MITIGATION OF PERIODIC JAMMING IN A SPREAD SPECTRUM SYSTEM BY ADAPTIVE FILTER SELECTION *

Bruce DeBruhl and Patrick Tague

Carnegie Mellon University, Pittsburgh, U.S.A.

Keywords: Adaptive Anti-jamming, Digital Filtering, Jamming Mitigation, Spread Spectrum.

Abstract: Jamming has long been a problem in wireless communication systems. Traditionally, defense techniques have looked to raise the cost of mounting an equally effective jamming attack. One technique to raise the cost of jamming is direct sequence spread spectrum (DSSS) which spreads data over a wider bandwidth and has built-in error correction. To work around this, attackers have developed intelligent jamming techniques to minimize the cost of mounting attacks on these systems. To lower the cost of attacking a DSSS system, an attacker can use periodic jamming which alternates between an attacking and sleeping state. Previously, a digital filter has been used to mitigate a periodic jamming attack at the center frequency of the attacker. In this work, we expand this previous attack model by allowing an attacker to jam at any frequency and even to move to different frequencies in the channel. To defend against the more general attack, we propose the use of an adaptive filter selection technique. This technique monitors packet delivery ratio (PDR) at the receiver and uses this information to infer whether it is being attacked. If the receivers PDR is low, it activates a filter from a pre-defined filter bank and tests if performance improves. This process continues by activating different filters from the filter bank until adequate PDR performance is achieved. We show that this approach can search through a small set of filters and recover over 90% of packets with a search time of less than 3 seconds on average for an attacker who randomly chooses its center frequency.

1 INTRODUCTION

The open nature of wireless communications allows for transferring data without expensive and bulky wired connections but also opens the communications channel to any user. This open nature, allows for malicious users to affect legitimate users experience by intentionally broadcasting interference onto the wireless medium, an attack known as jamming (Torrieri, 1992).

Traditionally, to defend against jamming, spread spectrum techniques are used (Torrieri, 1992; Molisch, 2005). Two common variants of spread spectrum are frequency hopping spread spectrum (FHSS) and direct sequence spread spectrum. In FHSS a transmitter and receiver synchronously "hops" from channel to channel using a secret sequence. If the number of channels is sufficiently high, this is effective at raising the cost of jamming, since the jammer either needs to know the secret sequence or be able to mount a wide-band jamming attack (Pelechrinis et al., 2009). In DSSS the receiver maps each bit to many chips which are sent at rates greater than the data rate. This makes the legitimate signal hard to detect and allows for easier bit recovery by providing bit level error correction. Spread spectrum technique do not eliminate jamming but forces attackers to use more energy to mount an equally effective attack. Another approach to deter jamming, is to use detection technique and retreat from the jammer. At the MAC layer, jamming detection can be done by monitoring the packet delivery ratio (PDR) and flagging unexpected changes (Çakıroğlu and Özcerit, 2008). To improve accuracy of jamming detection in mobile systems, consistency checks of PDR with physical layer information like received signal strength can be used (Xu et al., 2006). Once a jamming attack is detected, a legitimate sys-

tem can retreat in space or in the spectrum and try to regain the ability to communicate.

To alleviate the additional cost of jamming a spread spectrum system or to stealthily avoid detection, intelligent jamming techniques have been studied (Thuente and Acharya, 2006; Law et al., 2009; Pelechrinis et al., 2011). These attacks can range from using cross-layer information (Chan et al., 2007; Tague et al., 2009) to signal conditioning (Xu et al., 2006; Bayraktaroglu et al., 2008). In this work, we focus on periodic jamming (Bayraktaroglu et al., 2008) in which the attacker continually alternates between a sleeping and attacking state. Periodic jamming is able to effectively attack common DSSS systems with lower energy usage than tone jamming.

It has been shown that a periodic jammer modulated to the same frequency as the legitimate system can be mitigated using a digital filtering technique (DeBruhl and Tague, 2011). This work added a digital high-pass filter at the base-band of an IEEE 802.15.4 (IEEE 802.15.4, 2006) receiver architecture to mitigate the effects from a periodic jamming attacks. In this work, we extend this previous work by considering an expanded attack in which the jamming center frequency is not necessarily the same as that of the legitimate signal and may occasionally change.

To defend against an attack that uses an arbitrary center frequency and occasionally changes its center frequencies we propose an adaptive filter selection mechanism. This filter selection mechanism monitors the system's performance and uses this information to infer if it is performing well. If it is not performing well it tries a different filter and tests to see if it improves performance. Thus for any single-carrier, narrow-band, periodic jamming attack a filter is eventually found. The major contributions of our work include the following.

- We propose the inclusion of a adaptive filter selection technique that searches through a set of filters to mitigate the effect of periodic jamming to a minimal level.

- We show three possible filter selection algorithms to use in our jamming mitigation technique.

- We evaluate the effect of our proposed jamming mitigation technique against various attack models using a software defined radio implementation.

The remainder of this paper is organized as follows. In Section 2, we introduce the models we use for the attacker and defender. In Section 3, we motivate the filter selection approach and introduce our proposed architecture. In Section 4, we present three filter selection algorithms and in Section 5, we present three attack models to test these algorithms. Finally,

In Section 6, we show implementation details and empirical results for our architecture and we conclude the paper in Section 7.

2 MODELS

In this section, we introduce our models for the legitimate system and attacker, building on that of existing work.

We consider a system where the sender and receiver are using the DSSS-based IEEE 802.15.4 physical layer with the 2450 MHz PHY specification (IEEE 802.15.4, 2006). DSSS is achieved in this protocol by mapping each group of 4 bits into a 32 chip sequence. Half these chips are modulated on the in-phase and half on the quadrature channel using offset quadrature phase shift keying (O-QPSK). At the receiver, the signal is demodulated and passes through a low-pass filter to select the appropriate channel. The receiver recovers chips using maximum likelihood estimation and then de-spreads 32 chips to the most probable 4-bit symbol.


Figure 1: This figure we show the modulated frequency domain model of both the attacker and the IEEE 802.15.4 signal. Although, the attacker's signal is much narrower then that of the legitimate 802.15.4 signal, it is sufficient to force packets to be dropped.


Figure 2: We show the cumulative sum of normalized spectral power for IEEE 802.15.4 and a periodic jammer. Because the jamming power in contained in a much smaller region than the 802.15.4 signal, the jamming signal can be filtered out without significantly degrading the legitimate signal.

Figure 3: We show the packet delivery ratio for a vanilla receiver and a receiver with a jamming mitigation filter at the base-band. The use of only a high-pass filter defends against attacks at the channels center frequency, while the filter is detrimental to the receiver when attacks are centered at other frequencies.

The 802.15.4 protocol uses a 2 byte cyclic redundancy check to test for packet errors (IEEE 802.15.4, 2006). This is calculated at the transmitter and appended to the packet. The receiver raises a flag if there is an error and passes the data up the stack if no error is detected. The use of the 2 byte CRC can pose a vulnerability since an intelligent attacker can completely jam an 802.15.4 system using the two byte CRC by only jamming one symbol in every packet.

We consider a periodic jamming attack which alternates between a sleeping and attacking state. This attack is advantageous because it looks to take advantage of the weakness introduced in using a two-byte CRC. The periodic jamming attack also reduces energy to allow for attacks to be mounted longer from low-power devices. This attack is often assumed to occur at the center frequency of the legitimate channel. We do not make this assumption but rather allow the attacker to choose any center frequency and to occasionally change its center frequency within the channel.

We revisit the model which is used to motivates using a digital filter to mitigate a baseband jamming attack (DeBruhl and Tague, 2011). Using this model they arrive at the plot shown in Figure 1, which motivates the fact that a filter can be used to mitigate the effect of periodic jamming at the base band.

To further motivate this fact, we present Figure 2, showing the cumulative normalized power for 802.15.4 and a periodic jammer in for a bandwidth centered around the carrier frequency. This plot shows 85% of of the attacker's power is contained in a region with only 10% of the signal power of 802.15.4. The related work showed that a filter at the center frequency could thus increase packet delivery ratio (PDR) for a periodic jammer at the center frequency of the desired channel from under 5% to over 90%.

However, the previous approach fails if the attacker adjusts the jamming center frequency beyond

the base-band filter. In Figure 3, we empirically show how the single base-band filter yields worse performance than an unfiltered receiver when the jamming center frequency is shifted from the center of the channel. In this work, we thus propose a technique to mitigate the effects of the proposed jammer with an arbitrary center frequency.

3 ADAPTIVE FILTER SELECTION

As discussed in Section 2, it has previously been shown that a filter can mitigate a jamming attack with a center frequency at the center of the channel. In this work, we allow that attacker to modulate to an arbitrary center frequency. We consider two types of attackers: one chooses its center frequency at the beginning of the attack and never changes it, and the other periodically changes its center frequency.

To mitigate such attacks, we propose to redesign the 802.15.4 receiver as shown in Figure 4. In this figure the light colored blocks represent the normal 802.15.4 receiver and the dark blocks represent proposed modifications to the receiver, which include filter-banks on the in-phase (I) and quadrature (Q) channels as well as a controller to select the filter.

Our research hypothesis is that there exists a set of filters Φ such that at least one filter $\phi \in \Phi$ can allow a receiver to achieve a high PDR under a periodic jamming attack at any center frequency assuming the attack is of equal power to the legitimate signal. We consider this through an empirical study in Section 6.

Since 802.15.4 uses a 2 byte CRC we propose using information derived from these consistency checks to keep track of packet delivery ratio for a given time period. These calculations are performed in the "Packet error detection" block of Figure 4.

Given a set of filters Φ and the PDR information, we now have to derive the "Filter controller" block from Figure 4. We propose a filter controller which tests if PDR is higher then a threshold δ, if it is than there is no need to attempt a defense, as there is no effective attack being mounted. On the other hand, if the PDR is lower than a threshold δ, there is either an attack being mounted or environmental conditions affecting the system. Since the system is already receiving a high amount of error, it is in the receiver's interest to try and mitigate this attack. To mitigate the attack the receiver chooses a filter ϕ from the set Φ and try receiving for a fixed amount of time τ. After τ seconds the receiver then determines if PDR is again under δ, if so it selects another $\phi \in \Phi$ to try to increase PDR. Once it finds a filter ϕ to maximize PDR, nor-
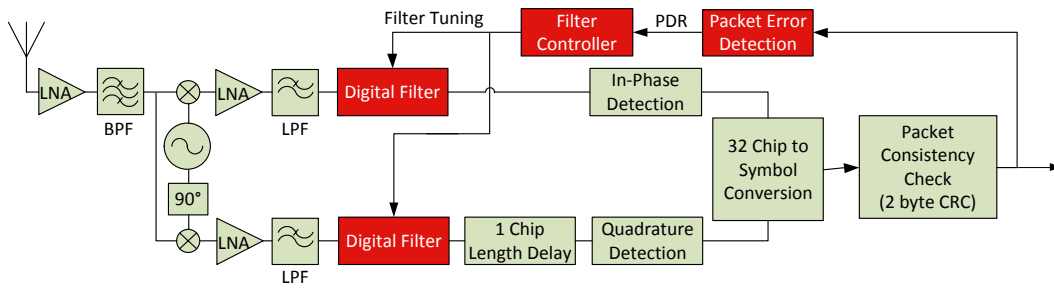
Figure 4: Our proposed approach incorporates a controller used to select from a digital filter bank mitigating the effect of a periodic jamming attack. In previous work, a single filter was proposed which is shown insufficient if an attacker is not constrained to a single center frequency.

mal communication can continue. In Section 4, we show various algorithms to select which filter to try once a reduced PDR is detected.

It is also possible to use environmental information to avoid unnecessary searches for filters. For example, the received signal strength could be used to test if a transmitter is relatively close or far away. If the received signal strength is low, it is likely that there is in fact not an attack occurring but simply poor communication conditions, in which case testing filters is not necessary. We leave such derivations of consistency checks as future work.

4 FILTER SELECTION ALGORITHMS

In this section, we introduce three methods to implement the filter search controller introduced in Section 3. The main goal in designing these algorithms is to allow for efficient searches while not opening easy attacks against the algorithms. We first introduce a straightforward filter selection technique, which searches the space by always choosing the next higher frequency filter. The second filter selection technique looks to decrease the search time by arranging the filters in order of decreasing filter width. Assuming a uniform likelihood of the attack occurring at any frequency, this decreases search time by having narrow filters that are less likely to find an attacker at the end of the search. The last filter selection algorithm is random filter search, which aims to increase entropy, making an intelligent frequency hopping attack difficult.

The first filter search algorithm we consider is *Incremental Filter Selection* as seen in Figure 5. This algorithm selects the next higher frequency filter whenever it is not receiving packets properly. If it is at the maximum frequency than it turns off the filter to see if it can obtain better performance, if this fails it con-
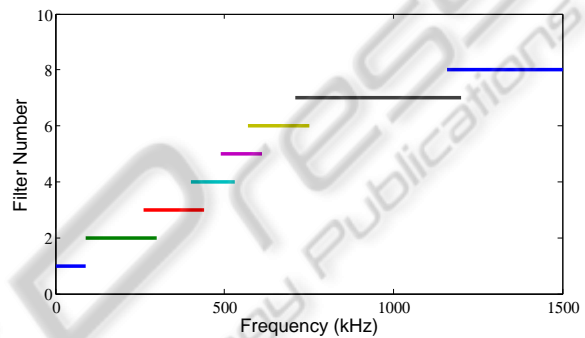

(a)

(b)

Figure 5: This figure illustrates the incremental filter selection algorithm. In (a) we see the ordering of filters by increasing filter center frequency. In (a) $0 Hz$ represents baseband. In (b) we see the transition diagram, showing the filter continuously shifting the filter to higher frequency unless it is maxed out.

tinues searching with the first filter again.

The second filter selection algorithm we propose is called *Ordered Filter Selection*. This algorithm, shown in Figure 6, orders filters from widest to narrowest. The reason for such an ordering is to allow for a quicker search of the bandwidth assuming an attacker is using a uniform distribution. This algorithm always start searching from the widest filter which we call filter 1. When its PDR first drops below δ it sets a "searching" variable we denote as x to 1. This variable remains at one until a solution is found such that

(a)
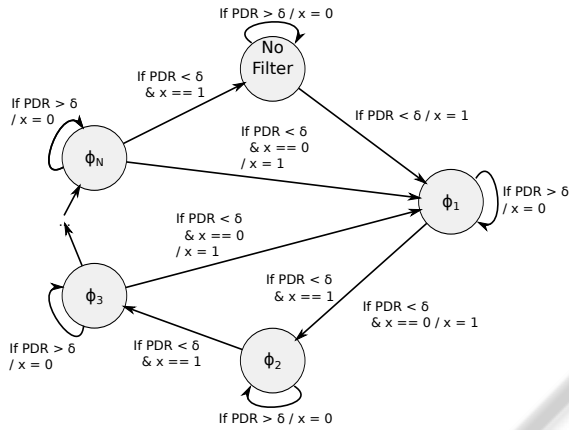
(b)

Figure 6: This figure shows the ordered filter selection algorithm. In (a) we see the selected filters arranged by decreasing bandwidth noting that the right most filter extends to 2.5 MHz. In (a) $0Hz$ represents baseband. In (b) we see the state diagram which uses the variable x to indicate whether it has found a solution. We use the notation "A/B" to indicate a transition is taken provided A is satisfied and B is executed during the transition.

the PDR is over δ at which point x is set to 0. When a jammer changes frequency and causes the PDR to drop, the receiver restarts the search at the widest filter and resets x to 1. Against an attacker choosing center frequencies randomly, this allows for a quick search on average by covering the spectrum as quick as possible.

The last filter search algorithm, shown in Figure 7, is called *Random Filter Selection*. This algorithm again uses a searching variable x like in the ordered filter selection algorithm. We add a permute state to this algorithm which is entered whenever $x = 0$ and PDR is below δ. During the permute state the ordered of the filters is randomly rearranged. We denote the mapping from the permuted set to the incrementally ordered set as $\phi_{\Pi(k)} = \phi_n$. For example, in Figure 6(a) the permutation is ordered $[7, 6, 3, 8, 2, 1, 4, 5]$ so $\phi_{\Pi(4)} = \phi_8$ indicates the 4^{th} filtered used is the

(a)
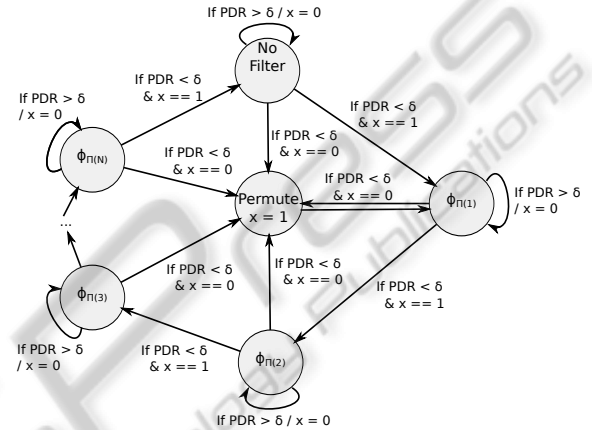
(b)

Figure 7: This figure shows the random filter selection algorithm. In (a) we show the order of filters at any point is unknown since it randomly permutes the filters whenever it leaves the not jammed state. In (a) $0Hz$ represents baseband. In (b) we see its behavior is similar to ordered jamming, except the added permute state changes the order of filters every time the jamming center frequency changes. We use the notation "A/B" to indicate a transition is taken provided A is satisfied and B is executed during the transition.

eighth highest in frequency. Continually changing the order of the search makes designing an attack against this algorithm particularly difficult.

5 ATTACK ALGORITHMS

In this section, we introduce three attacks to test our filter search algorithms. Each of these algorithms select a new center frequency every Ξ seconds. We propose one base-line attack and two attacks designed against particular filter selection algorithms.

The base-line attack we present is called *random center jammer* and is shown in Algorithm 1. In this attack a jammer randomly selects a new center frequency every Ξ seconds from a uniform distribution ranging from the center of the channel to an offset of

435

f_{max}. This attack represents an attacker who is aiming to attack an adaptive filtering but does not know what strategy and set of filters will be used. If the attacker had this knowledge, this attack should be the least effective.

An attacker aiming to disrupt an ordered filter selection with a known set of filters Φ could use an attack we call *Bi-Modal Jamming*. Since ordered filter selection always tries filter in the same order there is always two filters at the end of the order which take the longest to find, we denote these filters as ϕ_N and ϕ_{N-1} and denote the frequency at the middle of these filters as f_N and f_{n-1} respectively. Bi-modal jamming, shown in Algorithm 2, alternates its center frequency between f_N and f_{N-1} every Ξ seconds, forcing the algorithm to search through almost every filter.

The last attack algorithm is designed to attack an incremental filter search. In this algorithm, called *Decremental jamming*, a jammer aims to step its center frequency down by f_{dec} every Ξ seconds as shown in Algorithm 3. This should be most detrimental to incremental jamming since it searches in increasing frequencies. We use the constants f_{max} and f_{min} to indicate the maximum and minimum allowable frequency. Whenever the jammers frequency is decreased below f_{min} the frequency is increased by the difference in f_{max} and f_{min}. To find a slightly lower frequency it has a high search time as it goes through many different filters. In practice, the selection of f_{dec} is difficult since the width of filters is variable.

Algorithm 1: This algorithm performs random center jamming.

1: **while** True **do**
2:     $f_{offset} = random(0, f_{max})$
3:     $f_{center} = 2.425GHz + f_{offset}$
4:     wait(Ξ seconds)
5: **end while**

Algorithm 2: This algorithm performs random center jamming which randomly chooses a new center frequency every Ξ seconds.

1: **while** True **do**
2:     $f_{center} = f_N$
3:     wait(Ξ seconds)
4:     $f_{center} = f_{N-1}$
5:     wait(Ξ seconds)
6: **end while**

---

6 IMPLEMENTATION AND RESULTS

In this section, we present our implementation details as well as empirical results for our jamming mitigation techniques and attack models.

6.1 Implementation

We implement a proof-of-concept system using adaptive filter selection for jamming mitigation on the USRP2 platform (Ettus, 2011) using GNUradio (GNURadio, 2011) and an IEEE 802.15.4 implementation (Schmid et al., 2007). We allow a filter to be tested for half a second or $\tau = .5s$ and aim to keep a PDR of at least $\delta = 80\%$. We test all three jamming strategies suggested in Section 5 and set the amount of time an attacker stays on a center frequencies to $\Xi = 10$ seconds.

For implementation, a filter set Φ can be selected by using either analytical or empirical methods. The design of an optimal set of filters is out of the scope of this work, so we selected a set of 8 filters using empirical methods. The lowest frequency filter is a high-pass filter, the highest frequency filter is a low-pass filter, and all other filters are band-stop filters. To empirically tune the filters, we use our SDR setup to determine filter widths that would give adequate performance to our system in benign conditions and overlap to account for roll off. The selected filters are outlined in Table 1, all FIR filters (Tan, 2007) with a transition bandwidth of 40 kHz.

6.2 Results

In Figure 8, we show the adaptation for an incremental filter search against a random center jamming attack. This figure highlights the general operation of adaptive filter selection as well as some of the strengths and weaknesses than can be seen in it. The top plot in this figure shows the PDR for every half

---

Algorithm 3: This algorithm performs decremental jamming which decreases the jammers center frequency every Ξ seconds.

1: $f_{center} = 2.425GHz$
2: **while** True **do**
3:     $f_{center} = f_{center} - f_{dec}$
4:     **if** $f_{center} < f_{min}$ **then**
5:         $f_{center} = f_{center} + f_{max} - f_{min}$
6:     **end if**
7:   wait(Ξ seconds)
8: **end while**

Figure 8: We show a filter selection algorithm (incremental) running in real time.

Table 1: This table shows the empirically tuned filters.

Type	Low Cutoff	High Cutoff
High-Pass	–	90 kHz
Band-Stop	90 kHz	300 kHz
	260 kHz	440 kHz
	400 kHz	530 kHz
	490 kHz	610 kHz
	570 kHz	725 kHz
	710 kHz	1200 kHz
Low-Pass	1160 kHz	–

second time step during the experiment. The middle plot shows which of the eight filters is being used. If the value is zero than no filter is being used in that time step. The bottom plot shows the center frequency of the random center jammer. Figure 8 gives us a strong intuition to the strengths and weaknesses of our proposed jamming mitigation technique. This technique does provide for a system to mitigate the effects of a jammer that chooses one frequency and never changes quite well. It also allows for our system to mitigate the effects of an attacker who changes at a rate that is slower than our adaptation rate. When a random center attacker change frequency every 10 seconds our system is able to effectively recover over 70% of the data, increasing the amount of good data

from under 1%. To determine the bounds of our system, we consider whether our filter set is robust enough to cover the whole spectrum and the average search time for each filter selection algorithm.

We first consider whether our filters work at all frequencies. To do this, we programmed an adaptive filter selection scheme and used a static frequency jammer. We tested the the packet delivery ratio of the receiver once the adaptive filtering scheme had found an appropriate filter. The results for the PDR are shown in Figure 9. This graph indicates that a transmitter and jammer with equal sending power and distance from the receiver, our filtering technique yields over 90% packet delivery ratio once the appropriate filter is found.

Secondly, we consider the search time for each filter selection algorithm given various attacks. We implemented random center jamming, bi-modal jamming, and decremental jamming with $\Xi = 10$. We also implement all three filter selection algorithms and test the performance of all 9 combinations. Figure 10 shows the results for each search algorithm and attack model combination. It is clear that a random center jammer has less effect than the other two types of attacks. We also see that ordered filter selection has a slight advantage in a random attack. Bi-modal jamming is most effective against ordered filter selec-

Figure 9: We show the performance of our the appropriate filter from our filter-bank F at various frequencies compared to a receiver without filters and a single filter approach. Our approach is the only one offering high packet reception for attackers with an arbitrary center frequency.


Figure 10: We show how each filter selection algorithm performs under various attack models. Ordered filter selection does well except against a bi-modal attack designed against it.

tion because it is designed in such a way to make the algorithm always choose the last two filters, taking the greatest amount of time. With a decremental jamming attack ordered jamming again has the best performance. Random and incremental filter selection has similar performance. This is because choosing to a value for f_{dec}, the frequency step, to optimally attack incremental filter selection is not trivial. A large value will allow for quick searches and a small value makes it probable filters do not have to change every time step.

7 CONCLUSIONS

Intelligent jamming techniques have been shown to mount effective attacks against spread spectrum systmes with low energy. One of these attacks is periodic jamming which alternates between an attacking and sleeping state and is effective against DSSS systems. In this work, we consider how to mitigate the effects of a periodic jammer which can choose an ar-

bitrary center frequency and occasionally changes its center frequency. To do this, we monitor packet delivery at the receiver and if it is below a certain threshold activate a filter from a preselected filter bank. This process is repeated until an appropriate solution is found. We introduce three filter selection algorithms and three attack algorithms to consider the effects of different filter selection methods in our system. Once a filter is found, it is shown that over 90% of packets can be recovered, regardless of the jammers center frequency, offering a considerable gain over previous approaches. We also show the average search times of our filter selection algorithms range from 2-5 seconds.

REFERENCES

Bayraktaroglu, E., King, C., Liu, X., Noubir, G., Rajaraman, R., and Thapa, B. (2008). On the performance of IEEE 802.11 under jamming. *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE.*

Çakıroğlu, M. and Özcerit, A. T. (2008). Jamming detection mechanisms for wireless sensor networks. In *Proc. 3rd International Conference on Scalable Information Systems (InfoScale'08)*, pages 1–8, Vico Equense, Italy.

Chan, A., Liu, X., Noubir, G., and Thapa, B. (2007). Control channel jamming: Resilience and identification of traitors. In *Proc. IEEE International Symposium on Information Theory (ISIT'07)*, Nice, France.

DeBruhl, B. and Tague, P. (2011). Digital filter design for jamming mitigation in 802.15.4 communication. In *Proc. Intl. Conf. on Communications and Computer Networks*.

Ettus (2011). Ettus research LLC. http://www.ettus.com/.

GNURadio (2011). GNU radio. http://gnuradio.org/.

IEEE 802.15.4 (2006). IEEE 802.15.4-2006. http://standards.ieee.org/getieee802/download/802.15.4-2006.pdf.

Law, Y. W., Palaniswami, M., van Hoesel, L., Doumen, J., Hartel, P., and Havinga, P. (2009). Energy-efficient link-layer jamming attacks against wireless sensor network MAC protocols. *ACM Transactions on Sensor Networks*, 5(1):1–38.

Molisch, A. (2005). *Wireless Communications*. John Wiley & Sons, Inc.

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

Pelechrinis, K., Koufogiannakis, C., and Krishnamurthy, S. V. (2009). 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.

Schmid, T., Sekkat, O., and Srivastava, M. (2007). An experimental study of network performance impact of

increased latency in software defined radios. In *Proc. 2nd ACM workshop on Wireless network testbeds, experimental evaluation and characterization*, Montreal, Quebec, Canada.

Tague, P., Li, M., and Poovendran, R. (2009). Mitigation of control channel jamming under node capture attacks. *IEEE Transactions on Mobile Computing*, 8(9).

Tan, L. (2007). *Digital Signal Processing: Fundamentals and Applications*. Academic Press.

Thuente, D. J. and Acharya, M. (2006). Intelligent jamming in wireless networks with applications to 802.11b and other networks. In *Proc. 25th IEEE Communications Society Military Communications Conference (MILCOM'06)*, pages 1–7, Washington, DC.

Torrieri, D. J. (1992). *Principles of Secure Communication Systems*. Artech House, Boston, 2nd edition.

Xu, W., Ma, K., Trappe, W., and Zhang, Y. (2006). Jamming sensor networks: Attack and defense strategies. *IEEE Network*, 20(3):41–47.