Detection of Reactive Jamming in Sensor Networks

MARIO STRASSER, BORIS DANEV, and SRDJAN ČAPKUN
ETH Zurich, Switzerland

An integral part of most security- and safety-critical applications is a dependable and timely alarm notification. However, owing to the resource constraints of wireless sensor nodes (i.e., their limited power and spectral diversity), ensuring a timely and jamming-resistant delivery of alarm messages in applications that rely on wireless sensor networks is a challenging task. With current alarm forwarding schemes, blocking of an alarm by jamming is straightforward and jamming is very likely to remain unnoticed. In this work, we propose a novel jamming detection scheme as a solution to this problem. Our scheme is able to identify the cause of bit errors for individual packets by looking at the received signal strength during the reception of these bits and is well-suited for the protection of reactive alarm systems with very low network traffic. We present three different techniques for the identification of bit errors based on: predetermined knowledge, error correcting codes, and limited node wiring. We perform a detailed evaluation of the proposed solution and validate our findings experimentally with Chipcon CC1000 radios. The results show that our solution effectively detects sophisticated jamming attacks that cannot be detected with existing techniques and enables the formation of robust sensor networks for dependable delivery of alarm notifications. Our scheme also meets the high demands on the energy efficiency of reactive surveillance applications as it can operate without introducing additional wireless network traffic.

Categories and Subject Descriptors: C.2.0 [**General**]: Security and Protection

General Terms: Algorithms, Reliability, Security

Additional Key Words and Phrases: Jamming detection, reactive jamming, sensor networks

ACM Reference Format:

Strasser, M., Danev B., and Čapkun, S. 2010. Detection of reactive jamming in sensor networks. ACM Trans. Sensor Netw. 7, 2, Article 11 (August 2010), 29 pages.
DOI = 10.1145/1824766.1824772 http://doi.acm.org/10.1145/1824766.1824772

16

Author's address: M. Strasser, ETH Zurich, Inst. F. Techn. Informatik u. Kommunik. Netze. Gloriastrasse 35, 8092 Zurich, Switzerland; email: strasser@tik.ee.ethz.ch; B. Danev and S. Capkun, ETH Zurich, System Security Group, Universitätsstrasse 6, 8092 Zurich, Switzerland; email: {boris.danev, capkuns}@inf.ethz.ch.

DOI 10.1145/1824766.1824772 http://doi.acm.org/10.1145/1824766.1824772

1. INTRODUCTION

Initially motivated by battlefield intelligence, wireless sensor networks (WSNs) have expanded into a number of security and safety critical civilian applications including emergency response support, fire and burglar alarm systems, and the protection of critical infrastructures. Common to these applications is that they rely on dependable and timely delivery of alarm notifications. These alarms are typically raised by sensor nodes upon the detection of a sensed event (e.g., presence of an intruder) and must subsequently be forwarded to the network authority in a hop-by-hop manner. A sensor network that supports these applications must therefore guarantee the timely delivery of alarms even under jamming attacks.

The expected lifespan of such sensor network applications ranges from months to years and, given the limited power supply of sensor nodes, places high demands on the energy efficiency of the running algorithms. To meet these demands, existing surveillance applications [Gu and He 2007; He et al. 2006; Dutta et al. 2005; Strasser et al. 2007] combine low duty-cycling with reactive notification. Here, alarms are only transmitted upon detection of an event (i.e., for the network authority "no news is good news"). While such behavior is highly desirable in energy-constraint sensor networks, in conjunction with the low output power and limited spectral diversity of sensor node transceivers, it makes the alarm forwarding highly vulnerable to jamming-based denial-of-service attacks (i.e., alarm masking); these attacks have been shown to come at a low cost for the attacker while being particularly harmful to timely delivery of critical information [Wood and Stankovic 2002; Xu et al. 2005; Čagalj et al. 2007].

In principle, there are two solutions to counter jamming attacks on alarm forwarding: jamming mitigation and jamming detection. Common spread-spectrum-based jamming mitigation techniques such as FHSS or DSSS are, however, beyond the capabilities of current sensor nodes. Moreover, existing jamming detection techniques rely on the packet-delivery-ratio (PDR) and/or the received ambient signal strength as their main decision criteria [Poisel 2004; Xu et al. 2005; Noubir and Lin 2003; Li et al. 2007]. They have been shown to be well-suited for the detection of mid- or long-term jamming [Xu et al. 2005; Li et al. 2007], but detecting reactive packet or single-bit jamming remains a challenge: To begin with, jamming detection techniques that rely only on the CRC of a packet to decide whether it was received correctly can (in general) not distinguish between packet failures due to weak radio links and interference. Furthermore, assessing an accurate PDR is not practical in a reactive forwarding scheme as messages are sent very rarely and typically only on request or triggered by an external event. Finally, jamming does not necessarily cause a steady and high received signal strength value as only a small fraction of a packet has to be interfered with in order for the packet to be invalid [Poisel 2004; Gamma 2001; Noubir and Lin 2003]. A reactive jammer can thus keep the increase in the effective RSS value very low and hence avoid being detected.

In this work, we propose a novel jamming detection scheme as a solution to these problems. Our scheme is able to identify the cause of bit errors for

individual packets with high probability by looking at the received signal strength (RSS) during the reception of these bits; bit errors are detected either based on predetermined knowledge, error correcting/detecting codes, or limited node wiring in the form of wired node chains (n-tuples). The intuition behind this process is that if there was a bit error although the RSS value was high, this indicates external interference (intentional or unintentional); if the bit error was due to a weak signal (e.g., due to fast fading or shadowing), the RSS value should be low. This additional information allows an accurate differentiation of packet errors due to intentional interference from errors due to weak links, even in the case of a sophisticated (reactive) attacker that jams only a small portion of a packet.

We discuss the strengths and weaknesses of the proposed bit-error identification techniques and evaluate their jamming-detection performance analytically, by simulations, and experimentally with an implementation on BTnodes [BTnodes]. The evaluation results confirm that our solution meets the performance and accuracy requirements of (reactive) alarm forwarding protocols and enables the detection of advanced jamming attacks in which the attacker can freely choose the duration, strength, and beam width of the jamming signal. To the best of our knowledge, this work is the first to present a jamming detection scheme for sensor networks that enables the detection of reactive (single-bit) jamming on a per-packet basis. In summary, the contributions of this article are as follows.

—We present a novel jamming detection scheme for countering advanced (reactive single bit) jamming attacks in wireless sensor networks.
—We develop three different techniques for the identification of bit errors based on: predetermined knowledge, error correcting codes, and limited node wiring (n-tuples).
—We evaluate our solution by simulations and experimentally with a COTS sensor node platform (BTnodes).
—We analyze the threats on limited wiring and develop a low-power wire compromise detection scheme for the detection of malicious attacks on wires.
—We elaborate an analytical model for random, manual, and airdrop-based deployment of wired n-tuples.

The remainder of this article is organized as follows: After specifying the system and attacker model in Section 2, we discuss the impact and mitigation of reactive jamming in Section 3 and highlight the need for an efficient detection of reactive jamming. Our novel jamming detection scheme is presented in Sections 4 and evaluated in Section 5. We present our wire integrity verification protocol in Section 6, discuss related work in Section 7, and conclude in Section 8.

2. SYSTEM AND ATTACKER MODEL

As a typical application scenario for our jamming detection scheme, we consider the following setting: A wireless sensor network system is deployed in area $\mathcal{A}$

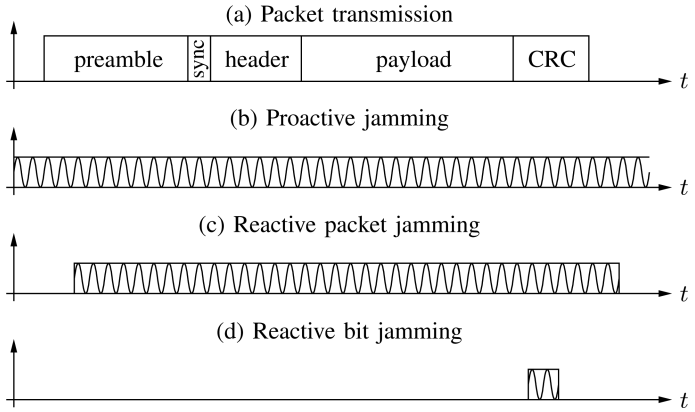

Fig. 1. Jamming types. (b) Proactive jammers keep the channel permanently occupied so that no transmissions [Fig. 1(a)] are possible whereas reactive jammers only jam once an ongoing transmission has been detected. (c) With reactive packet jamming, the attacker emits the jamming signal as soon as the transmission is detected and typically jams for an entire packet length. (d) With reactive bit jamming, the attacker targets its jamming signal at a specific part of the packet and keeps the jamming duration to a minimum.

(e.g., for the surveillance of a critical infrastructure). The main purpose of the network is to, upon the detection of an exceptional event (e.g., presence of an intruder), raise an alarm and forward it to the network authority. The network behavior is reactive, that is, alarms are sent when an exceptional event is sensed and they are resent if they are not acknowledged by the intended receivers. We assume that the node deployment is dense enough to ensure that alarm messages reach several neighbors and that, for security reasons, all traffic is encrypted and authenticated.

In this system, the attacker's goal is to interrupt or delay the alarm notification process by means of jamming. We assume that the attacker is in control of one/several static/mobile jamming devices but is unable to destroy or deactivate nodes without being noticed (e.g., tamper-responsive packaging triggers alarm upon misuse); otherwise, she could simply disable all nodes. To achieve her goal, the attacker can either proactively jam the intrusion area (Figure 1(b)) or reactively jam an alarm message once it is sent (Figure 1(c) and (d)). More specifically, we consider an attacker J that can freely choose its jamming location, frequency, rate, and strategy. We further assume that the maximal transmission power P_J of the attacker is finite, but we do not impose any restrictions on the attacker's energy supply. At each point in time, the attacker can freely choose the power P_j^i and beam width θ_j^i for a set $\{(P_j^1, \theta_j^1), (P_j^2, \theta_j^2), \ldots, (P_j^k, \theta_j^k)\}$ of emitted jamming signals, provided that $\frac{1}{2\pi}\sum_{i=1}^{k} P_j^i \theta_j^i \leq P_J$. The jammer can either be proactive or reactive [Poisel 2004; Xu et al. 2005]: Proactive jammers do not sense for ongoing transmissions but jam the channel permanently whereas reactive jammers initially solely sense for ongoing transmissions and start jamming only when a packet transfer has been detected. In order to remain undetected for as long as possible the attacker might decide to jam only a

certain fraction λ_j of all packets, to vary the beam direction and width, and/or to move between individual attacks.

The purpose of our jamming detection scheme is to detect jamming attacks on the message/alarm forwarding process. We assume that at least a fraction of all deployed nodes runs our detection algorithm and participates in the jamming detection. Upon the detection of a jamming attack, the nodes raise a jamming alert which is then either locally stored (e.g., as evidence for later investigations) or reported to the network authority by means of the existing alarm forwarding scheme. In the latter case, the attacker might extend the jammed region during her attack in order to also block these new alarms caused by her jamming. Consequently, if the jamming alarms raised by the initial set of nodes cannot escape the jammed area, these blocked jamming alarms must in turn also be detected by neighboring nodes and so forth. This means that the actual jamming detection can be composed of several iterative detection steps. However, since the same rules and conditions apply to all these steps, we consider only one such detection step and do not further discuss the jamming alarm reporting.

3. IMPACT AND MITIGATION OF REACTIVE JAMMING

Before we present our solution for the detection of reactive jamming in Section 4, we first highlight the importance of such a detection scheme by demonstrating how a reactive jammer can block communication in current wireless sensor networks with minimal exposure.

The reason for the weak jamming resistance of current MAC protocols for WSNs roots in the their dependency on a preamble/sync-byte header to mark the start of the packet header [Langendoen 2008]. This dependency makes the protocols extremely vulnerable to bit errors in the preamble, sync-byte, or packet header.

To demonstrate this vulnerability, we investigated the impact of reactive bit jamming on the performance of typical MAC protocols for WSNs. The experiments were conducted with BTnode sensor nodes that use an Atmel ATmega 128L microcontroller running at 8 MHz and a Chipcon CC1000 radio [BTnodes], the preamble and header length were set to 96 and 8 bytes, respectively. In order to enable the jamming of single bits without having to use expensive hardware, the transmission rate of the sender and receiver was reduced to 2.4 kBaud. We then implemented the jammer using an additional BTnode sending random data at a rate of 38.4 kBaud.

Our results clearly show that reactive bit jamming can efficiently block communication that uses current sensor network protocols with minimal exposure for the attacker. Specifically, in our experiments, we were able to achieve a jamming success rate of 90% by jamming only three bits in the packet header or in the sync byte (see Figure 2). Even more important, if the jamming was targeted at the sync-byte, the jammed packet transmissions were not even recognized as such by the network stack and were thus completely ignored by the nodes and not counted as packet losses. As a first step towards a better jamming resistance, we therefore introduce a more robust packet header detection technique

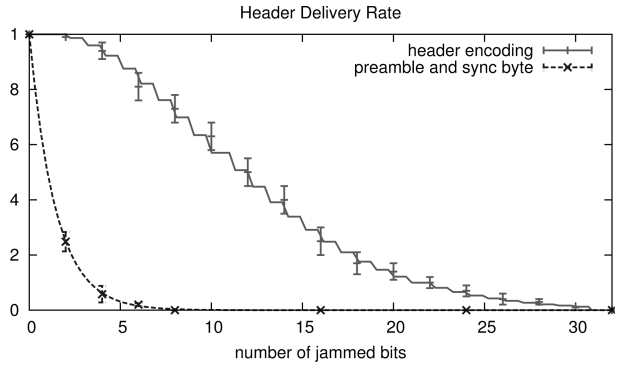

Fig. 2. Header delivery rate for our coding-based header detection algorithm and a common preamble/sync-byte-based approach. The lines show the expected theoretical results, the points the measured results; for each transmission the bit shuffling was based on a new (secret) seed. According to these results, a jammer that wants to mask a packet transmission with a probability of $> 90\%$ would have to jam only 3 bits if the common preamble/sync-byte method is used and more than 21 bits if our coding-based detection is used. Hence, the coding-based header detection is not only much more robust than common approaches, but also facilitates the detection of a jammer as it enforces a longer jamming duration.

that significantly increases the minimal time during which the jammer must interfere with a packet to block it.

With our header detection technique, before a packet is transmitted, the sender applies error correcting codes to the header and shuffles the encoded bits according to a pseudo random sequence based on the secret key shared by the sender and the receiver. As we shall see, this process ensures that a substantial part of the packet header must be jammed to prevent being decoded. Note that otherwise the MAC protocol in use is not modified; in particular a possibly required preamble—for example to account for imprecision in the nodes synchronization or to announce a transmission if low-power listening is used—is still transmitted.

The receiver (periodically) samples the channel according to the schedule of the employed MAC protocol and uses the received signal strength to assess whether a transmission is taking place. If a transmission is detected, the sender starts receiving it. However, as opposed to common practices, the sender will not wait for a predefined sync-byte to mark the start of the packet but will try do decode the header itself. Specifically, the receiver will receive a complete header length, unscramble the data, and try to decode it. Upon success, he receives the remainder of the packet; otherwise, he drops the first (i.e., oldest) bit, appends the newest received bit, and tries to decode the new input. This process is repeated until a packet is detected or the transmission ends (i.e., the received signal strength drops to the noise level for some time). The radio can then be turned off again and the node can fall back into sleep mode.

The main drawback of this packet detection algorithm is that the error correcting codes increase the header length and thus the required energy for a packet transmission. Whether this increase has a noticeable impact on the

overall energy consumption of the sender and/or receiver depends on the application and the MAC protocol in use. The well-known B-MAC protocol, for instance, uses preamble lengths which are several times longer than the packet header [Polastre et al. 2004]. A coding rate of up to 0.5 (i.e., an encoded header length that is twice as long as the unencoded header) has thus only a small impact on the duty-cycle and can usually be neglected. We point out that the above mentioned scheme is only related to the packet header detection; existing MAC wake-up schedules, LPL mechanisms, and preamble requirements are not affected.

We evaluated the performance of our proposed header detection scheme experimentally with the aforementioned setup based on BTnode sensor nodes. For the header encoding we used a Hamming (8,4) code that allows for correcting single bit errors, detecting all two bit errors, and detecting some three bit errors; the bit shuffling was performed with a linear feedback shift register.

The efficiency of our coding-based header detection algorithm compared to a preamble/sync-byte-based approach is shown in Figure 2. Assuming a flipping probability of 0.5 for the jammed bits and a jamming duration of x bits, in theory, the expected header delivery rate is 2^{-x} for the preamble/sync-byte method and about $\sum_{i=0}^{x} \binom{x}{i} 2^{-x} (1 - \frac{8-1}{8 \cdot 16})^{\binom{i}{2}}$ for the coding-based approach. This has been confirmed by our experiments: A jammer that wants to block a packet transmission with a probability of $> 90\%$ has to jam only 3 bits if the common preamble/sync-byte method is used but more than 21 bits if our coding-based detection is used. The results thus show that the coding-based header detection is not only much more robust than a preamble/sync-byte-based detection, but also increases the detection probability of a potential jammer as it enforces a longer jamming duration. However, as we shall see in the next section, even such a significantly longer jamming duration does not suffice to recognize jamming with current detection schemes.

4. DETECTION OF REACTIVE JAMMING

Traditional approaches for the detection of jamming in wireless sensor networks use the packet-delivery-ratio (PDR) and the received ambient signal strength as the main decision criteria. Jamming is detected as soon as the (averaged) PDR and/or the ambient signal strength exceeds a pre-defined threshold (see Section 7). Although these approaches are well-suited for the detection of mid- or long-term jamming, protecting the considered applications against targeted reactive packet or single-bit jamming remains a challenge: First, existing schemes rely only on the CRC of a packet to decide whether it was received correctly and thus can (in general) not distinguish between packet failures due to weak radio links and interference. Second, assessing an accurate PDR is not practical in a reactive forwarding scheme as messages are sent very rarely. Third, jamming does not necessarily cause a steady and high received signal strength (RSS) value as only a small fraction of a packet has to be interfered with in order for the packet to be invalid [Poisel 2004; Gamma 2001; Noubir and Lin 2003]. A (reactive) jammer can thus keep the increase in the

Algorithm 1. Jamming detection algorithm

```
resetJammingTest()
while true do
    (e, s) := getErrorSample()                    (A)
    x := detectInterference(e, s)                 (B)
    y := updateJammingTest(x)                     (C)
    if y = jamming then
        raise jamming suspicion
    else if y = no jamming then
        resetJammingTest()
    else
        do nothing as we need more evidence
    end if
end while
```

effective RSS value very low and can hence avoid being detected with current approaches.

Our novel jamming detection scheme does not suffer from these limitations. The central idea of our approach is to identify the cause of individual bit errors within a packet and to deduce therefrom whether the packet was jammed or just sent over a weak link. This is achieved as follows: Whenever a node receives a packet transmission, it not only receives the packet, but also records the RSS for each received bit of the packet.[1] Given a bit error, a node then deduces the root cause of this error by looking at the respective RSS value that was sampled during the reception of this bit. The intuition behind this process is that if there was a bit error although the RSS value was high, this indicates external interference (intentional or unintentional); if the error was due to a weak signal (e.g., due to fast fading or shadowing), the RSS value should be low. This additional information allows an accurate differentiation of packet errors that are caused by (un)intentional interference from errors that are caused by weak links, even in the case of a sophisticated attacker that jams only a small portion of a packet. We note that this also holds for m-ary modulation techniques where more than one bit is transmitted per symbol.

Our jamming detection algorithm comprises three steps (see Algorithm 1): (A) error sample acquisition (i.e., packet reception with RSS recording and identification of bit errors), (B) interference detection (i.e., error cause analysis), and (C) sequential jamming test. The function of the last step is to decide whether a detected interference was malicious or due to an unintentional packet collision and is only required if the probability of such a collision cannot be neglected. Next, we describe each of these steps in detail; the overall packet reception and jamming detection process that also considers proactive jamming is outlined in Figure 3.

[1]Some radios do not provide this accuracy but compute the averaged RSS value over a sequence of k bits (e.g., one byte). In these cases the algorithm as described might not detected jamming that affects less than k bits. This issue and possible mitigation strategies are further discussed in Section 5.

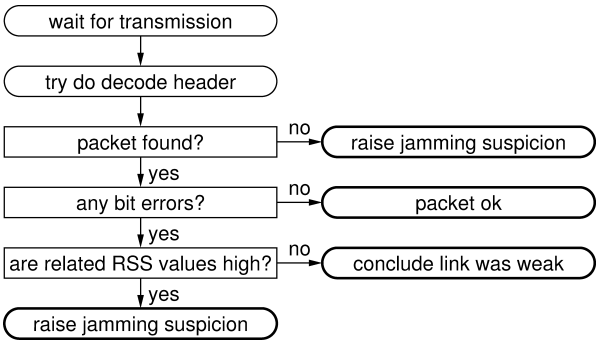

Fig. 3. Enhanced packet reception and jamming detection. Once a transmission signal is detected, the receiver tries to decode the presumed packet header. If it fails (i.e., if there is no transmission although the channel is busy) this might indicate (proactive) jamming and the sequential jamming test is updated. If the received packet contains bit errors, the root cause of this errors is analyzed and either the jamming test updated or the packet ignored.

4.1 Error Sample Acquisition

Whenever a node receives a packet transmission by radio, it receives the packet (even if it is not the intended receiver) and also records the RSS value for each received bit of the packet. That is, a node associates to each packet m a sequence s of RSS values corresponding to packet bits. Hence, for each bit in a packet the RSS at the time of its reception is also known; we denote by $m[i]$ and $s[i]$ the ith bit in the packet m and the ith RSS value in the RSS sequence s, respectively. An important aspect regarding the RSS sampling is the required storage overhead. For a packet size of 40 bytes and a typical RSS resolution of 8 bits per value, for example, this results in an overall overhead of $40 \cdot 8 = 320$ bytes. We point out that this RSS buffer is not required for each packet but only once. After the jamming detection for a received packet has been completed, the RSS values are no longer required and we can reuse the buffer.

The next and generally more challenging task is the identification of bit errors. In the simplest case, the content of m is predetermined (at least to a large extent) and known by the receiver. Finding bit errors thus reduces to comparing m with the reference packet $\hat{m}$. More formally, the error vector e can be computed as $e[i] := m[i] \oplus \hat{m}[i]$, where $\oplus$ is the exclusive or and $e[i] = 1$ if and only if the i-th bit is false. This error vector and the RSS sequence are combined to an error sample (e, s) which is then used as the base for the interference detection. The main drawback of this approach is that, because the packet content must be known to the receiver, the information conveyed by the packet is virtually limited to at best a few bits.

This limitation can be overcome by means of error detecting/correcting codes. These codes allow for detecting or even correcting bit errors in arbitrary messages. Where the original data can be recovered, the bit errors can be precisely detect by comparing the received and recovered data. If, however, a code word can be identified as being faulty but cannot be corrected, all bits in the word might equally be wrong and are thus marked as false. In addition to this possible loss of precision, the second drawback of using error correcting codes is

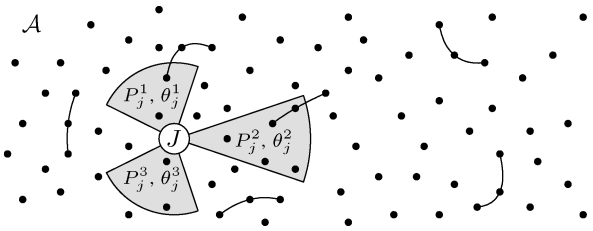

Fig. 4. System and attacker model. The network consists of N sensor nodes and—if detection is supported by limited wiring—of a number of wired chains of n sensor nodes (n-tuples), randomly deployed in $\mathcal{A}$. The attacker J can be located anywhere in $\mathcal{A}$ and freely choose the power P_j^i and beam width θ_j^i of the emitted signals. Her only constraint is her maximal transmission power P_J.

the overhead that the codes introduce. Depending on the strength of the code, the packet length (and thus the energy required for its transmission) might be several times higher than the length of the original packet.

A third, more elaborate way to acquire error samples is based on limited, short-range sensor node wiring in the form of wired node chains (n-tuples) as depicted in Figure 4. This method leverages the link redundancy (wired/wireless) provided by these n-tuples. Note that for simplicity, we assume that errors on the wired links can be neglected or are corrected (e.g., by forward error correction). If two nodes of a n-tuple are in the transmission range of a sending node, they will both receive the same packet transmission and record the corresponding RSS values. Two such independently received packet/RSS-sequence pairs (m_1, s_1) and (m_2, s_2) from the same packet transmission are then combined into an error sample (e, s), where $e[i] := m_1[i] \oplus m_2[i]$ and $s[i] := \min\{s_1[i], s_2[i]\}$. In general, error samples can be obtained in two ways:

In active monitoring, a node in an n-tuple sends a packet first over the wired and then over the wireless channel. The other nodes in the tuple receive both packets and record the RSS values of the packet received by radio; the RSS values of the (faultless) packet received by wire are set to infinity. Note that here the nodes in the tuple know when a packet is being transmitted and thus can still try to receive and compare its (payload) data, even if they fail to decode (part of) the header or payload. In passive monitoring, whenever a node in an n-tuple receives a packet that does not originate from a node in the tuple and that has not yet been received by wire, the node broadcasts the packet over the wire together with its respective RSS value sequence to all the nodes in the tuple. Each node in the n-tuple that receives a packet over the wired and over the wireless channel can then combine them to an error sample. Note that since the algorithm does not make any assumptions regarding the content of the packets, any regular application packet can be used to form a sample.

Being able to work with passive and active monitoring, our scheme allows to trade off n-tuple deployment density against energy consumption: In a passive system where the n-tuples do not introduce any additional network traffic, at least two nodes of an n-tuple must be in the transmission range of the sending node to detect jamming. In a (partially) active system where (some of) the

Algorithm 2. Error Sample Acquisition

```
function getErrorSample()
    while true do
        receive (m1, s1) by wire
        if the related packet (m2, s2) has already been received by radio then
            ∀i : e[i] := m1[i] ⊕ m2[i]
            ∀i : s[i] := min{s1[i], s2[i]}
            return (e, s)
        else if neighbor in the tuple will send it next then
            receive m2 by radio and record RSS into s2
            ∀i : e[i] := m1[i] ⊕ m2[i]
            ∀i : s[i] := min{s1[i], s2[i]}
            return (e, s)
        end if
    end while
end function
```

wired nodes periodically exchange probe messages,[2] only one node of such an n-tuple must be included in the jammed region. Moreover, in this case the wired neighbors of the sender know not only the content of the packet but also when its transmission takes place. Attacks where the data signal is entirely overshadowed by the attacker's signal or where the packet transmission is not (or only partially) recognized by the receiver's radio can thus also be detected.

4.2 Interference Detection

If a received packet contains at least one bit error, a node uses the measured RSS values to decide whether the identified errors are due to interference or due to a weak signal. The main intuition behind this approach is that if there was a bit error although the RSS value was high, this indicates external interference (intentional or unintentional); if the error was due to a weak signal (e.g., due to fast fading or shadowing), the RSS value should be low.

Here, we present a simple threshold-based mechanism to decide whether a packet error is due to interference. Let q be the counter for the number of recently observed packet errors due to interference. For each bit error in a packet, the respective RSS value is compared with a threshold S. If for at least one such case the RSS value is above the threshold S, q is increased, otherwise it is left unchanged. More formally, given an error sample (e, s), if $\exists i : e[i] = 1 \wedge s[i] > S$ then $q := q + 1$. With this simple decision algorithm, a reasonable selection of the threshold is important as it directly affects the false positive rate of the interference detection. If the threshold is set too low (i.e., too close to the noise level) all errors will be interpreted as jamming and the false positive rate will be 100%. If it is set too high (i.e., too close to the maximal RSS value) all bit errors are interpreted as being caused by weak links. The choice of (an optimal) S depends on the used radio and modulation scheme; it can be

[2]Note that since all traffic is encrypted, the attacker cannot distinguish probe from alarm messages.

Algorithm 3. Interference Detection

```
function DETECTINTERFERENCE(e, s)
    if ∃i : e[i] = 1 and s[i] > S then
        return 1
    else
        return 0
    end if
end function
```

predefined (e.g., as the result of experiments) or be computed on-the-fly (e.g., as a function of the RSS values of correctly received bits). If only links of poor quality are available, more sophisticated (but also more expensive) decision methods such as likelihood-ratio tests or Bayes factors can also be used [Baron 2007]. In our experiments, we achieved good results by adaptively changing S to the average signal strength of the last 10 successfully received packets.

4.3 Jamming Test

If the probability of packet collisions can be neglected, a node rises an alarm whenever it detects bit errors that were caused by interference. Otherwise, the result of the interference detection is taken as an input to a sequential probability ratio test (SPRT) [Zhang 1989] which is used to decide whether the recent packet errors (if any) were due to unintentional packet collisions or due to jamming. We assume that the nodes can assess the expected local interference (which is supposed to be low if the MAC works properly), either based on their knowledge about the used MAC and neighborhood or by using more sophisticated procedures such as those proposed in Zhou et al. [2005]. Let p_c be an upper-bound on the expected collision probability, τ_{FP} (τ_{FN}) be the targeted probability for a false alarm (missed attack), and q be the number of identified packet errors that were due to interference during the last k error samples. Given the probability p that the transmission of a packet fails, the probability that q out of k transmissions fail is $\binom{k}{q}p^q(1-p)^{k-q}$. The marginal likelihood that the observed packet errors were solely due to unintentional collisions (i.e., $0 \le p \le p_c$, hypothesis H_0) is then $p_0(k) := \int_{p=0}^{p_c}\binom{k}{q}p^q(1-p)^{k-q}\,\mathrm{d}p$; the marginal likelihood that there was jamming (i.e., $p_c \le p \le 1$, hypothesis H_1) is $p_1(k) := \int_{p=p_c}^{1}\binom{k}{q}p^q(1-p)^{k-q}\,\mathrm{d}p$. Hence, the log-likelihood ratio for H_0 and H_1 after k samples is

$$\eta(k) = \log\frac{p_1(k)}{p_0(k)} = \log\frac{\int_{p=p_c}^{1} p^q(1-p)^{k-q}\,\mathrm{d}p}{\int_{p=0}^{p_c} p^q(1-p)^{k-q}\,\mathrm{d}p}. \tag{1}$$

Now, if $\eta(k) \le \log\frac{\tau_{FN}}{1-\tau_{FP}}$ the nodes decide that there is no jamming and reset the sequence (i.e., set k and q to zero), if $\eta(k) \ge \log\frac{1-\tau_{FN}}{\tau_{FP}}$ jamming is detected and the nodes raise an alarm, finally if $\log\frac{\tau_{FN}}{1-\tau_{FP}} < \eta(k) < \log\frac{1-\tau_{FN}}{\tau_{FP}}$ no conclusive decision can be made yet and is deferred until there is more conclusive evidence available.

Algorithm 4. Jamming Test

```
function resetJammingTest
    k := 0; q := 0
end function

function updateJammingTest(x)
    k := k + 1; q := q + x;
    η(k) := SPRT(k, q)
    if η(k) ≥ log (1−τ_FN)/τ_FP then
        return jamming
    else if η(k) ≤ log τ_FN/(1−τ_FP) then
        return no jamming
    else
        return undefined
    end if
end function
```

5. PERFORMANCE EVALUATION

In this section, we evaluate the performance of the proposed jamming detection techniques. We therefore implemented them and conducted a series of experiments using COTS BTnodes (Atmel ATmega 128L microcontroller @ 8 MHz, Chipcon CC1000 radio) and Tmote Sky sensor nodes (TI MSP430F1611 microcontroller @ 8 MHz, Chipcon CC2420 radio). The experimental setup consisted of four nodes: One sender (node A), two receivers (node B and C), and one jammer (node J). For the wire-aided jamming detection, node B and C were connected over the I^2C bus, forming a two-tuple; a detailed performance and cost analysis of the wired communication is presented in Section 6.3.

A compilation of exemplary measurements for an undisturbed, jammed, and weak link between A and C is shown in Figure 5. The results confirm the validity or our approach and show that decoding errors caused by jamming can clearly be distinguished from errors caused by a weak radio signal by looking at the corresponding RSS values. However, these initial experiments also revealed some hardware constraints that limit the accuracy of the proposed detection techniques: In cases where the used radios do not provide an RSS value per bit but instead provide an averaged RSS value for a set of k bits, the algorithm might not be able to detected jamming that affects less than k bits. To mitigate the impact of this limitation, error correcting codes that enforce a minimal required jamming duration of $> k$ bits can be applied (see Section 3). Even if the attacker is able to split this duration into $> k$ one-bit jamming pulses and spread them over the entire packet, chances increase that at least one of these pulses is recognized by the detection scheme. Furthermore, packet based radio transceivers such as the Chipcon CC2420 typically rely on a particular synchronization preamble or training sequence to detect packet transmissions. If this preamble or training sequence is jammed, the corresponding transmission is simply ignored (an automatic CRC verification is usually not an issue as it can mostly be disabled). Simple bit- or byte-oriented radio transceivers such as the Chipcon CC1000 that provide a continuous data demodulation and RSS estimation are thus better suited for our purposes. Hence, we focus in the

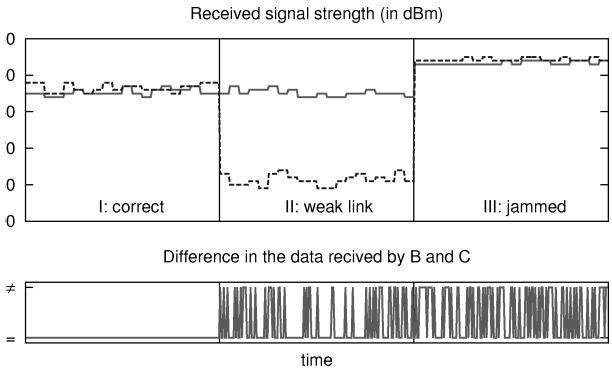

Fig. 5. Sample results obtained with our implementation and a CC1000 radio for three cases. I (adequate links and no jamming): both receivers are able to decode the packets and the packets do not differ. II (weak link from A to C): node C receives incorrect bits and thus the packets do not match; however, since the RSS of node C associated with the bit errors is low, the errors are correctly identified as non jamming related. III (with jamming): the RSS values for the observed bit errors are high for both receivers and thus the interference is correctly detected.

remainder of this section on bit or byte oriented radios and present the results obtained with our implementation for the CC1000 radio (i.e., the BTnodes) only. Nevertheless, we would like to point out that our basic considerations apply in general and thus our detection techniques in principle also work with packet-based radios.

The BTnode implementation uses our advanced header detection introduced in Section 3. To allow for the jamming of single bits with the jammer node J, the transmission rate of the sender and receiver was reduced to 2.4 kBaud whereas the jammer was sending random data at a rate of 38.4 kBaud.

We performed our experiments in two different scenarios: In the first scenario, the wireless connection between the sender and the receivers was fairly good, that is, the RSS of A's signal at B and C was about -55 dBm; in the second scenario, the connection between A and B was rather weak, that is, the RSS of A's signal at B and C was about -70 dBm. To make the jamming detection most challenging, the transmission power of the jammer was set to the lowest possible value for which the jamming was still effective (i.e., $>1\%$), which was 3 dBm for the scenario with the strong links and -5 dBm for the scenario with the weak link.

In both scenarios, we measured the performance of the four bit error detection techniques introduced in Section 4.1. Each technique was evaluated with a series of 1000 undisturbed packet transmissions, five times 2000 packet transmissions were a fraction of 2, 4, 8, 16, or 24 bit was jammed, and three times 2000 transmissions were a fraction of 8, 16, or 24 bit was suppressed (i.e., the transmission power at the sender was reduced to the minimum during their transmission in order to simulate a temporarily weak signal). The obtained results are summarized in Tables I and II. The second column shows the results for the case where the received packet is already known by the receiver, that is, the two techniques were the content of a packet is either predetermined or

Table I. Jamming Detection Performance for a Strong Link: True Positives/False Negatives—False Positives/True Negatives

number of jammed bits	message known or predetermined		
2	100% / 0%	—	0% / 100%
4	100% / 0%	—	0% / 100%
8	100% / 0%	—	0% / 100%
≥16	100% / 0%	—	0% / 100%
number of jammed bits	message encoded with ECCs		
2	100% / 0%	—	0% / 100%
4	100% / 0%	—	0% / 100%
8	99.9% / 0.1%	—	0% / 100%
≥16	100% / 0%	—	0% / 100%
number of jammed bits	comparison of two receptions		
2	84.9% / 15.1%	—	0% / 100%
4	94.1% / 5.9%	—	0% / 100%
8	98.8% / 1.2%	—	0% / 100%
≥16	100% / 0%	—	0% / 100%

Table II. Jamming-Detection Performance for a Weak Link: True Positives/False Negatives—False Positives/True Negatives

number of jammed bits	message known or predetermined		
2	100% / 0%	—	0% / 100%
4	100% / 0%	—	0% / 100%
8	100% / 0%	—	0% / 100%
≥16	100% / 0%	—	0% / 100%
number of jammed bits	message encoded with ECCs		
2	100% / 0%	—	0% / 100%
4	100% / 0%	—	0% / 100%
8	99.9% / 0.1%	—	0% / 100%
≥16	100% / 0%	—	0% / 100%
number of jammed bits	comparison of two receptions		
2	85.2% / 14.8%	—	0% / 100%
4	94.1% / 5.9%	—	0% / 100%
8	98.7% / 1.3%	—	0% / 100%
≥16	100% / 0%	—	0% / 100%

was first transmitted over the wire (active probing). The results in the third column represent the bit error location technique based error correcting codes and were obtained with a Hamming (8,4) code that allows for correcting single bit errors, detecting all two bit errors, and detecting some three bit errors. The results in the fourth column, finally, show the results for the case in which two wired nodes exchange their individual receptions.

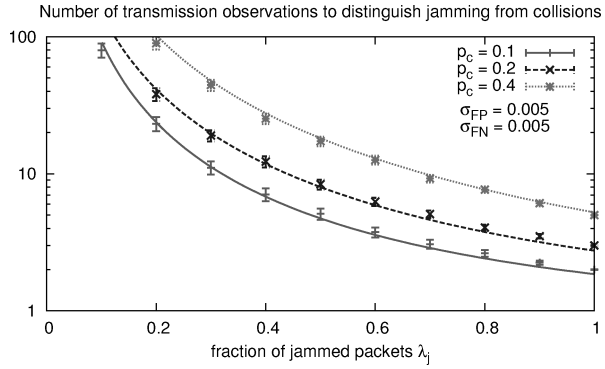

Fig. 6. Performance of the sequential jamming test. We observe that the larger the collision probability p_c and the lower the fraction of packets λ_j that the jammer jams, the longer it takes to detect the jammer; however, the lesser is also the impact of the jammer. If the attacker blocks an alarm (i.e., if $\lambda_j = 1$) the jamming will be detected after only five channel samples (for $p_c \leq 0.4$). Since alarm packets are immediately repeated if not acknowledged and because the attacker has to jam all alarms, this number will usually be reached after only a few seconds.

First of all, we notice that throughout our extensive experiments, no single false positive occurred (i.e., no bit error was erroneously identified as being caused by jamming). Furthermore, all false negatives (i.e., jamming-caused errors that were not identified as such) were due to inaccuracies in the bit error localization. More precisely, with some small probability it happens that the bit errors result again in a valid code word or that the two wired nodes observe exactly the same bit flips, respectively. We point out that in this respective, the measured results for the 2-tuple are actually worst case results because the more nodes are connected by wire, the less likely it is that all observe exactly the same bit flips.

5.1 Sequential Jamming Test

We next analyze the performance of the sequential testing which is required in cases where unintentional packet collisions cannot be neglected. Let λ_j be the fraction of all transmissions within the attacker's jamming range that she actually jams (i.e., the aggressiveness of the attacker) and p_c be the expected collision that a node observes. The expected number of channel samples that is faulty due to interference after k samples is thus $q = (1 - (1 - \lambda_j)(1 - p_c))k$. Inserting this expression into (1) and solving the equation $\eta(k) = \log \frac{1-\sigma_{FN}}{\sigma_{FP}}$ for k then yields the expected number of channel samples that must be processed before the jamming is detected. The resulting jamming detection performance as a function of p_c and λ_j is shown in Figure 6. The lines show the theoretical value, the points and σ-confidence intervals the results of our experiments. In a typical alarm forwarding scenario, the most relevant situation is one where the attacker intends to mask an alarm (i.e., $\lambda_j = 1$). We observe that in this case the jamming will on average be detected after only five channel samples (for reasonable collision probabilities $p_c \leq 0.4$). Due to the fact that alarm

packets are repeated if not acknowledged and because the attacker has to jam all alarms, we argue that this number will usually be reached after only a few seconds. We point out that although the packet retransmissions help to improve the detection time, they have no impact on the accuracy of the detection and are neither required nor part of our scheme.

5.2 Impact of the Node Density

Having evaluated the detection performance of our scheme if run on a node or n-tuple, we finally analyze the probability that the attacker's jamming activities are observed by a node or by an n-tuple in her proximity.

5.2.1 *Monitoring by Unwired Nodes.* Ideally, a node would receive and analyze every packet it overhears. However, given the stringent energy constraints of current sensor nodes, not all nodes in the transmission range of a sender usually receive a packet but only the set of intended receivers.

Let N_a be the average number of neighbors of a node, p_r be the probability that a neighbor which is not an intended receiver of a packet still receives and analyzes it, and p_d be the probability that potential jamming is correctly detected. As each packet has at least one receiver, the probability that the jammer is detected by the neighbors of the sender is $\geq 1-(1-p_d)^{1+(N_a-1)p_r}$. Let further N be the total number of nodes deployed in the deployment area $\mathcal{A}$ and $R(\cdot)$ be a function that maps transmission power levels to distances. The function $R(\cdot)$ depends on the nodes' radios and the environment they are deployed in. For the well-known physical communication model [Gupta and Kumar 2000], for instance, we have $R(P) := \sqrt[\alpha]{\frac{P}{\beta N_0}}$, where α, $2 < \alpha \leq 6$, is the so-called path-loss exponent, N_0 is the ambient noise power level, and β is the minimal required signal-to-noise-ratio to receive a packet. Given the transmission power P_a of a node and assuming (roughly) uniform node density, the node's expected number of neighbors N_a can then be estimated as $N_a \approx N\frac{R(P_a)^2\pi}{\mathcal{A}}$.

Figure 7 depicts the probability that the jammer is detected by the neighbors as a function of p_d, N_a and p_r. We observe that given the fairly high accuracy of typically $p_d \gtrapprox 0.9$ for the nodes' jamming detection (see above) an overall detection probability of ≥ 0.99 is already achieved with a single additional receiver (i.e., if $(N_a-1)p_r \geq 1$). Note that this result applies to a single transmission. Since alarm messages are repeated if not acknowledged, all blocked alarms will eventually be detected by at least one neighbor in practice.

5.2.2 *Monitoring by n-Tuples.* As mentioned in Section 4, the jamming detection performance depends not only on the n-tuple density, but also on whether the monitoring is passive or active. In a passive system at least *two nodes* of an n-tuple must be *in the transmission range of the sending node* for a minimum of two independent packet receptions are required. In an active system, the sender is part of the n-tuple and thus only *one (additional) node* of an n-tuple must be included *in the jammed region*. We next evaluate both scenarios for the case of a manual or airdrop-based node deployment where the

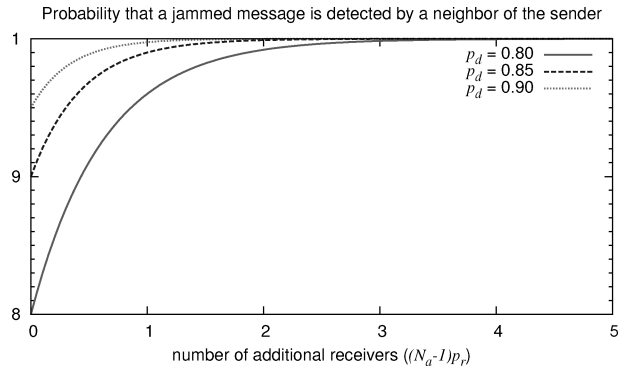

Fig. 7. Probability that a jammed packet is detected by at least one of the nodes in the proximity of the sender. Since the accuracy of the local jamming detection is already fairly high (i.e., $p_d \gtrapprox 0.9$) an overall detection probability of ≥ 0.99 is achieved with only one additional receiver (i.e., if $(N_a - 1)p_r \geq 1$).

position (orientation) of the n-tuples is chosen uniformly at random from the deployment area $\mathcal{A}$ (the interval $[0, 2\pi)$). The deployment orientation specifies the direction in which the nodes are located with respect to the first node. More precisely, the considered n-tuple deployment is as follows:

Let $u_i := (u_{i,1}, u_{i,2}, \ldots, u_{i,n})$ denote an n-tuple that is deployed in the deployment area $\mathcal{A}$. The order of the nodes in a tuple also determines their wiring: that is, for an n-tuple u_i and $1 \leq j < n$, node $u_{i,j}$ is connected to node $u_{i,j+1}$. We assume that all nodes are connected with wires of the same length l_w. The position of a node $u_{i,j}$ in the deployment area is denoted by $q_{i,j} \in \mathcal{A}$.

Given the deployment orientation ϕ and the position $q_{i,1}$ of the first node of the n-tuple u_i, the deployment of the remaining nodes $u_{i,2}$ to $u_{i,n}$ can be modeled as follows: For each node $u_{i,j}$, $2 \leq j \leq n$, imagine a disc $\mathcal{D}_{j-1} \subset \mathcal{A}$ centered at $u_{i,j-1}$ and of radius l_w. The position of $u_{i,j}$ is then chosen from $\mathcal{D}_{j-1}$ according to a random distribution defined by the (conditional) probability density function $f_D(q_{i,j}|\phi, q_{i,j-1})$. More formal, let $r_{i,j}$ be the euclidean distance between node $u_{i,j-1}$ and $u_{i,j}$, and $\alpha_{i,j}$ be the deviation of $u_{i,j}$'s position with respect to the deployment direction ϕ. To each n-tuple $u_i = (u_{i,1}, u_{i,2}, \ldots, u_{i,n})$ we can then associate a $(2n-1)$-dimensional (continuous) random variable $\mathcal{Q}_i := (\mathcal{Q}_{i,1}, \mathcal{R}_{i,2}, \Lambda_{i,2}, \ldots, \mathcal{R}_{i,n}, \Lambda_{i,n})$ taking values from the set $\{(q_{i,1}, r_{i,2}, \alpha_{i,2}, \ldots, r_{i,n}, \alpha_{i,n}) \,|\, q_{i,1} \in \mathcal{A} \wedge \forall j, 2 \leq j \leq n : 0 < r_{i,j} \leq l_w \wedge \forall j, 2 \leq j < n : -\pi \leq \alpha_{i,j} \leq \pi\}$ according to a random distribution defined by the (joint) probability density function $f_i(q_{i,1}, r_{i,2}, \alpha_{i,2}, \ldots, r_{i,n}, \alpha_{i,n}) = f_Q(q_1) f_\Phi(\phi) f_{R,\Lambda}(r_2, \alpha_2) f_{R,\Lambda}(r_3, \alpha_3) \cdots f_{R,\Lambda}(r_n, \alpha_n)$. Here, $f_Q(q_1)$ and $f_\Phi(\phi)$ represent the distributions on $\mathcal{A}$ and the interval $[0, 2\pi)$, respectively, and $f_{R,\Lambda}(r_j, \alpha_j)$ is the (joint) probability density function for the distance $r_{i,j}$ between node $u_{i,j-1}$ and node $u_{i,j}$ as well as the deviation α_j of $u_{i,j}$ from the deployment direction ϕ.

The probability density function $f_{R,\Lambda}(r_j, \alpha_j)$ reflects the actual deployment conditions and depends on the kind of deployment (random, manual, or airdrop-based) and on several physical parameters (e.g., the local terrain conditions or the rigidity of the wires). Usually the density function can be approximated

using appropriately parametrized (two-dimensional) Beta distributions (scaled to the interval $[0, l_w]$ and $[-\pi, +\pi]$). If $f_Q(\cdot)$, $f_\Phi(\cdot)$, and $f_{R,\Lambda}(\cdot)$ represent uniform distributions on $\mathcal{A}$, the interval $[0, 2\pi)$, and a disc of radius l_w, respectively, the resulting deployment is truly random. In any case, once the probability density functions are determined, the deployment of any set of n-tuples $\{u_1, u_2, \ldots, u_m\}$ can be formally described by the set of random variables $\{Q_1, Q_2, \ldots, Q_m\}$.

In our simulations, the deployment area $\mathcal{A}$ is a square with a side length of $a = 500\,\text{m}$. The position of the nodes in the tuples is chosen according to the probability density function $f_{R,\Lambda}(r_j, \alpha_j) = f_R(r) f_\Lambda(\alpha)$, where $f_R(r) = B(a_r, b_r)^{-1}(\frac{r}{l_w})^{a_r-1}(1-(\frac{r}{l_w}))^{b_r-1}$ and $f_\Lambda(\alpha) = B(a_\alpha, b_\alpha)^{-1}(\frac{\alpha+\pi}{2\pi})^{a_\alpha-1}(1-(\frac{\alpha+\pi}{2\pi}))^{b_\alpha-1}$ are two beta distributions and $B(a,b) = \int_0^1 t^{a-1}(1-t)^{b-1}\,\mathrm{d}t$. In order to assess a realistic parametrization, we conducted experiments using a dummy 3-tuple of 20 m length. Based on our (admittedly limited) experimental results, we chose the parameters $a_r = 10$, $b_r = 1.76$ and $a_\alpha = b_\alpha = 124$. For a given wire length l_w, this results in an expected n-tuple length of $0.85(n-1)l_w$, $\sigma \approx 0.1(n-1)l_w$ and an expected deviation of $\alpha = 0$, $\sigma \approx 0.14\pi$. In the analytical evaluation, we assume for simplicity that all nodes of an n-tuple lie on a straight line and that any two consecutive nodes in a tuple have the same distance $l/(n-1)$, where l is the expected length of the n-tuple.

(1) *Active Monitoring*. In order to determine the probability p_a that the jammed area is monitored by an n-tuple, we first compute the probability of the event X that at least one node of an n-tuple lies within a disc of radius r that is centered at the jammer. In a second step, this result is then generalized to the case where the attacker does not emit a single, omnidirectional signal but a set of directional signals. For simplicity, we assume that all nodes of an n-tuple lie on a straight line and that any two consecutive nodes in a tuple have the same distance.

Given the disc with radius r around the jammer, let z be the distance between the first node of an n-tuple and the center of the disc. The probability that at least one node of the tuple lies within the disc is then

$$p_X(r,l) = \int_{x=0}^{\infty} \mathsf{P}[X|z=x]\mathsf{P}[z=x] = \int_{x=0}^{r+l} \mathsf{P}[X|z=x]\frac{2x\pi\,\mathrm{d}x}{|\mathcal{A}|}\,. \tag{2}$$

Let $d_i = \frac{i-1}{n-1}l$ be the distance between the first and the ith node in the n-tuple. Now imagine a circle of radius d_i centered at the first node in the tuple. Given that the direction of a tuple is chosen uniformly at random, the probability that the ith node lies within the disc is then proportional to the central angle subtended by the two intersection points of this circle with the perimeter of the disc to the first node (see Figure 8(a)). As illustrated in Figure 8(b), this angle is

$$\alpha_X(d_i,x) := \begin{cases} 0 & \text{if } x<0 \text{ or } x>r+d_i \text{ or } x+r<d_i, \\ 2\pi & \text{if } 0 \le x+d_i \le r, \\ 2\arccos\left(\frac{x^2+d_i^2-r^2}{2d_i x}\right) & \text{otherwise.} \end{cases} \tag{3}$$

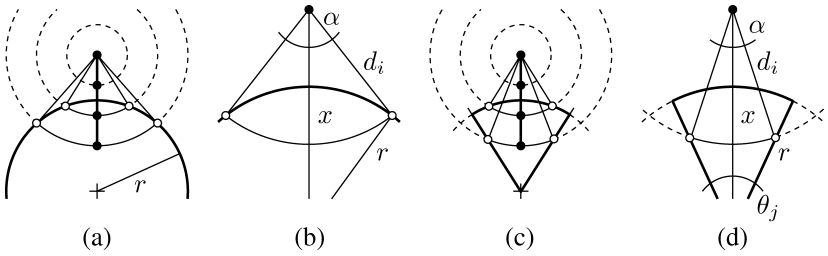

Fig. 8. Geometric relations between the position/orientation of an n-tuple and its possible intersections with a disc or a circular sector of radius r; the figures are not true to scale.

Hence,

$$\mathsf{P}[X|z=x] = \max_{1\le i<n} \alpha_X\left(\tfrac{i}{n-1}l, x\right)\frac{1}{2\pi} \tag{4}$$

and thus

$$p_X(r,l) = \frac{r^2\pi}{|\mathcal{A}|} + \int_{x=r}^{r+l} \frac{\max_{1\le i<n} \alpha_X\left(\tfrac{i}{n-1}l, x\right)}{2\pi}\frac{2x\pi\,\mathrm{d}x}{|\mathcal{A}|}\ . \tag{5}$$

For an omnidirectional jammer, the probability that the jammed area is monitored by a wired node is

$$p_a \ge 1 - \left(1 - p_X(R(P_j^1), l)\right)^m\ , \tag{6}$$

where m is the number of deployed n-tuples and $R(P_j^1)$ is the radius of the jammed area (i.e., the area in which jamming is effective and thus also detectable).

In the case of a general attacker that emits several directional signals specified by the set $\{(\theta_j^1, P_j^1), (\theta_j^2, P_j^2), \ldots, (\theta_j^k, P_j^k)\}$ (see Section 2), the probability that the jammed area is monitored by a wired node is equal to the probability that at least one node of an n-tuple lies within one of the respective circular sectors of central angle θ_j^i and radius $R(P_j^i)$. Considering only those n-tuples whose first node is enclosed by one of these sectors or their extension to a radius of length $R(P_j^i) + l$, this probability can be approximated as

$$p_w \gtrapprox 1 - \prod_{i=1}^{k}\left(1 - \left(1 - \frac{\theta_j^i}{2\pi}p_X(R(P_j^i), l)\right)^m\right)\ . \tag{7}$$

In order to account for those cases where the intersection of the (virtual) circles with radii d_i and r are outside of the circular sector given by θ_j^i (i.e., if $\theta_j^i < \pi$, see Figure 8(c)) the function $\alpha_X(d_i, x)$ has additionally to be substituted with $\min(\alpha_X(d_i, x), \alpha'_X(d_i, x))$, where

$$\alpha'_X(d_i, x) := 4\arcsin(\sqrt{y}/(2d_i)), \tag{8}$$
$$y = r^2 + (x - d_i)^2 - 2\cos(\theta_j^i/2)r(x - d_i)$$

is the angle subtended by the two intersection points of the circle with radius d_i and the radii of the circular sector to the first node (see Figure 8(d)).

(2) *Passive Monitoring*. Recall that with passive monitoring at least two nodes of an n-tuple must be in the transmission range of the sending node to detect a jamming attack. Let Z denote the event that at least two nodes of an n-tuple lie within a disc of radius r centered at the sender and let z denote the distance between the first node of an n-tuple and the center of this disc. The probability that at least two nodes of the tuple lie within the disc is then

$$p_Z(r,l) = \int_{x=0}^{\infty} \mathsf{P}[Z|z=x]\mathsf{P}[z=x] = \int_{x=0}^{r+l} \mathsf{P}[Z|z=x]\frac{2x\pi\,\mathrm{d}x}{|\mathcal{A}|}\ . \tag{9}$$

The probability that two nodes lie within the disc is proportional to the smaller of the two respective center angels. Let $\overset{2}{\max}$ be a function that returns the second largest value. We obtain

$$\mathsf{P}[Z|z=x] = \overset{2}{\underset{0\le i<n}{\max}}\, \alpha_X\left(\tfrac{i}{n-1}l, x\right)\frac{1}{2\pi} \tag{10}$$

and thus

$$p_Z(r,l) = \int_{x=0}^{r+l} \frac{\overset{2}{\underset{0\le i<n}{\max}}\, \alpha_X\left(\tfrac{i}{n-1}l, x\right)}{2\pi}\frac{2x\pi\,\mathrm{d}x}{|\mathcal{A}|}\ . \tag{11}$$

The probability that the neighborhood of a node is (passively) monitored by an n-tuple is

$$p_p \ge 1 - \left(1 - p_Z(R(P_a), l)\right)^m\ , \tag{12}$$

where m is the number of deployed n-tuples and $R(P_a)$ the nodes' transmission range.

The influence of the number of nodes n per tuple, the wire length l_w, the number of deployed tuples m, the node's transmission power P_a, and the size of the jammed area (i.e., P_j^i and θ_j^i) on the jamming detection performance of active and passive monitoring is depicted in Figure 9 and 10, respectively. The results show that even for short wires of about 3 m and a moderate jamming range of 100 m, only 80 3-tuples must be deployed per 1 km^2 in order that the jamming is (actively) monitored by at least one wired node with $p_a > 95\%$. In a purely passive scenario, about $(R(P_j)/R(P_a))^2$ times as many n-tuples have to be deployed to achieve the same protection as in an active scenario.

6. WIRE INTEGRITY PROTECTION

Our jamming detection scheme can handle sophisticated jamming attacks as demonstrated in the previous sections. However, it could be compromised by tampering the wired connectivity of n-tuples. For example, an attacker with physical access to the network could disconnect the wires, even if the sensor nodes themselves are tamper-resistant or not accessible. In such a way, she

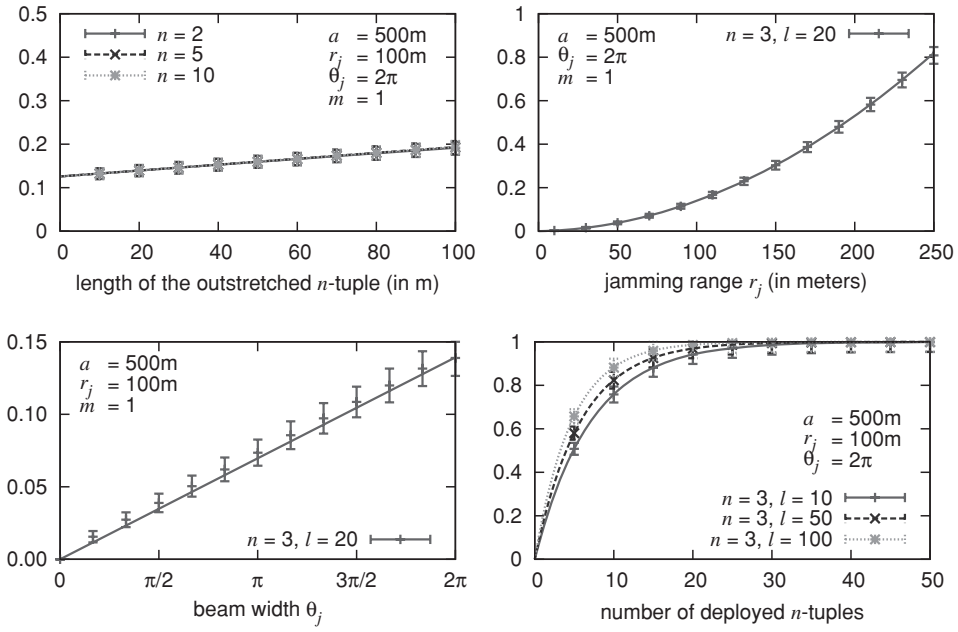

Fig. 9. Probability p_a that the jammed area contains at least one wired node. Full lines show the expected result according to our analysis, the points and σ-confidence intervals the simulation results. We observe that even very long tuples cannot benefit from more than 3 wired nodes per tuple whereas increasing the length of a tuple results in a linear increase of p_a. As p_a is proportional to the size of the jammed area, it increases exponentially with the jamming range and approximately linear with the beam width. Our findings also show that even for short wires with a length of about 3 m and a moderate jamming range of 100 m only 80 3-tuples must be deployed per 1 km^2 in order that the jamming is monitored by at least one wired node with $p_a > 95\%$.

could increase her chances to successfully jam the system alarms without being detected.

In particular, our scheme is vulnerable to four different attacks: disconnection, bridging, wiretapping and wire removal. In a disconnection attack the attacker achieves permanent disruption by cutting or unplugging. In a bridging attack, the attacker inserts a rogue (sensor) node between two wired sensor nodes. She therefore can control the wire and stay transparent to the system. In a wiretapping attack, the attacker instruments a wiretap (e.g., by direct electrical connection or by induction) and eavesdrops the communication. It can be used by the attacker to monitor if her malicious activity has been detected. A wire removal attack consists of removing wires or capturing entire n-tuples which is similar to a disconnection and classical node capture attack [Conti et al. 2008] respectively.

Given the vulnerability of our system to attacks on the wire connectivity, we propose a wire integrity verification protocol suitable for energy-constraint devices and analyze its security implications. In Section 6.3, we demonstrate the energy efficiency of our protocol in a real-life implementation.

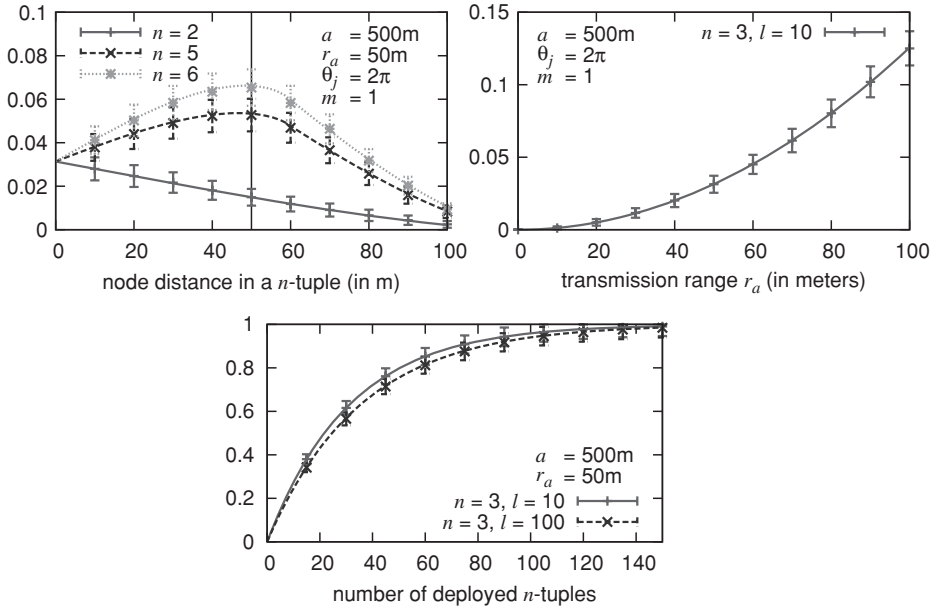

Fig. 10. Probability p_p that the neighborhood of a node is monitored by an n-tuple. We observe that for 2-tuples p_p decreases for longer wires and becomes zero if $l_w >$ twice the transmission range $r_a = R(P_a)$, whereas for $n > 2$ the probability p_p becomes maximal if the distance between two wired nodes in a tuple is $\approx r$. As p_p is proportional to the size of a node's neighborhood, it increases exponentially with the transmission range. Also, for $R(P_a) = R(P_j)/k$, in a passive scenario about k^2 times as many n-tuples have to be deployed than in an active scenario to achieve the same protection (e.g., $p_p >$95%).

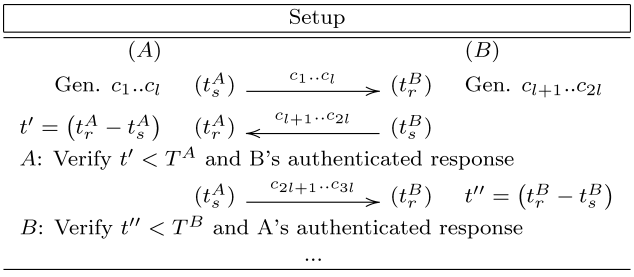

Fig. 11. Single session mutual authentication between A and B based on synchronous stream cipher. Note that the protocol is executed repetitively at regular intervals and that A and B alternate as protocol initiators.

6.1 Stream Cipher-based Wire Integrity Verification (SC-WIV)

The SC-WIV protocol provides mutual authentication by exchanging a sequence c of size l bits generated by a synchronized stream cipher. Node/wire compromise is detected by using appropriate local timeouts. The protocol is summarized in Figure 11 where A and B denote the two sensor nodes.

(1) *Setup*. Prior to deployment, A and B are preloaded with a shared symmetric key K_{AB} and an initialization vector IV_{AB}. These values are used to initialize the respective stream ciphers. The local timeouts for wire/node compromise detection, T^A and T^B, respectively, are calibrated before deployment based on measured round-trip times of messages and session frequency. We assume that both nodes A and B enter a special preemptive mode when sending and receiving bits. This mode postpones all other tasks in order to process the protocol messages as quickly as possible. We further assume that the stream cipher can generate one random bit upon request for a very long time. Trivium [De Cannière 2006] stream cipher can be used for such a purpose: it can be efficiently implemented in hardware and generate one bit per clock cycle (up to 2^{64} bits) from an 80-bit key and an 80-bit initialization vector.

(2) *Protocol Description*. A starts the first authentication session as follows: A sends a sequence of l bits to B and records the sending time t_s^A. Upon reception, B sends its response bits immediately and also records the sending time t_s^B. When A receives B's response, it calculates the elapsed time $t' = (t_r^A - t_s^A)$ and verifies that the sequence of bits is from B and is within the acceptable timeout $t' \leq T^A$. If the verification succeeds, A sends immediately the next sequence of bits to B. On timeout or if the authentication fails, an alarm is raised. Similarly, B receives A's response, calculates the elapsed time $t'' = (t_r^B - t_s^B)$, verifies that the sequence of bits is correct and with time $t'' \leq T^B$. Again, if the verification succeeds, B schedules to perform the next mutual authentication session after time T^F, where T^F is a shared time between successive protocol rounds.

The SC-WIV protocol is also used to signal whether application data will be transferred between the current and next protocol round: if the received protocol bit sequence c is inverted then data will be transmitted; if the received sequence is regular, there will be no data transmission.

6.2 Security Analysis

The presented SC-WIV protocol achieves disconnection and bridging attack. Some types of bridging attacks can still be possible, but will force the attacker to forward the protocol bit sequences. Wiretapping cannot be detected by our protocol but packet encryption on the application level is an appropriate mechanism in this case.

In the SC-WIV protocol, if the cipher cannot generate bits for a sufficiently long time, an initialization vector renewal procedure must be devised. We acknowledge that the protocol gives the attacker the possibility to guess the sequence of bits and thus to delay detection. More precisely, the attacker must guess one sequence of bits if the protocol session is initiated by the other entity and two sequences of bits when she initiates the protocol session. Therefore, the probability to delay detection for i runs (1 bit exchange) is $2^{-(i+\lfloor \frac{i}{2} \rfloor)}$ and thus decreases exponentially for successive protocol runs.

To summarize, there is a tradeoff between energy consumption and timely detection of wire compromise. Optimizing the energy consumption by reducing the number of bits increases the attacker's chances to delay detection. The actual delay gain depends on the frequency of successive protocol runs.

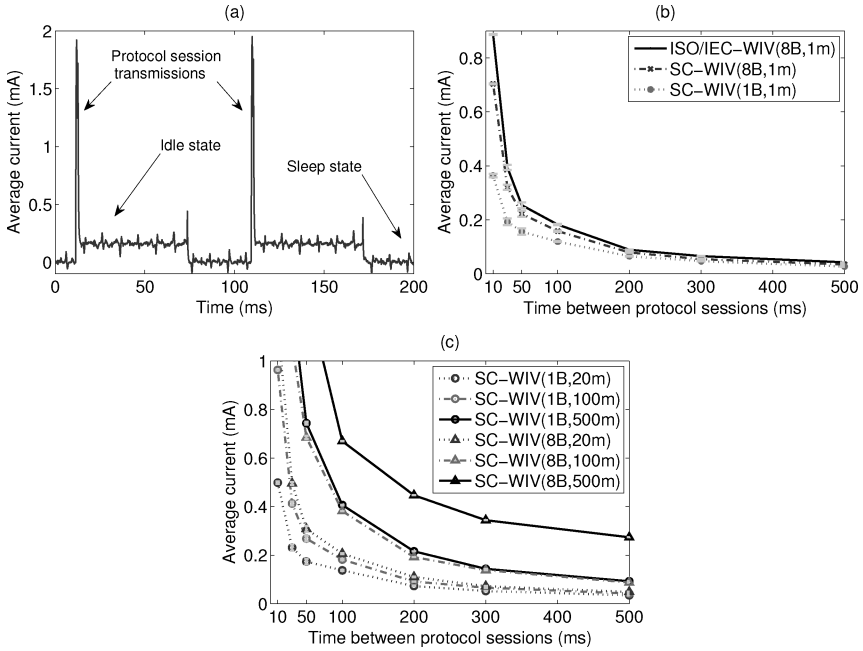

Fig. 12. Energy consumption: (a) A typical I^2C communication for SC-WIV. Current draw gradually returns to idle and sleep state after a session execution. (b) A comparison of the energy consumption of SC-WIV vs. ISO/IEC-WIV. SC-WIV is more efficient as it does not require encryption operations. (c) The SC-WIV energy efficiency for different wire lengths and bit sequence sizes.

6.3 Implementation Results

We implemented the SC-WIV protocol using Tmote Sky devices running TinyOS 2.x [TinyOS]. We used the pseudo-random number function in TinyOS with shared initialization vector to generate bits, emulating a stream cipher functionality. The I^2C bus on the Tmote Sky platform was used for wired communication with shielded and foiled Ethernet cables.

For comparison, we also implemented a standard challenge-response protocol for mutual entity authentication as described in ISO/IEC 9798-2. In this protocol, two sensor nodes A and B mutually authenticate themselves at regular intervals and alternate as session initiators. For the random nonce generation, we again used the pseudo-random number function in TinyOS. Symmetric key encryption/decryption was achieved with Skipjack block-cipher adapted from TinySec [Karlof et al. 2004] for the TinyOS platform. We refer to that solution as ISO/IEC-WIV.

Power consumption was measured at 3V using an Agilent 34411A multimeter. Figure 12(b) compares the ISO/IEC and SC-WIV protocols. As ISO/IEC-WIV uses Skipjack block cipher, the minimum size of an authenticated message is the minimum block size (8 bytes). So, we compared 8 bytes payload for both protocols and the SC-WIV protocol with 1 byte—the minimum possible payload in the I^2C implementation of the Tmote Sky platform.

Table III. Wire Length vs. Clock Speed

		SC-WIV Session Tx Time	
Wire Length	Clock Frequency	1 byte	8 bytes
1 m	95 kHz	1 ms	2.2 ms
20 m	39 kHz	2.1 ms	7.5 ms
50 m	21 kHz	3.5 ms	14 ms
100 m	12 kHz	6 ms	22.5 ms
500 m	3 kHz	25 ms	90 ms

From the results in Figure 12(b), we can see that at high session frequencies (i.e., if the time between protocol sessions is <100 ms) SC-WIV is more efficient than the ISO/IEC-WIV protocol and that SC-WIV with 1 byte additionally reduces the power consumption. For larger session frequencies, the average power consumption converges towards the idle/sleeping state power consumption of the sensor. This is expected as the sleeping state will be longer for lower frequencies, which tends to equalize the average power consumption (see Figure 12(a)).

It should be noted that the current I^2C specification minimum packet size consists of 7 bits address + 1 byte payload. Thus, SC-WIV protocol could be optimized to transmit one bit by changing the I^2C protocol. This would further improve the power consumption of the protocol.

Additionally, we tested SC-WIV over wires of 5, 20, 50, 100, and 500 meters length. (see Figure 12(c)). We found that I^2C can effectively be used for communication over longer wires. However, long distance wires significantly increase the energy consumption for both high (10 to 100 ms) and low (300 to 500 ms) frequencies. These findings also support the design of our SC-WIV protocol which optimizes the bit transmissions. The higher power consumption for longer wires is due to the decrease of clock speeds in order to compensate the resistance and capacitance in longer wires. Table III summarizes the clock speeds and transmission times required in the SC-WIV protocol session execution for 1 and 8 bytes.

In summary, our protocol is more energy efficient than standard solutions when a fast detection of wire integrity compromise (e.g., 10−50 ms) and/or an increased area coverage with longer wires (e.g., 20−100 m) is required.

7. RELATED WORK

The detection and mapping of jammed areas in the realm of wireless sensor networks has been studied by Wood et al. [2003]. Xu et al. [2005] advocate the usage of packet delivery ratio (PDR) along with either signal strength at the receiver (RSS) or location information as a consistency check for jamming detection. In the former case, jamming is detected if the PDR is low although the RSS value is high, in the latter case if the PDR is low although the senders are close. Unlike our work, in their scheme the RSS is not measured for each received bit but only when the PDR drops below a certain threshold and only after the last packet transmission that caused this to happen. Hence, reactive jamming during the packet transmission is not detected and the RSS measurement is only

related to the last packet transmission. The aim of our work is the detection of reactive jamming that might affect only a few bits of a packet. We therefore measure the RSS simultaneously to the packet transmission, for each bit (or with the granularity achievable), and we use it to identify the root cause of bit-errors, not as a general consistency check for a series of packets. Çakiroğlu and Özcerit [2008] propose two jamming detection schemes based on the PDR, the bad packet ratio (BPR), and the energy consumption amount (ECA) of the radio. In the basic scheme, jamming is detected if the PDR, BPR, or ECA values rise above or fall below specified thresholds. Altogether, five rules are specified, each focusing on a different set of jammer types. In the extended scheme, the nodes base their decision not only on their local view but exchange query and alarm messages with their neighbors in order to reduce the number of false positives at the expense of an increased communication overhead. A major drawback shared by the aforementioned schemes is that they depend on an accurate measurement of the current PDR and on an accurate assessment of its lower bound. Obtaining these values is difficult in most (reactive) sensor networks as messages are sent very infrequently and typically only on request or triggered by an external event. Moreover, as argued in this and previous work [Poisel 2004; Gamma 2001], jamming does not necessarily cause a steady and high RSS value as only a small fraction of a packet has to be interfered with in order for the packet to be invalid [Noubir and Lin 2003]. A (reactive) jammer can thus keep the increase in the effective RSS value very low and can hence avoid being detected by these schemes. Also, the proposed detection algorithms cannot distinguish between intentional and unintentional interference and timely delivery of alarm notifications is not considered.

A sequential jamming detection technique based on the number of erroneously received messages has been presented by Li et al. [2007]. The key idea of the proposed algorithm is that an increased number of observed message collisions during an observation window compared to the learned long-term average indicates a jamming attack. Using Wald's Sequential Probability Ratio Test they present optimal jamming attacks as well as network defense policies with respect to detection and notification time. Unlike our algorithm, this approach cannot distinguish between packet failures due to weak links and collisions. Thus, it is sensitive to changes in nodes' environment that influence the observed PDR and to (temporary) link failures or unanticipated changes in the traffic pattern which are likely to cause false alarms. Finally, none of the above mentioned schemes considers overshadowing, where the original packet is covered by a (maliciously inserted) second message.

The application of additional infrastructure in the form of wired shortcuts has been proposed before by Chitradurga and Helmy [2004] as well as by Sharma and Mazumdar [2005] with an objective to improve the energy efficiency of wireless networks. Čagalj et al. [2007] showed how wired node pairs can be used to build a wormhole in order to establish communication out of a jammed area. The main difference between our and their work is that theirs does not consider the use of wired tuples for jamming detection; they did neither consider the security of the wired links nor evaluated to what extent the proposed wirings are feasible.

8. CONCLUSION

In this work, we presented a novel jamming detection scheme for countering advanced (reactive single bit) jamming attacks in sensor networks. Our detection scheme is able to identify the cause of bit errors for individual packets by looking at the received signal strength during the reception of these bits. The scheme is thus well-suited for the protection of reactive alarm systems with very low network traffic. We presented and discussed three different techniques for the detection and localization of bit errors based on: predetermined knowledge, error correcting/detecting codes, and limited node wiring in the form of wired node chains (n-tuples). We further analyzed the threats on limited wiring and developed a low-power wire compromise detection scheme for the detection of malicious attacks on wires. The presented protocols and algorithm were evaluated analytically, by simulations, and experimentally on COTS BTnodes and Tmote Sky nodes. Since our scheme can operate without introducing additional wireless network traffic, it also meets the high-energy efficiency demand of reactive surveillance applications. To the best of our knowledge, this work is the first to present a jamming detection scheme for sensor networks that allows for the detection of advanced (reactive) single bit jamming on a per-packet basis. We further believe that this work provides useful insights into the utility of limited wiring as a means for securing wireless sensor networks.

REFERENCES

BARON, M. 2007. *Probability and Statistics for Computer Scientists*. Chapman & Hall/CRC.

BTnodes. BTnodes. http://www.btnode.ethz.ch/.

ČAGALJ, M., ČAPKUN, S., AND HUBAUX, J.-P. 2007. Wormhole-based antijamming techniques in sensor networks. *IEEE Trans. Mobile Comput.*

ÇAKIROĞLU, M. AND ÖZCERIT, A. T. 2008. Jamming detection mechanisms for wireless sensor networks. In *Proceedings of the 3rd International Conference on Scalable Information Systems (InfoScale)*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), ICST, Brussels, Belgium, Belgium, 1–8.

CHITRADURGA, R. AND HELMY, A. 2004. Analysis of wired short cuts in wireless sensor networks. In *Proceedings of the IEEE/ACS International Conference on Pervasive Services (ICPS)*. IEEE Computer Society Press, Los Alamitos, CA, 167–176.

CONTI, M., DI PIETRO, R., MANCINI, L. V., AND MEI, A. 2008. Emergent properties: Detection of the node-capture attack in mobile wireless sensor networks. In *Proceedings of the 1st ACM Conference on Wireless Network Security (WiSec)*. ACM, New York, 214–219.

DE CANNIÈRE, C. 2006. Trivium: A stream cipher construction inspired by block cipher design principles. In *Proceedings of the 9th International Conference (ISC)*.

DUTTA, P., GRIMMER, M., ARORA, A., BIBYK, S., AND CULLER, D. 2005. Design of a wireless sensor network platform for detecting rare, random, and ephemeral events. In *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks (IPSN)*. IEEE Computer Society Press, Los Alamitos, CA, 70.

GAMMA, D. 2001. *EW101: A First Course in Electronic Warfare*. Artech House.

GU, Y., AND HE, T. 2007. Data forwarding in extremely low duty-cycle sensor networks with unreliable communication links. In *Proceedings of the 5th International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, New York, 321–334.

GUPTA, P. AND KUMAR, P. R. 2000. The capacity of wireless networks. *IEEE Trans. Inform. Theory 46*, 2, 388–404.

HE, T., KRISHNAMURTHY, S., LUO, L., YAN, T., GU, L., STOLERU, R., ZHOU, G., CAO, Q., VICAIRE, P., STANKOVIC, J. A., ABDELZAHER, T. F., HUI, J., AND KROGH, B. 2006. VigilNet: An integrated sensor network system for energy-efficient surveillance. *ACM Trans. Sensor Netw. 2*, 1, 1–38.

KARLOF, C., SASTRY, N., AND WAGNER, D. 2004. TinySec: A link layer security architecture for wireless sensor networks. In *Proceedings of the ACM Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, New York, 162–175.

LANGENDOEN, K. 2008. *Medium Access Control in Wireless Networks*. Nova Science Publishers (Chapter Medium Access Control in Wireless Sensor Networks).

LI, M., KOUTSOPOULOS, I., AND POOVENDRAN, R. 2007. Optimal jamming attacks and network defense policies in wireless sensor networks. In *Proceedings of the 26th IEEE Conference on Computer Communications (INFOCOM)*. IEEE Computer Society Press, Los Alamitos, CA, 1307–1315.

NOUBIR, G. AND LIN, G. 2003. Low-power DoS attacks in data wireless LANs and countermeasures. *SIGMOBILE Mobile Comput. Comm. Rev. 7*, 3, 29–30.

POISEL, R. A. 2004. *Modern Communications Jamming Principles and Techniques*. Artech House.

POLASTRE, J., HILL, J., AND CULLER, D. 2004. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd ACM Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, New York, 95–107.

SHARMA, G. AND MAZUMDAR, R. 2005. Hybrid sensor networks: A small world. In *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*. ACM, New York, 366–377.

STRASSER, M., MEIER, A., LANGENDOEN, K., AND BLUM, P. 2007. Dwarf: Delay-aWAre Robust forwarding for Energy-Constrained Wireless Sensor Networks. In *Distributed Computing in Sensor Systems (DCOSS)*. Lecture Notes in Computer Science, vol. 4549/2007. Springer-Verlag, Berlin/Heidelberg, 64–81.

TinyOS. Tiny OS 2.x. http://www.tinyos.net/.

WOOD, A. D. AND STANKOVIC, J. A. 2002. Denial of service in sensor networks. *Computer 35*, 10 (Oct.), 54–62.

WOOD, A. D., STANKOVIC, J. A., AND SON, S. 2003. JAM: A jammed-area mapping service for sensor networks. In Proceedings of the *IEEE Real-Time Systems Symposium (RTSS)*. IEEE Computer Society Press, Los Alamitos, CA, 286–297.

XU, W., TRAPPE, W., ZHANG, Y., AND WOOD, T. 2005. The feasibility of launching and detecting jamming attacks in wireless networks. In *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*. ACM, New York, 46–57.

ZHANG, X. J. 1989. *Auxiliary Signal Design in Fault Detection and Diagnosis*. Springer-Verlag, Berlin.

ZHOU, G., HE, T., STANKOVIC, J., AND ABDELZAHER, T. 2005. RID: Radio interference detection in wireless sensor networks. In *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*. IEEE Computer Society Press, Los Alamitos, CA, 891–901.

Received August 2009; revised December 2009 and February 2010; accepted February 2010