

Hierarchical Coding for Distributed Computing

Hyegyong Park, Kangwook Lee, Jy-yong Sohn, Changho Suh and Jaekyun Moon
 School of Electrical Engineering, Korea Advanced Institute of Science and Technology (KAIST)
 Email: {parkh, kw1jjang, jysohn1108, chsuh}@kaist.ac.kr, jmoon@kaist.edu

Abstract—Coding for distributed computing supports low-latency computation by relieving the burden of straggling workers. While most existing works assume a simple master-worker model, we consider a hierarchical computational structure consisting of groups of workers, motivated by the need to reflect the architectures of real-world distributed computing systems. In this work, we propose a *hierarchical coding scheme* for this model, as well as analyze its decoding cost and expected computation time. Specifically, we first provide upper and lower bounds on the expected computing time of the proposed scheme. We also show that our scheme enables efficient parallel decoding, thus reducing decoding costs by orders of magnitude over non-hierarchical schemes. When considering both decoding cost and computing time, the proposed hierarchical coding is shown to outperform existing schemes in many practical scenarios.

I. INTRODUCTION

Enabling large-scale computations for big data analytics, distributed computing systems have received significant attention in recent years [1]. The distributed computing system divides a computational task to a number of subtasks, each of which is allocated to a different worker. This helps reduce computing time by exploiting parallel computing options and thus enables handling of large-scale computing tasks.

In a distributed computing system, the “stragglers”, which refers to the computing nodes that slow down in some random fashion due to a variety of factors, may increase the total runtime of the computing system. To address this problem, the notion of *coded computation* is introduced in [2] where an (n, k) maximum distance separable (MDS) code is employed to speed up distributed matrix multiplications. The authors show that for linear computing tasks, one can design n distributed computing tasks such that *any* k out of n tasks suffice to complete the assigned task. Since then, coded computation has been applied to a wide variety of task scenarios such as matrix-matrix multiplication [3], [4], distributed gradient computation [5]–[8], convolution [9], Fourier transform [10] and matrix sparsification [11], [12].

While the idea of coded computation has been studied in various settings, existing works have not taken into account the underlying *hierarchical* nature of practical distributed systems [13]–[15]. In modern distributed computing systems, each group of workers is collocated in the same rack, which contains a Top of Rack (ToR) switch, and cross-rack communication is available only via these ToR switches. Surveys on real cloud computing systems show that cross-rack communication through the ToR switches is highly unstable due to the limited bandwidth, whereas intra-rack communication is faster and more reliable [14], [15]. A natural question is whether one




Fig. 1. Illustration of the hierarchical computing system

can devise a coded computation scheme that exploits such hierarchical structure.

A. Contribution

In this work, we first model a distributed computing system with a tree-like hierarchical structure illustrated in Fig. 1, which is inspired by the practical computing systems in [13]–[15]. The workers (denoted by “W”) are divided into groups, each of which has a submaster (denoted by “SM”). Each submaster sends the computational result of its group to the master (denoted by “M”). The suggested model can be viewed as a generalization of the existing non-hierarchical coded computation.

In this framework, we propose a hierarchical coding scheme which employs an $(n_1^{(i)}, k_1^{(i)})$ MDS code within group i and another (n_2, k_2) outer MDS code across the groups as depicted in Fig. 2. We also develop a parallel decoding algorithm which exploits the concatenated code structure and allows low complexity.

Moreover, we analyze the latency performance of our proposed solution. It turns out that the latency performance of our scheme cannot be analyzed via simple order statistics as in other existing schemes. Here we resort to find lower and upper bounds on the average latency performance: Our upper bound relies on concentration inequalities, and our lower bound is obtained via constructing and analyzing an auxiliary Markov chain.

B. Related Work

Previous works on coded computation have rarely considered the inherent hierarchical structure of most real-world systems. Whereas a very recent work [16] deals with the multi-rack computing system reflecting imbalance between intra- and cross-rack communications, it is based on the settings of the coded MapReduce architecture which do not include general linear computation tasks that we focus on in this work. Another distinction is that the analysis of [16] includes only

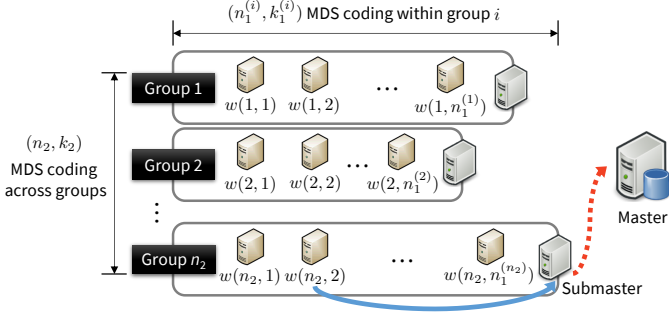


Fig. 2. Illustration of the proposed coding scheme applied to the hierarchical computing system. An $(n_1^{(i)}, k_1^{(i)})$ MDS code is employed within group i , and an (n_2, k_2) MDS code is applied across the groups. $w(i, j)$ denotes worker j in group i .

the cross-rack redundancy whereas our analysis considers both intra- and cross-group coding.

C. Notations

We use boldface uppercase letters for matrices and boldface lowercase letters for vectors. The transpose of a matrix \mathbf{A} is denoted by \mathbf{A}^T . For a matrix \mathbf{A} satisfying $\mathbf{A}^T = [\mathbf{A}_1^T \ \mathbf{A}_2^T]$, we write $\mathbf{A} = [\mathbf{A}_1; \mathbf{A}_2]$. For a positive integer n , the set $\{1, 2, \dots, n\}$ is denoted by $[n]$. The j^{th} worker in group i is represented by $w(i, j)$ for $i \in [n_2]$ and $j \in [n_1^{(i)}]$. The symbol $\lfloor r \rfloor$ indicates the largest integer less than or equal to a real number r .

II. HIERARCHICAL CODED COMPUTATION

A. Proposed Coding Scheme

Consider a matrix-vector multiplication task, i.e., computing $\mathbf{A}\mathbf{x}$ for a matrix $\mathbf{A} \in \mathbb{R}^{m \times d}$ and a vector $\mathbf{x} \in \mathbb{R}^{d \times 1}$. The input matrix \mathbf{A} is split into k_2 submatrices as $\mathbf{A} = [\mathbf{A}_1; \mathbf{A}_2; \dots; \mathbf{A}_{k_2}]$, where $\mathbf{A}_i \in \mathbb{R}^{\frac{m}{k_2} \times d}$ for $i \in [k_2]$. Here we assume that m is divisible by k_2 for simplicity. Then, we apply an (n_2, k_2) MDS code to set $\{\mathbf{A}_i\}_{i \in [k_2]}$ in obtaining $\{\tilde{\mathbf{A}}_i\}_{i \in [n_2]}$. Then, each coded matrix $\tilde{\mathbf{A}}_i$ is further divided into $k_1^{(i)}$ submatrices as $\tilde{\mathbf{A}}_i = [\tilde{\mathbf{A}}_{i,1}; \tilde{\mathbf{A}}_{i,2}; \dots; \tilde{\mathbf{A}}_{i,k_1^{(i)}}]$ where $\tilde{\mathbf{A}}_{i,j} \in \mathbb{R}^{\frac{m}{k_1^{(i)}k_2} \times d}$ for $j \in [k_1^{(i)}]$ and m divisible by $k_1^{(i)}k_2$. Afterwards, for each $i \in [n_2]$, we apply an $(n_1^{(i)}, k_1^{(i)})$ MDS code to set $\{\tilde{\mathbf{A}}_{i,j}\}_{j \in [k_1^{(i)}]}$ to obtain $\{\hat{\mathbf{A}}_{i,j}\}_{j \in [n_1^{(i)}]}$. Then, for each $i \in [n_2]$ and $j \in [n_1^{(i)}]$, worker $w(i, j)$ computes $\hat{\mathbf{A}}_{i,j}\mathbf{x}$. Fig. 2 illustrates the proposed coding scheme for the hierarchical computing system with a different number of workers in each group. In the case of $n_1^{(i)} = n_1$ and $k_1^{(i)} = k_1$ for all $i \in [n_2]$, we will refer this coding scheme as $(n_1, k_1) \times (n_2, k_2)$ coded computation.

We present our code in Fig. 3 via a toy example. In this example, $(n_1, k_1) = (n_2, k_2) = (3, 2)$. That is, the input matrix $\mathbf{A} = [\mathbf{A}_1; \mathbf{A}_2]$ is encoded via $(n_2, k_2) = (3, 2)$ MDS code, yielding $[\tilde{\mathbf{A}}_1; \tilde{\mathbf{A}}_2; \tilde{\mathbf{A}}_3 := \tilde{\mathbf{A}}_1 + \tilde{\mathbf{A}}_2]$. Afterwards, the matrix $\tilde{\mathbf{A}}_i = [\tilde{\mathbf{A}}_{i,1}; \tilde{\mathbf{A}}_{i,2}]$ is encoded via an $(n_1, k_1) = (3, 2)$ MDS code, producing $[\hat{\mathbf{A}}_{i,1}; \hat{\mathbf{A}}_{i,2}; \hat{\mathbf{A}}_{i,3} := \hat{\mathbf{A}}_{i,1} + \hat{\mathbf{A}}_{i,2}]$. For notational simplicity, we define $\hat{\mathbf{A}}_3 = \hat{\mathbf{A}}_1 + \hat{\mathbf{A}}_2$ and $\hat{\mathbf{A}}_{i,3} = \hat{\mathbf{A}}_{i,1} + \hat{\mathbf{A}}_{i,2}$.



Fig. 3. Allocation of the computational task to workers in a $(3, 2) \times (3, 2)$ coded computation

For $i, j \in \{1, 2, 3\}$, worker $w(i, j)$ computes $\hat{\mathbf{A}}_{i,j}\mathbf{x}$. Note that group i is assigned a subtask with respect to $\tilde{\mathbf{A}}_i$.

We now describe the decoding algorithm for our proposed coding scheme. When a worker completes its task, it sends the result to its submaster. With the aid of the (n_1, k_1) MDS code, submaster i (in group i) can compute $\tilde{\mathbf{A}}_i\mathbf{x}$ as soon as the task results from any k_1 workers within group i are collected. Once $\tilde{\mathbf{A}}_i\mathbf{x}$ is computed, it is sent to the master. The master can obtain $\mathbf{A}\mathbf{x}$ by retrieving $\tilde{\mathbf{A}}_i\mathbf{x}$ from any k_2 submasters. For each worker, we define *completion time* as the sum of the runtime of the worker and the time required for delivering its computation result to the submaster. For each group, we further define *intra-group latency* as the time for completing its assigned subtask. The total computation time is defined as the time from when the workers start to run until the master completes computing $\mathbf{A}\mathbf{x}$. The proposed computation framework can be applied to practical multi-rack systems where the input data \mathbf{A} is coded and distributed into n_2 racks; the i^{th} rack contains $\tilde{\mathbf{A}}_i$. For instance, in the Facebook's warehouse cluster, data is encoded with a $(14, 10)$ MDS code, and then the 14 encoded chunks are stored across different racks [17]. Once \mathbf{x} is given from the master, the i^{th} rack can compute $\tilde{\mathbf{A}}_i\mathbf{x}$ using the coded data $\tilde{\mathbf{A}}_i$ that it contains.

B. Application: Matrix-Matrix Multiplications

Our scheme can be also applied to matrix-matrix multiplications. More specifically, consider computing $\mathbf{A}^T\mathbf{B}$ for given matrices \mathbf{A} and $\mathbf{B} = [\mathbf{b}_1 \ \mathbf{b}_2 \ \dots \ \mathbf{b}_{k_2}]$. After applying an (n_2, k_2) MDS code to \mathbf{B} , we have $\tilde{\mathbf{B}} = [\tilde{\mathbf{b}}_1 \ \tilde{\mathbf{b}}_2 \ \dots \ \tilde{\mathbf{b}}_{n_2}]$. Moreover, group i divides \mathbf{A} into $k_1^{(i)}$ equal-sized submatrices as $\mathbf{A} = [\mathbf{A}_{i,1} \ \mathbf{A}_{i,2} \ \dots \ \mathbf{A}_{i,k_1^{(i)}}]$, and we apply an $(n_1^{(i)}, k_1^{(i)})$ MDS code, resulting in $\tilde{\mathbf{A}}_i = [\tilde{\mathbf{A}}_{i,1} \ \tilde{\mathbf{A}}_{i,2} \ \dots \ \tilde{\mathbf{A}}_{i,n_1^{(i)}}]$. The computation $\tilde{\mathbf{A}}_{i,j}^T\tilde{\mathbf{b}}_i$ is assigned to worker $w(i, j)$. Using an $(n_1^{(i)}, k_1^{(i)})$ MDS code, submaster i can compute $\mathbf{A}^T\tilde{\mathbf{b}}_i$ when any $k_1^{(i)}$ workers within its group delivered their computation results. The master can calculate $\mathbf{A}^T\mathbf{B}$ by gathering $\mathbf{A}^T\tilde{\mathbf{b}}_i$ results from any k_2 submasters, using the (n_2, k_2) MDS code. Under the homogeneous setting of $n_1^{(i)} = n_1$ and $k_1^{(i)} = k_1$ for all $i \in [n_2]$, the encoding algorithm of the proposed scheme reduces to that of the product coded scheme [3]. However, the suggested scheme with the homogeneous setting is shown to reduce the decoding cost compared to the product coded

to the master, i.e., when the Markov chain visits the states with $v = 2$ for the first time. We see that u increases by one when a worker completes its computation, and v increases by one when master receives the computation result from a group. The rightward transition (to increase u) rate is determined by the product of μ_1 and the number of remaining workers. The upward transition (to increase v) rate is the product of μ_2 and the number of groups that have not delivered their computation results to the master. The proposed lower bound \mathcal{L} can be easily computed from the first-step analysis [20] of the Markov chain produced by Lemma 1.

B. Upper Bound

We here provide two upper bounds on the expected total computation time. The first bound in the following lemma is applicable for all values of n_1 and k_1 .

Lemma 2: The expected total computation time of the $(n_1, k_1) \times (n_2, k_2)$ coded computation is upper bounded as $\mathbb{E}[T] \leq H_{n_1 n_2} / \mu_1 + (H_{n_2} - H_{n_2 - k_2}) / \mu_2$.

Proof: See [19] for the proof. ■

We now establish another upper bound using the following two steps. First we find an upper bound on the maximum intra-group latency among n_2 groups. Afterwards, adding this value to the expected latency of the group-master communication yields an upper bound on the expected total computation time. For given n_1 and k_1 , we use $\delta_1 > 0$ which satisfies $n_1 = (1 + \delta_1)k_1$. We now present the asymptotic upper bound as follows.

Theorem 2: For a fixed constant $\delta_1 > 0$, the expected total computation time of the $(n_1, k_1) \times (n_2, k_2)$ coded computing system is upper bounded as $\mathbb{E}[T] \leq \lceil \log(1 + \delta_1) / \delta_1 \rceil / \mu_1 + (H_{n_2} - H_{n_2 - k_2}) / \mu_2 + o(1)$ in the limit of k_1 .

Proof: Due to the space constraint, here we provide only a sketch of the proof; interested readers are referred to [19] for details. First, we focus on the intra-group latency. The expected latency of the k_1^{th} fastest worker out of n_1 is given by $(H_{n_1} - H_{n_1 - k_1}) / \mu_1$, which can be rewritten as $\lceil \log(1 + \delta_1) / \delta_1 \rceil / \mu_1$ for a fixed constant $\delta_1 > 0$ and a sufficiently large n_1 . Define $t_0 := \lceil \log(1 + \delta_1) / \delta_1 \rceil / \mu_1 + \alpha \sqrt{\log k_1 / k_1}$ for some constant $\alpha > 0$. Further, let E_i denote the event that group i has *not* finished its assigned subtasks by time t_0 . Then, one can show that $\Pr[E_i] \leq k_1^{-\frac{2\delta_1^2 \mu_1^2 \alpha^2}{1 + \delta_1}}$ by using the Hoeffding's inequality [21]. Let T_I be the time when all n_2 groups finish their assigned subtasks. Then, one can see that $\{T_I > t_0\} = \cup_{i=1}^{n_2} E_i$. Hence, using the union bound, we have

$$\Pr[T_I > t_0] \leq \sum_{i=1}^{n_2} \Pr[E_i] \leq n_2 k_1^{-\frac{2\delta_1^2 \mu_1^2 \alpha^2}{1 + \delta_1}} = o(k_1^{-1}), \quad (5)$$

where the last equality holds since α can be made arbitrarily large. Then the expected intra-group latency $\mathbb{E}[T_I]$ satisfies

$$\mathbb{E}[T_I] \leq \Pr[T_I \leq t_0] t_0 + \Pr[T_I > t_0] (H_{n_1 n_2} / \mu_1 + t_0) \quad (6)$$

$$= (1 - o(k_1^{-1})) t_0 + o(k_1^{-1}) (H_{n_1 n_2} / \mu_1 + t_0) \quad (7)$$

$$= \log \lceil (1 + \delta_1) / \delta_1 \rceil / \mu_1 + o(1), \quad (8)$$

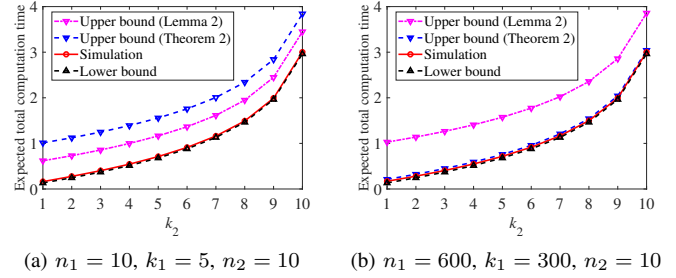


Fig. 6. The expected total computation time of the $(n_1, k_1) \times (n_2, k_2)$ coded computing with its bounds for varying k_2

where (6) is due to the fact that t_0 is the worst case latency for all events satisfying $T_I \leq t_0$, and $H_{n_1 n_2} / \mu_1 + t_0$ is an upper bound on $\mathbb{E}[T_I | T_I > t_0]$. From (8), we conclude that all n_2 groups embark on the group-master communication before time $\log \lceil (1 + \delta_1) / \delta_1 \rceil / \mu_1$, as k_1 grows large. Hence, adding the expected latency of the group-master communication $(H_{n_2} - H_{n_2 - k_2}) / \mu_2$ to (8) gives the claimed bound. ■

C. Evaluation of Bounds

Fig. 6 shows the behavior of the expected total computation time and its upper/lower bounds with varying k_2 . Here we consider two upper bounds proposed in Lemma 2 and Theorem 2. To see the impact of k_1 , the values of k_1 are fixed to 5 and 300 in Figs. 6a and 6b, respectively. The other code parameters are set to $n_1 = (1 + \delta_1)k_1$, $n_2 = 10$ for both figures, where δ_1 is fixed to 1. The rates of the completion time of the worker and group-master communication are set to $\mu_1 = 10$ and $\mu_2 = 1$. For a relatively small values of k_1 , the upper bound in Lemma 2 is tighter than that of Theorem 2. As can be seen in Fig. 6b, the asymptotic upper bound in Theorem 2 becomes tighter as k_1 grows, which concurs with Theorem 2. We also have numerically confirmed that the proposed lower bound is tight.

IV. DECODING COMPLEXITY

In this section, we compare decoding complexity of our hierarchical coding with the replication and non-hierarchical coding schemes including the $(n_1, k_1) \times (n_2, k_2)$ product code [3] and the (n, k) polynomial code [4]. For fair comparison, we set $n = n_1 n_2$ and $k = k_1 k_2$. We further assume that the decoding complexity of the (n, k) MDS code is $\mathcal{O}(k^\beta)$ for some $\beta > 1$.² In our framework, the n_2 intra-group codes can be decoded in parallel, followed by decoding of the cross-group code using the k_2 fastest results. Thus, the overall decoding procedure consists of 1) parallel decoding of (n_1, k_1) intra-group MDS codes and 2) decoding of the (n_2, k_2) cross-group MDS code, resulting in the total decoding cost of $\mathcal{O}(k_1^\beta + k_1 k_2^\beta)$.³ Similarly, one can show that the decoding cost of polynomial codes is $\mathcal{O}(k^\beta)$, and that of the product code is $\mathcal{O}(k_1 k_2^\beta + k_2 k_1^\beta)$. We note that the hierarchical code

²Note that this is the case for most practical decoding algorithms [22], [23]. Decoding with $\beta = 1$ requires a large field size [4].

³This is because the cross-group MDS code decoding requires a k_1 -fold symbol size compared to that of the intra-group MDS code decoding.

TABLE I
COMPARISONS OF VARIOUS CODING SCHEMES

Coding scheme	Computing time ($\mathbb{E}[T_{\text{comp}}]$)	Decoding cost (T_{dec})
Replication	$kH_k/(n\mu_2)$	0
Hierarchical code	$\mathbb{E}[T]$	$\mathcal{O}(k_1k_2^\beta + k_1^\beta)$
Product code [3]	$\frac{1}{\mu_2} \log \left(\frac{\sqrt{n/k} + \sqrt[4]{n/k}}{\sqrt{n/k}-1} \right)$	$\mathcal{O}(k_1k_2^\beta + k_2k_1^\beta)$
Polynomial code [4]	$(H_n - H_{n-k})/\mu_2$	$\mathcal{O}(k_1^\beta k_2^\beta)$

can have a substantial improvement, sometimes by an order of magnitude, in decoding complexity, compared to the product code. For instance, if $\beta = 2$ and $k_1 = k_2^2$, the decoding cost of hierarchical code becomes $\mathcal{O}(k_2^4)$ while that of the product code is $\mathcal{O}(k_2^5)$; if $k_1 = k_2^{1.5}$, the decoding costs are $\mathcal{O}(k_2^{3.5})$ and $\mathcal{O}(k_2^4)$, respectively. In general, if $k_1 = k_2^p$, one can show that the relative gain of the hierarchical codes in decoding cost monotonically increases as p increases, providing a guideline for efficient code designs. Table I summarizes the computing times and decoding costs of various coding schemes.

We now compare the expected total execution time defined as $T_{\text{exec}} := T_{\text{comp}} + \lambda T_{\text{dec}}$, where T_{comp} is the computing time, T_{dec} is the decoding cost, and $\lambda \geq 0$ is the relative weight of the decoding cost. We note that λ is a system-specific parameter that depends on 1) the relative CPU speed of the master compared to the workers and 2) dimension of the input data. Shown in Fig. 7 are the expected total execution times for parameters of $(n_1, k_1) = (800, 400)$, $(n_2, k_2) = (40, 20)$, $(\mu_1, \mu_2) = (10, 1)$ and $\beta = 2$.

Under the setting of parameters in Fig. 7, we observe that the hierarchical code strictly outperforms the product code for all values of λ . Although not shown, this observation is consistent for all practical values of μ_1 and μ_2 . In general, we observe that the optimal choice of coding scheme depends on the value of λ as follows:

- (moderate λ) when both T_{comp} and T_{dec} have to be minimized, the hierarchical code achieves the lowest T_{exec} by striking a balance between them;
- (low λ) when T_{dec} is negligible, the polynomial code achieves the lowest T_{exec} ; and
- (high λ) when T_{dec} dominates T_{exec} , the replication code is the best.

Note that the shaded area in Fig. 7 represents the additional achievable $(\lambda, \mathbb{E}[T_{\text{exec}}])$ region thanks to introducing the hierarchical code.

V. SUMMARY

We proposed a hierarchical coded computing scheme and provided bounds on its expected computing time. Specifically, we provided an upper bound using the concentration inequalities, and a lower bound based on the hitting time analysis of the auxiliary Markov chain. Further, we analyzed the decoding complexity of our scheme. In terms of total computation time combining the computing time and decoding cost, we showed that our scheme provides improvements by orders of magnitude compared to the non-hierarchical schemes in the practical regime where the decoding cost is significant.




Fig. 7. $\mathbb{E}[T_{\text{exec}}]$ of various coding schemes for parameters of $(n_1, k_1) = (800, 400)$, $(n_2, k_2) = (40, 20)$, $(\mu_1, \mu_2) = (10, 1)$ and $\beta = 2$

REFERENCES

- [1] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, A. Senior, P. Tucker, K. Yang, Q. V. Le *et al.*, “Large scale distributed deep networks,” in *Proc. NIPS*, 2012, pp. 1223–1231.
- [2] 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, 2018.
- [3] K. Lee, C. Suh, and K. Ramchandran, “High-dimensional coded matrix multiplication,” in *Proc. IEEE ISIT*, 2017, pp. 2418–2422.
- [4] Q. Yu, M. Maddah-Ali, and S. Avestimehr, “Polynomial codes: An optimal design for high-dimensional coded matrix multiplication,” in *Proc. NIPS*, 2017, pp. 4406–4416.
- [5] R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis, “Gradient coding: Avoiding stragglers in distributed learning,” in *Proc. ICML*, 2017, pp. 3368–3376.
- [6] W. Halbawi, N. Azizan-Ruhi, F. Salehi, and B. Hassibi, “Improving distributed gradient descent using Reed-Solomon codes,” *arXiv:1706.05436*, 2017.
- [7] N. Raviv, I. Tamo, R. Tandon, and A. G. Dimakis, “Gradient coding from cyclic MDS codes and expander graphs,” *arXiv:1707.03858*, 2017.
- [8] Z. Charles, D. Papailiopoulos, and J. Ellenberg, “Approximate gradient coding via sparse random graphs,” *arXiv:1711.06771*, 2017.
- [9] S. Dutta, V. Cadambe, and P. Grover, “Coded convolution for parallel and distributed computing within a deadline,” in *Proc. IEEE ISIT*, 2017, pp. 2403–2407.
- [10] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, “Coded Fourier transform,” *Proc. Allerton Conf.*, 2017.
- [11] S. Dutta, V. Cadambe, and P. Grover, “Short-Dot: Computing large linear transforms distributedly using coded short dot products,” in *Proc. NIPS*, 2016, pp. 2100–2108.
- [12] G. Suh, K. Lee, and C. Suh, “Matrix sparsification for coded matrix multiplication,” in *Proc. Allerton Conf.*, 2017.
- [13] J. Dean and S. Ghemawat, “MapReduce: Simplified data processing on large clusters,” *Commun. ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [14] F. Ahmad, S. T. Chakradhar, A. Raghunathan, and T. Vijaykumar, “ShuffleWatcher: Shuffle-aware scheduling in multi-tenant MapReduce clusters,” in *Proc. USENIX ATC*, 2014, pp. 1–12.
- [15] A. Vahdat, M. Al-Fares, N. Farrington, R. N. Mysore, G. Porter, and S. Radhakrishnan, “Scale-out networking in the data center,” *IEEE Micro*, vol. 30, no. 4, pp. 29–41, 2010.
- [16] S. Gupta and V. Lalitha, “Locality-aware hybrid coded MapReduce for server-rack architecture,” *arXiv:1709.01440*, 2017.
- [17] K. V. Rashmi, N. B. Shah, D. Gu, H. Kuang, D. Borthakur, and K. Ramchandran, “A solution to the network challenges of data recovery in erasure-coded distributed storage systems: A study on the Facebook warehouse cluster,” in *Proc. USENIX HotStorage*, 2013.
- [18] H. A. David and H. N. Nagaraja, *Order Statistics*. Wiley, New York, 2003.
- [19] H. Park, K. Lee, J. Sohn, C. Suh, and J. Moon, “Hierarchical coding for distributed computing,” *arXiv:1801.04686*, 2018.
- [20] P. Brémaud, *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*. Springer Science & Business Media, 2013, vol. 31.
- [21] W. Hoeffding, “Probability inequalities for sums of bounded random variables,” *J. Am. Stat. Assoc.*, vol. 58, no. 301, pp. 13–30, 1963.
- [22] W. Halbawi, Z. Liu, and B. Hassibi, “Balanced Reed-Solomon codes,” in *Proc. IEEE ISIT*, 2016, pp. 935–939.
- [23] —, “Balanced Reed-Solomon codes for all parameters,” in *Proc. IEEE ITW*, 2016, pp. 409–413.