RJ10511 (ALM1308-007) August 22, 2013
Computer Science

IBM Research Report

Construction of Sector-Disk (SD) Codes with Two Global Parity Symbols

Mario Blaum
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120-6099
USA

James S. Plank
EECS Department
University of Tennessee
Knoxville, TN 37996

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

Construction of Sector-Disk (SD) Codes with two Global Parity Symbols

Mario Blaum
IBM Almaden Research Center
San Jose, CA 95120

James S. Plank
EECS Department
University of Tennessee
Knoxville, TN 37996

August 8, 2013

Abstract

Sector-Disk (SD) codes are erasure codes that address the mixed failure mode of current RAID systems. Rather than dedicate entire disks to erasure coding, as done in RAID-5, RAID-6 and Reed-Solomon coding, an SD code dedicates entire disks, plus individual sectors to erasure coding. The code then tolerates combinations of disk and sector errors, rather than solely disk errors. It has been an open problem to construct general codes that have the SD property, and previous work has relied on Montecarlo searches. In this paper, we present a general construction that addresses the case of any number of failed disks and in addition, two erased sectors. This result generalizes previous constructions extending RAID 5 and RAID 6.

Keywords: Error-correcting codes, RAID architectures, MDS codes, array codes, Reed-Solomon codes, Blaum-Roth codes, PMDS codes, SD codes.

1 Introduction

Consider an $m \times n$ array whose entries are elements in a finite field $GF(2^b)$ [6] (in general, we could consider a field $GF(p^b)$, p a prime number, but for simplicity, we constrain ourselves to binary fields). The n columns represent storage devices like SSDs, HDDs or tapes. The arrays (often called stripes also) are repeated as many times as necessary. In order to protect against a device failure, a RAID 4 or RAID 5 type of scheme, in which one of the devices is the XOR of the other ones, can be implemented. During reconstruction, the failed device is recovered sector by sector. The problem with RAID 5 is, if an additional sector is defective in addition to the one corresponding to the failed device, data loss will occur. A solution to this problem is using a second device for parity (RAID 6), allowing for recovery against two failed devices. However, this scheme may be wasteful, and moreover, it is unable to correct

1	0	1	0	0
E	1	E	0	1
1	1	1	0	1
1	E	1	1	E

1	E	1	0	0
0	E	1	0	E
1	E	1	0	1
E	E	1	1	1

1	0	1	E	0
E	1	E	E	1
1	1	1	E	1
1	0	1	E	1

Figure 1: A 4×5 array with different types of failures

the situation in which in addition to the sector corresponding to the failed disk, we have two extra failed sectors in the row (we always assume that failed sectors can be identified, either by CRC or by other means, so the correcting scheme is an erasure correcting scheme). In order to overcome this problem, the so called Partial MDS (PMDS) codes [2] and Sector-Disk (SD) codes [7] were created. Very similar codes were presented in [5].

We start by giving the definition of PMDS and SD codes.

Definition 1.1 Let $\mathcal{C}$ be a linear $[rn, r(m-r)-s]$ code over a field such that when codewords are taken row-wise as $r \times n$ arrays, each row belongs in an $[n, n-m, m+1]$ MDS code. Then,

1. $\mathcal{C}$ is an $(m;s)$ partial-MDS (PMDS) code if, *for any* $(s_1, s_2, \ldots, s_t)$ such that each $s_j \geq 1$ and $\sum_{j=1}^{t} s_j = s$, and for any $i_1, i_2, \ldots, i_t$ such that $0 \leq i_1 < i_2 < \cdots < i_t \leq r-1$, $\mathcal{C}$ can correct up to $s_j + m$ erasures in each row i_j, $1 \leq j \leq t$, of an array in $\mathcal{C}$.

2. $\mathcal{C}$ is an $(m;s)$ sector-disk (SD) code if, for any $l_1, l_2, \ldots, l_m$ such that $0 \leq l_1 < l_2 < \cdots < l_m \leq n-1$, for any $(s_1, s_2, \ldots, s_t)$ such that each $s_j \geq 1$ and $\sum_{j=1}^{t} s_j = s$, and for any $i_1, i_2, \ldots, i_t$ such that $0 \leq i_1 < i_2 < \cdots < i_t \leq r-1$, $\mathcal{C}$ can correct up to $s_j + m$ erasures in each row i_j, $1 \leq j \leq t$, of an array in $\mathcal{C}$ provided that locations $l_1, l_2, \ldots l_m$ in each of the rows i_j have been erased.

SD codes satisfy a weaker condition than PMDS codes, but they may be sufficient in most applications. The case of $(r;1)$ PMDS codes has been solved in [2]. In this paper, we address the case of (1;2) PMDS and SD codes. Figure 1 illustrates the difference between (1;2) PMDS and SD codes for a 4×5 array (i.e., a code of length 20): the array in the left depicts a situation that can be handled by a (1;2) PMDS but not by a (1;2) SD code; the second and the fourth rows have two erasures (denoted by E) each and there is no column containing two of these erasures. The array in the middle illustrates a situation in which the second and fourth rows have two erasures each, but the second column contains two of those erasures, which correspond to a total failure of the second device. Individual erasures in a row can always be handled by single parity (like in the first and the third rows). This situation can be handled by both (1;2) PMDS and SD codes. Finally, the array in the right shows the situation of three erasures in a row, and at most one in the remaining ones. This situation can also be handled by both (1;2) PMDS and SD codes (but not by RAID 6).

In the next section we give the construction of an $(m;2)$ SD code. Constructions of (1;2) SD codes were given in [1] and of (2;2) codes in [3], so this result is a generalization of those constructions.

From now on, when we say SD codes, we refer to $(m;2)$ SD codes.

2 Code Construction

Consider the field $GF(2^w)$ and let α be an element in $GF(2^w)$. The (multiplicative) order of α, denoted $\mathcal{O}(\alpha)$, is the minimum ℓ, $0<\ell$, such that $\alpha^\ell=1$. If α is a primitive element [6], then $\mathcal{O}(\alpha)=2^w-1$. To each element $\alpha\in GF(2^w)$, there is an associated (irreducible) minimal polynomial [6] that we denote $f_\alpha(x)$.

Let $\alpha\in GF(2^w)$ and $rn\le\mathcal{O}(\alpha)$. We want to construct an SD-code consisting of $r\times n$ arrays over $GF(2^w)$, such that r of the columns correspond to parity (in RAID 5, $r=1$, while in RAID 6, $r=2$). In addition, two extra symbols also correspond to parity. When read row-wise, the codewords belong in an $[rn, r(n-m)-2]$ code over $GF(2^w)$. Specifically, let $\mathcal{C}(r,n,m,2;f_\alpha(x))$ be the $[rn, r(n-m)-2]$ code whose $(mr+2)\times rn$ parity-check matrix is given by

$$H \;=\; \left(\begin{array}{c|c|c|c} H_0 & \underline{0} & \ldots & \underline{0} \\ \hline \underline{0} & H_0 & \ldots & \underline{0} \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline \underline{0} & \underline{0} & \ldots & H_0 \\ \hline H_1 & H_2 & \ldots & H_r \end{array}\right) \qquad (1)$$

where

$$H_0 \;=\; \begin{pmatrix} 1 & 1 & 1 & \ldots & 1 \\ 1 & \alpha & \alpha^2 & \ldots & \alpha^{n-1} \\ 1 & \alpha^2 & \alpha^4 & \ldots & \alpha^{2(n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha^{m-1} & \alpha^{2(m-1)} & \ldots & \alpha^{(m-1)(n-1)} \end{pmatrix} \qquad (2)$$

and, for $1\le j\le r$

$$H_j \;=\; \begin{pmatrix} 1 & \alpha^m & \alpha^{2m} & \ldots & \alpha^{m(n-1)} \\ \alpha^{-(j-1)n} & \alpha^{-(j-1)n-1} & \alpha^{-(j-1)n-2} & \ldots & \alpha^{-(j-1)n-(n-1)} \end{pmatrix}. \qquad (3)$$

The main result in this paper is proving that code $\mathcal{C}(r,n,m,2;f_\alpha(x))$ is SD. Unless stated otherwise, for simplicity, let us denote $\mathcal{C}(r,n,m,2;f_\alpha(x))$ by $\mathcal{C}(r,n,m,2)$.

We start by giving some examples.

Example 2.1 Consider the finite field $GF(16)$ and let α be a primitive element, i.e., $\mathcal{O}(\alpha)=15$. Then, the parity-check matrix of $\mathcal{C}(3,5,1,2)$ is given by

$$\left(\begin{array}{ccccc|ccccc|ccccc}
1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ \hline
1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 \\
1 & \alpha^{14} & \alpha^{13} & \alpha^{12} & \alpha^{11} & \alpha^{10} & \alpha^9 & \alpha^8 & \alpha^7 & \alpha^6 & \alpha^5 & \alpha^4 & \alpha^3 & \alpha^2 & \alpha
\end{array}\right).$$

Similarly, the parity-check matrix of $\mathcal{C}(3,5,2,2)$ is given by

$$\left(\begin{array}{ccccc|ccccc|ccccc}
1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 \\ \hline
1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 & 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 & 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 \\
1 & \alpha^{14} & \alpha^{13} & \alpha^{12} & \alpha^{11} & \alpha^{10} & \alpha^9 & \alpha^8 & \alpha^7 & \alpha^6 & \alpha^5 & \alpha^4 & \alpha^3 & \alpha^2 & \alpha
\end{array}\right).$$

Let us point out that the construction of this type of codes is valid also over the ring of polynomials modulo $M_p(x) = 1 + x + \cdots + x^{p-1}$, p a prime number, as done with the Blaum-Roth (BR) codes [4]. In that case, $\mathcal{O}(\alpha) = p$, where $\alpha^{p-1} = 1 + \alpha + \cdots + \alpha^{p-2}$. The construction proceeds similarly, and we denote it $\mathcal{C}(r,n,m,2;M_p(x))$. Utilizing the ring modulo $M_p(x)$ allows for XOR operations at the encoding and the decoding without look-up tables in a finite field, which is advantageous in erasure decoding [4]. It is well known that $M_p(x)$ is irreducible if and only if 2 is primitive in $GF(p)$ [6]. Let us give an example similar to Example 2.1, but over the polynomials modulo $M_{17}(x)$.

Example 2.2 Consider the ring of polynomials modulo $M_{17}(x)$ and let α be an element in the ring such that $M_{17}(\alpha) = 0$, thus, $\mathcal{O}(\alpha) = 17$. Then, the parity-check matrix of $\mathcal{C}(3,5,1,2;M_{17}(x))$ is given by

$$\left(\begin{array}{ccccc|ccccc|ccccc}
1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ \hline
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ \hline
1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 \\
1 & \alpha^{16} & \alpha^{15} & \alpha^{14} & \alpha^{13} & \alpha^{12} & \alpha^{11} & \alpha^{10} & \alpha^9 & \alpha^8 & \alpha^7 & \alpha^6 & \alpha^5 & \alpha^4 & \alpha^3
\end{array}\right).$$

Similarly, the parity-check matrix of $\mathcal{C}(3,5,2,2;M_{17}(x))$ is given by

$$
\left(\begin{array}{ccccc|ccccc|ccccc}
1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\hline
0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & 0 & 0 & 0 & 0 & 0 \\
\hline
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 \\
\hline
1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 & 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 & 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 \\
1 & \alpha^{16} & \alpha^{15} & \alpha^{14} & \alpha^{13} & \alpha^{12} & \alpha^{11} & \alpha^{10} & \alpha^9 & \alpha^8 & \alpha^7 & \alpha^6 & \alpha^5 & \alpha^4 & \alpha^3
\end{array}\right).
$$

Let us give next a lemma that is key to proving that code $\mathcal{C}(r,n,m,2)$ is SD (we omit the proof).

Lemma 2.1 Let $\alpha \in GF(2^w)$, $rn \leq \mathcal{O}(\alpha)$, $1 \leq \ell \leq r-1$ and, if $1 \leq m \leq n-2$, let $0 \leq i_0 < i_1 < i_2 \ldots < i_{m-1} \leq n-1$ and $t, t' \not\in \{i_0, i_1, i_2 \ldots, i_{m-1}\}$. Consider the $(2m+2) \times (2m+2)$ matrix $M(i_0, i_1, i_2, \ldots, i_{m-1}; t, t'; r; n; \ell)$ given by

$$
\left(\begin{array}{ccccc|ccccc}
1 & 1 & \ldots & 1 & 1 & 0 & 0 & \ldots & 0 & 0 \\
\alpha^{i_0} & \alpha^{i_1} & \ldots & \alpha^{i_{m-1}} & \alpha^{t} & 0 & 0 & \ldots & 0 & 0 \\
\alpha^{2i_0} & \alpha^{2i_1} & \ldots & \alpha^{2i_{m-1}} & \alpha^{2t} & 0 & 0 & \ldots & 0 & 0 \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
\alpha^{(m-1)i_0} & \alpha^{(m-1)i_1} & \ldots & \alpha^{(m-1)i_{m-1}} & \alpha^{(m-1)t} & 0 & 0 & \ldots & 0 & 0 \\
\hline
0 & 0 & \ldots & 0 & 0 & 1 & 1 & \ldots & 1 & 1 \\
0 & 0 & \ldots & 0 & 0 & \alpha^{i_0} & \alpha^{i_1} & \ldots & \alpha^{i_{m-1}} & \alpha^{t'} \\
0 & 0 & \ldots & 0 & 0 & \alpha^{2i_0} & \alpha^{2i_1} & \ldots & \alpha^{2i_{m-1}} & \alpha^{2t'} \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & \ldots & 0 & 0 & \alpha^{(m-1)i_0} & \alpha^{(m-1)i_1} & \ldots & \alpha^{(m-1)i_{m-1}} & \alpha^{(m-1)t'} \\
\hline
\alpha^{mi_0} & \alpha^{mi_1} & \ldots & \alpha^{mi_{m-1}} & \alpha^{mt} & \alpha^{mi_0} & \alpha^{mi_1} & \ldots & \alpha^{mi_{m-1}} & \alpha^{mt'} \\
\alpha^{-i_0} & \alpha^{-i_1} & \ldots & \alpha^{-i_{m-1}} & \alpha^{-t} & \alpha^{-n\ell-i_0} & \alpha^{-n\ell-i_1} & \ldots & \alpha^{-n\ell-i_{m-1}} & \alpha^{-n\ell-t'}
\end{array}\right)
$$

Let $\Delta(i_0, i_1, i_2, \ldots, i_{m-1}; t, t'; r; n; \ell) = \det M(i_0, i_1, i_2, \ldots, i_{m-1}; t, t'; r; n; \ell)$. Then,

$$
\begin{aligned}
\Delta(i_0, i_1, i_2, \ldots, i_{m-1}; t, t'; r; n; \ell) &= \alpha^{-\sum_{u=0}^{m-1} i_u} \left(\prod_{0 \leq u < v \leq m-1} (\alpha^{i_u} \oplus \alpha^{i_v})^2 \right) \\
& \left(\prod_{u=0}^{m-1} (\alpha^{i_u} \oplus \alpha^{t})(\alpha^{i_u} \oplus \alpha^{t'}) \right) \left(\alpha^{-t} \oplus \alpha^{-n\ell-t'} \right) \qquad (4)
\end{aligned}
$$

Lemma 2.1 is valid also over the ring of polynomials modulo $M_p(x)$, p prime. We are ready now to state and prove our main result.

Theorem 2.1 Codes $\mathcal{C}(r,n,m,2;f_\alpha(x))$ and $\mathcal{C}(r,n,m,2;M_p(x))$ are SD.

Proof: Assume that m columns have been erased and in addition we have two random erasures. Assume first that these two random erasures have occurred in the same row ℓ of the stripe. The rows that are different from ℓ are corrected since each one of them has m erasures, which are handled by the horizontal code, that is, each horizontal code is given by the parity-check matrix H_0, which is the parity-check matrix of a RS code that can correct up to m erasures [6]. So, we have to solve a linear system with $m+2$ unknowns. Without loss of generality, assume that the erasures in row ℓ have occurred in locations $i_0, i_1, \ldots, i_m, i_{m+1}$, where $0 \le i_0 < i_1 < \cdots < i_m < i_{m+1} \le n$. According to the parity-check matrix of the code as given by (1), (2) and (3), there will be a unique solution if and only if the $(m+2)\times(m+2)$ matrix

$$\begin{pmatrix}
1 & 1 & \ldots & 1 & 1\\
\alpha^{i_0} & \alpha^{i_1} & \ldots & \alpha^{i_m} & \alpha^{i_{m+1}}\\
\alpha^{2i_0} & \alpha^{2i_1} & \ldots & \alpha^{2i_m} & \alpha^{2i_{m+1}}\\
\vdots & \vdots & \ddots & \vdots & \vdots\\
\alpha^{mi_0} & \alpha^{mi_1} & \ldots & \alpha^{mi_m} & \alpha^{mi_{m+1}}\\
\alpha^{-n\ell-i_0} & \alpha^{-n\ell-i_1} & \ldots & \alpha^{-n\ell-i_m} & \alpha^{-n\ell-i_{m+1}}
\end{pmatrix}$$

is invertible. By taking $\alpha^{-n\ell}$ in the last row as a common factor, and by multiplying each column j, $0 \le j \le m+1$, by α^{i_j}, this matrix is transformed into a Vandermonde matrix, which is always invertible in a field and also in the ring of polynomials modulo $M_p(x)$ [4].

Consider now the case in which the two random failures occur in different rows. Specifically, assume that columns $i_0, i_1, \ldots, i_{m-1}$ have been erased, where $0 \le i_0 < i_1 < \ldots < i_{m-1} \le n-1$, and in addition, entries (ℓ, t) and (ℓ', t') are erased, where $t, t' \notin \{i_0, i_1, \ldots, i_{m-1}\}$ and $0 \le \ell < \ell' \le r-1$. Again using the parity-check matrix of the code as given by (1), (2) and (3), there will be a unique solution if and only if the $(2m+2)\times(2m+2)$ matrix

$$\left(\begin{array}{ccccc|ccccc}
1 & 1 & \ldots & 1 & 1 & 0 & 0 & \ldots & 0 & 0\\
\alpha^{i_0} & \alpha^{i_1} & \ldots & \alpha^{i_{m-1}} & \alpha^{t} & 0 & 0 & \ldots & 0 & 0\\
\alpha^{2i_0} & \alpha^{2i_1} & \ldots & \alpha^{2i_{m-1}} & \alpha^{2t} & 0 & 0 & \ldots & 0 & 0\\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots\\
\alpha^{(m-1)i_0} & \alpha^{(m-1)i_1} & \ldots & \alpha^{(m-1)i_{m-1}} & \alpha^{(m-1)t} & 0 & 0 & \ldots & 0 & 0\\
\hline
0 & 0 & \ldots & 0 & 0 & 1 & 1 & \ldots & 1 & 1\\
0 & 0 & \ldots & 0 & 0 & \alpha^{i_0} & \alpha^{i_1} & \ldots & \alpha^{i_{m-1}} & \alpha^{t'}\\
0 & 0 & \ldots & 0 & 0 & \alpha^{2i_0} & \alpha^{2i_1} & \ldots & \alpha^{2i_{m-1}} & \alpha^{2t'}\\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots\\
0 & 0 & \ldots & 0 & 0 & \alpha^{(m-1)i_0} & \alpha^{(m-1)i_1} & \ldots & \alpha^{(m-1)i_{m-1}} & \alpha^{(m-1)t'}\\
\hline
\alpha^{mi_0} & \alpha^{mi_1} & \ldots & \alpha^{mi_{m-1}} & \alpha^{mt} & \alpha^{mi_0} & \alpha^{mi_1} & \ldots & \alpha^{mi_{m-1}} & \alpha^{mt'}\\
\alpha^{-n\ell-i_0} & \alpha^{-n\ell-i_1} & \ldots & \alpha^{-n\ell-i_{m-1}} & \alpha^{-n\ell-t} & \alpha^{-n\ell'-i_0} & \alpha^{-n\ell'-i_1} & \ldots & \alpha^{-n\ell'-i_{m-1}} & \alpha^{-n\ell'-t'}
\end{array}\right)$$

is invertible. Taking $\alpha^{-n\ell}$ as a common factor in the last row, we obtain the matrix $M(i_0, i_1, i_2, \ldots, i_{m-1}; t, t'; r; n; \ell' - \ell)$ as defined in Lemma 2.1, which is invertible since its determinant is a constant times a product of binomials, and each binomial is invertible. □

Let us point out that Lemma 2.1 and Theorem 2.1 not only prove that codes $\mathcal{C}(r, n, m, 2; f_\alpha(x))$ and $\mathcal{C}(r, n, m, 2; M_p(x))$ are SD, but also provide for efficient encoding and decoding algorithms. In effect, solving the linear systems corresponding to erasures, for instance, using Cramer's rule, involves inverting determinants of either Vandermonde type of matrices or determinants of the type $\Delta(i_0, i_1, i_2, \ldots, i_{m-1}; t, t'; r; n; \ell' - \ell)$, as defined in Lemma 2.1. Both types of determinants involve products of binomials, which are easily inverted both in $GF(q)$ and in the ring of polynomials modulo $M_p(x)$ [4].

Acknowledgements

This work is supported by the National Science Foundation, under grant CSR-1016636, and by an IBM Faculty Award.

References

[1] M. Blaum, "Construction of PMDS and SD Codes extending RAID 5," arXiv:1305.0032 [cs.IT], April 2013.

[2] M. Blaum, J. L. Hafner and S. Hetzler, "Partial-MDS Codes and their Application to RAID Type of Architectures," IEEE Trans. on Information Theory, vol.IT-59, pp. 4510-19, July 2013.

[3] M. Blaum and J. S. Plank, "Construction of two SD Codes," arXiv:1305.1221 [cs.IT], May 2013.

[4] M. Blaum and R. M. Roth, "New Array Codes for Multiple Phased Burst Correction," IEEE Trans. on Information Theory, vol. IT-39, pp. 66-77, January 1993.

[5] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li and S. Yekhanin, "Erasure Coding in Windows Azure Storage," 2012 USENIX Annual Technical Conference, Boston, Massachussetts, June 2012.

[6] F. J. MacWilliams and N. J. A. Sloane, "The Theory of Error-Correcting Codes," North Holland, Amsterdam, 1977.

[7] J. S. Plank, M. Blaum and J. L. Hafner, "SD Codes: Erasure Codes Designed for How Storage Systems Really Fail," FAST 13, San Jose, CA, February 2013.