

Integer Reconstruction Public-Key Encryption

Houda Ferradi² and David Naccache¹

¹ Département d’informatique, École normale supérieure, Paris, France
45 rue d’Ulm, 75230, Paris CEDEX 05, France
david.naccache@ens.fr

² NTT Secure Platform Laboratories
3–9–11 Midori-cho, Musashino-shi, Tokyo 180–8585, Japan
ferradi.houda@lab.ntt.co.jp

Abstract In [AJPS18], Aggarwal, Joux, Prakash & Santha described an elegant public-key encryption (AJPS-1) mimicking NTRU over the integers. This algorithm relies on the properties of Mersenne primes instead of polynomial rings.

A later ePrint [BCGN17] by Beunardeau et al. revised AJPS-1’s initial security estimates. While lower than initially thought, the best known attack on AJPS-1 still seems to leave the defender with an exponential advantage over the attacker [dBDJdW17]. However, this lower exponential advantage implies enlarging AJPS-1’s parameters. This, plus the fact that AJPS-1 encodes only a single plaintext bit per ciphertext, made AJPS-1 impractical. In a recent update, Aggarwal et al. overcame this limitation by extending AJPS-1’s bandwidth. This variant (AJPS-ECC) modifies the definition of the public-key and relies on error-correcting codes.

This paper presents a different high-bandwidth construction. By opposition to AJPS-ECC, we do not modify the public-key, avoid using error-correcting codes and use backtracking to decrypt. The new algorithm is *orthogonal* to AJPS-ECC as both mechanisms may be concurrently used *in the same ciphertext* and cumulate their bandwidth improvement effects. Alternatively, we can increase AJPS-ECC’s information rate by a factor of 26 for the parameters recommended in [AJPS18].

The obtained bandwidth improvement and the fact that encryption and decryption are reasonably efficient, make our scheme an interesting post-quantum candidate.

Keywords: KEM, Efficiency improvement, MERS assumption, implementation

1 Introduction

Public-key encryption schemes that are mostly used today are RSA [RSA78] and ElGamal [ElG85], their security are based on the problem of factoring large composite integers or computing the discrete logarithms.

However, Shor [Sho97] showed an algorithm that quantum computers can use to solve both discrete logarithms and factoring in the relevant group in

polynomial time. Therefore, it is urgent to construct new encryption schemes which are considered secure even against quantum computer attacks.

For this purpose, in 2016 the National Institute of Standards and Technology (NIST) has initiated a competition where they are selecting post-quantum cryptographic algorithms that are supposed to resist attacks from future quantum computers [NIS17].

One promising candidate selected for such quantum-resistant encryption scheme is the AJPS cryptosystem [AJPS17a,AJPS18,AJPS17b], introduced by Aggarwal, Joux, Prakash, and Santha, which is an interesting alternative to well-established encryption schemes of NTRU. The security of AJPS encryption relies on the hardness of *Mersenne Low Hamming Combination* (MERS) problem [AJPS17a,AJPS18]: Given a Mersenne prime in the form $p = 2^n - 1$ (where n is prime), samples of the $\text{MERS}_{h,n}$ distribution are constructed as $(a, b = as + e)$, where $a \in \mathbb{Z}_p$ is chosen uniformly at random, the secret s and the error e are chosen uniformly at random from the elements in \mathbb{Z}_p of the Hamming weight h . The decisional version of the MERS assumption states that any efficient adversary cannot distinguish the $\text{MERS}_{h,n}$ distribution from the uniform distribution over \mathbb{Z}_p^2 .

Despite the efficiency benefit of its reliance on mersenne primes. The cryptosystem introduced in [AJPS18], however, remains inefficient because of the constraint that $n = \Theta(h^2)$ from the correctness of the key-encapsulation mechanism.

In this paper, we present new KEM schemes for enhancing the information rate of [AJPS18] by a factor of 26, which corresponds to $\mathcal{O}(1)$ improvement in the security parameter.

1.1 Related Works

Lattice-based cryptography is the most popular candidate for post-quantum security. Its security relies on the hardness of basis reduction and other related problems in random lattices, like Learning with Errors (LWE) based cryptosystems [Reg06], Ring-LWE based cryptosystems [LPR10] and NTRU cryptosystem [HPS98].

Concurrently to the above, another proposal for post-quantum cryptosystem comes from the Mersenne cryptosystem, which belongs to the family of NTRU [HPS98,Reg05,LPR13,MTSB13]. It is based on some rings in which elements are hidden by adding some small noise, which we will measure by the Hamming weight [AJPS18].

1.2 Organization of the Paper

In Section 2, we review the basic notion and notations, key-encapsulation mechanism and backtracking techniques. In Section 3, we review the MERS problem and its hardness. We then in Section 4 review the original AJPS scheme and its variant. In 5 we propose our KEM schemes, which are the Bivariate KEM and Trivariate KEM and prove their security. In Section 7, we provide an instantiation for the backtracking used in our KEMs. Finally, we conclude the paper in Section 9.

2 Preliminaries

2.1 Notation

We denote by $\|x\|$ the Hamming weight of an n -bit string x , which is the total number of 1's in x . Let $\mathfrak{H}_{n,h}$ be the set of all n -bit strings of Hamming weight h and $\{0,1\}^n$ the set of all n -bit strings.

Let \mathbb{Z}_p be the integer ring modulo p , where p is a Mersenne prime, i.e., numbers of the form $p = 2^n - 1$. We have the following property:

Lemma 1 (we have $x, y \in \mathbb{Z}_p$ then).

- $\|x + y \pmod{p}\| \leq \|x\| + \|y\|$
- $\|x \cdot y \pmod{p}\| \leq \|x\| \cdot \|y\|$
- If $x \neq 0^n$, $\|-x \pmod{p}\| = n - \|x\|$

The proof can be found in [AJPS18].

2.2 Key-Encapsulation Mechanism Syntax

A key-encapsulation mechanism (KEM) consists in four algorithms: $\Pi = (\mathbf{Setup}, \mathbf{KeyGen}, \mathbf{Encap}, \mathbf{Decap})$.

- **Setup**: The set up algorithm takes as input a security parameter λ and outputs a public parameter \mathbf{pp} .
- **KeyGen**: The key generation algorithm takes as input a public parameter \mathbf{pp} and outputs a public-key \mathbf{pk} and a secret key \mathbf{sk} .
- **Encap**: The encapsulation algorithm takes as input a public key \mathbf{pk} and outputs a ciphertext \mathbf{C} and key \mathbf{K} .
- **Decap**: The decapsulation algorithm takes as input a ciphertext \mathbf{C} and \mathbf{sk} and outputs \perp or a key \mathbf{K} .

Definition 1. We say that $\Pi = (\mathbf{Setup}, \mathbf{KeyGen}, \mathbf{Encap}, \mathbf{Decap})$ has $(1-\lambda)$ -correctness if for any $(\mathbf{pk}, \mathbf{sk})$ generated by **KeyGen**, we have that:

$$\Pr[\mathbf{Decap}(\mathbf{sk}, \mathbf{C}) = \mathbf{K} : (\mathbf{C}, \mathbf{K}) \leftarrow \mathbf{Encap}(\mathbf{pk})] \geq 1 - \lambda.$$

2.3 Backtracking techniques

For solving constraint satisfaction problems, basically there are three main approaches: backtracking search, local search, and dynamic programming. In this work, we are considering the backtracking search approach (also known as *depth-first search*), which consists of searching every possible combination in order to solve an optimization problem. For that, it proceeds in three steps: it picks a solution as a sequence of choices to the first sub-problem, then recursively attempting to resolve other sub-problems based on the solution of the first sequence of choices, then it returns the best solution found. We refer the reader to [Tan] for further introductory background.

3 MERS Assumption: Notation and Definitions

Aggarwal et al. introduced a new assumption [AJPS18] mimicking NTRU over integers, this assumption relies on the properties of Mersenne primes in the ring \mathbb{Z}_p instead of polynomial rings $\mathbb{Z}_q[x]/(x^n - 1)$. The conjecture is based on the observation that given any number $a \in \mathbb{Z}_p$ we obtain this property: if we multiply a by any number $b = 2^x$ and $x \in [0, n - 1]$, then the result $c = a \cdot b$ is just a cyclic shift. Since the reduction modulo Mersenne prime cannot increase Hamming weights, then the result c looks pseudorandom i.e. it is hard to distinguish c from any random integer modulo Mersenne prime.

The security of our scheme are based on the following assumptions:

3.1 MERS assumption

For two integers $4h^2 < n$ and for n -bit Mersenne prime $p = 2^n - 1$, and for integer $s \in \mathbb{Z}_p$, we define a distribution $\text{MERS}_{s,n,h}$ as follows: choose $r \leftarrow \{0, 1\}^n$ and $b \leftarrow \mathfrak{H}_{n,h}$, return $(r, r \cdot s + b \bmod p)$. We also define a uniform distribution U as follows: choose $r \leftarrow \{0, 1\}^n$ and $b \leftarrow \{0, 1\}^n$, return $(r, b \bmod p)$.

Definition 2 (The (decisional) Mersenne Low Hamming Combination Assumption (MERS assumption) [AJPS18]). For two positive integers $4h^2 < n$ and for an adversary \mathcal{A} , we introduce the $\text{MERS}_{n,h}$ advantage as the quantity:

$$\text{Adv}^{\text{MERS}_{h,n}}(\mathcal{A}) = \left| \Pr[\mathcal{A}^{\text{MERS}_{s,n,h}(\cdot)} \Rightarrow \text{True}] - \Pr[\mathcal{A}^{U(\cdot)} \Rightarrow \text{True}] \right|,$$

where $s \xleftarrow{\$} \mathfrak{H}_{n,h}$. We say that the MERS problem $\text{MERS}_{n,h}$ is (t, q, ϵ) -hard if all attacker \mathcal{A} with time complexity t , making at most q queries, we have $\text{Adv}^{\text{MERS}}(\mathcal{A}) \leq \epsilon$.

3.2 RMERS assumption

Definition 3 (The (search) Mersenne Low Hamming Ratio Search Assumption (RMERS) [AJPS18]).

Given an n -bit Mersenne prime $p = 2^n - 1$, $4h^2 < n$ and an integer $H \in \mathbb{Z}_p$, find $F, G \in \mathfrak{H}_{n,h}$ such that

$$H = \frac{F}{G} \bmod p$$

3.3 Hardness of MERS problem

MEET-IN-THE-MIDDLE ATTACK. de Boer et al. [dBDJdW17] presented a meet-in-the-middle attack for solving the MERS problem. Their classical attack runs in the time $\tilde{O}\left(\binom{n-1}{h-1}^{1/2}\right)$. The quantum version runs in the time $\tilde{O}\left(\binom{n-1}{h-1}^{1/3}\right)$.

LLL-ATTACK. [BCGN17,dBDJdW17] present an LLL-based algorithm for solving the RMERS, whose the running time is $O(2^{2h})$ on Turing machine and $O(2^h)$ on quantum machine. This algorithm is the most efficient known approach for solving the RMERS. Assuming a quantum computer using Grover's algorithm, one obtains a quadratic speedup over [dBDJdW17] (i.e. $O(2^h)$). Therefore, as per [AJPS18], our scheme is secure when $h = \lambda = 256$.

PRIMALITY OF n IN MERSENNE PRIMES. Agrawal discussed that $p = 2^n - 1$ and n should be primes to avoid an attack on composite n . For the details, see Aggrawal et al. [AJPS17a,AJPS18].

4 The Original AJPS Cryptosystem

4.1 The AJPS-1 Encryption Scheme

The original AJPS-1 encryption scheme based on MERS assumption [AJPS17a], is defined by the following sub-algorithms:

- **Setup**(1^λ) \rightarrow pp . Chooses the public parameters $\text{pp} = \{n, h\}$ so that $p = 2^n - 1$ is an n -bit Mersenne prime achieving some λ -bit security level.
- **KeyGen**(pp) \rightarrow $\{\text{sk}, \text{pk}\}$. Picks $\{F, G\} \in_R \mathfrak{H}_{n,h}^2$ and returns:

$$\begin{cases} \text{sk} & \leftarrow G \\ \text{pk} & \leftarrow H = F/G \bmod p \end{cases}$$

- **Enc**($\text{pp}, \text{pk}, m \in \{0, 1\}$) $\rightarrow C$. Picks $\{A, B\} \in_R \mathfrak{H}_{n,h}^2$, and computes:

$$C \leftarrow (-1)^m (AH + B) \bmod p$$

- **Dec**(pp, sk, C) $\rightarrow \{\perp, 0, 1\}$, computes $d = \|GC \bmod p\|$ and returns:

$$\begin{cases} 0 & \text{if } d \leq 2h^2, \\ 1 & \text{if } d \geq n - 2h^2, \\ \perp & \text{otherwise} \end{cases}$$

The intuition behind the decryption formula is the observation that when $m = 0$ we get:

$$W = GC = G(AH + B) = FA + GB \Rightarrow W \text{ is of low Hamming weight}$$

The security of this cryptosystem is based on the (decisional) MERS problem. We refer the reader to [AJPS17a] for further details.

To increase bandwidth, Aggarwal et al. introduced the AJPS-ECC variant described hereafter.

4.2 The AJPS-ECC Encryption Scheme

In the second variant AJPS-ECC [AJPS18,AJPS17b] aims to extend AJPS-1's bandwidth, while requiring an ancillary error correction scheme $\{\mathcal{D}, \mathcal{E}\}$.

AJPS-ECC is formally defined by the following sub-algorithms:

- **Setup**(1^λ) \rightarrow pp. As in Section 4.1.
- **KeyGen**(pp) \rightarrow {sk, pk}. Picks $\{F, G\} \in_R \mathfrak{H}_{n,h}^2$, $R \in_R \{0, 1\}^n$ and returns:

$$\begin{cases} \text{sk} & \leftarrow F \\ \text{pk} & \leftarrow \{R, T\} = \{R, F \times R + G \bmod p\} \end{cases}$$

- **Enc**(pp, pk, $m \in \{0, 1\}^\lambda$) \rightarrow C . Picks $\{A, B_1, B_2\} \in_R \mathfrak{H}_{n,h}^3$ and computes the ciphertext:

$$C = \begin{cases} C_1 & \leftarrow A \times R + B_1 \bmod p \\ C_2 & \leftarrow (A \times T + B_2 \bmod p) \oplus \mathcal{E}(m) \end{cases}$$

- **Dec**(pp, sk, C) \rightarrow $\{\perp, m\}$ returns:

$$\mathcal{D}((F \times C_1 \bmod p) \oplus C_2).$$

For the sake of clarity, we keep the definition of C_1 unchanged but slightly depart from [AJPS17a]'s original formulae by modifying the definitions of T and C_2 as follows:

$$\begin{aligned} T & \leftarrow F \times R - G \bmod p \\ C_2 & \leftarrow (A \times T - B_2 \bmod p) \oplus \mathcal{E}(m) \end{aligned}$$

To understand the intuition behind **Dec** consider the quantity $W = FC_1 - C_2$ corresponding to the particular case $\mathcal{E}(m) = 0$:

$$W = FAR + FB_1 - AT + B_2 = FAR + FB_1 - A(FR - G) + B_2 = FB_1 + GA + B_2$$

As before, we see that $d = \|W\|$ is low. This means that the noise attached to $\mathcal{E}(m)$ after the clean-off operation $(F \times C_1 \bmod p) \oplus C_2$ is low and thus surmountable by the error-correcting code $\{\mathcal{E}, \mathcal{D}\}$. The security of this cryptosystem is based on the (decisional) MERS problem. We refer the reader, again, to [AJPS18] for further details about this cryptosystem and the parameter choices allowing successful decryption and sufficient security. Sticking only to the core idea, we purposely omit the hashing and re-encryption tests performed during the key de-encapsulation process.

5 Proposed Schemes

5.1 Overview of Our Approach

In this section we will describe two AJPS-ECC variants based on the new idea of *Randomness Reconstruction*.

Our idea departs from AJPS-1 in a direction orthogonal to the above.

We set by design $m = 0$ in AJPS-1 or $\mathcal{E}(m) = 0$ in AJPS-ECC and attempt to recover *the randomness*³ into which information (encapsulated keys and/or plaintext information) will be embedded.

The intuition is that the receiver might be able to recover the randomness if parameters are properly chosen *using his extra knowledge* of G, F and knowing that, in addition, the unknown randomness has a low Hamming weight.

We hence focus the rest of this paper on methods for solving the equations in the following forms:

$$W = Fx + Gy \quad \text{or} \quad W = Fx + Gy + z \pmod{p}$$

Where all parameters⁴ and unknowns are *randomly* chosen in $\mathfrak{H}_{n,h}$ and where a solution $\{x, y\}$ or $\{x, y, z\}$ is *known to exist*.

We do not introduce any modifications in **Setup** and **KeyGen**, nor do we modify **pp** or **sp**⁵. We thus focus on the encapsulation (encryption) and on the de-encapsulation (decryption) processes only, in KEM and PKE respectively.

In a non-KEM version, a plaintext m encoded in the unknowns $(x, y$ or $x, y, z)$ can be directly recovered upon decryption. Such an encryption mode must however be protected against active attacks using padding and randomization that we do not address here.

Remark 1. It is tempting but inadvisable to create dependencies between the variables F, G and/or the unknowns x, y, z . Consider an AJPS-ECC where $m \in_R \{0, 1\}^\lambda$ and $\{A, B_2\} \in_R \mathfrak{H}_{n,h}^2$ but where $B_1 \leftarrow \mathcal{H}(m)$ is obtained by hashing m into $\mathfrak{H}_{n,h}$. Given m , *anybody* can re-compute B_1 and algebraically infer A, B_2 . We hence see in this example that A, B_2 do not add extra entropy as security solely rests upon m .

We carefully distinguish between *security bandwidth* and *information rate*. An idea, unexplored in AJPS-ECC, may exploit remark 1 to transport more plaintext information in $\{C_1, C_2\}$ *without adding extra security*. To encrypt a τ -bit message μ , pick a key $m \in_R \{0, 1\}^\lambda$ and encrypt $c \leftarrow \mathcal{F}_m(\mu)$ using a block cipher \mathcal{F} . Set $B_1 \leftarrow \mathcal{H}(m)$. Let \mathcal{M} be any invertible public mapping $\mathcal{M} : \{0, 1\}^\tau \rightarrow \mathfrak{H}_{n,h}^2$. Encode: $\{A, B_2\} \leftarrow \mathcal{M}(c)$ and form $\{C_1, C_2\}$ using AJPS-ECC. To decrypt, recover m using error-correction, recompute B_1 , algebraically recover $\{A, B_2\}$ and retrieve the plaintext μ by:

³ A, B or A, B_1, B_2

⁴ Except W

⁵ Note that in AJPS-1/ECC given G one can compute F and *vice versa*.

$$\begin{aligned}
\mu &\leftarrow \mathcal{F}_m^{-1}(\mathcal{M}^{-1}(A, B_2)) \\
&= \mathcal{F}_m^{-1}\left(\mathcal{M}^{-1}\left(\frac{C_1 - B_1}{R} \bmod p, C_2 - \frac{(C_1 - B_1)T}{R} \bmod p\right)\right) \\
&= \mathcal{F}_m^{-1}\left(\mathcal{M}^{-1}\left(\frac{C_1 - \mathcal{H}(m)}{R} \bmod p, C_2 - \frac{(C_1 - \mathcal{H}(m))T}{R} \bmod p\right)\right)
\end{aligned}$$

Because in AJPS-ECC $\{n, h\} = \{756839, 256\}$ the potential encoding capacity of \mathcal{M} can be relatively high:

$$2 \log_2 \binom{756839}{256} = 6631 \text{ bits}$$

This increases AJPS-ECC's information rate by a factor of 26. Again, proper message padding may be necessary to resist active attacks (analysis underway). We stress, again, that this does not increase security bandwidth but information rate only. An attacker guessing m will determine μ .

5.2 The Bivariate KEM

Our Bivariate Cryptosystem $\Pi_1 = (\text{Setup}, \text{KeyGen}, \text{Encap}, \text{Decap})$ is defined as follows:

- **Setup**(1^λ) \rightarrow pp: As in Section 4.1.
- **KeyGen**(pp) \rightarrow {sk, pk} are identical to AJPS-1.
- **Encap**(pp, pk) \rightarrow C . Picks $\{A, B\} \in_R \mathfrak{H}_{n,h}^2$ and computes:

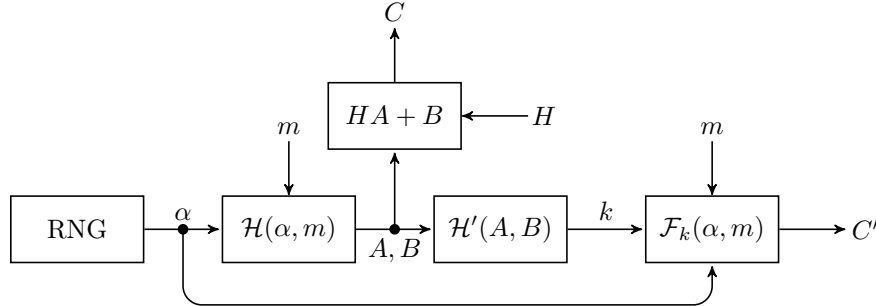
$$C \leftarrow AH + B \bmod p$$

- **Decap**(pp, sk, C) \rightarrow $\{\perp, \{A, B\}\}$ returns:

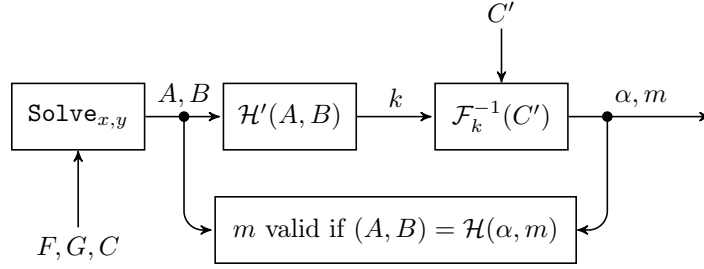
$$\{A, B\} \leftarrow \text{Solve}_{x,y}[GC = Fx + Gy \bmod p]$$

If $\{A, B\} \neq \perp$ use $\{A, B\}$ as KEM entropy for further encryption.

Example: Let μ be a plaintext and \mathcal{R} a redundancy function. Compute $m \leftarrow \mathcal{R}(\mu, \rho)$ where ρ is random. A typical KEM⁶ is shown here:



⁶ Where $\mathcal{H}, \mathcal{H}'$ s are hash functions and \mathcal{F} is a block-cipher.



Retrieve m from $\mu \leftarrow \mathcal{R}^{-1}(m)$.

5.3 The Trivariate KEM

Our Trivariate Cryptosystem $\Pi_2 = (\mathbf{Setup}, \mathbf{KeyGen}, \mathbf{Encap}, \mathbf{Decap})$ is defined as follows:

- **Setup**(1^λ) \rightarrow pp: As in Section 4.1
- **KeyGen**(pp) \rightarrow {sk, pk} are identical to AJPS-ECC but with the modified formula $T \leftarrow F \times R - G \bmod p$.
- **Encap**(pp, pk) \rightarrow C . Picks $\{A, B_1, B_2\} \in_R \mathfrak{H}_{n,h}^3$ and computes the ciphertext:

$$C = \begin{cases} C_1 & \leftarrow A \times R + B_1 \bmod p \\ C_2 & \leftarrow A \times T - B_2 \bmod p \end{cases}$$

- **Decap**(pp, sk, C) \rightarrow $\{\perp, \{A, B_1, B_2\}\}$ returns:

$$\{A, B_1, B_2\} \leftarrow \text{Solve}_{x,y,z}[FC_1 - C_2 = Fy + Gx + z \bmod p]$$

If $\{A, B_1, B_2\} \neq \perp$ use $\{A, B_1, B_2\}$ as KEM entropy for further encryption.

As noted before, the trivariate version may accommodate in the encryption formula an *independent* $\mathcal{E}(m)$ and thus cumulate the bandwidth improvements due to both mechanisms. This requires that the $\{n, h\}$ values of both schemes coincide and the enforcement of the condition $n \leq 16h^2$, not addressed here. We conjecture that such a meeting point exists.

6 Security Proof

Just as RSA, our encryption process is deterministic (*i.e.*, requires no nonce) if the message is encoded in A, B . As such, unless we use a proper padding before encryption, the encryption function itself cannot provide "native" indistinguishability against chosen plaintext attacks.

Hence provide a proof that breaking our KEM is equivalent to the solving the (search) RMERS Problem, defined in Section 3.2.

Theorem 1. *Inverting the KEM described in section 5.2 with parameter h is as hard as solving RMERS Problem for parameter $h/2$.*

Proof. We will show that any attack algorithm $\mathcal{A}(C, H, h, p) = \{A, B\}$ extracting the exchanged key $\{A, B\}$ without resorting to the secret key can be used to solve the RMERS for parameters $h/2, p$. There is hence a loss in the reduction of a factor of 2 in the security parameter h .

Assume that $\mathcal{A}(C, H, h, p)$ exists.

We note that \mathcal{A} resolves the equation $C = AH + B \pmod p$ without resorting to the secret elements F, G .

What happens if we invoke $\mathcal{A}(0, -H \pmod p, h, p)$?

In such a case \mathcal{A} will return A, B such that:

$$0 = -AH + B \pmod p \text{ that is: } H = \frac{B}{A} = \frac{F}{G} \pmod p$$

\mathcal{A} has thus solved the RMERS for parameters h, p .

This is, however, insufficient as \mathcal{A} may refuse to solve the equation $C = AH + B \pmod p$ when $C = 0$. We will hence mask the input C so that \mathcal{A} could not refuse to process it.

To do so, we sacrifice reduction tightness to force \mathcal{A} to solve arbitrary target instances H_t of the RMERS for parameters $h/2, p$.

- Generate two random numbers $r_A, r_B \in_R \mathfrak{S}_{n, h/2}^2$.
- Form the quantity:

$$C = H_t \times r_A - r_B \pmod p$$

- Invoke $\mathcal{A}(C, -H_t \pmod p, h, p)$

\mathcal{A} will return an A, B such that:

$$C = -A \times H_t + B = H_t \times r_A - r_B \pmod p$$

In other words:

$$H_t = \frac{B + r_B}{A + r_A} = \frac{F}{G} \pmod p$$

\mathcal{A} was hence instrumented to solve a RMERS instance of parameter $h/2$, as required.

□

The proof extends, *mutatis mutandis*, to the trivariate KEM of section 5.3.

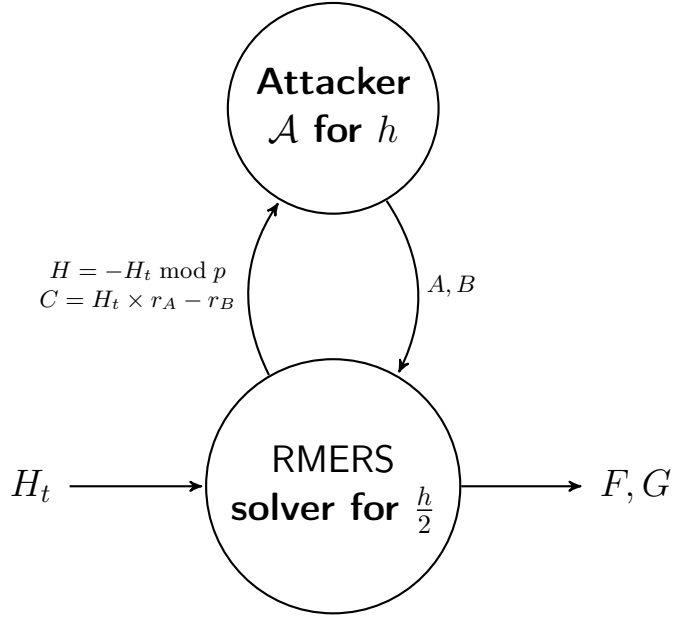


Figure 1: Security reduction: Turning an attacker into a RMERS solver.

7 Instantiating $\text{Solve}_{x,y}$ Using Backtracking

In this section explains how to instantiate $\text{Solve}_{x,y}$. The routine $\text{Solve}_{x,y,z}$ is obtained *mutatis mutandis*.

The intuition behind $\text{Solve}_{x,y}$ is the following: assume that we are given the quantity $W = GC = AF + BG \bmod p$ where $\|W\| \cong 2h^2$. Because multiplication modulo p is (somewhat) weight-preserving, we can test the hypothesis that the i -th bit of A is equal to one by looking at the quantity Δ :

$$\Delta = \|W\| - \|W - 2^i F \bmod p\|$$

Intuitively, a good guess should result in a weight decrease of $\simeq h$ whereas a wrong guess should re-blur W by triggering random carry propagations. Evidently, because there may be false positives during this process, we must be able to backtrack. To reduce the false positive error probability, n must be large enough with respect to h . The exact same idea applies to $\text{Solve}_{x,y,z}$.

7.1 Prerequisites & Subroutines

We start by introducing three necessary prerequisites.

Γ	50	51	52	53	54	55	56	57	58	59	60
Probability	24%	20%	26%	30%	32%	20%	22%	18%	14%	9%	4%

Table 1: Backtracking success chances for $\{t, h, n\} = \{1, 72, 19937\}$. 50 decryption simulations per entry.




Figure 2: Backtracking success chances for $\{t, h, n\} = \{1, 65, 19937\}$. 200 decryption simulations per entry. Fitted with $82.922 \exp(-(\Gamma - 45.7122)^2/34.6815)$

The ancillary function Confirm: Our algorithms require an ancillary function `Confirm`-ing a candidate solution $\{x, y\}$. e.g. given a candidate x , `Confirm`(x) may solve $GC = Fx + Gy \pmod p$ for y and return $\{y, \text{True}\}$ if $\|x\| = \|y\| = h$. Because in some cases several solutions may exist, a simpler implementation may just compare $\mathcal{H}(x, y)$ to a confirmation digest τ provided with the ciphertext and return $\{y, \text{True}\}$ if the purported solution hashes into τ . If $\mathcal{H}(x, y) \neq \tau$ then `Confirm`(x) returns $\{\perp, \text{False}\}$.

Determining the backtracking aperture Γ : Backtracking is parametrized by a constant Γ controlling the aperture of the exhausting process (*i.e.* the marginal tolerance allowing to exclude a search path from further investigation). Simulations indicate that for any given $\{n, h\}$ there is a Γ_{optimal} value minimizing the failure probability. We did not attempt a formal analysis of the dependency between $\{n, h\}$ and Γ_{optimal} but estimated Γ_{optimal} for various $\{n, h\}$ pairs using simulations as shown in Table 1 and Figure 2.

7.2 The Backtracking Algorithms

The deterministic backtracking algorithm \mathcal{B}_1 subtracts left-shifted F s from W to obtain candidate w s having smaller and smaller weights. \mathcal{B}_1 maintains a set of integers R containing the bit positions of x discovered so far. The deterministic algorithm is called by $\{A, B\} \leftarrow \mathcal{B}_1(W, \emptyset, 0)$ and the randomized version is called by $\mathcal{B}_2(W, \emptyset, 0, \phi_j)$ where ϕ_j s are random permutations of \mathbb{Z}_k . Code is available from the authors upon request.

ALGORITHM 1

Backtracking $\mathcal{B}_1(w, R, e)$

Input: w, R, e . The values $G, F, h, n, p = 2^n - 1, C$ are global and invariant.

Output: $\{x, y\} = \{A, B\}$ such that $W = CG = Fx + Gy \pmod p$ or Failure.

```

if  $e = n$  then return Failure
else
  if  $\#R = h$  then
     $x \leftarrow \sum_{i \in R} 2^i$ 
     $\{s, y\} \leftarrow \text{Confirm}(x)$ 
    if  $s$  then return  $\{x, y\}$ 
   $\bar{w} \leftarrow w - 2^e \times F \pmod p$ 
  if  $|\|\bar{w}\| - \|w\| + h| \leq \Gamma$  then
     $\mathcal{B}_1(\bar{w}, R \cup \{e\}, e + 1)$ 
  else
     $\mathcal{B}_1(w, R, e + 1)$ 

```

ALGORITHM 2

Backtracking $\mathcal{B}_2(w, R, e, \phi)$

Input: w, R, e, ϕ . The values $G, F, h, n, p = 2^n - 1, C$ are global and invariant.

Output: $\{x, y\} = \{A, B\}$ such that $W = CG = Fx + Gy \pmod p$ or Failure.

```

if  $e = n$  then return Failure
else
  if  $\#R = h$  then
     $x \leftarrow \sum_{i \in R} 2^{\phi(i)}$ 
     $\{s, y\} \leftarrow \text{Confirm}(x)$ 
    if  $s$  then return  $\{x, y\}$ 
   $\bar{w} \leftarrow w - 2^{\phi(e)} \times F \pmod p$ 
  if  $|\|\bar{w}\| - \|w\| + h| \leq \Gamma$  then
     $\mathcal{B}_2(\bar{w}, R \cup \{e\}, e + 1, \phi)$ 
  else
     $\mathcal{B}_2(w, R, e + 1, \phi)$ 

```

We conjecture that working with a fixed Γ during the entire backtracking process handicaps the algorithm. When the process starts the weight of W is

high, hence the probability to strike-out h bits by subtraction is high. However as subtractions make w sparser aperture should intuitively decrease. It may hence make sense to explore algorithms in which the constant T_{optimal} is replaced by a function $\Gamma(\|w\|, \|\bar{w}\|, n, h)$.

Best candidate search: \mathcal{B}_1 and \mathcal{B}_2 explore all the paths starting by an *a priori* promising Δ . However, \mathcal{B}_1 and \mathcal{B}_2 do not explore the most promising paths first. A more complex backtracking strategy (\mathcal{B}_3) trying with priority the paths starting by a Δ as close as possible to h was developed as well (information available from the authors upon request). We do not include this algorithm here for the sake of concision.

Dealing with decoding failures: Because we may discard seemingly uninteresting (but actually promising) exploration paths, backtracking may fail to decode W . As it seems complex to formally compute the algorithm's success probability, we estimated it by simulation. To deal with decryption failures we re-attempt backtracking after index randomization *i.e.* pick t random permutations $\{\phi_0, \dots, \phi_{t-1}\}$ of \mathbb{Z}_k and re-run $\mathcal{B}_2(W, \emptyset, 0, \phi_j)$ t times hoping that at least one of the t runs will succeed⁷. A more brutal approach consists in sending t encapsulated keys to increase the probability exponentially. This (conjectured) exponential probability gain only handicaps the information rate by a constant factor⁸. A simple idea for (conjectured) squaring the failure probability consists in trying to backtrack on A and, upon failure, re-launch the algorithm to backtrack on B .

Information leakage from decryption failures: Because decryption may fail, a possible cryptanalysis (that we did not investigate) might be to analyze, possibly adaptively, the ciphertexts causing failures and thereby extract information on $\{F, G\}$. We do not regard this as a major problem for the following reasons:

- Failure is highly dependent on the backtracking algorithm chosen by the receiver. The backtracking procedures that we give here are one possibility amongst many.
- An empirical protection consists in randomizing the backtracking process, *e.g.*, assume all the ϕ_i to be *randomly drawn per decryption and secret*.
- Another protection is to purposely fail decryption with some probability ϵ to prevent the cryptanalyst from identifying true failures. Note that the random tape used to simulate false failures must be derived from a fixed secret and *the ciphertext itself* to avoid replays and majority votes.

⁷ Note that $\mathcal{B}_1(W, \emptyset, 0) = \mathcal{B}_2(W, \emptyset, 0, \text{ID})$.

⁸ Link B_0, \dots, B_{t-1} in a way allowing the recovery of all the B_i if one of them is known (*e.g.* define $B_i = \mathcal{F}_i(\text{seed})$ where $\mathcal{F}_k(m)$ is a block-cipher encrypting into $\mathfrak{H}_{n,h}$). Use the A_i to transport entropy or information. One successful decryption reveals the seed \Rightarrow open all the B_i s \Rightarrow all the t information containers A_i .

Protection against side-channel attacks: It is reasonable to assume that, like most encryption schemes, the algorithms described in this paper are vulnerable to timing and side-channel attacks, an aspect that we did not investigate here.

7.3 Eccentric Reconstruction Strategies

Backtracking might be improved in a variety of ways. As examples, we are introducing in this section a list a few research ideas that we did not explore in detail.

Idea#1: Brittle encryption formulae. We may modify the bivariate encryption formula to $C \leftarrow HA + 3B$ and enforce by design that H, A, B do not contain the binary sequence 11. This means that the bit positions representing $3B$ will be “colored” by a pattern 11 making their isolation and identification easier. If $n \gg h$ we may even attempt to brutally reset all the isolated ones in W and divide the result by $3G$ to directly obtain B . For the trivariate version one may use:

$$C = \begin{cases} C_1 & \leftarrow AR + B_1 \pmod p \\ C_2 & \leftarrow AT - 3B_2 \pmod p \end{cases}$$

resulting in the decryption formula $W = FC_1 - C_2 = FB_1 + GA + 3B_2$. Here as well, we banish the pattern 11 from G, F, A, B_1, B_2 . We may thus attempt to identify in W the binary patterns 11, hinting the probable presence of B_2 to ease decoding. Note that the pattern 11 may result naturally from the multiplication, the addition or the reduction and hence mislead the decoder (backtrack). Similarly, an 11 due to $3B_2$ may disappear due to addition (backtrack). Note that marking B with 11s makes backtracking more efficient as this increases the SNR. e.g. if we replace each 1 in B by a 1111 we increase overall weight of W to $5h^2$ but a correct guess will cause a weight decrease of $\simeq 4h$ instead of h . In other words, while requiring a larger n , this improves the SNR:

$$\text{from SNR} = \frac{h}{2h^2} = \frac{1}{2h} \text{ to } \text{SNR} = \frac{4h}{5h^2} = \frac{4}{5h}$$

Idea#2: Dye tracing. In hydrogeology, *dye tracing* is a technique for tracking various flows using dye added to the water source. In other words, dye tracing uses dye as a flow tracer. It is an evolution of the ages-known float tracing method, which consists of throwing a buoyant object into a waterflow to see where it emerges. To simulate the effect of dye tracing, we inject into F 's digits a few low-weight binary patterns and track their appearance in W . For instance (toy example), generate an F of weight $h - 10$ not containing any of the ten sequences ℓ_i :

11	101	111	1001	1011	1101	10001	11001	10101	10011
----	-----	-----	------	------	------	-------	-------	-------	-------

randomly insert those ten ℓ_i s into F 's blank spaces (insert each ℓ_i once, this will increase the weight of F to $h - 10 + \sum \|\ell_i\| = h - 10 + 26 = h + 16$ and the weight of W to $\simeq 2h^2 + 16h$). To retrieve A , isolate the 10 dyes tracers in W and use majority voting on bit offsets to infer the probable positions of A 's bits.

Idea#3: Demodulation. We can attempt to “travel back in time” and infer $\omega = FA + GB \in \mathbb{Z}$ from W , or at least estimate the probability that a candidate bit in W originates from the number’s pre-reduced upper half. Given $\omega \in \mathbb{Z}$ decryption⁹ is immediate because:

$$A = \omega F^{-1} \bmod G = (W \text{ demod } p) \times F^{-1} \bmod G$$

To demodulate W we work modulo $p = 2^n - 3$ that “colors” the folded MSBs by turning them into LSB 11s. The process is error-prone¹⁰ but *actually works* for parameters that are large enough. We implemented the idea very brutally, by simply translating each 11 in W into a 1 in the MSB of ω without taking any further precautions. 100 demodulation attempts for $\{n = 6 \times 10^7, h = 55\}$ resulted in 29 successes. Although n is huge, the resulting information rate is not “that” catastrophic as we can pack:

$$2 \log_2 \binom{6 \times 10^7}{55} = 2356 \text{ plaintext bits into the ciphertext.}$$

In other words, each plaintext bit claims 25461 ciphertext bits and is successfully transmitted with probability 29%.

While $h = 55$ is not very large and $n = 6 \times 10^7$ is extremely large, our simulation shows that it is definitely possible to make ingredients meet at a workable parameter combination. We conjecture that with proper analysis and refined demodulation strategies k might be reduced by at least two orders of magnitude. It may also be possible to work modulo $2^n - \pi$ with a more distinguishable color $\pi \neq 3$ despite an extra weight due to a more complex π . $\pi = -1$ is interesting as well as -1 turns folded bits into long chains of 1s.

Note 1. (important) One of the features preventing lattice-based attacks in AJPS-1/ECC is the emergence of parasitic short vectors due to working modulo $2^n - 1$ (section 5.1.[?]¹¹). We did not evaluate the impact of $\pi \neq 1$ on the number and the norm of parasitic short vectors *and hence on security*.

⁹ Take F, G coprime in **Setup**.

¹⁰ Again, “natural” 11s may be already present in the LSBs of ω , 11 + 01 may destroy an 11, 10 + 01 may create fake 11s etc.

¹¹ ePrint version 20170530:072202.

Idea#4: Pattern identification. Another idea consists in exploiting the fact that $W = AF + BG \bmod p$ will naturally contain binary sequences of the form:

$$v_\ell = \underbrace{0, \dots, 0}_{\ell \text{ zeros}} | 1 | \underbrace{0, \dots, 0}_{\ell + 1 \text{ zeros}}$$

Let m be an ℓ -bit encapsulated key¹² and define $C = m(AH + B) \bmod p$ we get:

$$CG = m(AF + GB) = mW = m \times (w' | v_\ell | w) = u' | m | u \bmod p$$

m can thus be read¹³ on $CG \bmod p$. It remains to identify m . To do so, we can generate $\{A, B\} \leftarrow \mathcal{H}(m)$ and hence confirm proper decryption using n re-hashings and re-encryptions. This workload may be considerably reduced by sacrificing a few bits of m , e.g. 16 bits, to display a specific pattern (e.g. $0x\text{FFFF}$) allowing a quick identification of m . This divides the number of hashings and re-encryptions by 2^{16} . As a numerical example, $\{n, h\} = \{75 \times 10^4, 100\}$ corresponds to an $\ell \simeq 200$. If we sacrifice 20 bits for an identification pattern we can hope to decapsulate a $\simeq 160$ -bit key using one re-encryption only.

ℓ has a low variance as it is essentially determined by a max-min over the differences between the positions s_i of the bits equal to one in W :

$$\ell \cong \max_i \min(s_i - s_{i-1}, s_{i+1} - s_i)$$

For a subtle technical reason ℓ is actually higher than this crude estimate. To understand why we refer the reader to Figure 3 where we illustrate the expected distribution of $\tilde{h} = 2h^2$ bits amongst n potential positions. The least significant 1-bit \bullet is expected to appear at γ where:

$$\gamma = \frac{\tilde{h}}{(n-1)^{\tilde{h}}} \int_0^{n-1} x(n-1-x)^{\tilde{h}-1} dx = \frac{n-1}{\tilde{h}+1}$$

Similarly, the most significant 1-bit position \bullet is expected at $\simeq \tilde{h}\gamma$. The reason why two other points are singled-out by \bullet and \bullet will be clarified later. Now, because arithmetics modulo p wrap everything that overflows 2^n on 2^0 the actual gap between \bullet and \bullet is not γ but 2γ . This is illustrated in Figure 4. In other words, the primary formula for ℓ should be corrected to:

$$\ell \cong \max \left(\underline{\alpha}, \bar{\alpha}, \max_i \min(s_i - s_{i-1}, s_{i+1} - s_i) \right)$$

Where:

$$\underline{\alpha} = \min([\bullet, \bullet], [\bullet, \bullet]) \stackrel{\text{u}}{=} [\bullet, \bullet] \simeq \gamma$$

and

¹² We consider m to be beyond exhaustive search, typically 160 bits.

¹³ Note that reading is circular i.e. wrapping around $CG \bmod p$.

$$\bar{\alpha} = \min([\bullet, \bullet], [\bullet, \bullet]) \stackrel{\text{u}}{=} [\bullet, \bullet] \simeq \gamma$$

Where $\stackrel{\text{u}}{=}$ denotes “usually (or frequently) equal to”. We therefore see that wrapping due to modular arithmetic has an unexpected favorable effect on ℓ .



Figure 3: Dots show the expected positions of h bits picked randomly amongst n positions. Here $\gamma = \frac{n-1}{h+1}$.




Figure 4: The interval $[\min, \max] = [\gamma, h\gamma]$ of size 2γ created by wrapping mod p .

Pattern identification is somewhat homomorphic but with a very fast increasing noise: encode a zero as $m = 01$, a one as $m = 11$ and read the plaintext on the most significant bit of that encoding.

Idea#5: Prime embedding. A variant of the above, mostly of theoretical interest, is the following: because there is a number of natural leading and tailing zeros in $\eta = AF + GB \pmod p$ we can encode $C = m(AH + B) \pmod p$, recover $m(AF + GB) \pmod p$ and, provided that m is short enough, hope that η is small enough to get¹⁴ $\omega = m\eta \in \mathbb{Z}$. It remains to extract m from ω . To do so, pick m as a product of, say 64-bit random primes. The receiver can pull-out those primes from ω using ECM factorization¹⁵. When a candidate m was formed, confirm it using hashing and re-encryption as before. $\{n, h\} = \{2.5 \times 10^6, 100\}$ gives an average expected margin of $\simeq 240$ bits for encoding m . The main problem with

¹⁴ the possibly rotated

¹⁵ Accidental similar-size prime factors may come from $AF + GB$ as well, but those are few and hence easy to filter.

this variant would be the high variance in the size of m embeddable into the ciphertext which would make decryption very uncertain.

Note 2. (research note) To the above we add two ideas that we conjecture to be insecure (by opposition to the previous ideas that we do not conjecture to be secure)

Variant 9.1. Correlated As: To ease backtracking we wish to give the decoder several Δ s generated from the same B . Generate t independent keys $\{F_i, G_i, H_i\}$ (possibly modulo different p_i s). Pick $t-1$ public random permutations $\phi_1, \dots, \phi_{t-1}$ of \mathbb{Z}_n . Generate $\{B_0, \dots, B_{t-1}, A_0\} \in_R \mathfrak{H}_{n,h}^{t+1}$. Let

$$A_0 = \sum_{i=0}^{n-1} 2^i a_i$$

and form t ciphertexts $\{C_0, \dots, C_{t-1}\}$:

$$C_j = A_j H_j + B_j \pmod{p_j} \quad \text{where} \quad A_j = \sum_{i=0}^{n-1} 2^{\phi_j(i)} a_i$$

Simultaneous backtracking on the A_j s will reveal more information per bit guess to the decoder.

Variant 9.2. Correlated Hs: Set G and define $H_i = F_i/G$ for $i \geq 1$. We illustrate the idea with two H_i s. Encrypt $C = A_0 H_0 + A_1 H_1 + B$. We see that $W = GC = F_0 A_0 + F_1 A_1 + BG$. Linking A_0 and A_1 as in note 9.1 we see that the SNR¹⁶ in the case of a successful guess increases:

$$\text{from SNR} = \frac{h}{2h^2} = \frac{1}{2h} \quad \text{to} \quad \text{SNR} = \frac{2h}{3h^2} = \frac{2}{3h}$$

This modifies the complexity assumption as well.

Note 3. (a broken variant) We close this paper by attracting the reader's attention to the broken variant given in the appendix, that we mention as a target for fixing.

8 Security & Parameter Sizes

Brittle encryption, dye tracing, demodulation, pattern identification and prime embedding are only illustrative research directions that we consider interesting or curious but that we do not claim nor conjecture to be secure. This work did not cover the security of the proposed constructions and focused on the textbook modes in which data is encoded and decoded. Parameter sizes were not recommended and numerical examples are given for illustrative purposes.

¹⁶ Note that as backtracking proceeds the SNR improves. In this paper SNR stands for the SNR at the beginning of the backtracking process.

A careful trade-off must be established between ① the security, ② the (backtracking failure probability) and ③ the efficiency of the various Solve processes. So far, simulations indicate that there are ways to practically satisfy those three constraints at once.

9 Concluding Remarks

In this paper, we have introduced a new KEM schemes improving the AJPS cryptosystem and its variant. They are related to the hardness of MERS problem. Since our constructions rely on newly introduced assumptions, further cryptanalytic efforts are demanded in order to get more confidence about their exact security.

10 Acknowledgments

The authors thank Waïss Azizian, Sarah Houdaigoui and Quốc Tín Lê for the development and the simulation of different backtracking strategies.

References

- AJPS17a. Divesh Aggarwal, Antoine Joux, Anupam Prakash, and Miklos Santha. A new public-key cryptosystem via Mersenne numbers. Cryptology ePrint Archive, Report 2017/481, 2017. <http://eprint.iacr.org/2017/481>.
- AJPS17b. Divesh Aggarwal, Antoine Joux, Anupam Prakash, and Mikos Santha. Mersenne-756839. Technical report, National Institute of Standards and Technology, 2017. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions>.
- AJPS18. Divesh Aggarwal, Antoine Joux, Anupam Prakash, and Miklos Santha. A new public-key cryptosystem via mersenne numbers. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018*, pages 459–482, Cham, 2018. Springer International Publishing.
- BCGN17. Marc Beunardeau, Aisling Connolly, Rémi Géraud, and David Naccache. On the hardness of the mersenne low hamming ratio assumption. Cryptology ePrint Archive, Report 2017/522, 2017. <http://eprint.iacr.org/2017/522>.
- dBDJdW17. Koen de Boer, Léo Ducas, Stacey Jeffery, and Ronald de Wolf. Attacks on the ajps mersenne-based cryptosystem. Cryptology ePrint Archive, Report 2017/1171, 2017. <https://eprint.iacr.org/2017/1171>.
- ElG85. Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In George Robert Blakley and David Chaum, editors, *Advances in Cryptology*, pages 10–18, Berlin, Heidelberg, 1985. Springer Berlin Heidelberg.
- HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. Ntru: A ring-based public key cryptosystem. In *Lecture Notes in Computer Science*, pages 267–288. Springer-Verlag, 1998.

- LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, pages 1–23, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- LPR13. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43:1–43:35, November 2013.
- MTSB13. R. Misoczki, J. Tillich, N. Sendrier, and P. S. L. M. Barreto. Mdpcceliece: New mceliece variants from moderate density parity-check codes. In *2013 IEEE International Symposium on Information Theory*, pages 2069–2073, July 2013.
- NIS17. NIST. Post quantum crypto project (2017). <http://csrc.nist.gov/groups/ST/post-quantum-crypto/>, May 2017. Accessed 19 May 2017.
- Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 84–93, New York, NY, USA, 2005. ACM.
- Reg06. Oded Regev. Lattice-based cryptography. In Cynthia Dwork, editor, *Advances in Cryptology - CRYPTO 2006*, pages 131–141, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- RSA78. R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, February 1978.
- Sho97. Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, October 1997.
- Tan. Fang Tang. Backtracking. <https://www.cpp.edu/~ftang/courses/CS331/notes/backtracking.pdf>.

A Appendix: A Broken Scheme Target for Repair

While designing the algorithms in this paper we broke the following variant that we mention here as a fixing target. Let R be a secret n -bit number of the form: $R \leftarrow \text{random}|v_\ell|\text{random}$ where v_ℓ is defined as in Pattern identification, and define the auxiliary public-key $L \leftarrow R/G \bmod p$. Note that this modifies the complexity assumption on which the scheme rests.

Encrypt by $C \leftarrow AH + B + Lm \bmod p$ and decrypt by $W \leftarrow GC = AF + GB + Rm \bmod p$. This offers a (noisy) visibility window on m allowing to extract and error-correct m .

The problem here is the equation defining L from which G can be revealed using LLL. It is unclear if this can be fixed using modifications in the encryption process and/or in parameter sizes. It is also interesting to determine if secure H -less variants¹⁷ can be designed.

Assuming that the above could be repaired, the following is for readers fond of dangerous games.

A dangerous game, that we do not recommend, reduces noise in the visibility window. If A, B, G, F are generated in a biased way by shifting more Hamming weight into the MSBs and the LSBs as shown in Figure 5, then the result of the multiplication modulo p of two such numbers results in a number of the form shown in Figure 6. Those densities are illustrated in Figures 8 and 9. This reduces the noise in the reading window and allows an easier recovery of m . It must be stressed that this increases the vulnerability of the public-key to the partition attacks of [BCGN17] as the attacker can better zoom on the information-rich part of F and G . This might be compensated by a higher h . Note that weight shifting does not necessarily need to be similar in all four variables A, B, F, G and that several weight shifting schemes are circularly equivalent because of the multiplication modulo p .

weight = $h/4 + \Delta$	weight = $h/2 - 2\Delta$	weight = $h/4 + \Delta$
----------------------------	-----------------------------	----------------------------

Figure 5: Unbalanced weight of A, B, F, G before multiplication.

¹⁷ i.e. where the sender encrypts by $C \leftarrow Lm + B \bmod p$ (this is insecure) or a similar trick.

weight \cong $h^2/4 + 2\Delta^2$	weight \cong $h^2/2 - 4\Delta^2$	weight \cong $h^2/4 + 2\Delta^2$
---------------------------------------	---------------------------------------	---------------------------------------

Figure 6: Unbalanced weight of $AF \bmod p$ and $BG \bmod p$ after multiplication.

weight \cong $h^2/2 + 4\Delta^2$	weight \cong $h^2 - 8\Delta^2$	weight \cong $h^2/2 + 4\Delta^2$
---------------------------------------	-------------------------------------	---------------------------------------

Figure 7: Unbalanced weight of $AF + BG \bmod p$ after addition.




Figure 8: Unbalanced A and F for $\{n, h, \Delta\} = \{2^{14}, 32, 6\}$.




Figure 9: Unbalanced $AF \bmod p$ for $\{n, h, \Delta\} = \{2^{14}, 32, 6\}$. Note the relatively lower dot density at the middle of the diagram.