

MiniSec: A Secure Sensor Network Communication Architecture

Mark Luk Ghita Mezzour Adrian Perrig Virgil Gligor
CMU/CyLab CMU/CyLab CMU/CyLab University of Maryland

ABSTRACT

Secure sensor network communication protocols need to provide three basic properties: data secrecy, authentication, and replay protection. Secure sensor network link layer protocols such as TinySec and Zigbee enjoy significant attention in the community. However, TinySec achieves low energy consumption by reducing the level of security provided. In contrast, Zigbee enjoys high security, but suffers from high energy consumption.

MiniSec is a secure network layer that obtains the best of both worlds: low energy consumption and high security. MiniSec has two operating modes, one tailored for single-source communication, and another tailored for multi-source broadcast communication. The latter does not require per-sender state for replay protection and thus scales to large networks. We present a publicly available implementation of MiniSec for the Telos platform, and experimental results demonstrate our low energy utilization.

1. INTRODUCTION

Considerable attention had been paid to developing secure sensor network communication protocols. Unfortunately, existing technologies, such as TinySec and ZigBee, are unable to achieve low energy consumption while simultaneously providing the three important properties of secure communication: secrecy, authentication, and message replay protection.

TinySec, one of the most popular secure link layer protocols, achieves low energy consumption and memory usage. Unfortunately, it also sacrifices on the level of security. First, it employs a single network-wide key, such that every node in the network can masquerade as any other node. Second, TinySec does not attempt to protect against replay attacks.

ZigBee provides a higher level of security than TinySec since it is not restricted to using a network-wide key. By keeping a per-message counter as the Initialization Vector (IV), ZigBee protects against message replay attacks. However, ZigBee is an expensive protocol. First, ZigBee sends the entire 8-byte IV with each packet, resulting in high communication overhead and high energy consumption by the radio. Also, ZigBee requires per-sender state, which consumes a large amount of memory as the number of participants increases.

In this paper, we present MiniSec, a secure network layer protocol for wireless sensor networks. We achieve the best of both

worlds: lower energy consumption than TinySec, and a high level of security like ZigBee. We accomplish this feat by using three techniques. First, we employ a block cipher mode of operation that provides both secrecy and authenticity in only one pass over the message data. Second, we only sending a few bits of the IV, while retaining the security of a full-length IV. In contrast, previous approaches require two passes over the plaintext (one for encryption and one for authentication) and transmission of the full-length IV.

Third, we exploit the fundamental distinctions between unicast and broadcast communication, providing two energy-optimized communication modes. In unicast-mode, we reduce the radio's energy consumption by using synchronized counters and performing extra computation. Although radically different from conventional networking protocols, such a scheme is desirable in this setting because of the stringent energy constraints of sensor networks. In broadcast-mode, we employ a Bloom-filter based replay protection mechanism that precludes per-sender state.

Such improvement in energy consumption comes at the cost of a modest increase in memory size, which is a desirable tradeoff. Note that TinySec had been developed for the Mica2 motes in 2003, where memory constraints were much more stringent than they are today. Current trend in mote development reveals that memory size has been consistently increasing, while energy constraint remains as stringent as ever. Thus, the design tradeoffs in MiniSec makes it well-suited for current state-of-the-art sensor devices.

We present MiniSec as a complete solution to secure sensor network communication, and implement the protocol for the Telos motes [11]. To the best of our knowledge, MiniSec is the first general purpose security protocol available for this popular platform. Furthermore, MiniSec's source code is publicly available, and it can be easily ported to other platforms.

We evaluate MiniSec's performance against TinySec [7]. Our results show that under most circumstances, our energy consumption of security related tasks is about $\frac{1}{3}$ of the energy consumption in TinySec, yet we are still able to provide a higher level of security.

The main contributions of this paper are:

- We introduce MiniSec, the first fully-implemented general purpose security protocol for the Telos sensor motes, the latest generation of sensor motes from the Berkeley family.
- MiniSec simultaneously achieves low energy overhead and a high level of security (data secrecy, authentication, and replay protection). This is the first protocol to achieve replay protection in the sensor network broadcast setting.
- We measure the performance of MiniSec and show that under most real-world scenarios, MiniSec outperforms other comparable systems.
- The source code of MiniSec is publicly available. We provide a turnkey system for sensor network developers that seamlessly integrates into TinyOS applications.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

2. BACKGROUND

MiniSec uses two primitives to achieve low energy consumption: OCB encryption [13] and Bloom Filters [1]. We briefly review both technologies.

OCB. OCB, or Offset CodeBook, is a block-cipher mode of operation that features authenticated encryption. Given a plaintext of arbitrary length, OCB generates a ciphertext that simultaneously provides authenticity and data secrecy. OCB is provably secure, and is parameterized on a block cipher of blocksize n and a tag of length τ . τ is defined such that an adversary is able to forge a valid ciphertext with probability of $2^{-\tau}$.

OCB operates as follows. Let M be an arbitrary length message that needs to be encrypted and authenticated, H be an optional message header that only requires authentication, K be the encryption key (which is the key used by the underlying block cipher), and N be a non-repeating nonce. First, OCB takes in M , K , and N , and generates the ciphertext core C . Next, using the plaintext M , ciphertext C , and H , OCB generates tag of length τ . The final output of $OCB_K(N, M, H)$ is the tuple (C, tag) . To decrypt a ciphertext (C, tag) , the receiver performs the reverse process on C to arrive at plaintext M . Then, the receiver ensures that the tag is as expected. If the receiver computes a different tag than the one in the ciphertext, the ciphertext is considered to be invalid.

OCB is especially well suited for sensor nodes. First, OCB avoids ciphertext expansion. Given an arbitrary length messages M , OCB generates ciphertext with length of $|M| + \tau$. Disregarding the tag, the ciphertext core has the same length as the plaintext. This property cannot be achieved in most other encryption modes without padding and/or ciphertext stealing.

Second, OCB has superior performance, since it provides secrecy and authenticity in one pass of the block cipher. TinySec and ZigBee provide the same guarantees by encrypting and then authenticating the ciphertext. Consequently, the energy consumption is doubled, since they require one pass of CBC-encryption to provide secrecy, and another pass of CBC-MIC to provide authenticity. By comparing the number of block cipher calls, we see that OCB only requires $\lceil \frac{|M|}{n} \rceil$ block cipher calls [13], while CBC-encryption and CBC-MIC together take between $2\lceil \frac{|M|}{n} \rceil + 1$ to $2\lceil \frac{|M|+1}{n} \rceil + 4$ block cipher calls, depending on padding (recall that M is the message, and n is the blocksize).

Bloom filter. The Bloom filter is a space-efficient data structure used for fast probabilistic membership tests [1]. It features two functions: membership addition and membership query. It is possible to have false positives, where a query returns true even though the element is not in the set. However, false negatives, where a query returns false when the element is in fact a member of the set, is not possible.

A Bloom filter requires an array of m bits and k hash functions. The array is first initialized to all zeros. The hash function maps an element to one of the m array positions. To add an element, the element is hashed by all k hash functions, and the k subsequent array positions are set to one. To query for an element, we use the k hash functions to arrive at the k bit positions. If any of them were zero, then the element is not in the set. If all bits were set to one, then either the element is indeed in the set, or all k bits were set to one by insertion of other elements. The latter demonstrates a false positive. The more elements are added to a Bloom filter, the higher the probability of a false positive.

Bloom filters are well suited in the severe resource-constrained environment of sensor nodes because of space and time advantages. The space advantage is apparent, since it only requires $O(n \log n)$ bits to store a fingerprint of n messages. Bloom filters also exhibit the desirable property that adding and querying an element occurs in constant time.

3. PROBLEM DEFINITION

3.1 Assumptions

We assume that symmetric keys are already established between each sender and its receivers. We recommend a different key for each sender, but our protocol is by no means restricted to such a setup. We also assume a secure routing protocol that can route the packet to the intended destination with non-zero probability.

A list of sensor network keying and routing protocols is in Section 9. The goal of MiniSec is to leverage such existing primitives to provide for secure node-to-node communication at low energy cost.

3.2 Attacker Model

Sensor nodes rely on radio broadcasts for communication. We thus adopt the Dolev-Yao attacker model, where an attacker can overhear, intercept, alter, or inject any messages into the radio communication channel.

We do not restrict attackers to computationally bounded nodes. By using sufficiently long symmetric keys (i.e., 80 bits), we can defend against brute force attacks by a powerful adversary [?].

3.3 Desired Properties

We now design the desired properties of a secure sensor network communication architecture.

Data Authentication. Data authentication empowers legitimate nodes to verify whether a message indeed originated from another legitimate node and was unchanged during transmission. As a result, illegitimate nodes should not be able to participate in the network, either by injecting their own messages or by modifying legitimate messages.

Data authentication is one of the basic building blocks of a secure system. For example, nodes need to verify commands from the base station, and a base station needs to authenticate whether the data readings indeed originate from valid nodes.

Typically, data authentication is achieved by the sender computing a message-authentication code (MAC) over the payload and appending that to the message. Upon reception, the packet is considered to be valid if the receiver recomputes the MAC and it matches with the received MAC. ZigBee, TinySec and SNEP provide data authentication by using the CBC-MAC function, using Skipjack or RC5 as the block cipher.

Data Secrecy. Data secrecy, another basic requirement of any secure communication system, prevents unauthorized parties from discovering the plaintext. It is typically accomplished by setting up an encrypted communication channel.

Encryption schemes can be evaluated based on different criteria. A strong level of security is the notion of *semantic security* [?, ?]. Informally, it means that an eavesdropper cannot gain any information about the plaintext, even after observing many encryptions of the same plaintext. Under most circumstances, this can be accomplished through probabilistic encryption schemes that are resilient to chosen-plaintext attacks [?].

In secure communication protocols, data secrecy is provided by a cryptographic encryption function. Furthermore, to guarantee semantic security, encryption functions require a unique initialization vector (or IV) for each encryption to add variation to the ciphertext.

TinySec uses CBC-encryption, while SNEP and ZigBee employ counter-mode. Both schemes provide for semantic security using an IV. MiniSec provides both authentication and secrecy using OCB-encryption, while semantic security is provided by using a counter.

Replay Protection. A replay attack is when attackers record en-

tire packets and replay them at a later time. TinySec is not resilient to such an attack, while SNEP provides protection using a counter. Section 5 and Section 6 respectively demonstrate how MiniSec provides replay protection in the unicast and broadcast setting.

Freshness. Sensor nodes often stream time-varying measurements. Consequently, providing some guarantee of message freshness is an important property. MiniSec provides a mechanism to guarantee weak freshness, where a receiver can determine a partial ordering over received messages. Note that ZigBee and SNEP provide weak freshness, while TinySec does not provide any form of freshness.

Low Energy Overhead. Energy is an extremely scarce resource in sensor nodes. Thus, it is of paramount importance for the security protocol to retain a low energy overhead.

In particular, radio communication consumes the most amount of energy. On the Telos platform, sending a single byte is equivalent to executing about 4720 instructions. Thus, to reduce energy consumption, it is imperative to minimize communication overhead.

Although public key cryptography had enjoyed major advancement recently, it is still 3–4 orders of magnitude more expensive than symmetric cryptography. Therefore, security protocols that only employ symmetric cryptography are preferred.

Resilient to Lost Messages. The relatively high occurrence of dropped packets in wireless sensor networks requires a design that can remain in operation despite high message lost rate.

4. MINISEC OVERVIEW

We present MiniSec, a secure network layer that satisfies all the security properties outlined in Section 3. MiniSec has two operating modes: unicast and broadcast, henceforth known as MiniSec-U and MiniSec-B. Both schemes employ the OCB-encryption scheme to provide for data secrecy and authentication (see Section 2), while using a counter as a nonce to guarantee semantic security. The two modes differ in the way they manage the counters. In MiniSec-U, we employ synchronized counters, which require the receiver to keep a local counter for each sender. MiniSec-B has no such requirement for per-sender state. Instead, meta-data to prevent replay attacks is stored in a Bloom Filter. Both schemes will be explained in detail in Sections 5 and 6, respectively.

Notation. We use the following notation to describe our protocol and cryptographic operations:

A, B	Communicating nodes.
K_{AB}	OCB Encryption key used for communication channel from A to B . Note that key K_{BA} would be used to encrypt data from B to A .
C_{AB}	Monotonically increasing counter corresponding to K_{AB}
$(C, tag) =$ $OCB_K(N, M, H)$	Authenticated encryption under OCB, where M is the plaintext message, H is an optimal message header that only needs to be authenticated, N is a 64-bit nonce, and K is the OCB encryption key.
N_A	A nonce generated by device A .

5. MINISEC-U: UNICAST COMMUNICATION

5.1 Motivation

Both TinySec and SNEP had developed solutions for providing secure communication in the unicast setting, where we have one sender A and one receiver B . Although both protocols attempt to minimize energy consumption, there are aspects of both that demonstrate inefficient energy usage. TinySec uses an encrypted counter as part of its IV. This counter is appended to each message,

resulting in a 2-byte overhead per packet. SNEP, on the other hand, conserves radio energy consumption by not sending the counter with each packet. Instead, the counter is kept as internal state by both sender A and receiver B . However, dropped packets would cause the counters to become inconsistent, in which case both parties need to participate in an expensive counter resynchronization protocol.

TinySec and SNEP represent two extremes: sending the entire counter as opposed to not sending the counter at all. The key insight behind MiniSec-U is that optimal energy usage can be achieved by combining the best of both approaches. Similar to SNEP, MiniSec-U maintains a synchronized monotonically increasing counter C_{AB} between the sender and the receiver. However, MiniSec-U includes the last x bits of the counter along with each packet. We call this the *LB Optimization* (Last Bits), and the last x bits of the counter is called the *LB value*. By keeping x low, the radio’s energy consumption is almost as low as not sending the counter at all.

The LB optimization addresses one of the main flaws of SNEP, which is the high probability of running the expensive counter resynchronization protocol when packets are dropped. Instead, the LB optimization allows resynchronization to occur “implicitly.” Since sender A sends the last x bits of the counter, receiver B can compare the last x bits of its local counter C_{AB} to the LB value appended to the packet. As long as there are fewer than 2^x dropped packets since the last successfully received packet, the receiver can immediately increment his counter such that the final x bits match the LB value. For example, let x be 3, and the counter C_{AB} be synchronized on both parties at 0. Sender A sends six packets, but the first five packets were dropped. Receiver B successfully receives the 6th packet, which was sent using C_{AB} of 5. B would thus immediately increment hits counter C_{AB} to 5 and attempt decryption.

The LB optimization is useful even if more than 2^x packets were dropped, since the receiver could simply continue to increment C_{AB} by 2^x and reattempt decryption. In practice, the receiver B would set a maximum such that after y consecutive failed decryptions, B would run the expensive counter resynchronization protocol.

Lastly, by specifying the parameter x , we could arbitrary lower the probability of reverting to the resynchronization protocol. By monitoring the quality of the channel, it is possible to analytically solve for the optimal values of x and y such that energy consumption is minimized. We demonstrate this in Section 5.3.

5.2 Protocol Description

In this section, we describe the mechanics of MiniSec-U, in the context of sender A and receiver B . In MiniSec-U, each pair of sender and receiver share two keys, K_{AB} to protect communication from A to B , and K_{BA} to protect communication from B to A . Furthermore, a monotonically increasing counter is assigned to each key (C_{AB} used to for key K_{AB}), and is kept as internal state by both sender and receiver.

We employ OCB-encryption with the packet payload as M , packet header as H , counter C_{AB} as the nonce, and K_{AB} as the encryption key. We selected Skipjack to be the underlying block cipher with a blocksize of 64 bits. Since OCB requires the nonce to be the same length as the block size, counter C_{AB} can also be 64 bits. Alternatively, the counter could be of shorter length, and be padded out to 64-bits when requested by the OCB encryption function. The second parameter of OCB is the tag length, which we set to 32 bits, a length suitable for security in retail banking [13]. Finally, receiver B needs to maintain counter C_{AB} . Upon receiving and decrypting a valid packet, B would increment its local copy of C_{AB} accordingly so that it remains consistent with A .

The strictly monotonically increasing counter C_{AB} guarantees semantic security. Even if the sender A repeatedly sends the same message, each ciphertext is different since a different counter value

is used. Also, once a receiver observes a counter value C_{AB} , it rejects packets with an equal or smaller counter value. Therefore, an attacker cannot replay any packet that the receiver has previously received.

There are three interesting issues about counter C_{AB} . First, because of the nature of OCB encryption, the counter itself is not a secret. Even if an attacker knows the counter, the specified security properties are not compromised. This contrasts to CBC encryption used by TinySec, which requires a random and unpredictable nonce as the IV. Second, in order to provide semantic security, the counter cannot wrap around. A longer counter achieves higher level of security, yet adds additional overhead to each packet. Since we do not append the entire counter to each packet, MiniSec enjoys the security benefits of a longer counter without paying the communication cost. This allows us to use a 32-bit or longer counter while TinySec uses a 16-bit counter. Finally, despite the LB optimization, large number of dropped packets could still cause the counters to be desynchronized. Appendix A presents a counter resynchronization protocol similar to one used in SNEP.

5.3 Energy Analysis

Let x be the number of lowest counter bits to append to the packet, and y be the maximum number of trial decryptions the receiver performs as explained above. Let random variable I represent the time at which the node is able to decrypt the message (if the node successfully decrypts the message), where $1 \leq I \leq y$. The goal of this section is to determine the parameters x and y .

Generally, a receiver would perform i decryptions, if there were at least $(i-1)2^x$ and at most $i2^x - 1$ dropped packets. Thus, the probability that exactly i decryptions occur is (where P_{gar} is the probability that a message is garbled in transmission, and p_r is the probability that the message is received):

$$\begin{aligned} P(I=i) &= (1-P_{\text{gar}}) \sum_{k=(i-1)2^x}^{i2^x-1} (1-p_r)^k p_r \\ &= (1-P_{\text{gar}})(1-p_r)^{(i-1)2^x} [1-(1-p_r)^{2^x}] \end{aligned}$$

Thus, the expectation of I is:

$$\begin{aligned} E(I) &= \sum_{i=1}^y iP(I=i) \\ &= [1-(1-p_r)^{2^x}] \sum_{i=1}^y i[(1-p_r)^{2^x}]^{i-1} \\ &= [1-(1-p_r)^{2^x}] \frac{1-(y+1)(1-p_r)^{y2^x} + y(1-p_r)^{(y+1)2^x}}{[1-(1-p_r)^{2^x}]^2} \\ &= \frac{1+(1-p_r)^{y2^x} [1+y[1-(1-p_r)^{2^x}]]}{1-(1-p_r)^{2^x}} \end{aligned}$$

The expected energy consumed is the sum of the expected energy spent when an event happens times the probability of this event to occur plus the energy required for receiving x bits. At the reception of a message, two scenarios are possible for the receiver: (1) the receiver is able to decrypt the message at the i -th trial decryption. This event occurs with probability $P(I=i)$ and its cost is $E(i * E_{\text{OCB}}) = E(I)E_{\text{OCB}}$, where E_{OCB} represents the cost of an OCB decryption. (2) The receiver is unable to decrypt the message even after trying y times. This happens either because more than $y2^x$ packets were lost consecutively, or because the packet that was just received is garbled due to a transmission error (an unlikely event). Thus the expected energy consumption is E_{Energy} , where E_{rec} is the energy for receiving one bit, E_{resync} is the energy required to execute the counter resynchronization protocol, and $E(I)$ is as above:

$$E_{\text{Energy}} = E(I)E_{\text{OCB}}(1-P_{\text{gar}}) +$$

$$[yE_{\text{OCB}} + E_{\text{resync}}] [P_{\text{gar}} + (1-P_{\text{gar}})(1-p_r)^{y2^x}] + xE_{\text{rec}}$$

We need to find the ideal x and y for a given environment. A lossy channel with high packet loss requires larger values for x and y . In Section 8 we discuss how to select x and y in practice.

6. MINISEC-B: BROADCAST COMMUNICATION

MiniSec-U cannot be directly used to secure broadcast communication. First, it would be too expensive to run the counter resynchronization protocol among many receivers. Also, if a node B were to simultaneously receive packets from a large group of sending nodes (A_1, A_2, \dots, A_n) , B needs to maintain a counter for each sender, resulting in high memory overhead.

Similar to MiniSec-U, MiniSec-B uses OCB encryption to secure broadcast communication. Simply encrypting each packet with OCB provides secrecy and authenticity, while an increasing counter can still be used as a nonce to provide for partial ordering of messages. However, without synchronized counters, there is no defense against replay attacks. In fact, defending against replay attacks in a broadcast setting without per-sender state is currently an open challenge in the sensor network community.

This section describes two mechanisms used in MiniSec-B to provide replay protection. First, we present a timing-based approach that defends against replay attacks with a certain vulnerability window; replayed packets outside this window is always dropped. Next, we discuss a Bloom Filter-based approach which defends against attacks within the window.

6.1 Timing-based Approach

We define an epoch to be a finite time t_e , and we segment time into a series of t_e -length epochs E_1, E_2, E_3, \dots . Leveraging loose time synchronization [6, 14], each node agrees on the current epoch E_i . The maximum time synchronization between any pair of nodes is δ_t . Finally, let δ_n be the maximum network latency.

Simply using the current epoch number as the nonce for OCB-encryption defends against replay attacks from older epochs. Unfortunately, because of time synchronization error and network latency, such a scheme experiences high false positives at epoch transitions. For example, a node at time t_1 in Figure 1 would receive many packets from the previous epoch, and thus classify them as invalid packets.

The solution is as follows. First, we define epoch length t_e to be $2\delta_t + \delta_n$. Let E_i be the current epoch number. When a node B receives a packet within the range of $[t_e E_i, t_e E_i + \delta_t + \delta_n)$, it attempts decryption twice, with the nonce being the current epoch number and the immediately previous epoch number. The intuition is as follows. First, in the worst case scenario where the receiver's clock is ahead, we have the sender's clock at δ_t behind and the packet experiences maximum network latency of δ_n . Under this worst case scenario, the packet could have only originated from the current epoch E_i or previous epoch E_{i-1} . Similarly, the worst case scenario where the sender is ahead can be described as the sender's clock being δ_t ahead and the packet experiencing 0 network latency. In such a case, the packet could only have originated from the current epoch E_i or next epoch E_{i+1} .

Based on this technique, there are only two candidate epochs for any incoming packet. In addition, if a valid packet had been sent at the beginning of an epoch, an attacker can replay that packet for the remainder of the epoch as well as $\delta_t + \delta_n$ into the next epoch. Thus, the maximum window of vulnerability for replay is $3\delta_t + 2\delta_n$.

6.2 Bloom Filter-based Approach

We augment the timing-based approach with a counter C_A kept as internal state by the sender A , and two alternating Bloom Filters




Figure 1: Timeline.

stored by all nodes. In concert, they provide semantic security and replay protection within the vulnerability window.

The counter C_A is similar to the counter in MiniSec-U, with the exception that the sender resets the counter at the beginning of each epoch. This allows the counter itself to be short enough to be transmitted with each packet. If we can bound the number of broadcast messages sent by a node in each epoch to k , the length of the counter can be bounded to $\log_2(k)$.

Sender A increments C_A whenever it sends a message M , and uses the concatenation of its own nodeID, C_A and current epoch number E_i as the nonce in OCB encryption. Thus, the cipher-text sent is $(C, tag) = OCB_{K_A}(nodeID || C_A || E_i, M, H)$.

Receiver B keeps its Bloom Filters in the following manner. A Bloom Filter BF_i is assigned to each epoch E_i . All valid packets corresponding to epoch E_i are stored in BF_i , using $(C_A || sourceID)$ as the Bloom Filter key. At all times, each node keeps a Bloom Filter BF_i for the current epoch E_i , and either BF_{i-1} for the previous epoch E_{i-1} , or BF_{i+1} for the next epoch E_{i+1} . For example, starting at the beginning of epoch E_i until $\delta_t + \delta_n$ into the epoch, the node maintains Bloom Filters BF_i and BF_{i-1} . After $\delta_t + \delta_n$, all packets from the epoch E_{i-1} will be dropped, while packets from the next epoch E_{i+1} will be accepted. Thus, BF_{i-1} is reset and reused as BF_{i+1} .

When receiver B receives a packet, it first determines whether it is a valid packet. Since the entire short counter is included as plain-text, B can easily determine its validity after at most two OCB decryptions. This ensures that only packets within this window of two epochs will be accepted. After the packet is deemed to be a valid message, B also discovers its epoch number. B then queries the corresponding Bloom Filter for this packet. If the query returns true, the packet is considered to be a replay and is consequently dropped. If the Bloom Filter declares that this packet is new, the packet is considered to be a non-replayed packet. It is consequently accepted and added into the Bloom Filter.

Using such a counter provides for semantic security since sending the same message never results in the same cipher-text. Also, this scheme provides for replay protection. If receiver B were to receive a replay packet, hashing the source ID and counter C_A , would result in a match in all the corresponding bits in the Bloom Filter. Therefore, B would suspect a replay attack and reject the packet.

Note that such a replay policy would detect all replayed attacks, resulting a 0% false negative rate. However, because Bloom Filters may cause false positives; an innocent packet may be deemed to be a replayed packet. There are various trivial solutions to lowering false positives. For example, by selecting the size of the Bloom Filter m and duration of an epoch, the probability of false positives can be lowered arbitrarily [1].

7. SECURITY ANALYSIS

In this section, we provide an analysis on the level of security promised by both MiniSec-U and MiniSec-B. First, we discuss properties that are common across both protocols. Next, we discuss how these protocols are different.

Authentication. Both MiniSec-U and MiniSec-B use OCB encryption to provide for data authentication over the payload and packet header. The security of OCB’s authentication scheme is directly related to τ , the length of the *tag*. By setting τ to be 32 bits, an adversary has a 1 in 2^{32} chance of forging a correct *tag* for a particular message, which suffices for the majority of practical applications.

Secrecy and Semantic Security. Secrecy and semantic security rest upon the fact that nonces do not repeat. Since the sender uses a strictly monotonic counter as the nonce, each ciphertext would be different even if the plaintext were the same.

Avoiding repetition of nonce is easy. In MiniSec-U, the counter is kept as internal state, and thus can be made arbitrarily long. We choose 8 bytes, which means that the nonce would not repeat until after sending 2^{64} messages. In MiniSec-B, the nonce is the concatenation of sourceID, epoch number e , and a counter. By using all three variables, no two messages in the network would ever share the same nonce. `iiiiiii -mine`

Weak Freshness. In MiniSec-U, the receiver can arrive at the counter value used for each packet by verifying the validity of OCB decryption. While in MiniSec-B, the counter value is included in the packet as plaintext. In both cases, the receiver can use the counter value of two messages to enforce message ordering, thus providing weak freshness.

MiniSec-U and MiniSec-B exhibit different behavior in replay protection.

Replay Protection in MiniSec-U. Each sender and receiver keeps a synchronized counter that is being used as the nonce in OCB encryption. The receiver would only accept messages with higher counter values than the those maintained in the node state. Thus, replayed packets will all be rejected.

Replay Protection in MiniSec-B. The entire network lifetime is segmented into epochs. MiniSec-B leverages loose time synchronization to prevent replayed packets from previous epochs. Next, each receiver uses a Bloom Filter to track packet history for each epoch. Thus, MiniSec-B achieves protection against replay attacks.

8. IMPLEMENTATION

8.1 System Architecture

We have developed MiniSec for the Moteiv Telos motes - a popular architecture in the sensor network research community. It features the 8MHz TI MSP430 micro-controller, a 16-bit RISC processor that is well known for its low energy consumption. Although we implemented MiniSec on the Telos motes, our design principles are general enough such that porting MiniSec to different sensor platforms should yield similar performance results.

To implement MiniSec, we rewrote part of the TinyOS network stack. Specifically, we created `GenericCommMiniSec` based on `GenericComm`, a “generic” TinyOS network stack. Instead of using the interface `AMStandard` for Active Message transmission, `GenericCommMiniSec` directs all messages to `AMStandardMiniSec`, a custom-made `ActiveMessage` layer that encrypts outgoing messages and decrypts received packets. To use MiniSec, a developer simply needs to replace `GenericComm` with `GenericCommMiniSec` in the application’s configuration file.

Packet Format. We base MiniSec’s packet format on the current TinyOS packet header for Telos mote’s CC2420 radio. The Telos mote is one of the first wireless sensor devices to be equipped with

an IEEE 802.15.4 radio. Figure 2 shows the packet format for plain TinyOS, TinySec, and MiniSec.

The fields that MiniSec share with original TinyOS are: length, Frame Control Field, data sequence number, destination PAN address, destination address, and the AM number. MiniSec replaces the 2-byte CRC with a 4-byte tag, since the tag already protects the packet from being tampered. In the original TinyOS, the 1-byte groupID serves as a crude form of access control. This field is no longer necessary in MiniSec because access control is achieved through the use of different cryptographic keys. Finally, similar to TinySec, MiniSec requires a 2-byte source address, which is absent in a standard TinyOS packet. The net overhead of a MiniSec packet is 3-byte increase over a standard TinyOS packet.

8.2 MiniSec-U Design

MiniSec-U uses two security primitives: OCB encryption and Skipjack. We selected Skipjack to be the block-cipher because of efficient computation and low memory footprint [?]. To make encryption as flexible as possible, we set Skipjack’s block size to 64 bits. Furthermore, we use 80-bit symmetric keys, since Lenstra and Verheul recommended that such keys are considered to be secure until 2012, even against resourceful adversaries [?]. When 80-bit keys become insecure, we would use 128-bit AES keys, which is secure for at least the next 20 years.

OCB implementation is ported from Rogaway’s original OCB library¹ to the nesC interface. Similarly, the Skipjack implementation is ported to nesC from Yee Wei Law’s implementation [?]. In total, MiniSec-U requires about 4000 lines of nesC code, and consumes 874 bytes of RAM and 16 KB of code memory.

Packet Format. MiniSec-U uses the LB optimization by sending the last x bits of the sender’s counter along with each packet. Since TinyOS payloads are never greater than 29 bytes, we can safely overload the first 3 bits of the length field to store these bits, as shown in Figure 2(c). This is a significant advantage since we do not suffer any communication overhead for sending the last x bits of the counter.

Our empirical results show that by using the last 3 bits of the counter, even under high packet drop rate, the counter resynchronization protocol was rarely executed. Since we can send the last 3 bits of the counter for free, we use the default value of $x = 3$ for the remainder of the paper.

8.3 MiniSec-B Design

In addition to the security primitives in MiniSec-U, MiniSec-B utilizes loose time synchronization and Bloom filters. Here, we discuss practical issues in selecting epoch duration t_e and Bloom filter configurations.

Time Synchronization. As discussed in Section 6, epoch length t_e must be at least $2\delta_r + \delta_n$. Recent advancement in secure sensor network time synchronization [6, 14] enables pairwise time synchronization with error of mere μs . Transmission delay between neighboring nodes are on the order of ms. Even under extreme pessimistic conditions, epoch length of 1 second is longer than necessary according to the needs of MiniSec-B. For the remainder of our analysis, we will use epoch duration $t_e = 1s$.

Bloom filter configurations. A Bloom filter is defined by two configurations: size of the Bloom filter m , and number of hash functions h . We will show that under rather pessimistic assumptions of the hardware and network activity, we can achieve a 1% false positive rate by using $t_e = 1s$, $m = 18$ bytes, and $h = 8$ hash functions. We also show how the sensor network administrator can calculate custom values of m and h appropriate for a particular network pa-

rameterized on the network activity and underlying hardware.

The false positive rate of a Bloom filter can be calculated based on number of stored items. Thus, we first upper bound p_μ , the average number of packets received in one epoch of length t_e . If this is known a priori (e.g., regular heartbeats), the sensor network administrator can directly configure the Bloom filter accordingly. Else, we make the following argument.

Let t_l be a realistic lower bound of a node’s lifetime, E_c be energy capacity of the node’s battery, and E_p be energy consumption for receiving one packet. In the worse case, all energy provided by the battery will be used for packet reception. Thus, maximum possible packets received over the entire lifetime of the node is E_c/E_p . p_μ can be calculated as follows.

$$p_\mu = \frac{E_c t_e}{E_p t_l}$$

In our calculations, we set t_l to 12 months, E_c to 2850 mAH (AA Duracell Coppertop alkaline battery²), and E_p to 0.0441mAS (receiving a TinyOS packet on CC2420 radio [11]). Average packet reception rate p_u is 7.48 packets per second.

Each packet adds one item into our Bloom filter. In traditional networking fashion, we model packet reception as a Poisson process. Thus, the number of packets received within an epoch can be approximated by a Poisson distribution with mean of p_μ . This model allows us to bound the maximum number of received packets in an epoch with high probability. The cumulative distribution function of a Poisson process is $\frac{\Gamma(k+1, \lambda)}{k!}$, where k is number of occurrences, λ is the Poisson mean p_μ , and Γ is the incomplete gamma function. By setting the CDF to an arbitrarily high probability p and solving for k , we can conclude that one would not receive more than k packets in one epoch with probability p . In our case, we set p to 0.99, and arrived at $k = 14$. In other words, with probability of 0.99, we would never add more than 14 items to our Bloom filter.

Given this information, we can set a particular false positive rate and solve for appropriate configurations for the Bloom filter size m and number of hash functions h . This problem had been previously studied in Li Fan’s work in Summary Cache, where they evaluated the statistics behind Bloom filters [8]. The probability of a false positive after inserting n elements is

$$\left(1 - \left(1 - \frac{1}{m}\right)^{kn}\right)^k$$

Thus, with the worst case of $n = 14$ elements, we can achieve a 1% false positive rate with $m = 18$ bytes and $h = 8$ hash functions.

Note that this bound is an extremely pessimistic bound. In practice, the false positive rate should be significantly lower. There are several factors:

- False positive rate increases as more packets are added into the Bloom filter. However, our calculation is based on the false positive rate at the end of the epoch, which corresponds to the highest false positive rate. For example, we stated that we achieve a 1% false positive rate with our Bloom filter configuration. However, if we were in the middle of an epoch, our calculation actually shows a 0.014% false positive rate.
- We make the pessimistic assumption that all energy is consumed by the radio for packet reception. In practice, energy is consumed for sending packets, controlling other devices (LEDs, sensors), and computation.
- We model the packet reception as a Poisson process. However, we don’t use the mean packet arrival rate for false positive calculation. Instead, we use the 99th percentile based on

¹<http://www.cs.ucdavis.edu/rogaway/ocb/code-2.0.htm>, accessed June 26, 2006

²www.duracell.com/oem/Pdf/others/ATB-full.pdf

1	2	1	2	2	1	1	0...28	2
Len	FCF	DSN	DstPAN	DstAddr	AM	Grp	Data	CRC

(a) TinyOS

1	2	1	2	2	1	2	2	0...28	4
Len	FCF	DSN	DstPAN	DstAddr	AM	Src	Ctr	Enc Dta	MIC

(b) TinySec-AE

1	2	1	2	2	1	2	0...28	4
Len	FCF	DSN	DstPAN	DstAddr	AM	Src	Enc Dta	Tag/MIC

Ctr[0:2]

(c) MiniSec-U

1	2	1	2	2	1	2	0...28	4
Len	FCF	DSN	DstPAN	DstAddr	AM	Src	Enc Dta	Tag/MIC

Ctr[0:2]

Ctr[3:7]

(d) MiniSec-B

Figure 2: Packet Format. The size of each field is indicated in bytes. The packet format is based on the TinyOS packet for CC2420 radio. The diagonally hashed regions are authenticated; the checkered regions are encrypted; the gray region represents where we overload the corresponding field with the counter.

Table 1: Table comparing packet size, communication overhead, and total energy spent in transmitting one TinyOS packet. Our of all three security protocols, MiniSec achieves the lowest communication overhead with respect to a standard TinyOS network stack.

	Payload (B)	Packet Overhead (B)	Security Overhead (B)	Total Size (B)	Energy (mAs)	Increase over TinyOS
TinyOS	24	12	–	36	0.034	–
TinySec	24	17	5	41	0.0387	13.9%
SNEP	24	20	8	44	0.0415	22.2%
MiniSec	24	25	3	39	0.0368	8.3%

the CDF of the Poisson distribution. Thus, with probability of 0.99, we would have a false positive of at most 1%. As a matter of fact, using the same Bloom filter configuration, false positive rate would not exceed 0.0001% with probability of 0.5.

- We assume ideal conditions for batteries (constant current draw, constant temperature, no self-discharge prior to being loaded onto the motes). In practical settings, such ideal conditions are impossible to achieve. Thus, battery capacity and number of received packets, are significantly lower.

Packet Format Figure 2(d) shows the packet header for MiniSec-B where the sender sends the entire counter with each packet. The default counter is 8 bits long, which we claim to be sufficient in most networks based on the following reasoning. First, since the counter is reset at the beginning of each epoch, the length of the counter can be bounded by $\log_2(k)$, where k is the maximum numbers of packets sent in each epoch. Next, it is unlikely for k to be large, since such resource-constrained nodes are unlikely to continually broadcast large amounts of data. An 8-bit counter is already extremely generous since it allows for 255 packets per epoch of one second.

This 8-bit counter could be declared as an addition field. However, based on the TinyOS packet format, we propose an optimiza-

tion that overloads this counter with existing headers. First, 3 bits of this counter may be overloaded in the length field. Next, the remaining 5 bits of the counter may be embedded in the destination address, since the destination field is 2 bytes and it is unlikely for a network to have more than 2048 broadcasting participants. Furthermore, unlike the source address, the destination address is generally not needed for routing in broadcasts.

8.4 Benchmarks

To analytically evaluate the cost of security, we consider both the communication overhead on each packet as well as computational overhead from packet processing. Longer packets are extremely costly because of the extra energy consumption by the radio. Even sending one additional byte per packet would require significant amount of energy. As shown in Table 1, MiniSec was able to decrease a security overhead of 5 bytes by TinySec to 3 bytes. In addition, MiniSec employs OCB, which provides for authentication and secrecy with fewer block cipher calls than its cryptographic counterpart in TinySec.

MiniSec-U. MiniSec-U was able to save on packet header size by using synchronized counters between sender and receiver. First, as specified above, the LB optimization has the default value of $x = 3$. Using the expected energy equation from Section 5.3, it is possible to solve for the the optimal y , or max number of decryption attempts




Figure 3: Optimal y : Max number of decryption attempts before counter resynchronization. Also shows aggregate energy consumption of all security related features.




Figure 4: Drop Rate. Expected energy of security with respect to packet drop rate.

before counter resynchronization. Consequently, given the optimal y , one can solve for the expected energy consumed by all security related features.

Figure 3 illustrates our findings given different packet drop rates. The aggregate energy consumed by TinySec is constant. The energy overhead of security in TinySec is always the amount of energy consumed by receiving 5 extra bytes in the header, computing CBC-MAC, and decrypting the packet. MiniSec-U, on the other hand, behaves differently based on environmental conditions. resynchronization. Given a packet drop rate, a higher y value exhibits lower energy consumption, since it reduces the probability of running the counter resynchronization protocol. However, an increasing y exhibits diminishing marginal benefits, since the energy consumption of counter resynchronization becomes less and less of a factor in total energy consumption.

MiniSec-B. Figure 4 illustrates energy consumption between TinySec, MiniSec-B, and MiniSec-U using an optimal value y , while varying packet drop rate. The energy consumption of MiniSec-U is computed by first solving for the optimal y . Using this value, we were able to calculate the expected energy of security. MiniSec-B consumes a constant amount of energy, as its performance is not effected by lossiness of the communication channel. By leveraging an 8-bit counter and loose time synchronization, the receiver never needs to run any counter resynchronization protocol. Once the receiver successfully receives a packet, only two OCB decryptions and 8 hash functions needs to be performed.

Under normal circumstances, MiniSec consumes about $\frac{1}{3}$ amount

of energy of TinySec. As packet drop rate increases beyond 0.9, TinySec is more efficient than MiniSec-U because of the high number of counter resynchronizations. Such a scenario represents an extremely rare case. On the other hand, since MiniSec-B’s performance does not depend on the drop rate, MiniSec-B always outperforms TinySec.

Finally, we note that MiniSec-B can also be used for unicast. The advantage of MiniSec-B is the amount of work is constant regardless of the quality of the communication link. Under the condition that MiniSec-U does not execute counter resynchronization, MiniSec-U and MiniSec-B consumes comparable amount of energy.

Finally, we note that nothing prevents the use of MiniSec-B in unicasts. Since the energy consumption of MiniSec-B and MiniSec-U are comparable under normal conditions, while MiniSec-B far outperforms MiniSec-U under high packet loss, MiniSec-B is a much more robust protocol. If time synchronization is available, simply running MiniSec-B for all communication is an attractive solution.

9. RELATED WORK

Key establishment and management are considered to be prerequisites in secure sensor network communication, and have been extensively studied in the research community. There are numerous candidate solutions, such as Random Key Predistribution [3–5, 9, 12], Key Infection [2], and LEAP [16]. Asymmetric schemes based on elliptic curve cryptography [10] and Diffie-Hellman [15]. have also been proposed

The body of work most closely related to MiniSec consists of other secure communication protocols such as TinySec [7], ZigBee [17], and SNEP [?]. SNEP, part of the SPINS protocol suite, is one of the first attempts at a secure link layer protocol. It achieves low energy consumption by keeping a consistent counter between the sender and receiver, such that an IV is not required to be appended to each packet. However, packet loss would cause the counters to become inconsistent. Consequently, SNEP would have to execute a counter resynchronization protocol that is slow and expensive in terms of energy consumption. Inspired by SNEP, MiniSec makes various improvements to lower energy consumption. One such optimization is to reduce the probability that the resynchronization protocol needs to be executed.

TinySec is a link layer protocol designed by Karlof et al. [7]. It achieves low energy consumption by reusing part of the packet header as the IV. Thus, they were able to arrive at an 8-byte IV, but only adding a 2-byte counter overhead per packet. However, more serious drawbacks of TinySec are that (1) it only provides for authentication and data secrecy, but not replay protection; (2) it uses a single network-wide key, which is vulnerable to single node compromise.

ZigBee is a set of security standards proposed by a consortium interested in promoting embedded wireless technology. The security protocol is very similar to SNEP. However, the entire 8-byte counter is sent in the clear instead of being kept as consistent state between sender and receiver. Thus, ZigBee consumes significantly more energy than the other two protocols.

10. CONCLUSIONS

Battery energy is the main resource to protect in current wireless sensor networks. Researchers have proposed several approaches for securing communication; however, so far either optimizing either for a high level of security or for a low energy utilization. Our secure sensor network communication package is called MiniSec, which offers a high level of security while typically requiring less energy than previous approaches. We have implemented MiniSec

on Telos motes and distribute it for free under an open-source license.

11. REFERENCES

- [1] B. Bloom. Space/time trade-offs in hash coding with allowable errors. In *Communications of the ACM*, July 1970.
- [2] H. Chan, R. Anderson, and A. Perrig. Key infection: Smart trust for smart dust. In *IEEE ICNP*, May 2004.
- [3] Haowen Chan, Adrian Perrig, and Dawn Song. Random key predistribution schemes for sensor networks. In *IEEE Symposium on Security and Privacy*, May 2003.
- [4] Wenliang Du, Jing Deng, Yunghsiang S. Han, and Pramod K. Varshney. A pairwise key pre-distribution scheme for wireless sensor networks. In *ACM CCS*, October 2003.
- [5] L. Eschenauer and V. D. Gligor. A key-management scheme for distributed sensor networks. In *ACM CCS*, November 2002.
- [6] S. Ganeriwal, S. Capkun, C. Han, and M. B. Srivastava. Secure time synchronization service for sensor networks. In *WiSe*, 2005.
- [7] C. Karlof, N. Sastry, and D. Wagner. TinySec: A link layer security architecture for wireless sensor networks. In *ACM SenSys*, November 2004.
- [8] J. Almeida L. Fan, P. Cao and A. Broder. Summary cache: A scalable wide-area web cache sharing protocol. In *ACM SIGCOMM 98*, 1998.
- [9] D. Liu, P. Ning, and W. Du. Group-based key pre-distribution in wireless sensor networks. In *WiSe*, April 2005.
- [10] D. Malan, M. Welsh, and M. Smith. A public-key infrastructure for key distribution in TinyOS based on elliptic curve cryptography. In *IEEE SECON*, October 2004.
- [11] Joseph Polastre, Robert Szewczyk, and David Culler. Telos: Enabling ultra-low power wireless research. In *IPSN/SPOTS*, 2005.
- [12] M. Ramjumar and N. Memon. An efficient key predistribution scheme for ad hoc network security. In *IEEE Journal of Selected Areas in Communications*, March 2005.
- [13] P. Rogaway, M. Bellare, and J. Black. OCB: A block-cipher mode of operation for efficient authenticated encryption. In *ACM TISSEC*, November 2001.
- [14] K. Sun, P. Ning, C. Wang, A. Liu, and Y. Zhou. TinySeRSync: Secure and resilient time synchronization in wireless sensor networks. In *ACM CCS*, November 2006.
- [15] R. Watro, D. Kong, S.-F. Cuti, C. Gardiner, C. Lynn, and P. Kruus. TinyPK: Securing sensor networks with public key technology. In *SASN*, October 2004.
- [16] S. Zhu, S. Setia, and S. Jajodia. LEAP: efficient security mechanisms for large-scale distributed sensor networks. In *ACM CCS*, 2003.
- [17] ZigBee Alliance. Zigbee specification. Technical Report Document 053474r06, Version 1.0, ZigBee Alliance, June 2005.

APPENDIX

A. MiniSec-U COUNTER RESYNCHRONIZATION

Reliable Unicast. If the message type is a reliable unicast, this implies there exists some type of acknowledgment protocol. Thus, the sender could use the presence of ACKs to determine if the packets had been received and authenticated successfully, which implies whether the counter is consistent between sender and receiver.

However, MiniSec is intended to be a network layer protocol. Since reliable messaging is typically implemented at a higher layer, this approach might not be appropriate. Nevertheless, we note that if reliable messaging is used, cooperation between the transport layer and MiniSec would be a viable solution.

Best Effort Unicast. Without support for reliable message delivery, TinyOS packets are best effort. In the case of unicast messages, where there only exists one receiver B , B can directly query sender

A for the counter. Note that we use a nonce N_B to guarantee strong freshness.

$$\begin{aligned} B \rightarrow A : & \quad \langle N_B \rangle \\ A \rightarrow B : & \quad \langle C_A, OCB_{K_A}(C_A, N_B, C_A) \rangle \end{aligned} \quad (1)$$

Since the counter C_A is not a secret, A can send it in the clear. However, this value does need to be authenticated, using the pre-existing shared key K_A .