Location Leakage in Distance Bounding: Why Location Privacy does not Work

Aikaterini Mitrokotsa[1], Cristina Onete[2] and Serge Vaudenay[3]

[1] *Chalmers University of Technology*
Gothenburg, Sweden
aikaterini.mitrokotsa@chalmers.se

[2] *IRISA/INRIA, Univ. Rennes 1*
Rennes, France
maria-cristina.onete@irisa.fr

[3] *EPFL Lausanne, Switzerland*
serge.vaudenay@epfl.ch

Abstract

In many cases, we can only have access to a service by proving we are sufficiently close to a particular location (e.g., in automobile or building access control). In these cases, proximity can be guaranteed through signal attenuation. However, by using additional transmitters an attacker can relay signals between the prover and the verifier. Distance-bounding protocols are the main countermeasure against such attacks; however, such protocols may leak information regarding the location of the prover and/or the verifier who run the distance-bounding protocol.

In this paper, we consider a formal model for location privacy in the context of distance-bounding. In particular, our contributions are threefold: we first define a security game for location privacy in distance bounding; secondly, we define an adversarial model for this game, with two adversary classes; finally, we assess the feasibility of attaining location privacy for distance-bounding protocols. Concretely, we prove that for protocols with a beginning or a termination, it is theoretically impossible to achieve location privacy for either of the two adversary classes, in the sense that there always exists a polynomially-bounded adversary winning the security game. However, for so-called limited adversaries, who cannot see the location of arbitrary provers, carefully chosen parameters do, in practice, enable computational location privacy.

Preprint submitted to Computers and Security *February 24, 2014*

Keywords: location privacy, distance-bounding, authentication.

1. Introduction

Often, our location is critical in order to gain access to places and/or services. For instance, in applications such as automobile access control the key (prover) needs to be close enough to the car lock (verifier) in order to unlock it [17]. In some cases, unlocking the car may in fact also start the car (in passive keyless entry and start (PKES) systems [18]). If the proximity check is performed through signal attenuation, an adversary may easily perform man-in-the-middle attacks by relaying messages between the communicating parties (provers and verifiers), while these parties are situated far from each other. Thus, in the automobile example, an adversary may unlock the car even if the car key (the prover) is located very far. This type of attack (called mafia fraud [11]) can also be mounted against bankcards [13], mobile phones [19], proximity cards [20], and wireless ad-hoc networks [21, 27].

Distance-bounding (DB) protocols are meant to counteract man-in-the-middle relay attacks in authentication schemes. They are challenge-response authentication protocols, that allow the verifier, by measuring the time-of-flight of the messages exchanged, to calculate an upper bound on the prover's distance (as well as checking the validity of the responses which usually ensure authentication). DB protocols were first introduced by Brands and Chaum [6] to preclude relay attacks in ATM systems. Subsequently, numerous DB protocols were proposed [22, 30, 9] and many attacks against them have been published [2, 3, 15]. DB protocols have also been analysed for the case of noisy channels [23] and the optimal setting of security parameters [12, 25]. To the best of our knowledge [4, 5] describes the latest most secure distance-bounding protocol against all known attack modes. Another provably-secure protocol attaining quite strong terrorist-fraud resistance requirements has been recently published in [16].

Location privacy was introduced in the context of distance bounding by Rasmussen and Čapkun [28], who noted that distance-bounding protocols may leak further location-related information than just the fact that the prover is within the maximum allowed distance from the verifier. This information leakage follows from the measurement of the messages' arrival times.

To combat this, Rasmussen and Čapkun [28] proposed a privacy-preserving distance-bounding protocol (denoted here as the RČ protocol). Though the

protocol in [28] claims to preserve location privacy, we note that location privacy has never been formalized in the literature. Additionally, the RČ protocol has been shown to be susceptible to a non-polynomial dictionary attack which may reveal the prover's and verifier's locations [1] as well as to a *mafia fraud* attack [24]. Mitrokotsa *et al.* [24] have proposed a new distance-bounding protocol called *Location-Private Distance Bounding* (LPDB) that improves the basic construction of the RČ protocol and renders it secure against the latter attack.

Distance bounding can also be extended to location verification [32] (also known as secure positioning [31]), where multiple verifiers interact with a single prover. In that case the location of the prover can be determined using the intersection of the bounding spheres surrounding each verifier. This approach is also taken under consideration in the recent work regarding position-based cryptography [10]. Our approach here is different as we consider a single verifier and many provers, and we thus only achieve distance bounding, and not secure positioning. Moreover, in position-based cryptography all the adversaries have the same knowledge as the prover, including the secret key. However, in our model, we do not allow the adversary knowledge of the secret key, as that would allow it to trivially distinguish between the two provers in the location privacy game, without actually requiring any location data.

We also mention the recent work on localisation privacy by Burmester [7, 8], where location is used in a steganographic sense (such that provers are convinced that verifier-generated challenges are honest, and they do not reveal their presence to adversaries). However, very notably the constructions in [8] require provers to be aware of their position/location, which is a strong assumption in generic authentication/distance-bounding scenarios. In this case, location is used as a part of the verifier's challenge, and the prover verifies that the location is sufficiently close to its own location.

Contributions: In this paper, we address precisely the topics of location privacy in distance-bounding. Our contributions are threefold:

1. We first define a classical left-or-right *indistinguishability* game for location privacy in distance-bounding protocols. In this game, the adversary knows its distance to the verifier $\mathcal{V}$ and can create provers $\mathcal{P}$ at arbitrary distances from itself and $\mathcal{V}$.

2. For this location privacy game, we consider two main adversarial classes: *omniscient* and *limited* adversaries. *Omniscient* adversaries capture an

adversary that can measure the signal strength of the transmitted messages and is aware, for all transmissions along the timed channel, when the message is sent and when it arrives at its own interface. Unsurprisingly, no location privacy is feasible for omniscient adversaries. *Limited* adversaries, on the other hand, are only aware of the time at which they receive messages from other participants.

3. Finally, we show that achieving location privacy with respect to limited adversaries is impossible for protocols with a beginning or a termination, and which run in polynomial time. We prove that location privacy against limited adversaries minimally requires the prover and the verifier to introduce exponential delays between receiving and sending messages, and we give a lower bound for these delays. Since the transmission speed is high (e.g. the speed of light in the case of RFID transmissions), the delay can be implemented in practice. Finally, we show how to specify these delays in the LPDB protocol proposed in [24].

Organization: This paper is organized as follows. We begin by defining distance-bounding protocols and location privacy in Section 2, outlining also our adversarial classes. We then assess the feasibility of achieving location privacy for distance-bounding protocols in Section 3, for both *omniscient* and *limited* adversaries, giving a lower bound for the delays that each party must have between receiving a message and sending a response message. We apply our results and the obtained bound in Section 4, in order to modify the LPDB protocol [24] to attain location privacy with respect to limited adversaries.

2. Preliminaries

2.1. Communication Model

Our distance-bounding scenario resembles that of Dürholz *et al.* [14], but we consider multiple provers. Concretely, there is a single verifier $\mathcal{V}$, but many provers $\mathcal{P}_1, \ldots, \mathcal{P}_n$, such that $\mathcal{V}$ and $\mathcal{P}_i$ for every i share a secret key K_i output by a key generation algorithm Kg. We also assume that when it is initialised, the verifier $\mathcal{V}$ is also equipped with an upper bound on the maximum allowed communication time (or time distance) $t_{\max}$ between itself and the prover.

The communication model considered by [14] is round-based. However, e.g. the RČ [28] and the LPDB [24] distance-bounding protocols are *not*

round-based. Therefore, we consider a more generalised model, where the two parties $\mathcal{P}$ and $\mathcal{V}$ interact with no round-based restriction, via *two* types of channels: a *timeless* and a *timed* channel. Parties $\mathcal{P}$ and $\mathcal{V}$ may send messages m along each of the two channels (i.e., they are duplex channels). In order to make the model more realistic we consider the transmissions along the *timed* channel to be bit-by-bit.

More formally, the *timed* channel is associated with the global clock, such that each bit of an input message m will be associated with a time ts at which the sending party has *sent* the bit. The corresponding output bit of message m is associated with a time tr, which is the time at which the receiving party has *received* the bit. The bit-by-bit treatment of the transmission time is compulsory, as in practice, each bit of the message is transmitted sequentially or in smaller packets. However, for practical purposes we will often associate (in our proofs) the sending time of a message with the sending time of the first bit of this message, since this particular value is enough to leak significant information regarding the position of the honest protocol participants (prover and/or verifier). We discuss the soundness and the limitations of our model in Section 5.

For the sake of completeness, however, we associate in our model a message m with an $|m|$-dimensional vector of sending times $\bar{\mathsf{ts}}$ and an $|m|$-dimensional vector of transmission times $\bar{\mathsf{tr}}$. We also require that the values in $\bar{\mathsf{ts}}$ and those in $\bar{\mathsf{tr}}$ are monotone non-decreasing, i.e. for any message m and any $1 \leqslant i < j \leqslant m$, it holds that $\mathsf{ts}_i \leqslant \mathsf{ts}_j$ and $\mathsf{tr}_i \leqslant \mathsf{tr}_j$. Furthermore, if we consider the communication between two parties A and B and that a message m is sent from the party A to the party B at time $\bar{\mathsf{ts}}$ then the reception time $\bar{\mathsf{tr}}$ of the message m at the party B will satisfy the following equation for every $i = \{1, \ldots, |m|\}$[1]:

$$\mathsf{tr}_i = \mathsf{ts}_i + t_{AB},$$

where t_{AB} denotes the *time distance* between the parties A and B. More precisely, t_{AB} denotes the time (measured in time units TU) that every bit of a message m takes to travel between A and B.

Moreover, if the message m leaks off this channel to an adversary $\mathcal{A}$, each bit of the leaked message is associated with an $|m|$-dimensional timestamp

[1] In particular, we assume a perfect reliability of the transmission channel. We discuss the strength of this assumption in Section 5.

$\bar{\mathsf{tr}}_{\mathcal{A}}$. Note that this information alone may not suffice to learn the *sending* time of the message, as the adversary does not necessarily know the distance between it and the sending party.

Both channels allow the prover $\mathcal{P}$ and the verifier $\mathcal{V}$ to interact concurrently, i.e. it is possible that both the prover $\mathcal{P}$ and the verifier $\mathcal{V}$ transmit at the same time across the duplex channel. This is indeed the case for the RČ protocol [28].

We now define communication in distance-bounding protocols as being *slow* (or lazy) if it takes place on the timeless communication channel and *fast* (or time-critical) if it takes place on the timed communication channel. Note that it is possible to alternate fast and slow communication arbitrarily. We note that this approach is perfectly in-tune with the similar communication model of [14], but it is also compatible with protocols that are not round-based.

Definition 1. *We say that* $\mathcal{DB} = (\mathcal{V}, \mathcal{P}, \mathsf{Kg})$ *is a* distance-bounding protocol with parameters $(t_{\max}, \epsilon)$ *where* $t_{\max}$ *denotes the upper bound on transmission time in the fast phase and* ϵ *denotes the tolerance level for honest* $\mathcal{P}$*-*$\mathcal{V}$ *authentication failures if and only if:*

Key Generation: Kg *generates a secret key* $K \leftarrow \mathsf{Kg}(1^\ell)$ *for any* $\ell \in \mathbb{N}$.
Distance-Bounding Authentication: *The joint execution of the prover and verifier algorithms* $\mathcal{V}$ *and* $\mathcal{P}$ *for parameters* $(t_{\max}, \epsilon)$ *ends with a verifier-generated distance-bounding authentication bit* $b \in \{0, 1\}$.
We require ϵ*-completeness, i.e., the interaction of an honest prover* $\mathcal{P}$ *and an honest, fixed verifier* $\mathcal{V}$ *for parameters* $(t_{\max}, \epsilon)$ *is accepted by the verifier with probability at least* $1 - \epsilon$ *if* $t_{\mathcal{VP}} \leqslant t_{\max}$.

2.2. Adversarial Models

In our framework, the goal of the adversary is to break location privacy as defined below. In this section, we first show how adversaries interact with the communication channels and with the honest parties during an attack. Then, we define two adversarial classes depending on the strength of the adversary. Finally, we show the location privacy game.

We consider adversaries $\mathcal{A}$ that interact with the distance-bounding system as follows: (1) $\mathcal{A}$ may eavesdrop on the communication (across both the *timed* and the *timeless* channel) of an honest prover $\mathcal{P}$ and an honest verifier $\mathcal{V}$; and (2) $\mathcal{A}$ may interact with honest provers in prover-adversary sessions and with honest verifiers in adversary-verifier sessions. Note that

this behavior implies that an adversary can mount a full man-in-the-middle attack by simply opening concurrent prover-adversary and adversary-verifier sessions. This is again in agreement with the treatment given by Dürholz *et al.*; we refer to that paper for the more formal notions of session identifiers.

In view of [33, 26], we consider that frequency hopping (i.e. implementing a protocol such that the sender and the receiver hop from one frequency to another during the transmission) is not an effective countermeasure against eavesdropping adversaries. In particular, by simply eavesdropping all possible frequencies (in practice the prover and the verifier are unable to use too many different frequencies), the adversary can successfully "piece together" the communication.

We consider two types of adversaries: the *limited* and the *omniscient* adversaries, which are described as follows:

Limited adversaries: These adversaries may eavesdrop on honest prover-verifier sessions or communicate with provers and verifiers in prover-adversary and respectively adversary-verifier sessions. On eavesdropping the timed channel in honest prover-verifier sessions, limited adversaries learn the transmitted message m and the bit-by-bit time the message is received at, $\bar{\mathsf{tr}}_{\mathcal{A}} = \bar{\mathsf{ts}} + \bar{t}_{P\mathcal{A}}$, where P is the party that sent the message m and $\bar{t}_{P\mathcal{A}}$ is an $|m|$-dimensional vector with entries equaling the time distance $t_{P\mathcal{A}}$ between P and the adversary $\mathcal{A}$. Note that the adversary $\mathcal{A}$ is able to choose its location and knows $t_{\mathcal{A}\mathcal{V}}$ (i.e. its time distance from the verifier $\mathcal{V}$); consequently, $\mathcal{A}$ learns the sending times at which the verifier sends its messages.

Omniscient adversaries: These adversaries can also eavesdrop on honest prover-verifier sessions or communicate with provers and verifiers as above. Additionally, an omniscient adversary can measure the signal strength of the transmitted messages and is aware, for all transmissions along the timed channel, when the message is sent and when it arrives at its interface. More precisely, on eavesdropping on the timed channel during an honest prover-verifier session, omniscient adversaries learn the message m, the bit-by-bit time the message is received, $\bar{\mathsf{tr}}_{\mathcal{A}} = \bar{\mathsf{ts}} + \bar{t}_{P\mathcal{A}}$, and the bit-by-bit sending time $\bar{\mathsf{ts}}$. Thus, strong adversaries can trivially learn the distance between them and the party P that sent the message.

To justify that an omniscient adversary can also learn the sending time of messages, we could model this by distributed, *limited* adversaries, i.e. $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. The composite adversary $\mathcal{A}$ chooses the *locations* of $\mathcal{A}_1$ and $\mathcal{A}_2$ and can do triangulation of signals. This definition also extends to a *moving* adversary (i.e. an adversary that is able to change its location) as

discussed in Section 3.1.

We consider only polynomial adversaries, (i.e. having polynomial run-time and running polynomially many sessions with the provers and the verifier). The adversary's goal is to break the location privacy of the distance-bounding protocol, which we define by means of a left-or-right *indistinguishability* game as described below.

PHASE 1: In this phase, a limited adversary is given the security parameter (in unary) 1^λ. The adversary may now initialise provers $\mathcal{P}_i$ and the verifier $\mathcal{V}$ at arbitrary locations with respect to itself and the verifier, and may interact arbitrarily with the provers and the verifier. At the end of this phase, the adversary outputs two indices i, j such that $t_{\mathcal{P}_i\mathcal{V}}$ and $t_{\mathcal{P}_j\mathcal{V}}$ are both smaller than the threshold $t_{\max}$; the two indices are then forwarded to a challenger.

PHASE 2: The challenger checks that the two provers are both within the maximum distance $t_{\max}$, then closes all sessions that are open for these provers. The challenger flips a bit b and assigns the handle $\mathcal{P}_{\text{Chal}}$ as follows: $\mathcal{P}_{\text{Chal}} = \mathcal{P}_i$ if $b = 0$ and $\mathcal{P}_{\text{Chal}} = \mathcal{P}_j$ if $b = 1$.

PHASE 3: Finally, by interacting with the challenge prover $\mathcal{P}_{\text{Chal}}$, as well as all other provers with the exception of $\mathcal{P}_i$ and $\mathcal{P}_j$, the adversary must produce a decision bit d. Let $\mathsf{Exp}^{\mathsf{LocPriv}}_{\mathcal{DB}}(\mathcal{A}, 1^\lambda)$ be the output of a single run of the location privacy game. We say that the adversary *wins* if $d = b$, and we write it as $\mathsf{Exp}^{\mathsf{LocPriv}}_{\mathcal{DB}}(\mathcal{A}, 1^\lambda) = 1$. The adversary can be considered as a hypothesis test for the following hypotheses:

$\mathcal{H}_0$: the response sent from the prover $\mathcal{P}_{\text{Chal}}$ to $\mathcal{V}$'s challenge is actually from the prover $\mathcal{P}_0$.

and

$\mathcal{H}_1$: the response sent from the prover $\mathcal{P}_{\text{Chal}}$ to $\mathcal{V}$'s challenge is actually from the prover $\mathcal{P}_1$.

We define the advantage of the adversary in this game as:

$$\mathsf{Adv}^{\mathsf{LocPriv}}_{\mathcal{DB}}\mathcal{A} = \left|2\mathbb{P}\left[\mathsf{Exp}^{\mathsf{LocPriv}}_{\mathcal{DB}}(\mathcal{A}, 1^\lambda) = 1\right] - 1\right|$$

Definition 2. *We say that distance-bounding protocols provide* location privacy *if* $\forall loc_{\mathcal{P}_0}, loc_{\mathcal{P}_1}$, $\forall loc_{\mathcal{V}}$, $\forall \mathcal{A}$ *it holds:*

$$Adv^{\mathsf{LocPriv}}_{\mathcal{DB}}\mathcal{A} = negl(1^\lambda),$$

where $\mathsf{negl}(1^\lambda)$ *denotes a negligible function in the security parameters.*

We should note here that an adversary would select the location of the participants in such a way as to maximize his advantage. Thus, an adversary $\mathcal{A}$ would not select $\mathcal{P}_0$ and $\mathcal{P}_1$ at the same location or at equal distance to $\mathcal{A}$ and $\mathcal{V}$.

3. Why Location Privacy does not Work

In this section we first argue that *location privacy* cannot be achieved with respect to an *omniscient* adversary. Then, we show that *location privacy* can only be achieved with respect to *limited* adversaries if the honest parties running the protocol introduce (minimally) a delay in their transmissions; we furthermore give a lower bound on this delay.

3.1. Omniscient Adversary

It is trivial to see that no location privacy can be attained with respect to an omniscient adversary. Indeed, consider an omniscient adversary placed arbitrarily with respect to the verifier. Let this adversary $\mathcal{A}$ create two provers $\mathcal{P}_0$ and $\mathcal{P}_1$ such that the distance between this adversary and the provers is different i.e. $t_{\mathcal{P}_0\mathcal{A}} \neq t_{\mathcal{P}_1\mathcal{A}}$[2].

The adversary forwards $\mathcal{P}_0, \mathcal{P}_1$ to the challenger, receiving the handle $\mathcal{P}_{\text{Chal}}$, which is either $\mathcal{P}_0$ or $\mathcal{P}_1$. Now, the adversary eavesdrops on a session between $\mathcal{P}_{\text{Chal}}$ and $\mathcal{V}$, thus learning the sending time of the messages and the time the attacker receives them. The adversary thus calculates the time distance between itself and the two parties communicating and, since the distances are all different, it can identify the parties with probability 1.

A single, but *moving* adversary (i.e., an adversary than can change its position during the attack) could also infer some information about the location of the prover by standing between $\mathcal{P}_0$ and $\mathcal{P}_1$ and moving toward $\mathcal{P}_0$ due to the Doppler effect. If bits arrive with a higher frequency, they must be sent by $\mathcal{P}_0$ instead of $\mathcal{P}_1$.

[2]Obviously an adversary $\mathcal{A}$ would choose its location in order to maximise its advantage. Thus, choosing provers at equal distance to it would not be a good choice.

3.2. Limited Adversary

By eavesdropping on the duplex timed channel between the challenged prover and the verifier, the adversary will receive $\mathsf{tr}^i_{\mathcal{A}}$, the timestamp when $\mathcal{A}$ receives the first bit of message m_i. The adversary $\mathcal{A}$ also observes:

- $t_{\mathcal{V}} = \mathsf{tr}^1_{\mathcal{A}}$: the time $\mathcal{A}$ receives the first message bit from $\mathcal{V}$.
- $t_{\mathcal{P}} = \mathsf{tr}^2_{\mathcal{A}}$: the time $\mathcal{A}$ receives the first message bit from $\mathcal{P}$.

In what follows we show that the very first bit sent through the timed channel leaks. To be able to prove that, we make the following reasonable assumptions regarding how the sending time of this first bit is decided during the protocol. Note that similar observations hold for the final bit sent. For simplicity, we only treat the first one.

Assumption 1. *We assume that the distance-bounding phase of a distance-bounding protocol may have one of the following constructions:*

- ***Case 1:*** *The verifier $\mathcal{V}$ starts the distance-bounding phase after a reference time t_0 and a random delay, possibly equal to 0, which we denote $delay_{\mathcal{V}}$, while the prover $\mathcal{P}_b$ where $b \in \{0,1\}$ starts after receiving the first message from the verifier $\mathcal{V}$ and a random delay $delay_{\mathcal{P}_b}$.*
- ***Case 2:*** *The prover $\mathcal{P}_b$ starts the distance-bounding phase after a reference time t_0 and a random delay $delay_{\mathcal{P}_b}$, while the verifier $\mathcal{V}$ starts after receiving the first message from the prover $\mathcal{P}_b$ and a random delay $delay_{\mathcal{V}}$.*
- ***Case 3:*** *The prover $\mathcal{P}_b$ and the verifier $\mathcal{V}$ start sending messages independently. More precisely, the prover $\mathcal{P}_b$ starts sending messages after a reference time $T_{\mathcal{P}_b}$ and a random delay $delay_{\mathcal{P}_b}$, while the verifier $\mathcal{V}$ starts sending messages after a reference time $T_{\mathcal{V}}$ and a random delay $delay_{\mathcal{V}}$.*

We should note here that when we mention "random delay" we mean a delay of arbitrary distribution.

Assumption 2. *We also assume that $\mathcal{A}$ knows the times $T_{\mathcal{P}_b}$ (where $b \in \{0,1\}$) and $T_{\mathcal{V}}$; the latter value is defined only for Case 3 of* Assumption 1.

In Figure 1 are depicted the above described cases. Without loss of generality in Figure 1 the adversary $\mathcal{A}$ is located between the verifier $\mathcal{V}$ and the prover $\mathcal{P}$.

It is easy to see that in our model a limited adversary $\mathcal{A}$ knows and can even choose the locations of $\mathcal{P}_0$, $\mathcal{P}_1$ with respect to itself and the verifier $\mathcal{V}$, i.e. the values $t_{\mathcal{A}\mathcal{P}_0}, t_{\mathcal{A}\mathcal{P}_1}, t_{\mathcal{V}\mathcal{P}_0}, t_{\mathcal{V}\mathcal{P}_1}$. Also, $\mathcal{A}$ knows the distance $t_{\mathcal{A}\mathcal{V}}$ to $\mathcal{V}$. We will show how an adversary intercepting the values above can distinguish between the two hypotheses $\mathcal{H}_0, \mathcal{H}_1$ with non-negligible probability.

Lemma 1. *Under Assumptions 1 and 2 we assume that there exists ϵ and a bound B such that:*

$$\mathbb{P}[delay \leqslant B] = 1 - \epsilon,$$

where delay might represent the delays of the provers $delay_{\mathcal{P}_0}$, $delay_{\mathcal{P}_1}$, or the delay ($delay_{\mathcal{V}}$) of the verifier as defined in Assumption 1. Then there exists an adversary $\mathcal{A}$ against location indistinguishability which achieves a distinguishing advantage:

$$\mathsf{Adv}_{\mathcal{A}} \geqslant \left\lceil \frac{t_{\max}}{4B} \right\rceil (1 - 2\epsilon),$$

where $t_{\max}$ is the maximum allowed transmission time between a legitimate prover $\mathcal{P}$ and a verifier $\mathcal{V}$.

Moreover, this adversary does not need to take part in the actual protocol; the attack relies exclusively on eavesdropping. Assuming that the protocol is complete and polynomially bounded, there is a negligible ϵ such that B exists and is polynomially bounded. So, the advantage $\mathsf{Adv}_{\mathcal{A}}$ is not negligible. Consequently, a distance-bounding protocol as defined in Definition 1 does not provide *location privacy* as per Definition 2.

Proof. Based on *Assumption 1* we have three cases.
Case 1: *The verifier $\mathcal{V}$ starts the* distance bounding phase *after a reference time t_0 and a random delay (denoted as $delay_{\mathcal{V}}$), whereas the prover $\mathcal{P}_b$ starts after receiving the first message from the verifier $\mathcal{V}$ and a random delay (denoted as $delay_{\mathcal{P}_b}$).*
This case is depicted in figure 1 (a). More precisely, we consider that the following events take place:

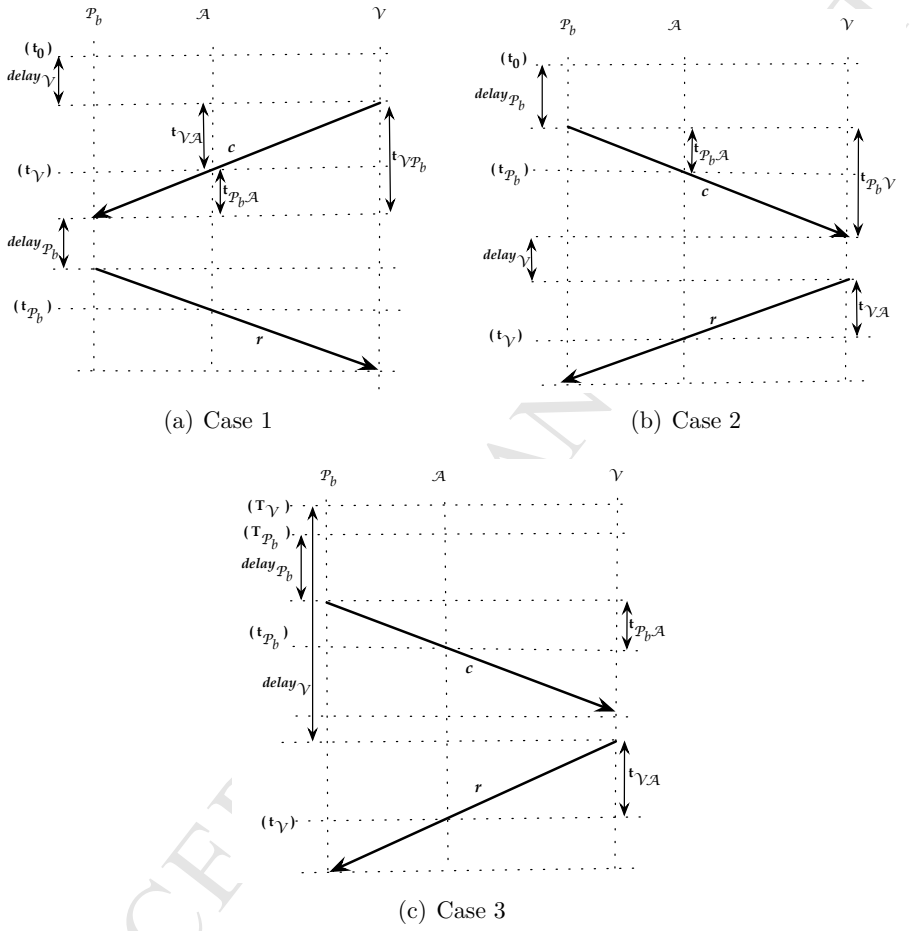

(a) Case 1 (b) Case 2

(c) Case 3

Figure 1: Transmission of messages between the verifier and the prover for the three different cases of the construction of a distance-bounding protocol.

1. After some time reference t_0 and a $delay_{\mathcal{V}}$ the verifier $\mathcal{V}$ sends a message c to the prover $\mathcal{P}_b$ where $b \in \{0,1\}$. The first bit of this message will arrive at the adversary $\mathcal{A}$ at time $t_{\mathcal{V}}$ such that:

$$t_{\mathcal{V}} = t_0 + delay_{\mathcal{V}} + t_{\mathcal{V}\mathcal{A}}, \tag{1}$$

where $t_{\mathcal{V}\mathcal{A}}$ denotes the time of flight for one bit from the verifier $\mathcal{V}$ to the adversary $\mathcal{A}$.

2. The prover $\mathcal{P}_b$ with $b \in \{0,1\}$ responds to the verifier $\mathcal{V}$ with a message r, after some delay ($delay_{\mathcal{P}_b}$). The first bit of r arrives at $\mathcal{A}$ at time $t_{\mathcal{P}_b}$ such that:

$$t_{\mathcal{P}_b} = t_0 + delay_{\mathcal{V}} + t_{\mathcal{V}\mathcal{P}_b} + delay_{\mathcal{P}_b} + t_{\mathcal{P}_b\mathcal{A}}, \tag{2}$$

where $t_{\mathcal{V}\mathcal{P}_b}$ denotes the time-of-flight for one bit from $\mathcal{V}$ to $\mathcal{P}_b$, and $t_{\mathcal{P}_b\mathcal{A}}$ denotes the time-of-flight for one bit from $\mathcal{P}_b$ to $\mathcal{A}$.

From equations (1) and (2) it is easy to see that:

$$t_{\mathcal{P}_b} - t_{\mathcal{V}} = t_{\mathcal{V}\mathcal{P}_b} - t_{\mathcal{V}\mathcal{A}} + delay_{\mathcal{P}_b} + t_{\mathcal{P}_b\mathcal{A}}.$$

We let d_b be the probability density function (pdf) of $delay_{\mathcal{P}_b}$, i.e. we consider the delay to be a random variable distributed according to d_b. If hypothesis $\mathcal{H}_0$ holds, then $t_{\mathcal{P}} = t_{\mathcal{P}_0}$, while if hypothesis $\mathcal{H}_1$ holds, then $t_{\mathcal{P}} = t_{\mathcal{P}_1}$. Since $t_{\mathcal{P}}$ and $t_{\mathcal{V}}$ depend on random delays, they can be perceived as random variables. Let:

$$T = t_{\mathcal{P}} - t_{\mathcal{V}} - t_{\mathcal{V}\mathcal{P}_0} + t_{\mathcal{V}\mathcal{A}} - t_{\mathcal{P}_0\mathcal{A}} \quad \text{and}$$
$$\Delta = t_{\mathcal{V}\mathcal{P}_1} + t_{\mathcal{P}_1\mathcal{A}} - t_{\mathcal{V}\mathcal{P}_0} - t_{\mathcal{P}_0\mathcal{A}}.$$

Note that whereas the value Δ is fixed and even chosen by the adversary, T is a random variable, depending on the delays. Indeed, if hypothesis $\mathcal{H}_0$ holds then $T = delay_{\mathcal{P}_0}$ has pdf d_0, while if hypothesis $\mathcal{H}_1$ holds, then $T = delay_{\mathcal{P}_1} + \Delta$ and we write $\mathbb{P}\big[T = t\big] = d_1(t - \Delta)$, i.e. T has a distribution equivalent to d_1, shifted by a fixed value Δ.

In the following, we often condition success probabilities on hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$ and use the notation $\mathbb{P}_{\mathcal{H}_b}[\text{event}]$ for $\mathbb{P}[\text{event} \mid \mathcal{H}_b \text{ holds}]$, i.e. the probability that *event* holds, conditioned on the fact that $\mathcal{H}_b$ holds.

We consider that $\mathcal{A}$ is implementing a best distinguisher based on the likelihood that $\mathbb{P}_{\mathcal{H}_0}[T = t] > \mathbb{P}_{\mathcal{H}_1}[T = t]$ for observed value t. If this holds, then $\mathcal{A}$ outputs 0, else it outputs 1. So $\mathcal{A}$ outputs 0 iff the observed value of $T = t_{\mathcal{P}} - t_{\mathcal{V}} - t_{\mathcal{V}\mathcal{P}_0} + t_{\mathcal{V}\mathcal{A}} - t_{\mathcal{P}_0\mathcal{A}}$ is $T = t$ such that:

$$\mathbb{P}[t = delay_{\mathcal{P}_0}] > \mathbb{P}[t = delay_{\mathcal{P}_1} + \Delta].$$

Then, it holds:

$$\begin{aligned} \mathsf{Adv} &= \mathbb{P}_{\mathcal{H}_0}[\mathcal{A} \to 0] - \mathbb{P}_{\mathcal{H}_1}[\mathcal{A} \to 0] \\ &= \frac{1}{2}\int_{-\infty}^{+\infty} |d_0(t) - d_1(t-\Delta)| dt, \end{aligned} \tag{3}$$

where d_0 and d_1 make $[0, B]$ have density at least $1 - \epsilon$. When $t_{\mathcal{P}_0\mathcal{V}} = t_{\mathcal{P}_1\mathcal{V}} = t_{\max}$, $\mathcal{P}_0$, $\mathcal{V}$ and $\mathcal{P}_1$ are aligned in this order, and the adversary $\mathcal{A}$ overlaps with the location of $\mathcal{P}_0$, then $\Delta = 2t_{\max}$.

Case 2: *The prover $\mathcal{P}_b$ starts the* distance-bounding phase *after a reference time t_0 and a random delay (denoted as $delay_{\mathcal{P}_b}$). While the verifier $\mathcal{V}$ starts after receiving the first message from the prover $\mathcal{P}_b$ and a random delay (denoted as $delay_{\mathcal{V}}$).*

This case is depicted in figure 1 (b). Now, we have:

$$\begin{aligned} t_{\mathcal{P}_b} &= t_0 + delay_{\mathcal{P}_b} + t_{\mathcal{P}_b\mathcal{A}} \\ t_{\mathcal{V}} &= t_0 + delay_{\mathcal{P}_b} + t_{\mathcal{P}_b\mathcal{V}} + delay_{\mathcal{V}} + t_{\mathcal{V}\mathcal{A}} \\ t_{\mathcal{V}} - t_{\mathcal{P}_b} &= t_{\mathcal{P}_b\mathcal{V}} + delay_{\mathcal{V}} + t_{\mathcal{V}\mathcal{A}} - t_{\mathcal{P}_b\mathcal{A}}. \end{aligned}$$

We let:

$$\begin{aligned} T &= t_{\mathcal{V}} - t_{\mathcal{P}} - t_{\mathcal{P}_0\mathcal{V}} - t_{\mathcal{V}\mathcal{A}} + t_{\mathcal{P}_0\mathcal{A}} \quad \text{and} \\ \Delta &= t_{\mathcal{P}_1\mathcal{V}} - t_{\mathcal{P}_1\mathcal{A}} - t_{\mathcal{P}_0\mathcal{V}} + t_{\mathcal{P}_0\mathcal{A}}. \end{aligned}$$

Similarly, if the adversary $\mathcal{A}$ is implementing a distinguisher for the two provers $\mathcal{P}_0$ and $\mathcal{P}_1$ then its advantage is given by:

$$\begin{aligned} \mathsf{Adv} &= \mathbb{P}_{\mathcal{H}_0}[\mathcal{A} \to 0] - \mathbb{P}_{\mathcal{H}_1}[\mathcal{A} \to 0] \\ &= \frac{1}{2}\int_{-\infty}^{+\infty} |d(t) - d(t-\Delta)|\, dt, \end{aligned} \tag{4}$$

where d denotes the pdf of the random variable $delay_{\mathcal{V}}$, such that $[0, B]$ has density at least $1 - \epsilon$. When $t_{\mathcal{P}_0\mathcal{V}} = t_{\mathcal{P}_1\mathcal{V}} = t_{\max}$, $\mathcal{P}_0$, $\mathcal{V}$ and $\mathcal{P}_1$ are aligned

and the location of the adversary $\mathcal{A}$ overlaps with the location of the prover $\mathcal{P}_1$, then $\Delta = 2t_{\max}$. Thus, from equations (3) and (4) we derive that in both cases it holds:

$$\mathsf{Adv} = \frac{1}{2}\int_{-\infty}^{+\infty} |q_0(t) - q_1(t-\Delta)|dt$$

for some functions q_0 and q_1 that make $[0, B]$ have density at least $1-\epsilon$. We further have a case where $\Delta = 2t_{\max}$. Let:

$$x_{b,i} = \int_{(i-1)|\Delta|}^{i|\Delta|} q_b(t)dt \quad \text{and } n = \left\lceil \frac{B}{|\Delta|} \right\rceil.$$

We have $x_{b,0} = 0$, $x_{b,n+1} = 0$, $x_{b,i} \geqslant 0$ and $x_{b,1} + \cdots + x_{b,n} \geqslant 1 - \epsilon$. Given $I \subseteq \{0, \ldots, n\}$ we let $T_I = \bigcup_{i \in I} \left[(i-1)|\Delta|, i|\Delta|\right]$. For $\Delta > 0$, we have:

$$\begin{aligned} \mathsf{Adv}_{T_I,\Delta} &= \sum_{i \in I}(x_{0,i} - x_{1,i-1}) \text{ and} \\ \mathsf{Adv}_{T_I,-\Delta} &= \sum_{i \in I}(x_{0,i} - x_{1,i+1}). \end{aligned} \tag{5}$$

Let:

$$\begin{aligned} \mathsf{Adv}_{\Delta} &= \max_I \mathsf{Adv}_{T_I,\Delta} = \frac{1}{2}\sum_{i=0}^{n} |x_{0,i} - x_{1,i-1}| \\ \mathsf{Adv}_{-\Delta} &= \max_I \mathsf{Adv}_{T_I,-\Delta} = \frac{1}{2}\sum_{i=0}^{n} |x_{0,i} - x_{1,i+1}|. \end{aligned}$$

We have:

$$\begin{aligned} \mathsf{Adv}_{\Delta} + \mathsf{Adv}_{-\Delta} &= \frac{1}{2}\sum_{i=0}^{n}(|x_{0,i} - x_{1,i-1}| + |x_{0,i} - x_{1,i+1}|) \\ &\geqslant \frac{1}{2}\sum_{i=0}^{n} |x_{1,i+1} - x_{1,i-1}|. \end{aligned} \tag{6}$$

Since $x_{1,i} \geqslant 0$ and $x_{1,1} + \cdots + x_{1,n} \geqslant 1 - \epsilon$, there exists some index j such that: $x_{1,j} \geqslant \frac{1-\epsilon}{n}$. Thus:

$$\begin{aligned} \mathsf{Adv}_{\Delta} + \mathsf{Adv}_{-\Delta} &\geqslant \frac{1}{2}(|x_{1,j} - x_{1,j-2}| + |x_{1,j-2} - x_{1,j-4}| + \ldots) \\ &\geqslant \frac{x_{1,j}}{2} \geqslant \frac{1-\epsilon}{2n}. \end{aligned} \tag{7}$$

Thus,

$$\max(\mathsf{Adv}_{\Delta}, \mathsf{Adv}_{-\Delta}) \geqslant \frac{1-\epsilon}{4n}.$$

So, there exists Δ such that:

$$\mathsf{Adv}_{\Delta} \geqslant \left\lceil \frac{|\Delta|}{4B} \right\rceil (1-\epsilon).$$

For $\Delta = 2t_{\max}$ there exists an adversary $\mathcal{A}$ such that:

$$\mathsf{Adv}_{\mathcal{A}} \geqslant \left\lceil \frac{t_{\max}}{2B} \right\rceil (1-\epsilon).$$

Case 3: *The prover $\mathcal{P}_b$ and the verifier $\mathcal{V}$ send messages independently. More precisely, the prover $\mathcal{P}_b$ starts sending messages after a reference time $T_{\mathcal{P}_b}$ and a random delay ($delay_{\mathcal{P}_b}$) while the verifier $\mathcal{V}$ starts sending messages after a reference time $T_{\mathcal{V}}$ and a random delay ($delay_{\mathcal{V}}$). We assume that for this case the adversary $\mathcal{A}$ knows the values $T_{\mathcal{P}_b} - T_{\mathcal{V}}$.*
This case is depicted in Figure 1 (c). We now have:

$$\begin{aligned} t_{\mathcal{V}} &= T_{\mathcal{V}} + delay_{\mathcal{V}} + t_{\mathcal{V}\mathcal{A}} \\ t_{\mathcal{P}_b} &= T_{\mathcal{P}_b} + delay_{\mathcal{P}_b} + t_{\mathcal{P}_b\mathcal{A}} \\ t_{\mathcal{P}_b} - t_{\mathcal{V}} &= delay_{\mathcal{P}_b} - delay_{\mathcal{V}} + T_{\mathcal{P}_b} + t_{\mathcal{P}_b\mathcal{A}} - T_{\mathcal{V}} - t_{\mathcal{V}\mathcal{A}}. \end{aligned}$$

We let:

$$T = t_{\mathcal{P}} - t_{\mathcal{V}} - T_{\mathcal{P}_1} - t_{\mathcal{P}_1\mathcal{A}} + T_{\mathcal{V}} + t_{\mathcal{V}\mathcal{A}} \text{ and} \tag{8}$$

$$\Delta = T_{\mathcal{P}_1} + t_{\mathcal{P}_1\mathcal{A}} - T_{\mathcal{P}_0} - t_{\mathcal{P}_0\mathcal{A}}. \tag{9}$$

We consider that the adversary $\mathcal{A}$ is implementing a best distinguisher based on the likelihood if $\mathbb{P}_{\mathcal{H}_0}[t_{\mathcal{P}} - t_{\mathcal{V}}] > \mathbb{P}_{\mathcal{H}_1}[t_{\mathcal{P}} - t_{\mathcal{V}}]$ then $\mathcal{A}$ outputs 0; otherwise it outputs 1. So, $\mathcal{A}$ outputs 0 iff $t_{\mathcal{P}} - t_{\mathcal{V}} - T_{\mathcal{P}_1} - t_{\mathcal{P}_1\mathcal{A}} + T_{\mathcal{V}} + t_{\mathcal{V}\mathcal{A}} = T = t$ such that:

$$\mathcal{P}[t = delay_{\mathcal{P}_0} - delay_{\mathcal{V}}] > \mathcal{P}[t = delay_{\mathcal{P}_1} - delay_{\mathcal{V}} + \Delta].$$

Then, it holds:

$$\begin{aligned} \mathsf{Adv} &= \mathbb{P}_{\mathcal{H}_0}[\mathcal{A} \to 0] - \mathbb{P}_{\mathcal{H}_1}[\mathcal{A} \to 0] \\ &= \frac{1}{2} \int_{-\infty}^{+\infty} |q_0(t) - q_1(t-\Delta)| dt, \end{aligned} \tag{10}$$

where q_b for $b \in \{0,1\}$ denotes the pdf of the random variable $delay_{\mathcal{P}_b} - delay_{\mathcal{V}}$ and the support of q_0 and q_1 make $[-B, B]$ have density at least $1-2\epsilon$. When $t_{\mathcal{P}_0\mathcal{V}} = t_{\mathcal{P}_1\mathcal{V}} = t_{\max}$, $\mathcal{P}_0$, $\mathcal{V}$ and $\mathcal{P}_1$ are aligned in this order and if $T_{\mathcal{P}_1} \geqslant T_{\mathcal{P}_0}$ the location of the adversary $\mathcal{A}$ overlaps with the location of $\mathcal{P}_0$ while if $T_{\mathcal{P}_1} < T_{\mathcal{P}_0}$ the location of the adversary $\mathcal{A}$ overlaps with the location of the prover $\mathcal{P}_1$. Thus, in both of these cases it holds that $|\Delta| \geqslant 2t_{\max}$. Let:

$$x_{b,i} = \int_{(i-1)|\Delta|}^{i|\Delta|} q_b(t)dt \quad \text{and } n = \left\lceil \frac{B}{|\Delta|} \right\rceil.$$

We have $x_{b,0} = 0$, $x_{b,n+1} = 0$, $x_{b,i} \geqslant 0$, $x_{b,-n+1} + ... + x_{b,n} \geqslant 1 - 2\epsilon$ and:

$$\begin{aligned} \mathsf{Adv}_{\Delta} + \mathsf{Adv}_{-\Delta} &= \frac{1}{2} \sum_{i=-n}^{n} \left(|x_{0,i} - x_{1,i-1}| + |x_{0,i} - x_{1,i+1}| \right) \\ &\geqslant \frac{1}{2} \sum_{i=0}^{-n} |x_{1,i+1} - x_{1,i-1}|. \end{aligned}$$

Since $x_{1,i} \geqslant 0$ and $x_{1,-n+1} + \cdots + x_{1,n} \geqslant 1 - 2\epsilon$, there exists some index j such that: $x_{1,j} \geqslant \frac{1-2\epsilon}{2n}$. Thus:

$$\begin{aligned} \mathsf{Adv}_{\Delta} + \mathsf{Adv}_{-\Delta} &\geqslant \frac{1}{2}(|x_{1,j} - x_{1,j-2}| + |x_{1,j-2} - x_{1,j-4}| + \ldots) \\ &\geqslant \frac{x_{1,j}}{2} \geqslant \frac{1-2\epsilon}{4n}. \end{aligned}$$

Thus,

$$\max(\mathsf{Adv}_{\Delta}, \mathsf{Adv}_{-\Delta}) \geqslant \frac{1-2\epsilon}{8n}.$$

So, there exists Δ such that:

$$\mathsf{Adv} \geqslant \left\lceil \frac{|\Delta|}{8B} \right\rceil \geqslant \frac{t_{\max}}{4B}(1-2\epsilon).$$

□

Lemma 2. *If Assumption 1 holds and d_b follows the uniform distribution in the range $[0, B]$ and denotes the pdf of the $delay_{\mathcal{P}_b}$ while $delay_{\mathcal{V}}$ is always equal to 0 then the best distinguisher based on $t_{\mathcal{P}} - t_{\mathcal{V}}$ and the locations satisfies:*

$$\mathsf{Adv}_{\mathcal{A}} = \frac{2t_{\max}}{B},$$

where $t_{\max}$ denotes the maximum allowed transmission time between a legitimate prover $\mathcal{P}$ and a verifier $\mathcal{V}$.

PROOF. Following the proof of the Lemma 1 on page 11 the best distinguisher based on $t_{\mathcal{P}} - t_{\mathcal{V}}$ and the locations (of the provers and the verifier) follows equations (3), (4) or (10). So, it satisfies:

$$\mathsf{Adv} = \frac{1}{2} \int_{-\infty}^{+\infty} \left| d_0(t) - d_1(-\Delta + t) \right| dt$$

since $delay_{\mathcal{V}} = 0$. Since d_b follows the uniform distribution in the range $[0, B]$, it holds:

$$\mathsf{Adv}_{\mathcal{A}} = \frac{1}{2} \int_0^{\Delta} \frac{dt}{B} + \frac{1}{2} \int_B^{B+\Delta} \frac{dt}{B} = \frac{\Delta}{B}$$

and Δ is bounded by $2t_{\max}$ in all three cases.

□

Practical Consequences. Although the attack is polynomial, we can still live with it in practice thanks to the very high celerity of light, since the time it takes to cover 10 m is 2^{-25} sec. Indeed, let:

$$h = \log_2 \frac{B}{2t_{\max}}$$

The best advantage is comparable to guessing h bits correctly. To have a privacy level of h bits (i.e., a best advantage of 2^{-h}), we shall thus have:

$$B \geqslant 2^{h+1} t_{\max} \tag{11}$$

For instance, when $t_{\max}$ is the time light takes to go through the distance of 10 m and $h = 20$ bits (i.e., an adversary cannot distinguish two provers, accept with one chance out of a million), we have $B \geqslant 0.07$ sec, which is still a reasonable delay, though not polynomially bounded due to equation (11).

However, note that adding such a delay does not immediately guarantee location privacy against arbitrary attackers. This delay only prevents the generic attack we showed, and can be extended to any passive attacker, but it is not trivial to know whether it also automatically prevents active limited-adversary attacks. This issue is left for future work.

4. Location Private Construction

In this section we apply our results from the previous section to achieve a location private distance-bounding protocol for limited adversaries. The proposed protocol is based on the LPDB protocol [24]. We assume that the verifier $\mathcal{V}$ and the prover $\mathcal{P}$ share a secret key K. As in the LPDB protocol, we have two phases: the *initialisation phase* and the *distance-bounding phase*.

- **Initialisation Phase:** The prover $\mathcal{P}$ generates a random nonce $N_{\mathcal{P}}$ and sends it to the verifier $\mathcal{V}$. The verifier $\mathcal{V}$ generates a random nonce $N_{\mathcal{V}}$ and sends it to the prover $\mathcal{P}$. Both the prover and the verifier use as input the concatenation of the nonces $N_{\mathcal{P}}$ and $N_{\mathcal{V}}$ as input to a keyed pseudorandom function (f_K) and divide the output of the PRF into two parts, i.e.:

$$M||R_{\mathcal{P}} \rightarrow f_K(N_{\mathcal{P}}||N_{\mathcal{V}}).$$

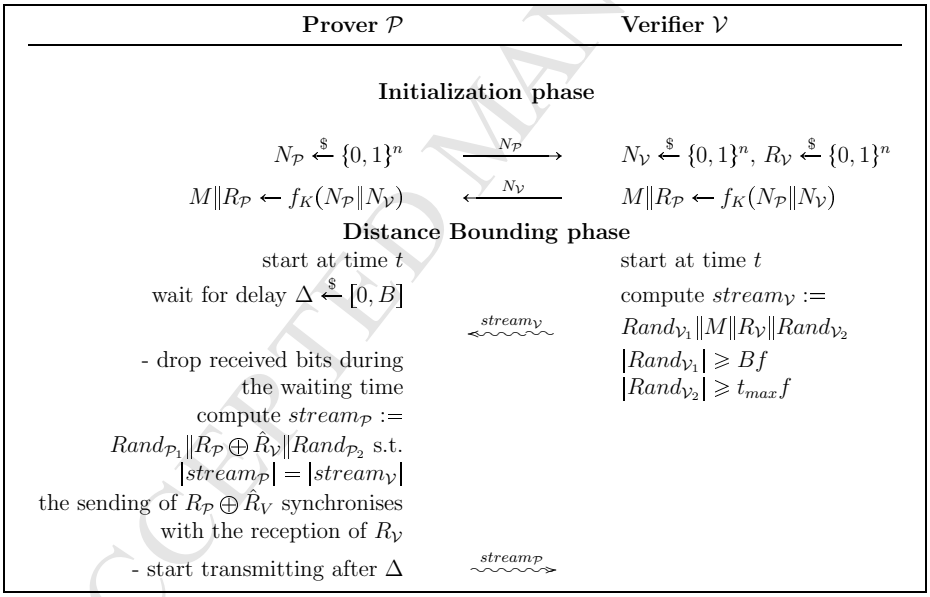

Figure 2: Proposed location-private distance-bounding protocol, secure against limited adversaries. Here $\xleftarrow{\$}$ denotes sampling uniformly at random, $\leftarrow$ denotes a simple message transmission, and $\leftsquigarrow$ denotes a continuous stream transmission at maximal bit rate.

Furthermore, $\mathcal{V}$ generates another random value $R_{\mathcal{V}}$ of length n.

- **Distance Bounding Phase:** Both the prover $\mathcal{P}$ and the verifier $\mathcal{V}$ start their actions at a commonly agreed time t. More precisely, at time t the verifier $\mathcal{V}$ starts transmitting the stream of bits $stream_{\mathcal{V}}$ such that: $stream_{\mathcal{V}} := Rand_{\mathcal{V}_1} \| M \| R_{\mathcal{V}} \| Rand_{\mathcal{V}_2}$. At time t the prover $\mathcal{P}$ starts waiting for a delay Δ that follows the uniform distribution with range $[0, B]$, where B satisfies the following condition as explained in section 3.2:

  $$B \geqslant 2^{h+1} t_{max}$$

  The prover $\mathcal{P}$ drops any bits received during the waiting time Δ. After this delay, the prover $\mathcal{P}$ starts transmitting the stream of bits $stream_{\mathcal{P}}$ such that:

  $$stream_{\mathcal{P}} := Rand_{\mathcal{P}_1} \| R_{\mathcal{P}} \oplus \hat{R}_{\mathcal{V}} \| Rand_{\mathcal{P}_2}$$

  where $\hat{R}_{\mathcal{V}}$ denotes the received value of $R_{\mathcal{V}}$ from the prover $\mathcal{P}$. The transmission of $R_{\mathcal{P}} \oplus \hat{R}_V$ must start as soon as $\mathcal{P}$ starts receiving the bits of $R_{\mathcal{V}}$.

  We note here that $Rand_{\mathcal{P}_1}$, $Rand_{\mathcal{P}_2}$, $Rand_{\mathcal{V}_1}$, $Rand_{\mathcal{V}_2}$ denote random values generated by the prover $\mathcal{P}$ and the verifier $\mathcal{V}$ respectively. Compared to the LPDB protocol [24], we further require that:

  $$|stream_{\mathcal{V}}| = |stream_{\mathcal{P}}| \text{ and } |Rand_{\mathcal{V}_1}| \geqslant Bf \text{ and } |Rand_{\mathcal{V}_2}| \geqslant t_{max} f.$$

  The verifier $\mathcal{V}$ could freely select the length of $Rand_{\mathcal{V}_1}$ and $Rand_{\mathcal{V}_2}$ satisfying these inequalities. It is easy to see that it holds:

  $$|Rand_{\mathcal{P}_1}| = |Rand_{\mathcal{V}_1}| + |M| + (t_{\mathcal{PV}} - \Delta) f,$$

  which is positive and

  $$|Rand_{\mathcal{P}_2}| = |Rand_{\mathcal{V}_2}| - (t_{\mathcal{PV}} - \Delta) f$$

  which is also positive.

4.1. Security of the Location Private Construction

We briefly sketch here the security proof for our new protocol.

Theorem 1. *For a passive limited adversary, if f is a PRF then:*

$$Adv_{\mathcal{DB}}^{\mathsf{LocPriv}}(\mathcal{A}) \leqslant 2^{-h} + negl$$

PROOF. Note that the maximal delay B is exponential in h due to equation (11). For a passive limited adversary $\mathcal{A}$, f_K can be replaced by a random function, then M and $R_{\mathcal{P}}$ can be assumed to be random. Then, the distribution of the view of the adversary $View_{\mathcal{A}}$ consists of $N_{\mathcal{P}}$, $N_{\mathcal{V}}$, $stream_{\mathcal{V}}$, $stream_{\mathcal{P}}$ and the time of reception of the two streams. The reception times of the first bits are $t_{\mathcal{V}}$ and $t_{\mathcal{P}}$. Since the streams have equal length, all other reception times can be obtained from $t_{\mathcal{V}}$ and $t_{\mathcal{P}}$.

We reduce the $\mathsf{LocPriv}$ game to a similar one where the PRF f is replaced by a random function. The difference between $\mathsf{Adv}_{\mathcal{DB}}^{\mathsf{LocPriv}}(\mathcal{A})$ and the new advantage Adv is negligible, thanks to the PRF property. Clearly, the messages are uniformly distributed.

The protocol belongs to Case 3 of assumption 2. Based on Lemma 5, we have:

$$\mathsf{Adv} \leqslant \frac{2t_{max}}{B} \leqslant 2^{-h}.$$

□

We should mention here that the security of the proposed protocol conforms with the bound given in Theorem 2 as already been proven for the LPDB protocol [24].

Theorem 2. *Assuming that f is a PRF, that $R_{\mathcal{V}}$ is uniformly distributed in a set of exponential size, that $R_{\mathcal{P}}$ is in a set of exponential size, the LPDB protocol [24] is a distance bounding protocol which provides resistance to distance fraud, and resistance to mafia fraud.*

5. Conclusions and Discussion

In this paper, we investigate the problem of location privacy in distance-bounding protocols. More precisely, we define a security game for location privacy in distance-bounding protocols and an adversarial model, composed of two classes of adversaries, an omniscient and a limited adversary. We prove that location privacy is information-theoretically impossible for any adversary of the two classes. In particular, a generic passive adversary can break the location privacy of any polynomial-time protocol. Nevertheless, we show that for limited adversaries, carefully chosen parameters enable computational, provable location privacy in practice. For those parameters we

propose a location private distance-bounding protocol based on the LPDB distance-bounding protocol [24].

We prove our results with respect to our game-based notion of location privacy, in which the communication between provers and verifiers takes place across a channel equipped with a timer. Adversaries may run man-in-the-middle attacks. They know their distance to the verifier, but not necessarily their distance to the prover. However, for each message of the protocol, the adversary learns the arrival time of the message, in a bitwise fashion. The goal of the adversary is to distinguish between two possible provers, which are within the proximity (associated with a bound $t_{\max}$) of the verifier.

In our model, we make two related, but distinct assumptions. The *first* is that the adversary is able to learn the (exact) time of arrival of messages at its interface. This is a reasonable assumption considering that in distance-bounding scenarios, the verifier has a clock that allows it to precisely measure the roundtrip transmission time (with a good granularity). In fact [29] describes an implementation of distance bounding, wherein the verifier pinpoints the location of a prover with a maximal distance error of 15 cm. An adversary has at least as much granularity in measuring the time of arrival as the verifier. Note that the more precise the adversary's clock is, the finer it can distinguish between two very close provers.

Our *second* assumption is that the transmission speed is constant and transmissions are (practically) collision-free. We equivalate transmission times with physical distances and assume that all bits sent by one party arrive at the others. Indeed, it is not unrealistic to model our attack in this way; an adversary can use the bits it *does* receive to correct any delays or errors. Once again, a reliable transmission only translates in a more fine-grained distinction for the adversary. The quality of the signal and the reliability of the channel depends on the hardware on which the protocol is deployed. Since most classical NFC and RFID hardware support reliable light speed transmission and run at very close proximity, our assumption accurately covers these scenarios.

The *omniscient* adversary model is very strong. We assume that the adversary may in fact represent a collusion of attackers, which can in fact triangulate signals. It is realistic to assume that such adversaries exist (e.g. governmental agencies and law enforcement institutions). Wireless transmissions are particularly vulnerable to triangulation. In this sense, our impossibility results state that one cannot stay off the radar and at the same time benefit from services requiring transmissions. However, it is still reassuring to

know that adding a large delay may at least prevent curious *limited* attackers from learning the sender's location.

Finally, we briefly comment on intermediate adversarial models. As mentioned in Section 2.2, an *omniscient* adversary can be realized either by a collusion of adversaries or by a single one who is able to move. Whereas it could make sense to consider intermediate adversary strengths, our results point out that location privacy is impossible to achieve in polynomial time even in the presence of the weakest adversaries we can define, i.e. *limited* ones.

Acknowledgment

This work was partially supported by the Marie Curie IEF Project "PPIDR: Privacy-Preserving Intrusion Detection and Response in Wireless Communications", Grant No. 252323.

We also thank the anonymous reviewers for their valuable and constructive comments.

References

[1] J.-P. Aumasson, A. Mitrokotsa, and P. Peris-Lopez. A Note on a Privacy-preserving Distance Bounding Protocol. In *Proceedings of the 13th International Conference on Information and Communications Security*. Springer, November 2011.

[2] A. Bay, I. Boureanu, A. Mitrokotsa, I. Spulber, and S. Vaudenay. The Bussard-Bagga and Other Distance Bounding Protocols under Man-in-the-Middle Attacks. In *Proceedings of Inscrypt'2012, 8th China International Conference on Information Security and Cryptology*, Lecture Notes in Computer Science, Beijing, China, 2012. Springer.

[3] I. Boureanu, A. Mitrokotsa, and S. Vaudenay. On the Pseudorandom Function Assumption in (Secure) Distance-Bounding Protocols - PRF-ness alone Does Not Stop the Frauds! In *LATINCRYPT*, volume 7533 of *Lecture Notes in Computer Science*, pages 100–120. Springer, 2012.

[4] I. Boureanu, A. Mitrokotsa, and S. Vaudenay. Practical and Provably Secure Distance-Bounding. In *the 16th Information Security Conference (ISC 2013)*, LNCS. Springer, 2013. To appear.

[5] I. Boureanu, A. Mitrokotsa, and S. Vaudenay. Secure & Lightweight Distance-Bounding. In *Proceedings of Second International Workshop on Lightweight Cryptography for Security & Privacy - LightSec 2013*, Gebze, Turkey, May 6-7 2013.

[6] S. Brands and D. Chaum. Distance-bounding Protocols. In *EUROCRYPT '93*, LNCS, pages 344–359. SPRINGER, 1993.

[7] M. Burmester. His Late Master's Voice: Barking for Location Privacy. In *Proceedings of Security Protocols Workshop*, pages 4–14, 2011.

[8] M. Burmester. Localization privacy. In D. Naccache, editor, *Cryptography and Security: From Theory to Applications*, volume 6805 of *Lecture Notes in Computer Science*, pages 425–441. Springer Berlin / Heidelberg, 2012.

[9] L. Bussard and W. Bagga. Distance-Bounding Proof of Knowledge Protocols to Avoid Terrorist Fraud Attacks. Technical Report RR-04-109, EURECOM, May 2004.

[10] N. Chandran, V. Goyal, R. Moriarty, and R. Ostrovsky. Position Based Cryptography. In *CRYPTO*, volume 5677 of *LNCS*, pages 391–407. Springer, 2009.

[11] Y. Desmedt. Major Security Problems with the Unforgeable' (Feige)-Fiat-Shamir Proofs of Identity and How to Overcome them. In *Proceedings of SecuriCom 1988*, pages 15–17. SEDEP Paris, France, 1988.

[12] C. Dimitrakakis, A. Mitrokotsa, and S. Vaudenay. Expected Loss Bounds for Authentication in Constrained Channels. In *Proceedings of INFOCOM 2012*, pages 478–485, Orlando, FL, USA, March 2012. IEEE press.

[13] S. Drimer and S. J. Murdoch. Keep your enemies close: distance bounding against smartcard relay attacks. In *Proceedings of 16th USENIX Security Symposium*, pages 7:1–7:16, Berkeley, CA, USA, 2007. USENIX Association.

[14] U. Dürholz, M. Fischlin, M. Kasper, and C. Onete. A Formal Approach to Distance Bounding RFID Protocols. In *Proceedings of the* 14^{th} *Information Security Conference ISC 2011*, LNCS, pages 47–62. SPRINGER, 2011.

[15] M. Fischlin and C. Onete. Subtle kinks in distance-bounding: an analysis of prominent protocols. In *6th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec) 2013*, pages 195 – 206. ACM, 2013.

[16] M. Fischlin and C. Onete. Terrorism in distance bounding: Modeling terrorist fraud resistance. In *Proceedings of the International Conference on Applied Cryptography and Network Security ACNS'13*, volume 7954 of *lncs*, pages 414 – 431. Springer, 2013.

[17] Ford. Safe and Secure *SecuriCode*^{TM} Keyless Entry. http://www.ford.com/technology/, 2011.

[18] A. Francillon, B. Danev, and S. Capkun. Relay attacks on passive keyless entry and start systems in modern cars. Cryptology ePrint Archive, Report 2010/332, 2010. EPRINTURL.

[19] L. Francis, G. Hancke, K. Mayes, and K. Markantonakis. Practical NFC peer-to-peer relay attack using mobile phones. In *Proceedings of the 6th international conference on Radio frequency identification: security and privacy issues*, RFIDSec'10, pages 35–49, Berlin, Heidelberg, 2010. Springer-Verlag.

[20] G. P. Hancke, K. E. Mayes, and K. Markantonakis. Confidence in Smart Token Proximity: Relay Attacks Revisited. *Computers & Security*, 28(7):404–408, October 2009.

[21] Y.-C. Hu, A. Perrig, and D. B. Johnson. Wormhole attacks in wireless networks. *IEEE Journal on Selected Areas in Communications*, 24:370–380, 2006.

[22] C. H. Kim, G. Avoine, F. Koeune, F.-X. Standaert, and O. Pereira. The Swiss-Knife RFID Distance Bounding Protocol. In *International Conference on Information Security and Cryptology – ICISC*, volume 5461 of *LNCS*, pages 98–115, Seoul, Korea, December 2008. SPR:full.

[23] A. Mitrokotsa, C. Dimitrakakis, P. Peris-Lopez, and J. C. H. Castro. Reid et al.'s distance bounding protocol and mafia fraud attacks over noisy channels. *IEEE Communications Letters*, 14(2):121–123, February 2010.

[24] A. Mitrokotsa, C. Onete, and S. Vaudenay. Mafia Fraud Attack against the RČ Distance-Bounding Protocol. In *Proceedings of the 2012 IEEE International Conference on RFID-Technology and Applications (IEEE RFID T-A 2012)*, 2012.

[25] A. Mitrokotsa, P. Peris-Lopez, C. Dimitrakakis, and S. Vaudenay. On selecting the nonce length in distance-bounding protocols. *The Computer Journal*, 56(10):1216–1227, 2013.

[26] K. Pelechrinis, C. Koufogiannakis, and S. V. Krishnamurthy. On the Efficacy of Frequency Hopping in Coping with Jamming Attacks in 802.11 Networks. *IEEE Transactions on Wireless Communications*, 9(10):3258–3271, October 2010.

[27] M. Poturalski, P. Papadimitratos, and J.-P. Hubaux. Secure neighbor discovery in wireless networks: formal investigation of possibility. In *Proceedings of the 2008 ACM symposium on Information, computer and communications security*, ASIACCS '08, pages 189–200, New York, NY, USA, 2008. ACM.

[28] K. Rasmussen and S. Čapkun. Location Privacy of Distance Bounding. In *Proceedings of the Annual Conference on Computer and Communications Security (CCS)*. ACM, 2008.

[29] K. B. Rasmussen and S. Čapkun. Realization of rf distance bounding. In *Proceedings of the 19th USENIX Conference on Security*, USENIX Security'10, pages 25–25, Berkeley, CA, USA, 2010. USENIX Association.

[30] J. Reid, J. M. Gonzalez Nieto, T. Tang, and B. Senadji. Detecting Relay Attacks with Timing-based Protocols. In *ASIACCS '07: Proceedings of the 2nd ACM Symposium on Information, Computer and Communications Security*, pages 204–213, Singapore, March 2007. ACM.

[31] N. Sastry, U. Shankar, and D. Wagner. Secure Verification of Location Claims. In *Proceedings of the 2nd ACM Workshop on Wireless Security (WiSe'03)*, pages 1–10, 2003.

[32] D. Singelée and B. Preneel. Location Verification Using Secure Distance Bounding Protocols. In *Proceedings of the IEEE International Confer-*

ence on Mobile Adhoc and Sensor Systems (MASS'05), pages 834–840, 2005.

[33] D. Spil and A. Bittau. Bluesniff: Eve Meets Alice and Bluetooth. In *Proceedings of the 1st USENIX Workshop on Offensive Technologies (WOOT'07)*. USENIX Association Berkley, CA, USA, 2007.