Research Article

Jean-Sébastien Coron* and Agnese Gini

Improved cryptanalysis of the AJPS Mersenne based cryptosystem

https://doi.org/10.1515/jmc-2019-0027
Received July 13, 2019; accepted September 23, 2019

Abstract: At Crypto 2018, Aggarwal, Joux, Prakash and Santha (AJPS) described a new public-key encryption scheme based on Mersenne numbers. Shortly after the publication of the cryptosystem, Beunardeau et al. described an attack with complexity $\mathcal{O}(2^{2h})$. In this paper, we describe an improved attack with complexity $\mathcal{O}(2^{1.75h})$.

Keywords: AJPS cryptosystem, Lattice attack, LLL

MSC 2010: 94A60

1 Introduction

The AJPS public-key encryption scheme

At Crypto 2018, Aggarwal, Joux, Prakash and Santha (AJPS) described a new public-key encryption scheme based on arithmetic modulo Mersenne numbers [2]. A Mersenne prime is a prime integer p of the form $p = 2^n - 1$, where n is a prime. The arithmetic modulo p has good properties, and one can establish a correspondence between integers modulo p and binary strings of length n, up to $0^n \sim 1^n$. In particular, one can define the *Hamming weight* of a number as the Hamming weight of the unique binary string associated to it, i.e. the number of ones in its binary representation. In the earliest version of their work, the authors presented a public-key encryption scheme (AJPS-1) somewhat similar to the NTRU cryptosystem, but based on a new assumption, the *Mersenne low Hamming ratio assumption*. Its security relies on the following assumption: given $H = F/G \bmod p$, where the binary representation of F and G modulo p has low Hamming weight, then H looks pseudorandom; namely, it is hard to distinguish H from a random integer modulo p.

The Beunardeau et al. attack

Even though the authors claimed that the known lattice attacks against NTRU would not apply, very soon, Beunardeau et al. [3] described a lattice-based attack against the first AJPS proposal. The attack complexity is $\mathcal{O}(2^{2h})$, where h is the Hamming weight of F and G. The attack was further analyzed in [4]; the authors also described a meet-in-the-middle attack against AJPS-1 based on locality-sensitive hash functions to obtain collisions; they showed that the lattice attack from [3] is more efficient.

Since AJPS-1 allows to encrypt only a single bit at a time, it is not very efficient. However, in a later version of the article, published at Crypto 2018 [2], Aggarwal et al. described a variant (AJPS-2) that encrypts many bits at a time, with much larger security parameters to prevent the lattice attack.

***Corresponding author: Jean-Sébastien Coron,** University of Luxembourg, Esch-sur-Alzette, Luxembourg, e-mail: jean-sebastien.coron@uni.lu. http://orcid.org/0000-0003-1021-3344
Agnese Gini, University of Luxembourg, Esch-sur-Alzette, Luxembourg, e-mail: agnese.gini@uni.lu

Our contribution

In this paper, we describe a variant of the Beunardeau et al. attack against AJPS-2, with improved complexity $\mathcal{O}(2^{1.75h})$ instead of $\mathcal{O}(2^{2h})$. Instead of recovering the private key, our attack only breaks the indistinguishability of ciphertexts.

2 The AJPS cryptosystems

In this section, we recall the two versions of the AJPS cryptosystems; see [2] for further details.

AJPS-1: bit-by-bit encryption

Let $p = 2^n - 1$ be a Mersenne prime, where n itself is prime. Let h be an integer. Let F and G be two random integers modulo p with Hamming weight h such that $4h^2 < n \leq 16h^2$. Then the public key is $\mathrm{pk} = H = F/G \bmod p$ and the private key is $\mathrm{sk} = G$. To encrypt, choose two random integers A and B of Hamming weight h. Encrypt the bit b as

$$C = (-1)^b \cdot (A \cdot H + B).$$

To decrypt, compute $d = \mathrm{Ham}(C \cdot G)$. Output 0 if $d \leq 2h^2$; otherwise, output 1.

Decryption works because

$$C \cdot G = (-1)^b \cdot (A \cdot H \cdot G + B \cdot G) = (-1)^b \cdot (A \cdot F + B \cdot G)$$

which has Hamming weight at most $2h^2$ if $b = 0$, and at least $n - 2h^2$ if $b = 1$. Namely, for any number x of Hamming weight h, the integer $x \cdot 2^z \bmod p$ for $z \geq 0$ is a cyclic shift of x, and therefore its Hamming weight remains unchanged. Therefore, the Hamming weight of $A \cdot F$ is at most h^2, and the Hamming weight of $B \cdot G$ is also at most h^2; therefore, the Hamming weight of $C \cdot G$ is at most $2h^2$ for $b = 0$.

AJPS-2: error correcting codes

Let n be a positive integer such that $p = 2^n - 1$ be a Mersenne prime. Let $h \in \mathbb{N}$ be such that $10h^2 < n \leq 16h^2$. Let F, G be two random integers modulo p with Hamming weight h, and let R be a random integer modulo p. Set

$$\mathrm{pk} = (R, F \cdot R + G) = (R, T) \quad \text{and} \quad \mathrm{sk} = F.$$

To encrypt a message $m \in \{0,1\}^h$, first generate three random integers A, B_1, B_2 modulo p, with Hamming weight h. Then, using the encoding algorithm $\mathcal{E}\colon \{0,1\}^h \to \{0,1\}^n$ of an error correcting code $(\mathcal{E}, \mathcal{D})$, compute the ciphertext

$$(C_1, C_2) = (A \cdot R + B_1, (A \cdot T + B_2) \oplus \mathcal{E}(m)).$$

To decrypt, compute $\mathcal{D}((F \cdot C_1) \oplus C_2)$, where $\mathcal{D}$ is the corresponding decoding algorithm.

Decryption works because

$$\begin{aligned} F \cdot C_1 &= A \cdot F \cdot R + F \cdot B_1 = A \cdot (T - G) + F \cdot B_1 \\ &= (A \cdot T + B_2) - A \cdot G - B_2 + B_1 \cdot F, \end{aligned}$$

and therefore the Hamming distance between $A \cdot T + B_2$ and $F \cdot C_1$ is expected to be low, which enables to recover m with good probability.

3 The Beunardeau et al. attack

Basic attack

Beunardeau et al. described an attack against AJPS-1 in [3] that recovers the private key from the public key. More precisely, they consider the following problem.

Definition 3.1 (Mersenne low Hamming ratio search problem (MLHSP)). Let $p = 2^n - 1$ be an n-bit Mersenne prime and h an integer. Let F, G be two n-bit random strings with Hamming weight h. Given $H = F/G \bmod p$, recover F and G.

Their basic attack is based on the following observation. With probability 2^{-2h}, we have both $F < \sqrt{p}$ and $G < \sqrt{p}$, and therefore, given $H = F/G \bmod p$, one can recover F and G by applying LLL in dimension 2. In the original proposal [1], it was recommended to take $h = 17$ for $\lambda = 120$ bits of security. However, here we have an attack that recovers the private key from the public key with probability 2^{-34}; see also [4] for a detailed analysis.

More precisely, one considers the lattice $\mathcal{L}$ generated by the rows of the matrix

$$\begin{bmatrix} 1 & H \\ 0 & p \end{bmatrix}.$$

We have that $\det \mathcal{L} = p$; hence, by the Gaussian heuristic, it contains a vector of norm $\simeq (\det \mathcal{L})^{\frac{1}{2}} = \sqrt{p}$. Moreover, (G, F) is a short vector of the lattice. Therefore, if both $F < \sqrt{p}$ and $G < \sqrt{p}$, we can recover F and G; since F and G have Hamming weight h, this happens with probability 2^{-2h}.

We note that a similar attack can also be applied to the encryption equation $C = (-1)^b \cdot (A \cdot H + B)$. Namely, if both $A < \sqrt{p}$ and $B < \sqrt{p}$, then we can recover A and B by applying LLL in dimension 3, hence the plaintext bit b. Indeed, we have that only one between (H, C) and $(-H, C)$ is an instance of the following problem.

Definition 3.2 (Mersenne low Hamming combination search problem (MLHCSP)). Let $p = 2^n - 1$ be an n-bit Mersenne prime, h an integer, R a uniformly random n-bit string, and let F, G have Hamming weight h. Given the pair $(R, F \cdot R + G \bmod p)$, find F, G.

Given R and $T = F \cdot R + G \bmod p$, a variant attack recovers F, G with probability 2^{-2h}. More precisely, the attack works by considering the lattice $\mathcal{L}$ of row vectors

$$\begin{bmatrix} 2^{\frac{n}{2}} & 0 & T \\ 0 & 1 & -R \\ 0 & 0 & p \end{bmatrix}.$$

We have that $(2^{\frac{n}{2}}, F, G)$ belongs to the lattice $\mathcal{L}$. Moreover, $\det \mathcal{L} = 2^{\frac{n}{2}} p \simeq 2^{\frac{3n}{2}}$. Hence, by the Gaussian heuristic, the lattice $\mathcal{L}$ contains a vector of norm $\simeq 2^{\frac{n}{2}}$. Therefore, if both $F < \sqrt{p}$ and $G < \sqrt{p}$, we can recover F and G by applying LLL to the lattice $\mathcal{L}$.

Extension with random partitions

The basic attack from [3] is only a weak-key attack that recovers the private key from the public key with probability 2^{-2h} over the set of possible public keys. Similarly, the above variant attack against the encryption equation can only decrypt a fraction 2^{-2h} of the ciphertexts. Therefore, the authors extended their attack by considering random partitions, with higher-dimensional lattices. In that case, the attack can recover the private key from any public key, solving MLHSP, with complexity $\mathcal{O}(2^{2h})$. The same partition strategy can be used for the MLHCSP with the same complexity. In our improved attack in the next section, we will also use random partitions.

4 Our new attack

We describe our new attack against AJPS-2. We consider the previous encryption equation

$$(C_1, C_2) = (A \cdot R + B_1, (A \cdot T + B_2) \oplus \mathcal{E}(m))$$

Given the public key (R, T) and a ciphertext (C_1, C_2), our attack can distinguish between $m = 0$ and $m \neq 0$. Assume that $m = 0$ and $\mathcal{E}(m) = 0$. In that case, we have

$$C_1 = A \cdot R + B_1,$$
$$C_2 = A \cdot T + B_2.$$

We claim that if A, B_1 and B_2 are less than $2^{\frac{2n}{3}}$, then we can recover A, B_1 and B_2 with LLL. Namely, we consider the lattice of row vectors

$$\begin{bmatrix} 2^{\frac{2}{3}n} & 0 & C_1 & C_2 \\ 0 & 1 & -R & -T \\ 0 & 0 & p & 0 \\ 0 & 0 & 0 & p \end{bmatrix}.$$

We have that $\det \mathcal{L} = 2^{\frac{2n}{3}} p^2 \simeq 2^{\frac{8n}{3}}$. Therefore, by the Gaussian heuristic, the lattice $\mathcal{L}$ contains vectors of norm $\simeq 2^{\frac{2n}{3}}$. Moreover, the lattice $\mathcal{L}$ contains the vector $(2^{\frac{2n}{3}}, A, B_1, B_2)$. Therefore if A, B_1 and B_2 are less than $2^{\frac{2n}{3}}$, we can recover A, B_1 and B_2 by applying LLL to $\mathcal{L}$.

Since A has Hamming weight h, the probability that $A < 2^{\frac{2n}{3}}$ is $\left(\frac{2}{3}\right)^h$; the same holds for B_1 and B_2. The success probability of the attack is therefore

$$\left(\frac{2}{3}\right)^{3h} \simeq 2^{-1.75h}$$

which gives a slightly better success probability than the original attack with 2^{-2h}. Therefore, using the same partition technique as in [3], the attack complexity to break the indistinguishability of any ciphertext is $\mathcal{O}(2^{1.75h})$ instead of $\mathcal{O}(2^{2h})$.

4.1 Working with random partitions

We show that, using the same random partition technique as in [3], we can break the indistinguishability property of *any* ciphertext (C_1, C_2), whereas the basic attack above only works when A, B_1 and B_2 are less than $2^{\frac{2n}{3}}$, which only happens with probability $\left(\frac{2}{3}\right)^{3h}$.

We consider the set $[n] = \{0, 1, \ldots, n-1\}$. We say that $P = \{P_i\}_{i=1}^k$ is an *interval-like partition* if it is a partition of $[n]$ such that the sets are of the form $P_i = \{y : c \leq y \leq d\}$ or $P_i = \{d, d+1, \ldots, 0, \ldots, c-1, c\}$ for $c \leq d \in [n]$. We define p_i as the least element of P_i, namely as c if the interval is of the first type and as d if it is of the second type. We can use a partition to represent a number E modulo p by a sequence of smaller integers. More precisely, letting $e_{n-1} \cdots e_0$ be the binary representation of e, we can divide it by the partition

$$e_{p_1-1} \cdots e_{p_k} \mid e_{p_k-1} \cdots e_{p_{k-1}} \mid \ldots \mid e_{p_2-1} \cdots e_{p_1},$$

and letting d_i the number represented by $e_{p_i-1} \cdots e_{p_{i-1}}$, we obtain

$$E = \sum_{i=1}^{k} d_i \cdot 2^{p_i}.$$

Consider P, Q, S three interval-like partitions of $[n]$ of cardinality k, ℓ and j, respectively. Let R, T, C_1, C_2, A, B_1, B_2 be as in AJPS-2. We define a family of embedded lattices parameterized with respect to β, P, Q, S as

$$\mathcal{L}_{\beta,P,Q,S} = \left\{ (\alpha\beta, \mathbf{x}, \mathbf{y}, \mathbf{z}) \in \mathbb{Z} \times \mathbb{Z}^k \times \mathbb{Z}^\ell \times \mathbb{Z}^j : \alpha \cdot C_1 \equiv R \cdot \sum_{i=1}^{k} x_i \cdot 2^{p_i} + \sum_{i=1}^{\ell} y_i \cdot 2^{q_i} \bmod p, \right.$$
$$\left. \alpha \cdot C_2 \equiv T \cdot \sum_{i=1}^{k} x_i \cdot 2^{p_i} + \sum_{i=1}^{j} z_i \cdot 2^{s_i} \bmod p \right\}$$

for some *scaling factor* $\beta \in \mathbb{Z}$. The dimension of $\mathcal{L}_{\beta,P,Q,S}$ is $d = k + \ell + j + 1$, and a basis of this lattice is given by rows of the matrix

$$M_{\beta,P,Q,S} = \left[\begin{array}{c|ccccc|ccc|c|ccc|c} \beta & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & C_1 \cdot 2^{-q_1} & 0 & \cdots & 0 & C_2 \cdot 2^{-s_1} \\ \hline 0 & 1 & 0 & \cdots & 0 & 0 & \cdots & 0 & -R \cdot 2^{p_k - q_1} & 0 & \cdots & 0 & -T \cdot 2^{p_k - s_1} \\ 0 & 0 & 1 & \cdots & 0 & 0 & \cdots & 0 & -R \cdot 2^{p_{k-1} - q_1} & 0 & \cdots & 0 & -T \cdot 2^{p_{k-1} - s_1} \\ 0 & & & \ddots & & 0 & \cdots & 0 & -R \cdot 2^{p_2 - q_1} & 0 & \cdots & 0 & -T \cdot 2^{p_2 - s_1} \\ 0 & 0 & 0 & \cdots & 1 & 0 & \cdots & 0 & -R \cdot 2^{p_1 - q_1} & 0 & \cdots & 0 & -T \cdot 2^{p_1 - s_1} \\ \hline 0 & 0 & 0 & \cdots & 0 & 1 & \cdots & 0 & -2^{q_\ell - q_1} & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & \cdots & 0 & & \ddots & & -2^{q_i - q_1} & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 1 & -2^{q_2 - q_1} & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & p & 0 & \cdots & 0 & 0 \\ \hline 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & 1 & \cdots & 0 & -2^{s_j - s_1} \\ 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & & \ddots & & -2^{s_i - s_1} \\ 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 1 & -2^{s_2 - s_1} \\ 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & p \end{array}\right].$$

We claim that we can recover A, B_1, B_2 by using a lattice of the family $\{\mathcal{L}_{\beta,P,Q,S}\}$. We define the *secret vector* to be

$$\mathbf{s} := (\beta, a_1, \ldots, a_k, b_1^{(1)}, \ldots, b_\ell^{(1)}, b_1^{(2)}, \ldots, b_j^{(2)}) \in \mathcal{L}_{\beta,P,Q,S},$$

where $0 \le a_i < 2^{p_i}$, $0 \le b_i^{(1)} < 2^{q_i}$, $0 \le b_i^{(2)} < 2^{s_i}$ and

$$A = \sum_{i=1}^{k} a_i \cdot 2^{p_i}, \quad B_1 = \sum_{i=1}^{\ell} b_i^{(1)} \cdot 2^{q_i}, \quad B_2 = \sum_{i=1}^{j} b_i^{(2)} \cdot 2^{s_i}.$$

We will use the notations $a = (a_1, \ldots, a_k)$, $b^{(1)} = (b_1^{(1)}, \ldots, b_\ell^{(1)})$, $b^{(2)} = (b_1^{(2)}, \ldots, b_j^{(2)})$, $\mathbf{e} = (a, b^{(1)}, b^{(2)})$ and $\mathbf{s} = (\beta, \mathbf{e})$.

In the following, we determine under which conditions the secret vector $\mathbf{s}$ is the unique shortest vector of the lattice $\mathcal{L}_{\beta,P,Q,S}$. Given A, B_1, B_2, we say that the triple (P, Q, S) of partitions of $[n]$ is a *lucky triple* if there exists a scaling factor $\beta \in \mathbb{N}$ such that the secret vector $\mathbf{s}$ is the unique shortest vector of $\mathcal{L}_{\beta,P,Q,S}$. In that case, $\mathcal{L}_{\beta,P,Q,S}$ will be said to be a *lucky lattice* respect to A, B_1, B_2. In other words, we aim to establish sufficient conditions under which a lattice $\mathcal{L}_{\beta,P,Q,S}$ is lucky given a ciphertext $C = (C_1, C_2)$ such that $\mathcal{E}(m) = 0$.

The volume of $\mathcal{L}_{\beta,P,Q,S}$ is

$$\mathrm{vol}(\mathcal{L}_{\beta,P,Q,S}) = |\det(M)| = p^2 \cdot \beta.$$

We write $\beta = 2^{tn}$; thus we have $\mathrm{vol}(\mathcal{L}_{\beta,P,Q,S}) \simeq 2^{(2+t)n}$. By the Gaussian heuristic, we obtain the following estimate of the length of the shortest vector of $\mathcal{L}_{\beta,P,Q,S}$:

$$\sqrt{\frac{d}{2\pi e}} \cdot \mathrm{vol}(\mathcal{L}_{\beta,P,Q,S})^{\frac{1}{d}} = \sqrt{\frac{d}{2\pi e}} \cdot 2^{\frac{(2+t)n}{d}}. \tag{4.1}$$

Since the Hamming weight of A, B_1, B_2 is the same, we take $k = j = \ell$. We note that the lattice $\mathcal{L}_{\beta,P,Q,S}$ contains intrinsic short vectors $\mathbf{u} = (0, \ldots, 0, 2^g, -1, 0, \ldots, 0)$ whose norm is $\simeq 2^g$ when g is of the form $p_i - p_{i-1}$ or $q_i - q_{i-1}$ or $s_i - s_{i-1}$. If we consider partitions with intervals of similar length, we obtain $\|\mathbf{u}\| \approx 2^{\frac{n}{k}}$. Therefore, we have to ensure that such vectors are not shorter than our target secret vector.

In low dimensions, we can assume that LLL recovers the shortest vector $\mathbf{s}$ of the lattice. From (4.1), we must therefore ensure

$$\|\mathbf{s}\| \le \sqrt{\frac{d}{2\pi e}} \cdot 2^{\frac{(2+t)n}{d}},$$

where $d = 3k + 1$ is the lattice dimension. We expect the entries of the secret vector to be about of the same

h	n	$\log_2(\bar{y})$	$\log_2(\bar{Y})$
3	127	6.5	7.4
6	521	13.0	14.5
7	607	14.6	16.5
9	1279	14.9	16.4

Table 1: Average number $\bar{y}$ of partitions required to recover the secret values A, B_1, B_2, compared to the average number $\bar{Y}$ required for the original attack. We used 70 samples for $h = 3, 6, 7$ and 9 samples for $h = 9$.

size for a lucky triple; hence we take the scaling factor β such that $\beta = 2^{tn} \simeq \|\mathbf{e}\|$. Then we have approximately

$$2^{tn+\frac{1}{2}} \leq 2^{\frac{(2+t)n}{3k+1}}$$

which gives $t \leq \frac{2}{3k} - \frac{3k+1}{6kn}$. Therefore, we have the approximative condition to have a lucky triple (P, Q, S) of partitions

$$\|\mathbf{e}\| < 2^{\frac{2n}{3k}}. \tag{4.2}$$

It remains to evaluate the probability to find a lucky triple of partitions (P, Q, S). It is actually easier to assume that the partitions (P, Q, S) are fixed and the ciphertext $C = (C_1, C_2)$ is random. In that case, from the bound (4.2), each of the h bits from the integers A, B_1 and B_2 must land in one of the subintervals of length $\frac{2n}{3k}$ of the k partition intervals. For a single bit, this happens with probability roughly $k \cdot \frac{2n}{3k} \cdot \frac{1}{n} = \frac{2}{3}$. Therefore, as in the basic attack, the success probability is roughly $\left(\frac{2}{3}\right)^{3h} \simeq 2^{-1.75h}$. Therefore, the number of partitions to try before finding a lucky one is approximately $\mathcal{O}(2^{1.75h})$ instead of $\mathcal{O}(2^{2h})$ in the original attack from [3].

Security parameter selection

In the latest version of the paper, the authors recommended to take for λ bit of security $h = \lambda$ in order to prevent possible improvements of the Beunardeau et al. attack. Then our attack does not affect the choice of parameter proposed in [2].

4.2 Practical experiments

We have performed some practical experiments for various values of bitsize n and Hamming weight h of AJPS-2 in order to compare our new attack with the original Beunardeau et al. attack. For both attacks, since we do not know a priori the optimal size of the partition k to recover the secret, we perform a repeated loop over all possible $1 \leq k \leq h$. We summarize our results in Table 1, showing that our attack indeed requires fewer partitions than the original attack.

References

[1] D. Aggarwal, A. Joux, A. Prakash and M. Santha, A new public-key cryptosystem via Mersenne numbers, Cryptology ePrint Archive (2017), https://eprint.iacr.org/2017/481.

[2] D. Aggarwal, A. Joux, A. Prakash and M. Santha, A new public-key cryptosystem via Mersenne numbers, in: *Advances in Cryptology—CRYPTO 2018*, Lecture Notes in Comput. Sci. 10993, Springer, Berlin (2018), 459–482.

[3] M. Beunardeau, A. Connolly, R. Géraud and D. Naccache, On the hardness of the Mersenne low Hamming ratio assumption, Cryptology ePrint Archive (2017), https://eprint.iacr.org/2017/522.

[4] K. de Boer, L. Ducas, S. Jeffery and R. de Wolf, Attacks on the AJPS Mersenne-based cryptosystem, in: *Post-Quantum Cryptography—PQCrypto 2018*, Lecture Notes in Comput. Sci. 10786, Springer, Cham (2018), 101–120.