SECURITY IN WIRELESS SENSOR NETWORKS

By ADRIAN PERRIG, JOHN STANKOVIC, *and* DAVID WAGNER

They are susceptible to a variety of attacks, including node capture, physical tampering, and denial of service, while prompting a range of fundamental research challenges.

Wireless sensor network applications include ocean and wildlife monitoring, manufacturing machinery performance monitoring, building safety and earthquake monitoring, and many military applications. An even wider spectrum of future applications is likely to follow, including the monitoring of highway traffic, pollution, wildfires, building security, water quality, and even people's heart rates. A major benefit of these systems is that they perform in-network processing to reduce large streams of raw data into useful aggregated information. Protecting it all is critical.

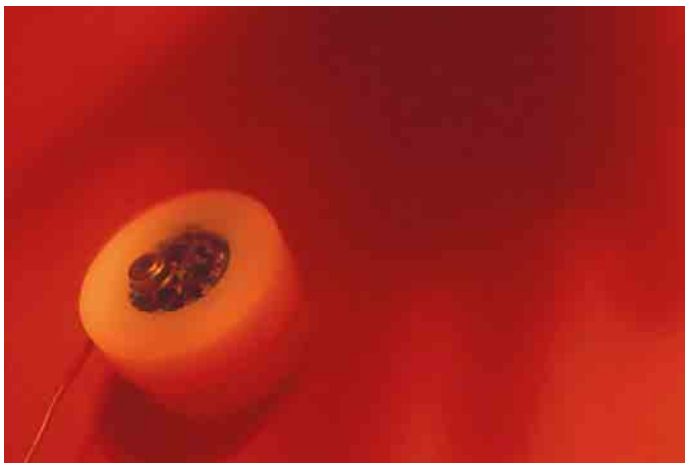

Because sensor networks pose unique challenges, traditional security techniques used in traditional networks cannot be applied directly. First, to make sensor networks economically viable, sensor devices are limited in their energy, computation, and communication capabilities. Second, unlike traditional networks, sensor nodes are often deployed in accessible areas, presenting the added risk of physical attack. And third, sensor networks interact closely with their physical environments and with people, posing new security problems. Consequently, existing security mechanisms are inadequate, and new ideas are needed. Fortunately, the new problems also inspire new research and represent an opportunity to properly address sensor network security from the start.

Here, we outline security issues in these networks, discuss the state of the art in sensor network security, and suggest future directions for research. We cover several important security challenges, including key establishment, secrecy, authentication, privacy, robustness to denial-of-service attacks, secure routing, and node capture. We also cover several high-level security services required for wireless sensor networks and conclude with future research challenges.

A Secure System

Security is sometimes viewed as a standalone component of a system's architecture, where a separate module provides security. This separation is, however,

CAREN ROSENBLATT

usually a flawed approach to network security. 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 must pervade every aspect of system design.

Key establishment and trust setup. When setting up a sensor network, one of the first requirements is to establish cryptographic keys for later use. Researchers have proposed a variety of protocols over several decades for this well-studied problem. Why can't the same key-establishment protocols be used in sensor networks? The inherent properties of sensor networks render previous protocols impractical. Many current sensor devices have limited computational power, making public-key cryptographic primitives too expensive in terms of system overhead. Key-establishment techniques need to scale to networks with hundreds or thousands of nodes. Moreover, 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 a networkwide shared key. Unfortunately, the compromise of even a single node in a network would reveal the secret key and thus allow decryption of all network traffic. 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, then erase the networkwide key after setting up the session keys. However, this variant of the key-establishment process does not allow addition of new nodes after initial deployment.

Public-key cryptography (such as Diffie-Hellman key establishment) is another option 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.

Yet another approach is to preconfigure the network with a shared unique symmetric key between each pair of nodes, though it doesn't scale well. 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 option. Here, each node needs to share only a single key with the base station and set up keys with other nodes through the base station [6]. This arrangement makes the base station a single point of failure, but because there is only one base station, the network may incorporate tamper-resistant packaging for the base station, ameliorating the threat of physical attack.

Researchers recently developed random-key predistribution protocols [3] in which 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 do, they use it to establish a session key. Not every pair of nodes shares a common key, but if the key-establishment probability is sufficiently great, nodes can still set up keys with sufficiently many nodes to obtain a fully connected network. This means of establishing keys avoids having to include a central trusted base station. The disadvantage of this approach is that attackers who compromised sufficiently many nodes could also reconstruct the complete key pool and break the scheme.

In the future, we expect to see research on better random-key predistribution schemes providing 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 allowing simple key establishment for large-scale sensor networks.

Secrecy and authentication. Like traditional networks, most sensor network applications require protection against eavesdropping, injection, and modification of packets. Cryptography is the standard defense. Interesting system trade-offs arise when incorporating cryptography into sensor networks. For point-to-point communication, end-to-end cryptography achieves a high level of security but requires that keys be set up among all end points and be incompatible with passive participation and local broadcast. Link-layer cryptography with a networkwide shared key simplifies key setup and supports passive participation and local broadcast, but intermediate nodes might eavesdrop or alter messages.

The earliest sensor networks are likely to use link-layer cryptography, because this approach provides the greatest ease of deployment among currently available network cryptographic approaches. Subsequent systems may respond to demand for more security with yet more sophisticated use of cryptography.

Cryptography entails a performance cost for extra computation that often increases packet size. Cryptographic hardware support increases efficiency but also increases the financial cost of implementing a network. Therefore, an important question facing sensor node researchers and practitioners is: Can reasonable security and performance levels be achieved with software-only cryptographic implementations, or is hardware support needed?

Recent research demonstrates that software-only cryptography is indeed practical with today's sensor technology; hardware support is not needed to

achieve acceptable security and performance levels. For instance, the University of California, Berkeley, implementation of TinySec incurs only an additional 5%–10% performance overhead using software-only methods. These experiments have also revealed an interesting phenomenon: Most of the performance overhead is attributable to the increase in packet size. In comparison, cryptographic computations have almost no effect on latency or throughput, since they can overlap with transmission. This puts a limit on how much dedicated hardware helps; hardware reduces only the computational costs, not packet size.

Privacy. Sensor networks have also thrust privacy concerns to the forefront. The most obvious risk is that ubiquitous sensor technology might 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 might spy on each other; and law enforcement agencies might spy on public places. This is certainly a valid concern; historically, as surveillance technology has become cheaper and more effective, it has increasingly been implicated in privacy abuses. Technology trends suggest the problem will only get worse with time. As devices get smaller, they will be easier to conceal; as devices get cheaper, surveillance networks will be more affordable.

Another risk is that sensor networks initially deployed for legitimate purposes might subsequently be used in unanticipated and even illegal ways. The notion of function creep is universal in the privacy literature. For instance, U.S. Social Security numbers were originally intended for use only by the Social Security program but have gradually come to be used as an all-purpose personal identification number.

The networked nature of sensor networks raises new threats that are qualitatively different from what private citizens worldwide 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 periods of time, with troubling implications.

Technology alone is unlikely to 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 might provide a reasonable guideline for how to build systems that better protect privacy. Providing awareness of the presence of sensor nodes and data acquisition is particularly important. Affected parties aware of the existence, form, and implications of surveillance are more likely to accept the technology. However, our current understanding of privacy in sensor networks is immature, and more research is needed.

> ULTIMATELY, *we need a secure and efficient key-distribution mechanism allowing simple key establishment for large-scale sensor networks.*

Robustness to communication denial of service. Adversaries can severely limit the value of a wireless sensor network through denial-of-service attacks [9]. In its simplest form, an adversary attempts to disrupt the network's operation by broadcasting a high-energy signal. If the transmission is powerful enough, the entire system's communication could be jammed. More sophisticated attacks are also possible; the adversary might inhibit communication by violating the

802.11 medium access control (MAC) protocol by, say, transmitting while a neighbor is also transmitting or by continuously requesting channel access with a request-to-send signal.

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

The networked nature of sensor networks allows new, automated defenses against denial of service. When the jamming affects only a portion of the network, a jamming-resistant network could defeat the attack by detecting the jamming, mapping the affected region, then routing around the jammed area [8]. Further progress in this area will hopefully allow for greater security against denial-of-service attacks.

Secure routing. Routing and data forwarding is an essential service for enabling communication in sensor networks. Unfortunately, current routing protocols suffer from many security vulnerabilities [5]. For example, an attacker might launch denial-of-service attacks on the routing protocol, preventing communication. The simplest attacks involve injecting malicious routing information into the network, resulting in routing inconsistencies. Simple authentication might guard against injection attacks, but some routing protocols are susceptible to replay by the attacker of legitimate routing messages [4].

Routing protocols are particularly susceptible to node-capture attacks. For instance, researchers have analyzed protocols for routing in sensor networks and found all are highly susceptible to node-capture attacks; in every case, the compromise of a single node suffices to take over the entire network or prevent any communication within it [5]. Network researchers would greatly improve sensor networks by devising secure routing protocols that are robust against such attacks.

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

Tamper-resistant packaging may be one defense, but it's expensive, since current technology does not provide a high level of security. Algorithmic solutions to the problem of node capture are preferable.

The challenge is to build networks that operate correctly even when, unbeknownst to us, several nodes have been compromised and thus might behave in an arbitrarily malicious way. A 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 researchers 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 the packets that were received [2].

A second direction for resilience is to gather multiple, redundant views of the environment and cross-check them for consistency. For instance, the network might require three reports of an interesting event before it responds to the event. Meanwhile, when many data values are collected, a histogram may be constructed; extreme outliers may indicate malicious spoofed data and hence should be ignored.

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

Network Security Services

So far, we've explored low-level security primitives for securing sensor networks. Here, we consider high-level security mechanisms, including secure group management, intrusion detection, and secure data aggregation.

Secure group management. Each node in a wireless sensor network is limited in its computing and communication capabilities. However, interesting in-network data aggregation and analysis can be performed by groups of nodes. For example, a group of nodes might be responsible for jointly tracking a vehicle through the 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, securely admitting new group members and supporting secure group communication. The outcome of the group's computation is normally transmitted to a base station. The output must be authenticated to ensure it comes from a valid group. Any solution must also be efficient in terms of time and energy (or

involve low computation and communication costs), precluding many classical group-management solutions.

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 in terms of the network's memory and energy consumption, as well as its inherently limited bandwidth. Wireless sensor networks require a solution that is fully distributed and inexpensive in terms of communication, energy, and memory requirements. In order to look for anomalies, applications and typical threat models must be understood. It is particularly important for researchers and practitioners to understand how cooperating adversaries might attack the system. The use of secure groups may be a promising approach for decentralized intrusion detection.

Secure data aggregation. One benefit of a wireless sensor network is the fine-grain sensing that large and dense sets of nodes can provide. The sensed values must be aggregated to avoid overwhelming amounts of traffic back to the base station. 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, aggregation may take place in many places in the network. All aggregation locations must be secured.

If the application tolerates approximate answers, powerful techniques are available; under appropriate trust assumptions, randomly sampling a small fraction of nodes and checking that they have behaved properly supports detection of many different types of attacks [7].

Research Challenges

The severe constraints and demanding deployment environments of wireless sensor networks make computer security for these systems more challenging than for conventional networks. However, several properties of sensor networks may help address the challenge of building secure networks. First, we have the opportunity to architect security solutions into these systems from the outset, since they are still in their early design and research stages. Second, many applications are likely to involve the deployment of 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 the solutions. If we build sensor networks so they continue operating even if some fraction of their sensors is compromised, we have an opportunity to use redundant sensors to resist further attack. Ultimately, the unique aspects of sensor networks may allow novel defenses not available in conventional networks.

Many other problems also need further research. One is how to secure wireless communication links against eavesdropping, tampering, traffic analysis, and denial of service. Others involve resource constraints. Ongoing directions include asymmetric protocols where most of the computational burden falls on the base station and on public-key cryptosystems efficient on low-end devices. Finally, finding ways to tolerate the lack of physical security, perhaps through redundancy or knowledge about the physical environment, will remain a continuing overall challenge. We are optimistic that much progress will be made on all of them. C

References

1. Adamy, D. *EW 101: A First Course in Electronic Warfare.* Artech House Publishers, Norwood, MA, 2001.
2. Deng, J., Han, R., and Mishra, S. A performance evaluation of intrusion-tolerant routing in wireless sensor networks. In *Proceedings of the 2nd IEEE International Workshop on Information Processing in Sensor Networks (IPSN 2003)* (Apr. 2003), 349–364.
3. Eschenauer, L. and Gligor, V. A key-management scheme for distributed sensor networks. In *Proceedings of the 9th ACM Conference on Computer and Communication Security* (Washington, D.C., Nov.). ACM Press, New York, 2002, 41–47.
4. Hu, Y.-C., Perrig, A., and Johnson, D. Packet leashes: A defense against wormhole attacks in wireless ad hoc networks. In *Proceedings of IEEE Infocom 2003* (San Francisco, Apr. 1–3, 2003).
5. Karlof, C. and Wagner, D. Secure routing in wireless sensor networks: Attacks and countermeasures. In *Proceedings of the 1st IEEE International Workshop on Sensor Network Protocols and Applications* (Anchorage, AK, May 11, 2003).
6. Perrig, A., Szewczyk, R., Wen, V., Culler, D., and Tygar, J. SPINS: Security protocols for sensor networks. *J. Wireless Nets. 8,* 5 (Sept. 2002), 521–534.
7. Przydatek, B., Song, D., and Perrig, A. SIA: Secure information aggregation in sensor networks. In *Proceedings of the 1st ACM International Conference on Embedded Networked Sensor Systems (SenSys 2003)* (Los Angeles, Nov. 5–7). ACM Press, New York, 2003, 255–265.
8. Wood, A., Stankovic, J., and Son, S. JAM: A mapping service for jammed regions in sensor networks. In *Proceedings of the IEEE Real-Time Systems Symposium* (Cancun, Mexico, Dec. 3–5, 2003).
9. Wood, A. and Stankovic, J. Denial of service in sensor networks. *IEEE Comput.* (Oct. 2002), 54–62.

Adrian Perrig (perrig@cmu.edu) is an assistant professor of electrical and computer engineering and engineering and public policy at Carnegie Mellon University in Pittsburgh, PA.
John A. Stankovic (stankovic@cs.virginia.edu) is the BP America Professor and Chair of the Computer Science Department at the University of Virginia in Charlottesville, VA.
David Wagner (daw@cs.berkeley.edu) is an assistant professor of electrical engineering and computer science at the University of California, Berkeley.