Diffusion Matrices from Algebraic-Geometry Codes with Efficient SIMD Implementation

Daniel Augot^{1,2}, Pierre-Alain Fouque^{3,4}, and Pierre Karpman^{1,2,5} (✉)

¹ Inria, Saclay, France

`{daniel.augot,pierre.karpman}@inria.fr`

² LIX — École Polytechnique, Palaiseau, France

³ Université de Rennes 1, Rennes, France

`pierre-alain.fouque@irisa.fr`

⁴ Institut Universitaire de France, Paris, France

⁵ Nanyang Technological University, Singapore, Singapore

Abstract. This paper investigates large linear mappings with very good diffusion and efficient software implementations, that can be used as part of a block cipher design. The mappings are derived from linear codes over a small field (typically $\mathbf{F}_{2^4}$) with a high dimension (typically 16) and a high minimum distance. This results in diffusion matrices with equally high dimension and a large branch number. Because we aim for parameters for which no MDS code is known to exist, we propose to use more flexible *algebraic-geometry* codes.

We present two simple yet efficient algorithms for the software implementation of matrix-vector multiplication in this context, and derive conditions on the generator matrices of the codes to yield efficient encoders. We then specify an appropriate code and use its automorphisms as well as random sampling to find good such matrices.

We provide concrete examples of parameters and implementations, and the corresponding assembly code. We also give performance figures in an example of application which show the interest of our approach.

Keywords: Diffusion matrix · Algebraic-geometry codes · Algebraic curves · SIMD · Vector implementation · SHARK

1 Introduction

The use of *MDS* matrices over finite fields as a linear mapping in block cipher design is an old trend, followed by many prominent algorithms such as the AES/Rijndael family [7]. These matrices are called MDS as they are derived from *maximum distance separable* linear error-correcting codes, which achieve the highest minimum distance possible for a given length and dimension. This notion of minimum distance coincides with the one of *branch number* of a mapping [7], which is a measure of the effectiveness of a diffusion layer. MDS matrices thus have an optimal diffusion, in a cryptographic sense, which makes them attractive for cipher designs.

The good security properties that can be derived from MDS matrices are often counter-balanced by the cost of their computation. The standard matrix-vector product is quadratic in the dimension of the vector, and finite field operations are not always efficient. For that reason, there is often a focus on finding matrices allowing efficient implementations. For instance, the AES matrix is circulant and has small coefficients. More recently, the PHOTON hash function [9] introduced the use of matrices that can be obtained as the power of a companion matrix, which sparsity may be useful in lightweight hardware implementations. The topic of finding such so-called recursive diffusion layers has been quite active in the past years, and led to a series of papers investigating some of their various aspects [2, 18, 23]. One of the most recent developments shows how to systematically construct some of these matrices from BCH codes [1]. This allows in particular to construct very large recursive MDS matrices, for instance of dimension 16 over $\mathbf{F}_{2^8}$. This defines a linear mapping over a full 128-bit block with excellent diffusion properties, at a moderate hardware implementation cost.

As interesting as it may be in hardware, the cost in software of a large linear mapping tends to make these designs rather less attractive than more balanced solutions. An early attempt to use a large matrix was the block cipher SHARK, a Rijndael predecessor [16]. It is a 64-bit cipher which uses an MDS matrix of dimension 8 over $\mathbf{F}_{2^8}$ for its linear diffusion. The usual technique for implementing such a mapping in software is to rely on a table of precomputed multiples of the matrix rows. However, table-based implementations now tend to be frowned upon as they may lead to *timing attacks* [21], and this could leave ciphers with a structure similar to SHARK's without reasonable software implementations when resistance to these attacks is required. Yet, such designs also have advantages of their own; their diffusion acts on the whole state at every round, and therefore makes structural attacks harder, while also ensuring that many S-Boxes are kept active. Additionally, the simplicity of the structure makes it arguably easier to analyze than in the case of most ciphers.

Our Contributions. In this work, we revisit the use of a *SHARK structure* for block cipher design and endeavour to find good matrices and appropriate algorithms to achieve both a linear mapping with very good diffusion and efficient software implementations that are not prone to timing attacks. To be more specific on this latter point, we target software running on 32 or 64-bit CPUs featuring an SIMD vector unit.

An interesting way of trying to meet both of these goals is to decrease the size of the field from $\mathbf{F}_{2^8}$ to $\mathbf{F}_{2^4}$. However, according to the *MDS conjecture*, there is no MDS code over $\mathbf{F}_{2^4}$ of length greater than 17, and no such code is known [15]. Because a diffusion matrix of dimension n is typically obtained from a code of length $2n$, MDS matrices over $\mathbf{F}_{2^4}$ are therefore restricted to dimensions less than 8. Hence, the prospect of finding an MDS matrix over $\mathbf{F}_{2^4}$ diffusing on more than $8 \times 4 = 32$ bits is hopeless. Obviously, 32 bits is not

enough for a large mapping *à la* SHARK. We must therefore search for codes with a slightly smaller minimum distance in the hope that they can be made longer.

Our proposed solution to this problem is to use *algebraic-geometry codes* [22], as they precisely offer this tradeoff. One way of defining these codes is as evaluation codes on algebraic curves; thus our proposal brings a nice connection between these objects and symmetric cryptography. Although elliptic and hyperelliptic curves are now commonplace in public-key cryptography, we show a rare application of an hyperelliptic curve to the design of block ciphers. We present a specific code of length 32 and dimension 16 over $\mathbf{F}_{2^4}$ with minimum distance 15, which is only 2 less than what an MDS code would achieve. This lets us deriving a very good diffusion matrix on $16 \times 4 = 64$ bits in a straightforward way. Interestingly, this matrix can also be applied to vectors over an extension of $\mathbf{F}_{2^4}$ such as $\mathbf{F}_{2^8}$, while keeping the same good diffusion properties. This allows for instance to increase the diffusion to $16 \times 8 = 128$ bits.

We also study two simple yet efficient algorithms for implementing the matrix-vector multiplication needed in a SHARK structure, when a *vector permute* instruction is available. From one of these, we derive conditions on the matrix to make the product faster to compute, in the form of a cost function; we then search for matrices with a low cost, both randomly, and by using automorphisms of the code and of the hyperelliptic curve on which it is based. The use of codes automorphisms to derive efficient encoders is not new [6, 11], but it is not generally applied to the architecture and dimensions that we consider in our case.

We conclude this paper by presenting examples of performance figures of assembly implementations of our algorithms when used as the linear mapping of a block cipher.

Structure of the Paper. We start with a few background notions in Sect. 2. We then present our algorithms for matrix-vector multiplication and their context in Sect. 3, and derive a cost function for the implementation of matrices. This is followed by the definition of the algebraic-geometry code used in our proposed linear mapping, and a discussion of how to derive efficient encoders in Sect. 4. We conclude with insights into the performance of the mapping when used over both $\mathbf{F}_{2^4}$ and $\mathbf{F}_{2^8}$ in Sect. 5.

2 Preliminaries

We note $\mathbf{F}_{2^m}$ the finite field with 2^m elements. We often consider $\mathbf{F}_{2^4}$, and implicitly use this specific field if not mentioned otherwise. W.l.o.g. we use the representation $\mathbf{F}_{2^4} \cong \mathbf{F}_2[\alpha]/(\alpha^4 + \alpha + 1)$. We freely use “integer representation” for elements of $\mathbf{F}_{2^4}$ by writing $n \in \{0 \dots 15\} = \sum_{i=0}^3 a_i 2^i$ to represent the element $x \in \mathbf{F}_{2^4} = \sum_{i=0}^3 a_i \alpha^i$.

Bold variables denote vectors (in the sense of elements of a vector space), and subscripts are used to denote their i^{th} coordinate, starting from zero. For instance, $\mathbf{x} = (1, 2, 7)$ and $\mathbf{x}_2 = 7$. If M is a matrix of n columns, we call

$\mathbf{m}^i = (M_{i,j}, j = 0 \dots n-1)$ the row vector formed from the coefficients of its i^{th} row. We use angle brackets “ $\langle \rangle$ ” and “ $\langle \rangle$ ” to write ordered sets.

Arrays, or tables, (in the sense of software data structures) are denoted by regular variables such as x or T , and their elements are accessed by using square brackets. For instance, $T[i]$ is the i^{th} element of the table T , starting from zero.

We conclude with two definitions.

Definition 1 (Systematic form and dual of a code). *Let the code $\mathcal{C}$ be an $[n, k, d]_{\mathbf{F}_{2^m}}$ code of length n , dimension k and minimum distance d with symbols in $\mathbf{F}_{2^m}$. A generator matrix for $\mathcal{C}$ is in systematic form if it is of the form $(I_k \ A)$, with I_k the identity matrix of dimension k and A a matrix of k rows and $n - k$ columns. A systematic generator matrix for the dual of $\mathcal{C}$ is given by $(I_{n-k} \ A^t)$.*

Definition 2 (Branch number [7]). *Let A be the matrix of a linear mapping over $\mathbf{F}_{2^m}$, and $w_m(\mathbf{x})$ be the number of non-zero positions of the vector $\mathbf{x}$ over $\mathbf{F}_{2^m}$. Then the differential branch number of A is equal to $\min_{\mathbf{x} \neq 0} (w_m(x) + w_m(A(x)))$, and the linear branch number of A is equal to $\min_{\mathbf{x} \neq 0} (w_m(x) + w_m(A^t(x)))$.*

Note that if A is such that $(I_k \ A)$ is a generator matrix of a code of minimum distance d which dual code has minimum distance d' , then A has a differential (resp. linear) branch number of d (resp. d').

3 Efficient Algorithms for Matrix-Vector Multiplication

This section presents software algorithms for matrix-vector multiplication over $\mathbf{F}_{2^4}$. We focus on square matrices of dimension 16. This naturally defines linear operations on 64 bits, which can also be extended to 128 bits, as it will be made clear in Sect. 5. Both cases are a common block size for block ciphers.

Targeted Architecture. The algorithms in this section target CPUs featuring vector instructions, including in particular a *vector shuffle* instruction such as Intel’s `pshufb` from the SSSE3 instruction set extension [12]. These instructions are now widespread and have already been used successfully in fast cryptographic implementations, see *e.g.* [4, 10, 20]. We mostly considered SSSE3 when designing the algorithms, but other processor architectures do feature vector instructions. This is for instance the case of ARM’s NEON extensions, which may also yield efficient implementations, see *e.g.* [5]. We do not consider these explicitly in this paper, however.

Because it plays an important role in our algorithms, we briefly recall the semantics of `pshufb`. The `pshufb` instruction takes two 128-bit inputs¹. The first (the destination operand) is an `xmm` SSE vector register which logically represents a vector of 16 bytes. The second (the source operand) is either a similar

¹ The instruction can actually also be used on 64-bit operands, but we do not consider this possibility here.

xmm register, or a 128-bit memory location. The result of calling `pshufb x y` is to overwrite the input x with the vector x' defined by:

$$x'[i] = \begin{cases} x[\lfloor y[i] \rfloor_4] & \text{if the most significant bit of } y[i] \text{ is not set} \\ 0 & \text{otherwise} \end{cases}$$

where $\lfloor \cdot \rfloor_4$ denotes truncation to the 4 least significant bits. This instruction allows to arbitrarily *shuffle* a vector according to a mask, with possible repetition and omission of some of the vector values². Notice that this instruction can also be used to perform 16 parallel 4-to-8-bit table lookups: let us call T this table; take as first operand to `pshufb` the vector $x = (T[i], i = 0 \dots 15)$, as second operand the vector $y = (a, b, c, d, \dots)$ on which to perform the lookup; then we see that the first byte of the result is $x[y[0]] = T[a]$, the second is $x[y[1]] = T[b]$, etc.

Finally, there is a three-operand variant of this instruction in the more recent AVX instruction set and onward [12], which allows not to overwrite the first operand.

Targeted Properties. In this paper we focus solely on algorithms that can easily be implemented in a way that makes them immune to timing-attacks [21]. Specifically, we consider the matrix as a known constant but the vector as a secret, and we wish to perform the multiplication without secret-dependent branches or memory accesses. It might not always be important to be immune (or even partially resistant) to this type of attacks, but we consider that it should be important for any cryptographic primitive or structure to possibly be implemented in such a way. Hence we try to find efficient such implementations for the SHARK structure and therefore for dense matrix-vector multiplications.

We now go on to describe the algorithms. In all of the remainder of this section, $\mathbf{x}$ and $\mathbf{y}$ are two (column) vectors of $\mathbf{F}_{2^4}^{16}$, and M a matrix of $\mathcal{M}_{16}(\mathbf{F}_{2^4})$. We first briefly recall the principle of table implementations, which are unsatisfactory when timing attacks are taken into account.

3.1 Table Implementation

We wish to compute $\mathbf{y} = M \cdot \mathbf{x}$. The idea behind this algorithm is to use table lookups to perform the equivalent multiplication $\mathbf{y}^t = \mathbf{x}^t \cdot M^t$, *i.e.* $\mathbf{y}^t = \sum_{i=0}^{15} \mathbf{x}_i \cdot (\mathbf{m}^t)^i$ (where $(\mathbf{m}^t)^i$ is the i^{th} row of M^t). This can be computed efficiently by tabulating beforehand the products $\lambda \cdot (\mathbf{m}^t)^i$, $\lambda \in \mathbf{F}_{2^4}$ (resulting in 16 tables, each of 16 entries of 64 bits), and then for each multiplication by accessing the table for $(\mathbf{m}^t)^i$ at the index $\mathbf{x}_i$ and summing all the retrieved table entries together. This only requires 16 table lookups per multiplication. However, the memory accesses depend on the value of $\mathbf{x}$, which makes this algorithm inherently vulnerable to timing attacks.

² We will use the word *shuffle* with this precise meaning in the remainder of this paper.

Note that there is a more memory-efficient alternative implementation of this algorithm which consists in computing each term $\lambda \cdot (\mathbf{m}^t)^i$ with a single `pshufb` instruction instead of using a table-lookup. In that case, only the 16 multiplication tables need to be stored, but their accesses still depend on the secret value $\mathbf{x}$.

3.2 A Generic Constant-Time Algorithm

We now describe our first algorithm, which can be seen as a variant of table multiplication that is immune to timing attacks. The idea consists again in computing the right multiplication $\mathbf{y}^t = \mathbf{x}^t \cdot M^t$, *i.e.* $\mathbf{y}^t = \sum_{i=0}^{15} \mathbf{x}_i \cdot (\mathbf{m}^t)^i$. However, instead of tabulating the results of the scalar multiplication of the matrix rows $(\mathbf{m}^t)^i$, those are always recomputed, in a way that does not explicitly depend on the value of the scalar.

Description of Algorithm 1. We give the complete description of Algorithm 1 in the full version of the paper [3], and focus here on the intuition. We want to perform the scalar multiplication $\lambda \cdot \mathbf{z}$ for an unknown scalar λ and a known, constant vector $\mathbf{z}$, over $\mathbf{F}_{2^4}$. Let us write λ as the polynomial $\lambda_3 \cdot \alpha^3 + \lambda_2 \cdot \alpha^2 + \lambda_1 \cdot \alpha + \lambda_0$ with coefficients in $\mathbf{F}_2$. Then, the result of $\lambda \cdot \mathbf{z}$ is simply $\lambda_3 \cdot (\alpha^3 \cdot \mathbf{z}) + \lambda_2 \cdot (\alpha^2 \cdot \mathbf{z}) + \lambda_1 \cdot (\alpha \cdot \mathbf{z}) + \lambda_0 \cdot \mathbf{z}$. Thus we just need to precompute the products $\alpha^i \cdot \mathbf{z}$, select the right ones with respect to the binary representation of λ , and add these together. This can easily be achieved thanks to a *broadcast* function defined as:

$$\text{broadcast}(x, i)_n = \begin{cases} \mathbf{1}_n & \text{if the } i^{\text{th}} \text{ bit of } x \text{ is set} \\ \mathbf{0}_n & \text{otherwise} \end{cases}$$

where $\mathbf{1}_n$ and $\mathbf{0}_n$ denote the n -bit binary string made all of one and all of zero respectively. The full algorithm then just consists in using this scalar-vector multiplication 16 times, one for each row of the matrix.

Implementation of Algorithm 1 with SSSE3 Instructions. We now consider how to efficiently implement Algorithm 1 in practice. The only non-trivial operation is the *broadcast* function, and we show that this can be performed with only one or two `pshufb` instructions.

To compute $\text{broadcast}(\lambda, i)_{64}$, with λ a 4-bit value, we can use a single `pshufb` with first operand x , such that $x[j] = 11111111_2$ if the i^{th} bit of j is set and 0 otherwise, and with second operand $y = (\lambda, \lambda, \lambda, \dots)$. The result of `pshufb` $x \ y$ is indeed $(x[\lambda], x[\lambda], \dots)$ which is $\mathbf{1}_{64}$ if the i^{th} bit of λ is set, and $\mathbf{0}_{64}$ otherwise, that is $\text{broadcast}(\lambda, i)_{64}$.

In practice, the vector x can conveniently be constructed offline and stored in memory, but the vector y might not be readily available before performing this computation³. However, it can easily be computed thanks to an additional `pshufb`. Alternatively, if the above computation is done with a vector

³ And because it depends on what we assume to be a secret value, it cannot either be fetched from memory.

$y = (\lambda, ?, ?, \dots)$ instead (with $?$ denoting unknown values) and call z its result $(x[\lambda], ?, ?, \dots)$, then we have $\text{broadcast}(\lambda, i)_n = \text{pshufb } z(0, 0, \dots)$.

In the specific case of matrices of dimension 16 over $\mathbf{F}_{2^4}$, one can take advantage of the 128-bit wide `xmm` registers by interleaving, say, $8 \cdot \mathbf{x}$ with $4 \cdot \mathbf{x}$, and $2 \cdot \mathbf{x}$ with $\mathbf{x}$, and by computing a slightly more complex version of the *broadcast* function $\text{broadcast}(x, i, j)_{2n}$ which interleaves $\text{broadcast}(x, i)_n$ with $\text{broadcast}(x, j)_n$. In that case, an implementation of one step of algorithm 1 only requires two *broadcast* calls, two logical *and*, folding back the interleaved vectors (which only needs a couple of logical shift and exclusive or), and adding the folded vectors together. We give a snippet of such an implementation in the full version [3].

3.3 A Faster Algorithm Exploiting Matrix Structure

The above algorithm is already reasonably efficient, and has the advantage of being completely generic w.r.t. the matrix. Yet, better solutions may exist in more specific cases. We present here an alternative that can be much faster when the matrix possesses a particular structure.

The idea behind this second algorithm is to take advantage of the fact that in a matrix-vector product, the same constant values may be used many times in finite-field multiplications. Hence, we try to take advantage of this fact by performing those in parallel. The fact that we now focus on multiplications by constants (*i.e.* matrix coefficients) allows us to compute these multiplications with a single `pshufb` instead of using the process from Algorithm 1.

Description of Algorithm 2. We give the complete description of Algorithm 2 in the full version [3], and focus here on the intuition. Let us first consider a small example, and compute $M \cdot \mathbf{x}$ defined as:

$$\begin{pmatrix} 1 & 0 & 2 & 2 \\ 3 & 1 & 2 & 3 \\ 2 & 3 & 3 & 2 \\ 0 & 2 & 3 & 1 \end{pmatrix} \cdot \begin{pmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \end{pmatrix}. \quad (1)$$

It is obvious that this is equal to:

$$\begin{pmatrix} x_0 \\ x_1 \\ 0 \\ x_3 \end{pmatrix} + 2 \cdot \begin{pmatrix} x_2 \\ x_2 \\ x_0 \\ x_1 \end{pmatrix} + 2 \cdot \begin{pmatrix} x_3 \\ 0 \\ x_3 \\ 0 \end{pmatrix} + 3 \cdot \begin{pmatrix} 0 \\ x_0 \\ x_1 \\ x_2 \end{pmatrix} + 3 \cdot \begin{pmatrix} 0 \\ x_3 \\ x_2 \\ 0 \end{pmatrix}, \quad (2)$$

where both the constant multiplications of the vector $(x_0 \ x_1 \ x_2 \ x_3)^t$ and the shuffles of its coefficients can be computed with a single `pshufb` instruction each, while none of these operations directly depends on the value of the vector. This type of decomposition can be done for any matrix, but the number of operations depends on the value of its coefficients.

We now sketch one way of obtaining an optimal decomposition as above. We consider a matrix product $M \cdot \mathbf{x}$ with M constant and $\mathbf{x}$ unknown, where $\mathbf{x}$ is

seen as the formal arrangement of variables $\mathbf{x}_i$. Let us define $\mathcal{S}(M, \gamma)$ as one of the minimal sets of shuffles of coefficients of $\mathbf{x}$, such that there exists a unique vector $\mathbf{z} \in \mathcal{S}(M, \gamma)$ with $\mathbf{z}_i = \mathbf{x}_j$ iff $M_{i,j} = \gamma$. For instance, in the above example, we have $\mathcal{S}(M, 2) = \{(x_2 \ x_2 \ x_0 \ x_1)^t, (x_3 \ 0 \ x_3 \ 0)^t\}$. Equivalently, we could have taken $\mathcal{S}(M, 2) = \{(x_3 \ x_2 \ x_3 \ 0)^t, (x_2 \ 0 \ x_0 \ x_1)^t\}$. These sets are straightforward to compute from this particular matrix, and so are they in the general case.

From the definition of $\mathcal{S}$, it is clear that we have:

$$M \cdot \mathbf{x} = \sum_{\gamma \in \mathbf{F}_{2^4}^*} \sum_{\mathbf{s} \in \mathcal{S}(M, \gamma)} \gamma \cdot \mathbf{s}. \quad (3)$$

Once the values of the sets $\mathcal{S}$ have been determined, it is clear that we only need to compute this sum to get our result, and this is precisely what this second algorithm does.

Cost of Algorithm 2. The cost of computing a matrix-vector product with Algorithm 2 depends on the coefficients of the matrix, since the size of the sets $\mathcal{S}(M, \gamma)$ depends both on the density of the matrix and of how its coefficients are arranged.

If we assume that a vector implementation of this algorithm is used, and if the dimension and the field of the matrix are well chosen, we can assume that both the scalar multiplication of $\mathbf{x}$ by a constant and its shuffles can be computed with a single `pshufb` and a few ancillary instructions. Hence, we can define a cost function for a matrix with respect to its implementation with Algorithm 2 to be $\text{cost2}(M) = (\sum_{\gamma \in \mathbf{F}_{2^4}^*} \mathbb{1}(\mathcal{S}(M, \gamma)) + \#\mathcal{S}(M, \gamma)) - \mathbb{1}(\mathcal{S}(M, 1))$, where $\mathbb{1}(\mathcal{E})$ with $\mathcal{E}$ a set is one if $\mathcal{E} \neq \emptyset$, and zero otherwise. We may notice that $\#\mathcal{S}(M, \gamma)$ is equal to the maximum number of occurrence of γ in a single row of M , and the cost2 function is therefore easy to compute. As an example the cost of the matrix M from Eq. 1 is 7.

In order to find matrices that minimize the cost2 function, we would like to minimize the sum of the maximum number of occurrence of γ for every $\gamma \in \mathbf{F}_{2^4}^*$. A simple observation is that for matrices with the same number of non-zero coefficients, this amount is minimal when every row can be deduced by permutation of a single one; an important particular case being the one of *circulant* matrices. More generally, we can heuristically hope that the cost of a matrix will be low if all of its rows can be deduced by permutation of a small subset thereof.

We can try to estimate the minimum cost for an arbitrary dense circulant matrix of dimension 16 over $\mathbf{F}_{2^4}$. It is fair to assume that nearly all of the values of $\mathbf{F}_{2^4}$ should appear as coefficients of such a matrix, 14 of them needing a multiplication. Additionally, 15–16 permutations are needed if all the rows are to be different. Hence we can assume that the cost2 function of such a matrix is about 30.

Finally, let us notice that special cases of this algorithm have already been used for circulant matrices, namely in the case of the AES MixColumn matrix [4, 10].

Implementation of Algorithm 2 with SSSE3 Instructions. The implementation of Algorithm 2 is straightforward. We refer to the full version for a small code snippet [3].

3.4 Performance

In Table 2 of Sect. 5, we give a few performance figures for ciphers with a SHARK structure using assembly implementations of Algorithms 1 and 2 for their linear mapping. From there it can be seen without surprise that Algorithm 2 is more efficient if the matrix is well chosen. However, Algorithm 1 still performs reasonably well, without imposing any condition on the matrix.

4 Diffusion Matrices from Algebraic-Geometry Codes

In this section, we present so-called algebraic-geometry codes and show how they can give rise to diffusion matrices with interesting parameters. We also focus on implementation aspects, and investigate how to find matrices with efficient implementations with respect to the algorithms of Sect. 3, and in particular Algorithm 2.

4.1 A Short Introduction to Algebraic-Geometry Codes

We first briefly present the concept of algebraic-geometry codes (or AG codes for short), which are linear codes, and how to compute their generator matrices. Because the codes are linear, these encoders are matrices. We do not give a complete description of AG codes, and refer to *e.g.* [22] for a more thorough treatment. We present a class of AG codes as a generalization of Reed-Solomon (RS) codes.

We see AG codes as *evaluation codes*: to build the codeword for a message w , we consider w as a function, and the codeword as a vector of values of this function evaluated on some “elements”. In our case, the elements are points of the two-dimensional affine space $\mathbf{A}^2(\mathbf{F}_{2^m})$, and the functions are polynomials in two variables, that is elements of $\mathbf{F}_{2^m}[x, y]$. The core idea of AG codes is to consider points of a (smooth) projective curve of the projective space $\mathbf{P}^2(\mathbf{F}_{2^m})$ and functions from the curve’s function space. Points at infinity are never included in the (ordered) set of points. However, points of the curve at infinity *are* useful in defining the curve’s function (sub)-space, which is why we do consider the curve in the projective space instead of the affine one.

We first give the definition of the Riemann-Roch space in the special case where it is defined from a divisor made of a single point at infinity. We refer to *e.g.* [19] or [22] for a more complete and rigorous definition.

Definition 3 (Riemann-Roch space). *Let $\mathcal{X}$ be a smooth projective curve of $\mathbf{P}^2(\mathbf{F}_{2^m})$ defined by the homogeneous polynomial $p(x, y, z)$, and let $p'(x, y)$ be the dehomogenized of p . We define $\mathbf{F}_{2^m}[\mathcal{X}] = \mathbf{F}_{2^m}[x, y]/p'$ as the coordinate*

ring of $\mathcal{X}$, and its corresponding quotient field $\mathbf{F}_{2^m}(\mathcal{X})$ as the function field of $\mathcal{X}$. Assume Q is the only point of $\mathcal{X}$ at infinity, and let r be a positive integer. The Riemann-Roch space $\mathcal{L}(rQ)$ is the set of all functions of $\mathbf{F}_{2^m}(\mathcal{X})$ with poles only at Q of order less than r . This is a finite-dimensional $\mathbf{F}_{2^m}$ -vector-space. Furthermore, let $\text{o}_Q(x)$ and $\text{o}_Q(y)$ be the order of the poles of x and y in Q ⁴, then a basis of $\mathcal{L}(rQ)$ is formed by all the monomial functions $x^i y^j$ that are such that $i \cdot \text{o}_Q(x) + j \cdot \text{o}_Q(y) \leq r$.

This space is particularly important because of the following theorem, which links its dimension with the genus of $\mathcal{X}$ [19].

Theorem 1 (Riemann and Roch). *Let $\mathcal{L}(rQ)$ be a Riemann-Roch space defined on $\mathcal{X}$, and g be the genus of $\mathcal{X}$. We have $\dim(\mathcal{L}(rQ)) \geq r + 1 - g$, with equality when $r > 2g - 2$.*

We have also mentioned earlier that a basis for a space $\mathcal{L}(rQ)$ can be computed as soon as the order of the poles of x and y in Q are known, and the dimension of the space can obviously be computed from the basis. In practice, computing $\text{o}_Q(x)$ and $\text{o}_Q(y)$ can be done from a local parameterization of x and y in Q . Both this parameterization and the values $\text{o}_Q(x)$ and $\text{o}_Q(y)$ can easily be obtained from a computational algebra software such as Magma. Again, we refer to [22] for more details.

We are now ready to define a simple class of AG codes.

Definition 4 (Algebraic-Geometry Codes). *Let $\mathcal{X}$ be a smooth projective curve of $\mathbf{P}^2(\mathbf{F}_{2^m})$ with a unique point Q at infinity, and call $\#\mathcal{X}$ its number of affine points (that is not counting Q). Assume that $\#\mathcal{X} \geq n$ and let r be s.t. $\dim(\mathcal{L}(rQ)) = k$, and call $(f_0, \dots, f_{k-1})$ one basis of this space. We define the codeword of the $[n, k, d]_{\mathbf{F}_{2^m}}$ algebraic-geometry code $\mathcal{C}_{AG}$ associated with the message $\mathbf{m}$ as the vector $\sum_{i=0 \dots k-1} (\mathbf{m}_i \cdot f_i(p_j), j = 0 \dots n-1)$, where $P = \langle p_0, \dots, p_{n-1} \rangle$ is an ordered set of points of $\mathcal{X}/\{Q\}$. The code $\mathcal{C}_{AG}$ is the set of all such codewords.*

These codes have the following properties: for fixed parameters n and k and a curve $\mathcal{X}$, there are $\binom{\#\mathcal{X}}{n} \cdot n!$ equivalent codes, which corresponds to the number of possible ordered sets P ; it is also obvious that the maximal length of a code over $\mathcal{X}$ is $\#\mathcal{X}$. We also have the following proposition [22]:

Proposition 1. *Let $\mathcal{C}_{AG}$ be a code of length n and dimension k , and let r be an integer such that $\dim(\mathcal{L}(rQ)) = k$. Then the minimum distance of $\mathcal{C}_{AG}$ is at least $n - r$. If $\mathcal{X}$ is of genus g and $r > 2g - 2$, this is equal to $n - ((k - 1) + g) = n - k - g + 1$. Therefore, the “gap” between this code and an MDS code of the same length is g . The same holds for the dual code.*

⁴ We slightly abuse the notations here and actually mean x/z and y/z . But we prefer manipulating their dehomogenized equivalents x and y . It is obvious that x/z and y/z indeed have poles in Q , which is at infinity and hence has a zero z coordinate.

The minimum distance of AG codes thus depends on the genus of the curves used to define them. Because the maximal number of points on a curve increases with its genus, there is a tradeoff between the length of a code and its minimum distance.

Construction of a Generator Matrix of an AG Code. Once the parameters of a code have been fixed, including the ordered set P , one just has to specify a basis of $\mathcal{L}(rQ)$, and to form the encoding matrix $M \in \mathcal{M}_{k,n}(\mathbf{F}_{2^m})$ obtained by evaluating this basis on P . A useful basis is one such that the encoding matrix is in systematic form, but it does not necessarily exist for any P . Note however that in the case of MDS codes (such as RS codes) this basis always exists whatever the parameters and the choice of P : this is because in this case every minor of M is of full rank [15]. When such a basis exists, it is easy to find as one just has to start from an arbitrary basis and to compute the reduced row echelon form of the matrix thus obtained.

Example 1: An AG Code from an Elliptic Curve. We give parameters for a code built from the curve defined on $\mathbf{P}^2(\mathbf{F}_{2^4})$ by the homogeneous polynomial $x^2z + xz^2 = y^3 + yz^2$, or equivalently defined on $\mathbf{A}^2(\mathbf{F}_{2^4})$ by $x^2 + x = y^3 + y$. It is of genus 1, and hence it is an elliptic curve. It has 25 points, including one point at infinity, the point $Q = [1 : 0 : 0]$; the order of the poles of x and y in Q are respectively 3 and 2. From this, a basis for the space $\mathcal{L}(12Q)$ can easily be obtained. This space has dimension $12 + 1 - g = 12$, and can be used to define a $[24, 12, 12]_{\mathbf{F}_{2^4}}$ code by evaluation over the affine points of the curve. This allows to define a matrix of dimension 12 over $\mathbf{F}_{2^4}$, which diffuses over $12 \times 4 = 48$ bits and has a differential and linear branch number of 12.

Example 2: An AG Code from an Hyperelliptic Curve. We increase the length of the code by using a curve with a larger genus. We give parameters for a rather well-known code, built from the curve defined on $\mathbf{P}^2(\mathbf{F}_{2^4})$ by the homogeneous polynomial $x^5 = y^2z + yz^4$. This curve has 33 points, including one point at infinity, the point $Q = [0 : 1 : 0]$; the order of the poles of x and y in Q are respectively 2 and 5. From this, a basis for the space $\mathcal{L}(17Q)$ can easily be defined. This space has dimension $17 + 1 - g = 16$, and can be used to define a $[32, 16, 15]_{\mathbf{F}_{2^4}}$ code by evaluation over the affine points of the curve. This code has convenient parameters for defining diffusion matrices: from a generator matrix in systematic form $(I_{16} \ A)$, we can extract the matrix A , which naturally diffuses over 64 bits and has a differential and linear branch number of 15. Furthermore, the code is *self-dual*, which means that A is orthogonal: $A \cdot A^t = I_{16}$. The inverse of A is therefore easy to compute.

We give the right matrix of two matrices of this code in systematic form in the full version [3], the latter further including an example of a basis of $\mathcal{L}(17Q)$ and the order of the points used to construct the matrix.

The problem for the rest of the section is now to find good point orders P for the hyperelliptic code of Example 2 such that efficient encoders can be

constructed thanks to Algorithm 2 of Sect. 3.3. For convenience, we name $\mathcal{C}_{\mathcal{HE}}$ any of the codes equivalent to the one of Example 2.

4.2 Compact Encoders Using Code Automorphisms

We consider matrices in systematic form $(I_{16} \ A)$. For dense matrices, Algorithm 2 tends to be most efficient when all the rows of a matrix can be deduced by permutation of one of them, or more generally of a small subset of them. Our objective is thus to find matrices of this form.

The main tool we use to achieve this goal are *automorphisms* of $\mathcal{C}_{\mathcal{HE}}$. Let us first give a definition. (In the following, $\mathfrak{S}_n$ denotes the group of permutations of n elements.)

Definition 5. (Automorphisms of a code). *The automorphism group $\text{Aut}(\mathcal{C})$ of a code $\mathcal{C}$ of length n is a subgroup of $\mathfrak{S}_n$ such that $\pi \in \text{Aut}(\mathcal{C}) \Rightarrow (c \in \mathcal{C} \Rightarrow \pi(c) \in \mathcal{C})$.*

Because we consider here the code $\mathcal{C}_{\mathcal{HE}}$ which is an evaluation code, we can equivalently define its automorphisms as being permutations of the points on which the evaluation is performed. If π is an automorphism of $\mathcal{C}_{\mathcal{HE}}$, if $\{O_0, \dots, O_l\}$ are its orbits, and if the code is defined with a point order P such that for each orbit all of its points are neighbours in the order P , then the effect of π on a codeword of $\mathcal{C}_{\mathcal{HE}}$ is to cyclically permute its coordinates along each orbit.

To see that this is useful, assume that there is an automorphism π with two orbits O_0 and O_1 of size $n/2$ each. Then, if $M = (I_{n/2} \ A)$ is obtained with point order $P = \langle O_0, O_1 \rangle$, each row of M can be obtained by the repeated action of π on, say, $\mathbf{m}^0$, and it follows that A is circulant (and therefore has a low cost w.r.t. Algorithm 2). More generally, if an automorphism can be found such that it has orbits of size summing up to $n/2$, the corresponding matrix M can be deduced from a small set of rows. We give two toy examples with Reed-Solomon codes, which can easily be verified.

$\pi : \mathbf{F}_{2^4} \rightarrow \mathbf{F}_{2^4}, x \mapsto 8x$. This automorphism has $O_0 = \langle 1, 8, 12, 10, 15 \rangle$ and $O_1 = \langle 2, 3, 11, 7, 13 \rangle$ for orbits, among others. The systematic matrix for the $[10, 5, 6]_{\mathbf{F}_{2^4}}$ code obtained with the points in that order is then such that A is circulant and obtained from the cyclic permutation of the row $(12, 10, 2, 6, 3)$.

$\pi : \mathbf{F}_{2^4} \rightarrow \mathbf{F}_{2^4}, x \mapsto 7x$. This automorphism has $O_0 = \langle 1, 7, 6 \rangle$, $O_1 = \langle 2, 14, 12 \rangle$, $O_2 = \langle 4, 15, 11 \rangle$, and $O_3 = \langle 8, 13, 5 \rangle$ for orbits, among others. The systematic matrix for the $[12, 6, 7]_{\mathbf{F}_{2^4}}$ code obtained with the points in that order is then of the form $\begin{pmatrix} I_3 & 0_3 & A & B \\ 0_3 & I_3 & C & D \end{pmatrix}$ with A, B, C and D circulant matrices. It can thus be obtained by cyclic permutation of only two rows.

Application to $\mathcal{C}_{\mathcal{HE}}$. Automorphisms of $\mathcal{C}_{\mathcal{HE}}$ are quite harder to find than ones of RS codes. They can however be found within automorphisms of the curve $\mathcal{X}$ on which it is based [19]. This is quite intuitive, as these will precisely permute points on the curve, which are the points on which the code is defined.

We mostly need to be careful to ensure that the point at infinity is fixed by these automorphisms. We considered the degree-one automorphisms of $\mathcal{X}$ described by Duursma [8]. They have two generators: $\pi_0 : \mathbf{F}_{2^4}^2 \rightarrow \mathbf{F}_{2^4}^2, (x, y) \mapsto (\zeta x, y)$ with $\zeta^5 = 1$, and $\pi_{1(a,b)} : \mathbf{F}_{2^4}^2 \rightarrow \mathbf{F}_{2^4}^2, (x, y) \mapsto (x + a, y + a^8 x^2 + a^4 x + b^4)$, with (a, b) an affine point of $\mathcal{X}$. These generators span a group of order 160. When considering their orbit decomposition, the break-up of the size of the orbits can only be of one of five types, given in Table 1.

Table 1. Possible combination of orbit sizes of automorphisms of $\mathcal{C}_{\mathcal{H}\mathcal{E}}$ spanned by π_0 and π_1 . A number n in col. c means that an automorphism of this type has n orbits of size c .

Orbit size	1	2	4	5	10
Type 1	32	0	0	0	0
Type 2	0	16	0	0	0
Type 3	0	0	8	0	0
Type 4	2	0	0	6	0
Type 5	0	1	0	0	3

From these automorphisms, it is possible to define a partitions of P in two sets of size 16, which are union of orbits. We may therefore hope to obtain systematic matrices of the type we are looking for. Unfortunately, after an extensive search⁵, it appears that ordering P in this fashion *never* results in obtaining a systematic matrix. We recall that indeed, because AG codes are not MDS, it is not always the case that computing the reduced row echelon form of an arbitrary encoding matrix yields a systematic matrix.

Extending the Automorphisms with the Frobenius Mapping. We extend the previous automorphisms with the Frobenius mapping $\theta : \mathbf{F}_{2^4}^2 \rightarrow \mathbf{F}_{2^4}^2, (x, y) \mapsto (x^2, y^2)$; this adds another 160 automorphisms for $\mathcal{X}$. However, these will not anymore be automorphisms for the *code* $\mathcal{C}_{\mathcal{H}\mathcal{E}}$ in general, and we will therefore obtain matrices of a form slightly different from what we first hoped to achieve.

The global strategy is still the same, however, and consists in ordering the points along orbits of the curve automorphisms. By using the Frobenius, new combinations of orbits are possible, notably 4 of size 8. We study below the result of ordering P along the orbits of one such automorphism. We take the example of $\sigma = \theta \circ \sigma_2 \circ \sigma_1$, with $\sigma_1 : (x, y) \mapsto (x + 1, y + x^2 + x + 7)$, $\sigma_2 : (x, y) \mapsto (12x, y)$, and θ the Frobenius mapping. The key observation is that in this case, only σ^0 and σ^4 are automorphisms of $\mathcal{C}_{\mathcal{H}\mathcal{E}}$. Note that not all orbits orderings of σ for P yield a systematic matrix. However, unlike as above, we were able to find some

⁵ Both on $\mathcal{C}_{\mathcal{H}\mathcal{E}}$ and on the smaller elliptic code of Example 1. However, we are not as yet able to explain this fact.

orders that do. In these cases, the right matrix “ A ” of the full generator matrix ($I_{16} A$) is of the form:

$$(\mathbf{a}^0, \dots, \mathbf{a}^3, \sigma^4(\mathbf{a}^0), \dots, \sigma^4(\mathbf{a}^3), \mathbf{a}^8, \dots, \mathbf{a}^{11}, \sigma^4(\mathbf{a}^8), \dots, \sigma^4(\mathbf{a}^{11}))^t,$$

with $\mathbf{a}^0, \dots, \mathbf{a}^3, \mathbf{a}^8, \dots, \mathbf{a}^{11}$ row vectors of dimension 16. For instance, the first and fifth row of one such matrix are:

$$\begin{aligned} \mathbf{a}^0 &= (5, 2, 1, 3, 8, 5, 1, 5, 12, 10, 14, 6, 7, 11, 4, 11) \\ \mathbf{a}^4 &= \sigma^4(\mathbf{a}^0) = (8, 5, 1, 5, 5, 2, 1, 3, 7, 11, 4, 11, 12, 10, 14, 6). \end{aligned}$$

We give the full matrix in the full version [3]. We have therefore partially reached our goal of being able to describe A from a permutation of a subset of its rows. However this subset is not small, as it is of size 8—half of the matrix dimension. Consequently, these matrices have a moderate cost according to the *cost2* function, when implemented with Algorithm 2, but it is not minimal. Interestingly, all the matrices of this form that we found have the same cost of 52.

4.3 Fast Random Encoders

We conclude this section by presenting the results of a very simple random search for efficient encoders of $\mathcal{C}_{\mathcal{HE}}$ with respect to Algorithm 2. Unlike the above study, this one does not exploit any kind of algebraic structure. Indeed, the search only consists in repeatedly generating a random permutation of the affine points of the curve, building a matrix for the code with the corresponding point order, tentatively putting it in systematic form ($I_{16} A$), and if successful evaluating the *cost2* function from Sect. 3.3 on A . We then collect matrices with a minimum cost.

Because there are $32! \approx 2^{117,7}$ possible point orders, we can only explore a very small part of the search space. However, matrices of low cost can be found even after a moderate amount of computation, and we found many matrices of cost 43, though none of a lower cost. We give the number of matrices of cost strictly less than 60 that we found during a search of 2^{38} encoders in the full version [3]. We also give in the full paper an example of a matrix of cost 43, which is only about a factor 1.5 away from the estimate of the minimum cost of a circulant matrix given in Sect. 3.3. We observe that the transpose of this matrix also has a cost of 43.

5 Applications and Performance

This last section presents the performance of straightforward assembly implementations of both of our algorithms when applied to a matrix of the code $\mathcal{C}_{\mathcal{HE}}$ from Sect. 4, of cost 43. For convenience, we denote $M_{\mathcal{H}16}$ this matrix. It is of dimension 16 over $\mathbf{F}_{2^4}$ and has a differential and linear branch number of 15.

We do this study in the context of block ciphers, by assuming that $M_{\mathcal{H}16}$ is used as the linear mapping of two ciphers with a **SHARK** structure: one with

4-bit S-Boxes and a 64-bit block, and one with 8-bit S-Boxes and a 128-bit block. What we wish to measure in both cases is the speed in cycles per byte of such hypothetical ciphers, so as to be able to gauge the efficiency of this linear mapping and of the resulting ciphers. In order to do this, we need to estimate how many rounds would be needed for the ciphers to be secure.

Basic Statistical Properties of a 64-bit Block Cipher with 4-bit S-Boxes and $M_{\mathcal{H}16}$ as a Linear Mapping. We use standard *wide-trail* considerations to study differential and linear properties of this cipher [7]. This is very easy to do thanks to the simple structure of the cipher. The branch number of $M_{\mathcal{H}16}$ is 15, which means that at least 15 S-boxes are active in any two rounds of a differential path or linear characteristic. The best 4-bit S-boxes have a maximum differential probability and a maximal linear bias of 2^{-2} (see *e.g.* [14, 17]). By using such S-boxes, one can upper-bound the probability of a single differential path or linear characteristic for $2n$ rounds by $2^{-2 \cdot 15n}$. This is smaller than 2^{-64} as soon as $n > 2$. Hence we conjecture that 6 to 8 rounds are enough to make a cipher resistant to standard statistical attacks.

If one were to propose a concrete cipher, a more detailed analysis would of course be needed, especially w.r.t. more dedicated structural attacks. However, it seems reasonable to consider at a first glance that 8 rounds would indeed be enough to bring adequate security. It is only 2 rounds less than AES-128, which uses a round with comparatively weaker diffusion. Also, 6 rounds might be enough. Consequently, we present software performance figures for 6 and 8 rounds of such an hypothetical 64-bit cipher with 4-bit S-Boxes in Table 2, on the left⁶. We include data both for a strict SSSE3 implementation and for one using AVX extensions, which can be seen to bring a considerable benefit. Note that the last round is complete and includes the linear mapping, unlike *e.g.* AES. Also, note that the parallel application of the S-Boxes can be implemented very efficiently with a single `pshufb`, and thus has virtually no impact on the speed.

Basic Statistical Properties of a 128-Bit Block Cipher with 8-Bit S-Boxes and $M_{\mathcal{H}16}$ as a Linear Mapping. Although the code $\mathcal{C}_{\mathcal{H}\mathcal{E}}$ from which the matrix $M_{\mathcal{H}16}$ is built was initially defined with $\mathbf{F}_{2^4}$ as an alphabet, this latter can be replaced by an algebraic extension of $\mathbf{F}_{2^4}$ such as $\mathbf{F}_{2^8}$, to yield a code $\mathcal{C}_{\mathcal{H}\mathcal{E}'}$ with the same parameters, namely a $[32, 16, 15]_{\mathbf{F}_{2^8}}$ code. Indeed, by using a suitable representation such as $\mathbf{F}_{2^4} \cong \mathbf{F}_2[\alpha]/(\alpha^4 + \alpha^3 + \alpha^2 + \alpha + 1)$; $\mathbf{F}_{2^8} \cong \mathbf{F}_{2^4}[t]/(t^2 + t + \alpha)$ ⁷, an element of $\mathbf{F}_{2^8}$ is represented as a degree-one polynomial $at + b$ over $\mathbf{F}_{2^4}$. It follows that the minimum weight of a codeword $w = (a_i t + b_i)$, $i \in \{0 \dots 31\}$ of $\mathcal{C}_{\mathcal{H}\mathcal{E}'}$ is at least equal to the minimum weight of words (a_i) , $i = 0 \dots 31$ and (b_i) , $i \in \{0 \dots 31\}$. If those are taken among codewords of $\mathcal{C}_{\mathcal{H}\mathcal{E}}$, their minimum weight is 15, and thus so is the one of w . It is also possible to efficiently compute the multiplication by $M_{\mathcal{H}16}$ over $\mathbf{F}_{2^8}$ from two computations over $\mathbf{F}_{2^4}$: because the coefficients of $M_{\mathcal{H}16}$ are in $\mathbf{F}_{2^4}$, we have

⁶ The figures are given for computation of a single block; better performance can be achieved for longer messages.

⁷ This representation is used by Hamburg in [10].

$M_{\mathcal{H}16} \cdot (a_i t + b_i) = (M_{\mathcal{H}16} \cdot (a_i) \cdot t) + (M_{\mathcal{H}16} \cdot (b_i))$. As a result, applying $M_{\mathcal{H}16}$ to $\mathbf{F}_{2^8}$ has only twice the cost of applying it to $\mathbf{F}_{2^4}$, while effectively doubling the size of the block.

From wide-trail considerations, the resistance of such a cipher to statistical attacks is comparatively even better than when using 4-bit S-Boxes, when an appropriate 8-bit S-Box is used. For instance, the AES S-box has a maximal differential probability and linear bias of 2^{-6} [7]. This implies that the probability of a single differential path or linear characteristic for $2n$ rounds is upper-bounded by $2^{-6 \cdot 15n}$, which is already much smaller than 2^{-128} as soon as $n > 1$. Again, 8 rounds of such a cipher should bring adequate security, and 6 rounds might be enough. We provide performance figures for both an SSSE3 and an AVX implementation in Table 2, on the right (see footnote 6). However, in the case of 8-bit S-Boxes, the S-Box application is rather more complex and expensive a step than with 4-bit S-Boxes. In these test programs, we decided to use the efficient vector implementation of the AES S-Box from Hamburg [10].

Table 2. Performance of software implementations of the hypothetical 64 and 128-bit cipher, in cycles per byte (cpb). Figures in parentheses are for an AVX implementation (when applicable).

Processor type	# rounds	64-bit Block		128-bit Block	
		cpb	cpb	cpb	cpb
		(Algorithm 1)	(Algorithm 2)	(Algorithm 1)	(Algorithm 2)
Intel Xeon E5-2650 @ 2.00 GHz	6	50 (45.5)	33 (24.2)	58 (52.3)	32.7 (26.5)
	8	66.5 (60.2)	44.5 (31.9)	76.8 (69.6)	43.8 (35.7)
Intel Xeon E5-2609 @ 2.40 GHz	6	72.3 (63.7)	45.3 (33.2)	79.8 (75.6)	47.1 (36.8)
	8	95.3 (84.7)	63.3 (45.6)	106.6 (97.1)	62.1 (50.3)
Intel Xeon E5649 @ 2.53 GHz	6	84.7	46	84.5	47
	8	111.3	59.8	111	61.9

Discussion. The performance figures given in Table 2 are average for a block cipher. For instance, it compares favourably with the optimized vector implementations of 64-bit ciphers LED and Piccolo in sequential mode from [4], which run at speeds between 70 and 90 cpb., depending on the CPU. It is however slower than Hamburg’s vector implementation of AES, with reported speeds of 6 to 22 cpb. (9 to 25 for the inverse cipher) [10, 20].

6 Conclusion

We revisited the SHARK structure by replacing the MDS matrix of its linear diffusion layer by a matrix built from an algebraic-geometry code. Although this code is not MDS, it has a very high minimum distance, while being defined over $\mathbf{F}_{2^4}$ instead of $\mathbf{F}_{2^8}$. This allows to reduce the size of the coefficients of the matrix from 8 to 4 bits, and has important consequences for efficient implementations of this linear mapping. We studied algorithms suitable for a vector implementation

of the multiplication by this matrix, and how to find matrices that are most efficiently implemented with those algorithms. Finally, we gave performance figures for assembly implementations of hypothetical SHARK-like ciphers using this matrix as a linear layer.

This work provided generalizations of SHARK that are not vulnerable to timing attacks as is the original cipher, and also a generalization to 128-bit blocks. It also showed that even if not the fastest, such potential design could be implemented efficiently in software.

As a future work, it would be interesting to investigate how to use the full automorphism group of the code to design matrices with a lower cost.

Acknowledgments. Pierre Karpman is partially supported by the Direction générale de l'armement and by the Singapore National Research Foundation Fellowship 2012 (NRF-NRFF2012-06).

References

1. Augot, D., Finiasz, M.: Direct construction of recursive MDS diffusion layers using shortened BCH codes. In: Cid, C., Rechberger, C. (eds.) FSE 2014. LNCS. Springer (2014). <https://eprint.iacr.org/2014/566>
2. Augot, D., Finiasz, M.: Exhaustive search for small dimension recursive MDS diffusion layers for block ciphers and hash functions. In: ISIT, IEEE, pp. 1551–1555 (2013)
3. Augot, D., Fouque, P.A., Karpman, P.: Diffusion Matrices from Algebraic-Geometry Codes with Efficient SIMD Implementation. IACR Cryptology ePrint Archive 2014, 551 (2014). <https://eprint.iacr.org/2014/551>
4. Benadjila, R., Guo, J., Lomné, V., Peyrin, T.: Implementing lightweight block ciphers on x86 architectures. In: Lange, T., Lauter, K., Lisoněk, P. (eds.) SAC 2013. LNCS, vol. 8282, pp. 324–352. Springer, Heidelberg (2014)
5. Bernstein, D.J., Schwabe, P.: NEON crypto. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 320–339. Springer, Heidelberg (2012)
6. Chen, J.P., Lu, C.C.: A serial-in-serial-out hardware architecture for systematic encoding of Hermitian codes via Gröbner Bases. IEEE Trans. Commun. **52**(8), 1322–1332 (2004)
7. Daemen, J., Rijmen, V.: The Design of Rijndael: AES – The Advanced Encryption Standard. Information Security and Cryptography. Springer, New York (2002)
8. Duursma, I.: Weight distributions of geometric Goppa codes. Trans. Am. Math. Soc. **351**(9), 3609–3639 (1999)
9. Guo, J., Peyrin, T., Poschmann, A.: The PHOTON family of lightweight hash functions. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 222–239. Springer, Heidelberg (2011)
10. Hamburg, M.: Accelerating AES with vector permute instructions. In: Clavier, C., Gaj, K. (eds.) CHES 2009. LNCS, vol. 5747, pp. 18–32. Springer, Heidelberg (2009)
11. Heegard, C., Little, J., Saints, K.: Systematic encoding via Gröbner bases for a class of algebraic-geometric goppa codes. IEEE Trans. Inf. Theor. **41**(6), 1752–1761 (1995)
12. Intel Corporation: Intel® 64 and IA-32 Architectures Software Developer's Manual, March 2012

13. Knudsen, L.R., Wu, H. (eds.): SAC 2012. LNCS, vol. 7707. Springer, Heidelberg (2013)
14. Leander, G., Poschmann, A.: On the classification of 4 bit S-Boxes. In: Carlet, C., Sunar, B. (eds.) WAIFI 2007. LNCS, vol. 4547, pp. 159–176. Springer, Heidelberg (2007)
15. MacWilliams, F.J., Sloane, N.J.A.: The Theory of Error-Correcting Codes. North-Holland Mathematical Library. North-Holland, Amsterdam (1978)
16. Rijmen, V., Daemen, J., Preneel, B., Bosselaers, A., Win, E.D.: The cipher SHARK. In: Gollmann, D. (ed.) FSE 1996. LNCS, vol. 1039, pp. 99–111. Springer, Heidelberg (1996)
17. Saarinen, M.-J.O.: Cryptographic analysis of all 4×4 -Bit S-Boxes. In: Miri, A., Vaudenay, S. (eds.) SAC 2011. LNCS, vol. 7118, pp. 118–133. Springer, Heidelberg (2012)
18. Sajadieh, M., Dakhilalian, M., Mala, H., Sepehrdad, P.: Recursive diffusion layers for block ciphers and hash functions. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 385–401. Springer, Heidelberg (2012)
19. Stichtenoth, H.: Algebraic Function Fields and Codes. Graduate Texts in Mathematics, vol. 254, 2nd edn. Springer, New York (2009)
20. Suzuki, T., Minematsu, K., Morioka, S., Kobayashi, E.: TWINE: A Lightweight Block Cipher for Multiple Platforms. [13], pp. 339–354
21. Tromer, E., Osvik, D.A., Shamir, A.: Efficient cache attacks on AES, and counter-measures. *J. Cryptol.* **23**(1), 37–71 (2010)
22. Van Lint, J.H.: Introduction to Coding Theory. Graduate Texts in Mathematics, vol. 86, 3rd edn. Springer, Berlin (1999)
23. Wu, S., Wang, M., Wu, W.: Recursive Diffusion Layers for (Lightweight) Block Ciphers and Hash Functions. [13], pp. 355–371