Security in Wireless Sensor Networks

Adrian Perrig John Stankovic David Wagner

March 22, 2004

1 Introduction

Today, wireless sensor networks are used for a wide variety of applications: ocean and wildlife monitoring, manufacturing, building safety and earthquake monitoring, and many military applications. An even wider spectrum of future applications may follow, such as monitoring of traffic, pollution, wildfires, building security, water quality, and even people's heart rates. A major benefit of these systems is that they can perform in-network processing to reduce large streams of raw data into useful aggregated information. It is critical to protect this information.

Sensor networks pose unique new challenges which prevent direct application of traditional security techniques. First, to make sensor networks economically viable, sensor devices are limited in their energy, computation, and communication capabilities. Second, in contrast to traditional networks, sensor nodes are often deployed in accessible areas, presenting a risk of physical attacks. Third, sensor networks interact closely with their physical environment and with people, posing new security problems. Consequently, existing security mechanisms are inadequate, and new ideas are needed. Fortunately, these are exciting new problems to inspire research, and we have an opportunity to properly address sensor network security from the start.

This article outlines the security issues in wireless sensor networks, discusses the current state of the art, and suggests future directions for research. We cover several important security challenges, including key establishment, secrecy and authentication, privacy, robustness to denial-of-service attacks, secure routing, and node capture. Then, we discuss several high level security services required for wireless sensor networks. We conclude with several research challenges for the future.

2 Security Goals

Security is sometimes considered a stand-alone component of an architecture, where a separate module provides security. This is usually a flawed approach. To achieve a secure system, security must be integrated into every component, since components designed without security can become a point of attack. Consequently, security pervades every aspect of system design. This section discusses some of the challenges in integrating security into sensor network architectures.

2.1 Key Establishment, Trust Setup

When setting up a sensor network, one of the first requirements is to establish cryptographic keys for later use. Key establishment is a well-studied problem—researchers have proposed a variety of protocols over the past decades. Why can't these key establishment protocols simply be used in sensor networks? The properties of sensor networks render previous protocols impractical. First, many current sensor devices have limited computational power, making public-key cryptographic primitives too expensive. Second, key establishment techniques need to scale to networks with hundreds or thousands of nodes. Third, the communication patterns of sensor networks differ from traditional networks; sensor nodes may need to set up keys with their neighbors and with data aggregation nodes.

The simplest solution for key establishment is to use a network-wide shared key. Unfortunately, compromise of any one node reveals the secret key and thus allows decryption of all traffic on the network. One variant on this idea is to use a single shared key to establish a set of link keys, one per pair of communicating nodes, and erase the network-wide key after the setup of session keys. However, with this variant we cannot add new nodes after initial deployment.

Public-key cryptography (e.g., Diffie-Hellman key establishment) is another option, but as discussed above, it is beyond the capabilities of today's sensor networks. Its main advantage is that a node can set up a secure key with any other node in the network.

Another approach is to pre-configure the network with a shared unique symmetric key between every pair of nodes, but this scales poorly. In a sensor network with n nodes, each node needs to store $n-1$ keys, and $n \cdot (n-1)/2$ keys need to be established in the network.

Bootstrapping keys using a trusted base station is another possibility. Here each node only needs to share a key with the base station and sets up keys with other nodes through the base station [6]. This makes the base station a single point of failure, but because there is only one base station, we may be able to afford tamper resistant packaging for the base station to ameliorate the threat of physical attacks.

Recently, researchers discovered *random-key predistribution* key establishment protocols [3]. In these protocols, a large pool of symmetric keys is chosen, and a random subset of the pool is distributed to each sensor node. Two nodes that want to communicate search their pools to determine whether they share a common key—if they share a key, they use that key to establish a session key. Not every pair of nodes share a common key, but if the key establishment probability is sufficiently large, nodes can still set up keys with sufficiently many nodes to obtain a connected network. This avoids the need for a central trusted base station. The disadvantage of this approach is that if an attacker can compromise sufficiently many nodes, he can reconstruct the complete key pool and break this scheme.

In the future, we expect to see continued research on better random-key predistribution schemes providing high resilience to node compromise, as well as investigation of hardware support for public-key cryptography and more efficient public-key schemes (such as elliptic curve cryptography). Ultimately, we need a secure and efficient key distribution mechanism that allows simple key establishment for large-scale sensor networks.

2.2 Secrecy and Authentication

Similar to traditional networks, most sensor network applications require protection against eavesdropping, injection, and modification of packets. The standard defense is cryptography.

Interesting systems tradeoffs arise when incorporating cryptography into sensor networks. For point-to-point communication, end-to-end cryptography achieves a high level of security, but requires keys set up between all end points and is incompatible with passive participation and local broadcast. Link-layer cryptography with a network-wide shared key simplifies key setup and supports passive participation and local broadcast, but intermediate nodes can eavesdrop or alter messages. We envision that the earliest sensor networks may use link-layer cryptography, because this provides greatest ease of deployment, but subsequent systems may respond to demands for more security with more sophisticated use of cryptography.

Cryptography comes at a performance cost, requiring extra computation and often increasing the packet size. Cryptographic hardware support increases efficiency, but also increases cost. Therefore, an important question facing sensor node design is: is it possible to achieve reasonable security and performance levels with software-only cryptographic implementations, or is hardware support needed?

Recent research demonstrates that software-only cryptography is practical with today's technology; hardware support is not needed to achieve acceptable security and performance levels. For instance, the Berkeley implementation of TinySec incurs only a 5-10% performance overhead using software-only methods. Those experiments also reveal an interesting phenomenon: most of the performance overhead can be attributed to the increase in packet size; in comparison, the cryptographic computations have almost no impact on latency or throughput, since they can overlap with transmission. This puts a limit on how much dedicated hardware will help, because hardware can only reduce the computational costs, not packet sizes.

2.3 Privacy

Sensor networks have thrust privacy concerns to the forefront. The most obvious risk is that ubiquitous sensor technology could allow ill-intentioned individuals to deploy secret surveillance networks for spying on unaware victims. Employers might spy on their employees, shop owners might spy on customers, neighbors may spy on each other, or law enforcement might spy on public places. This is certainly a valid concern; throughout history, as surveillance technology has become cheaper and more effective, it has increasingly been implicated in privacy abuses. Technology trends suggest that this problem is only going to get worse with time: as devices become smaller, they will be easier to conceal; as devices become cheaper, surveillance networks become more affordable.

Another risk is that sensor networks that are initially deployed for legitimate reasons might subsequently be used in unanticipated ways. The notion of "function creep" is universal in the privacy literature. For instance, in the United States, Social Security Numbers were originally intended for use only by the social security program, but since then have gradually come to be widely used as an all-purpose identification number.

The networked nature of sensor networks raises new threats that are qualitatively different from what we've faced before. Sensor networks allow data collection, coordinated analysis, and automated event correlation. For instance, networked systems of sensors enable routine tracking of people and vehicles over long time periods, which has troubling implications.

It is unlikely that technology alone will be able to solve the privacy problem; rather, a mix of societal norms, new laws, and technological responses are necessary. As a starting point, fair information practices may provide a reasonable guideline for how to build systems that better protect privacy. The notion of "notice" seems particularly important: if affected parties are aware of the existence, form, and implications of surveillance, they are more likely to accept the technology. However, our current understanding of privacy in sensor networks is not yet mature, and more research is needed.

2.4 Robustness to Communication Denial of Service

Adversaries can severely limit the value of a wireless sensor network by denial-of-service attacks [8]. In the simplest form of denial-of-service attack, an adversary attempts to disrupt operation by broadcasting a high-energy signal. If the transmission is strong enough, the entire system could be jammed. More sophisticated attacks are also possible: the adversary can inhibit communication by violating the MAC protocol, for instance by transmitting while a neighbor is also transmitting or by continuously requesting channel access with a RTS (request-to-send).

One standard defense against jamming employs spread-spectrum communications [1]. However, cryptographically secure spread-spectrum radios are currently not commercially available. Also, this defense is not secure against adversaries who can capture nodes and extract their cryptographic keys.

Interestingly, the networked nature of sensor networks allows new, automated defenses against denial of service. When the jamming only affects a portion of the network, a jamming-resistant network could defeat the attack by detecting the jamming, mapping the affected region, and then routing around the jammed area [9]. Further progress in this area may allow for greater security against denial-of-service attacks.

2.5 Secure Routing

Routing and data forwarding is an essential service in sensor networks to enable communication. Unfortunately, current routing protocols suffer from many security vulnerabilities [5]. For example, an attacker can easily perform denial-of-service attacks on the routing protocol, often preventing communication. The simplest attacks consist in injecting malicious routing information into the network that results in routing inconsistencies. Simple authentication can guard against such injection attacks, but some routing protocols are even susceptible to replay by the attacker of legitimate routing messages [4].

Routing protocols are particularly susceptible to node capture attacks, which we describe in more detail below. For instance, researchers have analyzed 14 protocols for routing in sensor networks and found that they are all highly susceptible to node capture attacks: in every case, compromise of a single node suffices to take over the entire network or to prevent communication [5]. It is an open research problem to devise secure routing protocols that are robust against such attacks.

2.6 Resilience to Node Capture

One of the most challenging issues facing sensor networks is resiliency against node capture attacks. In traditional computing, physical security is often taken for granted: attackers can be denied physical access to our computers. Sensor networks disrupt that paradigm. In most applications, sensor nodes are likely to be placed in locations accessible to an attacker. This raises the possibility that an attacker could capture sensor nodes, extract cryptographic secrets, modify their programming, or replace them with malicious nodes under the attacker's control.

Tamper-resistant packaging might be one defense, but this is expensive. Also, current tamper resistance technology does not provide a very high level of security. Therefore, algorithmic solutions to the problem of node capture would be preferable.

The challenge, then, is to build networks that operate correctly even when, unbeknownst to us, a few nodes have been compromised and thus might behave in an arbitrarily malicious way. One promising direction for building resilient networks is to replicate state across the network and use majority voting and other techniques to detect inconsistencies. For example, several authors have designed routing protocols that achieve some resilience against node capture by sending every packet along multiple, independent paths and checking at the destination for consistency among what is received [2].

A second direction for resilience is to gather multiple, redundant views of the environment and cross-check them for consistency. For instance, three reports of an interesting event might be required from three independent nodes before responding to this event. As another example, when many data values are collected, it is reasonable to construct a histogram; extreme outliers may indicate malicious spoofed data, and hence should be ignored.

Defenses based on redundancy are particularly well-suited to sensor networks, because a constellation of many cheap nodes may be able to provide more reliable network operation than a small group of a few, more sophisticated devices. Nonetheless, the node capture problem is one of the most challenging problems in sensor network security, and we are a long way from a good solution.

3 Network Security Services

So far, we discussed low-level security primitives. In this section, we discuss high-level security mechanisms such as secure group management, intrusion detection, and secure data aggregation.

3.1 Secure Group Management

Each node in a wireless sensor network is quite limited in capabilities. However, interesting in-network data aggregation and analysis can be performed by groups of nodes. For example, a group of nodes may be responsible for jointly tracking a vehicle through a wireless sensor network. The actual nodes comprising the group may change continuously and quickly. Many other key services in wireless sensor networks are also performed by groups.

Consequently, secure protocols for group management are required. This includes securely admitting new group members and supporting secure group communications. The outcome of the group's computation is normally transmitted to a base station. Such output must be authenticated to ensure that it is coming from a valid group. Any solution must be efficient in time and energy, which precludes many classical group management solutions.

3.2 Intrusion Detection

Wireless sensor networks are susceptible to many forms of intrusion. In wired networks, traffic and computation are typically monitored and analyzed for anomalies at various concentration points. This is often expensive. Wireless sensor networks require a solution that is fully distributed and inexpensive in communication, energy and memory requirements. In order to look for anomalies, applications and typical threat models must be understood. It is particularly important to understand how cooperating adversaries might attack the system. The use of secure groups may be one approach for decentralized intrusion detection.

3.3 Secure Data Aggregation

One benefit of a wireless sensor network is the fine-grained sensing that a large and dense set of nodes can provide. The sensed values must be aggregated to avoid overwhelming amounts of traffic back to the base station. There are many types of aggregation that can be performed. For example, the system may average the temperature or humidity of a geographic region, combine sensor values to compute the location and velocity of a moving object, or aggregate data to avoid false alarms in real world event detection. Depending on the architecture of the wireless sensor network, there may be many places in the network where aggregation takes place, and all aggregation locations must be secured. If the application is able to tolerate approximate answers, powerful techniques are available: under appropriate trust assumptions, randomly sampling a small fraction of nodes and checking that they have behaved properly allows detection of many attacks [7].

4 Summary and Research Challenges

Security is a difficult challenge for any system. The severe constraints and demanding deployment environments of wireless sensor networks make computer security for these systems even more challenging in many respects. On the other hand, several opportunities exist that may help with the challenge of building secure wireless sensor networks. First, we have the opportunity to architect security solutions into sensor systems from the outset, since these systems are still in the early design and research stages. Second, many applications envision deploying sensor networks under a single administrative domain, simplifying the threat model. Third, it may be possible to exploit redundancy, scale, and the physical characteristics of the environment in solutions. If we build sensor networks so that they will continue operating even if a fraction of sensors are compromised, we have an opportunity to use redundantly deployed sensors to resist attack. Ultimately, the unique aspects of sensor networks may allow novel defenses not available in conventional networks.

Many problems need further research. One challenge is how to secure wireless communication links against eavesdropping, tampering, traffic analysis, and denial of service. Also, resource constraints create many unique challenges for sensor network security; ongoing research directions include, for instance, work on asymmetric protocols where most of the computational burden falls on the base station and on public-key cryptosystems that are efficient on low-end devices. Finally, finding ways to tolerate the lack of physical security, perhaps through use of redundancy or knowledge about the physical environment, will remain a continuing research challenge. We are optimistic that much progress can be made on these problems.

References

[1] David Adamy. *EW 101: A First Course in Electronic Warfare.* Artech House, February 2001.

[2] J. Deng, R. Han, and S. Mishra. A Performance Evaluation of Intrusion-Tolerant Routing in Wireless Sensor Networks. In *2nd IEEE International Workshop on Information Processing in Sensor Networks (IPSN 2003)*, April 2003.

[3] Laurent Eschenauer and Virgil D. Gligor. A key-management scheme for distributed sensor networks. In *Proceedings of the 9th ACM Conference on Computer and Communication Security*, pages 41–47, November 2002.

[4] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Packet leashes: A defense against wormhole attacks in wireless ad hoc networks. In *Proceedings of IEEE INFOCOM 2003*, April 2003.

[5] Chris Karlof and David Wagner. Secure routing in wireless sensor networks: Attacks and countermeasures. In *Proceedings of First IEEE International Workshop on Sensor Network Protocols and Applications*, May 2003.

[6] Adrian Perrig, Robert Szewczyk, Victor Wen, David Culler, and J. D. Tygar. SPINS: Security protocols for sensor networks. In Journal of *Wireless Networks*, 8(5):521–534, September 2002.

[7] Bartosz Przydatek, Dawn Song, and Adrian Perrig. SIA: Secure information aggregation in sensor networks. In *Proceedings of the First ACM International Conference on Embedded Networked Sensor Systems (SenSys 2003)*, pages 255–265, November 2003.

[8] Anthony D. Wood and John A. Stankovic. Denial of service in sensor networks. *IEEE Computer*, pages 54–62, October 2002.

[9] Anthony D. Wood, John A. Stankovic, and Sang Son. JAM: A mapping service for jammed regions in sensor networks. *IEEE Real-Time Systems Symposium*, December 2003.