Secure Positioning with Non-Ideal Distance Bounding Protocols

Pericle Perazzo, Gianluca Dini
Department of Information Engineering, University of Pisa, Pisa, Italy
Email: {pericle.perazzo, gianluca.dini}@iet.unipi.it

***Abstract*—Distance bounding protocols are secure protocols to determine an upper bound to the distance between two devices. These protocols have shown to be useful for many tasks, from proximity verification to secure positioning. Unfortunately, real distance bounding protocols hardly fulfill the claimed property. Attacks at the PHY layer may cause significant reductions on the estimated upper bound. These attacks can be mitigated, not eliminated, by changing the receiver architecture and the PHY layer. Every distance bounding protocol is thus *non-ideal*.**

In this paper, we study the impact of non-ideal distance bounding on the reliability of secure positioning techniques. We show that a reduction of 10 meters, which is possible against a real PHY layer, allows the adversary to falsify a position of 21 meters. We also propose two countermeasures to mitigate the problem, and then estimate their efficacy by simulations.

I. Introduction

Secure positioning aims at determining the real position of a device in the presence of an adversary, both internal or external, determined upon falsifying it [13], [3], [11], [14]. We focus on *range-based* secure positioning, which acts by directly measuring distances (ranging) from a set of anchors whose positions are known. Distances are securely measured by means of wireless *distance-bounding protocols* [1]. These protocols measure a distance between two devices, a *verifier* and a *prover*, in such a way that an adversary cannot falsify the measurement to be shorter than the real one. In other words, ideally, no *reduction attack* is possible. They are usually realized on impulse-radio ultra-wide band (IR-UWB) technology, which is capable of sub-meter precision in ranging operations.

Recent research showed that ideal distance bounding protocols do not exist in practice [4], [8]. Attacks at the PHY layer may cause significant reductions on the round-trip time measurement, and thus on the estimated distance, so making reduction attacks actually possible. A proper design of the receiver and the PHY layer only makes it possible to mitigate these attacks.

In this paper, we perform a first analysis of the impact of non-ideal distance bounding protocol on the reliability of secure positioning techniques. We evaluate such an impact in terms of two metrics: namely the *uncertainty area* and the *spoofable distance*. We show that a reduction of 10 meters, possible against the standard IEEE 802.15.4a UWB PHY [8], can cause a position spoofing of 21 meters. We also propose two countermeasures that can mitigate the problem: (i) to discard the positions which are falsifiable over a certain limit, and (ii) to place the anchors following a particular scheme having certain geometrical properties. We evaluate these countermeasures by simulations.

The rest of the paper is organized as follows. Section II presents the current state of the art. Section III introduces the distance bounding protocols and the secure positioning technique. Section IV describes how to perform reduction attacks against the PHY layer of a distance bounding protocol. Section V analyzes the impact of non-ideality in secure positioning and the possible countermeasures. Section VI evaluates experimentally the non-ideality impact and the effectiveness of the countermeasures. Finally, Section VII concludes the paper.

II. Related Works

Brands and Chaum [1] proposed the first distance bounding protocol. Their solution is based on a rapid-bit exchange phase, in which one party sends single challenge bits to the other party, and the latter sends back single response bits. Many distance bounding protocols have been proposed [2], [6], enjoying different properties and defending against different adversary models.

Clulow et al. [4] first gave evidence of the non-ideality of all these protocols. They presented a collection of low-level attacks, that leverage the way that protocols transmits messages on the PHY channel. Among others, they presented "early detection" and "late commit" attacks, based on the fact that bits are modulated as signals of non-zero time duration. Poturalski et al. [8] studied the impact of these attacks on a standard PHY protocol: IEEE 802.15.4a UWB [7]. They showed that these attacks can have a big impact (251m of reduction), which can be mitigated (down to 10–12m of reduction) with proper countermeasures.

Čapkun and Hubaux [13] proposed the first range-based secure positioning method, called *verifiable multilateration*. In verifiable multilateration, the distances between the device and a set of anchors (with known positions) are measured by means of independent distance bounding protocols. The position of the device is then determined by multilateration and accepted as authentic only if it lies inside the polygon formed by the anchors. This prevents an adversary to cause the measurement of a false position. Verifiable multilateration makes the assumption that the employed distance bounding is ideal, meaning that a reduction attack is impossible. In this paper, we consider a non-ideal distance bounding protocol, against which a reduction is, to some extent, possible. We investigate the impact that such attacks have on the security of verifiable multilateration.

III. Basic Concepts

Distance bounding protocols involve the precise measurement of the round-trip time (T_{rtt}) between a challenge message and a response message both carrying cryptographic material. The round-trip time is proportional to the distance (d) between the devices: $d = T_{rtt} \cdot c/2$, where c is the speed of light. The security is based on the assumption that the adversary cannot predict the cryptographic material conveyed by the messages, and thus she cannot reduce the round-trip time by producing the correct messages in advance.

The following one is a simple distance bounding protocol taken from [1]:

CMT $\mathrm{P} \longrightarrow \mathrm{V} : \mathrm{hash}(m, \mathit{open})$
CHL $\mathrm{V} \longrightarrow \mathrm{P} : a$
RSP $\mathrm{P} \longrightarrow \mathrm{V} : a \oplus m$
SGN $\mathrm{P} \longrightarrow \mathrm{V} : \mathit{open}, \mathrm{sign}_k(V, P, a, m)$

where m and *open* are random numbers generated by the prover, a is a random number generated by the verifier, and k is a shared key. With the CMT message, the prover commits to using a particular quantity m without revealing it. The CHL message is the *challenge* and the RSP message is the *response*. The verifier measures the round-trip time between them. The response is computed from the challenge through a simple bit-a-bit XOR operation. This allows the prover to respond very fast. Finally, the SGN message opens the commitment and authenticates the whole communication. This protocol is effective against two adversary models: an *external adversary* trying to reduce the round-trip time, and a *dishonest prover* trying the respond in advance. The challenge and the response are transmitted by means of specialized ultra-wideband (UWB) protocols, which allow the receiver to precisely measure the time of arrival of the messages.

Note that a distance bounding protocol resists against reduction attacks, but it gives absolutely no assurances against *enlargement attacks*. For example, a dishonest prover can achieve an enlargement on the round-trip time by simply delaying the transmission of the response bits. Taponecco et al. [12] showed that, in case of external adversary and IEEE 802.15.4a UWB PHY layer, an enlargement attack is feasible but not always controllable. In this paper, we do not focus on a particular PHY layer and we take into account also internal adversaries. Thus, we make the conservative hypothesis that enlargement attacks are possible and controllable.

A. Verifiable multilateration

Verifiable multilateration is a secure positioning system proposed by [13], and based on distance bounding. The distances between the prover and $N \geq 3$ verifiers are measured by means of N independent distance bounding executions, and the position is determined by multilateration. Then, the system performs an additional security test, named *in-polygon test*. The in-polygon test verifies that the measured position is inside the polygon formed by the involved verifiers (*verifiable polygon*). If this is true, the system accepts the measured position as trusted, meaning that no spoofing attack has taken place. In fact, in order to falsify the position, the adversary should have performed at least a distance reduction, which is impossible if we consider the distance bounding protocol to be ideal. Figure 1 exemplifies the concept.

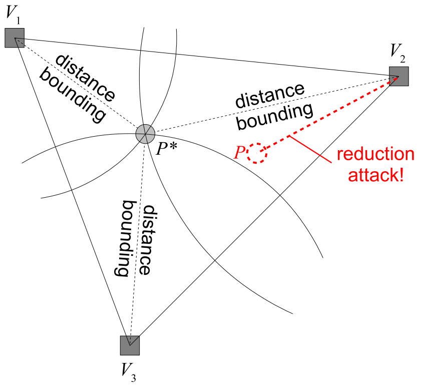

Fig. 1. Spoofing attack against verifiable multilateration. V_1, V_2, V_3 are the verifiers. P is the false position that the adversary wants to spoof, whereas P^* is its real position. In this scenario, the adversary has to reduce the distance measured between P and V_2, so the attack cannot succeed.

Like the underlying distance bounding protocol, verifiable multilateration is effective against both an external adversary and a dishonest prover.

IV. PHY-Layer Reduction Attacks

Real-life distance bounding protocols are vulnerable to PHY-layer reduction attacks. These attacks leverage the fact that bits are transmitted under the form of electro-magnetic signals of non-zero duration. The adversary uses these time latencies to anticipate the determination of an incoming bit, or to delay the decision about which bit to transmit. These techniques were originally proposed by [4], and studied against UWB protocols by [8]. They are called respectively *early detection* and *late commit*.

In early detection technique (Figure 2) a malicious receiver tries to infer the logic value of a received bit in advance, i.e. before having received it entirely.

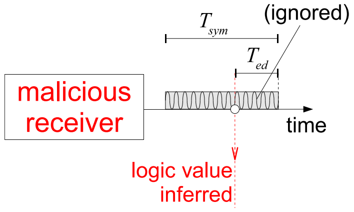

Fig. 2. Early detection technique.

In practice, she determines the logic value from the first part of the symbol only, and she ignores the second part. If $T_{sym} > 0$ is the symbol duration, we call *early detection time gain* ($T_{ed} \leq T_{sym}$) the part of the symbol duration that the receiver ignores. The larger it is, the more effective the technique. On the other hand, the time gain cannot be too large, because the probability of decoding a wrong logic value grows too. If the adversary decodes a wrong value, the protocol will fail and the distance measurement will be rejected by the verifier.

In late commit technique (Figure 3) a malicious transmitter starts transmitting a bit before having decided the logic value it will assume.

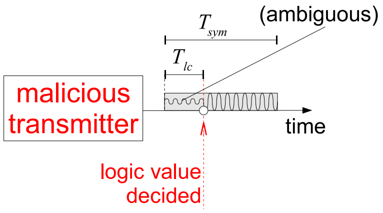

Fig. 3. Late commit technique.

She obtains this by transmitting an ambiguous signal in the first part of the symbol, and then, once the logic value is decided, she transmits the right one in the second part, possibly with more power to "overwrite" the precedent ambiguous value. We call *late commit time gain* ($T_{lc} \leq T_{sym}$) the part of the symbol duration that the transmitter sends with an ambiguous value. The more it is, the more effective the technique. On the other hand, the time gain cannot be too large, because the probability that the victim receiver decodes a wrong logic value grows too.

These basic techniques can be applied separately or together. When they are applied together, they sum their effects so as to obtain a more pronounced distance reduction. Figure 4 shows how a dishonest prover can employ both techniques.

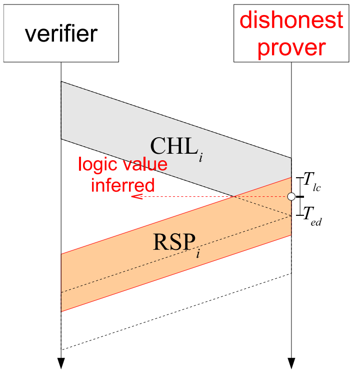

Fig. 4. PHY-layer reduction attack by a dishonest prover. The transmission of a single challenge bit (CHL_i) and a single response bit (RSP_i) are represented. The dashed lines represent the timing scheme that would take place in case the prover was honest.

The dishonest prover uses early detection to determine in advance the logic value of each challenge bit. Contemporaneously, it starts transmitting the correspondent response bit, using late commit. The obtained distance reduction is thus: $(T_{ed} + T_{lc}) \cdot c/2$. Early detection and late commit can be used as well by an external adversary, by establishing a *time-gaining relay link* between the prover and the verifier [4].

A. Mitigating PHY-layer attacks

Since the time gains obtainable with these techniques are bounded by the symbol duration, their effect can be mitigated by reducing it [4]. Ideally, a PHY layer with zero-length symbols is immune to both early detection and late commit. Such a protocol is physically infeasible, since it should transmit each bit with a finite energy in zero time, thus with an infinite power. In the practical case, the peak transmission power is strictly limited by the law [5].

A possibility is to modify the honest receivers to perform early detection. In this way a dishonest prover has less space to perform a "malicious" early detection. This countermeasure is equivalent to reducing the symbol duration, since only the first part of the signal is used to decode the symbol, and the last part is discarded. The countermeasures based on small symbol durations have the general drawback that they significantly increase the bit error rate. Thus, error-tolerant distance bounding protocols must be used [6]. In any case, such countermeasures can only mitigate the possibility of an attack, since it is not possible to reduce the symbol duration to zero.

Another common countermeasure is to use *analog circuitry* to implement the challenge-response part of the protocol [10]. In this way, the receiver is able to respond very fast, thus leaving very short time for an adversary to perform early detection or late commit. Though reduction attacks are still possible in theory (an adversary can still build quicker analog circuitry), this countermeasure can reach a close-to-ideal security. However, the analog circuitry is in general extremely sensitive to noisy wireless channels. In addition, implementing a digital protocol by means of analog circuitry could not be possible or cost-effective. Using off-the-shelf digital modules is cheaper and guarantees improved compatibility with existing standards, e.g. IEEE 802.15.4a UWB [7].

V. Security Analysis

We assume a *non-ideal* distance bounding protocol, meaning that an adversary is able, to a certain extent, to obtain a distance reduction by leveraging PHY-layer techniques. We measure the non-ideality of a distance bounding protocol with the concept of *reducible distance* (d_r), which is the largest distance reduction that a given adversary model can obtain with her most effective attack. An ideal distance bounding has a reducible distance equal to zero, which cannot be achieved in practice.

Note that the reducible distance is specific for the adversary model. So a distance bounding protocol has in general two different values for the reducible distance, one in case of external adversary, one in case of dishonest prover. From now on we will use the concept of reducible distance without specifying which adversary model. The same security analysis fits in both cases of external adversary and dishonest prover. We will say a "d_r-ideal distance bounding" to refer to a distance bounding protocol with reducible distance equal to d_r.

The concept of reducible distance is represented in an effective way by Figure 5.

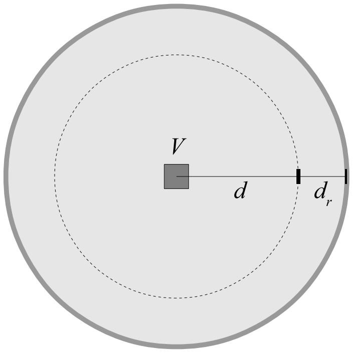

Fig. 5. Security assurance of a d_r-ideal distance bounding. Without attack, the prover is somewhere on the dashed circumference. With attack, the prover is somewhere inside the gray circle.

Let us suppose that the distance bounding protocol is executed without errors and gives the distance d as outcome. In an honest scenario, the prover will be exactly at a distance d from the verifier. In the presence of an adversary, the real distance could be shorter than d, meaning that an enlargement attack has taken place. However, if the distance bounding is not ideal, the real distance could be also longer than d, meaning that a reduction attack has taken place. In every case, we can assume that the real position of the prover is somewhere inside a circle which spans all the points whose distance from the verifier is less than or equal to $d + d_r$ (cfr. Figure 5).

A. Impact on verifiable multilateration

In the presence of a non-ideal distance bounding, a position accepted as trusted by verifiable multilateration could have been falsified to some extent. It is possible to define an *uncertainty area* (U) around the measured position (Figure 6), that surely contains the true position in case of spoofing attack.

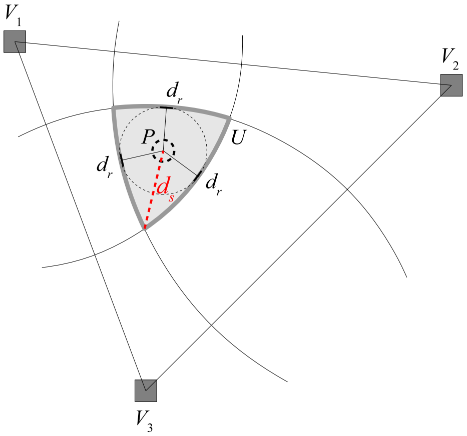

Fig. 6. Uncertainty area in verifiable multilateration with d_r-ideal distance bounding.

If we indicate by d_i the distance measured from the verifier V_i, the uncertainty area can be computed as the intersection of the circles with centers V_i and radii $d_i + d_r$. Even in case of attack, the true position cannot be outside the uncertainty area, since the adversary should have performed at least a reduction attack beyond the reducible distance. The shape and the size of the uncertainty area depends (i) on the reducible distance, (ii) on the number and positions of the verifiers, and (iii) on the measured position. Different positions inside the verifiable polygon exhibit different uncertainty areas. By geometry, the uncertainty area always contains the circle centered in P and with radius d_r. The size of the uncertainty area is thus lower-bounded by:

$$\text{size}(U) \geq d_r^2 \cdot \pi. \tag{1}$$

The size of the uncertainty area measures the power of the adversary. Another useful metric to express the adversary's capabilities is the *spoofable distance* (d_s). The spoofable distance is the maximum distance that the adversary could have spoofed. It is the distance from the measured position to the farthest point of the uncertainty area (cfr. Figure 6). By geometry, the spoofable distance is lower-bounded by the reducible distance:

$$d_s \geq d_r. \tag{2}$$

A smaller uncertainty area and a shorter spoofable distance are better for security. The best condition occurs in two special cases: when a verifier coincides exactly with the measured position (Figure 7a); and when a large number of verifiers surrounds it (Figure 7b).

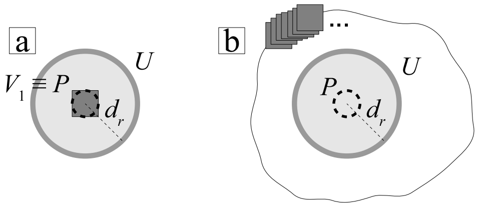

Fig. 7. Conditions under which the uncertainty area and the spoofable distance reach the minimum.

Only in these conditions we can reach the maximal assurances: $\text{size}(U) = d_r^2 \cdot \pi$ and $d_s = d_r$. These conditions are rarely reached in the practical case, so the spoofable distance will be always greater than the reducible distance.

Like the uncertainty area, the spoofable distance varies with the measured position. However, its maximum value depends only on the shape of the verifiable polygon, as the following theorem states.

Theorem 1. *In a verifiable multilateration system based on a d_r-ideal distance bounding, the spoofing distance is upper-bounded by:*

$$d_s \leq \sqrt{d_r^2 + d_r \cdot L_{max}}, \tag{3}$$

where L_{max} is the longest side of the verifiable polygon.

Because of space constraints, we do not include the complete proof, which is available in [9]. Intuitively, the upper bound stems from the fact that the spoofable distance exhibits the trend exemplified by Figure 8.

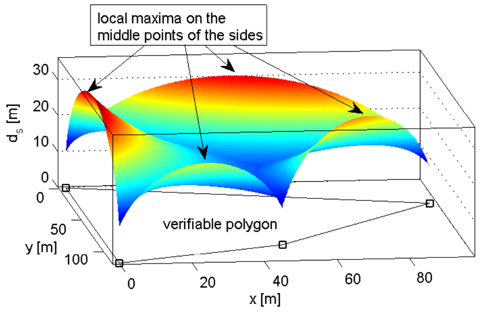

Fig. 8. Spoofable distance with respect to the measured position's coordinates with an example verifiable polygon. The reducible distance is 10m.

As we can see from the figure, the spoofable distance has three local maxima on the middle points of each side of the verifiable polygon. This is true for every number of sides and for every value of the reducible distance. Each local maximum grows with the correspondent side length. The global maximum is on the middle point of the longest side (L_{max}), and its value can be computed analytically:

$$\max_{x,y} d_s = \sqrt{d_r^2 + d_r \cdot L_{max}}. \tag{4}$$

Theorem 1 suggests that a verifiable polygon with shorter sides is more convenient, because it limits the spoofable distance inside it.

B. Countermeasures

In the presence of a non-ideal distance bounding, it is impossible to avoid spoofing attacks. Indeed, as Equation 2 implies, the spoofable distance cannot be zero. A good practice is thus to discard those positions in which the spoofable distance is beyond a given threshold (*spoofable distance threshold*, $\overline{d_s}$). This causes to leave some zones of the verifiable polygon uncovered (typically the neighborhoods of the sides' middle points).

Another good practice is to place the verifiers in such a way to avoid verifiable polygons with long sides. This helps limiting the spoofable distance, as stated by Theorem 1. A way to do this is to place the verifiers to form *regular* verifiable polygons, for example a grid of regular triangles. In fact, regular polygons maximize the covered area given a certain side length.

The optimal placement scheme uses regular triangles with side (L) equal to the communication range of the verifiers (R_{tx}). Unfortunately, this scheme (proposed by [13]) is not suitable in our case. In fact, if we impose a spoofable distance threshold, the zones near the sides' middle points could be left uncovered. A solution is to use the same scheme but with smaller triangles ($L \le R_{tx}$), in such a way that the uncovered zones are covered by the verifiers of the adjacent triangles (Figure 9). Note that the uncovered zone is now covered by a rhomboidal verifiable polygon (V_1, V_2, V_3, V_4), at the center of which the spoofable distance is low. As a consequence, the spoofable distance threshold is now respected.

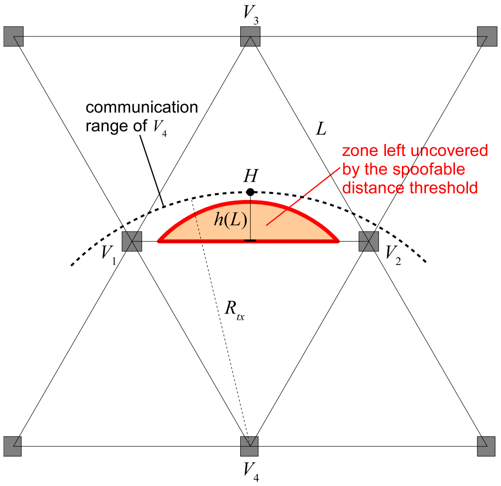
V_3
communication range of V_4
L
zone left uncovered by the spoofable distance threshold
H
$h(L)$
V_1
V_2
R_{tx}
V_4

Fig. 9. Regular placement scheme with spoofable distance threshold. The zones that would be left uncovered are covered by the verifiers of the adjacent triangles (V_4).

Given the reducible distance (d_r) and the spoofable distance threshold ($\overline{d_s}$), it is possible to determine analytically the optimal side length (L_{opt}) to obtain this:

$$L_{opt} = \max L, \text{ such that:}$$

$$L \le R_{tx} \tag{5}$$

$$\sqrt{[h(L)]^2 + d_r^2 + 2d_r\sqrt{[h(L)]^2 + \frac{L^2}{4}}} - h(L) \le \overline{d_s} \tag{6}$$

$$h(L) = R_{tx} - \frac{\sqrt{3}}{2}L \tag{7}$$

where $h(L)$ is the height of the top point H which depends on L (cfr. Figure 9). Equation 6 imposes that the spoofable distance at H is within the threshold (we skip the proof for the sake of brevity). With a grid of regular triangles with side L_{opt}, we reach the optimal coverage without uncovered zones, and we assure a bound on the spoofable distance at the same time.

VI. Experimental Results

We conducted a series of simulations to estimate the vulnerability of a verifiable multilateration system with a non-ideal distance bounding protocol. For each experiment, we placed a set of verifiers at random positions on a $200m \times 200m$ map, and a prover at a random position inside the verifiable polygon. Figures 10 and 11 show respectively the average size of the uncertainty area and the average spoofable distance with respect to the reducible distance and the number of verifiers. It can be noted that the security level considerably degrades with the non-ideality of the distance bounding, and it improves only marginally as the number of verifiers grows. In particular, the average spoofable distance is roughly twice the reducible distance even with 6 verifiers. A non-ideality of 10 meters, typical in IEEE 802.15.4a UWB [8], can leave space for a

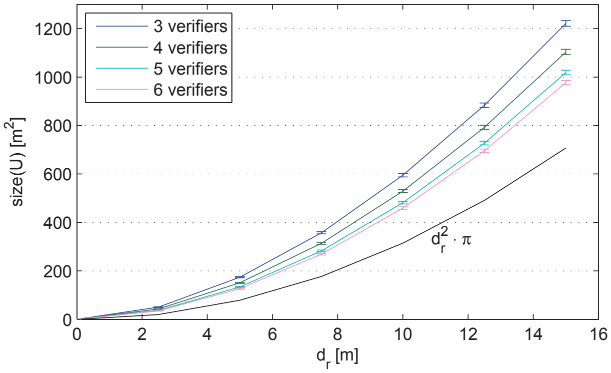

Fig. 10. Average size of the uncertainty area wrt the reducible distance and the number of verifiers. Each point stems from 1000 experiments. 95%-confidence intervals are displayed in error bars.

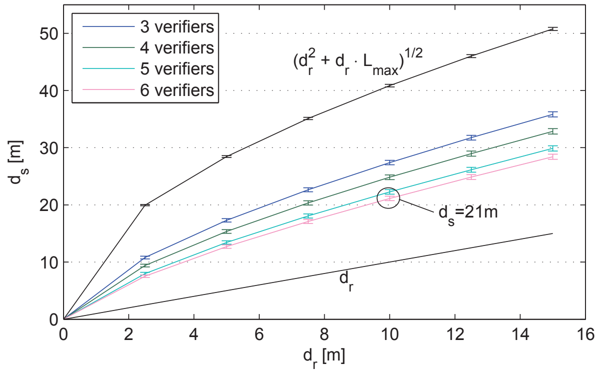

Fig. 11. Average spoofable distance wrt the reducible distance and the number of verifiers.

position spoofing of 21 meters. A threshold on the spoofable distance and a regular placement scheme limiting the side length are possible countermeasures.

We estimated the impact of these countermeasures on the coverage of verifiable multilateration. For each experiment, we placed a number of verifiers on a squared map, both at random or following the regular scheme described in Subsection V-B. Figure 12 shows the average number of verifiers needed to cover a squared area. We can see that the non-ideality of the distance bounding protocol increases the number of needed verifiers in case of random placement. On the other hand, a regular placement scheme dramatically reduces the number of needed verifiers, even in case of non-ideal distance bounding. Also, the impact of non-ideality is zero (with $d_r = 5m$) or negligible (with $d_r = 10m$) when a regular placement is used.

VII. Conclusions

In this paper, we performed a first analysis of the impact of non-ideal distance bounding protocol on the reliability of secure positioning techniques. We evaluated such an impact in terms of two metrics: namely the uncertainty area and the spoofable distance. We showed that a reduction of 10 meters, possible against the standard IEEE 802.15.4a UWB PHY, can cause a position spoofing of 21 meters. We also proposed two countermeasures that can mitigate the problem: (i) to discard the positions which are falsifiable over a certain limit, and (ii) to place the anchors following a particular scheme having certain geometrical properties. We evaluated these countermeasures by simulations.

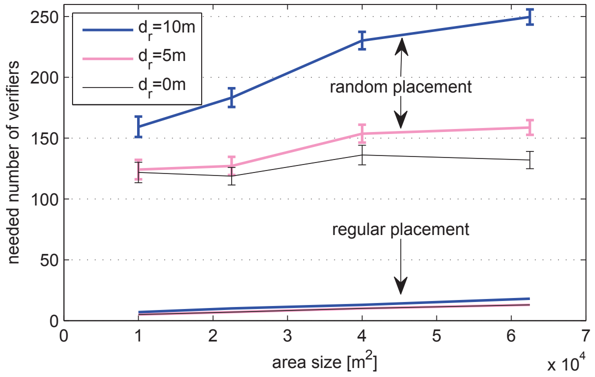

Fig. 12. Average number of needed verifiers to cover 90% of a squared area. The transmission range of the verifiers is 100 meters, and the spoofable distance threshold is 15 meters. Each point stems from 100 experiments. 95%-confidence intervals are displayed in error bars. (The curves of the regular placement for $d_r = 0m$ and $d_r = 5m$ coincide.)

References

[1] S. Brands and D. Chaum. Distance bounding protocols. In *Proceedings of EUROCRYPT'93*, pages 344–359, 1993.

[2] L. Bussard and W. Bagga. Distance-bounding proof of knowledge to avoid real-time attacks. *Proceedings of IFIP/SEC'05*, pages 223–238, 2005.

[3] J. T. Chiang, J. J. Haas, J. Choi, and Y.-C. Hu. Secure location verification using simultaneous multilateration. *IEEE Transactions on Wireless Communications*, 11(2):584–591, 2012.

[4] J. Clulow, G. P. Hancke, M. G. Kuhn, and T. Moore. So near and yet so far: Distance-bounding attacks in wireless networks. In *Proceedings of ESAS'06*, pages 83–97, 2006.

[5] FCC. First report and order, revision of part 15 of the commissions rules regarding ultra-wideband transmission systems. Technical report, FCC, Feb 2004.

[6] G. P. Hancke and M. G. Kuhn. An RFID distance bounding protocol. In *Proceedings of SecureComm'05*, pages 67–73, 2005.

[7] IEEE Computer Society. IEEE Std 802.15.4a-2007 (Amendment 1: Add Alternate PHYs), 2007.

[8] M. Poturalski, M. Flury, P. Papadimitrios, J.-P. Hubaux, and J.-Y. Le Boudec. Distance bounding with IEEE 802.15.4a: Attacks and countermeasures. *IEEE Transactions on Wireless Communications*, 10(4):1334–1344, 2011.

[9] Proof of Theorem 1. www.iet.unipi.it/g.dini/download/pubs/TR-2015-02-PD-appendix.pdf.

[10] K. B. Rasmussen and S. Čapkun. Realization of RF distance bounding. In *Proceedings of USENIX'10 Security Symposium*, pages 389–402, 2010.

[11] N. Sastry, U. Shankar, and D. Wagner. Secure verification of location claims. In *Proceedings of WiSe'03*, pages 1–10. ACM, 2003.

[12] L. Taponecco, P. Perazzo, A. D'Amico, and G. Dini. On the feasibility of overshadow enlargement attack on IEEE 802.15.4a distance bounding. *IEEE Communications Letters*, 18(2):257–260, 2014.

[13] S. Čapkun and J.-P. Hubaux. Secure positioning in wireless networks. *IEEE Journal on Selected Areas in Communications*, 24(2):221–232, 2006.

[14] A. Vora and M. Nesterenko. Secure location verification using radio broadcast. *IEEE Transactions on Dependable and Secure Computing*, 3(4):377–385, 2006.