

Computation Scheduling for Distributed Machine Learning with Straggling Workers

Mohammad Mohammadi Amiri and Deniz Gündüz

Abstract—We study scheduling of computation tasks across n workers in a large scale distributed learning problem with the help of a master. Computation and communication delays are assumed to be random, and redundant computations are assigned to workers in order to tolerate stragglers. We consider sequential computation of tasks assigned to a worker, while the result of each computation is sent to the master right after its completion. Each computation round, which can model an iteration of the stochastic gradient descent (SGD) algorithm, is *completed* once the master receives k distinct computations, referred to as the *computation target*. Our goal is to characterize the *average completion time* as a function of the *computation load*, which denotes the portion of the dataset available at each worker, and the *computation target*. We propose two computation scheduling schemes that specify the tasks assigned to each worker, as well as their computation schedule, i.e., the order of execution. Assuming a general statistical model for computation and communication delays, we derive the average completion time of the proposed schemes. We also establish a lower bound on the minimum average completion time by assuming prior knowledge of the random delays. Experimental results carried out on Amazon EC2 cluster show a significant reduction in the average completion time over existing coded and uncoded computing schemes. It is also shown numerically that the gap between the proposed scheme and the lower bound is relatively small, confirming the efficiency of the proposed scheduling design.

I. INTRODUCTION

The growing computational complexity and memory requirements of emerging machine learning applications involving massive datasets cannot be satisfied on a single machine. Thus, distributed computation across tens or even hundreds of computation servers, called *workers*, has been a topic of great recent interest [1], [2]. A major bottleneck in distributed computation is that the overall performance can significantly deteriorate due to slow servers, referred to as *stragglers*. To mitigate the limitation of stragglers, coded computation techniques, inspired by erasure codes against packet losses, have been proposed recently [3]–[8]. With coded computation, computations from only a subset of non-straggling workers are sufficient to complete the computation task, thanks to redundant computations performed by the faster workers. Coded distributed computation has also been studied for matrix-matrix multiplication [9]–[11], and distributed computing of a polynomial function [12], and linear regression problem [13].

Most existing coded computation techniques are designed to tolerate persistent stragglers, and discard computations performed by stragglers [4]–[13]. However, in practice we

often encounter *non-persistent stragglers*, which, despite being slower, complete a significant portion of the assigned tasks by the time faster workers complete all their tasks [14]. We argue in this paper that uncoded computing and communication can be even more effective in tackling stragglers and reducing the average computation time. We consider computation of an arbitrary function over a dataset, and introduce a centralized scheduling strategy for uncoded distributed computation, where the tasks are assigned to the workers by the *master*. Each worker has access to a limited number of data points, referred to as the *computation load*, and computes the tasks with respect to its local dataset. Computations are carried out sequentially, and the result of each computation is sent to the master right after it is completed. We assume that both the computation and communication delays are independent across workers, but may be correlated for different tasks carried out at the same worker. This sequential computation and communication framework allows the master to exploit partial computations by slow workers. The main computation is assumed to be completed when the master receives sufficient number of distinct computations from the workers, referred to as the *computation target*. In other words, computation target specifies the number of distinct computations the master is required to receive from the workers to complete the main computation task.

Assuming that the computation and communication delays are random variables, our goal is to characterize the minimum *average completion time*, which represents the time it takes the master to meet the computation target. We first provide a generic expression for the average completion time as a function of the *computation schedule*, which specifies both the tasks assigned to each worker and the order these tasks are computed at each worker. We propose two different computation scheduling schemes, and obtain closed-form expressions for their average completion times for a general statistical model of the random delays, which upper bound the minimum average completion time. We also establish a lower bound on the minimum average completion time. Experiments run on Amazon EC2 cluster illustrate a substantial reduction in the average completion time with the proposed uncoded computing schemes with task scheduling compared to coded computation schemes or uncoded computation without scheduling [15]. The improvement of the proposed uncoded schemes over the coded ones is due to the exploitation of partial computations performed by the slow workers, which is more beneficial when the variability of the delays at different workers is small. Also, the gain of the proposed schemes compared to the random scheduling approach studied in [15] is due to the proposed computation scheduling, which aims to minimize redundant

The authors are with the Information Processing and Communications Laboratory (IPC-Lab), Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2AZ, U.K. (e-mail: m.mohammadi-amiri15@imperial.ac.uk; d.gunduz@imperial.ac.uk).

This work has been supported by the European Research Council (ERC) through Starting Grant BEACON (agreement No. 677854).

repeated computations.

A. Related Works

While significant research efforts have been invested in designing coded computation techniques [4]–[13], recently, there have been efforts to exploit the computations carried out by non-persistent stragglers [14]–[18]. Techniques studied in [14], [16]–[18] are based on coding with associated encoding and decoding complexities, which require the availability and processing of all the data points at the master. In [18] a linear regression problem is studied, and the scheme in [13] is extended by allowing each worker to communicate multiple computations sequentially, where the computations are carried out using coded data. The authors in [14] propose to split the computation tasks into multiple levels, and code each level using maximum-distance separable (MDS) coding with the knowledge of the statistical behavior of the stragglers. Distributed matrix-vector multiplication is studied in [16], where random coding across the dataset is performed, and each worker completes the assigned tasks sequentially. To compute functions which are linear with respect to their arguments, e.g., matrix-vector multiplication, rateless codes are used in [17], requiring a large number of data points assigned to each worker to guarantee decodability of the function at the master.

Unlike coded computation, uncoded computing approach does not introduce any encoding and decoding delays and complexities; hence, can be particularly efficient for edge learning where the data is inherently distributed [19]–[21]. It also allows partial decoding, which can be exploited to reduce the communication load for distributed learning [22]–[24]. An uncoded computation approach is considered in [15], where the dataset is split into a limited number of mini-batches, and each worker is randomly assigned a mini-batch. This approach requires a large number of workers compared to the number of mini-batches to ensure that the master can recover all the data from the workers with high probability. The authors in [25] study dynamic computation allocation across the workers with feedback providing information about the workers' speeds. The proposed uncoded computation approach in this paper does not impose any constraint on the number of workers, and is designed without any prior knowledge or feedback on the computation and communication delays at the workers.

The problem under consideration is similar to the job scheduling problem [26]–[33], in which a set of tasks are to be executed by multiple workers given a partial ordering of task execution and the delay associated with each task. The goal is to find a schedule minimizing the total delay, which is shown to be NP-complete [34]. Our problem differs from the job scheduling one, since no ordering of task execution is imposed, and each task can be executed by an arbitrary number of workers. Also, in our model, no prior knowledge about the computation and communication delays is available.

B. Organization

The organization of the paper is as follows. We present the system model in Section II. In Section III, we analyze the performance of the minimum average completion time for the general case. We provide an upper and a lower bound

on the minimum average completion time in Section IV and Section V, respectively. In Section VI, we overview some of the alternative approaches in the literature, and compare their performances with the proposed uncoded schemes numerically. Finally, the paper is concluded in Section VII.

Notations: We represent sets of real values and integers by \mathbb{R} and \mathbb{Z} , respectively. For $i, j \in \mathbb{Z}$, $j \geq i$, $[i : j]$ denotes set $\{i, i + 1, \dots, j\}$. For $i \in \mathbb{Z}^+$, we define $[i] \triangleq [1 : i]$. $\mathcal{N}(0, \sigma^2)$ denotes a zero-mean normal distribution with variance σ^2 , and, for $a, b \in \mathbb{R}$, $\mathcal{U}(a, b)$ denotes a uniform distributed over $[a, b]$. $A(i, j)$ represents (i, j) -th entry of matrix A .

II. SYSTEM MODEL

We consider distributed computation of a function h over a dataset $\mathcal{X} = \{X_1, \dots, X_n\}$, which consists of n data points, across n workers. Function $h : \mathbb{V} \rightarrow \mathbb{U}$ is an arbitrary function, where \mathbb{V} and \mathbb{U} are two vector spaces over the same field \mathbb{F} , and data point X_i is an element of \mathbb{V} , $i \in [n]$. The dataset \mathcal{X} is distributed across the workers by the master, and a maximum number of $r \leq n$ data points are assigned to each worker, referred to as the *computation load*. We denote by \mathcal{E}_i the indices of the data points assigned to worker i , $i \in [n]$, where $\mathcal{E}_i \subset [n]$, $|\mathcal{E}_i| \leq r$.

The computations of the tasks assigned to each worker are carried out sequentially. We define the task ordering (TO) matrix C as an $n \times r$ matrix of integers, $C \in [n]^{n \times r}$, specifying the assignment of the tasks to the workers $\mathcal{E} \triangleq \{\mathcal{E}_i\}_{i=1}^n$, as well as the order these tasks are carried out by each worker $\mathcal{O} \triangleq \{\mathcal{O}_i\}_{i=1}^n$, where \mathcal{O}_i denotes the computing order of the tasks assigned to worker i . Each row of matrix C corresponds to a different worker, and its elements from left to right represent the order of computations. That is, the entry $C(i, j) \in \mathcal{E}_i$ denotes the index of the element of the dataset that is computed by worker i as its j -th computation, i.e., worker i first computes $h(X_{C(i,1)})$, then computes $h(X_{C(i,2)})$, and so on so forth until either it computes $h(X_{C(i,r)})$, or it receives the acknowledgement message from the master, and stops computations, $i \in [n]$, $j \in [r]$. Note that the task assignment \mathcal{E} and the order of computations \mathcal{O} are specified by a unique TO matrix C , and, in general, a data point can be assigned to different workers.

All the workers start at time $t = 0$, and each worker sends the result of each assigned task to the master right after its computation. We denote the time worker i spends to compute $h(X_j)$ by $T_{i,j}^{(1)}$, and the communication delay for sending $h(X_j)$ to the master by $T_{i,j}^{(2)}$, $j \in \mathcal{E}_i$, $i \in [n]$. Thus, the total delay of receiving $h(X_j)$ from worker i is $T_{i,j}^{(1)} + T_{i,j}^{(2)}$, $j \in \mathcal{E}_i$, $i \in [n]$. If $j \notin \mathcal{E}_i$, we set $T_{i,j}^{(l)} = \infty$, $\forall l \in [2]$, $i \in [n]$. We assume that the computation and communication delays, $T_{i,j}^{(1)}$ and $T_{i,j}^{(2)}$, $\forall i, j \in [n]$, are independent. We further assume that computation (communication) delays at different workers are independent. On the other hand, the computation (communication) delays associated with the tasks at the same worker can be dependent, and we denote the joint cumulative distribution function (CDF) of $T_{i,1}^{(l)}, \dots, T_{i,n}^{(l)}$ by $F_{i,[n]}^{(l)}$, and the joint probability density function (PDF) by $f_{i,[n]}^{(l)}$, $i \in [n]$, $l \in [2]$. We note that the statistical model of the computation

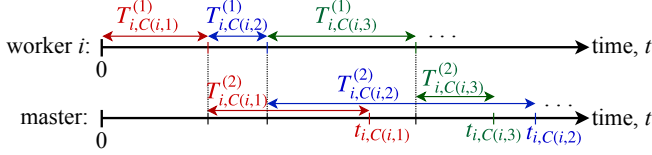


Fig. 1: Illustration of the computation and communication delays for the computations carried out by worker i .

(communication) delays at each worker does not depend on any specific order of computing (communicating) tasks, since we assume that the size and complexity of computing (communicating) each data point (computation) are the same.

Fig. 1 illustrates a graphical representation of a realization of the computation and communication delays from worker i to the master. Let $t_{i,j}$ denote the time the master receives $h(X_j)$ from worker i , for $i, j \in [n]$, where we set $t_{i,m} = \infty$ if $m \notin \mathcal{E}_i$. Then, the total computation delay of computing $h(X_{C(i,1)}), h(X_{C(i,2)}), \dots, h(X_{C(i,j)})$ sequentially plus the communication delay for receiving $h(X_{C(i,j)})$ is

$$t_{i,C(i,j)} = \sum_{m=1}^j T_{i,C(i,m)}^{(1)} + T_{i,C(i,j)}^{(2)}, \quad i, j \in [n], \quad (1)$$

As a result, the master receives computation $h(X_j)$ at time

$$t_j \triangleq \min_{i \in [n]} \{t_{i,j}\}, \quad j \in [n] \quad (2)$$

where the minimization is over the workers.

Example 1. Consider the TO matrix C for $n = 4$ and $r = 3$:

$$C = \begin{bmatrix} 1 & 2 & 3 \\ 3 & 2 & 1 \\ 3 & 4 & 1 \\ 4 & 3 & 1 \end{bmatrix}, \quad (3)$$

which dictates the following computation schedule:

- Worker 1 first computes $h(X_1)$, then $h(X_2)$, and $h(X_3)$.
- Worker 2 first computes $h(X_3)$, then $h(X_2)$, and $h(X_1)$.
- Worker 3 first computes $h(X_3)$, then $h(X_4)$, and $h(X_1)$.
- Worker 4 first computes $h(X_4)$, then $h(X_3)$, and $h(X_1)$.

Each worker sends the result of each computation to the master immediately after its completion. Accordingly, we have

$$t_{1,1} = T_{1,1}^{(1)} + T_{1,1}^{(2)}, \quad t_{1,2} = T_{1,1}^{(1)} + T_{1,2}^{(1)} + T_{1,2}^{(2)}, \quad (4a)$$

$$t_{1,3} = T_{1,1}^{(1)} + T_{1,2}^{(1)} + T_{1,3}^{(1)} + T_{1,3}^{(2)}, \quad t_{1,4} = \infty, \quad (4a)$$

$$t_{2,3} = T_{2,3}^{(1)} + T_{2,3}^{(2)}, \quad t_{2,2} = T_{2,3}^{(1)} + T_{2,2}^{(1)} + T_{2,2}^{(2)}, \quad (4b)$$

$$t_{2,1} = T_{2,3}^{(1)} + T_{2,2}^{(1)} + T_{2,1}^{(1)} + T_{2,1}^{(2)}, \quad t_{2,4} = \infty, \quad (4b)$$

$$t_{3,3} = T_{3,3}^{(1)} + T_{3,3}^{(2)}, \quad t_{3,4} = T_{3,3}^{(1)} + T_{3,4}^{(1)} + T_{3,4}^{(2)}, \quad (4c)$$

$$t_{3,1} = T_{3,3}^{(1)} + T_{3,4}^{(1)} + T_{3,1}^{(1)} + T_{3,1}^{(2)}, \quad t_{3,2} = \infty, \quad (4c)$$

$$t_{4,4} = T_{4,4}^{(1)} + T_{4,4}^{(2)}, \quad t_{4,3} = T_{4,4}^{(1)} + T_{4,3}^{(1)} + T_{4,3}^{(2)}, \quad (4d)$$

$$t_{4,1} = T_{4,4}^{(1)} + T_{4,3}^{(1)} + T_{4,1}^{(1)} + T_{4,1}^{(2)}, \quad t_{4,2} = \infty. \quad (4d)$$

□

For any TO matrix, the computation is considered completed once the master recovers k distinct tasks, referred to as the *computation target*. We allow partial computations, i.e., k can be smaller than n . Once the computation target is met, the master sends an acknowledgement message to all the workers

to stop computations. Given the TO matrix C , we denote the *completion time*; that is, the time it takes the master to receive k distinct computations, by $t_C(r, k)$, which is a random variable. We define the average completion time as

$$\bar{t}_C(r, k) \triangleq \mathbb{E}[t_C(r, k)], \quad (5)$$

where the randomness is due to the delays. We define the minimum average completion time

$$\bar{t}^*(r, k) \triangleq \min_C \{\bar{t}_C(r, k)\}, \quad (6)$$

where the minimization is over all TO matrices C . The goal is to characterize $\bar{t}^*(r, k)$. While any C matrix is a valid TO matrix, it is easy to see that the optimal TO matrix will have r distinct entries in each row, and k distinct entries overall.

Remark 1. We have defined each $X_i \in \mathbb{V}$ as a single data point, and assumed that the result of $h(X_i)$ at a worker is transmitted immediately to the master. It is possible to generalize this model by considering N data points instead, with $N \gg n$, and grouping them into n mini-batches, such that each X_i in our model corresponds to a mini-batch of $\lceil N/n \rceil$ data points. A worker sends the average of the gradients for all the data points in a mini-batch after computing all of them. For a mini-batch size of c data points, this corresponds to communicating once every c computations.

Remark 2. Most coded computation schemes in the literature, mainly targeting DGD, require the master to recover the gradients (or, their average) for the whole dataset at each iteration. However, convergence of stochastic gradient descent (SGD) is guaranteed even if the gradient is computed for a random portion of the dataset at each iteration [22], [24], [35]–[40]. This is indeed the case for the random straggling model considered here with $k < n$, where the straggling workers; hence, the uncomputed gradients, vary at each iteration.

Remark 3. When $k < n$, in order to prevent bias in the SGD algorithm, we need to make sure that the first k distinct computations received by the master are uniformly random across the mini-batches. If a few workers are significantly faster than the others, we may end up receiving computations corresponding to a few batches assigned to these workers. Alternatively, we can periodically re-index the mini-batches and their corresponding labels randomly after a fixed number of iterations, and provide the workers with the new mini-batches while the TO matrix is fixed. This introduces additional communication from the master to the workers to deliver the missing mini-batches after re-indexing.

III. AVERAGE COMPLETION TIME ANALYSIS

Here we analyze the average completion time $\bar{t}_C(r, k)$ for a given TO matrix C .

Theorem 1. For a given TO matrix C , we have

$$\begin{aligned} \Pr\{t_C(r, k) > t\} &= 1 - F_{t_C}(t) \\ &= \sum_{i=n-k+1}^n (-1)^{n-k+i+1} \binom{i-1}{n-k} \\ &\quad \sum_{\mathcal{S} \subset [n]: |\mathcal{S}|=i} \Pr\{t_j > t, \forall j \in \mathcal{S}\}, \quad (7) \end{aligned}$$

which yields

$$\bar{t}_C(r, k) = \sum_{i=n-k+1}^n (-1)^{n-k+i+1} \binom{i-1}{n-k} \sum_{\mathcal{S} \subset [n]: |\mathcal{S}|=i} \int_0^\infty \Pr \{t_j > t, \forall j \in \mathcal{S}\} dt. \quad (8)$$

Proof. See Appendix A. \square

Theorem 1 provides an exact expression for the average delay of receiving k distinct computations at the master for a given TO matrix C . Note that the dependence of the average completion time on the TO matrix is through the statistics of t_j . Hence, for any TO matrix C , the exact average completion time can be obtained using the joint statistics of t_1, \dots, t_n .

Remark 4. For $k = n$, we have

$$\Pr \{t_C(r, n) > t\} = \sum_{i=1}^n (-1)^{i+1} \sum_{\mathcal{S} \subset [n]: |\mathcal{S}|=i} \Pr \{t_j > t, \forall j \in \mathcal{S}\}, \quad (9)$$

and

$$\bar{t}_C(r, n) = \sum_{i=1}^n (-1)^{i+1} \sum_{\mathcal{S} \subset [n]: |\mathcal{S}|=i} \int_0^\infty \Pr \{t_j > t, \forall j \in \mathcal{S}\} dt. \quad (10)$$

The minimum average completion time $\bar{t}^*(r, k)$ can be obtained as a solution of the optimization problem $\bar{t}^*(r, k) = \min_C \bar{t}_C(r, k)$. Providing a general characterization for $\bar{t}^*(r, k)$ is elusive. Next, we will propose two scheduling schemes, and evaluate their average completion times.

IV. UPPER BOUNDS ON THE MINIMUM AVERAGE COMPLETION TIME

In this section we introduce two computation task assignment and scheduling schemes, namely *cyclic scheduling* (CS) and *staircase scheduling* (SS). The average completion time for these schemes will provide upper bounds on $\bar{t}^*(r, k)$.

A. Cyclic Scheduling (CS) Scheme

The CS scheme is motivated by the symmetry across the workers when we have no prior information on their computation speeds. CS makes sure that each computation task has the same order at different workers. This is achieved by a cyclic shift operator. The TO matrix is given by

$$C_{CS}(i, j) = g(i + j - 1), \quad \text{for } i \in [n] \text{ and } j \in [r], \quad (11)$$

where function $g: \mathbb{Z} \rightarrow \mathbb{Z}$ is defined as follows:

$$g(m) \triangleq \begin{cases} m, & \text{if } 1 \leq m \leq n, \\ m - n, & \text{if } m \geq n + 1, \\ m + n, & \text{if } m \leq 0. \end{cases} \quad (12)$$

Thus, we have

$$C_{CS} = \begin{bmatrix} g(1) & g(2) & \dots & g(r) \\ g(2) & g(3) & \dots & g(r+1) \\ \vdots & \vdots & \ddots & \vdots \\ g(n) & g(n+1) & \dots & g(n+r-1) \end{bmatrix}, \quad (13)$$

which, for $i \in [n]$ and $j \in [r]$, results in

$$t_{i, g(i+j-1)} = \sum_{m=1}^j T_{i, g(i+m-1)}^{(1)} + T_{i, g(i+j-1)}^{(2)}. \quad (14)$$

For $i \in [n]$, we can re-write (14) as follows:

$$t_{g(i-j+1), i} = \begin{cases} \sum_{m=1}^j T_{g(i-j+1), g(i-j+m)}^{(1)} + T_{g(i-j+1), i}^{(2)}, & \text{if } j \in [r], \\ \infty, & \text{if } j \notin [r], \end{cases} \quad (15)$$

which results in

$$t_i = \min_{j \in [r]} \left\{ \sum_{m=1}^j T_{g(i-j+1), g(i-j+m)}^{(1)} + T_{g(i-j+1), i}^{(2)} \right\}. \quad (16)$$

Example 2. Consider $n = 4$ and $r = 3$. We have

$$C_{CS} = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \\ 3 & 4 & 1 \\ 4 & 1 & 2 \end{bmatrix}, \quad (17)$$

and

$$t_{1,1} = T_{1,1}^{(1)} + T_{1,1}^{(2)}, \quad t_{1,2} = T_{1,1}^{(1)} + T_{1,2}^{(1)} + T_{1,2}^{(2)}, \quad (18a)$$

$$t_{1,3} = T_{1,1}^{(1)} + T_{1,2}^{(1)} + T_{1,3}^{(1)} + T_{1,3}^{(2)}, \quad t_{1,4} = \infty,$$

$$t_{2,2} = T_{2,2}^{(1)} + T_{2,2}^{(2)}, \quad t_{2,3} = T_{2,2}^{(1)} + T_{2,3}^{(1)} + T_{2,3}^{(2)}, \quad (18b)$$

$$t_{2,4} = T_{2,2}^{(1)} + T_{2,3}^{(1)} + T_{2,4}^{(1)} + T_{2,4}^{(2)}, \quad t_{2,1} = \infty,$$

$$t_{3,3} = T_{3,3}^{(1)} + T_{3,3}^{(2)}, \quad t_{3,4} = T_{3,3}^{(1)} + T_{3,4}^{(1)} + T_{3,4}^{(2)}, \quad (18c)$$

$$t_{3,1} = T_{3,3}^{(1)} + T_{3,4}^{(1)} + T_{3,1}^{(1)} + T_{3,1}^{(2)}, \quad t_{3,2} = \infty,$$

$$t_{4,4} = T_{4,4}^{(1)} + T_{4,4}^{(2)}, \quad t_{4,1} = T_{4,4}^{(1)} + T_{4,1}^{(1)} + T_{4,1}^{(2)}, \quad (18d)$$

$$t_{4,2} = T_{4,4}^{(1)} + T_{4,1}^{(1)} + T_{4,2}^{(1)} + T_{4,2}^{(2)}, \quad t_{4,3} = \infty, \quad \square$$

B. Staircase Scheduling (SS) Scheme

We can observe that CS imposes the same step size and direction in computations across all the workers. Alternatively, here we propose the SS scheme, which introduces inverse computation orders at the workers. The entries of the TO matrix C_{SS} for the SS scheme are given by, for $i \in [n]$, $j \in [r]$,

$$C_{SS}(i, j) = g(i + (-1)^{i-1}(j-1)). \quad (19)$$

It follows that

$$C_{SS} = \begin{bmatrix} g(1) & g(2) & \dots & g(r) \\ g(2) & g(1) & \dots & g(3-r) \\ \vdots & \vdots & \ddots & \vdots \\ g(n) & g(n+(-1)^{n-1}) & \dots & g(n+(-1)^{n-1}(r-1)) \end{bmatrix}, \quad (20)$$

which, for $i \in [n]$ and $j \in [r]$, results in

$$t_{i, g(i+(-1)^{i-1}(j-1))} = \sum_{m=1}^j T_{i, g(i+(-1)^{i-1}(m-1))}^{(1)} + T_{i, g(i+(-1)^{i-1}(j-1))}^{(2)}. \quad (21)$$

For $i \in [n]$, we can re-write (21) as follows:

$$t_{g(i+(-1)^{i+j-1}(j-1)),i} = \begin{cases} \sum_{m=1}^j T_{g(i+(-1)^{i+j-1}(j-1)),g(i+(-1)^{i+j-1}(j-m))}^{(1)} \\ \quad + T_{g(i+(-1)^{i+j-1}(j-1)),i}^{(2)}, & \text{if } j \in [r], \\ \infty, & \text{if } j \notin [r], \end{cases} \quad (22)$$

which results in

$$t_i = \min_{j \in [r]} \left\{ \sum_{m=1}^j T_{g(i+(-1)^{i+j-1}(j-1)),g(i+(-1)^{i+j-1}(j-m))}^{(1)} + T_{g(i+(-1)^{i+j-1}(j-1)),i}^{(2)} \right\}. \quad (23)$$

Example 3. Consider $n = 4$ and $r = 3$. We have

$$C_{SS} = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 1 & 4 \\ 3 & 4 & 1 \\ 4 & 3 & 2 \end{bmatrix}, \quad (24)$$

and

$$t_{1,1} = T_{1,1}^{(1)} + T_{1,1}^{(2)}, \quad t_{1,2} = T_{1,1}^{(1)} + T_{1,2}^{(1)} + T_{1,2}^{(2)}, \quad (25a)$$

$$t_{1,3} = T_{1,1}^{(1)} + T_{1,2}^{(1)} + T_{1,3}^{(1)} + T_{1,3}^{(2)}, \quad t_{1,4} = \infty, \quad (25a)$$

$$t_{2,2} = T_{2,2}^{(1)} + T_{2,2}^{(2)}, \quad t_{2,1} = T_{2,2}^{(1)} + T_{2,1}^{(1)} + T_{2,1}^{(2)}, \quad (25b)$$

$$t_{2,4} = T_{2,2}^{(1)} + T_{2,1}^{(1)} + T_{2,4}^{(1)} + T_{2,4}^{(2)}, \quad t_{2,3} = \infty, \quad (25b)$$

$$t_{3,3} = T_{3,3}^{(1)} + T_{3,3}^{(2)}, \quad t_{3,4} = T_{3,3}^{(1)} + T_{3,4}^{(1)} + T_{3,4}^{(2)}, \quad (25c)$$

$$t_{3,1} = T_{3,3}^{(1)} + T_{3,4}^{(1)} + T_{3,1}^{(1)} + T_{3,1}^{(2)}, \quad t_{3,2} = \infty, \quad (25c)$$

$$t_{4,4} = T_{4,4}^{(1)} + T_{4,4}^{(2)}, \quad t_{4,3} = T_{4,4}^{(1)} + T_{4,3}^{(1)} + T_{4,3}^{(2)}, \quad (25d)$$

$$t_{4,2} = T_{4,4}^{(1)} + T_{4,3}^{(1)} + T_{4,2}^{(1)} + T_{4,2}^{(2)}, \quad t_{4,1} = \infty, \quad (25d)$$

□

Remark 5. The main difference between CS and SS is that with CS all the workers have the same step size and direction in their computations, while with SS workers with even and odd indices have different directions (ascending and descending, respectively) in the order they carry out the computations assigned to them, but the same step size in their evaluations.

We highlight that the CS and SS schemes may not be the optimal schedules for certain realizations of the straggling behaviour, but our interest is in the average performance. We will see in Section VI that both perform reasonably well, and neither scheme outperforms the other at all settings.

C. Average Completion Time Analysis

Here we analyze the performance of CS and SS providing upper bounds on $\bar{t}^*(r, k)$. We represent the average completion time of CS and SS by $\bar{t}_{CS}(r, k)$ and $\bar{t}_{SS}(r, k)$, respectively. In order to characterize these average values through (8), we need to obtain $H_{S,\emptyset} = \Pr\{t_i > t, \forall i \in S\}$, for any set $S \subset [n]$, $n - k + 1 \leq |S| \leq n$, where t_1, \dots, t_n are given in (16) and (23), for CS and SS, respectively. For ease of presentation, we denote $H_{S,\emptyset}$ for CS and SS by $H_{S,\emptyset}^{CS}$ and $H_{S,\emptyset}^{SS}$, respectively.

For $S \subset [n]$ with $n - k + 1 \leq |S| \leq n$, we have

$$H_{S,\emptyset}^{CS} = \Pr \left\{ \sum_{m=1}^j T_{g(i-j+1),g(i-j+m)}^{(1)} + T_{g(i-j+1),i}^{(2)} > t, \right. \\ \left. \forall j \in [r], \forall i \in S \right\} = \Pr \{ \mathcal{T}_S^{CS}(t) \}, \quad (26)$$

where we define

$$\mathcal{T}_S^{CS}(t) \triangleq \left\{ \left(T_{1,1}^{(1)}, \dots, T_{n,n}^{(1)}, T_{1,1}^{(2)}, \dots, T_{n,n}^{(2)} \right) : \right. \\ \left. \sum_{m=1}^j T_{g(i-j+1),g(i-j+m)}^{(1)} + T_{g(i-j+1),i}^{(2)} > t, \forall j \in [r], \forall i \in S \right\}. \quad (27)$$

Similarly, for any set $S \subset [n]$, $n - k + 1 \leq |S| \leq n$, we have

$$H_{S,\emptyset}^{SS} = \Pr \left\{ \sum_{m=1}^j T_{g(i+(-1)^{i+j-1}(j-1)),g(i+(-1)^{i+j-1}(j-m))}^{(1)} \right. \\ \left. + T_{g(i+(-1)^{i+j-1}(j-1)),i}^{(2)} > t, \forall j \in [r], \forall i \in S \right\} = \Pr \{ \mathcal{T}_S^{SS}(t) \} \quad (28)$$

where we define

$$\mathcal{T}_S^{SS}(t) \triangleq \left\{ \left(T_{1,1}^{(1)}, \dots, T_{n,n}^{(1)}, T_{1,1}^{(2)}, \dots, T_{n,n}^{(2)} \right) : \right. \\ \left. \sum_{m=1}^j T_{g(i+(-1)^{i+j-1}(j-1)),g(i+(-1)^{i+j-1}(j-m))}^{(1)} \right. \\ \left. + T_{g(i+(-1)^{i+j-1}(j-1)),i}^{(2)} > t, \forall j \in [r], \forall i \in S \right\}. \quad (29)$$

It follows that, for $X \in \{CS, SS\}$,

$$H_{S,\emptyset}^X = \int \cdots \int_{\mathcal{T}_S^X(t)} f_{1,[r]}^{(1)}(\alpha_1^{(1)}) \cdots f_{n,[r]}^{(1)}(\alpha_n^{(1)}) f_{1,[r]}^{(2)}(\alpha_1^{(2)}) \cdots f_{n,[r]}^{(2)}(\alpha_n^{(2)}) d\alpha_1^{(1)} \cdots d\alpha_n^{(1)} d\alpha_1^{(2)} \cdots d\alpha_n^{(2)}. \quad (30)$$

By plugging (30) into (8), we can obtain, for $X \in \{CS, SS\}$,

$$\bar{t}_X(r, k) = \sum_{i=n-k+1}^n (-1)^{n-k+i+1} \binom{i-1}{n-k} \int_0^\infty H_{S,\emptyset}^X dt. \quad (31)$$

Note that we have obtained a general characterization of the average completion time of CS and SS in terms of the CDFs of the delays associated with different tasks at different workers. The numerical evaluation of the performances of CS and SS and the lower bound will be presented in Section VI.

V. LOWER BOUND

Here we present a lower bound on $\bar{t}^*(r, k)$ by considering an adaptive model. Note that the TO matrix, in general, may depend on the statistics of the computation and communication delays, i.e., $F_{i,[n]}^{(l)}$, $\forall l \in [2]$, but not on the realization of $T_{i,C(i,j)}^{(l)}$, $i \in [n]$, $j \in [r]$. Let $\hat{T}_{i,j}^{(1)}$ and $\hat{T}_{i,j}^{(2)}$, respectively, represent the computation and communication delays associated with the task worker i executes with its j -th computation, $i \in [n]$, $j \in [r]$. We note that $\hat{T}_{i,j}^{(l)}$ is a random variable




Fig. 2: Illustration of the arrival times of computations from the workers to the master.

independent of the TO matrix. We define

$$\mathbf{T} \triangleq \left(\hat{T}_{1,1}^{(1)}, \hat{T}_{1,1}^{(2)}, \dots, \hat{T}_{1,r}^{(1)}, \hat{T}_{1,r}^{(2)}, \dots, \hat{T}_{n,1}^{(1)}, \hat{T}_{n,1}^{(2)}, \dots, \hat{T}_{n,r}^{(1)}, \hat{T}_{n,r}^{(2)} \right). \quad (32)$$

For each realization of \mathbf{T} , we allow the master to employ a distinct TO matrix $C_{\mathbf{T}}$, and denote the completion time by $t_{C_{\mathbf{T}}}(\mathbf{T}, r, k)$, which is a random variable due to the randomness of \mathbf{T} . We define

$$t_{\text{LB}}(\mathbf{T}, r, k) \triangleq \min_{C_{\mathbf{T}}} \{t_{C_{\mathbf{T}}}(\mathbf{T}, r, k)\}, \quad (33)$$

and

$$\bar{t}_{\text{LB}}(r, k) \triangleq \mathbb{E}[t_{\text{LB}}(\mathbf{T}, r, k)], \quad (34)$$

where the expectation is taken over \mathbf{T} . It is easy to verify that

$$\begin{aligned} \bar{t}^*(r, k) &= \min_C \{\mathbb{E}[t_C(r, k)]\} \\ &\geq \mathbb{E}[\min_{C_{\mathbf{T}}} \{t_{C_{\mathbf{T}}}(r, k)\}] = \bar{t}_{\text{LB}}(r, k). \end{aligned} \quad (35)$$

Remark 6. We remark that \mathbf{T} does not depend on any specific order of computing and communicating tasks; that is, any realization of $(\hat{T}_{i,1}^{(1)}, \hat{T}_{i,1}^{(2)}, \dots, \hat{T}_{i,r}^{(1)}, \hat{T}_{i,r}^{(2)})$, i.e., the delays at worker i , is independent of any specific value $C(i, j)$ (index of the j -th task worker i computes), for $i \in [n], j \in [r]$. This is because we assume that the size and complexity associated with the computation of each data point are the same.

We denote by $\hat{t}_{i,j}$ the time at which the master receives the task computed by worker i with its j -th computation, $i \in [n], j \in [r]$. It follows that

$$\hat{t}_{i,j} = \sum_{l=1}^j \hat{T}_{i,l}^{(1)} + \hat{T}_{i,j}^{(2)}. \quad (36)$$

Fig. 2 illustrates the time instances computations from the workers are received by the master. For a realization of \mathbf{T} , $t_{\text{LB}}(\mathbf{T}, r, k)$ is the k -th order statistics of $\{\hat{t}_{1,1}, \dots, \hat{t}_{1,r}, \dots, \hat{t}_{n,1}, \dots, \hat{t}_{n,r}\}$, i.e., the k -th smallest value among $\{\hat{t}_{1,1}, \dots, \hat{t}_{1,r}, \dots, \hat{t}_{n,1}, \dots, \hat{t}_{n,r}\}$, denoted by $\hat{t}_{\mathbf{T},(k)}$. To prove that $t_{\text{LB}}(\mathbf{T}, r, k) = \hat{t}_{\mathbf{T},(k)}$, we note that $t_{\text{LB}}(\mathbf{T}, r, k)$ cannot be smaller than $\hat{t}_{\mathbf{T},(k)}$, since, according to the definition, for any time before $\hat{t}_{\mathbf{T},(k)}$ master has not received k computations. Also, since master receives the k -th computation exactly at time $\hat{t}_{\mathbf{T},(k)}$, knowing the realization of \mathbf{T} , one can design the TO matrix $C_{\mathbf{T}}$ such that the first k computations received by the master are all distinct. Since finding the statistics of $\hat{t}_{\mathbf{T},(k)}$ is analytically elusive, we obtain the lower bound on $\bar{t}^*(r, k)$ through Monte Carlo simulations.

VI. PERFORMANCE COMPARISONS

In this section, we evaluate the average completion time of the proposed CS and SS schemes, and compare them with different results in the literature. We will focus on distributed linear regression as the reference scenario.

A. Problem Scenario

We would like to compare the performance of the proposed uncoded computation schemes with coded computation techniques that have received significant interest in recent years. We will consider, in particular, the *polynomially coded* (PC) scheme [13] and the *polynomially coded multi-message* (PCMM) scheme [18]. PC and PCMM focus exclusively on linear computation tasks; and hence, we also consider a linear regression problem, in which the goal is to minimize

$$F(\boldsymbol{\theta}) = \frac{1}{N} \|\mathbf{X}\boldsymbol{\theta} - \mathbf{y}\|_2^2, \quad (37)$$

where $\boldsymbol{\theta} \in \mathbb{R}^d$ is the model parameter vector, $\mathbf{X} \in \mathbb{R}^{N \times d}$ is the data matrix, and $\mathbf{y} \in \mathbb{R}^N$ is the vector of labels. We split \mathbf{X} into n disjoint sub-matrices $\mathbf{X} = [X_1 \cdots X_n]^T$, where $X_i \in \mathbb{R}^{d \times N/n}$, and $\mathbf{y} = [y_1^T \cdots y_n^T]^T$, where $y_i \in \mathbb{R}^{N/n}$, $i \in [n]$. The gradient of loss function $F(\boldsymbol{\theta})$ is given by

$$\nabla F(\boldsymbol{\theta}) = \frac{2}{N} \mathbf{X}^T (\mathbf{X}\boldsymbol{\theta} - \mathbf{y}) = \frac{2}{N} \sum_{i=1}^n (X_i X_i^T \boldsymbol{\theta} - X_i y_i). \quad (38)$$

We perform gradient descent to minimize (37), in which the model parameters at the l -th iteration, $\boldsymbol{\theta}_l$, are updated as

$$\begin{aligned} \boldsymbol{\theta}_{l+1} &= \boldsymbol{\theta}_l - \eta_l \cdot \nabla F(\boldsymbol{\theta}_l) \\ &= \boldsymbol{\theta}_l - \eta_l \cdot \frac{2}{N} \sum_{i=1}^n (X_i X_i^T \boldsymbol{\theta}_l - X_i y_i), \end{aligned} \quad (39)$$

where η_l is the learning rate at iteration l . We consider a DGD algorithm, in which the computation of $\nabla F(\boldsymbol{\theta})$ is distributed across n workers, and the master updates the parameter vector according to (39) after receiving enough computations from the workers, and sends the updated parameter vector to the workers. In the following, we describe the computation tasks carried out by the workers and the master for different schemes, where for the l -th iteration, we set

$$h(X_i) = X_i X_i^T \boldsymbol{\theta}_l, \quad \text{for } i \in [n]. \quad (40)$$

Since $\mathbf{X}^T \mathbf{y} = \sum_{i=1}^n X_i y_i$ remains unchanged over iterations, we assume that its computation is carried out only once by the master node at the beginning of the learning task.

B. Distributed Computing Schemes

PC scheme [13]: At the l -th iteration of DGD, the task of computing $\mathbf{X}^T \mathbf{X} \boldsymbol{\theta}_l = \sum_{i=1}^n X_i X_i^T \boldsymbol{\theta}_l$ is distributed across the workers. For a computation load $r \geq 2$, worker i stores r distinct matrices $\tilde{X}_{i,1}, \dots, \tilde{X}_{i,r}$, where $\tilde{X}_{i,j} \in \mathbb{R}^{d \times N/n}$ is a linear combination of X_1, \dots, X_n , i.e., $\tilde{X}_{i,j} = \sum_{m=1}^n a_{i,j,m} X_m$, $a_{i,j,m} \in \mathbb{R}$, $j \in [r], i \in [n]$. Worker i , $i \in [n]$, then computes $\tilde{X}_{i,j} \tilde{X}_{i,j}^T \boldsymbol{\theta}_l$, $\forall j \in [r]$, and sends their sum, i.e., $\sum_{j=1}^r \tilde{X}_{i,j} \tilde{X}_{i,j}^T \boldsymbol{\theta}_l$, to the master. Thus, having set $h(\tilde{X}_{i,j}) = \tilde{X}_{i,j} \tilde{X}_{i,j}^T \boldsymbol{\theta}_l$, the computation task assigned to worker i is $\sum_{j=1}^r h(\tilde{X}_{i,j})$, and we denote the computation delay at worker

i by $T_{\text{PC},i}^{(1)}$, $i \in [n]$. We also denote the communication delay from worker i to the master by $T_{\text{PC},i}^{(2)}$, $i \in [n]$. The master receives the computation carried out by worker i at time

$$t_{\text{PC},i} = T_{\text{PC},i}^{(1)} + T_{\text{PC},i}^{(2)}, \quad \text{for } i \in [n]. \quad (41)$$

The PC scheme allows the master to recover $\mathbf{X}^T \mathbf{X} \boldsymbol{\theta}_l$ after receiving and processing the results from any $2 \lceil n/r \rceil - 1$ workers [13]. Thus, the completion time of PC is the $(2 \lceil n/r \rceil - 1)$ -th order statistics of $\{t_{\text{PC},1}, \dots, t_{\text{PC},n}\}$ denoted by $t_{\text{PC},(2 \lceil n/r \rceil - 1)}$. The average completion time of the PC is

$$\bar{t}_{\text{PC}}(r, n) = \mathbb{E} [t_{\text{PC},(2 \lceil n/r \rceil - 1)}], \quad (42)$$

where the expectation is taken over the computation and communication delay distributions. We note that, for PC, the master needs to further process the received computations to retrieve $\mathbf{X}^T \mathbf{X} \boldsymbol{\theta}_l$. This additional decoding delay is not taken into account here, but it can be significant.

Example 4. Consider $n = 4$ and $r = 2$. The following matrices are stored at worker i , for $i \in [4]$,

$$\tilde{X}_{i,1} = -(i-2)X_1 + (i-1)X_3, \quad (43a)$$

$$\tilde{X}_{i,2} = -(i-2)X_2 + (i-1)X_4. \quad (43b)$$

Worker i , $i \in [4]$, computes $(\tilde{X}_{i,1} \tilde{X}_{i,1}^T + \tilde{X}_{i,2} \tilde{X}_{i,2}^T) \boldsymbol{\theta}_l$, which is equivalent to evaluating a degree-2 polynomial

$$\begin{aligned} \phi_1(x) &= (X_1 X_1^T + X_2 X_2^T) \boldsymbol{\theta}_l (x-2)^2 \\ &\quad + (X_3 X_3^T + X_4 X_4^T) \boldsymbol{\theta}_l (x-1)^2 \\ &\quad - 2(X_1 X_3^T + X_2 X_4^T) \boldsymbol{\theta}_l (x-1)(x-2) \end{aligned} \quad (44)$$

at point $x = i$. The master can interpolate polynomial $\phi_1(x)$ after receiving computations from 3 workers. It then evaluates

$$\phi_1(1) + \phi_1(2) = \sum_{i=1}^4 X_i X_i^T \boldsymbol{\theta}_l = \mathbf{X}^T \mathbf{X} \boldsymbol{\theta}_l. \quad (45)$$

□

PCMM scheme [18]: PC is extended in [18] to exploit the partial computations carried out by the workers. For a computation load $r \geq 2$, worker i stores r distinct matrices $\hat{X}_{i,1}, \dots, \hat{X}_{i,r}$, where $\hat{X}_{i,j} = \sum_{m=1}^n b_{i,j,m} X_m$, $b_{i,j,m} \in \mathbb{R}$, $j \in [r]$, $i \in [n]$. Unlike PC, with PCMM proposed in [18], worker i , $i \in [n]$, computes $h(\hat{X}_{i,1}), \dots, h(\hat{X}_{i,r})$ sequentially, and sends the result of each computation to the master right after its execution, where $h(\hat{X}_{i,j}) = \hat{X}_{i,j} \hat{X}_{i,j}^T \boldsymbol{\theta}_l$. We denote the delay of computing task $h(\hat{X}_{i,j})$ and transmitting the computation to the master by $T_{\text{PCMM},i,j}^{(1)}$ and $T_{\text{PCMM},i,j}^{(2)}$, respectively, for $j \in [r]$ and $i \in [n]$. As a result, the master receives computation $h(\hat{X}_{i,j})$, $j \in [r]$, $i \in [n]$, at time

$$t_{\text{PCMM},i,j} = \sum_{m=1}^j T_{\text{PCMM},i,m}^{(1)} + T_{\text{PCMM},i,j}^{(2)}. \quad (46)$$

It is shown in [18] that the master can recover $\mathbf{X}^T \mathbf{X} \boldsymbol{\theta}_l$ after receiving and processing $2n - 1$ computations. Thus, the completion time of PCMM is the $(2n - 1)$ -th order statistics of $\{t_{\text{PCMM},i,j}, \forall j \in [r], \forall i \in [n]\}$ denoted by $t_{\text{PCMM},(2n-1)}$. The average completion time of the PC is given by

$$\bar{t}_{\text{PCMM}}(r, n) = \mathbb{E} [t_{\text{PCMM},(2n-1)}]. \quad (47)$$

Similarly to PC, PCMM also introduces an additional decoding delay, which will be ignored in our numerical comparisons.

Example 5. Consider $n = 4$ and $r = 2$. The following matrices are stored at worker i , for $i \in [4]$ and $j \in [2]$,

$$\hat{X}_{i,j} = \sum_{i=1}^4 X_i \prod_{m=1, m \neq i}^4 \frac{\beta_{i,j} - m}{i - m}, \quad (48)$$

where $\beta_{i,j}$, $\forall i \in [4]$, $\forall j \in [2]$, are different real values. Worker i , $i \in [4]$, computes $\hat{X}_{i,1} \hat{X}_{i,1}^T \boldsymbol{\theta}_l$ and $\hat{X}_{i,2} \hat{X}_{i,2}^T \boldsymbol{\theta}_l$ sequentially, and sends the result of each computation right after its completion. Computing $\hat{X}_{i,j} \hat{X}_{i,j}^T \boldsymbol{\theta}_l$, $i \in [4]$, $j \in [2]$, is equivalent to evaluating a degree-6 polynomial

$$\begin{aligned} \phi_2(x) &= \left(\sum_{i=1}^4 X_i \prod_{m=1, m \neq i}^4 \frac{x - m}{i - m} \right) \\ &\quad \left(\sum_{i=1}^4 X_i^T \prod_{m=1, m \neq i}^4 \frac{x - m}{i - m} \right) \boldsymbol{\theta}_l \end{aligned} \quad (49)$$

at point $x = \beta_{i,j}$. The master can interpolate polynomial $\phi_2(x)$ after receiving 7 computations from the workers. Having obtained $\phi_2(x)$, the master evaluates

$$\sum_{i=1}^4 \phi_2(i) = \sum_{i=1}^4 X_i X_i^T \boldsymbol{\theta}_l = \mathbf{X}^T \mathbf{X} \boldsymbol{\theta}_l. \quad (50)$$

□

Uncoded computing: Next, we focus on uncoded computation schemes. For a computation target k , the master updates the parameter vector after receiving k distinct computations, denoted by $h(X_{p_1}), \dots, h(X_{p_k})$, according to

$$\boldsymbol{\theta}_{l+1} = \boldsymbol{\theta}_l - \eta_l \cdot \frac{2n}{kN} \sum_{i=1}^k (h(X_{p_i}) - X_{p_i} y_{p_i}), \quad (51)$$

where we allow updating the parameter vector with partial computations, and if $k = n$, the update is equivalent to

$$\boldsymbol{\theta}_{l+1} = \boldsymbol{\theta}_l - \eta_l \cdot \frac{2}{N} \sum_{i=1}^n (h(X_i) - X_i y_i). \quad (52)$$

If $k < n$, the master stores $X_i y_i$, $\forall i \in [n]$, and at the l -th iteration computes $\sum_{i=1}^k X_{p_i} y_{p_i}$ to update the parameter vector as in (51). Whereas, if $k = n$, the master computes $\sum_{i=1}^n X_i y_i$ once, and updates the parameter vector as in (52).

We first consider the *random assignment* (RA) scheme [15]. For fairness of comparison, we assume that the training samples are divided into n batches. The computation load r for RA is $r = n$, i.e., the entire dataset is available at each worker. Each worker picks a distinct task (without replacement) independently at random, and sends its computation to the master immediately after its completion. In other words, each row of the TO matrix of RA, denoted by C_{RA} , is a random permutation of vector $[1 \cdots n]$. For a computation target k , the master updates the parameter vector according to (51) after receiving k distinct computations, and the corresponding average completion time is denoted by $\bar{t}_{\text{RA}}(n, k)$.

Example 6. Consider $n = r = 4$. Assume the RA scheme results in the following TO matrix:

$$C_{\text{RA}} = \begin{bmatrix} 2 & 1 & 4 & 3 \\ 2 & 4 & 1 & 3 \\ 1 & 4 & 3 & 2 \\ 4 & 3 & 1 & 2 \end{bmatrix}. \quad (53)$$

TABLE I: Characteristics of different schemes under consideration while performing iteration l of DGD.

Scheme	Computation load	Worker i	Computation target	Completion criteria	Master
<ul style="list-style-type: none"> • Cyclic Scheduling (CS) • Staircase Scheduling (SS) $X \in \{CS, SS\}$	$1 \leq r \leq n$	1. computes $h(X_{C_X(i,1)})$ and sends it 2. computes $h(X_{C_X(i,2)})$ and sends it ...	$1 \leq k \leq n$	receiving k distinct computations	$\theta_l - \frac{2\eta_l}{k} \sum_{i=1}^k (h(X_{p_i}) - X_{p_i} y_{p_i})$
Random Assignment (RA)	$r = n$	1. computes $h(X_{C_{RA}(i,1)})$ and sends it 2. computes $h(X_{C_{RA}(i,2)})$ and sends it ...	$1 \leq k \leq n$	receiving k distinct computations	$\theta_l - \frac{2\eta_l}{k} \sum_{i=1}^k (h(X_{p_i}) - X_{p_i} y_{p_i})$
Polynomially Coded (PC)	$r \geq 2$	1. computes $\sum_{j=1}^r h(\tilde{X}_{i,j})$ and sends it	$k = n$	receiving $2 \lceil n/r \rceil - 1$ computations	1. retrieves $\mathbf{X}^T \mathbf{X} \theta_l$ 2. $\theta_l - \frac{2\eta_l}{N} (\mathbf{X}^T \mathbf{X} \theta_l - \mathbf{X}^T \mathbf{y})$
Polynomially Coded Multi-Message (PCMM)	$r \geq 2$	1. computes $h(\hat{X}_{i,1})$ and sends it 2. computes $h(\hat{X}_{i,2})$ and sends it ...	$k = n$	receiving $2n - 1$ computations	1. retrieves $\mathbf{X}^T \mathbf{X} \theta_l$ 2. $\theta_l - \frac{2\eta_l}{N} (\mathbf{X}^T \mathbf{X} \theta_l - \mathbf{X}^T \mathbf{y})$

We have

$$\begin{aligned}
 t_{1,2} &= T_{1,2}^{(1)} + T_{1,2}^{(2)}, \quad t_{1,1} = T_{1,1}^{(1)} + T_{1,1}^{(1)} + T_{1,1}^{(2)}, \\
 t_{1,4} &= T_{1,2}^{(1)} + T_{1,1}^{(1)} + T_{1,4}^{(1)} + T_{1,4}^{(2)}, \\
 t_{1,3} &= T_{1,2}^{(1)} + T_{1,1}^{(1)} + T_{1,4}^{(1)} + T_{1,3}^{(1)} + T_{1,3}^{(2)}, \quad (54a)
 \end{aligned}$$

$$\begin{aligned}
 t_{2,2} &= T_{2,2}^{(1)} + T_{2,2}^{(2)}, \quad t_{2,4} = T_{2,2}^{(1)} + T_{2,4}^{(1)} + T_{2,4}^{(2)}, \\
 t_{2,1} &= T_{2,2}^{(1)} + T_{2,4}^{(1)} + T_{2,1}^{(1)} + T_{2,1}^{(2)}, \\
 t_{2,3} &= T_{2,2}^{(1)} + T_{2,4}^{(1)} + T_{2,1}^{(1)} + T_{2,3}^{(1)} + T_{2,3}^{(2)}, \quad (54b)
 \end{aligned}$$

$$\begin{aligned}
 t_{3,1} &= T_{3,1}^{(1)} + T_{3,1}^{(2)}, \quad t_{3,4} = T_{3,1}^{(1)} + T_{3,4}^{(1)} + T_{3,4}^{(2)}, \\
 t_{3,3} &= T_{3,1}^{(1)} + T_{3,4}^{(1)} + T_{3,3}^{(1)} + T_{3,3}^{(2)}, \\
 t_{3,2} &= T_{3,1}^{(1)} + T_{3,4}^{(1)} + T_{3,3}^{(1)} + T_{3,2}^{(1)} + T_{3,2}^{(2)}, \quad (54c)
 \end{aligned}$$

$$\begin{aligned}
 t_{4,4} &= T_{4,4}^{(1)} + T_{4,4}^{(2)}, \quad t_{4,3} = T_{4,4}^{(1)} + T_{4,3}^{(1)} + T_{4,3}^{(2)}, \\
 t_{4,1} &= T_{4,4}^{(1)} + T_{4,3}^{(1)} + T_{4,1}^{(1)} + T_{4,1}^{(2)}, \\
 t_{4,2} &= T_{4,4}^{(1)} + T_{4,3}^{(1)} + T_{4,1}^{(1)} + T_{4,2}^{(1)} + T_{4,2}^{(2)}, \quad (54d)
 \end{aligned}$$

□

The TO matrices for CS and SS are given in (13) and (20), respectively. At the l -th iteration of the DGD, the parameter vector is updated by the master according to (51).

Remark 7. An alternative approach to tackle the linear regression problem is that, besides $\mathbf{X}^T \mathbf{y}$, the master computes $\mathbf{W} \triangleq \mathbf{X}^T \mathbf{X}$ once at the beginning of the learning task. Accordingly, the problem reduces to computing matrix-vector multiplication $\mathbf{W} \theta_l$ at the l -th iteration in a distributed

manner, and after recovering $\mathbf{W} \theta_l$ the master updates the model parameter vector by

$$\theta_{l+1} = \theta_l - \eta_l \cdot \frac{2}{N} (\mathbf{W} \theta_l - \mathbf{X}^T \mathbf{y}). \quad (55)$$

In this case, the proposed CS and SS schemes can be updated accordingly to compute $\mathbf{W} \theta_l$ in a distributed manner. To the best of our knowledge, the coded computing scheme tolerating the highest number of straggling workers for the problem of computing $\mathbf{W} \theta_l$ distributively is the one proposed in [11]. Due to limited space we do not present the results for this setting here; however, the proposed CS and SS schemes outperform the coded computing scheme in [11] and approach the lower bound similarly to the results presented next.

We have summarized the characteristics of each of the schemes considered above in Table I, including the computation tasks conducted by the workers and the master. We note that the computational complexity of the CS, SS, PCMM, and RA schemes at the worker is the same, since with each of the schemes workers need to perform matrix-matrix-vector multiplications sequentially with dimensions $d \times N/n$, $N/n \times d$, and d , respectively. In the PC scheme, in addition to the matrix-matrix-vector multiplications, each worker needs to sum its results, r vectors of dimension d . With all the schemes, the master first needs to retrieve $\mathbf{X}^T \mathbf{X} \theta_l$, and then computes $\theta_l - \frac{2\eta_l}{N} (\mathbf{X}^T \mathbf{X} \theta_l - \mathbf{X}^T \mathbf{y})$. With the CS, SS, and RA schemes, the master retrieves $\mathbf{X}^T \mathbf{X} \theta_l$ by adding n d -dimensional vectors $h(X_1), \dots, h(X_n)$ (for fairness, we assume $k = n$),

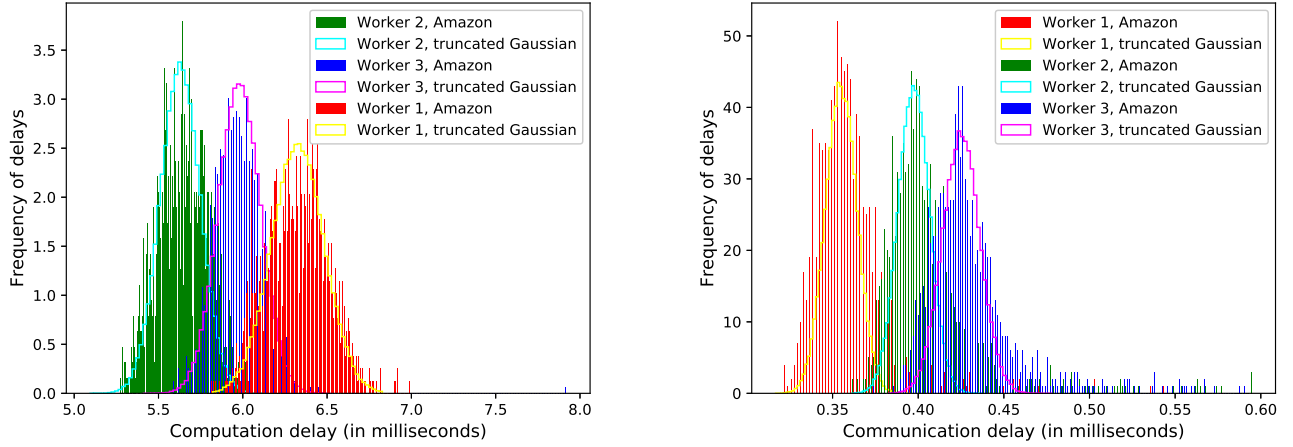


Fig. 3: Histogram of computation and communication delays of three different workers.

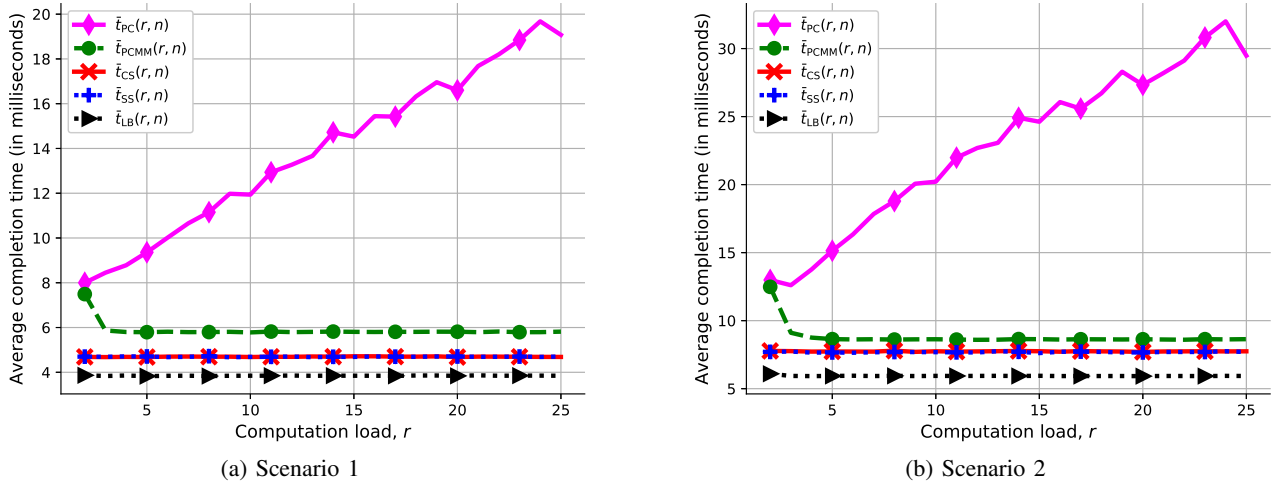


Fig. 4: Average completion time for the truncated Gaussian delay model in (56), when $k = n = 100$, and $2 \leq r \leq 25$.

which can be done in an online fashion as the computations from the workers arrive. With the PC scheme, the master should wait until it receives $2 \lceil n/r \rceil - 1$ computations, and, in order to retrieve $\mathbf{X}^T \mathbf{X} \theta_i$, it needs to interpolate d polynomials of degree $2 \lceil n/r \rceil - 2$, then evaluate it at $\lceil n/r \rceil$ points, and finally sum up the results from these $\lceil n/r \rceil$ points, each a vector of dimension N/n . Also, with PCMM, the master should first receive $2n - 1$ computations, then interpolate d polynomials of degree $2n - 2$, next evaluate it at n points, and finally sum up the results from these n points, each a vector of dimension N/n , to retrieve $\mathbf{X}^T \mathbf{X} \theta_i$. Accordingly, it can be concluded that the computational complexity at the master for the coded computing schemes is much higher than the uncoded ones. It is worth noting that, for this study, we do not take into account the computation delay at the master while evaluating the average completion time.

C. Numerical Experiments

For the numerical experiments we generate each entry of data matrix \mathbf{X} independently according to distribution $\mathcal{N}(0, 1)$. We also generate the labels as $y_i = (X_i + Z)^T U$, where $Z \in \mathbb{R}^{d \times N/n}$, with each entry distributed independently according to $\mathcal{N}(0, 0.01)$, and $U \in \mathbb{R}^d$ with each entry

distributed independently according to $\mathcal{U}(0, 1)$. For fairness we use the same dataset for all the schemes.

We train a linear regression model using the DGD algorithm described above with a constant learning rate $\eta_t = 0.01$. We run experiments on an Amazon EC2 cluster over t2.micro instance with $n + 1$ servers, where one of the servers is designated as the master and the rest serve as workers. We implement different schemes in Python and employ MPI4py library for message passing between different nodes.

At each iteration of the DGD algorithm, we measure the computation and communication delays of each task at each worker. We can then obtain the completion time of each scheme according to its completion criteria. We obtain the average completion time over 500 iterations.

In Fig. 3, we investigate the histograms of computation and communication delays experienced by three different workers. We carried out the experiment on an Amazon EC2 cluster with $N = 30000$, $d = 500$, $n = 3$, and set $r = 1$ and $k = n$ so that the master waits until it receives the computations from all the workers. Observe that both the computation and communication delays are not highly skewed across different workers. We also plotted the quantized PDF of a truncated

Gaussian distribution modelling the delays at each worker. As it can be seen, truncated Gaussian distribution provides a reasonably good estimate of the statistics of both computation and communication delays at different workers.

We first evaluate and compare the performances of different schemes assuming that both computation and communication delays follow truncated Gaussian distributions. For simplicity, we assume that the computation and communication delays of different tasks at the same worker are independent, i.e., $f_{i,[n]}^{(l)} = \prod_{j=1}^n f_{i,j}^{(l)}$, where $f_{i,j}^{(l)}$ denotes the PDF of $T_{i,j}^{(l)}$, $i, j \in [n]$, $l \in [2]$. For $t \in [\mu_i^{(l)} - a_i^{(l)}, \mu_i^{(l)} + b_i^{(l)}]$, $f_{i,j}^{(l)}$ is given by

$$f_{i,j}^{(l)}(t) = \frac{\phi\left((t - \mu_i^{(l)})/\sigma_i^{(l)}\right)}{\sigma_i^{(l)}\left(\Phi\left(b_i^{(l)}/\sigma_i^{(l)}\right) - \Phi\left(-a_i^{(l)}/\sigma_i^{(l)}\right)\right)}, \quad (56a)$$

and $f_{i,j}^{(l)}(t) = 0$, otherwise, for $i, j \in [n]$, $l \in [2]$, where

$$\phi(t) \triangleq \frac{1}{\sqrt{2\pi}} e^{-t^2/2}, \quad (56b)$$

$$\Phi(t) \triangleq \frac{1}{2} \left(1 + \operatorname{erf}(t/\sqrt{2})\right). \quad (56c)$$

For PC, we assume that $T_{PC,i}^{(1)}$ follows the same PDF as $\sum_{j=1}^r T_{i,j}^{(1)}$, for $i \in [n]$, where there is no loss of generality since the delays are independent and identically distributed (i.i.d.). We further assume that delays $T_{PCMM,i,j'}^{(1)}$, $\forall j' \in [r]$, follow PDF $f_{i,j}^{(1)}$, $j \in [n]$, for $i \in [n]$. Also, $T_{PC,i}^{(2)}$ and $T_{PCMM,i,j'}^{(2)}$, $\forall j' \in [r]$, have PDF $f_{i,j}^{(2)}$, $j \in [n]$, for $i \in [n]$.

For simplicity, we assume symmetric distributions for the delays, where $a_i^{(l)} = b_i^{(l)}$, $i \in [n]$, $l \in [2]$. We consider two scenarios in our simulations, and for both scenarios we set $a_i^{(1)} = 1.5E3$, $\sigma_i^{(1)} = 3E3$, $a_i^{(2)} = 5E4$, and $\sigma_i^{(2)} = 9E4$, $\forall i \in [n]$, where, for $\alpha, \beta \in \mathbb{R}$, we used the notation $\alpha E\beta$ to denote $\alpha \times 10^{-\beta}$. In Scenario 1, we set $\mu_i^{(1)} = 2E3$ and $\mu_i^{(2)} = 1E3$, $\forall i \in [n]$. In Scenario 2, $\{\mu_1^{(1)}, \dots, \mu_n^{(1)}\}$ is set as a random permutation of set $\{1.5E3, 1.53E3, \dots, (1.47 + 0.03n)E3\}$, and $\{\mu_1^{(2)}, \dots, \mu_n^{(2)}\}$ is a random permutation of set $\{1E3, 1.02E3, \dots, (0.98 + 0.02n)E3\}$. We note that, compared to Scenario 1, the computation and communication delays across workers are more diverse in Scenario 2. In Fig. 4 we compare the performances of different schemes for the truncated Gaussian model with $n = 100$ workers and $k = n$. For both scenarios, the performances of CS and SS almost coincide, and both CS and SS schemes outperform PCMM and PC for the whole range of r under consideration, $2 \leq r \leq 25$. Compared to Scenario 1, the gap between uncoded schemes CS and SS and coded schemes PCMM and PC is less significant in Scenario 2, in which the delays are more diverse. The improvement of both SS and CS schemes over PCMM is higher for $r \leq 4$ and $r \leq 5$ in Scenarios 1 and 2, respectively. Also, PCMM outperforms PC in both scenarios. For $r = n$, $\bar{t}_{RA}(n, n) = 11.66$ and $\bar{t}_{RA}(n, n) = 16.53$ milliseconds in Scenarios 1 and 2, respectively. SS reduces these average delays by %59.7 and %53.35 for Scenarios 1 and 2, respectively, showing that an efficient computation schedule for uncoded computing can reduce the latency significantly.

Next, we present the results of experiments carried out on

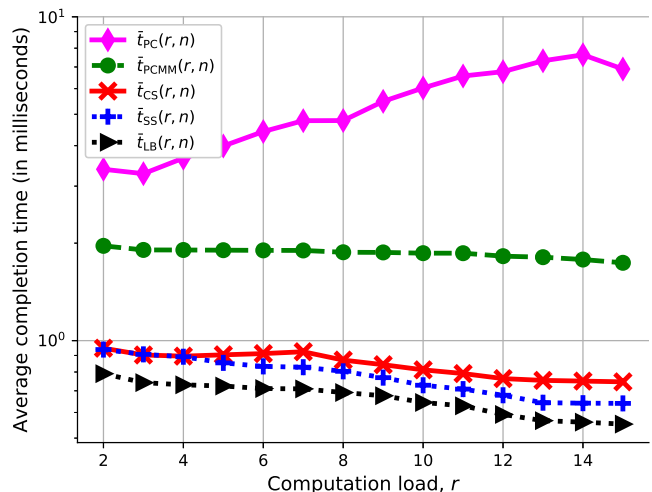


Fig. 5: Average completion time of different schemes with respect to computation load, $r \geq 2$, when $k = n = 15$.

Amazon EC2 cluster. We compare the average completion time of different schemes with respect to the computation load r , $r \geq 2$, in Fig. 5, where $n = 15$, $d = 400$, and $N = 900$. As it can be seen, CS and SS outperform PC and PCMM significantly; while PCMM improves upon PC. This result shows that standard coded computation framework cannot fully exploit the computing capabilities in the network, and splitting the computational tasks assigned to each worker and receiving partial computations performed by each worker can reduce the average completion time significantly. We also observe that the average completion time of PC increases with r . This is because the delays at different workers are not significantly different; and thus, increasing the computation load to reduce the number of received computations from different workers can increase the total delay. This is another limitation of the coded computation framework, as it requires careful tuning of the parameters based on the statistics of the delays in the system. We observe that the gap between the average completion time of SS and the lower bound is relatively small for the entire range, and reduces with r , and SS outperforms CS with the improvement slightly increasing with r . The average completion time of RA, which requires $r = n$, is $\bar{t}_{RA}(n, n) = 0.895$ millisecond, while SS achieves $\bar{t}_{SS}(n, n) = 0.64$ millisecond, i.e., around %28.5 reduction. Thus, designing the TO matrix, rather than random computations, can provide significant improvement in computation speed. We observe that the average completion time of each scheme considered in Figures 4 and 5 follows a similar pattern. This verifies that the truncated Gaussian model can reasonably capture the statistical behaviour of the delays.

In Fig. 6, we compare the performances of different schemes with respect to the number of workers, n . We consider $d = 500$, $N = 1000$, and $r = n$. When N/n is not an integer, we zero-pad the dataset. We observe that, except PCMM, the average completion time of different schemes reduce slightly with n when N is fixed. For PC, when $r = n$, the computation received from the fastest worker determines the completion time, and, with all other parameters fixed, the computation

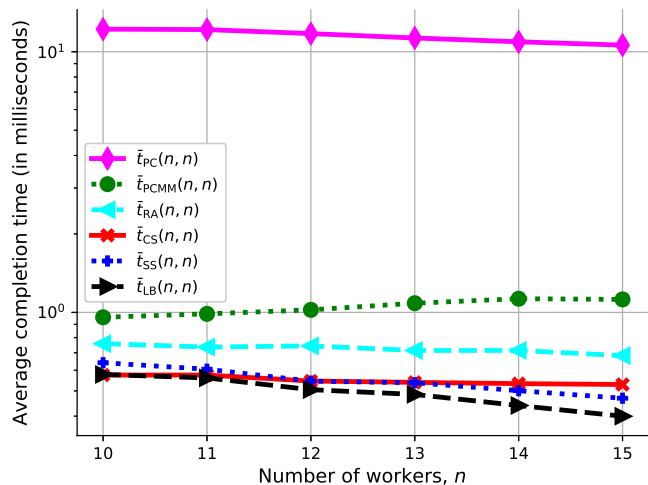


Fig. 6: Average completion time of different schemes versus number of workers, $10 \leq n \leq 15$, when $r = k = n$.

delay at each worker depends mostly on N . Thus, by introducing new workers when N is fixed, the average completion time is expected to decrease. Whereas, with PCMM, although the computation time of each task is expected to decrease with n , the average completion time increases. This is due to the increase in the number of communications required by a factor of two as we have $\bar{t}_{\text{PCMM}}(r, n) = \mathbb{E}[t_{\text{PCMM}, (2n-1)}]$. For uncoded computing schemes, RA, CS and SS, the average completion time decreases with n , as they allow a better utilization of the computing resources. As before, we observe that CS and SS improve the average completion time significantly compared to PC and PCMM. Also, based on the superiority of the CS and SS over RA, we conclude that the TO matrix design is essential in reducing the average delay of uncoded computing schemes. CS outperforms SS for small n values, but SS takes over as n increases. The relatively small gap between the average completion times of CS and SS and the lower bound illustrates their efficiency in scheduling the tasks despite the lack of any information on the speeds of the workers.

In Fig. 7, we compare the performance of different uncoded computation schemes and the lower bound with respect to the computation target, k . We set $n = 10$, $r = n$, $N = 1000$, and $d = 800$, and consider $k \in [2 : n]$. As expected, the average completion time increases with k . The gap between different schemes also increases with k , as the efficiency of scheduling tasks is more distinguishable for higher values of k . The average completion time of SS coincides with the lower bound for small and medium k values $k \in [2 : 6]$, and the gap between the two is negligible even for higher k values. Here we do not consider the coded computing schemes, PC and PCMM, as they are designed only for $k = n$.

VII. CONCLUSIONS

We have studied distributed computation across inhomogeneous workers. The computation here may correspond to each iteration of a DGD algorithm applied on a large dataset, and it is considered to be completed when the master receives k distinct computations. We assume that each worker has access

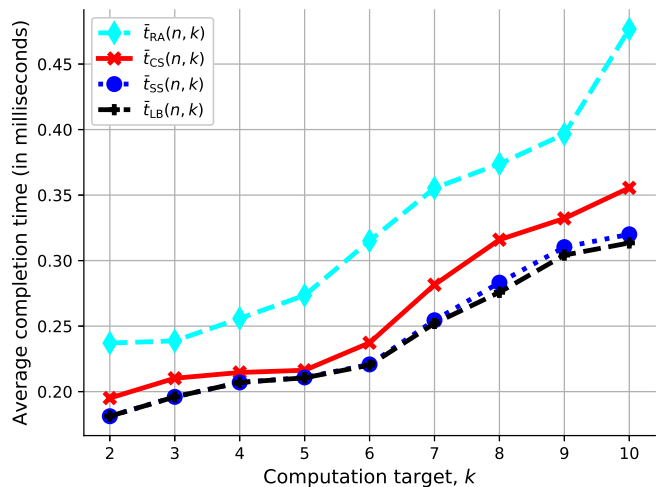


Fig. 7: Average completion time of different schemes versus computation target, $2 \leq k \leq n$, when $r = n = 10$.

to a limited portion of the dataset, defined as the computation load. In contrast to the growing literature on coded computation to mitigate straggling servers, here we have studied uncoded computations and sequential communication to the master in order to benefit from all the computations carried out by the workers, including the slower ones. Since the instantaneous computation speeds of the workers are not known in advance, allocation of the tasks to the workers and their scheduling become crucial in minimizing the average completion time. In particular, we have considered the assignment of data points to the workers with a predesigned computation order. Workers send the result of each computation to the master as soon as it is executed, and move on to compute the next task assigned to them. Assuming a general statistics for the computation and communication delays of different workers, we have obtained closed-form expressions for the average completion time of two particular computation allocation schemes, called CS and SS. The CS scheme dictates the same computation order at different workers, which is implemented by a cyclic shift operator. With SS, we introduce inverse computation orders at the workers. We have compared the performance of these proposed schemes with the existing ones in the literature, particularly the coded PC [13], PCMM [18], and uncoded RA [15] schemes. The results of the experiments carried out on Amazon EC2 cluster show that the CS and SS schemes provide significant reduction in the average completion time over these schemes. The poor performance of the PC scheme can be explained by the fact that when the delays associated with different workers are not highly skewed, utilizing the partial computations by the slower workers becomes beneficial. The superiority of CS and SS compared to the RA scheme, which randomly schedules the tasks, illustrates the importance of task scheduling in speeding up the computations.

We also remark that, unlike the proposed schemes, the PC and PCMM schemes introduce additional encoding and decoding complexities at the master, which have not been considered in the evaluations here. Moreover, in the case of DSGD, having computed the partial gradient on separate data points

may allow the workers to exploit more advanced methods to reduce their communication load, such as compression [22] or quantization [23], [24], and can be beneficial in the case of communications over noisy channels [19], which may not be applicable in the case of coded computations.

APPENDIX A
PROOF OF THEOREM 1

The event $\{t_C(r, k) > t\}$ is equivalent to the union of the events, for which the time to complete any arbitrary set of at least $n - k + 1$ distinct computations is greater than t , i.e.,

$$\Pr\{t_C(r, k) > t\} = \Pr\left\{\bigcup_{\mathcal{G} \subset [n]: n-k+1 \leq |\mathcal{G}| \leq n} \left\{t_j > t, t_{j'} \leq t, \forall j \in \mathcal{G}, \forall j' \in \mathcal{G}'\right\}\right\}, \quad (57)$$

where we define $\mathcal{G}' \triangleq [n] \setminus \mathcal{G}$. Since the events $\{t_j > t, t_{j'} \leq t, \forall j \in \mathcal{G}, \forall j' \in \mathcal{G}'\}$, for all distinct sets $\mathcal{G} \subset [n]$, are mutually exclusive (pairwise disjoint), we have

$$\begin{aligned} \Pr\{t_C(r, k) > t\} &= \sum_{i=n-k+1}^n \sum_{\mathcal{G} \subset [n]: |\mathcal{G}|=i} \Pr\left\{t_j > t, t_{j'} \leq t, \forall j \in \mathcal{G}, \forall j' \in \mathcal{G}'\right\} \\ &= \sum_{i=n-k+1}^n \sum_{\mathcal{G} \subset [n]: |\mathcal{G}|=i} H_{\mathcal{G}, \mathcal{G}'}, \end{aligned} \quad (58)$$

where, for $\mathcal{S}_1 \subset [n]$ and $\mathcal{S}_2 \subset [n]$, we define

$$H_{\mathcal{S}_1, \mathcal{S}_2} \triangleq \Pr\{t_{j_1} > t, t_{j_2} \leq t, \forall j_1 \in \mathcal{S}_1, \forall j_2 \in \mathcal{S}_2\}. \quad (59)$$

Lemma 1. *Given a particular set $\mathcal{G} \subset [n]$, $|\mathcal{G}| = i$, for $i \in [n - k + 1 : n]$, we have*

$$\begin{aligned} H_{\mathcal{G}, \mathcal{G}'} &= \sum_{m=i}^n (-1)^{i+m} \sum_{\hat{\mathcal{G}} \subset \mathcal{G}': |\hat{\mathcal{G}}|=m-i} H_{\mathcal{G} \cup \hat{\mathcal{G}}, \emptyset} \\ &= \sum_{m=i}^n (-1)^{i+m} \sum_{\hat{\mathcal{G}} \subset \mathcal{G}': |\hat{\mathcal{G}}|=m-i} \Pr\left\{t_j > t, \forall j \in \mathcal{G} \cup \hat{\mathcal{G}}\right\}. \end{aligned} \quad (60)$$

The importance of the results given in Lemma 1 is that it establishes a connection between the event $\{t_j > t, t_{j'} \leq t, \forall j \in \mathcal{G}, \forall j' \in \mathcal{G}'\}$ and all the events $\{t_j > t, \forall j \in \mathcal{G} \cup \hat{\mathcal{G}}\}$, $\forall \hat{\mathcal{G}} \subset \mathcal{G}'$, where for the special case $n = 1$ and $\mathcal{G} = \{1\}$, it is trivial to see that $\Pr\{t_1 > t\} = 1 - \Pr\{t_1 \leq t\}$. The proof of Lemma 1 can be found in Appendix B, where we use the fact that, for any $g \in \mathcal{G}'$, we have

$$H_{\mathcal{G}, \mathcal{G}'} = H_{\mathcal{G}, \mathcal{G}' \setminus \{g\}} - H_{\mathcal{G} \cup \{g\}, \mathcal{G}' \setminus \{g\}}. \quad (61)$$

According to Lemma 1, for $i \in [n - k + 1 : n]$, we have

$$\begin{aligned} &\sum_{\mathcal{G} \subset [n]: |\mathcal{G}|=i} H_{\mathcal{G}, \mathcal{G}'} \\ &= \sum_{\mathcal{G} \subset [n]: |\mathcal{G}|=i} \sum_{m=i}^n (-1)^{i+m} \sum_{\hat{\mathcal{G}} \subset \mathcal{G}': |\hat{\mathcal{G}}|=m-i} H_{\mathcal{G} \cup \hat{\mathcal{G}}, \emptyset} \\ &= \sum_{m=i}^n (-1)^{i+m} \sum_{\mathcal{G} \subset [n]: |\mathcal{G}|=i} \sum_{\hat{\mathcal{G}} \subset \mathcal{G}': |\hat{\mathcal{G}}|=m-i} H_{\mathcal{G} \cup \hat{\mathcal{G}}, \emptyset} \\ &\stackrel{(a)}{=} \sum_{m=i}^n (-1)^{i+m} \binom{m}{i} \sum_{\mathcal{S} \subset [n]: |\mathcal{S}|=m} H_{\mathcal{S}, \emptyset}, \end{aligned} \quad (62)$$

where (a) follows since, for each set $\mathcal{S} = \mathcal{G} \cup \hat{\mathcal{G}}$ with $|\mathcal{S}| = m$,

there are $\binom{m}{i}$ sets $\mathcal{G} \cup \hat{\mathcal{G}}$. Plugging (62) into (58) yields

$$\begin{aligned} &\Pr\{t_C(r, k) > t\} \\ &= \sum_{i=n-k+1}^n \sum_{m=i}^n (-1)^{i+m} \binom{m}{i} \sum_{\mathcal{S} \subset [n]: |\mathcal{S}|=m} H_{\mathcal{S}, \emptyset}. \end{aligned} \quad (63)$$

For a particular set $\mathcal{S} \subset [n]$ with $|\mathcal{S}| = s$, for some $s \in [n - k + 1 : n]$, the coefficient of $H_{\mathcal{S}, \emptyset}$ in (63) is given by

$$\begin{aligned} &\sum_{i=n-k+1}^s (-1)^{i+s} \binom{s}{i} \\ &= \sum_{i=0}^s (-1)^{i+s} \binom{s}{i} - \sum_{i=0}^{n-k} (-1)^{i+s} \binom{s}{i} \\ &= 0 - (-1)^{n-k+s} \binom{s-1}{n-k} = (-1)^{n-k+s+1} \binom{s-1}{n-k}, \end{aligned} \quad (64)$$

which results in

$$\begin{aligned} &\Pr\{t_C(r, k) > t\} \\ &= \sum_{i=n-k+1}^n (-1)^{n-k+i+1} \binom{i-1}{n-k} \sum_{\mathcal{S} \subset [n]: |\mathcal{S}|=i} H_{\mathcal{S}, \emptyset}. \end{aligned} \quad (65)$$

According to the definition of $H_{\mathcal{S}, \emptyset}$, (65) concludes the proof of (7). Furthermore, since $t_C(r, k) \geq 0$, we have

$$\bar{t}_C(r, k) = \int_0^\infty (1 - F_{t_C}(t)) dt, \quad (66)$$

which yields the expression in (8).

APPENDIX B
PROOF OF LEMMA 1

We prove the equation in (60) by induction. For the ease of presentation, we define set \mathcal{G}_i as $\mathcal{G}_i \subset [n]$ such that $|\mathcal{G}_i| = i$, and we denote $\mathcal{G}'_i = [n] \setminus \mathcal{G}_i$, $i \in [n]$. We first show that the equality in (60) holds for any \mathcal{G}'_{n-1} with $|\mathcal{G}'_{n-1}| = 1$ (note that the proof is trivial for \mathcal{G}'_n). According to (61), we have

$$H_{\mathcal{G}_{n-1}, \mathcal{G}'_{n-1}} = H_{\mathcal{G}_{n-1}, \emptyset} - H_{[n], \emptyset}, \quad (67)$$

which is identical to the equality in (60). Then, assuming that for any \mathcal{G}'_i with $|\mathcal{G}'_i| = n - i$, we have

$$H_{\mathcal{G}_i, \mathcal{G}'_i} = \sum_{m=i}^n (-1)^{i+m} \sum_{\hat{\mathcal{G}} \subset \mathcal{G}'_i: |\hat{\mathcal{G}}|=m-i} H_{\mathcal{G}_i \cup \hat{\mathcal{G}}, \emptyset}, \quad (68)$$

we prove that

$$\begin{aligned} H_{\mathcal{G}_{i-1}, \mathcal{G}'_{i-1}} &= \sum_{m=i}^n (-1)^{i+m-1} \\ &\quad \sum_{\hat{\mathcal{G}} \subset \mathcal{G}'_{i-1}: |\hat{\mathcal{G}}|=m-i+1} H_{\mathcal{G}_{i-1} \cup \hat{\mathcal{G}}, \emptyset}. \end{aligned} \quad (69)$$

Consider a fixed set \mathcal{G}_i . For $g_i \in \mathcal{G}_i$, let $\mathcal{G}_{i-1} = \mathcal{G}_i \setminus \{g_i\}$, which results in $\mathcal{G}'_{i-1} = \mathcal{G}'_i \cup \{g_i\}$. From (61), we have

$$\begin{aligned} H_{\mathcal{G}_{i-1}, \mathcal{G}'_{i-1}} &= H_{\mathcal{G}_{i-1}, \mathcal{G}'_{i-1} \setminus \{g_i\}} - H_{\mathcal{G}_{i-1} \cup \{g_i\}, \mathcal{G}'_{i-1} \setminus \{g_i\}} \\ &= H_{\mathcal{G}_i \setminus \{g_i\}, \mathcal{G}'_i} - H_{\mathcal{G}_i, \mathcal{G}'_i}. \end{aligned} \quad (70)$$

According to (68), it follows that

$$\begin{aligned} &H_{\mathcal{G}_i \setminus \{g_i\}, \mathcal{G}'_i} \\ &= \sum_{m=i-1}^{n-1} (-1)^{i+m-1} \sum_{\hat{\mathcal{G}} \subset \mathcal{G}'_i: |\hat{\mathcal{G}}|=m-i+1} H_{\mathcal{G}_{i-1} \cup \hat{\mathcal{G}}, \emptyset} \end{aligned}$$

$$= \sum_{m=i-1}^n (-1)^{i+m-1} \sum_{\hat{G} \subset G'_i: |\hat{G}|=m-i+1} H_{G_{i-1} \cup \hat{G}, \emptyset}, \quad (71)$$

and

$$\begin{aligned} H_{G_i, G'_i} &= \sum_{m=i}^n (-1)^{i+m} \sum_{\hat{G} \subset G'_i: |\hat{G}|=m-i} H_{G_i \cup \hat{G}, \emptyset} \\ &= \sum_{m=i-1}^n (-1)^{i+m} \sum_{\hat{G} \subset G'_i: |\hat{G}|=m-i} H_{G_i \cup \hat{G}, \emptyset} \\ &= \sum_{m=i-1}^n (-1)^{i+m} \sum_{\hat{G} \subset G'_i: |\hat{G}|=m-i} H_{G_{i-1} \cup \{g_i\} \cup \hat{G}, \emptyset}. \end{aligned} \quad (72)$$

By plugging (71) and (72) in (70), we have

$$\begin{aligned} H_{G_{i-1}, G'_{i-1}} &= \sum_{m=i-1}^n (-1)^{i+m-1} \left(\sum_{\hat{G} \subset G'_i: |\hat{G}|=m-i+1} H_{G_{i-1} \cup \hat{G}, \emptyset} \right. \\ &\quad \left. + \sum_{\hat{G} \subset G'_i: |\hat{G}|=m-i} H_{G_{i-1} \cup \{g_i\} \cup \hat{G}, \emptyset} \right) \\ &= \sum_{m=i-1}^n (-1)^{i+m-1} \sum_{\hat{G} \subset G'_i \cup \{g_i\}: |\hat{G}|=m-i+1} H_{G_{i-1} \cup \hat{G}, \emptyset} \\ &= \sum_{m=i-1}^n (-1)^{i+m-1} \sum_{\hat{G} \subset G'_{i-1}: |\hat{G}|=m-i+1} H_{G_{i-1} \cup \hat{G}, \emptyset}, \end{aligned} \quad (73)$$

which provides the proof of the equality in (69). This completes the proof of (60).

REFERENCES

- [1] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2012.
- [2] J. Choi, D. W. Walker, and J. J. Dongarra, "Pumma: Parallel universal matrix multiplication algorithms on distributed memory concurrent computers," *Concurrency: Practice and Experience*, vol. 6, no. 7, pp. 543–570, 1994.
- [3] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," *IEEE Trans. Inform. Theory*, vol. 64, no. 3, pp. 1514–1529, Mar. 2018.
- [4] R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *Proc. of ICML*, Sydney, Australia, Aug. 2017, pp. 3368–3376.
- [5] W. Halbawi, N. A. Ruhi, F. Salehi, and B. Hassibi, "Improving distributed gradient descent using Reed-Solomon codes," *arXiv:1706.05436v1 [cs.IT]*, Jun. 2017.
- [6] 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," *arXiv:1803.01113v3 [cs.IT]*, May 2018.
- [7] R. K. Maity, A. S. Rawat, and A. Mazumdar, "Robust gradient descent via moment encoding with LDPC codes," *arXiv:1805.08327v1 [cs.IT]*.
- [8] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," *arXiv:1802.03475v1 [cs.IT]*, Feb. 2018.
- [9] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "Straggler mitigation in distributed matrix multiplication: Fundamental limits and optimal coding," *arXiv:1801.07487v2 [cs.IT]*, May 2018.
- [10] S. Dutta, M. Fahim, F. Haddadpour, H. Jeong, V. Cadambe, and P. Grover, "On the optimal recovery threshold of coded matrix multiplication," *arXiv:1801.10292v2 [cs.IT]*, May 2018.
- [11] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "Polynomial codes: an optimal design for high-dimensional coded matrix multiplication," *arXiv:1705.10464 [cs.IT]*, Jan. 2018.
- [12] Q. Yu, N. Raviv, J. So, and A. S. Avestimehr, "Lagrange coded computing: Optimal design for resiliency, security and privacy," *arXiv:1806.00939v2 [cs.IT]*, Jun. 2018.
- [13] S. Li, S. M. Mousavi Kalan, Q. Yu, M. Soltanolkotabi, and A. S. Avestimehr, "Polynomially coded regression: Optimal straggler mitigation via data encoding," *arXiv:1805.09934v1 [cs.IT]*, May 2018.
- [14] N. Ferdinand and S. C. Draper, "Hierarchical coded computation," in *Proc. IEEE Int'l Symp. on Inform. Theory (ISIT)*, Vail, CO, USA, Jun. 2018, pp. 1620–1624.
- [15] S. Li, S. M. Mousavi Kalan, A. S. Avestimehr, and M. Soltanolkotabi, "Near-optimal straggler mitigation for distributed gradient methods," *arXiv:1710.09990v1 [cs.IT]*, Oct. 2017.
- [16] S. Kiani, N. Ferdinand, and S. C. Draper, "Exploitation of stragglers in coded computation," *arXiv:1806.10253v1 [cs.IT]*, Jun. 2018.
- [17] A. Mallick, M. Chaudhari, and G. Joshi, "Rateless codes for near-perfect load balancing in distributed matrix-vector multiplication," *arXiv:1804.10331v2 [cs.DC]*, Mar. 2018.
- [18] E. Ozfatura, D. Gündüz, and S. Ulukus, "Speeding up distributed gradient descent by utilizing non-persistent stragglers," *arXiv:1808.02240v2 [cs.IT]*, Aug. 2018.
- [19] M. Mohammadi Amiri and D. Gündüz, "Machine learning at the wireless edge: Distributed stochastic gradient descent over-the-air," *arXiv:1901.00844 [cs.DC]*, Feb. 2019.
- [20] —, "Federated learning over wireless fading channels," *arXiv:1907.09769 [cs.IT]*, Jul. 2019.
- [21] M. Mohammadi Amiri, T. M. Duman, and D. Gündüz, "Collaborative machine learning at the wireless edge with blind transmitters," in *Proc. IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, Ottawa, Canada, Nov. 2019.
- [22] J. Konecny, H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv:1610.05492v2 [cs.LG]*, Oct. 2017.
- [23] F. Seidel, H. Fu, J. Droppo, G. Li, and D. Yu, "1-bit stochastic gradient descent and its application to data-parallel distributed training of speech DNNs," in *INTERSPEECH*, Singapore, Sep. 2014, pp. 1058–1062.
- [24] N. Strom, "Scalable distributed DNN training using commodity GPU cloud computing," in *INTERSPEECH*, 2015, pp. 1488–1492.
- [25] M. A. Attia and R. Tandon, "Combating computational heterogeneity in large-scale distributed computing via work exchange," *arXiv:1711.08452 [cs.DC]*, Nov. 2017.
- [26] R. L. Graham, "Bounds on multiprocessing anomalies and packing algorithms," in *Proc. SJCC*, 1972, pp. 205–218.
- [27] S. Xue, M. Li, X. Xu, and J. Chen, "An ACO-LB algorithm for task scheduling in the cloud environment," *J. Software*, vol. 9, no. 2, pp. 466–473, Feb. 2014.
- [28] S. Zhan and H. Huo, "Improved PSO-based task scheduling algorithm in cloud computing," *J. Inf. Comput. Sci.*, vol. 9, no. 13, pp. 3821–3829, Nov. 2012.
- [29] G. Reddy, N. Reddy, and S. Phanikumar, "Multi objective task scheduling using modified ant colony optimization in cloud computing," *Int'l J. Intell. Eng. Sys.*, vol. 11, no. 3, pp. 242–250, 2018.
- [30] T. Choudhari, M. Moh, and T.-S. Moh, "Prioritized task scheduling in fog computing," in *Proc. ACMSE*, Richmond, Kentucky, Mar. 2018, pp. 401–405.
- [31] L. Yin, J. Luo, and H. Luo, "Tasks scheduling and resource allocation in fog computing based on containers for smart manufacturing," *IEEE Trans. ON Indus. Inform.*, vol. 14, no. 10, pp. 4712–4721, Oct. 2018.
- [32] R. Pedarsani, J. Walrand, and Y. Zhong, "Scheduling tasks with precedence constraints on multiple servers," in *Proc. Allerton*, Monticello, IL, 2014, pp. 1196–1203.
- [33] C.-S. Yang, R. Pedarsani, and A. S. Avestimehr, "Communication-aware scheduling of serial tasks for dispersed computing," *arXiv:1804.06468 [cs.DC]*, Apr. 2018.
- [34] J. D. Ullman, "NP-complete scheduling problems," *J. Comput. Syst. Sci.*, vol. 10, no. 3, pp. 384–393, Jun. 1975.
- [35] A. F. Aji and K. Heafield, "Sparse communication for distributed gradient descent," *arXiv:1704.05021v2 [cs.CL]*, Jul. 2017.
- [36] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Sparse binary compression: Towards distributed deep learning with minimal communication," *arXiv:1805.08768v1 [cs.LG]*, May 2018.
- [37] C.-Y. Chen, J. Choi, D. Brand, A. Agrawal, W. Zhang, and K. Gopalakrishnan, "AdaComp: Adaptive residual gradient compression for data-parallel distributed training," *arXiv:1712.02679v1 [cs.LG]*, Dec. 2017.
- [38] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally, "Deep gradient compression: Reducing the communication bandwidth for distributed training," *arXiv:1712.01887v2 [cs.CV]*, Feb. 2018.
- [39] Z. Tao and Q. Li, "eSGD: Communication efficient distributed deep learning on the edge," in *Workshop on Hot Topics in Edge Computing (HotEdge)*, Boston, MA, USA, Jul. 2018.
- [40] H. Wang, S. Sievert, S. Liu, Z. Charles, D. Papailiopoulos, and S. Wright, "ATOMO: Communication-efficient learning via atomic sparsification," *arXiv:1806.04090v2 [stat.ML]*, Jun. 2018.