How to Jam Without Getting Caught: Analysis and Empirical Study of Stealthy Periodic Jamming

Bruce DeBruhl and Patrick Tague
Carnegie Mellon University
{debruhl, tague}@cmu.edu

Abstract—**Despite the widespread commercial use of spread spectrum technology, advanced algorithms and modern hardware capabilities still allows efficient denial-of-service attacks against wireless communication systems using jamming. Much of the recent work on jamming mitigation has focused on how to adjust the transmitter-receiver system once a jamming attack has been detected. However, characterizing the detectability of certain classes of jamming attacks remains a largely unstudied problem. We aim to narrow this gap by analyzing the effect of a class of periodic jamming attacks on the attack detection metrics of packet delivery ratio (PDR) and received signal strength (RSS). We show that a well-designed jamming signal can effectively defeat RSS-based detection while causing a significant and often devastating reduction in PDR, demonstrating that RSS-based detection is insufficient. We further evaluate our claims through implementation of a periodic jammer using a wide range of signal parameters against a transmitter-receiver pair communicating using IEEE 802.15.4, demonstrating the validity of our analytical claims.**

I. Introduction

Simply defined, jamming is the broadcast of a signal on a wireless medium with the intention to interfere with legitimate traffic [1]. Jamming has the potential to reduce or eliminate the ability of neighboring nodes to communicate, thus making it an important issue to understand in our ubiquitous wireless age. The importance of availability in wireless communications has thus prompted the development of various techniques to mitigate jamming effects or detect jamming attacks.

Traditional jamming mitigation is done through the use of spread spectrum techniques, which aim to increase the cost of mounting an equally effective jamming attack, effectively pushing the cost-benefit ratio for the attacker to an unacceptable limit. Direct sequence spread spectrum (DSSS) and frequency hoping spread spectrum (FHSS) make up the typical set of spread spectrum techniques [1], through DSSS is deployed more widely in commercial systems, including IEEE 802.15.4 [2] and 802.11b [3]. In a typical DSSS implementation, each data bit is mapped to a chip sequence using a very low rate encoding and then combined into data symbols. The chip symbols are sent at a much higher rate than the desired bit rate to compensate for the low-rate encoding. The highly redundant DSSS encoding provides strong protection of each symbol against interference and jamming [4]. However, a simple constant tone-jammer is still able to degrade the performance of a DSSS-based system [5].

Because spread spectrum techniques do not solve the problem of mitigating or discouraging jamming, recent work has continued to study this problem in various contexts. One effective way to mitigate the effects of jamming is to detect the attack using a suitable collection of observable metrics and then respond by changing the system operation [5]. A promising detection technique suitable for even resource constrained devices such as wireless sensor nodes is to use a comparison of received signal strength (RSS) and packet delivery ratio (PDR). If the receiver observes a high RSS with a low PDR, a likely conclusion is that the receiver is under attack. This inexpensive detection mechanism is effective in detecting a wide range of jamming attacks. Another metric for jamming attack detection is the amount of time it takes a sender to gain access to the channel at the MAC layer, as a persistent jammer can cause the sender's carrier sense readings to be consistently above the threshold for channel occupancy. These RSS-based techniques can be used cooperatively to allow the sender and receiver to jointly detect a jamming attack. Once a jamming attack is detected, the sender and receiver can take a number of different actions, including hopping to a different channel, determining the jammer's location and retreating away from it [5], changing transmission signal or encoding parameters [6], or attempting to decommission the jamming device. The effectiveness of such techniques, however, relies on sufficiently accurate detection.

Another consideration of modern electronic warfare is the increasing availability of low-power ultra-portable devices that can provide incredible computation power in a small package. Instead of limiting the threat model to specialized radio equipment mounted on large vehicles or powered by a generator, jamming can now be mounted using commercially available, battery-powered platforms. With more advanced attack algorithms, these platforms enable highly effective jamming attacks that are more difficult to detect. Recent work has demonstrated a wide array of efficient jamming techniques using advanced algorithms and inference about the target system [7]–[11]. Though a lot of work has focused on the use of reactive and selective jammers [11], these require specialized hardware, which is often prohibitively expensive. We thus focus on non-reactive jammers and explore the results

This research was supported by CyLab at Carnegie Mellon University under grant DAAD19-02-1-0389 from the Army Research Office and by the National Science Foundation under grant CNS-1149582. Bruce DeBruhl is supported by an NDSEG fellowship. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of ARO, CMU, NSF, or the U.S. Government or any of its agencies.

that they can obtain.

In order to gain a deeper understanding of the capabilities of RSS-based jamming detection techniques as applied to efficient jamming attacks, we study the problem of characterizing the effect of a class of jamming attacks that offer efficiency and low probability of detection without significantly sacrificing efficacy. We show that a certain class of periodic jamming [12] using very short signal periods, hereafter referred to as *short form periodic jamming* (SFPJ), can seriously degrade communication capabilities without significantly altering RSS measurements. This is achieved by designing the jamming signal parameters with respect to the communication protocol used in the target system. Toward the desired analysis and design goals, we make the following contributions.

- We analyze the effects of SFPJ attacks on the packet delivery ratio (PDR) and received signal strength (RSS) observed by the receiving device as well as the energy expenditure of the attacker.
- We provide a design methodology describing how a SFPJ attacker can choose jamming signal parameters to balance its goals of efficacy, efficiency, and detectability.
- We demonstrate the effects of SFPJ in a system implementation in which the attacker targets an 802.15.4 link.

The remainder of this paper is organized as follows. We describe our system model and assumptions in Section II. In Section III, we present detailed analysis of the effectiveness, efficiency, and detectability of SFPJ attacks. We present a brief attack design methodology in Section IV. We show the empirical results from our system implementation in Section V. Finally, we summarize our contributions in Section VII.

II. System Model & Assumptions

As our goal in this work is to gain a better understanding of the impact of short form periodic jamming (SFPJ) on receiver-side attack detection, we consider a simplified network model that eliminates the influences of random noise, higher-layer protocols, and sender-side considerations. We consider a three-node network comprising a transmitter, receiver, and jammer, as shown in Figure 1. In our analysis, we assume symbol errors do not occur in the absence of attack. We make this assumption because of DSSS's heavy symbol error correction capabilities. We assume that the transmitter sends packets freely, without using carrier-sensing; this allows us to focus on receiver-side characteristics only. Moreover, this eliminates the potential impacts of MAC-layer jamming attacks [13] that aim to prevent the transmitter from gaining channel access.

We assume that the transmitter and receiver communicate using a standard packet communication protocol, packing sequences of bits or symbols into each packet, either using standard encoding techniques or anti-jamming methods such as direct sequence spread spectrum (DSSS). We allow for additional use of error correcting codes beyond the standard symbol encoding, and we define an error threshold E such that the receiver can correctly decode any packet with up to E symbol errors. We make no assumption about packet rates or periodicity, and we assume no correlation between timing of packet transmissions and jamming pulses. We assume that no additional anti-jamming signal processing is implemented beyond the protections of DSSS.

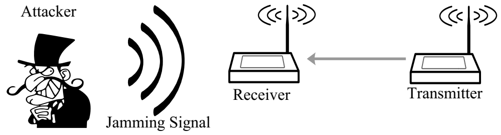

Fig. 1: We illustrate the assumed three-node system model in which the attacker jams the communication from the transmitter to receiver. We assume the sender does not use carrier sensing, so the transmitter is not affected by the jammer.

The jamming attack of interest in this work is referred to as periodic jamming, meaning that the attacker cycles the jamming signal on and off according to a regular period and duty cycle. We suppose that the period of the jamming signal is equal to T_j symbol durations. Within each attack period, the jammer is active for a duration of $\Delta_j \leq T_j$ symbol durations, corresponding to a duty cycle of $d_j = \Delta_j / T_j$. As previously defined, the SFPJ attack corresponds to T_j being relatively small compared to the packet duration, here assumed equal to Δ_{pkt} symbol durations, whereas long form periodic jamming uses T_j much larger than the packet duration. Due to the nature of the attack, we assume that $\Delta_j < \Delta_{pkt}$. In choosing its attack parameters T_j and d_j, we assume that the jammer has knowledge of the signal parameters, and communication protocols. Figure 2 illustrates the parameter choices corresponding to long form and short form periodic jamming. In the figure, the long form jammer misses a packet despite having a duty cycle of 50%, while the short form jammer hits every packet using only a 10% duty cycle. We show empirically, in Section V, that in most cases the short pulses of the short form jammer can be tuned to have a comparable effect on packet communication to that of the long form jammer, though often with a decrease in energy expenditure or detectability.

III. Evaluation of Jamming Attack Parameters

An adversary's design and the defender's understanding of efficient and stealthy jamming attacks using SFPJ rely on a firm characterization of the effects of SFPJ on packet communications for a variety of system configurations and attack parameters. In this section, we provide analysis to evaluate the effect of SFPJ using the previously defined model. First, to capture the effect of the attack on the transmitter-receiver system, we evaluate the packet delivery ratio (PDR), equal to the fraction of transmitted packets that are correctly decoded by the receiver. Next, to capture the detectability of the attack, characterizing the attacker's risk, we evaluate the effect of SFPJ on received signal strength (RSS), which is integral to many jamming detection algorithms. Finally, to capture the adversary's attack cost, we evaluate the energy consumption of the SFPJ attack. The metric of energy consumption directly affects the class of device that the attack can be launched from,

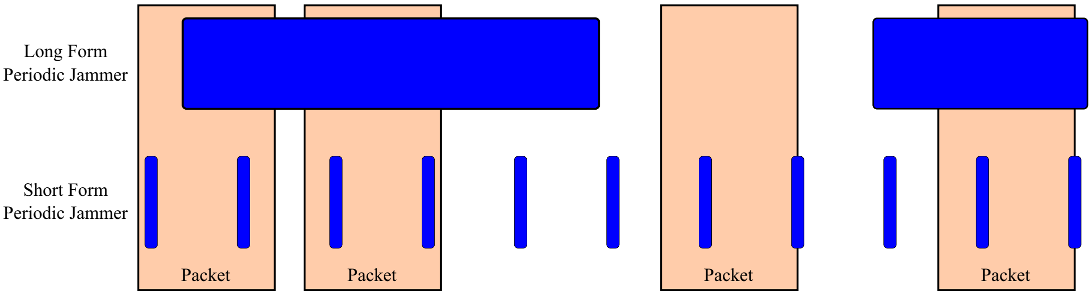

Fig. 2: The primary difference between long form and short form periodic jamming is the relative size of the jamming period T_j relative to the packet duration Δ_{pkt}. Because of the short jamming period in SFPJ, typically such that $T_j < \Delta_{pkt}$, an equally effective attack can be designed using a significantly smaller duty cycle, corresponding to less attack energy and lower detectability.

whether a large vehicle-mounted radio or a hand-held mobile device, as well as the detectability of the jamming attack.

A. Packet Delivery Ratio

Given the system model in Section II, we evaluate the probability of correct packet decoding at the receiver, known as the packet delivery ratio (PDR). Since correct decoding depends on a maximum of E symbol errors per packet, we denote this probability as $P_{pkt}(E)$ for a given jamming period T_j, jamming duty cycle d_j (or equivalently, the pulse duration Δ_j), and packet duration Δ_{pkt}. As previously mentioned, we neglect the effect of channel errors and assume symbol errors are only caused by the jammer. The probability $P_{pkt}(E)$ is therefore given by the probability that the jammer can induce more than E symbol errors in a single packet. As an auxiliary probability, we let $\varepsilon_s(x)$ denote the probability that exactly x symbol errors occur in a single packet, yielding

$$P_{pkt}(E) = 1 - \sum_{i>E} \varepsilon_s(i). \qquad (1)$$

Because of the periodic nature of the SFPJ attack, the symbol errors induced by the jammer must occur during the individual jamming pulses. To further assist in our analysis, we define the probability $\pi_1(x)$ as the probability that exactly x symbol errors are caused by a single jamming pulse of duration $\Delta_j = d_j T_j$. To estimate the probability $\pi_1(x)$, we count the number of symbols in the transmitted packet that overlap completely with the jamming pulse, effectively ignoring partially jammed symbols[1]. A jamming pulse of duration Δ_j symbols will interfere with either $\lfloor \Delta_j \rfloor - 1$ or $\lfloor \Delta_j \rfloor$ symbols, where $\lfloor x \rfloor$ is the maximum integer less than or equal to x, depending on the temporal alignment of the symbols and jamming pulse. Figure 3 illustrates the occurrence of each amount of symbol overlap for the same parameters. Assuming that the start time of the jamming pulse is uniformly distributed within the corresponding symbol duration, the probability of overlapping with $\lfloor \Delta_j \rfloor$ symbols is $p_\Delta = \Delta_j - \lfloor \Delta_j \rfloor$, while that of overlapping with $\lfloor \Delta_j \rfloor - 1$ symbols is $1 - p_\Delta$.

In many cases, especially when spread spectrum technology such as DSSS is used, the probability of decoding error due to jamming, even when the entire symbol is overlapped by a jamming signal, is less than unity. We thus define the probability of symbol error due to jamming as p_s, noting that p_s varies with jamming signal power and various system and attack parameters. Models to characterize this probability p_s exist in the literature [1], and we do not address this probability further. However, when symbols errors occur only with a probability p_s, the number of symbol errors due to the single jamming pulse becomes a binomial random variable. Letting $\beta(x,n,p)$ denote the binomial probability function [14], i.e. that x trials out of n are successful given that each trial is successful with probability p, the probability $\pi_1(x)$ that x symbol errors are induced by one jamming pulse is given by

$$\pi_1(x) = p_\Delta \beta(x, \lfloor \Delta_j \rfloor, p_s) + (1 - p_\Delta)\beta(x, \lfloor \Delta_j \rfloor - 1, p_s). \qquad (2)$$

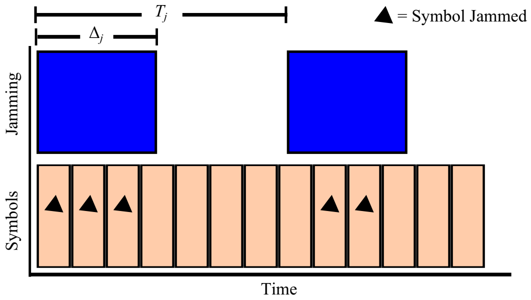

Fig. 3: Depending on the temporal alignment of the jamming pulse and the transmitted symbols, a single jamming pulse of duration Δ_j symbols will overlap completely with either $\lfloor \Delta_j \rfloor$ or $\lfloor \Delta_j \rfloor - 1$ symbols. This illustration shows how each of these can occur.

Using $\pi_1(x)$ as given in (2), we next derive equations for the symbol error probability $\varepsilon_s(x)$ and packet delivery ratio $P_{pkt}(E)$ for three different cases according to the number of jamming pulses that overlap with each packet transmission:

- **Case 1**: $T_j \geq \Delta_{pkt}$,
- **Case 2**: $T_j < \Delta_{pkt} < \lfloor \Delta_{pkt}/T_j \rfloor T_j + \Delta_j$,
- **Case 3**: $T_j < \Delta_{pkt}$ and $\Delta_{pkt} \geq \lfloor \Delta_{pkt}/T_j \rfloor T_j + \Delta_j$,

as illustrated in Figure 4. We use the notation $\varepsilon_{s,i}(x)$ and $P_{pkt,i}(E)$ to denote these probabilities for case $i \in \{1,2,3\}$. Each of the three cases is discussed individually as follows.

1) Case 1: In the first case, the jamming signal period is greater than or equal to the packet duration, i.e. $T_j \geq \Delta_{pkt}$.

[1] This approximation reflects the fact that partially jammed symbols in a typical DSSS system will still be correctly decoded with high probability, due to the despreading operation.

As illustrated in Figure 4, at most one jamming pulse will hit each packet in this case. Since not every packet will overlap with a jamming pulse, the symbol error rate is discounted by the probability $(\Delta_{pkt} - \Delta_j)/T_j$ that the packet will be hit by the jamming pulse. Therefore, the symbol error rate for this case is given by

$$\varepsilon_{s,1}(x) = \frac{\Delta_{pkt} - \Delta_j}{T_j} \pi_1(x), \tag{3}$$

and the packet delivery ratio $P_{pkt,1}(E)$ is given by (1) and (3).

2) Case 2: In the second case, the jamming signal period is such that either $k-1$ or k jamming pulses will hit each packet, where $k = \lfloor \Delta_{pkt}/T_j \rfloor$, as illustrated in Figure 4. In this case, the parameters satisfy the inequalities $T_j < \Delta_{pkt} < kT_j + \Delta_j$. As previously discussed, the difference between $k-1$ and k overlapping pulses relies on the temporal alignment between jamming pulses and transmitted symbols. If a packet starts within $\Delta_{pkt} - (k-1)T_j - \Delta_j$ symbol durations after the jammer's period starts, then k pulses will hit the packet. Otherwise, only $k-1$ pulses will hit the packet. Again assuming the time offset between the transmitted symbols and jamming pulses is uniform over the jamming period, the probabilities u_{k-1} that $k-1$ pulses hit and u_k that k pulses hit satisfy

$$u_k = 1 - u_{k-1} = \frac{\Delta_{pkt} - \Delta_j}{T_j} - k + 1. \tag{4}$$

The fact that multiple pulses now affect the symbol errors within the packet means that the total number of symbol errors must be aggregated over the $k-1$ or k jamming pulses. Since the number of errors due to each pulse is a random variable described using $\pi_1(x)$ in (2), the total number of errors due to any n pulses is the sum of n such random variables. Assuming these random variables are independent, the distribution of the total number of symbol errors due to n pulses is thus the n-fold convolution of the distribution $\pi_1(x)$ [15], which we denote as $\pi_1^{[n]}(x)$. The probability of symbol error $\varepsilon_{s,2}(x)$ for case two is thus given as the weighted sum of the $k-1$-fold convolution and k-fold convolution, weighted by u_{k-1} and u_k, as

$$\varepsilon_{s,3}(x) = u_{k-1}\pi_1^{[k-1]}(x) + u_k \pi_1^{[k]}(x), \tag{5}$$

and the corresponding packet delivery ratio $P_{pkt,2}(E)$ is then given by (1) and (5).

3) Case 3: In the third and final case, the jamming signal period is such that either k or $k+1$ jamming pulses will hit each packet, where $k = \lfloor \Delta_{pkt}/T_j \rfloor$, as illustrated in Figure 4. In this case, the parameters satisfy the inequalities $T_j < \Delta_{pkt}$ and $\Delta_{pkt} \geq kT_j + \Delta_j$. As before, the difference between k and $k+1$ overlapping pulses relies on the temporal alignment between jamming pulses and transmitted symbols. If a packet starts within $\Delta_{pkt} - kT_j - \Delta_j$ symbol durations after the jammer's period starts, then $k+1$ pulses will hit the packet. Otherwise, only k pulses will hit the packet. Using the same time offset assumption, the probabilities w_k that k pulses hit and w_{k+1} that $k+1$ pulses hit satisfy

$$w_{k+1} = 1 - w_k = \frac{\Delta_{pkt} - \Delta_j}{T_j} - k. \tag{6}$$

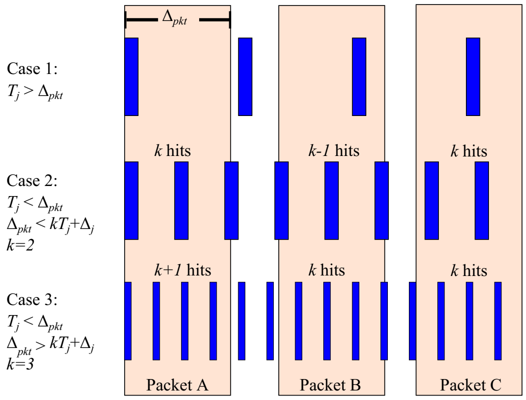

Fig. 4: The relationships between packet and jamming parameters are decomposed into three cases for analytical purposes. In each case, a different number of jamming pulses affect symbol reception in each packet.

Similar to the previous case, the symbol error $\varepsilon_{s,3}(x)$ is the weighted sum of convolution terms given by

$$\varepsilon_{s,3}(x) = w_k \pi_1^{[k]}(x) + w_{k+1}\pi_1^{[k+1]}(x). \tag{7}$$

The packet delivery ratio $P_{pkt,3}(E)$ in this case is then given by the combination of (1) and (5).

4) Unification of Cases: Despite the fact that we broke the analysis into three cases above, the overall equations for symbol error probability and packet delivery ratio can be unified into a single expression. In each of the three cases above, we observe that the number of jamming pulses that hit each packet can vary by 1: 0 or 1 in case 1, $k-1$ or k in case 2, and k or $k+1$ in case 3. Moreover, from the definitions of the three cases, we see that this variation is due to the relationship between Δ_{pkt}/T_j and $\lfloor \Delta_{pkt}/T_j \rfloor$. Namely, case 2 corresponds to

$$\frac{\Delta_{pkt}}{T_j} - \left\lfloor \frac{\Delta_{pkt}}{T_j} \right\rfloor < d_j, \tag{8}$$

while case 3 corresponds to

$$\frac{\Delta_{pkt}}{T_j} - \left\lfloor \frac{\Delta_{pkt}}{T_j} \right\rfloor \geq d_j. \tag{9}$$

Based on these case relationships, we define the variable m to take the appropriate values for the three cases such that either m or $m+1$ pulses hit each packet as

$$m = \left\lfloor \frac{\frac{\Delta_{pkt}}{T_j} - \left\lfloor \frac{\Delta_{pkt}}{T_j} \right\rfloor}{d_j} \right\rfloor + \left\lfloor \frac{\Delta_{pkt}}{T_j} \right\rfloor - 1. \tag{10}$$

By inspecting the weighting equations (4) and (6) used for cases 2 and 3, respectively, in the context of the variable m defined in (10), we see that both u_k and w_{k+1} can be replaced by a unifying weight ω given by

$$\omega = \frac{\Delta_{pkt}}{T_j} - m - d_j, \tag{11}$$

where m now differentiates between the three cases. Putting together the pieces, the unified equation for symbol error rate

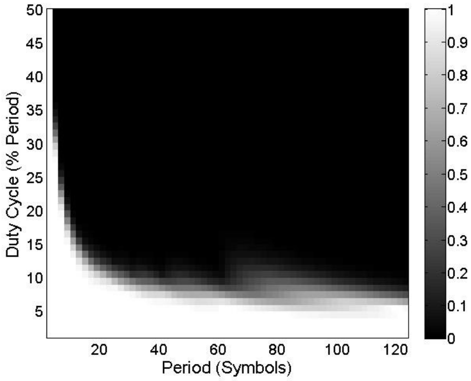

Fig. 5: The packet delivery ratio $P_{pkt}(E)$ given by (1) and (12) is plotted with $\Delta_{pkt} = 125$ symbols, $p_s = 0.8$, and $E = 7$, showing the effect of short form periodic jamming.

$\varepsilon_s(x)$ across the three cases can be expressed as

$$\varepsilon_s(x) = (1-\omega)\pi_1^{[m]}(x) + \omega\pi_1^{[m+1]}(x). \tag{12}$$

We note that (3) for case 1 is included as a special case of (12) because the 0-fold convolution $\pi_1^{[0]}(x)$ is zero.

Using the above analysis, the PDR $P_{pkt}(E)$ due to SFPJ can be generically expressed using the combination of (1) and (12). To illustrate the analytical result, we plot the equation directly in Figure 5 as a heat map with darker colors representing lower packet delivery ratio and lighter colors represent higher packet delivery ratio. The packet duration is fixed at $\Delta_{pkt} =$ 125 symbols, the probability that a symbol overlapping with a jamming pulse will be incorrectly decoded is fixed at $p_s = 0.8$, and the symbol error decoding threshold is fixed at $E = 7$. Both the jamming period T_j on the x-axis and the jamming duty cycle d_j on the y-axis are variable.

By inspection of Figure 5, we see that the jammer can meet a goal if inflicting significant packet error with a duty cycle of only 10% or lower, showing that SFPJ can have very high impact with a low duty cycle. We note that our analysis is conservative, in that we neglect the effect of random channel errors and partially jammed symbols, meaning the real-world PDR may be even lower than our analysis suggests.

B. Received Signal Strength

Based on the system model in Section II, we next evaluate the effect of short form periodic jamming on the received signal strength (RSS) measured when (correct or erroneous) packets are captured by the receiving radio. Since RSS has been assumed to be more-or-less orthogonal to PDR due to jamming in related literature [16], it has been proposed as an effective detector of jamming attacks. We evaluate the effect of SFPJ on RSS measurement, mainly as a function of the jamming duty cycle d_j (or equivalently the pulse duration Δ_j relative to the period T_j. We assume that the receiving radio uses a hardware- or OS-based service to sample the RSS values and keep a running distribution of RSS values. Since RSS is often used as a detection indicator, we assume the attacker would like to minimize its impact on RSS.

We consider an RSS measurement service that collects and averages m RSS samples at random intervals within each packet. The average RSS R comprises the weighted combination of an average RSS level R_0 in the absence of jamming with the level R_j in the presence of jamming. Since the SFPJ attack injects high-power pulses over a fraction d_j of the packet, the average RSS value is given by

$$R = (1-d_j)R_0 + d_jR_j. \tag{13}$$

Using existing knowledge of relative transmitter-receiver-jammer geometry and the RSS detection threshold τ, the attacker can thus estimate the probability distribution $\Pr[R < r]$ of the RSS R in (13). We do not elaborate on the details of the evaluation of the probability $\Pr[R > \tau]$, but we note that standard path-loss models [4] can be used to estimate the distributions of R_0 and R_j in (13) as a function of geometry, transmission and jamming signal powers, modulation and coding parameters, antenna gains, and other parameters.

C. Energy Efficiency

For a fixed average jamming power (corresponding to a fixed symbol error parameter p_s), the energy expenditure of a jamming attack is directly proportional to the duty cycle d_j. Hence, an efficient SFPJ attacker is one with the smallest duty cycle d_j that pushes the PDR below a desired target. By inspection of Figure 5, a goal of PDR near zero, for example, can be achieved with a jamming signal period of $T_j = 110$ symbols and a duty cycle of $d_j = 10\%$. Similarly, a less ambitious goal of PDR below 50% can be accomplished with $T_j = 82$ and $d_j = 7\%$. When lesser levels of error correction are employed, an equally effective attack becomes even more energy efficient, as a smaller number of symbol errors are required to cause erroneous packet decoding.

IV. Attack Design Considerations

Using our analysis from Section III as a guide, we briefly discuss the capability for an attacker to use SFPJ to design effective, efficient, and stealthy jamming attacks. Given that the goal of the attacker is to be as effective as possible subject to imposed constraints on the efficiency and detectability of the attack, we have provided all of the necessary analytical tools to allow the attacker to choose appropriate attack parameters, namely the signal period T_j and duty cycle d_j.

Our analysis of PDR under SFPJ using equations (1) and (12) provides one input into the design problem. The second component is the energy expenditure described previously as being directly proportional to the duty cycle d_j and the average signal power of the jammer. For a given power configuration, the enforcement of an upper bound on energy expenditure thus corresponds directly to an upper bound d_E on the allowable duty cycle d_j. The third component is the probability distribution of estimated RSS according to (13). In addition to these analytical components, the jammer is likely subjected to additional hardware constraints, such as a minimum switching time σ_j between the on and off states, effectively imposing a lower bound on Δ_j (and indirectly on

T_j). To avoid or bound detection, the attacker can choose a risk parameter δ and choose its attack parameters T_j and d_j subject to the risk-aversion constraint $\Pr[R > \tau] < \delta$.

Based on these various components and constraints, the design of an effective, efficient, and stealthy SFPJ attack with the optimal jamming period T_j^* and duty cycle d_j^* can be formulated as

SFPJ Attack Formulation

$$\begin{aligned} (T_j^*, d_j^*) &= \underset{\{(T_j, d_j)\}}{\arg\max}\, P_{pkt}(E) \\ \text{subject to} \quad & \frac{\sigma_j}{T_j} \le d_j \le d_E, \\ & \Pr[R > \tau] < \delta. \end{aligned} \tag{14}$$

Using this attack formulation, the adversary can select parameters that allow it to mount a highly effective attack with bounded resource cost and detection risk. We note that the attack formulation can be modified slightly if the RSS detection threshold is not linear. For example, if the RSS threshold τ is a function of the corresponding PDR, as suggested in previous approaches [5], then the constraint needs to be modified accordingly.

We could formulate our optimization problem for more advanced forms of jamming detection using richer RSS data. Since most production radio platforms only give very rough RSS data on the order of samples per packet or less, we decide to use the RSS versus PDR metric avoiding analysis that only works with expensive custom radios.

V. Empirical Study

In this section, we present an empirical study to show the efficacy of SFPJ attacks in terms of PDR and RSS for a wide variety of attack parameters. Our study is based on a hardware platform implementation using sensor motes for the transmitter and receiver and a software-defined radio for the attacker. After describing our platform setup, we present data to show the relationships between PDR and RSS for a variety of settings.

A. Evaluation Platform Setup

In our evaluation platform, the transmitter and receiver are implemented on Java SunSPOT nodes [17] which use the CC2420 chipset [18] to communicate according to the IEEE 802.15.4 standard [2]. The 802.15.4 protocol uses DSSS to encode symbols to be quite robust against interference, but does not specify any error correction at the packet level. Hence, if the receiver captures a packet with any non-zero number of symbol errors, the packet is discarded.

As described in Section II, we wish to ignore MAC-layer aspects in our study. We thus disabled the carrier sensing capability of the transmitter, so it can freely transmit a stream of packets. We set a default distance of 4 m between the transmitter and receiver with the jammer at a similar distance from the receiver, though these distances are varied in certain experiments. To characterize the average effect of the jammer on the communicating system, all of our experimental data reflects an average over a large number of data points. Specifically, each data point plotted in the figures in this section corresponds to an average over 1000 packets transmitted by the SunSPOT transmitter. We further note that the RSS values plotted in our figures are offset by 45 dB to translate the received signal strength indicator (RSSI) used on the CC2420 chip [18] to a standard mW reference using a dBm value. In our environment, the observed RSS in the absence of transmission or attack had an average value around -93 dBm.

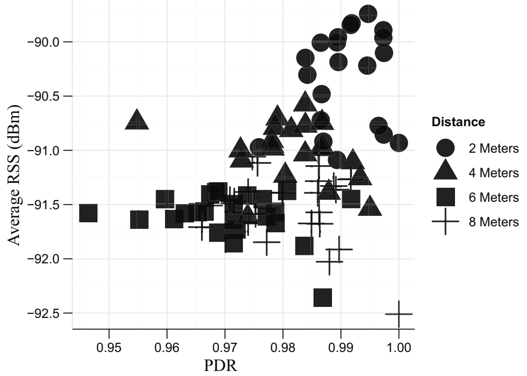

Fig. 6: We illustrate the implementation data for PDR and RSS for baseline performance in the absence of attack. The shape of each dot indicates the transmitter-receiver distance, and each data point corresponds to an average over 1000 measurements.

The computation of PDR is left to the receiver, as this is how the parameter would be evaluated in a realistic network scenario. Therefore, it is important to note that the PDR values reported in our figures are *observed* PDR and not *actual* PDR, meaning the receiver counts the number of correctly decoded packets from those that were received, not from those that were sent. In the presence of jamming attacks, this is an important distinction, as the receiver may not detect every packet transmission due to corruption of start- or end-of-frame sequences or other header information. In actuality, the observed PDR can be significantly higher than the actual PDR.

In the absence of the attacking signal, the baseline performance of this setup under benign conditions is illustrated in Figure 6. The figure shows the relationships between PDR and RSS for a variety of transmitter-receiver distances and over several implementation runs. For the distances that we tested, the PDR is typically above 95%, and the RSS readings are in the range of -89 dBm to -92.5 dBm, which is reasonable given the -93 dBm observed noise floor. It is interesting to note that at a distance of 2 m, the average RSS readings are considerably higher, which suggests that the jammer's attack tasks may be more difficult to attain in this case.

The adversarial component in our evaluation platform, namely the SFPJ attacker, was implemented on a USRP2 software-defined radio [19] using the GNURadio software package [20]. The set of jamming signal periods (measured in symbol durations) employed by the jammer is $T_j \in \{2^n * 22 : n = 1, \ldots, 6\}$. The transmission power emitted from the jamming radio was set empirically to the minimum power level

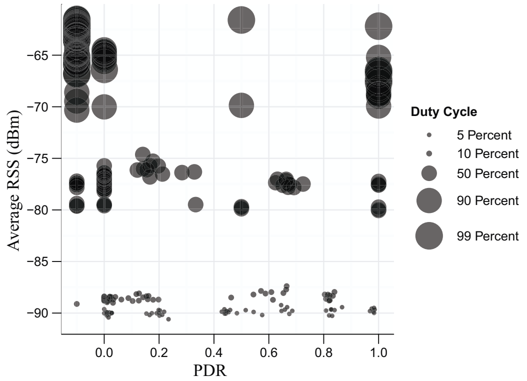

Fig. 7: We show the PDR and RSS data for a wide range of duty cycle parameters d_j. The size of each dot corresponds to the duty cycle parameter, and each dot represents the average over 1000 transmitted packets. The figure clearly demonstrates the relationship between duty cycle and RSS, showing that lower duty cycle jamming is harder to detect by RSS-based detection algorithms.

required to interfere effectively with packet communications.

B. Evaluating the Effect of Duty Cycle and Period

As a first experiment, we evaluate the effect of duty cycle and period parameters on the PDR and RSS measurements at the receiver. We implemented a jammer using the six values of T_j given above and varied the duty cycle d_j between 0.05 and 0.99. The results for several data points are illustrated in Figure 7, where the size of the data point indicates the duty cycle of the jammer. It is important to note that any data point with a PDR equal to -0.1 corresponds to a case where 0 correct packets and 0 erroneous packets were received during the sample period, making the observed PDR undefined. By observation of Figure 7, it is clear that there is a direct correlation between the duty cycle parameter d_j and the effect of the attack on RSS measurements at the receiver, as suggested by our analysis in Section III.

The samples at $d_j = 0.05$, given by the smallest dots, have low RSS readings, between -89 *dBm* and -91 *dBm*. This increases by about 1 *dBm* moving to $d_j = 0.1$, but increases drastically to above -80 *dBm* for $d_j = 0.5$. As expected, the RSS values for $d_j = 0.9$ and $d_j = 0.99$ demonstrate an even greater impact on RSS. It is interesting to note in Figure 7 that there are points where the receiver has a PDR of 100% even with $d_j = 0.99$. However, the reason for this is that the jammer blocks nearly all packets from even being observed by the receiver, and the 1 or 2 packets that get through are correctly decoded. Similarly, the points with PDR of 50% with $d_j = 0.99$ corresponds to one correct packet and one erroneous packet getting through to the receiver.

Since Figure 7 suggests that a low duty cycle attacker is often just as effective as a high duty cycle attacker, we next focus in on a duty cycle ranging only from $d_j = 0.02$ to $d_j = 0.2$ in Figure 8. This figure confirms that the duty cycle parameter is directly correlated to the observed RSS. Smaller dots again correspond to the lower duty cycle data points, again confirming the effectiveness of the low duty cycle attack.

As seen in Figure 8, the PDR still varies drastically for each duty cycle parameter d_j. To describe this phenomenon, we next evaluate the effect of the jamming period. In Figure 9, we decompose the data for $d_j = 0.02$ and $d_j = 0.04$ by jamming period parameter T_j. It can be clearly seen that a jammer with a 2% duty cycle is able to lower the packet delivery ratio to just 30% and a 4% duty cycle is able to lower the packet delivery ratio to only 5% provided that the attack period is chosen well. By inspection of the collection of figures, we can see that poor choice of attack period can also lead to higher PDR over a wide range of duty cycles, so an attacker must jointly consider the parameters.

C. Evaluating the Effect of Distance

Motivated by our earlier experiment studying the baseline performance as a function of the distance between the transmitter and receiver, we next study the effect of similar geometry under jamming. In this case, both the jammer and transmitter were set the same distance away from the receiver. Figure 10 illustrates the effect of jamming on PDR and RSS for duty cycle parameters ranging from $d_j = 0.04$ to $d_j = 0.16$. At the longer distances, the jammer with duty cycle $d_j = 0.04$ is able to reduce the PDR below 10% and keep the average RSS under -90 *dBm*. At shorter distance, the RSS increases significantly, indicating that the jammer has to work much harder at short distances, and it is difficult to keep the RSS low in such cases. All of the plots in Figure 10 have very similar characteristics, differing primarily by a slight shift in RSS. One interesting note is that in Figure 10(c), there is a cluster of points for $d_j = 0.12$ that appear out of place where neither benign operation nor attacked operation make sense; these points are due to an undiagnosed glitch in the data, but we kept them in the data set for completeness.

VI. Related Work

With the ever-increasing ubiquity of wireless devices in everyday life, making sure these devices are robust to attack is an important challenge. Denial of service (DoS) attacks are an important class of attacks that can be used to damage the operation of legitimate systems [21]. When considering system robustness, it is important to understand the most up-to-date attacks to be able to provide the best protection.

Jamming is a specific DoS attack type targeting availability of the wireless medium [1], [7]. Traditional jamming attacks assume an attacker with large amounts of energy resources, so there is no limit to the amount of energy an attacker may use, allowing for attacks that work very well against traditional spread spectrum technologies [22], [23]. However, as mobile devices expand into every aspect of our lives, the study of jamming feasibility in power-constrained devices introduces new challenges on the attack side and on the side of detection and defense [5], [10], [24].

Previous work has suggested that energy efficient jamming can be achieved by incorporating higher layer protocol information. Examples of such jamming attacks include control channel jamming [8], aiming to jam only the channels used

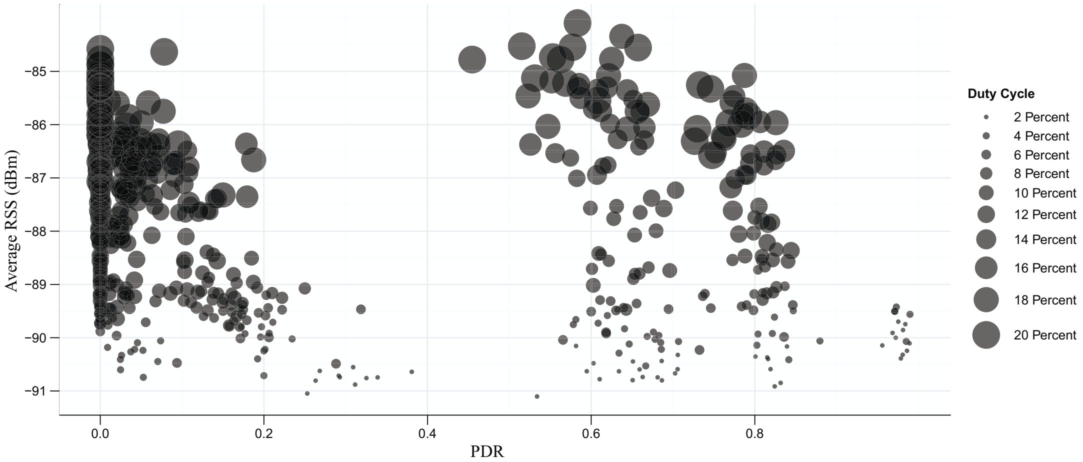

Fig. 8: We show the PDR and RSS data for a more focused range of duty cycle parameters d_j from 0.02 to 0.2. The size of each dot corresponds to the duty cycle parameter, and each dot represents the average over 1000 transmitted packets. The figure confirms the relationship between duty cycle and RSS.

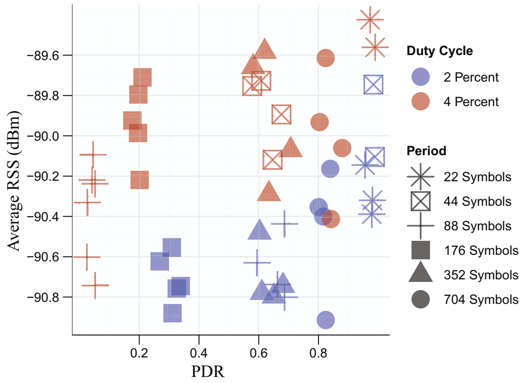

Fig. 9: We decompose the PDR and RSS data for $d_j = 0.02$ and $d_j = 0.04$ according to the jamming period parameter T_j. The data points cluster according to period, demonstrating the strong correlation between the attack parameters and the evaluation metrics.

for management and control in multi-channel systems; flow jamming [9], aiming to block end-to-end traffic flow at targeted locations in the network; and MAC-layer reactive jamming [10], which uses the structure of MAC layer CTS/RTS or ACK message exchanges to jam only when necessary. These approaches all depend on the attacker's ability to observe the target system and used learned information for attack formulation. This can be useful in avoiding detection, but it comes at the additional cost of frequent listening and observation. Our proposed technique, on the other hand, does not rely on observation or learning specific information beyond which communication protocols are used.

In addition to our technique, previous work has shown how to design optimal attack lengths and period for long form periodic jamming [25], where the jamming duration is on the order of packet duration. Reactive jamming [5], in which the attacker listens to the channel and jams whenever a transmission is detected, is another example of efficient jamming. Researchers have recently demonstrated the feasibility of reactive jamming using specialized software-defined radio hardware [11]. Since reactive jamming still requires the attacker to constantly sense the channel, it suffers the same energy expenditure problems as jamming using higher-layer protocol information. Random jamming [5], [22], in which the attacker switches the jamming signal on and off at random intervals, can also be used in both long and short form versions, and it has been shown to be effective against DSSS-based protocols. In our previous work, we showed that short form random jamming can be effective and efficient, but specialized signal processing can remove the jamming signal before the DSSS despreading operation [22].

VII. Conclusion

We have demonstrated that the precise design of a short form periodic jamming (SFPJ) attack can cause a devastating reduction of packet delivery ratio (PDR) without significantly increasing the received signal strength (RSS) observed at the receiver, even when spread spectrum based protocols are employed. Our analysis and attack methodology show how a jamming attacker can choose suitable attack parameters to effectively bypass RSS-based detection mechanisms. Moreover, our system implementation using a software-defined radio for the jammer and a commercial 802.15.4 based system for the transmitter and receiver validates our analytical claims. We believe that this study provides conclusive evidence that RSS-based detection is insufficient, thereby calling for future research into more robust jamming detection and mitigation.

References

[1] D. J. Torrieri, *Principles of Secure Communication Systems*, 2nd ed. Boston: Artech House, 1992.

[2] "IEEE 802.15.4-2006," 2006, http://standards.ieee.org/getieee802/download/802.15.4-2006.pdf.

[3] "IEEE 802.11b-1999," 1999, http://standards.ieee.org/getieee802/download/802.11b-1999.pdf.

[4] A. Molisch, *Wireless Communications*. John Wiley & Sons, Inc., 2005.

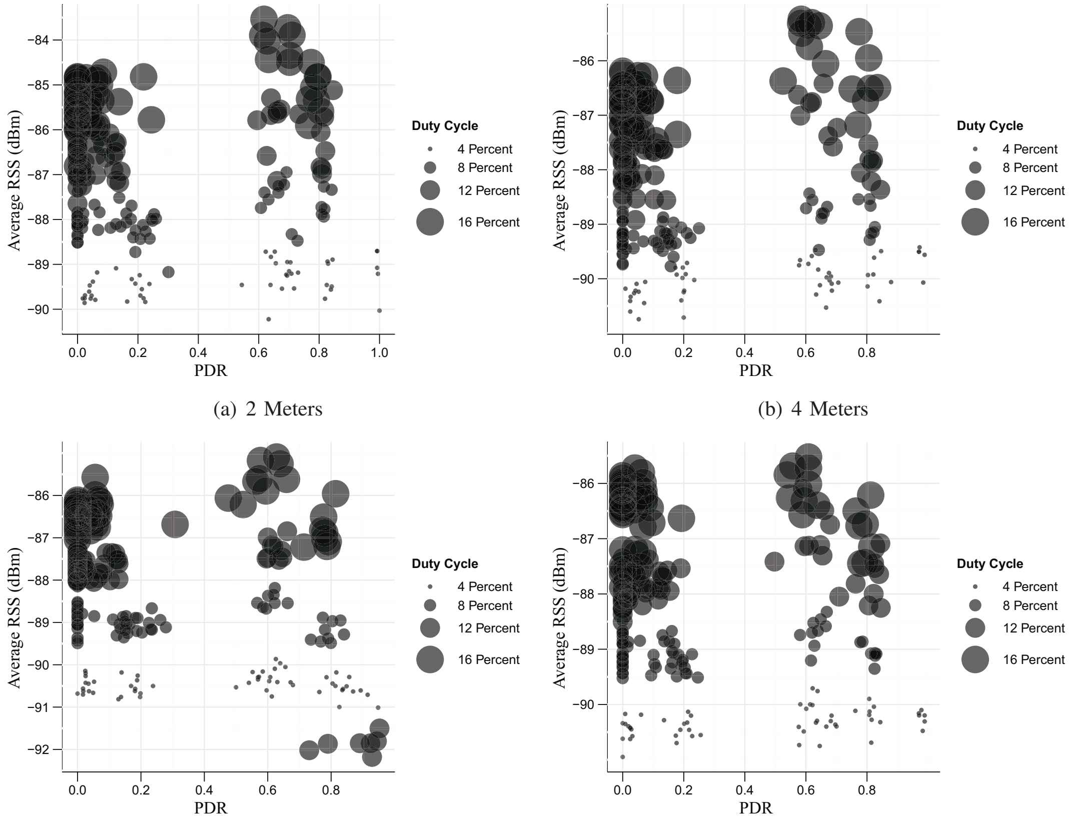

(a) 2 Meters

(b) 4 Meters

(c) 6 Meters

(d) 8 Meters

Fig. 10: We demonstrate how the PDR and RSS data varies with changes in the transmitter-receiver and jammer-receiver distances. In each plot, the dot size again corresponds to the duty cycle parameter d_j.

[5] W. Xu, K. Ma, W. Trappe, and Y. Zhang, "Jamming sensor networks: Attack and defense strategies," *IEEE Network*, vol. 20, no. 3, pp. 41–47, May/Jun. 2006.

[6] B. DeBruhl and P. Tague, "Adaptive filtering techniques for jamming mitigation," in *2nd International Conference on Pervasive and Embedded Computering and Communication Systems (PECCS'12)*, Feb. 2012.

[7] K. Pelechrinis, M. Iliofotou, and S. Krishnamurthy, "Denial of service attacks in wireless networks: the case of jammers," *IEEE Comm Surveys and Tutorials*, 2011.

[8] P. Tague, M. Li, and R. Poovendran, "Mitigation of control channel jamming under node capture attacks," *IEEE Transactions on Mobile Computing*, vol. 8, no. 9, Sep. 2009.

[9] P. Tague, D. Slater, G. Noubir, and R. Poovendran, "Linear programming models for jamming attacks on network traffic flows," in *Proc. 6th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'08)*, Berlin, Germany, Apr. 2008, pp. 207–216.

[10] D. J. Thuente and M. Acharya, "Intelligent jamming in wireless networks with applications to 802.11b and other networks," in *MILCOM'06*, Washington, DC, Oct. 2006.

[11] M. Wilhelm, I. Martinovic, J. Schmitt, and V. Lenders, "Reactive jamming in wireless networks: How realistic is the threat?" in *Proc. 4th ACM Conference on Wireless Network Security*, Hamburg, Germany, Jun. 2011.

[12] E. Bayraktaroglu, C. King, X. Liu, G. Noubir, R. Rajaraman, and B. Thapa, "On the performance of IEEE 802.11 under jamming," *INFOCOM 2008*, Apr. 2008.

[13] J. Bellardo and S. Savage, "802.11 denial-of-service attacks: Real vulnerabilities and practical solutions," in *Proc. USENIX Security Symposium*, Washington, DC, Aug. 2003, pp. 15–28.

[14] W. Feller, *An Introduction to Probability Theory and Its Applications*. John Wiley & Sons, Inc., 1957, vol. 1.

[15] B. Lathi, *Modern Digital and Analog Communication Systems*. Oxford University Press, 1998.

[16] W. Xu, W. Trappe, and Y. Zhang, "Channel surfing: Defending wireless sensor networks from interference," in *Proc. 6th International Conference on Information Processing in Sensor Networks (IPSN'07)*, Cambridge, MA, USA, Apr. 2007, pp. 499–508.

[17] "Java sunspot world," 2012, http://www.sunspotworld.com.

[18] "Chipcon cc2420 datasheet," 2011, http://focus.ti.com/lit/ds/symlink/cc2420.pdf.

[19] "Ettus research LLC," 2011, http://www.ettus.com/.

[20] "GNU radio," 2011, http://gnuradio.org/.

[21] A. D. Wood and J. A. Stankovic, "Denial of service in sensor networks," *IEEE Computer*, vol. 35, no. 10, pp. 54–62, Oct. 2002.

[22] B. DeBruhl and P. Tague, "Digital filter design for jamming mitigation in 802.15.4 communication," in *20th IEEE International Conference on Computer Communication Networks (ICCCN'11)*, Aug. 2011.

[23] K. Pelechrinis, C. Koufogiannakis, and S. V. Krishnamurthy, "Gaming the jammer: Is frequency hopping effective?" in *Proc. 7th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'09)*, Seoul, Korea, Jun. 2009.

[24] G. Thamilarasu, S. Mishra, and R. Sridhar, "A cross-layer approach to detect jamming attacks in wireless ad hoc networks," in *Proc. IEEE Military Communications Conference (MILCOM'06)*, Washington, DC, USA, Oct. 2006.

[25] Y. W. Law, M. Palaniswami, L. van Hoesel, J. Doumen, P. Hartel, and P. Havinga, "Energy-efficient link-layer jamming attacks against wireless sensor network MAC protocols," *ACM Transactions on Sensor Networks*, vol. 5, no. 1, pp. 1–38, 2009.