

On Known-Plaintext Attacks to a Compressed Sensing-Based Encryption: A Quantitative Analysis

Original

On Known-Plaintext Attacks to a Compressed Sensing-Based Encryption: A Quantitative Analysis / Cambareri, Valerio; Mangia, Mauro; Pareschi, Fabio; Rovatti, Riccardo; Setti, Gianluca. - In: IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY. - ISSN 1556-6013. - STAMPA. - 10:10(2015), pp. 2182-2195.
[10.1109/TIFS.2015.2450676]

Availability:

This version is available at: 11583/2696643 since: 2020-02-05T22:47:34Z

Publisher:

Institute of Electrical and Electronics Engineers Inc.

Published

DOI:10.1109/TIFS.2015.2450676

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

©2015 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)

On Known-Plaintext Attacks to a Compressed Sensing-Based Encryption: A Quantitative Analysis

Valerio Cambareri, *Student Member, IEEE*, Mauro Mangia, *Member, IEEE*,

Fabio Pareschi, *Member, IEEE*, Riccardo Rovatti, *Fellow, IEEE*, Gianluca Setti, *Fellow, IEEE*

Abstract—Despite the linearity of its encoding, compressed sensing may be used to provide a limited form of data protection when random encoding matrices are used to produce sets of low-dimensional measurements (ciphertexts). In this paper we quantify by theoretical means the resistance of the least complex form of this kind of encoding against known-plaintext attacks. For both standard compressed sensing with antipodal random matrices and recent multiclass encryption schemes based on it, we show how the number of candidate encoding matrices that match a typical plaintext-ciphertext pair is so large that the search for the true encoding matrix inconclusive. Such results on the practical ineffectiveness of known-plaintext attacks underlie the fact that even closely-related signal recovery under encoding matrix uncertainty is doomed to fail.

Practical attacks are then exemplified by applying compressed sensing with antipodal random matrices as a multiclass encryption scheme to signals such as images and electrocardiographic tracks, showing that the extracted information on the true encoding matrix from a plaintext-ciphertext pair leads to no significant signal recovery quality increase. This theoretical and empirical evidence clarifies that, although not perfectly secure, both standard compressed sensing and multiclass encryption schemes feature a noteworthy level of security against known-plaintext attacks, therefore increasing its appeal as a negligible-cost encryption method for resource-limited sensing applications.

Index Terms—Compressed sensing, encryption, security, secure communications

I. INTRODUCTION

THIS paper elaborates on the possibility of exploiting Compressed Sensing (CS) [1], [2] not only to reduce the resource requirements for signal acquisition, but also to protect the acquired data so that their information is hidden from unauthorised receivers. A number of prior analyses [3]–[7] show that, although the encoding performed by CS cannot be regarded as perfectly secure, practical encryption is still provided at a very limited cost, either at the analog-to-digital interface or immediately after it, in early digital-to-digital processing stages.

Such a lightweight encryption scheme may be particularly beneficial to acquisition systems within the framework of

wireless sensor networks [8] where large amounts of data are locally acquired by sensor nodes with extremely tight resource budgets, and afterwards transmitted to a remote node for further processing. When the security of these transmissions is an issue, low-resource techniques that help balancing the trade-off between encryption strength and computational cost may offer an attractive design alternative to the deployment of separate conventional encryption stages.

An encryption scheme based on CS leverages the fact that, in its framework, a high-dimensional signal is encoded by linear projection on a random subspace, thus producing a set of low-dimensional measurements. These can be mapped back to the acquired signal only under prior assumptions on its *sparsity* [9] and a careful choice of random subspaces such as those defined by antipodal random (also known as Bernoulli random [10], [11]) encoding matrices. In addition, suitable *sparse signal recovery algorithms* [12]–[14] are required to decode the original signal. These must be applied with an exact knowledge of the subspace on which the signal was projected. In complete absence of this information the acquired signal is unrecoverable. Hence, this subspace may be generated from a shared secret between the transmitter and intended receivers that enables their high-quality signal recovery.

If, on the other hand, the above subspace is only *partially* known, a low-quality version of the signal may be recovered from its measurements, with a degradation that increases gracefully with the amount of missing information on the projection subspace. By exploiting this effect, multiclass encryption schemes were devised [5], [7] in which *high-class users* are able to decode high-quality information starting from a complete knowledge of the shared secret, while *lower-class users* only recover a low-quality approximation of the acquired signal starting from partial knowledge of the secret. In order to take full advantage of this scheme, its security must be quantitatively assessed against potential cryptanalyses. The theoretical and empirical evidence provided in [7] dealt with statistical attacks on the measurements produced by universal random encoding matrices [10].

In this paper we address the resistance of an embodiment of CS against Known-Plaintext Attacks (KPA), *i.e.*, in threatening situations where a malicious eavesdropper has gained access to an instance of the signal (plaintext) and its corresponding random measurements (ciphertext), and from this information tries to infer the corresponding instance of an antipodal random encoding matrix. KPAs are more threatening than attacks solely based on observing the ciphertext. Yet, we will show how both simple and multiclass encryption based on

Copyright ©2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

V. Cambareri and R. Rovatti are with the Department of Electrical, Electronic and Information Engineering (DEI), University of Bologna, Italy (e-mail: valerio.cambareri@unibo.it, riccardo.rovatti@unibo.it).

M. Mangia is with the Advanced Research Center on Electronic Systems (ARCES), University of Bologna, Italy (e-mail: mmangia@arces.unibo.it).

F. Pareschi and G. Setti are with the Engineering Department in Ferrara (ENDIF), University of Ferrara, Italy (e-mail: fabio.pareschi@unife.it, gianluca.setti@unife.it).

CS exhibit a noteworthy level of resistance against this class of attacks due to the nature of the encoding.

The paper is organised as follows. In Section II we briefly review the fundamentals of CS and multiclass encryption in the two-class case, which distinguishes between *first-class receivers* authorised to reconstruct the signal with full quality and *second-class receivers* with reduced decoding quality.

Section III describes KPAs as delivered both by eavesdroppers and second-class receivers who aim at improving the quality of their signal recovery. There, it is shown that the expected number of candidate solutions matching a plaintext-ciphertext pair is enormous, thus implying that finding the true encoding matrix among such a huge solution set is practically infeasible. To extend this analysis, we also attack the two-class encryption scheme by using recovery algorithms that compensate encoding matrix perturbations [15], [16] as suffered by a second-class receiver. Their performances are shown to be equal to a standard decoding algorithm [13] that does not attempt such compensation, *i.e.*, that legitimately recovers the acquired signal at the prescribed quality level.

In Section IV the previous KPAs are exemplified for electrocardiographic tracks (ECG) and images containing sensitive identification text. For all these cases we give empirical evidence on how, even in favourable attack conditions, the encoding matrices produced by KPAs perform poorly when trying to decode any further ciphertext. Theoretical and empirical evidence allows us to conclude that compressed sensing-based encryption, albeit not perfectly secure [3], provides some security properties and defines a framework in which their violation is non-trivial. The Appendices report the proofs of the Propositions and Theorems given in Section III.

A. Relation to Prior Work

To prove how CS and multiclass encryption provide a satisfying level of privacy even against informed attacks, this work addresses the problem of finding all the instances of an antipodal random encoding matrix that map a known plaintext to the corresponding ciphertext, when both quantities are deterministic and digitally represented. Our analysis hinges on the connection between linear encoding by antipodal random matrices, the subset-sum problem [17] and its expected number of solutions [18]. While the authors of [3] proved how CS lacks perfect secrecy in the Shannon sense [19], both [3] and [4] contrasted this with computational security evidence substantially based on brute-force attacks. Our improvement in the specific, yet practically important case of antipodal random encoding matrices is in that our analysis predicts how the expected number of candidate solutions to a KPA varies with the plaintext dimensionality and its digital representation.

In addition, we evaluate specific attacks to multiclass encryption by CS in the case of lower-class users attempting to upgrade their recovery quality. To assess the resistance of this strategy against KPAs, we apply a similar theoretical analysis. Then, we extend the attacks to include sparse signal recovery under matrix uncertainty [15], [16] based on the idea that missing information [20], perturbations [21], [22] and basis mismatches [23] could be partially compensated, although we

verify that is not the case with the random perturbation entailed by multiclass encryption.

II. MULTICLASS ENCRYPTION BY COMPRESSED SENSING

A. A Brief Review of Compressed Sensing

The encryption schemes we consider in this paper are based on Compressed Sensing (CS) [1], [2], a mathematical framework in which a signal represented by a vector $x \in \mathbb{R}^n$ is acquired by applying a linear, dimensionality-reducing transformation $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ (*i.e.*, the *encoding matrix*) to generate a vector of *measurements* $y = Ax, y \in \mathbb{R}^m, m < n$. To enable the recovery of x given y , CS leverages the fact that x is known to be *sparse* in a proper basis D , *i.e.*, for any instance of x its representation is $x = Ds$ where $s \in \mathbb{R}^n$ has a number of non-zero entries at most $k \ll n$. The results presented in this paper are independent of D , which we consider an orthonormal basis for the sake of simplicity. In addition, the encoding matrix A must obey some information-preserving guarantees [24], [25] that we assume verified throughout this paper and essentially impose that $m = \mathcal{O}(k \log n)$. The most relevant fact here is that when A is a typical realisation of a random matrix with independent and identically distributed (*i.i.d.*) entries following a subgaussian distribution [26] we are reassured that signal recovery is possible regardless of the chosen basis D . In fact, some signal recovery algorithms exist for which guarantees can be given with very high probability [12] along with an ever-growing plethora of fast iterative methods capable of reconstructing x starting from y , A and D . An essential decoding scheme is the convex optimisation problem known as *basis pursuit with denoising*,

$$\hat{x} = \arg \min_{\xi \in \mathbb{R}^n} \|D^{-1}\xi\|_1 \quad \text{s.t.} \quad \|A\xi - y\|_2 \leq \omega \quad (\text{BPDN})$$

where the ℓ_1 -norm in the objective function promotes the sparsity of \hat{x} with respect to D , while the ℓ_2 -norm constraint enforces its fidelity to the measurements up to a threshold $\omega \geq 0$ that accounts for noise sources. In particular, we here concentrate on operators $A \in \{-1, 1\}^{m \times n}$ that are realisations of an antipodal random matrix with *i.i.d.* entries and equiprobable symbols $\{-1, 1\}$ [10]; such matrices are known to verify the above guarantees, and are remarkably (*i*) simple, and therefore suitable to be generated, implemented and stored in digital devices (*ii*) random in nature, thus suggesting the possibility of exploiting such randomness to generate an encryption mechanism using the linear encoding scheme of CS. Due to their limited set of possible symbols $\{-1, 1\}$, such antipodal random matrices are more easily subject to cryptanalysis; for this reason, we tackle them as a baseline for those defined by a larger set of symbols.

B. Security and Two-Class Encryption by Compressed Sensing

1) *A Security Perspective*: the knowledge of A is necessary in the recovery of x from y , since any error in its entries reflects on the quality of the recovered signal [21]. A number of security analyses leveraging this fundamental fact were introduced [3], [4], [7] in which CS is regarded as a symmetric encryption scheme, where the *plaintext* x is mapped to the

ciphertext y by means of the linear transformation operated by A , *i.e.*, the *encryption algorithm*. The ciphertext is then stored or transmitted, and its intended receivers may decrypt x by knowing y , the sparsity basis D , and by having a prior agreement on the *encryption key* or *shared secret* that is necessary to reproduce A .

The ideal requirement for a secure application of CS (as noted in [3], [27]) is that any encoding matrix instance is used for at most one plaintext-ciphertext pair; this implies the use of a potentially *infinite* sequence of encoding matrices $\{A^{[t]}\}_{t \in \mathbb{N}}$. In violation of this *non-repeatability* hypothesis, each $A^{[t]}$ could be simply recovered by collecting n linearly independent plaintext-ciphertext pairs related by it, *i.e.*, by solving a linear system of equations with the mn entries of $A^{[t]}$ as the unknowns.

In practice, the encoding matrices are obtained by algorithmic expansion of the shared secret, *e.g.*, by using the key as the seed of a pseudo-random number generator (PRNG) which outputs a reproducible bitstream. Due to its deterministic and finite-state nature, this stream yields a *periodic* sequence of encoding matrices $\{A^{[t \bmod P]}\}_{t \in \mathbb{N}}$ repeating with period P , where each $A^{[t]}$ is obtained by mapping mn distinct bits to antipodal symbols.

Thus, the non-repeatability hypothesis will be granted by a system-level choice of an encryption key and PRNG that makes P large enough to exceed any reasonable observation time.

However, such pseudo-random bitstreams may themselves be vulnerable to cryptanalysis if a few of their bits are exposed. As a simple example of this threat, assume that the encoding matrices are generated by a maximal-length shift register sequence [28, Chapter 4], for which a B_{key} bit seed grants $P = \lfloor \frac{2^{B_{\text{key}}} - 1}{mn} \rfloor$. Regrettably, such a sequence is easily cryptanalysed from only $2B_{\text{key}}$ of its bits by the well-known Berlekamp-Massey algorithm [29].

Hence, a successful KPA that retrieves even part of an encoding matrix, *e.g.*, one of its rows, may expose just enough information to reveal the key and therefore break a CS-based encryption. To contrast this type of threat, our analysis shows how KPAs are incapable of revealing missing information on the true encoding matrices, whose symbols remain undetermined.

2) *Two-Class Encryption*: in an extended version of this encryption framework, *i.e.*, *two-class encryption* by CS [5], [7], we consider a first sequence of matrices $\{A^{(0),[t]}\}_{t \in \mathbb{N}}$, $A^{(0),[t]} \in \{-1, 1\}^{m \times n}$ obtained by pseudo-random expansion of a seed $\text{Key}(A^{(0)})$. In parallel, a sequence of index pair sets $\{C^{(0),[t]}\}_{t \in \mathbb{N}}$, $C^{(0),[t]} \subset \{0, \dots, m-1\} \times \{0, \dots, n-1\}$ is obtained by pseudo-random expansion of a seed $\text{Key}(C^{(0)})$. We then generate a second sequence of matrices $\{A^{(1),[t]}\}_{t \in \mathbb{N}}$ whose elements $A^{(1),[t]}$ are obtained by combining $A^{(0),[t]}$, $C^{(0),[t]}$ as

$$A_{j,l}^{(1),[t]} = \begin{cases} A_{j,l}^{(0),[t]} & \text{if } (j, l) \notin C^{(0),[t]} \\ -A_{j,l}^{(0),[t]} & \text{if } (j, l) \in C^{(0),[t]} \end{cases} \quad (1)$$

with $C^{(0),[t]}$ indicating which entries of $A^{(0),[t]}$ must be sign-flipped to obtain $A^{(1),[t]}$, that is then used to encode x into

y . Thus, we consider a cardinality c for every $C^{(0),[t]}$, define $\eta = c/mn$ the sign flipping density, and let $A^{(0)}$, $A^{(1)}$, $C^{(0)}$ be generic, unique random matrix instances (that is, the matrix sequences will be implicitly considered from now on). Given any plaintext x , the corresponding ciphertext y is produced as $y = A^{(1)}x$, $A^{(1)}$ being the *true encoding matrix*. Two-class encryption is then achieved by distributing $\text{Key}(A^{(0)})$ to all authorised receivers and $\text{Key}(C^{(0)})$ only to first-class receivers. In fact, when y is communicated, receivers knowing both $\text{Key}(A^{(0)})$ and $\text{Key}(C^{(0)})$ are able to rebuild the corresponding $A^{(1)}$ used in the encoding and reconstruct x with full quality by solving BPDN with $\omega = 0$.

On the other hand, second-class receivers may only rebuild $A^{(0)}$ from their available information. For $0 < \eta \ll 1$ such a matrix is an approximation of the corresponding $A^{(1)}$, thus allowing signal recovery with lower quality than that achieved by first-class receivers. Furthermore, any receiver not knowing $\text{Key}(A^{(0)})$ has no information on the encoding matrix and is consequently unable to recover x , which remains encrypted.

In [7] we have characterised the effectiveness of this scheme by showing how eavesdroppers trying to compensate their ignorance of the key by means of straightforward statistical analysis of y are presented with approximately Gaussian-distributed ciphertexts (converging with rate $\mathcal{O}(n^{-1})$). In addition, if $A^{(0)}$ is an antipodal random matrix, the same can be said of $A^{(1)}$ since the statistics of its equiprobable symbols are unaltered by $C^{(0)}$ used to build the latter from the former. Hence, the ciphertext is statistically indistinguishable from the one that could be produced by encoding the same plaintext with $A^{(0)}$ instead of $A^{(1)}$, and second-class users will also be unable to exploit the statistical properties of y .

C. Signal Models and Assumptions

Since the attacks we present rely on deterministic knowledge of x and y , we assume throughout the paper that both plaintexts and ciphertexts are represented by digital words. For simplicity, we let $x = \{x_l\}_{l=0}^{n-1}$ be such that $x_l \in \{-L, \dots, -1, 0, 1, \dots, L\}$ for some integer $L > 0$. Note that the number of bits representing the plaintext in this fashion is at least $B_x = \lceil \log_2(2L + 1) \rceil$, so we may assume B_x is less than a few tens in typical embodiments (actually, $B_x \leq 32$ bit in typical signal processing applications). Consequently, the ciphertext will be represented by $\{y_l\}_{l=0}^{m-1}$, where each y_l is quantised with $B_y = B_x + \lceil \log_2 n \rceil$ bit that avoid any information loss.

III. KNOWN-PLAINTEXT ATTACKS

In view of quantifying the resistance of this scheme to threatening cryptanalyses, we now consider situations in which an attacker gains access to a given, exact value of the plaintext x corresponding to a ciphertext y . Based on this knowledge, the attacker aims at computing the true encoding $A^{(1)}$ such that $y = A^{(1)}x$. In the following we will consider a KPA by assuming that only one (x, y) pair is known for a certain $A^{(1)}$, consistently with the hypothesis that $A^{(1)}$ is never reused in the encoding (as detailed in Section II-B1). This type of attack gives rise to different strategies (see Fig. 1) whether




Fig. 1. A two-class encryption scheme and the known-plaintext attacks being analysed from an eavesdropper (Eve) and a second-class user (Steve).

the attacker knows nothing except the (x, y) pair (a pure eavesdropper, Eve) or it is a second-class receiver knowing also the partially correct encoding $A^{(0)}$ and attempting to complete its knowledge of $A^{(1)}$ (we will call this malicious second-class user Steve and its KPA a *class-upgrade*).

For the sake of simplicity, both KPAs are here characterised on a single row¹ of $A^{(1)}$, while a complete KPA will entail m of such attacks. Furthermore, we note that the analysis is carried out in full compliance with Kerckhoffs's principle [30], *i.e.*, the only information that the attackers are missing is their respective part of the encryption key, while any other detail on the sparsity basis, as well as two-class encryption specifications is here regarded as known.

A. Eavesdropper's Known-Plaintext Attack

Given a plaintext x and the corresponding ciphertext $y = A^{(1)}x$ we now assume the perspective of Eve and attempt to recover $A_j^{(1)}$ with a set of antipodal symbols $\hat{A}_j^{(1)} = \{\hat{A}_{j,l}^{(1)}\}_{l=0}^{n-1}$ such that

$$y_j = \sum_{l=0}^{n-1} \hat{A}_{j,l}^{(1)} x_l \quad (2)$$

Moreover, to favour the attacker² we assume all $x_l \neq 0$. We now introduce a combinatorial optimisation problem at the core of the analysed KPAs.

Definition 1 (Subset-Sum Problem). Let $\{u_l\}_{l=0}^{n-1}, u_l \in \{1, \dots, L\} \subset \mathbb{N}_+$ and $v \in \mathbb{N}_+$. We define *subset-sum problem* (SSP) [17, Chap. 4] the problem of assigning n binary variables $b_l \in \{0, 1\}, l = 0, \dots, n-1$ such that

$$v = \sum_{l=0}^{n-1} b_l u_l \quad (3)$$

We define *solution* any $\{b_l\}_{l=0}^{n-1}$ verifying (3). With the above definitions, the *density* of this problem is defined as [31]

$$\delta(n, L) = \frac{n}{\log_2 L} \quad (4)$$

Although in general a SSP is NP-complete, not all of its instances are equally *hard*. In fact, it is known that *high-density* instances (*i.e.*, $\delta(n, L) > 1$) have plenty of solutions found or

¹We denote with A_j the j -th row of a matrix A .

²If any $x_l = 0$ each corresponding summand would give no contribution to the sum (2), thus making $\hat{A}_{j,l}^{(1)}$ an undetermined variable in the attack.




Fig. 2. Sample average of the number of solutions for Eve's KPA compared to the theoretical value of (5) for $L = 10^4$.

approximated by, *e.g.*, dynamic programming, whereas *low-density* instances are typically hard, although for special cases polynomial-time algorithms have been found [31]. Moreover, such low-density hard SSP instances have been used in cryptography to develop the family of public-key knapsack cryptosystems [32], [33] although most have been broken with polynomial-time algorithms [34].

Proposition 1 (Eve's KPA). *The KPA to $A_j^{(1)}$ given (x, y) is equivalent to a SSP where each $u_l = |x_l|$, the variables $b_l = \frac{1}{2}(\text{sign}(x_l) \hat{A}_{j,l}^{(1)} + 1)$ and the sum $v = \frac{1}{2}(y_j + \sum_{l=0}^{n-1} |x_l|)$. This SSP has a true solution $\{\bar{b}_l\}_{l=0}^{n-1}$ that is mapped to the row $A_j^{(1)}$, and other candidate solutions that verify (3) but correspond to matrix rows $\hat{A}_j^{(1)} \neq A_j^{(1)}$.*

This mapping is explained in Appendix A, and we define $(x, y, A_j^{(1)})$ a *problem instance*. In our case we see that the density (4) is high since n is large and $\log_2 L$ is fixed by the digital representation of x (*e.g.*, so that $B_x \leq 64$). We are therefore operating in a region in which a solution of the SSP (3) is typically found in polynomial time. In fact, the resistance of the analysed embodiment of CS against KPAs is not due to the hardness of the corresponding SSP but, as we show below, to the huge number of candidate solutions as n increases, among which an attacker should find the only true solution to guess a single row of $A^{(1)}$. Since no *a priori* criterion exists to select them, we consider them *indistinguishable*. The next Theorem³ calculates the expected number of candidate solutions to Eve's KPA by applying the theory developed in [18].

Theorem 1 (Expected number of solutions for Eve's KPA). *For large n , the expected number of candidate solutions of the KPA in Proposition 1, in which (i) all the coefficients $\{u_l\}_{l=0}^{n-1}$ are i.i.d. uniformly drawn from $\{1, \dots, L\}$, and (ii) the true solution $\{\bar{b}_l\}_{l=0}^{n-1}$ is drawn with equiprobable and independent*

³ $\stackrel{n \rightarrow \infty}{\simeq}$ denotes asymptotic equality as $n \rightarrow \infty$.

binary values, is

$$\mathcal{S}_{\text{Eve}}(n, L) \stackrel{n \rightarrow \infty}{\approx} \frac{2^n}{L} \sqrt{\frac{3}{\pi n}} \quad (5)$$

The proof of Theorem 1 is given in Appendix A. This result (as well as the whole statistical mechanics framework from which it is derived) gives no hint on how much (5) is representative of finite- n behaviours. To compensate for that, we here enumerate by means of the binary programming solver in CPLEX [35] all the solutions to several small- n problem instances of Proposition 1 and verify that, even non-asymptotically, the expression (5) can be used to effectively estimate the expected number of candidate solutions to Eve's KPA. Such numerical evidence is reported in Fig. 2, where the sample average of the number of solutions $\hat{\mathcal{S}}_{\text{Eve}}(n, L)$ to 50 randomly generated problem instances with $L = 10^4$ and $n = 16, \dots, 32$ is plotted and compared with (5).

The remarkable matching observed therein allows us to estimate, for example, that a KPA to the encoding of a grayscale image of $n = 64 \times 64$ pixel quantised with $B_x = 8$ bit (unsigned, *i.e.*, $L = 128, n = 4096$) would have to discriminate on the average between $1.25 \cdot 10^{1229}$ equally good candidate solutions for each of the rows of the encoding matrix. This number is not far from the total possible rows, $2^{4096} = 1.04 \cdot 10^{1233}$. Hence, any attacker using this strategy is faced with a deluge of candidate solutions, from which it would choose one presumed to be exact to attempt a guess on a single row of $A^{(1)}$.

A legitimate concern when the attacker is presented with such a set of solutions is that most of them could be good approximations of the true encoding matrix row $A_j^{(1)}$. To see whether this is the case, we quantify the difference between $A_j^{(1)}$ and the corresponding candidates $\hat{A}_j^{(1)}$ resulting from a KPA in terms of their Hamming distance, *i.e.*, as the number of entries in which they differ.

Theorem 2 (Expected number of solutions for Eve's KPA at Hamming distance h from the true one). *The expected number of candidate solutions at Hamming distance h from the true solution of the KPA in Proposition 1, in which (i) all the coefficients $\{u_l\}_{l=0}^{n-1}$ are i.i.d. uniformly drawn from $\{1, \dots, L\}$, (ii) the true solution $\{\hat{b}_l\}_{l=0}^{n-1}$ is drawn with equiprobable and independent binary values, is*

$$\mathcal{S}_{\text{Eve}}^{(h)}(n, L) = \binom{n}{h} \frac{P_h(L)}{2^h L^h} \quad (6)$$

where $P_h(L)$ is a polynomial in L whose coefficients are reported in Table I for $h = 2, \dots, 15$.

The proof of this Theorem and the derivation of Table I are reported in Appendix B. As before, we collect some empirical evidence that the expression (6) correctly anticipates the expected number of solutions at a given Hamming distance from the true one, noting that Theorem 2 holds for finite n . Figure 3 reports for $n = \{21, 23, \dots, 31\}$ the sample average, over the same 50 problem instances generated in the experimental evaluation of (5), of the number of solutions to Eve's KPA whose Hamming distance from the true one is a given value $h = \{2, \dots, 15\}$. This sample average is compared against




Fig. 3. Sample average of the number of solutions for Eve's KPA at Hamming distance h from the true one, compared to the theoretical value of (6) for $L = 10^4$ and $n = 21, 23, \dots, 31$.

the value predicted by (6) with the polynomial coefficients in Table I. The remarkable matching we observe allows us to estimate that, resuming the case of a grayscale image with $n = 4096, L = 128$, only $1.95 \cdot 10^{41}$ candidate solutions out of the average $1.25 \cdot 10^{1229}$ are expected to have a Hamming distance $h \leq 16$, while $6.33 \cdot 10^{76}$ attain a Hamming distance $h \leq 32$. Since these results apply to each row of the matrix being inferred, this indicates how the chance that a randomly chosen candidate solution is close to the true one is negligible.

Under repeated threat of Eve's KPA, a system-level perspective would impose a change of encryption key (*i.e.*, of encoding matrix sequence) whenever the probability of failure of repeated KPAs, p_{fail} , drops below a desired security level $\zeta \in (0, 1)$, *i.e.*, at any time $p_{\text{fail}} \geq \zeta$. Some insight on the *encryption key lifetime* T that guarantees this is then obtained by modelling the repeated KPAs as i.i.d. Bernoulli trials, each leading to a successful choice of the true solution with a probability that can be estimated with $\mathcal{S}_{\text{Eve}}(n, L)^{-1}$ in case of Eve's KPA. With this $p_{\text{fail}} = \mathbb{P}[T \text{ KPA fail}] = (1 - \mathcal{S}_{\text{Eve}}(n, L)^{-1})^T$, so we may choose the key lifetime as $T \leq \lceil \log(1 - \mathcal{S}_{\text{Eve}}(n, L)^{-1}) \rceil^{-1} \log \zeta$ to ensure the security level set by ζ . Thus, we measure the key lifetime T in *attack opportunities* for Eve; however, since $\mathcal{S}_{\text{Eve}}(n, L)$ is typically huge, the resulting T is also very large. As an example, by plugging $n = 4096, L = 128$ in (5) and assuming $\zeta = 0.9999$, we obtain a key lifetime equivalent to at most $T = 1.25 \cdot 10^{1225}$ attack opportunities.

B. Class-Upgrade Known-Plaintext Attack

A known-plaintext attack may also be attempted by Steve, a second-class receiver aiming to improve its signal recovery performances with the intent of reaching the same quality of a first-class receiver. In this KPA, a partially correct encoding matrix $A^{(0)}$ that differs from $A^{(1)}$ in c entries is also known in addition to x and y . With this prior, Steve may compute $\varepsilon = y - A^{(0)}x = \Delta Ax$ where $\Delta A = A^{(1)} - A^{(0)}$ here is an unknown matrix with ternary entries in $\{-2, 0, 2\}$. Hence,

h	p_1^h	p_2^h	p_3^h	p_4^h	p_5^h	p_6^h	p_7^h	p_8^h	p_9^h	p_{10}^h	p_{11}^h	p_{12}^h	p_{13}^h	p_{14}^h
2	2													
3	-3	3												
4	$\frac{14}{3}$	-4	$\frac{16}{3}$											
5	$-\frac{15}{2}$	$\frac{65}{12}$	$-\frac{15}{2}$	$\frac{115}{12}$										
6	$\frac{62}{5}$	$-\frac{15}{2}$	11	$-\frac{27}{2}$	$\frac{88}{5}$									
7	-21	$\frac{959}{90}$	$-\frac{203}{12}$	$\frac{707}{36}$	$-\frac{301}{12}$	$\frac{5887}{180}$								
8	$\frac{254}{7}$	$-\frac{140}{9}$	$\frac{1226}{45}$	$-\frac{266}{9}$	$\frac{334}{9}$	$-\frac{422}{9}$	$\frac{19328}{315}$							
9	$-\frac{255}{4}$	$\frac{2613}{112}$	$-\frac{731}{16}$	$\frac{14701}{320}$	$-\frac{457}{8}$	$\frac{2233}{32}$	$-\frac{1415}{16}$	$\frac{259723}{2240}$						
10	$\frac{1022}{9}$	$\frac{2585}{72}$	$\frac{359105}{4536}$	$-\frac{7055}{96}$	$\frac{9869}{108}$	$-\frac{1725}{16}$	$\frac{28625}{216}$	$-\frac{48325}{288}$	$\frac{124952}{567}$					
11	$-\frac{1023}{5}$	$\frac{16973}{300}$	$-\frac{60775}{432}$	$\frac{5463953}{45360}$	$-\frac{435941}{2880}$	$\frac{7449761}{43200}$	$-\frac{19811}{96}$	$\frac{1091629}{4320}$	$-\frac{2764663}{8640}$	$\frac{381773117}{907200}$				
12	$\frac{4094}{11}$	$\frac{2277}{25}$	$\frac{687791}{2700}$	$-\frac{72523}{360}$	$\frac{3907067}{15120}$	$-\frac{341143}{1200}$	$\frac{599327}{1800}$	$-\frac{7909}{20}$	$\frac{1045349}{2160}$	$-\frac{2205833}{3600}$	$\frac{41931328}{51975}$			
13	$-\frac{1365}{2}$	$\frac{591721}{3960}$	$-\frac{2020421}{4320}$	$\frac{44385419}{129600}$	$-\frac{7815847}{17280}$	$\frac{116257063}{241920}$	$-\frac{3192163}{5760}$	$\frac{110721221}{172800}$	$-\frac{13148473}{17280}$	$\frac{19285357}{20736}$	$-\frac{20345507}{17280}$	$\frac{20646903199}{13305600}$		
14	$\frac{16382}{13}$	$-\frac{44863}{180}$	$\frac{34353347}{39600}$	$-\frac{38237381}{64800}$	$\frac{1292711}{1600}$	$-\frac{42972293}{51840}$	$\frac{122732801}{129600}$	$-\frac{92420419}{86400}$	$\frac{53508931}{43200}$	$-\frac{76095383}{51840}$	$\frac{77441609}{43200}$	$-\frac{588168119}{259200}$	$\frac{866732192}{289575}$	
15	$-\frac{16383}{7}$	$\frac{1074679}{2548}$	$-\frac{583763}{360}$	$\frac{113982839}{110880}$	$-\frac{12673507}{8640}$	$\frac{58584511}{40320}$	$-\frac{400088153}{241920}$	$\frac{1033251187}{564480}$	$-\frac{23927713}{11520}$	$\frac{193398181}{80640}$	$-\frac{98109773}{34560}$	$\frac{279340567}{80640}$	$-\frac{1060693411}{241920}$	$\frac{467168310097}{80720640}$

TABLE I
TABLE OF COEFFICIENTS OF THE POLYNOMIALS $P_h(L) = \sum_{j=1}^{h-1} p_j^h L^j$ IN (6) FOR $h = 2, \dots, 15$.

Steve performs a KPA by searching for a set of ternary symbols $\{\Delta A_{j,l}\}_{l=0}^{n-1}$ such that

$$\varepsilon_j = \sum_{l=0}^{n-1} \Delta A_{j,l} x_l \quad (7)$$

of which it is known that $\Delta A_{j,l} \neq 0$ only in c cases. Moreover, to ease the solution of this problem and make it row-wise separable, we assume that Steve has access to an even more accurate information, *i.e.*, the exact number c_j of non-zero entries for each row ΔA_j or equivalently the number of sign flips mapping $A_j^{(0)}$ into the corresponding $A_j^{(1)}$ (clearly, the total number of non-zero entries in ΔA is $c = \sum_{j=0}^{m-1} c_j$). By assuming this, we may prove the equivalence between Steve's KPA to each row of $A^{(1)}$ and a slightly adjusted SSP.

Definition 2 (γ -cardinality Subset-Sum Problem). Let $\{u_l\}_{l=0}^{n-1}, u_l \in \{1, \dots, Q\} \subset \mathbb{N}_+, \gamma \in \{1, \dots, n\} \subset \mathbb{N}_+$ and $v \in \mathbb{N}_+$. We define γ -cardinality subset-sum problem (γ -SSP) the problem of assigning n binary variables $b_l \in \{0, 1\}, l = 0, \dots, n-1$ such that

$$v = \sum_{l=0}^{n-1} b_l u_l \quad (8)$$

$$\gamma = \sum_{l=0}^{n-1} b_l \quad (9)$$

We define *solution* any $\{b_l\}_{l=0}^{n-1}$ verifying (8) and (9).

Proposition 2 (Steve's KPA). *The KPA to $A_j^{(1)}$ given $(x, y, A^{(0)}, c_j)$, is equivalent to a γ -SSP where $\gamma = c_j, Q = 2L, v = \frac{1}{2}\varepsilon_j + Lc_j, u_l = -A_{j,l}^{(0)}x_l + L$ and $b_l = \frac{1}{2} \left(1 - \frac{\hat{A}_{j,l}^{(1)}}{A_{j,l}^{(0)}}\right)$. This SSP has a true solution $\{\bar{b}_l\}_{l=0}^{n-1}$ that*

is mapped to the row $A_j^{(1)}$, and other candidate solutions that verify (8) and (9) but correspond to matrix rows $\hat{A}_j^{(1)} \neq A_j^{(1)}$.

The derivation of Proposition 2 is reported in Appendix C. We define $(x, y, A_j^{(0)}, A_j^{(1)})$ a *problem instance*. In the following, we will denote with $r = c_j/n$ the row-density of perturbations. Since in [18] the γ -cardinality SSP case is obtained as an extension of the results on the unconstrained SSP, we obtain the following Theorem.

Theorem 3 (Expected number of solutions for Steve's KPA). *For large n , the expected number of candidate solutions of the KPA in Proposition 2, in which (i) all the coefficients $\{u_l\}_{l=0}^{n-1}$ are i.i.d. uniformly drawn from $\{1, \dots, 2L\}$, and (ii) the true solution $\{b_l\}_{l=0}^{n-1}$ is drawn with equiprobable independent binary values, is*

$$\mathcal{S}_{\text{Steve}}(n, L, r) \stackrel{n \rightarrow \infty}{\simeq} \sqrt{\frac{3}{2}} \frac{r^{-1-nr} (1-r)^{-1-n(1-r)}}{2\pi nL} \quad (10)$$

The proof of Theorem 3 is reported in Appendix C. The number of candidate solutions found by Steve's KPA is by many orders of magnitude smaller than Eve's KPA, the reason being that Steve requires much less information to achieve complete knowledge of the true encoding $A^{(1)}$. In order to provide numerical evidence, we find all the solutions to Steve's KPA by means of the binary programming solver in CPLEX on a set of 50 randomly generated problem instances for $L = 5 \cdot 10^3$, a row-density of perturbations $r = 5/n, 10/n, 15/n$ and $n = 20, \dots, 32$ (except for $r = 5/n$, whose solution enumeration is still computationally feasible up to $n = 48$). The sample average of the number of solutions, $\hat{\mathcal{S}}_{\text{Steve}}(n, L, r)$, is reported in Fig. 4 and well predicted by the theoretical value in (10); note that this approximation is increasingly accurate for large n . Moreover, by resuming the previous example our $n = 64 \times$

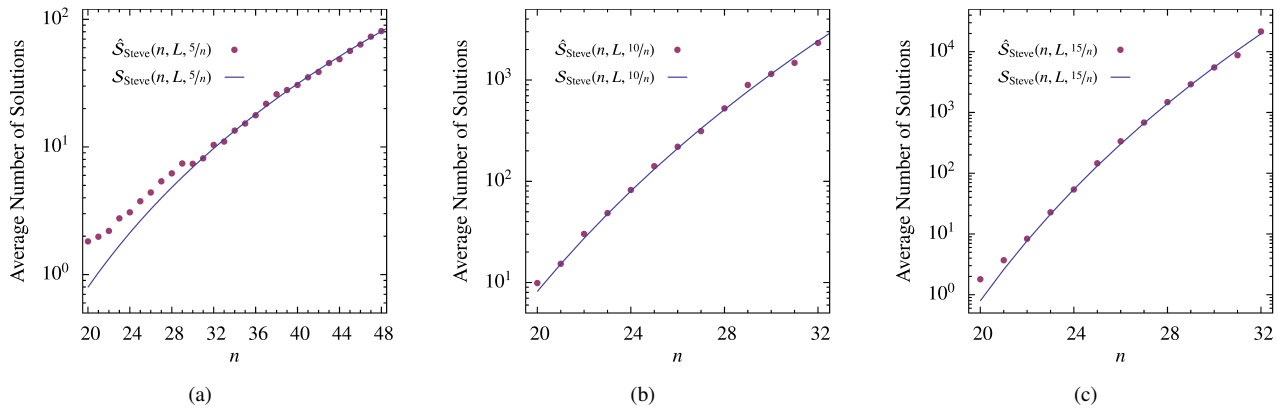


Fig. 4. Sample average of the number of solutions for Steve's KPA compared to the theoretical value of (10) for $L = 5 \cdot 10^3$ with row-density of perturbations $r = 5/n, 10/n, 15/n$.

64 pixel grayscale image quantised at $B_x = 8$ bit and encoded with two-class CS using ΔA with $r = 0.03$ will have on the average $6.25 \cdot 10^{234}$ candidate solutions of indistinguishable quality.

In terms of encryption key lifetime, leveraging the same considerations of Section III-A and simply replacing $\mathcal{S}_{\text{Eve}}(n, L)$ with $\mathcal{S}_{\text{Steve}}(n, L, r)$ yields the key lifetimes T with respect to class-upgrade attacks; as an example, plugging $n = 4096, L = 128, r = 0.03$ in (10) and assuming $\zeta = 0.9999$, yields at most $T = 1.25 \cdot 10^{231}$ attack opportunities for Steve.

The previous KPA analyses hinge on a counting argument in a general setting, without any other side information on the structure of $A^{(1)}$ or ΔA . As we will show in the experiments of Section IV, KPAs yield no advantage in terms of recovery performances to unintended receivers. Obviously, as further prior information becomes available (for example the knowledge that the unknown ΔA has additional structure, or that the original signal is distributed in a non-uniform fashion [36], [37]) revealing the hidden information may be easier. Yet, this is true for any encryption scheme in which either the encryption key or the plaintext have a non-uniform distribution and is out of the scope of this analysis.

C. Signal Recovery-Based Class-Upgrade Attacks

Class-upgrade attacks to two-class encryption schemes are closely related to a recovery problem setting that has attracted some attention in prior works, *i.e.*, *sparse signal recovery under matrix uncertainty*. To recast our problem in this setting, we may construct such a signal recovery-based attack by letting $A^{(1)} = A^{(0)} + \Delta A$ as the encoding matrix, where $A^{(0)}$ is known *a priori* and ΔA is an unknown random perturbation matrix. This information is paired with the knowledge of the ciphertext y and a prior on the unknown plaintext x , that is known to be sparse in a basis D . Thus, we attempt the *joint recovery* of x and ΔA , eventually just leading to a refinement of the estimated \hat{x} . Two main algorithms are capable of addressing specifically this problem setup for a generic ΔA , namely Generalised Approximate Message-Passing under Matrix Uncertainty (MU-GAMP [16]) and Sparsity-cognisant Total Least-Squares (S-TLS [15]).

Although appealing, this joint recovery approach can be anticipated to fail for multiple reasons. First, this attack is intrinsically harder than Steve's KPA in that the true plaintext x is here unknown. Whatever ΔA is a candidate solution to Steve's KPA given x , is also a possible solution of joint recovery with the same x as a further part of the solution. Since we know from Section III-B that Steve's KPA typically has a huge number of indistinguishable and equally-sparse candidate solutions, at least as many will verify the joint recovery problem when the plaintext is also unknown. Hence, this approach has negligible odds of yielding more information on ΔA than Steve's KPA.

Note that this relationship between the set of solutions to Steve's KPA and joint recovery-based attacks also prevents the latter from being of any use as a *refinement step* to improve ΔA after its guess by an initial KPA. In fact, recovering an estimate of x in this case would be to no avail, since the true x must be known *a priori* in the initial KPA.

Notwithstanding this, the above joint recovery approach estimates x along with a new ΔA ; thus, the best-case achievable signal recovery is the true x , for which the candidate solutions in ΔA are at best identical to those of the initial KPA, as by (7) they must verify $\varepsilon = \Delta A x$. No improvement is therefore obtained by applying joint recovery after Steve's KPA.

Furthermore, going back to simple joint-recovery, note that it amounts to solving $y = A^{(0)}x + \Delta A x$ with ΔA and x unknown, that is clearly a non-linear equality involving non-convex/non-concave operators. In general, this is a hard problem; both the aforementioned algorithms are indeed able to effectively compensate matrix uncertainties when ΔA depends on a low-dimensional, *deterministic* set of parameters. However, such a model does not apply to two-class encryption: even if ΔA is c -sparse, it has no deterministic structure – to make it so, one would need to know the exact set $C^{(0)}$ of c index pairs at which the sign flipping randomly occurred, which by itself entails a combinatorial search.

In fact, ΔA is *uniform* in the sense of [16] since it may be regarded as a realisation of a random matrix with i.i.d. zero-mean, bounded-variance entries (as also detailed in [7]). Hence, we expect the accuracy of the estimate \hat{x} with joint




Fig. 5. Average recovery signal-to-noise ratio performances of a class-upgrade attack using signal recovery under matrix uncertainty algorithms.

recovery (both using S-TLS and MU-GAMP) to agree with the uniform matrix uncertainty case of [16], where negligible improvement is shown with respect to the (non-joint) recovery algorithm GAMP [13]. The advocated reason is that the perturbation noise $\varepsilon = \Delta Ax$ is asymptotically Gaussian for a given x [16, Proposition 2.1].

We now provide some empirical evidence on the ineffectiveness of joint recovery as a class-upgrade attack for finite n , m and sparsity k . As an example, we let $n = 256$, $m = 128$, $k = 20$ and $\eta = \frac{c}{mn} \in [0.005, 0.1]$ and generate 100 random instances of $x = Ds$ with s which is k -sparse with respect to a randomly selected, known orthonormal basis D . For each η , we also generate 100 pairs of matrices $(A^{(0)}, A^{(1)})$ related as (1) and encode x by $y = A^{(1)}x$. Signal recovery is performed by MU-GAMP, S-TLS and GAMP. To maximise their performances, each of the algorithms is run with parameters provided by a “genie” revealing the exact value of the unknown features of x . In particular, MU-GAMP and GAMP are provided with an i.i.d. Bernoulli-Gaussian sparsity-enforcing signal model [13], [38] having the exact mean, variance and sparsity level of the instances s . As far as the perturbation ΔA is concerned, MU-GAMP is given the probability distribution of its i.i.d. entries. On the other hand, GAMP is initialised with the noise variance of $\varepsilon = \Delta Ax$, that is assumed Gaussian with i.i.d. entries. S-TLS is run in its locally-optimal, polynomial-time version [15, Section IV-B] and fine-tuned with respect to its regularisation parameter as η varies.

We here focus on measuring the Average⁴ Recovery Signal-to-Noise Ratio of the latter, ARSNR (dB) = $10 \log_{10} \hat{\mathbb{E}} \left(\frac{\|x\|_2^2}{\|x - \hat{x}\|_2^2} \right)$ reported in Fig. 5. The standard deviation from this average is less than 1.71 dB in all the reported curves. The maximum ARSNR performance gap between GAMP and MU-GAMP is 1.22 dB while S-TLS attains generally lower performances for high values of η . These observed performances confirm what is also found in [16], i.e., that GAMP, MU-GAMP and S-TLS substantially attain the same performances under uniform matrix uncertainty. As expected, class-upgrade attacks based on joint recovery are ineffective even for finite n and m , since GAMP under the same conditions is the reference case adopted in [7, Section IV] for the design of two-class encryption schemes.

⁴ $\hat{\mathbb{E}}(\cdot)$ denotes the sample average over a set of trials.

IV. NUMERICAL EXAMPLES

This Section aims at providing an intuitive appreciation of the poor quality obtained by signal recovery with KPA solutions. While the objective of KPAs is cryptanalysis the true encoding matrix to ultimately retrieve the encryption key, we here focus on the properties of KPA solutions as encoding matrix guesses that can, in the attackers’ belief, improve their signal recovery quality. Thus, we verify that this improvement does not occur by exemplifying practical cases of KPAs in a common framework, which follows this procedure:

- 1) *Attack*: an attacker performing a KPA gains access to a single plaintext-ciphertext pair (x', y') , and attacks the corresponding true encoding matrix $A^{(1)}$ row-by-row; we here infer each row $A_j^{(1)}$ by generating instances of an i.i.d. antipodal random vector until a large number of candidate solutions $\hat{A}_j^{(1)}$ that verify $y'_j = \hat{A}_j^{(1)}x'$ is found. Thus, the inferred $\hat{A}^{(1)}$ is composed by collecting the outputs of m Monte Carlo random searches for the corresponding matrix rows. This generation approach is preferable to solving each attacker’s KPA by means of CPLEX’s binary programming solver for two reasons. Firstly, it is known from Theorem 1 that the expected number of solutions is very large and thus the probability of finding one by random search is far from being negligible, while its computational cost is relatively low. Secondly, the theoretical conditions [24] that guarantee x' can be retrieved from y' despite the dimensionality reduction are applicable when $A^{(1)}$ is a typical realisation of an antipodal random matrix. On the contrary, integer programming solvers explore solutions in a systematic way, and tend to generate them in an ordered fashion. When only some of these solutions are considered (as obliged when n is large), this ordered approach yields non-typical sets of $\hat{A}_j^{(1)}$ that could be very distant from $A_j^{(1)}$;
- 2) *Signal Recovery*: to test its guess $\hat{A}^{(1)}$, the attacker may then pretend to ignore the *known* x' and recover an approximation \hat{x}' from $(y', \hat{A}^{(1)})$ by using a high-performance signal recovery algorithm such as GAMP [13], optimally tuned as in Section III-C. In this setting we measure its accuracy by the Recovery Signal-to-Noise Ratio, $\text{RSNR}' = 10 \log_{10} \frac{\|x'\|_2^2}{\|x' - \hat{x}'\|_2^2}$, which is the only quality indicator in the attacker’s perspective for $\hat{A}^{(1)}$. The RSNR' performances are here expected to match those of a (first-class) receiver fully informed on $A^{(1)}$, as the equality $y' = A^{(1)}x'$ is verified regardless of the exactness of $\hat{A}^{(1)}$;
- 3) *Verification*: as a further test of $\hat{A}^{(1)}$, the attacker attempts the recovery of a second, *unknown* plaintext x'' encoded as $y'' = A^{(1)}x''$, of which it is only *known* that it was obtained with the same encoding matrix as y' . The recovery \hat{x}'' is then obtained by means of GAMP, yielding a new $\text{RSNR}'' = 10 \log_{10} \frac{\|x''\|_2^2}{\|x'' - \hat{x}''\|_2^2}$ *unknown* to the attacker. If any point with high $\text{RSNR}'' \approx \text{RSNR}'$ is found, this will indicate the attacker’s success at guessing $\hat{A}^{(1)}$ close to the true $A^{(1)}$. We will show how this never occurs with a large number of candidate solutions,

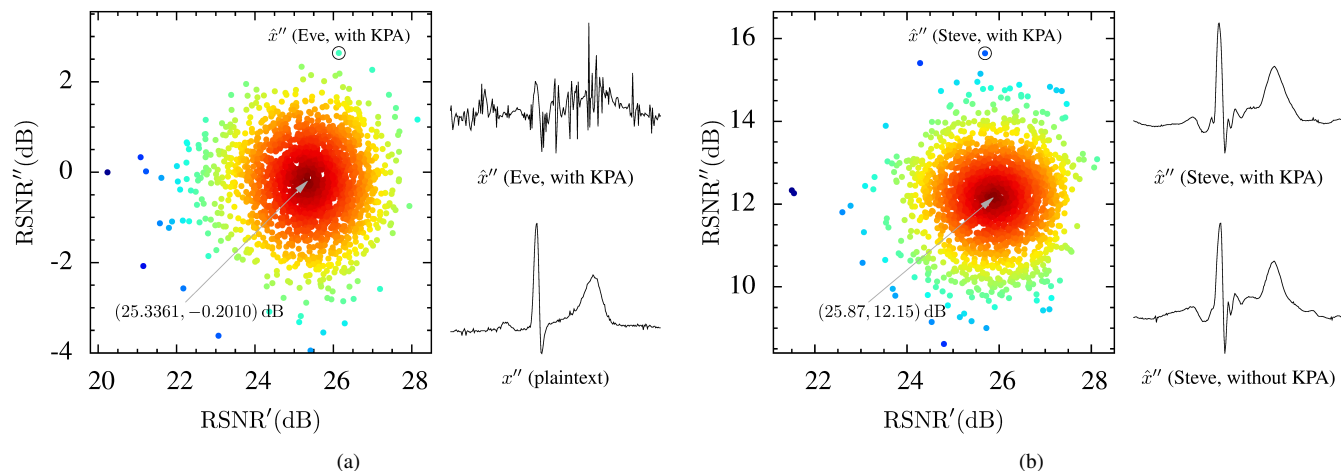


Fig. 6. Effectiveness of (a) Eve and (b) Steve's KPA in recovering a hidden ECG. Each point is a guess of the encoding matrix $A^{(1)}$ whose quality is assessed by decoding the ciphertext y' corresponding to the known plaintext x' (RSNR') and by decoding a new ciphertext y'' (RSNR''). The Euclidean distance from the average ($\text{RSNR}', \text{RSNR}''$) is highlighted by colour gradient.

and detail how the observed $(\text{RSNR}', \text{RSNR}'')$ pairs are distributed.

Both the practical examples of Eve and Steve's KPA follow the same procedure, with the exception that Eve directly generates $\hat{A}_j^{(1)}$, whereas Steve generates each row $\hat{A}_j^{(1)}$ by random search of the index set $C_j^{(0)}$ that maps the *known* $A_j^{(0)}$ to the guess $\hat{A}_j^{(1)}$ that verifies $y'_j = \hat{A}_j^{(1)} x'$. Coherently with the theoretical setting of Section III-B, we also assume that Steve knows that exactly c_j entries of $A^{(0)}$ have been flipped in each row of $A^{(1)}$. Repeating this search for m rows in both attacks provides Eve and Steve's candidate solutions $\hat{A}^{(1)}$, of which we will study how the corresponding $(\text{RSNR}', \text{RSNR}'')$ pairs are distributed as mentioned above.

A. Electrocardiographic Signals

We now consider ECG signals from the MIT PhysioNet database [39] sampled at $f_s = 256$ Hz and encoded as described, from two windows x', x'' of $n = 256$ samples (and quantised with $B_x = 12$ bit) into the measurement vectors y', y'' of dimensionality $m = 90$. Decoding is allowed by the sparsity level of the windowed signal when decomposed with D chosen as a Symmlet-6 orthonormal wavelet basis [40].

We generate 2000 candidate solutions for both Eve and Steve's KPA that correspond to the recovery performances reported in Fig. 6. While both malicious users are able to reconstruct the known plaintext x' with a relatively high average $\text{RSNR}' \approx 25$ dB (their KPAs indeed yield solutions to $y' = \hat{A}^{(1)} x'$), on the second window of samples x'' the eavesdropper achieves an average $\text{RSNR}'' \approx -0.20$ dB (Fig. 6a), whereas the second-class decoder achieves an average $\text{RSNR}'' \approx 12.15$ dB (Fig. 6b) when the two-class encryption scheme is set to a sign flipping density $\eta = c/mn = 0.03$ between $A^{(0)}$ and $A^{(1)}$. In this case, the nominal second-class $\text{RSNR} = 11.08$ dB when reconstructing x'' from y'' with $A^{(0)}$, while the correlation coefficient between RSNR' and RSNR'' is 0.0140; these figures clearly highlight the

ineffectiveness of KPAs at inferring $A^{(1)}$ in this case. This is also confirmed by the perceptual quality of \hat{x}'' corresponding to the maximum RSNR'' highlighted in Fig. 6.

B. Sensitive Text in Images

In this example we consider the same test images used in [7], *i.e.*, 640×512 pixel grayscale images of people holding a printed identification text concealed by means of two-class encryption. To reduce the computational burden of KPAs we assume a block size of 64×64 pixel, $B_x = 8$ bit per pixel, and encode the resulting $n = 4096$ pixels into $m = 2048$ measurements. Signal recovery is performed by assuming the blocks have a sparse representation on a 2D Daubechies-4 wavelet basis [40]. Two-class encryption is applied on the blocks containing printed text: we choose two adjacent blocks x', x'' containing some letters and encoded with the same $A^{(1)}$; in this case, the second-class decoder nominally achieves $\text{RSNR} = 12.57$ dB without attempting class-upgrade due to the flipping of $c = 251658$ entries (corresponding to a perturbation density $\eta = 0.03$) in the encoding matrix.

In order to test Eve and Steve's KPA we randomly generate 2000 solutions for the j -th row of the encoding given x', y' : it is worth noting that while in the previous case the signal dimensionality is sufficiently small to produce a solution set in less than two minutes, in this case generating 2000 different solutions for a single row may take up to several hours for some particularly hard instances.

By using these candidate solutions to find \hat{x}', \hat{x}'' we obtain the results of Figure 7: while both attackers attain an average $\text{RSNR}' \approx 33$ dB on x' , Eve is only capable of reconstructing x'' with an average $\text{RSNR}'' \approx 0.14$ dB where Steve reaches an average $\text{RSNR}'' \approx 12.80$ dB with $\eta = 0.03$. Note also that, although some lucky guesses exist with $\text{RSNR}'' > 12.57$ dB, it is impossible to identify them by looking at RSNR' since the correlation coefficient between RSNR' and RSNR'' is -0.0041 . Therefore, Steve cannot rely on observing the




Fig. 7. Effectiveness of (a) Eve and (b) Steve's KPA in recovering hidden image blocks. Each point is a guess of the encoding matrix $A^{(1)}$ whose quality is assessed by decoding the ciphertext y' corresponding to the known plaintext x' (RSNR') and by decoding a new ciphertext y'' (RSNR''). The Euclidean distance from the average (RSNR', RSNR'') is highlighted by colour gradient.

RSNR' to choose the best performing solution $\hat{A}^{(1)}$, so both Eve and Steve's KPAs are inconclusive. As a further perceptual evidence of this, the best recoveries according to the RSNR'' are reported in Fig. 7.

V. CONCLUSION

In this paper we have analysed known-plaintext attacks as they may be carried out on standard CS schemes with antipodal random encoding matrices as well as on the particular multiclass protocol developed in [7]. In particular, the analysis was carried out from the two perspectives of an eavesdropper and a second-class user trying to guess the true encoding matrix. In both cases we have mapped multiclass CS into a collection of subset-sum problems with the aim of counting the candidate encoding matrices that match a given plaintext-ciphertext pair. In the eavesdropper case we have found that for each row the expected number grows as $O(n^{-\frac{1}{2}} \cdot 2^n)$ – finding the true solution among such huge sets is infeasible. A further study of the candidate solutions' Hamming distance from the true one showed that, as the dimensionality n increases, the expected number of solutions close to the true one is only a small fraction of the solution set. As for the second-class user we have shown that depending on the available information on the true encoding matrix, the expected number of solutions is significantly smaller, yet sufficiently high for large n to reassure that a second-class user will not be able to perform class-upgrade. Moreover, other class-upgrade attacks based on signal recovery under matrix uncertainty were shown to yield almost identical performances to those of a standard decoding algorithm.

Finally, we showed some simulated cases of KPAs on real-world signals such as ECG traces and images by running a random search for a solution set corresponding to realistic plaintext-ciphertext pairs, and afterwards tested whether any of the returned candidate solutions could lead to finding the true encoding matrix by testing them on a successive ciphertext.

In all the observed cases, we have found that the decoding performances match the average RSNR level prescribed by the multiclass encryption protocol, *i.e.*, both malicious users are unable to successfully decode other plaintexts with significant and stable quality improvements with respect to their available prior information.

APPENDIX A

PROOFS ON EAVESDROPPER'S KPA

The following definition is used in Appendices A and C.

Definition 3. We define the functions

$$F_p(a, b) = \int_0^1 \frac{\xi^p}{1 + e^{a\xi - b}} d\xi \quad (11)$$

$$G_p(a, b) = \int_0^1 \frac{\xi^p}{(1 + e^{a\xi - b})(1 + e^{b - a\xi})} d\xi \quad (12)$$

Proof of Proposition 1. Define the binary variables $b_l \in \{0, 1\}$ so that $\text{sign}(x_l) \hat{A}_{j,l}^{(1)} = 2b_l - 1$ and the positive coefficients $u_l = |x_l|$. With this choice (2) is equivalent to $y_j = \sum_{l=0}^{n-1} (2b_l - 1)u_l$ which leads to a SSP with $v = \frac{1}{2} \left(y_j + \sum_{l=0}^{n-1} |x_l| \right)$. Since we know that each measurement y_j must correspond to the inner product between x and the row $A_j^{(1)}$, the latter's entries are straightforwardly mapped to the *true* solution of this SSP, $\{\bar{b}_l\}_{l=0}^{n-1}$. \square

Proof of Theorem 1. Let us first note that, for large n , v in Proposition 1 is an integer in the range $\{0, \dots, nL/2\}$, with the values outside this interval being asymptotically unachievable as $n \rightarrow \infty$ (see [18, Section 4]). We let $\tau = v/nL$, $\tau \in [0, 1/2]$, and $a(\tau)$ be the solution in a of the equation $\tau = F_1(a, 0)$ (*i.e.* [18, (4.2)]) that is unique since $F_p(a, 0)$ in (11) is monotonically decreasing in a .

From [18, (4.1)] the number of solutions of a SSP with integer coefficients $\{u_l\}_{l=0}^{n-1}$ uniformly distributed in $\{1, \dots, L\}$ is

$$\mathcal{S}_{\text{Eve}}(\tau, n, L) \stackrel{n \rightarrow \infty}{\simeq} \frac{e^{n[a(\tau)\tau + \int_0^1 \log(1 + e^{-a(\tau)\xi}) d\xi]} \sqrt{2\pi n L^2 G_2(a(\tau), 0)}}{}$$




Fig. 8. Gaussian approximation of $\mathcal{S}_{\text{Eve}}(\tau, n, L)$ for $n = 64, L = 10^4$ by letting $\sigma^2 \approx 1/12n$.

that we anticipate to have an approximately Gaussian profile (see Fig. 8). We now compute the average of $\mathcal{S}_{\text{Eve}}(\tau, n, L)$ in τ , that clearly depends on the probability of selecting any value of $v \in \{0, \dots, \frac{nL}{2}\}$, *i.e.*, of $\tau \in [0, \frac{1}{2}]$. Since it is the result of a linear combination, the probability that a specific value of v appears in a random instance of the SSP is proportional to the number of solutions associated to it. In normalised terms, the PDF of τ must be proportional to $\mathcal{S}_{\text{Eve}}(\tau, n, L)$, *i.e.*, τ is distributed as

$$f_\tau(t) = \frac{1}{\int_0^{\frac{1}{2}} \mathcal{S}_{\text{Eve}}(\xi, n, L) d\xi} \begin{cases} \mathcal{S}_{\text{Eve}}(t, n, L), & 0 \leq t \leq \frac{1}{2} \\ 0, & \text{otherwise} \end{cases}$$

With $f_\tau(t)$ we can compute the expected number of solutions:

$$\mathbb{E}_\tau[\mathcal{S}_{\text{Eve}}(\tau, n, L)] = \frac{\int_0^{\frac{1}{2}} \mathcal{S}_{\text{Eve}}^2(\xi, n, L) d\xi}{\int_0^{\frac{1}{2}} \mathcal{S}_{\text{Eve}}(\xi, n, L) d\xi} \quad (13)$$

Although we could resort to numerical integration, (13) can be simplified by exploiting what noted above, *i.e.*, that $\mathcal{S}_{\text{Eve}}(\tau, n, L)$ has an approximately Gaussian profile in τ (Fig. 8) with a maximum in $\tau = 1/4$. Hence, the expectation in τ becomes

$$\begin{aligned} \mathbb{E}_\tau[\mathcal{S}_{\text{Eve}}(\tau, n, L)] &\stackrel{n \rightarrow \infty}{\approx} \mathcal{S}_{\text{Eve}}\left(\frac{1}{4}, n, L\right) \frac{\int_{-\infty}^{\infty} \left(e^{-\frac{(\xi-\frac{1}{4})^2}{2\sigma^2}}\right)^2 d\xi}{\int_{-\infty}^{\infty} e^{-\frac{(\xi-\frac{1}{4})^2}{2\sigma^2}} d\xi} \\ &= \mathcal{S}_{\text{Eve}}\left(\frac{1}{4}, n, L\right) \frac{1}{\sqrt{2}} = \frac{2^n}{L} \sqrt{\frac{3}{\pi n}} \end{aligned} \quad (14)$$

that is actually independent of the σ^2 used in the Gaussian approximation, and in which we have exploited $a(1/4) = 0$ to obtain the statement of the theorem. \square

APPENDIX B

HAMMING DISTANCE OF KPA SOLUTIONS

Proof of Theorem 2. We here concentrate on counting the number of candidate solutions $\{b_l\}_{l=0}^{n-1}$ to Eve's KPA that

differ from the true one, $\{\bar{b}_l\}_{l=0}^{n-1}$, by exactly h components (at Hamming distance h). We assume that $K \subseteq \{0, \dots, n-1\}$ is the set of indexes for which there is a disagreement, *i.e.*, for all $l \in K$ we have $b_l = 1 - \bar{b}_l$; this set has cardinality h , and is one among $\binom{n}{h}$ possible sets. Since both $\{b_l\}_{l=0}^{n-1}$ and $\{\bar{b}_l\}_{l=0}^{n-1}$ are solutions to the same SSP, and that $b_l = \bar{b}_l$ are identical for $l \notin K$, $\sum_{l \in K} (1 - \bar{b}_l) u_l = \sum_{l \in K} \bar{b}_l u_l$ must hold, implying the equality

$$\sum_{\substack{l \in K \\ \bar{b}_l = 0}} u_l - \sum_{\substack{l \in K \\ \bar{b}_l = 1}} u_l = 0 \quad (15)$$

Although (15) recalls the well-known partition problem, in our case K is chosen by each problem instance that sets all u_l and \bar{b}_l . Thus, (15) holds in a number of cases that depends on how many of the $2^h L^h$ possible assignments of all u_l and \bar{b}_l satisfy it. The only feasible cases are for $h > 1$, and to analyse them we assume $K = \{0, \dots, h-1\}$ (the disagreements occur in the first h ordered indexes) without loss of generality.

Moreover, when (15) holds for some $\{\bar{b}_l\}_{l=0}^{n-1}$ it also holds for $\{1 - \bar{b}_l\}_{l=0}^{n-1}$. Hence, we may count the configurations that verify (15) with $\bar{b}_0 = 0$, knowing that their number will be only *half* of the total. With this, the configurations with $\bar{b}_0 = 0$ must have $\bar{b}_l = 1$ for *at least* one $l > 0$ in order to satisfy (15), giving $2^{h-1} - 1$ total cases to check.

The following paragraphs illustrate that, for $h < L$, the number of configurations that verify (15) can be written as a polynomial of order $h-1$. With this in mind we can start with the explicit computation for $h = \{2, 3\}$. For $h = 2$, there is only one feasible assignment for the $\{\bar{b}_l\}_{l=0}^{n-1}$, so $u_0 = u_1$ in (15), which makes $2L$ cases out of $2^2 L^2$. For $h = 3$, one has 3 feasible assignments for the $\{\bar{b}_l\}_{l=0}^{n-1}$. Due to the symmetry of (15) all the configurations have the same behaviour and we may focus on, *e.g.*, $\bar{b}_0 = \bar{b}_1 = 0$ and $\bar{b}_2 = 1 \Rightarrow u_0 + u_1 = u_2$; this can be satisfied only when $u_0 + u_1 \leq L$, *i.e.*, for $\frac{L(L-1)}{2}$ configurations. This makes a total of $2 \cdot 3 \cdot \frac{L(L-1)}{2} = 3L(L-1)$ over the $2^3 L^3$ possible configurations.

For $h > 3$, this procedure is much less intuitive; nevertheless, we can at least prove that the function $P_h(L)$ counting the configurations for which (15) holds is a polynomial in L of degree $h-1$. To show this, let us proceed in three steps.

- 1) Indicate with $\pi_{\bar{b}}$ the $(h-1)$ -dimensional subspace of \mathbb{R}^h defined by $\sum_{\substack{l \in K \\ \bar{b}_l = 0}} \xi_l - \sum_{\substack{l \in K \\ \bar{b}_l = 1}} \xi_l = 0, \xi \in \mathbb{R}^h$. The intersection $\alpha_{\bar{b}}(L) = \{1, \dots, L\}^h \cap \pi_{\bar{b}}$ is such that each assignment of $\{u_l\}_{l=0}^{h-1} \in \{1, \dots, L\}^h$ satisfying (15) is an integer point in $\alpha_{\bar{b}}$. To count those points define $\beta_{\bar{b}}(L) = \{0, \dots, L+1\} \cap \pi_{\bar{b}}$ and note that the number of integer points in $\alpha_{\bar{b}}$ is equal to the number of integer points in the interior of $\beta_{\bar{b}}$ (the points on the frontier of $\beta_{\bar{b}}$ have at least one coordinate that is either 0 or $L+1$). Note how $\{0, \dots, L+1\}^h$ scales linearly with $L+1$ while $\pi_{\bar{b}}$ is a subspace and therefore scale-invariant. Hence, their intersection $\beta_{\bar{b}}(L)$ is an $h-1$ -dimensional polytope that scales proportionally to the integer $L+1$, as required by Ehrhart's theorem [41]. The number $E_{\bar{b}}(L)$ of integer points in $\beta_{\bar{b}}(L)$ is then a polynomial in $L+1$ (and so L) of degree equal to the dimensionality of $\beta_{\bar{b}}(L)$, *i.e.*, $h-1$.

From Ehrhart-Macdonald's reciprocity theorem [42] we know that the number of integer points in the interior of $\beta_{\bar{b}}$ and thus in $\alpha_{\bar{b}}$ is $(-1)^{h-1}E_{\bar{b}}(-L)$, that is also a polynomial in L of degree $h-1$.

- 2) If two different assignments $\{\bar{b}_l\}_{l=0}^{h-1}$ and $\{\bar{b}'_l\}_{l=0}^{h-1}$ are considered, then $\alpha_{\bar{b}}(L) \cap \alpha_{\bar{b}'}(L) = \{1, \dots, L\}^h \cap \pi_{\bar{b}} \cap \pi_{\bar{b}'}$. The same argument we used above tells us that the number of integer points in such an intersection is a polynomial in L of degree $h-2$ and, in general that the number of integer points in the intersection of any number of polytopes $\alpha_{\bar{b}}(L)$ is a polynomial of degree not larger than $h-1$.
- 3) The number of configurations of $\{u_l\}_{l=0}^{h-1}$ and $\{\bar{b}_l\}_{l=0}^{h-1}$ that satisfy (15) with respect to the above K is the number of integer points in the union of all possible polytopes $\alpha_{\bar{b}}$, i.e., $\bigcup_{\{\bar{b}_l\}_{l=0}^{h-1}} \alpha_{\bar{b}}(L)$. Such a number can be computed by the inclusion-exclusion principle that amounts to properly summing and subtracting the number of integer points in those polytopes and their various intersections. Since sum and subtraction of polynomials yield polynomials of non-increasing degree, we know that number is the evaluation of a polynomial $P_h(L)$ with degree not greater than $h-1$.

Let us then write $P_h(L) = \sum_{j=0}^{h-1} p_j^h L^j$. In order to compute its coefficients p_j^h we may fix a binary configuration $\{\bar{b}_l\}_{l=0}^{h-1}$, count the points $\{u_l\}_{l=0}^{h-1} \in \mathbb{N}_+^h$ for which (15) is verified by means of integer partition functions (that also have a polynomial expansion), and subtract the points in which $\{u_l\}_{l=0}^{h-1} \notin \{1, \dots, L\}^h$. By summation over all binary configurations, one can extract the coefficients associated with L^j for each h . Table I reports the result of this procedure as carried out by symbolic computation for $h \leq 15$. \square

APPENDIX C

PROOFS ON THE CLASS-UPGRADE KPA

Proof of Proposition 2. In this case the attacker knows $(A^{(0)}, x, y)$, and is able to calculate $\varepsilon_j = y_j - \sum_{l=0}^{n-1} A_{j,l}^{(0)} x_l = \sum_{l=0}^{n-1} \Delta A_{j,l} x_l$ where the $\Delta A_{j,l}$ are unknown. For the j -th row, the attacker also knows there are c_j non-zero elements in $\Delta A_{j,l} = -2A_{j,l}^{(0)} b_l$ with $b_l \in \{0, 1\}$ binary variables that are 1 if the flipping occurred and 0 otherwise. Note that from the above information $c_j = \sum_{l=0}^{n-1} b_l$. With this we define a set of even weights $D_l = -2A_{j,l}^{(0)} x_l \in \{-2L, \dots, -2, 0, 2, \dots, 2L\}$ so the KPA is defined by satisfying the equalities

$$\varepsilon_j = \sum_{l=0}^{n-1} D_l b_l \quad (16)$$

$$c_j = \sum_{l=0}^{n-1} b_l \quad (17)$$

To obtain a standard γ -SSP with positive weights and $\gamma = c_j$ we sum $2L$ to all D_l so (16) becomes $\varepsilon_j + 2L \sum_{l=0}^{n-1} b_l = \sum_{l=0}^{n-1} (D_l + 2L) b_l$. Multiplying both sides by $1/2$ and using (17) yields $v = \frac{1}{2} \varepsilon_j + L c_j = \sum_{l=0}^{n-1} u_l b_l$ where $u_l = -A_{j,l}^{(0)} x_l + L \in \{0, \dots, Q\}$. $Q = 2L$. Finally, we exclude $u_l = 0$ to facilitate the attack. \square

Proof of Theorem 3. Assume $F_p(a, b)$ and $G_p(a, b)$ as in (11),(12). Define the normalised constraint $r = \frac{c_j}{n}$ and two quantities $a(\tau, r)$ and $b(\tau, r)$ that are the solutions of the following system of equalities

$$\begin{aligned} r &= F_0(a, b) \\ \tau &= F_1(a, b) \end{aligned}$$

that are respectively equivalent to [18, (5.3-4)]. We also define

$$\mathcal{G}(\tau, r) = \begin{pmatrix} G_0(a(\tau, r), b(\tau, r)) & G_1(a(\tau, r), b(\tau, r)) \\ G_1(a(\tau, r), b(\tau, r)) & G_2(a(\tau, r), b(\tau, r)) \end{pmatrix}$$

With this, [18, (5.8-9)] prove that the number of solutions of a γ -SSP with integer coefficients $\{u_l\}_{l=0}^{n-1}$ uniformly distributed in $\{1, \dots, Q\}$, $Q = 2L$, $\gamma = c_j$ is

$$\mathcal{S}_{\text{Steve}}(\tau, n, L, r) = \frac{e^{n(a(\tau, r)\tau - b(\tau, r)r)}}{4\pi n L \sqrt{\det(\mathcal{G}(\tau, r))}} \cdot e^{n \int_0^1 \log[1 + e^{b(\tau, r) - a(\tau, r)\xi}] d\xi} \quad (18)$$

Using the same arguments as in the proof of Theorem 1, we average on τ and obtain an expression identical to (13) for the computation of $\mathbb{E}_\tau[\mathcal{S}_{\text{Steve}}(\tau, n, L, r)]$. Since $\mathcal{S}_{\text{Steve}}(\tau, n, L, r)$ has once again an approximately Gaussian profile in τ with a maximum in $\tau = \frac{r}{2}$ we approximate the expectation in τ ,

$$\begin{aligned} \mathbb{E}_\tau[\mathcal{S}_{\text{Steve}}(\tau, n, L, r)] &\stackrel{n \rightarrow \infty}{\simeq} \mathcal{S}_{\text{Steve}}\left(\frac{r}{2}, n, L, r\right) \frac{1}{\sqrt{2}} \\ &= \sqrt{\frac{3}{2}} \frac{r^{-1-n\rho} (1-r)^{-1-n(1-r)}}{2\pi n L} \quad (19) \end{aligned}$$

by using the fact that $a\left(\frac{r}{2}, r\right) = 0$ and $b\left(\frac{r}{2}, r\right) = \log\left(\frac{r}{1-r}\right)$. \square

REFERENCES

- [1] D. L. Donoho, "Compressed Sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [2] E. J. Candes and M. B. Wakin, "An Introduction to Compressive Sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, Mar. 2008.
- [3] Y. Rachlin and D. Baron, "The secrecy of compressed sensing measurements," in *2008 Forty Sixth Annual Allerton Conference on Communication, Control, and Computing*. IEEE, 2008, pp. 813–817.
- [4] A. Orsdemir, H. O. Altun, G. Sharma, and M. F. Bocko, "On the security and robustness of encryption via compressed sensing," in *2008 IEEE Military Communications Conference (MILCOM 2008)*. IEEE, 2008, pp. 1–7.
- [5] V. Cambareri, J. Haboba, F. Pareschi, 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 (ISCAS)*. IEEE, 2013, pp. 1356–1359.
- [6] J. Barcelo-Llado, A. Morell, and G. Seco-Granados, "Amplify-and-Forward Compressed Sensing as a Physical-Layer Secrecy Solution in Wireless Sensor Networks," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 5, pp. 839–850, May 2014.
- [7] 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.
- [8] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "A survey on sensor networks," *IEEE Communications Magazine*, vol. 40, no. 8, pp. 102–114, 2002.
- [9] E. J. Candes and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse problems*, vol. 23, no. 3, p. 969, 2007.

- [10] E. J. Candes and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [11] J. Haboba, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "A pragmatic look at some compressive sensing architectures with saturation and quantization," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 2, no. 3, pp. 443–459, Sept 2012.
- [12] E. J. Candes and T. Tao, "Decoding by linear programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.
- [13] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *2011 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2011, pp. 2168–2172.
- [14] E. van den Berg and M. P. Friedlander, "Sparse optimization with least-squares constraints," *SIAM Journal on Optimization*, vol. 21, no. 4, pp. 1201–1229, 2011.
- [15] H. Zhu, G. Leus, and G. B. Giannakis, "Sparsity-cognizant total least-squares for perturbed compressive sampling," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2002–2016, 2011.
- [16] J. T. Parker, V. Cevher, and P. Schniter, "Compressive sensing under matrix uncertainties: An approximate message passing approach," in *2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*. IEEE, 2011, pp. 804–808.
- [17] S. Martello and P. Toth, *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc., 1990.
- [18] T. Sasamoto, T. Toyozumi, and H. Nishimori, "Statistical mechanics of an np-complete problem: subset sum," *Journal of Physics A: Mathematical and General*, vol. 34, no. 44, p. 9555, 2001.
- [19] C. E. Shannon, "Communication Theory of Secrecy Systems," *Bell System Technical Journal*, vol. 28, no. 4, pp. 656–715, Oct. 1949.
- [20] P.-L. Loh and M. J. Wainwright, "High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity," *Annals of Statistics*, vol. 40, no. 3, p. 1637, 2012.
- [21] M. Herman and T. Strohmer, "General deviants: An analysis of perturbations in compressed sensing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 342–349, 2010.
- [22] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "Average recovery performances of non-perfectly informed compressed sensing: with applications to multiclass encryption," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, April 2015, pp. 3651–3655.
- [23] Y. Chi, L. L. Scharf, A. Pezeshki, and A. R. Calderbank, "Sensitivity to basis mismatch in compressed sensing," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2182–2195, 2011.
- [24] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Comptes Rendus Mathématique*, vol. 346, no. 9, pp. 589–592, 2008.
- [25] D. Donoho and J. Tanner, "Precise undersampling theorems," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 913–924, June 2010.
- [26] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, 2008.
- [27] I. Drori, "Compressed video sensing," in *BMVA Symposium on 3D Video-Analysis, Display, and Applications*, 2008.
- [28] S. W. Golomb and G. Gong, *Signal Design for Good Correlation: For Wireless Communication, Cryptography, and Radar*. Cambridge University Press, Jul. 2005.
- [29] J. Massey, "Shift-register synthesis and BCH decoding," *IEEE Transactions on Information Theory*, vol. 15, no. 1, pp. 122–127, Jan. 1969.
- [30] A. Kerckhoffs, "La cryptographie militaire," *Journal des sciences militaires*, vol. IX, pp. 5–38, Jan. 1883.
- [31] J. C. Lagarias and A. M. Odlyzko, "Solving low-density subset sum problems," *Journal of the ACM (JACM)*, vol. 32, no. 1, pp. 229–246, 1985.
- [32] R. Merkle and M. Hellman, "Hiding information and signatures in trapdoor knapsacks," *IEEE Transactions on Information Theory*, vol. 24, no. 5, pp. 525–530, 1978.
- [33] B. Chor and R. L. Rivest, "A knapsack-type public key cryptosystem based on arithmetic in finite fields," *IEEE Transactions on Information Theory*, vol. 34, no. 5, pp. 901–909, 1988.
- [34] A. M. Odlyzko, "The rise and fall of knapsack cryptosystems," *Cryptology and computational number theory*, vol. 42, pp. 75–88, 1990.
- [35] ILOG, Inc., "ILOG CPLEX: High-performance software for mathematical programming and optimization," 2006.
- [36] M. Mangia, R. Rovatti, and G. Setti, "Rakeness in the design of analog-to-information conversion of sparse and localized signals," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 59, no. 5, pp. 1001–1014, May 2012.
- [37] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "A rakeness-based design flow for analog-to-information conversion by compressive sensing," in *2013 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2013, pp. 1360–1363.
- [38] J. Vila and P. Schniter, "Expectation-maximization bernoulli-gaussian approximate message passing," in *2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*. IEEE, 2011, pp. 799–803.
- [39] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000 (June 13).
- [40] S. Mallat, *A wavelet tour of signal processing*. Access Online via Elsevier, 1999.
- [41] E. Ehrhart, "Sur un probleme de géométrie diophantienne linéaire. ii. systemes diophantiens linéaires. (french)," *J. Reine Angew. Math.*, vol. 227, pp. 25–49, 1967.
- [42] I. G. Macdonald, "Polynomials Associated with Finite Cell-Complexes," *Journal of the London Mathematical Society*, vol. 2, no. 1, pp. 181–192, 1971.



Valerio Cambareri (S'13) received the B.S., M.S. (*summa cum laude*), and Ph.D. degrees in Electronic Engineering from the University of Bologna, Italy, in 2008, 2011 and 2015 respectively. From 2012 to 2015, he was a Ph.D. student in Electronics, Telecommunications and Information Technologies at DEI – University of Bologna, Italy. In 2014 he was a visiting Ph.D. student in the Integrated Imagers team at IMEC, Belgium. His current research activity focuses on statistical and digital signal processing, compressed sensing and computational imaging.



Mauro Mangia (S'09-M'13) received the B.S. and M.S. degree in Electronic Engineering from the University of Bologna, Italy, in 2004 and 2009 respectively; he received the Ph.D. degree in Information Technology from the University of Bologna in 2013. He is currently a post-doc researcher in the statistical signal processing group of ARCEN – University of Bologna, Italy. In 2009 and 2012 he was a visiting Ph.D. student at the École Polytechnique Fédérale de Lausanne (EPFL). His research interests are in nonlinear systems, compressed sensing, ultrawideband systems and system biology. He was recipient of the 2013 IEEE CAS Society Guillemin-Cauer Award and the best student paper award at IEEE ISCAS2011.



Fabio Pareschi (S'05-M'08) received the Dr. Eng. degree (with honours) in Electronic Engineering from University of Ferrara, Italy, in 2001, and the Ph.D. in Information Technology under the European Doctorate Project (EDITH) from University of Bologna, Italy, in 2007. He is currently an Assistant Professor in the Department of Engineering (EN-DIF), University of Ferrara. He is also a faculty member of ARCES – University of Bologna, Italy. He served as Associate Editor for the IEEE Transactions on Circuits and Systems – Part II (2010-2013).

His research activity focuses on analog and mixed-mode electronic circuit design, statistical signal processing, random number generation and testing, and electromagnetic compatibility. He was recipient of the best paper award at ECCTD2005 and the best student paper award at EMCZurich2005.



Riccardo Rovatti (M'99-SM'02-F'12) received the M.S. degree in Electronic Engineering and the Ph.D. degree in Electronics, Computer Science, and Telecommunications both from the University of Bologna, Italy in 1992 and 1996, respectively. He is now a Full Professor of Electronics at the University of Bologna. He is the author of approximately 300 technical contributions to international conferences and journals, and of two volumes. His research focuses on mathematical and applicative aspects of statistical signal processing and on the application

of statistics to nonlinear dynamical systems. He received the 2004 IEEE CAS Society Darlington Award, the 2013 IEEE CAS Society Guillemin-Cauer Award, as well as the best paper award at ECCTD 2005, and the best student paper award at EMC Zurich 2005 and ISCAS 2011. He was elected IEEE Fellow in 2012 for contributions to nonlinear and statistical signal processing applied to electronic systems.



Gianluca Setti (S'89-M'91-SM'02-F'06) received the Ph.D. degree in Electronic Engineering and Computer Science from the University of Bologna in 1997. Since 1997 he has been with the School of Engineering at the University of Ferrara, Italy, where he is currently a Professor of Circuit Theory and Analog Electronics and is also a permanent faculty member of ARCES – University of Bologna, Italy. His research interests include nonlinear circuits, implementation and application of chaotic circuits and systems, electromagnetic compatibility, statistical

signal processing and biomedical circuits and systems. Dr. Setti received the 2013 IEEE CAS Society Meritorious Service Award and co-recipient of the 2004 IEEE CAS Society Darlington Award, of the 2013 IEEE CAS Society Guillemin-Cauer Award, as well as of the best paper award at ECCTD 2005, and the best student paper award at EMC Zurich 2005 and at ISCAS 2011. He held several editorial positions and served, in particular, as the Editor-in-Chief for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS – PART II (2006-2007) and of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS – PART I (2008-2009). Dr. Setti was the Technical Program Co-Chair at ISCAS 2007, ISCAS 2008, ICECS 2012, BioCAS 2013 as well as the General Co-Chair of NOLTA 2006. He was a member of the Board of Governors of the IEEE CAS Society (2005-2008), served as its 2010 President, and he is a Distinguished Lecturer of CASS (2015-2016). He held several other volunteer positions for the IEEE and in 2013-2014 he was the first non North-American Vice President of the IEEE for Publication Services and Products.