

The Service Rate Region of Hamming Codes

Priyanka Choudhary and Maheshanand Bhaintwal

Department of Mathematics, Indian Institute of Technology Roorkee, Roorkee 247667, India

Abstract

The service rate region of a coded distributed storage system is the set of all achievable data access requests under the capacity constraints. This paper investigates the service rate regions of systematic Hamming codes using hypergraph theory and derives bounds for the maximal achievable service rate of individual data objects. We establish upper bounds on the sum of service rates of data symbols indexed by a subset of systematic nodes in a systematic binary Hamming code, and explore the achievability of these bounds. Additionally, for non-systematic binary Hamming codes, we conclude that the aggregate service rate is limited by the number of columns of odd weight in the associated generator matrix.

Index Terms

Service rate region, hypergraph, vertex cover, Hamming codes, recovery sets.

I. INTRODUCTION

DISTRIBUTED data storage systems have become integral to modern data management, ensuring reliability, scalability, and availability in diverse applications ranging from cloud computing to large-scale storage. To maintain accessibility and safeguard against potential failures, storage systems employ redundancy, either through data replication or coding techniques. While replication has traditionally been a popular method, it is no longer a practically feasible method for exponentially increasing data in modern times. As a result, erasure coding has emerged as a more efficient alternative. Unlike replication, erasure coding reduces storage overhead while maintaining reliability and improving data retrieval efficiency through local recovery, as detailed in works like [1], [3], [4], and [5]. Erasure coding also effectively enhances availability by reducing download latency for retrieving entire data, as highlighted in several research papers, see e.g., [6]–[8]. Among the erasure coding techniques, binary Hamming codes, renowned for their efficient error-correction properties and low overhead, present an interesting area to explore.

Whenever new erasure coding schemes are developed and presented, there is often a question regarding which one is the most effective for a given application. To assess the performance of different coded storage systems, various factors are compared to ensure the selection of a scheme that best aligns with the user's requirements. One critical metric for evaluating the performance of erasure-coded storage systems is the “*service rate region (SRR)*”. This region represents the set of download or access request rates that the system can fulfill without exceeding its service capacity. This concept was initially introduced by Noori et al. [9]. Later, Aktaş et al. [10] extended the framework of storage allocation by connecting it to the service rate region. Characterizing the SRR provides a clear understanding of the overall rate of requests that a system can handle. In simpler terms, the SRR metric serves as a measure of assessing the operational efficiency of distributed coded systems under the conditions of parallel access, while considering the constraints imposed by the limited service capacity of the network's nodes.

This metric is particularly important in scenarios where multiple users access shared content concurrently, often with uneven access patterns. It also highlights the system's resilience to varying request patterns (skewed demand distribution), where certain objects are accessed more frequently than others.

There are two primary problems associated with SRR:

- 1) Determining the SRR for a given coding scheme.
- 2) For a given region R , determining the coding scheme whose SRR contains R with the minimum number of storage nodes.

Research on SRR aims to find solutions to the problems mentioned above. Finding the service rate region of a scheme involves determining the optimal allocation of incoming data download requests to the recovery sets. This process effectively transforms the task of maximizing the service rate region into an optimization problem.

In [12], upper bounds for the SRR of binary simplex codes and first-order Reed-Muller codes were determined through a projective geometric approach. Ly et al. [21] generalized their approach to analyze the service rate region of higher-order RM codes, providing better inner and outer bounds.

It was realized in [2] that the problem of finding a valid allocation for SRR was related to identifying fractional matching in hypergraphs, and is similar to batch codes. Alfarano et al. [16] explored the polytope structure of the service rate region

Priyanka Choudhary is a Ph.D. scholar in the Department of Mathematics, Indian Institute of Technology Roorkee, Roorkee 247667, India (e-mail: priyanka_c@ma.iitr.ac.in).

Maheshanand Bhaintwal is with the Department of Mathematics, Indian Institute of Technology Roorkee, Roorkee 247667, India (e-mail: maheshanand@ma.iitr.ac.in).

and proved that every rational k -tuple of the service rate region has a valid rational allocation. In [18], the authors introduced two dual distance outer bounds for the SRR by using the structural properties of the generator matrix.

The second significant challenge is identifying the minimum number of storage nodes n_R and the best coding scheme for the given region R . This aspect has been analyzed in [13]. The development of a coding scheme for a fixed number of servers aimed at maximizing the SRR was investigated in [10] for the case when $k = 2$, and subsequently in [11] for $k = 3$. In 2023, Alfarano et al. [15] investigated the SRR of locally recoverable codes (LRCs) and batch codes. Kiliç et al. [17] have provided both upper and lower bounds for the minimum number of servers n_R corresponding to a fixed field size. Moreover, for a fixed number of servers, they have obtained bounds on the field size as well.

This research extends the idea of recent developments in the characterization of SRR for MDS-coded distributed systems. Ly and Soljanin [19] established a graph-theoretic framework that connects SRR to fractional matchings (or fractional vertex cover) in quasi-uniform hypergraphs, demonstrating that the SRR polytope expands with systematic columns in generator matrices. There are still many questions that remain unanswered regarding the SRR of various coding schemes, particularly when the generator matrix is non-systematic.

Compared to other coding schemes, characterizing the service rate region of binary Hamming codes using the hypergraph framework is not overly complex due to their structural properties and unique form of parity-check matrices. Recently, Ly and Soljanin [20] explored the connection between combinatorial design theory and SRR. They established that the maximum achievable demand of any data object in a linear code is attained when the minimum weight codewords of the dual code form a 2-design. Through this analysis, they proved that, for systematic binary Hamming codes, the maximum achievable demand of a data object is consistently 3, irrespective of the code's parameters. This result is also obtained in the present paper through a different approach by analyzing the recovery system structure of Hamming codes.

We explore the service rate region of binary Hamming codes thoroughly using the existing approach of hypergraphs [14] [19] in the context of SRR and attempt to find more accurate bounds to characterize their service rate region.

This paper is organized as follows. Section II includes some preliminaries and the fundamental concepts of the service rate region. It also examines the challenge of identifying the SRR from various angles, including combinatorial and geometric perspectives. In Section III, we present the principal findings of this paper, concentrating on the structure of the recovery G system and the maximal achievable service rate of a data object for systematic q -ary Hamming codes, denoted by $\text{Ham}(r, q)$. We also establish an upper bound for the cumulative service rate of non-systematic $\text{Ham}(r, 2)$. Furthermore, we outline an upper limit on the achievable demand of data files corresponding to a subset $\mathcal{I}$ of systematic nodes of systematic $\text{Ham}(r, 2)$ and demonstrate that this limit is indeed attainable. Section IV concludes the paper and proposes some potential directions for future research.

II. PRELIMINARIES

This section introduces the notations, definitions, and prior results that form the foundation of the service rate region of coded schemes in data storage systems. Unless otherwise stated, all notations and terminologies follow standard conventions.

Let $\mathbb{F}_q$ be a finite field of size q . A linear code C of length n and dimension k over $\mathbb{F}_q$ is a k -dimensional subspace of $\mathbb{F}_q^n$, and referred to as an $[n, k]$ -code. Given an information vector with k symbols, an n -symbol codeword can be generated using a valid encoding process. The set of coordinate positions with non-zero entries in a codeword $c \in C$ is known as the support of c , i.e., for any $c \in C$, the support of c , denoted by $\text{Supp}(c)$ is

$$\text{Supp}(c) = \{i \mid c_i \neq 0, 1 \leq i \leq n\}.$$

The minimum distance of a code C is defined as $d(C) = \min\{d(x, y) : x, y \in C, x \neq y\}$, where $d(x, y) = |\{i : x_i \neq y_i\}|$. For a linear code C , its orthogonal complement is referred to as the dual code of C and denoted by $C^\perp$. Thus,

$$C^\perp = \{x \in \mathbb{F}_q^n \mid x \cdot c = 0, \text{ for all } c \in C\},$$

where $x \cdot c$ denotes the usual inner product of x and c .

One can also characterize an $[n, k]$ linear code C by a $k \times n$ generator matrix G of rank k over $\mathbb{F}_q$. The rows of G form an $\mathbb{F}_q$ basis for C . We have $x = (x_1, x_2, \dots, x_k) \in \mathbb{F}_q^k$ as information (data) objects and $x \cdot G \in \mathbb{F}_q^n$ is the corresponding codeword where the k data objects are stored. Then, C can be described as the set $C = \{x \cdot G \mid x \in \mathbb{F}_q^k\}$. An $(n - k) \times n$ matrix H of rank $n - k$ over $\mathbb{F}_q$ is called a parity-check matrix of C if $G \cdot H^\top = 0$. A parity-check matrix H of C is also a generator matrix for the dual code of C .

We consider a coded distributed system consisting of n servers to store k information(data) files represented by elements of $\mathbb{F}_q$. Let C be the corresponding code over $\mathbb{F}_q$ and let G denote a generator matrix of C . One server stores exactly one element of $\mathbb{F}_q$ and has two major functions: one is to store the data, and the other is to process the incoming data access requests from the system's users. Each server possesses a service capacity denoted by μ , which indicates that it is capable of processing μ requests per unit of time. In this storage system, each server stores a linear combination of the k information files, and each column of the matrix G uniquely represents a particular server. A generator matrix G is said to be systematic if there are k columns in the matrix which are the ordered standard basis vectors $e_1, e_2, \dots, e_k$ of $\mathbb{F}_q^k$. The columns of G that are non-zero

scalar multiples of the standard basic vectors are known as systematic servers. For an $[n, k]$ linear code C over $\mathbb{F}_q$, the standard form of a systematic generator matrix is $G = [I_k \mid A]$, and it is well known that $H = [-A^\top \mid I_{n-k}]$ is the corresponding parity-check matrix for C . If another systematic matrix G' of C is obtained through permuting the columns of G , then the corresponding parity-check matrix H' is obtained by applying the same column permutation to H . This convention will be followed throughout this paper.

As the information files are saved across the n servers in coded or systematic servers, we can recover these files from multiple subsets $I \subset [n] = \{1, 2, \dots, n\}$. This leads to the following definition.

Definition 1. A set $R \subset [n]$ is called a recovery set for the i -th data object if the standard basis vector $e_i \in \mathbb{F}_q^k$ belongs to the $\mathbb{F}_q$ -linear span of the column vectors of G indexed by the elements in R .

The set of all recovery sets for the i -th information file is denoted by $\mathcal{R}_i^{\text{all}}$. We define the set $\mathcal{R}(G) = \{\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k\}$ as a recovery G system [14], where each component $\mathcal{R}_i \subseteq \mathcal{R}_i^{\text{all}}(G)$. The recovery system containing all the recovery sets corresponding to the generator matrix G is denoted by $\mathcal{R}^{\text{all}}(G)$.

An element R of $\mathcal{R}_i^{\text{all}}(G)$ is called a minimum recovery set for the i -th data object if there does not exist $R' \in \mathcal{R}_i^{\text{all}}(G)$ such that $R' \subsetneq R$. A minimum recovery G system is

$$\mathcal{R}^{\text{min}}(G) = \{\mathcal{R}_i^{\text{min}} \mid 1 \leq i \leq k\},$$

where $\mathcal{R}_i^{\text{min}}$ contains all the minimum recovery sets of the i -th object.

Example 1. A systematic code C has a systematic generator matrix. Therefore, a minimum recovery set for a data object consists only of the corresponding systematic column of the generator matrix.

Consider a linear code C characterized by a generator matrix G that can be transformed into a systematic form through a permutation of its columns. In this case, the parity-check matrix helps in completely characterizing the recovery system using the codewords of the corresponding dual code. Consequently, we have the following results.

Lemma 1. [18] Let G be a systematic generator matrix of a linear code C . A subset $R \subset [n]$ is a recovery set for the i -th coordinate if and only if one of the following conditions is satisfied:

- 1) The set R contains a systematic server of the i -th coordinate.
- 2) There exists a codeword $c' \in C^\perp$ such that the $\text{Supp}(c') \subseteq R \cup \{i\}$ and $i \in \text{Supp}(c')$.

Lemma 2. Let G be a systematic generator matrix of a linear code C or column permutation equivalent to a systematic generator matrix. If R is a recovery set for the i -th coordinate in C , then either $|R| = 1$ or $|R| \geq d^\perp - 1$, where $d^\perp$ is the minimum distance of the dual code of C .

Proof. When $R = \{i\}$, it is obvious that R is a recovery set for the i -th coordinate. Let us consider R a recovery set such that $R \neq \{i\}$. Using Lemma 1, we know that there exists a codeword $c' \in C^\perp$ such that the $\text{Supp}(c') \subseteq R \cup \{i\}$. Therefore we have,

$$d^\perp \leq |\text{Supp}(c')| \leq |R| + 1.$$

□

Continuing in this direction, we present the following result that helps in determining the number of recovery sets of a data file using the codewords of the dual code.

Lemma 3. [22] Consider an $[n, k, d]$ linear code C over $\mathbb{F}_q$. There are exactly $1/q$ of the codewords of C which have $\alpha \in \mathbb{F}_q$ at any given coordinate position i , or zero in all the codewords at the i -th coordinate position, $1 \leq i \leq n$.

For a code C with a systematic generator matrix G , let $S \subset [n]$ such that $|S| = k$ and all the columns of G indexed by the elements of S are the distinct k systematic columns. From Lemma 3, for each $i \in S$, there are q^{n-k-1} codewords in $C^\perp$ that have 1 in coordinate position i . Therefore, the number of recovery sets $|\mathcal{R}_i^{\text{min}}|$ of the data object at the i -th place satisfies $|\mathcal{R}_i^{\text{min}}| \geq q^{n-k-1}$ for all $i \in S$. This observation will lead to important results in Section III.

In the following subsection, we provide a formal definition of the service rate region associated with a given recovery G system, as well as analyze the bounds previously established for this region.

A. Service Rate Region

Let $\lambda_i \in \mathbb{R}_{\geq 0}$ be the download/access request rate for the i -th data file, where $\mathbb{R}_{\geq 0}$ denotes the set of positive real numbers. The k -tuple $(\lambda_1, \lambda_2, \dots, \lambda_k) \in \mathbb{R}_{\geq 0}^k$ is considered as an achievable service rate if the system can serve it, while satisfying the capacity constraints on the nodes.

Definition 2. [14] Consider a system where k data files are distributed among n servers using an $[n, k]$ linear code characterized by a generator matrix G with the server capacity $\mu \geq 1$. Let $\mathcal{R} = \{\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k\}$ be a recovery G system.

Then the associated service rate region (SRR) is the set of all such k -tuples $(\lambda_1, \lambda_2, \dots, \lambda_k) \in \mathbb{R}^k$ such that for each $i \in [k]$ there exists a list

$$\{\lambda_{iR} \in \mathbb{R} \mid R \in \mathcal{R}_i, \lambda_{iR} \geq 0\},$$

satisfying the following conditions:

$$\sum_{R \in \mathcal{R}_i} \lambda_{iR} = \lambda_i, \quad \forall i \in [k], \quad (1)$$

$$\sum_{i=1}^k \sum_{\substack{R \in \mathcal{R}_i \\ v \in R}} \lambda_{iR} \leq \mu, \quad \forall v \in [n]. \quad (2)$$

This service rate region is denoted by $\Lambda(\mathcal{R}(G), \mu)$.

Equation (1) guarantees that the demand for all data objects is adequately satisfied, while (2) asserts that the total number of requests handled by each server does not exceed its designated serving capacity. A collection of $\{\lambda_{iR}\}$ which satisfies (1) and (2) is called a valid allocation for the given recovery system $\mathcal{R}$ and service capacity μ . Here λ_{iR} is the fraction of λ_i assigned to the recovery set R of the i -th data symbol.

It is also important to note here that to determine the SRR for a given generator matrix G it is not necessary to consider all the recovery sets, because $\Lambda(\mathcal{R}^{all}(G), \mu)$ is the same as $\Lambda(\mathcal{R}^{\min}(G), \mu)$ [18], where

$$\mathcal{R}^{all}(G) = \{\mathcal{R}_1^{all}, \mathcal{R}_2^{all}, \dots, \mathcal{R}_k^{all}\}.$$

Additionally, $\Lambda(\mathcal{R}(G), \mu) = \mu \Lambda(\mathcal{R}(G), 1)$.

For the rest of the paper, we will examine the service rate region for the minimum recovery G system with uniform service capacity $\mu = 1$. We denote the corresponding SRR by $\Lambda(G)$.

It is clear from Definition (2) that the problem of finding the service rate region of a linear code is an optimization problem with λ_{iR} as variables. However, as the value of k increases, the complexity of solving this problem also increases due to the number of variables growing exponentially. To address this challenge, we turn to other mathematical tools to find the service rate region. Here, we briefly introduce two ways to interpret SRR due to [14]: one is using the geometric representation and the other is utilizing the concept of hypergraphs.

B. Geometric interpretation of SRR

Determining the SRR of a coded scheme can be framed as a linear optimization problem where the intersection of several half-spaces encloses the SRR. The feasible region of this problem results in a convex polytope in $\mathbb{R}_{\geq 0}^k$. A comprehensive study of SRR, proving that it is a complex polytope, was done by Alfarano et al. [16], and we have included some of the basic concepts and results from their work here.

Definition 3. A simplex in $\mathbb{R}^k$ is defined as the convex hull of the set $\{\delta_1 e_1, \delta_2 e_2, \dots, \delta_k e_k\}$, where $\delta_i \in \mathbb{R}$ and $e_1, e_2, \dots, e_k$ are the standard basis vectors in $\mathbb{R}^k$. In particular, a δ -simplex is a simplex for which $\delta_i = \delta$ for all $i \in [k]$.

We define $\lambda_i^*(G) = \max\{x \mid x e_i \in \Lambda(G)\}$. Then, the largest possible size of the δ -simplex in the SRR is given by

$$\delta(G) = \min\{\lambda_i^*(G), 1 \leq i \leq k\}.$$

The simplex structure plays a crucial role in characterizing SRR using the terms *maximal achievable simplex* and *maximal matching simplex* [19]. The maximal matching simplex refers to the smallest geometric simplex that contains the SRR of the relevant storage scheme, while the maximal achievable simplex is the largest simplex contained in the SRR. Additionally, the term $\delta(G)$ represents the size of the maximal achievable δ -simplex as it is given by the region of all service rate k -tuples satisfying $\sum_{i=1}^k \lambda_i \leq \delta$.

Remark 1. If there are t disjoint recovery sets present in a coded system with at least one systematic server present in the generator matrix for every information symbol, then we have $(t+1)e_i \in \Lambda(G), \forall i \in [k]$, implying that

$$\lfloor \delta(G) \rfloor \geq (t+1).$$

Lemma 4. [12] [16] For an $[n, k, d]$ linear code with generator matrix G , if $\delta e_i \in \Lambda(G)$ for all $i \in [k]$, then $\lceil \delta \rceil \leq d$. In particular, $\lceil \delta(G) \rceil \leq d$.

C. Hypergraph as recovery G system

Hypergraphs are a generalization of graphs, where an edge (hyperedge) may contain more than two vertices. A hypergraph is represented by $H = (V, E(V))$, where V is a set of vertices and $E(V)$ contains edges, and is a subset of the power set of V . The hypergraphs were first viewed as recovery G systems in [2], where each hyperedge represents a recovery set and is labeled by the corresponding standard basis vector of the data file it recovers. The service rate served by an edge is associated with the fractional matching assigned to that edge. A matching in the hypergraph H is a set of pairwise disjoint (non-adjacent) hyperedges. The matching number of a hypergraph H is the maximum cardinality of a matching in H , denoted by $\nu(H)$. We can also consider matching as assigning a number $w(e) \in \{0, 1\}$ to each edge e such that $\sum_{e \ni v} w(e) = 1$ for each vertex v of the hypergraph.

If we assign a number $w(e)$ to each edge e such that $w(e) \in [0, 1]$ and $\sum_{e \ni v} w(e) \leq 1$ for each vertex v , then this setup is known as fractional matching. The maximum sum of the assigned number $w(e)$ of all the edges in the hyperedge H is known as the *fractional matching number*, and it is denoted by $\mu_f(H)$. Therefore, the problem of splitting the request rates to different recovery sets is now reduced to assigning a fractional matching to the hyperedges. The equivalence of a demand vector being in SRR and the existence of a corresponding fractional matching in the graph representation is proven in [2].

To optimally assign request rates, we need to maximize the fractional matching to the edges. For a recovery G system and the associated recovery hypergraph H , the relation between the cumulative service rate and maximum fractional matching number is given by

$$\sum_{i=1}^k \lambda_i \leq \mu_f(H).$$

Similar to the concept of vertex cover in graphs, we also have the vertex cover/transversal in hypergraphs. A transversal is the set $T \subseteq V$ such that for every $e \in E, T \cap e \neq \emptyset$. The smallest size of a transversal T in H , denoted by $\tau(H)$, is known as the transversal number of H . For a hypergraph H , it stands that $\nu(H) \leq \mu_f(H) \leq \tau(H)$. Note that the size of the maximal matching simplex is equal to the fractional matching number $\mu_f(H)$.

Lemma 5. ([2] Theorem 2) *For a coding scheme with generator matrix G and uniform service capacity $\mu = 1$, if $(\lambda_1, \lambda_2, \dots, \lambda_k) \in \Lambda(G)$, and H is the associated recovery hypergraph, then the following relations are satisfied:*

$$\nu(H) \leq \max \left(\sum_{i=1}^k \lambda_i \right) = \mu_f(H) \leq \tau(H).$$

In the case where demand for some of the data symbols is zero, the service rate region is determined for the remaining data symbols. The partial hypergraph $H_{\mathcal{I}}, \mathcal{I} \subseteq [k]$, is formed by removing those edges from H which are not labeled by the standard basis vectors of the data symbols indexed by $\mathcal{I}$. This also satisfies the matching, fractional matching, and transversal number relation. We will use this fact to obtain new upper limits for the sum of service rates in our coding scheme.

The next subsection presents a demand-splitting technique also detailed in [20], and this technique is proven optimal for MDS codes in [10].

D. A Greedy (Waterfilling) algorithm for request splitting

One of the critical challenges in finding the service rate regions of distributed data storage systems is designing a request-splitting strategy that optimally distributes user requests to storage nodes. The approach is inspired by the well-known greedy algorithm, which involves making the best choice at each step, with the expectation that this approach will lead to a globally optimal solution: Always serve a request by first attempting to assign it to the smallest available recovery set (or resource) that can satisfy it. This way we efficiently use the small size recovery group, keeping in mind the load at every node.

The waterfilling algorithm, as a special case of the greedy algorithm, and inspired by concepts in information theory and signal processing, has emerged as a promising approach to address this challenge. In the context of the service rate region of coded schemes, this algorithm was initially introduced in [10] to characterize the SRR of systematic MDS codes. The algorithm allocates resources (such as power or bandwidth) among multiple channels to maximize overall system performance.

The waterfilling algorithm can be effectively adapted for distributed data storage by allocating user requests among storage nodes to minimize latency, balance load, and maximize throughput. It directs incoming requests to the systematic servers first, and once their capacity is reached, any remaining requests are sent to the least loaded recovery sets, excluding those with exhausted nodes. The load on a recovery set refers to the maximum load among the nodes it contains.

This strategy comprehensively characterizes the SRR of systematic MDS codes for $n \geq 2k$. We utilize the same request-splitting approach for systematic binary Hamming codes, demonstrating that it provides the largest service rate region compared to alternative strategies.

For non-systematic Hamming codes, we aim to use a greedy approach by directing requests to the smallest recovery sets with the least load. This involves keeping an index of all possible recovery sets for each object, sorting or organizing these sets by their cardinality, and potentially enhancing the approach with load-awareness to prevent hotspots.

III. SERVICE RATE REGION OF HAMMING CODES

In this section, we study the service rate region of Hamming codes, particularly of the binary Hamming code $[2^r - 1, k = 2^r - 1 - r, 3]$ denoted by $\text{Ham}(r, 2)$ where $r \geq 2$. For $r = 2$ the code $\text{Ham}(2, 2)$ is the repetition code of length 3 with $\Lambda = \{i \mid 0 \leq i \leq 3\}$. Therefore, the work in this paper focuses on the SRR of $\text{Ham}(r, 2)$ for $r \geq 3$.

We recall the definition of a q -ary Hamming code.

Definition 4. [23] For $r \geq 2$, a linear code C over a finite field $\mathbb{F}_q$ is called a q -ary Hamming code, denoted by $\text{Ham}(r, q)$, if the parity-check matrix of C is defined by choosing for its columns a non-zero vector from each one-dimensional subspace of $\mathbb{F}_q^r$. The code C is a $\left[\frac{q^r-1}{q-1}, \frac{q^r-1}{q-1} - r, 3\right]$ linear code.

The columns of a parity-check matrix of a binary Hamming code are all the non-zero vectors of $\mathbb{F}_2^r$ arranged in such a way that the i -th column is the binary expansion of the i -th number. Any other parity-check matrix of the binary Hamming code can be obtained by column permutation of the initial matrix. If the generator matrix of the Hamming code is systematic, then we say that the code is a systematic Hamming code.

Since there is exactly one systematic column (server) for each data symbol in a systematic generator matrix G of $\text{Ham}(r, q)$, the k -tuple $(1, 1, \dots, 1)$ is in the service rate region of systematic $\text{Ham}(r, q)$ for all r . The remaining r parity columns of G are of weight $q^{r-1} - 1$.

WLOG, we consider the standard form generator matrix for a systematic $\text{Ham}(r, q)$, represented as:

$$G = (I_k, P) = (\underbrace{e_1, e_2, \dots, e_k}_{\text{Standard basic vectors in } \mathbb{F}_q^k}, \underbrace{p_1, p_2, \dots, p_r}_{\text{parity-check columns}}).$$

Throughout this section, we adopt this notation for the systematic generator matrix.

Example 2. The parity-check matrix for the $[7, 4, 3]$ ($r = 3$) binary Hamming code is as follows

$$H = \begin{pmatrix} 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{pmatrix}.$$

Then the corresponding generator matrix for $\text{Ham}(3, 2)$ is

$$G = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 1 \end{pmatrix}.$$

The systematic generator matrix can be obtained by permuting the columns of this matrix.

The dual code of a Hamming code $\text{Ham}(r, q)$ is referred to as a simplex code $S(r, q)$ with the parameters $[\frac{q^r-1}{q-1}, r, q^{r-1}]$. Not only does the simplex code $S(r, q)$ have a minimum distance of q^{r-1} , but all the non-zero codewords in $S(r, q)$ have the same weight q^{r-1} .

From Lemma 1, the support of codewords of a simplex code completely determines the structure and elements of the recovery sets of the data symbols involving the corresponding systematic Hamming coding scheme.

The following Theorem 1 provides a general structure of these recovery sets in the systematic $\text{Ham}(r, q)$ coding scheme.

Theorem 1. Let C be a $\left[\frac{q^r-1}{q-1}, k = \frac{q^r-1}{q-1} - r, 3\right]$ q -ary Hamming code with the systematic generator matrix G . Then the recovery G system satisfies the following conditions:

- 1) A minimum recovery set has cardinality either 1 or $q^{r-1} - 1$.
- 2) There are exactly q^{r-1} non-singleton minimum local recovery sets of an information symbol.
- 3) For an i -th information symbol and any other node j , the cardinality of the set $\{R \mid j \in R\}$, where $R \in \mathcal{R}_i^{\min}$, is $(q-1)q^{r-2}$.

Proof. To prove the first statement, we use the fact that the minimum distance of the dual code of C is q^{r-1} . Due to Lemma 2, the cardinality of a recovery set R for the i -th information symbol is either 1 or at least $q^{r-1} - 1$. As all the codewords of the dual code have weight $q^{r-1} - 1$, we have $|R| = 1$ or $|R| \leq q^{r-1} - 1$. This proves the first statement.

Any codeword of $S(r, q)$ with a non-zero entry at the i -th place is associated with a recovery set for the i -th coordinate. From Lemma 3, we can say that there are exactly q^{r-1} codewords in the simplex code $S(r, q)$ which have 1 at the i -th place. These codewords provide exactly q^{r-1} distinct recovery sets of cardinality $q^{r-1} - 1$ for the i -th coordinate. Define

$$S = \{\text{Supp}(c) \setminus \{i\} \mid c \in S(r, q), c_i = 1\}.$$

Since any codeword with $\alpha \in \mathbb{F}_q^* \setminus \{1\}$ at the i -th place is an $\mathbb{F}_q$ -linear multiple of a codeword with $c_i = 1$, considering the i -th server systematic, we have

$$S \cup \{i\} = \mathcal{R}_i^{\min}.$$

Therefore, the cardinality of the set containing non-singleton recovery sets of the i -th information node is $|\mathcal{R}_i^{\min}| - 1 = q^{r-1}$.

Next, fix i , and consider all the codewords of $S(r, q)$ with the i -th coordinate zero; then the set of such codewords forms a subcode of $S(r, q)$, say D . The subcode D has q^{r-1} codewords, and

$$S(r, q) = D \bigcup_{\substack{c_i = \alpha \\ \alpha \in \mathbb{F}_q^*}} (c + D).$$

The number of codewords in D that have a zero in the j -th coordinate is q^{r-2} by Lemma 3. Consequently, the number of codewords in D with a non-zero j -th coordinate is $q^{r-1} - q^{r-2}$. Furthermore, due to the linearity of the code C , each additive coset of D also has $q^{r-1} - q^{r-2}$ codewords with a non-zero j -th coordinate.

Thus, the set of the support of the $q^{r-2}(q-1)$ codewords in the coset $c + D$ ($c_i = 1$), with the j -th coordinate non-zero, gives all the distinct recovery sets of the i -th data symbol containing the j -th node. Hence the result. $\square$

Remark 2. For a non-systematic generator matrix G of $\text{Ham}(r, q)$, recovery sets of data symbols can be of different sizes, but if there is a systematic server in G , then the recovery sets of the data symbol corresponding to this systematic column are determined in accordance with Theorem 1.

Example 3. For $r = 3$, we have the $[7, 4, 3]$ binary Hamming code whose generator and parity-check matrices are given in Example 2. Consider a message vector $u = (a, b, c, d)$. The associated codeword is $u \cdot G = (a+b+d, a+c+d, a, b+c+d, b, c, d)$. Then the recovery G system for this code is given as follows:

$$\begin{aligned} \mathcal{R}^{\min}(G) &= \{\mathcal{R}_a, \mathcal{R}_b, \mathcal{R}_c, \mathcal{R}_d\}, \quad \text{where} \\ \mathcal{R}_a &= \{(3), (1, 5, 7), (2, 6, 7), (2, 4, 5), (1, 4, 6)\}, \\ \mathcal{R}_b &= \{(5), (1, 3, 7), (4, 6, 7), (2, 3, 4), (1, 2, 6)\}, \\ \mathcal{R}_c &= \{(6), (1, 2, 5), (2, 3, 7), (4, 5, 7), (1, 4, 3)\}, \\ \mathcal{R}_d &= \{(7), (1, 3, 5), (2, 3, 6), (4, 5, 6), (1, 2, 4)\}, \end{aligned}$$

which satisfies the above theorem.

The next corollary is a consequence of Theorem 1 together with Lemma 4 and Remark 1. The upper bound is due to the distance, and the lower bound is due to the availability.

Corollary 1. In $\text{Ham}(r, q)$, the size $\delta(G)$ of the maximal achievable δ -simplex satisfies $\lceil \delta(G) \rceil \leq 3$. Moreover, in the case of systematic Hamming codes, we also have $\lfloor \delta(G) \rfloor \geq 2$.

It is proven in [20] that the upper bound $\delta(G) \leq 3$ is achievable for the systematic $\text{Ham}(r, 2)$ as the minimum weight codewords of $S(r, 2)$ form a 2-design. We also demonstrate that the upper bound on $\delta(G)$ is indeed achievable for systematic $\text{Ham}(r, 2)$ through optimal request splitting. The method of directing incoming requests toward the recovery sets is crucial for this achievability.

Lemma 6. In the systematic $\text{Ham}(r, q)$ with $r > 3$, an optimal request-splitting strategy is first directing all incoming requests to the systematic servers upon arrival, and then assigning the remaining requests to the recovery sets with the current least load.

Proof. Let $\lambda_i < \mu$ denote the incoming request rate for the i -th data object, and let t_i represent the load assigned to the corresponding systematic server. Consequently, $\lambda_i - t_i$ is the load redirected to another recovery set (i.e., towards $q^{r-1} - 1$ nodes present in the recovery set). By redistributing the load in this manner, the systematic server's load is reduced by $\lambda_i - t_i$, while each of the $q^{r-1} - 1$ coded nodes absorbs an additional $\lambda_i - t_i$ load.

Since $r > 3$, for $\text{Ham}(r, q)$, we have $n - k = r < q^{r-1} - 1$, implying that at least one server in any non-singleton recovery set of the i -th data file is a systematic server for a different j -th data file. The capacity of the systematic server serving the j -th file is exhausted by $\lambda_i - t_i$; therefore, this redistribution of $\lambda_i - t_i$ reduces the maximum service rate for the j -th file on its systematic node by $\lambda_i - t_i$. Consuming the non-singleton recovery set's capacity potentially reduces maximum service for other data files. Thus, it is optimal to prioritize sending incoming requests to systematic servers whenever they have available capacity.

Once the systematic server for a data object is saturated, requests are directed towards the recovery set with the least load. Consider two non-singleton recovery sets R_1 and R_2 of a data symbol. Let the loads on the nodes in set R_1 be $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_{q^{r-1}-1}$, and in the set R_2 be $\gamma'_1 \geq \gamma'_2 \geq \dots \geq \gamma'_{q^{r-1}-1}$, such that $\gamma_1 > \gamma'_1$. The set R_2 , having the smaller

load among the two, is capable of serving at most $\mu - \gamma'_1$ requests. This exceeds $\mu - \gamma_1$ requests that R_1 is capable of serving. In particular, it cannot be optimal to place the job on a more loaded recovery set while a less loaded one exists. Hence, the greedy “route-to-least-loaded” rule is optimal among recovery sets. If there are m recovery sets with the same load, say γ , then a maximum of $\left(\frac{1-\gamma}{q^{r-2}}\right)m$ request rate is distributed uniformly across these m recovery sets for maximum service.

Hence, the waterfilling algorithm ensures an optimal allocation of requests by first directing the requests to systematic servers and minimizing the impact on the overall service rate. $\square$

The following result indicates that the size of the maximal achievable simplex for systematic Ham(r, q) is $\left(1 + \frac{q}{q-1}\right)$. It has very recently come to our knowledge that this result can be obtained from Theorem 1 of [20], where the authors utilized a bound for the maximum achievable rate and demonstrated its achievability for the same coding scheme.

Theorem 2. *In the SRR of the systematic q -ary Hamming code for all $r \geq 3$, the maximal achievable demand for a single data object is $\left(1 + \frac{q}{q-1}\right)$, i.e., $\lambda_i^* = 1 + \frac{q}{q-1}$ for all $i \in [k]$.*

Proof. We show that the storage system with systematic Ham(r, q) coding scheme is able to serve the request rates $\left(1 + \frac{q}{q-1}\right)e_i$ for all $i \in [k]$. We use waterfilling algorithm to split the requests to the recovery sets. Since each server's capacity in our storage model is one, at most one request for an information symbol can be served using the corresponding systematic node.

From Theorem 1, we know that every node is in $q^{r-2}(q-1)$ recovery sets of a data object. The incoming request to download a data object is initially directed towards its systematic node. After the saturation of the systematic node, the optimal allocation is to send $\frac{1}{q^{r-2}(q-1)}$ of the request rate to every recovery set of the data object. This splitting strategy fulfills a maximum of $\frac{q^{r-1}}{q^{r-2}(q-1)}$ requests to download the required object without leaving any more service capacity in any node. Then we have

$$\left(1 + \frac{q}{q-1}, 0, \dots, 0\right), \left(0, 1 + \frac{q}{q-1}, 0, \dots, 0\right), \dots, \left(0, 0, \dots, 0, 1 + \frac{q}{q-1}\right) \in \Lambda(G)$$

Therefore, the maximum optimal allocation to serve a data object is $\lambda_i^* = 1 + \frac{q}{q-1}$ for any $i \in [k]$, when the request for the other data objects is zero. $\square$

The result indicates that the maximum demand does not depend on the parameter r but on the code alphabet. Since the service rate region is a convex polytope (see [16]), the next corollary follows.

Corollary 2. *The convex hull of the set $\{(3, 0, \dots, 0), (0, 3, 0, \dots, 0), \dots, (0, \dots, 0, 3)\} \subseteq \mathbb{R}^{2^r-1-r}$ lies in the SRR of the systematic binary Hamming code for all r , i.e., $\delta(G) = 3$.*

For a non-systematic matrix G with a systematic column for some $i \in [k]$, a parallel argument can be used to conclude that $\lambda_i^* = \frac{2q-1}{q-1}$ when $\lambda_j = 0$ for all $j \neq i \in [k]$ as discussed in Remark 2. Additionally, the allocation of requests to access such a data symbol can also be done using the waterfilling algorithm to get the maximum output.

A. Using Hypergraphs and matching to find SRR of systematic Ham($r, 2$)

In this section, we associate the binary Hamming codes with hypergraphs that illustrate the recovery structure of the code. In other words, the recovery G system of Ham($r, 2$) is seen as a hypergraph with $(2^r - 1)$ vertices. The recovery hypergraph is constructed with vertices corresponding to each column of the generator matrix G , where each recovery set of a data symbol forms a hyperedge.

Let the recovery hypergraph corresponding to the generator matrix G be denoted as Γ_G . Each hyperedge (recovery set) of Γ_G is labeled as the corresponding standard basis vector of the data symbol it recovers. A transversal and the transversal number of the hypergraph Γ_G play a crucial role in proving the following theorem.

Theorem 3. *Suppose G is a generator matrix for the code Ham($r, 2$). Let Γ_G be the associated recovery hypergraph and O_w be the number of odd weight columns in G . For a k -tuple $(\lambda_1, \lambda_2, \dots, \lambda_k) \in \Lambda(G)$, the following holds:*

- 1) $\sum_{i=1}^k \lambda_i \leq O_w$,
- 2) $\nu(\Gamma(G)) \geq \text{number of systematic columns in } G$,
- 3) if G is systematic, then

$$\sum_{i=1}^k \lambda_i \leq \begin{cases} 2^r - 1 - r, & \text{for } r > 3, \\ 5, & \text{for } r = 3, \end{cases} \quad (3)$$

and this bound is achievable.

Proof. All the standard basis vectors of $\mathbb{R}^k$ are of weight one. In order to obtain these vectors as a sum of other column vectors of G , there must be at least one column vector with an odd weight in each recovery set. This leads us to conclude that the intersection of any hyperedge of $\Gamma(G)$ with the set of odd-weighted column vectors of G is non-empty. Since the transversal

number is the smallest number of vertices needed to cover all hyperedges, we have $\tau(\Gamma(G)) \leq O_w$. Then, by Lemma 5, the inequality in (1) follows.

However, in the case of a systematic generator matrix, we have $O_w = n$, which does not lead to any concrete results.

The set of all systematic nodes in G forms a set of disjoint edges of the recovery hypergraph $\Gamma(G)$. Therefore, the matching number is greater than or equal to the number of systematic nodes in G .

If G is systematic, then we have each systematic node as an edge of the hypergraph, leading to the containment of all the systematic nodes in any vertex cover. According to Theorem 1, all the non-singleton recovery sets have cardinality $2^{r-1} - 1$, which is greater than r for $r > 3$. Therefore, there is no recovery set which contains only non-systematic columns. Consequently, it is impossible to form a recovery set that consists solely of non-systematic columns. This means that the set of all k systematic columns constitutes a vertex cover for the associated hypergraph, resulting in $\tau(\Gamma(G)) = k = 2^r - 1 - r$, for $r > 3$. To achieve the bound in (3), for each data symbol, a request is sent to the corresponding systematic node and served by it.

However, in the case of $r = 3$, the structure of the parity-check matrix allows for a recovery set made entirely of the three non-systematic nodes. As a result, any vertex cover includes one non-systematic column along with all systematic columns, leading to $\tau(\Gamma(G)) = k + 1 = 5$. Thus, the bound in (3) is achieved by serving one request for each data symbol by the corresponding systematic node, and serving one additional request through the recovery set with all the non-systematic nodes. This completes the proof. $\square$

Corollary 3. *For a non-systematic generator matrix G of $\text{Ham}(3, 2)$, the size of the maximal matching simplex for $\Lambda(G)$ is O_w .*

Remark 3. [14](Proposition 5) *The result 1 of Theorem 3 can also be concluded considering the hyperplane $\mathcal{H}$ of the projective space $PG_n(\mathbb{F}_2)$, defined by the equation*

$$x_1 + \dots + x_k = 0.$$

It is clear that none of the standard basis vectors e_i is in $\mathcal{H}$, as their weight is 1 (odd). Moreover, $|\mathcal{H}| = n - O_w$, and for any $\lambda \in \Lambda(G)$, it yields

$$\sum_{i=1}^k \lambda_i \leq |G \setminus \mathcal{H}| = O_w.$$

Example 4. *Let (a, b, c, d) be the data symbols and let G be a generator matrix of the $\text{Ham}(3, 2)$ code, given as follows:*

$$G = \begin{pmatrix} 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 1 & 0 & 0 \end{pmatrix}.$$




Fig. 1. Distributed storage system storing 4 files on 7 servers (nodes) using binary Hamming scheme (non-systematic generator matrix).

The odd weight columns of G are columns 3, 4, and 7. The recovery G system or the sets of edges of the recovery hypergraph $\Gamma(G)$ are presented below:

$$\mathcal{R}_a = \{(1, 7), (2, 4), (3, 5), (1, 4, 5, 6), (2, 5, 6, 7), (3, 4, 6, 7), (1, 2, 3, 6)\},$$

$$\mathcal{R}_b = \{(3, 4, 7), (1, 6, 7), (3, 5, 6), (2, 4, 6), (2, 5, 7), (1, 4, 5), (1, 2, 3)\},$$

$$\mathcal{R}_c = \{(7), (1, 2, 4), (1, 3, 5), (4, 5, 6), (2, 3, 6)\},$$

$$\mathcal{R}_d = \{(4), (1, 2, 7), (2, 3, 5), (5, 6, 7), (1, 3, 6)\}.$$

Since the set $\{3, 4, 7\}$ is a vertex cover of $\Gamma(G)$, the transversal number and the matching number are $\tau(\Gamma(G)) = 3$, and $\nu(\Gamma(G)) = 3$, respectively. It follows from Lemma 5 that $\sum_{i=1}^4 \lambda_i \leq 3$, and this bound is achieved by serving requests from edges $(3, 5)$, (7) , and (4) .

The requests are split using the greedy technique and are first directed towards the recovery sets of smallest cardinality with the least load for each of the data symbols. For data file a , the maximum demand served is $\lambda_a^{\max} = 3$. Similarly, for files c and d , we also have $\lambda_c^{\max} = \lambda_d^{\max} = 3$.

Since in $\mathcal{R}_b$ each node is contained in three edges, to get maximum service for file b , $1/3$ request is sent to each edge, recovering $\frac{7}{3}$ of file b . Note that $\delta(G) \neq 3$, since $\lambda_b^{\max}$ achievable is $\frac{7}{3}$.

Now we direct our attention to the service rate region of the systematic $\text{Ham}(r, 2)$, and generalize the second point of Theorem 3 for any subset of $[k]$, under the notion that the cumulative achievable service rate is less than the transversal




Fig. 2. The 3D-projection of service polytope considering $\lambda_d = 0$ (top), and service polytope when $\lambda_b = 0$ (bottom) in Example 4.

number of the associated recovery hypergraph. The recovery hypergraph Γ_G is defined by considering all data symbols and their recovery sets. Now, we consider partial hypergraph $\Gamma_{G_{\mathcal{I}}}$ of Γ_G with respect to a subset $\mathcal{I} \subseteq [k]$ such that the hyperedges labeled with e_j , where $j \in [k] \setminus \mathcal{I}$ are removed from $\Gamma(G)$. This formulation is analogous to considering that the demand/service rates for the data symbols indexed by nodes in $[k] \setminus \mathcal{I}$ are zero.

The fractional matching number of an s -uniform hypergraph with n vertices is less than or equal to $\frac{n}{s}$. We use this result as described in [14], and add $(2^{r-1} - 2)$ dummy vertices to each systematic node to represent the associated hyperedge in the recovery hypergraph (or partial hypergraph). As a result, the recovery hypergraph becomes $(2^{r-1} - 1)$ -uniform.

The following theorem is a consequence of the above discussion.

Theorem 4. Suppose G is the systematic generator matrix of $\text{Ham}(r, 2)$, and $(\lambda_1, \lambda_2, \dots, \lambda_k) \in \Lambda(G)$. Let $\mathcal{I} \subseteq [k = 2^r - 1 - r]$. Then

$$\sum_{i \in \mathcal{I}} \lambda_i \leq |\mathcal{I}| + 2 - \frac{|\mathcal{I}| - 1}{2^{r-1} - 1}.$$

Since the fractional part remains less than 1 for $|\mathcal{I}| > 1$, the size $\delta_{m_{\mathcal{I}}}$ of maximal matching simplex corresponding to the data symbols indexed by elements of $\mathcal{I}$ satisfies $\lceil \delta_{m_{\mathcal{I}}} \rceil \leq |\mathcal{I}| + 1$. We have shown later that this bound is attainable in numerous instances, and notably, it can be achieved without the nearest integer approximation.

The proposed lemma yields the number of recovery sets that contain a fixed number of non-systematic nodes of the systematic $\text{Ham}(r, 2)$. The recovery sets with cardinality $(2^{r-1} - 1)$ include both systematic and non-systematic nodes for $r > 3$.

Lemma 7. Suppose G is the systematic generator matrix of a binary Hamming code $\text{Ham}(r, 2)$, $r \geq 3$. If a recovery set contains t ($\leq r$) non-systematic nodes, then the total number of such recovery sets is $\binom{r}{t}(2^{r-1} - t)$.

Proof. Let R_a be a non-singleton recovery set with $n_1, n_2, \dots, n_t$ non-systematic nodes for some data symbol a . Since it contains t non-systematic nodes, the remaining $2^{r-1} - 1 - t$ nodes are systematic. In other words, the support of the sum of these t non-systematic columns has cardinality $2^{r-1} - t$. Let

$$\begin{aligned} S_p &= \text{Supp} \left(\sum (c_{n_1} + c_{n_2} + \dots + c_{n_t}) \right) \\ &= \{s_1, s_2, \dots, s_{2^{r-1}-t}\} \subset [k]. \end{aligned}$$

Among the nodes in S_p , any $(2^{r-1} - 1 - t)$ are the chosen systematic nodes to form a recovery set for the remaining data symbol. Thus, a recovery set of $s \in S_p \setminus \{s_{i_1}, s_{i_2}, \dots, s_{i_{(2^{r-1}-t-1)}}\}$ has the form

$$\{n_1, n_2, \dots, n_t, s_{i_1}, s_{i_2}, \dots, s_{i_{(2^{r-1}-t-1)}}\}. \quad (4)$$

Therefore, the total number of recovery sets with t non-systematic nodes is $\binom{r}{t}(2^{r-1} - t)$. $\square$

If the set $\mathcal{I}$ is chosen in such a way that in the associated partial hypergraph there is a hyperedge that does not contain any nodes from $\mathcal{I}$ or has only non-systematic nodes, then any vertex cover of the partial hypergraph cannot consist only of the nodes present in the set $\mathcal{I}$. For instance, consider the $2^{r-1} - r$ recovery sets of the form (4) with r non-systematic nodes and $2^{r-1} - r - 1$ systematic nodes. The sum of the columns of G corresponding to these r non-systematic nodes is a vector of weight $2^{r-1} - r$ in $\mathbb{F}_2^k$. Denote this vector by c_r . If for some $s \in \text{Supp}(c_r) \cap \mathcal{I}$, the set $\text{Supp}(c_r) \setminus \{s\}$ is not in $\mathcal{I}$, then there exists a recovery set say R_s of s which has r non-systematic nodes and $(2^{r-1} - r - 1)$ systematic nodes present in $\text{Supp}(c_r) \setminus \{s\}$. Therefore, a vertex cover of the partial hypergraph $\Gamma_{G_{\mathcal{I}}}$ also has an element from $[n] \setminus \mathcal{I}$ to cover the hyperedge associated with the recovery set R_s , along with the set $\mathcal{I}$.

The above argument raises a question about how large a vertex cover of the partial hypergraph $\Gamma_{G_{\mathcal{I}}}$ is, compared to the set $\mathcal{I}$. The following results aim to address this question.

Theorem 5. *Let G be a generator matrix of $\text{Ham}(r, 2)$ such that G contains systematic columns corresponding to two information nodes i and j . Then for a service request $\lambda \in \Lambda(G)$,*

$$\lambda_i + \lambda_j \leq 3.$$

Proof. Let $\mathcal{I} = \{i, j\}$. In the partial recovery hypergraph $\Gamma_{G_{\mathcal{I}}}$, there are two edges $\{i\}$ and $\{j\}$, each of which is assigned a weight 1. Then the capacity of these two nodes is exhausted. Therefore, now the hyperedges labeled e_i containing node j cannot be used for more request allotment. Similarly, the hyperedges labeled e_j containing node i cannot be used.

Let v_m represent the m -th column of the parity-check matrix of C corresponding to the matrix G . Note that each hyperedge labeled e_i containing the vertex j corresponds to a codeword $c = (c_1, c_2, \dots, c_n) \in C^\perp$ such that $c_i = c_j = 1$. Since all the columns of the parity-check matrix H of $\text{Ham}(r, 2)$ are distinct and linearly independent, for i and j , there is always a node $s \in [n]$ such that $v_i + v_j = v_s$. Consequently, in the partial hypergraph, there are 2^{r-2} hyperedges labeled e_i that contains s but not j , and there are also 2^{r-2} hyperedges labeled e_i that contains j but not node s . The same can be said for the hyperedges labeled e_j in the context of nodes i and s .

After assigning weight 1 to each node i and j , there are 2^{r-2} hyperedges e_i each containing node s , and 2^{r-2} hyperedges e_j each containing node s , along with all nodes having unused capacity. Since each available recovery set contains s , at most one request can be served in this system. A vertex cover for $\Gamma_{G_{\mathcal{I}}}$ is given by the set $\{i, j, s\}$. As a result, we conclude that $\lambda_i + \lambda_j \leq 3$. $\square$

Remark 4. *If the node s which is uniquely determined by the nodes i and j , is also a systematic node in G , then $\lambda_i + \lambda_j + \lambda_s \leq 3$.*

The argument of Theorem 5 has been extended to provide bounds on the sum of service rates for systematic $\text{Ham}(r, 2)$ corresponding to the data symbols indexed by any subset of $[k]$. In this theorem, we discuss the characteristics of the SRR for systematic binary Hamming codes and prove the achievability of the identified bounds.

Theorem 6. *Suppose G is the systematic generator matrix of a binary Hamming code $\text{Ham}(r, 2)$ where $r > 3$, and H is the corresponding parity-check matrix. Let $\mathcal{I}$ be a set consisting of information nodes (not all) with $|\mathcal{I}| \geq 2$, and v_j be the j -th column of H . If the demand for all information symbols not indexed by elements in $\mathcal{I}$ is zero, then the size $\delta_{m_{\mathcal{I}}}$ of the maximal matching simplex is given by:*

$$\sum_{i \in \mathcal{I}} \lambda_i \leq \delta_{m_{\mathcal{I}}} = \begin{cases} |\mathcal{I}|, & \text{if } \sum_{j \in \mathcal{I}} v_j = 0 \in \mathbb{F}_2^r, \\ |\mathcal{I}| + 1, & \text{otherwise.} \end{cases}$$

Additionally, if the service rate for all the information symbols is non-zero, i.e., $\mathcal{I} = [k]$, then $\delta_{m_{\mathcal{I}}} = k$.

Proof. The columns of H are all distinct. Also, since every non-zero vector of $\mathbb{F}_2^r$ is a column of H , for any subset $\mathcal{I}$ of information nodes, the sum $\sum_{j \in \mathcal{I}} v_j$ is again a column of H or the all-zero vector.

Consider $\sum_{j \in \mathcal{I}} v_j = 0$. This indicates that in any recovery set of an i -th data symbol, $i \in \mathcal{I}$, there is always at least one another node $s \in \mathcal{I}$ present in that recovery set. Moreover, all the hyperedges of the partial hypergraph $\Gamma_{G_{\mathcal{I}}}$ always contain at least one element of $\mathcal{I}$. This implies that the set $\mathcal{I}$ forms a vertex cover of $\Gamma_{G_{\mathcal{I}}}$. There are $|\mathcal{I}|$ distinct hyperedges in $\Gamma_{G_{\mathcal{I}}}$ which has the form $\{i\}$ with $i \in \mathcal{I}$, hence this set is the smallest vertex cover of $\Gamma_{G_{\mathcal{I}}}$ and $\tau(\Gamma_{G_{\mathcal{I}}}) = |\mathcal{I}| = \nu(\Gamma_{G_{\mathcal{I}}})$.

Now, suppose $\sum_{j \in \mathcal{I}} v_j \neq 0$. Then there exists a column of H , say v_s , such that

$$\sum_{j \in \mathcal{I}} v_j = v_s. \quad (5)$$

If a codeword $c = (c_1, c_2, \dots, c_n) \in C^\perp$ with $i \in \text{Supp}(c)$ and $s \notin \text{Supp}(c)$ for some $i \in \mathcal{I}$, then $|\text{Supp}(c) \cap \mathcal{I}| \geq 2$. Thus, the hyperedges associated with such codewords of the dual code are covered by $\mathcal{I}$.

Now we consider those hyperedges of $\Gamma_{G_{\mathcal{I}}}$ which do not contain any vertex from the set $\mathcal{I}$. These hyperedges correspond to the codewords $c \in C^\perp$ such that $|\text{Supp}(c) \cap \mathcal{I}| = 1$. Let $h_{c_a} \in \mathcal{R}_a$ be the hyperedge of $\Gamma_{G_{\mathcal{I}}}$ corresponding to a codeword $c \in C^\perp$ for $\text{Supp}(c) \cap \mathcal{I} = \{a\}$.

- **Case 1:** If the codeword c is a row of the parity-check matrix, then considering Equation (5), we have $c_s = 1$. Thus, $s \in h_{c_a}$ and $\mathcal{I} \cup \{s\}$ is the smallest transversal of $\Gamma_{G_{\mathcal{I}}}$.
- **Case 2:** If the codeword $c \in C^\perp$ is a linear combination of rows $\{x_1, x_2, \dots, x_t\}$, $1 \leq t \leq r$, of the parity-check matrix H . Then, from the assumption $\text{Supp}(c) \cap \mathcal{I} = \{a\}$, we have

$$x_{1a} + x_{2a} + \dots + x_{ta} = 1, \quad (6)$$

$$\sum_{i=1}^t x_{ij} = 0, \text{ for all } j \in \mathcal{I} \setminus \{a\}, \quad (7)$$

where x_{ij} is the j -th coordinate of the row vector x_i . After adding equations (6) and (7) we get

$$\sum_{j \in \mathcal{I}} \sum_{i=1}^t x_{ij} = \sum_{i=1}^t \sum_{j \in \mathcal{I}} x_{ij} = 1. \quad (8)$$

Now, applying the relation given in Equation (5) to Equation (8), we get

$$\sum_{i=1}^t \sum_{j \in \mathcal{I}} x_{ij} = \sum_{i=1}^t x_{is} = c_s = 1.$$

This indicates that the node s is in the hyperedge h_{c_s} . In this case as well, we conclude that the smallest vertex cover of the partial hypergraph $\Gamma_{G_{\mathcal{I}}}$ is formed by the set $\mathcal{I} \cup \{s\}$.

Note that the matching number of $\Gamma_{G_{\mathcal{I}}}$ in both the cases is $|\mathcal{I}| + 1$. Therefore, the size $\delta_{m_{\mathcal{I}}}$ of the maximal matching simplex is $|\mathcal{I}| + 1$.

For the systematic $\text{Ham}(r, 2)$ with $r > 3$ and the set $\mathcal{I} = [k]$, it is evident that all singleton hyperedges are included in a vertex cover. Also, there are no hyperedges that do not contain a systematic server; the set $[k]$ forms a complete smallest vertex cover of $\Gamma(G)$. Therefore, we conclude that $\delta_{m_{\mathcal{I}}} = k$. $\square$

Since any two columns of the parity-check matrix of $\text{Ham}(r, 2)$ are linearly independent, Theorem 5 is a special case of Theorem 6. The fractional matching number is equal to the transversal number for $\Gamma(G)$ of the systematic $\text{Ham}(r, 2)$, therefore, if the cumulative service rate exceeds the integral value $\tau(\Gamma(G))$, then the system is not able to complete the requested demand.

Corollary 4. *Let $r > 3$ be a positive integer, and consider a k -tuple $(\lambda_1, \lambda_2, \dots, \lambda_k)$, where $\lambda_i = 1$ for all $i \in [k]$ except one index $j \in [k]$, for which $\lambda_j > 1$. Then, the k -tuple $(\lambda_1, \lambda_2, \dots, \lambda_k)$ does not belong to the service rate region of systematic $\text{Ham}(r, 2)$.*

The next theorem summarizes the discussion presented above.

Theorem 7. *Let G denote the systematic generator matrix and H be the corresponding parity-check matrix for $\text{Ham}(r, 2)$ with the parameters $[2^r - 1, k = 2^r - 1 - r, d = 3]$. Consider a subset $\mathcal{I}$ of the set of systematic nodes in G such that $2 < |\mathcal{I}| \neq k$, and the demand for all data symbols indexed by elements in $[k] \setminus \mathcal{I}$ is zero. Let v_j denote the j -th column of H . Then the service rate region for the systematic $\text{Ham}(r, 2)$ is given by:*

$$\begin{cases} \lambda_i \geq 0, & \text{for all } i \in [k], \\ \lambda_i + \lambda_j \leq 3 & \text{for all } i, j \in [k], \\ \sum_{i \in \mathcal{I}} \lambda_i \leq \delta_{m_{\mathcal{I}}} = \begin{cases} |\mathcal{I}|, & \text{if } \sum_{j \in \mathcal{I}} v_j = 0 \in \mathbb{F}_2^r, \\ |\mathcal{I}| + 1, & \text{otherwise,} \end{cases} \\ \sum_{i=1}^k \lambda_i \leq \begin{cases} k, & \text{for } r > 3, \\ 5, & \text{for } r = 3. \end{cases} \end{cases}$$

Example 5. *Let G be the matrix as described in Example 3. Then the service rate region of systematic $\text{Ham}(3, 2)$ is the feasible region of the linear programming problem with the following constraints:*

$$\begin{aligned} \lambda_a + \lambda_b + \lambda_c + \lambda_d &\leq 5, \\ \lambda_a + \lambda_b + \lambda_c &\leq 3, \\ \lambda_i + \lambda_j + \lambda_d &\leq 4, \quad i \neq j \in \{a, b, c\}, \\ \lambda_i + \lambda_j &\leq 3, \quad i \neq j \in \{a, b, c, d\}, \\ 0 &\leq \lambda_i \leq 3, \quad \forall i \in \{a, b, c, d\}. \end{aligned}$$

Two 3-dimensional cross-sections of the corresponding 4-dimensional service polytope are illustrated in Figure 3.

This example demonstrates that the maximum service rate for a data symbol can differ based on the choice of the generator matrix, given that not all other data symbols have zero demand (service rate). Although the volume of the service rate region remains constant, the orientation of the convex service polytope may vary depending on the selected generator matrix. When the demand for a particular data symbol exceeds that of others, we select the generator matrix accordingly for the same code. For instance, let a, b, c , and d represent files stored in a coded storage system that employs a binary Hamming coding scheme. If the demand for file d is greater than the demand for the other files, then the generator matrix presented in Example 2 is appropriate, as $(1, 1, 1, 2) \in \Lambda(G)$.

In contrast, if more users are interested in accessing file a than the other files, a more suitable generator matrix would be




Fig. 3. The service polytope considering (i) the index set $\mathcal{I} = \{a, b, c\}$ (top), and (ii) when the index set is $\mathcal{I} = \{a, b, d\}$ (bottom).

$$G' = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

This matrix allows for the recovery of file a from the recovery set consisting of non-systematic nodes only, yielding $(2, 1, 1, 1) \in \Lambda(G')$.

Consider a situation where the cumulative demand for objects a , b , and c is high. In this case, the preferred generator matrix would be G' as we have $\lambda_a + \lambda_b + \lambda_c \leq 4$. However, as for generator matrix G , we have $\lambda_a + \lambda_b + \lambda_c \leq 3$.

Remark 5. Let $\mathcal{I} \subseteq [k]$. Based on the demand for the data symbols indexed by elements of the set $\mathcal{I}$, we can pick the generator matrix that maximizes the cumulative service rate by finalizing the associated parity-check matrix. The alternative approach is to rearrange the order of the data symbols in the message vector to selectively enhance the overall service rate for the set $\mathcal{I}$.

We can determine the number of equations that limit the sum of service rates for a set of three data objects to 3. This is useful for approximating the SRR when the demand for data files changes continuously, while the demand for any three specific files remains non-zero; that is, when $|\mathcal{I}| = 3$.

Counting the number of inequalities of the form $\lambda_i + \lambda_j + \lambda_k \leq 3$

Consider v_i to be a column of the parity-check matrix corresponding to the systematic server of the i -th data symbol in G . In other words, $v_i \in \mathbb{F}_2^r$, such that $\text{wt}(v_i) \geq 2$. There are $\binom{2^r - 1 - r}{2}$ inequalities of the form $\lambda_i + \lambda_j \leq 3$ for systematic $\text{Ham}(r, 2)$. However, not all sums $v_i + v_j$ correspond to a data symbol; some may relate to a parity-check.

Our goal is to determine the number of combinations such that $v_i + v_j = v_s$, where each column corresponds to distinct data symbols at coordinate places $i, j, s \in [k]$ respectively. We will denote this number by M_3 .

We first determine the number of combinations such that $v_i + v_j = e$, where e is a standard basis vector in $\mathbb{F}_2^r$. Suppose $v_i = \{x_1, x_2, \dots, x_r\}$ with $\text{Supp}(v_i) = \{i_1, i_2, \dots, i_t\}$, where $2 \leq t \leq r$. For $v_i + v_j = e$, we require

$$\text{wt}(v_j) = t - 1 \text{ or } t + 1,$$

as the vectors v_i and v_j differ in exactly one position.

- There are t vectors with weight $t - 1$ and support $\{i_1, i_2, \dots, i_t\} \setminus \{i_w\}$, where $w \in [t]$ respectively.
- There are $r - t$ vectors with weight $t + 1$ and support $\{i_1, i_2, \dots, i_t\} \cup \{i_u\}$, where $u \in [r] \setminus \text{Supp}(v_i)$.

Note that for $t = 2$, there are no vectors v_j of weight $t - 1$ that correspond to a systematic server in G . Therefore, we only consider vectors with weight $t + 1$. Similarly, any combination that involves vectors of weight t and $t - 1$ has already been accounted for when t was set to $(t - 1)$. Consequently, we conclude that there are $r - t$ vectors v_j for each vector v_i of weight t such that $v_i + v_j = e$.

In $\mathbb{F}_2^r$, the total number of vectors of weight t is given by $\binom{r}{t}$. This leads us to conclude that the total number of combinations of the form $v_i + v_j = e$ is equal to

$$\sum_{t=2}^{r-1} \binom{r}{t} (r - t).$$

Therefore, the total number of combinations of the type $v_i + v_j = v_s$ is

$$\begin{aligned} M_3 &= \frac{\binom{2^r-1-r}{2} - \sum_{t=2}^{r-1} \binom{r}{t}(r-t)}{3}, \\ &= \frac{\binom{2^r-1-r}{2} - r2^{r-1} + r^2}{3}. \end{aligned}$$

IV. CONCLUSION

In this paper, we studied the recovery G system of systematic $\text{Ham}(r, q)$, where G is a generator matrix of $\text{Ham}(r, q)$. Through optimal request splitting, it was demonstrated that the size of the maximal achievable simplex in SRR of systematic $\text{Ham}(r, q)$ is $\frac{2q-1}{q-1}$, which is independent of r . By considering the recovery sets as hyperedges and using elementary properties of the dual code, we identified limits on the sum of service rates of data symbols contained in a given subset of all data symbols of systematic binary Hamming codes. We showed that in systematic $\text{Ham}(r, 2)$, the cumulative service rate cannot surpass the total number of data symbols $2^r - 1 - r$ for $r > 3$, and for any two data symbols, i -th and j -th file, the combined service rate $\lambda_i + \lambda_j$ is less than 3. For future work, further analysis of the hypergraph and projective space context could lead to a complete characterization of the SRR of Hamming codes. This concept may be extended to codes with higher locality, such as locally recoverable codes and maximally recoverable codes. Investigating the link between the SRR of the code and the dual code's structure also presents an interesting research direction.

ACKNOWLEDGMENTS

The authors express their gratitude to the University Grants Commission (UGC), Government of India, for the financial support received under the UGC NET-JRF scheme.

REFERENCES

- [1] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin. "On the locality of codeword symbols," *IEEE Trans. Inf. Theory*, vol. 58, no. 11, pp. 6925–6934, Nov. 2012.
- [2] F. Kazemi, E. Karimi, E. Soljanin, and A. Sprintson. "A combinatorial view of the service rates of codes problem, its equivalence to fractional matching and its connection with batch codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2020, pp. 646–651.
- [3] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran. "Network coding for distributed storage systems," *IEEE Trans. Inf. Theory*, vol. 56, no. 9, pp. 4539–4551, Sep. 2010.
- [4] A. G. Dimakis, K. Ramchandran, Y. Wu, and C. Suh. "A survey on network codes for distributed storage," in *Proc. IEEE*, vol. 99, no. 3, pp. 476–489, Mar. 2011.
- [5] C. Huang, M. Chen, and J. Li. "Pyramid codes: Flexible schemes to trade space for access efficiency in reliable data storage systems," *ACM Trans. Storage*, vol. 9, no. 1, pp. 1–28, 2013.
- [6] G. Joshi, Y. Liu, and E. Soljanin. "Coding for fast content download," in *Proc. 50th Annu. Allerton Conf. Commun., Control, Comput.*, 2012, pp. 326–333.
- [7] G. Liang and U. C. Kozat. "Fast cloud: Pushing the envelope on delay performance of cloud storage with coding," *IEEE/ACM Trans. Netw.*, vol. 22, no. 6, pp. 2012–2025, Dec. 2013.
- [8] K. Gardner, S. Zbarsky, S. Doroudi, M. Harchol-Balter, and E. Hyttia. "Reducing latency via redundant requests: Exact analysis," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 43, no. 1, pp. 347–360, 2015.
- [9] M. Noori, E. Soljanin, and M. Ardakani. "On storage allocation for maximum service rate in distributed storage systems," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2016, pp. 240–244.
- [10] M. Aktaş, S. E. Anderson, A. Johnston, G. Joshi, S. Kadhe, G. L. Matthews, C. Mayer, and E. Soljanin. "On the service capacity region of accessing erasure coded content," in *Proc. 55th Annu. Allerton Conf. Commun., Control, Comput.*, 2017, pp. 17–24.
- [11] S. E. Anderson, A. Johnston, G. Joshi, G. L. Matthews, C. Mayer, and E. Soljanin. "Service rate region of content access from erasure coded storage," in *Proc. IEEE Inf. Theory Workshop (ITW)*, 2018, pp. 1–5.
- [12] F. Kazemi, S. Kurz, and E. Soljanin. "A geometric view of the service rates of codes problem and its application to the service rate of the first order Reed–Muller codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2020, pp. 66–71.
- [13] F. Kazemi, S. Kurz, E. Soljanin, and A. Sprintson. "Efficient storage schemes for desired service rate regions," in *Proc. IEEE Inf. Theory Workshop (ITW)*, 2021, pp. 1–5.
- [14] M. Aktaş, G. Joshi, S. Kadhe, F. Kazemi, and E. Soljanin. "Service rate region: A new aspect of coded distributed system design," *IEEE Trans. Inf. Theory*, vol. 67, no. 12, pp. 7940–7963, Dec. 2021.
- [15] G. N. Alfarano, A. B. Kılıç, and A. Ravagnani. "Service aspects of LRC and batch codes," *IEEE BITS Inf. Theory Mag.*, 2023.
- [16] G. N. Alfarano, A. B. Kılıç, A. Ravagnani, and E. Soljanin. "The service rate region polytope," *SIAM J. Appl. Algebra Geom.*, vol. 8, no. 3, pp. 553–582, 2024.
- [17] A. B. Kılıç, A. Ravagnani, and E. Soljanin. "On the parameters of codes for data access," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2024, pp. 819–824.
- [18] G. N. Alfarano, A. Ravagnani, and E. Soljanin. "Dual-code bounds on multiple concurrent (local) data recovery," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2022, pp. 2613–2618.
- [19] H. Ly and E. Soljanin. "Service rate regions of MDS codes and fractional matchings in quasi-uniform hypergraphs," 2025. [Online]. Available: <https://arxiv.org/abs/2504.17244>
- [20] H. Ly and E. Soljanin. "Maximal achievable service rates of codes and connections to combinatorial designs," 2025. [Online]. Available: <https://arxiv.org/abs/2506.16983>
- [21] H. Ly, E. Soljanin, and V. Lalitha. "On the service rate region of Reed–Muller codes," 2025. [Online]. Available: <https://arxiv.org/abs/2501.13105>
- [22] S. Ling and C. Xing. *Coding Theory: A First Course*, Cambridge Univ. Press, 2004.
- [23] W. C. Huffman and V. Pless. *Fundamentals of Error-Correcting Codes*, Cambridge Univ. Press, 2003.