

GRADIENT CODING WITH CLUSTERING AND MULTI-MESSAGE COMMUNICATION

Emre Ozfatura[†], Deniz Gündüz[†] and Sennur Ulukus[‡]

[†]Information Processing and Communications Lab,
Dept. of Electrical and Electronic Engineering, Imperial College London

[‡]Department of Electrical and Computer Engineering
Institute for Systems Research, University of Maryland, College Park
{m.ozfatura,d.gunduz}@imperial.ac.uk, ulukus@umd.edu

ABSTRACT

Gradient descent (GD) methods are commonly employed in machine learning problems to optimize the parameters of the model in an iterative fashion. For problems with massive datasets, computations are distributed to many parallel computing servers (i.e., *workers*) to speed up GD iterations. While distributed computing can increase the computation speed significantly, the per-iteration completion time is limited by the slowest *straggling* workers. Coded distributed computing can mitigate straggling workers by introducing redundant computations; however, existing coded computing schemes are mainly designed against persistent stragglers, and partial computations at straggling workers are discarded, leading to wasted computational capacity. In this paper, we propose a novel gradient coding (GC) scheme which allows multiple coded computations to be conveyed from each worker to the master per iteration. We numerically show that the proposed GC with multi-message communication (MMC) together with clustering provides significant improvements in the average completion time (of each iteration), with minimal or no increase in the communication load.

1. INTRODUCTION

In centralized synchronous distributed gradient descent (DGD), each worker computes a gradient estimate, called a *partial gradient*, using its own local dataset, which are then aggregated by the parameter server, i.e., the *master*, to update the model parameters [1]. In this implementation, the master waits until it receives the partial gradients from all the workers; hence, the overall computation speed is limited by the slowest, so-called *straggling* workers. Straggler tolerant DGD schemes have been introduced recently, where redundant computations are assigned to workers, either together with coded dataset [2–8], or combined with coded local computations [9–11], or by simply using backup computations [12–14].

Most of these works target *persistent stragglers*; that is, the straggling workers either cannot fulfill any computations, or their computations are discarded; hence, the computation capacity of the network is underutilized. However, in practice, we mostly encounter *non-persistent stragglers*, which are not capable of completing all their assigned tasks by a deadline, but may still compute a significant portion of them. Partial computations carried out by non-persistent stragglers can also be utilized by the master to reduce the per-iteration completion time [12, 13, 15–17].

This work has been funded by the European Research Council (ERC) through project BEACON (No. 725731) and by H2020-MSCA-ITN-2015 project SCAVENGE under grant number 675891.

In this paper, we introduce a *multi-message communication (MMC)* strategy for centralized synchronous DGD to better utilize the computational capacity and to reduce the computation time. In particular, we extend the gradient coding (GC) approach of [9, 11] to MMC in order to seek a balance between the communication load and the computation time. MMC has been previously considered with different straggler avoidance strategies [12, 13, 15–19]. Nevertheless, most of these studies are either limited to linear regression type problems, where the gradient computation boils down to matrix-matrix or matrix-vector multiplication [15–19], or simply assign the same computation to multiple workers [12, 13], which leads to high communication load. Hence, the aim of the proposed design is to reduce the computation time with MMC while avoiding high communication load, unlike the aforementioned uncoded repetition schemes, and to ensure that the design is applicable to a wide range of machine learning problems employing DGD. In this paper, we achieve these objectives by introducing a novel GC scheme, called GC-MMC, in which multiple coded partial gradients are generated and transmitted by each worker as computations go on. This allows the master to exploit the partial computations of straggling workers, which is particularly useful when the workers have similar speeds (i.e., non-persistent stragglers), and also allows early termination before the straggling threshold is met, even in the presence of persistent stragglers. The latter is beneficial when a few workers are significantly faster than the others. Assuming a shifted exponential computation time statistics, we numerically show the benefits of the proposed GC-MMC scheme compared to other coded and uncoded computation schemes with or without MMC.

2. GRADIENT CODING (GC)

Consider DGD implemented with one master and K workers. The dataset \mathbf{X} and the corresponding labels \mathbf{y} are divided into K non-overlapping equal-size mini-batches B_1, \dots, B_K . We use $g_k^{(t)}$ to denote the partial gradient for parameter vector θ_t evaluated with mini-batch B_k , i.e.,

$$g_k^{(t)} = \frac{1}{|B_k|} \sum_{(\mathbf{x}_i, y_i) \in B_k} \nabla l((\mathbf{x}_i, y_i), \theta_t). \quad (1)$$

The “full gradient” computed over the whole dataset, can be obtained by $\frac{1}{K} \sum_{i=1}^K g_i^{(t)}$. GC is a distributed computing framework to compute the full gradient across K workers while tolerating straggling workers thanks to redundant computations assigned to each worker [9] Let \mathcal{B}_k denote the set of mini-batches assigned to the k th worker,

i.e., if $B_i \in \mathcal{B}_k$, then $g_i^{(t)}$ is computed by the i th worker. The *computation load*, r , is defined as the number of mini-batches assigned to each worker, i.e., $|\mathcal{B}_k| = r, \forall k \in [K]$. After computing r partial gradients, one for each mini-batch available locally, each worker sends a linear combination of those computations, $c_k^{(t)} \triangleq \mathcal{L}_k(g_i^{(t)} : B_i \in \mathcal{B}_k)$, called the *coded partial gradient*. The master waits until it receives sufficiently many coded partial gradients to obtain the full gradient. It is shown in [9] that, for any set of non-straggler workers $\mathcal{W} \subseteq [K]$ with $|\mathcal{W}| = K - r + 1$, there exists a set of coefficients $\mathcal{A}_{\mathcal{W}} = \{a_k^{(t)} : k \in \mathcal{W}\}$ such that $\sum_{k \in \mathcal{W}} a_k^{(t)} c_k^{(t)} = 1/K \sum_{k=1}^K g_k^{(t)}$. Hence, GC can tolerate up to $r - 1$ persistent stragglers at each iteration t . Next, we introduce two strategies, namely *MMC* and *clustering* to reduce the per-iteration completion time with GC.

3. GRADIENT CODING WITH MULTI-MESSAGE COMMUNICATION (GC-MMC)

In the original GC scheme of [9], the number of messages transmitted to the master per-iteration per-worker is limited to one. Due to the synchronized model update, the workers that complete their computations stay idle until they receive the updated parameter vector to start the next iteration. Besides, since the workers send a single message only after completing all the assigned computations, computations carried out by the slow/straggling workers are discarded. To prevent the under-utilization of the computation resources, we will allow each worker to send coded partial gradients to the master; that is, at each iteration each worker sends multiple coded partial gradients instead of sending a single computation result. In the scope of this paper, we will present two different approaches to design coded partial gradients, namely *correlated code design* and *uncorrelated code design*, which are explained next.

3.1. Correlated code design

In GC, the number of partial gradients linearly combined to form the transmitted message from a worker is equal to the computation load r . In MMC, we allow each worker to compute and transmit multiple coded partial gradients, each of which will be generated by combining $m \leq r$ gradient computations. We will refer to m as the *order* of the corresponding partial gradient. In particular, each worker will be able to send up to $l = r - m + 1$ different messages, each of order m ; that is, each of the coded partial gradients will be a linear combination of the m most recently computed partial gradients.

As an example, let $r = 3$ and $m = 2$, and consider the worker with mini-batches B_1, B_2 and B_3 , which will compute partial gradients g_1, g_2 and g_3 in this order¹. After computing g_1 and g_2 the worker will send a linear combination of these two partial gradients as a coded message to the master, and after computing g_3 , it will send a linear combination of partial gradients g_2 and g_3 .

In general, the proposed scheme consists of two steps: coded message construction and message assignment. In the coded message construction step, structure of the coded messages are designed according to the order m by simply setting $r = m$ in the original GC scheme. Then, in the coded message assignment step, constructed messages are assigned to the workers based on the assigned partial gradients. We present the following example to clarify these steps.

Example 1: Let $K = 6, r = 3, m = 2$, and consider the assignment matrix \mathbf{M} , whose i th row indicates the mini-batches assigned to the i th worker; that is $\mathbf{M}(i, j) = 1$ means that partial gradient g_j

¹In the rest of the paper, to simplify the notation we will drop the time index from the gradients when we focus on a single iteration of the algorithm.

will be computed by the i th worker. In GC with $K = 6$ and $r = 3$, we have the following assignment matrix.

$$\mathbf{M} = \begin{bmatrix} 1 & 1 & \color{red}1 & 0 & 0 & 0 \\ 0 & 1 & 1 & \color{red}1 & 0 & 0 \\ 0 & 0 & 1 & 1 & \color{red}1 & 0 \\ 0 & 0 & 0 & 1 & 1 & \color{red}1 \\ \color{red}1 & 0 & 0 & 0 & 1 & 1 \\ 1 & \color{red}1 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (2)$$

When $m = 2$, coded gradients are obtained according to the assignment matrix $\tilde{\mathbf{M}}$, which is obtained by removing the last $r - m$ of the 1s in each row (shown in red above). When the assignment matrix $\tilde{\mathbf{M}}$ is used to design GC, a total of $K = 6$ coded partial gradients, each of order two, are constructed; and the full gradient can be obtained from any $K - m - 1 = 5$ coded partial gradients. Let c_1, \dots, c_6 denote the corresponding coded partial gradients. We remark that c_1 is a linear combination of g_1 and g_2 , while c_2 is a linear combination of g_2 and g_3 . Since g_1, g_2 and g_3 are assigned to the first worker, it can send both coded messages c_1 and c_2 . We call this approach *correlated code design*, since the same coded partial gradient can be computed and sent by more than one worker, e.g., in Example 1, c_2 can be sent by both worker 1 and worker 2. In Example 1, the original GC algorithm needs to receive computations from at least four workers in order to complete an iteration; whereas the proposed scheme can complete an iteration with results from only three workers. For instance, when workers 1, 2 and 4 each send two coded partial gradients, the master will obtain c_1, c_2, c_3, c_4, c_5 , and recover the full gradient. In the next section, we will analyze the uncorrelated coded design approach, where each coded gradient is assigned to exactly one worker.

3.2. Uncorrelated code design

As explained in [11], GC can be interpreted as a polynomial interpolation problem. In this model, the gradient assignment matrix is called the *mask matrix*, and the *support* of the i th row $\mathbf{M}(i, \cdot)$, denoted by $\text{supp}(\mathbf{M}(i, \cdot))$, gives the index of the partial gradients that are assigned to the i th worker. For a given mask matrix \mathbf{M} , GC is equivalent to interpolating a polynomial with *degree* h , where $h = N - \max_i |\text{supp}(\mathbf{M}(i, \cdot))|$; in other words, h is equal to the number of zeros in the most sparse column of \mathbf{M} . In a broad sense, each partial gradient g_k is embedded into a polynomial f_k , and each worker evaluates the polynomials f_1, \dots, f_K at preassigned points, and sends their sum to the master. The key design issue is that, if gradient g_k is not assigned to the j th user, then the value of polynomial f_k should be zero at the point assigned to that user. Hence, if partial gradient g_k is missing at h workers, which is the number of zeros in the j th column of the mask matrix \mathbf{M} , then the corresponding polynomial f_k should have h roots, which implies that f_k , and hence, the polynomial $\sum_k f_k$, both have degree h .

Remember that a polynomial with degree h can be interpolated from its values at $h + 1$ different points. Hence, as in the coded gradient approach in [11] non-straggling threshold $K_{th} = K - r + 1$ can be achieved utilizing a Reed-Solomon code structure. Please refer to [11] for the details².

In **Example 1**, the mask matrix \mathbf{M} has exactly 3 zeros in each column; therefore, at least 4 coded partial gradients must be received to recover the full gradient. As in GC, the main drawback of the scheme in [11] is that the workers send coded partial gradients only after computing all the assigned partial gradients. We will design a GC scheme with MMC, in which each coded partial gradient is

²Implementation of the encoding and decoding procedures also effect the completion time; however, this is ignored in the scope of this paper.

sent by at most one worker. If each worker is allowed to send l messages per iteration, we consider a setting with Kl workers, introducing $K(l - 1)$ “virtual” workers. Then, we design a GC scheme for the mask matrix of Kl users. We explain the proposed MMC strategy and the concept of virtual workers on Example 1.

Example 1 (continued): Since each worker can send $l = 2$ coded partial gradients, we create one virtual worker corresponding to each real worker. Partial gradient computations assigned to a virtual worker should be a subset of those assigned to the corresponding real worker. Since $r = 3$ partial gradients are assigned to each worker, we assign the first two partial gradients of each real worker to the corresponding virtual worker. The corresponding mask matrix for all the workers (including the virtual workers) is given below:

$$\mathbf{M} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad (3)$$

where the rows in red correspond to the real workers, while the rows in blue to the virtual workers. One can easily observe that in each column there are exactly 7 zeros; and hence, the master needs at least 8 coded gradients to recover the full gradient.

We remind that, in the original GC scheme, the master waits for 4 workers to recover the full gradient. However, in the proposed scheme each worker can send a message with a coded partial gradient as a virtual worker, after computing only two partial gradients, and the full gradient can be recovered from any 8 coded partial gradients, including those from the virtual workers. Consider, for example, three of the workers are non-straggler, and each of them sends 2 coded partial gradients, while two workers are non-persistent stragglers, each of them sends one coded gradient, and the last worker is a persistent straggler. In this case the full gradient can be obtained by the proposed approach but not with the original GC scheme. Hence, the proposed approach improves the per-iteration completion time.

Uncorrelated code design for GC with MMC is defined by the order vector $[m_0, \dots, m_l]$, where m_0 is the order of the coded partial gradient sent by the real worker, and m_i denotes the order of the coded partial gradient sent by the i th virtual worker. For a particular MMC strategy with order vector $\mathbf{m} = [m_0, \dots, m_l]$, the number of zeros in any column of the mask matrix \mathbf{M} is given by $Kl - (\sum_{i=0}^l m_i)$; and thus, $Kl - (\sum_{i=0}^l m_i) + 1$ coded partial gradients are required to recover the full gradient.

We emphasize here that the use of coded partial gradients with lower orders increases the number of messages required for recovering the full gradient; however, they can be obtained faster; and hence, we can exploit the computations carried out by non-persistent stragglers. The focus of the current paper is on identifying potential designs for GC with MMC; thus, the optimization of the partial gradient orders is left as future work.

Another important issue regarding MMC is the *communication load*, which denotes the average number of messages received by the master at each iteration. The communication load increases with the number of virtual workers; therefore, the optimal MMC strategy depends critically on the communication architecture of the network and the protocol used to transmit messages from the workers to the master as well as the computation speeds of the workers.

3.3. Clustering

Next we introduce clustering, which can further speed up the computation time. Let the workers be divided into P equal-size disjoint clusters, where the set of workers in cluster p is denoted by $\mathcal{K}_p \subset [K]$, $p \in [P]$. Mini-batches are also divided into equal-size disjoint subsets, and the set of mini-batches assigned to p th cluster is denoted by $\mathcal{B}_{\mathcal{K}_p}$. In the clustering approach, the workers in the p th cluster are responsible for computing

$$\frac{1}{|\mathcal{B}_{\mathcal{K}_p}|} \sum_{k: \mathcal{B}_k \in \mathcal{B}_{\mathcal{K}_p}} g_k^{(t)}, \quad (4)$$

and the GC scheme is applied to each cluster independently.

At this point, we remark that the fractional repetition scheme in [9] is a special case of the clustering approach where the size of a cluster is equal to r . Consider, $K = 40$, $r = 10$ and $P = 4$, where the workers are divided into $P = 4$ clusters, while the mini-batches are divided into 4 subsets, and each cluster is responsible for a different subset. In the fractional repetition scheme the master waits until at least one worker from each cluster completes and sends its partial gradient. One can observe that if there are $K - r + 1 = 31$ workers that have completed and sent their computations to the master, there must be at least one worker from each cluster among these workers; hence, the non-straggling threshold is $K - r + 1$. However, the non-straggling threshold represents a worst case scenario. Notice that, even 4 workers, each from a different cluster can be sufficient to obtain the full gradient. On the other hand, the other GC scheme proposed in [9], called the cyclic repetition scheme, always has to wait until receiving coded messages from $K - r + 1$ worker. Hence, we can state that although various GC scheme can achieve the same optimal non-straggling threshold, their average performance may differ.

While the fractional repetition scheme requires K to be an integer multiple of r , the clustering approach outlined above can be applied to any (K, r) pair. When GC is applied with clustering, it is possible to tolerate $r - 1$ stragglers in each cluster; thus, for a particular straggler realization, if the full gradient can be obtained with GC (without clustering) then it can also be obtained with clustering, while the converse is not true. Based on this observation, we can claim that with clustering, straggling threshold remains the same, but the average computation time can be reduced. However, when MMC is allowed, clustering may also be disadvantageous. On one hand, more straggling workers can be tolerated on average, on the other hand GC is applied to each cluster independently; hence, a coded partial gradient from a particular cluster cannot be utilized for another cluster. Consequently, the optimal clustering strategy with MMC depends on the computation statistics of the workers.

3.4. Hybrid implementation

The optimal DGD strategy depends critically on the computation time statistics of the workers. In particular, when the computation speeds of the workers are similar, MMC is expected to have a better performance as it can exploit all the computations carried out across the workers; however, when one of the workers is much faster compared to the others, fractional repetition can be preferred against MMC. To illustrate this trade-off, consider the case $K = 5$ and $r = 5$. With the fractional repetition scheme, the master waits for the fastest worker to finish all assigned computations; however, with GC-MMC with order vector $\mathbf{m} = [5, 3]$, the master waits for exactly 3 coded messages. Accordingly, we can propose a hybrid scheme, in which the workers initially behave as dictated by the GC-MMC scheme,

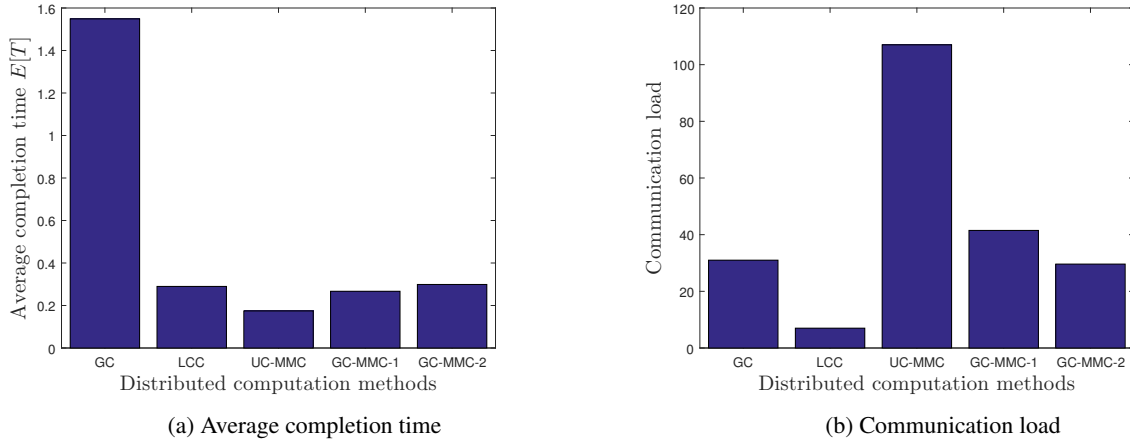


Fig. 1: Average per-iteration completion time and communication load under the shifted-exponential model.

but if a worker is fast enough to complete all its computations, then it switches to fractional repetition scheme, and sends the average gradient instead of a coded partial gradient.

4. NUMERICAL RESULTS AND DISCUSSIONS

4.1. Simulation Results

For the simulations, we consider the setup in [4], where there are $K = 40$ workers and the dataset is divided into $K = 40$ mini-batches. We set $r = 10$; that is, 10 partial gradient computations, each corresponding to a different mini-batch, are assigned to each worker. For the computation time statistics of the workers, we adopt the commonly used shifted exponential model [2, 18], which assumes that the probability of completing s computations (partial gradients) at any worker by time t is given by

$$F_s(t) \triangleq \begin{cases} 1 - e^{-\mu(\frac{t}{s} - \alpha)}, & \text{if } t \geq s\alpha, \\ 0, & \text{else.} \end{cases} \quad (5)$$

In our simulations, we used parameters $\mu = 10$ and $\alpha = 0.01$, which implies that, on average, $\mu = 10$ partial gradients can be computed by a worker in unit time (e.g., one second), and the computation time of a single partial gradient cannot be less than $\alpha = 0.01$.

For the performance analysis we design two different GC schemes with MMC. In the first scheme, called GC-MMC-1, four equal-size clusters are formed and correlated code design is employed in each cluster with order $m = 6$, so that each worker can send up to 5 coded partial gradients. In the second scheme, called GC-MMC-2, again four equal-size clusters are formed, but uncorrelated code design is employed in each cluster with order vector $\mathbf{m} = [6, 8, 10]$. We compare the performance of these schemes with GC, uncoded computation with MMC (UC-MMC) [19] and Langrange coded computation (LCC) [4] schemes, in terms of the average per-iteration completion time and the average communication load, i.e., the average number of messages sent to the master at each iteration.

The results are presented in Figure 1. We observe that the GC-MMC schemes can perform as well as LCC in terms of the average per-iteration completion time, although LCC introduces significant coding and decoding complexity. Furthermore, the GC-MMC

schemes can even outperform the LCC if the number of messages are increased. Indeed, notice that the UC-MMC scheme is actually a special case of the GC-MMC scheme with correlated code design with order $m = 1$. However, this improvement in the average per-iteration completion time performance comes at the expense of an increase in the communication load. The key advantage of the GC-MMC schemes is the more flexible trade-off they provide between the completion time and the communication load compared to UC-MMC and GC. To this end, an interesting observation is that the GC-MMC-2 achieves the same performance as LCC, a coded computation scheme, while still sending only a single message per-worker per-iteration on average.

We want to highlight that LCC, its multi-message variations [19] might be more efficient compared to GC-MMC schemes in terms of the completion time as well as the communication load statistics. However, LCC and other coded computation schemes have the following drawbacks in real applications i) they require encoding the whole dataset before computation; ii) they are limited to problems where the gradient computation can be reduced to matrix-vector or matrix-matrix multiplication (mainly linear regression problems); iii) other techniques that can increase the communication efficiency, such as gradient quantization [20, 21] cannot be incorporated with them. Therefore, we believe that speeding up coded computation schemes that compute uncoded data points is an important research direction for practical applications.

5. CONCLUSIONS

We introduced a novel GC scheme with MMC and clustering, allowing workers to send coded partial gradients before completing all of their assigned computations. This can potentially lead to early termination of an iteration as it can benefit from computations carried out at all the workers, including the relatively slower stragglers. The improvement is achieved at the expense of an increased communication load, and the proposed scheme seeks a balance between the communication load and the average per-iteration completion time, thanks to its flexible coding and clustering structure. We numerically show that, under certain completion time statistics, average completion time can be reduced dramatically with a very small or no increase in the communication load.

6. REFERENCES

- [1] J. Dean, G. S. Corrado, R. Monga, K. Chen, M. Devin, Q. V. Le, M. Z. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, and A. Y. Ng, "Large scale distributed deep networks," in *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, USA, 2012, NIPS'12, pp. 1223–1231, Curran Associates Inc.
- [2] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," *IEEE Transactions on Information Theory*, vol. 64, no. 3, pp. 1514–1529, March 2018.
- [3] R. K. Maity, A. S. Rawat, and A. Mazumdar, "Robust gradient descent via moment encoding with ldpc codes," *SysML Conference*, 2018.
- [4] S. Li, S. M. M. Kalan, Q. Yu, M. Soltanolkotabi, and A. S. Avestimehr, "Polynomially coded regression: Optimal straggler mitigation via data encoding," *CoRR*, vol. abs/1805.09934, 2018.
- [5] M. Fahim, H. Jeong, F. Haddadpour, S. Dutta, V. Cadambe, and P. Grover, "On the optimal recovery threshold of coded matrix multiplication," in *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Oct 2017, pp. 1264–1270.
- [6] Q. Yu, M. Maddah-Ali, and S. Avestimehr, "Polynomial codes: an optimal design for high-dimensional coded matrix multiplication," in *Advances in Neural Information Processing Systems 30*, pp. 4403–4413. Curran Associates, Inc., 2017.
- [7] H. Park, K. Lee, J. Sohn, C. Suh, and J. Moon, "Hierarchical coding for distributed computing," in *2018 IEEE International Symposium on Information Theory (ISIT)*, June 2018, pp. 1630–1634.
- [8] S. Kiani, N. Ferdinand, and S. C. Draper, "Exploitation of stragglers in coded computation," in *2018 IEEE International Symposium on Information Theory (ISIT)*, June 2018, pp. 1988–1992.
- [9] R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *Proceedings of the 34th International Conference on Machine Learning*, International Convention Centre, Sydney, Australia, 06–11 Aug 2017, vol. 70 of *Proceedings of Machine Learning Research*, pp. 3368–3376, PMLR.
- [10] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," in *Proceedings of the 35th International Conference on Machine Learning*, Stockholmssmassan, Stockholm Sweden, 10–15 Jul 2018, vol. 80 of *Proceedings of Machine Learning Research*, pp. 5610–5619, PMLR.
- [11] W. Halbawi, N. Azizan, F. Salehi, and B. Hassibi, "Improving distributed gradient descent using reed-solomon codes," in *2018 IEEE International Symposium on Information Theory (ISIT)*, June 2018, pp. 2027–2031.
- [12] S. Li, S. M. M. Kalan, A. S. Avestimehr, and M. Soltanolkotabi, "Near-optimal straggler mitigation for distributed gradient methods," in *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, May 2018, pp. 857–866.
- [13] M. Mohammadi Amiri and D. Gunduz, "Computation scheduling for distributed machine learning with straggling workers," *CoRR*, vol. abs/1810.09992, 2018.
- [14] J. Chen, R. Monga, S. Bengio, and R. Józefowicz, "Revisiting distributed synchronous SGD," *CoRR*, vol. abs/1604.00981, 2016.
- [15] A. Mallick, M. Chaudhari, and G. Joshi, "Rateless codes for near-perfect load balancing in distributed matrix-vector multiplication," *CoRR*, vol. abs/1804.10331, 2018.
- [16] E. Ozfatura, S. Ulukus, and D. Gunduz, "Distributed gradient descent with coded partial gradient computations," *CoRR*, vol. 1811.09271, 2018.
- [17] A. B. Das, L. Tang, and A. Ramamoorthy, "C3les: Codes for coded computation that leverage stragglers," in *2018 IEEE Information Theory Workshop (ITW)*, Nov 2018, pp. 1–5.
- [18] N. Ferdinand and S. C. Draper, "Hierarchical coded computation," in *2018 IEEE International Symposium on Information Theory (ISIT)*, June 2018, pp. 1620–1624.
- [19] E. Ozfatura, D. Gunduz, and S. Ulukus, "Speeding up distributed gradient descent by utilizing non-persistent stragglers," *CoRR*, vol. abs/1808.02240, 2018.
- [20] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "Qsgd: Communication-efficient sgd via gradient quantization and encoding," in *Advances in Neural Information Processing Systems 30*, pp. 1709–1720. Curran Associates, Inc., 2017.
- [21] P. Jiang and G. Agrawal, "A linear speedup analysis of distributed deep learning with sparse and quantized communication," in *Advances in Neural Information Processing Systems 31*, pp. 2529–2540. Curran Associates, Inc., 2018.