Physical-layer Intrusion Detection for Wireless Networks using Compressed Sensing

Alexandros Fragkiadakis*, Sofia Nikitaki*† and Panagiotis Tsakalides*†
*Institute of Computer Science, Foundation for Research and Technology-Hellas (FORTH-ICS)
† Department of Computer Science, University of Crete
Heraklion, Crete, Greece
email: {alfrag, nikitaki, tsakalid}@ics.forth.gr

***Abstract*—The broadcast nature of wireless networks has been widely exploited by adversaries in order to cause severe denial-of-service attacks. Several algorithms are proposed in the literature for the detection and mitigation of such attacks at the physical and medium access layers. In this work, we combine recent advances in compressed sensing theory, along with a cumulative-sum anomaly-based algorithm, for the detection of physical-layer attacks. The algorithm considers a metric based on the Signal-to-Interference-plus-Noise-Ratio (SINR). Compressed sensing makes feasible the use of far fewer SINR measurements for effective intrusion detection. The performance evaluation based on real experimental data shows that attacks are detected with high accuracy using a small number of measurements.**

***Index Terms*—jamming attacks, intrusion detection, signal-to-interference-plus-noise-ratio, compressed sensing, performance evaluation.**

I. Introduction

Wireless technology offers inexpensive ubiquitous broadband connectivity. The license-free nature of the industrial, scientific and medical (ISM) radio band along with the rapid proliferation of smart phones, has enabled ubiquitous broadband wireless internet access to millions of users worldwide. However, the broadband nature of this technology makes wireless networks (WNs) susceptible to a number of attacks. Adversaries can easily, using off-the-shelf devices, launch attacks and significantly impact the performance of a WN [1]. In the literature, adversaries are usually referred as *jammers*, and their corresponding attacks as *jamming attacks*. Following the definition from [2], a jammer is considered to be "*an entity who is purposefully trying to interfere with the physical transmission and reception of wireless communications*." A jammer can use different strategies such as (i) energy emission on the neighboring channels that legitimate nodes use for communication, causing the WN to suffer from external interference, or (ii) energy emission on the same channel, in which case WN suffers from internal interference. There are more intelligent attacks including *CTS corruption jamming*, *ACK corruption jamming*, *DATA corruption jamming*, *narrowband jamming*, and *DIFS waiting jamming* [3].

This work was funded by the CS-ORION (PIAP-GA-2009-251605) and HYDROBIONETS (ICT-GA-2011-287613) grants within the 7th Framework Program of the European Community.

In our previous studies, we proposed and evaluated several metrics based on the SINR for the detection of physical-layer jamming attacks. Using SINR and a cumulative-sum (Cusum) algorithm, which has the ability to detect abrupt changes, we can successfully detect the jamming attacks launched at the physical layer of a WN. This algorithm is part of the intrusion detection prototype described in [4].

Most contributions on intrusion detection (ID) for WNs mainly focus on the performance of the ID algorithm (*e.g.* true positives versus false positives) without considering its overhead that includes the processing (CPU), memory, and transmission bandwidth (in the case of a collaborative ID system), required for the algorithm to operate. Overhead minimization directly affects energy consumption as a smaller overhead leads to a more energy efficient ID. Energy efficiency in communication networks has been the focus of the research community as the latest EU directives have set the optimistic target to reduce carbon emissions by 20% until 2020.

In this paper, we propose an ID scheme that uses far fewer SINR measurements than existing techniques. Fewer measurements are expected to reduce the total energy consumption of the ID system. Supposing $N \in \mathbb{R}^+$ is the original number of SINR values used for ID, we investigate the performance of the ID algorithm when using only $M \in \mathbb{R}^+$ measurements, where $M \ll N$. The M measurements constitute a compressed version of the N-length original SINR signal obtained using the novel signal processing framework of *Compressed Sensing* (CS) [5], [6], [7]. CS presents a new method to capture and represent compressible signals at a rate significantly below the Nyquist rate. It has already been used in several areas in wireless communications including positioning [8], routing [9], video streaming [10], and signal reconstruction [11].

In this study, we explore the use of CS for effective ID considering an off-the-shelf jammer that emits energy on neighboring channels, attacking an IEEE 802.11 network. Our jammer uses ATH5K [12], an open source wireless driver. We have disabled several features of the IEEE 802.11 like backoff, CCA and carrier-sensing in the jammer, making it immune to the legitimate traffic; thus, it can freely perform jamming.

Our contributions are the following:

- we design a novel ID algorithm that uses measurements which are compressed and then reconstructed according to CS principles,
- we discuss how to select the various CS parameters for high performance ID,
- we evaluate the ID algorithm showing that high performance can be achieved even when far fewer measurements are used.

The remaining of this paper is organized as follows. Section II describes the related work in the field of ID for WNs. Section III provides an overview on CS theory. In Section IV, we introduce the new CS-based ID algorithm. The network testbed used, the jamming model, and the performance evaluation results are presented in Section V, while conclusions and further work appear in Section VI.

II. Related work

Significant efforts have been made by the research community for the detection of physical-layer jamming attacks. In [13], the authors describe several types of jammers, proposing two types of ID algorithms using as metrics the (i) packet delivery ratio, (ii) bad packet ratio, and (iii) energy consumption amount. Appropriate *if-else* statements are used by the ID algorithm along with the aforementioned metrics in order to detect the attacks. A more advanced scheme uses a distributed mechanism where information is collected from neighboring nodes.

The authors in [14] show how the errors from the physical layer propagate up to the application layer in the presence of jammers. They propose a distributed ID system based on simple thresholds using the Pearson's product of the received RSSI measurements. A major drawback of this method is that extensive control traffic (RSSI measurements) has to flow from the monitors to the sink node where the ID algorithm executes.

In [2], the authors consider four different types of jammers (constant, deceptive, random, and reactive), all focusing on the physical layer. The proposed detection algorithms are based on thresholds using the signal strength and location information as a consistency check to avoid false alarms.

ARES, an anti-jamming reinforcement system for IEEE 802.11 networks is presented in [15]. A random jammer emits energy into the network aiming to cause a DoS attack. The authors propose an attack mitigation scheme that tunes the parameters of rate adaptation and power control in order to improve network's performance when a jammer is present.

Li *et al.* [16] consider a scenario where a sophisticated jammer that controls the probability of jamming and the transmission range in order to cause maximal performance deterioration, jams an area in which a single-channel wireless network operates. For the detection of this type of attack, several observes are used that monitor the ratio of the corrupted packets over the correctly decoded ones. After a training period (that could be quite long), Walds Sequential Probability Ratio Test (SPRT) is used to signal when the network is under a jamming attack.

In [17], a distributed ID system is presented where the nodes monitor and then share among them several events such as transmission retries, transmitted packets, channel idle time, etc. Then, at each monitor, the events are matched and potential jamming attacks are detected. This ID scheme introduces a significant overhead to all nodes as they have to participate in the ID process and furthermore, it can create an excess control traffic within the network if the number of nodes is high.

Although all the aforementioned contributions are significant, none of them considers techniques in order to minimize the number of required measurements, and hence to reduce ID algorithmic overhead. Recently, there has been a number of studies that consider CS for event detection and energy efficient information collection in WNs. In [18], the authors propose a scheme for detecting M events (*e.g.*, liquid detection) using a network with N sensors, where $M \ll N$. The sensor measurements are correlated, and their relationship depends on the distance, the fading vector, and the thermal noise. This event detection problem is converted into a CS problem, where M sensors provide sensed data that are then reconstructed in a sink node.

In [19], a CS-based collection scheme is proposed. Here, a sink node generates a single packet that contains the appropriate coefficients for the linear projection of CS. As this packet traverses the nodes specified into the source route, each of the nodes multiplies its measurement by the corresponding coefficient and then adds it to the total sum contained within the packet. Using this method for the collection of data from N sensors, only M of them are polled and the N-length signal is reconstructed at the sink node.

In [20], a scheme for data acquisition though joint CS and principal component analysis (PCA) is presented. PCA is used in order to find the appropriate transformations for the CS operation. Again, the goal is to recover a given N-dimensional signal through the reception of a small number of samples. Finally, in [8], the authors exploit the signal correlation structure observed in an indoor localization environment in order to provide accurate position estimation by means of a limited amount of signal-strength measurements. They apply a distributed CS (DCS) scheme that rests on the joint sparsity of a signal ensemble and provides effective signal recovery by jointly reconstructing all the signals precisely.

All of the above contributions focus on principled ways to collect information from a WN in an energy efficient manner using distributed CS-based schemes. In this paper, we use CS in order to compress and then reconstruct the SINR signal at the location where the ID algorithm executes.

III. Compressed sensing background

The recently introduced theory of CS exploits the signal structure in order to enable a significant reduction in the sampling and computation costs at a central unit [6]. The key principles in the development of CS theory are *sparsity* and *incoherence*.

A signal $\mathbf{x} \in \mathbb{R}^N$ is called sparse if most of its elements are zero in a specific transform basis. In practice, we consider S-compressible signals where $N - S$ elements are very close to zero. Incoherence satisfies the fact that the sampling/sensing waveforms have an extremely dense representation in the basis. The discrete signal $\mathbf{x} \in \mathbb{R}^N$ can be expressed in terms of a sparsifying basis (dictionary) $\mathbf{\Psi}$ of $N \times 1$ vectors $\{\psi_{i=1}^N\}$ such that:

$$\mathbf{x} = \mathbf{\Psi}\mathbf{b} \tag{1}$$

where $\mathbf{b} \in \mathbb{R}^N$ is a sparse vector with S non-zero components ($\|\mathbf{b}\|_0 = S$). CS theory proves that an S-sparse signal $\mathbf{x}$ can be reconstructed exactly with high probability from M randomized linear projections of the signal $\mathbf{x}$ into a measurement matrix $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$. The general measurement model is expressed as follows:

$$\mathbf{y} = \mathbf{\Phi}\mathbf{x} = \mathbf{\Phi}\mathbf{\Psi}\mathbf{b} = \mathbf{\Theta}\mathbf{b} \tag{2}$$

where $\mathbf{\Theta} = \mathbf{\Psi}\mathbf{\Phi}$.

A necessary condition for accurate reconstruction of the original signal is that matrix $\mathbf{\Theta}$ must satisfy the so-called *restricted isometry property* (RIP). Particularly, incoherent matrices satisfy the RIP with very high probability. The coherence between the measurement matrix $\mathbf{\Phi}$ and the dictionary (or transformation matrix) $\mathbf{\Psi}$ is given by:

$$\mu(\mathbf{\Phi}, \mathbf{\Psi}) = \sqrt{N} \cdot \max_{1 \le k,j \le N} |\phi_k^T \psi_j|. \tag{3}$$

Coherence $\mu \in [1, \sqrt{N}]$ serves as a rough characterization of the degree of similarity between the transformation and measurement matrices. The two matrices $\mathbf{\Psi}$ and $\mathbf{\Phi}$ are highly incoherent if $\mu(\mathbf{\Phi}, \mathbf{\Psi}) \simeq 1$. In practice this means that the rows of $\mathbf{\Phi}$ cannot sparsely represent the columns of $\mathbf{\Psi}$ (and vice versa). It has been proved that both RIP and incoherence could be achieved by selecting the measurement matrix $\mathbf{\Phi}$ as a random matrix, *e.g.*, the elements of vectors ϕ_j could be independent and identically distributed (i.i.d.) random Gaussian variables.

When the above conditions hold, the original vector $\mathbf{b}$ and consequently the sparse signal $\mathbf{x}$, is estimated by solving the following ℓ_0-norm constrained optimization problem:

$$\hat{\mathbf{b}} = \arg\min \|\mathbf{b}\|_0 \quad s.t. \quad \mathbf{y} = \mathbf{\Theta}\mathbf{b} \tag{4}$$

where the $\|\mathbf{b}\|_0$ norm counts the number of non-zero components of $\mathbf{b}$. Note that the formulation of the optimization problem in (4) uses an l_0 norm that measures signal sparsity instead than the traditionally used in signal processing applications l_2 norm, which measures signal energy.

Unfortunately, solving (4) is both numerically unstable and NP-complete. Luckily, the ℓ_0 norm can be replaced by the ℓ_1 norm and problem (4) can be rephrased as the following ℓ_1 norm convex relaxation problem:

$$\hat{\mathbf{b}} = \arg\min \|\mathbf{b}\|_1 \quad s.t. \quad \mathbf{y} = \mathbf{\Theta}\mathbf{b}. \tag{5}$$

The ℓ_1 norm ($\|\mathbf{b}\|_1 := \sum_i |b_i|$) can exactly recover the S-sparse signal with high probability using only $M \geq CS\log(N/S)$ measurements ($C \in R^+$) [6]. Finally, the reconstructed signal is given by $\hat{\mathbf{x}} = \mathbf{\Psi}\hat{\mathbf{b}}$.

A variety of reconstruction algorithms based on linear programming, convex relaxation, and greedy strategies have been proposed to solve (5). Among them, greedy strategies (*e.g.* Orthogonal Matching Pursuit (OMP) [21]) are computationally efficient when the signal of interest is highly sparse.

IV. Intrusion detection using compressed sensing

A. Cumulative-sum algorithms for intrusion detection

A number of ID algorithms have been proposed in the literature. In our previous studies, we investigated several metrics based on the SINR (or SNR) [4], [22], [23]. We found that the max-min approach, which considers the maximum-minus-the minimum value of the SINR in a short window, and its average value in a long window, achieves the best performance.

Furthermore, we compared two types of algorithms: (i) simple threshold algorithms that trigger an alarm when an appropriately chosen metric deviates from its normal (expected) value by a certain amount, and (ii) Cusum algorithms that raise an alarm if the aggregated output exceeds a predefined threshold. We have found that Cusum algorithms achieve higher performance than the simple ones.

The Cusum algorithm belongs to the family of change-point detection techniques based on hypothesis testing, first introduced in [24]. It has the ability to detect abrupt changes and is of two types: parametric and non-parametric. Parametric Cusum requires the distribution of a metric (*e.g.* a metric based on the SINR) to be known in advance, so before a change takes place. On the other hand, non-parametric Cusum can detect abrupt changes even when the distribution of this metric is not known in advance. Since the SINR distribution cannot be known in advance, as it is volatile in nature and depends on several factors such as the network topology, the distance between the nodes, etc., we use the non-parametric Cusum for ID, referring to it as Cusum max-min (C_{mm}). C_{mm} is given by the following formula:

$$O_n = \begin{cases} O_{n-1} + Z_n - a & \text{if } O_n \geq 0 \\ 0 & \text{if } O_n < 0 \end{cases} \tag{6}$$

SINR samples are recorded sequentially by a monitor node

that executes the C_{mm} algorithm (cf. Figure 6). For each received SINR measurement n, the output O_n of the algorithm is affected by three parameters: (i) its previous output O_{n-1}, (ii) the so-called expectation Z_n, and (iii) $\alpha \in \mathbb{R}^+$.

Z depends on the SINR measured values x and is given by $Z_n = D(n) - \bar{D}(n)$, where

$$D(n) = \max_{n-K+1<i\leq n} x_i - \min_{n-K+1<i\leq n} x_i \quad (7)$$

and

$$\bar{D}(n) = \frac{\sum_{i=n-L+1}^{n} D(i)}{L}. \quad (8)$$

Z is expected to have a negative drift before a change and a positive drift after the change. In the context of this work, changes on the SINR-based metric (Z) take place when a jammer is active. Figure 1a shows how SINR drops during the jamming attacks, depicted by the orthogonal boxes. The effect on Z is shown in Figure 1b: Z's drift changes from negative to positive when an attack is taking place. Parameter α controls this drift; the larger it is, the smaller the drift becomes, as shown in (6).

K and L are the lengths of the short and long windows, respectively. An alarm is raised when $O_n \geq h$, where h is a predefined detection threshold. Figure 1c shows how C_{mm}'s output changes during the attacks. In this figure, as well as for the evaluation presented in Section V-B, we have empirically selected: $K = 10$, $L = 100$, and $a = 0.5$.

B. A CS-based cumulative-sum algorithm for intrusion detection

In this section, we introduce a novel C_{mm} CS-based ID algorithm. The block diagram of the proposed method is shown in Figure 2. The dashed shapes show the new operations that belong to CS, while the solid ones implement the basic ID functionalities. The SINR signal ($\mathbf{x} \in \mathbb{R}^N$) is initially projected through the measurement matrix $\mathbf{\Phi}$ to obtain signal $\mathbf{y} \in \mathbb{R}^M$ that is then sent to the C_{mm} algorithm through the communication medium. This medium can be of several types, depending on the implementation. For example, it can be the wireless medium in a collaborative ID system where several monitors collect and transmit information to a fusion center. It could also have the form of the kernel/user space Linux interface for an ID system implemented in a single monitor, as in [4]. However, regardless the type of the communication medium, valuable resources can be saved (*e.g.*, memory, CPU, bandwidth, energy, etc.) because CS provides a compressed version $\mathbf{y} \in \mathbb{R}^M$ of the original signal $\mathbf{x} \in \mathbb{R}^N$ ($M \ll N$), and successfully reconstructs it at the other end (just before C_{mm}).

C. Transformation matrix and compressibility

As mentioned in Section III, CS assumes that the signal is sparse in a certain domain in order to effectively reconstruct it from far fewer measurements. Hence, to efficiently apply

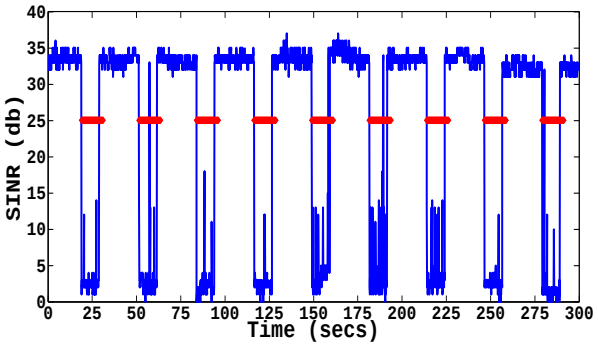

(a) SINR variations during the attacks.

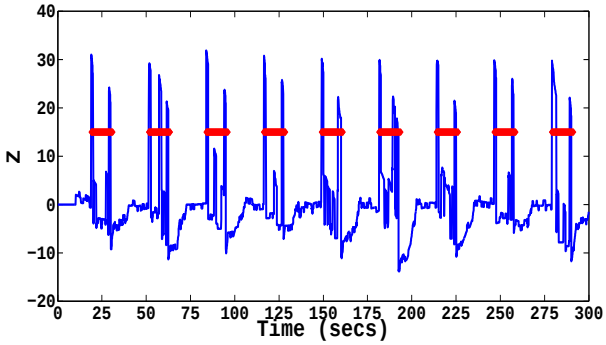

(b) Change of expectation's drift from negative to positive during the attacks.

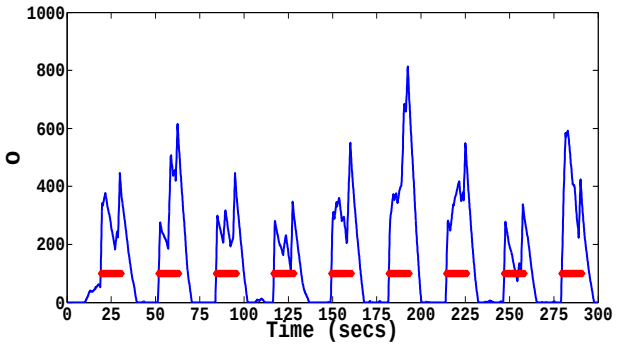

(c) C_{mm}'s output signalling the presence of a jammer.

Figure 1: Effect of the attacks on the SINR, expectation and C_{mm}'s output.

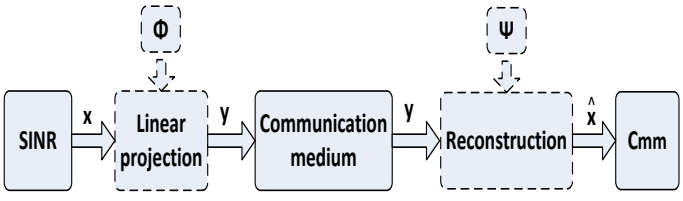

Figure 2: CS-based intrusion detection block diagram.

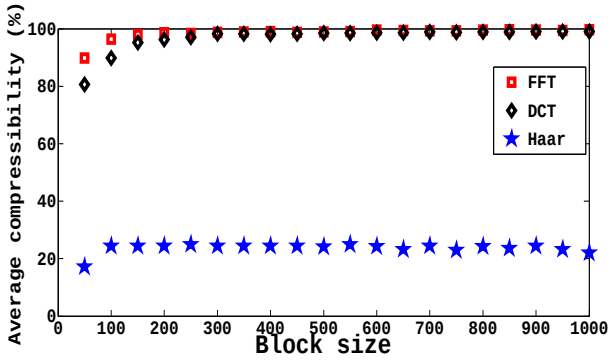

Figure 3: Average compressibility for the three different transforms and for an increasing block size.

CS, one should seek for the most appropriate transformation matrix $\mathbf{\Psi}$ for the SINR signal. Suppose that the SINR signal $\mathbf{x}$ can be expressed using (1). We define as *compressibility* index the percentage of the values of vector $\mathbf{b}$ (normalized by their maximum value) that are close or equal to zero. Compressibility is proportional to the sparseness of the signal.

Initially, we consider an SINR signal of a specific length collected using the network testbed shown in Figure 6. We divide this signal into equally-sized blocks of length N. For each block, we compute the compressibility of the signal using three different transforms: (i) the fast fourier transform (FFT), (ii) the discrete cosine transform (DCT), and (iii) the Haar wavelet transform. Then, we compute the average compressibility (ACO) for each of the transforms, repeating the procedure for different block sizes N. Figure 3 shows the ACO of actual SINR data for the three tested transforms as a function of the block size. We observe that the ACO index corresponding to the FFT is higher than the ACO indexes of both the DCT and the Haar wavelet transforms. Therefore, in the rest of this paper, we consider FFT as the transformation matrix $\mathbf{\Psi}$ of choice when implementing our CS-based scheme.

D. Measurement matrix and reconstruction error

CS performance (accurate signal reconstruction) depends also on the type of the measurement matrix $\mathbf{\Phi}$. Recent work has shown that when considering measurement matrices built using values selected independently from certain distributions, exact signal recovery can be achieved with high probability. One such choice is the Gaussian distribution used in several works (*e.g.* [8]). However, the generation of a Gaussian distribution may not be easily achieved in practical implementations (*e.g.*, in a ID system). Bajwa *et al.* [25] show that Toeplitz matrices with entries drawn from the same distributions (*e.g.* Gaussian) are also sufficient to recover a signal with high probability. The use of Toeplitz matrices can be attractive for a number of reasons [25]:

- a Toeplitz matrix requires the generation of $O(N)$ random variables, while i.i.d. matrices require the generation of $O(MN)$ variables.
- multiplication with a Toeplitz matrix can be performed using FFT and requires only $O(Nlog_2(N))$ operations compared to i.i.d. matrices that require $O(MN)$ operations.
- i.i.d. matrices are not easily applicable in certain scenarios (*e.g.*, linear-time invariant systems).

The characteristic of Toeplitz matrices is that all elements belonging to the same diagonal have the same value. As $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$, we use partial Toeplitz matrices of the form:

$$\mathbf{\Phi} = \begin{bmatrix} a_N & a_{N-1} & \cdots & a_2 & a_1 \\ a_{N+1} & a_N & \cdots & a_3 & a_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{N+M-1} & a_{N+M-2} & \cdots & a_{M+1} & a_M \end{bmatrix}.$$

The authors in [25] use synthetic signals to compare CS performance in terms of the empirical probability of success when (i) the entries of the measurement matrix are drawn independently from a Bernoulli distribution, and (ii) a Toeplitz matrix is used with entries drawn from the same distribution. In this work, we use real SINR signal values to investigate the performance of the C_{mm} algorithm, as well as the reconstruction error for two types of measurement matrices (i) a Gaussian matrix with entries drawn from an i.i.d. Gaussian distribution with zero mean and variance equal to $\frac{1}{M}$, and (ii) a Toeplitz matrix with entries drawn from the same distribution. The reconstruction error is defined as $e = \frac{||x-\hat{x}||_2}{||x||_2}$, where x and $\hat{x}$ are the original and reconstructed SINR signals, respectively.

Figure 4 shows the cumulative density function (CDF) of the reconstruction error for two different attack intensities (high, medium), two measurement matrices (Gaussian and Toeplitz), and for various compression ratios (defined as $100 \times (1 - \frac{M}{N})$), using the OMP method for reconstruction. The results were obtained by considering 100 Monte Carlo runs. We observe that as the compression ratio increases, the reconstruction error increases for both attack intensities. The use of a Toeplitz measurement matrix results in a higher reconstruction error than the one produced by a Gaussian matrix. This is because, as shown in Figure 5, the Toeplitz matrix exhibits higher coherence μ (cf. Eq. (3)) with the transformation matrix $\mathbf{\Psi}$ (FFT), than the Gaussian matrix does.

Furthermore, for all compression ratios, the reconstruction error is higher for the high intensity attack. This is due to the large variations of the SINR during this attack (Figure 1a) that affect its ACO.

V. Jamming model- experimental network testbed and performance evaluation

A. Jamming model and the experimental network testbed

Figure 6 shows the experimental network testbed used for SINR collection. We consider four nodes that communicate

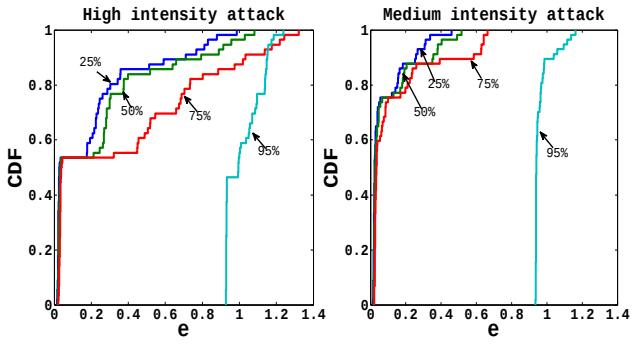

(a) Reconstruction error with the Gaussian matrix for the two intensity attacks.

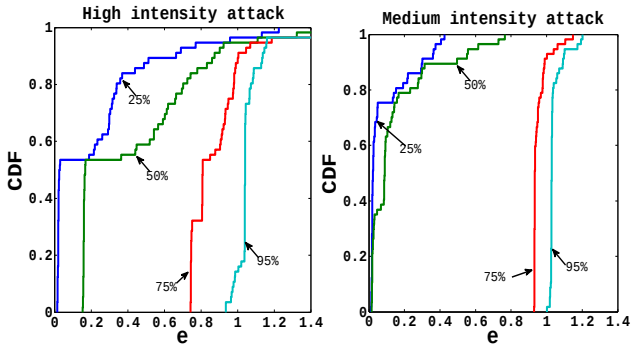

(b) Reconstruction error with the Toeplitz matrix for the two intensity attacks.

Figure 4: Reconstruction error for different measurement matrices, attack intensities, and compression ratios.

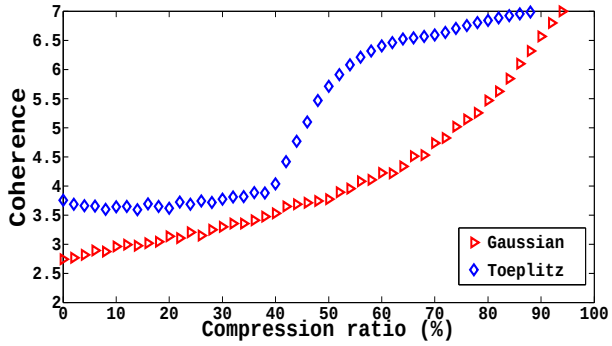

Figure 5: Coherence with the transformation matrix Ψ for two different measurement matrices.

through an access point (AP), a monitor node (MN) that collects the SINR measurements and executes the ID algorithm, and a jammer that emits energy on the immediately adjacent channel (IAC) that legitimate nodes use for communication (i.e., main channel). Nodes communicate in pairs through the AP at a constant rate of 1.5 Mbps.

We consider two types of attack intensities, namely, (i) high-intensity attacks, where the jammer broadcasts UDP traffic on the IAC with a transmission rate of 3 Mbps; and (ii) medium-intensity attacks, where the jammer emits energy with a rate of 1.5 Mbps. MN hosts five wireless interfaces in order to

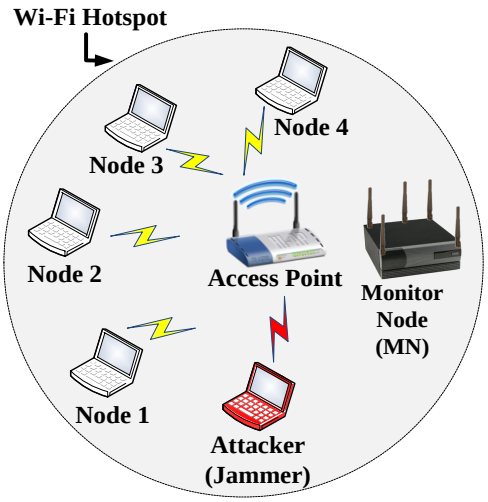

Figure 6: Experimental network testbed for SINR measurement collection.

record the main channel, the IACs, as well as the next adjacent channels. MN computes the SINR each time it captures a beacon packet transmitted by the AP. The interested reader can refer to [4] for more details on the SINR computation. All nodes use the same hardware based on a mini-ITX board that carries 512 MB of RAM and an 80 GB hard disk. This board is also equipped with an Atheros CM9-GP mini-PCI card, controlled by the ATH5K driver, running on Gentoo Linux. We note here that jammer operates in a periodic fashion, alternating between sleeping and jamming.

B. *Performance evaluation*

In this section, we present the performance evaluation of the CS-based C_{mm} algorithm. In general, the performance of a detection algorithm is evaluated by considering the RoC curves that show the trade-off between detection probability and false alarm rate for various detection thresholds. As this method is predicated on subjective criteria, we consider the F-measure, first introduced in [26] that allows for a quantitative analysis of the performance results. The F-measure $F \in [0, 1]$ is given by:

$$F = (1 + c^2) \times \frac{recall \times precision}{c^2 \times recall + precision} \qquad (9)$$

where $recall = \frac{t_p}{t_p + f_n}$ gives the ratio of true positives (t_p) over the false negatives (f_n); *precision* is the ratio of the true positives over the false positives (f_p); and $c \in \mathbb{R}^+$ controls the trade-off between precision and recall (we set $c = 1$). The higher the F-measure, the higher the performance achieved by the algorithm.

Table I shows the achieved F-measure for the two attack intensities (high and medium), the two measurement matrices (Gaussian and Toeplitz), and for various compression ratios. Note that compression ratio equal to 0 means that

TABLE I: The F-measure for different attack intensities, measurement matrices and compression ratios

Compression ratio	F-measure			
	High attack		Medium attack	
	Gaussian	Toeplitz	Gaussian	Toeplitz
0% (original SINR signal)	1	1	0.9987	0.9966
25%	1	0.9842	0.9307	0.8457
50%	0.9614	0.9370	0.8956	0.8342
75%	0.9121	0.8247	0.8453	0.8241
95%	0.8376	0.7357	0.7657	0.7082

no compression has been employed and the original SINR signal is used. We vary the detection threshold h from zero up to a maximum value that C_{mm} gives no alarms. The F-measure can vary depending on the threshold h, so in this table we show the maximum F-measure achieved. We observe that as the compression ratio increases, the performance of C_{mm} degrades gracefully. Also, when the Gaussian matrix is used, C_{mm} achieves higher performance than when a Toeplitz matrix is employed. This is expected because, as we saw in Section IV-D, the use of a Toeplitz measurement matrix results in a higher SINR signal reconstruction error.

We also note that C_{mm} achieves better performance during high intensity attacks, as compared to its performance during medium intensity attacks, although the reconstruction error is lower for the latter attacks. This can be explained considering that C_{mm} is used for ID because of its inherent ability to detect abrupt changes. The performance of the CS-based version depends on both the SINR variations and the reconstruction error. During the high intensity attack, there are large variations of the SINR and the reconstruction error is high. On the other hand, during the medium intensity attack, the SINR variations are not as large and the reconstruction error is low. From these results we conclude that the SINR variation dominates the reconstruction error, so C_{mm} achieves better performance in high intensity attacks. This is depicted in Figure 7 that shows the original SINR signal and its reconstructed version using the Gaussian matrix, with different compression ratios considered, and for the high intensity attack. The abrupt changes of the SINR are preserved through the CS process.

All in all, even for high compression ratios ($> 75\%$) the CS-based C_{mm} still achieves a high F-measure (> 0.7) for both measurement matrices considered.

VI. Conclusions-Further work

In this work, we investigated the feasibility of ID at the physical layer of a wireless network, using a change-point ID algorithm based on the SINR along with CS. CS allows the reconstruction of the SINR signal from far fewer measurements, thus minimizing algorithmic overhead as the transmission of the whole signal is not necessary. Using a real testbed for measurement collection, we showed that the SINR signal is highly sparse in the FFT domain. Furthermore, we studied ID's performance using two types of measurement matrices, for different compression ratios, and under two attack intensities.

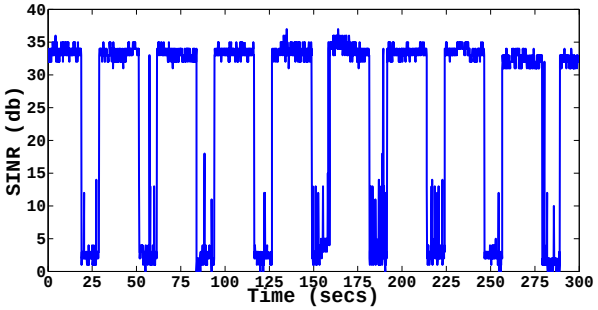

(a) Original SINR

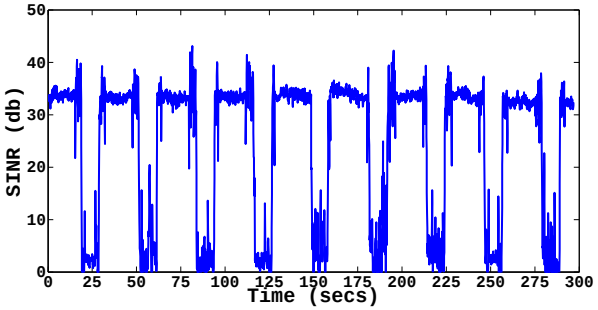

(b) Reconstructed SINR with a compression ratio of 75% .

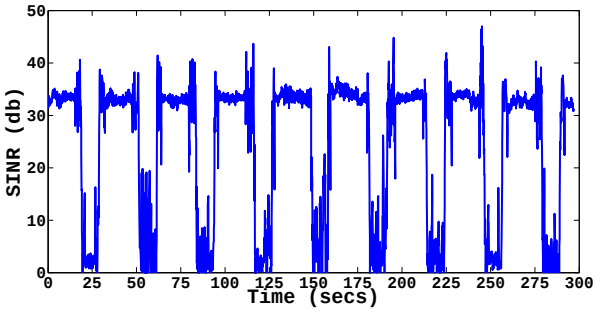

(c) Reconstructed SINR with a compression ratio of 50% .

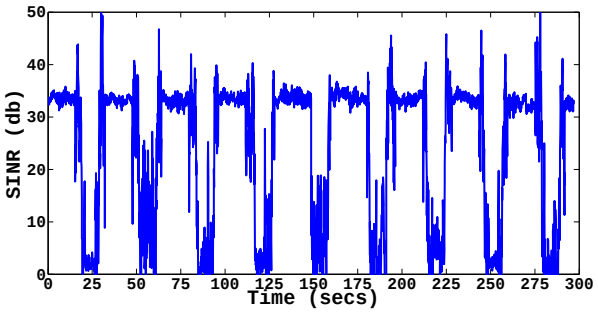

(d) Reconstructed SINR with a compression ratio of 25% .

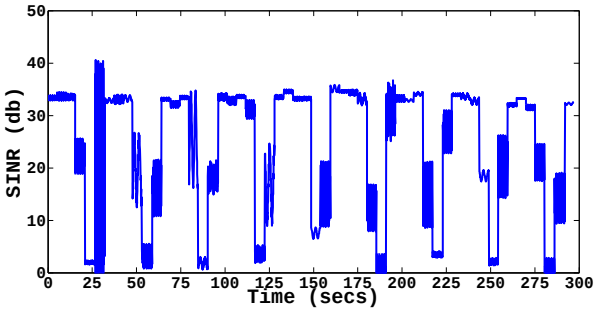

(e) Reconstructed SINR with a compression ratio of 5% .

Figure 7: SINR for different compression ratios when the Gaussian matrix is considered.

The results show that when the Gaussian matrix is used, the reconstruction error is smaller than when using the Toeplitz matrix. Also, the performance evaluation, in terms of the F-measure, shows that with both matrices, C_{mm} achieves high performance for both attack intensities, and for fairly high compression ratios. Furthermore, during high intensity attacks, although the SINR reconstruction error is larger than that of medium intensity attacks, C_{mm} achieves better performance because the large SINR variations dominate the reconstruction error.

Further work will include the performance evaluation of the CS-based ID algorithm when using more different measurement matrices such as the circulant and the structurally random ones. Also, we aim to investigate the trade-off between energy saving and performance in a real testbed for the various CS-based ID schemes.

REFERENCES

[1] A. Fragkiadakis, E. Tragos, and I. Askoxylakis, "Video streaming performance in wireless hostile environments," in *Proc. of the 5th FTRA International Conference on Multimedia and Ubiquitous Engineering*, June 2011, pp. 267–272.

[2] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," in *Proc. of ACM MobiHoc*, May 2005, pp. 47–57.

[3] K. Pelechrinis, M. Iliofotou, and S. Krishnamurthy, "Denial of service attacks in wireless networks: The case of jammers," *IEEE Communications Surveys & Tutorial*, vol. 13, no. 2, pp. 245–257, 2011.

[4] A. Fragkiadakis, E. Tragos, T. Tryfonas, and I. Askoxylakis, "Design and performance evaluation of a lightweight wireless early warning intrusion detection prototype," *EURASIP Journal on Wireless Communications and Networking*, vol. 2012, pp. 1–18, 2012.

[5] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, pp. 1289–1306, 2006.

[6] E. Candes and M. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.

[7] J. Haupt, W. Bajwa, M. Rabbat, and R. Nowak, "Compressed sensing for networked data," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.

[8] S. Nikitaki and P. Tsakalides, "Localization in wireless networks based on jointly compressed sensing," in *Proc. of EUSIPCO*, August 2011, pp. 1–5.

[9] G. Quer, R. Masiero, D. Munaretto, M. Rossi, J. Widmer, , and M. Zorzi, "On the interplay between routing and signal representation for compressive sensing in wireless sensor networks," in *Proc. of the Information Theory and Applications Workshop (ITA)*, February 2009, pp. 1–10.

[10] S. Pudlewski, A. Prasanna, and T. Melodia, "Compressed-sensing-enabled video streaming for wireless multimedia sensor networks," *IEEE Transactions on Mobile Computing*, vol. 11, pp. 1060–1072, 2011.

[11] G. Quer, D. Zordan, R. Masiero, M. Zorzi, and M. Rossi, "Wsn-control: signal reconstruction through compressive sensing in wireless sensor networks," in *Proc. of SenseApp*, October 2010, pp. 921–928.

[12] "Linux wireless drivers, ath5k, http://linuxwireless.org/en/users/Drivers/ath5k."

[13] M. Cakiroglou and T. Ozcerit, "Jamming detection mechanisms for wireless sensor networks," in *Proc. of 3rd Int. Conference on Scalable Information Systems*, Napoli, Italy, June 2008.

[14] A. Sheth, C. Doerr, D. Grunwald, R. Han, and D. Sicker, "MOJO: a distributed physical layer anomaly detection system for 802.11 WLANs," in *ACM MobiSys*, 2006.

[15] K. Pelechrinis, I. Broustis, S. Krishnamurthy, and C. Gkantsidis, "Ares: an anti-jamming reinforcement system for 802.11 networks," in *Proc. of CoNEXT*, 2009, pp. 181–192.

[16] M. Li, I. Koutsopoulos, and R. Poovendran, "Optimal jamming attack strategies and network defense policies in wireless sensor networks," *IEEE Transactions on Mobile Computing*, vol. 9, pp. 1119–1133, 2010.

[17] M. Aime, G. Calandriello, and A. Lioy, "A wireless distributed intrusion detection system and a new attack model," in *Proc. of ISCC*, 2006, pp. 35–40.

[18] Y. Liu, X. Zhu, C. Ma, and L. Zhang, "Multiple event detection in wireless sensor networks using compressed sensing," in *Proc. of ICT*, 2011, pp. 27–32.

[19] C. Chou, R. Rana, and W. Hu, "Energy efficient information collection in wireless sensor networks using adaptive compressive sensing," in *Proc. of LCN*, 2009, pp. 443–450.

[20] R. Masiero, G. Quer, D. Munaretto, M. Rossi, J. Widmer, and M. Zorzi, "Data acquisition through joint compressive sensing and principal component analysis," in *Proc. of Globecom*, 2009, pp. 1–6.

[21] J. Tropp and A. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Transactions on Information Theory*, vol. 53, pp. 4655–4666, 2007.

[22] A. Fragkiadakis, V. Siris, and A. Traganitis, "Effective and robust detection of jamming attacks," in *Proc. of the Future Network and Mobile Summit*, June 2010, pp. 1–8.

[23] A. Fragkiadakis, V. Siris, and N. Petroulakis, "Anomaly-based intrusion detection algorithms for wireless networks," in *Proc. of the of the 8th International Conference on Wired/Wireless Internet Communications*, June 2010, pp. 192–203.

[24] E. Page, *Continuous inspection schemes*. Biometrika, 1954.

[25] W. Bajwa, J. Haupt, G. Raz, S. Wright, and R. Nowak, "Toeplitz-structured compressed sensing matrices," in *Proc. of SSP*, 2007, pp. 295–298.

[26] C. van Rijsbergen, *Information Retrieval*. London: Butterworths, 1979.