Distributed Actuation Attacks in Wireless Sensor Networks: Implications and Countermeasures

Alexandra Czarlinska and Deepa Kundur
Department of Electrical and Computer Engineering
Texas A&M University, College Station, TX, 77843-3128
czlinska@ece.tamu.edu, deepa@ece.tamu.edu

Abstract

This paper investigates the loss of sensing fidelity in a wireless sensor network resulting from a proposed novel attack. The active attack is carried out by a distributed malicious Sensor Actuator Network (mSAN) which is able to actuate or change sensed parameters of the surrounding environment under observation. We show how the attack effectively produces a Denial of Service on the Sensing (DoSS) of a legitimate network, causing it to observe and record false intelligence about the environment. We demonstrate how a controlled level of random mobility in the network counters the attack under various deployments, network densities and actuation radii. We conclude that a random uniform distribution may be most resilient against these attacks and that a strictly deterministic grid deployment may be most vulnerable under certain circumstances. In general we note that in physically hostile environments where sensing fidelity is important, node location becomes as sensitive for dependability purposes as encryption information.

1. Introduction

Wireless Sensor Networks (WSNs) are enabling a new era of connectivity between our physical and digital environments with unprecedented industrial, public and military applications. Some of these applications include smart spaces, structural monitoring of buildings, fire or biohazard detection, as well as surveillance and target tracking. To perform the desired task in these applications the WSN must be able to sense its environment correctly and further process and report this data. If the WSN does not sense its environment correctly and reports faulty data, the application or end user relying on this data will take incorrect actions. Serious safety and security problems may occur as a result. Many applications rely on WSNs for monitoring of the physical environment in order to provide some form of security (for example biohazard detection). The need for a high WSN sensing fidelity is of tremendous importance in all such monitoring and security-related applications.

Dependable sensing (referred to as high sensing fidelity) is often achieved by deploying a very large number of sensors scattered densely throughout the region of interest. Hence the aim is to achieve high sensing fidelity and sufficient coverage of the phenomenon of interest through redundancy. Unfortunately WSNs face a unique security challenge due to their deployment in potentially harsh or even hostile environments. Nodes may be lost due to battery exhaustion, component failures or sensor miscalibrations. Typically it is assumed that only a few nodes might be lost in any given region and that the high density protects specifically against this type of loss. If a very significant number of sensors are lost in a given region, reporting of data from that region is typically visibly disrupted. The failure is thus detected and more nodes can be sent to the area as needed. Hence it is assumed that the sensing fidelity in this scenario is dependable.

Unfortunately nodes can also be attacked by an enemy especially when deployed for security reasons. It is typically assumed that the attacker may actively try to capture k nodes out of the N nodes in the WSN. With the capture of a node, the attacker may gain access to some cryptographic keys used in the WSN for encryption, authentication and message integrity. The attacker may then launch a wide variety of attacks on the data in the WSN or on the routing and control data used to support the WSN. Although these security problems have not been completely solved, a large body of research exists to mitigate these attacks. We note three key assumptions often present in this body of security research. 1) Only $k << N$ nodes get captured. Under current attack models, this might be reasonable since if too many nodes get captured the attack might become visible. 2) The attacker uses or breaks the cryptographic keys in a node. 3) The attacker targets either the information inside the WSN or the control and routing data used to support the

WSN.

In our work we address the complementary problem of physical attacks at the *sensing level* before the parameters of the observation region are recorded by the nodes. This attack targets data before it enters the WSN and can be viewed as a Denial of Service on Sensing (DoSS). Consequently the mechanisms devised to protect against attacks on the data *inside* the WSN as well as on routing and control data are ineffective. The attack results in a loss of WSN dependability while not requiring the physical capture of any nodes and while being immune to cryptographic strategies. Furthermore the attack is distributed in nature, potentially allowing any number of nodes to become affected within an area.

To alter the sensing of a legitimate network in such a manner, the attack is carried out by a malicious Sensor Actuator Network (mSAN for short) which is distributed either randomly or deterministically throughout the same physical space as the legitimate Sensor Actuator Network (lSAN for short). Examples of actuation which may be used for such attacks include but are not limited to, mobility, the release of chemicals, activation of fans to dissipate surrounding air or air-born agents, and the creation of noise for radio-jamming purposes. In order to carry out the attack successfully the mSAN may use various deployment distributions, actuation radii and any a priori knowledge of the location of the lSAN nodes. In order to remain undetected before the attack, the mSAN may deploy in the hostile region before or concurrently with the lSAN. Regardless of the competitive deployment strategy and evasion techniques used, the goal of the mSAN is to cause a DoSS on the lSAN. In this work we show that while actuation in sensor networks may be used for such malicious purposes, it can interestingly be also used by the legitimate network as an effective defense. We demonstrate this strategy of "fighting fire with fire" by endowing the legitimate network with mobility (as a form of actuation) and studying how a controlled amount of mobility counters the effects of malicious actuation.

2. Related work

The study of actuation as a method of attack and countermeasure bridges a number of different research areas such as sensor network security, coverage, localization, mobility and actuation. We briefly review the most salient results from these areas as applicable to the DoSS attack.

2.1. Sensor network security

Figures 1 and 2 show the relationship of the actuation attack to other forms of active attacks found in the field of sensor network security. We observe that the actuation attack occurs before the phenomenon P under observation is sensed and recorded by the lSAN. This happens when the

naturally noisy version of the phenomenon P, given by $\tilde{P}$, propagates through the environment and is actuated or altered by the mSAN to become an observable O. This possibly altered observable O is recorded by the lSAN as the true phenomenon and processed internally to produce some datum D. In contrast, other active attacks on sensor networks usually target data D flowing inside the network, or even control and routing data. References [2] and [5] provide comprehensive discussions of recent attacks and countermeasures while reference [24] focuses specifically on Denial of Service (DoS) attacks in sensor networks. A critical factor in sensor network security is the issue of *physical* vulnerability of the nodes deployed in an unattended and possibly hostile environment which poses extra security challenges that have not been fully addressed to date [16], [3], [23]. The actuation attack is a type of DoS but it affects the network at this "physical" or sensing level which current countermeasures do not address. This is illustrated in Figure 3 which shows traditional communication of data over a non-secure channel and in Figure 4 which shows "communication" of a phenomenon over a non-secure channel. In the traditional case, if we wish to transmit a datum D over a non-secure channel we can rely on encryption. In the case of a sensor network deployed to monitor a phenomenon P, the phenomenon travels through a hostile channel before it is observed and recorded. Clearly no cryptographic protection of the raw phenomenon is possible before it arrives at the sensor network. We are thus forced to consider another approach to securing sensing fidelity in the face of actuation attacks in a hostile environment.
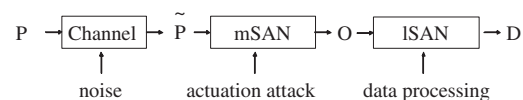

Figure 1. Flow of information during an actuation attack

2.2. Coverage and location uncertainty

The exact definition of coverage varies depending on the specific application and on the toolsets used to address it. Generally speaking however, coverage is a measure of how well the sensor network covers or observes all the points of a physically distributed phenomenon. In [14], [15], [25] and [4] the authors formulate the Best and Worst case coverage scenarios by calculating a path of Maximum Support and Maximum Breach for an object moving through the sensor field. References [8], [6], [9], [10] and [11] present other key results in coverage. In reference [12] the authors use
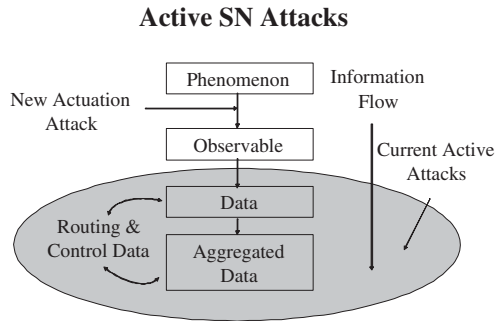
Active SN Attacks
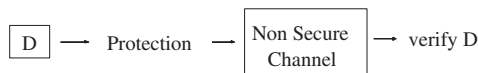

Figure 2. Actuation vs. other attacks

Figure 3. Traditional communication over a non-secure channel
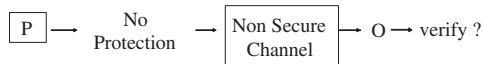

Figure 4. Phenomenon communication over a hostile channel

percolation theory to study the sensor network density required to achieve detection of a target with probability 1 almost surely. In [26] the authors consider the problem of coverage in the face of uncertainty in the sensor locations. In [20] the authors provide local algorithms for location discovery and coverage. This research is critical for understanding how an mSAN might find a path of least detection through the environment and how it might "cover" the lSAN nodes in the face of uncertainty of their locations.

2.3. Exposure and detection avoidance

In [19] the authors formulate exposure as a measure of how well an object moving on an arbitrary path can be observed by the sensor network over a period of time. The authors present an efficient algorithm for finding minimal exposure paths for the object to move along, which also simultaneously provides information about the worst case coverage of the sensor network. Simulation results show that for generally sparse fields with a random *uniform* spatial deployment, there exist many minimal exposure paths. Simulations also indicate that in general, deterministic deployments are better able to detect a moving object than purely random deployments. We will later compare these conclusions about various deployments with our results on the role of deployment in the face of actuation attacks. The authors also present a generalized sensing model of interest to the study of actuation.

References [1] and [7] provide an approach that allows a stealthy traverse through an *unknown* environment that contains dynamic objects and an observer. The key is to exploit the dynamic objects in the environment as they become known and use their shadow as cover to move undetected from an initial location to a target location. The observer is assumed to have infinite observational range in all directions. The traversing robots are assumed to also have omni-directional sensing but for finite ranges. Simulations and implementation results show that 100% stealth can be achieved at a tradeoff of taking a route which is 86% efficient compared with a direct route which is 100% efficient but only 36% stealthy. These studies of exposure and detection avoidance are key to understanding how an intruding mSAN can deploy in an environment undetected.

2.4. Actuation

We define actuation in sensor networks as the ability of a node to act upon, change or influence its environment using limited energy. The latter requirement is is in contrast with robotic actuation where the robot typically has access to a much larger battery or wired source of energy. The small size (especially height) of the node and of its components further restricts the type and range of actuation that it may perform, in contrast with much larger robots. The energy and size limitations imply that sensor nodes should employ distributed actuation to limit energy use while having a global effect on the environment.

Sensor actuation includes but is not limited to actions such as turning on external fans (possibly to disperse heat, chemicals or biological agents) and moving across the landscape (thereby re-shaping the topology of the environment). Thus far actuation in the sensor network literature has been mostly limited to studies of mobility.

In [18] the authors explore how mobility can be used by a sensor network as a type of actuation to repair its own coverage (called self-repair). In [17] the authors examine how mobile nodes can migrate to areas of high energy (solar for example) to charge themselves and then charge other starving nodes. In [22] the authors discuss how mobility can specifically help sensor network security by detecting misbehaving nodes. Reference [13] introduces the idea of parasitic mobility where nodes are able to catch a ride on any moving object and dislodge from it using an actuator. Hence adding actuation (including mobility) to sensor networks significantly expands their autonomy and fault-tolerance. In Section 3 we argue that mobility is a viable and effec-

COMPUTER
SOCIETY

tive countermeasure against attacks caused by other forms of actuation.

3. Problem formulation

3.1. General setup

This paper addresses DoSS attacks and countermeasures within the following framework. We consider a legitimate Sensor Actuator Network (lSAN) with N nodes deployed either deterministically or randomly throughout a finite physical region (which we call the "environment") to monitor a spatially distributed phenomenon of interest. A malicious Sensor Actuator Network (mSAN) is deployed either deterministically or randomly throughout the same environment with M nodes. We define the ratio of the number of mSAN nodes to the number of lSAN nodes as the Flooding Ratio $F.R$ where $F.R = M/N$.

3.2. Phenomenon and information flow

The lSAN is deployed in the environment to monitor a phenomenon of interest which can either be a point source phenomenon, such as a moving target, or a distributed phenomenon such as a temperature field. In our work we focus on spatially distributed phenomena. In order to facilitate studies of actuation, we distinguish between three levels of information that exist between the phenomenon of interest and a sensor network as shown in Figure 1. Let P denote the phenomenon of interest as it occurs in the environment. Let $\tilde{P}$ denote a possibly noisy version of P as it propagates through the environment from its source. Subsequently let O denote the observable that is sensed and recorded by a sensor node. When an actuation attack occurs O differs from $\tilde{P}$. Let D denote the data produced internally by a sensor node through the internal processing of O (such as averaging for instance). In general, P, $\tilde{P}$, and O are modeled as random processes and when time-sampled appropriately, can be treated as random variables.

3.3. Sensing model

We extend the general sensing model proposed in [19]. Equation 1 presents the general sensing model S of a node i located at l_i monitoring a point p in the environment at time t and distance d away. The parameters $\lambda_i(t)$ and $k_i(t)$ are technology-dependent and are generally allowed to vary with time due to errors and miscalibrations.

$$S_i(p, t) = \frac{\lambda_i(t)}{[d(l_i(t), p(t))]^{k_i(t)}} \qquad (1)$$

We note that when $d = 0$ (node i is taking readings of $\tilde{P}$ at its own location l_i) this model produces infinite sensing.

We propose a modified sensing model as shown in Equation 2 that produces finite sensing at $d = 0$. This omni-directional model diminishes exponentially with distance and the sharpness of this decay can be controlled through the parameter γ to resemble Equation 1 if desired. For simplicity we set $\lambda_i(t) = 1 \ \forall \ i$ and $\forall \ t$, set $\gamma = 1$ and restrict the sensing range d as shown.

$$
\begin{aligned}
S_i(p, t) &= \lambda_i(t) e^{-\gamma d(l_i(t), p(t))} \\
&= \begin{cases} e^{-d(l_i(t), p(t))} & \text{if } 0 \leq d \leq d_{max} \\ 0 & d > d_{max} \end{cases} \qquad (2)
\end{aligned}
$$

3.4. Deployment and detection assumptions

For an actuation attack to proceed, the mSAN should be present in the environment without first being detected by the lSAN. To achieve this the mSAN might deploy in the hostile environment before the lSAN, or it might deploy alongside the lSAN before the latter establishes its infrastructure and begins monitoring. Furthermore as discussed in Section 2.3, work in detection avoidance provides certain algorithms for moving through a sensor network undetected. We must also consider the fact that most detection algorithms are designed for 2D environments and that the optimal placement of surveillance in the 3D case has been shown to be NP-complete [14]. For realistic surveillance applications the lSAN is deployed in a 3D environment where opponents can hide in valleys, behind bushes, employ camouflage and move around to create network topology changes. We also note that mSAN nodes deployed by a reasonable attacker would most likely not be physically distinguishable from lSAN nodes (ie: visual surveillance through the use of camera sensors would not be sufficient) and that these nodes would most likely employ spread-spectrum techniques in their communications. Hence in most cases we cannot conclude that a hostile environment under the presence of an lSAN is free of the presence of a possibly actuating mSAN.

3.5. Mobility model

Given the limited computational abilities of the lSAN, we assume a simple mobility model where each lSAN node is capable of moving in a randomly chosen direction at a set speed and for a set duration of time determined in part by its available energy. We introduce a Mobility Threshold $M.T$ ranging from 0 to 1, which is defined as the minimum change in the observable O as sensed by a node i in order for node i to move. For instance, if we set $M.T = 0.3$ then node i will only move if it records a 30% change in the observable from the previous time instance.

COMPUTER SOCIETY

3.6. Actuation attack

We propose one of the simplest possible models of an actuation attack. The actuation carried out by an mSAN node j, is a point source phenomenon of amplitude A_j which we set to 1 for $\forall\, t$ and $\forall\, j$ for simplicity. The actuation begins at some arbitrary time t_0 at which point it is only present at node j's location denoted by l_j. The spatial propagation of the phenomenon is modeled by a decaying exponential given by Equation 3.

$$
\begin{aligned}
A_j(p,t) &= A_j(t)e^{-d(p(t),l_j(t))} \\
&= 1 \cdot e^{-d(p(t),l_j)}, \\
&\quad t_0 \le t \le t_F, \quad 0 \le d \le d_{max}, \forall j \quad (3)
\end{aligned}
$$

$A_j(p,t)$ denotes the actuation effect of node j at spatial point p at time t. The spatial point p can be for instance the location of an lSAN node and is in general allowed to be time-varying (mobile lSAN nodes) while for simplicity we assume that the mSAN nodes are stationary. We note that the actuation effect is negligible outside of the specified distance and time range. We assume that the only actuation performed by the lSAN is mobility while the only actuation performed by the mSAN is phenomenon actuation as described by Equation 3. We assume that the mSAN nodes do not dynamically coordinate with each other during the actuation but rather are programmed to start and continue actuating for a specified time. We also assume that in general the lSAN does not have an internal model of the actuation or its parameters.

A point p in the environment contains the possibly noisy version of the phenomenon $\tilde{P}$ and it might come under the actuation influence of several mSAN nodes. Hence the strength of the sensed Field at any point p and time t obtained through this superposition is given by Equation 4:

$$
F(p,t) = \Sigma_{j \in M} A_j(p,t) + \tilde{P}(p,t) \quad (4)
$$

Specifically if the point p is the location of a mobile lSAN node i and if the actuation attack is as given in Equation 3 then:

$$
\begin{aligned}
F(i,t) &= \Sigma_{j \in M} A_j(l_i,t) + \tilde{P}(l_i,t) \\
&= \Sigma_{j \in M} 1 \cdot e^{-d(l_i,l_j)} + \tilde{P}(l_i,t) \quad (5)
\end{aligned}
$$

where the distance d and the time t are constrained as stated earlier. Given a random spatial deployment of the mSAN nodes, F(p,t) will vary throughout the environment with p and t.

3.7. Energy considerations

Among the various resource constraints in WSNs, the energy constraint is considered one of the most serious and restrictive. It is of particular importance in the case of actuation, where malicious nodes use their energy not only to sense but also to act upon the environment. The feasibility of such an attack based on energy constraints must be considered.

We argue that energy constraints do not prevent an actuation attack from happening for several reasons. 1) While a legitimate WSN is expected to minimize its energy expenditure to ensure longevity of operation, the goal of the malicious network may be a direct short-lived attack after which the mSAN will stop operating. In this context the mSAN can afford to expand its energy in a collective effort to cause a DoSS in the lSAN. 2) The attack is distributed and hence each attacking node needs to contribute only a fraction of the overall required energy. 3) Research in sensor network energy suggests that nodes may harvest or replenish their energy from the environment [17]. 4) Certain types of actuation may require relatively little energy, such as transmitter and sensor jamming through noise generation. Furthermore, the amount of energy required by each node is proportional to the duration of actuation Δt and on the change in the phenomenon ΔP desired. For small ΔP and Δt, the amount of required energy may be small.

4. Countermeasures against actuation attacks

The goal of our countermeasure strategy shown in Figure 5 is to provide some copies of $\tilde{P}$ to the lSAN given that an actuation is occurring and that the observable O is generally not the same as $\tilde{P}$. Since the lSAN receives both copies of $\tilde{P}$ and of O, ideally we want to receive $k \ge N/2$ copies of $\tilde{P}$. In assessing the success of a countermeasure, we define the Average Sensing Error E, the Percent Improvement in Average Sensing Fidelity PI and the Percent lSAN nodes affected PA as follows:

$$
E(t) = \frac{1}{N}\Sigma_{i \in N} E_i(t) \quad (6)
$$

$$
PI = \frac{E(t_0) - E(t_f)}{E(t_0)} \cdot 100\% \quad (7)
$$

$$
PA(t) = AN(t)/N \cdot 100\% \quad (8)
$$

where AN is the Number of Affected nodes. In what follows we briefly examine two candidate approaches for a countermeasure : 1 - internal data processing (requiring some side information) and 2 - mobility-based approaches (requiring less side information). We argue that given limited side information, mobility-based approaches using the
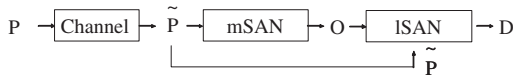
Figure 5. To mitigate an actuation attack the lSAN needs access to $\tilde{P}$
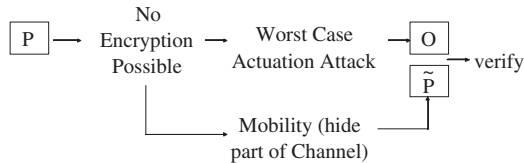

Figure 6. Use of mobility to hide part of the channel (akin to steganography)

ideology of "fighting fire with fire" are able to meet the countermeasure goal and hold significant potential for dependability against actuation attacks.

4.1. Internal data processing

Upon receiving N copies of the observable $O(t_k)$ at arbitrary time t_k (N is the number of lSAN nodes), an aggregator node may do internal data processing using a correction function c. The aim is to obtain a datum D that provides more "information content" about $\tilde{P}$ than we get from the Observables O alone. Let the entropy $H(X)$ be the information content of a random variable X. Hence the information content of $P(t_k)$ is simply $H(P(t_k))$.

Proposition 1 *Let $t = t_k$ be an arbitrary time instant. Let c be a deterministic correction function such that c: $\mathbb{R}^N \mapsto \mathbb{R}^N$ where $c(a(\tilde{P}(t = t_k))) = D(t_k)$ and where a is an arbitrary deterministic actuation function. Then $H(D(t_k)) \leq H(\tilde{P}(t_k))$ for all choices of c.*

Proof 1 *For all deterministic functions g and random variables X, $H(g(X)) \leq H(X)$ as in Reference [21]. Since $0(t_k) = a(\tilde{P}(t_k))$ we have that $H(O(t_k)) \leq H(\tilde{P}(t_k))$ and since $D(t_k) = c(a(\tilde{P}(t_k))$ then $H(D(t_k)) \leq H(\tilde{P}(t_k))$.*

Of course in general if we possess some side information such as an internal model of the attack a or some of its parameters, we are able to provide an estimate of $\tilde{P}(t_k)$ through data processing. In this work we assume the worst case scenario where the lSAN does not possess such a model and may not even be aware of the presence of an actuation mSAN, yet still requires high fidelity in the sensed observables.

We may be tempted to exploit the correlation between the sensed observables of the densely deployed nodes in order to detect actuation. This could be accomplished by having each lSAN node in a region send its sensed observable to an aggregator node and by applying a majority vote (or another rule). We note however that for certain deployments of the mSAN and lSAN nodes, the majority of the lSAN nodes in an area could be affected by the attack. Furthermore this technique requires communication between the nodes (which is generally costly) and requires a delay between observation and verification. This delay could be problematic for security applications that require timely detection (such as fires or biohazards).

4.2. Proposed mobility based techniques

Given that any spatial deployment (either deterministic or stochastic) of the lSAN and mSAN nodes is allowed, any number k of lSAN nodes may be affected by the actuation, where $k < N$ but possibly as large as $k \sim N$. This is particularly detrimental in the special case where the mSAN knows the locations of the lSAN nodes. Our countermeasure goal is to obtain a majority of $\tilde{P}$s rather than actuated Os. We argue that mobility of the lSAN nodes achieves this goal without directly having to detect the presence of any mSAN nodes and without the need to communicate and aggregate values.

To understand the potential effect of mobility in reducing sensing error, we consider the worst case attack where the mSAN is able to position itself so as to cover every lSAN node in the area. If the lSAN nodes simply compare their sensed observables or take a majority vote, they will find a high correlation among the set of Os and conclude that there is no attack. Hence data processing without additional side information yields an incorrect decision. However if some of the lSAN nodes move in a manner that is unpredictable to the mSAN, the latter will no longer be able to perfectly cover all of the lSAN nodes. As such, some of the lSAN nodes will move out of the region of actuation and collect true readings. This leads us to propose the following mobility based technique which makes use of a Mobility Threshold.

Upon sensing of P, each node independently determines if the change in P from the previous recorded value is sufficiently large (larger than the threshold $M.T$). If $M.T$ is exceeded, the node moves randomly to a new location while it continues sensing. As each node performs this operation, some of the nodes move outside of the actuation range of their neighboring mSAN nodes. In essence, although P cannot be encrypted as it travels over the non-secure channel, the channel from P to some lSAN nodes can be hidden in a stegonographic way through mobility. This is depicted in Figure 6. In this sense mobility corrects the Observable

IEEE
COMPUTER
SOCIETY

O without having to detect the presence and location of the attacker.

Mobility does expand energy, hence it is important to ask how much mobility is required to improve the sensing fidelity. Use of the Mobility Threshold $M.T$ (ranging from 0 to 1) allows a trade-off between sensing fidelity and the level of mobility. For simplicity we assume that all lSAN nodes have the same $M.T$. If the $M.T = 0$, all nodes will move as soon as there is any change in what is being sensed in order to verify it. On the other hand if $M.T = 0.6$, the sensed phenomenon will only be checked if the change exceeds 60%. In addition to not requiring extra node communication, this countermeasure also offers several side-benefits such as improving other security measures and allowing coverage repair as mentioned in Section 2.4.

4.3. Significance of node location and random mobility

Some of the most damaging active attacks on WSNs in current literature assume that the attacker is able to capture a number of nodes and obtain their cryptographic keys. In the case of an actuation attack the distributed attacker (mSAN) does not capture nodes or their keys. The success of the attack depends largely on the mSAN's ability to distribute itself correctly around the lSAN nodes. In this context the *location* of the lSAN nodes becomes the "secret key" which we wish to hide or at least render unpredictable for the attacker.

Results from [19] indicate that on average a WSN is better able to track a target moving through the field (have better exposure) in deterministic deployments than in random ones. Hence a deterministic deployment increases the coverage *service* that the WSN is able to *provide*. However a grid deployment renders the WSN more vulnerable to attacks against the WSN *itself*, hence placing the sensing fidelity at risk. This can be easily seen from the fact that given a deterministic grid deployment, knowledge of some node locations allows the complete determination of the remaining node locations. Hence in physically hostile environments location information holds special relevance for security and dependability. Figure 8a) shows simulation results for a deterministic lSAN deployment. We note that the mSAN nodes fully cover every lSAN node and through actuation, completely destroy its sensing fidelity.

5. Experimental results and insights

In this section we present simulation results to study the effect of deployment distributions, actuation radii and Flooding Ratios on the sensing error. We also study the effects of mobility under various Mobility Thresholds in reducing the sensing error. Specifically we focus on the

Gaussian and Uniform distributions with the former chosen to represent the family of related exponential distributions (Laplace, Exponential etc.) and with the latter as a distribution not fitting this family.

The simulations were performed using a Matlab-based simulator that was developed. In the simulations the travel distance of the lSAN nodes was taken as 3 spatial units/time interval for simplicity while the angle was chosen randomly according to a uniform distribution ranging from $[0, 2\pi]$. The simulation time interval was set to 300 time units. The sensing range of each lSAN node was taken as 3 spatial units while the actuation range $A.R$ of the mSAN was varied from 1 to 3 spatial units. To generate a (x, y) coordinate according to a uniform distribution, the x and y coordinates were each chosen from a uniform distribution. This method was also used for Gaussian distributions. The size of the simulation world was set to 50 x 50 units^2. The number of lSAN nodes was held constant at 300 while the number of mSAN nodes was varied from 1 to 600 to obtain various Flooding Ratios. Without loss of generality the phenomenon P was set to 0 in Equation 5 and hence any actuation by the mSAN recorded by the lSAN as an observable constituted an error. The Average Sensing Error of the lSAN was computed using Equation 6. The lSAN sensing was implemented using Equation 2, the actuation was implemented using Equation 3 and the superposition from actuating nodes was obtained using Equation 5. The figures of merit to evaluate the success of our DoSS countermeasures are: 1: Average Sensing Error E given by Equation 6, 2: Percent Improvement in Average Sensing Fidelity PI given by Equation 7 and 3: Percent lSAN nodes affected given by Equation 8.

First we would like to show how different deployments affect the severity of an actuation attack and how mobility helps regardless of the deployment used. Figure 7a) shows a typical uniform deployment of both networks with 300 lSAN nodes (circles) and 300 mSAN nodes (diamonds) giving a Flooding Ratio $F.R = 1$. The Actuation Radius $A.R = 2$ of each mSAN node is shown as a large open circle and lSAN nodes affected by one or more mSAN nodes are shown as filled circles. Figure 7b) depicts a typical Gaussian deployment for both networks where the lSAN is deployed with arbitrary mean $\mu_x = 25$, $\mu_y = 40$ and variance $\sigma_x^2 = 20$, $\sigma_y^2 = 20$. The mSAN is also deployed with a Gaussian distribution but given by $\mu_x = 40$, $\mu_y = 40$ and $\sigma_x^2 = 20$, $\sigma_y^2 = 20$ with $A.R = 2$. Figure 8a) shows a grid distribution for both networks and the corresponding actuation effects. From these three typical deployments we note the large number of affected lSAN nodes when no countermeasures are in place, with the most severe effect occurring for a deterministic grid deployment.

In a worst case attack the mSAN knows the distribution with which the lSAN is deployed (though the exact param-
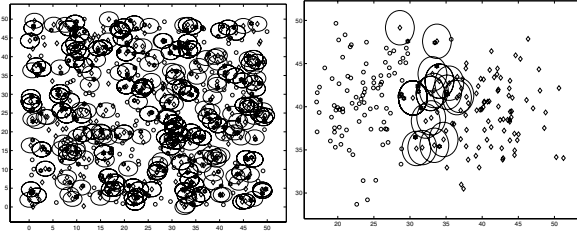
Figure 7. (a) 300 mSAN nodes (diamonds) vs. 300 lSAN nodes (circles) both with a Uniform distribution with $FR = 1$ **and** $A.R = 2$**. (b) 100 mSAN nodes vs. 100 lSAN nodes with a Gaussian with** $F.R = 1$ **and** $A.R = 2$**.**
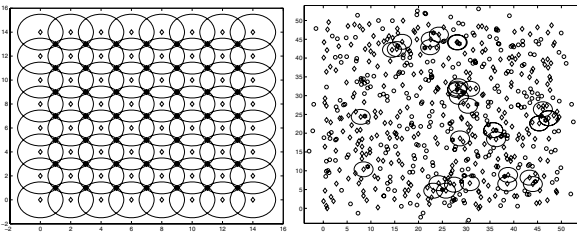

Figure 8. (a) 64 mSAN nodes vs. 49 lSAN nodes with a grid pattern and $A.R = 1.5$**. (b) Reduction in actuation attack after mobility with** $M.T = 0$ **for** $F.R = 1$ **for a Uniform distribution.**

mSAN knows the parameters well, that the attack is most severe (second only to a grid deployment).
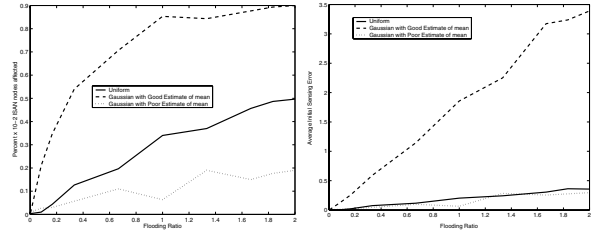

Figure 9. (a) Initial Percent x 10^{-2} **of affected nodes for various distributions and** $AR = 1$**. (b) Average Initial Sensing Error for Various Distributions with** $AR = 1$**.**
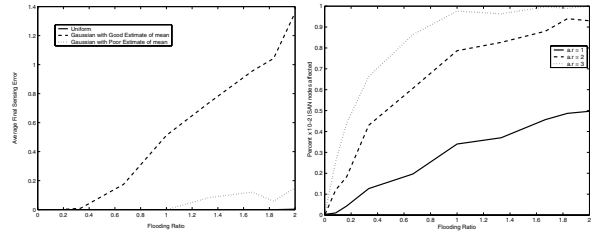

Figure 10. (a) Average Final Sensing Error for Various Distributions with $AR = 1$ $M.T = 0$**. (b) Percent x** 10^{-2} **initially affected nodes (before mobility) for a uniform distribution for various radii of actuation and various Flooding Ratios.**
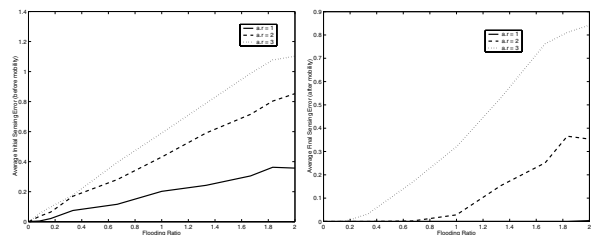

Figure 11. (a) Average Initial Error in Sensing Fidelity before mobility. (b) Average Final Sensing Error using mobility of $M.T = 0$**.**

eters may not be known as in the Gaussian example given above) and deploys with the same distribution. Figure 9 shows the resulting Percent (x 10^{-2}) nodes that are affected by the attack and the resulting Average Initial Sensing Error (before any countermeasures) for various Flooding Ratios. These results are shown for the case when 1-both networks deploy using a Uniform distribution, 2-both networks deploy using a Gaussian distribution and the mSAN has a good estimate of the lSAN distribution parameters and 3-where both deploy using a Gaussian distribution but the mSAN has a poor estimate. For cases 2 and 3 the lSAN was assigned the parameters mentioned earlier for a Gaussian. For case 2 the mSAN distribution was set as $\mu_x = 23$, $\mu_y = 38$ with $\sigma_x^2 = 20$, $\sigma_y^2 = 20$. For case 3 it was set as $\mu_x = 15$, $\mu_y = 15$ with $\sigma_x^2 = 20$, $\sigma_y^2 = 20$. Figure 10a) shows the effect of mobility with $M.T = 0$ in reducing the sensing error while Figure 8b) shows a typical reduction in the number of affected nodes in the case of a Uniform deployment. We see that the Average Final Sensing Error is decreased dramatically due to mobility and reduced by 100% in the Uniform case even when the Flooding Ratio is as high as 2. We also note that when the two networks are deployed according to a Gaussian distribution and when the

Next we would like to explore the effect of various actuation radii, where the larger the actuation radius of each node, the more spatially powerful the attack. Figure 10b) shows the percent (x 10^{-2}) of lSAN nodes affected for $A.R = 1$, $A.R = 2$ and $A.R = 3$ when deployed in a Uniform distribution. Figure 11a) shows the correspond-

COMPUTER SOCIETY

ing Average Initial Error in Sensing Fidelity before mobility and Figure 11b) shows the Average Final Sensing Error using mobility of $M.T = 0$. We conclude that in the case of $A.R = 1$, mobility reduces the average sensing error by as much as 100%. For a stronger attack of $A.R = 3$, the reduction is on the order of 20% for this choice of MT. We observe that the sensing fidelity degrades in a seemingly linear way with increasing Flooding Ratio regardless of the $A.R$ used. Mobility on the other hand appears to improve the sensing fidelity in a nonlinear way with respect to Flooding Ratio. Importantly, in the case of $A.R = 2$ (there are 2 malicious nodes for every legitimate node), mobility still improves the sensing fidelity for all Flooding Ratios.

Next we examine the trade-off between reducing the sensing error and the level of lSAN mobility required as controlled by the Mobility Threshold $M.T$. We note that the lower the $M.T$ the larger the percent of nodes that move at any time instant, hence a lower $M.T$ indicates higher mobility. Figure 12a) shows the improvement in sensing when $M.T = 0.3$ is used for a Uniform distribution for various actuation radii, and Figure 12b) shows these results for $M.T = 1$. We conclude that although in general a $M.T = 0$ provides best results (reducing the sensing error to 0% in the case of $A.R = 1$ for instance), that such a high mobility is not always required.

More specifically Tables 1, 2 and 3 show the trade-off between the *maximum* percent of nodes that move during any time instant within the simulation interval $t_0 \leq t \leq t_f$, and the average percent of nodes still affected by actuation at time t_f. Table 1 corresponds to $F.R = 1/3$, Table 2 to $F.R = 1$ and Table 3 to $F.R = 2$. Each entry (A, B) in a table corresponds to ($\%$ nodes affected after mobility, $\%$ max nodes move). Cells shaded in grey indicate cases where the countermeasure goal was fully achieved with a majority of lSAN nodes completely free of mSAN actuation. For the remaining cases we note two critical points: 1 - although most nodes remain affected, at least one clean copy of $\tilde{P}$ exists (in most cases dozens of copies exist). Hence the lSAN has at least one record of what really occurred in the environment at a particular instant. 2 - Although most nodes are still affected at the end of the time interval, most have *reduced* their sensing error by moving to a less affected area. The tables count *all* affected nodes even if the sensing error at those nodes is small.

6. Future research directions

This work presents one of the simplest actuation and mobility models in order to study the effects of actuation on sensing fidelity in the presence of DoSS attacks. Future research will examine the effect of the attack if the mSAN nodes are allowed mobility and cooperation, and will examine the effects of different actuation and mobility models.
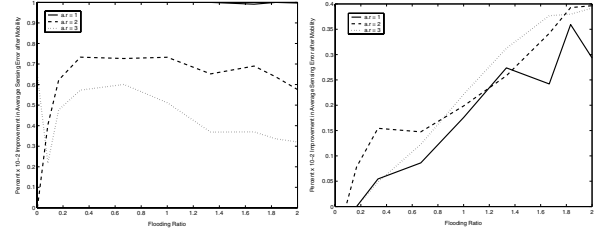

Figure 12. (a) Percent x 10^{-2} Improvement in Sensing after Mobility for a Uniform Distribution for various Actuation Radii and $M.T = 0.3$. (b) $M.T = 1$

Table 1. $F.R = 1/3$

	M.T = 0	M.T = 0.3	M.T = 0.6	M.T = 1
A.R = 1	(0.00, 43.0)	(18.3, 17.3)	(35.0, 1.00)	(38.6, 0.20)
A.R = 2	(0.00, 12.7)	(0.00, 13.3)	(11.6, 0.20)	(13.7, 0.30)
A.R = 3	(9.00, 66.3)	(61.3, 23.0)	(62.6, 8.00)	(68.3, 1.30)

Table 2. $F.R = 1$

	M.T = 0	M.T = 0.3	M.T = 0.6	M.T = 1
A.R = 1	(0.00, 34)	(0.00, 25.3)	(22.7, 0.90)	(26.6, 0.20)
A.R = 2	(6.30, 78.0)	(47.3, 53.0)	(64.3, 25.3)	(72.6, 0.80)
A.R = 3	(57.3, 97.6)	(82.6, 47.0)	(93.0, 25.3)	(96.0, 11.3)

Table 3. $F.R = 2$

	M.T = 0	M.T = 0.3	M.T = 0.6	M.T = 1
A.R = 1	(0.00, 50.9)	(0.00, 50.7)	(35.5, 27.6)	(50.6, 0.80)
A.R = 2	(38.0, 93.0)	(69.0, 71.0)	(84.6, 57.0)	(91.7, 27.0)
A.R = 3	(75.3, 100)	(95.0, 91.3)	(96.0, 71.6)	(99.0, 40.6)

7. Conclusions

We introduce a simple actuation attack and study the resulting loss in sensing fidelity (or sensing error) for a number of different deployments, actuation radii and flooding

ratios. We study mobility as a countermeasure for this attack using various Mobility Thresholds to trade-off between sensing fidelity and required mobility. We conclude that a uniform random deployment for the lSAN is most resilient to actuation attacks and argue that a deterministic grid deployment may be least resilient. We conclude that mobility reduces the average sensing error of the lSAN under a variety of conditions, even under heavy mSAN node flooding and large actuation radii.

References

[1] G. S. A. Tews, M. J. Mataric. Avoiding detection in a dynamic environment. *IEEE RSJ International Conference on Intelligent Robots and Systems*, 4:3773–3778, September-October 2004.

[2] L. Buttyan and J.-P. Hubaux. Report on a working session on security in wireless ad hoc networks. *Mobile Computing and Communications Review*, (4), 2002.

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

[4] L. L. D.P. Mehta, M.A. Lopez. Optimal coverage paths in ad-hoc sensor networks. In *IEEE International Conference on Communications*, volume 1, pages 507–511, May 2003.

[5] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proc. ACM/IEEE International Conference on Mobile Computing and Networking*, pages 263–270, August 1999.

[6] S. P. G. Kesidis, T. Konstantopoulos. Surveillance coverage of sensor networks under a random mobility strategy. In *Proceedings of IEEE Sensors*, volume 2, pages 961–965, October 2003.

[7] A. Howard, J. J. Mataric, and G. S. Sukhatme. An incremental self-deployment algorithm for mobile sensor networks. *Autonomous Robots, Special Issue on Intelligent Embedded Systems*, 13(2):113–126, September 2002.

[8] J. R. F. Z. J. Liu, X. Koutsoukos. Sensing field: Coverage characterization in distributed sensor networks. In *IEEE ICASSP*, volume 5, pages 173–176, April 2003.

[9] H. Q. E. C. K. Chakrabarty, S. S. Iyengar. Grid coverage for surveillance and target location in distributed sensor networks. *IEEE Transactions On Computers*, 51(12):1448–1453, December 2002.

[10] S. I. K. Chakrabaty. Sensor placement in distributed sensor networks using a coding theory, framework. page 157, June 2001.

[11] R. Kannan, S. Sarangi, S. Ray, and S. S. Iyengar. Minimal sensor integrity: Measuring the vulnerability of sensor deployments. *Information Processing Letters*, 86(1):49–55, April 2003.

[12] D. T. L. Benyuan. On the coverage and detectability of large-scale wireless sensor networks. 2003.

[13] J. P. M. Laibowitz. Parasitic mobility in dynamically distributed sensor networks. *Pervasive Computing: Third International Conference*, 3468:255–278, May 2005.

[14] S. Meguerdichian, F. Koushanfar, M. PotKonjak, and M. B. Srivastava. Coverage problems in wireless ad-hoc sensor networks. In *Proceedings of IEEE Infcom*, volume 3, pages 1380–1387, Anchorage, Ak, April 2001.

[15] S. Meguerdichian, F. Koushanfar, M. PotKonjak, and M. B. Srivastava. Worst and best-case coverage in sensor networks. In *IEEE Transactions On Mobile Computing*, volume 4, pages 84–92, January/February 2005.

[16] A. Perrig, J. Stankovic, and D. Wagner. Security in wireless sensor networks. *Communications of the ACM*, 47(6):53–57, June 2004.

[17] M. Rahimi, H. Shah, G. S. Sukhatme, J. Heideman, and D. Estrin. Studying the feasibility of energy harvesting in a mobile sensor network. In *Proc. IEEE International Conference on Robotics and Automation*, volume 1, pages 19–24, Taipei, Taiwan, September 2003.

[18] M. S. S. Ganeriwal, A. Kansal. In *IEEE Proceedings of International Conference on Robotics and Automation*, pages 5244–5249, April.

[19] G. Q. M. P. S. Meguerdichian, F. Koushanfar. Exposure in wireless ad hoc sensor networks. In *Proceedings of the Annual International Conference on Mobile Computing and Networking*, pages 139–150, Rome, July 2001.

[20] V. K. M. P. S. Meguerdichian, S. Slijepcevic. Localized algorithms in wireless ad-hoc networks: Location discovery and sensor exposure. *Proc. 2001 ACM Intern. Symp. Mobile Ad Hoc Netw. Comp. MobiHoc*, pages 106–116, October 2001.

[21] J. A. T. T. M. Cover. *Elements of Information Theory*. John Wiley & Sons Inc., 1991.

[22] S. Čapkun, J.-P. Hubaux, and L. Buttyàn. Mobility helps security in ad hoc networks. In *Proc. ACM International Symposium on Mobile Ad Hoc Network and Computing*, pages 46–56, Annapolis, Maryland, June 2003.

[23] D. Wagner. Resilient aggregation in sensor networks. In *SASN'04*, pages 78–87, October 2004.

[24] A. D. Wood and J. A. Stankovic. Denial of service in sensor networks. *IEEE Computer*, 35(10):54–62, October 2002.

[25] O. F. X. Li, Peng-JunWan. Coverage in wireless ad hoc sensor networks. *IEEE Transactions on Computers*, 52(6):753–763, June 2003.

[26] K. C. Y. Zou. Uncertainty-aware and coverage-oriented deployment for sensor networks. *Journal of Parallel and Distributed Computing*, 64(7):788–798, July 2004.