Coded Gradient Aggregation: A Tradeoff Between Communication Costs at Edge Nodes and at Helper Nodes

Birenjith Sasidharan

Dept. of Electronics and Communication Engineering

Govt. Engineering College, Barton Hill, India.

Email: birenjith@gecbh.ac.in

Anoop Thomas

School of Electrical Sciences

Indian Institute of Technology Bhubaneswar, India

Email: anoopthomas@iitbbs.ac.in

Abstract

The increasing amount of data generated at the edge/client nodes and the privacy concerns have resulted in learning at the edge, in which the computations are performed at edge devices and are communicated to a central node for updating the model. The edge nodes have low bandwidth and may be available only intermittently. There are helper nodes present in the network that aid the edge nodes in the communication to the server. The edge nodes communicate the local gradient to helper nodes which relay these messages to the central node after possible aggregation. Recently, schemes using repetition codes and maximum-distance-separable (MDS) codes, respectively known as Aligned MDS Coding (AMC) scheme and Aligend Repetition Coding (ARC) scheme, were proposed. It was observed that in AMC scheme the communication between edge nodes and helper nodes is optimal but with an increased cost of communication between helper and master. An upper bound on the communication cost between helpers and master was obtained. In this paper, a tradeoff between communication costs at edge nodes and helper nodes is established with the help of pyramid codes, a well-known class of locally repairable codes. The communication costs at both the helper nodes and edge nodes are exactly characterized. Using the developed technique, the exact communication cost at helper nodes can be computed for the scheme using MDS codes. In the end, we provide two improved aggregation strategies for the existing AMC and ARC schemes, yielding significant reduction in communication cost at helpers, without changing any of the code parameters.

Index Terms

gradient aggregation, locality, coded computing, tradeoff.

I. Introduction

A large amount of data is getting generated at the edge nodes of the network such as mobile devices and sensors. The generated data is required to train various deep learning models to improve the performance of intelligent applications. The edge nodes are often constrained by communication resources to share the collected data with a central server. In addition to that, the data generated may be sensitive, and sharing with a central server raises privacy concerns. Recently new techniques like federated learning and collaborative learning are considered, which enable the users to train the model at the edge node itself [1–6]. This alleviates the privacy concern since the raw data is decentralized and is not shared with the central server.

In this paper, synchronous gradient descent over a decentralized data set distributed over a fixed number of edge nodes is considered. Gradient descent with centralized data is well studied and many techniques to improve the completion time of the algorithm have been proposed [7–17]. When the data set is decentralized, the edge nodes are required to compute the local gradients which are passed to the central server to perform the global update as in federated learning. The edge nodes often have high latency, low bandwidth and may only be available intermittently. This results in straggling communication links that slow down the learning process. A technique to mitigate this problem is to consider a hierarchical setup in which reliable helper nodes are located close to the edge nodes [18]. These helper nodes can be used by the edge nodes for efficient communication with the server node. The communication between the edge nodes and the helper nodes is unreliable due to the presence of straggling links. The helper nodes transmit the received local gradients to the central node after possible aggregation over a reliable error-free link.

In [18], two schemes are considered one using repetition codes and the other using maximum-distance-separable (MDS) codes. The use of repetition code enables the reduction in communication costs at the helper nodes, however at the cost of an increased communication cost at the edge nodes. The MDS-code-approach reduces the communication cost at the edges at the expense of a larger communication cost at helpers. Inspired by this, we propose a scheme based on pyramid codes [19] which gives a tradeoff between the communication costs at edge nodes and helper nodes. In practice, the helper nodes can be shared by multiple types of edge nodes. The edge computing techniques discussed in [20] considers the use of cloudlets which are similar to the helper nodes in our architecture. Since these nodes are shared by multiple applications, there is a possibility that the helper nodes cannot offer entire resources for transmitting gradients associated to every application. In such scenarios, the proposed scheme enables us to adjust the communication costs at helper nodes. In the next section, we formalize the problem and present the system model.
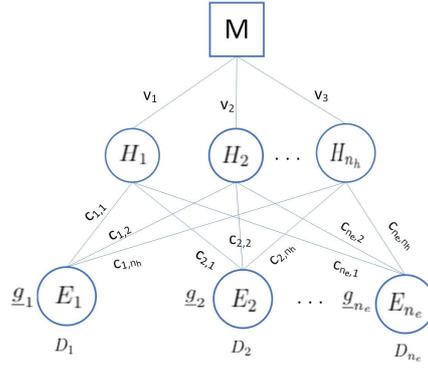
Fig. 1: Hierarchical distributed learning set up with n_e edge nodes and n_h helper nodes.

II. Coded Gradient Aggregation: Preliminaries

A. System Model

Gradient descent is an algorithm used in many machine learning applications to find an optimal parameter $\theta^* \in \mathbb{R}^{p \times 1}$ iteratively that minimizes a function $\sum_{x \in D} \ell(\theta, x) + \lambda R(\theta)$, where $\ell(\cdot)$ is the underlying loss function, $R(\cdot)$ denotes the regularization function, λ is the regularization parameter, and D is the data set. The gradient $\underline{g} = \sum_{x \in D} \nabla \ell(\theta^{(t)}, x)$ is calculated at each iteration t to update the parameter by the rule $\theta^{(t+1)} = \theta^{(t)} - \mu(\underline{g} + \lambda \nabla R(\theta^{(t)}))$, where μ is the learning rate. For large data sets, most of the computational effort is spent on calculating the gradient over entire dataset.

We consider the hierarchical distributed learning setup with n_e edge nodes $E_1, E_2, \ldots, E_{n_e}$ and n_h helper nodes $H_1, H_2, \ldots, H_{n_h}$ as illustrated in Fig. 1. For any positive integer n, let $[n]$ denote the set $\{1, 2, \ldots, n\}$. Each edge node $E_i, i \in [n_e]$ possesses an exclusive data set D_i and can compute the gradient $\underline{g}_i \in \mathbb{R}^p$ over the available data set D_i. The edge nodes make use of the helper nodes to communicate the partial gradients to the central master node M. The edge node E_i splits the computed gradient $\underline{g}_i$ into k vectors $\underline{g}_{i,j} \in \mathbb{R}^{(p/k)}, j \in [k]$. These gradient vectors are encoded using a linear block code (termed as *client code*) of length n_h and dimension k with generator matrix G to obtain $\underline{c}_i = [\underline{g}_{i,1} \ \underline{g}_{i,2} \ \cdots \underline{g}_{i,k}] \cdot G, \forall i \in [n_e]$, where $\underline{c}_i = [\underline{c}_{i,1} \ \underline{c}_{i,2} \ \cdots \underline{c}_{i,n_h}] \in \mathbb{R}^{(p/k) \times n_h}$. Each edge node $E_i, i \in [n_e]$ transmits the coded message $\underline{c}_{i,j}$ to the helper node H_j for all $j \in [n_h]$.

The communication between edge nodes and helper nodes are unreliable which is modelled by considering at most $s \in [n_h - 1]$ straggling links between each edge node and helper nodes. Each of the helper nodes communicate to the master the observed erasure (straggling) pattern, represented by a binary vector of length n_e. The master constructs an $n_e \times n_h$ binary matrix referred to as the observed erasure matrix (straggling matrix), where the value 1 in i, j position indicates the straggling of the communication link between edge node i and helper j. Let $\Omega(s)$ denote the set of all erasure matrices with s straggling links per edge node. Depending on the erasure matrix $\epsilon \in \Omega(s)$, the master node instructs the helper node to encode the messages received from edge nodes to form a coded message $v_j^\epsilon \in \mathbb{R}^{(p/k) \times m_{j,\epsilon}}$. These coded messages are communicated to the master over reliable error-free links. The master makes use of the encoded messages from all the helper nodes to compute $\underline{g}_D = \left[\sum_{i=1}^{n_e} \underline{g}_{i,1} \ \sum_{i=1}^{n_e} \underline{g}_{i,2} \ \cdots \ \sum_{i=1}^{n_e} \underline{g}_{i,k} \right]$. The communication cost at edge nodes is quantified by the normalized communication load C_{EH} where the normalization is over the local gradient vector size. Hence, we have $C_{EH} = \frac{(p/k) \times n_h}{p} = \frac{n_h}{k}$. The communication cost at the helper nodes depends on the straggling pattern. Hence it is quantified by considering the average helper to master communication load C_{HM}. The average helper to master communication load C_{HM} is the total size of messages from all the helpers to master averaged over all erasure patterns $\Omega(s)$, normalized by the gradient vector size p. We have, $C_{HM} = \frac{1}{|\Omega(s)|} \sum_{\epsilon \in \Omega(s)} \sum_{j=1}^{n_h} \frac{m_{j,\epsilon}}{k}$.

For s straggling links between each edge node and helper nodes, a tuple (C_{EH}, C_{HM}) is achievable if there exists a communication scheme between client nodes, helper nodes and the master node satisfying C_{EH} and C_{HM} as defined, that enables the master node to compute the gradient $\underline{g}_D$. The set of all achievable tuples (C_{EH}, C_{HM}) is denoted by $\mathcal{A}$. The objective is to characterize the minimum values of the communication loads defined as follows:

$$C_{EH}^* = \min_{(C_{EH}, C_{HM}) \in \mathcal{A}} C_{EH}, C_{HM}^* = \min_{(C_{EH}, C_{HM}) \in \mathcal{A}} C_{HM}.$$

In [18], it was shown that $C_{EH}^* = \frac{n_h}{n_h - s}$, and $1 \leq C_{HM}^* \leq s + 1$. The results on the optimum values of C_{EH}^* and C_{HM}^* are obtained from two different schemes. The bounds on the C_{HM}^* are obtained by using Aligned Repetition Coding (ARC) scheme. The communication cost at edge nodes however is high for the ARC scheme. The optimum value C_{EH}^* is obtained by using Aligned MDS Coding (AMC) scheme. The optimality is obtained however at an increased communication cost at the helper nodes. The results in [18] indicates the existence of a tradeoff between the communication costs which is explored in this paper.

B. Aligned MDS Coding Scheme

An $[n, k, d]$ code denotes a linear block code of length n, dimension k and minimum distance d. In the AMC scheme, the client code is an $[n_h, n_h - s, s + 1]$ MDS code. Since the dimension of the code is $n_h - s$, $C_{EH} = \frac{n_h}{n_h - s}$. For any straggling pattern, it is guaranteed that $n_h - s$ coded messages are delivered from every edge node. Consider the naive scheme in which each helper node forwards the received messages to the master. From the MDS property, the master node is able to recover the partial gradients and compute the total gradient. This naive scheme has $C_{HM} = n_e$. In the AMC scheme, the master identifies the largest set of rows in the straggling table that match exactly. The helpers corresponding to the columns of matching rows aggregates the gradient vectors and transmit the aggregated vector to the master. For the remaining rows, the helper nodes transmit the received messages to the master. From the aggregated messages the master is able to recover an aggregated partial gradient. Using the remaining received messages the master node is able to calculate the full gradient $\underline{g}_D$. The following bound on C_{HM} is obtained in [18] by modelling the largest set of rows as the maximum occupancy of a balls and bins problem: $C_{HM} \leq n_e - \frac{n_e}{n_h - s} + 1$.

C. Aligned Repetition Coding Scheme

Each edge node $E_i, i \in [n_e]$ partitions the computed gradient $\underline{g}_i$ into k components $\underline{g}_{i,1}, \underline{g}_{i,2}, \ldots, \underline{g}_{i,k}$, where $\underline{g}_{i,r} \in \mathbb{R}^{\frac{p}{k}}$ for $r \in [k]$. The number of components k satisfies the condition that $k = \frac{n_h}{s+1}$. Each component of the gradient vector is transmitted to $s + 1$ helper nodes. Since, the number of erasures is at most s, each partial gradient will be received by at least one helper node thus ensuring resiliency. The communication cost at the edge nodes, C_{EH} for the ARC scheme is $C_{EH} = \frac{n_h}{k} = s + 1$. The client code used in the ARC scheme has a generator matrix $G_{ARC} = [I_k \ldots (s+1) \text{ times} \ldots I_k]$. In the ARC scheme, after obtaining the straggling pattern from helper nodes, the master takes each component and assigns the first helper node which received that particular component for transmission to the master node. Each helper node aggregates all the assigned partitions by the master and transmits an aggregated gradient vector to the master node. Note that the communication cost at each transmitting helper node is $\frac{p}{k}$. The maximum communication cost at the helper nodes arise when all the helper nodes have to transmit and hence $C_{HM} \leq \frac{n_h}{k} = s + 1$.

D. Our Contributions

We propose a new scheme, termed as *pyramid scheme*, of coded gradient aggregation based on pyramid codes, that enables to achieve a tradeoff between the communication costs at the edge nodes and at the helper nodes. In contrast to the traditional application in efficient node-repair in distributed storage, our work expands the realm of codes with locality to the new area of data communication in collaborative learning. We derive an exact expression for C_{HM} and C_{EH} of the proposed scheme. We also derive an exact expression for C_{HM} for the existing AMC scheme. For an example parameter set $(n_e = 2048, n_h = 16, s = 5)$, we can reduce C_{HM} to 1925.5 from 2044.5 by increasing C_{EH} from 1.45 to 2 as illustrated in Fig. 2. Finally,
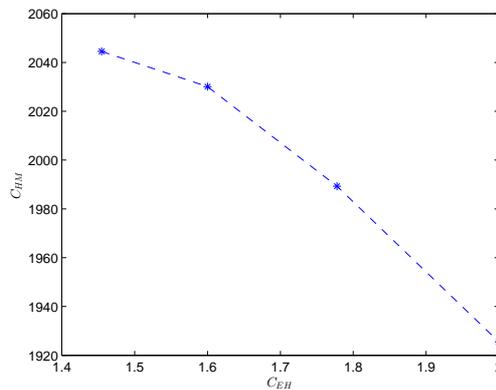
Fig. 2: Tradeoff between C_{EH} and C_{HM} for $(n_e, n_h, s) = (2048, 16, 5)$

we present two improved aggregation strategies for existing AMC and ARC schemes, both achieving significant reduction in C_{HM} as shown by simulations. An improvement in AMC is achieved by picking second, third, and upto m-th maximum when matching rows in the observed erasure matrix are identified by the master. This is in contrast to the present aggregation strategy that only considers the first maximum. An improvement in ARC is achieved by choosing helpers in a greedy manner for availing every partial gradient.

	H_1	H_2	H_3	H_4	H_5	H_6	H_7	H_8
E_1	c_{11}	~~c_{12}~~	c_{13}	~~c_{14}~~	c_{15}	c_{16}	~~c_{17}~~	c_{18}
E_2	c_{21}	~~c_{22}~~	c_{23}	~~c_{24}~~	c_{25}	c_{26}	c_{27}	~~c_{28}~~
E_3	c_{31}	~~c_{32}~~	c_{33}	~~c_{34}~~	c_{35}	c_{36}	c_{37}	~~c_{38}~~
E_4	~~c_{41}~~	c_{42}	~~c_{43}~~	~~c_{44}~~	c_{45}	c_{46}	c_{47}	c_{48}
E_5	c_{51}	~~c_{52}~~	~~c_{53}~~	~~c_{54}~~	c_{55}	c_{56}	c_{57}	c_{58}
E_6	c_{61}	c_{62}	~~c_{63}~~	~~c_{64}~~	~~c_{65}~~	c_{66}	c_{67}	c_{68}
E_7	c_{71}	c_{72}	~~c_{73}~~	~~c_{74}~~	~~c_{75}~~	c_{76}	c_{77}	c_{78}
E_8	c_{81}	c_{82}	c_{83}	~~c_{84}~~	c_{85}	c_{86}	~~c_{87}~~	~~c_{88}~~
E_9	c_{91}	c_{92}	~~c_{93}~~	c_{94}	c_{95}	c_{96}	~~c_{97}~~	~~c_{98}~~
E_{10}	$c_{10,1}$	$c_{10,2}$	~~$c_{10,3}$~~	~~$c_{10,4}$~~	$c_{10,5}$	$c_{10,6}$	$c_{10,7}$	~~$c_{10,8}$~~

Fig. 3: Aggregation strategy for $(n_e = 10, n_h = 8, s = 3)$ using $[8, 4, 4]$ pyramid code with 2 local codes.

III. An Example Scheme Based on Pyramid Codes

Consider a system with $n_e = 10$ client nodes, $n_h = 8$ helper nodes with at most $s = 3$ straggling links between client and helper nodes. We first construct an $[8, 4, 4]$ pyramid code $\mathcal{C}$ over $\mathbb{F}_2^3$ and use it as the client code. Starting from the systematic generator matrix $\tilde{G}_0$ of $[7, 4, 4]$ Reed-Solomon code over $\mathbb{F}_2^3$

$$\tilde{G}_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & a_{11} & a_{12} & a_{13} \\ 0 & 1 & 0 & 0 & a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 & 0 & a_{31} & a_{32} & a_{33} \\ 0 & 0 & 0 & 1 & a_{41} & a_{42} & a_{43} \end{bmatrix}$$

we can construct the generator matrix G_0 of $[8, 4, 4]$ pyramid code $\mathcal{C}_0$ as

$$G_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & a_{11} & 0 & a_{12} & a_{13} \\ 0 & 1 & 0 & 0 & a_{21} & 0 & a_{22} & a_{23} \\ 0 & 0 & 1 & 0 & 0 & a_{31} & a_{32} & a_{33} \\ 0 & 0 & 0 & 1 & 0 & a_{41} & a_{42} & a_{43} \end{bmatrix}.$$

Every real gradient vector can be quantized so that it is represented using a finite number of bits. In that case, it is possible to encode it using a client-code over finite-field. There is a parallel line of work [21–23] that attempts to establish that such quantizations do not affect the performance of algorithms making use of the gradients. Every client node splits the local gradient vector into 4 sub-vectors each of length $(p/4)$ and use the same code $\mathcal{C}_0$ to encode the gradient vector to produce a vector of length $(p/4) \cdot 8 = 2p$. Thus we obtain a (10×8) codeword array with every entry being a $(p/4)$-long vector. It is easy to find that $C_{EH} = \frac{8}{4} = 2$. The code $\mathcal{C}_0$ has two local single parity check codes of dimension 2 with supports $L_1 = \{1, 2, 5\}$ and $L_2 = \{3, 4, 6\}$. The global parities are over the support $P = \{7, 8\}$. Suppose the observed erasure matrix is as depicted by cross-marks in Fig. 3. Thus the observed erasure matrix is given by

$$\mathcal{M}_E = \begin{bmatrix} 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Given an erasure pattern, a local code is said to be unaffected by the erasure pattern, if every symbol in the local code can be decoded. Otherwise, the local code is said to be overwhelmed by the erasure pattern. The coded symbols are aggregated (summed up) at helper nodes before transmitting to the master. A set of columns T is picked such that

- they do not associate to global parities if no local code is overwhelmed, or else,
- they associate to overwhelmed local codes and global parities,

and hold unerased information symbols. Restricted to these columns, set of rows Ψ with a matching erasure pattern is picked. Symbols in every column of T are summed up over Ψ. Let us look at how it works as applied to the given example.

1) Over the rows $E1 - E3$, we aggregate symbols over the columns $\{1, 2, \ldots, 6\}$. This is because neither L_1 nor L_2 are overwhelmed. Thus we choose columns in $L_1 \cup L_2$ as candidate columns for aggregation. The information set within these columns is $\{1, 3, 5, 6\}$. Therefore the helpers $H_j, j = 1, 3, 5, 6$ send $c_{1j} + c_{2j} + c_{3j}$ to the master.

2) Over the rows $E4 - E7$, the local code L_2 is overwhelmed. So we choose $L_2 \cup P = \{3, 4, 6, 7, 8\}$ as columns for aggregation. Symbols belonging to unaffected local code L_1 are sent without aggregation. Once L_1 is completely decoded in every row, then symbols in columns $\{1, 2\}$ are available at every row $E4 - E7$. Since $\{1, 2, 7, 8\}$ is a information set, symbols are aggregated and sent for columns $\{7, 8\}$. Thus helpers H_j, $j = 7, 8$ send aggregated symbols $c_{4j} + c_{5j} + c_{6j} + c_{7j}$.

3) For $E8$, both the local codes are unaffected. Hence aggregation will be done over $L_1 \cup L_2$. However there are no other row with a matching erasure pattern. So the symbols from this row will be sent without aggregation.

4) The case of $E9$ is exactly similar to the case of $E8$.

5) For $E10$, L_2 is an overwhelmed code. Thus aggregation will be attempted over columns $L_2 \cup P$. However, there are no rows left out. For this reason, symbols from $E10$ will be sent without aggregation.

The above described aggregation is illustrated in Figure 3.

IV. THE PYRAMID SCHEME

The example described in Sec. III can be generalized to arrive at a new scheme for coded aggregation of gradient. In this section, we present the new scheme, referred to as *pyramid scheme*, for a fixed parameter set (n_e, n_h, s). Let $t \geq 2$ be an integer parameter. A pyramid scheme is a pair $(\mathcal{C}_t, \mathcal{A}_t)$ indexed by t, where $\mathcal{C}_t$ is a pyramid code and $\mathcal{A}_t$ is the associated aggregation strategy. For the degenerate case of $t = 1$, the scheme reduces to the AMC scheme. In the following, we describe the scheme and characterize its communication cost-pair $(C_{\text{EH,Pyr}}(t), C_{\text{HM,Pyr}}(t))$ for every possible value of t.

A. Construction of a Set of Pyramid Codes

Let the parameters be (n_e, n_h, s). We set $k_1 = n_h - s$ and define $k_t = k_1 - t + 1$, $2 \leq t \leq \lfloor \frac{k_1+1}{3} \rfloor$. Construct a $[k_t + s, k_t, s + 1]$-MDS code over $\mathbb{F}_q$ with generator matrix

$$
\begin{aligned}
G_{t,\text{MDS}} &= [I_{k_t} \ P] \\
&= \begin{bmatrix} I_{k_t} & \underline{p}_1 & \underline{p}_2 & \cdots & \underline{p}_s \end{bmatrix}
\end{aligned}
$$

where I_{k_t} is an identity matrix of size k_t and $\underline{p}_i, 1 \leq i \leq s$ are columns of the $(k_t \times s)$ Cauchy matrix. The field size q may be suitably chosen. We transform $G_{t,\text{MDS}}$ to derive the generator matrix G_t of a pyramid code $\mathcal{C}_t$ with blocklength n_h and dimenstion k_t. First, we write

$$
\begin{aligned}
k_t &= ta + b, \ 0 \leq b < t \\
&= b(a+1) + (t-b)a
\end{aligned}
$$

and split the parity vector $\underline{p}_1$ into t vectors $\underline{p}_{11}, \underline{p}_{12}, \dots, \underline{p}_{1t}$ such that the first b vectors $\underline{p}_{11}, \underline{p}_{12}, \dots \underline{p}_{1b}$ belong to $\mathbb{F}_q^{a+1}$ and the remaining $(t-b)$ vectors $\underline{p}_{1,b+1}, \underline{p}_{1,b+2}, \dots, \underline{p}_{1t}$ belong to $\mathbb{F}_q^a$. Then we construct $G_t = [I_{k_t} \ P]$ where P is a $k_t \times (n_h - k_t)$ matrix

$$
P = \begin{bmatrix}
\underline{p}_{11} & & & & & & & \\
& \underline{p}_{12} & & & & & & \\
& & \ddots & & & & & \\
& & & \underline{p}_{1b} & & & \underline{p}_2 & \cdots & \underline{p}_s \\
& & & & \underline{p}_{1,b+1} & & & \\
& & & & & \ddots & & \\
& & & & & & \underline{p}_{1t} &
\end{bmatrix}
$$

We put in place the required definitions and subsequently characterize the parameters of $\mathcal{C}_t$.

Definition 1:

1) An symbol c_i, $1 \leq i \leq n$ an $[n, k]$ code $\mathcal{C}$ is said to have locality r if there is a punctured code C_i of dimension r containing c_i.

2) An $[n, k]$ code $\mathcal{C}$ is said to be a code with locality r if every symbol has locality r. If every symbol c_i has locality r_i, and r_i is not necessarily equal to r_j for $i \neq j$, then the code is said to be a code with unequal locality.

3) The vector $(\kappa_1, \kappa_2, \dots, \kappa_r)$ is said to be the locality profile of an $[n, k]$ code $\mathcal{C}$ with unequal locality if

   - $\kappa_1 + \kappa_2 + \cdots + \kappa_r = k$
   - there are κ_j information symbols in $\mathcal{C}$ with locality j, $1 \leq j \leq r$.

The code $\mathcal{C}_t$ generated by G_t is a pyramid code [19] with t local codes. It is a code with unequl locality [24] [25] with every local code a single-parity-check code.

Lemma 4.1: The code $\mathcal{C}_t$, $2 \leq t \leq \lfloor \frac{k_1+1}{3} \rfloor$ is an $[n_h, k_t, s+1]$ code. It can thus correct any s erasures.

Proof: The block length of $\mathcal{C}_t$ can easily be computed as $k_t + t + (s-1) = k_1 + s = n_h$. By construction, the locality profile of the code

$$\underline{\kappa} \;=\; (\kappa_1, \kappa_2, \cdots, \kappa_a, \kappa_{a+1}) \;=\; (0, 0, \cdots, 0, (t-b)a, b(a+1)).$$

By Theorem 1 in [25], the minimum distance of the code

$$
\begin{aligned}
d \;&\leq\; n - k - \sum_{j=1}^{a+1} \left\lceil \frac{\kappa_j}{j} \right\rceil + 2 \\
&=\; n_h - k_t - \left\lceil \frac{(t-b)a}{a} \right\rceil - \left\lceil \frac{b(a+1)}{a+1} \right\rceil + 2 \\
&=\; n_h - k_1 + 1 \;=\; s + 1.
\end{aligned}
$$

It is shown in [25] that pyramid codes acheive the upper bound on minimum distance. This completes the proof. $\square$

Furthermore, we can explicitly identify the supports of all t local codes as $L_i = \{(i-1)(a+1)+1, (i-1)(a+1)+2, \ldots, i(a+1), k_t+i\}, 1 \leq i \leq b$ and $L_i = \{(i-1)a+b+1, (i-1)a+b+2, \ldots, ia+b, k_t+i\}, b+1 \leq i \leq t$. Then we define $Q = [n] \setminus \cup_{i \in [t]} L_i$ as the support of all global parities that do not participate in any local code. We shall also define $k_t(i), 1 \leq i \leq t$ as the dimension of the local code with support L_i. Clearly $\sum_{i=1}^{t} k_t(i) = k_t$. In the construction of $\mathcal{C}_t$, we do not wish to end up with a local code of trivial dimension 1. Thus we shall require $ta + b = k_1 - t + 1$ with $a \geq 2$, $0 \leq b < t$ and thereby it follows that

$$t \;=\; \frac{k_1 + 1 - b}{a+1} \;\leq\; \left\lfloor \frac{k_1+1}{3} \right\rfloor.$$

B. A Classification of Erasure Patterns

As discussed in Sec. II-A, an erasure (failure) pattern of links connecting edge clients to helpers can be represented as a binary vector of length n_h with Hamming weight s. Let $\mathcal{E}$ denote the set of all erasure patterns, and we have $|\mathcal{E}| = \binom{n_h}{s}$. Every edge node will be subjected to an erasure pattern, and there are a total of n_e erasure patterns witnessed by the system. The aggregation strategy $\mathcal{A}_t$ adopted by the helpers (as prescribed by the master) depends on these erasure patterns. Here we shall discuss a partition of $\mathcal{E}$ based on $\mathcal{C}_t$ that will become necessary in the subsequent section to describe $\mathcal{A}_t$.

For any binary vector $\underline{x} = (x_1, x_2, \ldots, x_n)$, we define $w_H(\underline{x})$ as the Hamming weight of $\underline{x}$, and $Supp(\underline{x})$ as the support of $\underline{x}$ i.e., $\{i \mid x_i \neq 0\}$. For the code $\mathcal{C}_t$, we say a local code or its support $L_i, 1 \leq i \leq t$ (by overloading of definition) is *overwhelmed* by an erasure pattern $\underline{e} \in \mathcal{E}$ if $Supp(\underline{e}) \cap L_i > 1$. Otherwise, if $Supp(\underline{e}) \cap L_i \leq 1$, we say the local code or L_i is *unaffected* by $\underline{e}$. For a given erasure pattern, let τ denote the number of unaffected local codes, and $i_1 < i_2 < \cdots < i_\tau$ denote the ordered indices of unaffected local codes. Similarly, $j_1 < j_2 \cdots < j_{t-\tau}$ denote the ordered indices of overwhelmed local codes. We define $f_i = |Supp(\underline{e}) \cap L_i|$ as the number of erasures within a local code L_i, $1 \leq i \leq t$, and $f_{t+1} = |Supp(\underline{e}) \cap Q|$ as the number of erasures within global parities. The vector $\underline{f} = (f_1, f_2, \ldots, f_{t+1})$ is referred to as the *code-erasure pattern*. Let $\lambda_i = |L_i|$, $1 \leq i \leq t$ denote the size of supports of local codes, and $\lambda_{t+1} = |Q| = (s-1)$ denote the number of global parities. Then there are

$$N(\underline{f}) \;=\; \prod_{i=1}^{t+1} \binom{\lambda_i}{f_i}$$

erasure patterns associated with a particular code-erasure pattern.

Let $\underline{u} \in \mathbb{F}_2^t$ and $\underline{v} \in \mathbb{F}_2^{t - w_H(\underline{u})}$. We define a subset $S_{(\underline{u}, \underline{v})} \subset \mathcal{E}$ indexed by the tuple $(\underline{u}, \underline{v})$ as follows. An erasure pattern $\underline{e}$ belongs to $S_{(\underline{u}, \underline{v})}$ if the following two conditions hold:
 (a) For every $i \in Supp(\underline{u})$, L_i is overwhelmed by $\underline{e}$, thereby implying $\tau = t - w_H(\underline{u})$,
 (b) If $i_1, i_2, \ldots i_\tau$ are the ordered indices of unaffected local codes, then $(f_{i_1}, f_{i_2}, \ldots, f_{i_\tau}) = \underline{v}$.
This means that (a) $\underline{u}$ determines which all local codes are overwhelmed by $\underline{e}$, and that (b) $\underline{v}$ determines the number of erasures (either 0 or 1) within unaffected local codes. We say that two erasure patterns in $\mathcal{E}$ are of *same type*, if they belong to the same subset $S_{(\underline{u}, \underline{v})}$. And we write $type(\underline{e}) = (\underline{u}, \underline{v})$ if $\underline{e} \in S_{(\underline{u}, \underline{v})}$. The size of $S_{(\underline{u}, \underline{v})}$ can be counted as:

$$|S_{(\underline{u}, \underline{v})}| \;=\; \binom{n_h - \sum_{\ell \notin \{i_1, i_2, \ldots i_\tau\}} \lambda_\ell}{s - w_H(\underline{v})} \prod_{\ell=1}^{\tau} \binom{\lambda_{i_\ell}}{v_\ell} - \sum_{\substack{\underline{x} \in \mathbb{F}_2^t \\ (x_{i_1}, x_{i_2}, \ldots, x_{i_\tau}) = \underline{v}}} \binom{s-1}{s - w_H(\underline{x})} \prod_{\ell=1}^{t} \binom{\lambda_\ell}{x_\ell}$$

Next, we shall define a relation $\simeq$ within $S_{(\underline{u}, \underline{v})}$. Define $A(\underline{u}) = \cup_{i \in Supp(\underline{u})} L_i \cup Q$ as the support of all overwhelmed local codes combined with that of global parities. Define $B(\underline{u}) = \cup_{i \in [t] \setminus Supp(\underline{u})} L_i$ as the support of unaffected local codes. Let $\underline{e}_1, \underline{e}_2 \in \mathcal{E}$. We say that $\underline{e}_1$ is *equivalent* to $\underline{e}_2$ or $\underline{e}_1 \simeq \underline{e}_2$ if they are of same type and

$$
\begin{aligned}
Supp(\underline{e}_1) \cap A(\underline{u}) \;&=\; Supp(\underline{e}_2) \cap A(\underline{u}) \text{ if } \underline{u} \neq \underline{0}, \\
Supp(\underline{e}_1) \cap B(\underline{u}) \;&=\; Supp(\underline{e}_2) \cap B(\underline{u}) \text{ if } \underline{u} = \underline{0}.
\end{aligned}
$$

Type $(\underline{u}, \underline{v})$	$|S_{(\underline{u},\underline{v})}|$	$\underline{f} = (f_1, f_2, f_3)$	$N(\underline{f})$	$\beta(\underline{u}, \underline{v})$
$(00, 00)$	0	$-$	0	$-$
$(00, 01)$	6	$(0, 1, 4)$	6	1
$(00, 10)$	6	$(1, 0, 4)$	6	1
$(00, 11)$	144	$(1, 1, 3)$	144	4
$(01, 0)$	246	$(0, 2, 3)$	60	1
		$(0, 3, 2)$	120	1
		$(0, 4, 1)$	60	1
		$(0, 5, 0)$	6	1
$(01, 1)$	1110	$(1, 2, 2)$	540	6
		$(1, 3, 1)$	480	6
		$(1, 4, 0)$	90	6
$(10, 0)$	246	$(2, 0, 3)$	60	1
		$(3, 0, 2)$	120	1
		$(4, 0, 1)$	60	1
		$(5, 0, 0)$	6	1
$(10, 1)$	1110	$(2, 1, 2)$	540	6
		$(3, 1, 1)$	480	6
		$(4, 1, 0)$	90	6
$(11, -)$	1500	$(2, 2, 1)$	900	1
		$(2, 3, 0)$	300	1
		$(3, 2, 0)$	300	1

Fig. 4: Classification of erasure patterns for a $[16, 10, 6]$ pyramid code with $t = 2$, $(\lambda_1, \lambda_2, \lambda_3) = (6, 6, 4)$.

In otherwords, equivalent erasure patterns remain the same at every location except possibly within unaffected local codes, and their code-erasure pattern remains the same within unaffected local codes. It is easy to verify that $\simeq$ is an equivalence relation, and hence partitions $S_{(\underline{u},\underline{v})}$ into equivalence classes. The set of equivalence classes is denoted by $\mathcal{B} = S_{(\underline{u},\underline{v})}/\simeq$, and we define $\mu_{(\underline{u},\underline{v})} = |\mathcal{B}|$. Whenever the type $(\underline{u}, \underline{v})$ is evident from the context, we simply write μ in place of $\mu_{(\underline{u},\underline{v})}$. The equivalence classes within $\mathcal{B}$ are denoted by $B_1(\underline{u}, \underline{v}), B_2(\underline{u}, \underline{v}), \ldots, B_\mu(\underline{u}, \underline{v})$. The size of every equivalence class $B_j(\underline{u}, \underline{v})$, $1 \le j \le \mu$ is the same and we define

$$
\begin{aligned}
\beta(\underline{u}, \underline{v}) &= |B_j(\underline{u}, \underline{v})| \\
&= \begin{cases} \prod_{\ell \in [t] \setminus Supp(\underline{u})} \binom{\lambda_\ell}{f_\ell} & \underline{u} \ne \underline{0} \\ \binom{\lambda_{t+1}}{f_{t+1}} & \underline{u} = \underline{0} \end{cases}
\end{aligned}
$$

where $(f_1, f_2, \ldots, f_{t+1})$ is the code-erasure pattern of $\underline{e}$. We refer to $\beta(\underline{e})$ as the *bunching factor* of every $\underline{e} \in B_j(\underline{u}, \underline{v})$, for reasons that will become apparent in later subsections. Since $(f_{i_1}, f_{i_2}, \ldots, f_{i_u})$ remains the same for every erasure pattern in $S_{(\underline{u},\underline{v})}$, the bunching factor too remains the same for them. Thus $\simeq$ partitions $S_{(\underline{u},\underline{v})}$ into equivalence classes of the same size. An example partition of $\mathcal{E}$ is illustrated in Fig. 4.

C. The Aggregation Strategy $\mathcal{A}_t$

As discussed in Sec. II-A, the codeword transmitted from all edges is an $(n_e \times n_h)$ matrix over $\mathbb{F}_q$

$$
C = \begin{bmatrix} c_{11} & c_{12} & \cdots & c_{1,n_h} \\ c_{21} & c_{22} & \cdots & c_{2,n_h} \\ \vdots & \vdots & \ddots & \vdots \\ c_{n_e,1} & c_{n_e,2} & \cdots & c_{n_e,n_h} \end{bmatrix}
$$

All n_e erasure patterns together can be represented as a matrix $\mathcal{M}_E$ over $\mathbb{F}_2$.

$$
\mathcal{M}_E = \begin{bmatrix} \epsilon_{11} & \epsilon_{12} & \cdots & \epsilon_{1,n_h} \\ \epsilon_{21} & \epsilon_{22} & \cdots & \epsilon_{2,n_h} \\ \vdots & \vdots & \ddots & \vdots \\ \epsilon_{n_e,1} & \epsilon_{n_e,2} & \cdots & \epsilon_{n_e,n_h} \end{bmatrix} := \begin{bmatrix} \underline{\epsilon}_1 \\ \underline{\epsilon}_2 \\ \vdots \\ \underline{\epsilon}_{n_e} \end{bmatrix}.
$$

We refer to $\mathcal{M}_E$ as the observed erasure matrix. The code symbol c_{ij} is unerased if $\epsilon_{ij} = 0$, and erased if $\epsilon_{ij} = 1$. Here $\underline{\epsilon}_i$, $1 \le i \le n_e$ is the erasure pattern observed by the i-th edge. It is possible that an erasure pattern can repeat at multiple edge-clients, and thus $E = \{\underline{\epsilon}_1, \underline{\epsilon}_2, \ldots, \underline{\epsilon}_{n_e}\}$ is a multi-set of observed erasure patterns.

Our scheme is inspired by the AMC scheme proposed in [18]. In AMC, code symbols from different rows of the codeword array are "aggregated" and transmitted to reduce the communication cost from helpers to the master. In our approach, we will carry out aggregation of symbols leveraging upon the local-decoding capabilities of $\mathcal{C}_t$. The aggregation strategy will take into

account the distribution of overwhelmed and unaffected local codes. We begin with formalizing what is meant by aggregation of code symbols.

Definition 2: Let the code $\mathcal{C}_t$ and the observed erasure matrix $\mathcal{M}_E$ be given. We say that the code-symbols are *aggregated* over a set of rows $R \subset [n_e]$ and a set of columns $T \subset [n_h]$ if the following operations are executed in order:

(a) Identify a minimal subset $I_{R,\overline{T}}$ of $C_{R,\overline{T}} = \{c_{i,j} \mid j \in R, i \in [n] \setminus T\}$ such that symbols in $I_{R,\overline{T}}$ are unerased, and it is possible to recover $C_{R,\overline{T}}$ given the knowledge of $I_{R,\overline{T}}$.

(b) Identify subsets $T_1 \subseteq T$, $\overline{T}_1 \subset [n] \setminus T$ such that (i) $T_1 \cup \overline{T}_1$ is an information set of $\mathcal{C}_t$, and (ii) $c_{i,j}$ is not erased for every $j \in T_1$ and every $i \in R$.

(c) Compute

$$d_j = \sum_{i \in R} c_{ij}, \quad j \in T_1$$

and generate the set $D = \{d_j \mid j \in T_1\} \cup I_{R,\overline{T}}$.

We refer to D as the *aggregated set of code symbols* over a set of rows R and a set of columns T.

In the following, we describe the aggregation strategy $\mathcal{A}_t$ associated with the pyramid code $\mathcal{C}_t$, $2 \leq t \leq \lfloor \frac{n_h - s + 1}{3} \rfloor$. The strategy takes in the codeword C and the observed erasure matrix $\mathcal{M}_E$ (or the multi-set E) as inputs.

1) Let $\{(\underline{u}_j, \underline{v}_j) \mid j = 1, 2, \ldots, j_{\max}\}$ be the set of types of erasure patterns in E. Let us write $\mu_j = \mu(\underline{u}_j, \underline{v}_j)$. Let $\Psi_{j,\ell} = \{i \mid type(\underline{e}_i) = (\underline{u}_j, \underline{v}_j), \ \underline{e}_i \in B_j(\underline{u}_j, \underline{v}_j)\}$, $j = 1, 2, \ldots, j_{\max}, \ell = 1, 2, \ldots, \mu_j$.

2) For every $j = 1, 2, \ldots, j_{\max}$, do the following:
   (a) compute $\ell^*(j) = \arg\max_{\ell=1,2,\ldots,\mu_j} |\Psi_{j,\ell}|$
   (b) Let $T(\underline{u}_j) = A(\underline{u}_j)$ if $\underline{u}_j \neq \underline{0}$, or else $T(\underline{u}_j) = B(\underline{u}_j)$. Aggregate code symbols over set of rows $\Psi_{j,\ell^*(j)}$ and columns $T(\underline{u}_j)$. Let D_j be the aggregated set of symbols.

3) For every row $\ell \in [n_e] \setminus \cup_{j=1}^{j_{\max}} \Psi_{j,\ell^*(j)}$, identify a set of unerased information symbols I_ℓ. Let $D_{j_{\max}+1} = \cup_\ell I_\ell$.

4) The master instructs helpers to transmit $D_{\mathrm{ag}} = \cup_{j=1}^{j_{\max}+1} D_j$.

D. Calculation of C_{EH} and C_{HM}

By Lemma 4.1, the code $\mathcal{C}_t$ has dimension $k_t = n_h - s - t + 1$. It follows that the communication cost by edge-nodes for the pyramid scheme $(\mathcal{C}_t, \mathcal{A}_t)$ is given by:

$$C_{\mathrm{EH,Pyr}}(t) = \frac{n_h}{n_h - s - t + 1} \tag{1}$$

We assume that the observed erasure matrix is a uniform random variable in $\mathbb{F}_2^{n_e \times n_h}$. Since the aggregation strategy depends on the erasure multi-set E, the size of aggregated set of symbols D is a random variable. Thus C_{HM} is the expected value of the random variable $|D_{\mathrm{ag}}|/k_t$:

$$C_{\mathrm{HM,Pyr}}(t) = \frac{\mathbb{E}[|D_{\mathrm{ag}}|]}{k_t}. \tag{2}$$

The problem thus reduces to finding expectation of $|D_{\mathrm{ag}}|$. For every $\underline{u} \in \mathbb{F}_2^t, \underline{v} \in \mathbb{F}_2^{t - w_H(\underline{u})}, \ell \in [\mu(\underline{u}, \underline{v})]$, define the random variable

$$M_\ell(\underline{u}, \underline{v}) = |\{j \in [n_e] \mid \underline{\epsilon}_j \in B_\ell(\underline{u}, \underline{v})\}|$$

as the number of rows subjected to an erasure pattern in $B_\ell(\underline{u}, \underline{v})$. Next we define

$$M(\underline{u}, \underline{v}) = \max_\ell M_\ell(\underline{u}, \underline{v}).$$

The strategy $\mathcal{A}_t$ attempts to aggregate code symbols over $M(\underline{u}, \underline{v})$ rows for every observed type $(\underline{u}, \underline{v})$. Therefore, if $M(\underline{u}, \underline{v}) = 0$, then it does not impact the value of $|D_{\mathrm{ag}}|$, but if $M(\underline{u}, \underline{v}) > 0$, aggregation will happen over those rows and columns given by $A(\underline{u})$. The size of aggregated set of code symbols for each $(\underline{u}, \underline{v})$ depends on the dimension accumulated by both the

unaffected local codes and the overwhelmed local codes. Let $\mathbb{1}(x)$ be an indicator function, i.e., $\mathbb{1}(x) = 1$ if $x \neq 0$, and $\mathbb{1}(x) = 1$ if $x = 0$. We can express D_{ag} and compute its expectation as given below:

$$|D_{\text{ag}}| = \sum_{\underline{u} \in \mathbb{F}_2^t} \sum_{\underline{v} \in \mathbb{F}_2^{t-w}} \mathbb{1}(M(\underline{u}, \underline{v})) \left[\sum_{i \in [t] \setminus Supp(\underline{u})} k_t(i) M(\underline{u}, \underline{v}) + \sum_{i \in Supp(\underline{u})} k_t(i) \right] + \left(n_e - \sum_{\underline{u} \in \mathbb{F}_2^t} \sum_{\underline{v} \in \mathbb{F}_2^{t-w}} M(\underline{u}, \underline{v}) \right) k_t \quad (3)$$

$$\mathbb{E}[|D_{\text{ag}}|] = \sum_{\underline{u} \in \mathbb{F}_2^t} \sum_{\underline{v} \in \mathbb{F}_2^{t-w}} \left[\sum_{i \in [t] \setminus Supp(\underline{u})} k_t(i) \mathbb{E}[M(\underline{u}, \underline{v})] + \sum_{i \in Supp(\underline{u})} k_t(i) \Pr(M(\underline{u}, \underline{v}) \geq 1) \right] +$$
$$\left(n_e - \sum_{\underline{u} \in \mathbb{F}_2^t} \sum_{\underline{v} \in \mathbb{F}_2^{t-w}} \mathbb{E}[M(\underline{u}, \underline{v})] \right) k_t \quad (4)$$

In order to compute the expression, it is required to determine both $\mathbb{E}[M(\underline{u}, \underline{v})]$ and $\Pr(M(\underline{u}, \underline{v}) \geq 1)$ for every $(\underline{u}, \underline{v})$. Following the approach in [18], we identify the random variable $M(\underline{u}, \underline{v})$ with a quantity associated with well-known balls-and-bins experiment. Every row of $(n_e \times n_h)$-codeword is identified as an unlabeled ball, and each erasure pattern in $\mathcal{E}$ as a labelled bin. The random experiment is to throw these n_e unlabeled balls at $\binom{n_h}{s}$ labeled bins. The number of balls in a bin, referred to as bin-size, is the number of rows subjected to the same erasure pattern. Without loss of generality, let us restrict focus to a single type, i.e., we focus on to bins associated with $S_{(\underline{u}, \underline{v})}$. Since we are interested in number of rows subjected to equivalent erasure patterns, we will bunch up multiple (real) bins to form an artificial bin. As the bunching factor remains the same for a given type, we bunch up a fixed number of real bins. It follows that $M(\underline{u}, \underline{v})$ is maximum over bin-sizes of all artificial bins of a fixed type. Thus the problem reduces to the following: Throw r unlabeled balls in n labeled bins uniformly at random. The first bm bins are bunched up to form m artificial bins each constituting of b real bins. Let $X_i, i = 1, 2, \ldots, n - bm + m$ denote the bin-sizes of all bins, and let $Z = \max_{i=1,2,\ldots,m} X_i$. Let us define $\rho(n, r, m, b) = \mathbb{E}[Z]$ and $\phi(n, r, m, b) = \Pr(Z \geq 1)$. Then we have

$$\mathbb{E}[M(\underline{u}, \underline{v})] = \rho\left(\binom{n_h}{s}, n_e, \frac{|S_{(\underline{u}, \underline{v})}|}{\beta(\underline{u}, \underline{v})}, \beta(\underline{u}, \underline{v}) \right) \quad (5)$$

$$\Pr[M(\underline{u}, \underline{v}) \geq 1] = \phi\left(\binom{n_h}{s}, n_e, \frac{|S_{(\underline{u}, \underline{v})}|}{\beta(\underline{u}, \underline{v})}, \beta(\underline{u}, \underline{v}) \right) \quad (6)$$

Both $\rho(n, r, m, b)$ and $\phi(n, r, m, b)$ are determined in App. A. Substituting (5), (6) back in (4) and then to (2), we obtain an expression for communication cost at helper-nodes for pyramid scheme $(\mathcal{C}_t, \mathcal{A}_t)$. Thus we have proved the theorem:

Theorem 4.2: Let (n_e, n_h, s) be given. The pyramid scheme $(\mathcal{C}_t, \mathcal{A}_t)$ for coded aggregation of gradient achieves the communication cost-pair $(C_{\text{EH,Pyr}}(t), C_{\text{HM,Pyr}}(t))$ as given in (1), (2) for $2 \leq t \leq \lfloor \frac{n_h - s + 1}{3} \rfloor$.

V. AN ACHIEVABLE TRADEOFF BETWEEN C_{EH} AND C_{HM}

The communication cost-pairs achieved by the pyramid scheme combined with that of AMC scheme permit to achieve a tradeoff between C_{EH} and C_{HM}.

A. C_{HM} of AMC

In [18], authors derived an upperbound on the communication cost at helpers C_{HM} of the AMC scheme. With no further modification, their bound can be converted to an equality with the help of $\rho(\cdot)$ function. Thus we obtain the following theorem.

Theorem 5.1: Let (n_e, n_h, s) be given. The communication cost-pair $(C_{\text{EH,AMC}}, C_{\text{HM,AMC}})$ of AMC scheme is given by

$$C_{\text{EH,AMC}} = \frac{n_h}{n_h - s}$$
$$C_{\text{HM,AMC}} = n_e + 1 - \rho\left(\binom{n_h}{s}, n_e, \binom{n_h}{s}, 1 \right)$$

B. Illustration of Tradeoff with An Example

For practical reasons, it is desirable to scale n_h logarithmically with n_e and to keep $s = \alpha n_h$ as a fraction of n_h. Following this rationale, we choose example parameters $(n_e, n_h, s) = (2048, 16, 5)$ to illustrate the tradeoff in Fig. 2.

VI. IMPROVED AMC CONSIDERING m MAXIMAS

In the AMC scheme [18], the master finds the maximum number of matching rows in the erasure matrix. Within these matching rows, the helper nodes can aggregate and transmit the aggregated vector to the master node. For the remaining rows, the helper nodes transmit the received messages without any aggregation to the master. We observe that the aggregation can be performed at the helper nodes not just for the maximum number of matching rows, but even for second maximum, third maximum and so on. To illustrate this, consider the system model with $n_e = 5, n_h = 4$ and $s = 1$. Let the observed erasure matrix be

$$
\mathcal{M}_E \; = \;
\begin{bmatrix}
0 & 0 & 0 & 1 \\
0 & 0 & 0 & 1 \\
0 & 0 & 0 & 1 \\
1 & 0 & 0 & 0 \\
1 & 0 & 0 & 0
\end{bmatrix}.
$$

Since the maximum matching rows is three, as per the AMC scheme, the helper nodes aggregates the first three rows and transmit the aggregated gradient. For the last two rows, it transmits all the received messages directly. However, the helper nodes can aggregate the messages in the last two rows and reduce the cost of transmission at helper nodes. As per the AMC scheme, for the observed erasure matrix $\mathcal{M}_E$ the communication cost at the helper nodes is 3. By considering the second maxima also, the communication cost at the helper nodes can be reduced to 2.

The idea is to generalize the AMC scheme by considering more than one set of maximum matching rows. We propose a scheme in which the helper nodes consider m maximas for some integer $m \geq 1$. For $m = 1$ the proposed scheme reduces to the AMC scheme. Let $M_i, i \in [m]$ denote the number of i-th maximum set of matching rows. The helper nodes perform aggregation at each of these matching sets and transmits aggregated gradient vectors. For $i \in [m]$, the helper nodes can aggregate M_i rows and transmit the aggregated partial gradients. The communication cost incurred for the transmission of aggregated partial gradient for each set of matching rows is $(n_h - s)\frac{1}{n_h - s} = 1$. For the remaining rows, similar to the AMC scheme, the helper nodes transmit the received messages without any aggregation to the master. Hence, the communication cost incurred at the helper nodes is equal to $(n_e - \sum_{i=1}^m M_i)(n_h - s)\frac{1}{n_h - s} = (n_e - \sum_{i=1}^m M_i)$. Hence, for a particular erasure matrix $\mathcal{M}_E$ the communication cost at the helper nodes for the scheme considering m maximas is

$$
n_e - \sum_{i=1}^m M_i + m.
$$

The normalized communication load is obtained by averaging over all erasure patterns. Hence, we have,

$$
C_{HM} = n_e - \sum_{i=1}^m \mathbb{E}[M_i] + 1.
$$

For a system with $n_e = 64$, $n_h = 6$ and $s = 1$, using Monte-Carlo technique, the C_{HM} values for various values of m are obtained and plotted in Fig. 5. It can be observed from Fig. 5 that as m increases, more and more aggregation opportunities arise and hence a reduction in C_{HM} is obtained. In the graph $m = 0$ indicates the naive scheme in which the helper nodes transmits all the received messages directly to the master node without any aggregation. We performed simulation to compute the normalized communication cost for different values of n_e, keeping $n_h = \lfloor \log(n_e) \rfloor$, and $s = 0.2 * n_h$. The communication cost at the helper nodes reduces with increasing values of m as evident from Fig. 6.

VII. A GREEDY APPROACH TO ARC

We begin with illustrating the ARC scheme with an example.

Example 1: Consider a system with $n_e = 4$ client nodes and $n_h = 6$ helper nodes with at most $s = 2$ straggling links between each client and helper nodes. The client node partitions the gradient vector into $k = \frac{n_h}{s+1} = 2$ parts and uses a generator matrix

$$
G \; = \;
\begin{bmatrix}
1 & 0 & 1 & 0 & 1 & 0 \\
0 & 1 & 0 & 1 & 0 & 1
\end{bmatrix}.
$$

Let the observed erasure matrix be

$$
\mathcal{M}_E \; = \;
\begin{bmatrix}
0 & 1 & 1 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 1 \\
1 & 0 & 1 & 0 & 0 & 0 \\
0 & 1 & 0 & 1 & 0 & 0
\end{bmatrix}.
$$

The observed erasure pattern is also depicted in Fig. 7. Following the aggregation strategy employed in ARC, it can be observed that all the helper nodes needs to transmit the messages and hence the communication cost at the helper nodes for this particular scenario is $s + 1 = 3$.
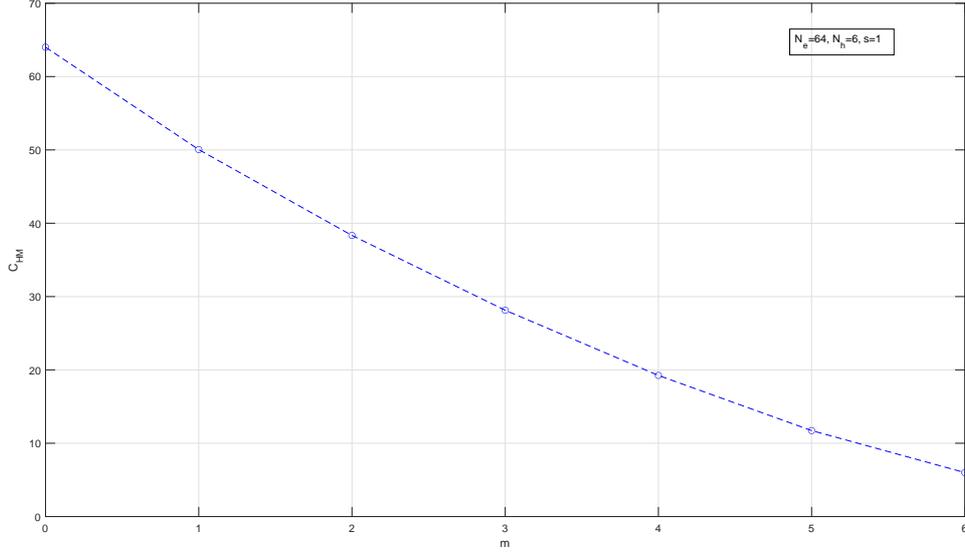
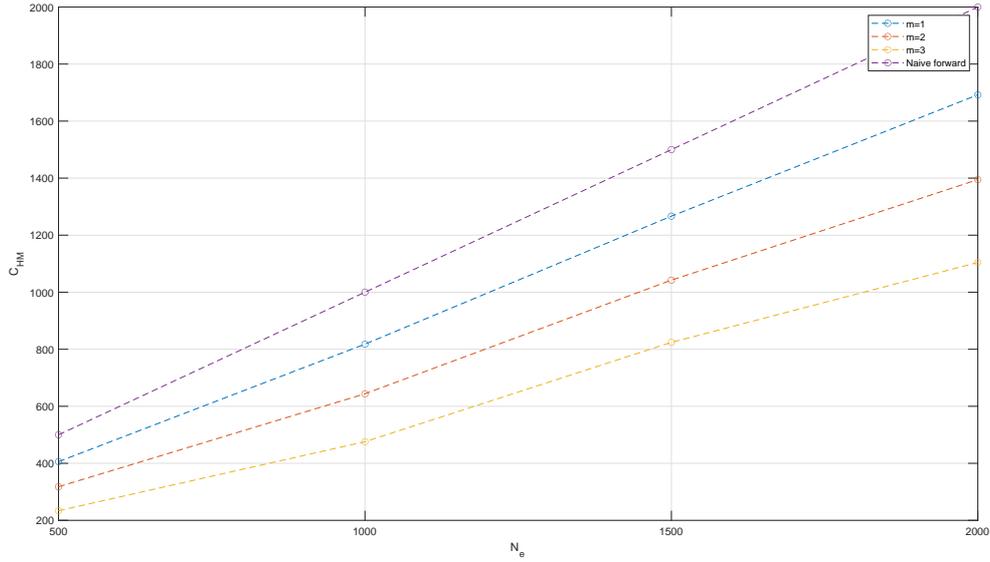Fig. 5: C_{HM} values corresponding to $m \in [6]$ for $(n_e, n_h, s) = (64, 6, 1)$.

Fig. 6: Comparison of variation of C_{HM} with n_e for different values of m.

The principal idea of ARC is to repeat the components of the computed gradient to helper nodes without coding. Therefore the gradient aggregation at the helper nodes for ARC can be viewed as a set cover problem [26]. Having recognized that, we can employ a greedy algorithm to solve it. We show by simulations the greedy algorithm performs better than the existing aggregation strategy followed in [18].

A. Greedy approach

The objective of the master is to compute the total gradient $\underline{g}_D = \left[\sum_{i=1}^{n_e} \underline{g}_{i,1} \ \sum_{i=1}^{n_e} \underline{g}_{i,2} \ \cdots \ \sum_{i=1}^{n_e} \underline{g}_{i,k} \right]$. Let $\mathcal{U} = \bigcup_{\substack{i \in [n_e] \\ j \in [k]}} \underline{g}_{i,j}$ denote a set of gradient vectors. Each helper node H_i receives a subset of partial gradients denoted by $\mathcal{S}_i$. The master nodes wants each element in $\mathcal{U}$ to be aggregated and transmitted by one helper node. The master can then combine the transmissions suitably to find the total gradient $\underline{g}_D$. The master node effectively wants to find a cover of the universal set

	H_1	H_2	H_3	H_4	H_5	H_6
E_1	g_{11}	g_{12}	g_{11}	g_{12}	g_{11}	g_{12}
E_2	g_{21}	g_{22}	g_{21}	g_{22}	g_{21}	g_{22}
E_3	g_{31}	g_{32}	g_{31}	g_{32}	g_{31}	g_{32}
E_4	g_{41}	g_{42}	g_{41}	g_{42}	g_{41}	g_{42}

Fig. 7: Aggregation strategy of the ARC scheme for $(n_e = 4, n_h = 6, s = 2)$.

of gradients $\mathcal{U}$ using minimum number of sets from $\mathcal{S}_i, i \in [n_h]$. The minimum number is desired since the communication cost at the helper nodes depends on the number of helper nodes participating in the aggregation phase. Hence we can model the aggregation phase of ARC to a set cover problem with universal set $\mathcal{U}$ and subsets $\mathcal{S}_i, i \in [n_h]$. The set covering problem is an NP-complete problem with a heuristic greedy algorithm [26]. The greedy algorithm for the aggregation phase in ARC scheme is described in Algorithm 1.

Algorithm 1 Greedy algorithm

1: Input: $\mathcal{U}$ and $\mathcal{S}_i, i \in [n_h]$
2: Initialize $U \leftarrow \mathcal{U}, C \leftarrow \phi$
3: **while** $U \neq \phi$ **do**
4:     Select an $i \in [n_h]$ such that $|\mathcal{S}_i \cap U|$ is maximum
5:     $\mathcal{A}_i = \mathcal{S}_i \cap U$
6:     $C \leftarrow C \cup i$
7:     $U \leftarrow U \setminus \mathcal{S}_i$
8: **end while**
9: For every $i \in C$, the master node asks helper node H_i to aggregate the partial gradients in $\mathcal{A}_i$ and transmit.

Example 2: For the system considered in Example 1 and the erasure patter as described in Fig. 7, the greedy algorithm outputs $C = \{5, 6, 2\}, \mathcal{A}_5 = \{\underline{g}_{1,1}, \underline{g}_{2,1}, \underline{g}_{3,1}, \underline{g}_{4,1}\}, \mathcal{A}_6 = \{\underline{g}_{1,2}, \underline{g}_{3,2}, \underline{g}_{4,2}\}$ and $\mathcal{A}_2 = \{\underline{g}_{2,2}\}$. The communication cost at the helper nodes for this particular scenario is equal to 1.5.

B. Simulation Results

The example described in Section VII-A considers only one particular erasure pattern. Recall that the normalized communication load $C_{HM} = \frac{1}{|\Omega(s)|} \sum_{\epsilon \in \Omega(s)} \sum_{j=1}^{n_h} \frac{m_{j,\epsilon}}{k}$. We consider a system with $n_e = 64$ client nodes and $s = 2$ straggling links per client node. The C_{HM} value for both greedy approach and the existing strategy is computed using Monte-Carlo technique by averaging over a large number of iterations. The performance improvement of the greedy approach over the existing ARC scheme can be observed from Fig. 8. The C_{HM} values are plotted for different n_h values.
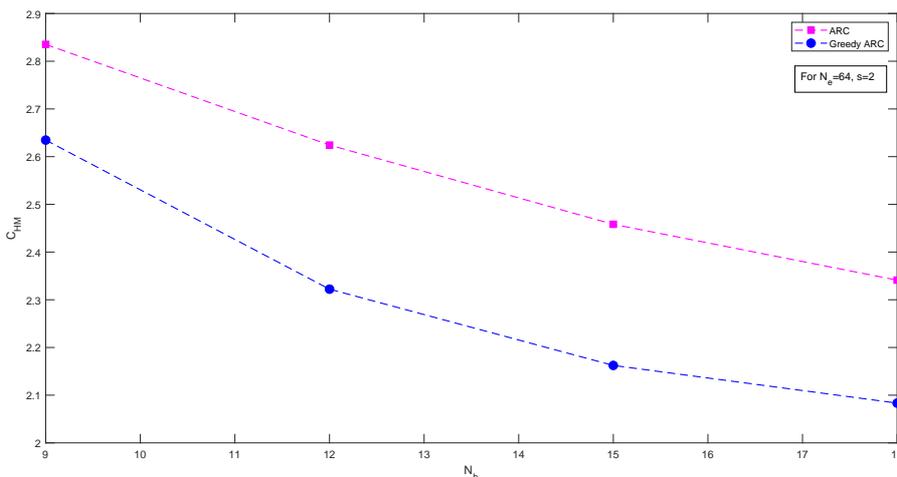

Fig. 8: Comparison of C_{HM} values of ARC scheme and the proposed greedy approach for $n_e = 64, s = 2$ and for different values of n_h

REFERENCES

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proc. of the 20th International Conference on Artificial Intelligence and Statistics*, vol. 54, Fort Lauderdale, FL, USA, 2017, pp. 1273–1282.

[2] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.

[3] A. Reisizadeh, A. Mokhtari, H. Hassani, A. Jadbabaie, and R. Pedarsani, "Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, vol. 108, 2020, pp. 2021–2031.

[4] L. Zhao, Q. Wang, Q. Zou, Y. Zhang, and Y. Chen, "Privacy-preserving collaborative deep learning with unreliable participants," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 1486–1500, 2020.

[5] A. Reisizadeh, H. Taheri, A. Mokhtari, H. Hassani, and R. Pedarsani, "Robust and communication-efficient collaborative learning," in *Advances in Neural Information Processing Systems*, vol. 32, 2019, pp. 8388–8399.

[6] Y. Liu, Y. Kang, X. Zhang, L. Li, Y. Cheng, T. Chen, M. Hong, and Q. Yang, "A communication efficient collaborative learning framework for distributed features," *arXiv preprint arXiv:1912.11187v6*, 2020.

[7] R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, Sydney, Australia, 2017, pp. 3368–3376.

[8] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80, Stockholm, Sweden, 2018, pp. 5610–5619.

[9] S. Li, S. M. M. Kalan, Q. Yu, M. Soltanolkotabi, and A. S. Avestimehr, "Polynomially coded regression: Optimal straggler mitigation via data encoding," *arXiv preprint arXiv:1805.09934v1*, 2018.

[10] 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 *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, vol. 84, 2018, pp. 803–812.

[11] S. Li, S. M. M. Kalan, A. S. Avestimehr, and M. Soltanolkotabi, "Near-optimal straggler mitigation for distributed gradient methods," in *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2018, pp. 857–866.

[12] W. Halbawi, N. Azizan, F. Salehi, and B. Hassibi, "Improving distributed gradient descent using reed-solomon codes," in *2018 IEEE International Symposium on Information Theory (ISIT)*, Vail, CO, USA, 2018, pp. 2027–2031.

[13] M. F. Aktaş and E. Soljanin, "Straggler mitigation at scale," *IEEE/ACM Transactions on Networking*, vol. 27, no. 6, pp. 2266–2279, 2019.

[14] A. Reisizadeh, S. Prakash, R. Pedarsani, and A. S. Avestimehr, "Tree gradient coding," in *2019 IEEE International Symposium on Information Theory (ISIT)*, Paris, France, 2019, pp. 2808–2812.

[15] E. Ozfatura, D. Gündüz, and S. Ulukus, "Speeding up distributed gradient descent by utilizing non-persistent stragglers," in *2019 IEEE International Symposium on Information Theory (ISIT)*, Paris, France, 2019, pp. 2729–2733.

[16] A. Badita, P. Parag, and V. Aggarwal, "Optimal server selection for straggler mitigation," *IEEE/ACM Transactions on Networking*, vol. 28, no. 2, pp. 709–721, 2020.

[17] S. Sasi, V. Lalitha, V. Aggarwal, and B. S. Rajan, "Straggler mitigation with tiered gradient codes," *IEEE Transactions on Communications*, vol. 68, no. 8, pp. 4632–4647, 2020.

[18] S. Prakash, A. Reisizadeh, R. Pedarsani, and A. S. Avestimehr, "Hierarchical coded gradient aggregation for learning at the edge," in *2020 IEEE International Symposium on Information Theory (ISIT)*, Los Angeles, CA, USA, 2020, pp. 2616–2621.

[19] C. Huang, M. Chen, and J. Li, "Pyramid codes: Flexible schemes to trade space for access efficiency in reliable data storage systems," in *Sixth IEEE International Symposium on Network Computing and Applications (NCA 2007)*, Cambridge, MA, 2007, pp. 79–86.

[20] M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 50, no. 1, pp. 30–39, 2017.

[21] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "Qsgd: Communication-efficient sgd via gradient quantization and encoding," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[22] J. Acharya, C. De Sa, D. Foster, and K. Sridharan, "Distributed learning with sublinear communication," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 97. PMLR, 09–15 Jun 2019, pp. 40–50.

[23] N. Shlezinger, M. Chen, Y. C. Eldar, H. V. Poor, and S. Cui, "Federated learning with quantization constraints," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 8851–8855.

[24] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, "On the locality of codeword symbols," *IEEE Transactions on Information theory*, vol. 58, no. 11, pp. 6925–6934, 2012.

[25] S. Kadhe and A. Sprintson, "Codes with unequal locality," in *2016 IEEE International Symposium on Information Theory (ISIT)*, Barcelona, 2016, pp. 435–439.

[26] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. The MIT Press, 2001.

APPENDIX A
MAXIMUM BIN-SIZE WITH BUNCHING

Suppose r unlabeled balls are thrown at n labeled bins uniformly at random. The first bm bins, $bm \leq n$ are *bunched* to form m artificial bins each consisting of b real bins. Now there are a total of $\nu = n - bm + m$ bins. Let random vector $\underline{X} = [X_1 \ X_2 \ \cdots \ X_\nu]$ denote the number of balls in each bin. When $k_1 + k_2 + \cdots + k_\nu = r$,

$$\Pr(\underline{X} = (k_1, k_2, \ldots, k_\nu)) = \frac{r!}{k_1! \, k_2! \cdots k_\nu!} \left(\frac{b}{n} \right)^{k_1 + k_2 + \cdots + k_m} \left(\frac{1}{n} \right)^{k_{m+1} + k_{m+2} + \cdots + k_\nu}$$

We define $M = \max_{i=1,2,\ldots,m} X_i$. We are interested in determining both $\rho(n, r, m, b) = \mathbb{E}[M]$ and $\phi(n, r, m, b) = \Pr[M \geq 1]$. Since M is a non-negative integer valued random variable,

$$
\begin{aligned}
\mathbb{E}[M] &= \sum_{i=1}^{r} \Pr(M \geq i) \\
&= \sum_{i=1}^{r} \Pr(\cup_{j=1}^{m} \{X_j \geq i\})
\end{aligned}
\tag{7}
$$

Define $E_{ij} = \{X_j \geq i\}$ and apply inclusion-exclusion principle to obtain $\rho(n, r, m, b) = \sum_{i=1}^{r} \sum_{\ell=1}^{m} A_{i\ell}$ where $A_{i\ell}$ is given in (8).

$$A_{i\ell} = (-1)^{\ell-1} \cdot \binom{m}{\ell} \cdot \sum_{\substack{k_1 \geq i, k_2 \geq i, \ldots, k_\ell \geq i \\ \sum_{i\in[\ell]} k_i \leq r}} \frac{r!}{k_1! \, k_2! \cdots k_\ell! \, (r - \sum_{i\in[\ell]} k_i)!} \left(\frac{b}{n}\right)^{\sum_{i\in[\ell]} k_i} \left(1 - \frac{\ell b}{n}\right)^{r - \left(\sum_{i\in[\ell]} k_i\right)} \tag{8}$$

Then it follows that

$$\rho(n, r, m, b) = \sum_{\ell=1}^{m} \sum_{i=1}^{\lfloor \frac{r}{\ell} \rfloor} A_{i\ell} \tag{9}$$

$$\phi(n, r, m, b) = \sum_{\ell=1}^{m} A_{1\ell} \tag{10}$$