This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

Swift Jamming Attack on Frequency Offset Estimation: The Achilles' Heel of OFDM Systems

Hanif Rahbari, Marwan Krunz, *Fellow, IEEE,* and Loukas Lazos

Abstract—Frequency offset (FO) refers to the difference in the operating frequencies of two radio oscillators. Failure to compensate for the FO may lead to decoding errors, particularly in OFDM systems. To correct the FO, wireless standards append a publicly known preamble to every frame before transmission. In this paper, we demonstrate how an adversary can exploit the known preamble structure of OFDM-based wireless systems, particularly IEEE802.11a/g/n/ac, to launch a very stealth (low energy/duty cycle) reactive jamming attack against the FO estimation mechanism. In this attack, the adversary quickly detects a transmitted OFDM frame and subsequently jams a tiny part of the preamble that is used for FO estimation at the legitimate receiver. By optimizing the energy and structure of the jamming signal and accounting for frame detection timing errors and unknown channel parameters, we empirically show that the adversary can induce a bit error rate close to 0.5, making the transmission practically irrecoverable. Such vulnerability to FO jamming exists even when the frame is shielded by efficient channel coding. We evaluate the FO estimation attack through simulations and USRP experimentation. We also propose three approaches to mitigate such an attack.

Index Terms—PHY-layer security, frequency offset, OFDM, reactive jamming, IEEE802.11, USRP implementation.

1 Introduction

Communication between two wireless devices involves several concerted functions at the physical (PHY) layer, including time synchronization, carrier frequency offset (FO) correction, channel estimation, channel coding, modulation, interleaving, and others [2]. PHY-layer functions are designed to combat oscillator imperfections and wireless channel impairments, and to decode wireless signals that are corrupted by a limited amount of interference. However, wireless transmissions still remain vulnerable to intentional interference attacks, commonly referred to as jamming.

One measure of the effectiveness of a jamming attack is its duty cycle, i.e., the fraction of the frame that needs to be jammed so that the frame is discarded at the receiver (Rx) [3], [4]. This metric is directly related to the jammer's distance to the Rx, energy budget, and the ability to disrupt concurrent transmissions. A jammer that remains active for a longer period can corrupt more bits and defeat stronger error correction codes (ECCs), at the expense of higher energy consumption and fewer targeted communications. This more potent jammer is also easier to detect [5], localize, and physically remove using jammer localization methods [4].

In this paper, we investigate an extremely low duty cycle jamming model that is facilitated by public knowledge of the frame structure and PHY-layer functions. Our goal is to demonstrate how an adversary can inflict the highest possible number of decoding errors at the Rx, without jamming the corresponding header or payload symbols. PHY-layer standards usually employ publicly known sequences, known as *preambles*, at the beginning of a frame to acquire important communication parameters, such as the transmission timing, channel, and FO [2]. These parameters are used to align received symbols. An adversary may exploit the publicity of the preamble to construct a reactive jamming attack and target the estimation of these critical parameters. In particular, we demonstrate the feasibility of an energy-efficient and low duty cycle attack against the FO estimation process of IEEE 802.11 OFDM-based devices (including 802.11a, .11g, .11n, .11ac, and 11ah), all of which exploit the same preamble structure. Our results can be extended to other OFDM-based systems, including 802.16e/m (WiMAX), LTE, and 5G.

The jamming of OFDM systems has recently been the subject of extensive research (e.g., [6]–[12]). These works often consider vulnerabilities in time synchronization or susceptibility to inter-carrier interference (ICI). For example, the authors in [8] proposed several jamming attacks against OFDM time synchronization, including barrage attacks, false preamble timing, and preamble warping. In the barrage attack, white noise is transmitted to decrease the SNR during synchronization. In false preamble timing, the jammer forges a preamble to fool the Rx about the true start time of the frame. A similar technique was used in [9] against an 802.11b Rx to hamper the network throughput. Preamble warping tries to destroy the time-domain correlation (used for time acquisition) within the preamble.

1.1 Frequency Offset Estimation Attacks

In OFDM systems, frequency synchronization errors are more devastating than timing errors [13]. When two radios are tuned to the same target frequency, their oscillators can-

- *H. Rahbari, M. Krunz, and L. Lazos are with the Department of Electrical and Computer Engineering, University of Arizona, Tucson, AZ 85721. E-mail: {rahbari,krunz,llazos}@email.arizona.edu*
- *An abridged version of this paper appeared in the Proceedings of the IEEE INFOCOM Conference, April 2014 [1].*

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

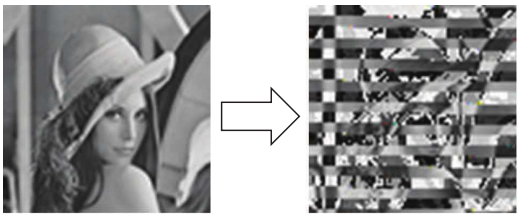

Fig. 1. Effect of uncompensated FO on a bitmap image over a noiseless channel (FO = 0.32% of the subcarrier spacing).

not be exactly aligned to that frequency due to hardware imperfections. FO is the inherent difference between the actual frequencies of these two oscillators. In OFDM, FO is usually normalized to the inter-subcarrier frequency interval, called *subcarrier spacing*. Without frequency synchronization, the performance of OFDM degrades severely because all subcarriers will move away from their expected frequencies, resulting in subcarriers' orthogonality violation, ICI [13], and channel estimation errors [1], [14].

To appreciate the significance of correct FO estimation, we conduct a simulation experiment in which a frame containing a bitmap image is transmitted between two nodes. Fig. 1 depicts the effect of a small FO estimation error (0.32% of subcarrier spacing) on the transmitted image (left) when 48 subcarriers are used at a rate of 6 Mbps. The received image (right) exhibits noticeable degradation in the form of image block misplacement. In practice, FO can be even larger than the subcarrier spacing [2].

A few jamming schemes have been proposed in the literature (e.g., [7], [9], [10]) with the goal of inflicting ICI. Phase warping and differential scrambling attacks [10] consider the preamble structure of Schmidl and Cox [15], which is different from the one used in 802.11 OFDM-based standards, and in essence try to alter preamble symbols in a heuristic fashion without providing any success guarantees. Gummadi *et al.* [9] showed the vulnerability of 802.11a clock (frequency) synchronization to a certain narrow-band jamming pattern that interferes with the entire preamble. In [7] the jammer transmits multiple asynchronous subcarriers to cause ICI in an OFDM symbol. These attacks may fail if robust ECC, interleaving methods, or additional FO estimation mechanisms are employed at the Rx.

1.2 Contributions

We design an energy-efficient jamming attack that interferes with a small portion of the preamble, i.e., one of the parts used for FO estimation, and causes one or two units *shift* of the subcarrier indices (e.g., every subcarrier takes the position of its next/previous subcarrier). To make this design possible, the adversary (Eve) must first estimate the FO between the legitimate transmitter (Alice) and intended receiver (Bob), and then quickly detect the transmission of a target frame. We provide an adaptive frame detection method to facilitate fast detection at Eve. The superposition of the jamming signal with the preamble are designed to delude Bob into estimating an FO that is sufficiently far from the true FO, so that Bob decodes wrong symbols, i.e., the symbols of adjacent subcarriers. The idea is to come up with a structure that is similar to the actual preamble so as to control the FO embedded in the jamming sequence. The superposition of these two signals with different FOs at the Rx achieves sufficient FO estimation error. We derive the amount of FO estimation error needed to guarantee erroneous OFDM demodulation and accordingly, develop an optimal attack strategy. To ensure that the jamming signal is independent of the Alice-Bob channel parameters (which are unknown to Eve), we propose a *pairing* scheme for the jamming sequence. The jamming attack should also account for timing errors in frame detection at Eve while keeping the jamming signal channel-independent. For this purpose, a *chaining* scheme is designed on top of the pairing scheme to account for other possible frame start times.

Consequently, not only the channel estimation is automatically corrupted at Bob, but more importantly, all the frequency subcarriers are shifted forward or backward. Hence, Bob will have a shifted version of the bitstream transmitted in every OFDM symbol. Combined with a faulty channel estimation and thus demodulation errors, the bits become irrecoverable. We further optimize the power of this jamming attack and experimentally evaluate its performance on a USRP testbed. In contrast to previous attacks on the frame preamble, ours in essence does not aim at necessarily causing ICI. It is also different from the attacks in [7], [9], [10] in that it is *channel-independent* and *energy-efficient*, i.e., only a small portion of the preamble is jammed irrespective of the jammer's location. This short-lived attack lasts for less than 3 μs per frame (equivalent to, for example, about 0.5% of 802.11a's maximum frame duration when the data rate is at its highest value). Note that this is even shorter than the duration of an OFDM symbol (4 μs). Our proposed attack also disarms all the provisioned FO estimation methods by just efficiently defeating one of them. Our work focuses on the 802.11 OFDM-based wireless systems, and efficiently exploits their FO vulnerability for the first time.

The paper is organized as follows. In Section 2, we provide background on frame detection, FO estimation, and channel estimation in OFDM-based 802.11 standards. The system model, assumptions, and evaluation metrics are given in Section 3. The proposed attack and the optimal jamming strategy are presented in Section 4 and related issues are discussed in Section 5. Section 6 demonstrates the effectiveness of the attack through simulations and experiments. Finally, we propose possible remedies and provide a summary of existing attacks in Sections 7 and 8, respectively.

2 Frame Detection and FO Correction in OFDM Systems

In OFDM, a bitstream is split into several substreams, each of which is digitally modulated and transmitted over one of the orthogonal frequency channels (subcarriers). For example, 802.11a/g defines 64 subcarriers with subcarrier spacing $f_\Delta = 312.5$ kHz within a bandwidth of 20 MHz. Only 48 of these subcarriers are used for data. Four other subcarriers carry pilot signals and the remaining 12 subcarriers are not used. So an 802.11a/g OFDM symbol is transmitted over 52 subcarriers.

ICI in OFDM systems creates significant BER at the Rx [16] (see Fig. 2). To prevent ICI, the Rx uses the PHY-layer preamble to estimate the FO (same for all subcarriers) and adjust the subcarriers to their expected orthogonal

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

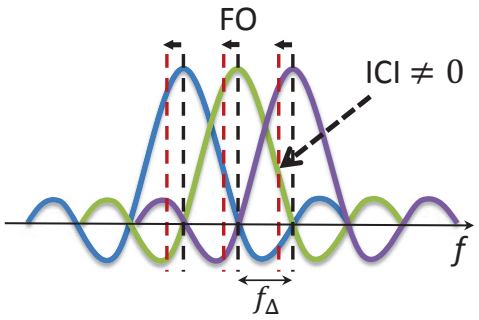

Fig. 2. Inter-carrier interference (ICI) as a result of uncorrected FO in a system with three subcarriers.

frequency bins. If the offset is less than half of the frequency distance between the subcarriers, the Rx can safely identify the frequency bin that each subcarrier belongs to.

Every PHY-layer frame starts with a preamble. In OFDM-based 802.11 systems, the preamble begins with two essential fields (see Fig. 3). The first field contains ten identical short training sequences (STSs), which represent ten replicas of a particular periodic function with period $\lambda_{STS} = 0.8\ \mu s$. The second field consists of two long training sequences (LTSs), which represent two cycles of another known periodic function with period $\lambda_{LTS} = 4\lambda_{STS}$, plus a $1.6\ \mu s$ cyclic prefix (GI)[1]. The periodic function in an STS is constructed by superposing only the subcarriers whose frequencies are integer multiples of $4f_\Delta$. As a result, the minimum subcarrier spacing between any two STS-enabled subcarriers is $4f_\Delta$, and hence their period is $\lambda_{STS} = \lambda_{LTS}/4$. STSs are used for frame detection and coarse FO correction. LTSs, on the other hand, employ all the data subcarriers and are used for channel estimation and fine-tuning the coarse STS-based FO estimation.

We briefly explain the channel estimation process in OFDM-based 802.11 systems because it is affected by the coarse FO estimation. LTSs are used for channel estimation, the task of estimating the response of the channel, because they are supposed to be almost FO-free after STS-based FO correction. There are two general approaches for channel estimation: Frequency domain and time domain [13]. In both approaches, the a priori known LTS symbols are compared with the received symbols in order to estimate the impulse or frequency response that results in the minimum mean-square-error (MSE). The MSE can grow quadratically as a function of the FO estimation error [14].

2.1 FO Estimation and Correction

Let Δf be the actual frequency offset between a transmitter (Tx) and an Rx. This FO translates into a phase offset of $\Delta\varphi(t) = 2\pi\Delta f t$ for the received signal, where t is the time elapsed since the start of the transmission. In addition to causing ICI, a linear increase in the phase offset during the LTSs due to FO results in incorrect channel phaser estimation. To compensate for channel impairments, the inverse of the phaser is multiplied to the received samples. As a result, all received modulated samples will be rotated equally on

1. In MIMO-OFDM systems, these two fields are followed by additional training sequences for MIMO channel estimation [17].

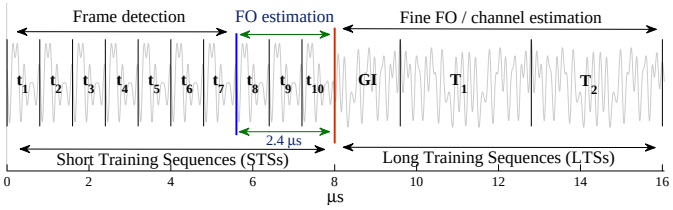

Fig. 3. Time-domain representation of a common preamble structure in 802.11a/g/n/ac systems (20 MHz bandwidth).

the constellation map, leading to more bit errors. Beyond channel estimation errors, accumulation of the phase offset can significantly change the phase of some of the symbols, especially in long frames.

The *de facto* time-domain FO estimation method used in OFDM systems is the one proposed by Schmidl and Cox [15]. We consider it as a representative FO estimation scheme. It assumes that the channel does not change during the preamble transmission. Having a sequence $\mathbf{r}$ with two identical halves is the key idea in this method. It works as follows. Assume that each half of the sequence has L samples with sampling period of t_s. Let r_i be the ith sample of the sequence $\mathbf{r}$, $i = 1, \ldots, 2L$. So $r_i = r_{L+i}$. Ignoring the noise, this equality also holds for the corresponding samples at the Rx as long as there is no FO. However, with an FO of Δf, the phase of r_{L+i} relative to r_i is rotated by $\Delta\varphi(t_s) = 2\pi\Delta f L t_s$. Multiplying the conjugate of r_i (i.e., r_i^*) by r_{L+i}, we obtain:

$$s_i \stackrel{\text{def}}{=} r_i^* r_{L+i} = |r_i|^2 e^{-j2\pi\Delta f L t_s} = |r_i|^2 e^{-j\Delta\varphi(t_s)}. \quad (1)$$

Taking into account the channel coefficient $h_i = h_{L+i}$ and the noise terms, n_i and n_{L+i}, the value of s_i at the Rx, denoted by $\widetilde{s_i}$, is:

$$\widetilde{s_i} = |h_i r_i|^2 e^{-j2\pi\Delta f L t_s} + \bar{n_i} \quad (2)$$

where $\bar{n_i} \stackrel{\text{def}}{=} r_i n_{L+i}^* + r_{L+i}^* n_i + n_i n_{L+i}^*$ has zero mean. To average out the $\bar{n_i}$'s, the estimated phase offset, $\widetilde{\Delta\varphi}$, is measured over the summation of all the $\widetilde{s_i}$'s, i.e.,

$$\widetilde{\Delta\varphi(t_s)} = \measuredangle\Big(\sum_{i=0}^{L-1} \widetilde{s_i}\Big) \quad (3)$$

where the notation $\measuredangle(x)$ indicates the phase of a complex quantity x. Thus, the estimated FO is:

$$\widetilde{\Delta f} = \frac{\widetilde{\Delta\varphi(t_s)}}{2\pi L t_s}. \quad (4)$$

Fig. 4 shows an example of a sequence of length $2L = 8$ samples. The more samples are used to estimate $\widetilde{\Delta\varphi}$, the more accurate the estimated FO is.

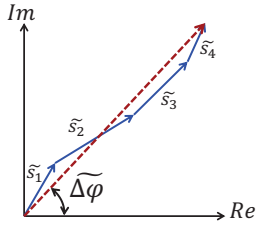

Fig. 4. Example of phase offset averaged over $L = 4$ $\widetilde{s_i}$ products.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

Regarding the phase of a complex number such as $\widetilde{s}_i$, the Rx observes a value between $-\pi$ and π. In other words, the Rx cannot distinguish $\Delta\varphi$ from $\Delta\varphi \pm 2k\pi$ in (4), for any integer k. The phase offset of 2π corresponds to $\frac{1}{Lt_s}$ offset, i.e., one subcarrier spacing. In particular, consider a subcarrier and two FOs from it, Δf_1 and Δf_2, where $|\Delta f_1| \leq \frac{1}{2Lt_s}$ and $|\Delta f_2| = |\Delta f_1| + \frac{1}{Lt_s}$. The corresponding phases are $2\pi|\Delta f_1|Lt_s$ and $2\pi|\Delta f_1|Lt_s + 2\pi$, respectively. Because the phases differ by 2π, there will be an ambiguity in distinguishing between them. The Rx interprets $\Delta f_1 + \frac{1}{Lt_s}$ as Δf_1 and will mistakenly adjust Δf_2 to the neighboring subcarrier bin. *In general, the phase is unambiguous and correctable as long as* $|\Delta f| < \frac{1}{2Lt_s}$ *(half a subcarrier spacing)*. This also implies that a longer period of a cycle reduces the range of FO that can be corrected unambiguously. Given a fixed sampling interval, a longer period results in higher L.

Let th_s and th_l be the maximum $|\Delta f|$ values that STSs and LTSs can correct unambiguously, respectively. In the 802.11a/g, two of the last three STSs are chosen to form a sequence with two identical halves for coarse FO estimation. Since the number of samples of an LTS is four times the number of samples of an STS, then $th_l = th_s/4 = f_\Delta/2$.

The above discussion reveals a tradeoff between the accuracy and range of the correctable FO. The goal of the STSs is to estimate a large FO value and compensate for it by multiplying the rest of the samples (including those obtained during the LTSs) by $e^{-j(-2\pi\widetilde{\Delta f_s}\, i\, t_s)}$, where $\widetilde{\Delta f_s}$ is the estimated FO in the STSs phase and i is the sample index. Using LTSs, the Rx then computes $\widetilde{\Delta f_l}$ to fine-tune the coarsely estimated FO. This explains one of the reasons for concatenating short and a long training fields in 802.11 systems. Consequently, if the actual FO is larger than th_s, this FO estimation method fails to fully compensate for it.

Even after the LTS-based FO correction, a small residual FO may remain due to noise. This error is typically too small to cause ICI, but it gradually rotates the phase of the received symbols on the constellation map and may increase the BER, specially in the long frames. A predetermined subset of subcarriers with known values (called *pilot subcarrier*) are used to track and compensate for these small phase changes. Theoretically, there is no frequency range limitation for FO estimation in pilot subcarriers [13]. In addition, known pilot subcarriers can be used for tracking channel variations.

2.2 Frame Detection

For a typical wireless Rx, an increase in the received power is a first indication of a new frame. To verify whether this increase is indeed due to a transmitted 802.11a/g/n/ac frame and then time synchronize with it, the Rx checks for the existence of successive identical sequences of a preset length [15]. In Schmidl and Cox's frame detection method, the Rx considers two non-overlapping intervals, each of duration $k\lambda_{STS}$ microseconds (equivalently, kL samples, where k is an integer) to represent two identical halves of a sequence. For example, three STSs with $t_s = 50\ ns$ sample period (owing to the Nyquist rate of 20 MHz) result in $L = 48$ samples. In the 802.11 standard, $1 \leq k \leq 5$. The correlation between the samples' conjugate in the first interval (window) and the corresponding samples in the second one is computed. Let $\mathcal{A}(n)$ be the summation of these correlations when the first window starts at the nth sample of the whole sequence:

$$\mathcal{A}(n) = \sum_{i=0}^{L-1} \widetilde{s}^*_{n+i} \widetilde{s}_{n+L+i}. \tag{5}$$

Using $\mathcal{A}(n)$, a normalized timing metric, $\mathcal{M}(n)$, is computed:

$$\mathcal{M}(n) = \frac{|\mathcal{A}(n)|^2}{\big(\mathcal{E}(n)\big)^2} \tag{6}$$

where $\mathcal{E}(n) \stackrel{\text{def}}{=} \sum_{i=0}^{L-1} |\widetilde{s}_{n+L+i}|^2$ is the received signal energy over the second window. $\mathcal{M}(n)$ is close to zero if either window does not contain any preamble sample. On the other hand, $\mathcal{M}(n)$ peaks when both windows contain only preamble samples. Ideally, $\mathcal{M}(n)$ should stay constant at the maximum value of 1, as long as both the windows are being moved inside the preamble boundaries. So the first time that $\mathcal{M}(n)$ hits the maximum is marked as the beginning of the frame. Because of noise, the maximum may occur later than the actual preamble start time. To account for this, the algorithm first finds $\hat{\mathcal{M}} = \max_n \mathcal{M}(n)$ and then searches for the earliest time before the occurrence of $\hat{\mathcal{M}}$ with an $\mathcal{M}$ value greater than $(1-\epsilon)\hat{\mathcal{M}}$, where $0 < \epsilon < 1$ is a system parameter. That time instant is taken as the beginning of the frame.

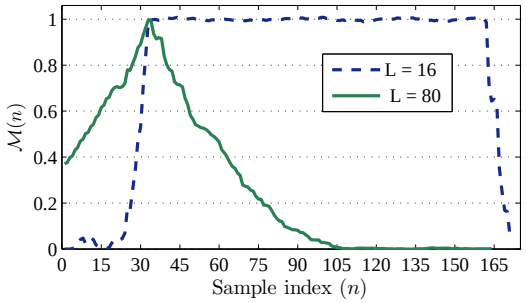

Fig. 5. $\mathcal{M}(n)$ vs. n for two extreme cases of window lengths (SNR = 42 dB, frame starts at $n = 31$, $t_s = 50\ ns$).

Fig. 5 shows two examples of the smallest and largest possible window sizes in the 802.11a frame detection scheme. When $L = 80$, the noise is averaged out, so the estimate $\hat{\mathcal{M}}$ is more reliable. In contrast, when $L = 16$, $\mathcal{M}(n)$ exhibits a higher fluctuation and $\hat{\mathcal{M}}$ is less reliable, requiring a larger ϵ to decrease the probability of misdetecting the frame start time. Even though the sharp increase of $\mathcal{M}(n)$ makes room to increase ϵ, it is unclear how much increase is sufficient.

3 Model and Assumptions

We consider a link between Alice (the Tx) and Bob (the Rx). The adversary (Eve) is in the transmission ranges of both Alice and Bob. Alice transmits an 802.11 OFDM frame and Bob uses a few of the first STSs for frame detection. He chooses two of the last three STSs, in conformity with the standard (see Fig. 3) and employs the Schmidl and Cox method for FO estimation. Once Bob estimates the coarse FO using STSs and compensates for $\widetilde{\Delta f_s}$, he assumes,

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

by default, that the residual FO is less than th_l and then estimates it using LTSs. According to the 802.11 standard, Bob does not perform any kind of boundary check during the LTS- and pilot-based FO estimation processes.

Eve aims at irrecoverably corrupting Alice's frame at Bob using the lowest possible jamming effort. Eve is aware of the PHY-layer protocol and the FO correction mechanism at Bob. She makes no assumptions about the channel parameters or Alice's transmission power. If the oscillators are either stable or accurate, Eve initially eavesdrops on Alice's and Bob's preamble transmissions (e.g., data-ACK exchanges) for a while. Through averaging, she estimates their FOs relative to Eve, denoted by Δf_{ae} and Δf_{be}, respectively[2].

The metrics of interest are coarse and final estimated FOs at Bob, Symbol error rate (SER), the BER after demodulation but before decoding, and the jamming effort (defined as the jammer's duty cycle [3]). These metrics will be studied with respect to the SNR, modulation scheme, and signal-to-jamming ratio (SJR) at Bob.

4 Frequency Offset Estimation Attack

In this section, we describe in detail an attack on the FO estimation. Eve launches this attack in two phases: (1) Eavesdropping on the channel to detect the start of Alice's frame transmission and acquire its timing information; and (2) jamming the last three STSs of the preamble, which are used for coarse FO estimation.

4.1 Phase 1: Adaptive Fast Frame Detection

To pinpoint the last three STSs in time and corrupt the FO estimation at Bob, Eve must detect Alice's frame and synchronize with its arrival at Bob. The detection should be fast enough to allow sufficient time for processing, switching to transmission mode, and jamming the last three STSs. Referring to the frame detection mechanism in Section 2, Eve chooses the minimum possible window size (one STS, $L = 16$) and reduces the capture time to $2.5\lambda_{STS} = 2\ \mu s$ to make sure that at least the first two STSs are captured.

To account for the higher detection inaccuracy due to the small window size, Eve assumes that the actual start time belongs to the first $V = \log_2(L)$[3] sample indices $i_0, i_1, \ldots, i_{V-1}$ that are greater than $(1-\epsilon)\hat{\mathcal{M}}$ and finds all of them, instead of just looking for the first one. She sets ϵ to a value less than $1/L$, the contribution of a preamble sample pair in $\mathcal{M}(n)$. This is an attempt to exclude the samples located more than one index before the actual frame start time. If there are less than V sample values greater than the threshold, Eve adaptively decreases the threshold by finding the smallest ϵ that guarantees the existence of V candidates[4].

2. In general, oscillators exhibit numerous instabilities, due to aging, temperature, acceleration, ionizing radiation, power supply voltage, etc. Thus, the Rx must update the FO estimate on a per-frame basis, even if the frame sender is already known. This is specially the case with non-stable oscillators. In this case, Eve can perform FO estimation along with fast frame detection to optimally design the jamming signal for each frame (see Section 4.1).

3. The reason of this specific number will be explained in Section 4.2.3.

4. Eve may also apply the synchronization method in [18] to improve the detection accuracy.

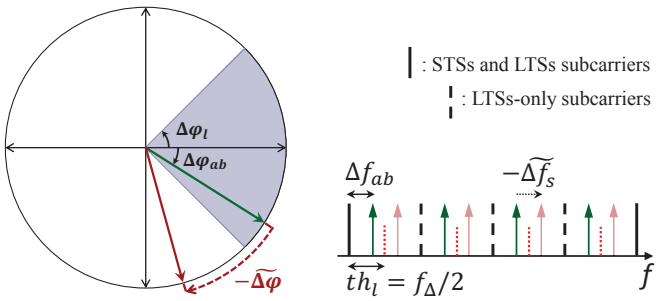

(a) Phase domain: The shaded area represents the LTS-based correctable range. A wrong phase estimation $\widetilde{\Delta\varphi}$ can move $\Delta\varphi_{ab}$ out of the correctable range.

(b) Frequency domain: Incorrect estimation of $\widetilde{\Delta f_s}$ can move Δf_{ab} out of the LTS-based correctable range.

Fig. 6. Phase and frequency offsets as observed during the STSs.

4.2 Phase 2: Preamble Jamming

Based on i_0, Eve computes the arrival time of the last three STSs of the preamble and generates a jamming signal that would be aligned with those STSs. The energy-efficient jamming sequence is designed to defeat all STS-, LTS-, and pilot-based FO corrections without jamming the LTSs and pilot subcarriers. For this attack to be successful, Eve has to account for unknown channel parameters and frame-detection timing errors. More specifically, the jamming sequence is designed to achieve the following goals:

4.2.1 Forcing Bob to make a destructive error

By default, Bob assumes that the FO to be estimated using LTSs is less than th_l. If Eve deceives Bob into erroneously push the FO beyond th_l after receiving the STSs instead of reducing it, then she achieves her goal without needing to jam the LTSs.

Without loss of generality, Eve assumes i_0 is the correct start time of the frame (we will relax this assumption later). Let $\Delta f_{eb} = -\Delta f_{be}$ and $\Delta f_{ab} = \Delta f_{ae} - \Delta f_{be}$ represent Bob's estimates of Eve-to-Bob and Alice-to-Bob FOs, respectively. Let $\Delta\varphi_{ab}$, $\Delta\varphi_{eb}$, and $\Delta\varphi_l = \pi/4$ be the phase offsets corresponding to Δf_{ab}, Δf_{eb}, and th_l, respectively, after a single STS ($0.8\ \mu s$). To cause incorrect FO estimation ($\widetilde{\Delta f_s}$) such that the updated FO after STSs ($\Delta f_{ab} - \widetilde{\Delta f_s}$) is higher than th_l, the following inequality should hold:

$$|\Delta\varphi_{ab} - \widetilde{\Delta\varphi}| > \Delta\varphi_l. \tag{7}$$

Fig. 6(a) and 6(b) show an example of such a situation in the polar coordinates and frequency domain, respectively.

Eve's jamming signal needs to satisfy (7). Let g be the Eve-to-Bob channel coefficient. We assume that during Eve's jamming period, g is the same for all the jamming samples that belong to the jamming sequence $\mathbf{u}$, denoted by $u_i, i = 1, \ldots, 2L$. Let $\tilde{r}_i = hr_i$ and $\tilde{u}_i = gu_i$. We consider two different approaches for generating the jamming sequence:

1) Random noise: A simple way to corrupt the FO estimation at Bob is to jam the last three STSs with a random signal. Recalculating the autocorrelation $\mathcal{A}$ at Bob after the

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

superposition and ignoring the noise term in (2), we have:

$$\begin{aligned}\mathcal{A}_{\text{random}} &\stackrel{\text{def}}{=} \sum_{i=0}^{L-1} \widetilde{s_i} = \sum_{i=0}^{L-1} (\tilde{r}_i + \tilde{u}_i)^* (\tilde{r}_i e^{-j\Delta\varphi_{ab}} + \tilde{u}_{L+i}) \\ &= \sum_{i=0}^{L-1} |\tilde{r}_i|^2 e^{-j\Delta\varphi_{ab}} + \sum_{i=0}^{L-1} \tilde{r}_i{}^* \tilde{u}_{L+i} \\ &\quad + \sum_{i=0}^{L-1} \tilde{u}_i^* (\tilde{r}_i e^{-j\Delta\varphi_{ab}} + \tilde{u}_{L+i}).\end{aligned} \tag{8}$$

The phase and amplitude of the 2nd and 3rd terms in (8) (and hence $\widetilde{\Delta\varphi}_{\text{random}} \stackrel{\text{def}}{=} \measuredangle\mathcal{A}_{\text{random}}$) are unknown because not only they include random complex numbers $\tilde{u_i}$, but also the phase and amplitude of $\tilde{r}_i$ are unknown after traversing the Alice-to-Bob channel. Hence, $\widetilde{\Delta\varphi}_{\text{random}}$ may not satisfy (7), so FO jamming with a random signal cannot provide any FO distortion guarantees to beat LTS-based FO estimation.

2) Fake preamble: A more effective jamming approach that exploits both knowledge of the FO estimation algorithm and Δf_{ab} is to construct a fake preamble with "identical halves". For now, assume that the samples of the jamming signal $u_i, i = 1, \ldots, 2L$ can take any arbitrary value as long as the signal conforms to the protocol bandwidth requirement. The preamble phase warping attack in [10] is a special case of this approach, where the jamming signal is a random frequency-shifted version of an arbitrary fake preamble. The advantage of having identical halves is that we can control and carefully calculate a desired FO for **u** based on how Bob estimates Δf_{ab}. Here, we also note that the channel response between Eve and Bob does not change the FO. Before we explain how a desired FO (and hence Δf_{eb}) is determined, consider the superposition of Alice's signal and Eve's jamming at Bob. Dropping the index i from (2) and ignoring the noise term, we have:

$$\begin{aligned}&\tilde{s} = (\tilde{r} + \tilde{u})^* (\tilde{r} e^{-j\Delta\varphi_{ab}} + \tilde{u} e^{-j\Delta\varphi_{eb}}) = e^{-j\Delta\varphi_{ab}} \times \\ &\Big[\underbrace{|\tilde{r}|^2 + |\tilde{u}|^2 e^{-j(\Delta\varphi_{eb} - \Delta\varphi_{ab})} + \tilde{r}^* \tilde{u} e^{-j(\Delta\varphi_{eb} - \Delta\varphi_{ab})} + \tilde{u}^* \tilde{r}}_{\mathcal{B}}\Big].\end{aligned} \tag{9}$$

Thus, the estimated phase offset at Bob is:

$$\widetilde{\Delta\varphi} = \measuredangle\tilde{s} = \Delta\varphi_{ab} + \measuredangle\mathcal{B} + \measuredangle\bar{n}. \tag{10}$$

Note that the phase estimation error $\varphi_e \stackrel{\text{def}}{=} \measuredangle\mathcal{B}$ is a function of SJR and $\Delta\varphi_{eb}$, and jamming will have no effect if $\varphi_e = 0$.

Upon calculating $\widetilde{\Delta\varphi}$ and $\widetilde{\Delta f_s}$, Bob changes the FO for the rest of the frame to $\Delta f_{ab} - \widetilde{\Delta f_s}$. According to (7), Eve is successful if she can ensure that $\Delta\varphi_{eb}$ satisfies the following:

$$|\Delta\varphi_{ab} - \widetilde{\Delta\varphi}| > \Delta\varphi_l \Rightarrow |\varphi_e + \measuredangle\bar{n}| > \Delta\varphi_l = \frac{\pi}{4}. \tag{11}$$

Eve can guarantee a desired φ_e only if SJR$\rightarrow -\infty$. Otherwise, even if she knows $\Delta\varphi_{ab}$ and $\tilde{u}$ and can also control $\Delta\varphi_{eb}$, she has no control over other channel-dependent parameters in $\mathcal{B}$. Specifically, the phase and amplitude of $\tilde{r}$ are channel-dependent and Eve cannot estimate the Alice-to-Bob channel coefficient h. That means that Eve is still unable to guarantee a successful attack, which is also the case in the preamble phase warping attack.

4.2.2 Designing a channel-independent jamming signal

To address the aforementioned challenge, Eve takes advantage of Alice's known preamble samples and the product sum in (3) to cancel out the terms with unknown phases. Eve first chooses $L/2$ non-overlapping pairs of samples. Without loss of generality, let Eve pair the samples in order and let (u_1, u_2) be the first pair of samples in the jamming sequence. By knowing the preamble sample values at Alice, u_2 can be designed such that when Bob sums up $\tilde{s}_1$ and $\tilde{s}_2$, all the terms that depend on $\tilde{r}$ (excluding $|\tilde{r}|$) in the term $\mathcal{B}$ in (9) are eliminated. Thus,

$$u_2 = -\frac{r_1^*}{r_2^*} u_1 \tag{12}$$

which implies that

$$\begin{aligned}&\widetilde{s_1} + \widetilde{s_2} = e^{-j\Delta\varphi_{ab}} \times \\ &\Big[|\tilde{r}_1|^2 + |\tilde{r}_2|^2 + (|\tilde{u}_1|^2 + |\tilde{u}_2|^2|)\, e^{-j(\Delta\varphi_{eb} - \Delta\varphi_{ab})}\Big].\end{aligned} \tag{13}$$

The requirement in (12) is similarly imposed on the rest of the even samples. We refer to this requirement as the *pairing rule*. Accordingly, the autocorrelation function $\mathcal{A}$ for this scheme, denoted by $\mathcal{A}_{\text{fake}}$, becomes:

$$\begin{aligned}\mathcal{A}_{\text{fake}} &= \sum_{i=0}^{L-1} \widetilde{s_i} = \\ &e^{-j\Delta\varphi_{ab}} \underbrace{\Big[\sum_{i=0}^{L-1} |\tilde{r}_i|^2 + \sum_{i=0}^{L-1} |\tilde{u}_i|^2\, e^{-j(\Delta\varphi_{eb} - \Delta\varphi_{ab})}\Big]}_{\mathcal{C}(\Delta\varphi_{eb} - \Delta\varphi_{ab})}.\end{aligned} \tag{14}$$

Now $\mathcal{A}_{\text{fake}}$ is a function of the difference between $\Delta\varphi_{ab}$ and $\Delta\varphi_{eb}$ only. So Eve can determine a desired value of $\Delta\varphi_{eb}$ in a way that makes $|\measuredangle\mathcal{C}(\Delta\varphi_{eb} - \Delta\varphi_{ab})| > \Delta\varphi_l$, which satisfies (11).

4.2.3 Robustness to errors in frame start time

We now relax the assumption that Eve can precisely determine the true frame start time and consider a scenario in which she compiles a short list of possible frame start times besides i_0, as explained in Section 4.1. Thus far, we have required the jamming sequence to have identical halves with a $\Delta\varphi_{eb}$ that satisfies (11) and the even samples to be a function of odd samples (pairing rule). Eve could still benefit from the remaining free, unassigned samples (i.e., odd samples) to cancel out channel-dependent terms for other possible start times. We generalize the pairing technique to larger sets of samples and define the following *chaining rule* to account for $V - 1$ other start times $i_1, i_2, \ldots, i_{V-1}$.[5]

Let $\mathbf{m} = \{m_1, \ldots, m_{V-1}\}$ where $m_j = i_j - i_0$. First, Eve extends her jamming sequence by appending (cyclically postfixing) the first m_{V-1} jamming samples to this sequence. So for any candidate frame start time i_j, the jamming signal will be fully superposed on Alice's three STSs because the jamming signal is cyclically extended already by $m_{V-1} > m_j$ samples. Next, Eve assumes that i_1 is the correct frame start. In this case, the superposition

5. Eve can precompute and then account for the propagation delays by timestamping the data-ACK exchanges between Alice and Bob and estimating the Eve-to-Bob distance. The chaining rule can also be leveraged to account for errors in estimating these delays.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

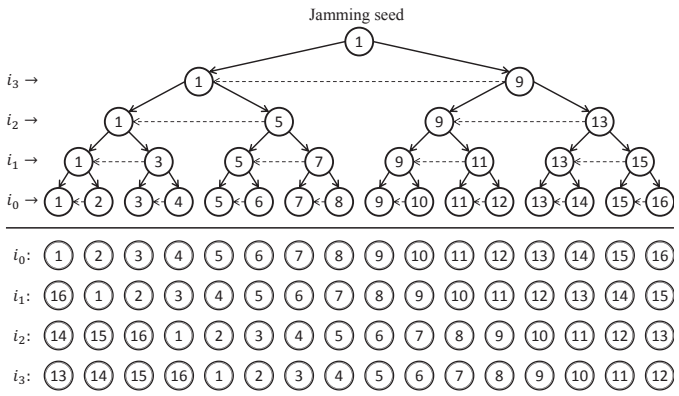

Fig. 7. Cascaded chaining and pairing of the samples towards the jamming seed. Jamming samples are shown on the tree and the shifted versions of Alice's preamble on the bottom. Horizontal dashed lines represent direct dependency between samples.

of the jamming signal on Alice's three STSs will be different from the previous case (i.e., the jamming sequence is slid with respect to Alice's STSs) and (12) is no longer sufficient to eliminate the last two channel-dependent terms within $\mathcal{B}$ in (9). Instead, Eve can find pairs of yet free samples and, similar to the pairing rule, define one of the samples of each of such pairs based on the other sample of that pair and also the corresponding samples in $\mathbf{r}$. After this step, half of the free samples will be given values. Eve repeats the same procedure for the rest of the frame start times and free samples. Based on these hierarchical dependencies among the samples u_i, Eve constructs a binary *chaining tree* in which the dependency between two samples is mapped to a parent-child relationship. Note that an unassigned (free) sample may already have a chain of other dependent sample(s). The value of the dependents will be updated whenever that sample takes a new value.

An example is depicted in Fig. 7 with $\mathbf{m} = \{0, 1, 3, 4\}$. Without loss of generality, we assume Alice's preamble sequence is shifted instead of the jamming sequence. The tree in this figure shows how the jamming samples are being chained together and used to construct the tree from the bottom to the top. A pair of free samples are considered as siblings. The left child specifies the value of its right sibling based on m_j and then the left child is copied to its parent node. So the right child depends on its sibling. To explicitly define the dependency between the two sibling samples, all their dependent samples must also be taken into account because their values in (9) are affected by their parents' values. For example, when $j = 1$, Eve may select two free samples u_1 and u_3 (together with their dependents u_2 and u_4) to eliminate the channel-dependent terms:

$$u_1^* r_{16} + u_2^* r_1 + u_3^* r_2 + u_4^* r_3 = 0 \tag{15}$$

which implies the dependency of u_3 to u_1 (u_2 and u_4 are substituted by their corresponding pairing rule dependencies on u_1 and u_3):

$$u_3^* = -\frac{r_4(r_2 r_{16} - r_1 r_1)}{r_2(r_2 r_4 - r_3 r_3)}\, u_1^*. \tag{16}$$

Now the value of the dependent of u_3 (u_4 in this example) is updated to maintain its dependency relationship with the right sibling u_3.

Algorithm 1 Chaining and pairing rules combined

1: **Input**: $L, V, \mathbf{r}[1 \dots L], \mathbf{m}[0 \dots V-1]$
2: **Initialize: u** $= \mathbf{0}$
3: **for** $j \leftarrow 1, V-1$ **do**
4: $\quad k \leftarrow 2^j$
5: $\quad$ **while** $k < L$ **do**
6: $\quad\quad \mathbf{t} =$ circularly shifted $\mathbf{r}$ by m_j
7: $\quad\quad x = -\sum_{i=k-2^j+1}^{k} u_i t_i^* \,/\, \sum_{i=k+1}^{k+2^j} u_i t_i^*$
8: $\quad\quad [u_k, \dots, u_{k+2^j-1}] = [u_k, \dots, u_{k+2^j-1}] * x$
9: $\quad\quad k = k + 2^{j+1}$
10: $\quad$ **end while**
11: **end for**
12: **Return u**

A pseudocode of the chaining rule, which also contains the pairing rule, is provided in Algorithm 1. The algorithm iterates for each m_j, $j = 1, \dots, V-1$. At each iteration and for each pair of free samples, the right subtree (the right siblings of all its $2^j - 1$ dependents) is multiplied by a coefficient x (defined in line 8) such that the summation of the corresponding 2^j product terms in (9) and the 2^j terms corresponding to the left subtree is zero. The horizontal arrows in Fig. 7 show the dependence of the right subtrees on their left subtrees. As a result, $L/2^j$ samples are assigned at each iteration and the algorithm terminates after $V = \log_2(L)$ iterations. In the end, all but one of the samples (u_1 in our example) will be a right sibling at least once at some point in the tree and so are assigned. We call the remaining free sample the *jamming seed*, to which all the samples are chained either directly or recursively. The jamming seed can be used to control the jamming power.

4.3 Effects of LTSs on FO and Channel Estimation

LTSs are used for fine FO and channel estimation. As explained in Section 2, the phase offset from the LTS-based FO corresction perspective is between $-\pi$ and π, which means that the true FO after STS-based correction has to be between $-th_l$ and th_l. So LTSs can correct up to $th_l = f_\Delta/2$ FO, and any remaining phase offset will be an integer multiple of 2π, which corresponds to $2k\,th_l = kf_\Delta$, $k = 1, 2, \dots$. In other words, the LTSs at Bob round up the FO manipulated by $\widetilde{\Delta f_s}$ to the nearest multiple of $2th_l$ and avoid ICI by adjusting the subcarriers to the closest, though incorrect, frequency bins. Consequently, in this attack all the subcarriers will be shifted forward or backward, replacing neighboring subcarriers. Bob eventually demodulates the bits of all OFDM symbols, but he is unaware that these symbols have been shifted and misplaced. A simple example with four subcarriers is provided in Fig. 8. Each subcarrier carries two bits (QPSK-modulated symbols). In the shifted version, two unknown bits are added in the beginning and the rest of the sequence is shifted to the right, although the bits are correctly demodulated. Therefore, when the bits of different OFDM symbols are concatenated to reconstruct the original bit sequence, the entire sequence will look shuffled and out-of-order compared to the original bit sequence. A shifted version of an arbitrary bit sequence will result in very high BER.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

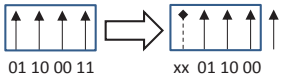

Fig. 8. Example of the FO attack on four subcarriers (left): The attack shifts the subcarriers and the corresponding bits to the right.

An STS-based FO estimation error also affects the channel estimation process, which is applied across the LTSs, specially if Bob estimates the channel irrespective to the outcome of the fine FO estimation. To elaborate, the phase offset accumulates over time, causing different LTS samples to have different phase offsets. However, Bob complacently tries to interpret this time-varying phase offset as a fixed-value channel phasor by minimizing the MSE. Hence, his attempt to model the FO as if it is a channel parameter results in an incorrect estimated channel phasor, which after equalization rotates the payload's modulation symbols on the constellation map.

4.4 Optimal Jamming Strategy (Optimal Jamming Sequence Design)

Let $\Phi_{eb} \stackrel{\text{def}}{=} \Delta\varphi_{eb} - \Delta\varphi_{ab}$. If the SJR at Bob is known, Eve can achieve the maximum possible $|\angle\mathcal{C}(\Phi_{eb})|$ value by optimally selecting $|\Phi_{eb}|$. This maximization allows Eve to inflict the maximum subcarrier shift and overcome possible FO estimation inaccuracies due to noise at Eve or Bob. To calculate the optimal $|\Phi_{eb}|$, we represent the total received jamming energy $|\tilde{u}|^2$ and signal energy $|\tilde{r}|^2$ in polar coordinates, as shown in Fig. 9. Using geometric arguments, we find the maximum $|\angle\mathcal{C}|$, where $\mathcal{C} = |\tilde{r}|^2 + |\tilde{u}|^2 e^{-j(\Phi_{eb})}$. Each circular contour in this figure shows the end points of the vector $\mathcal{C}$ for a given SJR but different Φ_{eb} values.

As long as $|\tilde{u}| < |\tilde{r}|$, $|\angle\mathcal{C}|$ reaches its maximum when the vector $\mathcal{C}$ is tangent to the contour circle. In a right triangle, this implies

$$|\angle\mathcal{C}| = \arcsin\frac{|\tilde{u}|^2}{|\tilde{r}|^2} \tag{17}$$

and

$$|\Phi_{eb}| = \pi/2 + \angle\mathcal{C}. \tag{18}$$

When $|\tilde{u}| \geq |\tilde{r}|$, the maximum $\angle\mathcal{C}$ equals to π, which is always achieved when $|\Phi_{eb}| = \pi$. In Fig. 10, we plot the corresponding optimal $|\Delta f_{eb} - \Delta f_{ab}|$ during the STSs for different SJR values. Based on $\angle\mathcal{C}$, we also derive the resulting number of subcarrier-spacings shift after LTSs. Note that phase offsets $\pi/2$ and π correspond to FOs of one and two f_Δ's, respectively. From the STSs perspective, LTSs adjust a phase offset to its closest multiple of $2\varphi_l$. So when $|\angle\mathcal{C}| > 3\varphi_l$, the attack results in a shift of two subcarriers.

The jamming sequence can be designed to minimize the total jamming energy $\sum_{i=0}^{L-1}|\tilde{u}_i|^2$, subject to the constraint of at least one subcarrier shift, i.e., $|\angle\mathcal{C}(\Phi_{eb})| \geq \Delta\varphi_l$. The shaded area in Fig. 9 shows the feasible region. According to (17) and the geometry in Fig. 9, we conclude that:

1) The energy minimization problem is feasible as long as

$$\text{SJR} = \frac{\sum_{i=0}^{L-1}|\tilde{r}_i|^2}{\sum_{i=0}^{L-1}|\tilde{u}_i|^2} \leq \frac{1}{\sin(\Delta\varphi_l)} = \sqrt{2} \approx 1.5 \text{ dB}. \tag{19}$$

2) The minimum jamming energy is achieved when

$$|\Delta\varphi_{eb} - \Delta\varphi_{ab}| = |\pi/2 + \Delta\varphi_l| = 3\pi/4, \tag{20}$$

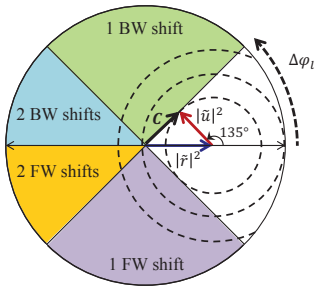

Fig. 9. Superposition of Alice's and Eve's signals at Bob and the resulting subcarrier shifts. The minimum feasible $|\tilde{u}|^2$ occur when the vector $|\tilde{u}|^2$ is perpendicular to an edge of the 1-shift regions. The parts of a contour crossing the shaded areas show the feasible phases for a given $|\tilde{u}|^2$.

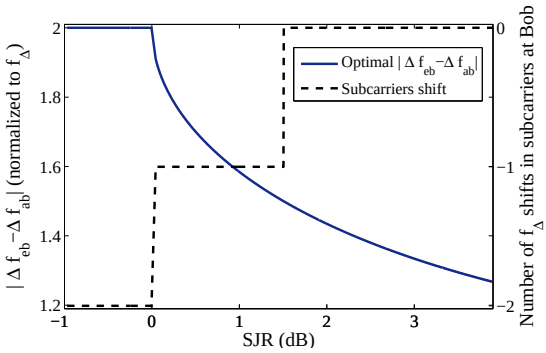

Fig. 10. Optimal $|\Delta f_{eb} - \Delta f_{ab}|$ and resulting amount of subcarrier shift for different SJR values.

or equivalently, $|\Delta f_{eb} - \Delta f_{ab}| = 1.5\, f_\Delta$.

Equation (20) says that the phase offset of Eve's signal as perceived by Bob should have phase difference of $|\pi/2 + \Delta\varphi_l|$ relative to Alice's signal. Even if $\Delta\varphi_{eb}$ does not satisfy (20), Eve can augment the hardware-dependent Δf_{eb} and obtain an *effective* Δf_{eb} by imposing an artificial FO of Δf_n on the jamming sequence before it is transmitted by the oscillator. This is achieved by multiplying the samples of the jamming sequence by $e^{-j2\pi\Delta f_n\, i\, t_s}$, where Δf_n is given by:

$$\Delta f_n = \pm 1.5\, f_\Delta - \Delta f_{eb} + \Delta f_{ab}. \tag{21}$$

The optimal $|\Phi_{eb}|$ that minimizes the jamming energy is particularly important in designing the optimal jamming strategy because the SJR at Bob is usually unknown to Eve. The optimal jamming strategy to deal with this situation is to consider the worst-case (highest) SJR under which the attack is successful and then set the effective FO according to (20). Therefore, Eve always sets Φ_{eb} to $\pm(\pi/2 + \Delta\varphi_l)$.

5 Discussion

OFDM-based 802.11 systems employ interleaving and adaptive modulation and coding (AMC) schemes to increase resiliency against jamming and bit errors. However, the achieved BER value of the aforementioned FO attack (~ 0.5) is high enough that the mutual information between the transmitted and received sequences is zero, and hence practical coding schemes cannot recover the frame. After an unsuccessful transmission and subsequent data rate reduction, Alice may increase her transmit power for the whole frame. In the case of the proposed FO attack, such an increase is unnecessary and inefficient for the payload, which constitutes up to 99.9% of a frame. In addition, an

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

intelligent jammer can track Alice's power increase (e.g., by overhearing management frames), adjust the jamming power to always achieve the optimal SJR, and force the dropping of subsequent transmissions.

It may also be argued that because pilot subcarriers are transmitted on known frequencies, Bob can compare the known symbols of the pilot subcarriers with the received symbols on different subcarriers to identify a possible subcarrier shift. However, because channel estimation is distorted, locating the corrupted pilot subcarriers at Bob is quite challenging. Furthermore, these pilot subcarriers cannot be easily used for channel estimation (we leave the investigation of this problem to a future work).

Moreover, we note that jamming the LTSs after jamming the STSs strengthens the attack by further distorting the channel estimation process. However, jamming the LTSs alone cannot lead to a subcarrier shift even though it involves more jamming effort (8-μs duration on 48 subcarriers) than jamming three STSs ($\leq$ 3 μs on 12 subcarriers). Furthermore, with LTSs jamming, pilot subcarriers can still be used to estimate the channel and correct any residual FO.

The system model in this paper assumes a single Tx-Rx-pair (i.e., Alice and Bob, and hence their FO, are known). In the case of multiple Tx-Rx pairs, Eve can construct a database of the FOs between different Tx-Rx pairs. Benefiting from CSMA/CA channel access mechanism, Eve can consider one transmission at a time and then leverage protocol semantics (e.g., data-ACK exchanges) to guess the Tx and Rx of an upcoming transmission. Further investigation of this issue is left for future work.

6 Performance Evaluation

In this section, we evaluate the effectiveness of the FO estimation attack through simulations and USRP experiments. We implemented the 802.11a/g preamble (including both short and long training sequences) by extending the PHY-layer library functions of LabVIEW. Alice appends 1500 modulated random bits to the frame preamble. Pilot-based channel and FO estimation and channel coding were not implemented to concentrate on the specific effects of the FO attack on received uncoded bits. The impact of coding and pilot subcarriers was discussed in Section 5.

We assume that Bob uses the STSs t_9 and t_{10}, as defined in Fig. 3, for coarse FO estimation, followed by fine FO estimation using LTSs. Channel estimation is performed over the first LTS using the time domain method [13]. We first evaluate the performance under a simulated AWGN channel model and later in a multi-path indoor environment. (More results are provided in [19].) We vary the SJR, the SNR (noise level), the modulation scheme, and Eve's effective FO, denoted by D_{eb}. In particular, we consider BPSK, QPSK, and 16-QAM modulation schemes. We measure $\widetilde{\Delta f_s}$ as well as final estimated FO, SER, and BER.

We compare three cases: 1) jamming the last STSs with a random signal (see Section 4.2.1), 2) FO attack with pairing rule only ($V = 1$ and $L = 16$ for frame detection), and 3) the entire FO attack including the chaining rule, with $L = 16$ and $V = \log_2 L$. The purpose of evaluating the second case is to study the impact of the chaining rule. The jamming duration for the second case is always equal to the sum of the durations of t_8 and t_9. However, it is not constant when the chaining rule is applied, and depends on m_{V-1}.

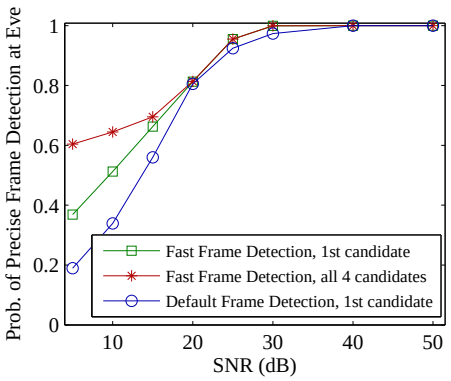

Fig. 11. Performance of different variants of frame detection vs. SNR (simulations).

6.1 Simulations

We consider an AWGN channel model without signal attenuation. In our simulations, the SJR is normalized to the energy of two full STSs. However, the chaining rule results in a variable-length cyclic postfix extension, which sometimes has a slightly higher sample power than the average sample power over an STS.

6.1.1 Frame Detection and Jamming Duration

Initially, we assess the accuracy of our adaptive fast frame detection method at Eve and also its impact on the jamming duration. Even though our adaptive detection method uses a small window of $L = 16$ compared to $L = 48$ for the default scheme, adapting ϵ based on finding V frame-start candidates increases the probability of precise frame detection even for the first candidate. This is shown in Fig. 11, where each probability is calculated based on more than 25000 runs. By including additional $V - 1$ candidate start times, we further increase the probability of including the true start time in V candidates, specially under high noise levels. The chaining rule benefits from V start times because it equally likely considers all the candidate start times to construct the jamming signal.

The jamming duration depends on m_{V-1} and the amount of postfix extension. In Table 1, we report the average index-distance between the first and the last samples (m_{V-1}) in the set of V start times when an STS contains $L = 16$ samples. The table shows that even at low SNR, the amount of cyclic extension due to the chaining rule is often less than half an STS. In particular, in 99.88% of the cases, $m_{V-1} \leq 8$, which means the jamming duration will be less than $3.5\lambda_{STS}$ or, equivalently, 0.7 of an OFDM-symbol duration. A 1500-bit BPSK-modulated payload lasts for 32 OFDM symbols, equivalent to $160\,\lambda_{STS}$. The durations of 16-QAM-modulated and QPSK-modulated signals of the same payload will be 40 and $80\,\lambda_{STS}$, respectively. So the jamming effort in our simulations is upper bounded by 2.0%, 3.5%, and 5.9% for BPSK, QPSK, and 16-QAM-modulated payloads, respectively. In general, an 802.11a frame lasts for $20 \times 10^{-6} + \lceil (22 + \text{LENGTH})/\text{DATARATE} \rceil$ seconds [2], where LENGTH and DATARATE denote the

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

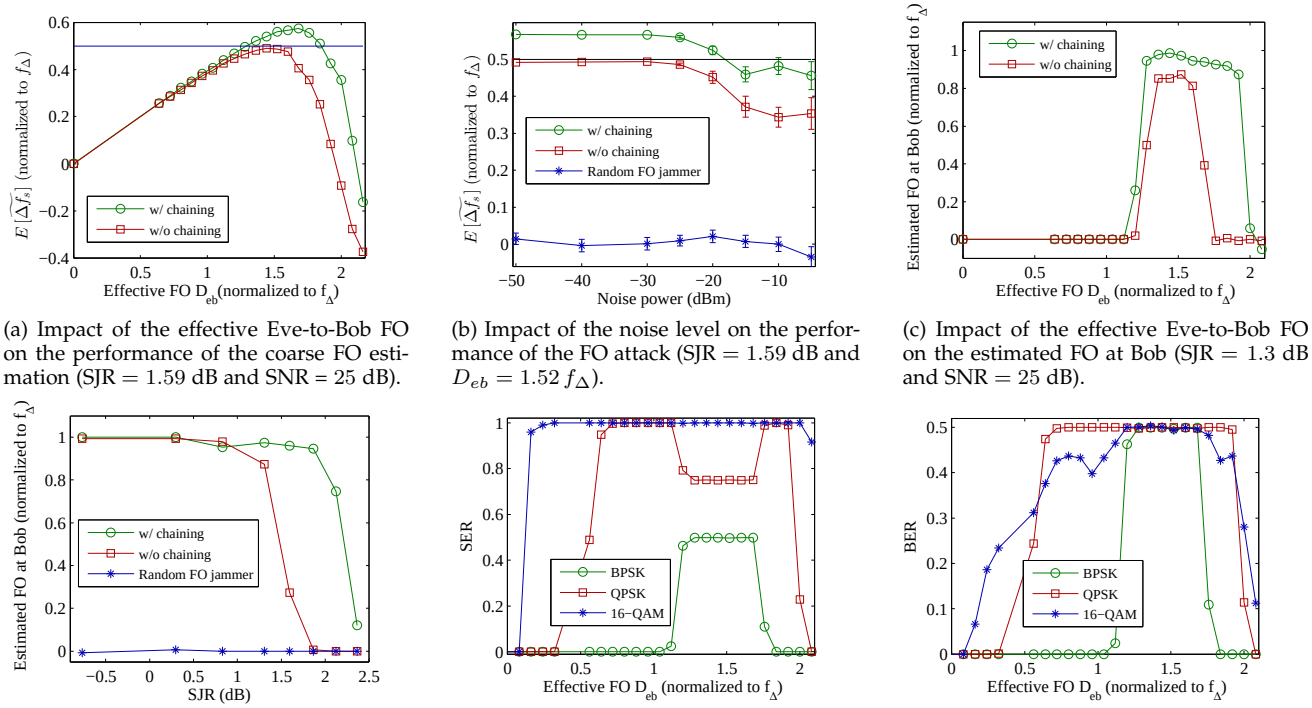

(a) Impact of the effective Eve-to-Bob FO on the performance of the coarse FO estimation (SJR $= 1.59$ dB and SNR = 25 dB).

(b) Impact of the noise level on the performance of the FO attack (SJR $= 1.59$ dB and $D_{eb} = 1.52\, f_\Delta$).

(c) Impact of the effective Eve-to-Bob FO on the estimated FO at Bob (SJR $= 1.3$ dB and SNR $= 25$ dB).

(d) Impact of SJR on the estimated FO at Bob (SNR $= 25$ dB and $D_{eb} = 1.52\, f_\Delta$).

(e) SER performance for different modulation schemes (SJR $= 1.59$ dB and SNR $= 30$ dB).

(f) BER performance for different modulation schemes (SJR$= 1.59$ dB and SNR = 30 dB).

Fig. 12. Performance of different variants of the FO attack and of random FO jamming under different noise levels, D_{eb} and SJR values, and modulation schemes. The transmission power is 0 dBm. (simulation results)

encoded payload size (in bits) and the data rate, respectively. For a typical 802.11a frame [3], the jamming effort varies between 0.07% and 0.88%, depending on the code and data rates. This is 30% less than the effort of the OFDM symbol jamming attack in [3].

SNR (dB)	5	10	15	20	25	30
$E(m_{V-1})$	4.05	3.74	3.54	3.23	3.02	3.0

TABLE 1
Average value of m_{V-1} in the chaining rule for different SNR levels.

6.1.2 FO Estimation

Fig. 12(a) depicts the average $\widetilde{\Delta f_s}$, measured after the corrupted STSs of 150 frames, when SJR$= 1.59$ dB, transmission power is 0 dBm, and noise level is -25dBm. The horizontal line represents th_l, normalized to f_Δ. The chaining rule improves the jamming effectiveness and guarantees a range of effective FO values for which the attack is successful ($\widetilde{\Delta\varphi} > \Delta\varphi_l$). When the chaining rule is not applied, the jamming attack is optimal at the optimal effective FO derived in Section 4.4, but is still insufficient to pass the threshold because of frame detection errors. When the chaining rule with V candidates is applied, the maximum average $\widetilde{\Delta\varphi}$ occurs later than the maximum for the no-chaining case because of slightly higher power during postfix samples. Fig. 12(b) shows the effect of noise on the STS-based estimated FO when SJR$= 1.59$ during the last three STSs and with $D_{eb} = 1.52\, f_\Delta$, a near-optimal value for this setup. The 90% confidence intervals are shown for each point. The increase in frame timing errors due to noise reduces the effectiveness of the attack, but this increase has less impact when the chaining rule is applied. When the noise level is higher than -20 dBm, the gap between the curves belonging to the two modes of the FO attack is wider, showing that the chaining rule is more robust in highly noisy channels .

When $|\widetilde{\Delta\varphi}| > \Delta\varphi_l$, the LTSs round the estimated FO to the nearest multiple of $2th_l$. Otherwise, LTSs try to round the FO to zero. In Fig. 12(c), we plot the average final estimated FO at Bob when SJR$= 1.3$ dB during the last three STSs and the noise level is -25 dBm throughout the frame. The chaining rule improves Eve's ability to shift the subcarriers by one f_Δ. With respect to the SJR, we can observe in Fig. 12(d) that when Eve'e D_{eb} is close to its optimal value, Eve is not able to guarantee a successful attack without the chaining rule even with the optimal SJR value of 1.5 dB.

6.1.3 Impact of the FO Attack on Modulation Performance

Under a relatively high SNR (e.g., 30 dB in our simulations) and without the FO attack, the SER is very close to zero. The FO attack impacts both the channel and FO estimations. We measure the overall impact for different modulation schemes by measuring the SER and BER. First, we consider the case when $\widetilde{\Delta\varphi} < \Delta\varphi_l$ and the LTSs are still able to correct the FO. In this case, Bob tries to minimize the error of estimating a channel phasor that is supposedly responsible for the phase shift accumulations over LTS samples. Because the phase shift $\Delta\varphi = 2\pi\Delta f t$ is linear in time, the best estimate is a phasor that equals to the average phase shifts. As long as $|\widetilde{\Delta\varphi}| \leq \Delta\varphi_l$ (i.e., the resulting FO is still

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

less than th_l), the maximum phase offset between the first and last samples in an LTS is π, which implies that the error in phasor estimation is always less than $\pi/2$. On the constellation map, This error will cause an identical rotation of all the payload's modulated samples [1]. We select to apply channel estimation to one LTS to limit the amount of rotation. Fig. 12(e) shows the SER for different modulation schemes. Clearly, BPSK is the most resilient scheme against channel phasor estimation error. Once $|\widetilde{\Delta\varphi}| > \Delta\varphi_l$ and the subcarriers are shifted, the sequence of modulated samples of any modulation scheme looks random relative to its original sequence, resulting in the highest possible SER, i.e., $(|M|-1)/|M|$, where $|M|$ is the modulation order.

The BER under higher-order modulation schemes, however, is less affected by the attack if the subcarriers are not shifted but the symbols are rotated to neighboring regions, as shown in Fig. 12(f). With the increase of D_{eb}, first the BER of 16-QAM starts to increase due to symbol errors. However, once QPSK also experiences symbol errors, its BER will be larger than the one for 16-QAM. Because of the Gray code structure, higher-order modulations guarantee lower BER when one of the neighboring symbols is mistakenly demodulated instead of the true symbol. Nonetheless, as long as the FO attack shifts the subcarriers, the BER stays at its maximum (0.5), irrespective of the modulation scheme.

6.2 USRP Experiments

We experimentally evaluate the impact of the proposed FO attack using an NI-USRP 2922 testbed, operated in an indoor environment in the 2.4 GHz band and controlled by Windows-based hosts. Our setup consists of three USRPs, acting as Alice, Bob, and Eve. To estimate the FOs between the USRPs, we connected Alice and Eve devices to Bob through an SMA cable and conduct 4000 FO estimations. Δf_{ab} and Δf_{eb} were measured to be 1086 and 340 Hz with standard deviations of 270 and 230 Hz, respectively. Based on the estimated FOs, effective FO D_{eb} was approximately found to be $715.5 + \Delta f_n$ Hz with standard deviation of 355 Hz. In our experiments, we fix the locations of Alice and Bob and move Eve to create two scenarios: LOS and non-LOS (see Table 2). In the non-LOS case, a metal shelf is placed between Eve and the other two. At each location, Eve launches the attack with different jamming powers and different values of Δf_n. In the experiments, Alice's transmission power is set to 7.9 dBm.

Scenario	Alice-Bob	Eve-Alice	Eve-Bob
LOS	1.5 m	1.77 m	1.77 m
NLOS	1.5 m	4.74 m	5.15 m

TABLE 2
Distances between Alice, Bob, and Eve for two different scenarios.

The USRP-based implementation of our reactive attack faced two challenges. First, the internal buffer size of the USRPs, which is used to store the samples before forwarding them to the host PC, is not big enough to store the samples captured at the nominal rate of 20 MHz. So we had to reduce the symbol rate to 0.2 MSPS. As a consequence, λ_{STS} expanded to 80 μs and f_Δ dropped to 3125 Hz (i.e., the total bandwidth of 200 kHz). Second, the USRP's reaction

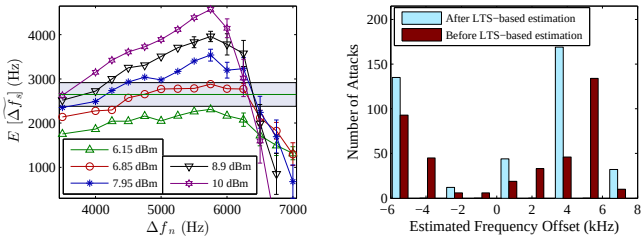

(a) Effect of Δf_n on STS-based FO estimation with different jamming powers.

(b) Effect of LTS-based FO estimation on $\widetilde{\Delta f_s}$ ($\Delta f_n = 7000$ Hz and jamming power $= 8.9$ dBm).

Fig. 13. USRP results: Performance of the FO attack in LOS scenario.

time (which consists of the communication delay between a USRP and its host PC through an Ethernet cable, the host's processing delay, and the time to initialize for transmission) exceeded several milliseconds. So Eve will miss the rest of the frame before she starts her jamming[6]. To overcome these challenges, we made the following modification in the implementation. We let Alice send several back-to-back frames periodically with a known period of T ms. Upon being triggered by a received power increase, Eve captures 2 μs worth of the sequence. If a frame is detected, she assumes that the next frame starts exactly T ms after the start of the this one. The host PC at Eve then constructs a jamming signal based on the timing information of the first detected frame and sends it to the USRP. After initialization, the USRP's onboard timer, which has nanosecond accuracy, waits for the remaining time before the next frame arrival. Once the timer expires, the device starts jamming the preamble of the new frame and other subsequent frames.

Fig. 13(a) shows the average STS-based estimation of Δf_{ab} in the LOS scenario for different values of Δf_n and jamming powers. Because our USRPs do not have stable oscillators and hence Δf_{ab} varies with time, we represent the probable value of $\Delta f_{ab} + f_\Delta/2$ by a shaded area whose height is twice the standard deviation of Δf_{ab}. Eve is able to shift the subcarriers by pushing $\widetilde{\Delta f_s}$ beyond the actual value of $\Delta f_{ab} + f_\Delta/2$. The results show that even though Eve-Bob distance is larger than Alice-Bob distance, Eve can shift the subcarriers using almost the same power as Alice's power if Δf_n is optimally selected. In particular, when the jamming signal is 7.95 dBm, Eve is successful in shifting the subcarreirs in 84% of the attacks if $\Delta f_n = 5500$ Hz. This validates our optimal Δf_n selection scheme (see Section 4.4) since the "estimated" optimal Δf_n in our setup is $715.5 + 4687.5 = 5403$ Hz. After STSs, LTS-based estimation rounds the FO to the nearest multiple of f_Δ. In Fig. 13(b), we depict a histogram to show/compare the number of jamming attacks that result in different ranges of FO estimates at Bob, before and after LTS-based estimation. It shows how LTSs can complacently exacerbate the FO estimation error. We show the results for the NLOS scenario in Fig. 14(a). As seen in this figure, the lower the jamming power, the smaller is the optimal value of Δf_n, which is inline with Fig. 10.

In the above results, we notice that the 95% confidence intervals at higher Δf_n values are noticeably larger than

6. This is not the case for an off-the-shelf reactive jammer, which usually has an onboard processor and dedicated hardware. In addition, implementing a correlation-based reactive jammer on the USRP's FPGA can achieve a reaction time of 2.56 μs [20].

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

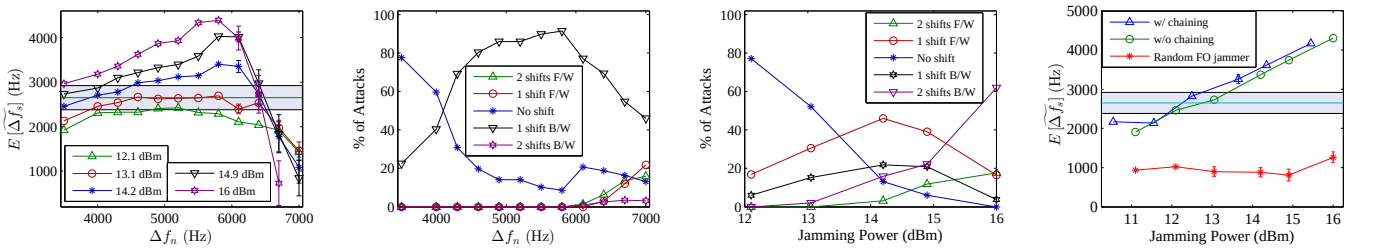

(a) Effect of Δf_n on STS-based FO estimation with different jamming powers.

(b) Effect of Δf_n on the amount of subcarrier shifts (jamming power $=$ 14.2 dBm).

(c) Effect of jamming power on the amount of subcarrier shifts ($\Delta f_n =$ 7000 Hz).

(d) Performance of different variants of the FO attack vs. jamming power ($\Delta f_n =$ 5700Hz).

Fig. 14. USRP results: Performance of the FO attack in the NLOS scenario. Alice's signal power is 7.9 dBm.

those at smaller Δf_n values. According to Fig. 9, higher values of effective Δf_{eb} may result in estimating a negative FO (when $\angle \mathcal{C} > \pi$) and thus the variance of $\widehat{\Delta f_s}$ increases. A negative FO estimate results in forward subcarrier shifts, instead of backward shifts. To illustrate this behavior, in Fig. 14(b) we plot the impact of Δf_n on the amount of subcarrier shift when the jamming power is 14.2 dBm. The attack achieves the highest success rate when $\Delta f_n = 5800$ Hz. As Δf_n increases further, the success rate slightly decreases, but Eve can impose various amounts of subcarrier shift, which can be leveraged to make it more difficult for Bob to guess the amount of subcarrier shift. Fig. 14(c) shows the effect of jamming power on the amount of subcarrier shift when Δf_n is high. As the jamming power increases, Eve not only can achieve a higher success rate, but can also impose more than one subcarrier shift (forward or backward).

We compare the FO attack, with and without the chaining rule, against a random FO jammer in Fig. 14(d). In this experiment, we configure Eve's USRP to start jamming zero, one, or two time indices before the estimated start of STSs. The random FO jammer generates uniform white noise. The results confirm that the chaining rule strengthens the attack while random jamming cannot manipulate Δf_s and overcome the LTSs even with high jamming power.

Finally, we launch the FO attack during the transmission of a packetized image. Specifically, Eve attacks 24 packets in the middle of the transmission of 44 QPSK-modulated 720-byte-long packets that represent the image in Fig. 15(a). In Fig. 15(b), we show the received image. The parts that experience FO jamming are completely destroyed.

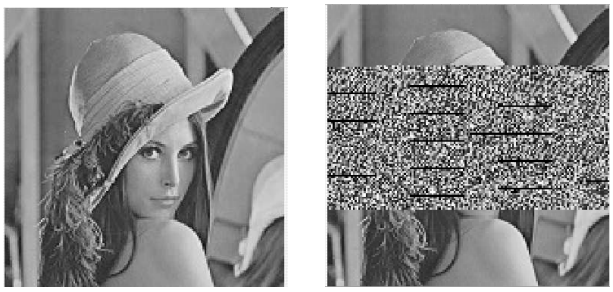

(a) Transmitted image

(b) Received image

Fig. 15. Applying the FO attack on a sequence of frames belonging to an image.

7 Defense Strategies

Alice and Bob may work cooperatively or independently to mitigate the previously presented FO attack. For example, they may prevent accurate estimation of Δf_{ab} at Eve by transmitting Tx-based friendly jamming. This, however, requires additional antennas. Bob may also use the power-spectral density of the captured signal after LTSs to identify the missing subcarriers, and thus determine the overall subcarrier shift. This technique, however, fails if Eve transmits only one or two bogus subcarriers to replace missing subcarriers. Assuming that Bob is not equipped with additional antennas, we propose three preliminary approaches for mitigating the FO attack. Analysing and evaluating these strategies are beyond the scope of this paper, and will be addressed in future research.

1) Randomizing FO Sequences (Sequence Hopping): Because of the redundancy in the STSs, Bob can choose any pair of the ten consecutive STSs, to perform FO estimation. Furthermore, due to the maximum FO requirement for 802.11 devices (212 kHz $= 1.3568\, th_s$ for devices operating in the 5 GHz band and 125 kHz $= 0.8\, th_s$ for devices in the 2.4 GHz band [17]), the two autocorrelation windows do not necessarily need to be contiguous. In fact, the two windows can be two or four STSs apart (i.e., each sample is three or five STSs away from its dual) in the 5 GHz and 2.4 GHz bands, respectively. This means that Bob has the flexibility to randomly hop to any pair of STSs for FO estimation, given that the STSs in this pair are not more than two or four STSs apart, depending on the frequency band. Even if Bob selects an STS that is corrupted by a jamming signal together with a jamming-free one, he is still able to estimate the same FO as if two jamming-free sequences are selected [1]. To implement sequence hopping, Bob can record the received signal (the ten STSs) while he is in the process of detecting the start of the frame. Once the frame has been detected, Bob randomly chooses two STSs for FO estimation, while satisfying the maximum STS-distance constraint.

2) Preamble Obfuscation: Preamble obfuscation aims at making the timing or FO features hard to extract by Eve. We provide one simple example for timing extraction. Alice can obfuscate the preamble by adding artificial noise that is only known to Bob. He, on the other side, modifies the denominator in (6) to account for the power of the artificial noise during a certain section of the received signal. For example, a signal identical to the first half of an STS may be added to the first half of the second STS in the preamble. So, Bob can still detect the frame by doubling the denominator in (6), but the increased power at Eve would decrease the

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

value of $\mathcal{M}(n)$ through (6) for the first $L/2$ samples. Hence, $\hat{\mathcal{M}}$ will be in the second half, making the chaining rule fail because Eve does not include the actual start time in the $V = \log_2 L$ start times.

3) STS Bypassing: Bob can simply disable the STS-based FO estimation mechanism to dodge the FO attack. However, he still has to meet the requirement for coarse FO estimation, i.e., unambiguous phase estimation (see Section 2.1). Under BPSK modulation, which is typically used to transmit the PHY header, Bob can tolerate channel estimation errors due to an FO estimation error of up to 15 kHz [13]. Hence, Bob can divide the range of possible FO values into several equal-size frequency bins, each of 30 kHz bandwidth. He can then try each of the possible bins and compensate for its center frequency before applying LTS-based FO estimation. The center frequency that results in the minimum MSE in channel estimation can be considered as $\widetilde{\Delta f_s}$. Bob may also suppress both STS- and LTS-based FO estimation, and instead rely on pilot subcarriers for FO and channel estimation. This approach, however, often gives rise to ICI because adjacent subcarriers interfere with the pilot subcarriers (which even are not yet channel-equalized) and the FO estimation will be erroneous.

8 Related Work

Vulnerabilities of wireless protocols and Denial-of-service (DoS) attacks have been studied in the literature since the early 2000s. DoS attacks can be applied at either MAC or PHY-layer. MAC layer attacks usually take the form of malicious packet insertion. For example, in the deauthentication deadlock attack [21], a specific packet is injected at a particular time during the EAPOL four-way handshake of 802.11, leading to DoS. In contrast, RF jamming is a form of PHY-layer attack.

RF Jamming techniques are categorized into constant, deceptive, random, reactive, and short noise-based (narrow-band) intelligent jamming methods [4], [22]. Constant, deceptive, and random jamming models achieve a high level of DoS by excessively transmitting over the channel, but exhibit poor energy efficiency and high detection probability [22]. On the other hand, energy-efficient reactive jamming attacks select and target a (part of a) packet based on traffic analysis, protocol semantics, or publicity of some fields [3], [4], [8]–[10], [20]. These attacks may fail to significantly corrupt ongoing transmissions if, for example, channel hopping, randomization, and coding are used to hide the transmission features.

The efficiency of reactive jamming is assessed by the effort needed to drop a packet. In [3], jamming efficiency is defined as the ratio of communication effort to jamming effort. The authors demonstrated the jamming efficiency of $50 \sim 500$ in 802.11a by jamming an OFDM symbol. Using a high duty-cycle jammer, Gummadi *et al.* [9] could disrupt a link when the jamming power is 1000 weaker than the signal power by targeting timing recovery, dynamic range selection (AGC), and header processing. The authors in [23] observed that 22 μs of jamming is sufficient to make a frame undecodable. In comparison, our FO attack can achieve a jamming efficiency of $136 \sim 1400$ in 802.11a and defeat any ECC by jamming for only 2.8 μs.

Jamming OFDM systems is of particular interest due to their widespread use in modern systems. Simple barrage jamming targets the entire spectrum/tones and corrupts more bits than the more power efficient but less destructive partial band, single- and multi-tone jamming [7], [24]. Asynchronous off-tone jamming attacks exploit the uncompensated FO between Eve and Bob to transmit one or multiple subcarriers that will be received between some of the data subcarriers [7]. This creates significant ICI for those subcarriers. Though energy-efficient, these attacks cannot achieve 50% BER. Furthermore, because coding and interleaving are employed in the 802.11 systems for robustness against narrow-band interference, Bob may still be able to recover the frame. Several pilot jamming attacks were proposed in [6], [7] in order to distort channel estimation. In contrast, our proposed attack lasts for less than the duration of a pilot symbol jamming and corrupts the channel estimation without jamming pilots. Jamming against timing acquisition in OFDM systems and some countermeasures were discussed in [8], [11], [12]. However, OFDM systems are more sensitive to FO than timing errors. This vulnerability was first revealed in [10]. The structured but essentially random jamming scheme in [10], however, does not provide any performance guarantee and may have higher jamming effort than ours. Interested readers are referred to [19], [25] for more details about jamming attacks against OFDM systems.

9 Conclusion

We demonstrated the vulnerability of OFDM systems against a highly disruptive but efficient-efficient DoS attack. This attack succeeds even when the PHY frame is shielded by interleaving and channel coding. The attack focuses on the frequency offset (FO) estimation process, and is channel-independent and robust to time-synchronization errors at Eve through applying the proposed pairing and chaining rules. Through this attack, a reactive jammer exploits and targets a small portion of the publicly known preamble used for FO estimation. The attack lasts for less than the duration of an OFDM symbol, i.e., less than 1% of a typical frame duration, and is at least 30% more efficient than previously reported attacks. Though short-lived, the attack results in a shift in subcarrier indices and the maximum possible BER (50%) even when the jamming signal at Bob is ~ 1.4 times weaker than Alice's signal. We verified via simulations and USRP experimentation. The simulation results show that different modulation schemes are equally susceptible if the FO attack can shift the subcarrier indices, and higher modulation orders are more affected when the attack impacts only the channel estimation. Finally, we sketched several possible mitigation approaches, whose detailed analysis and evaluation are left for future.

Acknowledgements

This research was supported in part by the NSF (grants # IIP-1265960 and CNS-1409172), in part by the ARO (grant # W911NF-13-1-0302), and in part by the AFOSR (grant # FA2386-13-1-3026). Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of NSF, ARO, or AFOSR.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMC.2015.2456916, IEEE Transactions on Mobile Computing

REFERENCES

[1] H. Rahbari, M. Krunz, and L. Lazos, "Security vulnerability and countermeasures of frequency offset correction in 802.11a systems," in *Proc. IEEE INFOCOM'14*, Toronto, ON, Canada, Apr. 2014, pp. 1015–1023.

[2] "IEEE Std 802.11a-1999," *Supplement to IEEE Standard for Information Technology— Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements: Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-Speed Physical Layer in the 5 GHz Band*, 1999.

[3] G. Lin and G. Noubir, "On link layer denial of service in data wireless LANs," *Wireless Communications and Mobile Computing*, vol. 5, no. 3, pp. 273–284, 2005.

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

[5] M. Strasser, B. Danev, and S. Čapkun, "Detection of reactive jamming in sensor networks," *ACM Trans. Sensor Networks*, vol. 7, no. 2, pp. 16:1–16:29, Aug. 2010.

[6] T. C. Clancy, "Efficient OFDM denial: Pilot jamming and pilot nulling," in *Proc. IEEE Int. Conf. Commun. (ICC'11)*, Kyoto, Japan, Jun. 2011.

[7] C. Shahriar, S. Sodagari, R. McGwier, and T. Clancy, "Performance impact of asynchronous off-tone jamming attacks against OFDM," in *Proc. IEEE Int. Conf. Commun. (ICC'13)*, Budapest, Hungary, Jun. 2013, pp. 2177–2182.

[8] M. J. LaPan, T. C. Clancy, and R. W. McGwier, "Physical layer orthogonal frequency-division multiplexing acquisition and timing synchronization security," *Wireless Communications and Mobile Computing*, 2014, to be published. [Online]. Available: http://dx.doi.org/10.1002/wcm.2500

[9] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, "Understanding and mitigating the impact of RF interference on 802.11 networks," in *Proc. ACM SIGCOMM'07*, Kyoto, Japan, 2007, pp. 385–396.

[10] M. J. LaPan, T. C. Clancy, and R. W. McGwier, "Phase warping and differential scrambling attacks against OFDM frequency synchronization," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP'13)*, Vancouver, BC, Canada, May 2013, pp. 2886–2890.

[11] C. Mueller-Smith and W. Trappe, "Efficient OFDM denial in the absence of channel information," in *Proc. Military Commun. Conf. (MILCOM'13)*, San Diego, CA, USA, Nov. 2013, pp. 89–94.

[12] L. Sanguinetti, M. Morelli, and H. Poor, "Frame detection and timing acquisition for OFDM transmissions with unknown interference," *IEEE Trans. Wireless Commun.*, vol. 9, no. 3, pp. 1226–1236, Mar. 2010.

[13] J. Heiskala and J. Terry, *OFDM Wireless LANs: A Theoretical and Practical Guide*. SAMS Publishing Indianapolis, 2002.

[14] L. Weng, R. Murch, and V. Lau, "SISO-OFDM channel estimation in the presence of carrier frequency offset," in *Proc. IEEE Wireless Commun. and Netw. Conf. (WCNC'06)*, vol. 3, Las Vegas, NV, USA, Apr. 2006, pp. 1444–1449.

[15] T. M. Schmidl and D. C. Cox, "Robust Frequency and Timing Synchronization for OFDM," *IEEE Trans. Commun.*, vol. 45, no. 12, pp. 1613–1621, 1997.

[16] T. Pollet, M. Van Bladel, and M. Moeneclaey, "BER sensitivity of OFDM systems to carrier frequency offset and Wiener phase noise," *IEEE Trans. Commun.*, vol. 43, no. 234, pp. 191–193, 1995.

[17] "IEEE Std 802.11n-2009," *IEEE Standard for Information Technology— Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements: Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput*, 2009.

[18] H. Rahul, H. Hassanieh, and D. Katabi, "SourceSync: a distributed wireless architecture for exploiting sender diversity," in *Proc. ACM SIGCOMM'10*, New Delhi, India, Sep. 2010, pp. 171–182.

[19] H. Rahbari, M. Krunz, and L. Lazos, "Jamming attack on frequency offset estimation in OFDM systems," University of Arizona, Department of ECE, Tech. Rep. TR-UA-ECE-2015-1, Jun. 2015. [Online]. Available: http://goo.gl/pAxxaR

[20] D. Nguyen, C. Sahin, B. Shishkin, N. Kandasamy, and K. R. Dandekar, "A real-time and protocol-aware reactive jamming framework built on software-defined radios," in *Proc. ACM Workshop Software Radio Implementation Forum (SRIF'14)*, Chicago, IL, USA, Aug. 2014, pp. 15–22.

[21] M. Eian and S. Mjølsnes, "A formal analysis of IEEE 802.11w deadlock vulnerabilities," in *Proc. IEEE INFOCOM'12*, Orlando, FL, USA, Mar. 2012, pp. 918–926.

[22] N. Sufyan, N. Saqib, and M. Zia, "Detection of jamming attacks in 802.11b wireless networks," *EURASIP J. Wireless Commun. and Networking*, vol. 2013, 2013.

[23] G. Noubir, R. Rajaraman, B. Sheng, and B. Thapa, "On the robustness of IEEE 802.11 rate adaptation algorithms against smart jamming," in *Proc. Fourth ACM Conf. Wireless Network Security (WiSec'11)*, Hamburg, Germany, Jun. 2011, pp. 97–108.

[24] J. Luo, J. Andrian, and C. Zhou, "Bit error rate analysis of jamming for OFDM systems," in *Wireless Telecommun. Symp. (WTS'07)*, Apr. 2007, pp. 1–8.

[25] C. Shahriar, M. LaPan, M. Lichtman, T. Clancy, R. McGwier, R. Tandon, S. Sodagari, and J. Reed, "PHY-Layer resiliency in OFDM communications: A tutorial," *Commun. Surveys Tuts.*, vol. 17, no. 1, pp. 292–314, 2015.

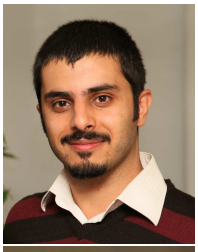

Hanif Rahbari is an electrical and computer engineering Ph.D. candidate at the University of Arizona. He has received his BSc from Sharif University of Technology and his MSc in computer networks from AmirKabir University of Technology, Iran. His research interests include wireless communications and networking, PHY-layer security, hardware experimentation, multimedia, and dynamic spectrum access networks.

Marwan Krunz (S'93-M'95-SM'04-F'10) received the Ph.D. degree in electrical engineering from Michigan State University in 1995. He is a Professor of ECE and CS at the University of Arizona. He is the Site Codirector at the US National Science Foundation (NSF) Broadband Wireless Access and Applications Center. He joined the University of Arizona in January 1997, after a brief postdoctoral stint at the University of Maryland. In 2010, he was a Visiting Chair of Excellence at the University of Carlos III de Madrid. He previously held other visiting research positions at INRIA-Sophia Antipolis, HP Labs, University of Paris VI, University of Paris V, and US West Advanced Technologies. His research interests include the areas of wireless communications and networking, with emphasis on resource management, adaptive protocols, and security issues. He has published more than 215 journal articles and peer-reviewed conference papers, and is a coinventor on five US patents. He received the 2012 IEEE TCCC Outstanding Service Award. He received the US NSF CAREER Award in 1998. He currently serves on the editorial board for the IEEE TRANSACTIONS ON COGNITIVE COMMUNICATIONS AND NETWORKING. Previously, he served on the editorial boards for the IEEE/ACM TRANSACTIONS ON NETWORKING, IEEE TRANSACTIONS ON MOBILE COMPUTING, IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT, Computer Communications Journal, and the IEEE Communications Interactive Magazine. He was the General Cochair for WiSec'12, and served as a TPC Chair for INFOCOM'04, SECON'05, and WoWMoM'06. He was the keynote speaker, an invited panelist, and a tutorial presenter at numerous international conferences. He is an Arizona Engineering Faculty fellow (2011-2014), and an IEEE Communications Society Distinguished lecturer (2013–14).

Loukas Lazos is an Associate Professor of Electrical and Computer Engineering at the University of Arizona. He received his Ph.D. in Electrical Engineering from the University of Washington in 2006. In 2007, he was the co-director of the Network Security Lab at the University of Washington. Dr. Lazos joined the University of Arizona in August 2007. His main research interests are in the areas of network security, user privacy, wireless communications, network performance analysis, and network visualization. He is a recipient of the NSF CAREER Award (2009), for his research in security of multi-channel wireless networks. He was the general co-chair for the ACM WiSec 2012 Conference and served as the TPC co-chair for the Communication and Information System Security Symposium at GLOBECOM 2013 and the 4th IEEE International Workshop on Data Security and Privacy in Wireless Networks (DSPAN) 2013. He has served on organizing and technical program committees of numerous international conferences and on panels for several funding agencies.