Wireless Client Puzzles in IEEE 802.11 Networks: Security by Wireless

Ivan Martinovic, Frank A. Zdarsky, Matthias Wilhelm,
Christian Wegmann, and Jens B. Schmitt
TU Kaiserslautern
Distributed Computers and Systems Lab
67653 Kaiserslautern, Germany
{martinovic,zdarsky,m_wilhel,c_wegman,jschmitt}@informatik.uni-kl.de

ABSTRACT

Resource-depletion attacks against IEEE 802.11 access points (APs) are commonly executed by flooding APs with fake authentication requests. Such attacks may exhaust an AP's memory resources and result in denied association service, thus enabling more sophisticated impersonation attacks accomplished by rogue APs.

This work introduces the concept of *wireless client puzzles,* a protection method which assists an AP to preserve its resources by discarding fake requests, while allowing legitimate clients to successfully join the network. Rather than conditioning a puzzle's solution on computational resources of highly heterogeneous clients, the puzzles utilize peculiarities of a wireless environment such as broadcast communication and signal propagation which provide more invariant properties. Using an implementation of the proposed scheme, we demonstrate its effectiveness within a realistic scenario. Based on the insights from the implementation a simulation is used to extend the threat model and to scale up the scenario. Simulations verify our implementation results and show that the impact of flooding rate is decreased by 75% even if an attacker changes its position or manipulates its signal strength, while $\approx$ 90% of the legitimate stations are still able to successfully associate during an attack.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*

General Terms

Design, Security, Performance

Keywords

Denial-of-Service (DoS), client puzzles, impersonation attacks, wireless security

WiSec'08, March 31–April 2, 2008, Alexandria, Virginia, USA.
Copyright 2008 ACM 978-1-59593-814-5/08/03 ...$5.00.

1. MOTIVATION

The association procedure in the infrastructure mode of an IEEE 802.11 network utilizes a stateful protocol execution which is prone to DoS attacks, especially to memory exhaustion of wireless access points (APs). After receiving an authentication request, an AP reserves memory for a client's connection state and proceeds with the next stage of the IEEE 802.11 state machine. Through sending a high number of fake authentication requests, an attacker can easily exhaust an AP's memory which may result in faulty operation or even a full device crash.

This type of attack is not new and due to its simplicity and impressive results, in 2002/2003 the first tools were available to facilitate its execution. In [7], Floeter briefly describes his experience of testing an AP by flooding it after which even a manual reset of the AP was required.

To further investigate this matter, we selected 12 modern APs from popular brands such as Belkin, Linksys, Cisco, D-Link, Proxim, Siemens, and Lancom, and analyzed their behavior after subjecting them to an authentication flooding attack. We were interested to find out if operational anomalies or crashes were still possible on today's equipment and how easily legitimate clients were able to associate to an AP during an attack.

Surprisingly, only 2 out of 12 APs showed the desired functionality, whereas all other APs exhibited various anomalies in their operation. In contrast to older devices, modern APs implement a simple resource protection based on blocking of any new request after a certain threshold is reached. The duration of these blocking periods is at least several minutes during which *any* authentication request is denied.

Although the IEEE 802.11 standard provides failure codes to respond to a client with an appropriate reason for a failed authentication, only 5 APs have actually sent responses indicating the failed request, while others discarded it silently. As a result, legitimate clients had no chance of joining the wireless network during the attack.

Another observation was that in allowing the attacker to flood the APs for a few more minutes, 2 APs fully crashed, 1 AP rebooted and another 2 APs increased their response delay to a magnitude of seconds. This of course not only affected new clients trying to associate, but the overall wireless network where an attacker was free to install a rogue AP on the same channel, under the same MAC address, and to hijack abandoned clients by offering them fake association service (see [12] for more details).

1.1 Client Puzzles

Denial of Service is a prominent attack for which a number of efficient solutions within wired networks exists [14]. Most of the solutions are based on mechanisms that require a certain investment of resources on the client's side before the resource reservation on the server's side is made. A common implementation of this concept is a cryptographic puzzle, also known as a client puzzle [11, 3, 10, 9]. It consists of a challenge with scalable complexity where a solution can only be found by brute-force computation, i.e., investment of the client's resources based on computational power.

However, in contrast to wired networks where devices have similar hardware characteristics, wireless networks consist of heterogeneous devices such as PDAs, smartphones, and laptops. These devices vary considerably in their hardware capabilities, especially in computational power. This implies that client puzzles which are suitable for one mobile device may be too simple or too complex for another. As a result, they can either be completely ineffective against stations with higher performance, or may negatively affect legitimate users on slower devices. In search for a more invariant performance factor, Abadi et al. [2] describe memory-bound functions whose computation time is dominated by the latency of memory access rather than CPU power. To ensure that fast caches are ineffective in solving them, puzzles based on memory-bound functions require a reservation of high amounts of memory (larger than available caches). Although such puzzles have less varying computational time, their memory requirement presents a major drawback for using them in wireless environments.

1.2 Contribution

Varying hardware characteristics of clients make traditional client puzzles ineffective in wireless environments. Nevertheless, the problem of resource-depletion attacks on APs demands a solution as provided by puzzles in wired networks, i.e.:

1. Resource conservation
2. Acceptance of legitimate requests

The first requirement is already achieved by limiting the number of accepted authentications as observed at some of the modern APs. However, such a conservative protection fully ignores the legitimate clients and it may even serve to extend a resource-depletion attack to a more sophisticated one. For example, during blocking periods every request is ignored and legitimate clients have no chance to associate. If blocking periods are long (in our empirical experiment most APs block for a few minutes), a rogue AP could provide fake responses instead.

Due to broadcast communication and the lack of physical control over the traffic, the challenge is not only to detect an attack, but more importantly to isolate it from legitimate requests. We call this property *attack containment.* It is achieved through estimating an attacker's location and by limiting the number of valid requests sent from it. Assuming that an attacker can easily change its transmission power, antenna orientation, or its physical position, the puzzle is conditioned on the signal strength relationship to other stations. Any alteration influencing a signal's vicinity re-defines the puzzle and imposes further costs for solving it. From an attacker's point of view, such a puzzle represents a tradeoff between *avoiding the attack containment* vs. *the cost of frequent position changes.* In this work we demonstrate that attack strategy can neither seriously harm an AP's resources nor hinder legitimate clients from joining the network.

The rest of the paper is structured as follows. Section 2 introduces the concept of wireless client puzzles together with its real-world implementation utilizing off-the-shelf hardware within an indoor environment. In Section 3, using insights from the implementation and by means of simulations, different configurations of wireless networks are analyzed and the threat model is extended. Section 4 covers related work and Section 5 concludes this paper and discusses future work.

2. CONCEPT AND IMPLEMENTATION

2.1 Wireless Client Puzzles - Initial Idea

The most prevailing invariant characteristic among wireless clients is the use of wireless communication itself. Although there is variation in their transmission capabilities, the broadcast communication as an inherent characteristic of wireless communication is common to all of them. Therefore, the idea is to base the client puzzle on estimating the wireless environment by monitoring the communication of already associated stations. The puzzle is a question about which other stations are in the client's signal proximity, and can thus be labeled as neighbors. The received signal strength of neighbors is strong, contrary to non-neighbors which are received weakly in relation to a certain signal value. The necessary resource to solve such a puzzle is the time required for monitoring the channel to estimate the neighborhood. The valid solution of the puzzle is a client's *region* containing only stations considered as neighbors. We define the signal level that must be reached to qualify as a neighbor as the *Neighborhood Signal Threshold* (NST), which is randomly chosen, broadcast and frequently changed by the AP.

By listening on a wireless channel, a joining client is able to explore its neighborhood as the set of those stations from which it receives a signal *equal or greater* to the current NST, whereby this set then forms a region. After a client has defined its region, it is included in an authentication request. Every associated station within the transmission range of the joining client is able to capture it and check whether the region fulfills the neighborhood relationship from its own view.

There are two cases where the views of a joining client and already associated stations can differ: 1) if an associated station receives a signal from a joining station which is stronger or equal to NST but it is not contained in the region, or, 2) on the contrary, if the signal strength is below the given NST but the station is included in the region. In both cases the associated station broadcasts a warning which results in the AP denying authentication. As a result each region is conditioned on a client's signal relationship to other stations. Changing the position or transmission power without previously monitoring the channel results in poor estimates and a high probability of being denied. On the other hand, if the AP detects many authentication requests coming from the same or similar regions, it can choose to reject them, while simultaneously allowing clients from other regions to join.

The idea behind randomly selecting and frequently chang-

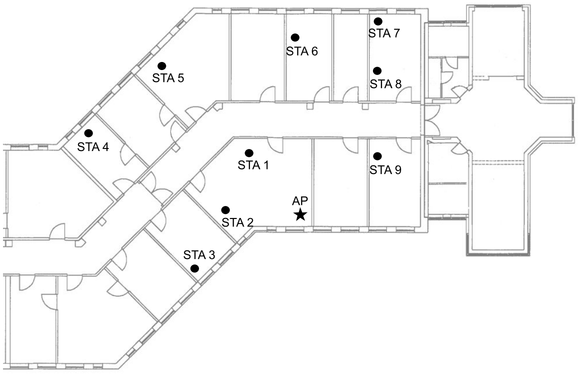

Figure 1: Testbed: university floor with 9 wireless stations placed in different rooms and with heterogeneous wireless characteristics.

STA	NIC
1	D-Link, DWL-G650
2	MacBook Pro Built-In
3	MSI, PC54G3 II
4	MSI, PC54G3 II
5	Orinoco 11a/b/g Gold
6	MSI, PC54G3 II
7	MSI, PC54G3 II
8	Orinoco 11a/b/g Gold
9	Siemens, PC Card 108
AP	Orinoco 11a/b/g Gold

Table 1: Wireless stations used in experiments.

ing the NST values is twofold. Even though the signal propagation contains randomness which requires channel monitoring to improve estimates of the client's neighborhood, randomly selecting NST values further increases uncertainty in predicting which regions are valid. Moreover, frequent NST changes also help those legitimate clients which, due to asymmetry of the received signal, suffer from, for them, an "unfortunate" NST that results in inconsistent regions (false positives). As this work show, the number of false positives forcing legitimate clients to re-attempt their associations remains small.

2.2 Testbed Implementation

We have set up a testbed by implementing a subset of the IEEE 802.11 distributed system services such as authentication and association procedures, and extended them with functionality for neighborhood monitoring, region building, and sending/processing of warning frames. The frames consist of an IEEE 802.11 frame header and an additional custom puzzle header that contains all required information (defined within a frame's custom Information Elements). All received frames are forwarded by the device driver in monitor mode to the driver via raw sockets, including a radiotap header which contains the frame's received signal strength.

For the experiments we used different WLAN cards with both built-in and external antennas to emulate heterogeneous wireless capabilities with different signal qualities (the WLAN cards used in the experiment are shown in Table 1).

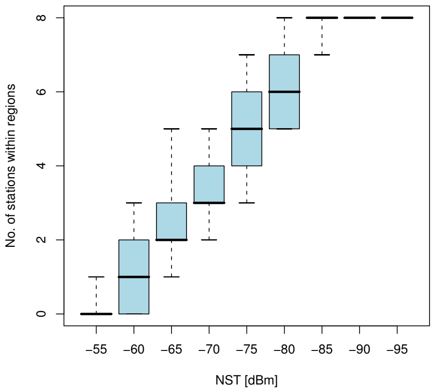

Figure 2: Box-plot (showing median, lower and upper quartiles, smallest and largest observation) of region sizes for various NST values within our testbed.

All cards are based on the Atheros chip-set but exhibit different transmission characteristics. For example, the D-Link card used by STA_1 highly varies in its transmission power (up to ± 15 dBm differences compared to other stations). The stations $STA_{3,4,6,7}$ have WLAN cards with external antennas that receive very well in general, although some of them have decreased transmission quality (e.g., $STA_{4,6}$ were positioned near the floor facing the wall).

The stations were placed in different rooms of the university building covering approximately 30x30 meters. The floor plan and the placement of the stations are shown in Figure 1.

The initial measurement of regions among all stations for several NSTs is shown in Figure 2. This figure provides an example of what region sizes to expect from different NST values.

2.3 Authentication Procedure

The experiment started with the AP broadcasting the NST within a beacon frame. The NST was randomly chosen by the AP from values between -55 dBm and -95 dBm in steps of 5 dBm and changed every 7 seconds. The joining station monitors the channel and computes the sample median (we choose a sample size of 20 received frames) that, after receiving a beacon frame and identifying the NST, is used to create a region by selecting those stations as neighbors whose signal strength is greater or equal to the current NST. The region is then sent along with the authentication requests to the AP. If no warnings arrive (the timer was set to 1 second) and no such region has already been used by another associated station, the AP responds with an authentication successful frame and proceeds with the association procedure. On the other hand, if a warning arrives the joining station is declined and it must wait for a different NST to re-attempt the authentication procedure.

For example, station STA_7 successfully joined at an NST of -80 dBm with a region $R_7 := \{1, 2, 3, 6, 8, 9\}$, while STA_2 joined at an NST of -60 dBm with a region $R_2 := \{1, 3, 6\}$. Station STA_9 had a problem joining at an NST of -75 dBm because STA_6 complained to be forgotten in its submitted region. Although the asymmetry in their received signal

STA	μ_{RSS}, σ^2_{RSS} (from STA_7)	μ_{RSS}, σ^2_{RSS} (to STA_7)
1	-62, 2.51	-68, 1.56
2	-70, 1.51	-65, 0.87
3	-76, 3.57	-77, 3.76
4	-76, 0.95	-83, 0.97
5	-81, 3.31	-84, 1.58
6	-62, 1.97	-69, 2.99
8	-39, 3.05	-39, 3.61
9	-61, 2.89	-62, 1.60

Table 2: An example of signal heterogeneity among stations: mean and std. deviation of signal strength from STA_7 to all other stations and vice versa (in dBm).

strengths was not high (STA_6 received STA_9 with -74 dBm, yet STA_9 received STA_6 with -78 dBm), the selected NST was exactly in-between which caused false positive rejection. Nevertheless, in its second attempt, STA_9 received an NST of -65 dBm and successfully joined with the region $R_9 := \{7, 8\}$.

We have repeated the measurement 6 times with different sequences of joining clients and interestingly, the average number of retries was at 1.5 (50% of stations were successful on the first authentication request while all others successfully joined within their second authentication attempt).

2.4 Tolerance Intervals

The problem of false positives, such as in the case of STA_9, occurs if the NST is chosen at a value which is influenced by differences in received signal strengths. An example is given in Table 2, which shows the received signal strength to and from STA_7. By only comparing the means given in the table, we can see that concerning STA_7, unfortunate NSTs are those of -65 dBm (where stations STA_1 and STA_6 would send a warning because from their view, they should be contained in the region of STA_7, on the other hand, STA_2 would complain that it does not belong to the region) and -80 dBm (where station STA_4 would complain to belong to the region). Thus, without considering signal variations, STA_7 has approximately a 78 % chance of successful authentication within the first attempt.

It can also be observed that differences between some of the stations are only 2-3 dBm, and yet still can result in a false positives. To mitigate the impact which small signal differences have on certain NSTs, we introduce *tolerance intervals* (TI). These are used by associated stations that verify the submitted region and tolerate signal differences up to a given TI. We have repeated our initial experiment using a TI of 5 dBm. The number of average authentication retries of the first experiment decreased from 1.5 to 1.1 ($\approx$ 87 % of all stations associated within the first authentication request, while all others succeeded within the second attempt).

2.5 Weak Positions

Although TIs increase the authentication probability of legitimate stations, there is a price to pay from a security point of view. If an attacker finds a physical position, or is able to find a signal strength for transmitting a region such that k stations tolerate it, it can generate as many as 2^k different regions that will not result in warnings. We refer to these as *Weak Positions* (WP). For example, if an attacker finds out that certain stations are tolerating its region, the attacker is able to create valid new regions by always including those stations that verify and combining them with stations that tolerate it.

On the other hand, every WP is conditioned by a signal neighborhood defined by the current NST. Therefore every time the NST is changed, the mapping of a WP to a physical position also changes (under the simplifying assumption of a free-space propagation model, the NST corresponds to a radius of a station's transmission circle which defines the neighborhood). Consequently, dynamically changing NST values changes the neighborhood relationships among stations. In a real-world scenario, the randomness of the realistic signal propagation and randomly chosen NST values leave the attacker without any knowledge where the next WP will physically be located.

To evaluate this property we use a simulation corresponding to our testbed (more details on the simulations are given in Section 4.1) and sample every 2 meters of the floor plan to detect which physical positions contain WPs and how often they occur.

The results are depicted in Figure 3, which shows the relative frequency of stations that tolerate regions of the joining client for all sampled physical positions. For example, in Subfigure (a), the selected NST is -65 dBm and there are physical positions where 60% of associated stations would tolerate received regions. Most of such positions are located towards the center of the floor plan. The reason is that the regions defined by an NST of -65 dBm contain only stations receiving each other with a very strong signal (using the analogy of a free-space propagation model, neighborhood radii of stations are small).

If the NST changes to -85 dBm, as given in Subfigure (b), although there are still WPs with 60% of all associated stations, this time their physical positions are located towards the border of the floor plan. Thus, by periodically selecting different NST values, the AP assures that the WPs are mapped to different physical positions. Furthermore, if such a position is found, it can only be used for the duration of the same NST value.

In Subfigure (c) an average of all physical positions and their WPs is depicted. The experiment shows that all positions are roughly equally likely (with an empirical probability of 0.2) to contain such a WP. Ensuring that the NST is randomly chosen and frequently changed, an attacker has no advantage of either staying at the same physical position or moving to another.

These results are obtained through simulations where an applied path loss model relies on an underlying distribution (we used Log-normal Shadowing with α = 9 dBm). In a real-world environment, the signal strength propagation is affected by various other environment-dependent parameters, which contribute to even more chaotic and erratic behavior of the wireless environment.

This should be considered as an advantage, as it contributes to uncertainty in predicting which positions can be used for an attack, or which regions are valid. It also demonstrates *security by wireless* paradigm followed in this work, where wireless communication and its chaotic properties are used to increase the security of the wireless networks.

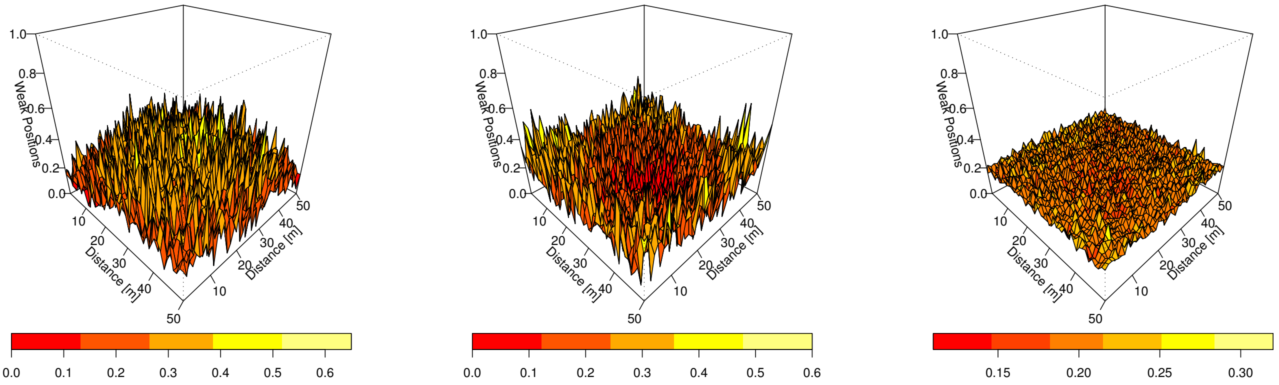

(a) NST= -65 dBm

(b) NST= -85 dBm

(c) Average of all NST values

Figure 3: (a), (b) occurrence of Weak Positions for different NST values, (c) occurrence of WPs in average over all available NST values.

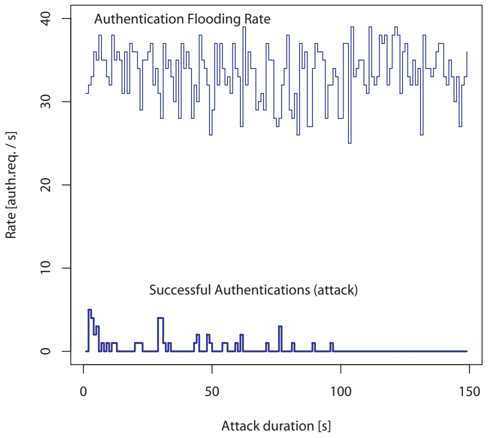

Figure 4: Flooding the AP with authentication requests and regions created from permutations of associated stations (brute-force attack).

3. REAL-WORLD IMPLEMENTATION RESULTS

This section discusses results from an implemented attacker who is ambitious to execute an:

1. *Attack on the AP* - a resource depletion attack by flooding the AP with authentication requests.
2. *Attack on the clients* - by exhausting available regions to deny further authentications of legitimate stations.

The attacker performs a brute-force attack by sending regions with all possible combinations of associated stations and is able to move around the floor plan. The associated stations are trusted and their identities have been authenticated (e.g., utilizing mutual authentication of the IEEE 802.11i security standard). They have no incentives to provide false warnings other than different views of the neighborhood.

The first attack is a primary attack as motivated at the beginning of this paper, while the second attack is a new front opened by our protection method. In executing the first attack we are interested in finding out if a resource-depletion attack is still possible, i.e., how many fake requests are allowed by an AP and what is the average flooding rate an attacker can achieve. In order to prevent new clients from joining the network, an attacker tries to block as many different regions as possible. Especially in small wireless networks with only a few associated stations, the number of possible regions is small enough to try all combinations of associated stations. Both attacks abuse weak positions and lossy channels, i.e., loss of warnings and loss of requests due to weak reception of stations positioned far away. The latter further aids the attacker to create more regions to flood with.

3.1 Resource Protection

The attack started with a flooding rate of ≈ 35 authentication requests per second. The testbed was as already described with a difference that one station (STA_2) was used as an attacker and displaced from the center. Due to the small scenario, the attacker easily tried all combinations of possible regions (there are only 256 possible regions in this scenario, thus it takes 6-7 seconds to try all of them). The attack lasted 2 minutes and 30 seconds and the results are shown in Figure 4. During a stationary attack where an attacker is not changing its positions, the AP only accepted 40 fake requests, while in a scenario where the attacker moved across the floor's hall, it succeeded in submitting only 53 authentications. The major reason for such a low number of accepted fake requests lies in our scheme's property of attack containment.

3.2 Attack Containment

As can be seen in Figure 4, attacking from a single position cannot provide a high rate of successful fake authentications. The reason is that associated stations in the attacker's signal proximity force it to include them in its regions, otherwise

Path loss model	Log-normal Shadowing
Path-loss exponent α	4.2
Log-normal σ	9 [dBm]
Mean difference μ	25%: $\mu_1 = 0$ 50%: $\mu_2 = 3$ 25%: $\mu_3 = 6$
No. of repetitions	10
Tolerance Interval	5 [dBm]
Surface	100x100 [m]

Table 3: Simulation Parameters.

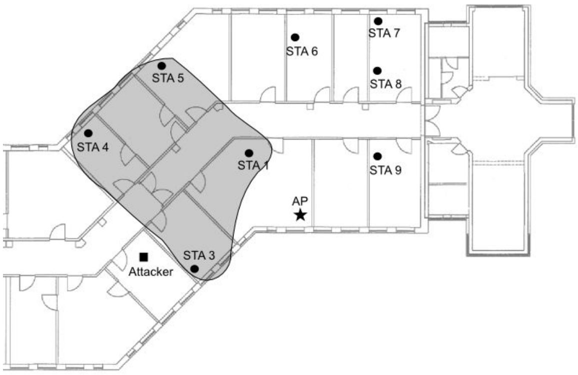

Figure 5: Attack containment: the attacker is successful only in blocking certain regions within its proximity like the region $\{1, 3, 4, 5\}$. All other regions were still unused and legitimate clients were able to join from them.

they send a warning. Due to the frequent NST changes, the attack is often interrupted and the lack of available regions causes its rate limitation. Hence, the attacker's strategy is to create regions with all the stations that would complain and then to combine them with stations placed far away (e.g., $\{1, 3, 6\}$, $\{1, 3, 7\}$, $\{1, 3, 6, 7\}$) which due to the weaker signal reception miss the fake regions and fail to validate them. Still, such an attack only marginally increased the number of accepted regions, and more importantly had no impact on blocking the legitimate clients from joining the network.

Predicting the stations which do not receive the fake regions while still being positioned within the AP's transmission range represents a significant problem to the attacker and random probing for such positions decreases its attack rate. On the other hand, such "artificial" regions, i.e., regions made only from certain stations with high spatial distribution, are never defined through monitoring of the wireless neighborhood, thus they are not used by legitimate clients (e.g. no NST can result in selecting only $\{3, 7\}$ as a valid region). The valid regions blocked during the attack were only those defined by stations in the attacker's signal proximity as depicted in Figure 5. The remaining free regions such as $\{5, 6\}$,$\{1, 5, 6\}$,$\{1, 4, 5, 6\}$ or any other defined by stations $STA_{7,8,9}$ were unused and clients were unhindered to join from them.

By walking along the university floor, an attacker increased the number of accepted valid regions and had more success in blocking legitimate clients from associating. Only small regions defined by NSTs like -55 dBm and -60 dBm such as $\{1\}, \{3\}, \{4\}, \ldots,$ were available and to use them legitimate clients had to be in the same room as the associated station. Nevertheless, requiring from an attacker to walk around the university floor is increasing its chance of being detected, and more importantly, this kind of blocking is hard to maintain. If the AP detects that many regions are used it can choose to periodically release some of them, requiring from the attacker to re-visit all previous physical positions and execute the same attack again.

4. SIMULATION AND ANALYSIS

Although the implementation has provided us with valuable insights, there are still open questions about the performance of this protection method in different scenarios, especially if more stations are involved and the attacker has more freedom of moving and searching for better positions to perform its attacks. Therefore, we used a packet-level simulation based on OMNeT++ [1] to analyze such cases.

4.1 Configuration

The simulation parameters are given in Table 3. For the path loss model we use Log-normal Shadowing with a standard deviation of 9 dBm and the difference that, instead of always using a zero mean, we explicitly model the discrepancies of the received signal strengths stemming from device heterogeneities. From the population of stations used in our simulations, 25% have no difference in mean of the received signal strength (thus having very similar transmission characteristics), 50% have a difference of 3 dBm and the remaining 25% represent stations with high differences in their transmission characteristics of 6 dBm.

Prior to starting with simulations of larger networks, we used exactly the same station placement as the one implemented on our university floor in order to adapt simulation parameters and to validate simulation results. As Figure 6 shows, the region sizes of stations are very similar to those from the implementation. The average number of authentication attempts without using TIs in the simulation is 2.1, while using TIs of 5 dBm reduces average to 1.3 attempts as shown in Figure 7 (in contrast to the implementation scenario with 1.5 and 1.1, respectively). The reason why simulation results where no TIs were used are slightly worse than the ones from the implementation is because of the path loss model used in simulations. Modeling high differences in means of received signal strength and adding σ of 9 dBm provides more frequent variation of the received signal as the one from the implementation. So, in a certain sense, simulation results can be considered conservative with respect to the performance of our protection scheme.

4.2 Extended Attacker Model

Using simulations we model a more powerful attacker with complete knowledge, i.e., although there may be stations not receivable within the attacker's transmission range, it can still use their identities to build regions.

Further, we model a *TX power attack,* which is an enhanced version of the stationary brute-force attack, during which different transmission powers for each authentication request can be selected. The values an attacker can choose from are between 1 mW and 100 mW. During this type of attack, the attacker counts the number of rejected authen-

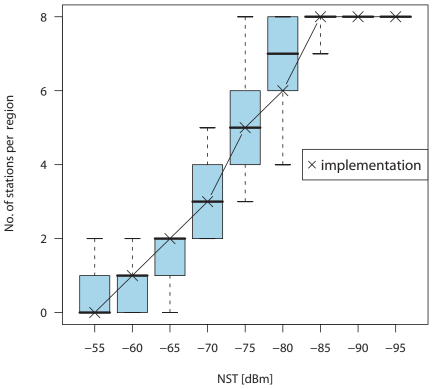

Figure 6: Box-plot of region sizes from simulation runs, compared with the median from the implementation scenario (see Figure 2).

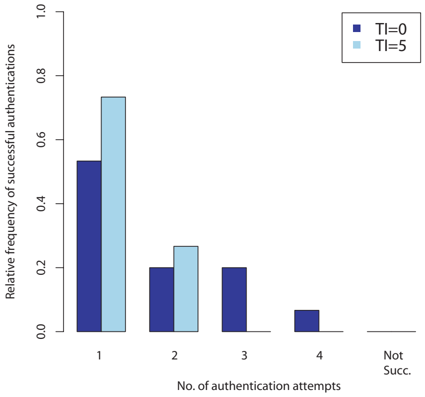

Figure 7: Testbed simulation: relative frequency of legitimate stations' attempts until successful authentication.

tication requests and changes the power only if a certain threshold has been reached.

The simulated *mobile attack* is an extended version of the implemented attack with more sinister characteristics. The attacker can walk through walls and uses a random walk algorithm to explore different physical positions. Again, it counts the number of unsuccessful authentication responses before moving to another position. To ensure that the random walk explores enough positions, we have increased the duration of these simulations.

4.3 Simulation Results

In every simulation run, the stations start their authentication procedure 30 seconds after the attack has begun. The starting times of the authentication attempts from legitimate stations are uniformly distributed between 30 and 60 seconds for the stationary attack scenario, and between 30 and 150 seconds for the mobile attack.

Every station has a maximum of 5 attempts to associate with the wireless network, otherwise it would be considered as not successful. The AP did not allow duplicated regions, i.e., every region could be used only once. The NST is chosen randomly from the same values as the ones implemented and is changed every 5 seconds. After each unsuccessful authentication attempt a legitimate station would wait 7 seconds before re-attempting.

4.3.1 Attacking Small Networks

Figure 8 shows the results of a stationary and TX power attack within the same number of stations as the one implemented but using 10 randomly chosen placements. Subfigures (a) and (b) are instances of one simulation run to illustrate one such attack as a timeline showing the attacker's flooding rate and its success rate similar to the implementation results given in Figure 4, while (c) shows the number of attempts of legitimate stations to authenticate averaged over 10 repetitions.

During the stationary attack only one station failed to associate. Similar to the implementation, the attacker had problems keeping its successful flooding rate. Again, its rate was interrupted by frequent NST changes and the lack of valid regions caused the attacker only to succeed in submitting 2-3 authentications every few seconds.

The TX power attack was more successful. Changing its transmission power after 50 unsuccessful authentication attempts provides an attacker with a much better brute-force method for detecting weak positions and discovering new regions. Although its successful attack rate is still very limited and discontinuous, 60 % of the stations could not authenticate within the first 5 attempts. Such a high number of blocked stations is due to the limited number of available regions which are now consumed by manipulating the transmission power. During this attack only small regions were left available for which an adequate NST had to be selected. This increased the number of required authentication re-attempts and resulted in a higher blocking success. Nevertheless, an improvement in such cases can be achieved if the AP periodically releases some of the blocked regions, which we analyze in Subsection 4.3.3.

4.3.2 Attacking Large Networks

In this experiment we increase the size of the network to 20 associated stations, which is a plausible number to various public hotspots such as airports, libraries, or conferences. Larger networks have much more regions to offer and thus result in an even higher authentication frequency. As depicted in Figure 9 Subfigure (c), without attack ≈ 80 % of stations joined within the first attempt. Similar to the results from previous subsection, the number of attacker's regions available for flooding rapidly decreased. On the other side, having many valid regions, even the TX power attack could not prohibit legitimate clients from joining. This time, less than 10% of stations failed to authenticate.

Interestingly, the attacker could not block large regions because the more station it includes in its fake region, the higher is the probability that some of those stations will receive it and respond with a warning. As a result, legitimate stations were able to authenticate under smaller NSTs (-80, -85,-90,-95 dBm) which resulted in selecting larger regions.

4.3.3 Mobile Attacker

In these simulations we investigate the impact of an attacker who is able to change its position. As a countermeasure against the TX power attacks and mobile attackers, the AP is configured to release its complete region list every 60s. The network size is set to 20 stations and their joining times are uniformly distributed between 30-150 seconds. Similar to the TX power attack, the mobile attacker counts the number of unsuccessful requests and changes its position

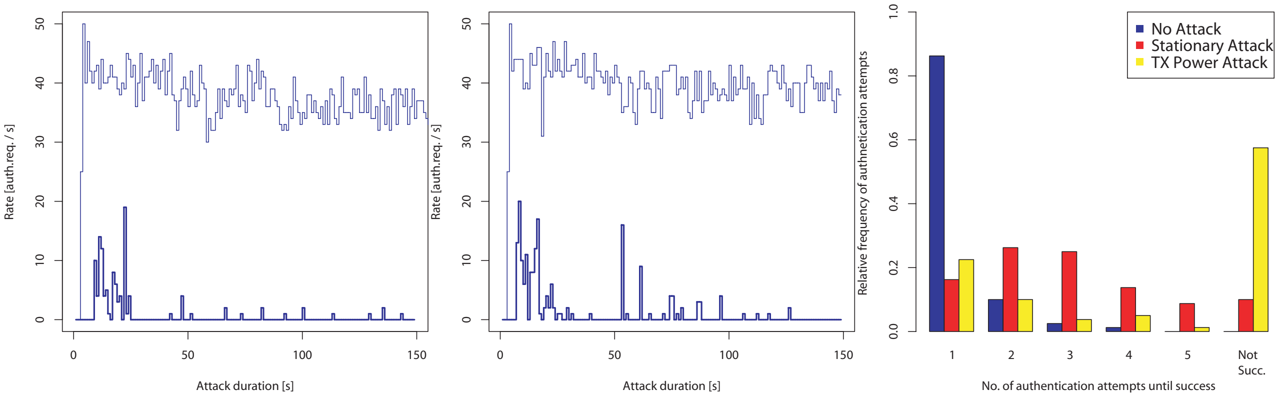

(a) (b) (c)

Figure 8: Simulation of the implemented scenario: (a) stationary attack, (b) TX power attack, (c) relative frequency of successful and failed authentication attempts of legitimate stations under both attacks.

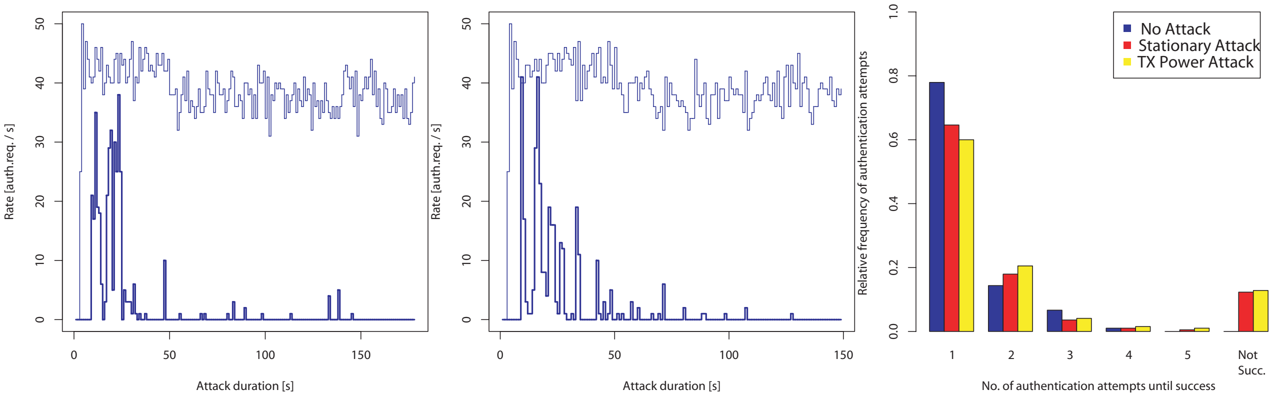

(a) (b) (c)

Figure 9: Simulation of a network with 20 associated stations: a) stationary attack, (b) TX power attack, (c) relative frequency of successful and failed authentication attempts of legitimate stations under attacks.

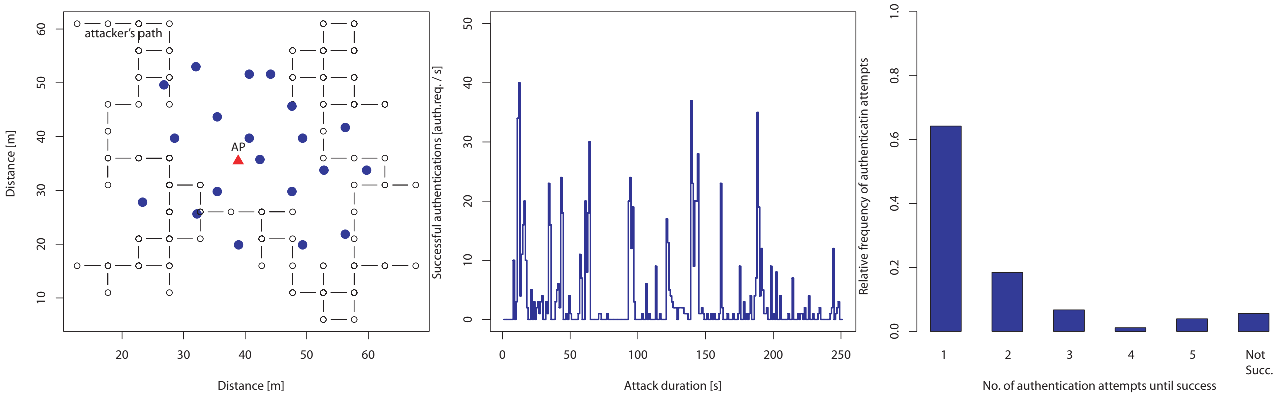

(a) (b) (c)

Figure 10: Scenario with a mobile attacker: (a) random walk taken by the attacker, (b) attacker's rate of accepted authentication requests, (c) relative frequency of successful and failed authentication attempts of legitimate stations during the attack.

after reaching 100 such rejections. Figure 10, Subfigure (a) gives an example of a random walk the attacker has taken and Subfigure (b) shows the rate of the attacker's successful authentication requests. This time, the rate is increased due to the periodic release of occupied regions, although it is still very limited, i.e, the flooding rate of 40 requests/s is reduced to 6 requests/s averaged over the simulation time. The number of successful authentications (as depicted in Subfigure (c)) is more than 90% for legitimate clients.

We have also experimented with multiple mobile attackers within a single scenario. Simulation required more than 4 simultaneous random walk attackers to achieve a successful flooding rate of more then 30 requests/s, and even then legitimate stations were able to associate.

The periodical release of used regions represents a tradeoff between two major objectives of wireless client puzzles – protecting the AP's resources vs. protecting new clients. This decision depends on the resources and capabilities of a dedicated AP and should be seen as a scenario-dependent parameter.

5. RELATED WORK

Substantial research contributions utilize properties of wireless communication to establish key management, authentication, and message integrity verification. For example, Perrig et al. present TESLA [15], a broadcast authentication protocol, which uses symmetric cryptography and a delayed key disclosure technique to achieve PKI properties offering protection against packet injection attacks. In [18], Čagalj et al., propose integrity codes (I-codes) for protecting message integrity without assuming shared authentication material or using cryptographic primitives. Using solely properties of the wireless channel and radio transmission, I-codes allow for broadcast message authentication and the so-called "authentication through presence" mechanism used for key establishment over insecure channels. Various other research papers are concerned with using physical properties of the radio environment to enhance security within different wireless networks (e.g.,[19], [17], [4],[16], [8])

However, measures against resource-depletion attacks, especially those considering early stages of protocol execution, are constrained by the number of available message exchanges and a lack of any prior knowledge of participating clients. In IEEE 802.11 networks, the authentication request is the first frame sent from a station after completing its network discovery process. Upon receiving it, the AP must decide whether to reserve resources or not. Therefore, protection against such attacks should avoid complex message exchanges and stateful protocol executions, but on the other hand, it can allow a certain amount of false negatives and false positives.

This is similar to sybil attacks, where a malicious client creates a large number of fake identities. In [5], Demirbas et al. describe a protection scheme against sybil attacks in WSNs using the ratio of different received signal strength measurements. By means of implementation they show that the received signal strengths provided by multiple receivers can be used to detect sybil attacks with a high precision and only few false negatives.

In [6], Faria et al. focus on identity attacks in IEEE 802.11 networks. The authors take advantage of radio signal strengths to build signalprints of clients and argue that an attacker has less control over them. They empirically show that signalprints are strongly correlated to physical location and that similar values of signalprints can mostly be found in close physical proximity.

Their findings are used against various attacks based on stealing identities (such as MAC addresses of legitimate clients). In our work, we consider a similar scenario, but follow a different approach to a solution. In contrast to [6], we avoid using absolute values of signal strengths and evaluate the protection under more realistic attacker models, like persignal strength manipulation and a mobile attacker.

The initial idea of wireless client puzzles was first described in [13]. In this paper we significantly improved and implemented the concept, and analyzed its performance under various aforementioned attack models.

6. CONCLUSION AND FUTURE WORK

In search for more performance invariant factors among devices, the wireless communication has shown to provide valuable properties for adapting the concept of client puzzles to the peculiarities of wireless networks.

We motivated this work by showing that certain attacks, although known for a long time, are still effective even against modern APs. At the moment, the only protection against authentication flooding attacks is a very conservative blocking of all authentication requests.

To assist APs to selectively block fake requests sent by an attacker, while at the same time allowing legitimate clients to successfully join the network, this work introduced the concept of wireless client puzzles. By using the property of attack containment, the attacker can be localized and its flooding attack limited to only a few requests per second. By means of a real-world implementation we showed the applicability of the introduced concept and verified its effectiveness in larger networks under more sophisticated attacker models through extensive simulations. Using wireless client puzzles in larger networks, more then 90% of stations were able to associate, while neither a mobile attacker nor manipulation of transmission power had an impact on an AP's resources.

Nevertheless, there are still various interesting issues left open for further investigation. For example, one of the most influencing factors of wireless client puzzles is the NST value. For legitimate clients, different NST values help to minimize the impact of signal asymmetry and for an attacker it limits the rate of successfully accepted requests. In this work the set of available NST values was large which sometimes required more re-attempts for a client to obtain an "appropriate" value. If an AP is able to estimate the distribution of associated stations, it could select more fine-grained NST values, while still minimizing the impact of an attack.

Furthermore, to avoid increasing the number of false positives in larger networks, an AP could select a subset of its associated stations to participate in wireless client puzzles. This would limit the number of warnings and false positives resulting from channel asymmetry. If such a subset is randomly mutated, an attacker can hardly predict which stations are currently monitoring the channel.

An alternative approach would be to use dedicated devices (e.g., sensors) installed by a network operator to implement wireless client puzzles instead of associated stations. Using a controlled deployment of such devices, a high number of available regions might be achieved even for a small number of stations.

7. REFERENCES

[1] OmNeT++: Discrete Event Simulation System. http://www.omnetpp.org/.

[2] M. Abadi, M. Burrows, M. Manasse, and T. Wobber. Moderately Hard, Memory-bound Functions. *ACM Transactions on Internet Technology*, 5(2):299–327, May 2005.

[3] T. Aura, P. Nikander, and J. Leiwo. DOS-Resistant Authentication with Client Puzzles. In *Revised Papers from the 8th International Workshop on Security Protocols*, pages 170–177, 2001.

[4] C. Castelluccia and P. Mutaf. Shake them up!: A Movement-based Pairing Protocol for CPU-constrained Devices. In *MobiSys '05: Proceedings of the 3rd International Conference on Mobile systems, Applications, and Services*, pages 51–64, June 2005.

[5] M. Demirbas and Y. Song. An RSSI-based Scheme for Sybil Attack Detection in Wireless Sensor Networks. In *WOWMOM '06: Proceedings of the 2006 International Symposium on World of Wireless, Mobile and Multimedia Networks*, pages 564–570, June 2006.

[6] D. B. Faria and D. R. Cheriton. Detecting Identity-based Attacks in Wireless Networks using Signalprints. In *WiSe '06: Proceedings of the 5th ACM workshop on Wireless Security*, pages 43–52, September 2006.

[7] R. Floeter. README file of void11 – Wireless LAN Security Framework. http://www.wlsec.net/void11.

[8] L. E. Holmquist, F. Mattern, B. Schiele, P. Alahuhta, M. Beigl, and H.-W. Gellersen. Smart-Its Friends: A Technique for Users to Easily Establish Connections between Smart Artefacts. In *UbiComp '01: Proceedings of the 3rd international conference on Ubiquitous Computing*, pages 116–122, September 2001.

[9] A. Juels and J. Brainard. Client Puzzles: A Cryptographic Countermeasure against Connection Depletion Attacks. In *Proceedings of the Network and Distributed Security Systems (NDSS'99)*, pages 151–165. IEEE Computer Society, February 1999.

[10] M. C. Lee and Chun-Kan Fung. A Public-key based Authentication and Key Establishment Protocol coupled with a Client Puzzle. *J. Am. Soc. Inf. Sci. Technol.*, 54(9):810–823, 2003.

[11] J. Leiwo, T. Aura, and P. Nikander. Towards Network Denial of Service Resistant Protocols. In *Proceedings of the IFIP TC11 Fifteenth Annual Working Conference on Information Security for Global Information Infrastructures*, pages 301–310, August 2000.

[12] I. Martinovic, F. A. Zdarsky, A. Bachorek, C. Jung, and J. B. Schmitt. Phishing in the Wireless: Implementation and Analysis. In *Proceedings of the 22nd IFIP International Information Security Conference (SEC 2007)*, pages 145–156, May 2007.

[13] I. Martinovic, F. A. Zdarsky, and J. B. Schmitt. Regional-based Authentication Against DoS Attacks in Wireless Networks. In *Q2SWinet '07: Proceedings of the 3rd ACM Workshop on QoS and Security for Wireless and Mobile Networks*, pages 176–179, September 2007.

[14] J. Mirkovic and P. Reiher. A Taxonomy of DDoS Attack and DDoS Defense Mechanisms. *SIGCOMM Computer Communication Review (CCR)*, 34(2), 2004.

[15] A. Perrig, R. Canetti, D. Tygar, and D. Song. The TESLA Broadcast Authentication Protocol. 2002. Cryptobytes, Volume 5, No. 2 (RSA Laboratories, pages 2–13. Available at www.rsa.com/rsalabs/cryptobytes/cryptobytes_v5n2.pdf.

[16] F. Stajano and R. Anderson. The Resurrecting Duckling: Security Issues for Ad-hoc Wireless Networks. In *Proceedings of the 7th International Workshop on Security Protocols*, pages 172–194, April 2000.

[17] A. Varshavsky, A. LaMarca, and E. de Lar. Enabling Secure and Spontaneous Communication between Mobile Devices using Common Radio Environment. In *Proceedings of the Eighth IEEE Workshop on Mobile Computing Systems and Applications*, February 2007.

[18] M. Čagalj, S. Čapkun, R. Rengaswamy, I. Tsigkogiannis, M. Srivastava, and J.-P. Hubaux. Integrity (I) Codes: Message Integrity Protection and Authentication Over Insecure Channels. In *SP '06: Proceedings of the 2006 IEEE Symposium on Security and Privacy (S&P'06)*, pages 280–294, May 2006.

[19] S. Čapkun, R. Rengaswamy, I. Tsigkogiannis, and M. Srivastava. Implications of Radio Fingerprinting on the Security of Sensor Networks. In *Proceedings of the 3rd International Conference on Security and Privacy in Communication Networks*, September 2007.