Improving Intrusion Detection Systems for Wireless Sensor Networks

Andriy Stetsko, Tobiáš Smolka, Vashek Matyáš, and Martin Stehlík

Masaryk University, Brno, Czech Republic
{stetsko,xsmolka,matyas,xstehl2}@fi.muni.cz

Abstract. A considerable amount of research has been undertaken in the field of intrusion detection in wireless sensor networks. Researchers proposed a number of relevant mechanisms, and it is not an easy task to select the right ones for a given application scenario. Even when a network operator knows what mechanism to use, it remains an open issue how to configure this particular mechanism in such a way that it is efficient for the particular needs. We propose a framework that optimizes the configuration of an intrusion detection system in terms of detection accuracy and memory usage. There is a variety of scenarios, and a single set of configuration values is not optimal for all of them. Therefore, we believe, such a framework is of a great value for a network operator who needs to optimize an intrusion detection system for his particular needs, e.g., attacker model, environment, node parameters.

Keywords: Intrusion detection, optimization, wireless sensor networks.

1 Introduction

A wireless sensor network (WSN) consists of sensor nodes – small devices equipped with sensors, microcontroller, wireless transceiver and battery. Each node monitors a physical phenomenon and sends the measurements to a base station. Since a node communication range is limited to tens of meters and it is not always feasible for the node to directly communicate with the base station, data are usually sent hop-by-hop from one node to another until they reach the base station. WSNs can support various applications for ecology and wildlife monitoring, military, building and industrial automation, energy management, agriculture, etc.

Sensor nodes are constrained in *processing power*, *memory* and mainly in *energy*. A MICAz sensor node is a typical sensor node. It is equipped with the 8 MHz Atmel Atmega128L microcontroller, 512 KB flash memory, 802.15.4 compliant Texas Instruments CC2420 transceiver and two AA batteries. The transceiver consumes 18.8 mA (with 3.3 V power supply) in the receiving mode [1], which is the most energy consuming mode. If we use two NiZn AA batteries with a nominal voltage of 1.65 V and a capacity of 1800 mAh, the estimated lifetime of a constantly receiving node is approximately 96 hours, i.e., 4 days. However, in general for WSNs one expects a functional network for the duration of time that ranges from several days to several years.

I. Boureanu, P. Owesarski, and S. Vaudenay (Eds.): ACNS 2014, LNCS 8479, pp. 343–360, 2014.

In this paper, we propose a *framework that semi-automatically optimizes the configuration of an intrusion detection system* (IDS) in terms of detection accuracy and memory usage for any given scenario, e.g., a network topology, the network stack of benign sensor nodes, and anticipated attacks. We do not aim to propose particular novel techniques for intrusion detection (as such) in sensor networks and our ultimate long-term aim is to provide a framework that does not depend on a particular attacker model (or a group of models). We focus on intrusion detection since it is an essential mechanism to protect a network against internal attacks that are relatively easy and not expensive to mount in WSNs. In comparison to conventional wired and wireless networks, an attacker can often easily access the deployment area of a WSN, capture some nodes, and launch a wide range of attacks (for the list of possible attacks, see [18]).

The paper roadmap is as follows. The conceptual architecture of our framework is described in Section 2. Section 3 contains high-level technical details of our proof-of-concept implementation. In Section 4, we describe our test case. We tested the framework using a static topology in three different scenarios – these scenarios were selected to illustrate the framework merits, Section 5 describes these scenarios and test results. We compare our approach to related work in Section 6. Finally, Section 7 concludes the paper and presents plans for our future work. Particular details of evolutionary algorithms that we used for optimization are then provided in our technical report [7].

2 Conceptual Architecture of the Framework

In this section, we present the conceptual architecture of our framework that semi-automatically optimizes the configuration of an IDS for a given application scenario. The framework includes an *optimization engine* and a general-purpose *network simulator* (see Figure 1). The whole process consists of five main steps. In the first step, a network operator defines a fitness function for the evaluation of an IDS configuration. We define a reasonable fitness function in Subsection 3.3. It integrates evaluation metrics from [2] such as true positives, true negatives and memory usage. In the second step, the network operator configures the network simulator in such way that it simulates a scenario in which an IDS should be deployed. This step is described in Subsection 2.1 in more detail. The remaining three steps are completely automatic. The third and fourth steps take place in an iterative manner. The optimization engine provides a candidate configuration of an IDS to the simulator. The simulator evaluates it according to predefined metrics, e.g., detection accuracy, memory usage, and returns information required to compute the fitness function back to the optimization engine. Based on the evaluation, the optimization engine changes the values of parameters and repeats the procedure until a predefined condition holds, e.g., parameters become optimal for a given scenario, or the maximum number of iterations is exceeded. Finally, in the fifth step, the optimization engine outputs the best found (hopefully, the optimal) IDS configuration.

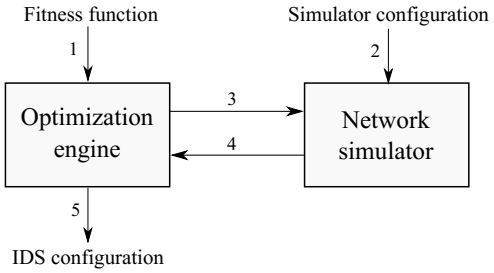

Fig. 1. Conceptual architecture. The arrows depict input (output) to (from) the components of our framework.

2.1 Simulator Configuration

The network operator provides a complete network model, which among others, includes a *network topology*, *models of benign and malicious nodes*, *wireless channel and energy consumption models* (see Figure 1).

Network Topology: The simulator provides a possibility to set a topology manually or to generate it automatically. In case the network operator knows the precise topology of the network, he can use the first option, and optimize the IDS for this particular topology. In case the topology is not known in advance, the network operator can use the second option, generate several random topologies, and optimize the IDS for all of them simultaneously. For more details, see the discussion on robustness of a found solution in Subsection 2.2.

Benign Sensor Node: Models of the node hardware and software should be provided, i.e., radio, IDS, medium access control (MAC) layer, network layer, and application layer models. There could be several types of a benign node in the network (e.g., a cluster head, a base station, a general-purpose sensor node), and they might have a different network stack. The network operator composes a benign node model from the available protocols (distributed within the simulator, or implemented by a third party). If the required protocol at a certain network layer is not available, the network operator should implement it. Further, the network operator configures the parameters of the models.

Malicious Sensor Node: Similarly, models of the node hardware and software should be provided. There could be several types of a malicious node in the network (e.g., internal/external [18], passive/active [21]). Usually, it is known to a network operator where the network will be deployed, and what the purpose of the network is. Based on this information, the network operator can estimate the risks of different attacks (e.g., selective forwarding, jamming, hello flood),

and include into a simulation only those that pose a serious threat. The network operator composes a malicious node model from the modified models available within the distribution of the simulator, or implements them by himself. For example, to implement a selective forwarder, it might be enough to modify the network layer of a benign node to drop a certain percentage of incoming packets.

Wireless Channel: The simulator can provide more than one model for radio propagation, and a network operator can choose the one that is more suitable for the environment where a network will be deployed.

Energy Consumption: The simulator can provide more than one energy consumption model.

2.2 Discussion

In this subsection, we discuss several issues related to the framework design choices and framework usage.

Simulator Versus Testbed: We decided to use a simulator since a testbed is slow for comparison of considered alternative configurations, labour-intensive and it does not produce comparable results due to the uncontrollable factors (e.g., wireless channel effects). Candidate configurations should be tested under the same network conditions, otherwise they cannot be compared. In a testbed, we are not able to *reproduce* the same environment each time a candidate configuration is tested because of the wireless channel effects. The simulator provides us with such a possibility. The usage of a simulator, however, does not mean that different wireless channel effects (similar to those in a real network) cannot be modelled in the simulator (see wireless channel models in [9]).

Simulator Calibration: In order to get realistic results, the operator needs to calibrate an energy consumption model and a wireless channel model in accordance to the environment where a network will be deployed (we calibrated the wireless channel for two specific environments in [4]). Their calibration can be done manually or automatically by the integration of a simulator and a testbed [5]. The calibration should take place before the optimization. The carefully calibrated wireless channel model (energy consumption model) can statistically reflect the wireless channel behaviour (energy consumption) in a real network.

Solution Robustness: A wireless environment is dynamic. It can change in an unexpected way which may result into conditions that were not observed during the calibration. The framework, however, should cope with that, i.e., a found solution should keep working (preferably decreasing its effectiveness only gradually) even if some network characteristics change in the network. In order to

achieve this, a candidate solution should be evaluated on networks with different topologies and a wireless channel model should be calibrated for different environments (e.g., network congestion, network maximum throughput). Moreover, the network operator can calibrate a wireless channel model for some pessimistic scenario (even though it has not been observed yet in a given environment). For example, he can deliberately increase the noise level in the noise model, increase the path loss or variation in the log-normal shadowing model. For more information on the log-normal shadowing model, see [9]. The evaluations obtained from different networks can be combined into a single evaluation score.

Framework Generality: The framework is generic, and it could be used to optimize different types of an IDS (misuse-based or anomaly-based, centralized or distributed). Further, we list several examples.

In [15], the authors proposed a scheme that activates IDS agents preloaded in sensor nodes with a certain probability. Our framework can be used to find an optimal value of the probability for a given application scenario, i.e., to solve the trade-off between the number of packets (links) left unmonitored (influences a detection accuracy) and the number of IDS agents being activated (influences energy consumption).

In [17], the authors proposed a distributed IDS that involves a set of rules to detect different types of an attack. Our framework can be used to automatically select the appropriate rules and optimize (among others) their detection thresholds.

In [16], the authors proposed an IDS to detect packet reception rate and receive power anomalies. The authors demonstrated that there is a trade-off between detection probability, detection delays and false positives for their technique. Our framework can help a network operator to find the optimal values of the IDS parameters for his/her application scenario.

3 Implementation of the Framework

Relevant high-level implementation details are provided in this section.

3.1 Optimization Engine

There are two classes of optimization algorithms – exact and approximate (heuristic) algorithms. Since the evaluation of candidate solutions cannot be done analytically in our case, the exact algorithms can hardly be applied. The heuristics are divided into population-based and single-solution based algorithms. We use a population-based algorithm because in comparison to a single-solution based algorithm it provides us with the ability to evaluate multiple candidate solutions in parallel and hence to speed up the convergence of an optimization.

There is a variety of population-based algorithms, e.g., evolutionary algorithms (EAs), particle swarm optimization, immune networks. We use EAs as we already have successfully applied these algorithms for the automatic generation of secrecy amplification protocols in WSNs [3].

EAs work with a population of candidate solutions (*individuals* in terms of EAs) and *evaluate* them using a *fitness function*. EAs generate new candidate solutions applying genetic operators of crossover and mutation to the solutions in the population. In each *generation*, EAs update the population with new candidate solutions. The process repeats until an optimal solution is found or the maximum number of iterations is exceeded. For more information on EAs, see [8].

The optimization engine is based on Evolving Objects [10], an advanced component-based framework with a high number of already implemented optimization algorithms. For the purpose of this work, we used a basic evolutionary algorithm *eoEasyEA* that is highly configurable and suits our needs. We reused existing operators for selection, replacement, termination and statistics collection, and implemented only problem specific parts of initialization, mutation, crossover, and evaluation. More details on the settings of EAs can be found in our technical report [7].

3.2 Network Simulator

We use the MiXiM network simulator [11], which is based on the OMNeT++ simulation framework [12]. MiXiM has a modular architecture with a high number of already implemented models for a WSN simulation. It inherits many advanced features from OMNeT++ and thus is very adaptable and configurable. The whole simulation is configured via a dedicated OMNeT++ configuration file. A candidate solution (an individual) is represented as a list of configuration values stored in a separate configuration file. Before the evaluation (simulation) starts, the file is included in the main configuration file.

The choice of a general-purpose WSN simulator allowed us to move one step forward towards more realistic simulations, since MiXiM provides more accurate simulation models, e.g., for wireless channel, radio, and MAC layers, in comparison to a very fast purpose-built simulator we used in our previous work [3]. However, the accuracy comes at the price of speed. We simulated a network operating for one hour, and it took about 5 minutes to simulate such network on a single CPU core. In order to get a solution in acceptable time, we decided to utilize distributed computing.

We chose the BOINC distributed computing platform [13] for our experiments. In cooperation with the Institute of Computer Science at Masaryk University, we attached about 200 CPU cores from the campus to our BOINC infrastructure and used them when they were idle. Other 700 cores were available from the National Grid Infrastructure project MetaCentrum.

3.3 Configuration Evaluation

A candidate configuration of an IDS is evaluated based on its *accuracy* and *memory usage*. In this paper, two terms IDS and monitoring node are used interchangeably.

Notation 1. *The set* $A = \{a_1, ..., a_{n_m}\}$ *is a set of malicious nodes in a network.*

Notation 2. *The set* $C = \{c_1, ..., c_{n_b}\}$ *is a set of all benign nodes in a network.*

Notation 3. *The function* $x : \mathbb{N} \to \mathbb{N}$ *takes a sensor node index as an argument, and returns a number of the neighbours that consider this node benign.*

Notation 4. *The function* $y : \mathbb{N} \to \mathbb{N}$ *takes a sensor node index as an argument, and returns a number of the neighbours that consider this node malicious.*

Notation 5. *The function* $n : \mathbb{N} \to \mathbb{N}$ *takes a sensor node index as an argument, and returns a number of the neighbours of this node.*

Notation 6. *The function* $m : \mathbb{N} \to \mathbb{N}$ *takes a sensor node index as an argument, and returns the amount of memory (in bytes) used by an IDS on this node.*

Accuracy: We measured accuracy based on the number of true positives and true negatives:

- A true positive occurs when a monitoring node $c \in C$ correctly considers its neighbour $a \in A$ malicious.
- A true negative occurs when a monitoring node $c_i \in C$ correctly considers its neighbour $c_j \in C$ $(i \neq j)$ benign.

A node $k \in C \cup A$ considers the node $l \in C \cup A$ as a neighbour if it received at least one packet from l during the simulation. We assume that sensor nodes are distributed in such a way, that every node in the network has at least one neighbour.

For a benign node c_i, we calculated the percentage of the neighbours that considered the node benign. Further, we found the average of such values over all benign nodes in the network and denoted the result as tn. Similarly, for a malicious node a_i, we calculated the percentage of the neighbours that considered the node malicious. Further, we found the average of such values over all malicious nodes in the network and denoted the result as tp.

The accuracy function is the weighted mean of tn and tp:

$\frac{w_1 * tn + w_2 * tp}{(w_1 + w_2)}$, where $tn = \frac{1}{|C|} * \sum_{c_i \in C} \frac{x(c_i)}{n(c_i)}$, $tp = \frac{1}{|A|} * \sum_{a_i \in A} \frac{y(a_i)}{n(a_i)}$.

We assume that $|C| > 0$ and $|A| > 0$.

The function values range from 0 to 1. If every malicious node in the network is detected by all of its neighbours, and every benign node in the network is not

considered malicious by any of its neighbours, the accuracy function is equal to 1, i.e., the maximum possible value. On the other hand, if none of malicious nodes is detected by at least one of its neighbours, and every benign node is considered malicious by all of its neighbours, the accuracy function is equal to 0, i.e., the minimum possible value.

The proposed function does not take the distribution of $\frac{x(c_i)}{n(c_i)}$ and $\frac{y(b_i)}{n(b_i)}$ into account. In certain cases, e.g., when a base station uses a majority voting scheme to make a final decision whether a node is benign or not, it might be preferable to have more values of $\frac{x(c_i)}{n(c_i)}$ and $\frac{y(b_i)}{n(b_i)}$ that are slightly above 0.5 instead of a few values that are extremely high.

Memory Usage: The effectiveness of memory usage by the IDS on a node $c_i \in C$ was evaluated using the formula: $\frac{1}{1+m(c_i)}$.

If the IDS is switched off at the node c_i, then $m(c_i) = 0$ and $\frac{1}{1+m(c_i)} = 1$. Furthermore, if $m(c_i)$ increases, then the effectiveness of memory usage decreases towards zero.

Further, we calculated the average value of $\frac{1}{1+m(c_i)}$ over all benign nodes in the network. More formally, it can be written as: $\frac{1}{|C|} * \sum_{c_i \in C} \frac{1}{1+m(c_i)}$. We assume that $|C| > 0$.

The designed function provides values that are not correlated to accuracy values. In certain cases, however, it might be useful to take into account that even a small amount of memory is a waste if the accuracy of an IDS is low. Yet a higher amount of used memory can be justified if the IDS is highly accurate.

Fitness Function: For the purpose of this work, we added both accuracy and memory usage metrics together, making the accuracy metric to contribute more to the value of the sum than the memory metric by introducing weights. The weight was set to 1 for the accuracy metric, and it was set to 0.1 for the memory usage metric. The weight for the memory usage should be carefully selected – if it is too high (i.e., it is more important to save memory than to detect attacks), the optimal solution is to switch all IDSs off, or set a maximum number of monitored nodes and buffer size to zero. We set $w_1 = w_2 = 1$ in the accuracy metric. The resulting fitness function is:

$$\frac{1}{2|C|} * \sum_{c_i \in C} \frac{x(c_i)}{n(c_i)} + \frac{1}{2|A|} * \sum_{a_i \in A} \frac{y(a_i)}{n(a_i)} + 0.1 \frac{1}{|C|} * \sum_{c_i \in C} \frac{1}{1+m(c_i)}.$$

4 Our Test Case

In this section, we describe the scenario for which we would like to test the framework on. The framework is used to find optimal parameters of an IDS (we implemented it in the MiXiM simulator for purposes of our previous work [6]) for a given scenario. A network operator (a person who uses our framework) knows behavior of benign nodes and assumes behavior of malicious nodes.

If another than assumed type of an attack appears in the deployed network, the parameters found by the framework might not be optimal for such network.

4.1 Topology

We generated a topology of 250 static sensor nodes uniformly randomly distributed over an 200 m × 200 m area. A single base station is placed in the center of the area. The topology together with a routing tree is depicted in Figure 2. There are 246 benign sensor nodes (white) including the base station, and 5 malicious nodes (black filled). According to [14], the terrestrial WSNs typically consist of hundreds to thousands of inexpensive sensor nodes. However, the purchase of a large network is not always feasible due to the current price of sensor nodes. A MICAz sensor node costs about €80. Therefore, we believe that medium-sized networks that consist of hundreds of nodes are more reasonable to consider.

In order to make the analysis of results (w.r.t. their optimality) from the framework simpler and more intuitive, we focus on static sensor nodes. However, more dynamic network scenarios can be modeled as well. MiXiM provides a variety of node mobility models [11].

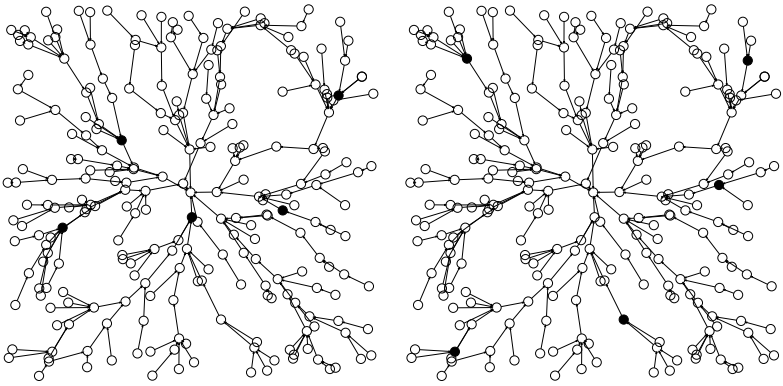

(a) 1st placement of mal. nodes. (b) 2nd placement of mal. nodes.

Fig. 2. Topology, routing tree, and placement of malicious nodes

4.2 Benign Node

We assume a benign node uses a network stack that consists of application, routing protocol, MAC protocol, and intrusion detection system.

Application Layer: We consider a standard application for a WSN, where every node sends one packet to a base station every 30 seconds. The application runs for one hour. In order to avoid collisions (due to node synchronization), the whole time-frame is divided into intervals of length 30 seconds. For every interval i, a node generates a random number r ($0 \leq r \leq 30$) and starts transmitting at $i + r$. The size of a packet is 152 B.

Network Layer: We assume that the network layer uses static routing. The routing tree was generated as follows. A base station broadcasts a packet containing its identification together with the value h set to 0. A node waits until it receives a packet from a neighbour that is the closest one (has the highest signal strength). Then the node sets the neighbour as its parent, increases value h by 1 and broadcasts the value together with its identification. Value h represents number of hops to the base station.

Medium Access Layer: We use CSMA-CA at this layer.

Physical Layer: We use a model for CC2420 radio that is commonly used in sensor nodes, e.g., TELOSB, MICAz sensor nodes.

Intrusion detection system: We use a simple IDS that we implemented in the MiXiM simulator for the purpose of our previous work [6]. The IDS uses a detection rule (more specifically, a retransmission rule) from [17]. It is not the goal of this paper to propose a complex IDS, but rather to test the framework as such. In the conventional networks, one can use a commercial IDS, to run the framework on, and compare the IDS before and after the optimization to see the improvement. For WSNs, however, to our best knowledge, there are no commercial IDSs available.

An IDS is running on a sensor node and it continuously analyzes sent and overheard packets. The IDS does not include responsive and collaborative components (see [20] for the conceptual architecture of an IDS in WSNs). Therefore, the IDS does not generate any additional traffic.

A monitoring node overhears to some extent both incoming and outgoing packets of a close enough monitored neighbour. An IDS stores a table, where each row corresponds to a certain monitored node. The table contains the number of packets received (PR) and forwarded (PF) by a monitored node. The number of rows is limited to a *number of monitored nodes*.

The detection exploits the fact that a monitoring node overhears (to some extent) both incoming and outgoing packets of a close enough monitored neighbour. If the IDS on a node $c_i \in C$ overhears a packet P sent to a node $b_j \in C \cup A$, b_j is close enough and b_j should forward the packet (e.g., b_j is not a base station), then the IDS stores P in the buffer and increments the PR counter of the monitored node b_j. The number of packets is limited by a *buffer size*. If a new packet arrives but the buffer is full, the oldest packet is removed from the buffer. When the IDS overhears the packet P being forwarded by the node b_j, it removes P from the buffer (if it is still there) and increments the PF counter of the node b_j. Since both the table and the buffer are limited, the IDS monitors only the closest nodes and the newest packets.

The detection is done at the end of the simulation based on the collected statistics. The node c_i considers the node b_j as a selective forwarder if the dropping ratio of b_j, i.e., ratio of a number of packets dropped to a number of packets

received, is higher than a predefined *detection threshold*. If the node c_i overheard less than the predefined *number of packets received* by b_j as overheard by c_i, c_i does not consider b_j malicious as the number of overheard packets is small and there is a high level of uncertainty. We cannot influence the number of overheard packet, but we can change our decision making process based on this value and potentially decrease a number of false positives.

The IDS running on a node c_i consumes $m(c_i) = p_1 * 8 + p_2 * 16$ B of memory. Each record in the table occupies 8 B (4 B for node ID, 2 B for PR, 2 B for PF), and there are p_1 such records. Each record in the buffer occupies 16 B (4 B for MAC source ID, 4 B for MAC destination ID, 4 B for MAC intermediate node ID, and 4 B for packet counter), and there are p_2 such records.

We would like to optimize the following four parameters: p_1 (number of nodes to be monitored), p_2 (a number of packets stored in a buffer), p_3 (number of packets received) and p_4 (detection threshold). The value of p_1 ranges from 0 to 54 (the maximum number of neighbours in our simulation scenario), p_2 – from 0 to 100, p_3 – from 0 to 2000, and p_4 from 0 to 100.

4.3 Malicious Node

Currently, for our proof-of-concept implementation of the framework, we assume a single type of malicious node – a selective forwarder, i.e., a node that drops a certain percentage of received packets. We assume the model is the same as the model of a benign node, except for the network layer that is modified to drop a certain percentage of received packets (in our case 50%), and an IDS that is omitted.

5 Testing Results

We tested our prototype on three optimization scenarios, each with the different size of the search space. Their description together with the obtained results are presented in the following subsections. The settings of EAs for each optimization scenario are described in the corresponding subsections of our technical report [7]. For the evaluation of a candidate configuration, we used the fitness function defined in Subsection 3.3. Time needed to complete the optimization is indicated in terms of a number of EA generations and evaluations.

5.1 Optimization Scenario No. 1

In this scenario, we assume that every benign node in the network runs an IDS, and the IDS is configured in the same way for these nodes. The goal is to optimize the configuration (p_1, p_2, p_3 and p_4), common for all sensor nodes, using our framework.

We performed both an exhaustive search and an EA-based search, and compared the results. In order to make the exhaustive search timely acceptable, we fixed $p_2 = 100$ and $p_3 = 0$. The reduced search space contained 5555 possible configurations (see the description of an IDS in Subsection 4.2).

Exhaustive Search: Fitness values for each possible combination of p_1 and p_4 are depicted in Figure 3. The maximum fitness value (0.8249276442) was achieved for the configuration with $p_1 = 27$ and $p_4 = 0.45$. The threshold is below 0.5 (the dropping rate of a malicious node, see Subsection 4.3), because a monitoring node cannot reliably overhear all packets sent to a malicious node, and hence the dropping rate of the malicious node as observed by the monitoring node may be lower than 0.5.

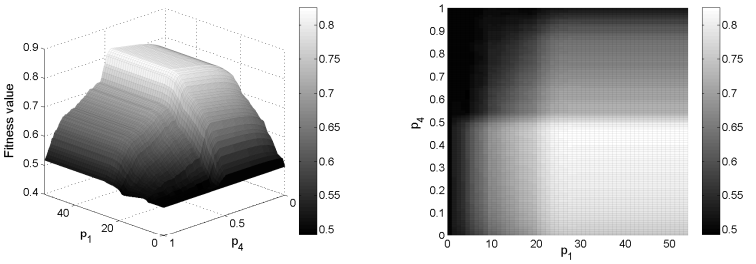

(a) View from the side. (b) View from the top.

Fig. 3. Fitness values for all possible combinations of p_1 and p_4

Evolutionary Algorithm: We ran the optimization process 30 times. The EA was able to find the best configuration (found by the exhaustive search) using 19 generations on average. The standard deviation was 9.64. On average, the EA required 144.87 evaluations. The standard deviation was 67.37. In the worst case, the EA required 271 evaluations, while the exhaustive search required 5400 evaluations.

5.2 Optimization Scenario No. 2

In this scenario, we assume that only a subset of benign nodes runs IDSs, which are configured in the same way on all selected nodes. As opposed to the node (IDS) placement problem, the framework should find the optimal placement as well as the optimal configuration of the IDSs.

There are 250 parameters to optimize: p_1, p_2, p_3, p_4, and 246 Boolean parameters, each indicating whether the IDS should be enabled or disabled on a given node. We fixed two parameters ($p_2 = 100$ and $p_3 = 0$). The search space contained $55 * 101 * 2^{246}$ possible configurations.

We ran the optimization process on two networks with the same topology but with the different placement of malicious nodes (see Figure 4(a) and Figure 4(b)). The malicious nodes are depicted with black filling. We repeated the optimization process 30 times for both placements of malicious nodes.

First Placement of Malicious Nodes: The best configuration found by the EA had the fitness value equal to 0.8940953359, $p_1 = 27$, and $p_4 = 0.45$ (the

same values were found for the first optimization scenario). The switched on/off IDSs are depicted in Figure 4(a). Nodes that ran an IDS are depicted with grey filling. The configuration generated 1723 false positives and 33 false negatives, i.e., 7.03 false positives per benign node, and 6.6 false negatives per malicious node. The configuration was found using 581.77 generations on average. The standard deviation was 88.59. On average, the EA required 10600.17 evaluations. The standard deviation was 1603.29. In the worst case, the EA required 13349 evaluations.

Although we did not perform an exhaustive search, we believe that the found configuration was optimal. The intuition behind this is as follows. If an IDS running on a node $c_i \in C$ detects a malicious neighbour $a_j \in A$, then switching it off reduces a number of false positives, increases memory usage effectiveness, but causes a false negative. As we discovered, it is natural for the EA to switch such an IDS on since the benefit (according to the designed evaluation function, see Subsection 3.3) is higher than from switching the IDS off. We verified (by setting $p_4 = 0$ and $p_1 = 54$) that the EA achieved the minimum possible number of false negatives, and the minimum number of false positives for the given number of false negatives.

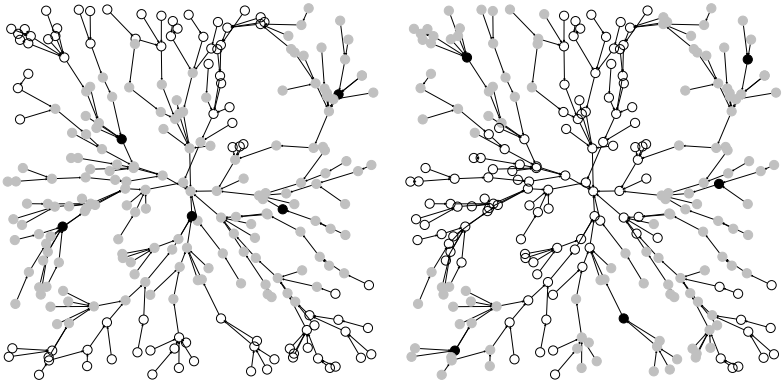

(a) 1st placement of mal. nodes. (b) 2nd placement of mal. nodes.

Fig. 4. The best found placements of IDSs in the network

Second Placement of Malicious Nodes: The best configuration found by the EA had the fitness value equal to 0.8780288967, $p_1 = 24$, and $p_4 = 0.5$. The switched on/off IDSs are depicted in Figure 4(b). The configuration generated 1173 false positives and 42 false negatives. We believe that the found configuration was optimal. The intuition behind this is the same as for the first placement of malicious nodes. We verified (by setting $p_4 = 0$ and $p_1 = 54$) that the achieved number of false negatives was the lowest possible, and the achieved number of false positives was the lowest possible for the given number of false negatives.

The best configuration was found by 14 optimization runs. The average number of generations was 499.57. The standard deviation was 82.185. On average,

the EA required 9096.1 evaluations. The standard deviation was 1501.5. Other 8 runs ended up with the configuration where the fitness value was equal to 0.8764108606. The rest of the runs found configurations with fitness values equal to or higher than 0.8656972029.

When comparing the best configurations found for both placements, an average number of false events per node (an average number of false positives per node added together with an average number of false negatives per node) was higher for the first placement (13.6 versus 13.17), but the fitness value was higher for the second one (0.8940953359 versus 0.8780288967). That might be caused by the fact that nodes falsely accused or falsely not detected in the first placement had a higher number of neighbours, which is also reflected in the selected value of p_1 (27 versus 24).

5.3 Optimization Scenario No. 3

In this scenario, we again assume that only a subset of benign nodes runs IDSs. In comparison to the previous optimization scenario, here each IDS may be configured in a different way. As opposed to the node (IDS) placement problem, the framework should find the optimal placement as well as the optimal configuration of each IDS.

The search space contained $(55 * 101 * 2001 * 101 * 2)^{246}$ configurations. The IDS implemented in our test case did not influence the network traffic, and hence did not influence other IDSs. Therefore, configurations of any two IDSs were independent of each other. We launched 246 independent optimizations and significantly reduced the search space, i.e., to $246 * (55 * 101 * 2001 * 101 * 2)$ possible configurations.

We did two experiments. In the first experiment, we launched a single optimization that searched for the best configuration for all sensor nodes together. In the second experiment, we launched 246 independent optimizations, each searching for the best configuration for a particular sensor node only.

First Experiment: We started with a randomly generated initial population. The evolution began to improve the configuration, but the speed of convergence was too slow. Therefore, we decided to start the optimization process again, using the population of the best configurations from the second optimization scenario (see Subsection 5.2). This optimization was gradually improving the configuration and reached the fitness value of 0.955133 after $136'765$ evaluations in the $21'439^{th}$ generation. The evolution was stopped when no improvement was found during next 500 generations.

Second Experiment: All 246 optimizations started with the initial population that contained the best configurations for the second optimization scenario. By combining optimized parameters from these optimizations, a configuration that reached the fitness value of 0.9604364302 was found. When compared to the second optimization scenario, this approach was able to significantly improve

the resulting IDS configuration by tweaking the parameters independently for each node. A configuration found in each independent optimization required 422.63 generations on average. The standard deviation was 63.84. On average, the EA required 2811.68 evaluations to find a configuration in each independent optimization. The standard deviation was 427.05. In the worst case, the EA required 5365 evaluations.

We ran another experiment to verify whether the configuration found by the EA was optimal. The experiment was based on the intuition we mentioned in Subsection 5.2. For an IDS, we fixed the parameters in such a way that it could detect as many malicious nodes as possible ($p_1 = 54$, $p_2 = 100$, $p_3 = 0$, $p_4 = 0$). Further, we found the minimum value of p_1 such that all malicious nodes can still be detected by the IDS. The same procedure was repeated for other three parameters. We confirmed that p_1, p_2 and the placement of IDSs were optimal. However, p_3 and p_4, as we discovered, could be further improved.

6 Related Work

The placement of nodes (IDSs) can be considered as a subproblem of a more general node (IDS) configuration problem where one of the node (IDS) configuration parameters may indicate where the node is placed (whether an IDS is enabled/disabled at this node).

The placement problem received high attention from research community. For example, [22,23,24] proposed techniques for IDS placement problem in WSNs. [19] surveyed 46 techniques that help to find optimal node placement in WSNs. In contrast, our framework is more generic and additionally provides a possibility to find optimal configuration of IDS parameters (e.g., detection threshold, buffer size). For more information, see Sections 4 and 5. To the best of our knowledge, we are the first optimizing IDS parameters in WSNs.

[22] proposed an algorithm finding a minimum number of activated IDSs such that every packet forwarded from a source towards the base station was analyzed at least once on its path. In comparison to our work, the authors did not consider IDS configuration. Furthermore, the algorithm considers only packets forwarded by the monitoring nodes and not packets overheard in a promiscuous mode.

[26] aimed at multi-objective optimization using an EA for deployment of a homogeneous WSN with the coverage and lifetime of the network as objectives. [25] presented a methodology of multi-objective optimization for self-organizing WSNs using an EA. Their fitness function took application-specific, connectivity and energy-related metrics into account. The goal was to find out optimized operation mode for each node in the network. The authors did not consider any IDS. Furthermore, they did not consider node parametrization.

Methodology how a multi-objective evolutionary algorithm for design-space exploration could be configured was presented in [27]. Lifetime, latency and reliability are used as three QoS (Quality of Service) metrics with several trade-offs. In our work, we used a different set of metrics. The authors used multi-objective optimization.

[28,29,30] used EAs for several optimization issues in WSNs.

[30] incorporated *local monitoring nodes* (LMNs) into the WSN. These nodes observe suspicious behavior and monitor data message patterns, message collisions, route traffic activity trends and sensor positioning in their neighbourhood. The fitness function measures optimality of the LMN positioning and accurate identification of malicious nodes. The authors do not consider optimization of IDSs parameters in this work.

More detailed treatment of related work is provided in our tech. report [7].

7 Conclusion and Future Work

To our best knowledge, there is no work that focuses on (semi) automatic and systematic configuration of intrusion detection systems for wireless sensor networks, which we believe is an important area to explore.

In this work, we describe procedures to optimize the configuration of an intrusion detection system for a given application of a wireless sensor network. Also, we discusses how solution robustness and solution realism can be achieved.

We presented a prototype of our framework that optimizes the configuration of an intrusion detection system in wireless sensor networks. The design and implementation of the framework leveraged our previous results in the field of simulators (particularly realism of simulators for wireless sensor networks), optimization and evaluation metrics.

We tested our framework on three carefully selected scenarios with a different size of their search space. Our results demonstrated that evolutionary algorithms can be potentially used to search for a solution of the given problem more effectively. However, this conclusion is not valid in general (since the hypothesis was tested only on three selected scenarios). More general conclusion can be made if the community use evolutionary algorithms on a bigger set of different scenarios (different application, routing, medium access control and physical layer, different topologies, different attacker models).

Our framework can find reasonable (if not optimal) configuration of an intrusion detection system for a given (arbitrary but specified in advance) scenario. Values found by our framework may not be reasonable if these values are used in a different scenario, i.e., intrusion detection system, topology, attacker behavior.

The obtained results, we believe, are more than promising. Hence, we plan to use our framework to optimize different techniques for detection of different attackers. Also, we plan to use our framework to optimize an intrusion detection system for our laboratory testbed, configure the intrusion detection system according to the output provided by the framework, deploy it, and analyze the results collected from the testbed.

The future version of our framework will also use multi-objective evolutionary algorithms.

While we focused on the optimization of an intrusion detection system (other layers were fixed), we believe that such a framework can be easily extended to optimize the whole network stack. This can be explored in the future.

Acknowledgment. This work was supported by the project GAP202/11/0422 of the Czech Science Foundation. We thank the National Grid Infrastructure project MetaCentrum and Institute of Computer Science at Masaryk University for the provided computational resources. Our thanks also belong to anonymous reviewers.

References

1. Texas Instruments. CC2420 – 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF transceiver, `http://focus.ti.com/`
2. Stetsko, A., Matyas, V.: Effectiveness metrics for intrusion detection in wireless sensor networks. In: European Conference on Computer Network Defense, pp. 21–28 (2009)
3. Svenda, P., Sekanina, L., Matyas, V.: Evolutionary design of secrecy amplification protocols for wireless sensor networks. In: ACM Conference on Wireless Network Security, pp. 225–236 (2009)
4. Stetsko, A., Stehlik, M., Matyas, V.: Calibrating and comparing simulators for wireless sensor networks. In: IEEE Conference on Mobile Adhoc and Sensor Systems, pp. 733–738 (2011)
5. Wen, Y., Zhang, W., Wolski, R., Chohan, N.: Simulation-based augmented reality for sensor network development. In: ACM Conference on Embedded Networked Sensor Systems, pp. 275–288 (2007)
6. Stetsko, A., Smolka, T., Jurnecka, F., Matyas, V.: On the credibility of wireless sensor network simulations: evaluation of intrusion detection system. In: Conference on Simulation Tools and Techniques, pp. 75–84 (2012)
7. Stetsko, A., Smolka, T., Matyas, V., Stehlik, M.: Improving intrusion detection systems for wireless sensor networks. Technical report FIMU-RS-2014-01: Masaryk University, Faculty of Informatics, Brno, Czech Republic (March 2014)
8. Talbi, E.-G.: Metaheuristics – From Design to Implementation. John Wiley & Sons, Inc. (2009)
9. Rappaport, T.: Wireless communications: Principles and practice, 2nd edn. Prentice Hall PTR (2001)
10. Keijzer, M., Merelo, J.J., Romero, G., Schoenauer, M.: Evolving objects: A general purpose evolutionary computation library. In: Conference on Evolution Artificielle, pp. 231–242 (2002)
11. Kopke, A., Swigulski, M., Wessel, K., Willkomm, D., Haneveld, P.T.K., Parker, T.E.V., Visser, O.W., Lichte, H.S., Valentin, S.: Simulating wireless and mobile networks in OMNeT++ the MiXiM vision. In: Conference on Simulation Tools and Techniques for Communications, Networks and Systems & Workshops (2008)
12. OMNeT++ Community, `http://www.omnetpp.org/`
13. Anderson, D.P.: BOINC: a system for public-resource computing and storage. In: IEEE/ACM Workshop on Grid computing, pp. 4–10 (2004)
14. Yick, J., Mukherjee, B., Ghosal, D.: Wireless sensor network survey. Computer Networks 52(12), 2292–2330 (2008)
15. Roman, R., Zhou, J., Lopez, J.: Applying intrusion detection systems to wireless sensor networks. In: IEEE Consumer Communications and Networking Conference, pp. 640–644 (2006)
16. Onat, I., Miri, A.: An intrusion detection system for wireless sensor networks. In: IEEE Conference on Wireless and Mobile Computing, Networking and Communications, pp. 253–259 (2005)

17. da Silva, A.P.R., Martins, M.H.T., Rocha, B.P.S., Loureiro, A.A.F., Ruiz, L.B., Wong, H.C.: Decentralized intrusion detection in wireless sensor networks. In: ACM International Workshop on Quality of Service & Security in Wireless and Mobile Networks, pp. 16–23 (2005)
18. Karlof, C., Wagner, D.: Secure routing in wireless sensor networks: attacks and countermeasures. Ad Hoc Networks 1(23), 293–315 (2003)
19. Younis, M., Akkaya, K.: Strategies and techniques for node placement in wireless sensor networks: A survey. Ad Hoc Networks 6(4), 621–655 (2008)
20. Zhang, Y., Lee, W.: Intrusion detection in wireless ad-hoc networks. In: Conference on Mobile Computing and Networking, pp. 275–283 (2000)
21. Roosta, T., Pai, S., Chen, P., Sastry, S., Wicker, S.: Inherent security of routing protocols in ad-hoc and sensor networks. In: Global Telecommunications Conference, pp. 1273–1278 (2007)
22. Anjum, F., Subhadrabandhu, D., Sarkar, S., Shetty, R.: On optimal placement of intrusion detection modules in sensor networks. In: Conference on Broadband Networks, pp. 690–699 (2004)
23. Liu, C., Cao, G.: Distributed monitoring and aggregation in wireless sensor networks. In: Conference on Computer Communications, pp. 1–9 (2010)
24. Hassanzadeh, A., Stoleru, R.: Towards optimal monitoring in cooperative IDS for resource constrained wireless networks. In: Conference on Computer Communications and Networks, pp. 1–8 (2011)
25. Ferentinos, K.P., Tsiligiridis, T.A.: Adaptive design optimization of wireless sensor networks using genetic algorithms. Computer Networks 51(4), 1031–1051 (2007)
26. Jourdan, D.B., de Weck, O.L.: Layout optimization for a wireless sensor network using a multi-objective genetic algorithm. In: IEEE Vehicular Technology Conference, pp. 2466–2470 (2004)
27. Nabi, M., Blagojevic, M., Basten, T., Geilen, M., Hendriks, T.: Configuring multi-objective evolutionary algorithms for design-space exploration of wireless sensor networks. In: ACM Workshop on Performance Monitoring and Measurement of Heterogeneous Wireless and Wired Networks, pp. 111–119 (2009)
28. Khanna, R., Liu, H., Chen, H.H.: Self-organization of sensor networks using genetic algorithms. In: IEEE Conference on Communications, pp. 3377–3382 (2006)
29. Khanna, R., Liu, H., Chen, H.H.: Dynamic optimization of secure mobile sensor networks: A genetic algorithm. In: IEEE Conference on Communications, pp. 3413–3418 (2007)
30. Khanna, R., Liu, H., Chen, H.H.: Reduced complexity intrusion detection in sensor networks using genetic algorithm. In: IEEE Conference on Communications, pp. 1–5 (2009)