



Research Article

Gustavo Banegas*, Paulo S. L. M. Barreto, Edoardo Persichetti, and Paolo Santini

Designing Efficient Dyadic Operations for Cryptographic Applications

<https://doi.org/10.1515/jmc-2015-0054>

Received Feb 05, 2020; accepted Feb 11, 2020

Abstract: Cryptographic primitives from coding theory are some of the most promising candidates for NIST's Post-Quantum Cryptography Standardization process. In this paper, we introduce a variety of techniques to improve operations on dyadic matrices, a particular type of symmetric matrices that appear in the automorphism group of certain linear codes. Besides the independent interest, these techniques find an immediate application in practice. In fact, one of the candidates for the Key Exchange functionality, called DAGS, makes use of quasi-dyadic matrices to provide compact keys for the scheme.

Keywords: post-quantum cryptography, code-based cryptography, dyadic matrices.

2010 Mathematics Subject Classification: 11T71, 94A60

1 Introduction

Post-Quantum Cryptography is the area of research that investigates cryptographic primitives that are deemed secure against attackers equipped with quantum technology. These include schemes based on a variety of mathematical problems, such as finding short vectors in a lattice, or decoding random linear codes. The latter is known as Code-based Cryptography and it relies more or less directly on the Syndrome Decoding Problem [4], which shows no vulnerabilities to quantum attacks. The first code-based scheme was introduced by McEliece in 1978 [10] and has resisted cryptanalysis, in its original form, for nearly 40 years.

McEliece's cryptosystem has often been ignored in favor of schemes based on number theory problems (such as RSA or El Gamal), mainly due to the size of its public key, which was deemed too large for practical use (especially at the time). However, Shor's algorithm [14] shows that, once quantum computers of an appropriate size are available, the cryptosystems currently in use will become obsolete. It is therefore important to offer a credible alternative to current cryptography, and, with this in mind, NIST has recently launched a call for papers to standardize the public-key primitives of the future [1].

Among the code-based candidates for NIST's call, DAGS [3] is a Key Encapsulation Mechanism (KEM) that uses Quasi-Dyadic (QD) matrices to considerably reduce the size of the public key, following a McEliece-like approach. The proposal builds on a line of work initiated by Misoczki and Barreto [11] and subsequently developed by Persichetti in [6, 12].

*Corresponding Author: **Gustavo Banegas:** Technische Universiteit Eindhoven, Netherlands; Email: souzag@chalmers.se

Paulo S. L. M. Barreto: Institute of Technology, University of Washington at Tacoma, United States of America

Edoardo Persichetti: Department of Mathematical Sciences, Florida Atlantic University, United States of America; Email: epersichetti@fau.edu

Paolo Santini: Università Politecnica delle Marche, Ancona, Italy; Email: p.santini@staff.univpm.it

Our Contribution.

We analyze two separate aspects of dyadic operations. First, we present three different algorithms for that are aimed specifically at computing multiplication of dyadic matrices. These are, respectively, a “standard” approach that makes use of dyadic signatures, a specialized Karatsuba-like algorithm, and a procedure based on the Fast Walsh-Hadamard Transform (FWHT) [7], also called dyadic convolution. We analyze the performance of all three methods and report our timings.

As a second contribution, we describe a procedure that applies the LUP decomposition [5] to the dyadic case. The method effectively factors every quasi-dyadic matrix into a product of two triangular matrices and a permutation matrix. This leads to the possibility of a very efficient algorithm for computing the inverse of a matrix, which is particularly useful in code-based cryptography, for instance for computing the systematic form of a parity-check (or generator) matrix. According to our measurements, this improved inversion procedure is extremely fast, and provides a very large speedup during DAGS Key Generation.

Organization of the Paper. This paper is organized as follows. We start with some preliminary definitions in Section 2. We then present our main contributions: the various multiplication techniques are described in Section 3 and the improved inversion algorithm is presented in Section 4. We conclude by showing the results obtained when applying our techniques to DAGS; this is done in Section 5.

2 Preliminaries

We introduce dyadic matrices and describe some of their general properties.

Definition 2.1. Given a ring \mathcal{R} and a vector $\mathbf{h} = (h_0, h_1, \dots, h_{n-1}) \in \mathcal{R}^n$, with $n = 2^r$ for some $r \in \mathbb{N}$, the dyadic matrix $\Delta(\mathbf{h}) \in \mathcal{R}^{n \times n}$ is the symmetric matrix with components $\Delta_{ij} = h_{i \oplus j}$ where \oplus stands for bitwise exclusive-or. Such a matrix is said to have order r . The sequence \mathbf{h} is called signature of the matrix $\Delta(\mathbf{h})$, and corresponds to its first row. The set of dyadic $n \times n$ matrices over \mathcal{R} is denoted $\Delta(\mathcal{R}^n)$.

One can alternatively characterize a dyadic matrix recursively: any 1×1 matrix is dyadic of order 0, and any dyadic matrix \mathbf{M} of order $r > 0$ has the form

$$\mathbf{M} = \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{pmatrix} \quad (1)$$

where A and B are two dyadic matrices of order $r - 1$. In other words, $\Delta(\mathcal{R}^n) = \Delta(\Delta(\mathcal{R}^{n/2}))$.

Definition 2.2. A dyadic permutation is a dyadic matrix $\Pi^i \in \Delta(\{0, 1\}^n)$ characterized by the signature $\boldsymbol{\pi}^i = (\delta_{ij} \mid j = 0, \dots, n - 1)$, where δ_{ij} is the Kronecker delta (hence $\boldsymbol{\pi}^i$ corresponds to the i -th row or column of the identity matrix).

A dyadic permutation is clearly an involution, i.e. $(\Pi^i)^2 = \mathbf{I}$. The i -th row, or equivalently the i -th column, of the dyadic matrix defined by a signature \mathbf{h} can be written as $\Delta(\mathbf{h})_i = \mathbf{h}\Pi^i$.

A dyadic matrix can be efficiently represented by its signature; in particular, all the operations between dyadic matrices can be referred only to the corresponding signatures. Indeed, for any two length- n vectors $\mathbf{a}, \mathbf{b} \in \mathcal{R}$, we have:

$$\Delta(\mathbf{a}) + \Delta(\mathbf{b}) = \Delta(\mathbf{a} + \mathbf{b}) \quad (2)$$

which means that, given two dyadic matrices \mathbf{A} and \mathbf{B} , with respective signatures \mathbf{a} and \mathbf{b} , their sum is the dyadic matrix described by the signature $\mathbf{a} + \mathbf{b}$.

In an analogous way, the multiplication between dyadic matrices can be done by considering only the corresponding signatures; we will discuss efficient ways for computing multiplications in Section 3.

Algorithm 1 Standard multiplication of dyadic matrices

INPUT: $r \in \mathbb{N}$, $n = 2^r$ and $\mathbf{a}, \mathbf{b} \in \mathbb{F}^n$.
OUTPUT: $\mathbf{c} \in \mathbb{F}^n$ such that $\Delta(\mathbf{c}) = \Delta(\mathbf{a})\Delta(\mathbf{b})$.

- 1: $\mathbf{c} \leftarrow$ vector of length n , initialized with null elements.
- 2: $c_0 \leftarrow a_0 \cdot b_0$
- 3: **for** $i \leftarrow 1$ **to** $n - 1$ **do**
- 4: $c_0 \leftarrow c_0 + a_i b_i$
- 5: $i^{(b)} \leftarrow$ binary representation of i , using n bits.
- 6: **for** $\{j = 0, 1, \dots, n - 1\}$ **do**
- 7: $j^{(b)} \leftarrow$ binary representation of j , using n bits.
- 8: $\pi^{(b)} \leftarrow i^{(b)} \oplus j^{(b)}$
- 9: $\pi \leftarrow$ conversion of $\pi^{(b)}$ into an integer.
- 10: $c_i \leftarrow c_i + a_i b_\pi$
- 11: **end for**
- 12: **end for**
- 13: **return** \mathbf{c}

Moreover, it is easy to see that the inverse of a dyadic matrix is also a dyadic matrix; this can be easily computed using Sylvester-Hadamard matrices (see Section 3.2). We will expand on this in Section 4.

Finally, we introduce a relaxed notion of dyadicity, which will be useful throughout the paper.

Definition 2.3. A quasi-dyadic matrix is a (possibly non-dyadic) block matrix whose elements are dyadic submatrices, i.e. an element of $\Delta(\mathcal{R}^n)^{d_1 \times d_2}$.

3 Multiplication of Dyadic Matrices

In this section we consider different methods for computing the multiplication between two dyadic matrices. In fact, we have just mentioned how some matrix operations, like the sum or the inversion, can be efficiently performed in the dyadic case just by considering the signatures. Multiplication can be strongly improved with similar methods, which exploit the particular structure of such matrices. In particular, we analyze three different algorithms and provide estimations for their complexities; we then compare the performance of the various algorithms.

For ease of notation, we will refer to the two $n \times n$ matrices that we want to multiply simply as \mathbf{A} and \mathbf{B} , with $\mathbf{a} = [a_0, a_1, \dots, a_{n-1}]$ and $\mathbf{b} = [b_0, b_1, \dots, b_{n-1}]$ being the respective signatures. Maintaining the same notation, the product matrix $\mathbf{C} = \mathbf{AB}$, which is also dyadic, will have signature $\mathbf{c} = [c_0, c_1, \dots, c_{n-1}]$.

In particular, we focus on the special case of quasi-dyadic matrices with elements belonging to a field \mathbb{F} of characteristic 2.

3.1 Standard Multiplication

The first algorithm we analyze is described in Algorithm 1; we refer to it as the *standard multiplication*. The element of \mathbf{C} in position (i, j) is obtained as the multiplication between the i -th row of \mathbf{A} and the j -th column of \mathbf{B} . Since dyadic matrices are symmetric, this is equivalent to the inner product between the i -th row of \mathbf{A} and the j -th one of \mathbf{B} . The signature \mathbf{c} (i.e., the first row of \mathbf{C}) is obtained by inner products involving only \mathbf{a} (i.e., the first row of \mathbf{A}). Thus, we can just construct the rows of \mathbf{B} , by permutations of the elements in \mathbf{b} , and then compute the inner products.

The complexity of the algorithm is due to two different types of operations:

1. In order to construct the rows of \mathbf{B} , we need the indexes of the corresponding permutations. Each index is computed as the modulo 2 sum of two binary vectors of length r , so can be obtained with a complexity of r binary operations. Thus, considering that we need to repeat this operation for $2^r - 1$ rows (for the first one, no permutation is needed), the complexity of this procedure can be estimated as $r \cdot 2^r \cdot (2^r - 1)$.
2. Each element of \mathbf{c} is obtained as the inner product between two vectors of 2^r elements, assuming values in \mathbb{F} . This operation requires 2^r multiplication and $2^r - 1$ sums in \mathbb{F} . If we denote as C_{mult} and C_{sum} the costs of, respectively, a multiplication and a sum in \mathbb{F} , the total number of binary operations needed to compute 2^r inner products can be estimated as $2^{2r} \cdot C_{\text{mult}} + (2^{2r} - 2^r) \cdot C_{\text{sum}}$.

The complexity of a standard multiplication between two dyadic signatures can be estimated as:

$$C_{\text{std}} = r \cdot (2^{2r} - 2^r) + 2^{2r} \cdot C_{\text{mult}} + (2^{2r} - 2^r) \cdot C_{\text{sum}} \quad (3)$$

3.2 Dyadic Convolution

Definition 3.1. The dyadic convolution of two vectors $\mathbf{a}, \mathbf{b} \in \mathcal{R}$, denoted by $\mathbf{a} \triangle \mathbf{b}$, is the unique vector of \mathcal{R} such that $\Delta(\mathbf{a} \triangle \mathbf{b}) = \Delta(\mathbf{a})\Delta(\mathbf{b})$.

Of particular interest to us is the case where \mathcal{R}^n is actually a field \mathbb{F} . Dyadic matrices over \mathbb{F} form a commutative subring $\Delta(\mathbb{F}^n) \subset \mathbb{F}^{n \times n}$, and this property gives rise to efficient arithmetic algorithms to compute the dyadic convolution. In particular, we here consider the fast Walsh-Hadamard transform (FWHT), which is well known [7] but seldom found in a cryptographic context. We describe it here for ease of reference. We firstly recall the FWHT for the case of a field \mathbb{F} such that $\text{char}(\mathbb{F}) \neq 2$, and then describe how this technique can be generalized to consider also the case of $\text{char}(\mathbb{F}) = 2$ (which, again, is the one we are interested in).

Definition 3.2. Let \mathbb{F} be a field with $\text{char}(\mathbb{F}) \neq 2$. The Sylvester-Hadamard matrix $\mathbf{H}_r \in \mathbb{F}^n$ is recursively defined as

$$\begin{aligned} \mathbf{H}_0 &= \begin{bmatrix} 1 \end{bmatrix}, \\ \mathbf{H}_r &= \begin{bmatrix} \mathbf{H}_{r-1} & \mathbf{H}_{r-1} \\ \mathbf{H}_{r-1} & -\mathbf{H}_{r-1} \end{bmatrix}, r > 0. \end{aligned}$$

One can show by straightforward induction that $\mathbf{H}_r^2 = 2^r \mathbf{I}_r$ and hence $\mathbf{H}_r^{-1} = 2^{-r} \mathbf{H}_r$, which can also be expressed recursively as

$$\begin{aligned} \mathbf{H}_0^{-1} &= \begin{bmatrix} 1 \end{bmatrix}, \\ \mathbf{H}_r^{-1} &= \frac{1}{2} \begin{bmatrix} \mathbf{H}_{r-1}^{-1} & \mathbf{H}_{r-1}^{-1} \\ \mathbf{H}_{r-1}^{-1} & -\mathbf{H}_{r-1}^{-1} \end{bmatrix}, r > 0. \end{aligned}$$

Lemma 3.1. Let \mathbb{F} be a field with $\text{char}(\mathbb{F}) \neq 2$. If $\mathbf{M} \in \mathbb{F}^{n \times n}$ is dyadic, then $\mathbf{H}_r^{-1} \mathbf{M} \mathbf{H}_r$ is diagonal.

Proof. The lemma clearly holds for $r = 0$. Now let $r > 0$, and write

$$\mathbf{M} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix}$$

where \mathbf{A} and \mathbf{B} are dyadic. It follows that

$$\mathbf{H}_r^{-1} \mathbf{M} \mathbf{H}_r = \frac{1}{2} \begin{bmatrix} \mathbf{H}_{r-1}^{-1} & \mathbf{H}_{r-1}^{-1} \\ \mathbf{H}_{r-1}^{-1} & -\mathbf{H}_{r-1}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix} \begin{bmatrix} \mathbf{H}_{r-1} & \mathbf{H}_{r-1} \\ \mathbf{H}_{r-1} & -\mathbf{H}_{r-1} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{r-1}^{-1} \mathbf{M}_+ \mathbf{H}_{r-1} & \mathbf{O} \\ \mathbf{O} & \mathbf{H}_{r-1}^{-1} \mathbf{M}_- \mathbf{H}_{r-1} \end{bmatrix},$$

and since both $\mathbf{M}_+ = \mathbf{A} + \mathbf{B}$ and $\mathbf{M}_- = \mathbf{A} - \mathbf{B}$ are dyadic, $\mathbf{H}_{r-1}^{-1} \mathbf{M}_+ \mathbf{H}_{r-1}$ and $\mathbf{H}_{r-1}^{-1} \mathbf{M}_- \mathbf{H}_{r-1}$ are diagonal by induction, as is thus also $\mathbf{H}_r^{-1} \mathbf{M} \mathbf{H}_r$. \square

Lemma 3.1 establishes that Sylvester-Hadamard matrices diagonalize all dyadic matrices. In particular, the factors in a product of dyadic matrices are thus simultaneously diagonalized, suggesting an efficient way to carry out the matrix multiplication, namely, computing $\mathbf{H}_r^{-1}(\mathbf{M}\mathbf{N})\mathbf{H}_r = (\mathbf{H}_r^{-1}\mathbf{M}\mathbf{H}_r)(\mathbf{H}_r^{-1}\mathbf{r}\mathbf{H}_r)$ given the diagonal forms $\mathbf{H}_r^{-1}\mathbf{M}\mathbf{H}_r$ and $\mathbf{H}_r^{-1}\mathbf{r}\mathbf{H}_r$ of two dyadic matrices \mathbf{M} and \mathbf{N} requires only n multiplications of the diagonal elements.

In fact, it is not necessary to compute $\mathbf{H}_r^{-1}\mathbf{M}\mathbf{H}_r$ in full to obtain the diagonal form of \mathbf{M} , as indicated by the following result:

Lemma 3.2. *Let \mathbb{F} be a field with $\text{char}(\mathbb{F}) \neq 2$. The diagonal form of a dyadic matrix $\mathbf{M} \in \mathbb{F}^{n \times n}$ is the first line of $\mathbf{M}\mathbf{H}_r$. In other words, $\mathbf{H}_r^{-1}\Delta(\mathbf{h})\mathbf{H}_r = \text{diag}(\mathbf{h}\mathbf{H}_r)$.*

Proof. The lemma clearly holds for $r = 0$. Now let $r > 0$, and with the notation of Lemma 3.1, the diagonal of $\mathbf{H}_r^{-1}\mathbf{M}\mathbf{H}_r$ is the concatenation of the diagonals of $\mathbf{H}_{r-1}^{-1}\mathbf{M}_+\mathbf{H}_{r-1}$ and $\mathbf{H}_{r-1}^{-1}\mathbf{M}_-\mathbf{H}_{r-1}$. Similarly, since

$$\mathbf{M}\mathbf{H}_r = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{bmatrix} \begin{bmatrix} \mathbf{H}_{r-1} & \mathbf{H}_{r-1} \\ \mathbf{H}_{r-1} & -\mathbf{H}_{r-1} \end{bmatrix} = \begin{bmatrix} \mathbf{M}_+\mathbf{H}_{r-1} & \mathbf{M}_-\mathbf{H}_{r-1} \\ \mathbf{M}_+\mathbf{H}_{r-1} & -\mathbf{M}_-\mathbf{H}_{r-1} \end{bmatrix},$$

the first line of $\mathbf{M}\mathbf{H}_r$ is the concatenation of the first lines of $\mathbf{M}_+\mathbf{H}_{r-1}$ and $\mathbf{M}_-\mathbf{H}_{r-1}$, which by induction are the diagonals of $\mathbf{H}_{r-1}^{-1}\mathbf{M}_+\mathbf{H}_{r-1}$ and $\mathbf{H}_{r-1}^{-1}\mathbf{M}_-\mathbf{H}_{r-1}$ respectively, yielding the claimed property. \square

Corollary 3.2.1. *Computing \mathbf{c} such that $\Delta(\mathbf{a})\Delta(\mathbf{b}) = \Delta(\mathbf{c})$ involves only three multiplications of vectors by Sylvester-Hadamard matrices.*

Proof. By Lemma 3.2, $\text{diag}(\mathbf{a}\mathbf{H}_r) \text{diag}(\mathbf{b}\mathbf{H}_r) = (\mathbf{H}_r^{-1}\Delta(\mathbf{a})\mathbf{H}_r)(\mathbf{H}_r^{-1}\Delta(\mathbf{b})\mathbf{H}_r) = \mathbf{H}_r^{-1}\Delta(\mathbf{a})\Delta(\mathbf{b})\mathbf{H}_r = \mathbf{H}_r^{-1}\Delta(\mathbf{c})\mathbf{H}_r = \text{diag}(\mathbf{c}\mathbf{H}_r)$. Now simply retrieve \mathbf{c} from $\mathbf{z} = \mathbf{c}\mathbf{H}_r$ as $\mathbf{c} = \mathbf{z}\mathbf{H}_r^{-1} = 2^{-r}\mathbf{z}\mathbf{H}_r$. \square

The structure of Sylvester-Hadamard matrices leads to an efficient algorithm to compute $\mathbf{a}\mathbf{H}_r$ for $\mathbf{a} \in \mathbb{F}^n$, which is known as the fast Walsh-Hadamard transform. Let $[\mathbf{a}_0, \mathbf{a}_1]$ be the two halves of \mathbf{a} . Thus

$$\mathbf{a}\mathbf{H}_r = [\mathbf{a}_0, \mathbf{a}_1] \begin{bmatrix} \mathbf{H}_{r-1} & \mathbf{H}_{r-1} \\ \mathbf{H}_{r-1} & -\mathbf{H}_{r-1} \end{bmatrix} = [(\mathbf{a}_0 + \mathbf{a}_1)\mathbf{H}_{r-1}, (\mathbf{a}_0 - \mathbf{a}_1)\mathbf{H}_{r-1}].$$

This recursive algorithm, which can be easily written in purely sequential fashion (Algorithm 2), has complexity $\Theta(n \log n)$, specifically, rn additions or subtractions in \mathbb{F} . It is therefore somewhat more efficient than the fast Fourier transform, which involves multiplications by n -th roots of unity, when they are available at all (otherwise working in extension fields is unavoidable, and more expensive).

The product of two dyadic matrices $\Delta(\mathbf{a})$ and $\Delta(\mathbf{b})$, or equivalently the dyadic convolution $\mathbf{a} \triangle \mathbf{b}$, can thus be efficiently computed as described in Algorithm 3. The total cost is $3rn$ additions or subtractions and $2n$ multiplications (half of these by the constant $2^{-r} = 1/n$) in \mathbb{F} , with an overall complexity $\Theta(n \log n)$. Notice that this is also the complexity of computing $\det \Delta(\mathbf{a})$.

The fast Walsh-Hadamard transform itself is not immediately possible on fields of characteristic 2, since it depends on Sylvester-Hadamard matrices which must contain a primitive square root of unity. Yet the FWHT algorithm can be lifted to characteristic 0, namely, from $\mathbb{F}_2 = \mathbb{Z}/2\mathbb{Z}$ to \mathbb{Z} , or more generally from $\mathbb{F}_{2^N} = (\mathbb{Z}/2\mathbb{Z})[x]/P(x)$ (for some irreducible $P(x)$ of degree N) to $\mathbb{Z}[x]$. Algorithm 3 can then be applied, and its output mapped back to the relevant binary field by modular reduction. This incurs a space expansion by a logarithmic factor, though. Each bit from \mathbb{F}_2 is mapped to intermediate values that can occupy as much as $3r + 1$ bits; correspondingly, each element from \mathbb{F}_{2^N} is mapped to intermediate values that can occupy as much as $(3r + 1)N$ bits. Thus the component-wise multiplication in Algorithm 3 becomes more complicated to implement for large N . However, the method remains very efficient for the binary case as long as each expanded integer component fits a computer word. For a typical word size of 32 bits and each binary component being expanded by a factor of $3r + 1$, this means that blocks as large as 1024×1024 can be tackled efficiently. On more restricted platforms where the maximum available word size is 16 bits, dyadic blocks of size 32×32 can still be handled with relative ease.

Algorithm 2 The fast Walsh-Hadamard transform (FWHT)INPUT: $r \in \mathbb{N}$, $n = 2^r$ and $\mathbf{a} \in \mathbb{F}^n$ with $\text{char}(\mathbb{F}) \neq 2$.OUTPUT: \mathbf{aH}_r .

```

1:  $v \leftarrow 1$ 
2: for  $j \leftarrow 1$  to  $n$  do
3:    $w \leftarrow v$ 
4:    $v \leftarrow 2v$ 
5:   for  $i \leftarrow 0$  to  $n - 1$  by  $v$  do
6:     for  $l \leftarrow 0$  to  $w - 1$  do
7:        $s \leftarrow a_{i+l}$ 
8:        $q \leftarrow a_{i+l+w}$ 
9:        $a_{i+l} \leftarrow s + q$ 
10:       $a_{i+l+w} \leftarrow s - q$ 
11:     end for
12:   end for
13: end for
14: return  $\mathbf{a}$ 

```

Algorithm 3 Dyadic convolution via the FWHTINPUT: $r \in \mathbb{N}$, $n = 2^r$ and $\mathbf{a}, \mathbf{b} \in \mathbb{F}^n$ with $\text{char}(\mathbb{F}) \neq 2$.OUTPUT: $\mathbf{a} \triangle \mathbf{b} \in \mathbb{F}^n$ such that $\Delta(\mathbf{a})\Delta(\mathbf{b}) = \Delta(\mathbf{a} \triangle \mathbf{b})$.

```

1:  $\mathbf{c} \leftarrow$  vector of length  $n$ , initialized with null elements.
2:  $\tilde{\mathbf{c}} \leftarrow$  vector of length  $n$ , initialized with null elements.
3: Compute  $\tilde{\mathbf{a}} \leftarrow \mathbf{aH}_r$  via Algorithm 2. ▷ expansion 1  $\rightarrow r + 1$ 
4: Compute  $\tilde{\mathbf{b}} \leftarrow \mathbf{bH}_r$  via Algorithm 2. ▷ expansion 1  $\rightarrow r + 1$ 
5: for  $j \leftarrow 0$  to  $n - 1$  do
6:    $\tilde{c}_j \leftarrow \tilde{a}_j \tilde{b}_j$  ▷ expansion  $r + 1 \rightarrow 2r + 1$ 
7: end for
8: Compute  $\mathbf{c} \leftarrow \tilde{\mathbf{cH}}_r$  via Algorithm 2. ▷ expansion  $2r + 1 \rightarrow 3r + 1$ 
9:  $\mathbf{c} \leftarrow 2^{-r}\mathbf{c}$ 
10: return  $\mathbf{c}$ 

```

3.3 Karatsuba Multiplication

In this section we propose a method which is inspired by Karatsuba's algorithm for the multiplication of two integers [9]. Let us denote by \mathbf{a}_0 and \mathbf{a}_1 , respectively, the first and second halves of \mathbf{a} , i.e.:

$$\begin{aligned}\mathbf{a}_0 &= [a_0, a_1, \dots, a_{\frac{n}{2}-1}] \\ \mathbf{a}_1 &= [a_{\frac{n}{2}}, a_{\frac{n}{2}+1}, \dots, a_{n-1}].\end{aligned}\quad (4)$$

The same notation is used for \mathbf{b}_0 and \mathbf{b}_1 and \mathbf{c}_0 and \mathbf{c}_1 , corresponding to the halves of \mathbf{B} and \mathbf{C} . Some straightforward computations show that the following relations hold:

$$\begin{aligned}\mathbf{c}_0 &= \mathbf{a}_0 \mathbf{b}_0 + \mathbf{a}_1 \mathbf{b}_1 \\ \mathbf{c}_1 &= (\mathbf{a}_0 + \mathbf{a}_1)(\mathbf{b}_0 + \mathbf{b}_1) + \mathbf{c}_0\end{aligned}\quad (5)$$

The iterative application of equation (5) allows to compute multiplications between dyadic matrices of any size. Let us denote as $C_{\text{mul}}^{(2^z)}$ and $C_{\text{sum}}^{(2^z)}$ the complexities of a multiplication and a sum between two signatures of length 2^z . For the sum of two dyadic signatures of size 2^z we have:

$$C_{\text{sum}}^{(2^z)} = 2^z \cdot C_{\text{sum}}, \quad (6)$$

where C_{sum} again denotes the complexity of a sum in the finite field.

The complexity of this algorithm can thus be estimated as:

$$\begin{aligned}C_{\text{Kar}} &= 3 \cdot C_{\text{mul}}^{(2^{r-1})} + 4 \cdot C_{\text{sum}}^{(2^{r-1})} = \\ &= 3 \cdot C_{\text{mul}}^{(2^{r-1})} + 4 \cdot 2^{r-1} \cdot C_{\text{sum}} = \\ &= 3 \cdot [3 \cdot C_{\text{mul}}^{(2^{r-2})} + 4 \cdot C_{\text{sum}}^{(2^{r-2})}] + 4 \cdot 2^{r-1} \cdot C_{\text{sum}} = \\ &= 3 [3 \cdot C_{\text{mul}}^{(2^{r-2})} + 4 \cdot 2^{r-2} \cdot C_{\text{sum}}] + 4 \cdot 2^{r-1} \cdot C_{\text{sum}} = \\ &= 3^2 \cdot C_{\text{mul}}^{(2^{r-2})} + 4 \cdot [3 \cdot 2^{r-2} + 2^{r-1}] C_{\text{sum}} = \\ &= 3^3 \cdot C_{\text{mul}}^{(2^{r-3})} + 4 \cdot [3^2 \cdot 2^{r-3} + 3 \cdot 2^{r-2} + 2^{r-1}] C_{\text{sum}} = \\ &= \dots = \\ &= 3^r \cdot C_{\text{mul}} + 4 \cdot \left[\sum_{j=1}^r 3^{j-1} 2^{r-j} \right] \cdot C_{\text{sum}} = \\ &= 3^r \cdot C_{\text{mul}} + \frac{4}{3} \cdot 2^r \cdot \left[\sum_{j=1}^r \left(\frac{3}{2} \right)^j \right] \cdot C_{\text{sum}}\end{aligned}\quad (7)$$

Taking into account the well known sum of a geometric series, we have:

$$\begin{aligned}\sum_{j=1}^r \left(\frac{3}{2} \right)^j &= -1 + \sum_{j=0}^r \left(\frac{3}{2} \right)^j = \\ &= -1 + \frac{1 - \left(\frac{3}{2} \right)^{r+1}}{1 - \frac{3}{2}} = \frac{3^{r+1}}{2^r} - 3.\end{aligned}$$

Considering this result, equation (7) leads to:

$$C_{\text{Kar}} = 3^r \cdot C_{\text{mul}} + 4 \cdot [3^r - 2^r] \cdot C_{\text{sum}} \quad (8)$$

4 Efficient Inversion of Dyadic and Quasi-Dyadic Matrices

In this section we propose an efficient algorithm for computing the inverse of quasi-dyadic matrices. The algorithm in principle is targeted to matrices that are not fully dyadic (even though, obviously, they have to be square). This is because, while it is of course possible to apply our procedure to fully dyadic matrices, these can in general be inverted much more easily, as we will see next.

To begin, remember that by definition of a quasi-dyadic matrix (Definition 2.3) we mean an element of $\Delta(\mathcal{R}^n)^{d_1 \times d_2}$.

4.1 Dyadic Matrices

The inverse of a dyadic matrix (i.e. $d_1 = d_2 = 1$) can be efficiently computed, using only the signature, as described by the following Lemma.

Lemma 4.1. *Let $n = 2^r$ for $r \in \mathbb{N}$ and let $\Delta(\mathbf{a}) \in \mathcal{R}^{n \times n}$ be a dyadic matrix with signature \mathbf{a} . Then the inverse $\Delta(\mathbf{a})^{-1}$ is the dyadic matrix $\Delta(\mathbf{b})$, for $\mathbf{b} = \frac{1}{2^r} \tilde{\mathbf{b}} \mathbf{H}_r$, where $\tilde{\mathbf{b}}$ is the vector such that $\text{diag}(\tilde{\mathbf{b}}) = [\text{diag}(\mathbf{a} \mathbf{H}_r)]^{-1}$.*

Proof We have $\Delta(\mathbf{b})\Delta(\mathbf{a}) = \mathbf{I}_n = \Delta([1, 0, \dots, 0])$. The diagonal form of \mathbf{I}_n corresponds to the first row of the product $\mathbf{I}_n \mathbf{H}_r$, and so it is equal to the first row of \mathbf{H}_r , that is the length- n vector made of all ones. According to Corollary 3.2.1, we can write:

$$\text{diag}(\mathbf{b}) \text{diag}(\mathbf{a}) = \text{diag}([1, 1, \dots, 1]).$$

We then define $\mathbf{a} \mathbf{H}_r = [\lambda_0, \lambda_1, \dots, \lambda_{n-1}]$, and obtain:

$$\text{diag}(\mathbf{b}) = \text{diag}([1, 1, \dots, 1]) \text{diag}^{-1}(\mathbf{a}) = \text{diag}([\lambda_0^{-1}, \lambda_1^{-1}, \dots, \lambda_{n-1}^{-1}]).$$

Because of Lemma 3.2, we finally have:

$$\mathbf{b} = \text{diag}^{-1}(\mathbf{a}) \mathbf{H}_r^{-1} = \frac{1}{2^r} \text{diag}^{-1}(\mathbf{a}) \mathbf{H}_r. \quad \square$$

As we mentioned before, the above Lemma yields a very simple way for computing the inverse of a dyadic matrix: given a signature \mathbf{a} , we just need to compute its diagonalized form as $\mathbf{a} \mathbf{H}_r$, compute the reciprocals of its elements and put it in a vector $\tilde{\mathbf{b}}$. Finally, the inverse of $\Delta(\mathbf{a})$ can be obtained as $\frac{1}{2^r} \tilde{\mathbf{b}} \mathbf{H}_r$. This property also leads to a very simple way to check the singularity of $\Delta(\mathbf{a})$: if its diagonalized form contains some null elements, then it is singular.

We now focus on the case of dyadic matrices over a field \mathbb{F} with characteristic 2. One can show by induction that in such a case a dyadic matrix $\Delta(\mathbf{a})$ of dimension n satisfies $\Delta(\mathbf{a})^2 = (\sum_i a_i)^2 \mathbf{I}$, and hence its inverse, when it exists, is $\Delta(\mathbf{a})^{-1} = (\sum_i a_i)^{-2} \Delta(\mathbf{a})$, which can be computed in $O(n)$ steps since it is entirely determined by its first row. It is equally clear that $\det \Delta(\mathbf{a}) = (\sum_i a_i)^n$, which can be computed with the same complexity (notice that raising to the power of $n = 2^r$ only involves r squarings). Basically, verifying whether a dyadic matrix has full rank or not can be easily done by checking whether the sum of the elements of the signature equals 0.

4.2 Quasi-Dyadic Matrices

Consider a quasi-dyadic matrix \mathbf{M} . Since the matrix has to be square, we have $d_1 = d_2 = d$, and the matrix has dimension $dn \times dn$. Such a matrix can be compactly represented just by the signatures of the dyadic blocks. To simplify notation, we can denote the signature of the dyadic-block in position (i, j) as $\hat{\mathbf{m}}_{i,j}$, and store all

such vectors in a matrix $\hat{\mathbf{M}} \in \mathcal{R}^{d \times dn}$:

$$\hat{\mathbf{M}} = \begin{pmatrix} \hat{\mathbf{m}}_{0,0} & \hat{\mathbf{m}}_{0,1} & \cdots & \hat{\mathbf{m}}_{0,d-1} \\ \hat{\mathbf{m}}_{1,0} & \hat{\mathbf{m}}_{1,1} & \cdots & \hat{\mathbf{m}}_{1,d-1} \\ \vdots & \vdots & \ddots & \vdots \\ \hat{\mathbf{m}}_{d-1,0} & \hat{\mathbf{m}}_{d-1,1} & \cdots & \hat{\mathbf{m}}_{d-1,d-1} \end{pmatrix}. \quad (9)$$

We focus again on the special case of quasi-dyadic matrices over a field \mathbb{F} with characteristic 2.

The LUP decomposition is a method which factorizes a matrix \mathbf{M} as \mathbf{LUP} , where \mathbf{L} and \mathbf{U} are lower triangular and upper triangular matrices, respectively, and \mathbf{P} is a permutation.

Exploiting this factorization, the inverse of \mathbf{M} can thus be expressed as:

$$\mathbf{M}^{-1} = \mathbf{P}^{-1} \mathbf{U}^{-1} \mathbf{L}^{-1}. \quad (10)$$

The advantage of this method is that the inverses appearing in (10) can be easily computed, because of their particular structures. In fact, the inverse of an upper (lower) triangular matrix is obtained via a simple backward (forward) substitution procedure, while the inverse of \mathbf{P} is its transpose.

In some cases, applying a block-wise LUP decomposition might lead to some complexity reduction; for instance, see [13] for the inversion of a sparse matrix. Here, we consider the case of a quasi-dyadic matrix; the corresponding procedure is shown in Algorithm 4.

Algorithm 4 LUP Decomposition of a Quasi-Dyadic Matrix

INPUT: $d, r \in \mathbb{N}$, $n = 2^r$ and $\hat{\mathbf{M}} \in \mathbb{F}^{d \times dn}$ with $\text{char}(\mathbb{F}) = 2$.

OUTPUT: $\hat{\mathbf{M}} \in \mathbb{F}^{d \times dn}$, $\hat{\mathbf{P}} \in \mathbb{N}^d$.

```

1:  $\hat{\mathbf{P}} \leftarrow [0, 1, \dots, d-1]$ 
2:  $u \leftarrow 0$ 
3: for  $j \leftarrow 0$  to  $d-1$  do
4:   Update  $u$ ,  $\hat{\mathbf{M}}$  and  $\hat{\mathbf{P}}$  via Algorithm 5. ▷ Pivoting of the signatures in the  $j$ -th column
5:   if  $u = 0$  then
6:     return  $u$  ▷  $\hat{\mathbf{M}}$  is singular
7:   end if
8:   for  $i \leftarrow j+1$  to  $d$  do
9:      $\hat{\mathbf{m}}_{i,j} \leftarrow \hat{\mathbf{m}}_{i,j} \hat{\mathbf{m}}_{j,j}^{-1}$ 
10:  end for
11:  for  $i \leftarrow j+1$  to  $d-1$  do
12:    for  $l \leftarrow j+1$  to  $d-1$  do
13:       $\hat{\mathbf{m}}_{i,l} \leftarrow \hat{\mathbf{m}}_{i,l} + \hat{\mathbf{m}}_{i,j} \hat{\mathbf{m}}_{j,l}$ 
14:    end for
15:  end for
16: end for
17: return  $\hat{\mathbf{M}}, \hat{\mathbf{P}}$ 

```

Our proposed procedure consists in using a block decomposition, which works directly on the signatures, in order to exploit the simple and efficient algebra of dyadic matrices. The operations in Algorithm 4 only refer to the signatures in $\hat{\mathbf{M}}$: for instance, the expression $\hat{\mathbf{m}}_{i,j} \hat{\mathbf{m}}_{j,l}$ means the product between the dyadics having as signatures $\hat{\mathbf{m}}_{i,j}$ and $\hat{\mathbf{m}}_{j,l}$. This choice may result in some abuse of notation, but is useful to emphasize the fact that, as we have explained in the previous sections, operations with dyadics can be efficiently computed just by taking into account their signatures. It can be easily shown that, for a quasi-dyadic matrix, its factors \mathbf{L} , \mathbf{U} and \mathbf{P} are in quasi-dyadic form as well: as we have done for the matrix \mathbf{M} , we refer to their compact representations as $\hat{\mathbf{L}}$, $\hat{\mathbf{U}}$ and $\hat{\mathbf{P}}$, respectively.

The algorithm takes as input a matrix $\hat{\mathbf{M}}$, as in (9), and computes its LUP factorization; outputs of the algorithm are the modified matrix $\hat{\mathbf{M}}$, having as elements the ones of its factors $\hat{\mathbf{L}}$ and $\hat{\mathbf{U}}$, and the permutation $\hat{\mathbf{P}}$. As in (9), we denote as $\hat{\mathbf{m}}_{i,j}$ the signature in position (i, j) in the output matrix $\hat{\mathbf{M}}$. The matrices $\hat{\mathbf{L}}$ and $\hat{\mathbf{U}}$ can then be expressed as:

$$\hat{\mathbf{L}} = \begin{pmatrix} \hat{\mathbf{1}} & \hat{\mathbf{0}} & \hat{\mathbf{0}} & \cdots & \hat{\mathbf{0}} \\ \hat{\mathbf{m}}_{1,0} & \hat{\mathbf{1}} & \hat{\mathbf{0}} & \cdots & \hat{\mathbf{0}} \\ \hat{\mathbf{m}}_{2,0} & \hat{\mathbf{m}}_{2,1} & \hat{\mathbf{1}} & \cdots & \hat{\mathbf{0}} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \hat{\mathbf{m}}_{d-1,0} & \hat{\mathbf{m}}_{d-1,1} & \hat{\mathbf{m}}_{d-1,2} & \cdots & \hat{\mathbf{1}} \end{pmatrix}, \quad \hat{\mathbf{U}} = \begin{pmatrix} \hat{\mathbf{m}}_{0,0} & \hat{\mathbf{m}}_{0,1} & \hat{\mathbf{m}}_{0,2} & \cdots & \hat{\mathbf{m}}_{0,d-1} \\ \hat{\mathbf{0}} & \hat{\mathbf{m}}_{1,1} & \hat{\mathbf{m}}_{1,2} & \cdots & \hat{\mathbf{m}}_{1,d-1} \\ \hat{\mathbf{0}} & \hat{\mathbf{0}} & \hat{\mathbf{m}}_{2,2} & \cdots & \hat{\mathbf{m}}_{2,d-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \hat{\mathbf{0}} & \hat{\mathbf{0}} & \hat{\mathbf{0}} & \cdots & \hat{\mathbf{m}}_{d-1,d-1} \end{pmatrix} \quad (11)$$

where $\hat{\mathbf{1}}$ and $\hat{\mathbf{0}}$ denote, respectively, the signature of the identity matrix and the one of the null matrix (i.e. the length- k vectors $[1, 0, \dots, 0]$ and $[0, 0, \dots, 0]$).

The matrix $\hat{\mathbf{P}}$ is represented through a length- d vector $[p_0, p_1, \dots, p_{d-1}]$, containing a permutation of the integers $[0, 1, \dots, d-1]$; the rows of $\hat{\mathbf{M}}$ get permuted according to the elements of $\hat{\mathbf{P}}$. In particular, the elements of $\hat{\mathbf{P}}$ are obtained through a block pivoting procedure, which is described in Algorithm 5.

Algorithm 5 Block pivoting

INPUT: $d, j, r \in \mathbb{N}$, $n = 2^r$, $\hat{\mathbf{P}} \in \mathbb{N}^d$ and $\hat{\mathbf{M}} \in \mathbb{F}^{d \times dn}$ with $\text{char}(\mathbb{F}) = 2$, .

OUTPUT: $u \in \mathbb{N}$.

```

1:  $u \leftarrow 0$ 
2:  $i \leftarrow j$ 
3: while  $i \leq d - 1$  do
4:    $w \leftarrow \text{sum}(\hat{\mathbf{m}}_{i,j})$  ▷ Sum of the elements in  $\hat{\mathbf{m}}_{i,j}$ 
5:   if  $w = 0$  then
6:      $z \leftarrow p_j$ 
7:      $p_j \leftarrow p_i$ 
8:      $p_i \leftarrow z$ 
9:     for  $l \leftarrow 0$  to  $d - 1$  do
10:       $z \leftarrow \hat{\mathbf{m}}_{j,l}$ 
11:       $\hat{\mathbf{m}}_{j,l} \leftarrow \hat{\mathbf{m}}_{i,l}$ 
12:       $\hat{\mathbf{m}}_{i,l} \leftarrow z$ 
13:       $i \leftarrow i + 1$ 
14:     end for
15:   else
16:      $i \leftarrow d$ 
17:      $u \leftarrow 1$ 
18:   end if
19: end while
20: return  $u$ 

```

This function takes as input $\hat{\mathbf{M}}$, $\hat{\mathbf{P}}$ and an integer j , and searches for a pivot (i.e., a non singular signature) in the j -th column of $\hat{\mathbf{M}}$, starting from $\hat{\mathbf{m}}_{j,j}$, and places it in position (j, j) . As the procedure goes on, every time a singular signature is tested, the rows of $\hat{\mathbf{M}}$ get permuted; the elements of $\hat{\mathbf{P}}$ are accordingly modified. If the j -th column contains all singular blocks, this means that the matrix $\hat{\mathbf{M}}$ is singular; in such a case, this event is notified by setting $u = 0$.

We point out that, for the matrices we are considering, we expect Algorithm 4 to be particularly efficient. First of all, as we have already said, this is due to the possibility of efficiently performing operations involving dyadic matrices; in addition, the dyadic structure should also speed-up the pivoting procedure. In fact, we can consider a signature in $\hat{\mathbf{M}}$ as a collection of k random elements picked from $GF(2^N)$: thus, their sum can be assumed to be a random variable with uniform distribution among the elements of the field $GF(2^N)$. So, the probability of it being equal to 0, which corresponds to the probability of the corresponding signature to be singular, equals 2^{-m} . This probability gets lower as m increases: this fact means that the expected number of operations performed by Algorithm 5 should be particularly low. Basically, most of the times the function will just compute the sum of the elements in $\hat{\mathbf{m}}_{j,j}$ and verify whether it is null or not.

Once the factorization of $\hat{\mathbf{M}}$ has been obtained, we just need to perform the computation of \mathbf{M}^{-1} through (10). Since the inverse of a triangular matrix maintains the original triangular structure, the computation of the inverses $\hat{\mathbf{L}}^{-1}$ and $\hat{\mathbf{U}}^{-1}$ can be efficiently performed. A possible way for computing these matrices is to store the elements of both matrices in just one output matrix $\hat{\mathbf{T}}$. We do this in Algorithm 6.

Algorithm 6 Computation of $\hat{\mathbf{T}}$

INPUT: $d, r \in \mathbb{N}, n = 2^r$ and $\hat{\mathbf{M}} \in \mathbb{F}^{d \times dn}$ with $\text{char}(\mathbb{F}) = 2$.

OUTPUT: $\hat{\mathbf{T}} \in \mathbb{F}^{d \times dn}$.

```

1:  $\hat{\mathbf{T}} \leftarrow \hat{\mathbf{I}}_d$ 
2: for  $j \leftarrow 0$  to  $d - 1$  do
3:   for  $i \leftarrow j + 1$  to  $d - 1$  do
4:     for  $l \leftarrow j$  to  $i - 1$  do
5:        $\hat{t}_{i,j} \leftarrow \hat{t}_{i,j} + \hat{m}_{i,k} \hat{t}_{k,j}$ 
6:     end for
7:   end for
8:   for  $i \leftarrow j$  to  $d - 1$  do
9:     for  $l \leftarrow j$  to  $i - 1$  do
10:       $\hat{t}_{j,i} \leftarrow \hat{t}_{j,i} + \hat{m}_{k,i} \hat{t}_{j,k}$ 
11:    end for
12:     $\hat{t}_{j,i} \leftarrow \hat{t}_{j,i} \hat{m}_{i,i}^{-1}$ 
13:  end for
14: end for
15: return  $\hat{\mathbf{T}}$ 

```

The matrix $\hat{\mathbf{I}}_d$ is the compact representation of a $dn \times dn$ identity matrix, and so is composed of signatures $\delta_{i,j} \hat{\mathbf{i}}$, where $\delta_{i,j}$ denotes the Kronecker delta.

If we denote as $\hat{t}_{i,j}$ the signature in position (i, j) , we have:

$$\hat{\mathbf{L}}^{-1} = \begin{pmatrix} \hat{\mathbf{i}} & \hat{\mathbf{o}} & \hat{\mathbf{o}} & \cdots & \hat{\mathbf{o}} \\ \hat{t}_{1,0} & \hat{\mathbf{i}} & \hat{\mathbf{o}} & \cdots & \hat{\mathbf{o}} \\ \hat{t}_{2,0} & \hat{t}_{2,1} & \hat{\mathbf{i}} & \cdots & \hat{\mathbf{o}} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \hat{t}_{d-1,0} & \hat{t}_{d-1,1} & \hat{t}_{d-1,2} & \cdots & \hat{\mathbf{i}} \end{pmatrix}, \quad \hat{\mathbf{U}}^{-1} = \begin{pmatrix} \hat{t}_{0,0} & \hat{t}_{0,1} & \hat{t}_{0,2} & \cdots & \hat{t}_{0,d-1} \\ \hat{\mathbf{o}} & \hat{t}_{1,1} & \hat{t}_{1,2} & \cdots & \hat{t}_{1,d-1} \\ \hat{\mathbf{o}} & \hat{\mathbf{o}} & \hat{t}_{2,2} & \cdots & \hat{t}_{2,d-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \hat{\mathbf{o}} & \hat{\mathbf{o}} & \hat{\mathbf{o}} & \cdots & \hat{t}_{d-1,d-1} \end{pmatrix}. \quad (12)$$

5 Performance Analysis: Application to DAGS

In this section we provide the results of the application of our techniques to DAGS. For completeness, we have included a specification of the three DAGS algorithms in Appendix A, but for our purpose, DAGS is essentially the McEliece cryptosystem, converted to a KEM via a standard transformation [8]. In particular, the Key Generation algorithm is the same as the QD-GS McEliece version described in [12]. In this algorithm, a key role is played by the systematization (i.e. reduced row echelon form) of a quasi-dyadic rectangular matrix, the result of which will in fact be the public key for the scheme. The cost of computing said systematic matrix dwarfs everything else in key generation: according to a static analysis, this takes over 98% of the total cost of key generation. Therefore, a fast procedure to compute the systematic form will have a substantial impact on the overall performance of the algorithm.

Implementation Details.

We developed a code in “C” to implement our procedures. In all cases, we use no optimizations apart from the optimization from the GCC compiler (“-O3”). The GCC version used was 7.3.1 20180406, the code was compiled for the processor Intel(R) Core(TM) i5-5300U CPU @ 2.30GHz with 16GB of memory and operating system Arch

linux version 2018.05.01 with Kernel 4.16.5. We ran 100 times each piece of code and computed the average of all measurements; to obtain the number of cycles, we used the file “cpucycles.h” from supercop¹.

Fast Multiplication.

To compare our methods, we fix a dyadic order r and measure the cost of a multiplication of two matrices of size $n = 2^r$. Relevant dyadic orders for DAGS are for instance $r = 4$ and $r = 5$. We do this over different fields to highlight the difference in performance when changing fields: we tested \mathbb{F}_{2^5} and \mathbb{F}_{2^6} which are the fields currently used by DAGS.

Table 1: Cost of Multiplication between Dyadic Matrices

		Standard	Karatsuba	Dyadic Convolution
\mathbb{F}_{2^5}	$r = 4$	4, 833	2, 194	3, 899
	$r = 5$	21, 285	5, 909	12, 045
\mathbb{F}_{2^6}	$r = 4$	5, 833	2, 194	4, 899
	$r = 5$	23, 231	6, 223	13, 568

Efficient Inversion.

We report here the results of the improved inversion procedure (Algorithms 4, 5 and 6). We compared our procedure with the equivalent portion of the DAGS implementation that we extrapolated from the publicly available source code [2]. In particular, we measured the piece of code that begins with the creation of the Cauchy matrix and ends with the generation of the systematic matrix. Table 2 shows the comparison, measured in cpu cycles.

Acknowledgement: Edoardo Persichetti and Paolo Santini were supported by NSF grant CNS-1906360.

¹ <https://bench.cr.yp.to/supercop.html>

Table 2: Comparison of Inversion Methods

	DAGS Implementation	LUP Inversion	LUP + Karatsuba
DAGS 1	1, 318, 973, 209	321, 771	108, 117
DAGS 3	2, 211, 076, 311	557, 822	198, 199
DAGS 5	17, 925, 330, 712	654, 713	431, 890

References

- [1] <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization>.
- [2] <https://git.dags-project.org/dags/dags>.
- [3] Gustavo Banegas, Paulo S. L. M. Barreto, Brice Odilon Boidje, Pierre-Louis Cayrel, Gilbert Ndollane Dione, Kris Gaj, Cheikh Thiécoumba Gueye, Richard Haeussler, Jean Belo Klamti, Ousmane Ndiaye, Duc Tri Nguyen, Edoardo Persichetti and Jefferson E. Ricardini, DAGS: Key Encapsulation using Dyadic GS Codes, *IACR Cryptology ePrint Archive* **2017** (2017), 1037.
- [4] E. Berlekamp, R. McEliece and H. van Tilborg, On the inherent intractability of certain coding problems (Corresp.), *Information Theory, IEEE Transactions on* **24** (1978), 384 – 386.
- [5] J. R. Bunch and J. E. Hopcroft, Triangular factorization and inversion by fast matrix multiplication, *Mathematics of Computation* **28** (1974), 231–236.
- [6] Pierre-Louis Cayrel, Gerhard Hoffmann and Edoardo Persichetti, Efficient Implementation of a CCA2-Secure Variant of McEliece Using Generalized Srivastava Codes, in: *Public Key Cryptography - PKC 2012 - 15th International Conference on Practice and Theory in Public Key Cryptography, Darmstadt, Germany, May 21-23, 2012. Proceedings* (Marc Fischlin, Johannes A. Buchmann and Mark Manulis, eds.), Lecture Notes in Computer Science 7293, pp. 138–155, Springer, 2012.
- [7] M. N. Gulamhusein, Simple matrix-theory proof of the discrete dyadic convolution theorem, *Electronics Letters* **9** (1973), 238–239.
- [8] Dennis Hofheinz, Kathrin Hövelmanns and Eike Kiltz, *A Modular Analysis of the Fujisaki-Okamoto Transformation*, Cryptology ePrint Archive, Report 2017/604, 2017, <http://eprint.iacr.org/2017/604>.
- [9] A. Karatsuba and Y. Ofman, *Multiplication of Multidigit Numbers by Automata*, 01 1963.
- [10] R. J. McEliece, A Public-Key Cryptosystem Based On Algebraic Coding Theory, *Deep Space Network Progress Report* **44** (1978), 114–116.
- [11] R. Misoczki and P. S. L. M. Barreto, Compact McEliece Keys from Goppa Codes, in: *Selected Areas in Cryptography*, pp. 376–392, 2009.
- [12] E. Persichetti, Compact McEliece keys based on quasi-dyadic Srivastava codes, *Journal of Mathematical Cryptology* **6** (2012), 149–169.
- [13] Lukas Polok and Pavel Smrz, Pivoting Strategy for Fast LU Decomposition of Sparse Block Matrices, in: *Proceedings of the 25th High Performance Computing Symposium, HPC '17*, pp. 14:1–14:12, Society for Computer Simulation International, San Diego, CA, USA, 2017.
- [14] P. W. Shor, Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer, *SIAM Journal on Computing* **26** (1997), 1484–1509.

A DAGS Algorithms

We briefly describe the three algorithms that define DAGS. Generalized Srivastava codes are defined by parameters s and t , where in our case $\log s$ is the dyadic order; the codes in use have length $n = n_0s$ and dimension $k = k_0s$ where n_0 and k_0 are the number of dyadic blocks. Other parameters are the cardinality of the base field q and the degree of the field extension m . In addition, we have $k = k' + k''$, where k' is arbitrary and set to be “small”.

The key generation process uses the following fundamental equation

$$\frac{1}{h_{i \oplus j}} = \frac{1}{h_i} + \frac{1}{h_j} + \frac{1}{h_0} \quad (13)$$

which guarantees we can build a dyadic matrix, with signature $\mathbf{h} = (h_0, h_1, \dots, h_{n-1})$, which is also a Cauchy matrix, i.e. a matrix $C(\mathbf{u}, \mathbf{v})$ with components $C_{ij} = \frac{1}{u_i - v_j}$. In [11] it is proved that we can use the fundamental

equation to choose a support and polynomial for a Goppa code such that this dyadic Cauchy matrix is a parity-check matrix for the code.

Key Generation

1. Generate dyadic signature \mathbf{h} according to the fundamental equation.
2. Build the vectors (\mathbf{u}, \mathbf{v}) that define the Cauchy matrix (again using the equation).
3. Form Cauchy matrix $\hat{H}_1 = C(\mathbf{u}, \mathbf{v})$.
4. Build \hat{H}_i , $i = 2, \dots, t$, by raising each element of \hat{H}_1 to the power of i .
5. Superimpose blocks \hat{H}_i in ascending order to form matrix \hat{H} .
6. Generate scaling vector \mathbf{z} by sampling elements z_i in \mathbb{F}_{q^m} with $z_{is+j} = z_{is}$ for $i = 0, \dots, n_0 - 1$, $j = 0, \dots, s - 1$.
7. Set $y_j = \frac{z_j}{\prod_{i=0}^{s-1} (u_i - v_j)^t}$ for $j = 0, \dots, n - 1$ and $\mathbf{y} = (y_0, \dots, y_{n-1})$.
8. Form $H = \hat{H} \cdot \text{Diag}(\mathbf{z})$.
9. Project H onto \mathbb{F}_q using the co-trace function: call this H_{base} .
10. Write H_{base} in systematic form $(A \mid \mathbf{I}_{n-k})$.
11. The public key is the generator matrix $G = (\mathbf{I}_k \mid A^T)$.
12. The private key is the pair (\mathbf{v}, \mathbf{y}) .

Note that all matrices involved in key generation are quasi-dyadic (with blocks of size $s \times s$), namely \hat{H} , H , H_{base} and its systematic form, and the final matrix G which is the public key. Step 10 is the systematization process which is impacted by our improved inversion algorithm.

The encapsulation and decapsulation algorithms make use of three hash functions $\mathcal{G} : \mathbb{F}_q^{k'} \rightarrow \mathbb{F}_q^k$, $\mathcal{H} : \mathbb{F}_q^{k'} \rightarrow \mathbb{F}_q^{k'}$ and $\mathcal{K} : \{0, 1\}^* \rightarrow \{0, 1\}^\ell$, where ℓ is the desired length of the key to be shared.

Encapsulation

1. Choose $\mathbf{m} \xleftarrow{\$} \mathbb{F}_q^{k'}$.
2. Compute $\mathbf{r} = \mathcal{G}(\mathbf{m})$ and $\mathbf{d} = \mathcal{H}(\mathbf{m})$.
3. Parse \mathbf{r} as $(\rho \parallel \sigma)$ then set $\mu = (\rho \parallel \mathbf{m})$.
4. Generate error vector \mathbf{e} of length n and weight w from σ .
5. Compute $\mathbf{c} = \mu G + \mathbf{e}$.
6. Compute $\mathbf{k} = \mathcal{K}(\mathbf{m})$.
7. Output ciphertext (\mathbf{c}, \mathbf{d}) ; the encapsulated key is \mathbf{k} .

The decapsulation algorithm is essentially a run of the decoding algorithm to decode the noisy codeword received as part of the ciphertext, plus a number of integrity checks.

Decapsulation

1. Recover parity-check matrix H' in alternant form from private key.
2. Use H' to decode \mathbf{c} and obtain codeword $\mu' G$ and error \mathbf{e}' .
3. Output \perp if decoding fails or $\text{wt}(\mathbf{e}') \neq w$
4. Recover μ' and parse it as $(\rho' \parallel \mathbf{m}')$.
5. Compute $\mathbf{r}' = \mathcal{G}(\mathbf{m}')$ and $\mathbf{d}' = \mathcal{H}(\mathbf{m}')$.
6. Parse \mathbf{r}' as $(\rho'' \parallel \sigma')$.
7. Generate error vector \mathbf{e}'' of length n and weight w from σ' .

8. If $\mathbf{e}' \neq \mathbf{e}'' \vee \rho' \neq \rho'' \vee \mathbf{d} \neq \mathbf{d}'$ output \perp .
9. Else compute $\mathbf{k} = \mathcal{K}(\mathbf{m}')$.
10. The decapsulated key is \mathbf{k} .