Low Field Size Constructions of Access-Optimal Convertible Codes

Saransh Chopra

CMU-CS-24-106

May 2024

Computer Science Department
School of Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Rashmi Vinayak, Chair
Ryan O'Donnell

Submitted in partial fulfillment of the requirements for the degree of Master of Science.

Keywords: Coding theory, distributed storage systems, redundancy tuning, code conversion, super-regularity

Abstract

Most large-scale storage systems employ erasure coding to provide resilience against disk failures. Recent work has shown that tuning this redundancy to changes in disk failure rates leads to substantial storage savings. This process requires *code conversion*, wherein data encoded using an $[n^I, k^I]$ initial code has to be transformed into data encoded using an $[n^F, k^F]$ final code. *Convertible codes* are a class of codes that enable efficient code conversion while maintaining other desirable properties. In this thesis, we focus on the *access cost* of conversion (corresponding to the total number of symbols accessed in the conversion process) and on an important subclass of conversions known as the merge regime (corresponding to combining multiple initial codewords into a single final codeword).

In this setting, explicit constructions are known for systematic access-optimal Maximum Distance Separable (MDS) convertible codes for all parameters in the merge regime. However, the existing construction for a key subset of these parameters, which makes use of Vandermonde parity matrices, requires a very large field making it unsuitable for practical applications. In this thesis, we provide (1) sharper bounds on the minimum field size requirement for such codes, and (2) explicit constructions for low field sizes for several parameter ranges. In doing so, we provide a proof of super-regularity of specially designed classes of Vandermonde matrices that could be of independent interest.

Contents

Chapter 1

Introduction

Erasure codes serve as a cornerstone of modern large scale distributed storage systems as a means to mitigate data loss in the event of disk failures. In this context, erasure coding involves dividing data into groups of k chunks that are each encoded into stripes of n chunks using an $[n, k]$ erasure code. These encoded chunks are then stored across n distinct storage nodes in the system. The code parameters n and k determine the amount of redundancy added to the system and the degree of durability guaranteed.

There are various classes of codes that are frequently used in real-world systems. For example, *systematic* codes are those in which the original message symbols are embedded among the code symbols. This is highly desirable in practice as in the event that there are no observed disk failures, there is no decoding process needed to recover the original data. Systematic codes with *Vandermonde parity matrices* (see §2.1) are even more advantageous as there are known efficient algorithms utilizing Fast Fourier Transform (FFT) for computing the product between vectors and Vandermonde matrices [5, 12], speeding up the encoding process. This attribute is becoming increasingly important given the recent trend to use wider (high k) and longer (high n) erasure codes [6, 10]. Additionally, *Maximum Distance Separable (MDS)* codes are a subset of erasure codes that require the least amount of additional storage in order to meet a specific failure tolerance. An $[n, k]$ MDS code can tolerate loss of any $n - k$ out of the n code symbols. In this thesis, our interest is on systematic MDS codes with Vandermonde parity matrices.

Recent findings by Kadekodi et al. [9] reveal the dynamic variability in disk failure rates over time due to changes in data usage patterns and hardware reliability. Their research highlights the potential for meaningful savings in storage and associated operational expenses through tuning code parameters to observed failure rates. However, the resource overhead associated with the *default approach* of re-encoding all of the data in order to modify n and k is prohibitively expensive [15].

The *code conversion* problem introduced in [15] formalizes the problem of efficiently transforming data that has been encoded under an $[n^I, k^I]$ initial code $\mathcal{C}^I$ to its new representation under an $[n^F, k^F]$ final code $\mathcal{C}^F$. One of the key measures of the cost of conversion is the *access cost*, which represents the total number of code symbols accessed (read/written) during conversion.

Convertible codes [15] are a class of codes that enable efficient conversion while maintaining other desirable properties such as being MDS and systematic (more details in §2.2).

Among various types of conversions, the *merge regime*, where $k^F = \lambda k^I$ for any integer $\lambda \geq 2$ (i.e., combining multiple initial codewords into a single final codeword), is the most important one. First, the merge regime requires the least resource utilization [17] among all types of conversions and hence are a highly favorable choice for practical systems. Second, constructions for the merge regime are key building blocks for the constructions for codes in the *general regime* which allows for any set of initial parameters and any set of final parameters [17]. In this thesis, our focus is on systematic MDS convertible codes in the merge regime.

In [15], the authors established lower bounds on the access cost of conversion and provided constructions of *access-optimal* convertible codes for all parameters in the merge regime. Let us denote $r^I := n^I - k^I$ and $r^F := n^F - k^F$, (which correspond to the number of parity symbols in the initial and final codes if the codes are systematic). For cases where $r^I > r^F$ (i.e., when the initial configuration has more parities than the final configuration),the authors provide explicit constructions of systematic MDS access-optimal convertible codes over fields of size linear in n^F. For cases where $r^I < r^F$ (i.e., when more parities are needed in the final configuration than in the initial), it was shown that the access cost of conversion for MDS erasure codes is lower bounded by that of the default approach to decode and re-encode all of the data. As a consequence, it is not possible to realize any savings with specialized code constructions.

However, in the case where $r^I = r^F$, the best-known construction requires a minimum field size of p^D for any prime p and some $D \in \Theta((n^F)^3)$ [15]. This field size is far too high for efficient practical implementations. Most current instruction-set architectures are optimized to operate on bytes of data at a time. Utilizing erasure codes defined over larger field sizes can hamper the encoding/decoding speed. Hence most (if not all) practical implementations of storage codes use $\mathbb{F}_{256}$ (which translates each field symbol to a one-byte representation). Thus, the problem of constructing low-field-size access-optimal convertible codes remains open for the case $r^I = r^F$.

In this thesis, we study the setting of systematic MDS access-optimal convertible codes in the merge regime in the case where $r^I = r^F$. Previously, the best known construction of codes in this setting was a systematic code with a very specific choice of Vandermonde parity matrix with a singular degree of freedom. In Chapter 3, we improve on this construction by allowing more freedom in the choice of *scalars* of the Vandermonde matrix. We then study the minimum field size, denoted $q^*(k, r)$, required for existence of the underlying $k \times r$ super-regular Vandermonde parity matrices of such codes (as will be detailed in §2.1). We establish the lower bounds $q^*(k, r) \geq \Omega(kr)$ (Theorem 1) and $q^*(k, r) \geq \Omega(2^r)$ (Theorem 2), with the latter bound pertaining to codes over fields of characteristic 2 where $k > r$. The first bound is tighter for regimes where $k \gg r$, while the second bound is tighter when $k \approx r$. Additionally, we establish an upper bound $q^*(k, r) \leq O(k^r)$ (Theorem 3), which in turn results in an improved upper bound $q \leq O((k^F)^{r^F})$ on the field size required for the existence systematic MDS access-optimal convertible codes in the merge regime in the case where $r^I = r^F$.

Furthermore, in Chapter 4, we provide the first explicit low-field-size constructions of codes in this setting for several parameter ranges via constructing their corresponding super-regular Vandermonde parity matrices. The proposed construction makes use of field automorphisms in designing the Vandermonde matrices. For any finite field $\mathbb{F}_q$ where $q = 2^w$, we find explicit constructions of $k \times 3$ super-regular Vandermonde matrices for all k such that $k < q$ (Theorem 4). This, in turn, gives us a construction of systematic MDS access-optimal convertible codes for all parameters in the merge regime such that $r^F = r^I \leq 3$ and $k^F < q$. We similarly resolve the regime over any general prime power field $\mathbb{F}_q$ where $q = p^w$, finding constructions of $k \times 3$ super-regular Vandermonde matrices for all k such that $k < w$ (Theorem 5).

These results are also of independent interest beyond the setting considered in this thesis as systematic MDS codes with Vandermonde parity matrices serve as the base codes for *bandwidth-optimal* convertible codes [14, 16] and have also been studied in various other settings [12, 19, 21].

Chapter 2

Background and Related Work

Let us begin with an overview of important concepts and notation referred to throughout this thesis, along with a literature review of previous related work.

2.1 Systematic MDS codes and Vandermonde matrices

An $[n, k]$ linear erasure code $\mathcal{C}$ with generator matrix $\mathbf{G} \in \mathcal{M}(\mathbb{F})_{k\times n}$ over a finite field $\mathbb{F}$ is said to be systematic, or in standard form, if $\mathbf{G} = [\mathbf{I_k} \mid \mathbf{P}]$ where $\mathbf{I_k}$ is the $k \times k$ identity matrix and $\mathbf{P}$ is a $k \times (n-k)$ matrix also known as the parity matrix. Let $\mathbf{m}$ be a message and $\mathbf{c}$ be its corresponding codeword under $\mathcal{C}$, where $\mathbf{m} = (m_i)_{i=1}^{k}$ and $\mathbf{c} = (c_i)_{i=1}^{n}$ are vectors of message and code symbols, respectively. As $\mathbf{m}$ is encoded under $\mathcal{C}$ via the multiplication $\mathbf{c} = \mathbf{m}^T\mathbf{G}$, it follows that $c_i = m_i$ for all $i \leq k$ if $\mathcal{C}$ is systematic.

An $[n, k]$ linear erasure code $\mathcal{C}$ is Maximum Distance Separable (MDS) if and only if every k columns of its generator matrix $\mathbf{G}$ are linearly independent; in other words, every $k \times k$ submatrix of $\mathbf{G}$ is non-singular [13]. As a result, data encoded by an $[n, k]$ MDS code can withstand any erasure pattern of $n - k$ out symbols in any codeword and still successfully recover the original data. If $\mathcal{C}$ is also systematic with parity matrix $\mathbf{P}$, this is equivalent to the property that every square submatrix of $\mathbf{P}$ is non-singular [13]. Such a matrix is also referred to as *super-regular*. It is useful to note that any submatrix of a super-regular matrix is also super-regular.

A systematic code with a Vandermonde parity matrix $\mathbf{P} \in \mathcal{M}(\mathbb{F}_{k\times r})$ is one where $\mathbf{P}$ is of the form

$$\begin{bmatrix} 1 & 1 & \dots & 1 \\ \xi_1 & \xi_2 & \dots & \xi_r \\ \xi_1^2 & \xi_2^2 & \dots & \xi_r^2 \\ \vdots & \vdots & \ddots & \vdots \\ \xi_1^{k-1} & \xi_2^{k-1} & \dots & \xi_r^{k-1} \end{bmatrix} \tag{2.1}$$

for some *scalars* $\xi = (\xi_i)_{i=1}^{r} \in \mathbb{F}^r$. Let us denote the above $k \times r$ Vandermonde matrix as $V_k(\xi)$. Such a matrix is not always guaranteed to be super-regular [13] and thus careful selection of the scalars is required to ensure the resulting systematic code is MDS.

2.2 Convertible Codes [15]

Recall that a *code conversion* is a procedure that converts data from its initial representation under an $[n^I, k^I]$ code $\mathcal{C}^I$ to its final representation under an $[n^F, k^F]$ code $\mathcal{C}^F$. In order to capture the potential change in dimension, let us denote $M := \mathrm{lcm}(k^I, k^F)$ and consider any message $\mathbf{m} \in \mathbb{F}_q^M$. This is equivalent to $\lambda^I := \frac{M}{k^I}$ stripes in the initial dimension and $\lambda^F := \frac{M}{k^F}$ stripes in the final dimension. Let $[i] := \{1, 2, \ldots, i\}$ and let $|S|$ denote the size of a set S. Let $\mathbf{m}[S]$ be the vector formed by projecting $\mathbf{m}$ onto the coordinates in the set S, and let $\mathcal{C}(\mathbf{m})$ stand for the encoding of $\mathbf{m}$ under the code $\mathcal{C}$. Let $r^I := n^I - k^I$ and $r^F := n^F - k^F$.

Definition 1 (Convertible Code [15])**.** *An $(n^I, k^I; n^F, k^F)$ convertible code over $\mathbb{F}_q$ is defined by: (1) a pair of codes $(\mathcal{C}^I, \mathcal{C}^F)$ over $\mathbb{F}_q$ such that $\mathcal{C}^I$ is an $[n^I, k^I]$ code and $\mathcal{C}^F$ is an $[n^F, k^F]$ code; (2) a pair of partitions $\mathcal{P}^I := \{P_i^I \mid i \in [\lambda^I]\}$ and $\mathcal{P}^F := \{P_j^F \mid j \in [\lambda^F]\}$ of $[M = \mathrm{lcm}(k^I, k^F)]$ such that $|P_i^I| = k^I$ for all $P_i^I \in \mathcal{P}^I$ and $|P_j^F| = k^F$ for all $P_j^F \in \mathcal{P}^F$; and (3) a conversion procedure which, for any $m \in \mathbb{F}_q^M$, maps the initial set of codewords $\{\mathcal{C}^I(\mathbf{m}[P_i^I]) \mid P_i^I \in \mathcal{P}^I\}$ to the corresponding set of codewords $\{\mathcal{C}^F(\mathbf{m}[P_j^F]) \mid P_j^F \in \mathcal{P}^F\}$ over the final code.*

Recall that access cost during code conversion refers to the number of code symbols that are read or written during conversion. Access-optimal convertible codes are those which meet the lower bounds on access cost established in [15] that are known to be tight. It is known that *any* $(n^I, k^I; n^F, k^F)$ convertible code for the merge regime where $r^I = r^F$ formed by a pair of systematic codes with Vandermonde parity matrices $\mathbf{P^I} = V_{k^I}(\xi)$ and $\mathbf{P^F} = V_{k^F}(\xi)$ over the same scalars is access-optimal [15]. This is due to $\mathbf{P^F}$ being r^F*-column block-constructible* from $\mathbf{P^I}$; in other words, each new parity of a merged codeword can directly be computed as a linear combination of the parities of the original codewords. If the parity matrices are super-regular, then the resulting convertible code is guaranteed to be MDS as well. The best known construction of a systematic MDS access-optimal convertible code for the merge regime where $r^I = r^F$ is formed by a pair of systematic codes with Vandermonde parity matrices over the scalars $\xi = (\theta^{i-1})_{i=1}^{r^I}$, for any primitive element $\theta \in \mathbb{F}$.

2.3 Additional Notation and Preliminaries

This section reviews terminology and notation used in this thesis that expands on the notation introduced in [15].

For any two sets I, J, let $I \triangle J$ denote the symmetric difference of I and J. For any two integers a, b, let $a \perp b$ denote that a and b are coprime. Let $\mathbf{x}$ denote the vector $(x_i)_{i=1}^r$ for some r. Let $\mathbf{M}_{i,j}$ denote the entry in the ith row and jth column of the matrix $\mathbf{M}$, with both indices 1-indexed. Let $\mathbf{M}_{I \times J}$ denote the submatrix of $\mathbf{M}$ formed by the intersection of the rows indexed by I and the columns indexed by J, with all indices 1-indexed. Let $\mathrm{row}_i(\mathbf{M})$ stand for the ith row vector of the matrix $\mathbf{M}$. Let χ_P be the indicator function for whether the proposition P is true.

Let $\mathbb{F}_p$ denote the prime field of size p, and let us reserve $\mathbb{F}_q$ for prime power fields of size $q = p^w$ for some prime p and $w > 1$. Let $\mathbb{F}^\times$ denote the multiplicative group of the field, or $\mathbb{F} \setminus \{0\}$. Let $\mathrm{ord}(a)$ denote the order of an element $a \in \mathbb{F}^\times$. Let $\mathbb{F}[x_1, \ldots, x_r]$ denote the ring of

polynomials in $x_1, \ldots, x_r$ over the field $\mathbb{F}$. Let $\mathrm{Aut}(\mathbb{F})$ denote the group of automorphisms over the field $\mathbb{F}$. Let S_n denote the group of permutations of $[n]$.

Recall that a field automorphism is a bijective map $\sigma : \mathbb{F} \to \mathbb{F}$ such that for all $x, y \in \mathbb{F}$, $\sigma(x + y) = \sigma(x) + \sigma(y)$ and $\sigma(xy) = \sigma(x)\sigma(y)$; in essence, the map preserves the structure of the field. Note also by definition, it must be the case that $\sigma(0) = 0$ and $\sigma(1) = 1$, which also gives us that $\sigma(-a) = -\sigma(a)$, $\sigma(a^{-1}) = \sigma(a)^{-1}$, and $\mathrm{ord}(a) = \mathrm{ord}(\sigma(a))$ for all $a \in \mathbb{F}^\times$. It is easy to verify that the set of fixed points of an automorphism form a sub-field of $\mathbb{F}$, denoted the *fixed field* of the automorphism. It is also a consequence of Galois theory that the fixed field of an automorphism over the field $\mathbb{F}_q$ where $q = p^w$ is always an extension of the base prime field $\mathbb{F}_p$ [4].

2.4 Related Work

The code conversion problem was first formulated in [15]. Previous work has yielded several optimal constructions of convertible codes, including pairs of systematic MDS codes with Vandermonde parity matrices, pairs of systematic MDS codes with parity matrices based on Hankel arrays, pairs of low-field-size non-systematic access-optimal convertible codes, and so on [11, 15, 17]. These constructions were analyzed under the lens of access cost, as defined in §2.2.

There also have been previous efforts to study the fundamental limits of existence of super-regular Vandermonde matrices. Shparlinski provided an upper bound on the total number of singular square submatrices of a Vandermonde matrix by showing that any $(q-1) \times m$ Vandermonde matrix $V_{q-1}(\xi_1, \ldots, \xi_m)$ over the field $\mathbb{F}_q$ has at most $3(m-1)(q-1)^m T^{\frac{-1}{m-1}}$ singular $m \times m$ square submatrices where $T := \min_{i \neq j \in [m]} \mathrm{ord}(\frac{\xi_i}{\xi_j})$; however, this bound has not shown to be very tight upon closer investigation [12, 21]. Additionally, Intel's Intelligent Storage Acceleration Library (ISA-L), commonly used to implement erasure coding in practice, has published bounds on the range of parameters $[n, k]$ over $\mathbb{F}_{256}$ for which its code supports generation of super-regular Vandermonde parity matrices, based on a very specific construction [8]. There is no proof provided alongside these bounds as they were likely determined by running a code script to test each submatrix for invertibility.

In addition, there has been independent work studying systematic linear MDS codes with various other constructions of super-regular parity matrices. For example, it is known that a Cauchy matrix $\mathbf{C}$, that is, one of the form $C_{i,j} = (a_i + b_j)^{-1}$ for all $i, j \in [n]$ given two vectors $(a_i)_{i=1}^n$ and $(b_j)_{j=1}^n$, is super-regular so long as the a_i's and b_j's are all distinct from each other [3, 18, 19]. Additionally, Lacan and Fimes introduced a construction of super-regular matrices formed by taking the product of two Vandermonde matrices [12]. To add on, there has been considerable progress in constructing super-regular Toeplitz matrices in the development of convolutional codes [1, 2, 7]. Nonetheless, none of these alternatives are suitable for the construction of access-optimal convertible codes.

To our knowledge, in this thesis we establish the best known bounds on the field size required for

the existence of systematic MDS access-optimal convertible codes for the merge regime where $r^F = r^I$. We are also the first to provide, with proof, explicit constructions of systematic MDS access-optimal convertible codes for the merge regime where $r^F = r^I$ over practically usable field sizes.

Chapter 3

Fundamental Limits on Field Size

In this section, we study a new construction of systematic MDS access-optimal convertible codes for the merge regime where $r^I = r^F$ that generalizes the construction introduced in [15]. The new construction is still based on systematic codes with super-regular Vandermonde parity matrices, but we allow the scalars to take on any distinct nonzero values, rather than being restricted to consecutive powers of a primitive element in the field. As detailed in §2.2, the new construction of convertible codes is still access-optimal. Thus, a proof of the existence of any $k \times r$ super-regular Vandermonde matrix yields $(n^I, k^I; n^F, k^F = \lambda k^I)$ systematic MDS access-optimal convertible codes for any $\lambda \geq 2$, $k^F \leq k$, and $r^I = r^F \leq r$.

We will establish both lower (Theorems 1 and 2) and upper (Theorem 3) bounds on the minimum field size required for existence of systematic MDS codes with Vandermonde parity matrices by studying super-regular Vandermonde matrices. We start with a result which provides a lower bound on the field size required for the existence of such matrices. This result draws upon intuition that an optimal choice of scalars for the Vandermonde matrix would avoid selecting elements with smaller order to avoid repetition along the corresponding columns.

Theorem 1. *Over the field $\mathbb{F}_q$, for every divisor m of $q-1$, for any r, k such that $k > m$, a $k \times r$ super-regular Vandermonde matrix can only exist if $q \geq rm + 1 \in \Omega(kr)$.*

Proof. Provided in **??**. □

The next lemma stems from the fact that finite prime power fields can be viewed as vector spaces over their base prime fields and have a fixed dimension. This implies that any collection of field elements larger than the field's dimension must be linearly dependent. Over fields of characteristic 2, this simply corresponds to a nonempty subset of elements that add to 0. This lemma will later be used to identify a corresponding singular submatrix in a proposed Vandermonde matrix and highlights the necessity of the linear independence of our selected scalars for fields of characteristic 2.

Lemma 1. *Over the field $\mathbb{F}_q$, where $q = 2^w$, for any $r > w$, for any $S = \{\xi_i\}_{i=1}^r \subseteq \mathbb{F}_q$, there must exist some nonempty subset $I \subseteq [r]$ such that $\sum_{i \in I} \xi_i = 0$.*

Proof. Provided in **??**. □

This results in another lower bound on the minimum field size required for the existence of super-regular Vandermonde matrices specific to fields of characteristic 2.

Theorem 2. *Over the field $\mathbb{F}_q$, where $q = 2^w$, for any r, k such that $k > r$, a $k \times r$ super-regular Vandermonde matrix with distinct, nonzero scalars can only exist if $q \geq 2^r$.*

Proof. Let $q < 2^r$, and consider the $k \times r$ Vandermonde matrix $V_k(\xi)$ for any distinct scalars $(\xi_i)_{i=1}^r \in (\mathbb{F}_q^\times)^r$ and r, k such that $k > r$. Then, it follows by Lemma 1, that $\exists I \subseteq [r]$ nonempty such that $\sum_{i \in I} \xi_i = 0$, and we must have $|I| > 2$ as the ξ_i's are nonzero and distinct. Let us define $\ell := |I|$ and $(c_i)_{i=1}^{\ell+1} \in \mathbb{F}_q^{\ell+1}$ to be the coefficient vector of the polynomial $f(x) := \prod_{i \in I}(x - \xi_i)$ such that $f(x) = \sum_{i=1}^{\ell+1} c_i x^{i-1}$, and note by construction $c_\ell = \sum_{i \in I} \xi_i = 0$. Now consider the square submatrix $\mathbf{H} := V_k(\xi)_{J \times I}$ where $J = [\ell + 1] \setminus \{\ell\}$. If we take the linear combination $\mathbf{m} = c_{\ell+1}\mathrm{row}_\ell(\mathbf{H}) + \sum_{i=1}^{\ell-1} c_i \mathrm{row}_i(\mathbf{H})$, it follows that $\mathbf{m} = (f(\xi_i))_{i \in I} = \mathbf{0}$. As $c_{\ell+1} = 1$, this is a nontrivial linear combination of the rows of $\mathbf{H}$, and thus $\mathbf{H}$ is singular. Therefore, in order for the matrix to be super-regular, we must have $q \geq 2^r$. □

The first lower bound for general fields is a tighter bound in many popular settings such as those which require wide codes, or very little storage overhead and thus $k \gg r$. For the field $\mathbb{F}_{256}$, for example, this bound tells us that $[90, 86]$ and $[58, 52]$ codes do not exist systematic MDS codes with Vandermonde parity matrices. On the other hand, the second lower bound specific to fields of characteristic 2 is more relevant in settings which demand narrow codes, such as storage in unreliable environments. Again, for $\mathbb{F}_{256}$, this bound informs us that there do not exist systematic MDS codes with Vandermonde parity matrices and more than 8 parities.

Finally, we will use the Schwartz–Zippel lemma[20, 22] to prove the existence of $k \times r$ super-regular Vandermonde matrices over all fields of size greater than a threshold in terms of k and r. We first start with a lemma to show that given any square submatrix $\mathbf{H}$ of a Vandermonde matrix, if we compared it to the submatrix $\mathbf{H}'$ formed by taking the row indices of $\mathbf{H}$ and "shifting" them all upwards by the same amount so that the first row of the Vandermonde matrix was included in the row index set of $\mathbf{H}'$, either both $\mathbf{H}$ and $\mathbf{H}'$ are singular or both are non-singular; more specifically, the determinant of $\mathbf{H}$ is just a non-zero multiple of that of $\mathbf{H}'$. This will be useful in narrowing down the submatrices that need to be tested for singularity to determine if the matrix is super-regular.

Lemma 2. *Over the field $\mathbb{F}_q$, for any r, k, ℓ such that $\ell \leq \min(r, k)$, for any $k \times r$ Vandermonde matrix $V_k(\xi)$ with $(\xi_i)_{i=1}^r \in (\mathbb{F}_q^\times)^r$, the submatrix $\mathbf{H} := V_k(\xi)_{I \times J}$ defined by $I := \{\alpha_1, \ldots, \alpha_\ell\} \subseteq [k]$ and $J := \{\beta_1, \ldots, \beta_\ell\} \subseteq [r]$, where $\alpha_i < \alpha_j$ for all $i < j$, is non-singular if and only if the submatrix $\mathbf{H}' := V_k(\xi)_{I' \times J}$ defined by $I' := \{1, \alpha_2 - (\alpha_1 - 1), \ldots, \alpha_\ell - (\alpha_1 - 1)\} \subseteq [k]$ and J is non-singular.*

Proof. Provided in **??**. □

We now utilize the Schwartz–Zippel lemma in a probabilistic argument for the existence of a super-regular Vandermonde matrix given a sufficiently large field size. This, in effect, establishes an *upper bound* on the *minimum* field size required for the existence such matrices.

Theorem 3. *Over the field $\mathbb{F}_q$, for any r, k, if $q > 1 + \binom{k}{2} \sum_{\ell=2}^{r} \binom{r}{\ell} \binom{k-2}{\ell-2} \in O(k^r)$, then there must exist scalars $(\xi_i)_{i=1}^r \in (\mathbb{F}_q^\times)^r$ such that the $k \times r$ Vandermonde matrix $V_k(\xi)$ is super-regular.*

Proof. Provided in **??**. □

Chapter 4

Low Field Size Constructions

In this section, we find several new families of explicit constructions of systematic MDS access-optimal convertible codes in the merge regime over low field sizes. Specifically, we provide explicit constructions of $(n^I, k^I; n^F, k^F = \lambda k^I)$ convertible codes where $\lambda \geq 2$ over fields $\mathbb{F}_q$ of characteristic 2 for all parameters such that $r^F = r^I \leq 3$ and $k^F < q$. In a related result for *general prime power fields* $\mathbb{F}_q$ where $q = p^w$, we provide constructions of $(n^I, k^I; n^F, k^F = \lambda k^I)$ convertible codes where $\lambda \geq 2$ for all parameters such that $r^F = r^I \leq 3$ and $k^F < w$. We do this by providing constructions of $k \times 3$ super-regular Vandermonde matrices given a sufficiently large field size: $q > k$ for finite fields of characteristic 2 (Theorem 4) and $q > p^k$ for general prime power fields (Theorem 5). These matrices serve as the parity matrices for the systematic MDS codes that underlie the aforementioned convertible codes. As every submatrix of a super-regular matrix is also super-regular, a valid parity matrix for 3 parities gives us one for any fewer than 3 parities as well.

We start with a lemma that builds on the intuition to choose primitive elements of the finite field for the scalars of the super-regular Vandermonde parity matrix.

Lemma 3. *Over the field $\mathbb{F}_q$, for all $k < q$, given any primitive element $\theta \in \mathbb{F}_q$, given $2 \leq e \leq q-1$ such that $e, e-1 \perp q-1$, the $k \times 3$ Vandermonde matrix $V_k(1, \theta, \theta^e)$ has no singular 2×2 square submatrices.*

Proof. Provided in **??**. □

Next, we introduce the idea of field automorphisms into our construction and choice of scalars, in particular as automorphisms are order preserving maps. Recall some key properties of field automorphisms from §2.3.

Lemma 4. *Over the field $\mathbb{F}_q$ where $q = p^w$, for all $k < q$, given any primitive element $\theta \in \mathbb{F}_q$ and nontrivial $\sigma \in \mathrm{Aut}(\mathbb{F}_q)$ with fixed field $\mathbb{F}_p$, the $k \times 3$ Vandermonde matrix $V_k(1, \theta, \sigma(\theta))$ has no 2×2 singular square submatrices.*

Proof. First, recall that $\mathrm{Aut}(\mathbb{F}_q)$ is a group generated by the Frobenius automorphism, or the map $\sigma : x \to x^p$, and thus any nontrivial element $\sigma \in \mathrm{Aut}(\mathbb{F}_q)$ is of the form $\sigma(x) = x^{p^n}$ for some $1 \leq n < w$. It follows that $p \leq p^n < p^w = q$, and because $q \equiv 0 \mod p$, $q - 1 \not\equiv 0 \mod p$ and clearly $p^n \perp q - 1$. Next, see that if σ has fixed field $\mathbb{F}_p$, this can only occur if $p_1(x) = x^{p^n} - x$, and consequently $p_2(x) = x^{p^n - 1} - 1$, have no roots in $\mathbb{F}_q$ outside of $\mathbb{F}_p$. This

implies that $p^n - 1 \perp q - 1$, and thus we can apply Lemma 3 to get that this matrix has no 2×2 singular submatrices. □

For the same construction of Vandermonde matrices as in Lemma 4, we next consider its 3×3 square submatrices and establish the necessary and sufficient conditions under which they are singular. We are able to show a significantly tighter end result for fields of characteristic 2 in particular, but a lot of the arguments used apply to all finite fields as well. Thus, we start with an intermediate result using the shared ideas.

Lemma 5. *Over the field $\mathbb{F}_q$ where $q = p^w$, for all $k < q$, given any primitive element $\theta \in \mathbb{F}_q$ and nontrivial $\sigma \in \mathrm{Aut}(\mathbb{F}_q)$ with fixed field $\mathbb{F}_p$, the $k \times 3$ Vandermonde matrix $V_k(1, \theta, \sigma(\theta))$ has a 3×3 singular square submatrix if and only if $\exists e_1, e_2 \in [k-1]$ and $c_1, c_2 \in \mathbb{F}_p^\times$ such that $e_1 < e_2$ and $\{1, \theta, \sigma(\theta)\}$ are all roots of the polynomial $f(x) = c_1 + c_2 x^{e_1} + x^{e_2}$.*

Proof. Provided in **??**. □

We now arrive at the first of our major results from this section on fields of characteristic 2.

Theorem 4. *Over the field $\mathbb{F}_q$ where $q = 2^w$, for all $k < q$, given any primitive element $\theta \in \mathbb{F}_q$ and a non-trivial automorphism $\sigma \in \mathrm{Aut}(\mathbb{F}_q)$ with fixed field $\mathbb{F}_2$, the $k \times 3$ Vandermonde matrix $V_k(1, \theta, \sigma(\theta))$ is super-regular.*

Proof. First, note that every 1×1 submatrix of $V_k(1, \theta, \sigma(\theta))$ is non-singular as every element is a power of a nonzero element of $\mathbb{F}_q$. Next, by Lemma 4, every 2×2 submatrix of $V_k(1, \theta, \sigma(\theta))$ is also non-singular. Finally, assume for sake of contradiction that $V_k(1, \theta, \sigma(\theta))$ has a singular 3×3 square submatrix. Then by Lemma 5, $\exists e_1, e_2 \in [k-1]$ and $c_1, c_2 \in \mathbb{F}_2^\times$ such that $\{1, \theta, \sigma(\theta)\}$ are all roots of the polynomial $f(x) = c_1 + c_2 x^{e_1} + x^{e_2}$. However, this implies $c_1 = c_2 = 1$, but then $f(1) = 1 + 1 + 1 = 1$, contradicting the fact that 1 is a root of f. Therefore, every 3×3 square submatrix is also non-singular and $V_k(1, \theta, \sigma(\theta))$ is super-regular, as desired. □

Using this result and the Frobenius automorphism, which is known to have fixed field $\mathbb{F}_p$ over any finite extension $K/\mathbb{F}_p$ [4], we show a family of constructions of super-regular Vandermonde matrices for all fields of characteristic 2. This is of particular interest as they are the most efficient choice for the representation of data in machines and on storage devices. We also give results specific to the field $\mathbb{F}_{256}$, which is most commonly used in practice.

Corollary 1. *Over the field $\mathbb{F}_q$ where $q = 2^w$, for all $k < q$, given any primitive element $\theta \in \mathbb{F}_q$, the $k \times 3$ Vandermonde matrix $V_k(1, \theta, \theta^2)$ is super-regular.*

Corollary 2. *Over the field $\mathbb{F}_{256}$, for all $k < 256$, given any primitive element $\theta \in \mathbb{F}_{256}$, the $k \times 3$ Vandermonde matrices $V_k(1, \theta, \theta^2)$, $V_k(1, \theta, \theta^8)$, $V_k(1, \theta, \theta^{32})$, and $V_k(1, \theta, \theta^{128})$ are super-regular.*

Finally, we show an analogous and more general result for super-regular Vandermonde matrices over any arbitrary field $\mathbb{F}_q$. Note that for fields of characteristic 2, the first result is stronger as it covers a larger range of possible dimensions of the Vandermonde matrix.

Theorem 5. *Over the field $\mathbb{F}_q$ where $q = p^w$, for all $k \leq w$, given any primitive element $\theta \in \mathbb{F}_q$ and a non-trivial automorphism $\sigma \in \mathrm{Aut}(\mathbb{F}_q)$ with fixed field $\mathbb{F}_p$, the $k \times 3$ Vandermonde matrix $V_k(1, \theta, \sigma(\theta))$ is super-regular.*

Proof. Again, note that every 1×1 submatrix of $V_k(1, \theta, \sigma(\theta))$ is non-singular as every element is a power of a nonzero element of $\mathbb{F}_q$. Next, by Lemma 4, every 2×2 submatrix of $V_k(1, \theta, \sigma(\theta))$ is also non-singular. Finally, assume for sake of contradiction that $V_k(1, \theta, \sigma(\theta))$ has a singular 3×3 square submatrix. Then by Lemma 5, $\exists e_1, e_2 \in [k-1]$ and $c_1, c_2 \in \mathbb{F}_p^\times$ such that $e_1 < e_2$ and $\{1, \theta, \sigma(\theta)\}$ are all roots of the polynomial $f(x) = c_1 + c_2 x^{e_1} + x^{e_2}$. However, as $f \in \mathbb{F}_p[x]$, it must be a multiple of the minimum polynomial of θ in $\mathbb{F}_p[x]$, which we know is of degree $w \geq k > e_2 = \deg(f)$ as θ is a generator of $\mathbb{F}_q^\times$, resulting in a contradiction. Thus, every 3×3 square submatrix is also non-singular and $V_k(1, \theta, \sigma(\theta))$ is super-regular, as desired. □

Corollary 3. *Over the field $\mathbb{F}_q$ where $q = p^w$, for all $k \leq w$, given any primitive element $\theta \in \mathbb{F}_q$, the $k \times 3$ Vandermonde matrix $V_k(1, \theta, \theta^p)$ is super-regular.*

The final result is one that shows for any binary field $\mathbb{F}_q$, when $k \leq 3$, in fact, almost every $3 \times r$ Vandermonde matrix for all $r < q$ is super-regular.

Theorem 6. *Over the field $\mathbb{F}_q$ where $q = 2^w$, for all $r < q$, any $3 \times r$ Vandermonde matrix $V_3(\xi)$ where $(\xi_i)_{i=1}^r \in \mathbb{F}_q^r$ and the ξ_i's are all distinct and nonzero is super-regular.*

Proof. Provided in **??**. □

Chapter 5

Conclusion

Bibliography

[1] P. Almeida, D. Napp, and R. Pinto. A new class of superregular matrices and mdp convolutional codes. *Linear Algebra and its Applications*, 439(7):2145–2157, 2013. ISSN 0024-3795. doi: https://doi.org/10.1016/j.laa.2013.06.013. 2.4

[2] Paulo José Fernandes Almeida and Diego Napp Avelli. Superregular matrices over small finite fields. *ArXiv*, abs/2008.00215, 2020. 2.4

[3] Joan-Josep Climent, Diego Napp, Carmen Perea, and Raquel Pinto. A construction of mds 2d convolutional codes of rate 1/n based on superregular matrices. *Linear Algebra and its Applications*, 437(3):766–780, 2012. ISSN 0024-3795. doi: https://doi.org/10.1016/j.laa.2012.02.032. 2.4

[4] D.S. Dummit and R.M. Foote. *Abstract Algebra*. Wiley, 2003. ISBN 9780471433347. 2.3, 4

[5] I. Gohberg and V. Olshevsky. Fast algorithms with preprocessing for matrix-vector multiplication problems. *Journal of Complexity*, 10(4):411–427, 1994. ISSN 0885-064X. 1

[6] Yuchong Hu, Liangfeng Cheng, Qiaori Yao, Patrick P. C. Lee, Weichun Wang, and Wei Chen. Exploiting combined locality for Wide-Stripe erasure coding in distributed storage. In *19th USENIX Conference on File and Storage Technologies (FAST 21)*, pages 233–248. USENIX Association, February 2021. ISBN 978-1-939133-20-5. 1

[7] Ryan Hutchinson, Roxana Smarandache, and Jochen Trumpf. On superregular matrices and mdp convolutional codes. *Linear Algebra and its Applications*, 428(11):2585–2596, 2008. ISSN 0024-3795. doi: https://doi.org/10.1016/j.laa.2008.02.011. 2.4

[8] Intel. Corrupted fragment on decode · issue #10 · intel/isa-l. URL `https://github.com/intel/isa-l/issues/10`. 2.4

[9] Saurabh Kadekodi, K. V. Rashmi, and Gregory R. Ganger. Cluster storage systems gotta have HeART: improving storage efficiency by exploiting disk-reliability heterogeneity. In Arif Merchant and Hakim Weatherspoon, editors, *17th USENIX Conference on File and Storage Technologies, FAST 2019, Boston, MA, February 25-28, 2019*, pages 345–358. USENIX Association, 2019. 1

[10] Saurabh Kadekodi, Shashwat Silas, David Clausen, and Arif Merchant. Practical design considerations for wide locally recoverable codes (lrcs). *ACM Trans. Storage*, 19(4), nov 2023. ISSN 1553-3077. doi: 10.1145/3626198. 1

[11] Xiangliang Kong. Locally repairable convertible codes with optimal access costs. *ArXiv*,

abs/2308.06802, 2023. 2.4

[12] J. Lacan and J. Fimes. Systematic mds erasure codes based on vandermonde matrices. *IEEE Communications Letters*, 8(9):570–572, 2004. 1, 2.4

[13] Florence Jessie MacWilliams and Neil James Alexander Sloane. *The theory of error correcting codes*, volume 16. Elsevier, 1977. 2.1, 2.1

[14] Francisco Maturana and K. V. Rashmi. Bandwidth cost of code conversions in the split regime. In *2022 IEEE International Symposium on Information Theory (ISIT)*, pages 3262–3267, 2022. doi: 10.1109/ISIT50566.2022.9834604. 1

[15] Francisco Maturana and K. V. Rashmi. Convertible codes: enabling efficient conversion of coded data in distributed storage. *IEEE Transactions on Information Theory*, 68:4392–4407, 2022. ISSN 1557-9654. doi: 10.1109/TIT.2022.3155972. (document), 1, 2.2, 1, 2.2, 2.3, 2.4, 3

[16] Francisco Maturana and K. V. Rashmi. Bandwidth cost of code conversions in distributed storage: Fundamental limits and optimal constructions. *IEEE Transactions on Information Theory*, 69(8):4993–5008, 2023. doi: 10.1109/TIT.2023.3265512. 1

[17] Francisco Maturana, V. S. Chaitanya Mukka, and K. V. Rashmi. Access-optimal linear MDS convertible codes for all parameters. In *IEEE International Symposium on Information Theory, ISIT 2020, Los Angeles, California, USA, June 21-26, 2020*, 2020. 1, 2.4

[18] R.M. Roth and A. Lempel. On mds codes via cauchy matrices. *IEEE Transactions on Information Theory*, 35(6):1314–1319, 1989. doi: 10.1109/18.45291. 2.4

[19] R.M. Roth and G. Seroussi. On generator matrices of mds codes (corresp.). *IEEE Transactions on Information Theory*, 31(6):826–830, 1985. doi: 10.1109/TIT.1985.1057113. 1, 2.4

[20] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, oct 1980. ISSN 0004-5411. doi: 10.1145/322217.322225. 3

[21] Igor E. Shparlinski. On the singularity of generalised vandermonde matrices over finite fields. *Finite Fields and Their Applications*, 11(2):193–199, 2005. ISSN 1071-5797. doi: https://doi.org/10.1016/j.ffa.2004.11.001. 1, 2.4

[22] Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation*, pages 216–226, Berlin, Heidelberg, 1979. Springer Berlin Heidelberg. ISBN 978-3-540-35128-3. 3