POLITECNICO DI TORINO
Repository ISTITUZIONALE

On the secrecy of compressive cryptosystems under finite-precision representation of sensing matrices

Original
On the secrecy of compressive cryptosystems under finite-precision representation of sensing matrices / Testa, M., Bianchi, T., Magli, E.. - ELETTRONICO. - (2018), pp. 1-4. (2018 IEEE International Symposium on Circuits and Systems (ISCAS) Florence, Italy 27-30 May 2018) [10.1109/ISCAS.2018.8351443].

Availability:
This version is available at: 11583/2709510 since: 2018-06-11T09:16:51Z

Publisher:
IEEE

Published
DOI:10.1109/ISCAS.2018.8351443

Terms of use:

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

Publisher copyright
IEEE postprint/Author's Accepted Manuscript

©2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

15 October 2026

On the secrecy of compressive cryptosystems under finite-precision representation of sensing matrices

Matteo Testa, Tiziano Bianchi and Enrico Magli
Department of Electronics and Telecommunications Engineering
Politecnico di Torino - Italy

***Abstract*—In recent years, the Compressed Sensing (CS) framework has been shown to be an effective private key cryptosystem. If infinite precision is available, then it has been shown that spherical secrecy can be achieved. However, despite its theoretically proven secrecy properties, the only practically feasible implementations involve the use of Bernoulli sensing matrices. In this work, we show that different distributions employing a much larger finite alphabet can be considered. More in detail, we consider the use of quantized Gaussian sensing matrices and experimentally show that, besides being suitable for practical implementation, they can achieve higher secrecy with respect to Bernoulli sensing matrices. Furthermore, we show that this approach can be used to tune the secrecy of the CS cryptosystems based on the available machine precision.**

I. Introduction

Compressed Sensing (CS) [1] has been extensively studied over the last decade as an attractive way to perform dimensionality reduction, and has recently gained popularity as an effective way to encrypt data. Accordingly to CS theory, a K-sparse signal, i.e. a signal with K non-zero entries, can be exactly recovered with overwhelming probability from its random measurements if some assumptions on the sensing matrix are satisfied [1]. One of the underlying assumptions is that the sensing matrix, i.e. the random projection matrix which defines the undetermined CS problem, is known at reconstruction stage. Conversely, the only knowledge of the measurements, does not allow to recover the original signal. Therefore, the CS framework can be seen as a private key cryptosytem where the sensing matrix entries are the secret, thus only shared among trusted parties, the original signal is the plaintext and the measurements are the ciphertext. The encryption is performed by means of CS acquisition while the decryption corresponds to CS recovery.

The seminal study of Rachlin and Baron [2] was the first effort in this direction. Further studies such as [3] focused on the asymptotic secrecy properties of compressive encryption, showing that measurements of equal energy signals becomes indistinguishable as the size of the original signals tends to infinity. Non-asymptotic analysis of the distinguishability of measurements sampled from Gaussian i.i.d. sensing matrices was carried out in [4]. In this latter work, the authors show that normalizing the signal to unit energy leads to perfect secrecy under the assumption of one time sensing (OTS) acquisition, i.e. the sensing matrix is re-generated at each encryption. A similar analysis was also extended to the case of circulant sensing matrices in [5], where the authors characterize the increased measurements' information leakage due to the structured nature of the sensing matrix.

In fact, from all the above works, it has become evident that, because of the linearity of the sensing process, the measurements will always reveal *at least* the energy of the original signal. The best case, in which only the original signal's energy is leaked in the non-asymptotic case, is that of sensing matrices made of real-valued Gaussian i.i.d. entries. In order to overcome this problem, proposed solutions consider either to normalize the signal to unit energy as in [4] or to obfuscate the energy as in [6]. In this latter work, the authors propose a method to obfuscate the energy of the original signal through scalar multiplication, avoiding the encryption and transmission burden related to the energy of the plaintext. Interestingly, the authors show that this method also allows trusted parties to perform basic signal processing operations in the encrypted domain, i.e. anomaly detection.

The vast majority of works we summarized up to this point provide theoretical guarantees and solve practical problems considering sensing matrices whose entries are sampled with *infinite* precision. However, characterizing the properties of *finite* precision cryptosystems is a problem of paramount importance. In fact, practical systems such as IoT devices which are power constrained and also need to cope with data confidentiality can find in CS cryptosystems a good fit. Indeed, this practical scenario was considered only in a few works. As an example, in [3], [7], the authors also consider practical Bernoulli sensing matrices and prove their asymptotical spherical secrecy. The same class of sensing matrices was also considered in [8]. In this latter work, similarly to standard private key cryptosystems, modes of operation for compressive encryption are introduced which, along with the use of Bernoulli sensing matrices, make the considered scenario suitable for practical implementations.

It is important to note that, to the best of our knowledge, the works which account for more practical scenarios make use of Bernoulli sensing matrices, e.g. [3], [8]. Also in [9], where the focus is put on practical and secure sensing matrix generation schemes, the results are based on the assumption of using sensing matrices for which the secrecy has already been proven. While suitable for practical implementations, the secrecy of Bernoulli sensing matrices has only been proved in the asymptotic case [3]. This means that when the plaintext is not sufficiently large, then more information, other than the original signal's energy, is leaked. Moreover, finite alphabets larger than the Bernoulli one, have not been considered yet. Intuitively, as the size of the finite alphabet for the sensing matrix is increased, higher secrecy is expected. Based on this intuition, and since there is still a gap between the practical Bernoulli sensing approach and the more theoretical one based

on Gaussian random variables, we consider the case of sensing matrices made of quantized entries drawn from truncated Gaussian distributions.

In more detail, the scope of this paper is twofold: first, we show that the use of sensing matrices with quantized and truncated Gaussian entries not only is practically feasible, but can achieve significantly higher secrecy than Bernoulli distributed ones; second, we discuss how to choose the system parameters in order to achieve the desired secrecy level.

II. Methodology

The main contribution of this work is to experimentally show the relationships existing between the parameters of a quantized Gaussian sensing matrix and the secrecy of the related cryptosystem. In this section we start by introducing some background and definitions related to the secrecy of compressive cryptosystems. Next, we discuss and describe how the employed metrics can be obtained under the considered settings.

A. Background and definitions

At first, let us recall the CS acquisition process which can be modeled as $\mathbf{y} = \mathbf{\Phi x}$, where $\mathbf{x} \in \mathbb{R}^{n\times 1}$ is the K-sparse original signal (plaintext), $\mathbf{\Phi} \in \mathbb{R}^{m\times n}$ is the sensing matrix, and $\mathbf{y} \in \mathbb{R}^{m\times 1}$ is the measurements vector with $m << n$.

At this point, it is important to highlight that, since in this work we are considering a finite precision representation, both the sensing matrix entries $\mathbf{\Phi}$ and the measurements $\mathbf{y}$ need to be quantized. For what concerns the measurements, by the data processing inequality [10], it is possible to show that the measurements quantization does not decrease the system secrecy. Conversely, as we will show in the remainder of this work, if the quantization is performed on the sensing matrix entries, the system secrecy can be reduced. It becomes thus evident that the most critical aspect of employing finite precision is related to the representation of sensing matrix entries. For this reason, from now on we will focus on this latter aspect.

Before proceeding with our discussion it is important to underline another important aspect: the CS cryptosystem we will consider from now on is based on the OTS assumption and we assume that the plaintexts have been normalized to unit energy. These two assumptions will allow us to exclude known plaintext and ciphertext attacks as well as to decouple the information leakage through the measurements due to non-normalized signals as discussed in [5].

We start with the definition of the truncated and quantized Gaussian distribution. This distribution can be defined over a one-dimensional lattice $\Lambda = \{qz \ : q \in \mathbb{R}, z \in \mathbb{Z}\}$, where q corresponds to the quantization bin width. Thus, the quantized Gaussian distribution over a lattice Λ can be defined as

$$\mathcal{G}_{\Lambda,\sigma}(z) = \int_{z-q/2}^{z+q/2} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{\iota^2}{2\sigma^2}} d\iota \ \text{ with } \ z \in \Lambda. \quad (1)$$

It is important to note that we are taking into account the physical limitations of a practical system which employs finite precision representations. Since we obtain samples from a quantized Gaussian distribution, these limitations also translate into tails truncation. In more detail, given a fixed amount of bits N_b we consider to truncate the tails at T_R, then we have that $q = 2T_R/2^{N_b}$ is a function of the number of available bits. We can now define the truncated quantized Gaussian distribution as $\overline{\mathcal{G}}_{\Lambda^C,\sigma}(z) = v\mathcal{G}_{\Lambda^C,\sigma}(z)$, where $v = \frac{1}{1-g_T}$ is a normalization factor to assure $\overline{\mathcal{G}}_{\Lambda^C,\sigma}(z)$ to be a probability density function, with $g_T = 2\sum_{z=T_R}^{+\infty} \mathcal{G}_{\Lambda^C,\sigma}(z)$ $z \in \Lambda$.

Next, in order to measure the secrecy of the system we introduce the metric we will use in the remainder of this paper: the θ-distinguishability. As introduced in [5] this metric is defined by means of a detection experiment. Given two signals $\mathbf{x}_1, \mathbf{x}_2$ we consider a simple detection test in which the attacker, by using a detector $D(\mathbf{y})$, has to guess whether $\mathbf{y}$ comes from $p(\mathbf{y}|\mathbf{x}_1)$ or $p(\mathbf{y}|\mathbf{x}_2)$. Therefore, we will say that CS measurements are θ-indistinguishable if, for every possible detector $D(\mathbf{y})$, $P_d - P_f \leq \theta$, where P_d and P_f are the probability of detection and false alarm of the detector, respectively. It is evident that $\theta = 0$ corresponds to perfect secrecy, namely no detector can distinguish the two signals. We now recall that, with a slight abuse of notation, according to Lemma 4 in [5] and the Pinker's inequality, CS measurements are at least $\delta_{\mathrm{KL}}(p(\mathbf{y}|\mathbf{x}_1), p(\mathbf{y}|\mathbf{x}_2))$-indistinguishable w.r.t. $\mathbf{x}_1, \mathbf{x}_2$ given that $P_d - P_f \leq \sqrt{1/2\,\delta_{\mathrm{KL}}(p(\mathbf{y}|\mathbf{x}_1), p(\mathbf{y}|\mathbf{x}_2))}$, where $\delta_{\mathrm{KL}}(\cdot,\cdot)$ corresponds to the KL divergence.

B. Methods

Here we derive the probability distributions which are needed to compute the θ-distinguishability.

The metric we will use to characterize the θ-distinguishability is the KL divergence. In fact, as shown above, it upper bounds the θ-distinguishability of any possible detector. To proceed with this characterization, we need to compute the conditional probability of having a specific measurement given the signal $\mathbf{x}$ which we denote as $p(\mathbf{y}_i|\mathbf{x})$. In order to find $p(\mathbf{y}_i|\mathbf{x})$, we can notice that $\mathbf{y}_i$ is a linear combination of n sensing matrix entries $\boldsymbol{\phi}$ and thus, its characteristic function can be written in product fashion as

$$\phi_{\mathbf{y}_i|\mathbf{x}}(t) = \prod_{k=1}^{n} \tilde{\phi}(\mathbf{x}_k t), \quad (2)$$

where $\tilde{\phi}(t)$ is the characteristic function of a truncated and quantized Gaussian distribution. According to [11], the characteristic function of a truncated Gaussian distribution whose realizations are quantized through area sampling can be written as $\tilde{\phi}(t) = \sum_{l=-\infty}^{+\infty} \phi_T(t + l\Psi)\mathrm{sinc}\left(\frac{q(t+l\Psi)}{2}\right)$, where $\Psi = \frac{2\pi}{q}$, $\phi_T(t)$ is the characteristic function of a truncated Gaussian distribution and q is the width of the quantization bin. Lastly, we can write $\phi_T(t) = \phi_G(t) * \frac{\sin(2T_R t)}{2T_R}$, where $2T_R$ is the truncation interval and $\phi_G(t)$ is the characteristic function of a Gaussian distribution. At this point, we can compute $p(\mathbf{y}_i|\mathbf{x})$ for a given $\mathbf{x}_1$ and $\mathbf{x}_2$ by using (2) and performing the inverse Fourier transform. Lastly, as done in [4], though hard to be obtained analytically the value of $P_d - P_f$ can be upper bounded by the KL divergence between $p(\mathbf{y}_i|\mathbf{x}_1)$ and $p(\mathbf{y}_i|\mathbf{x}_2)$ as

$$P_d - P_f \leq \sqrt{\frac{m}{2}\min(\delta_{\mathrm{KL}}([\mathbf{\Phi x}_1]_i, [\mathbf{\Phi x}_2]_i), \delta_{\mathrm{KL}}([\mathbf{\Phi x}_2]_i, [\mathbf{\Phi x}_1]_i))}, \quad (3)$$

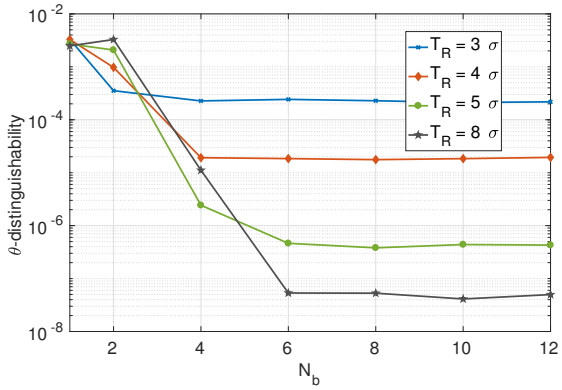

Fig. 1. Effects of the quantization bits N_b on the θ-distinguishability for different values of T_R

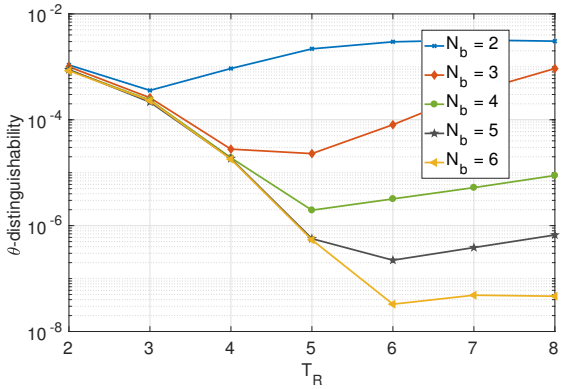

Fig. 3. Effects of the Gaussian truncation width T_R on the θ-distinguishability for different values of N_b.

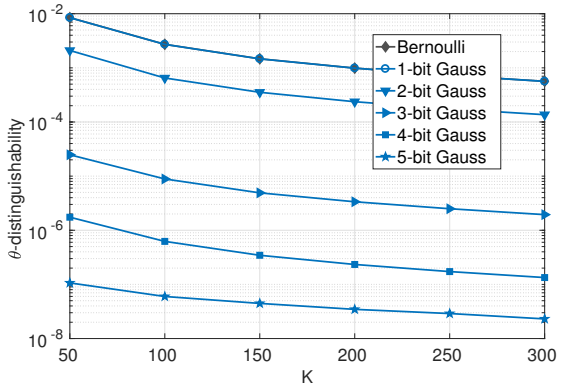

Fig. 2. Effects of the quantization bits N_b and sparsity K on the θ-distinguishability.

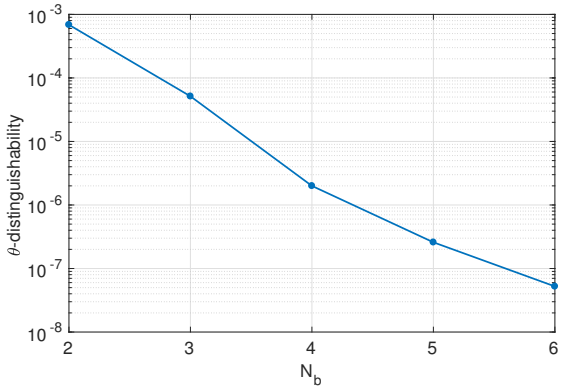

Fig. 4. Effects of the regime for both N_b and T_R on the θ-distinguishability for increasing values of N_b.

where the KL divergences can be evaluated numerically.

III. Experimental results

The experiments we perform aim to show the relationships existing between the θ-distinguishability and the system parameters, namely truncation factor T_R and quantization bits N_b. We consider these parameters at first separately, then jointly in order to correctly identify how their values affect the system secrecy.

Before starting with our characterization of the system parameters, it is important to underline that, despite being both practical alternatives, quantized Gaussian sensing significantly improves the secrecy with respect to Bernoulli sensing. This can be seen in Fig. 2 where the curve which corresponds to both 1-bit Gaussian and Bernoulli sensing is the one which achieves highest θ-distinguishability (lowest secrecy). This strongly motivates our interest to deeply investigate the relationships existing between the quantization parameters and the system secrecy.

In the first experiment we show the effects of the parameter N_b at different truncation levels T_R. The variance of the Gaussian distribution is fixed to $\sigma^2 = 1$, and we vary the number of bits employed for the quantization. The quantization bins are considered to be equally spaced in the truncation interval. In more detail, we consider 1000 signals $\mathbf{x}_1$ and $\mathbf{x}_2$ uniformly distributed on a unit norm sphere of size $n = 1000$ and $K = 100$ non-zero entries. This means that we consider increasingly small quantization bins of width $q = 2T_R/2^{N_b}$. The metric we consider is the KL distance, in more detail we show the value of the upper bound on the distinguishability as defined in (3). The result is depicted in Figure 1.

The first and most important result which can be immediately noted is that, as a general behavior, the θ-distinguishability exponentially decreases as the number of quantization bits N_b increases until a plateau is reached. This behavior, can be seen in Fig. 1 and in a following experiment in Fig. 2. This is an expected result since as the alphabet becomes larger, the quantized distribution becomes closer to continuous Gaussian distribution. Moreover, a linear increase in the number of bits results in an exponential increase in the alphabet size. Thus, according to the experimental results in Fig. 1-2 we can conjucture that the relationship existing between θ and N_b has the form $\theta \propto 2^{-N_b}$. This result is also confirmed by some preliminary theoretical findings which will be revealed in a forthcoming paper. This is a very important result, since it means that a limited number of bits, employed for the quantization of the sensing matrix, can suffice to achieve high secrecy. However, we need to consider that the exponential behavior cannot completely describe the results. As can be seen in Fig. 1 the parameters N_b and T_R can affect the θ-distinguishability in a joint fashion. When considering larger truncation intervals, and the number of employed bits is smaller, the value of θ is increased, namely the secrecy is

decreased. Conversely, even if the number of employed bits is higher, a shorter truncation interval can indeed limit the exponential decrease of θ which reaches a plateau. This can be explained by the fact that even though N_b is increased and thus the distribution of the measurements approaches the continuous Gaussian one, in practice it is limited by a small truncation interval. This joint effect will be explained more in detail when considering the effects of T_R.

In the next experiment, we jointly consider the effects of N_b and the sparsity of the original signal K. While this latter value is not a system parameter, it is interesting to consider how an increased number of linear combinations in the sensing process can affect the system secrecy. As for the previous experiment, we keep all the parameters fixed except for N_b and K. The KL divergence bound is shown in Figure 2.

Again, as in Fig. 1, it is possible to note that the value of θ exponentially decreases with N_b. Furthermore, it can also be noted that, as the sparsity K increases, the value of θ decreases. However, in this case the existing relationship seems to be a linear one as $\theta \propto \alpha/K$. This behavior can be explained through the central limit theorem. In fact, since the entries of the sensing matrices are i.i.d., as the size of the linear combination of i.i.d. elements increases, the result will tend to a Gaussian distribution. In the case of the limit $K \to \infty$, two signals will result in equally distributed measurements and thus they will become indistinguishable. However, while the sparsity parameter K can directly affect the secrecy of the system it is not a system design parameter, and its effects are negligible with respect to those due to N_b. For these reasons, the other experiments we performed consider a fixed value of $K = 100$.

For the third experiment, we change our focus. In fact, we consider the effects of the truncation width T_R. In more detail, we compute the KL divergence bound for different truncation intervals when employing a different number of quantization bits. The experiment settings are kept as described above: KL divergences are averaged over 1000 realizations of $K = 100$ sparse signals lying on a unit sphere. As can be seen in Fig. 3 and already pointed out previously, the effects of the truncation are paired with those due to the quantization. In fact, a larger truncation interval favors the secrecy of cryptosystems which make use of high N_b allowing to reach their optimal secrecy. Conversely, larger truncation intervals worsen the secrecy of those cryptosystems making use of coarse quantization. Thus, also in the light of the previous results, we expect that θ can be expressed as

$$\theta \propto f(N_b, T_R, K) = \frac{\alpha T_R^{\beta}}{K 2^{N_b}}, \tag{4}$$

where α, β are hyper-parameters. This function is increasing in T_R for fixed values of N_b and, is exponentially decreasing in N_b when T_R is fixed. Moreover, it is easy to see that the ratio between these two quantities must be carefully chosen in order to maximize the secrecy of a cryptosystem given the implementation constraints.

This lead us to conjecture that there exists a regime for N_b and T_R for which the value of θ achieves its minimum and thus decreases exponentially. The regime implies a relationship between N_b and T_R which is expected to be in the form $T_R \propto \alpha_1 N_b^{\beta_1}$. The effect of this regime on the distinguishability is shown in Fig. 4 where we considered a value of $\alpha_1 = 1.25$ and $\beta_1 = 1$. The result is that, if these two parameters are chosen accordingly, the value of θ decreases exponentially as N_b increases.

IV. Conclusion

In this paper, we showed that the practical compressive cryptosystems based on quantized Gaussian distributions can achieve high secrecy levels and outperform those based on Bernoulli sensing. We experimentally described how the quantization parameters affect the secrecy of a cryptosystem and showed that exists a regime which allows to exponentially increase the secrecy of the cryptosystem linearly with the number of employed bits. Our future works will prove, from a theoretical perspective, the relationship between secrecy and representation precision we experimentally demonstrated in this paper.

Acknowledgment

This work results from the research cooperation with the Sony Technology Center Stuttgart (Sony EuTEC). We would especially like to thank Lev Markhasin and Oliver Erdler from Sony EuTEC for their valuable feedback.

References

[1] D. L. Donoho, "Compressed sensing," *IEEE Transactions on information theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[2] Y. Rachlin and D. Baron, "The secrecy of compressed sensing measurements," in *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*. IEEE, 2008, pp. 813–817.

[3] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "Low-complexity multiclass encryption by compressed sensing," *IEEE transactions on signal processing*, vol. 63, no. 9, pp. 2183–2195, 2015.

[4] T. Bianchi, V. Bioglio, and E. Magli, "Analysis of one-time random projections for privacy preserving compressed sensing," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 2, pp. 313–327, 2016.

[5] T. Bianchi and E. Magli, "Analysis of the security of compressed sensing with circulant matrices," in *Information Forensics and Security (WIFS), 2014 IEEE International Workshop on*. IEEE, 2014, pp. 173–178.

[6] M. Testa, T. Bianchi, and E. Magli, "Energy obfuscation for compressive encryption and processing," in *Information Forensics and Security (WIFS), 2017 IEEE International Workshop on*. IEEE, 2017.

[7] V. Cambareri, J. Haboba, F. Pareschi, H. R. Rovatti, G. Setti, and K. w. Wong, "A two-class information concealing system based on compressed sensing," in *2013 IEEE International Symposium on Circuits and Systems (ISCAS2013)*, May 2013, pp. 1356–1359.

[8] R. Fay, "Introducing the counter mode of operation to compressed sensing based encryption," *Information Processing Letters*, vol. 116, no. 4, pp. 279–283, 2016.

[9] R. A. Djeujo and C. Ruland, "Secure matrix generation for compressive sensing embedded cryptography," in *2016 IEEE 7th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*, Oct 2016, pp. 1–8.

[10] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.

[11] B. Widrow, I. Kollar, and M.-C. Liu, "Statistical theory of quantization," *IEEE Transactions on instrumentation and measurement*, vol. 45, no. 2, pp. 353–361, 1996.