

Rescuing Wireless Sensor Networks Security from Science Fiction

Dieter Gollmann, Maryna Krotofil and Harald Sauff

{diego,maryna.krotofil,harald.sauff}@tu-harburg.de

Institute for Security in Distributed Applications

Hamburg University of Technology

21079 Hamburg, Germany

Abstract. We critically analyze the state of the art in research on wireless sensor network security. Assumptions about security requirements are not always consistent with the assumptions about the nature of sensor nodes. There are deficiencies in the specification of attacker models. Work on wireless sensor network security often fails to give proper definitions and justifications of what constitutes node misbehaviour. We analyze the merits and limitations of reputation-based routing protocols as a security mechanism, and observe that in wireless sensor networks there is a strong case for using application specific cross-layer optimizations and hence a diminished demand for generic security solutions.

1 Introduction

Early milestones in the research on wireless sensor networks are the Smartdust project¹, the NASA Sensor Webs project [4] and, on a related topic, Dynamic Source Routing (DSR) [11]. This work can be roughly dated to the second half of the 1990s. Since then a considerable body of work has been published on wireless sensor network security. Indeed, articles on wireless sensor networks figure prominently in the Citeseer list of most quoted papers².

It is inevitable that a promising new technology does not have many concrete applications in its early days. It has further been repeatedly observed that there is a considerable time lag between the conception of a new idea and its actual adoption, consider e. g. the following comment from [14]:

It typically takes at least 10 to 20 years for a good idea to move from initial research to final widespread practice.

Today we can look back at more than a dozen years of research on wireless sensor networks but it still remains a ‘promising’ technology with limited deployment. This is a problem for security research. Security requirements inherently depend on the application of a technology. When there are few real applications, and when these applications are not particularly security sensitive, take e.g. the ZebraNet project³, then security researchers have little choice but addressing

¹ <http://robotics.eecs.berkeley.edu/~pister/SmartDust/>

² <http://citeseer.ist.psu.edu/stats/citations>

³ <http://www.princeton.edu/~mrm/zebranet.html>

generic virtual problems. The best they can do is picking plausible problems based on consistent assumptions.

This aspect is too often neglected in research papers on wireless sensor network (WSN) security. Assumptions are simply copied from previous work without applying basic sanity checks. We will examine the major assumptions about wireless sensor networks, discuss some frequent fallacies, and point to research directions that might be followed whilst we are still waiting for concrete, security sensitive applications of wireless sensor networks (outside the military domain). In particular, we will comment on the use of reputation (trust) in security mechanisms.

Section 2 covers typical assumptions about sensor nodes, noting that under these assumptions some standard security mechanisms would be ineffective. Section 3 introduces the topic of ad-hoc routing. Section 4 deals with the definition of misbehaviour. Section 5 discusses issues arising when applying reputation systems in wireless sensor networks. Section 6 provides a critique of the mobility patterns typically used in WSN simulations. Section 7 makes the case for cross-layer considerations in the design of WSN security solutions. Section 8 concludes the paper.

2 Nodes

A typical sensor node has the following core properties: a sensor measures parameters from the environment, communicates on a short range radio channel, has limited energy supply, has limited computational resources, and is not tamper-resistant. We will now examine the implications of these assumptions on WSN security research.

2.1 Limited Power and Computational Resources

Limited power and limited computational resources are a popular motivation for research on light-weight cryptography. These are real limitations but their significance can be exaggerated. The following points must be noted.

- The main drain on power is not computation but communications; power savings due to light-weight cryptography may not be significant for an application overall.
- Modern cryptographic algorithms such as AES have been designed cognisant of current microprocessor instruction sets. New algorithms must show significant improvements to justify the switch from a thoroughly evaluated standard algorithm.
- Do not confuse temporary limitations with fundamental barriers. Experience in other fields, e. g. in the smart card sector, shows that resources will eventually become available if required by the applications.

It is thus more promising to consider the way an application uses a WSN and try to reduce communications while still meeting the goals of the application. This requires, of course, that there is an application in the first place.

2.2 Wireless Communications

The Unit Disc Model (UDM) is a common simplification when modelling wireless communications. UDM postulates that the sending and receiving range of the transceiver circuit on wireless nodes is equal in all directions, yielding a perfect circle of connectivity with the node at its centre. Within this circle nodes can receive the communication signals of other nodes; their own radio waves can be received by other nodes with enough signal quality for messages to be decoded by those nodes.

This model does not take several effects into account. One example is multi-path propagation. A wireless sensor network is rarely deployed on a flat plane devoid of any obstructions. Signals bounce off obstacles in their path, like e. g. buildings or hills, so that they may reach the receiver through more than one propagation path. Different paths have different effects on the signal: they are distorted depending on the environment and the reflections, they fade with different intensities and they arrive at different times so that interferences occur. These effects decrease or might even increase signal ranges in a way that the covered area cannot be considered circular anymore.

The position of the antenna can have an even bigger effect on the performance of radio communication. Most standard sensor node hardware uses a simple piece of wire as the antenna. The biggest fraction of the radio wave energy is emitted radially from the wire. Radio reception is best when all antennas are positioned orthogonally to the plane in which the nodes are in. But even when this rule is adhered to during deployment other problems might show up: the battery or the casing of the node might be in the way or influence radio communication.

Even when the UDM is accepted and the radio range is assumed equal in all directions a next problem might arise: the range is not equal for all nodes. Due to remaining energy resources, energy saving schemes or the position of the node in the environment two equal nodes could both have radio coverage in the shape of a circle, but with different diameters. Then, sending and receiving range for a node need not be identical. Measurements on our campus WSN conducted over an extended period of time show that it is an exception when a channel is measured from both sides and yields the same results [7].

The assumption of symmetric links is thus in general wrong. This is problematic for schemes that monitor how neighbouring nodes behave, e. g. as the basis for routing decisions. There is a limit to the ability to observe that something has not happened. The observer may simply remain unaware of an event. The local view held by a node may thus not correspond to the global view. This contradicts a standard assumption in reputation systems (Section 5) that direct information is always correct.

2.3 Short Range Communications

Short range wireless communication between sensor nodes suggests that an attacker has to be in the vicinity of a sensor to intercept or manipulate its traffic.

This in turn suggests that such an attacker would also be in a position to tamper with the sensor itself. There may then be little merit in using cryptographic protection. An attacker close enough to listen to traffic would be in a position to compromise the sensor generating the traffic.

The same question arises in the analysis of the Eschenauer-Gligor key distribution scheme [8] and its variants. In these schemes each node is equipped with a set of secrets and can establish session keys with nodes it shares a secret with. An attacker might use the secrets obtained from a compromised node to deduce session keys used by other nodes that happen to hold one of the compromised secrets. Such a session key is of value if the attacker is close enough to the node to hear it, and hence close enough to compromise the node directly.

These arguments do not imply that cryptographic protection is always unnecessary but care has to be taken when making the case for communications security.

2.4 Sensor Data

A sensor measures parameters from the environment. Environmental parameters are likely to be observable by any party in physical proximity of the sensor. Using the sensor data in a sensitive application thus does not automatically imply that the confidentiality of the data sent by a sensor node needs to be protected.

Consider a setting where temperature readings are transmitted via a few hops to a base station. An attacker close enough to the sensor nodes to listen to their short range wireless communications will be close enough to take temperature readings on its own. There is not much gained by encrypting sensor data. A remote attacker would only get access to traffic after it has gone through a base station. At this point, encryption may become advisable.

3 Routing

Wireless sensor networks provide a communications infrastructure for routing data from the sensors to some data sinks. The base stations mentioned above are one example for such a data sink. When nodes are deployed in an ad-hoc fashion there is no predefined routing infrastructure. Ad-hoc routing protocols, such as DSR, and their security have been extensively studied. Routing is a generic network service that can be examined independently of any specific WSN application.

A large portion of the work on WSN security addresses the security of routing protocols. A security analysis needs to state its threat model. Under the assumption that sensor nodes are not tamper resistant, the customary threat model assumes that the network may contain compromised, misbehaving nodes. The next section will explore the possible meanings of ‘misbehaviour’.

4 What is Misbehaviour

‘When I use a word,’ Humpty Dumpty said in rather a scornful tone, ‘it means just what I choose it to mean – neither more nor less.’

‘The question is,’ said Alice, ‘whether you can make words mean so many different things.’

‘The question is,’ said Humpty Dumpty, ‘which is to be master – that’s all.’

(Through the Looking Glass, Lewis Carroll, 1871)

In this section we intend to disambiguate the meaning of the term ‘misbehaving node’. In the WSN research literature the term ‘misbehaviour’ usually refers to nodes which do not behave in a proper way. However, in many cases it is not specified what kind of behaviour is considered as improper, leaving this to the reader’s imagination. Furthermore, several other terms denoting misbehaviour can also be found. In some cases these terms are used as synonyms for misbehaviour in general. However, often they indicate a particular form of misbehaviour. The terminology can hence become quite confusing. We will try to bring some structure into this discussion.

Misbehaviour can take many forms. According to the Oxford English Dictionary, if a person fails to conduct itself in an acceptable way or behaves badly, he/she misbehaves. In the realm of hardware, if a machine fails to function correctly, it misbehaves. For wireless sensor networks we may interpret this definition in the following way: if a node’s behaviour deviates from its specification, it misbehaves. With this definition, misbehaviour depends on the specification of intended behaviour. Any deviation from the specified behaviour would be considered as misbehaviour, regardless of the reason causing the deviation.

Continuing our linguistic endeavours, we have collected from the research literature a set of terms standing for node misbehaviour, viz. failed, malfunctioning, greedy, neglectful, selfish, free-rider, subverted, compromised, evil and malicious node. While some of these terms indicate distinct forms of misbehaviour (greedy vs. failed), others can be used interchangeably. Some of the terms have a strong anthropomorphic flavour, which can further complicate the discussion. Sensor nodes are computers. They are neither benevolent nor malicious, they do not react to incentives or punishments as a human might do; their intended behaviour is programmed.

We now propose a classification of node misbehaviour and provide directions for the use of this terminology. Depending on the nature of the deviating behaviour, a misbehaving node falls into one of the following categories:

- **Malfunctioning misbehaviour.** A node can malfunction/fail because of hard- or software problems, climate influence, radio channel interferences or link breakdown, bad location, accidental physical damage, etc. Nodes can fail once, repeatedly, randomly, short term or long term.
Suggested terms: failed, malfunctioning node.

- **Commercial misbehaviour.** A node can be programmed in a specific way in order, for instance, to save its own power and thereby prolong its own life expectancy. Such behaviour could manifest itself in the dropping of packets from certain other nodes, non-participation in route discovery or unfair channel occupancy. Such misbehaviour can be intended by a manufacturer to favour its own nodes in order to preserve energy and thus to outperform the competitors' nodes. On the other hand, greedy behaviour of nodes may be the unintended result of deviating from the specification given.
Suggested terms: selfish, greedy, neglectful nodes; free-rider.
- **Malicious misbehaviour.** A compromised node is re-programmed to execute a targeted attack. Malicious nodes can also be extraneous nodes, injected into the network by an adversary. An attacker might want to harm or severely disrupt communication, manipulate data or destroy the network.
Suggested terms: subverted, compromised, malicious, evil node.

Malfunctioning misbehaviour and commercial misbehaviour may be indistinguishable for an observer. This distinction may matter when assessing the impact of countermeasures. The main difference between commercial and malicious misbehaviour is the ultimate goal that drives a node to misbehave. Commercially misbehaving nodes try to maximize their own performance disregarding overall performance of the network. Malicious nodes primarily try to attack the network (and ultimately the application served by the network), potentially disregarding exhaustion of their own resources.

5 Reputation Systems for WSN

‘The way to gain a good reputation is to endeavour to be what you desire to appear.’

(Socrates)

‘You don’t build a reputation on what you’re going to do.’

(Henry Ford)

We will now present the principles and mechanisms of reputation schemes and discuss their ability to mitigate the negative influence of misbehaving nodes. Node cooperation is essential for the functioning of a multi-hop wireless network. In the absence of a fixed infrastructure, the sensor network forms a community of peers, which share the obligations of forwarding and processing gathered data. The success of the network in fulfilling its mission depends on the ability of all nodes to execute real-time routing functions in a coordinated manner, fairly use network resources, accurately measure, communicate, and process sensor data.

Non-cooperating nodes can limit the value of a wireless network via (i) non-participation in routing or packet forwarding, (ii) incorrect sensing or processing of data, (iii) preventing other nodes from executing their functions. It would thus be useful to have a mechanism to detect such nodes and exclude them from the network in order to keep network performance on a high level and to obtain correct application data.

A first step towards securing any type of network can be the use of cryptography for entity authentication and for message integrity protection. This would assure that only authorized nodes participate in the network and that messages sent by legitimate nodes have not been tampered with during transmission. Although sensor nodes are capable of symmetric cryptographic operations, the nodes themselves typically offer little tamper resistance. Considerations like this have induced doubts about the effectiveness of cryptographic methods to secure wireless sensor networks and look for alternatives.

5.1 Observations and Recommendations

Self-policing mechanisms based on reputation scores have been proposed for automatically estimating the quality of a node's behaviour and to deprive non-cooperative nodes of network services. Nodes observe each other to detect inconsistencies in the behaviour of their neighbours and then form an opinion about them. This opinion can be input to routing decisions, and also be passed as secondary information (recommendations) to other nodes. The mathematical foundation of forming reputation values of individual nodes has roots in statistics, belief and game theory. The major advantage of reputation- (a.k.a. trust)-based mechanisms is their relatively low overhead and their potential ability to successfully identify different types of misbehaving nodes, including nodes compromised by a strategic attacker, see e. g. [9, 2, 12].

One might arrive at the conclusion that this can be a good way to enhance WSN security. However, is this really the case? For a reputation system to work as intended certain conditions must hold (see also [3]):

1. Past behaviour predicts future behaviour reasonably well.
2. The reputation system's data is reasonably correct.
3. Reputation information is available.
4. Nodes within the network have unique identities.

The first premise holds for malfunctioning and selfish nodes. Even if their behaviour turns out to be erratic at a certain moment, inconsistent with 'good' past behaviour, it is still predictable, either in a statistical sense (malfunction) or because a behavioural pattern can be detected (commercial misbehaviour).

In contrast, strategic behaviour of an intelligent attacker is unpredictable. A compromised node may fully comply with the specification for a long time, but start to deviate at a time decided by the attacker. Moreover, an attacker may adapt her behaviour to the reputation mechanism, change a behaviour pattern, exhibit different behaviours with respect to different interaction partners, or simply alternate 'good' and 'bad' behaviour in the attempt to avoid detection.

Compliance with the second condition requires assurance that:

1. Reputation information observed and reported by others is truthful.
2. The integrity of the reputation scores stays intact while being forwarded.

In order to be useful, reputation values need to be accurate, at least to some degree [3]. This condition can again be violated by a cunning intruder. She can lie about the behaviour of other nodes and falsify or manipulate reputation values. Although a typical reputation system can identify ‘small’ lies, it does not help against sophisticated lies and colluding malicious nodes (see Figure 1). If two adjacent nodes on a route cooperate, they can launch a ‘conspiracy of silence’ attack. In this form of attack, one of the nodes drops or modifies the packets, whereas its ally does not disclose the fact of misbehaviour.

Distributed reputation systems have no dedicated server for storing node identities and their reputation ratings. This information is distributed among all the nodes. Usually each node has a partial view of the network only and just a subset of the node behaviours will be known to it. Another challenge is sharing observations. As discussed in Section 2.2, communications links are not symmetric. This can lead to wrong recommendations and to missed recommendations.

A node’s identity persistence over time is crucial for the effectiveness of a reputation system [3]. The distinctness of identities is violated by node replication and Sybil attacks. In a node replication attack the adversary attempts to add one or more nodes to the network that use the same ID as another node in the network [15]. In a Sybil attack, a single node illegitimately claims multiple identities [13]. A misbehaving node may continuously misbehave by cloning more and more new identities or use its multiple identities to badmouth a victim. Replicated nodes can confuse the reputation system by exhibiting inconsistent behaviour to different nodes. If node A observes cloned node B to behave properly for a long period, but gets negative recommendations for B (caused by a clone of B) from node C , it may lower its recommender rating for node C .

5.2 Effects of Reputation Systems

The usefulness of a reputation system suffers in the presence of a strategic adversary. For its accurate functioning, a reputation system then depends on the security mechanisms that provide protection against the attacks mentioned above. However, employing additional security mechanisms comes at the price of increased resources consumption, which is difficult to reconcile with the resources scarcity assumption in WSN. Moreover, we are not aware of any reputation system that copes with colluding malicious nodes.

Apart from detecting misbehaving nodes, a reputation system may provide an incentive mechanism to desist from misbehaviour by making misbehaviour unattractive [1]. Incentives work well in a system with human actors, e. g. in e-commerce systems like eBay. The desire to generate profit motivates individuals to provide a high quality service in order to keep their reputation score high thus attracting more buyers. Moreover, human behaviour is flexible and can be dynamically adapted to keep one’s own reputation rating high.

In contrast, sensors execute predefined code. Once deployed, a sensor node neither gets upset because it is suffering punishment nor can it adapt its own behaviour. A failed node would not be susceptible to incentives, as it is no longer capable of executing its functions properly. At best, the reputation scheme makes

the life of the attacker more difficult and motivates her to hide the misbehaviour in order to stay undetected as long as she needs in order to execute her attack.

A party who might possibly react to an incentive is the node manufacturer. A manufacturer does not want his nodes being excluded from the network. Therefore he has an incentive to program his devices in a way that they do not fall foul of the reputation system.

We thus conclude that reputation systems are able to deal with malfunctioning and commercial misbehaviour, but are not robust enough to curtail strategic attacks. Therefore, claims that a reputation system is suitable for mitigating **any** kind of misbehaviour overstate their case. It is indeed doubtful whether reputation mechanisms should be categorized as security mechanisms. A potential alternative to reputation-based routing protocols are solutions based on shadow pricing, as proposed for P2P networks [5, 6].

5.3 Attacker Models

Researchers need to make assumptions about the attackers' abilities when evaluating the performance of their security solutions. The attacker models typically adopted in sensor network simulations are very simplistic. Malicious actions are predominantly limited to selfish behaviour, dropping a predefined percentage of packets [9] or extreme lies, i. e. reporting either extremely negative or extremely positive recommendations for a peer [12, 9]. On the other hand, attackers never make 'intelligent' choices; they do not try to cut off selected nodes, do not insidiously blame other nodes or fake reputation messages, do not adapt their behaviour to the reactions of the security system, and they never collaborate.

Many security mechanisms can be circumvented when a group of malicious nodes works together, see e. g. [18]. We provide further illustrative examples from simulations of collusion attacks where CONFIDANT [1] was used for reputation-based routing. CONFIDANT promises to detect dropping attacks, some types of fabrication of messages, and big lies attacks.




Fig. 1. Selective forwarding attack by colluding malicious nodes

Our first example⁴ has a predefined stationary network deployment. By chaining two consecutive malicious nodes along the forwarding path, an attacker can successfully launch a dropping/selective forwarding attack (Figure 1) and a sinkhole attack without being detected by CONFIDANT.

Our second example assumes a random waypoint mobility model. Scenarios 1–5 capture different randomly generated movement scenarios. Table 1 refers to an attack where two nodes collude so that EVILNODE drops packets and its partner does not disclose fact to the reputation system. (In three of the scenarios no traffic was routed via EVILNODE and hence no packets were dropped.)

Table 1. Conspiracy of silence attack by colluding malicious nodes

Parameter	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 5
Number of packets dropped by EVIL-NODE	87	0	0	10	0
Identified as misbehaved	No	No	No	No	No

We have simulated an accusation attack where malicious nodes attempt to remove a node from the network by reporting false negative recommendations about the victim. Table 2 gives the identities of the attacking nodes and of the nodes marked as misbehaving. The attackers succeeded in all five scenarios to badmouth the victim node 4; only in one scenario one of the attackers was identified as misbehaving. We conclude that assessing security schemes only against trivial attacks can be misleading as it does not help to evaluate the effectiveness of the proposed solutions in the presence of strategic attackers.

Table 2. Accusation attack on victim node 4

Parameter	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 5
Evil nodes falsifying reputation info	2, 8	2	2	2	2, 8
Identified misbehaving nodes	4, 8	4	4	4	4

5.4 Priorities of Attacks

It is difficult to conduct a proper risk assessment in sensor networks because, in contrast to standard internet security, it is mostly unknown which problems/attacks occur most frequently. When deciding on the necessary security measures

⁴ Taken from Nguyen Dang: Simulation Intelligent Attacker on Wireless Sensor Network Routing Using GlomoSim Simulator, Master project, TUHH, 2008.

for a system on the internet the most common vulnerabilities and attacks are known: weak passwords in general, password authentication instead of public-key based authentication on SSH servers, SQL injection or Cross Site Scripting (XSS) in web applications, and so on. For more details see e. g. the CWE/SANS Top 25 Most Dangerous Software Errors list⁵. From this experience a risk analysis can create weighted lists with the most important problems which have to be fixed first.

Since so few real world sensor networks exist we have next to no documented experiences with attacks. Which aspect of the network is attacked? Which data is most valuable to the attacker? Which attack occurs most often? It is more or less impossible to judge which problems have to be addressed with the highest priority when it is unknown what real attackers – in comparison to academical attackers – really want.

Since the presentation of [10] this problem might have partly ‘fixed’ itself: Giannetsos et al. published a framework for an attack tool specialized in sensor networks. The focus of this extensible tool on certain security weaknesses decides which vulnerabilities will be the first ones to be exploited. The availability of an automated, easy-to-use tool now defines which vulnerabilities will be most critical; it does not have to wait for others to do so.

Nevertheless, we are in fact back to the problem of solving virtual problems: Solutions exist, but do they solve the right problems? Practical and useful solutions can only be present when a proper threat modelling has been done and the weaknesses are prioritized in the right way, but this is only possible when the analysis is made with a specific application in mind.

6 Mobility Patterns

In wireless sensor networks nodes may be mobile. Although some typical WSN scenarios like structural monitoring of buildings usually do not expect movement of sensor nodes, many others do. Mobility has an effect on the routing protocol due to the ever changing reachability of nodes, and therefore on the routing decisions as well as on data aggregation and reputation protocols.

Evaluation results should be reproducible. This can be ensured in different ways when simulating and evaluating WSN protocols. On the one hand the recorded movement of real-world motions, a so-called trace, can be replayed in every simulation run. Traces resemble realistic movement the closest. However, not many databases with traces exist and the recorded traces can only be used as-is: subsequent adjustment is difficult.

On the other hand the movements for every run can be calculated from a mathematical model. Mathematical mobility models are easy to obtain and easy to parameterize. That makes them the favourite choice when the mobility of nodes shall be taken into account during simulation runs. Mathematical mobility models can be categorized into several groups. The two main groups, entity

⁵ <http://cwe.mitre.org/top25/>

mobility models and group mobility models, can be further subdivided. Examples for the first group are the random waypoint model or the random walk model; examples for the second group are the reference point group model or the nomadic community model.

When analyzing results from simulation runs with reputation systems like CONFIDANT two observations can be made:

- The use of the reputation system does not increase network performance; it only limits degradation, i. e. the impact of misbehaving nodes on network performance is not as severe as without the reputation system [16].
- In combination with mobility, the performance of the network degrades even with no attacker present.

The second observation is due to the exclusion of nodes, although no node is actually deviating from the specified behaviour. The connection between ‘more mobility’ and ‘more falsely accused nodes’ is clear. Figure 2 shows the effect of mobility on the Packet Delivery Rate in comparison to DSR, i. e. with no reputation system applied. For the simulation setup, the simulation parameters used and more results concerning the effect of attacks against the reputation system see [17].

In the presence of misbehaving nodes, mobility has an even larger effect on network performance, since misbehaving and moving nodes are prone to exclusion from the network. We have tested two similar scenarios where 30 out of 50 nodes show selfish behaviour and drop all packets passed to them. In static scenario regular nodes still obtained a 73% throughput of their packets; in the mobile scenario this throughput dropped to 21%. The corresponding packet drop ratios were 33% and 85% respectively.

The mobility models most commonly used in MANET and WSN simulations typically do not reflect the situation nodes are exposed to in reality: most processes in nature stick to paths and group patterns (e. g. in the ZebraNet, the animals do not behave completely independent from each other) and are constrained by the environment (ocean currents, roads, . . .). Nevertheless most simulations that take mobility into consideration apply only the random waypoint mode or similar schemes. The choice of mobility models thus often constitutes an unreasonable assumption but has an impact on the performance of routing protocols. For an example of the improved success rates of CONFIDANT because a group-based mobility model was used, see [17].

7 Cross-Layer Design

Layering is a basic software design principle, applied e. g. in business applications (user interface, business logic layer, database layer) or in network stacks (see e. g. the ISO/OSI model). The layers introduce abstractions that make it more easy to design each layer independently and therefore exchangeable. Replacing e. g. a data link layer protocol by another should be ideally unnoted by the other layers, since each layers only depends on the correct implementation of the interfaces




Fig. 2. Comparison of Packet Delivery Ratio between DSR and CONFIDANT under the influence of mobility with no attack occurring (from [17])

to adjacent layers—the implementation is transparent. These advantages are bought by additional overheads.

An example for the application of this layered approach is the design of the WSN operating system TinyOS. Its architecture consists of components that are linked by interfaces. Hence, every component can be replaced as long as it implements the interface definition properly. This makes it possible to run the OS along with its applications on different sensor node hardware by replacing the hardware abstraction components. Another possible option is to replace the MAC protocol by a more energy-efficient one, or to replace the routing protocol. All this does not affect the application or the other layers above or below the exchanged building block.

Taking information from one layer and using violates this design principle. The borders between the layers get blurred and the easy replacement of the communications system’s building blocks is prevented. Although it is in general undesirable to take information from other layers into account due to the loss of generality and clarity this is sometimes done to improve performance or to facilitate certain features. This is called cross-layer optimization. There is always a trade-off between desired performance and the clarity of clearly defined communications layers (including the necessary resources which accompany the abstraction introduced by the layered models). We list a few examples of cross-layer optimization:

1. The routing layer may incorporate information from the transceiver chip to factor in data about the quality of the radio link, thus making better decisions about reliable and energy-saving links, so as to increase connectivity and battery lifetime of the network.
2. The reputation system may also have access to the transceiver chip. It makes a difference if a neighbouring node stopped routing packets for other nodes or just simply faded out of communication range because it was moving.
3. Sub-systems like the routing or reputation system may be specifically tailored to a concrete application instead of being designed for general usage (as in the internet). For example, in a sensor network used for monitoring forests for wildfires temperature (from the application layer) should be taken into account for routing decisions; the network should prepare for ‘hot’ nodes to be excluded soon because with high probability they are going to be lost to the fire.

8 The Way Forward?

‘Begin at the beginning and go on till you come to the end: then stop.’

(Through the Looking Glass, Lewis Carroll, 1871)

This quote captures an important principle of security engineering: when designing a security solution, one must first know the problem. Security requirements will depend on the application. The assets, their value, and the threats one has to defend against in a WSN are inherently specific to the purpose the WSN is being used for. Once the security requirements have been captured, we can start developing a solution. Often, the latter is the easier part.

Developing security solutions without a concrete reference problem gets us close to science fiction. Validating designs with simulations that make unrealistic assumptions about node movement and wireless communications gets us closer to science fiction. Making assumptions about potential threats that are inconsistent with assumptions about the nature of wireless sensor networks puts us squarely in the realm of science fiction.

Starting from the full set of standard assumptions makes it difficult to advance the state of the art in WSN security research. It is more promising to have a closer look at the characteristics of some specific application domain first. Take the example of a wireless network of sensor implants in the human body. The signals exchanged may be sensitive for medical and for privacy reasons. We had earlier cast doubt on the use of cryptography to protect traffic between sensor nodes, led by the observation that nodes are not tamper-resistant. To tamper with a sensor implant an attacker first has to tamper with the human body. Now it would be justified to rely on cryptography, and researchers can move on to the question of key management, which will have its own application specific idiosyncrasies.

As a second example, consider Smart Meter applications. Sensor nodes placed inside a Smart Meter have access to a continuous energy supply and may have

powerful microcontrollers. The standard assumption of battery life and computational resources as limiting factors would no longer be valid.

In the absence of concrete applications, research on reputation-based routing protocols will not make progress by tuning the algorithms for computing reputations and recommendations. There is no proper yardstick to decide which scheme performs better. Research on reputation-based routing protocols may make progress by investigating how liars might influence routing in a WSN. We can pose the question whether it is more advantageous for an attacker to take out a part of the network by jamming the signals (brute force) or by selectively removing nodes from the routing tables of their neighbours by spreading misleading recommendations. The advantage may be defined in terms of power consumption, or in terms of attack precision, or in the number of nodes that need to be compromised weighted by the cost of compromising an individual node. Validation by simulation would, of course, require a more sophisticated model of misbehaviour than dropping a fixed percentage of packets.

References

1. Buchegger, S., Le Boudec, J.Y.: Performance analysis of the CONFIDANT protocol: Cooperation of nodes - fairness in dynamic ad-hoc networks. In: MOBIHOC'02. pp. 226–236 (2002)
2. Buchegger, S., Le Boudec, J.Y.: Self-policing mobile ad-hoc networks by reputation systems. IEEE Communications Magazine pp. 101–107 (Jul 2005)
3. Buchegger, S., Mundinger, J., Le Boudec, J.Y.: Reputation systems for self-organized networks. IEEE Technology and Society Magazine pp. 41–47 (2008)
4. Delin, K., Jackson, S., Some, R.: Sensor webs. NASA Tech Brief 23 (1999)
5. Eger, K., Killat, U.: Resource pricing in Peer-to-Peer networks. IEEE Communications Letters 11(1), 82–84 (2007), <http://dl.comsoc.org/cocoon/comsoc/servlets/GetPublication?id=9019938>
6. Eger, K., Killat, U.: Bandwidth trading in BitTorrent-like P2P networks for content distribution. Computer Communications 31(2), 201–211 (2008)
7. Eichmann, J., Greßmann, B., Hackbarth, F., Klimek, H., Menrad, V., Meyerhoff, T., Pilsak, T., Sauff, H.: SomSeD – analysis of an experimental wireless sensor network. In: Proceedings of the Workshop Selbstorganisierende Sensor- und Datenfunknetze. pp. 11–17. Hamburg, Germany (Oct 2009)
8. Eschenauer, L., Gligor, V.: A key-management scheme for distributed sensor networks. In: Proceedings of 9th ACM Conference on Computer and Communications Security. pp. 41 – 47 (2002)
9. Ganeriwal, S., Srivastava, M.B.: Reputation-based framework for high integrity sensor networks. In: SASN'04 Proceedings of the 2nd ACM workshop on Security of ad hoc and sensor networks (2004)
10. Giannetsos, T., Dimitriou, T., Prasad, N.R.: Weaponizing wireless networks: An attack tool for launching attacks against sensor networks. In: Black Hat Technical Security Conference: Europe 2010 (2010)
11. Johnson, D.B.: Routing in ad hoc networks of mobile hosts. In: Proceedings of the Workshop on Mobile Computing Systems and Applications. pp. 158–163. IEEE Computer Society (Dec 1994)

12. Munding, J., Le Boudec, J.Y.: Analysis of a reputation system for mobile ad-hoc networks with liars. *Performance Evaluation* 65(3-4), 212–226 (2008)
13. Newsome, J., Shi, E., Song, D., Perrig, A.: The sybil attack in sensor networks: Analysis & defenses (2004)
14. Osterweil, L., Ghezzi, C., Kramer, F., Wolf, A.: Determining the impact of software engineering research on practice. *IEEE Computer* 41(3), 39 – 49 (2008)
15. Parno, B., Perrig, A., Gligor, V.: Distributed detection of node replication attacks in sensor networks. In: 2005 IEEE Symposium on Security and Privacy. pp. 49 –63 (2005)
16. Song, S.: Dynamic feed-back mechanisms in trust-based DSR. Master’s thesis, Informatics and Mathematical Modelling, Technical University of Denmark, DTU (2005), <http://www2.imm.dtu.dk/pubdb/p.php?3961>
17. Sun, J.: Analysis of Reputation-based Routing in Mobile Ad-hoc Networks. Master’s thesis, Hamburg University of Technology (2009)
18. Yu, B., Xiao, B.: Detecting selective forwarding attacks in wireless sensor networks. In: IPDPS’06 Proceedings of the 20th International Conference on Parallel and Distributed Processing. pp. 1–8. IEEE Computer Society (Apr 2006)