Public-key Cryptography from Different Assumptions
A Multi-bit Version

Hervé Chabanne[1,2], Gérard Cohen[1] and Alain Patey[1,2]
[1]*Télécom ParisTech, Paris, France*
[2]*Morpho, Issy-Les-Moulineaux, France*
Identity and Security Alliance (The Morpho and Télécom ParisTech Research Center)

Keywords: Public-key Cryptography, Homomorphic Encryption, Wire-Tap Channel.

Abstract: At STOC 2010, Applebaum, Barak and Wigderson introduced three new public-key cryptosystems based on combinatorial assumptions. In their paper, only encryption of bits has been considered. In this paper, we focus on one of their schemes and adapt it to encrypt a constant number of bits in a single ciphertext without changing the size of the public key. We add wire-tap channel techniques to improve the security level of our scheme, thus reaching indistinguishability. We show that it is homomorphic for the XOR operation on bit strings. We also suggest concrete parameters for a first instantiation of our scheme.

1 INTRODUCTION

Most public-key cryptosystems that have been proposed since the introduction of public-key cryptography (Diffie and Hellman, 1976) rely on hardness assumptions from number theory, lattices or error-correcting codes, e.g. (Rivest et al., 1978b; McEliece, 1978; Ajtai and Dwork, 1997). In (Applebaum et al., 2010a), Applebaum et al. introduce three new public-key cryptosystems based on combinatorial problems, following other works such as (Goldreich et al., 1988; Blum et al., 1993; Juels and Peinado, 2000; Achlioptas and Coja-Oghlan, 2008). Their three cryptosystems share the same simple encryption and decryption mechanisms, described in Figure 1, but the hardness assumptions and, consequently, the key generation algorithms differ.

We focus in this paper on the scheme of (Applebaum et al., 2010a) based on the 3LIN assumption, that we denote by ABW in the following. This assumption states that it is infeasible to solve a random set of 3-sparse linear equations modulo 2, of which a small fraction is noisy. It is quite close to the Learning Parity with Noise problem but considers a noise level much higher than those used before.

The public key of this scheme is a 3-sparse matrix M and the corresponding secret key is a small subset of the lines of this matrix that sum up to 0. To encrypt a bit $\mathfrak{m}$, the public key matrix is multiplied by a random vector x and noise e is added to the product $M \cdot x$, then the first bit of $M \cdot x + e$ is flipped if $\mathfrak{m} = 1$ and unmodified otherwise. To decrypt a message, the lines of the ciphertext corresponding to the secret key are added modulo 2, and the result is, with high probability, the original message. The principal components of the cryptosystem are summed up in Fig. 1.

We show that this scheme can be used to encrypt a constant number of bits, instead of a single bit, while keeping the same size for the parameters. In the ABW cryptosystem, one set of lines that sums up to 0 is used as secret key. To encrypt l bits, we suggest to use l such subsets of lines that sum up to 0 in the public key, each one having the same size as the subset of lines in the ABW scheme. Moreover, the ABW cryptosystem only guarantees a weak notion of privacy: the distributions of ciphertexts encrypting respectively 0 and 1 are at a statistical distance lower than 1/2. To improve the security level, we encode the message to be encrypted using coset coding techniques. Indeed, we draw a parallel between an adversary against our cryptosystem and an eavesdropper in the Wire-Tap channel model (Wyner, 1975), and security in the wire-tap channel can be ensured using coset coding techniques. The resulting scheme is described in Fig. 4. We prove that this scheme achieves indistinguishability under chosen-plaintext attack.

We propose parameters for a concrete instantiation of the protocol. They are derived from the security analysis of (Applebaum et al., 2010a; Applebaum et al., 2010b) and from our proofs. We estimate

Chabanne H., Cohen G. and Patey A..
Public-key Cryptography from Different Assumptions - A Multi-bit Version.
DOI: 10.5220/0004600205610567
In *Proceedings of the 10th International Conference on Security and Cryptography* (SECRYPT-2013), pages 561-567
ISBN: 978-989-8565-73-0

that we can achieve short-term security with a 5 GB public key and 18-element subsets in the secret key. We suggest to encrypt around 100 bits at the same time. The underlying code used for coset coding is a $[128,30,29]$ BCH code.

1.1 Homomorphic Properties

In 1978, Rivest et al. (Rivest et al., 1978a) introduced the notion of *homomorphic encryption*. A public-key cryptosystem is homomorphic for an operation $\bullet$ if, given any two ciphertexts c_1 and c_2 encrypting (resp.) m_1 and m_2, it is possible to compute a ciphertext c_3 encrypting $m_3 = m_1 \bullet m_2$, without knowing the secret key. An encryption scheme that is homomorphic for any operation is called a *fully homomorphic* scheme.

For instance, textbook RSA (Rivest et al., 1978b) and ElGamal (Gamal, 1984) cryptosystems are multiplicatively homomorphic: the product of two ciphertexts is a ciphertext of the product of the plaintexts. The most known additively homomorphic encryption scheme is the one of Paillier(Paillier, 1999): the product of two ciphertexts is a ciphertext of the sum of the plaintexts. A fully homomorphic cryptosystem has been first proposed in 2009 by Gentry (Gentry, 2009). Since then, intensive research is pursued in this domain to gain efficiency.

It is interesting to notice that the scheme of (Applebaum et al., 2010a) has been used, with modifications inspired by (McEliece, 1978), in an unsuccessful attempt to build a fully homomorphic cryptosystem based on codes (Bogdanov and Lee, 2011). It has been broken differently and independently by (Gauthier et al., 2012) and (Brakerski, 2012)

We here focus on the homomorphism w.r.t. the exclusive-OR (XOR) operation. The scheme of Goldwasser and Micali (Goldwasser and Micali, 1982), based on the quadratic residuosity problem, enables to homomorphically compute the XOR operation by multiplying ciphertexts. However, only one bit at a time can be encrypted, and thus it does not achieve "XOR-ly" homomorphism on bit strings.

The McEliece cryptosystem (McEliece, 1978) can be adapted to achieve XOR-ly homomorphism (see (Strenzke, 2011, Section 2.4)). The number of XOR's that can be performed is however limited, the scheme is in a sense "somewhat XOR-ly homomorphic".

Our modified cryptosystem, thanks to the linearity of the encryption/decryption operations of the initial ABW scheme and to the linearity of (linear) coset coding, is homomorphic for the exclusive-OR operation on bit strings. This can have nice applications such as the secure computation of Hamming distances.

2 THE ABW CRYPTOSYSTEM

We focus on the first of the three schemes introduced in (Applebaum et al., 2010a), that is based on the 3LIN assumption. As already noticed, all three schemes follow the same mechanisms for encryption and decryption, only secret key generation differs. The approach of our paper to enable multi-bit encryption and achieve indistinguishability also applies to the second scheme based on both dLIN and DUE (Decisional Unbalanced Expansion) assumptions, though we do not detail it here. We believe that the third scheme is not suited for our multi-bit techniques.

In the following, a matrix $M \in \mathbb{F}_2^{m\times n}$ is said to be *d-sparse* if each of its rows contains exactly d ones. We denote by $\mathcal{M}_{m,n,d}$ the uniform distribution over d-sparse $m \times n$ matrices and by $\mathcal{T}_{p,n,d}$ the distribution over 3-sparse matrices with n columns, where each possible 3-sparse row (out of the $\binom{n}{3}$ possibilities) is picked with probability p and placed randomly into the matrix.

2.1 Security Assumptions

The first scheme of (Applebaum et al., 2010a) and our proposal both rely on the 3LIN assumption. We sum up definitions and assumptions concerning 3LIN that are used in (Applebaum et al., 2010a).

Definition 1 (ε-Satisfiable dLIN Instance)**.** *A dLIN instance with m-clauses and n variables is described by an m-bit vector b and a d-sparse matrix $M \in \mathbb{F}_2^{m\times n}$. It is ε-satisfiable if there exists an assignment x for which the weight of $Mx - b$ is at most εm.*

Definition 2 (Search3LIN Problem)**.** *The Search3LIN(m,ε) problem is defined as follows:*
Input: *A random ε-satisfiable 3LIN instance (M,b) sampled as follows:*

$M \leftarrow \mathcal{M}_{m,n}$ and $b = Mx + e$, where $x \in_R \{0,1\}^n$ and $e \leftarrow \mathrm{Ber}_\varepsilon^m$ (where $\mathrm{Ber}_\varepsilon^m$ is the distribution over m-bit vectors where bits are independently distributed following Bernoulli distributions with probability p).
Output: *The assignment x.*

We say that Search3LIN(m(n),ε(n)) is intractable if for every probabilistic polynomial-time algorithm A, and every sufficiently large n, A solves Search-3LIN(m(n),ε(n)) with probability smaller than 2/3.

Assumption 1. *The problem Search3LIN($C_O n^{1.4}, C_1 n^{-0.2}$)) is intractable for every constants $C_0 > 0$ and $C_1 > 0$.*

Public Key: A 3-sparse $m \times n$ matrix M
Secret Key: A subset S of $[1,\ldots,m]$ such that $1 \in S$ and $\Sigma_{i\in S} M_i = 0$ (where M_i denotes the i-th line of M)
Encryption($\mathfrak{m} \in \{0,1\}$): Pick a random vector $x \in_R \{0,1\}^n$ and a random error vector $e \in_R \mathrm{Ber}_\varepsilon^m$. Let $c = M \cdot x + e$. If $\mathfrak{m} = 1$, flip the first bit of c. Output c.
Decryption(c): Output $\Sigma_{i\in S} c_i$

Figure 1: The ABW cryptosystem.

2.2 The ABW Cryptosystem based on 3LIN

The architecture of the cryptosystem, shared with the other schemes of (Applebaum et al., 2010a), is described in Fig. 1. We here describe the key generation algorithm, that is specific to the cryptosystem that we focus on.

Key Generation(m,n,q). Sample a matrix H from $H_{q,n}^{2,3}$, the uniform distribution over matrices with q rows and n columns, where each row contains exactly 3 ones and each column contains either zero or two ones.

Sample a matrix M from $\mathcal{T}_{m/\binom{n}{3},n,3}$.

Pick a random set $S \subset \{1,\ldots,m\}$ such that $1 \in S$ and $|S| = q$. Replace the lines of M indexed by S by the lines of H.

The public key is M, the secret key is S.

2.3 Security of the ABW Cryptosystem

We say that two sequences of distributions X and Y are ε-indsitinguishable if, for every probabilistic polynomial time algorithm A, $|Pr[A(X) = 1] - Pr[A(Y) = 1]| < \varepsilon$.

Definition 3 (β-Privacy)**.** *A single-bit encryption scheme E, with public key pk, is β-private if the distributions $(pk, E_{pk}(0))$ and $(pk, E_{pk}(1))$ are β-indistinguishable.*

The following theorem summarizes (Applebaum et al., 2010a, Thm 2) and (Applebaum et al., 2010b, Thm 5.5).

Theorem 1 (Privacy of the ABW Cryptosystem)**.** *Let $q = \Theta(n^{0.2})$, $m = O(n^{1.4})$ and $\varepsilon = \Omega(n^{-0.2})$. If Search3LIN is intractable, the ABW cryptosystem is $(1-\delta/2)$-private, for some absolute constant $0 < \delta < 1$ that does not depend on n.*

3 THE WIRE-TAP CHANNEL

3.1 Linear Coset Coding

Coset coding is a random encoding technique. This type of encoding uses a $[n,k,d]$ linear code C with a parity-check matrix H. Let $r = n-k$. To encode a message $\mathfrak{m} \in \mathbb{F}_2^r$, one randomly chooses an element among all $x \in \mathbb{F}_2^n$ such that $\mathfrak{m} = H^t x$. To decode a codeword x, one just applies the parity-check matrix H and obtains the syndrome of x for the code C, which is the message $\mathfrak{m}$. This procedure is summed up in Fig. 2.

Parameters: C a $[n, n-r, d]$ linear code with a $r \times n$ parity-check matrix H
Encode: $\mathfrak{m} \in \mathbb{F}_2^r \mapsto_R x \in \mathbb{F}_2^n$ s.t. $H^t x = \mathfrak{m}$
Decode: $x \in \mathbb{F}_2^n \mapsto \mathfrak{m} = H^t x$

Figure 2: Linear Coset-coding

3.2 The Wire-Tap Channel

The Wire-Tap Channel was introduced by Wyner (Wyner, 1975). In this model, a sender Alice sends messages over a potentially noisy channel to a receiver Bob. An adversary Eve listens to an auxiliary channel, the *wire-tap channel*, which is a noisier version of the main channel. It was shown that, with an appropriate coding scheme, the secret message can be conveyed in such a way that Bob has complete knowledge of the secret and Eve does not learn anything. In the special case where the main channel is noiseless, the secrecy capacity can be achieved through a linear coset coding scheme. The model of the Wire-Tap Channel is drawn in Fig. 3.

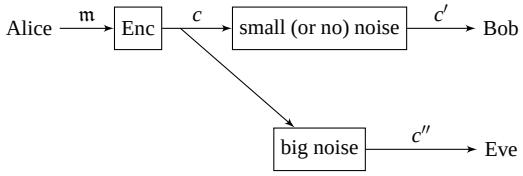

Figure 3: The Wire-Tap Channel.

Security is guaranteed if $\mathcal{H}(\mathfrak{m}|c') = 0$ and $\mathcal{H}(\mathfrak{m}|c'') = \mathcal{H}(\mathfrak{m})$, where $\mathcal{H}$ denotes entropy. Let us assume that the main channel is noiseless and that the wire-tap channel is a binary symmetric channel with probability of error equal to p. It has been proven (Wyner, 1975) that security in the Wire-Tap channel model can be achieved if the encoding technique is a linear coset coding achieving a rate at most $h(p) = -p\log(p) - (1-p)\log(1-p)$.

Public Key: A 3-sparse $m \times n$ matrix M and a $r \times l$ parity-check matrix H
Secret Key: l q-size subsets $S_1, \ldots, S_l$ of $[1, \ldots, m]$ such that $j \in S_j$ and $\Sigma_{i \in S_j} M_i = 0$ (where M_i denotes the i-th line of M)
Encryption($\mathfrak{m} \in \{0,1\}^r$): Pick a random vector $x \in_R \{0,1\}^n$, a random error vector $e \in_R \mathrm{Ber}_\varepsilon^m$, and a random vector $y \in \{0,1\}^l$ such that $H \cdot y = \mathfrak{m}$. Let $c = M \cdot x + e + (y_1, \ldots, y_l, 0, \ldots, 0)$. Output c.
Decryption(c): For $j = 1, \ldots, l$, let $y_j = \Sigma_{i \in S_j} c_i$. Output $\mathfrak{m} = H \cdot y$.

Figure 4: Our cryptosystem.

4 OUR PROPOSAL

4.1 Indistinguishability

Definition 4 (Indistinguishability)**.** *A (probabilistic) public-key cryptosystem E achieves indistinguishability under chosen-plaintext attack if, for any two messages m_0 and m_1, an adversary, given the public key pk, cannot distinguish between the distributions $E_{pk}(m_0)$ and $E_{pk}(m_1)$.*

4.2 The Cryptosystem

Our proposal for a multi-bit cryptosystem is described in Figure 4. We only need to specify the key generation algorithm.

Key Generation

Sample l matrices H_j from $H_{q,n}^{2,3}$, the uniform distribution over matrices with q rows and n columns, where each row contains exactly 3 ones and each column contains either zero or two ones.

Sample a matrix M from $\mathcal{T}_{m/\binom{n}{3},n,3}$.

Pick l random sets $S_j \subset \{1, \ldots, m\}$ such that $j \in S_j$ and $|S_j| = q$. Moreover, they must be such that $S_j \cap S_k = \emptyset$ for $j \neq k$.

For every $j \in \{1, \ldots, l\}$, replace the lines of M indexed by S_j by the lines of H_j.

The public key is M, the secret key is $\{S_j\}_{j \in \{1,\ldots,l\}}$.

Decryption Error

We recall that we insert in every ciphertext an error vector where errors occur following a Bernoulli distribution with low probability ε.

As in the ABW cryptosystem, it might happen, with low probability, that the error vector flips bits that are at positions indexed by the S_j's. If the number of flips in one S_j is odd, decryption of the j-th bit of x errs.

Each bit of the encoded message is erroneously decrypted with probability at most $\alpha = 1/2 - 1/2(1-2\varepsilon)^q < \varepsilon q$. Thus, the decryption of the message errs with probability at most $\beta = 1 - (1-\alpha)^l < \varepsilon q l$.

4.3 Proof of Security

We first prove that our distribution of public keys is close enough to $\mathcal{T}_{m/\binom{n}{3},n,3}$ and use arguments of (Applebaum et al., 2010a) to prove partial privacy of our cryptosystem. Then, we use the wire-tap model to prove indistinguishability of our full cryptosystem.

First, let us define the *j-th partial cryptosystem* as a single-bit cryptosystem, where the keys are generated in the same way as in our proposal, but where only one bit is encrypted; more precisely, the encryption and decryption algorithms are as follows:

Encryption($\mathfrak{m} \in \{0,1\}$). Pick a random vector $x \in_R \{0,1\}^n$ and a random error vector $e \in_R \mathrm{Ber}_\varepsilon^m$. Let $c = M \cdot x + e$. If $\mathfrak{m} = 1$, flip the j-th bit of c. Output c.

Decryption(c). Output $\Sigma_{i \in S_j} c_i$

Theorem 2. *Let $q = \Theta(n^{0.2})$, $m = O(n^{1.4})$ and $\varepsilon = \Omega(n^{-0.2})$. For $j = 1 \ldots l$, the j-th partial cryptosystem is $(1-\delta/2)$-private, where $0 < \delta < 1$ does not depend on n.*

Sketch. We rely on the proofs of (Applebaum et al., 2010a; Applebaum et al., 2010b) regarding the privacy of their first scheme, based on 3LIN. The privacy is proved by (Applebaum et al., 2010a, Thm 2) or (Applebaum et al., 2010b, Thm 5.5), which state that, if the distribution of public keys and $\mathcal{T}_{m/\binom{n}{3},n,3}$ are $(1-\delta)$-computationally indistinguishable, then encryption is $(1-\delta/2)$-private. The proof of this theorem is out of the scope of our paper, we only use its result.

We need to prove that our key distribution is $(1-\delta)$-computationally indistinguishable from $\mathcal{T}_{m/\binom{n}{3},n,3}$, *i.e.* we need an equivalent of (Applebaum et al., 2010a, Lemma 2) or (Applebaum et al., 2010b, Lemma 5.4), adapted to our new key distribution. The proof of (Applebaum et al., 2010b, Lemma 5.4) relies on arguments from (Feige et al., 2006). If the dimensions of the matrix and the size l of the private key are as in the hypotheses of Theorem 2, then each row of a 3-sparse matrix M sampled from $\mathcal{T}_{m/\binom{n}{3},n,3}$ has a probability $0 < \phi < 1$, independent of n, of belonging to a $q \times n$ sub-matrix where each column contains ei-

ther zero or two ones. (Such a sub-matrix is called an H-type sub-matrix in the following.)

Now let us assume that the first l rows all belong to H-type sub-matrices. We show that the probability that 2 (or more) of these rows belong to the same sub-matrix is negligible. Indeed, let us see the matrix M as a bipartite graph where the rows are on the left size, the H-type sub-matrices are on the right side and the edges connect the rows to the sub-matrices to which they belong. Let $q = \beta n^{0.2}$. With probability 1-$o(1)$, M contains at least $\alpha n^{1.4}$ H-type sub-matrices, each row of M participates in at most $\gamma n^{0.2}$ distinct H-type sub-matrices and M has at most $\theta n^{1.4}$ rows (see the proof of (Applebaum et al., 2010b, Lemma 5.4) and (Feige et al., 2006)). Let V be a right vertex. Let us consider the set of vertices that are at distance (smallest path) 4 from V. There are at most $\beta n^{0.2}$ (neighbours of V) $\times \gamma n^{0.2}$ (max. number of neighbours of a left vertex) $= \beta\gamma n^{0.4}$ vertices in this set. Since there are at least $\alpha n^{1.4}$ right vertices, the probability that two right vertices are at distance 4 is $\approx \frac{\beta\gamma}{\alpha n} = o(1)$. The probability that two (or more) out of the first l rows belong to the same H-type sub-matrix is thus negligible.

Consequently, with probability $\delta = \phi^l \in]0,1[$, independent of n, M is a possible public key. The statistical distance between $\mathcal{T}_{m/\binom{n}{3},n,3}$ and the uniform distribution over public keys is δ. Using (Applebaum et al., 2010a, Thm 2) or (Applebaum et al., 2010b, Thm 5.5), we obtain the $(1-\delta/2)$-privacy of the j-th partial cryptosystem, for every $1 < j < l$. □

Now we introduce linear coset coding as a way to enhance the security of the system and to prove indistinguishability

Theorem 3. *Assuming that all j-th partial cryptosystems are $(1-\delta/2)$-private, our cryptosystem using a linear coset coding with a $r/n = h(\delta/4)$-rate, achieves indistinguishability.*

Sketch. As proved by Theorem 2, the adversary can distinguish every bit of $Encode(m)$ (where $Encode$ means linear coset coding) with $1-\delta/2$ distinguishing advantage. The power of the adversary is thus equivalent to the power of an adversary in the wire-tap channel model where the wire-tap channel is a binary symmetric channel, with error probability p being at least $\delta/4$.

Thus, as in the wire-tap channel model, to prevent the adversary from learning information about the original message, one uses linear coset coding with a rate at most $h(\delta/4)$. The distributions of $Enc(m_0)$ and $E(m_1)$ are consequently indistinguishable, for any two messages m_0 and m_1 in $\mathbb{F}_2^r$. □

5 "XOR-ly" HOMOMORPHIC ENCRYPTION

5.1 Homomorphic Properties

The schemes of (Applebaum et al., 2010a) and the adaptations that we introduce are XOR-ly homomorphic. Indeed, let m_1 and m_2 be two r-bit messages and $y_1 = (Encode(m_1)||0\ldots0)$, $y_2 = (Encode(m_2)||0\ldots0)$. Let $c_1 = M.x_1 \oplus y_1 \oplus e_1$ and $c_2 = M.x_2 \oplus y_2 \oplus e_2$ be ciphertexts encrypting them. Then, we have $c_1 \oplus c_2 = M.(x_1 \oplus x_2) \oplus (y_1 \oplus y_2) \oplus e_1 \oplus e_2 = M.(x_1 \oplus x_2) \oplus (Encode(m_1 \oplus m_2)||0,\ldots,0) \oplus e_1 \oplus e_2$, due to linearity of coset coding. Thus $c_1 \oplus c_2$ encrypts $m_1 \oplus m_2$. We can see that the error term weight grows at every homomorphic XOR. Consequently, the probability of an erroneous decryption also grows with the number of XOR's.

5.2 Application to Secure Hamming Distance Computation

We want to homomorphically compute the Hamming distance between two c-bit strings X and Y, using the encryptions of X and Y. We first notice that obtaining the encryption of a random message m whose Hamming weight is the Hamming distance between X and Y is equivalent to obtaining an encryption of the Hamming distance. We use this trick to homomorphically compute the Hamming distance between X and Y

We know, from the previous section, that one easily obtains $E(X \oplus Y)$ from $E(X)$ and $E(Y)$ using the homomorphic properties of the cryptosystem. Now, if one randomly permutes the first c bits of the ciphertext, one obtains an encryption of the message $m = \sigma(X \oplus Y)$, where σ is the permutation. This message has the same Hamming weight as $X \oplus Y$, *i.e.* the Hamming distance between X and Y, but gives no more information about X and Y. This technique can for instance be used in the following applications:

Secure 2-Party Computation of Hamming Distance. Two parties P_1 and P_2 respectively hold inputs X and Y. P_1 sets a key pair sk, pk for one of the cryptosystems described in Section 4. P_2 is given pk and an encryption of X. P_2 can then homomorphically compute an encryption E of $X \oplus Y$. P_2 picks a random permutation σ and permutes the first c bits of the ciphertext E, he thus obtains E'. P_2 sends E' to P_1 who can decrypt it using sk. The Hamming weight of the decrypted result is the Hamming distance between X and Y. P_2 learns nothing about X and P_1 learns nothing more about Y than the Hamming distance between X and Y.

Secure Outsourcing of Hamming Distance Computation. Several data are stored encrypted on an external server (such as a cloud). The server holds $E(X)$ and $E(Y)$ and can homomorphically compute $E(X \oplus Y)$ then $E(\sigma(X \oplus Y))$, where σ is a randomly chosen permutation and sends it to the client holding the secret key. The client can then decrypt and retrieve the Hamming distance between X and Y without the server learning anything about the data involved

6 TOWARDS A CONCRETE IMPLEMENTATION

We here give a proposal for concrete parameters for our cryptosystem that would enable to achieve short-term security, *i.e.* the equivalent of a 80-bit symmetric encryption. We first recall the relations between parameters, as required by the security of proofs of Section 4.3 and of (Applebaum et al., 2010a; Applebaum et al., 2010b).

We assume that the privacy of the partial cryptosystems (see Section 4.3) is close to 1/2. Consequently, for coset coding, we require a code whose rate is approximately $h(1/4) \approx 0.81$. Since we consider linear binary codes, we suggest to employ BCH codes, but others could be used, as long as their rate is close to $h(1/4)$.

Finally, we would like to avoid naive attacks to recover the keys. Since all S_j sets are independent, we require S_j to be more than 80-bit sized, *i.e.* we require $q\log(m) > 80$. Moreover we would like to avoid a brute-force attack where the adversary picks every $(q/2)$-tuple of public key rows and then looks for $q/2$ other rows such that all these q rows sum up to 0. This attack requires at least $(3m/n)^{q/2}$ operations. We consequently require $(3m/n)^{q/2} > 2^{80}$.

If we combine everything together, we suggest to take $n = 2^{21}$, $m = 2^{29}$, $r = 98$, $q = 18$, $l = 128$, $\varepsilon = 10^{-6}$and a linear coset coding whose underlying code is a $[128, 30, 29]$ BCH code. This leads to a 5 GB public key and to a 128×540-bit private key.

We do not claim here to get a ready-to-use cryptosystem. Following [3] and their will to introduce new public-key cryptosystems relying on combinatorial assumptions despite the fact that they are not as efficient today as the more classical ones. Our work can be interpreted as a first step towards more practical implementation. Our suggestions for concrete parameters should therefore be seen as a challenge. We strongly support the idea of cryptanalysis of our scheme using these parameters, or evidences that we could reduce the parameters' size without impact on security.

ACKNOWLEDGEMENTS

The authors would like to thank Benny Applebaum, Boaz Barak and Avi Wigderson for their helpful comments. This work has been partially funded by the ANR SecuLar project.

REFERENCES

Achlioptas, D. and Coja-Oghlan, A. (2008). Algorithmic barriers from phase transitions. In *FOCS*, pages 793–802.

Ajtai, M. and Dwork, C. (1997). A public-key cryptosystem with worst-case/average-case equivalence. In *STOC*, pages 284–293.

Applebaum, B., Barak, B., and Wigderson, A. (2010a). Public-key cryptography from different assumptions. In *STOC*, pages 171–180.

Applebaum, B., Barak, B., and Wigderson, A. (2010b). Public-key cryptography from different assumptions (extended version). http://www.cs.princeton.edu/~boaz/Papers/ncpkcFull2.pdf.

Blum, A., Furst, M. L., Kearns, M. J., and Lipton, R. J. (1993). Cryptographic primitives based on hard learning problems. In *CRYPTO*, pages 278–291.

Bogdanov, A. and Lee, C. H. (2011). Homomorphic encryption from codes. *IACR Cryptology ePrint Archive*, 2011:622.

Brakerski, Z. (2012). When homomorphism becomes a liability. *IACR Cryptology ePrint Archive*, 2012:225.

Diffie, W. and Hellman, M. E. (1976). New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654.

Feige, U., Kim, J. H., and Ofek, E. (2006). Witnesses for non-satisfiability of dense random 3cnf formulas. In *FOCS*, pages 497–508.

Gamal, T. E. (1984). A public key cryptosystem and a signature scheme based on discrete logarithms. In *CRYPTO*, pages 10–18.

Gauthier, V., Otmani, A., and Tillich, J.-P. (2012). A distinguisher-based attack of a homomorphic encryption scheme relying on reed-solomon codes. *IACR Cryptology ePrint Archive*, 2012:168.

Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178.

Goldreich, O., Krawczyk, H., and Luby, M. (1988). On the existence of pseudorandom generators. In *CRYPTO*, pages 146–162.

Goldwasser, S. and Micali, S. (1982). Probabilistic encryption and how to play mental poker keeping secret all partial information. In *STOC*, pages 365–377.

Juels, A. and Peinado, M. (2000). Hiding cliques for cryptographic security. *Des. Codes Cryptography*, 20(3):269–280.

McEliece, R. J. (1978). A Public-Key Cryptosystem Based On Algebraic Coding Theory. *Deep Space Network Progress Report*, 44:114–116.

Paillier, P. (1999). Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT*, pages 223–238.

Rivest, R. L., Adleman, L., and Dertouzos, M. L. (1978a). On data banks and privacy homomorphisms. In DeMillo, R. A., Dobkin, D. P., Jones, A. K., and Lipton, R. J., editors, *Foundations of Secure Computation*, pages 165–179. Academic Press.

Rivest, R. L., Shamir, A., and Adleman, L. M. (1978b). A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126.

Strenzke, F. (2011). Message-aimed side channel and fault attacks against public key cryptosystems with homomorphic properties. *J. Cryptographic Engineering*, 1(4):283–292.

Wyner, A. D. (1975). The wire-tap channel. *The Bell System Technical Journal*, 54(8):1355–1387.