Security for the Mythical Air-dropped Sensor Network

Chandana Gamage, Kemal Bicakci, Bruno Crispo, and Andrew S. Tanenbaum
Department of Computer Science, Vrije Universiteit, Amsterdam, The Netherlands
{chandag,kemal,crispo,ast}@cs.vu.nl

Abstract

The research area of very large scale wireless sensor networks made of low-cost sensors is gaining a lot of interest as witnessed by the large number of published papers. The security aspects of such networks are addressed as well, and in particular many security papers investigating the security aspects of such networks make important assumptions about the capabilities of low-cost sensors. Consequently, the techniques proposed in the current literature to provide security properties for this low-cost wireless sensor networks are heavily shaped by such assumptions. In this position paper, we challenge such assumptions by presenting the results of an experiment we conducted using sensors representative of low cost units. And we show that the same security properties can be better provided using techniques based on application-specific knowledge, heuristics and statistical tests. Finally, we show that one of the most highly cited application scenarios to motivate such techniques, the air-dropped sensor network, is likely to be more a myth than a realistic scenario for low-cost sensors.

Keywords*: Sensor networks, security, practical issues*

1. Introduction

The low-cost wireless sensor technology has been described as a viable mechanism to rapidly build and deploy very large sensor networks consisting of many thousands of nodes for civilian applications such as taking environmental pollution measurements and for military applications such as tracking enemy movements over a harsh and inaccessible terrain. The main advantages of these proposed sensor networks are (1) the low-cost of sensors making it economical to construct higher node density networks to give greater reliability and coverage, (2) wireless communication and battery powered operation obviating the need for expensive infrastructure, and (3) maintenance-free operation of the network.

All references to sensor networks in this paper mean those constructed using low-cost sensors such as Crossbow Mica2 and Mica2Dot [1, 2]. This assumption is crucial to understand the research positions taken in this paper because *if* the low-cost sensors become more powerful and sophisticated, most of our arguments become less relevant. However, we also address this time vs. technology improvement aspect related to low-cost sensors.

A large amount of sensor network security research work has been done on making such low-cost sensor networks (1) suitable for *ad-hoc construction* through localization where physical location of nodes are securely self computed and notified to the base station, (2) *secure* against attacks on the sensors, on transmitted messages, and on its routing schemes, and (3) *energy efficient* using in-network processing for secure data aggregation. We review these techniques in the context of large sensor networks built using low-cost sensors and argue if they are needed and useful. We will support our arguments with several practical example applications that need security against motivated attackers and through data gathered from an experiment described in section 2.

In section 3, we discuss the sensor localization in the context of a large static sensor network. Our main argument is that if sensors are deployed in the field, it is not possible both for these sensors to function correctly *and* for the operators of the network to be unaware of the physical locations of sensors. For civilian applications, the sensor network will be installed by hand on the ground or if the sensors are dropped from an unmanned aerial vehicle (UAV) , that will be from a low altitude at a low speed allowing a GPS on board the UAV to record the drop locations and serial numbers of each sensor. For military applications in which sensors may have to be dropped under unfavorable operating conditions, we show in section 2 through experimental data that the density of the sensors on the ground will be so high as to be unsuitable for such deployment.

In section 4, we discuss the important issue of key management in a large sensor network. The main body of research work in this area for low-cost sensors has been in pair-wise key pre-distribution and this method of managing symmetric keys for sensor network is due to the random placement of sensors *and* the small amount of memory available on a sensor. If the neighbor nodes of a given sensor are not known a priori, then it is not possible to configure that sensor with the small number of keys it needs to share with its neighbors for secure inter-node communication. The simple scheme of sharing a single key among all sensors does not provide any meaningful security as the capture of a single sensor makes the entire network vulnerable. The more robust scheme of each sensor sharing a unique key with every other sensor, so that capture of a sensor only makes communications from that sensor to be compromised, is impractical due to the limited amount of sensor memory if the network is very large. Our arguments on localization in sensor networks, in that the hypothesis of not knowing where individual sensors are located is flawed, are also applicable for pair-wise key pre-distribution.

In section 5, we discuss the motivation for aggregating data in a large sensor network and what is meant by aggregation. Our main argument on data aggregation is that it is meaningful only if the aggregation is done for data collected by small number of sensors in local area and then this datum should be transmitted to the base station without further aggregation along the hop-by-hop forwarding path. Our position is that in such a scenario, secure data aggregation, discussed in section 6, is more effectively done by subjecting data received at the base station to application dependent tests for validation.

We conclude this position paper by first discussing possible counter arguments to the positions taken by us in section 7 and in section 8 by reiterating that the oft stated application of air dropping sensors to build a large sensor network, especially in a military application context, is an unrealistic example to motivate the development of such a network with low-cost sensors.

2. Experimental results

First we explain an experiment that was conducted to collect data on the radio reception range for low-cost wireless sensors using TNOdes, a unit similar to the Mica2Dot. A TNOde sensor, which is a base for a sensory unit, has an Atmel AtMega128 processor, 4 KB of RAM, 120 KB of ROM and a Chipcon CC1000 radio transmitter operating at 868 MHz. The test units were powered by 2xAA batteries and were equipped with an $8\ cm$ omni-directional wire antenna as shown in figure 1 with a 2 euro coin for comparison of its size.

The objective of the experiment was to determine the average distance between two sensor nodes for correct reception of a message at two different antenna heights: (1) sensor $1\ m$ above ground (to simulate manual installation) and (2) sensor on the ground (to simulate deployment by air). The tests were carried out in three different environments to simulate a forest (an area with trees and bushes), a desert (an open space with minimum above ground structures), and an urban alleyway (a long and wide corridor with moving people and several static objects) and the results are summarized in table 1.

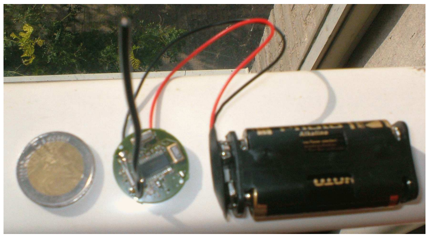

Figure 1. The TNOdes used in the experiment

Table 1. Average node-to-node radio transmission range values obtained in an experiment with TNOdes. The distances are in meters with measurements made in an environment of 22^oC and 68% relative humidity

Location	At $1\ m$ height	At ground level
Trees and bushes	25 m	12 m
Open space	35 m	7 m
Building corridor	42 m	35 m

The interesting observations from the experiment are:

- Foliage cover has no discernible effect on signal reception, irrespective of both the ground height and the ground location of the antennas, at the distances measured in the experiment. We tested correct receipt of a transmitted message in thick vegetation setting with line-of-sight (LOS) between nodes through the trees and also without LOS due to thick bushes on the transmission path.

- The reduction in signal reception range with decreasing ground height for the sensors in the open space setting was very rapid with an almost five-fold reduction from $1\ m$ height down to ground level. In contrast, there was only a minimal effect in this test for the building corridor.

If we consider the best case scenario for the air dropping of a 10,000 sensors uniformly on to a grid of 100x100 sensors on a desert, the sensor network will cover a maximum of $490,000\ m^2$ or approximately $0.5\ km^2$ based on data from our experiment. Therefore, even at 10,000 nodes, the purported battlefield will be very small and within the surveillance range of a naked eye and at 100,000 sensors, the area coverage is still less than $5\ km^2$. In contrast to the often cited air deployed battlefield observation network, a network of low-cost sensor can be more useful in an urban area installation. A 100 sensor network will provide a coverage of $175,000\ m^2$ (approximately 25 soccer fields), which makes it practical to use such a quickly deployable wireless sensor network for military or civilian surveillance purposes.

The above best case scenario for deployment of a low-cost sensor network is of theoretical interest only. For a practical network, the sensors may have to be spaced even closer together for reasons such as (1) to lower the probability of network disconnect due to sensor-to-sensor link failure, (2) to cope with those sensors that may be defective on installation or while in operation, (3) to accommodate sensory element range limitations (for example, a sensor for intruder detection may have a detection range much less than its radio communication range), etc.

From this experiment, we conclude that for effective and reliable coverage in applications such as battlefield monitoring, the sensor network will have to have a very high node density thus making it very vulnerable to detection and capture. A *passive attacker* who only listens to the radio transmissions of a sensor will invariably be in the close vicinity of that sensor due to the limited range of the transmitter. This fact combined with the high node density makes it very easy for an attacker to detect and capture sensors to become an *active attacker*. As the low-cost sensors are not hardened to be tamper-resistant, the attack model relevant to large sensor network is node capture attacks where the attacker can extract the cryptographic keys from the captured sensor.

3. Localization in sensor networks

Rather than having an end-to-end node to node communication pattern, a sensor network that sends data to a base station using a hop-by-hop transmission scheme is one of the simplest sensor network models. Such a model is also the most commonly applicable for a range of sensor network applications including intruder detection, environmental monitoring, and generating event alarms.

To install this type of sensor network in the field, it is necessary to place each sensor node at a location that provides radio connectivity with one or more neighbor nodes This work needs to be done by a trained and properly equipped technician to ensure: (1) an adequate received signal strength for the nodes that forward messages towards the base station, (2) for proper antenna height for the nodes to transmit messages, and (3) to satisfy application dependent requirements in safety of the device including physical security and camouflage After the correct installation of the sensor, the technician can easily record the physical location loc_i of a sensor (for example using a GPS device) against its unique identifier s_i. Thus, *localization* or position determination for sensors is essentially a manual task. It is impossible to satisfy the above listed practical requirements in signal reception and device safety by using other suggested approaches in the literature, such as dropping sensors from an aircraft or using mobile self-maneuverable nodes. Not withstanding the improvements in technology, *such units* cannot be low-cost, low-power and small all at the same time.

4. Key management in sensor networks

A sensor node in the above network can detect an event or obtain some measurement using its sensory unit and transmit a message to the base station. A simple way to achieve this goal is for the message to have the format $[s_i, count, data]$ where $count$ is a local counter maintained by the sensor. The base station will maintain a registry with a record $[s_i, loc_i, count]$ for each sensor and will update the $count$ value based on the messages received from the sensors. This counter allows the base station to detect both lost messages (when the counter sequence is broken) and multiple copies (for example, due to multipath propagation of messages).

A serious weakness in the above scheme is that an attacker can eavesdrop on sensor network communication to obtain valid sensor identifier values s_i and corresponding counter values $count$ from messages passing through his listening position. Thereafter, the attacker can inject false messages to the sensor network that will get forwarded to the base station. Also, an attacker can insert malicious sensor nodes to the network that accept messages from legitimate sensors but modify the data before forwarding the message to the next hop. Thus, it is necessary for the messages to have both integrity and authenticity protection.

For this sensor network model, integrity and authentication for messages transmitted by nodes can be easily achieved. Each sensor s_i can be configured with a symmetric key k_i shared with the base station and used to compute an authentication code value as $h = \texttt{MAC}_{k_i}(s_i, count, data)$ and the message $[s_i, count, data, h]$ transmitted. The SPINS security framework of Perrig, et al [3] for sensor networks is based on this trusted base station concept. The registry at base station will store the shared symmetric keys for all sensors as records $[s_i, loc_i, k_i, count]$. If the application requires confidentiality for the data transmitted over the sensor network to the base station, instead of computing an authentication code, the date can be encrypted as $c = \texttt{ENC}_{k_i}(s_i, count, data)$ to give integrity, authenticity and confidentiality for the transmitted message $[s_i, c]$.

In this security scheme, it is the task of the base station to detect false messages injected by attackers and message replays due to normal operating conditions or by attackers. The key management is limited to the preconfiguration of shared symmetric keys in each sensor and maintaining the key list at the base station. In contrast, the concept of pair-wise key pre-distribution introduced by Eschenauer and Gligor [4] and its many extensions [5, 6, 7, 8] in which sensors dynamically establish session keys is made relevant by the arguments that sensors are deployed over an area *randomly* in large numbers and the need for such sensors to securely communicate with other arbitrary selected sensors. The test data from the experiment discussed in section 2 clearly show the impracticality of such an application scenario. As a huge number of low-cost sensors area required to cover even a small geographical area, even with pair-wise key pre-distribution, the sensor will not have enough memory to store a subset of keys for arbitrary communication with another sensor.

The operating model described above and the unique shared secret key between a sensor and its base station preempts cloning attacks by an attacker who captures a sensor if the attacks are directed at the base station. For example, the base station of a sensor network to detect illegal release of hazardous effluents to the environment can determine if multiple sensory readings received by it are from the same sensor or different sensors easily as valid data needs a correct message authentication code. A plant operator who captures a sensor, make several clones and then program them to send *good* readings for his locality will not succeed in subverting the base station. However, an attacker can still succeed in disabling the sensor network by injecting spam messages in to the network to deplete the batteries of intermediary nodes as they needlessly forward the spam towards the base station [9].

The battery depletion through spamming can be significantly overcome by pair-wise key pre-distribution as a sensor will receive messages and forward them onwards to the base station only if the sender is authenticated through a shared key. However, this scheme is still not useful in a large network of low-cost sensors, as capturing a sensor is as easy as listening to transmitted messages due to their density, as explained in section 2. Once a sensor is captured, an attacker is able to extract the security data including keys and then clone many sensors and carry out a Sybil attack [10].

5. Data aggregation in sensor networks

In sensor network applications using low-cost sensors, the objective is to collect the data recorded by nodes at a central location, the base station, for processing and to initiate required responses. This is a many-to-one communication model. For example, a sensor network to detect forest fires will receive alarms from multiple sensors and then using the identifier values s_i of the message originating sensors and mapping it to the loc_i values, it can determine the actual physical location of the fires to send a response team. In general, it is unlikely that low-cost sensors will need to arbitrarily communicate with each other in a many-to-many model over the entire network. The applications suggested for large sensor networks built using low-cost nodes do not require such functionality.

However, in sensor networks, nodes are likely to communicate in a clustered hierarchical tree structure to allow in-network processing of data to provide *data aggregation*. For low-cost sensor nodes that are powered by batteries, the operational lifetime of the sensor is determined by the fixed capacity of the battery. The highest energy consuming component in a sensor is the wireless radio transmitter followed by the processor unit and finally the radio receiver unit. The data sheets of widely used low-cost sensors as well as experimental results [11] indicate that transmission of a single bit consumes as much energy as a thousand processor cycles. Therefore, it is important to reduce the number of bits transmitted over the wireless radio links in the hop-by-hop sensor network to maximize the operational lifetime of the network. A standard mechanism used for this purpose is to locally aggregate sensor data and transmit a single message towards the base station, thus saving transmission cost for multiple messages.

For many of the sensor network applications, it is meaningful to aggregate data only over a small local cluster of nodes and not over the entire network. For example, consider an environmental pollution monitoring sensor network that measure carbon monoxide levels in air over a 1 km by 10 km strip of land. While it is useful to know the pollutant concentration, for example, over every $100\ m^2$ area, it provides no meaningful scientific data if the sensor network computes an average value over the entire strip of land and transmit this to the base station. The base station operator will not know if there are any contaminated areas (for average concentration computation) or how many locations (for maximum concentration computation). The data from our experiments with low-cost sensors described in section 2 indicate that for a $100\ m^2$ area with vegetation cover only 10 to 15 low-cost sensors are required for full coverage and therefore a small local cluster.

For sensor network applications in which data aggregation can be useful and meaningful, the security requirements are authenticity and integrity for the data rather than confidentiality. For example, in the above environmental pollution monitoring scheme, an industrial company with a pollutant emitting plant above the legal limit can inject forged messages to the network to prevent the base station monitoring the system from obtain an accurate view of the pollutant concentration level over the area. However, in this example application, there is no requirement to provide confidentiality to the messages as anyone with a correct sensor device can obtain these values themselves. Therefore, the messages transmitted over the sensor network will have the format $[s_i, count, data, h]$ where $h = \texttt{MAC}_{k_i}(s_i, count, data)$ and s_i is the identifier of the message originating sensor.

6. Secure data aggregation in sensor networks

At the time a sensor network is installed, a subset of the sensors can be designated as aggregators. These aggregator nodes can be selected based on many different criteria, including (1) the location of the sensor node in the transmission path towards the base station, (2) the availability of a larger battery allowing for extra energy capacity for local processing, and (3) favorable geographical location providing better radio communication with neighboring sensor nodes. The presence of aggregator nodes can be advertised as part of the routing table setup for the sensor network in its initialization phase.

For example, consider a sensor network that use a simple distance vector routing scheme where the base station (`bs`) broadcast a `0-hop-to-bs` message and it's immediate neighbors broadcast `1-hop-to-bs` and so on, allowing the nodes to select the immediate neighbor to transmit a message for forwarding towards the base station base on the lowest hop count. An aggregator node (`an`) can append a message to this routing advertisement as `0-hop-to-an` and its immediate neighbors who are not aggregators will modify the message as `1-hop-to-an` and so on. Another aggregator node that receives this routing advertisement will reset the hop count to 0 and rebroadcast the message. This scheme allows a sensor node to select the next hop node for message transmission based on either shortest distance to the base station or nearest aggregator node.

A considerable amount of research literature [12, 13, 14, 15] is available on the issue of data aggregation in sensor networks with compromised nodes and many of the proposed schemes provide algorithms for computing a value by the aggregator over a set of data points received from other sensors. The suggested values for computation are a minimum, an average, a maximum, etc. However, the data aggregation in practical sensor network applications can be much simpler if done limited to a local cluster of sensors.

For example, if the application semantic is to forward the measured minimum or maximum value of a quantity from a locality, the aggregator node can simply cache the data it receives from it's neighbors and then select the message with the minimum or maximum value for onward transmission towards the base station. The message security is provided by the authentication code computed by the message originator and the aggregator does not have to compute any security related value. If an attacker were to inject a forged message that gives a false minimum or maximum value then the aggregator will forward this message to the base station. While the base station can detect the forged message it will not be able to obtain the correct data point for that particular locality. If the aggregator enclose the message to be forwarded in a message that it creates and compute an authentication code on the full message, the base station will be able to determine the aggregator node and therefore the locality under attack. A possible countermeasure is for the aggregator node to apply several heuristic/statistical tests to the received data to exclude both false and erroneous data. These tests would be application and context dependent We do not consider cryptographic solutions to authenticate neighbor nodes to the aggregator node as useful due to the fact pointed out in section 2, that capturing a sensor is as easy as listening to their radio communication.

Another example of data aggregation is in intruder detection. If the aggregator node receives several alarm messages from sensors that use it as the aggregator, it can generate a new alarm message itself for onward transmission to the base station. This message can contain the sensor identifiers that sent alarms to the aggregator node so that the base station can determine the actual physical locations at which intrusions occurred. As mentioned earlier, an attacker can insert forged alarm messages so that the integrator node sends a false alarm to the base station. A possible solution is again to use some heuristic such as a threshold of alarms from neighbors before sending an alarm to the base station.

The remaining type of data aggregation is where the aggregator node computes some function on the data received from other sensors in the neighborhood. This function could be as simple as computing the average or as complex as some form of matrix manipulation. Important point to investigate is if this application requires secure computation of the function. While there are straightforward reasons for an attacker to insert false messages into a forest fire detection network (by pranksters to send the fire fighters to a nonexistent event), an intruder detection sensor network (by the intruder trying to distract the guards) or environmental pollution monitoring network (by the industrial company contravening the laws), the research literature does not give example applications where a complex aggregation function needs to be computed securely. Additionally, the application code requiring the computation of a complex function at an aggregator node, such as a scientific experiment, will not be able to fit inside the small program/data memory of a low-cost sensor.

7. Effect of future technology improvements

A common argument against the above line of reasoning is that future improvements in processor fabrication, radio communication and battery technology will allow low-cost sensors to accomplish what they cannot do today. However, these predictions may turn out to be too optimistic or even unrealistic.

For example, let us consider the issue of radio transmission distance between two sensors. In the first study on optimum single-hop distance for wireless ad-hoc networks to minimize total system energy by Chan, O'Dea and Callaway [21], it is clearly shown that the optimum distance depends only on the propagation environment and analog hardware device parameters The energy consumption by a node for each transmission of a packet is given as

$$E_t(r) = k_1 r^{\omega} + k_2 \tag{1}$$

where r is the radio transmission range (hop-to-hop), ω is the path loss exponent, k_1 to denote the transmitter and channel characteristics, and k_2 to denote transceiver energy consumption not related to r. If E_r is the energy consumed by a node to receive, decode and process a packet at the receiver at a distance r, then the total energy consumed for a single transmission is $E_t(r) + E_r$.

The research work of Chan, O'Dea and Callaway [21] also shows that optimum distance is independent of the total hop-by-hop transmission distance, physical network topology, and the number of transmission sources. Therefore, while it is practical to build very large irregularly shaped multiple transmitter wireless sensor networks that transmit data hop-by-hop, the density of such networks will remain very high as the single-hop distance is not dependent on device parameters alone. Not withstanding improvements in electronics, parameters of the propagation environment such as path loss attenuation and path loss exponent are technology independent values.

Another important issue highlighted by Chan, O'Dea and Callaway [21] is that for short range and low bit rate transmissions to be energy efficient, the device electronics power dissipation needs to be in the *micro Watt* range while the radio is in active transmission and receive mode. In contrast the latest integrated radio chips targeted for low-cost sensor fabrication, such as the Chipcon CC1010, has a programmable output power range of -20 dBm to +10 dBm (equal to 0.01 mW to 10 mW)[22]. This is still 1 to 3 orders of magnitude higher than the *milli Watt* range rating required. While the study by Chan, O'Dea and Callaway [21] assume a dense wireless network so that intermediate nodes for hop-by-hop routing of messages is easily found, according to Deng, et al [23] even for a low density network the optimum transmission distance is still influenced more by nodal density than the coverage area. Therefore, our argument that the energy optimum single-hop distance is a critical network parameter and that it cannot be vastly improved through advances in technology is still valid.

Another counter argument made on our assertion that limitations of battery technology and the fact that the battery is the most critical system component in a battery powered sensor node is to point towards the Moore's Law and hypothesize future sensors that will be coupled to small yet higher capacity batteries. The problem with this reasoning is that Moore's Law is based on semiconductor lithography technology and the number of transistors on a silicon chip while improvements in battery technology will be determined by advances in electro-chemical technology. However, as noted by Gasman [24], the energy density of lithium-chemistry batteries have been increasing only at a few percentage points every year in comparison to the doubling of capacity every 18 months for microprocessor technology [25].

8. Conclusion

A consistently recurrent statement in research literature [16, 17, 18, 19, 20] is that battlefield surveillance with large number of sensors dropped from the air to be on of the most probable sensor network applications. While some of the research work consists of experimental work none provides directly relevant test data for feasibility or usefulness of an air deployed sensor network application. On the basis of these early discussions, other literature also continue to cite sensor network applications for a large network of low-cost sensors with an aircraft dropping *smart dust* style motes in to a battlefield so that the military can monitor enemy movements. This is a useful application in so far as it embodies many of the research questions:

- localization: on a large area, sensors are dropped from air and the operators need to know their actual physical location,
- key management: as the neighbor nodes of any given sensor cannot be predetermined, it is necessary to device a mechanism to exchange keys and establish secure links among these ad-hoc neighbors, and
- data aggregation: due to the fixed battery capacity and also the need to minimize the amount of transmissions to prevent the enemy from easily locating the sensors (for example, see Wood and Stankovic [26] on sporadic transmissions to defeat jamming and rate limiting at MAC layer to counter battery exhaustion).

However, in reality, if the battlefield is a desert, a satellite or a high-altitude surveillance aircraft is a better and cheaper solution as these assets are reusable over a long period. For forest areas where the tree canopy prevents effective surveillance by air, the air-dropped sensors will have to land on top of the trees to provide hop-by-hop radio coverage while lowering their sensor unit to ground level. It must also be able to withstand the strong forces associated with landing by air, must not land in a ditch or hole to allow correct antenna operation and have camouflage to prevent easy detection. Clearly, such a sensor cannot be classified as a low-cost unit that can be air-dropped by the thousands. If the battlefield is a built-up urban area or an area consisting of lakes, canals, etc; again this kind of sensor network establishment is neither practical nor desirable. In the fast-faced urban combat scenario, the sensors need to interact with the soldiers for them to be useful and must be maneuverable to locations of interest unlike a fixed sensor network . As this example sensor network cannot be established on desert, forest, city or lake, the example remains the mythical air-dropped sensor network.

The security requirements for very large sensor networks constructed using low-cost sensors need to be met with a combination of simple cryptographic techniques such as a message authentication code and application-level heuristics. In fact, for all the sample applications suggested in research literature for such sensor networks, the required security can be provided by simple techniques such as a unique secret key shared between each sensor and the base station and the base station applying application specific tests on received data sets to determine the validity of data. In contrast, for complex security schemes proposed in literature, such as for ad-hoc key establishment between two sensor nodes in a network, no example applications are cited .

The full paper is available at the one-time sensor project page http://www.cs.vu.nl/˜chandag/sensors/index.html.

References

[1] Crossbow Technology, Inc, *Motes, smart dust sensors, wireless sensor networks*

[2] J. Hill & D. Culler, *Mica: A wireless platform for deeply embedded networks*, IEEE MICRO, 22(6):12-24, 2002.

[3] A. Perrig, et al, *SPINS: Security protocols for sensor networks*, ACM/IEEE MobiCom ’01, pp 189-199, 2001.

[4] L. Eschenauer & V. D. Gligor, *A key-management scheme for distributed sensor networks*, ACM CCS ’02, 2002.

[5] H. Chan, A. Perrig, & D. Song, *Random key predistribution schemes for sensor networks*, IEEE SoSP, pp 197-213, 2003.

[6] W. Du, et al, *A pairwise key pre-distribution scheme for wireless sensor networks*, ACM CCS ’03, pp 42-51, 2003.

[7] D. Liu & P. Ning, *Establishing pairwise keys in distributed sensor networks*, ACM CCS ’03, pp 52-61, 2003.

[8] W. Du, et al, *A key management scheme for wireless sensor networks using deployment knowledge*, IEEE INFOCOM 2004.

[9] H. Chan & A. Perrig, *Security and privacy in sensor networks*, IEEE Computer, 36(10):103-105, 2003.

[10] J. Newsome, et al, *The Sybil attack in sensor networks: analysis and defenses*, ACM SIPSN, pp 259-268, 2004.

[11] V. Raghunathan, et al, *Energy-aware wireless microsensor networks*, IEEE Signal Processing, 19(2):40-50, 2002.

[12] L. Hu & D. Evans, *Secure aggregation for wireless network*, IEEE SAINT ’03: SAAN, pp 384-394, 2003.

[13] J. Zhao, R. Govindan, & D. Estrin, *Computing aggregates for monitoring wireless sensor networks*, IEEE SNPA, pp 139-148, 2003.

[14] B. Pryzdatek, D. Song, & A. Perrig, *SIA: Secure information aggregation in sensor networks*, SenSys ’03, pp 255-265, 2003.

[15] A. Perrig, J. Stankovic, & D. Wagner, *Security in wireless sensor networks*, CACM, 47(6):53-57, 2004.

[16] D. Estrin, et al, *Next century challenges: Scalable coordination in sensor networks*, ACM/IEEE MobiCom 99, pp 263-270, 1999.

[17] I. F. Akyildiz, et al, *Wireless sensor networks: A survey*, Computer Networks, 38(4):393-422, 2002.

[18] R. Malladi & D. P. Agrawal, *Current and future applications of mobile and wireless networks*, CACM, 45(10):144-146, 2002.

[19] R. Anderson, H. Chan, & A. Perrig, *Key infection: Smart trust for smart dust*, IEEE ICNP ’04, 2004.

[20] A. Arora, et al, *A line in the sand: A wireless sensor network for target detection, classification, and tracking*, Computer Networks, 46(5):605-634, 2004.

[21] P. Chen, B. O’Dea, & E. Callaway, *Energy efficient system design with optimum transmission range for wireless ad hoc networks*, IEEE ICC ’02, 25(1):945-952, 2002.

[22] Chipcon AS, *SmartRF CC1010 Data sheet*, 2004.

[23] J. Deng, et al, *Optimum transmission range for wireless ad hoc networks*, IEEE WCNC ’04, vol. 2, pp 1024-1029, 2004.

[24] L. Gasman, *Future opportunities in the mobile power IC business*, Advanced Battery Technology Magazine, Nov, 2003.

[25] R. L. Mitchell, *Energy crisis 1: Why batteries aren’t good enough - Moore’s law doesn’t work for power supplies*, Techworld, Jan, 2005.

[26] A. D Wood & J. A. Stankovic, *Denial of service in sensor networks*, IEEE Computer, 35(10):54-62, 2002.