Friendly Jamming on Access Points: Analysis and Real-World Measurements

Daniel S. Berger, *Student Member, IEEE*, Francesco Gringoli, *Member, IEEE*, Nicolò Facchi, *Member, IEEE*, Ivan Martinovic, *Member, IEEE*, and Jens B. Schmitt, *Member, IEEE*

Abstract—Frequency jamming is known as an efficient attack tool to disrupt wireless communication. This efficiency can also be exploited for the benefit of a network—an idea often referred to as friendly jamming. A prominent application case is the blocking of unauthenticated or malicious communication, such as injection attacks. In this paper, we propose access points as a natural place to implement friendly jamming functionality. We analyze this proposal using simulations, introduce an implementation on customer-grade access points, and report measurement results from the first real-world study of friendly jamming in an IEEE 802.11 campus network. We discover a fundamental tradeoff between the effectiveness of friendly jamming and the orthogonal aspect of having minimal side-effects to the campus network’s traffic. In particular, we observed what we call the power amplification phenomenon. This effect aggravates the known hidden station problem when the number of jammers increases. We also find evidence that the collaboration between jammers can enable friendly jamming, which is both effective and minimally invasive.

Index Terms—Friendly jamming, jamming for good, defensive jamming, reactive jamming, IEEE 802.11, Wi-Fi, WLAN.

I. INTRODUCTION

RADIO frequency jamming induces intentional interference that disrupts wireless communications on the physical layer. Therefore, jamming is hard to mitigate and commonly understood as a severe threat to wireless networks [21], [25], [26], [30], [44]. Unconventionally, jamming can also be used for the benefit of wireless communications. Two general application scenarios are usually promoted in the literature: (1) jamming as an intentional interference to prevent unauthenticated data from reaching the legitimate but resource-limited devices [5], [11], [16], [23], [24], [35], [42], and (2) jamming as a method to prevent eavesdropping on user communication by increasing the noise floor towards attacker [15], [17], [34], [39].

Manuscript received June 23, 2015; revised November 9, 2015 and February 29, 2016; accepted June 1, 2016. The associate editor coordinating the review of this paper and approving it for publication was K. Zeng.

D. S. Berger and J. B. Schmitt are with the Distributed Computer Systems Laboratory, Department of Computer Science, Kaiserslautern University of Technology, Kaiserslautern 67663, Germany (e-mail: berger@cs.uni-kl.de; jschmitt@cs.uni-kl.de).

F. Gringoli and N. Facchi are with the Department of Information Engineering, University of Brescia, Brescia 25123, Italy (e-mail: francesco.gringoli@ing.unibs.it; nicolo.facchi@ing.unibs.it).

I. Martinovic is with the Computer Science Department, University of Oxford, Oxford OX1 3QD, U.K. (e-mail: ivan.martinovic@cs.ox.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TWC.2016.2581165

Fig. 1. The friendly jamming scenario. The friendly jammer interferes with malicious communication of an attacker before it reaches the victim, while concurrent background traffic of legitimate users continues unaffected.

This work focuses on the following application scenario. An attacker seeks to send malicious messages (*attack frames*) to resource-constrained stations (*victims*). The victims are embedded devices such as sensor networks [23], [24], [42], implanted medical devices [11], or home automation and wearable gadgets [5], [16]. The victims cannot use sophisticated attack detection because they are running on battery [16], [23], [24], [42], or because updating them is unfeasible [11]. To protect such embedded devices, previous work has envisioned an external device called *friendly jammer*. Unlike the victim stations, the friendly jammer is able to detect attack frames, e.g., by real-time packet inspection [42], proximity information [16], or the absence of pre-shared tokens [11]. When detecting an attack frame, the friendly jammer emits physical-layer interference so that the victims cannot decode the attack frame and therefore discard it.

Friendly jamming promises to be a simple and effective protection mechanism for resource-constrained wireless devices. However, several challenges related to the physical-layer interference remain open. First, friendly jamming needs to be *effective* in a dynamic environment with multipath effects, attenuation, and fast channel fading. These effects can lead to undetected attack frames or insufficient interference with an attack frame. While prior work acknowledges these problems, the question remains how frequently this happens in a realistic environment. Second, friendly jamming also needs to be *minimally invasive*, i.e., only attack frames should be blocked, whereas the impact on legitimate transmissions should be minimized (see Fig. 1). Unfortunately, we are not aware of prior works that quantify the invasiveness of friendly jamming. Third, friendly jamming needs to be *pervasive*: friendly jamming should be available at low cost and for the most common communication protocol such as Wi-Fi (IEEE 802.11). Unfortunately, the technical challenges of friendly jamming meant that previous prototypes were based on powerful (and costly) specialized hardware platforms [5], [11], [42], [43] and do not support Wi-Fi (cf. Section II). In this work, as a part of real-world evaluation we implement

friendly jamming on off-the-shelf Wi-Fi access points. This approach solves the pervasiveness challenge, but entails new challenges for the effectiveness of friendly jamming. We analyze these challenges using both simulations and measurements with a realistic prototype based on the popular WRT54G access point. Specifically, we performed a friendly jamming experiment in a university IEEE 802.11g Wi-Fi network for three weeks: we measured both the effectiveness against the attacker and the invasiveness to the legitimate users of the network. We summarize our insights as follows:

- our analysis reveals that the WRT54G's jamming effectiveness depends not only on the power but also the symbol alignment of the jamming signal (with regard to the attack frame). The symbol alignment cannot be controlled by the friendly jammer so that multiple friendly jammers need to cooperate to increase their chance of the right alignment.
- our experimental evaluation shows that friendly jamming within a legitimate network is feasible. However, friendly jamming can increase legitimate station's packet loss due to a new type of hidden station problem, which has not been observed before and which we call *power amplification effect*. The packet loss of some stations doubled for seven jammers;
- jammers can collaborate to alleviate the power amplification effect. Our collaboration schemes can improve the packet loss by 20-30%, without compromising jamming effectiveness;
- we reproduce the previous insights with a different hardware setup. Friendly jamming is also effective for modern IEEE 802.11n networks in which the attacker can use antenna diversity, multiple spatial streams, and both the 2.4 and 5 GHz frequency band.

We discuss related work in Section II and jamming on access points (APs) in Section III; our analysis and simulation results are presented in Section IV and Section V describes our implementation. We present results from our measurement study for IEEE 802.11g and 11n in Sections VI and VII respectively. We conclude in Section VIII.

II. RELATED WORK

We review in this Section the main application cases of jamming in wireless networks, and we also discuss related works on jamming mitigation and general interference studies.

A. Blocking Harmful Communication

Table I gives an overview of prior works that use jammers to block attack frames similar to our friendly jamming scenario. With the exception of the works [24] (which does not evaluate reactive jamming) and [16] (which uses analysis and simulations), all of these works use the USRP2 software defined radio platform (e.g., [5], [11], [42], [43]) and do not target customer-grade APs. To the best of our knowledge, our system is the first to implement friendly jamming for IEEE 802.11 and on an AP. Most of the prior works used the IEEE 802.15.4 or industrial 400 MHz communication standard which have lower throughput and for which slower reaction times are sufficient

TABLE I
RELATED WORK ON BLOCKING HARMFUL COMMUNICATION

System	Technology and Evaluation
Jamming for good [24]	IEEE 802.15.4 sensor nodes, without reactive jamming
WiFire [42]	IEEE 802.15.4 USRP2, experiments with 2 friendly jammers
IMD shield [11]	400 MHz USRP2, experiments with 1 friendly jammer
Defend your home [5]	433 MHz USRP2, experiments with 1 friendly jammer
Considerate Jamming [16]	analysis and simulation

(cf. Section III). We also remark that unlike the USRP2 the WRT54G AP used in our work is likely to comply with legal regulations.¹

The performance metric considered in these works is the fraction of jammed attack frames (hit ratio): [24] report a median of 97%; [43] and [42] report 97.6% for one jammer and 99.9 for two concurrent jammers; [5] report 98.9 – 100%. All these experiments took place in a static setting (e.g., no other network, no moving people); we are not aware of a work that evaluates friendly jamming in a dynamic environment over several days as done in Section VI. The hit ratios in our experiments are significantly lower (e.g., 80 – 90% for two jammers) and we additionally consider the distribution of jammed (and not jammed) attack frames on shorter time scales.

B. Jamming for Confidentiality

In the second application case the jammer's interference ensures that an eavesdropper cannot decode network messages [15], [17], [20], [34], [39], [46]. The real-world feasibility of jamming for confidentiality has recently been thoroughly evaluated in [37].

Jammer placement schemes are considered in [17] and [32]. The setting of [17] differs from ours in that the jammers continuously emit a jam signal and thus there cannot be any co-existing networks. However, one of the goals mentioned [32] is to not interfere with any legitimate communication. This goal is similar to our collaboration schemes, but [32] relies on a theoretical propagation model and does not report experimental results. Also, our collaboration schemes address a different problem and they are based on different assumptions. Instead of the placement problem, we consider how jammers at fixed positions can collaborate efficiently. Instead of maximum interference at the attacker, we seek to maximize interference at the victim.

C. Jamming Mitigation

In the military context, jamming is an established technique to block an enemy's communication [1], [28]. Jamming has also been recognized as a threat to civilian networks [44].

¹Our prototype uses FCC-certified hardware with short low-power jamming signals that comply with limitations on maximum power and duty cycle [8].

Therefore, many works study ways to mitigate the effects of jamming [21], [25]–[27], [30], [44]. Many of the proposals (e.g., frequency hopping and robust modulation schemes) cannot be exploited by the attacker, as they would require cooperation of the victim stations. We assume that the victim stations follow the IEEE 802.11 protocol faithfully. This leaves the attackers a choice within the a set of allowed modulation and coding schemes (MCS). We evaluate the effect of this choice by considering all MCS configurations in our experiments.

D. Interference and Collisions in 802.11

Our analysis on jamming success in Section IV is related to the studies of packet error probabilities in face of non-intentional interference in IEEE 802.11 networks. These studies are motivated by the capture effect, which describes the successful reception of the stronger of two simultaneous transmissions [18]. Although intentional interference motivates a different perspective, the corresponding analysis often tries to answer the same questions. In particular, while initial capture models only considered the role of relative timing on the frame level and the received power [18], [19], [40], more recent models report observations that agree with our analysis: frame level temporal displacement and power ratio between interfering signals are insufficient predictors of interference probabilities. Our simulations in Section IV are most similar to [7], [10], [29], [33], and [41], which consider the effect of interference for different time and phase offset of the sender and receiver. However, these works only consider the temporal displacement at the frame preamble level [33]; focus on different technologies (IEEE 802.15.4) [41]; analyze the effects of jamming at the bit level when neither spreading or other explicit error detection and correction techniques are in use [29]; and consider only the relative time offset ignoring the relative power of the interfering signals [10].

What makes our analysis different from previous works is that we focus on IEEE 802.11 modulations and encoding schemes, we consider a scenario in which the receiver is already synchronized with the attack frame when the intentional interference starts and we take into account both the power ratio and the temporal displacement between interfering signals.

III. FRIENDLY JAMMING ON ACCESS POINTS

This section introduces our attacker model and highlights three fundamental challenges of friendly jamming which lead to our proposal of implementing friendly jamming on APs.

A. Attacker Model

The attacker seeks to send attack frames to victim stations, which are part of the wireless network. Victim stations faithfully follow the IEEE 802.11 protocol (uncompromised user stations), but do not implement any security functionality. Friendly jammers seek to stop each attack frame but emitting physical layer interference. Both the attacker and the friendly jammer are assumed to use standard hardware, which are

Fig. 2. Reactive jamming happens in four phases: 1) the jammer detects the attack frame (using the IEEE 802.11 training sequence), 2) analyzes the signal (to decide whether or not to jam), 3) incurs a reaction delay (due to switching from receiving to sending), and 4) emits the jamming frame. The interference is successful, if the superimposed signal's power P_{jammed} is significantly higher than P_{attack} .

power constrained to similar values. The attacker uses either an omnidirectional antenna (Section VI) or two omnidirectional antennas (Section VII).

B. System Challenges of Friendly Jamming

1) *The Reactive Jamming Challenge*: There are several jamming techniques that can block attack frames at the physical layer, e.g., continuously or randomly emitting a high power signal [26], [45]. However, *reactive jamming* is the most power-efficient technique and also the only one that can achieve minimal invasiveness: in reactive jamming, in fact, the interfering signal is selectively emitted only when another specific signal is detected on the channel.

Reactive jamming starts with the task of channel sensing in order to detect an attack frame (a signal in IEEE 802.11 terminology). The attack frame's detection and further decoding can only be achieved by jammers which are well positioned with regard to the arriving signal (to be specific: the attack signal arrives with high signal-to-noise-ratio (SNR)). This position is thus an important factor for jamming success. After having detected the signal, the jammer starts the decoding and analysis in order to make a jamming decision (cf. Fig. 2a). The jammer needs to distinguish between attack frames and legitimate frames. While the hit ratio ($\# jammed frames / \# sent frames$) needs to be maximized, the impact of jamming on legitimate frames shall remain small. This trade-off is further constrained by the fact that the decision has to be taken very fast. We recall these timing constraints by considering a frame comprising a medium-sized 100 Bytes UDP packet. If this frame is transmitted at the fastest 802.11g Modulation and Coding Scheme (MCS), its duration is only $48\mu s$. If we further assume that the jamming decision depends on bits in the MAC header (e.g., the sender's MAC address), the remaining part of the frame is another $24\mu s$ shorter. In the remaining time, the jammer's decision code has to be executed, and the hardware needs to switch from receive to transmission mode to emit the jamming signal. For faster communication standards like IEEE 802.11n the timing constraints become even more stringent. Once the signal is emitted, successful jamming depends on the attack frame's MCS and the jamming signal's power (cf. Section IV). While higher power drives high jamming success, this factor is subject to legal regulations and increases the detrimental impact on legitimate traffic.

2) *The Minimal-Invasiveness Challenge*: Although reactive jamming is key to achieving minimal invasiveness, there are

Fig. 3. The power amplification effect. (a) A legitimate user transmission is interference-free from the attacker's injections. (b) If a friendly jammer blocks the attack, it may accidentally interfere with the legitimate transmissions.

still side effects that can introduce interference to legitimate transmissions. We call the major such side effect the *power amplification phenomenon*.

Power amplification happens when friendly jammers significantly increase the interference radius of an attack transmission. This can be seen by considering a scenario in which a legitimate sender is located far away from the attack sender and transmissions of the two senders would be free of interference. Assume a jammer is in transmission (jamming) range of both senders: if the legitimate sender starts a transmission, the attack source could possibly start a concurrent transmission as well. Without jammers, both transmissions may proceed without problems (Fig. 3a); once the jammer interferes with the attack transmission, however, it also creates interference with the legitimate transmission (Fig. 3b). This interference can cause failure of the legitimate transmission. While in outdoor testbeds such constellations might be rare, we find that they happen more often in an indoor office environment.

Note that power amplification is different from the known hidden station problem. In hidden station problems, frames collide at a common receiver of two senders which cannot see each other. However, in the power amplification scenario in Fig. 3b the effect is due to the reactive jamming technique: both transmissions would proceed without problems, if the jammers were turned off and if the jammers were continuously jamming.

3) *The Pervasiveness Challenge*: Serving all users in a wireless network entails several sub-challenges. Firstly, friendly jamming needs to be available for those communication standards that are widely in use, such as Wi-Fi. Secondly, wireless communication underlies legal restrictions.² Such restrictions often define a maximum duty cycle (the fraction of one second a transmitter is active) that is commonly limited to small values and the maximum mean equivalent isotropic radiated power has to be less than a limit (e.g., 20 dBm for 2.4 GHz band in Europe [8]). In order to make the move from specialized use cases (e.g., military scenarios [1], [28]) to pervasive availability, friendly jamming needs to satisfy and work within these constraints.

Finally, friendly jamming needs good scalability to achieve coverage for every user in the network. This can only be achieved when a considerable number of well-positioned friendly jammers work together; in fact, a previous work even envisioned all network nodes to have friendly jamming functionality [24]. In any case this assumes that friendly jamming is made easily available, and at low cost. While previous reactive jamming and friendly jamming implementations exploited

the power of software-defined radios (SDR) (e.g., [2], [5], [42], [43]), scalability and low cost suggest to use customer-grade components. This is a great challenge as such components are typically underpowered which makes it hard to meet the timing constraints imposed by reactive jamming (which already pose a challenge for powerful software-defined platforms [2]). For customer-grade equipment, solutions to the underlying engineering challenges, such as reactive jamming, have begun to emerge only recently [3], [38].

C. Bringing Friendly Jamming to IEEE 802.11 Access Points

Our principal idea is to implement friendly jamming by modifying the software running on customer-grade APs. This has several advantages. Firstly, customer APs are optimized for the specific tasks (e.g., signal processing) necessary for Wi-Fi compatibility. They already implement current Wi-Fi features and can be expected to be kept up to date with future standards. Secondly, customer-grade APs are certified to be fully compliant with legal regulations. If we only modify the software, we are sure to encounter realistic hardware constraints. Thirdly and most importantly, APs are a natural deployment choice. Conceptually, APs should be the ones controlling a network's smooth operation and have much of the necessary network state information readily available. APs are also deployed in good positions to ensure coverage of all users in the network.

Modifying only the software of APs to implement friendly jamming functionality also poses several challenges. First, because of the real-time requirements of reactive jamming and delays in non-real-time architectures (e.g., bus delays), friendly jamming has to be implemented directly on the network interface card (NIC). This requires the ability to modify the firmware running on the NIC and detailed knowledge about the underlying hardware (chipset). Two critical changes to the firmware have to be made: the firmware needs to analyze a frame while still receiving it and it needs to be able to abort the current reception, switch to transmission mode, and emit the jam frame. As this behavior explicitly violates the IEEE 802.11 collision avoidance scheme, one needs to be able to work around such constraints. Additionally, to work as a friendly jammer, the system must support dynamic reconfiguration of the attack signatures.

Another part of the challenge is that APs are optimized for low cost: the WRT54GL AP – our choice – costs about 30-50 USD (5/2015 on ebay.com) and features a 200 MHz CPU and 16 MB of memory.

Finally, a major implication of the software only approach is that the jamming signal has to be a valid IEEE 802.11 frame: the encoding and modulation pipeline of Wi-Fi APs are hardwired and cannot be changed with a software update. It is not obvious that such a jamming signal can be used to interfere and thus block attack frames sent using the same technology.

IV. ANALYSIS OF FRIENDLY JAMMING ON APs

Jamming can be unsuccessful for many reasons, e.g., the friendly jammer might miss or react too late to a target frame,

²In particular, jamming is a sensitive topic for many country's legal regulators. For example, in the US the operation, marketing, or selling of continuous jammers such as GPS jammers, or cell phone blockers is prohibited [9].

Fig. 4. (a) Phase transition diagram for a real IEEE 802.11 frame (Broadcom). (b) The corresponding amplitudes are not constant over time but peak in the middle of a chip. The dashed line shows the amplitude of frame which is shifted by 1/2 a chip (Broadcom). (c) Amplitudes for an Atheros card.

or the jamming signal might not be powerful enough. This section focuses on the last case that is particularly challenging for friendly jamming on APs as they are power-constrained by legal regulations. In particular, we are interested in understanding how jamming success is affected by power ratio and temporal displacement between interfering signals. In the following, we consider DSSS and OFDM separately, and focus on the respective slowest data rates (1 Mb/s and 6 Mb/s) which are known to be the most resilient to interference. Also, in order to obtain realistic numbers, we ran trace-driven simulations by feeding IEEE 802.11 software receivers with the actual I/Q samples that we captured from real signals.

To the best of our knowledge, this is the first analysis of the relative power requirements for successful jamming as a function of chip or symbol alignment for IEEE 802.11abgn

A. Analysis of DSSS

We start by describing the key encoding steps defined by IEEE 802.11 [14] for the DSSS-PHY. This section assumes some prior knowledge about DSSS, which can be found in [28].

1) *Encoding DSSS*: First, bits are pre-processed to have a balanced number of 0s and 1s (called scrambling). Second, every bit is expanded to eleven chips with redundant information (called spreading gain, it also boosts the bandwidth of the transmitted signal to fit in the typical 20 MHz spectrum considered by the standard). Third, each chip is modulated with *Differential Binary Phase-Shift Keying* (DBPSK), which changes the phase by π for a 1 and leaves the phase unchanged for a 0. Fig. 4a shows the resulting phase transitions for a real IEEE 802.11 frame. The chip boundaries are marked with vertical lines. We observe that although DBPSK's phase changes are instantaneous in theory, they happen gradually in the real world due to hardware limitations. Figs. 4-8 (except Fig. 6) are based on traces of customer-grade IEEE 802.11 hardware captured with a real-time spectrum analyzer (Tektronix RSA3408). These traces have a resolution of four samples per chip (sampled with 12 bit at 44 MS/s).

2) *Decoding DSSS*: In order to recover the original chip and decode the frame, the receiver seeks to obtain the correct phase change: to this end it relies on the values of the phase occurring in the middle of each chip, i.e., on the “plateau” between two chip boundaries. We mark the ideal sampling time with arrows in Fig. 4. After sampling, the receiver assigns the chip values according to the phase variation, a phase change indicates a 1, no phase change indicates a 0. The receiver decides the sampling timing at the beginning of a frame, when decoding the preamble: it then keeps the same timing for decoding the rest of the frame.

3) *Chip Alignment and Jamming*: In DBPSK, the amplitude does not carry any information and it could therefore be expected to be constant over time. Nevertheless, our measurements show that the amplitude of real IEEE 802.11 frames change over time; Fig. 4 shows two examples for the popular Broadcom chipset in the WRT54G AP (4b) and for an Atheros chipset (4c). All the plots in Fig. 4 uses the same time scale and chip boundaries: this shows that the amplitude peaks in the middle of a chip, i.e., at the ideal sampling point of the phase transitions. The observation of this amplitude peak is significant to understand the probability of jamming success.

Friendly jamming uses a reactive jamming approach (cf. Section III). Therefore, when the interference of the friendly jammer starts, the receiver has already processed the preamble of the attack frame and fixed the sampling timing. As mentioned before, the receiver will keep the sampling timing to decode the signal, which consists of the superposition of the jamming and attack frame. A higher amplitude of the jamming signal occurring where the receiver is measuring the phase transition implies a higher chance that the receiver makes a wrong phase decision, which would lead to a wrong chip interpretation. Several such wrong chip interpretations are required to happen consecutively for successful jamming because of the eleven-fold redundancy of DSSS (spreading gain) [10]. For friendly jammers based on customer-grade IEEE 802.11 hardware, the jamming signals are valid IEEE 802.11 frames with a high amplitude in the middle of each jamming chip (and a low amplitude at the chip boundaries). If the interfering signals are aligned at the chip level, their amplitudes peak in the same points, exactly at that point when the receiver samples the phase of the attack frame. Therefore, the case of exact chip alignment yields the most effective jamming. On the contrary, when the two frames are off by 1/2 of a chip, jamming has the least success probability (dashed line in Fig. 4b).

4) *Chip Alignment Is Random*: Our measurement analysis of IEEE 802.11 I/Q traces indicates that consecutive frames transmitted by a specific node display a random alignment at the chip level. There are three independent reasons for this observation. First, the IEEE 802.11 standard [14] allows a 25 ppm tolerance of the chip clock frequency to avoid the circuitry for compensating temperature variations: as the clock skew differs between any two devices, the attacker and the friendly jammer will see a drifting chip alignment over time even if they start with a perfect alignment. Second, because of the reactive jamming technique jammers need to switch from listen mode to transmit mode: this switch depends on the device states and therefore adds a random delay. Third, the propagation delay depends on the (uncorrelated) positions

Fig. 5. The lowest gain (P_{Jam}/P_{Att} cf. Fig. 2b) at which jamming success is possible has a periodic dependency on the chip alignment. (a) Trace-driven simulation results show that alignments which are a multiple of the chip length ($\approx .09[\mu s]$) are particularly robust against jamming (i.e., need a high jamming gain). (b) The Fourier transform of the simulation data clearly shows the component of the periodic dependency.

of the attacker and jammer and changes over time in a dynamic environment.

5) *Trace-Based Simulation Results*: We have implemented a software receiver running on I/Q samples of real IEEE 802.11 data frames. Our receiver implements the DSSS decoding steps and actually outperforms various hardware receivers, i.e., our software receiver can decode frames for which respective hardware receivers have not sent an acknowledgment. To simulate jamming, we superimpose a jamming IEEE 802.11 1Mb/s frame at different chip alignments. We increase the peak power of the jamming frame until the attack frame cannot be decoded anymore.

This gives us the lowest gain in the jamming peak power with respect to the attacker signal peak power at which jamming is successful (in the following 0dB denotes peak powers are the same). We repeat this experiment for 27 frames and show the results in Fig. 5a which reports, for every alignment, the average gains (together with minimum and maximum) required for successful jamming. The figure shows a strong periodic dependency of the required jamming gain on the chip alignment. As predicted, this chip alignment corresponds to 1/2 of the length of a DSSS chip ($\approx .09[\mu s]$). A Fourier transform visualizes this dependency in Fig. 5b.

All simulations results reported so far were based on traces collected with a real-time spectrum analyzer. To verify that the chip alignment dependency is due to the non constant chip amplitudes (cf. Fig. 4b), we have implemented a synthetic trace generator. This trace generator encodes DSSS frames with the same payload as the traces considered so far. We found that the effect of the chip alignment happens only when the signal amplitude peaks in the middle of a chip as observed in the real traces. We reproduce an ideal version of the peaking amplitude behavior and show the results in Fig. 6. The effect on the chip alignment is very similar to the results for the trace-based simulations (yet, more pronounced). In fact, the Fourier transform analysis shows that the main components of the synthetic data match the trace-based simulations.

6) *Conclusions for Friendly Jamming*: The attacker can gain an advantage over a single friendly jammer by using the robust DSSS encoding. Due to the chip alignment effect, this encoding has a significant chance of not being jammed by the friendly jammer. This particularly applies to scenarios in which both attacker and friendly jammer are power constrained to similar magnitudes (e.g., use customer-grade hardware). This is the case in our experiments in Section VI.

The most practical solution to this challenge is to deploy more friendly jammers. Thanks to their independent chip

Fig. 6. Repetition of the experiment in Fig. 5 but with synthetic trace data. (a) Periodic behavior of jamming success as a function of the chip alignment occurs more pronounced but very similar. (b) Corresponding Fourier transform.

alignment, additional jammers would make the chance of a 1/2 chip alignment arbitrarily small. We corroborated this conjecture in additional simulations (not shown).

B. Analysis of OFDM

In contrast to DSSS, OFDM builds the transmitted 20MHz spectrum by shaping it rather than relying on spreading. To this end the transmitter sets each of the 52 orthogonal carriers separately before invoking an Inverse Fast Fourier Transform (IFFT) and getting a *signal symbol* in the temporal domain. Furthermore, to add resiliency to frequency-selective multipath channels, the transmitter prefixes each symbol with a repetition of its own end to create a *Guard Interval* (GI). This section investigates the alignment effect of this prefix (called *Cyclic Prefix* (CP)) based encoding. We call this the *jamming symbol alignment*, in analogy to the *chip alignment* of DSSS. This section assumes prior knowledge about OFDM, which can be found in [6].

1) *Encoding OFDM*: The encoding process defined for the ERP-OFDM-PHY starts by dividing the frame into fixed size blocks. Each block is then enriched with forward error correction (FEC) information and the resulting data is distributed over the 52 spectrum carriers according to specific constellations, e.g., $\{-1, +1\}$ for the basic BPSK scheme for 6Mbit/s. After padding the 52 coefficients with zeros, the encoder runs the IFFT and obtains 64 complex values that it finally prefixes with a CP of 16 samples. For more detailed explanations (e.g., different datarates), we refer to chapter 18 of [14]. The resulting 80 temporal I/Q samples transmitted at 20MS/sec lead to a symbol duration of 4μs.

2) *Decoding OFDM*: Key for a successful reception of a frame is the recovery of the original spectrum associated to each symbol. To this end the receiver starts by capturing from each symbol a window of 64 I/Q samples. Fig. 7a shows two examples of such windows. Thanks to the CP, any two windows are equivalent as each one is the cyclic shift of the other: after applying the FFT they differ only for a phase factor that is linear in frequency. The receiver can hence eliminate this factor by knowing the symbol boundary timing, which is learned during the processing of the PCLP preamble [22]. The reception of the physical header allows also to correct the Carrier Frequency Offset [36], and to estimate the channel response: the latter is necessary to equalize the recovered spectrum and mitigate the non-flat channel fading before deciding which value from the corresponding constellation, e.g., $\{-1, +1\}$, was carried by every carrier. Finally, FEC information is used to decode the original data associated to

Fig. 7. (a) Symbol alignments of jam signals. (b) Minimum power peak gain for successful jamming w/ and w/o FEC for synthetic traces and no channel fading. (c) Two cases of minimum power peak gain for successful jamming with real traces.

each symbol before putting all pieces together in the recovered frame.

3) *Symbol Alignment and Jamming*: As shown in Fig. 7a, the interaction between the attacker and the jamming frames changes according to their symbol alignment. Specifically, this gives rise to one of either two cases depending on how many jamming symbols fit in the decoding window when processing the attack's preamble. First, when the decoding window covers a single jamming symbol (Jamming frame Jam_1 in Fig. 7a), then the result of the FFT operation on the two overlapping symbols leads to the sum of the corresponding spectra. Assuming a perfectly flat channel and equal constellations $\{-1, +1\}$, at the output of this phase we have 52 values that are the sum of constellation symbols: the constellation of the jamming is however rotated of the relative delay. If the intensity of the jamming signal is below that of the attack frame, then perfect decoding of the latter is still possible, independently of the rotation. Otherwise jamming might induce decoding errors. Second, two consecutive jamming symbols can also partially fit inside each decoding window (Jamming frame Jam_2 in Fig. 7a). In this case, after the FFT operation the spectrum of the attack frame is corrupted with coefficients that do not belong to any constellation. The FFT applied to a window of samples containing the tail and the head of two consecutive jamming symbols lead to random data, not shaped according to any constellation. Then, some of the 52 values will not be decodable even if the intensity of the jamming signal is lower. Decoding issues will increase with more complex constellations and in the presence of non flat-fading channel: in this case the equalization restores the attacker carriers inside their relative constellation while it further damages those of the jamming signal increasing the probability of non correct decoding.

4) *Simulation Results*: We start by considering synthetic traces that we generate and decode using the gr-ieee802-11 framework [4]. We simulate jamming by adding to a synthetic attack frame I/Q samples belonging to another (jamming) frame. Fig. 7b shows, for increasing symbol alignment (multiple of $4\mu s$ means perfect alignment), the gain in the peak power of the jamming signal with respect to peak power of the attacker that is required for successful jamming. In the top graph we disable FEC to better show the effect of the

alignment and the periodicity over 10 symbols: there are regions long as the cyclic prefix which require high gain ($0dB$) separated by regions requiring less, due to the interaction of multiple jamming symbols inside the attacker decoding window. When FEC is enabled the periodicity is no more evident: anyways, the spectral analysis in both cases reveals a strong component centered at $f = 250KHz$ that correspond to the $4\mu s$ periodicity together with its harmonics. We also gather from the Fig. that FEC introduces better resiliency towards jamming in the longer regions: the required gain for successful jamming increases and is almost flat with the delay (very close to $0dB$).

We then repeat the same analysis with I/Q traces of real signals transmitted by Broadcom devices and captured with an USRP-N210: before using the traces we normalize all signals to the same energy. We simulate jamming by adding to an attack frame transmitted by one device a jamming signals transmitted by another one: we use data from different devices on purpose to avoid artifacts due to the same transmission clock. To attempt decoding we still use gr-ieee802-11. We report in Fig. 7c the gain required for successful jamming with two different jam signals (differently than with DSSS we cannot meaningfully report here the average plus confidence interval so we report results separately for set1 and set2): periodicity is even less evident with real traces. We also notice that with respect to the simulation of the synthetic traces with FEC, the minimum gain for successful jamming diminishes thanks to the effect of the channel equalization on the jamming samples (almost flat around $-2dB$). Then, we repeat the spectral analysis by averaging the behavior of different jamming signals captured for set 1 and set 2: we see that each jamming set leads to a specific behavior but both have a strong component at $f = 250KHz$ (we also see a few harmonics).

5) *Conclusions for Friendly Jamming*: When considering different alignments at the symbol level, our synthetic traces show the existence of robust OFDM symbol alignments in theory. Such alignments could mean an advantage for the attacker similar to the robust chip alignment cases for DSSS. However, our simulations with real traces show that this effect is less significant in practice. These observations substantiate the common knowledge that the IEEE 802.11ag OFDM

Fig. 8. Hit ratio for (a) different attack frame MCSs and for (b) a 1 Mb/s DSSS vs. a 6 Mb/s OFDM jam frame.

encodings are less robust against jamming attacks, which is advantageous for friendly jamming.

We also validated that 6 Mb/s is the most jamming-resilient OFDM encoding: Fig. 8a compares the susceptibility of four different OFDM data rates between 6 Mb/s and 18 Mb/s. Indeed, a jammer gain of -12 dB is enough to reach 100% hit ratio for 18 Mb/s but while for 6 Mb/s we need to have similar peak powers (0 dB).

As a final simulation experiment, we also consider the effectiveness of two different jamming signals. Fig. 8b compares whether jamming a 6 Mb/s frame with DSSS frames or with OFDM frames is more effective. As a fair comparison, we normalized both jamming signals to the same energy level. Nevertheless, the 1 Mb/s DSSS jamming frame has a slight energy advantage. We leave the further investigation of this effect open to further research and fix 1 Mb/s jamming signals for the following measurement studies.

C. Analysis of Carrier Phase Offset

We next study the impact of different carrier phase offsets (CPO) between the jamming signal and the attack frame. Without spreading or other forward error correction, CPO is known to require an increased power of the attack frame [29]. We find that this is not the case for the IEEE 802.11 DSSS and OFDM modulations. Specifically, we again use our software decoder with a synthetic trace, in which the CPO uniformly takes one of 32 values in the interval $\theta \in [0, 2\pi[$. We multiply the jamming signal for the corresponding CPO coefficient $\exp[j\theta]$ and again consider different chip and symbol alignments for DSSS and OFDM, respectively. Fig. 9 shows the power gain required for successful jamming: the thick lines in the three Figures corresponds to perfect phase alignment and, in fact, they are exactly the same gains as shown Figs. 6a and 7b. The thin lines in Fig. 9 correspond to other values of the CPO. We find that the CPO does not introduce major changes on the jamming behavior, i.e., the same periodic fluctuations are still visible for all the curves. We also calculated the overall change in the power required for successful jamming by averaging over the entire time and phase delay offsets. We find that in all the three considered cases the gain changes are negligible. This quantitatively confirms that the dynamics involved by complex modulations using either spreading or frequency multiplexing do not require explicit control of the phase of the jamming signal for modifying the information decoded by a receiver.

V. IMPLEMENTATION

The friendly jamming firmware was developed in three steps. First, a basic prototype for IEEE 802.11g (on the Broadcom 4138 chipset) showed the idea's feasibility;

second, we optimized it with extensive lab experiments; and third, we ported the friendly jamming project to a modern IEEE 802.11n platform to see how fast future hardware or standard changes can be incorporated.

Access to the network interface card's firmware is facilitated by the OpenFWWF open source project [12] (also see [31]). The basic OpenFWWF firmware can receive and send IEEE 802.11-compliant frames. As outlined in Section III-C, significant changes were necessary.

Our first prototype was validated using several micro-benchmarks, e.g., traces by real-time spectrum analyzers as in Fig. 2 and open-space experiments (Tbl. II). The second step included a variety of optimizations to shorten the jammer's reaction time and increase jamming performance: we implemented a prefetching mechanism that proactively pulled the jammer's (re)configuration from the system memory and stored it in fast-access registers on the NIC. We then improved the management of the transmission power: we ran comprehensive tests on corresponding hardware registers before settling on a final configuration. As motivated by the results in Section IV, we used 1 Mb/s DSSS jamming signals. As the jamming frame's length had negligible impact, we set it to the minimum allowed by the hardware (ten bytes) that led to a jam signal's duration of $272\mu s$ ($192\mu s$ for the PLCP header, other $80\mu s$ for the rest). We also measured the jammer's reaction time after matching a target MAC address: by comparing original and jammed frames' content, we verified this delay to be two bytes at 1 Mb/s, i.e., $16\mu s$.

The third step involved the porting of the software to a modern chipset (the Broadcom 43224) that supports IEEE 802.11n (2×2 MIMO). Because the selected chipset is supported by different kernel driver and firmware, we developed the required software from scratch making the new system back-compatible with the old configuration files.³

This concludes the discussion on the implementations. The basic IEEE 802.11g implementation (result of the second step) is used for the experiments in Section VI. The IEEE 802.11n implementation (result of the third step) is used for the experiments in Section VII.

VI. RESULTS: IEEE 802.11bg NETWORK

This section describes results from a measurement study of our friendly jammer implementation in a complex radio communication environment.

A. Deployment and Measurement Data

We selected a university office floor as a particularly heterogeneous and dynamic environment. The observed background traffic is generated by a dynamic user base of faculty members and students; it comprises a variety of traffic sources such as typical back-office activities, research-oriented applications, and multimedia applications. Besides the university's main campus network, there are various smaller APs with many overlapping Basic Service Sets, so that we are able to observe different network types in parallel.

³Our preliminary friendly jammer implementation is available on <http://www.ing.unibs.it/~openfwwf/friendlyjammer/>.

Fig. 9. Minimum gain required for jamming synthetic traces successfully. (a) DSSS, (b) OFDM without FEC (attack and jam frame 6Mb/s) and (c) OFDM with FEC (attack and jam frame 6Mb/s). The jamming power is less impacted by the CPO $\theta \in [0, 2\pi]$ than by the chip or symbol alignment. Also note that as in Fig. 7, FEC again destroys most of this effect.

Fig. 10. Experimental setup: the attacker and 13 jammers use WRT54GL aps and are located on a university floor of about $40 \times 85m$. Traffic is recorded with ALIX.2D2s on three positions, m_{close} , m_{middle} , and m_{far} .

We deployed one attack traffic source, 13 friendly jammers, and three monitors on the university floor, see Fig. 10. The locations of all devices were constrained by availability but several pre-runs showed the same qualitative outcomes for various locations. In these pre-runs, we experimented with different attack traffic rates, inter-frame spacing (e.g., SIFS) and backoff mechanisms. We found that the jamming performance is unaffected for all configurations tested (e.g., full saturation of the channel with attack traffic) because the jammers emit a very brief jamming signal, and go back to listening on the channel before an attack frame ends (cf. Fig. 2b). However, as discussed in Section V, the jamming performance is affected if the packet payload is less than 42 bytes. For the remainder of this section, we report on results for standard-compliant inter-frame spacing, a constant attack rate of 500 Kb/s, and 50 bytes packet payload. As channel 6 is the busiest of the university network, we chose it for studying the jamming system in order to see the negative impact of friendly jamming. The attack traffic used a round-robin selection of all IEEE 802.11b/g modulations without retransmissions, so that the friendly jammer could not lock onto any modulation. This way the impact of different modulations on the jamming performance could be fairly judged.

The 13 friendly jammers used the configuration from Section V and blocked traffic based on the attacker's MAC address. The attack frames are received by three monitors (m_{close} , m_{middle} , m_{far} , see Fig. 10). To accurately track sent, received, and valid-checksum statistics for the attack traffic, we implemented low-level counters (on the NIC of the monitors) that were not impacted by the operating system. The traffic that we captured from legitimate users was recorded as a full pcap trace and streamed to a central server over dedicated Ethernet lines. In summary, our setup was carefully crafted to obtain reliable attack traffic and accurate measurements.

All results in this section come from 24h-experiments over a three-week-period to observe different conditions on the wireless channel. Time was divided into short experiments

Fig. 11. (a) A single jammer's hit ratio is fully dependent on its position. (b) Using an increasing number of jammers (0 corresponds to packet loss without jamming) brings the hit ratio asymptotically close to 100%.

TABLE II
OPEN-SPACE HIT RATIO

dist	1 Jammer	2 Jammers	3 Jammers
5.4m	98.93% (± 0.19)	100.0% (± 0.00)	99.98% (± 0.02)
6.7m	98.57% (± 0.85)	99.95% (± 0.04)	100.0% (± 0.01)
8.1m	77.03% (± 12.3)	98.37% (± 0.29)	99.43% (± 0.14)

each of about 150s, where in each experiment a random number and selection of jammers were activated uniformly from all possible combinations. In total, this setup generated 13660 experiments with about 350GB of measurement data comprising almost 370 million target frames and 490 million legitimate-traffic frames. Throughout this section, we use 97.5% confidence intervals based on the t-distribution.

B. The Jamming Performance

We measure jamming performance via the hit ratio.

1) *Position-Dependence*: In open-space testbed experiments (e.g., [3], [42], [43]) the hit ratio monotonically decreases with the jammer's distance from the receiver. We also observe a monotonic behavior for our jammers, and we thus summarize only part of the measurements in Table II. In contrast to this, behavior in our in-door scenario is quite erratic and there are weak positions even close to a victim receiver (cf. Fig. 11a). As expected from the analysis in Section IV, the impact of weak positions is more pronounced for robust MCSs and small numbers of jammers. For example, the hit ratio when averaging over random placements of three jammers decreases by up to 30% for the robust MCS 1Mb/s; however, for seven jammers this issue mostly disappears (less than 5% difference between MCSs) as it is more likely that a good position is included in the active jammer set.

2) *Number of Jammers*: An increasing number of jammers increases the hit ratio due to improving the detection and emitted power of the jammers (Fig. 11b). Note, however, that one may grossly underestimate the number of jammers necessary

Fig. 12. We complement the hit ratio with an additional metric called *unpredictability of jamming misses*, which is explained in the legend on the left (if the attacker observed that two (#5, #10) out of twelve attack frames (#0-#11) were not jammed, can it predict its next attack opportunity?). (a) Increasing the number of jammers from 3 to 7 to 13 thins and spreads out the occurrence of not jammed attack frames. (b) Quantile-quantile plots show good correspondence with the geometric distribution, which indicates unpredictability for the attacker, for seven jammers but less so for three jammers.

when only considering open-space experiments (compare to Table of Fig. II). Due to the weak positions described above, a single jammer's hit ratio lies between 60 – 90% (averaged over all positions). At least three jammers are required to sustain an average hit ratio above 90%, seven jammers for 99%, and the hit ratio graph's slope only flattens out for more than eleven jammers. This appears to the case across receivers at different proximities, which yielded distinct packets loss rates without any jammers (19 – 48% at 12 – 75m, cf. Fig. 10).

3) *Unpredictability of Jamming Misses*: The hit ratio defines an aggregated view of the jamming performance, which might hide short-term glitches of the friendly jammers. Fig. 12a shows that this is not the case: not-jammed attack frames are distributed uniformly over time. Specifically, we consider the series of the sequence numbers of the attack traffic and we find that increasing the number of jammers from 3 to 7 to 13 makes this distribution increasingly uniform. Such a uniform distribution of not-jammed frames over time means that no statistical knowledge can be obtained for the attacker from observing past not-jammed attack frames, which makes their exploitation hard. We formalize this observation by considering the inter-miss distribution on the sequence number space. If this distribution is close to the geometric distribution, then jamming misses possess the memoryless property, i.e., are unpredictable even when an attacker is able to observe the friendly jamming system over a long time. Indeed, the quantile-quantile plots in Fig. 12b show that, qualitatively, the measured inter-miss distribution converges to the geometric distribution for seven jammers. However, three active jammers seem insufficient to enforce this property.

C. The Cost of Friendly Jamming

In this section, we investigate the side effects of friendly jamming on *legitimate* traffic in the network. Recall that legitimate frames have not been targeted by the frame-selective jammers; instead, any adverse effect on legitimate traffic is collateral – the cost of friendly jamming.

In order to assess this cost, we distinguish legitimate frames from attack frames (and jamming signals) in our measurement data. This required some attention because fragments resulting from the attacker's frames and the jammers' frames could not easily be distinguished from actual legitimate frames due to high bit error levels. We used an automated procedure, which we verified selectively down to the bit level of individual frames: we first discovered valid legitimate transmitter MAC

Fig. 13. Packet loss of legitimate traffic subdivided by RSSI levels. (a) The loss at monitor m_{close} is stable. (b) monitor m_{middle} observed an increase in the packet loss for senders with medium RSSI, which we attribute to a slight power amplification effect. (c) the packet loss of traffic from low RSSI senders to m_{far} doubles with an increasing jammer count.

addresses, and then matched frames with bit errors to the legitimate stations using the smallest Hamming distance.

1) *Loss of Legitimate Traffic and Power Amplification*: Friendly jamming can negatively affect legitimate traffic as suggested by the power amplification phenomenon discussed in Section III. This, however, happens only for certain constellations of a legitimate station to the attacker and to the jammers. We quantify this effect by exploiting the fact that our monitors m_{close} , m_{middle} , and m_{far} cover a large observation area, 85m in length (cf. the map in Fig. 10). For each of the monitors, we categorize the observed senders into three classes based on the received signal strength: low RSSI ($< -65dB$), medium RSSI ($-65dB < RSSI < -55dB$), and high RSSI ($> -55dB$). The packet loss for each sender in these classes is measured with the monitor as a receiver, based on each frame's checksum. Fig. 13 reveals that the packet loss increases with the number of jammers only for two classes (and is stable for the other classes). The most badly affected class is the traffic with low RSSI as observed by m_{far} (Fig. 13c): the packet loss doubles from 20.2% to 40.4%. The low-RSSI senders are located at positions to the right on the floor (Fig. 10), where they do not receive the attack messages but may suffer interference from the jammers in the center (i.e. the setting described in Fig. 3b). Accordingly, we attribute the observed increase of packet loss to the power amplification phenomenon, for which the likelihood increases with an increasing number of active jammers. The second affected class are m_{middle} 's medium-RSSI senders (Fig. 13b): the packet loss increases from 13.7% to 22.6%. We also explain this observation with the jammer's power amplification; however, the effect is less pronounced than before since there are fewer candidate jammers. It is also interesting to note that the low-RSSI senders at m_{middle} are not similarly affected by this phenomenon – we attribute this to erratic propagation effects of the indoor topology (cf. Fig. 10);

Fig. 14. (a) The RSSI (at m_{close}) surges when enabling the first jammer, indicating an increased noise level. (b) The adapter rate of background traffic becomes more volatile as soon as the jammers are enabled – indicating higher activity of the rate selection algorithm. (c) The mean number of (re)association requests becomes more volatile but remains approximately stable.

for example, there is a fire door to the left of this monitor that significantly alters signal propagation paths.

Summing up, we found that friendly jamming can substantially impair the channel of legitimate transmissions due to the power amplification effect. Future studies of friendly jamming should thus seriously consider a cost perspective besides traditional metrics of performance and security.

2) *Interaction With Rate Adaptation*: The packet loss of stations, which are not affected by the power amplification effect, does generally not depend on the number of active jammers. Some configurations (cf. Fig. 13) seem to have a slight impact; but we attribute this to increased variability when more jammers are active. This increased variability becomes evidence when considering the RSSI recorded at the close-by monitor: the RSSI increases significantly (Fig. 14a). As the RSSI indicates an increased level of channel noise, rate adaption algorithms can be expected to select a more robust modulation. Indeed, Fig. 14b shows that the adapter rate shows greater variance and tends towards more robust MCSs when the jammers are enabled.

3) *Effect on Associations*: As a final aspect of the cost of jamming, we consider the number of (re)associations requests under different numbers of jammers. Fig. 14c presents the change in this number with respect to the configurations without jammers. While the close-by stations seem to be affected by an increase of up to 40%, the farther-away stations exhibit approximately stable behavior. We also measured (re)associations of individual stations, and find that the number of significantly affected station is a small.

In summary, we again emphasize the trade-off between jamming performance and the cost of jamming: enabling a larger number of jammers is necessary to boost the hit ratio but also increases the collateral damage to legitimate traffic. While not unexpected, we obtained detailed quantitative statements on the cost of jamming in a real-world 802.11 network. These results were obtained under the assumption of random jammer placement. However, a friendly jamming system might be able to overcome this limitation as investigated next.

D. Jammer Collaboration Schemes

In this section, we show the potential of jammer collaboration, by which we mean that jammers are selected according to some deterministic rule instead of randomly as above. This promises to improve the hit ratio while remaining marginally invasive, based on the previous section’s finding that the

jammer’s position is critical to both the performance and cost of friendly jamming (cf. Fig. 15).

Our first scheme selects the jammer closest to the attacker (e.g., realized by trilateration), which is denoted *ACL*. Our second scheme selects the jammer closest to the victim, which is denoted *VCL*. This can be possible, if the jamming system can observe bidirectional communication attempts between the attacker and its victim. Compared to *ACL*, we expect *VCL* to improve the hit ratio because the friendly jammers would be closer to the receiver contributing to an higher energy jamming signal. Our third scheme is based on the idea to select jammers based on their reception quality. Specifically, the scheme *RCL* selects jammers that received the highest number of attack frames. While this heuristic only considers the channel from the attacker to the jammer, one may speculate that it also corresponds to better jamming positions in general.

We compare our schemes to random jammer selection (*RND*) and to the optimal configuration (*OPT*), which is found by an exhaustive search through all possible jammer selections for each repetition of the experiment. We use two metrics to compare the schemes: the hit ratio for the attack traffic (Fig. 15a) and the packet loss of legitimate traffic (Fig. 15b). We normalize the packet loss with regard to the *RND* scheme to show how our schemes improve upon the power amplification effect. Specifically, the legitimate packet loss is measured as the normalized packet loss where normalization is done to the *RND* configuration that yields the same hit ratio.⁴ The simple *ACL* scheme significantly increases the jamming performance and reduces the legitimate packet loss (power amplification effect). We attribute this improvement to the proximity of the receiver and the sender and to the fact that far-away jammers are efficiently excluded. The extension, the *VCL* scheme, yields a slightly better hit ratio than *ACL* and can reduce the packet loss even further. However, we remark that this scheme is less practical than *ACL*. The third practical scheme *RCL* can outperform *VCL* slightly, but with overlapping confidence intervals. While all collaboration schemes improve significantly over random jammer selection, there remains room for improvement as the hit ratios are 2.5 – 4% below the *OPT* case.

We also studied the schemes for seven jammers. In that case, the mean hit ratios are closer together: 98.95% for *ACL*, 99.41% for *VCL*, 99.62% for *RCL*, and 99.95% for *OPT*.

VII. RESULTS: IEEE 802.11n NETWORK

In this section, we extend our measurement scope to more broadly account for the Wi-Fi feature set occurring in practice. Specifically, we target the new features introduced with the IEEE 802.11n standard (e.g., spatial substreams), the impact of different receiver types, and measurements on 5 GHz channels.

All receivers were equipped with two antennas. We tested various different antenna types (including those specifically built for the respective frequency bands), and selected a type of dual-band standard-dipole antennas that showed the best overall performance for 2.4 GHz (channel 11) and 5.2 GHz

⁴Normalized packet loss of scheme x with hit ratio p : (packet loss of x) / (packet loss of *RND* with that number of jammers that yield a p hit ratio).

Fig. 15. When the position of the attacker and the victim is not known a priori, random jammer selection (RND) with several candidate positions is the only choice. However, a *collaboration scheme* can improve the jamming performance by selecting appropriate positions based on additional measurements: the attacker's position (ACL), the victim's position (VCL), or the reception ratio of individual friendly jammers (RCL). These three collaboration schemes lead to near-optimal improvements of the hit ratio (a) and reduce legitimate packet loss (due to the power-amplification effect) by 20-30% (b).

Fig. 16. (a) The traffic source and nine friendly jammers were located in a similar layout as in the previous experiment (cf. Fig. 10). (b) The frequency band has only a small impact on the jamming performance; the main difference can be seen in the packet loss without jamming due to reduced interference on the 5 GHz band. Also note that the new chipsets used here yield a higher hit ratio for our friendly jammers (cf. Fig. 11b).

(channel 44). To simplify the comparison to the previous IEEE 802.11bg setup, we deployed the new hardware on the same floor and on similar positions as before, cf. Fig. 16a. Unless noted otherwise, all measurements are due to the monitor m_{bcm} (Fig. 16a) which uses the same hardware configuration as the friendly jammers.

1) *Role of Chipset Implementation and Frequency Band:* Fig. 16b shows the hit ratio for the new friendly jammer implementation. In comparison to previous results, the hit ratio increases much faster with the number of jammers. The hit ratio curve flattens out for about six jammers. We attribute this to a measurably improved reception ratio of the more modern chipset used in the jammers. This results in fewer missed jamming opportunities by individual friendly jammers and thus increases the hit ratio at lower jammer counts. Another important factor is that for the 5 GHz band the IEEE 802.11an standard restricts the attacker's transmissions to OFDM modulations, which previously turned out to be more susceptible to jamming. To ensure comparable results, we also enforced the same restriction on the 2.4 GHz results reported in Fig. 16b. Overall, a significant difference in hit ratio between the two frequency bands only appears without jammers, for which the 2.4 GHz has a higher packet loss, and for 1 – 2 jammers, for which the hit ratio increases less steeply compared to the 5 GHz results. We attribute this difference mainly to a generally reduced interference on the 5 GHz band (increased channel spacing and fewer active stations). We conclude that the hit ratio analysis of the previous section essentially holds for all the IEEE 802.11abgn frequency bands. Nevertheless, comparing the previous friendly jammer implementation's results from Fig. 11b on 2.4 GHz to the new results in Fig. 16b reveals that specific wireless chipset technologies can have some impact.

Fig. 17. (a) In our setup, jammers and attackers are 2x2 systems, which can be either configured to exploit antenna diversity and one IEEE 802.11n spatial stream (top), or in a multiplexing configuration with two IEEE 802.11n spatial streams (bottom). (b)-(c) The box plots show that for two types of receivers (m_{bcm} and m_{ath} respectively) both one (top) and two (bottom) spatial streams are jammed successfully.

2) *Spatial Multiplexing and Receiver Type:* The IEEE 802.11n standard exploits MIMO antenna systems to provide independent spatial paths between the transmitter and receiver so that different streams of information can be sent in parallel along these spatial paths (roughly speaking, for an introduction see [13]). In contrast, IEEE 802.11abg exploits only the antennas' diversity, e.g., selecting the better one from a pair of antennas (cf. Fig. 17a).

We are interested in whether an attacker's transmission using the 11n spatial multiplexing feature is more or less susceptible to interference from our friendly jammers. Fig. 17b shows a bar plot of the hit ratio for 0, 3, 6, or 9 jammers with and without 11n spatial multiplexing. For "receiver 1" (identical to the friendly jammer's chipsets) there are only minor differences in the hit ratio behavior. The packet loss (without jammers) for the two spatial streams configuration is slightly higher, and the hit ratio for 3 jammers is slightly lower. This can be explained by a marginally decreased reception ratio of both the receiver and the friendly jammers.

We also show results for a second receiver in Fig. 17c. This receiver is based on an Atheros AR5418 chipset and is denoted as m_{ath} (close to the first receiver in the map, Fig. 16a). The hit ratio of this receiver for 3 jammers is somewhat lower, but comparable for 6 and 9 jammers. Interestingly, the packet loss (0 jammers) is higher for one spatial stream compared to two spatial streams. However, the hit ratio with activated jammers is similar for both receiver configurations.

3) *The Power Amplification Effect:* We revisit the phenomenon that was found in our measurements and which leads to increasing costs when more jammers are enabled. Section VI-C established that this effect actually happens in a real-world network. However, this was demonstrated indirectly

Fig. 18. (a) Example position of two legitimate stations that endure a severe power amplification effect when the right station communicates to the left one. (b) The corresponding decrease in throughput. (c) The corresponding increase in packet loss.

using RSSI to differentiate different stations' positions. In this experiment, we now moved two stations until we found a position, which experienced a significant power amplification phenomenon (cf. Fig. 18a). In fact, we were easily able to find such a constellation in our environment, as predicted by Fig. 3b. By taking measurements directly from these stations, we find that the throughput significantly decreases from over 60Mb/s to around 40Mb/s (Fig. 18b) and that the packet loss increases from below 30% to around 45% (Fig. 18c). This finding confirms the power amplification phenomenon and also shows an almost linear increase in the cost on legitimate transmissions when the number of friendly jammers is increased (cf. Fig. 13).

VIII. CONCLUSION

System proposals for friendly jamming applications have received substantial research interest in recent years. This work seeks to bridge the gap from small and static test bed evaluation to a more realistic scenario by investigating friendly jamming at work in a real-world 802.11 network over a three-weeks period. We find three main insights: 1) a rather high number of jammers is required to achieve a high hit ratio for different modulation schemes, 2) the identification and quantification of the power amplification phenomenon which represents a shadow cost of friendly jamming, and 3) the effectiveness of potential jammer collaboration schemes in achieving a good trade-off between jamming performance and cost. We extended our analysis to measurements on 5GHz and IEEE 802.11n and found that the key insight of a trade-off between jamming effectiveness and invasiveness carry over.

Besides the rather black-box oriented measurement results on friendly jamming, we also used simulations of the reception process of superimposed signals to investigate the question why perfect jamming is not always feasible. The analysis for DSSS revealed a strong dependence of the jamming success on the chip alignment between attack frame and jam frame(s). This insight opens up another interesting opportunity for future research: can friendly jammers boost their hit ratios by sending jam signals that are synchronized with the attack signal? Another interesting direction would be to quantify the effect if the friendly jammers could control the carrier frequency offset between their jamming signals and the attack frame.

REFERENCES

- [1] D. L. Adamy, *Ew 101: A First Course in Electronic Warfare*, vol. 1. Norwood, MA, USA: Artech House, 2000.
- [2] E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa, "On the performance of IEEE 802.11 under jamming," in *Proc. IEEE INFOCOM*, Apr. 2008, pp. 1265–1273.
- [3] D. S. Berger, F. Gringoli, N. Facchi, I. Martinovic, and J. Schmitt, "Gaining insight on friendly jamming in a real-world IEEE 802.11 network," in *Proc. ACM WiSec*, 2014, pp. 105–116.
- [4] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, "An IEEE 802.11a/g/p OFDM receiver for GNU radio," in *Proc. ACM SIGCOMM SRIIF*, Hong Kong, Aug. 2013, pp. 9–16.
- [5] J. Brown, I. E. Bagci, A. King, and U. Roedig, "Defend your home!: Jamming unsolicited messages in the smart home," in *Proc. ACM HotWiSec*, New York, NY, USA, 2013, pp. 1–6.
- [6] M. Debbah. (2004). *Short Introduction to OFDM*, accessed on Jun. 21, 2016. [Online]. Available: <http://goo.gl/Zoi2ID>
- [7] P. Dutta, S. Dawson-Haggerty, Y. Chen, C.-J. M. Liang, and A. Terzis, "Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless," in *Proc. ACM SenSys*, 2010, pp. 1–14.
- [8] European Telecommunications Standards Institute. (Aug. 2012). *EN 300 328 V1.8.1 Electromagnetic Compatibility and Radio Spectrum Matters (ERM); Wideband Transmission Systems; Data Transmission Equipment Operating in the 2.4 GHz ISM Band and Using Wide Band Modulation Techniques; Harmonized EN Covering the Essential Requirements of Article 3.2 of the R&TTE Directive*. [Online]. Available: <http://goo.gl/G8ksx8>
- [9] FCC. *Jamming Prohibition*, accessed on Feb. 13, 2014. [Online]. Available: <http://goo.gl/Z3Nvak>
- [10] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh, "Efficient network flooding and time synchronization with glossy," in *Proc. 10th Int. Conf. Inf. Process. Sensor Netw. (IPSN)*, Apr. 2011, pp. 73–84.
- [11] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu, "They can hear your heartbeats: Non-invasive security for implantable medical devices," in *Proc. ACM SIGCOMM*, 2011, pp. 2–13.
- [12] F. Gringoli and L. Nava. *OpenFWWF: Open Firmware for Wi-Fi Networks*, accessed on Jun. 21, 2016. [Online]. Available: <http://www.ing.unibs.it/>
- [13] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "802.11 with multiple antennas for dummies," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 1, pp. 19–25, Jan. 2010.
- [14] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Standard 802.11-2012, Mar. 2012.
- [15] M. L. Jorgensen, B. R. Yanakiev, G. E. Kirkelund, P. Popovski, H. Yomo, and T. Larsen, "Shout to secure: Physical-layer wireless security with known interference," in *Proc. IEEE GLOBECOM*, Nov. 2007, pp. 33–38.
- [16] Y. S. Kim and P. Tague, "Proximity-based wireless access control through considerate jamming," in *Proc. ACM SPME*, 2014, pp. 19–24.
- [17] Y. S. Kim, P. Tague, H. Lee, and H. Kim, "Carving secure Wi-Fi zones with defensive jamming," in *Proc. ACM ASIACCS*, 2012, pp. 53–54.
- [18] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. IEEE ICNP*, Oct. 2004, pp. 252–261.
- [19] J. Lee *et al.*, "An experimental study on the capture effect in 802.11a networks," in *Proc. ACM WinTECH*, New York, NY, USA, 2007, pp. 19–26.
- [20] Z. Li, W. Xu, R. Miller, and W. Trappe, "Securing wireless systems via lower layer enforcements," in *Proc. ACM WiSe*, 2006, pp. 33–42.
- [21] G. Lin and G. Noubir, "On link layer denial of service in data wireless LANs," *Wireless Commun. Mobile Comput.*, vol. 5, no. 3, pp. 273–284, May 2005.
- [22] C.-H. Liu, "On the design of OFDM signal detection algorithms for hardware implementation," in *Proc. IEEE GLOBECOM*, Dec. 2003, pp. 596–599.
- [23] I. Martinovic, N. Gollan, and J. B. Schmitt, "Firewalling wireless sensor networks: Security by wireless," in *Proc. IEEE LCN*, Oct. 2008, pp. 770–777.
- [24] I. Martinovic, P. Pichota, and J. B. Schmitt, "Jamming for good: A fresh approach to authentic communication in WSNs," in *Proc. ACM WiSec*, New York, NY, USA, Mar. 2009, pp. 161–168.
- [25] A. Mpitzopoulos, D. Gavalas, C. Konstantopoulos, and G. Pantziou, "A survey on jamming attacks and countermeasures in WSNs," *IEEE Commun. Surveys Tuts.*, vol. 11, no. 4, pp. 42–56, 4th Quart., 2009.
- [26] K. Pelechrinis, M. Iliofotou, and S. Krishnamurthy, "Denial of service attacks in wireless networks: The case of jammers," *IEEE Commun. Surveys Tuts.*, vol. 13, no. 2, pp. 245–257, May 2011.
- [27] K. Pelechrinis, C. Koufogiannakis, and S. V. Krishnamurthy, "Gaming the jammer: Is frequency hopping effective?" in *Proc. IEEE 7th Int. Symp. WiOPT*, Jun. 2009, pp. 1–10.
- [28] R. A. Poisel, *Modern Communications Jamming: Principles and Techniques*. Norwood, MA, USA: Artech House, 2011.

- [29] C. Pöpper, N. O. Tippenhauer, B. Danev, and S. Capkun, "Investigation of signal and message manipulations on the wireless channel," in *ESORICS*, Leuven, Belgium: Springer-Verlag, 2011, pp. 40–59.
- [30] D. R. Raymond and S. F. Midkiff, "Denial-of-service in wireless sensor networks: Attacks and defenses," *IEEE Pervasive Comput.*, vol. 7, no. 1, pp. 74–81, Jan./Mar. 2008.
- [31] P. Salvador, V. Mancuso, P. Serrano, F. Gringoli, and A. Banchs, "VoIPiggy: Analysis and implementation of a mechanism to boost capacity in IEEE 802.11 WLANs carrying VoIP traffic," *IEEE Trans. Mobile Comput.*, vol. 13, no. 7, pp. 1640–1652, Jul. 2014.
- [32] S. Sankararaman *et al.*, "Optimization schemes for protective jamming," *Mobile Netw. Appl.*, vol. 19, no. 1, pp. 45–60, Feb. 2014.
- [33] N. Santhapuri, S. Nelakuditi, and R. R. Choudhury, "On spatial reuse and capture in ad hoc networks," in *Proc. IEEE WCNC*, Mar./Apr. 2008, pp. 1628–1633.
- [34] A. Sheikholeslami, D. Goeckel, H. Pishro-Nik, and D. Towsley, "Physical layer security from inter-session interference in large wireless networks," in *Proc. IEEE INFOCOM*, Mar. 2012, pp. 1179–1187.
- [35] W. Shen, P. Ning, X. He, and H. Dai, "Ally friendly jamming: How to jam your enemy and maintain your own wireless connectivity at the same time," in *Proc. IEEE Symp. SP*, May 2013, pp. 174–188.
- [36] E. Sourour, H. El-Ghoroury, and D. McNeill, "Frequency offset estimation and correction in the IEEE 802.11a WLAN," in *Proc. IEEE VTC*, vol. 7, Sep. 2004, pp. 4923–4927.
- [37] N. O. Tippenhauer, L. Malisa, A. Ranganathan, and S. Capkun, "On limitations of friendly jamming for confidentiality," in *Proc. IEEE Symp. SP*, May 2013, pp. 160–173.
- [38] M. Vanhoef and F. Piessens, "Advanced Wi-Fi attacks using commodity hardware," in *Proc. ACM ACSAC*, 2014, pp. 256–265.
- [39] J. P. Vilela, P. C. Pinto, and J. Barros, "Position-based jamming for enhanced wireless secrecy," *IEEE Trans. Inf. Forensics Security*, vol. 6, no. 3, pp. 616–627, Sep. 2011.
- [40] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler, "Exploiting the capture effect for collision detection and recovery," in *Proc. IEEE EmNetS*, May 2005, pp. 45–52.
- [41] M. Wilhelm, V. Lenders, and J. B. Schmitt, "On the reception of concurrent transmissions in wireless sensor networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 12, pp. 6756–6767, Dec. 2014.
- [42] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "WiFire: A firewall for wireless networks," in *Proc. ACM SIGCOMM*, 2011, pp. 456–457.
- [43] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "Short paper: Reactive jamming in wireless networks: How realistic is the threat?" in *Proc. ACM WiSec*, 2011, pp. 47–52.
- [44] A. D. Wood and J. A. Stankovic, "Denial of service in sensor networks," *Computer*, vol. 35, no. 10, pp. 54–62, Oct. 2002.
- [45] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," in *Proc. ACM MobiHoc*, 2005, pp. 46–57.
- [46] X. Zhou and M. R. McKay, "Physical layer security with artificial noise: Secrecy capacity and optimal power allocation," in *Proc. IEEE ICSPCS*, Sep. 2009, pp. 1–5.

Francesco Gringoli received the Laurea degree in telecommunications in 1998, and the Ph.D. degree in information engineering in 2002. He has been an Assistant Professor of Telecommunications with the Department of Information Engineering, University of Brescia, Italy, since 2005.

Nicolò Facchi received the master's degree in information technology engineering from the University of Brescia, Italy, in 2012, and the Ph.D. degree in telecommunication engineering in 2016. He is currently a Post-Doctoral Researcher with the Department of Information Engineering, University of Brescia.

Ivan Martinovic received the Ph.D. degree from the Kaiserslautern University of Technology. Before coming to Oxford, he was a Post-Doctoral Researcher with the Security Research Laboratory, University of California, Berkeley, and the Secure Computing and Networking Centre, University of California at Irvine. He is currently an Associate Professor with the Department of Computer Science, University of Oxford.

Daniel S. Berger received the B.Sc. and M.Sc. degrees from the Distributed Computer Systems Laboratory, Kaiserslautern University of Technology, in 2011 and 2013, respectively, where he is currently pursuing the Ph.D. degree in computer science. He was with the German Cancer Research Center from 2008 to 2010.

Jens B. Schmitt received the Ph.D. degree from the Darmstadt University of Technology, Germany, in 2000. Since 2003, he has been heading the Distributed Computer Systems Laboratory. He is currently a Professor of computer science with the Kaiserslautern University of Technology, Germany. His research interests are broadly in performance and security aspects of networked and distributed systems.

Friendly Jamming on Access Points: Analysis and Real-World Measurements

Daniel S. Berger, *Student Member, IEEE*, Francesco Gringoli, *Member, IEEE*, Nicolò Facchi, *Member, IEEE*, Ivan Martinovic, *Member, IEEE*, and Jens B. Schmitt, *Member, IEEE*

Abstract—Frequency jamming is known as an efficient attack tool to disrupt wireless communication. This efficiency can also be exploited for the benefit of a network—an idea often referred to as friendly jamming. A prominent application case is the blocking of unauthenticated or malicious communication, such as injection attacks. In this paper, we propose access points as a natural place to implement friendly jamming functionality. We analyze this proposal using simulations, introduce an implementation on customer-grade access points, and report measurement results from the first real-world study of friendly jamming in an IEEE 802.11 campus network. We discover a fundamental tradeoff between the effectiveness of friendly jamming and the orthogonal aspect of having minimal side-effects to the campus network’s traffic. In particular, we observed what we call the power amplification phenomenon. This effect aggravates the known hidden station problem when the number of jammers increases. We also find evidence that the collaboration between jammers can enable friendly jamming, which is both effective and minimally invasive.

Index Terms—Friendly jamming, jamming for good, defensive jamming, reactive jamming, IEEE 802.11, Wi-Fi, WLAN.

I. INTRODUCTION

RADIO frequency jamming induces intentional interference that disrupts wireless communications on the physical layer. Therefore, jamming is hard to mitigate and commonly understood as a severe threat to wireless networks [21], [25], [26], [30], [44]. Unconventionally, jamming can also be used for the benefit of wireless communications. Two general application scenarios are usually promoted in the literature: (1) jamming as an intentional interference to prevent unauthenticated data from reaching the legitimate but resource-limited devices [5], [11], [16], [23], [24], [35], [42], and (2) jamming as a method to prevent eavesdropping on user communication by increasing the noise floor towards attacker [15], [17], [34], [39].

Manuscript received June 23, 2015; revised November 9, 2015 and February 29, 2016; accepted June 1, 2016. The associate editor coordinating the review of this paper and approving it for publication was K. Zeng.

D. S. Berger and J. B. Schmitt are with the Distributed Computer Systems Laboratory, Department of Computer Science, Kaiserslautern University of Technology, Kaiserslautern 67663, Germany (e-mail: berger@cs.uni-kl.de; jschmitt@cs.uni-kl.de).

F. Gringoli and N. Facchi are with the Department of Information Engineering, University of Brescia, Brescia 25123, Italy (e-mail: francesco.gringoli@ing.unibs.it; nicolo.facchi@ing.unibs.it).

I. Martinovic is with the Computer Science Department, University of Oxford, Oxford OX1 3QD, U.K. (e-mail: ivan.martinovic@cs.ox.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TWC.2016.2581165

Fig. 1. The friendly jamming scenario. The friendly jammer interferes with malicious communication of an attacker before it reaches the victim, while concurrent background traffic of legitimate users continues unaffected.

This work focuses on the following application scenario. An attacker seeks to send malicious messages (*attack frames*) to resource-constrained stations (*victims*). The victims are embedded devices such as sensor networks [23], [24], [42], implanted medical devices [11], or home automation and wearable gadgets [5], [16]. The victims cannot use sophisticated attack detection because they are running on battery [16], [23], [24], [42], or because updating them is unfeasible [11]. To protect such embedded devices, previous work has envisioned an external device called *friendly jammer*. Unlike the victim stations, the friendly jammer is able to detect attack frames, e.g., by real-time packet inspection [42], proximity information [16], or the absence of pre-shared tokens [11]. When detecting an attack frame, the friendly jammer emits physical-layer interference so that the victims cannot decode the attack frame and therefore discard it.

Friendly jamming promises to be a simple and effective protection mechanism for resource-constrained wireless devices. However, several challenges related to the physical-layer interference remain open. First, friendly jamming needs to be *effective* in a dynamic environment with multipath effects, attenuation, and fast channel fading. These effects can lead to undetected attack frames or insufficient interference with an attack frame. While prior work acknowledges these problems, the question remains how frequently this happens in a realistic environment. Second, friendly jamming also needs to be *minimally invasive*, i.e., only attack frames should be blocked, whereas the impact on legitimate transmissions should be minimized (see Fig. 1). Unfortunately, we are not aware of prior works that quantify the invasiveness of friendly jamming. Third, friendly jamming needs to be *pervasive*: friendly jamming should be available at low cost and for the most common communication protocol such as Wi-Fi (IEEE 802.11). Unfortunately, the technical challenges of friendly jamming meant that previous prototypes were based on powerful (and costly) specialized hardware platforms [5], [11], [42], [43] and do not support Wi-Fi (cf. Section II). In this work, as a part of real-world evaluation we implement

friendly jamming on off-the-shelf Wi-Fi access points. This approach solves the pervasiveness challenge, but entails new challenges for the effectiveness of friendly jamming. We analyze these challenges using both simulations and measurements with a realistic prototype based on the popular WRT54G access point. Specifically, we performed a friendly jamming experiment in a university IEEE 802.11g Wi-Fi network for three weeks: we measured both the effectiveness against the attacker and the invasiveness to the legitimate users of the network. We summarize our insights as follows:

- our analysis reveals that the WRT54G's jamming effectiveness depends not only on the power but also the symbol alignment of the jamming signal (with regard to the attack frame). The symbol alignment cannot be controlled by the friendly jammer so that multiple friendly jammers need to cooperate to increase their chance of the right alignment.
- our experimental evaluation shows that friendly jamming within a legitimate network is feasible. However, friendly jamming can increase legitimate station's packet loss due to a new type of hidden station problem, which has not been observed before and which we call *power amplification effect*. The packet loss of some stations doubled for seven jammers;
- jammers can collaborate to alleviate the power amplification effect. Our collaboration schemes can improve the packet loss by 20-30%, without compromising jamming effectiveness;
- we reproduce the previous insights with a different hardware setup. Friendly jamming is also effective for modern IEEE 802.11n networks in which the attacker can use antenna diversity, multiple spatial streams, and both the 2.4 and 5 GHz frequency band.

We discuss related work in Section II and jamming on access points (APs) in Section III; our analysis and simulation results are presented in Section IV and Section V describes our implementation. We present results from our measurement study for IEEE 802.11g and 11n in Sections VI and VII respectively. We conclude in Section VIII.

II. RELATED WORK

We review in this Section the main application cases of jamming in wireless networks, and we also discuss related works on jamming mitigation and general interference studies.

A. Blocking Harmful Communication

Table I gives an overview of prior works that use jammers to block attack frames similar to our friendly jamming scenario. With the exception of the works [24] (which does not evaluate reactive jamming) and [16] (which uses analysis and simulations), all of these works use the USRP2 software defined radio platform (e.g., [5], [11], [42], [43]) and do not target customer-grade APs. To the best of our knowledge, our system is the first to implement friendly jamming for IEEE 802.11 and on an AP. Most of the prior works used the IEEE 802.15.4 or industrial 400 MHz communication standard which have lower throughput and for which slower reaction times are sufficient

TABLE I
RELATED WORK ON BLOCKING HARMFUL COMMUNICATION

System	Technology and Evaluation
Jamming for good [24]	IEEE 802.15.4 sensor nodes, without reactive jamming
WiFire [42]	IEEE 802.15.4 USRP2, experiments with 2 friendly jammers
IMD shield [11]	400 MHz USRP2, experiments with 1 friendly jammer
Defend your home [5]	433 MHz USRP2, experiments with 1 friendly jammer
Considerate Jamming [16]	analysis and simulation

(cf. Section III). We also remark that unlike the USRP2 the WRT54G AP used in our work is likely to comply with legal regulations.¹

The performance metric considered in these works is the fraction of jammed attack frames (hit ratio): [24] report a median of 97%; [43] and [42] report 97.6% for one jammer and 99.9 for two concurrent jammers; [5] report 98.9 – 100%. All these experiments took place in a static setting (e.g., no other network, no moving people); we are not aware of a work that evaluates friendly jamming in a dynamic environment over several days as done in Section VI. The hit ratios in our experiments are significantly lower (e.g., 80 – 90% for two jammers) and we additionally consider the distribution of jammed (and not jammed) attack frames on shorter time scales.

B. Jamming for Confidentiality

In the second application case the jammer's interference ensures that an eavesdropper cannot decode network messages [15], [17], [20], [34], [39], [46]. The real-world feasibility of jamming for confidentiality has recently been thoroughly evaluated in [37].

Jammer placement schemes are considered in [17] and [32]. The setting of [17] differs from ours in that the jammers continuously emit a jam signal and thus there cannot be any co-existing networks. However, one of the goals mentioned [32] is to not interfere with any legitimate communication. This goal is similar to our collaboration schemes, but [32] relies on a theoretical propagation model and does not report experimental results. Also, our collaboration schemes address a different problem and they are based on different assumptions. Instead of the placement problem, we consider how jammers at fixed positions can collaborate efficiently. Instead of maximum interference at the attacker, we seek to maximize interference at the victim.

C. Jamming Mitigation

In the military context, jamming is an established technique to block an enemy's communication [1], [28]. Jamming has also been recognized as a threat to civilian networks [44].

¹Our prototype uses FCC-certified hardware with short low-power jamming signals that comply with limitations on maximum power and duty cycle [8].

Therefore, many works study ways to mitigate the effects of jamming [21], [25]–[27], [30], [44]. Many of the proposals (e.g., frequency hopping and robust modulation schemes) cannot be exploited by the attacker, as they would require cooperation of the victim stations. We assume that the victim stations follow the IEEE 802.11 protocol faithfully. This leaves the attackers a choice within the a set of allowed modulation and coding schemes (MCS). We evaluate the effect of this choice by considering all MCS configurations in our experiments.

D. Interference and Collisions in 802.11

Our analysis on jamming success in Section IV is related to the studies of packet error probabilities in face of non-intentional interference in IEEE 802.11 networks. These studies are motivated by the capture effect, which describes the successful reception of the stronger of two simultaneous transmissions [18]. Although intentional interference motivates a different perspective, the corresponding analysis often tries to answer the same questions. In particular, while initial capture models only considered the role of relative timing on the frame level and the received power [18], [19], [40], more recent models report observations that agree with our analysis: frame level temporal displacement and power ratio between interfering signals are insufficient predictors of interference probabilities. Our simulations in Section IV are most similar to [7], [10], [29], [33], and [41], which consider the effect of interference for different time and phase offset of the sender and receiver. However, these works only consider the temporal displacement at the frame preamble level [33]; focus on different technologies (IEEE 802.15.4) [41]; analyze the effects of jamming at the bit level when neither spreading or other explicit error detection and correction techniques are in use [29]; and consider only the relative time offset ignoring the relative power of the interfering signals [10].

What makes our analysis different from previous works is that we focus on IEEE 802.11 modulations and encoding schemes, we consider a scenario in which the receiver is already synchronized with the attack frame when the intentional interference starts and we take into account both the power ratio and the temporal displacement between interfering signals.

III. FRIENDLY JAMMING ON ACCESS POINTS

This section introduces our attacker model and highlights three fundamental challenges of friendly jamming which lead to our proposal of implementing friendly jamming on APs.

A. Attacker Model

The attacker seeks to send attack frames to victim stations, which are part of the wireless network. Victim stations faithfully follow the IEEE 802.11 protocol (uncompromised user stations), but do not implement any security functionality. Friendly jammers seek to stop each attack frame but emitting physical layer interference. Both the attacker and the friendly jammer are assumed to use standard hardware, which are

Fig. 2. Reactive jamming happens in four phases: 1) the jammer detects the attack frame (using the IEEE 802.11 training sequence), 2) analyzes the signal (to decide whether or not to jam), 3) incurs a reaction delay (due to switching from receiving to sending), and 4) emits the jamming frame. The interference is successful, if the superimposed signal's power P_{jammed} is significantly higher than P_{attack} .

power constrained to similar values. The attacker uses either an omnidirectional antenna (Section VI) or two omnidirectional antennas (Section VII).

B. System Challenges of Friendly Jamming

1) *The Reactive Jamming Challenge*: There are several jamming techniques that can block attack frames at the physical layer, e.g., continuously or randomly emitting a high power signal [26], [45]. However, *reactive jamming* is the most power-efficient technique and also the only one that can achieve minimal invasiveness: in reactive jamming, in fact, the interfering signal is selectively emitted only when another specific signal is detected on the channel.

Reactive jamming starts with the task of channel sensing in order to detect an attack frame (a signal in IEEE 802.11 terminology). The attack frame's detection and further decoding can only be achieved by jammers which are well positioned with regard to the arriving signal (to be specific: the attack signal arrives with high signal-to-noise-ratio (SNR)). This position is thus an important factor for jamming success. After having detected the signal, the jammer starts the decoding and analysis in order to make a jamming decision (cf. Fig. 2a). The jammer needs to distinguish between attack frames and legitimate frames. While the hit ratio ($\# jammed frames / \# sent frames$) needs to be maximized, the impact of jamming on legitimate frames shall remain small. This trade-off is further constrained by the fact that the decision has to be taken very fast. We recall these timing constraints by considering a frame comprising a medium-sized 100 Bytes UDP packet. If this frame is transmitted at the fastest 802.11g Modulation and Coding Scheme (MCS), its duration is only $48 \mu s$. If we further assume that the jamming decision depends on bits in the MAC header (e.g., the sender's MAC address), the remaining part of the frame is another $24 \mu s$ shorter. In the remaining time, the jammer's decision code has to be executed, and the hardware needs to switch from receive to transmission mode to emit the jamming signal. For faster communication standards like IEEE 802.11n the timing constraints become even more stringent. Once the signal is emitted, successful jamming depends on the attack frame's MCS and the jamming signal's power (cf. Section IV). While higher power drives high jamming success, this factor is subject to legal regulations and increases the detrimental impact on legitimate traffic.

2) *The Minimal-Invasiveness Challenge*: Although reactive jamming is key to achieving minimal invasiveness, there are

Fig. 3. The power amplification effect. (a) A legitimate user transmission is interference-free from the attacker's injections. (b) If a friendly jammer blocks the attack, it may accidentally interfere with the legitimate transmissions.

still side effects that can introduce interference to legitimate transmissions. We call the major such side effect the *power amplification phenomenon*.

Power amplification happens when friendly jammers significantly increase the interference radius of an attack transmission. This can be seen by considering a scenario in which a legitimate sender is located far away from the attack sender and transmissions of the two senders would be free of interference. Assume a jammer is in transmission (jamming) range of both senders: if the legitimate sender starts a transmission, the attack source could possibly start a concurrent transmission as well. Without jammers, both transmissions may proceed without problems (Fig. 3a); once the jammer interferes with the attack transmission, however, it also creates interference with the legitimate transmission (Fig. 3b). This interference can cause failure of the legitimate transmission. While in outdoor testbeds such constellations might be rare, we find that they happen more often in an indoor office environment.

Note that power amplification is different from the known hidden station problem. In hidden station problems, frames collide at a common receiver of two senders which cannot see each other. However, in the power amplification scenario in Fig. 3b the effect is due to the reactive jamming technique: both transmissions would proceed without problems, if the jammers were turned off and if the jammers were continuously jamming.

3) *The Pervasiveness Challenge*: Serving all users in a wireless network entails several sub-challenges. Firstly, friendly jamming needs to be available for those communication standards that are widely in use, such as Wi-Fi. Secondly, wireless communication underlies legal restrictions.² Such restrictions often define a maximum duty cycle (the fraction of one second a transmitter is active) that is commonly limited to small values and the maximum mean equivalent isotropic radiated power has to be less than a limit (e.g., 20 dBm for 2.4 GHz band in Europe [8]). In order to make the move from specialized use cases (e.g., military scenarios [1], [28]) to pervasive availability, friendly jamming needs to satisfy and work within these constraints.

Finally, friendly jamming needs good scalability to achieve coverage for every user in the network. This can only be achieved when a considerable number of well-positioned friendly jammers work together; in fact, a previous work even envisioned all network nodes to have friendly jamming functionality [24]. In any case this assumes that friendly jamming is made easily available, and at low cost. While previous reactive jamming and friendly jamming implementations exploited

the power of software-defined radios (SDR) (e.g., [2], [5], [42], [43]), scalability and low cost suggest to use customer-grade components. This is a great challenge as such components are typically underpowered which makes it hard to meet the timing constraints imposed by reactive jamming (which already pose a challenge for powerful software-defined platforms [2]). For customer-grade equipment, solutions to the underlying engineering challenges, such as reactive jamming, have begun to emerge only recently [3], [38].

C. Bringing Friendly Jamming to IEEE 802.11 Access Points

Our principal idea is to implement friendly jamming by modifying the software running on customer-grade APs. This has several advantages. Firstly, customer APs are optimized for the specific tasks (e.g., signal processing) necessary for Wi-Fi compatibility. They already implement current Wi-Fi features and can be expected to be kept up to date with future standards. Secondly, customer-grade APs are certified to be fully compliant with legal regulations. If we only modify the software, we are sure to encounter realistic hardware constraints. Thirdly and most importantly, APs are a natural deployment choice. Conceptually, APs should be the ones controlling a network's smooth operation and have much of the necessary network state information readily available. APs are also deployed in good positions to ensure coverage of all users in the network.

Modifying only the software of APs to implement friendly jamming functionality also poses several challenges. First, because of the real-time requirements of reactive jamming and delays in non-real-time architectures (e.g., bus delays), friendly jamming has to be implemented directly on the network interface card (NIC). This requires the ability to modify the firmware running on the NIC and detailed knowledge about the underlying hardware (chipset). Two critical changes to the firmware have to be made: the firmware needs to analyze a frame while still receiving it and it needs to be able to abort the current reception, switch to transmission mode, and emit the jam frame. As this behavior explicitly violates the IEEE 802.11 collision avoidance scheme, one needs to be able to work around such constraints. Additionally, to work as a friendly jammer, the system must support dynamic reconfiguration of the attack signatures.

Another part of the challenge is that APs are optimized for low cost: the WRT54GL AP – our choice – costs about 30-50 USD (5/2015 on ebay.com) and features a 200 MHz CPU and 16 MB of memory.

Finally, a major implication of the software only approach is that the jamming signal has to be a valid IEEE 802.11 frame: the encoding and modulation pipeline of Wi-Fi APs are hardwired and cannot be changed with a software update. It is not obvious that such a jamming signal can be used to interfere and thus block attack frames sent using the same technology.

IV. ANALYSIS OF FRIENDLY JAMMING ON APs

Jamming can be unsuccessful for many reasons, e.g., the friendly jammer might miss or react too late to a target frame,

²In particular, jamming is a sensitive topic for many country's legal regulators. For example, in the US the operation, marketing, or selling of continuous jammers such as GPS jammers, or cell phone blockers is prohibited [9].

Fig. 4. (a) Phase transition diagram for a real IEEE 802.11 frame (Broadcom). (b) The corresponding amplitudes are not constant over time but peak in the middle of a chip. The dashed line shows the amplitude of frame which is shifted by 1/2 a chip (Broadcom). (c) Amplitudes for an Atheros card.

or the jamming signal might not be powerful enough. This section focuses on the last case that is particularly challenging for friendly jamming on APs as they are power-constrained by legal regulations. In particular, we are interested in understanding how jamming success is affected by power ratio and temporal displacement between interfering signals. In the following, we consider DSSS and OFDM separately, and focus on the respective slowest data rates (1 Mb/s and 6 Mb/s) which are known to be the most resilient to interference. Also, in order to obtain realistic numbers, we ran trace-driven simulations by feeding IEEE 802.11 software receivers with the actual I/Q samples that we captured from real signals.

To the best of our knowledge, this is the first analysis of the relative power requirements for successful jamming as a function of chip or symbol alignment for IEEE 802.11abgn

A. Analysis of DSSS

We start by describing the key encoding steps defined by IEEE 802.11 [14] for the DSSS-PHY. This section assumes some prior knowledge about DSSS, which can be found in [28].

1) *Encoding DSSS*: First, bits are pre-processed to have a balanced number of 0s and 1s (called scrambling). Second, every bit is expanded to eleven chips with redundant information (called spreading gain, it also boosts the bandwidth of the transmitted signal to fit in the typical 20 MHz spectrum considered by the standard). Third, each chip is modulated with *Differential Binary Phase-Shift Keying* (DBPSK), which changes the phase by π for a 1 and leaves the phase unchanged for a 0. Fig. 4a shows the resulting phase transitions for a real IEEE 802.11 frame. The chip boundaries are marked with vertical lines. We observe that although DBPSK's phase changes are instantaneous in theory, they happen gradually in the real world due to hardware limitations. Figs. 4-8 (except Fig. 6) are based on traces of customer-grade IEEE 802.11 hardware captured with a real-time spectrum analyzer (Tektronix RSA3408). These traces have a resolution of four samples per chip (sampled with 12 bit at 44 MS/s).

2) *Decoding DSSS*: In order to recover the original chip and decode the frame, the receiver seeks to obtain the correct phase change: to this end it relies on the values of the phase occurring in the middle of each chip, i.e., on the “plateau” between two chip boundaries. We mark the ideal sampling time with arrows in Fig. 4. After sampling, the receiver assigns the chip values according to the phase variation, a phase change indicates a 1, no phase change indicates a 0. The receiver decides the sampling timing at the beginning of a frame, when decoding the preamble: it then keeps the same timing for decoding the rest of the frame.

3) *Chip Alignment and Jamming*: In DBPSK, the amplitude does not carry any information and it could therefore be expected to be constant over time. Nevertheless, our measurements show that the amplitude of real IEEE 802.11 frames change over time; Fig. 4 shows two examples for the popular Broadcom chipset in the WRT54G AP (4b) and for an Atheros chipset (4c). All the plots in Fig. 4 uses the same time scale and chip boundaries: this shows that the amplitude peaks in the middle of a chip, i.e., at the ideal sampling point of the phase transitions. The observation of this amplitude peak is significant to understand the probability of jamming success.

Friendly jamming uses a reactive jamming approach (cf. Section III). Therefore, when the interference of the friendly jammer starts, the receiver has already processed the preamble of the attack frame and fixed the sampling timing. As mentioned before, the receiver will keep the sampling timing to decode the signal, which consists of the superposition of the jamming and attack frame. A higher amplitude of the jamming signal occurring where the receiver is measuring the phase transition implies a higher chance that the receiver makes a wrong phase decision, which would lead to a wrong chip interpretation. Several such wrong chip interpretations are required to happen consecutively for successful jamming because of the eleven-fold redundancy of DSSS (spreading gain) [10]. For friendly jammers based on customer-grade IEEE 802.11 hardware, the jamming signals are valid IEEE 802.11 frames with a high amplitude in the middle of each jamming chip (and a low amplitude at the chip boundaries). If the interfering signals are aligned at the chip level, their amplitudes peak in the same points, exactly at that point when the receiver samples the phase of the attack frame. Therefore, the case of exact chip alignment yields the most effective jamming. On the contrary, when the two frames are off by 1/2 of a chip, jamming has the least success probability (dashed line in Fig. 4b).

4) *Chip Alignment Is Random*: Our measurement analysis of IEEE 802.11 I/Q traces indicates that consecutive frames transmitted by a specific node display a random alignment at the chip level. There are three independent reasons for this observation. First, the IEEE 802.11 standard [14] allows a 25 ppm tolerance of the chip clock frequency to avoid the circuitry for compensating temperature variations: as the clock skew differs between any two devices, the attacker and the friendly jammer will see a drifting chip alignment over time even if they start with a perfect alignment. Second, because of the reactive jamming technique jammers need to switch from listen mode to transmit mode: this switch depends on the device states and therefore adds a random delay. Third, the propagation delay depends on the (uncorrelated) positions

Fig. 5. The lowest gain (P_{Jam}/P_{Att} cf. Fig. 2b) at which jamming success is possible has a periodic dependency on the chip alignment. (a) Trace-driven simulation results show that alignments which are a multiple of the chip length ($\approx .09[\mu s]$) are particularly robust against jamming (i.e., need a high jamming gain). (b) The Fourier transform of the simulation data clearly shows the component of the periodic dependency.

of the attacker and jammer and changes over time in a dynamic environment.

5) *Trace-Based Simulation Results*: We have implemented a software receiver running on I/Q samples of real IEEE 802.11 data frames. Our receiver implements the DSSS decoding steps and actually outperforms various hardware receivers, i.e., our software receiver can decode frames for which respective hardware receivers have not sent an acknowledgment. To simulate jamming, we superimpose a jamming IEEE 802.11 1Mb/s frame at different chip alignments. We increase the peak power of the jamming frame until the attack frame cannot be decoded anymore.

This gives us the lowest gain in the jamming peak power with respect to the attacker signal peak power at which jamming is successful (in the following 0dB denotes peak powers are the same). We repeat this experiment for 27 frames and show the results in Fig. 5a which reports, for every alignment, the average gains (together with minimum and maximum) required for successful jamming. The figure shows a strong periodic dependency of the required jamming gain on the chip alignment. As predicted, this chip alignment corresponds to 1/2 of the length of a DSSS chip ($\approx .09[\mu s]$). A Fourier transform visualizes this dependency in Fig. 5b.

All simulations results reported so far were based on traces collected with a real-time spectrum analyzer. To verify that the chip alignment dependency is due to the non constant chip amplitudes (cf. Fig. 4b), we have implemented a synthetic trace generator. This trace generator encodes DSSS frames with the same payload as the traces considered so far. We found that the effect of the chip alignment happens only when the signal amplitude peaks in the middle of a chip as observed in the real traces. We reproduce an ideal version of the peaking amplitude behavior and show the results in Fig. 6. The effect on the chip alignment is very similar to the results for the trace-based simulations (yet, more pronounced). In fact, the Fourier transform analysis shows that the main components of the synthetic data match the trace-based simulations.

6) *Conclusions for Friendly Jamming*: The attacker can gain an advantage over a single friendly jammer by using the robust DSSS encoding. Due to the chip alignment effect, this encoding has a significant chance of not being jammed by the friendly jammer. This particularly applies to scenarios in which both attacker and friendly jammer are power constrained to similar magnitudes (e.g., use customer-grade hardware). This is the case in our experiments in Section VI.

The most practical solution to this challenge is to deploy more friendly jammers. Thanks to their independent chip

Fig. 6. Repetition of the experiment in Fig. 5 but with synthetic trace data. (a) Periodic behavior of jamming success as a function of the chip alignment occurs more pronounced but very similar. (b) Corresponding Fourier transform.

alignment, additional jammers would make the chance of a 1/2 chip alignment arbitrarily small. We corroborated this conjecture in additional simulations (not shown).

B. Analysis of OFDM

In contrast to DSSS, OFDM builds the transmitted 20MHz spectrum by shaping it rather than relying on spreading. To this end the transmitter sets each of the 52 orthogonal carriers separately before invoking an Inverse Fast Fourier Transform (IFFT) and getting a *signal symbol* in the temporal domain. Furthermore, to add resiliency to frequency-selective multipath channels, the transmitter prefixes each symbol with a repetition of its own end to create a *Guard Interval* (GI). This section investigates the alignment effect of this prefix (called *Cyclic Prefix* (CP)) based encoding. We call this the *jamming symbol alignment*, in analogy to the *chip alignment* of DSSS. This section assumes prior knowledge about OFDM, which can be found in [6].

1) *Encoding OFDM*: The encoding process defined for the ERP-OFDM-PHY starts by dividing the frame into fixed size blocks. Each block is then enriched with forward error correction (FEC) information and the resulting data is distributed over the 52 spectrum carriers according to specific constellations, e.g., $\{-1, +1\}$ for the basic BPSK scheme for 6Mbit/s. After padding the 52 coefficients with zeros, the encoder runs the IFFT and obtains 64 complex values that it finally prefixes with a CP of 16 samples. For more detailed explanations (e.g., different datarates), we refer to chapter 18 of [14]. The resulting 80 temporal I/Q samples transmitted at 20MS/sec lead to a symbol duration of 4μs.

2) *Decoding OFDM*: Key for a successful reception of a frame is the recovery of the original spectrum associated to each symbol. To this end the receiver starts by capturing from each symbol a window of 64 I/Q samples. Fig. 7a shows two examples of such windows. Thanks to the CP, any two windows are equivalent as each one is the cyclic shift of the other: after applying the FFT they differ only for a phase factor that is linear in frequency. The receiver can hence eliminate this factor by knowing the symbol boundary timing, which is learned during the processing of the PCLP preamble [22]. The reception of the physical header allows also to correct the Carrier Frequency Offset [36], and to estimate the channel response: the latter is necessary to equalize the recovered spectrum and mitigate the non-flat channel fading before deciding which value from the corresponding constellation, e.g., $\{-1, +1\}$, was carried by every carrier. Finally, FEC information is used to decode the original data associated to

Fig. 7. (a) Symbol alignments of jam signals. (b) Minimum power peak gain for successful jamming w/ and w/o FEC for synthetic traces and no channel fading. (c) Two cases of minimum power peak gain for successful jamming with real traces.

each symbol before putting all pieces together in the recovered frame.

3) *Symbol Alignment and Jamming*: As shown in Fig. 7a, the interaction between the attacker and the jamming frames changes according to their symbol alignment. Specifically, this gives rise to one of either two cases depending on how many jamming symbols fit in the decoding window when processing the attack's preamble. First, when the decoding window covers a single jamming symbol (Jamming frame Jam_1 in Fig. 7a), then the result of the FFT operation on the two overlapping symbols leads to the sum of the corresponding spectra. Assuming a perfectly flat channel and equal constellations $\{-1, +1\}$, at the output of this phase we have 52 values that are the sum of constellation symbols: the constellation of the jamming is however rotated of the relative delay. If the intensity of the jamming signal is below that of the attack frame, then perfect decoding of the latter is still possible, independently of the rotation. Otherwise jamming might induce decoding errors. Second, two consecutive jamming symbols can also partially fit inside each decoding window (Jamming frame Jam_2 in Fig. 7a). In this case, after the FFT operation the spectrum of the attack frame is corrupted with coefficients that do not belong to any constellation. The FFT applied to a window of samples containing the tail and the head of two consecutive jamming symbols lead to random data, not shaped according to any constellation. Then, some of the 52 values will not be decodable even if the intensity of the jamming signal is lower. Decoding issues will increase with more complex constellations and in the presence of non flat-fading channel: in this case the equalization restores the attacker carriers inside their relative constellation while it further damages those of the jamming signal increasing the probability of non correct decoding.

4) *Simulation Results*: We start by considering synthetic traces that we generate and decode using the gr-ieee802-11 framework [4]. We simulate jamming by adding to a synthetic attack frame I/Q samples belonging to another (jamming) frame. Fig. 7b shows, for increasing symbol alignment (multiple of $4\mu s$ means perfect alignment), the gain in the peak power of the jamming signal with respect to peak power of the attacker that is required for successful jamming. In the top graph we disable FEC to better show the effect of the

alignment and the periodicity over 10 symbols: there are regions long as the cyclic prefix which require high gain ($0dB$) separated by regions requiring less, due to the interaction of multiple jamming symbols inside the attacker decoding window. When FEC is enabled the periodicity is no more evident: anyways, the spectral analysis in both cases reveals a strong component centered at $f = 250KHz$ that correspond to the $4\mu s$ periodicity together with its harmonics. We also gather from the Fig. that FEC introduces better resiliency towards jamming in the longer regions: the required gain for successful jamming increases and is almost flat with the delay (very close to $0dB$).

We then repeat the same analysis with I/Q traces of real signals transmitted by Broadcom devices and captured with an USRP-N210: before using the traces we normalize all signals to the same energy. We simulate jamming by adding to an attack frame transmitted by one device a jamming signals transmitted by another one: we use data from different devices on purpose to avoid artifacts due to the same transmission clock. To attempt decoding we still use gr-ieee802-11. We report in Fig. 7c the gain required for successful jamming with two different jam signals (differently than with DSSS we cannot meaningfully report here the average plus confidence interval so we report results separately for set1 and set2): periodicity is even less evident with real traces. We also notice that with respect to the simulation of the synthetic traces with FEC, the minimum gain for successful jamming diminishes thanks to the effect of the channel equalization on the jamming samples (almost flat around $-2dB$). Then, we repeat the spectral analysis by averaging the behavior of different jamming signals captured for set 1 and set 2: we see that each jamming set leads to a specific behavior but both have a strong component at $f = 250KHz$ (we also see a few harmonics).

5) *Conclusions for Friendly Jamming*: When considering different alignments at the symbol level, our synthetic traces show the existence of robust OFDM symbol alignments in theory. Such alignments could mean an advantage for the attacker similar to the robust chip alignment cases for DSSS. However, our simulations with real traces show that this effect is less significant in practice. These observations substantiate the common knowledge that the IEEE 802.11ag OFDM

Fig. 8. Hit ratio for (a) different attack frame MCSs and for (b) a 1 Mb/s DSSS vs. a 6 Mb/s OFDM jam frame.

encodings are less robust against jamming attacks, which is advantageous for friendly jamming.

We also validated that 6Mb/s is the most jamming-resilient OFDM encoding: Fig. 8a compares the susceptibility of four different OFDM data rates between 6Mb/s and 18Mb/s. Indeed, a jammer gain of -12dB is enough to reach 100% hit ratio for 18Mb/s but while for 6Mb/s we need to have similar peak powers (0dB).

As a final simulation experiment, we also consider the effectiveness of two different jamming signals. Fig. 8b compares whether jamming a 6Mb/s frame with DSSS frames or with OFDM frames is more effective. As a fair comparison, we normalized both jamming signals to the same energy level. Nevertheless, the 1Mb/s DSSS jamming frame has a slight energy advantage. We leave the further investigation of this effect open to further research and fix 1Mb/s jamming signals for the following measurement studies.

C. Analysis of Carrier Phase Offset

We next study the impact of different carrier phase offsets (CPO) between the jamming signal and the attack frame. Without spreading or other forward error correction, CPO is known to require an increased power of the attack frame [29]. We find that this is not the case for the IEEE 802.11 DSSS and OFDM modulations. Specifically, we again use our software decoder with a synthetic trace, in which the CPO uniformly takes one of 32 values in the interval $\theta \in [0, 2\pi[$. We multiply the jamming signal for the corresponding CPO coefficient $\exp[j\theta]$ and again consider different chip and symbol alignments for DSSS and OFDM, respectively. Fig. 9 shows the power gain required for successful jamming: the thick lines in the three Figures corresponds to perfect phase alignment and, in fact, they are exactly the same gains as shown Figs. 6a and 7b. The thin lines in Fig. 9 correspond to other values of the CPO. We find that the CPO does not introduce major changes on the jamming behavior, i.e., the same periodic fluctuations are still visible for all the curves. We also calculated the overall change in the power required for successful jamming by averaging over the entire time and phase delay offsets. We find that in all the three considered cases the gain changes are negligible. This quantitatively confirms that the dynamics involved by complex modulations using either spreading or frequency multiplexing do not require explicit control of the phase of the jamming signal for modifying the information decoded by a receiver.

V. IMPLEMENTATION

The friendly jamming firmware was developed in three steps. First, a basic prototype for IEEE 802.11g (on the Broadcom 4138 chipset) showed the idea's feasibility;

second, we optimized it with extensive lab experiments; and third, we ported the friendly jamming project to a modern IEEE 802.11n platform to see how fast future hardware or standard changes can be incorporated.

Access to the network interface card's firmware is facilitated by the OpenFWWF open source project [12] (also see [31]). The basic OpenFWWF firmware can receive and send IEEE 802.11-compliant frames. As outlined in Section III-C, significant changes were necessary.

Our first prototype was validated using several micro-benchmarks, e.g., traces by real-time spectrum analyzers as in Fig. 2 and open-space experiments (Tbl. II). The second step included a variety of optimizations to shorten the jammer's reaction time and increase jamming performance: we implemented a prefetching mechanism that proactively pulled the jammer's (re)configuration from the system memory and stored it in fast-access registers on the NIC. We then improved the management of the transmission power: we ran comprehensive tests on corresponding hardware registers before settling on a final configuration. As motivated by the results in Section IV, we used 1Mb/s DSSS jamming signals. As the jamming frame's length had negligible impact, we set it to the minimum allowed by the hardware (ten bytes) that led to a jam signal's duration of $272\mu\text{s}$ ($192\mu\text{s}$ for the PLCP header, other $80\mu\text{s}$ for the rest). We also measured the jammer's reaction time after matching a target MAC address: by comparing original and jammed frames' content, we verified this delay to be two bytes at 1Mb/s, i.e., $16\mu\text{s}$.

The third step involved the porting of the software to a modern chipset (the Broadcom 43224) that supports IEEE 802.11n (2×2 MIMO). Because the selected chipset is supported by different kernel driver and firmware, we developed the required software from scratch making the new system back-compatible with the old configuration files.³

This concludes the discussion on the implementations. The basic IEEE 802.11g implementation (result of the second step) is used for the experiments in Section VI. The IEEE 802.11n implementation (result of the third step) is used for the experiments in Section VII.

VI. RESULTS: IEEE 802.11bg NETWORK

This section describes results from a measurement study of our friendly jammer implementation in a complex radio communication environment.

A. Deployment and Measurement Data

We selected a university office floor as a particularly heterogeneous and dynamic environment. The observed background traffic is generated by a dynamic user base of faculty members and students; it comprises a variety of traffic sources such as typical back-office activities, research-oriented applications, and multimedia applications. Besides the university's main campus network, there are various smaller APs with many overlapping Basic Service Sets, so that we are able to observe different network types in parallel.

³Our preliminary friendly jammer implementation is available on <http://www.ing.unibs.it/~openfwwf/friendlyjammer/>.

Fig. 9. Minimum gain required for jamming synthetic traces successfully. (a) DSSS, (b) OFDM without FEC (attack and jam frame 6Mb/s) and (c) OFDM with FEC (attack and jam frame 6Mb/s). The jamming power is less impacted by the $CPO \theta \in [0, 2\pi]$ than by the chip or symbol alignment. Also note that as in Fig. 7, FEC again destroys most of this effect.

Fig. 10. Experimental setup: the attacker and 13 jammers use WRT54GL aps and are located on a university floor of about $40 \times 85m$. Traffic is recorded with ALIX.2D2s on three positions, m_{close} , m_{middle} , and m_{far} .

We deployed one attack traffic source, 13 friendly jammers, and three monitors on the university floor, see Fig. 10. The locations of all devices were constrained by availability but several pre-runs showed the same qualitative outcomes for various locations. In these pre-runs, we experimented with different attack traffic rates, inter-frame spacing (e.g., SIFS) and backoff mechanisms. We found that the jamming performance is unaffected for all configurations tested (e.g., full saturation of the channel with attack traffic) because the jammers emit a very brief jamming signal, and go back to listening on the channel before an attack frame ends (cf. Fig. 2b). However, as discussed in Section V, the jamming performance is affected if the packet payload is less than 42 bytes. For the remainder of this section, we report on results for standard-compliant inter-frame spacing, a constant attack rate of 500 Kb/s, and 50 bytes packet payload. As channel 6 is the busiest of the university network, we chose it for studying the jamming system in order to see the negative impact of friendly jamming. The attack traffic used a round-robin selection of all IEEE 802.11b/g modulations without retransmissions, so that the friendly jammer could not lock onto any modulation. This way the impact of different modulations on the jamming performance could be fairly judged.

The 13 friendly jammers used the configuration from Section V and blocked traffic based on the attacker's MAC address. The attack frames are received by three monitors (m_{close} , m_{middle} , m_{far} , see Fig. 10). To accurately track sent, received, and valid-checksum statistics for the attack traffic, we implemented low-level counters (on the NIC of the monitors) that were not impacted by the operating system. The traffic that we captured from legitimate users was recorded as a full pcap trace and streamed to a central server over dedicated Ethernet lines. In summary, our setup was carefully crafted to obtain reliable attack traffic and accurate measurements.

All results in this section come from 24h-experiments over a three-week-period to observe different conditions on the wireless channel. Time was divided into short experiments

Fig. 11. (a) A single jammer's hit ratio is fully dependent on its position. (b) Using an increasing number of jammers (0 corresponds to packet loss without jamming) brings the hit ratio asymptotically close to 100%.

TABLE II
OPEN-SPACE HIT RATIO

dist	1 Jammer	2 Jammers	3 Jammers
5.4m	98.93% (± 0.19)	100.0% (± 0.00)	99.98% (± 0.02)
6.7m	98.57% (± 0.85)	99.95% (± 0.04)	100.0% (± 0.01)
8.1m	77.03% (± 12.3)	98.37% (± 0.29)	99.43% (± 0.14)

each of about 150s, where in each experiment a random number and selection of jammers were activated uniformly from all possible combinations. In total, this setup generated 13660 experiments with about 350GB of measurement data comprising almost 370 million target frames and 490 million legitimate-traffic frames. Throughout this section, we use 97.5% confidence intervals based on the t-distribution.

B. The Jamming Performance

We measure jamming performance via the hit ratio.

1) *Position-Dependence*: In open-space testbed experiments (e.g., [3], [42], [43]) the hit ratio monotonically decreases with the jammer's distance from the receiver. We also observe a monotonic behavior for our jammers, and we thus summarize only part of the measurements in Table II. In contrast to this, behavior in our in-door scenario is quite erratic and there are weak positions even close to a victim receiver (cf. Fig. 11a). As expected from the analysis in Section IV, the impact of weak positions is more pronounced for robust MCSs and small numbers of jammers. For example, the hit ratio when averaging over random placements of three jammers decreases by up to 30% for the robust MCS 1Mb/s; however, for seven jammers this issue mostly disappears (less than 5% difference between MCSs) as it is more likely that a good position is included in the active jammer set.

2) *Number of Jammers*: An increasing number of jammers increases the hit ratio due to improving the detection and emitted power of the jammers (Fig. 11b). Note, however, that one may grossly underestimate the number of jammers necessary

Fig. 12. We complement the hit ratio with an additional metric called *unpredictability of jamming misses*, which is explained in the legend on the left (if the attacker observed that two (#5, #10) out of twelve attack frames (#0-#11) were not jammed, can it predict its next attack opportunity?). (a) Increasing the number of jammers from 3 to 7 to 13 thins and spreads out the occurrence of not jammed attack frames. (b) Quantile-quantile plots show good correspondence with the geometric distribution, which indicates unpredictability for the attacker, for seven jammers but less so for three jammers.

when only considering open-space experiments (compare to Table of Fig. II). Due to the weak positions described above, a single jammer's hit ratio lies between 60 – 90% (averaged over all positions). At least three jammers are required to sustain an average hit ratio above 90%, seven jammers for 99%, and the hit ratio graph's slope only flattens out for more than eleven jammers. This appears to the case across receivers at different proximities, which yielded distinct packets loss rates without any jammers (19 – 48% at 12 – 75m, cf. Fig. 10).

3) *Unpredictability of Jamming Misses*: The hit ratio defines an aggregated view of the jamming performance, which might hide short-term glitches of the friendly jammers. Fig. 12a shows that this is not the case: not-jammed attack frames are distributed uniformly over time. Specifically, we consider the series of the sequence numbers of the attack traffic and we find that increasing the number of jammers from 3 to 7 to 13 makes this distribution increasingly uniform. Such a uniform distribution of not-jammed frames over time means that no statistical knowledge can be obtained for the attacker from observing past not-jammed attack frames, which makes their exploitation hard. We formalize this observation by considering the inter-miss distribution on the sequence number space. If this distribution is close to the geometric distribution, then jamming misses possess the memoryless property, i.e., are unpredictable even when an attacker is able to observe the friendly jamming system over a long time. Indeed, the quantile-quantile plots in Fig. 12b show that, qualitatively, the measured inter-miss distribution converges to the geometric distribution for seven jammers. However, three active jammers seem insufficient to enforce this property.

C. The Cost of Friendly Jamming

In this section, we investigate the side effects of friendly jamming on *legitimate* traffic in the network. Recall that legitimate frames have not been targeted by the frame-selective jammers; instead, any adverse effect on legitimate traffic is collateral – the cost of friendly jamming.

In order to assess this cost, we distinguish legitimate frames from attack frames (and jamming signals) in our measurement data. This required some attention because fragments resulting from the attacker's frames and the jammers' frames could not easily be distinguished from actual legitimate frames due to high bit error levels. We used an automated procedure, which we verified selectively down to the bit level of individual frames: we first discovered valid legitimate transmitter MAC

Fig. 13. Packet loss of legitimate traffic subdivided by RSSI levels. (a) The loss at monitor m_{close} is stable. (b) monitor m_{medium} observed an increase in the packet loss for senders with medium RSSI, which we attribute to a slight power amplification effect. (c) the packet loss of traffic from low RSSI senders to m_{far} doubles with an increasing jammer count.

addresses, and then matched frames with bit errors to the legitimate stations using the smallest Hamming distance.

1) *Loss of Legitimate Traffic and Power Amplification*: Friendly jamming can negatively affect legitimate traffic as suggested by the power amplification phenomenon discussed in Section III. This, however, happens only for certain constellations of a legitimate station to the attacker and to the jammers. We quantify this effect by exploiting the fact that our monitors m_{close} , m_{middle} , and m_{far} cover a large observation area, 85m in length (cf. the map in Fig. 10). For each of the monitors, we categorize the observed senders into three classes based on the received signal strength: low RSSI ($< -65dB$), medium RSSI ($-65dB < RSSI < -55dB$), and high RSSI ($> -55dB$). The packet loss for each sender in these classes is measured with the monitor as a receiver, based on each frame's checksum. Fig. 13 reveals that the packet loss increases with the number of jammers only for two classes (and is stable for the other classes). The most badly affected class is the traffic with low RSSI as observed by m_{far} (Fig. 13c): the packet loss doubles from 20.2% to 40.4%. The low-RSSI senders are located at positions to the right on the floor (Fig. 10), where they do not receive the attack messages but may suffer interference from the jammers in the center (i.e. the setting described in Fig. 3b). Accordingly, we attribute the observed increase of packet loss to the power amplification phenomenon, for which the likelihood increases with an increasing number of active jammers. The second affected class are m_{middle} 's medium-RSSI senders (Fig. 13b): the packet loss increases from 13.7% to 22.6%. We also explain this observation with the jammer's power amplification; however, the effect is less pronounced than before since there are fewer candidate jammers. It is also interesting to note that the low-RSSI senders at m_{middle} are not similarly affected by this phenomenon – we attribute this to erratic propagation effects of the indoor topology (cf. Fig. 10);

Fig. 14. (a) The RSSI (at m_{close}) surges when enabling the first jammer, indicating an increased noise level. (b) The adapter rate of background traffic becomes more volatile as soon as the jammers are enabled – indicating higher activity of the rate selection algorithm. (c) The mean number of (re)association requests becomes more volatile but remains approximately stable.

for example, there is a fire door to the left of this monitor that significantly alters signal propagation paths.

Summing up, we found that friendly jamming can substantially impair the channel of legitimate transmissions due to the power amplification effect. Future studies of friendly jamming should thus seriously consider a cost perspective besides traditional metrics of performance and security.

2) *Interaction With Rate Adaptation:* The packet loss of stations, which are not affected by the power amplification effect, does generally not depend on the number of active jammers. Some configurations (cf. Fig. 13) seem to have a slight impact; but we attribute this to increased variability when more jammers are active. This increased variability becomes evidence when considering the RSSI recorded at the close-by monitor: the RSSI increases significantly (Fig. 14a). As the RSSI indicates an increased level of channel noise, rate adaption algorithms can be expected to select a more robust modulation. Indeed, Fig. 14b shows that the adapter rate shows greater variance and tends towards more robust MCSs when the jammers are enabled.

3) *Effect on Associations:* As a final aspect of the cost of jamming, we consider the number of (re)associations requests under different numbers of jammers. Fig. 14c presents the change in this number with respect to the configurations without jammers. While the close-by stations seem to be affected by an increase of up to 40%, the farther-away stations exhibit approximately stable behavior. We also measured (re)associations of individual stations, and find that the number of significantly affected station is a small.

In summary, we again emphasize the trade-off between jamming performance and the cost of jamming: enabling a larger number of jammers is necessary to boost the hit ratio but also increases the collateral damage to legitimate traffic. While not unexpected, we obtained detailed quantitative statements on the cost of jamming in a real-world 802.11 network. These results were obtained under the assumption of random jammer placement. However, a friendly jamming system might be able to overcome this limitation as investigated next.

D. Jammer Collaboration Schemes

In this section, we show the potential of jammer collaboration, by which we mean that jammers are selected according to some deterministic rule instead of randomly as above. This promises to improve the hit ratio while remaining marginally invasive, based on the previous section's finding that the

jammer's position is critical to both the performance and cost of friendly jamming (cf. Fig. 15).

Our first scheme selects the jammer closest to the attacker (e.g., realized by trilateration), which is denoted *ACL*. Our second scheme selects the jammer closest to the victim, which is denoted *VCL*. This can be possible, if the jamming system can observe bidirectional communication attempts between the attacker and its victim. Compared to *ACL*, we expect *VCL* to improve the hit ratio because the friendly jammers would be closer to the receiver contributing to an higher energy jamming signal. Our third scheme is based on the idea to select jammers based on their reception quality. Specifically, the scheme *RCL* selects jammers that received the highest number of attack frames. While this heuristic only considers the channel from the attacker to the jammer, one may speculate that it also corresponds to better jamming positions in general.

We compare our schemes to random jammer selection (*RND*) and to the optimal configuration (*OPT*), which is found by an exhaustive search through all possible jammer selections for each repetition of the experiment. We use two metrics to compare the schemes: the hit ratio for the attack traffic (Fig. 15a) and the packet loss of legitimate traffic (Fig. 15b). We normalize the packet loss with regard to the *RND* scheme to show how our schemes improve upon the power amplification effect. Specifically, the legitimate packet loss is measured as the normalized packet loss where normalization is done to the *RND* configuration that yields the same hit ratio.⁴ The simple *ACL* scheme significantly increases the jamming performance and reduces the legitimate packet loss (power amplification effect). We attribute this improvement to the proximity of the receiver and the sender and to the fact that far-away jammers are efficiently excluded. The extension, the *VCL* scheme, yields a slightly better hit ratio than *ACL* and can reduce the packet loss even further. However, we remark that this scheme is less practical than *ACL*. The third practical scheme *RCL* can outperform *VCL* slightly, but with overlapping confidence intervals. While all collaboration schemes improve significantly over random jammer selection, there remains room for improvement as the hit ratios are 2.5 – 4% below the *OPT* case.

We also studied the schemes for seven jammers. In that case, the mean hit ratios are closer together: 98.95% for *ACL*, 99.41% for *VCL*, 99.62% for *RCL*, and 99.95% for *OPT*.

VII. RESULTS: IEEE 802.11n NETWORK

In this section, we extend our measurement scope to more broadly account for the Wi-Fi feature set occurring in practice. Specifically, we target the new features introduced with the IEEE 802.11n standard (e.g., spatial substreams), the impact of different receiver types, and measurements on 5 GHz channels.

All receivers were equipped with two antennas. We tested various different antenna types (including those specifically built for the respective frequency bands), and selected a type of dual-band standard-dipole antennas that showed the best overall performance for 2.4 GHz (channel 11) and 5.2 GHz

⁴Normalized packet loss of scheme x with hit ratio p : (packet loss of x) / (packet loss of *RND* with that number of jammers that yield a p hit ratio).

Fig. 15. When the position of the attacker and the victim is not known a priori, random jammer selection (RND) with several candidate positions is the only choice. However, a *collaboration scheme* can improve the jamming performance by selecting appropriate positions based on additional measurements: the attacker's position (ACL), the victim's position (VCL), or the reception ratio of individual friendly jammers (RCL). These three collaboration schemes lead to near-optimal improvements of the hit ratio (a) and reduce legitimate packet loss (due to the power-amplification effect) by 20-30% (b).

Fig. 16. (a) The traffic source and nine friendly jammers were located in a similar layout as in the previous experiment (cf. Fig. 10). (b) The frequency band has only a small impact on the jamming performance; the main difference can be seen in the packet loss without jamming due to reduced interference on the 5 GHz band. Also note that the new chipsets used here yield a higher hit ratio for our friendly jammers (cf. Fig. 11b).

(channel 44). To simplify the comparison to the previous IEEE 802.11bg setup, we deployed the new hardware on the same floor and on similar positions as before, cf. Fig. 16a. Unless noted otherwise, all measurements are due to the monitor m_{bcm} (Fig. 16a) which uses the same hardware configuration as the friendly jammers.

1) *Role of Chipset Implementation and Frequency Band:* Fig. 16b shows the hit ratio for the new friendly jammer implementation. In comparison to previous results, the hit ratio increases much faster with the number of jammers. The hit ratio curve flattens out for about six jammers. We attribute this to a measurably improved reception ratio of the more modern chipset used in the jammers. This results in fewer missed jamming opportunities by individual friendly jammers and thus increases the hit ratio at lower jammer counts. Another important factor is that for the 5 GHz band the IEEE 802.11an standard restricts the attacker's transmissions to OFDM modulations, which previously turned out to be more susceptible to jamming. To ensure comparable results, we also enforced the same restriction on the 2.4 GHz results reported in Fig. 16b. Overall, a significant difference in hit ratio between the two frequency bands only appears without jammers, for which the 2.4 GHz has a higher packet loss, and for 1 – 2 jammers, for which the hit ratio increases less steeply compared to the 5 GHz results. We attribute this difference mainly to a generally reduced interference on the 5 GHz band (increased channel spacing and fewer active stations). We conclude that the hit ratio analysis of the previous section essentially holds for all the IEEE 802.11abgn frequency bands. Nevertheless, comparing the previous friendly jammer implementation's results from Fig. 11b on 2.4 GHz to the new results in Fig. 16b reveals that specific wireless chipset technologies can have some impact.

Fig. 17. (a) In our setup, jammers and attackers are 2x2 systems, which can be either configured to exploit antenna diversity and one IEEE 802.11n spatial stream (top), or in a multiplexing configuration with two IEEE 802.11n spatial streams (bottom). (b)-(c) The box plots show that for two types of receivers (m_{bcm} and m_{ath} respectively) both one (top) and two (bottom) spatial streams are jammed successfully.

2) *Spatial Multiplexing and Receiver Type:* The IEEE 802.11n standard exploits MIMO antenna systems to provide independent spatial paths between the transmitter and receiver so that different streams of information can be sent in parallel along these spatial paths (roughly speaking, for an introduction see [13]). In contrast, IEEE 802.11abg exploits only the antennas' diversity, e.g., selecting the better one from a pair of antennas (cf. Fig. 17a).

We are interested in whether an attacker's transmission using the 11n spatial multiplexing feature is more or less susceptible to interference from our friendly jammers. Fig. 17b shows a bar plot of the hit ratio for 0, 3, 6, or 9 jammers with and without 11n spatial multiplexing. For "receiver 1" (identical to the friendly jammer's chipsets) there are only minor differences in the hit ratio behavior. The packet loss (without jammers) for the two spatial streams configuration is slightly higher, and the hit ratio for 3 jammers is slightly lower. This can be explained by a marginally decreased reception ratio of both the receiver and the friendly jammers.

We also show results for a second receiver in Fig. 17c. This receiver is based on an Atheros AR5418 chipset and is denoted as m_{ath} (close to the first receiver in the map, Fig. 16a). The hit ratio of this receiver for 3 jammers is somewhat lower, but comparable for 6 and 9 jammers. Interestingly, the packet loss (0 jammers) is higher for one spatial stream compared to two spatial streams. However, the hit ratio with activated jammers is similar for both receiver configurations.

3) *The Power Amplification Effect:* We revisit the phenomenon that was found in our measurements and which leads to increasing costs when more jammers are enabled. Section VI-C established that this effect actually happens in a real-world network. However, this was demonstrated indirectly

Fig. 18. (a) Example position of two legitimate stations that endure a severe power amplification effect when the right station communicates to the left one. (b) The corresponding decrease in throughput. (c) The corresponding increase in packet loss.

using RSSI to differentiate different stations' positions. In this experiment, we now moved two stations until we found a position, which experienced a significant power amplification phenomenon (cf. Fig. 18a). In fact, we were easily able to find such a constellation in our environment, as predicted by Fig. 3b. By taking measurements directly from these stations, we find that the throughput significantly decreases from over 60 Mb/s to around 40 Mb/s (Fig. 18b) and that the packet loss increases from below 30% to around 45% (Fig. 18c). This finding confirms the power amplification phenomenon and also shows an almost linear increase in the cost on legitimate transmissions when the number of friendly jammers is increased (cf. Fig. 13).

VIII. CONCLUSION

System proposals for friendly jamming applications have received substantial research interest in recent years. This work seeks to bridge the gap from small and static test bed evaluation to a more realistic scenario by investigating friendly jamming at work in a real-world 802.11 network over a three-weeks period. We find three main insights: 1) a rather high number of jammers is required to achieve a high hit ratio for different modulation schemes, 2) the identification and quantification of the power amplification phenomenon which represents a shadow cost of friendly jamming, and 3) the effectiveness of potential jammer collaboration schemes in achieving a good trade-off between jamming performance and cost. We extended our analysis to measurements on 5GHz and IEEE 802.11n and found that the key insight of a trade-off between jamming effectiveness and invasiveness carry over.

Besides the rather black-box oriented measurement results on friendly jamming, we also used simulations of the reception process of superimposed signals to investigate the question why perfect jamming is not always feasible. The analysis for DSSS revealed a strong dependence of the jamming success on the chip alignment between attack frame and jam frame(s). This insight opens up another interesting opportunity for future research: can friendly jammers boost their hit ratios by sending jam signals that are synchronized with the attack signal? Another interesting direction would be to quantify the effect if the friendly jammers could control the carrier frequency offset between their jamming signals and the attack frame.

REFERENCES

- [1] D. L. Adamy, *Ew 101: A First Course in Electronic Warfare*, vol. 1. Norwood, MA, USA: Artech House, 2000.
- [2] E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa, "On the performance of IEEE 802.11 under jamming," in *Proc. IEEE INFOCOM*, Apr. 2008, pp. 1265–1273.
- [3] D. S. Berger, F. Gringoli, N. Facchi, I. Martinovic, and J. Schmitt, "Gaining insight on friendly jamming in a real-world IEEE 802.11 network," in *Proc. ACM WiSec*, 2014, pp. 105–116.
- [4] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, "An IEEE 802.11a/g/p OFDM receiver for GNU radio," in *Proc. ACM SIGCOMM SRIIF*, Hong Kong, Aug. 2013, pp. 9–16.
- [5] J. Brown, I. E. Bagci, A. King, and U. Roedig, "Defend your home!: Jamming unsolicited messages in the smart home," in *Proc. ACM HotWiSec*, New York, NY, USA, 2013, pp. 1–6.
- [6] M. Debbah. (2004). *Short Introduction to OFDM*, accessed on Jun. 21, 2016. [Online]. Available: <http://goo.gl/Zoi2ID>
- [7] P. Dutta, S. Dawson-Haggerty, Y. Chen, C.-J. M. Liang, and A. Terzis, "Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless," in *Proc. ACM SenSys*, 2010, pp. 1–14.
- [8] European Telecommunications Standards Institute. (Aug. 2012). *EN 300 328 V1.8.1 Electromagnetic Compatibility and Radio Spectrum Matters (ERM); Wideband Transmission Systems; Data Transmission Equipment Operating in the 2.4 GHz ISM Band and Using Wide Band Modulation Techniques; Harmonized EN Covering the Essential Requirements of Article 3.2 of the R&TTE Directive*. [Online]. Available: <http://goo.gl/G8ksx8>
- [9] FCC. *Jamming Prohibition*, accessed on Feb. 13, 2014. [Online]. Available: <http://goo.gl/Z3Nvak>
- [10] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh, "Efficient network flooding and time synchronization with glossy," in *Proc. 10th Int. Conf. Inf. Process. Sensor Netw. (IPSN)*, Apr. 2011, pp. 73–84.
- [11] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu, "They can hear your heartbeats: Non-invasive security for implantable medical devices," in *Proc. ACM SIGCOMM*, 2011, pp. 2–13.
- [12] F. Gringoli and L. Nava. *OpenFWFW: Open Firmware for Wi-Fi Networks*, accessed on Jun. 21, 2016. [Online]. Available: <http://www.ing.unibs.it/>
- [13] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "802.11 with multiple antennas for dummies," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 1, pp. 19–25, Jan. 2010.
- [14] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Standard 802.11-2012, Mar. 2012.
- [15] M. L. Jorgensen, B. R. Yanakiev, G. E. Kirkelund, P. Popovski, H. Yomo, and T. Larsen, "Shout to secure: Physical-layer wireless security with known interference," in *Proc. IEEE GLOBECOM*, Nov. 2007, pp. 33–38.
- [16] Y. S. Kim and P. Tague, "Proximity-based wireless access control through considerate jamming," in *Proc. ACM SPME*, 2014, pp. 19–24.
- [17] Y. S. Kim, P. Tague, H. Lee, and H. Kim, "Carving secure Wi-Fi zones with defensive jamming," in *Proc. ACM ASIACCS*, 2012, pp. 53–54.
- [18] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. IEEE ICNP*, Oct. 2004, pp. 252–261.
- [19] J. Lee *et al.*, "An experimental study on the capture effect in 802.11a networks," in *Proc. ACM WinTECH*, New York, NY, USA, 2007, pp. 19–26.
- [20] Z. Li, W. Xu, R. Miller, and W. Trappe, "Securing wireless systems via lower layer enforcements," in *Proc. ACM WiSe*, 2006, pp. 33–42.
- [21] G. Lin and G. Noubir, "On link layer denial of service in data wireless LANs," *Wireless Commun. Mobile Comput.*, vol. 5, no. 3, pp. 273–284, May 2005.
- [22] C.-H. Liu, "On the design of OFDM signal detection algorithms for hardware implementation," in *Proc. IEEE GLOBECOM*, Dec. 2003, pp. 596–599.
- [23] I. Martinovic, N. Gollan, and J. B. Schmitt, "Firewalling wireless sensor networks: Security by wireless," in *Proc. IEEE LCN*, Oct. 2008, pp. 770–777.
- [24] I. Martinovic, P. Pichota, and J. B. Schmitt, "Jamming for good: A fresh approach to authentic communication in WSNs," in *Proc. ACM WiSec*, New York, NY, USA, Mar. 2009, pp. 161–168.
- [25] A. Mpitzopoulos, D. Gavalas, C. Konstantopoulos, and G. Pantziou, "A survey on jamming attacks and countermeasures in WSNs," *IEEE Commun. Surveys Tuts.*, vol. 11, no. 4, pp. 42–56, 4th Quart., 2009.
- [26] K. Pelechrinis, M. Iliofotou, and S. Krishnamurthy, "Denial of service attacks in wireless networks: The case of jammers," *IEEE Commun. Surveys Tuts.*, vol. 13, no. 2, pp. 245–257, May 2011.
- [27] K. Pelechrinis, C. Koufogiannakis, and S. V. Krishnamurthy, "Gaming the jammer: Is frequency hopping effective?" in *Proc. IEEE 7th Int. Symp. WiOPT*, Jun. 2009, pp. 1–10.
- [28] R. A. Poisel, *Modern Communications Jamming: Principles and Techniques*. Norwood, MA, USA: Artech House, 2011.

- [29] C. Pöpper, N. O. Tippenhauer, B. Danev, and S. Capkun, "Investigation of signal and message manipulations on the wireless channel," in *ESORICS*, Leuven, Belgium: Springer-Verlag, 2011, pp. 40–59.
- [30] D. R. Raymond and S. F. Midkiff, "Denial-of-service in wireless sensor networks: Attacks and defenses," *IEEE Pervasive Comput.*, vol. 7, no. 1, pp. 74–81, Jan./Mar. 2008.
- [31] P. Salvador, V. Mancuso, P. Serrano, F. Gringoli, and A. Banchs, "VoIPiggy: Analysis and implementation of a mechanism to boost capacity in IEEE 802.11 WLANs carrying VoIP traffic," *IEEE Trans. Mobile Comput.*, vol. 13, no. 7, pp. 1640–1652, Jul. 2014.
- [32] S. Sankararaman *et al.*, "Optimization schemes for protective jamming," *Mobile Netw. Appl.*, vol. 19, no. 1, pp. 45–60, Feb. 2014.
- [33] N. Santhapuri, S. Nelakuditi, and R. R. Choudhury, "On spatial reuse and capture in ad hoc networks," in *Proc. IEEE WCNC*, Mar./Apr. 2008, pp. 1628–1633.
- [34] A. Sheikholeslami, D. Goeckel, H. Pishro-Nik, and D. Towsley, "Physical layer security from inter-session interference in large wireless networks," in *Proc. IEEE INFOCOM*, Mar. 2012, pp. 1179–1187.
- [35] W. Shen, P. Ning, X. He, and H. Dai, "Ally friendly jamming: How to jam your enemy and maintain your own wireless connectivity at the same time," in *Proc. IEEE Symp. SP*, May 2013, pp. 174–188.
- [36] E. Sourour, H. El-Ghoroury, and D. McNeill, "Frequency offset estimation and correction in the IEEE 802.11a WLAN," in *Proc. IEEE VTC*, vol. 7, Sep. 2004, pp. 4923–4927.
- [37] N. O. Tippenhauer, L. Malisa, A. Ranganathan, and S. Capkun, "On limitations of friendly jamming for confidentiality," in *Proc. IEEE Symp. SP*, May 2013, pp. 160–173.
- [38] M. Vanhoef and F. Piessens, "Advanced Wi-Fi attacks using commodity hardware," in *Proc. ACM ACSAC*, 2014, pp. 256–265.
- [39] J. P. Vilela, P. C. Pinto, and J. Barros, "Position-based jamming for enhanced wireless secrecy," *IEEE Trans. Inf. Forensics Security*, vol. 6, no. 3, pp. 616–627, Sep. 2011.
- [40] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler, "Exploiting the capture effect for collision detection and recovery," in *Proc. IEEE EmNetS*, May 2005, pp. 45–52.
- [41] M. Wilhelm, V. Lenders, and J. B. Schmitt, "On the reception of concurrent transmissions in wireless sensor networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 12, pp. 6756–6767, Dec. 2014.
- [42] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "WiFire: A firewall for wireless networks," in *Proc. ACM SIGCOMM*, 2011, pp. 456–457.
- [43] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "Short paper: Reactive jamming in wireless networks: How realistic is the threat?" in *Proc. ACM WiSec*, 2011, pp. 47–52.
- [44] A. D. Wood and J. A. Stankovic, "Denial of service in sensor networks," *Computer*, vol. 35, no. 10, pp. 54–62, Oct. 2002.
- [45] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," in *Proc. ACM MobiHoc*, 2005, pp. 46–57.
- [46] X. Zhou and M. R. McKay, "Physical layer security with artificial noise: Secrecy capacity and optimal power allocation," in *Proc. IEEE ICSPCS*, Sep. 2009, pp. 1–5.

Daniel S. Berger received the B.Sc. and M.Sc. degrees from the Distributed Computer Systems Laboratory, Kaiserslautern University of Technology, in 2011 and 2013, respectively, where he is currently pursuing the Ph.D. degree in computer science. He was with the German Cancer Research Center from 2008 to 2010.

Francesco Gringoli received the Laurea degree in telecommunications in 1998, and the Ph.D. degree in information engineering in 2002. He has been an Assistant Professor of Telecommunications with the Department of Information Engineering, University of Brescia, Italy, since 2005.

Nicolò Facchi received the master's degree in information technology engineering from the University of Brescia, Italy, in 2012, and the Ph.D. degree in telecommunication engineering in 2016. He is currently a Post-Doctoral Researcher with the Department of Information Engineering, University of Brescia.

Ivan Martinovic received the Ph.D. degree from the Kaiserslautern University of Technology. Before coming to Oxford, he was a Post-Doctoral Researcher with the Security Research Laboratory, University of California, Berkeley, and the Secure Computing and Networking Centre, University of California at Irvine. He is currently an Associate Professor with the Department of Computer Science, University of Oxford.

Jens B. Schmitt received the Ph.D. degree from the Darmstadt University of Technology, Germany, in 2000. Since 2003, he has been heading the Distributed Computer Systems Laboratory. He is currently a Professor of computer science with the Kaiserslautern University of Technology, Germany. His research interests are broadly in performance and security aspects of networked and distributed systems.