

A Practical Algorithm Design and Evaluation for Heterogeneous Elastic Computing with Stragglers

Nicholas Woolsey¹, Jörg Kliewer², Rong-Rong Chen¹ and Mingyue Ji¹

¹University of Utah, ²New Jersey Institute of Technology

Email: {nicholas.woolsey@utah.edu, jkliewer@njit.edu, rchen@ece.utah.edu, mingyue.ji@utah.edu}

Abstract—Our extensive real measurements over Amazon EC2 show that the virtual instances often have different computing speeds even if they share the same configurations. This motivates us to study heterogeneous Coded Storage Elastic Computing (CSEC) systems where machines, with different computing speeds, join and leave the network arbitrarily over different computing steps. In CSEC systems, a Maximum Distance Separable (MDS) code is used for coded storage such that the file placement does not have to be re-defined with each elastic event. Computation assignment algorithms are used to minimize the computation time given computation speeds of different machines. While previous studies of heterogeneous CSEC do not include stragglers – the slow machines during the computation, we develop a new framework in heterogeneous CSEC that introduces straggler tolerance. Based on this framework, we design a novel algorithm using our previously proposed approach for heterogeneous CSEC such that the system can handle any subset of stragglers of a specified size while minimizing the computation time. Furthermore, we establish a trade-off in computation time and straggler tolerance. Another major limitation of existing CSEC designs is the lack of practical evaluations using real applications. In this paper, we evaluate the performance of our designs on Amazon EC2 for applications of the power iteration and linear regression. Evaluation results show that the proposed heterogeneous CSEC algorithms outperform the state-of-the-art designs by more than 30%.

I. INTRODUCTION

Coded Storage Elastic Computing (CSEC) was introduced by Yang et al. [1] as an effective method to perform distributed computing on elastic cloud systems where machines have limited storage capacity. Here, elasticity means that machines can join and leave the network periodically as machines may be reserved or freed from other higher priority jobs. The main advantage of CSEC is that machines store coded data such that a computation on the data can be recovered as long as enough machines are available. This is in contrast to storing uncoded data where if some machine leaves with a particular subset of data, then a computation on the entire data set may not be recovered; or, in order to tolerate the same number of preempted machines, the required storage/computation overhead is larger. The use of coded storage alleviates the need to re-distribute the storage with each elastic event which could greatly slow down the cloud computing system.

While coded computing has been studied in the literature for straggler mitigation [2]–[14], CSEC is different from these works because, at the beginning of a computation step, we know which machines are available *a priori*. This allows us

to design an effective computation assignment to the machines. Assume we use a Maximum Distance Separable (MDS) code to encode the data matrix with a recovery threshold of L and each machine stores one of coded sub-matrices (cs-matrices). If there are exactly L available machines for the given time step, then each machine will perform computations on all locally available data and the master machine will collect all computation results to recover the original computation task. However, if there are more than L machines, we can strategically assign computations such that each machine performs computations on a fraction of its local coded data, effectively reducing the overall computation time. The overall computation is split into many steps, where in each step we know which machines are available, even though their availability may change over computation steps. The original CSEC design [1] studied homogeneous CSEC systems where machines have the same computation speed. It proposed a cyclic homogeneous computation assignment design to minimize the computation load at each machine such that no computation redundancy is present in the network.

Based on our measurements of the computation speeds on Amazon EC2 instances, we found that these virtual instances often have different computation speeds even if they share the same configuration. Hence, in [15], for heterogeneous elastic computing systems where machines have different computation speeds, we reformulated the CSEC problem as a combinatorial optimization problem with the objective of minimizing the overall computation time, instead of per machine computation load as in [1]. This combinatorial optimization problem was solved using a novel approach, which decomposed this problem into two sub-optimization problems consisting of a convex optimization problem and a combinatorial “filling” problem, where an optimal solution was found. Then, in [16], Dau et al. introduced a new concept called Transition Waste (TW) that measures the difference between the total number of changes and the necessary changes of the computation assignment. The authors proposed novel algorithms to find the minimum TW for some parameter settings in homogeneous CSEC systems.

However, a question often arises for CSEC designs. *What happens if an elastic event occurs during a computation step?* Computation assignments are designed such that the set of available machines is known during a computation step. If a machine does not respond during a computation step, it is

deemed as a failed node and we have to re-assign computations to recover the intended computation. In addition, a machine may not respond for reasons other than an elastic event such as becoming unpredictably slow, unresponsive or otherwise simply fails. In general, we will label any one of these events during a computation step as a *straggler*. Adding straggler tolerance in CSEC is an interesting problem that has not been studied in the literature. Given the MDS coded storage placement used in [1], Kiani et al. in [17] proposed hierarchical computation assignment algorithms for homogeneous CSEC systems. However, the performance evaluation in [17] were done using computer simulations, where stragglers are modeled to take place with certain probabilities, instead of using practical cloud computing systems such as Amazon EC2 or Microsoft Azure. In addition, the computation task is only matrix-matrix multiplication rather than real applications.

In this paper, our focus is to 1) incorporate straggler toleration into the heterogeneous CSEC framework and 2) conduct actual performance evaluations over Amazon EC2 without any assumptions or models of stragglers. In particular, using the existing MDS coded storage design, we will show that straggler tolerance can be incorporated into our proposed heterogeneous CSEC optimization framework [15]. The main advantage of our proposed approach is to provide straggler tolerance by adjusting computation assignments. Furthermore, we show that there is an interesting trade-off between the computation time of heterogeneous CSEC system and the number of stragglers that can be tolerated. One extreme case is the heterogeneous CSEC with no straggler tolerance and the other extreme case is to maximize straggler tolerance using the coded computing design of [2]. We will present a general framework that includes these designs and show there is a straggler tolerance and computation time trade-off.

Our contributions are summarized as follows:

- 1) We develop a new framework in heterogeneous CSEC that introduces straggler tolerance. Based on this framework, we design a novel algorithm using our previously proposed approach for heterogeneous CSEC such that the system can handle any subset of stragglers of a specified size and the computation time is minimized.
- 2) The proposed new heterogeneous CSEC framework introduces a trade-off between computation time and straggler tolerance of CSEC systems. An analytical expression is derived to characterize the optimized computation time as a function of the computation speeds and the straggler tolerance.
- 3) We validate our proposed CSEC designs by actual implementations on Amazon EC2 instances for applications of both the power iteration algorithm and linear regression. In particular, the computation assignment for the proposed CSEC is based on real-time measurements of the computation speeds of Amazon EC2 instances.

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

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

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

of size $\frac{q}{L} \times r$.

The machines collectively perform matrix-vector computations over multiple computation steps. In a given step only a subset of the N machines are available to perform matrix computations. More specifically, in computation 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.

The machines of \mathcal{N}_t do not compute \mathbf{y}_t directly. Instead, each machine $n \in \mathcal{N}_t$ computes the set $\mathcal{V}_n = \{v = \tilde{\mathbf{X}}_n^{(i)} \mathbf{w}_t : i \in \mathcal{W}_n\}$, 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 .

Definition 1: (Computation load) The computation load vector, $\boldsymbol{\mu}$, is defined as

$$\mu[n] = \frac{|\mathcal{W}_n|}{\left(\frac{q}{L}\right)}, \quad \forall n \in \mathcal{N}_t, \quad (3)$$

which is the fraction of rows of the corresponding stored cs-matrix computed by machine n in time step t . \diamond

Note that, $\boldsymbol{\mu}$, \mathcal{V}_n and \mathcal{W}_n change with each computation 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 known for each computation step and defined as follows.

Definition 2: (Computation Speed) The computation speed vector \mathbf{s} is a length- N vector with elements $s[n]$, $n \in [N]$, where $s[n]$ is the speed of machine n measured as the inverse of the time it takes machine n to compute all rows of its cs-matrix. The normalized computation speed is the computation speed divided by the average computation speed among all available machines. \diamond

The computation time is dictated by the machine that takes the most time to perform its assigned computations, defined as follows.

Definition 3: (Computation Time) The computation time in a particular computation step is defined as

$$c(\boldsymbol{\mu}) = \max_{n \in \mathcal{N}_t} \frac{\mu[n]}{s[n]}. \quad (4)$$

\diamond

In a given computation step, for each $i \in [\frac{q}{L}]$, at least L machines perform the vector-vector multiplication with the i -th row of their local cs-matrix and w_t . The results are sent to a master machine which can resolve the elements of y_t by the MDS code design.

A computation assignment is defined by 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 machines, $\mathcal{P}_t = (\mathcal{P}_1, \dots, \mathcal{P}_F)$, are defined such 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 . In other words, the rows computed by machine $n \in \mathcal{N}_t$ in time step t are in the set $\mathcal{W}_n = \bigcup \{\mathcal{M}_f : f \in [F], n \in \mathcal{P}_f\}$ and therefore, μ is a function of $(\mathcal{M}_t, \mathcal{N}_t)$. The sets $\mathcal{M}_1, \dots, \mathcal{M}_F$ and $\mathcal{P}_1, \dots, \mathcal{P}_F$ and F may vary with each computation step based on machines' availability.

In a given computation step t , our goal is to design the task assignments, \mathcal{M}_t and \mathcal{N}_t , such that the computation $y_t = X w_t$ can be recovered even when some machines are stragglers that do not provide their assigned computations to the master machine. We define S as the straggler tolerance of the network, such that the computation can be recovered even when S of the available machines become stragglers. Furthermore, let $\mathcal{S}_t \subset \mathcal{N}_t$ be the subset of the available machines that become stragglers. When designing the computation assignment, \mathcal{S}_t is not known and in general can be any subset of S available machines.

Then, we aim to design the computation assignment that minimizes the computation time of (4) resulting from the computation load vector defined in (3). In computation step t , given \mathcal{N}_t and s , the optimal computation time, c^* , is the minimum of computation times defined by all possible task assignments, $(\mathcal{M}_t, \mathcal{N}_t)$, such that S stragglers can be tolerated and the computation can be recovered. In particular c^* is the optimal value of the following combinatorial optimization problem.

$$\min_{(\mathcal{M}_t, \mathcal{N}_t)} c(\mu(\mathcal{M}_t, \mathcal{N}_t)) \quad (5a)$$

$$\text{s.t.} \quad \bigcup_{\mathcal{M}_f \in \mathcal{M}_t} \mathcal{M}_f = [\frac{q}{L}], \quad (5b)$$

$$|\mathcal{P}_f \setminus \mathcal{S}_t| \geq L \quad \forall \mathcal{P}_f \in \mathcal{P}_t, \quad \forall \mathcal{S}_t \subset \mathcal{N}_t, \quad |\mathcal{S}_t| = S \quad (5c)$$

$$|\mathcal{M}_t| = |\mathcal{P}_t|. \quad (5d)$$

Remark 1: The main difference between this problem formulation and that of [15] is the consideration of the straggler machine set \mathcal{S}_t . The computation assignment is designed such that the computation on each row set is guaranteed to be computed at least L times across the machines (excluding stragglers). In general, all possibilities of \mathcal{S}_t must be considered which could be any subset of S available machines.

III. EXAMPLES

We first compare the CSEC designs [1], [15] to the straggler tolerant design of [2]. Then, we present examples of the proposed design which bridge the gap between the two design methodologies. Homogeneous systems are discussed first, followed by heterogeneous systems where machines have varying computation speeds.

A. Homogeneous Examples

Consider a set of $N_t = 5$ available machines in a computation step. Each machine n has the same computation speed of $s[n] = 1$. The storage at these machines adopts a $(5, 3)$ -MDS code where the machines store the coded matrices of $\tilde{X}_1 = X_1$, $\tilde{X}_2 = X_2$, $\tilde{X}_3 = X_3$, $\tilde{X}_4 = X_1 + X_2 + X_3$, and $\tilde{X}_5 = X_1 + 2X_2 + 4X_3$. Here $L = 3$, is the recovery threshold of the MDS code. With these 5 available machines, we can assign computations such that each row is computed by at least L machines.

Homogeneous Cyclic CSEC Design [1], ($S = 0$): Using the cyclic homogeneous design, the rows of the coded matrices are split into $F = 5$ disjoint subsets. The first row set is computed at machines 1 through 3, the next row set is computed at machines 2 through 4 and so on. Following the notation of the problem formulation, the row sets are $\mathcal{M}_i = \{1 + (i-1)\frac{q}{5L}, \dots, i\frac{q}{5L}\}$ for $i \in [5]$ and the machine sets are $\mathcal{P}_1 = \{1, 2, 3\}$, $\mathcal{P}_2 = \{2, 3, 4\}$, $\mathcal{P}_3 = \{3, 4, 5\}$, $\mathcal{P}_4 = \{4, 5, 1\}$, and $\mathcal{P}_5 = \{5, 1, 2\}$. The main advantage of this design is that each machine n computes $\mu[n] = \frac{3}{5}$ of its local cs-matrix. Since each machine has a speed of 1, we see that the computation time is $c = \frac{5}{3}$. However, this design is susceptible to stragglers. Since each row set is assigned to exactly L machines, if any one machine becomes a straggler, then the matrix-vector multiplication cannot be recovered.

Straggler Design of [2], ($S = 2$): Alternatively, each machine computes the entire local cs-matrix. In other words, the assignment is $F = 1$, $\mathcal{M}_1 = \{1, \dots, \frac{q}{L}\}$, $\mathcal{P}_1 = \{1, 2, 3, 4, 5\}$. In this case, it is clear that $S = 2$ stragglers can be tolerated, since we only need computations recovered from any $L = 3$ machines. However, this is at a cost of computation time which increases to $c = 1$, or the time for a machine to compute all rows of its local cs-matrix.

A Homogeneous CSEC Straggler Design, ($S = 1$): The question remains if the computation time can be reduced when the straggler tolerance is only $S = 1$ straggler. It turns out there is a computation-straggler tolerance trade-off if we use another computation assignment. For example, using the principles of the cyclic assignment, we can assign computations to $L + S = 4$ machines instead of just L machines. That is, we define $F = 5$ row sets, $\mathcal{M}_i = \{1 + (i-1)\frac{q}{5L}, \dots, i\frac{q}{5L}\}$ for $i \in [5]$, and assign these to the machine sets of $\mathcal{P}_1 = \{1, 2, 3, 4\}$, $\mathcal{P}_2 = \{2, 3, 4, 5\}$, $\mathcal{P}_3 = \{3, 4, 5, 1\}$, $\mathcal{P}_4 = \{4, 5, 1, 2\}$, and $\mathcal{P}_5 = \{5, 1, 2, 3\}$, respectively. Note that each row set is assigned to 4 machines, such that if any one machine becomes a straggler, the computation on that row set can be recovered. Also, note that the computation time is now $c = \frac{4}{5}$ since each machine performs computations on 4 row sets each containing $\frac{1}{5}$ of the rows.

B. Heterogeneous Examples

Assume that there are $N_t = 5$ available machines which have the same storage as in the homogeneous examples. Again, $L = 3$ and each row computation required computations from $L = 3$ machines to be recovered. However, this time, the machines have heterogeneous computation speeds defined by

$s = [1, 1, 2, 2, 3]$ such that machine 5 is the fastest with a computation speed of 3 and machines 1 and 2 are the slowest with a computation speed of 1.

Heterogeneous Algorithm CSEC Design [15], ($S = 0$): We define a computation assignment that minimizes the computation time. It is ideal to assign each machine an amount of computations relative to the speed of the machine. In this case, we define the computation load of the machines to be $\boldsymbol{\mu} = [\frac{1}{3}, \frac{1}{3}, \frac{2}{3}, \frac{2}{3}, 1]$, where, for example, machines 1 and 3 will compute a $\frac{1}{3}$ and $\frac{2}{3}$ fraction, respectively, of their local cs-matrix. Moreover, the computation load vector sums to $L = 3$ such that it allows each row to be computed 3 times. A computation assignment that yields this computation load is $F = 3$ row sets each with $\frac{1}{3}$ of the rows assigned to the machines of $\mathcal{P}_1 = \{1, 4, 5\}$, $\mathcal{P}_2 = \{2, 3, 5\}$ and $\mathcal{P}_3 = \{3, 4, 5\}$. Each machine takes the same time to perform their assigned computations and the computation time of the network is $c = \frac{1}{3}$. However, there is no straggler tolerance.

Straggler Design of [2], ($S = 2$): Each machine computes the entirety of its local cs-matrix. In this case, $S = 2$ stragglers are tolerated. However, the computation time is dictated by the slowest machines which are machines 1 and 2. Since we plan for the possibility of any machine straggling, we may need to wait for machines 1 and 2 to finish their computations. This means the computation time is $c = 1$.

Proposed Heterogeneous CSEC Straggler Design, ($S = 1$): The computation time can be improved if we reduce the straggler tolerance to $S = 1$. We assign computations to $L + S = 4$ machines yielding a computation load vector of $\boldsymbol{\mu} = [\frac{1}{2}, \frac{1}{2}, 1, 1, 1]$, where machines 3 through 5 have a computation load of 1 and compute the entirety of their local matrix. Adapting the assignment algorithm of [15], the following computation assignment yields this computation load. There are $F = 2$ row sets each containing half of the rows assigned to the machines of $\mathcal{P}_1 = \{1, 3, 4, 5\}$ and $\mathcal{P}_2 = \{2, 3, 4, 5\}$, respectively. Notice that each row set is assigned to 4 machines, such that if any one machine straggles, the result can be recovered. Furthermore, the computation time is defined by the local computation time of machines 1 through 4 which each take the same time to finish. The computation time is $c = \frac{1}{2}$.

IV. PROPOSED HETEROGENEOUS CSEC DESIGN

Our proposed straggler tolerant CSEC design, given by Algorithm 1, is obtained by solving the combinatorial optimization (5) in a similar fashion as in [15] (line 6 in Algorithm 1), and measuring (line 14 in Algorithm 1) and updating (line 4 in Algorithm 1) the speed vector at each iteration. This algorithm adapts previous CSEC designs [15] to assign computations to $L + S$ machines instead of just L machines. We will explain the proposed heterogeneous CSEC design in the following. For completeness, we start with the homogeneous CSEC design first.

Homogeneous CSEC Straggler Tolerant Design: Given N_t available machines, $\mathcal{N}_t = \{1, 2, \dots, N_t\}$, define a computation assignment with $F = N_t$ row sets. There are N_t

disjoint equally-sized row sets that collectively span all rows: $\mathcal{M}_i = \{1 + (i - 1)\frac{a}{LN_t}, \dots, i\frac{a}{LN_t}\}$ for $i \in [N_t]$. Then, define a cyclic assignment such that machine set $\mathcal{P}_i = \{i\%N_t, \dots, (i + L + S - 1)\%N_t\}$ for $i \in [N_t]$, where we define $a\%N_t \triangleq a - \lfloor \frac{a-1}{N_t} \rfloor N_t$ to facilitate the cyclic design.

Proposed Heterogeneous CSEC Straggler Tolerant Design: Assume there are N_t available machines, $\mathcal{N}_t = \{1, 2, \dots, N_t\}$ with computation speeds in ascending order $s[1] \leq s[2] \leq \dots \leq s[N_t]$. Define the computation load vector to be

$$\boldsymbol{\mu}^*[n] = \begin{cases} c^*s[n] & \text{if } 1 \leq n \leq k^* \\ 1 & \text{if } k^* + 1 \leq n \leq N_t, \end{cases} \quad (6)$$

where k^* is the largest integer in $[N_t]$ such that

$$\frac{1}{s[k^* + 1]} < c^* = \frac{L + S - N_t + k^*}{\sum_{n=1}^{k^*} s[n]} \leq \frac{1}{s[k^*]}. \quad (7)$$

Given the computation load vector $\boldsymbol{\mu}^*$, we can obtain the computation assignment by applying the assignment algorithm of [15] to assign computations to $L + S$ machines instead of L machines (line 6 in Algorithm 1). This naturally includes straggler tolerance and yields the following theorem.

Theorem 1: The computation times of the straggler tolerant CSEC designs are:

Homogeneous design:

$$c^* = \frac{L + S}{N_t} \cdot \frac{1}{\min_n s[n]} \quad (8)$$

Heterogeneous design:

$$c^* = \frac{L + S - N_t + k^*}{\sum_{n=1}^{k^*} s[n]} \quad (9)$$

where k^* is defined in the heterogeneous design. \square

Remark 2: For both designs, we observe that the computation time c increases with the straggler tolerance, S . This demonstrates a trade-off between the computation time and straggler tolerance of the system.

V. EVALUATION ON AMAZON EC2

A. Evaluation Setup

We evaluate the proposed algorithm using two representative real applications including the power iteration and linear regression algorithms on Amazon EC2 instances as follows.

Power Iteration: The power iteration algorithm computes the largest eigenvalue and the corresponding eigenvector of a large matrix \mathbf{X} . In particular, it starts with a vector \mathbf{b}_0 , which may be an approximation to the dominant eigenvector or a random vector. The method is described by the recursive relation, $\mathbf{b}_{k+1} = \frac{\mathbf{X}\mathbf{b}_k}{\|\mathbf{X}\mathbf{b}_k\|}$. The sequence \mathbf{b}_k converges to an eigenvector associated with the dominant eigenvalue. It can be seen that at each iteration, we can directly apply the proposed Algorithm 1. In particular, a dense 60,000-by-60,000 symmetric matrix is row-wise split into $L = 10$ sub-matrices which are used to define the coded matrices stored at each machine. A vector of length 60,000 is updated by performing a matrix-vector multiplication in a distributed

Algorithm 1 Adaptive Straggler Tolerant Coded Storage Elastic Computing

Input: $\hat{s}, \gamma, S, T, w_1$

- 1: $\nu \leftarrow \hat{s}$
- 2: **for** $t \in [T]$ **do**
- 3: **At Master Machine:**
- 4: $\hat{s} \leftarrow \gamma\nu + (1 - \gamma)\hat{s}$ (update estimate of speed vector).
- 5: $\mathcal{N}_t \leftarrow$ list of available machines
- 6: $(F, \mathcal{M}_1, \dots, \mathcal{M}_F, \mathcal{P}_1, \dots, \mathcal{P}_F) \leftarrow$ Results of computation assignment algorithm with straggler tolerance of S for available machines \mathcal{N}_t with speeds of \hat{s}
- 7: Send w_t and $(F, \mathcal{M}_1, \dots, \mathcal{M}_F, \mathcal{P}_1, \dots, \mathcal{P}_F)$ to Worker Machines
- 8: **At Worker Machines:**
- 9: $n \leftarrow$ index of worker machine
- 10: $\mu \leftarrow$ computation load of worker machine n
- 11: $\tau_1 \leftarrow$ current time
- 12: Perform assigned computations based on $(F, \mathcal{M}_1, \dots, \mathcal{M}_F, \mathcal{P}_1, \dots, \mathcal{P}_F)$
- 13: $\tau_2 \leftarrow$ current time
- 14: $\nu[n] \leftarrow \mu[n]/(\tau_2 - \tau_1)$ (calculate speed based on current computation step)
- 15: Send computations and $\nu[n]$ to Master Machine
- 16: **At Master Machine:** after receiving results from L workers.
- 17: $w_{t+1} \leftarrow$ Decode and combine worker results
- 18: **end for**

Output: w_T

manner on the available worker machines. The master machine combines the results and normalizes the vector. This process is repeated such that the vector converges to the eigenvector associated with the largest eigenvalue.

Linear Regression: A linear regression problem is to solve $\min_{\mathbf{b}} \|\mathbf{X}\mathbf{b} - \mathbf{y}\|^2$. One typical approach to solve such algorithm is to use gradient descent $\mathbf{b}_{k+1} = \mathbf{b}_k - \eta \mathbf{X}^T (\mathbf{X}\mathbf{b}_k - \mathbf{y})$. Letting $\mathbf{X}\mathbf{b}_k - \mathbf{y} = \mathbf{z}_k$, we see that in each iteration, the algorithm consist of 2 matrix-vector multiplications, which are $\mathbf{X}\mathbf{b}_k$ and $\mathbf{X}^T \mathbf{z}_k$. We apply Algorithm 1 to both of them. In particular, we use a 200,000-by-5,000 matrix representing a dataset with 200,000 entries each with 5,000 features. We follow the distributed gradient descent design of [1], [2]. Each worker machine stores two coded data matrices, one coded row-wise and the other coded column-wise and $L = 10$. These are used to perform matrix multiplication with the data matrix and its transpose in a distributed fashion to update the weights.

The network has one `t2.x2large` master machine with 8 vCPUs and 32 GiB of memory. The worker machines consist of 10 `t2.large` instances, each with 2 vCPUs and 8 GiB of memory, and 10 `t2.xlarge` instances, each with 4 vCPUs and 16 GiB of memory. Each machine is either a *stable machine* or an *elastic machine* such that with each iteration, the 12 stable machines (6 `t2.large` and 6 `t2.xlarge`) are always available for computation assignment and the

elastic machines (4 `t2.large` and 4 `t2.xlarge`) are each available with a probability of 0.5. The availability of the elastic machines are independent across machines and also each computation step. Note that, any available machine that is assigned computations in a computation step could become a straggler (slowest machine). In this case, the master machine will not wait for the 2 slowest machines (stragglers) to finish their computations (“with straggler” scenario in Figs. 1 and 2). We define the *uncoded* and *no straggler* scenarios as follows.

Uncoded: No straggler tolerance is considered. The stable machines include 5 `t2.large` and 5 `t2.xlarge` machines which store uncoded matrices of the dataset without redundancy. We do not consider elastic machines in this case.

No Stragglers: No straggler tolerance is considered. The stable machines consist of 5 `t2.large` and 5 `t2.xlarge` machines and the elastic machines consist of 4 `t2.large` and 4 `t2.xlarge` machines for a total of 18 worker machines.

B. Heterogeneous Computation Speed

In Table I, we give an example of the measured relative computation speeds of Amazon EC2 instances during a trial of the power iteration and gradient descent algorithm for linear regression, resp., using CSEC. The slowest speed is shown in blue and the highest speed is shown in red. One interesting observation is that the speeds can vary widely even for the same instance type with the same configurations. For example, in the linear regression application, the normalized computation speed of the `t2.large` instances varies from 0.48 to 0.90 and the `t2.xlarge` instances vary from 0.58 to 1.37. Furthermore, while not shown here, we observed the computation speeds remained nearly constant on the order of minutes and the speeds are predictable between computation steps which are on the order of seconds.

TABLE I
NORMALIZED COMPUTATION SPEED

Power Iteration				Linear Regression			
t2.large		t2.xlarge		t2.large		t2.xlarge	
0.85	0.84	1.28	1.20	0.70	0.53	1.30	1.26
0.76	0.88	0.92	1.27	0.90	0.90	1.26	1.34
0.59	0.64	1.20	0.91	0.88	0.77	0.88	1.26
0.59	0.88	1.34	0.85	0.90	0.48	1.23	1.37
0.90	0.70	1.29	0.90	0.52	0.48	0.83	0.58

C. Evaluation Results

The results of using the CSEC designs on Amazon EC2 are shown in Figs. 1 and 2 for power iteration and linear regression, respectively. We observe similar trends for both applications. When no straggler tolerance is considered, the results are shown as solid lines in Figs. 1 and 2. We can see that the proposed heterogeneous CSEC algorithm converges faster than other cases. In particular, for power iteration, nearly 40% gain can be obtained compared to the uncoded design and about 30% gain can be obtained compared to the homogeneous design. This occurs because each computation step, which updates the vector or weights, is performed faster

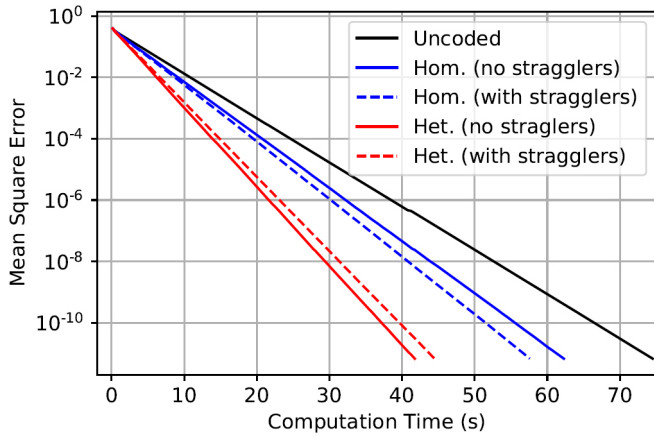


Fig. 1. **Power Iteration:** Results using CSEC designs on Amazon EC2 without stragglers and with 2 stragglers each iteration. The y-axis represents the normalized mean square error between the true dominant eigenvector and the estimated eigenvector.

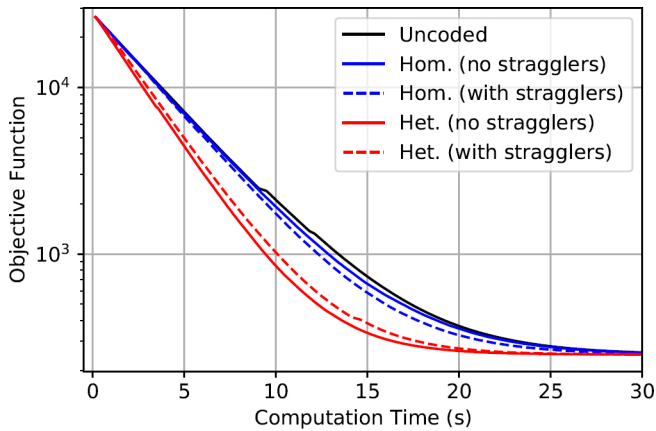


Fig. 2. **Linear Regression:** Results using CSEC designs on Amazon EC2 without stragglers and with 2 stragglers each iteration.

when compared to the other designs. Both the uncoded design and homogeneous cyclic design assign the same computation load to each machine. These designs are hindered by slower machines with lower computation speed, even though they may share the same configurations. Faster machines will complete and then wait for slower machines to finish. However, for the proposed heterogeneous design, computation load is assigned relative to computation speed so that ideally each machine will finish at approximately the same time regardless of its speed.

There are several key observations from the experiments with straggler tolerance whose results are shown as dashed lines in Figs. 1 and 2. First, the heterogeneous CSEC design converges faster than the homogeneous CSEC design for similar reasons as those of the experiments without straggler tolerance. Second, the straggler tolerance algorithms improve the performance of homogeneous designs. This is because the slowest machines are slow enough so that ignoring them while assigning more tasks to other machines can help. However,

interestingly, for the proposed heterogeneous design, since the slowest machines are optimally assigned less tasks and the speed of machines changes slowly in our experiments, the algorithms without straggler tolerance can be better.

VI. CONCLUSION

We introduce straggler tolerance in CSEC systems and bridge the gap between CSEC and straggler tolerant distributed computing. We demonstrate a trade-off between computation time and straggler tolerance under our new general framework. We validate our straggler tolerant CSEC designs on Amazon EC2 for the applications of the power iteration and gradient descent for linear regression. We find the proposed heterogeneous CSEC designs to be particularly effective and robust in heterogeneous environments where machine speeds can vary widely even though they share the same configurations.

REFERENCES

- [1] 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.
- [2] 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.
- [3] S. Dutta, V. Cadambe, and P. Grover, "Short-dot: Computing large linear transforms distributedly using coded short dot products," in *Advances In Neural Information Processing Systems*, 2016, pp. 2100–2108.
- [4] N. S. Ferdinand and S. C. Draper, "Anytime coding for distributed computation," in *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Sept 2016, pp. 954–960.
- [5] 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.
- [6] Malihe Aliasgari, Jörg Kliewer, and Osvaldo Simeone, "Coded computation against processing delays for virtualized cloud-based channel decoding," *IEEE Transactions on Communications*, vol. 67, no. 1, pp. 28–38, 2019.
- [7] C. Karakus, Y. Sun, S. Diggavi, and W. Yin, "Straggler mitigation in distributed optimization through data encoding," in *Advances in Neural Information Processing Systems*, 2017, pp. 5434–5442.
- [8] W. Halbawi, N. Azizan, F. Salehi, and B. Hassibi, "Improving distributed gradient descent using reed-solomon codes," in *2018 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2018, pp. 2027–2031.
- [9] G. Suh, K. Lee, and C. Suh, "Matrix sparsification for coded matrix multiplication," in *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Oct 2017, pp. 1271–1278.
- [10] R. K. Maity, A. S. Rawat, and A. Mazumdar, "Robust gradient descent via moment encoding with ldpc codes," in *SysML Conference*, 2018.
- [11] M. F. Aktas, P. Peng, and E. Soljanin, "Straggler mitigation by delayed relaunch of tasks," *SIGMETRICS Perform. Eval. Rev.*, vol. 45, no. 3, pp. 224–231, Mar. 2018.
- [12] S. Wang, J. Liu, N. Shroff, and P. Yang, "Fundamental limits of coded linear transform," *arXiv:1804.09791*, 2018.
- [13] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," *arXiv:1802.03475*, 2018.
- [14] K. Wan, H. Sun, M. Ji, and G. Caire, "Distributed linearly separable computation," *arXiv preprint arXiv:2007.00345*, 2020.
- [15] N. Woolsey, R. Chen, and M. Ji, "Heterogeneous computation assignments in coded elastic computing," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020.
- [16] H. Dau, R. Gabrys, Y.C. Huang, C. Feng, Q.H. Luu, E. Alzahrani, and Z. Tari, "Optimizing the transition waste in coded elastic computing," in *2020 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2020, pp. 174–178.
- [17] S. Kiani, T. Adikari, and S. C. Draper, "Hierarchical coded elastic computing," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 4045–4049.