

# Distributed Black-Box Optimization via Error Correcting Codes

Burak Bartan and Mert Pilanci

Department of Electrical Engineering, Stanford University

Email: {bbartan, pilanci}@stanford.edu

**Abstract**—We introduce a novel distributed derivative-free optimization framework that is resilient to stragglers. The proposed method employs coded search directions at which the objective function is evaluated, and a decoding step to find the next iterate. Our framework can be seen as an extension of evolution strategies and structured exploration methods where structured search directions were utilized. As an application, we consider black-box adversarial attacks on deep convolutional neural networks. Our numerical experiments demonstrate a significant improvement in the computation times.

## I. INTRODUCTION

Derivative-free optimization is an important computational task in many areas such as machine learning, statistics, design optimization and decision making [1]. In many applications, it is common to encounter black-box optimization problems where the gradient of the objective function is not available. It is also possible that the gradient information is unreliable, or impractical to obtain. Derivative-free optimization methods employ only the function values to make progress toward an optimal solution. They can be used to directly tackle complex and diverse problems without an analytical form of the objective function. Additionally, since the function values can be evaluated independently in parallel, derivative-free methods can exploit parallelism to a great extent.

In the recent years, cloud computing has offered inexpensive and scalable solutions to computational problems, and is widely adopted and used by the community. However, cloud-based serverless computing systems typically suffer from system noise, latency and variability in computation times [2], [3]. This results in a subset of slower worker nodes called stragglers. In this work, we propose a distributed derivative-free optimization method that is resilient to stragglers. Our method forms gradient estimates by evaluating the function at several points by employing coding theoretic ideas and tools. We first encode ordinary basis perturbations using Hadamard transform into perturbation directions along the rows of the Hadamard matrix  $H$ . The objective function is evaluated at the encoded points in parallel, and then a gradient direction is determined from the available workers via decoding. We establish a connection between derivative free optimization via the well-known finite differences method, recently proposed structured evolution strategies [4], and polar codes [5]. The structured evolution strategies approach corresponds to averaging the outputs of function evaluations with variables perturbed along the rows of  $HD$  where  $H$  is the Hadamard

matrix and  $D$  is a diagonal matrix with random diagonal entries. When the decoding step is omitted, encoding the unit perturbations leads to the structured evolution strategies which is simple to implement. If we decode the outputs instead of averaging them using the modified successive cancellation decoder of polar codes, however, we obtain partial derivatives as in the finite differences method. Because the function evaluations are scalar valued and the decoder can run in near linear time, the time it takes to decode these scalar outputs is negligible. For a large class of problems, including the decoding step yields faster convergence in optimization. The proposed method also enables a way to switch between finite differences and structured evolution strategies and provide a more favorable trade-off.

## A. Related Work

This work draws ideas from coded computation literature where the goal is to speed up distributed matrix multiplication. The authors in [6] introduced the concept of using Maximum Distance Separable (MDS) codes to compute matrix-vector multiplications faster in distributed computing platforms. The way coding is integrated into the computation of a matrix vector product  $Ax$ , where  $A \in \mathbb{R}^{n \times d}$  and  $x \in \mathbb{R}^d$  is as follows: First  $A$  is partitioned into submatrices along the rows and then these submatrices are encoded using MDS codes. The workers multiply their assigned submatrix with  $x$ . Because the encoded  $A$  contains redundancy, it is possible to decode  $Ax$  via a standard decoder without accessing the outputs of all workers. Works that introduce variants and extensions of this idea include [7], [8] and [9]. There are many other recent studies on coded computation that explore the partitioning of both matrices being multiplied, e.g., [10], [11]. This work is more related to the case where only one of the matrices is partitioned.

Another line of study that this work is connected to is the study of black-box optimization methods for reinforcement learning. The authors in [12] explore the use of the evolution strategies method in reinforcement learning and show that distributed training with evolution strategies can be very fast because of its scalability. The work in [4] shows that using orthogonal exploration directions leads to lower errors and present the *structured* evolution strategies method which is based on a special way of generating random orthogonal exploration directions, as we discuss later in detail.

## B. Main Contributions

We introduced a novel framework based on coded gradient estimates that leads to straggler-resilient distributed optimization of black-box functions. Our method introduces a novel connection between the finite differences method and the structured evolution strategies method by using polar codes [5]. Once function evaluations at perturbed parameters are returned, it is possible to use them as they are without decoding, which is equivalent to the structured evolution strategies method or decode them to obtain the gradient estimate and use it in a gradient-based optimization algorithm. Our method offers this flexibility at the expense of doing some additional inexpensive computing for decoding because of the efficient polar code decoder.

## II. PRELIMINARIES

### A. Notation

We are interested in solving the black-box optimization problem given by

$$\underset{\theta}{\text{minimize}} \quad f(\theta)$$

where  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ . Throughout the text, unless otherwise stated, we assume that we do not have access to the analytical form of  $f(\theta)$  or its gradient, and we assume that we are only able to make queries for function evaluations. We use  $\nabla_v$  to denote directional derivative along the direction  $v$ . We use  $e_i$  to denote the  $i$ 'th unit column vector with the appropriate dimension. Throughout the text,  $H$  refers to the Hadamard matrix (its dimension can be determined from the context), and  $\mathbf{H}$  refers to the Hessian of a function.

### B. Finite Differences

For black-box optimization problems, approximate gradients can be obtained by using the finite differences estimator. The derivative of a function  $f(x)$  with respect to the variable  $x_i$  can be approximated by using

$$\frac{\partial f(x)}{\partial x_i} \approx \frac{f(x + \delta e_i) - f(x - \delta e_i)}{2\delta}, \quad (1)$$

where  $e_i$  is the  $i$ 'th unit vector and  $\delta \in \mathbb{R}$  is a small scalar that determines the perturbation amount. One can obtain an approximate gradient using (1) and feed these estimate gradients to gradient-based optimization methods such as gradient descent.

The finite differences estimator can easily exploit parallelism since the partial approximate derivatives can be computed in different workers in a distributed setting. This method can be modified to be straggler-resilient by using only the available derivative estimates and ignoring the outputs of the slower workers. However, this typically leads to slower convergence. We note that the proposed method results in a mechanism that always returns all entries of the gradient estimate and is straggler-resilient at the same time. Another alternative is to replicate finite difference calculations, which is not optimal from a coding theory perspective.

### C. Evolution Strategies

Let us consider the following evolution strategies (ES) gradient estimator ([12], [4])

$$\nabla f(\theta) \approx \frac{1}{2N\delta} \sum_{i=1}^N (f(\theta + \delta \epsilon_i) \epsilon_i - f(\theta - \delta \epsilon_i) \epsilon_i) \quad (2)$$

where  $\delta$  is the scaling coefficient for the random perturbation directions  $\epsilon_i$  and  $N$  is the number of perturbations. This estimator is referred to as antithetic evolution strategies gradient estimator. There are other gradient estimators slightly different from the one in (2). In this work, we always consider the antithetic version given in (2) because it leads to better empirical results in our numerical simulations. The random perturbation directions  $\epsilon_i$  may be sampled from a standard multivariate Gaussian distribution  $\mathcal{N}(0, I)$ . Alternatively,  $\epsilon_i$ 's may be generated as the rows of  $HD$  where the  $D$  is a diagonal matrix with entries distributed according to Rademacher distribution. The authors in [4] propose generating perturbation directions by repetitive multiplications  $HD_1 \dots HD_k$  where  $D_i$ 's are independent.

It is shown in [4] that if exploration (or perturbation) directions  $\epsilon_i$  are orthogonal, the gradient estimators lead to lower error. We note that the rows of  $HD$  are orthogonal with the appropriate scaling factor. We omit this scaling factor by absorbing it in the  $\delta$  term which scales perturbation directions  $\epsilon_i$ . This work considers the case where random perturbation directions are generated according to  $HD$  and we will refer to it as *structured evolution strategies* as they do in [4].

Evolution strategies can exploit parallelism as well since workers need to communicate only scalars which are the function evaluations and the random seeds used when generating the random perturbation directions.

### D. Coded Computation

Coded computation has recently been introduced in [6] for speeding up distributed matrix multiplication. The idea is to add redundancy to computations by using linear codes in order to make it possible to recover the output of the multiplication without having all worker outputs. The fact that matrix multiplication is a linear operation plays a key role in coded computation. The scheme in [6] is based on MDS codes. More recently, the authors in [9] proposed a coded computation mechanism for serverless computing platforms using polar codes instead of MDS codes. Polar codes have low encoding and decoding complexity  $O(N \log N)$ .

This work draws ideas from the coded computation literature. In coded computation, for computing  $Ax$ , one encodes submatrices  $A_i$  and decodes worker outputs  $A_i x$ . In coded optimization, the perturbation directions are encoded and function evaluations are decoded.

Many of the coded computation schemes have the restriction that they are designed to work with finite field data and not applicable to full-precision inputs. The decoding algorithm introduced in [9] for polar codes can however work with full-precision data. Since we do not restrict the output of the black-box function  $f$  to take values in a finite field, we

need a decoder that can operate on full-precision real values. Quantization is an option, though it would introduce noise, which can be avoided by using a decoding algorithm that can work with full-precision data. We note that it is possible to use the decoder in [9] with a slight modification to account for the difference in the encoding kernels ( $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$  for polar codes, and  $\begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix}$  for the Hadamard transform). We discuss the decoding in detail in the next section.

### III. CODED BLACK-BOX OPTIMIZATION

In this section, we present the proposed method for speeding up distributed black-box optimization. We start by introducing some more notation and definitions that we will be using. The derivative of a differentiable function  $f$  at a point  $\theta$  along the unit vector direction  $e_i$  is the  $i$ 'th component of the gradient  $\nabla f$ . More precisely, we have

$$\frac{\partial f}{\partial \theta_i} = e_i^T \nabla f. \quad (3)$$

The directional derivative along  $v$  is defined as follows

$$\nabla_v f = \lim_{\delta \rightarrow 0} \frac{f(\theta + \delta v) - f(\theta)}{\delta}. \quad (4)$$

If the function is differentiable at a point  $\theta$ , then the directional derivative exists along any direction  $v$  and is a linear map [13]. In this case we have

$$\nabla_v f = v^T \nabla f. \quad (5)$$

When we do not have access to exact gradients, we can employ a *numerical directional derivative* by choosing a small  $\delta$  in

$$\nabla_v f \approx \frac{f(\theta + \delta v) - f(\theta)}{\delta}. \quad (6)$$

Note that the approximation in (6) is not symmetric, that is, we only perturb the parameters along  $+v$ . We instead use the symmetric version of (6) for approximating derivatives in which the parameters are perturbed along both the directions  $-v$  and  $+v$ :

$$\nabla_v f \approx \frac{f(\theta + \delta v) - f(\theta - \delta v)}{2\delta}. \quad (7)$$

If we consider the Taylor series expansion for  $f(\theta + \delta v)$ , we obtain

$$f(\theta + \delta v) = f(\theta) + \delta \nabla f^T v + \frac{\delta^2}{2} v^T \mathbf{H} v + \mathcal{O}(\delta^3) \quad (8)$$

and for  $f(\theta - \delta v)$ , we have

$$f(\theta - \delta v) = f(\theta) - \delta \nabla f^T v + \frac{\delta^2}{2} v^T \mathbf{H} v + \mathcal{O}(\delta^3) \quad (9)$$

where  $\mathcal{O}(\delta^3)$  is a third order error term and  $\mathbf{H}$  is the Hessian matrix for  $f$ . Substituting these expansions in (7), we obtain

$$\begin{aligned} \nabla_v f &\approx \frac{2\delta \nabla f^T v + \mathcal{O}(\delta^3)}{2\delta} \\ &= \nabla f^T v + \mathcal{O}(\delta^2). \end{aligned} \quad (10)$$

This shows that by choosing  $\delta$  small, we can make the numerical directional derivative approximately linear in  $\nabla f$ .

Our proposed method makes use of this assumption that gradient estimates are approximately linear in their directions  $v$  to ensure that coding can be applied to directional derivative estimates. To make this more concrete, let us consider the construction given in Figure 1. The block  $H$  simply corresponds to the linear transformation  $H = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$ .

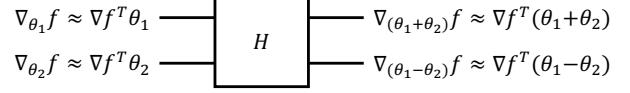


Fig. 1. 2-by-2 construction based on Hadamard transformation.

In Figure 1, if we know the estimates for  $\nabla_{(\theta_1+\theta_2)} f$  and  $\nabla_{(\theta_1-\theta_2)} f$ , we can compute the estimates for  $\nabla_{\theta_1} f$  and  $\nabla_{\theta_2} f$  because the directional derivative estimates are approximately linear in their corresponding directions. Furthermore, if we know the estimate for  $\nabla_{\theta_1} f$ , then it is sufficient to know only one of the estimates for  $\nabla_{(\theta_1+\theta_2)} f$  or  $\nabla_{(\theta_1-\theta_2)} f$  in order to be able to compute the estimate for  $\nabla_{\theta_2} f$ . This would happen if, for example,  $\theta_1$  is the zero vector (i.e. frozen direction) because the estimate for  $\nabla_{\theta_1} f$  would be zero and either of the directional derivative estimates from the right-hand side would be enough for us to obtain the estimate of  $\nabla_{\theta_2} f$ .

We refer to the directions  $\theta_1$ ,  $\theta_2$ ,  $(\theta_1 + \theta_2)$ ,  $(\theta_1 - \theta_2)$  as perturbation directions.

We now summarize the proposed method and detailed description of each item will follow:

- Encode all the unit vectors in  $\mathbb{R}^d$  to obtain the encoded perturbation directions.
- Assign each perturbation direction to a worker and have workers compute their directional derivative estimates using (7).
- Master starts collecting worker outputs.
- When a decodable set of worker outputs (i.e. directional derivative estimates) is available, the master node decodes these outputs to obtain an estimate for the gradient.
- The master node then makes a gradient update to find the next iterate for the parameter  $\theta$ .
- Repeat until convergence or for a desired number of iterations.

The above procedure assumes that we want to estimate all entries of the gradient, but it is possible estimate only a portion of gradient entries by encoding only the unit vectors corresponding to the desired entries.

We note that we can always check whether decoding helps in obtaining a better objective function compared to the the structured evolution strategies and make the update accordingly. Because the decoding step is quite fast thanks to the polar decoder, having decoded the outputs but ended up not using the decoded estimate is not a computational burden.

#### A. Encoding

Encoding is done based on the Hadamard transformation whose kernel is shown in Figure 2. Encoding is done the

same way as polar encoding is done. Hence, the computational complexity of encoding is only  $N \log N$  operations.

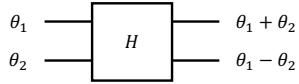


Fig. 2. 2-by-2 Hadamard transformation of the perturbation directions.

The way polar coding for the erasure channel works is that one computes the erasure probabilities of the transformed channels and freezes the channels with the highest erasure probabilities and the remaining channels are used for information. We refer the reader to [5] and [9] for more details on polar codes. In channel coding, freezing channels corresponds to sending known bits, e.g., zeros, at the receiver. In this work, freezing channels corresponds to setting it to zero coordinates such that the directional derivative along these vectors is zero. This makes it possible, when decoding, to take the value of the derivative estimates for frozen channels to be zero. For the information channels, we simply send in unit vectors.

Encoding for a function of 3 variables using  $N = 4$  workers is shown in Figure 3. The resulting 4 output vectors are the perturbation directions. In this construction, the rate

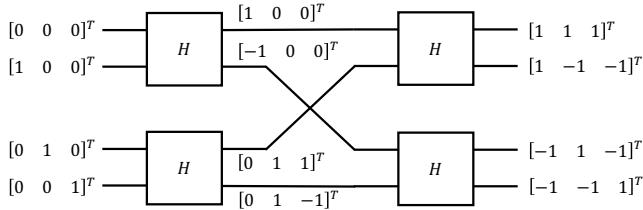


Fig. 3. Example encoding.

is  $\frac{3}{4}$  since 3 out of 4 channels are used for sending in unit vectors. When the erasure probabilities of the transformed channels are computed (see [5]), one will see that the worst one corresponds to the first index. Hence, the first channel is the frozen channel and the remaining 3 channels are the information channels. The frozen channel is set to the zero vector and the information channels are set to unit vectors. Sending in the zero vector for the frozen channel makes sense because perturbing by the zero vector is the same as not perturbing the variables and hence we get  $f(\theta + 0) - f(\theta - 0) = 0$ . This is important because during decoding, we will not have to do any computations to evaluate the value of the frozen inputs as we know they are zero.

If we wish to use  $N = 8$  workers instead of 4, we would set the worst 5 channels to zero vectors and the remaining best three channels would be set to the unit vectors  $e_1, e_2, e_3$ . The rate in this case would be  $\frac{3}{8}$ , and this construction would be more straggler-resilient since we can recover the gradient estimate in the presence of even more stragglers compared to the  $N = 4$  case.

**Embedding interpretation:** It is also possible to perform the encoding step slightly differently for a different view on

freezing channels. Instead of setting frozen channels to zero vectors, one can increase the dimension of the inputs from  $d$  to  $N$  and set the frozen channels to unit vectors  $e_j$  where  $j \in \{d+1, d+2, \dots, N\}$ . Because the function  $f(\theta)$  accepts  $d$ -dimensional inputs, we could embed  $f(\theta)$  into a higher dimension, that is, we could define  $\tilde{f}(\tilde{\theta})$  where  $\tilde{\theta} \in \mathbb{R}^N$  and  $\tilde{f}(\tilde{\theta}) = f(\theta)$  if  $\tilde{\theta}_i = \theta_i$  for  $i \in \{1, 2, \dots, d\}$ . The advantage of this approach is that the output of the encoding step will be equal to the  $H$  matrix with permuted rows.

### B. Decoding

The sequential polar code decoder given in [9] works for linear operations such as matrix-vector multiplication with full-precision data (i.e. does not require finite field data). Because we assume that the gradient estimates  $\nabla_{h_i} f$  are approximately linear, we use the same decoder in [9] by only modifying it slightly so that it decodes according to the Hadamard kernel and not the polar code kernel. The decoder is given in Alg. 1 along with the subroutine *decodeRecursive(i,j)* in Alg. 2. The decoder uses the notation  $D_{ij}$  to denote the value corresponding to the node  $i$  in the  $j$ 'th level in the coding circuit. Figure 4 shows this notation in the coding circuit for  $N = 4$ . The blocks  $W$  in Figure 4 represent workers. The notation  $I_{D_{ij}}$  in Alg. 1 indicates whether the data for the node  $i$  in the  $j$ 'th level is available or erased (i.e., not computed yet). *checkDecodability* is another subroutine that checks whether it is possible to decode the available outputs given the indicators of output availabilities. It checks decodability of outputs by running the decoder on the indicators and if decoding fails, it returns false, otherwise, it returns true.

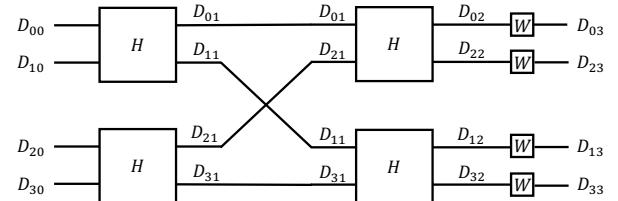


Fig. 4.  $N = 4$  construction.

### C. Multiplication by $D$

Structured evolution strategies method is usually based on exploring the parameter space along the rows of  $HD$  instead of  $H$  alone.  $D$  is a diagonal matrix with random diagonal entries distributed according to Rademacher distribution ( $\pm 1$  with probability 0.5). So far, we have only considered the  $H$  matrix for perturbation directions. It is possible to incorporate the diagonal matrix  $D$  into our method as well.

Multiplying  $H$  by  $D$  from the right corresponds to multiplying all the entries of the  $i$ 'th column of  $H$  by  $D_{ii}$  for all  $i$ . In this case, instead of computing an estimate for the directional derivative along a direction of  $v$ , we will approximate the directional derivative with respect to the direction  $Dv$ . The derivations we have done for  $v$  also apply for these randomized directions  $Dv$ . Decoding will work the

---

**Algorithm 1:** Decoding algorithm.

---

**Input:** Indices of the frozen channels  
**Result:** estimate of  $\nabla f$  ▷ Part I  
 Initialize  $I_{D_{:,0}} = [I_{D_{0,0}}, I_{D_{1,0}}, \dots, I_{D_{N-1,0}}] = [0, \dots, 0]$

**while**  $I_{D_{:,0}}$  not decodable **do**  
 | update  $I_{D_{:,0}}$   
 | checkDecodability( $I_{D_{:,0}}$ )  
**end**  
 Initialize an empty list  $y$  ▷ Part II  
**for**  $i \leftarrow 0$  **to**  $N - 1$  **do**  
 |  $D_{i,0} = \text{decodeRecursive}(i, 0)$   
 | **if** node  $i$  is a data node **then**  
 | |  $y = [y; D_{i,0}]$   
 | **end**  
 | **if**  $i \bmod 2 = 1$  **then** ▷ forward prop  
 | | **for**  $j \leftarrow 0$  **to**  $\log_2 N$  **do**  
 | | | **for**  $l \leftarrow 0$  **to**  $i$  **do**  
 | | | | compute  $D_{lj}$  if unknown  
 | | | **end**  
 | | **end**  
 | **end**  
**end**  
 return  $y$

---

same as before, except, after the decoding algorithm is run on the outputs, we need to divide the  $i$ 'th entry of the decoded gradient estimate by  $D_{ii}$  to obtain the entries of the gradient estimate.

#### IV. SPEEDING UP DISTRIBUTED COMPUTATION OF ADVERSARIAL ATTACKS

Deep neural networks have found success in many different tasks. Their robustness is a critical performance criterion for many applications. In the literature, there are many works that propose methods to generate adversarial inputs to attack neural networks to investigate their robustness. Attack settings may differ in terms of levels of knowledge on the neural network model. Some works assume knowledge on the architecture of the neural network and this setting is referred to as open-box setting. On the other hand, if we only assume we have access to the outputs of the neural network, then this is called the black-box setting. In this section, we are interested in applying our proposed method to the black-box setting.

We apply our proposed method to the setting of finding adversarial examples based on black-box optimization. We adopt the settings in [14] where they develop black-box attacks to neural network based classifiers and the goal is to find a misclassified image that looks like a given image  $\theta_0$ . They pose the following optimization problem for finding adversarial examples: minimize  $f(\theta)$ , where

$$f(\theta) = \|\theta - \theta_0\|_2^2 + c \max\{\max_{i \neq t} [F(\theta)]_i - \log[F(\theta)]_t, -\kappa\}. \quad (11)$$

---

**Algorithm 2:** decodeRecursive( $i, j$ )

---

**Input:** Node  $i \in [0, N - 1]$ , level  $j \in [0, \log_2 N]$   
**Result:**  $I_{D_{ij}}$  and modifies  $D$   
**if**  $j = \log_2 N$  **then** return  $I_{D_{i,\log_2 N}}$  ▷ base case 1  
**if**  $I_{D_{ij}} = 1$  **then** return 1 ▷ base case 2  
 $I_{D_{i,(j+1)}} = \text{decodeRecursive}(i, j + 1)$   
 $I_{D_{\text{pair}(i),(j+1)}} = \text{decodeRecursive}(\text{pair}(i), j + 1)$   
**if**  $i$  is upper node **then**  
 | **if**  $I_{D_{i,(j+1)}} \text{ AND } I_{D_{\text{pair}(i),(j+1)}} = 1$  **then**  
 | | compute  $D_{ij}$   
 | | return 1  
**end**  
**else**  
 | **if**  $I_{D_{i,(j+1)}} \text{ OR } I_{D_{\text{pair}(i),(j+1)}} = 1$  **then**  
 | | compute  $D_{ij}$   
 | | return 1  
**end**  
**end**  
 return 0

---

$F(\theta)$  represents the class probabilities of the neural network when we feed in  $\theta$  as the input. The subscript  $i$  in  $[F(\theta)]_i$  means the class probability for the  $i$ 'th class.  $c$  and  $\kappa$  are non-negative parameters for tuning the attack. Note that the objective in (11) is for a targeted attack, meaning we wish to find an image that is classified as class  $t$ . If one is interested in finding an example misclassified but not necessarily to a given class, then this is called an untargeted attack. The optimization problem to be solved for untargeted attacks is given in [14] as

$$f(\theta) = \|\theta - \theta_0\|_2^2 + c \max\{\log[F(\theta)]_{t_0} - \max_{i \neq t_0} \log[F(\theta)]_i, -\kappa\} \quad (12)$$

where  $t_0$  is the correct class for the image that we wish to misclassify. Note that since we assume we have no information about the neural network  $F(\theta)$ , we cannot compute the gradient of the functions in (11) and (12).

The proposed approach in [14] involves selecting a given number of parameters in every iteration and estimating the derivatives only for these parameters. Hence, in every iteration, only a number of selected parameters are updated. The updates are done using the ADAM algorithm [15]. We present results regarding the use of our method in black-box adversarial attacks problem in numerical results section.

#### V. NUMERICAL RESULTS

We now present results showing the performance of our proposed framework.

##### A. An $l_1$ -norm Based Objective as the Black-Box Function

In this subsection, we present experiment results for optimizing a known objective function  $f(\theta) = \|A\theta - b\|_1$  as a test case. Even though we experiment with a known function, we still use the black-box optimization algorithms assuming we do not have the exact gradients.

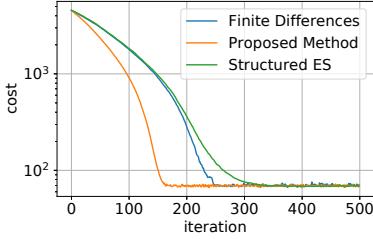


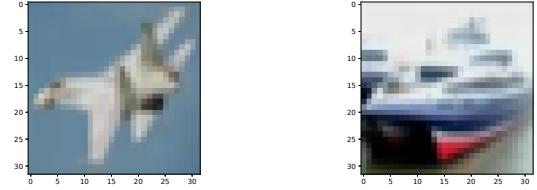
Fig. 5. We optimize the function  $f(\theta) = \|A\theta - b\|_1$  where  $A \in \mathbb{R}^{200 \times 32}$  and  $y \in \mathbb{R}^{200}$  using black-box optimization methods.

Figure 5 shows the cost as a function of iterations when we use gradient descent algorithm with gradient estimates obtained by the finite differences method, the proposed method, and the structured evolution strategies method. For fairness in comparing these methods, we used straggler-resilient versions of the finite differences method and the structured evolution strategies method in obtaining these results. To clarify, for the finite differences method, we wait for the first arriving 16 worker outputs out of the total 32 outputs to make a gradient update. For the structured evolution strategies method, the perturbation directions come from the rows of  $HD$  matrix which is a  $32 \times 32$  matrix, so we wait for the first arriving 16 outputs out of the 32. Finally, the way we implement the proposed method is that we use a rate of  $\frac{1}{2}$  with a total of 64 workers and wait for the first decodable set of outputs out of a total of 64 outputs. Figure 5 illustrates that having all the entries of the gradient estimate through decoding leads to faster convergence than having only a half of the entries of the gradient estimate. It also shows that the proposed method results in faster convergence compared to the structured evolution strategies method.

### B. Black-Box Adversarial Attacks

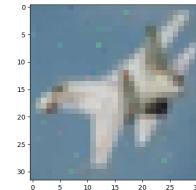
We now present some experiment results for our proposed method applied to finding adversarial examples. We have tested our proposed method for solving the optimization problem for targeted black-box attacks (objective of which is given in (11)) to a deep learning model. The model we considered is the VGG-16 architecture [16] trained on CIFAR-10 dataset [17] for image classification. VGG-16 architecture is a 16-layer convolutional neural network. CIFAR-10 is an image dataset with a total of 60000 images of size  $32 \times 32 \times 3$  (50000 training and 10000 test images) with 6000 images per class for a total of 10 classes.

Before we delve into the details of the experiments, we show two original input images from CIFAR-10 and the generated adversarial images that look almost the same as the original images but are misclassified. Figure 6 shows the original images  $\theta_0$  which are correctly classified by the neural network  $F(\theta)$ . Figure 7 shows the generated adversarial images obtained by solving the black-box optimization problem for targeted attacks where the target class is the dog class. The adversarial images in Figure 7 are both misclassified as dog. These adversarial images were obtained by using our proposed method.



(a) Airplane classified correctly (b) Ship classified correctly

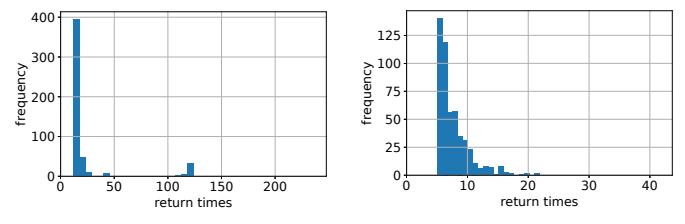
Fig. 6. Original images  $\theta_0$ .



(a) Airplane classified as dog (b) Ship classified as dog

Fig. 7. Adversarial images generated using the proposed method.

We now discuss details of how we obtained these adversarial images. For obtaining these results, we have solved the problem of minimizing  $f(\theta)$  where  $f(\theta)$  is given in (11). We have set  $\kappa = 0$ ,  $c = 0.1$ . We have solved this optimization problem using three different methods and compared them in Figure 9 via a run time simulation. The simulation was done by sampling worker run times from the distributions shown in Figure 8. Figure 8(a) shows the empirical histogram of worker times which were obtained for the AWS Lambda platform, which is a serverless computing platform. Figure 8(b) shows the histogram of worker times generated from the shifted exponential distribution.



(a) AWS Lambda return times (b) Exponential distribution

Fig. 8. Run time distributions used in simulations.

Figure 9 illustrates that it is possible to eliminate the straggling workers effect by using the proposed method. The way finite differences and structured evolution strategies methods were implemented is not straggler-resilient in these simulations, that is, all worker outputs are waited for in every gradient update. We note that this is different from the results on the optimization of the  $l_1$ -norm based function given in previous subsection because we used straggler-resilient versions of the methods in those experiments. The point of the simulation results given in Figure 7 is not necessarily

to compare the run times of each method but instead to showcase that our proposed method works for the recent problem of adversarial attacks to deep neural networks.

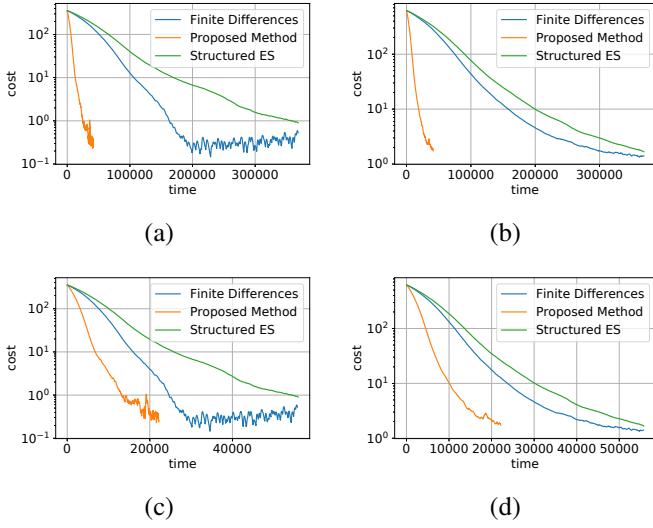


Fig. 9. Simulated plots of cost as a function of time (all methods were run for 3000 iterations). a,c: Airplane image classified as dog. b,d: Ship image classified as dog. a,b: Worker return times sampled from AWS Lambda return times shown in Figure 8(a). c,d: Worker return times sampled from the exponential distribution shown in Figure 8(b).

## VI. CONCLUSION

We have introduced a novel distributed black-box optimization mechanism based on coded perturbation directions. Because of the redundancy inserted during encoding, the proposed method is resilient against straggling workers in distributed computing platforms. This makes our method suitable for large scale serverless cloud-based systems where nodes are unreliable. Our work utilizes coding theoretic tools and ideas from coded distributed computation methods. Furthermore, the proposed method provides a connection between two black-box optimization methods: finite differences and structured evolution strategies, in a way that it is possible to go from one to the other and choose the one that drives the cost down more effectively. We have provided an argument why coded gradient estimates would work for a variety of cost functions and shown this through numerical examples. There are several open directions for future work. In particular, our coding Hadamard transform structured coding method is analogous to subsampled Hadamard constructions in random projections [18], [19], [20]. An interesting direction is to combine these structured embeddings for faster distributed optimization [21], [22].

## ACKNOWLEDGEMENTS

This work was supported in part by the National Science Foundation under grant IIS-1838179.

## REFERENCES

- [1] A. R. Conn, K. Scheinberg, and L. N. Vicente, *Introduction to derivative-free optimization*. Siam, 2009, vol. 8.
- [2] J. Dean and L. A. Barroso, "The tail at scale," *Communications of the ACM*, vol. 56, no. 2, pp. 74–80, 2013.
- [3] T. Hoefer, T. Schneider, and A. Lumsdaine, "Characterizing the influence of system noise on large-scale applications by simulation," in *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Computer Society, 2010, pp. 1–11.
- [4] K. Choromanski, M. Rowland, V. Sindhwani, R. E. Turner, and A. Weller, "Structured evolution with compact architectures for scalable policy optimization," *arXiv preprint arXiv:1804.02395*, 2018. [Online]. Available: <http://arxiv.org/abs/1804.02395>
- [5] E. Arikan, "Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels," *IEEE Transactions on Information Theory*, vol. 55, pp. 3051–3073, 2009.
- [6] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," *IEEE Transactions on Information Theory*, vol. 64, no. 3, pp. 1514–1529, March 2018.
- [7] A. Severinsson, A. G. i Amat, and E. Rosnes, "Block-diagonal and lt codes for distributed computing with straggling servers," *IEEE Transactions on Communications*, 2018.
- [8] A. Mallick, M. Chaudhari, and G. Joshi, "Rateless codes for near-perfect load balancing in distributed matrix-vector multiplication," *CoRR*, vol. abs/1804.10331, 2018. [Online]. Available: <http://arxiv.org/abs/1804.10331>
- [9] B. Bartan and M. Pilanci, "Polar coded distributed matrix multiplication," *arXiv preprint arXiv:1901.06811*, 2019.
- [10] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "Polynomial codes: An optimal design for high-dimensional coded matrix multiplication," *Adv. in Neural Info. Proc. Systems (NIPS) 30*, pp. 4406–4416, 2017.
- [11] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "Straggler mitigation in distributed matrix multiplication: Fundamental limits and optimal coding," *CoRR*, vol. abs/1801.07487, 2018. [Online]. Available: <http://arxiv.org/abs/1801.07487>
- [12] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, "Evolution strategies as a scalable alternative to reinforcement learning," *arXiv preprint arXiv:1703.03864*, 2017.
- [13] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [14] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh, "Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models," *arXiv preprint arXiv:1708.03999*, 2017.
- [15] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [16] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [17] A. Krizhevsky, "Learning multiple layers of features from tiny images," 2009.
- [18] M. Pilanci and M. J. Wainwright, "Randomized sketches of convex programs with sharp guarantees," *IEEE Trans. Info. Theory*, vol. 9, no. 61, pp. 5096–5115, September 2015.
- [19] —, "Iterative hessian sketch: Fast and accurate solution approximation for constrained least-squares," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1842–1879, 2016.
- [20] Y. Yang, M. Pilanci, M. J. Wainwright, *et al.*, "Randomized sketches for kernels: Fast and optimal nonparametric regression," *The Annals of Statistics*, vol. 45, no. 3, pp. 991–1023, 2017.
- [21] M. Pilanci and M. J. Wainwright, "Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence," *SIAM Journal on Optimization*, vol. 27, no. 1, pp. 205–245, 2017.
- [22] M. Pilanci, "Fast randomized algorithms for convex optimization and statistical estimation," Ph.D. dissertation, UC Berkeley, 2016.