

Heterogeneous Computation Assignments in Coded Elastic Computing

Nicholas Woolsey, Rong-Rong Chen, and Mingyue Ji
Department of Electrical and Computer Engineering, University of Utah
Salt Lake City, UT, USA
Email: {nicholas.woolsey@utah.edu, rchen@ece.utah.edu, mingyue.ji@utah.edu}

Abstract—We study the optimal design of a heterogeneous coded elastic computing (CEC) network where machines have varying relative computation speeds. CEC introduced by Yang *et al.* is a framework which mitigates the impact of elastic events, where machines join and leave the network. A set of data is distributed among storage constrained machines using a Maximum Distance Separable (MDS) code such that any subset of machines of a specific size can perform the desired computations. This design eliminates the need to re-distribute the data after each elastic event. In this work, we develop a process for an arbitrary heterogeneous computing network to minimize the overall computation time by defining an optimal computation load, or number of computations assigned to each machine. We then present an algorithm to define a specific computation assignment among the machines that makes use of the MDS code and meets the optimal computation load.

I. INTRODUCTION

Coding is an effective tool to speed up distributed computing networks. Examples include Coded Distributed Computing (CDC) for MapReduce-like frameworks [1] and coded data shuffling for distributed machine learning [2]–[4], where code designs minimize the communication load by trading increased computation and/or storage on each machine. Another example is to use codes to mitigate the straggler effect in applications such as matrix multiplications [5], [6], where any subset of machines with a cardinality larger than the recovery threshold can recover the matrix multiplication. This eliminates the need to wait for the computation of slow machines.

Coded Elastic Computing (CEC) was designed by Yang *et al.* in 2019 to mitigate the impact of *preempted* machines on a storage limited computing network [7]. As opposed to stragglers, whose identities are unknown when computations are assigned, the preempted (unavailable) machines are known. In elastic computing, computations are performed over many times steps and between each time step an *elastic event* may occur where machines become preempted or available again. Computations are performed on the same set of data, for example a matrix, but the computations change each time step. For example, in each time step the data matrix may be multiplied with a different vector. In each time step, the goal becomes to assign computations among the available

machines. A naive approach is to assign each machine a non-overlapping part of the data. However, this is inefficient as the storage has to be redefined with each elastic event.

In CEC the storage of each machine is defined once using a Maximum Distance Separable (MDS) code and remains unchanged between elastic events. The data is split into L equal sized, disjoint data sets and each machine stores a coded combination of these sets. In this way, each machine only stores an equivalent of a $\frac{1}{L}$ fraction of the data. Furthermore, any computation can be resolved by combining the coded computation results of L machines. Then, given a set of available machines, each computation is assigned to L machines. In the original CEC scheme of [7], the authors proposed a “cyclic” computation assignment such that each machine is assigned the same number computations.

The recent work [8] also studies CEC and aims to maximize the overlap of the task assignments between computation time steps. With each elastic event, the computation assignment must change. In the cyclic approach in [7], the assignments in the current time step are independent of assignments in previous time steps. In [8], the authors design assignment schemes to minimize the changes in the assignments between time steps. In some cases, the proposed assignment schemes were shown to achieve zero transition waste, or minimize the amount of new local computations at the machines. However, both [7] and [8] only study homogeneous computing networks.

In this paper, we propose a CEC framework optimized for a heterogeneous network where machines have varying computation speeds. In this setting, more computations are assigned to faster machines and less computations to slower machines to minimize the maximum local computation time among the machines. This assignment problem is non-trivial since by the MDS code design we still require that each computation is assigned to L machines. We propose and solve an optimization problem to find the optimal computation load, or amount of computations assigned to each machine. We then show an assignment exists that yields this computation load and design a low complexity algorithm to find such an assignment.¹ Our proposed CEC design works for an arbitrary set of machine speeds and requires a number of computation assignments at most equal to the number of available machines.

Work supported through the National Science Foundation grants CCF-1817154 and SpecEES-1824558 and the INL Laboratory Directed Research and Development (LDRD) Program under DOE Idaho Operations Office Contract DE-AC07-05ID14517.

¹The CEC assignment algorithm is adapted from our heterogeneous private information retrieval (PIR) storage placement algorithm of [9].

Notation Convention: We use $|\cdot|$ to represent the cardinality of a set or the length of a vector and $[n] := [1, 2, \dots, n]$.

II. NETWORK MODEL AND PROBLEM FORMULATION

We consider a set of N machines. Each stores a coded matrix derived from a $q \times r$ data matrix, \mathbf{X} . The coded matrices are defined by an $N \times L$ MDS generator matrix $\mathbf{G} = (g_{n,\ell})$ such that any L rows of \mathbf{G} are invertible. The data matrix, \mathbf{X} , is row-wise split into L disjoint, $\frac{q}{L} \times r$ matrices, $\mathbf{X}_1, \dots, \mathbf{X}_L$. Each machine $n \in [N]$ stores the $\frac{q}{L} \times r$ coded matrix

$$\tilde{\mathbf{X}}_n = \sum_{\ell=1}^L g_{n,\ell} \mathbf{X}_\ell. \quad (1)$$

The machines collectively perform matrix-vector computations² over multiple times steps. In a given time step only a subset of the N machines are available to perform matrix computations. More specifically, in time step t , a set of available machines $\mathcal{N}_t \subseteq [N]$ aims to compute

$$\mathbf{y}_t = \mathbf{X} \mathbf{w}_t \quad (2)$$

where \mathbf{w}_t is some vector of length r . The machines of $[N] \setminus \mathcal{N}_t$ are preempted and we assume the number of available machines is $N_t = |\mathcal{N}_t| \geq L$.

The machines of \mathcal{N}_t do not compute \mathbf{y}_t directly. Instead, each machine $n \in \mathcal{N}_t$ computes the set

$$\mathcal{V}_n = \left\{ v = \tilde{\mathbf{X}}_n^{(i)} \mathbf{w}_t : i \in \mathcal{W}_n \right\} \quad (3)$$

where $\tilde{\mathbf{X}}_n^{(i)}$ is the i -th row of $\tilde{\mathbf{X}}_n$ and $\mathcal{W}_n \subseteq [\frac{q}{L}]$ is the set of rows assigned to machine n in time step t . Furthermore, we define the computation load vector, $\boldsymbol{\mu}$, such that

$$\mu[n] = \frac{|\mathcal{W}_n|}{\left(\frac{q}{L}\right)}, \quad \forall n \in \mathcal{N}_t \quad (4)$$

is the fraction of rows computed by machine n in time step t . Note that, $\boldsymbol{\mu}$, \mathcal{V}_n and \mathcal{W}_n change with each time step, but reference to t is omitted for ease of disposition. Moreover, the machines have varying computation speeds defined by the strictly positive vector, \mathbf{s} , which is fixed over all time steps. Here, computation speed is the number of row multiplications per unit time. The computation time is dictated by the machine that takes the most time to perform its assigned computations such that the computation time in a particular time step is

$$c(\boldsymbol{\mu}) = \max_{n \in \mathcal{N}_t} \frac{\mu[n]}{s[n]}. \quad (5)$$

In a given time step, for each $i \in [\frac{q}{L}]$, L machines perform the vector-vector multiplication with the i -th row of their local coded matrix and \mathbf{w}_t . The results are sent to a master node which can resolve the elements of \mathbf{y}_t by the MDS code design. To assign each row to L machines, we define F disjoint sets of rows, $\mathcal{M}_t = (\mathcal{M}_1, \dots, \mathcal{M}_F)$ whose union is $[\frac{q}{L}]$. Then, F sets of L machines, $\mathcal{P}_t = (\mathcal{P}_1, \dots, \mathcal{P}_F)$, are defined such

²It can be shown that our CEC designs can also operate on the other applications outlined in [10] and not just matrix-vector multiplications.

that $\mathcal{P}_f \subseteq \mathcal{N}_t$ and $|\mathcal{P}_f| = L$ for all $f \in [F]$. The rows of \mathcal{M}_f are assigned to the machines of \mathcal{P}_f . The rows computed by machine $n \in \mathcal{N}_t$ in time step t are in the set

$$\mathcal{W}_n = \bigcup \{ \mathcal{M}_f : f \in [F], n \in \mathcal{P}_f \} \quad (6)$$

and $\boldsymbol{\mu}$ is a function of $(\mathcal{M}_t, \mathcal{P}_t)$. The sets $\mathcal{M}_1, \dots, \mathcal{M}_F$ and $\mathcal{P}_1, \dots, \mathcal{P}_F$ and F may vary with each time step.

In a given time step t , our goal is to define the computation assignments, \mathcal{M}_t and \mathcal{P}_t , such that the resulting computation load vector defined in (4) has the minimum computation time. In time step t , given \mathcal{N}_t and \mathbf{s} , the optimal computation time, c^* , is the infimum of computation time defined by all possible computation assignments, $(\mathcal{M}_t, \mathcal{P}_t)$, such that

$$\begin{aligned} c^* &= \inf_{(\mathcal{M}_t, \mathcal{P}_t)} c(\boldsymbol{\mu}(\mathcal{M}_t, \mathcal{P}_t)) \\ \text{s.t.} \quad &\bigcup_{\mathcal{M}_f \in \mathcal{M}_t} \mathcal{M}_f = \left[\frac{q}{L} \right], \\ &|\mathcal{P}_f| = L \quad \forall \mathcal{P}_f \in \mathcal{P}_t, \\ &|\mathcal{M}_t| = |\mathcal{P}_t|. \end{aligned} \quad (7)$$

Remark 1: In Sections IV and V, we precisely solve the combinatorial optimization problem of (7). We decompose the problem into two sub-problems: 1) a convex optimization problem to find an optimal $\boldsymbol{\mu}$ without the consideration of a specific computation assignment and 2) a computation assignment problem. Moreover, we show that an optimal assignment, $(\mathcal{M}_t, \mathcal{P}_t)$, can be found via a low complexity algorithm.

III. AN EXAMPLE

There are a total of $N = 6$ machines where each has the storage capacity to store $\frac{1}{3}$ of a data matrix \mathbf{X} . In time step t , the machines have the collective goal of computing $\mathbf{y}_t = \mathbf{X} \mathbf{w}_t$ where \mathbf{w}_t is some vector. In order to allow for preempted machines, \mathbf{X} is split row-wise into $L = 3$ sub-matrices, \mathbf{X}_1 , \mathbf{X}_2 , and \mathbf{X}_3 and an MDS code is used to define the matrices $\{\tilde{\mathbf{X}}_n : n \in [N]\}$ which are stored among the machines. This placement is designed such that any element of \mathbf{y}_t can be recovered by obtaining the corresponding coded computation from any 3 machines. To recover the entirety of \mathbf{y}_t , we split the coded matrices into sets of rows, such that each set is used for computation at $L = 3$ machines.

The machines have relative computation speeds defined by

$$\mathbf{s} = [2, 2, 3, 3, 4, 4]. \quad (8)$$

Machines 5 and 6 are the fastest machines and can perform row computations twice as fast as machines 1 and 2. Machines 3 and 4 are the next fastest machines and can perform matrix computations 1.5 times as fast as machines 1 and 2. Our goal is to assign computations, or rows of the coded matrices, to the machines to minimize the overall computation time with the constraint that each computation is assigned to 3 machines.

In time step 1, there are no preempted machines and $\mathcal{N}_1 = \{1, \dots, 6\}$. We assign fractions of the rows to the machines defined by the computation load vector

$$\boldsymbol{\mu} = \left[\frac{1}{3}, \frac{1}{3}, \frac{1}{2}, \frac{1}{2}, \frac{2}{3}, \frac{2}{3} \right] \quad (9)$$

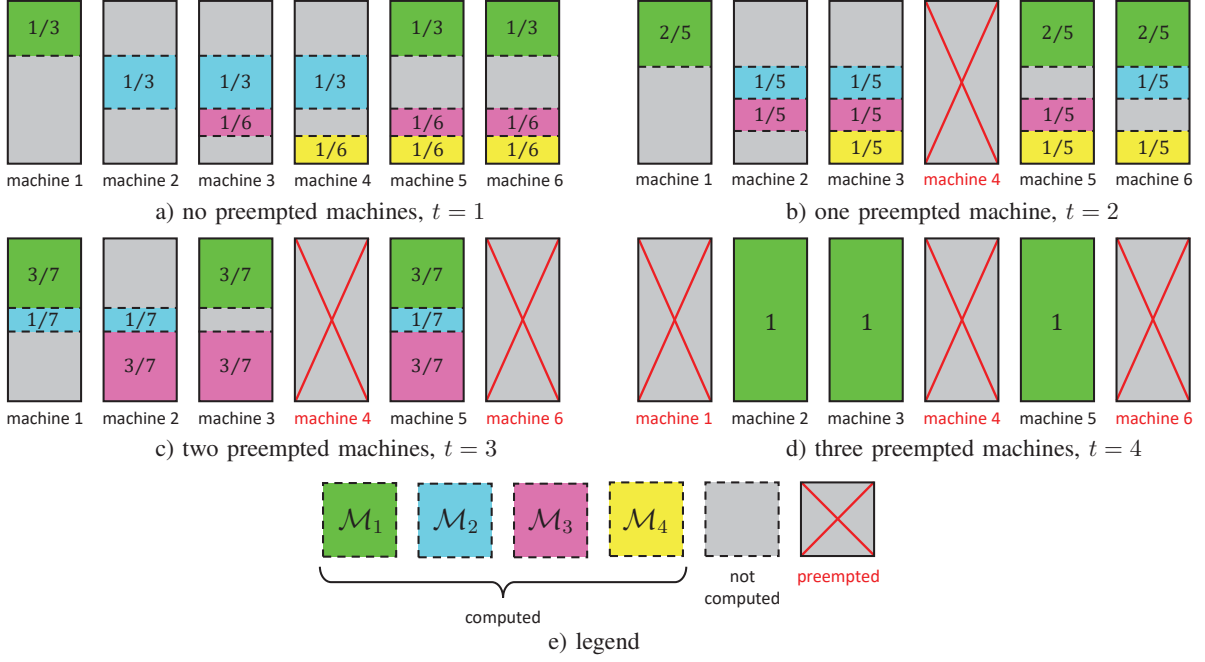


Fig. 1. Optimal computation assignments over 4 times steps on a heterogeneous CEC network.

such that machines 1 and 2 are assigned $\frac{1}{3}$, machines 3 and 4 are assigned $\frac{1}{2}$ and machines 5 and 6 are assigned $\frac{2}{3}$ of the rows of their respective coded matrices. We define μ such that it sums to $L = 3$ and each row can be assigned to 3 machines. Furthermore, based on the machine computation speeds, the machines finish at the same time to minimize the overall computation time. In Section IV, we outline the systematic approach to determine μ . Next, given μ , the rows of the coded matrices must be assigned. We define sets of rows, \mathcal{M}_1 , \mathcal{M}_2 , \mathcal{M}_3 , and \mathcal{M}_4 which are assigned to sets of machines \mathcal{P}_1 , \mathcal{P}_2 , \mathcal{P}_3 , and \mathcal{P}_4 , respectively. These sets are depicted in Fig. 1(a) where, for example, \mathcal{M}_1 contains the first $\frac{1}{3}$ of the rows assigned to machines $\mathcal{P}_1 = \{1, 5, 6\}$. Moreover, \mathcal{M}_2 contains the next $\frac{1}{3}$ of the rows assigned to machines $\mathcal{P}_2 = \{2, 3, 4\}$, \mathcal{M}_3 contains the next $\frac{1}{6}$ of the rows assigned to machines $\mathcal{P}_3 = \{3, 5, 6\}$ and \mathcal{M}_4 contains the final $\frac{1}{6}$ of the rows assigned to machines $\mathcal{P}_4 = \{4, 5, 6\}$. In Section V, we present Algorithm 1, which defines the computation assignment for general μ . By this assignment, the fraction of rows assigned to machine n sums to $\mu[n]$ and each row is assigned to $L = 3$ machines to recover the entirety of \mathbf{y}_1 .

In time step 2, $N_t = 5$ as machine 4 is preempted and is no longer available to perform computations. Therefore, the computations must be reassigned. First, we define

$$\mu = \left[\frac{2}{5}, \frac{2}{5}, \frac{3}{5}, 0, \frac{4}{5}, \frac{4}{5} \right] \quad (10)$$

which sums to $L = 3$ and minimizes the overall computation time. Given μ , we then use Algorithm 1, which aims to assign computations to a machine with the least remaining rows to

be assigned and $L - 1 = 2$ machines with the most remaining rows to be assigned. For example, in the first iteration, \mathcal{M}_1 is defined to contain the first $\frac{2}{5}$ of the rows and is assigned to machines $\mathcal{P}_1 = \{1, 5, 6\}$. After this iteration, machines 2, 5 and 6 require $\frac{2}{5}$ of the total rows to still be assigned to them and machine 3 requires $\frac{3}{5}$ of the total rows. In the next iteration, \mathcal{M}_2 contains the next $\frac{1}{5}$ of the rows and is assigned to $\mathcal{P}_2 = \{2, 3, 6\}$. Note that, only $\frac{1}{5}$ of the rows could be assigned in this iteration otherwise there would only be two machines, 3 and 5, which still require assignments and therefore, the remaining rows cannot be assigned to three machines. In the final two iterations, \mathcal{M}_3 and \mathcal{M}_4 contain $\frac{1}{5}$ of the previously unassigned rows and are assigned to the machines of $\mathcal{P}_3 = \{2, 3, 5\}$ and $\mathcal{P}_4 = \{3, 5, 6\}$, respectively. These assignments are depicted in Fig. 1(b).

Next, in time step 3, machines 4 and 6 are preempted. Similar to previous examples it is ideal to have machines 3 and 5 compute $1.5\times$ and $2\times$ the number of computations, respectively, compared to machines 1 and 2. However, this is not possible since each machine can be assigned at most a number of rows equal to the number of rows of the coded matrices. In this case, we assign all rows to the fastest machine, machine 5, and assign fractions of the rows to the remaining machines which sum up to 2. As a result, we define

$$\mu = \left[\frac{4}{7}, \frac{4}{7}, \frac{6}{7}, 0, 1, 0 \right]. \quad (11)$$

Then, Algorithm 1 defines, \mathcal{M}_1 , \mathcal{M}_2 and \mathcal{M}_3 , disjoint sets containing $\frac{3}{7}$, $\frac{1}{7}$ and $\frac{3}{7}$ of the rows respectively. Moreover, these row sets are assigned to the machines of $\mathcal{P}_1 = \{1, 3, 5\}$,

$\mathcal{P}_2 = \{1, 2, 5\}$ and $\mathcal{P}_3 = \{2, 3, 5\}$, respectively. These assignments are depicted in Fig. 1(c).

Finally, in time step 4, machines 1, 4 and 6 are preempted. To assign all the rows to $L = 3$ machines, each available machine is assigned all of the rows and

$$\boldsymbol{\mu} = [0, 1, 1, 0, 1, 0]. \quad (12)$$

In other words, \mathcal{M}_1 contains all rows and $\mathcal{P}_1 = \{2, 3, 5\}$. This is depicted in Fig. 1(d).

IV. OPTIMAL COMPUTATION LOAD VECTOR

We solve the exact optimization problem of (7). We start by introducing a relaxed optimization problem of (7) without considering an explicit computation assignment $(\mathcal{M}_t, \mathcal{P}_t)$. In this domain, we find the optimal computation load vector $\boldsymbol{\mu}^*$ and computation time $\hat{c}^* \leq c^*$. Then, in Section V, we show there is no gap between this relaxed optimization problem and (7) because there exists a computation assignment $(\mathcal{M}_t, \mathcal{P}_t)$ that yields the computation load vector $\boldsymbol{\mu}^*$ and computation time \hat{c}^* . Therefore, we find that $c^* = \hat{c}^*$.

Henceforth, WLOG, we assume $N_t = \{1, 2, \dots, N_t\}$ and ignore the computation load of any preempted machine which is simply 0.

A. A Relaxed Convex Optimization Problem

Given a computation speed vector \mathbf{s} , we define the optimal computation load vector $\boldsymbol{\mu}^*$ to be the solution to the following relaxed optimization problem:

$$\begin{aligned} \boldsymbol{\mu}^* &= \underset{\boldsymbol{\mu}}{\operatorname{argmin}} \max_{n \in [N_t]} \frac{\mu[n]}{s[n]} \\ \text{s.t. } &\sum_{n \in [N_t]} \mu[n] = L \\ &0 \leq \mu[n] \leq 1, \forall n \in [N_t], \end{aligned} \quad (13)$$

which can be shown to be convex. While computation assignments, $(\mathcal{M}_t, \mathcal{P}_t)$, are not explicitly considered in (13), we note that the key constraint of $\sum_{n \in [N_t]} \mu[n] = L$ is a relaxed version of that requirement on the computation assignment that each row should be assigned to L machines. When $N_t = L$, the solution to (13) is $\boldsymbol{\mu}^* = [1, \dots, 1]$. The analytical solution to (13) when $N_t > L$ is presented in Theorem 1.

Theorem 1: Assume that $N_t > L$ and $s[1] \leq s[2] \leq \dots \leq s[N_t]$. The optimal solution $\boldsymbol{\mu}^*$ to the optimization problem of (13) must take the following form

$$\boldsymbol{\mu}^*[n] = \begin{cases} \hat{c}^* s[n] & \text{if } 1 \leq n \leq k^* \\ 1 & \text{if } k^* + 1 \leq n \leq N_t, \end{cases} \quad (14)$$

where k^* is the largest integer in $[N_t - L + 1, N_t]$ such that

$$\frac{1}{s[k^* + 1]} < \hat{c}^* = \frac{k^* + L - N_t}{\sum_{n=1}^{k^*} s[n]} \leq \frac{1}{s[k^*]}. \quad (15)$$

Here, $\hat{c}^* = c(\boldsymbol{\mu}^*)$ is the maximum computation time among the N_t machines given the computation load assignment $\boldsymbol{\mu}^*$. The left side of (15) is ignored when $k^* = N_t$. \square

Theorem 1 is proved in our online version [11].

Remark 2: The two cases in (14) are determined by whether a machine n satisfies $\mu^*[n] = \hat{c}^* s[n]$ or $\mu^*[n] < \hat{c}^* s[n]$. For $1 \leq n \leq k^*$, the equality is achieved and we must have $0 < \mu^*[n] \leq 1$. When $k^* + 1 \leq n \leq N_t$, we have the strict inequality and $\mu^*[n] = 1$. The equality in (15) ensures that $\sum_{n=1}^{N_t} \mu^*[n] = L$; the right-most inequality ensures that $\mu^*[n] \leq \mu^*[k^*] = \hat{c}^* s[k^*] \leq 1$, for any $1 \leq n \leq k^*$; the left-most inequality ensures that for any $k^* + 1 \leq n \leq N_t$, we have $\mu^*[n] < \hat{c}^* s[n]$. Hence, the computation time \hat{c}^* is induced by the k^* slowest machines.

Since the optimization problem of (13) aims to minimize a convex function on a closed and convex set, the existence of an optimal solution is guaranteed. This ensures the existence of some $k^* \in [N_t - L + 1, N_t]$ such that (15) is satisfied. In the following, we provide a numerical procedure to find k^* . First, it is straightforward to verify that if the right-hand-side (RHS) inequality “ \leq ” of (15) is violated for $k^* = i$, then the left-hand-side (LHS) inequality “ $<$ ” of (15) must hold for $k = i - 1$. In other words, for any $i = N_t, \dots, N_t - L + 2$,

$$\text{If } \hat{c}_i^* > \frac{1}{s[i]}, \text{ then } \frac{1}{s[i]} < \hat{c}_{i-1}^*. \quad (16)$$

We first check $k^* = N_t$. If the RHS of (15) holds, then we have $k^* = N_t$. Otherwise, it follows from (16) that the LHS of (15) must hold for $k^* = N_t - 1$. If the RHS of (15) also holds for $k^* = N_t - 1$, then we have $k^* = N_t - 1$. Otherwise, it follows from (16) that the LHS of (15) must hold for $k^* = N_t - 2$. We continue this process by decreasing k^* until we find one value of k^* for which both sides of (15) hold. This process is guaranteed to terminate before reaching $k^* = N_t - L + 1$ for which the RHS of (15) always holds. Hence, this establishes the procedure to find k^* directly using (15).

B. Computation Load Examples

We return to the first example and explain how to find the optimal computation load vector. When $t = 1$, we have $N_t = 6, L = 3$. Given $\mathbf{s} = [2, 2, 3, 3, 4, 4]$, the largest k^* that satisfies (15) is $k^* = 6$, and thus $\hat{c}^* = 1/6$, $\boldsymbol{\mu}^* = \hat{c}^* \mathbf{s} = [\frac{1}{3}, \frac{1}{3}, \frac{1}{2}, \frac{1}{2}, \frac{2}{3}, \frac{2}{3}]$. Similarly, for $t = 2$, since machine 4 preempts, we have now $N_t = 5$, and $\mathbf{s} = [2, 2, 3, 4, 4]$ (we ignore any preempted machines). In this case, we have $k^* = 5$, and thus $\hat{c}^* = 1/5$, $\boldsymbol{\mu}^* = \hat{c}^* \mathbf{s} = [\frac{2}{5}, \frac{2}{5}, \frac{3}{5}, \frac{4}{5}, \frac{4}{5}]$. Similarly, for $t = 3$, we have $N_t = 4$, and $\mathbf{s} = [2, 2, 3, 4]$ because machines 4 and 6 preempt. Here, we have $k^* = 3$, $\hat{c}^* = 2/7$, and $\boldsymbol{\mu}^* = [\frac{4}{7}, \frac{4}{7}, \frac{6}{7}, 1]$. Note that, similar to the optimization problem of (13), the computation load of the preempted machines are ignored since they are simply 0, presenting a slight difference between the optimal computation load vectors presented in Section III.

V. OPTIMAL COMPUTATION ASSIGNMENT

We present a computation assignment $(\mathcal{M}_t, \mathcal{P}_t)$ that solves the optimization problem of (7). First, we show the existence of a computation assignment that yields the computation load vector $\boldsymbol{\mu}^*$ and computation time \hat{c}^* . Therefore, we show there is no gap between (7) and the relaxed optimization problem of

(13). Then, we provide an iterative algorithm that converges to such an assignment in N_t iterations.

Our goal is to assign computations among the machines such that each computation is assigned to L machines and the assignments satisfy μ^* . This is equivalent to the filling problem (FP) introduced in [9]. In particular, it was shown that a FP solution exists *if and only if*

$$\mu^*[n] \leq \frac{\sum_{i=1}^{N_t} \mu^*[i]}{L} \quad (17)$$

for all $n \in [N_t]$. In this case, we see that $\sum_{i=1}^{N_t} \mu^*[i] = L$ and $\mu^*[n] \leq 1$ for all $n \in [N_t]$. Therefore, an optimal computation assignment exists. Moreover, we provide Algorithm 1 to define an optimal computation assignment, $(\mathcal{M}_t, \mathcal{P}_t)$.

Algorithm 1 Computation Assignment: Heterogeneous CEC

Input: μ^* , N_t , L , and q

- 1: $\mathbf{m} \leftarrow \mu^*$
- 2: $f \leftarrow 0$
- 3: **while** \mathbf{m} contains a non-zero element **do**
- 4: $f \leftarrow f + 1$
- 5: $L' \leftarrow \sum_{n=1}^{N_t} m[n]$
- 6: $N' \leftarrow$ number of non-zero elements in \mathbf{m}
- 7: $\ell \leftarrow$ indices that sort the non-zero elements of \mathbf{m} from smallest to largest³
- 8: $\mathcal{P}_f \leftarrow \{\ell[1], \ell[N' - L + 2], \dots, \ell[N']\}$
- 9: **if** $N' \geq L + 1$ **then**
- 10: $\alpha_f \leftarrow \min\left(\frac{L'}{L} - m[\ell[N' - L + 1]], m[\ell[1]]\right)$
- 11: **else**
- 12: $\alpha_f \leftarrow m[\ell[1]]$
- 13: **end if**
- 14: **for** $n \in \mathcal{P}_f$ **do**
- 15: $m[n] \leftarrow m[n] - \alpha_f$
- 16: **end for**
- 17: **end while**
- 18: $F \leftarrow f$
- 19: Partition rows $[\frac{q}{L}]$ into F disjoint row sets: $\mathcal{M}_1, \dots, \mathcal{M}_F$ of size $\frac{\alpha_1 q}{L}, \dots, \frac{\alpha_F q}{L}$ rows respectively

Remark 3: In [9], Algorithm 1 was shown to take at most N_t iterations to complete. Therefore, $F \leq N_t$ and at most there are N_t computations assignments.

A. Example Using Algorithm 1

We return to the example of Section III and use Algorithm 1 to derive the computation (row) assignments for $t = 2$. The steps of the algorithm are shown in Fig. 2. In the first iteration, $f = 1$, $\mathbf{m} = \mu$ as no computations have been assigned yet. Rows of the respective coded matrices are assigned to machine 1, which is a machine with the least remaining computations to be assigned, and machines 5 and 6 with the most remaining computations to be assigned. Moreover,

$$m[1] = \frac{2}{5} \leq \frac{L'}{L} - m[3] = 1 - \frac{3}{5} = \frac{2}{5} \quad (18)$$

³ ℓ is an N' -length vector and $0 < m[\ell[1]] \leq m[\ell[2]] \leq \dots \leq m[\ell[N']]$.

f	α_f	$m[1]$	$m[2]$	$m[3]$	$m[4]$	$m[5]$	$m[6]$	L'
1	$\frac{2}{5}$	$\frac{2}{5}$	$\frac{2}{5}$	$\frac{3}{5}$	0	$\frac{4}{5}$	$\frac{4}{5}$	3
2	$\frac{1}{5}$	0	$\frac{2}{5}$	$\frac{3}{5}$	\vdots	$\frac{2}{5}$	$\frac{2}{5}$	$\frac{9}{5}$
3	$\frac{1}{5}$	\vdots	$\frac{1}{5}$	$\frac{2}{5}$	\vdots	$\frac{2}{5}$	$\frac{1}{5}$	$\frac{6}{5}$
4	$\frac{1}{5}$	\vdots	0	$\frac{1}{5}$	\vdots	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{3}{5}$
-		\downarrow	\downarrow	0	\downarrow	0	0	0

Fig. 2. Task assignment by Algorithm 1 for example of III ($t = 2$).

where machine 3 is the machine with the most remaining rows to be assigned that is not included in $\mathcal{P}_1 = \{1, 5, 6\}$. Therefore, a fraction $\alpha_1 = \frac{2}{5}$ of the rows are assigned to machines 1, 5, 6. Then, \mathbf{m} is adjusted to reflect the remaining computations to be assigned and $L' = 3 - 3\alpha_1 = \frac{9}{5}$.

In the second iteration, $f = 2$, machine 2 is a machine with the least remaining rows to be assigned. Computations are assigned to machine 2 and machines 3 and 6 which are a pair of machines with the most remaining computations to be assigned. Ideally, we would like to assign all the remaining rows to machine 2. However,

$$m[2] = \frac{2}{5} > \frac{L'}{L} - m[5] = \frac{3}{5} - \frac{2}{5} = \frac{1}{5} \quad (19)$$

and assigning the remaining rows to machine 2 in this iteration will prevent a valid solution going forward. Therefore, $\alpha_2 = \frac{1}{5}$ and after this iteration \mathbf{m} and L' are adjusted accordingly.

In the third iteration, $f = 3$,

$$m[2] = \frac{1}{5} \leq \frac{L'}{L} - m[6] = \frac{2}{5} - \frac{1}{5} = \frac{1}{5} \quad (20)$$

and an $\alpha_3 = \frac{1}{5}$ of the rows are assigned to machines 2, 3, 5. \mathbf{m} and L' are adjusted accordingly. Finally, in the fourth iteration, $f = 4$, the three machines with remaining assignments, machines 3, 5, 6 are assigned an $\alpha_4 = \frac{1}{5}$ of the rows. After the 4 iterations, $m[n] = 0$ for all $n \in 1, 2, 3, 4$ and the computation assignment is complete.

VI. CONCLUSION

In this work, we study CEC where machines store MDS coded data and have varying computation speed. Given a set of available machines with arbitrary relative computation speeds, we derive an optimal computation load among the machines. Then, we show the existence of a computation assignment which yields the optimal computation load. The assignment makes use of the MDS code design by assigning computations to L machines. Moreover, we present a low complexity algorithm to define the computation assignments with at most a number of iterations equal to the number of available machines. Our CEC design has the potential to perform computations faster than the state-of-the-art design which was developed for a homogeneous computing network.

REFERENCES

- [1] S. Li, M. A. Maddah-Ali, Q. Yu, and A. S. Avestimehr, "A fundamental tradeoff between computation and communication in distributed computing," *IEEE Transactions on Information Theory*, vol. 64, no. 1, pp. 109–128, 2018.
- [2] M. Adel Attia and R. Tandon, "Near optimal coded data shuffling for distributed learning," *IEEE Transactions on Information Theory*, vol. 65, no. 11, pp. 7325–7349, Nov 2019.
- [3] A. Elmahdy and S. Mohajer, "On the fundamental limits of coded data shuffling," in *2018 IEEE International Symposium on Information Theory (ISIT)*, June 2018, pp. 716–720.
- [4] K. Wan, D. Tuninetti, M. Ji, G. Caire, and P. Piantanida, "Fundamental limits of decentralized data shuffling," *IEEE Transactions on Information Theory*, 2020.
- [5] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," *IEEE Transactions on Information Theory*, vol. PP, no. 99, pp. 1–1, 2017.
- [6] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "Straggler mitigation in distributed matrix multiplication: Fundamental limits and optimal coding," *IEEE Transactions on Information Theory*, vol. 66, no. 3, pp. 1920–1933, 2020.
- [7] Y. Yang, M. Interlandi, P. Grover, S. Kar, S. Amizadeh, and M. Weimer, "Coded elastic computing," in *2019 IEEE International Symposium on Information Theory (ISIT)*, July 2019, pp. 2654–2658.
- [8] H. Dau, R. Gabrys, Y. Huang, C. Feng, Q. Luu, E. Alzahrani, and Z. Tari, "Optimizing the transition waste in coded elastic computing," *arXiv preprint arXiv:1910.00796*, 2019.
- [9] N. Woolsey, R. Chen, and M. Ji, "An optimal iterative placement algorithm for pir from heterogeneous storage-constrained databases," in *GLOBECOM 2019 IEEE Global Communications Conference*. IEEE, 2019.
- [10] Y. Yang, M. Interlandi, P. Grover, S. Kar, S. Amizadeh, and M. Weimer, "Coded elastic computing," *arXiv preprint arXiv:1812.06411*, 2018.
- [11] N. Woolsey, R. Chen, and M. Ji, "Heterogeneous computation assignments in coded elastic computing," *arXiv preprint arXiv:2001.04005*, 2020.