

Security Challenges in Next Generation Cyber Physical Systems

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

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

1 Introduction

The advent of low-powered wireless networks of embedded sensors has spurred the development of new applications at the interface between the real world and its digital manifestation. Following this trend, the next generation Supervisory Control And Data Acquisition (SCADA) system is expected to replace traditional data gathering – a distributed network of Remote Terminal Units (RTU) or Programmable Logic Controllers (PLC), with devices such as the wireless sensing devices.

Before these intelligent systems can be deployed in critical infrastructure 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.

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 enlist challenges for sensor networks – security obstacles that, when overcome, move us closer to deploying them in large numbers for monitoring and protecting critical infrastructures.

2 Security Challenges and Research Agenda

2.1 Measuring Confidentiality

Existing literature has proposed the use of computationally inexpensive cryptographic techniques to handle message confidentiality and authenticity in sensor networks. 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. However, the compromise of some nodes need not result in a total loss of security. Rather than providing all-or-nothing guarantees about privacy or security, there is a need for providing *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. One possible approach to this problem could be to quantify the threat under the assumption that some nodes may be compromised [1].

2.2 Context 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. Sensor networks must therefore be aware of these metadata and their role in security. **Challenge 2** is therefore 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.

2.3 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. 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. 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.

2.4 Topology Obfuscation

Aggregation of data by the intermediate nodes leads to a non-uniform distribution of information across nodes. Therefore, 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. 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 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.

2.5 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. **Challenge 5** is therefore 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 [2].

2.6 Aggregation with Privacy

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.

Challenge 6 is to develop new anonymity techniques to handle such requirements.

Illustrative Example Application Scenarios

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). Preventing the adversary from learning the state of the network would involve hiding the timing and topology (Challenges 2, 4). In a large-scale 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 [3]. 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 lightweight, secure trust management solutions (Challenge 5) that permits addition and

removal of sensor nodes.

Many sensor network applications involve collecting personally identifiable information (PII) [3], 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).

3 Conclusion

Existing systems and methodologies have largely applied the internet model of security to cyber physical systems. However, these approaches fail to make the distinction between physical messages and logical information, focus on all-or-nothing security guarantees, and are ill-equipped to deal with the inherent asymmetry in the distribution of information across the network. In this position paper we have proposed a more appropriate security model and identified 6 research challenges for building secure cyber physical systems. We believe that once these challenges are surmounted, applications with intrinsic security considerations will become immediately realizable.

References

- [1] 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.
- [2] 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>.
- [3] Steve Wicker. Privacy and security: Technology & challenges. <http://robotics.eecs.berkeley.edu/~sinopoli/SCADA/wicker.ppt>.

Madhukar Anand is a PhD candidate in Computer and Information Sciences at the University of Pennsylvania. His research interests include embedded systems, real-time computing, formal methods, and sensor networks. Anand has an MSE in computer and information science from the University of Pennsylvania and a MS in Mathematics and Computing from Indian Institute of Technology, Kharagpur. He is a member of the ACM and SIGBED. Contact him at anandm@cis.upenn.edu, Tel: 215 746 3160

Eric Cronin is a PhD candidate in Computer and Information Sciences at the University of Pennsylvania. His research interests include network security, privacy, and distributed systems. Cronin has an MS in computer science and engineering from the University of Michigan. He is a member of ACM, IEEE, and USENIX. Contact him at ecronin@cis.upenn.edu, Tel: 215 746 3160

Micah Sherr is a PhD candidate in Computer and Information Sciences at the University of Pennsylvania. His research interests include network security, protocol design and analysis, network intrusion detection and prevention, and privacy and data confidentiality. Sherr has a MSE in computer and information science from the University of Pennsylvania. He is a member of IEEE and USENIX. Contact him at msherr@cis.upenn.edu, Tel: 215 746 3160

Matt Blaze is an associate professor of computer and information sciences and director of the Trusted Network Eavesdropping and Countermeasures project at the University of Pennsylvania. His research interests include secure systems, cryptology and cryptographic protocols, and large-scale systems. Blaze has a PhD in computer science from Princeton University. He is a member of ACM, IACR and IEEE, and is a director of the USENIX association. Contact him at mab@crypto.com.

Zachary Ives is an assistant professor of computer and information sciences at the University of Pennsylvania. His research interests

lie in the areas of databases and distributed systems, especially as they relate to the Web, Web-scale information sharing, and distributed networks of devices. Ives has a PhD in computer science from the University of Washington. He is a member of ACM, and the SIGMOD. Contact him at zives@cis.upenn.edu, Tel: 215 746 2789.

Insup Lee is the Cecilia Fidler Moore Professor in the Department of Computer and Information Science at the University of Pennsylvania, where he has been since 1983. His research interests include embedded systems, real-time computing, formal methods, wireless network, and software engineering. He is IEEE fellow and IEEE Distinguished Visitor Speaker. He was Chair of IEEE Computer Society Technical Committee on Real-Time Systems (TCRTS) during 2003-2004. He has served on numerous program committees and chaired several international conferences and workshops, and on various steering and advisory committees of technical societies, including Steering Committee of ACM SIGBED. He has also served on the editorial boards on the several scientific journals. Contact him at lee@cis.upenn.edu, Tel: 215 898 3532.