Secure Compressed Sensing over Finite Fields

Original

Secure Compressed Sensing over Finite Fields / Bioglio, Valerio; Bianchi, Tiziano; Magli, Enrico. - (2014), pp. 191-196. (Intervento presentato al convegno 2014 IEEE International Workshop on Information Forensics and Security tenutosi a Atlanta, GA, USA nel December 3-5, 2014) [10.1109/WIFS.2014.7084326].

Availability:

This version is available at: 11583/2580548 since:

Publisher:

IEEE - INST ELECTRICAL ELECTRONICS ENGINEERS INC

Published

DOI:10.1109/WIFS.2014.7084326

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Secure Compressed Sensing over Finite Fields

V. Bioglio, T. Bianchi, E. Magli

Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino

Abstract—In this paper, we analyze the security of compressed sensing (CS) defined over finite fields. First, we prove that acquiring signals using dense sensing matrices may provide almost perfect secrecy. Then, we prove that using sparse sensing matrices, which admit efficient recovery algorithms motivated by coding theory, reveals information only on the sparsity of the sensed signal, and that such information is conveyed only by the sparsity of the measurements. Finally, we introduce an operational definition of security, based on the error probability in estimating the signal sparsity, and show that there is a tradeoff between the sparsity of the sensing matrix and the security of the CS system.

Index Terms—Compressed Sensing, Security, Finite Fields

I. INTRODUCTION

Compressed Sensing (CS) has recently emerged as a novel signal acquisition technique at a rate well below that predicted by the classical Shannon-Nyquist theory [1], [2]. The key intuition behind CS is that, under the hypothesis that the signal admits a sparse representation under some basis, a small number of linear measurements enables signal recovery with high probability, as long as such measurements can be modeled as linear projection over a second basis which is incoherent with respect to the sparsity basis.

Most of the results in the CS literature deal with signals represented over a real field. However, if the signal belongs to a finite alphabet, the recovery problem can be addressed exploiting finite fields operations. The study of CS over finite fields is an emerging topic [3], [4], motivated by the fact that it can provide some advantages with respect to classical CS. Namely, while sensing and measurement quantization of a real signal may cause loss of accuracy, performing operations over finite fields avoids this issue. Moreover, techniques derived from the decoding of linear codes can be exploited for efficient signal recovery [4]–[6].

As to CS over a real field, researchers have hinted the possibility that acquiring signals via random linear projection may provide some notion of security. In [7], the authors argue that CS does not provide information theoretic secrecy [8], however it can be viewed as a cryptosystem offering computational secrecy. The security of CS is also investigated in [9], showing that CS is computationally secure against a systematic search of the sensing matrix, even if the sparsity is known. Low complexity encryption systems based on CS have been recently proposed and analyzed in [10], [11].

In this paper, we analyze the security of CS over finite fields as a cryptosystem. Even if the notion of secrecy emerges in a natural way for CS over finite fields, to the best of our knowledge there is no precise study in the literature

addressing the security of such systems. We consider the case of dense as well as sparse sensing matrices. This second case is particularly interesting, since it admits signal recovery via efficient linear decoding algorithms [6]. While using a dense sensing matrix over finite fields provides almost perfect security, we will show that a sparse sensing matrix reveals information only on the *sparsity* of the sensed signal, and that such information is conveyed only by the sparsity of the linear measurements. In order to characterize this information leakage, we will introduce an operational security definition based on the performance of practical estimators, showing that there is a tradeoff between the sparsity of the sensing matrix and the security of the system. Simulation results are included to evaluate such a tradeoff in simple scenarios.

II. BACKGROUND

A. Compressed Sensing and Finite Fields

A signal x of length n is called k -sparse if there exists a basis Φ such that $x = \Phi\theta$ and θ has exactly k nonzero entries, denoted as $\|\theta\|_0 = k$. According to the CS framework [1], [2], a k -sparse signal can be exactly recovered from $m < n$ linear measurements $y = Ax$ by solving the minimization problem

$$\hat{\theta} = \arg \min_{\theta} \|\theta\|_0 \quad \text{subject to } A\Phi\theta = y, \quad (1)$$

as long as $m \geq 2k$ and the sensing matrix $A \in \mathbb{R}^{m \times n}$ satisfies certain properties.

In the conventional CS framework, the signal is defined over the real field $\mathbb{R}$, but nothing prevents to employ signals defined over a finite field $\mathbb{F}$ of size q . The authors in [3] prove that if all the elements of the sensing matrix are randomly drawn over a finite field, (1) admits a unique k -sparse solution with high probability. However, even if the theoretical performance of these dense random matrices over finite fields appears to be promising, efficient algorithms for the recovery of the signal in this scenario are not currently available.

Actually, solving (1) over a finite field can be seen as a syndrome decoding of an error correcting code, where θ is the channel error vector, y is the syndrome of the received word and $A\Phi$ is the parity check matrix of the code [12]. When $A\Phi$ is a fully random matrix, syndrome decoding is believed to be NP-hard [13]. However, the sensing matrix A can be chosen such that $A\Phi$ is the parity check matrix of a channel code for which an efficient decoding algorithm is known. In this sense, an interesting class of codes are the so called low-density parity check (LDPC) codes. LDPC codes admit efficient recovery algorithms using belief propagation.

To simplify the notation, in the following we can consider $\Phi = I$, i.e., the signal is sparse in the sensing domain, without loss of generality.

B. Security definitions, scenarios, and attack models

Let us define the set of possible plain texts $\mathcal{P}$, the set of cipher texts $\mathcal{C}$ and a key generation function providing a key K . A private key cryptosystem is a pair of functions $e_K : \mathcal{P} \rightarrow \mathcal{C}$, $d_K : \mathcal{C} \rightarrow \mathcal{P}$ such that, given a plain text $p \in \mathcal{P}$, $d_K(e_K(p)) = p$ and such that, given a cipher text $c \in \mathcal{C}$, it is unfeasible to determine p such that $e_K(p) = c$, without knowing the key K .

A cryptosystem is perfectly secure [8] if the posterior probability of the cipher text given the plain text p is independent of p . For a perfectly secure cryptosystem, any attack can not be more successful than guessing the plain text at random. Practical cryptosystems are usually computationally secure, meaning that breaking the cryptosystem is equivalent to solve a computationally hard problem, that is, a problem whose solution can not be computed in polynomial time with respect to the size of the key.

Given the CS model $y = Ax$ over a finite field, we can define the following equivalences between CS and a private key cryptosystem: the signal x is the plain text, the sensing matrix A is the secret key and the measurement vector y is the cipher text. The encryption function is matrix multiplication, whereas decryption is achieved by solving the problem in (1). As a consequence, the notion of perfect security can be extended to the cryptosystem defined by the CS framework:

Definition 1. A CS system provides *perfect secrecy* if $\mathbb{P}(y|x) = \mathbb{P}(y)$.

Definition 2. A CS system provides *conditional perfect secrecy* if $\mathbb{P}(y|x, a) = \mathbb{P}(y|a)$, where a is a system parameter.

The last is equivalent to saying that observing y does not reveal anything more about x than what can be inferred by the knowledge of the parameter a .

The security of the CS-based cryptosystem will be affected by the policies regarding the generation of the sensing matrix in the case of multiple measurements. On the one hand, using the same sensing matrix for multiple measurements limits the overhead due to the transmission of the sensing matrix. On the other hand, generating a different and possibly independent sensing matrix for each measurement is somewhat analogous to a one-time pad cryptosystem and may offer greater security. In this paper, we will focus on the one-time sensing matrix (OTS) scenario [11]. We will assume that each sensing matrix is used only once, and that different sensing matrices are statistically independent. Under this scenario, it is sufficient to consider the security of $y = Ax$, since multiple measurements will be statistically independent. Such a scenario can be implemented by producing consecutive sensing matrices via a secure random number generator (SRNG) [14] and using the seed of the SRNG as a secret key.

The security of a cryptosystem depends also on the resources of the adversary. In this paper, we will focus on a

ciphertext-only attack (COA) scenario where the adversary has only knowledge of the measurements y .

C. Niederreiter's Cryptosystem

CS-based cryptosystems defined over finite fields have a close relationship with a public key cryptosystem proposed by Niederreiter [15], which is a dual version of the popular McEliece cryptosystem [16], [17]. Both cryptosystems rely on the assumption that decoding a random code is a hard problem and that it is computationally unfeasible to distinguish a scrambled code from a random code. Namely, given a parity check matrix H of an error correcting code that can correct up to k errors, the Niederreiter's cryptosystem is based on the scrambled parity check matrix $A = SHP$, where S is a random invertible matrix and P is a random permutation. The public key is A , the secret key is (S, H, P) . A message is first encoded as a k -sparse vector x and the corresponding ciphertext is obtained as $y = Ax$. For the decryption, the receiver applies syndrome decoding for the code H to the vector $y' = S^{-1}y$ and recovers x by inverting P .

The original version of Niederreiter's cryptosystem based on Reed-Solomon codes can be broken in polynomial time [18]. However, when implemented with Goppa codes [19] the cryptosystem is believed to be secure [20]. As to CS over finite fields, a disadvantage of McEliece's and Niederreiter's cryptosystems is that the matrix A and the sparsity parameter k must satisfy some constraints in order to be secure. For example, the authors of [20] suggest using (2960, 2288) Goppa codes and $k = 56$ in order to have 128-bit security.

In the following, we will consider a different scenario with respect to Niederreiter's cryptosystem. We will assume that the choice of A is mainly dictated by the properties of the sensed signal x and we will explore whether sensing matrices with good recovery properties can also guarantee a security layer if used as secret keys of a symmetric cryptosystem.

III. SECURITY OF DENSE SENSING MATRICES

In this section we will prove that, under some hypotheses, when using randomly drawn dense sensing matrices, as proposed in [3], the CS system is almost perfectly secure. First, we need the following preliminary result:

Lemma 1. Given a vector $r \in \mathbb{F}^n$, if the elements of r are uniformly drawn over $\mathbb{F}$ and $k \geq 1$, the scalar product $r \cdot x$ does not depend on x and it is uniformly distributed over $\mathbb{F}$, i.e., $\mathbb{P}(r \cdot x = a | x) = \mathbb{P}(r \cdot x = a) = \frac{1}{q}$.

Proof: Given $x \in \mathbb{F}^n$, we call $Q_a(x) = \{v \in \mathbb{F}^n \text{ s.t. } v \cdot x = a\}$, where $a \in \mathbb{F}$. For $a = 0$, $Q_0(x)$ is a (additive) subgroup of $\mathbb{F}^n$ (and we write $Q_0(x) < \mathbb{F}^n$) since for all $v, w \in Q_0(x)$ we have that $v - w \in Q_0(x)$. If $k \geq 1$, it is possible to prove that $Q_a(x)$ is a coset of $Q_0(x)$, since for all $v, w \in Q_a(x)$ we have that $v - w \in Q_0(x)$. $\mathbb{F}^n$ is an abelian group for the sum, so the right and left cosets are equal. Through Lagrange's theorem, we have that $q^n = |\mathbb{F}^n| = [\mathbb{F}^n : Q_0(x)] \cdot |Q_0(x)| = q |Q_0(x)|$, and hence $|Q_a(x)| = |Q_0(x)| = q^{n-1}$ for all $a \in \mathbb{F}$. This implies that, if

r is randomly drawn, $\mathbb{P}(r \cdot x = a | x) = \mathbb{P}(r \cdot x = a) = \mathbb{P}(r \in Q_a(x))$, and the last probability is equal to $\frac{1}{q}$. ■

Hence, we are ready to state the main result of this section:

Proposition 1. If $k \neq 0$ and the elements of A are uniformly drawn over $\mathbb{F}$, the CS system provides perfect secrecy.

Proof: If y is randomly drawn over $\mathbb{F}^m$, $\mathbb{P}(y) = \frac{1}{q^m}$. On the contrary, if we call A_i the i -th row of A , for Lemma 1

$$\begin{aligned} \mathbb{P}(y | x) &= \mathbb{P}(y = A \cdot x | x) = & (A \text{ i.i.d}) \\ &= \prod_{i=1}^m \mathbb{P}(y_i = A_i \cdot x | x) = & (\text{Lemma 1}) \\ &= \prod_{i=1}^m \mathbb{P}(y_i = A_i \cdot x) = \frac{1}{q^m}. \end{aligned}$$

Following Def. 1, the proposition is proved. ■

The previous proposition proves that the use of a random sensing matrix provides a high level of security. However, for this kind of matrices no efficient recovery algorithm is known. To overcome this problem, the use of sparse matrices is emerging as an alternative to the dense ones [5], [6]. In the following section, we will study the security performance of sparse sensing matrices.

IV. SECURITY OF SPARSE SENSING MATRICES

In this section we will prove two security properties of sparse sensing matrices satisfying certain conditions. Firstly, we prove that, under a known sparsity k , every signal x is equiprobable given a measurement vector y , implying that the CS system is conditionally perfectly secure. Secondly, we consider that a sparse A implies a sparse y and we prove that, if the sensing matrix has uniformly drawn nonzero elements, all measurements vectors with a constant sparsity $h = \|y\|_0$ are equiprobable, implying that y does not reveal anything more about x than what can be inferred by the knowledge of h . Finally, these properties are used together to prove that, given y , the only information we can leak is k , and this information can be obtained through the observation of h .

A. Equiprobability of the signals

We suppose that the sensing matrix belongs to a subset of the complete field, i.e., $A \in \mathcal{F} \subseteq \mathbb{F}^{m \times n}$, closed under the elementary matrix operations of column switching and multiplication. This means that if $A \in \mathcal{F}$, also the matrices A' , obtained switching two columns of A , and A'' , obtained multiplying a column of A by an element of $\mathbb{F}$, belong to $\mathcal{F}$. Parity check matrices of linear codes are closed under these operations.

We call $\mathcal{A}(x, y) \subseteq \mathcal{F}$ the set of the *admissible matrices* defined as $\mathcal{A}(x, y) = \{A \in \mathcal{F} \text{ s.t. } A \cdot x = y\}$. As a consequence, $\mathbb{P}(y | x) = \frac{|\mathcal{A}(x, y)|}{|\mathcal{F}|}$, and the system provides perfect secrecy only if the sets of the admissible matrices have the same size. Unfortunately, these sets have different size; however, with Prop. 2 we will prove that the size of an admissible set only depends on the sparsity of x , hence the system provides conditional perfect security given k .

We call S_x the support of x , i.e., $S_x \subseteq \{1, \dots, n\}$ such that $x_i \neq 0$ iff $i \in S_x$; obviously, $|S_x| = k$. The vector $\pi(x)$ is a permutation of x , applied by the permutation $\pi \in \mathcal{S}_n$. In particular, $\pi_{a,b}$ is the swap between the elements in position a and b . Before proving the main statement, we need the following preliminary results:

Lemma 2. $|\mathcal{A}(x, y)| = |\mathcal{A}(\pi_{a,b}(x), y)|$ for all $a, b \in \{1, \dots, n\}$.

Proof: We consider the map $\varphi : \mathcal{A}(x, y) \rightarrow \mathcal{A}(\pi_{a,b}(x), y)$ defined as the function that swaps the a -th and the b -th columns of A . If we prove that φ is bijective, the lemma holds. To be bijective, a function must be well defined, injective and surjective. The function is well defined because if $A \in \mathcal{A}(x, y)$ then $\varphi(A) \in \mathcal{A}(\pi_{a,b}(x), y)$ by definition. The function is injective because only two identical matrices can be mapped into the same matrix. Finally, the function is surjective because if $B \in \mathcal{A}(\pi_{a,b}(x), y)$ there exists a matrix $A \in \mathcal{A}(x, y)$ such that $\varphi(A) = B$. ■

Corollary 1. $|\mathcal{A}(x, y)| = |\mathcal{A}(\pi(x), y)|$ for all $\pi \in \mathcal{S}_n$.

Proof: The corollary holds since every permutation admits a decomposition in swaps and Lemma 2 can be applied for every swap. ■

Lemma 3. If $S_x = S_{x'}$ and $\|x - x'\|_0 = 1$ then $|\mathcal{A}(x, y)| = |\mathcal{A}(x', y)|$.

Proof: If $S_x = S_{x'}$ and $\|x - x'\|_0 = 1$ then x and x' differ for one element in position a . We consider the map $\varphi_\alpha : \mathcal{A}(x, y) \rightarrow \mathcal{A}(x', y)$ defined as the function that multiplies by α all the elements of the a -th column of A . If $\alpha = x'_a \cdot x_a^{-1}$, the function is a bijection between the two sets. ■

Corollary 2. If $S_x = S_{x'}$ then $|\mathcal{A}(x, y)| = |\mathcal{A}(x', y)|$.

Proof: We create a succession $x = x^{(0)}, x^{(1)}, \dots, x^{(k-1)} = x'$ where $S_{x^{(i)}} = S_{x^{(i+1)}}$ and $\|x^{(i)} - x^{(i+1)}\|_0 = 1$. This succession is created by substituting one by one the nonzero elements of x with the corresponding nonzero elements of x' in the same position. Now, it is possible to exploit Lemma 3 at each step to prove the corollary. ■

Proposition 2. If $|S_x| = |S_{x'}|$ then $|\mathcal{A}(x, y)| = |\mathcal{A}(x', y)|$.

Proof: If $|S_x| = |S_{x'}|$ there exists a permutation $\pi \in \mathcal{S}_n$ such that $S_{\pi(x)} = S_{x'}$, and hence $|\mathcal{A}(x, y)| = |\mathcal{A}(\pi(x), y)| = |\mathcal{A}(x', y)|$, where the first and the second equalities are due to Corollaries 1 and 2 respectively. ■

We are now ready to state the main result:

Proposition 3. If A is randomly drawn from a family $\mathcal{F}$ closed under the operations of column switching and multiplication, then the CS system provides conditional perfect secrecy, i.e., $\mathbb{P}(y | x, k) = \mathbb{P}(y | k)$.

Proof: Under the hypothesis of the proposition, $\mathbb{P}(y | x) = \frac{|\mathcal{A}(x, y)|}{|\mathcal{F}|}$. Given a signal x' such that $|S_x| = |S_{x'}| = k$, for Prop. 2 we have that $\mathbb{P}(y | x) = \mathbb{P}(y | x')$, hence $\mathbb{P}(y | x) = \mathbb{P}(y | x, k) = \mathbb{P}(y | k)$. ■

This is equivalent to saying that observing y does not

reveal anything more about x than what can be inferred by the knowledge of k . Moreover, since the set of the parity check matrices of a nonbinary linear codes is a subset of $\mathcal{F}$ closed under the matrix operations of column switching and multiplication, Prop. 2 holds for these families of matrices.

B. Equiprobability of the measurements

We begin proving an extension of Lemma 1 for the nonzero measurements calculated through sparse matrices.

Lemma 4. Given a sparse vector $r \in \mathbb{F}^n$, if the nonzero elements of r are uniformly drawn over $\mathbb{F} \setminus \{0\}$ and $b = r \cdot x \neq 0$, then b does not depend on x and it is uniformly distributed over $\mathbb{F} \setminus \{0\}$, i.e., $\mathbb{P}(r \cdot x = b \mid r \cdot x \neq 0, x) = \mathbb{P}(r \cdot x = b \mid r \cdot x \neq 0) = \frac{1}{q-1}$.

Proof: We call l_j the d positions where both r_{l_j} and x_{l_j} are nonzero. Since $r \cdot x \neq 0$, we have $d > 0$ and $b = \sum_{j=1}^n r_j x_j = \sum_{j=1}^d r_{l_j} x_{l_j}$. If $d = 1$, b is the result of the multiplication of a value uniformly drawn over $\mathbb{F} \setminus \{0\}$ and a generic element of $\mathbb{F} \setminus \{0\}$, hence it is distributed uniformly over $\mathbb{F} \setminus \{0\}$. On the contrary, if $d > 1$ we have that $b = \sum_{j=1}^d r_{l_j} x_{l_j} = r_{l_1} x_{l_1} + \sum_{j=2}^d r_{l_j} x_{l_j} = r_{l_1} (x_{l_1} + \sum_{j=2}^d r_{l_1}^{-1} r_{l_j} x_{l_j}) = r_{l_1} x'$. As before, b can be written as the result of the multiplication of a value uniformly drawn over $\mathbb{F} \setminus \{0\}$ and a generic elements of $\mathbb{F} \setminus \{0\}$, hence it is distributed uniformly over $\mathbb{F} \setminus \{0\}$ also in this case. ■

Hence, we can prove the following statement:

Proposition 4. If the values of the nonzero elements of A are drawn uniformly over $\mathbb{F} \setminus \{0\}$ then the values of the nonzero elements of y are drawn uniformly over $\mathbb{F} \setminus \{0\}$, and $\mathbb{P}(x \mid y, h) = \mathbb{P}(x \mid h)$.

Proof: Each element of y is the result of the scalar product between the signal and a row of A . Each row of A is a vector that matches the hypotheses of Lemma 4, hence the values of the nonzero elements of y are drawn uniformly over $\mathbb{F} \setminus \{0\}$. As a consequence, the values of the nonzero elements of y do not give any information on x ; the only usable information carried by y is its sparsity h , hence $\mathbb{P}(x \mid y, h) = \mathbb{P}(x \mid h)$. ■

The consequence of this Proposition is that the value of a nonzero entry of y does not give any further information on x than the information given by the knowledge of $h = \|y\|_0$.

C. Information leakage

Thanks to Propositions 3 and 4 above, the distribution of y given k does not depend on the values of x , and conversely the distribution of x given h does not depend on the values of y . We can use this to calculate the mutual information between x and y , proving the following statement:

Proposition 5. If the hypotheses of Prop. 3 and 4 hold, the mutual information between x and y is equal to the mutual information between k and h , i.e., $I(x; y) = I(k; h)$.

Proof:

$$\begin{aligned} I(x; y) &= I(x, k; y) = I(k; y) + I(x; y \mid k) = \\ &= I(k; y) \end{aligned} \quad (\text{Prop. 3})$$

$$\begin{aligned} I(k; y) &= I(k; y, h) = I(k; h) + I(k; y \mid h) = \\ &= I(k; h) \end{aligned} \quad (\text{Prop. 4})$$

This means that the only feature of y that reveals some information is its sparsity h , and the only information we can leak is the sparsity of the signal k . We point out that the parity check matrices of linear codes match the hypotheses of Prop. 5, hence the discussion above is valid for this class of matrices.

V. OPERATIONAL DEFINITION OF SECURITY

In this section, we want to characterize the information leakage about the signal given by the knowledge of the measurements in the case of sparse sensing matrices. Since, under the hypothesis of Prop. 5, the sparsity of the measurements h gives some information about the sparsity of the signal k , we develop an operational definition of security based on the error committed by attempting to estimate k from h . Let us define the error probability (EP) as $\mathbb{P}(\hat{k}(h) \neq k)$, where $\hat{k}(h)$ is an estimator of k based on h .

Definition 3. A CS system is said η -EP secure if

$$\eta = \frac{\mathbb{E}[\mathbb{P}(\hat{k}(h) \neq k)]}{1 - \max_k \mathbb{P}(k)}.$$

It is easy to check that a perfectly secure system is 1-EP secure. Conversely, a 0-EP secure system reveals everything about k . In general, the security of a CS system can be assessed by using the following lemma:

Lemma 5. A CS system is η^* -EP secure, where

$$\eta^* = \frac{\mathbb{E}[1 - \max_k \mathbb{P}(k \mid h)]}{1 - \max_k \mathbb{P}(k)}. \quad (2)$$

Proof: According to Bayesian estimation theory, the EP is minimized by the maximum a posteriori (MAP) estimator of k , given by

$$\hat{k}_{MAP}(h) = \arg \max_k \mathbb{P}(k \mid h). \quad (3)$$

Hence, for each h we have $\min_{\hat{k}(h)} \mathbb{P}(\hat{k}(h) \neq k) = 1 - \max_k \mathbb{P}(k \mid h)$. ■

In order to evaluate the security of a given CS system using lemma 5, we need to explicitly compute the probability distribution $\mathbb{P}(k \mid h)$. Since $\mathbb{P}(k \mid h) = \mathbb{P}(h \mid k) \mathbb{P}(k)$, if we assume that the prior distribution of k is known, the problem is equivalent to determining the probability distribution $\mathbb{P}(h \mid k)$ induced by a particular sensing matrix.

A. Distribution of $\mathbb{P}(h \mid k)$

In the following, we explicitly calculate the distribution of $\mathbb{P}(h \mid k)$ in the case of matrices with independent rows, i.e., if each row of A is generated independently from the others. Moreover, the nonzero entries have to be drawn uniformly over $\mathbb{F} \setminus \{0\}$. In this case, Prop. 5 holds. This leads to more restrictive hypotheses than those requested in Prop. 5, but allows one to write $\mathbb{P}(h \mid k)$ as a function of known distributions.

We begin by calculating the probability P that an entry of y is equal to zero. We call d_i the number of positions such that an entry a_{ij} of the i -th row of the sensing matrix and x_j are both nonzero. In this case,

$$P = \sum_{j=0}^k \mathbb{P}(y_i = 0 | d_i = j) \mathbb{P}(d_i = j) = \sum_{j=0}^k p_j t_j,$$

where $p_j = \mathbb{P}(d_i = j)$ and $t_j = \mathbb{P}(y_i = 0 | d_i = j)$. Row independence implies that each measurement is independent from the others, hence $\mathbb{P}(h | k)$ is distributed according to a binomial distribution of parameter $1 - P$, i.e., $\mathbb{P}(h | k) \sim \mathcal{B}(1 - P, m)$. We note that if $q \rightarrow \infty$, $P \rightarrow p_0 t_0 = p_0$. In fact, h is a random variable that counts the number of nonzero entries of y . Since the rows of A are independent, each measurement is independent on the others, and it is equal to zero with probability P . As a result, h is the sum of m Bernoulli processes, hence it is distributed according to a binomial distribution. The parameter P depends on the field and on A . In the following, we separately study these factors that determine P .

$\mathbb{P}(y_i = 0 | d_i)$ depends on the size q of the finite field. Given the arrangement of the nonzero elements in the matrix, we naturally have that $t_0 = \mathbb{P}(y_i = 0 | d_i = 0) = 1$ and $t_1 = \mathbb{P}(y_i = 0 | d_i = 1) = 0$. The general case $t_j = \mathbb{P}(y_i = 0 | d_i = j)$ can be recursively calculated as follows. If $d_i = 2$, t_2 can be seen as the probability for two nonzero elements to be complementary, hence $t_2 = \frac{1}{q-1}$. When $d_i = 3$, t_3 can be seen as the probability that, given three nonzero elements, the sum of the first two is complementary to the third. They are complementary if the sum is a nonzero (that happens with probability $1 - t_2$) and they assume the same value, hence $t_3 = (1 - t_2) \frac{1}{q-1}$. This procedure can be generalized, obtaining $t_j = \mathbb{P}(y_i = 0 | d_i = j) = (1 - t_{j-1}) \frac{1}{q-1}$, for $j \geq 2$.

The distribution of d_i depends on the arrangement of the nonzero elements in the matrix. Since the rows are independent by definition, d_i does not depend on i . In the following, we calculate the distribution of d_i for some kind of sparse matrices.

Rows with Constant Degree: Each row of A has exactly r nonzero entries in positions randomly drawn. Each nonzero entry is drawn uniformly over $\mathcal{F} \setminus \{0\}$. In this case, d_i is distributed according to an hypergeometric distribution, i.e., $d_i \sim \mathcal{H}(n, k, r)$. In fact, d_i counts the number of collisions between the signal x and the i -th row of the sensing matrix. These are two n -length vectors, where the first has k nonzero elements while the second has r nonzero entries. A collision occurs where both the row and the signal have a nonzero value in the same entry. As a consequence, the random variable that counts these collisions is distributed according to a hypergeometric distribution.

Probabilistically Sparse Matrices: Each entry a_{ij} of A is set to zero with probability $1 - \alpha$, where α is a tunable parameter. With probability α the entry will be a nonzero, and its value is drawn uniformly over $\mathcal{F} \setminus \{0\}$. In this case, d_i is distributed according to a binomial distribution,

i.e., $d_i \sim \mathcal{B}(k, \alpha)$. In fact, d_i counts the number of collisions as in the previous case, but only depends on the values of the entries of the i -th row in correspondence to the k nonzero elements of x . Since each entry is independent and assumes a nonzero value with probability α , d_i is the result of the sum of k Bernoulli processes, hence it is distributed according to a binomial distribution.

According to (4), we can use the distribution of $\mathbb{P}(h | k)$ calculated above to find a MAP estimator for k as

$$\hat{k}_{MAP}(h) = \arg \max_k \left(\frac{\binom{m}{h} (1 - P(k))^h (P(k))^{m-h} \cdot \mathbb{P}(k)}{\mathbb{P}(h)} \right) \quad (4)$$

where $P(k) = P$ since k is the only free parameter that generated P . In the following section, we will exploit this estimator in (2) to evaluate the accuracy of the proposed security model.

VI. EXPERIMENTAL RESULTS

In this section, we evaluate the security of a CS system defined over finite fields for different values of the involved parameters. For each experiment, the average EP of the MAP estimator in (4) is compared to the theoretical η -EP value obtained through (2). In order to have a MAP estimate of k , we need to assume an a priori distribution for k . We propose two models for the signal x that naturally result in a distribution of k . Both models provide a bound on the possible values of k , i.e., $k_{min} \leq k \leq k_{max}$.

Uniform sparsity: k takes value uniformly between k_{min} and k_{max} . Each value is picked with probability $\frac{1}{k_{max} - k_{min} + 1}$.

Binomial sparsity: $k = k_{min} + B$, where B is distributed according to the binomial distribution $B \sim \mathcal{B}(k_{max} - k_{min}, \beta)$, with β a tunable parameter.

In Fig. 1, the security behavior of probabilistically sparse matrices is plotted as a function of the matrix sparsity parameter α . The curves refer to the two distributions of k just presented, where $n = 200$, $m = 100$, $q = 256$ and $2 \leq k \leq 20$. The simulated EP values, obtained by averaging over 10000 independent tests, prove the accuracy of the theoretical model. As expected, in general a denser sensing matrix provides a better security. The system appears to have a high security even if the matrix is very sparse: in this case, however, the matrix is so sparse that $y = 0$ with high probability, and the measurements do not carry enough information for the recovery of x . In Fig. 2, the EP is plotted as a function of the number of measurements m , for $\alpha = 0.3$, while all the other parameters are the same as in the previous figure. This plot shows that the security of the system decreases as the number of measurements increases. This means that the more measurements we have, the better will be the sparsity estimation. We can also note that a uniform distribution of k is more secure than a binomial one, since in the latter case the probability is more concentrated on few values. The performance in the case of matrix having rows with constant degree is similar, hence the plots are not shown.

Fig. 1. Value of η -EP vs. α for uniform and binomial distribution of k , where $n = 200$, $m = 100$, $q = 256$ and $2 \leq k \leq 20$.

Fig. 2. Value of η -EP vs. m for uniform and binomial distribution of k , where $n = 200$, $\alpha = 0.3$, $q = 256$ and $2 \leq k \leq 20$.

VII. CONCLUSIONS

In this paper, we have analyzed the security of a CS system defined over a finite field. We have proved that an important class of sparse sensing matrices, which includes the parity check matrices of linear codes, provides conditional perfect secrecy if the sparsity of the signal is known. In other words, the measurements taken with such matrices leak only the sparsity of the signal. Moreover, under additional assumptions, only the sparsity of the measurements leaks information about the sparsity of the signal. We have then introduced an operational definition of security, based on the probability of error of a MAP estimator, showing that there is a tradeoff between the security of the CS system and the sparsity of the sensing matrix. Since the signal may not be recoverable in the case of very sparse sensing matrices, an interesting open

problem is determining for what sparsity parameters of the sensing matrix the legitimate user has indeed an information advantage with respect to the attacker. Beyond the security aspects, it is worth noting that the above results also suggest the possibility of using a sparse sensing matrix to estimate the signal sparsity without recovering the signal, which could be an interesting application in compressed sensing.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC Grant agreement n. 279848.

REFERENCES

- [1] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [2] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [3] S. C. Draper and S. Malekpour, "Compressed sensing over finite fields," in *IEEE Int. Symp. Inf. Theory (ISIT)*, 2009, pp. 669–673.
- [4] A. K. Das and S. Vishwanath, "On finite alphabet compressive sensing," in *2013 IEEE Int. Conf. Acoustics, Speech and Signal Proc. (ICASSP)*, 2013, pp. 5890–5894.
- [5] J. T. Seong and H. N. Lee, "On the compressed measurements over finite fields: Sparse or dense sampling," in *arXiv*, 2012.
- [6] V. Bioglio, G. Coluccia, and E. Magli, "Sparse image recovery using compressed sensing over finite alphabets," in *2014 IEEE Int. Conf. Image Proc. (ICIP)*, 2014, accepted.
- [7] Y. Rachlin and D. Baron, "The secrecy of compressed sensing measurements," in *2008 46th Annual Allerton Conference on Communication, Control, and Computing*. IEEE, 2008, pp. 813–817.
- [8] C. E. Shannon, "Communication theory of secrecy systems," *Bell. Syst. Tech. J.*, vol. 28, pp. 656–715, 1949.
- [9] A. Orsdemir, H. Altun, G. Sharma, and M. Bocko, "On the security and robustness of encryption via compressed sensing," in *IEEE Military Communications Conference, 2008 (MILCOM 2008)*, 2008, pp. 1–7.
- [10] V. Cambareri, J. Haboba, F. Pareschi, H. Rovatti, G. Setti, and K.-W. Wong, "A two-class information concealing system based on compressed sensing," in *2013 IEEE Int. Symp. Circuits and Systems (ISCAS)*, 2013, pp. 1356–1359.
- [11] T. Bianchi, V. Bioglio, and E. Magli, "On the security of random linear measurements," in *2014 IEEE Int. Conf. Acoustics, Speech and Signal Proc. (ICASSP)*, 2014, pp. 4020–4024.
- [12] C. E. J. and T. T., "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [13] E. Berlekamp, R. McEliece, and H. C. A. Van Tilborg, "On the inherent intractability of certain coding problems (corresp.)," *IEEE Trans. Inf. Theory*, vol. 24, no. 3, pp. 384–386, May 1978.
- [14] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone, *Handbook of Applied Cryptography*. CRC Press, 1996.
- [15] H. Niederreiter, "Knapsack-type cryptosystems and algebraic coding theory," *Problems of Control and Information Theory/Problemy Upravleniya i Teorii Informacii*, vol. 15, no. 2, pp. 157–166, 1986.
- [16] R. J. McEliece, "A public-key cryptosystem based on algebraic coding theory," *DSN progress report*, vol. 42-44, pp. 114–116, Jan.-Feb. 1978.
- [17] Y. X. Li, R.-H. Deng, and X.-M. Wang, "On the equivalence of McEliece's and Niederreiter's public-key cryptosystems," *IEEE Trans. Inf. Theory*, vol. 40, no. 1, pp. 271–273, Jan 1994.
- [18] V. M. Sidel'nikov and S. O. Shestakov, "On the insecurity of cryptosystems based on generalized Reed-Solomon codes," *Discrete Mathematics and Applications*, vol. 2, no. 4, p. 439444, 1992.
- [19] E. Berlekamp, "Goppa codes," *IEEE Trans. Inf. Theory*, vol. 19, no. 5, pp. 590–592, Sep 1973.
- [20] D. J. Bernstein, T. Lange, and C. Peters, "Attacking and defending the McEliece cryptosystem," in *Post-Quantum Cryptography*, ser. Lecture Notes in Computer Science, J. Buchmann and J. Ding, Eds. Springer Berlin Heidelberg, 2008, vol. 5299, pp. 31–46.