

We first consider $n_2 = 15$, c = 1, and k = 5, and vary n_1 from 5 to 15. The tradeoff between the server utilization cost and service completion time for both the proposed codes and the codes in [1] are depicted in Fig. 3a, 3b, and 3c for SE1, SE2, and Pa, respectively. In all three cases, the point with lowest service completion time and server utilization cost corresponds to $n_1 = 2(k-1)$. Thus, the decrease in task size more than compensates the increase in expected completion time due to the delayed launching of $n_2 - n_1$ tasks. The use of efficient tiered gradient codes decrease both the metrics significantly for $n_1 = 2(k-1)$ as compared to the gradient codes which corresponds to $n_1 = n_2$. We also consider a different case - $n_2 = 12$, c = 1, and k = 3 and plot the trade off between the server utilization cost and service completion time in Fig. 5a, 5b, and 5c for SE1, SE2, and Pa, respectively, and achieve the same conclusions. We note that there is no monotonically relation with the parameters n_1 for the two metrics which are in part due to the code constructions having discrete changes. The proposed codes help choose parameters that can help system designer trade off the two metrics more efficiently. In Fig. 3a, we see more than 25% decrease in the both the metrics for tiered gradient codes at $n_1 = 2(k-1)$ as compared to the gradient codes thus showing that delayed relaunching is helpful and the code construction reduces the amount of computation efficiently.

So far, we assumed c = 1. We next consider the impact of general c. We let $n_2 = 15$, $n_1 = 8$, k = 5 in Fig. 6.

TABLE I	
TABLE THAT COMPARES GRADIENT CODING AND	TIERED GRADIENT CODING

	Service Completion Time	Server Utilization Cost
Gradient Coding	67.6776	197.6694
Tiered Gradient Coding	38.4224	71.0327

We note that the server utilization cost decreases with c since $n_2 - n_1$ servers are not started at t = 0, and wait till completion of c servers. However, the service completion time increases with c since the delayed starting of tasks lead to a delay in waiting for k tasks to finish. However, for c = 1and c = 2, both the metrics are significantly lower in the proposed approach as compared to the gradient codes. For c = 3, the server utilization cost for the proposed codes is significantly lower for the proposed codes at an expense of the service completion time. Thus, both the metrics may need to be taken into account together for deciding the code parameters for the tiered gradient codes. The proposed codes give additional degrees of flexibility in the design that can lead to significantly improved performance in the different metrics of the use of distributed servers, including the task per server, server utilization cost, and service completion time.

A. Real Data Implementation

Intel DevCloud is a cloud computing service made available by Intel [28] for several profiles of researchers, students and professional engineers. Intel DevCloud is a compute cluster, consisting of multiple servers called compute nodes, storage servers, and the login node. Each node has Intel Xeon processor of the Skylake architecture (Intel Xeon Scalable Processors family), an Intel Xeon Gold 6128 CPU, on-platform memory of 192 GB and a Gigabit Ethernet interconnect. In order to do evaluations of gradient descent, we launched the tasks of training ResNet on CIFAR10 dataset.¹ The training samples were augmented to allow for large batch sizes. As a proxy to the amount of computation per node, batch-size was varied in proportion. We evaluated the proposed codes for $k = 3, n_1 =$ 2(k-1), and $n_2 = 12$, where the batch size of 80,000 was used for gradient codes and 48,000 for tiered gradient codes. Running the codes on n_2 different nodes, we calculated the training time per epoch for each of the batch-sizes. These run-times were then post-processed to compute the metrics for service completion time and server utilization cost. Further, the results were computed for 1000 epochs thus giving 1000 data points for computing the average of the two metrics. The two metrics for the gradient codes and the tiered gradient codes can be seen in Table I, which demonstrates that the proposed framework reduces the service completion time by 43% and the server utilization cost by 64% as compared to gradient codes in the considered scenario on realistic evaluation on Intel DevCloud.

VII. TIERED GRADIENT CODES FOR c = 1

In this section, we provide the code constructions for Theorem 1. Precisely, the achievability part for line 1 in Theorem 1 is discussed in Section VII-A, line 2 in in Section VII-B, line 3 in Sections VII-C and VII-D, and line 4 in Section VII-E.

A. Tiered Gradient Codes for $n_1 = k, n_2 - n_1 = 1$ (Even k)

In this section, we provide tiered gradient codes for $n_1 = k, n_2 - n_1 = 1$, where k is even. The computation per server required is $\frac{2(k-1)}{k^2}$. Let t = k - 1. We split the data into $\frac{k^2}{2}$ partitions. Each user is assigned t partitions of data. The code construction is as follows.

Construction 1: $(n_1 = k, n_2 - n_1 = 1, k \text{ even})$. The support structure of the matrix F is as follows: $supp(\mathbf{f_i}) = [(i-1)\left\lceil \frac{t}{2} \right\rceil, (i-1)\left\lceil \frac{t}{2} \right\rceil + (t-1)] \mod n_1$. If server $m \in [1, n_1]$ finishes the task first in the first phase, the support of the B_M matrix, which is a row vector is as follows: $Z_1 = \bigcup_{j \in [0,k-1] \setminus (m-1)} (j-1) \left\lceil \frac{t}{2} \right\rceil + t \mod n_1$, *Proof:* We split data into $\frac{k^2}{2}$ parts, namely $\{x_0, x_1, \dots, x_{\frac{k^2}{2}-1}\}$. The support of the first row of

Proof: We split data into $\frac{k^2}{2}$ parts, namely $\{x_0, x_1, \ldots, x_{\frac{k^2}{2}-1}\}$. The support of the first row of the F matrix is [0, t - 1]. Each row of the F matrix is obtained by shifting the previous row by $\lceil \frac{t}{2} \rceil$ towards right. Any two consecutive servers have exactly $\lfloor \frac{t}{2} \rfloor$ partitions of data in common. Server i and i + 1 have $\{x_{(i)} \lceil \frac{t}{2} \rceil, x_{(i)} \lceil \frac{t}{2} \rceil + 1, \ldots, x_{(i+1)} \lceil \frac{t}{2} \rceil - 1\}$ in common. Initially the first n_1 servers are launched. Let us assume that server m finishes the task first. Then one more server is launched. The content of this server includes the partitions of data which are unique to each of the first n_1 servers except server m. The data which is unique to the server $i, i \in [0, n_1]$ is $x_{(i-1)} \lceil \frac{t}{2} \rceil + \lfloor \frac{t}{2} \rfloor$. Hence, k - 1 = t partitions of data are included in the newly added server.

Theorem 3: The code given in Construction 1 is a $(n_1, n_2 = n_1 + 1, k = n_1, c = 1)$ tiered gradient code where k is even.

Proof: We need to prove that the support condition given in Condition 1 is satisfied by the code in Construction 1. Here, $M = \{m\}$. The F matrix is a circulant matrix with each row shifted by $\lceil \frac{t}{2} \rceil$ towards right from the previous row. Hence, if $\mathcal{T}_2 = \phi$, Condition 1 holds. We will now consider the case when $|\mathcal{T}_2| = 1$, $m \notin \mathcal{T}_1$, $|\mathcal{T}_1| = k - 2$. That is precisely when we have taken all the servers from the first n_1 servers except server m and one more server which is referred as server b. When we picked the coordinates for the server n_2 , we have included the coordinate which is unique to the server b. Hence, Condition 1 is satisfied. Thus, Condition 1 is satisfied for all cases and hence the code is a $(n_1, n_2 = n_1 + 1, k = n_1, c = 1)$ tiered gradient code for even k.

B. Tiered Gradient Codes for $k \le n_1 \le 2(k-1)$

In this subsection, we assume that $k \leq n_1 \leq 2(k-1)$ and provide a construction of tiered gradient codes based on fractional repetition gradient codes. The goal is to design tiered

¹The used source-code can be seen at https://github.rcac.purdue.edu/Clan-labs/TieredGradient/

gradient codes which perform smaller computation per server than that is demanded by an (n_2, k) gradient code. Here we assume that $n_2 > 2(k-1)$.

Construction 2: We pick $p \triangleq 2(k-1) - n_1$. Let Q = 2(k-1). 1). Consider the matrices F_1 and F_2 , where $F_1 = 1_{\lceil \frac{n_1}{2} \rceil \times \lfloor \frac{Q}{2} \rfloor}$ and $F_2 = 1_{\lfloor \frac{n_1}{2} \rfloor \times \lfloor \frac{Q}{2} \rfloor}$. The matrix F corresponding to the tiered gradient code is given by $F = \begin{bmatrix} F_1 & 0 \\ 0 & F_2 \end{bmatrix}$. The matrix B_M is as follows:

- If $M \subseteq \{1, \dots, \lceil \frac{n_1}{2} \rceil\}$, then, $B_M = \left[0 \ 1_{(n_2 n_1) \times (\frac{Q}{2})}\right]$ • For all other possibilities of M, we set, B_M
- For all other possibilities of M, we set, $B_M = \begin{bmatrix} 1_{(n_2-n_1)\times(\frac{Q}{2})} & 0 \end{bmatrix}$,

where $1_{A \times B}$ is $A \times \overline{B}$ matrix with all entries as 1.

Theorem 4: The code given in Construction 2 is a (n_1, n_2, k, c) tiered gradient code where $k \le n_1 \le 2(k-1)$.

Proof: To prove the theorem, we will check that the code satisfies the span condition given in Lemma 1. It is clear that whenever $I_2 \neq \phi$, span condition is satisfied. Consider the case when $I_2 = \phi$. In this case, we need that $1_{1 \times Q}$ to lie in the span of any k rows of the F matrix. Since $k \ge \frac{n_1}{2} + 1 > \lceil \frac{n_1}{2} \rceil$, we will have at least one row each from the two types of rows (corresponding to the F_1 and F_2 matrices) and hence the span condition is satisfied.

The proof of Theorem 1 for $k \le n_1 \le 2(k-1)$ is as follows: The number of samples per partition is $\frac{d}{Q}$. The computation per server for the (n_2, k) gradient code is proportional to $\frac{d}{n_2}(n_2 - k + 1)$. The computation per server for the tiered gradient code given in Construction 2 is proportional to $\frac{d}{n_1+p}(n_1 + p - k + 1)$. The theorem follows by noting that when $n_1 + p < n_2$,

$$\frac{1}{2} = \frac{n_1 + p - k + 1}{n_1 + p} < \frac{n_2 - k + 1}{n_2}.$$

Example 1: As an example, let $n_1 = 7$, k = 5, c = 1, and $n_2 = 10$. In this example, we have a total of $n_2 = 10$ servers and we launch $n_1 = 7$ servers initially. For this case, we divide data into eight partitions $\{D_1, D_2, \dots, D_8\}$. The first four servers, out of the 7 servers launched initially, compute the sum of gradients of the first four partitions (D_1, D_2, D_3, D_4) and the next three servers, out of the 7 servers launched initially, compute the sum of gradients of the first four partitions the last four partitions (D_5, D_6, D_7, D_8) .

After launching the first 7 servers, we wait for any one server to finish its task. Suppose if the first server finishes the computation, then we have got the sum of gradients of the first four partitions. What we require now is the the sum of gradients of the last four partitions. So what we do is, we assign the task of computing the sum of the gradients of the last four partitions to the remaining $n_2 - n_1 = 3$ servers as well. As of now, our position is one server has completed its task and we have launched the remaining servers also. A total of 9 servers are busy computing their tasks (since one server has already completed its work). Since our recovery threshold is k = 5 for this particular example, we just need outputs from any 4 servers out of the remaining 9 servers working now. If any 4 servers complete their task, it is ensured that at least one of them is computing the sum of the gradients of

the last four partitions since, out of the remaining 9 servers, 3 servers compute the sum of the gradients of the first four partitions and 6 servers compute the sum of the gradients of the last four partitions.

With the server 1 results available which provides the sum of gradients of the first four partitions, any 4 of the remaining 9 servers will provide the sum of gradients of the last four partitions, thus giving the overall computation result. Here, each server is computing the gradients of either the first four partitions or the last four partitions. Each server performs a computation on 4 out of 8 partitions, and thus on 1/2 of the data. This is in contrast to each server performing computation on 6/10 of the data in case of the (n_2, k) gradient code.

C. Tiered Gradient Codes for $n_1 \ge 3(k-1)$

In this section, we construct tiered gradient codes in which a subset of servers under consideration will be allocated a cyclic repetition gradient code of suitable parameters. We will construct codes for all values of n_2 , with $n_1 - k + 1$ computations per server. We note that this is the best possible, since the gradient code restricted to first n_1 servers has to be an (n_1, k) gradient code and $n_1 - k + 1$ is the lower bound on the number of computations per server of an (n_1, k) gradient code [1]. We make the following observation with respect to the condition which the tiered gradient code has to necessarily satisfy. These will in turn be used to construct certain tiered gradient codes.

Lemma 2: Consider an $(n_1, n_2, k, c = 1)$ tiered gradient code. Suppose $Q = n_1$ and the support of the F matrix is picked as those given by the cyclic repetition gradient code. Let the m^{th} server finish its job in the first phase, i.e. $M = \{m\}$. Let L_m denote the support of \mathbf{f}_m , and Z_j denote the support of \mathbf{b}_j (\mathbf{b}_j is the j^{th} row of matrix B_M). Then, the following holds: $[0, n_1 - 1] \setminus L_m \subset Z_j, \forall j \in [n_2 - n_1]$.

Proof: Suppose not. Consider $r \in \{0, 1, \ldots, n_1 - 1\} \setminus L_m$ and $r \notin Z_j$ for some $j \in [n_2 - n_1]$. Based on the structure of the cyclic repetition gradient code, there are k-1 rows in the Fmatrix including m where the r^{th} column is zero. Considering these k - 1 rows and adding the j^{th} row of the B_M matrix, we have a set of k rows which is required to satisfy the span property. However, since the r^{th} coordinate is zero in all these rows, $1_{1 \times Q}$ cannot be in the span of these k rows. Hence, Lemma 2 is necessary for the code to be $(n_1, n_2, k, c = 1)$ gradient code.

We consider $Q = n_1$ and $C_{n_1} = \lfloor \frac{n_1 - k + 1}{k - 1} \rfloor$. Initially, the first n_1 servers are launched. We assume (without loss of generality) that the server 1 has finished the job in the first phase. Then the remaining $n_2 - n_1$ servers are launched. We will now construct codes for the case where $n_2 = n_1 + C_{n_1}$. Let B_M be a $C_{n_1} \times n_1$ matrix with \mathbf{b}_i representing the i^{th} row and Z_i representing the support of \mathbf{b}_i , where $i \in [C_{n_1}]$. Let the columns of the F and B_M matrices be indexed by $[0, n_1 - 1]$.

Construction 3 $(n_2 = n_1 + C_{n_1})$: The support structure of the matrix F is as follows:

$$supp(\mathbf{f_i}) = [i - 1, i + (n_1 - k - 1)] \mod n_1.$$

The procedure to design the support of each row of the B_M matrix is as follows.

If $C_{n_1} = 1$, the B_M matrix is a row matrix. The k - 1 coordinates of Z_1 are given by $[n_1 - k + 1, n_1 - 1] \subset Z_1$. We pick the remaining $n_1 - 2(k - 1)$ coordinates as a subset of $L_1 = [0, n - k]$ such that at least one of every pair of consecutive coordinates (modulo n_1) is present in the set.

If $C_{n_1} > 1$, do the following. Let $l = (n_1 - k + 1) - (k - 1)C_{n_1}$. The l + k - 1 coordinates of $Z_j, j \in [C_{n_1}]$, are given by $[n_1 - (l + k - 1), n_1 - 1] \subset Z_j, j = [C_{n_1}]$. Let

$$B_M = \begin{bmatrix} B_{M_1} & B_{M_2} & \dots & B_{M_{k-1}} & B_{M'} \end{bmatrix}$$

Each submatrix B_{M_j} , where $j \in [k-1]$, is of size $C_{n_1} \times C_{n_1}$ and the $B_{M'}$ matrix is of size $C_{n_1} \times (l+k-1)$. The $B_{M'}$ matrix constitutes the l+k-1 columns - $[n_1-(l+k-1), n_1-1]$ of the B_M matrix filled with non zero entries. The support structure for the remaining coordinates of the B_M matrix is obtained from the design of the support structure corresponding to the matrices B_{M_j} , $j \in [k-1]$. The support of the i^{th} row of each matrix B_{M_j} , where $j \in [k-1]$, is of the form $[i-1, i+C_{n_1}-3]$ mod C_{n_1} .

Now, we will present a randomized construction of the matrices F and B_M . The matrix H of size $(n_1 - k) \times n_1$ is picked at random as given in Section III-D. Each vector $\mathbf{f_i}$ is calculated by solving the following equation

$$\mathbf{f_i}|_{L_i} H(:, L_i)^T = 0,$$

 \mathbf{b}_i is calculated by solving the following equation

$$\mathbf{b}_i|_{Z_i} H(:, Z_i)^T = 0.$$

Proof: Initially, the first n_1 servers are launched. Without loss of generality, let us assume that the server 1 finishes the job first. Then the remaining $n_2 - n_1$ servers are launched. The procedure to design the support of each row of the B_M matrix is as follows. It is necessary that $|Z_i| = n_1 - k + 1$, for each $j \in [C_{n_1}]$. From Lemma 2, we have $[0, n_1 - 1] \setminus L_1 \subset$ $Z_j, \quad j = [C_{n_1}] \text{ and } |[0, n_1 - 1] \setminus L_1| = k - 1.$ Thus k - 1coordinates are included in each $Z_j, j = [C_{n_1}]$. We have to add exactly $|Z_i| - (k-1) = n_1 - 2(k-1) > 0$ coordinates from L_1 to the set to complete the specification of Z_j . We pick these $n_1 - 2(k-1)$ coordinates as a subset of $L_1 = [0, n_1 - k]$ such that at least one of every pair of consecutive coordinates $(modulo n_1)$ is present in the set. We will refer to this condition as consecutive coordinate property. It is possible to pick $n_1 - 2(k-1)$ coordinates satisfying the consecutive coordinate property only if $n_1 - 2(k-1) \ge \lfloor \frac{n_1 - k + 1}{2} \rfloor$. We can easily see that the above property is satisfied when $n_1 \ge 3(k-1)$. In addition, these are also picked so that $|Z_j \cup Z_i| = n_1$, for any $j, i \in [C_{n_1}]$.

The *l* coordinates from $L_1 - [(n_1 - k) - (l - 1), (n_1 - k)]$ are also included in Z_j . Thus, totally, l + k - 1 coordinates are included in each Z_j . We have to add $|Z_j| - (l + k - 1) =$ $n_1 - 2(k - 1) - l$ more coordinates to Z_j from $L_1 \setminus [(n_1 - k) - (l - 1), (n_1 - k)] = [0, n_1 - k - l]$. That is, we need to pick $n_1 - 2(k - 1) - l$ from $n_1 - k + 1 - l$ locations available.

Let $B_{M'}$ be the matrix obtained by taking the l + k - 1 coordinates corresponding to $[n_1 - (l + k - 1), n_1 - 1]$ from

each row in the B_M matrix, i.e., by taking the last l + k - 1columns $-[n_1 - (l + k - 1), (n_1 - 1)]$ from the B_M matrix. B_{M_1} constitutes the first C_{n_1} columns of the B_M matrix, B_{M_2} constitutes the next C_{n_1} columns and so on. Hence, each B_{M_j} , $j \in [k - 1]$, is a $C_{n_1} \times C_{n_1}$ matrix which is obtained by taking distinct and consecutive C_{n_1} columns from the B_M matrix sequentially. $n_1 - 2(k - 1) - l$ more coordinates to be added to $Z_j, j = [C_{n_1}]$ is obtained from the design of the support structure corresponding to the matrices $B_{M_j}, j \in [k-1]$. We have the support structure of the B_M matrix. The support structure for the remaining coordinates of the B_M matrix is obtained from the design of the support structure corresponding to the matrices $B_{M_j}, j \in [k-1]$.

The support of the i^{th} row of each matrix B_{M_j} , where $j \in [k-1]$, is of the form $[i-1, i+C_{n_1}-3] \mod C_{n_1}$. The cardinality of the support of each row of the B_{M_j} matrix is $C_{n_1}-1$, i.e, there is exactly one zero in each row of the B_{M_j} matrix at disjoint locations. Hence, the number of zeros in each row of the B_M matrix is exactly k-1, which is exactly what we needed. The cardinality of the support of union of any two rows of the B_{M_j} matrix is C_{n_1} . Hence if we take union of support of any two rows in the B_M matrix, then it has cardinality n_1 . That is, $|Z_r \cup Z_s| = n_1$, for any $r, s \in [C_{n_1}]$. Hence the support structure of the B_M matrix satisfies all the required conditions.

Theorem 5: The code given in Construction 3 is a $(n_1, n_2 = n_1 + C_{n_1}, k, c = 1)$ tiered gradient code where $n_1 \ge 3(k-1)$ and $C_{n_1} = \frac{n_1 - k + 1}{k-1}$.

Proof: We have to show that Condition 1 is satisfied by the code in Construction 3 with probability 1. Here $M = \{1\}$, assuming that the server 1 finished its task first. If \mathcal{T}_2 = ϕ , the above condition follows from the support structure of the cyclic repetition code. If $|\mathcal{T}_2| = 1$ and $M \in \mathcal{T}_1$, then the support of the union of \mathcal{T}_1 and \mathcal{T}_2 is $[n_1]$ and hence Condition 1 is satisfied. Now, we will consider the case when $|\mathcal{T}_2| = 1, M \notin \mathcal{T}_1, |\mathcal{T}_1| = k - 2$ and \mathcal{T}_1 is such that $|\cup_{i\in\mathcal{T}_1} L_i| = n_1 - k + (k-2) = n_1 - 2$. Based on the cyclic support structure of the F matrix, the above condition is true whenever (k-2) consecutive rows (modulo n_1) are picked. Hence, the two coordinates which are not included in the union are consecutive. Since the support of rows in B_M matrix satisfies consecutive coordinate property, at least one of the coordinates of the two coordinates which are not picked up before will be included after adding the new row. So we have, $|\bigcup_{i \in \mathcal{T}_1} L_i \bigcup_{i \in \mathcal{T}_2} Z_i| \ge (n_1 - 1)$. Hence Condition 1 is satisfied.

For the cases when $|T_2| \ge 2$, since Z_i and Z_j are chosen such that $Z_i \cup Z_j = [n_1]$, for any $i, j \in [C_{n_1}]$, we have that the condition 1 being trivially satisfied. Thus, Condition 1 is satisfied for all cases and hence the code is a $(n_1, n_2 = n_1 + C_{n_1}, k, c = 1)$ tiered gradient code.

Example 2: Let $n_1 = 9, k = 3, c = 1$ and $n_2 = 12$. We split data into 9 partitions $-\{x_0, x_1, \ldots, x_8\}$. The server *i* is assigned data $\{x_j, j \in [i - 1, i + 5]\}$. Each server computes the gradients on their respective data. We assume that server 1 finishes its task first when n_1 servers are launched. After that the remaining $n_2 - n_1 = 3$ servers are launched. Since server 1 doesn't have $\{x_7, x_8\}$ as its content, we have to

TABLE II

TABLE THAT ILLUSTRATES THE VALUE OF $C_{n_1}^*$ for any $k \ge p+4$ and $n_1 = 3(k-1)+p$, for Some Integer p. Let k' = p+4 and $n'_1 = 3(k'-1)+p$. If $k \ge p+4$ and $n_1 = 3(k-1)+p$, we can write k and n_1 in terms of k' and n'_1 as k = k' + p' and $n_1 = n'_1 + 3p'$, where $p' = \{0, 1, 2, \ldots\}$

0*			
C_{n_1}	All the possible cases when p is even	All the possible cases when p is odd	
2	p = 0	$p' > \frac{3(p+1)}{2}$	
	p' > 3P	-	
3	$\frac{3p}{2} \le p' \le 3p$	$\frac{3(p-1)}{2} \le p' \le \frac{3(p+1)}{2}$	
	$p' = 1 \mod 3$ and	$p' = 1 \mod 3$ and $3 \lfloor \frac{p-1}{4} \rfloor < p' \le \frac{3(p-1)}{2}$	
4	$\max\left\{0, 3\lfloor \frac{p}{4} \rfloor - 1\right\} < p' < \frac{3p}{2}$	$p'=2 \mod 3, \ 2 \le p' \le \lceil rac{3(p-1)}{2} \rceil$ and	
		$p' \neq 3\lfloor \frac{p-1}{4} \rfloor - 1$	
	$p' = 0, 2 \mod 3$ and $2 < p' \le \lceil \frac{p-1}{2} \rceil$	$p' = 2 \mod 3, p' = 7 \mod 8$ and	
		$\frac{3(p-7)}{4} < p' < \frac{3(p-1)}{2}$	
	$p' = 0, 2 \mod 3$ and $3 \lceil \frac{p}{4} \rceil - 1 < p' < \frac{3p}{2}$	$p' = 2 \mod 3, \ p' \neq 7 \mod 8$ and	
		$6\lfloor \frac{p}{8} \rfloor - 3 < p' < \frac{3(p-1)}{2}$	
		$p' = 1 \mod 3$ and $3 \lceil \frac{p-2}{6} \rceil < p' \le 3 \lfloor \frac{p-1}{4} \rfloor$	
5		$p'=2 \mod 3, \ 2 \le p' \le \lceil \frac{3(p-1)}{2} \rceil$ and	
	p' is a multiple of 3 and $\lceil \frac{p-1}{2} \rceil < p' < \frac{3p}{2}$	$p' = 3\lfloor \frac{p-1}{4} \rfloor - 1$	
		$p' = 0 \mod 3, p' = 7 \mod 8$ and	
		$0 < p' \le \frac{3(p-7)}{4}$	
		$0 < p' \le 6\lfloor \frac{\hat{p}}{8} \rfloor - 3$	
6	$p' = 2 \mod 3$ and $\lceil \frac{p-1}{2} \rceil < p' < \frac{3p}{2}$	$p' = 1 \mod 3$ and $0 < p' \le 3 \lceil \frac{p-2}{6} \rceil$	
	$p' = 0 \mod 3$ and $0 < p' < \max\left\{\overline{0}, 3\lfloor \frac{p}{4} ight\} - 1 ight\}$		
0	For all other cases not discussed above		

include $\{x_7, x_8\}$ in the content of three added servers. Here, l = 1. Hence, x_6 needs to be included in the content of three added servers. The last three columns of the B_M matrix are filled with non zero entries. Let $B_M = [B_{M_1} B_{M_2} B_{M'}]$. The $B_{M'}$ matrix is obtained by taking the last three columns of B_M . Hence it is a 3×3 matrix. Both the B_{M_1} and B_{M_2} matrices are 3×3 matrices. B_{M_1} is the submatrix formed by the first three columns of the B_M matrix and B_{M_2} is formed by the next three columns of the B_M matrix. The support of the i^{th} row of each of the matrices B_{M_1} and B_{M_2} is of the form $\{i-1,i\}$ mod 3. Hence the structures of B_{M_1}, B_{M_2} and B_M are of the form

$$B_{M_1} = B_{M_2} = \begin{bmatrix} * & 0 & * \\ * & * & 0 \\ 0 & * & * \end{bmatrix} \text{ and } \\B_M = \begin{bmatrix} * & 0 & * & * & 0 & * & * & * \\ * & * & 0 & * & * & 0 & * & * & * \\ 0 & * & * & 0 & * & * & * & * & * \end{bmatrix}.$$

The symbol * in the above matrices implies non zero entries in those locations. The content of the three added servers are $\{x_0, x_2, x_3, x_5, x_6, x_7, x_8\}$, $\{x_0, x_1, x_3, x_4, x_6, x_7, x_8\}$ and $\{x_1, x_2, x_4, x_5, x_6, x_7, x_8\}$ respectively. Out of the 8 servers which haven't finished the job earlier and the three added servers, any two servers can give the sum of the gradients along with server 1. Each server does 7/9 computations compared to 10/12 required for the (n_2, k) gradient code.

Remark 1: If $n_1 < n_2 < n_1 + C_{n_1}$, we take the support structure of any $n_2 - n_1$ rows of the B_M matrix constructed using Construction 3 $(n_2 = n_1 + C_{n_1})$ to generate the

support structure for the B_M matrix in this case. The support structure of the matrix F, the construction of the B_M and F matrices using the above support structures are same as in Construction 3.

D. Tiered Gradient Code for $k \ge 4, n_1 \in [3(k-1), 3(k-1) + (k-4)]$

In this subsection, we consider the case where $k \ge 4$ and $n_1 \in [3(k-1), 3(k-1) + (k-4)]$. For such cases we provide construction for $n_2 = n_1 + C_{n_1}^*$, where $C_{n_1}^* \ge C_{n_1}$. The value of $C_{n_1}^*$ is given in Table II.

For any integer p, if $k \ge p + 4$ and $n_1 = 3(k-1) + p$, the code construction is provided below. Let k' = p + 4 and $n'_1 = 3(k'-1) + p$, for some integer p. We can write k and n_1 in terms of k' and n'_1 as k = k' + p' and $n_1 = n'_1 + 3p'$, where $p' = \{0, 1, 2, \ldots\}$. The value $C^*_{n_1}$ varies from 2 to 6 depending upon k and n_1 , which is given in Table II. Let $(0 * *)^y$ represent the sequence $\{0 * *\}$ repeated y times, i.e,

$$(0**)^y = \begin{bmatrix} 0 & * & * & 0 & * & * & \dots & 0 & * & * \end{bmatrix}_{1 \times 3y}$$
(4)

where * represents some non zero entry. Similarly let $(**0)^y$ represent the sequence $\{**0\}$, $(0*)^y$ represent the sequence $\{0*\}$, $(*0)^y$ represent the sequence $\{*0\}$ and $(*)^y$ represent the sequence $\{*\}$ repeated y times. We will now construct codes for the case where $n_2 = n_1 + C_{n_1}^*$. Let B_M is a $C_{n_1}^* \times n_1$ matrix with \mathbf{b}_i representing the i^{th} row and Z_i representing the support of \mathbf{b}_i , where $i \in [C_{n_1}^*]$.

TABLE III
FABLE THAT ILLUSTRATES THE SUPPORT STRUCTURE OF THE B_M MATRIX ASSOCIATED WITH CONSTRUCTION 4

$C_{n_1}^*$	p'	The support structure of the B_M matrix if p is even	The support structure of the B_M matrix if p is odd
3	any p'	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ *(0*)^{\lceil \frac{n_1}{2}\rceil+1}(0**)^{p-1}(0*)^{\lfloor \frac{n_1}{2}\rfloor+2}(*)^{k-1} \end{bmatrix}$	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ *(0*)^{\lceil \frac{n_1}{2}\rceil+1}(0**)^{p-1}(0*)^{\lfloor \frac{n_1}{2}\rfloor+2}(*)^{k-1} \end{bmatrix}$
4	$1,2 \mod 3$	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ (0**)^p(0*)^{n_1+2}(*)^{k-1} \\ (*0)^{n_1+2}(**0)^p(*)^{k-1} \end{bmatrix}$	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ (0**)^p(0*)^{n_1+2}(*)^{k-1} \\ (*0)^{n_1+2}(**0)^p(*)^{k-1} \end{bmatrix}$
4	0 mod 3	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ 0*(**0)^{\frac{p}{2}-1}(*0)^{n_1+1}*(0**)^{\frac{p}{2}}*0(*)^{k-1} \\ 0*(**0)^{\frac{p}{2}}*(0*)^{n_1+1}(0**)^{\frac{p}{2}-1}*0(*)^{k-1} \end{bmatrix} $ if $p' \in [3, \lceil \frac{p-1}{2}\rceil]$	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ 0*(**0)^{p-1}(*0)^{n_1+2}*(*)^{k-1} \\ *(0*)^{n_1+2}(0**)^{p-1}*0(*)^{k-1} \end{bmatrix}$
5	1,2 mod 3	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ 0*(**0)^{\frac{p}{2}-1}(*0)^{n_1+1}*(0**)^{\frac{p}{2}}*0(*)^{k-1} \\ 0*(**0)^{\frac{p}{2}}*(0*)^{n_1+1}(0**)^{\frac{p}{2}-1}*0(*)^{k-1} \\ *(0*)^{\lceil \frac{n_1}{2}\rceil+1}(0**)^{p-1}(0*)^{\lfloor \frac{n_1}{2}\rfloor+2}(*)^{k-1} \end{bmatrix}$	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ (0**)^p(0*)^{n_1+2}(*)^{k-1} \\ (*0)^{n_1+2}(**0)^p(*)^{k-1} \\ 0*(**0)^{\frac{p-1}{2}}(*0)^{n_1+1}*(0**)^{\frac{p-1}{2}}*0(*)^{k-1} \end{bmatrix}$
5	0 mod 3	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ 0*(**0)^{\frac{p}{2}-1}(*0)^{n_1+1}*(0**)^{\frac{p}{2}}*0(*)^{k-1} \\ 0*(**0)^{\frac{p}{2}}*(0*)^{n_1+1}(0**)^{\frac{p}{2}-1}*0(*)^{k-1} \\ *(0*)^{\lceil\frac{n_1}{2}\rceil+1}(0**)^{p-1}(0*)^{\lfloor\frac{n_1}{2}\rfloor+2}(*)^{k-1} \end{bmatrix}$	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ *(0*)^{\lceil \frac{n_1}{2}\rceil+1}(0**)^{p-1}(0*)^{\lfloor \frac{n_1}{2}\rfloor+2}(*)^{k-1} \\ 0*(**0)^{u}(*0)^{n_1+1}1(0**)^{p-1-u}*0(*)^{k-1} \\ 0*(**0)^{p-1-u}(*0)^{n_1+1}1(0**)^{u}*0(*)^{k-1} \end{bmatrix},$
			where $u = \begin{cases} \frac{p'-7}{3} & \text{if } p' = 1 \mod 6\\ \frac{p'-3}{3} & \text{if } p' = 3 \mod 6\\ \frac{p'-5}{3} & \text{if } p' = 5 \mod 6 \end{cases}$
6	any p'	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ (0**)^p(0*)^{n_1+2}(*)^{k-1} \\ (*0)^{n_1+2}(**0)^p(*)^{k-1} \\ 0*(**0)^{\frac{p}{2}-1}(*0)^{n_1+1}*(0**)^{\frac{p}{2}}*0(*)^{k-1} \\ 0*(**0)^{\frac{p}{2}}*(0*)^{n_1+1}(0**)^{\frac{p}{2}-1}*0(*)^{k-1} \end{bmatrix}$	$\begin{bmatrix} (0**)^{p+1}(0*)^{n_1}0(*)^{k-1} \\ 0(*0)^{n_1}(**0)^{p+1}(*)^{k-1} \\ (0**)^p(0*)^{n_1+2}(*)^{k-1} \\ (*0)^{n_1+2}(**0)^p(*)^{k-1} \\ 0*(**0)^u(*0)^{n_1+1}*(0**)^{p-1-u}*0(*)^{k-1} \\ 0*(**0)^{p-1-u}(*0)^{n_1+1}*(0**)^u*0(*)^{k-1} \end{bmatrix},$ where $u = \begin{cases} \frac{p'+2}{3} & \text{if } p' = 1 \mod 6 \\ \frac{p'-3}{3} & \text{if } p' = 3 \mod 6 \end{cases}$
			$\int \frac{p'-2}{3} \text{if } p' = 5 \mod 6$

Construction 4 $(n_2 = n_1 + C_{n_1}^*)$: The support structure of the matrix F is as follows:

$$supp(\mathbf{f_i}) = [i - 1, i + (n_1 - k - 1)] \mod n_1.$$

If $C_{n_1}^* = 2$, the support structure of the B_M matrix is same as in Construction 3. The support structure of the B_M matrix for all other values of $C_{n_1}^*$ is given in Table III. The construction of the B_M and F matrices using the above support structures are same as in Construction 3.

Discussion on Construction 4: The procedure to design the support of each row of the B_M matrix is as follows. k-1 coordinates of Z_j are given by $[n_1 - k + 1, n_1 - 1] \subset Z_j, j \in C_{n_1}^*$, which corresponds to the partitions that are not included in server 1.

We pick the remaining $n_1 - 2(k-1)$ coordinates as a subset of $L_j = [0, n-k]$ such that at least one of every pair of consecutive coordinates is present in the set. This is termed as the consecutive coordinate property. And also, the support structure of the B_M matrix is designed in such a way that for every \mathcal{T}_2 , which is a subset of $C^*_{n_1}$ of size ℓ $(1 \le \ell \le k - 1)$, we have $|\bigcup_{i \in \mathcal{T}_2} Z_i| \ge (n_2 - k) + \ell$.

Theorem 6: The code given in Construction 4 is a $(n_1, n_2 = n_1 + C^*_{n_1}, k, c = 1)$ tiered gradient code where $k \ge 4, n_1 \in [3(k-1), 3(k-1) + (k-4)].$

Proof: We have to show that Condition 1 is satisfied by the code in Construction 4. Here $M = \{1\}$, assuming that server 1 finished its task first. If $\mathcal{T}_2 = \phi$, Condition 1 follows from the support structure of cyclic repetition code. If $|\mathcal{T}_2| = 1$ and $M \in \mathcal{T}_1$, then the support of the union of \mathcal{T}_1 and \mathcal{T}_2 is $[n_1]$ and hence Condition 1 is satisfied. Now, we will consider the case when $|\mathcal{T}_2| = 1$, $M \notin \mathcal{T}_1$, $|\mathcal{T}_1| = k - 2$ and \mathcal{T}_1 is such that $|\bigcup_{i \in \mathcal{T}_1} L_i| = n_1 - k + (k - 2) = n_1 - 2$. Based on the cyclic support structure of the F matrix, Condition 1 is true whenever (k - 2) consecutive rows (modulo n_1) are picked. Hence, the two coordinates which are not included in the union are consecutive. Since the support of the rows in the B_M matrix satisfies consecutive coordinate property, at least one of the coordinates of the two coordinates which are not picked up before will be included after adding the new row. So we have, $|\bigcup_{i \in T_1} L_i \bigcup_{j \in T_2} Z_j| \ge (n_1 - 1)$. Hence Condition 1 is satisfied. The support structure of the B_M matrix is designed in such a way that Condition 1 is satisfied. Hence, for the cases when $|T_2| \ge 2$, Condition 1 is satisfied.

Thus, Condition 1 is satisfied for all cases and hence the code is a $(n_1, n_2 = n_1 + C_{n_1}^*, k, c = 1)$ tiered gradient code.

Construction 5: (General $n_1, n_2 \ge 3(k-1)$). If $n_1, n_2 \ge 3(k-1)$,

1) we need to find a p such that the following condition is satisfied, i.e., for $C_{n_1} = \lfloor \frac{n_1 - k + 1}{k - 1} \rfloor$,

$$n_{2} \leq n_{1} + p + C_{n_{1}+p} = n_{1} + p + \frac{n_{1} + p - k + 1}{k - 1}$$
$$p \geq n_{2} - n_{1} - \frac{n_{2} - k + 1}{k}.$$

Hence the minimum value of p possible is $p^* = \lceil n_2 - n_1 - \frac{n_2 - k + 1}{k} \rceil$. 2) we need to find $n^+ =$

- 2) we need to find $n^+ = \min\{n' \in [\max\{n_1, n_2 6\}, n_2 1]\}$ such that $C^*_{n^+} = \max_{n' \in [\max\{n_1, n_2 6\}, n_2 1]} C^*_{n'}$.
- 3) we also need to find $n_{min} = \min\{n'' \in [\max\{n_1, n_2 6\}, n_2 1]\}$ such that $n_2 \le n'' + C^*_{n''}$.

Our objective is to maximize G_1 in Theorem 1. Considering point 1 alone, which is mentioned above, the maximum G_1 possible is $G_1^1 = \min\{n_2 - (n_1 + p^*), C_{n_1+p^*}\}$. If we consider the point 2, the maximum G_1 possible is $G_1^2 =$ $\min\{n_2 - n^+, C_{n^+}^*\}$. Similarly, if we consider the point 3, the maximum G_1 possible is $G_1^3 = n_2 - n_{min}$. So summarizing all the three points, the maximum G_1 is $\max\{G_1^1, G_1^2, G_1^3\}$. If $G_1 = G_1^1$, let $n'_1 = n_1 + p^*$, else if $G_1 = G_1^2$, let $n'_1 = n^+$, else, let $n'_1 = n_{min}$. For $n_2 \ge 3(k-1)$, we can use Construction 3 where we replace n_1 with n'_1 servers. Initially, we launch n_1 of n'_1 servers. After one of those servers finish their tasks, the remaining $n_2 - n_1$ servers are launched.

Remark 2: In this case, for general n_2 , the amount of computation per server required is $\frac{n_2-k+1-G_1}{n_2-G_1}$, where $G_1 = \max\{\min\{n_2 - (n_1 + p^*), \lfloor \frac{n_1+p^*-k+1}{k-1} \rfloor\}, \min\{n_2 - n^+, C_{n^+}^*\}, n_2 - n_{min}\}, p^* = \lceil n_2 - n_1 - \frac{n_2-k+1}{k} \rceil, n^+ = \min\{n' \in \lceil \max\{n_1, n_2 - 6\}, n_2 - 1 \rceil\}$ such that $C_{n^+}^* = \max_{n' \in \lceil \max\{n_1, n_2 - 6\}, n_2 - 1 \rceil} C_{n'}^*, n_{min} = \min\{n'' \in \lceil \max\{n_1, n_2 - 6\}, n_2 - 1 \rceil\}$ such that $n_2 < n'' + m_2 < n_2 < n_2$

 $\min \{n'' \in [\max \{n_1, n_2 - 6\}, n_2 - 1]\} \text{ such that } n_2 < n'' + C^*_{n''}. \text{ This proves Theorem 1 for } n_1, n_2 \ge 3(k - 1).$

E. Tiered Gradient Codes for $2(k-1) < n_1 < 3(k-1), n_2 \ge 3(k-1)$

In this subsection, we provide tiered gradient codes for the case where $2(k-1) < n_1 < 3(k-1), n_2 \ge 3(k-1)$. The

construction is in similar lines to that of Construction 5, except that,

- for the point 1 mentioned in Construction 5, the 'p' should satisfy one more condition, i.e., $n_1 + p \ge 3(k-1)$. So the minimum value of p possible is $p^* = \max\{3(k-1) n_1, \lceil n_2 n_1 \frac{n_2 k + 1}{k} \rceil\}$.
- we have to replace n_1 with 3(k-1) in point 2 of Construction 5, i.e., we need to find $n^+ = \min\{n' \in [\max\{3(k-1), n_2 - 6\}, n_2 - 1]\}$ such that $C_{n^+}^* = \max_{n' \in [\max\{3(k-1), n_2 - 6\}, n_2 - 1]} C_{n'}^*$.
- for point 3 of Construction 5 also, we need to replace n_1 with 3(k-1), i.e., we also need to find $n_{min} = \min\{n'' \in [\max\{3(k-1), n_2 6\}, n_2 1]\}$ such that $n_2 \leq n'' + C_{n''}^*$.

Our aim is to maximize G_2 in Theorem 1. So considering all the points mentioned above, the maximum G_2 is $\max \{G_2^1, G_2^2, G_2^3\}$, where $G_2^1 = \min \{n_2 - (n_1 + p^*), C_{n_1+p^*}\}, G_2^2 = \min \{n_2 - n^+, C_{n^+}^*\}$ and $G_2^3 = n_2 - n_{min}$. If $G_2 = G_2^1$, let $n'_1 = n_1 + p^*$, else if $G_2 = G_2^2$, let $n'_1 = n^+$, else, let $n'_1 = n_{min}$. For $n_2 \ge 3(k-1)$, we can use Construction 3 where we replace n_1 with n'_1 servers. Initially, we launch n_1 of n'_1 servers. After one of those servers finish their tasks, the remaining $n_2 - n_1$ servers are launched.

VIII. TIERED GRADIENT CODES FOR c > 1

In this section, we deal with the case where initially we launch the first n_1 servers and wait for c > 1 servers to complete their tasks. After that the remaining $n_2 - n_1$ servers are launched.

This section is organized as follows. Initially, we discuss about tiered gradient codes for $k \leq n_1 \leq 2(k-1)$ and $n_2 > 2(k-1)$. Then we move on to $n_1, n_2 \geq 2(k-1)+(k-c)$. Towards the end, we provide a discussion on codes for $2(k-1) < n_1 < 2(k-1)+(k-c)$ and $n_2 \geq 2(k-1)+(k-c)$. For $k \leq n_1 \leq 2(k-1)$ and $n_2 > 2(k-1)$, we use Construction 2, where we wait for c servers to complete their tasks instead of one server.

For any specific $n_1 > 2(k-1)$, we use the unique cyclic repetition gradient code for the first n_1 servers. We need to show that the Condition 1 holds for the codes under consideration. Let $M = \{i_1, i_2, \ldots, i_c\}$. If $\mathcal{T}_2 = \phi$, Condition 1 follows from the support structure of cyclic repetition code. If $|\mathcal{T}_2| = 1$ and some subset of M is included in \mathcal{T}_1 , the support of the union is $[n_1]$ and hence Condition 1 is satisfied. Consider the case where $|\mathcal{T}_2| = 1, M \notin \mathcal{T}_1, |\mathcal{T}_1| = k - c - 1$ and \mathcal{T}_1 is such that $|\bigcup_{i \in \mathcal{T}_1} L_i| = n_1 - k + (k - c - 1) = n_1 - c - 1$. Based on the cyclic support structure of the F matrix, Condition 1 is true whenever (k - c - 1) consecutive rows (modulo n_1) are picked. Hence, the c + 1 coordinates which are not included in the union are consecutive. So, each row in the B_M matrix should be designed in such a way that at least one of every c + 1 consecutive coordinates should be non zero.

Now, we will construct codes where at least one of every two consecutive coordinates is non zero in the B_M matrix. We consider $Q = n_1$. Let $C'_{n_1} = \lfloor \frac{n_1 - k + c}{k - 1} \rfloor$. Initially, we launch n_1 servers. Let \mathbf{f}_i represent the i^{th} row of the Fmatrix and L_i represent the support of \mathbf{f}_i , where $i \in [n_1]$. Let \mathbf{b}_i represent the i^{th} row of the B_M matrix and Z_i represent the support of \mathbf{b}_i , where $i \in [C'_{n_1}]$. Let the columns of the F and B_M matrices be indexed by $[0, n_1 - 1]$. Let $\{i_1, i_2, \ldots, i_c\}$ be the c servers who complete their tasks first. Then, $\{L_{i_1}, L_{i_2}, \ldots, L_{i_c}\}$ is the support of $\{\mathbf{f}_{i_1}, \mathbf{f}_{i_2}, \ldots, \mathbf{f}_{i_c}\}$ respectively.

The code construction for $n_1 \ge 2(k-1) + (k-c)$ and $n_2 = n_1 + C'_{n_1}$ is as follows.

Construction 6: $(n_1 \ge 2(k-1) + (k-c), n_2 = n_1 + C'_{n_1})$ The support structure of the matrix F is as follows: $supp(\mathbf{f_i}) = [i-1, i+(n_1-k-1)] \mod n_1$. For some $t \in [0, n_1-1]$, let

$$n_1 - |\cup_{i \in \{i_1, i_2, \dots, i_c\}} L_i| = g,$$

$$[0, n_1 - 1] \setminus \bigcup_{i \in \{i_1, i_2, \dots, i_c\}} L_i = [t, t + g - 1],$$

$$l = (n_1 - k + c) - (k - 1)C'_{n_1}.$$

The l+k-c elements of Z_j , for each $j \in [C'_{n_1}]$, are given by $[t-l,t+k+c-1] \subset Z_j, \quad j = [C'_{n_1}].$

If $C'_{n_1} = 1$, pick the remaining coordinates so that consecutive coordinate property is satisfied, i.e, at least one coordinate from every possible pair of two consecutive coordinates are picked up. Else if $C'_{n_1} > 1$, do the following. Let B^*_M be the matrix obtained by shifting all the columns (say, y number of shifts done to each column towards right) in the B_M matrix in such a way that the columns -[t - l, t + k - c - 1] of the B_M matrix become the last l + k - c columns in B^*_M , $B^*_M = [B_{M_1} B_{M_2} \dots B_{M_{k-1}} B_{M'}]$. Let Z^*_i represent the support of the i^{th} row of the B^*_M matrix. $B_{M'}$ is the matrix obtained by taking the last l + k - c columns of the B^*_M matrix. All the entries in the $B_{M'}$ matrix are non zero, i.e., $[n_1 - (l + k - c), n_1 - 1] \subset Z^*_j, j \in [C'_{n_1}]$.

 B_{M_1} constitutes of the first C'_{n_1} columns of the B^*_M matrix, $B^*_{M_2}$ constitutes of the next C'_{n_1} columns and so on. Each B_{M_j} is a $C'_{n_1} \times C'_{n_1}$ matrix which is obtained by taking distinct and consecutive C'_{n_1} columns from the B^*_M matrix sequentially.

The support of the i^{th} row of each matrix B_{M_j} , where $j \in [k-1]$, is of the form $[i-1, i+C'_{n_1}-3] \mod C'_{n_1}$. The support structure of the B_M matrix is same as that of the B^*_M matrix with each column of the B^*_M matrix shifted towards left by y. The construction of the B_M and F matrices using the above support structures is same as in Construction 3.

Proof: From Lemma 2, we have $[0, n_1 - 1] \setminus \bigcup_{i \in \{i_1, i_2, \dots, i_c\}} L_i \subset Z_j$, $j = [C'_{n_1}]$. Thus g coordinates are included in each Z_j , for each $j \in [C'_{n_1}]$. We have to add $|Z_j| - g = n_1 - k + 1 - g$ more coordinates to Z_j from $\bigcup_{i \in \{i_1, i_2, \dots, i_c\}} L_i$. That is, we need to pick $n_1 - k + 1 - g$

from $n_1 - g$ locations available. Hence for the consecutive coordinate property to be satisfied, $n_1 - k + 1 - g \ge \lfloor \frac{n_1 - g}{2} \rfloor$, i.e., $n_1 \ge 2(k - 1) + g$. The maximum value that g can take is k - c, which is basically when all the c servers who finish first are consecutive ones. Considering the worst case scenario, the consecutive coordinate property is satisfied when $n_1 \ge 2(k - 1) + (k - c)$, which is our range of n_1 for which the code is constructed.

The procedure to design the support of each row of the B_M matrix is as follows. It is required that $|Z_j| = n_1 - k + 1$, for each $j \in [C'_{n_1}]$. No two servers among the first n_1 servers can have disjoint data set. It comes from the fact that $2(n_1 - k + 1) = 2n_1 - 2(k - 1) > 2n_1 - n_1 = n_1$. The inequality in the second step is satisfied since $n_1 \ge 2(k-1) + (k-c)$. Hence the set $[0, n_1 - 1] \setminus \bigcup_{i \in \{i_1, i_2, \dots, i_c\}} L_i$ contains consecutive g coordinates.

Pick any consecutive k-c coordinates from $[0, n_1-1]$ which includes the above g coordinates. Let it be $\{t, t+1, \ldots, t+k-c-1\}$. Let $[t, t+k-c-1] \subset Z_j$, $j = [C'_{n_1}]$. Thus k-c coordinates are included in each Z_j , $j = [C'_{n_1}]$. Rest of the coordinates of Z_j , are picked to satisfy the consecutive coordinate property. And also, these are picked so that $|Z_j \cup Z_i| = n_1$, for any $j, i \in [C'_{n_1}]$. The l coordinates - [t-l, t-1]are also included in Z_j . Thus, totally, l+k-c coordinates are included in each Z_j . We have to add $|Z_j| - (l+k-c) = n_1 - 2k + 1 + c - l$ more coordinates to Z_j from $[0, n_1 - 1] \setminus [t-l, t+k-c-1]$. That is, we need to pick $n_1 - 2k + 1 + c - l$ from $n_1 - k + c - l$ locations available.

If $C'_{n_1} = 1$, we pick the remaining coordinates so that consecutive coordinate property is satisfied. Else if $C'_{n_1} > 1$, we do the following. Let B^*_M be the matrix obtained by shifting all the columns in the B_M matrix in such a way that the columns -[t-l, t+k-c-1] in the B_M matrix become the last l+k-c columns in the B^*_M matrix. Let y be the number of shift done to each column of the B_M matrix towards right to obtain the B^*_M matrix. Let Z^*_i represent the support of the i^{th} row of the B^*_M matrix.

Let $B_{M'}$ be the matrix obtained by taking the last l + k - ccolumns of the B_M^* matrix. All the entries in the $B_{M'}$ matrix are non zero, i.e., $[n_1 - (l + k - c), n_1 - 1] \subset Z_j^*, j \in [C'_{n_1}]$. B_{M_1} constitutes of the first C'_{n_1} columns of the B_M^* matrix, $B_{M_2}^*$ constitutes of the next C'_{n_1} columns and so on. Hence, each B_{M_j} is a $C'_{n_1} \times C'_{n_1}$ matrix which is obtained by taking distinct and consecutive C'_{n_1} columns from the B_M^* matrix. We have the support structure of the $B_{M'}$ matrix. The support structure for the remaining coordinates of the B_M^* matrix, i.e., $n_1 - 2k + c + 1 - l$ more coordinates to be added to Z_j^* , is obtained from the design of the support structure corresponding to the matrices $B_{M_j}, j \in [k-1]$.

The support of the i^{th} row of each matrix B_{M_j} , where $j \in [k-1]$, is of the form $[i-1, i+C'_{n_1}-3] \mod C'_{n_1}$. The cardinality of the support of each row of the B_{M_j} matrix is $C'_{n_1}-1$, i.e, there is exactly one zero in each row of the B_{M_j} matrix at disjoint locations. Hence, the number of zeros in each row of the B^*_M matrix is exactly k-1, which is exactly what we needed.

The cardinality of the support of union of any two rows of the B_{M_i} matrix is C'_{n_1} . Hence if we take union of support of any two rows in the B_M^* matrix, then it has cardinality n_1 . That is, $|Z_r \cup Z_s| = n_1$, for any $r, s \in [C'_{n_1}]$. Since the B_M^* matrix is obtained by column shift of the B_M matrix, the above property holds for the B_M matrix also. Hence the support structure of the B_M matrix satisfies all the required conditions.

Example 3: Let $n_1 = 9, k = 4, c = 2$ and $n_2 = 11$. We split data into 9 partitions $-\{x_0, x_1, \ldots, x_8\}$. The server *i* is assigned partitions $\{x_j, j \in [i-1, i+4]\}$. Each server computes the gradients on their respective data. Suppose server 1 and 3 finish their tasks first when n_1 servers are launched. After that the remaining $n_2 - n_1 = 2$ servers are launched. Server 1 or 3 do not have $\{x_8\}$ as their contents. Here, k - c = 2 and l = 1. Hence we have to include $\{x_7, x_8, x_0\}$ in the content of the two added servers. The first column and the last two columns of the B_M matrix is filled with non zero entries. Shift each column of the B_M matrix by 8 units towards right to obtain the B_M^* matrix, $B_M^* = \begin{bmatrix} B_{M_1} & B_{M_2} & B_{M_3} & B_{M'} \end{bmatrix}$. The $B_{M'}$ matrix is obtained by taking the last three columns of the B_M^* matrix. Hence it is a 2 × 3 matrix. B_{M_1}, B_{M_2} and B_{M_3} are 3 × 3 matrices. B_{M_1} is the submatrix formed by the first columns of the B_M^* matrix, B_{M_2} is formed by the next three columns of the B_M^* matrix and B_{M_3} by the next three columns. The support of the i^{th} row of each of the matrices B_{M_1}, B_{M_2} and B_{M_3} is of the form $\{i-1\}$ mod 2. Hence the structures of the $B_{M_1}, B_{M_2}, B_{M_3}$ and B_M^* matrices are of the form

$$B_{M_1} = B_{M_2} = B_{M_3} = \begin{bmatrix} 0 & * \\ * & 0 \end{bmatrix} \text{ and}$$
$$B_M^* = \begin{bmatrix} 0 & * & 0 & * & 0 & * & * & * \\ * & 0 & * & 0 & * & 0 & * & * & * \\ \end{bmatrix}.$$

The entries * represent non zero values. Shift each column, in the B_M^* matrix, 8 units towards left to obtain the support structure of the B_M matrix,

$$B_M = \begin{bmatrix} * & 0 & * & 0 & * & 0 & * & * \\ * & * & 0 & * & 0 & * & 0 & * & * \end{bmatrix}.$$

The content of the two added servers are $\{x_0, x_2, x_4, x_6, x_7, x_8\}$ and $\{x_0, x_1, x_3, x_5, x_7, x_8\}$ respectively. Out of the 8 servers which haven't finished the job earlier and the two added servers, any three servers can give the sum of the gradients along with server 1 and 3. Each server does 6/9 computations compared to 8/11 required for the (n_2, k) gradient code.

For $n_1 \ge 2(k-1) + (k-c)$, $n_2 < n_1 + C_{n_1}$, we take the support structure of any $n_2 - n_1$ rows of the B_M matrix constructed using Construction 6 $(n_2 = n_1 + C'_{n_1})$ to generate the support structure for the B_M matrix in this case. The support structure of the matrix F, the construction of the B_M and F matrices using the above support structures is same as in Construction 3.

Example 4: Consider Example 3 with $n_1 = 9, k = 4$, c = 2. Consider $n_2 = 10$. Here, $C'_{n_1} = 2$. Hence, $n_2 - n_1 = 1 < 2$. The setting is same as in Example 3. The only difference is that $n_2 - n_1 < C'_{n_1}$. Hence we can use any one row of the B_M matrix from Example 3 to generate the B_M matrix for this case. Let us take the first row. Hence,

 $B_M = [* \ 0 \ * \ * \ 0 \ * \ * \ *]$, where the symbol * represent non zeros entries. Hence the content of the one added server is $\{x_0, x_2, x_3, x_5, x_6, x_7, x_8\}$. Out of the 8 servers which haven't finished the job earlier and the one added servers, any three servers can give the sum of the gradients along with server 1. Each server does 7/9 computations compared to 8/10 required for the (n_2, k) gradient code.

For $n_1 \ge 2(k-1) + (k-c)$, $n_2 > n_1 + C'_{n_1}$, we need to find a p such that the following condition holds, i.e.,

$$n_{2} \leq n_{1} + p + C'_{n_{1}+p} = n_{1} + p + \frac{n_{1} + p - k + c}{k - 1}$$
$$p \geq n_{2} - n_{1} - \frac{n_{2} - k + c}{k}.$$

Hence the minimum value of p possible is $p^* = \lceil n_2 - n_1 - \frac{n_2 - k + c}{k} \rceil$. So, for $n_2 > n_1 + C'_{n_1}$, we can use Construction 6 where we replace n_1 with $n_1 + p^*$ servers. Initially, we launch n_1 of $n_1 + p^*$ servers. After c of those servers finish their tasks, the remaining $n_2 - n_1$ servers are lauched.

Example 5: Consider $n_1 = 9, k = 4, c = 2$ and $n_2 = 12$. Here, $C'_{n_1} = 2$. Hence $n_2 > n_1 + C_{n_1}$. Here, $p^* = 1$ and $C_{n_1+p^*} = 2$. So, we split data into 10 partitions - $\{x_0, x_1, \ldots, x_9\}$. The server *i* is assigned data $\{x_j, j \in [i - 1, i+5]$. Each server computes the gradients on their respective data.

Initially the first n_1 servers are launched. Suppose server 2 and 4 finish their tasks first. After that all the remaining servers are launched. Since server 2 or 4 doesn't have $\{x_1\}$ as their content, k-c = 2 and $l = n_1 + p^* - k + 1 - (k-1)C_{n_1+p^*} = 2$, we have to include $\{x_0, x_9, x_1, x_2\}$ in the content of server 11 and 12. The last three and the first column of the B_M matrix are filled with non zero entries. Shift each column of the B_M matrix by 9 units towards right to obtain B_M^* matrix. $B_M^* = |B_{M_1} B_{M_2} B_{M_3} B_{M'}|$. $B_{M'}$ is a 2×3 matrix obtained by taking the last four columns of B_M^* . B_{M_1}, B_{M_2} and B_{M_3} are 2×2 matrices. B_{M_1} is the submatrix formed by the first two columns of the B_M^* matrix, B_{M_2} is formed by the next two columns of the B_M^* matrix and B_{M_3} by the next two columns. The support of the i^{th} row of each of the matrices B_{M_1}, B_{M_2} and B_{M_3} is of the form $\{i-1\} \mod 2$. Hence the structures of the $B_{M_1}, B_{M_2}, B_{M_3}$ and B_M^* matrices are of the form

$$B_{M_1} = B_{M_2} = B_{M_3} = \begin{bmatrix} 0 & * \\ * & 0 \end{bmatrix} \text{ and}$$
$$B_M^* = \begin{bmatrix} 0 & * & 0 & * & 0 & * & * & * & * \\ * & 0 & * & 0 & * & 0 & * & * & * & * \\ \end{bmatrix}.$$

The entries * represent non zero values. Shift each column, in the B_M^* matrix, 9 units towards left to obtain the support structure of the B_M matrix,

$$B_M = \begin{bmatrix} * & * & * & 0 & * & 0 & * & 0 & * & * \\ * & * & * & * & 0 & * & 0 & * & 0 & * \end{bmatrix}$$

The content of the two added servers are $\{x_0, x_1, x_2, x_4, x_6, x_8, x_9\}$ and $\{x_0, x_1, x_2, x_3, x_5, x_7, x_9\}$ respectively. Out of the 9 servers which haven't finished the job earlier and the two added servers, any three servers can give the sum of the gradients along with servers 2 and 4. Each

server does 7/10 computations compared to 9/12 required for the (n_2, k) gradient code.

Theorem 7: The code given in Construction 6 is a tiered gradient code where $n_1, n_2 \ge 2(k-1) + (k-c)$.

Proof: We need to prove that Condition 1 is satisfied by the code in Construction 6. Let $M = \{i_1, i_2, \ldots, i_c\}$. If $T_2 = \phi$, Condition 1 follows from the support structure of the cyclic repetition code. If $|T_2| = 1$ and some subset of M is included in T_1 , then the support of the union is $[n_1]$ and hence Condition 1 is satisfied. For the case of $|T_2| \ge 2$, since Z_i and Z_j are chosen such that $Z_i \cup Z_j = [n_1]$, for any $i, j \in [C'_{n_1}]$, we have that Condition 1 is trivially satisfied.

Construction 7: (General $n_1, n_2 \ge 2(k-1) + (k-c)$). If $n_1, n_2 \ge 2(k-1) + (k-c)$, we need to find a p such that the following condition is satisfied, i.e.,

$$n_{2} \leq n_{1} + p + C'_{n_{1}+p} = n_{1} + p + \frac{n_{1} + p - k + c}{k - 1}$$
$$p \geq n_{2} - n_{1} - \frac{n_{2} - k + c}{k}.$$

Hence the minimum value of p possible is $p^* = \lceil n_2 - n_1 - \frac{n_2 - k + c}{k} \rceil$. Our objective is to maximize G_3 in Theorem 1. So considering the above mentioned point, the maximum G_3 possible is min $\{n_2 - (n_1 + p^*), C'_{n_1+p^*}\}$. So, for $n_2, n_1 \ge 2(k-1) + (k-c)$, we can use Construction 6 where we replace n_1 with $n_1 + p^*$ servers. Initially, we launch n_1 of $n_1 + p^*$ servers. After c of those servers finish their tasks, the remaining $n_2 - n_1$ servers are launched.

Remark 4: For $n_1, n_2 \ge 2(k-1) + (k-c)$, the computation per server required is proportional to $\frac{n_2-k+1-G_3}{n_2-G_3}$, where $G_3 = \min\{n_2 - (n_1 + p^*), \lfloor \frac{n_1+p^*-k+c}{k-1} \rfloor\}, p^* = \lceil n_2 - n_1 - \frac{n_2-k+c}{k} \rceil$. This proves Theorem 1 for general c and $n_1 \ge 2(k-1) + (k-c)$.

For $2(k-1) < n_1 < 2(k-1)+(k-c)$ and $n_1 \ge 2(k-1)+(k-c)$, the construction is in similar lines to that of Construction 7, except that, for the condition mentioned in Construction 7, the 'p' should satisfy one more condition, i.e., $n_1 + p \ge 2(k-1) + (k-c)$. So the minimum value of p possible is $p^* = \max\{(2(k-1)+(k-c))-n_1, \lceil n_2-n_1-\frac{n_2-k+c}{k}\rceil\}$.

Our aim is to maximize G_4 in Theorem 1. So noting the point mentioned above, the maximum G_4 possible is $\min \{n_2 - (2(k-1) + (k-c)), C'_{n_1+p^*}\}$. So, for $n_2 \ge 2(k-1) + (k-c)$, we can use Construction 6, where we replace n_1 with n_1+p^* servers. Initially, we launch n_1 of n_1+p^* servers. After c servers finish their tasks, the remaining n_2-n_1 servers are launched.

Remark 5: For $2(k-1) < n_1 < 2(k-1) + (k-c), n_2 \ge 2(k-1) + (k-c), n_1 < 2(k-1) + (k-c), \text{ the computation per server required is proportional to <math>\frac{n_2-k+1-G_4}{n_2-G_4}$, where $G_4 = \min\{n_2 - (2(k-1) + (k-c)), \lfloor \frac{n_1+p^*-k+c}{k-1} \rfloor\}, p^* = \max\{(k, k), (k-c)\}$

 $\{(2(k-1)+(k-c))-n_1, \lceil n_2-n_1-\frac{n_2-k+c}{k}\rceil\}.$ This proves Theorem 1 for general c and $n_1 \ge 2(k-1)+(k-c).$

IX. CONCLUSION

This paper provides a framework for tiered gradient codes where all redundant gradient computation servers are not launched at the same time. The framework assumes that when c out of n_1 launched servers finish execution, $n_2 - n_1$ additional servers can be launched, with a property that any k of the servers can be used to compute the gradients. The framework allows for asynchronous launching of servers, and speculative execution by delayed launching of certain servers. Improvement in task computations per server is shown as compared to the case where all n_2 servers are launched without waiting for the results from c out of n_1 servers. It can be noted that the number of data partitions for the tiered gradient coding scheme is always different from that of the gradient coding scheme. In fact, our advantage over gradient coding is achieved by changing the number of partitions.

Coding across epochs, and allowing multiple messages per iteration have been shown to improve the computation [29], [30], while their impact with tiered gradient codes is open. This paper does not consider intermediate feedback while the job is in operation, while impact of such feedback based on [14], [15] is an important next step. Extension of the proposed framework to approximate gradient computations as in [17], [18] is another important problem. Finally, the servers may be heterogenous and considering such impact in multi-tiered gradient code designs is open (Such problem for gradient codes has been studied in [16]).

ACKNOWLEDGMENT

The authors would like to thank Sneh Gupta, who was dual degree student at IIIT Hyderabad for discussions regarding the topic of this article.

REFERENCES

- R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 3368–3376.
- [2] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 5606–5615.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] K. Wan, D. Tuninetti, M. Ji, and P. Piantanida, "Fundamental limits of distributed data shuffling," in *Proc. 56th Annu. Allerton Conf. Commun.*, *Control, Comput. (Allerton)*, Oct. 2018, pp. 662–669.
- [7] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "Straggler mitigation in distributed matrix multiplication: Fundamental limits and optimal coding," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2018, pp. 2022–2026.
- [8] S. Dutta, G. Joshi, S. Ghosh, P. Dube, and P. Nagpurkar, "Slow and stale gradients can win the race: Error-runtime trade-offs in distributed sgd," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2018, pp. 803–812.
- [9] M. F. Aktas and E. Soljanin, "Straggler mitigation at scale," *IEEE/ACM Trans. Netw.*, vol. 27, no. 6, pp. 2266–2279, Dec. 2019.
- [10] 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.
- [11] A. Mallick, M. Chaudhari, and G. Joshi, "Fast and efficient distributed matrix-vector multiplication using rateless fountain codes," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2019, pp. 8192–8196.
- [12] J. Zhang and O. Simeone, "Improved latency-communication trade-off for map-shuffle-reduce systems with stragglers," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2019, pp. 8172–8176.

- [13] Q. Yan, M. Wigger, S. Yang, and X. Tang, "A fundamental storagecommunication tradeoff in distributed computing with straggling nodes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2019, pp. 2803–2807.
- [14] E. Ozfatura, D. Gündüz, and S. Ulukus, "Speeding up distributed gradient descent by utilizing non-persistent stragglers," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2019, pp. 2729–2733.
- [15] E. Ozfatura, S. Ulukus, and D. Gündüz, "Distributed gradient descent with coded partial gradient computations," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2019, pp. 3492–3496.
- [16] H. Wang, S. Guo, B. Tang, R. Li, and C. Li, "Heterogeneity-aware gradient coding for straggler tolerance," in *Proc. IEEE 39th Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Jul. 2019, pp. 555–564.
- [17] H. Wang, Z. B. Charles, and D. S. Papailiopoulos, "Erasurehead: Distributed gradient descent without delays using approximate gradient coding," 2019, arXiv:1901.09671. [Online]. Available: http://arxiv.org/abs/1901.09671
- [18] S. Wang, J. Liu, and N. Shroff, "Fundamental limits of approximate gradient coding," in *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 3, no. 3, pp. 1–22, Dec. 2019.
- [19] Z. Charles and D. Papailiopoulos, "Gradient coding via the stochastic block model," 2018, arXiv:1805.10378. [Online]. Available: http://arxiv.org/abs/1805.10378
- [20] S. Li, S. M. M. 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.
- [21] 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.
- [22] N. Raviv, R. Tandon, A. Dimakis, and I. Tamo, "Gradient coding from cyclic MDS codes and expander graphs," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 4302–4310.
- [23] A. Badita, P. Parag, and V. Aggarwal, "Optimal server selection for straggler mitigation," *IEEE/ACM Trans. Netw.*, vol. 28, no. 2, pp. 709–721, Apr. 2020.
- [24] A. Badita, P. Parag, and V. Aggarwal, "Sequential addition of coded sub-tasks for straggler mitigation," in *Proc. IEEE Infocom*, Jul. 2020.
- [25] Y. Xiang, T. Lan, V. Aggarwal, and Y.-F.-R. Chen, "Joint latency and cost optimization for erasure-coded data center storage," *IEEE/ACM Trans. Netw.*, vol. 24, no. 4, pp. 2443–2457, Aug. 2016.
- [26] V. Aggarwal, Y.-F.-R. Chen, T. Lan, and Y. Xiang, "Sprout: A functional caching approach to minimize service latency in erasure-coded storage," *IEEE/ACM Trans. Netw.*, vol. 25, no. 6, pp. 3683–3694, Dec. 2017.
- [27] S. Chen et al., "When queueing meets coding: Optimal-latency data retrieving scheme in storage clouds," in Proc. IEEE Conf. Comput. Commun. (INFOCOM), Apr. 2014, pp. 1042–1050.
- [28] A Development Sandbox for Data Center to Edge Workloads. Accessed: May 10, 2020. [Online]. Available: https://software.intel. com/content/www/us/en/develop/tools/devcloud.html
- [29] E. Ozfatura, D. Gündüz, and S. Ulukus, "Gradient coding with clustering and multi-message communication," in *Proc. IEEE Data Sci. Workshop* (*DSW*), Jun. 2019, pp. 42–46.
- [30] J. Zhang and O. Simeone, "LAGC: Lazily aggregated gradient coding for straggler-tolerant and communication-efficient distributed learning," *IEEE Trans. Neural Netw. Learn. Syst.*, early access, 2020, doi: 10.1109/TNNLS.2020.2979762.



V. Lalitha (Member, IEEE) received the B.E. degree in electronics and communication engineering from Osmania University, Hyderabad, in 2003, and the M.E. and Ph.D. degrees from the Indian Institute of Science (IISc), Bengaluru, in 2005 and 2015, respectively. She worked as an Engineer with Qualcomm, Hyderabad, in 2006, and a Design Engineer with Conexant Systems, Noida, from 2006 to 2008. Prior to joining IIIT Hyderabad, she worked as a Research Intern at Microsoft Research, Bengaluru. Since May 2015, she has been working as an Assis-

tant Professor with IIIT Hyderabad, where she is currently with the Signal Processing and Communications Research Center. Her research interests include coding for distributed storage and computing, index coding, polar codes, and learning-based codes. She was a recipient of Prof. I. S. N. Murthy medal from IISc in 2005, and the TCS Research Scholarship for the year 2011. Her article won the runner up best paper award at NCC 2019. She is the Present Chapter Chair of the Joint Chapter of Communications and Signal Processing Societies and the IEEE Hyderabad Section.



Vaneet Aggarwal (Senior Member, IEEE) received the B.Tech. degree from the Indian Institute of Technology, Kanpur, India, in 2005, and the M.A. and Ph.D. degrees from Princeton University, Princeton, NJ, USA, in 2007 and 2010, respectively, all in electrical engineering.

He was a Senior Member of Technical Staff Research with the AT&T Labs-Research, Florham Park, NJ, from 2010 to 2014, an Adjunct Assistant Professor with Columbia University, New York City, NY, USA, from 2013 to 2014, and the VAJRA

Adjunct Professor with IISc, Bengaluru, from 2018 to 2019. He is currently an Associate Professor with Purdue University, West Lafayette, IN, USA, where he has been since January 2015. He is also a Distinguished Adjunct Professor with KAU, Saudi Arabia. His current research interests are in communications and networking, cloud computing, and machine learning.

Dr. Aggarwal received Princeton University's Porter Ogden Jacobus Honorific Fellowship in 2009, the AT&T Vice President Excellence Award in 2012, the AT&T Key Contributor Award in 2013, the AT&T Senior Vice President Excellence Award in 2014, the 2017 Jack Neubauer Memorial Award recognizing the Best Systems Paper published in IEEE TRANS-ACTIONS ON VEHICULAR TECHNOLOGY, and the 2018 Infocom Workshop HotPOST Best Paper Award. He is on the Editorial Board of IEEE TRANSACTIONS ON COMMUNICATIONS, IEEE TRANSACTIONS ON GREEN COMMUNICATIONS AND NETWORKING, and IEEE/ACM TRANSACTIONS ON NETWORKING.



B. Sundar Rajan (Fellow, IEEE) was born in Tamil Nadu, India. He received the B.Sc. degree in mathematics from Madras University, Madras, India, in 1979, the B.Tech. degree in electronics from the Madras Institute of Technology, Madras, in 1982, and the M.Tech and Ph.D. degrees in electrical engineering from the Indian Institute of Technology, Kanpur, in 1984 and 1989, respectively.

He was a Faculty Member with the Department of Electrical Engineering, Indian Institute of Technology, Delhi, from 1990 to 1997. Since 1998, he has

been a Professor with the Department of Electrical Communication Engineering, Indian Institute of Science, Bengaluru. His primary researches interests include space-time coding for MIMO channels, distributed space-time coding and cooperative communication, coding for multiple-access and relay channels, and network coding.

Dr. Rajan is a member of the American Mathematical Society. He is a fellow of the Indian National Academy of Engineering, the Indian National Science Academy, the Indian Academy of Sciences, and the National Academy of Sciences, India. He was a recipient of Prof. Rustum Choksi award by I.I.Sc., for excellence in research in Engineering for the year 2009, the IETE Pune Center's S.V.C. Aiya Award for Telecom Education in 2004, the Best Academic Paper Award at the IEEE WCNC 2011, and the J. C. Bose National Fellowship since 2016. He has served as a Technical Program Co-Chair of the IEEE Information Theory Workshop (ITW'02), held in Bengaluru, in 2002. He was an Editor of IEEE WIRELESS COMMUNICATIONS LETTERS from 2012 to 2015, an Associate Editor of Coding Theory of IEEE TRANSAC-TIONS ON INFORMATION THEORY from 2008 to 2011 and 2013 to 2015, and an Editor of IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS from 2007 to 2011.



Shanuja Sasi (Student Member, IEEE) was born in Kerala, India. She received the B.Tech. degree in electronics and communication engineering from the National Institute of Technology, Calicut, in 2016. She is currently pursuing the Ph.D. degree with the Department of Electrical Communication Engineering, Indian Institute of Science. She was a Visiting Scholar with Purdue University in 2019 under Science and Engineering Research Board (SERB) Overseas Visiting Doctoral Fellowship (OVDF) Programme. Her primary research interests include

index coding, coded caching, and gradient coding.