

# List-Decodable Coded Computing: Breaking the Adversarial Toleration Barrier

Mahdi Soleymani, Ramy E. Ali, Hessam Mahdavifar, and A. Salman Avestimehr

**Abstract**—We consider the problem of coded computing, where a computational task is performed in a distributed fashion in the presence of adversarial workers. We propose techniques to break the adversarial toleration threshold barrier previously known in coded computing. More specifically, we leverage list-decoding techniques for folded Reed-Solomon codes and propose novel algorithms to recover the correct codeword using side information. In the coded computing setting, we show how the master node can perform certain carefully designed extra computations to obtain the side information. The workload of computing this side information is negligible compared to the computations done by each worker. This side information is then utilized to prune the output of the list decoder and uniquely recover the true outcome. We further propose folded Lagrange coded computing (FLCC) to incorporate the developed techniques into a specific coded computing setting. Our results show that FLCC outperforms LCC by breaking the barrier on the number of adversaries that can be tolerated. In particular, the corresponding threshold in FLCC is improved by a factor of two asymptotically compared to that of LCC.

**Index Terms**—Coded computing, secure computing, Byzantine adversaries, list-decoding, folded Reed-Solomon codes.

## I. INTRODUCTION

Recently, ideas from the coding theory literature have been widely leveraged in large-scale distributed computing and learning problems to alleviate major performance bottlenecks including latency in computations, communication overheads, and stragglers [1]–[4]. This has led to the emergence of the *coded computing* paradigm by combining coding theory and distributed computing, also addressing critical issues such as security and privacy in distributed settings. More specifically, there has been an increasing interest in recent years toward adopting coded computing techniques in computationally-demanding machine learning tasks that give rise to several privacy and security issues [5]–[10]. In such settings, the underlying dataset must remain private from the cloud and the contributing computational workers, as it may contain highly sensitive information such as biometric data of patients in a

This material is based upon work supported by Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001117C0053, ARO award W911NF1810400, NSF grants CCF-1703575, CCF-1763673, CCF-1763348, CCF-1909771 and CCF-1941633, MLWINS-2002874, ONR Award No. N00014-16-1-2189, and a gift from Intel/Avast/Borsetta via the PrivateAI institute. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

This article was presented in part at the proceedings of the 2021 IEEE International Symposium on Information Theory. The associate editor coordinating the review of this article and approving it for publication was P. Grover. (Corresponding author: Mahdi Soleymani.)

M. Soleymani and H. Mahdavifar are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48104 (email: mahdy@umich.edu and hessam@umich.edu).

R.E. Ali and A. Salman Avestimehr are with the Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089 USA (e-mail: reali@usc.edu and avestimehr@ee.usc.edu).

hospital [11] or customers’ data of a company [12]. Moreover, the outcome of a distributed learning scheme, e.g., model parameters trained on a dataset, must be *secured* against Byzantine (malicious) adversaries that attack the cloud or are present as adversarial workers aiming at altering the outcome either for their benefits or to deceive the other users.

A well-established architecture often considered in coded computing consists of a *master* node and a set of workers having communication links with the master node. The goal for the master is to perform a certain computational job, e.g., training a model on its dataset, with the help of the workers. To this end, the master disperses its dataset among the workers that operate in parallel and return their results to the master to recover the outcome of the computational job efficiently. Then a problem of significant interest is the following: what fraction of adversarial workers can be tolerated, i.e., the master is still able to recover the true outcome even though the adversaries have returned corrupted results, in such coded computing schemes? To answer this question, the well-known classical results on the error correction capability of linear codes and, in particular, maximum distance separable (MDS) codes such as Reed-Solomon (RS) codes are leveraged, see, e.g., [3]. The tightness of such results on adversarial toleration is based on certain assumptions on the underlying code and the corresponding decoder employed by the master node. For instance, it is implicitly assumed that the master node performs the decoding only given the returned results by the workers and does not perform any extra computations to gain *side information* about the computation outcome. Also, the decoder employed by the master is assumed to be the classical decoder that recovers errors up to half the minimum distance bound. However, the list-decoding paradigm offers the potential to decode errors beyond this bound [13]. In fact, there is a long history on list-decoding algorithms for RS codes with the end result of achieving the information-theoretic Singleton bound  $1 - R$ , where  $R$  is the code rate, on the decoding radius for a variant of RS codes, called *folded* RS (FRS) codes [14]–[17]. This improves upon the half the minimum distance bound, expressed as  $(1 - R)/2$  for MDS codes, by a factor of 2 closing the gap with the Singleton bound [17].

We consider the following fundamental question in this paper: is it possible to break the adversarial toleration threshold barrier established in the coded computing literature? We show that the answer to this question is yes. To this end, we leverage the advances in the list-decoding literature as well as the particular coded computing setting that naturally allows the master node to have access to side information and uniquely determine the computation outcome.

### A. Our contributions

In this paper, we show how to adapt the *folding* technique in algebraic coding to the realm of coded computing, where the

underlying computational job is a polynomial evaluation over the dataset. Then it is shown how the master node can employ an off-the-shelf FRS list-decoding algorithm, e.g., the linear-algebraic algorithm proposed in [18], to the results returned from the workers. This results in a low-dimensional linear subspace that contains the true outcome of the computation assuming a certain bound on the number of adversaries. In order to uniquely recover the true outcome, we propose two schemes which involve the master node performing certain carefully designed extra computations to obtain side information about the outcome. This side information is then utilized to prune the subspace of possible outcomes, i.e., the output of the FRS list-decoding algorithm, to uniquely recover the true outcome. In both schemes, the cost of computing the side information is negligible compared to the computation load of each worker. Specifically, our contributions are as follows.

- 1) We propose a deterministic pruning algorithm, in which the master node waits until the results are returned by the worker nodes and the FRS list-decoding algorithm is applied to the returned results. Then the master node carefully selects a certain small subset of evaluation points and computes the polynomial evaluation over these points to obtain the side information. It is shown that this can be done in such a way that the true outcome is uniquely recovered from the output list.
- 2) We also propose a probabilistic pruning algorithm in which the side information is obtained by computing the polynomial evaluation over a randomly selected set of evaluation points. This can be done in parallel to the tasks being performed by the workers resulting in a lower latency compared to the first approach. Then it is shown that the true outcome can be uniquely recovered with a *high* probability. Moreover, if the unique recovery is not possible in this case, the master node can identify it as a decoding failure. We show results outperforming the state-of-the-art schemes in terms of the lower bound on the probability of successful decoding and the amount of side information needed for unique decoding.
- 3) To illustrate how our proposed protocols break the adversarial toleration thresholds in a certain class of coded computing schemes, we consider the Lagrange coded computing (LCC) scheme. We introduce a *folded* version of LCC, referred to as FLCC. Similar to other mainstream coded computing schemes, the master node in LCC attempts to decode the computation outcome merely based on the computations of the workers. By relaxing this restriction, the master node in FLCC invokes our proposed pruning algorithms together with list-decoding FRS codes and uniquely recovers the outcome. The performance of FLCC with both the deterministic and the probabilistic pruning algorithms is characterized and compared to that of LCC. Our results indicate that the cost of overcoming a Byzantine worker in FLCC can be reduced to be almost the same as that of a straggler worker, in the characterization of recovery thresholds, as opposed to LCC in which Byzantine adversaries cost twice as stragglers.

## B. Related work

The problem of list-decoding with side information was initially considered in [19] for binary codes in a communication setup. In this setting, a clean noise-free channel is assumed over which a *small* amount of side information, compared to the length of the message, is provided to the receiver. This side information consists of a random hash function along with its value over the message. Another variant of RS codes, known as *derivative* codes that also achieve the *optimal* performance have been studied in [18], [20]. In [18], a linear-algebraic list-decoding approach along with pruning algorithms with side information specific to derivative codes were proposed to uniquely recover the codeword either deterministically or with high probability. This approach also can list-decode FRS codes with side information to recover the codeword with high probability. However, this approach was not extended to deterministically recover the codeword. The reason for this is that to deterministically recover the codeword using this approach, the side information may not be decided before starting to decode. While this is inconvenient in communication setups, it is possible in coded computing as the master node plays the role of both the encoder and the decoder. List-decoding of FRS codes has been also incorporated in the context of secret sharing to enhance the security [21]–[23]. However, these works do not consider computations over data and are only concerned with recovering the data from the secret shares.

Coded computing has recently gained much interest due to its promise to overcome several issues raised in large-scale distributed computing and machine learning. It has been utilized for straggler mitigation in various distributed computing tasks [24]–[30]. Several schemes for distributed matrix-matrix multiplication, which is one of the main building blocks for various machine learning algorithms, have been also proposed in the literature [31]–[35]. Moreover, certain protocols have been introduced for computations over real-valued data [36]–[39]. Also, recently, analog coded computing protocols have been introduced to enable privacy for large-scale distributed machine learning in the analog domain [8], [40]. Improving the adversarial toleration threshold of LCC has been also considered recently in [41] for Boolean computations and sparse polynomials and in [42] for matrix-matrix multiplication. In addition, this problem was also studied for the probabilistic noise model in [43], [44]. Our work, however, considers any polynomial-based computations and the worst-case adversarial model. Moreover, none of these prior works incorporated list-decoding ideas into the coded computing protocols.

The rest of this paper is organized as follows. In Section II, some background on list-decoding of RS codes and their variants is provided. The system model considered in this paper is discussed in Section III and FLCC is proposed. Our results on list-decoding of FRS codes with side information are shown in Section IV. In Section V it is shown how our techniques applied to FLCC improve upon the security of LCC against Byzantine adversaries. Finally, the paper is concluded in Section VI.

## II. BACKGROUND

In this section, we first provide a brief background on list-decoding. Then, we briefly overview Lagrange coded comput-

ing (LCC) that will be used later in Section V to showcase how the list decoding ideas are leveraged to improve the adversarial toleration threshold in a well-established coded computing scheme.

### A. List decoding FRS codes

We begin by introducing the notations that are used throughout this paper. For a positive integer  $i$ , the set  $\{1, 2, \dots, i\}$  is denoted by  $[i]$ . The number of positions at which two strings of length  $n$ ,  $\mathbf{y}$  and  $\mathbf{y}'$ , differ is denoted by  $\Delta(\mathbf{y}, \mathbf{y}') \stackrel{\text{def}}{=} |\{i : y_i \neq y'_i\}|$  and their relative distance is denoted by  $\delta(\mathbf{y}, \mathbf{y}') \stackrel{\text{def}}{=} \Delta(\mathbf{y}, \mathbf{y}')/n$ . A code  $C : \Sigma^k \rightarrow \Sigma^n$  of length  $n$  over alphabet  $\Sigma$ , where  $|\Sigma| = q$ , is denoted by  $[n, k]_q$ . A finite field of size  $q$  is denoted by  $\mathbb{F}_q$  and the set of all non-zero elements of  $\mathbb{F}_q$  is represented by  $\mathbb{F}_q^*$ .

An  $[n, k]_q$  MDS code such as a Reed-Solomon code can always correct up to  $\rho_U(R) = (1 - R)/2$  normalized number of errors, also referred to as decoding radius, where  $R \stackrel{\text{def}}{=} k/n$  is the rate of the code and normalization is done by dividing the number of errors by  $n$ . In order to correct errors beyond this bound, list-decoding [13], [14], [45], [46] relaxes the unique decoding requirement and allows the decoder to output a list of codewords. Specifically, given  $0 \leq \rho \leq 1$ , an  $[n, k]_q$  code  $C \subseteq \Sigma^n$  is said to be  $(\rho, L)$ -list decodable if for every  $\mathbf{y} \in \Sigma^n$ , the set  $\mathcal{L} \stackrel{\text{def}}{=} \{c \in C | \delta(\mathbf{y}, c) \leq \rho n\}$  has at most  $L$  elements. Based on this relaxation, Guruswami and Sudan [15] showed that Reed-Solomon codes can be list-decoded up to the decoding radius of  $\rho_{GS}(R) = 1 - \sqrt{R}$ . However, it is well-known that there exist codes that can be list-decoded up to a decoding radius of  $1 - R - \epsilon$  with a list size of at most  $O(1/\epsilon)$  [47], [48]. Parvaresh and Vardy further improved Guruswami-Sudan decoding radius by introducing a variant of Reed-Solomon codes, also referred to as Parvaresh-Vardy codes [16], followed by Guruswami and Rudra [17] who showed that such variations can be more efficiently realized by folding the Reed-Solomon codes, thereby improving the decoding radius to approach the ultimate Singleton bound  $1 - R$ . Next, the definition of folded Reed-Solomon codes is provided.

*Definition 1:* ( $m$ -Folded Reed-Solomon Code [17]) Let  $\gamma$  be a primitive element of  $\mathbb{F}_q$ ,  $n \leq q - 1$  be a multiple of  $m$ , and  $k$  with  $1 \leq k < n$  be the degree parameter. The folded Reed-Solomon (FRS) code  $\text{FRS}_q^{(m)}[n, k]$  is a code over alphabet  $\mathbb{F}_q^m$  that encodes a polynomial  $f \in \mathbb{F}_q[X]$  of degree at most  $k - 1$  as

$$\left( \begin{bmatrix} f(1) \\ f(\gamma) \\ \vdots \\ f(\gamma^{(m-1)}) \end{bmatrix}, \begin{bmatrix} f(\gamma^m) \\ f(\gamma^{m+1}) \\ \vdots \\ f(\gamma^{(2m-1)}) \end{bmatrix}, \dots, \begin{bmatrix} f(\gamma^{n-m}) \\ f(\gamma^{n-m+1}) \\ \vdots \\ f(\gamma^{(n-1)}) \end{bmatrix} \right),$$

where the block length of  $\text{FRS}_q^{(m)}[n, k]$  is  $N = n/m$ , and its rate is  $R = k/n$ .

Guruswami and Rudra [17] showed that FRS codes can be efficiently list-decoded up to the decoding radius  $\rho_{GR}(R) = 1 - R - \epsilon$  with a list size of  $L = n^{O(1/\epsilon)}$ . Later, it was shown in [18] that this can be done using an alternative linear-algebraic approach. Next, we recall this result.

*Lemma 1 (List-Decoding of FRS codes [18]):* For the FRS code  $\text{FRS}_q^{(m)}$  of block length  $N = n/m$  and rate  $R = k/n$ , the following holds for all integers  $s \in [m]$ . Given a received word  $\mathbf{y} \in (\mathbb{F}_q^m)^N$ , using  $O(n^2 + sk^2)$  operations over  $\mathbb{F}_q$ , one can find a subspace of dimension at most  $s - 1$  that contains all encoding polynomials  $f \in \mathbb{F}_q[X]$  of degree less than  $k$  whose FRS encoding differs from  $\mathbf{y}$  in at most a fraction

$$\frac{s}{s+1} \left( 1 - \frac{mR}{m-s+1} \right) \quad (1)$$

of the  $N$  codeword positions.

Note that choosing  $s = m = 1$  in Lemma 1 corresponds to the unique decoding radius of  $(1 - R)/2$ , while choosing  $s \approx 1/\epsilon$  and  $m \approx 1/\epsilon^2$  ensures a decoding radius of  $1 - R - \epsilon$ .

We note that the work of Guruswami and Wang [49] only considers errors and not symbol erasures. In other words, it is assumed that all coordinates of the received word are either error-free or corrupt and none of them is erased during the transmission. However, it can be observed that the same algorithm also works if  $S$  number of coordinates in  $\mathbf{y}$  are erased, as provided in the following lemma.

*Lemma 2:* The result of Lemma 1 still holds if  $N$  is replaced by  $N - S$  for the case where  $S$  out of  $N$  symbols are erased.

We provide the proof of Lemma 2 in the Appendix.

While list-decoding allows for correcting more errors as compared to unique decoding, it is often required to output a unique codeword in certain applications. If the decoder has access to some noise-free side information, then it can use it to prune the list and output a unique codeword. This problem was studied in [19] in a communication setup. More specifically, a probabilistic scheme was proposed in [19] in which the transmitter sends a random error-free symbol generated by a random hash function along with this hash function as a side information. The receiver then checks whether there is a unique message in the list that is consistent with the side information or not. If yes, the receiver outputs this message. Otherwise, the receiver declares a decoding failure.

Guruswami and Wang [18] also developed an alternative linear-algebraic list-decoding approach with side information for derivative codes. In this approach, a subspace of candidate polynomials is first determined. The unique message can be then found with high probability by pruning this subspace using the side information. The main advantage of this approach compared to the hashing approach developed in [19] is that the decoder does not need to compute the full list, which may have an exponential size in  $s$ , and then prune it to get the unique solution. Kopparty *et. al* [50] also provide a randomized decoding algorithm for FRS codes that returns a constant-size list which is much smaller than the number of candidate polynomials in the subspace returned by the approach of Guruswami-Wang described in Lemma 1. As a subroutine, one can prune the list by looking at one random symbol of FRS code, i.e.,  $(f(a), f(a\gamma), \dots, f(a\gamma^{m-1}))$  for a random  $a \in \mathbb{F}_q^*$ , and recover the true polynomial with high probability. However, using a similar approach and following the result of [51, Lemma 12], one can show that the unique polynomial can

be decoded successfully with a probability  $p_d$  that is at least

$$p_d \geq 1 - \frac{k(s-1)}{q}, \quad (2)$$

using  $s-1$  extra evaluations instead of  $m$  evaluations. However, this approach requires the field size to be larger than a certain threshold, i.e.,  $q > k(s-1)$ . To avoid such a constraint for the case where the field size is not relatively *large*, the result of [52, Lemma 2] implies that by using a similar approach of [50] the unique codeword can be found with probability at least

$$p_d \geq 1 - \binom{t}{l-1} \left(\frac{k}{n}\right)^{t-l+1}, \quad (3)$$

where  $l < s$  is the dimension of subspace returned by the list decoding algorithm provided that  $t$  extra evaluations  $f(a_1), f(a_2), \dots, f(a_t)$  for randomly picked  $a_i \in \mathbb{F}_q^*$  with  $a_i \neq a_j$  for all  $i, j \in [t]$  are available at the decoder.

In this work, we provide a new lower bound on the probability of successful decoding of FRS codes using at most  $s-1$  extra evaluations of  $f(\cdot)$  and compare our result with the aforementioned existing works in Section IV.

### B. Lagrange coded computing

Consider a coded computing setup consisting of a master node and a set of  $N$  workers. We consider the worst-case adversarial model with up to  $A$  computationally-unbounded Byzantine (or malicious) adversarial workers, up to  $S$  stragglers and a privacy model where any set of up to  $T$  workers can collude. Let  $(\mathbf{X}_1, \dots, \mathbf{X}_K)$  denote a batch of  $r \times h$  matrices over  $\mathbb{F}_q$ . The goal is to compute a  $D_2$ -degree polynomial function  $g(\cdot) : \mathbb{F}_q^{r \times h} \rightarrow \mathbb{F}_q^{r' \times h'}$ , over this dataset, i.e.,  $g(\mathbf{X}_i)$  for all  $i \in [K]$ , where  $r'$  and  $h'$  are the dimensions of the output matrix. More specifically, we say  $g(\cdot)$  is a  $D_2$ -degree polynomial function if all entries of the output matrix are multivariate polynomial functions of the entries of the input with total degree at most  $D_2$ , i.e.,  $\mathbf{Y} = g(\mathbf{X})$  implies that

$$y_{ij} = g_{ij}(x_{11}, x_{12}, \dots, x_{rh}), \quad (4)$$

where  $y_{ij}$  is the  $(i, j)$ -th entry of  $\mathbf{Y}$ , for  $i \in [r']$  and  $j \in [h']$ ,  $x_{lk}$  is the  $(l, k)$  entry of  $\mathbf{X}$ , for  $l \in [r]$  and  $k \in [h]$ , and,  $g_{ij}$  is a multivariate polynomial of total degree at most  $D_2$ . Let  $E \stackrel{\text{def}}{=} \{\alpha_1, \dots, \alpha_N\}$  for some distinct elements  $\alpha_1, \dots, \alpha_N \in \mathbb{F}_q$  and  $I \stackrel{\text{def}}{=} \{\beta_1, \dots, \beta_{K+T}\}$  for some other distinct elements  $\beta_1, \dots, \beta_{K+T} \in \mathbb{F}_q$ . The sets  $E$  and  $I$  are referred to as the set of evaluation points and the set of interpolation points, respectively. Note that  $E$  and  $I$  do not intersect, i.e.,  $E \cap I = \emptyset$ . The underlying encoding polynomial in LCC is the Lagrange interpolation polynomial of degree  $D_1 = K+T-1$  constructed as

$$u(z) = \sum_{j=1}^K \mathbf{X}_j \ell_j(z) + \sum_{j=K+1}^{K+T} \mathbf{Z}_j \ell_j(z), \quad (5)$$

where  $\mathbf{Z}_j$ 's for  $j \in \{K+1, \dots, K+T\}$  are random matrices whose entries are independent and uniformly distributed over  $\mathbb{F}_q$  and  $\ell_j(z)$ 's are called Lagrange monomials specified as

$$\ell_j(z) = \prod_{l \in [K+T] \setminus \{j\}} \frac{z - \beta_l}{\beta_j - \beta_l}, \quad (6)$$

for  $j \in [K+T]$ . The master node offloads the coded matrix  $\tilde{\mathbf{X}}_i \stackrel{\text{def}}{=} u(\alpha_i)$  to the worker node  $i$  whose task is to compute  $\tilde{\mathbf{Y}}_i \stackrel{\text{def}}{=} g(u(\alpha_i)) = g(\tilde{\mathbf{X}}_i)$ . The composed polynomial  $f(z) \stackrel{\text{def}}{=} g(u(z))$  can be recovered provided that at least  $D_1 D_2 + 2A + 1$  workers return their computation results, i.e.,  $\tilde{\mathbf{Y}}_i$ 's, to the master node. This can be done by invoking Berlekamp-Welch (BW) decoder [53] individually for all entries of the output matrix of the polynomial function  $f(z)$ . Note that BW decoder can successfully reconstruct a polynomial of degree  $D_1 D_2$  provided that  $D_1 D_2 + 2A + 1$  evaluations of the polynomial are available and up to  $A$  of the evaluations can be erroneous. Note that since  $\tilde{\mathbf{X}}_i = u(\beta_j)$  for all  $j \in [K]$ , then  $f(\beta_j) = g(u(\beta_j)) = g(\tilde{\mathbf{X}}_j)$ . The master then evaluates  $g(\cdot)$  over the interpolation points  $\beta_1, \dots, \beta_K$  to recover the desired computation outcome, i.e.,  $g(\mathbf{X}_1), \dots, g(\mathbf{X}_K)$ . We say that LCC is  $S$ -resilient and  $A$ -secure if it is robust against  $S$  stragglers and  $A$  Byzantine adversaries, respectively, and that it is  $T$ -private if any set of size up to  $T$  workers remain oblivious to the content of dataset. It is shown in [3] that the number of required workers for an  $S$ -resilient,  $A$ -secure, and  $T$ -private LCC to compute  $\{g(\mathbf{X}_i)\}_{i=1}^K$  for a  $D_2$ -degree polynomial  $g(\cdot)$  is lower bounded as

$$N \geq (K+T-1)D_2 + S + 2A + 1, \quad (7)$$

where  $D = (K+T-1)D_2$  is the degree of the composed polynomial  $f(\cdot)$  to be interpolated at the master node during the decoding step. The lower bound provided in (7) implies that tolerating Byzantine adversaries in LCC is twice as costly as stragglers, i.e., the additional number of workers required to tolerate each Byzantine worker is equal to what is needed to tolerate two stragglers.

### III. SYSTEM MODEL

We consider a coded computing setup consisting of a master node and a set of  $N$  workers. The goal is to design a coded computing scheme that is  $S$ -resilient,  $A$ -secure and  $T$ -private, where we consider the worst-case adversarial model with up to  $A$  computationally-unbounded Byzantine adversarial workers. Unlike LCC [54], where the master node does not do any computations except encoding and decoding, we allow the master node to do a tiny amount of extra computations that is negligible compared to the computations of each worker. Specifically, the normalized extra computation cost of the master node with respect to the workload of each worker must go to 0 as the workload of each worker grows.

A well-known class of coded computing schemes extensively studied in the literature for this setup employs polynomial evaluations to encode data. We refer to them as *polynomial-based* coded computing schemes. In such schemes, the shares sent to the worker nodes are evaluations of a certain polynomial over a finite field  $\mathbb{F}_q$ . The worker nodes perform a predefined computation task over their share(s), e.g., polynomial evaluation, matrix multiplication, etc., and return the results to the master node. The master node then follows a decoding process involving a polynomial interpolation to recover the overall computation outcome. We denote the  $D$ -degree polynomial to be interpolated at the decoding step by  $f(\cdot)$ . Such a class of

coded computing schemes includes, but is not limited to, LCC [3], polynomial codes [55] and MatDot codes [56]. It is well-known that the polynomial  $f(\cdot)$  can be uniquely recovered provided that up to  $(N - D - 1)/2$  evaluations, out of  $N$  available evaluation points, are erroneous. This can be done by utilizing efficient Reed-Solomon decoding algorithms at the master node. This, in a high level, imposes a limit on the maximum number of Byzantine workers that can be tolerated in prior works on polynomial-based coded computing. In this work, we break this barrier by employing folded Reed-Solomon (FRS) codes instead of RS codes often used in the polynomial-based schemes together with leveraging their list-decoding algorithms instead of the unique decoding algorithm for RS codes. The output of the list decoder is then pruned using a tiny amount of side information to recover the unique computation outcome. The computation cost of this side information is negligible compared to the computation cost of each worker. This will be clarified more later in Section V.

In order to illustrate the key idea of our method more explicitly, we consider the LCC scheme [3] described earlier in Section II-B. We propose a variant of LCC, referred to as folded LCC (FLCC), inspired by the folded RS code construction. Consider  $m$  batches of size  $K$  which can be considered as a larger batch of size  $k \stackrel{\text{def}}{=} mK$ , i.e.,  $(\mathbf{X}_1, \dots, \mathbf{X}_k)$ , where the goal is to compute  $g(\mathbf{X}_i)$  for all  $i \in [k]$ . The parameter  $m$  is an arbitrary integer, referred to as the folding parameter. We also assume  $N = (K + T - 1)D_2 + S + 2A + 1$  which is the minimum number of workers required in LCC. Let  $E_m \stackrel{\text{def}}{=} \{\alpha_i = \alpha^{i-1} | i \in [N]\}$  denote the set of evaluation points in our proposed scheme, where  $\alpha$  is a primitive element of  $\mathbb{F}_q$  and  $q > N$ . Note that the evaluation points here are picked more specifically compared to those of LCC, but that does not impose any limitations on the scheme. Furthermore, instead of the polynomial in (5), we construct the following encoding polynomial

$$u_m(z) = \sum_{j=1}^{Km} \mathbf{X}_j \ell_j(z) + \sum_{j=Km+1}^{(K+T)m} \mathbf{Z}_j \ell_j(z), \quad (8)$$

where  $\mathbf{Z}_j$ 's for  $j \in \{Km + 1, \dots, (K + T)m\}$  are random matrices whose entries are independent and uniformly distributed over  $\mathbb{F}_q$  and  $\ell_j(z)$ 's are Lagrange monomials defined as

$$\ell_j(z) = \prod_{l \in [(K+T)m] \setminus j} \frac{z - \beta_l}{\beta_j - \beta_l}, \quad (9)$$

where  $I_m \stackrel{\text{def}}{=} \{\beta_1, \dots, \beta_{m(K+T)}\}$  for distinct  $\beta_i$ 's belonging to  $\mathbb{F}_q$  such that  $E_m \cap I_m = \emptyset$ . We refer to this scheme as FLCC. In FLCC, the share of encoded data sent to the worker  $i$  consists of the evaluations of  $u_m(\cdot)$  over the points  $\alpha_{m(i-1)+1}, \dots, \alpha_{mi}$ , i.e.,  $u_m(\alpha_{m(i-1)+1}), \dots, u_m(\alpha_{mi})$ . Let  $f_m(z) \stackrel{\text{def}}{=} g(u_m(z))$  denote the composed polynomial to be interpolated at the decoder in our scheme. The task of each worker node is then to compute  $g(\cdot)$  on all of its associated evaluations separately, i.e., node  $i$  computes  $f_m(\alpha_{m(i-1)+1}), \dots, f_m(\alpha_{mi})$  and returns the results to the master node.

Intuitively speaking, the encoding polynomial considered in our scheme is similar to that of LCC in which  $m$  batches of

data are regarded as a single dataset of size  $m$ -times larger. In other words, in the encoding step of our protocol, we first encode the data according to the RS encoder and the coded symbols are then *folded* with parameter  $m$ , resembling the FRS encoding procedure. Consequently, we can apply list-decoding algorithms developed for FRS in the literature to gain a better resiliency-security-privacy trade-off in FLCC compared to that of LCC. A linear-algebraic list-decoding algorithm for FRS codes is considered in the next section along with our proposed methods to improve it when certain side information is available at decoder. We then discuss in Section V how this result can be utilized to improve upon the performance of LCC decoder in terms of the achievable triples  $(S, A, T)$ .

#### IV. LIST DECODING FRS CODES WITH SIDE INFORMATION

In this section, we provide our approach to adapt list-decoding techniques to make the polynomial-based coded computing protocols more *robust* against malicious adversaries. This is done in such a way that it can be used as a *black box* and regardless of the technical details associated with the encoder and decoder of the underlying coded computing scheme. We then illustrate in Section V how the proposed techniques can be applied to LCC.

In particular, we consider the linear-algebraic FRS list decoder introduced in [49]. Let  $\mathcal{W}$  denote the linear space of univariate polynomials of degree at most  $k - 1$  over  $\mathbb{F}_q$  and  $\mathcal{L}$  denote the list of candidate polynomials at the output of the list decoder that can be represented by an affine subspace  $U$ . The elements of  $U$  can be represented as  $f = \mathbf{M}\mathbf{x} + \mathbf{z}$  for  $\mathbf{x} \in \mathbb{F}_q^l$ , where  $l < s$ ,  $\mathbf{M} \in \mathbb{F}_q^{k \times l}$  and  $\mathbf{z} \in \mathbb{F}_q^k$ . The vector  $\mathbf{f} = (f_0, f_1, \dots, f_{k-1})^T$  denotes the coefficients of the corresponding polynomial in  $U$ . It is also shown in [49] that  $\mathbf{M}$  can be assumed to have  $l \times l$  identity matrix  $\mathbf{I}_l$  as a submatrix, without any extra computation. Note that the location of the identity submatrix is not known prior to applying the list-decoding algorithm at the decoder.

Now, suppose that the decoder of the FRS code can request to have access to  $l < s$  additional error-free evaluations of  $f(\cdot)$  as a side information. This can be done with the aim of pruning the output of the list-decoding algorithm specified in Lemma 1 to uniquely recover  $f(\cdot)$ . In the next theorem, we provide a result that this is always possible provided that the  $l < s$  evaluation points can be decided after applying the list-decoding algorithm to the received word  $\mathbf{y}$ .

*Theorem 1:* For the FRS code of length  $N = \frac{n}{m}$  and rate  $R = \frac{k}{n}$  and for all  $s \in [m]$ , the polynomial  $f(\cdot)$  can be uniquely recovered if

- i) the received word  $\mathbf{y} \in (\mathbb{F}_q^m)^N$  differs from the FRS code-word corresponding to  $f(\cdot)$  in at most  $\frac{s}{s+1}(1 - \frac{mR}{m-s+1})$  fraction of the  $N$  symbols and
- ii) up to  $s - 1$  additional evaluations of  $f(\cdot)$  can be requested and are provided error-free, and assuming that the corresponding evaluation points can be decided after  $\mathbf{y}$  is received.

Moreover, the entire algorithm is run with  $O(n^2 + sk^2)$  complexity.

*Proof:* Let  $\tilde{V} \in \mathbb{F}_q^{n \times k}$  be an arbitrary Vandermonde matrix. Let also the affine subspace  $U$  that contains the true polynomial be represented by  $f = Mx + z$  for  $x \in \mathbb{F}_q^l$ , where  $l < s$ ,  $M \in \mathbb{F}_q^{k \times l}$  and  $z \in \mathbb{F}_q^k$ , as discussed above. Then,  $\tilde{V}M$  is full rank since both  $\tilde{V}$  and  $M$  are full rank and  $n \geq k \geq l$ . By using the fast multiplication algorithm available for Vandermonde matrices [57],  $\tilde{V}M$  can be formed in  $O(nl \log n)$ . Furthermore, an  $l \times l$  full-rank submatrix of  $\tilde{V}M$  can be found in time  $O(nl^2)$  by Gaussian elimination. Then, the rows in  $\tilde{V}$  associated to this  $l \times l$  submatrix form an  $l \times k$  Vandermonde matrix, namely,

$$V \stackrel{\text{def}}{=} \begin{bmatrix} 1 & \lambda_1 & \cdots & \lambda_1^{k-1} \\ 1 & \lambda_2 & \cdots & \lambda_2^{k-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \lambda_l & \cdots & \lambda_l^{k-1} \end{bmatrix}, \quad (10)$$

such that  $\tilde{V}M$  is full rank. The evaluations of  $f(\cdot)$  over the points  $\lambda_1, \dots, \lambda_l$  associated to the Vandermonde matrix suffice to uniquely recover  $f(\cdot)$ . To see that, let  $y_e \stackrel{\text{def}}{=} Vf$  denote the vector of extra evaluations of  $f(\cdot)$  over  $\lambda_1, \dots, \lambda_l$ . Therefore, one can write

$$y_e = VMx + Vz. \quad (11)$$

This implies that if  $VM$  is full rank, then (11) can be solved for  $x$  which is then utilized to determine  $f$ , thereby uniquely recovering the polynomial  $f(\cdot)$ . The computational complexity of the entire algorithm is dominated by that of the list-decoding algorithm described in Lemma 1 which is  $O(n^2 + sk^2)$  by noting that  $l < s$ . ■

The result of Theorem 1 implies that if the set of extra evaluation points  $\Lambda \stackrel{\text{def}}{=} \{\lambda_1, \dots, \lambda_l\}$  can be decided by the decoder after observing the entire received vector, the output subspace provided by the list-decoding algorithm, as specified in Lemma 1, can be efficiently pruned to uniquely recover the polynomial  $f(\cdot)$  in a deterministic fashion. This provides a pruning algorithm with a computational complexity that is dominated by that of the corresponding list-decoding algorithm. Note that in a communication setting, with the encoder and the decoder being separate entities, such an assumption on the side information would necessitate multiple rounds of communication that may not be desirable in practice. However, in coded computing settings, the encoding and decoding are both done by the same entity, i.e., the master node. Hence, obtaining error-free side information after the results are received from the workers does not impose any major hurdle to the protocol. This comes only at the cost of extra computation complexity and latency, which can be characterized and optimized based on the limitations of the master node.

In order to mitigate the latency of computing the side information, we propose an alternative probabilistic algorithm. In this algorithm, the master node does not have to wait till the results are received from the workers and can compute the side information in parallel to them. This is described in the following theorem. It is shown that the polynomial  $f(\cdot)$  can be uniquely recovered using this algorithm with *high probability* if the size of the underlying finite field  $\mathbb{F}_q$  is large enough. Moreover, if the unique recovery is not possible, then the

decoder can identify it as a decoding failure, i.e., the output of this decoding scheme with the probabilistic pruning algorithm is either the true outcome or a decoding failure, provided that the number of errors is bounded by a certain threshold.

To provide our result for the probabilistic scheme, we need the following definition and Lemma that are provided below.

*Definition 2:* A vector  $v \in \mathbb{F}_q^k$  is called a Vandermonde-type vector, or a V-vector in short, if  $v = (1, \lambda, \lambda^2, \dots, \lambda^{k-1})^T$  for some  $\lambda \in \mathbb{F}_q \setminus \{0\}$ .

The following lemma is used to prove our main result in this section.

*Lemma 3:* For any arbitrary matrix  $A_{k \times l}$  ( $k > l$ ) over  $\mathbb{F}_q$  of rank  $r \leq l$ , there exist at most  $r$  distinct V-vectors of length  $k$  that lie in the column space of  $A$ .

*Proof:* Assume to the contrary that there exist  $r+1$  distinct V-vectors that lie in the subspace spanned by the columns of  $A$ , namely,  $v_1, \dots, v_{r+1}$ . Let  $V_{k \times (r+1)} \stackrel{\text{def}}{=} [v_1 | v_2 | \dots | v_{r+1}]$ , which is a Vandermonde matrix. Then, the column space of  $V$  is a subspace of the linear space spanned by the columns of  $A$ . On the other hand, it is well-known that a Vandermonde matrix with distinct columns is always full rank, i.e., the column space of  $V$  has dimension  $r+1$  while  $\text{rank}(A) = r$  which is a contradiction. ■

*Theorem 2:* For the FRS code of length  $N = \frac{n}{m}$  and rate  $R = \frac{k}{n}$  and for any  $s \in [m]$ , the polynomial  $f(\cdot)$  can be uniquely recovered with probability

$$p_d \geq \frac{\sum_{i=l}^t \binom{q-k+l-1}{i} \binom{k-l}{t-i}}{\binom{q-1}{t}}, \quad (12)$$

where  $l < s$  is the dimension of subspace returned by the list decoding algorithm, provided that

- i) the received word  $y \in (\mathbb{F}_q^m)^N$  differs from the FRS code-word corresponding to  $f(\cdot)$  in at most  $\frac{s}{s+1}(1 - \frac{mR}{m-s+1})$  fraction of  $N$  symbols, and
- ii) the side information  $f(\lambda_1), \dots, f(\lambda_t)$  are given to the decoder, where  $\lambda_i$  is drawn uniformly at random from  $\mathbb{F}_q \setminus \{0, \lambda_1, \dots, \lambda_{i-1}\}$ , for  $i \in [t]$ .

Furthermore, if the unique recovery is not possible, the decoder can identify it.

*Proof:* Recall that by applying the linear-algebraic list-decoding algorithm, discussed in Lemma 1, we have  $f = Mx + z$  for which  $M$  has  $I_l$  as its submatrix for some  $l < s$ . Then, without loss of generality, we can assume

$$M = \begin{bmatrix} \tilde{M}_{(k-l) \times l} \\ I_l \end{bmatrix}, \quad (13)$$

where  $\tilde{M}$  is an  $(k-l) \times l$  matrix over  $\mathbb{F}_q$  since the locus of identity submatrix is determined after the list-decoding algorithm is applied. Let  $V$  and  $y_e$  be as characterized in (10) and (11), respectively.

Then, if  $VM$  is full rank, (11) can be solved for  $x$  which is then utilized to determine  $f$ , thereby uniquely recovering the polynomial  $f(\cdot)$ . Furthermore, note that all the following are

equivalent:

$$VM \text{ is full rank.} \iff M^T V^T \text{ is full rank.} \quad (14)$$

$$\iff \langle V^T \rangle \cap \langle N \rangle = \{0\}, \quad (15)$$

where  $N_{k \times (k-l)}$  is a matrix whose columns span the null-space of  $M_{l \times k}^T$ , i.e.,  $M^T N = 0$ . Moreover, note that  $\langle V^T \rangle \cap \langle N \rangle = \{0\}$ , where  $0$  denotes the all-zero vector, if and only if  $[N_{k \times (k-l)} | V_{k \times l}]$  is full rank. According to the result of Lemma 3, there exist at most  $k-l$  V-vectors of length  $k$  that lie in  $\langle N \rangle$ . The field elements that correspond to these vectors are referred to as *bad* choices for  $\lambda_i$ 's. Hence, there are at most  $k-l$  bad choices for  $\lambda_i$ 's in  $\mathbb{F}_q^*$ . Suppose that we randomly pick  $t$  distinct  $\lambda_i$ 's from  $\mathbb{F}_q^*$ . The number of cases where at least  $l$  out of  $t$  random choices of  $\lambda_i$ 's do not fall into this set of size at most  $k-l$  is equal to  $\sum_{i=l}^t \binom{q-1-k+l}{i} \binom{k-l}{t-i}$ . In such cases, we set  $V$  to be the corresponding Vandermonde matrix and  $[N_{k \times (k-l)} | V]$  would be full rank. There are  $\binom{q-1}{t}$  different choices of  $t$ -subsets of  $\mathbb{F}_q^*$ . Consequently, the probability of

$$[N_{k \times (k-l)} | V] \text{ being full rank is at least } \frac{\sum_{i=l}^t \binom{q-1-k+l-1}{i} \binom{k-l}{t-i}}{\binom{q-1}{t}} \text{ by picking } t \text{ extra evaluation points uniformly at random from } \mathbb{F}_q^*.$$

Note that in the event that  $[N_{k \times (k-l)} | V]$  is not full-rank, the system of linear equations characterized in (11) does not have a unique solution. In such a case, the decoder verifies the unique recovery is not possible and declares recovery failure.  $\blacksquare$

We now compare the lower bound on  $p_d$  provided in Theorem 2 with (2) and (3) that follow from [51, Lemma 12] and [52, Lemma 2], respectively, as described in Section II-A. The lower bound provided in (2) is compared with our result in Figure 1 for  $k = 1000$  and  $k = 10000$ . We set  $t = s-1$  in our proposed scheme to have a fair comparison with the lower bound provided in (2). For  $k = 1000$ , our result is almost the same as what is guaranteed by (2) as  $k(s-1)$  is well below the field size  $q = 10003$ . For the case where  $k = 10000$ , the lower bound in (2) is trivial for  $s \geq 10$  where our result still provides a non-trivial bound on  $p_d$ . In other words, the lower bound characterized in Theorem 2 is non-trivial for a wider range of encoding parameters. Moreover, our result allows having  $t \geq s$  extra evaluation points and the lower bound on  $p_d$  provided in (12) improves as the number of extra evaluations  $t$  increases. Consequently, our result does not necessarily require  $q \gg k(s-1)$  as one can improve the lower bound on  $p_d$  with a few more extra evaluations. For instance, with  $t = s+3$ , i.e., having 4 more evaluations, the lower bound provided in Theorem 2 on  $p_d$  significantly improves as demonstrated in Figure 1. This illustrates the superiority of our result over the one that is established based on the result in [51, Lemma 12], provided in (2). Roughly speaking, the advantage of our approach over the one established based on the result in [51, Lemma 12] is due to the difference in how the extra evaluation points are picked. In our approach, we pick all extra evaluation points independently and uniformly at random from  $\mathbb{F}_q^*$  where the extra evaluation points in the latter approach keeps the same structure as FRS symbols, i.e., they are equal to  $a, a\gamma, \dots, a\gamma^{s-1}$ , for some random  $a \in \mathbb{F}_q^*$ . In a sense,

we require more randomness than what is needed in the latter approach.

As discussed in Section II-A, the lower bound on  $p_d$  provided in (3) that follows from the results in [52, Lemma 2] is also improved when more than  $s-1$  extra evaluations are available. Consequently, the large field size is not required for the lower bound on  $p_d$  characterized in (3) as in our result in Theorem 2. In Figure 2, we compare the lower bound provided in (3) with our result for two different field sizes  $q = 10007, 100003$ . The bounds are plotted versus  $t$ , the number of extra evaluations available at the decoder, for  $s = 10$ . It is illustrated that our result requires a significantly smaller number of extra evaluations to guarantee a reasonable successful decoding probability  $p_d$ . The advantage of our approach might be justified by noting that the result in [52, Lemma 2] is valid for any linear code and not necessarily FRS codes. Intuitively, we leverage the certain structure of the encoding matrix of FRS codes to arrive at a better bound on the decoding probability compared to the one provided in [52, Lemma 2] for a general linear code.

## V. FOLDED LAGRANGE CODED COMPUTING

In this section, we demonstrate how the FRS list-decoding algorithm together with the pruning algorithms proposed in Section IV can be utilized to break the barrier on the number of Byzantine workers that can be tolerated in LCC.

Let the parameters  $m$ ,  $N$  and  $K$  be associated with our proposed FLCC, specified in Section III. The main result of this section is that the lower bound on  $N$  in FLCC can be well-approximated by  $(K+T)D_2 + A + S - 1$  for sufficiently large, but fixed, folding parameter  $m$ . This implies that Byzantine adversaries are *as costly as* stragglers in terms of the number of additional workers required in FLCC reducing their effect by a factor of 2 compared to LCC.

Similar to LCC, the decoding algorithm is performed over the received matrices element-wise. In theory, one can apply the decoding algorithm for FRS codes with side information as discussed in Section IV individually for all  $r' \times h'$  elements of the returned matrices. In the rest of this Section, we assume that all the steps discussed below are performed element-wise on the entries of  $\bar{Y}_i$ 's individually. In practice, one might be able to perform the algorithm collaboratively on all entries at once and amortize the computational cost over all entries. A similar approach is provided recently in [44] to decode polynomial codes by utilizing a collaborative decoder for interleaved generalized RS codes. It is shown that up to  $N - K - 1$  Byzantine workers can be tolerated under the additive Gaussian model with *high* probability. The main difference between our proposed method and that of [44] is that we do not make assumptions on the error model and its probability distribution. That is, we consider the worst-case adversarial model.

In FLCC, the master node first finds the linear subspace  $U$  of dimension at most  $s-1$  containing the legitimate polynomial by applying the linear-algebraic list-decoding algorithm. Then it determines the extra evaluation points needed to uniquely identify  $f(\cdot)$  in  $U$  according to the procedure described in the proof of Theorem 1. The master node then performs extra computations to evaluate  $f(\cdot)$  over these points which results

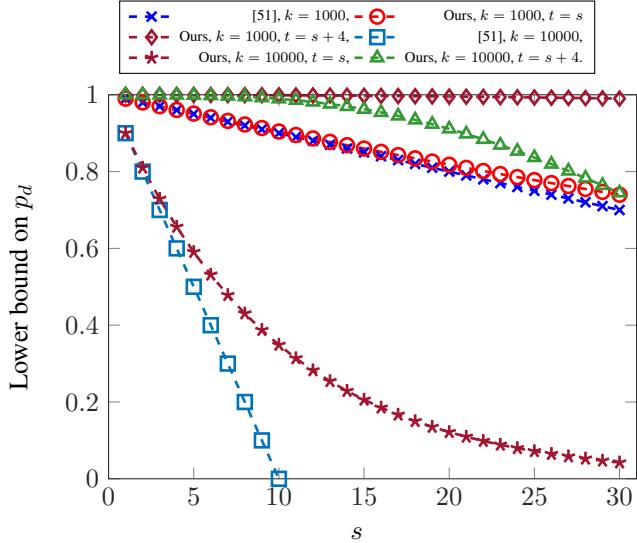


Fig. 1: Comparison of the lower bounds provided in (2) and (12) for  $k = 1000$  and  $k = 10000$ . The field size is  $q = 100003$ . Note that in both case, our proposed lower bound improves significantly by adding 4 more evaluations, i.e.,  $t = s + 3$ .

in uniquely determining  $f(\cdot)$  by solving the system of linear equations in (11). Let

$$r \stackrel{\text{def}}{=} \frac{(K + T - \frac{1}{m})D_2 + 1}{N - S},$$

which is referred to as the *modified rate* of FLCC. The fraction of adversaries tolerated in FLCC is characterized in the following theorem.

*Theorem 3:* An FLCC with folding parameter  $m$  and the number of worker nodes equal to  $N$  is  $S$ -resilient,  $A$ -secure and  $T$ -private to compute  $\{g(\mathbf{X}_i)\}_{i=1}^{mK}$  for a  $D_2$ -degree polynomial  $g(\cdot)$  as long as

$$\frac{A}{N - S} \leq \frac{s^*}{s^* + 1} \left(1 - \frac{mr}{m - s^* + 1}\right), \quad (16)$$

where  $s^*$  is equal to  $\lceil \tilde{s} \rceil - 1$  or  $\lceil \tilde{s} \rceil$ , where  $\tilde{s} = \frac{\sqrt{m(m+1)(m(1-r)+2)r} - (m+1)}{mr-1}$ , depending on which one results in a larger RHS in (16).

*Proof:* We first note that the degree of the composed polynomial interpolated at FLCC decoder is  $(m(K+T)-1)D_2$ . By the result of Lemma 2, we can use the result of Lemma 1 with  $S$  straggling nodes by replacing  $N$  with  $N - S$ , i.e., replacing  $R$  in (1) by the modified rate  $r$ . Let  $a(s) \stackrel{\text{def}}{=} \frac{s}{s+1} \left(1 - \frac{mr}{m-s+1}\right)$ . Then, the result of Theorem 1 implies that FLCC is  $S$ -resilient and  $A$ -secure as long as

$$\frac{A}{N - S} \leq a(s), \quad (17)$$

for any arbitrary integer  $s \in [m]$ . Let  $s^* \stackrel{\text{def}}{=} \arg \max_{s \in [m]} a(s)$  and  $\tilde{s}$  denote the solution to the same optimization problem with a difference that the underlying variable  $s$  is assumed to be continuous, i.e.,  $\tilde{s} \stackrel{\text{def}}{=} \arg \max_{s \in \mathbb{R}, 0 \leq s \leq m} a(s)$ . One can check that  $a(s)$

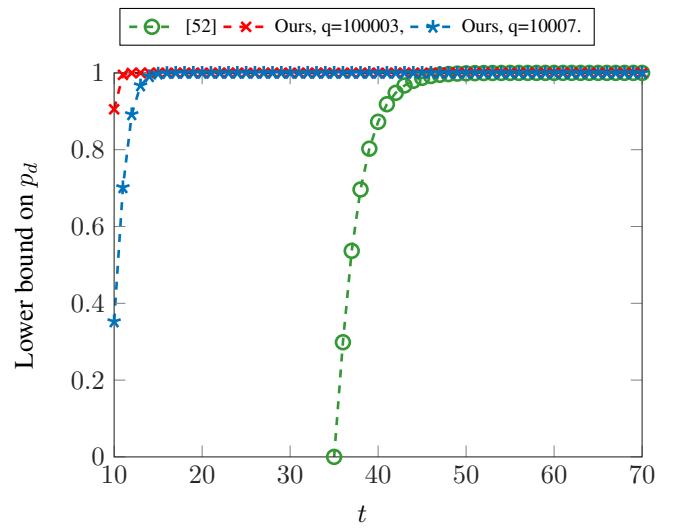


Fig. 2: Comparison of the lower bounds provided in (3) and (12). The plots show the lower bounds on  $p_d$  versus the number of extra evaluation points available at the decoder for  $s = 10$ . Other parameters are  $q = 10007, 100003$  and  $n = 2000$ .

is a concave function which implies that  $s^*$  is either equal to  $\lceil \tilde{s} \rceil - 1$  or  $\lceil \tilde{s} \rceil$ , whichever maximizes  $a(s)$  and also belongs to  $[m]$ . The concavity of  $a(s)$  also implies that  $\tilde{s}$  is a root of  $\frac{da(s)}{ds}$  or it is equal to one of the boundary values. The roots of  $\frac{da(s)}{ds}$  are

$$\tilde{s}_{\pm} = \frac{\sqrt{m(m+1)(m(1-r)+2)r} \pm (m+1)}{mr-1}. \quad (18)$$

One can check that  $\tilde{s}_+$  is not a feasible solution since it does not satisfy the constraints of the continuous optimization problem, i.e., it is always the case that  $\tilde{s}_+ < 0$  or  $m < \tilde{s}_+$ . Furthermore, we have  $0 < \tilde{s}_- \leq m$  for  $\frac{1}{m} \leq r$ , and,  $m < \tilde{s}_- \leq m+1$ , otherwise. Note that for the latter case the boundary condition implies  $\tilde{s} = m$ . Then, at least one of  $\lceil \tilde{s} \rceil - 1$  and  $\lceil \tilde{s} \rceil$  is always a feasible solution for the discrete optimization problem and, consequently,  $s^*$  is either equal to  $\lceil \tilde{s} \rceil - 1$  or  $\lceil \tilde{s} \rceil$ , whichever is feasible and returns a larger value for  $a(s)$ . The proof of  $T$ -privacy is similar to the one in LCC by noting that both the number of random mask matrices in the encoding polynomial of FLCC, characterized in (8), and the number of shares available at each worker node are larger than those in LCC by a multiplicative factor of  $m$ . Hence, the dataset is perfectly masked by  $mT$  random matrices  $\mathbf{Z}_j$ 's against any coalition of  $T$  worker nodes, each having  $m$  evaluations of  $u_m(\cdot)$ . ■

In order to compare the performance of FLCC with LCC, we consider evaluating  $g(\cdot)$  over  $m$  batches of data where each batch contains  $K$  input matrices, as explained in Section III. To this end, LCC is run  $m$  times in the first scenario, each time computing  $g(\cdot)$  over a single batch of matrices. Then, the total amount of computations performed at the master node for decoding is  $m$  times the decoding complexity of running LCC once. More specifically, the overall decoding complexity when LCC is employed is  $O(m(N - S) \log^2(N - S) \log \log(N -$

$S)r'h'$ ). Furthermore, (7) implies that the maximum number of Byzantine workers tolerated in this scenario can be expressed as follows

$$A_{\text{LCC}} \stackrel{\text{def}}{=} \left\lfloor \frac{N - (K + T - 1)D_2 - S - 1}{2} \right\rfloor. \quad (19)$$

In FLCC, the computations performed at the master node can be considered as two separate procedures. The first one is running the list-decoding algorithm and pruning it as well as solving (11) to uniquely interpolate  $f(\cdot)$ , referred to as the interpolation step. The second one corresponds to computing evaluations of  $g(\cdot)$  over the set of extra evaluation points returned by the underlying pruning algorithm used and is referred to as the extra computation step. The computational complexity of the interpolation step is  $O((N^2m^2 + m^2K^2s)r'h')$  according to Theorem 1. In the extra computation step, the master node evaluates  $g(\cdot)$  over at most  $s^* - 1$  points. This implies that the amount of extra computations performed over the master node normalized by the computational complexity of each worker node, referred to as *normalized extra computation*, is at most  $\frac{s^*-1}{m}$ . Moreover, the computation load of the worker nodes in both FLCC and LCC is the same, i.e., each worker node evaluates  $g(\cdot)$  over a batch of data consisting of  $m$  matrices in either of these scenarios. Also, the amount of commutations required, referred to as the *communication complexity*, in FLCC is equal to that of running LCC  $m$  times as well. The result of Theorem 3 implies that the number of adversaries tolerated in FLCC is expressed as follows

$$A_{\text{FLCC}} \stackrel{\text{def}}{=} \left\lfloor \frac{s^*}{s^* + 1} \left( N - \frac{mD_2}{m - s^* + 1} \left( K + T - \frac{1}{m} \right) - S - 1 \right) \right\rfloor, \quad (20)$$

where  $s^*$  is characterized in Theorem 3. Note that by setting  $m = 1$ , (20) is reduced to (19) since an FRS code with  $m = 1$  is an RS code, implying that LCC and FLCC are in fact identical for this special case, as expected. Note also that the decoding complexity in LCC and interpolation complexity in FLCC grow linearly with the dimensions of the output, i.e.,  $r'$  and  $h'$ , as in both schemes the decoding procedure must be performed element-wise for all the elements of the output matrix and also both are independent of the size of input matrices, i.e.,  $r \times h$ , as well as the degree of polynomial function evaluated over the dataset, i.e.,  $D_2$ . However, the decoding complexity in FLCC is quadratic in the number of worker nodes  $N$  while it is *almost* linear in  $N$  in LCC. In Figure 3, the normalized extra computation is plotted versus the folding parameter  $m$ , for a certain set of parameters. It is illustrated that the ratio of the computational cost of extra evaluations of  $f(\cdot)$  at the master node to the workload of a worker node approaches zero as  $m$  grows. In particular, by using the result of Theorem 3, it can be observed that this ratio approaches zero in  $O(\frac{1}{\sqrt{m}})$ . For instance, when  $m = 100$  and  $N = 1000$ , one can observe that the amount of extra computations needed is less than 5% of the computational task of each individual worker node or, equivalently, is less than 0.005% of the total computational job. Figure 4 demonstrates the advantage of FLCC over LCC by comparing the maximum number of adversaries tolerated in each scheme for the same set of parameters. The ratio of the maximum number of adversaries tolerated in FLCC to that

of LCC is plotted versus the folding parameter  $m$ . It shows that FLCC can tolerate *almost* twice as many as adversaries tolerated in LCC for the same parameters  $N, S, K$  and  $T$ .

The result of Theorem 3 is simplified for large enough  $m$  in the following corollary.

**Corollary 1:** The FLCC specified in Theorem 3 can tolerate up to

$$A_{\text{FLCC}} = \lfloor (1 - \epsilon)(N - S) - (K + T)D_2 - 1 \rfloor \quad (21)$$

Byzantine adversaries for  $m = O(\frac{1}{\epsilon^2})$  and  $s^* = O(\frac{1}{\epsilon})$ .

**Remark 1:** Note that (21) implies the marginal cost of tolerating one more Byzantine adversary in FLCC is one additional worker node, the same as that of tolerating one more straggler. This demonstrates the advantage of FLCC over LCC in which two additional worker nodes are needed to tolerate one more Byzantine adversary. In other words, FLCC improves the trade-off between the number of adversaries and stragglers can be tolerated while other parameters are fixed by removing the factor 2 in (7), thereby providing a scheme in which both adversaries and stragglers cost evenly, as opposed to LCC.

**Remark 2:** In this paper, we have assumed that the extra computations needed to obtain the side information are done at the master node. The numerical results shown in Figure 3 confirms the soundness of this approach in practice. However, for scenarios where even this tiny amount of extra computation must be avoided, the master node can employ a few trusted nodes that can perform this computation without error, e.g., by using software guard extensions (SGX) technology implemented in Intel central processing units (CPU).

The decoding algorithm provided for FLCC in this section always guarantees uniquely recovering the computation outcome, i.e., the computation result is deterministically provided by the decoder. This algorithm is established upon the novel deterministic pruning algorithm for FRS code characterized in Theorem 1 in Section IV. In this algorithm, the extra evaluation points are determined *after* the list-decoding algorithm is applied. In other words, it is assumed that the side information symbols are allowed to be constructed based on the output of the list-decoding algorithm. In a practical scenario where parallelization of tasks is preferred to reduce the latency, the side information can be specified simultaneously by the master node as the workers perform computations, as shown by Theorem 2. In this case, the evaluation of  $f(\cdot)$  over  $s^* - 1$  points picked uniformly at random from  $\mathbb{F}_q$  are provided as the side information to the list decoder. Theorem 2 implies that the system of linear equations specified in (11) has a unique solution with *high* probability, establishing that each element of the computation outcome can be uniquely determined with the same probability.

## VI. CONCLUSION

In this work, we considered a coded distributed computing setting with a master node and a set of workers. We proposed a coding-theoretic approach that boosts the adversarial toleration threshold in such systems. In particular, we adapted the folding technique in coding theory to the context of coded computing and leveraged the list-decoding algorithms for FRS

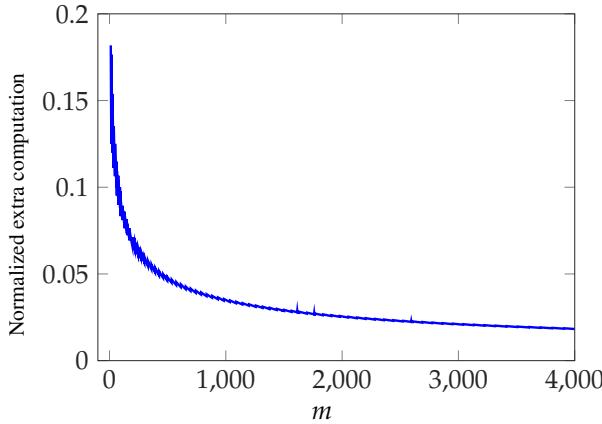


Fig. 3: Demonstration of the ratio between the extra computations performed at the master node to the workload of each worker ( $\frac{s^*-1}{m}$ ), referred to as *normalized extra computation*, versus the folding parameter  $m$  in FLCC. The relative computational cost of evaluating  $f(\cdot)$  over the set of extra points at the master node approaches zero in  $O(\frac{1}{\sqrt{m}})$ . Other parameters are  $N = 1000$ ,  $K = 180$ ,  $T = 11$ ,  $S = 20$ , and  $D = 2$ .

codes for recovering the overall computation outcome at the master node. Furthermore, in order to guarantee unique recovery of the outcome, we proposed novel deterministic and probabilistic pruning algorithms for list-decoding FRS codes with side information that are of independent interest in the list-decoding literature. By utilizing our proposed techniques, we introduced the folded Lagrange coded computing (FLCC) protocol that outperforms LCC by improving the number of adversaries that can be tolerated *almost* by a factor of two. More specifically, we showed that in FLCC adversaries and stragglers cost almost evenly in terms of the number of workers required, compared to LCC in which tolerating one adversary costs twice as overcoming one straggler.

#### ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewer for providing a simple and short proof for Theorem 1.

#### REFERENCES

[1] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, “Speeding up distributed machine learning using codes,” *IEEE Trans. Inf. Theory*, vol. 64, no. 3, pp. 1514–1529, 2018.

[2] S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “Coding for distributed fog computing,” *IEEE Commun. Mag.*, vol. 55, no. 4, pp. 34–40, 2017.

[3] Q. Yu, S. Li, N. Raviv, S. M. M. Kalan, M. Soltanolkotabi, and S. A. Avestimehr, “Lagrange coded computing: Optimal design for resiliency, security, and privacy,” in *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 2019, pp. 1215–1225.

[4] S. Li and S. Avestimehr, “Coded computing: Mitigating fundamental bottlenecks in large-scale distributed computing and machine learning,” *Found. Trends Commun. Inf. Theory*, vol. 17, no. 1, pp. 1–148., 2020.

[5] J. So, B. Guler, A. S. Avestimehr, and P. Mohassel, “CodedPrivateML: A fast and privacy-preserving framework for distributed machine learning,” *IEEE Journal on Selected Areas in Information Theory (JSAIT): Special issue on Privacy and Security of Information Systems*, 2021.

[6] J. So, B. Guler, and A. S. Avestimehr, “Turbo-aggregate: Breaking the quadratic aggregation barrier in secure federated learning,” *IEEE Journal on Selected Areas in Information Theory (JSAIT): Special issue on Privacy and Security of Information Systems*, 2021.

[7] J. So, B. Guler, and S. Avestimehr, “A scalable approach for privacy-preserving collaborative machine learning,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.

[8] M. Soleymani, H. Mahdavifar, and A. S. Avestimehr, “Privacy-preserving distributed learning in the analog domain,” *arXiv preprint:2007.08803*, 2020.

[9] T. Jahani-Nezhad and M. A. Maddah-Ali, “Berrut approximated coded computing: Straggler resistance beyond polynomial computing,” *arXiv preprint arXiv:2009.08327*, 2020.

[10] S. Prakash, S. Dhakal, M. R. Akdeniz, Y. Yona, S. Talwar, S. Avestimehr, and N. Himayat, “Coded computing for low-latency federated learning over wireless edge networks,” *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 1, pp. 233–250, 2020.

[11] W. Raghupathi and V. Raghupathi, “Big data analytics in healthcare: promise and potential,” *Health information science and systems*, vol. 2, no. 1, p. 3, 2014.

[12] A. McAfee, E. Brynjolfsson, T. H. Davenport, D. Patil, and D. Barton, “Big data: the management revolution,” *Harvard business review*, vol. 90, no. 10, pp. 60–68, 2012.

[13] P. Elias, “List decoding for noisy channels,” 1957.

[14] M. Sudan, “Decoding of Reed Solomon codes beyond the error-correction bound,” *Journal of complexity*, vol. 13, no. 1, pp. 180–193, 1997.

[15] V. Guruswami and M. Sudan, “Improved decoding of Reed-Solomon and algebraic-geometric codes,” in *39th Annual Symposium on Foundations of Computer Science (Cat. No. 98CB36280)*. IEEE, 1998, pp. 28–37.

[16] F. Parvaresh and A. Vardy, “Correcting errors beyond the Guruswami-Sudan radius in polynomial time,” *Proceedings of 46th Annual Symposium on Foundations of Computer Science*, pp. 285–294, 2005.

[17] V. Guruswami and A. Rudra, “Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy,” *IEEE Trans. Inf. Theory*, vol. 54, no. 1, pp. 135–150, 2008.

[18] V. Guruswami and C. Wang, “Linear-algebraic list decoding for variants of Reed-Solomon codes,” *IEEE Trans. Inf. Theory*, vol. 59, no. 6, pp. 3257–3268, 2013.

[19] V. Guruswami, “List decoding with side information,” in *18th IEEE Annual Conference on Computational Complexity*, 2003, pp. 300–309.

[20] S. Kopparty, “List-decoding multiplicity codes,” *Theory of Computing*, vol. 11, no. 1, pp. 149–182, 2015.

[21] R. Cramer, I. B. Damgård, N. Döttling, S. Fehr, and G. Spini, “Linear secret sharing schemes from error correcting codes and universal hash functions,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2015, pp. 313–336.

[22] R. Safavi-Naini and P. Wang, “A model for adversarial wiretap channels and its applications,” *Journal of information processing*, vol. 23, no. 5, pp. 554–561, 2015.

[23] M. Cheraghchi, “Nearly optimal robust secret sharing,” *Designs, Codes and Cryptography*, vol. 87, no. 8, pp. 1777–1796, 2019.

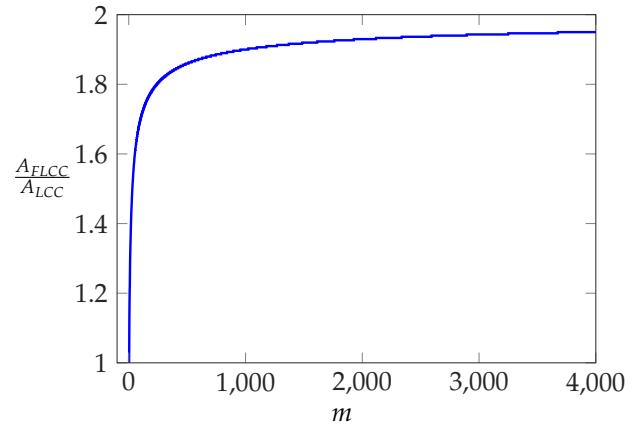


Fig. 4: Demonstration of the ratio of the number of Byzantine adversaries in FLCC to that of LCC for  $N = 1000$ ,  $K = 180$ ,  $T = 11$ ,  $S = 20$ , and  $D = 2$ . The plot indicates that as  $m$  grows, FLCC can tolerate *almost* as twice as the number of adversaries in LCC with the same set of parameters.

[24] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, “Speeding up distributed machine learning using codes,” *IEEE Trans. Inf. Theory*, vol. 64, no. 3, pp. 1514–1529, 2018.

[25] S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “A unified coding framework for distributed computing with straggling servers,” in *2016 IEEE Globecom Workshops (GC Wkshps)*, pp. 1–6.

[26] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, “Straggler mitigation in distributed matrix multiplication: Fundamental limits and optimal coding,” *IEEE Trans. Inf. Theory*, vol. 66, no. 3, pp. 1920–1933, 2020.

[27] A. Reisizadeh, S. Prakash, R. Pedarsani, and A. S. Avestimehr, “Coded computation over heterogeneous clusters,” *IEEE Trans. Inf. Theory*, vol. 65, no. 7, pp. 4227–4242, 2019.

[28] M. Aliasgari, J. Kliewer, and O. Simeone, “Coded computation against straggling decoders for network function virtualization,” in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 711–715.

[29] M. V. Jamali, M. Soleymani, and H. Mahdavifar, “Coded distributed computing: Performance limits and code designs,” in *2019 IEEE Information Theory Workshop (ITW)*, pp. 1–5.

[30] C.-S. Yang, R. Pedarsani, and A. S. Avestimehr, “Edge computing in the dark: Leveraging contextual-combinatorial bandit and coded computing,” *IEEE/ACM Transactions on Networking*, 2021.

[31] Q. Yu and A. S. Avestimehr, “Entangled polynomial codes for secure, private, and batch distributed matrix multiplication: Breaking the “cubic” barrier,” *arXiv preprint arXiv:2001.05101*, 2020.

[32] M. Aliasgari, O. Simeone, and J. Kliewer, “Private and secure distributed matrix multiplication with flexible communication load,” *IEEE Trans. on Information Forensics and Security*, vol. 15, pp. 2722–2734, 2020.

[33] R. G. D’Oliveira, S. El Rouayheb, and D. Karpuk, “GASP codes for secure distributed matrix multiplication,” *IEEE Trans. Inf. Theory*, vol. 66, pp. 4038–4050, 2020.

[34] R. Bitar, Y. Xing, Y. Keshtkarjahromi, V. Dasari, S. E. Rouayheb, and H. Seferoglu, “Private and rateless adaptive coded matrix-vector multiplication,” *arXiv preprint arXiv:1909.12611*, 2019.

[35] H. A. Nodehi and M. A. Maddah-Ali, “Secure coded multi-party computation for massive matrix operations,” *arXiv preprint:1908.04255*, 2019.

[36] M. Fahim and V. R. Cadambe, “Numerically stable polynomially coded computing,” in *IEEE Trans. Inf. Theory*, 2021.

[37] A. Ramamoorthy and L. Tang, “Numerically stable coded matrix computations via circulant and rotation matrix embeddings,” *arXiv preprint arXiv:1910.06515*, 2019.

[38] 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.

[39] N. Charalambides, H. Mahdavifar, and A. O. Hero III, “Numerically stable binary gradient coding,” *arXiv preprint arXiv:2001.11449*, 2020.

[40] M. Soleymani, H. Mahdavifar, and A. S. Avestimehr, “Analog Lagrange coded computing,” *IEEE Journal on Selected Areas in Information Theory (JSAT): Special issue on Privacy and Security of Information Systems*, 2021.

[41] 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.

[42] T. Tang, R. E. Ali, H. Hashemi, T. Gangwani, S. Avestimehr, and M. Annavaram, “Verifiable coded computing: Towards fast, secure and private distributed machine learning,” *arXiv preprint arXiv:2107.12958*, 2021.

[43] S. Dutta, Z. Bai, H. Jeong, T. M. Low, and P. Grover, “A unified coded deep neural network training strategy based on generalized polydot codes,” in *2018 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2018, pp. 1585–1589.

[44] A. M. Subramaniam, A. Heidarzadeh, and K. R. Narayanan, “Collaborative decoding of polynomial codes for distributed computation,” in *2019 IEEE Information Theory Workshop (ITW)*. IEEE, 2019, pp. 1–5.

[45] J. M. Wozencraft, “List decoding,” *Quarterly Progress Report*, vol. 48, pp. 90–95, 1958.

[46] O. Goldreich and L. A. Levin, “A hard-core predicate for all one-way functions,” in *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, 1989, pp. 25–32.

[47] P. Elias, “Error-correcting codes for list decoding,” *IEEE Trans. Inf. Theory*, vol. 37, no. 1, pp. 5–12, 1991.

[48] V. Guruswami, J. Hastad, M. Sudan, and D. Zuckerman, “Combinatorial bounds for list decoding,” *IEEE Trans. Inf. Theory*, vol. 48, no. 5, pp. 1021–1034, 2002.

[49] V. Guruswami, “Linear-algebraic list decoding of folded Reed-Solomon codes,” in *2011 IEEE 26th Annual Conference on Computational Complexity*, pp. 77–85.

[50] S. Kopparty, N. Ron-Zewi, S. Saraf, and M. Wootters, “Improved decoding of folded reed-solomon and multiplicity codes,” in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2018, pp. 212–223.

[51] V. Guruswami and S. Kopparty, “Explicit subspace designs,” *Combinatorica*, vol. 36, no. 2, pp. 161–185, 2016.

[52] S. Saraf and S. Yekhanin, “Noisy interpolation of sparse polynomials, and applications,” in *2011 IEEE 26th Annual Conference on Computational Complexity*. IEEE, 2011, pp. 86–92.

[53] R. E. Blahut, *Algebraic codes on lines, planes, and curves: an engineering approach*. Cambridge University Press, 2008.

[54] Q. Yu, S. Li, N. Raviv, S. M. M. Kalan, M. Soltanolkotabi, and S. A. Avestimehr, “Lagrange coded computing: Optimal design for resiliency, security, and privacy,” in *The 22nd International Conference on Artificial Intelligence and Statistics*, 2019, pp. 1215–1225.

[55] 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, pp. 4403–4413, 2017.

[56] S. Dutta, M. Fahim, F. Haddadpour, H. Jeong, V. Cadambe, and P. Grover, “On the optimal recovery threshold of coded matrix multiplication,” *IEEE Trans. Inf. Theory*, vol. 66, no. 1, pp. 278–301, 2019.

[57] V. Pan, *Structured matrices and polynomials: unified superfast algorithms*. Springer Science & Business Media, 2001.

## APPENDIX

*Proof of Lemma 2:* The first step in the algebraic list decoding algorithm when  $S$  symbols are erased is to interpolate the following multivariate polynomial by using non-erased symbols:

$$Q(X, Y_1, \dots, Y_s) = A_0(X) + A_1(X)Y_1 + \dots + A_s(X)Y_s, \quad (22)$$

where  $\deg(A_i) \leq D$  for  $i = 1, \dots, s$ , and  $\deg(A_0) \leq D + k - 1$ , with degree parameter  $D = \left\lfloor \frac{(N-S)(m-s+1)-k+1}{s+1} \right\rfloor$ . Then, the number of unknown coefficients in  $Q$  is  $(D+1)s + D + k = (D+1)(s+1) + k - 1 > (N-S)(m-s+1)$ . The interpolation requirements are

$$Q(\gamma^{im+j}, y_{im+j}, \dots, y_{im+j+s-1}) = 0 \quad (23)$$

for all indices  $i$  corresponding to the non-erased symbols and  $j = 0, 1, \dots, m - s$ . Since  $S$  symbols are erased, the number of interpolation conditions is  $(N - S)(m - s + 1)$  which is less than the number of unknown coefficients in  $Q$ . Hence, a solution for  $Q(\cdot)$  can be found by solving a homogeneous linear system over  $\mathbb{F}_q$  with the same complexity claimed in Lemma 1. By using the interpolation requirements, one can show that  $Q(X, f(X), f(\gamma X), \dots, f(\gamma^{s-1}X)) = 0$  if  $f(\cdot)$  is a polynomial of degree at most  $k - 1$  whose FRS encoding agrees with the received word in  $t > \frac{D+k-1}{m-s+1}$  locations. For our choice of  $D$ , the requirement on  $t$  is met if  $t(m - s + 1) > \frac{(N-S)(m-s+1)+s(k-1)}{s+1}$ . Alternatively, the fractional disagreement is at most  $\frac{s}{s+1}(1 - \frac{mk}{(N-S)(m-s+1)})$ .



**Mahdi Soleymani** (Student Member, IEEE) received his B.S. and M.S. degrees in Electrical Engineering at Sharif University of Technology, Tehran, Iran, in 2014 and 2016, respectively. He is currently pursuing his Ph.D. degree in Electrical Engineering and Computer Science at University of Michigan, Ann Arbor. His research interests lie in the area of algebraic coding theory with applications to distributed storage systems, wireless networks, distributed computing and machine learning.



**Ramy E. Ali** (Member, IEEE) is a Postdoctoral Research Associate at the Electrical and Computer Engineering Department at the University of Southern California. He received the B.Sc. degree in electrical engineering from Alexandria University, Alexandria, Egypt, in 2011, the M.Sc. degree in wireless communications from the Wireless Intelligent Networks Center (WINC), Nile University, Giza, Egypt, in 2014 and the PhD degree in electrical engineering from Penn State University, University Park, USA, in 2020.

He was a research intern at Osram Sylvania (Siemens Business), Beverly, MA, USA, in 2015. He was also a research intern at the Math of Communications Department, Bell Labs, Crawford Hill, NJ, USA, in 2017 and 2018 and a research intern at the IP Networking Department, Bell Labs, Murray Hill, NJ, USA, in 2019.

Dr. Ali was a recipient of the Young Innovators Award (YIA) in 2011, Nile University's Graduate Fellowship in 2011, the Pennsylvania State University's Graduate Fellowship in 2015, the Newport Research Excellence Award in 2016 and the Unix 50 Presentation Award from Bell Labs in 2019. His research interests include distributed systems, machine learning, information theory and coding theory.



**Hessam Mahdavifar** (Member, IEEE) is an Assistant Professor in the Department of Electrical Engineering and Computer Science at the University of Michigan Ann Arbor. He received the B.Sc. degree from the Sharif University of Technology, Tehran, Iran, in 2007, and the M.Sc. and the Ph.D. degrees from the University of California San Diego (UCSD), La Jolla, in 2009, and 2012, respectively, all in electrical engineering. He was with the Samsung US R&D between 2012 and 2016, in San Diego, US, as a staff research engineer.

He received the NSF career award in 2020. He also received Best Paper Award in 2015 IEEE International Conference on RFID, and the 2013 Samsung Best Paper Award. He also received two Silver Medals at International Mathematical Olympiad in 2002 and 2003, and two Gold Medals at Iran National Mathematical Olympiad in 2001 and 2002. His main area of research is coding and information theory with applications to wireless communications, security, privacy, and machine learning.



**A. Salman Avestimehr** (Fellow, IEEE) is a Professor and director of the Information Theory and Machine Learning (vITAL) research lab at the Electrical and Computer Engineering Department of University of Southern California. He received his Ph.D. in 2008 and M.S. degree in 2005 in Electrical Engineering and Computer Science, both from the University of California, Berkeley. Prior to that, he obtained his B.S. in Electrical Engineering from Sharif University of Technology in 2003. His research interests include information theory and coding theory, and large-scale

distributed computing and machine learning, secure and private computing, and blockchain systems.

Dr. Avestimehr has received a number of awards for his research, including the James L. Massey Research & Teaching Award from IEEE Information Theory Society, an Information Theory Society and Communication Society Joint Paper Award, a Presidential Early Career Award for Scientists and Engineers (PECASE) from the White House (President Obama), a Young Investigator Program (YIP) award from the U. S. Air Force Office of Scientific Research, a National Science Foundation CAREER award, the David J. Sakrison Memorial Prize, and several Best Paper Awards at Conferences. He has been an Associate Editor for IEEE Transactions on Information Theory and a general Co-Chair of the 2020 International Symposium on Information Theory (ISIT).