Secrecy Beyond Encryption: Obfuscating Transmission Signatures in Wireless Communications

Hanif Rahbari and Marwan Krunz, University of Arizona

Abstract—The privacy of a wireless user and the operation of a wireless network can be threatened by the leakage of transmission signatures, even when encryption and authentication services are employed. In this article, we describe various passive (traffic analysis) and active (jamming) attacks that are facilitated by side-channel information (SCI). Our goal is to highlight the need for novel PHY-layer security techniques that can be used to complement classical encryption methods. We discuss several of these techniques along with advanced hardware that exhibits promising capabilities for countering privacy and SCI-related attacks.

1 Introduction

In 1960, the British Secret Intelligence Service (MI6) was under pressure to break a cipher related to the French position on the issue of Britain's membership in the European Economic Community. The officers were unable to break the code. However, Peter Wright, an MI6 scientist, noticed that the intercepted encrypted telex coming out of the French embassy in London carried a faint secondary signal. This signal turned to be an electromagnetic "echo" of the plaintext message that was being entered to the cipher machine. Wright exploited this signal to disclose the content of the ciphertext without having to break the code. Decades later, *features* of communicated traffic in the form of echoes or footprints of encrypted messages have been widely used to disclose clues about the traffic content, (e.g., the spoken language in a VoIP session).

In computer networks, a given layer in the protocol stack is secured independently of other layers. Encryption is the common way for providing message confidentiality. For example, at the application layer, encryption algorithms and protocols, such as *HTTPS* and *SSH*, provide message confidentiality. At the transport and network layers, the corresponding headers and payloads are encrypted using protocols such as *TLS* and *IPSec*. At the data link (MAC) layer, *WPA2* is used for 802.11 frames. 3G/UMTS and 4G LTE technologies for wireless communications also provide message confidentiality through encryption.

However, even when the payload of any protocol data unit (PDU) is encrypted, packet size, inter-packet times, and various communication features can still be determined by eavesdropping on the physical (PHY) layer frame. Wireless traffic is particularly vulnerable to eavesdropping because of the broadcast nature of wireless communications. For example, the 128-bit AES block cipher used in WPA2 preserves the size of the plaintext and does not impact PHY-layer parameters, such as modulation scheme. At the transmitter (Tx) side, the PHY layer is responsible for receiving the (encrypted) payload from the MAC layer, prepending the required unencrypted PHY header and preamble, converting the entire frame to an analog signal through a channel-dependent modulation, and then transmitting it over the air. Moreover, in many wireless security standards such as WPA2 (802.11i), MAC and PHY headers are not encrypted (see Fig. 1). Therefore, various transmission features remain visible to eavesdroppers. These features include the received-signal strength (RSS), modulation scheme, traffic direction (uplink/downlink), and traffic statistics (e.g., frame size, inter-frame time, data rate). Collectively, these features are referred to as *side-channel information* (SCI).

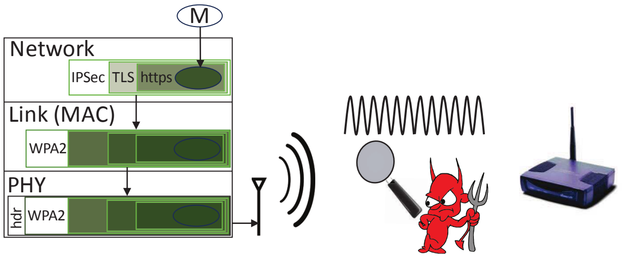

Fig. 1. Encryption of a message (shown in shaded area) at different layers of the protocol stack. An upper-layer packet is considered the payload for the next lower-layer.

In this article, we explain how adversaries can exploit SCI of encrypted wireless traffic to launch various attacks against user privacy and functionality of a practical wireless network. We then discuss some PHY-layer solutions that have been proposed to counter such attacks and complement conventional message encryption at upper layers.

2 SCI-based Attacks in Wireless Networks

In this section, we present two types of SCI-enabled attacks: passive and active attacks. Passive attacks refer to SCI analysis performed by an eavesdropper (Eve) to disclose some private information about a user. Active attacks refer to *selective* jamming of specific packets or parts of a packet, where "significance" is determined based on leaked SCI. The mechanisms for acquiring and analyzing SCI will be discussed in Section 3.

2.1 Passive (Privacy) Attacks

The privacy of a wireless user can be violated by overhearing and analyzing encrypted traffic at the PHY layer. We categorize the types of leaked information and privacy violation into two groups:

Device identification and user tracking– An eavesdropper can fingerprint a wireless device or its user by exploiting device identifiers embedded in unencrypted headers, the device's intrinsic signature, or captured SCI. Using a device's fingerprint, the adversary can easily track the user's geographical location or determine his online activity. For example, *Snoopy* is a software program that can be deployed on a low-altitude flying drone to track users based on their fingerprints, steal their confidential information, or launch man-in-the-middle attack by spoofing already trusted access points. It does not require a visual sensor; instead, it uses an antenna to observe WiFi encrypted communications.

Background activities of installed Apps on a smart phone/tablet or the specific implementation of its wireless card driver can be used to construct a fingerprint. For instance, Eve can create a device-specific traffic fingerprint by analysing only the SCI of background software activities on a 3G smartphone for 6 hours [1]. This is because more than 70% of a smartphone's traffic is independent of user interactions and depends only on installed Apps. In fact, by monitoring 15 minutes worth of traffic bursts, it is possible to identify a particular device with 90% success rate among 20 devices running different sets of Apps [1]. Similarly, traffic statistics can characterize an 802.11 device with high probability. Apart from Apps/user-generated traffic, different vendors often have different protocol implementations and data rate distributions on their wireless cards that result in vendor-specific inter-frame times, medium access wait (backoff) times, and transmission times [2]. Together, these parameters constitute a vendor-specific fingerprint of the device.

Beside traffic statistics, hardware-specific and electromagnetic characteristics of an RF emitter form a "radiometric" identity of a particular Tx. The analog components of a wireless card's transmit path (e.g., oscillator, baseband filter, amplifier, and antenna) exhibit inherent manufacturing impairments that differ from one card to another. Small variations in these components create distinct artifacts in the emitted signal (e.g., frequency offset and amplitude clipping). The distortions in the captured modulation symbols due to hardware impairments can be exploited to detect a signal's device-of-origin [3].

User's activities and browsing interests– An eavesdropper can also exploit SCI to discern the online activities of a user, his interests, or his search queries. For example, through captured SCI, Eve can identify not only the website that a user is browsing, but also the currently active page within a specific website. A typical website is characterized by a nominal uplink/downlink traffic volume and duration. These coarse-grain traffic features are sufficient to classify websites [4]. Even within a given website in which different pages are designed for different users, analyzing the packets size distribution allows for identifying a specific page. As a result, the attacker may be able to conclude the user's product of interest and may overwhelm him with many commercial ads.

The leakage of private information is not limited to online browsing. An adversary can determine with 80% accuracy the type of user activity (gaming, video streaming, Skype, browsing, etc.) by only eavesdropping for 5 seconds on an 802.11 WLAN traffic [5]. Differences between the traffic statistics of different applications are often large enough to distinguish these applications. Further, the adversary can find out the user's specific actions during an activity, such as posting a status on Facebook or opening a chat window in Gmail, based on the statistics of the sequence of packets generated by the user. Along the same lines, tracking the traffic of two users can reveal if they are communicating with each other.

The sizes (in bytes) and directionality (uplink/downlink) of a sequence of packets exchanged between a mobile user and an access point can also reveal what the user is searching for. Google, Bing, and other search engines provide users with suggestions for a searched phrase, i.e., *auto-suggestion* feature. When a user types the first letter of a keyword, the search engine quickly responds with a list of suggested words. Typing the second letter updates the list of suggestions, and so on. The size of the packet that contains the list of suggestions is highly correlated with the typed letters [6]. An adversary can construct a table of different keywords and associate them with the sizes of per-keystroke suggested lists. He can then match the sizes of an observed sequence of packets to one of the entries in the table, and determine the queried word [6]. Finally, the message length and the language used in an encrypted instant messaging application can be determined based on packet sizes only.

2.2 Active Attacks (Selective Jamming)

Besides breaching user privacy, SCI can be used by malicious users to disrupt wireless communications by selectively jamming transmissions and preventing correct decoding at the receiver (Rx). Jamming includes random attacks, persistent attacks (barrage jamming), and smart/selective attacks in which only a certain packets or parts of a particular packet are jammed. In selective (reactive) jamming, a targeted packet (or parts of) is selected based on the amount of disruption caused by not delivering this packet to its intended Rx. For example, TCP Acknowledgement (ACK) packets are much shorter in duration than TCP data packets, but are critical for maintaining high TCP throughout by preventing a significant reduction in the congestion window size. Jamming these packets not only requires less energy than jamming a data packet, but also can deceive the TCP sender into thinking that the last data packet was not successfully received due to network congestion. Consequently, the source may unnecessarily reduce its packet transmission rate and retransmit the last packet, which was already received correctly. The attacker can identify the TCP ACK by analyzing the sequence of inter-arrival times and packet sizes. In the case of link-layer ACK packets, the packet type can also be identified by inspecting the unencrypted MAC header (see Fig. 2).

The unencrypted PHY-layer header (Fig. 2) can be used to detect and jam data packets transmitted at high rates. Wireless devices adapt their transmission rates based on

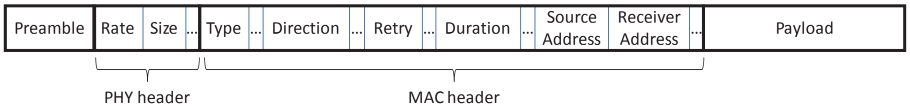

Fig. 2. Typical 802.11 frame preamble, PHY header, and MAC header.

channel conditions. A good channel prompts the Tx to use a higher-order modulation scheme, hence a higher data rate. When a packet is not successfully received, the Tx attributes that to channel conditions and accordingly retransmits the packet at a lower rate. This can be exploited by the attacker to jam only high data-rate packets, making the Tx mistakenly reduce its rate and waste communication resources.

In addition to the PHY header, the payload's modulation scheme may disclose the data rate. Furthermore, the frame preamble can be exploited to detect the arrival of a packet and launch reactive attacks. The frame preamble is a publicly known signal, prepended to the beginning of a frame to help the Rx detect the frame and estimate various communication parameters (e.g., frequency offset, channel response). Correct decoding of a frame depends on correct estimation of these parameters. Once a frame is detected using the publicly known preamble, an attacker can jam a vulnerable part of the preamble to disrupt the parameter estimation functions at the Rx [7].

3 SCI Extraction and Analysis Methods

In this section, we discuss the techniques used to acquire and process SCI.

3.1 Extraction of Traffic Attributes

Unencrypted PHY and MAC headers are the main sources for acquiring traffic parameters. At the packet level, the PHY header contains the packet size, transmission rate, and the modulation scheme fields. Parameters such as source and destination MAC addresses, direction of the packet, and packet type (e.g., a retransmission) are specified in the MAC header (see Fig. 2).

In addition to header fields, SCI can be used to extract certain packet parameters and radiometric features. For example, the modulation scheme used for the frame payload reveals the packet size and the data rate. In digital communications, a bit sequence is modulated into symbols before transmission over the air. The number of possible symbols of a modulation scheme (known as modulation order) relates to the number of bits that can be modulated into a single symbol. Because of channel noise, symbols must be sufficiently separated so that the Rx can distinguish them from one another. Consequently, a more noisy channel can support fewer bits per symbol, and the transmission of a fixed-size payload can take different durations under different channel conditions. By measuring the frame duration (in seconds) and detecting the modulation scheme, an adversary can estimate the packet size (in bytes).

Flow-level parameters and statistical distributions can be calculated based on packet-level parameters. These include (but not limited to) the total traffic volume in each direction and the number of unique packets. To reduce the effect of packet collisions on traffic statistics, retransmission packets are identified and removed from the observation set. When normalized to the total session duration, the traffic volume may provide a reliable data rate statistic, despite variations in the instantaneous data rate.

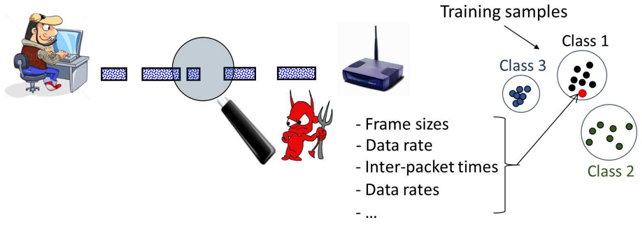

Fig. 3. Traffic classification based on the training samples of three known classes. The closest class to the observed features set is likely the correct class.

3.2 Analysis of Traffic Attributes

Although acquiring traffic attributes is often sufficient to launch selective active attacks, classification-based passive attacks require further traffic analysis. Supervised machine learning is the main approach used for classification and traffic identification. In this approach, features of known traffic types, e.g., specific websites or activities, are used to train a classifier. These features include, for example, frame size, traffic direction, inter-packet times, etc. (see Fig. 3).

Different types of classifiers can be employed to identify the class of a features set, including (multinomial) naive Bayes, support vector machine (SVM), k-nearest neighbor algorithm, decision tree, neural networks, and hidden Markov models (HMM). Samples belonging to different known classes are used to train the classifier. Once the attributes of an unknown traffic have been extracted, the classifier tries to find the most similar traffic type to the observed features. SVM and HMM usually provide better classification accuracy than others [5].

4 Prevention and Remedies

Several defense mechanisms have been proposed to prevent SCI-enabled passive and active attacks. Some countermeasures obfuscate SCI to distort traffic statistics. Others employ a PHY-layer specific approach, whereby potential eavesdroppers are deafened through friendly jamming (FJ) to prevent correct decoding of unencrypted headers. SCI obfuscation through PHY-level cryptography has also been considered. In this section, we first discuss the limitations of PHY-layer encryption and then present other solutions. Note that although spreading techniques (FHSS and DSSS), often used to combat narrow-band and pulse jamming attacks, can be used against non-selective and random jamming, other preventive approaches are needed to counter SCI-enabled selective attacks. Moreover, spread spectrum techniques used in common wireless standards (e.g., 802.11b/g) rely on known spreading patterns, and hence cannot prevent the leakage of SCI.

4.1 Limitations of Header Encryption

Given that a significant amount of SCI is leaked through PHY/MAC headers, a natural question to ask is “Why not encrypt these headers?” However, this is usually not a viable option for the following reasons:

1) **Transmitter authentication**: Encryption is based on a shared secret key. In a network of nodes, different pairs of nodes establish distinct keys for different sessions during the association process at the MAC layer. Session participants are identified by globally unique MAC addresses. Each node maintains a table of session keys that are associated with the MAC addresses of the participants of each session. This means that before decoding the MAC addresses in an incoming frame, a node does not know the sender and intended receiver of that frame; hence, it cannot immediately look up the corresponding decryption key. Instead of the MAC address, the Tx-Rx channel or other radiometric features can be used as a PHY-level identifier to look up the key. However, mobility and inaccuracy of low-end RF receivers limit the applicability of these identifiers [8].
2) **Broadcast operation**: Certain fields in the header are to be broadcasted to every node in the vicinity of the transmitting node (e.g., the “duration” field in the 802.11 MAC header). In a multi-user environment, while a user is transmitting, other users should remain silent to avoid collision. Sensing the carrier before a transmission is a common collision avoidance approach, which has been adopted in 802.11 schemes. Devices may also perform virtual carrier sense by overhearing the duration field and updating their network allocation vectors (NAVs) accordingly. Thus, if the MAC header is encrypted, other users cannot overhear the duration field.
3) **Delay and complexity**: The decryption process of an encrypted header incurs additional delay and complexity, especially when block ciphering. Specifically, the Rx needs to set its buffer timer and initiate its demodulator according to PHY-header fields. Delay in decrypting the PHY header may prevent timely operation at the Rx.

Note also that header encryption (if possible) cannot prevent the leakage of certain SCI, such as the modulation scheme.

4.2 Obfuscation of Traffic Features

The methods used to thwart traffic analysis attacks can be divided into two subcategories, based on whether or not the attacker can be prevented from tracking the user.

4.2.1 Identifier concealment

To accurately extract traffic statistics pertaining to a given device, Eve needs to filter out packets of other devices from the set of captured packets. Packets belonging to a user or to the traffic between a user and an AP are identified by their MAC addresses. With *traffic reshaping* [9], several virtual MAC interfaces that have different MAC addresses are configured for the same device. Generated packets are dynamically divided among these interfaces so as to create different traffic patterns on each MAC interface. This prevents Eve from linking the packets to the same sender and measuring the true traffic statistics.

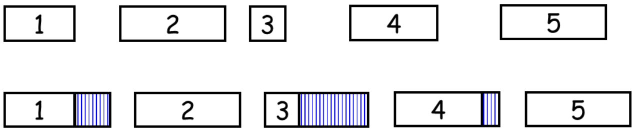

Fig. 4. A simple padding scheme: The packets of the actual traffic (top) are padded to have the same, indistinguishable size and inter-arrival time (bottom).

Another way of decoupling packets that belong to the same traffic flow is to change the MAC address based on a chain of secure identifiers; hence, obfuscating the traffic statistics. In this approach, Tx and Rx agree on a sequence of bogus identifiers, before they start to exchange data packets. *SlyFi* [10] is one such method in which the Tx and Tx true addresses are encrypted together with the elapsed time of the session to generate a set of time-rolling identifiers. A chain of hash values can also be employed to change the MAC addresses on a per-packet basis. These obfuscation methods, however, cannot conceal PHY layer identifiers, including device fingerprints, or even packet-level parameters.

4.2.2 Padding Techniques

SCI can be distorted by appending bogus bits to packet payloads (padding) or adding dummy packets (see Fig. 4). This obfuscates the inter-arrival times, the packet sizes, and the number of packets (hence, total traffic volume). Different methods have been proposed to calculate amount of padding needed to efficiently hide the type of the underlying traffic. For example, traffic morphing techniques modify the traffic pattern by altering packet sizes and inserting dummy packets. However, these techniques can be extremely inefficient, incurring traffic overhead as high as 400% of the original traffic volume (for example, in defending against a website identification attack [4]).

4.3 Eavesdropper Deafening

The challenges of PHY and MAC header encryption along with the overhead and limitations of traffic feature obfuscation techniques necessitate a complementary PHY-layer approach based on friendly jamming. Jamming involves transmitting a noisy signal that interferes with the data signal, making it undecodable. It is typically considered as an adversarial act against a legitimate Rx; however, it can be employed to degrade the eavesdropper’s channel without affecting the legitimate reception.

4.3.1 Friendly Jamming

The idea of friendly jamming (FJ) is that when two identical signals of opposite signs arrive simultaneously at the Rx, they will cancel out (nullify) each other. If each friendly jammer knows the phase and amplitude of its signal at the Rx (i.e, by knowing its channel to the legitimate Rx), then several friendly jammers can cooperatively adjust their signals such that they are collectively nullified only at the Rx. This idea can be generalized to multiple jammers or multiple antennas at a node, i.e., MIMO jammer (see Fig. 5).

MIMO-capable friendly jammers can be placed in various locations with respect to the Tx and Rx. For example,

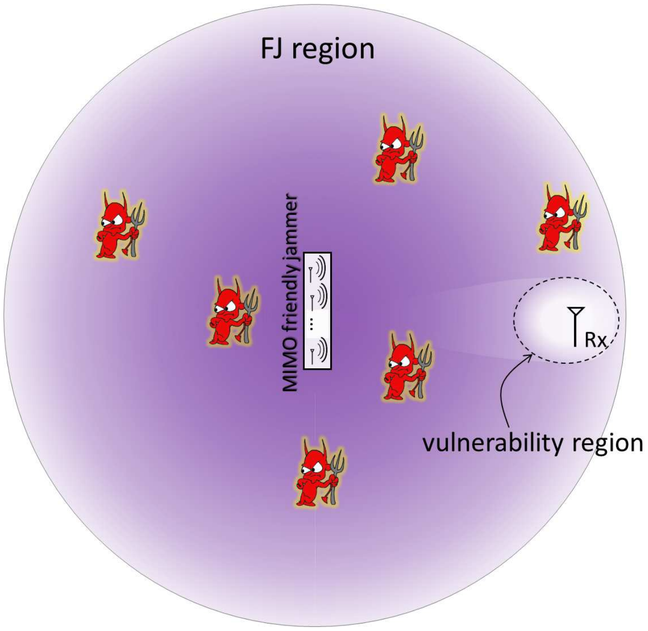

Fig. 5. FJ region of a MIMO friendly jammer. The FJ signals are nullified at and (sometimes up to several wavelengths) around the legitimate Rx and along the LOS direction. In other places, eavesdroppers experience high jamming.

when FJ is generated by the Tx of the information signal, it is called *Tx-based* FJ. Likewise, the placement of FJ close to the legitimate Rx results in *Rx-based* FJ. Using *full-duplex* radios, Rx-based FJ is possible even with a single jamming antenna. In full-duplex communications, the self interference caused by a node's own transmission (in this case, a jamming signal) is suppressed during the reception of the information signal. One use-case of Rx-based FJ is to secure the unencrypted communications of implantable medical devices (IMDs) by exploiting Rx-based FJ at the access point [11]. Whenever an IMD sends a signal, the access point receives the IMD's transmission while simultaneously generating a jamming signal. Rx-based FJ can complement Tx-based FJ if the latter fails to deafen the eavesdropper, who may reside in a "vulnerability region" around the Rx. This region contains the set of locations whose CSI is highly correlated with the Rx CSI. As a result, the Tx-based FJ signal is weak in this region. These points can be along the LOS direction and close to the Rx (see Fig. 5).

As a result of superposing the FJ signal, eavesdroppers are unable to decode the PHY and MAC headers. However, in some practical scenarios, Eve may be able to estimate and remove the FJ signal from the received signal. Specifically, wireless systems often rely on transmitting a publicly known signal (e.g., preamble) in each frame to be used for various purposes, such as channel estimation, frame detection, and frequency offset estimation. A MIMO-capable eavesdropper can estimate the Tx-based FJ signal by exploiting one of the known parts of the information signal and then subtracting the estimated FJ from the received signal [12]. Furthermore, if the FJ signal is transmitted on antennas different than the antennas of the information signal, Eve can tune her antennas to receive the same FJ signal at different antennas and then subtract the received combined signal at one antenna from the received signal received at another to remove the FJ signal [13]. Even if the FJ signal is not removed at Eve, hypothesis-test cross-correlation attacks can reveal the content of a header field that takes one of a few known values [14]. The intuition is that a random FJ signal averages out when it is cross-correlated with another independent signal. So Eve may try to correlate the received composite signal against each possible value of the information signal and detect with high confidence the true value of the information field. In addition to capturing such header fields, Eve may extract SCI in the presence of FJ (which is usually a form of additive random noise) by using frame and modulation classification techniques designed for noisy channels. In the case of Rx-based FJ, Eve may use a directional antenna to suppress the FJ signal.

4.3.2 *Friendly CryptoJam*

The robustness of FJ against the aforementioned attacks can be significantly improved if the FJ signal is transmitted from the same antenna as the information signal and is intermixed cryptographically with it, i.e., via stream ciphering. *Friendly CryptoJam* (FCJ) [14] does that by using a secret modulated FJ signal to encrypt the modulated headers of an 802.11 frame. In contrast to classic FJ, this signal is known to the Rx and is not a function of the Tx-Rx channel. Furthermore, FCJ obfuscates the payload's modulation scheme (hence, packet size and data rate) by hiding it in the highest-order modulation scheme.

FCJ combines the jamming signal with the modulated data signal before transmission. Instead of nullifying the FJ signal at the Rx, in FCJ a secret sequence of modulated symbols is generated at both the Tx and Rx, based on a shared secret key. This signal is used for two purposes. First, it encrypts the modulated symbols by securely relocating each modulated symbol within the same constellation map (see Fig. 6). Rx uses the same sequence to recover the original modulated symbols from the received encrypted symbols. Eve cannot decrypt the header because she cannot generate the same FCJ signal without knowing the secret key. Second, the FCJ signal is used to embed the payload's symbols that have been modulated with any one of several available modulation schemes into the constellation map of the highest-order modulation scheme. Different FCJ symbols map the same encrypted data symbol into different locations on the constellation map of the highest-order modulation scheme. In the example in Fig. 6, encrypted symbol 1 can be mapped to 16-QAM symbols 13 and 5 when the corresponding FCJ symbols are 4 and 2, respectively. This embedding tries to preserve the original "distances" between symbols on the constellation map so as to maintain the same performance of the original modulation scheme. From Eve's standpoint, the frame payload will always seem to have been generated using one modulation scheme (16-QAM in the example). As a result, Eve cannot identify the true modulation scheme. The generation of the FJ signal is location-independent and is robust to node mobility, time-synchronization errors, and packet losses. This is achieved by generating this signal on a per-frame basis, based on a per-frame identifier that is embedded in the preamble (without disrupting normal preamble functions). This identifier can also be used for sender identification, in place

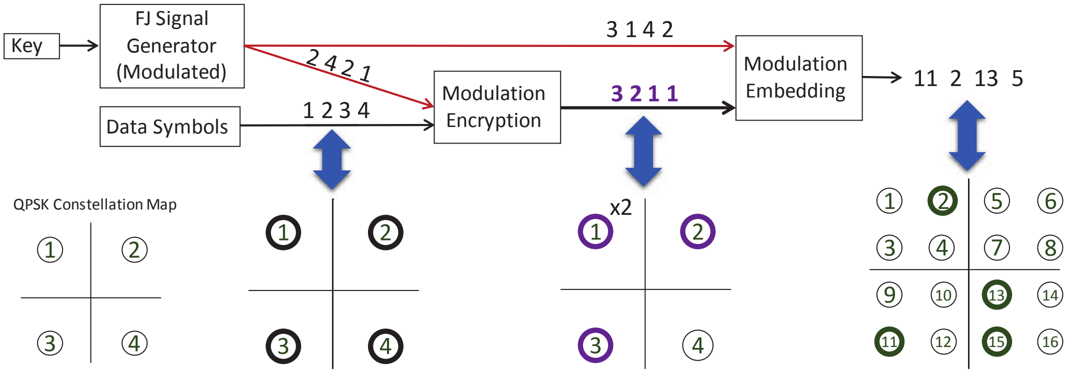

Fig. 6. Example of applying FCJ on a QPSK-modulated signal. The FJ signal is divided into two parts: one for modulation encryption and another for stealthily embedding the modulated QPSK symbols into a 16-QAM modulation scheme. The numbers represent the decimal values of modulated symbols, shown with thick circles on the constellation maps.

of the encrypted Tx MAC address. Changing the jamming signal on a per-frame basis also prevents Eve from detecting retransmitted frames.

5 Conclusions and Future Directions

SCI leaked from encrypted wireless communications can be exploited to violate user privacy using various traffic analysis techniques. Moreover, software-defined radios (SDRs) have been used to experimentally demonstrate SCI-enabled reactive jamming attacks. In particular, it has recently been shown that FPGA implementation of a reactive jammer can achieve extremely fast reaction time [15]. On the other hand, emerging (MU-)MIMO and 5G systems and upcoming 802.11 standards, such as 802.11ai/aq/ax, have not yet incorporated PHY-layer security in their designs for new application trends and still leave the header fields, modulation schemes, spreading patterns, and management frames exposed to eavesdroppers. To prevent the leakage of SCI, PHY-layer technologies (including full-duplex MIMO Tx/Rx) and joint design of FJ and cryptography (e.g., FCJ) have been proposed. SDRs have also been used to implement Rx/Tx-based FJ and FCJ security schemes, and to demonstrate the limitations of FJ. However, so far these techniques cannot completely prevent the leakage of SCI without incurring high overhead, which calls for further research in this area.

Acknowledgements

This research was supported in part by the Army Research Office (grant # W911NF-13-1-0302) and in part by the National Science Foundation (grants # IIP-1265960 and CNS-1409172). Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the author(s) and do not necessarily reflect the views of ARO or NSF.

References

[1] T. Stöber, M. Frank, J. Schmitt, and I. Martinovic, “Who do you sync you are?: Smartphone fingerprinting via application behaviour,” in *Proc. Sixth ACM Conf. Security and Privacy in Wireless and Mobile Networks (WiSec’13)*, Budapest, Hungary, 2013, pp. 7–12.

[2] C. Neumann, O. Heen, and S. Onno, “An empirical study of passive 802.11 device fingerprinting,” in *32nd IEEE Int. Conf. Distributed Computing Syst. Workshops (ICDCSW’12)*, June 2012, pp. 593–602.

[3] V. Brik, S. Banerjee, M. Gruteser, and S. Oh, “Wireless device identification with radiometric signatures,” in *Proc. 14th ACM Int. Conf. Mobile Computing and Networking (MobiCom’08)*, San Francisco, California, USA, 2008, pp. 116–127.

[4] K. Dyer, S. Coull, T. Ristenpart, and T. Shrimpton, “Peek-a-boo, I still see you: Why efficient traffic analysis countermeasures fail,” in *Proc. IEEE Symp. Security and Privacy (SP’12)*, May 2012, pp. 332–346.

[5] F. Zhang, W. He, X. Liu, and P. G. Bridges, “Inferring users’ online activities through traffic analysis,” in *Proc. Fourth ACM Conf. Wireless Network Security (WiSec’11)*, Hamburg, Germany, 2011, pp. 59–70.

[6] S. Chen, R. Wang, X. Wang, and K. Zhang, “Side-channel leaks in web applications: A reality today, a challenge tomorrow,” in *Proc. IEEE Symp. Security and Privacy (SP’10)*, May 2010, pp. 191–206.

[7] H. Rahbari, M. Krunz, and L. Lazos, “Security vulnerability and countermeasures of frequency offset correction in 802.11a systems,” in *Proc. IEEE INFOCOM’14*, Toronto, ON, Canada, April 2014, pp. 1015–1023.

[8] S. U. Rehman, K. W. Sowerby, and C. Coghill, “Analysis of impersonation attacks on systems using RF fingerprinting and low-end receivers,” *J. Computer and System Sciences*, vol. 80, no. 3, pp. 591 – 601, 2014, special Issue on Wireless Network Intrusion.

[9] F. Zhang, W. He, and X. Liu, “Defending against traffic analysis in wireless networks through traffic reshaping,” in *31st IEEE Int. Conf. Distributed Computing Syst. (ICDCS’11)*, June 2011, pp. 593–602.

[10] B. Greenstein, D. McCoy, J. Pang, T. Kohno, S. Seshan, and D. Wetherall, “Improving wireless privacy with an identifier-free link layer protocol,” in *Proc. 6th Int. Conf. Mobile Syst., Appl., and Services*, Breckenridge, CO, USA, 2008, pp. 40–53.

[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 Conf.*, Toronto, Ontario, Canada, Aug. 2011, pp. 2–13.

[12] M. Schulz, A. Loch, and M. Hollick, “Practical known-plaintext attacks against physical layer security in wireless MIMO systems,” in *Proc. Network and Distributed Syst. Security Symp. (NDSS’14)*, February 2014.

[13] N. Tippenhauer, L. Malisa, A. Ranganathan, and S. Capkun, “On limitations of friendly jamming for confidentiality,” in *Proc. IEEE Symp. Security and Privacy (SP’13)*, May 2013, pp. 160–173.

[14] H. Rahbari and M. Krunz, “Friendly CryptoJam: A mechanism for securing physical-layer attributes,” in *Proc. ACM Conf. Security and Privacy in Wireless & Mobile Networks (WiSec’14)*, Oxford, United Kingdom, July 2014, pp. 129–140.

[15] D. Nguyen, C. Sahin, B. Shishkin, N. Kandasamy, and K. R. Dandekar, “A real-time and protocol-aware reactive jamming framework built on software-defined radios,” in *Proc. 2014 ACM Workshop Software Radio Implementation Forum*, Chicago, Illinois, USA, 2014, pp. 15–22.

BIOGRAPHIES

HANIF RAHBARI (rahbari@email.arizona.edu) is currently an electrical and computer engineering Ph.D. candidate at the University of Arizona. He has received his BSc in information technology from Sharif University of Technology and his MSc in computer networks from AmirKabir University of Technology, Tehran, Iran. His research interests include wireless communications and networking, PHY-layer security, hardware implementation, dynamic spectrum access networks, and multimedia networking.

MARWAN KRUNZ (krunz@email.arizona.edu) (S'93-M'95-SM'04-F'10) received the Ph.D. degree in electrical engineering from Michigan State University in 1995. He is a Professor of ECE and CS at the University of Arizona and the Site Codirector of the NSF Broadband Wireless Access and Applications Center. His research interests are in wireless communications and networking, with emphasis on resource management, adaptive protocols, and security issues. He has published more than 215 journal articles and peer-reviewed conference papers, and is a coinventor on five US patents. He received numerous awards, including the 2012 IEEE TCCC Outstanding Service Award and the NSF CAREER Award. He was an Arizona Engineering Faculty Fellow (2011-2014) and an IEEE Communications Society Distinguished lecturer (2013 and 2014). He served on the editorial boards of several IEEE journals. He was the general and program chair for numerous conferences, including INFOCOM'04, SECON'05, WoWMoM'06, and WiSec'12.

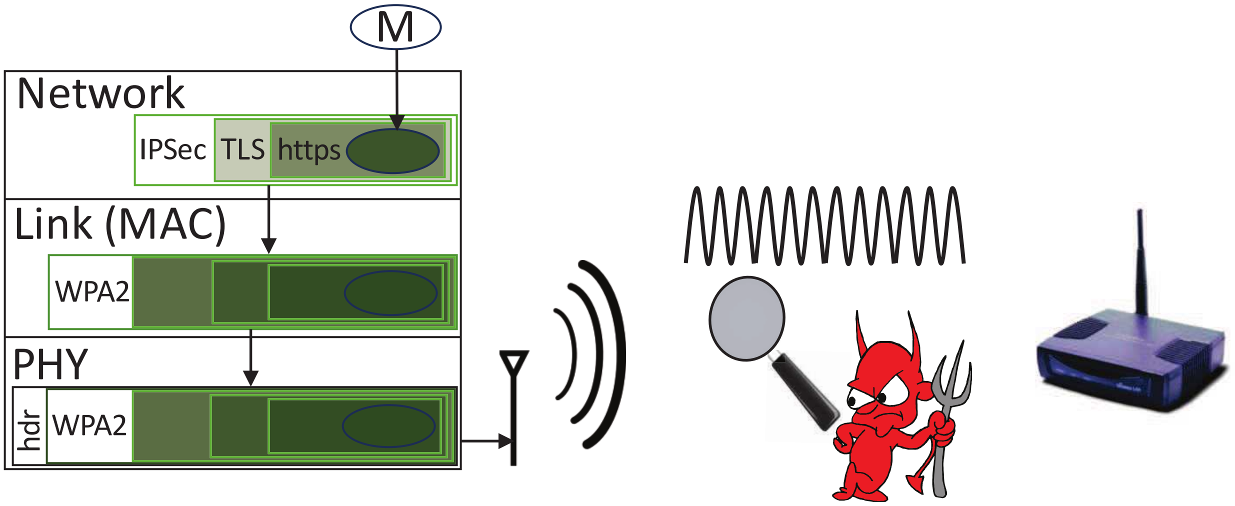

Fig. 7. Encryption of a message (shown in shaded area) at different layers of the protocol stack. An upper-layer packet is considered the payload for the next lower-layer.

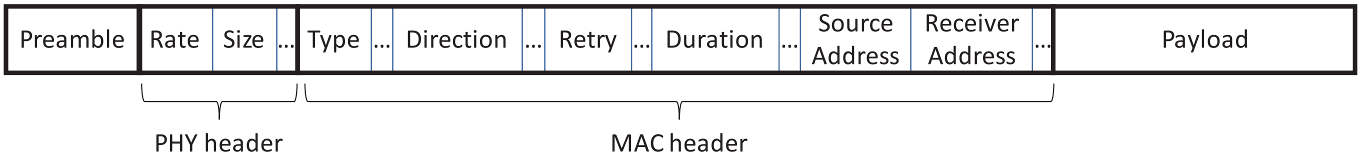

Fig. 8. Typical 802.11 frame preamble, PHY header, and MAC header.

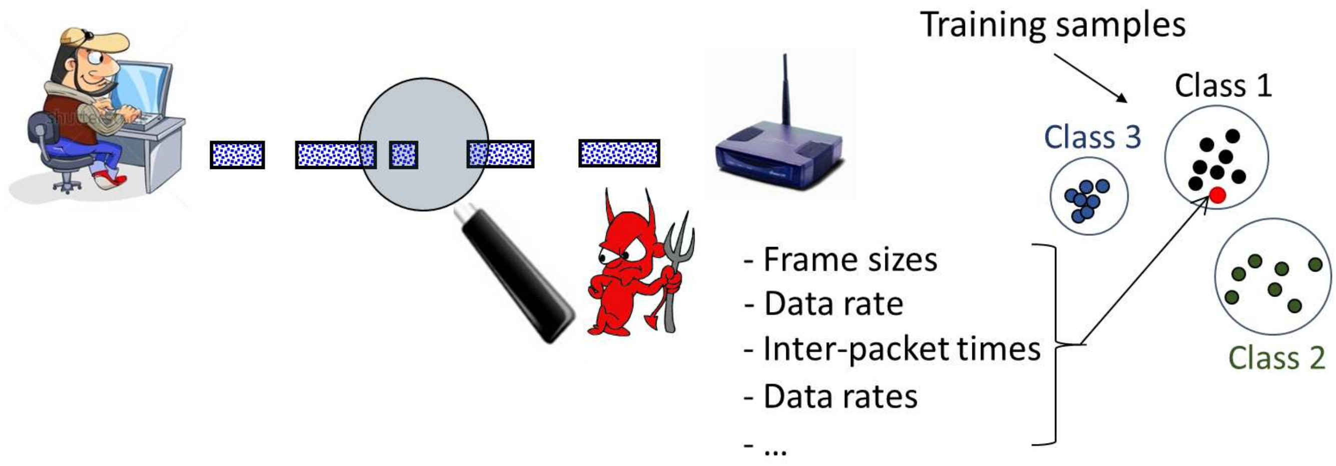

Fig. 9. Traffic classification based on the training samples of three known classes. The closest class to the observed features set is likely the correct class.

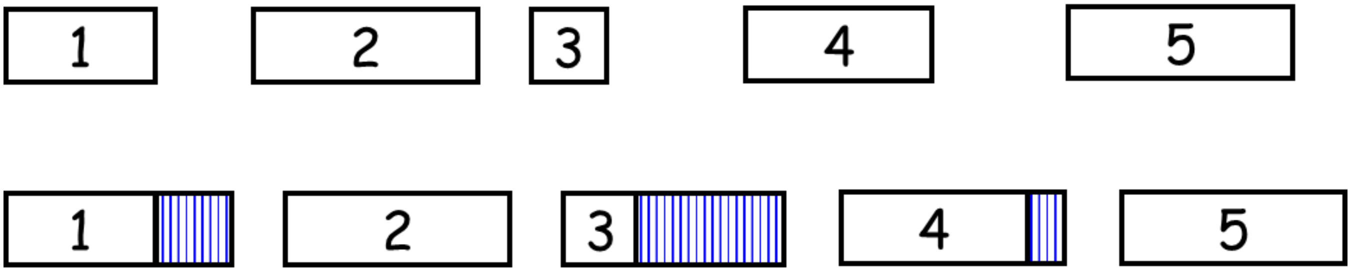

Fig. 10. A simple padding scheme: The packets of the actual traffic (top) are padded to have the same, indistinguishable size and inter-arrival time (bottom).

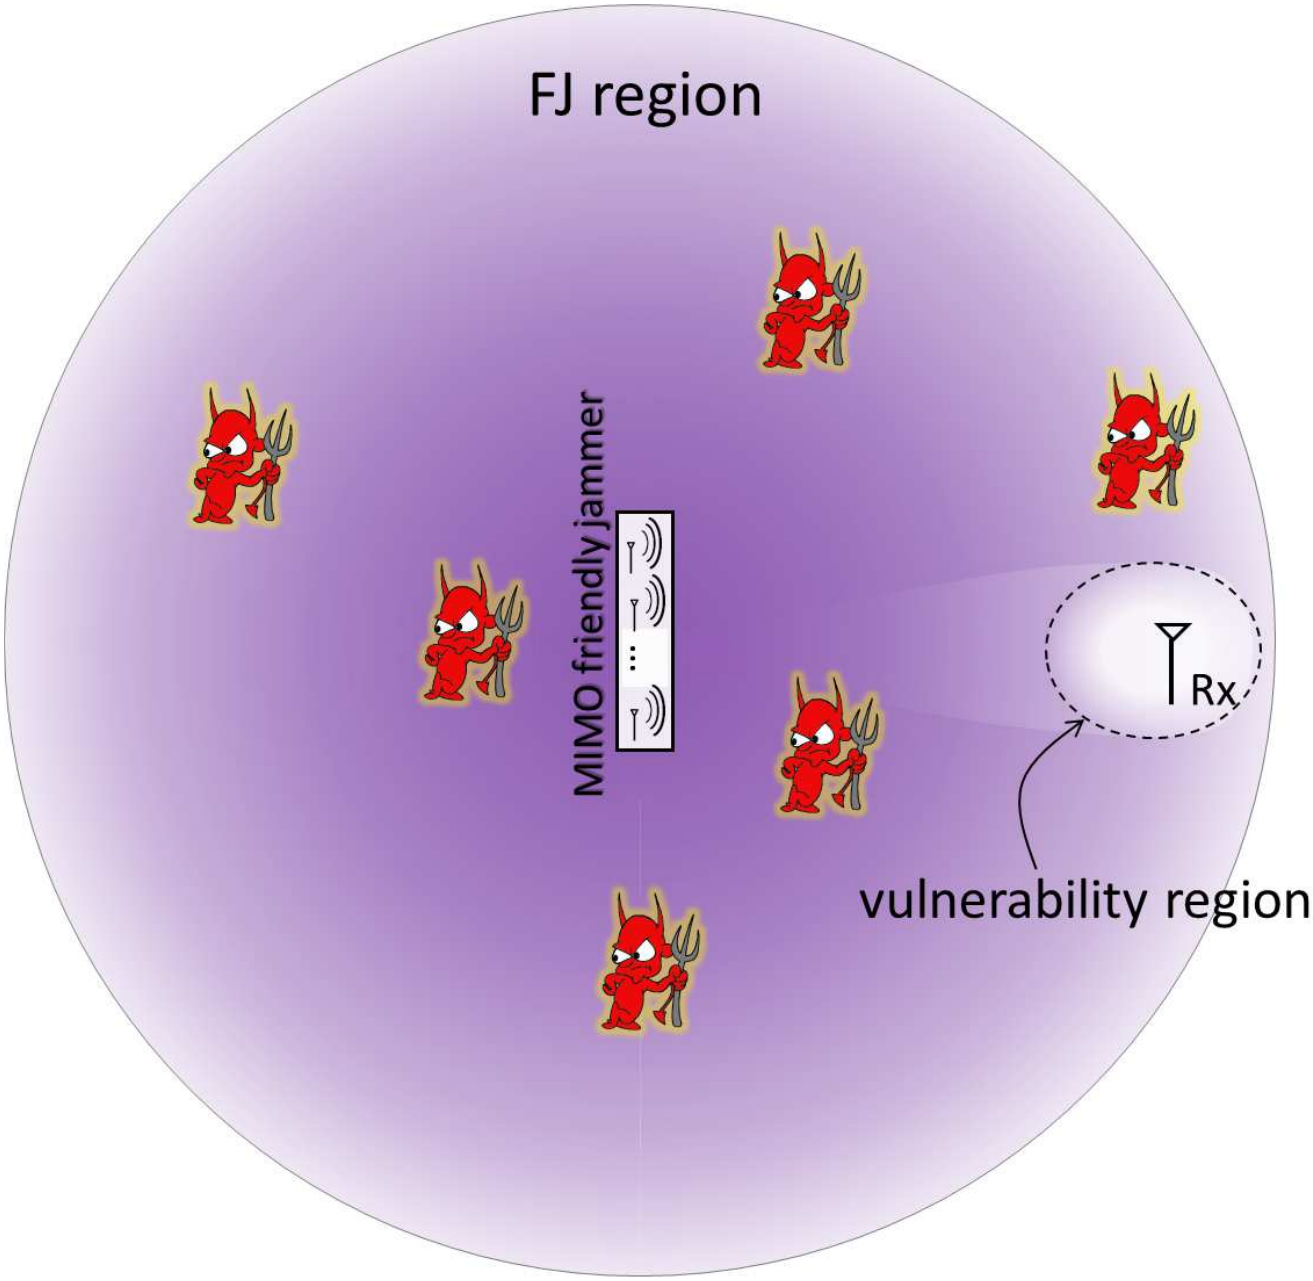

Fig. 11. FJ region of a MIMO friendly jammer. The FJ signals are nullified at and (sometimes up to several wavelengths) around the legitimate Rx and along the LOS direction. In other places, eavesdroppers experience high jamming.

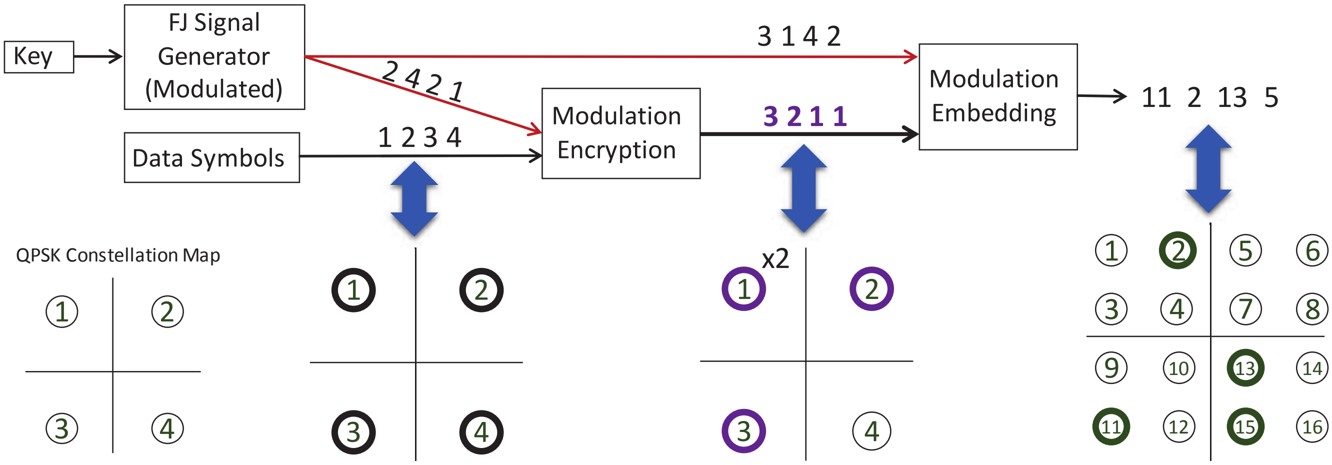

Fig. 12. Example of applying FCJ on a QPSK-modulated signal. The FJ signal is divided into two parts: one for modulation encryption and another for stealthily embedding the modulated QPSK symbols into a 16-QAM modulation scheme. The numbers represent the decimal values of modulated symbols, shown with thick circles on the constellation maps.