Shortfall-based Optimal Placement of Security Resources for Mobile IoT Scenarios

Antonino Rullo¹, Edoardo Serra², Elisa Bertino³ and Jorge Lobo⁴

¹ DIMES Department, Universita della Calabria, 87036 Rende, Italy,
n.rullo@dimes.unical.it

² Department of Computer Science, Boise State University, Boise, ID 83725, USA,
edoardoserra@boisestate.edu

³ Lawson Computer Science Department, Purdue University, West Lafayette, IN
47907, USA, bertino@purdue.edu

⁴ ICREA and Department of Information and Communication Technologies,
Universitat Pompeu Fabra, 08018 Barcelona, Spain, jorge.loba@upf.edu

Abstract. We present a method for computing the best provisioning of security resources for Internet of Things (IoT) scenarios characterized by a high degree of mobility. The security infrastructure is specified by a security resource allocation plan computed as the solution of an optimization problem that minimizes the risk of having IoT devices not monitored by any resource. Due the mobile nature of IoT devices, a probabilistic framework for modeling such scenarios is adopted. We adapt the concept of *shortfall* from economics as a risk measure and show how to compute and evaluate the quality of an allocation plan. The proposed approach fits well with applications such as vehicular networks, mobile ad-hoc networks, smart cities, or any IoT environment characterized by mobile devices that needs a monitoring infrastructure.

Keywords: Network security, Internet of Things, Stochastic allocation.

1 Introduction

The Internet of Things (IoT) will increase the ubiquity of the Internet by integrating every object for interaction via embedded systems leading to a highly distributed network of devices communicating with human beings as well as other devices [37]. The International Telecommunication Union defines the Internet of Things as “a global infrastructure for the Information Society, enabling advanced services by interconnecting (physical and virtual) things based on, existing and evolving, inter-operable information and communication technologies”.⁵ These definitions give us a hint to the role IoT will play in the everyday life, and the impact it will have in several areas: the inclusion of developing countries in global trade, the use of search engines to the benefit of civil society, combating

⁵ <http://www.itu.int/en/ITU-T/Pages/default.aspx>.

product counterfeiting, tackling environmental concerns, improving health conditions, securing food supply and monitoring compliance with labor standards [36]. From these observations, it also follows that security will play a very important role in IoT. A security infrastructure will be mandatory for ensuring the integrity of data and reliability of network participants. However, the mobility and heterogeneity of IoT devices make the process of security provisioning more complicated than for traditional computer networks, where both the network topology and the number of network devices are assumed to be static. IoT networks continuously change topology because of mobility and/or (dis)appearance of devices.

In such IoT environments, the geographical distribution and the number of connected devices are not stationary, and vary from spot to spot, according to the different activities taking place in the different areas. A security infrastructure must be able to address such a dynamic nature of IoT networks. To this end, a security resource allocation plan must take into account the numerous shapes the device mobility confers to the system of interest, and ensure a certain security level in the majority of cases.

In this paper we present a method for computing the best allocation plan of security resources for IoT scenarios characterized by a high degree of mobility. By security resources we mean passive monitors of wireless traffic, and by plan, the number and location of such resources. We formalize a model for such scenarios, and provide a heuristic for computing allocation plans that minimize the risk of having IoT devices not monitored by any resource. We employ the key-concept of *shortfall* [3] as a risk measure. Shortfall is mostly used in economics to model the risk associated with an investment by combining in a single (risk) value the return of the investment in the worst scenarios together with the expected return. In fact, an investment might provide a return much more scarce than the expected one due to the significant changes that can affect the market. We adapt the concept of shortfall to model the risk associated with a security solution for which its effectiveness depends on how well the solution is able to address the continuous topology changes that affect the system of interest.

Our contributions can be summarized as follows:

- the formalization of a model for describing a mobile IoT scenario;
- the formalization of the problem of finding an appropriate set of candidate allocation plans of security resources for an IoT environment as an optimization problem;
- the adaptation of the concept of shortfall to our security domain as a risk measure of a security solution.

2 The Framework

The main characteristics of IoT environments are the heterogeneity, due to the simultaneous presence of different types of signals and networks, and the continuous topology changes of the network composed by the devices moving in the area of interest. Simple devices such as RFID tags broadcast messages for

others to use, devices paired with other devices use bluetooth or WiFi as communication medium, people equipped with smartphones, laptops, VoIP devices, PDAs, smartwatches, well-being devices etc., move from place to place within a common macro-area. As an example of IoT environment we can think of a university campus where thousand of people with several mobile connected devices continuously change their position. Another well-fitting example is a vehicular network (VANET) that, although more homogeneous in the type of devices, assumes different shapes and different sizes according to the hour and day of the week. In such environments, the traditional security monitoring must be enhanced by security tools *(i)* capable of listening to different types of signals, *(ii)* capable of working in conjunction with other security resources to cover all possible signals, and *(iii)* smartly deployed in the network area in order to monitor as many devices as possible despite continuous changes of the network topology. In this paper we focus on the third aspect, and propose a method for intelligently placing security monitors, such that a certain level of security is ensured even when the system of interest assumes shapes and sizes that are uncommon.

2.1 Threat Model

In such complex IoT ecosystem, different entities have different capabilities and potential to cause security incidents. While our defensive approach is independent from the actual attacks, we assume that the attacker may perform attacks at different layers, and can take any of the steps commonly used to carry out attacks: capturing and reprogramming devices, adding malicious entities to the network to overhear data communications, inject false data and control traffic, intercept and drop data packets, introduce interference, claim multiple identities, exfiltrate data, security credentials or encryption keys, compromise other IoT devices in the network with the help of a compromised device. Also attacks aimed at Internet services are carried through IoT devices. Recent news have in fact reported the use of IoT devices for botnet attacks [10].

2.2 Defender Side

The goal of the defender is to passively monitor all wireless communications the devices moving in the area of interest do. To this end, (s)he needs to choose a set of security resources to monitor traffic and presence of devices, and decide where to deploy them in order to minimize the number of non-monitored devices. There are different kinds of security services the defender may have to use. Intrusion detection systems (IDS) [20, 27] as well as attack prevention systems (APS) [29] can be installed in the network area to, respectively, detect ongoing attacks, and prevent future ones; recovery systems [34] can be installed in order to address security problems reported by IDSs and APSs; physical tools, like directional antennas and highly sensitive transceivers can be adopted for enabling security resources collecting data for the services to widen their action range [24]. There are several characteristics of IoT that can be leveraged to design a security resource well fit for IoT. While heterogeneous, most IoT devices communicate and

operate on standard mediums and protocols, such as IEEE 802.15.4 [14], WiFi or Bluetooth for mediums, and ZigBee [39] or 6LoWPAN [13] for protocols. Therefore, as long as a device communicates by using any of these mediums and protocols, effective techniques such as promiscuous overhearing and watchdog-based mechanisms [12, 21] can be deployed (see e.g. [20]). Notice that a security resource may embed one or more of these services, or may just collect data to be sent to security services located in remote locations. We assume that different kinds of security resources act independently and hence, are independently deployed.

The geographical distribution and the number of connected devices in the area of interest are key parameters for the task of computing the best placement of security resources. Such distributions can be learned during a training period long enough to be a representative sample of the monitored eco-system. Following the two examples of IoT environments described above, in a campus area, most students follow patterns that repeat every day, and can be learned in a few months of an academic year: during the morning and the afternoon they are located at the classrooms; at lunch/dinner time they move to the dining areas; during the night they stay in the dorms or other in-campus accommodations. Such patterns can be learned by analyzing the network traces produced by the interaction between users' devices, or between access points and the devices. In a VANET, where cars move according to the hour and day of the week, such patterns can be learned from data collected by road side units (RSUs). Commuting patterns are also predictable by stochastic processes that capture local mobility decisions. Such processes help analytically derive commuting and mobility fluxes that require as input only information on the population distribution. The resulting model predicts mobility patterns in good agreement with mobility and transport patterns observed in a wide range of phenomena, from long-term migration patterns to communication volume between different regions [33].

In the next sections we show that for a security manager, the problem of computing the best allocation plan of security monitors for a changing IoT environment can be interpreted as the problem an investor faces when choosing an investment that maximizes the return, considering the continuous changes that affect the market.

3 Shortfall

In this section we briefly introduce the concept of shortfall, clarify the motivations that have lead us to adopt it in our formalization, and show how we manipulate it in order to obtain a more appropriate solution to our situation.

3.1 Definition

Shortfall is a risk measure used in economics which has conceptual, computational and practical advantages over other commonly used risk measures [3]. The shortfall, or more precisely, the *shortfall at level α* , measures how large losses,

below the expected return, can be expected if the return of the investment drops below its α -quantile. Given all the possible market trends, we can compute the set of all possible returns of our investment. The shortfall at level α is the difference between the expected return and the average of the returns of the worst $\alpha\%$ of cases. Given an investment x and a value $\alpha \in (0, 1)$, the shortfall $s_\alpha(x)$ is defined as follows:

$$s_\alpha(x) = E[R] - E[R|R \leq q_\alpha(R)] \quad (1)$$

where $E[R]$ is the expected return of the investment, and $q_\alpha(X)$ is the α -quantile of a random variable X :

$$q_\alpha(X) = \inf\{x|P(X \leq x) \geq \alpha\} \quad (2)$$

According to Levy and Kroll [18], for an investment x chosen to minimize $s_\alpha(x)$ for a fixed α and a given target mean μ_t , there is no other investment with the same mean which would be preferred to x , because less profitable in the worst $\alpha\%$ of cases. Thus, one is naturally led to minimize the quantity $s_\alpha(x)$ for some $\alpha \in (0, 1)$ as follows:

$$\begin{aligned} \min \quad & s_\alpha(x) \\ \text{subject to} \quad & E[R] = \mu_t \end{aligned} \quad (3)$$

3.2 From the minimum shortfall to the best choice

The solution we obtain by solving Problem 3 depends on the target mean μ_t we choose. In fact, given a target mean, we obtain the best investment (i.e., the one with minimum shortfall) among those with the same expected value, but we do not know whether there exists a better investment for different values of μ_t . In other words, there might exist a *dominating* investment, i.e., of a higher mean and a smaller shortfall, or alternatively, a *non-dominated* investment, i.e., of a higher (resp. lower) mean and a greater (resp. smaller) shortfall, which may be preferred by the investor. A possible solution is that of solving Problem 3 with different μ_t , and finally select the investment that best meets our needs. However, this solution is not efficient because it might compute some dominated investment, i.e., which would not be preferred to any other because of a lower expected value and a greater shortfall than those of the other ones. Consequently the time needed to compute such dominated solutions would be wasted. As alternative, we propose to adapt Problem 3 in such a way that it computes only non dominated solutions. This is possible by turning Problem 3 into a bi-objective optimization problem. First of all, we start considering the following equivalent problem:

$$\begin{aligned} \max \quad & E_\alpha[R] \\ \text{subject to} \quad & E[R] = \mu_t \end{aligned} \quad (4)$$

where $E_\alpha[R] = E[R|R \leq q_\alpha(R)]$ is the mean of the worst $\alpha\%$ of cases, namely,

the subtrahend in Equation 1. Problem 4 is equivalent to Problem 3 because for a target mean μ_t , minimizing $E[R] - E_\alpha[R]$ is equivalent to maximizing $E_\alpha[R]$. Since we want to compute a set of non dominated investments based on their $E[R]$ and $E_\alpha[R]$ values, we place the expected value as an objective of the optimization problem, and we impose two inequality constraints on the values of $E[R]$ and $E_\alpha[R]$, in order to restrict the computation only to the cases we are interested in. The resulting problem is as follows:

$$\begin{aligned} \max \quad & E[R], E_\alpha[R] \\ \text{subject to} \quad & E[R] \geq \mu_t \\ & E_\alpha[R] \geq \mu_\alpha \end{aligned} \tag{5}$$

Problem 5 is a less restricted version of Problem 3, where we split the two terms of the shortfall, $E[R]$ and $E_\alpha[R]$, and optimize them separately. The output is a Pareto frontier PF [19] whose points are non dominated investments x with values of $E[R]$ and $E_\alpha[R]$ below the desired thresholds. Formally, PF is defined as follows:

$$PF = \{(E[R], E_\alpha[R]) \mid \nexists (E'[R], E'_\alpha[R]) \in PF \text{ such that} \\ (E[R] \geq E'[R] \wedge E_\alpha[R] \geq E'_\alpha[R]) \wedge (E[R] > E'[R] \vee E_\alpha[R] > E'_\alpha[R])\}$$

An investor can thus conduct a cost-benefit analysis on the solutions of the Pareto frontier, and finally choose the one that best fits his/her requirements. In the next sections we show that the shortfall is well suited for describing the risk associated with an allocation plan, such that a certain level of security is ensured for every variation in the topology of the monitored environment.

4 Problem Definition

4.1 Preliminaries

We define an IoT environment as a set of n geographic areas $R_k, k = 1, \dots, n$, that we call *regions*. We divide each region in *locations*, i.e., space units where IoT devices or security resources could reside.

Definition 1 (region). A region R is a tuple $\langle L_R, P_R \rangle$, where:

- L_R is the set of locations of R ;
- P_R is the probability distribution over a discrete random variable X that takes integer values in the interval $[0, \infty)$. $P_R(X = x)$ is the probability that there are x devices in R .

Definition 2 (security resource). A security resource sr is a tuple $\langle c, r, P_{sr} \rangle$, where:

- c is the cost;
- r (from radius) is the maximum action range;

- P_{sr} is the probability distribution over a discrete random variable Y that takes values in the interval $[0, r]$. $P_{sr}(Y = y)$ is the probability that sr is able to monitor a device located at a distance y .

Furthermore, we define $\mathcal{R} = \bigcup_{k=1}^n R_k$ as the set of all regions, and $\mathcal{L} = \{loc : loc \in \mathcal{R}\}$, as the set of all locations in $\mathcal{R}$. In the rest of the paper we will use the notation $sr.c$ to denote the cost c of a security resource sr .

Security resources have an associated probability distribution, that describes how likely a device at a certain distance is seen by a resource. This choice is driven by the fact that the typical assumptions about all radios having circular range and perfect coverage in that range is far from real [17]. More realistic models take into account antenna height and orientation, terrain and obstacles, surface reflection and absorption, and so forth. It is often difficult in reality to estimate whether or not one has a functioning radio link between nodes, because signals fluctuate greatly due to mobility and fading as well as interference. Several signal attenuation models have thus been proposed [9, 25], and there are a few that can be used in concrete implementations of our model. For example, the Okumura model, the Hata model for urban areas, the Hata model for open areas, are models for outdoor attenuation; instead, the ITU model and the Log-distance path loss model are models for indoor attenuation. One of these models, or a combination of them, can be adopted for modeling the attenuation of a device according to its radio characteristics.

4.2 From economics to the IoT domain

The problem we face is that of computing the optimal security resource allocation plan for IoT scenarios characterized by a high degree of mobility. This problem can be reduced to Problem 5 if we consider an allocation plan along with its cost as an investment, and all the possible configurations a set of IoT devices may assume in the area of interest as all the possible market trends. A security manager (the investor) knows that the number of devices moving within a region R_k follows a certain probability distribution P_{R_k} , and that the resource allocation plan to choose must be able to provide a good security level not only in the average cases of P_{R_k} , but also in the rare ones. In fact, rare device configurations may result in high losses if not addressed by an adequate security infrastructure. We use $E[R]$ and $E_\alpha[R]$ as evaluation metrics of an allocation plan for an IoT environment, where the multitude of device configurations is wide because mobility. The optimal plan is the one that minimizes the expected number of IoT devices not reached by any security resource. We call this quantity the *risk* associated with the allocation plan (AP), denoted by $risk_{AP}$. We adopt the risk as a negative return of an allocation plan, such that the smaller the risk is, the more effective the plan. Being the risk a negative return, we need to adapt the definition of α -quantile given in Equation 2 as follows:

$$q_\alpha(risk) = \inf\{x | P(risk \geq x) \geq \alpha\} \quad (6)$$

Given a set of regions $\mathcal{R}$ (i.e., an IoT environment), and a fixed security budget b , we define an allocation plan for $\mathcal{R}$ as the set

$$AP = \{(sr_i, loc_j) | \sum_i sr_i.c \leq b, loc_j \in \mathcal{L}\} \quad (7)$$

Given the set of all possible device configurations, we can compute the set of all possible returns of an AP in terms of risk, and thus compute $E[risk_{AP}]$ and $E_\alpha[risk_{AP}]$. The adaptation of Problem 5 to our application domain is as follows:

$$\begin{aligned} & \max_{AP \in \mathcal{AP}} -E[risk_{AP}], -E_\alpha[risk_{AP}] \\ & \text{subject to} \\ & E[risk_{AP}] \leq \mu_t \\ & E_\alpha[risk_{AP}] \leq \mu_\alpha \\ & \sum_{sr \in AP} sr.c \leq b \end{aligned} \quad (8)$$

where $\mathcal{AP}$ is the set of all possible allocation plans. The minus sign before the two objectives means that we are actually minimizing the two measures, being the risk a negative return. For the same reason, μ_t and μ_α are upper bounds for $E[risk_{AP}]$ and $E_\alpha[risk_{AP}]$, respectively, and not lower bounds as in Problem 5. Note that without the constraint on the maximum budget, Problem 8 becomes meaningless because it would compute a unique allocation plan possibly with an unreasonable number of resources, providing risk 0 for any device configuration, which is surely the most effective solution, but also unlikely the most efficient because it would be very expensive in terms of cost.

Once a security manager has obtained the set of solutions to Problem 8, (s)he can choose the allocation plan that best fits his/her security requirements. However, there may be cases for which there is no solution, i.e., solving Problem 8 outputs an empty set. This can happen for two reason:

1. no allocation plan exists able to satisfy the constraints, because the security budget is too small to provide a value of $E[risk_{AP}]$ and $E_\alpha[risk_{AP}]$ below the desired thresholds;
2. the α -quantile is too small to admit $E_\alpha[risk_{AP}] \leq \mu_\alpha$.

In case (1), the security manager can try to increase the security budget until at least one solution is returned, or alternatively can turn Problem 8 in a three-objective optimization problem by placing the cost as a further objective as follows:

$$\begin{aligned}
& \max_{AP \in \mathcal{AP}} -E[risk_{AP}], -E_\alpha[risk_{AP}], -cost_{AP} \\
& \text{subject to} \\
& E[risk_{AP}] \leq \mu_t \\
& E_\alpha[risk_{AP}] \leq \mu_\alpha \\
& minCost \leq cost_{AP} \leq maxCost
\end{aligned} \tag{9}$$

where $cost_{AP} = \sum_{sr \in AP} sr.c$. This way, the security manager can have a comprehensive view of cost-benefit, by determining the minimum cost required to ensure a certain security level. However, this problem takes more time to solve than Problem 8, proportionally to the quantity $maxCost - minCost$.

In case (2), the security manager may decide to increase μ_α , or to tolerate a higher risk by having a larger quantile. In fact, α can be intended as a “measure of the risk tolerance”: by increasing α , the value of $q_\alpha(risk)$ decreases, and Problem 8 computes $E_\alpha[risk]$ over a larger set of cases, which means that cases with high risk are more tolerated. On the contrary, when α decreases, $q_\alpha(risk)$ tends to the maximum value of risk, thus restricting the cases over which $E_\alpha[risk]$ is computed, meaning that we tolerate fewer cases with high risk.

5 Evaluation Algorithm

To solve Problem 8 an algorithm should enumerate all possible allocation plans $\mathcal{AP}$, and evaluate each $AP \in \mathcal{AP}$ over all possible configurations the IoT devices can take within the geographic area of interest $\mathcal{R}$. It is easy to see that such an algorithm would take an unreasonable amount of time to solve the problem. In fact, for each region $R_k \in \mathcal{R}$, the number of possible configurations is the cardinality of the power set of L_{R_k} , and the size of $\mathcal{AP}$ is upper-bounded by the number of all subsets of $\mathcal{L}$ of cardinality at most s , where s is the maximum number of security resources allowed by the security budget. This leads us to look for approximations in place of an exact solution. A standard tool for approximation is the use of genetic algorithms. There exist several genetic algorithms for computing a Pareto frontier, one of the most commonly used is NSGA-II [8], an evolutionary genetic algorithm able to find an approximation of the Pareto frontier for multi-objective optimization problems. The main difference between NSGA-II and other evolutionary genetic algorithms is the selection phase: NSGA-II has no unique fitness function but one for each objective. During the selection phase, the selected points are only the non-dominated ones. When some constraints exist the selection phase will remove all points that do not satisfy the constraints.

In Problem 8 there are two objectives $-E[risk_{AP}]$ and $-E_\alpha[risk_{AP}]$. To compute their fitness at each iteration, the NSGA-II algorithm needs to evaluate the current AP over the set of all possible device configurations. An AP is implemented as an individual of the population of the genetic algorithm, where each gene is the *id* of the location where a security resource has been placed.

Algorithm 1

```
1: procedure COMPUTEFITNESS( $AP, n$ )
2:    $\vec{risk}_{AP} = \text{MONTECARLO}(AP, n)$ ;
3:    $q_\alpha = \alpha$ -quantile of  $\vec{risk}_{AP}$ ;
4:    $E[risk_{AP}] = 0$ ;
5:    $E_\alpha[risk_{AP}] = 0$ ;
6:    $count = 0$ ;
7:   for  $r \in \vec{risk}_{AP}$  do
8:      $E[risk_{AP}] += r$ ;
9:     if  $r \geq q_\alpha$  then
10:       $E_\alpha[risk_{AP}] += r$ ;
11:       $count ++$ ;
12:     end if
13:   end for
14:    $E[risk_{AP}] = E[risk_{AP}] / \vec{risk}_{AP}.length$ ;
15:    $E_\alpha[risk_{AP}] = E_\alpha[risk_{AP}] / count$ ;
16:   return  $E[risk_{AP}], E_\alpha[risk_{AP}]$ ;
17: end procedure
```

Algorithm 2

```
1: procedure MONTECARLO( $AP, n$ )
2:    $\vec{risk}_{AP} = []$ ;
3:   while  $n > 0$  do
4:      $devicesLocations = \emptyset$ ;
5:     for  $R_k \in \mathcal{R}$  do
6:        $numberOfDevices = P_k.nextInt()$ ;
7:       while  $numberOfDevices > 0$  do
8:          $loc = \text{a location} \in R_k$  chosen randomly;
9:          $devicesLocations.add(loc)$ ;
10:         $numberOfDevices --$ ;
11:      end while
12:    end for
13:     $risk = \text{COMPUTERISK}(AP, devicesLocations)$ ;
14:    add  $risk$  to  $\vec{risk}_{AP}$ ;
15:     $n --$ ;
16:  end while
17:  return  $\vec{risk}_{AP}$ ;
18: end procedure
```

As stated before, for each region R_k the set of different configurations C_k is too big to be employed in the evaluation process without incurring in scalability problems. In its place we generate a subset of C_k with a Monte Carlo method [28], whose elements are generated according to the probability distribution P_{R_k} associated with the region R_k . More precisely, for a given region R_k , a device configuration is a set of locations chosen randomly, with cardinality equal to an integer randomly generated from P_{R_k} . Algorithms 1 and 2 illustrate the fitness function NSGA-II used for evaluating an allocation plan. Algorithm 1 takes as input an allocation plan AP and an integer n . In Line 2, the MONTECARLO procedure returns a vector of length n with the risks associated with AP evaluated over n different device configurations. In Line 3, q_α is the value of the α -quantile of the vector $\vec{risk}_{AP}$. In Algorithm 2, the MONTECARLO procedure produces n different device configurations over which it evaluates AP . The number of devices for each region R_k is chosen randomly according to the probability distribution P_{R_k} (Line 6). In Line 13, the procedure COMPUTERISK computes the number of

devices not monitored by any security resource $sr \in AP$ for each configuration, according to P_{sr} (the probability that sr is able to monitor a device located at a certain distance). The exact positions where to place the sensors are identified by each single gene of the individual corresponding to the chosen Pareto point.

6 Results Analysis

In this section we show the steps a security manager has to follow for computing the allocation plan that best fits his/her security requirements, according to his/her security budget. To this end, we report results related to the simulation of an area of 2000 square meters and consisting of three regions divided in locations of 1 square meter. The details are shown in the following figure:

$\text{Gamma}(\alpha, \beta)$ is the Gamma probability distribution, adopted for simulating the devices distribution over each region.

Suppose the security manager has at his/her disposal a security budget of 30 monetary units, and only one type of security resource, based on the watchdog mechanism, with action range of 8 space units, and unitary cost. The security level (s)he wants to achieve is at most 30 unmonitored devices in no more than 10% of cases. This is formally translated as:

- $\forall$ security resource $sr, sr = \langle 1, 8, LDPL \rangle$;
- $\alpha = 0.1$;
- $b = 30$;
- $\mu_t = \infty, \mu_\alpha = 30$.

where $LDPL$ is the Log-distance path loss model [9], adopted as the signal attenuation model for the security resources. In order to compute the set of plans that satisfy his/her security constraints, the security manager solves Problem 8 with an NSGA-II algorithm, and set the Monte Carlo simulation with $n = 100K$ (see the MonteCarlo procedure shown in Algorithm 2) such that each plan is evaluated over a set of 100K different device configurations. Figure 1 shows the Pareto frontier (left), and the risk distribution of its plans (right). Looking at the Pareto frontier the security manager would choose the AP having the best combination of $E[risk]$ and $E_\alpha[risk]$. $AP_{30.1}$ would be a better plan than $AP_{30.2}$ and $AP_{30.3}$ because $E[risk]$ is almost the same in the three cases and

$E_\alpha[risk]$ is 25% lower. Looking at the risk distributions of Figure 1 (right), (s)he can also base his/her decision on additional information and conduct a more in-depth analysis. Over 100K different device configurations, each plan provides a minimum risk of 0 and a maximum risk of 45. $AP_{30.2}$ is the plan that provides more cases with risk 0, thus a lower probability of an attack to happen in the network area. $AP_{30.1}$, although it provides the lowest number of cases with risk in the range bounded by the α -quantile, it also provides the lowest number of cases with risk 0, thus the highest probability to have an attack.⁶ The trade-off between $E[risk]$, $E_\alpha[risk]$, and the probability of having an attack, is at the basis of the reasoning a security manager may want to follow to choose a plan. In this situation, a security manager interested in minimizing the cases with $risk > 0$ would choose $AP_{30.2}$, since $AP_{30.3}$ provides similar values of $E[risk]$ and $E_\alpha[risk]$ but fewer cases with $risk = 0$.

Fig. 1. Pareto frontier (left), and the risk distribution of its allocation plans (right).

Fig. 2. The relation between cost, $E[risk]$, $E_\alpha[risk]$, and the probability to have an attack, of plans computed with $\alpha = 0.1$, $\mu_t = \infty$, $\mu_\alpha = 30$, $minCost = 15$ and $maxCost = 132$.

Now suppose that the security manager wants to know how much to rise the budget to achieve higher security performances. In this case (s)he can solve Problem 9 and plot the values of $E[risk]$ and $E_\alpha[risk]$ along with the attack

⁶ The probability to have an attack is computed as the ratio between the number of cases with $risk > 0$ and the total number of cases.

probability as shown in Figure 2. It can be noticed that, although $E[risk]$ and $E_\alpha[risk]$ decrease of two orders of magnitude from $cost = 15$ to $cost = 40$, the probability of having an attack remains around 1, which means that there are very few cases (possibly none) with $risk = 0$. This way, the security manager would know (s)he needs a security budget greater than 40 for ensuring a higher level of security.

7 Case Study

In this section we show how we applied our method to a real case scenario, and demonstrate its validity by comparing it with other heuristics.

7.1 The Dartmouth College Data

In [16] Henderson et al. analyze an extensive network trace from a mature 802.11 WLAN, including more than 550 Access Points (AcP) and 7000 users over seventeen weeks in the Dartmouth College area. This work reports several statistics about the movement of people equipped with various IoT devices within the campus area. All these data plus some additional data collected covering more than 160 weeks over the same population have been made publicly available.⁷ We used this dataset for testing our method on a simulation of a real IoT environment, intended as a set of regions (the area covered by each AcP), and with real device distributions. Those were computed as follows. Each location data entry is associated with the MAC address of a mobile device that generated the data point. By grouping the devices that were connected at the same hour of the day in the same AcP, we counted how many devices were observed in each hour over the entire period in each AcP. We also added up the number of different devices observed during a day (a period of 24 hours) in each AcP. In this way we could compute the device distribution for each region/AcP of the campus area. Other experimental settings are as follows: we divided each region/AcP in locations of 1 square meter, for a total of 183279 square meters; we adopted watchdog-like devices as security resources with an action range of 15 meters and unitary cost; we adopted *LDPL* [9] as the signal attenuation model for the security resources; we set $\alpha = 0.1$; we used the data of the first 3 months as training set for computing the Pareto frontier. Finally, we tested the plan of the Pareto frontier with the lowest $E[risk]$ on the data relative to the remaining 21 months.

7.2 Evaluation Methodology

In order to evaluate our method, we compare the allocation plans obtained by running Problem 8 on the IoT environment described above, with those obtained by using other heuristics that follow different approaches for solving the placement problem. We implemented four heuristics: (1) *square lattice*, (2) *triangular lattice*, (3) *max coverage*, and (4) *greedy*.

⁷ <http://crawdad.org/dartmouth/campus/20090909>

Fig. 3. Values of $E[risk]$ (top-left), $E_{\alpha}[risk]$ (top-right), and attack probability (bottom), computed for different values of cost, with five different heuristics.

Methods (1) and (2) are placements typically used by system engineers to model the structure of a set of base stations (BSs) in cellular networks [4]. We have chosen those models because the BSs placement problem is similar to the problem of placing security monitors in an IoT environment, since both mobile phones and IoT devices are mobile, while BSs and security monitors are not. Method (3) is implemented as a single-objective optimization problem solved by an NSGA-II algorithm with the objective of maximizing the number of locations that fall under the action range of security resources, given a fixed budget. Method (4) is an algorithm that computes the placement in a greedy manner, i.e., given a fixed budget b , at each iteration it chooses the security resource which maximizes the number of covered locations without exceeding b . We ran experiments with different security budgets on the network area built, and with the device distributions computed from the Dartmouth College dataset.⁸ The results are reported in Figure 3. It can be noticed that our method computes allocation plans with $E[risk]$, $E_{\alpha}[risk]$ and an attack probability lower than those provided by the other methods. Furthermore, we recall that evolutionary algorithms compute solutions that are an approximation of the optimal one.

⁸ Method (4) was run with budgets not greater than 1298, i.e., the cheapest budget for covering all locations; running such a greedy algorithm with a budget greater than 1298 is meaningless because the objective of covering as much locations as possible would be already reached.

The results of Figure 3 were obtained by running the NSGA-II algorithm with 50 generations, thus one can obtain better plans by just rising the number of generations, or also by adjusting other parameters like the mutation method, mutation probability, and algorithm seed.

8 Related Work

We are not the first to use shortfall to characterize risk in a security context. Molloy et al. [23] adapt shortfall to take security decisions (such as access-control decisions or spam filtering decisions) under uncertainty when the benefit of doing so outweighs the need to absolutely guarantee that these decisions are correct. Molloy et al. have also put forward a more general vision on economic models for security [22]. We also would like to mention the risk-aware security solutions by Chen et Crampton incorporating the notion of risk into the RBAC model [5]. The problem of computing the best placement of network devices for mobile scenarios has been mostly addressed in the area of cellular networks. In such networks, the optimal placement of base stations (BSs) is crucial for the correct functioning of the communication system. The similarity with our scenarios is in the presence of mobile devices that have to be covered by the action range of static entities. However, the roaming problem, that has to be taken into account in the computation of the optimal BS placement, makes those solutions not suitable for the IoT scenario. We refer the reader to [11] for a more comprehensive discussion on this topic. The problem of finding efficient security solutions in the domain of computer networks with the help of the Pareto analysis has been extensively investigated. In [31] the Pareto analysis has been used to compute the best combination between cost of patching vulnerabilities and cost of deactivating products within an enterprise system. Chang and Zhuang propose an approach based on a node clustering algorithm, with effective tax-based sub-carrier allocation, tailored for wireless mesh networks with QoS support [6]. Here, the Pareto analysis is used for the optimal resource management. Signal attenuation models have also been widely investigated. Shen et al. propose an indoor wireless propagation model in WiFi radio-over-fiber network architecture for received signal strength (RSS) based localization in the IoT [32]. The proposed model adds attenuation terms of obstacles in each sub-space by dividing the room into several sub-spaces according to the obstacles' distribution. Alwajeeh et al. [2] propose an intelligent method to associate known models with spatial zones according to the electromagnetic interactions. Past work on IoT security focuses on protection mechanisms against specific attacks [1, 35, 38], investigates ISO/OSI layer-related security problems [15], or proposes architectures for intrusion detection, attack prevention, or recovery systems [27, 29]. Raza et al. [26] propose an IPsec extension of 6LoWPAN, and show that IPsec is a feasible option for securing the IoT in terms of packet size, energy consumption, memory usage, and processing time. Chigan et al. [7] propose a resource-aware self-adaptive network security provisioning scheme for the resource constraint Mobile Ad-hoc Networks (MANET), in order to avoid security provisioning De-

nial of Service (SPDoS) attack. While we model efficiency and effectiveness of a security solution as cost and risk, Chigan et al. model them as two indexes, performance index (PI) and security index (SI), respectively. SI quantitatively reflects the security contribution of a secure protocol to a MANET system, while PI quantitatively reflects network performance perspectives of a secure protocol. The problem of finding the optimal security resource allocation plan for IoT networks has already been investigated by Rullo et al. [30]. However, such approach only works for static networks for the topology changes due to the occasional (dis)appearance of devices.

9 Conclusion

In this paper we built upon the concept of security resources specialized for the IoT. These security resources monitor traffic generated by devices possibly using different protocols and wireless networks. Our proposed framework takes into consideration the possibility that the security resources are not able to monitor all the devices at all times. Our framework provides different security resource deployment plans based on cost and risk. We measure the risk associated with an allocation plan as the expected number of devices left uncovered. Considering the cost of the plan as an investment, and the expected risk as the expected return of the investment, we were able to borrow the concept of shortfall from economics, that looks at the worst possible outcomes of an investment to make better decisions from investments with similar expected returns. Hence, our framework produces resource allocation plans that are the outcomes of bi-objective optimization of expected risks and expected worst risk scenarios. We have developed a reference implementation of the framework and ran simulations that incorporate mobility patterns from real data. The reference implementation uses standard genetic algorithms to find approximate solutions to the optimization problem. In our experimental setting we employed just one type of security resource for simplicity, but with some modeling effort, NSGA-II algorithm can be tuned to work with individuals of different types, such that a wider range of security scenarios can be addressed.

Acknowledgment. Jorge Lobo was partially supported by the Secretaria de Universitats i Recerca de la Generalitat de Catalunya, the Maria de Maeztu Units of Excellence Programme and the Spanish Ministry for Economy and Competitiveness (MINECO) under Grant Ref.: TIN2016-81032-P.

References

1. Altman, E., Avrachenkov, K., Gamaev, A.: Jamming in wireless networks: The case of several jammers. In: Proc. of the First ICST Int'l Conf. on Game Theory for Networks (2009)

2. Alwajeeh, T., Combeau, P., Bounceur, A., Vauzelle, R.: Efficient method for associating radio propagation models with spatial partitioning for smart city applications. In: *Proceedings of the International Conference on Internet of things and Cloud Computing*. p. 8. ACM (2016)
3. Bertsimas, D., Lauprete, G.J., Samarov, A.: Shortfall as a risk measure: properties, optimization and applications. *Journal of Economic Dynamics and control* 28(7), 1353–1381 (2004)
4. Charoen, P., Ohtsuki, T.: Codebook based interference mitigation with base station cooperation in multi-cell cellular network. In: *Vehicular Technology Conference (VTC Fall)*, 2011 IEEE. pp. 1–5. IEEE (2011)
5. Chen, L., Crampton, J.: Risk-aware role-based access control. In: *International Workshop on Security and Trust Management*. pp. 140–156. Springer (2011)
6. Cheng, H.T., Zhuang, W.: Pareto optimal resource management for wireless mesh networks with qos assurance: joint node clustering and subcarrier allocation. *IEEE Trans. on Wireless Communications* (2009)
7. Chigan, C., Li, L., Ye, Y.: Resource-aware self-adaptive security provisioning in mobile ad hoc networks. In: *IEEE Wireless Communications and Networking Conference*, 2005. vol. 4, pp. 2118–2124. IEEE (2005)
8. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast elitist multi-objective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation* 6, 182–197 (2000)
9. Goldhirsh, J., Vogel, W.J.: *Handbook of propagation effects for vehicular and personal mobile satellite systems*. NASA Reference Publication 1274, 40–67 (1998)
10. Gonsalves, A.: New toolkit seeks routers, internet of things for ddos botnet. <http://www.csoonline.com/article/2687653/data-protection/new-toolkit-seeks-/routers-internet-of-things-for-ddos-botnet.html> (2014), accessed: May 2016
11. Guo, A., Haenggi, M.: Spatial stochastic models and metrics for the structure of base stations in cellular networks. *IEEE Transactions on Wireless Communications* 12(11), 5800–5812 (2013)
12. Huang, Y.a., Lee, W.: A cooperative intrusion detection system for ad hoc networks. In: *Proceedings of the 1st ACM Workshop on Security of Ad Hoc and Sensor Networks*. pp. 135–147. SASN '03, ACM, New York, NY, USA (2003), <http://doi.acm.org/10.1145/986858.986877>
13. Hui, J., Culler, D., Chakrabarti, S.: 6lowpan: Incorporating ieee 802.15. 4 into the ip architecture. *IPSO Alliance White Paper* 3 (2009)
14. IEEE: Ieee 802.15 wpan task group 4 (tg4). <http://www.ieee802.org/15/pub/TG4.html>
15. Jinwala, D., Patel, D., Dasgupta, K.: Flexisec: a configurable link layer security architecture for wireless sensor networks (2012)
16. Kotz, D., Henderson, T., Abyzov, I., Yeo, J.: CRAWDAD dataset dartmouth/campus (v. 2009-09-09) (Sep 2009)
17. Kotz, D., Newport, C., Gray, R.S., Liu, J., Yuan, Y., Elliott, C.: Experimental evaluation of wireless simulation assumptions. In: *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*. pp. 78–82. ACM (2004)
18. Levy, H., Kroll, Y.: Ordering uncertain options with borrowing and lending. *The Journal of Finance* 33(2), 553–574 (1978)
19. Messac, A., Ismail-Yahaya, A., Mattson, C.A.: The normalized normal constraint method for generating the pareto frontier. *Structural and multidisciplinary optimization* 25(2), 86–98 (2003)

20. Midi, D., Rullo, A., Mudgerikar, A., Bertino, E.: Kalis: A system for knowledge-driven adaptable intrusion detection for the internet of things. *IEEE 37th International Conference on Distributed Computing Systems (ICDCS)* (2017)
21. Mishra, A., Nadkarni, K., Patcha, A.: Intrusion detection in wireless ad hoc networks. *IEEE Wireless Communications* 11(1), 48–60 (Feb 2004)
22. Molloy, I., Cheng, P.C., Rohatgi, P.: Trading in risk: Using markets to improve access control. In: *Proceedings of the 2008 workshop on New security paradigms*. pp. 107–125. *ACM* (2009)
23. Molloy, I., Dickens, L., Morisset, C., Cheng, P.C., Lobo, J., Russo, A.: Risk-based security decisions under uncertainty. In: *Proceedings of the second ACM conference on Data and Application Security and Privacy*. pp. 157–168. *ACM* (2012)
24. Nasipuri, A., Li, K.: A directionality based location discovery scheme for wireless sensor networks. In: *Proc. of 1st ACM Int'l workshop on Wireless sensor networks and applications*. *ACM* (2002)
25. Rappaport, T.S., et al.: *Wireless communications: principles and practice*, vol. 2. Prentice Hall PTR New Jersey (1996)
26. Raza, S., Duquennoy, S., Höglund, J., Roedig, U., Voigt, T.: Secure communication for the internet of things: a comparison of link-layer security and ipsec for 6lowpan. *Security and Communication Networks* (2012)
27. Raza, S., Wallgren, L., Voigt, T.: Svelte: Real-time intrusion detection in the internet of things. *Ad hoc networks* (2013)
28. Robert, C., Casella, G.: *Monte Carlo statistical methods*. Springer Science & Business Media (2013)
29. Roman, R., Alcaraz, C., Lopez, J., Sklavos, N.: Key management systems for sensor networks in the context of the internet of things. *Computers & Electrical Engineering* 37(2), 147–159 (2011)
30. Rullo, A., Midi, D., Serra, E., Bertino, E.: A game of things: Strategic allocation of security resources for iot. *ACM/IEEE 2nd International Conference on Internet of Things Design and Implementation (IoTDI 2017)* p. 6 (2017)
31. Serra, E., Jajodia, S., Pugliese, A., Rullo, A., Subrahmanian, V.: Pareto-optimal adversarial defense of enterprise systems. *ACM Transactions on Information and System Security (TISSEC)* (2015)
32. Shen, X., Xu, K., Sun, X., Wu, J., Lin, J.: Optimized indoor wireless propagation model in wifi-rof network architecture for rss-based localization in the internet of things. In: *Microwave Photonics, 2011 International Topical Meeting on & Microwave Photonics Conference, 2011 Asia-Pacific, MWP/APMP*. pp. 274–277. *IEEE* (2011)
33. Simini, F., González, M.C., Maritan, A., Barabási, A.L.: A universal model for mobility and migration patterns. *Nature* 484(7392), 96–100 (2012)
34. Sultana, S., Midi, D., Bertino, E.: Kinesis: a security incident response and prevention system for wireless sensor networks. In: *Proc. of ACM SensSys* (2014)
35. Tumrongwittayapak, C., Varakulsiripunth, R.: Detecting sinkhole attack and selective forwarding attack in wireless sensor networks. In: *ICICS 2009*. (2009)
36. Weber, R.H., Weber, R.: *Internet of Things*, vol. 12. Springer (2010)
37. Xia, F., Yang, L.T., Wang, L., Vinel, A.: Internet of things. *International Journal of Communication Systems* 25(9), 1101 (2012)
38. Zhu, Q., Li, H., Han, Z., Basar, T.: A stochastic game model for jamming in multi-channel cognitive radio systems. In: *IEEE ICC* (2010)
39. ZigBee Alliance and others: *Zigbee specification* (2006)