

June 2024

Optimal Selection of Good Polynomials and Constructions of Locally Recoverable Codes via Galois Theory

Austin Dukes
University of South Florida

Follow this and additional works at: <https://digitalcommons.usf.edu/etd>



Part of the [Mathematics Commons](#)

Scholar Commons Citation

Dukes, Austin, "Optimal Selection of Good Polynomials and Constructions of Locally Recoverable Codes via Galois Theory" (2024). *USF Tampa Graduate Theses and Dissertations*.
<https://digitalcommons.usf.edu/etd/10503>

This Dissertation is brought to you for free and open access by the USF Graduate Theses and Dissertations at Digital Commons @ University of South Florida. It has been accepted for inclusion in USF Tampa Graduate Theses and Dissertations by an authorized administrator of Digital Commons @ University of South Florida. For more information, please contact digitalcommons@usf.edu.

Optimal Selection of Good Polynomials and Constructions of Locally Recoverable Codes
via Galois Theory

by

Austin Dukes

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Mathematics & Statistics
College of Arts and Sciences
University of South Florida

Co-Major Professor: Giacomo Micheli, Ph.D.
Co-Major Professor: Lukas Kölsch, Ph.D.
Jean-François Biasse, Ph.D.
Xiang-Dong Hou, Ph.D.
Joachim Rosenthal, Ph.D.

Date of Approval:
June 12, 2024

Keywords: good polynomial, locally recoverable code, hierarchical locally recoverable code,
chebotarev density theorem

Copyright © 2024, Austin Dukes

DEDICATION

To Peter, the love of my life.

TABLE OF CONTENTS

List of Tables	iii
List of Figures	iv
Abstract	v
Chapter 1 Introduction	1
1.1 History of Coding Theory	1
1.2 Subject of Dissertation	3
Chapter 2 Coding Theory and Codes	4
2.1 Preliminaries	4
2.2 Some Bounds on Codes	5
2.3 Linear Algebra	9
Chapter 3 Locally Recoverable Codes	11
3.1 Introduction	11
3.2 Singleton-like Bound	12
3.3 Tamo-Barg Construction	13
3.4 Modified Constructions	19
3.4.1 Optimal LRCs Without $r \mid k$	19
3.4.2 Almost Optimal LRCs Without $(r + 1) \mid n$	20
Chapter 4 Algebraic Number Theory and Chebotarev Density Theorem	24
4.1 Preliminaries	24
4.1.1 Field Theory	24
4.1.2 Algebraic Number Theory	25
4.2 Galois Theory and Chebotarev Density Theorem	32
4.2.1 Polynomial Maps	32
4.2.2 Chebotarev Density Theorem	33
Chapter 5 Good Polynomials of Degree Up to Five	38
5.1 Introduction	38
5.2 Monodromy Groups and Totally Split Places	40
5.3 Degrees Up to 4	44
5.4 Degree 5: $\text{AGL}_1(\mathbb{F}_5)$ Never Occurs	47
5.5 Degree 5: If D_5 or A_5 Occurs, Then $5 \mid (q^2 - 1)$	52
5.6 Computational Examples	53
Chapter 6 Hierarchical Good Polynomials	55
6.1 Introduction	55
6.1.1 Definitions	56
6.1.2 Motivation	57

6.1.3 Our Contribution	57
6.2 An Improved Bound for Hierarchical Locally Recoverable Codes	58
6.2.1 The Singleton Bound for $[n, k, d, b, a, \lambda]$ HLRCs With $\lambda \leq b$	58
6.3 Our Construction of Optimal HLRCs Using Nested f -adic Expansions	60
6.3.1 Main Tool for the Construction	60
6.3.2 The Main Construction	61
6.3.3 Locality	63
6.3.4 Optimality of the Code	64
6.3.5 Comparison With Another Optimal Hierarchical RS-like Code	66
6.3.6 Toy Example	67
6.4 Existential Results via Chebotarev Density Theorem	69
6.4.1 The Number of Totally Split Places t_0 of $f(h) - t$	70
6.5 Practical Choice of Parameters to Construct Optimal HLRCs	72
References	73
Appendix A The Genus of a Function Field	78

LIST OF TABLES

Table 1: Comparing the asymptotic and actual number of totally split places for a few degree 5 polynomials defined over varying base fields.	54
--	----

LIST OF FIGURES

Figure 1: The nest hierarchy for $f(X) = X^2$ and $h(X) = X^3$	68
--	----

ABSTRACT

To keep up with the ever-growing demand for reliable and efficient availability of data, locally recoverable codes (LRCs) have been the focus of much study due to their applications in cloud and distributed storage systems. A fundamental construction of LRCs was given in [36] based on polynomials which relied on the existence of r -good polynomials. In the same paper some constructions of good polynomials were given, but these constructions did not cover every configuration of parameters. Naturally this led to research into constructing good polynomials for what was not addressed in [36], but new ponderings were also posed, such as the following: for a locally recoverable code with fixed parameters, are some good polynomials *better* than others? In [29], Micheli approached this topic using machinery from Galois theory and defined more generally (r, ℓ) -good polynomials and showed that maximizing ℓ (for a given locality r) yields a code with higher dimension. In this dissertation, we determine the optimum value of ℓ for any (r, ℓ) -good polynomials of degree up to 5. We then provide an explicit construction of a newer class of LRCs called hierarchical locally recoverable codes (HLRCs) and prove that our construction yields codes with higher dimension than existing literature in some regimes of parameters. This dissertation is intended to be self-contained. As a result, the first few chapters address the fundamental background in coding theory, algebraic number theory, etc., needed to understand the aforementioned results.

CHAPTER 1: INTRODUCTION

1.1 History of Coding Theory

As technology continues to rapidly develop, the idea of being “connected” has become somewhat commonplace nowadays. The average person can access data in a multitude of ways, from making a call or sending a text, to watching a movie, to accessing a file from a cloud storage system. In each of these situations, the reliability of the content being transmitted is of vital import. Discrepancies (typically erasures or errors) can occur in data transmission for various reasons, such as physical interference, environmental factors, or equipment malfunctions, to name a few. Also, bits or entire packets of the transmitted data can be affected, though this dissertation will be concerned only with the effect of noise on the bits of the data. Increasing the integrity of received data by accounting for and correcting errors due to noise is the primary goal of coding theory.

Suppose data is transmitted from a sender to a recipient over a noisy channel. The fundamental problem is as follows:

Problem 1.1. *How can we efficiently increase the integrity of the data obtained by the recipient?*

Since this statement is rather vague, we will begin by discussing a couple simple potential solutions and the obstacles they face before addressing the more intricate methods which are the topic of this dissertation.

The approach which is perhaps the most obviously impractical is to request re-transmission of the data from the sender. Not only can re-transmission be costly and time-consuming,

but there are many cases where the sender no longer has access to the data, temporarily or even permanently. For example, a server might be down due to a power outage, or the data may have been stored in a data center that experienced a fire. To combat this, a way of correcting errors in the data (up to a certain extent) is incorporated along with the data so that the error correction, if it is possible, depends only on the data obtained by the recipient. This method is called forward error correction.

A trivial but classical example of forward error correction is to simply repeat the bits within the data. As an example, if the sender wishes to transmit the string 1010 then they could send the string 111000111000. As such, if an error occurs during transmission and the recipient obtains 011000111000, they can be reasonably confident that the first symbol is an error and hence that the data was 1010. However, this comes at the cost of sending three times as much as data, so for practical applications this method requires far too much overhead. Furthermore, one is only guaranteed to be able to recover the data in the presence of a single error. If again the string 111000111000 is sent and there are two errors in what is received, it may or may not be possible to recover the data. This can be seen from the strings 011000111001 (for which it is possible) and 001000111000 (for which it is not possible).

To summarize, we seek a way to transmit data of an arbitrary size in such a way that the recipient of the transmission is able to correct a maximal number of erasures or errors as efficiently as possible, i.e., with as little overhead as possible. To this end, modern coding theory has focused on error-correcting codes. Error-correcting codes as a whole aim to address the discrepancies, or the errors, which result from the transmission of data over a channel by encoding the original data in a way which allows the user to successfully recover the data in the presence of errors.

There are many classes of error-correcting codes, and we will discuss a few particularly relevant ones in [Chapter 2](#).

1.2 Subject of Dissertation

This dissertation is organized as follows. The current chapter serves as an overview of the subject area and provides some motivation into its practical applications. Chapters 2 and 3 provide the necessary background in coding theory for the techniques and results in each of Chapter 5 and Chapter 6 but is not otherwise necessary. In a similar manner Chapter 4 provides the necessary background in algebraic number theory for the techniques and results in Chapters 5 and 6 and, in particular, provides greater detail about the Chebotarev Density Theorem which is employed in those following chapters. Finally, the main results of the dissertation are contained in the last two chapters. In Chapter 5 we classify all good polynomials of degree up to 5 and, in turn, provide an explicit estimate on the maximal length and dimension of a Tamo-Barg code (see [36]). The last chapter, Chapter 6, generalizes the ideas of [36] via nested polynomials to allow for one to simultaneously obtain the efficiency in decoding granted by locality in the case of a single erasure *and* in the case of multiple erasures.

Note that each chapter in this dissertation is intended to be readable on its own, so any necessary notation will be restated at the beginning of each chapter.

For the reader's convenience, in this section we collect any notation which is used throughout the entire thesis.

- The cardinality of the set S will be denoted $|S|$.
- The union of two disjoint sets A and B is denoted $A \sqcup B$.
- \mathbb{F}_q denotes the finite field with q elements.
- Given a function $f : D \rightarrow R$ and some $A \subseteq D$, the restriction of f to A will be denoted by $f|_A$.

CHAPTER 2: CODING THEORY AND CODES

2.1 Preliminaries

Modern coding theory has been heavily interested in the theory of error-correcting codes. For the sake of being self-contained, we dedicate this chapter to addressing any necessary technical background needed to understand the techniques and results in Chapters 5 and 6. The experienced reader may skip this chapter as there are no original results contained within.

A code of length n over the alphabet \mathbb{F}_q is a set \mathcal{C} of elements of \mathbb{F}_q^n , and an element $c \in \mathcal{C}$ is called a *codeword* of \mathcal{C} . The (Hamming) weight $w(c)$ of the codeword $c \in \mathcal{C}$ is the number of nonzero entries in c , and the hamming distance $d(c_1, c_2)$ between two codewords $c_1, c_2 \in \mathcal{C}$ is the number of distinct components between them; equivalently, $d(c_1, c_2) = w(c_1 - c_2)$. If the minimum distance between distinct codewords of \mathcal{C} is d , we say that \mathcal{C} has distance d . A code \mathcal{C} is said to be linear if every linear combination of codewords in \mathcal{C} is also a codeword in \mathcal{C} , that is, if every $c_1, c_2 \in \mathcal{C}$ and $\alpha, \beta \in \mathbb{F}_q$ satisfy $\alpha c_1 + \beta c_2 \in \mathcal{C}$. One can easily check that a linear code forms a vector subspace of \mathbb{F}_q^n (and hence a vector space over \mathbb{F}_q), and we say that \mathcal{C} has dimension k if $\dim_{\mathbb{F}_q} \mathcal{C} = k$. If \mathcal{C} is a code over \mathbb{F}_q of length n which has minimum distance d , then we call \mathcal{C} an $(n, d)_q$ code (or an (n, d) code if the choice of alphabet is clear). If \mathcal{C} is further known to be linear (of dimension, say, k), then we will instead call \mathcal{C} an $[n, k, d]_q$ code (or an $[n, k, d]$ code if, again, the choice of alphabet is clear).

2.2 Some Bounds on Codes

In general, we want to maximize k and d while minimizing n . Codes with larger dimensions are able to encode more information, and codes with larger distances permit increased error-correction. On the other hand, shorter codes require less storage overhead. From this perspective, a vital concern in practical applications is how to optimize each of the parameters simultaneously. One class of codes which accomplishes this is *maximum distance separable* (MDS) codes. Before we can formally define this class of codes, we will first state and prove a fundamental theorem known as the Singleton bound.

Theorem 2.1 (Singleton bound, [25, Theorem 5.2.1]). *For any positive integers n, d, q with $q \geq 2$, an $(n, d)_q$ code \mathcal{C} satisfies*

$$|\mathcal{C}| \leq q^{n-d+1}.$$

Proof. Let $\mathcal{C}' = \{(c_1, c_2, \dots, c_{n-(d-1)}) : (c_1, c_2, \dots, c_n) \in \mathcal{C}\}$ be the (punctured) code obtained by deleting the last $d - 1$ components of each codeword in \mathcal{C} . The length of \mathcal{C}' is clearly $n - (d - 1) = n - d + 1$. Since the minimum distance of \mathcal{C} is d , no two distinct codewords in \mathcal{C} can agree on all of the first $n - d + 1$ coordinates, so we see that $|\mathcal{C}'| = |\mathcal{C}|$. We conclude that $|\mathcal{C}| \leq |\mathbb{F}_q^{n-d+1}| = q^{n-d+1}$. \square

If \mathcal{C} is known further to be linear of dimension k , then $|\mathcal{C}| = q^k$, yielding the following corollary:

Corollary 2.2 ([25, Corollary 5.2.2]). *If \mathcal{C} is an $[n, k, d]_q$ code, then*

$$k \leq n - d + 1.$$

A (linear) $[n, k, d]_q$ code \mathcal{C} is called a maximum distance separable (MDS) code if its parameters satisfy the Singleton bound, that is, if $k = n - d + 1$.

We next state and prove the Hamming bound, which is commonly called the sphere-packing bound. The aptness of this name becomes clear after considering the bound's proof along with the following definition.

For $x \in \mathbb{F}_q^n$, let $B(x, r) \subseteq \mathbb{F}_q^n$ be the ball of radius r centered at x , i.e., $B(x, r) = \{y \in \mathbb{F}_q^n : d(x, y) \leq r\}$. Observe immediately that

$$\begin{aligned} |B(x, r)| &= \sum_{i=0}^r |\{y \in \mathbb{F}_q^n : d(x, y) = i\}| \\ &= \sum_{i=0}^r \binom{n}{i} (q-1)^i. \end{aligned}$$

Theorem 2.3 (Hamming bound, [25, Theorem 5.2.7]). *For any positive integers n, d, q with $q \geq 2$ and $e = \lfloor \frac{d-1}{2} \rfloor$, an $(n, d)_q$ code \mathcal{C} satisfies*

$$|\mathcal{C}| \leq \frac{q^n}{\sum_{i=0}^e \binom{n}{i} (q-1)^i}.$$

Proof. Since the distance of \mathcal{C} is d , one can easily see by the triangle inequality that the balls $B(c_1, e)$ and $B(c_2, e)$ are disjoint for any two distinct codewords $c_1, c_2 \in \mathcal{C}$. Thus

$$\begin{aligned} |\mathbb{F}_q^n| &\geq \left| \bigcup_{c \in \mathcal{C}} B(c, e) \right| \\ &= |\mathcal{C}| \sum_{i=0}^e \binom{n}{i} (q-1)^i. \end{aligned}$$

The bound now follows. □

Our last upper bound on the size of \mathcal{C} is known as the Plotkin bound. It was initially proved for binary codes, but the proof we present works for any q .

Theorem 2.4 (Plotkin bound, [25, Theorem 5.2.4]). *For any positive integers n, d, q with $q \geq 2$ and $\theta = \left(1 - \frac{1}{q}\right)$, if \mathcal{C} is an $(n, d)_q$ code satisfying $d > \theta n$, then*

$$|\mathcal{C}| \leq \frac{d}{d - \theta n}.$$

Proof. Write $|\mathcal{C}| = M$. Since the distance of \mathcal{C} is d , we have $\sum_{x \neq y \in \mathcal{C}} d(x, y) \geq M(M - 1)d$. On the other hand, for $1 \leq i \leq n$ and $0 \leq j \leq q - 1$ define $m_{i,j}$ to be the number of codewords in \mathcal{C} whose i th component is j . Then we have

$$\sum_{x \neq y \in \mathcal{C}} d(x, y) = \sum_{i=1}^n \sum_{j=0}^{q-1} m_{i,j}(M - m_{i,j}) = \sum_{i=1}^n \left(M^2 - \sum_{j=0}^{q-1} m_{i,j}^2 \right),$$

where the last equality follows from the fact that $\sum_{j=0}^{q-1} m_{i,j} = M$ for any i . Further, for any fixed i , Cauchy-Schwarz gives us

$$M^2 - \sum_{j=0}^{q-1} m_{i,j}^2 \leq M^2 - \frac{1}{q} \left(\sum_{j=0}^{q-1} m_{i,j} \right)^2 = M^2 \theta.$$

Hence

$$\sum_{x \neq y \in \mathcal{C}} d(x, y) \leq \sum_{i=1}^n \theta M^2 = n\theta M^2.$$

Combining this with the first inequality yields

$$M(M - 1)d \leq n\theta M^2.$$

The bound now follows. □

The next bound is quite similar in appearance to the Hamming bound, but this one provides a *lower* bound for the size of a maximal $(n, d)_q$ code \mathcal{C} . We note that \mathcal{C} is a *maximal* $(n, d)_q$ code if for any $(n, d)_q$ code \mathcal{C}' we have $|\mathcal{C}| \geq |\mathcal{C}'|$.

Theorem 2.5 (Gilbert-Varshamov bound, [25, Theorem 5.1.7]). *For any positive integers n, d, q with $q \geq 2$ and $d \leq n$, if \mathcal{C} is a maximal $(n, d)_q$ code \mathcal{C} then*

$$|\mathcal{C}| \geq \frac{q^n}{\sum_{i=0}^{d-1} \binom{n}{i} (q-1)^i}.$$

Proof. If \mathcal{C} is maximal, then there is no element $x \in \mathbb{F}_q^n \setminus \mathcal{C}$ such that $d(x, c) \geq r$ for all $c \in \mathcal{C}$ since otherwise the code $\mathcal{C}' = \mathcal{C} \cup \{x\}$ is an $(n, r)_q$ code with $|\mathcal{C}'| > |\mathcal{C}|$, which contradicts the maximality of \mathcal{C} . Thus

$$\begin{aligned} |\mathbb{F}_q^n| &\geq \left| \bigcup_{c \in \mathcal{C}} B(c, d-1) \right| \\ &= |\mathcal{C}| \sum_{i=0}^{d-1} \binom{n}{i} (q-1)^i. \end{aligned}$$

The bound now follows. □

We note that the bounds given in this section (besides that of Corollary 2.2) apply to arbitrary codes. In this dissertation, however, we will be concerned solely with linear codes. As such, we now turn our focus to linear codes in particular. Moreover, from this point onward (unless otherwise specified) when we say that \mathcal{C} is a code, we mean that \mathcal{C} is a linear code. If \mathcal{C} is a linear code over \mathbb{F}_q with length n and dimension k , we will briefly write that \mathcal{C} is an $[n, k]_q$ code. If the distance of \mathcal{C} is known to be d , then we may write that \mathcal{C} is an $[n, k, d]_q$ code instead.

2.3 Linear Algebra

Since a linear code \mathcal{C} forms a vector space over \mathbb{F}_q , linear algebra becomes an invaluable tool in the study of these codes. We will identify the codeword $c \in \mathcal{C}$ with the row vector $[c_1 \ c_2 \ \cdots \ c_n]$, where $c_i \in \mathbb{F}_q$ are the components of the codeword c and n is the code length. If the codeword in question is already denoted by c_i , then we will enumerate its components as $c_i(1), c_i(2), \dots, c_i(n)$ to prevent any overlapping of notation.

Suppose \mathcal{C} is an $[n, k, d]_q$ code. Then one can select k linearly independent codewords $c_1, c_2, \dots, c_k \in \mathcal{C} \subseteq \mathbb{F}_q^n$ which span \mathcal{C} . (In other words, the c_i form a basis for \mathcal{C} .) With these, we define the $k \times n$ matrix G to be the matrix whose i th row is the codeword c_i . This allows us to express any codeword $x \in \mathcal{C}$ as

$$x = [x_1 \ x_2 \ \cdots \ x_n] = [x_1^* \ x_2^* \ \cdots \ x_k^*] \underbrace{\begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_k \end{bmatrix}}_G = x^* G,$$

where the element $x^* = [x_1^* \ x_2^* \ \cdots \ x_k^*] \in \mathbb{F}_q^k$ is seen to be unique by comparing the dimensions of \mathcal{C} and \mathbb{F}_q^k . The matrix G defined above is called a generator matrix for the code \mathcal{C} , and we say that G is in standard form if $G = \begin{bmatrix} I_k & | & G' \end{bmatrix}$, where I_k is the $k \times k$ identity matrix and G' is a $k \times (n-k)$ matrix over \mathbb{F}_q . Two codes \mathcal{C} and \mathcal{C}' are called equivalent if there is a monomial matrix T such that $T : \mathcal{C} \rightarrow \mathcal{C}'$ is a vector space isomorphism. For an arbitrary linear code there may be no generator matrix in standard form, but it is well-known that every linear code is equivalent to a code with a generator matrix in standard form.

We will not provide an explicit formula for the map $\text{Dec}(\cdot) : \mathcal{C} \rightarrow \mathbb{F}_q^k$ given by $x \mapsto x^*$ (though such a formula can be obtained using one of various elementary techniques, such as Gaussian elimination). Nevertheless, since G has full rank, this map is a bijection, and

we define Enc to be the inverse map to Dec . Observe immediately that $\text{Enc}(x) = xG$. We *encode* a message vector $x \in \mathbb{F}_q^k$ by computing $\text{Enc}(x)$, and, similarly, we *decode* a codeword $c \in \mathcal{C}$ by computing $\text{Dec}(c)$.

Now, given an $[n, k]_q$ code \mathcal{C} , the dual code \mathcal{C}^\perp to \mathcal{C} is defined as the set of $x \in \mathbb{F}_q^n$ satisfying $xy^T = 0$ for every $y \in \mathcal{C}$. The dual code is easily seen to be an $[n, n - k]_q$ code, and a generator matrix for \mathcal{C}^\perp is called a parity-check matrix for \mathcal{C} .

Example 2.6. Let \mathcal{C} be the $(5, 4, d)_2$ code obtained by the generator matrix

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

For $x = [x_1 \ x_2 \ x_3 \ x_4] \in \mathbb{F}_2^4$ we have

$$\text{Enc}(x) = xG = [x_1 \ x_2 \ x_3 \ x_4 \ x_1 + x_2 + x_3 + x_4],$$

so we see that

$$\mathcal{C} = \{(c_1, c_2, c_3, c_4, c_5) \in \mathbb{F}_2^5 : c_1 + c_2 + c_3 + c_4 = c_5\}.$$

Further, \mathcal{C} is easily seen to have minimum distance $d = 2$.

In the next chapter, we discuss some of the fundamental notions and results pertaining to a class of codes known as locally recoverable codes. Such codes are a primary focus of applications of the original theoretical results written in this dissertation.

CHAPTER 3: LOCALLY RECOVERABLE CODES

3.1 Introduction

In this chapter we will restrict our attention to codes with locality, so we begin with the basic definitions. A (linear) $[n, k, d]$ code \mathcal{C} is said to have locality parameter r if every codeword $c \in \mathcal{C}$ is such that each component of c can be recovered by accessing at most r other components of c . More formally, \mathcal{C} has locality r if for every $c \in \mathcal{C}$ and $1 \leq i \leq n$ there is a set A_i of indices such that $i \in A_i$ and the component c_i is a function of the components c_j for $j \in A_i \setminus \{i\}$. Note that the set A_i depends only on the index i and not on the choice of codeword. We refer to the set A_i as a locality set for the component in position i (or simply for position i).

We define and provide a construction of a class of codes having *multiple* localities in Chapter 6. In a practical application, it is far more likely that a single erasure will occur in a codeword than that two or more erasures will occur, and this is particularly what locally recoverable codes help to address. However, locally recoverable codes do not necessarily provide an advantage in the case of multiple erasures. It may be the case in a specific application that having exactly one erasure is the most likely scenario but that having exactly two erasures is common enough for one to desire a code which allows for efficient repairing in both cases. This is the typical scenario the construction in Chapter 6 aims to address.

In this chapter we will only discuss locally recoverable codes with a single locality parameter, and we will write that a code \mathcal{C} is an $[n, k, d, r]$ code to mean that \mathcal{C} is an $[n, k, d]$

code with locality r . Our first goal is to define *optimal* LRCs, which are locally recoverable codes whose parameters satisfy with equality a particular bound similar in nature to that of the Singleton bound.

3.2 Singleton-like Bound

Recall the Singleton bound for linear codes, which states that the parameters of an $[n, k, d]$ code \mathcal{C} satisfy $d \leq n - k + 1$. We state and prove the following bound, which is commonly called the Singleton-like bound.

Theorem 3.1 (Singleton-like bound). *Let \mathcal{C} be an $[n, k, d, r]$ LRC. Then*

$$k - 1 + \left\lfloor \frac{k - 1}{r} \right\rfloor \leq n - d.$$

To prove the above theorem, we first state the following well-known lemma, and we include a proof for completeness. We note that the first proof of Theorem 3.1 is due to Gopalan et al. in [19]. In what follows we write $M_{m \times n}(q)$ to denote the set of all matrices of dimension $m \times n$ defined over \mathbb{F}_q .

Lemma 3.2 ([13, Proposition 2.1]). *Let \mathcal{C} be an $[n, k, d]_q$ code with generator matrix $\mathcal{G} \in M_{k \times n}(q)$ and let $S \in M_{k \times t}(q)$ be a submatrix of \mathcal{G} . If $\text{rank}(S) \leq k - 1$ then $t \leq n - d$.*

Proof. Let $S = [S_1, \dots, S_t]$, where S_i is a column of \mathcal{G} for $i \in \{1, \dots, t\}$. Define $\tilde{S}: \mathbb{F}_q^k \rightarrow \mathbb{F}_q^t$ such that $x \mapsto \tilde{S}(x) = xS = [xS_1, \dots, xS_t]$. Since $\text{rank}(S) \leq k - 1$ and we can write $\tilde{S}(x) = x_1R_1 + \dots + x_kR_k$, where R_i are the rows of S , there exists $x' \in \mathbb{F}_q^k$ such that $\tilde{S}(x') = 0$. Assuming without loss of generality that S consists of the first t columns of \mathcal{G} , there exists a codeword $c = [\tilde{S}(x'), y_{t+1}, \dots, y_n]$ whose weight equals $n - t$. Hence $d \leq n - t$. \square

We now prove Theorem 3.1 and note that our proof differs from other proofs in the literature in that it does not make formal use of an algorithm.

Proof. We know that every set of $r + 1$ columns has rank r by the locality condition. This means that we can choose a set S of $\lfloor \frac{k-1}{r} \rfloor (r + 1) + \{ \frac{r-1}{r} \} r$ columns in such a way that $\text{rank}(S) \leq k - 1$ (here $\{x\} = x - \lfloor x \rfloor$). Thus, by applying Lemma 3.2, we have the following:

$$\left(\left\lfloor \frac{k-1}{r} \right\rfloor + \left\{ \frac{k-1}{r} \right\} \right) r + \left\lfloor \frac{k-1}{r} \right\rfloor = k - 1 + \left\lfloor \frac{k-1}{r} \right\rfloor \leq n - d. \quad \square$$

We also present the following convenient reformulation of the Singleton-like bound which is often used in the literature.

Corollary 3.3. *The parameters of an $[n, k, d, r]$ locally recoverable code \mathcal{C} satisfy*

$$d \leq n - k - \left\lceil \frac{k}{r} \right\rceil + 2. \quad (3.1)$$

The obvious question is whether codes exist (for all, or for almost all, values of r) which satisfy the bound in Corollary 3.3 with equality. Some ad hoc constructions were provided in [34, 38], for example, but these constructions have the problem that the alphabet size grows exponentially in the length of the code, making these codes impractical for most applications. On top of that, these examples are only defined for a few choices of the parameters, most notably the cases $r = 1$ and $r = k$. In the next section, we will explore the first general construction that does not have these shortcomings.

3.3 Tamo-Barg Construction

Recall that a locally recoverable code \mathcal{C} with parameters n, k, r, d is called *optimal* if the parameters satisfy the bound in Corollary 3.3 with equality. In 2014, Itzhak Tamo and Alexander Barg constructed a large family of optimal locally recoverable codes with locality r in [36], and we present this construction here. For simplicity we will assume $r \mid k$ and $(r + 1) \mid n$, but in Section 3.4 we present modified constructions which do not require these conditions. A fundamental part of all the constructions is the existence of a polynomial $g(x) \in \mathbb{F}_q[x]$ satisfying the following conditions:

- The degree of g is $r + 1$, and
- There exists a set $A \subseteq \mathbb{F}_q$ of size n and a partition $\mathcal{A} = \{A_1, \dots, A_{\frac{n}{r+1}}\}$ of A into sets, each of size $r + 1$, such that g is constant on each set A_i in the partition.

We call a polynomial g satisfying the two conditions above a *good* polynomial. Such a polynomial can be used to construct a locally recoverable code as follows.

Construction 3.4. [36, Construction 1] *Let $n \leq q$ be the desired code length. Let $A = \{a_1, \dots, a_n\} \subseteq \mathbb{F}_q$, and let $g(x)$ be a good polynomial for the partition \mathcal{A} of the set A . To obtain the codeword for the message vector $m \in \mathbb{F}_q^k$, first write $m = (m_{i,j})$, where $0 \leq i \leq r-1$ and $0 \leq j \leq \frac{k}{r} - 1$. Define the encoding polynomial associated to m by*

$$f_m(x) = \sum_{i=0}^{r-1} \left(\sum_{j=0}^{\frac{k}{r}-1} m_{i,j} g(x)^j \right) x^i.$$

Now, the codeword $\text{Enc}(m)$ for m is exactly the evaluation vector of f_m at all the points of A . In other words, the $[n, k, r]$ locally recoverable code \mathcal{C} is defined to be the set of vectors

$$\mathcal{C} = \{(f_m(a_1), \dots, f_m(a_n)) : m \in \mathbb{F}_q^k\} \subset \mathbb{F}_q^n.$$

The elements of the set A are called *locations*, and the $f_m(a_1), \dots, f_m(a_n)$ are the *symbols* of the codeword.

In short, a message vector $m \in \mathbb{F}_q^k$ is encoded by first associating m with the encoding polynomial f_m . One then evaluates f_m at each location in A to obtain the codeword $(f_m(a_1), \dots, f_m(a_n))$.

Note that when $r = k$, the Tamo-Barg construction becomes that of the well-known Reed-Solomon error correction codes.

We now describe the recovery process: given the codeword $c \in \mathcal{C}$, assume that the erasure occurs at location $\tilde{i} \in A_{\tilde{j}}$, where $A_{\tilde{j}}$ is a set in the partition \mathcal{A} as described above. To recover

the symbol $c_{\tilde{i}}$, we determine the (unique) polynomial $L(x) \in \mathbb{F}_q[x]$ of degree less than r that satisfies $L(j_\alpha) = c_{j_\alpha}$ for all $j_\alpha \in A_{\tilde{j}} \setminus \{\tilde{i}\}$. This can be done in a simple and straightforward manner using Lagrange interpolation; that is,

$$L(x) = \sum_{j_\alpha \in A_{\tilde{j}} \setminus \{\tilde{i}\}} c_{j_\alpha} \left(\prod_{j'_\alpha \in A_{\tilde{j}} \setminus \{\tilde{i}, j_\alpha\}} \frac{x - j'_\alpha}{j_\alpha - j'_\alpha} \right).$$

We then set $c_i = L(i)$.

To see that this correctly recovers the symbol c_i , observe in particular that since $g(x)$ is constant on $A_{\tilde{j}}$, observe that for any $j_\alpha \in A_{\tilde{j}} \setminus \{\tilde{i}\}$ we have

$$L(\tilde{i}) = \sum_{i=0}^{r-1} \left(\sum_{j=0}^{\frac{k}{r}-1} m_{i,j} g(j_\alpha)^j \right) \tilde{i}^i = \sum_{i=0}^{r-1} \left(\sum_{j=0}^{\frac{k}{r}-1} m_{i,j} g(\tilde{i})^j \right) \tilde{i}^i = f_m(\tilde{i}).$$

Thus the r known symbols in the recovery set $A_j \setminus \{i\}$ can be used to recover the symbol c_i .

Example 3.5. We will use Construction 3.4 to construct a $[10, 8, 4, 2]$ LRC over \mathbb{F}_{11} . First, note that the polynomial $g(x) = x^5$ is constant on each of the sets $A_1 = \{1, 3, 4, 5, 9\}$ and $A_2 = \{2, 6, 7, 8, 10\}$, and these sets clearly partition $A = A_1 \sqcup A_2$. Suppose we wish to encode the message vector $m = (m_{0,0}, m_{0,1}, m_{0,2}, m_{0,3}, m_{1,0}, m_{1,1}, m_{1,2}, m_{1,3}) \in \mathbb{F}_{11}^8$. The encoding polynomial for m is defined to be

$$f_m(x) = (m_{0,0} + m_{0,1}g(x) + m_{0,2}g(x)^2 + m_{0,3}g(x)^3) + (m_{1,0} + m_{1,1}g(x) + m_{1,2}g(x)^2 + m_{1,3}g(x)^3)x,$$

which, for $g(x) = x^5$, simplifies to

$$f_m(x) = m_{0,0} + m_{1,0}x + m_{0,1}x^5 + m_{1,1}x^6 + m_{0,2}x^{10} + m_{1,2}x^{11} + m_{0,3}x^{15} + m_{1,3}x^{16}.$$

Then the codeword $c = \text{Enc}(m)$ is obtained by evaluating $f_m(x)$ at each of the elements in A_1 and A_2 , respectively. That is,

$$c = (f_m(1), f_m(3), f_m(4), f_m(5), f_m(9), f_m(2), f_m(6), f_m(7), f_m(8), f_m(10)).$$

Suppose now that $m = (1, 1, 0, 0, 1, 0, 1, 0)$. Then $f_m = 1 + x + x^5 + x^{11}$, so we compute $c = (4, 8, 10, 1, 9, 4, 1, 3, 5, 9)$. To demonstrate the recovery process, let us assume that the symbol $f_m(2) = 4$ is erased. Since $2 \in A_2$, we will use the set $A_2 \setminus \{2\}$ and the corresponding values of f_m to compute $L(x)$. Explicitly, we have

$$\begin{aligned} L(x) &= \sum_{j_\alpha \in A_2 \setminus \{2\}} c_{j_\alpha} \prod_{j'_\alpha \in A_2 \setminus \{2, j_\alpha\}} \frac{x - j'_\alpha}{j_\alpha - j'_\alpha} \\ &= \frac{(x-7)(x-8)(x-10)}{(6-7)(6-8)(6-10)} + 3 \frac{(x-6)(x-8)(x-10)}{(7-6)(7-8)(7-10)} \\ &\quad + 5 \frac{(x-6)(x-7)(x-10)}{(8-6)(8-7)(8-10)} + 9 \frac{(x-6)(x-7)(x-8)}{(10-6)(10-7)(10-8)} \\ &= 4(x-7)(x-8)(x-10) + (x-6)(x-8)(x-10) \\ &\quad + 7(x-6)(x-7)(x-10) + 10(x-6)(x-7)(x-8). \end{aligned}$$

Finally, we compute $L(2) = 4 = f_m(2)$, so we have recovered the erased symbol.

Now, since the encoding in Construction 3.4 is linear, a lower bound for the distance of the resulting code is given by

$$d \geq n - \max_{m \in \mathbb{F}_q^k} \{\deg(f_m)\}. \quad (3.2)$$

Observe that for any $m \in \mathbb{F}_q^k$ we have

$$\deg(f_m) \leq \left(\frac{k}{r} - 1\right)(r+1) + r - 1 = k + \frac{k}{r} - 2$$

and hence

$$d \geq n - k - \frac{k}{r} + 2.$$

With the Singleton-like bound given in Corollary 3.3, the above shows that these codes are optimal locally recoverable codes. We state this formally in the following theorem.

Theorem 3.6. *[36, Theorem 3.1] The Tamo-Barg codes obtained from Construction 3.4 are optimal locally recoverable codes, i.e., satisfy the Singleton-like bound in Corollary 3.3 with equality.*

We note that the Tamo-Barg construction yields optimal LRCs which are not subject to the same restrictions as previously known locally recoverable codes. In particular, the alphabet size is no longer exponential in the code length, and such a construction can be defined for a large selection of parameters (which satisfy the stated divisibility requirements) so long as one can find a good polynomial $g(x)$. We show how to remove the aforementioned divisibility requirements in Section 3.4, and we completely classify good polynomials up to degree 5 (over an arbitrary finite field) in Chapter 5.

The construction of the Tamo-Barg codes naturally leads to the following questions:

1. Are there simple ways of obtaining many good polynomials, in particular for different parameters?
2. Is it possible to classify all good polynomials of arbitrary degree over \mathbb{F}_q for a fixed prime power q ?

We now present some constructions of good polynomials from the literature. These constructions depend on the desired locality of the code (and hence on the degree of the good polynomial) as well as on the characteristic and size of the field \mathbb{F}_q . As such, for what follows let $q = p^s$ be the size of the underlying field and write $r + 1 = m \cdot p^t$, where $\gcd(m, p) = 1$ and t is a nonnegative integer. (Recall that Construction 3.4 uses a good polynomial of degree $r + 1$ to construct a code with locality r .)

1. If $g(x)$ is a good polynomial, then $\lambda g(x) + \alpha$ is clearly a good polynomial for any $\lambda \in \mathbb{F}_q^*$ and $\alpha \in \mathbb{F}_q$.
2. If $t = 0$, i.e., if $\gcd(r + 1, p) = 1$, then the power function

$$P(x) = x^m$$

is a good polynomial. This follows from the multiplicative structure in \mathbb{F}_q^* , as $P(x)$ is constant on the multiplicative cosets of $\{x \in \mathbb{F}_q^* : P(x) = 1\}$. (See [36, Proposition 3.2].)

3. If $t > 0$ and $m = 1$, then for $a_i \in \mathbb{F}_q$ satisfying $a_0, a_t \neq 0$, the linear function

$$\ell(x) = \sum_{i=0}^t a_i x^{p^i}$$

is a good polynomial so long as $\ell(x)$ splits over \mathbb{F}_q . This follows from the additive structure in \mathbb{F}_q , as $\ell(x)$ is constant on the additive cosets of $\{x \in \mathbb{F}_q : \ell(x) = 0\}$. (See [36, Proposition 3.2].)

4. If $t > 0$, $m > 1$, and $p^s \equiv 1 \pmod{m}$, then for $e \mid t$ such that $p^e \equiv 1 \pmod{m}$ and $a_i \in \mathbb{F}_q$ satisfying $\sum_{i=0}^{t/e} a_i = 0$ and $a_0, a_{t/e} \neq 0$, the function

$$F(x) = \left(\sum_{i=0}^{t/e} a_i x^{p^{ei}} \right)^m$$

is a good polynomial so long as $\sum_{i=0}^{t/e} a_i x^{p^{ei}}$ splits over \mathbb{F}_q . This follows jointly from the multiplicative and additive structures in the previous two examples, as $F(x)$ is constant on the additive cosets of $\{x \in \mathbb{F}_q : F(x) = 0\}$. (See [36, Theorem 3.3].)

5. If $t > 0$, $q = p^s \equiv 1 \pmod{m}$, $t \mid s$, and $s > t(m-1)$, then the function

$$\gamma(x) = x^{p^t} - \alpha^{p^t-1}x$$

is a good polynomial for any $\alpha \in \mathbb{F}_q^*$. (See [26, Corollary 26].)

Remark 3.7. One of the functions $P(x)$ or $\gamma(x)$ above yields a good polynomial for any choice of locality r . If $p \nmid (r+1)$, then $P(x)$ can be used to construct an optimal LRC over \mathbb{F}_q with locality r . On the other hand, if $p \mid (r+1)$, then one can instead use $\gamma(x)$ to construct an optimal LRC over \mathbb{F}_q . In other words, there is a known good polynomial for any choice of locality r .

As mentioned above, as an original part of this dissertation, we give a complete classification of all good polynomials up to degree 5 in Chapter 5.

3.4 Modified Constructions

3.4.1 Optimal LRCs Without $r \mid k$

Recall that Construction 3.4 as stated requires that $r \mid k$ and $(r+1) \mid n$. In this section, we present some modifications which allow for the removal of these constraints. The constructions in this section were originally detailed in [36], but we adapt the notation to our setting.

The assumption $r \mid k$ can be removed without any effect on the optimality of the code. This construction is very similar to that of Construction 3.4; the difference between the two can be seen in the way the encoding polynomial $f_m(x)$ is defined. Recall that in Construction 3.4 we have

$$f_m(x) = \sum_{i=0}^{r-1} \left(\sum_{j=0}^{\frac{k}{r}-1} m_{i,j} g(x)^j \right) x^i.$$

If $r \nmid k$ the internal sum is clearly no longer defined, so for this case the authors of [36] redefine $f_m(x)$ by

$$f_m(x) = \sum_{i=0}^{r-1} \left(\sum_{j=0}^{s(k,r,i)} m_{i,j} g(x)^j \right) x^i, \quad \text{where } s(k,r,i) = \begin{cases} \lfloor \frac{k}{r} \rfloor & \text{if } i < k \pmod{r}, \\ \lfloor \frac{k}{r} \rfloor - 1 & \text{if } i \geq k \pmod{r}. \end{cases}$$

The dimension of the resulting code is seen to be k by a nearly identical argument as for Construction 3.4. The maximum possible degree of $f_m(x)$ with this definition is attained when $k \equiv r - 1 \pmod{r}$, and the particular term from the sum which attains this degree corresponds to $i = r - 2 < k \pmod{r}$ and $j = s(k,r,i) = \lfloor \frac{k}{r} \rfloor$. Observe that the degree of this term is exactly

$$\begin{aligned} \lfloor \frac{k}{r} \rfloor (r+1) + (r-2) &= k - r + 1 + \lfloor \frac{k}{r} \rfloor + r - 2 \\ &= k + \lfloor \frac{k}{r} \rfloor - 2, \end{aligned}$$

and this agrees with the maximum degree of $f_m(x)$ in the original definition given in Construction 3.4. Thus Construction 3.4 can easily be adjusted to yield an optimal LRC in the case $r \nmid k$ without any effect on the optimality of the code.

3.4.2 Almost Optimal LRCs Without $(r+1) \mid n$

In the previous subsection we saw that the condition $r \mid k$ was not necessary for Construction 3.4 to produce an optimal LRC. This is not exactly the case for the condition $(r+1) \mid n$. Before we address the issue of optimality, we first present the construction.

Construction 3.8. [36, Construction 5, Construction 6] *Let $A \subseteq \mathbb{F}_q$ be a subset such that $|A| = n$ and write $n \pmod{r+1} = s \neq 0, 1$. Let $N = \lceil \frac{n}{r+1} \rceil$ and assume that $(r+1) \mid k$ (this assumption is not necessary but allows for simpler notation). Let $\mathcal{A} = \{A_1, \dots, A_N\}$ be a partition of A such that $|A_i| = r+1$ for $1 \leq i \leq N-1$ and $|A_N| = s < r+1$, and let*

$g(x)$ be a polynomial which is constant on each A_i . Assume without loss of generality that g vanishes on A_N (since otherwise we can use $g'(x) = g(x) - g(A_N)$). To obtain the codeword for the message vector $m \in \mathbb{F}_q^k$, first write $m = (m_{i,j})$, where $0 \leq i \leq r-1$ and the bounds for j are as follows: for $i \neq s-1$ we have $0 \leq j \leq \frac{k+1}{r} - 1$ and for $i = s-1$ we have $1 \leq j \leq \frac{k+1}{r} - 1$. Define the encoding polynomial associated to m by

$$f_m(x) = \sum_{i=0}^{s-2} \left(\sum_{j=0}^{\frac{k+1}{r}-1} m_{i,j} g(x)^j \right) x^i + \sum_{j=1}^{\frac{k+1}{r}-1} m_{s-1,j} g(x)^j x^{s-1} + h(x) \sum_{i=s}^{r-1} \left(\sum_{j=0}^{\frac{k+1}{r}-1} m_{i,j} g(x)^j \right) x^{i-s},$$

where $h(x) = \prod_{\alpha \in A_N} (x - \alpha) \in \mathbb{F}_q[x]$ is the annihilator of the set A_N . As in Construction 3.4, the $[n, k, r]$ locally recoverable code \mathcal{C} is now defined to be the set of vectors

$$\mathcal{C} = \{(f_m(a_1), \dots, f_m(a_n)) : m \in \mathbb{F}_q^k\} \subset \mathbb{F}_q^n.$$

In this case it is not as obvious that the dimension is k , so we begin by showing that if $d_{i,j}$ is the degree of the polynomial obtained from the pair (i, j) , then $d_{i,j} = d_{i',j'}$ implies $(i, j) = (i', j')$. For $0 \leq i \leq s-2$ we have $d_{i,j} = j \cdot \deg(g) + i = j(r+1) + i$, so $d_{i,j} \pmod{(r+1)} = i \leq s-1$. For $i = s-1$ we have $d_{s-1,j} = j(r+1) + s-1$, so $d_{s-1,j} \equiv s-1 \pmod{(r+1)}$. Finally, for $s \leq i \leq r-1$ we have $d_{i,j} = \deg(h) + j(r+1) + i - s = j(r+1) + i$, so once again we have $d_{i,j} \equiv i \pmod{(r+1)}$. Now, if $d_{i,j} = d_{i',j'}$ then $d_{i,j} \pmod{(r+1)} = i = i' = d_{i',j'} \pmod{(r+1)}$ so that $i = i'$. From this it follows that $j(r+1) = j'(r+1)$ so that $j = j'$ and hence $(i, j) = (i', j')$. Since these polynomials are all of distinct degrees, we have

$$\begin{aligned} \dim_{\mathbb{F}_q} \mathcal{C} &= (s-1) \frac{k+1}{r} + \left(\frac{k+1}{r} - 1 \right) + (r-1-s+1) \frac{k+1}{r} \\ &= r \frac{k+1}{r} - 1 \\ &= k. \end{aligned}$$

Next we consider the maximum possible degree of $f_m(x)$. Clearly $\deg(f_m) \leq \max_{(i,j)} \{d_{i,j}\}$. Further, in the above we saw that $d_{i,j} = j(r+1)+i$ for any pair (i, j) , so we have the following:

$$\begin{aligned} \deg(f_m) &\leq \max_{(i,j)} \{j(r+1) + i\} \\ &= \left(\frac{k+1}{r} - 1 \right) (r+1) + (r-1) \\ &= k + \frac{k+1}{r} - 1 \\ &= k + \left\lceil \frac{k}{r} \right\rceil - 1. \end{aligned}$$

Combining the above with Equation 3.2 yields the following theorem:

Theorem 3.9. [36, Theorem 5.2] *The code obtained in Construction 3.8 is an $[n, k, r]$ LRC whose minimum distance satisfies*

$$d \geq n - k - \left\lceil \frac{k}{r} \right\rceil + 1.$$

We note that the recovery process is still quite similar to that of Construction 3.4. Since $g(x)$ is constant on A_i for each i , the degree of the restriction $f_m|_{A_i}$ of f_m to the set A_i for $0 \leq i \leq N-1$ satisfies

$$\deg(f_m|_{A_i}) = \max\{s-2, s-1, r-1\} = r-1,$$

so recovery can be done exactly as in Construction 3.4. For $i = N$, however, only the first of the three summations in f_m is not identically 0, so

$$\deg(f_m|_{A_N}) = s-2.$$

If the symbol $f_m(\beta)$ is erased for some $\beta \in A_N$, then we need to interpolate a polynomial of degree at most $s-2$ on the set $A_N \setminus \{\beta\}$, which consists of $|A_N| - 1 = s-1$ locations. Thus recovery is possible in this case as well since $g(x)$ is constant on A_N .

Finally, notice that the bound in Theorem 3.9 does not match the bound given in Corollary 3.3, i.e., the codes obtained from Construction 3.8 are not optimal LRCs. However, the distance of these codes differs from being optimal by at most 1, so they are often referred to as being *almost optimal* LRCs. Moreover, it was shown in [21, Corollary 10] that no $[n, k, r]$ LRC satisfies the Singleton-like bound with equality when $2 < d < r + 3$ and $r \mid k$. Equivalently, there are no (n, k, r) LRCs with $r \mid k$ and $d = n - k - \frac{k}{r} + 2$ which satisfy $\frac{k}{r} < \frac{n}{r+1} < \frac{k}{r} + 1$. This is not restrictive when $(r + 1) \mid n$ since these inequalities clearly cannot both be satisfied. So we conclude that any $[n, k, r]$ LRC satisfying $r \mid k$, $(r + 1) \nmid n$, and $\frac{k}{r} < \frac{n}{r+1} < \frac{k}{r} + 1$ is subject to $d < n - k - \frac{k}{r} + 2$, i.e., $d \leq n - k - \frac{k}{r} + 1$. Thus the codes obtained from Construction 3.8 are optimal in some sense for at least one regime of parameters n, k, r .

We end this chapter by recalling that each of the aforementioned constructions relies entirely on the existence of a good polynomial of degree $r + 1$. For any choice of locality r we have explicitly provided a good polynomial which allows for the construction of an $[n, k, r]$ LRC over some field \mathbb{F}_q with $q \geq n$, and we detailed multiple constructions in which the corresponding LRCs are proved to be optimal (or at least almost optimal when $(r + 1) \nmid n$). We next cover the background in algebraic number theory which is needed for the optimal selection results in Chapter 5 and the existential results in Chapter 6.

CHAPTER 4:
ALGEBRAIC NUMBER THEORY AND CHEBOTAREV DENSITY
THEOREM

4.1 Preliminaries

For the sake of being self-contained, in this section we define the terms we will use most often in the rest of this dissertation. The field-theoretical definitions in Subsection 4.1.1 can be found in any graduate-level textbook on abstract algebra (see [22], for example). The notions in Subsection 4.2 closely follow the notation and terminology in [35].

4.1.1 Field Theory

For q a power of a prime, let \mathbb{F}_q be the finite field with q elements and let $\mathbb{F}_q^* = \mathbb{F}_q \setminus \{0\}$ be its multiplicative subgroup. It is known that this subgroup is cyclic, that is, generated by a single element for any q .

For a field K , we will write $K[X]$ to denote the polynomial ring in the indeterminate X over K .

Recall that if $K \subseteq F$ is a field extension, then F forms a vector space over K . If $\dim_K F < \infty$, then the degree $[F : K]$ of the extension is defined to be $[F : K] = \dim_K F$.

Definition 4.1. Given an extension of fields $K \subseteq F$, the group of all automorphisms of F/K is denoted by $\text{Aut}(F/K)$. When $|\text{Aut}(F/K)| = [F : K] < \infty$, we say that the extension F/K is Galois with Galois group $\text{Gal}(F/K) = \text{Aut}(F/K)$.

Definition 4.2. For a polynomial $f \in K[X]$ of positive degree, an extension $F \supseteq K$ is said to be the splitting field of f over K if f splits into linear factors in $F[X]$ and $F = K(x_1, x_2, \dots, x_n)$, where the x_i are the roots of f in F .

Definition 4.3. Given a field K and a polynomial $f \in K[X]$, we define the Galois group $\text{Gal}(f(X) | K)$ of f over K to be the group $\text{Aut}(F/K)$, where F is a splitting field of f over K .

Definition 4.4. Given a field F , an algebraic closure \overline{F} of F is a field extension of F which is algebraically closed. In other words, if $f \in F[X]$ is any polynomial over F , then f splits into linear factors in \overline{F} .

4.1.2 Algebraic Number Theory

Definition 4.5. A polynomial $f \in F[X]$ is said to be separable over F if its roots are distinct in an algebraic closure \overline{F} of F .

We note that when $F = \mathbb{F}_q$, Definition 4.5 allows us to write the following: an *irreducible* polynomial $f \in \mathbb{F}_q[X]$ is separable if $f \notin \mathbb{F}_q[X^p]$, where $p = \text{char } \mathbb{F}_q$.

Definition 4.6. Let t be transcendental over \mathbb{F}_q . We will denote by $\mathbb{F}_q(t)$ the field of rational functions in t over \mathbb{F}_q . A finite-dimensional field extension F of $\mathbb{F}_q(t)$ is called a (global) function field over \mathbb{F}_q .

Note immediately that if $f \in \mathbb{F}_q[X]$ is separable over \mathbb{F}_q , then $f - t$ is a separable and irreducible polynomial over $\mathbb{F}_q(t)$.

If M is defined to be the splitting field of $f(X) - t$ over $\mathbb{F}_q(t)$, then M can equivalently be defined as the Galois closure of the extension $\mathbb{F}_q(x)/\mathbb{F}_q(t)$, where x is any root of $f(X) - t$ in the algebraic closure $\overline{\mathbb{F}_q(t)}$ of $\mathbb{F}_q(t)$.

Definition 4.7. Given a function field $F \supseteq \mathbb{F}_q(t)$, the field of constants k_F of F consists of the elements in F which are algebraic over \mathbb{F}_q . We may more simply denote k_F by k if there is no risk of confusion.

In the definition above, we clearly have $\mathbb{F}_q \subseteq k$ for any choice of F . However, we note that it is possible to have $\mathbb{F}_q \subsetneq k$. We provide a brief example below.

Example 4.8. Let $q \equiv 2 \pmod{3}$ and $f(X) = X^3 \in \mathbb{F}_q[X]$ and consider the field of constants k of M , where M is a splitting field of $f(X) - t$ over \mathbb{F}_q . One can easily see that $[M : \mathbb{F}_q(t)] = 6$ and that $M \cong \mathbb{F}_q(t)(\eta, \sqrt[3]{t}) = \mathbb{F}_q(\eta, \sqrt[3]{t})$ for some element $\eta \in \overline{\mathbb{F}_q} \setminus \mathbb{F}_q$ satisfying $\eta^3 = 1$. We have $\eta \notin \mathbb{F}_q$ since $q - 1 \equiv 1 \pmod{3}$ and hence $3 \nmid (q - 1)$; moreover, we see that actually $\eta \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$ since $q^2 - 1 \equiv 0 \pmod{3}$. Thus $[\mathbb{F}_q(\eta, t) : \mathbb{F}_q(\eta, t)] = [\mathbb{F}_q(\eta) : \mathbb{F}_q] = 2$, so $[k : \mathbb{F}_q] \geq 2$. Further, since we clearly have $[\mathbb{F}_q(\eta, \sqrt[3]{t}) : \mathbb{F}_q(t)] = 3 = 6/2$, it follows that $[k : \mathbb{F}_q] = 2$. Hence $k = \mathbb{F}_{q^2} \supsetneq \mathbb{F}_q$.

Definition 4.9. A valuation ring of a function field F/K is a ring \mathcal{O} such that $K \subsetneq \mathcal{O} \subsetneq F$ and which contains at least one of z or z^{-1} for every $z \in F$.

Definition 4.10. A place P of F/K is the unique maximal ideal of some valuation ring \mathcal{O} of F/K . We will write \mathbb{P}_F to denote the set of all places of F/K . An element $t \in P$ such that $P = t\mathcal{O}$ is called a prime element for P .

It turns out that each place $P \in \mathbb{P}_F$ is actually the maximal ideal of a *unique* valuation ring \mathcal{O} of F/K , namely the ring $\mathcal{O}_P = \{z \in F : z^{-1} \notin P\}$. Hence there is a one-to-one correspondence between the places of F/K and the valuation rings \mathcal{O} of F/K , and so we will sometimes conveniently write \mathcal{O}_P to denote the valuation ring whose maximal ideal is P .

We can also examine places and valuation rings via valuations, which are defined as follows.

Definition 4.11. A (discrete) valuation of a function field F/K is a function $v : F \rightarrow \mathbb{Z} \cup \{\infty\}$ satisfying the following:

- (a) $v(x) = \infty$ if and only if $x = 0$,
- (b) $v(xy) = v(x) + v(y)$ for all $x, y \in F$,

- (c) $v(x + y) \geq \min\{v(x), v(y)\}$ for all $x, y \in F$.
- (d) there exists an element $z \in F$ with $v(z) = 1$, and
- (e) $v(a) = 0$ for all $0 \neq a \in K$.

To connect these functions to places and valuation rings, we first need the following theorem.

Theorem 4.12 ([35, Theorem 1.1.6]). *Let \mathcal{O} be a valuation ring of the function field F/K and let $P \in \mathbb{P}_F$ be its maximal ideal. Then \mathcal{O} is a discrete valuation ring. In other words, the following hold:*

- (a) P is a principal ideal, that is, $P = t\mathcal{O}$ for some $t \in \mathcal{O}$.
- (b) If $P = t\mathcal{O}$, then each $0 \neq z \in F$ has a unique representation of the form $z = t^n u$ for some $n \in \mathbb{Z}$ and some unit $u \in \mathcal{O}^\times$.
- (c) \mathcal{O} is a principal ideal domain. More precisely, if $P = t\mathcal{O}$ and $\{0\} \neq I \subseteq \mathcal{O}$ is an ideal, then $I = t^n \mathcal{O}$ for some $n \in \mathbb{N}$.

Now, suppose we are given a place P and/or its corresponding valuation ring \mathcal{O}_P . By the above theorem, we can find an element $t \in P$ such that every element $z \in F$ can be written as $z = t^n u$ for some $n \in \mathbb{N}$ and unit $u \in \mathcal{O}^\times$. Define the function $v_P : F \rightarrow \mathbb{N} \cup \{\infty\}$ given by $v_P(z) = v_P(t^n u) = n$ for $z \neq 0$ and $v_P(0) = \infty$. Then v_P is a discrete valuation of F/K . We will revisit these functions in Appendix A.

If we consider the function field $K(X)/K$, where X is an indeterminate over K , then a more explicit representation of the places and valuation rings of $K(X)/K$ can be stated as follows. For an irreducible monic polynomial $p(X) \in K[X]$, the ring

$$\mathcal{O}_{p(X)} = \left\{ \frac{f(X)}{g(X)} \mid f(X), g(X) \in K[X] \text{ coprime, } p(X) \nmid g(X) \right\} \quad (4.1)$$

is a valuation ring of $K(X)/K$, and its maximal ideal $P_{p(X)}$ is given by

$$P_{p(X)} = \left\{ \frac{f(X)}{g(X)} \mid f(X), g(X) \in K[X] \text{ coprime, } p(X) \mid f(X), p(X) \nmid g(X) \right\}.$$

In fact, all valuation rings of $K(X)/K$ *but one* are of the form 4.1. The exceptional valuation ring is given by

$$\mathcal{O}_\infty = \left\{ \frac{f(X)}{g(X)} \mid f(X), g(X) \in K[X] \text{ coprime, } \deg(f) \leq \deg(g) \right\}. \quad (4.2)$$

The maximal ideal $P_\infty \subseteq \mathcal{O}_\infty$, which is given by

$$P_\infty = \left\{ \frac{f(X)}{g(X)} \mid f(X), g(X) \in K[X] \text{ coprime, } \deg(f) < \deg(g) \right\},$$

is called the infinite place, or the place at infinity, of $K(X)/K$.

Definition 4.13. Let F/K be a function field, and let $P \in \mathbb{P}_F$.

- (a) $F_P = \mathcal{O}_P/P$ is called the residue class field of P , and the map $x \mapsto x(P)$ from F to $F_P \cup \{\infty\}$ is called the residue class map with respect to P .
- (b) $\deg(P) = [F_P : K]$ is called the degree of P . A place of degree one is called a rational place of F/K , and the set of all rational places of F/K will be denoted by $\mathbb{P}_F^1 \subseteq \mathbb{P}_F$.

Now, the above definition requires some clarification. Since the place P is the maximal ideal of \mathcal{O}_P , the quotient \mathcal{O}_P/P (often called a residue class ring) forms a field. Then the residue class map induces a (ring) homomorphism from K into F_P since $K \subseteq \mathcal{O}_P$ by definition. Moreover, we must have $K \cap P = \{0\}$ since otherwise P contains an element which is invertible in \mathcal{O}_P and hence P is equal to \mathcal{O}_P , which is a contradiction. Thus the residue class map actually embeds K in F_P , so we can treat K as a subfield of F_P under this embedding, and so the quantity $[F_P : K]$ is defined.

Definition 4.14. For places $P \in \mathbb{P}_K$ and $Q \in \mathbb{P}_{F'}$, we say that Q lies over P (and write $Q | P$) if $P \subseteq Q$. We may equivalently say that Q divides P or that $Q | P$ is an extension of places.

Definition 4.15. Let F'/K' be an algebraic extension of F/K , and let $Q \in \mathbb{P}_{F'}$ be a place of F'/K' lying over $P \in \mathbb{P}_F$.

- (a) The integer $e(Q|P) = e$ satisfying $v_Q(x) = e \cdot v_P(x)$ for all $x \in F$ is called the ramification index of Q over P . We say that $Q | P$ is ramified if $e(Q|P) > 1$, and $Q | P$ is unramified if $e(Q|P) = 1$.
- (b) $f(Q|P) = [F'_Q : F_P]$ is called the relative degree of Q over P .

The next theorem, commonly called the Fundamental Equality, relates the quantities defined above.

Theorem 4.16 ([35, Theorem 3.1.11]). *Let F'/K' be a finite extension of F/K , let P be a place of F/K , and let Q_1, \dots, Q_m be all the places of F'/K' lying over P . Then we have*

$$\sum_{i=1}^m e(Q_i|P) f(Q_i|P) = [F' : F].$$

In the above, if we know further that F'/F is a Galois extension, then $e(Q_i|P) = e(P)$ and $f(Q_i|P) = f(P)$, i.e., the ramification index and relative degree do not depend on the index i . In this case, the above becomes

$$\sum_{i=1}^m e(Q_i|P) f(Q_i|P) = m \cdot e(P) \cdot f(P) = [F' : F],$$

and this shows that each of $e(P)$, $f(P)$, and m divides $[F' : F]$ when F'/F is Galois.

Definition 4.17. Let F'/F be an algebraic extension of functions fields and let $P \in \mathbb{P}_F$.

- (a) An extension Q of P in F' is said to be tamely (resp. wildly) ramified if $e(Q|P) > 1$ and the characteristic of K does not divide $e(Q|P)$ (resp. $\text{char } K$ divides $e(Q|P)$).

- (b) We say that P is ramified (resp. unramified) in F'/F if there is at least one $Q \in \mathbb{P}_{F'}$ lying over P such that $Q | P$ is ramified (resp. if $Q | P$ is unramified for all $Q | P$). The place P is tamely ramified in F'/F if it is ramified in F'/F and no extension of P in F' is wildly ramified. If there is at least one wildly ramified place $Q | P$ we say that P is wildly ramified in F'/F .
- (c) P is totally ramified in F'/F if there is only one extension $Q \in \mathbb{P}_{F'}$ of P in F' , and the ramification index is $e(Q|P) = [F' : F]$.
- (d) F'/F is said to be ramified (resp. unramified) if at least one $P \in \mathbb{P}_F$ is ramified in F'/F (resp. if all $P \in \mathbb{P}_F$ are unramified in F'/F).
- (e) F'/F is said to be tame if no place $P \in \mathbb{P}_F$ is wildly ramified in F'/F .

Theorem 4.18 (Abhyankar's Lemma, [35, Theorem 3.9.1]). *Let F'/F be a finite separable extension of function fields. Suppose that $F' = F_1F_2$ is the compositum of two intermediate fields $F \subseteq F_1, F_2 \subseteq F'$. Let $R \in \mathbb{P}_{F'}$ be an extension of $P \in \mathbb{P}_F$, and set $Q_i = R \cap F_i$ for $i = 1, 2$. Assume that at least one of the extensions $Q_1 | P$ or $Q_2 | P$ is tame. Then*

$$e(R|P) = \text{lcm}\{e(Q_1|P), e(Q_2|P)\}.$$

Definition 4.19. Let F'/F be a Galois extension of algebraic function fields with Galois group $G = \text{Gal}(F'/F)$. Let P be a place of F and let Q be an extension of P to F' .

- (a) The decomposition group $D(Q|P)$ of Q over P is defined to be

$$D(Q|P) = \{\sigma \in G : \sigma(Q) = Q\}.$$

- (b) The inertia group $I(Q|P)$ of Q over P is defined to be

$$I(Q|P) = \{\sigma \in G : v_Q(\sigma(z) - z) > 0 \text{ for all } z \in \mathcal{O}_Q\}.$$

The final definition we need is that of the genus of a function field. This requires a number of tools outside the scope of this dissertation, so we defer formally defining this quantity to Appendix A. For now we will simply mention that for a given function field F/K , the genus (denoted g_F) is the most essential invariant of this function field.

Theorem 4.20 ([35, Proposition 1.6.3]). *If F/K is a rational function field, that is, if $F = K(x)$ for some x which is transcendental over the field K , then F/K has genus 0.*

We present the next theorem in a form which will be especially useful in the rest of this dissertation.

Theorem 4.21 (Riemann-Hurwitz Formula, [35, Theorem 3.4.13]). *Let F/K be an algebraic function field of genus g_F and let F'/F be a finite separable extension. Let K' denote the constant field of F' and $g_{F'}$ the genus of F'/K' , respectively. Then we have*

$$2g_{F'} - 2 = \frac{[F' : F]}{[K' : K]}(2g_F - 2) + \sum_{P \in \mathbb{P}_F} \sum_{\substack{Q \in \mathbb{P}_{F'} \\ Q|P}} (e(Q|P) - 1).$$

Theorem 4.22 (Castelnuovo's Inequality, [35, Theorem 3.11.3]). *Let F/K be a function field with constant field K . Suppose there are given two subfields F_1/K and F_2/K of F/K satisfying*

- (1) $F = F_1F_2$ is the compositum of F_1 and F_2 , and
- (2) $[F : F_i] = n_i$ and F_i/K has genus g_i for $i = 1, 2$.

Then the genus g of F/K is bounded by

$$g \leq n_1g_1 + n_2g_2 + (n_1 - 1)(n_2 - 1).$$

4.2 Galois Theory and Chebotarev Density Theorem

A significant portion of the content in this section is based on the published extended abstract “Understanding Polynomial Maps over Finite Fields,” written by the author of this dissertation together with Giacomo Micheli (see [12]). We also adapt some parts of [11] and [13] so that the framework set up in this section can be easily applied to the following chapters of this dissertation.

4.2.1 Polynomial Maps

In this section we explain how to use algebraic number theory to study polynomial maps over finite fields, which occur virtually everywhere in cryptography and coding theory (such as in APN functions, Reed-Solomon codes, locally recoverable codes, etc.). The method we will discuss, which is based on techniques from Galois theory and algebraic geometry, has been a particularly useful tool in these areas in recent literature.

Our method was initially used for constructing locally recoverable codes (LRCs) where the known constructions did not work (see [29] for details). This method has also been used to classify functions with low differential uniformity, such as *perfect nonlinear* functions (PN) and *almost perfect nonlinear* functions (APN), which have been studied extensively (see [4] and [28], for example) for their applications in cryptography. Much work has been dedicated to classifying such functions, and nonexistence results for some exceptional monomial PN and APN (and their recent generalizations to PcN and APcN) were obtained in [4] using similar ideas to the ones here. We begin by summarizing the context in which our method finds relevance and then describing the method.

Let $q = p^r$ be a power of a prime and let \mathbb{F}_q be the finite field of order q . For any map f from \mathbb{F}_q to \mathbb{F}_q , by using Lagrange interpolation we can write

$$f(x) = \sum_{a \in \mathbb{F}_q} f(a) \left(\prod_{b \in \mathbb{F}_q \setminus \{a\}} \frac{x - b}{a - b} \right).$$

Since f is completely determined by the values $f(a)$ as a runs through all elements of \mathbb{F}_q , the above shows that any map from \mathbb{F}_q to \mathbb{F}_q can be identified uniquely with a polynomial of degree at most $q - 1$. Hence we need only consider polynomial maps of bounded degree in this framework.

Where algebraic number theory most often finds application, however, is when considering a polynomial $f(X) \in \mathbb{F}_q[X]$ of degree n in the regime $n \ll q$. With this context in mind, we now describe the method we use to study f as a map.

Let $f \in \mathbb{F}_q[X]$ be a polynomial of degree n . Let $\mathbb{F}_q(t)$ be the field of rational functions in the transcendental t , and denote by M the splitting field of $f(X) - t$ over $\mathbb{F}_q(t)$. Many of the properties of the (polynomial) map $f : \mathbb{F}_q \rightarrow \mathbb{F}_q$ are encoded in the Galois group $A := \text{Gal}(M/\mathbb{F}_q(t))$. In particular, analyzing the cycle structures of particular elements of A allows us to obtain asymptotic estimates on the number of $t_0 \in \mathbb{F}_q$ such that $f(X) - t_0$ splits in some desired way.

4.2.2 Chebotarev Density Theorem

In what follows, let F/\mathbb{F}_q be a global function field with full constant field \mathbb{F}_q , and let M/F be a finite Galois extension with Galois group $A := \text{Gal}(M/F)$. Let $k = \overline{\mathbb{F}_q} \cap M$ be the field of constants of M . Notice that $\text{Gal}(kF/F) \cong \text{Gal}(k/\mathbb{F}_q) \cong \mathbb{Z}/[k:\mathbb{F}_q]\mathbb{Z}$, where the final isomorphism follows from the fact that every finite extension of a finite field is cyclic. Let $\gamma \in A$ be such that $\phi = \gamma|_k$ is the Frobenius automorphism of the extension k/\mathbb{F}_q . Let $R \subseteq M$ be a place of degree 1 of M lying above a place $P \subseteq F$. Let \mathcal{O}_P and \mathcal{O}_R be the valuation rings of P and of R , respectively. Write $G := \text{Gal}(M/kF)$. It is known that $D(R|P)/I(R|P) \cong \text{Gal}(\mathcal{O}_R/\mathcal{O}_P)$ (for a proof, see [35, Theorem 3.8.2(c)]), and we define $D_\phi(R|P)$ to be the coset in $D(R|P)$ of elements which are mapped to ϕ via this isomorphism. For an element σ in the coset $G\gamma$, let Γ_σ be the conjugacy class of σ in A and notice that since A/G is cyclic we have $\Gamma_\sigma \subseteq G\gamma$. Further, recall that for an element $\sigma \in A$, we say σ is a Frobenius at the place P if there exists R lying above P such that $\sigma \in D(R|P)$ and

the induced map of σ in $\text{Gal}(\mathcal{O}_R/R/\mathcal{O}_P/P)$ is $x \mapsto x^{q^{\deg(P)}}$. Finally, for $\sigma \in G\gamma$ we define the quantity

$$w_P(\sigma) = \frac{|D_\phi(R|P) \cap \Gamma_\sigma|}{|D_\phi(R|P)| \cdot |\Gamma_\sigma|}.$$

We give the definition of the genus g_M of the extension M/F in Appendix A.

We are now ready to state the Chebotarev Density Theorem.

Theorem 4.23 (Chebotarev Density Theorem). *Let F/\mathbb{F}_q be a global function field whose constant field is exactly \mathbb{F}_q . Let M/F be a finite Galois extension and define $k = \overline{\mathbb{F}}_q \cap M$ to be the field of constants of M . Let $G = \text{Gal}(M/kF)$ be the geometric Galois group of M/F , g_M be the genus of M , and $\sigma \in G\gamma$. Define $w_P(\sigma)$ as above. Then*

$$\left| \sum_{\substack{P \in \mathbb{P}^1(F/\mathbb{F}_q) \\ P \text{ ramified}}} w_P(\sigma) + \sum_{\substack{P \in \mathbb{P}^1(F/\mathbb{F}_q) \\ P \text{ unramified} \\ \sigma \text{ is a Frobenius at } P}} \frac{1}{|\Gamma_\sigma|} - \frac{1}{|G|}(q+1) \right| \leq \frac{2}{|G|} g_M \sqrt{q},$$

where $\mathbb{P}^1(F/\mathbb{F}_q)$ is the set of places of degree 1 of F .

A particularly convenient corollary obtained from the above theorem (by ignoring the first sum and multiplying throughout by $|\Gamma_\sigma|$) is as follows:

Corollary 4.24. *The number of places $P \in \mathbb{P}^1(F/\mathbb{F}_q)$ such that σ is a Frobenius at P is $\frac{|\Gamma_\sigma|}{|G|}(q+1) + O(\sqrt{q})$, where the implied constant can be chosen independently of q .*

For more information regarding Theorem 4.23 and Corollary 4.24, see [15], and for a full exposition and proof of Theorem 4.23, see [24]. We now state the theorem which is the key to obtaining the asymptotic estimate mentioned above.

Theorem 4.25. *Let $f(X) \in \mathbb{F}_q[X]$ be a polynomial of degree n with geometric Galois group G . Let ℓ be a positive integer and d_1, \dots, d_ℓ be positive integers such that $\sum_{i=1}^\ell d_i = n$. Then the number of $t_0 \in \mathbb{F}_q$ such that there exist some distinct irreducible polynomials $p_i \in \mathbb{F}_q[X]$*

(depending on t_0) with $f(X) - t_0 = \prod_{i=1}^{\ell} p_i(X)$ and $\deg(p_i) = d_i$ is $(|S|/|G|) \cdot q + O(\sqrt{q})$, where S is the subset of elements of G having cycle decomposition

$$\underbrace{(- \cdots -)}_{d_1} \underbrace{(- \cdots -)}_{d_2} \cdots \underbrace{(- \cdots -)}_{d_\ell}.$$

We temporarily defer the proof of this theorem as we wish to first demonstrate how we can employ it in some concrete applications.

Let $q = 100003$ and let $f \in \mathbb{F}_q[X]$ be a polynomial of degree 4. Suppose we are interested in finding the number T of elements $t_0 \in \mathbb{F}_q$ for which $f(X)$ has exactly four preimages as a map from \mathbb{F}_q to \mathbb{F}_q . Notice that this is the same as the number T of elements $t_0 \in \mathbb{F}_q$ such that $f(X) - t_0$ has four zeroes. We first compute the Galois group $A = \text{Gal}(f(X) - t \mid \mathbb{F}_q(t))$ and suppose that the splitting field M of $f(X) - t$ over $\mathbb{F}_q(t)$ has field of constants $k = \mathbb{F}_q$. For the sake of simplicity of notation (and since it is the generic case), we assume $A = G = S_4$.

Clearly the number of elements t_0 in \mathbb{F}_q such that $f(X) - t_0$ has four zeros is the same as the number of t_0 such that $f(X) - t_0 = (X - a)(X - b)(X - c)(X - d)$. Since the identity element of A is the only element with four fixed points, we see that $|S| = 1$. Therefore Theorem 4.25 gives that the number T of t_0 having 4 preimages is roughly $100003/24 \approx 4167$.

As a more elaborate example, suppose we are in the situation above but are now interested in finding the number T' of $t_0 \in \mathbb{F}_q$ for which $f(X) - t_0$ has exactly two zeros. We compute $A = \text{Gal}(f(X) - t \mid \mathbb{F}_q(t))$ and verify that the splitting field of $f(X) - t$ has field of constants \mathbb{F}_q . Again, we assume $A = G = S_4$. Note immediately that the number T' is the same as the number of t_0 such that $f(X) - t_0$ factors over \mathbb{F}_q as $f(X) - t_0 = (X - a)(X - b)g(X)$ for some irreducible polynomial $g \in \mathbb{F}_q[X]$ of degree 2. Let $S' = \{(1, 2), (1, 3), (1, 4), (2, 3), (2, 4), (3, 4)\} \subseteq A$ be the set of elements in A fixing exactly 2 points. We see that $|S'| = 6$, so Theorem 4.25 gives that the number T' of t_0 having exactly 2 preimages is roughly $100003/6 \approx 16667$.

We now prove Theorem 4.25.

Proof of Theorem 4.25. Suppose $f \in \mathbb{F}_q[X]$ is a polynomial of degree n , let t be transcendental over \mathbb{F}_q , and fix natural numbers $d_1, d_2, \dots, d_\ell \geq 1$ such that $\sum_{i=1}^\ell d_i = n$. Define \mathcal{T} to be the set of $t_0 \in \mathbb{F}_q$ for which the polynomial $f(X) - t_0$ can be written in the form

$$f(X) - t_0 = \prod_{i=1}^{\ell} p_{t_0,i}(x),$$

where the $p_{t_0,i} \in \mathbb{F}_q[X]$ are distinct irreducible polynomials satisfying $\deg(p_{t_0,i}) = d_i$. We wish to count the size of \mathcal{T} , so we let $T = |\mathcal{T}|$. Write $F = \mathbb{F}_q$ and observe that $t_0 \leftrightarrow P_{t_0} = \langle t - t_0 \rangle$ is a one-to-one correspondence between the elements $t_0 \in \mathbb{F}_q$ and the places $P_{t_0} \subset F$ of degree 1. Hence the number of places P_{t_0} such that $f(X) - t_0$ admits the desired factorization is exactly T , and $P_{t_0} \in \mathbb{P}^1(F/\mathbb{F}_q)$ for each t_0 , where $\mathbb{P}^1(F/\mathbb{F}_q)$ is the set of places of degree 1 of F . Notice that the t_0 's for which $f(X) - t_0$ has a multiple root are only $O(1)$, as they are at most the zeros of the discriminant $\Delta(t)$ of $f - t$, so we may restrict to the unramified places.

Next, fix a place $P := P_{t_1} \subset F$ for some $t_1 \in \mathbb{F}_q$ as described above. Let M be the splitting field of $f(X) - t$ over $\mathbb{F}_q(t)$, fix some root $x \in M$ of $f(X) - t$, and let $L = \mathbb{F}_q(x)$ (note that $\mathbb{F}_q(t) \subseteq \mathbb{F}_q(x)$ since $f(x) - t = 0$). Then since the irreducible factors of $f(X) - t_1$ have degrees d_1, d_2, \dots, d_ℓ , there are corresponding places $Q_1, Q_2, \dots, Q_\ell \subset L$ lying over P such that $e(Q_i|P)f(Q_i|P) = d_i$.

Now fix any place $R \subset M$ lying above P . By observing the form of the factorization of $f(X) - t_1$, we see that $R | P$ is unramified since the polynomials $p_{t_1,i}$ are irreducible and pairwise distinct. It follows that $e(R|P) = 1$, and this implies $e(Q_i|P) = 1$ for each i since M is Galois over F . Thus by Lemma [5, Lemma 2.1] the orbit of $D(R|P)$ corresponding to Q_i has size $e(Q_i|P)f(Q_i|P) = f(Q_i|P) = d_i$ for each i . On the other hand, that $R | P$ is unramified yields $D(R|P) \cong \text{Gal}(\mathcal{O}_R/R/\mathcal{O}_P/P) \cong \text{Gal}(k/\mathbb{F}_q) \cong \mathbb{Z}/([k:\mathbb{F}_q]\mathbb{Z})$. From this it follows that we can write $D(R|P) \cong \langle \gamma \rangle$ for some $\gamma \in A$, and since the orbits of $D(R|P)$ are the

same as those of γ , we obtain the decomposition $\gamma = \varphi_1 \circ \cdots \circ \varphi_\ell$ of γ into disjoint cycles φ_i such that the length of φ_i is d_i .

Since $R \mid P$ is unramified and $\gamma \in D(R \mid P)$, we may assume without loss of generality that γ is a Frobenius at P , i.e., that $\gamma|_k$ is the map $a \mapsto a^q$. Note that for any place $P' \in \mathbb{P}^1(F/\mathbb{F}_q)$ we have that if $\sigma \in A$ is a Frobenius at P' , then $\sigma \in G\gamma$ and every Frobenius at P' is conjugate to σ . In particular, $\Gamma_\sigma \subseteq G\gamma$.

Finally, let $S \subseteq G$ be the elements of G that have the same cycle decomposition as γ . Then S consists of exactly the automorphisms $\sigma \in G$ for which σ is a Frobenius at the place P_{t_0} for some $t_0 \in \mathcal{T}$. Note that $\Gamma_\sigma \subseteq S$ for every $\sigma \in S$ and let $\sigma_1, \dots, \sigma_m$ be a list of representatives for the distinct conjugacy classes $\Gamma_{\sigma_1}, \dots, \Gamma_{\sigma_m} \subseteq S$. For each i , let U_i be the set of unramified places $P' \in \mathbb{P}^1(F/\mathbb{F}_q)$ such that σ_i is a Frobenius at P' , and this allows us to write

$$T = \sum_{\substack{P_{t_0} \in \mathbb{P}^1(F/\mathbb{F}_q) \\ t_0 \in \mathcal{T}}} 1 = \sum_{i=1}^m |U_i|.$$

Because the conjugacy classes Γ_{σ_i} partition S , applying Corollary 4.24 to each U_i thus yields

$$\begin{aligned} T &= \sum_{i=1}^m |U_i| \\ &\approx \sum_{i=1}^m \frac{|\Gamma_{\sigma_i}|}{|G|} (q+1) \\ &= \frac{1}{|G|} (q+1) \sum_{i=1}^m |\Gamma_{\sigma_i}| \\ &= \frac{|S|}{|G|} (q+1), \end{aligned}$$

and this completes the proof. □

CHAPTER 5: GOOD POLYNOMIALS OF DEGREE UP TO FIVE

This chapter is based on the original and published work “Optimal selection for good polynomials of degree up to five,” written by the author of this dissertation together with Andrea Ferraguti and Giacomo Micheli (see [11]).

5.1 Introduction

We remind the reader that a code \mathcal{C} with alphabet \mathbb{F}_q , length n , and dimension k is called a locally recoverable code (LRC) with locality parameter r , or an $(n, k, r)_q$ LRC, if for any $c = (c_1, \dots, c_n) \in \mathcal{C}$ and any $1 \leq i \leq n$, the coordinate c_i is a function of at most r other coordinates $c_{i_1}, c_{i_2}, \dots, c_{i_r}$ of c . In other words, the value of any symbol in a particular codeword can be recovered by accessing at most r other symbols of the codeword.

Given the linear (n, k, r) LRC \mathcal{C} , Gopalan *et al.* [21] and Papailiopoulos and Dimakis [32] proved that the minimum distance $d = d(\mathcal{C})$ of \mathcal{C} satisfies the upper bound $d \leq n - k - \lceil k/r \rceil + 2$. As in the literature, we will say \mathcal{C} is an *optimal* LRC if the minimum distance d of \mathcal{C} achieves this bound, that is, if $d = n - k - \lceil k/r \rceil + 2$.

A powerful approach to constructing LRCs was given by Tamo and Barg in [36], and it can be accomplished by constructing polynomials of degree $r + 1$ which are constant on pairwise disjoint subsets of \mathbb{F}_q of size $r + 1$. Such polynomials are called *good* polynomials. This construction and some families of good polynomials are discussed in detail in Section 3.3.

A more formal definition is as follows. For a positive integer ℓ , we will say that the polynomial $f \in \mathbb{F}_q[X]$ is (r, ℓ) -good if

- the degree of f is $r + 1$, and
- there are pairwise disjoint sets $A_1, \dots, A_\ell \subseteq \mathbb{F}_q$, each of cardinality $r + 1$, such that $f(A_i) = \{t_i\}$ for some $t_i \in \mathbb{F}_q$, i.e., f is constant on each A_i .

Given a good polynomial, one can construct an optimal linear LRC as follows (we use the notation of [26], which is the most convenient for our purposes). Fix $r \geq 1$, and let $f(X) \in \mathbb{F}_q$ be a good polynomial. Write $n = (r + 1)\ell$ and $k = rt$, where $t \leq \ell$. For $a = (a_{ij} \mid i = 0, \dots, r - 1; j = 0, \dots, t - 1) \in \mathbb{F}_q^k$, define the encoding polynomial

$$f_a(X) = \sum_{i=1}^{r-1} \sum_{j=0}^{t-1} a_{ij} f(X)^j X^i.$$

Let $A = \bigcup_{i=1}^{\ell} A_i$ and define

$$\mathcal{C} = \{(f_a(x), x \in A) \mid a \in \mathbb{F}_q^k\}.$$

Then \mathcal{C} is an optimal linear (n, k, r) LRC code over \mathbb{F}_q .

In the rest of this chapter, we classify all (r, ℓ) -good polynomials up to $r = 4$ as follows: for any fixed prime power q (even or odd) and a fixed r up to 4, we provide an explicit estimate (of the form $cq + O(\sqrt{q})$, where $c \in [0, 1)$, and the implied constant in the error term is explicitly computable) of the maximal ℓ such that a polynomial of degree $r + 1$ is (r, ℓ) -good. Moreover, we provide examples of polynomials achieving these values for ℓ , showing that the estimate is the best possible.

The machinery we use involves Galois theory, the classification of transitive subgroups of the symmetric group \mathcal{S}_n up to $n = 5$, and the theory of function fields, using the results and techniques of [29, 30], then further developed in [4, 5, 7, 8, 14, 15].

5.2 Monodromy Groups and Totally Split Places

We continue with the notation introduced in Chapter 4, particularly that of Subsection 4.2.2. Recall that, given a separable polynomial $f \in \mathbb{F}_q[X]$, the arithmetic and geometric monodromy groups of f are denoted by $A(f)$ and $G(f)$, respectively.

For any fixed n , it is possible to construct a polynomial f of degree n having $A(f) = G(f) = S_n$. Hence one can define a function

$$G_n(\cdot): \{\text{prime powers}\} \rightarrow \mathbb{N}$$

that assigns to every prime power q the least positive integer such that there exists a separable $f \in \mathbb{F}_q[X]$ of degree n with $|G(f)| = |A(f)| = G_n(q)$. Notice that $G_n(q) \geq n$ for every q as a group with order strictly less than n cannot act transitively on a set of n elements.

Thanks to the techniques introduced in [29], given a separable polynomial $f \in \mathbb{F}_q[X]$ such that $A(f) = G(f)$, one can obtain an explicit estimate on the cardinality of the set

$$\mathcal{T}_{split}^1(f) := \{t_0 \in \mathbb{F}_q: f(X) - t_0 \text{ splits into } \deg(f) \text{ distinct linear factors}\}.$$

This is done via the following result:

Proposition 5.1 ([29, Proposition 3.1]). *Let $f \in \mathbb{F}_q[X]$ be a separable polynomial of degree n with $G(f) = A(f)$ and let g_f be the genus of the splitting field of $f(X) - t$. Then*

$$\frac{q+1-2g_f\sqrt{q}}{|G(f)|} - \frac{n}{2} \leq |\mathcal{T}_{split}^1(f)| \leq \frac{q+1+2g_f\sqrt{q}}{|G(f)|}.$$

The genus g_f can be bounded solely in terms of $\deg(f)$ by using, for example, Castelnuovo's inequality (see Theorem 4.22). As noticed in [29, Proposition 3.3], if $\text{char}(\mathbb{F}_q) \nmid |G(f)|$ then we have

$$g_f \leq \frac{(n-2)|G(f)|+2}{2}.$$

It is clear from the above proposition that for a fixed n , minimizing $|G(f)|$ maximizes the expected number of totally split places, which in turn maximizes the dimension of the Tamo-Barg code.

In this chapter, we compute the function G_n for every $n \in \{2, \dots, 5\}$. The simpler cases $n = 2, 3, 4$ are completely treated in Section 5.3. When $n = 5$ the problem becomes more difficult as, up to conjugation, there are 5 transitive subgroups of the symmetric group S_5 :

- The cyclic group C_5 , generated by a 5-cycle;
- The dihedral group D_5 , generated by a 5-cycle and a product of two disjoint transpositions;
- The affine general linear group $\text{AGL}_1(\mathbb{F}_5)$ (which is isomorphic to $C_5 \rtimes C_4$), generated by a 5-cycle and a 4-cycle;
- The alternating group A_5 ;
- The symmetric group S_5 .

Nevertheless, we prove the following theorem for good polynomials of degree 5.

Theorem 5.2. *Let q be a prime power. Then*

$$G_5(q) = \begin{cases} 5 & \text{if } 5 \mid q(q-1), \\ 10 & \text{if } 5 \mid (q+1), \\ 120 & \text{otherwise.} \end{cases}$$

We note that the estimate of Proposition 5.1 can be made explicit, leading to a formula for the maximal dimension of a Tamo-Barg code of locality 4 as in the next three theorems.

Theorem 5.3. *Let q be a prime power. Let $f \in \mathbb{F}_q[X]$ of degree 5 with $G(f) = A(f) \cong C_5$.*

Then

$$\frac{q+1}{5} - \frac{5}{2} \leq |\mathcal{T}_{split}^1(f)| \leq \frac{q+1}{5}.$$

Proof. Let $x \in \overline{\mathbb{F}_q(t)}$ be a root of $f(X) - t$. Let $F = \overline{\mathbb{F}_q}(x)$ and let M be the splitting field of $f - t$ over $\overline{\mathbb{F}_q}(t)$. Note that $n = \deg(f) = 5$ and $|G(f)| = 5$, so the only quantity remaining to address is $g_f = g_M$.

Since $G(f) \cong C_5$ is cyclic, then we immediately see that $F = M$, and so by Theorem 4.20 we have $g_M = 0$. With this, the desired inequalities follow immediately from Proposition 5.1. \square

Theorem 5.4. *Let q be a prime power for some prime $p \neq 5$. Let $f \in \mathbb{F}_q[X]$ of degree 5 with $G(f) = A(f) \cong D_5$. Then*

$$\frac{q+1}{10} - \frac{5}{2} \leq |\mathcal{T}_{split}^1(f)| \leq \frac{q+1}{10}.$$

Proof. Let $x \in \overline{\mathbb{F}_q(t)}$ be a root of $f(X) - t$. Let $F = \overline{\mathbb{F}_q}(x)$ and let M be the splitting field of $f - t$ over $\overline{\mathbb{F}_q}(t)$. Note that $n = \deg(f) = 5$ and $|G(f)| = 10$, so the only quantity remaining to address is $g_f = g_M$.

Since $G(f) \cong D_5$ is isomorphic to a subgroup of A_5 , the discriminant of $f(X) - t$ (of degree 4) must be a square in $\overline{\mathbb{F}_q}(t)$. From this we deduce that there are distinct places $P, P' \subseteq \overline{\mathbb{F}_q}(t)$ which ramify in M . With the infinite place $P_\infty \subseteq \overline{\mathbb{F}_q}(t)$ (which clearly ramifies), these comprise all the ramified places.

We first consider the place $P \subseteq \overline{\mathbb{F}_q}(t)$. Observe that there must be at least two places of F lying above P . If not, then P is totally ramified, and this implies that $f - t$ is of the form $(x - a)^5 - b - t$, which yields $G(f) \cong C_5$, a contradiction. By examining the cycle decompositions of the elements in D_5 , one concludes that there are places Q_1, \dots, Q_5 lying above P and that these are the only such places. In particular, by Theorem 4.16 we must have $e(Q_i|P) = 1$ for each i since $[F : \overline{\mathbb{F}_q}(t)] = \deg(f - t) = 5$. For each $1 \leq i \leq 5$, let $R_i \subseteq M$ be a place lying above Q_i (and hence above P). Because P ramifies in M we necessarily have $e(R_i|P) \geq 2$ for any $R_i \subseteq M$ lying over P . With $[M : \overline{\mathbb{F}_q}(t)] = 10$ and the fact that $e(R_i|P)$ is independent of the choice of $R_i | P$ since M is Galois over $\overline{\mathbb{F}_q}(t)$, we see

that $e(R|P) = 2$ for any such R . Further, R_1, \dots, R_5 are all the places of M lying above P since $\sum_{i=1}^5 e(R_i|P) = 5 \cdot 2 = [M : \overline{\mathbb{F}}_q(t)]$.

By an identical argument to the above, we have the following: there are places $Q'_1, \dots, Q'_5 \subseteq F$ lying over P' ; for each Q'_i there is a place $R'_i \subseteq M$ lying over Q'_i ; and the Q'_i and R'_i are the only such places lying over P' .

Next we address P_∞ . Let $x' \in M$ be another root of $f(X) - t$ which is distinct from x . Then since $[M : F] = 2$ we must have $M = F(x') = \overline{\mathbb{F}}_q(x, x') = \overline{\mathbb{F}}_q(x)\overline{\mathbb{F}}_q(x')$. Since P_∞ is totally ramified in each of $\overline{\mathbb{F}}_q(x)$ and $\overline{\mathbb{F}}_q(x')$, there is exactly one place $Q \subseteq \overline{\mathbb{F}}_q(x)$ (resp. $Q' \subseteq \overline{\mathbb{F}}_q(x')$) lying over P , and $e(Q|P) = 5$ (resp. $e(Q'|P) = 5$). By Theorem 4.18, if $R \subseteq M$ is a place lying above both of Q and Q' , we have $e(R|P) = \text{lcm}\{e(Q|P), e(Q'|P)\} = 5$. So by Theorem 4.16 there must be exactly two such places $R, R' \subseteq M$.

Finally, noting that the genus of $\overline{\mathbb{F}}_q(t)$ is 0 and using Theorem 4.21 with the above data yields the following:

$$\begin{aligned} 2g_M - 2 &= 10(0 - 2) + \underbrace{5(2 - 1)}_P + \underbrace{5(2 - 1)}_{P'} + \underbrace{2(5 - 1)}_{P_\infty} \\ &= -2. \end{aligned}$$

Thus $g_M = 0$. The desired inequalities now follow immediately from Proposition 5.1. \square

In the next theorem we give an explicit estimate for the remaining cases, restricting for simplicity to the case $2, 3, 5 \nmid q$.

Theorem 5.5. *Let q be a prime power with $2, 3, 5 \nmid q$. Let $f \in \mathbb{F}_q[X]$ of degree 5 with $G(f) = A(f) \cong S_5$. Then*

$$\frac{q + 1 - 72\sqrt{q}}{120} - \frac{5}{2} \leq |\mathcal{T}_{split}^1(f)| \leq \frac{q + 1 + 72\sqrt{q}}{120}.$$

Proof. Let $x \in \overline{\mathbb{F}}_q(t)$ be a root of $f(X) - t$. Let $F = \overline{\mathbb{F}}_q(x)$ and M be the splitting field of $f - t$ over $\overline{\mathbb{F}}_q(t)$. By Proposition 5.1, all we have to do is bound the genus g_M of M .

We will do this via the Riemann-Hurwitz formula (see Theorem 4.21) applied to the degree $24 = 5!/5$ extension M/F . Notice that F has genus g_F equal to zero. We have that

$$2g_M - 2 = 24(2g_F - 2) + \sum_{P \in \mathbb{P}_F} \sum_{\substack{Q \in \mathbb{P}_{F'} \\ Q|P}} (e(Q|P) - 1)$$

because, by our assumptions on q , the ramification is tame. Here the external sum is over all places P of F , while the internal one is over all places Q of M dividing P , and $e(Q|P)$ is the ramification index. Since f has degree 5, its derivative has degree 4 and therefore there are at most 5 places of F that can ramify in M (notice that a place of $\overline{\mathbb{F}}_q(t)$ ramifies in M if and only if it ramifies in F). Now since M/F is a Galois extension the ramification index $e(Q|P)$ depends only on P , and it is at most 24. On the other hand, there are at least two places of F that ramify in M , since there are at least two places of $\overline{\mathbb{F}}_q(t)$ that ramify in F : the infinite place and a finite one, since the derivative of f has positive degree. All in all, we have that

$$\sum_{P \in \mathbb{P}_F} \sum_{\substack{Q \in \mathbb{P}_{F'} \\ Q|P}} (e(Q|P) - 1) = \sum_{P \in \mathbb{P}_F} e(Q|P) - \sum_{P \in \mathbb{P}_F} \sum_{\substack{Q \in \mathbb{P}_{F'} \\ Q|P}} 1 \leq 5 \cdot 24 - 2 = 118,$$

and substituting in the above equation yields $g_M \leq 36$. □

5.3 Degrees Up to 4

In this section we compute G_2, G_3 and G_4 . We start with two general lemmas.

Lemma 5.6. *Let p be a prime and $q = p^m$ for some $m \geq 1$. Let $f = X^q - X \in \mathbb{F}_q[X]$. Then $A(f) = G(f) \cong (C_p)^m$.*

Proof. Let x be a root of f and let $F := \mathbb{F}_q(x)$. Then F is a Galois extension of $\mathbb{F}_q(t)$, because $x + \alpha$ is a root of $f(X) - t$ for every $\alpha \in \mathbb{F}_q$, and therefore $|A(f)| = q$. Since $f(X) - t$ is absolutely irreducible, both $A(f)$ and $G(f)$ act transitively on the set of roots of $f(X) - t$

and therefore it must be that $G(f) = A(f)$. If $\sigma \in A(f)$ and $r = x + \alpha$ is a root of $f - t$ for some $\alpha \in \mathbb{F}_q$, then $\sigma(r) = r + \beta$ for some $\beta \in \mathbb{F}_q$, and therefore $\sigma^p(r) = r$. It follows that σ^p is the identity, and therefore $A(f) = (C_p)^m$. \square

Lemma 5.7. *Let ℓ be a prime and q a prime power with $\ell \nmid q$. Let $f \in \mathbb{F}_q(X)$ be a degree ℓ polynomial. Then $G(f) = A(f) \cong C_\ell$ if and only if $\ell \mid q - 1$ and $f = (X - a)^\ell + b$ for some $a, b \in \mathbb{F}_q$.*

Proof. Necessity is obvious.

Conversely, suppose that $G(f) = A(f) \cong C_\ell$. Let x be a root of $f(X) - t$ and $F := \mathbb{F}_q(x)$. Then the ramification in F is always tame, and hence Theorem 4.21 implies that there must be a finite place of $\overline{\mathbb{F}_q}(t)$ that ramifies in F . Let this place correspond to $b \in \overline{\mathbb{F}_q}$. Then $f(X) - b$ must factor as $(X - a)^\ell$ for some $a \in \overline{\mathbb{F}_q}$. Comparing the coefficients of the linear terms, it follows immediately that $a \in \mathbb{F}_q$, and hence $b \in \mathbb{F}_q$. But then $f = (X - a)^\ell + b$, and in order to have $G(f) = A(f)$ the field of constants $F \cap \overline{\mathbb{F}_q}$ must be \mathbb{F}_q . This immediately implies that $\ell \mid q - 1$ because certainly F contains a primitive ℓ -th root of unity. \square

Theorem 5.8. *The following hold:*

1. $G_2(q) = 2$ for every q , and
2. $G_3(q) = \begin{cases} 3 & \text{if } 3 \mid q(q - 1), \\ 6 & \text{otherwise.} \end{cases}$

Proof. When $n = 2$ and q is odd, every quadratic $f \in \mathbb{F}_q[X]$ has $G(f) = A(f) \cong C_2$. When q is even, by Lemma 5.6 if $f = X^2 + X \in \mathbb{F}_q(X)$ we have $G(f) = A(f) = 2$.

When $n = 3$ and $3 \mid q$, by Lemma 5.6 for $f = X^3 - X$ we have $G(f) = A(f) \cong C_3$. When $3 \mid (q - 1)$ for $f = X^3$ we have $G(f) = A(f) \cong C_3$. When $3 \nmid q(q - 1)$ by Lemma 5.7 we cannot have $G(f) = A(f) = C_3$. The only other transitive group inside S_3 is S_3 itself, and hence $G_3(q) = 6$. \square

Theorem 5.9. *The following holds:*

$$G_4(q) = \begin{cases} 24 & \text{if } q = 2, \\ 4 & \text{if } 4 \mid (q - 1) \text{ or } q = 2^m \text{ for some } m > 1, \\ 8 & \text{otherwise.} \end{cases}$$

Proof. Recall that the transitive groups of degree 4 are $C_4, C_2 \times C_2, D_4, A_4$ and S_4 . Here the non-trivial elements of $C_2 \times C_2$ are products of two disjoint transpositions, and therefore this copy of $C_2 \times C_2$ is contained in A_4 .

If q is even and greater than 2, then there exist distinct elements $\alpha_1, \alpha_2, \alpha_3 \in \mathbb{F}_q^*$ such that $\alpha_1 + \alpha_2 + \alpha_3 = 0$. Now let $f = X(X + \alpha_1)(X + \alpha_2)(X + \alpha_3)$; this is of the form $X^4 + aX^2 + bX$ for some a, b with $b \neq 0$. Therefore if x is a root of $f(X) - t$, then all other roots are of the form $x + \alpha_i$ for some $i \in \{1, 2, 3\}$. It follows that $\mathbb{F}_q(x)/\mathbb{F}_q(t)$ is a Galois extension of degree 4, and therefore $G(f) = A(f)$ and $G_2(q) = 4$. In fact it is easy to see that these monodromy groups are isomorphic to $C_2 \times C_2$: if σ is any element of $G(f)$ and r is any root of $f - t$, then $\sigma(r) = r + \alpha$ for some $\alpha \in \mathbb{F}_q$, and hence $\sigma^2(r) = r$, showing that non-trivial elements have order 2.

If $q = 2$, a quick search shows that the only two separable polynomials f with $A(f) \neq S_4$ are $X^4 + X$ and $X^4 + X^2 + X$. However, the first one has $A(f) \cong D_8$ and $G(f) \cong C_2 \times C_2$, while the second one has $A(f) \cong A_4$ and $G(f) \cong C_2 \times C_2$. Hence $G_4(2) = 24$.

If $4 \mid (q - 1)$, then for $f = X^4$ we have $G(f) = A(f) \cong C_4$, and therefore $G_4(q) = 4$.

Finally, suppose q is odd and $4 \nmid (q - 1)$. If $A(f) \subseteq A_4$, then the discriminant of $f - t$ is a square in $\mathbb{F}_q(t)$. However, this discriminant is always a polynomial of degree 3 in t , so $A(f) \subseteq A_4$ can never happen (and therefore, in particular, we cannot have $A(f) \cong C_2 \times C_2$). If we assume $A(f) = G(f) \cong C_4$, then since the ramification in the splitting field F of $f - t$ is tame, by Theorem 4.21 there is a finite place P of $\overline{\mathbb{F}}_q(t)$ that ramifies in F . Let R be a place of F lying over it. Then the decomposition group $D(R|P)$ is either C_4 or C_2 . In the former case, for some $t_0 \in \overline{\mathbb{F}}_q$ the polynomial $f - t_0$ factors as $(X - a)^4$, and this leads to a contradiction as in the proof of Lemma 5.7. In the latter case, since the element of

order 2 in C_4 is a product of two disjoint transpositions, then up to translations we have $f - t_0 = X^2(X - a)^2$ for some $a, t_0 \in \overline{\mathbb{F}}_q$ with $a \neq 0$. This implies that f equals the composition $g \circ h$, where $g = X^2$ and $h = X(X - a)$, and it is a well-known fact (see for example [16]) that for the Galois group of $g \circ h - t$ to be smaller than D_4 one needs that t and $a^2/4 - t$ are linearly dependent in the \mathbb{F}_2 -vector space $\mathbb{F}_q(t)^*/(\mathbb{F}_q(t)^*)^2$, and this clearly does not hold since $a \neq 0$. Hence $G_4(q) \geq 8$. On the other hand, for the same well-known reasons one has that if $f = X^4 + bX^2$ for some $b \neq 0$, then $G(f) = A(f) = D_4$. Hence $G_4(q) = 8$. \square

5.4 Degree 5: $\text{AGL}_1(\mathbb{F}_5)$ Never Occurs

In this section we will prove that if $5 \nmid q$ then there exists no degree 5 polynomial f with $G(f) \cong \text{AGL}_1(\mathbb{F}_5)$. From now on, we let M be the splitting field of $f - t$ over $\mathbb{F}_q(t)$. If P is a place of $\overline{\mathbb{F}}_q(t)$ and R is a place of M lying above it, we denote by $D(R|P)$ the corresponding decomposition group. For every $t_0 \in \overline{\mathbb{F}}_q$, we denote by P_{t_0} the corresponding place of $\overline{\mathbb{F}}_q(t)$.

We start with a preliminary lemma.

Lemma 5.10. *Let q be a prime power with $5 \nmid q$. Let $f \in \mathbb{F}_q[X]$ be a degree 5 polynomial and assume that $G(f) \cong \text{AGL}_1(\mathbb{F}_5)$. Then there are $a, b, t_0 \in \overline{\mathbb{F}}_q$ with $a \neq b$ such that $f - t_0 = (X - a)^4(X - b)$.*

Proof. We start by showing that $M = \overline{\mathbb{F}}_q(x, x') = \overline{\mathbb{F}}_q(x)\overline{\mathbb{F}}_q(x')$ for any two roots $x \neq x'$ of $f(X) - t$ in M . Observe that $\overline{\mathbb{F}}_q(t)(x) = \overline{\mathbb{F}}_q(x)$ since $t = f(x) \in \overline{\mathbb{F}}_q(x)$ (and similarly for x') and write $F = \overline{\mathbb{F}}_q(x)$ and $F' = \overline{\mathbb{F}}_q(x')$. Clearly we have $[F : \overline{\mathbb{F}}_q(t)] = [F' : \overline{\mathbb{F}}_q(t)] = \deg(f) = 5$. Because $G(f) \cong \text{AGL}_1(\mathbb{F}_5)$ is 2-transitive, the stabilizer $G_x \subseteq G$ of x acts transitively on the four other roots of $f - t$. In particular, since G_x is the Galois group of $(f(X) - t)/(X - x)$ over $\overline{\mathbb{F}}_q(t)$, and since the orbit of x' under the action of G_x is a set of size 4, we have that $[FF' : F] = [G : G_x] = 4$. By definition $M \supseteq FF'$, so since $|G| = 20$ we have $M = FF' = \overline{\mathbb{F}}_q(x, x')$.

Now, let P_∞ be the place at infinity of $\overline{\mathbb{F}_q}(t)$ and let R_∞ be a place of M lying over P_∞ . Let $Q_\infty = R_\infty \cap F$ and consider the ramification index $e(Q_\infty|P_\infty)$ of Q_∞ over P_∞ . Recalling that $f(x) = t$ in F , we have

$$v_{Q_\infty}(f(x)) = v_{Q_\infty}(t) = v_{P_\infty}(t) \cdot e(Q_\infty|P_\infty) = -e(Q_\infty|P_\infty),$$

since P_∞ is a pole of order 1 of t . On the other hand, by the strict triangle inequality we also have $v_{Q_\infty}(f(x)) = \deg(f) \cdot v_{Q_\infty}(x) = -5$. This yields $e(Q_\infty|P_\infty) = 5$, and an identical argument applied to $Q'_\infty = R_\infty \cap F'$ yields $e(Q'_\infty|P_\infty) = 5$. Since we have seen that M is the compositum of the fields F and F' (both of which are tame extensions of $\overline{\mathbb{F}_q}(t)$ as we are working in characteristic $\neq 5$), it now follows from Lemma 4.18 that $e(R_\infty|P_\infty) = \text{lcm}\{e(Q_\infty|P_\infty), e(Q'_\infty|P_\infty)\} = 5$. Thus the decomposition group $D(R_\infty|P_\infty)$ is a group of order 5, and hence it is isomorphic to C_5 .

Next, we claim that there must be some $t_0 \in \overline{\mathbb{F}_q}$ such that for any place R of M lying over P_{t_0} , the decomposition group $D(R|P_{t_0})$ is isomorphic to C_4 . To see this, notice that there must be some $t_0 \in \overline{\mathbb{F}_q}$ such that the decomposition group of any place of M lying above it contains a cycle of order 4. In fact, consider the subset of $G(f)$ of elements of even order that belong to some decomposition group: this contains no transpositions because the transitive copy of $\text{AGL}_1(\mathbb{F}_5)$ inside S_5 contains no transpositions, and on the other hand if all such elements had order 2 then they would all be products of two transpositions. However the decomposition groups generate $G(f)$ ¹, and in this latter case it would follow that $G(f) \subseteq A_5$, which is false once again. So let $t_0 \in \overline{\mathbb{F}_q}$ be such that for some place R of M lying above P_{t_0} , the decomposition group contains a cycle of order 4. If we had $C_4 \subsetneq D(R|P_{t_0})$, it would follow that $D(R|P_{t_0}) \cong G(f)$ by the maximality of C_4 in $G(f)$. But then $R | P_{t_0}$ would be totally ramified, and hence $f(X) - t_0 = (X - a_0)^5$ for some $a_0 \in \overline{\mathbb{F}_q}$. Since the field of

¹This is because if G is the subgroup of $G(f)$ generated by all the decomposition subgroups, then M^G is an unramified extension of $\overline{\mathbb{F}_q}(t)$, and there are no non-trivial such extensions.

constants of $M/\overline{\mathbb{F}}_q(t)$ is trivial, this factorization implies $5 \mid (q - 1)$. But then $G \cong C_5$, an immediate contradiction. Thus $D(R|P_{t_0}) \cong C_4$.

To conclude the proof, notice that specializing at the place P_{t_0} and applying the Dedekind-Kummer Theorem [35, Theorem 3.3.7] allows us to write $f(X) - t_0 = (X - a)^4(X - b)$ for $a, b \in \overline{\mathbb{F}}_q$ with $a \neq b$. \square

We are now ready to prove that $\text{AGL}_1(\mathbb{F}_5)$ cannot occur as a geometric monodromy group. The proof will require separate arguments for even and odd characteristics.

Theorem 5.11. *Let q be a prime power with $5 \nmid q$ and $f \in \mathbb{F}_q[X]$ a polynomial of degree 5. Then $G(f) \not\cong \text{AGL}_1(\mathbb{F}_5)$.*

Proof. Assume by contradiction that $G(f) \cong \text{AGL}_1(\mathbb{F}_5)$. By Lemma 5.10, there are elements $t_0, a, b \in \overline{\mathbb{F}}_q$ such that $f - t_0 = (X - a)^4(X - b)$. We can assume without loss of generality that $t_0 = a = 0$ and $b \neq 0$ since $\text{Gal}(f(X) - t \mid \overline{\mathbb{F}}_q(t)) \cong \text{Gal}(f(X - c) - (t - d) \mid \overline{\mathbb{F}}_q(t - d))$ for every $c, d \in \overline{\mathbb{F}}_q$.

First, assume the characteristic of \mathbb{F}_q is *odd*. Computing $f'(X) = X^3(5X - 4b)$ shows that $f'(4b/5) = 0$, so for $t_1 = f(4b/5) \in \overline{\mathbb{F}}_q$, we see that $f(X) - t_1$ is divisible by $(X - 4b/5)^2$. Furthermore, since $X = 4b/5$ is not a root of $f''(X) = 4X^2(5X - 3b)$, it follows that $X = 4b/5$ is precisely a double root of $f(X) - t_1$. Notice that

$$t_1 = f(4b/5) = (4b/5)^4(-b/5) \neq 0$$

and hence $t_1 \neq 0$, so $X = 0$ is not a root of $f(X) - t_1$. This implies that the only repeated root of $f(X) - t_1$ is $X = 4b/5$ since the only roots of $f'(X)$ are $X = 0$ and $X = 4b/5$. In other words, we can write $f(X) - t_1 = (X - 4b/5)^2(X - x_1)(X - x_2)(X - x_3)$ for pairwise distinct elements $4b/5, x_1, x_2, x_3 \in \overline{\mathbb{F}}_q$. Finally, let R be any place of M lying over P_{t_1} . Then the previous factorization shows that there is a transposition in $D(R|P_{t_1}) \subseteq G(f)$. But $G(f)$ contains no transpositions, and we have a contradiction.

Next, assume the characteristic of \mathbb{F}_q is *even*. From now on, we let x be some fixed root of $f(X) - t$ in M and let $F = \overline{\mathbb{F}_q}(x)$. Fix a place R of M lying above the place P_0 of $\overline{\mathbb{F}_q}(t)$. The natural action of $C_4 \cong D(R|P_0) \subseteq G$ on the set of roots of $f(X) - t$ yields orbits of sizes 4 and 1, so there must be two places Q_0 and Q_1 of F lying over P_0 with ramification indices $e(Q_0|P_0) = 4$ and $e(Q_1|P_0) = 1$, respectively, by [5, Lemma 2.1]. Let R_0 be a place of M lying over Q_0 . We have just seen that $4 = |D(R_0|P_0)| = e(R_0|P_0) = e(R_0|Q_0) \cdot e(Q_0|P_0)$, so it follows that $e(R_0|Q_0) = 1$.

Before proceeding, we introduce the following notation: given a function field K and a place P of K , we will write \hat{K}_P to denote the completion of K at P with respect to the P -adic metric. In particular, $\hat{F}_{Q_0} = \overline{\mathbb{F}_q}((x))$ and $\hat{\mathbb{F}_q}(t)_{P_0} = \overline{\mathbb{F}_q}((t))$. Using a well-known number theoretical fact (see for example [31, Proposition II.9.6]), we have $\text{Gal}(\hat{M}_{R_0}/\overline{\mathbb{F}_q}((t))) \cong D(R_0|P_0) \cong C_4$. Observe that $\hat{M}_{R_0} \supseteq \hat{F}_{Q_0} \supseteq \overline{\mathbb{F}_q}((t))$, so since $[\hat{F}_{Q_0} : \overline{\mathbb{F}_q}((t))] = e(Q_0|P_0) = 4$ and $[\hat{M}_{R_0} : \overline{\mathbb{F}_q}((t))] = e(R_0|P_0) = 4$ we have $\hat{M}_{R_0} = \hat{F}_{Q_0}$. In particular, $\hat{F}_{Q_0}/\overline{\mathbb{F}_q}((t))$ is a Galois extension. Denoting the local Galois group $\text{Gal}(\hat{F}_{Q_0} | \overline{\mathbb{F}_q}((t)))$ by \hat{G} , we have $\hat{G} \cong C_4$.

The above shows that every root of $f(X) - t$ in \hat{M}_{R_0} can be expressed as an element of $\hat{F}_{Q_0} = \overline{\mathbb{F}_q}((x))$, that is, as a Laurent series in x . We proceed by showing that if $z \neq x$ is any other root of $f(X) - t$, then we can write $z = x + ux^i$ for some $i \geq 2$ and $u \in \overline{\mathbb{F}_q}[[x]]$. First, recall that $f(X) - t = X^4(X - b) - t$ so that b is a simple root of $f(X)$. Then by Hensel's lifting lemma (see (4.6) in [31]) there is some $\bar{b} \in \overline{\mathbb{F}_q}((t))$ such that we can write $f(X) - t = \bar{f}(X)(X - \bar{b})$ over $\overline{\mathbb{F}_q}((t))$, where $\bar{f}(X) \in \overline{\mathbb{F}_q}((t))[X]$ and $\deg(\bar{f}) = 4$. Further, the polynomial \bar{f} must be irreducible over $\overline{\mathbb{F}_q}((t))$ since otherwise we could write $\bar{f}(X) = \bar{g}(X)(X - r)$ for an irreducible $\bar{g} \in \overline{\mathbb{F}_q}((t))[X]$ and some $r \in \overline{\mathbb{F}_q}((t))$, or we could write $\bar{f}(X) = \bar{h}_1(X)\bar{h}_2(X)$ for two irreducible quadratic polynomials $\bar{h}_1, \bar{h}_2 \in \overline{\mathbb{F}_q}((t))[X]$ having distinct roots in F_{Q_0} (since $f - t$ is separable). The former factorization implies $|\hat{G}| = 4$ divides $3!$, a clear contradiction, so assume the latter factorization holds and let \hat{H}_1 and \hat{H}_2

be the splitting fields of $\bar{h}_1(X)$ and $\bar{h}_2(X)$, respectively, in F_{Q_0} . Then $F_{Q_0} = H_1H_2$ so that

$$C_4 \cong \hat{G} = \text{Gal}(H_1H_2 \mid \bar{\mathbb{F}}_q((t))) \subseteq \text{Gal}(H_1 \mid \bar{\mathbb{F}}_q((t))) \times \text{Gal}(H_2 \mid \bar{\mathbb{F}}_q((t))) \cong C_2 \times C_2,$$

another contradiction. Thus we conclude that \hat{G} is the Galois group of $\bar{f}(X)$ over $\bar{\mathbb{F}}_q((t))$; in particular, we have that \hat{G} acts transitively on the roots of $\bar{f}(X)$ in F_{Q_0} and hence there is some automorphism $\tau \in \hat{G}$ of M satisfying $\tau(x) = z$.

Observe that $\langle x \rangle$ is the unique maximal ideal of the ring $\bar{\mathbb{F}}_q[[x]]$, so since τ is an automorphism of $\bar{\mathbb{F}}_q((x))$ (and hence τ preserves maximal ideals) we must have $\langle \tau(x) \rangle = \tau(\langle x \rangle) = \langle x \rangle$. Then $z \equiv 0 \pmod{\langle x \rangle}$ if and only if $\tau(x) \equiv 0 \pmod{\langle \tau(x) \rangle}$, and the latter clearly holds. This allows us to write $z = cx + ux^i$ for some $c \in \bar{\mathbb{F}}_q$, $u \in \bar{\mathbb{F}}_q[[x]]$, and some $i \geq 1$. We can assume further that $i \geq 2$ since otherwise we could replace cx by $c'x$ for an appropriate $c' \in \bar{\mathbb{F}}_q^*$ so that this holds. Now computing $\tau^4(x)$ by using $\tau(x) = cx + ux^i$ and comparing coefficients with $\tau^4(x) = x$ yields $c^4 = 1$ and hence $c = 1$ (as we are working over a field with characteristic 2). Putting everything together, we can now write $z = x + ux^i$ for some $i \geq 2$ and some $u \in \bar{\mathbb{F}}_q[[x]]^*$.

Observe the following:

$$\begin{aligned} f(z) - t &= z^4(z - b) - t \\ &= (x + ux^i)^4(x + ux^i - b) - t \\ &= (x^4 + u^4x^{4i})(x - b + ux^i) - t \\ &= x^4(x - b) - t + ux^{4+i} + u^4x^{4i}(x - b) + u^5x^{5i} \\ &= ux^{4+i} + u^4x^{4i}(x - b) + u^5x^{5i}, \end{aligned}$$

where the last equality holds since x is a root of $f(X) - t$. Let $A = ux^{4+i} + u^4x^{4i}(x - b) + u^5x^{5i} = f(z) - t$. Since z was chosen to be another root of $f(X) - t$, we must have $A = 0$.

But $v_{Q_0}(A) = v_{Q_0}(x) \cdot \min_{i \geq 2} \{4 + i, 4i, 5i\} = 4 + i \neq \infty = v_{Q_0}(0)$, a contradiction since $u \neq 0$.

Thus our initial assumption was false, so $G(f) \not\cong \text{AGL}_1(\mathbb{F}_5)$. \square

5.5 Degree 5: If D_5 or A_5 Occurs, Then $5 \mid (q^2 - 1)$

Assume q is a prime power and $f \in \mathbb{F}_q[X]$ is a separable polynomial of degree 5. Let t be transcendental over \mathbb{F}_q and let M be the splitting field of $f(X) - t$. Let $A(f)$ and $G(f)$ be the arithmetic and geometric monodromy groups of f , respectively.

Theorem 5.12. *Suppose that $5 \nmid q$ and $A(f) \subseteq A_5$. Then $5 \mid (q^2 - 1)$. In particular, if $A(f) \subseteq D_5$, then $5 \mid (q^2 - 1)$.*

Proof. The second assertion follows immediately from the fact that the transitive copy of D_5 inside S_5 lies inside A_5 .

If q is odd, just use the fact that $A(f) \subseteq A_5$ if and only if the discriminant of $f(X) - t$ is a square in $\mathbb{F}_q(t)$. When $f(X)$ is monic of degree 5, the discriminant has the form $5^5 t^4 + \sum_{i=0}^3 a_i t^i$. Hence 5 needs to be a square in \mathbb{F}_q ; this implies that either q is an even power of a prime p or, by quadratic reciprocity, that $q \equiv \pm 1 \pmod{5}$. In any case, $5 \mid q^2 - 1$.

If $q = 2^n$ for some $n \geq 1$, one needs to use the *Berlekamp discriminant* (see [9]), which is the characteristic two analogue of the discriminant. If k is a field of characteristic 2 and $g \in k[X]$ is of degree n , the Berlekamp discriminant of g is an element $\Delta \in k$ which can be effectively computed using the coefficients of g . Further, Δ has the property that $\text{Gal}(g) \subseteq A_n$ if and only if the polynomial $X^2 + X + \Delta$ has a root in k .

Now let $f = X^5 + aX^4 + bX^3 + cX^2 + dX \in \mathbb{F}_{2^n}[X]$. We will show that if $A(f) \subseteq A_5$ then $2 \mid n$, and consequently $5 \mid (q^2 - 1)$ once again. One can compute the Berlekamp discriminant Δ_f of $f - t$, viewed as a polynomial over $\mathbb{F}_{2^n}(t)$, and see that this is given by an expression of the form $r(t)/s(t)^2$, where $r, s \in \mathbb{F}_{2^n}(t)$ are two monic polynomials with $\deg(r) = 4$ and $\deg(s) = 2$. Suppose that $A(f) \subseteq A_5$ and hence that $X^2 + X + \Delta_f$ has a root in $\mathbb{F}_{2^n}(t)$. Then there are coprime polynomials $u(t), v(t) \in \mathbb{F}_{2^n}[t]$, with $v(t)$ monic, such that $\Delta_f = (u(t)^2 + u(t)v(t))/v(t)^2$. Therefore if r, s share a common factor, this can only

have degree 2 or 4, and if it has degree 2 then it must be of the form $(t + \alpha)^2$ for some α . Clearly if they share a factor of degree 4 then $\Delta_f = r(t)/s(t)^2 = 1$, so in this case $X^2 + X + 1$ has a root in $\mathbb{F}_{2^n}(t)$ and consequently $2 \mid n$. Otherwise, by comparing degrees we must have $\deg(u) = \deg(v)$ and $\deg(u^2 + uv) = 2 \deg(v)$. If the leading coefficient of u is $\delta \in \mathbb{F}_{2^n}$, these two conditions, together with the fact that $r(t)$ is monic of degree 4, imply that $\delta^2 + \delta + 1 = 0$ and consequently that $2 \mid n$. \square

Proof of Theorem 5.2. First, suppose that $5 \mid q(q - 1)$. Then by Lemmas 5.6 and 5.7 we have $G_5(q) = 5$.

Now suppose that $5 \nmid q(q - 1)$. Then by Lemma 5.7 we have $G_5(q) > 5$. If $5 \mid (q + 1)$, it is known (see for example [10, Section 3]) that degree 5 Dickson polynomials of the first kind, e.g. $f = X^5 - 5X^3 + 5X$, satisfy $G(f) = A(f) = D_5$. Hence $G_5(q) = 10$.

Finally, suppose that $5 \nmid q(q^2 - 1)$. Then by Lemma 5.7 and Theorems 5.11 and 5.12 we cannot have $G_5(q) = 5, 10, 20$ or 60 . Hence $G_5(q) = 120$. \square

Remark 5.13. Notice that the $q/10$ asymptotic when $5 \mid (q + 1)$ was in fact obtained in [27] using Dickson Polynomials and an independent approach.

5.6 Computational Examples

Let us show with a couple of explicit examples how the number of totally split places compares to the theoretical estimate given by Proposition 5.1. We pick examples with $G_5(q) = 120$; for each of these values of q we pick polynomials f with $G(f) = A(f) \cong S_5$ for $5 \nmid q(q^2 - 1)$. As proved in Theorem 5.2, it is not possible to do better for these q 's.

In order to construct polynomials whose geometric monodromy (and therefore also arithmetic monodromy) group is S_5 , one can use the following well-known group-theoretical fact (see [18]): if $G \subseteq S_5$ is a transitive subgroup containing a transposition and a cycle of prime length $\ell > 2$, then $G = S_5$. In order to force the geometric monodromy group to contain two such elements, it is enough, by ramification arguments, to pick $g(X) \in \mathbb{F}_q[X]$ irreducible of degree 3 and set $f = X^2g(X)$.

q	$ \mathcal{T}_{split}^1(f) $	$\lfloor \frac{q}{120} \rfloor$
2^{13}	78	68
2^{15}	278	273
2^{17}	1088	1092
2^{19}	4332	4369

(a) $f = X^2(X^3 + X + 1)$

q	$ \mathcal{T}_{split}^1(f) $	$\lfloor \frac{q}{120} \rfloor$
3^7	21	18
3^9	159	164
3^{11}	1474	1476
3^{13}	13338	13286

(b) $f = X^2(X^3 - X + 1)$

q	$ \mathcal{T}_{split}^1(f) $	$\lfloor \frac{q}{120} \rfloor$
19583	156	163
19597	163	163
19687	155	164
19753	194	164
19793	179	164
19913	189	165
19927	160	166
19963	162	166
19993	156	166
19997	161	166

(c) $f = X^2(X^3 + X + 3)$

Table 1. Comparing the asymptotic and actual number of totally split places for a few degree 5 polynomials defined over varying base fields.

CHAPTER 6: HIERARCHICAL GOOD POLYNOMIALS

This chapter is based on the original and published work “Optimal locally recoverable codes with hierarchy from nested F -adic expansions,” written by the author of this dissertation together with Giacomo Micheli and Vincenzo Pallozzi Lavorante (see [13]).

6.1 Introduction

Various classes of locally recoverable codes have received great attention in recent times due to their applications to cloud and distributed storage systems [2, 3, 6, 11, 20, 23, 26, 34, 36, 37].

In this chapter, we produce new optimal hierarchical locally recoverable codes (HLRCs). HLRCs are suitable solutions for the problem of recovering lost information in a distributed storage system, and they have been widely studied in [1, 17, 33, 38].

Hierarchical locally recoverable codes allow to recover certain patterns of erasures by gradually looking at more components of a codeword depending on the number of erasures that occurred in that codeword. One can then design codes which recover one erasure by looking at at most b other components; λ erasures by looking at a other components; and $d - 1$ erasures by looking at at most k components, where k is the dimension of the code. This is impactful from a practical perspective, as one can deal with the most likely scenario (1 erasure) in the optimal way, with the less likely scenario (λ erasures) in an acceptable way, and still be able to recover $d - 1$ erasures by accessing a large number of nodes. Tuning these parameters in an efficient way depends on the reliability of the servers and the required efficiency of the system in terms of node retrieval. One of the features that one would desire

from this kind of code is that λ is not too large, as the second most likely scenario is the failure of just a few other nodes more than 1 (and not too many others since the worse cases can be recovered using the minimum distance). We address this problem by proving a sharper Singleton-like bound for this regime of parameters and then constructing codes which achieve our bound. Let us now define the main objects we will be treating in this chapter.

6.1.1 Definitions

In the rest of the chapter we will consider the occurrence of either 1, λ , or $d - 1$ erasures, as these arise most commonly from applications (instead of the more general setting where one allows λ_1 , λ_2 , or $d - 1$ erasures). Let n, k, b be positive integers with $k \leq n$. A locally recoverable code (LRC) \mathcal{C} having parameters $[n, k, b]$ is an \mathbb{F}_q -subspace of \mathbb{F}_q^n of dimension k such that if one deletes one component of any $c \in \mathcal{C}$, this can be recovered by accessing at most b other components of c . If d is the minimum distance of the code, we will write that \mathcal{C} is an $[n, k, d, b]$ LRC.

We now give the following definition which will be useful in the rest of the chapter.

Definition 6.1. Let n be a positive integer, $\mathcal{C} \subseteq \mathbb{F}_q^n$ be a linear code, and S be a subset of the set of indices $\{1, \dots, n\}$. We say that \mathcal{C} can tolerate x erasures on S if, whenever there are x erasures on components of a codeword with indices belonging to S , the missing components can be recovered by looking at $|S| - x$ other coordinates in S .

In this chapter we construct new locally recoverable codes with hierarchy of locality sets. Our Definition 6.2 is equivalent to the one of hierarchical codes in [1], but we find it slightly easier to employ ours in practical situations as we keep direct track of the size of the “hierarchy.”

Definition 6.2. Let n, k, d, b, a, λ be positive integers with $n > k$ and $2 \leq \lambda \leq b$. An $[n, k, d, b, a, \lambda]$ hierarchical locally recoverable code (HLRC) is an $[n, k, d]$ -linear code such that

- $(a + \lambda) \mid n$,
- $(b + 1) \mid (a + \lambda)$,
- the codeword indices are partitioned into $\ell \geq 1$ distinct sets A_i , each of size $a + \lambda$, such that \mathcal{C} tolerates λ erasures on A_i for every $i \in \{1, \dots, \ell\}$, and
- each A_i can be partitioned into distinct sets $B_{i,j}$, each of size $b + 1$, such that \mathcal{C} tolerates 1 erasure on each $B_{i,j}$ for every $i \in \{1, \dots, \ell\}$ and every $j \in \{1, \dots, (a + \lambda)/(b + 1)\}$.

6.1.2 Motivation

Let us now briefly explain the motivation behind codes with hierarchical locality. Let T be the time needed to replace a failed node. Suppose that a second node fails in the same locality set as the first node during the time T . An $[n, k, d, b]$ LRC will still need to access k information symbols, as the 1-locality procedure is not guaranteed to work anymore. However, an $[n, k, d, b, a, \lambda]$ HLRC only requires accessing at most a information symbols. Since the failure of only a few nodes, say $\lambda < d - 1$, is significantly more likely than the failure of $d - 1$ nodes in the span of time T , it is convenient to have a code which addresses separately the case in which only λ nodes fail. The codes in [1] address this issue, but they are restricted to certain λ 's (as we explain in subsection 6.3.5). Moreover, in many cases they require restrictions on the arithmetic of q and the size of the hierarchy (see for example the case of power functions in [1, Section IV.A, Example]).

6.1.3 Our Contribution

In this chapter we provide new constructions of optimal codes with hierarchical locality and an improved bound for HLRCs for a special set of parameters. Our construction is based on the ideas in [1] combined with powerful techniques from algebraic number theory, allowing us to remove arithmetic restrictions on the size of the hierarchy compared with q or $q - 1$.

Structure of the chapter:

- In Section 6.1 and its subsections we explain the basic coding theoretical definitions and provide the practical motivations for the study of such codes.
- In Section 6.2, for some regime of parameters, we provide a stronger Singleton bound than the one already present in the literature for HLRCs [33]. Our bound beats the previous bound for an infinite set of parameters (see for example Remark 6.5).
- In Section 6.3 we achieve our new bound with a new construction of HLRC that covers a set of parameters that are not available using previous constructions (see subsection 6.3.5). In Subsection 6.3.6 we construct one of our codes to show what a generator matrix looks like in practice.
- In Section 6.4 we show that our codes can be constructed without any arithmetic restrictions on $q, q - 1$, the locality parameters, or the sizes of the sets in the hierarchy.
- Using the existential results provided in Section 6.4, in Section 6.5 we provide some practical choice of parameters for codes with large length.

6.2 An Improved Bound for Hierarchical Locally Recoverable Codes

6.2.1 The Singleton Bound for $[n, k, d, b, a, \lambda]$ HLRCs With $\lambda \leq b$.

To help the reader understand the more complex bound we propose on HLRCs in what follows, we encourage the reader to refer first to the standard Singleton-like bound for LRCs and its proof given in Section 3.2.

We aim to generalize the bound in Theorem 3.1 when \mathcal{C} is an $[n, k, d, b, a, \lambda]$ HLRC. The key observation is that one can partition the columns of the generator matrix into ℓ sets of $a + \lambda$ columns so that each set has rank strictly less than a , and each set of $a + \lambda$ columns can be partitioned further into sets of $b + 1$ columns so that each of these sets has rank at most b . To see this, note that each set S_i of $a + \lambda$ columns (corresponding to the indices in

A_i) can be divided into $\beta = (a + \lambda)/(b + 1)$ sets, say $S_{i,j}$ for $j \in \{1, \dots, \beta\}$, of $b + 1$ columns (corresponding to $B_{i,j}$) with rank at most b for $i \in \{1, \dots, \ell\}$ by the definition of the code. Now, in the first set $S_{i,1}$ we have λ columns which are in the span of the other a columns in S_i . This means that we can choose $b + 1 - \lambda$ columns from $S_{i,1}$ and b columns from each of the other $S_{i,j}$, with $j \neq 1$, and be able to recover any λ of the $a + \lambda$ columns in S_i . Therefore, the rank of each S_i is at most

$$\rho := \underbrace{\left[\frac{(a + \lambda)}{(b + 1)} - 1 \right]}_{\beta - 1} b + (b + 1 - \lambda) \leq a.$$

Theorem 6.3. *Let \mathcal{C} be an $[n, k, d, b, a, \lambda]$ HLRC with $\lambda \leq b$, and let $\rho = b(a + \lambda)/(b + 1) - (\lambda - 1)$. Then*

$$\left\lfloor \frac{k - 1}{\rho} \right\rfloor (a + \lambda) + k_1 + \left\lfloor \frac{k_1}{b} \right\rfloor \leq n - d, \quad (6.1)$$

where $k - 1 \equiv k_1 \pmod{\rho}$ and $0 \leq k_1 < \rho$.

Proof. Given $\left\lfloor \frac{k-1}{\rho} \right\rfloor$ locality sets A_i , say $i \in \{1, \dots, \left\lfloor \frac{k-1}{\rho} \right\rfloor\}$, denote by S the set of the corresponding columns of the generator matrix \mathcal{G} of the code. Then $|S| = \left\lfloor \frac{k-1}{\rho} \right\rfloor (a + \lambda)$, and, by the above discussion, we have $\text{rank}(S) = \rho \left\lfloor \frac{k-1}{\rho} \right\rfloor \leq k - 1$. This allows us to add more columns to S until the rank equals $k - 1$ using a smaller locality set. More precisely, we can always choose a set of the remaining columns of \mathcal{G} , say S_1 , of size $\left\lfloor \frac{k_1}{b} \right\rfloor (b + 1) + \left\{ \frac{k_1}{b} \right\} b$, such that $\text{rank}(S_1) = k_1$ (explicitly, S_1 is the union of columns which correspond to $B_{i,j}$, for $j \in \{1, \dots, \left\lfloor \frac{k_1}{b} \right\rfloor\}$). Hence

$$\text{rank}(S \cup S_1) = \left\lfloor \frac{k - 1}{\rho} \right\rfloor \rho + k_1 = k - 1$$

by the definition of k_1 . Applying Proposition 3.2 we have

$$\left\lfloor \frac{k - 1}{\rho} \right\rfloor (a + \lambda) + \left\lfloor \frac{k_1}{b} \right\rfloor (b + 1) + \left\{ \frac{k_1}{b} \right\} b \leq n - d.$$

Now, since $\frac{k_1}{b} = \lfloor \frac{k_1}{b} \rfloor + \{\frac{k_1}{b}\}$ we have

$$\left\lfloor \frac{k-1}{\rho} \right\rfloor (a+\lambda) + k_1 + \left\lfloor \frac{k_1}{b} \right\rfloor \leq n-d. \quad \square$$

Definition 6.4. We say that an $[n, k, d, b, a, \lambda]$ HLRC is *optimal* if its minimum distance attains the upper bound in (6.1), i.e., if

$$d = n - \left(\left\lfloor \frac{k-1}{\rho} \right\rfloor (a+\lambda) + k_1 + \left\lfloor \frac{k_1}{b} \right\rfloor \right),$$

for $k-1 \equiv k_1 \pmod{\rho}$ and $0 \leq k_1 < \rho$.

Remark 6.5. Note that our bound improves upon the bound in [33] for infinitely many parameters, but ours holds only for $\lambda \leq b$. In fact, for any length n , and for parameters $k=6$, $a=4$, $r_1=\rho=3$, $r_2=b=2$, $\delta_1=\lambda+1=3$, and $\delta_2=2$, [33, Theorem 2.1] gives $d \leq n-8$ when instead our bound gives $d \leq n-9$. The moral reasons for this are that we are taking into account a finer arithmetic of the parameters which involves the reduction of the dimension modulo the upper level hierarchical locality, and we are restricting to the case in which the number of nodes that we simultaneously erase is strictly smaller than the size of the smaller locality set.

6.3 Our Construction of Optimal HLRCs Using Nested f -adic Expansions

6.3.1 Main Tool for the Construction

Lemma 6.6. *Let $f, h \in \mathbb{F}_q[X]$ be non-constant polynomials. Suppose there is some $t_0 \in \mathbb{F}_q$ such that $f(h(X)) - t_0$ splits completely (i.e., factors into $\deg(f)\deg(h)$ distinct factors) over \mathbb{F}_q . Then the set of roots of $f(h(X)) - t_0$, say A_0 , can be partitioned into sets $B_1, \dots, B_{\deg(f)} \subseteq \mathbb{F}_q$ which satisfy the following:*

- $h(B_i) = c_i \in \mathbb{F}_q$ for each $1 \leq i \leq \deg(f)$,
- the cardinality of each B_i is $\deg(h)$, and

- $h(B_i) \neq h(B_j)$ whenever $i \neq j$.

Proof. By the hypothesis we may write $f(h(X)) - t_0 = \prod_{i=1}^{\deg(f)\deg(h)} (X - x_i)$ for distinct elements $x_1, \dots, x_{\deg(f)\deg(h)} \in \mathbb{F}_q$. Notice now that if $f(h(X)) - t_0$ splits completely, then $f(X) - t_0$ splits completely. If we let $\alpha_1, \dots, \alpha_{\deg(f)} \in \overline{\mathbb{F}_q}$ be the (distinct) roots of $f(X) - t_0$, then we may also write $f(h(X)) - t_0 = \prod_{i=1}^{\deg(f)} (h(X) - \alpha_i) \in \overline{\mathbb{F}_q}[X]$. Combining these two factorizations and relabeling the x_i appropriately yields

$$\prod_{i=1}^{\deg(f)\deg(h)} \prod_{j=1}^{\deg(h)} (X - x_{i,j}) = \prod_{i=1}^{\deg(f)} (h(X) - \alpha_i),$$

so that $\prod_{j=1}^{\deg(h)} (X - x_{i,j}) = h(X) - \alpha_i$ for each $1 \leq i \leq \deg(f)$. In particular, it follows that $\alpha_i \in \mathbb{F}_q$ for each i . Write $B_i = \{x_{i,j} : 1 \leq j \leq \deg(h)\}$. Then we have $h(B_i) = \alpha_i$ for each i , proving the first statement. The second and third statements both follow from the fact that the $x_{i,j}$'s are pairwise distinct and the B_i 's are pairwise disjoint. \square

Definition 6.7. For $f, h \in \mathbb{F}_q[X]$, we say that a set $A \subset \mathbb{F}_q$ is a *nest* for (f, h) if A is the set of preimages of $t_0 \in \mathbb{F}_q$ such that $f(h(X)) - t_0$ is totally split.

Furthermore, we say that $B \subset A$ is a *sub-nest* if h is constant on B and $|B| = \deg(h)$.

6.3.2 The Main Construction

We present a general method of constructing linear codes with the nested locality property. Later we will show that these codes are optimal in the sense of Section 6.2. In line with the notion of (r, ℓ) -good polynomials in [29], we now begin defining our *nested* polynomials.

Definition 6.8 (*ℓ -nested*). Let $f, h \in \mathbb{F}_q[X]$ and let ℓ be a positive integer. Then f and h are said to be ℓ -nested if $f(h(X)) - t_0$ splits completely over \mathbb{F}_q for at least ℓ elements $t_0 \in \mathbb{F}_q$.

Remark 6.9. Note that if f and h are ℓ -nested, then from Lemma 6.6 there exist A_1, \dots, A_ℓ distinct nests for (f, h) such that

- for any $i \in \{1, \dots, \ell\}$, $f(h(A_i)) = \{t_i\}$ for some $t_i \in \mathbb{F}_q$,
- $|A_i| = \deg(f) \deg(h)$,
- $A_i \cap A_j = \emptyset$ for any $i \neq j$, and
- each A_i can be partitioned into sub-nests $B_{i,j}$ for (f, h) .

Those properties will be fundamental in the following construction.

Construction 6.10 (Nested HLRC). Let $f, h \in \mathbb{F}_q[X]$ be ℓ -nested, with $3 \leq \deg(h) = b + 1$ and $\deg(f) = \frac{a+\lambda}{b+1}$ for some integer $2 \leq \lambda \leq b$, and let $\mathcal{A} = \cup_{i=1}^{\ell} A_i$, where $\{A_1, \dots, A_\ell\}$ is a set of nests for (f, h) .

For a positive integer $s \geq 1$, consider the set \mathcal{D} of polynomials of the form

$$\partial(X) = \sum_{i=0}^s \left[\left(\sum_{j=0}^{\deg(f)-2} g_{i,j}(X) h(X)^j \right) + \tilde{g}_i(X) h(X)^{\deg(f)-1} \right] f(h(X))^i, \quad (6.2)$$

where $g_{i,j} \in \mathbb{F}_q[X]_{\leq \deg(h)-2}$ and $\tilde{g}_i \in \mathbb{F}_q[X]_{\leq \deg(h)-\lambda-1}$.

Let $n = \deg(f) \deg(h) \ell$ and let k be the dimension of \mathcal{D} as an \mathbb{F}_q -vector space.

Define

$$\mathcal{D} = \{(\partial(x), x \in \mathcal{A}) \mid \partial \in \mathcal{D}\}. \quad (6.3)$$

We will prove that \mathcal{C} is an optimal $[n, k, b, a, \lambda]$ HLRC over \mathbb{F}_q .

6.3.3 Locality

Since we evaluate at n distinct points of \mathbb{F}_q , we need $q \geq n$. Write $n = (a + \lambda)(s + 1)$ and recall that $b + 1$ divides $a + \lambda$.

Take $m \in \mathbb{F}_q^k$ and write $\text{Enc}_{\mathcal{C}}(m) = c = c_{i,j_1,j_2}$ for $1 \leq i \leq s + 1$, $1 \leq j_1 \leq (a + \lambda)/(b + 1)$, $1 \leq j_2 \leq b + 1$. Note that the first index i determines a nest A_i , the second index j_1 refers to a sub-nest B_{i,j_1} , and the final index j_2 specifies the particular element of the sub-nest in question, which we denoted by c_{i,j_1,j_2} . We begin by showing that the code \mathcal{C} described in Construction 6.10 allows one to recover a single missing component of c by accessing at most b other components of c .

Fix $\ell \geq 1$ and let $f, h \in \mathbb{F}_q[X]$ be the ℓ -nested polynomials from which \mathcal{C} is obtained. Write $\mathcal{A} = \{A_1, \dots, A_\ell\}$ with $A_i = \bigsqcup_{j_1=1}^{\deg(f)} B_{i,j_1}$ and $B_{i,j_1} = \{x_{i,j_1,j_2} : 1 \leq j_2 \leq b + 1\}$ as in Remark 6.9.

Without loss of generality, assume that the missing component is $c_{1,1,b+1} = \partial_m(x_{1,1,b+1})$, where $\partial_m \in \mathcal{D}$. Observe immediately that because both of $f \circ h$ and h are constant on $B_{1,1}$, the restriction $\partial_m|_{B_{1,1}}$ can be written as a polynomial of degree $\max\{\deg(h) - 2, \deg(h) - \lambda - 1\} = \deg(h) - 2 = b - 1$. Since $x_{1,1,j_2} \in B_{1,1}$ for each j_2 , we have that $\partial_m|_{B_{1,1}}(x_{1,1,j_2}) = \partial_m(x_{1,1,j_2}) = c_{1,1,j_2}$. Using Lagrange interpolation on the points $(x_{1,1,j_2}, c_{1,1,j_2})$ for $1 \leq j_2 \leq b$, we obtain a polynomial $\Delta^{B_{1,1}}$ of degree $b - 1$ which agrees with $\partial_m|_{B_{1,1}}$ at b distinct points, so the two polynomials must be equal. Thus we can recover $c_{1,1,b+1}$ by evaluating $\Delta^{B_{1,1}}$ at the element $x_{1,1,b+1}$.

Let us now consider the case of λ erasures (in the practical example we will take $\lambda = 2$, as that is the second most likely scenario of failures). Among these λ erasures, the erasures which are isolated in locality sets $B_{i,j}$ can be recovered by using the 1-locality, so the interesting case is when multiple erasures occur in a single $B_{i,j}$. Let us assume that $\lambda \geq 2$ erasures occur in the same locality set $B_{i,j}$. In this case, since $f \circ h$ is constant on A_i , the restriction $\partial_m|_{A_i}$ is a polynomial of degree $\deg(f) \deg(h) - \lambda - 1 = a + \lambda - \lambda - 1 = a - 1$.

Thus Lagrange interpolation on a set of a points of A_i on which no erasure occurred yields a polynomial Δ^{A_i} which agrees with ∂_m on all of A_i . Hence the missing components can be obtained by evaluating Δ^{A_i} at each of the corresponding locations in A_i .

6.3.4 Optimality of the Code

We dedicate this subsection to proving the optimality of our code \mathcal{C} . Therefore, we will be computing the values of k and d .

Lemma 6.11. *Let \mathcal{C} be the code in (6.3). Then*

$$k = (s + 1)((\deg(f) - 1)(\deg(h) - 1) + \deg(h) - \lambda).$$

Proof. Since in particular $\deg(g_{i,j}h^j) + \deg(\tilde{g}_i h^{\deg f - 1}) \leq \deg(f \circ h)$ and $\deg(g_{i,j}), \deg(\tilde{g}_i) \leq \deg(h)$, by uniqueness of F -adic expansion both for $F = f \circ h$ and for $F = h$, we have

$$k = \dim_{\mathbb{F}_q} \mathcal{D} = (s + 1)((\deg(f) - 1)(\deg(h) - 1) + (\deg(h) - \lambda)),$$

as we wanted to prove. □

Lemma 6.12. *Let \mathcal{C} be the code in (6.3). Then $d \geq n - \delta$, where*

$$\delta = (s + 1) \deg(f) \deg(h) - \lambda - 1.$$

Proof. A lower bound for the minimum distance is obtained by subtracting δ from n , where δ is the upper bound for the maximum number of zeros of $\partial \in \mathcal{D}$. We compute

$$\begin{aligned} \delta &= (\deg(f) \deg(h)s + (\deg(f) - 1) \deg(h) + \deg(h) - \lambda - 1) \\ &= (s + 1) \deg(f) \deg(h) - \lambda - 1, \end{aligned} \tag{6.4}$$

and this proves the claim. □

Theorem 6.13. *Let \mathcal{C} be the code obtained by using Construction 6.10. Then \mathcal{C} is an optimal $[n, k, b, a, \lambda]$ HLRC.*

Proof. Let $\rho = (a + \lambda)/(b + 1)b - (\lambda - 1)$ and $k_1 = k - 1 - \left\lfloor \frac{k-1}{\rho} \right\rfloor \rho$. Moreover, we recall that $a + \lambda = \deg(f) \deg(h)$ and $\deg(h) = b + 1$. Let d' denote the optimal distance, such that

$$\delta' = n - d' = \left(\left\lfloor \frac{k-1}{\rho} \right\rfloor (a + \lambda) + k_1 + \left\lfloor \frac{k_1}{b} \right\rfloor \right).$$

Note that

$$\left\lfloor \frac{k-1}{\rho} \right\rfloor = \left\lfloor \frac{1}{-\deg(f) \deg(h) + \deg(f) + \lambda - 1} + s + 1 \right\rfloor = s,$$

since $\lambda \leq \deg(h) - 1$, and

$$\left\lfloor \frac{k_1}{b} \right\rfloor = \left\lfloor \deg(f) - \frac{\lambda}{\deg(h) - 1} \right\rfloor = \deg(f) - \left\lceil \frac{\lambda}{\deg(h) - 1} \right\rceil,$$

in fact $k_1 = \deg(f)(-(b + 1)s + \deg(h)(s + 1) - 1) - \lambda = \deg(f)(\deg(h) - 1) - \lambda$. By using the results of Lemma 6.11, 6.12, we have

$$\begin{aligned} \delta - \delta' &= (s + 1) \deg(f) \deg(h) - \lambda - 1 - s \deg(f) \deg(h) - k_1 - \left\lfloor \frac{k_1}{b} \right\rfloor \\ &= \deg(f) \deg(h) - \lambda - 1 - (\deg(f)(\deg(h) - 1) - \lambda) - \deg(f) + \left\lceil \frac{\lambda}{\deg(h) - 1} \right\rceil \\ &= \left\lceil \frac{\lambda}{\deg(h) - 1} \right\rceil - 1, \end{aligned}$$

and since $\left\lceil \frac{\lambda}{\deg(h) - 1} \right\rceil - 1 = 0$ for $\lambda \leq \deg(h) - 1$, the code is optimal. \square

6.3.5 Comparison With Another Optimal Hierarchical RS-like Code

A construction of optimal HLRCs for a certain set of parameters is presented in [1, Proposition IV.2]. Let us fix the parameters for which that construction exists, i.e., $r_1 = sr_2$ (we note that we do not require such a constraint, but that even in this scenario we show that we can construct codes that are not available from [1, Proposition IV.2]). The set of parameters of the codes in [1, Proposition IV.2], also given in our notation, is as follows:

- the length of the codes in both settings is n ,
- each small locality set (at the bottom level of the hierarchy) has size $r_2 + 1$, so in our case each has size $b + 1$,
- their ν is our $a + \lambda$,
- the middle code has distance $r_2 + 3$ and hence can tolerate $r_2 + 2$ erasures, so their $r_2 + 2$ corresponds to our λ ,
- their r_1 is our ρ ,
- the code is optimal, with distance $d = n - t(r_1 + r_2 + 1 + s) + r_2 + 3$, for some t, s , and
- the two-level hierarchy has locality parameters $(r_1, r_2 + 3)$ and $(r_2, 2)$.

This shows immediately that our class of codes is different from the codes in [1, Proposition IV.2]. In fact, the optimality of our codes strongly relies on the assumption $\lambda \leq r_2$, which is not the case in the construction in [1, Proposition IV.2], in which instead $\lambda = r_2 + 2$. It follows that our class of codes contains codes which are not covered by this construction, as we can construct optimal hierarchical codes with two-level hierarchy which have locality parameters (r_1, λ) and $(r_2, 2)$ for any $\lambda \leq r_2$, such as for $\lambda = 2$.

We emphasize that in [33] it is necessary to set a fixed $\lambda = r_2 + 2$ since in this way one can reach optimality using the bound in [33, Theorem 2.1], while, using our improved bound and enhancing the construction in [1], one is allowed more flexibility as we explained. Moreover,

we will see in Section 6.4 how to construct our codes without the arithmetic restrictions appearing in the examples which use monomials or linearized functions.

For a better comparison and to simplify the presentation, in the next paragraph we will still use monomials for the toy example, even though it is not a requirement as we explain in Section 6.4.

6.3.6 Toy Example

Suppose one desires a code over \mathbb{F}_{19} of dimension 6 which can recover 1, 2, and 8 lost nodes by accessing at most 2, 4, and 6 other nodes, respectively (i.e., the distance of the code is equal to 9). This is not possible using the standard Tamo-Barg construction since, to recover more than 1 node, one would need to access as many nodes as the dimension of the code, that is, 6 nodes. Another option is to consider codes with availability using an orthogonal partition of the multiplicative group of \mathbb{F}_{19} that includes C_3 (as one wants the locality to be 3). But this does not work in this case either as the only other option is C_9 and $C_3 \subseteq C_9$ (since \mathbb{F}_q^* is cyclic for any prime power q). Moreover [33, Proposition IV.2] does not hold for $\lambda = 2$.

Our construction instead provides a code that allows these recovery capabilities and is information-theoretically optimal in the sense of the Singleton bound in Section 6.2.

Suppose we choose $f(X) = X^2$ and $h(X) = X^3$ (so $b = 2$ and $a = 4$). A general information polynomial is given by

$$\partial(X) = \sum_{i=0}^1 \left[g_i(X) + \tilde{g}_i(X)h(X) \right] f(h(X))^i,$$

for $g_i \in \mathbb{F}_q[X]_{\leq 1}$ and $\tilde{g}_i \in \mathbb{F}_q[X]_{\leq 0}$. In particular the $\tilde{g}_i(X) = \tilde{g}_i$ are constants (notice that the internal sum in j in (6.2) disappears since $\deg(f) = 2$). Therefore, by evaluating the messages at the preimage of the the 3 totally split places of $x^6 = f \circ h$, we get a code of length $n = 18$ and dimension $k = 6$ with $a = 4$ and $b = 2$. Notice that this code can recover 1 erasure by looking at $b = 2$ other nodes. Moreover, if two erasures occur, we have two

possibilities: either the erasures occur in the same nest for (f, h) , in which case one needs to access (in the worst case scenario) at most 4 other nodes, or the erasures occur in different nests, in which case one can use twice the locality (that is, 2) to recover each node so that one again needs to access at most 4 other nodes. Since we are evaluating polynomials of degree at most 9, the distance of the code is $18 - 9 = 9$ and therefore one has a fault tolerance of 8 erasures. Practically, given those 18 nodes, we are looking at the disposition of hierarchy in Figure 1. To simplify the presentation, we label each part of the hierarchy corresponding to $t_1 = 1$ only.

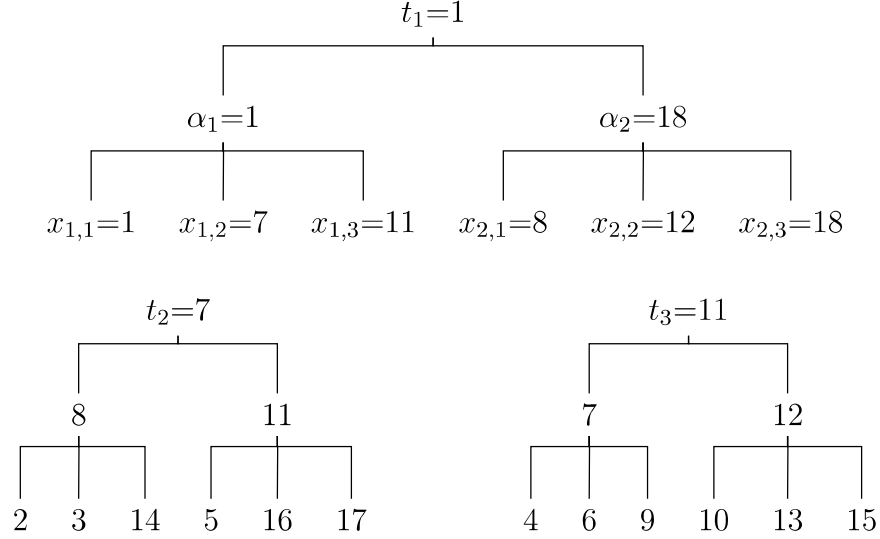


Figure 1. The nest hierarchy for $f(X) = X^2$ and $h(X) = X^3$.

The above corresponds to the following matrix:

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 7 & 11 & 8 & 12 & 18 & 2 & 3 & 14 & 5 & 16 & 17 & 4 & 6 & 9 & 10 & 13 & 15 \\ 1 & 1 & 1 & 18 & 18 & 18 & 8 & 8 & 8 & 11 & 11 & 11 & 7 & 7 & 7 & 12 & 12 & 12 \\ 1 & 1 & 1 & 1 & 1 & 1 & 7 & 7 & 7 & 7 & 7 & 7 & 11 & 11 & 11 & 11 & 11 & 11 \\ 1 & 7 & 11 & 8 & 12 & 18 & 14 & 2 & 3 & 16 & 17 & 5 & 6 & 9 & 4 & 15 & 10 & 13 \\ 1 & 1 & 1 & 18 & 18 & 18 & 18 & 18 & 18 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 18 & 18 & 18 \end{pmatrix},$$

where the rows correspond to (the evaluations of) the basis $\{1, x, x^3, x^6, x^7, x^9\}$ and the columns to the elements of \mathbb{F}_{19}^* ordered as in Figure 1. This means that to check the locality of each set one just needs to check the rank of the corresponding set of columns in the above matrix. For example, suppose we want to recover the third column, which corresponds to the symbol 11. We can do that using only the first two columns, since the matrix

$$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 7 & 11 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 7 & 11 \\ 1 & 1 & 1 \end{pmatrix}$$

has rank equal to 2. Similarly, we can recover any two lost symbols using either 3 (if they belong to the same large orbit) or 4 (otherwise, if they belong to distinct large orbits) other symbols.

6.4 Existential Results via Chebotarev Density Theorem

In this section we explain how to apply Chebotarev Density Theorem to count the places $t_0 \in \mathbb{F}_q$ such that $f(h) - t_0$ is totally split. A lower bound on this quantity determines directly a lower bound on the size of the hierarchy in our construction. This determines completely the range of parameters of our hierarchical codes, and in turn it shows that they always exist for q large enough, without arithmetic restrictions on the localities and the size of the base field.

6.4.1 The Number of Totally Split Places t_0 of $f(h) - t$

We will appeal to the Chebotarev Density Theorem as in Proposition 3.1 of [29] since this formulation is the most convenient for our purposes. We provide a full exposition in this section, but we briefly describe in the next paragraph the general procedure and ideas.

For polynomials $f, h \in \mathbb{F}_q[X]$, consider the composition $f(h)$. By the lower bound in [29, Proposition 3.1] on the number ℓ of $t_0 \in \mathbb{F}_q$ such that $f(h) - t_0$ splits into linear factors over \mathbb{F}_q , we have that for large enough q it is guaranteed to have a large number of totally split places of degree 1 of $\mathbb{F}_q(x)/\mathbb{F}_q(t)$ when $f(h)$ is chosen correctly. Now, we may assume that the field of constants $k_{f(h)}$ of $M_{f(h)}$ is trivial since otherwise there cannot be a totally split place of degree 1. Since we want ℓ to be as large as possible, one quickly sees from the lower bound in [29, Proposition 3.1] that minimizing the size of the monodromy group $G_{f(h)}$ of $f(h)$ achieves this goal. Thus our construction always effectively results in an optimal code as long as the size of the alphabet verifies a certain lower bound.

For the extension $M/\mathbb{F}_q(t)$, let $G = \text{Gal}(M/\mathbb{F}_q(t))$ be its arithmetic Galois group and let N be its geometric Galois group. Since we are interested in the number ℓ of places $P \subseteq \mathbb{F}_q(t)$ of degree 1 which are totally split in M , by Proposition 3.4 of [29] we may assume that $M \cap \overline{\mathbb{F}_q} = \mathbb{F}_q$ is the field of constants of the extension $M/\mathbb{F}_q(t)$ since otherwise $\ell = 0$. Hence $G = N$.

Lemma 6.14. *Let $f, h \in \mathbb{F}_q[X]$ be nonzero polynomials having positive degrees. Define $G_f = \text{Gal}(f(X) - t/\mathbb{F}_q(t))$ and similarly for G_h . Then the number of $t_0 \in \mathbb{F}_q$ such that $f(h(X)) - t_0$ splits completely into distinct (linear) factors over \mathbb{F}_q is at least $\frac{1}{|G_h|^{\deg(f)}|G_f|}q + O(\sqrt{q})$, where the implied constant can be chosen explicitly and is independent of q .*

Proof. Denoting the number of $t_0 \in \mathbb{F}_q$ we are considering by $|T_{split}^1(f \circ h)|$, from Proposition 3.1(ii) of [29] we immediately have

$$|T_{split}^1(f \circ h)| \geq \frac{q + 1 - 2g\sqrt{q}}{|G|} - \frac{|\text{Ram}^1(M : \mathbb{F}_q(t))|}{2}. \quad (6.5)$$

We proceed by proving an upper bound on the size of G , which in turn gives the wanted lower bound for $|T_{split}^1(f \circ h)|$

Let \mathcal{T} be the rooted tree of height 2 with $\deg(f)$ branches and $\deg(h)$ roots adjacent to each branch. One can easily see that $G \subseteq \text{Aut}(\mathcal{T})$, so because $\text{Aut}(\mathcal{T})$ is isomorphic to the wreath product $\underbrace{(G_h \times \cdots \times G_h)}_{\deg(f)} \rtimes G_f$, we have $|G| \leq |G_h|^{\deg(f)} |G_f|$.

Combining (6.5) with the bound on $|G|$, we obtain

$$\ell \geq \frac{q+1-2g\sqrt{q}}{|G_h|^{\deg(f)}|G_f|} - \frac{|\text{Ram}^1(M : \mathbb{F}_q(t))|}{2}. \quad \square$$

Note that the bound given in the previous lemma can be written more explicitly as

$$\ell \geq \frac{(q+1)-2g\sqrt{q}}{|G_h|^{\deg(f)}|G_f|} - \frac{\deg(f)\deg(h)}{2}.$$

Proposition 6.15. *Let $f, h \in \mathbb{F}_q[x]$ be polynomials such that $f(h)-t$ has Galois group G and the splitting field M of $f(h)-t$ has constant field equal to \mathbb{F}_q . Then there exists an optimal HLRC with parameters $[\deg(f(h))\ell, k, d, \deg(h)-1, \deg(f(h))-\lambda, \lambda]$ for any $\lambda < \deg(h)$,*

$$\ell \geq \frac{q}{|G_h|^{\deg(f)}|G_f|} + O(\sqrt{q}),$$

where the implied constant can be made explicit, G_f (resp. G_h) is the Galois group of $f-t$ (resp. $h-t$), and k is as in Lemma 6.11.

Remark 6.16. Notice that the condition of having trivial constant field extension is automatic once there is a single totally split place, and it is the generic situation if the polynomials are chosen at random.

Proof. Since M has trivial constant field \mathbb{F}_q , Lemma 6.14 guarantees that there exist at least

$$\ell \geq \frac{1}{|G_h|^{\deg(f)}|G_f|} q + O(\sqrt{q})$$

totally split places, i.e., elements t_0 of \mathbb{F}_q such that $f(h) - t_0$ is totally split. Now Lemma 6.6 guarantees that the evaluation set T consisting of the preimages of the t_0 's forms a nest for the pair (f, h) (see Remark 6.9). Construct the code by evaluating the polynomials in (6.2) at the subset \mathcal{A} of preimages of T via $f(h)$, i.e. $\mathcal{A} = (f \circ h)^{-1}(T)$, which has size $\deg(f(h))\ell$. The hierarchy is now given by the nest structure in the sense of Remark 6.9 and the parameters obtained from Section 6.3. \square

6.5 Practical Choice of Parameters to Construct Optimal HLRCs

The construction we presented in the previous sections allows us to exhibit some interesting examples of HLRCs. To begin with, we consider the case \mathbb{F}_{64} . Choosing f and h such that $\deg(f) = \deg(h) = 3$ and $\ell = 7$, our construction gives rise to a $(63, k, d, 2, 5, 2)$ HLRC, where the values of k and d depend on the choice of s in Construction 6.10. In fact, the first locality b equals $\deg(h) - 1$, whereas the second locality ($a = 5$) can be computed by following the passages of Section 6.3.3. This means that we are able to recover 1 (resp. 2) lost node(s) by looking at 2 (resp. 5) other nodes. We point out that the Tamo-Barg construction for availability over the field of size 64, under the same first locality assumption ($b = 2$), forces to have length 21 (with locality sets of size 3 and 7), whereas ours permits to have length 63, leading to a much better minimum distance and a larger number of servers allowed. More precisely, the Tamo-Barg construction requires the use of two orthogonal partitions, and this can be achieved by using 21 symbols corresponding to the action of x^3 and of x^7 on $\mathbb{F}_{64} \setminus \{0\}$. Note further that their construction has a larger second locality: 7, against our better parameter $a = 5$.

REFERENCES

- [1] Sean Ballentine, Alexander Barg, and Serge Vlăduț. “Codes with hierarchical locality from covering maps of curves”. In: *IEEE Transactions on Information Theory* 65.10 (2019), pp. 6056–6071.
- [2] Alexander Barg, Itzhak Tamo, and Serge Vlăduț. “Locally recoverable codes on algebraic curves”. In: *IEEE Transactions on Information Theory* 63.8 (2017), pp. 4928–4939.
- [3] Alexander Barg et al. “Locally recoverable codes from algebraic curves and surfaces”. In: *Algebraic Geometry for Coding Theory and Cryptography*. Springer, 2017, pp. 95–127.
- [4] Daniele Bartoli and Marco Calderini. “On construction and (non) existence of c - (almost) perfect nonlinear functions”. In: *Finite Fields and Their Applications* 72 (2021), p. 101835.
- [5] Daniele Bartoli and Giacomo Micheli. “Algebraic Constructions of Complete m -Arcs”. In: *Comb.* 42.5 (2022), pp. 673–700. DOI: [10.1007/s00493-021-4712-5](https://doi.org/10.1007/s00493-021-4712-5). URL: <https://doi.org/10.1007/s00493-021-4712-5>.
- [6] Daniele Bartoli, Maria Montanucci, and Luciane Quoos. “Locally recoverable codes from automorphism group of function fields of genus $g \geq 1$ ”. In: *IEEE Transactions on Information Theory* 66.11 (2020), pp. 6799–6808.
- [7] Daniele Bartoli, Giovanni Zini, and Ferdinando Zullo. “Investigating the exceptionality of scattered polynomials”. In: *arXiv preprint arXiv:2103.04591* (2021).
- [8] Daniele Bartoli et al. “ r -fat linearized polynomials over finite fields”. In: *arXiv preprint arXiv:2012.15357* (2020).

- [9] E. R. Berlekamp. “An analog to the discriminant over fields of characteristic two”. In: *J. Algebra* 38.2 (1976), pp. 315–317. ISSN: 0021-8693. DOI: [10.1016/0021-8693\(76\)90222-2](https://doi.org/10.1016/0021-8693(76)90222-2). URL: [https://doi.org/10.1016/0021-8693\(76\)90222-2](https://doi.org/10.1016/0021-8693(76)90222-2).
- [10] Stephen D. Cohen and Rex W. Matthews. “Monodromy groups of classical families over finite fields”. In: *Finite fields and applications (Glasgow, 1995)*. Vol. 233. London Math. Soc. Lecture Note Ser. Cambridge Univ. Press, Cambridge, 1996, pp. 59–68. DOI: [10.1017/CB09780511525988.007](https://doi.org/10.1017/CB09780511525988.007). URL: <https://doi.org/10.1017/CB09780511525988.007>.
- [11] Austin Dukes, Andrea Ferraguti, and Giacomo Micheli. “Optimal selection for good polynomials of degree up to five”. In: *Designs, Codes and Cryptography* 90.6 (2022), pp. 1427–1436.
- [12] Austin Dukes and Giacomo Micheli. “Understanding Polynomial Maps over Finite Fields”. In: *cryptTORino 2021*. Ed. by Laura Capuano et al. Collectio CiphRARum. Rome: Aracne, 2023, pp. 39–42.
- [13] Austin Dukes, Giacomo Micheli, and Vincenzo Pallozzi Lavorante. “Optimal Locally Recoverable Codes With Hierarchy From Nested F-Adic Expansions”. In: *IEEE Transactions on Information Theory* 69.11 (2023), pp. 6981–6988. DOI: [10.1109/TIT.2023.3298401](https://doi.org/10.1109/TIT.2023.3298401).
- [14] Andrea Ferraguti and Giacomo Micheli. “Exceptional scatteredness in prime degree”. In: *Journal of Algebra* 565 (2021), pp. 691–701.
- [15] Andrea Ferraguti and Giacomo Micheli. “Full classification of permutation rational functions and complete rational functions of degree three over finite fields”. In: *Des. Codes Cryptogr.* 88.5 (2020), pp. 867–886. ISSN: 0925-1022. DOI: [10.1007/s10623-020-00715-0](https://doi.org/10.1007/s10623-020-00715-0). URL: <https://doi.org/10.1007/s10623-020-00715-0>.
- [16] Andrea Ferraguti, Carlo Pagano, and Daniele Casazza. *The inverse problem for arboreal Galois representations of index two*. <https://arxiv.org/abs/1907.08608>. 2019.

- [17] Ragnar Freij-Hollanti, Thomas Westerbäck, and Camilla Hollanti. “Locally repairable codes with availability and hierarchy: matroid theory via examples”. In: *International Zurich Seminar on Communications-Proceedings*. ETH Zurich. 2016, pp. 45–49.
- [18] P. X. Gallagher. “The large sieve and probabilistic Galois theory”. In: *Analytic number theory (Proc. Sympos. Pure Math., Vol. XXIV, St. Louis Univ., St. Louis, Mo., 1972)*. 1973, pp. 91–101.
- [19] P. Gopalan et al. “On the Locality of Codeword Symbols”. In: *IEEE Transactions on Information Theory* 58.11 (2012), pp. 6925–6934. DOI: [10.1109/TIT.2012.2208937](https://doi.org/10.1109/TIT.2012.2208937).
- [20] Parikshit Gopalan et al. “Explicit Maximally Recoverable Codes with Locality.” In: *IEEE Trans. Information Theory* 60.9 (2014), pp. 5245–5256.
- [21] Parikshit Gopalan et al. “On the locality of codeword symbols”. In: *IEEE Transactions on Information theory* 58.11 (2012), pp. 6925–6934.
- [22] Thomas W. Hungerford. *Algebra*. 1st. Graduate Texts in Mathematics. Springer New York, NY, 1980. ISBN: 9780387905181.
- [23] Govinda M Kamath et al. “Codes with local regeneration and erasure correction”. In: *IEEE Transactions on information theory* 60.8 (2014), pp. 4637–4660.
- [24] Michiel Kusters. “A short proof of the Chebotarev density theorem for function fields”. In: *Mathematical Communications* 22.2 (2017), pp. 227–233.
- [25] J. H. van Lint. *Introduction to Coding Theory*. 3rd. Graduate Texts in Mathematics. Springer-Verlag, 1998. ISBN: 9783642636530.
- [26] Jian Liu, Sihem Mesnager, and Lusheng Chen. “New Constructions of Optimal Locally Recoverable Codes via Good Polynomials”. In: *IEEE Transactions on Information Theory* 64.2 (2018), pp. 889–899. DOI: [10.1109/TIT.2017.2713245](https://doi.org/10.1109/TIT.2017.2713245).
- [27] Jian Liu, Sihem Mesnager, and Deng Tang. “Constructions of optimal locally recoverable codes via Dickson polynomials”. In: *Designs, Codes and Cryptography* (2020), pp. 1–22.

- [28] Sihem Mesnager et al. “Investigations on c-(Almost) Perfect Nonlinear Functions”. In: *IEEE Transactions on Information Theory* 67.10 (2021), pp. 6916–6925. DOI: [10.1109/TIT.2021.3081348](https://doi.org/10.1109/TIT.2021.3081348).
- [29] Giacomo Micheli. “Constructions of locally recoverable codes which are optimal”. In: *IEEE Trans. Inform. Theory* 66.1 (2020), pp. 167–175. ISSN: 0018-9448. DOI: [10.1109/TIT.2019.2939464](https://doi.org/10.1109/TIT.2019.2939464). URL: <https://doi.org/10.1109/TIT.2019.2939464>.
- [30] Giacomo Micheli. “On the selection of polynomials for the DLP quasi-polynomial time algorithm for finite fields of small characteristic”. In: *SIAM J. Appl. Algebra Geom.* 3.2 (2019), pp. 256–265. DOI: [10.1137/18M1177196](https://doi.org/10.1137/18M1177196). URL: <https://doi.org/10.1137/18M1177196>.
- [31] Jürgen Neukirch. *Algebraic Number Theory*. Vol. Grundlehren der mathematischen Wissenschaften. Springer, 1999.
- [32] D. S. Papailiopoulos and A. G. Dimakis. “Locally Repairable Codes”. In: *IEEE Transactions on Information Theory* 60.10 (2014), pp. 5843–5855. DOI: [10.1109/TIT.2014.2325570](https://doi.org/10.1109/TIT.2014.2325570).
- [33] Birenjith Sasidharan, Gaurav Kumar Agarwal, and P Vijay Kumar. “Codes with hierarchical locality”. In: *2015 IEEE International Symposium on Information Theory (ISIT)*. IEEE. 2015, pp. 1257–1261.
- [34] Natalia Silberstein et al. “Optimal locally repairable codes via rank-metric codes”. In: *2013 IEEE International Symposium on Information Theory*. IEEE. 2013, pp. 1819–1823.
- [35] Henning Stichtenoth. *Algebraic Function Fields and Codes*. 2nd. Springer Publishing Company, Incorporated, 2008. ISBN: 3540768777.

- [36] Itzhak Tamo and Alexander Barg. “A Family of Optimal Locally Recoverable Codes”. In: *IEEE Transactions on Information Theory* 60.8 (Aug. 2014), pp. 4661–4676. ISSN: 1557-9654. DOI: [10.1109/tit.2014.2321280](https://doi.org/10.1109/TIT.2014.2321280). URL: <http://dx.doi.org/10.1109/TIT.2014.2321280>.
- [37] Itzhak Tamo, Alexander Barg, and Alexey Frolov. “Bounds on the parameters of locally recoverable codes”. In: *IEEE Transactions on information theory* 62.6 (2016), pp. 3070–3083.
- [38] Itzhak Tamo, Dimitris S Papailiopoulos, and Alexandros G Dimakis. “Optimal locally repairable codes and connections to matroid theory”. In: *IEEE Transactions on Information Theory* 62.12 (2016), pp. 6661–6671.

**APPENDIX A:
THE GENUS OF A FUNCTION FIELD**

This appendix is dedicated to formally defining the genus of a function field. We closely follow the exposition from [35].

Let F/K be a function field with full constant field K , where K is algebraically closed. Recall that a place $P \subseteq F$ of F/K is the unique maximal ideal of a valuation ring $\mathcal{O} \subsetneq F$ in F and that given P we can define a discrete valuation $v_P : F \rightarrow \mathbb{N} \cup \{\infty\}$.

Definition A.1. The divisor group $\text{Div}(F/K)$ of F/K is defined as the (additively written) free abelian group which is generated by the places of F/K . The elements of $\text{Div}(F/K)$ are called divisors of F/K . More explicitly, a divisor $D \in \text{Div}(F/K)$ is a formal sum

$$D = \sum_{P \in \mathbb{P}_F} n_P P,$$

where each $n_P \in \mathbb{Z}$ and only finitely many n_P are nonzero.

For $Q \in \mathbb{P}_F$ and $D = \sum n_P P \in \text{Div}(F)$, we define $v_Q(D) = n_Q$.

A partial ordering on $\text{Div}(F)$ is defined by

$$D_1 \leq D_2 \iff v_P(D_1) \leq v_P(D_2) \text{ for all } P \in \mathbb{P}_F.$$

The degree of a divisor is defined to be

$$\deg(D) = \sum_{P \in \mathbb{P}_F} v_P(D) \cdot \deg(P).$$

A divisor $D \in \text{Div}(F)$ is called principal if there is some $x \in F$ such that $D = \sum_{P \in \mathbb{P}_F} v_P(x)P$. If such an element $x \in F$ exists, we will write $D = (x)$.

Our next definition is of great import in algebraic function field theory.

Definition A.2. For a divisor $D \in \text{Div}(F)$, we define the Riemann-Roch space associated to D by

$$\mathcal{L}(D) = \{x \in F : (x) \geq -D\} \cup \{0\}.$$

Given a divisor $D \in \text{Div}(F)$, the Riemann-Roch space $\mathcal{L}(D)$ forms a vector space over K . In particular, this vector space is of finite dimension.

Definition A.3. For $D \in \text{Div}(F)$, we define the dimension $\ell(D)$ of the divisor D by $\ell(D) = \dim_K \mathcal{L}(D)$.

Now, for a divisor $D \in \text{Div}(F)$ we have defined both the degree and the dimension of D . The following proposition shows that the degree $\deg(D)$ and the dimension $\ell(D)$ cannot be “too far apart,” and this property is independent of the choice of D .

Proposition A.4 ([35, Proposition 1.4.14]). *There is a constant $\gamma \in \mathbb{Z}$ such that for all divisors $D \in \text{Div}(F)$ the following holds:*

$$0 \leq \deg(D) - \ell(D) + 1 \leq \gamma.$$

We are now ready to define the genus of a function field.

Definition A.5. The genus g of the function field F/K is defined to be

$$g = \max_{D \in \text{Div}(F)} \{\deg(D) - \ell(D) + 1\}.$$