Realization of RF Distance Bounding

Kasper Bonne Rasmussen
Department of Computer Science
ETH Zurich
8092 Zurich, Switzerland
kasperr@inf.ethz.ch

Srdjan Čapkun
Department of Computer Science
ETH Zurich
8092 Zurich, Switzerland
capkuns@inf.ethz.ch

Abstract

One of the main obstacles for the wider deployment of radio (RF) distance bounding is the lack of platforms that implement these protocols. We address this problem and we build a prototype system that demonstrates that radio distance bounding protocols can be implemented to match the strict processing that these protocols require. Our system implements a prover that is able to receive, process and transmit signals in less than $1ns$. The security guarantee that a distance bounding protocol built on top of this system therefore provides is that a malicious prover can, at most, pretend to be about $15cm$ closer to the verifier than it really is. To enable such fast processing at the prover, we use specially implemented concatenation as the prover's processing function and show how it can be integrated into a distance bounding protocol. Finally, we show that functions such as XOR and the comparison function, that were used in a number of previously proposed distance bounding protocols, are not best suited for the implementation of radio distance bounding.

1 Introduction

Distance bounding denotes a class of protocols in which one entity (the verifier) measures an upper-bound on its distance to another (untrusted) entity (the prover). In recent years, distance bounding protocols have been extensively studied: a number of protocols were proposed [3, 13, 10, 19, 30, 15, 25, 17, 12, 29] and analyzed [8, 26, 11, 23]. The use of distance bounding was suggested for secure localization [28], location verification [25], wormhole detection [16, 27], key establishment [22, 32] and access control [22].

Regardless of the type of distance bounding protocol, the distance bound is obtained from a rapid exchange of messages between the verifier and the prover. The verifier sends a challenge to the prover, to which the prover replies after some processing time. The verifier measures the round-trip time between sending its challenge and receiving the reply from the prover, subtracts the prover's processing time and, based on the remaining time, computes the distance bound between the devices. The verifier's challenges are unpredictable to the prover and the prover's replies are computed as a function of these challenges. In most distance bounding protocols, a prover XORs the received challenge with a locally stored value [3] or uses the received challenge to determine which of the locally stored values it will return [13, 29]. Thus, the prover cannot reply to the verifier sooner than it receives the challenge, it can only delay its reply. The prover, therefore, cannot pretend to be closer to the verifier than it really is; only further away.

One of the main assumptions on which the security of distance bounding protocols relies is that the time that the prover spends in processing the verifier's challenge is negligible compared to the propagation time of the signal between the prover and the verifier. If the verifier overestimates the prover's processing time (i.e., the prover is able to process signals in a shorter time than expected), the prover will be able to pretend to be closer to the verifier. If the verifier underestimates this time (i.e., the prover needs more time to process the signals than expected), the computed distance bounds will be too large to be useful.

The challenge in implementing distance bounding protocols is therefore to implement a prover that is able to receive, process and transmit signals in negligible time. This requirement can be easily met with ultrasonic distance bounding implementations where the prover's processing needs to be in the order of μs. However, because ultrasonic distance bounding is vulnerable to RF wormhole attacks [16, 27],

its application is limited to few specific applications (e.g., [22]). For most applications, radio distance bounding is the main viable way of verifying proximity to or a location of a device. In this case, the prover's processing time needs to be about $1ns$ which would, in the worse case, allow a malicious prover to pretend to be closer to the verifier by approx. $15cm$ (assuming that the malicious prover is able to process signals instantaneously). Currently available platforms do not support such fast processing. This strict processing requirement has been, so far, one of the main obstacles for the wider deployment of RF distance bounding protocols and related solutions.

In this work, we address this problem. We make the following contributions. We build a prototype system that demonstrates that radio (RF) distance bounding protocols can be implemented to match the prover's strict processing requirements (i.e., that the prover's processing time is below $1ns$). We use concatenation as the prover's processing function and implement it using a scheme that we call Challenge Reflection with Channel Selection (CRCS). Our implementation eliminates the need for signal conversion and demodulation since it does not require that the received challenges are interpreted by the prover before the prover responds to them. Our prover is therefore able to receive, process and transmit signals in less than $1ns$. We design a distance bounding protocol that uses concatenation, implemented with CRCS, as the prover's processing function and we analyze its security; we base this protocol on Brands and Chaum's original distance bounding protocol [3].

We further show that processing functions such as XOR and the comparison function, that were used in a number of proposed distance bounding protocols, are not best suited for the implementation of radio distance bounding. The main reason is that, although XOR and comparison can be executed fast, these functions require that the radio signal that carries the verifier's challenge is demodulated, which, with today's state-of-the-art hardware, results in long processing times (typically $\geq 50ns$). The design and implementation of the distance bounding protocol based on concatenation shows that the use of functions which require that the prover demodulates (interprets) the verifier's challenge before responding to it is not necessary for the implementation of distance bounding.

To our knowledge this work is the first to propose a realizable distance bounding protocol using radio communication, with a processing time at the prover that is low enough to provide a useful distance granularity.

The rest of the paper is organized as follows. In Section 2 we describe the basic operation of distance bounding protocols. In Section 3, we discuss prover's processing functions and their appropriateness for the implementation of radio distance bounding. In Section 4 we describe the design of our distance bounding protocol and in Section 5 we analyze its security. In Section 6 we present our implementation and our measurement results. In Section 7 we discuss related work and we conclude in Section 8.

2 Background on Distance Bounding Protocols

Distance bounding protocols were first introduced by Brands and Chaum [3] for the prevention of mafia-fraud attacks on Automatic Teller Machines (ATMs). The purpose of Brands and Chaum's distance bounding protocol was to enable the user's smart-card (verifier) to check its proximity to the legitimate ATM machine (prover).

The core of all distance bounding protocols is the distance measurement phase (shown in Figure 1), wherein the verifier measures the round-trip time between sending its challenge and receiving the reply from the prover. More precisely, the verifier challenges the prover with a b-bit freshly generated nonce N_v (typically $b = 1$). Upon reception of the challenge, the prover computes a response $f^P(N_v)$, and sends it to the verifier. This process is repeated k times. After the challenge-response exchange the verifier verifies the authenticity of the replies (in this step distance bounding protocols differ) and measures the time $t_s^V - t_r^V$ between the challenge and the response. Based on the measured times, the verifier estimates the upper-bound on the distance to the prover. The time $t_s^P - t_r^P$ between the reception of the challenge and the transmission of the response at the prover is either negligible compared to the propagation time $t_r^P - t_s^V$ or is lower bounded by the prover's processing and communication capabilities δ, i.e., $t_s^P - t_r^P \geq \delta$.

After the execution of a distance bounding protocol the verifier knows that the prover is within a certain distance, namely:

$$dist = \frac{t_s^V - t_r^V - \delta}{2} \cdot c$$

where δ is the processing time of the prover (ideally 0) and c is the propagation of the radio signal.

Although the designs of distance bounding protocols differ [3, 13, 10, 19, 30, 15, 25, 17, 12, 29], given their common distance measurement phase,
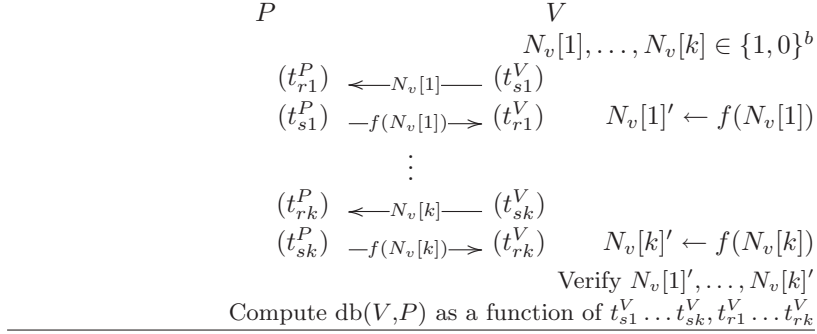
$$
\begin{array}{ll}
P & V \\
 & N_v[1], \ldots, N_v[k] \in \{1,0\}^b \\
(t_{r1}^P) \quad \xleftarrow{\;\;N_v[1]\;\;} (t_{s1}^V) & \\
(t_{s1}^P) \quad \xrightarrow{\;f(N_v[1])\;} (t_{r1}^V) & N_v[1]' \leftarrow f(N_v[1]) \\
\qquad\qquad \vdots & \\
(t_{rk}^P) \quad \xleftarrow{\;\;N_v[k]\;\;} (t_{sk}^V) & \\
(t_{sk}^P) \quad \xrightarrow{\;f(N_v[k])\;} (t_{rk}^V) & N_v[k]' \leftarrow f(N_v[k]) \\
 & \text{Verify } N_v[1]', \ldots, N_v[k]' \\
\end{array}
$$
$$\text{Compute db}(V,P) \text{ as a function of } t_{s1}^V \ldots t_{sk}^V, t_{r1}^V \ldots t_{rk}^V$$

Figure 1: The distance measurement phase of distance bounding protocols consists of a rapid exchange of messages where the verifier measures the round-trip time between sending its challenges and receiving the replies from the prover.

their security relies on the same underlying ideas. We briefly summarize them here. Distance fraud attacks [3], in which the prover tries to pretend to be closer to the verifier, are prevented by the following: (i) the prover cannot generate the reply before it receives the challenge and (ii) the duration of time the verifier accounts that the prover will process the reply is not longer than the prover's actual processing time. The Mafia-fraud (or man-in-the-middle - MITM) attack [9], by which an attacker convinces the verifier that the prover is closer than it really is, is prevented since the attacker cannot predict exchanged challenges/replies and since it cannot speed-up the propagation of messages (the messages propagate at the speed of light over a radio channel). Given this, the attacker cannot shorten the distance measured between the verifier and the prover.

Distance bounding protocols therefore provide the verifier with an upper-bound on its physical distance to the prover.

3 Functions Appropriate for Distance Bounding Realization

As discussed in Section 2, one of the main assumptions on which the security of distance bounding protocols relies is that the time that the prover is allowed to spend in processing the verifier's challenge is negligible compared to the propagation time $t_r^P - t_s^V$ of the signal between the prover and the verifier. In most applications, the prover's processing time would therefore need to be around $1ns$. This would, in the worse case, allow a malicious prover to pretend to be closer to the verifier by approx. $15cm$ (assuming that the malicious prover is able to process signals instantaneously). Such short processing time cannot be achieved with existing platforms.

The main challenge is therefore to design distance bounding protocols which use prover processing functions $f(N_v)$ that can implemented such that they can be executed in $\leq 1ns$. Before presenting a function that is well suited for this purpose, we first discuss functions that were used in distance bounding protocols that are proposed in the open literature.

The first (obvious) candidate processing functions are various encryption functions, hash functions, message authentication codes and digital signatures; the use of digital signatures for this purpose was proposed by Beth and Desmedt in [1]. The use of such functions would largely simplify the design of distance bounding protocols; it would be sufficient to use well studied challenge-response authentication protocols [2] where the verifier would measure the round-trip time between the issued challenge and the received response. However, the processing time for these functions even with the fastest available implementations by far exceeds the required processing time.

In [3] Brands and Chaum proposed a distance bounding protocol that uses XOR as a processing function. In this protocol the prover XORs the verifiers challenge with the value that the prover wants to transmit back and sends the result back to the verifier. The main reasoning behind this choice was that XOR is a fast operation and that it should be feasible to execute it within the required processing time. Hancke and Kuhn [13] propose a distance bounding protocol where the prover, based on the verifier's challenge chooses from which of the two local registers it should send a value back. Again, one of the main reasons for choosing this function was that such a function (comparison and access) can be executed fast.

Although XOR and comparison can be executed fast, these functions require that the radio signal that carries the verifier's challenge is converted from an analog to a digital signal (ADC) and demodulated. Only when it is demodulated, the challenge can be used by the prover in an XOR function or for the selection of the register. Equally, in order to communicate the reply back to the verifier, the prover needs to modulate the signal and convert it from the digital to the analog signal (DAC). These steps, signal detection, ADC/DAC conversion and signal modulation/demodulation, increase the provers processing delay by approx. 170ns [24], not including possible RX/TX switching costs[1]. The implementations of an XOR or of a comparison function that require the signals to be digitalized and demodulated therefore require such processing which, using today's state-of-the-art hardware, is not sufficiently fast to meet the security requirements of distance bounding protocols. Even if some processing steps can be sped-up or removed, the prover will still need a way of (reliably) detecting if it received a challenge that corresponds to a bit "0" or a bit "1", which requires some processing and thus reduces the security guarantees of the protocol. Namely, every nanosecond of additional processing in the implementation of the prover means that a malicious prover with a faster implementation shorten the measured distance even further.

In what follows, we show that the choice of a concatenation function as the prover's processing function, when implemented using a scheme that we call *Challenge Reflection with Channel Selection (CRCS)* eliminates the need for signal conversion and demodulation since it does not require that the received challenges are interpreted by the prover before the prover responds to them. The prover, implemented using CRCS is therefore able to receive, process and transmit signals in less than $1ns$.

3.1 Prover: Concatenation Implemented using Challenge Reflection With Channel Selection

In this section we describe our implementation of concatenation as the prover's processing function.

Bit concatenation $CAT : N_p[i] \times N_v[i] \to r[i] = N_v[i]||N_p[i]$ takes as input the verifier's challenge bit $N_v[i]$ and the prover's input bit $N_p[i]$ and returns a two-bit reply $r[i] = N_v[i]||N_p[i]$. CAT is therefore

---

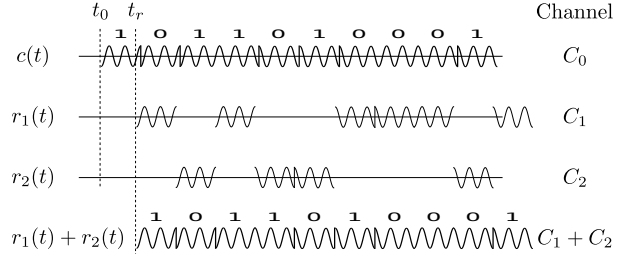[1]We are not aware of the radio design that can perform these operations faster.

Figure 2: The verifier measures the time between sending a challenge signal $c(t)$ and receiving the reply signal $r(t) = r_1(t) + r_2(t)$. If $c(t) = r(t)$, the distance bound to the prover is then given by $(t_r - t_0) \cdot c$, where c is the speed of light.

given by the following table.

$$CAT : \quad \begin{array}{c|cc} N_p[i]\backslash^{N_v[i]} & 0 & 1 \\ \hline 0 & 00 & 10 \\ 1 & 01 & 11 \end{array}$$

3.2 Verifier: Calculation of the Distance Bound

In order for concatenation to be useful for distance bounding, we implement it by Challenge Reflection with Channel Selection. Our implementation uses three (non-overlapping) communication channels. The verifier sends its challenge bits to the prover using one communication channel (C_0), whereas the prover replies using two communication channels (C_1, C_2) (Figure 2). While it is receiving the verifier's challenge bit (i.e., the signal that encodes it), the prover is responding with the same signal (bit), but it is sending it on either channel C_1 or channel C_2, depending on its current input bit $N_p[i]$. For every challenge bit that it received from the verifier, the prover therefore transmits two bits of the reply back to the verifier, encoded in the form of the signal (it reflect back the same signal that it received) and of the response channel (it chose the channel on which to reply). The response $r = 10$ is then interpreted as: the challenge bit 1 is reflected on channel C_1, where the channel C_1 denotes bit 0, and channel C_2 denotes bit 1). The prover therefore implements challenge reflection with channel selection. Note that, although the prover replies with two bits for each challenge bit, the duration of transmission of those two bits is the same as for a single bit of the verifier's challenge, since the second bit of the prover's reply is encoded in the form of channel selection. This is illustrated on Figure 2.

The schematic of our prover implementing CRCS is shown on Figure 3. The figure shows the signal in
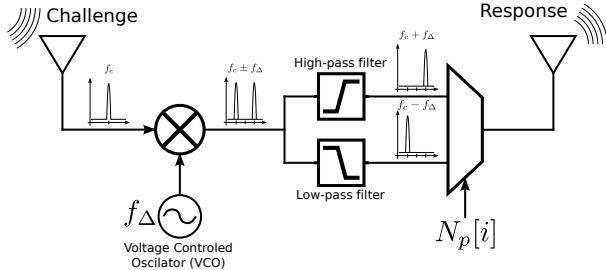
Figure 3: Schematic of the prover (i.e., of the implementation of concatenation as its processing function using CRCS). The figure shows the signal in the frequency domain at various stages of the circuit. The challenge-signal (with center frequency f_c) is received by the receiving antenna (on the left) and multiplied by f_Δ. This multiplication shifts the signal by $\pm f_\Delta$ to the channels on two sides of the original channel. The bit of the prover's nonce $N_p[i]$ determines which of the two channels is used to send the response on the transmitting antenna (on the right).

the frequency domain as it passes through various stages of the prover's circuit. The prover receives the challenge-signal (centered at the frequency f_c) on the receiving antenna. The received signal is then multiplied by f_Δ which creates two signals on two channels each with central frequencies $f_c + f_\Delta$ and $f_c - f_\Delta$, respectively. The current bit of the prover's nonce $N_p[i]$ determines which of the two channels are used to send the response signal on the transmitting antenna. The verifier's signal is thus reflected back on the channel selected by the prover. Here, the verifier's challenge bit can be encoded in the challenge signal using e.g., Pulse Amplitude Modulation (PAM) or Binary Phase Shift Keying Modulation (both of which are used with Ultra-Wide-Band ranging systems). The prover's response carries two bits, one encoded in the signal that it sends back (the same bit that it received by the verifier), and the other encoded in the channel on which it responds (i.e., $N_p[i]$).

Here, signal multiplication and selection are done using analog components only. Namely, the challenge signal passes through an analog mixer where it is multiplied with a local oscillator signal with a frequency f_Δ. This mixer outputs two signals on frequencies $f_c + f_\Delta$ and $f_c - f_\Delta$, which are separated by a high-pass and a low-pass filter, respectively. Finally, the $N_p[i]$ bit (which the prover have committed to), determines which of the two signals will be transmitted back to the verifier.

Figure 2 shows the calculation of the distance

bound by the verifier (the signals are shown in the time domain). The verifier notes the exact time t_0 when it starts transmitting the challenge bits $N_v[i],...N_v[k]$ encoded in the signal $r_1(t)$, and then listens on the two reply channels C_1 and C_2 (that correspond to the frequencies $f_c + f_\Delta$ and $f_c - f_\Delta$). When a reply comes back (e.g., on channel C_1) the verifier will mark the exact time t_r of the arrival of the signal. The verifier will then wait for the arrival of the entire challenge, noting for every time slot on which channel the reply was sent. After the entire nonce has been received and processed by the radio, the verifier checks that the data bits in the reply are the same as those sent in the challenge, i.e., that $c(t) = r_1(t) + r_2(t)$. If that is the case, the distance bound is then computed as $(t_r - t_0) \cdot c$, where c is the speed of light. This bit comparison is important for the security of our distance bounding protocol (as we detail in Section 4); it can be efficiently done using autocorrelation, which can then simultaneously be used to calculate the time difference (e.g., as it is used in GPS [20]).

4 Distance Bounding Realization

In this section we present our distance bounding protocol and its realization. The protocol uses concatenation implemented using CRCS as the prover's processing function. The main security properties that we want our protocol to achieve are resilience to distance fraud and Mafia fraud attacks.

Our protocol is shown in Figure 4. It closely resembles the original protocol of Brands and Chaum [3], except that it does not use rapid bit exchange, but instead uses full duplex communication with signal streams. XOR is replaced with the concatenation function, and additional checks by the prover and the verifier are added to make sure the implementation of concatenation using CRCS does not introduce vulnerabilities.

The prover starts the protocol by picking a fresh (large) nonce N_p. The prover then sends a commitment to the nonce (e.g., a signed hash of the nonce) to the verifier. Already now, the prover will activate its distance bounding hardware and set the output channel according to the opposite of the first bit of the nonce N_p. From this moment, any signal that the prover receives on channel C_0 will be reflected on the output channel that is set. However, the prover does not yet start switching between output channels.

Upon receiving the commitment, the verifier picks a fresh (large) nonce N_v and prepares to initiate the distance bounding phase in which it will measure
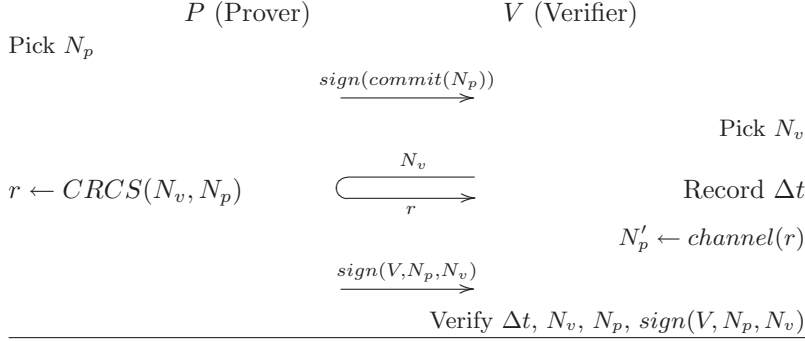
P (Prover)		V (Verifier)
Pick N_p		
	$\xrightarrow{sign(commit(N_p))}$	
		Pick N_v
$r \leftarrow CRCS(N_v, N_p)$	$\xrightarrow[\ r\ ]{N_v}$	Record Δt
		$N_p' \leftarrow channel(r)$
	$\xrightarrow{sign(V,N_p,N_v)}$	
		Verify Δt, N_v, N_p, $sign(V, N_p, N_v)$

Figure 4: RF distance bounding protocol.

the distance bound to the prover. The verifier starts a high precision clock to measure the (roundtrip) time of flight of the signal and begins to transmit his nonce N_v on channel C_0. From this point on, the verifier will also listen on the two reply channels C_1 and C_2 and will keep listening on the two channels until he either receives the expected response from the prover or until he detects an error and aborts the protocol.

As soon as the prover receives (and demodulates) the first bit of N_v on C_0, he starts switching reply channels according to the bits of his nonce N_p. Here, we note that while the first few bits are being demodulated, the prover is still reflecting the input (challenge) bits, but he did not start the switching of the channels (i.e., he did not start sending back N_p). The demodulation of the bits is not done within the distance bounding hardware (that we call the distance bounding extension), but is done in the prover's regular radio. It is not important how long it takes for the prover's radio to demodulate the first bits, since the prover does not need to begin to switch the output channels within any predefined time (as long as the switching starts within the duration of N_v and allows the transmission of N_p). Equally, the first part of N_v could be known and constitute a public, fixed-length preamble upon the detection of which the prover would start switching the channels (i.e., would start sending N_p).

When the prover starts sending N_p, he will send the bits of N_p with a fixed frequency (e.g., every $500ms$) by switching channels depending on the value of the current bit (Figure 2). In each interval, the prover will therefore reflect back several bits of N_v and a single bit of N_p. The bit of N_p is encoded in the choice of the reply channel. The prover will, in parallel, also receive the challenge on channel C_0 using his regular radio and will demodulate it.

When the verifier has sent all the bits of his nonce, he waits for the prover to complete the reflection of

the signal and then both the prover and verifier disable their distance bounding extensions. The verifier can then use an auto-correlation detector like the ones used in GPS receivers [20] to determine the exact time of flight of the reflected signal. This can also be done during the distance bounding phase, i.e., in parallel to the analog distance bounding circuit.

After the (time-critical) distance bounding phase is complete the prover sends a signed message containing his nonce N_p, the identity of the verifier V and the verifier's nonce N_v to the verifier. The verifier must then check five things:

- That all the bits of N_p reflected by the prover are of the same width (time duration). This is necessary to prevent mafia fraud and is described in more detail in Section 5.3.

- The data that was reflected back from the prover must be *exactly* the same as what was sent. I.e., when the signal $r(t) = r_1(t) + r_2(t)$ is demodulated, the message must contain N_v. This is visualized in Figure 2.

- The value of N_p' obtained during the distance bounding phase must match the commitment sent in the first protocol message.

- The signature of the final message must be valid and it must correspond to the expected identity of the prover.

- The time of flight of the signal Δt must be less than some predefined upper limit t_{max}. The upper limit is application dependent. E.g., it can be the radius of some region of interest, or it can be the (estimated) maximum transmission range of the radio.

The order is which these checks are performed is not important but all checks must pass for the distance bound to be accepted. If all the checks pass,

6

the verifier calculates the distance to the prover as

$$d = \frac{\Delta t - \delta_p}{2} \cdot c \qquad (1)$$

Where c is the speed of light and δ_p is the very small processing delay of the prover. In our implementation $\delta_p < 1ns$ resulting in a maximum error on about $15cm$.

5   Security Analysis

In this section we analyze the resistance of our protocol to distance fraud and mafia fraud, as well as attacks against CRCS.

5.1   System And Attacker Model

We consider three nodes, the prover P, the verifier V and the attacker M. The goals for the three participants are as follows: the verifier wants to acquire an upper bound on the distance to the prover, i.e., the verifier wants to know that the prover is closer than a certain distance. The prover wants to prove to the verifier that he is within a certain distance. The goal of the attacker is to disrupt this process such that the verifier obtains an incorrect distance bound. The verifier holds an authentic public key of the prover. The attacker and the prover do not collude. The attacker corresponds to the standard Dolev-Yao attacker that controls the network and thus can eavesdrop on all the communication between the prover and the verifier, can arbitrary insert and remove messages to/from the communication channel. She is equally free to transmit nonsensical signals. The attacker knows the public parameters of the distance bounding protocol and the type of hardware used by the nodes and thus the processing times of the prover's and verifier's radios. She is only limited by the fact that it does not have access to the secrets that are held by the prover and the verifier and cannot break cryptographic primitives.

We consider two attacks: Distance fraud, where the prover tries to shorten the measured distance bound, and Mafia fraud where the attacker tries to shorten the bound (but does not collude with the prover). We show that our protocol resists both attacks. There is a third type of attack in which the attacker colludes with the prover and has access to some, but not all, of the secret key material of the prover (e.g., only nonces and short-term secrets). This attack is often called the terrorist attack. We do not specifically address terrorist attacks, but it has been shown [4] that if needed, distance bound-

ing protocols can be extended to generally protect against this attack.

5.2   Distance Fraud

Distance fraud is an attack performed by a malicious prover and consists of the prover trying to shorten the distance measured by the verifier.

The verifier uses equation (1) to calculate the distance to the prover. For the prover to "shorten" the distance to the verifier (without actually moving closer) he must manipulate the verifiers calculation and the only thing the prover can influence is Δt. For the prover to reduce the Δt measured by the verifier, thereby reducing the distance, he must make his replies arrive at the verifier sooner than they otherwise would, i.e., he must guess the correct reply (i.e., guess the challenge) and send it before the verifier expects. In our protocol, the reply which the prover must send back is the signal he receives on channel C_0. In order to do this, the prover must guess the content of the challenge signal since the content of the reply is checked by the verifier as a part of the verification process. The content of the challenge is N_v and the probability of successfully guessing that is given by $\frac{1}{2^{|N_v|}}$.

Attacks that rely on manipulation of the modulation scheme, e.g., "late commit" attacks described by Hancke and Kuhn [14] will not work on this protocol because the verifier uses auto-correlation to find the exact time-of-flight of the signal (as it is done in GPS receivers [20]) rather than using a peak or energy detector. This means that any manipulation done to, say, the first symbol of the response will not have any effect unless all subsequent symbols are also shifted forward. This would require the malicious prover to guess all the symbols in advance and can therefore only be done with negligible probability of $\frac{1}{2^{|N_v|}}$.

The same argument applies to attacks where the prover tries to guess the first bit of the nonce [8]. Because the prover doesn't store and forward the nonce, but instead must reflect it directly, the prover would have to guess all the bits of the verifier's nonce to perform the attack. We can therefore conclude that the prover can commit distance fraud only with probability $\frac{1}{2^{|N_v|}}$.

5.3   Mafia Fraud

Mafia fraud is an attack performed by an external attacker that physically resides closer to the verifier than the prover. The attack aims to make one of the parties (either the prover or the verifier or both) believe that the protocol was successfully executed

when, in fact, the attacker shortened the distance measurement. The requirement that the attacker be closer to the verifier than the prover is only necessary because, if the attacker is further away the attack is trivially defeated by the protection against distance fraud attacks.

In order for an external attacker to shorten the distance measured by the verifier, the attacker must respond before the prover during the distance bounding phase. However, because of the checks performed by the verifier at the end of (or during) the distance bounding phase, it is not sufficient to just reply before the prover, the attacker must also make the value of his nonce match the commitment sent by the prover in the beginning of the protocol. Since the attacker can not find a nonce to match the commitment sent by the prover, e.g., find a collision for the hash function used to generate the commitment, the attacker is forced to replace the provers commitment with his own, thereby passing the commitment check. However, the attacker cannot fake the prover's signature in the final message so he cannot confirm the nonce.

The attacker can get the prover to reply before the prover receives N_v, e.g., by sending his own early signal to the prover, however, this will result in the prover getting $N_v' \neq N_v$ which will be detected by the verifier in the final message. This assumes that any malicious change to the signal will result in a change in the demodulated nonce N_v. If that cannot be guarantied, e.g., because of the sample rate at the prover or the modulation scheme used for communication, the prover can record the raw incoming signal and send it back to the verifier. The verifier can then, e.g., use autocorrelation to make sure the signal received by the prover is the same as what the verifier sent.

We can therefore conclude that an attacker can only commit mafia fraud if he can break, either the commitment scheme or the signature scheme used in the protocol.

Because of the way the distance bounding radio extension is designed it is possible for an attacker to get the current bit of the provers nonce. As explained in Section 3.1, the prover's radio extension will shift any signal that arrives on the center channel to either channel C_1 or channel C_2 depending on the current bit of the provers nonce. An attacker can exploit this to get the current bit of the prover's nonce without the prover's knowledge. If the attacker sends a very weak signal, e.g., a DSSS [21] signal with a spreading code known only to the attacker, the attacker can determine what channel the response is sent back on, and therefore the current
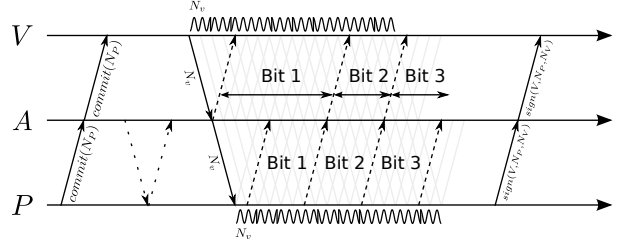

Figure 5: Man in the middle attack (Mafia fraud). The figure shows the timing of the messages sent by the verifier (V), the attacker (M) and the prover (P). Even if the attacker is able to learn the value of the first bit on the prover's nonce, the attack will fail because the attacker is forced to make the first bit longer than the subsequent bits if he wants to reply early.

bit of the prover's nonce. Unless this is prevented, the attacker can use this information to perform a successful mafia fraud attack.

In order to prevent this attack the prover must make sure not to expose all the bits of his nonce before they are needed. There are two ways this can be ensured: Either the prover must only enable his distance bounding hardware once he is sure that the verifier has started his transmission or he must make sure that his reply bits (of N_p) are of exactly the same duration. Of course the time duration must also be known and later checked by the verifier. Our protocol uses the second method. Figure 5 illustrates how this measure prevents the attack. In the example of this figure the attacker obtains the value of the first bit of the provers nonce, and uses it to reply early to the verifier's challenge. However, because the prover doesn't expose the second bit of his nonce until after the duration of the first bit has expired, the attacker is forced to make the first bit 'too long', thus getting detected.

In order to perform this attack, the attacker would need to guess all the bits of N_p, which she can do only with the probability $\frac{1}{2^{|N_p|}}$.

6 Implementation and Measurements

In this section, we describe our implementation of the prover and the related measurement results.

Our prototype can be seen on Figure 6. The central part of the prototype is the mixer (1) which is responsible for shifting the received challenge up and down in frequency. The signal from the receiving antenna comes in from the right (A) and passes
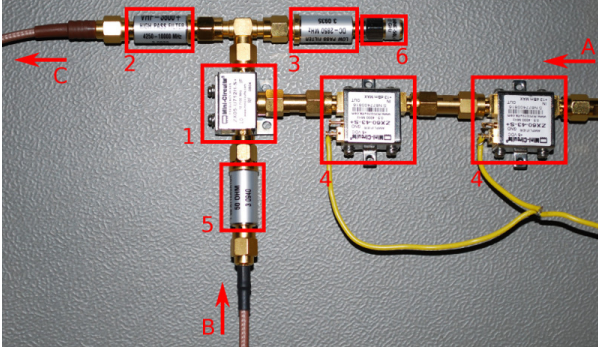
Figure 6: This picture shows the prototype implementation of the prover. It consists of a mixer (1), a high-pass filter (2), a low-pass filter (3), four amplifiers (4) (only two visible), a 1dB attenuator (5) and a terminating resistor (6). The signal from the receiving antenna (A) is mixed with the local oscillator (B) and sent to the transmitting antenna (C). The yellow wires are power (+5V). This prototype is an implementation of the scheme described in Figure 3.

through four amplifiers (4) to bring it up to a power level where is can be mixed by our mixer. The local 500MHz sine wave used for the mixing, comes in from the bottom of the figure (B) and is passed through a 1dB attenuator (5) to bring it to the same level as the radio signal before mixing. The output of the mixer is split in two and each is passed through either a high-pass filter (2) or a low-pass filter (3) to eliminate the unwanted channel. In this prototype we did not implement the switching mechanism. Instead channel C_2 is fed directly to the transmission antenna (C). In order for the signal to split properly, both sides must have a similar load. for this reason we added a 50Ω resistor (6) to terminate the unused channel C_1. The implementation of the switching mechanism can be done using a simple transistor based switch. We note, that the switch can only marginally increase the processing delay since, once set to a particular channel, the switch essentially acts as a piece of very short wire connecting the setup to the antenna. This prototype is an implementation of the scheme described in Figure 3.

6.1 Delay At The Prover

We first wanted to see if our prototype implementation could receive a signal, shift it to another channel and transmit it back to the verifier in $\leq 1ns$.

In order to test this, we first transmit the challenge and response signals through cables so as to better be able to control signal strength and reduce
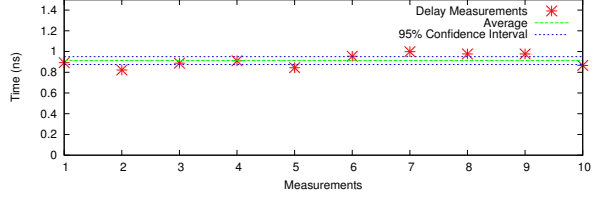

Figure 8: Processing time at the prover. The ten different delay measurements where done using our measurement setup described in Section 6.1. The figure shows that the variation in processing time is small ($\sigma = 61.22ps$) and that the average processing delay is $\mu = 912.92ps$. I.e., less than $1ns$.

noise (later we show that the same setup works using wireless communication as well). The challenge signal sent on channel C_0 is a 3.5GHz sine, modulated by a 1Hz pulse so it is easy to see and capture the start of a new "bit". Our response signal is sent back on channel C_2 at 4.0GHz (i.e., $f_c = 3.5GHz$ and $f_\Delta = 0.5GHz$). We generated the 3.5GHz challenge using a function generator. The generated signal is split by a power splitter and one end is fed, via a 1 meter cable, into our prototype. The other end was connected to a 40Gs/s oscilloscope, via another 1 meter cable, to provide the ground truth signal to which we compare the delay of our prototype. Because both cables have the same length, the 3.5GHz signal (the challenge) will arrive at the same time at the oscilloscope and at the reception point of our prototype. The output (the response) from the prototype is plugged directly into another input of the same oscilloscope (keeping the signal path as short as we could make it using this setup).

Figure 7(a) shows the two signals. The top (yellow) signal is coming directly from the function generator. It is an exact copy of the signal that arrives at the input of our prototype (this signal arrives at the oscilloscope and at the prototype input at the same time). The bottom (green) signal is what comes out of our prototype implementation. It is a 4.0GHz signal, i.e., the original signal shifted up by 500MHz. We see that the difference in arrival times between these two signals (i.e., the processing time of the prover) is $0.888ns$. As described in Section 2 the delay at the prover determines the theoretical advantage a powerful attacker might get. If we translate $0.888ns$ into distance, the maximum theoretical distance by which an attacker will be able to shorten its distance is about $12cm$.

We repeated this measurement 10 times, using the same setup. Figure 8 shows all 10 measured processing times along with their average value and a 95%
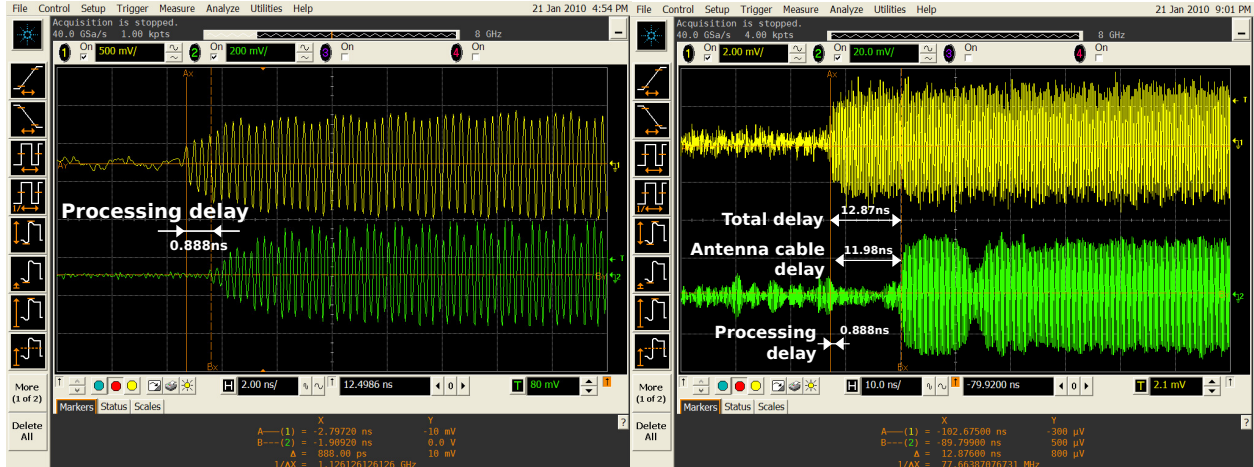
(a) Cable	(b) Wireless

Figure 7: The delay of the prover's distance bounding radio extension. The top signal is measured at the reception antenna of the provers radio and is transmitted on channel C_0 at 3.5GHz. The bottom signal is measured at the transmission antenna and is being transmitted at the C_2 channel at 4.0GHz. The delay between them, and thus the prover's processing time is $0.888ns$.

confidence interval. We see from the figure that the processing time of the prover is stable between 0.8ns and 1ns.

Note that if the same setup would have been implemented in an integrated circuit, the signal path would be a lot shorter and consequently the processing time would have been smaller. We therefore do not claim that our prototype is the best that can be achieved, rather it shows the processing time that can be achieved using standard SMA components.

6.2 Wireless Implementation

Since distance bounding protocols are primarily useful in wireless environments, in this section we show that our prototype equally enables distance bounding using wireless communication (instead of wires). The basic construction of the prover is the same as in the wired setup, except that the prototype input and output are connected to antennas. The function generator that generates the verifiers signal and the oscilloscope used to measure the round trip time are likewise connected to antennas.

The result of the wireless implementation can be seen in Figure 7(b). Unfortunately we had to use SMA cables of about $1m$ to connect the antennas because of the way the antennas are mounted. In addition there was about $.1m$ between the transmission antenna and the receiving antenna. This results in a delay introduced by the cables and the space between the antennas referred to on Figure 7(b) as

"antenna cable delay". The output of the prototype was passed through a high-pass filter and the input passed through a low-pass filter to prevent the transmitting antenna from feeding back into the receiving antenna. The oscilloscope used to measure the difference in arrival time also had filters to separate the ground truth signal, i.e., the signal coming directly from the function generator from the one being transmitted by the prototype. The filters allowed for a full duplex wireless channel to be created between our wireless prototype and the function generator and oscilloscope.

It should be noted that the channel switching mechanism of our prototype is ideal for a wireless implementation. Any wireless distance bounding protocol needs more than one channel (i.e., full duplex) in order to reply as fast as possible. Encoding the prover's reply in the choice of channel means that the solution is strait forward to apply without causing interference between the prover and verifier.

7 Related Work

Distance bounding, as a concept, was first proposed by Brands and Chaum in [3] who introduced techniques enabling a verifier to determine an upperbound on the physical distance to a prover (as summarized in Section 2). In addition, they considered the case where the verifier also authenticates the prover in addition to establishing the distance bound.

Several optimizations and studies of distance bounding were subsequently proposed for wireless networks, including [28, 30, 5] and for sensor networks [18, 5, 27]. Distance bounding protocols have also been proposed in other contexts, e.g., for RFIDs [13, 10, 19] and ultra wide band (UWB) devices [17, 12].

In [23] the authors studied information leakage in distance bounding protocols. A mutual distance bounding protocol using interleaved challenges and responses was proposed in [31] and in [28] and [5] the authors investigated the use of distance bounding protocols for location verification and secure localization. Sastry, Shankar and Wagner [25] proposed the so-called "in-region verification" appropriate for certain applications, such as location-based access control. Collusion attacks on distance bounding location verification protocols where considered in [7, 6]. Ultrasonic distance bounding was used for access control [25] and for key establishment [32]. In [22] ultrasonic distance bounding was further used for proximity based access control to implementable medical devices. Other attacks have been proposed against distance bounding protocols in general. The so-called "late-commit" attacks where proposed in [14], where the attacker exploits the modulation scheme in order to manipulate the distance. Bit guessing attacks [8] that accomplish the same thing where also proposed. These attacks were further studied in practical implementations in [11].

Until now, most of the work done in this field has been theoretical. To our knowledge our work is the first to propose a realizable distance bounding protocol using radio communication, with a processing time at the prover that is low enough to provide a useful distance granularity.

8 Conclusion

We demonstrated that radio distance bounding protocols can be implemented to match the strict processing that these protocols require (i.e., that the prover receives, processes and transmits signals in $\leq 1ns$). This can be achieved using a specially implemented concatenation as the prover's processing function. Through this we showed that the use of processing functions which require that the prover demodulates (interprets) the verifier's challenge before responding to it, is not desirable or necessary for distance bounding. Finally, we showed that other processing functions such as XOR and the comparison function, that were used in a number of proposed distance bounding protocols, are not best suited for the implementation of radio distance bounding.

References

[1] Thomas Beth and Yvo Desmedt. Identification tokens - or: Solving the chess grandmaster problem. In *CRYPTO '90: Proceedings of the 10th Annual International Cryptology Conference on Advances in Cryptology*, pages 169–177, London, UK, 1991. Springer-Verlag.

[2] Colin Boyd and Anish Mathuria. *Protocols for authentication and key establishment.* Springer, 1998.

[3] Stefan Brands and David Chaum. Distance-bounding protocols. In *EUROCRYPT '93*, pages 344–359, Secaucus, NJ, USA, 1994. Springer-Verlag New York, Inc.

[4] Laurent Bussard and Walid Bagga. Distance-bounding proof of knowledge protocols to avoid terrorist fraud attacks. Technical report, Institut Eurecom, France, 05 2004.

[5] Srdjan Capkun and Jean-Pierre Hubaux. Secure positioning of wireless devices with application to sensor networks. In *IEEE INFOCOM*, 2005.

[6] Nishanth Chandran, Vipul Goyal, Ryan Moriarty, and Rafail Ostrovsky. Position based cryptography. In *CRYPTO '09: Proceedings of the 29th Annual International Cryptology Conference on Advances in Cryptology*, pages 391–407, Berlin, Heidelberg, 2009. Springer-Verlag.

[7] Jerry T. Chiang, Jason J. Haas, and Yih-Chun Hu. Secure and precise location verification using distance bounding and simultaneous multilateration. In *ACM WiSec '09*, pages 181–192, New York, NY, USA, 2009. ACM.

[8] Jolyon Clulow, Gerhard P. Hancke, Markus G. Kuhn, and Tyler Moore. So near and yet so far: Distance-bounding attacks in wireless networks. In *Proceedings of the European Workshop on Security and Privacy in Ad-hoc and Sensor Networks (ESAS)*, 2006.

[9] Yvo Desmedt. Position statement in rfid s&p panel: From relative security to perceived secure. In *Financial Cryptography*, pages 53–56, 2007.

[10] Saar Drimer and Steven J. Murdoch. Keep your enemies close: Distance bounding against smartcard relay attacks. In *Proceedings of the USENIX Security Symposium 2007*, 2007.

[11] Manuel Flury, Marcin Poturalski, Panos Papadimitratos, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. Effectiveness of Distance-Decreasing Attacks Against Impulse Radio Ranging. In *3rd ACM Conference on Wireless Network Security (WiSec)*, 2010.

[12] S. Gezici, Zhi Tian, G.B. Giannakis, H. Kobayashi, A.F. Molisch, H.V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *Signal Processing Magazine, IEEE*, 22(4):70–84, July 2005.

[13] Gerhard P. Hancke and Markus G. Kuhn. An rfid distance bounding protocol. In *SecureComm '05: Proceedings of the First International Conference on Security and Privacy for Emerging Areas in Communications Networks*, pages 67–73, Washington, DC, USA, 2005. IEEE Computer Society.

[14] Gerhard P. Hancke and Markus G. Kuhn. Attacks on time-of-flight distance bounding channels. In *WiSec '08: Proceedings of the first ACM conference on Wireless net work security*, pages 194–202, New York, NY, USA, 2008. ACM.

[15] Y.-C. Hu, A. Perrig, and D. B. Johnson. Packet Leashes: A Defense against Wormhole Attacks in Wireless Networks. In *Proceedings of the IEEE Conference on Computer Communications (InfoCom)*, San Francisco, USA, April 2003.

[16] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Ariadne: a secure on-demand routing protocol for ad hoc networks. *Wirel. Netw.*, 11(1-2):21–38, 2005.

[17] J.-Y. Lee and R.A. Scholtz. Ranging in a Dense Multipath Environment Using an UWB Radio Link. *IEEE Journal on Selected Areas in Communications*, 20(9), December 2002.

[18] Catherine Meadows, Paul Syverson, and LiWu Chang. Towards more efficient distance bounding protocols for use in sensor networks. *Securecomm*, pages 1–5, Aug. 28 2006-Sept. 1 2006.

[19] Jorge Munilla, Andres Ortiz, and Alberto Peinado. Distance bounding protocols with void-challenges for RFID. Printed handout at the Workshop on RFID Security – RFIDSec 06, July 2006.

[20] National Space-Based Positioning, Navigation, and Timing Coordination Office. Global positioning system. http://www.gps.gov/.

[21] Maxim Integrated Products. An introduction to direct sequence spread spectrum communications. http://www.maxim-ic.com/, 2003.

[22] Kasper Bonne Rasmussen, Claude Castelluccia, Thomas S. Heydt-Benjamin, and Srdjan Čapkun. Proximity-based access control for implantable medical devices. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*. ACM, 2009.

[23] Kasper Bonne Rasmussen and Srdjan Čapkun. Location privacy of distance bounding protocols. In *CCS '08: Proceedings of the 15th ACM conference on Computer and communications security*, pages 149–160, New York, NY, USA, 2008. ACM.

[24] Qingchun Ren and Qilian Liang. Throughput and energy-efficiency-aware protocol for ultrawideband communication in wireless sensor networks: A cross-layer approach. *IEEE Transactions on Mobile Computing*, 7:805–816, 2007.

[25] Naveen Sastry, Umesh Shankar, and David Wagner. Secure verification of location claims. In *WiSe '03: Proceedings of the 2nd ACM workshop on Wireless security*, New York, NY, USA, 2003. ACM.

[26] Patrick Schaller, Benedikt Schmidt, David Basin, and Srdjan Capkun. Modeling and verifying physical properties of security protocols for wireless networks. In *CSF '09: Proceedings of the 2009 22nd IEEE Computer Security Foundations Symposium*, pages 109–123, Washington, DC, USA, 2009. IEEE Computer Society.

[27] S. Sedighpour, S. Capkun, S. Ganeriwal, and M. Srivastava. Implementation of attacks on ultrasonic ranging systems, nov 2005.

[28] D. Singelee and B. Preneel. Location verification using secure distance bounding protocols. In *Mobile Adhoc and Sensor Systems Conference, 2005. IEEE International Conference on*, Nov. 2005.

[29] Nils Ole Tippenhauer and Srdjan Čapkun. Id-based secure distance bounding and localization. In *In Proceedings of ESORICS (European Symposium on Research in Computer Security)*, 2009.

[30] S. Čapkun, L. Buttyán, and J.-P. Hubaux. SECTOR: Secure Tracking of Node Encounters in Multi-hop Wireless Networks. In *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, Washington, USA, October 2003.

[31] Srdjan Čapkun, Levente Buttyán, and Jean-Pierre Hubaux. Sector: secure tracking of node encounters in multi-hop wireless networks. In *ACM SASN '03*, pages 21–32, New York, NY, USA, 2003. ACM.

[32] Srdjan Čapkun and Mario Čagalj. Integrity regions: authentication through presence in wireless networks. In *WiSe '06: Proceedings of the 5th ACM workshop on Wireless security*, pages 1–10. ACM, 2006.