ID-based Secure Distance Bounding and Localization

Nils Ole Tippenhauer and Srdjan Čapkun

Department of Computer Science
ETH Zürich
8092 Zürich, Switzerland
{tinils, capkuns}@inf.ethz.ch

Abstract. In this paper, we propose a novel ID-based secure distance bounding protocol. Unlike traditional secure distance measurement protocols, our protocol is based on standard insecure distance measurement as elemental building block, and enables the implementation of secure distance bounding using commercial off-the-shelf (COTS) ranging devices. We use the proposed protocol to implement secure radio frequency (RF) Time-of-Arrival (ToA) distance measurements on an ultra-wideband (UWB) ranging platform. Based on this, we implement Verifiable Multilateration — a secure localization scheme that enables the computation of a correct device location in the presence of an adversary. To the best of our knowledge, this is the first implementation of an RF ToA secure localization system.

1 Introduction

A number of secure distance bounding ([1–4]) and secure localization protocols ([5–11]) have been proposed in the recent years. Secure distance bounding protocols were first described in [1] to protect against *mafia fraud* attacks [12]. Secure distance bounds can be derived in scenarios in which the measurement target B is either trusted [9], or untrusted [1] by the measuring entity A. In both cases, a third entity (the attacker $\mathcal{M}$) cannot shorten the measured distance, but prolong it by delaying the sent messages. The established distance bound can be used in many applications, including the prevention of relay (wormhole) attacks [2] and physical proximity verification (e.g., for access control purposes) [13]. Using these established distance bounds, secure time-off-arrival (ToA)-based secure localization systems (e.g. [8, 9]) can be realized.

Secure localization protocols were proposed to provide trusted location information in security- and safety-critical applications like location-based access control, asset monitoring, protection of critical infrastructures, emergency and rescue, and to enable secure networking functions (i.e., location-based routing, secure data harvesting). Secure localization systems such as [5, 9], and [7] rely on ToA measurements.

One of the main problems that prevents a wider deployment of secure distance measurement protocols is the requirement that devices process messages with minimal delays — ideally instantaneously. As existing insecure commercial off-the-shelf (COTS) distance measurement platforms are not designed to provide this feature, nor the required cryptographic operations, the implementation of secure distance bounding based on these platforms would require extensive redesign of their hardware and software.

In this paper, we address this problem, and we propose a novel ID-based secure distance bounding protocol. Our protocol is based on insecure ranging as elemental building block, and enables the implementation of secure distance bounding using COTS ranging devices.

Our main contributions are as follows:

- We propose a new secure distance bounding protocol that can be implemented on available distance measurement platforms. The proposed protocol lowers the complexity of the implementation and does not require modifications of existing ranging platforms.
- We implement the proposed protocol using ultra-wideband radio frequency ranging devices, show that it enables secure and accurate distance bounding and discuss possible design choices.
- Based on our secure distance bounding implementation, we implement a Verifiable Multilateration-based secure localization protocol; we show that our implementation enables accurate and secure localization of a trusted target.
- We further show several new attacks on secure localization, specifically those that can be performed by untrusted mobile targets, and propose solutions to these attacks.

To the best of our knowledge, this paper presents the first implementation of an RF ToA-based secure localization system.

The rest of the paper is organized as follows. Background on secure distance bounding protocols and the used hardware is given in Section 2. Our secure distance bounding protocol is motivated and described in Section 3. Section 4 discusses our implementation of a secure localization system. Related work is described in Section 5. We conclude the paper in Section 6.

2 Background

We will now introduce secure distance bounding in more detail and then present the COTS hardware platform for the secure distance bounding and secure localization implementation.

2.1 Secure Distance Bounding

Secure distance bounding aims at detecting attacks on distance bound measurements in scenarios in which the target devices are either trusted or untrusted. If we assume that the ranging target B cannot be compromised by an attacker and if the measuring node A trusts B to follow the protocol honestly, a *trusting* *distance bounding* (tDB) protocol such as the authenticated ranging protocol proposed in [9] can be used by A to determine the upper bound on its distance to B. If A does not trust B, it has to use an *untrusting distance bounding* (uDB) protocol, e.g. [1], to compute an upper bound on the distance to B. In both cases, the goal of A is to obtain an upper bound on the distance to B. Note that in both cases, the attacker is always able to delay messages between A and B and thus enlarge their measured distance by jamming/replaying or overshadowing the signals, but she cannot reduce the measured distance since the attacker cannot advance the RF signals between A and B.

Please note that we will use the notion of untrusting and trusting distance bounding (uDB/tDB) throughout this paper to refer to protocols aiming to establish distance bounds with an untrusted or trusted target B. The notion of secure distance bounding will be used to summarize both variants.

We now briefly describe the original uDB protocol by Brands and Chaum [1]; in this protocol (shown in Figure 1), an untrusted target B starts by committing to a message m of size b bits and by sending this commitment to A. A then generates b secret challenge bits $|\alpha_1 \dots \alpha_b|$, after which both parties perform b rounds of rapid bit exchange (RBE). In each round, A sends the current challenge α_i, B then computes $\beta_i = \alpha_i \oplus m_i$ and immediately sends β_i to A. After b rounds, B concatenates the received challenges into a bit string m, opens the initial commit to A and sends a signed m to A. A now verifies the commitment and the signature of m. If both verifications are successful, A computes the round-trip time RTT_i for each challenge and response. The distance bounding is considered successful if each distance $d_i = \frac{\mathrm{RTT}_i \cdot v}{2}$ was shorter than the maximal possible distance between A and B (v is the signal propagation speed, approximately speed of light). This maximal distance could for example be determined by A and B's power ranges.

In the case of trusting distance bounding (authenticated ranging in [9]), A trusts that B will correctly execute the protocol and will not cheat in the ranging process. As a consequence, the reply by B is not required instantaneously anymore; instead A trusts B to process the challenge in a constant or known time, after which B will send the reply. A can then compute the distance by subtracting the known processing delays from the measured RTT.

Theoretically, the only way that an attacker can compromise secure distance bounding protocols to reduce the measured distance is to either guess all the challenge bits sent by A or all the replies sent by B in the RBE phase. The probability of a successful attack therefore depends on the amount of rounds of RBE b and is equal to 2^{-b}.

2.2 The MSSI UWB Ranging System

The ranging devices by MSSI [14] operate in the frequency range of 6.1-6.6 GHz both for communication and for ToA ranging measurements. Their serial interface currently only provides a very limited set of operations, of which only one is of special interest for us: the *ranging* command that allows one device to

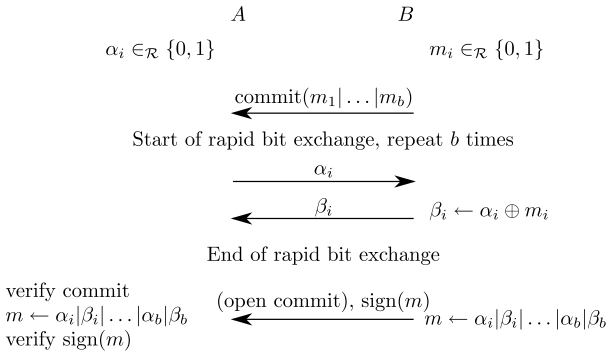

A
B
$\alpha_i \in_{\mathcal{R}} \{0,1\}$
$m_i \in_{\mathcal{R}} \{0,1\}$
commit($m_1|\ldots|m_b$)
Start of rapid bit exchange, repeat b times
α_i
β_i
$\beta_i \leftarrow \alpha_i \oplus m_i$
End of rapid bit exchange
verify commit
$m \leftarrow \alpha_i|\beta_i|\ldots|\alpha_b|\beta_b$
verify sign(m)
(open commit), sign(m)
$m \leftarrow \alpha_i|\beta_i|\ldots|\alpha_b|\beta_b$

Fig. 1. Brands and Chaum's untrusting distance bounding protocol [1].

measure its distance to another device. Every radio has a unique address which consists of an 8 bit subnet number and of an 8 bit unit identifier, which can be changed fast via the serial interface. To perform a ranging operation, device A broadcasts a request containing the ID of a device that it wants to range (e.g., B's ID). Upon reception of this message, B processes the message in constant time and sends back a reply message. A measures the RTT between transmitting the request and receiving the reply, and computes its distance to B.

In the request messages for the distance measurement, no additional data can be transmitted from A to B, which prevents the transmission of a challenge, needed in all secure distance bounding protocols. This means that none of the existing secure distance bounding protocols can be implemented on this platform. This limitation is common in insecure ranging devices, which motivated us to propose a protocol that enables secure distance bounding with such commercial-off-the-shelf devices.

3 The ID-Based Secure Distance Bounding Protocol

In this section, we present our ID-based secure distance bounding protocol. This protocol can be implemented on existing commercial off-the-shelf ranging platforms like the one of MSSI, as described in Section 2.2. We then discuss its security, performance, present our implementation, and propose further performance improvements.

3.1 ID-Based Secure Distance Bounding

Our ID-based secure distance bounding protocol enables devices which *cannot* add binary challenges to the ranging messages and *cannot* compute XOR ($\oplus$) operations on the challenge to still perform secure ranging. The only requirement for the ranging devices is that they can be instructed to change their IDs. We

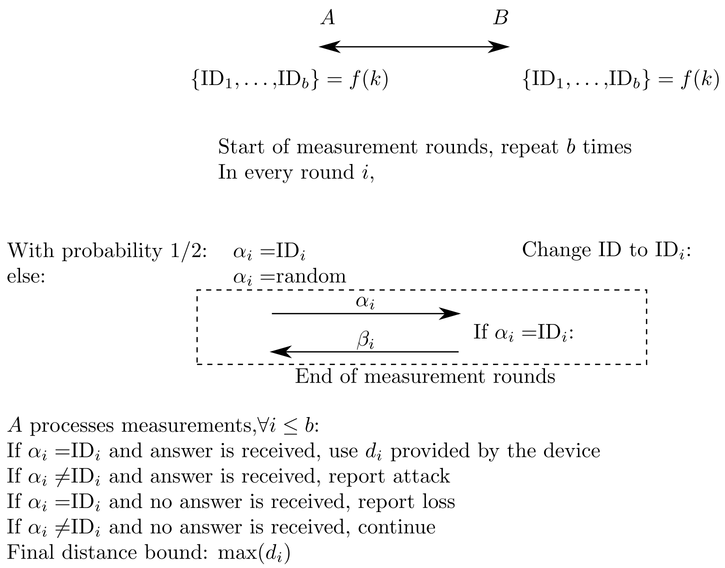
A
B
$\{\text{ID}_1, \ldots, \text{ID}_b\} = f(k)$
$\{\text{ID}_1, \ldots, \text{ID}_b\} = f(k)$
α_i
β_i

Start of measurement rounds, repeat b times
In every round i,

With probability 1/2: $\alpha_i = \text{ID}_i$
else: $\alpha_i = \text{random}$

Change ID to ID_i:

If $\alpha_i = \text{ID}_i$:

End of measurement rounds

A processes measurements,$\forall i \leq b$:
If $\alpha_i = \text{ID}_i$ and answer is received, use d_i provided by the device
If $\alpha_i \neq \text{ID}_i$ and answer is received, report attack
If $\alpha_i = \text{ID}_i$ and no answer is received, report loss
If $\alpha_i \neq \text{ID}_i$ and no answer is received, continue
Final distance bound: $\max(d_i)$

Fig. 2. ID-based secure distance bounding protocol: Initial setup, the measurement rounds and postprocessing. The steps in the dashed box are executed on ranging devices, requiring only standard ranging commands.

assume that A and B each control one ranging device (in the case of MSSI devices via their serial interfaces) can communicate directly (e.g., using their IEEE 802.11 interfaces) and that they share a secret key or hold each other's valid public keys before the start of the protocol. The key establishment process itself is outside the scope of this work.

The ID-based secure distance bounding protocol is executed as follows (Figure 2). In the protocol initialization phase, A and B agree on a shared key k, from which they derive a secret ID sequence $\text{ID}_1, \ldots, \text{ID}_b$. A and B then run b rounds of the ID-based secure distance bounding primitive. In the ith round, A sends a ranging request to ID_i with probability 1/2, else it ranges a random ID. An honest B will reply only to the ranging requests sent to ID_i, the ID corresponding to the ith protocol round. After b rounds, the distance bound is computed by taking the maximum of all valid measured distances.

Unlike B, an external attacker $\mathcal{M}$ can only guess which ID to reply to, since she does not know the ID sequence shared between A and B. The attacker will therefore be able to shorten the range between A and B only with probability 1/2 in each round; in case that the attacker answers to the random ID, A will not accept the range and will detect the attack. In addition, an untrusted B will only be able to shorten its distance to A with probability 1/2 by sending an early reply message because it does not know if its current ID_i or a random ID will be queried.

In summary, in every round $i \leq b$, A can distinguish between the following cases:

1. A ranges ID_i and receives a reply from ID_i. A concludes that the distance computed by this measurement is a valid upper bound on B's distance.
2. A ranges ID_i and receives no reply. A concludes that a transmission error or an attack could be the cause. The handling of this event depends on the quality of the communication channel; if no signal losses are to be expected, we can assume an attack.
3. A ranges a random ID and receives a reply from this ID. A concludes that an attacker replied, as no honest B would reply to a random ID.
4. A ranges a random ID and receives a reply from ID_i. A concludes that a dishonest B tried to shorten the distance by sending an early reply.
5. A ranges a random ID and no reply is received. A concludes that no attack was attempted this round.

After b rounds, the distance bound is computed by taking the maximum of all valid measured distances. Depending on the security policy, A can decide not to accept the upper bound if it detects attempted attacks such as case 2, 3, or 4 in one (or more) rounds of the protocol.

3.2 Communication Cost

In the original Brands and Chaum's proposal, only single bits of information are transmitted between A and B in each round of the protocol. In the ID-based secure ranging protocol, ℓ-bit IDs are being transmitted in each round. From this, it might seem that the ID-based protocol incurs ℓ-times higher communication cost than Brands and Chaum's protocol. However, in existing UWB ranging systems, $\approx$ 10 byte long preambles need to be sent with each message for the receiver to recognize (i.e., synchronize to) the ranging signals of the sender. With the IDs of size $\ell = 16$ bit, ID-based secure distance bounding protocol will therefore have about 20 % higher communication overhead than the original Brands and Chaum's protocol using the same UWB message format.

The number of rounds in the rapid bit exchange depends on the chances an attacker has to cheat successfully in each individual round. We will discuss this chance in the next section and show that it is only marginally greater than in the original protocol of Brands and Chaum, therefore the number of rounds needed are almost the same. This value is determined by the size of the ID space and other implementation details as discussed in Section 3.3.

3.3 Security Analysis

In this section, we discuss the security of the ID-based secure ranging protocol and our specific implementation assuming a trusted B.

Attacker Model In our analysis we will only discuss attacks by an external attacker $\mathcal{M}$ and assume that B is honest and trusted by A to correctly follow the protocol. The goal of these two attackers are the same: to shorten the measured distance between A and B and thus to make A believe that B is closer than it really is.

We assume that $\mathcal{M}$ controls the communication channel in the sense that she can eavesdrop, jam, replay, insert and modify transmitted messages. However, the attacker cannot transmit messages at a speed higher than the speed of light. We further assume that $\mathcal{M}$ cannot obtain the secret key shared between A and B. We do not specifically address side-channel information leaks in the analysis — trusted nodes are assumed to not leak information, and malicious nodes already have access to all information which could leak. In addition, we do not consider denial of service attacks — like most wireless communications systems, denial of service attacks, e.g. through jamming, are possible. The goal of our protocol is to obtain a correct distance bound, and not to guarantee availability.

Protocol Analysis As we showed earlier, our protocol prevents external attackers and even dishonest users from sending early replies to A's challenges by randomizing the challenges. Since $\mathcal{M}$ does not know the ID sequence shared between A and B, it can only guess which ID she should reply to in order to impersonate B. The attacker will therefore be able to shorten the range between A and B only with probability 1/2 in each round; in case the attacker answers to the random ID, A will not accept the range and will detect the attack. Equally, an untrusted B will not be able to shorten its distance to B by sending an early reply message, because it does not know if the current ID_i or a random ID will be queried.

Privacy: Existing secure distance bounding protocols only authenticate the challenges after the *rapid bit exchange*. An attacker exploiting this to find her distance to B will only be detected after she obtained her distance. This would enable the attacker to easily obtain the same information as A, implications of this attack are discussed in detail in [15]. Our protocol prevents this by effectively authenticating each challenge, because these are derived from the shared secret. Therefore, the attacker cannot send her own ranging messages to B before A sends the legitimate requests.

Implementation Analysis Although we have shown the resistance of our protocol to attacks from external attackers, different implementations of secure distance bounding protocols can be vulnerable to physical layer attacks [16]. We will now describe three possible attacks on our implementation of trusting distance bounding, discuss their effectiveness and how to prevent them. The first attack concerns packet level latencies, whereas the other two are based on scanning the space of possible ID values. If we do not trust B, more attacks by a malicious B' are possible.

External early-send late-commit attacks: As Clulow et. al. pointed out in [16], a malicious B can exploit packet level latencies to his advantage. When

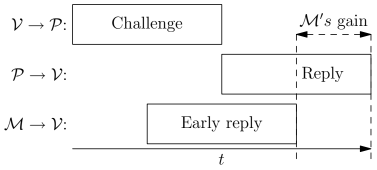
$\mathcal{V} \to \mathcal{P}$:
Challenge
$\mathcal{M}$'s gain
$\mathcal{P} \to \mathcal{V}$:
Reply
$\mathcal{M} \to \mathcal{V}$:
Early reply
t

Fig. 3. External early-send late-commit attack by $\mathcal{M}$: While B is still receiving the challenge, $\mathcal{M}$ is already sending a reply. If B reacts to the challenge, $\mathcal{M}$ completes its early reply. Otherwise, $\mathcal{M}$ interrupts its early reply, making the attack harder to detect. If the attack was successful, $\mathcal{M}$ shortened the distance by the time its reply started earlier.

using the ID-based secure distance bounding, the reply of B carries basically one bit of information (to reply or not), this enables early-send late-commit attacks by a malicious B'. In trusting distance bounding, A trusts B, but a similar attack is possible by $\mathcal{M}$. When using MSSI's devices, which use packets with a length of 56 μs, $\mathcal{M}$ could start a reply early (replying to A's challenge), but only finishing the reply (i.e., completing it) if it observes the answer of B. If $\mathcal{M}$ does not receive the answer from B, she knows that A sent the challenge to a random ID and she will stop the early response, as displayed in Figure 3. This way, the attacker could shorten the distance up to the length of one packet, which is 56μs for our devices.

To detect this attack, A has to listen for incomplete packet transmissions. If A is able to detect a single UWB signal on the channel, the early-send late-commit attack is defeated, and all that remains is the same attack on the signal level, only yielding a gain of half the signal length as described in [16].

Preemptive challenge attack: Our protocol relies on the fact that the current IDs of A and B are unknown to $\mathcal{M}$ until they send messages. This implies that we have to make sure that there is no *efficient* way to query the current ID from one of the two entities. The external attacker $\mathcal{M}$ could try to send out distance bounding challenges to random addresses, trying to hit the right ID of B. The chance for this is $2^{-\ell}$, in our case 2^{-16}. As the attacker has to use the normal message format with messages of length 50μs, the maximum frequency with which it can query the devices is $f_q = \frac{1}{50\mu s} = \frac{20}{ms}$. Hence, the chances of success for this attack depend on the delay between B changing its ID and A's distance measurement. In our implementation, this takes less than 20 ms, which means that in the worst case, the attacker is able to query 400 ($< 2^9$) IDs between two rounds of the protocol. 2^{-7} is therefore an upper bound for the attacker's success chance.

A generalized formula for $\mathcal{M}$'s gain using the preemptive challenge attack is the following: given an ID space of size 2^ℓ, a round length t_r, and $\mathcal{M}$'s ID scanning ratio $f_s = \frac{\text{IDs scanned}}{\text{time}}$, the gain is $\frac{t_r f_s}{2^\ell}$ per round, in our case $< 2^{-1} + 2^{-7}$. We conclude that the preemptive challenge attack seems *inefficient* compared to $\mathcal{M}$'s chance of simply guessing the answer with 50 % chance per round. If the devices report successful rangings to the controlling PC, both are easily detected.

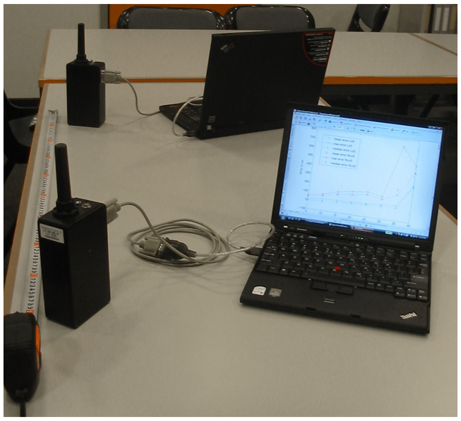

Fig. 4. The setup for the trusting distance bounding implementation.

3.4 Implementation and Measurement Results

We implemented our secure distance bounding protocol to allow authenticated ranging (assuming a trusted B) using two UWB ranging devices controlled by PCs over serial connections; our implementation setup is shown in Figure 4.

A client program running on a PC initiates a trusting distance bounding session and specifies the number of protocol rounds. All communication between the programs besides the ranging is done over standard TCP/IP sockets, using IEEE 802.11 wireless channels. This communication consists of the initial authentication of the involved parties, secure key establishment, and the synchronization of the individual protocol rounds. For simplicity in our experiments, keys were manually preloaded in the PCs.

In our implementation, individual protocol rounds are about 20 ms long; this is mainly due to the slow serial connection to the devices. Upon reception of the signal to start the next round, B's PC sets the ranging devices ID over the serial connection. A then commands its ranging device over the serial connection to perform the ranging operation with either ID_i (in round i) or with a random ID. The results of the successful distance measurements are computed internally in the ranging devices. The controlling program on the PC queries the ranging device for results, which are provided to the PC as the message RTT (expressed in nanoseconds).

We tested the accuracy of our secure distance bounding protocol on MSSI platforms. We performed 1000 measurements in a line-of-sight (LoS) outdoor environment and 1000 in non-line-of-sight (NLoS) environment (indoor office area), for distances up to 40 meters. The results are listed in Table 1.

Compared with insecure distance bounding on our platform, the additional effort in our implementation is the following: First, the frequent changing of the device's ID requires a control program to handle the initial protocol setup and the actual ID changes. Second, instead of performing b measurements subse-

	LoS			NLoS		
d	σ	$\bar{d}-d$	d_m-d	σ	$\bar{d}-d$	d_m-d
in m	in cm	in cm	in cm	in cm	in cm	in cm
5	10.23	-5.00	9.25	8.64	40.81	57.10
10	9.60	8.25	30.65	11.54	63.61	82.10
15	9.05	17.32	36.75	19.46	105.57	132.60
20	9.66	24.41	38.95	16.37	123.23	158.35
25	9.54	31.94	48.20	14.92	148.54	177.65
30	9.97	39.30	58.50	14.41	120.06	147.15
35	9.31	44.22	65.65	253.33	240.68	722.35
40	10.23	289.99	304.40	52.78	448.13	527.37

Table 1. Secure distance bounding results of 1000 measurements: d is the correct distance between A and B, σ the standard deviation of the measurements, $\bar{d}$ the mean of the measurements and d_m the maximum value of all measurements.

quently as for insecure distance bounding, in secure distance bounding we have to split those operations in multiple rounds. In our current implementation, one measurement takes about 40 ms on average (each round will only have a measurement in 50% of the cases), while unauthenticated ranging can perform up to 16 measurements in 53 ms. The difference in runtimes is mainly due to the slow serial communication with the ranging device, and the fact that secure distance bounding requires many commands to be sent while insecure distance bounding can perform 16 rangings with a single command sent to the radio.

4 Secure Localization

Based on our secure distance bounding protocol presented in Section 3.1 and its implementation presented in Section 3.4, secure localization can be implemented using Verifiable Multilateration as proposed in [9]. In the following section we will introduce Verifiable Multilateration, present the implementation and discuss further improvements to its performance and security when localizing moving targets.

4.1 Background: Verifiable Multilateration

The goal of Verifiable Multilateration (VM) with a trusted target is to determine the correct location of B in the presence of an external adversary using secure distance bounding (untrusting distance bounding or trusting distance bounding). It consists of measurements from at least three reference points (localizers) to B's device and of subsequent computations performed by an authority. In this description, we will assume that the verification is performed with trusting distance bounding. For simplicity, we discuss the algorithm for 2-D localization. The intuition behind the VM algorithm is the following: due to the trusting distance bounding properties, the attacker can only increase the measured distance

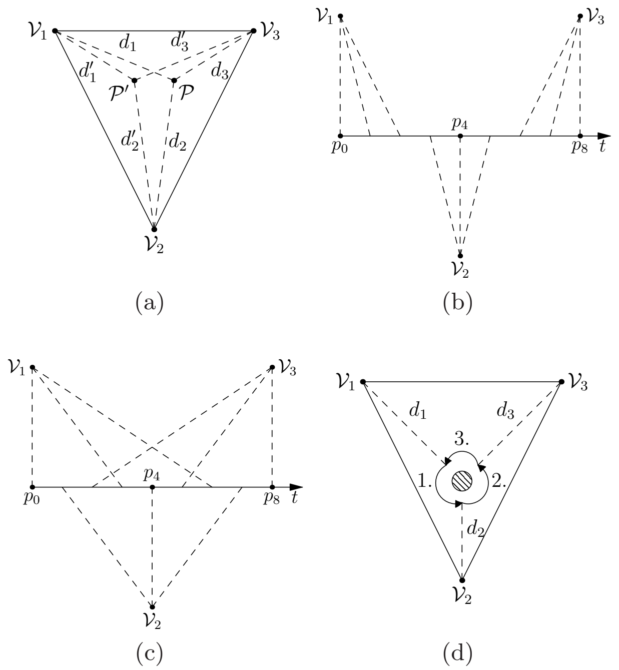

Fig. 5. (a) Verifiable Multilateration: Basic localization setup, three localizers A_1, A_2, A_3 measure the distance to B and localize it within the verification triangle. If $\mathcal{M}$ wants to influence the measurements to result in a location B', she would have to reduce at least one measured range, which she cannot due to trusting distance bounding as it prevents distance reduction attacks; (b,c) Secure localization of an object moving from p_0 to p_8: dashed lines represent sequential rangings (b) and interleaved rangings (c). In this simple example, each A executes only 3 ranging rounds. (d) Movement attack on localization: The attacker moves B changing its location between range measurements to claim a location that is otherwise for him unreachable (in this example, the shaded region located in the middle of the triangle).

between B and A. If $\mathcal{M}$ increases the measured distance to one of the As, she needs to prove that at least one of the measured distances to other As is shorter than it actually is in order to keep the position consistent, which she cannot because of the trusting distance bounding. This property holds only if the position of B is determined within the *verification triangle* formed by the As. This can be explained with a simple example: if an object is located within the triangle, and it moves to a different position within the triangle, she will certainly reduce its distance to at least one of the triangle vertices (Figure 5(a)). Verifiable Multilateration guarantees the following property: an external attacker performing a distance enlargement attack cannot trick the As into believing that a target, which is located at a location in the verification triangle, is located at some other location in the triangle. Equally, the attacker cannot trick the As into believing that a target located outside of the verification triangle is located within the triangle. Verifiable Multilateration therefore prevents attacks on localization within an area covered by the localization infrastructure (i.e., by the verification triangles). More details and a security analysis can be found in [9].

4.2 Implementation

We implemented Verifiable Multilateration as a natural extension of our secure distance bounding implementation. We assume that B is trusted in our implementation and we therefore can use authenticated ranging to determine its location. Our implementation consists of a set of three verifying MSSI ranging devices, controlled by a PC, and a target also using a ranging device. Secure localization can be initialized with a variable number of RBE rounds in each individual secure ranging. In our implementation, the resulting distances from the localization are processed by the controlling PC in Matlab [17] to display a visual representation of the position and provide statistical information. If required, the localization process itself can be executed in a loop to continuously update the location plot, providing real time location information.

4.3 Results and Further Improvements of the Aggregation Function

To evaluate our implementation, we used the accumulated squared error between the individual measured ranges and the final position $e = \sum_{i=1}^{3}(\hat{d}_i - d_i)^2$. Since secure distance bounding protocols take the maximum measured distance d_m (over all protocol rounds) as an upper bound on the distance between each A and B, high measurement variance will lead to decreased accuracy. Using the mean or other aggregation functions, however, would make secure distance bounding more vulnerable to attacks; if the attacker (e.g., by guessing a reply) shortens a distance in only one round, she could significantly affect the computed mean. The trade-off between the influence an attacker can have on the alternative aggregation function's outcome and the probability of the detection of the attack is visualized for mean and median aggregators in Figure 6(a). We define the influence i by $\tilde{d} = (1 - i)d$, with the original distance d and $\tilde{d}$ the influenced distance, assuming that the attacker is able to reply instantaneously in a successful attack.

We compared the performance of the maximum, median and mean function when aggregating a variable number of measurements. The distances between the localizers and B were in the range of 10 to 20 meters. To see the influence of the number of measurements in the secure ranging protocol, we measured these values for 1, 2,4,8,16, and 32 measurements. The results are given in Figure 6(b) and show that in secure distance bounding, more measurements do not necessarily decrease the error, as the accumulated error of the max function is 3 times higher than the error of the median and mean aggregation function if more than 10 rounds of secure distance bounding are performed each. The median performs similar to the mean, without any influence by the attacker if less than 50% of the samples are compromised. If the attacker can compromise more, then she

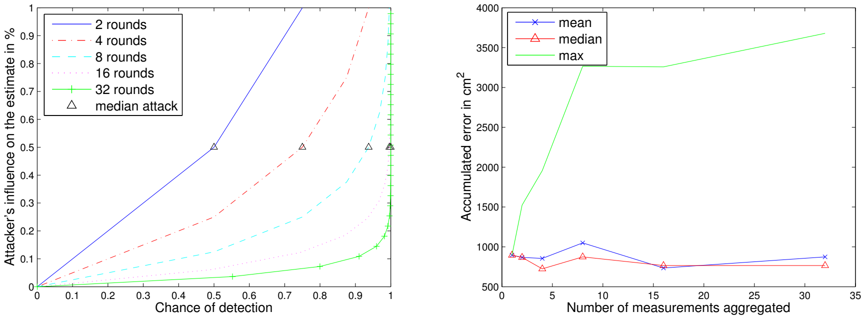

(a)

(b)

Fig. 6. (a) Attacks on the aggregation: Attacker's influence on the aggregated range as a function of the probability of attack detection, for 2, 4, 8, 16 and 32 secure distance bounding rounds. Continuous lines show values for mean aggregation, triangles denote the respective chances to fully compromise the aggregated ranging result if the median is used (half of the measurements needs to be compromised).(b) Implementation of Verifiable Multilateration: the accumulated squared error e in cm^2 for different range aggregation functions, for 1, 2, 4, 8, 16, and 32 measurements during the secure distance bounding protocol.

trivially controls the final result of the median computation, effectively halving the number of rounds from a security perspective. The max aggregation will result in a much higher error. We discuss this trade off between security and accuracy further in our technical report [18].

4.4 Performance Improvements for Moving Targets

Depending on the time that one secure distance bounding protocol run takes, the accuracy a moving target localization can suffer, as B's position can be different for each of the three secure distance boundings with the As. Figure 5(b) illustrates the localization of a target moving from p_0 to p_8 if only three rounds of measurements are performed by each A. If we assume that 10 rounds of secure distance bounding are performed each, the total duration of each localization is about 600 ms. This means that an object travelling at a speed of 10 km/h or 2.78 m/s already moved 1.5 m during the localization process.

To improve the accuracy of the sequential measurements, we modify the localization protocol. Instead of performing full runs of secure ranging between each A and B, we run rounds of single localizations. Each localization round consists of three ranging runs, one between each A and B, as illustrated in Figure 5(c). Each localization round gives B's location at a certain time. These single localizations can now be used to track B, (e.g., using a Kalman filter [19]). When a new range is measured, the error between the predicted distance and the measurement result can be computed. If this error exceeds a certain threshold, an attack can be detected. This way, the attacker would have to continuously and successfully compromise the measured distances, the probability of which is small (i.e., $\leq 2^{-b}$).

4.5 Moving Target Attacks and Countermeasures

Here, we consider the following attack on Verifiable Multilateration: in addition to controlling the communication channel, the attacker can move the trusted target B without being noticed. Based on this, the attacker can defeat Verifiable Multilateration by changing $B's$ location between the two range measurements. In the case of sequential Verifiable Multilateration, after each ranging run, the attacker can move B (e.g., closer) to A_i with which it will range next and thus violate the assumption of non-reduceable distances. This attack could be used by the attacker, for example, to claim a location in the middle of the verification triangle, which is otherwise not reachable by B. This attack is illustrated in Figure 5(d). After the first range measurement from A_1 the attacker changes B's position closer to A_2 (1. in the figure). After A_2 has completed secure distance bounding, $\mathcal{M}$ moves B closer to A_3 (step 2.). If there is another round of ranging, the attacker will then move B back to its initial position.

To prevent this attack, we randomize the ranging sequence from the As. The attacker can therefore only guess the next A with $\frac{1}{3}$ chance. Failure to predict the next A to move closer to will lead to larger distances being measured, and a resulting higher e, which will indicate to the authority that there is an attack on the localization process.

5 Related work

The first untrusting distance bounding protocol was proposed by Brands and Chaum [1]; this protocol was later applied to a wireless scenario and extended to provide mutual authentication in [4]. A noise resilient version of this mutual authentication protocol was proposed in [20]. To support more resource constrained devices like RFID tags, an alternative untrusting distance bounding protocol was proposed in [3]. The first trusting distance bounding protocol was proposed in [10]. Several protocols address the thread of *terrorist fraud* attacks, in which B partially cooperates with the attacker ([21–23]).

The first implementation of distance bounding over a wired channel was presented in [24]. Implementations of wireless distance bounding for RFID tags with very short range and low accuracy appeared in [25] and [21]. Attacks on possible implementations of secure distance bounding protocols were discussed in [26]. Our work extends this previous work by presenting a highly accurate secure distance bounding system, as demonstrated in our implementation.

A system for secure localization based on US and RF communications was proposed in [5], attacks on ultrasonic ranging systems were discussed in [27]. [7] proposes a set of techniques for secure positioning of a network of sensors based on directional antennas, with an extension in [8] to cope with the replay of navigation signals. In [10], a secure localization scheme based on hidden and mobile base stations is presented. In [11], a system for broadcast localization and time-synchronization was proposed and implemented. In [28], the authors propose a multilateration system based on multiple simultaneous distance bounding measurements to prevent movement based attacks as discussed in Section 4.5. This solution requires a large number of high bandwidth channels for the range measurements. None of related work above implement a localization system based on ToA measurements of RF signals.

6 Conclusion

In this paper, we propose a novel ID-based secure distance bounding protocol, and implemented this protocol on a COTS UWB ranging platform. Unlike traditional secure distance bounding protocols, our protocol is constructed using insecure distance measurement operations as basic building block. Thus, the proposed protocol lowers the complexity of the implementation and does not require modifications of existing ranging platforms. We discussed possible attacks on the protocol and implementation level and argued about their negligible impact. Based on this implementation of secure distance bounding, we further implemented a secure localization system that enables the correct computation of a device location in the presence of an adversary. We analyzed the implemented localization protocol and we discussed a number of improvements that increase its security and accuracy. To the best of our knowledge, this is the first implementation of an RF Time-of-Arrival (ToA) secure localization system. We are also the first to discuss the design choices related to different aggregation functions in the distance computation.

7 Acknowledgements

This work was partially supported by the Zurich Information Security Center. It represents the views of the authors.

References

1. Brands, S., Chaum, D.: Distance-bounding protocols. In: Proceedings of EUROCRYPT. (1994)
2. Hu, Y.C., Perrig, A., Johnson, D.B.: Packet leashes: a defense against wormhole attacks in wireless networks. In: Proceedings of IEEE InfoCom. (2003)
3. Hancke, G.P., Kuhn, M.G.: An RFID Distance Bounding Protocol. In: Proceedings of IEEE SecureComm. (2005)
4. Čapkun, S., Buttyan, L., Hubaux, J.P.: Sector: Secure tracking of node encounters in multi-hop wireless networks. In: Proceedings of ACM SASN. (2003)
5. Sastry, N., Shankar, U., Wagner, D.: Secure verification of location claims. In: Proceedings of ACM WiSe. (2003)

6. Kuhn, M.G.: An asymmetric security mechanism for navigation signals. In: Proceedings of the Information Hiding Workshop. (2004)
7. Lazos, L., Poovendran, R.: Serloc: secure range-independent localization for wireless sensor networks. In: Proceedings of ACM WiSe. (2004)
8. Lazos, L., Poovendran, R., Čapkun, S.: Rope: robust position estimation in wireless sensor networks. In: Proceedings of IPSN. (2005)
9. Čapkun, S., Hubaux, J.P.: Secure positioning in wireless networks. IEEE Journal on Selected Areas in Communications (2006)
10. Čapkun, S., Čagalj, M., Srivastava, M.: Secure localization with hidden and mobile base stations. In: Proceedings of IEEE InfoCom. (2006)
11. Rasmussen, K.B., Čapkun, S., Čagalj, M.: Secnav: secure broadcast localization and time synchronization in wireless networks. In: Proceedings of ACM/IEEE MobiCom. (2007)
12. Desmedt, Y.G.: Major security problems with the 'unforgeable' (feige-)fiat-shamir proofs of identity and how to overcome them. In: Proceedings of Securicom. (1988)
13. Papadimitratos, P., Poturalski, M., Schaller, P., Lafourcade, P., Basin, D., Čapkun, S., Hubaux, J.P.: Secure neighborhood discovery: A fundamental element for mobile ad hoc networking. IEEE Communications Magazine (2008)
14. Multispectral Solutions, Inc: UPS (Urban positioning system) *www.multispectral.com*.
15. Rasmussen, K.B., Čapkun, S.: Location privacy of distance bounding protocols. In: Proceedings of ACM CCS. (2008)
16. Clulow, J., Hancke, G.P., Kuhn, M.G., Moore, T.: So near and yet so far: Distance-bounding attacks in wireless networks. In: Proceedings of ESAS. (2006)
17. The MathWorks, Inc: Matlab – a numerical computing environment *www.mathworks.com*.
18. Tippenhauer, N.O., Čapkun, S.: UWB-based secure ranging and localization. Technical Report 586, ETH Zurich (January 2008)
19. Kalman, R.E.: A new approach to linear filtering and prediction problems. Transactions of the ASME Journal of Basic Engineering (1960)
20. Singelée, D., Preneel, B.: Distance Bounding in Noisy Environments. In: Proceedings of ESAS. (2007)
21. Reid, J., Nieto, J.M.G., Tang, T., Senadji, B.: Detecting relay attacks with timing-based protocols. In: Proceedings of ACM ASIACCS. (2007)
22. Bussard, L., Bagga, W.: Distance-bounding proof of knowledge to avoid real-time attacks. In: in proceedings of SEC. (2005)
23. Singelee, D., Preneel, B.: Location verification using secure distance bounding protocols. In: in Proceedings of MASS, Society Press (2005) 834–840
24. Drimer, S., Murdoch, S.J.: Keep your enemies close: Distance bounding against smartcard relay attacks. In: Proceedings of the USENIX Security Symposium. (2007)
25. Munilla, J., Ortiz, A., Peinado, A.: Distance bounding protocols with void-challenges for RFID. Printed handout at RFIDSec (July 2006)
26. Hancke, G., Kuhn, M.G.: Attacks on 'Time-of-Flight' Distance Bounding Channels. In: Proceedings of WiSeC. (2008)
27. Sedihpour, S., Čapkun, S., Ganeriwal, S., Srivastava, M.: Implementation of Attacks on Ultrasonic Ranging Systems, demo at ACM SENSYS'05 (2005)
28. Chiang, J.T., Haas, J.J., Hu, Y.C.: Secure and precise location verification using distance bounding and simultaneous multilateration. In: ACM-WISEC. (2009)