Public Review for

Proactive Attacker Localization in Wireless LAN

Chuan Han, Siyu Zhan, and Yaling Yang

The paper presents a technique to localize WLAN intruders. Traditionally, this problem has been solved by assuming that multiple observers (usually Access Points) can simultaneously observe the intruder's transmissions, and use time delays, angle of arrival, or signal strength information to localize the intruder. The authors of this work consider a more capable intruder, who can, for example, beamform its transmissions to be heard by one or few Access Points. The fewer the number of such observers, the less accurate can be the localization process. The novelty of this work is a proactive technique that forces that intruder to expose its transmissions to more Access Point observers. Essentially, the authors propose a coordinated system in which the current Access Point serving an intruder observes the latter's transmission characteristics for a short while and then dissociates it. At that time, a different Access Point, located elsewhere but part of the same WLAN system, accepts this intruder, provides access to its traffic for a little while, and then dissociates it again. As the process repeats, the intruder transmission characteristics get exposed to many observers allowing localization. A core part of the paper is focused on determining the sequence of Access Points that should serve the intruder for faster and more accurate localization.

The core idea is quite interesting and was appreciated by all reviewers. Clearly there are many interesting next questions that need careful exploration. The system works on a somewhat long timescale over which the intruder is assumed to be fairly stationary. So can this be adapted to mobile intruders? What happens when the intruder is aware of the WLAN's strategy of localization and tries to throw off the localization process in some way? The work is simulation-based, and clearly an implementation of the system will throw up some new challenges such a system will have to address. The localization system depends on a somewhat high density of observers in the environment. What is the trade-off between observers and accuracy of localization in this manner?

Overall, a fairly interesting piece of work, that clearly warrants further exploration.

Public review written by
Suman Banerjee
University of Wisconsin

acm sigcomm

Proactive Attacker Localization in Wireless LAN

Chuan Han, Siyu Zhan, Yaling Yang
ECE Department, Virginia Tech
Email: hanc@vt.edu, mvpzhansy@hotmail.com and yyang8@vt.edu

ABSTRACT

This paper addresses the open problem of locating an attacker that intentionally hides or falsifies its position using advanced radio technologies. A novel attacker localization mechanism, called Access Point Coordinated Localization (APCL), is proposed for IEEE 802.11 networks. APCL actively forces the attacker to reveal its position information by combining access point (AP) coordination with the traditional range-free localization. The optimal AP coordination process is calculated by modeling it as a finite horizon discrete Markov decision process, which is efficiently solved by an approximation algorithm. The performance advantages are verified through extensive simulations.

Categories and Subject Descriptors

C.2.1 [**Computer Systems Organization**]: COMPUTER-COMMUNICATION NETWORKS—*Network Architecture and Design*

General Terms

Security

Keywords

Secure localization, wireless LAN

1. INTRODUCTION

With the pervasive deployment of the IEEE 802.11 wireless local area networks (WLAN), it is quite easy for an attacker to launch network attacks from a wireless terminal to remote critical network infrastructures, such as national, financial, energy, transportation and military network systems. Figure 1 shows a scenario where an attacker uses a laptop to attack remote critical network infrastructures. Unlike in wired networks, where the attacker has to be physically close to an Ethernet port to connect to the network, an attacker in a wireless network is able to launch an attack without a fixed position. In addition, an attacker can also easily fake its MAC and IP addresses, invalidating any attempt to identify the attacker's identity through these addresses. The highly mobile, anonymous and stealthy nature of wireless communications makes the countermeasure of wireless attack much more challenging than in wired networks. Therefore, to guarantee the security of these critical infrastructures and completely eliminate the threat of a remote wireless attacker, we must design a defense mechanism that can effectively locate the attacker in a wireless network and make it liable for its offense.

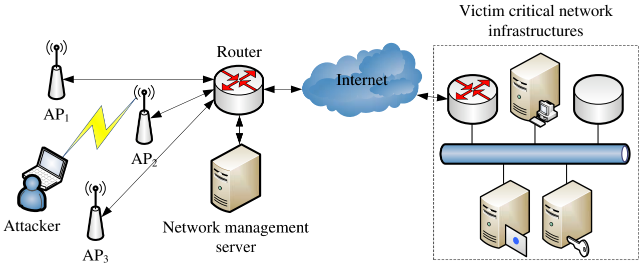
Victim critical network infrastructures
Router
Internet
AP$_1$
AP$_2$
AP$_3$
Attacker
Network management server

Figure 1: Wireless attacker's threat to remote critical infrastructures

The design of such a mechanism involves two stages. First, an attacker is traced back to its *home access point (home AP)*, which is the AP the attacker currently connects to. Existing attack traceback [1, 2], traffic monitoring [3] and device identification [4] techniques can successfully complete this stage. Then, wireless localization schemes locate the attacker so that law enforcement agencies can catch and penalize the attacker for its misdemeanor. This second stage, unfortunately, is far from a well solved problem. Existing localization methods [5–7], including both range-free and range-based schemes, rely on passive observation of the attacker's signal features, such as its connectivity with neighboring nodes, received signal strength (RSS), time of arrival (TOA), angle of arrival (AOA), and etc. An intelligent attacker equipped with advanced radio technologies, like directional antennas and software defined radios (SDRs), can change its beam direction and radio parameters to distort these signal features so that it can falsify or hide its position with great ease and anonymity. In the presence of such an intelligent attacker, the best that existing secure localization schemes [7, 8] can do is only to verify that an attacker's position claim is false. None of the existing efforts can reveal the true position of such an attacker.

To address this critical challenge, in this paper, we propose a novel *Access Point Coordinated Localization (APCL)* scheme, which is the first method that can locate an attacker equipped with directional antennas and SDRs. APCL coordinates APs around the attacker to force the attacker to reveal undistorted signal features unintentionally. Traditional localization techniques can then be used to capture these features and locate the attacker. To ensure optimal localization of the attacker, the AP coordination process is

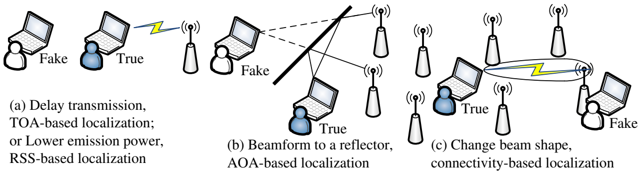
Fake
True
(a) Delay transmission, TOA-based localization; or Lower emission power, RSS-based localization
Fake
True
(b) Beamform to a reflector, AOA-based localization
True
Fake
(c) Change beam shape, connectivity-based localization

Figure 2: Fake attacker position scenarios

modeled as a finite horizon discrete Markov decision process (MDP), and an approximation algorithm is proposed to efficiently find the quasi-optimal coordination process. APCL only imposes a negligible query load on APs, and does not require any realtime computation and special hardware.

The rest of this paper is organized as follows. The threat model is shown in Section 2. APCL is described in Section 3, and is modeled as a finite horizon discrete MDP in Section 4. An efficient approximation algorithm is proposed in Section 5. The simulation results are shown in Section 6. Finally, Section 7 concludes the whole paper.

2. THREAT MODEL

Using directional antennas and SDRs, an attacker can easily fake a false position by fooling traditional localization schemes. A few such scenarios are listed as follows:

- In traditional TOA-based or RSS-based localization systems, an AP computes its distance to the attacker based on round trip time or the received signal strength from the attacker. If the attacker intentionally delays its transmission or lowers its transmission power to the AP, distance estimation becomes overly large, resulting in a false location estimation as shown in Figure 2 (a).

- In traditional AOA-based localization systems, the arrival angles of the attacker's signal to multiple APs are used to estimate the attacker's position. If the attacker intentionally beamforms its antenna beam to a strong reflector. AOA-based localization schemes locate the attacker to its mirrored image as shown in Figure 2 (b).

- In traditional connectivity-based localization systems, the connectivity information of the attacker with multiple APs is used to estimate the attacker's position. For example, in [9], the centroid of APs which can receive the attacker's signal is used as the position estimation. If the attacker uses a directional antenna, it can intentionally distorts the number and the distribution of APs that can receive its signal, which results in a false position estimation as shown in Figure 2 (c).

As the above examples show, an intelligent attacker equipped with directional antennas and SDRs can easily fool traditional localization schemes that are based on passive signal feature measurements. In the next Section, we present our APCL scheme to locate such an intelligent attacker.

3. APCL MODEL

APCL is based on range-free localization schemes. We select range-free localization schemes because they can be directly implemented over off-the-shelf Wi-Fi devices and require no modification to the existing IEEE 802.11 standard [10]. APCL assumes that the only reliable measurement at an AP is the connectivity between the attacker and the AP. Such a low requirement on the hardware devices ensures the widest applicability of APCL since any Wi-Fi device can provide such measurement. In addition, APCL can locate the attacker even if it is intelligent enough that it can control its beam direction and transmission power to only allow one AP to receive its signal. This is the worst case for the localization purpose. In the remainder of this section, we discuss the design of APCL.

3.1 Basic assumptions

The design of APCL is based on the following assumptions. Multiple APs can be coordinated to locate the attacker. Using existing network attack traceback, traffic monitoring, and wireless device identification techniques, the attacker has been successfully identified and traced down to its home AP before APCL is started to locate the attacker. In addition, we assume that there is only one attacker in the coverage region of its home AP and the attacker is relatively stationary during the localization process. In our future work, we will further study the cases of locating multiple attackers and mobile attackers. Each AP is assumed to have correct knowledge about its position. We further assume that we can get a rough estimation about the maximum possible communication range between the AP and the attacker. This assumption is based on the fact that the duration of ACK timeout in the IEEE 802.11 standard [10] limits the maximum distance between an AP and its client to be around 150 m. This limit cannot be changed by larger transmission power or a more sensitive antenna. Instead, long range communications beyond 150 m have to modify the value of ACK timeout in both the AP and the attacker [11], which is impossible since the attacker cannot modify AP's configurations.

3.2 Mechanism description

Based on the assumptions in Section 3.1, APCL works as follows. The initial estimation of the attacker's position is the coverage region of its home AP, which is the region determined by the maximum communication range between an AP and an attacker. Since the initial estimation region may be too large to locate the attacker, APCL disassociates the attacker from its home AP. This is done by sending a disassociation frame to the attacker to terminate this individual connection. This operation is only targeted at the attacker and does not interfere existing legitimate WLAN users. A specified Reason Code is filled in the Reason Code field of the disassociation frame. Several legitimate reasons can be used to disassociate the attacker [10], e.g., bad link quality, overcrowded wireless access, etc. To continue its attack, the attacker has to reconnect to one of its neighboring APs. Once the reconnection is established and the attacker starts to send its attacking traffic through this new home AP again, this new home AP can be quickly identified through attack traffic traceback, traffic analysis or wireless device identification techniques. After successful identification, APCL narrows the position estimation region down to the intersection of the previously connected and the new home APs' coverage regions. This process continues until there are no neighboring APs of the attacker that can help APCL to narrow the position estimation.

In the above localization process, if there is no control over the possible APs which the attacker can reconnect to after its disassociation, the attacker may pick the AP which cannot contribute to the narrow down process or the AP which contributes little to the process. For instance, the attacker can pick the AP whose coverage region encompasses the current estimation region. In this case, the reconnection of the attacker does not narrow down the attacker. To guarantee effective narrow down of the estimation region in each disassociation step, APCL has to control the possible APs which the attacker may reconnect to. We define that a set of APs are *activated* for the attacker, when they are selected to be available to the attacker for communication. This means that in the passive scanning mode [10], only those activated APs reply Association Response frames to Association Request frames from the attacker, and in the active scanning mode [10], only activated APs reply Probe Response frames to Probe Request frames from the attacker. It is important to note that the selected activation of APs is targeted only to the attacker. For legitimate users, all APs are available for communications and reply to Association Request frames or Probe Request frames from the legitimate users. Therefore, the AP activation process does not affect legitimate users' operations. As disassociation and association operations are common phenomena in normal WLAN environments, it is difficult for the attacker to detect the existence of APCL.

To illustrate APCL, a simple example is shown in Figure 3. Suppose there are five APs, AP_0, AP_1, AP_2, AP_3 and AP_4. In Figure 3, let the black dot denote the attacker, the empty dots be the APs, the circles be the coverage regions of the APs, and the activated APs' coverage regions are marked as solid-line circles. Assume, firstly, the attacker initially connects to AP_0, and the attacker is located in the coverage region of AP_0 as shown in the shadowed region in Figure 3 (a). Next, APCL activates AP_1 and AP_2, and disassociates the attacker from AP_0 as shown in Figure 3 (b). To continue its attack, the attacker reconnects to AP_1, and then APCL narrows down the attacker to the intersection of coverage regions of AP_0 and AP_1 as shown in Figure 3 (c). By the similar process, APCL activates AP_2, and disassociates the attacker from AP_1 as shown in Figure 3 (d). The attacker connects to AP_2, and APCL narrows down the attacker to the intersection of coverage regions of AP_0, AP_1 and AP_2 as shown in Figure 3 (e).

In the APCL localization process, it is possible the attacker cannot find an activated AP within its communication range. In this case, the attacker is termed as *alerted* and the process terminates. The attacker's position is estimated within the previous estimated region but outside of the union of the activated AP coverage regions. As shown in Figure 3 (f), before disassociating the attacker from AP_0, suppose APCL activates AP_3. Then, the attacker cannot find any activated APs within its communication range, and the estimation region of the attacker's position is the shadowed region in Figure 3 (f).

APCL is initiated and controlled by a network management server as shown in Figure 1, which is in command of the whole process of tracing back and locating the attacker. The flow chart of APCL is shown in Figure 4. The effectiveness of APCL, to some extent, relies on the existence of multiple APs around the attacker. Given the wide deployment of WLAN, this is very likely to be the case. Figure 5 shows a sample measurement of the number of APs at 13 randomly picked locations in Virginia Tech campus buildings. In the measurement, a normal omnidirectional antenna is used. The fact that more than 12 APs are present at all the 13 locations validates the assumptions of APCL.

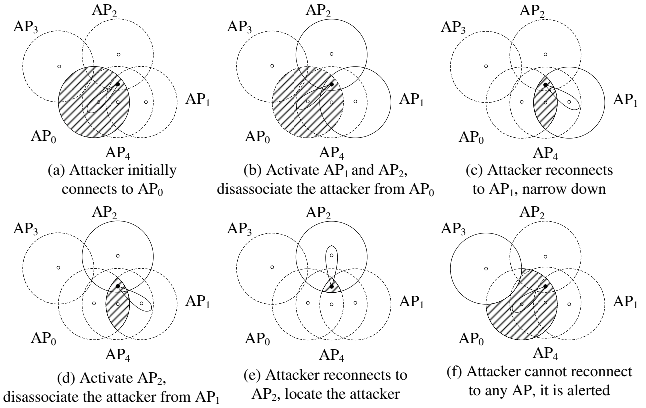

(a) Attacker initially connects to AP_0

(b) Activate AP_1 and AP_2, disassociate the attacker from AP_0

(c) Attacker reconnects to AP_1, narrow down

(d) Activate AP_2, disassociate the attacker from AP_1

(e) Attacker reconnects to AP_2, locate the attacker

(f) Attacker cannot reconnect to any AP, it is alerted

Figure 3: APCL example

4. OPTIMAL AP ACTIVATION SEQUENCE

In each activation and disassociation step, APCL consumes certain network resources. To minimize the network resource consumption, it is desirable to locate the attacker within the least number of steps. Moreover, to minimize the final localization error, APCL also needs to find the optimal AP coordination process. Therefore, there arises an interesting problem of identifying the optimal AP activation sequence.

4.1 MDP model for the AP activation process

The optimal AP activation sequence problem can be modeled as a finite horizon discrete Markov decision process (MDP) based on the following two reasons. Firstly, APCL only has a partial control over the whole process. While it is possible to control which APs to activate in one step, it is impossible to control which activated AP the attacker actually reconnects to. Secondly, the activation process has the Markov property, i.e., given any current state, the transition to the next state is only dependent on the current estimation region and is independent of the previous localization process. Hence, MDP is the appropriate model for computing the optimal AP activation sequence.

4.1.1 Definition

The MDP model for the APCL activation process is defined as a tuple

$$\left[S, A, A(s), P_a(s, s'), r_a(s)\right], \tag{1}$$

where S is the state space, A is the action space, $A(s)$ is the action space for state $s \in S$, $P_a(s, s')$ is the transition probability of a given action $a \in A(s)$ from state s to state s', $r_a(s)$ is the expected immediate reward received after taking an action $a \in A(s)$ at state s.

Each state s in S corresponds to a possible estimated region of the attacker. To simplify our reference to each state, without loss of generality, we index the first home AP of the attacker as 0, and assume that AP 0 has N neighboring APs, which are indexed as $1, 2, ..., N$. Since each estimated region

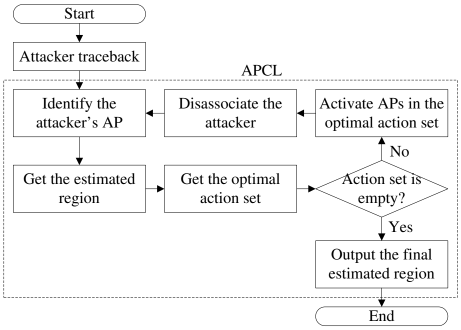
Start
Attacker traceback
APCL
Identify the attacker's AP
Disassociate the attacker
Activate APs in the optimal action set
No
Get the estimated region
Get the optimal action set
Action set is empty?
Yes
Output the final estimated region
End

Figure 4: APCL flow chart

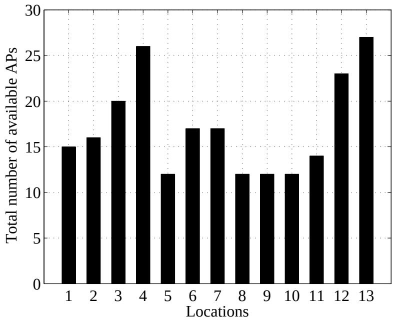

Figure 5: AP measurement

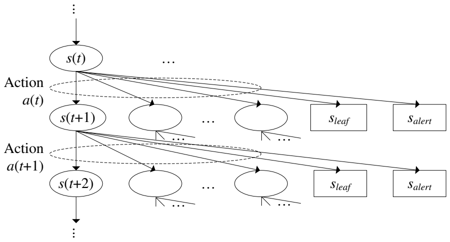

Figure 6: State transition diagram

is an intersection of multiple AP coverage regions, we denote a state in S as s_x, where the subscript $x = \{x_1, x_2, \ldots, x_k\}$ represents that the estimated region in this state is the intersection of the coverage regions of APs $x_1, x_2, \ldots,$ and x_k. The only exception in the state space notation is the state that corresponds to situations where the attacker is alerted. This happens when the attacker cannot find any activated APs to reconnect to within its communication range after it is disassociated from its previous home AP. This state is defined as an *alert state* s_{alert}. A state s whose action space $A(s) = \emptyset$ is called a *leaf state* s_{leaf}. A leaf state cannot transit to any other states.

Each action a in A corresponds to the set of APs activated by APCL in one step. Given the current estimated region of the attacker as s, note that APCL can only narrow down the estimated region if the attacker connects to a new activated AP whose coverage region intersects with but does not encompass s. Hence, by denoting the set of all such APs by $I(s)$, s's action space $A(s)$ is $2^{I(s)} - \emptyset$, where $2^{I(s)}$ is the power set of $I(s)$ and $\emptyset$ is the empty set.

A state transition happens when APCL takes an action by activating a set of APs and the attacker subsequently chooses a new activated AP to connect to. A transition from state s_x to s_y is possible under an action a if the following three conditions are satisfied.

- $x \subset y$ and $|y - x| = 1$. This is because when the attacker connects to a new home AP, APCL computes the estimated region as the intersection of coverage regions of all the previous home APs and this new home AP.
- $(y - x) \in a$. This is because the attacker can only connect to home APs which are activated in action a.
- $s_x \neq s_{alert}$ and $s_x \neq s_{leaf}$. This is because the localization process cannot continue when there are no APs to activate or the attacker is alerted.

4.1.2 Transition probability calculation

Given states s_x, s_y and an action a, if the transition from s_x to s_y is possible, APCL calculates the transition probability based on the following assumptions. At state s_x, we assume the probability density of the attacker's position is uniform over the estimated region of state s_x, and the attacker connects to any activated AP in its communication range with equal opportunity. Note that, these assumptions are derived by assuming that we have no knowledge of the position distribution and reconnection preference of the attacker, which is the worst case. Once we have some knowledge of the attacker's position distribution and reconnection preference, we can get a better and faster estimation of the attacker's position. With these assumptions, it can be proved that the state transition probability from state s_x to s_y under action a is

$$P_a(s_x, s_y) = \frac{\sum\limits_{\substack{v \subset (2^a - \emptyset) \\ (y-x) \in v}} \frac{(-1)^{|v|-1}}{|v|} Area(s_{x \cup v})}{Area(s_x)}, \qquad (2)$$

where $Area(s)$ represents the area of state s's estimated region, and 2^a is the power set of the set a. Correspondingly, the probability that the attacker is alerted due to its inability to connect to any activated AP is

$$P_a(s_x, s_{alert}) = 1 - \sum_{i \in a} P_a(s_x,\ s_{x \cup \{i\}}). \qquad (3)$$

4.1.3 Action reward function calculation

The reward for transition from state s_x to state s_y can be represented by $r(s_x, s_y) = Area(s_x) - Area(s_y) - C$, where $Area(s_x) - Area(s_y)$ represents the enhanced accuracy of position estimation, and constant cost C is used to reflect the fact that the more transitions, the longer the localization process and hence the higher risk that the attacker may move away before APCL finalizes the process. Hence, the action reward function in (1) becomes

$$r_a(s_x) = \sum_{(y-x) \in a,\ \text{or}\ s_y = s_{alert}} P_a(s_x,\ s_y) r(s_x,\ s_y). \qquad (4)$$

Since leaf states and alert states cannot transit to any other states, the expected aggregated rewards for these two types of states are defined as

$$r_a(s_{leaf}) = 0,\ r_a(s_{alert}) = -C_{alert}, \qquad (5)$$

where C_{alert} is the constant cost for alerting the attacker.

4.2 Optimization objective design

The optimization objective is defined as follows,

$$\max_{\pi} \sum_{t=1}^{N} \gamma^{t-1} r_{\pi(s(t))}(s(t)), \qquad (6)$$

where π is the MDP policy that is essentially a sequence of actions, $\pi(s(t))$ is the action taken at state $s(t)$ according to policy π, N is the horizon length that is the length of the MDP policy, $0 < \gamma \leq 1$ is the discounting factor for the future reward, and it captures the fact that the future reward is less important due to the chance that the attacker may move away in the future. Correspondingly, the optimal policy that achieves the optimal objective is,

$$\pi^* = \arg\max_{\pi} \sum_{t=1}^{N} \gamma^{t-1} r_{\pi(s(t))}(s(t)). \tag{7}$$

For state $s(t)$ where future transition is possible, by (4), the expected aggregated reward is defined as follows

$$\begin{aligned} r_{\pi(s(t))}(s(t)) \quad = & \sum_{\substack{s(t+1)-s(t)\in\pi(s(t)),\\ \text{or } s(t+1)=s_{alert}}} [P_{\pi(s(t))}(s(t), s(t+1)) \\ & \times r(s(t), s(t+1))], \end{aligned} \tag{8}$$

where $s(t+1)-s(t)$ is the AP the attacker connects to given action $\pi(s(t))$. The AP activation process of APCL, hence, can be depicted in the state transition diagram in Figure 6. At time t, the state is $s(t)$ and APCL activates APs in $a(t)$ based on π^* calculated by (7). Then, the attacker selects its new AP at time $t+1$ and the localization process transits to a new state $s(t+1)$. This process continues until APCL reaches a leaf state or an alert state.

4.3 Optimal decision calculation

Given the MDP model in Section 4.1, it is natural to solve the optimization problem by the traditional backward induction method for MDPs. Given any state s_x, the optimal action, i.e., the optimal set of APs to activate is as follows

$$\pi^*(s_x) \triangleq \arg\max_{a\in A(s_x)} \{r_a(s_x) + \gamma \sum_{\substack{(y-x)\in a,\\ \text{or } s_y=s_{alert}}} P_a(s_x, s_y)V(s_y)\}, \tag{9}$$

where $V(s_x)$, which represents the maximum expected aggregated reward for state s_x, is

$$V(s_x) \triangleq r_{\pi^*(s_x)}(s_x) + \gamma \sum_{\substack{(y-x)\in\pi^*(s_x),\\ \text{or } s_y=s_{alert}}} P_{\pi^*(s_x)}(s_x, s_y)V(s_y). \tag{10}$$

With (9) and (10), after network deployment, the optimal AP activation decision and the position estimate for each possible state can be precomputed. By storing the precomputed optimal decisions and position estimates to a database, APCL can easily decide which APs to activate based on the current estimated region of the attacker.

5. EFFICIENT IMPLEMENTATION

Although the optimal activation decision can be precomputed using the model in the previous section, its computation complexity may still be very high and undesirable. This is especially true if APs are densely deployed, which creates a large action space in the MDP. To reduce the computation overhead, this section develops an efficient approximation algorithm to the MDP.

The design of the algorithm is based on the observation that the MDP algorithm has two objectives: minimizing the estimation error and minimizing the total number of action steps. To identify the condition that minimizes the estimation error, we define a region that cannot be further divided by the boundaries of APs' coverage regions as an *unsplittable region.* If a region is the intersection of multiple APs' coverage regions, this region is defined as a *convex region.* Otherwise, it is a *non-convex region.* Clearly, an attacker can be either in a convex unsplittable region or in a non-convex unsplittable region. If the attacker is in a convex unsplittable region, APCL can eventually narrow the attacker down to this convex unsplittable region through the activation process. The estimation error, hence, is the determined by the size of the convex unsplittable region. If the attacker is in a non-convex unsplittable region, then in the activation process, APCL will eventually reach an alert state. In this case, the minimum estimation error is achieved when APCL ends in an alert state whose parent state corresponds to the minimum size convex region encompassing the non-convex unsplittable region. Based on these observations, the following Proposition can be proved.

PROPOSITION 1. *Given AP positions and coverage regions, the minimum localization error is determined by the shape, size and distribution of the unsplittable regions. A policy achieves the minimum localization error if it satisfies the following two conditions: 1) if the attacker is in a convex unsplittable region, no alerts happen, 2) or if the attacker is in a non-convex unsplittable region, the only alert happens when the current state corresponds to a minimum size convex region encompassing this non-convex unsplittable region.*

One simple method to guarantee the above two conditions is to keep the alert probability minimum at each step. Note that the alert probability in (3) can also be calculated as $P_a(s_x, s_{alert}) = A_n(a, s_x)/Area(s_x)$, where $A_n(a, s_x)$ is the area of the part of s_x that is not covered by any AP in action a. Hence, if the action a activates all the APs that intersect but do not encompass s_x, the alert probability is minimized. If we remove one or more APs from a without affecting $A_n(a, s_x)$, the alert probability remains unchanged and minimum.

Based on these observations, we design a three-step approximation algorithm. This algorithm first guarantees that the estimation error is minimized and then tries to minimize the total number of steps. The approximation algorithm starts from the AP set a_0 whose APs intersect with but do not encompass the current estimation region s_x. It iteratively removes AP j from $a_i, i = 0, 1, ...$, generating a_{i+1}. The selected AP j satisfies the following three conditions:

- $j \in J \triangleq \{j : j \in a_i, A_n(a_i - \{j\}, s_x) = A_n(a_0, s_x)\}$, i.e., the removal does not change the area of the part in s_x that is not covered by APs in a_0;
- $(j, k) \in [J', K] \triangleq \{(j, k) : j \in J, k \in a_i, k \neq j, s_{x\cup\{k\}} \subseteq s_{x\cup\{j\}}\}$, i.e., the intersection region between s_x and AP j encompasses the intersection region between s_x and another AP $k \in a_i$;
- $j \in J'' \triangleq \arg\max_{(j,k)\in[J',K]} \{Area(s_{x\cup j\cup k})\}$, i.e., the intersection region $s_{x\cup j\cup k}$ is the largest among all the possible (j, k) pairs in $[J', K]$.

The removal process stops when there is no AP satisfying the above conditions. The AP set consisting of the remaining APs is the quasi-optimal action set. In the above approximation algorithm, the first condition guarantees that the final estimation error does not degrade since the localization process will not terminate prematurely. The combination of the second and the third condition is a greedy method that tries to minimize the number of action steps

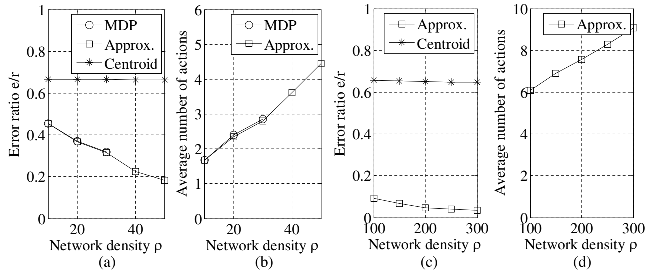

Figure 7: APCL performance

by eliminating APs that can potentially introduce unnecessarily more steps.

6. PERFORMANCE SIMULATION

We evaluate the performance of APCL in the random topology, where APs are uniformly distributed. The network density ρ is defined as the average number of APs per km^2. In our simulation, the discounting factor in (6) of APCL is set as $\gamma = 0.9$, the action cost C equals $0.2r^2$, and $r = 150$ m is the maximum range the attacker can communicate to an AP, which is determined by the ACK timeout configuration at the AP. The alert cost is $C_{alert} = 10^2C$. The position estimation of APCL is treated as the centroid of the estimated region. The localization performance is measured in error ratio e/r, where the localization error e is defined as the distance between the true position and the estimated position. The position estimation errors are averaged over 10000 possible attacker positions within the coverage region of the first home AP of the attacker. For each density value, the performance is averaged over 50 possible AP topologies. The APCL performance is compared with the traditional centroid method [9], which locates the target to the centroid of the neighboring APs.

Figure 7 (a) shows the error ratio for a low AP density that ranges from 10 APs per km^2 to 50 APs per km^2. This corresponds to 0.7 APs to 3.5 APs in a normal wireless card's communication range, which is far less than our measurement of AP densities in Virginia Tech campus (See Figure 5), demonstrating that there are more than enough APs to support APCL in current Wi-Fi networks. The MDP performance overlaps with the approximation algorithm performance, validating the effectiveness of the approximation algorithm. The APCL greatly outperforms the traditional centroid method. While the performance of the traditional method remains unchanged when the network density increases, the performance of APCL improves dramatically. It is because APCL gets a more accurate position estimation by luring the attacker to change its beam direction. When the network density increases, APCL uses more neighboring APs to locate the attacker and achieves more accurate localization. Figure 7 (b) shows the average number of actions under the low AP density. Again, the MDP performance overlaps with the approximation algorithm performance. The average number of actions increases when the network density increases. This is because there are more neighboring APs to locate the attacker, when the network density increases. This results in more actions to finally locate the attacker. This overhead of localization, however, is not significant since the dissociation and reassociation processes are fast operations that take around one second to complete [12]. Hence, the overhead of APCL is only several seconds, which is reasonable for the localization purpose.

In the case of high density, MDP cannot solve the optimization problem because of the great computation overhead, but our approximation algorithm can successfully solve it. The error ratio performance with respect to the network density is shown in Figure 7 (c). As shown in Figure 7 (c), APCL outperforms the traditional centroid method. The error ratio decreases when the network density increases. The average number of actions with respect to the network density is shown in Figure 7 (d). The average number of actions increases when the network density increases. The average localization time is still in the range of several seconds.

7. CONCLUSIONS AND FUTURE WORKS

Existing research of wireless localization focuses on either robust legitimate node localization or position claim verification. No efforts have been dedicated to locate the attacker, which may hide or falsify its position using advanced radio technologies such as directional antennas and SDRs. This paper addresses this threat in the context of the IEEE 802.11 WLAN by proposing a range-free localization scheme, termed APCL. In the future, we will extend APCL to track mobile attackers and locate multiple attackers simultaneously.

8. REFERENCES

[1] T. Baba and S. Matsuda. Tracing network attacks to their sources. *IEEE Internet Computing*, vol. 6, 2002.

[2] M. Snow and J.-M. Park. Link-layer traceback in ethernet networks. In *IEEE Workshop on Local and Metropolitan Area Networks (LANMAN)*. 2007.

[3] S. Northcutt and J. Novak. *Network Intrusion Detection*. Sams, 2002.

[4] V. Brik, *et al.* Wireless device identification with radiometric signatures. In *MobiCom'08*. 2008.

[5] J. Hightower and G. Borriello. A survey and taxonomy of location sensing systems for ubiquitous computing. UW CSE 01-08-03, Univ. of Washington, Dept. of Computer Sci. and Engineering, Aug. 2001.

[6] Y. Zhang, *et al.* Secure localization and authentication in ultra-wideband sensor networks. *IEEE JSAC*, vol. 24, no. 4, 2006.

[7] A. Srinivasan and J. Wu. A survey on secure localization in wireless sensor networks. In *Encyclopedia of Wireless and Mobile Communications*. CRC Press, Taylor and Francis Group, 2008.

[8] S. Capkun, *et al.* Secure location verification with hidden and mobile base stations. *IEEE Transactions on Mobile Computing*, 2008.

[9] N. Bulusu, J. Heidemann, and D. Estrin. GPS-less low-cost outdoor localization for very small devices. *IEEE Personal Communications*, vol. 7, no. 5, 2000.

[10] IEEE 802.11 standard, June 2007.

[11] madwifi.org. Long distance links with madwifi. http://madwifi.org/wiki/UserDocs/LongDistance, October 2008.

[12] V. Bychkovsky, *et al.* A measurement study of vehicular internet access using in situ wifi networks. In *MobiCom'06*. 2006.