Towards Attack-Resistant Peer-Assisted Indoor Localization

Jingyu Hua(✉), Shaoyong Du, and Sheng Zhong

State Key Laboratory for Novel Software Technology,
Department of Computer Science and Technology,
Nanjing University, Nanjing, China
{huajingyu,zhongsheng}@nju.edu.cn, shaoyong.du.cs@gmail.com

Abstract. Peer-assisted smartphone localization, which leverages pairwise acoustic ranging among nearby peer phones to refine location estimation, significantly pushes the accuracy limit of WiFi-based indoor localization. Unfortunately, this technique is designed for non-adversarial settings. Dishonest peers may cheat in their distance measurements. Outside attackers may interfere with the acoustic ranging by continually broadcasting interference signals. In this paper, we propose countermeasures against each of these attacks. We first present an algorithm that can identify peers that are not cheating in the current localization, by searching for devices that can be embedded into the same plane according to their pairwise distances. We also design a robust acoustic ranging method exploiting signal modulation, which can defend effectively against intentional interference of outside attackers. Experimental results demonstrate that our countermeasures can greatly improve the robustness of peer-assisted localization.

Keywords: Peer-assisted localization · Acoustic ranging · Attack resistance · Smartphone

1 Introduction

Outdoor localization with smartphones is being widely used in our daily life. Indoor localization, however, remains in the elementary stage. Although there do exist many accurate indoor location mechanisms [3,8,16], they require either special hardware not yet supported by smartphones, or infrastructures expensive to deploy. Compared with them, WiFi-based localization, which leverages radio signals of existing WiFi access points, is much cheaper to implement on smartphones. Whereas current proposals [1,17,21,23,24] can only achieve room-level accuracy. For example, according to the experiments of Liu et al. [13], the errors of Fingerprinting Based Localization [2,17,23], which is one of the most

This work was supported by NSFC-61321491, NSFC-61425024, and NSFC-61300235.
The extended version of this paper is available at http://cl.ly/3d3V1z0D2d45.

G. Pernul et al. (Eds.): ESORICS 2015, Part II, LNCS 9327, pp. 417–437, 2015.
DOI: 10.1007/978-3-319-24177-7_21

popular WiFi localization technologies, may exceed 8 m. This is far from enough for indoor localization.

Targeting this problem, Liu et al. [13] propose an interesting approach of peer-phone assisted acoustic ranging to push the accuracy limit of WiFi based localization on smartphones. This proposal mainly takes advantage of the high accuracy of acoustic ranging (Measurement error can be confined below 5 cm [15]) to eliminate large WiFi localization errors. Specifically, when a target phone wants to improve its location accuracy, a group of nearby peer phones (including itself) are made to emit sound signals according to the schedule of a central server. They also make recordings in this process, and all the recorded sound files are sent back to the server. The server analyzes these files to calculate pairwise distances among these peers based on the Time of Flight (ToF) approach, and then uses the obtained relative positions of nearby phones as physical constraints to refine the WiFi-based location estimation of the target phone. Their experiments show that this approach can reduce the maximum and 80 % errors to 2 m and 1 m, respectively.

While such peer-phone assisted localization (PAL) is effective for non-adversarial settings, it is vulnerable to various attacks that can significantly reduce its high accuracy or even prevent it from working properly. First, PAL relies on a group of peers that are not under the control of the server. It is hard to guarantee that all of them are honest. Instead, they may cheat by emitting their signals without following the server's schedule (which we call *emission attacks* – see Sect. 2.2), or by directly manipulating the uploaded sound files (which we call *tampering attacks* – see Sect. 2.2), thus altering the distance measurements, and so disrupt the final location estimation of the target phone. In addition, current acoustic ranging can be easily interfered with by even outsider attacks: The server in PAL has no ability to associate ranging signals detected from the recorded data to their emitters except based on their present order. As a result, if an attacker continuously broadcasts his interference signals during the ranging process, the server may mistake the interference signals for legitimate ones and then obtain false distance measurements. We call these attacks *saturation attacks* (please see Sect. 2.2).

Location information is a critical input to a wide variety of high-level location-based applications. Compromised localization results are a serious threat because of their impacts on applications [4]. For example, indoor navigation application may bring users to wrong ways and advertising applications may deliver unmatched ads to users if localization results are compromised. So, in this paper, we aim to achieve a secure PAL resistant to the three attacks we mentioned above. Specifically, we make the following contributions:

We first study emission attacks. We show that when a peer launches this kind of attack, all the distance measurements between him and other phones are increased or decreased by the same value, which in theory makes this peer no longer embeddable in the same plane with any three honest ones. We leverage this observation to identify those peers having not performed emission attacks. In particular, we prove that if we find greater than or equal to $k+3$ peers (here, k is the number of dishonest peers) embeddable in the same plane according to

their distance measurements, we can guarantee that none of them is launching emission attacks.

Next, we consider the scenario involving tampering attacks. Dishonest peers launch this kind of attack could manipulate any distance measurement involving them to arbitrary values. We show that so long as we synchronize the clocks of peers in advance, the above result for emission attacks also applies to this scenario and the distance measurements among the $k + 3$ or more peers, which can be embedded into the same plane, do not suffer from any emission attacks.

We thus present an algorithm in search of no fewer than $k + 3$ peers that can be embedded in the same plane if they really exist. These phones are considered not cheating in the current localization. This algorithm has a worst-case computational complexity polynomial in n – the total number of peers. Since n is usually very small due to the limitation of the transmission range of beep signals, this algorithm is extremely fast. In addition, to apply this algorithm to the real world, we take ranging errors into consideration. We propose additional mechanisms to reduce false positives and false negatives due to these errors.

After that, we propose a new correlation-based beep detection approach that can well defend against saturation attacks during acoustic ranging. In this approach, beep signals assigned to peers are produced by modulating distinct pseudonoise (PN) codes on a sine carrier wave. Such modulations guarantee that these signals are poorly cross-correlated. The server can then precisely identify a specific beep from a recorded signal by searching for the earliest sharp peak of the cross-correlation function between them. If attackers have no knowledge of the PN codes, they have small chance of producing highly correlated beeps to interfere with the beep detection.

We finally perform extensive experiments to demonstrate the real effects of the above countermeasures. For the algorithm against dishonest peers, we show that it can achieve a high detection rate of honest peers while produce very few false positives. For the new correlation-based acoustic ranging method, we show that it confines the ranging errors to the same level (below 20 cm) before and after we introduce the saturation attack. By contrast, the errors of the existing energy-based method may exceed 1 m facing this attack.

2 Peer Assisted Localization and Attacks

2.1 Review of Peer Assisted Localization

Peer Assisted Localization (PAL) proposed by Liu et al. [13] uses nearby phones as reference anchors to push the limit of WiFi-based indoor localization. It exploits the high accuracy of acoustic ranging. There can be many possible designs of PAL protocols. To be specific, we use [13] as an example in this subsection and present attacks against it in Sect. 2.2. This technique includes the following four steps:

(1) WiFi-Based Localization: Smartphones use traditional WiFi-based localization techniques to roughly estimate their locations.

(2) Peer Recruitment: When a target phone wants to refine his location, he has to first broadcast a special audio signal to recruit a group of nearby peers. All the phones receiving this signal will report themselves to a central server.
(3) Relative Acoustic Ranging Among Peers: The server creates a time schedule to specify which device should emit a beep signal for ranging at what time. Involved devices send beeps accordingly while also turn on their recording function at the outset. All the recorded sound files are uploaded to the server, which will compute the relative distances among peers by estimating the sound travel time among them, and then construct a graph based these distances.
(4) Location Refining: the server then refines the location estimation of the target by superimposing the graph based on the relative distances among peers onto the graph base on the WiFi localization. The final result is sent back to the target.

We now review more details on the third step since most of our work below focuses on the acoustic ranging process in this step. We first want to mention that peer phones in this system are only responsible for emitting and recording beeps, and all the signal processing and computation are carried out on the server. This could avoid the inconvient peer-to-peer communication among smartphones. Next, the high accuracy of the acoustic ranging is based on an assumption that the server could precisely detect the earliest position of each beep signal in the recorded sound files of peer phones, which is corresponding to the arrival time of each beep signal on these phones. The server uses the difference between the specified emission time and the detected arrival time to estimate the distances between two phones.

There exist two methods to detect beep signals hidden in the sound files [13,22]. The first one is correlation-based. It computes the cross-correlation (CC) function of an emitted beep signal and a recorded signal. The first sharp peak in this function is considered with a high probability to be corresponding to the arrival time. The second method is energy-based. It generates beep signals with stronger energy than the background noises. Thereby, the point before and after which the energy distribution differs significantly is regarded as the arrival point of a beep signal. Through extensive experiments, Liu et al. [13] employ the second method due to its higher accuracy.

2.2 Attacks Against Peer Assisted Localization

As the current PAL system is designed for non-adversarial settings, it is highly vulnerable to both insider and outsider attacks. We now analyze the possible vulnerabilities of PAL and present the details of the attacks this research aims to address.

Insider Attacks. As peers in PAL are recruited randomly from the neighbors of the target and are beyond the control of the system, their behaviors are

hard to predict. We mentioned earlier in this section that peer devices in PAL are mainly responsible for signal emitting and recording. Dishonest peers may launch attacks by cheating in either of them.

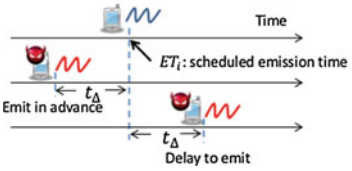

(a) An example of emission attacks

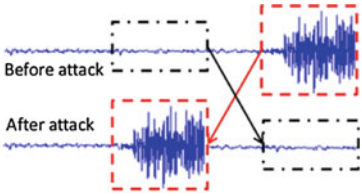

(b) An example of tampering attacks

Fig. 1. Attacks from dishonest peers

In the emission task, they may intentionally bring forward or delay their beep emissions rather than follow the schedule of the server. The server computes the distance between two devices by estimating the sound travel time between them. Suppose that the server schedules peer P_i to emit his beep signal b_i at time ET_i. Then, if P_i follows this schedule and really emits b_i at time ET_i, the server can learn the true distance between P_i and P_j by computing $d_ij = c(RT_{ij} - ET_i)$ provided that their clocks are synchronized. Here, c is the sound speed, and $RT_{ij}$is the arrival time of b_i at P_j, which is obtained by analyzing the audio record uploaded by P_j. However, if P_i sends his beep (ranging) signal t_δ earlier or later than ET_i as we show in Fig. 1(a), all the values of $RT_{ij}(j = 1, 2, \cdots, n)$ will be t_δ smaller or larger than the true value, respectively. As a result, all the distances d_{ij} from P_i to other peers are decreased or increased by the same value ct_δ.

In the recording task, dishonest peers may manipulate their recorded signals before sending them back to the server. Since the server learns RT_{ij} based on the detected position of the related beep signal in the recorded data uploaded by P_j, if P_j intentionally modifies the position of this signal (e.g., swaps the positions of this signal and a nearby noise window as we show in Fig. 1(b)), RT_{ij} will diverge from its real value and the obtained distance will also be changed. We name such kind of attacks *tampering attacks*. Compared with emission attacks, tampering attacks are relatively more flexible: a dishonest peer could freely choose one or several phones to change his distances to them without affecting other distances.

In addition, we assume that dishonest peers know their own locations in advance and may collude with each other.

Outsider Attacks. Outsider attacks are mainly caused by another vulnerability exists in the energy-based beep detection approach employed in acoustic-ranging. As we show in Fig. 2, the arrival of a beep signal will significantly change

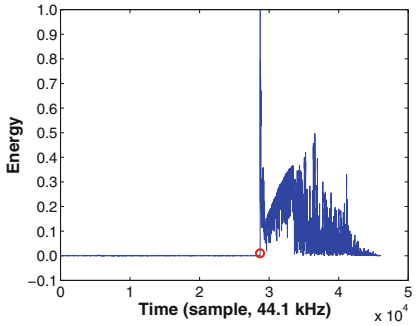

Fig. 2. Detecting the arrival time of a signal by identifying the energy saltation point (circled in red) (Color figure online)

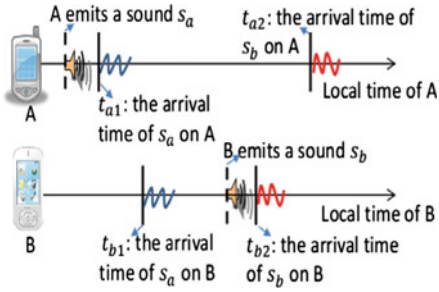

Fig. 3. Event sequence in the acoustic ranging algorithm of Beepbeep

the energy distribution of the recorded signal. This approach then locates ranging beeps in the recorded signal by identifying the earliest saltation points (marked by a red circle in Fig. 2) from where the energy distribution changes severely. It has no way to tell the difference between beep signals emitted by different devices but based on the assumption that the server's schedule makes these beep signals touch every device in a pre-defined order. Consequently, this technique can be easily fooled by interference signals from attackers even outside the system. For example, if an attacker in the proximity emits a strong signal at the same time when a peer emits his beeps, other peers may mistake the arrival of the interference signal for the legitimate one.

We may alleviate this problem by encrypting the communications between the server and the peers with individualized keys to prevent them and outside attackers from knowing the emission time of other devices. However, attackers may still obstruct the normal ranging with *saturation attacks*, i.e., constantly emitting interference signals. Such attacks will make the audio files recorded by peers full of energy saturation points. It is hard for the server to tell which one of them is due to the arrival of a real ranging signal. In this research, we focus on improving the PAL system to resist the above three attacks.

3 Countermeasure Against Insider Attacks

In this section, we first present countermeasures against insider attacks in an ideal scenario without measurement errors in the acoustic ranging. We then consider measurement errors to make our countermeasures more practical.

3.1 Countermeasure Against Emission Attacks

We first consider emission attacks, in which dishonest peers violate the server's schedule to bring forward or delay their beep emissions. We temporarily assume that dishonest peers do not perform tampering attacks except those altering distance measurements among themselves.

A naive way to defeat emission attacks is to employ the acoustic ranging algorithm of Beepbeep [15], which can reach an accuracy of 5 cm when there are *unintentional faults* in the timing of emitting signals. Below we briefly explain why this naive approach does not work when there are *intentional attacks*.

In this algorithm, the distance between two peers P_1 and P_2 is computed by

$$d_{P_1P_2} = \frac{c}{2} \cdot (t_{b1} - t_{a1} + t_{a2} - t_{b2}) + \frac{d_{AA} + d_{BB}}{2}, \tag{1}$$

where c is the speed of sound and d_{xy} is the distance between device x's speaker and device y's microphone. Other notations are illustrated in Fig. 3. If P_1 delays his emission for t_Δ due to unintentional faults, then both t_{a1}, the arrival time of his beep on P_1, and t_{b1}, the arrival time of the same beep on P_2, are increased by t_Δ. These two increments will cancel each other out in Eq. 1, and we will obtain the correct distance measurement between P_1 and P_2. In fact, the biggest advantage of this mechanism is that many uncertainties including those due to the lack of clock synchronization between devices can be eliminated in the calculation.

Nevertheless, if P_1 is an intentional attacker, he can easily bypass this countermeasure by manipulating his recorded sound file to change the value of t_{a1} before uploading it. For instance, P_1 can move a noise window of length t_Δ after the position of his ranging signal to its front in his recorded data, which will result in an error of $c/2 \cdot t_\Delta$ in every distance measurement involving P_1. Thus, we require a more advanced mechanism to defend against this type of attack.

Our new mechanism aims to filter out false distance measurements due to emission attacks. For simplicity, we present our theoretical analysis in a two-dimensional scenario. This is reasonable because the peers in PAL are on the same floor. In addition, we first assume the estimated distances among honest peers are exactly equal to their real values and we will consider measurement errors in the final algorithm design. Under these assumptions, we have the following lemma:

Lemma 1. *For four peers in the same plane but not in the same line, if one of them launches the emission attack while the other three keep honest, these four peers cannot be embedded in the same plane according to their distance measurements.*

Please find the proof in the extended version. This leads to our first theorem:

Theorem 1. *Let k be the maximum possible number of dishonest peers. Assume that the target phone is not in the same line with any three peers, and there are only emission attacks and no other attacks. If we can find $m \geq k + 3$ peers (including the target phone) that can be embedded into the same plane according to distance measurements among them, none of them performs emission attacks.*

Please find the proof in the extended version. We may design an algorithm based on Theorem 1 to identify a group of peers that do not launch emission attacks. However, Theorem 1 does not consider tampering attacks. If dishonest peers are allowed to perform tampering attacks, Theorem 1 is valid only when $k < 3$. We explain the reasons below.

3.2 Countermeasure Against Tampering Attacks

In PAL, the emission schedule in PAL guarantees that the arrival sequence of beep signals on each peer is exactly the same as the emission sequence. The server leverages this property to distinguish among different beep signals. Unfortunately, this property may also be exploited by dishonest peers. Two of them with synchronized clocks could associate their recorded data to calculate the time difference that the same beep touched them. Therefore, if there exist three or more dishonest peers that know their own locations, they could cooperate with each other to precisely locate every honest peer with TDoA localization technique.

Once the dishonest peers know the exact locations of the honest peers, they can further invalidate Lemma 1 and Theorem 1 by tampering with their recorded sound signals. For instance, if P_4 knows the positions of P_1, P_2 and P_3, he can easily predict the false distance measurements between him and these peers due to his emission attack. He can then adjust these measurements by altering the positions of corresponding beep signals in his recorded data to make them consistent in the same plane again. Therefore, due to the presence of tampering attacks, even if a peer can be embedded into the same plane with three honest peers, he may still perform emission attacks without being detected. Lemma 1 works iff $k < 3$ in the presence of tampering attacks. When $k < 3$, since the number of dishonest peers is not enough to position honest peers, they do not know how to adjust their distance measurements to make them consistent in the same plane.

We now present countermeasures against tampering attacks. Our proposal requires that all the peers synchronize their clocks before the localization. A possible solution is to use NTP (Network Time Protocol). There exists a free Android application, ClockSync, which can synchronize system clocks of Android devices with atomic time from local or remote NTP servers. If the user can use the root mode, the accuracy can reach milliseconds based on the NTP server.

Since the server knows the scheduled emission time of every peer, he can derive two distance estimations for each pair of peers (P_i and P_j): $d_{ij} = c \cdot (t_{ij} - t_{eP_i})$ and $d_{ji} = c \cdot (t_{ji} - t_{eP_j})$, where t_{ij} is the detected arrival time of P_i's ranging signal on P_j, and t_{eP_i} is the emission time of P_i. Once all the peers' clocks are synchronized, the two estimations for the same pair should be very close: if not exactly the same because of other local uncertainties of smartphones, the difference will, at least, be much less than the error due to attacks. If they are inconsistent (i.e. the difference between them is beyond some predefined

threshold ϵ), we can conclude that at least one peer has lied. On the other hands, however, if the two values are consistent, we cannot simply claim that both phones are honest because they may have colluded with each other. Under this assumption, we have the following theorem:

Theorem 2. *Let k be the maximum possible number of dishonest peers. Assume that there are Tampering Attacks in addition to Emission Attacks, but no other attacks. If we can find $m \geq k + 3$ peers (including the target phone) that can be embedded into the same plane according to the distance measurements among them, and if any three involved peers are not in the same line, then none of them has performed any attack that affects the distance measurements among them.*

Please find the proof in the extended version.

We design an algorithm based on Theorem 2 to identify a group of peers that do not lie about the distances among them. This algorithm can always succeed when the total number of peers (including the target phone) $n \geq 2k+3$. Its basic idea is to transverse triangles including the target phone (The total number of such triangles is $C(n-1,2)$) until we find one that can be embedded into the same plane with at least additional k peers. Specifically, for each triangle, we first test whether it can be embedded into the same planes with $R > k$ peers (We name them candidate peers) separately. If so, we place this triangle into a two-dimensional coordinate system by assigning the three peers coordinates consistent with their distance measurements. Once we do like this, the coordinates of the R candidate peers are also determined based on their distances to the triangle vertexes. We compute the required lengths of edges between each pair of the candidate peers and then remove those peers that the derived distance measurements based on their uploaded sound files are contradicting to the corresponding edge lengths. Afterwards, if the number of remained candidate peers is greater than k, the algorithm succeeds. Otherwise, it tries the next triangle. Due to the space limit, please find the pseudocode of this algorithm in Appendix.

It is easy to find that the worst-case time complexity of Algorithm 1 is $O(n^4)$. Since n is very small in PAL (usually below 10), this algorithm can be fast enough as you can see in Sect. 5. Once we identify these correct distance measurements, we can execute the last step of PAL to precisely locate the target phone.

Our discussion has assumed that there is no measurement error in acoustic ranging, which is obviously too ideal for the real world. Thus, to apply Algorithm 1 into the real world, we have to consider how to tolerate measurement errors. Due to the space limit, please find this part in Appendix.

4 Countermeasure Against Saturation Attacks

So far we have presented the countermeasures against insider attacks. We now consider countermeasures against Saturation Attacks.

As we pointed out in Sect. 2.2, the energy-based beep detection technique in current PAL is a major reason for the existence of saturation attacks. Beepbeep [15] uses a correlation-based technique, which is completely different from the energy-based method, to detect the location of a specific beep signal within a recorded signal. This technique has the potential to be extended to defend against saturation attacks: so long as we can prevent attackers from producing interference beeps that are highly correlated with ranging signals, they cannot affect the normal beep detection in theory because non-correlated interferences will not introduce noisy sharp peaks in the correlation functions with ranging signals.

Beepbeep does not fully solve this challenge since security is not its major focus. It makes all the ranging participants simply share the same ranging signal, which leads to that even if we prevent the outside attackers from knowing this signal, malicious peers inside can still launch saturations attacks to interfere with the server. In this section, we aim to present a new correlation-based method that can better resist saturation attacks.

4.1 Modulation-Based Beep Generation

Specifically, to resist the saturation attacks, we need beep signals that satisfy the following requirements:

- Each beep signal is only assigned to one peer. Aside from this peer and the server, it is infeasible for others to guess it in a short time.
- Beep signals have bad cross-correlation with each other or background noises. It is also hard to create a signal that is highly cross-correlated with a beep signal without knowing it.
- Each beep signal has a good auto-correlation property, which is critical for countering multi-path effects.

We find that the modulation technique in Direct Sequence Spread Spectrum (DSSS) [20], which is widely used in digital radio communication systems, can help us generate our required signals. The basic idea is to produce beep signals by using pseudonoise (PN) codes to modulate a sine sound-wave. For simplicity, we use Binary Phase Shifting Key (BPSK) as our modulation strategy. The correlation properties of the obtained signals are completely determined by PN codes (i.e., binary sequences in BPSK). If we can find a family of PN codes that satisfy the requirements above, the resulting signals hold similar properties.

We find that Maximum Length Sequences (M-Sequences) [20], which is a special class of pseudo-random binary sequences generated with maximal linear feedback shift registers, are ideal for such PN codes. An M-Sequence has a good autocorrelation property: the autocorrelation function $R_A(\tau)$ reaches its peak when $\tau = 0$, and as τ deviates from 0, $R_A(\tau)$ drops quickly. As a result, if we choose non-overlapped subsequences from the same M-Sequence as our PN codes for modulation, they must satisfy the requirements of R2 and R3. In addition,

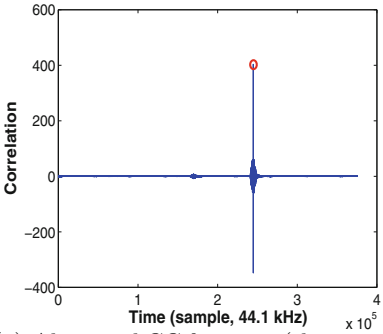

(a) Abnormal CC function (the maximal peak is due to a strong noise or interference signal)

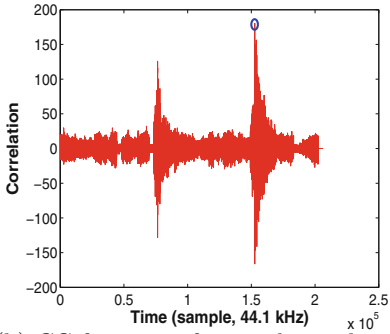

(b) CC function after applying the normalization

Fig. 4. An illustration of the effects of random signal attenuation on the results of CC functions (the peaks in red circles are due to the presence of ranging signals) (Color figure online)

to guarantee R1, we can use an extremely long M-Sequence that can produce a huge number of PN codes. For instance, if we use an M-Sequence of length $2^{48} - 1$ and suppose the length of the final PN codes is 256 bits (this length is long enough according to our experiments), we can obtain a family of 2^{40} PN codes. Given such a huge space, it is infeasible for an attacker to guess a specific code assigned to a peer.

4.2 Beep Detection

To detect a specific ranging signal from a recorded signal, we have to compute their CC function and then search for the sharp peak of this function. This task is not trivial due to the possible existence of some noise peaks, which are mainly caused by the correlation noises of the ranging signal with background noises, interference signals and the same signals due to the multi-path effect.

Because of the careful design of ranging signals, noise peaks due to background noises and interference signals are usually much lower than the desired peaks due to real ranging signals. However, there are abnormal cases where noise peaks suppress the true one when the strength of a ranging signal has become very weak when it arrives at another peer. We show an example in Fig. 4(a). We solve this problem by normalizing recorded signals before correlation. Specifically, when we compute the CC value of a recorded signal χ at time t with a ranging beep, we first find the maximal signal power of χ within a window of length $2d$ around t. Here, d is the length of the ranging beep. We then use this maximum value to normalize the signal segment involved in computing the CC value at t. Figure 4(b) shows the CC function after applying such normalization for the abnormal case in Fig. 4(a). We see that its maximum peak now becomes the one we desire.

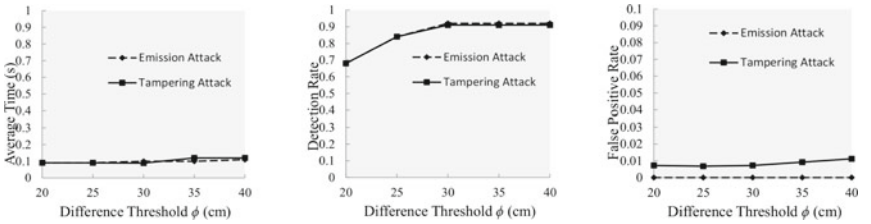

(a) Average time costs (b) Average detection rates of peers who did not cheat in the ranging (c) Average rates of false positives

Fig. 5. Performance of Algorithm 1 in experiments with real smartphones (ϕ is the difference threshold to etermine whether two distance values between the same pair of peers are consistent in the presence of measurement errors.)

Due to the existence of multi-path effects, the maximum peak we identify now may still not correspond to the earliest time that a ranging signal touches a phone. We deal with this problem with a simple method. We first locate the maximum peak whose correlation value is C_m. We then compute all the correlation values in a small window (500 samples) before the maximum peak and the first one whose value is larger than 85 %C_m is regarded as the earliest presence point of the ranging signal.

5 Evaluation

We have performed extensive experiments to evaluate the effects of our proposed countermeasures against the three attacks. We develop an Android application responsible for the beep emission and recording, and deploy it on five different Android smartphones: HTC G14, HTC G7, Motome 600 , HTC G12 and Coolpad 7260. All of them are equipped with two built-in speakers and one microphone that support 44.1 kHz sampling rate. In all the experiments, we use the back speaker and the microphone on every phone. We generate distinct beep signals for each device based on the design in Sect. 4.1. Due to the space limit, we have to put the detailed parameters for this process in the extended version.

To measure the distance between two phones, we make them emit their beep signals at a random order. All their recorded files are then manually copied to a desktop for analysis with a MatLab application that implements the automatic beep detection and distance calculation. We do not implement the last step of PAL because it depends on what WiFi localization technique that the peers use and is also beyond the scope of this paper. We only aim to verify whether our proposals can guarantee that all the distance measurements input into the last step are true.

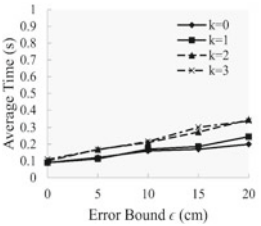

(a) Average time costs

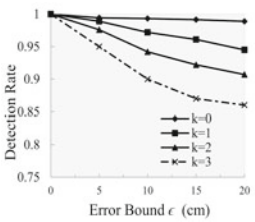

(b) Average detection rates of peers who did not cheat in the ranging

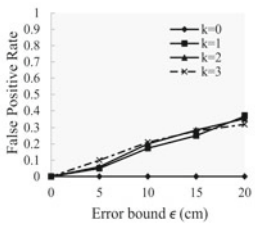

(c) Average rates of false positives

Fig. 6. Performance of Algorithm 1 in simulations with different numbers of dishonest peers (the measurement error is uniformly distributed over $[-\epsilon, \epsilon]$ and $\phi = 3\epsilon$)

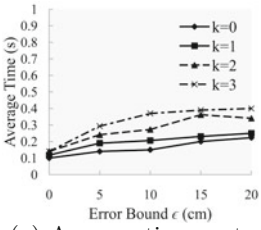

(a) Average time costs

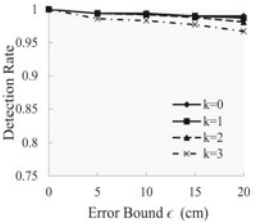

(b) Average detection rates of peers who did not cheat in the ranging

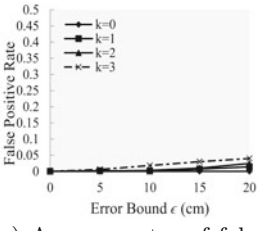

(c) Average rates of false positives

Fig. 7. Performance of Algorithm 1 (employing the patch in Appendix to filter false positives) in simulations with different number of dishonest peers (The measurement error is uniformly distributed over $[-\epsilon, \epsilon]$ and $\phi = 3\epsilon$)

5.1 Evaluation of Algorithm 1

We first evaluate the real performance of Algorithm 1 against dishonest peers. We conduct experiments in an empty room that is about 10 m × 6 m. We make five students carrying smartphones stand inside a circle of 2 m radius. Their topology is random but ensures line-of-sight between any two devices. Due to the limitation of the penetrating power of the used ranging signal, we do not consider the scenarios where some students stand in the corridor and some students stand in the room. We make phones emit their assigned ranging signals in a random order, and all their recorded data are uploaded to a desktop for analysis. We repeat this process for five times and each time all the students change their positions (i.e. topology). Therefore, we will obtain five groups of recorded signals.

Since we use Formula (1) to calculate distances, both emission and tampering attacks are carried out by manipulating dishonest peers' recorded signals, which are collected earlier. For emission attacks, we move the signal window containing the dishonest peer's ranging signal m samples ahead. For tampering attacks, we simply insert a noise window of n samples immediately after the ranging

signal. Here, both m and n are random values over $[260, 780]$, which will produce ranging errors over $[1\,\text{m}, 3\,\text{m}]$. Since the total number of peers is five, the maximum number of dishonest peer Algorithm 1 can tolerate is one. For each group of recorded data, we launch 100 emission attacks and 100 tampering attacks, respectively. Each attack randomly selects one device as the dishonest peer and another as the target peer. We perform the pairwise ranging with our matlab application and can get 500 inputs for Algorithm 1 for each kind of attack.

We then run Algorithm 1 for each input and each value of ϕ, which is the threshold for determining whether two distances are consistent, from 20 cm to 40 cm in steps of 5 cm. We have applied the patches for reducing FPs and FNs. To filter out FPs, we use 2.6 m as the upper bound for the average distance of an honest peer. The average detection rate of peers who did not cheat in the ranging, and the average rate of false positives are plotted in Fig. 5(b) and (c), respectively. We can see that the average detection rate exceeds 90 % in both two attacks when ϕ is larger than 30 cm. The false positive rates are always small enough to ignore. In addition, the average time cost is below 0.15 s and increases slightly in ϕ.

Due to the limited number of smartphones available for experiments, the above experiments only consider the scenario with three honest peers and one dishonest peer. To better evaluate the performance of Algorithm 1 with more peers, especially more dishonest ones, we do further simulations using Java language programs. We assume that there are 10 peer phones within an area of 4 m $\times$ 4 m and one of them is the target phone. We think it is difficult and also meaningless to employ more peers in the real world. The positions of each node is selected uniformly over the 4 m $\times$ 4 m area. Since the total number of nodes is fixed to ten, the maximum number of dishonest nodes this algorithm can tolerate is three. These dishonest peers are uniformly selected, and they are made to perform emission attacks and tampering attacks concurrently: bidirectional distance measurements between dishonest and honest nodes are enlarged by the same value δ, which is a random value over $[1\,\text{m}, 3\,\text{m}]$. We also assume the measurement error is uniformly distributed over $[-\epsilon, \epsilon]$, and two distances measurements between the same pair of nodes are thought consistent if and only if their difference is within $[-3\epsilon, 3\epsilon]$. We make Algorithm 1 try all the possible coordinate combinations of nodes in S_1 to reduce false negatives.

We run this algorithm 1000 times for each value of ϵ from 0 cm to 20 cm in steps of 4 cm and each possible value of k. In each run, all the nodes are assigned new positions. Figure 6(a) shows that the average simulation time increases in both k and ϵ. The reason for the first observation is obvious: a larger number of dishonest peers makes the algorithm harder to find enough number of nodes that can be embedded into the same plane. The reason for the second observation, however, is not so straightforward. According to our analysis, the increase is due to the fact that larger ranging errors usually bring more ambiguous nodes as Fig. 9(a) shows, which are extremely time consuming to deal with.

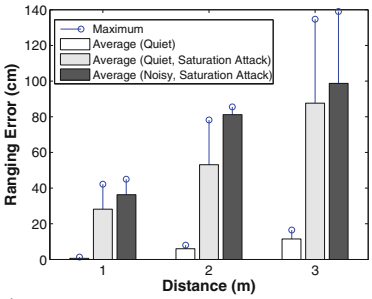

(a) Errors of the existing energy-based acoustic ranging approach

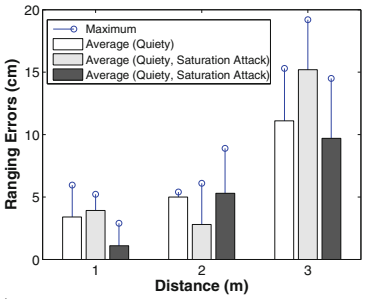

(b) Errors of the proposed correlation-based acoustic ranging approach

Fig. 8. Comparing the errors of the existing acoustic ranging approach in PAL and our proposal under different environments

We plot the average detection rate, and the average false positive rate in Fig. 6(b) and (c), respectively. We can see that Algorithm 1 works exactly the same as we expect in the cases without measurement errors: it can identify all the nodes that did not cheat in the ranging without bringing about any false positives. However, when we introduce measurement errors, Algorithm 1 produces both FNs and FPs. As we have applied our measure for avoiding FNs, the detection rate exceeds 90 % for all the cases except the one with three dishonest peers and $\epsilon = 20$ cm. The rate of FPs, however, is a little bit too high.

We then apply the patch mentioned in Appendix to reduce false positives: we consider a node to be malicious if its average distance to other nodes exceeds 3.7 m. We determine this threshold with extensive experiments. The new result is plotted in Fig. 7. The false positive rates are now confined below 4 %, which are much smaller that those in Fig. 6(c).

5.2 Evaluation of Countermeasures Against Saturation Attacks

We next evaluate the real effects of our acoustic ranging mechanism against saturation attacks. We conduct ranging between HTC G7 and Coolpad 7260 in the three indoor environments: (1) quiet, (2) quiet with Saturation Attack and (3) Noisy with Saturation Attack. Please find the detailed information about these environments in the extended version.

In all the experiments, we place two phones parallel to each other and back to back. For each environment, we vary the distance between two phones among 1 m, 2 m and 3 m, and repeat each experiment for four times. Besides our proposed correlation-based acoustic ranging, we also implement the energy-based method proposed in [13] as the reference. We present the average and the maximum ranging errors of these two methods in Fig. 8(b) and (a), respectively. We

see that the ranging errors of the existing energy-based method could exceed 1 m in the latter two cases when two phones are placed 3 m apart. For our proposal, however, we do not observe any big difference in the ranging errors between the environments with and without saturation attacks or noises. All the measurement errors are below 20 cm. We obtain similar results among other smartphones. This well demonstrates that our modulation-based acoustic ranging could well defend against the saturation attack.

In addition, the total signal processing time (i.e., compute the TOAs of the two ranging signals within the two recorded signals) in our proposal is 0.73 s on average. Although this value is much higher than 0.24 s in the energy-based approach, it is still acceptable. According to our analysis, most of the time is spent on the computations of the cross-correlation values. We may leverage the emission time of each ranging signal to reduce this time.

6 Related Work

Our work is towards robust peer-assisted indoor localization by defending against different kinds of attacks. Although PAL is novel, robust localization and ranging are not new. Related theory and systems have been developed for a long time, especially in the context of wireless sensor networks.

Most of current robust localization algorithms are designed for beacon-based localization systems. These systems require the presence of special nodes, so-called beacons or anchors, that know their own locations. Other nodes estimate their locations by measuring their distances to a set of beacons. Robust localization algorithms [10–12,18,25] then aims to enable a node to locate himself precisely even if some beacons are malicious. However, these algorithms have a premise that most of the beacons are still honest. For instance, Misra et al. [14] prove that the minimum number of honest beacons required for exact localization of the target in the presence of dishonest beacons is $\lceil n/2 \rceil + 2$, where n is the total number beacons. In our work, if we regard peer nodes as beacons, since the errors of their rough locations from WiFi localization reach 4 m on average, all of them can be regarded to have lied about their locations considering the strict requirement of indoor localization on the accuracy. As a result, we cannot directly use robust localization algorithms in this area. Compared with beacon-based localization, the scenario of beacon-less localization [5,9,19] is much closer to our problem. However, few of them consider security issues.

The last attack that we focus on is due to vulnerabilities in acoustic ranging. Girod et al. [6,7] propose a robust acoustic ranging mechanism that cleverly exploits signal modulation. Specially, the system is composed by a transmitter and a receiver. The transmitter produces a distinct sound by modulating a sine sound-wave with some special PN code. After the transmitter plays this sound, the receiver detects the arrival time of this sound by searching for the first sharp peak in the cross-correlation function between this sound and the recorded signal. Using a known sound speed and the emission time, the distance travelling from the transmitter and the receiver can be computed. While their work can work

very well even in very obstructed or noisy environments, they only consider the interference from background noises and reflections, and do not consider intentional interference from attackers.

7 Conclusion

Peer-assisted localization (PAL) through acoustic ranging could significantly improve the accuracy of WiFi localization. In this paper, we have studied the problem of robust PAL in the presence of dishonest peers and outside attackers. We first show that so long as the number of peers that can be embedded in the same plane according to their distance measurements exceeds some threshold, we can guarantee that none of them lies on these distances. We then present an algorithm based on this principle to identify peers having not cheated in the current localization, which can finish in polynomial time even in the worst case. We also present a robust acoustic ranging mechanism that leverages signal modulation to resist saturating interference from outside attackers. Extensive experiment on real smartphones have demonstrated that our countermeasures can greatly improve the robustness of peer-assisted localization.

Appendix

Practical Consideration of Measurement Errors

Our discussion has assumed that there is no measurement error in acoustic ranging, which is obviously too ideal for the real world. Thus, to apply Algorithm 1 into the real world, we consider how to tolerate measurement errors in this subsection.

Our solution is straightforward: facing measurement errors, Algorithm 1 regards two different distance measurements between the same pair of peers, or a distance measurement and its expected value, as consistent so long as their difference is below some pre-defined threshold ϕ. We empirically set $\phi = 3\epsilon$, where ϵ is the upper bound of the measurement error. Nevertheless, this mechanism has a side effect that it can produce both false negatives and false positives.

(a) Source of false negatives (b) Source of false positives

Fig. 9. False positives and negatives of Algorithm 1

False negatives (FNs) refer to that some peers which did not cheat in the ranging are falsely classified as dishonest by Algorithm 1. They mainly occur

Input: $P_{1,2,\cdots,n}$: n peer points
P_t: the target point
$\{d_{ij}|i,j \in \{1,\cdots,n,t\}\}$: d_{ij} is the distance between P_i and P_j based on the recorded data of P_j

for $i = 0,\cdots,n$ **do**
  if d_{it} *is conflicting with* d_{ti} **then** continue;
  for $j = i+1,\cdots,n$ **do**
    if $j - i > Malicous_{max} + 1$ **then** break ;
    if d_{jt} *is conflicting with* d_{tj} *or* d_{ij} *is conflicting with* d_{ji} **then** continue;
    Assign P_t, P_i and P_j two-dimension coordinates that meet their side-length requirements ;
    $failCount = 0$;
    Define an empty set S_1;
    foreach $P_k(k \notin \{i,j\})$ **do**
      if P_k *can be embedded in the plane of* $\triangle P_t P_i P_j$ **then**
        compute the coordinates of P_k;
        $S_1 = S_1 \cup \{P_k\}$;
      end
      else
        $failCount++$;
        if $failCount > n - 2 - Malicious_{max}$ **then** break;
      end
    end
    if $Size(S_1) < Malicious_{max}$ **then** continue;
    Define another empty set S_2;
    foreach $P_a \in S_1$ **do**
      if S_2 *is empty* **then**
        $S_2 \cup \{P_a\}$;
        Continue;
      end
      foreach $P_b \in S_2$ **do**
        Compute $d'_{ab} = \sqrt{(P_a.x - P_b.x)^2 + (P_a.y - P_b.y)^2}$;
        if d_{ba} *is consistent with* d'_{ab} **then** $S_2 \cup \{P_a\}$;
        if d_{ab} *is conflicting with* d'_{ab} **then** $S_2 - \{P_b\}$;
      end
    end
    if $Size(S_2) >= Malicious_{max}$ **then**
      Output $S_2 \cup \{P_i, P_j\}$;
      Stop the Algorithm;
    end
  end
end

Algorithm 1. Algorithm to identify true distances in the presence of emission attacks and tampering attacks

in the special cases that the three vertices of the winning triangle $\Delta P_i P_j P_t$ in Algorithm 1 are either too close to each other or approximately in the same straight line, which makes the algorithm determine false positions for some nodes in the presence of measurement errors. We show an example in Fig. 9(a). Suppose $\Delta P_1 P_2 P_3$ is the winning triangle and the algorithm is computing the coordinates of P_4 at line 11. We also assume that another node P_4' satisfies the condition: $d_{P_4 P_1} = d_{P_4' P_1}$ and $d_{P_4 P_2} = d_{P_4' P_2}$. Then, since P_1, P_2, P_3 are almost collinear, the distance measurement between P_4' and P_3 can be even closer to $d_{P_4 P_3}$ than the measurement between P_4 and P_3 in the presence of measurement errors. As a result, the algorithm may assign the coordinates of P_4' to P_4, which will lead to contradictories at line 25 or 26 and then falsely classify P_4 as a dishonest node. We can solve this problem by recording both coordinates of such special nodes, and then executing Line 21–31 for each possible coordinate combination of the nodes in S_1.

False positives (FPs) refer to that some dishonest peers which launched attacks are falsely reported as honest by Algorithm 1. They are mainly caused by dishonest nodes that are located on one side of the other nodes (i.e., not surrounded by any triangles formed by other nodes), launching emission attacks. We show a typical example in Fig. 9(b). Suppose P_1, P_2 and P_3 are honest, while P_4 is dishonest and delayed his emission for some time. So, the three distance measurements between P_4 and the other three nodes are increased by the same value, which is impossible in theory. However, when we move P_4 further from $\Delta P_1 P_2 P_3$, for example to the new position of P_4', the real increments of the three distances are very close even if their absolute values are very large. Thus, in the presence of measurement errors, these different changes may be approximated to be equal, which leads to the fact that P_4 can be accepted to be at some position in the same plane of $\Delta P_1 P_2 P_3$.

We find that dishonest nodes causing FPs usually do not choose to shorten their distance measurements. This is because the peer phones in PAL should be in the vicinity in order to receive each other's ranging signals. If the dishonest peers not surrounded by other nodes shorten their distance measurements, they are very likely to be falsely positioned at a place surrounded by some honest peers, and so they can be captured. In addition, some distance measurements in this case may even become minus, which is obviously ridiculous. Therefore, these dishonest peers usually choose to enlarge their distance measurements. However, since they are located on one side of the other nodes, their real average distances to other nodes are already larger than those of normal nodes. If they further enlarge their distance measurements, they will expose a larger anomaly. We leverage this observation to add a patch to Algorithm 1 to reduce false positives: before we check the size of S_2 at Line 32, we first remove each node whose average distance measurement to other nodes exceeds some threshold.

References

1. Azizyan, M., Constandache, I., Choudhury, R.R.: Surroundsense: mobile phone localization via ambience fingerprinting. In: Proceedings of the 15th MOBICOM, pp. 261–272. ACM (2009)
2. Bahl, P., Padmanabhan, V.: Radar: an in-building RF-based user location and tracking system. In: Proceedings of the 19th INFOCOM, pp. 775–784. IEEE (2000)
3. Borriello, G., Liu, A., Offer, T., Palistrant, C., Sharp, R.: WALRUS: wireless acoustic location with room-level resolution using ultrasound. In: Proceedings of the 3rd MobiSys, pp. 191–203. ACM (2005)
4. Chen, Y., Trappe, W., Martin, R.P.: Attack detection in wireless localization. In: Proceedings of the 26th INFOCOM, pp. 1964–1972. IEEE (2007)
5. Fang, L., Du, W., Ning, P.: A beacon-less location discovery scheme for wireless sensor networks. In: Proceedings of the 24th INFOCOM, vol. 1, pp. 161–171. IEEE (2005)
6. Girod, L., Estrin, D.: Robust range estimation using acoustic and multimodal sensing. In: Proceedings of the 2001 IROS, vol. 3, pp. 1312–1320. IEEE (2001)
7. Girod, L.: A self-calibrating system of distributed acoustic arrays. Ph.D. thesis, University of California Los Angeles (2005)
8. Hazas, M., Kray, C., Gellersen, H., Agbota, H., Kortuem, G., Krohn, A.: A relative positioning system for co-located mobile devices. In: Proceedings of the 3rd MobiSys, pp. 177–190. ACM (2005)
9. Ji, X., Zha, H.: Sensor positioning in wireless ad-hoc sensor networks using multidimensional scaling. In: Proceedings of the 23rd INFOCOM, vol. 4, pp. 2652–2661. IEEE (2004)
10. Li, Z., Trappe, W., Zhang, Y., Nath, B.: Robust statistical methods for securing wireless localization in sensor networks. In: Proceedings of the 4th IPSN, pp. 91–98. IEEE (2005)
11. Liu, D., Ning, P., Du, W.: Detecting malicious beacon nodes for secure location discovery in wireless sensor networks. In: Proceedings of the 25th ICDCS, pp. 609–619. IEEE (2005)
12. Liu, D., Ning, P., Liu, A., Wang, C., Du, W.: Attack-resistant location estimation in wireless sensor networks. ACM Trans. Inf. Syst. Secur. (TISSEC) **11**(4), 22 (2008)
13. Liu, H., Gan, Y., Yang, J., Sidhom, S., Wang, Y., Chen, Y., Ye, F.: Push the limit of WiFi based localization for smartphones. In: Proceedings of the 18th MOBICOM, pp. 305–316. ACM (2012)
14. Misra, S., Bhardwaj, S., Xue, G.: Rosetta: robust and secure mobile target tracking in a wireless ad hoc environment. In: Proceedings of the 2006 MILCOM, pp. 1–7. IEEE (2006)
15. Peng, C., Shen, G., Zhang, Y., Li, Y., Tan, K.: Beepbeep: a high accuracy acoustic ranging system using cots mobile devices. In: Proceedings of the 5th SenSys, pp. 1–14. ACM (2007)
16. Priyantha, N., Chakraborty, A., Balakrishnan, H.: The cricket location-support system. In: Proceedings of the 6th MOBICOM, pp. 32–43. ACM (2000)
17. Rai, A., Chintalapudi, K., Padmanabhan, V., Sen, R.: Zee: zero-effort crowdsourcing for indoor localization. In: Proceedings of the 18th MOBICOM, pp. 293–304. ACM (2012)
18. Ray, S., Ungrangsi, R., Pellegrini, D., Trachtenberg, A., Starobinski, D.: Robust location detection in emergency sensor networks. In: Proceedings of the 22nd INFOCOM, vol. 2, pp. 1044–1053. IEEE (2003)

19. Shang, Y., Rumi, W., Zhang, Y., Fromherz, M.: Localization from connectivity in sensor networks. IEEE Trans. Parallel Distrib. Syst. **15**(11), 961–974 (2004)
20. Simon, M., Omura, J., Scholtz, R., Levitt, B.: Spread Spectrum Communications Handbook, vol. 2. McGraw-Hill, New York (1994)
21. Wu, C., Yang, Z., Liu, Y., Xi, W.: Will: wireless indoor localization without site survey. In: Proceedings of the 31st INFOCOM, pp. 64–72. IEEE (2012)
22. Yang, J., Sidhom, S., Chandrasekaran, G., Vu, T., Liu, H., Cecan, N., Chen, Y., Gruteser, M., Martin, R.: Detecting driver phone use leveraging car speakers. In: Proceedings of the 17th MOBICOM, pp. 97–108. ACM (2011)
23. Yang, Z., Wu, C., Liu, Y.: Locating in fingerprint space: wireless indoor localization with little human intervention. In: Proceedings of the 18th MOBICOM, pp. 269–280. ACM (2012)
24. Ye, H., Gu, T., Zhu, X., Xu, J., Tao, X., Lu, J., Jin, N.: Ftrack: infrastructure-free floor localization via mobile phone sensing. In: Proceedings of the 10th PerCom, pp. 2–10. IEEE (2012)
25. Zhong, S., Jadliwala, M., Upadhyaya, S., Qiao, C.: Towards a theory of robust localization against malicious beacon nodes. In: Proceedings of the 27th INFOCOM, pp. 1391–1399. IEEE (2008)