Proceedings of the 2006 IEEE
Workshop on Information Assurance
United States Military Academy, West Point, NY

Effects of Denial of Sleep Attacks on Wireless Sensor Network MAC Protocols

David Raymond, Randy Marchany, Michael Brownfield, and Scott Midkiff

Abstract – **As wireless platforms get less expensive and more powerful, the promise of wide-spread use for everything from health monitoring to military sensing continues to increase. Like other networks, sensor networks are vulnerable to malicious attack, however, the hardware simplicity of these devices makes defense mechanisms designed for traditional networks infeasible. This paper explores the *denial-of-sleep* attack, in which a sensor node's power supply is targeted. Attacks of this type can reduce sensor lifetime from years to days and have a devastating impact on a sensor network. This paper classifies sensor network denial-of-sleep attacks in terms of an attacker's knowledge of the MAC layer protocol and ability to bypass authentication and encryption protocols. Attacks from each classification are then modeled to show the impacts on three sensor network MAC protocols: S-MAC, T-MAC, and G-MAC. A framework for preventing denial-of-sleep attacks in sensor networks is also introduced. With full protocol knowledge and an ability to penetrate link-layer encryption, all wireless sensor network MAC protocols are susceptible to a full domination attack which reduces network lifetime to the minimum possible by maximizing the power consumption of the nodes' radio subsystem. Even without the ability to penetrate encryption, subtle attacks can be launched that reduce network lifetime by orders of magnitude. If sensor networks are to live up to current expectations, they must be robust in the face of network attacks, to include denial-of-sleep.**

Index Terms–**Wireless Security, Wireless Sensor Networks, Medium Access Control (MAC)**

I. Introduction

Wireless sensor networks (WSN) are becoming increasingly attractive for a variety of application areas to include industrial automation, security, weather analysis, and a broad range of military scenarios. The challenge of designing these systems to be robust in the face of myriad security threats is a priority issue. One such threat is the *denial-of-sleep* attack, a specific type of denial-of-service (DoS) attack that targets a battery-powered device's power supply in an effort to exhaust this constrained resource. If a large percentage of network nodes, or a few critical nodes, are attacked in this way, the network lifetime can be reduced from months or years to days.

The impacts of denial-of-sleep attacks on WSN MAC protocols have not been thoroughly addressed. This paper introduces a system for classifying denial-of-sleep attacks in WSNs. The impact of various denial-of-sleep attacks on current wireless sensor devices are modeled and a framework for defending against these potentially devastating attacks is presented.

To make nodes small and inexpensive so that they can be economically deployed in large numbers, they generally have very limited processing capability and memory capacity. Because the design of these devices usually favors decreased cost over increased capabilities, we cannot expect Moore's Law to lead to enhanced performance. Another challenge unique to sensor node platforms is their extremely limited and often non-replenishable power supply. For example, if sensors are deployed via aircraft in a military scenario for sensing enemy movements, replacing or recharging batteries is not feasible. Two examples of widely available sensor node platforms are the MICAz™ [1] and the TMote™ Sky [2]. Both devices are configured to run for a year or more on a pair of AA batteries, relying on long periods of sleep in order to save power. The dominant source of power loss in these sensor platforms is the radio subsystem. Table I shows power consumption during receive, transmit, and sleep periods for these devices. The data-link layer, specifically the medium access control (MAC) protocol, is responsible for managing the radio. Therefore the MAC protocol must keep the radio in a low-power sleep mode as much as possible. As a result, most research in the area of sensor node power conservation is focused on MAC protocols.

The rest of the paper is organized as follows: Section II explores sources of energy loss in sensor networks and briefly describes three leading WSN MAC protocols. Section III discusses related work in the area of sensor network security. Section IV outlines a framework for classifying denial-of-sleep attacks in these networks and Section V explores the impact of a selection of denial-of-sleep attacks against the MAC protocols presented in Section II. Finally, Section VI provides a framework for defending against denial-of-sleep attacks in sensor networks and Section VII concludes.

TABLE I
SENSOR PLATFORM POWER CONSUMPTION AND SLEEP TRANSITION DATA

		MICAz™ [1]	TMote™ Sky [2]
Power Draw	Receive (mW)	65.91	64.68
	Transmit (mW)	59.10	55.20
	Sleep (mW)	0.570	0.114
Avg Sleep Transition Current (mW)		9.60	5.64
Sleep Transition Time (ms)		5.87	6.81
Data rate (kbps)		250	250

D. Raymond and M. Brownfield are Ph.D. students, R. Marchany is the Director of the IT Security Lab, and Scott Midkiff is a Professor of Electrical and Computer Engineering at Virginia Tech. (Emails: raymondd, michael.brownfield, marchany, midkiff@vt.edu)

II. Sources of Energy Loss in Sensor Nodes

MAC layer protocols designed for WSNs use various algorithms to save battery power by placing the radio in low-power modes when not actively sending or receiving data. Table I illustrates the importance of maximizing a node's sleep ratio because the transmit and receive power can be up to three orders of magnitude greater than the sleep power. Let the sleep ratio, or R_{sleep}, equal $T_{sleep}/(T_{active}+T_{sleep})$, where T_{active} and T_{sleep} are active time and sleep time. A node's lifetime is:

$$T_{sensorlife} = \frac{C_{battery(mWh)}}{(R_{sleep})(P_{sleep(mW)})+(1-R_{sleep})(P_{active(mW)})}, \quad (1)$$

where P_{active} and P_{sleep} are active mode power draw and sleep mode power draw and $C_{battery}$ is the total amount of available energy. P_{active} is almost 3 orders of magnitude greater than P_{sleep}, so it is important to keep nodes in sleep mode as much as possible. The TMote™ Sky consumes 64.68 mW in receive mode and 0.114 mW in sleep mode [2]. Using two standard 3,000 mAh AA batteries, it will last 3,300 days in sleep mode, but only 5.8 days in receive mode. The disparity between receive cost and sleep cost leads to an exponential increase in network lifetime as sleep time increases, suggesting that an attack that decreases sleep time by even a few percentage points can have dramatic impact on network lifetime.

A. Sources of Energy Loss

The amount of power that can be saved depends largely on the MAC protocol's ability to overcome the radio's four primary sources of energy loss: *collisions*, *control packet overhead, overhearing*, and *idle listening*.

Collisions. Collision loss refers to energy wasted due to packet collisions on the wireless medium. If a transmission of sufficient signal strength interferes with a data packet being sent, the data will be corrupted at the receiving end. Some collisions can be mitigated using error correcting codes (ECC), however ECCs add transmission overhead, which is contrary to the goal of reducing radio transmit time.

Control Packet Overhead. Examples of control packets are the *request to send* (*RTS*) and *clear to send* (*CTS*) messages used by the IEEE 802.11 protocols. Depending on the MAC protocol used, these control packets may have to be received by all nodes within radio range of the sender, as is the case with RTS/CTS packets, resulting in power drain in a potentially large numbers of nodes. If nodes can be forced to stay awake for spurious control packets, battery life can be greatly impacted.

Overhearing. Overhearing loss refers to energy wasted by a node having its radio in receive mode while a packet is being transmitted to another node. Most WSN MAC protocols reduce overhearing by trying to ensure that a node is only awake when there is traffic destined for it. One way to prevent overhearing is to ignore packets destined for other nodes after hearing an RTS/CTS exchange. After overhearing the RTS and CTS, nodes set a network allocation vector (NAV) interrupt based on the message duration field in the CTS message and go to sleep [3]. Fig. 1 depicts a typical NAV scenario. Opportunities for NAV sleep are significantly reduced on new platforms because the time required to transition to sleep and back is shorter than packet transmit times for even the longest packets. The TMote™ Sky, for example, takes 6.81 ms to transition the radio from receive to sleep mode and back while the time required to send a maximum-sized IEEE 802.15.4-compliant frame of 128 bytes is only 4.09 ms.

Idle Listening. A node's radio consumes the same amount of power simply monitoring the channel as it does when it is receiving data. If a node can be made to listen even when there is no traffic destined for it, power is wasted.

B. WSN MAC Protocols

Section V analyzes the impact of denial-of-sleep attacks against three WSN MAC protocols: Sensor MAC (S-MAC)[4], Timeout MAC (T-MAC) [5], and Gateway MAC (G-MAC) [6].

S-MAC. The S-MAC protocol uses a fixed duty cycle, usually set at 10%, during which traffic is exchanged between nodes. Radios in networks using this protocol will be asleep 90% of the time, thereby producing an almost 10-fold improvement in sensor life. In S-MAC, sensor nodes organize themselves into virtual clusters using periodic broadcast synchronization (SYNC) messages. Upon deployment, a node will listen for a SYNC message. If it does not hear one before timeout, it will broadcast a SYNC message announcing its sleep cycle. Nearby nodes overhear this message and synchronize their schedules to the sending node. SYNC messages are repeated at the beginning of each frame to correct time drift and keep virtual clusters' sleep cycles synchronized. If a node overhears two SYNC messages, it will adapt both duty cycles to maintain network connectivity. Fig. 2 depicts the S-MAC frame architecture.

T-MAC. T-MAC improves on S-MAC by concentrating all traffic at the beginning of the duty period as depicted in Fig. 3, thus trading network latency for power conservation. The arrows in the figure indicate transmitted and received messages. T-MAC uses the same SYNC mechanism to form virtual clusters as S-MAC. The improvement in network

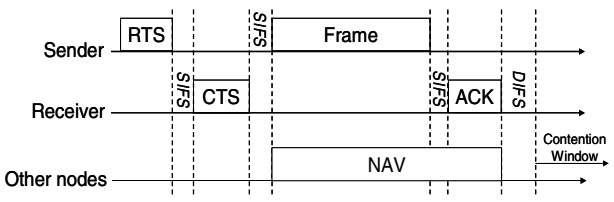

Fig. 1. Typical NAV scenario.

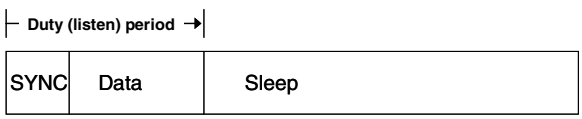

Fig. 2. S-MAC frame structure.

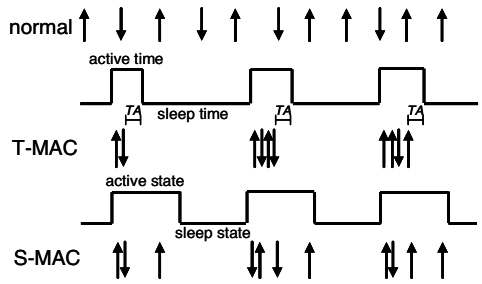

Fig. 3. T-MAC adaptive timeout [5].

lifetime using this protocol is dependent on the amount of traffic in the network since nodes cannot go to sleep until they are assured, through an adaptive timeout (TA), that there will be no more traffic for them. TA is set based on the longest time that a hidden node would have to wait before hearing the beginning of a CTS response message as follows:

$$TA = 1.5 \times \left(T_{Max\_CW} + T_{RTS} + T_{SIFS} \right), \quad (2)$$

where T_{Max_CW} is the largest contention window, T_{RTS} is the time to send an RTS and T_{SIFS} is the short interframe space. In [5], T-MAC is shown to have up to a 5-fold increase in network lifetime over S-MAC.

G-MAC. G-MAC [6] is an energy-efficient MAC protocol designed to coordinate transmissions within a cluster. It divides frames into a collection period and a contention-free distribution period. During the collection period, nodes that have unicast or broadcast traffic to send transmit a future-RTS (FRTS) message to a gateway node. At the end of the contention period the cluster head, or gateway, transmits a gateway traffic indication message (GTIM) which provides a mechanism for cluster synchronization while broadcasting a schedule of message transactions between nodes. Nodes then exchange data during the contention-free period. The gateway is elected using a periodic, resource-adaptive election process in which nodes volunteer based on current resource levels. New elections are indicated by a flag in the GTIM message.

G-MAC eliminates overhearing, except for a minimum amount of control traffic that a node might overhear while waiting to transmit an FRTS during the contention period.

III. Current Research in Sensor Network Security

Most research on sensor network security focuses on integrity and confidentiality. This section first introduces basic WSN security mechanisms and then reviews current research on denial-of-service in sensor networks.

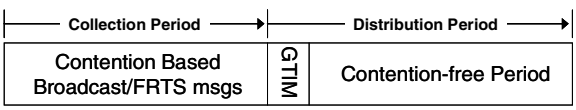

Fig. 4. G-MAC frame structure.

A. General Sensor Network Security Mechanisms

Perrig, et al. [7], introduce two building blocks for security in WSN: *μTesla*, and *Secure Network Encryption Protocol* (*SNEP*). *μ*Tesla is a "micro" version of *Timed, Efficient, Streaming, Loss-tolerant Authentication Protocol* [8] modified for the resource-constrained environment of WSN. It uses delayed disclosure of symmetric keys to overcome the inability to use asymmetric security schemes for authentication due to their high processing overhead and long key lengths. SNEP provides data confidentiality and two-party authentication using symmetric encryption. SNEP also provides a mechanism to support data freshness using monotonically increasing counter values shared by sender and receiver. This anti-replay mechanism requires every node to maintain a table of counter values listing every node from which it receives a packet. Each node must also share a secret key with every communication partner. The memory requirements for storing such information make it unrealistic in memory-constrained sensor nodes even in a moderately-sized network of 25 nodes [9]. This table can also be targeted. If an attacker can overhear traffic in one part of the network and replay in another, it can fool sensors into adding new entries in their neighbor tables and consuming more resources. If the size of the neighbor table is restricted or if neighbor information is fixed on deployment, old entries must be purged for new ones, which defeats the purpose of the table. If tables are fixed and entries cannot be purged, new nodes cannot be added to the network. The packet counter mechanism also requires that nodes receive replayed packets in their entirety before dropping them due to repeated counter values, thus wasting power by overhearing.

TinySec is a link-layer security architecture designed specifically for sensor networks [9]. It provides support for both authenticated encryption and authentication only. Authenticated encryption is achieved with a *message authentication code* computed over an encrypted message while pure authentication is provided using an authentication code computed over the entire packet. TinySec uses *Skipjack* as the default encryption mechanism and *cipher-block chaining* (CBC) as the block cipher mode of operation. Energy consumption overhead in TinySec is relatively low and is caused primarily by longer transmission times due to increased packet length. Power consumption is increased by 10% for authenticated encryption and 3% for authentication only. While TinySec provides a good architecture for data confidentiality and partial integrity support, it does not prevent message replay, nor does it provide specific protection against *resource consumption*, or denial-of-sleep, attacks.

The IEEE 802.15.4 specification [10], also known as *ZigBee*, details architectural requirements for a particular class of wireless radios and protocols for personal area network devices and wireless sensor nodes. The specification provides hardware support for data confidentiality and integrity in compliant devices. It mandates the use of *Advanced Encryption Standard* (AES) encryption and CBC-MAC to provide support for access control, data encryption, and frame

integrity. Support for defense against replay attack, in the form of frame counters, is optional according to the standard.

B. Denial-of-Service in Wireless Sensor Networks

Physical-layer jamming can simultaneously prevent traffic flow on a WSN and rapidly drain sensor batteries. A potential defense against power consumption caused by jamming is to go to a low-power state when such attacks are in progress, waking only periodically to sense the channel [11]. A prerequisite for such a mechanism is for nodes to identify that a jamming attack is ongoing. In [12], jamming attacks are classified as either *constant*, *deceptive*, *random*, or *reactive*. Constant jamming normally involves a constant high-power transmission which requires maximum energy by an attacker and may not be feasible if the attacker is under similar power constraints as the target network. A *deceptive jammer* sends a constant stream of packets into the network to make it appear that the medium is being filled with legitimate traffic. A *random jammer* randomly alternates between sleep and jamming to save power. Finally, a *reactive jammer* only sends a jam signal when it senses traffic to cause collisions. The further the jammer is from the target nodes, the longer the average packet length must be in order for the jammer to sense traffic and send a signal in time to cause a collision.

Techniques for identifying jamming attacks explored in [12] include statistical analysis of received signal strength (RSSI), average time required to sense an idle channel (*carrier sense time)*, and packet delivery ratio (PDR). All of these techniques require that the network not be jammed upon deployment so that baseline measurements can be taken and none of them alone identify all types of jamming. By combining techniques and introducing the notion of a consistency check, however, all four types of jamming can be reliably identified. One such algorithm first identifies poor link utility through PDR analysis, and then uses a statistical RSSI analysis as a consistency check to determine whether the poor network performance is due to jamming. A second technique is to compare PDR values with those expected based on location of neighbor nodes as a consistency check, assuming neighbor locations are known.

The *denial-of-sleep broadcast attack* is presented in [13], where the impact of a malicious host obeying the MAC layer protocol and broadcasting unauthenticated traffic into the network is modeled. Even though the malicious broadcast traffic is dropped due to authentication failure, network lifetime is significantly reduced for networks using the S-MAC and T-MAC protocols. The authors introduce the G-MAC protocol, which weathers this type of malicious broadcast attack particularly well. In G-MAC, requests to broadcast traffic must be authenticated by the gateway node before the traffic can be sent to other nodes, so only the gateway suffers power loss due to unauthenticated broadcast.

In a *replay attack*, network traffic is recorded and replayed. The replayed data will be treated just as it was the first time it was sent over the network and it will be received by a subset of nodes. If there is no replay protection, the traffic will be accepted as legitimate and forwarded to the destination, thus depleting power on each node along the path. Even if replay protection is enabled as described in [7] or [10], the replayed traffic must be received by a node before it is rejected, thus wasting power at that node.

One-way hash chains (OHC) [14] can be used to prevent replay attacks by allowing nodes to detect replayed packets and reject them. In OHC, each node S maintains a set of one-way hash values $<HS_0, HS_1, \ldots, HS_n>$ such that for all i, $HS_i = F(HS_{i-1})$, where $F(x)$ is a one-way hash function [15]. As S sends packets into the network, it attaches the next hash value in its chain, so that packet number i includes HS_i. Each node on the path to the base station maintains a verifier V_S for node S, initially set to HS_0. When a packet reaches an intermediate node, that node computes $F(HS_i)$ and confirms that it matches V_S, otherwise the packet is dropped. The overhead of using OHC includes 8 bytes of data added to each transmitted packet and an OHC table at each node containing 16 bytes of data for each routing path that might include it.

IV. Classifying MAC Layer Denial-of-Sleep Attacks in Wireless Sensor Networks

MAC layer denial-of-sleep attacks on WSNs can be categorized based on the level of protocol knowledge required to initiate them and on the level of network penetration achieved by an attacker. Penetration refers to an attacker's ability to read and send trusted traffic. A network is easily penetrated if the networking protocols are known and if cryptographic mechanisms are not used for communication. While there are mechanisms available for secure communication in WSN, they are not as robust as those found in traditional networks due to resource constraints.

Any shared medium can be attacked with physical layer jamming. Jamming, however, is a blunt instrument for executing a denial-of-sleep attack on a WSN. Depending on the MAC protocol, the lifetime of a WSN can be significant even in the face of jamming, requiring that an attacker jam the network for an extended period to render it ineffective. Network lifetime could be further extended if protocols were modified to put nodes to sleep for long periods when a jamming attack is ongoing. As discussed in Section III, jamming attacks can be reliably identified through analysis of packet delivery ratios and neighbor radio signal strengths (RSSI) [28]. Furthermore, conducting a jamming attack requires considerable resources. A more effective attack strategy is to use knowledge of MAC protocols to initiate an assault aimed at draining power from sensor platforms, thereby rendering the network unusable and nullifying any other security mechanisms. In the ensuing discussion, the following classifications of MAC layer denial-of-sleep attacks are used:

Class 1: No protocol knowledge, no ability to penetrate network. With no knowledge of MAC layer protocols, attacks are limited to physical layer jamming and unintelligent replay attacks. A constant jamming attack normally involves

transmitting a high-power signal on the frequency bands used by the sensor network to prevent any communication between nodes. In a *reactive jamming attack*, the attacker transmits a jam signal only when traffic is detected on a network to cause collisions [12]. Another type of attack that can be mounted with no knowledge of MAC protocols or ability to penetrate network is an *unintelligent replay attack*. Simply replaying all, or randomly selected, traffic causes nodes to waste energy receiving and processing these extra packets. If nodes in the network do not implement an anti-replay mechanism, this attack causes the replayed traffic to be forwarded through the network, consuming power at each node on the path to the destination. Undetected replay has the added benefit (to the attacker) of causing the network to resend data that could subvert the network's purpose. For example, replaying traffic in a military sensor network deployed to sense enemy movement could cause combat units to be misdirected.

Class 2: Full protocol knowledge, no ability to penetrate network. Determining which protocol a sensor network is using may be discernible through traffic analysis. With this knowledge, an attacker could expand his attack types beyond those listed above to include *intelligent jamming*, injecting unauthenticated unicast or broadcast traffic into the network, or being more selective about replaying previous traffic. Intelligent jamming [16] uses knowledge of link-layer protocols to reduce network throughput without relying on a constant jam signal, for example, by jamming only RTS packets. Such attacks improve over constant physical-layer jamming in that they preserve attacker energy, which can be important if attacking nodes have constraints similar to those of the target nodes. Even when attacker power consumption is not a factor, intelligent jamming might be used to make it more difficult for a network to detect an attack.

If valid source and destination addresses are inserted by an attacker, unauthenticated traffic requires that nodes stay awake to receive packets even if they are later discarded due to invalid authentication. If packets are encrypted, a node must receive the entire packet before decrypting and discarding it. The number of nodes impacted by unauthenticated broadcast traffic is dependent upon the MAC protocol. For example, if the protocol uses a cluster-head or gateway node to authenticate broadcast traffic before other nodes are compelled to receive it then only the gateway node energy is impacted. Replay attacks can also be more cleverly executed with knowledge of the protocol, even if the messages cannot be deciphered. It has been shown that if the MAC layer protocol is known, traffic analysis can be used to distinguish data from control traffic [17]. In some cases, even specific control messages can even be identified. Depending on the protocol, effective denial-of-sleep attacks can be mounted by replaying specific control messages even without the ability to decrypt the traffic. For example, properly timed SYNC retransmission in the S-MAC protocol could potentially prevent nodes from entering their duty/sleep cycle and could keep all nodes in receive mode until their batteries were depleted.

Class 3: Full protocol knowledge, network penetrated. Attacks in this category could be devastating to a WSN. With full knowledge of the MAC protocol and the ability to send trusted traffic, an attacker can produce traffic to gain maximum possible impact from denial-of-sleep attack. The types of attacks that could be executed against each MAC protocol and the impact of such attacks are analyzed in Section V.

Table II classifies the denial-of-sleep attacks available based on the attacker's protocol knowledge and ability to penetrate the network. A fourth case, no knowledge of protocol but an ability to penetrate network, is not considered. The ability to penetrate the network assumes full knowledge of the MAC layer protocol used.

V. Impact of Denial-of-Sleep Attacks on Selected MAC Protocols

In this section, attacks from each of the classifications and their impacts on S-MAC, T-MAC, and G-MAC are analyzed. The only study that focuses on denial-of-sleep in WSN is [13], which models network lifetime under routine traffic patterns for a representative set of MAC protocols and models the impact of a *denial-of-sleep broadcast attack* on these protocols. The following paragraphs describe a more potent unauthenticated broadcast attack on each of the MAC protocols in addition to exploring the impacts of constant physical layer jamming, intelligent replay, and a full domination attack for each of the three protocols considered. A full domination attack assumes that the attacker has penetrated the network and has full knowledge of the MAC protocol. In each case, a full domination attack can reduce network lifetime to 6 days, equivalent to network lifetime under IEEE 802.11 with no power saving features. Results of each of the attacks are given in Table III.

A. Network Model

Each network is modeled in MATLAB using similar configurations. All protocol models use a 500 ms frame size. S-MAC uses a fixed 10% duty cycle. The T-MAC adaptive sleep timeout is set to 13.5 ms. The G-MAC frame is broken into a collection period, a GTIM broadcast, and a distribution

TABLE II
CLASSIFICATION OF WSN DENIAL OF SLEEP ATTACKS

Attack	*Class I* *No protocol knowledge, no network penetration*	*Class II* *Full protocol knowledge, no network penetration*	*Class III* *Full protocol knowledge, network penetrated*
Constant jammer	✓	✓	✓
Random or reactive jammer	✓	✓	✓
Intelligent jammer		✓	✓
Untrusted unicast/broadcast		✓	✓
Trusted rogue unicast/bcast			✓
Unintelligent replay	✓	✓	✓
Intelligent replay		✓	✓
Full domination			✓

TABLE III
IMPACT OF DENIAL OF SLEEP ON NETWORK LIFETIME FOR SELECTED MAC PROTOCOLS
(NETWORK LIFETIME GIVEN IN DAYS)

Routine Network Traffic				Attack Traffic (see Table II for classifications)			
				Class I	*Class II*		*Class III*
MAC Protocol	Empty Network	Unicast Traffic	Broadcast Traffic	PHY Layer Jamming	DoS Broadcast	Intelligent Replay	Full Domination
802.11	6	6	6	6	6	6	6
S-MAC	56	56	56	56	56	6	6
T-MAC	194	111	133	6	6	6	6
G-MAC	1024	828	295	237	371	160	6

period. The size of the GTIM is 14 bytes, plus (3 bytes × number of packets per frame). The size of the G-MAC collection and distribution periods are both 1/2 frame, or 250 ms. The system models 50 TMote™ Sky nodes in a single-hop neighborhood and operates at 250 kbps. The network lifetime is based on the TMote™ Sky power consumption for receive (64.68 mW), transmit (55.20 mW) and radio power-down sleep (0.114 mW). IEEE 802.11 with no power management provides the baseline case. In regular IEEE 802.11, the radio is always in receive mode unless transmitting and the node therefore draws maximum power from the batteries, resulting in a 6-day network lifetime. For the routine network traffic scenarios in Table III, traffic is modeled as four 64-byte unicast or broadcast packets per second.

B. Denial-of-Sleep Attacks and Impacts

Physical layer jamming attack. The first attack classification in Section IV considers an attacker with no protocol knowledge and no ability to penetrate the network. This classification of attack is modeled using a constant physical-layer jamming attack. Under this attack, S-MAC is unable to transmit data and nodes remain awake during the entire 10% duty cycle because they are not able to enter NAV sleep. T-MAC fares much worse under this type of attack because nodes must sense an idle channel for the period dictated by the network's adaptive sleep timeout (TA) before going to sleep. Under constant physical layer jamming, nodes will never sense an idle channel and will remain in receive mode constantly, resulting in a network lifetime of 6 days. In the GMAC protocol, the gateway node will remain constantly awake because there is no network idle time to allow it to go to sleep and will therefore last for 6 days. Other nodes will wake up and timeout once during each frame listening for a GTIM. Waking up for these small GTIM messages results in 0.16% duty cycle and a lifetime of 1287 days (or battery shelf-life) for all of the other nodes in the network. A more effective attack against G-MAC would be to periodically lift the jamming attack so that a new gateway is elected, thereby distributing the maximum power draw among all nodes. This would cause the average node per frame power consumption to be

$$P_{NodeAverage} = \frac{\left(P_{Gateway}\right) + \left(n-1\right)\left(P_{OtherNodes}\right)}{n}, \quad (3)$$

where n is the number of nodes, $P_{Gateway}$ is the gateway power consumption while always awake and $P_{OtherNodes}$ is the power consumption of the rest of the nodes. Under this attack, G-MAC network lifetime is reduced to the 237 days.

DoS Unauthenticated Broadcast Attack. The second attack classification considers an attacker with full protocol knowledge but no ability to penetrate the network. In this case, the attacker broadcasts traffic into the network following all MAC protocol rules for timing and collision avoidance. Under S-MAC and T-MAC, these messages are received by all nodes, but discarded because they cannot be authenticated. Even though the broadcast messages are not authenticated, the fact that all nodes stay awake to receive the messages has significant impact on network lifetime. Sensor nodes using the S-MAC protocol are unable to save power using NAV sleep, keeping them in receive mode during their entire 10% duty cycle and resulting in a network lifetime of 56 days. To minimize network lifetime for networks running the T-MAC protocol, short broadcast messages are sent at a period just short of the adaptive timeout (TA) to prevent nodes from going to sleep. This attack will keep the sensor nodes awake during the entire frame and reduce lifetime to 6 days while keeping the attacker's power requirements to a minimum. Under G-MAC, only the gateway receives the broadcast FRTS during the collection period. Since it cannot be authenticated, the broadcast message is not scheduled during the distribution period. To maximize the impact of this attack on G-MAC, the gateway should be kept awake during the entire collection period. The G-MAC gateway uses the same adaptive timeout mechanism as T-MAC to go to sleep during the contention period if there is no more traffic for it. An attacker should therefore send short broadcast messages at a rate just short of the adaptive timeout period to prevent the gateway from transitioning to sleep mode. Assuming no other traffic in the network, the other nodes would only wake up to receive an empty GTIM, and then sleep for the remainder of the time, resulting in an overall network lifetime of 371 days. Any legitimate network traffic in addition to the unauthenticated broadcast packets further reduces this lifetime.

Intelligent Replay Attack. Another attack in the category of full protocol knowledge but no network penetration is an intelligent replay attack. If an attacker can distinguish control traffic from data traffic under S-MAC and T-MAC, SYNC packets can be replayed at an interval short of the sensor

cluster's duty cycle, effectively restarting the duty cycle and pushing back the sleep period each time. This would keep all nodes awake until they run out of power. In G-MAC, FRTS messages should be replayed such that the corresponding NAV periods fill the contention-free portion of each frame. The number of FRTS messages required to do this is dependent on the size of the corresponding messages. For a message size of 64 bytes, 75 FRTSs would fill the contention-free period ensuring that at least one node is awake at all times. This effect, combined with a longer GTIM message which all nodes must receive, results in a network lifetime of 160 days, assuming all of the FRTSs are for unicast packets. If any of the replayed FRTS messages happen to be broadcast FRTSs, network lifetime is further degraded because all nodes must wake up during the contention-free period to listen for the broadcasts. If only 10% of the FRTS messages, or 7 FRTS messages per frame, are for broadcasts, the network life is cut by almost 50%, dropping to 83 days. The worst case is if all FRTSs are for broadcast messages. In this case, network lifetime is reduced to 12 days as discussed below. Even if the message size is not known, the attacker could simply attempt to resend recorded FRTS messages until the gateway quits accepting them. The maximum number of FRTSs that an attacker can send can be determined based on the length of the collection period as follows:

$$NumberFRTS = \left(\frac{T_{CollectionPeriod}}{T_{Contention} + T_{DIFS} + T_{FRTS} + T_{SIFS} + T_{ACK}} \right), \quad (4)$$

where $T_{Contention}$ is the average contention period, T_{DIFS} and T_{SIFS} are the 802.11 distributed and short interframe space periods, T_{FRTS} is the time required to send a 13 byte FRTS message, and T_{ACK} is the time required for the gateway to send a 5 byte acknowledgement. With a 250 ms collection period, a maximum of 138 FRTS messages can be sent. With the potential for 138 FRTSs, the attacker can easily maximize traffic during the contention-free period.

Full Domination Attack. The final attack classification is one in which an attacker has full protocol knowledge and has penetrated the network. This type of attack might be mounted using one or more compromised nodes in the network. Once this level of network penetration is achieved, all of the MAC protocols are susceptible to worst-case power consumption. An attack against S-MAC is simply to send a SYNC message at a frequency just short of the duty cycle to keep pushing back the sleep time. T-MAC network lifetime is minimized by continually sending packets at an interval slightly shorter than the adaptive timeout (TA) so that none of the nodes can ever transition to sleep. A full domination attack against G-MAC has the attacker broadcasting a GTIM message before the gateway node by waiting for less than the required PIFS backoff normally required before a GTIM. If the attacker fills the GTIM with broadcast messages that fill the entire frame up to the next GTIM, all nodes will remain in receive mode during the entire frame waiting for the broadcast traffic. By repeating this pattern for each frame, all nodes are kept awake and the network lifetime is reduced to 6 days. A simpler full domination attack against G-MAC would simply have the attacker send broadcast FRTSs to the gateway such that the contention-free period is filled with broadcast messages. With 89 64-byte packets, the 250ms contention-free period would be filled, resulting in a 50% duty cycle for all nodes and a network lifetime of 12 days.

C. Discussion

The analysis of these attacks shows that with knowledge of the MAC protocol, even without the ability to penetrate encryption, attacks can be constructed that have more significant impact on the network than even constant physical layer jamming. These attacks not only reduce network lifetime significantly, but they are subtle enough that the network may not even be able to identify that it is under attack. Furthermore, these attacks can be sustained longer because the attacker can conserve power by not transmitting a constant jam signal.

VI. A Framework for Defending Against Denial-of-Sleep Attacks in WSN

In this section, a framework for defending against denial-of-sleep attacks is presented. To prevent the attacks across the spectrum of the classifications presented in Section IV, a defensive framework must incorporate five key components: strong link-layer authentication, anti-replay protection, jamming identification and mitigation, broadcast attack defense, and resilience to compromised nodes.

Strong Link-layer Authentication. This is the first and most important component of denial-of-sleep defense and must be incorporated into any WSN that might be vulnerable to abuse. Authentication at higher protocol layers can be effective for providing data integrity and confidentiality but still fails to ensure service availability. An attacker's ability to send trusted MAC-layer traffic on the network leaves it open to the types of full-domination attacks that can reduce network lifetime from a year or more to less than a week. Existing options for implementing link-layer authentication in WSN include TinySec, which is incorporated into current releases of TinyOS [18], and the authentication algorithms built in to IEEE 802.15.4-compliant devices.

Anti-replay Protection. An attacker's ability to replay messages, even without being able to read them, can force nodes to forward old traffic through the network and can significantly increase power consumption for all nodes on the path from sender to receiver. Traffic analysis makes it possible to distinguish control traffic from data traffic. Replaying control packets, like RTS messages, prevents nodes from sleeping and results in wasted power. If this period is short enough that other nodes in the network do not have time to transition to sleep mode and back again, such power loss can be forced on all nodes within transmission range of the attacker. Existing techniques for protecting against replay attack at the link layer have the disadvantage of requiring

resource-constrained sensor nodes to maintain a neighbor table of packet sequence numbers, a requirement that can become unwieldy even in moderately-sized networks. The neighbor table can also be exploited by an attacker if packets from other clusters in the network can be replayed, thereby increasing a node's neighbor table and consuming more resources. If replay protection is only provided at the link layer, inserting packets in this way can cause replayed data to be propagated to it's destination from the new insertion point, thus defeating the anti-replay mechanism. Another weakness in this technique is the requirement for a node to receive the entire replayed packet before dropping it due to a repeated counter, causing another source of overhearing power loss. Better techniques for providing replay protection under resource-constrained conditions are needed. A cross-layer solution in which neighbor information from the network layer is provided to the MAC layer to prevent an attack against the neighbor table is one potential improvement.

Jamming Identification and Mitigation. A strong jamming attack can prevent all sensor nodes' access to the wireless medium and can shut down the network. To reduce costs, sensor nodes are usually equipped with single-channel radios that are not designed to use spread-spectrum techniques to defend against jamming. A logical reaction to jamming is for nodes to go into low-power mode, waking only periodically to sense the medium. With techniques available to reliably identify jamming attacks, such a mechanism is now feasible. Processing and storage requirements for jamming identification, however, may still be beyond the capabilities of the most resource-constrained of sensor platforms. Further work in reducing these requirements is needed to fully realize the potential of these mechanisms.

Broadcast Attack Protection. Most MAC protocols are susceptible to a simple unauthenticated broadcast attack. Long messages can be broadcast and must be received in full by all network nodes before the nodes discard them due to authentication failure. MAC protocols must therefore be designed to prevent all nodes from receiving broadcast messages before these messages are authenticated. An example of such a design is the G-MAC protocol, in which a gateway node receives and authenticates broadcast messages before forwarding them on to other nodes.

Resilience against Compromised Nodes. This is perhaps the most difficult security component to implement. Tamper-resistant cases may mitigate tampering, but will increase the cost of each device. Another option is to develop protocols that can detect compromise, perhaps by identifying a node that is on the network but is not performing its expected mission or is sending unexpected data for its node type. Finally, a key management mechanism that avoids network-wide shared symmetric keys would help to contain the threat of node compromise.

VII. Conclusion and Future Work

Providing security in sensor networks is critical if they are to realize the potential of wide-spread deployment. Current WSN security work focuses on data confidentiality and integrity, largely ignoring availability. Without the ability to secure the physical medium over which communication takes place, sensor networks are susceptible to an array of potential attacks focused on rapidly draining sensor node batteries, thereby rendering the network unusable. This work makes three significant contributions to the area of sensor network security. First, it provides a system for classifying denial-of-sleep attacks on WSN MAC protocols. Second, it explores potential attacks from each attack classification and models their impacts on sensor networks running three leading WSN MAC protocols. Finally, it proposes a framework for defending against denial-of-sleep attacks in sensor networks.

Future work in this area will involve exploring the defensive framework provided here and finding ways to apply it to currently available sensor devices in order to develop specific mechanisms to protect them against these attacks.

References

[1] "Micaz datasheet." Online reference available at: http://www.xbow.com.
[2] "Tmote sky low power wireless sensor module." Online reference available at: http://www.moteiv.com.
[3] S. Singh and C. S. Raghavendra, "PAMAS: Power aware multi-access protocol with signaling for ad hoc networks," In *ACM Comput. Commun. Rev.*, pp. 5-26, 1999.
[4] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient MAC protocol for wireless sensor networks," In *IEEE INFOCOM*, pp. 1567-1576, 2002.
[5] T. VanDam and K. Langendoen, "An adaptive energy-efficient MAC protocol for wireless sensor networks," In *ACM SENSYS*, pp. 171-180, 2003.
[6] M. Brownfield, K. Mehrjoo, A. Fayez, and N. Davis, "Wireless sensor network energy-adaptive MAC protocol," In *IEEE CCNC*, pp. 778-782, 2006.
[7] A. Perrig, R. Szewczyk, J. D. Tygar, V. Wen, and D. E. Culler, "SPINS: Security protocols for sensor networks," In *Wireless Networks*, pp. 521-534, 2002.
[8] A. Perrig, R. Canetti, D. Song, and J. D. Tygar, "Efficient and secure source authentication for multicast," In *NDSS*, pp. 35-46, 2001.
[9] C. Karlof, N. Sastry, and D. Wagner, "TinySec: A link layer security architecture for wireless sensor networks," In *ACM SENSYS*, pp. 162-175, 2004.
[10] "Part 15.4: Wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wpans," IEEE Std. 802.15.4 - 2003 edition, 2003.
[11] A. D. Wood and J. A. Stankovic, "Denial of service in sensor networks," In *IEEE Computer*, pp. 54-62, 2002.
[12] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," In *ACM MOBICOM*, pp. 46-57, 2005.
[13] M. Brownfield, Y. Gupta, and N. Davis, "Wireless sensor network denial of sleep attack," In *IEEE IA Workshop*, pp. 356-364, 2005.
[14] J. Deng, R. Han, and S. Mishra, "Defending against path-based DoS attacks in wireless sensor networks," In *ACM SASN*, pp. 89-96, 2005.
[15] L. Lamport, "Constructing digital signatures from a one-way function," SRI International, Technical Report CSL-98, 1979.
[16] R. Negi and A. Perrig, "Jamming analysis of MAC protocols," Carnegie Mellon Technical Memo, 2003.
[17] Y. W. Law, L. vanHoesel, J. Doumen, and P. Havinga, "Energy-efficient link-layer jamming attacks against wireless sensor network MAC protocols," In *ACM SASN*, pp. 76-88, 2005.
[18] "Tinyos community forum." Online reference available at: http://www.tinyos.net/.