

BARTER: Behavior Profile Exchange for Behavior-Based Admission and Access Control in MANETs

Vanessa Frias-Martinez^{1*}, Salvatore J. Stolfo², and Angelos D. Keromytis²

¹ Telefónica Research, Madrid, Spain

² Computer Science Department, Columbia University, New York, USA
vanessa@tid.es, {sal,angelos}@cs.columbia.edu

Abstract. Mobile Ad-hoc Networks (MANETs) are very dynamic networks with devices continuously entering and leaving the group. The highly dynamic nature of MANETs renders the manual creation and update of policies associated with the initial incorporation of devices to the MANET (*admission control*) as well as with anomaly detection during communications among members (*access control*) a very difficult task. In this paper, we present *BARTER*, a mechanism that automatically creates and updates admission and access control policies for MANETs based on behavior profiles. *BARTER* is an adaptation for fully distributed environments of our previously introduced *BB-NAC* mechanism for NAC technologies. Rather than relying on a centralized NAC enforcer, MANET members initially exchange their behavior profiles and compute individual local definitions of normal network behavior. During admission or access control, each member issues an individual decision based on its definition of normalcy. Individual decisions are then aggregated via a threshold cryptographic infrastructure that requires an agreement among a fixed amount of MANET members to change the status of the network. We present experimental results using content and volumetric behavior profiles computed from the ENRON dataset. In particular, we show that the mechanism achieves true rejection rates of 95% with false rejection rates of 9%.

1 Introduction

Mobile Ad-Hoc Networks (MANETs) are composed of devices that enter and leave the network dynamically, quickly changing the network topology and administrative domain membership. MANETs differ from wired/wireless networks in that there is no central control, no base station, and no wireless switches. As a result, any task in the network must be distributed and executed by all its members. These tasks include manually creating and updating policies for the admission as well as the access control of devices over time. Admission control refers to the decision process prior to the incorporation of devices to the

* Work performed while a PhD student at Columbia University.

MANET. On the other hand, access control involves the membership update of devices that are already part of the MANET. In general, admission and access control policies are difficult to create manually unless one has a profound understanding of the resource that needs to be controlled. Additionally, the update of policies is even more difficult given the highly dynamic nature of MANETs.

In our previous work, we introduced *BB-NAC*, a behavior-based network admission and access control mechanism for NAC technologies that centralized the decision process on a unique NAC enforcer located at the edge of the network [3] [4] [5]. Behavior was intended to represent the typical communications of network devices *i.e.*, the traffic payload observed or specific volumetric measurements of the traffic such as average number of packets. In this paper, we present BARTER, a behavior-based admission and access control mechanism for MANETs. BARTER is an adaptation of BB-NAC for fully distributed networks. As in the BB-NAC mechanism [5], a newcomer would present its behavior profile to the MANET members during admission control. If an agreement is reached among the members, the newcomer is admitted into the MANET. Analogously, during access control, the traffic exchanged would be checked against the behavior profiles of similar MANET members to perform anomaly detection.

Unlike BB-NAC, the admission and access control decisions in BARTER are distributed among the MANET members rather than being centrally performed by a NAC enforcer. The decision of each individual MANET member is based on the accumulation of knowledge gathered from the behavior profiles of other members. Ultimately, the final admission or access control decision is achieved by building BARTER on top of a threshold cryptographic infrastructure that guarantees not only distributed decision making but also secure communications among MANET members. Due to the limited computational resources of many MANET platforms (such as cellphones or PDAs), the calculation of *clusters of behavior profiles* similar to the one implemented in the BB-NAC mechanism would not be feasible. Instead, BARTER takes advantage of the restrictions imposed by the threshold cryptographic infrastructure as a way to approximate groups of similar behavior within the network.

Apart from the full description of the mechanism, we present an experimental evaluation of BARTER based on content and volumetric behavior profiles computed from the ENRON dataset [2]. Throughout the paper, we assume that there exists a tamper resistance scheme [6] [11] running in the MANET that prevents devices from having multiple identifications (each device has a unique, identifiable ID_i) and that detects manipulations in the packets exchanged between MANET members.

The main contributions of the BARTER mechanism are the following:

- A mechanism that provides automatic and fully distributed creation of admission and access policies for MANETs. Individual decisions are made by each MANET member based on the knowledge accumulated from previous profile exchanges among members. The final admission or access control decision is determined from the aggregation of individual decisions using a threshold cryptographic layer that runs under the BARTER mechanism.

- A mechanism that is robust against attacks from MANET members. The mechanism adjusts over time in order to maintain its robustness even in the presence of malicious devices within the MANET.
- An extensive evaluation of the mechanism using hundreds of content and volumetric behavior profiles computed from the ENRON dataset.

The paper is organized as follows: in Section 2 we describe the foundations of the BARTER mechanism. Section 3 discusses possible attacks to the mechanism and analyzes the costs incurred by the threshold cryptographic infrastructure. Section 4 and Section 5 describe the experimental evaluation for content and volumetric profiles respectively. Section 6 summarizes related work. Finally, Section 7 presents conclusions and future work.

2 The BARTER Mechanism

We start with the assumption that each device in the MANET is running an Anomaly Detection (AD) sensor that allows the device to compute a behavior profile that models its typical behavior. BARTER consists of an initial setup and two main phases: *admission control* and *access control*. Initially, MANET members exchange their behavior profiles in order to build their own individual definition of normal behavior which will later be used during admission and access control. During admission or access control, each MANET member emits an individual decision based on its definition of normal behavior. Individual decisions are aggregated using a threshold cryptographic scheme (t, n) that requires at least t out of the total n MANET members to change the status of the network. Next, we describe each of the phases in detail as well as the interaction with the cryptographic infrastructure that runs underneath the mechanism.

2.1 Initial Setup

The principal goal of the setup is for each MANET member to build their own individual definition of normal behavior, which will be ultimately used during admission control. MANET members are not clients or servers but rather peers *i.e.*, all members are considered equal and can execute client or server activities simultaneously. As a result, MANET members can have both input and output behavior profiles for the same service (port). Throughout, we assume that the behavior profiles are computed from previous interactions of the device or alternatively are provided as built-in profiles from the vendor. We further assume that the profiles of the initial MANET members are clean and provide an accurate representation of the typical behavior in the MANET.

During setup, all the initial members broadcast their output behavior profiles to all the other MANET members. Each member proceeds to calculate the distance between its own input behavior profile and the output behavior profiles received from the other devices. Given the distributed nature of the mechanism, MANET members are only required to provide their output behavior profiles.

This step prevents any member from crafting attacks based on the knowledge of the input profiles of the others. The distance between a member i and each of the other MANET members $j = 1, \dots, n$ is given by $d_{i,j} = d(P_{i,in}, P_{j,out})$, where $P_{i,in}$ is the input behavior profile for i and $P_{j,out}$ is the output behavior profile received from node j . Each pair $(P_{j,out}, d_{i,j})$ computed by member i is then stored as an entry $Q_i[j]$ in its local table Q_i . The entries are sorted by member i according to their distance values such that the closest profiles to $P_{i,in}$ are placed at the top of the table. In general, distances can be interpreted as a measure of confidence between a local device and the rest of the MANET members. Specifically, profiles at shorter distances would be trusted more than their more distant counterparts.

Armed with its sorted table, each MANET member proceeds to calculate its local threshold τ_i that will determine acceptance or rejection of new devices during admission control. The threshold τ_i is defined as the maximum distance between its input profile $P_{i,in}$ and its top $t-1$ most similar/trusted profiles. In this context, t corresponds to the value from the (t,n) threshold cryptographic scheme. Thus, $\tau_i = Q_i[t-1]$, where $Q_i[t-1]$ represents the $t-1$ th entry at the local table of member i .

Simultaneously, the members of the MANET are also responsible for setting up the threshold cryptographic scheme (t,n) . This scheme guarantees that all communications among the n MANET members are encrypted using group keys, which can only be reconstructed by any t members of the MANET. The threshold cryptographic scheme also ensures that all decisions within the MANET must meet the approval of at least t members.

The initial setup of the threshold cryptographic scheme is executed in a distributed fashion without a central authority (CA) following the approach proposed by Narasimha *et al.* [9] using the cryptosystem theory without initial trusted parties from Pedersen [12]. In the approach by Narasimha *et al.*, the group of all MANET members $(M_i, i=1..n)$ uses Shamir's secret sharing [13] to divide a group secret S into n shares. Specifically, the secret is represented as a polynomial $f(z) = f_1(z) + \dots + f_n(z)$, where each $f_i(z)$ is generated by each individual MANET member M_i . Each MANET member M_i computes its share as follows. First, each M_i chooses a random polynomial $f_i(z) \in Z_q$ (where q is a prime number) of degree $t-1$ such that $f_i(0) = S_i$. Next, each M_i computes M_j 's share as $s_i^j = f_i(j)$ (for $j=1..n$) and securely transmits these values to each j through a secure channel. Finally, M_j computes its share s_j of the secret S (partial signature) by summing all the shares received as $s_j = \sum_{i=1}^n (s_i^j)$ and computes its Group Membership Certificate (GMC_i).

Under this scheme, any group of t members among the total n will be able to jointly recover the secret S via Lagrange interpolation. Subsequently, this threshold cryptographic scheme will play a principal role during the admission and access control discussed in the following sections. It is important to note that the merger between BARTER and the threshold cryptographic layer creates a robust mechanism that guarantees distributed admission and access control decisions as well as secure communications among the MANET members.

Cross-Validation The initial cross-validation seeks to find the ratio t/n that yields the best results for the admission and access control mechanism. In particular, the performance of BARTER for each ratio t/n is measured in terms of false rejection (FR) *i.e.*, number of normal profiles wrongly rejected from entering the MANET, true rejection (TR) *i.e.*, number of anomalous profiles detected as such, cryptographic costs (CC) and possibility of Distributed Denial of Service (DDoS) attacks. The values of the ratios t/n are ranked according to their performance $r = (1 - FR) + TR + (1 - CC) + DDoS$ and the highest ranked value is selected. Here, the cryptographic costs (CC) quantify the total time involved during key (re)generation by the MANET members. For practical purposes, the value of CC is normalized between 0 and 1. On the other hand, $DDoS$ evaluates the robustness against MANET members lying about their decisions in order to manipulate the admission and access control. At the end of the setup and cross-validation, each device will have a sorted local table Q_i , a local threshold τ_i as well as the best t/n ratio for the MANET. The actual computation of the parameters used for the ranking is discussed in more detail in Section 3. Experimental results are presented in Section 4 and Section 5.

2.2 Admission Control

Whenever a new device attempts to enter the MANET, it needs to broadcast its own local output behavior profile to the current members. Initially, the members will check whether the new device is blacklisted. If it passes this check, the members proceed to compute the distance between their own input profile and the output profile of the newcomer. If the distance is within its own local threshold of normalcy, the member emits a favorable vote $v_i = 1$. The final MANET vote v can be expressed as:

$$v = \frac{1}{n} \sum_{i=0..n} v_i$$

$$v_i = 0 \text{ if } d(P_{i,in}, P_{new,out}) > \tau_i$$

$$v_i = 1 \text{ if } d(P_{i,in}, P_{new,out}) \leq \tau_i$$

where n is the number of members in the MANET, τ_i is the threshold of member i , $P_{i,in}$ is the input behavior profile of member i and $P_{new,out}$ is the output behavior profile of the newcomer. If t or more members of the MANET emit a favorable vote $v_i = 1$, the newcomer is admitted. Otherwise, the newcomer is rejected and added to a grey list that keeps track of the number of admission attempts by the device. If a device exceeds a fixed number of attempts, it will be added to a blacklist. In order to keep the latest updates, grey and blacklists are exchanged among MANET members.

Upon acceptance of the newcomer, all the members of the MANET submit their output behavior profiles to the new member. The new member stores the profiles together with the distance measures between its own input profile and the output profiles of the remaining members of the MANET in its local table. Then,

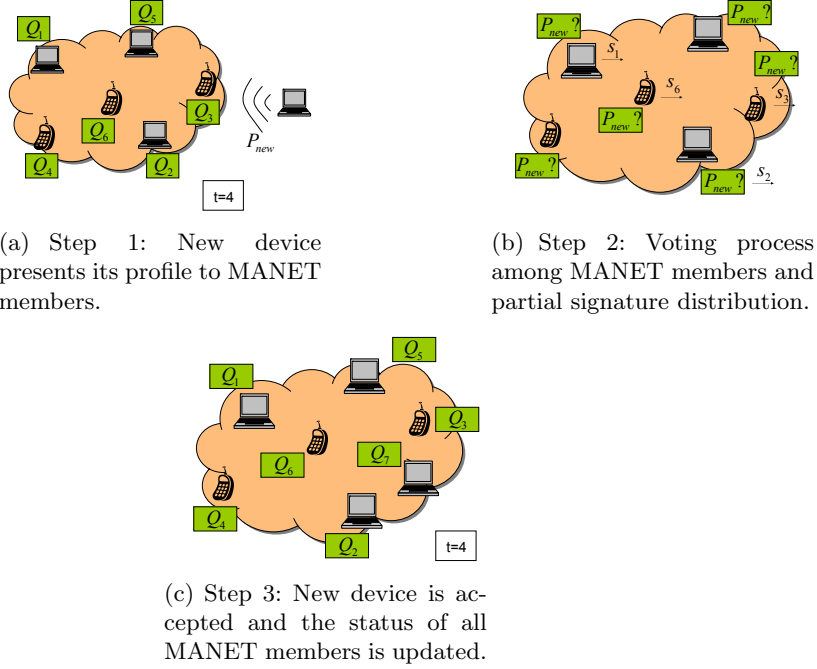


Fig. 1. Admission Control of a Newcomer new to the MANET.

it proceeds to sort the values in its local table according to their distance. The maximum distance value among its top $t-1$ profiles determines its local threshold $\tau_{P_{new}}$. The original members of the MANET store the output behavior profile of the newcomer in their local tables and update their distance computations as well as thresholds accordingly.

Whenever a new device enters or leaves the MANET (n increases or decreases), the ratio of t/n will also change. As a result, the mechanism must make the proper adjustment to restore the ratio to its original value that yielded the best performance for the admission control. In order to avoid recalculating t every time that the value of n changes, we set an update window w such that the value of t is changed only when the ratio exceeds the range $t/n \pm w$. If we consider t_0 to be the initial value of t and n_0 the initial number of MANET members, t would be updated as $t = \lceil (t_0/n_0) * n \rceil$, where n is the final size of the MANET. Throughout, we assume the members can *easily* calculate or approximate the total size of the MANET (n).

Every time a newcomer attempts to be admitted into the MANET, BARTER combines with the threshold cryptographic admission control by Narasimha *et al.* [9] as follows:

1. The newcomer M_{new} broadcasts its public key certificate PKC_{new} and its behavior profile $P_{new,out}$ to the MANET members.

2. Members that deem the behavior profile of the newcomer normal ($d(P_{i,in}, P_{new,out}) \leq \tau_i$) reply with their Group Membership Certificates (GMC_i).
3. M_{new} forms a list of signers SL_{new} and sends it back to each of the members M_j that replied initially.
4. Each M_j computes its partial signature s_j and submits it to M_{new} .
5. M_{new} computes the complete signature s by summing t partial signatures s_j and obtains its own GMC_{new} as well as its partial share s_{new} . In addition, it updates its local table with the behavior profiles of the MANET members.

At the end of the process, if t or more members in the MANET agree on the normal nature of the profile, the newcomer can compute its own GMC and start communications with the MANET. Otherwise, the newcomer will not be able to participate or even eavesdrop because the communications are encrypted. Figure 1 depicts an example of the admission control in a MANET with six initial members and a value of $t=4$.

2.3 Access Control

During access control, communications among the MANET members are continuously screened to ensure that users do not deviate from their declared behavior profiles. In practical terms, each MANET member continuously checks the incoming traffic from any device against its local input behavior profile as well as against the output behavior profile of the sender that was originally saved during the admission of the latter into the MANET. If a device considers some traffic to be anomalous, it requires at least t members of the MANET in order to act against the sender. Thus, the receiver of the anomalous traffic submits the anomaly to its *top $t-1$ most similar* members drawn from its local table. If the other $t-1$ members agree on the anomalous nature of the traffic, the sender is expelled immediately from the MANET. This process relies on the assumption that there exists an scheme that prevents data tampering within the MANET and prevents users from falsifying alerts or replay attacks.

From a threshold cryptographic point of view, if a MANET member detects an anomaly, it adds the anomalous member to its local Certificate Revocation List (CRL) and proceeds to broadcast its own CRL to all the MANET members. In order to ensure that the anomalous member no longer has a vote in the distributed admission and access control, each member will generate new partial signatures that will be submitted to each of the MANET members outside the CRL via point-to-point communications (individual cryptographic channels). This proactive key sharing [7] combines the approaches introduced by Ostrovsky and Yung [10] and by Luo and Lu [8] as follows:

1. Each member M_i defines a polynomial $f_i(z) = f_1z^1 + f_2z^2 + \dots + f_{t-1}z^{(t-1)}$ with $f_i(0) = S_i$, where $f_1..f_{t-1} \in Z_q$ are randomly selected and q is a prime number.
2. M_i secretly sends $s_i^j = f_i(j)(mod\ q)$ to the MANET members M_j outside the CRL. The members are assumed to have established point-to-point encrypted channels.

3. M_j would reply if and only if it has received t revocation lists (CRL) from different MANET members.
4. M_i decrypts the s_j^i received from the other MANET members and computes its new share s_i .

3 Attacks and Cryptographic Costs

Due to the fully distributed nature of the BARTER mechanism, the main source of attacks derives from MANET members lying about their admission control decisions either to admit members with malicious profiles or else to prevent members with normal profiles from joining the MANET. Small values of the t/n ratio would permit attackers to get a hold of the admission control by compromising only a few MANET members. In contrast, larger t/n ratios could be attacked by compromising a few nodes that would prevent the other members of the MANET from reaching a decision. We quantify the robustness against DDoS attacks via a numerical factor $DDoS$ in the ranking index formula such that $r = (1 - FR) + TR + (1 - CC) + DDoS$. Our assumption is that a $t/n = 0.5$ represents the optimal value to minimize the risk of potential DDoS attacks. As a result, we set $DDoS = 0.5$ for $t/n = 0.5$. We set $DDoS = t/n$ for smaller ratios ($t/n < 0.5$) and assume $DDoS = 1 - (t/n)$ for larger ratios ($t/n > 0.5$). While some t/n ratios may yield better FR or TR rates, the value of $DDoS$ serves as a counterweight to estimate the robustness against DDoS attacks for a certain configuration.

The cryptographic costs (CC) involved in this process are quantified in terms of the total time spent during key regeneration and enter in the evaluation of the ranking index such that $r = (1 - FR) + TR + (1 - CC) + DDoS$. As written, the ranking index penalizes high values of CC and favors more economic key regenerations. During the initial cryptographic setup, each MANET member exchanges shares with all the other members in order to compute its own GMC as well as its partial signature. If we assume that all MANET members exchange their shares in parallel, we can approximate the initial setup cost as $CC = K \times (n_0 - 1)$, where K represents the cost of a single exchange and n_0 is the number of initial members in the MANET. Every time a device enters or leaves the MANET, the mechanism checks that the ratio t/n is within the window w . If the value of t needs to be adjusted, the cost incurred in the key regeneration can be approximated by,

$$CC = K \times \sum_{n_0}^{n_{final}} (update \times n) - 1$$

$$update = \begin{cases} 1 & \text{if } t/n < (t_0/n_0 - w) \quad \text{or if } t/n > (t_0/n_0 + w) \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

where n_{final} is the final number of MANET members, n is the current number of MANET members, and $update$ is a boolean variable that determines whether

an update in the value of t is required. Equation 1 describes the boolean variable that is only *true* whenever the ratio t/n falls below the lower bound of the range ($t/n - w$) or exceeds the upper bound ($t/n + w$).

4 Evaluation of BARTER with Content Profiles

In this section, we begin with a description of the AD sensor responsible for the computation of the content behavior profiles and provide experimental results of the BARTER mechanism using this type of profile.

4.1 Semi-supervised Content AD Sensor

We have implemented a content-based AD sensor that represents an adaptation of Shanner’s ideas [14]. Shanner proposes an algorithm that incorporates only the most heavily weighted grams to the behavior profile. These grams are the ones that best discriminate between two or more classes of data. Although Shanner is more expensive than other AD sensors, we chose it because it rapidly captures the significant information of the traffic being exchanged.

In our sensor, we consider two classes of data (content): good samples (*goodS*) and bad samples (*badS*). The content of the traffic exchanged is captured as 3-grams. This choice is less computationally expensive than higher n-grams and appropriately captures the specifics of email traffic (as it will be shown in Section 4.3). The weight (frequency) of each 3-gram observed during training is calculated using Shanner’s Formula (see Equation 2), where the frequency W of each 3-gram i ($W(i)=F(i) \times U(i) \times A(i)$) is expressed as,

$$W(i) = \log\left(\frac{x_i}{N_g}\right) \times \left(\frac{1}{\log N_g}\right) \sum_{j=1}^{N_g} (p_{ij} \log\left(\frac{1}{p_{ij}}\right)) \times \left(1 - \left(\frac{1}{\log L}\right) \sum_j^{goodS, badS} (p_{ij} \log\left(\frac{1}{p_{ij}}\right))\right) \quad (2)$$

where $F(i)$ measures the frequency of occurrence of each distinct 3-gram i over all the good samples N_g ; $U(i)$ measures how uniformly distributed each unique 3-gram i is spread among the set of good samples N_g (p_{ij} represents the probability of seeing 3-gram i in good sample j); and $A(i)$ measures how uniformly distributed each unique 3-gram i is spread across all good and bad samples types ($L=2$). Once all the weights have been calculated, a top percentage of the 3-grams are selected to represent the content profile of the device. We assert that this is a semi-supervised learning technique since the devices store an initial collection of bad 3-grams drawn from known malware samples.

4.2 Behavior Profile Privacy

Due to the fact that the output behavior profiles are exchanged among devices, it may be the case that certain users do not feel comfortable sharing the content they exchange. In order to deal with this possibility, the BARTER mechanism

hashes the content behavior profile into Bloom Filters (BF) [1]. Input behavior profiles, although not exchanged, are also converted into BFs so that the comparison with output profiles is fast and straight forward. The way the traffic is mapped to a BF depends on the AD sensor used. The only requirement is that all devices use the same sensor with the same type of mapping. Behavior profiles in BARTER are then easily comparable, since boolean operators allow us to discern similarities or differences between profiles (BFs). For this type of content profiles, the distance $d(P_{i,in}, P_{j,out})$ between two profiles is computed using an exclusive OR (XOR) operator that quantifies the amount of entries that differ between the two profiles: $d(P_{i,in}, P_{j,out}) = |P_{i,in} \oplus P_{j,out}|$, where \oplus represents the XOR operator and $||$ denotes the total number of entries with different values.

4.3 Evaluation Experiments

Having extensively described the foundations of the BARTER mechanism, we proceeded to test the admission control of the mechanism with real content behavior profiles. We focus on the admission control because our main aim is to prove the functionality of the mechanism together with the threshold cryptographic layer. For that purpose, we used the publicly available ENRON dataset [2] that contains 125,218 emails from 140 ENRON employees. The reason for choosing email as a testing dataset is justified based on its use as a primary application on handhelds. Moreover, email constitutes a good approximation of other popular text messaging applications that could be used in MANETs. For each of the 140 users, we computed input and output behavior profiles using Shanner’s algorithm [14]. In particular, we first calculated the frequency for each 3-gram in the body of each user’s emails and selected the top 5000 most heavily weighted grams. Our choice of 5000 worked well for our experiments, however, other values might be more appropriate for different datasets. The bad samples used to execute Shanner’s algorithm were drawn from the signature content of the Snort rules (a total of 58) [15] and from 600 virus samples of *vxheavens* [17]. Finally, the top 5000 3-grams were hashed into Bloom filters in order to provide privacy to the profiles.

For experimental purposes, we refer to the set of behavior profiles modeled with the content of emails from the ENRON dataset as *pool of normal users* (140 users). The pool of normal users was considered to be clean, composed of normal behavior profiles and *ground truth*. On the other hand, we refer to *pool of bad users* as a set of 60 profiles that represent anomalous behavior (content) that should be rejected from entering the MANET. In order to compute the bad profiles, we used content (3-grams) from 12000 executable files and code files (C or Java). These type of files were chosen because their content is dramatically different from email. Each bad profile was computed with 200 executable/code files with an average of 39 3-grams per file. Once again, behavior profiles were computed using Shanner’s algorithm by selecting the top 5000 3-grams and hashing them into Bloom filters.

We defined a group of 80 behavior profiles randomly selected from the pool of normal users as our training set (initial MANET members). The remaining 60

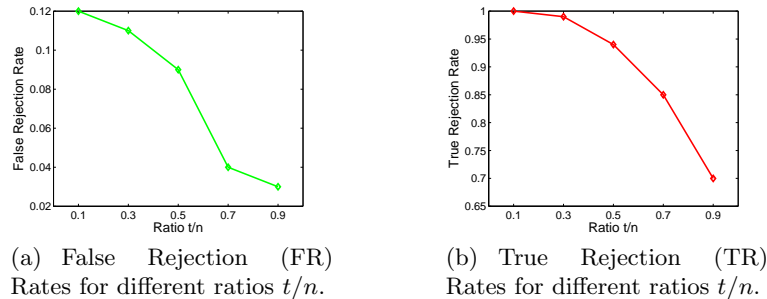


Fig. 2. FR and TR for different ratio values t/n .

profiles were divided into a cross-validation set (30 profiles) and a testing set (30 profiles). The cross-validation set was used to calculate the best ratio t/n for our dataset determined as the highest ranking index $r = (1 - FR) + TR + (1 - CC) + DDoS$. One by one, each randomly selected profile in the cross-validation set was presented to the MANET members as a newcomer attempting to be admitted into the MANET. From these attempts, we measured the false rejection rate (FR) as the percentage of normal profiles wrongly rejected as anomalous. Next, 30 randomly selected profiles from the pool of bad users were also presented as new devices attempting to gain access into the MANET. The latter experiment allowed us to measure the true rejection rate (TR) by determining the percentage of profiles correctly rejected as anomalous by the MANET members. Finally, we proceeded to determine the cryptographic costs (CC) incurred during the generation of new signatures for the MANET members as well as the robustness against DDoS attacks ($DDoS$).

The experiments were repeated 60 times for each value of t/n to cover all different evolutions resulting from the random selection of the initial profiles. The results presented here constitute an average over all runs. Five different values of t/n were considered, namely 0.1, 0.3, 0.5, 0.7, and 0.9. These values represent MANETs whereby 10%, 30%, 50%, 70% or 90% of the total members respectively are needed in order to emit an admission control decision. For each ratio, the initial value of t was calculated as $t/n \times 80$ with a window $w = 0.02$. Here, 80 represents the number of total profiles in the training set (initial MANET members). Alternative values of w would produce different numerical results but would follow the same trends observed in our experiments.

Figures 2(a) and 2(b) show the FR and TR rates for different ratio values. As can be seen, smaller ratio values produced larger FR and TR rates. This is most likely related to the fact that smaller ratios reflect a larger number of small sets of profiles. Consequently, the thresholds for the admission control become very restrictive, which results in a larger rejection of normal profiles as well as a larger detection of anomalous profiles. In contrast, larger ratios result in smaller FR and TR rates possibly associated with fewer sets of profiles that define less restrictive thresholds.

We also computed the ranking indices $r = (1 - FR) + TR + (1 - CC) + DDoS$ for different ratio values t/n for the ENRON dataset, and we found that the highest ranked index corresponded to a ratio $t/n = 0.1$. Such result probably captures the wide variety of content contained in the email exchanges among users. In other words, small sets of normalcy provide a better characterization of the behaviors shared by the users. Armed with the highest ranked index $t/n = 0.1$ obtained from cross-validation, we proceeded to simulate the admission control with randomly selected profiles from the testing set acting as newcomers to the MANET. In order to compute TR and FR, we used the remaining 30 normal and bad profiles drawn from the pool of normal and bad users respectively. For the performance of BARTER, we obtained a false rejection FR=13%, a true rejection TR=100% and cryptographic costs CC=179 \times K.

5 Evaluation of BARTER with Volumetric Profiles

Rather than content, volumetric profiles capture the typical characteristics of the communications such as number of emails exchanged, number of different people contacted (clique), and frequency of usage. In this Section, we describe the volumetric AD sensor used to compute the profiles. This is followed by an actual evaluation of the BARTER mechanism using volumetric behavior profiles computed from the ENRON dataset.

5.1 Histogram-based Volumetric AD Sensor

In order to compute volumetric input and output behavior profiles, we used the EMT tool (Email Mining Toolkit) [16]. The behavior profile of each user is represented as a daily histogram that reflects the behavior of a user exchanging emails. In order to be able to compute the initial behavior profiles of the MANET members, we presume that members have an archive of stored emails from previous interactions in other environments. Alternatively, one can always provide the members with an initial set of training samples chosen according to the type of user.

EMT computes two types of daily histograms: *hourly histograms* and *grouped histograms*. Hourly histograms divide the day in 24 bins where each bin represents the average number of emails (sent or received) per hour. Grouped histograms, on the other hand, divide the day in 4 bins of 6 hours each, where each bin is the average number of emails sent or received during a 6-hour period. Hereafter, we will refer to the number of bins in which the day is divided as bin granularity (bg). In particular, we shall use $bg = 24$ for hourly modeling and $bg = 4$ for grouped modeling. Each profile $P_{i,d}$ is a vector with bg entries, where d represents the direction of the traffic *i.e.*, either input ($P_{i,in}$) or output emails ($P_{i,out}$). Each histogram entry represents a bin b_j that contains the average value a and standard deviation σ for the number of emails sent or received by user i during a time frame j . The average and standard deviation values for each time frame j are averaged throughout the duration of the training period. Hence,

$P_{i,d} = \{b_1, \dots, b_{b_g}\}$ and $b_j = \{(a, \sigma)\}$ where $j \in [1..24]$ for hourly histograms or $j \in [1..4]$ for grouped histograms.

The *grouped histograms* are intended to save bandwidth usage by exchanging smaller behavior profiles among MANET members. Nonetheless, it is only through cross-validation tests that an appropriate bin granularity that both minimizes the behavior profile size while maximizing BARTER performance can be properly selected.

5.2 Evaluation Experiments

Our evaluation of the BARTER mechanism for volumetric profiles is similar to the one presented for content profiles. Again, we used the publicly available ENRON dataset to compute the volumetric behavior profiles of 140 users. For each user, we computed its input and output volumetric profiles in two formats: hourly and grouped histograms. Behavior profiles were computed by calculating the average number of emails sent or received by a user throughout the duration of the training period. For experimental purposes, the set of 140 volumetric profiles modeled with emails from the ENRON dataset is referred to as *pool of normal users* and is considered *ground truth*. In order to compute a *pool of bad users*, we produced volumetric behavior profiles one, two, and three standard deviations away from the top $t-1$ entries in the local table of each MANET member. Our assumption is that the $t-1$ top entries represent the most similar counterpart to a particular profile. As a result, behavior profiles separated by one or more standard deviations from this set constitute potential anomalous profiles. We repeated this process for all the members of the MANET and obtained a final pool of bad users for each ratio t/n .

As in the experiments with content profiles, we simulated an environment where a number of profiles attempt to gain admission into an already formed MANET. The pool of normal users (140 profiles) was divided into three sets: the training set (80 randomly selected profiles), the cross-validation set (30 randomly selected profiles), and the testing set (the remaining 30 profiles). Armed with these sets, we measured the FR rate of the BARTER mechanism with volumetric profiles. Next, we created additional cross-validation and testing sets with 30 profiles each randomly selected from the pool of bad users. The latter sets were used to compute the TR rate of the mechanism.

The purpose of the cross-validation experiments is to determine the combination of t/n ratio and type of histogram that yields the highest ranking index r . We experimented with five different values for the ratio t/n : 0.1, 0.3, 0.5, 0.7, 0.9 and two types of histograms: hourly and grouped. Simulations were repeated 60 times to account for the random draw of the initial profiles, and the results were averaged among the 60 simulations. For each combination of parameters, we computed the ranking index r and selected the highest ranked.

Figures 3(a) and 3(b) depict our TR and FR results for a wide variety of ratio values as well as two types of histograms (hourly and grouped). As can be seen, grouped histograms outperform hourly histograms in terms of FR and TR rates. Our interpretation is that hourly histograms likely produce a too fine grained

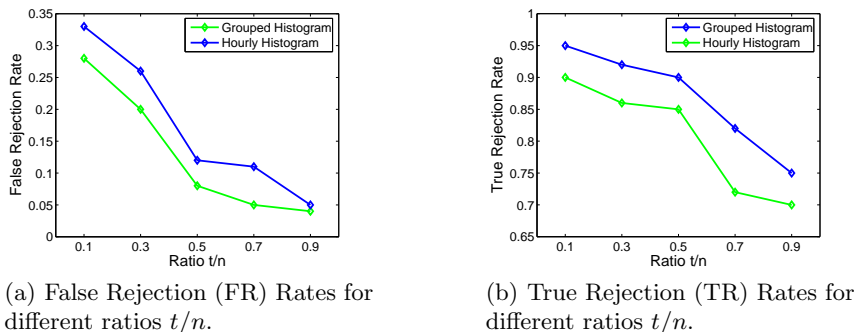


Fig. 3. FR and TR for different ratio values t/n .

modeling for our dataset. In contrast, grouped histograms can identify behaviors more effectively thus improving the performance of BARTER. In general, higher t/n ratios translate into smaller FR and TR rates.

We note that the highest ranked index occurs for grouped histograms and a ratio $t/n = 0.5$. Such ratio indicates that the best admission results take place when 50% of the MANET members are needed to make a decision. We can also interpret this result as an indication of few distinct behaviors within the ENRON dataset. Taking the highest ranked index parameters ($t/n = 0.5$ and grouped histograms), each randomly selected profile from the testing set was presented to the MANET members as a newcomer attempting to be admitted into the MANET. From these admission control experiments, we measured FR=8%, TR=90% and cryptographic costs $CC=1571 \times K$. These results demonstrate the feasibility of BARTER using volumetric as well as content behavior profiles.

6 Related Work

There is a body of work about the use of threshold cryptography for admission control in ad-hoc networks. Narasimha et al. [9] and Ostrovsky et al. [10] studied possible adaptations of existing threshold cryptographic schemes to MANETs. However, none of the previous works have discussed the implementation of the decision process during admission control. BARTER enhances threshold cryptographic approaches by automatizing the individual admission decision at each device. AD sensors have been widely used to implement access control in MANETs. The main idea is that profiles computed from audit data can be used as a representation of the normal behavior. As a result, any behavior that deviates from the profile is considered anomalous [18]. However, the current literature does not offer a satisfactory explanation on the interaction of ADs with secure cryptographic platforms. BARTER provides an access control that uses AD sensors at an application level rather than at the routing level. Additionally, our work describes the interaction between the AD sensors and the cryptographic layer.

7 Conclusions and Future Work

We have presented BARTER, a mechanism that automatically creates admission and access control policies for MANETs. Individual decisions regarding admission and access control are issued based on a local definition of normal behavior computed from the knowledge of the behavior profiles from other members. A threshold cryptographic layer (t,n) that runs underneath the mechanism aggregates the individual decisions by requiring at least t devices to participate in the decision. We have discussed experimental results using both content and volumetric behavior profiles computed from the ENRON dataset. Our results show that the mechanism can successfully perform under both types of behavior profiles with FR rates ranging from 9% to 12% and TR rates between 95% and 100%. Future work will evaluate how to best determine the most defining behavioral characteristics of a host using techniques such as bagging or boosting.

References

1. B. H. Bloom. “Space/Time tradeoffs in hash coding with allowable errors”, Communications of the ACM, Volume 13 (7), 1970.
2. ENRON Dataset. “www.cs.cmu.edu/enron”, 2004.
3. V. Frias-Martinez and S. J. Stolfo and A. D. Keromytis. “Behavior-Based Network Access Control: A Proof-of-Concept”, ISC, 2008.
4. V. Frias-Martinez and S. J. Stolfo and A. D. Keromytis. “Behavior-Profile Clustering for False Alert Reduction in Anomaly Detection Sensors”, ACSAC, 2008.
5. V. Frias-Martinez, et al. “A Network Access Control Mechanism Based on Behavior Profiles”, ACSAC, 2009.
6. J. Hastad et al. “Funkspiel Schemes: An Alternative to Conventional Tamper Resistance”, In Proc. of the 7th ACM Conf. on Computer Communn. Security, 2000.
7. A. Herzberg et al. “Proactive Secret Sharing Or: How to Cope with the Perpetual Leakage”, In Proceedings of Advances in Cryptology (CRYPTO), 1995.
8. H. Luo and S. Lu. “Ubiquitous and Robust Authentication Services for Ad Hoc Wireless Networks”, Technical Report, UCLA, 2000.
9. M. Narasimha et al. “On the utility of Distributed Cryptography in P2P and MANETs: the case of Membership Control”, In Proc. of the 11th ICNP, 2003.
10. R. Ostrovsky and M. Yung. “How To Withstand Mobile Virus Attacks”, In Proc. of the 10th ACM Symp. on the Principles of Distributed Computing, 1991.
11. P. Papadimitratos and Z.J. Haas. “Secure Data Transmission in Mobile Ad Hoc Networks”, In Proceedings of the ACM Workshop on Wireless Security, WiSe, 2003.
12. T. P. Pedersen. “A Threshold Cryptosystem without a Trusted Party”, In Proceedings of Eurocrypt, LNCS 547, 1991.
13. A. Shamir. “How to share a secret”, Communications ACM, Vol. 22 (11), 1979.
14. R. A. Shaner. “US Patent No. 5,991,714”, November, 1999.
15. Snort Rulesets. “<http://www.snort.org/pub-in/downloads.cgi>”
16. S. J. Stolfo et al. “Behavior-based Modeling and its Application to Email Analysis”, ACM Transactions on Internet Technology (TOIT), Volume 6(2), 2006.
17. VXHeavens. “vx.netlux.org”
18. Y. Zhang and W. Lee and Y. Huang. “Intrusion Detection Techniques for Mobile Wireless Networks”, Mobile Networks and Applications, Volume 9 (5), 2003.