King's Research Portal

DOI:

[10.1109/TNNLS.2020.2979762](https://doi.org/10.1109/TNNLS.2020.2979762)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Zhang, J., & Simeone, O. (2020). LAGC: Lazily Aggregated Gradient Coding for Straggler-Tolerant and Communication-Efficient Distributed Learning. *IEEE Transactions on Neural Networks and Learning Systems*, 1-13. <https://doi.org/10.1109/TNNLS.2020.2979762>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

LAGC: Lazily Aggregated Gradient Coding for Straggler-Tolerant and Communication-Efficient Distributed Learning

Jingjing Zhang^{id} and Osvaldo Simeone^{id}, *Fellow, IEEE*

Abstract—Gradient-based distributed learning in parameter server (PS) computing architectures is subject to random delays due to straggling worker nodes and to possible communication bottlenecks between PS and workers. Solutions have been recently proposed to separately address these impairments based on the ideas of gradient coding (GC), worker grouping, and adaptive worker selection. This article provides a unified analysis of these techniques in terms of wall-clock time, communication, and computation complexity measures. Furthermore, in order to combine the benefits of GC and grouping in terms of robustness to stragglers with the communication and computation load gains of adaptive selection, novel strategies, named lazily aggregated GC (LAGC) and grouped-LAG (G-LAG), are introduced. Analysis and results show that G-LAG provides the best wall-clock time and communication performance while maintaining a low computational cost, for two representative distributions of the computing times of the worker nodes.

Index Terms—Adaptive selection, coding, distributed learning, gradient descent (GD), grouping.

I. INTRODUCTION

IN ORDER to scale machine learning in order to cope with large volumes of input data, distributed implementations of gradient-based methods that leverage the parallelism of first-order optimization techniques are commonly adopted [1]–[3]. A standard large-scale distributed computing architecture relies on many parallel worker nodes to perform iterative computations of the gradients and on a central parameter server (PS) to aggregate the computed gradients and communicate with the workers [4], [5]. The PS computing architecture is subject to two key impairments. First, the potentially high tail of the distribution of the computing times at the workers can cause significant slowdowns in wall-clock runtime per iteration due to straggling workers [6]. Second, the communication overhead resulting from intensive two-way communications between the PS and the workers may require significant networking resources to be available in order not to dominate the overall runtime [7].

Manuscript received May 22, 2019; revised December 9, 2019; accepted March 2, 2020. This work was supported by the European Research Council through the European Union’s Horizon 2020 Research and Innovation Program under Grant 725731. (Corresponding author: Jingjing Zhang.)

The authors are with the Department of Informatics, King’s College London, London WC2R 2LS, U.K. (e-mail: jingjing.l.zhang@kcl.ac.uk; osvaldo.simeone@kcl.ac.uk).

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

Digital Object Identifier 10.1109/TNNLS.2020.2979762

TABLE I
QUALITATIVE COMPARISONS WITH RESPECT
TO STANDARD (DISTRIBUTED) GD

	Coding	Grouping	Adaptive selection
Robustness to stragglers	better	better	same
Communication load	same	same	better
Computation load	worse	worse	better

Recently, solutions have been developed that aim at improving robustness to stragglers—namely, gradient coding (GC) and grouping [8], [9]—or communication load—namely, adaptive selection [10] (see Table I for a summary). GC, introduced in [8], increases robustness to stragglers by leveraging storage and computation redundancy at the worker nodes as compared with standard (distributed) gradient descent (GD) [11]. With a redundancy factor $r > 1$, each worker stores, and computes on, r times more data than with GD. Under GC, given a redundancy factor $r > 1$, up to $r - 1$ stragglers can be tolerated, while still allowing the PS to exactly compute the gradient at any iteration. GC requires coding the computed gradients prior to communication from the workers to the PS and decoding at the PS.

As a special case of GC, given a redundancy factor r equal to the number M of workers, each worker can store the entire data set. Hence, the gradient can be obtained from any worker without requiring any coding or decoding operation. In the typical case in which r is smaller than M , the same simple procedure can be applied to groups of workers. In particular, given a redundancy factor r , the data set can be partitioned so that each partition is available to all nodes of a group of r workers. The PS can then recover the gradient upon receiving the computations of any server for each group. The outlined grouping scheme can, hence, tolerate up to $(M/r)(r - 1)$ stragglers, which may be significantly larger than $r - 1$ when $M \gg r$ [9].

While GC and grouping aim at reducing wall-clock time per iteration by leveraging storage and computation redundancy, the goal of adaptive selection is to reduce the communication and computation loads. This is done by selecting at each iteration a subset of workers to be active [10]. Selection is made by predicting at the PS (or in a distributed way at the servers) the servers that are more likely to have informative updates as compared with their latest communicated gradients. The approach termed lazily aggregated gradient (LAG) is shown

in [10] to have approximately the same iteration complexity as GD at substantially reduced communication and computation loads.

In this article, we provide a unified study of coding, grouping, and adaptive selection techniques in terms of wall-clock runtime, communication load and computation load. Furthermore, in order to combine the benefits of GC and grouping in terms of robustness to stragglers with the communication and computation load gains of adaptive selection of LAG, novel strategies, named lazily aggregated GC (LAGC) and Grouped-LAG (G-LAG), are introduced.

A. Related Work

The original work [8] on GC described earlier has been extended in a number of directions. By coding across the elements of gradient vectors, rather than only across different gradient vectors as in [8], [12] proposed a variant of GC that provides a generalized tradeoff in terms of straggler tolerance, computation, and communication loads. By using Reed–Solomon codes, [13] improves the computational complexity of GC. GC techniques that enable the approximate, rather than exact computations of the gradient, were studied in [14]–[16]. References [17]–[19] considered GC for stochastic gradient methods. GC is also part of an active line of research that aims at more generally improving the robustness of distributed computing (see [20]–[24]).

Another related line of work addresses the communication overhead of distributed GD. This can be done by using various methods, such as quantization and sparsification [25]–[27], which generally entail a performance loss in terms of accuracy; duality-based methods [28], [29] and approximate Newton-type methods [30], [31], both of which increase the computational load in order to guarantee a given accuracy; and adaptive batch size selection [32], which can reduce iteration complexity and communication load at a larger computation cost per iteration than conventional gradient methods.

Finally, computation complexity in large-scale machine learning applications has also attracted extensive attention. The standard approach is stochastic GD (SGD), which trains using a random mini-batch of samples at each iteration [33]. A large number of variants of SGD have been proposed that present different tradeoffs among accuracy, communication load, and computation complexity [34]. Typical solutions include quantization and sparsification [35], [36]. Of particular, practical relevance is also asynchronous schemes [37], whereby updates by different workers can be aggregated as they are produced with limited need for locking mechanism [38], [39]. Alternatives to SGD are explored in [40], [41].

B. Main Contributions

This article studies gradient-based optimization in a distributed PS architecture that enables GC, grouping, and adaptive selection. An analysis of wall-clock runtime complexity, communication complexity, and computation complexity is provided to measure the performance of GC, LAG, and of

the newly proposed LAGC and G-LAG methods. The main contributions are summarized as follows.

- 1) A novel strategy, named LAGC, is proposed that is able to leverage the advantages summarized in Table I of GC and LAG via a specific integration of the two techniques. As a special case, we also consider a scheme that only uses grouping and adaptive selection, hence not requiring coding, which is referred to as G-LAG.
- 2) We provide an analysis of time complexity, communication complexity, and computation complexity for the discussed strategies, namely, GD, GC, and G-LAG, under the standard assumption of smooth convex loss. We specifically illustrate the tradeoff among these metrics under Pareto and exponential distributions for the random computing times of the workers. These distributions are representative of high and low tails, respectively, for computing times.
- 3) Finally, we present numerical results on baseline regression tasks in order to study the evolution of the accuracy and communication and computation loads, as functions of wall-clock time.

The rest of this article is organized as follows. Section II describes the PS architecture and the adopted performance metrics. Under this framework, GC and LAG are reviewed in Section III, and the proposed strategies are presented in Section IV. Section V provides the analysis of the adopted metrics along with some numerical illustrations. Section VI presents some numerical examples for a regression task. Finally, Section VII concludes the work and also highlights future research directions.

II. SYSTEM MODEL

We are given a training data set $\mathcal{D} = \{\mathbf{z}_n = (\mathbf{x}_n, y_n)\}_{n=1}^N$, where the explanatory vector $\mathbf{x}_n \in \mathbb{R}^d$ contains d covariates and the label $y_n \in \mathbb{F}$ takes values in a finite discrete set $\mathbb{F}$. The objective is to learn a model parameter vector $\boldsymbol{\theta} \in \mathbb{R}^p$ by minimizing the training loss

$$L(\mathcal{D}; \boldsymbol{\theta}) = \sum_{\mathbf{z}_n \in \mathcal{D}} \ell(\mathbf{z}_n; \boldsymbol{\theta}) \quad (1)$$

where $\ell(\mathbf{z}_n; \boldsymbol{\theta})$ is a loss function that depends on the hypothesis class and on the performance criterion of interest, e.g., quadratic error or cross entropy. To tackle the minimization of function $L(\mathcal{D}; \boldsymbol{\theta})$ over vector $\boldsymbol{\theta}$, we consider methods based on approximate GD steps, whereby parameter $\boldsymbol{\theta}$ is updated iteratively by following the rule:

$$\boldsymbol{\theta}^{i+1} = \boldsymbol{\theta}^i - \alpha \hat{\mathbf{g}}(\boldsymbol{\theta}^i) \quad (2)$$

with α being the stepsize, superscript i indicating the iteration index, and $\hat{\mathbf{g}}(\boldsymbol{\theta}^i)$ being an estimate of the exact gradient $\mathbf{g}(\boldsymbol{\theta}^i) = \nabla L(\mathcal{D}; \boldsymbol{\theta}^i)$. Note that we focus on full gradient techniques that aim at linear convergence rates in terms of number of iterations, and we do not consider SGD methods, which instead can only achieve sublinear convergence rates (see [11]).

A PS framework is commonly adopted to run GD using parallel workers. As shown in Fig. 1, the PS has access to the

Fig. 1. PS model with storage redundancy r .

entire data set $\mathcal{D}$, and it can communicate with M workers, which are denoted by set $\mathcal{M} = \{1, 2, \dots, M\}$. Note that the workers are part of the same computation system, and hence, communication is not subject to privacy constraints. In the following, we describe a framework for a PS-based implementation of updates (2) that allows us to study in a unified way GC [8] and (an equivalent variant of) LAG [10], and to generalize both via the proposed LAGC strategy.

Prior to the start of training, the data set $\mathcal{D}$ is partitioned into S subsets $\mathcal{D}_1, \dots, \mathcal{D}_S$ of equal size, with the partial gradient of each data partition $\mathcal{D}_s$ defined as

$$\mathbf{g}_s(\boldsymbol{\theta}^i) = \sum_{\mathbf{z}_n \in \mathcal{D}_s} \nabla \ell(\mathbf{z}_n; \boldsymbol{\theta}^i). \quad (3)$$

As seen in Fig. 1, to parallelize the computation of the gradient in (2), each worker $m \in \mathcal{M}$ is assigned $S_r = rS/M$ partitions for some integer $1 \leq r \leq M$. The partitions assigned to worker m are denoted as $\mathcal{D}_{m(1)}, \dots, \mathcal{D}_{m(S_r)}$, where $m(j) \in \{1, \dots, S_r\}$ for $j = 1, \dots, S_r$. Integer r is referred to as the *storage redundancy*, since, with an even arrangement of the data partitions, each partition is replicated r times across workers. Note that the choice $r = M$ implies that the entire data set $\mathcal{D}$ can be stored at each worker.

To elaborate on the communication and computation protocol, we define $\boldsymbol{\theta}_m^i$ as the version of the model parameter that worker $m \in \mathcal{M}$ has available at iteration i prior to computation. At each iteration i , a subset of workers, denoted by $\mathcal{M}_D^i \subseteq \mathcal{M}$, is selected to download the model parameter $\boldsymbol{\theta}^i$ from the PS. As shown in Fig. 2, each selected worker $m \in \mathcal{M}_D^i$ sets $\boldsymbol{\theta}_m^i = \boldsymbol{\theta}^i$, and computes the local partial gradients $\mathbf{g}_{m(1)}(\boldsymbol{\theta}^i), \dots, \mathbf{g}_{m(S_r)}(\boldsymbol{\theta}^i)$ over all assigned data partitions. In contrast, each nonselected worker $m \in \mathcal{M} \setminus \mathcal{M}_D^i$ does not download the current model parameter from the PS, and it reuses the locally available parameter vector $\boldsymbol{\theta}_m^{i-1}$ from the previous iteration by setting $\boldsymbol{\theta}_m^i = \boldsymbol{\theta}_m^{i-1}$. As a result, the local parameter vector $\boldsymbol{\theta}_m^i$ is different from global parameter vector $\boldsymbol{\theta}^i$ for all workers $m \in \mathcal{M} \setminus \mathcal{M}_D^i$.

The wall-clock time T_m^i required to complete the computation of all local gradients at each worker $m \in \mathcal{M}_D^i$ is random, with mean ηr , for some $\eta > 0$, proportional to the workload, which is, in turn, proportional to the storage redundancy r . Variables $\{T_m^i\}_{m \in \mathcal{M}_D^i}$ are assumed to be i.i.d. across the worker index m and iteration index i . This standard assumption is motivated by the fact that, in most deployments,

Fig. 2. Illustration of the training protocol with $M = 4$ workers, $\mathcal{M}_D^i = \{1, 2, 3\}$ and $\mathcal{M}_U^i = \{1, 3\}$ at the i th iteration.

the workers are similar machines that operate independently (see [5], [8], [10], [12]). In this article, we will consider as representative examples an exponential distribution with mean ηr and a Pareto distribution with scale-shaped parameter pair $(\eta r(\beta - 1)/\beta, \beta)$, where $\eta > 0$ and $\beta > 1$ are constants. The latter has a higher tail than the former, implying a larger probability of straggling workers.

In order to reduce the time per iteration, as shown in Fig. 2, the PS may only wait for a subset $\mathcal{M}_U^i \subseteq \mathcal{M}_D^i$ of active workers to complete their computations. Each worker $m \in \mathcal{M}_U^i$ uploads a function $\mathbf{f}_m^i \in \mathbb{R}^p$ of the computed partial gradients $\{\mathbf{g}_{m(j)}(\boldsymbol{\theta}^i)\}_{j=1}^{S_r}$ at current iteration back to the PS. Based on the received functions $\{\mathbf{f}_m^i\}_{m \in \mathcal{M}_U^i}$, and possibly also previously received functions $\{\mathbf{f}_m^j\}_{m \in \mathcal{M}_U^i}$, with $j < i$, the PS computes an estimate $\hat{\mathbf{g}}(\boldsymbol{\theta}^{i-1})$ of the gradient. Note that each worker m in subset $\mathcal{M}_D^i \setminus \mathcal{M}_U^i$ downloads the model parameter $\boldsymbol{\theta}^i$ but it does not upload a local update due to the excessive time elapsed for its computation. Similar to each worker in subset $\mathcal{M} \setminus \mathcal{M}_D^i$, the workers in subset $\mathcal{M}_D^i \setminus \mathcal{M}_U^i$ also reuse the outdated gradients $\{\mathbf{g}_{m(j)}(\boldsymbol{\theta}^{i-1})\}_{j=1}^{S_r}$ for the local parameter vector $\boldsymbol{\theta}_m^i = \boldsymbol{\theta}_m^{i-1}$.

The PS then updates the model parameter $\boldsymbol{\theta}^i$ via the rule in (2). The PS also keeps track of the workers' parameter vectors $\{\boldsymbol{\theta}_m^i\}_{m \in \mathcal{M}}$. The next iteration $i + 1$ then starts with the workers in subset $\mathcal{M}_D^{i+1}$ downloading the updated model $\boldsymbol{\theta}^{i+1}$ from the PS. The training continues until a convergence criterion is satisfied or a maximum number of iterations is reached.

Throughout this article, by following [10], [38], and [42], we make the following standard assumptions on the local loss functions $L_s(\boldsymbol{\theta}) = \sum_{\mathbf{z}_n \in \mathcal{D}_s} \ell(\mathbf{z}_n; \boldsymbol{\theta})$ for each data partition $\mathcal{D}_s$ and on the overall loss function $L(\boldsymbol{\theta}) = \sum_{s=1}^S L_s(\boldsymbol{\theta}) = L(\mathcal{D}; \boldsymbol{\theta})$.

Assumption 1: The local loss functions $L_s(\boldsymbol{\theta})$ are L_s -smooth for $s = 1, \dots, S$, for some $L_s > 0$, that is, we have the inequalities

$$\|\nabla L_s(\boldsymbol{\theta}) - \nabla L_s(\boldsymbol{\theta}')\|^2 \leq L_s^2 \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2 \quad (4)$$

for all $\boldsymbol{\theta}$ and $\boldsymbol{\theta}' \in \mathbb{R}^p$; and for $s = 1, \dots, S$, and the loss function $L(\boldsymbol{\theta})$ is L -smooth with $L \leq \sum_{s=1}^S L_s$. From (4), the smoothness parameters L_s determine the rate of

variability of the gradient of each local loss function $L_s(\theta)$. The parameters $\{L_s\}$ are assumed to be known, and they can be practically estimated using techniques, such as those discussed in [43].

Assumption 2: $L(\theta)$ is μ -strongly convex, or more generally, it satisfies the Polyak Łojasiewicz condition for some $\mu > 0$, i.e., it satisfies the inequality

$$2\mu(L(\theta) - L(\theta^*)) \leq \|g(\theta)\|^2 \quad (5)$$

for all $\theta \in \mathbb{R}^p$, where θ^* is a minimum of $L(\theta)$. Note that the minimum θ^* is unique for strongly convex function. From (5), the norm of the gradient can be used as a measure of distance to the optimal value of the loss function.

A. Performance Metrics

We are interested in studying the performance in terms of training accuracy, communication load, and computation load as a function of the wall-clock time. To this end, we start by defining the number I of iterations of the SGD rule (2) carried out by time t as

$$I(t) = \max \left\{ I : \sum_{i=1}^I \max_{m \in \mathcal{M}_U^i} \{T_m^i\} \leq t \right\}. \quad (6)$$

Note that $I(t)$ is a random variable due to the randomness of the times T_m^i and of the subsets $\mathcal{M}_U^i$. Furthermore, the quantity $\max_{m \in \mathcal{M}_U^i} \{T_m^i\}$ represents the time per iteration since the PS waits for all the workers in subset $\mathcal{M}_U^i$ to finish their respective computations. For each iteration i , the *communication load* is defined as the sum of the numbers $|\mathcal{M}_D^i|$ of workers that download model parameters from the PS and of the numbers $|\mathcal{M}_U^i|$ of workers that upload the computed gradients to the PS. It follows that the *communication load* $C(t)$ as a function of time t is given as

$$C(t) = \sum_{i=1}^{I(t)} (|\mathcal{M}_D^i| + |\mathcal{M}_U^i|). \quad (7)$$

We also define the *computation load* as the total number of gradients per data point computed at the workers. The computation load at time t is given by

$$P(t) = \sum_{i=1}^{I(t)} \frac{r}{M} |\mathcal{M}_D^i| \quad (8)$$

since each worker in set $\mathcal{M}_D^i$ computes rN/M local gradients. Note that the workers in $\mathcal{M}_U^i \setminus \mathcal{M}_D^i$ may compute only partially the local gradients, but here, we do not make this distinction since, in practice, this would require additional signaling from PS to the workers. Finally, the *training loss optimality gap* at time t is given by

$$L(t) = L(\theta^{I(t)}) - L(\theta^*). \quad (9)$$

Beside the random tuple $(L(t), C(t), P(t))$, we also consider the *time complexity*, the *communication complexity*, and the *computation complexity*, which are summary metrics that measure as the average time, communication load, and computation load needed to ensure an optimality gap equal to $\epsilon > 0$.

Accordingly, defining as the (random) number of iterations needed to obtain an ϵ -optimality gap, also known as iteration complexity [10], $I_\epsilon = \min\{I : \|L(\theta^I) - L(\theta^*)\|^2 \leq \epsilon\}$, the wall clock time complexity is defined as

$$\bar{T}_\epsilon = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon} \max_{m \in \mathcal{M}_U^i} \{T_m^i\} \right] \quad (10)$$

the communication complexity as

$$\bar{C}_\epsilon = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon} |\mathcal{M}_D^i| + |\mathcal{M}_U^i| \right] \quad (11)$$

and the computation complexity as

$$\bar{P}_\epsilon = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon} \frac{r}{M} |\mathcal{M}_D^i| \right]. \quad (12)$$

We note that we have included in the wall-clock time only the durations of the computation steps, hence excluding the contribution of communications. This allows to more clearly highlighting the tradeoff between computing and communication. A compound wall-clock runtime metric that accounts for both computation and communication can be easily derived from the results in this article.

III. BACKGROUND

In this section, we review the state-of-art techniques GC [8] and LAG [10].

A. Gradient Coding

GC, introduced in [8], is an exact full-GD approach, implementing rule (2) with $\hat{g}(\theta^i) = g(\theta^i)$, that aims at mitigating straggling workers by leveraging storage and computational redundancy. Prior to training, GC replicates each data partition $r > 1$ times across the workers. To this end, data set $\mathcal{D}$ is first divided into $S = M$ partitions $\{\mathcal{D}_s\}_{s=1}^S$. Each partition $\mathcal{D}_s$ is stored at r workers, and we have $S_r = r$ partitions at each worker. Specifically, worker m stores partition $\mathcal{D}_{[m+i]_M}$ with $i = 0, \dots, r-1$, with $[m]_M = \text{mod}_M(m-1) + 1$ and $\text{mod}_M(\cdot)$ being the modulo- M operation (see Fig. 3 for example). At each iteration i , all the workers download the model θ^i from the PS to execute the computations, i.e., we have $\mathcal{M}_D^i = \mathcal{M}$ (recall Fig. 2). Note that, for GC, we, hence, have the local parameter given as $\theta_m^i = \theta^i$ for all the workers. The PS waits only the fastest F workers to finish their computations, yielding the subset $\mathcal{M}_U^i = \{m \in \mathcal{M} : T_m^i \leq T_{F:M}^i\}$, where $T_{F:M}^i$ is the F th order statistic of the variables $\{T_m^i\}_{m=1}^M$. To enable the recovery of the gradient $g(\theta^i)$ at the PS, each of the worker in $\mathcal{M}_U^i$ sends a designed linear combination f_m^i of the computed partial gradients $\{g_{m(i)}(\theta^i)\}_{i=1}^{S_r}$ to the PS. The PS then computes a linear combination of the vectors $\{f_m^i\}_{m \in \mathcal{M}_U^i}$ in order to recover the full gradient $g(\theta^i)$. In order to guarantee the existence of linear encoding and decoding functions that enable the recovery of the full gradient $g(\theta^i)$ for any set of F workers, the inequality $F \geq M - r + 1$ is necessary and sufficient [8].

As a special case of GC, if $r = M$, the entire data set can be stored at each worker. In this case, the PS can wait for

Fig. 3. GC: at current iteration i , the PS can recover the full gradient $g(\theta^i)$ by aggregating the functions f_1^i and f_3^i received from the fastest workers 1 and 3.

the fastest worker only ($F = 1$), and no coding and decoding operations are needed.

Example: Consider $M = 3$ workers, $S_r = 2$ data partitions at each worker, and storage redundancy $r = 2$. GC allows the PS to wait only for the $F = 2$ fastest workers since $F = 2 \geq M - r + 1$. As shown in Fig. 3, this is done by splitting data set $\mathcal{D}$ into $S = M = 3$ partitions $\mathcal{D}_1, \mathcal{D}_2$, and $\mathcal{D}_3$. Worker $m = 1, 2$ stores $S_r = r = 2$ partitions $\mathcal{D}_m$ and $\mathcal{D}_{m+1}$, while and worker 3 stores $\mathcal{D}_3$ and $\mathcal{D}_1$. The three workers compute the linear functions indicated in Fig. 3. It can be easily seen that, by summing functions received from any two workers, the PS can recover the exact full gradient $g(\theta^i)$.

B. Lazily Aggregated Gradient

LAG, proposed in [10], is an approximate GD scheme that judiciously selects the subset $\mathcal{M}_D^i$ of active workers at each iteration in order to reduce communication and computation loads. Unlike GC, LAG does not require storage redundancy, and hence, we have $r = 1$. We can also set without loss of generality $S = M$ and assign each worker m a disjoint data partition $\mathcal{D}_m$ so that $S_r = 1$. The PS determines the subset $\mathcal{M}_D^i$ of active workers at each iteration i on the basis of the estimated change in the gradient for the local loss function $L_m(\theta)$ corresponding to the data partition at worker m as compared with the latest communicated gradient from the worker. Note that, since the subset $\mathcal{M}_D^i$ is determined irrespective of the realization of the computation times, unlike GC, LAG is not tolerant to stragglers.

In the following, we provide a more detailed description of LAG. We specifically follow the LAG-PS strategy introduced in [10]. More precisely, in order to highlight the common elements with GC, the scheme described here is functionally equivalent to LAG-PS, but it differs from it in terms of the way operations are split between encoding functions computed at the workers and decoding functions evaluated at the PS (see Remark 1 in the following for details).

At each iteration i , the PS first determines the subset $\mathcal{M}_D^i$ by checking the following condition for each worker m

$$L_m^2 \|\theta_m^{i-1} - \theta^i\|^2 \geq \frac{\zeta}{\alpha^2 M^2 D} \sum_{d=1}^D \|\theta^{i+1-d} - \theta^{i-d}\|^2 \quad (13)$$

where we recall that L_m is the smoothness constant of the local function $L_m(\theta)$, while $\zeta < 1$ is some constant. By (4),

Fig. 4. LAG: At current iteration i , only worker 1 satisfies Condition (13), and hence only partial gradient $f_1^i = g_1^i$ is communicated to the PS for update.

the left-hand side of (13) represents a bound on the change in the gradient squared norm expected for the local loss at worker m as compared with the last available gradient from the worker. The right-hand side of (13), by the rule (2), represents the per-server average contribution over the most recent D iterations to the approximate gradient norm squared $(1/D) \sum_{d=1}^D \|\hat{g}(\theta^{i-d})\|^2$, scaled by a parameter ζ . The PS selects the workers that satisfy condition (13), i.e., the workers that are expected to have a sizeable difference between the current local gradient and their more recently computed gradients, and hence, may contribute more significantly to the model update (2). This yields the subset $\mathcal{M}_D^i = \{m \in \mathcal{M} : (13) \text{ holds}\}$.

With the received parameter θ^i , each worker $m \in \mathcal{M}_D^i$ computes the local gradient $g_m(\theta^i)$ in (3) and then sends the result $f_m^i = g_m(\theta^i)$ to the PS. In contrast, for each unselected user in subset $\mathcal{M} \setminus \mathcal{M}_D^i$, which does not compute an update, the outdated partial gradients $\{g_m(\theta_m^{i-1})\}_{m \in \mathcal{M} \setminus \mathcal{M}_D^i}$ from iteration $i - 1$ is reused. Note that we have $\mathcal{M}_U^i = \mathcal{M}_D^i$. By combining both the computed gradients and the outdated gradients, the PS estimates the full gradient

$$\hat{g}(\theta^i) = \sum_{m \in \mathcal{M}_D^i} g_m(\theta^i) + \sum_{m \in \mathcal{M} \setminus \mathcal{M}_D^i} g_m(\theta_m^{i-1}). \quad (14)$$

The PS computes θ^{i+1} , and updates the variables $\{\theta_m^i = \theta^i\}_{m \in \mathcal{M}_D^i}$ and $\{\theta_m^i = \theta_m^{i-1}\}_{m \in \mathcal{M} \setminus \mathcal{M}_D^i}$ before moving to the next iteration.

Example: Consider again $M = 3$ workers. As shown in Fig. 4, at each iteration i , the PS checks condition (13) for each worker. Assume that only worker 1 satisfies it, and hence, we have $\mathcal{M}_D^i = \mathcal{M}_U^i = \{1\}$. Therefore, only worker 1 downloads the current model θ^i from the PS. It then computes the local gradient $g_1(\theta^i)$ and uploads $f_1^i = g_1(\theta^i)$ to the PS. Combining with the outdated partial gradients $g_2(\theta_2^{i-1})$ and $g_3(\theta_3^{i-1})$ of workers 2 and 3, the PS recovers the estimate $\hat{g}(\theta^i) = g_1(\theta^i) + g_2(\theta_2^{i-1}) + g_3(\theta_3^{i-1})$. The PS then updates θ^{i+1} and $\theta_1^i = \theta^i$ and $\theta_m^i = \theta_m^{i-1}$ for $m = 2, 3$. The next iteration $i + 1$ then continues with a check of condition (13) by the PS in the same way.

Remark 1: In the original LAG in [10], instead of uploading $g_m(\theta^i)$, each worker $m \in \mathcal{M}_U^i$ transmits to the PS the gradient change $f_m^i = g_m(\theta^i) - g_m(\theta_m^{i-1})$. The PS then estimates the full gradient $g(\theta^i)$ by summing the vectors $\{f_m^i\}_{m \in \mathcal{M}_U^i}$ to the

previous gradient estimate $\hat{\mathbf{g}}(\boldsymbol{\theta}^{i-1})$ as [10, eq. (6)]

$$\begin{aligned}\hat{\mathbf{g}}(\boldsymbol{\theta}^i) &= \hat{\mathbf{g}}(\boldsymbol{\theta}^{i-1}) + \sum_{m \in \mathcal{M}_D^i} \mathbf{f}_m^i \\ &= \hat{\mathbf{g}}(\boldsymbol{\theta}^{i-1}) + \sum_{m \in \mathcal{M}_D^i} (\mathbf{g}_m(\boldsymbol{\theta}^i) - \mathbf{g}_m(\boldsymbol{\theta}_m^{i-1})).\end{aligned}\quad (15)$$

The update (14) considered here yields by direct computation the equalities

$$\begin{aligned}\hat{\mathbf{g}}(\boldsymbol{\theta}^i) &= \sum_{m \in \mathcal{M}} \mathbf{g}_m(\boldsymbol{\theta}_m^{i-1}) + \sum_{m \in \mathcal{M}_D^i} (\mathbf{g}_m(\boldsymbol{\theta}^i) - \mathbf{g}_m(\boldsymbol{\theta}_m^{i-1})) \\ &\stackrel{(a)}{=} \left(\sum_{m \in \mathcal{M} \setminus \mathcal{M}_D^{i-1}} \mathbf{g}_m(\boldsymbol{\theta}_m^{i-2}) + \sum_{m \in \mathcal{M}_D^{i-1}} \mathbf{g}_m(\boldsymbol{\theta}^{i-1}) \right) \\ &\quad + \sum_{m \in \mathcal{M}_D^i} (\mathbf{g}_m(\boldsymbol{\theta}^i) - \mathbf{g}_m(\boldsymbol{\theta}_m^{i-1})) \\ &= \hat{\mathbf{g}}(\boldsymbol{\theta}^{i-1}) + \sum_{m \in \mathcal{M}_D^i} (\mathbf{g}_m(\boldsymbol{\theta}^i) - \mathbf{g}_m(\boldsymbol{\theta}_m^{i-1}))\end{aligned}\quad (16)$$

where (a) holds because we have $\boldsymbol{\theta}_m^{i-2} = \boldsymbol{\theta}_m^{i-1}$ for any $m \in \mathcal{M} \setminus \mathcal{M}_D^{i-1}$ at iteration $i-1$. From (15) and (16), we can see that the described LAG and the original LAG in [10] are equivalent since they use the same gradient estimate in the update (2).

IV. LAZILY AGGREGATED GRADIENT CODING

In this section, we propose a strategy, named LAGC, that aims at exploring the tradeoff between the robustness to stragglers of GC and the computation and communication efficiency of LAG by generalizing both schemes. The idea is to cluster all the workers into groups, treat each group as a single worker in LAG, and, within each group, mitigate the effect of stragglers by applying for computation redundancy as in GC. In a manner similar to LAG, the PS selects only groups of workers that have collectively large expected new contributions to the gradient. The tradeoff between the robustness to stragglers of GC and the computation and communication efficiency of LAG can be controlled by selecting the size of the groups, with GC and LAG being two extreme special cases. In particular, increasing the size of the groups enhances the capability of LAGC to mitigate stragglers by utilizing storage redundancy within each group. Conversely, reducing the size of the groups gives the PS more flexibility on the selection of the subset of workers to activate at each iteration, and hence potentially reducing the computation complexity.

To elaborate, LAGC divides all the M workers into $G = M/M_G$ groups $\mathcal{G}_1, \dots, \mathcal{G}_G$, each having M_G workers, where design parameter M_G is an integer divisor of M . Data set $\mathcal{D}$ is split into $S = G$ partitions $\mathcal{D}_1, \dots, \mathcal{D}_G$ of equal size, with each partition $\mathcal{D}_g$ assigned exclusively to group $\mathcal{G}_g$, for any $g \in [G]$. In each group, partition $\mathcal{D}_g$ is assigned to the M_G workers in $\mathcal{G}_g$ by following GC. Accordingly, we further split $\mathcal{D}_g$ into M_G equal-size batches $\mathcal{D}_{g,1}, \dots, \mathcal{D}_{g,M_G}$. Each partition is then stored at $r_G = \min\{r, M_G\}$ workers in each group. Specifically, the m th worker stores partition $\mathcal{D}_{g,[m+i]_{M_G}}$ for $i = 0, \dots, r_G - 1$. Note that, when the design parameter

M_G is selected as $M_G < r$, the workers' storage redundancy is underused. Furthermore, if $M_G \leq r$, all workers in a group $\mathcal{G}_g$ can fully store the partition $\mathcal{D}_g$, and hence, as in LAG and as further discussed in the following, no coding is needed.

We denote as $\boldsymbol{\theta}_g^i$ the model parameter available at all workers in group $\mathcal{G}_g$ at iteration i , i.e., we have $\{\boldsymbol{\theta}_m^i = \boldsymbol{\theta}_g^i\}_{m \in \mathcal{G}_g}$. At each iteration i , the PS determines the subset of groups to be activated first. To this end, the PS evaluates the condition

$$L_g^2 \|\boldsymbol{\theta}_g^{i-1} - \boldsymbol{\theta}^i\|^2 \geq \frac{M_G^2 \zeta}{\alpha^2 M^2 D} \sum_{d=1}^D \|\boldsymbol{\theta}^{i+1-d} - \boldsymbol{\theta}^{i-d}\|^2 \quad (17)$$

for all groups $\{\mathcal{G}_g\}_{g=1}^G$, where we write L_g for the smoothness constant of the local function $L_g(\boldsymbol{\theta}) = \sum_{\mathbf{z}_n \in \mathcal{D}_g} \ell(\mathbf{z}_n; \boldsymbol{\theta})$ of each group $\mathcal{G}_g$. In a manner similar to (13), condition (17) is satisfied by groups that are expected to have a large new contribution to the model update (2). This is because, the right-hand side of (17), by the rule (2), represents the per-group average contribution over the most recent D iterations to the approximate gradient norm squared $(1/D) \sum_{d=1}^D \|\hat{\mathbf{g}}(\boldsymbol{\theta}^{i-d})\|^2$; while, by (4), the left-hand side of (17) represents a bound on the change in the gradient squared norm expected for the local loss at group $\mathcal{G}_g$ as compared with the available gradient of the group at last iteration. The PS selects all the groups that satisfy condition (17), i.e., the subset of groups $\mathcal{T}^i = \{g \in [G] : (17) \text{ hold}\}$. All the workers in each group $\mathcal{G}_g$, with $g \in \mathcal{T}^i$, download the parameter $\boldsymbol{\theta}^i$ from the PS, and the subset $\mathcal{M}_D^i$ of active workers is, hence, given as $\mathcal{M}_D^i = \bigcup_{g \in \mathcal{T}^i} \mathcal{G}_g$.

For each selected group with index $g \in \mathcal{T}^i$, the PS only waits for the fastest F workers to finish their computations. As in GC, in order to guarantee the recovery of the full local gradient $\mathbf{g}_g(\boldsymbol{\theta}^i) = \sum_{\mathbf{z}_n \in \mathcal{D}_g} \nabla L(\mathbf{z}_n; \boldsymbol{\theta}^i)$ for each group g , the condition $F \geq M_G - r_G + 1$ is necessary and sufficient [8]. As a result, we have the subset $\mathcal{M}_U^i = \bigcup_{g \in \mathcal{T}^i} \{m \in \mathcal{G}_g : T_m^i \leq T_{F:M_G}^{g,i}\}$ of workers uploading their gradients to the PS at iteration i , where $T_{F:M_G}^{g,i}$ is the F th order statistic of the variables $\{T_m^i\}_{m \in \mathcal{G}_g}$. Note, in particular, that, if $M_G = r$ yielding $r_G = M_G$, the PS can wait for the fastest worker in the group, i.e., $F = 1$. In fact, each worker can compute directly the gradient $\mathbf{g}_g(\boldsymbol{\theta}^i)$ for the partition $\mathcal{D}_g$ allocated to group $\mathcal{G}_g$. For the more general case that M_G is not equal to r , each of the $F > 1$ nonstragglers m uploads a linear combination $\mathbf{f}_m^i$ of the computed partial gradients $\{\mathbf{g}_{\mathcal{D}_{g,m(j)}}(\boldsymbol{\theta}^i)\}_{j=1}^{r_G}$ and the PS decodes the local gradient $\mathbf{g}_g(\boldsymbol{\theta}^i)$ by following GC.

Finally, by combining with the outdated gradients $\{\mathbf{g}_g(\boldsymbol{\theta}_g^{i-1})\}_{g \in [G] \setminus \mathcal{T}^i}$ from the inactive groups, the PS estimates the full global gradient $\mathbf{g}(\boldsymbol{\theta}^i)$ as

$$\hat{\mathbf{g}}(\boldsymbol{\theta}^i) = \sum_{g \in \mathcal{T}^i} \mathbf{g}_g(\boldsymbol{\theta}^i) + \sum_{g \in [G] \setminus \mathcal{T}^i} \mathbf{g}_g(\boldsymbol{\theta}_g^{i-1}) \quad (18)$$

which is used in the update rule (2). The PS also updates the variables $\{\boldsymbol{\theta}_g^i = \boldsymbol{\theta}_g^{i-1}\}_{g \in [G] \setminus \mathcal{T}^i}$ and $\{\boldsymbol{\theta}_g^i = \boldsymbol{\theta}^i\}_{g \in \mathcal{T}^i}$. The full algorithm is summarized in Algorithm 1.

Remark 2: As discussed, LAGC generalizes both LAG and GC: when choosing $M_G = M$ and $\zeta = 0$, LAGC reduces to

Algorithm 1 LAGC

-
- 1: **Input:** number of groups $G = M/M_G$, stepsize $\alpha > 0$, smoothness constants $\{L_g\}_{g=1}^G$, number of nonstraggling servers per group $F \geq M_G - r_G + 1$ ($r_G = \min\{r, M_G\}$)
 - 2: **Initialize:** $\theta, \{\theta_g^0\}_{g=1}^G$
 - 3: repeat $i = 1$
 - 4: for each group $\mathcal{G}_g$ that satisfies (17)
 - 5: $\triangleright$ all the workers in $\mathcal{G}_g$ download θ^i from the PS
 - 6: $\triangleright$ all workers in $\mathcal{G}_g$ compute the gradient $\mathbf{g}_g(\theta^i)$
 - 7: $\triangleright$ the F fastest workers send GC-encoded functions $\{f_m^i\}$ to the PS
 - 8: the PS recovers gradients $\{\mathbf{g}_g(\theta^i)\}_{g \in \mathcal{I}^i}$ using GC decoding and estimates $\hat{\mathbf{g}}(\theta^i)$ using (18)
 - 9: the PS updates θ^i via (2) and sets $\{\theta_g^i = \theta_g^{i-1}\}_{g \in [G] \setminus \mathcal{I}^i}$
 - 10: until convergence criterion is satisfied
-

Algorithm 2 G-LAG

-
- 1: **Input:** number of groups $G = M/M_G$, stepsize $\alpha > 0$, smoothness constants $\{L_g\}_{g=1}^G$, number of nonstraggling servers per group $F \geq 1$ ($r_G = M_G$)
 - 2: **Initialize:** $\theta, \{\theta_g^0\}_{g=1}^G$
 - 3: repeat $i = 1$
 - 4: for each group $\mathcal{G}_g$ that satisfies (17)
 - 5: $\triangleright$ all the workers in $\mathcal{G}_g$ download θ^i from the PS
 - 6: $\triangleright$ all workers in $\mathcal{G}_g$ compute the gradient $\mathbf{g}_g(\theta^i)$
 - 7: $\triangleright$ the fastest worker sends $\mathbf{g}_g(\theta^i)$ directly to the PS
 - 8: the PS aggregates gradients $\{\mathbf{g}_g(\theta^i)\}_{g \in \mathcal{I}^i}$ and estimates $\hat{\mathbf{g}}(\theta^i)$ using (18)
 - 9: the PS updates θ^i via (2) and sets $\{\theta_g^i = \theta_g^{i-1}\}_{g \in [G] \setminus \mathcal{I}^i}$
 - 10: until convergence criterion is satisfied
-

GC, while setting $M_G = 1$ recovers LAG. Intermediate values of M_G yield novel schemes.

Remark 3: Setting the number of groups to be smaller or equal to the storage redundancy, i.e., $M_G \leq r$, yields a novel scheme that does not require coding within each group, while still benefiting from both robustness to stragglers and reduced computation complexity. With this choice, each worker $m \in \mathcal{G}_g$ stores the entire data partition $\mathcal{D}_g$ for group $g \in [G]$. Hence, for each selected group $\mathcal{G}_g$ with $g \in \mathcal{I}^i$, the PS only needs to wait for the fastest worker (i.e., $F = 1$), since the latter can send the desired gradient $\mathbf{f}_m^i = \mathbf{g}_g(\theta^i)$ directly to the PS. To highlight the fact that no coding is involved, we refer to this set of schemes as G-LAG (Algorithm 2). We note that setting $\zeta = 0$, and hence selecting all groups at all times, G-LAG reduces to grouped-GD (G-GD) [9].

Example: Consider $M = 6$ workers and storage redundancy $r = 2$. In this case, M_G can take values $M_G = 2$ or 3 , apart from the cases $M_G = 1$ and 6 , corresponding to LAG and GC, respectively. For $M_G = 2$, we have the G-LAG scheme described in Remark 3, whereby each worker in each group can store the entire data partition of that group and no coding is necessary, as shown in Fig. 5(a). In the example of Fig. 5(a), the PS selects only group $\mathcal{G}_1$ using condition (17), and the PS obtains $\mathbf{g}_1(\theta^i) = \mathbf{f}_1^i$ directly from the fastest worker in $\mathcal{G}_1$.

Fig. 5. LAGC with two different grouping strategies. (a) $M_G = 2$. (b) $M_G = 3$.

For $M_G = 3$, the workers are clustered into two groups and the partition for each group is divided into two parts, each stored at two workers with redundancy $r_G = 2$, as shown in Fig. 5(b). In the illustration of Fig. 5(b), only group 1 satisfies condition (17). Hence, workers 1, 2, and 3 in group $\mathcal{G}_1$ download the model θ^i from the PS, and the PS waits for the fastest $F = 2 \geq M_G - r_G + 1$ workers to finish their computations. GC is used to recover the gradient $\mathbf{g}_1(\theta^i)$ from group $\mathcal{G}_1$. By summing up the outdated partial $\mathbf{g}_2(\theta_2^{i-1})$, the PS estimates the global gradient as $\hat{\mathbf{g}}(\theta^i) = \mathbf{g}_1(\theta^i) + \mathbf{g}_2(\theta_2^{i-1})$.

V. ANALYSIS

In this section, we analyze the *time complexity* (10), *communication complexity* (11), and *computation complexity* (12) of all the schemes considered earlier, which are summarized in Table II. Note that, as a reference, we also study standard GD, which is a special case of GC without data redundancy, i.e., with $r = 1$; and G-GD, which, as seen, is an extreme case of G-LAG where all the groups are active at each iteration. The results of the analysis in this section are summarized in Section V-D via numerical illustrations.

In order to derive the mentioned metrics, we first discuss the iteration complexity I_ϵ in (6). A standard result is that GD has iteration complexity

$$I_\epsilon^{\text{GD}} = \bar{I}_\epsilon = \kappa \log\left(\frac{\Delta L^0}{\epsilon}\right) \quad (19)$$

with $\kappa = L/\mu$ being the condition number for the training empirical loss (1) and $\Delta L^0 = L(\theta^0) - L(\theta^*)$ being the difference between the loss $L(\theta^0)$ at the initial iteration θ^0 and the loss at the optimal point θ^* [44]. We note that this result is obtained by choosing a stepsize $\alpha < 1/L$. By constructions, GC and G-GD have the same iteration complexity

TABLE II
SUMMARY OF TECHNIQUES CONSIDERED IN THIS WORKER

	Coding	Adaptive selection	Grouping
GD	×	×	×
GC [8]	✓	×	×
LAG [10]	×	✓	×
G-GD [9]	×	×	✓
G-LAG [this paper]	×	✓	✓
LAGC [this paper]	✓	✓	✓

$I_{\epsilon}^{\text{GC}} = I_{\epsilon}^{\text{G-GD}} = \bar{I}_{\epsilon}$. For LAG, as shown in [10], by choosing the stepsize as $\alpha < 1/L$, we have the iteration complexity

$$I_{\epsilon}^{\text{LAG}} = \frac{\bar{I}_{\epsilon}}{\alpha L} \quad (20)$$

which shows the same scaling with ϵ and κ as GD. Finally, for LAGC, the iteration complexity is by construction equivalent to that of LAG with G workers. In fact, since LAGC treats each group of workers as a single worker, its operation is equivalent to a system with a smaller number of workers. Given that the iteration complexity of LAG does not depend on the number of workers, we have $I_{\epsilon}^{\text{LAGC}} = I_{\epsilon}^{\text{LAG}} = \bar{I}_{\epsilon}/(\alpha L)$.

A. Wall-Clock Time Complexity

We proceed to analyze the average wall-clock time complexity (10). To elaborate, we define as $T_{a:b}$ the a th order statistics of i.i.d. variables $\{T_i\}_{i=1}^b$, that is, the a th smallest value in the set $\{T_i\}_{i=1}^b$, where a and b are two integers satisfying $a \leq b$. The average $\bar{T}_{a:b} = \mathbb{E}[T_{a:b}]$ can be written in closed form for the two representative distributions considered here. In particular, we have

$$\bar{T}_{a:b}(r) := \mathbb{E}[T_{a:b}] = \eta r (H_b - H_{b-a}) \quad (21)$$

for the case of exponential distribution with mean ηr [45], where $H_a = \sum_{k=1}^a 1/k$ is the a th harmonic number; and

$$\begin{aligned} \bar{T}_{a:b}(r) &:= \mathbb{E}[T_{a:b}] \\ &= \frac{\eta r (1 - \beta) \Gamma(b - a + 1 - 1/\beta) \Gamma(n + 1)}{\beta \Gamma(b - a + 1) \Gamma(b + 1 - 1/\beta)} \end{aligned} \quad (22)$$

for the case of Pareto distribution with scale-shape pair $(\eta r(\beta - 1)/\beta, \beta)$ with $\beta > 1$, where $\Gamma(x)$ is Gamma function given by $\Gamma(x) = \int_0^{\infty} t^{x-1} e^{-t} dt$ [46]. Note that both averages (21) and (22) increase with the mean ηr of each variable T_i and with b . To ease the notation, we also write $\bar{T}_a(r) = \bar{T}_{a:a}(r)$ in the following.

GD: The average runtime of each iteration for GD is given as $\bar{T}_M(1)$, since the PS waits for all M servers to complete their computations at each iteration and no computational redundancy is leveraged, i.e., $r = 1$. This yields the overall runtime

$$\bar{T}_{\epsilon}^{\text{GD}} = \bar{I}_{\epsilon} \bar{T}_M(1). \quad (23)$$

GC: With GC, at each iteration $i \in I_{\epsilon}^{\text{GC}}$, the PS waits only for the set $\mathcal{M}_U^i$ of the fastest $F \geq M - r + 1$ workers to finish their computations, yielding the average runtime $\bar{T}_{F:M}$. As a result, the overall runtime is given as

$$\bar{T}_{\epsilon}^{\text{GC}} = \bar{I}_{\epsilon} \bar{T}_{F:M}(r). \quad (24)$$

Comparing with (23), GC can reduce the time complexity if $\bar{T}_{F:M}(r) < \bar{T}_M(1)$.

LAG: At each iteration $i \in I_{\epsilon}^{\text{LAG}}$ of the LAG scheme which assumes $r = 1$, all the selected workers in subset $\mathcal{M}_U^i$ have to complete their computations, yielding the average runtime $\bar{T}_{|\mathcal{M}_U^i|}(1)$. From [10, Lemma 4], given an integer $d \in \{0, 1, \dots, D\}$, a worker with smoothness parameter L_m satisfying the inequalities

$$\bar{L}_{d+1}^2 < L_m^2 < \bar{L}_d^2, \quad \text{with } \bar{L}_d^2 = \frac{\xi}{D d \alpha^2 M^2} \quad (25)$$

is selected in at most $I_{\epsilon}^{\text{LAG}}/(d+1)$ iterations. A larger smoothness constant, and hence, a less sensitive gradient cause a worker to be selected less frequently. Using this result, we can bound the average number $(1/I_{\epsilon}^{\text{LAG}}) \sum_{i=1}^{I_{\epsilon}^{\text{LAG}}} |\mathcal{M}_U^i|$ of selected servers per iteration as

$$\frac{1}{I_{\epsilon}^{\text{LAG}}} \sum_{i=1}^{I_{\epsilon}^{\text{LAG}}} |\mathcal{M}_U^i| \leq M \sum_{d=0}^D \frac{h(d)}{d+1} := \bar{M} \leq M \quad (26)$$

where we have defined the function $h(d) = (1/M) \sum_{m \in \mathcal{M}} \mathbb{1}(\bar{L}_{d+1}^2 < L_m^2 < \bar{L}_d^2)$ with $\bar{L}_0 = \bar{L}_{D+1} = 0$. Note that function $h(d)$ indicates the fraction of the workers satisfying condition (25). Therefore, parameter $\bar{M}$ increases when all smoothness constant $\{L_m\}$ decrease. From [10], we have that $\sum_{d=0}^D h(d)/(d+1) \leq 1$, i.e., $\bar{M} \leq M$. As proved in the Appendix, the bound (26) can be used in turn to bound the overall time as

$$\bar{T}_{\epsilon}^{\text{LAG}} = \sum_{i=1}^{I_{\epsilon}^{\text{LAG}}} \bar{T}_{|\mathcal{M}_U^i|} \leq \frac{\bar{I}_{\epsilon}}{\alpha L} \bar{T}_{\bar{M}}(1). \quad (27)$$

Comparing with (23), we see that LAG can only decrease the time complexity as compared with GD if $\alpha \approx 1/L$ and the smoothness constant $\{L_m\}$ are large.

G-GD: Based on grouping, at each iteration i , the PS waits for the fastest $F \geq M_G - r_G + 1$ workers in each group to finish their computations. Hence, the runtime T_g^G for each group is the F th order statistic of the random times $\{T_i\}_{i \in \mathcal{G}_g}$ of the workers in group $\mathcal{G}_g$. In a manner similar to (21) and (22), we define $\bar{T}_{a:M_G}^G(r)$ to be the expectation of the a th order statistics of the random variables $\{T_g^G\}_{g=1}^{M_G}$, and we denote $\bar{T}_{M_G:M_G}^G(r) = \bar{T}_{M_G}^G(r)$. Unlike (21) and (22), this expectation does not generally have a closed form, but it can be computed numerically as

$$\bar{T}_{a:M_G}^G(r) = \int_0^{+\infty} \left(1 - (F^G(x))^a\right) dx, \quad (28)$$

where we have defined the cumulative distribution function (CDF) $F^G(x) = \sum_{j=F}^{M_G} \binom{M_G}{j} (F(x))^j (1 - F(x))^{M_G-j}$ of each variable T_g^G [47]. With these definitions, the overall runtime of G-GD is given as

$$\bar{T}_{\epsilon}^{\text{G-GD}} = I_{\epsilon}^{\text{G-GD}} \bar{T}_{M_G}^G = \bar{I}_{\epsilon} \bar{T}_{M_G}^G. \quad (29)$$

Based on (29), G-GD can reduce the time complexity if $\bar{T}_{M_G}^G(r) < \bar{T}_M(r)$.

LAGC: At each iteration $i \in I_{\epsilon}^{\text{LAGC}}$, LAGC selects a subset $\mathcal{I}^i$ of groups $\mathcal{G}_g$ with $g \in \mathcal{I}^i$. Using the same

approach discussed earlier from [10], the average group size $(1/I_\epsilon^{\text{LAGC}}) \sum_{i=1}^{I_\epsilon^{\text{LAGC}}} |\mathcal{I}^i|$ can be upper bounded as

$$\frac{1}{I_\epsilon^{\text{LAGC}}} \sum_{i=1}^{I_\epsilon^{\text{LAGC}}} |\mathcal{I}^i| \leq G \sum_{d=0}^D \frac{h_G(d)}{d+1} := \bar{G} \leq G \quad (30)$$

where we have defined the function $h_G(d) = (1/G) \sum_{g \in [G]} \mathbb{1}(\bar{L}_{G,d+1}^2 < L_g^2 < \bar{L}_{G,d}^2)$, with $\bar{L}_{G,d}^2 = \xi/(Dd\alpha^2 G^2)$ and $\bar{L}_{G,0} = \bar{L}_{G,D+1} = 0$. This distribution has the same interpretation given earlier for function $h(d)$, with groups replacing workers. Similarly, we have that $\sum_{d=0}^D h_G(d)/(d+1)$, i.e., $\bar{G} \leq G$. To allow all the selected groups in subset $\mathcal{I}^i$ to complete the computations, the average runtime of each iteration is given as $\bar{T}_{|\mathcal{I}^i|:M_G}^G(r)$. As a result, the overall average runtime can be upper bounded as

$$\bar{T}_\epsilon^{\text{LAGC}} = \sum_{i=1}^{I_\epsilon^{\text{LAGC}}} \bar{T}_{|\mathcal{I}^i|}^G \leq \frac{\bar{I}_\epsilon}{\alpha L} \bar{T}_G^G(r). \quad (31)$$

Comparing with (23), LAGC can hence outperform GD in terms of time complexity if $\bar{T}_G^G(r)/(\alpha L) \leq \bar{T}_M(1)$.

B. Communication Complexity

We proceed to investigate communication complexity.

GD: With the conventional GD, the overall communication load is given as

$$\bar{C}_\epsilon^{\text{GD}} = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{GD}}} |\mathcal{M}_D^i| + |\mathcal{M}_U^i| \right] = 2M\bar{I}_\epsilon \quad (32)$$

since all M workers download and upload model parameters from the PS.

GC: At each iteration $i \in I_\epsilon^{\text{GC}}$ of GC, there are $|\mathcal{M}_D^i| = M$ workers downloading the model parameter and $|\mathcal{M}_U^i| = F$ workers uploading the computed results. Hence, the overall communication load is given as

$$\bar{C}_\epsilon^{\text{GC}} = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{GC}}} |\mathcal{M}_D^i| + |\mathcal{M}_U^i| \right] = \bar{I}_\epsilon(M + F). \quad (33)$$

Comparing (33) with (35), we observe that GC reduces the communication complexity by a factor $2/(1 + F/M)$, which can be as large as 2 if F/M is small enough.

LAG: At each iteration $i \in I_\epsilon^{\text{LAG}}$ of LAG, all the workers in subset $\mathcal{M}_D^i$ download from the PS and also upload model parameters to the PS. Hence, the overall communication load is given as

$$\begin{aligned} \bar{C}_\epsilon^{\text{LAG}} &= \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{LAG}}} |\mathcal{M}_D^i| + |\mathcal{M}_U^i| \right] \\ &= 2\mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{LAG}}} |\mathcal{M}_U^i| \right] \leq 2\bar{M} \frac{\bar{I}_\epsilon}{\alpha L} \end{aligned} \quad (34)$$

with $\bar{M}$ defined in (26). Since we have $\bar{M} \leq M$, LAG can reduce the communication complexity as long as αL is not too small.

G-GD: With G-GD, the overall communication load within $I_\epsilon^{\text{G-GD}}$ iterations is given as

$$\begin{aligned} \bar{C}_\epsilon^{\text{G-GD}} &= \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{G-GD}}} |\mathcal{M}_D^i| + |\mathcal{M}_U^i| \right] \\ &= (M + G(M_G - r_G + 1))\bar{I}_\epsilon \end{aligned} \quad (35)$$

since all M workers download the model parameters but only the fastest $M_G - r_G + 1$ workers in each group upload the computed results. Hence, G-GD reduces the communication complexity as compared with GD by a factor that increases with the computation redundancy factor r_G .

LAGC: At each iteration $i \in I_\epsilon^{\text{LAGC}}$ of LAGC, $|\mathcal{I}^i|$ groups of workers are chosen to download the model parameter, which amount to $|\mathcal{M}_D^i| = |\mathcal{I}^i|M_G$. In each group, $F \geq M_G - r_G + 1$ workers upload their computations, yielding $|\mathcal{M}_U^i| = |\mathcal{I}^i|F$. As a result, the overall communication complexity within I_ϵ^{LAGC} iterations is given as

$$\begin{aligned} \bar{C}_\epsilon^{\text{LAGC}} &= \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{LAGC}}} |\mathcal{M}_D^i| + |\mathcal{M}_U^i| \right] \\ &= \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{LAGC}}} |\mathcal{I}^i|(M_G + F) \right] \\ &\leq (M_G + F)\bar{G} \frac{\bar{I}_\epsilon}{\alpha L} \end{aligned} \quad (36)$$

with $\bar{G}$ defined in (30). Since the average number $M_G\bar{G} \leq M_G G = M$ of selected workers for download and the number $F\bar{G} \leq M$ of selected workers for upload are both no larger than the total number M of workers, LAGC can outperform GD in terms of communication complexity when αL is close to one.

C. Computation Complexity

Finally, we evaluate the computation complexity (12).

GD: The overall computation complexity over I_ϵ^{GD} iterations for GD is given as

$$\bar{P}_\epsilon^{\text{GD}} = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{GD}}} \frac{1}{M} |\mathcal{M}_D^i| \right] = \bar{I}_\epsilon \quad (37)$$

since the gradient for each data point is computed as each iteration.

GC: With a data redundancy $r \leq 1$ at the workers, the overall computation complexity of GC is given as

$$\bar{P}_\epsilon^{\text{GC}} = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{GC}}} \frac{r}{M} |\mathcal{M}_D^i| \right] = r\bar{I}_\epsilon. \quad (38)$$

The computation load of GC is, hence, r times that of GD.

LAG: The overall computation complexity of LAG is similarly bounded as

$$\bar{P}_\epsilon^{\text{LAG}} = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{\text{LAG}}} \frac{1}{M} |\mathcal{M}_D^i| \right] \leq \frac{\bar{M}}{M} \frac{\bar{I}_\epsilon}{\alpha L}. \quad (39)$$

Therefore, LAG can reduce the computation complexity if αL is close to one.

G-GD: The overall computation complexity of G-GD is given as

$$\bar{P}_\epsilon^{G-GD} = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{G-GD}} \frac{r_G}{M} |\mathcal{M}_D^i| \right] = r_G \bar{I}_\epsilon \quad (40)$$

yielding an increase equal to r_G in computation complexity.

LAGC: Finally, the overall computation complexity of LAGC is bounded as

$$\bar{P}_\epsilon^{LAGC} = \mathbb{E} \left[\sum_{i=1}^{I_\epsilon^{LAGC}} \frac{r_G}{M} |\mathcal{M}_D^i| \right] \leq \frac{r_G M_G \bar{G}}{M} \frac{\bar{I}_\epsilon}{\alpha L}. \quad (41)$$

Since we have the inequality $M_G \bar{G}/M \leq 1$ due to adaptive selection, LAGC can reduce the computation complexity of GC when the average number $\bar{G}$ of groups satisfies the inequality $\bar{G} \leq M^2/(M_G r_G)$ as long as $\alpha L \approx 1$.

D. Numerical Illustration

We now provide an illustration of the relative performance of the considered schemes. To this end, we consider a linear regression model $y = \theta^T \mathbf{x}$, where input $\mathbf{x}$ is an input vector obtained by pre-processing a vectorized 784×1 handwritten digit image from the Modified National Institute of Standards and Technology (MNIST) data set. Specifically, the MNIST training data set of 60 000 examples is divided into $S = 20$ partitions of equal size, each including 3 000 images, for each digit $0, 1, \dots, 9$. Each partition $\mathcal{D}_s$, with $s = 1, \dots, S$ is processed as in [10] such that the smoothness constant of the corresponding loss function is set to be equal $L_s = (1.3^{s-1} + 1)^2$. The loss function $L_s(\theta)$ is given as $L_s(\theta) = \|\mathbf{X}_s \theta - \mathbf{y}_s\|^2$, where $\mathbf{X}_s$ is the 3000×784 input matrix for the s th partition $\mathcal{D}_s$ (with each input corresponding to a row of $\mathbf{X}_s$) and the target vector $\mathbf{y}_s$ is given as $\mathbf{y}_s = \mathbf{X}_s \theta^*$ for a randomly generated ground-truth parameter vector θ^* with i.i.d. zero-mean unitary power Gaussian entries. We consider the most recent $D = 10$ iterations to approximate the gradient norm square. Furthermore, for the computing times, the Pareto distribution has shape parameter $\alpha = 1.1$ and scale parameter $0.005/\alpha$, and the exponential distribution has mean parameter 0.05. We evaluate the complexity measures derived above—using the bounds (27), (31), (39), and (41) for LAG and LAGC—for redundancy $r = 4$ for GC and LAGC and hyperparameters $\eta = 0.05$, $\beta = 1.1$, and $\xi = 1$. We also set $F = 17$ for GC and $F = M_G - r_G + 1$ for G-GD and LAGC with M_G set to different values. Time, communication and computation complexities of GD, GC, LAG, G-GD, and LAGC for different values of M_G with Pareto distribution and exponential distribution are shown in Figs. 6 and 7, respectively. As discussed, in Figs. 6 and 7, when $M_G = 1$ and $M_G = M$, LAGC coincides with LAG and GC, respectively, while, when for $M_G \leq r = 4$, we have the (uncoded) G-LAG scheme in Remark 3.

We first consider computing times with Pareto distribution. The high tail of the Pareto distribution entails a high probability that some workers are significantly slower than the rest. As seen in Fig. 6, in this case, both G-GD and GC have a

Fig. 6. Time, communication, and computation complexity measures for GD, GC, LAG and LAGC that guarantees the ϵ -optimality gap with $\epsilon = 10^{-8}$ under the Pareto distribution for the workers' computing times.

Fig. 7. Time, communication, and computation complexity measures for GD, GC, LAG, and LAGC that guarantees the ϵ -optimality gap with $\epsilon = 10^{-8}$ under the exponential distribution for the workers' computing times.

lower time complexity than both GD and LAG thanks to their robustness to stragglers (recall Table I): Although each active worker executes more computations, the reduced requirements on the number of workers that need to complete their computations offset the increased per-server computation load. However, this wall-clock time saving implies a tradeoff with the computation complexity, which is increased. Furthermore, G-GD outperforms GC in all metrics, since more stragglers can be tolerated thanks to grouping for without increasing the computational redundancy. LAG has a larger wall-clock time complexity as compared with G-GD and GC, but it can significantly reduce both communication and computation complexities by selecting a reduced number of workers to be active. The proposed LAGC scheme is seen to be able to harness both the robustness to stragglers of GC and G-GD, which requires a larger M_G and the reduced communication and computation complexity of LAG, which requires a smaller M_G . In fact, in line with the comparison between G-GD and GC, we observe that G-LAG—the special case of LAGC that only with grouping and adaptive selection—yields the best overall performance.

Fig. 8. Loss function $L(t)$, communication load $C(t)$, and computation load $P(t)$ versus wall-clock time t with the Pareto distribution for the computing times.

We now consider the performance under the exponential distribution for the workers' computing times. This distribution has a lower tail, and hence, the workers have comparable computing times with a higher probability than for the Pareto distribution. In this case, in stark contrast to Figs. 6 and 7 shows that GC does not improve the time complexity since the cost resulting from computational redundancy does not offset the savings accrued thanks to the mitigation of stragglers. As compared with LAG, G-GD provides a reduction in wall-clock time due to its stronger ability to tolerate stragglers, implying a tradeoff with the increasing computation complexity. LAG outperforms schemes based solely on grouping or coding in terms of communication and computation complexity. Finally, the proposed G-LAG outperforms all other schemes in terms of computation and time complexities while requiring a larger computation complexity than LAG.

VI. NUMERICAL RESULTS

In this section, we present numerical examples in order to illustrate the loss function $L(t)$ in (9), communication load

Fig. 9. Loss function $L(t)$, communication load $C(t)$, and computation load $P(t)$ versus wall-clock time t with the exponential distribution for the computing times.

$C(t)$ in (7), and computation load $P(t)$ in (8) as function of wall-clock time t for the considered training strategies. We adopt the same linear regression setup described in Section V-D with Pareto and exponential distributions for the random computing times of the workers. In Figs. 8 and 9, we plot the mentioned metrics averaged over 100 random realization of the computing times for GD, GC, LAG, G-GD, LAGC with $M_G = 5 > r$, and G-LAG with $M_G = r$. We note that each curve terminates at the time when the ϵ -optimality is achieved.

Confirming the conclusion from the analysis in Section V-D, under Pareto computing times, the loss functions of GC, G-GD, LAGC, and G-LAG decrease significantly faster than GD and LAG thanks to their robustness to stragglers, as shown in Fig. 8, with G-LAG yielding the steepest descent. LAG is seen to be effective in reducing the communication and computation loads per unit of time. However, G-LAG yields a small overall communication complexity, at the cost of a large computation due to the smaller time for iteration afforded by its robustness to stragglers.

We now turn to consider exponential computing times. In this regard, Fig 9 verifies the conclusion based on the analysis in Section V-D that coding is not advantageous in terms of any performance metric with respect to GC. In contrast, G-GD can be significantly more time efficient due to its stronger robustness to stragglers. By using adaptive selection, the communication loads of LAG, LAGC and G-LAG increase with similar rates as a function of t , but G-LAG has a smaller communication complexity due to the smaller time complexity. Finally, due to computational redundancy, G-LAG, GC, and LAGC have higher computation loads than the other schemes.

VII. CONCLUSION

In this article, we explored the trade-off among wall-clock time, communication, and computation requirements for gradient-based distributed learning by leveraging coding, grouping, and adaptive selection. As summarized in Table I, both coding and grouping provide robustness to stragglers, while the adaptive selection is beneficial to reduce communication and computation loads. We proposed two novel strategies that aim at integrating the benefits of both types of approaches. Through analysis and numerical results, we have concluded that, when the distribution of the computing times of the workers has a low tail, the advantage of straggler mitigation via coding does not compensate for the increased computation load even in terms of wall-clock runtime. In contrast, for both high- and low-tail distributions of the computing times, the proposed G-LAG was seen to strike a desirable balance in terms of wall-clock time and communication overhead, with only a limited increase in computation cost.

This work leaves open a number of research directions. First, it would be interesting to combine stochastic GC [14], [15], [17] with grouping and adaptive selection. Second, the potential advantages of the techniques considered here should be reconsidered for asynchronous implementations, where any server can compute the gradient and send an update to the PS without waiting for the other servers [38]. Lastly, a related issue would be to introduce data privacy requirements [48].

APPENDIX

PROOF OF BOUNDS (27) AND (31)

Define as $F(x)$ the cumulative distribution function (CDF) of each variable T_i . From [49, Lemma 2], we have the equality

$$\bar{T}_a - \bar{T}_{a-1} = \int_{-\infty}^{+\infty} F^{a-1}(x)[1 - F(x)]dx. \quad (42)$$

From (42), since $0 \leq F(x) \leq 1$, the function $\bar{T}_a$ has decreasing increments in a , and hence, it is discrete concave [50]. It follows that for any integers $a_i \leq b$ with $\sum_{i=1}^K a_i = S_a$, we have Jensen's inequality:

$$\bar{T}_{\frac{S_a}{K}} \geq \frac{1}{K} \sum_{i=1}^K \bar{T}_{a_i} \quad (43)$$

as long as S_a/K is an integer. In order to apply (43), we define $a_i = |\mathcal{M}_i^U|$ and $K = |\mathcal{T}_e^{\text{LAG}}|$. We then have the desired inequality in (27). Bound (31) follows from the same argument.

REFERENCES

- [1] J. Dean *et al.*, "Large scale distributed deep networks," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, Lake Tahoe, NV, USA, Dec. 2012, pp. 1223–1231.
- [2] M. Abadi *et al.* (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems*. [Online]. Available: <http://download.tensorflow.org/paper/whitepaper2015.pdf>
- [3] M. Zinkevich, M. Weimer, L. Li, and A. J. Smola, "Parallelized stochastic gradient descent," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, Vancouver, BC, Canada, 2010, pp. 2595–2603.
- [4] A. Smola and S. Narayanamurthy, "An architecture for parallel topic models," *Proc. VLDB Endowment*, vol. 3, nos. 1–2, pp. 703–710, Sep. 2010.
- [5] M. Li, "Scaling distributed machine learning with the parameter server," in *Proc. Symp. Oper. Syst. Design Implement. (OSDI)*, Denver, CO, USA, 2014, pp. 583–598.
- [6] J. Dean and L. A. Barroso, "The tail at scale," *Commun. ACM*, vol. 56, no. 2, pp. 74–80, Feb. 2013.
- [7] L. A. Barroso and U. Hoelzle, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. San Rafael, CA, USA: Morgan & Claypool, 2009.
- [8] R. Tandon, Q. Lei, A. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *Proc. 34th Int. Conf. Mach. Learn.*, Sydney, NSW, Australia, Aug. 2017, pp. 3368–3376.
- [9] E. Ozfatura, D. Gündüz, and S. Ulukus, "Gradient coding with clustering and multi-message communication," 2019, *arXiv:1903.01974*. [Online]. Available: <http://arxiv.org/abs/1903.01974>
- [10] T. Chen, G. Giannakis, T. Sun, and W. Yin, "LAG: Lazily aggregated gradient for communication-efficient distributed learning," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, Montreal, QC, Canada, Dec. 2018, pp. 5050–5060.
- [11] J. Watt, R. Borhani, and A. K. Katsaggelos, *Machine Learning Refined: Foundations, Algorithms, and Applications*. New York, USA: Cambridge Univ. Press, 2016.
- [12] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Stockholm, Sweden, Jul. 2018, pp. 5606–5615.
- [13] W. Halbawi, N. Azizan, F. Salehi, and B. Hassibi, "Improving distributed gradient descent using Reed–Solomon codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Vail, CO, USA, Jun. 2018, pp. 2027–2031.
- [14] N. Raviv, R. Tandon, A. Dimakis, and I. Tamo, "Gradient coding from cyclic MDS codes and expander graphs," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Stockholm, Sweden, Jul. 2018, pp. 4302–4310.
- [15] Z. Charles, D. Papailiopoulos, and J. Ellenberg, "Approximate gradient coding via sparse random graphs," 2017, *arXiv:1711.06771*. [Online]. Available: <http://arxiv.org/abs/1711.06771>
- [16] S. Kadhe, O. Ozan Koyluoglu, and K. Ramchandran, "Gradient coding based on block designs for mitigating adversarial stragglers," 2019, *arXiv:1904.13373*. [Online]. Available: <http://arxiv.org/abs/1904.13373>
- [17] R. Bitar, M. Wootters, and S. El Rouayheb, "Stochastic gradient coding for straggler mitigation in distributed learning," 2019, *arXiv:1905.05383*. [Online]. Available: <http://arxiv.org/abs/1905.05383>
- [18] H. Wang, Z. Charles, and D. Papailiopoulos, "ErasureHead: Distributed gradient descent without delays using approximate gradient coding," 2019, *arXiv:1901.09671*. [Online]. Available: <http://arxiv.org/abs/1901.09671>
- [19] R. K. Maity, A. S. Rawat, and A. Mazumdar, "Robust gradient descent via moment encoding with LDPC codes," 2018, *arXiv:1805.08327*. [Online]. Available: <http://arxiv.org/abs/1805.08327>
- [20] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," *IEEE Trans. Inf. Theory*, vol. 64, no. 3, pp. 1514–1529, Mar. 2018.
- [21] K. Lee, C. Suh, and K. Ramchandran, "High-dimensional coded matrix multiplication," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Aachen, Germany, Jun. 2017, pp. 2418–2422.
- [22] Q. Yu, M. Maddah-Ali, and S. Avestimehr, "Polynomial codes: An optimal design for high-dimensional coded matrix multiplication," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, Long Beach, CA, USA, 2017, pp. 4403–4413.
- [23] M. Fahim, H. Jeong, F. Haddadpour, S. Dutta, V. Cadambe, and P. Grover, "On the optimal recovery threshold of coded matrix multiplication," in *Proc. Allerton Conf. Commun., Control Comput.*, Chicago, IL, USA, Oct. 2017, pp. 1264–1270.

- [24] S. Dutta, V. Cadambe, and P. Grover, "Short-dot: Computing large linear transforms distributedly using coded short dot products," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, Barcelona, Spain, 2016, pp. 2100–2108.
- [25] H. Zhang, J. Li, K. Kara, D. Alistarh, J. Liu, and C. Zhang, "ZipML: Training linear models with end-to-end low precision, and a little bit of deep learning," in *Proc. Int. Conf. Mach. Learn.*, Sydney, NSW, Australia, Aug. 2017, pp. 4035–4043.
- [26] A. T. Suresh, F. X. Yu, S. Kumar, and H. B. McMahan, "Distributed mean estimation with limited communication," in *Proc. Int. Conf. Mach. Learn.*, Sydney, NSW, Australia, Aug. 2017, pp. 3329–3337.
- [27] J. Wangni, J. Wang, J. Liu, and T. Zhang, "Gradient sparsification for communication-efficient distributed optimization," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS) Workshop*, 2018, pp. 1306–1316.
- [28] M. Jaggi *et al.*, "Communication-efficient distributed dual coordinate ascent," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, Montreal, QC, Canada, 2014, pp. 3068–3076.
- [29] C. Ma *et al.*, "Distributed optimization with arbitrary local solvers," *Optim. Methods Softw.*, vol. 32, no. 4, pp. 813–848, Jul. 2017.
- [30] O. Shamir, N. Srebro, and T. Zhang, "Communication-efficient distributed optimization using an approximate Newton-type method," in *Proc. Int. Conf. Mach. Learn.*, Beijing, China, Jun. 2014, pp. 1000–1008.
- [31] Y. Zhang and X. Lin, "DiSCO: Distributed optimization for self-concordant empirical loss," in *Proc. Int. Conf. Mach. Learn.*, Lille, France, Jul. 2015, pp. 362–370.
- [32] M. Jahani *et al.*, "Efficient distributed hessian free algorithm for large-scale empirical risk minimization via accumulating sample strategy," 2018, *arXiv:1810.11507*. [Online]. Available: <http://arxiv.org/abs/1810.11507>
- [33] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proc. COMPSTAT*. Heidelberg, Germany: Physica-Verlag, 2010, pp. 177–186.
- [34] L. Bottou and O. Bousquet, "The tradeoffs of large scale learning," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS) Workshop*, 2007, pp. 161–168.
- [35] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "QSGD: Communication-efficient SGD via gradient quantization and encoding," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS) Workshop*, Long Beach, CA, USA, 2017, pp. 1709–1720.
- [36] J. Bernstein, Y.-X. Wang, K. Azizzadenesheli, and A. Anandkumar, "SIGNSGD: Compressed optimisation for non-convex problems," in *Proc. 35th Int. Conf. Mach. Learn. (ICML)*, Stockholm, Sweden, 2018, pp. 1709–1720.
- [37] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Autom. Control*, vol. AC-31, no. 9, pp. 803–812, Sep. 1986.
- [38] S. Dutta, G. Joshi, S. Ghosh, P. Dube, and P. Nagpurkar, "Slow and stale gradients can win the race: Error-runtime trade-offs in distributed SGD," in *Proc. Int. Conf. Artif. Intell. Statist. (AISTATS)*, Lanzarote, Canary Islands, 2018, pp. 803–812.
- [39] B. Recht, C. Re, S. Wright, and F. Niu, "Hogwild: A lock-free approach to parallelizing stochastic gradient descent," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS) Workshop*, J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, Eds., 2011, pp. 693–701.
- [40] J. Martens and R. Grosse, "Optimizing neural networks with Kronecker-factored approximate curvature," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Lille, France, 2015, pp. 2408–2417.
- [41] J. Ba, R. Grosse, and J. Martens, "Distributed second-order optimization using Kronecker-factored approximations," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, Toulon, France, 2017.
- [42] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," 2016, *arXiv:1606.04838*. [Online]. Available: <http://arxiv.org/abs/1606.04838>
- [43] G. R. Wood and B. P. Zhang, "Estimation of the Lipschitz constant of a function," *J. Global Optim.*, vol. 8, no. 1, pp. 91–103, Jan. 1996.
- [44] M. Schmidt, *Lecture Slides: Rates of Convergence*. Accessed: Jan. 12, 2018. [Online]. Available: <https://www.cs.ubc.ca/~schmidtm/Courses/540-W18/L5.pdf>
- [45] A. Mallick, M. Chaudhari, U. Sheth, G. Palanikumar, and G. Joshi, "Rateless codes for near-perfect load balancing in distributed matrix-vector multiplication," 2018, *arXiv:1804.10331*. [Online]. Available: <http://arxiv.org/abs/1804.10331>
- [46] K. Vannman, "Estimators based on order statistics from a Pareto distribution," *J. Amer. Stat. Assoc.*, vol. 71, no. 355, pp. 704–708, Sep. 1976.
- [47] B. C. Arnold, N. Balakrishnan, and H. N. Nagaraja, *A First Course in Order Statistics*. Philadelphia, PA, USA: SIAM, 2008.
- [48] K. Bonawitz *et al.*, "Practical secure aggregation for privacy-preserving machine learning," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, Dallas, TX, USA, 2017, pp. 1175–1191.
- [49] H. A. David, "Augmented order statistics and the biasing effect of outliers," *Statist. Probab. Lett.*, vol. 36, no. 2, pp. 199–204, Dec. 1997.
- [50] K. Murota, *Discrete Convex Analysis*. Philadelphia, PA, USA: SIAM, 2003.

Jingjing Zhang received the B.Sc. degree from the Harbin Institute of Technology, Harbin, China, in 2010, the M.Sc. degree from the Beijing University of Posts and Telecommunications, Beijing, China, in 2013, and the Ph.D. degree from Télécom Paris, Paris, France, in 2017, all in electrical engineering.

She was a Post-Doctoral Researcher with the King's College London, London, U.K., where she was working with Prof. O. Simeone. Her research interests include distributed computing/learning, caching, feedback, topological networks, and information theory.

Dr. Zhang was a recipient of the Prix de These for her Ph.D. thesis.

Osvaldo Simeone (Fellow, IEEE) received the M.Sc. degree (Hons.) and the Ph.D. degree from the Politecnico di Milano, Milan, Italy, in 2001 and 2005, respectively.

From 2006 to 2017, he was a Faculty Member of the Department of Electrical and Computer Engineering, New Jersey Institute of Technology, Newark, NJ, USA. He is currently a Professor of information engineering with the Department of Engineering, King's College London, London, U.K., where he is the Director of the King's Communications, Learning and Information Processing Laboratory. He received the Consolidator Grant from the European Research Council (ERC) in 2016. His research has been supported by the U.S. National Science Foundation (NSF), the ERC, the Vienna Science and Technology Fund, and by a number of industrial collaborations.

Dr. Simeone currently serves on the Editorial Board of the IEEE SIGNAL PROCESSING. He was a Distinguished Lecturer of the IEEE Information Theory Society in 2017 and 2018. He was a co-recipient of the Best Paper Awards at the IEEE SPAWC 2007 and the IEEE WRECOM 2007, the 2015 IEEE Communication Society Best Tutorial Paper Award, the 2017 JCN Best Paper Award, the 2018 IEEE Signal Processing Best Paper Award, and the 2019 IEEE Communication Society Best Tutorial Paper Award.