

Tackling Heterogeneous Traffic in Multi-access Systems via Erasure Coded Servers

Tuhinangshu
Choudhury
Carnegie Mellon University
Pittsburgh, PA

Weina Wang
Carnegie Mellon University
Pittsburgh, PA

Gauri Joshi
Carnegie Mellon University
Pittsburgh, PA

ABSTRACT

Since most data generated by modern applications is stored in the cloud, there is an exponential growth in the volume of jobs to access these data and perform computations using them. The volume of data access or computing jobs can be heterogeneous across different files stored in the cloud and can unpredictably change over time. Cloud service providers cope with this demand heterogeneity and unpredictability by over-provisioning the number of servers hosting each job type. In this paper, we propose the addition of erasure-coded servers that can flexibly serve multiple job types without additional storage cost. We analyze the service capacity region and the response time of such erasure-coded systems and compare them with standard uncoded replication-based systems currently used in the cloud. We show that coding expands the service capacity region, thus enabling the system to handle variability in demand for different data types. Moreover, we characterize the response time of the coded system in various arrival rate regimes. This analysis reveals that adding even a small number of coded servers can significantly reduce the mean response time, with a drastic reduction in regimes where the demand is skewed across different job types.

ACM Reference Format:

Tuhinangshu Choudhury, Weina Wang, and Gauri Joshi. 2021. Tackling Heterogeneous Traffic in Multi-access Systems via Erasure Coded Servers. In . ACM, New York, NY, USA, 22 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Modern-day cloud computing systems are used for various big-data applications such as performing machine learning (ML) inference tasks [8], hosting large files for web services [21], computing the PageRank of web graphs [19], and other data-intensive computations. Since these jobs require access to the specific data stored on the cloud server(s), each server is dedicated to one job type depending on the availability of data and high-performance hardware required for it. Therefore, the massive volume of users performing such cloud computing jobs have to contend for the server(s) storing data relevant to their job. For example, an ML inference system

may store one trained model on each server, users that need to perform inference using that model are directed to that server. Or a cloud storage system hosting files for Netflix might use each server to store one movie, and users requesting that movie are assigned to that server. Therefore, the massive volume of users performing different cloud computing jobs have to contend for the server(s) storing data relevant to their job.

The user traffic for different job types can vary due to diurnal variations or unpredictable demand fluctuations. These traffic variations often exhibit a negative correlation, i.e., if one job type is experiencing high traffic, another job type experiences low traffic. For example, in an ML inference system hosting different specialized models to process pictures captured by users, the inference traffic for each model can vary periodically depending on the time zone of the users accessing them. Such computing and inference jobs are latency-sensitive. Thus, a larger number of servers needs to be allocated to job types experiencing high traffic to satisfy the latency requirements of those users. The ideal solution for such situations is to enable dynamic allocation of servers where the number of servers provided to a job type depends on its current traffic. However, the servers often host large files (for example, the size of the Google Web crawler is over 10 million GB) and have specialized computing hardware. Dynamically reconfiguring a server for a different job type would require the movement of large amounts of data and may even be possible due to hardware constraints. Hence, there is a critical need to design multi-access computing systems resilient to traffic variations that do not require dynamic reconfiguration of servers.

One standard solution to handle traffic variations is to over-provision the number of servers dedicated to various job types to meet their peak demand. Overprovisioning can be achieved by adding replicas of servers dedicated to a job type [7, 23] in proportion to the maximum historical demand for that job type over a large enough time horizon. While over-provisioning can meet latency requirements under traffic variations, it comes at the cost of severe underutilization of expensive computing resources and a massive energy footprint. Another solution is to supplement job-type-specific servers with flexible general servers that can serve more than one job type. However, adding such servers can be expensive because they need access to data, memory, and computing capabilities relevant to multiple job types. One such model was studied by Tsitsiklis and Xu in [24] which showed that the addition of a small number of flexible servers could give a dramatic reduction in the queueing delay.

While the solutions mentioned above provide some robustness to variations in traffic, none of them considers the correlation between traffic for various job types. For negatively correlated arrival

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

rates, the overall traffic in the system varies slowly over time, even though individual jobs may experience significant changes. A novel approach that is especially well-suited for such skewed traffic patterns was proposed in [1, 2], where they supplement replicas of the servers of each type by a set of *erasure coded servers*. When a job is sent to an erasure-coded server, the output is a linear combination of the outputs of two or more of the regular (uncoded) servers. For example, in a matrix computation task where job type 1 (or type 2) seeks to compute the product Ax (or Bx) of an incoming vector x with the matrix A (or B) stored at a server, a coded server can store $A + B$ such that its output is $Ax + Bx$. Thus, a type 1 job can be served and its output Ax can be obtained using an uncoded server storing B and a coded server storing $A + B$. Using this property, erasure-coded servers can serve multiple job types when combined with regular servers. Furthermore, unlike the flexible servers proposed in [24], the coded servers do not require extra resources such as larger memory and specialized hardware. For various encoding schemes, [1, 2] showed an improvement in the service capacity region of the system, the set of arrival rates for which the system is stable. The addition of coded servers significantly expands the service capacity region, especially in regions where the traffic for different job types is negatively correlated.

1.1 Main Contributions

While [1, 2] proposed the idea of using coded servers to handle traffic variations and showed an expansion of the service capacity region, these works did not analyze the impact of coded servers on other performance metrics such as mean response time or tail latency. Since a coded server has to be used with one or more other servers, coding can increase the system load and result in a higher response time in some traffic regimes. However, this effect on the mean response time is not yet well-understood.

In this paper, we build upon the model provided in [1, 2] and generalize it to consider ≥ 2 job types and an arbitrary allocation of servers to each job type and the coded servers. We compare the coded system with an uncoded system with the same total number of servers and corroborate the insight that the coded system significantly improves the system's stability by increasing the volume of the service capacity region. In addition, we characterize the mean response times of the coded and uncoded systems in several traffic regimes. We show that for a large number of servers, our coded system has a comparable or significantly smaller mean response time in most traffic regimes. The reduction in mean response time is prominent when the traffic of various job types is heavy and skewed, i.e., some job types are experiencing significantly higher traffic than other job types. The only regime where the uncoded system is better is when all job types simultaneously have high traffic, which is unlikely to occur in practical applications. Thus, we show our coded system can better handle heterogeneous traffic for different job types, and it improves the stability as well as the latency of multi-access systems.

1.2 Related Work

Improving the latency of multi-access cloud systems has been extensively studied in the literature. In the context of storage systems, it has been studied in the caching literature [4, 17, 20, 27], where

the number and location of replicas are dynamically adjusted. This paper focuses on data-intensive computing and storage tasks where such dynamic reconfiguration is not feasible.

Another way to tackle the problem is to use replication, as done in MapReduce [9] and Hadoop [21]. More generally, several literatures proposed the usage of coded storage systems to improve latency. Chen *et al.* in [6] considered general servers using (n, k) error correction codes in content replication systems and showed delay optimality for $k \geq 1$. Li *et al.* in [16] considered general servers partitioned to store chunks of coded files encoded using MDS code and showed that coding strictly improves over replication. Erasure codes for fast content downloading have been studied in [13, 14]. Comparison of various coding schemes for delay reduction was studied in [3, 12, 25]. Various scheduling algorithms were also proposed to minimize the system's latency. A comparison between replication-based schedulers and traditional schedulers in a distributed storage system is studied in [7]. Delay optimal scheduling in the centralized and decentralized replication-based system is studied in [23].

The above works use coding mainly for straggler mitigation, where a job is divided into many parallel tasks and waiting for the slowest task is the bottleneck. Coding introduces redundancy into the job to reduce the effect of stragglers. In contrast, we use coding to handle skewed traffic of different job types.

2 PROBLEM FORMULATION

We consider a multi-access cloud system consisting of n servers that are used to provide service to k types of jobs. Jobs of each type i arrive into the system according to a Poisson process. We assume that the traffic variation happens at a slower time-scale such that the system reaches a steady state before the traffic pattern changes. Therefore, it suffices to consider the setting where type i jobs have a constant arrival rate λ_i . The service times of n servers are unit-rate exponential random variables and are i.i.d. across servers and the jobs assigned to the server.

Uncoded System. In most current implementations, the n servers are divided into k disjoint subsets S_i , $i = 1, 2, \dots, k$, each consisting of $|S_i| = n_i = \alpha_i n$ servers for fractions $0 < \alpha_i < 1$ such that $\sum_{i=1}^k \alpha_i = 1$. The fraction α_i of servers dedicated to type i jobs is a hyperparameter that can be set by the system designer. For example, α_i can be set to be proportional to the long-term average of past values of the arrival rate λ_i . Since future arrival rates are not known beforehand or they may vary with time, the number of servers dedicated to each job type may not be sufficient to satisfy its arrival rate. Dynamically re-configuring a server to serve a different job type is expensive and slow in cloud systems, because it involves the movement of large amounts of data. Thus, we consider that the server types are fixed beforehand and cannot be modified on the fly, that is, type i systematic servers cannot be quickly reconfigured to serve jobs of type j for $j \neq i$.

Coded System. In this paper, we consider a more generalized setup where out of n servers, we reserve a set S_{coded} of $n_{\text{coded}} \triangleq |S_{\text{coded}}|$ servers to host *erasure coded* versions of the job types. The remaining $n - n_{\text{coded}}$ systematic servers are split across the k job-types such that $\alpha_i(n - n_{\text{coded}})$ servers are the systematic set S_i of

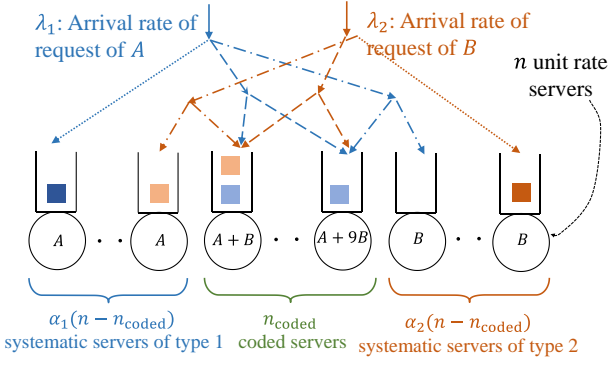


Figure 1: An example of a multi-access computing system hosting two matrices A and B , and consisting of n unit-rate servers, out of which n_{coded} are coded servers. The value written inside each server is the matrix that is stored in the server. Tasks corresponding to A is represented by blue squares, while the one for B is shown using orange squares. Different line types represents different ways in which a job can be served.

job type i , with $0 < \alpha_i < 1$ for all $i = 1, \dots, k$ and $\sum_{i=1}^k \alpha_i = 1$. We illustrate this server assignment for the $k = 2$ case in Fig. 1.

Erasure codes [10] were originally developed in communication systems to add redundancy to transmitted messages in order to provide resilience against noise in the communication channel. In this paper, we employ them for a new purpose – to handle skews in the arrival rate of various job types. When the arrival rate λ_i of type i jobs exceeds the cumulative service rate of the type i systematic servers, the excess arrivals can be served using combination of the coded servers and systematic servers of type(s) $j \neq i$. To illustrate the benefit of coded servers, we first give two concrete examples of erasure coded systems for $k = 2$ job types. Then we describe the proposed coded system for general k .

Example 1. (Coded storage system) Consider an online storage platform with two video files, V_1 and V_2 of equal size, and 3 servers that can host one file each. Consider that the servers store V_1 , V_2 and $(V_1 \oplus V_2)$ respectively, where $(V_1 \oplus V_2)$ is the bit-wise XOR of V_1 and V_2 . Note that the data-size stored on the coded server $(V_1 \oplus V_2)$ is the same as the file size V_1 and V_2 . The jobs of type 1 and 2 are download requests for files V_1 and V_2 . A type 1 job can be served in two ways: 1) it can be sent to the systematic server storing V_1 , or 2) it can be sent to the systematic server storing V_2 and the coded server storing $(V_1 \oplus V_2)$, and after downloading these two files, V_1 can be recovered by taking their bit-wise XOR, $(V_1 \oplus V_2) \oplus V_2 = V_1$. Thus, if the server storing V_1 cannot meet the demand for file V_1 , the excess requests can be served using V_2 and $(V_1 \oplus V_2)$.

To design more general coded storage systems with more than two files and more than one coded server, the files can be represented in a higher alphabet size than bits which will allow more flexible linear combinations. Such erasure coding of files is currently used in commercial cloud storage systems such as RAID [5] for the purpose of resilience against disk failures.

Besides coded storage systems where the jobs represent download requests, erasure coding can be applied to computing systems, as illustrated in Example 2 below.

Example 2. (Coded computing system) Consider an online computing system, where a computing job seeks to find the product of an input vector x with the matrix A (or B). Such jobs are common in machine learning inference systems, where the vector x is an inference query and the matrix-vector product Ax (or Bx) is the predicted output of a linear model. Suppose we have a system consisting of 5 servers, where each server stores one of two large matrices A and B , or their linear combination. Consider that the 5 servers store A , B , $(A + B)$, $(A - B)$, and $(A + 2B)$, respectively. Given an input vector x , the server multiplies the vector with the stored matrix and outputs the matrix-vector product. A type 1 job that seeks to find the product Ax can be served in the following ways: 1) sending the job to the server storing A , 2) sending the job to any two of the three coded servers $(A + B)$, $(A - B)$ and $(A + 2B)$ and obtain Ax by taking a linear combination of the resulting matrix-vector products, or 3) sending the job to the systematic server storing B and any one of the three coded servers, and solving for Ax from the resulting matrix-vector products.

Since erasure codes are inherently linear, the coded computing framework described above can be directly applied only to linear computations such as matrix-vector multiplication. However, some recent research in coding theory is designing ways to apply erasure codes to non-linear computations [15, 18] such as kernel methods and neural networks. The queueing and scheduling insights presented in this paper can be extended to the non-linear coded computing frameworks proposed in these works.

Maximum-Distance Separable Codes and Recovery Sets. The examples shown above considered just 2 types of jobs. More generally, when there are k types of jobs, we propose a coded multi-access system that employs a class of erasure codes called *maximum-distance-separable (MDS) codes*. MDS codes are often used in distributed storage systems to provide resilience against disk failures. If there are k files that need to be stored on l disks, an (l, k) MDS code constructs l independent linear combinations of the k files such that a file i can be recovered from any set of k coded files. A commonly used MDS code is the Reed-Solomon code, which constructs the linear combinations by evaluating a $k - 1$ -th degree polynomial at l points. A special case of an (l, k) MDS code is the *systematic MDS code*, where k of the l combinations are uncoded copies of the k files, and the remaining $l - k$ are other independent linear combinations. In our coded multi-access cloud system, we consider such a systematic MDS code. We have $|S_i| = \alpha_i(n - n_{\text{coded}})$ servers storing uncoded copies of each job type i and n_{coded} independent linear combinations of the k job types.

In our MDS coded system, we can serve a type i job using one of the following options: 1) one of the systematic servers from set S_i , 2) any k coded combinations from the set S_{coded} of coded servers, or 3) k_1 coded servers and $k - k_1$ distinct systematic servers from the sets S_j , where $j \in \{1, 2, \dots, k\} \setminus \{i\}$ and for some integer k_1 such that $1 \leq k_1 < k$. We denote the union of all these possible subsets of servers that can serve a type i job by R_i , and refer to it as the *set of recovery sets*. The size $|R_i|$ is the number of possible recovery sets for job type i , and R_{ij} denotes the set of servers corresponding to the

j -th recovery set, for $j = 1, \dots, |R_i|$. For instance, in Example 1, the set of recovery sets for the file V_1 is $R_1 = \{\{1\}, \{2, 3\}\}$, consisting of the two possible combinations that can be used to recover V_1 .

Queueing model. Next, we describe the queueing model for our system, which is also illustrated in Fig. 1. We consider a first-come-first-served (FCFS) queue at each server. When a type i job arrives, it needs to be assigned to a recovery set in R_i immediately. Specifically, if the chosen set consists of a single systematic server of type i , then the job needs to receive service from that systematic server, and thus we say that *the job consists of one task*. If the chosen set contains $(k - k_1)$ systematic servers and k_1 coded servers for some $k_1 \geq 0$, then the job needs to receive service from all the k servers, and thus we say that *the job consists of k tasks*. This service model induces the following queueing dynamics: when a job arrives, we send a task to the queue at each server in the chosen set in R_i . The job is completed when all of its tasks are completed. We call the policy that determines which recovery set to assign to each arriving job as the *routing policy*.

Performance Metrics. Since the coded system provides the flexibility of having multiple ways of serving a job, it improves load balancing. However, this flexibility comes at the cost of redundancy because to serve a job using a coded combination, we need responses from k servers. To compare the coded and uncoded systems in terms of flexibility vs redundancy trade-off, we use two performance metrics: (1) the service capacity region, the set of arrival rate vectors for which the system is stable, and (2) mean response time of jobs. We define these metrics formally below.

Definition 1 (Service Capacity Region). *The service capacity region of a multi-access cloud system, denoted by Λ , is the region such that for any arrival rate vector $\lambda = (\lambda_1, \dots, \lambda_k)$ in the interior of the region, there exists a routing policy under which*

$$\lim_{c \rightarrow \infty} \lim_{t \rightarrow \infty} \mathbb{P}(\text{Number of jobs in the system at time } t > c) = 0. \quad (1)$$

The service capacity region consists of all supportable throughput. Therefore, it is a notion independent of policies, and it measures the fundamental limit of the system.

Definition 2 (Mean Response time). *The response time T of a job is the total time that a job spends in the system from its arrival until all its tasks are completed. If the system is stable, then the mean response time of the system, denoted by $\mathbb{E}[T]$, is defined as follows when the limit on the right-hand-side exists with probability 1:*

$$\mathbb{E}[T] = \lim_{m \rightarrow \infty} \frac{1}{m} \sum_{j=1}^m T_j. \quad (2)$$

where T_j denotes the response time of the j -th job that departs from the system. For an unstable system, we define the mean response time $\mathbb{E}[T]$ to be ∞ .

The mean response time is a performance metric specific to the routing policy. In this paper, we consider probabilistic job routing policies. In particular, for any R_i , we fix a probability vector \mathbf{p}_i of length $|R_i|$, and any incoming job of type i is assigned to R_{ij} with probability p_{ij} , the j -th element of \mathbf{p}_i . Then, the total arrival rate of type- i job to the set R_{ij} is given by $\lambda_{ij} = \lambda_i p_{ij}$.

When comparing the mean response times of the coded and uncoded systems, we compare the best achievable performances by considering the optimal probabilistic routing policy that minimizes the mean response time. For example, for the uncoded system, the optimal routing policy corresponds to sending a job of type i to one of the systematic servers of type i chosen uniformly at random. Obtaining the optimal routing policy for the coded system can be difficult. We provide some routing policies that perform well, both theoretically and practically in Section 4 and Section 5.

For the analysis of the mean response time of the coded and uncoded system, we assume that the number of coded servers is $n_{\text{coded}} = o(n)$. The case of $n_{\text{coded}} = \Theta(n)$ is both feasible and interesting. However, the comparison of the mean response times of the uncoded and coded systems becomes intractable in this regime. That is why we focus on the $n_{\text{coded}} = o(n)$ for the response time characterization presented in Section 4. However, we provide some insights for on the $n_{\text{coded}} = \Theta(n)$ regime in our simulations in Section 5.

We also assume that $\lambda_i = \Theta(n)$ for all i . We believe that this is a mild assumption because the traffic experienced by a job is usually proportional to the servers allocated to the job type.

Organization of the paper. In Section 3, we analyze the service capacity region of the coded and uncoded system. Section 4 compare the mean response times of the coded and uncoded system. In Section 5.1, we present simulation results that corroborate our theoretical result given in Theorem 2 for a fixed arrival rate vector. In Section 5.2, we provide simulations with variable arrival rate vectors and show how negative correlation in traffic benefits the coded system. Finally, in Section 6 we provide a conclusion and directions for future work.

3 SERVICE CAPACITY REGION

The service capacity region is the set of job arrival rates for which the system is stable, that is, the cumulative arrival rate to any server does not exceed its service rate, as defined in Definition 1. Lemma 1 and Theorem 1 below provide the service capacity regions for the uncoded and coded multi-access cloud systems. We show that the service capacity region of the coded system is significantly larger than that of the uncoded system.

Lemma 1 (Uncoded Service Capacity Region). *Consider an uncoded multi-access cloud system consisting of n servers and k job types, with $|S_i| = \alpha_i n$ servers allocated to job type i . Then, the service capacity region is the set*

$$\Lambda_{\text{uncoded}} = \{(\lambda_1, \dots, \lambda_k) : 0 \leq \lambda_i \leq \alpha_i n, \forall i \in \{1, \dots, k\}\}. \quad (3)$$

The service capacity region of the uncoded system is the hyper-cuboid of size k and length of the i -th side being $\alpha_i n$. This comes from the fact that the total service capacity of systematic servers of type i is $\alpha_i n$. The orange region in Fig. 2 illustrates the service capacity region of the uncoded system for $k = 2, 3$.

Next, we characterize the service capacity region of the coded system in Theorem 1. We first characterize the service capacity region of the coded system using (4). However, for any vector λ , it is difficult to verify whether λ satisfies (4) or not. Hence, we provide a simpler characterization via (5) that is necessary and sufficient condition for an arrival rate to be stabilizable.

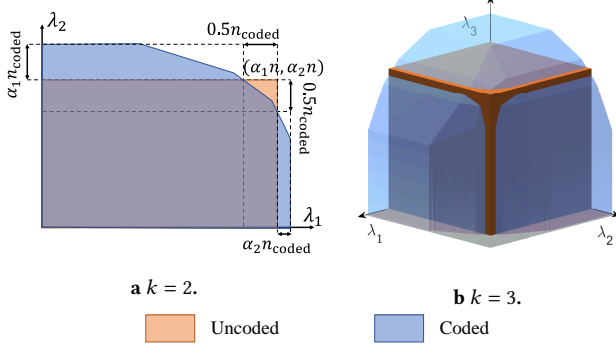


Figure 2: Service Capacity region of our system for $k = 2, 3$. The blue region represents the service capacity region for the coded system, and the orange one shows the service capacity region for the uncoded system. The service capacity region of the coded system expands in regimes where the traffic is skewed. However, it loses some area when all job types have similar and high arrival rates.

Theorem 1 (Coded Service Capacity Region). *Consider a coded multi-access cloud system consisting of n servers and k job types. Let R_{ij} be the j -th element of R_i , the set of recovery sets of type i . Then, the service capacity region is the set*

$$\Lambda_{\text{coded}} = \left\{ (\lambda_1, \dots, \lambda_k) : \exists \lambda_{ij} \geq 0, \text{ s.t. } \lambda_i = \sum_{j=1}^{|R_i|} \lambda_{ij}, \forall i \in \{1, \dots, k\}, \right. \\ \left. \sum_{i=1}^k \sum_{j: \ell \in R_{ij}} \lambda_{ij} \leq 1, \forall \ell \in \{1, \dots, n\} \right\} \quad (4)$$

Let r_i denote the residual capacity for type i jobs, given as $r_i \triangleq \alpha_i(n - n_{\text{coded}}) - \lambda_i$. Also define $r_i^+ \triangleq \max\{r_i, 0\}$, and $r_i^- \triangleq -\min\{r_i, 0\}$. Without loss of generality, assume that $r_1 \leq r_2 \leq \dots \leq r_k$. Then, an arrival rate vector $(\lambda_1, \dots, \lambda_k) \in \Lambda_{\text{coded}}$ if and only if

$$\min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+}{k_0} \right\} \geq \sum_{i=1}^k r_i^-. \quad (5)$$

The proof sketch of Theorem 1 is given in Section 3.1 below, and the complete proof is given in the Appendix B.

For the special case of $k = 2$ and 3, the blue regions in Fig. 2 illustrate the service capacity region of the coded system. The service capacity region for $k = 2$ was previously derived in [1, 2] and is also included in Appendix B. However, [1, 2] did not consider the general case when the number of job types $k > 2$.

From Fig. 2a, observe that there is an expansion in the service capacity region in the top-left and bottom-right areas, the regions where the traffic is heavy and skewed. The routing policy reduces the traffic in the higher loaded systematic servers by redirecting some of their load to the coded servers and less loaded systematic servers. The induced load balancing then allows the traffic of a job type to extend beyond the capacity of its systematic servers, leading to an expansion in the service capacity region.

However, the coded system loses a portion of the service capacity region in the top right corner, where the traffic is heavy but

not skewed. In this region, the traffic of all job types is usually greater than the total capacity of their systematic servers, hence the systematic servers are not enough on their own to serve all job types. In addition, the coded servers add a significant amount of redundancy which further increases the traffic, effectively making it unstable.

For $k = 2$, the total area of the regions gained is $\Theta(nn_{\text{coded}})$ since the maximum possible arrival rate of a job of type i can increase by $\Theta(n_{\text{coded}})$. Likewise, the area of lost region is $\Theta((n_{\text{coded}})^2)$ which makes the total area of service capacity region in the coded system to be $(\alpha_1 \alpha_2 n^2 + \Theta(nn_{\text{coded}}) - \Theta((n_{\text{coded}})^2))$. If the number of coded servers is small, i.e., $n_{\text{coded}} = o(n)$, then there is an effective gain in the service capacity region for the coded system.

3.1 Proof Sketch of Theorem 1

The proof of Theorem 1 contains two parts, the proofs of (4) and (5), and both follow a similar pattern. We first prove that if any arrival rate vector satisfies the property (the requirement in (4) or (5)), there is a way to stabilize the system. We then prove that if the property is not satisfied, no policy can stabilize the system.

We first present the exact proof of (4) and then the proof sketch of (5). The detailed proof of (5) is provided in Appendix B.

Proof of equation (4). We first prove that if any arrival rate vector λ satisfies (4), then there exists a way to stabilize the system. By the definition of the set Λ_{coded} given in Theorem 1, for any arrival rate vector $\lambda \in \Lambda_{\text{coded}}$, there exists a set of λ_{ij} 's that satisfies (4). One routing policy for this arrival rate vector is to serve a job of type- i using servers in the recovery set R_{ij} with probability (λ_{ij}/λ_i) . Under this policy, the total arrival rate to ℓ -th server is $\sum_{i=1}^k \sum_{j: \ell \in R_{ij}} \lambda_{ij}$, where the inner sum represents the total arrival rate of tasks corresponding to the job of type i . Equation (4) ensures that for all $\ell \in \{1, \dots, n\}$, the total arrival rate to ℓ -th server is less than 1, the service rate of the ℓ -th server, which implies that the system is stable.

The other direction of the result, i.e., any λ outside Λ_{coded} makes the system unstable, can be proven using standard techniques involving the *Strict separation theorem* and the *Strong law of large numbers* (see Chapter 4.2 of [22] for an example). We omit the proof here for conciseness.

Proof sketch of equation (5). We first prove that for any arrival rate vector satisfying (5), there exists a policy that stabilizes the system. We use a water-filling argument where we initially serve a job only using its systematic servers. We start using the coded servers when the systematic servers reach their maximum capacity. Then, the right-hand side of (5) is the total excess service requirement after systematic servers are fully utilized. The left-hand side is the total service capacity available at the coded servers plus the excess service capacity at the underutilized systematic servers. The system is stable if the excess service capacity exceeds or equals the excess service requirement. A pictorial representation of the water-filling argument is given in Fig. 3.

We then show that every arrival rate vector $\lambda \in \Lambda_{\text{coded}}$ satisfies (5). The key idea is that for any $\lambda \in \Lambda_{\text{coded}}$, there exist λ_{ij} 's satisfying (4) such that for any i either a job of type i only use its own systematic servers or type i systematic servers is used to serve only

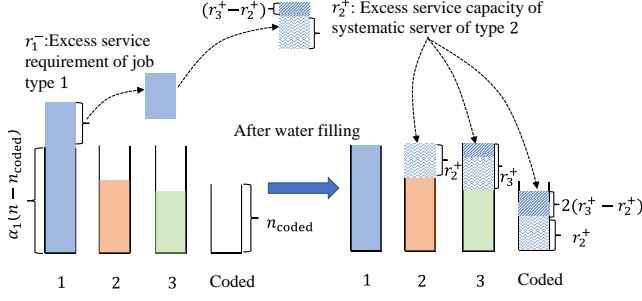


Figure 3: An illustration of the water filling method for $k = 3$ used in the proof of (5). The water-filling strategy involves first filling the systematic servers to the brink. The remaining service requirement is then served using the coded servers and the under-utilized systematic servers.

jobs of type i . This decomposition of λ ensures that the service requirement from the coded servers exceeds $(\sum_{i=1}^k r_i^-)$. Moreover, the service capacity available at the coded servers plus the under-utilized systematic servers is less than $(n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+)/k_0$ for any $k_0 \in \{1, \dots, k\}$. The proof concludes by using the fact that the system is stable which implies equation (5) is true.

4 RESPONSE TIME CHARACTERIZATION

In this section, we compare the minimum mean response times of the coded and uncoded systems, denoted by $\mathbb{E}[T_{\text{coded}}^{(n)}]$ and $\mathbb{E}[T_{\text{uncoded}}^{(n)}]$ respectively. For the remainder of the paper, we use the notation superscript (n) to denote that the quantity depends on n . To characterize and compare the response times, we consider five different traffic regimes listed and illustrated in Fig. 4. We formally define these regimes in Section 4.1, and then we compare the response times in each of these regimes in our main theorem in Section 4.2.

Note that Fig. 4 illustrates the *orders* of the edges of different areas as n becomes large. Hence, the slanted edges in Fig. 2a translates to vertical and horizontal lines in Fig. 4.

4.1 Traffic Regimes

Recall that jobs of type i arrive at the system according to a Poisson process with rate $\lambda_i^{(n)}$. Let $\beta_i^{(n)}$, referred to as the *slack capacity* of systematic servers of type i , be defined as

$$\beta_i^{(n)} = \alpha_i n - \lambda_i^{(n)}. \quad (6)$$

We define five traffic regimes based on the orders of the slack capacities of all job types. For ease of exposition, we first describe the regimes for the case where the system has two job types, i.e., $k = 2$, and then we generalize them to system with any value of k .

Traffic regimes for $k = 2$. Without loss of generality, we assume that $\beta_1^{(n)} \leq \beta_2^{(n)}$.

Light regime: $\beta_i^{(n)} \geq 0$ for all $i \in \{1, 2\}$, and

$$\beta_i^{(n)} = \omega\left(\sqrt{nn_{\text{coded}}^{(n)}}\right), \quad \text{for all } i \in \{1, 2\}.$$




Figure 4: A pictorial representation of the traffic regimes for the case $k = 2$. The star represents a possible value of arrival rate vector and $\beta_i^{(n)}$ is represented by its distance for the edge. The response time comparison of the coded and uncoded system for these regimes is provided in Theorem 2.

This is the regime where the slack capacities are large for both job types.

Inner-heavy regime: $\beta_i^{(n)} \geq 0$ for all $i \in \{1, 2\}$, and

$$\beta_1^{(n)} = o\left(\sqrt{nn_{\text{coded}}^{(n)}}\right), \beta_2^{(n)} = \Omega\left(n_{\text{coded}}^{(n)}\right), \beta_2^{(n)} = \omega\left(\beta_1^{(n)}\right).$$

In this regime, type 1 jobs experience heavier traffic than type 2 jobs do, and thus type 1 jobs can benefit from a coded system.

Outer-heavy regime: $\beta_i^{(n)} \geq 0$ for all $i \in \{1, 2\}$, and

$$\beta_1^{(n)} = o\left(n_{\text{coded}}^{(n)}\right), \beta_2^{(n)} = \omega\left(n_{\text{coded}}^{(n)}\right).$$

Compared to the inner-heavy regime, the outer-heavy regime has an even heavier traffic for type 1 jobs, so the traffic intensities are further skewed.

Uncoded-unstable: This is the regime where the uncoded system is unstable while the coded system is stable.

Coded-unstable: This is the regime where the coded system is unstable while the uncoded system is stable.

Traffic regimes for a general k . Without loss of generality, we assume that $\beta_1^{(n)} \leq \beta_2^{(n)} \leq \dots \leq \beta_k^{(n)}$. To generalize the definitions of the five traffic regimes, we divide the k job types into two groups based on their slack capacities, referred to as the *beneficiaries* and the *helpers*.

Intuitively speaking, the beneficiaries are the job types that experience heavier traffic, and the helpers are those who experience lighter traffic. Therefore, the beneficiaries can benefit from sending jobs to coded servers and systematic servers of the helpers. In contrast, the helpers only need their own systematic servers. Formally, consider the index i^* defined as

$$i^* = \max \left\{ j : \left(j \sum_{i=1}^j \alpha_i \right) < 1 \right\}. \quad (7)$$

Then we call job type in $\{1, 2, \dots, i^*\}$ the beneficiaries, and job types in $\{i^* + 1, \dots, k\}$ the helpers. Note that in the special case where

$k = 2$, the index $i^* = 1$, and thus job type 1 is the beneficiary and job type 2 is the helper. In another special case where $\alpha_i = 1/k$ for all i , the index i^* is around \sqrt{k} .

The choice of i^* in (7) has an intuitive explanation based on the stability of coded servers as follows. For any j , if there are j beneficiaries, then their total traffic is given by $\sum_{i=1}^j (\alpha_i n - \beta_i^{(n)})$. We demonstrate in Appendix C that a good routing policy is as follows. Each helper assigns its jobs only to its own systematic servers. Each beneficiary assigns its jobs to its own systematic servers with probability $(1 - \Theta(n_{\text{coded}}^{(n)}/n))$, and to a recovery set consisting of $(k - j)$ systematic servers from helpers and j coded servers with the rest of the probability. Under this routing policy, the total traffic to the coded servers is $j \cdot \Theta(n_{\text{coded}}^{(n)}/n) \left(\sum_{i=1}^j \alpha_i n - \beta_i^{(n)} \right)$ which is $\Theta(n_{\text{coded}}^{(n)} \left(j \sum_{i=1}^j \alpha_i \right))$. The choice of the index i^* in (7) then intuitively ensures the stability of coded servers.

Now, to define the traffic regimes, it is sufficient to look at the orders of $\beta_1^{(n)}$ and $\beta_{i^*+1}^{(n)}$. We define the traffic regimes as follows.

Light regime: $\beta_i^{(n)} \geq 0$ for all $i \in \{1, 2, \dots, k\}$, and

$$\beta_i^{(n)} = \omega \left(\sqrt{nn_{\text{coded}}^{(n)}} \right), \text{ for all } i \in \{1, 2, \dots, k\}.$$

Inner-heavy regime: $\beta_i^{(n)} \geq 0$, for all $i \in \{1, 2, \dots, k\}$, and

$$\beta_1^{(n)} = o \left(\sqrt{nn_{\text{coded}}^{(n)}} \right), \beta_1^{(n)} = \Omega \left(n_{\text{coded}}^{(n)} \right), \beta_{i^*+1}^{(n)} = \omega \left(\beta_1^{(n)} \right).$$

Outer-heavy regime: $\beta_i^{(n)} \geq 0$, for all $i \in \{1, 2, \dots, k\}$, and

$$\beta_1^{(n)} = o \left(n_{\text{coded}}^{(n)} \right), \beta_{i^*+1}^{(n)} = \omega \left(n_{\text{coded}}^{(n)} \right).$$

The *coded-unstable* and *uncoded-unstable* regimes are defined in the same way as those given in Section 4.1. Note that the assumption of $n_{\text{coded}}^{(n)} = o(n)$ is necessary for defining the light, inner-heavy and outer-heavy regimes, but it is not required to define the uncoded-unstable and coded-unstable regimes. A sufficient condition for an arrival rate vector to lie inside the coded-unstable regime is given as $\beta_{i^*+1}^{(n)} = \omega(n_{\text{coded}}^{(n)})$ and $\beta_{i^*+1}^{(n)} \geq 0$. Likewise, a sufficient condition for an arrival rate vector to lie inside the uncoded-unstable regime is given as $|\beta_{i^*}^{(n)}| = o(n_{\text{coded}}^{(n)})$, $\beta_{i^*}^{(n)} \leq 0$, $\beta_{i^*+1}^{(n)} \geq 0$, and $\beta_{i^*+1}^{(n)} = \omega(n_{\text{coded}}^{(n)})$. The proof of the sufficient conditions are given in Appendix D.

4.2 Main Result

We state our main result in Theorem 2, which compares the response time in a coded system with that in an uncoded system in the five traffic regimes.

Theorem 2 (Response Time Comparison). *Consider a multi-access cloud system consisting of n servers and k job types. Consider the minimum response times in a coded system and in an uncoded system over all probabilistic job assigning policies, denoted by $\mathbb{E} \left[T_{\text{coded}}^{(n)} \right]$ and $\mathbb{E} \left[T_{\text{uncoded}}^{(n)} \right]$, respectively. Then we have the following comparison in*

the five traffic regimes:

$$\text{Light regime: } \left| \mathbb{E} \left[T_{\text{coded}}^{(n)} \right] - \mathbb{E} \left[T_{\text{uncoded}}^{(n)} \right] \right| = o(1); \quad (8)$$

$$\text{Inner-heavy regime: } \mathbb{E} \left[T_{\text{coded}}^{(n)} \right] \leq \mathbb{E} \left[T_{\text{uncoded}}^{(n)} \right] - \omega(1); \quad (9)$$

$$\text{Outer-heavy regime: } \mathbb{E} \left[T_{\text{coded}}^{(n)} \right] = o \left(\mathbb{E} \left[T_{\text{uncoded}}^{(n)} \right] \right); \quad (10)$$

$$\text{Uncoded-unstable regime: } \mathbb{E} \left[T_{\text{coded}}^{(n)} \right] < \infty, \mathbb{E} \left[T_{\text{uncoded}}^{(n)} \right] = \infty; \quad (11)$$

$$\text{Coded-unstable regime: } \mathbb{E} \left[T_{\text{coded}}^{(n)} \right] = \infty, \mathbb{E} \left[T_{\text{uncoded}}^{(n)} \right] < \infty. \quad (12)$$

4.2.1 Proof sketch of Theorem 2. In this subsection, we provide a proof sketch for Theorem 2. We first analyze the mean response time in an uncoded system, and then analyze the mean response time in a coded system.

Uncoded system. The arrival rate of jobs of type i is $\lambda_i^{(n)}$, which is served using $\alpha_i n$ servers. One can prove that the optimal probabilistic routing policy is to serve an incoming job of type i using one of the systematic server of type i chosen uniformly at random. Under this optimal policy, all systematic servers of type i behave like independent $M/M/1$ queues with arrival rate $\lambda_i^{(n)}/(\alpha_i n)$ and unit service rate. The mean response time of a type i job is then $1/\left(1 - \frac{\lambda_i^{(n)}}{\alpha_i n}\right) = \frac{\alpha_i n}{\beta_i^{(n)}}$. Therefore, the mean response time of the uncoded system is given by

$$\mathbb{E} \left[T_{\text{uncoded}}^{(n)} \right] = \sum_{i=1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}}. \quad (13)$$

Coded system. For the coded system, the routing probability is the key component deciding the mean response time. To better understand job routing policies in the coded system, we first give the Property 1 we observe for any stabilizing policy.

Property 1. *Consider any routing policy and let $q_{i0}^{(n)}$ be the total probability of sending a job of type i to a systematic server of type i . Then, the policy can stabilize the system only if*

$$1 - q_{i0}^{(n)} = O \left(\frac{n_{\text{coded}}^{(n)}}{n} \right), \text{ for all } i. \quad (14)$$

Property 1 states that the fraction of traffic that any job type can divert away from its own systematic servers is limited to a $O(n_{\text{coded}}^{(n)}/n)$ fraction. This provides a lower bound on the traffic that has to be served by systematic servers, which further leads to a lower bound on the mean response time. We use this lower bound in the analysis of the light regime. Moreover, in the heavy regimes where traffic is more skewed, we need to divert the traffic of heavily loaded job types as much as we can. In this case, Property 1 also provides a guideline for choosing a good job routing policy. The analysis of the individual regimes then proceeds as follows.

(1) *Light regime:* To prove (8), it suffices to show a lower bound,

$$\mathbb{E} \left[T_{\text{coded}}^{(n)} \right] \geq \mathbb{E} \left[T_{\text{uncoded}}^{(n)} \right] + o(1), \quad (15)$$

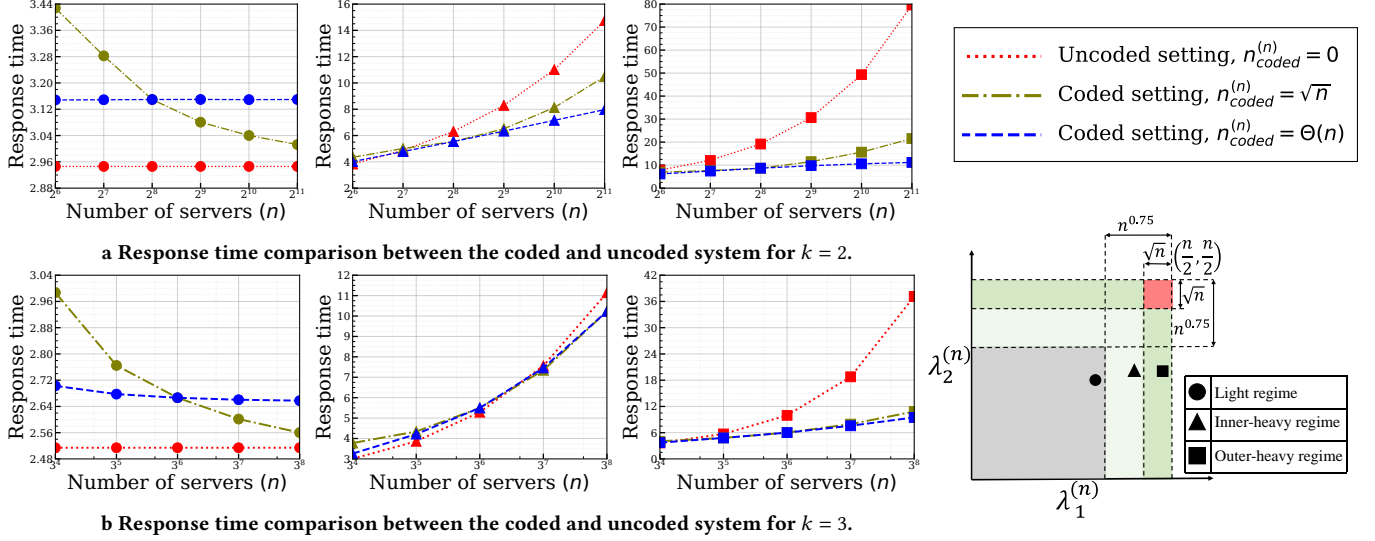


Figure 5: Comparison of the response times of the coded and uncoded systems. The different marker types (circle, triangle, square) represent arrival rates from different traffic regimes (light, inner-heavy and outer-heavy), as shown in the bottom-right illustration. The empirical standard error of the plotted points is $O(10^{-3})$. The coded servers significantly improve the system's performance in the inner-heavy and outer-heavy regimes. The coded and uncoded systems have similar response times in the light regime.

and an upper bound given by

$$\mathbb{E} \left[T_{\text{coded}}^{(n)} \right] \leq \mathbb{E} \left[T_{\text{uncoded}}^{(n)} \right] + o(1). \quad (16)$$

As mentioned earlier, the lower bound is proved using Property 1. To show the upper bound, note that it suffices to focus on a particular routing policy and show that its mean response satisfies (16). In the light regime, the slack capacity for every job type is large enough. Therefore, we consider the routing policy that assigns every job to its own systematic server. Computing the corresponding mean response time verifies (16).

(2) *Heavy regimes*: The analysis of the inner-heavy and outer-heavy regimes follow the same structure. In these regimes, the beneficiaries experience heavy traffic and have small slack capacities. To reduce the mean response times, we divert the traffic of beneficiaries from their systematic servers to the recovery sets that utilize coded servers as much as possible.

In fact, we consider the following routing policy. For some appropriate index $k^* \leq i^*$, we choose the routing probability $q_{i0}^{(n)} = 1$ for all $i > k^*$. For any $i \leq k^*$, we use the routing probability $q_{i0}^{(n)} = 1 - vn_{\text{coded}}^{(n)}/n$ and $q_{ik^*}^{(n)} = vn_{\text{coded}}^{(n)}/n$, for some constant v .

For any $i > k^*$, if the routing option is chosen corresponding to the probability $q_{ik^*}^{(n)}$, then a recovery set consisting of k^* coded servers and $(k - k^*)$ systematic servers of type $k^* + 1, \dots, k$ respectively is chosen uniformly at random from all recovery sets satisfying the property. For any i , if the routing option is chosen corresponding to the probability $q_{i0}^{(n)}$, then a systematic server of type i is chosen uniformly at random. Upper-bounding the mean response time for this policy gives the upper bounds in (9) and (10)

5 SIMULATION RESULTS

In this section, we present our simulation results to demonstrate the performance comparison between the uncoded and coded system under various traffic settings. In Section 5.1, we focus on arrival rates that are time-invariant. Our main goal is to demonstrate the performance comparison given in Theorem 2, but we have also investigated the choice of the $n_{\text{coded}}^{(n)}$ not covered in Theorem 2. In Section 5.2, we consider arrival rates that are time-varying, with a traffic pattern commonly observed in practical systems. A detailed description of the simulation setups are provided in Appendix E.

Before we get into the simulation settings, we first describe the routing policy we use in the simulations for the coded system.

Pseudo-optimal routing policy. Each queue behaves like a $M/M/1$ queue; hence the response time of each task is an exponentially distributed random variable. However, finding the optimal routing policy is non-trivial since the response time of a job is the maximum of the response time of its tasks, and the queues at each server are not independent. Moreover, the routing policy discussed in Section 4.2.1 does not perform well empirically for smaller values of n even though it works well asymptotically.

The difficulty in obtaining the optimal routing policy is the dependence among queues. We derive a policy that we call the *pseudo-optimal* routing policy by treating the queues as if they were independent. This approximation is based on the commonly observed phenomenon that queues are asymptotically independent in large systems [26]. With the independence assumption, one can calculate a job's response time as the expectation of maximum of independent exponential random variables is known. The pseudo-optimal routing policy is then the policy that minimizes the approximated




Figure 6: A comparison of the coded and uncoded systems for time-varying traffic in terms of total job in the system. The leftmost picture shows the search frequency of the words *news* and *movies* on Google, demonstrating a real-world example of negatively correlated traffic. The central figure shows the arrival rate vector as a function of time for a simplified system with two job types that has a negative correlation in traffic. The rightmost figure shows the number of jobs in the system. The presence of coded servers helps reduce the load from the heavier loaded systematic servers, making the coded system enjoy a lesser total job in the system on average.

mean response time. In our simulations, we find the pseudo-optimal routing policy numerically using Scipy Optimization libraries.

5.1 Time-invariant Arrival Rates

In this subsection, we experimentally demonstrate the performance comparison between a coded system and an uncoded system. We provide simulations for systems with 2 and 3 job types. For these simulations, apart from using $n_{\text{coded}}^{(n)} = o(n)$, we also consider $n_{\text{coded}}^{(n)} = \Theta(n)$ and show the effect of large number of coded servers. We simulate until 10^8 jobs leave the system and average it over 50 runs to calculate the mean response time. Finally, based on our simulation results, we provide intuitions on how our main result would change for the case of $n_{\text{coded}}^{(n)} = \Theta(n)$.

Simulation for systems with 2 job types. We consider $n = 2^m$ servers, where we vary $m \in \{6, 7, \dots, 11\}$. For the uncoded system, we calculate the response time theoretically. For the coded system, we consider two cases of $n_{\text{coded}}^{(n)} = o(n)$ and $n_{\text{coded}}^{(n)} = \Theta(n)$.

Fig. 5a shows numerical comparison between mean response of the coded and uncoded system for $k = 2$. The results for the case of $n_{\text{coded}}^{(n)} = o(n)$ resembles the theoretical results provided in Theorem 2. In the light regime, the coded systems with $o(n)$ coded servers perform similar to the uncoded system with a diminishing performance gap as n increases. In the inner-heavy and outer-heavy regimes, the coded system with $o(n)$ servers outperforms the uncoded system and the gap increases as n increases. Moreover, the performance gap between the coded and uncoded systems increases with the skewness in arrival rate, i.e., the coded system performs significantly better in the outer-heavy regime.

However, as illustrated in Fig. 5a, the results are slightly different when the number of coded servers increases as $n_{\text{coded}}^{(n)} = \Theta(n)$. In the light regime, the coded system with $\Theta(n)$ coded server performs worse than the uncoded system. The redundancy added by $\Theta(n)$ coded servers worsens the system. However, the performance gap

does not change much as n increases as the traffic is light enough. The cost of redundancy is not substantial, and the coded system is worse only by an $\Theta(1)$ term. Compared to the light regime, the coded system with $\Theta(n)$ servers performs considerably better than both uncoded and coded systems with $o(n)$ servers in the inner-heavy and outer-heavy regimes. Because of the higher skew, the coded system with $\Theta(n)$ coded servers allows more uniform load balancing, thus greatly improving the performance.

Simulation for systems with 3 job types. We also perform experiments for system with three types of jobs. We consider $n = 3^m$ servers, where we vary $m \in \{4, 5, \dots, 8\}$.

Fig. 5b shows the response time comparison of the coded and uncoded system for this simulation setup. For the light and outer-heavy regime, the trends in Fig. 5b is similar to the simulation results for the two-job type system and hence follows a similar reasoning. However, for the inner-heavy regime, the trends are slightly different. The coded system outperforms the uncoded system asymptotically; however, the difference is not as significant as seen in the simulation with two job types. One plausible reason is that the system has two beneficiaries and only one helper based on our arrival rate choice and allocation of servers. Hence, a slight skew in the arrival rate vector is insufficient to reduce the beneficiaries' load. However, there is enough skew in the outer-heavy region such that the coded system outperforms the uncoded system regardless of the number of coded servers.

Based on the simulations for systems with two or three job types, we conjecture that for $n_{\text{coded}}^{(n)} = \Theta(n)$, the inner-heavy and outer-heavy regimes would merge into a single heavy regime where the coded system would outperform the uncoded system. The light regime would also change, and instead of $o(1)$, the mean response time's of the coded and uncoded system can differ by $\Theta(1)$.

5.2 Time-Varying Arrival Rates

In practical systems such as Google's search system, the traffic of a job type usually varies across time and often shows a periodic nature. For example, as shown in Fig. 6a, search for news in Google [11] usually peaks during early hours, while the search for movies peaks during the night.

To model this situation, we consider a system with two job types where the arrival rate of each job type behaves like a pulse wave as illustrated in Fig. 6b. The arrival rates of the two job types have a similar period but are negatively correlated, i.e., if one job experiences higher traffic, the other job type should experience lesser traffic. We consider a total of 60 servers for both the coded and uncoded system, out of which the coded system consists of 7 coded servers. The remaining servers are distributed in the same proportion to the two job types for the coded and uncoded systems.

Fig. 6c shows a comparison between the mean number of jobs in the coded and uncoded system. The top plot shows the total number of jobs in the system as a function of time, while the second and the third plot shows the number of jobs of type 1 and 2, respectively. One can observe that whenever a job is experiencing low traffic, the number of jobs of that type in the uncoded system is slightly less than that in the coded system. However, whenever a job is experiencing heavy traffic, the number of jobs in the uncoded system is significantly higher than the coded system. The uncoded system is underutilizing the lighter loaded systematic servers, while those underutilized systematic servers are used in the coded system to serve heavier traffic jobs.

6 CONCLUSION

This paper proposes the usage of erasure-coded servers in a multi-access cloud system. We characterize the service capacity region of the coded system showing significant improvements in regions of skewed traffic. The addition of erasure-coded servers improves load balancing as the coded system can handle excess demand compared to the uncoded system. We also compare the latency of the coded and uncoded systems and show that the coded system is better or at least comparable in most cases. The erasure-coded servers improve the system's flexibility, making it an excellent solution in skewed traffic regimes by enabling the usage of servers experiencing low traffic, which would otherwise be idle. Thus, at a slight cost of redundancy, our coded solution provides a general framework to improve the system's stability and latency.

There are substantial directions for future work. Our analysis of the coded system considered any arbitrary value of k . However, a large k is usually not preferred as the decoding cost scales as $O(k^3)$. One approach for a system with large k is to divide it into r subsystems with k/r job types in each subsystem. A large r will save the decoding cost, but it loses some flexibility offered by coding. Hence, one can analyze the optimal choice of r and optimal strategies to group the k job types into r bins. Another open direction can be to analyze our system for queue length based routing policies instead of probabilistic ones. Then the technical challenges in comparing the latency involve proving a complicated state-space collapse and lower bounding the mean response time.

ACKNOWLEDGMENTS

This work was supported in part by the NSF CCF grants #2007834 and #2045694, NSF CNS grants #2007733 and #2112471, and a Carnegie Bosch Institute Research Award. The authors are grateful to Mor Harchol-Balter and Isaac Grosz for helpful discussions.

REFERENCES

- [1] Mehmet Aktaş, Sarah E. Anderson, Ann Johnston, Gauri Joshi, Swanand Kadhe, Gretchen L. Matthews, Carolyn Mayer, and Emina Soljanin. 2017. On the service capacity region of accessing erasure coded content. In *Proc. Ann. Allerton Conf. Communication, Control and Computing*. Monticello, IL, USA, 17–24.
- [2] Mehmet Aktaş, Gauri Joshi, Swanand Kadhe, Fatemeh Kazemi, and Emina Soljanin. 2021. Service Rate Region: A New Aspect of Coded Distributed System Design. *IEEE Trans. Inf. Theory* 67, 12 (2021), 7940–7963.
- [3] Ganesh Ananthanarayanan, Ali Ghodsi, Scott Shenker, and Ion Stoica. 2012. Why Let Resources Idle? Aggressive Cloning of Jobs with Dolly. In *Proc. USENIX Conf. Hot Topics in Cloud Computing (HotCloud)* (Boston, MA), 17.
- [4] L. Breslau, Pei Cao, Li Fan, G. Phillips, and S. Shenker. 1999. Web caching and Zipf-like distributions: evidence and implications. In *Proc. IEEE Int. Conf. Computer Communications (INFOCOM)*, Vol. 1. New York, NY, USA, 126–134.
- [5] Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, and David A. Patterson. 1994. RAID: high-performance, reliable secondary storage. *ACM Comput. Surv.* 26 (1994), 145–185.
- [6] Shengbo Chen, Yin Sun, Ulas, C. Kozat, Longbo Huang, Prasun Sinha, Guanfang Liang, Xin Liu, and Ness B. Shroff. 2014. When queueing meets coding: Optimal-latency data retrieving scheme in storage clouds. In *Proc. IEEE Int. Conf. Computer Communications (INFOCOM)*. 1042–1050.
- [7] Walfredo Cirne, Francisco Brasileiro, Daniel Paranhos, Luis Fabricio W. Góes, and William Voorsluys. 2007. On the efficacy, efficiency and emergent behavior of task replication in large distributed systems. *Parallel Comput.* 33, 3 (2007), 213–234.
- [8] William Dally. 2015. High-performance Hardware for Machine Learning. *NIPS Tutorial* (2015).
- [9] Jeffrey Dean and Luiz André Barroso. 2013. The Tail at Scale. *ACM Commun.* 56, 2 (Feb. 2013), 74–80.
- [10] E. Berlekamp. 1968. *Algebraic coding theory*. McGraw-Hill, New York, USA.
- [11] Google. [n.d.]. *Google Trends*. <https://trends.google.com/trends>
- [12] Sushant Jain, Michael Demmer, Rabin Patra, and Kevin Fall. 2005. Using Redundancy to Cope with Failures in a Delay Tolerant Network. *SIGCOMM Comput. Commun. Rev.* 35, 4 (Aug. 2005), 109–120.
- [13] Gauri Joshi, Yanpei Liu, and Emina Soljanin. 2013. On the Delay-Storage Trade-Off in Content Download from Coded Distributed Storage Systems. *IEEE J. Sel. Areas Commun.* 32 (05 2013).
- [14] Gauri Joshi, Emina Soljanin, and Gregory Wornell. 2017. Efficient Redundancy Techniques for Latency Reduction in Cloud Systems. *ACM Trans. Model. Perform. Eval. Comput. Syst.* 2, 2 (April 2017), 30.
- [15] Jack Kosaiian, K. V. Rashmi, and Shivaram Venkataraman. 2019. Parity Models: Erasure-Coded Resilience for Prediction Serving Systems. In *Proc. ACM Symp. Operating Systems Principles (SOSP)* (Huntsville, Ontario, Canada), 30–46.
- [16] Bin Li, Aditya Ramamoorthy, and R. Srikant. 2018. Mean-Field Analysis of Coding Versus Replication in Large Data Storage Systems. *ACM SIGMETRICS Perform. Evaluation Rev.* 3, 1, Article 3 (Feb. 2018), 28 pages.
- [17] Mohammad Ali Maddah-Ali and Urs Niesen. 2016. Coding for caching: fundamental limits and practical challenges. *IEEE Communications Magazine* 54, 8 (Aug. 2016), 23–29.
- [18] Ankur Mallick, Sophie Smith, and Gauri Joshi. 2021. Rateless Codes for Distributed Non-linear Computations. In *IEEE Int. Symp. Topics in Coding (ISTC)*. 1–5.
- [19] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank Citation Ranking: Bringing order to the Web*. Technical Report. Stanford InfoLab.
- [20] Michael Rabinovich and Oliver Spatscheck. 2002. *Web caching and replication*. Vol. 67. Addison-Wesley Boston, USA.
- [21] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, and Robert Chansler. 2010. The Hadoop Distributed File System. In *IEEE Symp. Mass Storage Systems and Technologies (MSST)*. IEEE Computer Society, USA, 1–10.
- [22] R. Srikant and Lei Ying. 2014. *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge University Press, USA.
- [23] Yin Sun, C. Emre Koksal, and Ness B. Shroff. 2017. On Delay-Optimal Scheduling in Queueing Systems with Replications.
- [24] John N. Tsitsiklis and Kuang Xu. 2017. Flexible Queueing Architectures. *Operations Research* 65, 5 (2017), 1398–1413.
- [25] Ashish Vulimiri, Philip Brighten Godfrey, Radhika Mittal, Justine Sherry, Sylvia Ratnasamy, and Scott Shenker. 2013. Low Latency via Redundancy. In *Proc. Conf. Emerging Networking Experiments and Technologies (CoNEXT)* (Santa Barbara,

- California, USA). 283–294.
- [26] Weina Wang, Mor Harchol-Balter, Haotian Jiang, Alan Scheller-Wolf, and R. Srikant. 2019. Delay asymptotics and bounds for multitask parallel jobs. *Queueing Systems* 91, 3–4 (April 2019), 207–239.
- [27] Gala Yadgar, Oleg Kolosov, Mehmet Fatih Aktas, and Emina Soljanin. 2019. Modeling The Edge: Peer-to-Peer Reincarnated. In *2nd USENIX Workshop on Hot Topics in Edge Computing (HotEdge 19)*. USENIX Association, Renton, WA.

A BACHMANN-LANDAU NOTATION

Definition 3 (Bachmann–Landau notation). *For any functions f and g , the Bachmann–Landau notations are given below*

(1) *Small o :*

$$f = o(g) \triangleq \lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 0. \quad (17)$$

(2) *Big O :*

$$f = O(g) \triangleq \lim_{n \rightarrow \infty} \sup \frac{f(n)}{g(n)} < \infty. \quad (18)$$

(3) *Small omega:*

$$f = \omega(g) \triangleq \lim_{n \rightarrow \infty} \left| \frac{f(n)}{g(n)} \right| = \infty. \quad (19)$$

(4) *Big Theta:*

$$f = \Theta(g) \triangleq \exists k_1, k_2 > 0 \text{ s.t. } k_1 \leq \lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} \leq k_2. \quad (20)$$

B PROOF OF SERVICE-CAPACITY REGION

In this section, we prove the service capacity region of system given in (5). We then determine the service capacity region for the special case of $k = 2$. The proof for the special case is taken from [1, 2] and is included for completion.

B.1 Proof of Equation 5

PROOF. Recall the notations given in (5). The residual capacity for a job of type i is given by $r_i = \alpha_i(n - n_{\text{coded}}) - \lambda_i$. Also, r_i^+ and r_i^- are defined as $\max\{r_i, 0\}$ and $-\min\{r_i, 0\}$, respectively. We also assume $r_1 \leq r_2 \leq r_3 \leq \dots \leq r_k$. We need to prove that the system is stable if and only if

$$\min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+}{k_0} \right\} \geq \sum_{i=1}^k r_i^-. \quad (21)$$

The proof consists of two parts.

- (1) *Proof of “if” statement:* Using the water-filling argument, we prove that if (21) is true, the system is stable.
- (2) *Proof of “only if” statement:* We prove that if (21) is not true, then the system is unstable. To prove this, we prove that if an arrival rate vector lies in the service capacity region, it must satisfy (21).

Proof of if statement. We first fill the capacity of the systematic servers, i.e., serve a job using only its dedicated systematic server to ensure minimum redundancy in the system. Once the systematic servers reach full capacity, we gradually fill the coded servers.

Based on the notations, the volume of unserved jobs after we fill the systematic servers is $\sum_{i=1}^k r_i^-$. The remaining capacity of systematic servers of type i is r_i^+ and the capacity of the coded servers is n_{coded} . The system is stable if we can serve the remaining unserved jobs using a coded combination. We prove that using the coded combination, we can serve a volume of $\min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+}{k_0} \right\}$.

If $r_1 \geq 0$, then (21) is satisfied because each term on the left hand side is greater than or equal to zero and the right hand side is exactly zero. Hence, for the remaining proof, we assume that $r_1 < 0$. Let m be the minimum index such that $r_m > 0$.

Initially, we start filling the systematic servers $m, m+1, \dots, k$. If we add v volume of water in the systematic servers, we add $(m-1)v$ volume of water in the coded server. Either the coded servers or the systematic server of type m will fill up. If the coded servers reach the maximum capacity, we stop. Else, we start filling the systematic servers of type $m+1, \dots, k$. For the coded servers, we add water at a m time faster rate. We continue this way until the coded servers fill up. If all systematic servers fill up, we then start filling the coded servers by a multiple of k . We now prove that we can add $\min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+}{k_0} \right\}$ volume of water using this method.

Observe that after the first step, if the coded servers fill up, then the total volume of water added is $n_{\text{coded}}/(m-1)$. If not, the remaining capacity of the coded servers is $(n_{\text{coded}} - (m-1)r_m^+)/m$. In the next step, if the coded servers fill up, the total volume water added is $r_m^+ + (n_{\text{coded}} - (m-1)r_m^+)/m = (n_{\text{coded}} + r_m^+)/m$. If not, the remaining capacity of the coded servers is $(n_{\text{coded}} - (m-1)r_m^+ - m(r_{m+1}^+ - r_m^+)) = (n_{\text{coded}} + r_m^+ - mr_{m+1}^+)/m$. If the coded servers fill up in the third step, the total volume of added $r_{m+1}^+ + (n_{\text{coded}} + r_m^+ - mr_{m+1}^+)/m + 1 = (n_{\text{coded}} + r_m^+ + r_{m+1}^+)/(m+1)$. Following a similar argument, if the coded servers fill up at the n -th step, then the total volume of water added is $(n_{\text{coded}} + r_m^+ + r_{m+1}^+ + \dots + r_{m+n-2}^+)/(m+n-2)$. Moreover, at the end of step $p < n$, the remaining capacity of coded servers is $(n_{\text{coded}} + r_m^+ + r_{m+1}^+ + \dots + r_{m+p-2}^+ - (m+p-2)r_{m+p-1}^+)/(m+p-1)$ which is positive. This implies, for all $p < n$,

$$\frac{n_{\text{coded}} + \sum_{j=m}^{m+p-1} r_j^+}{m+p-1} < \frac{n_{\text{coded}} + \sum_{j=m}^{m+p-2} r_j^+}{m+p-2}. \quad (22)$$

Moreover, since the coded servers fill up at the n -th step, for all $p > n$,

$$\frac{n_{\text{coded}} + \sum_{j=m}^{m+p-1} r_j^+}{m+p-1} > \frac{n_{\text{coded}} + \sum_{j=m}^{m+p-2} r_j^+}{m+p-2}. \quad (23)$$

This proves that the volume of water added is exactly

$$\min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+}{k_0} \right\} = \min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=m}^{k_0} r_i^+}{k_0} \right\} = \frac{n_{\text{coded}} + r_m^+ + r_{m+1}^+ + \dots + r_{m+n-2}^+}{m+n-2}. \quad (24)$$

We also show that the maximum volume of water filled using r_i^+ capacity of systematic servers of type i , for all i and n_{coded} coded capacity cannot exceed $\min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+}{k_0} \right\}$. To prove this, assume V_{max} is the maximum volume of water that can be added. Then the total volume of water added in servers of type $1, 2, \dots, k_0$ and coded servers is at least $k_0 V_{\text{max}}$. This is because anytime one of the systematic servers in $\{1, 2, \dots, k_0\}$ is not used, the rate at which water is added in the coded servers increases by 1, keeping the net rate of water addition at least k_0 . But since total volume of systematic servers of type i for all $i \in \{1, 2, \dots, k_0\}$ and coded servers is $(n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+)$, we have $k_0 V_{\text{max}} \leq (n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+)$. Hence, for any k_0 , $V_{\text{max}} \leq \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+}{k_0}$, which completes the proof.

Proof of only if statement. We now prove that if an arrival rate vector lies in the service capacity region, it must satisfy (21). Assume that the arrival rate vector $\lambda = (\lambda_1, \dots, \lambda_k)$ lies in the interior of the service capacity region. Using the definition of service capacity region in Theorem 1, we know there exists $\lambda_{ij} \geq 0$ for all $i \in \{1, \dots, k\}$, $j \in \{1, 2, \dots, |R_i|\}$ satisfying

$$\lambda_i = \sum_{j=1}^{|R_i|} \lambda_{ij}, \forall i \in \{1, \dots, k\}, \quad (25)$$

$$\sum_{i=1}^k \sum_{j: \ell \in R_{ij}} \lambda_{ij} \leq 1, \forall \ell \in \{1, \dots, n\} \quad (26)$$

We can assume that there exists no i, i' s.t. there exists some $j \in \{1, \dots, |R_i|\}$ and $j' \in \{1, \dots, |R_{i'}|\}$, for which $\lambda_{ij} > 0$, $\lambda_{i'j'} > 0$, R_{ij} contains a systematic server of type i' and $|R_{i'j'}| > 1$. If there exists any such i, i' , then we can use the following operation multiple times such that the condition becomes true.

- (1) If $\lambda_{ij} < \lambda_{i'j'}$: Change $\lambda_{ij} = 0$, $\lambda_{i'j'} = \lambda_{i'j'} - \lambda_{ij}$. Let ℓ be the systematic server of type i' that lies in R_{ij} . Change $\lambda_{i'm}$ to $\lambda_{i'm} + \lambda_{ij}$ where $R_{i'm} = \{\ell\}$. If $R_{i'j'}$ contains a systematic server of type i , say n then change $\lambda_{ik} = \lambda_{ik} + \lambda_{ij}$ where $R_{ik} = \{n\}$, else change $\lambda_{ik} = \lambda_{ik} + \lambda_{ij}$, where $R_{ik} = R_{i'j'}$.
- (2) If $\lambda_{ij} > \lambda_{i'j'}$: Change $\lambda_{i'j'} = 0$, $\lambda_{ij} = \lambda_{ij} - \lambda_{i'j'}$, $\lambda_{i'k} = \lambda_{i'k} + \lambda_{i'j'}$ where $R_{i'k}$ contains only the systematic server of type i' that is in R_{ij} . If $R_{i'j'}$ contains a systematic server of type i say ℓ , change $\lambda_{im} = \lambda_{im} + \lambda_{i'j'}$, where $R_{im} = \{\ell\}$. Else, change $\lambda_{im} = \lambda_{im} + \lambda_{i'j'}$, where $R_{im} = R_{i'j'}$.

Note, the above operation does not increase the net arrival rate of tasks to any server. Hence, the conditions in (25) and (26) are still valid. A consequence of the above assumption is that if a job of type i uses a coded combination, it can only use a systematic server of type j if the job of type j only uses its systematic server.

Let Γ be the set of indices i 's such that jobs of type i only use its systematic servers to serve its job. For all i , let \hat{r}_i^+ be the remaining service capacity of systematic servers of type i . Note, for all $i \in \Gamma^C$, $\hat{r}_i^+ = 0$, since the remaining capacity cannot be used. Similarly, for all i , let \hat{r}_i^- be the volume of the job of type i that is served using a coded combination. Using the fact that the arrival rate vector is in the service capacity region and arguments in the proof of the *Forward* part, we know that

$$\min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} \hat{r}_i^+}{k_0} \right\} \geq \sum_{i=1}^k \hat{r}_i^-. \quad (27)$$

Since the arrival rate vector may not fill the systematic servers, $\hat{r}_i^- \geq r_i^-$. Also, since for all $i \in \Gamma^C$, $\hat{r}_i^+ = 0$, we have for all i , $r_i^+ \geq \hat{r}_i^+$. Hence,

$$\min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} r_i^+}{k_0} \right\} \geq \min_{k_0 \in \{1, \dots, k\}} \left\{ \frac{n_{\text{coded}} + \sum_{i=1}^{k_0} \hat{r}_i^+}{k_0} \right\} \geq \sum_{i=1}^k \hat{r}_i^- \geq \sum_{i=1}^k r_i^-. \quad (28)$$

This implies that if an arrival rate lies within the service capacity region, it must satisfy (21). \square

B.2 Service capacity region for $k = 2$

PROOF. (1) $0 \leq \lambda_1 \leq \alpha_1 n - (1 + \alpha_1)n_{\text{coded}}^{(n)}$: The total service capacity of S_1 servers is $\alpha_1(n - n_{\text{coded}}^{(n)})$. If all requests for jobs of type 1 are sent to S_1 servers, then remaining capacity of S_1 servers is at least $n_{\text{coded}}^{(n)}$. Total service capacity of S_2 servers is $\alpha_2(n - n_{\text{coded}}^{(n)})$. Therefore, the maximum arrival rate that can be supported for jobs of type 2 is $\alpha_2(n - n_{\text{coded}}^{(n)})$ plus the minimum of total service capacity of the coded servers and remaining service capacity of S_1 servers. Hence, $\lambda_2^{(n)} \leq \alpha_2 n + \alpha_1 n_{\text{coded}}^{(n)}$.

(2) $\alpha_1 n - (1 + \alpha_1)n_{\text{coded}}^{(n)} \leq \lambda_1^{(n)} \leq \alpha_1(n - n_{\text{coded}}^{(n)})$: If all requests for jobs of type 1 are sent to S_1 servers, then the remaining capacity of S_1 servers is $\alpha_1(n - n_{\text{coded}}^{(n)}) - \lambda_1^{(n)}$. Hence, the maximum arrival rate for jobs of type 2 that may be sent to a recovery set consisting of one coded server and one systematic server of type 1 is $(\alpha_1(n - n_{\text{coded}}^{(n)}) - \lambda_1^{(n)})$. The remaining capacity of S_{coded} servers is $n_{\text{coded}}^{(n)} - \alpha_1(n - n_{\text{coded}}^{(n)}) + \lambda_1^{(n)}$. Hence, the maximum arrival rate request that may be sent to the recovery sets consisting of two coded servers is $(n_{\text{coded}}^{(n)} - \alpha_1(n - n_{\text{coded}}^{(n)}) + \lambda_1^{(n)})/2$. Total service capacity of S_2 servers is $\alpha_2(n - n_{\text{coded}}^{(n)})$. Hence,

$$\lambda_2^{(n)} \leq \alpha_2(n - n_{\text{coded}}^{(n)}) + \alpha_1(n - n_{\text{coded}}^{(n)}) - \lambda_1^{(n)} + (n_{\text{coded}}^{(n)} - \alpha_1(n - n_{\text{coded}}^{(n)}) + \lambda_1^{(n)})/2 \quad (29)$$

$$= (n - n_{\text{coded}}^{(n)}) \left(1 - \frac{\alpha_1}{2}\right) + \frac{n_{\text{coded}}^{(n)}}{2} - \frac{\lambda_1^{(n)}}{2}. \quad (30)$$

(3) $\alpha_1(n - n_{\text{coded}}^{(n)}) \leq \lambda_1^{(n)} \leq \alpha_1 n + (1/2 - \alpha_1)n_{\text{coded}}^{(n)}$: If the entire service capacity of S_1 is used for jobs of type 1, then the remaining required service capacity for jobs of type 1 is $\lambda_1^{(n)} - \alpha_1(n - n_{\text{coded}}^{(n)})$, which would be served by S_{coded} servers. Hence, the remaining service capacity of S_{coded} servers is $n_{\text{coded}}^{(n)} - 2(\lambda_1^{(n)} - \alpha_1(n - n_{\text{coded}}^{(n)}))$. Hence, the maximum arrival rate request possible for jobs of type 1 is

$$\lambda_2^{(n)} \leq \alpha_2(n - n_{\text{coded}}^{(n)}) + \frac{n_{\text{coded}}^{(n)} - (2\lambda_1^{(n)} - 2\alpha_1(n - n_{\text{coded}}^{(n)}))}{2} = \frac{(2n - n_{\text{coded}}^{(n)})}{2} - \lambda_1^{(n)}. \quad (31)$$

For the remaining portion, we find the service capacity region by fixing $\lambda_2^{(n)}$ and analyzing the range of value of $\lambda_1^{(n)}$. \square

C PROOF OF THEOREM 2

Recall, our system consists of n servers that provide service to k job types. The n servers are divided into two parts, $n_{\text{coded}}^{(n)}$ coded servers, and the remaining are systematic servers. α_i represents the fraction of systematic servers dedicated to the job of type i . Hence, $n_i^{(n)}$, the number of systematic servers for the i -th job type is given by

$$n_i^{(n)} = \alpha_i(n - n_{\text{coded}}^{(n)}). \quad (32)$$

The arrival rate of job type i is $\lambda_i^{(n)}$ and, the residual capacity for job type i is $\beta_i^{(n)}$, defined as

$$\beta_i^{(n)} = \alpha_i n - \lambda_i^{(n)}. \quad (33)$$

Let $q_{ij}^{(n)}$ denote the probability of serving a job of type i using j coded servers. $\mathbb{E}[T_{\text{coded}}^{(n)}]$ and $\mathbb{E}[T_{\text{uncoded}}^{(n)}]$ denotes the response time of the coded and uncoded system, respectively. Let $\nu_i^{(n)}$ denote the arrival rate of tasks to any systematic server of type i . With the notations defined, we first determine the response time of the uncoded system.

In the uncoded system, queues corresponding to systematic servers of type i behave like independent $M/M/1$ systems with arrival rate $(\lambda_i^{(n)}/\alpha_i n) = (1 - (\beta_i^{(n)}/\alpha_i n))$. Hence, the response time of jobs of type i is $(\alpha_i n / \beta_i^{(n)})$. Moreover, the fraction of jobs that belongs to type i out of all incoming jobs is $(\lambda_i^{(n)} / \sum_{\ell=1}^k \lambda_{\ell}^{(n)})$. Hence, the mean response time of the uncoded system is given by,

$$T_{\text{uncoded}}^{(n)} = \sum_{i=1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}}. \quad (34)$$

Because of our assumption that $\beta_1^{(n)} \leq \beta_2^{(n)} \leq \dots \leq \beta_k^{(n)}$, the dominant term in the response time of the uncoded system is $(n/\beta_1^{(n)})$.

In the remainder of the proof, we bound the response time of the coded system. Notice that the response time of the coded system depends on the routing probability $q_{ij}^{(n)}$'s. Hence, the optimal response time of the coded system depends on the optimal routing probability. However, because of the combinatorial nature, solving the optimal routing probability is a complicated problem. In addition, the fork-join nature of the coded system makes the queues dependent. Hence, there is no closed-form expression of the response time in terms of routing probability and arrival rate vectors. We overcome this problem by using an upper bound expression of the response time for a *reasonable* routing probability vector. The proof consists of three separate parts, one for each regime.

C.1 The Light Regime

PROOF OF THEOREM 2 IN THE LIGHT REGIME. The light regime is the set of arrival rate vectors that satisfy $\beta_i^{(n)} \leq \beta_j^{(n)}$, for all $i \leq j$ and $\beta_1^{(n)} = \omega \left(\sqrt{nn_{\text{coded}}^{(n)}} \right)$. For this region, we prove that

$$\left| T_{\text{coded}}^{(n)} - T_{\text{uncoded}}^{(n)} \right| = o(1). \quad (35)$$

In order to prove (35), we lower bound and upper bound the response time of the coded system. In particular, we prove the following properties.

(1) For any routing policy, the response time of the coded system can be bounded by

$$T_{\text{coded}}^{(n)} \geq T_{\text{uncoded}}^{(n)} + o(1). \quad (36)$$

(2) There exists a routing policy for which the response time of the coded system can be bounded by

$$T_{\text{coded}}^{(n)} \leq T_{\text{uncoded}}^{(n)} + o(1). \quad (37)$$

Combining (36) and (37) completes the proof of (35).

First, we prove (36) and lower bound the response time of the coded system. For this proof, we use Property 1 that states that $q_{i0}^{(n)} = 1 - O(n_{\text{coded}}^{(n)}/n)$, for $i = 1, 2$. Then, the arrival rate of tasks to any server in S_i is bounded as follows

$$v_i^{(n)} \geq \frac{q_{i0}^{(n)} \lambda_i^{(n)}}{n_i^{(n)}} = \frac{\left(1 - O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)\right) (\alpha_i n - \beta_i^{(n)})}{\alpha_i (n - n_{\text{coded}}^{(n)})}. \quad (38)$$

Using these, we lower bound the response time of the coded system as follows

$$T_{\text{coded}}^{(n)} = \sum_{i=1}^k \sum_{j=0}^k \frac{\lambda_i^{(n)} q_{ij}^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \times (\text{Mean response time of jobs served using } j \text{ coded servers}) \quad (39)$$

$$\geq \sum_{i=1}^k \frac{\lambda_i^{(n)} q_{i0}^{(n)}}{\lambda_1^{(n)} + \lambda_2^{(n)}} \times (\text{Mean response time of jobs served using 0 coded servers}) \quad (40)$$

$$= \sum_{i=1}^k \frac{\lambda_i^{(n)} q_{i0}^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{1 - v_i^{(n)}} \quad (41)$$

$$\geq \sum_{i=1}^k \left(1 - O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)\right) \left(\frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{1 - \frac{(1 - O(n_{\text{coded}}^{(n)}/n))(\alpha_i n - \beta_i^{(n)})}{\alpha_i (n - n_{\text{coded}}^{(n)})}} \right) \quad (42)$$

$$\geq \sum_{i=1}^k \left(1 - O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)\right) \left(\frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{1 - (1 - O(n_{\text{coded}}^{(n)}/n))(1 - (\beta_i^{(n)}/\alpha_i n))(1 + (n_{\text{coded}}^{(n)}/n))} \right) \quad (43)$$

$$= \sum_{i=1}^k \left(1 - O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)\right) \left(\frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{\frac{\beta_i^{(n)}}{\alpha_i n} + O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)} \right) \quad (44)$$

$$\geq \sum_{i=1}^k \left(1 - O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)\right) \left(\frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \left(1 + O\left(\frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}}\right)\right) \right) \quad (45)$$

$$= \sum_{i=1}^k \left(1 - O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)\right) \left(\frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} + O\left(\frac{nn_{\text{coded}}^{(n)}}{(\beta_i^{(n)})^2}\right) \right) \quad (46)$$

$$= \sum_{i=1}^k \left(1 - O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)\right) \left(\frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} + o(1) \right) \quad (47)$$

$$= \sum_{i=1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} + o(1) \quad (48)$$

$$= T_{\text{uncoded}}^{(n)} + o(1), \quad (49)$$

where

- (42) uses (38) and, Property 1.
- (43) uses $1 + (n_{\text{coded}}^{(n)}/n) \leq (1/(1 - (n_{\text{coded}}^{(n)}/n)))$ for any $0 \leq n_{\text{coded}}^{(n)} \leq n$.
- (44) uses the fact that for $i = 1, 2$

$$O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right) + O\left(\frac{\beta_i^{(n)} n_{\text{coded}}^{(n)}}{n^2}\right) + O\left(\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)^2\right) + O\left(\frac{\beta_i^{(n)} (n_{\text{coded}}^{(n)})^2}{n^3}\right) = O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right), \quad (50)$$

since $n_{\text{coded}}^{(n)} = o(n)$ and $\beta_i^{(n)} = O(n)$ for all i .

- (47) and, (48) uses the fact that $\beta_i^{(n)} \leq \beta_{i+1}^{(n)}$ for all i and $\beta_1^{(n)} = \omega\left(\sqrt{nn_{\text{coded}}^{(n)}}\right)$.

To complete the proof, we now prove (37) and upper bound the response time. To upper bound the response time of the coded system, we show that there exists a routing probability for which $T_{\text{coded}}^{(n)} \leq T_{\text{uncoded}}^{(n)} + o(1)$. Consider the routing probability $q_{i0}^{(n)} = 1$ for all i . The routing probability given above is a sub-optimal routing policy as it never uses the coded servers. Hence

$$T_{\text{coded}}^{(n)} = \sum_{i=1}^k \sum_{j=0}^k \frac{\lambda_i^{(n)} q_{ij}^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \times (\text{Mean response time of jobs served using } j \text{ coded servers}) \quad (51)$$

$$\leq \sum_{i=1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{1 - \frac{\alpha_i n - \beta_i^{(n)}}{\alpha_1 (n - n_{\text{coded}}^{(n)})}} \quad (52)$$

$$\leq \sum_{i=1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{1 - \left(1 - \frac{\beta_i^{(n)}}{\alpha_i n}\right) \left(1 + \frac{2n_{\text{coded}}^{(n)}}{n}\right)} \quad (53)$$

$$= \sum_{i=1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{\frac{\beta_i^{(n)}}{\alpha_i n} - \frac{2n_{\text{coded}}^{(n)}}{n} + \frac{2\beta_i^{(n)} n_{\text{coded}}^{(n)}}{\alpha_i n^2}} \quad (54)$$

$$\leq \sum_{i=1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \left(1 + O\left(\frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}}\right)\right) \quad (55)$$

$$= \sum_{i=1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} + \sum_{i=1}^k O\left(\frac{nn_{\text{coded}}^{(n)}}{(\beta_i^{(n)})^2}\right) \quad (56)$$

$$= T_{\text{uncoded}}^{(n)} + o(1), \quad (57)$$

where

- (53) uses $(1/(1 - (n_{\text{coded}}^{(n)}/n))) \leq 1 + (2n_{\text{coded}}^{(n)}/n)$ for $n_{\text{coded}}^{(n)} = o(n)$, and large enough n .
- (57) uses the fact that $\beta_i^{(n)} \leq \beta_{i+1}^{(n)}$ for all i and $\beta_1^{(n)} = \omega\left(\sqrt{nn_{\text{coded}}^{(n)}}\right)$.

This completes the proof for the light region. \square

C.2 The Inner-heavy Regime

PROOF OF THEOREM 2 IN THE INNER-HEAVY REGIME. The inner-heavy regime is the set of arrival rate vectors that satisfy $\beta_i^{(n)} \leq \beta_j^{(n)}$ for any $i \leq j$, $\beta_1^{(n)} = \Omega(n_{\text{coded}}^{(n)})$, $\beta_1^{(n)} = o\left(\sqrt{nn_{\text{coded}}^{(n)}}\right)$, and $\beta_{i^*+1}^{(n)} = \omega\left(\beta_1^{(n)}\right)$. For this regime, we prove that

$$T_{\text{coded}}^{(n)} \leq T_{\text{uncoded}}^{(n)} - \omega(1). \quad (58)$$

The key idea to upper bound the response time is to use a *reasonable* routing policy that distributes the load uniformly across all servers. From our assumption $\lambda_i^{(n)} = \Theta(n)$, we know that arrival rate of tasks to the systematic servers is $\Theta(1)$. Hence, we want the arrival rate of tasks to the coded servers to be $\Theta(1)$. From Property 1, we know that $1 - q_{i0}^{(n)} = O(n_{\text{coded}}^{(n)}/n)$ for all i . However, to ensure that the arrival rate of tasks to the coded server is $\Theta(1)$, we want $1 - q_{i0}^{(n)} = \Theta(n_{\text{coded}}^{(n)}/n)$ for at least some i . Let k^* be the max k s.t. $k \leq i^*$ and $\beta_k^{(n)} = O\left(\sqrt{nn_{\text{coded}}^{(n)}}\right)$.

Since $\beta_i^{(n)} \leq \beta_j^{(n)}$ for all $i \leq j$, the systematic servers of type i , for $i \in \{1, \dots, k^*\}$, experience heavy traffic. Hence, we want to decrease their load by serving some of those job types using the coded servers. Therefore, we would like to choose the routing probability where $1 - q_{i0}^{(n)} = \Theta(n_{\text{coded}}^{(n)}/n)$, for all $i \leq k^*$. Consider the following routing probability. Define,

$$v = \frac{1}{k^* \sum_{j=1}^{k^*} \alpha_j} \quad (59)$$

For all jobs of type i , such that $i > k^*$,

$$q_{ij}^{(n)} = \begin{cases} 1 & j = 0 \\ 0 & \text{else} \end{cases} \quad (60)$$

For all jobs of type i , such that $i \leq k^*$,

$$q_{ij}^{(n)} = \begin{cases} 1 - \frac{vn_{\text{coded}}^{(n)}}{n} & j = 0 \\ \frac{vn_{\text{coded}}^{(n)}}{n} & j = i^* \\ 0 & \text{else} \end{cases} \quad (61)$$

Moreover, when serving a job of type i , for $i \leq k^*$, according to routing probability $q_{1i^*}^{(n)}$, only systematic servers in S_j , $j > i^*$ are used. Then, $v_i^{(n)}$ and $v_{\text{coded}}^{(n)}$ the arrival rate to servers in S_i , for all $i \leq k$ and, S_{coded} servers respectively, are bounded as follows. For all $i \leq k^*$

$$v_i^{(n)} = \frac{\left(1 - \frac{vn_{\text{coded}}^{(n)}}{n}\right) (\alpha_i n - \beta_i^{(n)})}{\alpha_i (n - n_{\text{coded}}^{(n)})} \quad (62)$$

$$= \frac{1 - \frac{\beta_i^{(n)}}{\alpha_i n} - \frac{vn_{\text{coded}}^{(n)}}{n} + o\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)}{1 - \frac{n_{\text{coded}}^{(n)}}{n}} \quad (63)$$

$$\leq \left(1 - \frac{\beta_i^{(n)}}{\alpha_i n} - \frac{vn_{\text{coded}}^{(n)}}{n} + o\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)\right) \left(1 + \frac{1+v}{2} \frac{n_{\text{coded}}^{(n)}}{n}\right) \quad (64)$$

$$= 1 - \frac{\beta_i^{(n)}}{\alpha_i n} - \frac{v-1}{2} \frac{n_{\text{coded}}^{(n)}}{n} + o\left(\frac{n_{\text{coded}}^{(n)}}{n}\right), \quad (65)$$

In the uncoded system, the servers in S_i for $i \leq k^*$ experiences a higher traffic and arrival rate of tasks in those servers is $(1 - O(\beta_i^{(n)}/n))$. However, in the coded system, because we distribute the load more uniformly, we decrease the arrival rate to $(1 - O(\beta_i^{(n)}/n) - O(n_{\text{coded}}^{(n)}/n))$. For all i s.t. $k^* < i \leq i^*$, the arrival rate is bounded as

$$v_i^{(n)} = \frac{\alpha_i n - \beta_i^{(n)}}{\alpha_i (n - n_{\text{coded}}^{(n)})} \quad (66)$$

$$= \frac{1 - \frac{\beta_i^{(n)}}{\alpha_i n}}{1 - \frac{n_{\text{coded}}^{(n)}}{n}} \quad (67)$$

$$\leq \left(1 - \frac{\beta_i^{(n)}}{\alpha_i n}\right) \left(1 + \frac{1+v}{2} \frac{n_{\text{coded}}^{(n)}}{n}\right) \quad (68)$$

$$= 1 - \frac{\beta_i^{(n)}}{\alpha_i n} + O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right), \quad (69)$$

where (64) uses $1/(1-x) \leq (1+(1+\epsilon)x)$, for any $x, \epsilon > 0$ when $x = o(1)$. The bound for these servers is similar to the light arrival regime, i.e., although the arrival rate is slightly higher compared to the uncoded system, the effect on the mean response time is $o(1)$. Similarly, for all $i > i^*$

$$v_i^{(n)} = \frac{\alpha_i n - \beta_i^{(n)} + \sum_{j=1}^{k^*} \frac{vn^{(n)}_{\text{coded}}}{n} (\alpha_j n - \beta_j^{(n)})}{\alpha_i (n - n_{\text{coded}}^{(n)})} \quad (70)$$

$$= \frac{1 - \frac{\beta_i^{(n)}}{\alpha_i n} + \frac{vn^{(n)}_{\text{coded}}}{\alpha_i n} \sum_{j=1}^{k^*} \alpha_j - o\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)}{1 - \frac{n_{\text{coded}}^{(n)}}{n}} \quad (71)$$

$$\leq \left(1 - \frac{\beta_i^{(n)}}{\alpha_i n} + \frac{vn^{(n)}_{\text{coded}}}{\alpha_i n} \sum_{j=1}^{k^*} \alpha_j - o\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)\right) \left(1 + \frac{2n_{\text{coded}}^{(n)}}{n}\right) \quad (72)$$

$$= 1 - \frac{\beta_i^{(n)}}{\alpha_i n} + \Theta\left(\frac{n_{\text{coded}}^{(n)}}{n}\right), \quad (73)$$

Because of uniform load balancing, the arrival rate of tasks to the servers in S_i for all $i > i^*$ increases, but the order of the load still remains the same. Although the traffic loss for servers of heavy job types is similar to the traffic gain for servers of light jobs, their effect on the response time is not the same. Finally, for the coded servers,

$$v_{\text{coded}}^{(n)} = \frac{k^* \sum_{j=1}^{k^*} \frac{vn^{(n)}_{\text{coded}}}{n} (\alpha_j n - \beta_j^{(n)})}{n_{\text{coded}}^{(n)}} \quad (74)$$

$$= \sum_{j=1}^{k^*} \left(vk^* \alpha_j - \frac{k^* v \beta_j^{(n)}}{n} \right) \quad (75)$$

$$= 1 - \sum_{j=1}^{k^*} \frac{k^* v \beta_j^{(n)}}{n} \quad (76)$$

$$\leq 1 - \frac{(k^*)^2 v \beta_1^{(n)}}{n}. \quad (77)$$

The coded servers experience higher traffic because of our choice of routing policy. However, because we use the coded servers with $O(n_{\text{coded}}^{(n)}/n) = o(1)$ probability, the impact on the mean response time is small. Using this bound on the arrival rate of tasks, we bound the mean response time of the coded system as follows.

$$T_{\text{coded}}^{(n)} = \sum_{i=1}^k \sum_{j=0}^k \frac{\lambda_i^{(n)} q_{ij}^{(n)}}{\lambda_1^{(n)} + \lambda_2^{(n)}} \times (\text{Mean response time of jobs served using } j \text{ coded servers}) \quad (78)$$

$$\leq \sum_{i=1}^{k^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \left(\frac{1 - \frac{vn^{(n)}_{\text{coded}}}{n}}{1 - v_i^{(n)}} + \frac{vn^{(n)}_{\text{coded}}}{n} \left(\sum_{\ell=i^*+1}^k \frac{1}{1 - v_{\ell}^{(n)}} + \frac{i^*}{1 - v_{\text{coded}}^{(n)}} \right) \right) + \sum_{i=k^*+1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{1 - v_i^{(n)}} + \sum_{i=i^*+1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{1 - v_i^{(n)}} \quad (79)$$

$$\leq \sum_{i=1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \left(\frac{1}{\frac{\beta_i^{(n)}}{\alpha_i n} + \frac{v-1}{2} \frac{n_{\text{coded}}^{(n)}}{n} + o\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)} + \frac{vn^{(n)}_{\text{coded}}}{n} \left(\sum_{\ell=i^*+1}^k \frac{1}{\Theta\left(\frac{\beta_{\ell}^{(n)}}{n}\right)} + \frac{i^*}{\frac{(i^*)^2 v \beta_1^{(n)}}{n}} \right) \right) + \sum_{i=k^*+1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{\frac{\beta_i^{(n)}}{\alpha_i n} - O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)} + \sum_{i=i^*+1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{\frac{\beta_i^{(n)}}{\alpha_i n} - \Theta\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)} \quad (80)$$

$$\begin{aligned} &\leq \sum_{i=1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_\ell^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \left(\frac{1}{1 + \frac{\alpha_i(v-1)}{2} \frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}} + o\left(\frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}}\right)} \right) + \Theta\left(\frac{n_{\text{coded}}^{(n)}}{\beta_1^{(n)}}\right) + \\ &\quad + \sum_{i=k^*+1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_\ell^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \frac{1}{1 - O\left(\frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}}\right)} + \sum_{i=i^*+1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_\ell^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \left(\frac{1}{1 - \Theta\left(\frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}}\right)} \right) \end{aligned} \quad (81)$$

$$\begin{aligned} &\leq \sum_{i=1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_\ell^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \left(1 - \frac{\alpha_i(v-1)}{4} \frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}} - o\left(\frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}}\right) \right) + \Theta\left(\frac{n_{\text{coded}}^{(n)}}{\beta_1^{(n)}}\right) + \\ &\quad + \sum_{i=k^*+1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_\ell^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \left(1 + O\left(\frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}}\right) \right) + \sum_{i=i^*+1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_\ell^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \left(1 + \Theta\left(\frac{n_{\text{coded}}^{(n)}}{\beta_i^{(n)}}\right) \right) \end{aligned} \quad (82)$$

$$\leq \sum_{i=1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_\ell^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} - \sum_{i=1}^{k^*} \Theta\left(\frac{nn_{\text{coded}}^{(n)}}{(\beta_i^{(n)})^2}\right) + \sum_{i=k^*+1}^k \Theta\left(\frac{nn_{\text{coded}}^{(n)}}{(\beta_i^{(n)})^2}\right) \quad (83)$$

$$\leq T_{\text{uncoded}}^{(n)} - \omega(1), \quad (84)$$

where

- (79) upper bounds mean of the max response time of tasks by the mean of their sum,
- (82) uses $1/(1-x) \leq (1+(1+\epsilon)x)$, for any $x = o(1)$ and any $\epsilon > 0$,
- (84) uses $\beta_1^{(n)} = o\left(\sqrt{nn_{\text{coded}}^{(n)}}\right)$. Moreover, from the definition of k^* , if $k^* < i^*$, then $\beta_j^{(n)} = \omega\left(\sqrt{nn_{\text{coded}}^{(n)}}\right)$ for all $j > k^*$ in which case the statement is true. If $k^* = i^*$, then we use the fact that $\beta_1^{(n)} = o(\beta_{i^*+1}^{(n)})$.

□

C.3 The Outer-heavy Regime

PROOF OF THEOREM 2 IN THE OUTER-HEAVY REGIME. The outer-heavy regime is the set of arrival rate vectors $\lambda_i^{(n)}$ that satisfies the property that $\beta_i^{(n)} \leq \beta_j^{(n)}$ for all $i \leq j$ and $\beta_i^{(n)} = o(n_{\text{coded}}^{(n)})$ and, $\beta_{i^*+1}^{(n)} = \omega(n_{\text{coded}}^{(n)})$. For this regime, we prove that

$$T_{\text{coded}}^{(n)} = o\left(T_{\text{uncoded}}^{(n)}\right). \quad (85)$$

The proof for the outer-heavy regime is similar to the proof of inner-heavy regime, i.e., we figure out a good routing probability that distributes the load uniformly across the various servers. We consider a similar routing probability that we considered for the inner-heavy regime. We define k^* be the max k s.t. $k \leq i^*$ and $\beta_k^{(n)} = O(n_{\text{coded}}^{(n)})$. Define,

$$v = \frac{1}{k^* \sum_{j=1}^{k^*} c_j} \quad (86)$$

For all jobs of type i such that $i > k^*$,

$$q_{ij}^{(n)} = \begin{cases} 1 & j = 0 \\ 0 & \text{else} \end{cases} \quad (87)$$

For all jobs of type i such that $i \leq k^*$,

$$q_{ij}^{(n)} = \begin{cases} 1 - \frac{vn_{\text{coded}}^{(n)}}{n} & j = 0 \\ \frac{vn_{\text{coded}}^{(n)}}{n} & j = i^* \\ 0 & \text{else} \end{cases} \quad (88)$$

For the routing probability given above, we can reuse the bound for $v_i^{(n)}$ and $v_{\text{coded}}^{(n)}$ given in (65), (69), (73) and, (77) respectively, i.e., for all i s.t. $i \leq k^*$,

$$v_i^{(n)} \leq 1 - \frac{\beta_i^{(n)}}{\alpha_i n} - \frac{v-1}{2} \frac{n_{\text{coded}}^{(n)}}{n} + o\left(\frac{n_{\text{coded}}^{(n)}}{n}\right). \quad (89)$$

For all i s.t. $k^* < i \leq i^*$,

$$v_i^{(n)} \leq 1 - \frac{\beta_i^{(n)}}{\alpha_i n} + O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right). \quad (90)$$

For all i s.t. $i^* < i$,

$$v_i^{(n)} \leq 1 - \frac{\beta_i^{(n)}}{\alpha_i n} + \Theta\left(\frac{n_{\text{coded}}^{(n)}}{n}\right), \quad (91)$$

For the coded servers,

$$v_{\text{coded}}^{(n)} \leq 1 - \frac{(k^*)^2 v \beta_1^{(n)}}{n}. \quad (92)$$

$$v_{\text{coded}}^{(n)} = 1 - \frac{\beta_1^{(n)}}{\alpha_1 n}. \quad (93)$$

The above bounds only use the fact that $n_{\text{coded}}^{(n)} = o(n)$ and $\beta_i^{(n)} = O(n_{\text{coded}}^{(n)})$ for all $i \leq k^*$, which is also true for the outer-heavy regime. However, one key difference to note is that unlike the inner-heavy regime, for the outer-heavy regime there is an order-wise decrease in traffic experienced by the systematic servers of type $i \in \{1, 2, \dots, k^*\}$. The traffic decreased from $(1 - O(\beta_i^{(n)}/n))$ to $(1 - O(n_{\text{coded}}^{(n)}/n))$. This enables the coded system to achieve an order-wise improvement in the mean response time. The response time of the coded system is then upper bounded as follows.

$$T_{\text{coded}}^{(n)} = \sum_{i=1}^k \sum_{j=0}^k \frac{\lambda_i^{(n)} q_{ij}^{(n)}}{\lambda_1^{(n)} + \lambda_2^{(n)}} \times (\text{Mean response time of jobs served using } j \text{ coded servers}) \quad (94)$$

$$\begin{aligned} &\leq \sum_{i=1}^{k^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \left(\frac{1 - \frac{vn_{\text{coded}}^{(n)}}{n}}{1 - v_i^{(n)}} + \frac{vn_{\text{coded}}^{(n)}}{n} \left(\sum_{\ell=i^*+1}^k \frac{1}{1 - v_{\ell}^{(n)}} + \frac{i^*}{1 - v_{\text{coded}}^{(n)}} \right) \right) + \\ &\quad + \sum_{i=k^*+1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{1 - v_i^{(n)}} + \sum_{i=i^*+1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{1 - v_i^{(n)}} \end{aligned} \quad (95)$$

$$\begin{aligned} &\leq \sum_{i=1}^{k^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \left(\frac{1}{\frac{\beta_i^{(n)}}{\alpha_i n} + \frac{v-1}{2} \frac{n_{\text{coded}}^{(n)}}{n} + o\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)} + \frac{vn_{\text{coded}}^{(n)}}{n} \left(\sum_{\ell=i^*+1}^k \frac{1}{\Theta\left(\frac{\beta_{\ell}^{(n)}}{n}\right)} + \frac{i^*}{\frac{(i^*)^2 v \beta_1^{(n)}}{n}} \right) \right) + \\ &\quad + \sum_{i=k^*+1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{\frac{\beta_i^{(n)}}{\alpha_i n} - O\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)} + \sum_{i=i^*+1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{1}{\frac{\beta_i^{(n)}}{\alpha_i n} - \Theta\left(\frac{n_{\text{coded}}^{(n)}}{n}\right)} \end{aligned} \quad (96)$$

$$\begin{aligned} &\leq \sum_{i=1}^{k^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \Theta\left(\frac{n}{n_{\text{coded}}^{(n)}}\right) + \Theta\left(\frac{n_{\text{coded}}^{(n)}}{\beta_1^{(n)}}\right) + \\ &\quad + \sum_{i=k^*+1}^{i^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \frac{1}{1 - o(1)} + \sum_{i=i^*+1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} \left(\frac{1}{1 - o(1)} \right) \end{aligned} \quad (97)$$

$$\leq \sum_{i=1}^{k^*} \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \Theta\left(\frac{n}{n_{\text{coded}}^{(n)}}\right) + \sum_{i=k^*+1}^k \frac{\lambda_i^{(n)}}{\sum_{\ell=1}^k \lambda_{\ell}^{(n)}} \frac{\alpha_i n}{\beta_i^{(n)}} (1 - o(1)) \quad (98)$$

$$\leq \sum_{i=1}^{k^*} o\left(\frac{n}{\beta_i^{(n)}}\right) \quad (99)$$

$$\leq o\left(T_{\text{uncoded}}^{(n)}\right) \quad (100)$$

where

- (95) upper bounds the maximum of response times of tasks by their sum.
- (100) uses the fact that $T_{\text{coded}}^{(n)} = O(n/\beta_1^{(n)})$. This is because of $\beta_1^{(n)} = o(n_{\text{coded}}^{(n)})$ and (34).

Hence, by including coded servers, the order of the response time decreases from $\Theta(n/\beta_1^{(n)})$ to $\Theta(n/n_{\text{coded}}^{(n)})$. \square

D PROOF OF SUFFICIENT CONDITIONS FOR UNSTABLE REGIMES

D.1 Proof of Sufficient Condition for Uncoded-Unstable Regime

PROOF. The region is defined as $|\beta_{i^*}^{(n)}| = o(n_{\text{coded}}^{(n)})$, $\beta_{i^*+1}^{(n)} = \omega(n_{\text{coded}}^{(n)})$ and $\beta_{i^*}^{(n)} < 0, \beta_{i^*+1}^{(n)} \geq 0$. Clearly, the uncoded system is unstable in this regime. To prove the theorem, it suffices to show that the coded system is stable. Let $k^* = \max j$ such that $j \leq i^*$ and $\beta_j^{(n)} < 0$. Then, all jobs of type $i^* + 1, i^* + 2, \dots, k$ can be served by their systematic servers. For a job of type i for $i \leq k^*$, we send $\alpha_i(n - n_{\text{coded}}^{(n)})$ amount of traffic to their systematic servers. For the remaining traffic, we send i^* tasks to the coded server and $(k - i^*)$ tasks to servers, one each to S_{i^*+1}, \dots, S_k . The servers in set S_j for all $j \leq k^*$ are stable.

For $i > i^*$, the total traffic to systematic servers of type i is $(\alpha_i n - \beta_i^{(n)} + \sum_{j=1}^{k^*} (\alpha_j n_{\text{coded}}^{(n)} - \beta_j^{(n)})) = (\alpha_i n - \omega(n_{\text{coded}}^{(n)}))$ which is less than its total service capacity $\alpha_i(n - n_{\text{coded}}^{(n)})$. The traffic to the coded server is $i^* (\sum_{j=1}^{k^*} (\alpha_j n_{\text{coded}}^{(n)} - \beta_j^{(n)}))$ which is less than $n_{\text{coded}}^{(n)}$ using definition of i^* and $|\beta_1^{(n)}| = o(n_{\text{coded}}^{(n)})$. Hence, the system is stable. \square

D.2 Proof of Sufficient Condition for Coded-Unstable Regime

PROOF. Given

$$\beta_{i^*+1}^{(n)} = \omega(n_{\text{coded}}^{(n)}). \quad (101)$$

Total service capacity required by files $\beta_1^{(n)}, \dots, \beta_{i^*+1}^{(n)}$ is $\sum_{i=1}^{i^*+1} (\alpha_i n - \beta_i^{(n)})$. Let p be the fraction of the jobs that are serviced completely by servers not dedicated to the first $(i^* + 1)$ servers. For stability, we require

$$(1 - p) \sum_{i=1}^{i^*+1} (\alpha_i n - \beta_i^{(n)}) \leq \sum_{i=1}^{i^*+1} \alpha_i (n - n_{\text{coded}}^{(n)}). \quad (102)$$

Simplifying the equation, we get,

$$p \geq \frac{\sum_{i=1}^{i^*+1} (\alpha_i n_{\text{coded}}^{(n)} - \beta_i^{(n)})}{\sum_{i=1}^{i^*+1} (\alpha_i n - \beta_i^{(n)})} \quad (103)$$

Hence, the total service required from the coded servers for the first $(i^* + 1)$ job types is at least

$$(i^* + 1) \left(\sum_{i=1}^{i^*+1} (\alpha_i n_{\text{coded}}^{(n)} - \beta_i^{(n)}) \right) \geq n_{\text{coded}}^{(n)}. \quad (104)$$

where (104) uses the definition of i^* given in (7). Hence, the system is unstable. \square

E SIMULATION SETUPS

E.1 Simulations for Fixed Arrival Rate

Simulation for system with two job types. We consider $n = 2^m$ servers, where we vary $m \in \{6, 7, \dots, 11\}$. For the uncoded system, we calculate the response time theoretically. We consider a symmetrical allocation of servers, i.e., $\alpha_1 = \alpha_2 = 0.5$. For each experiment, we run the experiment until 10^8 jobs leave the system and average it over 50 runs to calculate the mean response time. For the coded system, we consider two different values for the number of coded servers $n_{\text{coded}}^{(n)} = n/16$, and $n_{\text{coded}}^{(n)} = \sqrt{n}$. The arrival rate vectors for the various traffic regimes are chosen as follows.

- (1) Light regime: $\lambda^{(n)} = [n/2 - 5n/32, n/2 - 6n/32]$.
- (2) Inner-heavy regime: $\lambda^{(n)} = [n/2 - (n/2)^{0.55}, n/2 - 6n/32]$.
- (3) Outer-heavy regime: $\lambda^{(n)} = [n/2 - (n/2)^{0.3}, n/2 - 6n/32]$.

Simulation for system with three job types. We consider $n = 3^m$ servers, where we vary $m \in \{5, 7, \dots, 8\}$. For the uncoded system, we calculate the response time theoretically. We consider an asymmetrical allocation of servers, i.e., $\alpha_1 = 2/11, \alpha_2 = 0.3/11$, and $\alpha_3 = 6/11$. For each experiment, we run the experiment until 10^8 jobs leave the system and average it over 50 runs to calculate the mean response time. For the coded system, we consider two different values for the number of coded servers $n_{\text{coded}}^{(n)} = n/27$, and $n_{\text{coded}}^{(n)} = \sqrt{n}$. The arrival rate vectors for the various traffic regimes are chosen as follows.

- (1) Light regime: $\lambda^{(n)} = [\alpha_1 n - 5n/48, \alpha_2 n - 7n/48, \alpha_3 n - 9n/48]$.
- (2) Inner-heavy regime: $\lambda^{(n)} = [\alpha_1 n - (\alpha_1 n)^{0.55}, \alpha_2 n - (\alpha_2 n)^{0.65}, \alpha_3 n - 9n/48]$.
- (3) Outer-heavy regime: $\lambda^{(n)} = [\alpha_1 n - (\alpha_1 n)^{0.3}, \alpha_2 n - (\alpha_2 n)^{0.7}, \alpha_3 n - 9n/48]$.

E.2 Simulations for Variable Arrival Rate

We consider a system with 60 servers and two job types. For the uncoded system, the 26 servers are provided to job type 1 and the remaining to job type 2. We allocate 7 coded servers for the coded system and distribute 22 and 31 systematic servers to job types 1 and 2. For the arrival rate of job type 1, we consider a pulse wave with a minimum amplitude of 6 jobs/unit, a maximum amplitude of 18 jobs/unit, and a period of 2000 units. Similarly, for the arrival rate of job type 2, we consider a pulse wave with a minimum amplitude of 12 and a maximum amplitude of 30 and a period of 2000 units. For job type 1, we consider the same lengths for high amplitude and low amplitude duration. For job type 2, we consider 1200 and 800 units of length for high amplitude and low amplitude duration, respectively. Moreover, we consider a 100 units of rightward shift in the period of job type 2 compared to the period of job type 1.