



July 2006

Sensor Network Security: More Interesting Than You Think

Madhukar Anand

University of Pennsylvania, anandm@cis.upenn.edu

Eric Cronin

University of Pennsylvania, ecronin@cis.upenn.edu

Micah Sherr

University of Pennsylvania, msherr@cis.upenn.edu

Matthew A. Blaze

University of Pennsylvania, blaze@cis.upenn.edu

Zachary G. Ives

University of Pennsylvania, zives@cis.upenn.edu

See next page for additional authors

Follow this and additional works at: http://repository.upenn.edu/cis_papers

Recommended Citation

Madhukar Anand, Eric Cronin, Micah Sherr, Matthew A. Blaze, Zachary G. Ives, and Insup Lee, "Sensor Network Security: More Interesting Than You Think", . July 2006.

Copyright 2006 USENIX Association. Reprinted from *1st USENIX Workshop on Hot Topics in Security (HotSec 2006)*.

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/cis_papers/249
For more information, please contact libraryrepository@pobox.upenn.edu.

Sensor Network Security: More Interesting Than You Think

Abstract

With the advent of low-power wireless sensor networks, a wealth of new applications at the interface of the real and digital worlds is emerging. A distributed computing platform that can measure properties of the real world, formulate intelligent inferences, and instrument responses, requires strong foundations in distributed computing, artificial intelligence, databases, control theory, and security.

Before these intelligent systems can be deployed in critical infrastructures such as emergency rooms and powerplants, the security properties of sensors must be fully understood. Existing wisdom has been to apply the traditional security models and techniques to sensor networks. However, sensor networks are not traditional computing devices, and as a result, existing security models and methods are ill suited. In this position paper, we take the first steps towards producing a comprehensive security model that is tailored for sensor networks. Incorporating work from Internet security, ubiquitous computing, and distributed systems, we outline security properties that must be considered when designing a secure sensor network. We propose challenges for sensor networks – security obstacles that, when overcome, will move us closer to decreasing the divide between computers and the physical world.

Comments

Copyright 2006 USENIX Association. Reprinted from *1st USENIX Workshop on Hot Topics in Security* (HotSec 2006).

Author(s)

Madhukar Anand, Eric Cronin, Micah Sherr, Matthew A. Blaze, Zachary G. Ives, and Insup Lee

Sensor Network Security: More Interesting Than You Think *

Madhukar Anand, Eric Cronin, Micah Sherr, Matt Blaze, Zachary Ives, and Insup Lee
Department of Computer and Information Science
University of Pennsylvania

{anandm, ecronin, msherr, blaze, zives, lee}@cis.upenn.edu

Abstract

With the advent of low-power wireless sensor networks, a wealth of new applications at the interface of the real and digital worlds is emerging. A distributed computing platform that can measure properties of the real world, formulate intelligent inferences, and instrument responses, requires strong foundations in distributed computing, artificial intelligence, databases, control theory, and security.

Before these intelligent systems can be deployed in critical infrastructures such as emergency rooms and powerplants, the security properties of sensors must be fully understood. Existing wisdom has been to apply the traditional security models and techniques to sensor networks. However, sensor networks are not traditional computing devices, and as a result, existing security models and methods are ill suited. In this position paper, we take the first steps towards producing a comprehensive security model that is tailored for sensor networks. Incorporating work from Internet security, ubiquitous computing, and distributed systems, we outline security properties that must be considered when designing a secure sensor network. We propose challenges for sensor networks – security obstacles that, when overcome, will move us closer to decreasing the divide between computers and the physical world.

1 Introduction

The advent of low-powered wireless networks of embedded sensors [HSW⁺00, MFHH03, ABC⁺04] has spurred the development of new applications at the interface between the real world and its digital manifestation. A distributed computing platform that can measure properties of the real world, formulate intelligent inferences, and instrument responses, requires a new class of techniques in distributed computing, artificial intelligence, databases, control theory, and (the focus of this position paper) security.

Before these intelligent systems can be deployed in critical infrastructures such as emergency rooms and power plants, the security properties of sensors must be fully understood. Existing wisdom has been to apply the traditional security models and techniques to sensor networks: as in conventional computing environments, the goal has been to protect *physical* entities: devices, packets, links, and ultimately networks.

However, sensor networks are not traditional computing devices, and as a result, existing security models and methods are insufficient. Sensors have unique characteristics that warrant novel security considerations: the geographic distribution of the devices allows an attacker to physically capture nodes and learn secret key material, or to intercept or inject messages; the hierarchical nature of sensor networks and their route maintenance protocols permit the attacker to determine where the root node is placed. Perhaps most importantly, most sensor networks rely on *redundancy* (followed by aggregation) to accurately capture environmental information even with poorly calibrated and unreliable devices. This results in a fundamental distinction between a *physical message* in a sensor network and a *logical unit* of sensed information: a message with a single sensor reading may reveal very little information about the real environment, whereas a message containing an aggregate or collection of readings may reveal a great deal more.

These characteristics open the door for an entirely new security paradigm: one that acknowledges that there is a fundamental distinction between physical messages and logical information, and that focuses on how to minimize the correlation between the two in order to limit opportunities for compromise. In this position paper, we take the first steps towards producing a comprehensive security model that is tailored for these low-powered distributed devices. We begin with a discussion of the unique properties of sensor networks, and then introduce an attack model that addresses these unique properties. Incorporating work from Internet security, ubiquitous computing, and distributed systems, we outline security properties that must be considered when designing a secure sensor network. Finally, we propose challenges for sensor networks – security obstacles that, when overcome, move us closer to decreasing the divide between com-

*This research was supported in part by the following grants: ONR MURI N00014-04-1-0735; NSF CNS-0509327, IIS-0477972, IIS-0513778; ARO W911NF-05-1-0182; and DARPA HR0011-06-1-0016.

puters and the physical world.

2 Attacker Goals for Sensor Networks

In traditional networks such as the Internet, attackers target physical systems and packets, and this is reflected in today's common security techniques and practices. In contrast, the redundancy and aggregation intrinsic to sensor networks limit the systemwide impact of attacks against individual nodes: sensor devices themselves are dispensable and vary in their impact on the network. To discern useful information or to accomplish a change in network output, a sensor network attacker must carefully target his attack to those devices with the most influence.

However, the potentially hostile environment in which sensors are located also introduces new challenges in defending the network, e.g., sensor devices may be physically captured, and nodes near the root of the sensor network are of high value if captured or compromised. It is therefore useful to establish a threat model that considers the unique properties of sensor networks. We briefly enumerate three basic categories of attacks based on our earlier work [AIL05]:

1. *Eavesdropping.* The adversary (eavesdropper) seeks to determine what data is being output by the sensor network. The adversary either listens to messages transmitted by the nodes, or directly compromises nodes. Eavesdropping may take two forms. A *passive* eavesdropper conceals her presence from the sensor nodes. She passively intercepts messages. An *active* eavesdropper sends queries to sensors or aggregation points, or attacks sensor nodes, in order to gain more information.

In either passive or active eavesdropping, the adversary's goal is to ascertain logical information about the sensed environment. Because individual sensor readings vary in their level of contribution to an aggregate value, the eavesdropper's location in the sensor network determines the amount of information that she can accurately obtain. This differs significantly from traditional eavesdropping threat models, where although data may be distributed there is no redundancy or aggregation to be considered.

2. *Disruption.* The adversary aims to disrupt the sensor application. To be most effective, the adversary must direct her attack against locations in the sensor network that significantly influence the logical output of the network. She can conduct a disruption attack using a combination of two techniques. *Semantic* disruption injects messages, corrupts data, or changes values in order to render the aggregated

data corrupt, useless, or incomplete. *Physical* disruption upsets sensor readings by directly manipulating the environment, e.g., by generating heat in the vicinity of temperature sensors.

3. *Hijacking.* The adversary subverts the sensor application output by gaining control over sensors. By hijacking a carefully chosen set of sensors, both eavesdropping and disruption attacks can be accomplished from within the sensor network. These attacks are hardest to counter since they come from trusted nodes.

This is not the first attack model on sensor security (e.g., [WS02, KW03]), but it is unique in two ways. First, the organization of this taxonomy is a classification based on adversary's *goals*, not on particular methods. Second, the focus is on the overall logical output of the network, assuming that compromise of individual nodes is a certainty.

Many sensor networks do not just measure their environment, but also interact with it through actuators. When sensors are coupled with actuator devices, care must be taken that disruption attacks cannot also be mounted against the actuators (a potentially catastrophic attack in medical or defense applications). For example, even if an attacker is unable to read or inject messages into the sensor network, they may still be able to disable nodes by exhausting their batteries with bogus queries [Sta02]. Even though the sensor/actuator is able to discard these requests, it must expend energy to process them.

3 Unique Properties of Sensor Networks

The sensor network domain is characterized by large numbers of limited-computation, often unreliable and low-powered devices embedded within an environment. As a result, sensor networks exhibit unique properties not present in more traditional network configurations. We briefly recap the chief distinctions that lead to new challenges and opportunities in security, and give each a label that we will later reference.

- P1: **Tree-structured routing** is the basis of most current sensor networks (e.g., [MFHH03]), with the base station at the root. While recent work [NGSA04] has begun to consider DAG-structured networks with redundant transmission of values, such approaches are limited in the functions they can compute (since complex schemes must be used to avoid double-counting readings).
- P2: **Aggregation** is used not only to monitor conditions across a wide area of coverage, but also to compen-

sate for unreliability, miscalibration of sensor devices, and intermittent connectivity.

- P3: **Tolerable failures:** the critical component in sensor networks is the sensed data, not the physical devices. Sensors are typically low-cost devices, and the loss or corruption of a sensor can either be mitigated by redundant sensors or tolerated by the network. This sharply contrasts with services on the Internet, in which the compromise of a host is often catastrophic. The redundancy of sensors and tolerance for a limited quantity of noisy (or malicious) data makes individual sensor nodes less critical.
- P4: **In-network filtering and computation** allows work (especially aggregation and computation) to be “pushed” as close as possible to the devices that originate specific sensor readings. This enables greater power efficiency, since fewer data packets must be transmitted.
- P5: **Sensors as routers:** in a typical sensor network, there is no distinction between sensing nodes, compute nodes, and routing nodes. This, combined with the characteristics described above, reduces network traffic.
- P6: **Phased transmission periods** are an integral component of most sensor network routing protocols (even, in many cases, those that use CDMA or other techniques for avoiding collisions): within a sensor network epoch, each node has a phase in which it senses, a phase in which it receives messages from its children, and a phase in which it forwards its (filtered or aggregated) data to its parent¹. This approach allows each device to deactivate its radio for a significant portion of each epoch.

These sensor properties lead to a number of constraints and characteristics that have security implications. Below, we consider the impact of these features on sensor network security.

4 Sensor Network Security Challenges

To protect against the attacks outlined above, system designers must be cognizant of the security properties that accompany sensor networks. Some of these properties, such as tolerable failures (Property P1) present opportunities for designing protocols for sensor networks that are infeasible in other types of networks. Below, we take a first step towards establishing a comprehensive set of security challenges for sensor networks. Some challenges are similar to those faced in more traditional environments, but with additional constraints; others are unique

¹Sometimes one or more of these time phases may be combined.

to sensor networks and similar technologies (e.g., mobile ad hoc networks [Sta02]). When steps have already been made towards a challenge, we place the related work in context.

Challenge 1: Measuring Confidentiality

Existing literature has proposed the use of computationally inexpensive cryptographic techniques to handle message confidentiality and authenticity in sensor networks [AUJP03, PSW⁺01]. The difficulty of ensuring confidentiality and authenticity is not, however, due solely to the energy constraints imposed on sensors. A sensor network is comprised of many small computing devices, each of which is subject to physical capture. Any cryptosystem must therefore tolerate the compromise of sensors and their keys. New cryptographic approaches must be developed that are geared towards this failure model.

However, the compromise of some nodes need not result in a total loss of security. Unlike traditional networks in which logical information is often conveyed as single messages or packets, sensor networks rely on redundancy and aggregation (Properties P1, P2), and therefore some messages may be more influential than others. In an earlier paper [AIL05], we presented an initial framework for *quantifying* the privacy and security of sensor network applications under the assumption that some nodes may be compromised. Rather than providing all-or-nothing guarantees about privacy or security, we examined *probabilistic guarantees* with respect to compromise. **Challenge 1** is to define models and metrics along these lines, for different protocols’ logical-level information privacy and security properties.

Challenge 2: Timing Obfuscation

For a sensor value to have meaning, context is needed. Where the value was recorded, and at what time, are necessary for interpretation. Conversely, if the time and location of one reading are known, it may be possible for an adversary to infer a great deal about other readings nearby (Properties P5, P6). Sensor networks must therefore be aware of these metadata and their role in security.

It may be possible for an eavesdropper to correlate public data to infer confidential information. Deshpande et al have proposed incorporating a probabilistic model for data aggregation in a sensor network [DGM⁺04]. By exploiting the correlation between different values and between different attributes, they report significant energy savings in query processing. Such a model also implies that an adversary could pose innocuous-looking queries on certain attributes to obtain confidential data.

The timing of sensor messages may also reveal confidential data. In applications where anonymity is desired (see Challenge 6), we might limit the ability of an

eavesdropper (or even the aggregating node) to infer the identity of the sensor node. **Challenge 2** is to identify cost-effective schemes for hiding sensor network timing. Possible solutions might be based on sending messages at regular intervals, disassociating a reading from a physical event by adding a random delay to message transmission, or adding spurious messages to mask the legitimate send times.²

Challenge 3: Secure Aggregation

In sensor networks where aggregation occurs at intermediary nodes, end-to-end encryption from sensors to the base station is not possible because each node must be able to compute with the data. Although cryptosystems have been proposed that allow computation on ciphertexts [GHY87], such approaches require significant computational cost and may be infeasible in low powered devices. The standard security doctrine that the network should not be trusted and that all messages should be encrypted and decrypted at the source and destination is incompatible with aggregation (due to Property P4). Unfortunately, the alternative of trusting each link between the sensor and the base station is unappealing. **Challenge 3** is to develop novel cryptographic approaches that allow the aggregation of messages while ensuring adequate security.

An alternative to employing secure techniques to collect data is to use more robust statistical aggregation functions. Common aggregation functions such as average, sum, minimum/maximum are not resilient and are vulnerable to easy attacks [Wag04]. On the other hand, count, median and root mean squared error are better estimators of the data being aggregated as they are more robust.

Challenge 4: Topology Obfuscation

Unlike traditional networks, where intermediate nodes in the routing tree simply relay messages, nodes in sensor networks often carry out computation on messages before passing them along (Property P3). This computation leads to a non-uniform distribution of information across nodes: different nodes carry differing amounts of influence on the final computed value. Attacking a leaf node in a tree-structured network gains little influence (for disruption) or information (for eavesdropping); attacking a node near the root gains significant influence and information about the aggregate value (Property P1). For eavesdropping, there is an interesting third case of attacking nodes in the middle of the tree: intermediary nodes perform enough aggregation to compensate for inaccurate sensors, but their values may be local enough

²Masking timing information does not necessarily imply that aggregation cannot be performed on the data. Aggregation is performed on data that have the same *logical* timestamp whereas hiding the timing interferes with the ability to discern *physical* time.

to reveal private data (see Challenge 6). **Challenge 4** is to hide the routing infrastructure of the sensor network. If an adversary can attack a few chosen nodes, the obvious strategy is to compromise sensors (and their keys) that logically reside in high value locations in the routing tree.

Challenge 5: Scalable Trust Management

In the domain of sensor networks, trust management is the problem of identifying which nodes are legitimate and which are not to be trusted. The threat of physical compromise (and need to revoke trust when detected), the energy constraints, the number of nodes which must be considered, and the difficulty in re-establishing trust once sensors are deployed are all unique challenges to trust management in sensor networks.

Due to the power and energy constraints of many of the nodes, it may not be possible to run expensive key generation algorithms, or to run them pairwise between every node. Even if this is feasible once, it may not be practical to run them frequently. Since there is the assumption that the physical compromise of some nodes (and therefore their shared keys) is unavoidable, limitations must be placed on the number of nodes sharing keys to limit the impact of compromise.

Key management is one of the better studied areas of sensor network security, but many of the proposed approaches are practical only under certain conditions. **Challenge 5** is to develop “lightweight” key management and distribution schemes appropriate for large-scale sensor networks. Due to space constraints, it is impossible to enumerate all the proposed key management systems in this paper, but the reader is referred to [WLSC].

Challenge 6: Aggregation with Privacy

The interaction between sensors and the physical world leads to new challenges in privacy and anonymity for those being sensed. Unlike traditional computing platforms, end users who are identified by sensor nodes have little ability to set policy. When browsing the Internet, for example, users can use anonymizing proxies to protect their privacy. When being sensed by a sensor, however, the end user has no input as to the level of information disclosure, and must trust in the decisions made by the sensor network. Since being sensed can be a passive act and can be done without the knowledge of the observed party, designing networks with privacy guarantees is an arduous task.

Anonymity may be desired in some sensor network applications. If the objective is to be anonymous with respect to an external observer, then techniques such as Onion Routing [DMS04] could be extended to achieve anonymity. However, onion routing may be expensive

here, and in some cases, it may be desirable to protect individual readings while still computing the aggregate over all readings. **Challenge 6** is to develop new anonymity techniques to handle such requirements.

Illustrative Example Applications

In this section, we present example applications to illustrate the challenges that we have introduced. Our first example is the next generation Supervisory Control And Data Acquisition (SCADA) system. Currently, the system consists of a central controller and a distributed network of Remote Terminal Units (RTU) or Programmable Logic Controllers (PLC). Data Acquisition in the SCADA system begins at the RTU or PLC which collect data such as meter readings and equipment status and communicate it to the central controller where a supervisory decision is made using a human-machine interface. With maturing wireless sensor network technology, it is envisaged that the network of RTU and PLCs will be replaced by devices such as the wireless sensor nodes [SCA]. Sensor networks could be deployed to monitor and protect power grids, transportation, water and fuel infrastructure. In such a system, it is critical to ensure that the readings collected be robust (Challenge 3) and the degree of robustness be quantified so that appropriate degree of control can be exercised (Challenge 1). By hiding the timing information, we can hide the state of the system (Challenge 2). This helps prevent the adversary from knowing what information is being acquired (Challenge 4). In the SCADA network, each sensor will be assumed to be active for a certain lifetime. The lifetime will be estimated using a probabilistic model of network activity and the resources at each node. With such a model, it would be possible to define the coverage offered by a sensor node and therefore, to devise replenishment strategies to replace dead sensors [Wic]. Given a large number of sensors, some of which are periodically replaced, management of encryption keys can be quite difficult; thus it becomes necessary to develop trust management solutions that are lightweight and scale to a large number of sensors (Challenge 5). Such a scheme must also permit addition and removal of sensor nodes.

Many sensor network applications involve collecting personally identifiable information (PII) [Wic], such as (1) sensing persons in buildings as part of embedded sensors for disaster preparedness or power savings, (2) monitoring activities of the elderly so they can safely live at home, (3) monitoring automobiles' FastTRAK on the highway transponders in automobiles. In such applications, in addition to challenges 1-5, there is also a need to protect the privacy and in some cases, ensure anonymity (Challenge 6).

5 Conclusions and Research Agenda

Existing literature on sensor network security has largely applied the Internet security model to sensor networks. Prior work tends to concentrate exclusively on the low-power aspect of sensor networks, often neglecting these other unique properties that further distinguish them from more traditional computing systems.

Although there are some similarities, sensor network topologies and functions introduce a range of considerations different from those found of the Internet. These unique characteristics, e.g., tree-structured routing, aggregation, in-network filtering, etc., have important security implications. This position paper proposes a more appropriate attack taxonomy and looks at how the security model must be tailored for sensor networks. By more carefully considering the threats posed to sensor networks, applications with intrinsic security considerations become immediately realizable. We conclude by summarizing the list of security challenges for sensor networks.

- **Challenge 1 [Measuring Confidentiality]** : is to define models and metrics for information privacy and security properties of sensor network protocols.
- **Challenge 2 [Timing Obfuscation]**: is to identify cost-effective schemes for hiding the timing information in sensor networks.
- **Challenge 3 [Secure Aggregation]**: is to develop novel cryptographic solutions that allow aggregation of messages while ensuring adequate security.
- **Challenge 4 [Topology Obfuscation]**: is to hide the routing infrastructure so as to offset the non-uniform node information in a sensor network.
- **Challenge 5 [Scalable Trust Management]**: is to develop "lightweight" key management and distribution schemes appropriate for large-scale sensor networks.
- **Challenge 6 [Aggregation with Privacy]**: is to develop new techniques to handle the privacy and anonymity while ensuring meaningful aggregation of sensor data.

References

- [ABC⁺04] T. Abdelzaher, B. Blum, Q. Cao, D. Evans, J. George, S. George, T. He, L. Luo, S. Son, R. Stoleru, J. Stankovic, and A. Wood. Envirotrack: Towards an environmental computing paradigm for distributed sensor networks. In *IEEE International Conference on Distributed Computing Systems*, March 2004.

- [AIL05] Madhukar Anand, Zachary Ives, and Insup Lee. Quantifying eavesdropping vulnerability in sensor networks. In *DMSN '05: Proceedings of the 2nd international workshop on Data management for sensor networks*, pages 3–9, New York, NY, USA, 2005. ACM Press.
- [AUJP03] Sasikanth Avancha, Jeffrey L Undercoffer, Anupam Joshi, and John Pinkston. Secure sensor networks for perimeter protection. *Computer Networks*, 43(4):421–435, November 2003.
- [DGM⁺04] Amol Deshpande, Carlues Guestrin, Samuel Madden, Joseph M. Hellerstein, and Wei Hong. Model-driven data acquisition in sensor networks. In *VLDB '04*, 2004.
- [DMS04] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The Second-Generation Onion Router. In *Proc. of the 13th USENIX Security Symposium*, pages 303–320, Aug 2004.
- [GHY87] Zvi Galil, Stuart Haber, and Moti Yung. Cryptographic computation: Secure fault-tolerant protocols and the public-key model. *LNCS: A Conference on the Theory and Applications of Cryptographic Techniques on Advances in Cryptology*, 293:135–155, 1987.
- [HSW⁺00] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister. System architecture directions for network sensors. In *ASPLOS*, November 2000.
- [KW03] Chris Karlof and David Wagner. Secure routing in wireless sensor networks: Attacks and countermeasures. *Elsevier's Ad Hoc Networks Journal, Special Issue on Sensor Network Applications and Protocols*, 1(2-3):293–315, May 2003.
- [MFHH03] Samuel Madden, Michael J. Franklin, Joseph M. Hellerstein, and Wei Hong. Design of an acquisitional query processor for sensor networks. In *SIGMOD '03*, pages 491–502, 2003.
- [NGSA04] Suman Nath, Phillip B. Gibbons, Srinivasan Seshan, and Zachary R. Anderson. Synopsis diffusion for robust aggregation in sensor networks. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 250–262, New York, NY, USA, 2004. ACM Press.
- [PSW⁺01] Adrian Perrig, Robert Szewczyk, Victor Wen, David E. Culler, and J. D. Tygar. SPINS: security protocols for sensor networks. In *Mobile Computing and Networking*, pages 189–199, 2001.
- [SCA] Beyond SCADA planning meeting. <http://trust.eecs.berkeley.edu/scada/wiki/Scada/Main>.
- [Sta02] Frank Stajano. *Security for Ubiquitous Computing*. John Wiley and Sons, February 2002.
- [Wag04] David Wagner. Resilient aggregation in sensor networks. In *SASN '04: Proceedings of the 2nd ACM workshop on Security of ad hoc and sensor networks*, pages 78–87, New York, NY, USA, 2004. ACM Press.
- [Wic] Steve Wicker. Privacy and security: Technology & challenges. <http://robotics.eecs.berkeley.edu/~sinopoli/SCADA/wicker.ppt>.
- [WLSC] John Paul Walters, Zhengqiang Liang, Weisong Shi, and Vipin Chaudhary. Wireless sensor network security: A survey. <http://www.cs.wayne.edu/~weisong/papers/walters05-wsn-security-survey.pdf>.
- [WS02] Anthony D. Wood and John A. Stankovic. Denial of service in sensor networks. *IEEE Computer*, 35(10):54–62, October 2002.