



HAL
open science

Multi-User Linearly-Separable Distributed Computing

Ali Khalesi, Petros Elia

► **To cite this version:**

Ali Khalesi, Petros Elia. Multi-User Linearly-Separable Distributed Computing. *IEEE Transactions on Information Theory*, 2023, 69 (10), pp.6314-6339. 10.1109/TIT.2023.3283967 . hal-04360278

HAL Id: hal-04360278

<https://hal.science/hal-04360278>

Submitted on 21 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Multi-User Linearly-Separable Distributed Computing

Ali Khalesi[✉], *Student Member, IEEE*, Petros Elia[✉], *Member, IEEE*,

Abstract—In this work, we explore the problem of multi-user linearly-separable distributed computation, where N servers help compute the desired functions (jobs) of K users, and where each desired function can be written as a linear combination of up to L (generally non-linear) subtasks (or sub-functions). Each server computes some of the subtasks, communicates a function of its computed outputs to some of the users, and then each user collects its received data to recover its desired function. We explore the computation and communication relationship between how many servers compute each subtask vs. how much data each user receives.

For a matrix F representing the linearly-separable form of the set of requested functions, our problem becomes equivalent to the open problem of sparse matrix factorization $F = DE$ over finite fields, where a sparse decoding matrix D and encoding matrix E imply reduced communication and computation costs respectively. This paper establishes a novel relationship between our distributed computing problem, matrix factorization, syndrome decoding and covering codes. To reduce the computation cost, the above D is drawn from covering codes or from a here-introduced class of so-called ‘partial covering’ codes, whose study here yields computation cost results that we present. To then reduce the communication cost, these coding-theoretic properties are explored in the regime of codes that have low-density parity check matrices. The work reveals — first for the commonly used one-shot scenario — that in the limit of large N , the optimal normalized computation cost $\gamma \in (0, 1)$ is in the range $\gamma \in (H_q^{-1}(\frac{\log_q(L)}{N}), H_q^{-1}(K/N))$ — where H_q is the q -ary entropy function — and that this can be achieved with normalized communication cost that vanishes as $\sqrt{\log_q(N)/N}$. The above reveals an unbounded coding gain over the uncoded scenario, as well as reveals the role of a certain *functional rate* $\log_q(L)/N$ and *functional capacity* $H_q(\gamma)$ of the system. In the end, we also explore the multi-shot scenario, for which we derive bounds on the computation cost.

Keywords—Distributed computing, Linearly separable functions, Coding theory, Sparse matrix factorization, Covering codes, Straggler mitigation, Distributed gradient coding.

I. INTRODUCTION

Distributed computing plays an ever-increasing role in speeding up non-linear and computationally-hard computing tasks. As the complexity of these tasks increases, research seeks novel parallel processing techniques to efficiently offload computations to groups of distributed servers, under various

Part of this work was presented at the 2022 IEEE Information Theory Workshop (ITW).

This work was supported by the European Research Council (ERC) through the EU Horizon 2020 Research and Innovation Program under Grant 725929 (Project DUALITY).

The authors are with the Communication Systems Department at EU-RECOM, 450 Route des Chappes, 06410 Sophia Antipolis, France (email: ali.khalesi@eurecom.fr; elia@eurecom.fr).

frameworks such as MapReduce [1] and Spark [2]. Distributed computing naturally entails several challenges that involve accuracy [3]–[5], scalability [6]–[10], privacy and security [11]–[23], as well as latency and straggler mitigation [24]–[32]. For a detailed survey of related research works, the reader is referred to [33], [34].

This aforementioned effort to efficiently distribute computation load across multiple servers, is intimately intertwined with the concept of communication complexity which refers to the amount of communication required to solve a computation problem when the desired task is distributed among two or more parties [37]. This celebrated computation-vs-communication relationship has been studied in a variety of different forms and scenarios [26], [38]–[49] for various types of problems.

a) Preliminary description of setting: This same relationship between computation and communication costs, is the topic of interest in our work here for the very broad and practical setting of multi-user, multi-server computation of linearly-separable functions. Such functions appear in several classes of problems such as for example in training large-scale machine learning algorithms and deep neural networks with massive data [38], where indeed both computation and communication costs are crucial [50], [51].

In particular, our setting here considers a master node that manages N server nodes that must contribute in a distributed manner to the computation of the desired functions of K different users. Under the linearly-separable assumption (cf. [27]), we consider that user $k \in \{1, 2, \dots, K\}$ demands a function $F_k(D_1, D_2, \dots, D_L)$ that takes as input L datasets D_1, D_2, \dots, D_L , and that each such requested function takes the basic form

$$F_k(D_1, \dots, D_L) = \sum_{\ell=1}^L f_{k,\ell} f_{\ell}(D_{\ell}) = \sum_{\ell=1}^L f_{k,\ell} W_{\ell} \quad (1)$$

where in the above, $W_{\ell} = f_{\ell}(D_{\ell})$ denotes the computed output of a subfunction when the input is D_{ℓ} , and where $f_{k,\ell}$ are the combining coefficients which belong, together with the entries of W_{ℓ} , in some finite field¹. Upon notification of the users’ requests — where these requests are jointly described by the $K \times L$ matrix F that contains the different coefficients $f_{k,\ell}$ — the master instructs the servers to compute some of the subfunctions $f_{\ell}(D_{\ell})$. Each server may naturally

¹The setting nicely includes the case where each F_k itself is a linear combination of some linearly separable functions, i.e., where F_k can itself be written as $F_k(D_1, \dots, D_L) = \sum_{\ell=1}^L \sum_{i=1}^M f_{k,\ell,i} f_{\ell,i}(D_{\ell}) = \sum_{\ell=1}^L \sum_{i=1}^M f_{k,\ell,i} W_{\ell,i}$, corresponding to some set of basis subfunctions $f_{\ell,i}(D_{\ell})$. For simplicity we will henceforth refer to the model in (1).

compute a different number of functions. Upon completing its computations, each server communicates linear combinations of its locally computed outputs (files) to carefully selected subsets of users. Each user can then only linearly combine what it receives from the servers that have transmitted to it, and the goal is for each user to recover its desired function. The problem is completed when every user k retrieves its desired $F_k(D_1, \dots, D_L)$.

We note that there is a clear differentiation between the server nodes that are asked to compute hard (generally non-linear) component functions (subfunctions), and the users that can only linearly combine their received outputs. Generating the so-called output files $W_\ell = f_\ell(D_\ell), \ell \in \{1, 2, \dots, L\}$, can be the result of a computationally intensive task that may for example relate to training a deep learning model on a dataset, or it can relate to the distributed gradient coding problem [24], [52]–[54], the distributed linear-transform computation problem [40], [55], or even the distributed matrix multiplication and the distributed multivariate polynomial computation problems [25], [28]–[30], [44], [56]–[60].

b) *Brief summary of the basic ingredients of the problem:*

Our setting brings to the fore the following crucial questions.

- How many and which servers must compute each subfunction $f_\ell(D_\ell)$?
 - This decision defines the computation cost: the more the servers that compute a subfunction, the higher the computation cost. The extreme centralized scenario where each active server would compute all L subfunctions, would imply a maximal computation cost, but a minimal communication cost, equal to (as we can see) one transmission received per user. The other extreme scenario (for the case of $L = N$) would imply a minimal computation cost of 1 subfunction per server, but a maximal communication cost of N shots received per user.
- What linear combinations of its computed outputs must each server generate?
 - These linear combination coefficients in question, define an $N \times L$ matrix \mathbf{E} that describes which servers compute each subfunction, and how each server combines its computed outputs in order to transmit them. This matrix must be designed in consideration of the requested functions, which are themselves described by the aforementioned $K \times L$ matrix \mathbf{F} .
 - The number of non-zero elements in \mathbf{E} reflects the computation cost on the collective of servers.
- What fraction of the servers must each user get data from, and from which servers?
 - This defines the communication cost. The more data each user gets, the higher the cost.
- How must each user combine (linearly decode) the computed outputs arriving from the servers?
 - This step is determined by a $K \times N$ decoding matrix \mathbf{D} that must be carefully designed. The number of non-zero elements of \mathbf{D} reflects our communication cost. Having a non-sparse \mathbf{D} , implies the need to activate

a substantial fraction of the existing communication links.

- How sparse can \mathbf{D} and \mathbf{E} be so that each user recovers their desired function?
 - This defines the overall costs in computation and communication. As one might expect, the larger the number L of possible subtasks/datasets, the higher the worst-case costs. Having a larger L allows the computing service to provide more refined computations, conceivably though at a higher cost.

To answer these questions, we take a novel approach that employs coding theory. The general idea behind our approach is described as follows.

c) *Brief summary of the new connection to sparse matrix factorization and covering codes:*

- *Connection to the problem of matrix factorization into sparse components:* First, when exploring our distributed computing problem, one can see that the feasibility conditions that ensure that each user recovers its desired function, constitute in fact a (preferably sparse) matrix factorization problem of the form

$$\mathbf{D}\mathbf{E} = \mathbf{F} \quad (2)$$

where the problem is over some q -sized finite field \mathbb{F} , and where any potential sparsity of \mathbf{D} and \mathbf{E} translates to savings in communication and computation costs respectively.

- *Connection to coding theory and syndrome decoding:* To then resolve this problem in a manner that yields non-trivial sparse factors, we notice that — if for example, we were to fix the above matrix \mathbf{D} , and associate this to the parity-check matrix of some linear code — then for each column \mathbf{E}_ℓ of \mathbf{E} and associated column \mathbf{F}_ℓ of \mathbf{F} , the corresponding equation $\mathbf{D} \cdot \mathbf{E}_\ell = \mathbf{F}_\ell$ would tell us that the desired sparse \mathbf{E}_ℓ can be the lowest-weight coset leader whose syndrome is equal to \mathbf{F}_ℓ . Hence, under this analogy, the columns of \mathbf{E} are associated to error vectors, the columns of \mathbf{F} to the corresponding syndromes, and \mathbf{D} is assigned the role of a parity check matrix, and the question is of which code?
- *Connection to covering codes and the new class of partial covering codes:* The above connection with syndromes, in turn brings about the concept of covering codes that refer to codes with good covering properties, which in turn entail low weight \mathbf{E}_ℓ , which is what we need. In (error-control) coding theory though — which generally considers that any error vector is possible — such covering codes consider a full space of possible syndromes, i.e., consider the case where any appropriately-dimensioned vector can indeed be a syndrome. To account for the fact that \mathbf{F} corresponds to a *restricted* set of syndromes (only those that correspond to the columns of our \mathbf{F}), we here consider a new class of *partial covering codes*, the analysis of which is part of this work.
- *Connection with codes having low-density parity-check matrices:* The above effort yields a sparse \mathbf{E} . Our effort is concluded when the aforementioned exploration of

covering codes and partial covering codes (which yielded a sparse \mathbf{E}), is extended to involve analysis of codes with a sparse \mathbf{D} as well.

- *Extending the one-shot scenario:* Our framework allows us to address but also extend the one-shot scenario which is the scenario of choice in various works (see for example [27]) and which, in our case, asks that each server can send only one linear combination to one set of users. We extend this model to the practical and realistic scenario where, for a fixed subset of subfunctions/files $\{f(D_\ell)\}$ computed locally at each server, the server can communicate linear combinations to various sets of users.

d) Highlights of contributions: Our focus is on establishing the normalized computation² cost $\gamma = \frac{1}{N} \max_{l \in \{1, \dots, L\}} \omega(\mathbf{E}(:, l))$, and the normalized communication cost $\delta = \omega(\mathbf{D})/KN$. In our setting, $\gamma \in (0, 1]$ represents the maximum fraction of all servers that must compute any one subfunction, while $\delta \in (0, 1]$ represents the average fraction of servers that each user gets data from, which in turn simply implies an average number of $\Delta = \delta N$ ‘symbols’ received by each user.

We first consider the one-shot case. We proceed to highlight some of the derived results, whose exact statement can be found in the following sections.

- Theorem 1 makes the connection between coding theory and our distributed computing problem, by showing that a (γ, δ) -feasible distributed computing scheme exists if and only if the decoding matrix \mathbf{D} has degree of sparsity δ and is the parity check matrix of an N -length code $C \subset \mathbb{F}^N$ over a field \mathbb{F} where this code has minimum normalized distance from each vector $\{\mathbf{x} \in \mathbb{F}^N | \mathbf{D}\mathbf{x} = \mathbf{F}(:, \ell), \ell \in \{1, \dots, L\}\}$ that is at most γN . This brings to the fore the concept of covering and partial covering codes, where covering codes are codes that guarantee a minimum distance to each vector of the entire vector space, while partial covering codes must guarantee a minimum distance to only a specific subset of the entire space. Establishing the properties of such codes is key to our problem.
- Theorem 2 shows that in the limit of large N , the optimal computation cost per server is in the range $\gamma \in (H_q^{-1}(\frac{\log_q(L)}{N}), H_q^{-1}(K/N))$, where H_q is the entropy function over our field of size q . This theorem reveals the role of what one might refer to as the *functional rate* $R_f = \log_q(L)/N$. The higher this rate, the more ‘involved’ is the space of functions we can compute. In this sense — given that, from the above, $\frac{\log_q(L)}{N} \leq H_q(\gamma)$ — the expression $H_q(\gamma)$ plays the role of an upper bound on what one might call the *functional capacity* of the system.
- Extending the famous covering codes theorem of Blinovskii from [61], we established our bounds on partial covering codes to the setting of codes with low density parity check matrices, revealing that any aforementioned achievable computation cost γ , can be achieved

with *normalized* communication cost that vanishes³ as $\delta \doteq \sqrt{\log_q(N)}/N$. This latter cost will be unboundedly lower than in the uncoded approach of resource-sharing between the two extreme regimes discussed previously in the introduction (See Figure 4 in Section IV-D). As a consequence, we can talk of an unbounded coding gain in our distributed computing problem.

- We also consider the multi-shot scenario where, for the same fixed subset of subtasks/files $\{f_\ell(D_\ell)\}$ computed locally at each server, now the server can communicate different linear combinations to different sets of users. This ability offers a certain degree of refinement that the single-shot scenario may lack. This is exploited, and Theorem 4 reveals a range of parameters for which the multi-shot approach provides computation savings over the single-shot scenario. Interestingly, these computational savings are shown to be unbounded.

A. Paper Organization

The rest of the paper is organized as follows. Section II introduces the model for multi-user distributed computing of linearly separable functions. Section III formulates our problem, focusing on the single-shot scenario, for which Section IV presents the main results. This latter section first makes the connection to coding theory, and then presents the converse and achievability for the computation cost, as well as the achievable bound on the communication cost. Section IV-D offers some insights including a discussion on the gains due to coding. Subsequently, in Section V, we present our proposed achievable multi-shot scheme and the corresponding results, and finally we conclude in Section VI. The appendices are in the subsequent sections.

Notations: We define $[n] \triangleq \{1, 2, \dots, n\}$. For matrices \mathbf{A} and \mathbf{B} , $[\mathbf{A}, \mathbf{B}]$ indicates the horizontal concatenation of the two matrices. For any matrix $\mathbf{X} \in \mathbb{F}^{m \times n}$, then $\mathbf{X}(i, j)$, $i \in [m]$, $j \in [n]$, represents the entry in the i th row and j th column, while $\mathbf{X}(i, :)$, $i \in [m]$, represents the i th row, and $\mathbf{X}(:, j)$, $j \in [n]$ represents the j th column of \mathbf{X} . For two index sets $\mathcal{I} \subset [m]$, $\mathcal{J} \subset [n]$, then $\mathbf{X}(\mathcal{I}, \mathcal{J})$ represents the sub-matrix comprised of the rows in \mathcal{I} and columns in \mathcal{J} . We will use $\omega(\mathbf{X})$ to represent the number of nonzero elements of some matrix (or vector) \mathbf{X} . We denote the finite field $\mathbf{GF}(q)$ as \mathbb{F} . For any code $C \subseteq \mathbb{F}^n$ and any vector $\mathbf{x} \in \mathbb{F}^n$, we use $d(\mathbf{x}, C)$ to represent the Hamming distance of \mathbf{x} to the nearest codeword in C . We will dedicate the use of the letter ρ when referring to normalized covering radii, and we will often use $\rho(C)$ to indicate that this is the normalized covering radius of a specific code $C \subset \mathbb{F}^n$. We will often use the notation $C_{\mathbf{H}}$ to refer to a code whose parity check matrix is \mathbf{H} , and similarly, we will use \mathbf{H}_C to refer to a matrix that serves as the parity-check matrix of a specific linear code C . For some $k \leq n$, $k, n \in \mathbb{N}$, we will also often use the notation $C(k, n)$ to emphasise that a linear code has message length k and codeword length n . For any two codes C_1 and C_2 , we will use $[C_1, C_2]$ to represent the code resulting from their direct

²Both communication and computation costs will be defined in more detail later on. Also, in the following, $\omega(\cdot)$ represents the well known Hamming weight of the argument vector or matrix.

³We will henceforth use \doteq to denote asymptotic optimality. This will be clarified later on.

product. For some vector $\mathbf{x} \in \mathbb{F}^n$, we will use $C_2 = \langle \mathbf{x}, C_1 \rangle$ to represent a code whose basis is the union of \mathbf{x} with the basis of C_1 . Furthermore $V_q(n, \rho)$ will represent the volume of a Hamming ball in \mathbb{F}^n of radius ρn . For $0 \leq x \leq 1 - \frac{1}{q}$, $x \in \mathbb{R}$, we will use $H_q(x) \triangleq x \log_q(q-1) - x \log_q(x) - (1-x) \log_q(1-x)$ to represent the q -ary entropy function, while when $q = 2$ we will use the simplified notation $H(x)$. We will use $\text{sup}(\mathbf{x}^\top)$ to represent the support of some vector $\mathbf{x}^\top \in \mathbb{F}^n$, describing the set of indices of non-zero elements. We will also use the notation $\epsilon(n)$ to represent an expression which, in the limit of large n , vanishes to zero.

II. SYSTEM MODEL

We consider the multi-user linearly-separable distributed computation setting (cf. Fig. 1), which consists of K users/clients, N active (non-idle) servers, and a master node that coordinates servers and users. A main characteristics of this setting is that the tasks performed at the servers, substantially outweigh in computation cost the linear operations performed at the different users. Another defining characteristic is that the cost of having the servers communicate to the users is indeed non-trivial. We consider the setting where each server can use T consecutive time slots to communicate different messages to different subsets of users, where in particular, during time-slot (shot) $t \in [T]$, server n communicates to some arbitrary user-set $\mathcal{T}_{n,t} \subset [K]$, via a dedicated broadcast channel.

In our setting, each user asks for a (generally non-linear) function from a space of linearly separable functions, where each such function takes several datasets as input. Each desired function can be decomposed into a different linear combination of individual (again generally non-linear, and computationally hard) sub-functions $f_\ell(D_\ell)$ that each take a single dataset D_ℓ as input. Consequently the demanded function $F_k(D_1, \dots, D_L)$ of each user $k \in [K]$, is a function of L independent datasets D_ℓ , $\ell \in [L]$, and it takes the general linearly-separable form

$$F_k(D_1, D_2, \dots, D_L) \triangleq f_{k,1}f_1(D_1) + f_{k,2}f_2(D_2) \quad (3)$$

$$+ \dots + f_{k,L}f_L(D_L), \quad k \in [K] \quad (4)$$

$$= f_{k,1}W_1 + f_{k,2}W_2 + \dots \quad (5)$$

$$+ f_{k,L}W_L, \quad k \in [K] \quad (6)$$

where, as previously discussed, $W_\ell = f_\ell(D_\ell) \in \mathbb{F}$, $\ell \in [L]$ is a so-called ‘file’ output, and $f_{k,\ell} \in \mathbb{F}$, $k \in [K]$, $\ell \in [L]$ are the linear combination coefficients. As also mentioned before, F_k itself can be a linear combination of some linearly separable functions.

A. Phases

The model involves three phases, with the first being the *demand phase*, then the *assignment and computation phase* and then the *transmission and decoding phase*. In the demand phase, each user $k \in [K]$ sends the information of its desired function $F_k(\cdot)$ to the master node, who then deduces the linearly-separable decomposition of this function according to (6). Then based on these K desired functions, during the

assignment and computation phase, the master assigns some of the subfunctions to each server, who then proceeds to compute these and produce the corresponding files $W_\ell = f_\ell(D_\ell)$. In particular, each subfunction $f_\ell(D_\ell)$ will be assigned to the servers belonging to some carefully chosen server-set $\mathcal{W}_\ell \subset [N]$.

During the transmission phase, each server $n \in [N]$ broadcasts during time slots $t = 1, 2, \dots, T$, different linear combinations of the locally computed output files, to different subsets of users $\mathcal{T}_{n,t}$. In particular, during time slot t , each server n transmits

$$z_{n,t} \triangleq \sum_{\ell \in [L]} e_{n,\ell,t} W_\ell, \quad n \in [N], t \in [T] \quad (7)$$

where the so-called encoding coefficients $e_{n,\ell,t} \in \mathbb{F}$ are determined by the master. Finally during the decoding part, each user k linearly combines the received signals as follows

$$F'_k \triangleq \sum_{n \in [N], t \in [T]} d_{k,n,t} z_{n,t} \quad (8)$$

for some decoding coefficients $d_{k,n,t} \in \mathbb{F}$, $n \in [N]$, $t \in [T]$, determined again by the master node. Naturally $d_{k,n,t} = 0, \forall k \notin \mathcal{T}_{n,t}$. Decoding is successful when $F'_k = F_k$ for all $k \in [K]$.

B. Computation and Communication Costs

Remembering that $|\mathcal{W}_\ell|$ indicates the number of servers that compute a subfunction $W_\ell = f_\ell(D_\ell)$, $\ell \in [L]$, our *normalized computation cost* metric takes the form

$$\gamma \triangleq \frac{\max_{\ell \in [L]} |\mathcal{W}_\ell|}{N} \quad (9)$$

and represents the maximum fraction of all servers that must compute any subfunction.

We also formally define the *normalized communication cost* as

$$\delta \triangleq \frac{\sum_{t=1}^T \sum_{n=1}^N |\mathcal{T}_{n,t}|}{KN} \quad (10)$$

to represent the average fraction of servers that each user gets data from^{4,5}. Hence in our setting,

$$\Delta \triangleq \delta N \quad (11)$$

represents the average number of transmitted ‘symbols’ received by each user. We wish to provide schemes that correctly compute the desired functions, at reduced computation and communication costs.

⁴We here clarify that our setting implies that any link can be exploited, and our metric simply captures how many of these links are engaged when communicating. Reducing the communication cost implies activating fewer of these links, leaving the rest to be used for other responsibilities of the computing network.

⁵The observant reader may notice the computational cost being a worst-case cost, unlike the communication cost which refers to the average case. This choice is essential in making the connection to coding theory. This same choice though has an advantage; it allows us to better capture the effect of having some subfunctions that are much harder to compute than others.




Fig. 1. The figure represents the K -user, N -server, linearly separable computation setting. In this problem after each user informs the master of its desired function $F_k(\cdot)$, each component subfunction output $W_\ell = f_\ell(D_\ell)$ is computed at each server in $\mathcal{W}_\ell \subset [N]$. During slot t , each server n broadcasts a linear combination $z_{n,t}$ (of the locally available computed files) to all users in $\mathcal{T}_{n,t}$. This combination is defined by the coefficients $e_{n,\ell,t}$. Finally, to decode, each user $k \in [K]$ linearly combines (based on decoding vectors \mathbf{d}_k) all the received signals from all the slots and servers it has received from. Decoding must produce for each user its desired function $F_k(D_1, \dots, D_L)$.

III. PROBLEM FORMULATION: ONE-SHOT SETTING

In this single-shot setting of $T = 1$, we will remove the use of the index t . Thus the transmitted value from (7) will take the form

$$z_n = \sum_{\ell \in [L]} e_{n,\ell} W_\ell, \quad n \in [N] \quad (12)$$

where $e_{n,\ell} \in \mathbb{F}$ will denote the corresponding encoding coefficients, and where each such transmitted value at server n will now be destined for the users in set \mathcal{T}_n . Similarly, the decoding value at each user k (cf. (8)) will take the form $F'_k \triangleq \sum_{n \in [N]} d_{k,n} z_n$, where now $d_{k,n}, n \in [N]$, are the decoding coefficients. The desired functions $F_k(\cdot)$ (cf. (6)), their linear decomposition coefficients $f_{k,\ell}$ (cf. (6)), and the decoded functions $F'_k(\cdot)$ in (8), remain as previously described. With the above in place, we will use

$$\mathbf{f} \triangleq [F_1, F_2, \dots, F_K]^\top \quad (13)$$

$$\mathbf{f}_k \triangleq [f_{k,1}, f_{k,2}, \dots, f_{k,L}]^\top, \quad k \in [K] \quad (14)$$

$$\mathbf{w} \triangleq [W_1, W_2, \dots, W_L]^\top \quad (15)$$

where \mathbf{f} represents the vector of the output demanded functions (cf. (6)), \mathbf{f}_k the vector of function coefficients for user k (cf. (6)), and \mathbf{w} the vector of output files. We also have

$$\mathbf{e}_n \triangleq [e_{n,1}, e_{n,2}, \dots, e_{n,L}]^\top, \quad n \in [N] \quad (16)$$

$$\mathbf{z} \triangleq [z_1, z_2, \dots, z_N]^\top \quad (17)$$

respectively representing the encoding vector at server n , and the overall transmitted vector across all the servers (cf. (12)). Furthermore, we have

$$\mathbf{d}_k \triangleq [d_{k,1}, d_{k,2}, \dots, d_{k,N}]^\top, \quad k \in [K] \quad (18)$$

$$\mathbf{f}' \triangleq [F'_1, F'_2, \dots, F'_K]^\top \quad (19)$$

respectively representing the decoding vector at user k , and the vector of the decoded functions across all the users. In addition, we have

$$\mathbf{F} \triangleq [\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_K]^\top \in \mathbb{F}^{K \times L} \quad (20)$$

$$\mathbf{E} \triangleq [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_N]^\top \in \mathbb{F}^{N \times L} \quad (21)$$

$$\mathbf{D} \triangleq [\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_K]^\top \in \mathbb{F}^{K \times N} \quad (22)$$

where \mathbf{F} represents the $K \times L$ matrix of all function coefficients across all the users, where \mathbf{E} represents the $N \times L$ encoding matrix across all the servers, and where \mathbf{D} represents the $K \times N$ decoding matrix across all the users.

Directly from (6), we have that

$$\mathbf{f} = [\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_K]^\top \mathbf{w} \quad (23)$$

and from (7) we have the overall transmitted vector taking the form

$$\mathbf{z} = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_N]^\top \mathbf{w} = \mathbf{E} \mathbf{w}. \quad (24)$$

Furthermore, directly from (8) we have that

$$F'_k = \mathbf{d}_k^\top \mathbf{z} \quad (25)$$

and thus we have

$$\mathbf{f}' = [\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_K]^\top \mathbf{z} = \mathbf{D} \mathbf{z}. \quad (26)$$

Recall that we must guarantee that

$$\mathbf{f}' = \mathbf{f}. \quad (27)$$

After substituting (23), (24) and (26) into (27), we see that the above feasibility condition in (27) is satisfied iff

$$\mathbf{DE}\mathbf{w} = \mathbf{F}\mathbf{w}. \quad (28)$$

For this to hold for any \mathbf{w} , we must thus guarantee

$$\mathbf{DE} = \mathbf{F}. \quad (29)$$

At this point, since $\mathcal{W}_\ell = \sup(\mathbf{E}(:, \{\ell\})^\top)$, and since $|\mathcal{W}_\ell| = \omega(\mathbf{E}(:, \{\ell\}))$, we have that

$$\max_{\ell \in [L]} \omega(\mathbf{E}(:, \ell)) = \max_{\ell \in [L]} |\mathcal{W}_\ell| \quad (30)$$

which simply tells us that our computation cost γ from (9) takes the form

$$\gamma = \frac{1}{N} \max_{\ell \in [L]} \omega(\mathbf{E}(:, \ell)). \quad (31)$$

Similarly, directly from (8) and (11), we see that

$$\delta = \frac{\omega(\mathbf{D})}{KN} \quad (32)$$

which simply says (cf. (11)) that

$$\Delta = \frac{\omega(\mathbf{D})}{K}. \quad (33)$$

It is now clear that decomposing \mathbf{F} into the product of two relatively sparse matrices \mathbf{D} and \mathbf{E} , implies reduced communication and computation costs respectively.

We here provide a simple example to help clarify the setting and the notation.

A. Simple Example

As illustrated in Figure 2, we consider the example of a system with a master node, $N = 8$ servers, $K = 4$ users, $L = 6$ subfunctions/datasets, and a field of size $q = 7$.

Let us assume that the users ask for the following functions:

$$F_1 = 2f_1(D_1) + 4f_2(D_2) + 4f_3(D_3) + 5f_4(D_4) + 5f_5(D_5) \\ = \mathbf{f}_1^\top \mathbf{w},$$

$$F_2 = 3f_1(D_1) + 4f_2(D_2) + 5f_3(D_3) + 2f_4(D_4) + 6f_5(D_5) \\ + 6f_6(D_6) = \mathbf{f}_2^\top \mathbf{w},$$

$$F_3 = 2f_1(D_1) + 4f_2(D_2) + 6f_3(D_3) + 5f_4(D_4) + 2f_5(D_5) \\ = \mathbf{f}_3^\top \mathbf{w},$$

$$F_4 = 3f_1(D_1) + 5f_2(D_2) + 2f_4(D_4) + 3f_5(D_5) + f_6(D_6) \\ = \mathbf{f}_4^\top \mathbf{w}$$

where $F_k, \mathbf{f}_k, k \in [4]$, and \mathbf{w} , are respectively defined in (6), (15) and (14). Consequently from (20), our function matrix takes the form

$$\mathbf{F} = \begin{bmatrix} 2 & 4 & 4 & 5 & 5 & 0 \\ 3 & 4 & 5 & 2 & 6 & 6 \\ 2 & 4 & 6 & 5 & 2 & 0 \\ 3 & 5 & 0 & 2 & 3 & 1 \end{bmatrix}.$$

In the assignment phase, the master allocates the computation of $f_1(D_1), f_2(D_2), \dots, f_6(D_6)$ to the 8 servers according to

$$\begin{aligned} \mathcal{W}_1 &= \{1, 2, 3, 5, 8\}, \mathcal{W}_2 = \{1, 2, 3, 4, 6, 7\}, \\ \mathcal{W}_3 &= \{1, 2, 3\}, \mathcal{W}_4 = \{1, 4, 5, 7\} \\ \mathcal{W}_5 &= \{1, 2, 4, 5, 6, 8\}, \mathcal{W}_6 = \{3, 4, 5, 6, 7, 8\} \end{aligned}$$

so that for example subfunction $f_3(D_3)$ is assigned to servers $\{1, 2, 3\}$, while we can also see that for example server 2 has to compute $W_1 = f(D_1), W_2 = f(D_2), W_3 = f(D_3)$, and $W_5 = f(D_5)$. A quick inspection shows that the normalized computation cost (cf. (9)) is equal to

$$\gamma = \frac{\max_{\ell \in [6]} |\mathcal{W}_\ell|}{8} = 6/8. \quad (34)$$

After computing their designated output files, each server n transmits z_n as follows

$$\begin{aligned} z_1 &= 2W_1 + 6W_2 + 3W_3 + W_4 + 2W_5, \\ z_2 &= 4W_1 + 5W_2 + 2W_3 + 3W_5, \\ z_3 &= W_1 + 2W_2 + W_3 + 2W_6, \\ z_4 &= W_2 + 2W_4 + 4W_5 + W_6, \\ z_5 &= 2W_1 + W_4 + 3W_5 + 2W_6, \\ z_6 &= 2W_2 + 5W_5 + 3W_6 \\ z_7 &= W_2 + 2W_4 + 4W_6, \\ z_8 &= 2W_1 + 4W_5 + 5W_6 \end{aligned}$$

corresponding to an encoding matrix (cf. (24)) of the form

$$\mathbf{E} = \begin{bmatrix} 2 & 6 & 3 & 1 & 2 & 0 \\ 4 & 5 & 2 & 0 & 3 & 0 \\ 1 & 2 & 1 & 0 & 0 & 2 \\ 0 & 1 & 0 & 2 & 4 & 1 \\ 2 & 0 & 0 & 1 & 3 & 2 \\ 0 & 2 & 0 & 0 & 5 & 3 \\ 0 & 1 & 0 & 2 & 0 & 4 \\ 2 & 0 & 0 & 0 & 4 & 5 \end{bmatrix}.$$

We can quickly verify (cf. (34)) that indeed $\max_{\ell \in [6]} \omega(\mathbf{E}(:, \ell))/8 = 6/8 = \gamma$.

Subsequently, the master asks each server n to send its generated z_n to its designated receiving users, where for each server, these user-sets are:

$$\begin{aligned} \mathcal{T}_1 &= \{2, 4\}, \mathcal{T}_2 = \{1, 3\}, \mathcal{T}_3 = \{3, 4\}, \mathcal{T}_4 = \{1, 2, 3, 4\}, \\ \mathcal{T}_5 &= \{1, 2, 3, 4\}, \mathcal{T}_6 = \{1, 2\}, \mathcal{T}_7 = \{1, 4\}, \mathcal{T}_8 = \{4\} \end{aligned}$$

so now, for example, server 2 will broadcast z_2 to users 1 and 3. A quick inspection also shows that users 1 and 4 will receive 5 different symbols, whereas users 2 and 3 will receive 4 symbols each. The above corresponds to a normalized communication cost (cf. (11)) equal to

$$\delta = \frac{\sum_{n=1}^8 |\mathcal{T}_n|}{4 \cdot 8} = (5 + 4 + 4 + 6)/32 = 19/32$$

corresponding to an average of $\Delta = \frac{19}{4}$ symbols received per user.




Fig. 2. Multi-user distributed computing setting with 8 servers, 4 users, and 6 datasets/subfunctions.

To decode, each user $k \in [4]$ computes the linear combination F'_k as

$$\begin{aligned} F'_1 &= 2z_2 + 3z_4 + 4z_5 + 2z_6 + z_7, \\ F'_2 &= 4z_1 + 2z_4 + z_5 + 3z_6, \\ F'_3 &= 4z_2 + 5z_3 + 2z_4 + z_5, \\ F'_4 &= 4z_1 + 2z_3 + z_4 + 2z_5 + 4z_7 + 5z_8 \end{aligned}$$

adhering to a decoding matrix of the form

$$\mathbf{D} = \begin{bmatrix} 0 & 2 & 0 & 3 & 4 & 2 & 1 & 0 \\ 4 & 0 & 0 & 2 & 1 & 3 & 0 & 0 \\ 0 & 4 & 5 & 2 & 1 & 0 & 0 & 0 \\ 4 & 0 & 2 & 1 & 2 & 0 & 4 & 5 \end{bmatrix}.$$

A quick verification⁶ reveals the correctness of decoding, and that indeed $F'_k = F_k$ for all $k = 1, 2, 3, 4$. For example, for the first user, we see that $F'_1 = 2z_2 + 3z_4 + 4z_5 + 2z_6 + z_7 = 2(4W_1 + 5W_2 + 2W_3 + 3W_5) + 3(W_2 + 2W_4 + 4W_5 + W_6) + 4(2W_1 + W_4 + 3W_5 + 2W_6) + 2(2W_2 + 5W_5 + 3W_6) + (W_2 + 2W_4 + 4W_6) = 2W_1 + 4W_2 + 4W_3 + 5W_4 + 5W_5 + 0W_6$ which indeed matches F_1 . In this example, each user recovers their desired function, with a corresponding normalized computation cost $\gamma = 3/4$ and a normalized communication cost $\delta = 19/32$. This has just been an example to illustrate the setting. The effort to find a solution with reduced computation and communication costs, follows in the subsequent section.

IV. COMPUTATION AND COMMUNICATION COSTS FOR THE SINGLE-SHOT SETTING

In this section we present the results for the one-shot setting. We first rigorously establish the bridge between our problem,

⁶Let us recall that each decoded symbol takes the form $F'_k = \mathbf{d}_k^T \mathbf{z}$ where \mathbf{d}_k^T is the k th row of \mathbf{D} , and where $\mathbf{z} = [z_1 \ z_2 \ \dots \ z_N]^T$.

coding theory, covering and partial covering codes. The main results — focusing first on the computational aspects — are presented in Section IV-B which derives bounds on the optimal computation cost in the large N setting. With these results in place, the subsequent Section IV-C extends our consideration to the communication cost as well. Finally, Section IV-D offers some intuition on the results of this current section.

We briefly recall (cf. [62]) that an n -length code $C \subset \mathbb{F}^n$ is called a ρ -covering code if it satisfies

$$d(\mathbf{x}, C) \leq \rho n, \quad \forall \mathbf{x} \in \mathbb{F}^n \quad (35)$$

for some $\rho \in (0, 1)$ which is referred to as the normalized covering radius.

A. Establishing a Relationship to Covering Codes and Partial Covering Codes

We will first seek to decompose \mathbf{F} into $\mathbf{F} = \mathbf{D}\mathbf{E}$ under a constrained computation cost γ which will generally imply a sparsity constraint on \mathbf{E} . For $\mathbf{E}_\ell \triangleq \mathbf{E}(:, \ell)$ and $\mathbf{F}_\ell \triangleq \mathbf{F}(:, \ell)$ denoting the ℓ th column of \mathbf{E} and \mathbf{F} respectively, we can rewrite our decomposition as

$$\mathbf{D}\mathbf{E}_\ell = \mathbf{F}_\ell, \quad \forall \ell \in [L]. \quad (36)$$

As suggested before, if we viewed $\mathbf{D} \in \mathbb{F}^{K \times N}$ as a parity check matrix $\mathbf{H}_C = \mathbf{D}$ of a code $C \subset \mathbb{F}^N$, then we could view $\mathbf{E}_\ell \in \mathbb{F}^N$ as an arbitrary error pattern, and $\mathbf{F}_\ell \in \mathbb{F}^K$ as the corresponding syndrome. Since we wish to sparsify \mathbf{E}_ℓ , we are interested in having \mathbf{E}_ℓ be the minimum-weight coset leader whose syndrome is \mathbf{F}_ℓ . This is simply the output of the minimum-distance syndrome decoder⁷. To get a first handle

⁷Naturally our viewing \mathbf{D} as a parity check matrix, does not limit the scope of options in choosing \mathbf{D} . Similarly, associating \mathbf{E}_ℓ the role of an error pattern, or a minimum-weight coset leader, is again not a limiting association.

on the weights of \mathbf{E}_ℓ , we can refer to the theory of covering codes which bounds the weights of coset leaders, where these weights are bounded by the code's covering radius $\rho(C)N$, for some normalized radius $\rho(C) \in (0, 1)$. Since the covering radius ρN upper bounds the weights of the coset leaders⁸, it upper bounds our computation cost. A covering radius γN would reflect our computation constraint γ .

To capture some of the coding-theoretic properties, we will transition to the traditional coding-theoretic notation which speaks of an n -length code $C \subseteq \mathbb{F}^n$ of rate k/n , where for us $n = N$ and $k = N - K$. The parity check matrix $\mathbf{H}_C \in \mathbb{F}^{(n-k) \times n}$ will generally be associated to our decoding matrix $\mathbf{D} \in \mathbb{F}^{K \times N}$, the received (or error) vectors $\mathbf{x} \in \mathbb{F}^n$ will be associated to the encoding vectors $\mathbf{E}_\ell \in \mathbb{F}^N$, and its syndrome $\mathbf{s}_x \in \mathbb{F}^{n-k}$ (or just \mathbf{s} , depending on the occasion) will be associated to $\mathbf{F}_\ell \in \mathbb{F}^K$. Please recall that when we write $C_{\mathbf{D}}$ (or $C_{\mathbf{H}}$), we will refer to the code whose parity check matrix is \mathbf{D} (or \mathbf{H}).

As a first step, we extend the concept of covering codes to the following class.

Definition 1. For some $\rho \in (0, 1]$, we say that a set $\mathcal{X} \subseteq \mathbb{F}^n$ is ρ -covered by a code $C \subseteq \mathbb{F}^n$ iff

$$d(\mathbf{x}, C) \leq \rho n, \quad \forall \mathbf{x} \in \mathcal{X} \quad (37)$$

in which case we say that C is a (ρ, \mathcal{X}) -partial covering code.

Naturally when $\mathcal{X} = \mathbb{F}^n$, such a (ρ, \mathcal{X}) -partial covering code is simply the traditional covering code. We are now able to link partial covering codes to our distributed computing problem.

Theorem 1. For the setting of distributed computing with K users, N servers and L subfunctions, a solution to the linearly separable function computation problem $\mathbf{DE} = \mathbf{F}$ with normalized computation cost γ exists if and only if \mathbf{D} is the parity check matrix to a (γ, \mathcal{X}) -partial covering code $C_{\mathbf{D}}$ for some existing set

$$\mathcal{X} \supset \mathcal{X}_{\mathbf{F}, \mathbf{D}} \triangleq \{\mathbf{x} \in \mathbb{F}^N \mid \mathbf{D}\mathbf{x} = \mathbf{F}(:, \ell), \text{ for some } \ell \in [L]\}. \quad (38)$$

With such \mathbf{D} in place, each $\mathbf{E}(:, \ell)$ is the output of the minimum-distance syndrome decoder of $C_{\mathbf{D}}$ for syndrome $\mathbf{F}(:, \ell)$ ⁹

Proof. To first prove that the computation constraint $\gamma = \rho$ indeed requires \mathbf{D} to correspond to a partial covering code that covers \mathcal{X} , let us assume that \mathbf{D} does not have this property and that there exists an $\mathbf{x} \in \mathcal{X}$ such that $d(\mathbf{x}, C_{\mathbf{D}}) > \rho n$. Let \mathbf{c}_{\min} be the closest codeword to \mathbf{x} in the sense that $d(\mathbf{x}, \mathbf{c}_{\min}) = d(\mathbf{x}, C_{\mathbf{D}})$. Now let $\mathbf{e}_{\min} = \mathbf{x} - \mathbf{c}_{\min}$, and note, directly from the above assumption, that $\omega(\mathbf{e}_{\min}) > \rho n$. Naturally $\mathbf{D}\mathbf{x} = \mathbf{D}(\mathbf{e}_{\min} + \mathbf{c}_{\min}) = \mathbf{D}\mathbf{e}_{\min}$ by virtue of the fact that \mathbf{D} is the parity check matrix of $C_{\mathbf{D}}$. Since $\mathbf{x} \in \mathcal{X}$, we know that $\exists \ell \in [L]$ such that $\mathbf{D}\mathbf{x} = \mathbf{F}(:, \ell)$, which directly means that $\exists \ell \in [L]$ such that $\mathbf{D}\mathbf{e}_{\min} = \mathbf{F}(:, \ell)$. This \mathbf{e}_{\min} is the coset leader associated to syndrome $\mathbf{F}(:, \ell)$.

Since though $\mathbf{DE} = \mathbf{F}$, we also have that $\mathbf{DE}(:, \ell) = \mathbf{F}(:, \ell)$. Since $\mathbf{E}(:, \ell)$ and \mathbf{e}_{\min} are in the same coset (of the same

syndrome $\mathbf{F}(:, \ell)$), and since \mathbf{e}_{\min} is the minimum-weight coset leader, we can conclude that $\omega(\mathbf{E}(:, \ell)) \geq \omega(\mathbf{e}_{\min})$. Thus the assumption that $\omega(\mathbf{e}_{\min}) > \rho n$ implies that $\omega(\mathbf{E}(:, \ell)) > \rho n$ which contradicts the computation-cost requirement that $\omega(\mathbf{E}(:, \ell)) \leq \rho n$ from (31). Thus if \mathbf{D} does not correspond to a partial covering code (with $\rho = \gamma$) that covers $\mathcal{X}_{\mathbf{F}, \mathbf{D}}$, the complexity constraint is violated.

On the other hand, recalling that $C_{\mathbf{D}}$ is a partial covering code for \mathcal{X} , we get that for any $\mathbf{x} \in \mathcal{X}$ then $d(\mathbf{x}, C_{\mathbf{D}}) \leq \rho n$. For the same $\mathbf{x} \in \mathcal{X}$, let \mathbf{c}_{\min} be again its closest codeword, and let $\mathbf{e}_{\min} = \mathbf{x} - \mathbf{c}_{\min}$, where again by definition of the partial covering code, $\omega(\mathbf{e}_{\min}) \leq \rho n$. Since, like before, $\mathbf{D}\mathbf{e}_{\min} = \mathbf{F}(:, \ell)$ for some $\ell \in [L]$, then we simply set $\mathbf{E}(:, \ell) = \mathbf{e}_{\min}$ whose weight is indeed sufficiently low to guarantee the computation constraint. We recall that for each $\mathbf{F}(:, \ell)$, this coset leader $\mathbf{E}(:, \ell) = \mathbf{e}_{\min}$ can be found by using the minimum-distance syndrome decoder. \square

Intuitively, a smaller \mathcal{X} could potentially — depending on \mathcal{X} and the code — be covered in the presence of a smaller covering radius. Now that we have established the connection with partial covering codes, we proceed to present computation bounds. The following result, as well as all subsequent results, assumes a large N .

B. Bounds on the Optimal Computation Cost

The following theorem bounds the optimal computation cost of the multi-user linearly-separable computation setting.

Theorem 2. For the setting of distributed-computing of linearly-separable functions, with K users, N servers and any number of L subfunctions, in the limit of large N (with K/N and $\log_q(L)/N$ fixed) the optimal computation cost is bounded as

$$\gamma \in (H_q^{-1}(\frac{\log_q(L)}{N}), H_q^{-1}(\frac{K}{N})). \quad (39)$$

Proof. The proof of the converse (lower bound in (39)) employs sphere-covering arguments and can be found in Appendix A. The proof of achievability follows from covering- and partial covering-code arguments, and can be found in Appendix B. \square

Remark 1. The two bounds meet when $L = q^K$.

Remark 2. An observant reader may notice that while the communication cost has an average-case interpretation, the computational cost metric retains a worst-case nature. This is out of necessity, as the coding-theoretic bound that supports the bound on γ is based on a worst-case analysis of the Hamming weight of each column of \mathbf{E} while on the other hand, for bounding the communication cost, we are forced to consider an average metric simply because again the coding-theoretic bound that we have derived, now draws from a bound on the Hamming weight of the entire matrix \mathbf{D} .

Theorem 2 suggests a range of computation costs. In the next corollary, we will describe the conditions under which a reduced normalized computation cost, strictly inside this range, can be achieved. This reduced cost will relate to (our

⁸Let us recall (cf. [63]) that the preferred coset leaders are the minimum-weight vectors in each row of the standard array.

⁹ $\mathcal{X}_{\mathbf{F}, \mathbf{D}}$ simply means that $\mathbf{x} \in \mathcal{X}_{\mathbf{F}, \mathbf{D}} \iff \exists \ell \in [L] : \mathbf{D}\mathbf{x} = \mathbf{F}(:, \ell)$.

ability to choose) a set $\mathcal{X} \subset \mathbb{F}^N$. As we will see, a smaller \mathcal{X} will imply a smaller γ . To understand the connection between our problem and this set \mathcal{X} , and thus to better understand the following theorem whose proof will be presented in Appendix C, we provide the following sketch of some crucial elements in the proof of Theorem 2. In particular, we will here sketch an algorithm that iterates in order to converge to the aforementioned \mathcal{X} , and then to the corresponding decoding matrix \mathbf{D} , that will eventually provide reduced normalized complexity γ . Before describing the algorithm, it is worth noting that a crucial ingredient can be found in Lemma 1 (see Appendix D), which modifies the approach in [64] in order for us to design — for any set $\mathcal{X}' \in \mathbb{F}^N$ — a (ρ, \mathcal{X}') -partial covering code for some $\rho = H_q^{-1}\left(\frac{K}{N} - \left(1 - \frac{\log_q(|\mathcal{X}'|)}{N}\right)\right)$.

With this in place, the algorithm starts by picking an initial set $\mathcal{X}_0 \in \mathbb{F}^N, |\mathcal{X}_0| = Lq^{N-K}$, and then applies Lemma 1 to construct a (ρ_0, \mathcal{X}_0) -partial covering code, \mathcal{C}_0 , where $\rho_0 = H_q^{-1}\left(\frac{K}{N} - \left(1 - \frac{\log_q(|\mathcal{X}_0|)}{N}\right)\right)$. With this code \mathcal{C}_0 in place, we create — as a function of \mathcal{C}_0 — the set $\mathcal{X}_{\mathbf{F}, \mathbf{D}, 0}$ as defined in (38) where $\mathbf{D} = \mathbf{H}_{\mathcal{C}_0}$, and then we check if $\mathcal{X}_0 \supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}, 0}$. If so, then the algorithm terminates, else it goes to the next iteration which starts by picking a new larger set $\mathcal{X}_1 \in \mathbb{F}^N, |\mathcal{X}_1| = Lq^{N-K} + 1$, then uses Lemma 1 to create a new (ρ_1, \mathcal{X}_1) -partial covering code for $\rho_1 = H_q^{-1}\left(\frac{K}{N} - \left(1 - \frac{\log_q(|\mathcal{X}_1|)}{N}\right)\right)$, and then compares if $\mathcal{X}_1 \supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}, 1}$. This procedure terminates during some round m where this terminating round is the first round for which the chosen set \mathcal{X}_m (now of cardinality $|\mathcal{X}_m| = Lq^{N-K} + m$) and the corresponding (ρ_m, \mathcal{X}_m) -partial covering code with $\rho_m = H_q^{-1}\left(\frac{K}{N} - \left(1 - \frac{\log_q(|\mathcal{X}_m|)}{N}\right)\right)$, yield $\mathcal{X}_m \supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}, m}$.

In the following corollary, the mentioned \mathcal{X} refers to the terminating¹⁰ \mathcal{X}_m , and the decoding matrix \mathbf{D} will be the parity-check matrix of the aforementioned (ρ_m, \mathcal{X}_m) -partial covering code that covers the terminating $\mathcal{X} = \mathcal{X}_m$, while the normalized computation cost in the theorem will take the form $\gamma = \rho = \rho_m$.

With the above in place, the following speaks of a set \mathcal{X} that is ρN -covered by a code $\mathcal{C}_{\mathbf{D}}$ that generates — as described in (38) — its set $\mathcal{X}_{\mathbf{F}, \mathbf{D}}$.

Corollary 1. In the multi-user linearly separable computing problem $\mathbf{DE} = \mathbf{F}$, if there exists a set

$$\mathcal{X} \supset \mathcal{X}_{\mathbf{F}, \mathbf{D}} \triangleq \{\mathbf{x} \in \mathbb{F}^N \mid \mathbf{D}\mathbf{x} = \mathbf{F}(:, \ell), \text{ for some } \ell \in [L]\}$$

that is ρN -covered by a code $\mathcal{C}_{\mathbf{D}}$ for $\rho = H_q^{-1}\left(\frac{K}{N} - \left(1 - \frac{\log_q(|\mathcal{X}|)}{N}\right)\right)$, then the computation cost

$$\gamma = H_q^{-1}\left(\frac{K}{N} - \left(1 - \frac{\log_q(|\mathcal{X}|)}{N}\right)\right)$$

is achievable. If $\mathcal{X} = \mathcal{X}_{\mathbf{F}, \mathbf{D}}$, then $\gamma = H_q^{-1}\left(\frac{\log_q(L)}{N}\right)$ is achievable and optimal.

Proof. The proof can be found in Appendix C. \square

¹⁰Note that in the worst case this termination will happen when $\mathcal{X}_m = \mathbb{F}^N$, in which case the output code will be a covering code.

As suggested before, the above reflects that covering a smaller \mathcal{X} could entail a smaller covering radius and thus a smaller computation cost.

C. Jointly Considering Computation and Communication Costs

The following theorem combines computation and communication considerations. Theorem 3 builds on Theorem 1, where now we recall that any chosen decoding matrix \mathbf{D} will automatically yield a normalized communication cost $\delta = \frac{\omega(\mathbf{D})}{KN}$ corresponding to $\Delta = \delta N = \frac{\omega(\mathbf{D})}{K}$. The following bounds this communication cost.

Theorem 3. For the setting of distributed-computing of linearly-separable functions, with K users, N servers and L subfunctions, in the limit of large N (with K/N and $\log_q(L)/N$ fixed) the optimal computation cost is bounded as

$$\gamma \in \left(H_q^{-1}\left(\frac{\log_q(L)}{N}\right), H_q^{-1}\left(\frac{K}{N}\right)\right)$$

and for any achievable computation cost $\gamma \leq \min\{\frac{\sqrt{5}-1}{2}, 1 - \frac{1}{q}\}$, then the corresponding achievable communication cost takes the form

$$\delta \doteq \frac{\sqrt{\log_q(N)}}{N}. \quad (40)$$

Proof. The proof can be found in Appendix E. \square

We here offer a quick sketch of the proof of the above theorem. The proof first employs a modified version of the famous result by Blinovskii in [61] which proved that, as n goes to infinity, almost all random linear codes $\mathcal{C}(k, n)$ are covering codes, as long as the normalized covering radius satisfies $\rho \geq H_q^{-1}\left(\frac{n-k}{n}\right)$. This modification of Blinovskii's theorem is presented in Theorem 5, whose proof is found in Appendix E. With this modification in place, we prove that almost all (k, n) random linear codes with

$$\rho = H_q^{-1}\left(\frac{\log_q(|\mathcal{X}|) - k}{n}\right) \quad (41)$$

are (ρ, \mathcal{X}) -partial covering codes, each for their own set $\mathcal{X} \in \mathbb{F}^n$. This is again in Theorem 5. With this theorem in place, we then employ a concatenation argument (which can be found in the proof of Theorem 3 in Appendix E), to build a sparse parity-check matrix \mathbf{H} of a partial covering code, which — by virtue of the connection made in Theorem 1 — allows us to complete the proof of Theorem 3.

To show that sparse parity check codes can indeed offer reduced computation costs, we had to show that sparse codes can indeed offer good partial covering properties. To do that, we followed some of the steps described below. In particular, we designed an algorithm that begins with constructing a sparse parity check code that can cover, for a given radius ρ_0 , a minimum necessary cardinality set \mathcal{X}_0 , where this minimum cardinality of $|\mathcal{X}_0| = Lq^{N-K}$ is imposed on us by \mathbf{F} . The parity-check matrix of this first code is \mathbf{H}_0 . Then following the steps in the proof of Theorem 1, we set $\mathbf{D} = \mathbf{H}_0$ and check if $\mathcal{X}_0 \supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}}$ holds. If it indeed holds, the algorithm outputs

\mathbf{D} and \mathcal{X}_0 , and the corresponding cost is $\gamma = \rho_0$, where this ρ value is derived from (41) by setting $\mathcal{X} = \mathcal{X}_0$. Otherwise the algorithm constructs another sparse partial covering code with a new parity check matrix \mathbf{H}_1 , now covering a set \mathcal{X}_1 with cardinality $|\mathcal{X}_1| = Lq^{N-K} + 1$, and then checks again the same inclusion condition as above. The procedure continues until it terminates, with some covered set \mathcal{X}_m of cardinality $|\mathcal{X}_m| = Lq^{N-K} + m$. As before, reaching $\mathcal{X}_m = \mathbb{F}^N$ will terminate the algorithm (if it has not terminated before that). In the proposition below, the set \mathcal{X} is exactly our terminating set \mathcal{X}_m we referred to above.

Proposition 1. *After adopting the achievable scheme proposed in Theorem 3 together with adopting the corresponding conditions on ρ, δ and its corresponding \mathbf{D} that was designed as a function of \mathbf{F} , then if there exists a subset $\mathcal{X} \supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}}, \mathcal{X} \subseteq \mathbb{F}^N$, that is ρN -covered by $\mathbf{C}_{\mathbf{D}}$ for some $\rho = H_q^{-1}(\frac{K}{N} - (1 - \frac{\log_q(|\mathcal{X}|)}{N}))$, we can conclude that the computation cost $\gamma = H_q^{-1}(\frac{K}{N} - (1 - \frac{\log_q(|\mathcal{X}|)}{N}))$ is achievable. If $\mathcal{X} = \mathcal{X}_{\mathbf{F}, \mathbf{D}}$, then the computation cost converges to the optimal $H_q^{-1}(\frac{\log_q(L)}{N})$. The above remains in place for any \mathbf{D} which yields communication cost no less than $\Delta = O(\sqrt{\log_q(N)})$.*

Proof. The proof can be found in Appendix G. \square

D. Discussing the Results of the Current Section

Theorem 3 reveals that the optimal computation cost lies in the region $\gamma \in (H_q^{-1}(\frac{\log_q(L)}{N}), H_q^{-1}(\frac{K}{N}))$, and that this cost can be achieved with a communication cost that vanishes as $\delta \doteq \sqrt{\log_q(N)}/N$. To get a better sense of the improvements that come from our coded approach, let us compare this to the uncoded case. Looking at Figure 3, this uncoded performance is described by (the line connecting) point 1 and point 2. Point 1, located at point $(\gamma = 1/N, \delta = 1)$, corresponds to the fully decentralized scenario where each server must compute just one subfunction¹¹, but which in turn implies that each server must communicate to all K users. This scenario corresponds to the decomposition $\mathbf{DI} = \mathbf{F}$ where we maximally¹² sparsify \mathbf{E} by setting it equal to $\mathbf{E} = \mathbf{I}_{N \times N}$.

On the other hand, point 2, located at $(\gamma = 1, \delta = 1/N)$, corresponds to the fully centralized scenario where each of the K activated servers¹³ is asked to compute all L subfunctions, but where now each server need only transmit to a single user.

To clarify, a scheme achieves a certain (δ, γ) when all the elements in $\mathcal{X}_{\mathbf{F}, \mathbf{D}}$ are γ -covered by $\mathbf{C}_{\mathbf{D}}$ independently of \mathbf{F} , and that the above holds for a \mathbf{D} that has only a fraction δ of non-zero elements. In particular, from Theorem 3 we now know that point 3 at $(\gamma = H_q^{-1}(\frac{K}{N}), \delta \doteq \sqrt{\log_q(N)}/N)$ is always achievable irrespective of \mathbf{F} . So is any point inside the triangle defined by points 1, 2, 3.

¹¹Due to the single-shot assumption, this corresponds to having $N(q-1) = L$. This matches the converse — in our large N setting — because after writing $L = N(q-1) = \binom{N}{1}(q-1) \simeq q^{NH_q(1/N)} = q^{NH_q(\gamma)}$, we see that $H_q^{-1}(\frac{\log_q(L)}{N}) = \gamma = \frac{1}{N}$.

¹²Note that the stated $\delta = 1$ accounts for the worst-case scenario where \mathbf{F} contains no zero elements.

¹³This number of activated servers is again a consequence of the single-shot assumption.

To go beyond point 3, assuming that $\mathbf{C}_{\mathbf{D}}$ is a (γ, \mathcal{X}) -partial covering code, then we achieve γ if $\mathcal{X} \supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}}$, where this reduced γ derived in Proposition 1. Note that the satisfaction of such a condition is dependent on \mathbf{F} . Also, note that the result of Theorem 5 allows us to construct a Low-Density Parity Check partial-covering code $\mathbf{C}_{\mathbf{D}}$ that satisfies $\mathcal{X} \supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}}$, in an algorithmic procedure described in the proof of Theorem 3. The algorithmic procedure offers a way to give instances of some conditional achievable schemes dependent on the choice of \mathbf{F} , which though may not be optimal for all \mathbf{F} . If the suggested procedure terminates at the first round of the procedure then point 4 is achieved (see Appendix E).

With the above in mind, we see that any point inside the region defined by points 1, 4, 2, 3, can be achieved under some conditions over \mathbf{F} . The converse also tells us that no point to the left of point 5, i.e., no point with $\gamma < H_q^{-1}(\frac{\log_q(L)}{N})$, can be achieved. Therefore point 5 is simply a converse point. Finally, the achievability of the points inside the triangle defined by points 5, 2, 4 remains an open problem.

V. DISTRIBUTED COMPUTING OF LINEARLY-SEPARABLE FUNCTIONS WITH MULTI-SHOT COMMUNICATIONS ($T > 1$)

In this section we present our results for the multi-shot setting where each server is able to broadcast T consecutive transmissions to T potentially different subsets of users. This is mainly motivated by the fact that having $T > 1$, naturally allows us to employ fewer servers, but it is also motivated — as we will discuss later on — by an additional coding flexibility and refinement that multiple transmissions can provide. We briefly note that we assume as before that K and N are sufficiently large.

A. Problem Formulation

The notation of the parameters that characterize the system will now generally follow directly from Section III, sometimes after clarifying the corresponding time-slot t of interest. For example, as before we will have

$$\mathbf{f} \triangleq [F_1, F_2, \dots, F_K]^\top, \quad (42)$$

$$\mathbf{f}_k \triangleq [f_{k,1}, f_{k,2}, \dots, f_{k,L}]^\top, \quad k \in [K], \quad (43)$$

$$\mathbf{w} \triangleq [W_1, W_2, \dots, W_L]^\top, \quad (44)$$

$$\mathbf{f}' \triangleq [F'_1, F'_2, \dots, F'_K]^\top, \quad (45)$$

$$\mathbf{F} \triangleq [\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_K]^\top. \quad (46)$$

On the other hand, the notation for the encoding coefficients and the corresponding transmitted symbols during slot t , will now take the slightly modified form

$$\mathbf{e}_{n,t} \triangleq [e_{n,1,t}, e_{n,2,t}, \dots, e_{n,L,t}]^\top, \quad n \in [N], \quad t \in [T], \quad (47)$$

$$\mathbf{z}_t \triangleq [z_{1,t}, z_{2,t}, \dots, z_{N,t}]^\top, \quad t \in [T], \quad (48)$$

$$\mathbf{z} \triangleq [\mathbf{z}_1^\top, \mathbf{z}_2^\top, \dots, \mathbf{z}_T^\top]^\top \quad (49)$$

with the corresponding modified

$$\mathbf{E}_t \triangleq [\mathbf{e}_{1,t}, \mathbf{e}_{2,t}, \dots, \mathbf{e}_{N,t}]^\top, \quad t \in [T] \quad (50)$$

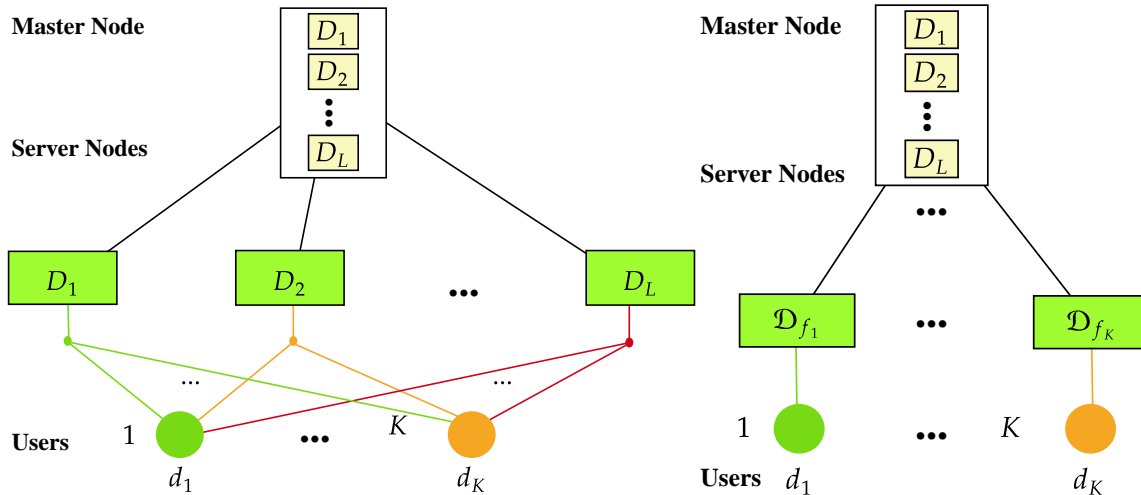


Fig. 3. (Left. Fully decentralized): Uncoded scheme for point 1 corresponding to $(\gamma = 1/N, \delta = 1)$. Each of the $N(q - 1) = L$ servers, computes one subfunction, but must send to all K users. (Right. Fully centralized): Uncoded scheme for point 2 corresponding to $(\gamma = 1, \delta = 1/K)$. K activated servers, each computing L subfunctions, and each transmitting to a single user.




Fig. 4. The figure summarizes the results of Theorem 3. Recall that while N is asymptotically large, both K/N and $\log_q(L)/N$ are fixed.

while the corresponding decoding coefficients will now take the form

$$\mathbf{d}_{k,t} \triangleq [d_{k,1,t}, d_{k,2,t}, \dots, d_{k,N,t}]^\top, \quad k \in [K], t \in [T], \quad (51)$$

$$\mathbf{d}_k \triangleq [\mathbf{d}_{k,1}^\top, \mathbf{d}_{k,2}^\top, \dots, \mathbf{d}_{k,T}^\top]^\top, \quad k \in [K]. \quad (52)$$

We note that the decoding coefficients are decided as a function of all received signals throughout all T transmissions.

As before, (cf. (6)), we have that

$$\mathbf{f} = [\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_K]^\top \mathbf{w} \quad (53)$$

and now we use

$$\mathbf{z}_t = \mathbf{E}_t \mathbf{w} = [\mathbf{e}_{1,t}, \mathbf{e}_{2,t}, \dots, \mathbf{e}_{N,t}]^\top \mathbf{w} \quad (54)$$

to denote the t -th slot transmission vector across all servers. The set of all transmissions now takes the form

$$\mathbf{z} = \mathbf{E} \mathbf{w} \quad (55)$$

where now the encoding matrix takes the form

$$\mathbf{E} \triangleq [\mathbf{E}_1^\top, \mathbf{E}_2^\top, \dots, \mathbf{E}_T^\top]^\top \in \mathbb{F}^{NT \times L}. \quad (56)$$

Upon decoding, each user k generates

$$\mathbf{F}'_k = \mathbf{d}_k^\top \mathbf{z} \quad (57)$$

and the cumulative set of all decoded elements across the users takes the form

$$\mathbf{f}' = [\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_K]^\top \mathbf{z}. \quad (58)$$

Now we note that our decoding matrix

$$\mathbf{D} \triangleq [\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_K]^\top \in \mathbb{F}^{K \times NT} \quad (59)$$

is of dimension $K \times NT$. Naturally, correct decoding requires

$$\mathbf{f} = \mathbf{f}' \quad (60)$$

and after substituting (53), (54), (58) into (60), we can conclude as before that computing succeeds if and only if

$$\mathbf{D} \mathbf{E} = \mathbf{F}. \quad (61)$$

The problem remains similar to the one in the single-shot scenario, except that now our decoding matrix $\mathbf{D} \in \mathbb{F}^{K \times NT}$ and encoding matrix $\mathbf{E} \in \mathbb{F}^{NT \times L}$ are bigger¹⁴ and can have a certain restrictive structure.

Again similar to before, each server $n \in [N]$ is asked to compute all the subfunctions in $\cup_{t=1}^T \text{sup}(\mathbf{e}_{n,t})$, and thus equivalently the set of servers \mathcal{W}_ℓ that must compute subfunction $f_\ell(D_\ell)$, takes the form

$$\begin{aligned} \mathcal{W}_\ell &= \cup_{t=1}^T \text{sup}(\mathbf{E}([(t-1)N+1 : tN], \{\ell\})^\top), \\ &\forall \ell \in [L], \forall t \in [T]. \end{aligned} \quad (62)$$

The following theorem provides an achievable upper bound on the computation cost γ of our distributed computing setting for the multi-shot scenario.

Theorem 4. *For the setting of distributed-computing of linearly-separable functions, with K users, N servers, L subfunctions and T shots, in the limit of large N (with $K|NT$*

¹⁴The size of $\mathbf{F} \in \mathbb{F}^{K \times L}$ remains the same, and thus again we have $L \leq q^K$.

and T fixed) the optimal computation cost γ is upper bounded by

$$\gamma \leq TH_q^{-1}\left(\frac{K}{NT}\right). \quad (63)$$

Proof. We first note that, directly from (62) and the union bound, we have that

$$\max_{\ell \in [L]} \omega(\mathbf{E}(:, \ell)) \geq \max_{\ell \in [L]} |\mathcal{W}_\ell| \quad (64)$$

and thus our normalized computation cost will be upper bounded as

$$\gamma \leq \max_{\ell \in [L]} \omega(\mathbf{E}(:, \ell))/N.$$

To bound $\max_{\ell \in [L]} \omega(\mathbf{E}(:, \ell))$, we apply covering code arguments as in the single-shot case, after though accounting for the dimensionality change from having larger matrices. In particular, this means that now the corresponding covering code $C(n, k)$ will have $n = NT$ and again $K = n - k$ (now we only ask that $NT \geq K$). To account for this increase in n , we note that while the computation cost must still be normalized by the same number of servers N , when considering our covering code¹⁵, we must consider a radius $\frac{\gamma}{T}n = \frac{\gamma}{T}NT = \gamma N$ to guarantee our computation constraint. In other words, the ρ -covering codes that will guarantee the computation constraint, will be for $\rho = \gamma/T$. Consequently, combined with the aforementioned union bound, we now see that ρT serves as an achievable upper bound on γ . The rest follows directly from the proof of the corresponding theorem in the single-shot scenario. \square

The following two propositions help us make sense of the computational effect of having $T > 1$.

Proposition 2. *In the distributed computing setting of interest in the limit of large T , the normalized computation cost γ vanishes to zero.*

Proof. The proof is direct once we prove that for any fixed c , then

$$\lim_{T \rightarrow \infty} TH_q^{-1}(c/T) = 0. \quad (65)$$

This property will be proved in Appendix I-A. \square

In a system with an unchanged number of users and servers, the above reveals the notable (unbounded) computational advantage of allowing a large number T of distinct transmissions per server. This advantage of the multi-shot approach must be seen in light of the fact that in the single-shot approach, the computation cost γ was always bounded below by a fixed $\gamma \geq H_q^{-1}\left(\frac{\log_q(L)}{N}\right)$, irrespective of the communication cost. Consequently we can deduce that the computational gains that we see in the regime of larger T , are — at least partly — a result of the increased refinement in transmission that a larger T allows, and it should not be solely attributed to an increased communication cost.

¹⁵Let us quickly recall that in the previous single-shot scenario, a covering code with covering radius $\rho n = \rho N$ implied a computation cost of $\gamma N = \rho N$ and thus a normalized computation cost of $\gamma = \rho$.

The following proposition discusses the non-asymptotic computational effect of increasing T beyond 1. Recall that our results hold for sufficiently large K and N .

Proposition 3. *For $q = 2$, then γ monotonically decreases in T , while for $q > 2$ then γ monotonically decreases in T after any $T \geq \lceil \frac{K}{NH_q^{-1}(1/q)} \rceil$.*

Proof. The proof is based on the fact that the derivative of $f = TH_q^{-1}(c/T)$, $0 \leq c/T \leq 1 - 1/q$, with respect to T , satisfies

$$\frac{\partial f}{\partial T} = \frac{H_q(f/T)}{\log_q\left(\frac{f/T}{1-f/T}(q-1)\right)} + f/T. \quad (66)$$

This is proved in Appendix I-B. From the above, and after observing that $\frac{\partial f}{\partial T} \leq 0$ where $0 \leq H_q^{-1}(K/NT) = f/T \leq 1/q$, we can conclude that since $0 \leq H_q^{-1}(K/NT) = f/T \leq 1/2 = 1 - 1/q$, then for $q = 2$, increasing T always strictly reduces γ . On the other hand, when $q > 2$, this reduction happens — as we see above — when $T \geq [T_0]$ for some real T_0 for which $H_q(K/NT_0) = 1/q$. \square

VI. CONCLUSIONS

In this work we have introduced a new multi-user distributed-computation setting for computing from the broad class of linearly-separable functions.

Our work revealed the link between distributed computing and the problem of factorizing a ‘functions’ matrix \mathbf{F} into a product of two preferably sparse matrices, these being the encoding matrix \mathbf{E} and the decoding matrix \mathbf{D} . The work then made the new connection to the area of covering codes, revealing for the first time the importance of these codes in distributed computing problems, as well as in sparse matrix factorization over finite fields. Furthermore, this work here brought to the fore the concept of partial covering codes, and the need for codes that cover well smaller subsets of the ambient vector space. For this new class of codes — which constitute a generalization of covering codes — we have provided some extensions and generalizations of well-studied results in the literature.

Our two metrics — γ , representing the maximum fraction of all servers that must compute any subfunction, and δ , representing the average fraction of servers that each user gets data from — capture the computation and communication costs, which are often at the very core of distributed computing problems. The observant reader might notice that the creation of \mathbf{E} entails a complexity equal to that of syndrome decoding. Our results hold unchanged when we consider — as suggested before — that the computational cost of evaluating the various subfunctions, far exceeds all other costs. What the results reveal is that in the large N regime, the optimal computation cost lies in the region $\gamma \in (H_q^{-1}(\frac{\log_q(L)}{N}), H_q^{-1}(\frac{K}{N}))$, and that this entails the use of a vanishingly small fraction $\delta \doteq \sqrt{\log_q(N)}/N$ of all communication resources. What we show is that our coded approach yields unbounded gains over the uncoded scenario, in the sense that the ratio $\frac{\delta_{un}(\gamma)}{\delta(\gamma)}$ between the uncoded and coded communication costs, is unbounded.

We have also studied the multi-shot setting, where we have explored the gains over the single-shot approach. What we now know is that the gains from increasing T , are unbounded (and strictly increasing) in the regime of large T , whereas in the regime of finite T , the gains are strictly increasing after some threshold value of T . We are thus able to conclude, as suggested before, that computation reductions due to larger T , are — at least partly — a result of the increased refinement in a transmission that a larger T allows, and that these gains should not be interpreted as being purely the result of an increased communication load.

Our work naturally relates partly to the recent results in [27] that considered the single-user linearly-separable distributed computing scenario, where a single user may request multiple linearly-separable functions. In this setting in [27], as well as in the extended works in [65] and [66], a key ingredient is the presence of straggling servers, while another key ingredient is that the subfunction-assignment is fixed and oblivious of the actual functions requested by the user. In this context, the coefficients of the functions are assumed to be distributed uniformly and *i.i.d.*, and the decodability is probabilistic. There is also an interesting connection (cf. [67], [68]) between compressed sensing and coding theory. Naturally this connection entails no link to covering codes, as the problem of compressed sensing relates to decodability and is very different from the existence problem that we are faced with.

As suggested above, our setting can apply to a broad range of ‘well-behaved’ functions, and thus can enjoy several use cases, some of which are suggested in our introduction (see also [27] for additional motivation of the linearly separable function computation problem). When considering problems over the real numbers, we may consider a very large q . An additional new scenario that our work can extend is the so-called hierarchical or tree-like scenario introduced in [45], [69] whose purpose is to ameliorate bandwidth limitations and straggler effects in distributed gradient coding [25]. In this hierarchical setting, each user¹⁶ is connected to a group of servers in a hierarchical manner¹⁷ that allows for a hierarchical aggregation of the sub-gradients. Our approach can extend the hierarchical model by allowing the users to connect to any subset of servers, as well as by allowing them to deviate from the single-shot assumption.

Additional considerations that involve stragglers, channel unevenness or computational heterogeneity, are all interesting research directions.

¹⁶In [45], these users are referred to as master nodes.

¹⁷In particular, each user computes a linearly separable function based on its locally available data, and then sends this to the ‘Aggregator’ that finally computes the gradient.

APPENDIX A
PROOF OF CONVERSE IN THEOREM 2

To prove the converse in (39), we modify the sphere-covering bound for the case of partial covering codes. We wish to show that for a set \mathcal{X} that satisfies $\mathcal{X} \subseteq \mathbb{F}_q^n$, $|\mathcal{X}| = q^k L$, $k \in \mathbb{N}$, a (ρ, \mathcal{X}) -partial covering code $\mathcal{C}(k, n)$ has to satisfy

$$\log_q(L) \leq \log_q(V_q(n, \rho)). \quad (67)$$

This is easy to show because having q^k codewords directly means that the maximum number of points they can jointly ρn -cover is equal to $q^k V_q(n, \rho)$. This in turn implies that

$$Lq^k \leq V_q(n, \rho)q^k \quad (68)$$

which yields (67) after taking the logarithm on both sides of the inequality.

Now letting the above \mathcal{X} be the \mathcal{X} found in Theorem 1, we note that if $|\mathcal{X}| = Lq^k$ then $\mathcal{X} = \mathcal{X}_F$. Then by substituting $N = n$, $K = n - k$, we see that $\log_q(L) \leq \log_q(V_q(N, \rho))$. Since $q^{NH_q(\rho) - o(N)} \leq V_q(N, \rho) \leq q^{NH_q(\rho)}$, we can conclude that $\log_q(L) \leq NH_q(\rho)$ and thus that $H_q^{-1}(\frac{\log_q(L)}{N}) \leq \rho$, which concludes the proof. \square

APPENDIX B
PROOF OF ACHIEVABILITY IN THEOREM 2

Directly from [64], we know that there exists at least one ρ -covering code $\mathcal{C}_X(k, n)$ that satisfies

$$n - k \geq \log_q(V_q(n, \rho)) - 2 \log_2(n) + \log_q(n) - O(1). \quad (69)$$

Then applying Theorem 1 with $\mathbf{D} = \mathbf{H}_C$, $N = n$, $K = n - k$ and $\mathcal{X} = \mathbb{F}^n$, allows us to conclude that there exists a feasible scheme for the distributed computing problem, with computation cost $\gamma = \rho$, that satisfies

$$\begin{aligned} K/N &\geq \log_q(V_q(N, \rho))/N \\ &\quad - 2 \log_2(N)/N + \log_q(N)/N - O(1)/N. \end{aligned}$$

Combining this with the fact that $q^{NH_q(\rho) - o(N)} \leq V_q(N, \rho) \leq q^{NH_q(\rho)}$, yields

$$K/N \geq H_q(\rho) - \epsilon(N) \quad (70)$$

which tells us that $\rho \leq H_q^{-1}(K/N + \epsilon(N))$, which in turn proves the result in the limit of large N . \square

APPENDIX C
PROOF OF COROLLARY 1

We first start with the following lemma which proves the existence of a (ρ, \mathcal{X}) -partial covering linear code \mathcal{C} , for a properly-sized set $\mathcal{X} \subseteq \mathbb{F}^n$ that encloses $\mathcal{B}_q(0, \rho)$. Before proceeding with the lemma, we note that the lemma is an outcome of involving a linear greedy algorithm. Let us also briefly recall from Theorem 2 and its proof in Appendix A, that $\log_q(L) \leq \log_q(V_q(n, \rho))$.

Lemma 1. Let $\mathcal{X} \subseteq \mathbb{F}_q^n$ be a set of size $|\mathcal{X}| = L'q^k$ that satisfies $\mathcal{X} \supseteq \mathcal{B}_q(0, \rho)$. Then as long as

$$\log_q(L') \geq \log_q(V_q(n, \rho)) \quad (71)$$

$$- 2 \log_2(n) + \log_q(n) - O(1) \quad (72)$$

there exists a (ρ, \mathcal{X}) -partial covering code.

Proof. The proof is found in Appendix D. \square

With this lemma in place, let us define $\mathcal{A}_m \triangleq \{\mathcal{X} \subseteq \mathbb{F}^n \mid |\mathcal{X}| = m, \mathcal{X} \supseteq \mathcal{B}_q(0, \rho)\}$ to be the family of all subsets of \mathbb{F}^n which have cardinality m and which enclose $\mathcal{B}_q(0, \rho)$. Consider the following algorithm.

- 1) Assign $m = Lq^{N-K}$.
- 2) For each \mathcal{X} in \mathcal{A}_m , find a (ρ, \mathcal{X}) -partial covering code \mathcal{C}_X via (the algorithm corresponding to) Lemma 1.
- 3) For each \mathcal{X} in \mathcal{A}_m , set $\mathbf{D} = \mathbf{H}_{\mathcal{C}_X}$, and create $\mathcal{X}_{\mathbf{F}, \mathbf{D}} = \{\mathbf{x} \in \mathbb{F}^N \mid \mathbf{D}\mathbf{x} = \mathbf{F}(:, \ell), \text{ for some } \ell \in [L]\}$.
- 4) If there exists an \mathcal{X} in \mathcal{A}_m , for which $\mathcal{X} \subseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}}$, then output this \mathcal{X} and its corresponding $\mathbf{D} = \mathbf{H}_{\mathcal{C}_X}$ from the above step.
- 5) If there exists no \mathcal{X} in \mathcal{A}_m for which $\mathcal{X} \subseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}}$, then increase m by one and go back to step 2.

Let us continue now by supposing that the scheme terminates, outputting \mathbf{D} and \mathcal{X} at the fourth step, before m reaches $m = q^N$. Lemma 1, which guarantees (cf. (72)) that

$$\log_q(|\mathcal{X}|q^{-k}) \geq \log_q(V_q(n, \rho)) \quad (73)$$

$$- 2 \log_2(n) + \log_q(n) - O(1) \quad (74)$$

also guarantees that

$$\frac{\log_q(|\mathcal{X}|) - (N - K)}{N} \geq \log_q(V_q(N, \rho))/N \quad (75)$$

$$- 2 \log_2(N)/N \quad (76)$$

$$+ \log_q(N)/N - O(1)/N \quad (77)$$

where this last inequality holds after setting $N = n$, $K = n - k$, and after we divide both sides of (73) by N , and then apply Theorem 1 after recalling that \mathcal{X} is indeed ρn -covered by $\mathcal{C}_\mathbf{D}$. Applying that $q^{NH_q(\rho) - o(N)} \leq V_q(N, \rho) \leq q^{NH_q(\rho)}$ into (75), gives

$$\left(\frac{K}{N} - 1 + \frac{\log_q(|\mathcal{X}|)}{N}\right) \geq H_q(\rho) - \epsilon(N) \quad (78)$$

telling us that the algorithm yields a scheme with computation cost¹⁸

$$\gamma = \rho \leq H_q^{-1}(K/N - 1 + \log_q(|\mathcal{X}|/N) + \epsilon(N)) \quad (79)$$

¹⁸Here it is worth elaborating on a fine point regarding our metric. As the reader may recall, γ describes the fraction of *active* (non-idle) servers that compute any subfunction. Then the observant reader may wonder if our proposed scheme indeed activates all existing servers. This corresponds to having a scheme with an \mathbf{E} matrix that has no all-zero rows. In the (rare) degenerate scenario where a row of \mathbf{E} may contain only zeros, then our derived computation cost γ would — by definition — have to be recalculated (to account for having idle servers) and would be higher than stated here. To account for this degenerate case, we add a small step in our algorithm which reduces the recorded computation cost by guaranteeing that all servers are active. This step simply says that if a row in \mathbf{E} contains only zeros, then this row is substituted by an arbitrary non-zero row (let's say, the first row) of \mathbf{E} , except that, if that (first) row contains a non-zero element in the position $\ell_{\max} = \arg \max \omega(\mathbf{E}(:, \ell))$, then this element is substituted by a zero. Then the two servers (the first server and the previously idle server) will split their communication load, except that the server corresponding to the originally all-zero row, will not send any linear combination that involves $w_{\ell_{\max}}$. This small modification guarantees that whatever γ we declare here as being achievable, is indeed achievable even in degenerate scenarios. Finally, this degenerate scenario does not affect the algebraic converse.




Fig. 5. Map of lemmas, theorems and appendices.

which matches the stated result in the regime of large N . Note that when $\mathcal{X} = \mathcal{X}_{\mathbb{F}, \mathbb{D}}$, then naturally $|\mathcal{X}| = Lq^{N-K}$ which, directly from (79), yields $\gamma = \rho = H_q^{-1}(\log_q(L)/N + \epsilon(N))$. At the other extreme, when the algorithm terminates at the very end when $\mathcal{X} = \mathbb{F}^N$, then the corresponding code will be the standard ρ -covering code (see Appendix B), and the computation cost will correspond to $\gamma = H_q^{-1}(K/N)$. \square

APPENDIX D PROOF OF LEMMA 1

We here start by employing the recursive construction approach of Cohen and Frankl in [64]. This recursive approach builds an $(n, j+1)$ code C_{j+1} from a previous (n, j) code C_j , by carefully adding a vector \mathbf{x} on the basis of C_j , so that now the new basis span is bigger. Our aim will be to recursively construct ever bigger codes that cover an ever increasing portion of our set \mathcal{X} .

Let us start by setting $C_0 = \{\mathbf{0}\}$. Let us then make the assumption that the aforementioned integer L' in Lemma 1, takes the form

$$L' = q^{n-k'} \quad (80)$$

for some real $k' \geq k$. Let $Q(C)$ denote the set of points in \mathcal{X} that are not ρn -covered by C , and let

$$q(C) \triangleq \frac{|Q(C)|}{q^{n+k-k'}} \quad (81)$$

where naturally

$$|Q(C_0)| = q^{n+k-k'} - V_q(n, \rho) \quad (82)$$

and

$$q(C_0) = 1 - V_q(n, \rho)q^{-(n+k-k')}. \quad (83)$$

To proceed, we need the following lemma from [64].

Lemma 2 ([64]). Let $\mathcal{Y} \subseteq \mathbb{F}^n$, $\mathcal{Z} \subseteq \mathbb{F}^n$, and consider $\mathcal{Y} + \mathbf{x} = \{\mathbf{y} + \mathbf{x} : \mathbf{y} \in \mathcal{Y}\}$ for some $\mathbf{x} \in \mathbb{F}^n$. Then

$$\mathbb{E}(|(\mathcal{Y} + \mathbf{x}) \cap \mathcal{Z}|) = q^{-n}|\mathcal{Y}||\mathcal{Z}| \quad (84)$$

where the average is taken, with uniform probability, over all $\mathbf{x} \in \mathbb{F}^n$.

Now, we develop the proof in two parts.

- 1) **Binary Case:** The proof for $q = 2$ where $k = k'$ (corresponding to the singular case of maximal $L = 2^K$) has been presented in [70] and [64] in two different ways. We will modify the latter approach to establish our claim for any $k' \geq k$ (which will allow us to also handle L values that are smaller than 2^K). First let us easily deduce from Lemma 2 that there exists an $\mathbf{x} \in \mathbb{F}^n$ for which $|(\mathcal{Y} + \mathbf{x}) \cap \mathcal{Z}| \leq \frac{|\mathcal{Y}||\mathcal{Z}|}{q^n}$. Now let us set $\mathcal{Y} = \mathcal{Z} = Q(C_j)$, and let us append a vector \mathbf{x} to the generator matrix of C_j to create C_{j+1} , where \mathbf{x} is chosen to minimize $|Q(C_{j+1})|$. Now we can directly verify that

$$|Q(C_{j+1})| = |Q(C_j) \cap Q(C_j + \mathbf{x})| \quad (85)$$

$$= |\mathcal{Q}(C_j) \cap (\mathcal{Q}(C_j) + \mathbf{x})| \quad (86)$$

$$\leq |\mathcal{Q}(C_j)|^2 / 2^n \quad (87)$$

which implies that

$$q(C_{j+1}) \leq q(C_j) 2^{k-k'} \leq q(C_j)^2 \quad (88)$$

where the latter inequality holds because $k' \geq k$. Combining (83) and (88), gives

$$q(C_k) \leq q(C_0)^{2^k} \leq (1 - V_q(n, \rho) 2^{-(n-k'+k)})^{2^k} \quad (89)$$

where the latter inequality again holds due to the fact that $k' \geq k$. Now let us continue this recursion until k is such that

$$2^k = \lceil (n - k' + k) 2^{(n-k'+k)} \ln(2) / V_2(n, \rho) \rceil \quad (90)$$

in which case — given that $(1 - \frac{1}{x})^x \leq e^{-1}$, $\forall x \geq 1$ — we get that

$$q(C_k) < 2^{-(n+k-k')} \quad (91)$$

which automatically yields that $Q(C_k) = 0$. This, again with the choice of k in (90), tells us that for a set \mathcal{X} that satisfies $\mathcal{B}_q(0, \rho) \subseteq \mathcal{X} \subseteq \mathbb{F}_q^n$, $|\mathcal{X}| = Lq^k$, then indeed there exists a (ρ, \mathcal{X}) -partial covering code $C(n, k)$ satisfying

$$0 \leq \log_q(L/V_q(n, \rho)) + 2 \log_2(\log_q(|\mathcal{X}|)) \quad (92)$$

$$- \log_q(\log_q(|\mathcal{X}|)) + O(1). \quad (93)$$

This conclusion can be considered as a tighter version of Lemma 1. After a few very basic algebraic manipulations we get the proof of Lemma 1, for the binary case of $q = 2$.

2) **Non-Binary Case:** Considering first an arbitrary $\mathcal{Z} \subset \mathbb{F}^n$, we have that

$$\mathbb{E}(1 - (q^{-n+k'-k} |(\mathcal{Z} + \mathbf{x}) \cup \mathcal{Z}|)) \quad (94)$$

$$= \mathbb{E}(1 - q^{-n+k'-k} (|\mathcal{Z} + \mathbf{x}| + |\mathcal{Z}|)) \quad (95)$$

$$- |(\mathcal{Z} + \mathbf{x}) \cap \mathcal{Z}|) \quad (96)$$

$$= \mathbb{E}(1 - 2q^{-n+k'-k} |\mathcal{Z}| + q^{-n+k'-k} |(\mathcal{Z} + \mathbf{x}) \cap \mathcal{Z}|) \quad (97)$$

$$\stackrel{(a)}{=} 1 - 2q^{-n+k'-k} |\mathcal{Z}| + q^{-2n+k'-k} |\mathcal{Z}|^2 \quad (98)$$

$$\stackrel{(b)}{\leq} 1 - 2q^{-(n-k'+k)} |\mathcal{Z}| + q^{-2(n-k'+k)} |\mathcal{Z}|^2 \quad (99)$$

$$= (1 - \frac{|\mathcal{Z}|}{q^{(n-k'+k)}})^2 \quad (100)$$

where (a) is directly from Lemma 2, and where (b) holds since $k' \geq k$. Similarly to the binary case, we begin with $C_0 = \{\mathbf{0}\}$, and again recursively extend as

$$C_{j+1} = \langle C_j; \mathbf{x} \rangle \quad (101)$$

where \mathbf{x} is chosen so that $|\mathcal{Z}|$ is maximized. We do so, after again setting $\mathcal{Z} = Q(C_j)$.

At this point, from (100) we have that

$$q(C_{j+1}) \leq q(C_j)^2. \quad (102)$$

We now consider the following lemma from [64].

Lemma 3. ([64, Lemma 2]) For any fixed $\mathcal{Z} \subseteq \mathcal{X} \subset \mathbb{F}^n$ where $|\mathcal{Z}| q^{-(n-k'+k)} = \epsilon < (q(n - k' + k))^{-1}$, then

$$\mathbb{E}_{\mathbf{x} \in \mathbb{F}^n} (1 - q^{-(n-k'+k)} | \cup_{\alpha \in \mathbb{F}_q} \mathcal{Z} + \alpha \mathbf{x} |) \leq \quad (103)$$

$$(1 - \epsilon)^{q^{(1-(2(n-k'+k))^{-1})}}. \quad (104)$$

Continuing from $\mathcal{Z} = \mathcal{X} \cap (\cup_{\mathbf{c} \in C_j} \mathcal{B}_q(\mathbf{c}, \rho))$, where

$$|\mathcal{Z}| < \frac{1}{n} q^{(n-k'+k-1)}, \quad q(C_{j+1}) \leq q(C_j)^{q^{(1-(2(n-k'+k))^{-1})}}$$

we have that

$$q(C_{j+1}) \quad (105)$$

$$\leq (1 - q^{n-k'+k} V_q(n, \rho))^{(q^{(1-(2(n-k'+k))^{-1})})^j} \quad (106)$$

$$\leq (1 - q^{n+k-k'} V_q(n, \rho))^{e^{-0.5} q^j} \quad (107)$$

since $(1 - (2(n - k' + k))^{-1}) \geq (1 - (2(n - k' + k))^{-1})^{n-k'+k-1} \geq e^{-0.5}$. For

$$j_1 \triangleq \arg \min_j \{q(C_j) \leq 1 - (q(n + k - k'))^{-1}\} \quad (108)$$

we see that

$$j_1 \leq n - \log_q(q^{k'-k} V_q(n, \rho)) \quad (109)$$

$$- \log_q(n + k - k') + O(1) \quad (110)$$

where the inequality holds by first observing that Lemma 3 yields

$$1 - (q(n - k' + k))^{-1} \quad (111)$$

$$\leq q(C_j) \quad (112)$$

$$\leq (1 - q^{(n-k'+k)} V_q(n, \rho))^{q^{j_1-1} e^{-1/2}} \quad (113)$$

and then by comparing the upper and lower bounds in (113).

We now have an (n, j_1) code C and we have (102). We are now looking for the minimum number j_2 of generators \mathbf{x} that have to be appended to the generator of C in order to get a $(n, j_1 + j_2)$ code with $q(C_{j_1+j_2}) \leq q^{-(n-k'+k)}$. We note that $q(C_{j_1}) \leq 1 - (q(n - k' + k))^{-1}$, so by (113) we only need to ensure that $(1 - (q(n - k' + k))^{-1})^{2^{j_2}} \leq q^{-(n-k'+k)}$, which can be achieved by using

$$j_2 = 2 \log_2(n - k' + k) + O(1). \quad (114)$$

Hence for $k = j_1 + j_2$, there indeed exist (n, k) codes with normalized covering radius no bigger than ρ . Applying (109), (114), (80), and the fact that $|\mathcal{X}| = Lq^k$, proves (72) and thus proves Lemma 1. \square

APPENDIX E PROOF OF THEOREM 3

We quickly note that the converse (lower bound on γ) holds directly from the converse arguments in Theorem 2.

Let us start with the following definition.

Definition 2. Let $\rho \in (0, 1 - \frac{1}{q}]$, and let $\tau \in (0, 1]$. A code $C \subseteq \mathbb{F}^n$ is said to be a (ρ, τ) -partial covering code if there exists a set $\mathcal{X} \subseteq \mathbb{F}^n$, with $\frac{1}{n} \log_q(|\mathcal{X}|) = 1 - \tau$, that is ρ -covered by C .

We now present a theorem that extends the famous Theorem of Blinovskii in [61], which proved that almost all linear codes satisfy the sphere-covering bound. We recall that $C_{k,n}$ denotes

the ensemble of all linear codes generated by all possible $k \times n$ matrices in $\mathbb{F}^{k \times n}$.

Theorem 5. *Let $\rho \in (0, 1 - \frac{1}{q}]$. Then there exists an infinite sequence k_n that satisfies*

$$\frac{k_n}{n} \leq 1 - \tau - H_q(\rho) + O(n^{-1} \log_q(n)) \quad (115)$$

for $\tau \in [0, 1 - H_q(\rho) - \frac{k}{n}]$ so that the fraction of codes $C_n \in C_{k_n, n}$ that are (ρ, τ) -partial covering, tends to 1 as n grows to infinity. Thus in the limit of large n , almost all codes of rate less than $1 - \tau - H(\rho)$ will be (ρ, τ) -partial covering.

Proof. The proof can be found in Appendix F. \square

Now let us design such covering codes. In the following we will consider the set of codes in $C_{k_n, n}$ that are (ρ, τ) -partial covering, for the claimed sequence k_n of Theorem 5, and for some real τ . We will also consider $g(n)$ to be the fraction of such (ρ, τ) -partial covering codes among all codes in $C_{k_n, n}$. The scheme design is defined by the following steps.

- 1) Assign $m = L$.
- 2) Set $\tau = \frac{K - \log_q(m)}{N}$.
- 3) Noticing that the value

$$m_n \triangleq g(n)q^{k_n n} \quad (116)$$

serves as a lower bound on the number of (ρ, τ) -partial covering codes in the ensemble $C_{k_n, n}$, we now create $\mathcal{B} \triangleq \{C_1, C_2, \dots, C_{m_n}\}$ to be the set of the first m_n such codes. Now let

$$\mathbf{D}_n \triangleq \begin{bmatrix} \mathbf{H}_{C_1} & & & \\ & \mathbf{H}_{C_2} & & \\ & & \ddots & \\ & & & \mathbf{H}_{C_{m_n}} \end{bmatrix} \quad (117)$$

and accordingly set $K = m_n(n - k_n)$, and $N = m_n n$. Now design $C_{\mathcal{W}_n} = [C_1, C_2, \dots, C_{m_n}]$, and then create the set

$$\mathcal{X}_{\mathbf{F}, \mathbf{D}} \triangleq \{\mathbf{x} \in \mathbb{F}^N \mid \mathbf{D}\mathbf{x} = \mathbf{F}(:, \ell), \text{ for some } \ell \in [L]\}. \quad (118)$$

Then create the set

$$\mathcal{X} \triangleq \{\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{m_n}] \mid \mathbf{x}_i \in \mathcal{X}_i\} \quad (119)$$

where $\mathcal{X}_i, i \in [m_n]$, is the set of all n -length vectors that are ρn -covered by C_i . Then note that

$$|\mathcal{X}_i| \geq q^{n(1-\tau)}, \forall i \in [m_n] \quad (120)$$

because of Definition 2. We now note that for any $\mathbf{x} \in \mathcal{X}$, it is the case that

$$\mathbf{d}(\mathbf{x}, C)/N = \sum_{i=1}^{m_n} \mathbf{d}(\mathbf{x}_i, C_i)/N \leq \sum_{i=1}^{m_n} \frac{\rho n}{m_n n} \quad (121)$$

$$= \sum_{i=1}^{m_n} \rho \frac{1}{m_n} = \rho \quad (122)$$

which means that $C_{\mathbf{D}_n}$ is also a (ρ, \mathcal{X}) -partial covering code. Now if $\mathcal{X} \not\supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}}$, then m has to be increased by one, and the procedure starts again from Step 2.

- 4) Let us define $k'_n \triangleq n - k_n$. From (115), we know that

$$\frac{k'_n}{n} \geq \tau + H_q(\rho) - O(n^{-1} \log_q(n)). \quad (123)$$

We now see that $R \triangleq \frac{K}{N} = \frac{k'_n}{n}$ since $K = k'_n m_n, N = n m_n$. Thus, directly from the above, we have that

$$K/N = R = H_q(\rho) + \tau - \epsilon(N). \quad (124)$$

We note that as n (and thus N) goes to infinity, the term $O(n^{-1} \log_q(n))$ vanishes, and thus from the above we have that

$$\rho = H_q^{-1}\left(\frac{\log_q(m)}{N} + \epsilon(N)\right). \quad (125)$$

We also have that

$$\frac{\omega(\mathbf{D}_n)}{K} \stackrel{(a)}{\leq} \frac{m_n n k'_n}{m_n k'_n} = n \quad (126)$$

where (a) holds since $\omega(\mathbf{D}_n) = m_n k_n n$ is the maximum number of nonzero elements that \mathbf{D} can have, due to the block-diagonal design.

After taking the logarithm on both sides of the above, and since $N = m_n n$ and $k_n = (1 - R)n$, and after considering (116), we have that

$$\log_q(n) + n^2(1 - R) + \log_q(g(n)) = \log_q(N) \quad (127)$$

and thus we have that $n^2(1 - R) \leq \log_q(N)$ and $n \leq \sqrt{\frac{\log_q(N)}{(1-R)}}$. Combining this with (126) and Theorem 1, we have that

$$\Delta \leq \sqrt{\frac{\log_q(N)}{(1-R)}} \quad (128)$$

where, as mentioned before, R is constant.

We can also see that the above design terminates, since reaching $m = q^K$ implies that $\tau = 0$. Then we will have $\mathcal{X}_i = \mathbb{F}^n$ since $|\mathcal{X}_i| = q^n$ from Definition 2. Therefore from (119), we will have that $\mathcal{X} = \mathbb{F}^N = \mathbb{F}^{m_n n}$, which means that $C_{\mathcal{D}_n}(N, N - K)$ is a ρ -covering code, and that $\mathcal{X} \supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}}$, and thus the scheme would terminate at Step 4 with $\gamma = \rho = H_q^{-1}(\frac{K}{N} + \epsilon(N))$ from (125), and with communication cost as shown in (128). \square

APPENDIX F PROOF OF THEOREM 5

Before offering the formal proof, we provide a quick sketch of the proof to help the reader place the different steps in context.

First we consider the ensemble¹⁹ of codes $C_{k^*, n}$, and we prove that with a consistent enumeration of codewords, each nonzero point in \mathbb{F}^m has the same chance to be a codeword of a certain index, as we move across the code ensemble.

Second, we pick a code $C \in C_{k^*, n}$ at random, and fix it. Then, based on this code, and for a specific choice of τ (to be described later on), we introduce a random so-called ‘covered set’ \mathcal{X}_C of size $2^{n(1-\tau)}$ that includes code C .

¹⁹The details about the choice of k^* will be described later on.

Then we will see that every point in $\mathbb{F}^n \setminus \mathcal{B}(\mathbf{0}, \rho)$ has an equal probability — as we go through the choices of $C \in C_{k^*, n}$ — of belonging to this subset. To analyze the ρ -coverage of points inside \mathcal{X}_C , we derive $\mathbb{P}(\mathbf{c}_i = \mathbf{x} | \mathbf{x} \in \mathcal{X}_C)$, where \mathbf{c}_i describes the codeword indexed by a fixed i , as we move across the codes (and the corresponding generator matrices) in the ensemble.

Toward showing that \mathcal{X}_C is covered by C , we first note that $\mathcal{B}(\mathbf{0}, \rho)$ is covered since $\mathbf{0} \in C$. To prove that the remaining part, $\mathcal{X}_C \setminus \mathcal{B}(\mathbf{0}, \rho)$, is also covered by C , we prove that if codes in the ensemble are sufficiently large, then there is, for almost all codes C , a large number (polynomial in n) of codewords that covers each specific point in $\mathbf{x} \in \mathcal{X}_C$. With this in place, we will be able to conclude that almost all codes come close to being (ρ, τ) -partial covering.

Finally, we utilize a linear greedy algorithm and successive appending of a very small number of $\lfloor \log_q n(1 - \tau) \rfloor$ carefully selected vectors (cf. Lemma 2) to each of these almost (ρ, τ) -partial covering codes, to render them fully (ρ, τ) -partial covering codes.

We proceed with the formal proof.

Let $k^* \triangleq k - \lfloor \log_q(n(1 - \tau)) \rfloor$, $k, n \in \mathbb{N}, 0 \leq \tau \leq 1$ and let $C_{k^*, n}$ be the ensemble of codes generated by $k^* \times n$ generator matrices whose elements are chosen randomly and independently with probability $\frac{1}{q}$ from \mathbb{F}_q . Naturally, any fixed non-zero linear combination of rows of the generator matrix, will generate — as we move across the ensemble of generator matrices — all possible q^n vectors in \mathbb{F}^n . The zero codeword corresponds to the void linear combination of rows, and is present in all generated codes. Also let us assume a consistent enumeration of the codewords, in the sense that a word's index is defined by the linear combination of rows of the generator matrix, that generate that codeword, in each code. For example, the word indexed by 5, will vary in value across the different codes, but it will always be defined as the output of a specific (the fifth) linear combination of the corresponding generator matrix. The first codeword in all codes will be the zero word. We proceed with the following lemma.

Lemma 4. For any fixed $i \in [2 : 2^{k^*}]$, and for any fixed $\mathbf{x} \in \mathbb{F}^n$, then

$$\mathbb{P}(\mathbf{c}_i = \mathbf{x}) = q^{-n} \quad (129)$$

where the probability is over all codes $C \in C_{k^*, n}$.

Proof. The proof is presented in Appendix H-A. \square

Let us set $\tau \in [0, 1 - H_q(\rho) - k^*/n]$, and let us note that for sufficiently²⁰ large n , we can guarantee that

$$q^{n(1-\tau)} \geq V_q(n, \rho) + q^{k^*}. \quad (130)$$

Let us now go over the ensemble of codes $C \in C_{k^*, n}$, and for each code, let us create the covered set \mathcal{X}_C such that

$$|\mathcal{X}_C| = q^{n(1-\tau)} \quad (131)$$

$$C \cup \mathcal{B}(\mathbf{0}, \rho) \subseteq \mathcal{X}_C \subseteq \mathbb{F}^n. \quad (132)$$

We can see that (130) is a necessary condition for the above to happen. The procedure for designing \mathcal{X}_C , simply

²⁰We quickly remind the reader that our results here will hold for sufficiently large n .

starts by taking the union $C \cup \mathcal{B}(\mathbf{0}, \rho)$, and then proceeds by appending on this union, a sufficiently large number of vectors, chosen uniformly and independently at random from $\mathbb{F}^n \setminus C \cup \mathcal{B}(\mathbf{0}, \rho)$. The following lemma simply says that every point $\mathbf{x} \in \mathbb{F}^n \setminus \mathcal{B}(\mathbf{0}, \rho)$ has an equal probability — as we go through the choices of $C \in C_{k^*, n}$ — of belonging to this subset \mathcal{X}_C .

Lemma 5. For any fixed $\mathbf{x} \in \mathbb{F}^n$, then

$$\mathbb{P}(\mathbf{x} \in \mathcal{X}_C) = \begin{cases} 1 & \omega(\mathbf{x}) \leq \rho n \\ \frac{q^{n(1-\tau)} - V(\rho, n)}{q^n - V(\rho, n)} & \omega(\mathbf{x}) > \rho n. \end{cases} \quad (133)$$

Proof. The proof is presented in Appendix H-B. \square

With the above lemma in place, we will now calculate the following conditional probability. The following asks us to first pick and fix a vector $\mathbf{x} \in \mathbb{F}^n$, and then pick an index $i \in [1 : 2^{k^*}]$. Recall — from the above discussion on the consistent enumeration of codewords — that this index will define a codeword \mathbf{c}_i , which changes as we go across all the codes C in the ensemble $C_{k^*, n}$. The following conditional probability is again calculated over the code ensemble.

Lemma 6. Pick any vector $\mathbf{x} \in \mathbb{F}^n$ and any index $i \in [1 : 2^{k^*}]$. Then

$$\begin{aligned} \mathbb{P}(\mathbf{c}_i = \mathbf{x} | \mathbf{x} \in \mathcal{X}_C) & \quad (134) \\ & = \begin{cases} 0 & i = 1, \mathbf{x} \neq \mathbf{0} \\ 1 & i = 1, \mathbf{x} = \mathbf{0} \\ q^{-(n)} & i \in [2 : K^*], \mathbf{x} \neq \mathbf{0}, \omega(\mathbf{x}) \leq \rho n \\ q^{-n(1-\tau)} \zeta(n) & i \in [2 : K^*] \text{ and } \omega(\mathbf{x}) > \rho n \end{cases} \quad (135) \end{aligned}$$

where the term $\zeta(n)$ converges to 1 as n approaches to infinity.

Proof. The proof is presented in Appendix H-C. \square

Let us now discuss the coverage of any vector $\mathbf{x} \in \mathcal{X}_C$, as we go along the ensemble $C \in C_{k^*, n}$. First of all, it is clear that any $\mathbf{x} \in \mathcal{B}(\mathbf{0}, \rho)$ is both ρ -covered by the code C (because $\mathbf{0} \in C$) as well as is included in \mathcal{X}_C (because $\mathbf{x} \in \mathcal{B}(\mathbf{0}, \rho) \subset \mathcal{X}_C$). Thus for each code $C \in C_{k^*, n}$, for the purposes of the current proof, we can focus on the set

$$\mathcal{X}'_C \triangleq \mathcal{X}_C \setminus \mathcal{B}(\mathbf{0}, \rho). \quad (136)$$

For every $\mathbf{x} \in \mathcal{X}'_C$, let us define the random variable $\eta_{\mathbf{x}, i}$ which takes the value 1 if \mathbf{c}_i ρn -covers \mathbf{x} , and which takes the value 0 otherwise. Thus

$$\eta_{\mathbf{x}} \triangleq \sum_{i=1}^{2^{k^*}} \eta_{\mathbf{x}, i} \quad (137)$$

describes the number of codewords that cover $\mathbf{x} \in \mathcal{X}'_C$. For any fixed $\mathbf{x} \in \mathbb{F}^n$, the following lemma describes the conditional average $\eta_{\mathbf{x}}$, where again the average is taken over the code ensemble.

Lemma 7. For any fixed $\mathbf{x} \in \mathbb{F}^n$, then

$$\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C) = |\{0\} \cap \mathcal{B}(\mathbf{x}, \rho)| \times 1 \quad (138)$$

$$+ (q^{k^*} - 1) [|\mathcal{B}(\mathbf{0}, \rho) \setminus \{0\}|] \quad (139)$$

$$\cap \mathcal{B}(\mathbf{x}, \rho) | q^{-n} \quad (140)$$

$$+ |\mathcal{B}(\mathbf{x}, \rho) \setminus \mathcal{B}(0, \rho)| q^{-n(1-\tau)} \zeta(n) \quad (141)$$

where again $\zeta(n) \rightarrow 1$ as n increases.

Proof. The proof is in Appendix H-D, and it involves an extension of Blinovskii's Theorem [61], from evaluating $\mathbb{E}(\eta_{\mathbf{x}})$ to evaluating the conditional $\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)$. \square

Before proceeding, we need the following lemma which is an extension of a related lemma found in [61]. The following considers as before the set $\mathcal{X}'_C \triangleq \mathcal{X}_C \setminus \mathcal{B}(0, \rho)$, and considers again the variance and expectation, over the aforementioned code ensemble.

Lemma 8. For any fixed $\mathbf{x} \in \mathbb{F}^n$, then

$$\frac{\text{Var}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)}{\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)^2} \leq 1. \quad (142)$$

Proof. The proof is presented in Appendix H-E. \square

Combining (142) with Chebyshev's inequality, gives

$$\mathbb{P}(|\eta_{\mathbf{x}} - \mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)| \quad (143)$$

$$> q^{\epsilon+1} \sqrt{\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)} | \mathbf{x} \in \mathcal{X}'_C) \quad (144)$$

$$< \frac{\text{Var}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)}{q^{2\epsilon+2} \mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)} \leq q^{-2\epsilon}. \quad (145)$$

Our aim is to show that, for any $\mathbf{x} \in \mathcal{X}'_C$, — under some conditions on k^* and ϵ — $\eta_{\mathbf{x}}$ will be, with high probability, bigger than 0 which in turn implies that any $\mathbf{x} \in \mathcal{X}'_C$ will, with high probability, be covered by C . To see this, we continue from (145). We first see from (141) that $\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C) > 0$. Hence whenever we have $\eta_{\mathbf{x}} > \mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)$, we also have that $\eta_{\mathbf{x}} > 0$. Let us now focus on the remaining scenario where $\eta_{\mathbf{x}} \leq \mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)$. In this case, from (145), we have that

$$\mathbb{P}(\eta_{\mathbf{x}} \geq \mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C) \quad (146)$$

$$- q^{\epsilon+1} \sqrt{\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)} | \mathbf{x} \in \mathcal{X}'_C) \quad (147)$$

$$\geq 1 - q^{-2\epsilon}. \quad (148)$$

We also have that

$$\mathbb{P}(\eta_{\mathbf{x}} > 0 | \mathbf{x} \in \mathcal{X}'_C) \quad (149)$$

$$\geq \mathbb{P}(\eta_{\mathbf{x}} \geq \mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C) \quad (150)$$

$$- q^{\epsilon+1} \sqrt{\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)} | \mathbf{x} \in \mathcal{X}'_C) \quad (151)$$

under the assumption that

$$\beta(\epsilon) \triangleq \mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C) - q^{\epsilon+1} \sqrt{\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)} > 0. \quad (152)$$

This assumption will be guaranteed — as we will see later on — by a proper choice of k^* and ϵ .

Now combining (148) with (151), we will show that

$$\mathbb{P}(\eta_{\mathbf{x}} \geq \mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C) \quad (153)$$

$$- q^{\epsilon+1} \sqrt{\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C)} | \mathbf{x} \in \mathcal{X}'_C) \rightarrow 1 \quad (154)$$

as n grows to infinity.

To guarantee that $\beta(\epsilon) > 0$, we must guarantee that

$$\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C) > q^{2\epsilon+2}. \quad (155)$$

To do this, given Lemma 7, we must prove that

$$|\{0\} \cap \mathcal{B}(\mathbf{x}, \rho)| \times 1 \quad (156)$$

$$+ (q^{k^*} - 1)[(\mathcal{B}(0, \rho) \setminus \{0\}) \cap \mathcal{B}(\mathbf{x}, \rho)] q^{-n} \quad (157)$$

$$+ |\mathcal{B}(\mathbf{x}, \rho) \setminus \mathcal{B}(0, \rho)| q^{-n(1-\tau)} \zeta(n) > q^{2\epsilon+2} \quad (158)$$

again for some properly chosen k^* and ϵ . The following applies toward this effort.

Lemma 9. For any $C \in \mathcal{C}_{k^*, n}$, any $\mathbf{x} \in \mathcal{X}'_C$, and any $\rho \in (0, \min\{1 - 1/q, \frac{\sqrt{5}-1}{2}\}]$, then

$$|\mathcal{B}(\mathbf{x}, \rho) \setminus \mathcal{B}(0, \rho)| > q^{nH_q(\rho) - o(n)}. \quad (159)$$

Proof. The proof is in Appendix I. \square

We now combine (141) and (159) to get

$$E(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C) > (q^{k^*} - 1) q^{nH_q(\rho) - o(n) - n(1-\tau)} \zeta(n) \quad (160)$$

and we also choose k^*, ϵ , to guarantee (cf. (155)) that the inequality

$$(q^{k^*} - 1) q^{nH_q(\rho) - o(n) - n(1-\tau)} \zeta(n) \geq q^{2\epsilon+2} \quad (161)$$

holds for large n . Thus with (160) and (161) in place — something that will indeed be validated by the end of the proof (cf. (185)) — we can guarantee (155). Thus we know that

$$\mathbb{P}(\eta_{\mathbf{x}} < n^\alpha | \mathbf{x} \in \mathcal{X}'_C) < q^{-2\epsilon}. \quad (162)$$

Following the approach in [61], we consider points in \mathcal{X}'_C that are called ‘partial-remote points’, which are the points that are ρn -covered by fewer than n^α , $\alpha > 1$ codewords. Now let $\mathcal{Q}_0(\mathcal{X}'_C) \subset \mathcal{X}'_C$ be the set of partial remote points in \mathcal{X}'_C , and let

$$q_0(\mathcal{X}'_C) \triangleq \frac{|\mathcal{Q}_0(\mathcal{X}'_C)|}{q^{n(1-\tau)} - V_q(n, \rho)}. \quad (163)$$

Now applying (162), gives

$$\sum_{\mathbf{x} \in \mathcal{X}'_C} \mathbb{P}(\eta_{\mathbf{x}} < n^\alpha | \mathbf{x} \in \mathcal{X}'_C) \leq (q^{n(1-\tau)} - V_q(n, \rho)) q^{-2\epsilon} \quad (164)$$

and thus we see that

$$\sum_{\mathbf{x} \in \mathcal{X}'_C} \mathbb{P}(\eta_{\mathbf{x}} < n^\alpha | \mathbf{x} \in \mathcal{X}'_C) \quad (165)$$

$$\stackrel{(a)}{=} \sum_{\mathbf{x} \in \mathcal{X}'_C} \mathbb{E}[\mathbb{1}(\eta_{\mathbf{x}} < n^\alpha | \mathbf{x} \in \mathcal{X}'_C)] \quad (166)$$

$$\stackrel{(b)}{=} \mathbb{E}[\sum_{\mathbf{x} \in \mathcal{X}'_C} \mathbb{1}(\eta_{\mathbf{x}} < n^\alpha | \mathbf{x} \in \mathcal{X}'_C)] \quad (167)$$

$$\stackrel{(c)}{=} \mathbb{E}[|\mathcal{Q}_0(\mathcal{X}'_C)|] \quad (168)$$

where now the expectation in (a) is over the codes in the ensemble $\mathcal{C}_{k^*, n}$, where (b) results from interchanging the expectation with the summation, and where (c) is by definition of \mathcal{Q}_0 .

Now combining (163), (164) and (168), we have

$$\mathbb{E}(q_0) \leq q^{-2\epsilon} \quad (169)$$

which bounds the average (over the code ensemble) number of partial-remote points in \mathcal{X}'_C . Then Markov's inequality directly tells us that

$$\mathbb{P}(q_0 > q^\epsilon \mathbb{E}(q_0)) < q^{-\epsilon} \quad (170)$$

which means that the expression $q_0 \leq q^\epsilon \mathbb{E}(q_0)$ holds for a proportion greater than $1 - q^{-\epsilon}$ of all codes.

Now, in the footsteps of [62], we apply a procedure that successively appends cosets to an initial code $C' \in \mathcal{C}_{k^*,n}$ that belongs in this above family of codes that indeed satisfies $q_0 \leq q^\epsilon \mathbb{E}(q_0)$. Let us quickly remember that the optimal successive appending linear greedy method resulting from Lemma 2 and enclosed in Appendix D, allowed us to prove (96)–(100) which yielded

$$q(C_{j+1}) \leq q(C_j)^2 \quad (171)$$

where $C_{j+1} = \langle C_j; \mathbf{x} \rangle$, and where $q(C_j)$ represented the number of remote points of the code C_j , normalized by $q^{n(1-\tau)} - V_q(n, \rho)$.

With the above in mind, let us now set this first initializing code C_0 to be equal to $C_0 = C'$, where C' is one of the aforementioned 'good' codes that satisfy

$$q_0 \leq q^\epsilon \mathbb{E}(q_0). \quad (172)$$

Then we will design $C_1 = \langle C_0; \mathbf{x} \rangle$ where \mathbf{x} is a guaranteed-to-exist vector (cf. (171)) that increases the span of C_0 .

Now we calculate the same quantity we calculated in (163), but we do so for \mathcal{X}'_{C_1} . In other words, we calculate

$$q_1(\mathcal{X}'_{C_1}) \triangleq \frac{|\mathcal{Q}_0(\mathcal{X}'_{C_1})|}{q^{n(1-\tau)} - V_q(n, \rho)} \quad (173)$$

where similar to before, now q_1 represents the average normalized remote points of the code C_1 with respect to its associated \mathcal{X}_{C_1} . We can now see that directly from (171), we have that

$$\mathbb{E}(q_1) \leq q_0^2 \quad (174)$$

where the above average is taken over all C_1 codes, meaning over all codes that can take the role of our aforementioned C_1 , going over all possible initializing codes $C_0 = C'$, and over all possible base-expanding vectors \mathbf{x} . Now we apply again Markov's inequality, this time over the expanded codes C_1 , to get

$$\mathbb{P}(q_1 > q^\lambda \mathbb{E}(q_1)) < q^{-\lambda}, \quad \lambda \in \mathbb{R} \quad (175)$$

which tells us — as before — that the proportion of codes C_1 that satisfy

$$q_1 \leq q^{\lambda-2\epsilon} \quad (176)$$

is at least $1 - q^{-\lambda}$. This proportion of codes that achieve (176), is over all generated C_1 that were built over all 'good' $C_0 = C'$ that already satisfied (172). Thus we now know (cf. (170)) that the proportion of codes C_1 — among all codes in the entire ensemble $\mathcal{C}_{k^*,n}$ — that satisfy (176), is equal to $(1 - q^{-\epsilon})(1 - q^{-\lambda})$.

We now go from our step 1, to an arbitrary step i , and following the same logic as before, we conclude that the

proportion of codes C_i — where this proportion is among all codes in the entire ensemble $\mathcal{C}_{k^*,n}$ — that satisfy

$$q_i < q^{-2^i(\epsilon-\lambda)-\lambda} \quad (177)$$

is equal to $(1 - q^{-\epsilon})(1 - q^{-\lambda})^i$.

Now let us go to some step m which will allow us to terminate. We explain when will this termination happen. Consider, for this step $i = m$, as before, the quantity

$$q_m(\mathcal{X}'_{C_m}) \triangleq \frac{|\mathcal{Q}_0(\mathcal{X}'_{C_m})|}{q^{n(1-\tau)} - V_q(n, \rho)} \quad (178)$$

where similar to before, now q_m represents the average normalized remote points of the code C_m with respect to its associated \mathcal{X}_{C_m} . We want the corresponding \mathcal{Q}_m (cf. (163)) to be empty. This will be guaranteed when m is such that

$$q_m < q^{-n(1-\tau)} \quad (179)$$

where the above guarantee can be provided given that the cardinality of a set is a non-negative integer.

This will be achieved by setting $m = \lceil \log_2(n(1-\tau)) \rceil$ and $\lambda = \epsilon - 1$. This can be indeed verified by considering (177) after setting $i = m, \lambda = \epsilon - 1$. In conclusion, the proportion of 'good' codes (among the entire ensemble) designed at this stage m , is no less than

$$(1 - q^{-\epsilon})(1 - q^{-\epsilon+1})^{\lceil \log_2 n(1-\tau) \rceil} \quad (180)$$

and for each such code C , every point $\mathbf{x} \in \mathcal{X}_C$ will be ρn -covered by at least n^α codewords of that same code.

With the above in place, let us return to (152) where we wish to guarantee that $\beta(\epsilon) > 0$. Toward this, let us consider (180) and in this equation, let us set $\epsilon = 2 \log_q \log_2(n(1-\tau))$.

Let us now prove that there exists an $\alpha > 1$ such that

$$\mathbb{E}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}'_C) \geq (n(1-\tau))^\alpha \quad (181)$$

$$= (q^{k^*} - 1)q^{nH_q(\rho) - o(n)}q^{-(n-n\tau)}. \quad (182)$$

This can be readily shown (cf. (160)) by noting that, for large n , then the expression

$$\exists \alpha > 1 : (n - n\tau)^\alpha = (q^{k^*} - 1)q^{nH_q(\rho) - o(n)}q^{-(n-n\tau)} \quad (183)$$

holds. With (181) in place, we take the logarithm on both sides of the above, and after dividing by n , we get

$$\frac{k^*}{n} = 1 - \tau - H_q(\rho) + \frac{\alpha \log_q(n - n\tau) + o(n)}{n}. \quad (184)$$

Let us now recall that we had conditionally accepted (161), by saying that (161) holds for some properly chosen ϵ and k^* . We will use the aforementioned $\epsilon = 2 \log_q \log_2(n(1-\tau))$ and the k^* from (184). Let us apply these values in the LHS of (161), and note, after employing (183), that for these values in place, it holds that

$$(q^{k^*} - 1)q^{nH_q(\rho) - o(n) - n(1-\tau)} \zeta(n) = (n - n\tau)^\alpha \zeta(n).$$

Let us now note that for sufficiently large n then

$$(n - n\tau)^\alpha \zeta(n) \geq q^2 \log_2(n(1-\tau))^4 \quad (185)$$

simply because the RHS is logarithmic in n . At the same time though, we also note that $q^2 \log_2(n(1-\tau))^4 =$

$q^{2(2\log_q(\log_2(n(1-\tau))))+2}$ and then, by applying the chosen ϵ , we get that $q^2 \log_2(n(1-\tau))^4 = q^{2\epsilon+2}$. Thus we now know that $(n-n\tau)^\alpha \zeta(n) \geq q^{2\epsilon+2}$ which, after applying (185), gives that

$$(q^{k^*} - 1)q^{nH_q(\rho) - o(n) - n(1-\tau)} \zeta(n) \geq q^{2\epsilon+2}$$

which is exactly (161). Thus (161) is validated, and consequently, directly, we can also guarantee (155), which in turn guarantees $\beta(\epsilon) \geq 0$, which in turn proves that (151) indeed holds.

Following the logic immediately before (151), and with (151) now in place, we can conclude that for any $\mathbf{x} \in \mathcal{X}'_C$, — given our chosen k^* and ϵ — $\eta_{\mathbf{x}}$ will be, with high probability, bigger than 0 which in turn implies that any $\mathbf{x} \in \mathcal{X}'_C$ will, with high probability, be covered by C . As it can be seen, the proportion of partial-covering codes (180) approaches 1, as n increases.

The only thing that remains now to be verified is the rate k_n/n of these codes, which was declared in the theorem to be as in (115). To verify that indeed this is our rate, we recall that we have started with a code from $C_{k^*,n}$ and that we have performed the appending procedure m times, where we chose $m = \log_2(n - n\tau)$. This means that the current message length becomes

$$k_n = k^* + \log_2(n(1 - \tau)).$$

Now directly adding m/n on both sides of the equation in (184), we have that

$$\frac{k_n}{n} = \frac{k^* + \log_2(n - n\tau)}{n} \quad (186)$$

$$= 1 - \tau - H(\rho) + \frac{(\alpha + 1) \log_2(n - n\tau) + o(n)}{n} \quad (187)$$

which simply says that $\frac{k_n}{n} \leq 1 - \tau - H_q(\rho) + O(n^{-1} \log_q(n))$ as in (115). This concludes the proof of the theorem. \square

APPENDIX G PROOF OF PROPOSITION 1

Referring to the proof of Theorem 3 in Appendix E, let us suppose that $L \leq m < q^K$ and that $\mathcal{X} \supseteq \mathcal{X}_{\mathbf{F}, \mathbf{D}}$. From (119) we see that

$$|\mathcal{X}| \stackrel{(a)}{=} \prod_{i=1}^m |\mathcal{X}_i| \quad (188)$$

$$\stackrel{(b)}{\geq} q^{nm_n(1-\tau)} \quad (189)$$

$$\stackrel{(c)}{=} q^{N(1-\tau)} \quad (190)$$

where (a) comes from the definition of \mathcal{X} (cf. (119)), where (b) holds due to (120), and where (c) holds since $N = nm_n$. Then from (124) and (190), we can conclude that

$$\rho = H_q^{-1}\left(\frac{K}{N} - \tau + \epsilon(N)\right) \quad (191)$$

$$\leq H_q^{-1}\left(\frac{K}{N} - \left(1 - \frac{\log_q(|\mathcal{X}|)}{N}\right) + \epsilon(N)\right). \quad (192)$$

Setting $m = L$ gives $\rho = H_q^{-1}\left(\frac{\log_q(L)}{N} + \epsilon(N)\right)$ (cf. (125)). The communication cost is as described in (128). \square

APPENDIX H VARIOUS PROOFS

A. Proof of Lemma 4

For a fixed index $i \neq 0$, there is a fixed information vector $\mathbf{d}_i \in \mathbb{F}^{k^*} \setminus \mathbf{0}$ that generates — as we move across the generator matrices \mathbf{G} in the ensemble of codes $C_{k^*,n}$ — the codewords \mathbf{c}_i that take the form

$$\mathbf{c}_i = \mathbf{d}_i \mathbf{G} = \sum_{j=1}^n d_i(j, 1) \mathbf{G}(j, :). \quad (193)$$

Given that the elements of $\mathbf{G} \in \mathbb{F}^{k^* \times n}$ are chosen uniformly and independently at random from \mathbb{F} , directly implies that the same holds for the elements of \mathbf{c}_i , since in the above linear combination, the elements of \mathbf{d}_i are fixed, and naturally because the operations are over a *finite* field. \square

B. Proof of Lemma 5

We can first see that whenever $\omega(\mathbf{x}) \leq \rho n$, then (133) automatically holds simply because such \mathbf{x} must belong in $\mathcal{B}(0, \rho)$ which in turn is a subset of \mathcal{X}_C .

Let us now consider the case of $\omega(\mathbf{x}) > \rho n$, $\mathbf{x} \in \mathbb{F}^n$. Let us also recall the element selection process²¹ that was described right underneath equation (132). Let $\mathcal{S} \subset \mathbb{F}^n$ be the set of all vectors \mathbf{x} that are not codewords but are selected randomly in the aforementioned process. At this point, we can see that

$$\mathbb{P}(\mathbf{x} \in \mathcal{X}_C) = \mathbb{P}(\mathbf{x} \in C) + \mathbb{P}(\mathbf{x} \notin C) \mathbb{P}(\mathbf{x} \in \mathcal{S} \mid \mathbf{x} \notin C) \quad (194)$$

where C is the code that has been chosen uniformly at random from $C_{k^*,n}$. Consider a vector $\mathbf{y} \in \mathbb{F}^n$ with $\omega(\mathbf{y}) \geq \rho n$. We clearly see that $\mathbb{P}(\mathbf{x} \in C) = \mathbb{P}(\mathbf{y} \in C)$, and we also see that $\mathbb{P}(\mathbf{x} \in \mathcal{S} \mid \mathbf{x} \notin C) = \mathbb{P}(\mathbf{y} \in \mathcal{S} \mid \mathbf{y} \notin C)$ as a direct outcome of the aforementioned vector selection process, and of the fact that $\omega(\mathbf{x}) > \rho n$, $\omega(\mathbf{y}) > \rho n$, which yields

$$\mathbb{P}(\mathbf{x} \in \mathcal{X}_C) = \mathbb{P}(\mathbf{y} \in \mathcal{X}_C). \quad (195)$$

Now let us note that

$$\sum_{\mathbf{y} \in \mathbb{F}^n \setminus \mathcal{B}(0, \rho)} \mathbb{P}(\mathbf{y} \in \mathcal{X}_C) = \sum_{\mathbf{y} \in \mathbb{F}^n \setminus \mathcal{B}(0, \rho)} \mathbb{E}[\mathbb{1}(\mathbf{y} \in \mathcal{X}_C)] \quad (196)$$

$$\stackrel{(a)}{=} \mathbb{E}\left[\sum_{\mathbf{y} \in \mathbb{F}^n \setminus \mathcal{B}(0, \rho)} \mathbb{1}(\mathbf{y} \in \mathcal{X}_C)\right] \quad (197)$$

$$\stackrel{(b)}{=} \mathbb{E}[q^{n(1-\tau)} - V_q(n, \rho)] \quad (198)$$

$$= q^{n(1-\tau)} - V_q(n, \rho) \quad (199)$$

where the average is over the codes in the ensemble $C_{k^*,n}$ and over the randomness in constructing \mathcal{X}_C once $C \in C_{k^*,n}$ is picked. In the above, (a) follows by interchanging summation and expectation, and (b) holds since for every occurrence of

²¹We recall that the procedure for designing \mathcal{X}_C , simply starts by taking the union $C \cup \mathcal{B}(0, \rho)$, and then appending on this union, a sufficiently large number of vectors, chosen uniformly and independently at random from $\mathbb{F}^n \setminus (C \cup \mathcal{B}(0, \rho))$.

$C \in C_{k^*, n}$, there exist $q^{n(1-\tau)} - V_q(n, \rho)$ elements of $\mathbb{F} \setminus \mathcal{B}(0, \rho)$ that are in \mathcal{X}_C . Finally, (195) and (199) jointly imply that

$$(q^n - V_q(n, \rho))\mathbb{P}(\mathbf{x} \in \mathcal{X}_C) = \sum_{\mathbf{y} \in \mathbb{F}^n \setminus \mathcal{B}(0, \rho)} \mathbb{P}(\mathbf{y} \in \mathcal{X}_C) \quad (200)$$

$$= q^{n(1-\tau)} - V_q(n, \rho) \quad (201)$$

which completes the proof. \square

C. Proof of Lemma 6

For any $i \neq 1$ and $\mathbf{x} \in \mathbb{F}^n$, $\mathbf{x} \neq \mathbf{0}$, then

$$\mathbb{P}[\mathbf{c}_i = \mathbf{x} | \mathbf{x} \in \mathcal{X}_C] \mathbb{P}[\mathbf{x} \in \mathcal{X}_C] \quad (202)$$

$$\stackrel{(a)}{=} \mathbb{P}[\{\mathbf{c}_i = \mathbf{x}\} \cap \{\mathbf{x} \in \mathcal{X}_C\}] \quad (203)$$

$$\stackrel{(b)}{=} \mathbb{P}[\{\mathbf{c}_i = \mathbf{x}\} \cap \{\mathbf{c}_i \in \mathcal{X}_C\}] \quad (204)$$

$$\stackrel{(c)}{=} \mathbb{P}[\mathbf{c}_i = \mathbf{x}] \quad (205)$$

$$\stackrel{(d)}{=} q^{-n} \quad (206)$$

where (a) is directly from the definition of conditional probability [71], (b) is true since the LHS requirement that $\mathbf{x} = \mathbf{c}_i$ is maintained in the RHS, (c) is true since $C \subset \mathcal{X}_C$, and (d) is from (129) in Lemma 4.

Thus for $i \neq 1$ and $\mathbf{x} \neq \mathbf{0}$, we have that

$$\begin{aligned} \mathbb{P}(\mathbf{c}_i = \mathbf{x} | \mathbf{x} \in \mathcal{X}_C) &= q^{-n(1-\tau)} \frac{1 - q^{-n}V(\rho, n)}{1 - q^{-(n-n\tau)}V(\rho, n)} \\ &= q^{-(n-n\tau)}\zeta(n), \quad i \in [2 : K^*], \quad \omega(\mathbf{x}) > \rho n \end{aligned}$$

and the proof is concluded by noting that in the limit of large n , the expression

$$\zeta(n) \triangleq \frac{1 - q^{-n}V(\rho, n)}{1 - q^{-(n-n\tau)}V(\rho, n)}$$

converges to 1. \square

D. Proof of Lemma 7

From the definition in (137), let us recall that

$$\eta_{\mathbf{x}} \triangleq \sum_{i=1}^{q^{k^*}} \eta_{\mathbf{x}, i} \quad (207)$$

describes the number of codewords that cover $\mathbf{x} \in \mathcal{X}'_C$. Consider a Hamming ball of radius ρ centered around $\mathbf{x} \in \mathcal{X}'_C$.

- Considering (135), we know that if $|\{0\} \cap \mathcal{B}(\mathbf{x}, \rho)| = 1$, then with probability one, $\mathbf{0}$ covers \mathbf{x} . Hence now the assumption that $\mathbf{x} \in \mathcal{X}'_C$, contradicts the above, and thus we can conclude that $|\{0\} \cap \mathcal{B}(\mathbf{x}, \rho)| = 0$.
- Now consider some vector \mathbf{x}' in $(\mathcal{B}(0, \rho) \setminus \{0\}) \cap \mathcal{B}(\mathbf{x}, \rho)$, i.e., some vector that covers our aforementioned $\mathbf{x} \in \mathcal{X}'_C$. We are interested in the probability $\mathbb{P}(\mathbf{c}_i = \mathbf{x}' | \mathbf{x}' \in \mathcal{X}'_C)$ which is the probability that \mathbf{x}' is equal to \mathbf{c}_i , for our fixed $i, i \neq 1$. From (135), we know that this probability is equal to q^{-n} . Now going over all $q^{k^*} - 1$ codewords of interest, we can conclude that our current case of interest, contributes to the sum $\eta_{\mathbf{x}}$, by an amount equal to $(q^{k^*} - 1)[(\mathcal{B}(0, \rho) \setminus \{0\}) \cap \mathcal{B}(\mathbf{x}, \rho)]q^{-n}$.

- Now let us consider the dominant case where \mathbf{x}' in $\mathcal{B}(\mathbf{x}, \rho) \setminus \mathcal{B}(0, \rho)$. In this case, directly from (135), the aforementioned probability $\mathbb{P}(\mathbf{c}_i = \mathbf{x}' | \mathbf{x}' \in \mathcal{X}'_C)$ takes the form $q^{-n(1-\tau)}\zeta(n)$, and thus — similarly to above — yields a contribution to the sum $\eta_{\mathbf{x}}$ by an amount equal to

$$(q^{k^*} - 1)[(\mathcal{B}(\mathbf{x}, \rho) \setminus \mathcal{B}(0, \rho))q^{-n(1-\tau)}\zeta(n).$$

\square

E. Proof of Lemma 8

We prove the lemma in two steps.

In the first step, after defining $\bar{\eta} \triangleq \mathbb{E}(\eta_{\mathbf{x}, i} | \mathbf{x} \in \mathcal{X}'_C)$ and $\overline{\eta^{(2)}} \triangleq \mathbb{E}(\eta_{\mathbf{x}, i}\eta_{\mathbf{x}, j} | \mathbf{x} \in \mathcal{X}'_C)$ for any $i, j \in [q^{k^*}]$, we can see that

$$\text{Var}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}_C) \leq (q^{k^*} - 1)(q - 1)\bar{\eta}(1 - \frac{q-2}{q-1}\bar{\eta}) \quad (208)$$

since

$$\text{Var}(\eta_{\mathbf{x}} | \mathbf{x} \in \mathcal{X}_C) = \mathbb{E}(\left(\sum_{i=1}^{q^{k^*}} \eta_{\mathbf{x}, i}\right)^2 | \mathbf{x} \in \mathcal{X}'_C) \quad (209)$$

$$- \mathbb{E}^2\left(\sum_{i=1}^{q^{k^*}} \eta_{\mathbf{x}, i} | \mathbf{x} \in \mathcal{X}'_C\right) \quad (210)$$

$$= \mathbb{E}\left(\sum_{i=1}^{q^{k^*}} \eta_{\mathbf{x}, i}^2\right) \quad (211)$$

$$+ \sum_{p, i=1, i \neq p}^{q^{k^*}} \eta_{\mathbf{x}, i}\eta_{\mathbf{x}, p} | \mathbf{x} \in \mathcal{X}'_C) \quad (212)$$

$$- \mathbb{E}^2\left(\sum_{i=1}^{q^{k^*}} \eta_{\mathbf{x}, i} | \mathbf{x} \in \mathcal{X}'_C\right) \quad (213)$$

$$= \sum_{i=1}^{q^{k^*}} \mathbb{E}(\eta_{\mathbf{x}, i} | \mathbf{x} \in \mathcal{X}'_C) \quad (214)$$

$$+ \sum_{p, i=1, i \neq p}^{q^{k^*}} \mathbb{E}(\eta_{\mathbf{x}, i}\eta_{\mathbf{x}, p} | \mathbf{x} \in \mathcal{X}'_C) \quad (215)$$

$$- \mathbb{E}^2\left(\sum_{i=1}^{q^{k^*}} \eta_{\mathbf{x}, i} | \mathbf{x} \in \mathcal{X}'_C\right) \quad (216)$$

$$= q^{k^*}\bar{\eta} + q^{k^*}(q^{k^*} - 1)\overline{\eta^{(2)}} - q^{2k^*}\bar{\eta}. \quad (217)$$

Combining now the above with (208), allows us to modify the main claim as

$$q^{k^*}\bar{\eta} + q^{k^*}(q^{k^*} - 1)\overline{\eta^{(2)}} - q^{2k^*}\bar{\eta} \quad (218)$$

$$- (q^{k^*} - 1)(q - 1)\bar{\eta}(1 - \frac{q-2}{q-1}\bar{\eta}) \leq 0 \quad (219)$$

so we now need to prove (219). To prove this, we focus on the LHS and show that

$$q^{k^*}\bar{\eta} + q^{k^*}(q^{k^*} - 1)\overline{\eta^{(2)}} \quad (220)$$

$$- q^{2k^*}\bar{\eta} - (q^{k^*} - 1)(q - 1)\bar{\eta}(1 - \frac{q-2}{q-1}\bar{\eta}) \quad (221)$$

$$\stackrel{(a)}{\leq} q^{k^*}\bar{\eta} + q^{k^*}(q^{k^*} - 1)\overline{\eta^{(2)}} - q^{2k^*}\bar{\eta} \quad (222)$$

$$-(q^{k^*} - 1)(q - 1)\bar{\eta}\left(1 - \frac{q-2}{q-1}\bar{\eta}\right) \quad (223)$$

$$\stackrel{(b)}{\leq} -(q^{k^*+1} - 2q^{k^*} - q + 1)\bar{\eta} \quad (224)$$

$$+ (q^{k^*+1} - 3q^{k^*} - q + 2)\bar{\eta}^2 \quad (225)$$

$$\stackrel{(c)}{\leq} 0 \quad (226)$$

where (a) holds because $\overline{\eta^{(2)}} \leq \bar{\eta}^2$, (b) follows after simple rearranging of terms, and (c) holds because $0 \leq \bar{\eta} \leq 1$ and because $q^{k^*+1} - 2q^{k^*} - q + 1 > q^{k^*+1} - 3q^{k^*} - q + 2$ for any $q^{k^*} \geq 1$.

In the second step, we will prove that

$$\frac{(q^{k^*} - 1)(q - 1)\bar{\eta}\left(1 - \frac{q-2}{q-1}\bar{\eta}\right)}{\mathbb{E}(\eta_{\mathbf{x}}|\mathbf{x} \in \mathcal{X}'_C)q^2} \leq 1. \quad (227)$$

To see this, we note that

$$\frac{(q^{k^*} - 1)(q - 1)\bar{\eta}\left(1 - \frac{q-2}{q-1}\bar{\eta}\right)}{\mathbb{E}(\eta_{\mathbf{x}}|\mathbf{x} \in \mathcal{X}'_C)q^2} \quad (228)$$

$$\stackrel{(a)}{=} \frac{(q^{k^*} - 1)(q - 1)\left(1 - \frac{q-2}{q-1}\bar{\eta}\right)}{q^{k^*}q^2} \quad (229)$$

$$\stackrel{(b)}{=} \left(\frac{q^{k^*} - 1}{q^{k^*}}\right)\left(\frac{q-1}{q^2}\right)\left(1 - \frac{q-2}{q-1}\bar{\eta}\right) \quad (230)$$

$$\stackrel{(c)}{\leq} 1 \quad (231)$$

where (a) holds because $\mathbb{E}(\eta_{\mathbf{x}}) = \sum_{i=1}^{q^{k^*}} \mathbb{E}(\eta_{\mathbf{x},i}) = q^{k^*}\bar{\eta}$, where (b) holds by rearranging terms, and where (c) holds because each multiplicative element in the RHS of (b) is non-negative and less than 1.

Now combining the two steps by bringing together (208) with (227), yields the desired (142). \square

APPENDIX I PROOF OF LEMMA 9

Let us define

$$\mathcal{I}_{(\omega(\mathbf{x}),\rho)} \triangleq |\mathcal{B}(\mathbf{x},\rho) \cap \mathcal{B}(0,\rho)| \quad (232)$$

and let us note that

$$\rho n = \arg \max_{\omega(\mathbf{x}):\mathbf{x} \in \mathcal{X}_C \setminus \mathcal{B}(0,\rho)} |\mathcal{I}_{(\omega(\mathbf{x}),\rho)}| \quad (233)$$

since the distance between $\mathbf{x} \in \mathcal{X}_C \setminus \mathcal{B}(0,\rho)$ and $\mathbf{0}$ is minimized when $\omega(\mathbf{x}) = \rho n$. We also know that

$$|\mathcal{B}(\mathbf{x},\rho) \setminus \mathcal{B}(0,\rho)| = \text{Vol}(n,\rho) - |\mathcal{I}(\rho,\rho)|. \quad (234)$$

Let us focus on the case where $q = 2$ and $0 \leq \rho \leq \frac{1}{2}$, where from [72] we have that

$$|\mathcal{I}(\rho,\rho)| = \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \sum_{j=0}^i \binom{n\rho}{i} \binom{n-n\rho}{j} \quad (235)$$

$$+ \sum_{i=\lfloor \frac{n\rho}{2} \rfloor + 1}^{n\rho} \sum_{j=0}^{n\rho-i} \binom{n\rho}{i} \binom{n-n\rho}{j}. \quad (236)$$

We know that $n\rho \leq n - n\rho$ and that

$$V_q(n,\rho) = \sum_{i=0}^{n\rho} \binom{n}{i} = \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \sum_{j=0}^{n\rho-i} \binom{n\rho}{i} \binom{n-n\rho}{j} \quad (237)$$

$$+ \sum_{i=\lfloor \frac{n\rho}{2} \rfloor + 1}^{n\rho} \sum_{j=0}^{n\rho-i} \binom{n\rho}{i} \binom{n-n\rho}{j} \quad (238)$$

and thus after substituting (236) and (238) into (234), we conclude that

$$|\mathcal{B}(\mathbf{x},\rho) \setminus \mathcal{B}(0,\rho)| = \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \sum_{j=i+1}^{n\rho-i} \binom{n\rho}{i} \binom{n-n\rho}{j}. \quad (239)$$

Now considering that $0 < \rho \leq \frac{1}{2}$, we have

$$\binom{n}{n\rho} = \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \binom{n\rho}{i} \binom{n-n\rho}{n\rho-i} + \sum_{i=\lfloor \frac{n\rho}{2} \rfloor + 1}^{n\rho} \binom{n\rho}{i} \binom{n-n\rho}{n\rho-i} \quad (240)$$

$$= \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \binom{n\rho}{i} \binom{n-n\rho}{n\rho-i} \quad (241)$$

$$+ \sum_{i=\lfloor \frac{n\rho}{2} \rfloor + 1}^{n\rho} \binom{n\rho}{i} \binom{n-n\rho}{n\rho-i} \quad (242)$$

$$= \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \binom{n\rho}{i} \binom{n-n\rho}{n\rho-i} + \sum_{i=0}^{n\rho - \lfloor \frac{n\rho}{2} \rfloor - 1} \binom{n\rho}{i} \binom{n-n\rho}{i}. \quad (243)$$

Now let us note that

$$|\mathcal{B}(\mathbf{x},\rho) \setminus \mathcal{B}(0,\rho)| \stackrel{(a)}{=} \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \sum_{j=i+1}^{n\rho-i} \binom{n\rho}{i} \binom{n-n\rho}{j} \quad (244)$$

$$\stackrel{(b)}{=} \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \sum_{j=i+2}^{n\rho-i-1} \binom{n\rho}{i} \binom{n-n\rho}{j} \quad (245)$$

$$+ \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \binom{n\rho}{i} \binom{n-n\rho}{n\rho-i} (q-1)^{n\rho} \quad (246)$$

$$+ \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor - 1} \sum_{j=\lfloor \frac{n\rho}{2} \rfloor}^{n\rho} \binom{n\rho}{i} \binom{n-n\rho}{i+1} \quad (247)$$

$$\stackrel{(c)}{\geq} \sum_{i=0}^{\lfloor \frac{n\rho}{2} \rfloor} \binom{n\rho}{i} \binom{n-n\rho}{n\rho-i} \quad (248)$$

$$+ \sum_{i=0}^{n\rho - \lfloor \frac{n\rho}{2} \rfloor - 1} \binom{n\rho}{i} \binom{n-n\rho}{i} \quad (249)$$

$$\stackrel{(d)}{=} \binom{n}{n\rho} \stackrel{(e)}{\geq} 2^{nH(\rho) - o(n)} \quad (250)$$

where (a) holds from (239), (b) follows after expanding the inner summation, (c) holds because the first summation of the RHS in (b) is non-negative, the second summation is present on the RHS of (c) after considering that $q = 2$, and because the third summation of the RHS in (b) is present on the RHS of (c) after considering Lemma 11 found in Appendix I-D. Furthermore (d) follows from (243), and (e) follows from the

Stirling inequality that applies because $0 < \rho \leq \frac{1}{2}$. The proof is concluded for the binary case of $q = 2$.

Let us now consider the more involved case of $q > 2$. For this we will need the following lemma, whose proof is found in Appendix I-C.

Lemma 10. For $\mathcal{I}_{(\omega(x), \rho)} \triangleq |\mathcal{B}(\mathbf{x}, \rho) \cap \mathcal{B}(0, \rho)|$ (cf. (232)), then

$$|\mathcal{I}(\rho, \rho)| = \sum_{i+j=\rho n, i \leq j} \binom{n-\rho n}{i} \binom{\rho n}{j} (q-1)^{\rho n} + \theta(q) \quad (251)$$

where we can guarantee that $\theta(q) \leq V_q(n, \rho - 1/n)$.

Proof. See Appendix I-C. \square

With this lemma in place, we now note that

$$|\mathcal{B}(\mathbf{x}, \rho) \setminus \mathcal{B}(0, \rho)| \quad (252)$$

$$\stackrel{(a)}{=} V_q(n, \rho) - |\mathcal{I}(\rho, \rho)| \quad (253)$$

$$\stackrel{(b)}{=} V_q(n, \rho - 1/n) + \binom{n-\rho n}{\rho n} \quad (254)$$

$$- \sum_{i+j=\rho n, i \leq j} \binom{n-\rho n}{i} \binom{\rho n}{j} \quad (255)$$

$$\times (q-1)^{\rho n} \quad (256)$$

$$- \theta(q) \quad (257)$$

$$\stackrel{(c)}{\geq} \sum_{i+j=\rho n, i > j} \binom{n-\rho n}{i} \binom{\rho n}{j} (q-1)^{\rho n} \quad (258)$$

$$\stackrel{(d)}{=} \sum_{j=\min\{0, 2\rho n - n\}}^{\lfloor \frac{\rho n}{2} \rfloor - 1} \binom{n-\rho n}{\rho n - j} \binom{\rho n}{j} (q-1)^{\rho n} \quad (259)$$

where (a) follows from a basic set cardinality rule, (b) follows from (251), (c) follows from Lemma (10) which tells us that $\theta(q) \leq V_q(n, \rho - 1/n)$, and from the fact that $\binom{n-\rho n}{\rho n} = \sum_{i+j=\rho n} \binom{n-\rho n}{i} \binom{\rho n}{j} (q-1)^{\rho n}$, and (d) holds since $j \geq 0$ and at the same time $i \leq n - \rho n$ which gives $j \geq 2\rho n - n$.

Now from Stirling's bound, we know that for any $j \in [\min\{0, 2\rho n - n\}, \lfloor \frac{\rho n}{2} \rfloor - 1]$, the following holds

$$\binom{n-\rho n}{\rho n - j} \binom{\rho n}{j} \geq \sqrt{\frac{n-\rho n}{8(\rho n - j)(n - 2\rho n + j)}} \quad (260)$$

$$\times 2^{nH((\rho n - j)/(n - \rho n))} \quad (261)$$

$$\times \sqrt{\frac{\rho n}{8(j)(\rho n - j)}} 2^{nH(j/\rho n)} \quad (262)$$

$$= 2^{n[H((\rho n - j)/\rho n)]} \quad (263)$$

$$\times 2^{n[H((\rho n - j)/(n - \rho n))] - o(n)}. \quad (264)$$

After defining

$$\kappa \triangleq \frac{\rho n - j}{\rho n} \quad (265)$$

the exponent in (264) takes the form

$$n[H(\kappa) + H(\kappa \frac{\rho}{1 - \rho})] - o(n). \quad (266)$$

With this exponent in place, let us consider a large n and let us assume without loss of generality²² that $n\rho^2 \in \mathbb{N}$. For the case where $0 < \rho \leq \frac{1}{2}$, let us set j such that $j = n\rho^2 \leq \lfloor \frac{\rho n}{2} \rfloor - 1$, in which case we get $\kappa = 1 - \rho$. On the other hand, for the case where $\frac{1}{2} < \rho \leq -1/2 + \sqrt{5}/2$, let us set j such that $2\rho n - n \leq j = n\rho(1 - \rho) \leq \lfloor \frac{\rho n}{2} \rfloor - 1$, in which case $\kappa = \rho$. In each case, κ is plugged in (266), and after utilizing (264), we see that

$$\sum_{j=\min\{0, 2\rho n - n\}}^{\lfloor \frac{\rho n}{2} \rfloor - 1} \binom{n-\rho n}{\rho n - j} \binom{\rho n}{j} (q-1)^{\rho n} \geq q^{nH_q(\rho) - o(n)} \quad (267)$$

which holds for the range $0 < \rho \leq -1/2 + \sqrt{5}/2$ which is the union of the above two regions in ρ . \square

A. Proof of Proposition 2

For $h(x) \triangleq -x \log_q(x)$, we know that

$$h(x) \leq H_q(x), \quad 0 \leq x \leq 1 - 1/q \quad (268)$$

and thus that

$$H_q^{-1}(x) \leq h^{-1}(x). \quad (269)$$

We also know that if $y = x \ln(x)$, then $x = e^{W(y)}$ where $W(\cdot)$ is the Lambert function. Also note that since $h(x) = -\log_q(e)x \ln(x)$, we have that

$$h^{-1}(x) = e^{W(-\ln(q)x)}. \quad (270)$$

Furthermore, for $c > 0$ being a positive real number, we have that

$$\lim_{T \rightarrow \infty} T e^{W(-c/T)} \stackrel{(a)}{=} \lim_{T \rightarrow \infty} \frac{e^{W(-c/T)}}{1/T} \quad (271)$$

$$\stackrel{(b)}{=} \lim_{T \rightarrow \infty} \frac{e^{W(-c/T)} \frac{1}{-c/T + e^{W(-c/T)}} c T^{-2}}{-T^{-2}} \quad (272)$$

$$\stackrel{(c)}{=} \lim_{T \rightarrow \infty} \frac{c e^{W(-c/T)}}{c/T - e^{W(-c/T)}} \quad (273)$$

$$\stackrel{(d)}{=} \lim_{T \rightarrow \infty} \frac{c T e^{W(-c/T)}}{c - T e^{W(-c/T)}} \quad (274)$$

$$\stackrel{(e)}{=} \frac{\lim_{T \rightarrow \infty} c T e^{W(-c/T)}}{\lim_{T \rightarrow \infty} c - T e^{W(-c/T)}} \quad (275)$$

where (a), (c) and (d) follow by basic algebraic rearranging, (b) follows from L'Hopital's rule, and (e) follows from the Algebraic Limit Theorem. The above implies that

$$\lim_{T \rightarrow \infty} T e^{W(-c/T)} = 0 \quad (276)$$

which allows us to conclude that

$$\lim_{T \rightarrow \infty} T H_q^{-1}(c/T) \stackrel{(a)}{\leq} \lim_{T \rightarrow \infty} T h^{-1}(c/T) \quad (277)$$

$$\stackrel{(b)}{\leq} \lim_{T \rightarrow \infty} T e^{W(-\ln(q)c/T)} \quad (278)$$

$$\stackrel{(c)}{=} 0 \quad (279)$$

where (a) follows from (269), (b) from (270), and (c) from (276). Thus the proof is concluded. \square

²²This assumption, along with the assumption that $n\rho$ is an integer, has no impact on the result, because any non-integer residual will vanish in importance as n increases.

B. Proof of Proposition 3

Starting from

$$c/T = H_q(f/T) \quad (280)$$

we take the derivative with respect to T on both sides, to get

$$c = \log_q\left(\frac{f/T}{1-f/T}(q-1)\right)\left(\frac{\partial f}{\partial T}T - f\right). \quad (281)$$

After applying (280) into (281), and after some basic algebraic rearranging, the proof is concluded. \square

C. Proof of Lemma 10

Let us consider two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{F}^n$, where $\omega(\mathbf{a}) = \rho n$ and where $\mathbf{b} \in \mathcal{B}(\mathbf{0}, \rho) \cap \mathcal{B}(\mathbf{a}, \rho)$. Let

$$\mathcal{B} \triangleq \{\mathbf{b}(i), i \in [n] : \mathbf{a}(i) = 0 \text{ \& \ } \mathbf{b}(i) \neq 0\} \quad (282)$$

$$\mathcal{C} \triangleq \{\mathbf{b}(i), i \in [n] : \mathbf{a}(i) \neq 0 \text{ \& \ } \mathbf{b}(i) = \mathbf{a}(i)\} \quad (283)$$

and let $x \triangleq |\mathcal{B}|, y \triangleq |\mathcal{C}|$. Let $b_j, j \in [x]$ denote the j th element of \mathcal{B} , and let $c_j, j \in [y]$ denote the j th element of \mathcal{C} . Note that ordering does not matter. Let $a_j, j \in [\rho n]$ be the non-zero elements of \mathbf{a} , and without loss of generality, let

$$\mathbf{a} = [0, \dots, 0, a_1, a_2, \dots, a_y, \dots, a_{\rho n}] \quad (284)$$

as well as let

$$\mathbf{b} = [b_1, b_2, \dots, b_x, 0, 0, \dots, 0, \quad (285)$$

$$c_1, c_2, \dots, c_y, 0, \dots, 0]. \quad (286)$$

Since $\omega(\mathbf{b}) \leq \rho n$ and $d(\mathbf{a}, \mathbf{b}) \leq \rho n$, we have that

$$x + y \leq \rho n \quad (287)$$

$$x \leq y \quad (288)$$

and we have that for any fixed pair $(x, y) \in [\rho n]^2$, there are $\binom{n-\rho n}{x} \binom{\rho n}{y} (q-1)^{x+y}$ such points (that satisfy that given (x, y)) in $\mathcal{B}(\mathbf{0}, \rho) \cap \mathcal{B}(\mathbf{a}, \rho)$. Therefore, accumulating over all (x, y) for which $x + y = \rho n$, the overall intersection gains cardinality $|\mathcal{B}(\mathbf{0}, \rho) \cap \mathcal{B}(\mathbf{a}, \rho)| = \sum_{i+j=\rho n, i \leq j} \binom{n-\rho n}{i} \binom{\rho n}{j} (q-1)^{\rho n}$, while when $x + y < \rho n$ — due to (288) — this same accumulated intersection has at most $V_q(n, \rho - 1/n)$ points. This concludes the proof. \square

D. Statement and Proof of Lemma 11

Lemma 11. If $0 < \rho \leq 1/2$ and $\rho n \in \mathbb{N}$, then

$$\begin{aligned} & \sum_{i=0}^{\lfloor \frac{\rho n}{2} \rfloor - 1} \binom{\rho n}{i} \binom{n-\rho n}{i+1} \\ & \geq \sum_{i=0}^{\rho n - \lfloor \frac{\rho n}{2} \rfloor - 1} \binom{\rho n}{i} \binom{n-\rho n}{i}. \end{aligned}$$

Proof. We solve the problem by considering the following cases.

Case 1: ($2 \mid \rho n$). We note that if $2 \mid \rho n$, then $\lfloor \frac{\rho n}{2} \rfloor = \frac{\rho n}{2}$, which gives

$$\sum_{i=0}^{\frac{\rho n}{2} - 1} \binom{\rho n}{i} \binom{n-\rho n}{i+1} \geq \sum_{i=0}^{\frac{\rho n}{2} - 1} \binom{\rho n}{i} \binom{n-\rho n}{i}. \quad (289)$$

Since $0 < \rho \leq \frac{1}{2}$, we have that $i + 1 \leq (n - \rho n)/2, \forall i \in [0 : \rho n/2 - 1]$, which gives

$$\binom{n-\rho n}{i+1} \geq \binom{n-\rho n}{i} \quad (290)$$

to conclude the proof for this case.

Case 2: ($2 \nmid \rho n$). Here we note that having $2 \nmid \rho n$, implies $\lfloor \frac{\rho n}{2} \rfloor = \frac{\rho n}{2} - 0.5$ which gives

$$\sum_{i=0}^{\frac{\rho n}{2} - 0.5} \binom{\rho n}{i} \binom{n-\rho n}{i+1} \geq \sum_{i=0}^{\frac{\rho n}{2} - 0.5} \binom{\rho n}{i} \binom{n-\rho n}{i}. \quad (291)$$

For this case, we consider the following two subcases.

Case 2a: (n is odd). We first note that if n is an odd number, then $n - \rho n$ is an even number. Also since $\rho n \in \mathbb{N}, 0 < \rho < 1/2$ and $\rho n \leq \frac{n-1}{2}$, then $2(\frac{\rho n}{2} + 0.5) \leq n - \rho n$. Thus as before we can say that $\forall i \in [0 : \rho n/2 - 0.5], i + 1 \leq (n - \rho n)/2$, which gives

$$\binom{n-\rho n}{i+1} \geq \binom{n-\rho n}{i} \quad (292)$$

which in turn concludes the proof for this case.

Case 2b: (n is even). If n is even then $n - \rho n$ is odd, and thus we can again say that having $0 < \rho \leq 1/2$ gives $i + 1 \leq (n - \rho n + 1)/2, \forall i \in [0 : \rho n/2 - 0.5]$, which gives

$$\binom{n-\rho n}{i+1} \geq \binom{n-\rho n}{i} \quad (293)$$

which in turn completes the proof for this final case also. \square

REFERENCES

- [1] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [2] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica, "Spark: Cluster computing with working sets," in *2nd USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 10)*, 2010.
- [3] T. Jahani-Nezhad and M. A. Maddah-Ali, "Codedsketch: A coding scheme for distributed computation of approximated matrix multiplication," *IEEE Transactions on Information Theory*, vol. 67, no. 6, pp. 4185–4196, 2021.
- [4] J. Wang, Z. Jia, and S. A. Jafar, "Price of precision in coded distributed matrix multiplication: A dimensional analysis," in *2021 IEEE Information Theory Workshop (ITW)*, pp. 1–6, IEEE, 2021.
- [5] E. Ozfatura, S. Ulukus, and D. Gündüz, "Coded distributed computing with partial recovery," *IEEE Transactions on Information Theory*, vol. 68, no. 3, pp. 1945–1959, 2022.
- [6] S. Li, Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "A scalable framework for wireless distributed computing," *IEEE/ACM Transactions on Networking*, vol. 25, no. 5, pp. 2643–2654, 2017.
- [7] F. Haddadpour, M. M. Kamani, M. Mahdavi, and V. Cadambe, "Trading redundancy for communication: Speeding up distributed sgd for non-convex optimization," in *International Conference on Machine Learning*, pp. 2545–2554, PMLR, 2019.
- [8] C.-S. Yang, R. Pedarsani, and A. S. Avestimehr, "Coded computing in unknown environment via online learning," in *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 185–190, IEEE, 2020.
- [9] N. Charalambides, H. Mahdavi, and A. O. Hero, "Numerically stable binary gradient coding," in *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 2622–2627, 2020.
- [10] M. Soleymani, H. Mahdavi, and A. S. Avestimehr, "Analog lagrange coded computing," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 283–295, 2021.
- [11] H. Sun and S. A. Jafar, "The capacity of private computation," *IEEE Transactions on Information Theory*, vol. 65, no. 6, pp. 3880–3897, 2018.

- [12] M. Soleymani and H. MahdaviFar, "Distributed multi-user secret sharing," *IEEE Transactions on Information Theory*, vol. 67, no. 1, pp. 164–178, 2020.
- [13] A. Khaledi, M. Mirmohseni, and M. A. Maddah-Ali, "The capacity region of distributed multi-user secret sharing," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 3, pp. 1057–1071, 2021.
- [14] M. Soleymani, H. MahdaviFar, and A. S. Avestimehr, "Privacy-preserving distributed learning in the analog domain," *arXiv preprint arXiv:2007.08803*, 2020.
- [15] M. Soleymani, R. E. Ali, H. MahdaviFar, and A. S. Avestimehr, "List-decodable coded computing: Breaking the adversarial toleration barrier," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 3, pp. 867–878, 2021.
- [16] C. Hofmeister, R. Bitar, M. Xhemrishi, and A. Wachter-Zeh, "Secure private and adaptive matrix multiplication beyond the singleton bound," *IEEE Journal on Selected Areas in Information Theory*, vol. 3, no. 2, pp. 275–285, 2022.
- [17] C. Hofmeister, R. Bitar, M. Xhemrishi, and A. Wachter-Zeh, "Secure private and adaptive matrix multiplication beyond the singleton bound," *arXiv preprint arXiv:2108.05742*, 2021.
- [18] H. Akbari-Nodehi and M. A. Maddah-Ali, "Secure coded multi-party computation for massive matrix operations," *IEEE Transactions on Information Theory*, vol. 67, no. 4, pp. 2379–2398, 2021.
- [19] Z. Jia and S. A. Jafar, "On the capacity of secure distributed batch matrix multiplication," *IEEE Transactions on Information Theory*, vol. 67, no. 11, pp. 7420–7437, 2021.
- [20] Z. Chen, Z. Jia, Z. Wang, and S. A. Jafar, "Gcsa codes with noise alignment for secure coded multi-party batch matrix multiplication," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 306–316, 2021.
- [21] C.-S. Yang and A. S. Avestimehr, "Coded computing for secure boolean computations," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 326–337, 2021.
- [22] Q. Yu and A. S. Avestimehr, "Coded computing for resilient, secure, and privacy-preserving distributed matrix multiplication," *IEEE Transactions on Communications*, vol. 69, no. 1, pp. 59–72, 2020.
- [23] M. Xhemrishi, R. Bitar, and A. Wachter-Zeh, "Distributed matrix-vector multiplication with sparsity and privacy guarantees," in *2022 IEEE International Symposium on Information Theory (ISIT)*, pp. 1028–1033, 2022.
- [24] N. Raviv, I. Tamo, R. Tandon, and A. G. Dimakis, "Gradient coding from cyclic mds codes and expander graphs," *IEEE Transactions on Information Theory*, vol. 66, no. 12, pp. 7475–7489, 2020.
- [25] 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, 2017.
- [26] M. Egger, R. Bitar, A. Wachter-Zeh, and D. Gündüz, "Efficient distributed machine learning via combinatorial multi-armed bandits," in *2022 IEEE International Symposium on Information Theory (ISIT)*, pp. 1653–1658, 2022.
- [27] K. Wan, H. Sun, M. Ji, and G. Caire, "Distributed linearly separable computation," *IEEE Transactions on Information Theory*, vol. 68, no. 2, pp. 1259–1278, 2022.
- [28] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "Straggler mitigation in distributed matrix multiplication: Fundamental limits and optimal coding," *IEEE Transactions on Information Theory*, vol. 66, no. 3, pp. 1920–1933, 2020.
- [29] Q. Yu, M. Maddah-Ali, and S. Avestimehr, "Polynomial codes: an optimal design for high-dimensional coded matrix multiplication," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [30] Z. Jia and S. A. Jafar, "Cross subspace alignment codes for coded distributed batch computation," *IEEE Transactions on Information Theory*, vol. 67, no. 5, pp. 2821–2846, 2021.
- [31] A. Behrouzi-Far and E. Soljanin, "Efficient replication for straggler mitigation in distributed computing," *arXiv preprint arXiv:2006.02318*, 2020.
- [32] E. Parrinello, E. Lampiris, and P. Elia, "Coded distributed computing with node cooperation substantially increases speedup factors," in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 1291–1295, 2018.
- [33] J. S. Ng, W. Y. B. Lim, N. C. Luong, Z. Xiong, A. Asheralieva, D. Niyato, C. Leung, and C. Miao, "A survey of coded distributed computing," *arXiv preprint arXiv:2008.09048*, 2020.
- [34] S. Li and S. Avestimehr, *Coded Computing: Mitigating Fundamental Bottlenecks in Large-Scale Distributed Computing and Machine Learning*, vol. 17. 2020.
- [35] F. Brunero and P. Elia, "Fundamental limits of combinatorial multi-access caching," *IEEE Transactions on Information Theory*, vol. 69, no. 2, pp. 1037–1056, 2023.
- [36] F. Brunero and P. Elia, "Unselfish coded caching can yield unbounded gains over selfish caching," *IEEE Transactions on Information Theory*, vol. 68, no. 12, pp. 7871–7891, 2022.
- [37] A. C.-C. Yao, "Communication complexity and its applications," in *International Workshop on Frontiers in Algorithmics*, pp. 2–2, Springer, 2009.
- [38] J. Verbraeken, M. Wolting, J. Katzy, J. Kloppenburg, T. Verbelen, and J. S. Rellermeier, "A survey on distributed machine learning," *ACM Computing Surveys (CSUR)*, vol. 53, no. 2, pp. 1–33, 2020.
- [39] S. Ulukus, S. Avestimehr, M. Gastpar, S. Jafar, R. Tandon, and C. Tian, "Private retrieval, computing and learning: Recent progress and future challenges," *IEEE Journal on Selected Areas in Communications*, 2022.
- [40] S. Wang, J. Liu, N. Shroff, and P. Yang, "Fundamental limits of coded linear transform," *arXiv preprint arXiv:1804.09791*, 2018.
- [41] S. Li, M. A. Maddah-Ali, Q. Yu, and A. S. Avestimehr, "A fundamental tradeoff between computation and communication in distributed computing," *IEEE Transactions on Information Theory*, vol. 64, no. 1, pp. 109–128, 2017.
- [42] K. Wan, H. Sun, M. Ji, D. Tuninetti, and G. Caire, "Cache-aided matrix multiplication retrieval," *IEEE Transactions on Information Theory*, vol. 68, no. 7, pp. 4301–4319, 2022.
- [43] Q. Yu, M. Maddah-Ali, and S. Avestimehr, "Polynomial codes: an optimal design for high-dimensional coded matrix multiplication," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [44] S. Dutta, M. Fahim, F. Haddadpour, H. Jeong, V. Cadambe, and P. Grover, "On the optimal recovery threshold of coded matrix multiplication," *IEEE Transactions on Information Theory*, vol. 66, no. 1, pp. 278–301, 2019.
- [45] A. Reiszadeh, S. Prakash, R. Pedarsani, and A. S. Avestimehr, "Codereduce: A fast and robust framework for gradient aggregation in distributed learning," *IEEE/ACM Transactions on Networking*, 2021.
- [46] N. Woolsey, R.-R. Chen, and M. Ji, "A new combinatorial coded design for heterogeneous distributed computing," *IEEE Transactions on Communications*, vol. 69, no. 9, pp. 5672–5685, 2021.
- [47] N. Woolsey, R.-R. Chen, and M. Ji, "Coded elastic computing on machines with heterogeneous storage and computation speed," *IEEE Transactions on Communications*, vol. 69, no. 5, pp. 2894–2908, 2021.
- [48] N. Woolsey, J. Kliever, R.-R. Chen, and M. Ji, "A practical algorithm design and evaluation for heterogeneous elastic computing with stragglers," in *2021 IEEE Global Communications Conference (GLOBECOM)*, pp. 1–6, 2021.
- [49] M. Chen, D. Gündüz, K. Huang, W. Saad, M. Bennis, A. V. Feljan, and H. V. Poor, "Distributed learning in wireless networks: Recent progress and future challenges," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 12, pp. 3579–3605, 2021.
- [50] M. Zinkevich, M. Weimer, L. Li, and A. Smola, "Parallelized stochastic gradient descent," in *Advances in Neural Information Processing Systems*, vol. 23, 2010.
- [51] T. Chilimbi, Y. Suzue, J. Apacible, and K. Kalyanaraman, "Project Adam: Building an efficient and scalable deep learning training system," in *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, pp. 571–582, 2014.
- [52] R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *International Conference on Machine Learning*, pp. 3368–3376, PMLR, 2017.
- [53] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," in *International Conference on Machine Learning*, pp. 5610–5619, PMLR, 2018.
- [54] 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)*, pp. 2027–2031, IEEE, 2018.
- [55] S. Dutta, V. Cadambe, and P. Grover, "Short-dot: Computing large linear transforms distributedly using coded short dot products," *Advances In Neural Information Processing Systems*, vol. 29, 2016.
- [56] A. Ramamoorthy, L. Tang, and P. O. Vontobel, "Universally decodable matrices for distributed matrix-vector multiplication," in *2019 IEEE International Symposium on Information Theory (ISIT)*, pp. 1777–1781, IEEE, 2019.
- [57] A. B. Das and A. Ramamoorthy, "Distributed matrix-vector multiplication: A convolutional coding approach," in *2019 IEEE International Symposium on Information Theory (ISIT)*, pp. 3022–3026, IEEE, 2019.
- [58] F. Haddadpour and V. R. Cadambe, "Codes for distributed finite alphabet matrix-vector multiplication," in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 1625–1629, IEEE, 2018.

- [59] S. Wang, J. Liu, and N. Shroff, "Coded sparse matrix multiplication," in *International Conference on Machine Learning*, pp. 5152–5160, PMLR, 2018.
- [60] A. Ramamoorthy, A. B. Das, and L. Tang, "Straggler-resistant distributed matrix computation via coding theory: Removing a bottleneck in large-scale data processing," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 136–145, 2020.
- [61] V. M. Blinovskii, "Lower asymptotic bound on the number of linear code words in a sphere of given radius in $(f_q)^n$," *Problemy Peredachi Informatsii*, vol. 23, no. 2, pp. 50–53, 1987.
- [62] G. Cohen, I. Honkala, S. Litsyn, and A. Lobstein, *Covering codes*. Elsevier, 1997.
- [63] R. M. Roth, "Introduction to coding theory," *IET Communications*, vol. 47, 2006.
- [64] G. Cohen and P. Frankl, "Good coverings of hamming spaces with spheres," *Discrete Mathematics*, vol. 56, no. 2-3, pp. 125–131, 1985.
- [65] K. Wan, H. Sun, M. Ji, and G. Caire, "On the tradeoff between computation and communication costs for distributed linearly separable computation," *IEEE Transactions on Communications*, vol. 69, no. 11, pp. 7390–7405, 2021.
- [66] K. Wan, H. Sun, M. Ji, and G. Caire, "On secure distributed linearly separable computation," *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 3, pp. 912–926, 2022.
- [67] A. K. Das and S. Vishwanath, "On finite alphabet compressive sensing," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 5890–5894, IEEE, 2013.
- [68] A. G. Dimakis, R. Smarandache, and P. O. Vontobel, "Ldpc codes for compressed sensing," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3093–3114, 2012.
- [69] A. Reiszadeh, S. Prakash, R. Pedarsani, and A. S. Avestimehr, "Tree gradient coding," in *2019 IEEE International Symposium on Information Theory (ISIT)*, pp. 2808–2812, IEEE, 2019.
- [70] G. Cohen, "A nonconstructive upper bound on covering radius," *IEEE Transactions on Information Theory*, vol. 29, no. 3, pp. 352–353, 1983.
- [71] A. N. Kolmogoroff, *Foundations of the Theory of Probability*. Chelsea Publishing Company, 1956.
- [72] H. E. Danoyan, "On some properties of intersection and union of spheres in hamming metric," *Mathematical Problems of Computer Science*, vol. 39, pp. 119–124, 2013.

Ali Khalesi was born in Glasgow, UK, in 1996. He received the bachelor's degree (Hons.) in electrical engineering from K. N. Toosi University of Technology, in 2018, and the master's degree (Hons.) in secure communications & cryptography from the Sharif University of Technology in 2020. He is currently pursuing the Doctoral degree from Sorbonne University with the Communication Systems Department of EURECOM Institute. His research interests include coding theory, information theory, Cryptography, and distributed computing.

Petros Elia (M '09) received the B.Sc. degree from the Illinois Institute of Technology and the M.Sc. and Ph.D. degrees in electrical engineering from the University of Southern California (USC), Los Angeles, in 2001 and 2006, respectively. He is currently a Professor with the Department of Communication Systems, EURECOM, Sophia Antipolis, France. His latest research deals with distributed computing and with the intersection of caching and communications in multi-user settings. He has also worked in the area of complexity-constrained communications, MIMO, queueing theory and cross-layer design, coding theory, information-theoretic limits in cooperative communications, and surveillance networks. He is a Fulbright Scholar, a co-recipient of the NEWCOM++ Distinguished Achievement Award (2008–2011) for a sequence of publications on the topic of complexity in wireless communications, and a recipient of the ERC Consolidator Grant (2017–2022) on cache-aided wireless communications.