Technische Universität München
Professur für Coding and Cryptography
Prof. Dr.-Ing. Antonia Wachter-Zeh

Master's Thesis

Secure Coding for Distributed Data Storage with Sum-Rank Metric Codes

Vorgelegt von:
Tim Janz

München, April 2023

Betreut von:
Dr. Rawad Bitar [♭], Hedongliang Liu [♭], Prof. Frank R. Kschischang [♯]
[♭] Chair for Coding and Cryptography, Technical University of Munich
[♯] The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, University of Toronto

Master's Thesis an der
Professur für Coding and Cryptography (COD)
der Technischen Universität München (TUM)
Titel: Secure Coding for Distributed Data Storage with Sum-Rank Metric Codes
Autor: Tim Janz

Tim Janz
tim.janz@tum.de

Ich versichere hiermit wahrheitsgemäß, die Arbeit bis auf die dem Aufgabensteller bereits bekannte Hilfe selbstständig angefertigt, alle benutzten Hilfsmittel vollständig und genau angegeben und alles kenntlich gemacht zu haben, was aus Arbeiten anderer unverändert oder mit Abänderung entnommen wurde.

Antibes, 16.04.2023

Ort, Datum (Tim Janz)

Contents

Abstract

Distributed storage systems (DSSs) using maximally recoverable locally repairable codes (MR-LRCs) are considered. A new global repair scheme for MR-LRCs based on linearized Reed-Solomon codes that uses so-called local polynomials to distribute the repair process is suggested. Two different schemes that use local polynomials are introduced, namely direct global repair and forwarded global repair. The secrecy capacity of a system, i.e., the number of information symbols that can securely be stored, using direct and forwarded global repair given an eavesdropper is determined.

1 Introduction

The increasing demand of cloud-based applications such as cloud computing, video streaming and cloud storage increases the required storage capacity and stresses the importance of efficient storage solutions. In addition, DSS should be protected against the possibility of data loss due to disk failure. Since disk failures are very likely in large servers, erasure-correcting codes are used to prevent data loss. The most simple erasure-correcting code is the replication of data. However, this has the drawback of a large storage overhead. Therefore, more complex coding schemes with higher code rates are used to protect the data against loss. [RBS+22] gives a good overview of codes for distributed storage.

One popular choice of codes that have already been implemented by Facebook [SAP+13] and Microsoft [HSX+12] are locally repairable codes (LRCs). In case of a node failure, only a small number of nodes, r, in the same local server rack are contacted for the repair; r is called the code locality. An interesting class of locally repairable codes (LRCs) are maximally recoverable locally repairable codes (MR-LRCs) [GHJY14], also called partial MDS (or PMDS) codes. They can correct any erasure pattern that is information-theoretically correctable given the parameters of the code. The maximally recoverable property can be achieved by a two step encoding procedure with an outer code of size n and dimension k and a local encoding as for LRCs. After the outer encoding, the codeword is split into g parts, which are distributed among g groups. The g parts are then further encoded with a local code such that the global codeword, i.e., the concatenation of all local parts, has size N. If too many erasures occur in one local group such that the local code cannot correct them, they can be corrected using the outer code. Such a repair is called global repair. An MR-LRC with sub-exponential field size was introduced by Martínez-Peñas and Kschischang in [MPK19], which is the main MR-LRC construction considered in this work. Linearized Reed-Solomon codes, which are based on polynomials with a non-commutative product, so-called skew polynomials, are used for this construction.

Another aspect of DSSs is secrecy. There are two different types of attacks that can threat the security and secrecy of a system: active and passive attacks. Active attacks include maliciously reconfiguring the system, modifying packets or injecting data, while passive attacks only involve eavesdropping on the stored or transmitted data. This work

only considers the latter and investigates the consequences of the maximal recoverability property on the ability to store data in the presence of an eavesdropper, similar to [RKSV14].

The main problem of MR-LRCs in the presence of an eavesdropper is that the outer code, which is used for global repairs, is a maximum distance separable (MDS) code, hence the name partial MDS. If a global repair is performed in a node that simply downloads as many symbols as needed for the repair, in this case any k out of n symbols, an eavesdropper that can observe the downloads of this node would gain knowledge about all stored symbols. This would mean that no information can be stored securely, given the described repair process. To solve this issue, we introduce a distributed global repair process that allows MR-LRCs to have a nonzero secrecy capacity, i.e., the number of symbols that can be stored securely is in general greater than zero, in the presence of an eavesdropper. The distributed global repair uses so-called *local polynomials* generated by each group, whose sum recovers the global encoding polynomial. To perform a global repair, each group calculates the evaluation of its local polynomial at the code locator of the failed node and sends it to the failed node. Sending the evaluation of the local polynomials to the failed node can be realized in a direct or forwarded way. In the case of *direct global repair*, each group sends its local polynomial evaluation directly to the group with the failed node, whereas in the case of *forwarded global repair*, the evaluations are forwarded along a forwarding list, at the end of which is the group with the failed node. At each group, the received evaluation is added to the contribution of the group. Therefore, the failed node only receives one symbol, the sum of all local polynomial evaluations, instead of each evaluation separately. The two schemes, direct and forwarded global repair, are compared and their secrecy capacities are derived. The forwarded global repair achieves in general larger secrecy capacities since each group receives at most one symbol.

The structure of the thesis is as follows. The first part, Chapter 2, explains the considered distributed storage system (DSS) and gives an overview of the assumptions that are made about the system. Moreover, it defines the secrecy threat that is considered throughout the work. Chapter 3, Preliminaries, reviews all concepts to which the other parts refer. It reviews secret sharing, locally repairable codes (LRCs) and sum-rank metric codes. Furthermore, linearized Reed-Solomon codes (LRSCs) are summarized and two well known concepts are applied and adapted to skew polynomials, namely secret sharing and Lagrange polynomials. At the end of the Preliminaries chapter, an important information-theoretic concept is stated. The third part, Chapter 4, summarizes important constructions of a secure locally repairable coding scheme and MR-LRC coding schemes. In the fourth part, Chapter 5, a distributed global repair for MR-LRCs is suggested. The

repair uses local polynomials to distribute the recovery to each group of a system. Two different schemes implementing distributed global repair are considered and discussed. A secrecy capacity for the schemes is derived and a construction which achieves capacity is given. In the last part, Chapter 7, all the results are summarized and an outlook for possible further investigations is given. Appendix A summarizes essential properties of skew polynomials.

Notation

Throughout this work, the following notation is used. The set of $m \times n$ matrices with entries in the field $\mathbb{K}$ is denoted as $\mathbb{K}^{m\times n}$. Matrices are written as uppercase bold letters, e.g., $\mathbf{A} \in \mathbb{K}^{m\times n}$, with its transpose $\mathbf{A}^\mathsf{T} \in \mathbb{K}^{n\times m}$. The rank of a matrix $\mathbf{M}$ is written as $\mathrm{rank}(\mathbf{M})$. Vectors are written as lowercase bold letters, e.g., $\mathbf{a} = (a_1, a_2, \ldots, a_n) \in \mathbb{K}^n$. The set $\mathbb{K}^*$ denotes the field without zero, i.e., $\mathbb{K}^* = \mathbb{K}\backslash\{0\}$.

Only in the motivation of Section 3.3 can the field $\mathbb{K}$ be any field; otherwise the considered field $\mathbb{K}$ is a finite extension field $\mathbb{F}_{q^m}$ of degree m with base field $\mathbb{F}_q$, where q is a prime power. When the size is not relevant, we simply write $\mathbb{F}$. In some sections, the field $\mathbb{F}_{2^2} = \mathbb{F}_4$ is used to provide examples. The primitive element generating the field $\mathbb{F}_{2^2}$ is denoted as ω, i.e.,

$$\mathbb{F}_4 = \{0, 1, \omega, \bar{\omega}\} \text{ with } \omega^2 = \bar{\omega} = 1 + \omega. \tag{1.1}$$

Sets are written in calligraphic font or as uppercase Greek letters, e.g., $\mathcal{M}$ or Ω with cardinality denoted as $|\mathcal{M}|$ and $|\Omega|$. The set of natural numbers excluding zero is denoted as $\mathbb{N}$ and $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. Let $[n] := \{m \in \mathbb{N} \mid m \leq n\} = \{1, 2, \ldots, n\}$ and let $\mathcal{I}_t \subseteq [n]$ denote a set with cardinality t. For example, $\mathcal{I}_2 = \{1, 3\} \subseteq [4]$ is an example of a set with $t = 2$. For a vector $\mathbf{s}$ denote $\mathbf{s}_{\mathcal{I}_t} := (s_i \mid i \in \mathcal{I}_t)$.

A code $\mathcal{C}$ is a nonempty subset of $\mathbb{F}_{q^m}^n$, i.e., $\mathcal{C} \subseteq \mathbb{F}_{q^m}^n$, where each codeword $\mathbf{c} \in \mathcal{C}$ is a vector of length n with components from $\mathbb{F}_{q^m}$.

For a nonempty set $\mathcal{R} \subseteq [n]$ and a vector $\mathbf{x} \in \mathbb{F}_{q^m}^n$, define a projection map $\pi_\mathcal{R} : \mathbb{F}_{q^m}^n \longrightarrow \mathbb{F}_{q^m}^{|\mathcal{R}|}$ taking $(x_1, x_2, \ldots, x_n) = (x_i \mid i \in [n])$ to the vector $(x_j \mid j \in \mathcal{R})$. Then, $\mathcal{C}|_\mathcal{R} = \{\pi_\mathcal{R}(c) \mid c \in \mathcal{C}\}$ is the punctured code with respect to $\mathcal{R}$.

Discrete random variables are written in sans serif font; for example A or R. The probability of an event is denoted by $\Pr[\bullet]$; e.g., $\Pr[\mathsf{A} = a]$ denotes the probability that the random variable A takes the value a. For ease of notation, $\mathrm{P}_\mathsf{A}(a)$ is used interchangeably for $\Pr[\mathsf{A} = a]$. The support of a probability distribution, $\mathrm{supp}(\mathrm{P}_\mathsf{X})$, is the set of x such that $\mathrm{P}_\mathsf{X}(x) > 0$. The conditional probability of the event $\mathsf{A} = a$ given $\mathsf{B} = b$ is written as $\Pr[\mathsf{A} = a \mid \mathsf{B} = b] = \mathrm{P}_{\mathsf{A}|\mathsf{B}}(a \mid b)$. The entropy of a discrete random variable X with range

$\mathcal{X}$ is defined as

$$\mathrm{H}(\mathsf{X}) = \sum_{a \in \mathrm{supp}(\mathrm{P}_\mathsf{X})} - \mathrm{P}_\mathsf{X}(a) \log_{|\mathcal{X}|} \mathrm{P}_\mathsf{X}(a).$$

The entropy $\mathrm{H}(\mathsf{X})$ is bounded by $0 \leq \mathrm{H}(\mathsf{X}) \leq 1$. The mutual information between two discrete random variables X and Y is written as $\mathrm{I}(\mathsf{X};\mathsf{Y})$. A summary of used rules and (in-)equalities of mutual information and entropy is given in Section B.1.

2 System Model

In this chapter, the underlying structure for which the codes should be designed is summarized. DSSs are explained in Section 2.1. They consist of g racks, which are also referred to as groups. Each rack has a fixed number of nodes that can store data. The main goal of this thesis is to analyze the secrecy of the considered DSS. Therefore, another question naturally arises: what kind of secrecy threat is considered? This question is answered in Section 2.2, which describes the considered eavesdropper model. The eavesdropper is assumed to be a passive eavesdropper that can only read the stored symbols of l_1 nodes and in addition read the symbols of l_2 groups and the symbols downloaded for the repair of these l_2 groups.

2.1 Distributed Storage System Model

As the name suggests, a DSS is a connected group of storage units. The largest unit of such a system is a server rack. It has multiple rack slots in each of which there is a storage node. The nodes are connected to a Top-of-Rack switch [TCS19] with a rack processing unit (RPU). The RPU is responsible for local computations and manages the whole rack. All the racks in a system are connected to aggregation-layer modules which are responsible for switching and provide backend functions such as Layer 2 domain definitions, load balancing and firewall features [Cis07]. Each storage node in a rack consists of multiple disks that are managed by a processing unit. The processing unit of each node sends a heartbeat request to its disks and checks whether they are still available. In case of a working disk, the disk controller returns an acknowledgement to the processing unit of the node which marks the disk as available. If no acknowledgement from the disk controller follows, the corresponding disk is marked as dead or off [Bor08, Cep]. Disks that are marked as off for a longer time, i.e., disks that have not sent a heartbeat acknowledgement several times, can easily be identified and fixed or replaced by the administrator of the system.

In this work, the data storage topology is considered in a simplified form with only two hierarchical layers as shown in Figure 2.1. There are g racks with n_i storage nodes in each rack $i \in [g]$. The storage nodes can be seen as black boxes which are able to store

data. It is assumed that all the nodes are connected to a rack processing unit which is the connecting device to other racks and responsible for local computations. The global switch has a global processing unit, which is coordinating the whole system. Each rack can be seen as an independent storage system but the global processing unit provides useful features such as a map of the stored data or metadata of the system. This allows, for example, a direct switching to the local unit after receiving a request to access specific data without a broadcast request to every rack. In practical examples, all the racks have the same number of nodes, i.e., $n_i = n$ for all $i \in [g]$.

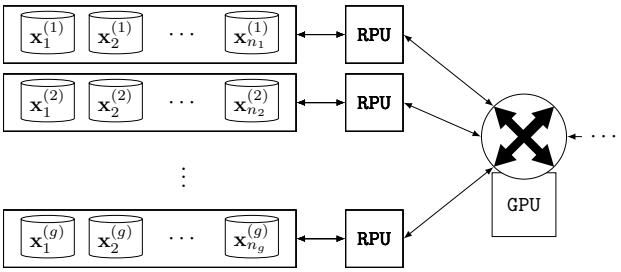

Figure 2.1: Distributed storage system model with g racks. The j-th node in the i-th rack is denoted by $\mathbf{x}_j^{(i)}$. Each rack has a processing unit, RPU, which is responsible for the communication between nodes and local computations. The processing unit is connected to a switch which allows communication with other racks. The model is adapted from [TCS19].

A unified notation for DSS is used throughout this work. The nodes of each rack are depicted by a square box. Connected square boxes represent a rack. In each square box, one symbol $x_j^{(i)} \in \mathbb{F}_{q^m}$ is representative for all symbols stored on the j-th node in the i-th group.

Example 2.1. Consider a DSS with three racks and four nodes in each of the racks, i.e., $g = 3$ and $n_i = n = 4$ for $i \in [g]$. The system is illustrated in Figure 2.2.

In [ACRV14] the authors report a bandwidth between nodes of different racks that was by a factor of 5 smaller than between nodes of the same rack. If the transmission time is now seen as an indicator for the communication costs, the communication costs between racks are higher than within a rack. The estimation of Ahmad et al. in [ACRV14] is that the factor 5 is on the lower end of the discrepancy between intra-rack and inter-rack bandwidth ranging from 5 to 20. Besides the measured link capacity in practical systems, the higher communication costs can also be motivated by the fact that the switch connecting the processing units with each other is more likely to be the bottleneck of the system [IPC+09]. This motivates to take the different communication costs into account

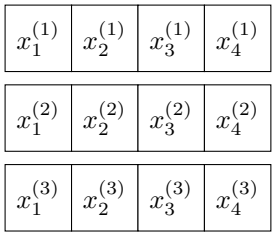

Figure 2.2: Illustration of a DSS with 3 racks and 4 nodes in each rack. Each rack consists of connected square boxes depicting the nodes. The symbol $x_j^{(i)}$ is representative for all symbols stored on the j-th node in the i-th group.

for the design of codes for DSSs. An example are LRCs which recover failed nodes locally, i.e., only nodes that are in the same rack are used for the repair. LRCs are discussed in Section 3.2.

Remark: It is useful to have in mind that the same system structure could also be used to model a single rack where each node again has multiple disks. In this case, $x_j^{(i)}$ is a single disk, the global processing unit (GPU) from Figure 2.1 is the rack processing unit (RPU), and the rack processing unit is now a node processing unit (NPU). The same system model would hold but for a different layer depth.

2.2 Eavesdropper Model

The secrecy threat for DSSs considered in this work is a passive eavesdropper. A passive eavesdropper is an adversary that can only read the stored data or downloaded data, needed for a repair, but cannot actively or maliciously change data or protocols of the system [PERR11]. However, it is assumed that the eavesdropper knows the system parameters.

The eavesdropper can read the data stored on l_1 nodes, in addition observe downloaded symbols and read all the nodes of l_2 groups and is therefore called (l_1, l_2)-eavesdropper. The corresponding sets that denote the nodes that can be observed are $\mathcal{E}_1$ with cardinality $|\mathcal{E}_1| = l_1$ and $\mathcal{E}_2$ with cardinality $|\mathcal{E}_2| = l_2 r$, where r denotes the number of independent symbols in each group. Furthermore, it is assumed that the eavesdropper does not observe nodes twice, i.e., $\mathcal{E}_2 \cap \mathcal{E}_1 = \emptyset$. The eavesdropper model is an adjusted version of the model presented in [RKSV14] which is similar to the model from [SRK11] with $l = l_1 + l_2$ and $l' = l_2$. The difference is that in [RKSV14] single nodes and their downloaded symbols are observed in an l_2-manner rather than whole groups. This implies that the repair computations are done in the corresponding node. In this work, a hierarchical model is assumed which is motivated by global repair discussed in Section 3.2.

The illustration of a DSS that was introduced in Section 2.1 and shown in Example 2.1 is now extended to cover the eavesdropper model as well as coding schemes with parities. In the following, local parity symbols are denoted by light grey boxes. The l_1 nodes that can be read by an eavesdropper are marked with a blue dot in the top left corner. The groups that are observed in an l_2-manner by the eavesdropper are marked with a red triangle at the top left corner of the group.

Example 2.2. Consider the DSS from Example 2.1 with three racks and four nodes per rack.

Let each group be encoded with a [4,3,2] single-parity check code. The parity symbol in each group is $x_4^{(i)} = x_1^{(i)} + x_2^{(i)} + x_3^{(i)}$. Let $l_1 = 2$ with $\mathcal{E}_1 = \{x_1^{(1)}, x_4^{(3)}\}$ be denoted by blue dots and let the group observed in an l_2-manner be indicated by a red triangle with the corresponding set $\mathcal{E}_2 = \{x_1^{(2)}, x_2^{(2)}, x_3^{(2)}\}$. The set $\mathcal{E}_2$ consists only of three nodes since the fourth node would be redundant. The example is illustrated in Figure 2.3.

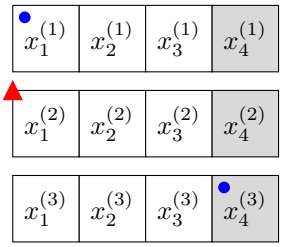

Figure 2.3: Illustration of a DSS with 3 racks and 4 nodes in each rack where each group is encoded with a single-parity check code. The blue dots indicate single nodes that can be read by the eavesdropper, i.e., $l_1 = 2$. The red triangle indicates that the second group is observed in an l_2-manner and is therefore fully known by the eavesdropper.

3 Preliminaries

The Preliminaries chapter summarizes the basics that are needed to understand the related work and the suggestions in the following sections.

It starts with a brief summary of secure communication and introduces secret sharing in Section 3.1. Secret sharing schemes are secrecy systems where a secret message is encoded, split into n parts, the shares, and distributed to n groups. Given $z < k \leq n$ shares, no information about the message is revealed. If $k \leq n$ shares are known, the secret message can be retrieved.

In Section 3.2, locally repairable codes (LRCs) are explained. The idea is to split a file into g parts and distribute it among g server racks. In each of these g groups, a local code is used to protect the data against erasures. In case of a node failure, only a small number of nodes, r, in the same local server rack are contacted for the repair. Maximally recoverable locally repairable codes (MR-LRCs), which is a class of LRCs, are also studied in this section. They can correct any erasure pattern that is information-theoretically correctable given the parameters of the code [GHJY14].

Important properties of sum-rank metric codes are summarized in Section 3.3. The sum-rank metric is a generalization of the rank metric and the Hamming metric.

Recently, constructions of MR-LRCs with sum-rank metric codes were introduced [MPK19, CMST22]. The constructions use linearized Reed-Solomon codes (LRSCs) which are summarized in Section 3.5. LRSCs are based on skew polynomials, which is a family of polynomials whose product is non-commutative. If the reader is not yet familiar with skew polynomials, they are extensively summarized in Appendix A following [MPSK22].

Moreover, two applications of skew polynomials are presented. In Section 3.6, the concept of secret sharing is applied to skew polynomials, which is useful to prove the secrecy of code constructions in the following parts. The concepts of skew polynomials fulfilling Lagrange constraints, i.e., vanishing on a P-independent set of points except for one, is explored in Section 3.7.

Finally, an information theory based lemma is stated in Section 3.8. It is later used to derive the secrecy capacity of different global repair schemes. The idea is to bound the entropy of the information collected by a data collector with the rank of a matrix.

3.1 Secret Sharing

3.1.1 Secure Communication

In secrecy systems, one party wants to transmit a message to another party in such a way that potential eavesdroppers are not able to recover the message. In the following, only messages chosen from a finite field $\mathbb{F}$ are considered. The keys to encrypt the message are chosen from the same field $\mathbb{F}$. On an abstract level, secrecy systems can be seen, as described by Shannon in [Sha49], as a transformation from the set of possible messages $m \in \mathbb{F}$ to a set of possible cryptograms or ciphertexts $c \in \mathbb{F}$. In this case, each particular transformation is characterized by a key $k \in \mathbb{F}$. The transformation should be reversible so that, given the key, a deciphering is possible. The random variables corresponding to the message m and the ciphertext c are denoted by M and C, respectively.

Definition 3.1 (**Perfect Secrecy**)**.** A secrecy system is perfectly secret if for every *a priori* probability distributions over the message space $\mathbb{F}$, every message $m \in \mathbb{F}$ and every ciphertext $c \in \mathbb{F}$, it holds that [Sha49]:

$$\Pr[\mathsf{M} = m \mid \mathsf{C} = c] = \Pr[\mathsf{M} = m].$$

That is, the mutual information $\mathrm{I}(\mathsf{M};\mathsf{C})$ between the random variables M and C must be zero: $\mathrm{I}(\mathsf{M};\mathsf{C}) = 0$.

Remark: In this work, information theoretical secrecy is the goal, which is realizable for DSS under the eavesdropper assumptions presented in the following section. However, perfectly secure systems are not feasible in many cases. Take for example an interceptable wireless channel where two parties want to transmit information in a perfectly secret way using the One-Time Pad. A secure channel would be needed to transmit the keys, for instance, by an in-person messenger carrying a hard disk drive with randomly generated key symbols that are only used once. Such efforts are rarely undertaken. Therefore, the goal in many applications is to design systems that have a computational security [GM84]. This means that any attack on the secrecy system should be so computationally complex that it cannot be done in a certain time correlating to the security level, assuming that certain computational tasks are "hard".

3.1.2 Secret Sharing Systems

Secret sharing systems are a special kind of secrecy system. In secret sharing systems, a message is not transmitted via a channel to a second party, but rather it is encoded, split

into pieces and shared with multiple parties. The encrypting or encoding involves randomly generated symbols. The message can be recovered by contacting multiple parties. The following definition is based on the secret sharing scheme requirement introduced by Shamir in [Sha79].

Definition 3.2 (**Secret sharing**)**.** The secret message m is encoded into n pieces $s_1, \ldots, s_n$, the so-called *shares*, such that

1. knowledge of any $k \leq n$ or more shares s_i allows to compute the message m (decodability).
2. knowledge of any $z < k$ shares s_i reveals no information about the secret message m (privacy).

Such a system is called an (n, k, z) secret sharing scheme.

In other words, a secret sharing system is perfectly secure if for every *a priori* probability distribution over the message space $\mathbb{F}$, every message $m \in \mathbb{F}$ and every vector of $z < k$ shares $\mathbf{s}_{\mathcal{I}_z} = (s_1, \ldots, s_z)$ with its corresponding random variable $\mathsf{S}_{\mathcal{I}_z}$, it holds that

$$\Pr[\mathsf{M} = m \mid \mathsf{S}_{\mathcal{I}_z} = (s_1, \ldots, s_z)] = \Pr[\mathsf{M} = m].$$

Equivalently, the mutual information between the random variables M and $\mathsf{S}_{\mathcal{I}_z}$ representing the message and the set of known shares must be zero, i.e.,

$$\mathrm{I}(\mathsf{M}; \mathsf{S}_{\mathcal{I}_z}) = 0 \text{ or } \mathrm{H}(\mathsf{M} \mid \mathsf{S}_{\mathcal{I}_z}) = \mathrm{H}(\mathsf{M}).$$

In addition, for every vector of k shares $\mathbf{s}_{\mathcal{I}_k} = (s_1, \ldots, s_k)$ it should hold that

$$\mathrm{H}(\mathsf{M} \mid \mathsf{S}_{\mathcal{I}_k}) = 0,$$

which means that the message can be recovered from the set of shares $\mathcal{S}_{\mathcal{I}_k}$.

To show that the mutual information of the eavesdropped shares and the message is zero, i.e., $\mathrm{I}(\mathsf{M}; \mathsf{S}_{\mathcal{I}_z}) = 0$, the following lemma can be used. It follows the steps described in [SRK11].

Lemma 3.1 (**Secrecy lemma**)**.** *Consider a secrecy system with message symbols m and random symbols r that are used in the encrypting/encoding process to generate the ciphertext and their random variables M and R. An eavesdropper is observing the symbols e at positions $\mathcal{I}_e \subseteq [n]$, a subset of the stored shares of the ciphertext $\mathbf{c} = (c_1, \ldots, c_n)$. The eavesdropper's observation is represented by the random variable E.*

If $\mathrm{H}(\mathsf{E}) \leq \mathrm{H}(\mathsf{R})$ *and* $\mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{E}) = 0$*, then the information leaked to the eavesdropper is zero:* $\mathrm{I}(\mathsf{M}; \mathsf{E}) = 0$.

Proof. Consider the mutual information

$$\begin{aligned}
\mathrm{I}(\mathsf{M}; \mathsf{E}) &= \mathrm{H}(\mathsf{E}) - \mathrm{H}(\mathsf{E} \mid \mathsf{M}) \\
&\overset{(a)}{\leq} \mathrm{H}(\mathsf{E}) - \mathrm{I}(\mathsf{E}; \mathsf{R} \mid \mathsf{M}) \\
&\overset{(b)}{\leq} \mathrm{H}(\mathsf{R}) - \mathrm{I}(\mathsf{E}; \mathsf{R} \mid \mathsf{M}) \\
&= \mathrm{H}(\mathsf{R}) - [\mathrm{H}(\mathsf{R} \mid \mathsf{M}) - \mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{E})] \\
&\overset{(c)}{=} \mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{E}) \\
&\overset{(d)}{=} 0
\end{aligned}$$

where (a) follows from the inequality $\mathrm{I}(\mathsf{X}; \mathsf{Y} \mid \mathsf{Z}) \leq \min(\mathrm{H}(\mathsf{X} \mid \mathsf{Z}), \mathrm{H}(\mathsf{Y} \mid \mathsf{Z}))$ (B.3), (b) is the assumption $\mathrm{H}(\mathsf{E}) \leq \mathrm{H}(\mathsf{R})$, (c) is due to independence of R and M yielding $\mathrm{H}(\mathsf{R}) = \mathrm{H}(\mathsf{R} \mid \mathsf{M})$ and (d) is the condition given above that $\mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{E}) = 0$. □

Example 3.1. The most simple secret sharing system is a (2,2,1) secret sharing scheme. Given the message $m \in \mathbb{F}_q$, the two shares are generated as follows:

1. generate a random number $r \in \mathbb{F}_q$ (uniformly distributed) and take it as the first share: $s_1 = r$
2. the second share is the sum of the message symbol and the random symbol: $s_2 = m + r$

It can be proven that the suggested scheme is a (2,2,1) secret sharing scheme fulfilling the requirements from Definition 3.2 as follows.

Proof. Given $k = 2$ shares, the message can be recovered: $m = s_2 - s_1$. To show that the privacy constraint is fulfilled, $\mathrm{P}_{\mathsf{M}|\mathsf{S}}(m \mid s_i) = \mathrm{P}_{\mathsf{M}}(m)$ for $i \in \{1, 2\}$ has to hold. Consider $\mathrm{P}_{\mathsf{M}|\mathsf{S}}(m \mid s_i)$ and apply Bayes' theorem

$$\begin{aligned}
\mathrm{P}_{\mathsf{M}|\mathsf{S}}(m \mid s_i) = \frac{\mathrm{P}_{\mathsf{S}|\mathsf{M}}(s_i \mid m)\,\mathrm{P}_{\mathsf{M}}(m)}{\mathrm{P}_{\mathsf{S}}(s_i)} &= \frac{\mathrm{P}_{\mathsf{S}|\mathsf{M}}(s_i \mid m)\,\mathrm{P}_{\mathsf{M}}(m)}{\sum_{\tilde{m} \in \mathbb{F}} \mathrm{P}_{\mathsf{S}|\mathsf{M}}(s_i \mid m)\,\mathrm{P}_{\mathsf{M}}(\tilde{m})} \\
&\overset{(a)}{=} \frac{\mathrm{P}_{\mathsf{R}}(r)\,\mathrm{P}_{\mathsf{M}}(m)}{\mathrm{P}_{\mathsf{R}}(r) \sum_{\tilde{m} \in \mathbb{F}} \mathrm{P}_{\mathsf{M}}(\tilde{m})} = \mathrm{P}_{\mathsf{M}}(m)
\end{aligned}$$

where (a) follows from $\mathrm{P}_{\mathsf{S}|\mathsf{M}}(s_1 \mid m) = \mathrm{P}_{\mathsf{S}|\mathsf{M}}(m + r \mid m) = \mathrm{P}_{\mathsf{R}}(r)$ and $\mathrm{P}_{\mathsf{S}|\mathsf{M}}(s_2 \mid m) = \mathrm{P}_{\mathsf{S}|\mathsf{M}}(r \mid m) = \mathrm{P}_{\mathsf{R}}(r)$ (r and m independent). This shows that the described system is a secret sharing scheme by Definition 3.2. □

To illustrate secret sharing, an example of a $(3, 3, 2)$ secret sharing scheme is given. Note that the key, i.e., the set of random symbols, is not needed to recover the message but access to k shares.

Example 3.2. Consider a secret sharing system with $n = 3$, $k = 3$ and $z = 2$. All three shares are required to retrieve the message. If an eavesdropper can access 2 or fewer shares, no information is revealed. The system is illustrated in Figure 3.1.

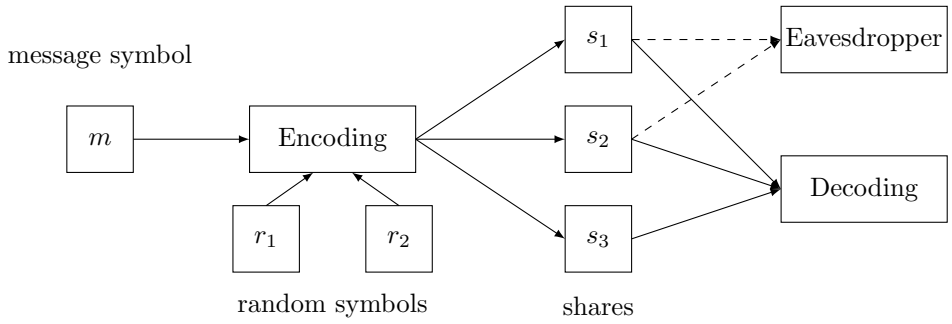

Figure 3.1: Illustration of a (3,3,2) secret sharing scheme. One message symbol $m \in \mathbb{F}$ is encoded with two random symbols $r_1, r_2 \in \mathbb{F}$. The result are 3 shares. If two or fewer shares are known, no information about the message is revealed. Given 3 shares, the message m can be decoded.

3.1.3 Shamir's Secret Sharing

A secret sharing scheme for an arbitrary number of parties was introduced by Shamir in [Sha79] and is based on polynomial interpolation.

Construction 3.1 (**Shamir's secret sharing**)**.** Fix the following integers n, k, $z = k-1$ and a prime power $q > n$. Given a message symbol $m \in \mathbb{F}_q$, generate z random numbers $r_1, \ldots, r_z$ independently and uniformly distributed over $\mathbb{F}_q$. The shares $s_1, \ldots, s_n$ can be calculated as evaluations of the polynomial

$$p(x) = m + r_1 x + r_2 x^2 + \cdots + r_z x^z = m + \sum_{j=1}^{z} r_j x^j$$

with $s_i = p(a_i)$, where all $a_i \in \mathbb{F}_q^*$ are pairwise distinct elements for all $i \in [n]$.

Theorem 3.1 (**Shamir's Secret Sharing**)**.** *Shamir's secret sharing scheme presented above is an $(n, k, z = k - 1)$ secret sharing scheme.*

Proof. Decodability: Given any k shares, Lagrange interpolation can be used to recover the polynomial $\tilde{p}(x) = p(x)$ of degree $k-1$. The message can then be retrieved with $\tilde{p}(0) = m$.
Privacy: Lemma 3.1 is used. The eavesdropped information is the set of shares $\mathcal{S}_{\mathcal{I}_z}$ given by a vector $\mathbf{s}_{\mathcal{I}_z}$ with $\mathcal{I}_z \subseteq [n]$, $|\mathcal{I}_z| = z$ and its random variable $\mathsf{S}_{\mathcal{I}_z}$. It is obvious that $\mathrm{H}(\mathsf{S}_{\mathcal{I}_z}) \leq \mathrm{H}(\mathsf{R})$ since R represents z independently and uniformly at random generated numbers $r_1, \ldots, r_z$. It remains to show that $\mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{S}_{\mathcal{I}_z}) = 0$. Rewriting the polynomial $p(x)$ yields

$$p(x) = m + x \underbrace{\sum_{j=1}^{z} r_j x^{j-1}}_{\hat{p}(x)}. \tag{3.1}$$

Given m and $s_i = p(a_i)$ for $i \in \mathcal{I}_z$, the evaluations of the polynomial $\hat{p}(x)$ of degree $(k-2)$, i.e., $\hat{s}_i = \hat{p}(a_i)$, can be calculated with

$$\hat{s}_i = \frac{p(a_i) - m}{a_i} = \frac{s_i - m}{a_i}.$$

With the z shares $\hat{s}_i$, $\hat{p}(x)$ can be retrieved using Lagrange interpolation, and the random numbers $r_1, \ldots, r_z$ are the coefficients of $\hat{p}(x)$. Thus, $\mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{S}_{\mathcal{I}_z}) = 0$. □

Example 3.3. Consider the field $\mathbb{F}_{q^m} = \mathbb{F}_{2^2}$ with $q = 2$ and $m = 2$ with the elements $a \in \{0, 1, \omega, \bar{\omega}\}$ as defined in (1.1). As parameters of the secret sharing scheme, take $n = 3$, which fulfills the constraint $q > n$, $k = 3$ and $z = k - 1 = 2$. This means that there are three parties and only access to all the shares allows retrieval of the message symbol. Let the message symbol be $m = \omega$ and the two randomly generated symbols $r_1 = \bar{\omega}$ and $r_2 = 1$. The shares are then evaluations of the polynomial

$$p(x) = \omega + \bar{\omega}x + x^2.$$

Thus, $s_1 = p(1) = 0$, $s_2 = p(\omega) = 0$ and $s_3 = p(\bar{\omega}) = \omega$ with $a_1 = 1$, $a_2 = \omega$ and $a_3 = \bar{\omega}$. Two cases are considered. The first case is a reconstruction of the polynomial given the evaluations. The polynomial received by Lagrange interpolation (see B.2) is

$$\begin{aligned}
\tilde{p}(x) &= \sum_{j=1}^{k} s_j \ell_j(x) \\
&= 0 + 0 + \omega \frac{(x-1) \cdot (x-\omega)}{(\bar{\omega}-1) \cdot (\bar{\omega}-\omega)} \\
&= \omega + \bar{\omega}x + x^2
\end{aligned}$$

which is, as required, the same as $p(x)$.

The second case is that the message m and a set of z shares are known. The goal is to recover the random symbols r_1, r_2. Let s_1 and s_3 be known. First, calculate the shares $\hat{s}_1$ and $\hat{s}_3$ as evaluations of the polynomial $\hat{p}(x)$ by removing the contribution of the message symbol m to s_1 and s_3 (Equation (3.1)).

$$\hat{s}_1 = \frac{s_1 - m}{a_1} = \frac{0 - \omega}{1} = \omega$$
$$\hat{s}_3 = \frac{s_3 - m}{a_3} = \frac{\omega - \omega}{\bar{\omega}} = 0$$

With $\hat{s}_1$ and $\hat{s}_3$, $\hat{p}(x)$ can be calculated using Lagrange interpolation:

$$\begin{aligned}\hat{p}(x) &= \sum_{j \in \{1,3\}} \hat{s}_j \ell_j(x) \\ &= \omega \frac{x - \bar{\omega}}{1 - \bar{\omega}} + 0 = \bar{\omega} + x\end{aligned}$$

and r_1, r_2 can be retrieved directly as the coefficients of $\hat{p}(x)$.

3.1.4 McEliece-Sarwate Secret Sharing

A more general secret sharing scheme, which is linked to Reed-Solomon codes, was introduced by McEliece and Sarwate in [MS81].

Construction 3.2 (**McEliece-Sarwate secret sharing**)**.** Fix the following integers n, k, $z < k$ and a prime power $q > n$. Given $k - z$ message symbols $m_1, \ldots, m_{k-z} \in \mathbb{F}_q$, generate z random numbers $r_1, \ldots, r_z$ independently and uniformly distributed over $\mathbb{F}_q$. The shares $s_1, \ldots, s_n$ can be calculated as evaluations of the polynomial

$$p(x) = r_1 + r_2 x + \cdots + r_z x^{z-1} + m_1 x^z + \cdots + m_{k-z} x^{k-1} = \sum_{j=1}^{z} r_j x^{j-1} + \sum_{i=1}^{k-z} m_i x^{i+z-1}$$

with $s_i = p(a_i)$, where all $a_i \in \mathbb{F}_q^*$ are pairwise distinct elements for all $i \in [n]$.

Remark: It is also possible to use a polynomial for the construction where the message symbols are the first $k - z$ coefficients and the random numbers follow as coefficients.

Theorem 3.2 (**McEliece-Sarwate secret sharing**)**.** *The McEliece-Sarwate secret sharing scheme presented above is an (n, k, z) secret sharing scheme.*

Proof. Decodability: Given any k shares, Lagrange interpolation can be used to recover the polynomial $\tilde{p}(x) = p(x)$ of degree $k - 1$. The message can then be retrieved by taking

the coefficients $\tilde{p}_i$ of $\tilde{p}(x)$ for $i \in \{k-z, k-z+1, \ldots, k-1\}$.
Privacy: Lemma 3.1 is used and the notation is the same as for the proof of Theorem 3.1. As R represents z independently and uniformly at random generated numbers $r_1, \ldots, r_z$, $\mathrm{H}(\mathsf{S}_{\mathcal{I}_z}) \leq \mathrm{H}(\mathsf{R})$ holds. It remains to show that $\mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{S}_{\mathcal{I}_z}) = 0$. Consider the polynomial $p(x)$

$$p(x) = \underbrace{\sum_{j=1}^{z} r_j x^{j-1}}_{p_r(x)} + \underbrace{\sum_{i=1}^{k-z} m_i x^{i+z-1}}_{p_m(x)}.$$

Given $m_1, m_2, \ldots m_{k-z}$ and $s_i = p(a_i)$ for $i \in \mathcal{I}_z$, $p_r(x)$ of degree $(z-1)$ can be retrieved using Lagrange interpolation with the shares $\hat{s}_i = p_r(a_i)$ that can be calculated as follows

$$\hat{s}_i = p(a_i) - p_m(a_i).$$

The random numbers $r_1, \ldots, r_z$ are the coefficients of the polynomial $p_r(x)$.
Thus, $\mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{S}_{\mathcal{I}_z}) = 0$. □

Remark: The encoding of the McEliece-Sarwate secret sharing scheme can also be written as a vector matrix multiplication:

$$(s_1, \ldots, s_n)^{\mathsf{T}} = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & 2 & \cdots & 2^{k-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & n & \cdots & n^{k-1} \end{pmatrix} (r_1, \ldots, r_z, m_1, \ldots, m_{k-z})^{\mathsf{T}}$$

The matrix is a Vandermonde matrix (corresponding to the generator matrix of a Reed-Solomon code) with full rank. The decodability can therefore also be easily shown for this representation since a $k \times k$ block of the Vandermonde matrix has rank k and thus one can recover the message symbols and the random symbols with k shares.

3.2 Locally Repairable Codes

As described in Section 2.1, the communication cost between nodes of the same rack is much less than between nodes of different racks. Therefore, it is beneficial to repair failed nodes with data stored on nodes of the same rack. A popular proposed solution are locally repairable codes (LRCs) [GHSY12],[HCL07],[OD11]. They allow to repair a failed node by contacting only a small number, r, of other nodes, where r is called the locality. For this purpose the code has several local groups in which such a local repair can be

performed. In practice, they have already been implemented by Facebook [SAP+13] and by Microsoft [HSX+12].

The following formal definition of LRCs follows the notation in [MPK19, Def. 4].

Definition 3.3 (Locally repairable code). Let $\Gamma_1, \Gamma_2, \ldots, \Gamma_g$ be a partition of $[n]$, i.e., $\Gamma_i \cap \Gamma_j = \varnothing$ for $i \neq j$, $[n] = \bigcup_{i=1}^{g} \Gamma_i$. For a fixed number of groups $g \geq 1$ and fixed integers r_i, δ_i for all $i \in [g]$, a code $\mathcal{C} \subseteq \mathbb{F}_{q^m}^n$ is said to be an (n, k) locally repairable code (LRC) with $(r_i, \delta_i)_{i=1}^{g}$-localities and $k = \log_{q^m} |\mathcal{C}|$ if it holds that

$$|\Gamma_i| \leq r_i + \delta_i - 1 \quad \text{and} \quad \mathrm{d_H}(\mathcal{C}|_{\Gamma_i}) \geq \delta_i,$$

for all $i \in [g]$. The set Γ_i is called i-th local group and δ_i is called the i-th local distance.

Thus, each local group can tolerate up to $\delta_i - 1$ erasures that can be recovered by contacting the r_i remaining nodes. In most cases, the groups are chosen to be of equal size, i.e., $r_1 = r_2 = \cdots = r_g$ and $\delta_1 = \delta_2 = \cdots = \delta_g$. In the following, an example is given to illustrate LRCs.

Example 3.4. Consider an LRC with two local groups ($g = 2$). The two local groups have $(3, 3)$-locality and can therefore tolerate $\delta - 1 = 2$ erasures that can be recovered with the $r = 3$ remaining nodes in each group. In Figure 3.2, the described code is illustrated. The light grey nodes indicate local parity symbols. In Figure 3.3, the erasure pattern cannot be entirely corrected. The first group can be repaired but there are too many erasures in the second group since $3 > \delta - 1 = 2$.

$x_1^{(1)}$	$x_2^{(1)}$	$x_3^{(1)}$	$x_4^{(1)}$	$x_5^{(1)}$
$x_1^{(2)}$	$x_2^{(2)}$	$x_3^{(2)}$	$x_4^{(2)}$	$x_5^{(2)}$

Figure 3.2: Illustration of an LRC with two local groups ($g = 2$) and $(r_i, \delta_i) = (3, 3)$-localities for $i \in [2]$. Each row forms a local group Γ_i that can correct up to $\delta - 1 = 2$ erasures. The light grey nodes store the parity symbols.

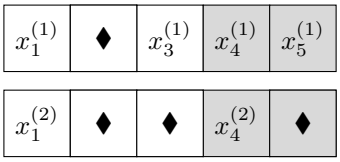

Figure 3.3: Illustration of an LRC with two local groups ($g = 2$) and $(r_i, \delta_i) = (3, 3)$-localities for $i \in [2]$. Each row forms a local group Γ_i that can correct up to $\delta - 1 = 2$ erasures. Four node failures have occurred, each denoted by a black diamond. The first local group can be repaired. The second group has too many erasures and can therefore not be repaired locally.

Note that the code discussed in Example 3.4 has 4 parity symbols, yet it cannot correct the in Figure 3.3 illustrated erasure pattern with 4 erasures. This illustrates that the definition of LRCs does not make any statement about the global distance of the code. LRCs that attain Singleton-type bounds on their global distance are called optimal LRCs. One bound is briefly discussed in Chapter 4 in the context of MR-LRCs on which this work focuses. MR-LRCs are a strictly stronger class of LRCs than "optimal" LRCs. Given the localities, MR-LRCs can correct any information theoretically correctable erasure pattern.

Definition 3.4 (**Maximal recoverability** [MPK19, Def. 5])**.** Let $\mathcal{C} \subseteq \mathbb{F}_{q^m}^n$ be a code with $(r_i, \delta_i)_{i=1}^g$-localities. It is said to be maximally recoverable (MR), if for any $\mathcal{R}_i \subseteq \Gamma_i$ with $|\Gamma_i \setminus \mathcal{R}_i| = \delta_i - 1$ for $i = 1, 2, \ldots, g$, the code $\mathcal{C}|_{\mathcal{R}} \subseteq \mathbb{F}_{q^m}^{|\mathcal{R}|}$ with $\mathcal{R} = \bigcup_{i=1}^g \mathcal{R}_i$ is MDS.

To achieve maximal recoverability, global parities that can correct erasures of the global code are needed. The number of global parities is denoted by $h = \sum_{i=1}^g r_i - k$. Any MR-LRC can therefore correct $\delta_i - 1$ in each local group and in addition h erasures anywhere.

Example 3.5. Consider an MR-LRC with five local groups ($g = 5$). The five local groups have $(3, 3)$-locality and thus can tolerate $\delta - 1 = 2$ erasures that can be recovered with the $r = 3$ remaining nodes. One of these local groups consists of 3 global parities, i.e., $k = 12$ and $h = 3$. If one local group has more than two erasures, they can be corrected with the global parities. In Figure 3.4 the described code is illustrated. The light grey nodes indicate local parity symbols, the dark grey nodes global parities.

In Figure 3.5, the black diamonds and black stars denote erasures. In two steps, all the erasures can be corrected.

1. Correct the three erasures denoted by black stars with the global erasure capability. Definition 3.4 tells us that after puncturing two symbols in each group, e.g., the

two erasures that have occurred denoted by black diamonds, the remaining nodes form an MDS code. With the given code parameter $k = 12$ and 15 symbols left after the puncturing, which means that $n_{\text{mds}} = 15$, the code can tolerate another $d - 1 = n - k = 3$ erasures, for example, the ones denoted by black stars.

2. Correct the other erasures, denoted by black diamonds (not more than two per group), locally in each group.

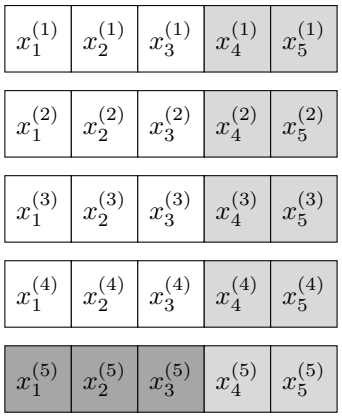

Figure 3.4: Illustration of an MR-LRC with $(r_i, \delta_i) = (3, 3)$-localities for $i \in [5]$ and five local groups ($g = 5$). Each row forms a local group Γ_i. This code can correct up to $\delta - 1 = 2$ erasures locally in each group and in addition 3 erasures located anywhere with the global parities (dark grey).

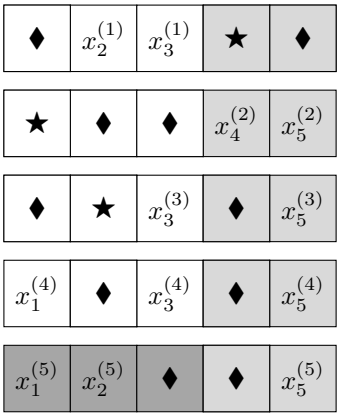

Figure 3.5: Illustration of an MR-LRC with $(r_i, \delta_i) = (3, 3)$-localities for $i \in [5]$ and five local groups ($g = 5$). Each row forms a local group Γ_i. This code can correct up to $\delta - 1 = 2$ erasures locally in each group and in addition 3 global erasures with the global parities (dark grey). The given erasure pattern can be corrected with the two steps described in Example 3.5.

Constructions of MR-LRCs are presented in Section 4.2.

At this point, it is important to link the DSS model described in Section 2.1 with LRCs. Usually, the system perspective is not given or not presented in detail and DSSs are considered from a theoretical perspective only. The local repair in a DSS is often abstracted such that each node can communicate with all the other nodes directly. Given a node failure, the node sends requests to the nodes needed for a repair, e.g., it requests the symbols of r nodes given an LRC with (r, δ)-locality. The download of symbols needed for a repair is illustrated in Figure 3.6.

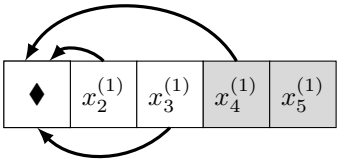

Figure 3.6: Repair download of a failed node in an LRC group with $(3, 3)$-locality under the assumption that the local repair takes place in the failed node. The node downloads 3 symbols for repair.

In contrast to that, the model assumed in this work has a hierarchical structure. Given a failed node, the node sends a request to the rack processing unit (RPU). The RPU knows which of the nodes in the group are still running and sends a request to a sufficient number of nodes such that the erased symbols of the failed node can be

recovered. The RPU can now reconstruct the symbols and send them back to the node that has failed. The described model is illustrated in Figure 3.7.

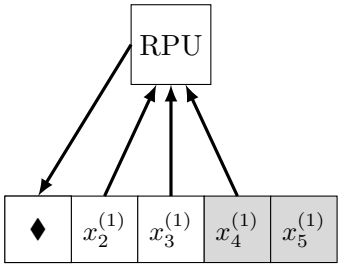

Figure 3.7: Repair download of a failed node in an LRC group with $(3, 3)$-locality managed by a rack processing unit (RPU). The RPU downloads the 3 symbols needed to repair the failed node, calculates the symbol and sends it to the failed node for repair.

In this thesis, the second model was chosen since the focus is on secure coding. If only local repair is considered, the two models behave equivalently. In case of the first model, an eavesdropper observing the group in an l_2-manner gains knowledge of the whole group since it can read the downloaded symbols for repair such as in [RKSV14]. In the second model, an eavesdropper observing the group in an l_2-manner has access to the whole group including the RPU and has therefore direct access to all the nodes. For local repair only, the two system models illustrated in Figure 3.6 and Figure 3.7 behave equivalently.

For global repair, the models behave differently. Assume that in an MR-LRC there is a group with more erasures than the local group can correct. A global correction would be performed. In case of the first model, the failed node would contact as many nodes in the same and in other groups such that a global correction would be possible. Thus, it gains global knowledge. Given an eavesdropper which observes the group in an l_2-manner, global knowledge would be revealed which contradicts the possibility of storing symbols securely. On the other hand, the second model allows global repair schemes which do not reveal global knowledge to the group where a global repair is performed given an eavesdropper which observes the group in an l_2-manner. Therefore, the second model is chosen in this work and discussed in detail in Chapter 5. It also has the advantage that local group pre-computations can reduce the load of the global switch which would have higher communication costs as discussed in Section 2.1.

3.3 Sum-Rank Metric Codes

The sum-rank metric and sum-rank metric codes are motivated by communication scenarios over channels that involve the action of a block diagonal matrix [MPSK22]. Consider the multiplicative-additive matrix channel

$$\mathbf{Y} = \mathbf{AX} + \mathbf{Z} = \operatorname{diag}(\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_g)\mathbf{X} + \mathbf{Z}$$

with $\mathbf{Y} \in \mathbb{K}^{r\times m}$ as the received symbols, the additive noise denoted by $\mathbf{Z} \in \mathbb{K}^{r\times m}$, the sent symbols $\mathbf{X} \in \mathbb{K}^{k\times m}$ and the multiplicative behavior of the channel represented by $\mathbf{A} \in \mathbb{K}^{r\times k}$. In this example, $\mathbb{K}$ can be any field. The block diagonal matrix with matrices $\mathbf{A}_i \in \mathbb{K}^{r_i\times k_i}$ for $i \in [g]$ on the main diagonal is denoted by

$$\mathbf{A} = \operatorname{diag}(\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_g) = \begin{pmatrix} \mathbf{A}_1 & 0 & \cdots & 0 \\ 0 & \mathbf{A}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A}_g \end{pmatrix}.$$

Sum-rank metric codes, which are a natural coding solution for this channel model, are useful in multiple coding disciplines such as network coding [NUF10], space-time coding [SK22] and coding for DSSs [MPK19]. For the latter, each subchannel represents a rack of a server. The thesis focuses on DSSs only and therefore $\mathbb{K} = \mathbb{F}_{q^m}$ is a finite field with extension degree m.

Since $\mathbb{F}_{q^m}$ is an m-dimensional vector space over $\mathbb{F}_q$, a vector $\mathbf{x} \in \mathbb{F}_{q^m}^n$ can also be written as a matrix $\mathbf{X}_i \in \mathbb{F}_q^{m\times n}$, as follows.

Definition 3.5. For an ordered basis $\boldsymbol{\beta} = (\beta_1, \beta_2, \ldots, \beta_m)$ of $\mathbb{F}_{q^m}$ over $\mathbb{F}_q$ with $\beta_i \in \mathbb{F}_{q^m}$, any vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{F}_{q^m}^n$ of length n can be represented by a matrix $M_\beta^n(\mathbf{x})$. The matrix map $M_\beta^n(\mathbf{x})$: $\mathbb{F}_{q^m}^n \longrightarrow \mathbb{F}_q^{m\times n}$ takes $\mathbf{x}$ to

$$M_\beta^n(\mathbf{x}) = \begin{pmatrix} x_{1,1} & x_{1,2} & \cdots & x_{1,n} \\ x_{2,1} & x_{2,2} & \cdots & x_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ x_{m,1} & x_{m,2} & \cdots & x_{m,n} \end{pmatrix} \in \mathbb{F}_q^{m\times n}$$

with $x_{i,j} \in \mathbb{F}_q$ for $i \in [m]$, $j \in [n]$. The vector $\mathbf{x}$ can be retrieved by the operation $\boldsymbol{\beta} M_\beta^n(\mathbf{x}) = \mathbf{x}$. The matrix representation of a vector $\mathbf{x} \in \mathbb{F}_{q^m}^n$ is written as $\mathbf{X} = M_\beta^n(\mathbf{x})$.

Example 3.6. Consider $\mathbb{F}_4$ with $q = 2$ and $m = 2$ as defined in (1.1). An ordered basis of $\mathbb{F}_{2^2}$ over $\mathbb{F}_2$ is $\boldsymbol{\beta} = (1, \omega)$. Take now the following vector as an example

$$\mathbf{x} = (\bar{\omega}, 1, \omega) \in \mathbb{F}_4^3.$$

The corresponding matrix representation with respect to the basis $\boldsymbol{\beta}$ is

$$\mathbf{X} = M_{\beta}^{n}(\mathbf{x}) = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix} \in \mathbb{F}_2^{2\times 3}.$$

Now, let the vector $\mathbf{x} \in \mathbb{F}_{q^m}^n$ be partitioned into g groups and let the i-th part $\mathbf{x}^{(i)}$ of the vector $\mathbf{x}$ be of length r_i such that $\sum_{i=1}^{g} r_i = n$. Thus, the vector $\mathbf{x}$ can be written as the concatenation of all parts $\mathbf{x}^{(i)} \in \mathbb{F}_{q^m}^{r_i}$ for $i \in [g]$, i.e., $\mathbf{x} = \left(\mathbf{x}^{(1)} \mid \ldots \mid \mathbf{x}^{(g)}\right) \in \mathbb{F}_{q^m}^n$.

With its corresponding matrix representation $\mathbf{X}_i$, the rank weight of $\mathbf{x}^{(i)}$ can be defined as

$$\mathrm{wt}_{\mathrm{rk}} : \mathbb{F}_{q^m}^{r_i} \longrightarrow \mathbb{N}_0$$

with

$$\mathrm{wt}_{\mathrm{rk}}(\mathbf{x}^{(i)}) = \mathrm{rank}(\mathbf{X}_i).$$

Definition 3.6 (Sum-rank metric). Let $\mathbf{x} = \left(\mathbf{x}^{(1)} \mid \ldots \mid \mathbf{x}^{(g)}\right) \in \mathbb{F}_{q^m}^n$ be partitioned into g groups of length r_i where $n = \sum_{i=1}^{g} r_i$ and $\mathbf{x}^{(i)} \in \mathbb{F}_{q^m}^{r_i}$. The sum-rank weight of $\mathbf{x}$ is defined as

$$\mathrm{wt}_{\mathrm{SR}} : \mathbb{F}_{q^m}^{r_1} \times \mathbb{F}_{q^m}^{r_2} \times \cdots \times \mathbb{F}_{q^m}^{r_g} \longrightarrow \mathbb{N}_0$$

with

$$\mathrm{wt}_{\mathrm{SR}}(\mathbf{x}) = \sum_{i=1}^{g} \mathrm{wt}_{\mathrm{rk}}(\mathbf{x}^{(i)}) = \sum_{i=1}^{g} \mathrm{rank}(\mathbf{X}_i).$$

With the sum-rank weight function, a sum-rank distance between two vectors $\mathbf{x}, \mathbf{y}$ with the same sum-length partition, i.e., the same number of groups g and lengths of the groups r_i, can be defined as

$$\mathrm{d}_{\mathrm{SR}} : \mathbb{F}_{q^m}^n \times \mathbb{F}_{q^m}^n \longrightarrow \mathbb{N}_0$$

with

$$\mathrm{d}_{\mathrm{SR}}(\mathbf{x}, \mathbf{y}) = \mathrm{wt}_{\mathrm{SR}}(\mathbf{x} - \mathbf{y}).$$

The reader may verify that the sum-rank distance is indeed a metric. By choosing $r_1 = r_2 = \ldots = r_g = 1$, the sum-rank metric is equal to the Hamming metric with Hamming weight wt_{H} and Hamming distance d_{H} [Ham50]. For $g = 1$, it is equal to the rank metric [Gab85], [Rot91].

Definition 3.7 (**Sum-rank metric code**)**.** Given a code $\mathcal{C} \subseteq \mathbb{F}_{q^m}^n$ with a fixed sum length partition $r_1, \ldots, r_g$, its minimum sum-rank distance is

$$\mathrm{d_{SR}}(\mathcal{C}) = \min_{\mathbf{c}\neq\mathbf{d}\in\mathcal{C}}\{\mathrm{d_{SR}}(\mathbf{c},\mathbf{d})\} \overset{(a)}{=} \min_{\mathbf{c}\in\mathcal{C}\setminus\{\mathbf{0}\}}\{\mathrm{wt_{SR}}(\mathbf{c})\}$$

where (a) holds if $\mathcal{C}$ is linear.

An interesting case is sum-rank metric codes multiplied with a block diagonal matrix. Let $\mathcal{C} \subseteq \mathbb{F}_{q^m}^n$ be a code, then $\mathcal{C}\mathbf{A} \triangleq \{\mathbf{cA} \mid \mathbf{c} \in \mathcal{C}\}$ for $\mathbf{A} \in \mathbb{F}_q^{n\times n}$ which is needed for the following theorem.

Theorem 3.3 ([MPK19, Th. 1])**.** *For a code $\mathcal{C} \subseteq \mathbb{F}_{q^m}^n$, it holds that*

$$\begin{aligned}\mathrm{d_{SR}}(\mathcal{C}) = \min\{\mathrm{d_H}(\mathcal{C}\mathbf{A}) \mid \mathbf{A} = \mathrm{diag}(\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_g),\\ \mathbf{A}_i \in \mathbb{F}_q^{r_i\times r_i} \text{ invertible}, \forall i \in [g]\}.\end{aligned}$$

Proof. The matrices $\mathbf{A}_i \in \mathbb{F}_q^{r_i\times r_i}$ have full rank following from their invertibility. Therefore, $\mathrm{d_{SR}}(\mathcal{C}) = \mathrm{d_{SR}}(\mathcal{C}\mathbf{A})$. The sum-rank distance is upper bounded by the Hamming distance as follows. Consider the Hamming weights of the codeword $\mathbf{x} \in \mathcal{C}\mathbf{A}$. Every nonzero column of $\mathbf{X}_i$ ($\mathbf{x}^{(i)}$ in matrix representation) contributes to the Hamming weight but is not necessarily linearly independent from the other columns of $\mathbf{X}_i$. This yields $\mathrm{d_{SR}}(\mathcal{C}) = \mathrm{d_{SR}}(\mathcal{C}\mathbf{A}) \leq \mathrm{d_H}(\mathcal{C}\mathbf{A})$ which shows that $\mathrm{d_{SR}}(\mathcal{C})$ is upper bounded by $\mathrm{d_H}(\mathcal{C}\mathbf{A})$. Choose $\mathbf{x}, \mathbf{y} \in \mathcal{C}$, such that $\mathrm{d_{SR}}(\mathbf{x},\mathbf{y}) = \mathrm{d_{SR}}(\mathcal{C})$. Let $\mathbf{x} = \left(\mathbf{x}^{(1)} \mid \ldots \mid \mathbf{x}^{(g)}\right) \in \mathbb{F}_{q^m}^n$ with matrix representation $\mathbf{X} = (\mathbf{X}_1 \mid \ldots \mid \mathbf{X}_g) \in \mathbb{F}_q^{m\times n}$, similarly $\mathbf{y}$ and $\mathbf{Y}$, where $\mathbf{x}^{(i)}, \mathbf{y}^{(i)} \in \mathbb{F}_{q^m}^{r_i}$ and $\mathbf{X}_i, \mathbf{Y}_i \in \mathbb{F}_q^{m\times r_i}$ for all $i \in [g]$. There exists an invertible matrix $\mathbf{A}_i \in \mathbb{F}_q^{r_i\times r_i}$ transforming $(\mathbf{X}_i - \mathbf{Y}_i)$ by basic linear algebra operations such that

$$(\mathbf{X}_i - \mathbf{Y}_i)\,\mathbf{A}_i = (\mathbf{B}_i \mid 0_{r_i - s_i}) \in \mathbb{F}_q^{m\times r_i}, \tag{3.2}$$

where $\mathbf{B}_i \in \mathbb{F}_q^{m\times s_i}$ is a full-rank matrix, with rank s_i, spanning the column space of $(\mathbf{X}_i - \mathbf{Y}_i)$. Therefore, it holds that $s_i = \mathrm{rank}(\mathbf{B}_i) = \mathrm{rank}(\mathbf{X}_i - \mathbf{Y}_i)$ for all $i \in [g]$. Let $\mathbf{A} = \mathrm{diag}(\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_g) \in \mathbb{F}_q^{n\times n}$. With Equation (3.2) applied to all the g blocks of $\mathbf{X}$ and $\mathbf{Y}$ it follows that $\mathrm{wt_{SR}}((\mathbf{x}-\mathbf{y})\mathbf{A}) = \sum_{i=1}^g \mathrm{rank}(\mathbf{B}_i) = \sum_{i=1}^g s_i = \mathrm{wt_H}((\mathbf{x}-\mathbf{y})\mathbf{A})$. Therefore,

$$\begin{aligned}\mathrm{d_{SR}}(\mathcal{C}) = \mathrm{d_{SR}}(\mathcal{C}\mathbf{A}) &= \mathrm{d_{SR}}(\mathbf{xA}, \mathbf{yA}) = \mathrm{wt_{SR}}((\mathbf{x}-\mathbf{y})\,\mathbf{A})\\ &= \mathrm{wt_H}((\mathbf{x}-\mathbf{y})\,\mathbf{A}) = \mathrm{d_H}(\mathbf{xA}, \mathbf{yA})\\ &\geq \mathrm{d_H}(\mathcal{C}\mathbf{A})\end{aligned}$$

which proves that there exists a matrix $\mathbf{A}$ such that $\mathrm{d_{SR}}(\mathcal{C})$ is lower bounded by $\mathrm{d_H}(\mathcal{C}\mathbf{A})$. □

From this theorem, an erasure correction corollary can be derived.

Corollary 3.1 (**Erasure correction** [MPK19, Cor. 1])**.** *Let $0 \leq t < n$ and $\mathcal{C} \subseteq \mathbb{F}_{q^m}^n$ be the considered code. Also let $\mathbf{A}_i \in \mathbb{F}_q^{r_i \times n_i}$ for all $n_i \geq 1$ and $i \in [g]$. It is equivalent that*

$$t < d_{\mathrm{SR}}(\mathcal{C}),$$

and that for

$$n - \sum_{i=1}^{g} \mathrm{rank}(\mathbf{A}_i) \leq t,$$

any codeword $\mathbf{x} \in \mathcal{C}$ can be uniquely recovered from $\mathbf{x}' = \mathbf{x}\,\mathrm{diag}(\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_g)$.

For sum-rank metric codes, there is a similar terminology to MDS codes in the Hamming metric.

Definition 3.8 (**Maximum sum-rank distance codes**)**.** A linear code $\mathcal{C}$ is said to be a maximum sum-rank distance (MSRD) code if one of the following equivalent conditions hold:

1. $$\mathrm{d_{SR}}(\mathcal{C}) = n - \dim(\mathcal{C}) + 1$$
2. $\mathcal{C}\mathbf{A} \subseteq \mathbb{F}_{q^m}^n$ is MDS, for all $\mathbf{A} = \mathrm{diag}(\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_g) \in \mathbb{F}_q^{n \times n}$ with $\mathbf{A}_i \in \mathbb{F}_q^{r_i \times r_i}$ is invertible for all $i \in [g]$.

Maximum sum-rank distance (MSRD) is the equivalent in the sum-rank metric to MDS in the Hamming metric for codes that achieve the Singleton bound. The Singleton bound for sum-rank metric codes reads smilar to the the Singleton bound for codes in the Hamming metric:

Definition 3.9 (**Singleton bound for sum-rank metric codes**)**.** Let $\mathcal{C} \subseteq \mathbb{F}_{q^m}^n$ be a code. Then,

$$|\mathcal{C}| \leq q^{m(n - \mathrm{d_{SR}}(\mathcal{C}) + 1)}.$$

Linearized Reed-Solomon codes were the first suggested codes that are MSRD codes and therefore achieve the bound with equality. They are a generalization of Reed-Solomon codes and Gabidulin codes and use skew polynomials which are characterized in Appendix 7.2.

3.4 Skew Reed-Solomon Codes

Similar to Reed-Solomon codes which are linked to Vandermonde matrices, skew Reed-Solomon codes are linked to skew Vandermonde matrices. They go back to [LL88],[Lam85] and are introduced in the following.

Definition 3.10 (**Skew Vandermonde matrix**)**.** Given a vector $\mathbf{b} = (b_1, b_2, \ldots, b_n) \in \mathbb{F}_{q^m}^n$, the skew Vandermonde matrix of order $d \in \mathbb{N}$ on $\mathbf{b}$, with respect to σ, is defined as

$$\mathbf{V}_d^\sigma(\mathbf{b}) = \begin{pmatrix} N_0(b_1) & N_0(b_2) & \cdots & N_0(b_n) \\ N_1(b_1) & N_1(b_2) & \cdots & N_1(b_n) \\ \vdots & \vdots & \ddots & \vdots \\ N_{d-1}(b_1) & N_{d-1}(b_2) & \cdots & N_{d-1}(b_n) \end{pmatrix} \in \mathbb{F}_{q^m}^{d\times n}$$

with $N_i(b) = \sigma^{i-1}(b)\sigma^{i-2}(b)\cdots\sigma(b)b$ for $b \in \mathbb{F}_{q^m}$ and $i \in \mathbb{N}$. The field automorphism σ is chosen as

$$\sigma(a) = a^q, \quad \forall\, a \in \mathbb{F}_{q^m},$$

and the i-th composition of σ is

$$\sigma^i(a) = \underbrace{\sigma(\sigma(\cdots\sigma(a)))}_{\text{i times}} = a^{q^i},$$
$$x^i a = \sigma^i(a)x = a^{q^i}x.$$

The usual Vandermonde matrix is a representation of the conventional polynomial evaluation map. Similarly, this is the case for skew polynomials and skew Vandermonde matrices. Observe that if $F = F_0 + F_1 x + \cdots F_{d-1}x^{d-1} \in \mathbb{F}_{q^m}[x;\sigma]$ is a skew polynomial of degree $d-1$ with coefficients $F_0, F_1, \ldots, F_{d-1} \in \mathbb{F}_{q^m}$, then

$$(F_0, F_1, \ldots, F_{d-1}) \cdot \mathbf{V}_d^\sigma(\mathbf{b}) = (F(b_1), F(b_2), \ldots, F(b_n)).$$

Therefore, many results that are stated in terms of the evaluation of skew polynomials can also be expressed by their skew Vandermonde matrix equivalent. The following theorem can be derived directly from the Lagrange theorem for skew polynomials (Theorem A.7).

Theorem 3.4 ([Lam85, Th. 8])**.** *Let $\Omega = \{b_1, b_2, \ldots, b_n\} \subseteq \mathbb{F}_{q^m}$ be a set with cardinality $|\Omega| = n$ and $\mathbf{b} = (b_1, b_2, \ldots, b_n)$ the corresponding vector. The following statements are equivalent:*

1. *The set Ω is P-independent in $\mathbb{F}_{q^m}[x;\sigma]$.*
2. *For some $d \geq n$ the matrix $\mathbf{V}_d^\sigma(\mathbf{b}) \in \mathbb{F}_{q^m}^{d\times n}$ has rank n.*

3. *The $n \times n$ matrix $\mathbf{V}_n^\sigma(\mathbf{b}) \in \mathbb{F}_{q^m}^{n\times n}$ is invertible.*

Proof. The items 2. and 3. in this theorem are equivalent to the items 2. and 3. in Theorem A.7 but in a different notation. Thus, they hold as well. □

Based on the skew Vandermonde matrix and correspondingly the skew polynomial evaluation, a skew Reed-Solomon code can be defined.

Definition 3.11 (**Skew Reed-Solomon codes** [BU13, Def. 7], [MPSK22, Def. 2.10])**.** Let $\Omega = \{b_1, b_2, \ldots, b_n\} \subseteq \mathbb{F}_{q^m}$ be a P-independent set with cardinality $|\Omega| = n$ and $\mathbf{b} = (b_1, b_2, \ldots, b_n) \in \mathbb{F}_{q^m}$ the corresponding vector. The k-dimensional skew Reed-Solomon code on $\mathbf{b}$ with respect to σ for $k \in [n]$ is

$$\mathcal{C}_{\mathrm{SRS}}^{\sigma,k}(\mathbf{b}) = \left\{ \mathbf{x}\mathbf{V}_k^\sigma(\mathbf{b}) \mid \mathbf{x} \in \mathbb{F}_{q^m}^k \right\}$$

where the generator matrix is given by the skew Vandermonde matrix $\mathbf{V}_k^\sigma(\mathbf{b}) \in \mathbb{F}_{q^m}^{k\times n}$. For a skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$, denote

$$F(\mathbf{b}) = (F(b_1), F(b_2), \ldots, F(b_n)) \in \mathbb{F}_{q^m}^n$$

as the vector of evaluations of F at $\mathbf{b}$. The skew Reed-Solomon code can then be also defined as

$$\mathcal{C}_{\mathrm{SRS}}^{\sigma,k}(\mathbf{b}) = \{F(\mathbf{b}) \mid F \in \mathbb{F}_{q^m}[x;\sigma],\ \deg(F) < k\}\,.$$

From Theorem 3.4, it follows that skew Reed-Solomon codes, like their conventional counterpart, attain the Singleton bound in the Hamming metric. The following result was introduced in [BU13].

Theorem 3.5 ([BU13, Prop. 2])**.** *Let $\Omega = \{b_1, b_2, \ldots, b_n\} \subseteq \mathbb{F}_{q^m}$ be a P-independent set with cardinality $|\Omega| = n$ and $\mathbf{b} = (b_1, b_2, \ldots, b_n) \in \mathbb{F}_{q^m}^n$ the corresponding vector. The skew Reed-Solomon code $\mathcal{C}_{\mathrm{SRS}}^{\sigma,k}(\mathbf{b})$ has dimension k and is MDS, i.e.,*

$$\mathrm{d_H}\left(\mathcal{C}_{\mathrm{SRS}}^{\sigma,k}(\mathbf{b})\right) = n - k + 1.$$

Proof. The MDS property of $\mathcal{C}_{\mathrm{SRS}}^{\sigma,k}(\mathbf{b})$ and that it has dimension k can be proven by showing that every $k \times k$ submatrix of the generator matrix is invertible. The $k \times k$ submatrix of $\mathbf{V}_k^\sigma(\mathbf{b})$ is $\mathbf{V}_k^\sigma(\tilde{\mathbf{b}}) \in \mathbb{F}_{q^m}^{k\times k}$ with $\tilde{\mathbf{b}} = (b_{i_1}, b_{i_2}, \ldots, b_{i_k}) \in \mathbb{F}_{q^m}^k$ and $1 \le i_1 < i_2 < \ldots < i_k \le n$. From Theorem A.7, it follows that every subset of a P-independent set is also P-independent (Corollary A.2). Thus, the set $\tilde{\Omega} = \{b_{i_1}, b_{i_2}, \ldots, b_{i_k}\} \subseteq \Omega$ is also P-independent and with Theorem 3.4, the matrix $\mathbf{V}_k^\sigma(\tilde{\mathbf{b}})$ is invertible. □

Skew Reed-Solomon codes recover Reed-Solomon codes for $\sigma = \mathrm{Id}$. In that case, the skew Vandermonde matrix becomes a conventional Vandermonde matrix. However, they do not recover generalized Reed-Solomon [RS60] codes or Gabidulin codes [Gab85],[Rot91]. The skew Reed-Solomon codes can be adapted such that they recover the former by using a conventional Vandermonde matrix and column multipliers as follows

$$\mathbf{V}_k^{\sigma=\mathrm{Id}}(\mathbf{b}) \cdot \mathrm{diag}(\alpha_1, \alpha_2, \ldots, \alpha_n) \in \mathbb{F}_{q^m}^{k \times n}$$

with $\alpha_i \in \mathbb{F}_{q^m}^*$ arbitrary for $i \in [n]$. Gabidulin codes can as well be recoverd for a specific choice of the column multipliers α_i. This specific choice yields linearized Reed-Solomon codes which are discussed in the next section.

3.5 Linearized Reed-Solomon Codes

The goal is to construct codes that can be multiplied by a block diagonal matrix $\mathbf{A}$ and remain MDS (see Definition 3.8). To achieve this, skew polynomials are modified so that they become $\mathbb{F}_q$-linear with respect to the evaluation.

Definition 3.12 ([Ler95, Ex. 2.6])**.** Let $\sigma : \mathbb{F}_{q^m} \longrightarrow \mathbb{F}_{q^m}$, for all $a \in \mathbb{F}_{q^m}$, be given by $\sigma(a) = a^q$. To each $a \in \mathbb{F}_{q^m}$, a map is associated:

$$\begin{aligned} D_a \; : \; \mathbb{F}_{q^m} &\longrightarrow \mathbb{F}_{q^m} \\ \beta &\longmapsto \sigma(\beta)a. \end{aligned}$$

The operators D_a are F_q-linear, since for $a, \beta_1, \beta_2 \in \mathbb{F}_{q^m}$ and $\lambda_1, \lambda_2 \in \mathbb{F}_q$ it holds that

$$\begin{aligned} D_a(\lambda_1\beta_1 + \lambda_2\beta_2) &= \sigma(\lambda_1\beta_1 + \lambda_2\beta_2)a \\ &= \sigma(\lambda_1\beta_1)a + \sigma(\lambda_2\beta_2)a = \lambda_1\sigma(\beta_1)a + \lambda_2\sigma(\beta_2)a \\ &= \lambda_1 D_a(\beta_1) + \lambda_2 D_a(\beta_2). \end{aligned}$$

The i-th composition of D_a with itself is given by

$$\begin{aligned} D_a^i(\beta) &= \sigma^i(\beta)N_i(a) \\ &= \sigma^i(\beta)\sigma^{i-1}(a) \cdots \sigma(a)a \end{aligned}$$

for all $\beta \in \mathbb{F}_{q^m}$ and $i \in \mathbb{N}$, where $N_i(a)$ is defined as in Definition 3.10. The zeroth power of D_a is $D_a^0 = \mathrm{Id}$ with Id denoting the identity map.

With this family of operators, a set of $\mathbb{F}_q$-linear polynomials can be defined.

Definition 3.13 (**Linear operator polynomials** [Ler95, Ex. 2.6])**.** The ring of polynomials in D_a is defined as

$$\mathbb{F}_{q^m}[D_a] = \left\{ \sum_{i=0}^{d} F_d D_a^i \,\middle|\, d \in \mathbb{N}_0, F_i \in \mathbb{F}_{q^m} \ \forall i \in [d] \cup \{0\} \right\}$$

The identity element is the identity map $\mathrm{Id} = D_a^0$. For two linear operator polynomials

$$F = F_0 D_a^0 + F_1 D_a^1 + \cdots + F_d D_a^d \in \mathbb{F}_{q^m}[D_a]$$
$$G = G_0 D_a^0 + G_1 D_a^1 + \cdots + G_d D_a^d \in \mathbb{F}_{q^m}[D_a]$$

with $d \in \mathbb{N}$, $F_i, G_i \in \mathbb{F}_{q^m}$ for $i \in \{0\} \cup [d]$ and a scalar $a \in \mathbb{F}_{q^m}$, $F+G$ and aF are defined as follows

$$F + G = (F_0 + G_0) D_a^0 + (F_1 + G_1) D_a^1 + \cdots + (F_d + G_d) D_a^d \in \mathbb{F}_{q^m}[D_a],$$
$$aF = (aF_0) D_a^0 + (aF_1) D_a^1 + \cdots + (aF_d) D_a^d \in \mathbb{F}_{q^m}[D_a].$$

For every skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$ of degree d the associated linear operator polynomial is

$$F^{D_a} = F_0 \,\mathrm{Id} + F_1 D_a + \cdots + F_d D_a^d \in \mathbb{F}_{q^m}[D_a].$$

An alternative notion of the evaluation of the skew polynomial F at an element β by linearly combining the powers of the operator D_a to β is

$$F^{D_a}(\beta) = F_0 D_a^0(\beta) + F_1 D_a^1(\beta) + \cdots + F_d D_a^d(\beta)$$

for all $a, \beta \in \mathbb{F}_{q^m}$. Since $a \in \mathbb{F}_{q^m}$ is also variable, there is a variety of possible evaluations of F characterized by the pair $(a, \beta) \in \mathbb{F}_{q^m}^2$. Note that the evaluation map that sends β to $F^{D_a}(\beta)$ is $\mathbb{F}_q$-linear since D_a is $\mathbb{F}_q$-linear.

The normal skew polynomial evaluations are linked to the operator evaluations as follows.

Theorem 3.6. *Given $a \in \mathbb{F}_{q^m}$, $\beta \in \mathbb{F}_{q^m}^*$ and $F \in \mathbb{F}_{q^m}[x;\sigma]$,*

$$F^{D_a}(\beta) = F(\sigma(\beta) a \beta^{-1})\beta = F({}^{\beta}a)\beta,$$

where ${}^{\beta}a$ denotes the β-conjugate of a with respect to σ (see Definition A.3).

Proof. By the product rule (Theorem A.5), it follows that $F(\sigma(\beta)a\beta^{-1})\beta = (F\beta)(a)$. Therefore, by the linearity rule (Theorem A.4) it is only necessary to show that for $F = x^i$, $x^i\beta = \sigma^i(\beta)x^i$ since $D_a^i(\beta) = \sigma^i(\beta)N_i(a)$. This follows directly from (A.5). □

For $\beta = 1$, the usual evaluation is recovered, i.e., for all $a \in \mathbb{F}_{q^m}$ and all $F \in \mathbb{F}_{q^m}[x;\sigma]$,

$$F(a) = F^{D_a}(1).$$

Example 3.7. Consider the field $\mathbb{F}_4$ with $q = 2$ and $m = 2$ and the skew polynomial $F = \bar{\omega} + x + \omega x^2$. For $a = \omega$, the corresponding linear operator polynomial is

$$F^{D_\omega} = \bar{\omega}\,\mathrm{Id} + D_\omega + \omega D_\omega^2.$$

The evaluation of this linear operator polynomial at β is

$$\begin{aligned} F^{D_\omega}(\beta) &= \bar{\omega}\,\mathrm{Id}(\beta) + D_\omega(\beta) + \omega D_\omega^2(\beta) \\ &= \bar{\omega}\beta + \beta^2 + \omega\beta^4. \end{aligned}$$

Thus, the evaluations of F^{D_ω} at $\beta \in \mathbb{F}_4^*$ are

$$\begin{aligned} F^{D_\omega}(1) &= \bar{\omega}\cdot 1 + 1^2 + \omega\cdot 1^4 = 0, \\ F^{D_\omega}(\omega) &= \bar{\omega}\omega + \omega^2 + \omega\omega^4 = 1, \\ F^{D_\omega}(\bar{\omega}) &= \bar{\omega}\bar{\omega} + \bar{\omega}^2 + \omega\bar{\omega}^4 = 1. \end{aligned}$$

Observe that $\omega + 1 = \bar{\omega}$ is an $\mathbb{F}_2$-linear combination and therefore

$$F^{D_\omega}(1+\omega) = F^{D_\omega}(1) + F^{D_\omega}(\omega) = 0 + 1 = 1,$$

which is the same as the direct evaluation above.

It is important to point out that linear operator polynomials recover linearized polynomials, which are used for Gabidulin codes, for the choice $a = 1$. The ring $\mathbb{F}_{q^m}[D_1]$ coincides with the ring of linearized polynomials, i.e., for

$$F = F_0 + F_1 x + \cdots + F_d x^d \in \mathbb{F}_{q^m}[x;\sigma]$$

with $F_i \in \mathbb{F}_{q^m}$ for $i \in [d] \cup \{0\}$, F^{D_1} is the linearized polynomial $F^{D_1} : \mathbb{F}_{q^m} \longrightarrow \mathbb{F}_{q^m}$ with

$$F^{D_1}(x) = F_0 x + F_1 x^q + \cdots + F_d x^{q^d}.$$

With the linear operator D_a, a linearized version of skew Vandermonde matrices (Definition 3.10) can be defined. The matrices were introduced in [MP17].

Definition 3.14 (**Linearized Vandermonde matrix** [MP17])**.** Let $\boldsymbol{\beta} = (\beta_1, \beta_2, \ldots, \beta_n) \in \mathbb{F}_{q^m}^n$ and $a \in \mathbb{F}_{q^m}$. Define the linearized Vandermonde matrix on $(a, \boldsymbol{\beta})$ of order d with

respect to σ as

$$\mathbf{V}_d^D(a, \boldsymbol{\beta}) = \begin{pmatrix} D_a^0(\beta_1) & D_a^0(\beta_2) & \cdots & D_a^0(\beta_n) \\ D_a^1(\beta_1) & D_a^1(\beta_2) & \cdots & D_a^1(\beta_n) \\ \vdots & \vdots & \ddots & \vdots \\ D_a^{d-1}(\beta_1) & D_a^{d-1}(\beta_2) & \cdots & D_a^{d-1}(\beta_n) \end{pmatrix} \in \mathbb{F}_{q^m}^{d \times n}.$$

Now let $\mathbf{r} = (r_1, r_2, \ldots, r_g)$ with $\sum_{i=1}^{g} r_i = n$ be a length n sum-rank partition of order g. In addition, let $\boldsymbol{\beta} = (\boldsymbol{\beta}^{(1)}, \boldsymbol{\beta}^{(2)}, \ldots, \boldsymbol{\beta}^{(g)}) \in \mathbb{F}_{q^m}^n$ with $\boldsymbol{\beta}^{(i)} \in \mathbb{F}_{q^m}^{r_i}$ for $i \in [g]$ and $\mathbf{a} \in \mathbb{F}_{q^m}^g$. The linearized Vandermonde matrix on $(\mathbf{a}, \boldsymbol{\beta})$ of order d with respect to σ is defined as

$$\mathbf{V}_d^D(\mathbf{a}, \boldsymbol{\beta}) = \left(\mathbf{V}_d^D\left(a_1, \boldsymbol{\beta}^{(1)}\right), \mathbf{V}_d^D\left(a_2, \boldsymbol{\beta}^{(2)}\right), \ldots, \mathbf{V}_d^D\left(a_g, \boldsymbol{\beta}^{(g)}\right)\right) \in \mathbb{F}_{q^m}^{k \times n}$$

with g matrices $\mathbf{V}_d^D\left(a_i, \boldsymbol{\beta}^{(i)}\right) \in \mathbb{F}_{q^m}^{k \times r_i}$ for $i \in [g]$ that are appended so that $\mathbf{V}_d^D(\mathbf{a}, \boldsymbol{\beta})$ is a $k \times n$ matrix.

The linearized Vandermonde matrix $\mathbf{V}_d^D(a, \boldsymbol{\beta})$ is again a representation of a polynomial evaluation map. Observe that if $F = F_0 + F_1 x + \cdots F_{d-1} x^{d-1}$ is a skew polynomial of degree $d-1$ with coefficients $F_0, F_1, \ldots, F_{d-1} \in \mathbb{F}_{q^m}$, then

$$(F_0, F_1, \ldots, F_{d-1}) \cdot \mathbf{V}_d^D(\boldsymbol{\beta}) = (F^{D_a}(\beta_1), F^{D_a}(\beta_2), \ldots, F^{D_a}(\beta_n)).$$

Since F^{D_a} is $\mathbb{F}_q$-linear for every $a \in \mathbb{F}_{q^m}$, it is equivalent to perform $\mathbb{F}_q$-linear transformations on the evaluations or on $(\beta_1, \beta_2, \ldots, \beta_n)$ prior to the evaluation. This important property is expressed in the following proposition.

Proposition 3.1 ([MPSK22, Prop. 2.9])**.** *Let $\mathbf{V}_d^D(a, \boldsymbol{\beta}) \in \mathbb{F}_{q^m}^{d \times n}$ be a linearized Vandermonde matrix for $a \in \mathbb{F}_{q^m}$, $\boldsymbol{\beta} \in \mathbb{F}_{q^m}^n$ and $d \in \mathbb{N}$. Moreover, let $\mathbf{A} \in \mathbb{F}_q^{n \times s}$ be any $n \times s$ matrix with entries in $\mathbb{F}_q$ for $s \in \mathbb{N}$. Then*

$$\mathbf{V}_d^D(a, \boldsymbol{\beta}) \cdot \mathbf{A} = \mathbf{V}_d^D(a, \boldsymbol{\beta} \cdot \mathbf{A})$$

holds.

Proof. Since the map $D_a^i : \mathbb{F}_{q^m} \to \mathbb{F}_{q^m}$ is $\mathbb{F}_q$-linear for all $i \in \mathbb{N}_0$ and $a \in \mathbb{F}_{q^m}$, this follows directly.

Let $a_{i,j} \in \mathbb{F}_q$ be the entry of $\mathbf{A}$ in the i-th row and j-th column and $\tilde{\boldsymbol{\beta}} = \boldsymbol{\beta} \cdot \mathbf{A} = (\tilde{\beta}_1, \tilde{\beta}_1, \ldots, \tilde{\beta}_s)$ with

$$\tilde{\beta}_j = \sum_{l=1}^{n} \beta_l a_{l,j}$$

for $j \in [s]$.

Therefore, the entry of $\mathbf{V}_d^D(a, \boldsymbol{\beta} \cdot \mathbf{A})$ in the i-th row and the j-th column is

$$D_a^{i-1}(\tilde{\beta}_j) = D_a^{i-1}\left(\sum_{l=1}^{n} \beta_l a_{l,j}\right) = \sum_{l=1}^{n} D_a^{i-1}(\beta_l) a_{l,j}.$$

The right hand side is equivalent to the entry of $\mathbf{V}_d^D(a, \boldsymbol{\beta}) \cdot \mathbf{A}$ in the i-th row and the j-th column which proves the proposition. □

The codes that are generated by such linearized Vandermonde matrices are called linearized Reed-Solomon codes. They were introduced by Martínez-Peñas in [MP17].

Definition 3.15 (**Linearized Reed-Solomon codes** [MP17, Def. 31], [MPSK22, Def. 2.14])**.** Let $\mathbf{r} = (r_1, \ldots, r_g)$ be a length n sum-rank partition of order g with $\sum_{i=1}^{g} r_i = n$, let $\mathbf{a} = (a_1, a_2, \ldots, a_g) \in \mathbb{F}_{q^m}^g$ and let $\boldsymbol{\beta} = \left(\boldsymbol{\beta}^{(1)}, \boldsymbol{\beta}^{(2)}, \ldots, \boldsymbol{\beta}^{(g)}\right) \in \mathbb{F}_{q^m}^n$ with $\boldsymbol{\beta}^{(i)} \in \mathbb{F}_{q^m}^{r_i}$ for $i \in [g]$. Furthermore, assume that

1. the elements $a_1, a_2, \ldots, a_g$ are nonzero and pairwise non-conjugate in $\mathbb{F}_{q^m}$ with respect to σ, which imposes the constraint

   $$1 \leq g \leq q - 1.$$

   Take for instance a primitive element $\gamma \in \mathbb{F}_{q^m}$ and choose $a_i = \gamma^{i-1}$ for $i \in [g]$ (Theorem A.11).

2. the vectors $\boldsymbol{\beta}^{(i)} = \left(\beta_1^{(i)}, \beta_2^{(i)}, \ldots, \beta_{r_i}^{(i)}\right) \in \mathbb{F}_{q^m}^{r_i}$ where $\beta_1^{(i)}, \beta_2^{(i)}, \ldots, \beta_{r_i}^{(i)}$ are $\mathbb{F}_q$-linearly independent for $i \in [g]$ and therefore

   $$\max\{r_1, r_2, \ldots, r_g\} \leq m.$$

   Take for example a primitive element $\gamma \in \mathbb{F}_{q^m}$ and choose $\beta_j^{(i)} = \gamma^{j-1}$ for $j \in [r_i]$.

For $k \in [n]$, the k-dimensional linearized Reed-Solomon code on $(\mathbf{a}, \boldsymbol{\beta})$ with respect to σ is given by

$$\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta}) = \left\{\mathbf{x}\mathbf{V}_k^D(\mathbf{a}, \boldsymbol{\beta}) \mid \mathbf{x} \in \mathbb{F}_{q^m}^k\right\} \subseteq \mathbb{F}_{q^m}^n$$

with the generator matrix $\mathbf{V}_k^D(\mathbf{a}, \boldsymbol{\beta}) \in \mathbb{F}_{q^m}^{k \times n}$ from Definition 3.14.

In polynomial form the k-dimensional linearized Reed-Solomon code on $(\mathbf{a}, \beta)$ is

$$\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta}) = \left\{F^{D_a}(\boldsymbol{\beta}) \mid F \in \mathbb{F}_{q^m}[x; \sigma], \deg(F) < k\right\}$$

where

$$F^{D_{a_i}}\left(\boldsymbol{\beta}^{(i)}\right) = \left(F^{D_{a_i}}\left(\beta_1^{(i)}\right), F^{D_{a_i}}\left(\beta_2^{(i)}\right), \ldots, F^{D_{a_i}}\left(\beta_{r_i}^{(i)}\right)\right) \in \mathbb{F}_{q^m}^{r_i}$$

for $i \in [g]$ and

$$F^{D_a}(\boldsymbol{\beta}) = \left(F^{D_{a_1}}\left(\boldsymbol{\beta}^{(1)}\right), F^{D_{a_2}}\left(\boldsymbol{\beta}^{(2)}\right), \ldots, F^{D_{a_g}}\left(\boldsymbol{\beta}^{(g)}\right)\right) \in \mathbb{F}_{q^m}^{n}.$$

For specific parameter choices, linearized Reed-Solomon codes recover known codes.

Theorem 3.7 ([MPSK22, Th. 2.17])**.** *Let the notation and assumption be the same as in Definition 3.15.*

1. *For $\sigma = \mathrm{Id}$ and $\mathbf{r} = (1, 1, \ldots, 1)$ (therefore $g = n$), the linearized Reed-Solomon code $\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$ is a generalized Reed-Solomon code with distinct nonzero evaluation points $\mathbf{a} \in (\mathbb{F}_q^*)^n$ and with column multipliers $\boldsymbol{\beta} \in (\mathbb{F}_q^*)^n$.*

2. *For $g = 1$, the linearized Reed-Solomon code $\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$ is a Gabidulin code with $\mathbb{F}_q$-linearly independent evaluation points $\boldsymbol{\beta} \in (\mathbb{F}_q^*)^n$.*

Thus, linearized Reed-Solomon codes coincide with generalized Reed-Solomon codes and with Gabidulin codes whenever the sum-rank metric recovers the Hamming metric and the rank metric, respectively.

Next, skew and linearized Reed-Solomon codes are related.

Theorem 3.8 ([MP17, Prop. 33])**.** *Let $\mathbf{r} = (r_1, r_2, \ldots, r_g)$ be a length n sum rank partition of order g, let $\mathbf{a} \in \mathbb{F}_{q^m}^g$ and $\boldsymbol{\beta} = \left(\boldsymbol{\beta}^{(1)}, \boldsymbol{\beta}^{(2)}, \ldots, \boldsymbol{\beta}^{(g)}\right) \in \mathbb{F}_{q^m}^n$ with $\boldsymbol{\beta}^{(i)} \in \mathbb{F}_{q^m}^{r_i}$ for $i \in [g]$. Define* $\mathrm{diag}(\boldsymbol{\beta})$ *as*

$$\mathrm{diag}(\boldsymbol{\beta}) = \begin{pmatrix} \beta_1^{(1)} & 0 & \cdots & 0 \\ 0 & \beta_2^{(1)} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \beta_{r_g}^{(g)} \end{pmatrix} \in \mathbb{F}_{q^m}^{n \times n}$$

and define

$$b_j^{(i)} = \sigma\left(\beta_j^{(i)}\right) \cdot a_i \cdot \left(\beta_j^{(i)}\right)^{-1} = {}^{\beta_j^{(i)}}a_i$$

for all $j \in [r_g]$ and for all $i \in [g]$, and where ${}^{\beta_j^{(i)}}a_i$ denotes the $\beta_j^{(i)}$-conjugate of a_i with respect to σ (see Definition A.3). Moreover, set $\mathbf{b} = \left(\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \ldots, \mathbf{b}^{(g)}\right) \in \mathbb{F}_{q^m}^n$ with $\mathbf{b}^{(i)} \in \mathbb{F}_{q^m}^{r_i}$ for $i \in [g]$. Then, it holds that

$$\mathbf{V}^D(\mathbf{a}, \boldsymbol{\beta}) = \mathbf{V}^\sigma(\mathbf{b}) \cdot \mathrm{diag}(\boldsymbol{\beta})$$

for any $d \in \mathbb{N}$. This directly implies that

$$\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta}) = \mathcal{C}_{\mathrm{SRS}}^{\sigma,k}(\mathbf{b}) \cdot \mathrm{diag}(\boldsymbol{\beta}).$$

This above theorem that if skew Reed-Solomon codes are modified by column multipliers, which are chosen $\mathbb{F}_q$-linearly independent in each group, they recover linearized Reed-Solomon codes.

The next theorem is one of the main results of the Preliminaries. It combines Theorem 3.5 and Theorem 3.8 and shows that linearized Reed-Solomon codes are MSRD codes. It was given in [MP17, Th. 4]. First, a lemma is stated to simplify the proof.

Lemma 3.2 ([MPSK22, Lem. 2.19])**.** *Let the notation be the same as in Definition 3.15. In addition, let $\mathbf{A} = \mathrm{diag}(\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_g) \in \mathbb{F}_q^{n \times n}$ be a block diagonal matrix where $\mathbf{A}_i \in \mathbb{F}_q^{r_i \times r_i}$ is an invertible matrix for $i \in [g]$. Then,*

$$\mathbf{V}^D\left(a_i, \boldsymbol{\beta}^{(i)}\right) \cdot \mathbf{A}_i = \mathbf{V}^D\left(a_i, \boldsymbol{\beta}^{(i)} \cdot \mathbf{A}_i\right) \tag{3.3}$$

holds for all $k \in [n]$ and for $i \in [g]$. Note that the components $\boldsymbol{\beta}^{(i)} \cdot \mathbf{A}_i$ are again $\mathbb{F}_q$-linearly independent since $\mathbf{A}_i$ is an invertible matrix. As a result,

$$\mathbf{V}^D(\mathbf{a}, \boldsymbol{\beta}) \cdot \mathbf{A} = \mathbf{V}^D(\mathbf{a}, \boldsymbol{\beta} \cdot \mathbf{A})$$

and

$$\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta}) \cdot \mathbf{A} = \mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta} \cdot \mathbf{A}).$$

Proof. From Proposition 3.1, Equation (3.3) follows and the rest can be easily deduced from that. □

Theorem 3.9 ([MP17, Th. 4])**.** *The linearized Reed-Solomon code $\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$ with dimension k, as defined in Definition 3.15, is an MSRD code. In other words, the minimum sum-rank distance of the linearized Reed-Solomon code $\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$ satisfies*

$$\mathrm{d}_{\mathrm{SR}}\left(\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})\right) = n - k + 1.$$

Proof. From Theorem 3.8 and Theorem 3.5, it can be deduced that the linearized Reed-Solomon code $\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$ is an MDS code. By Definition 3.8, it remains to show that $\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta}) \cdot \mathbf{A}$ is again an MDS code for all block diagonal matrices $\mathbf{A} = \mathrm{diag}(\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_g) \in \mathbb{F}_q^{n \times n}$ with $\mathbf{A}_i \in \mathbb{F}_q^{r_i \times r_i}$ invertible for $i \in [g]$. By Lemma 3.2,

$$\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta}) \cdot \mathbf{A} = \mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta} \cdot \mathbf{A})$$

where the components $\boldsymbol{\beta}^{(i)} \cdot \mathbf{A}_i$ are $\mathbb{F}_q$-linearly independent for $i \in [g]$ since $\mathbf{A}_i$ is an invertible matrix. Thus, the code $\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta}) \cdot \mathbf{A}$ is an MDS code since it is again a linearized Reed-Solomon code. Therefore, $\mathcal{C}_{\mathrm{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$ is an MSRD code since $\mathbf{A}$ is an arbitrary invertible block diagonal matrix over $\mathbb{F}_q$. □

3.6 Secret Sharing with Skew Polynomials

Two secret sharing schemes with conventional polynomials were discussed in Section 3.1. This section briefly shows that secret sharing can also be realized with skew polynomials. In [Zha10], skew polynomials were already used to construct a secret sharing scheme, but with only one message symbol. The following construction introduces a McEliece-Sarwate type secret sharing scheme for multiple message symbols.

Construction 3.3 (**Secret sharing with skew polynomials**)**.** Fix the following integers n, k, $z < k$ and let $\mathbb{F}_{q^m}$ be an extension field of the prime power q of degree m with $(q-1)m \geq n$. Let $\Omega = \{a_1, a_2, \ldots, a_n\} \subseteq \mathbb{F}_{q^m}$ be a P-independent set with $|\Omega| = n$. Given $k - z$ message symbols $m_1, \ldots, m_{k-z} \in \mathbb{F}_{q^m}$, generate z random numbers $r_1, \ldots, r_z$ independently and uniformly distributed over $\mathbb{F}_{q^m}$. The shares $s_1, \ldots, s_n$ can be calculated as evaluations of the skew polynomial

$$F = r_1 + r_2 x + \cdots + r_z x^{z-1} + m_1 x^z + \cdots + m_{k-z} x^{k-1} = \sum_{j=1}^{z} r_j x^{j-1} + \sum_{i=1}^{k-z} m_i x^{i+z-1}$$

with $s_i = F(a_i)$ for $a_i \in \Omega$ with $i \in [n]$.

Proposition 3.2 (**Skew polynomial secret sharing**)**.** *The secret sharing scheme with skew polynomials presented above is an (n, k, z) secret sharing scheme.*

Proof. Decodability: Given any k shares, Lagrange interpolation can be used to recover the polynomial $\tilde{F} = F$ of degree $k-1$ (Theorem A.7). The message can then be retrieved by taking the coefficients $\tilde{F}_i$ of $\tilde{F}$ for $i \in \{k-z, k-z+1, \ldots, k-1\}$.
Privacy: Lemma 3.1 is used and the notation is the same as in Theorem 3.1. As R represents z independently and uniformly at random generated numbers $r_1, \ldots, r_z \in \mathbb{F}_{q^m}$, $\mathrm{H}(\mathsf{S}_{\mathcal{I}_z}) \leq \mathrm{H}(\mathsf{R})$ holds. It remains to show that $\mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{S}_{\mathcal{I}_z}) = 0$. Consider the polynomial F

$$F = \underbrace{\sum_{j=1}^{z} r_j x^{j-1}}_{F_r} + \underbrace{\sum_{i=1}^{k-z} m_i x^{i+z-1}}_{F_m}.$$

Given $m_1, m_2, \dots m_{k-z}$ and $s_i = F(a_i)$ for $i \in \mathcal{I}_z$, F_r of degree $z-1$ can be retrieved using Lagrange interpolation with the shares $\hat{s}_i = F_r(a_i)$ that can be calculated as follow

$$F_r(a_i) = F(a_i) - F_m(a_i)$$

which holds by the linearity rule (Theorem A.4). The random numbers $r_1, \dots, r_z$ are the coefficients of the polynomial F_r. Thus, $\mathrm{H}(\mathsf{R} \mid \mathsf{M}, \mathsf{S}_{\mathcal{I}_z}) = 0$. □

Note that this secret sharing scheme recovers the McEliece-Sarwate secret sharing scheme which uses conventional polynomials for $\sigma = \mathrm{Id}$ (Construction 3.2). In this case the P-independent set consists of n distinct elements. As for the McEliece-Sarwate secret sharing scheme, which is related to Reed-Solomon codes, the skew polynomial secret sharing scheme is related to skew Reed-Solomon codes encoded with a skew Vandermonde matrix.

This secret sharing scheme can be slightly adjusted so that it uses linear operator polynomials which also recover linearized polynomials.

Construction 3.4 (**Secret sharing with linear operator polynomials**)**.** Let $\mathbf{r} = (n_1, \dots, n_g)$ be a length n sum-rank partition of order g such that $\sum_{i=1}^{g} n_i = n$. Fix the following integers $k \leq n$, $z < k$ and let $\mathbb{F}_{q^m}$ be an extension field of the prime power q of degree m with $(q-1)m \geq n$. Let $\Omega = \{a_1, a_2, \dots, a_g\} \subseteq \mathbb{F}_{q^m}$ be a set of nonzero and pairwise non-conjugate elements in $\mathbb{F}_{q^m}$ with $|\Omega| = g$. Let $\boldsymbol{\beta} = \left(\boldsymbol{\beta}^{(1)}, \dots, \boldsymbol{\beta}^{(g)}\right) \in \mathbb{F}_{q^m}^n$ with $\boldsymbol{\beta}^{(i)} \in \mathbb{F}_{q^m}^{n_i}$ where $\boldsymbol{\beta}^{(i)} \in \mathbb{F}_{q^m}^{n_i}$ consists of $\mathbb{F}_q$-linearly independent elements. Given $k-z$ message symbols $m_1, \dots, m_{k-z} \in \mathbb{F}_{q^m}$, generate z random numbers $r_1, \dots, r_z$ independently and uniformly distributed over $\mathbb{F}_{q^m}$. The shares $s_1, \dots, s_n$ can be calculated as evaluations of the linear operator polynomial

$$F^{D_a} = r_1 \,\mathrm{Id} + r_2 D_a + \cdots + r_z D_a^{z-1} + m_1 D_a^z + \cdots + m_{k-z} D_a^{k-1} = \sum_{j=1}^{z} r_j D_a^{j-1} + \sum_{i=1}^{k-z} m_i D_a^{i+z-1}$$

with $s_{i,j} = F^{D_{a_i}}(\beta_j^{(i)})$ for $a_i \in \Omega$ with $i \in [g]$ and $j \in [n_i]$.

Proposition 3.3 (**Linear operator polynomial secret sharing**). *The secret sharing scheme with linear operator polynomials presented above is an* (n, k, z) *secret sharing scheme.*

Proof. By Theorem 3.8, the secret sharing scheme coincides with Construction 3.3 for the P-independent set Ω with elements $b_j^{(i)} = {}^{\beta_j^{(i)}}a_i$ for all $j \in [n_i]$ and $i \in [g]$, where each share $s_{i,j} = F(b_j^{(i)})\beta_j^{(i)}$ is an evaluation of the skew polynomial $F = \sum_{j=1}^{z} r_j x^{j-1} + \sum_{i=1}^{k-z} m_i x^{i+z-1}$. □

The linear operator polynomial secret sharing scheme recovers conventional polynomial secret sharing for $r_1 = \cdots = r_g = 1$, $\sigma = \mathrm{Id}$ and $g = n$ where the elements in Ω are distinct elements. It also recovers secret sharing with linearized polynomials for $g = 1$, $r_1 = n$ and $m \geq n$ with n $\mathbb{F}_q$-linearly independent elements $\beta_j^{(1)}$ for $j \in [n]$. The two cases follow directly from Theorem 3.7 and the fact that the secret sharing scheme is an encoding of a linearized Reed-Solomon code or in other words the multiplication of the vector $\mathbf{u} = (r_1, \ldots, r_z, m_1, \ldots, m_{k-z}) \in \mathbb{F}_{q^m}^k$ with a linearized Vandermonde matrix.

3.7 Skew Lagrange Polynomials

In Appendix A, a Newton interpolation algorithm was discussed (see Definition A.8). Given a P-independent set of n evaluation points $a_1 \ldots, a_n$ with corresponding values $b_1, \ldots, b_n$, it returns the unique skew polynomial G with $\deg(G) < n$ which fulfills the conditions $G(a_i) = b_i$ for all $i \in [n]$. The following section covers a construction of Lagrange type polynomials $\ell_i(x)$ with $\ell_i(a_i) = 1$ and $\ell_i(a_j) = 0$ for $i \neq j$. Skew polynomials fulfilling the Lagrange conditions were discussed in [Zha10].

For conventional polynomials, Lagrange polynomials can be constructed with a simple formula (Definition A.5). Given n evaluation constraints $F(a_i) = b_i$, the polynomial F of degree less than n fulfilling these constraints can be constructed as

$$F(x) = \sum_{i=1}^{n} b_i \ell_i(x)$$

with ℓ_i defined as

$$\ell_i(x) = \prod_{\substack{0 \leq m < n \\ m \neq i}} \frac{x - a_m}{a_i - a_m}.$$

For skew polynomials, the construction of Lagrange type polynomials is more complicated and an easy term such as the one for conventional polynomials does not yet exist to the best of our knowledge. However, the construction can be done by Newton

interpolation for which an algorithm for skew polynomials is already known. This idea was introduced in [Zha10].

Definition 3.16 (**Skew Lagrange polynomials**)**.** Let $\Omega = \{a_1, a_2, \ldots, a_k\} \subseteq \mathbb{F}_{q^m}$ be a P-independent set with cardinality $|\Omega| = k$. The skew Lagrange polynomial ℓ_i fulfilling the constraints $\ell_i(a_i) = 1$ and $\ell_i(a_j) = 0$ can be constructed using the Newton interpolation algorithm for skew polynomials (Definition A.8). The i-th Lagrange polynomial is also written as ℓ_i^{Ω} to indicate the set it is constrained on.

As for conventional polynomials, every skew polynomial can be written in a Lagrange basis instead of the usual monomial basis. The following definition shows the transformation between monomial and Lagrange basis similarly as in [Gan05].

Definition 3.17 (**Skew polynomials - monomial and Lagrange basis**)**.** Let $F = F_0 + F_1 x + \ldots + F_{k-1} x^{k-1} \in \mathbb{F}_{q^m}[x;\sigma]$ be a skew polynomial of degree $k-1$ in monomial basis with coefficient vector $\mathbf{f} = (F_0, F_1, \ldots, F_{k-1}) \in \mathbb{F}_{q^m}^k$. Let $\Omega = \{a_0, a_1, \ldots, a_{k-1}\} \subseteq \mathbb{F}_{q^m}$ be a P-independent set with cardinality $|\Omega| = k$ and $\Phi = \{p_0, p_1, \ldots, p_{k-1}\} \subseteq \mathbb{F}_{q^m}$ the set of evaluations of F such that $F(a_i) = p_i$ for $i = 0, 1, \ldots, k-1$ and its vector representation $\mathbf{p} = (p_0, p_1, \ldots, p_{k-1}) \in \mathbb{F}_{q^m}^k$. Furthermore, let $\mathcal{L} = \{\ell_0, \ell_1, \ldots, \ell_{k-1}\}$ be a Lagrange basis on Ω as defined in Definition 3.16, so that the skew polynomial F can be written as

$$F = p_0\ell_0 + p_1\ell_1 + \ldots + p_{k-1}\ell_{k-1}.$$

Thus, we have two representations

$$F = \mathbf{f} \cdot \mathbf{m}(x) = \mathbf{p} \cdot \boldsymbol{\ell}(x)$$

where $\mathbf{m}(x) = (1, x, \ldots, x^{k-1})^{\mathsf{T}}$ and $\boldsymbol{\ell}(x) = (\ell_0, \ell_1, \ldots, \ell_{k-1})^{\mathsf{T}}$. By Definition 3.10,

$$\mathbf{f}\mathbf{V}_k^{\sigma}(\mathbf{a}) = \mathbf{p} \tag{3.4}$$

holds with $\mathbf{V}_k^{\sigma}(\mathbf{a})$ being the $k \times k$ skew Vandermonde matrix on the vector $\mathbf{a} = (a_0, a_1, \ldots, a_{k-1})$. Therefore, the transformation between monomial and Lagrange basis is

$$\mathbf{m}(x) = \mathbf{V}_k^{\sigma}(\mathbf{a})\boldsymbol{\ell}(x)$$

with an invertible transformation matrix.

With these two basis of skew polynomials in mind, a lemma is given, which deals with the rank of a matrix with a special structure. It is an essential tool deriving secrecy capacities in Chapter 6.

Lemma 3.3. *Let $F = F_0 + F_1 x + \ldots + F_{k-1}x^{k-1} \in \mathbb{F}_{q^m}[x;\sigma]$ be a skew polynomial of degree $k-1$ in monomial basis with coefficient vector $\mathbf{f} = (F_0, F_1, \ldots, F_{k-1}) \in \mathbb{F}_{q^m}^k$. Let $\Omega = \{a_0, a_1, \ldots, a_{n-1}\} \subseteq \mathbb{F}_{q^m}$ be a P-independent set with cardinality $|\Omega| = n$. Let the sets Ω be split into two subsets $\Omega_k = \{a_i \mid i \in 0 \cup [k-1]\}$ and $\Omega_d = \{a_i \mid i \in [n-1] \setminus [k-1]\}$ with $d = n - k$. Let $\mathcal{L} = \{\ell_0, \ell_1, \ldots, \ell_{k-1}\}$ be a Lagrange basis on Ω_k as defined in Definition 3.16. The matrix*

$$\mathbf{M} = \begin{pmatrix} \ell_0^{\Omega_k}(a_k) & \ell_1^{\Omega_k}(a_k) & \cdots & \ell_{k-1}^{\Omega_k}(a_k) \\ \ell_0^{\Omega_k}(a_{k+1}) & \ell_1^{\Omega_k}(a_{k+1}) & \cdots & \ell_{k-1}^{\Omega_k}(a_{k+1}) \\ \vdots & \vdots & \ddots & \vdots \\ \ell_0^{\Omega_k}(a_{n-1}) & \ell_1^{\Omega_k}(a_{n-1}) & \cdots & \ell_{k-1}^{\Omega_k}(a_{n-1}) \end{pmatrix}$$

has full rank, i.e., $\operatorname{rank}(\mathbf{M}) = \min(k, d)$.

Proof. It is shown that the matrix $\mathbf{M}$ has full rank by decomposing it into several matrices that are proven to have full rank. By Definition 3.10, it holds that

$$\mathbf{M} = (\mathbf{V}_k^\sigma(\mathbf{a}_d))^\mathsf{T} \underbrace{\begin{pmatrix} \ell_{0,0}^{\Omega_k} & \ell_{1,0}^{\Omega_k} & \cdots & \ell_{k-1,0}^{\Omega_k} \\ \ell_{0,1}^{\Omega_k} & \ell_{1,1}^{\Omega_k} & \cdots & \ell_{k-1,1}^{\Omega_k} \\ \vdots & \vdots & \ddots & \vdots \\ \ell_{0,k-1}^{\Omega_k} & \ell_{1,k-1}^{\Omega_k} & \cdots & \ell_{k-1,k-1}^{\Omega_k} \end{pmatrix}}_{=:\mathbf{L}}$$

with $\mathbf{a}_d = (a_k, \ldots, a_{n-1})$ and where $\mathbf{L}$ is the monomial representation of the Lagrange skew polynomials with $\ell_{i,j}$ being the j-th coefficient of the i-th polynomial. The transformation from Lagrange to monomial basis can be achieved by multiplying with the inverse of the Vandermonde matrix $(\mathbf{V}_k^\sigma)^\mathsf{T}$ (see Equation (3.4)). Thus, it holds that $\mathbf{L} = ((\mathbf{V}_k^\sigma(\mathbf{a}_k))^\mathsf{T})^{-1}$ with $\mathbf{a}_k = (a_0, \ldots, a_{k-1})$. Overall, we have

$$\mathbf{M} = (\mathbf{V}_k^\sigma(\mathbf{a}_d))^\mathsf{T}((\mathbf{V}_k^\sigma(\mathbf{a}_k))^\mathsf{T})^{-1}.$$

By Theorem 3.4, both matrices are full rank matrices and it holds that $\operatorname{rank}(\mathbf{M}) = \min(k, d)$. □

Remark: The above lemma also implies that submatrices of $\mathbf{M}$ have full rank since they are also a product of two Vandermonde matrices.

3.8 Information Theory

We turn now to information-theoretic concepts needed for this work. Further relevant equalities are summarized in Section B.1. Here, only one important consideration is made that is later used to bound the entropy of a random vector from above.

Lemma 3.4. *Let* $\mathsf{K}^k = (\mathsf{K}_1, \mathsf{K}_2, \ldots, \mathsf{K}_k) \in \mathbb{K}^k$ *be a random vector and* $f : \mathbb{K}^k \longrightarrow \mathbb{K}^{n_f}$ *a function. It holds that*

$$\mathrm{H}(f(\mathsf{K}^k)) \leq \mathrm{H}(\mathsf{K}^k).$$

Proof. Applying the chain rule (B.1) to the joint entropy $\mathrm{H}(\mathsf{K}^k, f(\mathsf{K}^k))$ yields

$$\begin{aligned}\mathrm{H}(\mathsf{K}^k, f(\mathsf{K}^k)) &= \mathrm{H}(\mathsf{K}^k) + \mathrm{H}(f(\mathsf{K}^k) \mid \mathsf{K}^k) \\ &= \mathrm{H}(f(\mathsf{K}^k)) + \mathrm{H}(\mathsf{K}^k \mid f(\mathsf{K}^k)).\end{aligned}$$

Since K^k essentially determines $f(\mathsf{K}^k)$, it holds that $\mathrm{H}(f(\mathsf{K}^k) \mid \mathsf{K}^k) = 0$.
With $\mathrm{H}(\mathsf{K}^k \mid f(\mathsf{K}^k)) \geq 0$, we have

$$\mathrm{H}(f(\mathsf{K}^k)) \leq \mathrm{H}(\mathsf{K}^k)$$

with equality if, and only if, f is bijective. □

Consider the random vectors $\mathsf{K}^k \in \mathbb{K}^k$, $\mathsf{X}^{n_x} \in \mathbb{K}^{n_x}$ and $\mathsf{Y}^{n_y} \in \mathbb{K}^{n_y}$ such that

$$\begin{aligned}\mathsf{X}^{n_x} &= \mathbf{A}(\mathsf{K}^k)^\mathsf{T} \\ \mathsf{Y}^{n_y} &= \mathbf{B}(\mathsf{K}^k)^\mathsf{T},\end{aligned} \tag{3.5}$$

where $\mathbf{A} \in \mathbb{K}^{n_x \times k}$ and $\mathbf{B} \in \mathbb{K}^{n_y \times k}$.

Lemma 3.5. *For two random vectors* X^{n_x} *and* Y^{n_y}*, that have the above described properties, it holds that*

$$\begin{aligned}\mathrm{H}(\mathsf{X}^{n_x}) &\leq k, \\ \mathrm{H}(\mathsf{Y}^{n_y}) &\leq k.\end{aligned}$$

Proof. The proofs for X^{n_x} and Y^{n_y} follow the same arguments. Therefore, it is only shown that $\mathrm{H}(\mathsf{X}^{n_x}) \leq k$ and $\mathrm{H}(\mathsf{Y}^{n_y}) \leq k$ follows analogously.
We know that

$$\mathrm{H}(\mathsf{K}^k, \mathsf{X}^{n_x}) = \mathrm{H}(\mathsf{X}^{n_x}) + \mathrm{H}(\mathsf{K}^k \mid \mathsf{X}^{n_x})$$

and thus it holds that

$$\mathrm{H}(\mathsf{X}^{n_x}) \leq \mathrm{H}(\mathsf{K}^k, \mathsf{X}^{n_x})$$

with equality if, and only if, X^{n_x} essentially determines K^k, i.e., $\mathrm{H}(\mathsf{K}^k \mid \mathsf{X}^{n_x}) = 0$. Furthermore, it holds that $\mathsf{X}^{n_x} = f(\mathsf{K}^k) = \mathbf{A}(\mathsf{K}^k)^\mathsf{T}$ which yields

$$\mathrm{H}(\mathsf{K}^k, \mathsf{X}^{n_x}) = \mathrm{H}(\mathsf{K}^k, f(\mathsf{K}^k)) = \mathrm{H}(\mathsf{K}^k, \mathbf{A}(\mathsf{K}^k)^\mathsf{T}) = \mathrm{H}(\mathsf{K}^k) \leq k.$$

As a result,

$$\mathrm{H}(\mathsf{X}^{n_x}) = \mathrm{H}(\mathbf{A}(\mathsf{K}^k)^\mathsf{T}) \leq \mathrm{H}(\mathsf{K}^k) \leq k$$

holds and we are done. □

Now let us assume that the random variable K^k consists of independent and uniformly distributed symbols over $\mathbb{K}$, such that $\mathrm{H}(\mathsf{K}^k) = k$ holds. The entropy of the random variables X^{n_x} and Y^{n_y} and the joint entropy can then be expressed in terms of the rank of the matrices $\mathbf{A}$ and $\mathbf{B}$.

Lemma 3.6. *Let $\mathsf{K}^k \in \mathbb{K}^k$, $\mathsf{X}^{n_x} \in \mathbb{K}^{n_x}$ and $\mathsf{Y}^{n_y} \in \mathbb{K}^{n_y}$ be three random vectors as defined in (3.5). Furthermore, let the random vector K^k consist of independent and uniformly distributed symbols over $\mathbb{K}$. The entropies of the random vectors X^{n_x} and Y^{n_y} are*

$$\mathrm{H}(\mathsf{X}^{n_x}) = \mathrm{rank}(\mathbf{A})$$

and

$$\mathrm{H}(\mathsf{Y}^{n_y}) = \mathrm{rank}(\mathbf{B}).$$

Proof. The proof is only shown for X^{n_x} and the proof for Y^{n_y} follows analogously. We know that $\mathrm{H}(\mathsf{K}^k) = k$. The entropy of X^{n_x} is

$$\mathrm{H}(\mathsf{X}^{n_x}) = \mathrm{H}(\mathbf{A}(\mathsf{K}^k)^\mathsf{T}) \leq k.$$

The rank of the matrix $\mathbf{A}$ determines how many symbols of X^{n_x} are independent and this can be expressed by $\mathrm{H}(\mathsf{X}^{n_x}) = \mathrm{rank}(\mathbf{A})$. □

For the conditional entropy of X^{n_x} given Y^{n_y} a similar expression can be derived.

Lemma 3.7. *Let $\mathsf{K}^k \in \mathbb{K}^k$, $\mathsf{X}^{n_x} \in \mathbb{K}^{n_x}$ and $\mathsf{Y}^{n_y} \in \mathbb{K}^{n_y}$ be three random vectors as defined in (3.5). Furthermore, let the random vector K^k consist of independent and uniformly distributed symbols over $\mathbb{K}$. The conditional entropy of X^{n_x} given Y^{n_y} is*

$$\mathrm{H}(\mathsf{X}^{n_x} | \mathsf{Y}^{n_y}) = \mathrm{H}(\mathsf{X}^{n_x}, \mathsf{Y}^{n_y}) - \mathrm{H}(\mathsf{Y}^{n_y}) = \mathrm{rank}(\mathbf{C}) - \mathrm{rank}(\mathbf{B}), \tag{3.6}$$

where $\mathbf{C} = \begin{pmatrix} \mathbf{A} \\ \mathbf{B} \end{pmatrix} \in \mathbb{K}^{(n_x+n_y)\times k}$ *is the stacked matrix of $\mathbf{A}$ and $\mathbf{B}$.*

Proof. The first equality $\mathrm{H}(\mathsf{X}^{n_x}|\mathsf{Y}^{n_y}) = \mathrm{H}(\mathsf{X}^{n_x}, \mathsf{Y}^{n_y}) - \mathrm{H}(\mathsf{Y}^{n_y})$ follows from the chain rule of entropy (B.1). From Lemma 3.6, we know that $\mathrm{H}(\mathsf{Y}^{n_y}) = \mathrm{rank}(\mathbf{B})$. It remains to show that $\mathrm{H}(\mathsf{X}^{n_x}, \mathsf{Y}^{n_y}) = \mathrm{rank}(\mathbf{C})$ holds. Plugging in the definition of X^{n_x} and Y^{n_y} yields

$$\mathrm{H}(\mathsf{X}^{n_x}, \mathsf{Y}^{n_y}) = \mathrm{H}(\mathbf{A}(\mathsf{K}^k)^\mathsf{T}, \mathbf{B}(\mathsf{K}^k)^\mathsf{T}).$$

The two random vectors can then be stacked since they are both expressed in terms of K^k, which yields

$$\mathrm{H}(\mathbf{A}(\mathsf{K}^k)^\mathsf{T}, \mathbf{B}(\mathsf{K}^k)^\mathsf{T}) = \mathrm{H}(\mathbf{C}(\mathsf{K}^k)^\mathsf{T}),$$

where $\mathbf{C} = \begin{pmatrix} \mathbf{A} \\ \mathbf{B} \end{pmatrix} \in \mathbb{K}^{n_x+n_y}$ is the stacked matrix of $\mathbf{A}$ and $\mathbf{B}$. We can then apply Lemma 3.6 to receive $\mathrm{H}(\mathbf{C}(\mathsf{K}^k)^\mathsf{T}) = \mathrm{rank}(\mathbf{C})$ and we are done. □

The idea on how the above derived bounds are applied is illustrated with an example.

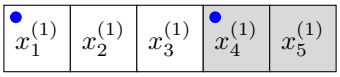

Figure 3.8: Illustration of one group of a DSS which uses an LRC with $(3, 3)$-locality. The two blue dots indicate an eavesdropper that is observing two nodes of the group, i.e., $l_1 = 2$ and $l_2 = 0$.

Example 3.8. Consider one group of a DSS with $r = 3$ and $\delta = 3$ and let an eavesdropper observe two nodes of the group such that $l_1 = 2$ and $l_2 = 0$ as shown in Figure 3.8. Since any 3 of the 5 nodes need to be downloaded to recover the data, $\mathcal{K} = \{x_1^{(1)}, x_2^{(1)}, x_3^{(1)}\}$ is a possible set for a data collector denoted by K. It is assumed that the stored data are uniformly distributed and independent. The symbols observed by the eavesdropper are a realization of the random vector E. If we want to calculate $\mathrm{H}(\mathsf{K} \mid \mathsf{E})$, Equation (3.6) can be used. First, write E and K in their matrix form in terms of K:

$$\mathsf{K} = \mathbf{K} \cdot \mathsf{K}^\mathsf{T} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \mathsf{K}^\mathsf{T}$$

and

$$\mathsf{E} = \mathbf{E} \cdot \mathsf{K}^\mathsf{T} = \begin{pmatrix} 1 & 0 & 0 \\ \ell_{1,1}^{\mathcal{K}}(\alpha_4^{(1)}) & \ell_{1,2}^{\mathcal{K}}(\alpha_4^{(1)}) & \ell_{1,3}^{\mathcal{K}}(\alpha_4^{(1)}) \end{pmatrix} \mathsf{K}^\mathsf{T}$$

where $\alpha_j^{(i)}$ denotes the code locator of the j-th node in the i-th group and $\ell_{i,j}^{\mathcal{K}}$ the Lagrange polynomial on $\mathcal{K}$ which is 1 at $\alpha_j^{(i)}$.

By Lemma 3.7, it holds that

$$\mathrm{H}(\mathsf{K} \mid \mathsf{E}) = \mathrm{H}(\mathsf{K}, \mathsf{E}) - \mathrm{H}(\mathsf{E}) = \operatorname{rank}\begin{pmatrix}\mathbf{K}\\ \mathbf{E}\end{pmatrix} - \operatorname{rank}(\mathbf{E}) = 3 - 2 = 1.$$

This means that the eavesdropper needs to observe one more symbol to be able to recover all the data stored in the group.

The data collector $\bar{\mathcal{K}} = \{x_1^{(1)}, x_2^{(1)}, x_4^{(1)}\}$ is another possible choice. It makes the representation of E easier with

$$\bar{\mathbf{E}} = \begin{pmatrix}1 & 0 & 0\\ 0 & 0 & 1\end{pmatrix}.$$

When calculating the conditional entropy with (3.6), it is therefore beneficial to think about the choice of the data collector first making the calculations of the rank of the corresponding matrices as easy as possible.

4 Related Work

In this chapter, two MR-LRC constructions from [MPK19] and [RKSV14] are given and it is shown that they are achieving a Singleton-type bound.

A secrecy capacity for LRCs that was introduced in [RKSV14] is summarized in Section 4.3. This secrecy capacity only considers an (l_1, l_2)-eavesdropper threat without the possibility of observing global repairs. It is the starting point for Chapter 6, which derives the decrease of the capacity due to global repairs.

4.1 Bounds on Maximally Recoverable Locally Repairable Codes

In the literature, "optimal" LRCs are defined as codes whose Hamming distance is as large as possible and attains a Singleton-type bound, given the locality parameters. Such a general Singleton-like bound was given in [GHSY12] for $\delta = 2$, in [PKLK12] for arbitrary δ, and in [RKSV14] for vector codes with arbitrary δ. Vector codes are codes where each symbol is in $\mathbb{F}_{q^m}^{\kappa}$ instead of in $\mathbb{F}_{q^m}$. For scalar LRCs, it reads as follows.

Proposition 4.1. *Let $\mathcal{C} \subseteq \mathbb{F}^N$ be an (r, δ)-locality LRC, as in Definition 3.3. The dimension of the code is $k = \log_{|\mathbb{F}|} |\mathcal{C}|$, which is assumed to be an integer. The Hamming distance of the code $\mathcal{C}$ is bounded by*

$$\begin{aligned} \mathrm{d_H}(\mathcal{C}) &\leq N - k + 1 - \left(\left\lceil \frac{k}{r} \right\rceil - 1 \right) (\delta - 1) \\ &= h + \left(\left\lfloor \frac{h}{r} \right\rfloor + 1 \right) (\delta - 1) + 1, \end{aligned} \tag{4.1}$$

where $h = gr - k = N - g(\delta - 1) - k$.

The proof is omitted for brevity and can be found in [RKSV14] in a more general form for vector codes.

The bound is attained by any MR-LRC which was shown in [MPSK22, Th. 3.3].

Theorem 4.1 ([MPSK22, Th. 3.3])**.** *Let $\mathcal{C} \subseteq \mathbb{F}^N$ be a linear MR-LRC with (r, δ)-localities as in Definition 3.4. Then $\mathcal{C}$ has optimal Hamming distance with respect to the bound given by (4.1), i.e.,*

$$\mathrm{d_H}(\mathcal{C}) = h + \left(\left\lfloor \frac{h}{r} \right\rfloor + 1 \right) (\delta - 1) + 1, \tag{4.2}$$

where $k = \dim(\mathcal{C})$ and $h = gr - k = N - g(\delta - 1) - k$.

Proof. Let $d = \mathrm{d_H}(\mathcal{C})$. Moreover, let $\mathcal{E} \subseteq [N]$ be an erasure pattern such that $|\mathcal{E}| = d$. $\mathcal{E}$ is therefore a non-correctable erasure pattern with smallest possible cardinality. This means that if an erasure is removed from $\mathcal{E}$, then the pattern is correctable by $\mathcal{C}$. It either holds that $|\mathcal{E} \cap \Gamma_i| = 0$ or that $|\mathcal{E} \cap \Gamma_i| \geq \delta$ for $i \in [g]$, since assuming that $|\mathcal{E} \cap \Gamma_i| < \delta$, would imply that the erasures can be repaired locally. However, the remaining erasures may not be correctable, so an erasure pattern with smaller $|\mathcal{E}|$ which is still not correctable could be found. This would be a contradiction to the assumed minimality of the erasure pattern. Denote the number of affected groups, i.e., the number of groups where $|\mathcal{E} \cap \Gamma_i| \geq \delta$, as κ. From Definition 3.4, it can be deduced that $\mathcal{E}$ contains $\delta - 1$ elements per affected local group and in addition at least $h + 1$ elements anywhere in the κ affected groups. Thus,

$$|\mathcal{E}| \geq \kappa(\delta - 1) + h + 1$$

and moreover, $|\mathcal{E}|$ is upper bounded as follows

$$\kappa(\delta - 1) + h + 1 \leq |\mathcal{E}| \leq \kappa(r + \delta - 1)$$

since each affected local group has size $r + \delta - 1$. As a result,

$$\kappa \geq \left\lceil \frac{h+1}{r} \right\rceil = \left\lfloor \frac{h}{r} \right\rfloor + 1$$

holds. With (4.2), it can be concluded

$$d = |\mathcal{E}| \geq \kappa(\delta - 1) + h + 1 \geq \left(\left\lfloor \frac{h}{r} \right\rfloor + 1 \right) (\delta - 1) + h + 1$$

and the proof is complete. □

Therefore, MR-LRCs are optimal LRCs with respect to (4.1). Explicit MR-LRC constructions are given in the next section.

4.2 Maximally Recoverable Locally Repairable Code Constructions

In this section, constructions of maximally recoverable locally repairable codes (MR-LRCs) for DSSs are presented. MR-LRCs are able to tolerate the information theoretical maximum of erasures given the parameters of the system (Definition 3.4).

In [RKSV14], an MR-LRC construction is introduced using Gabidulin codes. The construction is given in a simplified version for a scalar LRC instead of a vector LRC.

Construction 4.1 ([RKSV14, Constr. I])**.** Let N, r and δ be positive integers such that $r+\delta-1 < N$ and $q \geq (r+\delta-1)$. Consider an information vector $\mathbf{u} \in \mathbb{F}_{q^m}^k$ with $k \geq r$. Let $g = \left\lceil \frac{N}{r+\delta-1} \right\rceil$ be the number of groups. There are two cases depending on whether or not $r+\delta-1$ divides N.

Case 1 ($(r+\delta-1)|N$): Let $n = gr$, $m \geq n$ and $\mathcal{C}_{\text{out}}$ be an $[n, k, D = n-k+1]_{q^m}$ Gabidulin code. The encoding follows two steps:

1. Encode $\mathbf{u}$ with the Gabidulin code yielding $\mathbf{c}_{\text{out}} \in \mathcal{C}_{\text{out}}$ and partition $\mathbf{c}_{\text{out}}$ into $g = \frac{n}{r}$ disjoint groups.

2. Apply an $[(r+\delta-1), r, \delta]_q$ MDS code on each local group with r nodes to generate $\delta-1$ parities, respectively.

Case 2 ($(r+\delta-1) \nmid N$): Let $t \in [r-1]$ be an integer such that $N = (g-1)(r+\delta-1)+(t+\delta-1)$. Let $n = (g-1)r+t$, $m \geq N$ and $\mathcal{C}_{\text{out}}$ be an $[n, k, D = n-k+1]_{q^m}$ Gabidulin code. The encoding follows two steps:

1. Encode $\mathbf{u}$ with the Gabidulin code yielding $\mathbf{c}_{\text{out}} \in \mathcal{C}_{\text{out}}$ and partition $\mathbf{c}_{\text{out}}$ into $g-1$ disjoint groups of size r and one additional group of size t.

2. Apply an $[(r+\delta-1), r, \delta]_q$ MDS code on the first $g-1$ local groups with r nodes and an $[(t+\delta-1), t, \delta]_q$ MDS code on the last local group with t nodes to generate $\delta-1$ parities.

The encoding steps are illustrated in Figure 4.1. A proof that the construction above yields an MR-LRC is given in Theorem 4.2 for a more generalized construction which is equivalent for specific parameters.

Example 4.1 (MR-LRC with Gabidulin code)**.** Consider a DSS with $g = 3$ groups where an information vector of length 7 is stored in a maximally recoverable manner. The system has $h = 2$ global parities and each local group is able to tolerate one erasure locally, i.e., $\delta = 2$, $r = 3$. The required field size is therefore $q \geq r+\delta-1 = 4$ and $m \geq n = 9$. The encoding is illustrated in Figure 4.2.

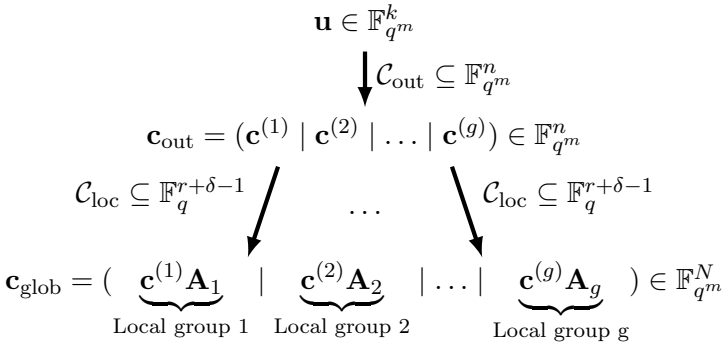

Figure 4.1: Illustration of the two-step encoding procedure generating the global codeword $\mathbf{c}_{\text{glob}}$ (Case 1 of Construction 4.1). The information vector is $\mathbf{u} \in \mathbb{F}_{q^m}^k$. The first step is the encoding with the outer code $\mathcal{C}_{\text{out}}$ resulting in $\mathbf{c}_{\text{out}} \in \mathbb{F}_{q^m}^n$, which is then partitioned into g parts $\mathbf{c}^{(i)} \in \mathbb{F}_{q^m}^r$. In the final step, $\mathbf{c}^{(i)}$ is encoded using the MDS code generator matrix $\mathbf{A} \in \mathbb{F}_q^{r \times (r+\delta-1)}$. The resulting codeword part $\mathbf{c}^{(i)}\mathbf{A} \in \mathbb{F}_{q^m}^{r+\delta-1}$ is then stored in the nodes of the i-th local group.

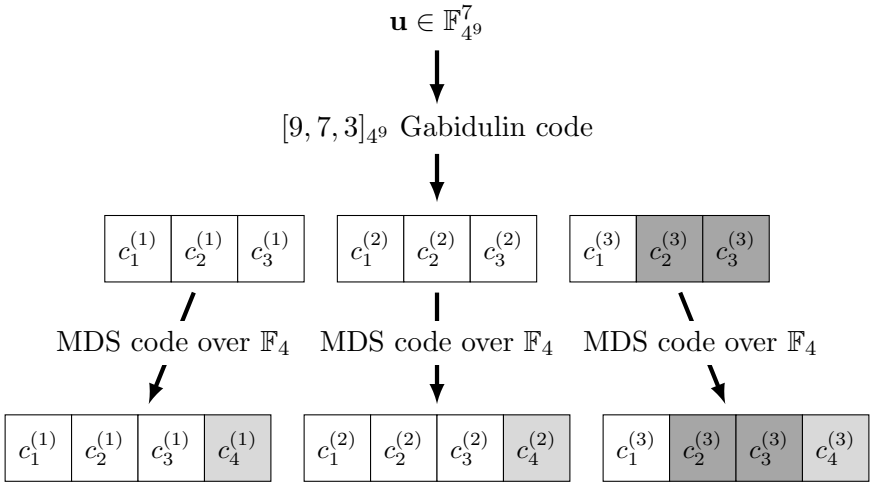

Figure 4.2: Illustration of an MR-LRC with three local groups ($g = 3$). Each local group has $(3, 2)$-localities. This code can correct up to $\delta - 1 = 1$ erasures in each group and additionally $h = 2$ erasures globally. (*Remark*: To simplify the illustration it is assumed that the MDS codes are systematic).

A more general construction using linearized Reed-Solomon codes was suggested by Martínez-Peñas and Kschischang in [MPK19]. The construction allows an arbitrary distribution of global parity symbols over the groups. Additionally, arbitrary group sizes can be chosen. The construction is an advancement of Construction 4.1 that was introduced in [RKSV14].

Construction 4.2 ([MPK19, Constr. 1])**.** Let g be the number of local groups with their corresponding (r_i, δ_i) -localities for $i \in [g]$. Let $\mathbb{F}_{q^m}$ be the extension field with base size q and extension degree m. The construction of the code has 2 steps:

1. Outer code: Choose any (n, k) code $\mathcal{C}_{\text{out}} \subseteq \mathbb{F}_{q^m}^n$ that has maximum sum-rank distance for the length n sum-rank partition $r_1, r_2, \ldots, r_g$ of order g, e.g., a linearized Reed–Solomon code which restricts $\mathbb{F}_{q^m}$ to fulfill the constraints $q > g$ and $m \geq \max_{i\in[g]} r_i$.

2. Local codes: Choose any $(r_i + \delta_i - 1, r_i)$ MDS code $\mathcal{C}_{\text{loc},i} \subseteq \mathbb{F}_q^{r_i+\delta_i-1}$ which is linear over the local field $\mathbb{F}_q$ for $i \in [g]$. The MDS codes require $q \geq \max\{r_i + \delta_i - 1 \mid i \in [g], \delta_i > 2\}$.

The global code $\mathcal{C}_{\text{glob}} \subseteq \mathbb{F}_{q^m}^N$ with $N = n + \sum_{i=1}^{g}(\delta_i - 1) = \sum_{i=1}^{g}(r_i + \delta_i - 1)$ is then defined by

$$\mathcal{C}_{\text{glob}} = \mathcal{C}_{\text{out}} \operatorname{diag}(\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_g),$$

with $\mathbf{A}_i \in \mathbb{F}_q^{r_i \times (r_i+\delta_i-1)}$ being the generator matrix of $\mathcal{C}_{\text{loc},i}$ for $i \in [g]$.

The encoding procedure of the global code is illustrated in Figure 4.3. The field size constraints for LRSCs from the first step results from the constraint to have a P-independent set of evaluation points. The number of conjugacy classes for a given $\mathbb{F}_{q^m}$ is $q - 1$. For each group a separate conjugacy class element is needed which gives the first constraint. The second constraint comes from the maximum number of $\mathbb{F}_q$-linearly independent elements that are used to construct a P-independent set as described in Theorem A.10.

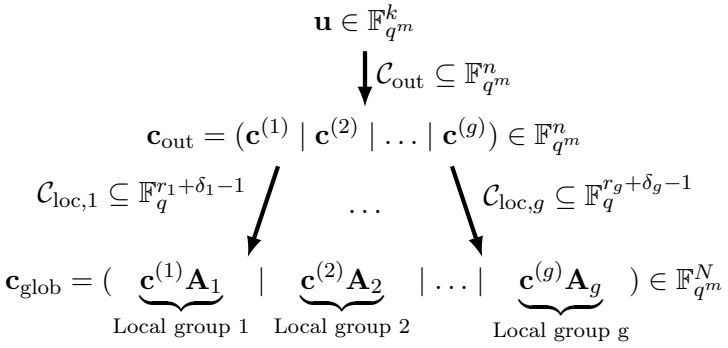

Figure 4.3: Illustration of the two-step encoding procedure generating the global codeword $\mathbf{c}_{\text{glob}}$ (Construction 4.2). The vector of symbols that should be stored with redundancy is$\mathbf{u} \in \mathbb{F}_{q^m}^k$. The first step is the encoding with the outer code $\mathcal{C}_{\text{out}}$ resulting in $\mathbf{c}_{\text{out}} \in \mathbb{F}_{q^m}^n$, which is then partitioned into g parts $\mathbf{c}^{(i)} \in \mathbb{F}_{q^m}^{r_i}$. In the final step, $\mathbf{c}^{(i)}$ is encoded using the MDS code generator matrix $\mathbf{A}_i \in \mathbb{F}_q^{r_i \times (r_i+\delta_i-1)}$ of the i-th local group. $\mathbf{c}^{(i)}\mathbf{A}_i \in \mathbb{F}_{q^m}^{r_i+\delta_i-1}$ is then stored in the nodes of the i-th local group. The figure illustrates the same as [MPK19, Fig. 4].

Theorem 4.2 ([MPK19, Th. 2])**.** *Let $\mathcal{C}_{\text{glob}} \subseteq \mathbb{F}_{q^m}^N$ be the global code from Construction 4.2 with local groups $\Gamma_i \subseteq [N]$ for $i \in [g]$. Then the code $\mathcal{C}_{\text{glob}} \subseteq \mathbb{F}_{q^m}^N$ is an MR-LRC with (r_i, δ_i)-localities for $i \in [g]$ as in Definition 3.3 and 3.4.*

Proof. Since the local codes $\mathcal{C}_{\text{loc}}^{(i)}$ have Hamming distance δ_i and the restricted global code $\mathcal{C}_{\text{glob}}|_{\Gamma_i}$ is contained in the row space of $\mathbf{A}_i$, the bound $\mathrm{d_H}\left(\mathcal{C}_{\text{glob}}|_{\Gamma_i}\right) \geq \delta_i$ holds. Therefore, the locally repairable property is fulfilled and it remains to show that the code $\mathcal{C}$ obtained after puncturing the global code $\mathcal{C}_{\text{glob}} \subseteq \mathbb{F}_{q^m}^N$ at any $\delta_i - 1$ positions in the i-th local group for each $i \in [g]$ is an MDS code. Let Δ_i be an arbitrary subset of Γ_i for $i \in [g]$ satisfying $|\Delta_i| = r_i$. Every $r_i \times r_i$ submatrix of the local generator matrices $\mathbf{A}_i \in \mathbb{F}_q^{r_i \times (r_i+\delta_i-1)}$ is invertible since $\mathbf{A}_i$ is the generator matrix of an MDS code. Thus, $\mathbf{A}_i|_{\Delta_i}$ is invertible as well, where $\mathbf{A}_i|_{\Delta_i}$ denotes the submatrix of $\mathbf{A}_i$ after restricting it to the columns indexed by Δ_i. Since the outer code is an MSRD code, the code

$$\mathcal{C} = \mathcal{C}_{\text{out}} \operatorname{diag}\left(\mathbf{A}_1|_{\Delta_1}, \ldots, \mathbf{A}_g|_{\Delta_g}\right) \subseteq \mathbb{F}_{q^m}^N$$

is an MDS code by Definition 3.8. □

This also proves that Construction 4.1 is an MR-LRC since Construction 4.2 recovers it for $\delta_1 = \cdots = \delta_g = \delta$, $r_1 = \cdots = r_{g-1} = r$ and $r_g = g$ in case 1 or $r_g = t$ in case 2. The linearized Reed-Solomon recovers the Gabidulin as described in Theorem 3.7.

Example 4.2 (MR-LRC with linearized Reed-Solomon code)**.** Consider a DSS with $g = 3$ groups where an information vector of length $k = 7$ is stored in a maximally recoverable manner. The system has $h = 2$ global parities and each local group is able to tolerate one erasure locally, i.e., $\delta = 2$, $r = 3$. The required field size is therefore $q \geq \max\{r + \delta - 1, g + 1\} = 4$ and $m \geq \max\{r_1, r_2, r_3\} = 3$. The encoding is illustrated in Figure 4.4.

The example shows that Construction 4.2 has a significantly smaller required field size compared to the construction that uses Gabidulin codes.

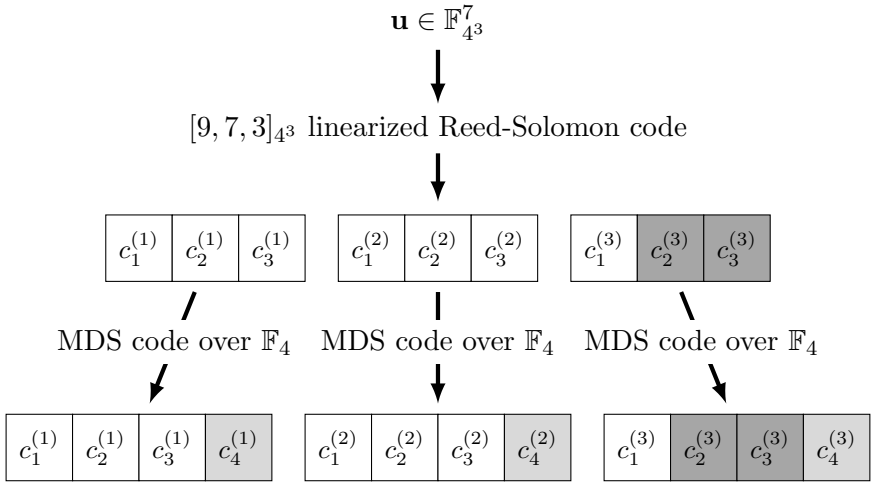

Figure 4.4: Illustration of an MR-LRC with three local groups ($g = 3$). Each row forms a local group Γ_i with localities and local distances $(r_i, \delta_i) = (3, 2)$ for $i \in [3]$. This code can correct up to $\delta - 1 = 1$ erasures in each group and additionally $h = 2$ erasures globally. (*Remark*: To simplify the illustration it is assumed that the MDS codes are systematic).

4.3 Secrecy Bound on Locally Repairable Codes

In [RKSV14], the secrecy of LRCs was analyzed when colluding eavesdroppers have access to a part of the system. An upper bound on the amount of symbols that can be stored securely on an LRC system under the influence of an (l_1, l_2)-eavesdropper was given. First, a general upper bound is derived. In the second step, the bound is specialized following the steps in [RKSV14].

Let $\mathcal{C} \subseteq \mathbb{F}^N$ be an LRC that has (r, δ)-locality and is $\mathrm{d}_{\min}$-optimal, i.e., it fulfills (4.1) and has equal group parameters $\delta_i = \delta$, $r_i = r$ for all $i \in [g]$. Associate the set of indices $\mathcal{K}$ to a data collector that can contact $N - \mathrm{d}_{\min} + 1$ nodes to reconstruct the stored data. Let $\mathcal{K}_i$ denote the indices of nodes that are contacted by the data collector in the i-th local group such that $\mathcal{K} = \bigcup_{i=1}^{g} \mathcal{K}_i$ with $|\mathcal{K}| = N - \mathrm{d}_{\min} + 1$. The indices of the eavesdropper are denoted by $\mathcal{E}_1$ and $\mathcal{E}_2$ for nodes eavesdropped in an l_1 and l_2-manner, respectively as introduced in Section 2.2. In the i-th local group the eavesdropper indices are denoted by $\mathcal{E}_1^i$ and $\mathcal{E}_2^i$ with $\mathcal{E}_1 = \bigcup_{i=1}^{g} \mathcal{E}_1^i$, $\mathcal{E}_2 = \bigcup_{i=1}^{g} \mathcal{E}_2^i$, $l_1^i = |\mathcal{E}_1^i|$, $l_2^i = |\mathcal{E}_2^i|/r$, $l_1 = \sum_{i=1}^{g} l_1^i$ and $l_2 = \sum_{i=1}^{g} l_2^i$. The set of tuples $\{(\mathcal{E}_1^i, \mathcal{E}_2^i, \mathcal{K}_i)\}_{i=1}^{g}$, that satisfy the system requirements, is denoted by $\mathcal{X}$. Lemma 32 in [RKSV14] that gives an upper bound on the secrecy capacity reads as follows.

Lemma 4.1 (**Secrecy capacity of LRCs**[RKSV14, Lem. 32])**.** *For an (r, δ)-LRC that is secure against an (l_1, l_2)-eavesdropper, the following holds*

$$k^{(\mathrm{s})} \leq \sum_{i=1}^{g} \mathrm{H}(\mathsf{K}_i \mid \mathsf{E}_1^i, \mathsf{E}_2^i),$$

where K_i denotes the random variable corresponding to the nodes contacted by a data collector $\mathcal{K}_i$. The random variables corresponding to the sets of nodes in $\mathcal{E}_1$ and $\mathcal{E}_2$ are denoted by E_1 and E_2, respectively. $k^{(\mathrm{s})}$ is the number of information symbols that can be stored securely given the eavesdropper parameters.

Proof. Assume k symbols can be encoded by the LRC without an eavesdropper. Since eavesdropping on r nodes in a group gives the eavesdropper all the information of the group, without loss of generality the focus is on indices $\{\mathcal{E}_1^i\}_{i=1}^{g}$ and $\{\mathcal{E}_2^i\}_{i=1}^{g}$ such that $|\mathcal{E}_1^i \cup \mathcal{E}_2^i| \leq r$. It is assumed that $\mathcal{K}$ is chosen such that either $\mathcal{E}_1^i \cup \mathcal{E}_2^i \subseteq \mathcal{K}_i$ or $\mathcal{K}_i = \emptyset$. To have a non-empty secure file size, $|\mathcal{E}_1| + |\mathcal{E}_2| = l_1 + l_2 r < k$. The data of size $k^{(\mathrm{s})}$ is

denoted by $\mathbf{u}^{(\mathrm{s})}$ with the corresponding random variable $\mathsf{U}^{(\mathrm{s})}$.

$$\begin{aligned}
\mathrm{H}(\mathsf{U}^{(\mathrm{s})}) &\overset{(a)}{=} \mathrm{H}(\mathsf{U}^{(\mathrm{s})} \mid \mathsf{E}_1, \mathsf{E}_2) \\
&\overset{(b)}{=} \mathrm{H}(\mathsf{U}^{(\mathrm{s})} \mid \mathsf{E}_1, \mathsf{E}_2) - \mathrm{H}(\mathsf{U}^{(\mathrm{s})} \mid \mathsf{E}_1, \mathsf{E}_2, \mathsf{K}) \\
&= \mathrm{I}(\mathsf{U}^{(\mathrm{s})}; \mathsf{K} \mid \mathsf{E}_1, \mathsf{E}_2) \\
&\leq \mathrm{H}(\mathsf{K} \mid \mathsf{E}_1, \mathsf{E}_2) \\
&= \mathrm{H}(\mathsf{K}_1, \ldots, \mathsf{K}_g \mid \mathsf{E}_1^1 \ldots, \mathsf{E}_1^g, \mathsf{E}_2^1, \ldots, \mathsf{E}_2^g) \\
&\overset{(c)}{\leq} \sum_{i=1}^{g} \mathrm{H}(\mathsf{K}_i \mid \mathsf{E}_1^i, \mathsf{E}_2^i)
\end{aligned} \tag{4.3}$$

where (a) follows from the secrecy constraint, which can be written as $\mathrm{I}(\mathsf{U}^{(\mathrm{s})}; \mathsf{E}_1, \mathsf{E}_2) = \mathrm{H}(\mathsf{U}^{(\mathrm{s})}) - \mathrm{H}(\mathsf{U}^{(\mathrm{s})} \mid \mathsf{E}_1, \mathsf{E}_2) = 0$ and (b) follows from the ability of the data collector to recover the file $\mathbf{u}^{(\mathrm{s})}$. (c) is a result of the chain rule (B.1) where equality holds if, and only if, K_i and K_j are i.i.d. for $i \neq j$.

For each choice from $\mathcal{X}$ such an upper bound holds and therefore

$$k^{(\mathrm{s})} = \mathrm{H}(\mathsf{U}^{(\mathrm{s})}) \leq \sum_{i=1}^{g} \mathrm{H}(\mathsf{K}_i \mid \mathsf{E}_1^i, \mathsf{E}_2^i).$$

□

This general bound is not very useful to compute the number of securely storable symbols. Therefore, this question is explored from a system perspective in the following. Consider a $\mathrm{d}_{\min}$-optimal LRC with g groups and equal group sizes. Let $\mathbf{l}_1 = (l_1^1, l_1^2, \ldots, l_1^g)$ and $\mathbf{l}_2 = (l_2^1, l_2^2, \ldots, l_2^g)$ be the vectors representing the pattern of the eavesdropper.

First, consider the minimum number of groups from which a data collector needs to collect all independent symbols to recover the stored file. The minimum number of groups is denoted by μ and is given as

$$\mu = \left\lfloor \frac{N - \mathrm{d}_{\min} + 1}{r + \delta - 1} \right\rfloor \overset{(a)}{=} \left\lfloor \frac{k + \left(\left\lceil \frac{k}{r} \right\rceil - 1\right)(\delta - 1)}{r + \delta - 1} \right\rfloor,$$

where (4.1) is plugged in at (a). The number of nodes that need to be downloaded in addition to the nodes of μ groups to recover the stored file is

$$\nu = N - \mathrm{d}_{\min} + 1 - (r + \delta - 1)\mu \overset{(a)}{=} k + \left(\left\lceil \frac{k}{r} \right\rceil - 1\right)(\delta - 1) - (r + \delta - 1)\mu,$$

where again (4.1) is plugged in at (a). If a data collector has the symbols of all independent

nodes of μ groups and additionally the symbols of ν nodes, it can recover the stored file.

The two parameters μ and ν can be best illustrated with an example.

Example 4.3. Consider the same LRC as in Example 4.2 with $k = 7$, $n = 9$, $N = 12$, $r = 3$, $\delta = 2$ which is illustrated in Figure 4.5. If the file should be retrieved, the symbols of k nodes that store k linearly independent symbols (over $\mathbb{F}_{q^m}$) has to be downloaded. If three symbols from one group are downloaded, the knowledge of the whole group is revealed. Calculate the two parameters μ and ν, yielding $\mu = 2$ and $\nu = 1$. This means that downloading r symbols from 2 groups and additionally downloading the symbol from one node in the third group is required to recover the stored file.

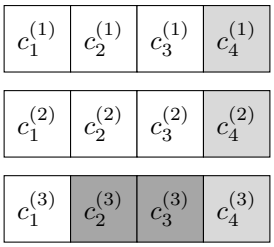

Figure 4.5: Illustration of an MR-LRC with three local groups ($g = 3$) and $(r_i, \delta_i) = (3, 2)$-localities for $i \in [3]$. Each row forms a local group Γ_i. To recover the file, the symbols of $\mu = 2$ groups and $\nu = 1$ additional node in the remaining group need to be downloaded.

Note that the notion of data collection is the same for an authorized recovery and an eavesdropper. This automatically gives an upper bound on the eavesdropper pattern. An eavesdropper pattern can only have access to fewer than $\mu r + \nu$ nodes. Otherwise, it would have global knowledge. The parameters μ and ν therefore give a naive bound on the file size that can be stored on the system without an eavesdropper. Namely, $k = \mu r + \nu$. With increasing l_1 and l_2, this size reduces and a precoding step such as a similar step to secret sharing is required since otherwise the eavesdropper would directly gain partial information of the stored file. A bound for the file size under the influence of an (l_1, l_2)-eavesdropper is deduced in the following. It follows the steps from [RKSV14, Sec. VI].

First, note that $\mathrm{H}(\mathsf{K}_i) = r$ and therefore a data collector would contact at most r nodes in a local group. Also note that if $l_2^i > 0$ for some $i \in [g]$, then the information of the whole group is revealed to the eavesdropper, i.e., $\mathrm{H}(\mathsf{K}_i|\mathsf{E}_2^i) = 0$ for a fixed $i \in [g]$ and $l_2^i > 0$. With these remarks in mind the following theorem (Theorem 33 [RKSV14]) can be proven.

Theorem 4.3 (**Secrecy capacity**). *The secrecy capacity of an (r, δ)-LRC against an (l_1, l_2)-eavesdropper is*

$$k^{(\mathrm{s})} = [\mu r + \nu - (l_2 r + l_1)]^+ \tag{4.4}$$

where $[\xi]^+$ denotes $\max\{\xi, 0\}$.

Proof. The first step is to show that the right hand side of (4.4) is an upper bound of the secrecy capacity with the help of Lemma 4.1. Consider a data collector with $\mathcal{K}_1 = \Gamma_1$, $\mathcal{K}_1 = \Gamma_1$, ..., $\mathcal{K}_\mu = \Gamma_\mu$, $\mathcal{K}_{\mu+1} = \ldots = \mathcal{K}_g = \emptyset$ and $\mathcal{K}_{\mu+1} \subset \Gamma_{\mu+1}$ such that $|\mathcal{K}_{\mu+1}| = \nu$. The eavesdropper pattern is a worst case estimation, i.e., $\mathbf{l}_2 = (1, 1, \ldots, 1, 0, \ldots, 0)$ with ones at the first l_2 positions and $\mathbf{l}_1 = (0, \ldots, 0, l_1^{l_2+1}, \ldots, l_2^g)$ with zeros at first l_2 positions.

Case 1: $l_2 r + l_1 \geq \mu r + \nu$ For this case the eavesdropper is able to reconstruct the file and therefore $\mathrm{H}(\mathsf{U}^{(\mathrm{s})} \mid \mathsf{E}_1, \mathsf{E}_2) = 0$ and (4.4) holds.

Case 2: $l_2 r + l_1 < \mu r + \nu$ Without loss of generality, the given eavesdropper pattern is concentrated on the first $\mu + 1$ groups with $l_2 < \mu + 1$, $\sum_{i=l_2}^{\mu+1} l_1^i < (\mu - l_2) r + \nu$, $l_1^i \leq r$ for all $i \in \{l_2 + 1, l_2 + 2, \ldots, g\}$, $l_1^{l_2+1} \geq l_1^{l_2+2} \geq \ldots \geq l_1^{\mu+1}$ and therefore $l_1^i = 0$ for all $i \in \{\mu + 2, \mu + 3, \ldots, g\}$. For the $(\mu + 1)$-th group, the eavesdropper pattern is, if applicable, a subset of the data collector indices, i.e., $\mathcal{E}_1^{\mu+1} \subset \mathcal{K}_{\mu+1}$. With these restrictions and (4.4), it follows that

$$\begin{aligned}
k^{(\mathrm{s})} &\leq \sum_{i=1}^{g} \mathrm{H}(\mathsf{K}_i \mid \mathsf{E}_1^i, \mathsf{E}_2^i) \overset{(a)}{=} \sum_{i=1}^{l_2} \mathrm{H}(\mathsf{K}_i \mid \mathsf{E}_2^i) + \sum_{i=l_2+1}^{g} \mathrm{H}(\mathsf{K}_i \mid \mathsf{E}_1^i) \\
&\overset{(b)}{=} \sum_{i=1}^{l_2} \mathrm{H}(\mathsf{K}_i, \mathsf{E}_2^i) - \left(\sum_{i=1}^{l_2} \mathrm{H}(\mathsf{E}_2^i) \right) + \sum_{i=l_2+1}^{\mu} \mathrm{H}(\mathsf{K}_i, \mathsf{E}_1^i) - \left(\sum_{i=l_2+1}^{\mu} \mathrm{H}(\mathsf{E}_1^i) \right) \\
&\quad + \mathrm{H}(\mathsf{K}_{\mu+1}, \mathsf{E}_1^{\mu+1}) - \mathrm{H}(\mathsf{E}_1^{\mu+1}) \\
&\overset{(c)}{=} \sum_{i=1}^{l_2} \max\left\{ \mathrm{H}(\mathsf{K}_i), \mathrm{H}(\mathsf{E}_2^i) \right\} - \left(\sum_{i=1}^{l_2} \mathrm{H}(\mathsf{E}_2^i) \right) + \sum_{i=l_2+1}^{\mu} \max\left\{ \mathrm{H}(\mathsf{K}_i), \mathrm{H}(\mathsf{E}_1^i) \right\} \\
&\quad - \left(\sum_{i=l_2+1}^{\mu} \mathrm{H}(\mathsf{E}_1^i) \right) + \max\left\{ \mathrm{H}(\mathsf{K}_{\mu+1}), \mathrm{H}(\mathsf{E}_1^{\mu+1}) \right\} - \mathrm{H}(\mathsf{E}_1^{\mu+1}) \\
&\overset{(d)}{=} l_2 r - l_2 r + (\mu - l_2) r - \sum_{i=l_2+1}^{\mu} l_1^i + \nu - l_1^{\mu+1} = \mu r + \nu - (l_2 r + l_1)
\end{aligned}$$

where (a) and (d) follow from the eavesdropper pattern, (b) follows from $\mathrm{H}(\mathsf{A} \mid \mathsf{B}) = \mathrm{H}(\mathsf{A}, \mathsf{B}) - \mathrm{H}(\mathsf{B})$, (c) holds since $\mathrm{H}(\mathsf{A}, \mathsf{B}) = \max\{\mathrm{H}(\mathsf{A}), \mathrm{H}(\mathsf{B})\}$ if, and only if, $\mathsf{B} = f(\mathsf{A})$ or $\mathsf{A} = f(\mathsf{B})$ which is the case here.

Remark: Note that for a data collector that only observes cleverly chosen $\mu r + \nu$ nodes

and the assumption that the eavesdropper choses the same nodes so that $\mathcal{K}$ and $\mathcal{E}$ overlap as much as possible, the result will be the same. This assumption is also implicitly made later in Theorem 6.1.

As a next step, a construction of a secure coding scheme that shows tightness of the upper bound on the secrecy capacity was presented in [RKSV14].

Construction 4.3. Consider an (l_1, l_2)-eavesdropper and integers μ, r, ν such that $\mu r + \nu - (l_2 r + l_1) > 0$.

1. Given the secure file size $k^{(\mathrm{s})} = \mu r + \nu - (l_2 r + l_1)$, generate $l_2 r + l_1$ independent random symbols uniformly distributed over $\mathbb{F}_{q^m}$, $\mathbf{r} = (r_1, r_2, \ldots, r_{(l_2 r + l_1)})$, and append with $\mathbf{u}^{(\mathrm{s})} = (u_1, u_2, \ldots, u_{k^{(\mathrm{s})}})$ to obtain $\mathbf{u} = (\mathbf{r}, \mathbf{u}^{(\mathrm{s})})$.

2. Encode the $k = \mu r + \nu$ symbols of $\mathbf{u}$ with an $[k, k, 1]_{q^m}$ Gabidulin code.

3. Encode the k symbols of the Gabidulin codeword $\mathbf{c}_{\mathrm{gab}}$ with a $\mathrm{d}_{\min}$-optimal LRC that has (r, δ)-locality, e.g., the second step of an MR-LRC construction as given in Construction 4.1 or 4.2.

Lemma 4.2. *Construction 4.3 is information-theoretically secure against the (l_1, l_2)-eavesdropper and achieves the secrecy capacity stated in Theorem 4.3, if no global repairs are taken into account.*

Proof. To prove secrecy of the coding scheme, the secrecy lemma (Lemma 3.1) is used. An (l_1, l_2)-eavesdropper can observe at most $l_2 r + l_1$ symbols. There are $l_2 r + l_1$ randomly generated symbols that are uniformly distributed. Thus, the first condition is fulfilled, i.e., $\mathrm{H}(\mathsf{E}) \leq \mathrm{H}(\mathsf{R})$ where E and R are the random variables of the eavesdropper's observation and the randomly generated symbols, respectively. It remains to prove that $\mathrm{H}(\mathsf{R} \mid \mathsf{U}^{(\mathrm{s})}, \mathsf{E}) = 0$. The eavesdropped information E is assumed to be chosen in the best possible way, i.e., the eavesdropper will not observe a node in an already known l_2 eavesdropped group in an l_1-manner or a redundant node in a group. Thus, E consists of $l_2 r + l_1$ evaluations at linearly independent (over $\mathbb{F}_q$) points of a linearized polynomial. The evaluations are at $\mathbb{F}_q$-linearly independent points since the local codes are MDS codes over $\mathbb{F}_q$ and the linearized polynomial F^{D_1} is $\mathbb{F}_q$-linear. In other words, the evaluations of a linear operator polynomial F^{D_1}, which is a linearized polynomial, are given at $l_2 r + l_1$ P-independent points since in this case P-independence for $g = 1$ coincides with $\mathbb{F}_q$-linear independence. In addition the symbols of $\mathbf{u}^{(\mathrm{s})}$ are known. This recovers the problem of an $(n, \mu r + \nu, l_2 r + l_1)$ secret sharing scheme for linear operator polynomials that was proven to be information-theoretically secret in Section 3.6. □

The secrecy lemma ensures that $\mathrm{I}(\mathsf{U}^{(\mathrm{s})};\mathsf{E}) = 0$. This implies that the coding scheme ensures that there is no plaintext symbol in the codeword. In the following, an example is given to illustrate this.

Example 4.4. Let $k^{(\mathrm{s})}$ denote the number of message symbols with $\mathbf{u}^{(\mathrm{s})} = (m_1, \ldots, m_k^{(\mathrm{s})}) \in \mathbb{F}^{k^{(\mathrm{s})}}$ and z the number of random symbols that are uniformly distributed over $\mathbb{F}$ with $\mathbf{r} = (r_1, \ldots, r_z) \in \mathbb{F}$ such that $k^{(\mathrm{s})} + z = k$. Let $\mathcal{C} \subseteq \mathbb{F}^n$ be a code that is defined by the coding scheme Enc: $\mathbb{F}^k \longrightarrow \mathbb{F}^n$. Let E be the random variable corresponding to $\mathcal{E}$, i.e., the symbols observed of the codeword $\mathbf{c} \in \mathcal{C}$ that are observed by the eavesdropper with $|\mathcal{E}| = z$ such that $\mathrm{H}(\mathsf{E}) \leq \mathrm{H}(\mathsf{R})$. Now assume that the coding scheme yields a plaintext message symbol m_p for some $p \in [k^{(\mathrm{s})}]$ in the codeword $\mathbf{c} \in \mathcal{C}$. This obviously means that $\mathrm{I}(\mathsf{U}^{(\mathrm{s})};\mathsf{E}) \neq 0$. It can now be shown that $\mathrm{H}(\mathsf{R} \mid \mathsf{U}^{(\mathrm{s})}, \mathsf{E}) \neq 0$:

$$\begin{aligned}
\mathrm{H}(\mathsf{R} \mid \mathsf{U}^{(\mathrm{s})}, \mathsf{E}) &\overset{(a)}{=} \mathrm{H}(\mathsf{R} \mid \mathsf{U}^{(\mathrm{s})}) - \mathrm{H}(\mathsf{E} \mid \mathsf{U}^{(\mathrm{s})}) + \mathrm{H}(\mathsf{E} \mid \mathsf{R}, \mathsf{U}^{(\mathrm{s})}) \\
&\overset{(b)}{=} \mathrm{H}(\mathsf{R}) - \mathrm{H}(\mathsf{E} \mid \mathsf{U}^{(\mathrm{s})}) + \mathrm{H}(f(\mathsf{R}, \mathsf{U}^{(\mathrm{s})}) \mid \mathsf{R}, \mathsf{U}^{(\mathrm{s})}) \\
&\overset{(c)}{=} \mathrm{H}(\mathsf{R}_1)z - \mathrm{H}(\mathsf{E} \mid \mathsf{U}^{(\mathrm{s})}) \\
&\overset{(d)}{\geq} \mathrm{H}(\mathsf{R}_1)z - \mathrm{H}(\mathsf{R}_1)(z-1) \\
&= \mathrm{H}(\mathsf{R}_1) \\
&> 0
\end{aligned}$$

where (a) follows from $\mathrm{I}(\mathsf{R};\mathsf{E} \mid \mathsf{U}^{(\mathrm{s})})$ being written in two ways following $\mathrm{I}(\mathsf{X};\mathsf{Y}) = \mathrm{H}(\mathsf{X}) - \mathrm{H}(\mathsf{X} \mid \mathsf{Y}) = \mathrm{H}(\mathsf{Y}) - \mathrm{H}(\mathsf{Y} \mid \mathsf{X})$ (see (B.2)), (b) is a result from R and $\mathsf{U}^{(\mathrm{s})}$ being independent and that E is a subset of the codeword which is a function of the random symbols and the message symbols. (c) follows since R and $\mathsf{U}^{(\mathrm{s})}$ essentially determine $f(\mathsf{R}, \mathsf{U}^{(\mathrm{s})})$ and it follows from R being a sequence of z uniformly distributed random variables. (d) follows from the fact that one eavesdropped symbols is a message symbol and that $\mathrm{H}(\mathsf{E} \mid \mathsf{U}^{(\mathrm{s})}) \leq \mathrm{H}(\mathsf{E}) \leq \mathrm{H}(\mathsf{R})$.

In other words, $\mathrm{H}(\mathsf{R} \mid \mathsf{U}^{(\mathrm{s})}, \mathsf{E}) = 0$ and $\mathrm{H}(\mathsf{E}) \leq \mathrm{H}(\mathsf{R})$ ensure that an encoding "scrambles" the random symbols with the information symbols such that secrecy is ensured.

5 Global Repair of MR-LRCs

In this chapter, the distributed global repair is motivated and we introduce a concept that allows us to realize a global repair in a distributed way. A global repair can be realized by evaluating the outer code polynomial. The idea is now to split this polynomial into a sum of, what we call, local polynomials. Each local polynomial corresponds to a group and allows to calculate the contribution of this group to the global repair process.

Two schemes realizing such a distributed global repair are presented in Section 5.3 and Section 5.4. In direct global repair, the contribution to the global repair of each group is directly send to the group where the repair is performed. In forwarded global repair, the contribution is forwarded such that each group only receives one symbol. The group where the global repair is performed is then at the end of the forwarding list and only receives one symbol which is a sum of all contributions.

5.1 Global Repair Introduction and Definitions

The big advantage of MR-LRCs is that they can correct the information theoretical maximum of erasures given the code parameters as seen in Section 3.2. If a local group has more failed nodes than it can handle locally, a global repair with the help of other groups is possible. For Construction 4.2, the global repair involves an erasure correction by the outer code, i.e., the linearized Reed-Solomon code. For a global repair of the j-th node in the i-th group, the linear operator polynomial $P^{D_{a_i}}$ is evaluated at the position $\beta_j^{(i)}$. One way to get the evaluation is to generate the corresponding skew polynomial P of degree $k-1$ by Newton interpolation. For the Newton interpolation, k evaluations of the polynomial at P-independent points are needed to generate the unique polynomial (Theorem A.7). This means that k symbols of the n outer code symbols are needed. The global repair is illustrated with the following example.

Example 5.1. Consider a DSS with three local groups $g = 3$, $(3, 2)$-locality and code parameters $k = 7$, $n = 9$ and $N = 12$. If two nodes in the second group fail for example $c_2^{(2)}$ and $c_4^{(2)}$, a global repair, followed by a local repair, is needed to recover the nodes. The failed node could for instance download the symbols $c_1^{(1)}, c_2^{(1)}, c_3^{(1)}, c_1^{(2)}, c_3^{(2)}, c_1^{(3)}$ and $c_2^{(3)}$ which is shown in Figure 5.1. The second node of the second group can now generate

the skew polynomial P using Definition A.8. It can then recover the symbol evaluating $P^{D_{a_2}}$ at $\beta_2^{(2)}$, i.e., $P^{D_{a_2}}(\beta_2^{(2)}) = c_2^{(2)}$.

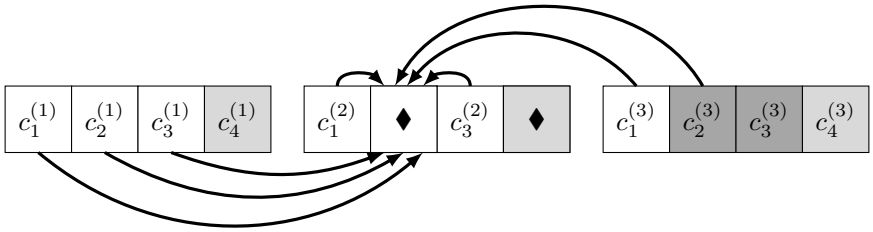

Figure 5.1: Illustration of a DSS with three groups, $(3, 2)$-locality and $k = 7$. It shows the repair download for a global repair of the second node in the second group in a non-hierarchical DSS. The failed node downloads as many symbols as needed for its repair.

The above example raises the question of which sets of nodes can be used for a global repair. For simplicity, only equal group sizes are considered in the following.

Proposition 5.1. *Let $\mathcal{C}_{\text{glob}} \subseteq \mathbb{F}_{q^m}^N$ be the global code from Construction 4.2 with local groups $\Gamma_i \subseteq [N]$ and equal localities, i.e., $r_i = r$ and $\delta_i = \delta$ for $i \in [g]$. A global repair can be performed by downloading the symbols of the nodes in any subset Δ with $\Delta_i \subseteq \Gamma_i$, $|\Delta_i| \leq r$ and $\sum_{i=1}^{g} |\Delta_i| \geq k$ of intact nodes.*

Proof. If an arbitrary subset Δ_i of each local group Γ_i with $|\Delta_i| = k_i \leq r$ is chosen, every submatrix $\mathbf{A}_i|_{\Delta_i} \in \mathbb{F}_{q^m}^{r \times k_i}$ of the local generator matrix $\mathbf{A}_i \in \mathbb{F}_{q^m}^{r \times r+\delta-1}$ has rank k_i since $\mathcal{C}_{\text{loc},i}$ is MDS. Given that linearized Reed-Solomon codes are MSRD (Definition 3.8) and Corollary 3.1, the erasure can be corrected only if

$$n - \sum_{i=1}^{g} \mathbf{A}_i|_{\Delta_i} < n - k + 1$$

is fulfilled, which means that

$$\sum_{i=1}^{g} \mathbf{A}_i|_{\Delta_i} \geq k$$

has to hold. This is the case by definition. □

Given the outer code $\mathcal{C}_{\text{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$, the P-independent set of evaluation points $\hat{b}_j^{(i)} = {}^{\hat{\beta}_j^{(i)}} a_i$, which is used for the global repair, is given by

$$\hat{\boldsymbol{\beta}} = \boldsymbol{\beta} \cdot \mathbf{A}|_{\Delta} = (\hat{\boldsymbol{\beta}}^{(1)}, \hat{\boldsymbol{\beta}}^{(2)}, \ldots, \hat{\boldsymbol{\beta}}^{(g)})$$

with $\hat{\boldsymbol{\beta}}^{(i)} = (\hat{\beta}_j^{(i)} \mid j \in \Delta_i) \in \mathbb{F}_{q^m}^{k_i}$ for $i \in [g]$ and the group elements a_i for $i \in [g]$ which are the same as for the encoding with Construction 4.2. The points ${}^{\hat{\beta}_j^{(i)}} a_i$ yield a P-independent set of evaluations by Theorem A.10 since the matrices $\mathbf{A}_i|_{\Delta_i}$ have full rank and the elements of $\boldsymbol{\beta}^{(i)}$ are $\mathbb{F}_q$-linearly independent.

To illustrate possible choices of nodes for the global repair, Example 5.1 is continued.

Example 5.2. Consider the same parameters as in Example 5.1. The failed node could for instance download the symbols $c_1^{(1)}, c_2^{(1)}, c_3^{(1)}, c_1^{(2)}, c_3^{(2)}, c_1^{(3)}$ and $c_2^{(3)}$ which is shown in Figure 5.1. It could also download the symbols $c_2^{(1)}, c_3^{(1)}, c_4^{(1)}, c_1^{(2)}, c_1^{(3)}, c_2^{(3)}$ and $c_3^{(3)}$ or $c_2^{(1)}, c_4^{(1)}, c_1^{(2)}, c_3^{(2)}, c_2^{(3)}, c_3^{(3)}$ and $c_4^{(3)}$. It is just important that no more than r nodes from each group are used since the local parities are only a linear combination of the information nodes in a group. Therefore, the parity symbol would not bring any innovation if the nodes which it is a linear combination of are already taken into account.

It is beneficial to choose k_i for $i \in [g]$ such that $\sum_{i=1}^{g} k_i = k$ so that only the minimum number of symbols needed for repair is used.

Definition 5.1 (**Minimal global repair set**)**.** Let $\mathcal{C}_{\text{glob}} \subseteq \mathbb{F}_{q^m}^N$ be the global code from Construction 4.2 with local groups $\Gamma_i \subseteq [N]$ of equal group size $r + \delta - 1 = r_i + \delta_i - 1$ for $i \in [g]$. Let $\Delta = \bigcup_{i=1}^{g} \Delta_i$ be a global repair set which can perform a global repair by Proposition 5.1. The global repair set Δ is said to be minimal if $\sum_{i=1}^{g} |\Delta_i| = k$.

Another representation of the set $\Delta \subseteq [N]$ is in the following denoted by Δ_{glob} and is a set of tuples $(i, j) \in \mathbb{N} \times \mathbb{N}$. There exists a bijective mapping $\Phi : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ which maps $m \in \Delta$ to its tuple representation $(i, j) \in \Delta_{\text{glob}}$ where i denotes the group number and j the number of the node in the i-th group.

Given a non-hierarchical DSS model, as illustrated in Figure 3.6 and Figure 5.1, a failed node would contact k nodes with linearly independent symbols to recover itself. However, such an approach would yield a secrecy capacity of zero since global knowledge is revealed to one node which could be observed by an eavesdropper in an l_2-manner. Therefore, the important question is whether a node can be repaired globally while the eavesdropper gets only partial knowledge about the symbols stored in other groups. Such a scheme only works if the system is hierarchical such that in each group computations can be performed for the global repair in another group and only a limited number of symbols is sent to the group where the global repair is performed.

Assume, without loss of generality, that a global repair is performed in the first group. The first group uses its nodes that are still intact for the global repair process. In addition, the other groups send their contribution to the global repair to the first group. In the first

scheme, each group sends its information directly to the group where a global repair is performed. In the second scheme, each group forwards its contribution to the next group where the contributions are combined. The two schemes are illustrated in Figure 5.2. The two schemes are explained in detail in Sections 5.3 and 5.4.

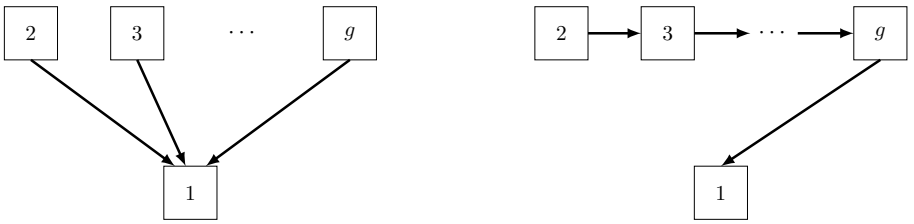

(a) Direct global repair (b) Forwarded global repair

Figure 5.2: Illustration of two different global repair schemes where an erasure in the first group is repaired.

5.2 Local Polynomials

Before the global repair schemes are discussed in detail, we introduce a principle which allows the two types of global repair to be performed. The goal is to implement a global repair process such that only one symbol is sent from each group as a contribution to the global repair of one symbol. It is assumed that the DSS uses a linearized Reed-Solomon code as a global code as seen in Construction 4.2.

Definition 5.2 (Local polynomial). Let $\mathcal{C}_{\text{glob}} \subseteq \mathbb{F}_{q^m}^N$ be the global code from Construction 4.2 and $\mathcal{C}_{\text{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$ be the outer code. Fix a minimal global repair set Δ_{glob} of nodes with $|\Delta_{\text{glob}}| = k$ as proposed in Proposition 5.1. The local polynomial of the i-th group L_i has the following properties:

- $L_i^{D_{a_i}}(\beta_j^{(i)}) = c_j^{(i)}$ for all $(i,j) \in \Delta_{\text{glob}}$
- $L_i^{D_{a_s}}(\beta_m^{(s)}) = 0$ for all $s \neq i$ and $(s,m) \in \Delta_{\text{glob}}$

Thus, $|\Delta_{\text{glob}}| = k$ constraints are imposed on L_i which has degree $k-1$. By Theorem 3.6, the constraints can also be written as:

- $L_i(^{\beta_j^{(i)}} a_i) = c_j^{(i)}/\beta_j^{(i)}$ for all $(i,j) \in \Delta_{\text{glob}}$
- $L_i(^{\beta_m^{(s)}} a_s) = 0$ for all $s \neq i$ and $(s,m) \in \Delta_{\text{glob}}$

The local polynomials can be generated by Newton interpolation.

Theorem 5.1. *Let $\mathcal{C}_{\text{glob}} \subseteq \mathbb{F}_{q^m}^N$ be the global code from Construction 4.2 and $\mathcal{C}_{\text{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$ be the outer code $\mathcal{C}_{\text{out}}$. Fix a minimal global repair set Δ_{glob} of nodes with $|\Delta_{\text{glob}}| = k$ as proposed in Proposition 5.1. Let P be the encoding polynomial of the linearized Reed-Solomon code $\mathcal{C}_{\text{out}}$. It then holds that*

$$P = \sum_{i=1}^{g} L_i. \tag{5.1}$$

Proof. By Theorem A.7 the sum of local polynomials is equivalent to the encoding polynomial if their evaluations at k points are the same. It holds that

$$P^{D_{a_i}}(\beta_j^{(i)}) = c_j^{(i)} \text{ for all } (i,j) \in \Delta_{\text{glob}}.$$

For the sum of local polynomials and $(i,j) \in \Delta_{\text{glob}}$, we have

$$\sum_{m=1}^{g} L_m^{D_{a_i}}(\beta_j^{(i)}) = \sum_{\substack{1 \le m \le g \\ m \neq i}} L_m^{D_{a_i}}(\beta_j^{(i)}) + L_i^{D_{a_i}}(\beta_j^{(i)}).$$

By Definition 5.2, it holds that

$$\sum_{\substack{1 \le m \le g \\ m \neq i}} L_m^{D_{a_i}}(\beta_j^{(i)}) + L_i^{D_{a_i}}(\beta_j^{(i)}) = \sum_{\substack{1 \le m \le g \\ m \neq i}} 0 + c_j^{(i)} = c_j^{(i)}.$$

Therefore, (5.1) holds. □

The theorem is the essential result of this section. It means that a global repair which can be seen as the evaluation of the global encoding polynomial can be distributed. Each group can calculate its contribution to the global repair as the evaluation of its local polynomial.

Definition 5.3. Let $\mathcal{C}_{\text{glob}} \subseteq \mathbb{F}_{q^m}^N$ be the global code from Construction 4.2 and $\mathcal{C}_{\text{LRS}}^{\sigma,k}(\mathbf{a}, \boldsymbol{\beta})$ be the outer code $\mathcal{C}_{\text{out}}$. Let (s,m) be the node to repair. Fix a minimal global repair set Δ_{glob} of nodes with $|\Delta_{\text{glob}}| = k$ and $(s,m) \notin \Delta_{\text{glob}}$ as proposed in Proposition 5.1. The global repair, which is an evaluation of the encoding polynomial P of the linearized Reed-Solomon code $\mathcal{C}_{\text{out}}$, can be realized by

$$P^{D_{a_s}}(\beta_m^{(s)}) = \sum_{i=1}^{g} L_i^{D_{a_s}}(\beta_m^{(s)}). \tag{5.2}$$

It is important to point out that such a distributed global repair is only feasible if

all the groups have knowledge about the parameters of the code, which should be given anyway, and more important the used global repair set. The repair can be realized in two ways which is discussed in the following sections.

5.3 Direct Global Repair

For the direct global repair each group calculates the evaluations of its local polynomial. The evaluation is then sent directly to the global repair group. All the summands of (5.2) are known to the group where the global repair is performed. The scheme is illustrated with the following example.

Example 5.3. Consider a DSS with three local groups $g = 3$ and the code parameters $k = 7$, $n = 9$, $N = 12$, $r = 3$ and $\delta = 2$. If two nodes in the second group fail for example $c_2^{(2)}$ and $c_4^{(2)}$, a global repair followed by a local repair is needed to recover the nodes. The global repair is performed with the minimal global repair set $\Delta_{\text{glob}} = \{(1,1),(1,2),(1,3),(2,1),(2,3),(3,1),(3,2)\}$. The local polynomials L_1, L_2 and L_3 have the evaluation constraints as summarized in the table of Figure 5.3. Note that for each group, a different group element a_i is used, i.e., the first constraint of L_1 means that $L_1^{D_{a_1}}(\beta_1^{(1)}) = c_1^{(1)}$ whereas the fourth constraint means that $L_1^{D_{a_2}}(\beta_1^{(2)}) = 0$.

Direct global repair has one disadvantage. Given an eavesdropper observing a group with global repairs in an l_2-manner, the eavesdropper can read all the symbols that are sent to the group. This problem and its resulting secrecy capacity is discussed in Section 6.2.

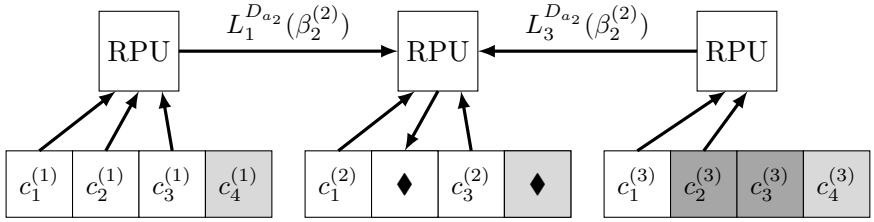

Code Locator	$\beta_1^{(1)}$	$\beta_2^{(1)}$	$\beta_3^{(1)}$			$\beta_1^{(2)}$	$\beta_2^{(2)}$	$\beta_3^{(2)}$			$\beta_1^{(3)}$	$\beta_2^{(3)}$	$\beta_3^{(3)}$	
L_1	$c_1^{(1)}$	$c_2^{(1)}$	$c_3^{(1)}$	–		0	□	0	–		0	0	–	–
L_2	0	0	0	–		$c_1^{(2)}$	□	$c_3^{(2)}$	–		0	0	–	–
L_3	0	0	0	–		0	□	0	–		$c_1^{(3)}$	$c_2^{(3)}$	–	–

Figure 5.3: Global repair process of a failed node in the second group of an MR-LRC with $(3, 2)$-localities managed by rack processing units (RPUs). Each of the RPUs computes the local polynomial of its group with the constraints that are summarized in the table. The symbol "□" means that the code locator cannot be constrained since it corresponds to the code symbol that should be recovered.

5.4 Forwarded Global Repair

The forwarded global repair, as illustrated in Figure 5.2, is not organized as a tree structure where each group sends its evaluation to the root but as a line. Every group, except for one group and the receiving global repair group, receives one symbol and adds the evaluation of its local polynomial to the received symbol, which is then forwarded to the next group. The sum in (5.2) is therefore calculated iteratively, where each group adds its summand to the already calculated partial sum. Without loss of generality, the following convention is applied throughout this work.

Definition 5.4. Let $\mathcal{G} = [g]$ be the set of groups in ascending order that are involved in a global repair process where a node in a fixed group at index $m \in \mathcal{G}$ is repaired. The forwarded repair starts at the group with index $i_1 \in \mathcal{G} \setminus m$ which denotes the first entry of $\mathcal{G} \setminus m$. The symbols are then forwarded in ascending order in $\mathcal{G} \setminus m$. The forwarding tree is summarized in a forwarding tuple $\mathbf{f}$ with

$$\mathbf{f} = (i_1, i_2, \ldots, i_{g-1}, m)$$

where $i_j \in \mathcal{G} \setminus m$ for $j \in [g-1]$. For a given group which is at position s in the forwarding list $\mathbf{f}$, the list can be split into two parts, the upstream part $\mathbf{f}_{\text{up}} = (f_1, f_2, \ldots, f_{s-1})$ corresponding to the set $\mathcal{G}_{\text{up}} \subseteq \mathcal{G}$ and the downstream part $\mathbf{f}_{\text{down}} = (f_{s+1}, f_{s+2}, \ldots, f_g)$ corresponding to the set $\mathcal{G}_{\text{down}} \subseteq \mathcal{G}$.

The scheme is illustrated with the following example.

Example 5.4. Consider a DSS with three local groups $g = 3$ and the code parameters $k = 7$, $n = 9$, $N = 12$, $r = 3$ and $\delta = 2$. If two nodes in the second group fail, for example $c_2^{(2)}$ and $c_4^{(2)}$, a global repair followed by a local repair is needed to recover the nodes. The global repair is performed with the minimal global repair set $\Delta_{\text{glob}} = \{(1,1),(1,2),(1,3),(2,1),(2,3),(3,1),(3,2)\}$. The local polynomials L_1, L_2 and L_3 have the same evaluation constraints as summarized in the table of Figure 5.3. The forwarding tuple is $(1,3,2)$. For $s = 2$ the forwarding tuple is split into $\mathbf{f}_{\text{up}} = (1)$ and $\mathbf{f}_{\text{down}} = (2)$ with the corresponding sets $\mathcal{G}_{\text{up}} = \{1\}$ and $\mathcal{G}_{\text{down}} = \{2\}$.

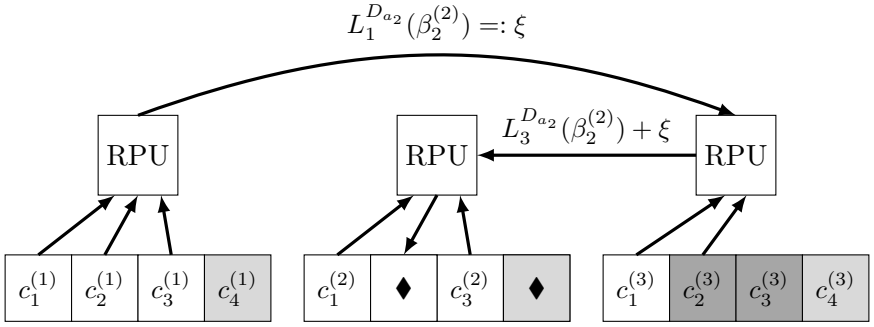

Code Locator	$\beta_1^{(1)}$	$\beta_2^{(1)}$	$\beta_3^{(1)}$			$\beta_1^{(2)}$	$\beta_2^{(2)}$	$\beta_3^{(2)}$			$\beta_1^{(3)}$	$\beta_2^{(3)}$	$\beta_3^{(3)}$	
L_1	$c_1^{(1)}$	$c_2^{(1)}$	$c_3^{(1)}$	–		0	□	0	–		0	0	–	–
L_2	0	0	0	–		$c_1^{(2)}$	□	$c_3^{(2)}$	–		0	0	–	–
L_3	0	0	0	–		0	□	0	–		$c_1^{(3)}$	$c_2^{(3)}$	–	–

Figure 5.4: Global repair process of a failed node in the second group of an MR-LRC with $(3,2)$-localities managed by rack processing units (RPUs). The repair scheme is forwarded global repair. Each of the RPUs computes the local polynomial of its group with the constraints that are summarized in the table. The symbol "□" means that the code locator cannot be constrained since it corresponds to the code symbol that should be recovered.

6 Secrecy Capacity of MR-LRCs

The main goal of this chapter is to derive a secrecy capacity for the two global repair schemes, i.e., direct global repair and forwarded global repair, that were introduced in Chapter 5.

First, the construction of an MR-LRC that is achieving the secrecy capacities for the two global repair schemes is introduced. It also achieves the secrecy capacity that was stated in Theorem 4.3 when no global repairs are considered. Second, a general upper bound on the secrecy capacity given global repair is derived.

Third, the secrecy capacities of direct and forwarded global repair are derived following the same steps. Preliminary considerations are made on fundamental limits of the repair schemes. For direct global repair such a limit is that one of the received symbols for a global repair is not independent of the others, given the assumption that the eavesdropper is able to read the failed node before the erasure occurs. For forwarded repair, the most important preliminary considerations is that if the eavesdropper observes the first or last group of a forwarding list in an l_2-manner, no new knowledge is gained in the context of global repair. After the preliminary considerations, the upper bounds on the secrecy capacity for the repair schemes are derived using the general upper bound. The main tool that is used deriving the secrecy capacity is Lemma 3.7. We write all the eavesdropper knowledge in a matrix whose rank indicates how much independent equations about the stored data the eavesdropper has. This number is then subtracted by the number of equations that the eavesdropper has by directly reading nodes. In the end of the chapter, the two repair schemes and their capacities are compared. Forwarded global repair has in general a larger secrecy capacity compared to direct global repair. However, a drawback of forwarded global repair is that the latency might be quite high since each node is waiting for the previous contribution to arrive before it can send its contribution.

6.1 MR-LRC Secrecy Construction and Secrecy Bound

In Section 4.3, a construction for an LRC that attains the secrecy capacity of Theorem 4.3 was given. This coding scheme can be slightly modified using linear operator polynomials instead of linearized polynomials. The following construction is a combination of Construction 4.2 and Construction 4.3.

Construction 6.1. Let $k^{(e)}$ be the number of independent constraints that an (l_1, l_2)-eavesdropper has on the stored file of size k. Assume that $k - k^{(e)} > 0$ and let $r|n$. The construction has the following steps:

1. Given the secure file size $k^{(\mathrm{s})} = k - k^{(e)}$, generate $k^{(e)}$ independent random symbols uniformly distributed over $\mathbb{F}_{q^m}$, $\mathbf{r} = (r_1, r_2, \ldots, r_{k^{(e)}})$, and append with $\mathbf{u}^{\mathrm{s}} = (u_1, u_2, \ldots, u_{k^{(\mathrm{s})}})$ to obtain $\mathbf{u} = (\mathbf{r}, \mathbf{u}^{\mathrm{s}})$.

2. Encode the k symbols of $\mathbf{u}$ with an $[n, k, n-k+1]_{q^m}$ linearized Reed-Solomon code.

3. Divide the n symbols of the linearized Reed-Solomon codeword $\mathbf{c}_{\mathrm{out}}$ into g groups such that $n = gr$ and encode each group with an $\mathbb{F}_q$-linear MDS code as in Construction 4.2.

Proposition 6.1. *Construction 6.1 is information-theoretically secure against an (l_1, l_2)-eavesdropper and achieves the secrecy capacity $k^{(\mathrm{s})} = k - k^{(e)}$ with $k^{(e)} = l_2 r + l_1$ from Theorem 4.3, if global repair is not considered.*

Proof. To prove secrecy of the coding scheme, the secrecy lemma (Lemma 3.1) is used. An (l_1, l_2)-eavesdropper can observe at most $k^{(e)} = l_2 r + l_1$ symbols. There are $k^{(e)} = l_2 r + l_1$ randomly generated symbols. Thus, the first condition is fulfilled, i.e., $\mathrm{H}(\mathsf{E}) \leq \mathrm{H}(\mathsf{R})$ where E and R are the random variables of the eavesdropper's observation and the randomly generated symbols, respectively. It remains to prove that $\mathrm{H}(\mathsf{R} \mid \mathsf{U}^{(\mathrm{s})}, \mathsf{E}) = 0$. The eavesdropper's $k^{(e)} = l_2 r + l_1$ observations are chosen in the best possible way, i.e., the eavesdropper will not observe a node in an l_1-manner if the whole group is already known from the l_2 observations. The crucial point is that the evaluation points of the eavesdropper's observations are P-independent. By Theorem A.10, this is equivalent to showing that all the evaluation points in a group are $\mathbb{F}_q$-linearly independent since each group uses evaluation points from a different conjugacy class. The local groups are MDS codes that are $\mathbb{F}_q$-linear which yields the desired property. In addition, the symbols of $\mathbf{u}^{(\mathrm{s})}$ are known. This recovers the problem of a $(n, k, k^{(e)})$ secret sharing scheme for linear operator polynomials that was explored in Section 3.6 and the proof is done. □

If global repair is considered, the capacity from Theorem 4.3 is decreased by a term $\lambda_{\mathrm{MR}} \geq 0$, i.e.,

$$k^{(\mathrm{s})}{}_{\mathrm{MR-LRC}} = \left[k - (l_2 r + l_1) - \lambda_{\mathrm{MR}}\right]^+, \tag{6.1}$$

where $k^{(e)} = (l_2 r + l_1) + \lambda_{\mathrm{MR}}$.

Before deriving an upper bound on λ_{MR}, we first explain the notations that are used. The set of nodes contacted by a data collector are denoted $\mathcal{K} = \bigcup_{i=1}^{g} \mathcal{K}_i = \Delta_{\mathrm{glob}}$ with $|\mathcal{K}| = k$. The observation of the eavesdropper from the groups that are observed in an l_2-manner, $\mathcal{E}_2$, is split into two different subsets. The nodes that are directly read by the eavesdropper in an l_2-manner are represented by the set $\mathcal{E}_2^{\mathrm{sto}}$. The set of evaluations of the local polynomial L_i that are sent to a group with a global repair is denoted by $\mathcal{E}_2^{\mathrm{rep}}$.

Proposition 6.2 (**Global repair - secrecy capacity upper bound**). *Consider a DSS with g groups, parameters $h < r$ and fixed eavesdropper parameters $l_2 \geq 1$, l_1. Furthermore, let the eavesdropper observe the DSS in such a way that $\mathcal{E}_1 \cap \mathcal{E}_2^{\mathrm{sto}} = \emptyset$ and that the eavesdropper is not observing more nodes than needed to recover all the stored data. The secrecy capacity of a DSS with global repair is bounded by*

$$\mathrm{H}(\mathsf{U}^{(\mathrm{s})}) \leq \mathrm{H}(\mathsf{K} \mid \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}, \mathsf{E}_2^{\mathrm{rep}}) = \mathrm{H}(\mathsf{K}) - \mathrm{H}(\mathsf{E}_1) - \mathrm{H}(\mathsf{E}_2^{\mathrm{sto}}) - \mathrm{H}(\mathsf{E}_2^{\mathrm{rep}} \mid \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}). \tag{6.2}$$

Proof. From (4.3), we have

$$\mathrm{H}(\mathsf{U}^{(\mathrm{s})}) \leq \mathrm{H}(\mathsf{K} \mid \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}, \mathsf{E}_2^{\mathrm{rep}}).$$

We can express the term $\mathrm{H}(\mathsf{K} \mid \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}, \mathsf{E}_2^{\mathrm{rep}})$ applying the chain rule of entropy (B.1) on $\mathrm{H}(\mathsf{K}, \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}, \mathsf{E}_2^{\mathrm{rep}})$:

$$\mathrm{H}(\mathsf{K}, \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}, \mathsf{E}_2^{\mathrm{rep}}) = \mathrm{H}(\mathsf{E}_1) + \mathrm{H}(\mathsf{E}_2^{\mathrm{sto}} \mid \mathsf{E}_1) + \mathrm{H}(\mathsf{E}_2^{\mathrm{rep}} \mid \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}) + \mathrm{H}(\mathsf{K} \mid \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}, \mathsf{E}_2^{\mathrm{rep}}) \tag{6.3}$$

Since $\mathsf{E}_1 = f(\mathsf{K})$,$\mathsf{E}_2^{\mathrm{sto}} = f(\mathsf{K})$ and $\mathsf{E}_2^{\mathrm{rep}} = f(\mathsf{K})$, it holds that

$$\mathrm{H}(\mathsf{K}, \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}, \mathsf{E}_2^{\mathrm{rep}}) = \mathrm{H}(\mathsf{K}). \tag{6.4}$$

Combining (6.4) and (6.3) yields

$$\mathrm{H}(\mathsf{K} \mid \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}, \mathsf{E}_2^{\mathrm{rep}}) = \mathrm{H}(\mathsf{K}) - \mathrm{H}(\mathsf{E}_1) - \mathrm{H}(\mathsf{E}_2^{\mathrm{sto}} \mid \mathsf{E}_1) - \mathrm{H}(\mathsf{E}_2^{\mathrm{rep}} \mid \mathsf{E}_1, \mathsf{E}_2^{\mathrm{sto}}).$$

Consider the term $\mathrm{H}(\mathsf{E}_2^{\mathrm{sto}} \mid \mathsf{E}_1)$. The two assumption mean that $\mathcal{E}_1 \cap \mathcal{E}_2^{\mathrm{sto}} = \emptyset$ and that

$\mathcal{E}_1$ and $\mathcal{E}_2^{\text{sto}}$ are independent sets. Thus, it holds that $\mathrm{H}(\mathsf{E}_2^{\text{sto}} \mid \mathsf{E}_1) = \mathrm{H}(\mathsf{E}_2^{\text{sto}})$ and we have

$$\mathrm{H}(\mathsf{K} \mid \mathsf{E}_1, \mathsf{E}_2^{\text{sto}}, \mathsf{E}_2^{\text{rep}}) = \mathrm{H}(\mathsf{K}) - \mathrm{H}(\mathsf{E}_1) - \mathrm{H}(\mathsf{E}_2^{\text{sto}}) - \mathrm{H}(\mathsf{E}_2^{\text{rep}} \mid \mathsf{E}_1, \mathsf{E}_2^{\text{sto}}).$$

□

With Equation (6.2), the term λ_{MR} from (6.1) can be derived.

Proposition 6.3 (**Global repair - Secrecy capacity decrease upper bound**). *Consider a DSS with g groups, parameters $h < r$ and fixed eavesdropper parameters $l_2 \geq 1$, l_1. Furthermore, let the eavesdropper observe the DSS in such a way that $\mathcal{E}_1 \cap \mathcal{E}_2^{\text{sto}} = \emptyset$ and that the eavesdropper is not observing more nodes than needed to recover all the stored data. For the term λ_{MR} from Equation (6.1) it holds that*

$$\lambda_{\text{MR,dir}} \leq \mathrm{H}(\mathsf{E}_2^{\text{rep}} \mid \mathsf{E}_1, \mathsf{E}_2^{\text{sto}}) = \mathrm{H}(\mathsf{E}_1, \mathsf{E}_2^{\text{sto}}, \mathsf{E}_2^{\text{rep}}) - \mathrm{H}(\mathsf{E}_1) - \mathrm{H}(\mathsf{E}_2^{\text{sto}}). \tag{6.5}$$

Proof. We know that $\mathrm{H}(\mathsf{K}) = k$, $\mathrm{H}(\mathsf{E}_1) = l_1$ and $\mathrm{H}(\mathsf{E}_2^{\text{sto}}) = l_2 r$, since this corresponds to the number of symbols directly read by a data collector and the eavesdropper, respectively. Combining this with (6.1) and (6.2) yields

$$\lambda_{\text{MR,dir}} \leq \mathrm{H}(\mathsf{E}_2^{\text{rep}} \mid \mathsf{E}_1, \mathsf{E}_2^{\text{sto}}).$$

By the chain rule of entropy (B.1), it holds that

$$\mathrm{H}(\mathsf{E}_2^{\text{rep}} \mid \mathsf{E}_1, \mathsf{E}_2^{\text{sto}}) = \mathrm{H}(\mathsf{E}_1, \mathsf{E}_2^{\text{sto}}, \mathsf{E}_2^{\text{rep}}) - \mathrm{H}(\mathsf{E}_2^{\text{sto}} \mid \mathsf{E}_1) - \mathrm{H}(\mathsf{E}_1).$$

Following the same arguments as in the proof of Proposition 6.2 with the assumptions that $\mathcal{E}_1 \cap \mathcal{E}_2^{\text{sto}} = \emptyset$ and $\mathcal{E}_1$ and $\mathcal{E}_2^{\text{sto}}$ are independent sets, we have

$$\mathrm{H}(\mathsf{E}_2^{\text{rep}} \mid \mathsf{E}_1, \mathsf{E}_2^{\text{sto}}) = \mathrm{H}(\mathsf{E}_1, \mathsf{E}_2^{\text{sto}}, \mathsf{E}_2^{\text{rep}}) - \mathrm{H}(\mathsf{E}_2^{\text{sto}}) - \mathrm{H}(\mathsf{E}_1).$$

□

With the help of Equation (6.5), λ_{MR} is investigated for direct and forwarded global repair in the following sections. For simplicity, the local encoding of the MR-LRCs is not considered. Only a set of n symbols of the outer codeword $\mathbf{c}_{\text{out}} \in \mathcal{C}_{\text{out}}$ from Construction 4.2 is considered. A subset with cardinality k will then be the minimal global repair set given at most $h = n - k$ erasures, which need to be repaired globally, in the set of n nodes. For any erasure pattern that is correctable by the MR-LRC, it can be guaranteed that after the global repair, the rest of the erasures can be repaired

by the local MDS code. Since the local parities are a $\mathbb{F}_q$-linear combination of the code symbols in the group, a relabeling can be performed locally so that any r nodes of a local group which correspond to a P-independent set of evaluations by Theorem A.8 can be considered for the global repair. More than r nodes of a group would only be a redundant set of symbols.

The number of erasures, which need to be repaired globally, is in the following called global erasures.

6.2 Secrecy Capacity for Direct Global Repair

Before the decrease of the secrecy capacity given a direct global repair in the presence of an (l_1, l_2)-eavesdropper is quantified, some preliminary considerations should be made.

First, note that the secrecy capacity is only decreased if an eavesdropper is observing the group where the global repair is performed in an l_2-manner. The direct global repair does not reveal any information to the other groups which are only sending information. Second, the number of globally repairable erasures is bounded from above by the number of global parities h. If more than h global erasures occur, part of the data cannot be recovered. Third, if the number of global parities h is greater or equal than the number of nodes in a group r, the secrecy capacity will be zero, which is shown in the following proposition.

Proposition 6.4. *Consider a DSS with g groups, r nodes per group and $h \geq r$ global parities. If the system uses direct global repair and $l_2 \geq 1$, its secrecy capacity is zero, i.e.,*

$$k^{(\mathrm{s})}{}_{\mathrm{MR-LRC}} = 0.$$

Proof. Without loss of generality, consider r node failures in a group that is observed by the eavesdropper in an l_2-manner, i.e., all the nodes in the group fail. Every other group sends r symbols to repair the global erasures. The local polynomials L_i have by definition at most r degrees of freedom since the nonzero constraints are only in the respective group with size r. Given r evaluations of the local polynomials L_i, the local polynomial can be generated by Newton interpolation. Evaluating the local polynomials at the respective code locators of their group, the eavesdropper can recover all groups, which yields $k^{(\mathrm{s})}{}_{\mathrm{MR-LRC}} = 0$. □

The preliminary considerations including Proposition 6.4 motivate the following convention.

Direct global repair systems are considered with the constraint $h < r$ since the capacity is zero otherwise. Without loss of generality, the first group is always observed by the

eavesdropper in an l_2-manner where at most h erasures, which have to be repaired globally, occur, i.e., $l_2 \geq 1$. This assumption can be made since it does not matter in which l_2-observed groups the global repairs are performed. The eavesdropper is only able to gain innovation from groups that are not yet fully observed. For $l_2 = 0$ the secrecy capacity $k^{(\mathrm{s})}$ coincides with Theorem 4.3.

The following example illustrates further properties and limitations of direct global repair.

Example 6.1 (Direct global repair with two groups)**.** Consider a DSS with $g = 2$, $r = 2$ and $h = 1$ as shown in Figure 6.1. An eavesdropper observes the first group, where a global erasure occurs, in an l_2-manner. It is assumed that the eavesdropper has read the symbol $c_1^{(1)}$ prior to its erasure. To recover the symbol, the encoding skew polynomial P is evaluated as a linear operator polynomial such that $c_1^{(1)} = P^{D_{a_1}}(\beta_1^{(1)})$. For the direct global repair, the evaluation is split into two evaluations of local polynomials L_1 and L_2. The first local polynomial L_1 is generated in the first group with $c_2^{(1)}$ as a constraint, the second local polynomial L_2 is generated in the second group with $c_1^{(2)}$ and $c_2^{(2)}$ as a constraint. Both are evaluated as linear operator polynomials such that

$$c_1^{(1)} = L_1^{D_{a_1}}(\beta_1^{(1)}) + L_2^{D_{a_1}}(\beta_1^{(1)})$$

where red indicates that the symbols are known to the eavesdropper. Note that the eavesdropper can calculate $L_2^{D_{a_1}}(\beta_1^{(1)})$ given $c_1^{(1)}$ and $L_1^{D_{a_1}}(\beta_1^{(1)})$. Therefore, the global repair process does not reveal any new knowledge to the eavesdropper.

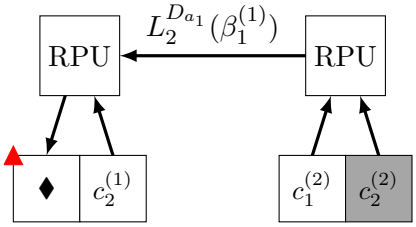

Figure 6.1: Illustration of a DSS with two groups and one global erasure. The global erasure is in the first group which is observed by an eavesdropper in an l_2-manner. The second group sends the evaluation of its local polynomial at the code locator of the failed node to the first group for global repair.

The above example shows that not every symbol which is sent to the eavesdropper reveals new information. Another insight is the following:

Groups that are already fully observed by the eavesdropper cannot reveal any new information.

If the group is fully known to the eavesdropper, the local polynomial can be generated with the information of the group. Thus, the evaluations of the local polynomial will not reveal any new information.

Therefore, the secrecy capacity is crucially linked to the number of groups that are not already fully observed by the eavesdropper and the number of symbols that are not yet known to the eavesdropper these groups.

Proposition 6.5. *Let $\mathcal{G} = [g]$ denote the set of groups of a DSS. A group $j \in \mathcal{G}$ can reveal at most*

$$\min(h, r - e_j)$$

symbols containing new information to the eavesdropper, given h global repairs in a group $m \in \mathcal{G}$, where e_j denotes the number of nodes in the j-th group that are observed by the eavesdropper.

Proof. The j-th group has r symbols and sends h evaluations of its local polynomial L_j to group m. Since there are r independent symbols in each group, $r - e_j$ symbols of the group are unknown to the eavesdropper. Thus, it can reveal at most $\min(h, r - e_j)$ symbols to the eavesdropper within the global repair. □

Another effect that was shown in Example 6.1 is that not every symbol from a group, that is not fully observed by the eavesdropper, reveals new information. If an erased symbol at $(s, m) \in \Delta_{\text{glob}}$, which is about to be globally repaired, was read by the eavesdropper before the erasure occurs, then one summand of

$$c_m^{(s)} = \sum_{i=1}^{g} L_i^{D_{a_s}}(\beta_m^{(s)})$$

will not bring any innovation. Without loss of generality under the assumption that the eavesdropper observes the first group where the global repair is performed, the above equation can be written as

$$c_m^{(1)} - L_1^{D_{a_1}}(\beta_m^{(1)}) = \sum_{i=2}^{g} L_i^{D_{a_1}}(\beta_m^{(1)})$$

where the left hand side is known to the eavesdropper. Therefore, after at most $g - 2$ symbols, which are revealed to the eavesdropper within one global repair, the last symbol

is a result of the other observed symbols with

$$c_m^{(1)} - \sum_{i=1}^{g-1} L_i^{D_{a_1}}(\beta_m^{(1)}) = L_g^{D_{a_1}}(\beta_m^{(1)}).$$

The observation can be generalized to the following proposition.

Proposition 6.6. *Consider a DSS with g groups. Let $h < r$ and fix the eavesdropper parameters l_1 and l_2 with $l_2 \geq 1$ such that $k^{(\mathrm{s})} > 0$. Without loss of generality, it is assumed that the first group is observed by the eavesdropper in an l_2-manner and that h global erasures occur in the first group. For the m-th global repair,*

$$c_m^{(1)} = \left[\sum_{i=1}^{g} L_i({}^{\beta_m^{(1)}} a_1) \right] \beta_m^{(1)},$$

where the sum on the right hand side consists of g linearly independent summands.

Proof. Consider the evaluations in the Lagrange basis over Δ_{glob}, i.e., over $\mathcal{A}_{\Delta_{\mathrm{glob}}} = \{c_{h+1}^{(1)}, \ldots, c_r^{(1)}, \ldots, c_1^{(g)}, \ldots, c_r^{(g)}\}$ with $|\mathcal{A}_{\Delta_{\mathrm{glob}}}| = k$. Rewriting $c_m^{(1)}/\beta_m^{(1)}$ yields

$$c_m^{(1)}/\beta_m^{(1)} = \sum_{i=1}^{g} L_i({}^{\beta_m^{(1)}} a_1) = \sum_{(i,j)\in\Delta_{\mathrm{glob}}} c_j^{(i)} \ell_{i,j}^{\Delta_{\mathrm{glob}}}({}^{\beta_m^{(1)}} a_1).$$

The g evaluations of local polynomials L_i can therefore be written as a product of $(c_{h+1}^{(1)}, \ldots, c_r^{(1)}, \ldots, c_1^{(g)}, \ldots, c_r^{(g)})$ and the matrix $\mathbf{M}_L$ defined below:

$$\underbrace{\begin{pmatrix} \ell_{1,h+1}^{\Delta_{\mathrm{glob}}}({}^{\beta_m^{(1)}} a_1) & \cdots & \ell_{1,r}^{\Delta_{\mathrm{glob}}}({}^{\beta_m^{(1)}} a_1) & 0 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ 0 & \cdots & 0 & \ell_{2,1}^{\Delta_{\mathrm{glob}}}({}^{\beta_m^{(1)}} a_1) & \cdots & \ell_{2,r}^{\Delta_{\mathrm{glob}}}({}^{\beta_m^{(1)}} a_1) & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 0 & \cdots & 0 & \cdots & \ell_{g,1}^{\Delta_{\mathrm{glob}}}({}^{\beta_m^{(1)}} a_1) & \cdots & \ell_{g,r}^{\Delta_{\mathrm{glob}}}({}^{\beta_m^{(1)}} a_1) \end{pmatrix}}_{=: \mathbf{M}_L}.$$

Inspecting the rows of the matrix, we see that the vector representation in terms of $\mathcal{A}_{\Delta_{\mathrm{glob}}}$ of the evaluations of the local polynomials are pairwise orthogonal since their dot product is zero. Therefore, the evaluations are linearly independent with respect to $\mathcal{A}_{\Delta_{\mathrm{glob}}}$. It follows that

$$\mathrm{rank}(\mathbf{M}_L) = g$$

which means that all the summands are linearly independent. □

Remark: For $h < r$, it always holds that $g \leq k$ since $k = n - h > rg - r = g - 1$.

Proposition 6.7. *Let the assumptions be the same as in Proposition 6.6. If an eavesdropper reads the symbols before the global erasure occurs, the maximum knowledge gain for one global repair will be $g-1$ symbols.*

Proof. It holds that

$$c_m^{(1)}/\beta_m^{(1)} = \sum_{i=1}^{g} L_i(\beta_m^{(1)} a_l) = \sum_{(i,j)\in\Delta_{\text{glob}}} c_j^{(i)} \ell_{i,j}^{\Delta_{\text{glob}}}(\beta_m^{(1)} a_l)$$

which means that $c_m^{(1)}$ can be represented by a linear combination of the g independent evaluations of local polynomials. Stacking $\mathbf{M}_L \in \mathbb{F}_{q^m}^{g\times k}$ for the m-th global repair from the lemma above and

$$c_m^{(1)} = \left(\ell_{1,h+1}^{\Delta_{\text{glob}}}(\beta_m^{(1)} a_1) \quad \ell_{1,h+2}^{\Delta_{\text{glob}}}(\beta_m^{(1)} a_1) \quad \cdots \quad \ell_{g,r}^{\Delta_{\text{glob}}}(\beta_m^{(1)} a_1)\right)$$

yields $\mathbf{M}_{L,c} \in \mathbb{F}_{q^m}^{g+1\times k}$. For the rank of the matrix $\mathbf{M}_{L,c}$, we have $\operatorname{rank}(\mathbf{M}_{L,c}) = g$ because the codeword $c_m^{(1)}$ is the sum of all rows of the matrix $\mathbf{M}_L$. Thus, it does not change the rank. Since the codeword symbols $c_j^{(i)}$ are uniformly distributed and independent within the global repair set Δ_{glob}, (3.6) can be used, which yields

$$\begin{aligned}\mathrm{H}(\mathsf{M}_{L,c} \mid \mathsf{X}_{c_m^{(1)}}) &= \mathrm{H}(\mathsf{M}_{L,c}, \mathsf{X}_{c_m^{(1)}}) - \mathrm{H}(\mathsf{X}_{c_m^{(1)}}) \\ &= \operatorname{rank}(\mathbf{M}_{L,c}) - \operatorname{rank}(c_m^{(1)}) = g-1,\end{aligned}$$

where $\mathsf{M}_{L,c}$ and $\mathsf{X}_{c_m^{(1)}}$ denote the random variables corresponding to the repair symbols of the m-th global repair and the repair symbol represented by the matrix $\mathbf{M}_{L,c}$ and the m-th symbol, respectively. □

This effect can be later on seen in the term for the secrecy capacity.

We now want to find an expression for the upper bound based on the parameters of the system. As a starting point (6.5) from Proposition 6.3, establishing an upper bound on the knowledge that can be gained by an eavesdropper due to global repair, can be used. The idea is to express the entropies with matrices using Lemma 3.7.

Let $\mathbf{M}_{\mathcal{E}_1}$ represent the symbols accessed by an eavesdropper in an l_1-manner, let $\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}$ represent the symbols accessed by an eavesdropper in an l_2-manner, let $\mathbf{M}_L$ represent the symbols sent to the groups observed by the eavesdropper in an l_2-manner and denote by $\mathbf{M}_{\text{joint}}$ the stacked matrix with all the symbols of $\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}$,$\mathbf{M}_{\mathcal{E}_1}$ and $\mathbf{M}_L$.

Remark: The matrices are written with respect to a Lagrange basis in the set Δ_{glob}. The Lagrange basis consists of Lagrange skew polynomials. Therefore to transform the matrices to matrices representing the linearized Reed-Solomon codes, they have to be

multiplied by a diagonal matrix. However, the diagonal matrices are neglected in the following to ease readability. This is possible since the ranks of the matrices are of interest and the diagonal matrices have full rank. In Example 6.2, the diagonal matrices are written out explicitly.

Definition 6.1. Consider a DSS which is observed by an (l_1, l_2)-eavesdropper. All matrices are expressed in terms of Δ_{glob} with cardinality $|\Delta_{\text{glob}}| = k$. The matrices $\mathbf{M}_{\mathcal{E}_1} \in \mathbb{F}_{q^m}^{l_1 \times k}$ and $\mathbf{M}_{\mathcal{E}_2^{\text{sto}}} \in \mathbb{F}_{q^m}^{l_2 r \times k}$ represent the sets $\mathcal{E}_1$ and $\mathcal{E}_2^{\text{sto}}$, respectively. The set of symbols that are sent to groups observed in an l_2-manner, i.e., $\mathcal{E}_2^{\text{rep}}$, is represented by the matrix $\mathbf{M}_L \in \mathbb{F}_{q^m}^{gh \times k}$. If only the first $m \in [h]$ global repairs are considered, we write $\mathbf{M}_L^m \in \mathbb{F}_{q^m}^{gm \times k}$. The matrices can be stacked to form a matrix, representing all symbols that the eavesdropper observes, yielding

$$\begin{pmatrix} \mathbf{M}_{\mathcal{E}_2^{\text{sto}}} \\ \mathbf{M}_{\mathcal{E}_1} \\ \mathbf{M}_L \end{pmatrix} =: \mathbf{M}_{\text{joint}} \in \mathbb{F}_{q^m}^{(l_2 r + l_1 + gh) \times k}.$$

If only a part of $\mathbf{M}_L$ is stacked, i.e., $\mathbf{M}_L^m$, the corresponding stacked matrix is $\mathbf{M}_{\text{joint}}^m$.

The stacked matrix $\mathbf{M}_{\text{joint}}$, representing the joint entropy of all observations of the eavesdropper, can have at most rank k. If the rank of the matrix $\mathbf{M}_{\text{joint}}^m$ is k for $m < h$, it means that the eavesdropper has global knowledge after m global repairs. This observation motivates the following definition.

Definition 6.2. Consider a DSS with g groups, parameters $h < r$ and fixed eavesdropper parameters $l_2 \geq 1$, l_1. Without loss of generality, it is assumed that the first group is observed by the eavesdropper in an l_2-manner and that h global erasures occur in the first group. Let $h_{\min}$ be the minimal number of repairs that need to be performed until either $h_{\min} = h$ or $\operatorname{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) = k$ with $h_{\min} < h$.

The structure of the matrices stated in Definition 6.1 is now analyzed. The matrix $\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}$ has the structure

$$\mathbf{M}_{\mathcal{E}_2^{\text{sto}}} = \begin{pmatrix} & & \mathbf{M}_{h_{\min}} & & \\ \mathbf{I}_{r-h_{\min}} & \cdots & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{I}_r & \cdots & \mathbf{0} \end{pmatrix} \in \mathbb{F}_{q^m}^{l_2 r \times k}$$

with full rank identity matrices of size r at the indices of groups observed in an l_2-manner. The first group is split in an identity matrix of size $r - h_{\min}$ representing the notes that

are in the global repair set Δ_{glob} and the $h_{\min}$ nodes that are being repaired. The matrix $\mathbf{M}_{h_{\min}}$ is the representation of the failed nodes in terms of the repair set Δ_{glob} with

$$\mathbf{M}_{h_{\min}} = \begin{pmatrix} \ell_{1,h_{\min}+1}^{\Delta_{\text{glob}}}(\beta_1^{(1)} a_1) & \ell_{1,h_{\min}+2}^{\Delta_{\text{glob}}}(\beta_1^{(1)} a_1) & \cdots & \ell_{g,r}^{\Delta_{\text{glob}}}(\beta_1^{(1)} a_1) \\ \vdots & \vdots & \ddots & \vdots \\ \ell_{1,h_{\min}+1}^{\Delta_{\text{glob}}}(\beta_{h_{\min}}^{(1)} a_1) & \ell_{1,h_{\min}+2}^{\Delta_{\text{glob}}}(\beta_{h_{\min}}^{(1)} a_1) & \cdots & \ell_{g,r}^{\Delta_{\text{glob}}}(\beta_{h_{\min}}^{(1)} a_1) \end{pmatrix} \in \mathbb{F}_{q^m}^{h_{\min} \times k}.$$

The matrix $\mathbf{M}_{\mathcal{E}_1}$ with l_1 rows has l_1 nonzero entries at the positions of the observed nodes in the global repair set Δ_{glob}. The matrix $\mathbf{M}_L$ has a row for each symbol that is sent to the group where the global repair is performed. The elements of each row are a subset of the corresponding rows in $\mathbf{M}_{h_{\min}}$.

The structure of the matrices can be best illustrated with an example.

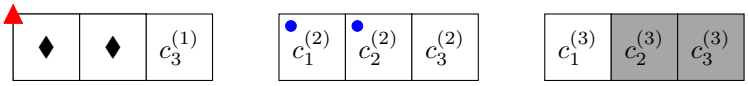

Figure 6.2: Illustration of a DSS with $g = 3$ groups, eavesdropper parameters $l_1 = 2$, $l_2 = 1$ and two global erasures in the first group.

Example 6.2. Consider a DSS using the coding scheme from Construction 4.2 and parameters $r = 3$, $g = 3$ and $h = 2$. An eavesdropper is observing the first group with two global erasures in an l_2-manner. The number of nodes that are in addition observed by the eavesdropper is $l_1 = 2$ and they are both in the second group. The system is depicted in Figure 6.2. The vector corresponding to the global repair set Δ_{glob} is

$$\begin{pmatrix} c_3^{(1)} & c_1^{(2)} & c_2^{(2)} & c_3^{(2)} & c_1^{(3)} & c_2^{(3)} & c_3^{(3)} \end{pmatrix}.$$

The following matrices are with respect to this vector. The matrix corresponding to the nodes observed in an l_2-manner is

$$\mathbf{M}_{\mathcal{E}_2^{\text{sto}}} = \begin{pmatrix} \beta_1^{(1)} & 0 & 0 \\ 0 & \beta_2^{(1)} & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} \ell_{1,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) \\ \ell_{1,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

where $b_j^{(i)} = {}^{\beta_j^{(i)}} a_i$. The matrix corresponding to the l_1-observations is

$$\mathbf{M}_{\mathcal{E}_1} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix},$$

and

$$\mathbf{M}_L = \mathbf{D}_L \begin{pmatrix} \ell_{1,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \ell_{2,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \ell_{3,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) \\ \ell_{1,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \ell_{2,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \ell_{3,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) \end{pmatrix}$$

with $\mathbf{D}_L = \text{diag}(\beta_1^{(1)}, \beta_1^{(1)}, \beta_1^{(1)}, \beta_2^{(1)}, \beta_2^{(1)}, \beta_2^{(1)})$ is the matrix corresponding to the symbols sent to the first group for the global repairs. Thus, the matrix $\mathbf{M}_{\text{joint}}$ has the form

$$\mathbf{M}_{\text{joint}} = \mathbf{D} \begin{pmatrix} \ell_{1,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) \\ \ell_{1,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ \ell_{1,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \ell_{2,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{2,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \ell_{3,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) \\ \ell_{1,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \ell_{2,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{2,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \ell_{3,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) \end{pmatrix}$$

with

$$\mathbf{D} = \text{diag}(\beta_1^{(1)}, \beta_2^{(1)}, 1, 1, 1, \beta_1^{(1)}, \beta_1^{(1)}, \beta_1^{(1)}, \beta_2^{(1)}, \beta_2^{(1)}, \beta_2^{(1)}).$$

By Equation (6.5) and Lemma 3.7, $\lambda_{\text{MR,dir}}$ can be bounded from above by

$$\lambda_{\text{MR,dir}} \leq \text{H}(\mathsf{E}_1, \mathsf{E}_2^{\text{sto}}, \mathsf{E}_2^{\text{rep}}) - \text{H}(\mathsf{E}_1) - \text{H}(\mathsf{E}_2^{\text{sto}}) = \text{rank}(\mathbf{M}_{\text{joint}}) - \text{rank}(\mathbf{M}_{\mathcal{E}_1}) - \text{rank}(\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}),$$

since the matrices $\mathbf{M}_{\text{joint}}$,$\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}$ and $\mathbf{M}_{\mathcal{E}_1}$ represent $\mathsf{E} = (\mathsf{E}_1, \mathsf{E}_2^{\text{sto}}, \mathsf{E}_2^{\text{rep}})$, $\mathsf{E}_2^{\text{sto}}$ and E_1, respectively.

Clearly, the first two rows of $\mathbf{M}_{\text{joint}}$ are a linear combination of the repair symbols

and thus they do not contribute to the rank. In addition, the rows of the eavesdropper observation with only one nonzero entry are independent of the other rows and contribute 3 to the rank. This can be seen by using Gaussian elimination yielding

$$\mathbf{M'}_{\text{joint}} = \mathbf{D} \begin{pmatrix} 0 & 0 & 0 & \ell_{2,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) \\ 0 & 0 & 0 & \ell_{2,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \ell_{2,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \ell_{3,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \ell_{2,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \ell_{3,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_2^{(1)}). \end{pmatrix}.$$

The rank of $\mathbf{M'}_{\text{joint}}$ can then be derived by calculating the rank of the two submatrices from column 1 to 3 and column 4 to 7. Obviously, for

$$\mathbf{M'}_{\text{joint}}|_{\Delta_{\text{glob}} \cap (\mathcal{E}_2^{\text{sto}} \cup \mathcal{E}_1)} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

it holds that $\text{rank}(\mathbf{M'}_{\text{joint}}|_{\Delta_{\text{glob}} \cap (\mathcal{E}_2^{\text{sto}} \cup \mathcal{E}_1)}) = 3$. For

$$\mathbf{M'}_{\text{joint}}|_{\Delta_{\text{glob}} \setminus (\mathcal{E}_1, \mathcal{E}_2^{\text{sto}})} = \mathbf{D}_x \begin{pmatrix} \ell_{2,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) & 0 & 0 & 0 \\ 0 & \ell_{3,1}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_1^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_1^{(1)}) \\ \ell_{2,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) & 0 & 0 & 0 \\ 0 & \ell_{3,1}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,2}^{\Delta_{\text{glob}}}(b_2^{(1)}) & \ell_{3,3}^{\Delta_{\text{glob}}}(b_2^{(1)}) \end{pmatrix} \tag{6.6}$$

with

$$\mathbf{D}_x = \text{diag}(\beta_1^{(1)}, \beta_1^{(1)}, \beta_2^{(1)}, \beta_2^{(1)}),$$

we also have $\text{rank}(\mathbf{M'}_{\text{joint}}|_{\Delta_{\text{glob}} \setminus (\mathcal{E}_1, \mathcal{E}_2^{\text{sto}})}) = 3$. For the second matrix, the first and third row are linearly dependent. Therefore, the two rows contribute only 1 to the rank while the second and fourth row contribute 2 to the rank since they are independent by Lemma 3.3. The rows correspond to evaluations of the same polynomial but at two P-independent points and can therefore be decomposed into two Vandermonde matrices with full rank. Following the same arguments, we have $\text{rank}(\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}) = 3$, where the first two rows con-

tribute 2 to the rank of $\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}$ by Lemma 3.3.

Overall, it holds that

$$\begin{aligned}\lambda_{\text{MR,dir}} &\leq \text{rank}(\mathbf{M}'_{\text{joint}}|_{\Delta_{\text{glob}}\cap(\mathcal{E}_2^{\text{sto}}\cup\mathcal{E}_1)}) + \text{rank}(\mathbf{M}'_{\text{joint}}|_{\Delta_{\text{glob}}\setminus(\mathcal{E}_1,\mathcal{E}_2^{\text{sto}})}) - \text{rank}(\mathbf{M}_{\mathcal{E}_1}) - \text{rank}(\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}) \\ &= 3+3-2-3=1.\end{aligned}$$

This means that the secrecy capacity of the DSS using direct global repair is decreased by 1.

With this example in mind, a general expression for the secrecy capacity decrease $\lambda_{\text{MR,dir}}$ can be derived.

Lemma 6.1 (**Direct global repair - secrecy capacity upper bound**). *Consider a DSS with g groups, parameters $h < r$ and an eavesdropper with $l_2 \geq 1$, l_1. Without loss of generality, it is assumed that the first group is observed by the eavesdropper in an l_2-manner and that h global erasures occur in the first group. Let $h_{\min}$ be the minimal number of global repairs as defined in Definition 6.2. The secrecy capacity decrease for direct global repair $\lambda_{\text{MR,dir}}$ is upper bounded by*

$$\lambda_{\text{MR,dir}} \leq \text{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) - l_2 r - l_1 =: \bar{\lambda}_{\text{MR,dir}}.$$

Moreover, it holds that

$$\bar{\lambda}_{\text{MR,dir}} = \left(\sum_{i=1}^{g} \min(h_{\min}, r-e_i)\right) - h_{\min} \tag{6.7}$$

where e_i denotes the number of symbols that the eavesdropper is observing in the i-group.

Proof. The upper bound $\bar{\lambda}_{\text{MR,dir}}$ can be deduced from Equation (6.5). The proof focuses on showing that (6.7) holds.
The first $h_{\min}$ rows of $\mathbf{M}_{\text{joint}}^{h_{\min}}$ are a linear combination of the rows of matrix $\mathbf{M}_L^{h_{\min}}$ and can therefore be neglected calculating the rank. The contribution to the rank of $\mathbf{M}_{\text{joint}}^{h_{\min}}$ of nodes read directly by the eavesdropper, i.e., the number of rows with a single nonzero entry in the matrix $\mathbf{M}_{\text{joint}}^{h_{\min}}$, is $(l_2-1)r + r - h_{\min} + l_1$. This can be inferred by using Gaussian elimination. Thus, it holds that

$$\text{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) = (l_2-1)r + r - h_{\min} + l_1 + \text{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}|_{\Delta_{\text{glob}}\setminus(\mathcal{E}_1,\mathcal{E}_2^{\text{sto}})}),$$

where

$$\mathbf{M}_{\text{joint}}^{h_{\min}}|_{\Delta_{\text{glob}}\setminus(\mathcal{E}_1,\mathcal{E}_2^{\text{sto}})} = \mathbf{M}_L^{h_{\min}}|_{\Delta_{\text{glob}}\setminus(\mathcal{E}_1,\mathcal{E}_2^{\text{sto}})}$$

holds. The matrix $\mathbf{M}_{\text{joint}}^{h_{\min}}|_{\Delta_{\text{glob}}\setminus(\mathcal{E}_1,\mathcal{E}_2^{\text{sto}})}$ that is left after puncturing the nonzero entries has a similar structure to the matrix in (6.6). Consider the matrix groupwise for the i-th group. If the i-th group is fully punctured, we have $e_i = r$ and there is no column in $\mathbf{M}_{\text{joint}}^{h_{\min}}|_{\Delta_{\text{glob}}\setminus(\mathcal{E}_1,\mathcal{E}_2^{\text{sto}})}$ corresponding to the i-th group. Otherwise, there are still $r - e_i$ columns corresponding to the i-th group. The rows corresponding to the i-th row are of the structure investigated in Lemma 3.3. They can be represented by the product of two Vandermonde matrices since they correspond to evaluations of the same polynomial at P-independent points. We can therefore bound the contribution of the i-th group, which has full rank, by the number of rows or columns, i.e., $\min(h_{\min}, r - e_i)$ (see Proposition 6.5). Thus, it holds that

$$\operatorname{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}|_{\Delta_{\text{glob}}\setminus(\mathcal{E}_1,\mathcal{E}_2^{\text{sto}})}) = \sum_{i=1}^{g} \min(h_{\min}, r - e_i)$$

and we have

$$\begin{aligned}\operatorname{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) &= (l_2 - 1)r + r - h_{\min} + l_1 + \sum_{i=1}^{g} \min(h_{\min}, r - e_i) \\ &= l_2 r + l_1 - h_{\min} + \sum_{i=1}^{g} \min(h_{\min}, r - e_i).\end{aligned}$$

This yields

$$\begin{aligned}\bar{\lambda}_{\text{MR,dir}} &= \operatorname{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) - l_2 r - l_1 = l_2 r + l_1 - h_{\min} + \sum_{i=1}^{g} \min(h_{\min}, r - e_j) - l_2 r - l_1 \\ &= \left(\sum_{i=1}^{g} \min(h_{\min}, r - e_j)\right) - h_{\min}.\end{aligned}$$

□

Example 6.2 can now also be verified with Equation (6.7). For the considered DSS and eavesdropper parameters, we have

$$\begin{aligned}\bar{\lambda}_{\text{MR,dir}} &= \left(\sum_{i=1}^{3} \min(2, 3 - e_i)\right) - 2 \\ &= \min(2,0) + \min(2,1) + \min(2,3) - 2 = 1 + 2 - 2 = 1,\end{aligned}$$

which is in accordance with the derivations.

The upper bounds, that have been derived, can be achieved with equality as the

following theorem shows.

Theorem 6.1 (**Direct global repair - secrecy capacity**). *Consider a DSS with g groups, parameters $h < r$ and an (l_1, l_2)-eavesdropper with $l_2 \geq 1$ and fixed l_1 such that $l_2 r + l_1 \leq k$. Without loss of generality, it is assumed that the first group is observed by the eavesdropper in an l_2-manner and that h global erasures occur in the first group. Let $h_{\min}$ be the minimal number of global repairs as defined in Definition 6.2. The secrecy capacity for direct global repair is*

$$k^{(\mathrm{s})}{}_{\mathrm{dir}} = k - (l_2 r + l_1) - \left(\left[\sum_{i=1}^{g} \min(h_{\min}, r - e_i) \right] - h_{\min} \right), \tag{6.8}$$

where e_i denotes the number of symbols that the eavesdropper is observing in the i-th group.

Proof. It follows from Lemma 6.1 that the right hand side of (6.8) is an upper bound. The proof is done by showing that Construction 6.1 achieves equality (6.8). □

Theorem 6.2. *Construction 6.1 is information-theoretically secure against an (l_1, l_2)-eavesdropper and achieves the secrecy capacity $k^{(\mathrm{s})}{}_{\mathrm{dir}} = k - k^{(e)}$ with $k^{(e)} = (l_2 r + l_1 + \bar{\lambda}_{\mathrm{MR,dir}})$ from Theorem 6.1 if direct global repair is used.*

Proof. To prove secrecy of the coding scheme, the secrecy lemma (Lemma 3.1) is used. An (l_1, l_2)-eavesdropper can at most observe $l_2 r + l_1 + \bar{\lambda}_{\mathrm{MR,dir}}$ symbols. There are $l_2 r + l_1 + \bar{\lambda}_{\mathrm{MR,dir}}$ randomly generated symbols. Thus, the first condition is fulfilled, i.e., $\mathrm{H}(\mathsf{E}) \leq \mathrm{H}(\mathsf{R})$ where E and R are the random variables corresponding to the eavesdropper's observation and the randomly generated symbols, respectively. It remains to prove that $\mathrm{H}(\mathsf{R} \mid \mathsf{U}^{(\mathrm{s})}, \mathsf{E}) = 0$. Without loss of generality, it can be assumed that the global repairs are performed in the first group which is observed in an l_2-manner by the eavesdropper. The global repair set Δ_{glob} is chosen in such a way that it overlaps as much as possible with the static eavesdropper observations $\mathcal{E}_1$ and $\mathcal{E}_2^{\mathrm{sto}}$. Since the local encoding is MDS, the punctured local encoding matrix $\mathbf{A}|_{\Delta_{\mathrm{glob}}} = \mathrm{diag}(\mathbf{A}_1, \ldots, \mathbf{A}_g)|_{\Delta_{\mathrm{glob}}}$ has rank k. The eavesdropped information can be summarized in the matrix $\mathbf{M}_{\mathrm{joint}}^{h_{\min}}$. Thus, we have

$$\underbrace{\mathbf{A}|_{\Delta_{\mathrm{glob}}} \mathbf{D}^{h_{\min}} \mathbf{M}_{\mathrm{joint}}^{h_{\min}}}_{=:\mathbf{E}_{\mathrm{joint}}} \mathbf{c}_{\Delta_{\mathrm{glob}}} = \mathbf{e}_{\mathcal{E}},$$

where $\mathbf{c}_{\Delta_{\mathrm{glob}}}$ represents the set of code symbols after the first encoding step with a skew Reed-Solomon code which yields a linearized Reed-Solomon code with the column multiplier matrix $\mathbf{D}^{h_{\min}}$. The eavesdropped symbols are denoted by $\mathbf{e}_{\mathcal{E}}$. Since $\mathbf{A}|_{\Delta_{\mathrm{glob}}}$ and

$\mathbf{D}^{h_{\min}}$ have full rank k, the matrix $\mathbf{E}_{\text{joint}}$ has the same rank as $\mathbf{M}_{\text{joint}}^{h_{\min}}$ which is bounded by $\operatorname{rank}(\mathbf{E}_{\text{joint}}) = \operatorname{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) = l_2 r + l_1 + \bar{\lambda}_{\text{MR,dir}} \leq k$ by Lemma 6.1. The matrix $\mathbf{E}_{\text{joint}}$ can be transformed in the domain of the coefficients of the encoding skew polynomial by Definition 3.17 yielding

$$\mathbf{E}_{\text{joint}}(\mathbf{V}_k^\sigma(\mathbf{b}))^{-1}\mathbf{f} = \mathbf{e}_{\mathcal{E}}(\mathbf{V}_k^\sigma(\mathbf{b}))^{-1},$$

where $\mathbf{f}$ denotes the coefficients of the encoding skew polynomial F. The matrix $\mathbf{E}_{\text{joint}}(\mathbf{V}_k^\sigma(\mathbf{b}))^{-1}$ gives $l_2 r + l_1 + \bar{\lambda}_{\text{MR,dir}}$ independent constraints on the polynomial coefficient. Together with the ${k^{(\text{s})}}_{\text{dir}}$ coefficients of the information symbols we have k constraints on k coefficients. It remains to show that the k constraints are linearly independent. The ${k^{(\text{s})}}_{\text{dir}}$ coefficients are known and can therefore be written as a row with one nonzero entry expanding the matrix $\mathbf{E}_{\text{joint}}(\mathbf{V}_k^\sigma(\mathbf{b}))^{-1}$. If the Vandermonde matrix $(\mathbf{V}_k^\sigma(\mathbf{b}))^{-1}$ is punctured at ${k^{(\text{s})}}_{\text{dir}}$ columns, it still has rank $k - {k^{(\text{s})}}_{\text{dir}}$ due to its structure. Multiplied with $\mathbf{E}_{\text{joint}}$ the overall rank of $\mathbf{E}_{\text{joint}}(\mathbf{V}_k^\sigma(\mathbf{b}))^{-1}$ is therefore still $k - {k^{(\text{s})}}_{\text{dir}} = l_2 r + l_1 + \bar{\lambda}_{\text{MR,dir}}$. Thus, the system of equations has enough linearly independent equations to determine the randomly generated symbols $k - {k^{(\text{s})}}_{\text{dir}}$ and $\mathrm{H}(\mathsf{R} \mid \mathsf{U}^{(\text{s})}, \mathsf{E}) = 0$ holds. This shows that the bound from 6.1 can be achieved with equality. □

6.3 Secrecy Capacity for Forwarded Global Repair

Before quantifying of the secrecy capacity given a forwarded global repair in the presence of an (l_1, l_2)-eavesdropper, some preliminary considerations should be made. In the following, the forwarded global repair is only considered for $l_2 = 1$.

First, note that the secrecy capacity is only decreased if an eavesdropper is observing a group in an l_2-manner, where the global repair is not performed, which is formulated in the following proposition.

Proposition 6.8. *Consider a DSS with g groups. Let $l_2 = 1$ be in the m-group where global erasure occurs. The eavesdropper is able to read the node before the erasure occurs. Then, the eavesdropper will not gain any new knowledge within the repair.*

Proof. In the corresponding forwarding tuple, the m-th group is by Definition 5.4 at the very end. Thus it receives one symbol which is the contribution of the other groups in Δ_{glob} to the repair process. The repair of the j-th node is characterized by

$$c_j^{(m)} = L_m({}^{\beta_j^{(m)}} a_m)\beta_j^{(m)} + \underbrace{\sum_{i \in \mathcal{G}\setminus\{m\}} L_i({}^{\beta_j^{(m)}} a_m)\beta_j^{(m)}}_{x_{\text{rx}}}.$$

The eavesdropper receives x_{rx} but the uncertainty of x_{rx} given the symbol that is about to be repaired $c_j^{(m)}$ and the contribution of the m-th group to the repair processes is zero, i.e.,

$$x_{\text{rx}} = c_j^{(m)} - L_m({}^{\beta_j^{(m)}} a_m)\beta_j^{(m)}.$$

□

Another insight is that the eavesdropper may only gain new knowledge if it is not in the first group of the forwarding list. By Definition 5.4, the first group of the forwarding list does not receive any symbol. This means that if the system has only two groups and one or multiple global repairs are needed, the gain of knowledge given the global repairs is zero. Forwarded global repair therefore behaves equivalently to direct global repair for $g = 2$; see Example 6.1.

For $g \geq 2$, only a particular case is considered in this work, so that the result is comparable to the result of the direct global repair. Let $h < r$ and let the global repairs only be in one group, so that the forwarding list does not change between different global repairs. Let $l_2 = 1$ not be located in the first group of the forwarding list.

We now want to be able to find an expression for the upper bound based on the parameters of the system similarly as for direct global repair in the previous section. As a starting point (6.5) from Proposition 6.3, establishing an upper bound on the knowledge that can be gained by an eavesdropper due to global repair, can be used. The idea is to express the entropies with matrices using Lemma 3.7.

The notation is the same as for direct global repair. Let $\mathbf{M}_{\mathcal{E}_1}$ represent the symbols accessed by an eavesdropper in an l_1-manner, let $\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}$ represent the symbols accessed by an eavesdropper in an l_2-manner, let $\mathbf{M}_L$ represent the symbols sent to the groups observed by the eavesdropper in an l_2-manner and denote by $\mathbf{M}_{\text{joint}}$ the stacked matrix with all the symbols of $\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}$,$\mathbf{M}_{\mathcal{E}_1}$ and $\mathbf{M}_L$. The only difference to the matrices in Definition 6.1 is that $\mathbf{M}_L$ is now a matrix of size $h \times k$ instead of $gh \times k$.

The stacked matrix $\mathbf{M}_{\text{joint}}$, representing the joint entropy of all observations of the eavesdropper, can have at most rank k. If the rank of the matrix $\mathbf{M}_{\text{joint}}^m$ is k for $m < h$, it means that the eavesdropper has global knowledge after m global repairs. This observation motivates the following definition.

Definition 6.3. Consider a DSS with g groups, parameters $h < r$ and fixed eavesdropper parameters $l_2 \geq 1$, l_1. Without loss of generality, it is assumed that h global erasures occur in the same group and that this group is not observed in an l_2-manner. Let $h_{\min}$ be the minimal number of repairs that need to be performed until either $h_{\min} = h$ or $\text{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) = k$ with $h_{\min} < h$.

The structure of the matrices stated in Definition 6.1 is now analyzed. The global repair set Δ_{glob} is chosen to be from the set of nodes, which were not erased. Therefore, the matrices $\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}$ and $\mathbf{M}_{\mathcal{E}_1}$ only have rows with single nonzero entries. The matrix $\mathbf{M}_L$ is also slightly different and has the structure

$$\mathbf{M}_L = \begin{pmatrix} \ell_{1,1}^{\Delta_{\text{glob}}}(\beta_1^{(g)} a_g) & \ell_{1,2}^{\Delta_{\text{glob}}}(\beta_1^{(g)} a_g) & \cdots & \ell_{s-1,r}^{\Delta_{\text{glob}}}(\beta_1^{(g)} a_g) & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \ell_{1,1}^{\Delta_{\text{glob}}}(\beta_{h_{\min}}^{(g)} a_g) & \ell_{1,2}^{\Delta_{\text{glob}}}(\beta_{h_{\min}}^{(g)} a_g) & \cdots & \ell_{s-1,r}^{\Delta_{\text{glob}}}(\beta_{h_{\min}}^{(g)} a_g) & 0 & \cdots & 0 \end{pmatrix} \in \mathbb{F}_{q^m}^{h_{\min} \times k},$$

given that the erasures occur in group g and group s is observed in an l_2-manner.

The structure of the matrices can be best illustrated with an example.

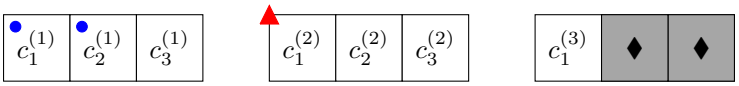

Figure 6.3: Illustration of a DSS with $g = 3$ groups, eavesdropper parameters $l_1 = 2$, $l_2 = 1$ and two global erasures in the third group.

Example 6.3. Consider a DSS with $r = 3$, $g = 3$ and $h = h_{\min} = 2$. An eavesdropper is observing the second group in an l_2-manner while two global erasures occur in the third group. The number of l_1 eavesdropped nodes is two and both nodes are in the first group. The forwarding list for the global repair is $\mathbf{f} = (1, 2, 3)$. The system is displayed in Figure 6.3. The vector corresponding to the global repair set Δ_{glob} is

$$\begin{pmatrix} c_1^{(1)} & c_2^{(1)} & c_3^{(1)} & c_1^{(2)} & c_2^{(2)} & c_3^{(2)} & c_1^{(3)} \end{pmatrix}.$$

The following matrices are with respect to this vector. The matrix that corresponds to the nodes observed in an l_2-manner is

$$\mathbf{M}_{\mathcal{E}_2^{\text{sto}}} = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

The matrix corresponding to the l_1 nodes that are observed in addition is

$$\mathbf{M}_{\mathcal{E}_1} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix},$$

and

$$\mathbf{M}_L^1 = \begin{pmatrix} \beta_2^{(3)} & 0 \\ 0 & \beta_3^{(3)} \end{pmatrix} \begin{pmatrix} \ell_{1,1}^{\Delta_{\text{glob}}}(b_2^{(3)}) & \ell_{1,2}^{\Delta_{\text{glob}}}(b_2^{(3)}) & \ell_{1,3}^{\Delta_{\text{glob}}}(b_2^{(3)}) & 0 & 0 & 0 & 0 \\ \ell_{1,1}^{\Delta_{\text{glob}}}(b_3^{(3)}) & \ell_{1,2}^{\Delta_{\text{glob}}}(b_3^{(3)}) & \ell_{1,3}^{\Delta_{\text{glob}}}(b_3^{(3)}) & 0 & 0 & 0 & 0 \end{pmatrix}$$

is the matrix corresponding to the symbols that are received by the second group for global repairs in the third group. Thus, the matrix $\mathbf{M}_{\text{joint}}$ has the form

$$\mathbf{M}_{\text{joint}} = \mathbf{D} \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \ell_{1,1}^{\Delta_{\text{glob}}}(b_2^{(3)}) & \ell_{1,2}^{\Delta_{\text{glob}}}(b_2^{(3)}) & \ell_{1,3}^{\Delta_{\text{glob}}}(b_2^{(3)}) & 0 & 0 & 0 & 0 \\ \ell_{1,1}^{\Delta_{\text{glob}}}(b_3^{(3)}) & \ell_{1,2}^{\Delta_{\text{glob}}}(b_3^{(3)}) & \ell_{1,3}^{\Delta_{\text{glob}}}(b_3^{(3)}) & 0 & 0 & 0 & 0 \end{pmatrix}$$

with

$$\mathbf{D} = \text{diag}(1, 1, 1, 1, 1, \beta_2^{(3)}, \beta_3^{(3)}).$$

By Equation (6.5) and Lemma 3.7, $\lambda_{\text{MR,for}}$ can be bounded from above by

$$\lambda_{\text{MR,for}} \leq \text{H}(\mathsf{E}_1, \mathsf{E}_2^{\text{sto}}, \mathsf{E}_2^{\text{rep}}) - \text{H}(\mathsf{E}_1) - \text{H}(\mathsf{E}_2^{\text{sto}}) = \text{rank}(\mathbf{M}_{\text{joint}}) - \text{rank}(\mathbf{M}_{\mathcal{E}_1}) - \text{rank}(\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}),$$

since the matrices $\mathbf{M}_{\text{joint}}$,$\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}$ and $\mathbf{M}_{\mathcal{E}_1}$ represent $\mathsf{E} = (\mathsf{E}_1, \mathsf{E}_2^{\text{sto}}, \mathsf{E}_2^{\text{rep}})$, $\mathsf{E}_2^{\text{sto}}$ and E_1, respectively.

The rank of the matrix $\mathbf{M}_{\text{joint}}$ is

$$\text{rank}(\mathbf{M}_{\text{joint}}) = 6.$$

this can be derived by applying Gaussian elimination yielding

$$\mathbf{M'}_{\text{joint}} = \mathbf{D} \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \ell_{1,3}^{\Delta_{\text{glob}}}(b_2^{(3)}) & 0 & 0 & 0 & 0 \\ 0 & 0 & \ell_{1,3}^{\Delta_{\text{glob}}}(b_3^{(3)}) & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Each row with only a single one contributes with 1 to the rank whereas the two last rows

are linearly dependent which yields $\text{rank}(\mathbf{M}'_{\text{joint}}) = 6$. For the matrices $\mathbf{M}\mathcal{E}_2^{\text{sto}}$ and $\mathbf{M}_{\mathcal{E}_1}$, it holds that $\text{rank}(\mathbf{M}\mathcal{E}_2^{\text{sto}}) = 3$ and $\text{rank}(\mathbf{M}_{\mathcal{E}_1}) = 2$. Overall, we have

$$\lambda_{\text{MR,for}} \leq \text{rank}(\mathbf{M}_{\text{joint}}) - \text{rank}(\mathbf{M}_{\mathcal{E}_1}) - \text{rank}(\mathbf{M}_{\mathcal{E}_2^{\text{sto}}}) = 6 - 2 - 3 = 1.$$

This means that the secrecy capacity of the DSS using forwarded global repair is decreased by 1.

With this example in mind, a general expression for the secrecy capacity decrease $\lambda_{\text{MR,for}}$ can be derived.

Lemma 6.2 (**Forwarded global repair - secrecy capacity upper bound**). *Consider a DSS with g groups, parameters $h < r$ and an eavesdropper with $l_2 = 1$, l_1. Without loss of generality, it is assumed that the global erasures occur in one group, which is not the group observed by the eavesdropper in an l_2-manner. Let $h_{\min}$ be the minimal number of global repairs as defined in Definition 6.3. The secrecy capacity decrease in the presence of forwarded global repairs $\lambda_{\text{MR,for}}$ is upper bounded by*

$$\lambda_{\text{MR,for}} \leq \text{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) - l_2 r - l_1 =: \bar{\lambda}_{\text{MR,for}}.$$

For forwarded global repair, it holds that

$$\bar{\lambda}_{\text{MR,for}} = \min(h_{\min}, \sum_{j \in \mathcal{G}_{\text{up}}} (r - e_j)). \tag{6.9}$$

Proof. The upper bound $\bar{\lambda}_{\text{MR,for}}$ can be deduced from Equation (6.5). The proof focuses on showing that (6.9) holds.
Consider the matrix $\mathbf{M}_{\text{joint}}^{h_{\min}} \in \mathbb{F}_{q^m}^{(r+l_1+h_{\min})\times k}$. We now make the assumption that the nodes in $\mathcal{E}_1$ are, without loss of generality, nodes from Δ_{glob}, i.e., $|\mathcal{E}_1 \cap \Delta_{\text{glob}}| = l_1$. For the derivations, it does not make a difference where the l_1 observed nodes are and we could simply reorder them such that the condition is fulfilled and the parameters are still the same. By Gaussian elimination, the matrix has $l_2 r + l_1$ rows that only have a single one as an entry. These rows contribute $l_2 r + l_1$ to the rank similarly to Example 6.3. The rows with the global repair symbols, that are left after the elimination, consist of evaluations of polynomials at a P-independent set of points where the structure is the same as in Lemma 3.3. By this lemma, the submatrix has full rank with size $h_{\min} \times (\sum_{j \in \mathcal{G}_{\text{up}}} r - e_j)$, i.e., the rows are linearly independent. Thus, the overall rank of

$\mathbf{M}_{\text{joint}}^{h_{\min}}$ is $\text{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) = l_2 r + l_1 + \min(h_{\min}, \sum_{j\in\mathcal{G}_{\text{up}}}(r-e_j))$ and we have

$$\bar{\lambda}_{\text{MR,for}} = \text{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) - l_2 r - l_1 = \min(h_{\min}, \sum_{j\in\mathcal{G}_{\text{up}}}(r-e_j)).$$

□

Example 6.3 can now also be verified with Equation (6.9). For the considered DSS and eavesdropper parameters, we have

$$\bar{\lambda}_{\text{MR,for}} = \min(2, (3-2)) = \min(2,1) = 1$$

which is in accordance with the derivations.

At the end of this section, the secrecy capacity for DSSs is stated and it is shown that Construction 6.1 achieves the capacity with equality as seen before for the direct global repair (Theorem 6.1).

Theorem 6.3 (**Forwarded global repair - secrecy capacity**). *Consider a DSS with g groups, parameters $h < r$ and an eavesdropper with $l_2 = 1$, fixed l_1 such that $l_2 r + l_1 \leq k$. Without loss of generality, it is assumed that the global erasures occur in the same group, which, furthermore, is not observed by the eavesdropper in an l_2-manner. Let the l_1-observations of the eavesdropper be distributed in such a way that $|\mathcal{E}_1 \cap \Delta_{\text{glob}}| = l_1$. Let $h_{\min}$ be the minimal number of global repairs as defined in Definition 6.3. The secrecy capacity for forwarded global repair is*

$$k^{(\text{s})}{}_{\text{for}} = k - (l_2 r + l_1) - \min\left(h_{\min}, \sum_{j\in\mathcal{G}_{\text{up}}} r - e_j\right). \tag{6.10}$$

Proof. It follows from Lemma 6.2 that the right hand side of (6.10) is an upper bound. The proof is done by giving a construction that achieves equality in (6.10). □

Theorem 6.4. *Construction 6.1 is information-theoretically secure against an (l_1, l_2)-eavesdropper and achieves the secrecy capacity $k^{(\text{s})}{}_{\text{for}} = k - k^{(e)}$ with $k^{(e)} = (l_2 r + l_1 + \bar{\lambda}_{\text{MR,for}})$ from Theorem 6.3 if forwarded global repair is used.*

Proof. To prove secrecy of the coding scheme, the secrecy lemma (Lemma 3.1) is used. An (l_1, l_2)-eavesdropper can observe at most $l_2 r + l_1 + \bar{\lambda}_{\text{MR,for}}$ symbols. There are $l_2 r + l_1 + \bar{\lambda}_{\text{MR,for}}$ randomly generated symbols. Thus, the first condition is fulfilled, i.e., $\text{H}(\mathsf{E}) \leq \text{H}(\mathsf{R})$ where E and R are the random variables of the eavesdropper's observation and the randomly generated symbols, respectively. It remains to prove that $\text{H}(\mathsf{R} \mid \mathsf{U}^{(\text{s})}, \mathsf{E}) =$

0. Since the local encoding is MDS, the punctured local encoding matrix $\mathbf{A}|_{\Delta_{\text{glob}}} = \text{diag}(\mathbf{A}_1, \ldots, \mathbf{A}_g)|_{\Delta_{\text{glob}}}$ has rank k. The eavesdropped information can be summarized in the matrix $\mathbf{M}_{\text{joint}}^{h_{\min}}$. Thus, we have

$$\underbrace{\mathbf{A}|_{\Delta_{\text{glob}}}\mathbf{D}^{h_{\min}}\mathbf{M}_{\text{joint}}^{h_{\min}}}_{=:\mathbf{E}_{\text{joint}}}\mathbf{c}_{\Delta_{\text{glob}}} = \mathbf{e}_{\mathcal{E}}$$

where $\mathbf{c}_{\Delta_{\text{glob}}}$ represents the set of code symbols after the first encoding step with a skew Reed-Solomon code which yields a linearized Reed-Solomon code with the column multiplier matrix $\mathbf{D}^{h_{\min}}$. The eavesdropped symbols are denoted by $\mathbf{e}_{\mathcal{E}}$. Since $\mathbf{A}|_{\Delta_{\text{glob}}}$ and $\mathbf{D}^{h_{\min}}$ have rank k and are therefore full rank, the matrix $\mathbf{E}_{\text{joint}}$ has the same rank as $\mathbf{M}_{\text{joint}}^{h_{\min}}$ which is bounded by $\text{rank}(\mathbf{E}_{\text{joint}}) = \text{rank}(\mathbf{M}_{\text{joint}}^{h_{\min}}) = l_2 r + l_1 + \bar{\lambda}_{\text{MR,for}} \leq k$ by Lemma 6.2. The matrix $\mathbf{E}_{\text{joint}}$ can be transformed in the domain of the coefficients of the encoding skew polynomial by Definition 3.17 yielding

$$\mathbf{E}_{\text{joint}}(\mathbf{V}_k^{\sigma}(\mathbf{b}))^{-1}\mathbf{f} = \mathbf{e}_{\mathcal{E}}(\mathbf{V}_k^{\sigma}(\mathbf{b}))^{-1},$$

where $\mathbf{f}$ denotes the coefficients of the encoding skew polynomial F. The matrix $\mathbf{E}_{\text{joint}}(\mathbf{V}_k^{\sigma}(\mathbf{b}))^{-1}$ gives $l_2 r + l_1 + \bar{\lambda}_{\text{MR,for}}$ independent constraints on the polynomial coefficient. Together with the ${k^{(\text{s})}}_{\text{for}}$ coefficients of the information symbols we have k constraints on k coefficients. It remains to show that the k constraints are linearly independent. The ${k^{(\text{s})}}_{\text{for}}$ coefficients are known and can therefore be written as a row with one nonzero entry expanding the matrix $\mathbf{E}_{\text{joint}}(\mathbf{V}_k^{\sigma}(\mathbf{b}))^{-1}$. If the Vandermonde matrix $(\mathbf{V}_k^{\sigma}(\mathbf{b}))^{-1}$ is punctured at ${k^{(\text{s})}}_{\text{for}}$ columns, it still has rank $k - {k^{(\text{s})}}_{\text{for}}$ due to its structure. Multiplied with $\mathbf{E}_{\text{joint}}$ the overall rank of $\mathbf{E}_{\text{joint}}(\mathbf{V}_k^{\sigma}(\mathbf{b}))^{-1}$ is therefore still $k - {k^{(\text{s})}}_{\text{for}} = l_2 r + l_1 + \bar{\lambda}_{\text{MR,for}}$. Thus, the system of equations has enough linearly independent equations to determine the randomly generated symbols $k - {k^{(\text{s})}}_{\text{for}}$ and $\text{H}(\mathsf{R} \mid \mathsf{U}^{(\text{s})}, \mathsf{E}) = 0$ holds. This shows that the bound from 6.2 can be achieved with equality. □

6.4 Comparison of Direct and Forwarded Global Repair

In this short section, the secrecy capacities of DSSs with direct and forwarded global repair from Theorem 6.1 and Theorem 6.3 are briefly compared.

To have comparable parameters, it is assumed that $l_2 = 1$ and $h < r$ since then in both cases only one group is observed in an l_2-manner. The two capacities for the parameters are then

$$k^{(\mathrm{s})}{}_{\mathrm{dir}} = k - (r + l_1) - \lambda_{\mathrm{MR,dir}},$$
$$k^{(\mathrm{s})}{}_{\mathrm{for}} = k - (r + l_1) - \lambda_{\mathrm{MR,for}},$$

and the difference between them is

$$\lambda_{\mathrm{MR,dir}} - \lambda_{\mathrm{MR,for}} = \left(\left[\sum_{i=1}^{g} \min(h_{\mathrm{min,dir}}, r - e_i) \right] - h_{\mathrm{min,dir}} \right) - \min \left(h_{\mathrm{min,for}}, \sum_{j \in \mathcal{G}_{\mathrm{up}}} r - e_j \right).$$

It is important to point out that $h_{\mathrm{min,for}}$ and $h_{\mathrm{min,dir}}$ are not necessarily the same, since the matrices $\mathbf{M}_{\mathrm{joint}}$ for forwarded and direct global repair are not the same and therefore not necessarily have rank k for the same number of global repair given the same parameters. For a comparison, let us assume that the parameters of the DSS and the eavesdropper are chosen such that $h_{\mathrm{min,for}} = h_{\mathrm{min,dir}} = h$. The following example shows the differences for a set of parameters.

Example 6.4. Consider a DSS with $g = 4$, $r = 3$ and $h = 2$. An eavesdropper is observing only the second group in an l_2-manner, i.e., $l_2 = 1$ and $l_1 = 0$. The system, which is equivalent for forwarded and direct global repair, is illustrated in Figure 6.4.

The secrecy capacity decrease in such a system would be

$$\lambda_{\mathrm{MR,for}} = \min \left(h, \sum_{j \in \mathcal{G}_{\mathrm{up}}} r - e_j \right) = \min(2, 3) = 2$$

for forwarded global repair, and

$$\lambda_{\mathrm{MR,dir}} = \left(\left[\sum_{i=1}^{g} \min(h, r - e_i) \right] - h \right) = \min(2,3) + \min(2,0) + \min(2,3) + \min(2,3) - 2 = 4$$

for direct global repair. Thus, an eavesdropper would get twice as many independent symbols for the considered system when it is using direct global repair instead of forwarded global repair.

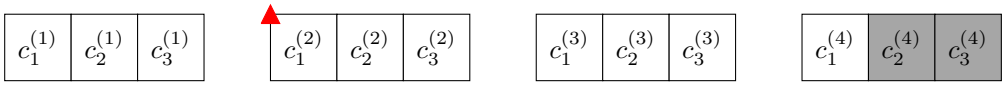

Figure 6.4: Illustration of a DSS with $g = 4$ groups and eavesdropper parameters $l_1 = 0$, $l_2 = 1$. The system has a different secrecy capacity for direct global repair and forwarded global repair.

Intuitively, we see that for $l_2 = 1$, $l_1 = 0$ and $g \geq 3$, it holds that $\lambda_{\mathrm{MR,for}} < \lambda_{\mathrm{MR,dir}}$.

While the eavesdropper gets at most h linearly independent symbols during the global repairs in a forwarded repair scheme, no matter how many groups there are, this looks different for the direct repair, where the number of symbols that the eavesdropper gets, during the global repairs, increases with the number of groups.

It can be concluded that the forwarded global repair has a higher secrecy capacity than the direct global repair for the considered constraints. However, there are also other factors to be compared when it comes to implementing global repair on a system. The forwarded global repair might have a larger latency than the direct global repair since the groups, except for the first group, have to wait to send their contribution to the next group until receiving the contribution from the previous group to the global repair. With a large number of groups and the inter-rack communication being the bottleneck, this might be a serious issue.

7 Summary and Outlook

7.1 Summary

The secrecy capacity of distributed storage systems (DSSs) with locally repairable codes (LRCs) having maximal recoverability was investigated. We considered the threat of an (l_1, l_2)-eavesdropper, which belongs to the class of passive attacks. The eavesdropper can read some l_1 nodes, and read all the nodes of additional l_2 groups while also being able to read the downloaded symbols for repair of these groups. Given an (l_1, l_2)-eavesdropper and optimal LRCs, a secrecy capacity, i.e., the file size which can be stored secretly on the system without revealing any information to the eavesdropper, was derived in [RKSV14]. The thesis links to this work with the goal to derive a secret file size for MR-LRCs, which are constructed as suggested in [MPK19].

Crucial for the construction of MR-LRCs are skew polynomials. Besides the recapitulation of important properties of skew polynomials in the Appendix A and codes that can be constructed using skew polynomials, two known concepts from conventional polynomials are adapted for skew polynomials. The first concept is secret sharing as described in Section 3.1. The secret sharing scheme with skew polynomials, described in Section 3.6, has no advantages compared to secret sharing with conventional polynomials. In fact, for a special parameter choice it recovers the conventional secret sharing. However, it is useful to prove the secrecy of LRC constructions that involve secret sharing such as Construction 4.2. This construction is a modification of Construction 4.1 presented in [RKSV14] but requiring smaller field sizes. However, it still does not take global repairs into account.

When taking them into account, the first insight was that a nonzero capacity is only possible if the system has a hierarchical structure. Maximal recoverability means that besides the local erasure correction capability, that is determined by the local code, the system has global parities to repair additional erasures that cannot be handled by the local code. A hierarchical structure allows to distribute such a global repair over the groups of the system. If this were not possible, the single node, which is repaired globally, would acquire global knowledge to repair itself. In the presence of an eavesdropper with $l_2 \geq 1$, the secrecy capacity would therefore be zero. To resolve this problem, local polynomials

were introduced. Their sum returns the global encoding polynomial, i.e., for the encoding polynomial P, it holds that

$$P = \sum_{i=1}^{g} L_i$$

for a DSS with g groups and local polynomials L_i. The idea of the local polynomial L_i is that it vanishes on the nodes that contribute to the global repair, i.e., the nodes that are in the global repair set Δ_{glob}, but that are not in the i-th group. Knowing the global repair set Δ_{glob}, the local polynomials can be generated by Newton interpolation.

The easiest way to perform a global repair is the direct global repair. Each group sends its contribution, i.e., the evaluation of the local polynomial at the code locator that needs repair, directly to the group with the erasure. If an eavesdropper observes the group where the global repair is performed in an l_2-manner, the secrecy capacity is further decreased. The information revealed to the eavesdropper, besides the static (l_1, l_2) eavesdropped nodes, can be upper-bounded by representing all the eavesdropped symbols in a matrix with respect to a basis representing the information vector (see Lemma 3.7 and Example 3.8). It turns out that choosing the global repair set Δ_{glob} as a basis is beneficial for the derivations of the secrecy capacity. As a result, the secrecy capacity for an MR-LRC is

$$k^{(\text{s})}{}_{\text{dir}} = k - (l_2 r + l_1) - \left[\left(\sum_{i=1}^{g} \min(h_{\min}, r - e_i) \right) - h_{\min} \right],$$

given the assumption that $h < r$. Otherwise for $h \geq r$, the secrecy capacity is zero.

A construction which shows the tightness of the upper bound is given. It essentially relies on the same secret-sharing-based approach given in [RKSV14], and seems to be a universal approach to the given problem. Here again, representing the eavesdropped information in a matrix is very useful to prove secrecy using the secrecy lemma (Lemma 3.1).

Another way to perform the global repair, forwarded global repair, is also analyzed but only for a special choice of parameters. For $h < r$ and $l_2 = 1$, the secrecy capacity, which can be derived with the same ideas as for the direct global repair, is

$$k^{(\text{s})}{}_{\text{for}} = k - (l_2 r + l_1) - \min \left(h_{\min}, \sum_{j \in \mathcal{G}_{\text{up}}} r - e_j \right).$$

Therefore, it seems to have a better secrecy compared to direct global repair. This coincides with the intuition since the eavesdropper can only observe one symbol for each global repair instead of $g - 1$ symbols.

7.2 Outlook

It would be very interesting to investigate the secrecy capacity of an MR-LRC with forwarded global repair further. However, the analysis of the secrecy capacity for more general parameters seems to be tricky. For $h \geq r$, the forwarding list changes for global repairs from different groups. For $l_2 \geq 2$, the knowledge gain is also a function of their position in the forwarding list. Nevertheless, the method of representing the symbols known to the eavesdropper in a matrix might still work for more general parameters.

The idea of using linearized Reed-Solomon codes or more generally skew polynomials for designing MR-LRCs is rather new. Therefore, a lot of interesting questions arise. One aspect that has not yet been investigated, to the best of our knowledge, is the performance of MR-LRCs using skew polynomials compared to already implemented LRCs. To compare the performance, the skew arithmetic would have to be efficiently implemented in software, such as ceph [WBM$^+$06], that is used to manage DSSs. However, the downside of applying MR-LRCs is that the overhead of updating the global parities with each adjustment of the stored data might be too large compared to the gain of the possibility to perform a global repair in the unlikely event of too many failures.

Appendix A

A.1 Skew Polynomials

The ring of skew polynomials is a generalization of the ring of conventional polynomials. The addition of two skew polynomials and conventional polynomials, for example, is the same. The main difference is that the product of two skew polynomials is not commutative.

Skew polynomials were first studied by Ore in [Ore33]. In the following, they are examined and their properties are summarized. The introduction and notation in this section follows [MPSK22].

Let

$$\mathcal{P} = \left\{ \sum_{i=0}^{d} F_i x^i \,\middle|\, d \in \mathbb{N}_0, F_i \in \mathbb{F}_{q^m}, \text{ for all } i \in \{0, 1, \ldots, d\} \right\}$$

be the set of polynomials in the variable x with coefficients in $\mathbb{F}_{q^m}$.
In $\mathcal{P}$ addition of two polynomials and the multiplication of a scalar on the left with a polynomial is defined as follows. For two polynomials

$$\begin{aligned} F &= F_0 x^0 + F_1 x^1 + \cdots + F_d x^d \in \mathcal{P} \\ G &= G_0 x^0 + G_1 x^1 + \cdots + G_d x^d \in \mathcal{P} \end{aligned}$$

with $d \in \mathbb{N}$, $F_i, G_i \in \mathbb{F}_{q^m}$ for $i \in \{0\} \cup [d]$ and a scalar $a \in \mathbb{F}_{q^m}$, $F + G$ and aF are defined as follows

$$\begin{aligned} F + G &= (F_0 + G_0)x^0 + (F_1 + G_1)x^1 + \cdots + (F_d + G_d)x^d \in \mathcal{P}, \\ aF &= (aF_0)x^0 + (aF_1)x^1 + \cdots + (aF_d)x^d \in \mathcal{P}. \end{aligned} \tag{A.1}$$

Thus, skew polynomials and conventional polynomials behave identically with respect to addition and multiplication with a scalar on the left.

Define the degree of a polynomial $F = \sum_{i=0}^{d} F_i x^i \in \mathcal{P}$ with $F_d \neq 0$ as $\deg(P) = d$ and $\deg(F) = -\infty$ for $F = 0$. Since the addition of two polynomials behaves as with

conventional polynomials it holds that

$$\deg(F+G) \leq \max\{\deg(F), \deg(G)\} \tag{A.2}$$

for $F, G \in \mathcal{P}$. If $\deg(F) \neq \deg(G)$, then $\deg(F+G) = \max\{\deg(F), \deg(G)\}$. In [Ore33], it is shown that a product in $\mathcal{P}$ turns $\mathcal{P}$ into a ring with

$$\deg(FG) = \deg(F) + \deg(G) \tag{A.3}$$

for all $P, Q \in \mathcal{P}$ and multiplicative identity $1 = x^0$ where $x^{i+j} = x^i x^j, \forall i, j \in \mathbb{N}_0$, if and only if,

$$xa = \sigma(a)x + \delta(a) \tag{A.4}$$

for all $a \in \mathbb{F}_{q^m}$ with the automorphism σ: $\mathbb{F}_{q^m} \longrightarrow \mathbb{F}_{q^m}$ (see Definition B.2) and the σ-derivation δ: $\mathbb{F}_{q^m} \longrightarrow \mathbb{F}_{q^m}$ (see Definition B.3). Without loss of generality [MP17], the σ-derivation δ can be set to zero: $\delta(a) = 0$ for all $a \in \mathbb{F}_{q^m}$. The field automorphism σ is in the following chosen as

$$\sigma(a) = a^q, \quad \forall\, a \in \mathbb{F}_{q^m}.$$

As a result, the i-th composition of σ is

$$\begin{aligned} \sigma^i(a) &= \underbrace{\sigma(\sigma(\cdots\sigma(a)))}_{\text{i times}} = a^{q^i}, \\ x^i a &= \sigma^i(a)x = a^{q^i}x. \end{aligned} \tag{A.5}$$

Definition A.1 (**Skew polynomials** [MPSK22, Def. 2.1])**.** The set $\mathcal{P}$ with multiplication as in (A.4) (here $\delta(a) = 0$) and with addition as in (A.1) is called the ring of skew polynomials denoted by $\mathbb{F}_{q^m}[x;\sigma]$. Its elements are called skew polynomials.

For σ being the identity automorphism, i.e., $\sigma(a) = a$, skew polynomials coincide with regular, conventional polynomials over finite fields namely $\mathbb{F}_{q^m}[x]$. This is for instance the case if the extension degree of the field is $m = 1$.

Example A.1. Consider the field $\mathbb{F}_{q^m} = \mathbb{F}_4$ with $q = 2$ and $m = 2$. Therefore, $\mathbb{F}_4[x;\sigma]$ has the automorphism $\sigma(a) = a^2$ for all $a \in \{0, 1, \omega, \bar{\omega}\}$ as defined in (1.1) with

$$\begin{aligned} \sigma(0) &= 0, \\ \sigma(1) &= 1^2 = 1, \\ \sigma(\omega) &= \omega^2 = \bar{\omega}, \\ \sigma(\bar{\omega}) &= \bar{\omega}^2 = \omega. \end{aligned}$$

To illustrate the polynomial multiplication, consider $F = \omega + \bar{\omega}x + x^2 \in \mathbb{F}_4[x;\sigma]$ and $G = \bar{\omega} + \omega x \in \mathbb{F}_4[x;\sigma]$ then

$$\begin{aligned} FG &= (\omega + \bar{\omega}x + x^2)(\bar{\omega} + \omega x) \\ &= (\omega + \bar{\omega}x + x^2)\bar{\omega} + (\omega + \bar{\omega}x + x^2)\omega x \\ &= \omega\bar{\omega} + \bar{\omega}x\bar{\omega} + x^2\bar{\omega} + \omega\omega x + \bar{\omega}x\omega x + x^2\omega x \\ &= 1 + x + \bar{\omega}x^2 + \bar{\omega}x + \omega x^2 + \omega x^3 \\ &= 1 + (1 + \bar{\omega})x + (\bar{\omega} + \omega)x^2 + \omega x^3 \\ &= 1 + \omega x + x^2 + \omega x^3, \end{aligned}$$

whereas

$$\begin{aligned} GF &= (\bar{\omega} + \omega x)(\omega + \bar{\omega}x + x^2) \\ &= \bar{\omega}(\omega + \bar{\omega}x + x^2) + \omega x(\omega + \bar{\omega}x + x^2) \\ &= \bar{\omega}\omega + \bar{\omega}\bar{\omega}x + \bar{\omega}x^2 + \omega x\omega + \omega x\bar{\omega}x + \omega x x^2 \\ &= 1 + \omega x + \bar{\omega}x^2 + x + \bar{\omega}x^2 + \omega x^3 \\ &= 1 + (\omega + 1)x + (\bar{\omega} + \bar{\omega})x^2 + \omega x^3 \\ &= 1 + \bar{\omega}x + \omega x^3. \end{aligned}$$

This shows that multiplication in $\mathbb{F}_4[x;\sigma]$ is in general not commutative.

For the division of skew polynomials, the non-commutativity implies that there is a difference between division on the left and on the right. In this work, only division on the right is considered. For the division, the following result was given in [Ore33].

Theorem A.1 (**Euclidean division** [Ore33, Sec. 2])**.** *Given $F, G \in \mathbb{F}_{q^m}[x;\sigma]$ with $G \neq 0$, there exists a unique decomposition such that*

$$F = QG + R$$

with $Q, R \in \mathbb{F}_{q^m}[x;\sigma]$ and $\deg(R) < \deg(G)$. Q is called the quotient and R is called the remainder.

Proof. The first part is to prove the existence of the polynomials Q and R.
Case 1: $F = 0$, then the proposition is true for $Q = R = 0$.
Case 2: Suppose $\deg(F) < \deg(G)$, then the proposition is true for $Q = 0$ and $R = F$.
Case 3: If $\deg(F) \geq \deg(G)$, the existence of the decomposition can be shown by induction on the degree of F.

I. If $\deg(F) = 0$, then also $\deg(G) = 0$. Hence $F = a \in \mathbb{F}_{q^m}$ and $G = b \in \mathbb{F}_{q^m}$. The theorem is true for $Q = ab^{-1}$ and $R = 0$.

II. Now it is assumed that the proposition is true for $\deg(F) < n$ and it must be shown that the proposition then also holds for $\deg(F) = n$. Let $F = F_n x^n + \cdots + F_1 x + F_0 \in \mathbb{F}_{q^m}[x;\sigma]$ and $G = G_m x^m + \cdots + G_1 x + G_0 \in \mathbb{F}_{q^m}[x;\sigma]$ be the considered polynomials with $F_n \neq 0$, $G_m \neq 0$ and $m \leq n$. Multiply the divisor G by $F_n \sigma^{n-m}(G_m)^{-1} x^{n-m}$ to obtain

$$\begin{aligned} &F_n \sigma^{n-m}(G_m)^{-1} x^{n-m} G = F_n \sigma^{n-m}(G_m)^{-1} x^{n-m}(G_m x^m + \cdots + G_1 x + G_0) \\ &= F_n x^n + F_n \sigma^{n-m}(G_m)^{-1} x^{n-m} G_{m-1} x^{m-1} + \cdots + F_n \sigma^{n-m}(G_m)^{-1} x^{n-m} G_0 \\ &= F_n x^n + F_n \sigma^{n-m}(G_m)^{-1} \sigma^{n-m}(G_{m-1}) x^{n-1} + \cdots + F_n \sigma^{n-m}(G_m)^{-1} \sigma^{n-m}(G_0) x^{n-m}. \end{aligned}$$

The leading term of this polynomial is identical to the one of F. Therefore,

$$F - F_n \sigma^{n-m}(G_m)^{-1} x^{n-m} G$$

is a polynomial of degree less than n.
Now either Case 1 applies for $F - F_n \sigma^{n-m}(G_m)^{-1} x^{n-m} G = 0$ or given the hypothesis, there exist polynomials Q' and R such that

$$F - F_n \sigma^{n-m}(G_m)^{-1} x^{n-m} G = Q'G + R$$

with $\deg(R) < \deg(G)$. F can then be written as

$$F = (F_n \sigma^{n-m}(G_m)^{-1} x^{n-m} + Q')G + R$$

with $\deg(R) < \deg(G)$. Thus, the proposition holds with $Q = F_n \sigma^{n-m}(G_m)^{-1} x^{n-m} + Q'$ when $\deg(F) = n$, which completes the induction and shows that Q and R exist for any divided F and any divisor G.

The second part is to prove that Q and R are unique. Suppose that there are two decompositions with

$$F = QG + R$$

and

$$F = \tilde{Q}G + \tilde{R}$$

with $\deg(R) < \deg(G)$ and $\deg(\tilde{R}) < \deg(G)$. Then it would hold that

$$QG + R = \tilde{Q}G + \tilde{R}$$

or

$$(Q - \tilde{Q})G = \tilde{R} - R. \tag{A.6}$$

If $Q - \tilde{Q} \neq 0$, then $\deg((Q - \tilde{Q})G) \geq \deg(G)$. But given the assumption that the degrees of R and $\tilde{R}$ are strictly less than the degree of G, (A.6) can only hold for $Q - \tilde{Q} = 0$, which is a contradiction. Therefore, $Q - \tilde{Q} = 0$ and as a result also $\tilde{R} - R = 0$ holds which means that $Q = \tilde{Q}$ and $R = \tilde{R}$. Thus, the polynomials are unique. □

The induction step in the proof uses the Euclidean algorithm for long division which is illustrated in the following example.

Example A.2. Consider the skew polynomials from Example A.1, $F = x^2 + \bar{\omega}x + \omega \in \mathbb{F}_4[x;\sigma]$ and $G = \omega x + \bar{\omega} \in \mathbb{F}_4[x;\sigma]$. A long division is performed to calculate quotient and remainder:

The first step is to calculate $F_2\sigma(G_1)^{-1}x = 1 \cdot \sigma(\omega)^{-1}x = \omega x$ which is the first and, in this case also, only summand of Q. The calculations are shown in the following division table.

$$\begin{array}{rl} (x^2 + \bar{\omega}x \ + \omega) \div (\omega x + \bar{\omega}) = \omega x + \dfrac{\omega}{\omega x + \bar{\omega}} \\ \underline{-(x^2 + \bar{\omega}x)} \\ 0 \ + \omega \\ \underline{-\quad 0} \\ \omega \end{array}$$

Thus, we can write

$$F = QG + R = (\omega x)(\omega x + \bar{\omega}) + \omega$$

which can be verified by the reader.

Another operation that is different compared to conventional polynomials is the polynomial evaluation. It is not defined by simply "plugging" a value in, i.e., $F(a) = \sum_{i=0}^{n} F_i a^i$ with $a \in \mathbb{F}_{q^m}$, but defined by forcing a remainder theorem as suggested by Lam and Leroy in [LL88].

Theorem A.2 (Remainder theorem). *Given a skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$ and an element $a \in \mathbb{F}_{q^m}$, it holds that $F(a)$ is the only element in $\mathbb{F}_{q^m}$ such that there exists a skew polynomial $Q \in \mathbb{F}_{q^m}[x;\sigma]$ fulfilling*

$$F = Q \cdot (x - a) + F(a). \tag{A.7}$$

For $F(a) = 0$, F is said to be a left-multiple of $x - a$.

From the Euclidean division (Theorem A.1) it follows directly that there exists a unique $F(a)$ which is an element in $\mathbb{F}_{q^m}$ since it has degree 0 ($\deg(R) < \deg(G)$). This fact gives the motivation to define the evaluation of a skew polynomial as follows.

Definition A.2 (**Evaluation** [MPSK22, Def. 2.2])**.** Given a skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$ and an element $a \in \mathbb{F}_{q^m}$, the evaluation of F at a is defined as the remainder of the Euclidean division of F on the right by the skew polynomial $x-a \in \mathbb{F}_{q^m}[x;\sigma]$. The evaluation is denoted by $F(a) \in \mathbb{F}_{q^m}$.

There is another way to calculate the evaluation with an explicit formula which was given in [LL88, Lem. 2.4].

Theorem A.3. *Given an element $a \in \mathbb{F}_{q^m}$, define $N_0(a) = 1$ and*

$$N_i(a) = \sigma^{i-1}(a)\sigma^{i-2}(a)\cdots\sigma(a)a \tag{A.8}$$

or recursively

$$N_{i+1} = \sigma^i(a) \cdot N_i(a) \tag{A.9}$$

for all $i \in \mathbb{N}$. The evaluation of a skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$ of degree d can be calculated with

$$F(a) = F_0N_0(a) + F_1N_1(a) + \cdots + F_dN_d(a). \tag{A.10}$$

Proof. The theorem will be proven by strong induction on $d \geq 0$ and follows the proof in [MPSK22, Prop. 2.1]. For $d = 0$, $F = F_0$ it holds trivially. Assume that (A.10) holds for $\deg(F) < d$ and take $F \in \mathbb{F}_{q^m}[x;\sigma]$ with $\deg(F) = d$. The skew polynomial F can be written as

$$\begin{aligned} F =& F_0 + F_1x + \cdots + F_dx^d \\ =& F_0 + (F_1 + F_2x + \cdots + F_dx^{d-1})(x-a) + \\ & (F_1a + F_2xa + \cdots + F_dx^{d-1}a) \\ =& (F_1 + F_2x + \cdots + F_dx^{d-1})(x-a) + \\ & \underbrace{(F_0 + F_1a + F_2\sigma(a)x + \cdots + F_d\sigma^{d-1}(a)x^{d-1})}_{\tilde{F}}. \end{aligned}$$

In the last equality (A.5) is used. The skew polynomial $\tilde{F}$ is of degree less than d. By the strong induction hypothesis,

$$\begin{aligned} \tilde{F}(a) &= (F_0 + F_1a)N_0(a) + F_2\sigma(a)N_1(a) + \cdots + F_d\sigma^{d-1}(a)N_{d-1}(a) \\ &= F_0N_0(a) + F_1N_1(a) + F_2N_2(a) + \cdots + F_dN_d(a). \end{aligned}$$

With (A.7), $\tilde{F}$ can be written as

$$\begin{aligned} \tilde{F} &= \tilde{Q}(x-a) + \tilde{F}(a) \\ &= \tilde{Q} \cdot (x-a) + F_0N_0(a) + F_1N_1(a) + F_2N_2(a) + \cdots + F_dN_d(a) \end{aligned}$$

with $\tilde{Q} \in \mathbb{F}_{q^m}[x;\sigma]$. Defining

$$Q = \tilde{Q} + F_1 + F_2 x + \cdots + F_d x^{d-1} \in \mathbb{F}_{q^m}[x;\sigma],$$

we can write F as

$$F = Q \cdot (x - a) + F_0 N_0(a) + F_1 N_1(a) + \cdots + F_d N_d(a).$$

By Theorem A.2, the induction step holds which proves (A.10). □

An example is given to illustrate the two ways in which the evaluation of a skew polynomial can be calculated.

Example A.3. Consider again $\mathbb{F}_4$ with $q = 2$ and $m = 2$ as defined in (1.1) and the skew polynomial $F = \omega x^2 + x + \bar{\omega} \in \mathbb{F}_4[x;\sigma]$. For the evaluation of F at $a = \omega$, F can be written as

$$F = \omega x(x - \omega) + \bar{\omega}$$

and therefore

$$F(\omega) = \bar{\omega}.$$

The evaluation can also be calculated as follows

$$\begin{aligned} F(\omega) &= \omega N_2(\omega) + 1 \cdot N_1(\omega) + \bar{\omega} N_0(\omega) \quad = \omega\bar{\omega}\omega + \omega + \bar{\omega} \\ &= \omega + \omega + \bar{\omega} \\ &= \bar{\omega}, \end{aligned}$$

which is the same result.

It is important to note that the evaluation of skew polynomials is not $\mathbb{F}_q$-linear, i.e., for $F \in \mathbb{F}_{q^m}[x;\sigma]$, $\lambda_1, \lambda_2 \in \mathbb{F}_q$ and $a, b \in \mathbb{F}_{q^m}$,

$$F(\lambda_1 a + \lambda_2 b) \neq F(\lambda_1 a) + F(\lambda_2 b)$$

in general.

Example A.4. Consider the same skew polynomial $F = \omega x^2 + x + \bar{\omega} \in \mathbb{F}_4[x;\sigma]$ as in Example A.3. For this polynomial it holds that

$$\begin{aligned} F(\omega) &= \bar{\omega}, \\ F(\bar{\omega}) &= \omega, \\ F(1) &= 0. \end{aligned}$$

Obviously, $F(\bar{\omega}) = F(1+\omega) \neq F(1) + F(\omega)$ which shows that the evaluation of skew polynomials is not $\mathbb{F}_q$-linear.

A useful property of skew polynomials is the linearity of the evaluation of two polynomials.

Theorem A.4 (**Linearity rule** [MPSK22, Prop.2.3])**.** *Given two skew polynomials $F, G \in \mathbb{F}_{q^m}[x;\sigma]$, for all scalars $a, \lambda, \mu \in \mathbb{F}_{q^m}$ it holds that*

$$(\lambda F + \mu G)(a) = \lambda F(a) + \mu G(a).$$

Proof. Follows from the distributive properties of skew polynomials and is left to the reader. □

This means that the evaluation of a sum of polynomials is the sum of their evaluations. For the evaluation of a product $(FG)(a)$, a similar concept is not applicable, i.e., $(FG)(a) \neq F(a) \cdot G(a)$. However, there is a rule for the evaluation of the product of two skew polynomials. It is helpful to first introduce the notion of conjugacy which is later also needed to explain concepts about the roots of skew polynomials. Conjugacy for skew polynomials was first introduced in [LL88], [Lam85]. The notation follows the notation in [GG22] and [MPSK22].

Definition A.3 (**Conjugacy**)**.** Let $a \in \mathbb{F}_{q^m}$ and $c \in \mathbb{F}_{q^m}^*$. The c-conjugate of a with respect to the field automorphism σ is defined as

$$^c a = \sigma(c) a c^{-1}.$$

The element b is said to be a conjugate of a, with respect to the field automorphism σ, written $b \sim_\sigma a$, if there exists some $c \in \mathbb{F}_{q^m}^*$ such that

$$b = {}^c a.$$

For the automorphism $\sigma(a) = a^q$, it follows that

$$^c a = \sigma(c) a c^{-1} = c^{q-1} a.$$

It also follows that for $c \in \mathbb{F}_q^*$,

$$^c a = a.$$

Also note that $^c ac = \sigma(c)a$.

Lemma A.1. *It holds that $^y(^x a) = {}^{yx} a$.*

Proof. Consider ${}^{y}({}^{x}a)$ with

$$\begin{aligned} {}^{y}({}^{x}a) &\overset{(a)}{=} \sigma(y) \cdot {}^{x}a \cdot y^{-1} \\ &\overset{(b)}{=} \sigma(y) \left(\sigma(x) a x^{-1}\right) y^{-1} \\ &\overset{(c)}{=} \sigma(yx) a (yx)^{-1} \\ &= {}^{yx}a, \end{aligned}$$

where (a) and (b) follow from Definition A.3 and (c) holds by 3. of Definition B.2. □

Proposition A.1. *For $a, b \in \mathbb{F}_{q^m}$, $\sim_\sigma$ is an equivalence relation in $\mathbb{F}_{q^m}$.*

Proof. For $a, b, c \in \mathbb{F}_{q^m}$ three properties need to be shown:

1. Reflexivity: $a \sim_\sigma a$.
2. Symmetry: If $a \sim_\sigma b$, then $b \sim_\sigma a$.
3. Transitivity: If $a \sim_\sigma b$ and $a \sim_\sigma c$, then $a \sim_\sigma c$.

Reflexivity: For $x = 1 \in \mathbb{F}_{q^m}^*$, $a \sim_\sigma a$ since

$${}^{1}a = \sigma(1) a 1^{-1} = a.$$

Symmetry: Suppose $a \sim_\sigma b$ with $x \in \mathbb{F}_{q^m}^*$ such that $a = {}^{x}b$. Then it holds that

$${}^{x^{-1}}a = {}^{x^{-1}}({}^{x}b) = {}^{x^{-1}x}b = b,$$

where Lemma A.1 is applied. Thus, $a \sim_\sigma b$ since $b = {}^{y}a$ for $y = x^{-1} \in \mathbb{F}_{q^m}^*$.
Transitivity: Suppose $a \sim_\sigma b$ with $x \in \mathbb{F}_{q^m}^*$ such that $a = {}^{x}b$ and $b \sim_\sigma c$ with $y \in \mathbb{F}_{q^m}^*$ such that $b = {}^{y}c$, then

$$a = {}^{x}b = {}^{x}({}^{y}c) = {}^{xy}c,$$

where Lemma A.1 is applied. So $a \sim_\sigma c$ since $a = {}^{z}c$ for $z = xy \in \mathbb{F}_{q^m}^*$. □

Therefore, $\mathbb{F}_{q^m}$ is partitioned into distinct classes of the equivalence relation $\sim_\sigma$. These classes are called conjugacy classes.

Definition A.4 (**Conjugacy classes** [MPSK22, Def. 2.8])**.** For $a \in \mathbb{F}_{q^m}$, the conjugacy class with respect to σ is defined as

$$C_\sigma(a) = \{b \in \mathbb{F}_{q^m} \mid b \sim_\sigma a\} = \{{}^{\beta}a \in \mathbb{F}_{q^m} \mid \beta \in \mathbb{F}_{q^m}^*\}.$$

With this notion of conjugacy in mind, a product rule for skew polynomials can be given. It was introduced by Lam and Leroy in [LL88].

Theorem A.5 (**Product rule** [LL88, Th. 2.7])**.** *Consider $F, G \in \mathbb{F}_{q^m}[x;\sigma]$ and an element $a \in \mathbb{F}_{q^m}$. Denote $\beta = G(a)$. If $\beta = 0$, then*

$$(FG)(a) = 0.$$

If $\beta \neq 0$, then

$$(FG)(a) = F\left(\sigma(\beta)a\beta^{-1}\right) \cdot \beta = F\left({}^{\beta}a\right)\beta = F\left({}^{G(a)}a\right)G(a). \tag{A.11}$$

Proof. By Theorem A.2, the skew polynomial G can be written as

$$G = Q \cdot (x - a) + G(a) \tag{A.12}$$

with $Q \in \mathbb{F}_{q^m}[x;\sigma]$. If $G(a) = 0$, then $G = Q \cdot (x - a)$ yielding

$$FG = F \cdot (Q \cdot (x - a)) = (FQ) \cdot (x - a),$$

and $(FG)(a) = 0$, which follows directly from Theorem A.2. Let $b = {}^{\beta}a = \sigma(\beta)a\beta^{-1}$ such that

$$(x - b)\beta = \sigma(\beta)(x - a)$$

holds. The skew polynomial F can by Theorem A.2 be written as

$$F = P \cdot (x - b) + F(b) \tag{A.13}$$

with $P \in \mathbb{F}_{q^m}[x;\sigma]$. Thus, with (A.12) and (A.13), FG can be written as follows

$$\begin{aligned} FG &= F \cdot (Q \cdot (x - a) + G(a)) \\ &= (FQ) \cdot (x - a) + F \cdot G(a) \\ &= (FQ) \cdot (x - a) + (P \cdot (x - b) + F(b))G(a) \\ &= (FQ) \cdot (x - a) + P \cdot (x - b)\beta + F(b)G(a) \\ &= (FQ) \cdot (x - a) + P \cdot \sigma(\beta)(x - a) + F(b)G(a) \\ &= (FQ + P \cdot \sigma(\beta))(x - a) + F(b)G(a), \end{aligned}$$

and (A.11) follows by Theorem A.2. □

This means that the evaluation of a product of skew polynomials is the product of

the evaluation of the left skew polynomial F at ${}^{G(a)}a$ multiplied with the evaluation of the skew polynomial on the right G at a.

A crucial property that is linked to evaluation is Lagrange interpolation. Interpolation is used to decode Reed-Solomon codes. Since the goal is to define codes that have similar properties to Reed-Solomon codes but defined over $\mathbb{F}_{q^m}[x;\sigma]$, it is important to understand the differences between conventional polynomials and skew polynomials regarding interpolation.

Definition A.5 (**Lagrange interpolation** [MPSK22, Def. 2.3])**.** It is possible to perform Lagrange interpolation in $\mathbb{F}_{q^m}[x;\sigma]$ on the evaluation points $a_1,\dots,a_n \in \mathbb{F}_{q^m}$ if it holds that:
For all evaluation values $b_1,\dots,b_n \in \mathbb{F}_{q^m}$, there exists a skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$ with $F(a_i) = b_i$, for $i \in [n]$.

For regular Reed-Solomon codes it is sufficient to have a set of *distinct* evaluation points $a_1,\dots,a_n$ with their corresponding evaluation values $b_1,\dots,b_n$ to recover the corresponding polynomial of degree $n-1$. However, for skew polynomials this is more complex, since the notion of distinctness is not enough as the following example, that was given in [MPSK22, Ex. 2.8], shows.

Example A.5. Consider $\mathbb{F}_4$ with $q = 2$ and $m = 2$ as defined in (1.1). Take $a_1 = 1$, $a_2 = \omega$ and $a_3 = \bar{\omega}$, which are distinct elements of $\mathbb{F}_4$. With (A.8) and (A.9), it can be shown that

$$N_i(\bar{\omega})\bar{\omega} = N_i(\omega)\omega + N_i(1)$$

for all $i \in \mathbb{N}$. Therefore, for any skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$,

$$F(\bar{\omega})\bar{\omega} = F(\omega)\omega + F(1)$$

which can be deduced from (A.10). If we consider $F(a_1) = F(1) = b_1$ and $F(a_2) = F(\omega) = b_2$, then $F(a_3) = F(\bar{\omega})$ follows directly from $F(\bar{\omega})\bar{\omega} = b_2\omega + b_1$,

$$F(a_3) = F(\bar{\omega}) = b_2\bar{\omega} + b_1\omega$$

which means that $F(a_3)$ is determined by b_1 and b_2 and can therefore not be chosen arbitrarily. This means that Lagrange interpolation cannot be performed on the points a_1, a_2 and a_3 even though they are distinct elements in $\mathbb{F}_4$.

For a well defined Lagrange interpolation for skew polynomials, it is crucial to understand how the number of zeros or roots is bounded by the degree of a skew polynomial. For this reason, the following part deals with zeros of skew polynomials.

Definition A.6 ([MPSK22, Def. 2.4])**.** For a skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$, the set of roots or zeros is defined as

$$Z(F) = \{a \in \mathbb{F}_{q^m} \mid F(a) = 0\} \subseteq \mathbb{F}_{q^m}.$$

Given a set $\Omega \subseteq \mathbb{F}_{q^m}$, the set of skew polynomials vanishing on Ω is defined as

$$I(\Omega) = \{F \in \mathbb{F}_{q^m}[x;\sigma] \mid F(a) = 0 \; \forall a \in \Omega\} \subseteq \mathbb{F}_{q^m}[x;\sigma].$$

Note that $I(\Omega)$ is a left ideal of the ring $\mathbb{F}_{q^m}[x;\sigma]$, which means that

1. $I(\Omega)$ is a subgroup of $\mathbb{F}_{q^m}[x;\sigma]$,
2. $I(\Omega)$ is closed under inside-outside multiplication from the left,

which mean that

$$\begin{aligned} F + G &\in I(\Omega) \; \forall F, G \in I(\Omega) \\ FG &\in I(\Omega) \; \forall F \in \mathbb{F}_{q^m}[x;\sigma] \text{ and } G \in I(\Omega). \end{aligned} \tag{A.14}$$

Theorem A.6 (**Minimal skew polynomial** [MPSK22, Th. 2.5])**.** *Given a skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$, define the the set*

$$[F] = \{HF \mid H \in \mathbb{F}_{q^m}[x;\sigma]\}.$$

Consider a nonempty set $\Omega \subseteq \mathbb{F}_{q^m}$. Then, there exists a unique monic skew polynomial $F_\Omega \in \mathbb{F}_{q^m}[x;\sigma]$ such that

$$I(\Omega) = [F_\Omega].$$

This means that a skew polynomial $G \in \mathbb{F}_{q^m}[x;\sigma]$ vanishes on Ω if, and only if, G is a left-multiple of F_Ω. The skew polynomial F_Ω is called minimal skew polynomial of Ω in $\mathbb{F}_{q^m}[x;\sigma]$.

Proof. First, the existence of F_Ω for a nonempty set Ω is proven by induction on $|\Omega| \geq 1$ by showing that $I(\Omega) \neq \emptyset$. For $|\Omega| = 1$, $I(\Omega)$ contains the skew polynomial $F = x - a$ for $\Omega = \{a\}$. It is now assumed that $I(\Omega) \neq \emptyset$ for $|\Omega| = n - 1 \geq 1$. Consider $\Omega = \{a_1, a_2, \ldots, a_n\} \subseteq \mathbb{F}_{q^m}$ with $|\Omega| = n$. It follows from the induction hypothesis that there exists a skew polynomial G with $G \in I(\{a_1, a_2, \ldots, a_{n-1}\})$. There are now two possibilities for the n-th element a_n: If $G(a_n) = 0$, then $G \in I(\Omega)$ and $I(\Omega)$ is a nonempty set. Else, set $\beta = G(a_n) \neq 0$. Then for the skew polynomial

$$F = \left(x - \sigma(\beta) a_n \beta^{-1}\right) G \in \mathbb{F}_{q^m}[x;\sigma],$$

it holds that $F \in I(\Omega)$ by Theorem A.5 and therefore $I(\Omega) \neq \emptyset$. Thus, the set $I(\Omega)$ is not empty and it contains a nonzero skew polynomial $G \in I(\Omega)$ that is monic and of minimum possible degree among all nonzero skew polynomials in $I(\Omega)$. Let F be a skew polynomial in $I(\Omega)$. For F, there exists a decomposition with $Q, R \in \mathbb{F}_{q^m}[x;\sigma]$ such that

$$F = QG + R$$

with $\deg(R) < \deg(G)$ by the Euclidean division theorem A.1. It follows that $R = F - QG \in I(\Omega)$ (by (A.14)), but with $\deg(R) < \deg(G)$ and G being the nonzero skew polynomial of minimum possible degree in $I(\Omega)$ this would be a contradiction for $R \neq 0$. Thus $R = 0$ and $F = QG \in [G]$. If F is also nonzero, monic and of minimum degree among nonzero skew polynomials in $I(\Omega)$, by (A.2) $Q = 1$ with $\deg(Q) = 0$ and $F = G$. It can be concluded that

$$I(\Omega) = [G]$$

since $I(\Omega) \subseteq [G]$ and $[G] \subseteq I(\Omega)$, due to $I(\Omega)$ being a left ideal. Furthermore, G is unique and $F_\Omega = G$. □

Proposition A.2 ([MPSK22, Prop. 2.4])**.** *Given any set $\Omega \subseteq \mathbb{F}_{q^m}$, it holds that*

$$\deg(F_\Omega) \leq |\Omega|$$

with $F_\Omega \in \mathbb{F}_{q^m}[x;\sigma]$ being the minimal skew polynomial in $\mathbb{F}_{q^m}[x;\sigma]$ of the set Ω.

Proof. In the proof of Theorem A.6, it was shown in the first part that there exists $F \in I(\Omega)$ with $\deg(F) \leq |\Omega|$. It can be deduced that $\deg(F_\Omega) \leq \deg(F) \leq |\Omega|$ since F is a left multiple of F_Ω. □

This might be a bit confusing since for conventional polynomials $\mathbb{F}_{q^m}[x]$ the degree is greater or equal than the number of distinct zeros. For skew polynomials there is another notion of distinctness, P-independence. P-independent means polynomially independent and was introduced in [Lam85].

Definition A.7 (**P-independence**)**.** A set Ω is said to be P-independent in $\mathbb{F}_{q^m}[x;\sigma]$ if

$$\deg(F_\Omega) = |\Omega|.$$

Having this notion in mind, a bound for distinct zeros at a set of P-independent points exists for skew polynomials. It is similar to the number of distinct zeros of conventional polynomials that is bounded by its degree.

Corollary A.1 ([MPSK22, Cor. 2.6]). *Let $\Omega \subseteq \mathbb{F}_{q^m}$ be a set with cardinality $|\Omega|$ and $F \in \mathbb{F}_{q^m}[x;\sigma]$ a nonzero skew polynomial vanishing on Ω. The set Ω is P-independent if, and only if, it holds that*

$$\deg(F) \geq |\Omega|. \tag{A.15}$$

Proof. Both implications are proven.
$\Longrightarrow$: Suppose Ω is P-independent which means that $\deg(F_\Omega) = |\Omega|$ holds. By Theorem A.6, $F \in [F_\Omega]$. Thus, it can be concluded that

$$|\Omega| = \deg(F_\Omega) \leq \deg(F).$$

$\Longleftarrow$: The minimal skew polynomial F_Ω vanishes on Ω by definition where $F \in [F_\Omega]$. From (A.15), it is given that $\deg(F_\Omega) \geq |\Omega|$. By Proposition A.2, it also holds that $\deg(F_\Omega) \leq |\Omega|$. Therefore, $\deg(F_\Omega) = |\Omega|$ which means that Ω is P-independent. □

For conventional polynomials ($\sigma = \mathrm{Id}$) P-independence is equivalent to distinctness which follows directly from the corollary above. P-independence, zeros of a skew polynomial, and minimal skew polynomials are illustrated in the following example.

Example A.6 ([MPSK22, Ex. 2.9]). Consider $\mathbb{F}_4$ with $q = 2$ and $m = 2$ as defined in (1.1). Take $a_1 = 1$, $a_2 = \omega$ and $a_3 = \bar{\omega}$ as in Example A.5 and consider the set

$$\Omega = \{a_1, a_2, a_3\} = \{1, \omega, \bar{\omega}\}.$$

For the set Ω, $F_\Omega = x^2 + 1$ is the minimal skew polynomial. There is no polynomial of degree 1 which vanishes on 1 and on ω. Since F_Ω can be written as

$$x^2 + 1 = (x+1)(x+1) = (x+\omega)(x+\bar{\omega}) = (x+\bar{\omega})(x+\omega),$$

$F_\Omega \in I(\Omega)$. But 1, ω and $\bar{\omega}$ are not P-independent since $\deg(F_\Omega) < |\Omega|$. The same was shown in Example A.5. The evaluations of a skew polynomial at the points in Ω are not independent. However, consider $\tilde{\Omega} = \{1, \omega\}$. The minimal skew polynomial is again $F_{\tilde{\Omega}} = x^2 + 1$. Here it holds that $\deg(F_{\tilde{\Omega}}) = |\tilde{\Omega}|$ which means that $\tilde{\Omega}$ is P-independent in F.

Equipped with the terminology of P-independence, the Lagrange interpolation theorem for skew polynomials can be provided.

Theorem A.7 (**Lagrange interpolation** [MPSK22, Th. 2.7]).
Let $\Omega = \{a_1, a_2, \ldots, a_n\} \subseteq \mathbb{F}_{q^m}$ be a set with $|\Omega| = n$. The following statements are equivalent:

1. *The set Ω is P-independent in $\mathbb{F}_{q^m}[x;\sigma]$.*

2. *Lagrange interpolation, as defined in Definition A.5, can be performed on $a_1, a_2, \ldots, a_n$, i.e., there exists $F \in \mathbb{F}_{q^m}[x;\sigma]$ for all $b_1, b_2, \ldots, b_n \in \mathbb{F}_{q^m}$ such that $F(a_i) = b_i$ for all $i \in [n]$.*

3. *There exists a unique skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$ for all $b_1, b_2, \ldots, b_n \in \mathbb{F}_{q^m}$ such that $deg(F) < n$ and $F(a_i) = b_i$ for all $i \in [n]$.*

Proof. 1. $\iff$ 3. : Let $\mathbb{F}_{q^m}^{<n}[x;\sigma]$ be the set of skew polynomials in $\mathbb{F}_{q^m}[x;\sigma]$ with degree less than n. Since $\mathbb{F}_{q^m}^{<n}[x;\sigma]$ is a vector space over $\mathbb{F}_{q^m}$ of dimension n (follows from (A.1), (A.2) and (A.3)), a map

$$\phi : \mathbb{F}_{q^m}^{<n}[x;\sigma] \longrightarrow \mathbb{F}_{q^m}^n$$

can be defined with

$$\phi(F) = (F(a_1), F(a_2), \ldots, F(a_n))$$

for all $F \in \mathbb{F}_{q^m}^{<n}[x;\sigma]$. The evaluation map ϕ is linear over $\mathbb{F}_{q^m}$ by Theorem A.4. Statement 3. is equivalent to ϕ being bijective. Since ϕ is linear and the dimensions of the domain and the codomain are equal, i.e., $\dim\left(\mathbb{F}_{q^m}^{<n}[x;\sigma]\right) = \dim\left(\mathbb{F}_{q^m}^n\right) = n$, ϕ is bijective if, and only if, ϕ is injective. Now ϕ is injective by definition if, and only if, any nonzero skew polynomial $F \in \mathbb{F}_{q^m}[x;\sigma]$ vanishing in Ω satisfies $\deg(F) \geq n = |\Omega|$. This is equivalent to Ω being P-independent by Corollary A.1 and the first part is done.

2. $\iff$ 3. : Here only 2. $\implies$ is shown since clearly 3. implies 2. Let $F \in \mathbb{F}_{q^m}[x;\sigma]$ such that $F(a_i) = b_i$ holds for all $i \in [n]$. The Euclidean division theorem (Theorem A.1) implies that there exist $Q, R \in \mathbb{F}_{q^m}[x;\sigma]$ such that

$$F = Q \cdot F_\Omega + R$$

with $\deg(R) < \deg(F_\Omega)$. It can be deduced that $R(a_i) = F(a_i) = b_i$ for all $i \in [n]$ by the product rule (Theorem A.5) and the fact that F_Ω vanishes on Ω. From Proposition A.2, it follows that $\deg(R) < \deg(F_\Omega) \leq |\Omega| = n$ which shows that 3. holds. □

From the above theorem, a corollary follows.

Corollary A.2 ([MPSK22, Cor. 2.8])**.** *Given a P-independent set, any subset of it is P-independent.*

P-independent sets are discussed after introducing Newton interpolation.

The following algorithm is an extension of the classical Newton interpolation algorithm. It provides the minimal skew polynomial as well as the only polynomial that

fulfills the evaluation constraints given a P-independent set of evaluation points with corresponding values. The minimal polynomial is constructed in the same way as the polynomial F_Ω in the first part of the proof of Theorem A.6.

Definition A.8 (**Newton interpolation** [MPSK22, Prop. 2.6])**.** Choose any n evaluation values $b_1, \ldots, b_n \in \mathbb{F}_{q^m}$ and let $\Omega = \{a_1, a_2, \ldots, a_n\} \subseteq \mathbb{F}_{q^m}$ be a P-independent set with $|\Omega| = n$. First, initialize

$$F_1 := x - a_1,$$
$$G_1 := b_1,$$

and then, for $i = 2, 3, \ldots, n$, perform the iteration steps

$$F_i := \left(x - \sigma\left(F_{i-1}(a_i)\right) a_i F_{i-1}(a_i)^{-1}\right) \cdot F_{i-1} = \left(x - {}^{F_{i-1}(a_i)}a_i\right) \cdot F_{i-1}$$
$$G_i := G_{i-1} + \left(b_i - G_{i-1}(a_i)\right) \cdot F_{i-1}(a_i)^{-1} \cdot F_{i-1}$$

Newton's algorithm returns the polynomials F_n and G_n which have the following properties:

$$F_n(a_i) = 0, \ \forall\, i \in \{1, \ldots, n\}$$
$$G_n(a_i) = b_i, \ \forall\, i \in \{1, \ldots, n\}$$

where F_n is the minimal skew polynomial of the set Ω and G_n is the only skew polynomial with degree less than n fulfilling the mentioned evaluation constraints.

In the following example, the Newton interpolation algorithm is illustrated.

Example A.7. Consider $\mathbb{F}_4$ with $q = 2$ and $m = 2$ as defined in (1.1). Take $a_1 = 1$ and $a_2 = \bar{\omega}$, set $\Omega = \{a_1, a_2\}$ and let $b_1 = \bar{\omega}$ and $b_2 = \omega$. Initialize F_1 and G_1 to

$$F_1 := x + 1,$$
$$G_1 := \bar{\omega}$$

as the first step in Newton's algorithm. The second and only additional step returns

$$\begin{aligned} F_2 :&= \left(x - \sigma\left(F_1(a_2)\right) a_2 F_1(a_2)^{-1}\right) \cdot F_1 \\ &= \left(x - \sigma(\omega)\bar{\omega}\omega^{-1}\right)(x+1) \\ &= (x+1) \cdot (x+1) = x^2 + 1 \\ G_2 :&= G_1 + \left(b_2 - G_1(a_2)\right) \cdot F_1(a_2)^{-1} \cdot F_1 \\ &= \bar{\omega} + (\omega - \bar{\omega}) \cdot \omega^{-1} \cdot (x+1) \\ &= \bar{\omega} + \bar{\omega} \cdot (x+1) = \bar{\omega} x \end{aligned}$$

It can be verified that $G_2(1) = \bar{\omega}$ and $G_2(\bar{\omega}) = \omega$. F_2 is the minimal skew polynomial

$F_\Omega = x^2 + 1$ that was discussed in Example A.5.

It is now of interest for the codes that are making use of skew polynomials to link conjugacy, zeros and P-independence.

Theorem A.8 ([MPSK22, Th. 2.9])**.**
Let $a, \beta_1, \beta_2, \ldots, \beta_n \in \mathbb{F}_{q^m}^*$. The set $\Omega = \{a_1, a_2, \ldots, a_n\} \subseteq \mathbb{F}_{q^m}$ is P-independent with $a_i = {}^{\beta_i}a$ for $i \in [n]$ if, and only if, the elements β_i for $i \in [n]$ are $\mathbb{F}_q$-linearly independent.

Proof. Please refer to the proof of Theorem 2.9 in [MPSK22]. □

Theorem A.9 ([MPSK22, Th. 2.10])**.** *Let $\Omega_1, \Omega_2 \subseteq \mathbb{F}_{q^m}$ be nonempty subsets where no element in Ω_1 is conjugate in $\mathbb{F}_{q^m}$ with respect to σ to an element in Ω_2 and both sets are P-independent. Then the union of the sets $\Omega_1 \cup \Omega_2$ is P-independent.*

Proof. Let $\Omega_1 = \{a_1, a_2, \ldots, a_n\}$ and $\Omega_2 = \{b_1, b_2, \ldots, b_m\}$ with cardinality $|\Omega_1| = n$ and $|\Omega_1| = m$ and define $l = n + m$. The proof of the theorem is done by induction on $l \geq 2$ since $n \geq 1$ and $m \geq 1$. For $l = 2$, i.e., $\Omega_1 \cup \Omega_2 = \{a, b\}$, a and b are always P-independent if they are distinct elements in $\mathbb{F}_{q^m}$. Consider $F_{\{a\}} = x - a$. Since $F_{\{a\}}(b) = b - 1 \neq 0$, the skew polynomial that vanishes on a and b must have degree greater or equal than 2, i.e., $\deg(F_{\{a,b\}}) \geq 2$. Thus, by Corollary A.1 the set $\Omega_1 \cup \Omega_2$ is P-independent. It is now assumed that the result holds whenever $n + m < l$, with $l > 2$ and it is in the following shown that the result follows for $n + m = l$. Since $l > 2$ either n or m is greater than 1. Without loss of generality it is assumed that $m > 1$ and therefore

$$\tilde{\Omega}_2 = \{b_1, b_2, \ldots, b_{m-1}\}$$

is not empty. Since Ω_2 is P-independent, $\tilde{\Omega}_2$ is as well P-independent by Corollary A.2. The set $\tilde{\Omega} = \Omega_1 \cup \tilde{\Omega}_2$ is P-independent by the induction hypothesis. From the Newton interpolation (Definition A.8), it is known that there exist $\beta_1, \beta_2, \ldots, \beta_n \in \mathbb{F}_{q^m}^*$ such that

$$F_{\tilde{\Omega}} = \left(x - {}^{\beta_n}a_n\right) \cdots \left(x - {}^{\beta_1}a_1\right) F_{\tilde{\Omega}_2}.$$

Now assume that $\Omega = \Omega_1 \cup \Omega_2$ is not P-independent. Therefore, $\deg(F_\Omega) \leq l - 1$ which follows from Corollary A.1. By Theorem A.6, $F_{\tilde{\Omega}}$ is a left multiple of F_Ω and $\deg(F_{\tilde{\Omega}})$. It follows that $F_\Omega = F_{\tilde{\Omega}}$. For $F_{\tilde{\Omega}_2}$, it holds that $F_{\tilde{\Omega}_2}(b_m) \neq 0$ because if $F_{\tilde{\Omega}_2}(b_m) = 0$, $F_{\tilde{\Omega}_2}$ would be left multiple of F_{Ω_2} but $\deg(F_{\Omega_2}) = m$ while $\deg(F_{\tilde{\Omega}_2}) = m - 1$. In addition, $F_\Omega(b_m) = 0$ since $b_m \in \Omega$. Thus, it can be deduced that there exist $\xi \in \mathbb{F}_{q^m}^*$ and $j \in [n]$ such that

$${}^{\xi}b_m - {}^{\beta_j}a_j = 0$$

by Theorem A.5. This means that b_m is conjugate to a_j in $\mathbb{F}_{q^m}$ with respect to σ which contradicts the hypothesis. Therefore, $\Omega = \Omega_1 \cup \Omega_2$ is P-independent. □

Theorem A.8 and Theorem A.9 together give the following structure theorem for P-independent sets which is directly applied later for linearized Reed-Solomon codes.

Theorem A.10 ([MPSK22, Th. 2.11])**.** *Let Ω be a nonempty set that is a disjoint union of subsets of conjugacy classes, i.e., for pairwise non-conjugate elements $a_1, a_2, \ldots, a_g$, let*

$$\Omega = \Omega_1 \cup \Omega_2 \cup \ldots \Omega_g$$

with $i \in [g]$ and

$$\Omega_i = \Omega \cap C_\sigma(a_i) \neq \emptyset.$$

Denote the cardinality of Ω_i as $n_i = |\Omega_i| = n_i$, for $i \in [g]$. Let the elements in Ω_i be of the form

$$b_j^{(i)} = {}^{\beta_j^{(i)}} a_i$$

for some $\beta_j^{(i)} \in \mathbb{F}_{q^m}^*$ for $j \in [n_i]$. The set Ω is P-independent in $\mathbb{F}_{q^m}[x;\sigma]$ if, and only if, the elements $\beta_1^{(i)}, \beta_2^{(i)}, \ldots, \beta_{n_i}^{(i)} \subseteq \mathbb{F}_{q^m}$ are $\mathbb{F}_q$-linearly independent for each $i \in [g]$.

As $\mathbb{F}_q$-linearly independent elements, a basis of $\mathbb{F}_{q^m}$ over $\mathbb{F}_q$ can be taken. For the application of the theorem above, it is interesting to study the structure of conjugacy classes. A simple way to find distinct conjugacy classes is by taking consecutive powers of a primitive element of $\mathbb{F}_{q^m}$ [LMK17].

Theorem A.11 (**Conjugacy classes with primitive elements** [MPSK22, Th. 2.12])**.** *Let $\gamma \in \mathbb{F}_{q^m}^*$ be a primitive element of $\mathbb{F}_{q^m}$, so that*

$$\mathbb{F}_{q^m}^* = \{\gamma^0, \gamma^1, \gamma^2, \ldots, \gamma^{q^m-2}\}.$$

For all $a \in \mathbb{F}_{q^m}^*$, the corresponding conjugacy class is

$$C_\sigma(a) = \{\gamma^{i(q-1)} \mid 0 \leq i < \frac{q^m-1}{q-1}\} \tag{A.16}$$

with the cardinality

$$|C_\sigma(a)| = \frac{q^m-1}{q-1}. \tag{A.17}$$

In addition, $C_\sigma(0) = \{0\}$ is a special conjugacy class with cardinality 1. Furthermore, the first $q-1$ powers of the primitive element

$$1, \gamma, \gamma^2, \ldots, \gamma^{q-2} \in \mathbb{F}_{q^m}$$

are pairwise non-conjugate, i.e., $\gamma^i \nsim_\sigma \gamma^j$ *for* $i \neq j \in [q-2] \cup \{0\}$*. Moreover, it holds that*

$$\mathbb{F}_{q^m} = C_\sigma(0) \cup C_\sigma(\gamma^0) \cup C_\sigma(\gamma^1) \cup \ldots \cup C_\sigma(\gamma^{q-2}), \tag{A.18}$$

where the union is disjoint. Therefore, $\mathbb{F}_{q^m}$ *can be divided into* $q-1$ *nonzero conjugacy classes with respect to* σ*.*

Proof. First, the proof of (A.16) and (A.17) is given where $C_\sigma(0) = \{0\}$ with $|C_\sigma(0)| = 1$ is trivial. Fix $a \in \mathbb{F}_{q^m}^*$. For $\beta \in \mathbb{F}_{q^m}^*$, the elements in $C_\sigma(a)$ are

$$^\beta a = \beta^{q-1} a.$$

The element β can also be represented by a power of the primitive element, i.e., there exists an integer $i \in \{0, 1, \ldots, q^m - 2\}$ such that $\beta = \gamma^i$. Thus $C_\sigma(a)$ can be written as

$$C_\sigma(a) = \{\gamma^{i(q-1)} \mid 0 \leq i < q^m - 2\}.$$

It holds that $\gamma^{i(q-1)}a = a$ if, and only if, i is multiple of $(q^m-1)/(q-1)$. Therefore, i can be restricted to $0 \leq i < (q^m-1)/(q-1)$. Second, the non-conjugacy of the first $q-1$ powers of a primitive element is proven. Assume, to the contrary, that there exist $i, j \in [q-2] \cup \{0\}$ with $i \neq j$ such that $\gamma^i \sim_\sigma \gamma^j$. Thus, $\gamma^j \in C_\sigma(\gamma^i)$ and by (A.16) there exists a $l \in [(q^m-1)/(q-1)-1] \cup \{0\}$ such that

$$\gamma^j = \gamma^{i+t(q-1)}.$$

This means that $\gamma^{i+t(q-1)}$ can only be a primitive element for $t(q-1) = 0$ since $i, j \leq q-2$, which is a contradiction. Finally, observe that the disjoint union

$$C_\sigma(0) \cup C_\sigma(\gamma^0) \cup C_\sigma(\gamma^1) \cup \ldots \cup C_\sigma(\gamma^{q-2}) \subseteq \mathbb{F}_{q^m}$$

has the following cardinality

$$|C_\sigma(0) \cup C_\sigma(\gamma^0) \cup C_\sigma(\gamma^1) \cup \ldots \cup C_\sigma(\gamma^{q-2})| = 1 + (q-1)\frac{q^m-1}{q-1} = q^m,$$

which shows (A.18). □

Example A.8. Consider the field $\mathbb{F}_{q^m} = \mathbb{F}_4$ with $q = 2$ and $m = 2$ as defined in equation (1.1), where ω is a primitive element of the field. The one conjugacy class besides $C_\sigma(0)$ is

$$C_\sigma(1) = \{1, \omega, \bar{\omega}\}$$

with cardinality

$$|C_\sigma(1)| = \frac{q^m - 1}{q - 1} = 3,$$

which is in accordance to Theorem A.11. The corresponding $\beta_i^{(1)}$ for the conjugacy class $C_\sigma(1)$ are $\beta_1^{(1)} = \omega$ and $\beta_2^{(1)} = \bar{\omega}$ since

$$\omega = {}^{\beta_1^{(1)}}1 = {}^{\bar{\omega}}1 \text{ and } \bar{\omega} = {}^{\beta_2^{(1)}}1 = {}^{\omega}1.$$

The elements $\{1, \omega\}$ are P-independent as seen in Example A.6.

Appendix B

B.1 Information Theory

This section summarizes important rules and (in-)equalities that are used throughout the work. All the following definitions and theorems are taken from [Kra20].

Given a joint entropy of multiple variables $\mathrm{H}(\mathsf{X}_1, \mathsf{X}_2, \ldots, \mathsf{X}_n)$, it holds that

$$\mathrm{H}(\mathsf{X}_1, \mathsf{X}_2, \ldots, \mathsf{X}_n) = \mathrm{H}(\mathsf{X}_1) + \mathrm{H}(\mathsf{X}_2 \mid \mathsf{X}_1) + \ldots + \mathrm{H}(\mathsf{X}_n \mid \mathsf{X}_1, \mathsf{X}_2, \ldots, \mathsf{X}_{n-1}). \tag{B.1}$$

The joint entropy of two random variables $\mathrm{H}(\mathsf{X}, \mathsf{Y})$ can be bounded by

$$\max\left(\mathrm{H}(\mathsf{X}), \mathrm{H}(\mathsf{Y})\right) \leq \mathrm{H}(\mathsf{X}, \mathsf{Y})$$

with equality if X essentially determines Y or if Y essentially determines X, i.e., $\mathsf{X} = f(\mathsf{Y})$ or $\mathsf{Y} = f(\mathsf{X})$ for some function f.

The mutual information $\mathrm{I}(\mathsf{X}; \mathsf{Y})$ can be written as

$$\begin{aligned} \mathrm{I}(\mathsf{X}; \mathsf{Y}) &= \mathrm{H}(\mathsf{X}) - \mathrm{H}(\mathsf{X} \mid \mathsf{Y}) \\ &= \mathrm{H}(\mathsf{Y}) - \mathrm{H}(\mathsf{Y} \mid \mathsf{X}) = \mathrm{I}(\mathsf{Y}; \mathsf{X}) \\ &= \mathrm{H}(\mathsf{X}) + \mathrm{H}(\mathsf{Y}) - \mathrm{H}(\mathsf{X}, \mathsf{Y}) \\ &= \mathrm{H}(\mathsf{X}, \mathsf{Y}) - \mathrm{H}(\mathsf{X} \mid \mathsf{Y}) - \mathrm{H}(\mathsf{Y} \mid \mathsf{X}). \end{aligned} \tag{B.2}$$

The conditional mutual information $\mathrm{I}(\mathsf{X}; \mathsf{Y} \mid \mathsf{Z})$ is upper bounded by

$$\mathrm{I}(\mathsf{X}; \mathsf{Y} \mid \mathsf{Z}) \leq \min\left(\mathrm{H}(\mathsf{X} \mid \mathsf{Z}), \mathrm{H}(\mathsf{Y} \mid \mathsf{Z})\right) \tag{B.3}$$

following from

$$\begin{aligned} \mathrm{I}(\mathsf{X}; \mathsf{Y} \mid \mathsf{Z}) &= \mathrm{H}(\mathsf{X} \mid \mathsf{Z}) - \mathrm{H}(\mathsf{X} \mid \mathsf{Y}, \mathsf{Z}) \\ &= \mathrm{H}(\mathsf{Y} \mid \mathsf{Z}) - \mathrm{H}(\mathsf{Y} \mid \mathsf{X}, \mathsf{Z}) \end{aligned}$$

and the non-negativity of entropy.

B.2 Lagrange Interpolation

Lagrange interpolation is a useful tool to recover polynomials given evaluation points and their corresponding values.

Definition B.1. Given a set of k evaluation points $\{a_0, a_1, ..., a_{k-1}\} \subseteq \mathbb{K}$, which are distinct, i.e., $a_i \neq a_j$ for all $i \neq j$, and given k corresponding evaluation values $\{b_0, b_1, ..., b_{k-1}\} \subseteq \mathbb{K}$, the minimal polynomial fulfilling the constraints $p(a_i) = b_i$ for all $i \in [k-1] \cup \{0\}$ can be recovered by

$$p(x) = \sum_{i=0}^{k-1} b_i \ell_i(x)$$

with $\ell_i(x)$ defined as

$$\ell_i(x) = \prod_{\substack{0 \leq m < k \\ m \neq i}} \frac{x - a_m}{a_i - a_m}.$$

B.3 Automorphism and Derivation

Let $\mathbb{K}$ be a field.

Definition B.2. (Automorphism) A map $\sigma : \mathbb{K} \longrightarrow \mathbb{K}$ is called an automorphism if:

1. σ is bijective,
2. σ is a linear map, i.e., for all $a, b \in \mathbb{K}$ it holds that $\sigma(a + b) = \sigma(a) + \sigma(b)$ and
3. $\sigma(ab) = \sigma(a)\sigma(b)$ for all $a, b \in \mathbb{K}$.

If σ only fulfills 2. and 3., it is called an endomorphism.

Example B.1. Consider the field $\mathbb{F}_{q^m}$ with the Frobenius automorphism $\sigma(x) = x^q$. This automorphism fixes $a \in \mathbb{F}_q$, i.e., $\sigma(a) = a$ for all $a \in \mathbb{F}_q$. For the field $\mathbb{F}_4$ with $q = 2$ and $m = 2$, σ fixes $\mathbb{F}_2 = \{0, 1\}$, which can be seen in Example A.1.

Definition B.3. (Derivation) A map $\delta : \mathbb{K} \longrightarrow \mathbb{K}$ is called a σ-derivation if:

1. δ is a linear map, i.e., for all $a, b \in \mathbb{K}$ it holds that $\delta(a + b) = \delta(a) + \delta(b)$ and
2. $\delta(ab) = \sigma(a)\delta(b) + \delta(a)b$ for all $a, b \in \mathbb{K}$.

List of Abbreviations

DSS	distributed storage system.
GPU	global processing unit.
LRC	locally repairable code.
LRSC	linearized Reed-Solomon code.
MDS	maximum distance separable.
MR	maximally recoverable.
MR-LRC	maximally recoverable locally repairable code.
MSRD	maximum sum-rank distance.
NPU	node processing unit.
RPU	rack processing unit.

Bibliography

[ACRV14] F. Ahmad, S. T. Chakradhar, A. Raghunathan, and T. N. Vijaykumar, "Shufflewatcher: Shuffle-aware scheduling in multi-tenant mapreduce clusters," in *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, ser. USENIX ATC'14. USA: USENIX Association, 2014, p. 1–12.

[Bor08] D. Borthakur. (2008) Hdfs architecture guide. [Online]. Available: https://hadoop.apache.org/docs/r1.2.1/hdfs_design.pdf

[BU13] D. Boucher and F. Ulmer, "Linear codes using skew polynomials with automorphisms and derivations," *Designs, Codes and Cryptography*, vol. 70, 01 2013.

[Cep] Ceph. Ceph configuring monitor/osd interaction. [Online]. Available: https://docs.ceph.com/en/latest/rados/configuration/mon-osd-interaction/

[Cis07] Cisco. (2007) Cisco data center infrastructure 2.5 design guide. San Jose, CA, USA. [Online]. Available: https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.233.2957&rep=rep1&type=pdf

[CMST22] H. Cai, Y. Miao, M. Schwartz, and X. Tang, "A construction of maximally recoverable codes with order-optimal field size," *IEEE Transactions on Information Theory*, vol. 68, no. 1, pp. 204–212, 2022.

[Gab85] E. Gabidulin, "Theory of codes with maximum rank distance (translation)," *Problems of Information Transmission*, vol. 21, pp. 1–12, 01 1985.

[Gan05] W. Gander, "Change of basis in polynomial interpolation," *Numerical Linear Algebra with Applications*, vol. 12, no. 8, pp. 769–778, 2005. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/nla.450

[GG22] S. Gopi and V. Guruswami, "Improved maximally recoverable lrcs using skew polynomials," *IEEE Transactions on Information Theory*, vol. 68, no. 11, pp. 7198–7214, 2022.

[GHJY14] P. Gopalan, C. Huang, B. Jenkins, and S. Yekhanin, "Explicit maximally recoverable codes with locality," *IEEE Transactions on Information Theory*, vol. 60, no. 9, pp. 5245–5256, 2014.

[GHSY12] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, "On the locality of codeword symbols," *IEEE Transactions on Information Theory*, vol. 58, no. 11, pp. 6925–6934, 2012.

[GM84] S. Goldwasser and S. Micali, "Probabilistic encryption," *Journal of Computer and System Sciences*, vol. 28, no. 2, pp. 270–299, 1984. [Online]. Available: https://www.sciencedirect.com/science/article/pii/0022000084900709

[Ham50] R. W. Hamming, "Error detecting and error correcting codes," *The Bell System Technical Journal*, vol. 29, no. 2, pp. 147–160, 1950.

[HCL07] C. Huang, M. Chen, and J. Li, "Pyramid codes: Flexible schemes to trade space for access efficiency in reliable data storage systems," in *Sixth IEEE International Symposium on Network Computing and Applications (NCA 2007)*, 2007, pp. 79–86.

[HSX+12] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin, "Erasure coding in windows azure storage," in *2012 USENIX Annual Technical Conference (USENIX ATC 12)*. Boston, MA: USENIX Association, Jun. 2012, pp. 15–26. [Online]. Available: https://www.usenix.org/conference/atc12/technical-sessions/presentation/huang

[IPC+09] M. Isard, V. Prabhakaran, J. Currey, U. Wieder, K. Talwar, and A. Goldberg, "Quincy: Fair scheduling for distributed computing clusters," in *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles*, ser. SOSP '09. New York, NY, USA: Association for Computing Machinery, 2009, p. 261–276. [Online]. Available: https://doi.org/10.1145/1629575.1629601

[Kra20] G. Kramer, *Information Theory - Lecture notes WS 2020/21*, Munich, 2020.

[Lam85] T. Lam, "A general theory of vandermonde matrices," *Expositiones Mathematicae*, vol. 4, 01 1985.

[Ler95] A. Leroy, "Pseudolinear transformations and evaluation in ore extensions," *Bulletin of the Belgian Mathematical Society - Simon Stevin*, vol. 2, 01 1995.

[LL88] T. Lam and A. Leroy, "Vandermonde and wronskian matrices over division rings," *Journal of Algebra*, vol. 119, no. 2, pp. 308–336, 1988. [Online]. Available: https://www.sciencedirect.com/science/article/pii/0021869388900634

[LMK17] S. Liu, F. Manganiello, and F. R. Kschischang, "Matroidal structure of skew polynomial rings with application to network coding," *Finite Fields and Their Applications*, vol. 46, pp. 326–346, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1071579717300394

[MP17] U. Martínez-Peñas, "Skew and linearized reed-solomon codes and maximum sum rank distance codes over any division ring," *Journal of Algebra*, vol. 504, 10 2017.

[MPK19] U. Martínez-Peñas and F. Kschischang, "Universal and dynamic locally repairable codes with maximal recoverability via sum-rank codes," *IEEE Transactions on Information Theory*, vol. PP, pp. 1–1, 06 2019.

[MPSK22] U. Martínez-Peñas, M. Shehadeh, and F. R. Kschischang, "Codes in the sum-rank metric: Fundamentals and applications," *Foundations and Trends® in Communications and Information Theory*, vol. 19, no. 5, pp. 814–1031, 2022. [Online]. Available: http://dx.doi.org/10.1561/0100000120

[MS81] R. J. McEliece and D. V. Sarwate, "On sharing secrets and reed-solomon codes," *Commun. ACM*, vol. 24, no. 9, p. 583–584, sep 1981. [Online]. Available: https://doi.org/10.1145/358746.358762

[NUF10] R. W. Nóbrega and B. F. Uchôa-Filho, "Multishot codes for network coding using rank-metric codes," in *2010 Third IEEE International Workshop on Wireless Network Coding*, 2010, pp. 1–6.

[OD11] F. Oggier and A. Datta, "Self-repairing homomorphic codes for distributed storage systems," in *2011 Proceedings IEEE INFOCOM*, 2011, pp. 1215–1223.

[Ore33] O. Ore, "Theory of non-commutative polynomials," *Annals of Mathematics*, vol. 34, no. 3, pp. 480–508, July 1933.

[PERR11] S. Pawar, S. El Rouayheb, and K. Ramchandran, "Securing dynamic distributed storage systems against eavesdropping and adversarial attacks," *IEEE Transactions on Information Theory*, vol. 57, no. 10, pp. 6734–6753, 2011.

[PKLK12] N. Prakash, G. M. Kamath, V. Lalitha, and P. V. Kumar, "Optimal linear codes with a local-error-correction property," in *2012 IEEE International Symposium on Information Theory Proceedings*, 2012, pp. 2776–2780.

[RBS+22] V. Ramkumar, S. B. Balaji, B. Sasidharan, M. Vajha, M. N. Krishnan, and P. V. Kumar, "Codes for distributed storage," *Foundations and Trends® in Communications and Information Theory*, vol. 19, no. 4, pp. 547–813, 2022. [Online]. Available: http://dx.doi.org/10.1561/0100000115

[RKSV14] A. S. Rawat, O. O. Koyluoglu, N. Silberstein, and S. Vishwanath, "Optimal locally repairable and secure codes for distributed storage systems," *IEEE Transactions on Information Theory*, vol. 60, no. 1, pp. 212–236, 2014.

[Rot91] R. Roth, "Maximum-rank array codes and the their application to crisscross error correction," *Information Theory, IEEE Transactions on*, vol. 37, pp. 328 – 336, 04 1991.

[RS60] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," *Journal of the Society for Industrial and Applied Mathematics*, vol. 8, no. 2, pp. 300–304, 1960. [Online]. Available: https://doi.org/10.1137/0108018

[SAP+13] M. Sathiamoorthy, M. Asteris, D. Papailiopoulos, A. G. Dimakis, R. Vadali, S. Chen, and D. Borthakur, "Xoring elephants: Novel erasure codes for big data," *Proc. VLDB Endow.*, vol. 6, pp. 325–336, 2013.

[Sha49] C. E. Shannon, "Communication theory of secrecy systems," *The Bell System Technical Journal*, vol. 28, no. 4, pp. 656–715, 1949.

[Sha79] A. Shamir, "How to share a secret," *Commun. ACM*, vol. 22, no. 11, p. 612–613, nov 1979. [Online]. Available: https://doi.org/10.1145/359168.359176

[SK22] M. Shehadeh and F. R. Kschischang, "Space–time codes from sum-rank codes," *IEEE Transactions on Information Theory*, vol. 68, no. 3, pp. 1614–1637, 2022.

[SRK11] N. B. Shah, K. V. Rashmi, and P. V. Kumar, "Information-theoretically secure regenerating codes for distributed storage," in *2011 IEEE Global Telecommunications Conference - GLOBECOM 2011*, 2011, pp. 1–5.

[TCS19] A. Tebbi, T. H. Chan, and C. W. Sung, "Multi-rack distributed data storage networks," *IEEE Transactions on Information Theory*, vol. 65, no. 10, pp. 6072–6088, 2019.

[WBM+06] S. A. Weil, S. A. Brandt, E. L. Miller, D. D. E. Long, and C. Maltzahn, "Ceph: A scalable, high-performance distributed file system," in *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*, ser. OSDI '06. USA: USENIX Association, 2006, p. 307–320.

[Zha10] Y. Zhang, "A secret sharing scheme via skew polynomials," in *Proceedings of the 2010 International Conference on Computational Science and Its Applications*, ser. ICCSA '10. USA: IEEE Computer Society, 2010, p. 33–38. [Online]. Available: https://doi.org/10.1109/ICCSA.2010.32

Acknowledgements

A very special thank you goes to Prof. Frank Kschischang who supervised this thesis after hosting my master's research project in Toronto. It has been a pleasure and an inspiration to work with him and benefit from his excellent ideas and comments. I would like to thank Dr. Rawad Bitar very much. He contributed significantly to my way of doing scientific research. It has been great that he guided me through my master's studies and gave me a lot of valuable feedback. Many thanks also to Lia Liu who helped me a lot with practical hints and ideas throughout this thesis. I also want to thank Rawad and Dr.-Ing. Sven Puchinger for offering me a seminar project which started the master's research trias of seminar work, research project and master's thesis. A big thank you goes to Prof. Antonia Wachter-Zeh for giving me the opportunity to conduct the three research parts of my master's degree as one unit and arranging the contact with Prof. Kschischang.

Last but not least, I would like to express my sincere gratitude to my family and friends for their support throughout my masters. Vielen Dank!