

An approach to RAID-6 based on cyclic groups of a prime order.

Robert Jackson, Dmitriy Rumynin and Oleg V. Zaboronski

Abstract—As the size of data storing arrays of disks grows, it becomes vital to protect data against double disk failures. An economic way of providing such protection consists of adding two disks to the array. The increased storage capacity of the array is used to store the information necessary for the recovery of all data in the case any two disks fail (RAID-6). A popular method of generating recovery information from data utilizes linear operations in the Galois field $GF(2^8)$. Mathematically, this solution is equivalent to using the Reed-Solomon (RS) code with two parity words. Its principle advantage is based on simplicity of the basic operations in a field extension of the primary field $GF(2)$: addition is just bitwise XOR, while multiplication can be constructed using shifts, AND's and additions.

The RS code with two parity words is an example of a linear block code capable of correcting up to two errors in known positions per block. In the paper we construct alternative examples of linear block codes which can be used in RAID-6. Our construction is based on a matrix representation of the roots of unity. We analyze an assembly implementation of the corresponding RAID-6 on a general purpose processor and find that in certain situations it outperforms the RS-based RAID-6 scheme in terms of write and recovery speed and also in terms of the memory usage.

Index Terms—RAID, Reed-Solomon code, cyclic group, Sylvester matrix, Fermat prime.

I. INTRODUCTION

THE volume of information accumulated and stored by a typical small-size information technology company amounts to fifty 100-gigabyte drives. The specified mean time between failures (MTBF) for a modern desktop drive is about 500,000 hours, [1]. Assuming that such an MTBF is actually achieved and that the drives fail independently, the probability of a disk failure in the course of a year is $1 - e^{-50/57} \approx 0.5!$ Therefore, even a small company can no longer avoid the necessity of protecting its data against disk failures. The use of redundant arrays of independent disks (RAID's) enables such a protection in a cost efficient manner.

To protect an array of N disks against a single disk failure it is sufficient to add one more disk to the array. For every N bits of user data written on N disks of the array, a parity bit equal to an exclusive OR (XOR) of these bits is written on the $(N + 1)$ -st disk. Binary content of any disk can be then recovered as a bitwise XOR of contents of remaining N disks. The corresponding system for storing data and distributing parity between disks of the array is referred to as RAID-5, see [2] for a review. Today, RAID-5 constitutes the most popular solution for protected storage.

As the amount of data stored by humanity on magnetic media grows, the danger of *multiple* disk failures within a single array becomes real. In [2], Maddock, Hart and Kean argue that for a storage system consisting of one hundred $8+P$ RAID-5 arrays the rate of failures amounts to losing one array every six months. Because of this danger, RAID-5 is currently being replaced with RAID-6, which offers protection against double failure of drives within the array. RAID-6 refers to any technique to add two strips of redundant data to strips of user data in such a way that all the information can be restored if any two strips are lost.

A number of RAID-6 techniques is reviewed in [2] and [3]. The most well-known RAID-6 scheme is based on the rate-255/257 Reed-Solomon code, see [4] for details. In this scheme two extra disks are introduced for up to 255 disks of data and two parity bytes are computed per 255 data bits. Hardware implementation of RS-based RAID-6 is simple as operations in $GF(256)$ are byte-based. Addition of bytes is just a bitwise XOR. Multiplication of bytes corresponds to multiplication of boolean polynomials modulo an irreducible polynomial. Multiplication can be implemented using XOR's, AND's and shifts.

Some RAID-6 schemes use only bitwise XOR for the computation of parity bits by exploiting a two-dimensional striping of disks of the array. An example is a proprietary RAID-DP developed by Network Appliances, [5]. Some other RAID-6 methods use a non-trivial striping and employ only XOR operation for parity calculation and reconstruction. Examples include X-code, ZZS-code and Park-code, [2].

In all the cases mentioned above, the problem dealt with is inventing an error correcting block code capable of correcting up to two errors in known positions per block (we assume that it is always known which disks have failed). In the present paper we describe a general approach to the solution of this problem, which allows one to develop the most optimal RAID-6 scheme for given technological constraints (e. g. available hardware, the number of disks in the array, the required read and write performance). We also consider an assembly implementation of an exemplary RAID-6 system built using our method and show that it outperforms the Linux kernel implementation of RS-based RAID-6.

The paper is organized as follows. In Section 2 we discuss RAID-6 in the context of systematic linear block codes and construct simple examples of codes capable of correcting two errors in known positions. In Section 3 we identify a mathematical structure common to all such codes and use it to construct RAID-6 schemes starting with elements of a cyclic group of a prime order. In Section 4 we compare an assembly

Manuscript prepared on November 21, 2006.

R. Jackson is with Arithmetica Limited. D. Rumynin and O. Zaboronski are with the University of Warwick.

implementation of RAID-6 based on Z_{17} with its RS-based counterpart implemented as a part of Linux kernel.

II. RAID-6 FROM THE VIEWPOINT OF LINEAR BLOCK CODES.

Suppose that information to be written on the array of disks is broken into words of length n bits. What is the best rate linear block code, which can protect data against the loss of two words?

Altogether, there are 2^{2n} possible pairs of words. In order to distinguish between them, one needs at least $2n$ distinct syndromes, see [6] for an introduction to the theory of linear codes. Therefore, any linear block code capable of restoring 2 lost words in known locations must have at least $2n$ parity checks. Suppose the size of the information block is Nn bits or N words. In the context of RAID, N is the number of information disks to be protected against the failure. Then the code's block size must be at least $(2 + N)n$ and the rate is

$$R \leq \frac{N}{N+2}$$

This result is intuitively clear: to protect N information disks against double failure, we need at least 2 parity disks. Note however, that in order to achieve this optimal rate, the word length n must grow with the number of disks N . Really, the size of parity check matrix is $2n \times (N+2)n$. All columns of the matrix must be distinct and nonzero. Therefore, $(N+2)n \leq (2^{2n} - 1)$, i. e.

$$\frac{1}{n} \left(2^{2n} - 1 - 2n \right) \geq N$$

If in particular $n = 1$, then $N \leq 1$. Therefore, if one wants to protect information written on the disks against double failures using just two parity disks, the word size $n \geq 2$ is necessary. If $n = 2$, we get $N \leq 5$. In reality, the lower bound on n (or, equivalently, the upper bound on N) is more severe, as the condition that all columns of parity check matrix are distinct leads to a code with minimal distance $d_{min} = 2$. However, in order to build a code which corrects up to $2n$ errors in the known location we need $d_{min} = 2n$.

In the following Subsections we will construct explicit examples of linear codes for RAID-6 for small values of n and N . These examples both guide and illustrate our general construction of RAID-6 codes presented in Section 3.

A. Redundant array of four independent disks, which protects against the failure of any two disks.

We restrict our attention to *systematic* linear block codes. These are determined by the parity matrix. To preserve a backward compatibility with RAID-5 schemes, we require half of the parity bits to be the straight XOR's of the information bits. Hence the general form of the parity check matrix for $N = 2$ is

$$P = \begin{pmatrix} I_{n \times n} & I_{n \times n} & I_{n \times n} & 0_{n \times n} \\ H & G & 0_{n \times n} & I_{n \times n} \end{pmatrix} \quad (1)$$

where $I_{n \times n}$ and $0_{n \times n}$ are $n \times n$ identity and zero matrix correspondingly; G and H are some $n \times n$ binary matrices. The corresponding parity check equations are

$$\begin{aligned} d_1 + d_2 + \pi_1 &= 0, \\ H \cdot d_1 + G \cdot d_2 + \pi_2 &= 0. \end{aligned} \quad (2)$$

Here d_1, d_2 are n -bit words written on disks 1 and 2, π_1 and π_2 are n -bit parity check words written on disks 3 and 4; "·" stands for binary matrix multiplication.

Matrices G and H defining the code are constrained by the condition that the system of parity check equations must have a unique solution with respect to *any* pair of variables. To determine these constraints we need to consider the following particular cases.

(π_1, π_2) are lost. The system (2) always has a unique solution with respect to lost variables: we can compute parity bits in terms of information bits.

(d_1, π_2) are lost. The system (2) always has a unique solution with respect to lost variables: compute d_1 in terms of π_1 and d_2 using the first equation of (2) as in RAID-5. Then compute π_2 using the second equation.

(d_2, π_2) are lost. The system (2) always has a unique solution with respect to lost variables: compute d_2 using π_1 and d_1 as in RAID-5. Then compute π_1 using (2).

(π_1, d_1) are lost. The system (2) always has a unique solution with respect to lost variables provided the matrix H is invertible.

(π_1, d_2) are lost. The system (2) always has a unique solution with respect to lost variables provided the matrix G is invertible.

(d_1, d_2) are lost. The system (2) always has a unique solution with respect to lost variables provided the matrix

$$\begin{pmatrix} I_{n \times n} & I_{n \times n} \\ H_{n \times n} & G_{n \times n} \end{pmatrix} \quad (3)$$

is invertible.

As it turns out, one can build a parity check matrix satisfying all the non-degeneracy requirements listed above for $n = 2$. The simplest choice is

$$H = I_{2 \times 2}, \quad G = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}. \quad (4)$$

Non-degeneracy of the three matrices H , G and (3) is evident. For instance,

$$\det \begin{pmatrix} I_{n \times n} & I_{n \times n} \\ H_{n \times n} & G_{n \times n} \end{pmatrix} = -1 = 1.$$

We conclude that the linear block code with a 4×8 parity check matrix (1, 4) gives rise to RAID-6 consisting of four disks. The computation of parity dibits π_1, π_2 in the described DP RAID is almost as simple as the computation of regular parity bits: Let $d_1 = (d_{11}, d_{12})$ and $d_2 = (d_{21}, d_{22})$ be the

dibits to be written on disks one and two correspondingly. Then

$$\begin{aligned}\pi_{11} &= d_{11} + d_{21}, \\ \pi_{12} &= d_{12} + d_{22}, \\ \pi_{21} &= d_{11} + d_{22}, \\ \pi_{22} &= d_{12} + d_{21} + d_{22}.\end{aligned}\quad (5)$$

The computations involved in the recovery of lost data are bitwise XOR's only. As an illustration, let us write down expressions for lost data bits in terms of parity bits explicitly:

$$\begin{aligned}d_{22} &= \pi_{11} + \pi_{12} + \pi_{21} + \pi_{22} \\ d_{12} &= \pi_{11} + \pi_{21} + \pi_{22} \\ d_{11} &= \pi_{11} + \pi_{12} + \pi_{22} \\ d_{21} &= \pi_{12} + \pi_{22}.\end{aligned}$$

It is interesting to note that RAID-6 code described here is equivalent to Network Appliances' horizontal-diagonal parity RAID-DPTM with two data disks, [5]. Really, diagonal-horizontal parity system for two info disks is

$$\begin{array}{cccc} A & B & HP & DP1 \\ C & D & HP2 & DP2,\end{array}$$

where strings (A, C) are written on information disk 1, strings (B, D) are written on disk 2, $(HP, HP2)$ is horizontal parity, $(DP1, DP2)$ is diagonal parity. By definition, $HP = A + B$, $HP2 = C + D$, $DP1 = A + D$, $DP2 = B + C + D$, which coincides with parity check equations (5).

On the other hand, the code (1) is a reduction of the RS code based on $GF(4)$ which we will describe in the next Subsection.

B. Redundant array of five independent disks, which protects against the failure of any two disks.

The code (1) can be extended to a scheme providing double protection of user data written on three disks. The parity check matrix is

$$P = \begin{pmatrix} I_{2 \times 2} & I_{2 \times 2} & I_{2 \times 2} & I_{2 \times 2} & 0_{2 \times 2} \\ I_{2 \times 2} & G & G^2 & 0_{2 \times 2} & I_{2 \times 2}, \end{pmatrix} \quad (6)$$

where 2×2 matrix G was defined in the previous subsection. The corresponding parity check equations are

$$d_1 + d_2 + d_3 + \pi_1 = 0, \quad (7)$$

$$d_1 + G \cdot d_2 + G^2 \cdot d_3 + \pi_2 = 0. \quad (8)$$

The solubility of these equations w. r. t. any pair of variables from the set $\{d_1, d_2, d_3, \pi_1, \pi_2\}$ requires two extra conditions of non-degeneracy in addition to non-degeneracy conditions listed in the previous subsection. Namely, matrices

$$\begin{pmatrix} I_{2 \times 2} & I_{2 \times 2} \\ I_{2 \times 2} & G^2 \end{pmatrix} \text{ and } \begin{pmatrix} I_{2 \times 2} & I_{2 \times 2} \\ G & G^2 \end{pmatrix}$$

must be invertible. It is possible to check the invertibility of these matrices via a direct computation. However, in the next section we will construct a generalization of the above example and find a more economical way of proving non-degeneracy.

The code (1) is a reduction (6) corresponding to $d_3 = 0$. Note also that the code (6) is equivalent to rate-3/5 Reed-Solomon code based on $GF(4)$: a direct check shows that the set of 2×2 matrices $0, 1, G, G^2$ is closed under multiplication and addition and all non-zero matrices are invertible. Thus this set forms a field isomorphic to $GF(4)$. On the other hand, as we established in the previous subsection, the code (1) is equivalent to RAID-DPTM with four disks. Therefore, RAID-DPTM with four disks is a particular case of the RS-based RAID-6. It would be interesting to see if RAID-DPTM can be reduced to the RS-based RAID-6 in general.

We are now ready to formulate general properties of linear block codes suitable for RAID-6 and construct a new class of such codes.

III. RAID-6 BASED ON THE CYCLIC GROUP OF A PRIME ORDER.

A. RAID-6 and cones of $GL_n(GF(2))$.

In this Subsection we will define a general mathematical object underlying all existing algebraic RAID-6 schemes. We refer reader to [7] for basic facts about finite groups used below.

Definition III-A.1. Let $GL_n(GF(2))$ be the set of $n \times n$ non-degenerate binary matrices. A cone¹ C is a subset of $GL_n(GF(2))$ such that $\forall g \neq h \in C \ g + h \in GL_n(GF(2))$.

The usefulness of cones for RAID-6 is explained by the following

Lemma III-A.2. Let $C = \{g_1, g_2, \dots, g_N\} \subseteq GL_n(GF(2))$ be a cone of N elements. Then the system of parity equations

$$\begin{aligned}d_{N+1} &= \sum_{k=1}^N d_k \\ d_{N+2} &= \sum_{k=1}^N g_k d_k\end{aligned} \quad (9)$$

has a unique solution w. r. t. any pair of variables $(d_i, d_j) \in GF(2)^n \times GF(2)^n$, $1 \leq i < j \leq N + 2$. Here d_i 's are binary n -dimensional vectors.

Proof. The fact that system (9) has a unique solution w. r. t. (d_{N+1}, d_{N+2}) is obvious.

The system has a unique solution w. r. t. (d_{N+1}, d_j) for any $j \leq N$: from the second of equations (9), $d_j = g_j^{-1}(d_{N+2} + \sum_{k \neq j} g_k d_k)$, where we used invertibility of $g_j \in GL_n(GF(2))$. With d_j known, d_{N+1} can be computed from the first of equations (9).

The system has a unique solution w. r. t. (d_{N+2}, d_j) for any $j \leq N$: from the first of equations (9), $d_j = d_{N+1} + \sum_{k \neq j} g_k d_k$. With d_j known, d_{N+2} can be computed from the second of equations (9).

¹ This terminology is slightly questionable. If one asks $g + h \in C$ then $C \cup \{0\}$ is a convex cone in the usual mathematical sense. Our choice of the term is influenced by this analogy. Quasi-cone or RAID-cone could be more appropriate scientifically but would pay a heavy linguistic toll.

The system has a unique solution w. r. t. any pair of variables d_i, d_j for $1 \leq i < j \leq N$: multiplying the first of equations (9) with g_i and adding the first and second equations, we get $d_j = (g_i + g_j)^{-1}(g_i d_{N+1} + d_{N+2} + \sum_{k \neq i, j}^N (g_k + g_i) d_k)$. Here we used the invertibility of the sum $g_i + g_j$ for any $i \neq j$, which follows from the definition of the cone. With d_j known, d_i can be determined from any of the equations (9). **QED**

In the context of RAID-6, d_i 's for $1 \leq i \leq N$ can be thought of as n -bit strings of user data, d_{N+1}, d_{N+2} - as n -bit parity strings. The lemma proven above ensures that any two strings can be restored from the remaining N strings.

We conclude that any cone can be used to build RAID-6. The following lemma gives some necessary conditions for a cone.

Lemma III-A.3. *Let $C \subset GL_n(GF(2))$ be a cone.*

- (i) *For all $g, h \in C$ such that $g \neq h$ and for all $x \in GF(2^m)^n$, $gx = hx$ if and only if $x = 0$.*
- (ii) *No two elements of the same cone can share an eigenvector in $GF(2)^n$.*
- (iii) *the cone C can contain no more than one permutation matrix.*

Proof. To prove (i), assume that there is $x \neq 0 : gx = hx$. Then $(g + h)x = 0$, which contradicts the fact that $g + h$ is non-degenerate. Therefore, $x = 0$. Let us prove (ii) now. As elements of C are non-degenerate, the only possible eigenvalue in $GF(2)$ is 1, thus for any two elements sharing an eigenvector x , $x = hx = gx$, which again would imply degeneracy of $h + g$ unless $x = 0$. The statement (iii) follows from (ii) if one notices that any two permutation matrices share an eigenvector whose components are all equal to one. **QED**

The notion of the cone is convenient for restating well understood conditions for a linear block code to be capable of recovering up to two lost words. Our main challenge is to find examples of cones with sufficiently many elements, which lead to easily implementable RAID-6 systems. We will now construct a class of cones starting with elements of a cyclic subgroup of $GL_n(GF(2))$ of a prime order.

B. RAID-6 based on matrix generators of Z_N .

Theorem III-B.1. *Let N be an odd prime number. Let g be an $n \times n$ binary matrix such that $Id + g$ is non-degenerate and $g^N = Id$. Then the elements of cyclic group $Z_N = \{Id, g, g^2, \dots, g^{N-1}\}$ form a cone.*

The proof of the Theorem III-B.1. is based on the following two Lemmas:

Lemma III-B.2. *Let g be a binary matrix such that $Id + g$ is non-degenerate and $g^N = Id$, where N is an*

integer. Then

$$\sum_{k=0}^{N-1} g^k = 0 \quad (10)$$

Proof. Let us multiply the left hand side of (10) with $(Id + g)$ and simplify the result using that $h + h = 0$ for any binary matrix:

$$\begin{aligned} (Id + g) \sum_{k=0}^{N-1} g^k &= Id + g + g + g^2 + \dots \\ &+ g^{N-1} + g^{N-1} + g^N = Id + g^N = Id + Id = 0. \end{aligned}$$

As $Id + g$ is non-degenerate, this implies that $\sum_{k=0}^{N-1} g^k = 0$. **QED**

Lemma III-B.2 is a counterpart of a well-known fact from complex analysis that roots of unity add to zero.

Lemma III-B.3. *Let g be a binary matrix such that $Id + g$ is non-degenerate and $g^N = Id$, where N is an odd prime. Then the matrix $g^l + g^k$ is non-degenerate for any $k, l : 0 \leq k < l < N$.*

Proof. As $g^N = Id$, the matrix g is invertible. To prove the lemma, it is therefore sufficient to check the non-degeneracy of $Id + g^k$ for $0 < k < N$.

As the order of the group $Z_N = \{1, g, g^2, \dots, g^{N-1}\}$ is a prime number, Z_N has no nontrivial subgroup. (Recall that the order of a subgroup must divide the order of the group.) As a result, any element $g_k = g^k$ for $0 < k < N$ generates the whole group. Since the matrix g satisfies all the conditions of Lemma III-B.2, the sum of all elements of Z_N is zero. Therefore,

$$\begin{aligned} 0 &= \sum_{m=0}^{N-1} g_k^m = (Id + g_k) + g_k^2(Id + g_k) + \dots \\ &+ g_k^{N-3}(Id + g_k) + g_k^{N-1} = 0. \end{aligned} \quad (11)$$

The grouping of terms used in (11) is possible as N is odd. Assume that matrix $1 + g_k$ is degenerate. Then there exists a non-zero binary vector x such that $(1 + g_k)x = 0$. Applying both sides of (11) to x we get $g_k^{N-1}x = g^{k(N-1)}x = 0$. This contradicts non-degeneracy of g . Thus the non-degeneracy of $1 + g^k$ is proven for all $0 < k < N$. **QED**

The proof of Theorem III-B.1. The matrix g described in the statement of the theorem satisfies all requirements of Lemma III-B.3. The statement of the theorem follows from Definition III-A.1 of the cone. **QED**

Theorem III-B.1 allows one to determine whether elements of Z_N belong to the same cone by verifying a single non-degeneracy conditions imposed on the generator.

The following corollary of Theorem III-B.1 makes an explicit link between the constructed cone and RAID-6:

Corollary III-B.4. *Let g be an $n \times n$ binary matrix such that $Id + g$ is non-degenerate and $g^N = Id$, where N is an odd prime. The systematic linear block code defined by the parity check matrix*

$$P = \begin{pmatrix} I_{n \times n} & I_{n \times n} & I_{n \times n} & \cdots & I & I & 0 \\ I_{n \times n} & g & g^2 & \cdots & g^{N-1} & 0 & I_{n \times n} \end{pmatrix}$$

can recover up to 2 n -bit lost words in known positions. Equivalently, the system of the parity check equations

$$\begin{aligned} d_1 + d_2 + \cdots + d_N + d_{N+1} &= 0 \\ d_1 + g d_2 + \cdots + g^{N-1} d_N + d_{N+2} &= 0 \end{aligned} \quad (12)$$

has a unique solution w. r. t. any pair of variables (d_i, d_j) , $1 \leq i < j \leq N + 2$.

Proof. It follows from Theorem III-B.1. that first N powers of g belong to a cone. The statement of the corollary is an immediate consequence of Lemma III-A.2 for $g_k = g^{k-1}$, $1 \leq k \leq N$. **QED.**

As a simple application of Theorem III-B.1, let us show that the parity check matrix (6) does indeed satisfy all non-degeneracy requirements: The matrix

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

is non degenerate and has order 3. Also, the matrix

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

is non-degenerate. Hence in virtue of Corollary III-B.4, the parity check matrix (6) determines a RAID system consisting of five disks, that protects against the failure of any two disks.

C. Extension of Z_N -based cones for Fermat primes $N = 2^{2^m} + 1$.

Fermat primes are prime numbers of the form $2^{2^m} + 1$. Currently, only four Fermat primes are known ($N = 3$ ($m = 0$), $N = 5$ ($m = 1$), $N = 17$ ($m = 2$), $N = 257$ ($m = 3$), 65537 ($m = 4$)), but the conjecture stands that there are infinitely many Fermat primes, [8].

We will now show that if N is a Fermat number greater than 3, the cone constructed in the previous subsection can be extended by elements of the form $Id + g^k$, $0 < k < N$. Namely, we will prove the following

Lemma III-C.1. *Let $N = 2^q + 1 > 3$ be a prime number. Let g be an $n \times n$ binary matrix such that $Id + g$ is non-degenerate and $g^N = Id$. Then the set of $2N - 1$ matrices $\{Id, g, g^2, \dots, g^{N-1}, Id + g, Id + g^2, \dots, Id + g^{N-1}\}$ is a cone.*

Proof. There is a slight issue while all these matrices are distinct. This formally follows from the non-degeneracy of the sums in the definition of the cone.

Due to Theorem III-B.1, it is sufficient to check only the non-degeneracy conditions involving the matrices $h_k = Id +$

g^k . The matrices $h_k + h_j = g^k + g^j$ are non degenerate, as elements of the cyclic group of prime order constitute a cone. Matrices $Id + h_k = g^k$ are non-degenerate as the matrix g is non-degenerate.

It remains to check the non-degeneracy of matrices $g^k + h_j = Id + g^k + g^j$. As any power of g between one and $N - 1$ generates the whole group Z_N , we can redefine the generator g^k to be g . Then it suffices to prove the non-degeneracy of all the matrices of the form $Id + g + g^j$, $0 \leq j \leq N - 1$. Suppose that such a matrix is degenerate. Then there is a non-zero binary vector x such that

$$g^j x = (Id + g)x. \quad (13)$$

This relation implies that

$$g^{Nj} x = (Id + g)^N x. \quad (14)$$

Recall the identity $(Id + h)^{2^q} = Id + h^{2^q}$, which holds true for any binary matrix h . Recall also that $g^N = Id$. Therefore,

$$\begin{aligned} (Id + g)^N &= (Id + g)(Id + g)^{2^q} = (Id + g)(Id + g^{2^q}) \\ &= (Id + g)(Id + g^{-1}) = g + g^{-1}. \end{aligned}$$

Using this result, (14) can be re-written as

$$(Id + g + g^{-1})x = 0.$$

Multiplying both sides of the above equality with g , we conclude that degeneracy of $Id + g + g^j$ implies degeneracy of $Id + g + g^2$.

As N is an integer, Lemma III-B.2 holds. Hence the sum of all powers of g up to g^{N-1} is zero. Let us re-write this sum as follows:

$$0 = \sum_{k=0}^{N-1} g^k = (Id + g + g^2)(Id + g^3 + g^6 + \cdots + g^M) + R,$$

where $M = N - 4$, $R = g^{N-1}$ if $N \equiv 1 \pmod{3}$ and $M = N - 5$, $R = g^{N-1} + g^{N-2}$ if $N \equiv 2 \pmod{3}$. Thus the degeneracy of $Id + g + g^2$ would imply degeneracy of either g or $Id + g$, which is a contradiction. The degeneracy of $Id + g^i + g^j$ for any i, j is therefore proven. **QED.**

Applying Lemma III-A.2, we find that starting with matrix generator of the cyclic group of prime order $N = 2^q + 1$ we can build a RAID-6 system protecting up to $2N - 1$ information disks. The explicit expression for Q-parity is

$$Q = \sum_{k=0}^{N-1} g^k d_k + \sum_{k=N}^{2N-2} (Id + g^{k-N+1}) d_k, \quad (15)$$

where $d_0, d_1, \dots, d_{2N-2}$ are information words.

D. Specific examples of matrix generators of Z_N and the corresponding RAID-6 systems.

Now we are ready to construct explicit examples of RAID-6 basing on the theory of cones developed in the above subsections.

Lemma III-D.1. Let S_N be the $(N - 1) \times (N - 1)$ Sylvester matrix,

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

Then

- (i) S_N has order N .
- (ii) Matrix $Id + S_N$ is non-degenerate if N is odd and is degenerate if N is even.

Proof.

(i) An explicit computation shows, that for any $(N - 1)$ -dimensional binary vector x and for any $1 \leq k \leq N$,

$$S_N^k \begin{pmatrix} x_{N-1} \\ x_{N-2} \\ x_{N-3} \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ x_1 \end{pmatrix} = \begin{pmatrix} x_k \\ x_k \\ x_k \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ x_k \end{pmatrix} + \begin{pmatrix} x_{k-1} \\ x_{k-2} \\ \cdot \\ x_1 \\ 0 \\ x_{N-1} \\ x_{N-2} \\ \cdot \\ x_{k+1} \end{pmatrix}. \quad (16)$$

In the above formula $x_j \equiv 0$, unless $1 \leq j \leq (N - 1)$. Therefore, $S_N^k \neq Id$, for any $1 \leq k \leq N - 1$. Setting $k = N$ in the above formula, we get $S_N^N x = x$ for any x , which implies that $S_N^N = Id$. Therefore, the order of the matrix S_N is N .

(ii) The characteristic polynomial of S_N is $f(x) = \sum_{k=0}^{N-1} x^k$. (In order to prove this it is sufficient to notice that the matrix S_N is the companion matrix of the polynomial $f(x)$, [9]. As such, $f(x)$ is both the characteristic and the minimal polynomial of the matrix S_N .) Therefore,

$$f(S_N) = \sum_{k=0}^{N-1} S_N^k = 0.$$

Notice that the matrix S_N is non-degenerate as it has a positive order. If N is odd, we can re-write the characteristic polynomial as

$$f(S_N) = (Id + S_N)(1 + S_N + S_N^3 + \dots + S_N^{(N-3)}) + S_N^{N-1}$$

Therefore, the degeneracy of $Id + S_N$ will contradict the non-degeneracy of S_N . If N is even, the sum of all rows of $Id + S_N$ is zero, which implies degeneracy. **QED**

Lemma III-D.1 states that the matrix S_N generates the cyclic group Z_N and that the matrix $Id + S_N$ is non-degenerate for any odd N . Given that N is an odd prime, Corollary III-B.4 implies that using parity equations (12) with $g = S_N$, it is possible to protect N data disks against the

failure of any two disks. Furthermore, if $N > 3$ is a Fermat prime, $2N - 1$ data disks can be protected against double failure due to Lemma III-A.2.

We will refer to the RAID-6 system based on Sylvester matrix S_N as Z_N -RAID. Let us give several examples of such systems.

(1) Z_3 -RAID has been considered in subsections II-A, II-B. It can protect up to 3 information disks against double failure. As $N = 3$, protection of 5 information disks using extended Q -parity (15) is impossible.

(2) Using Z_{17} -RAID, one can protect up to $N = 17$ disks using Q -parity (12) and up to $2N - 1 = 33$ disks using extended Q -parity (15).

(3) Using Z_{257} -RAID, one can protect up to $N = 257$ disks using Q -parity (12) and up to $2N - 1 = 513$ disks using extended Q -parity (15).

It can be seen from (16), that the multiplication of data vectors with any power of the Sylvester matrix S_N requires one left and one right shift, one n -bit XOR and one AND only. Thus the operations of updating Q -parity and recovering data within Z_N -RAID does not require any special instructions, such as Galois field look-up tables for logarithms and products. As a result, the implementation of Z_N -RAID can in some cases be more efficient and quick than the implementation of the more conventional Reed-Solomon based RAID-6. In the next Section we will demonstrate the advantage of Z_N -RAID using an example of Linux kernel implementation of Z_{17} -RAID system.

IV. LINUX KERNEL Z_N -RAID IMPLEMENTATION

A. Syndrome Calculation for the Reed-Solomon RAID-6.

First, let us briefly recall the RAID-6 scheme based on Reed-Solomon code in the Galois field $GF(2^8)$, see [4] for more details. Let D_0, \dots, D_{N-1} be the bytes of data from N information disks. Then the parity bytes P and Q are computed as follows

$$\begin{aligned} P &= D_0 + D_1 + \dots + D_{N-1} \\ Q &= D_0 + gD_1 + \dots + g^{N-1}D_{N-1} \end{aligned} \quad (17)$$

where² $g = \{02\} \in GF(2^8)$.

The multiplication by $g = \{02\}$ can be viewed as the

² Algebraically, we use the standard representation in electronics: $GF(2^8) = GF(2)[x]/I$ where the ideal I is generated by $x^8 + x^4 + x^3 + x^2 + 1$ and $g = x + I$

following matrix multiplication.

$$\begin{bmatrix} y_0 \\ y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \end{bmatrix} = \begin{bmatrix} x_7 \\ x_0 \\ x_1 \oplus x_7 \\ x_2 \oplus x_7 \\ x_3 \oplus x_7 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} \quad (18)$$

Given (18), parity equations (17) become similar to (12). Indeed, the element g generates a cyclic group, so a 2-error correcting Reed-Solomon code is a partial case of a cone based RAID. However, Z_N -RAID has several advantages. For instance, using Sylvester matrices one can achieve a simpler implementation of matrix multiplication.

B. Linux Kernel Implementation of Syndrome Calculation

To compute the Q -parity, we rewrite (17) as

$$Q = D_0 + g(D_1 + g(\dots + g(D_{N-2} + gD_{N-1})\dots)) \quad (19)$$

which requires $(N - 1)$ multiplications by $g = \{02\}$.

The product y of a single byte x and $g = \{02\}$ can be implemented as follows.

```
uint8_t x, y;
y = (x << 1) ^ ((x & 0x80) ? 0x1d : 0x00);
```

Notice that $(x \& 0x80)$ picks out x_7 from x , so

$$((x \& 0x80) ? 0x1d : 0x00)$$

selects between the two bit patterns 00011101 and 00000000 depending on x_7 . Since the carry is discarded from $(x \ll 1)$,

$$\begin{aligned} (x \ll 1) &= [x_6, x_5, x_4, x_3, x_2, x_1, x_0, 0] \\ ((x \& 0x80) ? 0x1d : 0x00) &= \quad (20) \\ &= [0, 0, 0, x_7, x_7, x_7, 0, x_7]. \end{aligned}$$

We can also implement the multiplication as follows.

```
int8_t x, y;
y = (x + x) ^ (((x < 0) ? 0xff : 0x00) & 0x1d);
```

Here we treat the values as signed, rather than unsigned. Whilst this implementation appears more complex than the first (since it uses addition and comparison), it can efficiently be implemented using SIMD instructions on modern processors, such as MMX/SSE/SSE2/Altivec.

In particular, we will use the following four SSE2 instructions, which store the result in place of second operand.

```
pxor x, y      : y = x ^ y;
pand x, y      : y = x & y;
paddb x, y     : y = x + y;
pcmpgtb x, y   : y = (y > x) ? 0xff : 0x00;
```

Therefore we can implement a single multiplication with the following pseudo SSE2 assembler code. We assume that the variables y and c are initialized as $y = 0$ and $c = 0x1d$.

```
pcmpgtb x, y   : y = (x < 0) ? 0xff : 0x00; // (x < 0) ? 0xff : 0x00
paddb x, x     : x = x + x;
// x + x
pand c, y      : y = y & 0x1d;
// ((x < 0) ? 0xff : 0x00) & 0x1d
pxor x, y      : y = x ^ y;
// (x + x) ^
// (((x < 0) ? 0xff : 0x00) & 0x1d)
```

The comparison operation overwrites the constant 0 stored in y . Therefore, when we implement the complete algorithm we must recreate the constant before each multiplication. We can do it as follows.

```
pxor y, y      : y = y ^ y;
// y ^ y = 0
```

Besides the five instruction above we need three other instructions to complete the inner loop of the algorithm. They are multiply, fetch a new byte of data D and update the parity variables P and Q .

$$\begin{aligned} P &= D + P \\ Q &= D + gQ \end{aligned} \quad (21)$$

The complete algorithm requires the following eight instructions.

```
pxor y, y      : y = y ^ y;
// y ^ y = 0
pcmpgtb q, y   : y = (q < 0) ? 0xff : 0x00; // (q < 0) ? 0xff : 0x00
paddb q, q     : q = q + q;
// q + q
pand c, y      : y = y & 0x1d;
// ((q < 0) ? 0xff : 0x00) & 0x1d
pxor y, q      : q = q ^ y;
// g.q = (q + q) ^
// (((q < 0) ? 0xff : 0x00) & 0x1d)
movdqa d[i], d : d = d[i]
// d[i]
pxor d, q      : q = d ^ q;
// d[i] ^ p
pxor d, p      : p = d ^ p;
// d[i] ^ g.q
```

We can gain a further increase in speed by partially unrolling the 'for' loop around the inner loop.

Below are the results of the Linux kernel RAID-6 algorithm selection program, which aims to select the fastest implementation of the algorithm. Algorithms using the CPU/MMX/SSE/SSE2 instructions with various levels of unrolling are compared.

```
raid6: int32x1      694 MB/s
```

```

raid6: int32x2    939 MB/s
raid6: int32x4    635 MB/s
raid6: int32x8    505 MB/s
raid6: mmxx1     1893 MB/s
raid6: mmxx2     2025 MB/s
raid6: sselx1    1200 MB/s
raid6: sselx2    2000 MB/s
raid6: sse2x1    1850 MB/s
raid6: sse2x2    2702 MB/s

```

Results were obtained from a 2.8 GHz Intel Pentium 4 (x86).

C. Reconstruction

We consider a situation that two data disks D_x and D_y have failed. We must reconstruct D_x and D_y from the remaining data disks D_i ($i \neq x, y$) and the parity disks P and Q , see (17). Let us define P_{xy} and Q_{xy} as the syndromes under an assumption that the failed disks were zero.

$$P_{xy} = \sum_{i \neq x, y} D_i \quad Q_{xy} = \sum_{i \neq x, y} g^i D_i \quad (22)$$

Let us rewrite (17) in the light of (22).

$$\begin{aligned} D_x + D_y &= P + P_{xy} \\ g^x D_x + g^y D_y &= Q + Q_{xy} \end{aligned} \quad (23)$$

Let us define

$$\begin{aligned} A &= (1 + g^{y-x})^{-1} \\ B &= g^{-x} (1 + g^{y-x})^{-1}. \end{aligned} \quad (24)$$

Now we eliminate D_x from equations (23).

$$\begin{aligned} (1 + g^{y-x})D_y &= (P + P_{xy}) + g^{-x}(Q + Q_{xy}) \\ D_y &= (1 + g^{y-x})^{-1}(P + P_{xy}) + \\ &\quad + g^{-x}(1 + g^{y-x})^{-1}(Q + Q_{xy}) \\ D_y &= A(P + P_{xy}) + B(Q + Q_{xy}) \end{aligned} \quad (25)$$

Finally, D_x is computed from D_y by the back substitution into (23).

$$D_x = D_y + (P + P_{xy}) \quad (26)$$

D. Linux Kernel Implementation of Reconstruction

We need to compute the following values in $GF(256)$.

$$\begin{aligned} A &= (1 + g^{y-x})^{-1} \\ B &= g^{-x} (1 + g^{y-x})^{-1} = (g^x + g^y)^{-1} \\ D_y &= A(P + P_{xy}) + B(Q + Q_{xy}) \\ D_x &= D_y + (P + P_{xy}) \end{aligned} \quad (27)$$

It is worth pointing out that for specific x and y , we only need to compute A and B once.

The Linux kernel provides the following lookup tables.

```

raid6_gfmul[256][256] : xy
raid6_gfexp[256]      : g^x
raid6_gfinv[256]     : x^-1
raid6_gfexi[256]     : (1 + g^x)^-1

```

Using this, we compute A and B as follows.

```

A = raid6_gfexi[y-x]
B = raid6_gfinv[raid6_gfexp[x] ^
raid6_gfexp[y]]

```

To reconstruct D_x and D_y we start by constructing P_{xy} and Q_{xy} using the standard syndrome code. Then we execute the following code.

```

dP = P ^ Pxy;
// P + Pxy
dQ = Q ^ Qxy;
// Q + Qxy
Dy = raid6_gfmul[A][dP] ^
raid6_gfmul[B][dQ];
// A(P + Pxy) + B(Q + Qxy)
Dx = Dy ^ dP;
// Dy + (P + Pxy)

```

V. Z_{17} RAID IMPLEMENTATION

The multiplication by the Sylvester matrix g looks like

$$\begin{bmatrix} y_0 \\ y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \\ y_9 \\ y_{10} \\ y_{11} \\ y_{12} \\ y_{13} \\ y_{14} \\ y_{15} \end{bmatrix} = g \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \\ x_9 \\ x_{10} \\ x_{11} \\ x_{12} \\ x_{13} \\ x_{14} \\ x_{15} \end{bmatrix} = \begin{bmatrix} x_{15} \\ x_0 \oplus x_{15} \\ x_1 \oplus x_{15} \\ x_2 \oplus x_{15} \\ x_3 \oplus x_{15} \\ x_4 \oplus x_{15} \\ x_5 \oplus x_{15} \\ x_6 \oplus x_{15} \\ x_7 \oplus x_{15} \\ x_8 \oplus x_{15} \\ x_9 \oplus x_{15} \\ x_{10} \oplus x_{15} \\ x_{11} \oplus x_{15} \\ x_{12} \oplus x_{15} \\ x_{13} \oplus x_{15} \\ x_{14} \oplus x_{15} \end{bmatrix} \quad (28)$$

where

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

We implement the multiplication of a double byte $y = gx$ as follows.

```
int16_t x, y;
```

```
y = (x + x) ^ ((x < 0) ? 0xffff : 0x0000);
```

We can implement this in assembler using the following seven instructions.

```
pxor y, y      : y = y ^ y;
// y ^ y = 0
pcmpgtw q, y   : y = (q < 0) ? 0xffff :
0x0000; // (q < 0) ? 0xffff : 0x0000
paddw q, q     : q = q + q;
// q + q
pxor y, q      : q = q ^ y;
// g.q = (q + q) ^

// ((q < 0) ? 0xffff : 0x0000)
movdqa d[i], d : d = d[i]
// d[i]
pxor d, q      : q = d ^ q;
// d[i] ^ p
pxor d, p      : p = d ^ p;
// d[i] ^ g.q
```

Below are the results of the Linux kernel DP RAID algorithm selection program, modified to support Z_{17} RAID, which aims to select the fastest implementation of the algorithm. Algorithms using CPU/MMX/SSE/SSE2 instructions with various levels of unrolling are compared.

raid6: int32x1	766 MB/s
raid6: int32x2	854 MB/s
raid6: int32x4	838 MB/s
raid6: int32x8	604 MB/s
raid6: mmxx1	2117 MB/s
raid6: mmxx2	2301 MB/s
raid6: sselx1	1284 MB/s
raid6: sselx2	2263 MB/s
raid6: sse2x1	2357 MB/s
raid6: sse2x2	3160 MB/s

Results were obtained from a 2.8 GHz Intel Pentium 4 (x86).

Comparing the above results against the standard Linux kernel results shows an average of 14.5% speed increase and an increase of 16.9% for the fastest sse2x2 implementation. This is consistent with the theoretical increase of 14.3% for seven instructions instead of eight instructions. It is worth mentioning that no look-up tables have been used to implement Z_{17} -RAID.

A. Z_N RAID Reconstruction

We need to compute the following matrices and vectors.

$$\begin{aligned}
 A &= (1 + g^{y-x})^{-1} \\
 B &= g^{-x}(1 + g^{y-x})^{-1} = (g^x + g^y)^{-1} \\
 D_y &= A(P + P_{xy}) + B(Q + Q_{xy}) \\
 D_x &= D_y + (P + P_{xy})
 \end{aligned} \tag{29}$$

We rewrite them as follows.

$$\begin{aligned}
 z &= y - x \\
 \Delta P &= P + P_{xy} \\
 \Delta Q &= Q + Q_{xy} \\
 D_y &= (1 + g^z)^{-1} \Delta P + g^{-x}(1 + g^z)^{-1} \Delta Q \\
 D_x &= D_y + \Delta P
 \end{aligned} \tag{30}$$

Let us notice the following identities.

$$\begin{aligned}
 g^{17} &= 1 \\
 g^{-k} &= g^{17-k} \\
 (1 + g)^{-1} &= 1 + g^2 + g^4 + \dots + g^{16}
 \end{aligned} \tag{31}$$

Using them we derive new identities

$$(1 + g^z)^{-1} = 1 + g^{2z} + g^{4z} + \dots + g^{16z} \tag{32}$$

and

$$\begin{aligned}
 g^{-x}(1 + g^z)^{-1} &= g^{17-x}(1 + g^z)^{-1} \\
 &= g^{17-x}(1 + g^{2z} + g^{4z} + \dots + g^{16z}).
 \end{aligned} \tag{33}$$

Consequently, we need to compute

$$\begin{aligned}
 (1 + g^z)^{-1} \Delta P &= (1 + g^{2z} + g^{4z} + \dots + g^{16z}) \Delta P \\
 &= 1 + g^{2z}(1 + g^{2z}(1 + g^{2z}(1 + g^{2z}(1 + g^{2z}(1 + g^{2z} \Delta P))))))
 \end{aligned} \tag{34}$$

and

$$\begin{aligned}
 g^{-x}(1 + g^z)^{-1} \Delta Q &= g^{17-x}(1 + g^{2z} + g^{4z} + \dots + g^{16z}) \Delta Q \\
 &= g^{17-x}(1 + g^{2z}(1 + g^{2z}(1 + g^{2z}(1 + g^{2z}(1 + g^{2z}(1 + g^{2z} \Delta Q))))))
 \end{aligned} \tag{35}$$

Both (34) and (35) require only one principle operation, multiplication by g^k .

The multiplication of a single word $y = g^k x$ for $1 \leq k \leq 16$ can be implemented as follows.

```
int16_t x, y;
y = (x << k) ^ (x >> (17 - k)) ^ (((x << (k - 1)) < 0) ? 0xffff : 0x0000);
```

We precompute $m = k - 1$ and $n = 17 - k$ as they remain constant during reconstruction. This leads to the following assembler implementation.

```
pxor y, y      : y = 0
movdqa x, z    : z = x
psllw m, z     : z = x << (k-1)
pcmpgtw z, y   : y = (((x << (k-1)) < 0) ?
0xffff : 0x0000)
paddw z, z     : z = x << k
pxor z, y      : y = (((x << (k-1)) < 0) ?
0xffff : 0x0000) ^ (x << k)
psrlw n, x     : x = x >> (17-k)
pxor x, y      : y = g ^ k x
```

Below is a table showing benchmark results of complete reconstruction algorithm implemented using SSE2 assembler

and the standard Linux kernel lookup table reconstruction implementation, for the cases of double data disk failure, double disk failure of one data disk and the P-parity disk, and double parity disk failure. Note the data represents time taken to complete benchmark, so lower is better.

Failure	DD	DP	PQ
DP-RAID	2917	2771	905
Z_{17} -RAID	2711	1274	809

Comparing the complete reconstruction algorithm implemented using SSE2 assembler against the standard Linux kernel lookup table implementation, shows approximately 7% speed increase for *DD* failure, 54% speed increase for *DP* failure and 11% speed increase for *PQ* failure.

VI. CONCLUSIONS.

In this paper we have demonstrated that *cones* provide a natural framework for the design of RAID. They provide a flexible approach that can be used to design a part system. It is worth further theoretical investigation what other examples of cones can be constructed or what the maximal possible size of a cone is.

We have also demonstrated that cyclic groups give rise to natural and convenient to operate examples of cones. One particular advantage is that Z_N -RAID does not require support of the Galois field operations.

On the practical side, Z_{17} -RAID and Z_{257} -RAID are breakthrough techniques that show at least 10% improvement during simulations compared to DP-RAID.

ACKNOWLEDGMENT

The authors would like to thank Arithmatica Limited for the opportunity to use its research facilities. The authors would also like to thank Bob Maddock of Xyratex Technology Limited for valuable information.

REFERENCES

- [1] Next Generation Mobile Hard Disk Drives, *white paper*, Fujitsu Inc., http://www.fujitsu.com/downloads/COMP/fcpa/hdd/sata-mobile-ext-duty_wp.pdf.
- [2] Robert Maddock, Nigel Hart and Tom Kean, Surviving two disk failures. Introducing various 'RAID-6' implementations, *white paper*, Xyratex Ltd., May 2005, http://www.xyratex.com/pdfs/whitepapers~Xyratex_White_Paper_Raid_6_Implementations.pdf.
- [3] Wikipedia, the free online encyclopedia, *RAID*, http://en.wikipedia.org/wiki~Redundant_array_of_independent_disks#RAID_6.
- [4] H. Peter Anvin, The mathematics of RAID-6, *online paper*, <http://kernel.org/pub/linux/kernel/people/hpa/raid6.pdf>.
- [5] Chris Lueth, RAID- DP^{TM} : NetApp implementation of RAID double parity for data protection, *online paper*, <http://www.netapp.com/library/tr/3298.pdf>.
- [6] David J. C. Mackay, *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, 2003.
- [7] Serge Lang, *Algebra*, Addison-Wesley Publishing Company, 1993.
- [8] R. M. Robinson, Mersenne and Fermat Numbers, *Proc. Amer. Math. Soc.* vol. 5, 842-846, 1954.
- [9] R. Lidl and H. Niederreiter, *Finite Fields*, Second Edition, Cambridge University Press, 1997.