A Communications Jamming Taxonomy

Marc Lichtman, Jeffrey D. Poston, SaiDhiraj Amuru, Chowdhury Shahriar, T. Charles Clancy, R. Michael Buehrer, and Jeffrey H. Reed | Virginia Tech

With the now widespread availability of software-defined radio technology for wireless networks, the distinction between jamming in the original electronic warfare sense and wireless cybersecurity attacks becomes hazy. A taxonomy delineates these concepts in the rapidly expanding field of wireless security, classifying communication jammers' theoretical behaviors and characteristics.

The wireless medium's inherent openness makes it susceptible to adversarial attacks. A wireless system's vulnerabilities can be broadly classified based on an adversary's capabilities; for example, a passive adversary might eavesdrop on the wireless channel and try to infer information, an active adversary might transmit energy to jam reliable data transmission, and a higher-layer active adversary might threaten a link's integrity and confidentiality. In this article, we focus on *jamming attacks*, in which attackers transmit signals interfering with victims' communications, principally those at the physical (PHY) layer, intended to cause a denial of service (DoS) and thus compromise a link's availability.

Jammers use a wide range of behaviors to cause DoS; the jamming literature shows numerous jamming models and assumptions (for more information on related work in jamming taxonomies, see the sidebar). These models or behaviors span in complexity from continuous wave interference to sensing and real-time decision making to increase an attack's effectiveness and covertness.

We propose a taxonomy that covers the communications side of jamming (as opposed to radar jamming or attacks on radio navigation).

Background

Research on electronic warfare (EW) and jamming dates to World War II. At this time, jammers were categorized by signal type because each was constructed from distinct radio circuitry. However, in the present era of software-defined radio (SDR), the historical approach unnecessarily restricts the categorization of jamming. Today, the important questions to answer are: What information does the jammer possess? And what's the jammer's capacity to act on this information?

Our jamming taxonomy aims to help researchers place newly discovered jamming or antijamming strategies in a larger context of known strategies in a way that is consistent with modern EW.

Our work's technological theme is similar to the Common Attack Pattern Enumeration and Classification (CAPEC; https://capec.mitre.org)—a catalog and taxonomy of cyberattack patterns created to help build secure software. Each attack pattern provides a challenge that an attacker must overcome, common methods used to overcome that challenge, and recommended methods for mitigating the attack. At the top level, the taxonomy is organized by attack mechanisms (for instance, abuse of functionality, exploitation of authentication, or malicious code execution) and domains (for instance, hardware, software, or social engineering). Although our jamming taxonomy is fundamentally different in structure, CAPEC represents a cybersecurity equivalent.

Jamming versus Cyberattack

An early jamming technique was barrage jamming that, qualitatively, resembled the approaches of early

1540-7993/16/$33.00 © 2016 IEEE Copublished by the IEEE Computer and Reliability Societies

Related Works in Jamming Taxonomies

David L. Adamy's[1] and Richard Poisel's[2] comprehensive references generally reflect the historical tradition of distinguishing jamming by signal type (for example, noise, tone, and pulse). Poisel's work focuses more on communications than Adamy's and includes smart jamming techniques that, in the main text, we call *protocol-aware jamming*. In contrast, we emphasize a jammer's potential behaviors and attributes and discuss the specific jamming techniques characterizing a given behavior. Consequently, we place jammers with one or multiple tones in the same category; they're just expressing different parameters of the same jamming behavior.

"The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks" provides another categorization of jammers.[3] The authors use *constant*, *deceptive*, *random*, and *reactive jammers* for categorization. A constant jammer sends out random bits without following any media access control (MAC)-layer protocols. The deceptive jammer (called *spoofing* in the main text) transmits regular packets into the channel, following the PHY- and MAC-layer protocol the target uses. Random jamming refers to a jammer turning on and off with a random or fixed period. Finally, reactive jamming (called *time-correlated jamming* in the main text) senses the target channel and transmits only when there's activity. Although these categories suit the authors' analysis, they don't distinguish between whether the jammer is adapting its signal based on a priori or acquired information. A similar concern applies to jamming literature surveys such as "Denial of Service Attacks in Wireless Networks: The Case of Jammers."[4]

Most previous work studied only specific aspects of the jamming problem and didn't provide a complete overview of the potential jamming attacks that can be performed depending on the information available to the jammer. In this regard, our taxonomy is not only more comprehensive than those but also unique in that it's based on the information the jammer possesses and the jammer's capacity to act on that information.

References

1. D.L. Adamy, *EW 101: A First Course in Electronic Warfare*, Artech House, 2001.
2. R. Poisel, *Modern Communications Jamming: Principles and Techniques*, Artech House, 2011.
3. W. Xu et al., "The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks," *Proc. 6th ACM Int'l Symp. Mobile Ad Hoc Networking and Computing* (MobiHoc 05), 2005, pp. 46–57.
4. K. Pelechrinis, M. Iliofotou, and S.V. Krishnamurthy, "Denial of Service Attacks in Wireless Networks: The Case of Jammers," *IEEE Comm. Surveys & Tutorials*, vol. 13, no. 2, 2011, pp. 245–257.

Internet DoS flooding attacks. However, more recent EW and cyberattacks tend to begin with a reconnaissance phase to better understand the target's technical characteristics and tailor the attack. The EW literature, reflecting its military heritage, called this preliminary stage *signals intelligence*. As with cyberattacks, jamming can serve a larger purpose than just denying communications. For example, it could deny wireless users access to a network with strong authentication and privacy mechanisms but permit access to another network with inadequate security measures, thereby setting the stage for confidentiality, integrity, and identity breaches. However, a full discussion of the parallels between EW and cyberattacks is beyond this article's scope.

Because this taxonomy covers only jamming, we distinguish jammers and cyberattacks on the basis of the adversary's intended failure mode and attack type. Traditionally, jamming is performed using an RF attack vector, whereas a cyberattack is launched through a network attack vector. These lines are blurred when dealing with *correlated jamming*, wherein a jammer both receives and transmits a signal. We'll assume that a jamming signal isn't a valid frame or packet, because such attacks are rarely classified as jamming. However, we don't limit the jammer's receiving capabilities. For example, a jammer could process the received waveform at the media access control (MAC) and network (NET) layers to target a certain type of frame or packet. However, to align with the common definition of jamming, the jammer's transmitter must either inject noise into the communications link or transmit what looks like a real PHY-layer signal. Otherwise, the attack should be classified as a cyberattack. This classification isn't meant to limit the jammer's capability but rather to put a label on a given attack and better define this taxonomy's scope.

We don't discuss malicious node detection, antijamming strategies, jammer detection, or jammer localization in this article. Likewise, we don't cover radar jamming or radio navigation jamming or spoofing (that is, positioning navigation and timing), such as attacks on GPS. Our goal is to shed light on jammers' broad characteristics and provide the right references for someone interested in pursuing jamming-related research.

Key Jammer Capabilities

Our taxonomy primarily delineates jammers by capabilities that define their fundamental behavior. A jammer can have one or more of the following major capabilities:

- time correlation,
- protocol awareness,
- ability to learn, or
- signal spoofing.

Figure 1 shows how the jammer capabilities interrelate.

We chose these four capabilities based on our survey of jammer models that emphasized complex forms of jamming. For example, a *learning* (or *cognitive*) jammer might not represent the majority of what's found in current-day operations, but it's a topic of interest in recent research and likely to become more prevalent over the next decade. We discuss correlation in the time domain specifically, because it's implicit that, to be successful, a jammer's signal must have some correlation in the frequency domain with the victim's desired signal (that is, the jammer must at least be aware of the spectrum the victim uses to perform jamming).

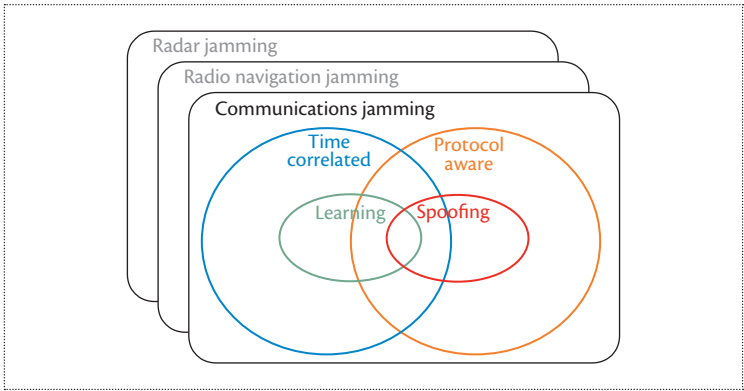

Figure 1. A jammer's key capabilities and their relations. As the Venn diagram illustrates, some embodiments of communications jamming have only one of the key capabilities (time-correlated, protocol-aware, learning, or spoofing) discussed in the article, whereas others incorporate multiple capabilities.

Time Correlation

A *time-correlated jammer* (or *reactive jammer* in some literature) transmits a jamming signal correlated to the target signal in time. It can listen to the transmitter's signal, leading to the geometrical configuration shown in Figure 2. It can alternately receive then transmit or, for simultaneous receive and transmit operations, cancel its own signal or use separate directional antennas.

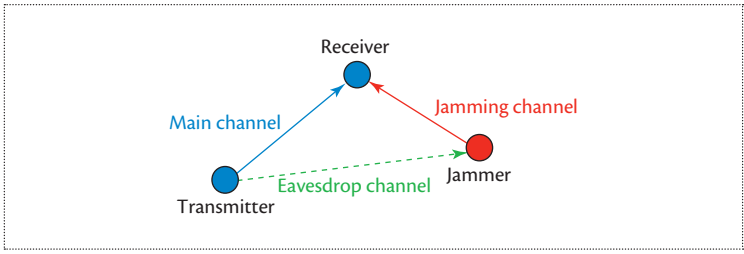

Figure 2. Geometrical configuration of a time-correlated jamming scenario, showing the three channels involved. This idealized representation shows the intrinsic physical distinction of the three channels: the victim's intended communications link from transmitter to receiver via the main channel, the jammer's observation of the victim's transmitter via the eavesdrop channel, and the path from the jammer's signal to the victim's receiver via the jamming channel.

This class of jammer can take on a wide range of specific behaviors. For example, it might sense a block of subchannels and jam those containing energy significantly above the noise floor, or it might retransmit a manipulated replica of what it receives, as in a digital RF memory (DRFM) or repeater jammer. Although time-correlated jamming is a very broad category, it quickly identifies a jammer's complexity, because a time-correlated jammer must have some form of receiver. Because significant engineering goes along with receiving capability (for instance, a full RF chain, sampling, and processing), any time-correlated jamming attack corresponds to a more complex attack. In this article, we refer to a jammer that isn't time correlated as *noncorrelated*.

How is jamming possible without a receiver? And how does a jammer know which signals to jam? When we discuss a jamming attack, we're referring to a specific attack being launched against a signal. An adversary must perform the following steps before launching an attack:

1. signal awareness—sense and detect signals across the spectrum of interest;
2. threat assessment—decide whether each signal will be jammed; and
3. attack selection—select the best attack for each chosen signal.

Time correlation comes into play during the actual attack, not the signal awareness step. Obtaining signal awareness requires receiving capability, but a time-correlated attack requires a jammer tightly synchronized to the target signal.

Protocol Aware

Protocol awareness simply means that a jammer knows the target signal's protocol. An adversary obtains information about the signal's protocol during the signal awareness step and uses this information in the attack selection step. For example, a jammer might discover that a particular signal is Wi-Fi or the current, fourth generation of cellular technology known as Long Term Evolution (LTE), which, due to the specifications' open nature, lets the jammer know almost everything about the PHY and MAC layers. A jammer could use a priori protocol knowledge to exploit its weaknesses and launch a jamming attack that's more effective and possibly harder to detect than a non-protocol-aware jamming attack.

A signal doesn't have to belong to a specific technology to be open to a protocol-aware attack. For example, a jammer might only know that a signal uses orthogonal frequency division multiplexing (OFDM) with pilots in certain locations, but it is considered protocol aware if it knows exactly where the pilots are. If a jammer knows the specific protocol being used, it can increase its effectiveness by jamming a PHY- or MAC-layer mechanism instead of the data. In most wireless protocols, the data takes up the largest portion of time and frequency resources. Thus, if a jammer targets something other than the data, the resulting attack will likely use less power and be harder to detect (as long as the targeted mechanism is essential for communications). Mechanisms that could be targeted in a protocol-aware attack include

- control channels or subchannels,
- control frames or packets (for instance, acknowledgments [ACKs]),
- pilots (or reference symbols), and
- synchronization signals.

For a survey of protocol-aware jamming attacks against Wi-Fi and LTE, see "Intelligent Jamming in Wireless Networks with Applications to 802.11 b and Other Networks," and "Vulnerability of LTE to Hostile Interference."[1,2]

Ability to Learn

In this article, we use the term *learning* in the machine-learning sense—systems that learn from data rather than only following explicit instructions. A jammer that can learn is one that can modify its behavior in real time in response to its experiences (that is, instances of successful or unsuccessful jamming actions or decisions).[3] However, a learning system has capabilities beyond an adaptive system that's limited to following a preprogrammed change sequence in response to stimuli. A simple test to determine whether a jammer can learn is to see whether its behavior evolves in response to a target's behavior and adaptation. Learning jammers go beyond detecting a target's waveform type and choosing from a preprogrammed set of jamming waveforms. Rather, they might detect that a target has initiated an antijam strategy and then explore different strategies to circumvent this defense.

This category includes some of the most complex jammers. Learning algorithms (for instance, supervised learning algorithms such as the popular Support Vector Machine or artificial neural networks) are complex, with high computational complexity during training. In addition, a jammer might have difficulty determining a certain jamming attack's success, because it might not have access to channel feedback information. In this case, it could use traffic analysis.

The ability to learn often leads to a "cognitive" label. However, a cognitive jammer capable of learning shouldn't be confused with "cognitive radio jamming" wherein a jammer is designed to deny a cognitive radio network (for example, the primary user emulation attack[4]). Some cognitive radio-jamming literature uses the term "cognitive jammer" even though the primary user emulation attack rarely involves learning and often isn't time correlated.

In some situations, a learning jammer might target radios that can also learn, such as cognitive radios in the Mitola sense,[5] as opposed to dynamic spectrum-sharing radios. A jammer can exploit this fact using a *belief manipulation attack*, which alters the victim radio's internal model of the world so that the targeted system's adaptation process seeks a poor operating point.[6] If you can metaphorically convince a radio that up is down and down is up, you can deceive the radio into rejecting optimal operating configurations and accepting poor operating configurations.

A jammer capable of learning is almost surely time correlated because learning involves observing the target signal. We consider learning and protocol awareness as independent features, leading to the relationship shown in Figure 1.

Spoofing

Spoofing, or *protocol emulation*, is broadly defined as a situation in which one entity successfully masquerades as another by falsifying data or signals to gain an illegitimate advantage. Typically, spoofing targets a PHY-layer mechanism by emulating a signal. In terms of jamming, which is assumed to be a PHY-layer adversary, we define spoofing as a signal transmission meant to look like a legitimate signal. To distinguish PHY-layer spoofing from, for example, transmitting fake frames or packets, we confine spoofing to the PHY layer. In other words, the spoofed signal need not have properties that make it look like a valid frame or packet; rather, it must be intended to fool the target's signal processing. Spoofing might or might not be time-correlated, although in literature it's correlated more often than not.

Protocol-aware jamming might or might not involve spoofing. But if spoofing occurs, the jammer is almost surely protocol aware because it needs to know what to spoof. Determining whether a given adversary is spoofing is simple: check whether it's transmitting noise or transmitting something that looks legitimate to a target's PHY layer. The difference between PHY-layer spoofing and higher-layer spoofing is less clear. However, if an adversary is transmitting what looks like a valid packet or frame, the attack is definitely not

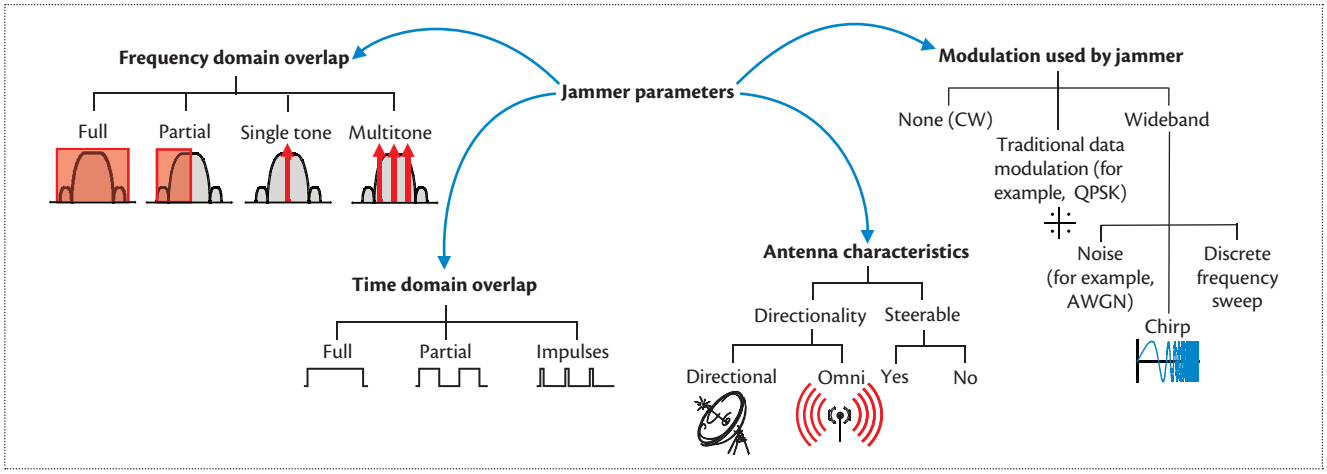

Figure 3. Jammer parameters organized into trees. A jamming signal has several operating parameters, including its frequency domain overlap with the target signal, its transmission pattern in time, and its modulation scheme. The jammer itself must have a transmit antenna, which has its own characteristics. The settings of each parameter largely depend on the jammer's desired objectives. Organizing these settings into subgroups is a lower-level refinement of the jamming taxonomy.

confined to the PHY layer and falls under the category of cyberattack.

Cognitive radio-jamming techniques, such as primary user emulation, could be considered spoofing depending on a jammer's specific waveform transmissions. In primary user emulation in which secondary users utilize only an energy detector, a jammer must transmit noise only at a particular frequency for the secondary users to evacuate the band and avoid using the spectrum. Other forms of primary user emulation could involve a jammer transmitting a signal meant to look like the primary user's signal (for example, a radio station's pilots), in which case it's PHY-layer spoofing.

Jammer Parameters

Building on the definition of major categories of jammer capabilities, we add a second-tier refinement of physical parameter values. As Figure 3 illustrates, example parameters include frequency and time overlap with the jamming target, antenna directionality, and the jammer's waveform or modulation. In this way, jammer types that were treated as distinct technologies in early literature can be understood now as minor variations on a common algorithm.

Specific Jamming Attacks

Figure 4 shows several example jamming attacks; we discuss how they're classified with respect to our taxonomy.

Barrage Jamming

Barrage jamming is the simplest form of jamming and is usually defined as a jammer transmitting noise-like energy across the entire portion of spectrum occupied by the target with 100 percent duty cycle in time (that is, transmitting until the attack ends). Thus, it is noncorrelated and non–protocol aware. Barrage jamming has been shown game theoretically and information theoretically to be the best a jammer can do in the absence of any knowledge of the target signal.[7]

Partial-Band Jamming

When jamming a single-carrier signal, gains can be achieved by jamming only a fraction of the entire signal in the frequency domain. Partial-band jamming is usually considered a noncorrelated jamming attack because the jammer transmits continuously in time. Performing partial-band jamming against an OFDM waveform doesn't make sense because strong forward error correction could still permit the OFDM receiver to reconstruct data from the unjammed subcarriers.

Automatic Gain Control Jamming

A receiver's automatic gain control (AGC) mechanism adjusts the input gain such that the received signal comes in at the best level to utilize the dynamic range of analog-to-digital converters. A jamming attack that targets an AGC mechanism uses a very low duty cycle (for instance, 2 percent) but with extremely high instantaneous power. By not transmitting continuously, a jammer can save power and remain harder to detect in some situations.[8] AGC jamming is noncorrelated, although the specific period and duty cycle used are important parameters. Aside from the assumption or knowledge that the target receiver uses AGC, it's non–protocol aware.

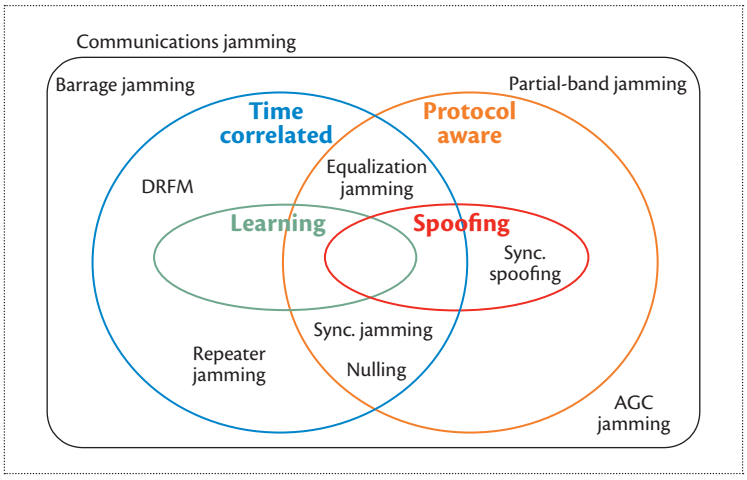

Figure 4. Specific jamming techniques mapped according to key jammer capabilities. As the Venn diagram shows, a jammer may incorporate just one of the key capabilities (time-correlated, protocol-aware, learning, or spoofing) or it might incorporate several capabilities. For example, a repeater jammer (shown at bottom left) has only time-correlated capability, but an equalization jamming (shown at middle top) has both time-correlated and protocol-aware capabilities.

Equalization Jamming

Equalization jamming involves targeting any mechanism related to *equalization*, the processing a wireless receiver applies to a received signal to compensate for distortions the signal suffered during propagation over the wireless link. Wireless transmitters insert known data symbols (or *reference symbols*) into the transmitted waveform to estimate the channel's frequency response and equalize the channel's effect at the receiver prior to demodulation. These known symbols are called *pilot symbols* in multicarrier communications, such as OFDM or single-carrier frequency division multiple access (SC-FDMA), and *channel sounding symbols* in multiple-input and multiple-output (MIMO) systems.[9]

In OFDM, pilot tone jamming is simply the process of jamming pilot tones, which might reside on certain subcarriers (in the case of IEEE 802.11) or multiplexed in time and frequency with data (in the case of LTE). Pilot jamming is protocol aware because the jammer must know where the pilots are. If the pilots are on a dedicated subcarrier, the attack is noncorrelated, but if they're multiplexed in time, the attack on the pilots must be time correlated to surgically jam the pilots. Pilot jamming can be energy efficient; similar degradation in the targeted receiver's bit error rate (BER) can be achieved using roughly one-tenth of the energy.[9] The pilot-jamming process is similar in SC-FDMA—the single-carrier variant of OFDM used in the uplink of the LTE air interface.

In MIMO systems, adversaries use known reference signals for channel sounding, thus, these can be jammed as well if they're known by the jammer a priori.

Synchronization Jamming

For a communications link to function, the receiver must synchronize to the incoming signal in both time and frequency. To aid in this task, a synchronization signal or synchronization symbols are usually designed into the PHY-layer protocol. For example, in LTE, two different synchronization signals appear every 5 ms. Synchronization jamming (also called *synchronization signal jamming*) is simply the process of surgically jamming one or more synchronization signals. This jamming technique is unique in that it might only prevent radios from establishing a communications link, and thus, it won't cause immediate DoS.[2] However, synchronization signals tend to be very sparse with respect to the entire signal, thus providing a significant jamming gain. Synchronization jamming must be protocol aware to know where the synchronization signal is located. It must be time correlated, assuming the synchronization signal is multiplexed in time with data and other signaling.

Nulling

Instead of transmitting noise as the jamming waveform, a nulling (or *phase-coherent*) attack involves transmitting a structured waveform so the target's received energy is driven as close to zero as possible.[9] This is done by causing the jamming signal to be received as the target signal's π-radian phase shift, thus nulling out the target signal and leaving only channel noise for the demodulator.

Nulling attacks are extremely challenging and considered infeasible in real-world scenarios because they require extremely accurate knowledge of the three channels involved, which can be difficult to obtain considering wireless channels' varying nature. Even if the target signal includes pilots and synchronization symbols, this would provide accurate knowledge of only the channel between the transmitter and jammer. To perform nulling, the jammer would also need to know the jammer-receiver and transmitter-receiver channels, as Figure 2 shows. Nulling also requires a priori knowledge of the signal, which is possible in some circumstances (for example, the value of pilot sequences in Wi-Fi and LTE is openly published). Thus, a nulling attack is protocol aware but impractical to implement due to the need for real-time knowledge of what are, in reality, random characteristics of a wireless channel.

Even though nulling attacks are believed to be infeasible in practice, we include them in this taxonomy due to their presence in academic literature. Pilot nulling against OFDM, which was introduced in "Efficient OFDM Denial: Pilot Jamming and Pilot Nulling,"[9] involves nulling the pilots at the target receiver.

Repeater Jamming

Repeater jamming (also called *DRFM* or *follower jamming*) is the simplest form of time-correlated jamming

when a jammer has no knowledge of the protocol. In repeater jamming, a jammer transmits only when there is energy on the channel. It might transmit what it receives with noise added, or it could sense a series of subchannels and transmit noise when it senses energy on one or more subchannels. Regardless of the specific behavior used, repeater jamming can "follow" a signal if it hops around in frequency, negating the antijam gains associated with frequency-hopping spread spectrum (FHSS). When there are large distances between the transmitter, receiver, and jammer, a repeater jamming attack might fail to achieve time correlation with the target signal due to propagation delay.

Protocol-Aware Jamming against Wi-Fi

Likely owing to Wi-Fi's popularity and longevity, several intelligent jamming attacks against Wi-Fi (IEEE 802.11) are described in open literature.[1] Clear-to-send jamming occurs when a jammer waits for a request-to-send packet to be transmitted over the channel, then transmits a burst of noise after waiting for the short interframe space (SIFS) interval, which is defined in the specifications.[10] ACK jamming works the same way, except a jammer waits for a data packet to be transmitted then, after an SIFS interval, transmits a burst of energy with the intent to jam the ACK.[10]

These two attacks are protocol aware and time correlated; however, a protocol-aware and noncorrelated attack in 802.11 is possible by transmitting periodic pulses that repeat with a frequency based on the 802.11 DCF interframe space (DIFS) duration.[1] This duration determines how long a node senses the channel to decide whether it's idle; thus, the DIFS jamming attack causes a "busy channel" while saving power in a noncorrelated manner.

A protocol-aware jamming strategy that incorporates learning is proposed in "Optimal Jamming Using Delayed Learning."[11]

Protocol-Aware Jamming against LTE

In LTE, data is multiplexed with control information in both time and frequency, due to the use of OFDM. Especially in the downlink, symbols often contain data combined with control information, which can be surgically jammed by targeting the specific subcarriers (that is, frequencies) they occupy. Researchers have recently investigated several protocol-aware jamming attacks against LTE and found them to be significantly more effective than barrage jamming.[2] For example, a downlink control channel called the physical control format indicator channel (PCFICH) holds only two bits of information but is transmitted within every subframe and vital to the downlink control channel operation.[12] Because the PCFICH is so sparse in time and frequency, jamming it is approximately 20 dB more effective than barrage jamming in terms of the overall jammer-to-signal ratio.[2] The PCFICH attack is time correlated because the PCFICH is multiplexed in time with other physical channels.

One noncorrelated jamming attack against LTE is jamming the physical uplink control channel (PUCCH), which is always at the edges of the uplink bandwidth, meaning a jammer can use the open specifications to predict the PUCCH's frequency. This is a noncorrelated attack because no other physical channels are multiplexed in time with the PUCCH. For more information on jamming attacks against LTE, see "Vulnerability of LTE to Hostile Interference."[2]

As communications systems' sophistication increases, complex jamming will likely become a bigger threat in public safety, military, and other mission-critical domains. Our jammer taxonomy organizes jammer classes by the information they possess and their capacity to act on that information. This new view of jammers emerges naturally from present-day wireless technology's extensive reliance on software-driven behavior. Understanding the key capabilities that distinguish major classes of jamming, as well as the multidimensional parameter space, can aid in the correct application of antijam and detection strategies.

Further research includes the design of radar jamming and radio navigation jamming taxonomies. It might even be possible to formulate a taxonomy that applies to all forms of jamming. ■

References

1. D. Thuente and M. Acharya, "Intelligent Jamming in Wireless Networks with Applications to 802.11 b and Other Networks," *Proc. IEEE Conf. Military Comm.* (MILCOM 06), vol. 6, 2006, pp. 1075–1081.
2. M. Lichtman et al., "Vulnerability of LTE to Hostile Interference," *IEEE Global Conf. Signal and Information Processing* (GlobalSIP 13), 2013.
3. S. Amuru and R.M. Buehrer, "Optimal Jamming Strategies in Digital Communications—Impact of Modulation," *IEEE Global Comm. Conf.* (GLOBECOM 14), 2014, pp. 1619–1624.
4. R. Chen, J.-M. Park, and J.H. Reed, "Defense against Primary User Emulation Attacks in Cognitive Radio Networks," *IEEE J. Selected Areas in Comm.*, vol. 26, no. 1, 2008, pp. 25–37.
5. J. Mitola and G.Q. Maguire Jr., "Cognitive Radio: Making Software Radios More Personal," *IEEE Personal Comm.*, vol. 6, no. 4, 1999, pp. 13–18.
6. T.C. Clancy and N. Goergen, "Security in Cognitive Radio Networks: Threats and Mitigation," *Cognitive Radio*

Oriented Wireless Networks and Communications (CrownCom 08), 2008, pp. 1–8.
7. T. Basar, "The Gaussian Test Channel with an Intelligent Jammer," *IEEE Trans. Information Theory*, vol. 29, no. 1, 1983, pp. 152–157.
8. D.L. Adamy, *EW 101: A First Course in Electronic Warfare*, Artech House, 2001.
9. T.C. Clancy, "Efficient OFDM Denial: Pilot Jamming and Pilot Nulling," *IEEE Int'l Conf. Comm.* (ICC 11), 2011, pp. 1–5.
10. K. Pelechrinis, M. Iliofotou, and S.V. Krishnamurthy, "Denial of Service Attacks in Wireless Networks: The Case of Jammers," *IEEE Comm. Surveys & Tutorials*, vol. 13, no. 2, 2011, pp. 245–257.
11. S. Amuru and R.M. Buehrer, "Optimal Jamming Using Delayed Learning," *Proc. IEEE Conf. Military Comm.* (MILCOM 14), 2014, pp. 1528–1533.
12. S. Sesia, I. Toufik, and M. Baker, *LTE: The UMTS Long Term Evolution*, Wiley, 2009.

Marc Lichtman is a PhD student in electrical engineering at Virginia Tech. His research interests include electronic warfare, machine learning, cognitive radio, and wireless communication system design. Lichtman received an MS in electrical engineering from Virginia Tech. Contact him at marcll@vt.edu.

Jeffrey D. Poston is a PhD student in electrical engineering at Virginia Tech. His research interests include the intersection of cybersecurity and electronic warfare (EW) to form a new field, cyber EW; machine learning; cognitive radio; and novel indoor geolocation techniques and their privacy implications. Poston received an MS from George Washington University. Contact him at poston@vt.edu.

SaiDhiraj Amuru was a PhD student in electrical engineering at Virginia Tech at the time of this writing. His research interests include cognitive radio, statistical signal processing, and machine learning. Amuru received a PhD in electrical engineering from Virginia Tech in 2015. Contact him at adhiraj@vt.edu.

Chowdhury Shahriar was a PhD student in electrical engineering at Virginia Tech and a graduate research assistant in Virginia Tech's Hume Center for National Security and Technology at the time of this writing. His research interests include wireless communications, communications security, signal processing, and spectrum management and policy. Shahriar received a PhD in electrical engineering from Virginia Tech in 2015. Contact him at cshahria@vt.edu.

T. Charles Clancy is an associate professor of electrical and computer engineering at Virginia Tech and director of the Hume Center for National Security and Technology. His current research interests include cognitive communications and spectrum security. Clancy received a PhD in computer science from the University of Maryland. He's a senior member of IEEE. Contact him at tcc@vt.edu.

R. Michael Buehrer is a professor of electrical engineering at Virginia Tech and director of Wireless @ Virginia Tech, a comprehensive research group focusing on wireless communications. His research interests include geolocation, position location networks, iterative receiver design, electronic warfare, dynamic spectrum sharing, cognitive radio, communication theory MIMO communications, intelligent antenna techniques, ultra wideband, spread spectrum, interference avoidance, and propagation modeling. Buehrer received a PhD in electrical engineering from Virginia Tech. Contact him at rbuehrer@vt.edu.

Jeffrey H. Reed is the founder of Wireless @ Virginia Tech and served as its director until 2014. His research interests include software-defined radio, cognitive radio, smart antennas, and wireless security. Reed received a PhD in electrical engineering from the University of California, Davis. He's an IEEE Fellow, a distinguished lecturer for the IEEE Vehicular Technology Society, and a member of the Commerce Spectrum Management Advisory Committee, which advises the National Telecommunications and Information Administration on spectrum issues. Contact him at reedjh@vt.edu.