

A Formal Approach to Distance-Bounding RFID Protocols

Ulrich Dürholz² Marc Fischlin¹ Michael Kasper² Cristina Onete¹

¹Darmstadt University of Technology & CASED, Germany
www.cryptoplexity.de

²Fraunhofer Institute for Secure Information Technology (SIT) and CASED, Germany

Abstract. Distance-Bounding identification protocols aim at impeding man-in-the-middle attacks by measuring response times. There are three kinds of attacks such protocols could address: (1) Mafia attacks where the adversary relays communication between honest prover and honest verifier in different sessions; (2) Terrorist attacks where the adversary gets limited active support from the prover to impersonate. (3) Distance attacks where a malicious prover claims to be closer to the verifier than it actually is. Many protocols in the literature address one or two such threats, but no rigorous cryptographic security models —nor clean security proofs— exist so far. For resource-constrained RFID tags, distance-bounding is more difficult to achieve. Our contribution here is to formally define security against the above-mentioned attacks and to relate the properties. We thus refute previous beliefs about relations between the notions, showing instead that they are independent. Finally we use our new framework to assess the security of the RFID distance-bounding scheme due to Kim and Avoine, and enhance it to include impersonation security and allow for errors due to noisy channel transmissions.

Keywords. RFID distance-bounding protocols, formal models, provable security

1 Introduction

Man-in-the-middle attacks are a powerful strategy for an adversary to fool identification schemes: by relaying communication between provers and verifiers, the adversary makes verifiers accept. Following [17] quintessential relaying is called Mafia fraud. Environments with no central authority and certificates, like RFID identification, are particularly subject to such attacks. Practical set-ups [25, 16, 18, 24, 21] indicate their feasibility, and several works investigate attacks on the HB protocol [29, 22, 19, 10, 33, 32], which is designed for low-power devices e.g. RFIDs. For a more general overview of RFID security issues see [30].

1.1 Distance-Bounding Protocols

Distance-bounding protocols, proposed initially by Brands and Chaum [9], suggest a countermeasure against man-in-the-middle attacks. The basic idea is that relaying communication takes longer than genuine responses. Thus, if verifiers measure the time elapsed between sending a value and receiving the reply, man-in-the-middle attacks should become infeasible. In practice, verifiers check round-times for many so-called *fast* or *time-critical* communication phases, (as opposed to *slow* or *lazy* phases, where round times do not matter).

We mainly address RFID authentication, but our new framework applies to general distance-bounding protocols where provers and verifiers may interact and at the end the verifier outputs a bit indicating whether the prover has been authenticated or not. For RFID authentication, the verifiers are readers and the provers are RFID tags; we use these terms interchangeably for provers and verifiers. RFID distance-bounding has been investigated quite extensively [1, 2, 3, 4, 9, 11, 13, 14, 16, 17, 18, 24, 25, 28, 31, 35, 36, 37]. See also [26] for a comprehensive overview. The three main threats that need to be avoided are: (1) Mafia fraud, where the adversary tries to impersonate to the reader while communicating with the genuine tag (the timing prevents it from using pure relay though); (2) Terrorist fraud where tags may leak useful information to the adversary in offline phases to help it authenticate (the restriction being that tags should not reveal trivial information like the secret key); (3) Distance Fraud, where the tag claims to be closer than it actually is. We also consider the basic (often neglected) requirement for identification, i.e. slow-round impersonation resistance, independent of the limited number of fast phases.

We exemplify the three attacks as follows: consider a gym locker with an inbuilt RFID reader, for which Alice holds the unique pass key (an RFID tag). One evening, Alice is not at the gym, but at a party. In the *Mafia fraud* scenario, Bob *is* at the gym; his accomplice, Bobette, is at the party with Alice. Bob wants to open the locker (without Alice’s consent for Mafia fraud). In this attack, Bob and Bobette relay messages between the locker and Alice’s tag. If, on the contrary, Alice *wants* Bob to use her locker (for this night only) we have *Terrorist fraud*. Alice may now give Bob information to help him use her locker, but she doesn’t want Bob to abuse her kindness and open the locker on his own, this or any other time. For Terrorist attacks thus, Alice helps Bob herself: Bobette is not needed. Finally, if Alice parked her car in a bad spot, she might want to “prove” that she was at the gym instead (this is *distance fraud*) by opening the locker, which can be opened only if the unique key is in direct proximity.

Several existing protocols implement resistance against one (or more) of the above threats. However, the two protocols due to Brands and Chaum [9] and to Bussard et al. [11] rely on signatures and zero-knowledge proofs, and are too computationally expensive for devices like (low-cost) RFID tags. The protocol due to Hancke and Kuhn [28] is very efficient and Mafia fraud resistant (though allowing impersonation with probability about $\frac{3}{4}$ per fast round, unlike $\frac{1}{2}$ as in [9]), but not terrorist fraud resistant. Also, authentication in [28] is achieved during fast phases only, thus authentication is reduced to

the limited number of rounds that tags can support.

Avoine et al. ([31, 4]) also show several efficient distance bounding protocols. In particular, Kim and Avoine [31] offer partial reader authentication with mixed fast-round challenges and thus provide distance fraud resistance and nearly-optimal Mafia fraud resistance. Terrorist fraud is not addressed here. By contrast, [4] requires more storage capacity for tags, and the inter-dependency of fast-phase challenges makes the false acceptance rate hard to compute.

The protocols due to Bussard et al. [11] and Reid et al. [35] aim to resist terrorist fraud. The schema in [35] builds on the Mafia resistant protocol of [28], inheriting its lack of impersonation security. Bussard et al. use public key cryptography, which is too expensive for RFID systems. The same holds for the mutual authentication work of Singelée and Preneel [36] and Capkun et al. [13], which use public-key and elliptic-curve cryptography. Also, the security of these constructions have never been formally assessed.

Also, most existing work permits adversaries to impersonate the reader to the tag, thus leaking information about fast-phase response times. If only bits are transmitted in fast phases, the ideal impersonation bound would be 2^{N_c} for N_c critical rounds; however, most protocols allow impersonation and thus reach a lower than ideal bound. To account for this Mafia fraud attack, under “Rounds”, we give the number N_c of time-critical rounds required for a Mafia resistance of about 2^{-k} . We round down the number of rounds in [28, 35] to $2k$. Note that [4] shows a construction with reduced complexity, at the expense of security.

	[9]	[28]	[4]	[35]	[31]
Mafia	✓	✓	✓	✓	✓
Terror	×	×	×	(✓) ¹	×
Distance	✓	✓	✓	✓	✓
Impersonation	×	×	✓	×	×
Rounds N_c	k	$> 2k$	k	$> 2k$	k
Storage	N_c	$2N_c$	$O(2^{N_c})$	$2 N_c$	$4N_c$
Private-key	×	✓	✓	✓	✓

Figure 1: Claimed Security and Actual Efficiency of Distance Bounding Protocols at a glance (¹only special terrorists, no formal proof.)

We lastly outline some related cryptographic concepts from the literature. Most prominently, the recent position-based cryptography work [15] aims to determine if a prover is (exactly) at a claimed position — but in a single protocol run, with many verifiers. As we discuss in more detail in Appendix A, this scenario is closest to our distance fraud model, where tags must prove they are closer to the reader than they really are. We note also that exact positioning is impossible in practice for RFID, requiring too many readers to deal with the high variance in response time.¹

Also related to our topic is the approach describing self-delegation, such as the work of [23] and [12]. We discuss in Appendix A that self-delegation resembles terrorist fraud, but

¹Recent work due to Hancke [27] in fact suggests designing a distance bounding channel limiting channel-specific variations of response times.

with certain limitations, such as restricting the scenario only to the public-key setting and disallowing online help of any kind. Also, [23] relies on public-key cryptography and non-interactive zero-knowledge proofs, primitives that are unsuitable for RFID. We discuss this topic further in the Appendix, and also address the relationship between this work and transferability of anonymous credentials.

1.2 Our Contributions

Our contributions are threefold: (1) We give rigorous cryptographic security models for Mafia, Terrorist, and Distance fraud, thus (2) relating the security properties formally. We also refute the claim in [35] that terrorist attacks resistance implies distance-fraud resistance. Finally, we (3) use our framework to formally assess the security of the prominent scheme in [31], and enhance it to allow for noisy channels and implement impersonation resistance.

The Practice behind the Theory. Practical investigations [14, 16, 35, 28, 34, 27] indicate some design issues for RFID distance-bounding protocols. As such considerations apply for all low-power devices, we provide for them in our framework.

BANDWIDTH. Time measurements are very fragile when the parties send other messages than bits [16, 27, 35]. The reason is that on the one hand, fresh noise is introduced in the communication, and on the other hand, the unreliability of the transmission increases with the size of the transmitted message. Distance-bounding protocols should thus only exchange bits for time-critical steps (we still formulate models for arbitrary transmissions).

COMPUTATIONS. Fast step computations should be very simple, otherwise they dominate the round-time. It is also important that the computation-time for each round is constant.

STORAGE. To be suitable for low-power devices, RFID distance bounding protocols must require only little storage.

NOISE. Both transmissions and time measurements are subject to noise [16, 27, 35]. Considering, as above, that only bits are exchanged in measured phases, immediate error correction for example is impossible. Hence, our model and protocols take into account threshold levels for failures during timed steps. The models are thus more profound, as the adversary can now run the man-in-the-middle strategy for some phases.

EARLY BIT DETECTION. Depending on the physical implementations the adversary may be able to predict a transmitted bit “halfway through the signal” [18]. Similarly, the computation time of the parties in a measured phase may depend on the actual value received. In other words, the adversary may occasionally be able to derive

information from the reader or tag faster than expected. Hence, our model allows the adversary to relay information as long as it is not purely duplicated.

OFFLINE AUTHENTICATION. Distance bounding is often achieved by timing the communication between the reader and the tag; this is implemented in a few fast communication rounds [4]. Due to hardware constraints, one cannot use many such rounds, and therefore further authentication should be used. As pointed out in [4] it is preferable that the basic authentication step be done before the fast phase begins. Some protocols do not have this property [9, 28], whereas we give a strong definition of it and suggest it as an enhancement of solutions such as [31].

The Models. A sound modeling of the above attacks is crucial to assessing protocol security. Confusions appear especially with attack modes and successful man-in-the-middle attacks, e.g. for the HB protocol [29, 22, 19, 10, 33, 32]. As another example, the allegedly secure Hitomi and NUS protocols were recently proved insecure [1]. We formalize game-based models while also considering practical conditions. This enables us to formally prove that, contrary to the remark of [35], terrorist fraud resistance does not imply distance fraud resistance. In fact, we show that Mafia resistance, terrorist resistance, and distance-fraud resistance are all independent. More precisely, we present protocols that are vulnerable to one attack, but resistant to all others (including the basic authentication-protocol requirement of impersonation resistance). In particular, terrorist fraud resistance also does not imply Mafia fraud resistance, nor vice versa.

Some groundwork has already been laid in this field by Avoine et al. [2], who model Mafia, distance, and terrorist fraud in both a black-box and a white-box sense i.e. giving adversaries access to the implementation of the primitive or not. Distance-bounding protocols here have two main goals: authentication and distance checking; each type of fraud is also more formally defined. Adversaries may choose from three main strategies: pre-ask (query prover before being queried by verifier), post-ask (query prover after being queried by verifier), and early-reply (respond before verifier sends query, without querying the prover). In the black box model, Mafia and terrorist fraud are proved equivalent, whereas terrorist fraud resistance is said to imply distance fraud resistance. In the white-box model, terrorist fraud resistance implies both Mafia and distance fraud. Mafia fraud resistance is equivalent in the black and white box models, and white-box terrorist and distance fraud are strictly stronger than the black-box notions.

By contrast, our definitions are much more concrete and formal. Protocols have many rounds (lazy or time-critical), and adversaries choose (possibly different) strategies at each round, unlike [2]. Our Mafia adversaries may relay *parts of the communication*, e.g. flip bits, or purely relay (taint) some rounds. Our Mafia and terrorist provers may be anywhere, unlike [2], where provers are outside the target distance from the verifier. By using a simulator, we concretely define “advantage for future attacks” [2] for terrorist fraud. Hence, we prove that all security notions are independent. We also extend impersonation resistance to lazy-phase authentication, thus preventing information leaks to fake provers.

Using our Framework. We use our framework to assess the security of the protocol due to Kim and Avoine [31], which relies on mutual tag-to-reader and reader-to-tag fast-phase authentication to achieve good Mafia and Distance Fraud (but not impersonation-) resistance. If reader authentication fails, the tag generates random responses every round. In its original form, this scheme offers no impersonation resistance, but is both Mafia and Distance fraud resistant. Since it relies on mutual tag-to-reader and reader-to-tag fast-phase authentication, the protocol [31] limits the adversary’s ability to query the tag in advance and thus learn the correct responses in fast phases. If an impersonation attempt is detected, the tag generates random responses every round. We first make the construction in [31] impersonation resistant, then formally assess its security in our framework. We also prove that it is not terrorist-fraud resistant in our model. Additionally, we illustrate the notions of Mafia, Terrorist, and Distance Fraud and Impersonation security by showing how to achieve them in practice.

2 Preliminaries

We consider a single reader \mathcal{R} and a single tag \mathcal{T} , sharing a secret key generated through Kg . To the reader we associate a clock and a database entry storing the tag’s secret key. We assume that the identification scheme $\mathcal{ID} = (\text{Kg}, \mathcal{T}, \mathcal{R})$ marks (consecutive) steps of the identification protocol as *lazy* or *time-critical*: in time-critical steps, one party—usually the reader—measures the round-time Δt and compares it to a predetermined threshold t_{\max} ; else the phase is called lazy. A protocol run can consist of arbitrary *non-overlapping* sequences of lazy and time-critical phases, with time-critical phases possibly following one another. Denote by N_c the number of time-critical phases. Errors due to time-measurement noise are modelled by allowing T_{\max} -many round-times to exceed t_{\max} . Similarly, E_{\max} is the maximum number of time-critical phases with erroneous transmissions.

Definition 2.1 *An identification scheme for timing parameters $(t_{\max}, T_{\max}, E_{\max}, N_c)$ is a triplet of efficient algorithms $\mathcal{ID} = (\text{Kg}, \mathcal{T}, \mathcal{R})$ with:*

KEY GENERATION. *For parameter $n \in \mathbb{N}$, Kg generates a secret key sk .*

IDENTIFICATION. *The joint execution of algorithms $\mathcal{T}(sk)$ and $\mathcal{R}(sk)$ generates, depending on $t_{\max}, T_{\max}, E_{\max}, N_c$, a verifier output $b \in \{0, 1\}$.*

We assume that the scheme is complete: for any $n \in \mathbb{N}$ and any key $sk \leftarrow \text{Kg}(1^n)$, the decision bit b produced by honest party $\mathcal{R}(sk)$ interacting with honest party $\mathcal{T}(sk)$ under the requirements following from the timing parameters, is 1 with probability (negligibly close to) 1.

Note that although most definitions of distance-bounding protocols omit t_{\max} , this parameter is a crucial difference between distance-bounding and authentication protocols, where t_{\max} is by default infinitely large. The parameters E_{\max} and T_{\max} are intrinsic

to communication over noisy channels (e.g. RF channels between readers and passive and semi-passive RFID tags²). In distance bounding, it is unreasonable to separate the *reliability* of the communication from its *security*; these properties are connected by the importance of round-time measurements towards acceptance or rejection. Bit errors are unavoidable in RF communication, as stated in point 4 of Clulow et al.’s principles for secure time-of-flight distance-bounding [16]. As described in section 1, RF communication noise implies that transmissions between readers and tags are not always reliable, possibly reaching the reader outside the time bound. We can, however, set $T_{\max} = E_{\max} = 0$ for extremely reliable scenarios. It is also easy to extend the above definition to capture RFID systems relying on some shared public information, where readers use public-key schemes. We can add an algorithm `SysPar` generating a key pair (K_S, K_P) such that K_S is given to the reader and the key generation algorithm `Kg` for the tags takes K_P . Completeness then requires that the honest verifier accepts honest provers for all pairs (K_S, K_P) .

3 Security Model

3.1 Communication Model

The adversary can access: a reader instance to which it impersonates the tag (a *reader-adversary session*), a tag instance to which it impersonates the reader (*adversary-tag session*), and an interface observing a genuine reader-tag protocol for which the adversary cannot change transmissions (*reader-tag session*). The adversary can access all interfaces concurrently and in many sessions (sessions share a secret key, but have different random tapes). Each session has an identifier `sid` (given to the adversary, but not to protocol participants). We assume that the adversary knows if an authentication attempt succeeded or not.³

In our concurrent single-reader-single-tag scenario (as opposed to a single reader and multiple tags), many instances of the single tag may exist in parallel, sharing the secret key, but not the random tape. The key is *static*, i.e., not updated after executions. For many independent keys (multiple tags), adversaries can always pick a tag to attack in our model. Three factors are crucial to multiple-tag scenarios: the interdependency of the keys; the noise in the communication due to tag-to-reader collisions (a factor modeled by E_{\max}); and key management. A formal approach for key update is, however, beyond the scope of this paper.

We assume message-driven attacks, i.e., honest parties reply as soon as they receive a (protocol) message. The adversary schedules message delivery to honest parties. We assume a global clock, assigning an integer `clock(sid, k)` to the k -th protocol mes-

²Passive RFID tags have no power source of their own and are very sensitive to their environment, in particular metals and liquids. Semi-passive tags use their own power source for computation, but rely on readers for communication, and are also vulnerable to interference by metals and liquids.

³This is not a strong requirement. In practice the success of an authentication attempt is marked by a physical event: a beep, the opening of a door, a green light etc.

sage, delivered in session sid to an honest party. The honest party's reply is assigned $\text{clock}(\text{sid}, k + 1) = \text{clock}(\text{sid}, k) + 1$.⁴ Furthermore, $\text{clock}(\text{sid}, k) < \text{clock}(\text{sid}^*, k)$ if the adversary delivers the k -th message in session sid^* after the k -th message in session sid . Denote by $\Pi_{\text{sid}}[i \dots j]$ messages i to j exchanged in session sid and by $\Pi_{\text{sid}}[1 \dots]$ all messages exchanged in sid . Let $\text{view}_{\mathcal{A}}$ denote the adversary's view in an attack, containing its internal randomness and all the transcripts (of communication with and among other parties).

Let t denote the adversary's running time, including steps of honest parties. Denote by $q_{\mathcal{R}}$ (resp. $q_{\mathcal{T}}$ and q_{OBS}) the maximal number of reader-adversary (resp. adversary-tag and reader-tag) sessions. Below we refine the attacks and define winning conditions for the adversary (who must non-trivially impersonate the tag in a reader-adversary session). For an attack att we write $\text{Adv}_{\mathcal{ID}}^{\text{att}}(\mathcal{A})$ for the probability that the $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -adversary \mathcal{A} wins.

3.2 Mafia Fraud Detection Model

Mafia fraud adversaries can communicate arbitrarily with tag and reader, *except for purely relaying time-critical transmissions*. We exclude only attacks where the adversary relays *exact* transmissions, calling such time-critical phases *tainted*:

Definition 3.1 (Tainted Time-Critical Phase (Mafia)) *A time-critical phase $\Pi_{\text{sid}}[k \dots k + 2\ell - 1] = (m_k, \dots, m_{k+2\ell-1})$ for $k, \ell \geq 1$ of a reader-adversary session sid , with the k -th message being received by the adversary, is tainted by the phase $\Pi_{\text{sid}^*}[k \dots k + 2\ell - 1] = (m_k^*, \dots, m_{k+2\ell-1}^*)$ of an adversary-tag session sid^* if for all $i = 0, 1, \dots, \ell - 1$ we have:*

$$\begin{aligned} (m_k, \dots, m_{k+2\ell-1}) &= (m_k^*, \dots, m_{k+2\ell-1}^*), \\ \text{clock}(\text{sid}, k + 2i) &< \text{clock}(\text{sid}^*, k + 2i), \\ \text{and } \text{clock}(\text{sid}, k + 2i + 1) &> \text{clock}(\text{sid}^*, k + 2i + 1). \end{aligned}$$

Our notion, shown in Figure 2, is conservative with respect to the following:

- We do not exclude phases where the adversary changes the content before relaying. since the content may determine the response time if say, the computation is less involved for a 0-bit than for a 1-bit. This may allow the adversary to receive the tag's answer for a different value in time, before it is required to answer the reader.

As a consequence, if a protocol sends redundancy like an extra 0-bit, then an adversary can easily flip this bit and not taint the phase, albeit simply relaying the crucial information. We nevertheless grant this freedom to the adversary, as it coincides with the similar idea of matching sessions in key exchange protocols [7, 6]: protocols with obvious redundancy can be easily modified; also, it is common cryptographic practice to err on the safe side.

⁴We could also allow adversaries to delay message delivery *from* honest parties. Our model and results are robust with respect to this idea, but this contradicts the implementation of reliable time measurements and enable denial-of-service attacks.

- According to our definition, a time-critical phase becomes tainted if there is another session in which the adversary relays *all* transmitted messages in the two sessions. If the adversary changes the content of a single transmission in such a phase or the order of a single step only, then the phase is not tainted by the other session anymore. This provides again more freedom to the adversary and strengthens the security notion.

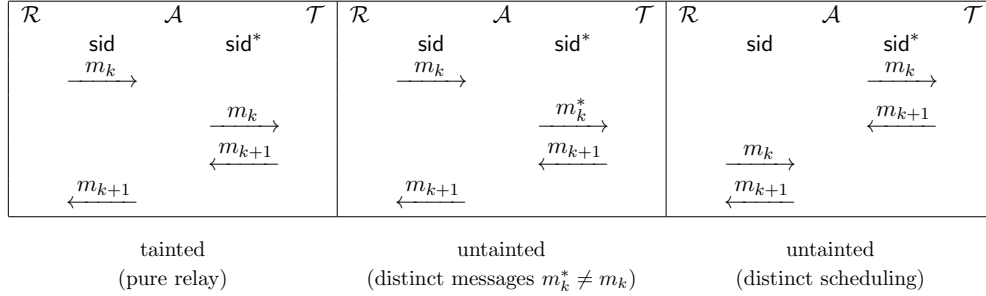


Figure 2: Examples of Tainted and Untainted Time-Critical Phases.

The adversary must now make the reader accept in session sid such that for each adversary-tag session sid^* at most T_{\max} phases of sid are tainted by sid^* :

Definition 3.2 (Mafia Fraud Resistance) *For a distance-bounding identification scheme \mathcal{ID} with parameters $(t_{\max}, T_{\max}, E_{\max}, N_c)$, a $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -Mafia-fraud adversary \mathcal{A} wins against \mathcal{ID} if the reader accepts in a reader-adversary session sid such that any adversary-tag session sid^* taints at most T_{\max} time-critical phases of sid . Let $\text{Adv}_{\mathcal{ID}}^{\text{mafia}}(\mathcal{A})$ denote the probability that \mathcal{A} wins.*

Different adversary-tag sessions may taint different rounds of reader-adversary session sid . As we count T_{\max} over all adversary-tag sessions the adversary wins if it taints at most T_{\max} *distinct* phases. Protocols must prevent such attacks to be Mafia fraud secure in concurrent settings. Further session interdependencies should also be avoided so that messages from another session do not taint sid .

3.3 Terrorist Attack Model

In a terrorist attack the tag aids the adversary in all short of revealing its secret key, in fact wanting to ensure that the adversary only wins with the tag’s aid (the dishonest prover controls the adversary’s access). Desmedt [17] concretely describes the tag’s involvement as offline help in a single impersonation attempt. The adversary now wins if the reader accepts, but the adversary cannot use the help given by tag \mathcal{T}' to impersonate further.

We formalize the idea by using ideas from proofs of computational ability [38, 5], which exactly capture the intuition of terrorist attacks: given support from a prover e.g. \mathcal{T}' , one can solve a hard problem e.g. identifying to the reader. This is independent of

how the prover gives support. We are not, however, interested in the cases where \mathcal{T}' yields the entire key (or large parts of it) and mark certain auxiliary data given by \mathcal{T}' as trivial, i.e. the data is trivial if it allows one to successfully complete a “fresh” identification attempt *without help from \mathcal{T}'* . This includes the case when \mathcal{T}' gives the secret key, but circumvents the problem of determining which parts of the key are helpful. Data is trivial if it aids identification beyond the dedicated help in the session where \mathcal{T}' helps.

We formalize the latter by demanding that no algorithm \mathcal{S} , called simulator, can use the data passed by \mathcal{T}' to \mathcal{A} to authenticate without the help of \mathcal{T}' (to be fair, we allow \mathcal{S} the same number $q_{\mathcal{R}}$ of attempts as \mathcal{A}). This is in line with well-known simulation paradigms, and allows to compare the respective success probabilities of the adversary \mathcal{A} aided by \mathcal{T}' , and the simulator \mathcal{S} using \mathcal{A} 's information to authenticate. If \mathcal{A} is significantly more successful than \mathcal{S} , the attack is non-trivial and the protocol is insecure against terrorist attacks. Note that “unsophisticated” adversaries may do worse than simulators for secure schemes, thus yielding negative advantages.

For terrorist fraud, \mathcal{A} acts as for Mafia fraud, but may query the “malicious” interface \mathcal{T}' in lazy phases. Sessions sid' with \mathcal{T}' are arbitrary, not following protocol. In fact we may consider only one session sid' when \mathcal{T}' helps \mathcal{A} . The tag may *not* aid \mathcal{A} in time-critical phases, a fact which we model by defining tainted time-critical phases as pure-relay phases or rounds where \mathcal{A} queries \mathcal{T}' .

Definition 3.3 (Tainted Time-Critical Phase (Terror)) *A time-critical phase $\Pi_{\text{sid}}[k \dots k+2\ell-1] = (m_k, \dots, m_{k+2\ell-1})$ for $k, \ell \geq 1$ of a reader-adversary session sid , with the k -th message being received by the adversary, is tainted if there is a session sid' between the adversary and \mathcal{T}' such that, for some i ,*

$$\text{clock}(\text{sid}, k) < \text{clock}(\text{sid}', i) < \text{clock}(\text{sid}, k + 2\ell - 1).$$

For the new definition of tainted phases, terrorist fraud resistance demands that for any terrorist fraud attacker \mathcal{A} there exists a simulator \mathcal{S} such that for any supporting \mathcal{T}' , \mathcal{S} is essentially as successful as \mathcal{A} . We use concrete security statements and omit quantification over \mathcal{A} , \mathcal{S} , and \mathcal{T}' algorithms; this quantification is included in subsequent security claims in the usual form (i.e., for any adversary there exists a simulator such that for all tags the advantage is small).

Definition 3.4 (Terrorist Fraud Resistance) *Let \mathcal{ID} be a distance-bounding identification scheme with parameters $(t_{\max}, T_{\max}, E_{\max}, N_c)$. Let \mathcal{A} be a $(t, q_{\mathcal{R}}, q'_{\mathcal{T}})$ -terrorist-fraud adversary, \mathcal{S} be an algorithm running in time $t_{\mathcal{S}}$, and \mathcal{T}' be an algorithm running in time t' . Denote*

$$\text{Adv}_{\mathcal{ID}}^{\text{terror}}(\mathcal{A}, \mathcal{S}, \mathcal{T}') = p_{\mathcal{A}} - p_{\mathcal{S}}$$

where $p_{\mathcal{A}}$ is the probability that the reader accepts in one of the $q_{\mathcal{R}}$ reader-adversary sessions sid such that at most T_{\max} time-critical phases of sid are tainted, and $p_{\mathcal{S}}$ is the probability that, given $\text{view}_{\mathcal{A}}$ in an attack of \mathcal{A} , \mathcal{S} makes the reader accept in one of $q_{\mathcal{R}}$ subsequent executions.

Again, if the advantage is negative, \mathcal{A} performs worse than \mathcal{S} . Our notion is quite strong: the simulator only gets to see \mathcal{A} 's transcript in an offline phase, instead of communicating with \mathcal{T}' online. This guarantees stronger security and saves us from dealing with issues related to the number of queries and successful attacks (adversary vs. simulator).

How does our definition fit into previous efforts? Previous protocols [35, 4] claim to achieve a security of $(1/2)^{-N_c}$. This, however, corresponds to a tailor-made strategy of \mathcal{T}' ; other strategies may still exist. Proving that the advantage in Definition 3.4 is negligible, then we *prove* that \mathcal{T}' can only help trivially.

3.4 Distance-Fraud Model

For distance fraud an adversary must reply ahead of a time-critical phase or it cannot respond in time. In practice this is enforced by a tight value of t_{\max} . For any time-critical phase, with possibly many communication rounds, the adversary must commit to the *first* message to be sent. For any later rounds in the phase, the adversary has time to reply even from farther away.

The order of committed and sent values is determined by on oracle `CommitTo` with a single session $\text{sid}_{\text{CommitTo}}$, taking tuples (sid, i, m_i) from the adversary and giving empty responses. The adversary commits to the first message of time-critical phase i of session sid (message j in sid) at time $\text{clock}(\text{sid}_{\text{CommitTo}}, j)$. As the adversary may repeatedly commit to this message, we take the last commitment before phase i begins. A time-critical phase is tainted if the adversary returns an answer it has not committed to.

Definition 3.5 (Tainted Time-Critical Phase (Distance)) *A time-critical phase $\Pi_{\text{sid}}[k \dots k + 2\ell - 1] = (m_k, \dots, m_{k+2\ell-1})$ for $k, \ell \geq 1$ of a reader-adversary session sid , with the k -th message being received by the adversary, is tainted if the maximal j with $\Pi_{\text{sid}_{\text{CommitTo}}}[j] = (\text{sid}, k + 1, m_{k+1}^*)$ for some m_{k+1}^* and $\text{clock}(\text{sid}, k) > \text{clock}(\text{sid}_{\text{CommitTo}}, j)$ satisfies $m_{k+1}^* \neq m_{k+1}$ (or no such j exists).*

Definition 3.6 (Distance Fraud Resistance) *For an identification scheme \mathcal{ID} with parameters $(t_{\max}, T_{\max}, E_{\max}, N_c)$, a $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -distance-fraud adversary \mathcal{A} wins against \mathcal{ID} if the reader accepts in one of $q_{\mathcal{R}}$ reader-adversary sessions sid with at most T_{\max} tainted time-critical phases. Let $\text{Adv}_{\mathcal{ID}}^{\text{dist}}(\mathcal{A})$ be the probability of \mathcal{A} winning.*

3.5 Impersonation Resistance

We suggest a simple, but very strong definition of impersonation security as a basic requirement of identification in our concurrent setting. Thus even adversaries who actively take part in intertwined prover and verifier runs cannot impersonate the prover. Whereas the previous properties concern time-critical phases, impersonation security requires that an adversary cannot impersonate a tag in *lazy* phases. This ensures that the reader leaks no time-critical information to an invalid tag. Following the idea that parties should authenticate even if the time-critical phases are not executed, we consider projections $\Pi_{\text{sid}}^{\text{lazy}}[1 \dots]$ of $\Pi_{\text{sid}}[1 \dots]$ containing lazy phases transmissions only, and (not necessarily

consecutive) indices $\iota_{\text{sid}}^{\text{lazy}} = (i_1, i_2, \dots)$ of lazy phase messages. The adversary wins if a reader-adversary session succeeds and no adversary-tag session has the same “lazy transcript”, created via pure relaying.

Definition 3.7 (Impersonation Resistance) *In a distance-bounding identification scheme \mathcal{ID} with parameters $(t_{\max}, T_{\max}, E_{\max}, N_c)$ where \mathcal{R} always go first, a $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -impersonation adversary \mathcal{A} wins against \mathcal{ID} if \mathcal{R} accepts in a reader-adversary session sid such that no adversary-tag session sid^* has*

$$\Pi_{\text{sid}}^{\text{lazy}}[1 \dots] = \Pi_{\text{sid}^*}^{\text{lazy}}[1 \dots],$$

and

$$\text{clock}(\text{sid}, i) < \text{clock}(\text{sid}^*, i)$$

for any $i \in \iota_{\text{sid}}^{\text{lazy}} \cap \iota_{\text{sid}^*}^{\text{lazy}}$ s.t. \mathcal{R} has sent the i -th message to \mathcal{A} in sid , and

$$\text{clock}(\text{sid}, j) > \text{clock}(\text{sid}^*, j)$$

for any $j \in \iota_{\text{sid}}^{\text{lazy}} \cap \iota_{\text{sid}^*}^{\text{lazy}}$ such that the adversary has sent the j -th message to the reader in sid . Let $\text{Adv}_{\mathcal{ID}}^{\text{imp}}(\mathcal{A})$ be the probability that \mathcal{A} wins.

4 Relationship between Fraud Types

Impersonation security concerns lazy protocol phases, while Terrorist, Mafia, and distance fraud attack time-critical phases. In our framework we refute the idea in [35] that terrorist fraud resistance implies distance fraud resistance and show that all properties are independent. The formal proofs for each statement are shown in Appendix B.

Theorem 4.1 (Security Diagram — Informal) *If pseudorandom functions exist, the following holds:*

1. *There exists a distance-bounding identification scheme that is impersonation-secure, Mafia and distance fraud resistant, but not terrorist fraud resistant.*
2. *There exists a distance-bounding identification scheme that is impersonation-secure, Terrorist and Mafia fraud resistant, but not distance fraud resistant. Thus, terrorist fraud resistance does not imply distance fraud resistance.*
3. *There exists a distance-bounding identification scheme that is impersonation-secure, Terrorist and distance fraud resistant, but not Mafia fraud resistant. Thus, terrorist fraud resistance does not imply Mafia fraud resistance.*

Terrorist-Fraud Resistance. The enhanced Kim-Avoine scheme in Section 5 has all properties except for terrorist-fraud resistance. The reason it fails against terrorist attacks is that time-critical messages are predetermined by the lazy phase and can be revealed without disclosing the secret key (thus providing sufficient, but non-trivial offline help). In general, terrorist attacks are thwarted by interlinking authentication sessions, such that malicious tags (partially) reveal long-term secrets if they help the adversary. The difficulty in designing terrorist-fraud resistant schemes is formally ensuring that the simulator can extract the secret from the adversary and thus authenticate. The simulator’s only advantage is that it can rewind executions and get responses for different challenges.

Distance-Fraud Resistance. We separate distance-fraud resistance from the other properties by giving the tag a special key which makes time-critical responses predictable. Honest parties never use this key, but malicious tags may use it to commit distance fraud. Other security properties are unaffected, as the special key is never used by honest parties. Distance-fraud resistance depends on the unpredictability of each round’s answer. This is easily achieved by adding some time-critical rounds where tags echo random bits.

Mafia-Fraud Resistance. We show Mafia fraud resistance independence by starting with a protocol having all other security properties; the tag may use a bit to indicate that time-critical bits are flipped. Then a man-in-the-middle adversary can flip replies from an adversary-tag session and authenticate to the reader without tainting the phases. There are two options to prevent Mafia fraud attacks. Assume that in each fast phase the reader sends a random challenge. If the adversary correctly predicts the challenge in a reader impersonation, it can use the reply in the reader-adversary session without tainting the phase; for a wrong prediction, the adversary guesses the answer instead. The overall success is $\frac{3}{4}$ per round as in, e.g., the Hancke and Kuhn protocol [28]. The other option is to authenticate the reader by the fast phase challenges. Now the adversary-tag session in the above attack aborts for a wrong prediction, dropping the adversary’s success probability in the reader-adversary execution to $\frac{1}{2}$ for subsequent rounds. This is the strategy of the Kim-Avoine as discussed next.

5 Case Study: The Construction due to Kim and Avoine

The scheme in [31] is Mafia and Distance fraud resistant. We tweak it to add impersonation security, provide for noisy channels as in Section 2, then prove it secure in our framework. The proof relies on the fact that the nonce pairs exchanged in each run are quasi unique; also for any efficient adversary \mathcal{A}' the advantage $\text{Adv}_{\text{PRF}}^{\text{dist}}(\mathcal{A}')$ of distinguishing a pseudorandom function from a truly random one is small (see Appendix C for a formal proof).

Theorem 5.1 (Security Properties) *The distance-bounding identification scheme \mathcal{ID} in Fig. 3 with parameters $(T_{\max}, t_{\max}, E_{\max}, N_c)$ has the following properties:*

- It is not terrorist-fraud resistant.
- For any $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -impersonation adversary \mathcal{A} against \mathcal{ID} there exists a (t', q') -distinguisher \mathcal{A}' against PRF (with $t' = t + O(n)$ and $q' = q_{\mathcal{R}} + q_{\mathcal{T}} + q_{\text{OBS}}$) such that,

$$\mathbf{Adv}_{\mathcal{ID}}^{\text{imp}}(\mathcal{A}) \leq q_{\mathcal{R}} \cdot 2^{-|I|} + \mathbf{Adv}_{\text{PRF}}^{\text{dist}}(\mathcal{A}') + \binom{q_{\mathcal{R}} + q_{\text{OBS}}}{2} \cdot 2^{-|N_{\mathcal{R}}|} + \binom{q_{\mathcal{T}}}{2} \cdot 2^{-|N_{\mathcal{T}}|}.$$

- For any $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -distance-fraud adversary \mathcal{A} against \mathcal{ID} there is a (t', q') -distinguisher \mathcal{A}' against PRF (where $t' = t + O(n)$ and $q' = q_{\mathcal{R}} + q_{\mathcal{T}} + q_{\text{OBS}}$) such that, for $N_t = T_{\max} + E_{\max}$

$$\mathbf{Adv}_{\mathcal{ID}}^{\text{dist}}(\mathcal{A}) \leq q_{\mathcal{R}} \cdot \binom{N_c}{N_t} \left(\frac{7}{8}\right)^{N_c - N_t} + \mathbf{Adv}_{\text{PRF}}^{\text{dist}}(\mathcal{A}') + \binom{q_{\mathcal{R}} + q_{\text{OBS}}}{2} \cdot 2^{-|N_{\mathcal{R}}|}$$

- For any $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -Mafia-fraud adversary \mathcal{A} against \mathcal{ID} there exists a (t', q') -distinguisher \mathcal{A}' against PRF (where $t' = t + O(n)$ and $q' = q_{\mathcal{R}} + q_{\mathcal{T}} + q_{\text{OBS}}$) such that, for $N_t = T_{\max} + 2E_{\max}$

$$\begin{aligned} \mathbf{Adv}_{\mathcal{ID}}^{\text{mafia}}(\mathcal{A}) \leq & \frac{5}{8} \cdot q_{\mathcal{R}} \binom{N_c}{N_t} \cdot (N_c - N_t + 2) \cdot 2^{-(N_c - N_t)} + \mathbf{Adv}_{\text{PRF}}^{\text{dist}}(\mathcal{A}') \\ & + \binom{q_{\mathcal{R}} + q_{\text{OBS}}}{2} \cdot 2^{-|N_{\mathcal{R}}|} + \binom{q_{\mathcal{T}}}{2} \cdot 2^{-|N_{\mathcal{T}}|} \end{aligned}$$

For a single impersonation attempt and $T_{\max} = E_{\max} = 0$ we have up to small terms the (almost optimal) bound $\frac{1}{2}(N_c + 2) \cdot 2^{-N_c}$ for Mafia-Fraud resistance. The distance fraud resistance of $\frac{7}{8}$ per round is tight, corresponding to an adversary who sends v_i^0 in round i (v^0 is precomputed in the lazy phase).

Acknowledgments

We thank the anonymous reviewers for valuable comments.

The second author was supported by grants Fi 940/2-1 and Fi 940/3-1 of the German Research Foundation (DFG). This work was also supported by CASED (www.cased.de).

References

- [1] Abyneh, M.R.S.: Security analysis of two distance-bounding protocols. In: Proceedings of RFIDSec 2011. Lecture Notes in Computer Science, Springer (2011)
- [2] Avoine, G., Bingol, M.A., Karda, S., Lauradoux, C., Martin, B.: A formal framework for analyzing rfid distance bounding protocols. In: Journal of Computer Security - Special Issue on RFID System Security, 2010 (2010)

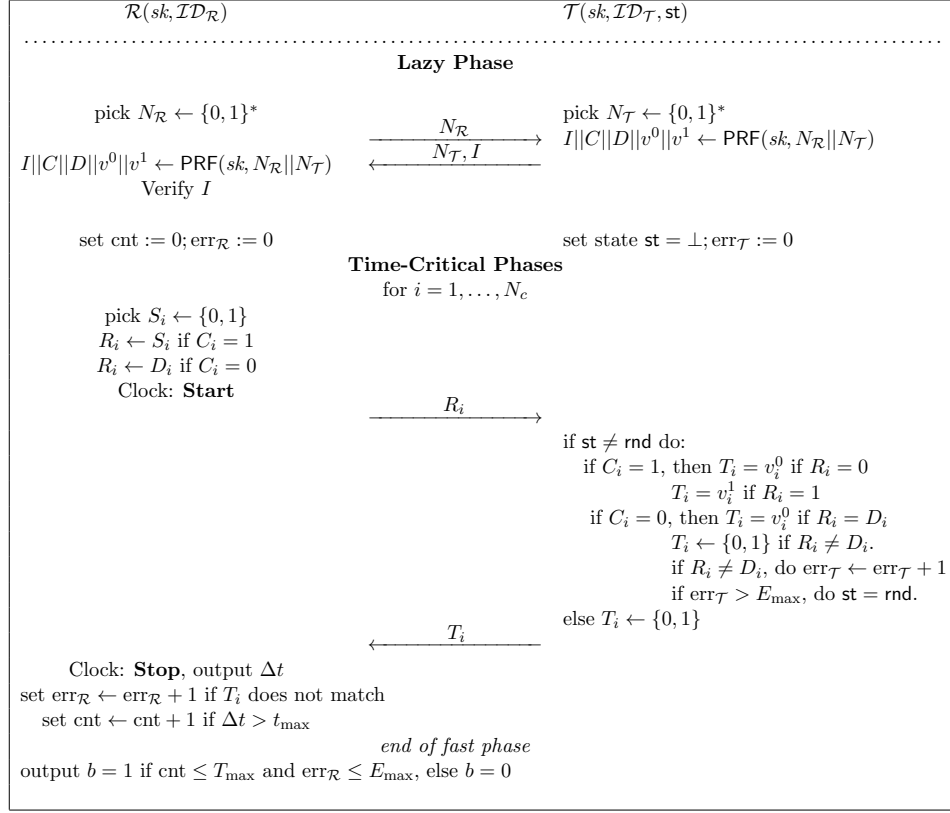


Figure 3: Enhanced Kim/Avoine protocol.

- [3] Avoine, G., Martin, B., Martin, T.: Optimal security limits of rfid distance bounding protocols. In: RFIDSec 2010. pp. 220 – 238
- [4] Avoine, G., Tchamkerten, A.: An efficient distance bounding rfid authentication protocol: Balancing false-acceptance rate and memory requirement. In: Information Security. Lecture Notes in Computer Science, vol. 5735, pp. 250–261. Springer-Verlag (2009)
- [5] Bellare, M., Goldreich, O.: Proving computational ability. <http://www.wisdom.weizmann.ac.il/~oded/PS/poa.ps> (1992)
- [6] Bellare, M., Pointcheval, D., Rogaway, P.: Authenticated key exchange secure against dictionary attacks. In: Advances in Cryptology — Eurocrypt 2000. Lecture Notes in Computer Science, vol. 1807, pp. 139–155. Springer-Verlag (2000)
- [7] Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Advances in Cryptology — Crypto '93. Lecture Notes in Computer Science, vol. 773, pp. 232–249. Springer-Verlag (1994)

- [8] Brands, S.: Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy. MIT Press (2000)
- [9] Brands, S., Chaum, D.: Distance-bounding protocols. In: Advances in Cryptology — Eurocrypt’93. pp. 344–359. Lecture Notes in Computer Science, Springer-Verlag (1993)
- [10] Bringer, J., Chabanne, H.: Trusted-hb: A low-cost version of hb⁺ secure against man-in-the-middle attacks. Transactions on Information Theory 54(9), 4339–4342 (2008)
- [11] Bussard, L., Bagga, W.: Distance-bounding proof of knowledge to avoid real-time attacks. Security and Privacy in the Age of Ubiquitous Computing 181, 222–238 (2005)
- [12] Camenisch, J., Lysyanskaya, A.: An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In: Advances in Cryptology — Eurocrypt. Lecture Notes in Computer Science, vol. 2045, pp. 93–118. Springer-Verlag (2001)
- [13] Capkun, S., Butty’an, L., Hubaux, J.P.: Sector: Secure tracking of node encounters in multi-hop wireless networks. In: ACM Workshop on Security of Ad Hoc and Sensor Networks - SASN. pp. 21 – 32. ACM Press (2003)
- [14] Carluccio, D., Kasper, T., Paar, C.: Implementation details of a multi purpose iso 14443 rfidtool. In: Printed handout of Workshop on RFID Security - RFIDSec 06 (July 2006)
- [15] Chandran, N., Goyal, V., Moriarty, R., Ostrovsky, R.: Position based cryptography. In: Advances in Cryptology — Crypto. Lecture Notes in Computer Science, vol. 5677, pp. 391–407. Springer-Verlag (2009)
- [16] Clulow, J., Hancke, G.P., Kuhn, M.G., Moore, T.: So near and yet so far: Distance-bounding attacks in wireless networks. In: European Workshop on Security and Privacy in Ad-Hoc and Sensor Networks. Lecture Notes in Computer Science, vol. 4357, pp. 83–97. Springer-Verlag (2006)
- [17] Desmedt, Y.: Major security problems with the ‘unforgeable’ (feige)-fiat-shamir proofs of identity and how to overcome them. In: SecuriCom. pp. 15–17. SEDEP Paris, France (1988)
- [18] Drimer, S., Murdoch, S.J.: Keep your enemies close: distance bounding against smartcard relay attacks. In: Proc. of the 16-th USENIX Security Symposium on USENIX Security Symposium, article no. 7. ACM Press (2007)
- [19] Duc, D., Kim, K.: Securing hb⁺ against grs man-in-the-middle attack. In: Symposium on Cryptography and Information Security (SCIS). The Institute of Electronics, Information and Communication Engineers (2007)

- [20] Dwork, C., Lotspiech, J.B., Naor, M.: Digital signets: Self-enforcing protection of digital information (preliminary version). In: STOC. pp. 489–498 (1996)
- [21] Francillon, A., Danev, B., Capkun, S.: Relay attacks on passive keyless entry and start systems in modern cars. Cryptology ePrint Archive, Report 2010/332 (2010), ePRINTURL
- [22] Gilbert, H., Robshaw, M., Sibert, H.: An active attack against hb+ - a provably secure lightweight authentication protocol. Cryptology ePrint Archive, Report 2005/237 (2005), ePRINTURL
- [23] Goldreich, O., Pfitzmann, B., Rivest, R.L.: Self-delegation with controlled propagation - or - what if you lose your laptop. In: Advances in Cryptology — Crypto. Lecture Notes in Computer Science, vol. 1462, pp. 153–168. Springer-Verlag (1998)
- [24] Haataja, K., Toivanen, P.: Two practical man-in-the-middle attacks on bluetooth secure simple pairing and countermeasures. Transactions on Wireless Communications 9(1), 384–392 (2010)
- [25] Hancke, G.P.: A practical relay attack on iso 14443 proximity cards. <http://www.cl.cam.ac.uk/gh275/relay.pdf> (2005)
- [26] Hancke, G.: Distance bounding publication database. <http://www.rfidblog.org.uk/db.html> (2010)
- [27] Hancke, G.P.: Design of a secure distance-bounding channel for rfid. Journal of Network and Computer Applications (2010)
- [28] Hancke, G.P., Kuhn, M.G.: An rfid distance bounding protocol. In: SECURECOMM. pp. 67–73. ACM Press (2005)
- [29] Hopper, N.J., Blum, M.: Secure human identification protocols. In: Proceedings of the 7th International Conference on the Theory and Application of Cryptology and Information Security ADVCRYPTO. Lecture Notes in Computer Science, vol. 2248, pp. 52–66. Springer-Verlag (2001)
- [30] Juels, A.: Rfid security and privacy: a research survey. IEEE Journal on Selected Areas in Communications 24(2), 381–394 (2006)
- [31] Kim, C.H., Avoine, G.: Rfid distance bounding protocol with mixed challenges to prevent relay attacks. In: Proceedings of the 8th International Conference on Cryptology and Networks Security (CANS 2009). Lecture Notes in Computer Science, vol. 5888, pp. 119–131. Springer-Verlag (2009)
- [32] Leng, X., Mayes, K., Markantonakis, K.: Hb-mp+ protocol: An improvement on the hb-mp protocol. In: International Conference on RFID. pp. 118–124. IEEE Computer Society Press (2008)

- [33] Ouafi, K., Overbeck, R., Vaudenay, S.: On the security of hb# against a man-in-the-middle attack. In: Advances in Cryptology — Asiacrypt 2008. Lecture Notes in Computer Science, vol. 5350, pp. 108–124. Springer (2008)
- [34] Rasmussen, K.B., Čapkun, S.: Realization of rf distance bounding. USENIX Security Symposium (2010)
- [35] Reid, J., Nieto, J.M.G., Tang, T., Senadji, B.: Detecting relay attacks with timing-based protocols. In: ASIACCS. pp. 204–213. ACM Press (2007)
- [36] Singele, D., Preneel, B.: Distance bounding in noisy environments. In: European Workshop on Security in Ad-hoc and Sensor Networks – ESAS. Lecture Notes in Computer Science, vol. 4572, pp. 101 – 115. IEEE Computer Society Press (2007)
- [37] Trujillo-Rasua, R., Martin, B., Avoine, G.: The poulidor distance-bounding protocol. In: RFIDSec 2010. pp. 239 – 257
- [38] Yung, M.: Zero-knowledge proofs of computational power. In: Advances in Cryptology — Eurocrypt ’89. Lecture Notes in Computer Science, vol. 434, pp. 196–207. Springer-Verlag (1990)

A Related Cryptographic Concepts

Position-Based Cryptography. Chandran et al. [15] recently introduced the notion of position-based cryptography, where a set of verifiers wishes to check whether a prover is at a position P or not. Apart from broadcasting and sending directional messages to the prover, verifiers can also securely communicate with one other. Provers can only broadcast or send directional messages to a verifier. The so-called vanilla model also considers several colluding adversaries, which may broadcast, send directional messages, and communicate securely with each other. Communication time is measured according to distance.

Chandran et al. [15] prove an important impossibility result in this setting, namely that secure positioning cannot be achieved in the vanilla model (with collusions). In the proof, colluding adversaries are placed centric around point P , where the prover should be, closer to the verifiers than P . The adversaries can communicate with one another within time α . Then each adversary impersonates a copy of the prover for the verifier in proximity and answers each message of this verifier on the prover’s behalf, but faster since it is closer. The time gained is used to “synchronize” with the other adversaries such that each prover copy is always up-to-date, with a time delay of α . Overall, this collusion strategy successfully simulates the presence of a prover at P .

How does this model (and impossibility result) relate to RFID distance bounding attacks? The main difference to Mafia attacks is that the Mafia adversary doesn’t know the key shared between prover and verifier. Even colluding adversaries are thus unable to answer on the tag’s behalf quickly and then “synchronize” (the impossibility result

requires that the adversaries have all the information the prover has). Also, Mafia attacks are rather malleability attacks, as the adversary has access to other copies of the protocol (but with switched roles) and tries to take advantage of this, with some restrictions due to the distance. Position-based cryptography considers instead a single protocol run with multiple verifiers.

If the collusion adversary in [15] holds the secret key in a setting resembling our distance fraud attacks, each prover copy is in fact closer than point P . The success of such an adversary is in tune with our assumption that legitimate provers may authenticate within close proximity.

It is unclear how position-based cryptography compares to terrorist attacks, where the adversary has some limited help. This idea is closer related to the concept of non-transferability (see below), discouraging users to reveal parts of their secrets [20, 23, 8, 12]. Also, establishing exact position is impossible in practice for RFID tags, as their response times have a high variance. Only by using an extremely large number of readers is this possible.

From Useful Help to Undesired Disclosure. Several (public-key based) approaches in the literature [20, 23, 8, 12] associate the distribution of parts of a secret to leaking (external) personal information like a credit card number [20, 8], or the entire (internal) secret itself [23, 12]. We discuss the latter case, which resembles our terrorist attacks.

Goldreich et al [23] introduce self-delegatable schemes, where users generate secondary keys with restricted rights, usable in more vulnerable environments e.g. laptops; the secondary keys are authenticated through a long-term key. Losing a few such keys should not harm the security of other keys; however, since only self-delegation is supported, leaking too many secondary keys endangers the security of the user’s long-term key.

This idea is mirrored in our terrorist attack resistance: if the support the adversary receives from the tag allows authentication, then this leaks essential information about the tag’s secret. The main differences in the model are: that [23] consider the public-key setting only (where server certification of secondary keys is used), that they investigate signature-leakage only, and that no online help (with restriction due to the distance) is available. Also, all the schemes in [23] rely on public-key cryptography and non-interactive zero-knowledge proofs, and are unsuitable for RFID.

Camenisch and Lysanskaya [12] model transferability of anonymous credentials. This “all-or-nothing” approach associates sharing secret information (pseudonyms or credentials only) to recovery of users’ full secret. This is again similar to terrorist resistance but [12] do not formally model attacks and security. The use of public-key infrastructures here also makes the idea inapplicable to RFID.

B Relationships of Security Notions

Proposition B.1 *If a pseudorandom function $\mathcal{PRF}=(\mathsf{Kg}, \text{PRF})$ exists, then there exists a distance-bounding identification scheme $\mathcal{ID}_{\text{terror}}^{\text{imp}, \text{dist}, \text{mafia}} = (\mathsf{Kg}_{\text{PRF}}, \mathcal{R}_{\text{PRF}}, \mathcal{T}_{\text{PRF}})$ with*

parameters $(t_{\max}, T_{\max}, E_{\max}, N_c)$ that is Mafia fraud and distance fraud resistant, and secure against impersonations, but that is vulnerable to terrorist fraud attacks.

Proof. This scheme is the enhanced Kim/Avoine scheme shown in Section 5. \square

Proposition B.2 *If a distance-bounding identification scheme \mathcal{ID}' exists such that it is resistant to mafia and terrorist fraud, and impersonation resistant, then there exists a distance-bounding identification scheme $\mathcal{ID} = \mathcal{ID}_{\text{dist}}^{\text{imp,mafia,terror}} = (\text{Kg}, \mathcal{R}, \mathcal{T})$ with parameters $(t_{\max}, T_{\max}, E_{\max}, N_c)$ that is still secure against impersonations and terrorist and Mafia fraud resistant, but that is vulnerable to distance fraud adversaries. In particular, terrorist fraud resistance does not imply distance fraud resistance.*

Proof. Consider an identification scheme \mathcal{ID}' that is Mafia and Terrorist fraud resistant and secure against impersonations. Modify \mathcal{ID}' to obtain \mathcal{ID} as follows: apart from any secret(s) the tag and the reader share in \mathcal{ID}' , add another secret key sk^* . In the lazy phase, the tag's first message to the reader will now be preceded by a bit b and a bitstring V of length $\|sk^*\|$. An honest tag always sends $b = 0$ and $V = \mathbf{0}$, the all-zero vector. The reader parses the beginning bit and checks it. If the received bit is 0, protocol \mathcal{ID}' is followed exactly, in its original form (and V is ignored). Else, if the received bit is 1, the reader skips any checks on the tag's lazy phase messages that appear in \mathcal{ID} and checks that $V = sk^*$. If so, the reader goes on to the time-critical phases and always expects a 0 response from the tag if the response comes in time.

The following statements hold for \mathcal{ID} :

- For any $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -impersonation adversary \mathcal{A} against \mathcal{ID} there exists a $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -impersonation adversary \mathcal{A}' against \mathcal{ID}' such that

$$\text{Adv}_{\mathcal{ID}}^{\text{imp}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{ID}'}^{\text{imp}}(\mathcal{A}') + 2^{-|sk^*|}.$$

- For any $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -Mafia adversary \mathcal{A} against \mathcal{ID} there exists a $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -Mafia adversary \mathcal{A}' against \mathcal{ID}' such that

$$\text{Adv}_{\mathcal{ID}}^{\text{mafia}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{ID}'}^{\text{mafia}}(\mathcal{A}') + 2^{-|sk^*|}.$$

- For any $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q'_{\mathcal{T}}, q_{\text{OBS}})$ -terrorist-fraud adversary \mathcal{A} against \mathcal{ID} there exists a $t_{\mathcal{S}}$ -simulator \mathcal{S} such that for any \mathcal{T}' running in time $t_{\mathcal{T}'}$ it holds that,

$$\text{Adv}_{\mathcal{ID}}^{\text{terror}}(\mathcal{A}, \mathcal{S}, \mathcal{T}) \leq 0.$$

- This scheme is not resistant to distance fraud.

The first bound follows from the bound of the underlying construction. Honest users do not send sk^* , therefore an adversary against impersonation does not know sk^* . An impersonation adversary $\mathcal{A}_{\mathcal{ID}}^{\text{imp}}$ against the distance-bounding identification scheme \mathcal{ID}

must either break the underlying schema \mathcal{ID}' or guess the secret key sk^* . This adds a term $2^{-|sk^*|}$ to the impersonation security bound.

Similarly, honest participants will never send 1 in the first message and subsequently $V = sk^*$. Therefore, a Mafia fraud adversary $\mathcal{A}_{\mathcal{ID}}^{\text{mafia}}$ against \mathcal{ID} can either guess sk^* (and succeed with probability 1) or can try to break the underlying construction. Therefore, the advantage of a Mafia adversary against \mathcal{ID} is the advantage of a Mafia adversary against \mathcal{ID}' with an added term $2^{-|sk^*|}$.

Let $\mathcal{A}_{\mathcal{ID}}^{\text{terror}}$ be an adversary against scheme \mathcal{ID} . We build a simulator $\mathcal{S}_{\mathcal{ID}}$ as follows. The simulator looks first in the transcripts of $\mathcal{A}_{\mathcal{ID}}^{\text{terror}}$; if sk^* was used, the simulator reuses this value and wins with probability 1. If sk^* was not used, the simulator removes the 0 bit from each execution, and also the string $V = \mathbf{0}$, then following the protocol of \mathcal{ID}' . The adversary $\mathcal{A}_{\mathcal{ID}}^{\text{terror}}$ is now an adversary against \mathcal{ID}' and therefore there must exist a simulator \mathcal{S}^* such that $p_{\mathcal{A}_{\mathcal{ID}}^{\text{terror}}} - p_{\mathcal{S}^*} \leq 0$. The simulator \mathcal{S} against \mathcal{ID} uses \mathcal{S}^* against \mathcal{ID}' as a black box, and adds a 0 bit and the string $V = \mathbf{0}$ to the execution. Clearly, $p_{\mathcal{S}} = p_{\mathcal{S}^*}$. Therefore the bound in this statement is achieved.

An adversary $\mathcal{A}_{\mathcal{ID}}^{\text{dist}}$ against \mathcal{ID} is a legitimate tag, which therefore knows the value of sk^* . This adversary will send as its first message 1 and later $V = sk^*$. During the fast phase, the adversary will commit to each round a response of 0. This adversary succeeds with probability 1. □

Proposition B.3 *If a distance-bounding identification scheme \mathcal{ID}' exists such that it is secure against impersonations and resistant to terrorist and distance fraud, then there exists a distance-bounding identification scheme $\mathcal{ID} = \mathcal{ID}_{\text{mafia}}^{\text{imp,dist,terror}} = (\text{Kg}, \mathcal{R}, \mathcal{T})$ with parameters $(t_{\max}, T_{\max}, E_{\max}, N_c)$ that is still secure against impersonations and resistant to terrorist and distance fraud, but that is vulnerable to Mafia fraud adversaries. In particular, terrorist fraud resistance does not imply Mafia fraud resistance.*

Proof. Consider an identification scheme \mathcal{ID}' that is secure against impersonations and resistant to terrorist and distance fraud. Modify \mathcal{ID}' to obtain \mathcal{ID} as follows: apart from any secret(s) the tag and the reader share in \mathcal{ID}' , add another secret key sk^* . In the lazy phase, the tag's first message to the reader will now be preceded by a bit b . An honest tag always sends $b = 0$. The reader parses the beginning bit and checks it. If the received bit is 0, protocol \mathcal{ID}' is followed exactly, in its original form. Assume that in the N_c time-critical rounds in \mathcal{ID}' , the tag sends responses T_i , which are verified by the reader. If the bit received by the reader during the lazy phase of \mathcal{ID} is $b = 1$, in each of the time-critical rounds of this protocol, the reader will expect response \bar{T}_i , i.e. the bit(s) of the response are flipped.

The following statements hold for \mathcal{ID} :

- For any $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -impersonation adversary \mathcal{A} against \mathcal{ID} there exists a $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -impersonation adversary \mathcal{A}' against \mathcal{ID}' such that

$$\text{Adv}_{\mathcal{ID}}^{\text{imp}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{ID}'}^{\text{imp}}(\mathcal{A}').$$

- For any $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -distance adversary \mathcal{A} against \mathcal{ID} there exists a $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q_{\text{OBS}})$ -distance adversary \mathcal{A}' against \mathcal{ID}' such that

$$\mathbf{Adv}_{\mathcal{ID}}^{\text{dist}}(\mathcal{A}) \leq \mathbf{Adv}_{\mathcal{ID}'}^{\text{dist}}(\mathcal{A}').$$

- For any $(t, q_{\mathcal{R}}, q_{\mathcal{T}}, q'_{\mathcal{T}}, q_{\text{OBS}})$ -terrorist-fraud adversary \mathcal{A} against \mathcal{ID} there exists a $t_{\mathcal{S}}$ -simulator \mathcal{S} such that for any \mathcal{T}' running in time $t_{\mathcal{T}'}$ it holds that,

$$\mathbf{Adv}_{\mathcal{ID}}^{\text{terror}}(\mathcal{A}, \mathcal{S}, \mathcal{T}) \leq 0.$$

- This scheme is not resistant to mafia fraud.

The first bound follows from the bound of the underlying construction, as the same adversary that succeeds in an impersonation attack against \mathcal{ID}' will succeed in an impersonation attack against \mathcal{ID} .

We consider the second statement. Let $\mathcal{A}_{\mathcal{ID}}^{\text{dist}}$ be a distance-fraud adversary that succeeds against the identification scheme \mathcal{ID} . We show how to use this adversary to construct a distance-fraud adversary $\mathcal{A}_{\mathcal{ID}'}^{\text{dist}}$ against the underlying scheme \mathcal{ID}' . If $\mathcal{A}_{\mathcal{ID}}^{\text{dist}}$ forwards a 0 as its first lazy-phase message, the adversary $\mathcal{A}_{\mathcal{ID}'}^{\text{dist}}$ forwards the exact responses T_i which $\mathcal{A}_{\mathcal{ID}}^{\text{dist}}$ commits to in each round. Otherwise, if a 1 is used during the lazy-phase, the adversary simply flips the bits.

Let $\mathcal{A}_{\mathcal{ID}}^{\text{terror}}$ be an adversary against scheme \mathcal{ID} . We build a simulator $\mathcal{S}_{\mathcal{ID}}$ as follows. The simulator removes the 0/1 bit from each execution, following the protocol of \mathcal{ID}' . The adversary $\mathcal{A}_{\mathcal{ID}}^{\text{terror}}$ is now an adversary against \mathcal{ID}' and therefore there must exist a simulator \mathcal{S}^* such that $p_{\mathcal{A}_{\mathcal{ID}}^{\text{terror}}} - p_{\mathcal{S}^*} \leq 0$. The simulator \mathcal{S} against \mathcal{ID} uses \mathcal{S}^* against \mathcal{ID}' as a black box, and adds a 0 bit to the execution. Clearly, $p_{\mathcal{S}} = p_{\mathcal{S}^*}$. Therefore the bound in this statement is achieved.

An adversary $\mathcal{A}_{\mathcal{ID}}^{\text{mafia}}$ against \mathcal{ID} opens a reader-adversary session sid and an adversary-tag session sid^* . During the lazy phase, the adversary relays messages from one session to another (this is permitted during the lazy phase), but in the first message flips the 0 bit sent by the honest tag in sid^* to a 1 bit. During each of the time-critical rounds, the adversary relays the challenge bits of the reader, but flips the return bits T_i sent by the tag. Since the communication is not simply relayed, this is a valid mafia adversary, which wins with probability 1. \square

C Security Proof of the Protocol of Kim and Avoine

Proof. The protocol is not terrorist-fraud resistant: \mathcal{T}' can forward adversary \mathcal{A} the value $I||C||S||v^0||v^1$. Now \mathcal{A} authenticates successfully; a simulator can't authenticate, however, as a fresh session has new nonces in the lazy phase.

We prove Mafia-fraud resistance as follows:

1. (1) Show that we can safely replace the honest parties' PRF output by independent random values $I||C||D||v^0||v^1$ for new nonces $(N_{\mathcal{R}}, N_{\mathcal{T}})$;

2. (2) Show quasi-uniqueness of nonce pairs except in 1 adversary-tag session and one reader-adversary session s.t. \mathcal{A} relays the nonces;
3. (3) Bound \mathcal{A} 's winning probability in time-critical phases for at most one adversary-tag interaction.

For the first step we claim that replacing the PRF-values by random (but consistent) values can at most decrease the adversary's success probability by the distinguishing advantage for PRF. This can be seen easily by construction adversary \mathcal{A}' against PRF via black-box simulation of \mathcal{A} , each time applying the random or pseudorandom oracle to nonce pairs on behalf of the honest parties. Finally, \mathcal{A}' checks if \mathcal{A} succeeds in some reader-adversary session (such that there is no adversary-tag session with more than T_{\max} tainted time-critical phases for this session) and outputs 1 if this happens. The distinguishing advantage of \mathcal{A}' then corresponds to the decrease of the success probability of \mathcal{A} when switching to random values $I||C||D||v^0||v^1$.

Next consider the adversary \mathcal{A} mounting a Mafia fraud attack and all the pairs of nonces appearing the attack. Assume that there exist two sessions (between adversary and tag or reader, or between both honest parties) with the same pair $(N_{\mathcal{R}}, N_{\mathcal{T}})$. Then we claim that this can only be a reader-adversary session and an adversary-tag session, except with probability

$$\binom{q_{\mathcal{R}} + q_{\text{OBS}}}{2} \cdot 2^{-|N_{\mathcal{R}}|} + \binom{q_{\mathcal{T}}}{2} \cdot 2^{-|N_{\mathcal{T}}|}.$$

This holds as for each two executions for the reader resp. tag the nonce of this party is picked at random. If there were three identical nonce pairs in some executions then two of them would be either in the at most $q_{\mathcal{R}} + q_{\text{OBS}}$ executions with the reader, or in the $q_{\mathcal{T}}$ executions with the tag. Such collisions can only occur with the above probability.

Declare the adversary now to lose if such a collision appears elsewhere, decreasing \mathcal{A} 's success probability only be the above negligible term, but allowing us to consider such collision-free executions from now on. In particular, except for the matching session all other values $I||C||D||v^0||v^1$ appearing in the attack are independent.

Let sid be a reader-adversary session where \mathcal{A} successfully impersonates to \mathcal{R} . By assumption at most one other adversary-tag session sid^* has the same nonce pair. If sid^* exists, it taints sid with high probability (if sid^* doesn't exist, \mathcal{A} can't benefit from sid^*). Suppose now that sid^* taints at most T_{\max} time-critical phases of sid . Assume for the moment that $E_{\max} = 0$; we make provisions for $E_{\max} > 0$ later.

Consider an untainted time-critical phase of sid where \mathcal{R} sends R_i and expects T_i , i.e. assume \mathcal{A} successfully passed the first $i - 1$ time-critical phases. There are four strategies for the adversary in this i -th phase:

GO-EARLY. In session sid^* \mathcal{A} sends bit R_i^* to \mathcal{T} before receiving R_i (i.e., $\text{clock}(\text{sid}, i+2) > \text{clock}(\text{sid}^*, i+2)$). As R_i is random and independently chosen, $R_i^* \neq R_i$ w.p. $\frac{1}{2}$ – then \mathcal{A} doesn't receive T_i in sid^* and must guess T_i in sid . Also, session sid^* becomes invalid with probability $\frac{1}{4}$.

GO-LATE. In session sid , \mathcal{A} replies to R_i with T_i before receiving T_i^* in session sid^* ($\text{clock}(\text{sid}, i + 3) < \text{clock}(\text{sid}^*, i + 3)$). Now \mathcal{A} wins the phase w.p. $\frac{1}{2}$.

MODIFY-IT. \mathcal{A} receives R_i in sid , sends R_i^* in sid^* , gets T_i^* in sid^* , and forwards T_i in sid . This scheduling is pure relay, but $R_i \neq R_i^*$ or $T_i \neq T_i^*$. If R_i^* is wrong then T_i^* was never sent by \mathcal{T} in sid^* and \mathcal{A} can only guess T_i w.p. $\frac{1}{2}$; if $R_i = R_i^*$ then $T_i \neq T_i^*$ makes the reader reject.

TAINT-IT. The adversary taints this phase of sid through sid^* .

Tainting the phase makes \mathcal{R} accept with probability 1, deducting 1 from the remaining taintable phases. The Go-Late and Modify-it Strategy both succeed w.p. at most $\frac{1}{2}$. Go-Early succeeds w.p. $\frac{3}{4}$, inactivating sid^* w.p. $\frac{1}{2}$. Assume that \mathcal{A} taints the last T_{\max} time-critical phases (else we renumber the phases). For the other $P := N_c - T_{\max}$ phases let pass_i denote the event that \mathcal{A} passes phase i of sid . We have

$$\text{Prob} \left[\bigwedge_{j=i}^P \text{pass}_j \mid \bigwedge_{j=1}^{i-1} \text{pass}_j \right] \leq \frac{5}{8} \cdot \text{Prob} \left[\bigwedge_{j=i+1}^P \text{pass}_j \mid \bigwedge_{j=1}^i \text{pass}_j \right] + \frac{1}{2} \cdot \frac{1}{2} \cdot 2^{-P+i+1}$$

The first term captures the success of Go-Late, Modify-It, and correct Go-Early-prediction. The second term covers incorrect Go-Early prediction (w.p. $\frac{1}{4}$); now sid^* is inactivated, and \mathcal{A} must guess T_i for this and the next $P - i - 1$ rounds (the responses are independent). Expanding the probabilities we obtain

$$\text{Prob} \left[\bigwedge_{j=1}^P \text{pass}_j \right] \leq 2^{-P} + \sum_{j=0}^{P-1} \frac{5}{8} \cdot 2^{-j} \cdot 2^{-P+j} = \frac{5}{8} \cdot (P + 2) \cdot 2^{-P}.$$

We sum over $q_{\mathcal{R}}$ reader-adversary sessions, distribute $T_{\max} + E_{\max}$ “jokers” on the reader side and E_{\max} on the tag side, and obtain the claimed bound.

For impersonation security, the only way to generate colliding nonce pairs (and produce authentication string I) is by lazy phase relay, which is an invalid impersonation attack. For distinct nonce pairs, the probability that \mathcal{A} sends a correct I in a reader-adversary session is: $q_{\mathcal{R}} \cdot 2^{-|I|}$ plus the distinguishing advantage for the PRF plus the probability of colliding nonces.

Distance-bounding (the third statement) is proved as above: once the pseudorandom values are replaced by truly random ones, the probability that $C_i = 1$ and $v_i^0 \neq v_i^1$ is at least $\frac{1}{4}$ for round i . Since \mathcal{A} can commit only then, \mathcal{A} fails with probability at least $\frac{1}{8}$. Overall, \mathcal{A} succeeds only w.p. $\frac{7}{8}$ per round, except for a number $T_{\max} + E_{\max}$ of phases. \square