Security in Sensor Networks
Yang Xiao (Eds.) pp. – - –
©2006 CRC Press

Chapter 12

MAC Layer Attacks in 802.15.4 Sensor Networks

Vojislav B. Mišić, Jun Fung, and Jelena Mišić
Department of Computer Science
University of Manitoba
Winnipeg, Manitoba, Canada R3T 2N2
E-mail: {vmisic, umfungj, jmisic}@cs.umanitoba.ca

1 Abstract

In this chapter, we consider the networks compliant with the recent IEEE 802.15.4 standard and describe a number of possible attacks at the MAC layer. Several of these attacks can be easily launched with devices that are fully compliant with the 802.15.4 standard, and we show that such attacks can introduce serious disruption. We also discuss some remedial measures that may help defend against those attacks, or at least alleviate their impact on the performance of the network.

2 Introduction

Wireless sensor networks pose many challenges with respect to security [2, 14, 16]. From the networking perspective, security threats may occur at different layers of the ISO/OSI model [2]:

- **Routing layer** attacks include spoofed, altered, or replayed routing information spread by an adversary, selective forwarding of packets, sinkhole attacks that attract traffic from a specific area to a compromised node (or nodes), Sybil attacks in which a compromised node assumes many identities, acknowledgement spoofing, injecting corrupted packets, neglecting routing information, or forward messages along wrong paths [6, 12].

- **MAC layer** attacks typically focus on disrupting channel access for regular nodes, thus disrupting the information flow both to and from the sensor node; this leads to a DoS condition at the MAC layer [16]. Security at the MAC layer has been mostly studied in the context of 802.11 MAC layer [4, 8, 12, 17] but sometimes also in the more general context of different types of attacks [1, 11, 16].

- Finally, **physical layer** (jamming) attacks consist of the attacker sending signals that disrupt the information flow through radio frequency interference. Jamming at the MAC level may be accomplished through sending large size packets with useless information.

Recently, IEEE has adopted the 802.15.4 standard for low rate Wireless Personal Area Networks (WPANs) [7, 5]. As 802.15.4-compliant WPANs use small, cheap, energy-efficient devices operating on battery power that require little infrastructure to operate, or none at all, they appear particularly well suited for building wireless sensor networks [5]. Hence, all aspects of 802.15.4 network operation and performance—security included—should be investigated and thoroughly analyzed.

In the discussions that follow, we identify a number of security threats that might affect the operation of 802.15.4 networks. Since the IEEE Std 802.15.4.[7] defines only the physical (PHY) and medium access layers (MAC) of the low rate WPAN, we focus on the MAC layer attacks, with particular emphasis on DoS attacks [16].

The chapter is organized as follows. In Section 3, we describe the operation of an IEEE 802.15.4-compliant sensor network, including the security provisions prescribed by the standard. Section 4 lists and briefly describes two classes of possible attacks at the MAC level, distinguished by compliance to the operation of the MAC protocol or lack thereof. Section 5 discusses the impact of some of those attacks, while Section 6 describes the manner in which some of those attacks could be prevented and/or their effects minimized. Section 7 concludes the chapter and outlines some directions for future work.

3 Operation of the 802.15.4 MAC

In an IEEE 802.15.4-compliant WPAN, a central controller device (commonly referred to as the PAN coordinator) builds a WPAN with other devices within a small physical space known as the personal operating space. Two topologies are supported: in the star topology network, all communications, even those between the devices themselves, must go through the PAN coordinator. In the peer-to-peer topology, the devices can communicate with one another directly – as long as they are within the physical range – but the PAN coordinator must be present nevertheless. The standard also defines two channel access mechanisms, depending on whether a beacon frame (which is sent periodically by the PAN coordinator) is used to synchronize communications or not. Beacon enabled networks use slotted carrier sense multiple access mechanism with collision avoidance (CSMA-CA), while the non-beacon enabled networks use simpler, unslotted CSMA-CA.

In beacon enabled networks, the channel time is divided into superframes which are bounded by beacon transmissions from the coordinator [7]. All communications in the cluster take place during the active portion of the superframe, the duration of which is referred to as the superframe duration *SD*. During the (optional) inactive portion, nodes may enter a low power mode, or engage in other activities at will.

The active portion of each superframe is divided into equally sized slots; the beacon transmission commences at the beginning of slot 0, and the contention access period (CAP) of the active portion starts immediately after the beacon. Slots are further subdivided into backoff periods, the basic time units of the MAC protocol to which all transmissions must be synchronized. The actual duration of the backoff period depends on the frequency band in which the 802.15.4 WPAN is operating; in the highest, 2.4GHz band, the maximum data rate is 250kbps [7].

A part of the active portion of the superframe may be reserved by the PAN coordinator for dedicated access by some devices; this part is referred to as the contention-free period (CFP), while the slots within are referred to as the guaranteed time slots (GTS).

3.1 The CSMA-CA algorithm

During the CAP period, individual nodes access the channel using the CSMA-CA algorithm, the operation of which is schematically shown in Fig. 1. The algorithm begins by initializing NB to zero and CW to 2;

the variable $NB = 0\,..\,macMaxCSMABackoff - 1$ represents the index of the backoff attempt, while the variable $CW = 0, 1, 2$ represents the index of the Clear Channel Assessment (CCA) phase counter.

If the device operates on battery power, as indicated by the attribute *macBattLifeExt*, the parameter *BE* (the backoff exponent which is used to calculate the number of backoff periods before the node device attempts to assess the channel) is set to 2 or to the constant *macMinBE*, whichever is less; otherwise, it is set to *macMinBE* (the default value of which is 3). The algorithm then locates the boundary of the next backoff period; as mentioned above, all operations must be synchronized to backoff time units.

In step (2), the algorithm generates a random waiting time k in the range $0\,..\,2^{\mathrm{BE}} - 1$ backoff periods. The value of k is then decremented at the boundary of each backoff period. Note that the counter will be frozen during the inactive portion of the beacon interval, and the countdown will resume when the next superframe begins.

When this counter becomes zero, the device must make sure the medium is clear before attempting to transmit a frame. This is done by listening to the channel to make sure no device is currently transmitting. This procedure, referred to as Clear Channel Assessment (CCA), has to be done in two successive backoff periods, as shown by steps (3) and (5) in Fig. 1. If both CCAs report that the channel is idle, packet transmission may begin.

If the channel is busy at the first CCA, the values of i and BE are increased by one, while c is reset to 2, and another random wait is initiated; this is step (4) in the flowchart. In this case, when the number of retries is below or equal to *macMaxCSMABackoffs* (the default value of which is 5), the algorithm returns to step (2), otherwise it terminates with a channel access failure status; it is up to the higher protocol layers to decide whether to re-attempt the transmission as a new packet or not. However, if the channel is found busy at the second CCA, the algorithm simply repeats the two CCAs starting from step (3).

Before undertaking step (3), the algorithm checks whether the remaining time within the CAP area of the current superframe is sufficient to accommodate the CCAs, the data frame, the proper interframe spacing, and the acknowledgment. If this is the case, the algorithm proceeds with step (3); otherwise it will simply pause until the next superframe, and resume step (3) immediately after the beacon frame.

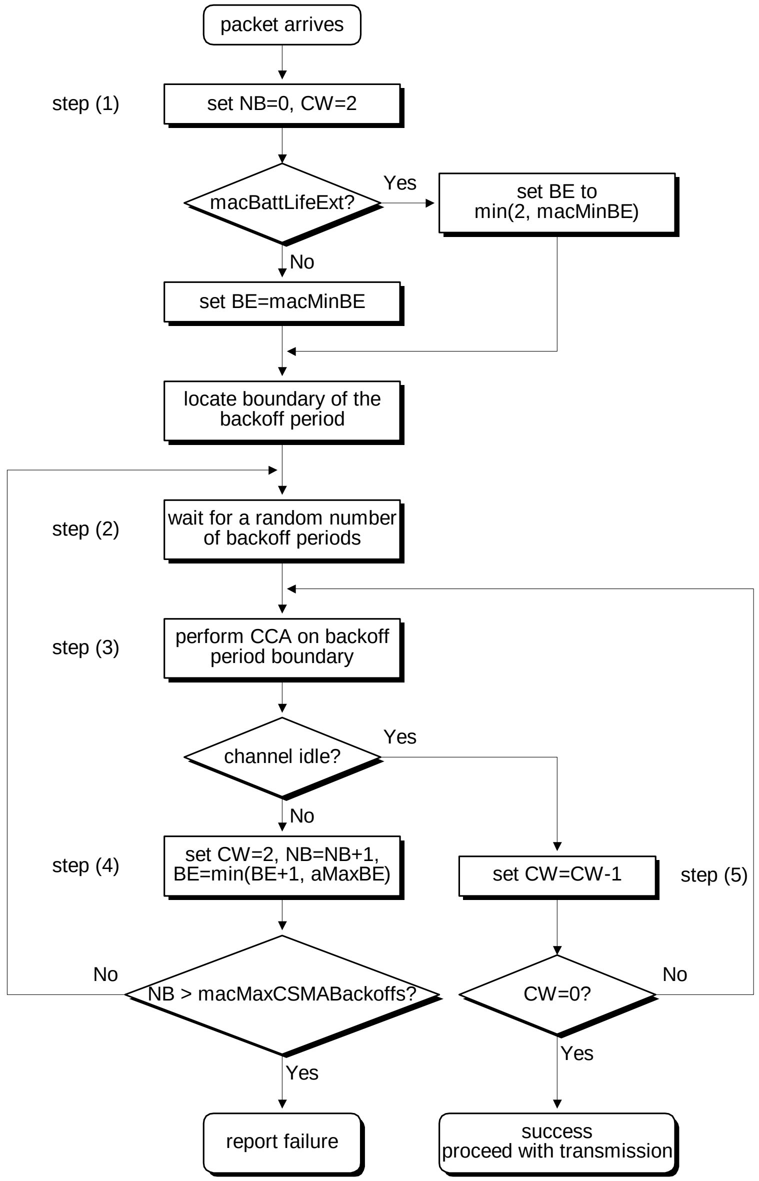
packet arrives
step (1)
set NB=0, CW=2
macBattLifeExt?
Yes
set BE to min(2, macMinBE)
No
set BE=macMinBE
locate boundary of the backoff period
step (2)
wait for a random number of backoff periods
step (3)
perform CCA on backoff period boundary
channel idle?
Yes
No
step (4)
set CW=2, NB=NB+1, BE=min(BE+1, aMaxBE)
set CW=CW-1
step (5)
No
NB > macMaxCSMABackoffs?
CW=0?
No
Yes
Yes
report failure
success proceed with transmission

Figure 1: Operation of the slotted CSMA-CA MAC algorithm in the beacon enabled mode (adapted from [7]).

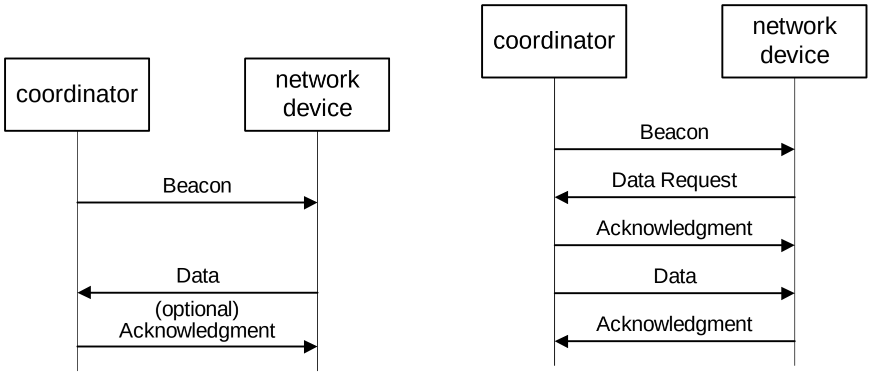

(a) Uplink transmission.

(b) Downlink transmission.

Figure 2: Uplink and downlink data transfers in beacon enabled PAN.

3.2 On uplink and downlink communication

According to the 802.15.4 standard, uplink data transfers from a node to the coordinator are synchronized with the beacon, in the sense that both the original transmission and the subsequent acknowledgment must occur within the active portion of the same superframe, as shown in Fig. 2(a). Uplink transmissions always use the CSMA-CA mechanism outlined above.

Data transfers in the downlink direction, from the coordinator to a node, are more complex, as they must first be announced by the coordinator. In this case, the beacon frame will contain the list of nodes that have pending downlink packets, as shown in Fig. 2(b). When the node learns there is a data packet to be received, it transmits a MAC command requesting the data. The coordinator acknowledges the successful reception of the request by transmitting an acknowledgement. After receiving the acknowledgement, the node listens for the actual data packet for the period of *aMaxFrameResponseTime*, during which the coordinator must send the data frame.

According to the standard, it is allowed to send the data frame 'piggybacked' after the request acknowledgment packet, i.e., without using CSMA-CA. However, two conditions have to be fulfilled: the coordinator must be able to commence the transmission of the data packet between *aTurnaroundTime* and *aTurnaroundTime* + *aUnitBackoffPeriod*, and there must be sufficient time in the CAP for the message, appropriate inter-frame spacing, and acknowledgement; if either of these is not possible, the data frame must be sent using the CSMA-CA mechanism [7]. While the first condition depends on the implementation platform, the second depends on the actual traffic, and *some* data frames may have to be sent using CSMA-CA.

Note that acknowledgments are sent only if explicitly requested by the transmitter; the acknowledgment must be received within a prescribed time interval, otherwise the entire procedure (starting from the announcement in the beacon frame) has to be repeated.

3.3 Security services and suites

The 802.15.4 standard specifies several security suites which consist of a 'set of operations to perform on MAC frames that provide security services' [7]. Specified security services include the following:

- Any device can maintain an Access Control List (ACL) – a list of trusted devices from which the device wishes to receive data; this mechanism is intended to filter out unauthorized communications.

- Data Encryption service helps a device encrypt a MAC frame payload using the key shared between two peers, or among a group of peers. If the key is to be shared between two peers, it is stored with each entry in the ACL list; otherwise, the key is stored as the default key. (The MAC layer provides the symmetric encryption security systems using application-provided key, or keys.) Thus, the device can make sure that the data cannot be read by devices that do not possess the corresponding key. However, device addresses are always transmitted in the clear (i.e., unencrypted), which makes attacks that rely on device identity somewhat easier to launch.

- Frame Integrity service ensures that a frame cannot be modified by a receiver device that does not share a key with the sender, by appending a message integrity code (MIC) generated from blocks of encrypted message text.

- Finally, Sequential Freshness uses the frame counter and key sequence counter to ensure the freshness of the incoming frame and guard against replay attacks.

The services listed above are typically implemented in hardware for performance reasons, and their use is optional. A device can choose to operate un-secured mode, secured mode, and ACL mode. In unsecure mode, none of the services mentioned above are available. In secured mode, the device may use one of the security suits supported by the standard [7], all of which use the Data Encryption service explained above.

A device operating in ACL mode can maintain a list of trusted devices from which it expects to receive packets, but the only security service available is access control service which enables the receiver to filter received frames according to the source address listed in the frame. However, since no encryption of is used, it is not possible to authenticate the true source of the data packet, or to ascertain that the packet payload has not been modified in any way.

We note that the procedures for key management, device authentication, and freshness protection are not specified by the standard; they are left to be implemented by the applications running on 802.15.4 devices.

4 MAC Layer Attacks

Attacks can be broadly classified in two categories, depending on whether the attacker follows the rules of the 802.15.4 MAC protocol, either fully or only to a certain extent, or not. While the attacks from the latter category are potentially more dangerous, defence against them is much more difficult, as might be expected; in this case, the attacker can use a separate 802.15.4-compliant device, possibly modified to loosed the adherence to the MAC protocol. Alternatively, an existing 802.15.4 device may be captured and subverted so as to be used for malicious purposes.

We note that an adversary with appropriate resources might develop and use dedicated hardware which is compatible but not compliant with the 802.15.4 standard. The discussion of such attacks is beyond the scope of this chapter.

4.1 Attacks that follow the MAC protocol

It may come as a surprise that a number of attacks may be conducted by an adversary which follows the IEEE 802.15.4 slotted CSMA-CA protocol to the letter. All of those attacks may be conducted by an adversary which acts as a legitimate member of the PAN.

A simple but not very efficient attack against network availability is to flood the network by simply transmitting a large number of packets. Packets should be large in size, perhaps the largest size allowed by the standard. In this manner, an adversary may degrade the network performance and drastically reduce throughput; our previous work indicates that the performance of an 802.15.4 network can be seriously affected by high packet arrival rates or by nodes operating in saturation regime [10, 9].

An adversary may target different destination devices with unnecessary packets, possibly in other PANs, regardless of whether the destination PAN and/or device actually exist or not. If the goal of the attack is the depletion of the power source for a specific node (and the PAN coordinator), all injected packets may target that node. Since the downlink packets have to be explicitly requested from the PAN coordinator, this will keep the both the PAN coordinator and the chosen destination device busy and eventually exhaust their respective power sources.

A malicious node can simply pretend to run in battery life extension mode, by setting the *aMacBattLifeExt* variable to true. In that case, the CSMA-CA algorithm will choose the initial value of the backoff exponent as 2 instead of 3, as explained in Section 3.1, and the random number foo the backoff countdown will be in the range $0 \mathinner{\ldotp\ldotp} 3$ rather than in the range $0 \mathinner{\ldotp\ldotp} 7$. Shorter backoff countdown means that the probability to access the medium is much higher than for a regular node. On top of that, a regular node would have to wait for the malicious one to finish its transmission, and waste power in the process.

Note also that the node that succeeds in getting access to the medium will not increase its backoff exponent for the next transmission, while the unsuccessful one will increase it by one. Therefore, if the first attempt succeeded, the second one is even more likely to do so, which again clearly favors malicious nodes.

It should be noted that power consumption during packet reception is about one-half to two-thirds of the corresponding power consumption required for packet transmission [5]. Therefore, while transmitting does consume lots of energy (relatively speaking), receiving is not terribly efficient either, and the best way to conserve power is to turn off the radio subsystem whenever possible.

4.2 Attacks that use a modified MAC protocol

The attacks mentioned above can be conducted using a completely functional 802.15.4 device that follows the protocol to the letter – it suffices for the malicious node to control the application that executes on the sensor device. However, a number of additional attacks may be launched by simply modifying or disregarding certain features of the protocol. This can be accomplished either through dedicated hardware or by controlling an otherwise fully compliant 802.15.4 hardware device, as follows.

We have already mentioned the possibility to decrease the backoff exponent and shorten the random backoff countdown, by falsely reporting that battery life extension is enabled. Similar effects may be achieved by *not* incrementing the backoff exponent after an unsuccessful transmission attempt.

The random number generator can be modified to give preference to shorter backoff countdowns. Again, this allows the malicious node to capture the channel in a disproportionately high number of cases, and gives it an unfair advantage over regular nodes.

The number of required CCA attempts can be reduced to one instead of two, which would give the malicious node an unfair advantage over the regular nodes.

The CCA check can be omitted altogether, in which case the node will start transmitting immediately after finishing the random backoff countdown. Even worse, the node can omit the random backoff countdown itself. In this manner, the malicious node can transmit its packets more often than a regular one. While not all of the messages will be sent successfully—there will be collisions in many cases—the malicious node probably doesn't even care, as long as the transmissions from regular nodes end up garbled and thus have to be repeated. Moreover, some of the attacker's transmissions may collide with acknowledgments. Again, this wastes power of the devices affected, but also the bandwidth of the entire network.

In case the acknowledgment is requested by the data frame or the beacon frame, a malicious node may simply refuse to send it. The PAN coordinator will retry transmission (up to a maximum of *aMaxFrameRetries*) and thus waste power and bandwidth.

5 Impact of attacks that follow the MAC protocol

In order to assess the impact of attacks, we have built the simulator of an IEEE 802.15.4-compliant sensor cluster at the MAC level, using the object oriented Petri Net-based simulation engine Artifex by RSoft Design, Inc. [13]. The cluster operates in the ISM band at 2.4GHz, in a star configuration with the PAN coordinator acting as the network sink, a total of twenty regular nodes, and one or two attacker nodes. Slotted, beacon-enabled communication with CSMA-CA mechanism described in Section 3.1 is used for all communications, for reasons outlined in Section 3.2. Regular nodes generate Poisson distributed traffic with packets of three backoff periods, which corresponds to the payload of thirty bytes, with the average arrival rate of 120 packets per minute.

The following attack scenarios from Sections 4.1 and 4.2 were considered:

- The attacker node follows the MAC protocol to the letter but generates spurious traffic with varying packet size and arrival rate. This is the simplest form of attack: flooding the cluster with useless packets.
- The attacker node tries to gain unfair advantage by operating with the *aMacBattLifeExt* variable set; this reduces the random backoff exponent and reduces the range for random backoff countdown. Packet length and arrival rate are variable.
- The attacker node uses one CCA check instead of two required by the standard; it also has the *aMacBattLifeExt* variable set and uses variable packet length and arrival rate.

In each scenario, we have measured the probability that the regular node will succeed in transmitting its packets, as well as the mean packet delay for those packets. In this manner, we can obtain a quantitative measure of the disruption caused by the attacker node, or nodes. Note that the traffic and network parameter values are chosen so that the cluster operates in non-saturation mode.

Fig. 3 shows the success probability for the regular nodes in the three scenarios outlined above, in clusters with one and two attacker nodes, respectively. As can be seen, the success probability rapidly decreases with the increase in length and/or injection rate of the packets generated by the attacker node(s). As can be expected, two attacker nodes generate more disruption than one. This is particularly pronounced in the second and third scenarios described above. In the worst case, Figs. 3(b) and 3(d), the success probability for regular transmissions drops to about 0.5, as opposed to the 'normal' value of about 0.95. In other words, the energy expenditure per packet will be twice the value needed in normal operation, which means that the cluster lifetime—if operating on battery power—will be cut in half. Note that those values were obtained in a cluster with twenty regular nodes which is attacked by two nodes only – and those nodes still follow the MAC protocol to the letter!

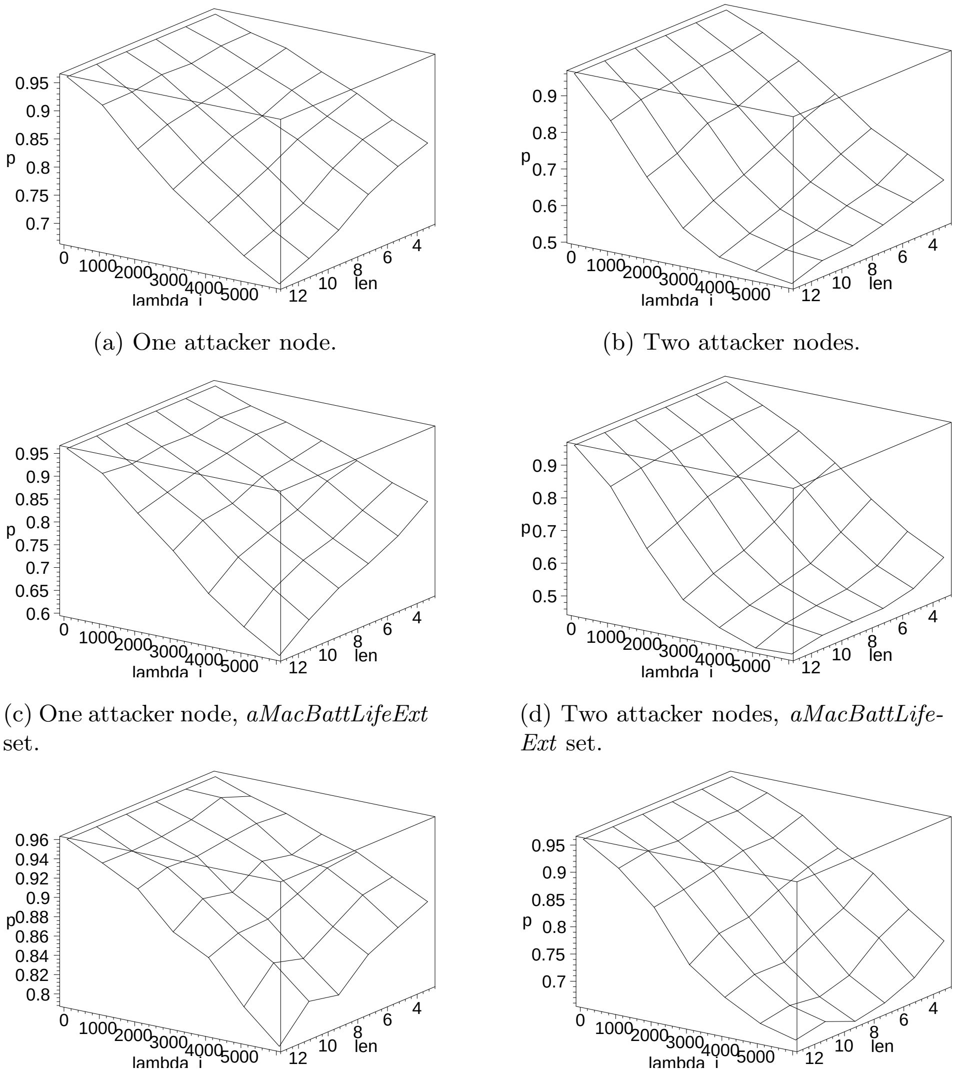

(a) One attacker node.

(b) Two attacker nodes.

(c) One attacker node, *aMacBattLifeExt* set.

(d) Two attacker nodes, *aMacBattLifeExt* set.

(e) One attacker node, *aMacBattLifeExt* variable set and a single CCA.

(f) Two attacker nodes, *aMacBattLifeExt* variable set and a single CCA.

Figure 3: Attack impact: success probability for regular nodes.

Somewhat unexpectedly, the use of one CCA check instead of two does not decrease the success probability for regular nodes – actually, it improves it instead. Namely, collisions happen when two or more nodes finish their random backoff countdown at the same time, check that the medium is idle during two CCAs, and then start transmissions simultaneously. By skipping the second CCA, the attacker node is able to gain access to the medium faster than a regular node. In such cases, however, the regular node will see the medium busy in *its* second CCA and it will restart the CSMA-CA algorithm with the backoff exponent increased by one. Consequently, the probability of collision between transmissions from a regular node and an attacker node is reduced. The decrease in collision probability means that the success probability increases, as witnessed by measurements in Figs. 3(e) and 3(f). Note that collisions still occur when the attacker node finishes its backoff countdown in the backoff period in which a regular node performs its first CCA; and the probability of collisions between two regular nodes is not affected at all.

The argument described above is illustrated through the diagrams of success probabilities for the first and second CCA for regular nodes shown in Fig. 4. For brevity, we show only the case with two attacker nodes. As can be seen from Figs. 4(a), 4(c), and 4(e), the probability that the first CCA is successful does deteriorate with increased packet arrival rate and packet length. However, it is not particularly affected by the introduction of reduced backoff exponent or by the attacker nodes skipping one CCA; the differences between successive diagrams is rather small.

Similarly, the success probability for the second CCA does not change much from the case where attackers follow the MAC protocol to the letter, Fig. 4(b), to the one in which they use the reduced backoff exponent, Fig. 4(d). In fact, the second CCA even shows some recovery at high packet arrival rates and large packet length; note, however, that this is the *second* CCA and only the transmissions that have successfully passed the first one get to this point.

But when the attacker nodes skip the second CCA themselves, the situation changes. As can be seen from Fig. 4(f), the probability that a regular node will succeed in its second CCA is radically reduced, compared to the diagrams above. (Note that success in both CCAs does not guarantee a successful transmission, as collisions may still occur.) Therefore, a large fraction of transmissions from regular nodes will fail the second CCA and will have to wait for at least another pass of the CSMA-CA algorithm. The net result is that the regular packets will experience much larger transmission delays.

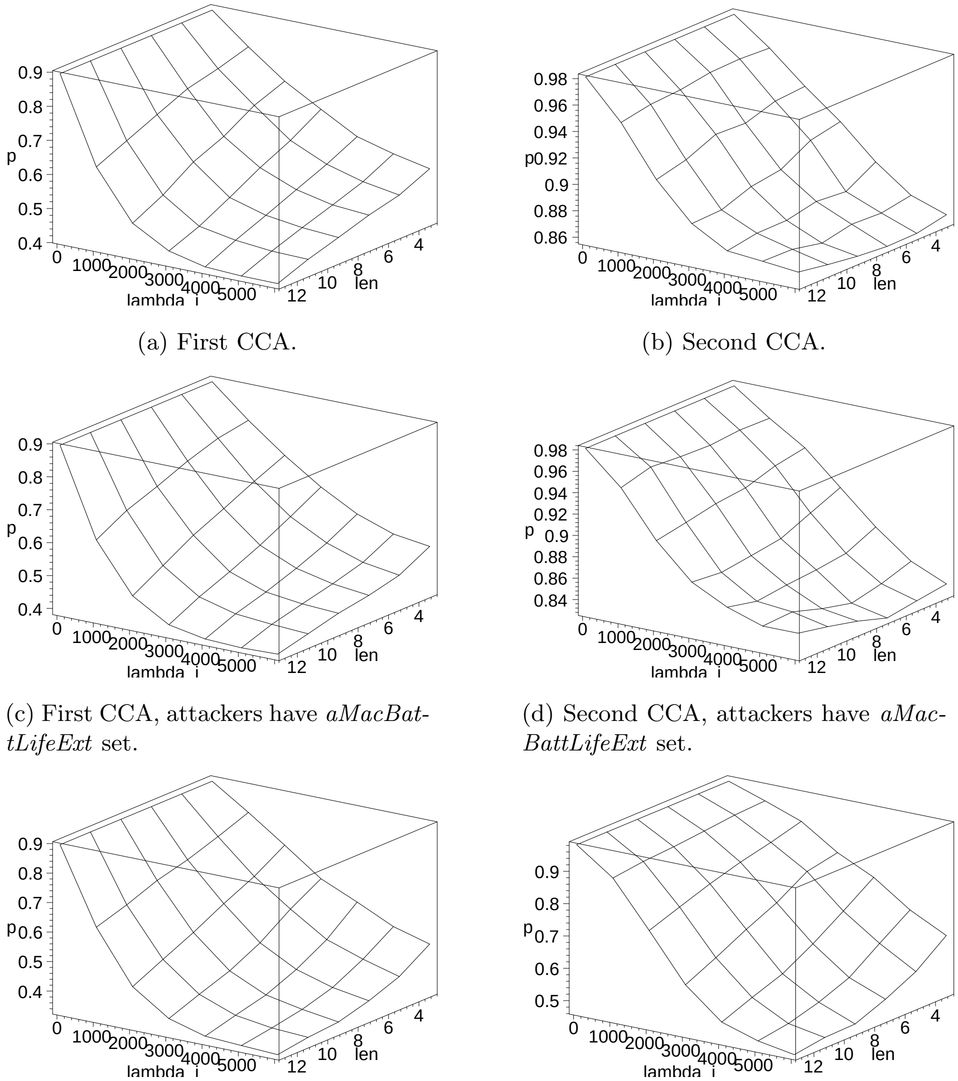

(a) First CCA.

(b) Second CCA.

(c) First CCA, attackers have *aMacBattLifeExt* set.

(d) Second CCA, attackers have *aMacBattLifeExt* set.

(e) First CCA, attackers have *aMacBattLifeExt* set and use a single CCA.

(f) Second CCA, attackers have *aMacBattLifeExt* set and use a single CCA.

Figure 4: Attack impact: success probabilities at first and second CCA for regular nodes (two attacker nodes).

This behavior is obvious from the diagrams in Fig. 5, which show mean end-to-end delays for the packets sent by regular nodes in the three scenarios outlined above. All delays are expressed in units of backoff periods, which last for $0.32ms$ if the network operates in the 2.4GHz ISM band. We consider sensor clusters with one and two attacker nodes, respectively. Delays generally increase when the length and/or injection rate of the packets generate by the attacker node(s) increase. When two attacker nodes are present, delays are about 70 to 80 percent larger than in the case with only one attacker node, as can be expected. Furthermore, the impact of different attack mechanisms on packet delay is cumulative – i.e., each attack mechanism makes the delays longer.

We end by noting that some applications may actually tolerate the increased delays, provided the increase in success probability suffices; in other cases, delay is a critical factor and increased delays can't be tolerated.

6 Defending the 802.15.4 PAN

While the attacks listed in the last Subsection may indeed pose formidable risks to normal operation of an 802.15.4 WPAN, it should be noted that they are probably not very cost-effective to launch. Since individual 802.15.4 sensor nodes are small, low power, low cost devices, the development of dedicated compatible-but-not-compliant devices with modified behavior is likely to be prohibitively expensive – the potential attacker would probably find the use of simple devices for jamming at the PHY layer to be much more attractive. Let us consider just the 2450MHz (the so-called ISM band), with the raw data rate of 250kbps, which is already used by other wireless LAN/PAN standards such as 802.11b and 802.15.1 (Bluetooth), and high interference may be expected. From the specifications of the 802.15.4 standard [7], the processing gain is only around 8, and the Bit Error Rate is given with $BER = Q\left(\sqrt{\frac{E_b}{N_0}}\right)$ [3], where $Q(u) \approx e^{-u^2/2}(\sqrt{2\pi}u)$, $u \gg 1$. Therefore, in an interference-free environment, we should expect the BER values slightly below 10^{-4}. As the packet error rate is $PER = 1 - (1 - BER)^X$, where X is the total packet length (expressed in bits) including MAC and physical layer headers, the probability that a given packet (data or acknowledgment) is much higher: Section 6.1.6 of the standard states that PER of 1% is expected on packets with length of 20 octets [7]. The use of shorter packets will increase the resiliency of the transmission.

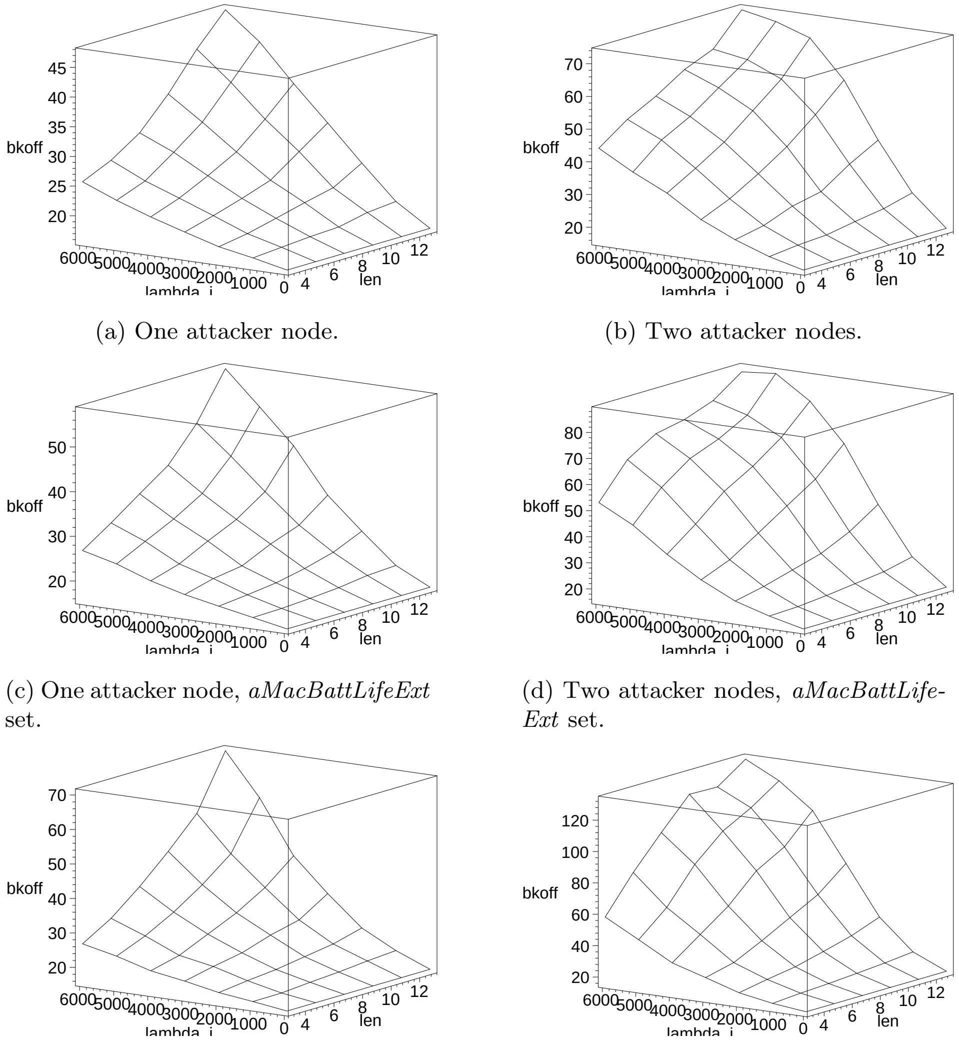

(a) One attacker node.

(b) Two attacker nodes.

(c) One attacker node, *aMacBattLifeExt* set.

(d) Two attacker nodes, *aMacBattLifeExt* set.

(e) One attacker node, *aMacBattLifeExt* variable set and a single CCA.

(f) Two attacker nodes, *aMacBattLifeExt* variable set and a single CCA.

Figure 5: Attack impact: mean packet delay for regular nodes.

However, in the presence of interference and noise in the ISM band, higher BER values, and much higher PER values, may be expected; for example, at BER around 10^{-3}, the PER for a 20-octet packet increases to around 15%! Consequently, an attacker that chooses to jam the transmissions in an 802.15.4 sensor network can cause quite a lot of damage with modest energy expenditure.

It is worth noting that the provisions of Section 6.7.9 of the standard allow 802.15.4 devices to perform CCA checks in one of three different modes: by energy only (mode 1), by carrier sense only (mode 2), or by energy and carrier sense (mode 3) [7]. A legitimate node which uses mode 1 may experience CCA failures in a high interference environment, without an attacker specifically using 802.15.4 modulated signal. Obviously, CCA mode 3 should be used for highest resilience.

In terms of encryption, the obvious weakness of the standard is that the encryption is applied to packet/frame payload only, but not to the information in packet headers; this holds for ordinary data packets but also for data request and beacon frames. While the address list field—the list of addresses of devices that have pending downlink packets—is considered as part of the beacon frame payload according to the standard, it is not encrypted even when the security is enabled for the beacon frame. It would not be too big a change of the standard to encrypt not just the packet payload but also the destination address and the address list fields. This simple measure would make some of the attacks much more difficult to accomplish.

Replay attacks can be identified by the PAN coordinator – provided the Sequential Freshness service is used. While this would not prevent a malicious node from sending such packets, at least the coordinator could filter them out and avoid any further processing. However, this feature may not be very efficient in terms of memory required. Namely, each device has to maintain a table for storing the monotonically increasing counter values associated with message streams sent by each entry in the ACL. If the Sequential Freshness service is enabled in conjunction with the Access Control List service, there is a risk that the available memory of the PAN coordinator will be exhausted.

While the standard does not prescribe, or indeed even recommend, any particular key management scheme, the overall effectiveness of the security services provided by the standard depend very much on the choice of a suitable scheme [15].

Intrusion detection techniques could help identify malicious nodes that might be trying to disrupt the normal operation of the PAN. By analyzing the traffic patterns, the PAN coordinator may become aware of the activity of such nodes, so that appropriate measures can be taken to minimize the disruption. Suspicious activities include amounts of traffic well above the average, traffic intensity that increases over time, possibly in an abrupt fashion, and sending packets to many destinations, possibly in different PANs. In more critical applications, devices with substantially higher computational capabilities (and operating on mains power, rather than battery power) could analyze the activities of individual nodes at the MAC level and identify potential intruder(s).

We note that sensing applications often involve operation with very low duty cycle of individual nodes; this makes intrusion detection comparatively easier to accomplish, and does not help the attackers which might be eager to achieve their objectives.

The standard does not provide for periodical checking of presence and/or integrity of individual devices. However, a sensing application might establish such checks on its own. Simple time-out counters, one per each associated device, would enable the coordinator to check for their continued presence; in addition, a simple challenge/response scheme could allow the coordinator to verify their integrity as well.

Of course, it is impossible to physically isolate an unwanted device in a wireless network; but at least the application should be made aware of the presence of such devices so that their impact on the normal operation of the network could be minimized.

7 Conclusions

We have investigated the security of IEEE 802.15.4 sensor networks at the MAC level. A number of vulnerabilities have been identified, and some simple attack mechanisms that exploit those vulnerabilities have been described and analyzed. More work is needed to address the possible defences against those and other attack mechanisms.

References

[1] J. Bellardo and S. Savage. 802.11 denial-of-service attacks: Real vulnerabilities and practical solutions. In *Proc. 12th USENIX Security Symposium*, Washington, DC, August 2003.

[2] H. Chan and A. Perrig. Security and privacy in sensor networks. *IEEE Computer*, 36(10):103–105, October 2003.

[3] V. K. Garg, K. Smolik, and J. E. Wilkes. *Applications of CDMA in Wireless/Personal Communications*. Prentice Hall, Upper Saddle River, NJ, 1998.

[4] V. Gupta, S. Krishnamurthy, and M. Faloutsos. Denial of service attacks at the MAC layer in wireless ad hoc networks. In *Proc. MILCOM 2002*, Anaheim, CA, October 2002.

[5] J. A. Gutiérrez, E. H. Callaway, Jr., and R. L. Barrett, Jr. *Low-Rate Wireless Personal Area Networks*. IEEE Press, New York, NY, 2004.

[6] Y.-C. Hu and A. Perrig. A survey of secure wireless ad hoc routing. *IEEE Security & Privacy Magazine*, 2(3):28–39, May-June 2004.

[7] Standard for part 15.4: Wireless MAC and PHY specifications for low rate WPAN. IEEE Std 802.15.4, IEEE, New York, NY, October 2003.

[8] C. Karlof, N. Sastry, and D. Wagner. TinySec: A link layer security architecture for wireless sensor networks. In *Proc. SenSys 2004*, pages 162–175, Baltimore, MD, November 2004.

[9] J. Mišić and V. B. Mišić. Access delay and throughput for uplink transmissions in IEEE 802.15.4 pan. *Elsevier Computer Communications*, to appear, 2005.

[10] J. Mišić, S. Shafi, and V. B. Mišić. Analysis of 802.15.4 beacon enabled PAN in saturation mode. In *Proc. SPECTS 2004*, San Jose, CA, July 2004.

[11] J. Newsome, E. Shi, D. Song, and A. Perrig. The Sybil attack in sensor networks: Analysis and defenses. In *Proc. IPSN 2004*, pages 259–268, Berkeley, CA, April 2004.

[12] Q. Ren and Q. Liang. Secure media access control (MAC) in wireless sensor networks: Intrusion detections and countermeasures. In *Proc. PIMRC 2004*, volume 4, pages 3025–3029, Barcelona, Spain, September 2004.

[13] RSoft Design, Inc. *Artifex v.4.4.2*. San Jose, CA, 2003.

[14] E. Shi and A. Perrig. Designing secure sensor network. *IEEE Wireless Communications*, 11(6):38–43, December 2004.

[15] W. Stallings. *Cryptography and Network Security – Principles and Practice*. Prentice Hall, Upper Saddle River, NJ, 3rd edition, 2003.

[16] A. D. Wood and J. A. Stankovic. Denial of service in sensor networks. *IEEE Computer*, 35(10):54–62, October 2002.

[17] Y. Zhou, D. Wu, and S. M. Nettles. Analyzing and preventing MAC-layer denial of service attacks for stock 802.11 systems. In *Proc. BroadWISE 2004*, pages 22–88, San Jose, CA, October 2004.