

Expected loss analysis for authentication in constrained channels

Christos Dimitrakakis. E-mail: chrdimi@chalmers.ch^{a,*} Aikaterini Mitrokotsa^a and Serge Vaudenay^b

^a *Chalmers University of Technology, Gothenburg, 41296, Sweden*

Email: {chrdimi,aikmitr}@chalmers.se

^b *EPFL, Switzerland*

E-mail: serge.vaudenay@epfl.ch

Abstract. We derive bounds on the expected loss for authentication protocols in channels which are constrained due to noisy conditions and communication costs. This is motivated by a number of authentication protocols, where at least some part of the authentication is performed during a phase, lasting n rounds, with no error correction. This requires assigning an acceptable threshold for the number of detected errors and taking into account the cost of incorrect authentication and of communication. This paper describes a framework enabling an expected loss analysis for all the protocols in this family. Computationally simple methods to obtain nearly optimal values for the threshold, as well as for the number of rounds are suggested and upper bounds on the expected loss, holding uniformly, are given. These bounds are tight, as shown by a matching lower bound. Finally, a method to adaptively select both the number of rounds and the threshold is proposed for a certain class of protocols.

Keywords: authentication, decision theory, distance bounding, expected loss, noise

1. Introduction

Traditionally, authentication is assumed to be taking place on an error-free channel [1,2], and error analysis is performed separately from cryptographic analysis of protocols. However, a number of authentication protocols have been proposed [3,4,5,6,7,8,9,10,11,12,13], where at least some part of the authentication is performed during a challenge-response phase lasting n rounds with no error correction. These protocols are known as *distance-bounding* protocols and are employed as the main countermeasure against relay attacks by timing transmission delays. However, channel noise necessitates the acceptance of some responses which are not perfectly correct. In addition, increasing n carries a significant cost. The problem is to strike an optimal balance between the increased number of rounds, the probability of falsely authenticating an illegitimate party, and failing to authenticate a legitimate party.

For example the *rapid-bit exchange* phase in an RFID distance-bounding protocol (i.e. [9]), satisfies these criteria. This is a phase, lasting n rounds, where an equal number of challenges and responses,

*Corresponding author

consisting of single bits, are sent. Since there is no error correction during that phase, some responses received from a legitimate user may be erroneous, due to transmission errors in either direction. This necessitates the use of a tolerance threshold τ , such that if the number of erroneous responses is lower than τ , the communicating party is nevertheless authenticated. Unfortunately, while increasing the number of rounds decreases the probability of incorrect authentication, it also increases communication costs, which is particularly undesirable in power-constrained devices such as RFID tags. The problem then is how to select n, τ in an optimal manner and how to define optimality in the first place.

This paper performs an expected loss analysis of the authentication problem. This is necessary, because of the non-trivial cost of increasing the number of rounds n , the lack of an error-free channel and the need to trade off optimally the costs of incorrectly authenticating an illegitimate entity, or failing to authenticate a legitimate one.

We place the problem in a decision-theoretic framework. We assign a loss ℓ_A to the event that we authenticate a malicious party A —which we call the *attacker*—a loss ℓ_U to the event that we fail to authenticate a valid party U —which we call the *user*—and a loss ℓ_B for each round of the challenge-response phase, such that the total communication cost is $n\ell_B$. Adding a non-negligible cost to the communications is of fundamental importance in resource-constrained environments. Otherwise, n can be made as large as necessary to make the probability of authentication mistakes infinitesimal. Our goal is to select n and τ so as to minimise the *expected loss* $\mathbb{E}L$ of the authentication system. This is achieved through a finite sample analysis.

We should note here that a preliminary version of this paper was presented at INFOCOM [14] and as a technical report at [15]. The current paper features corrections, an extended theoretical analysis, additional experiments and a more realistic attacker model. The main differences are that we consider a stronger attacker model, present an improved choice for the threshold when the gap between the attacker and user is large, as well as new experimental results where we detail the effect of parameters on the choice of threshold and number of rounds. Finally, we present a more extensive experimental comparison with a likelihood-ratio threshold and the exact minimax solution.

The paper is organised as follows. Section 2 introduces notation and our framework. Section 3 contains the *expected loss* analysis under noise. In particular, Sec. 3.1 suggests a method to calculate the threshold accompanied by finite sample upper loss bounds and a matching lower bound, while Sec. 3.2 provides a further loss bound by selecting a near-optimal number of rounds n . These results only require two reasonable assumptions: that the expected error of the attacker is higher than that of the user and that the errors are independent in each round, something that can be achieved by appropriate protocol design. In particular, we analyse methods for selecting a nearly-optimal threshold and number of rounds, prove bounds on the expected loss for our choices, show their dependence on the problem parameters and discuss their relation with the optimal (minimax) solution. Section 4 applies the above analysis to two existing RFID protocols, and shows the effect of channel noise on the expected loss, as well as on our bounds. Since the above analysis requires bounds on the errors made by users and attackers to be known, Section 5 suggests a high-probability method for estimating the channel noise in order to obtain them, under a powerful attacker model. Section 6 presents two sets of simulation experiments. In the first, we investigate the stability of the noise evaluation method under varying noise. This is done both for our method of selecting the threshold with a Bayesian likelihood ratio test, which requires additional assumptions. In the second set of experiments, we directly compare the performance of those two methods as noise and the other problem parameters vary with the minimax solution. Finally, Sec. 7 concludes the paper with a discussion of related work. For completeness, the appendix provides known auxiliary results for the finite sample and the Bayesian likelihood ratio test derivations.

2. Preliminaries

We consider sequences $x = x_1, \dots, x_n$ with all x_i in some alphabet \mathcal{X} and $x \in \mathcal{X}^n$. We write $\mathcal{X}^* \triangleq \bigcup_{n=0}^{\infty} \mathcal{X}^n$ for the set of all sequences. We use \triangleq to indicate a definition. $\mathbb{P}(A)$ denotes the probability of event A , while \mathbb{E} denotes expectations so that $\mathbb{E}(X | A) = \sum_{u \in \Omega} u \mathbb{P}(X = u | A)$ denotes the conditional expectation of a random variable $X \in \Omega$ when A is true. In addition, $\mathbb{I}\{A\}$ is an indicator function equal to 1 when A is true and 0 otherwise. Finally, we use $[k] \triangleq \{0, 1, \dots, k\}$ to denote the set of integers up to k .

We consider additive-error challenge-response authentication protocols. In such protocols, a verifier \mathcal{V} grants access to a prover \mathcal{P} , if the latter can demonstrate its identity via possession of a shared secret. The protocol has three phases: (i) An initialisation phase. (ii) A *challenge-response* phase, lasting n rounds, performed without error correction under noisy conditions. (iii) A termination phase. During the *challenge-response* phase the verifier \mathcal{V} sends n challenges c_1, \dots, c_n , with $c_k \in \mathcal{X}$, to the prover \mathcal{P} , which responds by transmitting n responses r_1, \dots, r_n , with $r_k \in \mathcal{X}$. We use $c = (c_k)_{k=1}^n$, and $r = (r_k)_{k=1}^n$ to denote the complete challenge and response sequences respectively. The verifier \mathcal{V} uses an error function $\mathcal{E} : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$ to calculate an error $\varepsilon_i = \mathcal{E}(r_i, c_i)$ for the i -th round. The errors when interacting with a valid user may be non-zero due to noise constraints in the channel. However, as the attacker has to resort to guessing the responses, the expected error of the attacker should be higher than that of the user.

In order to trade off false acceptances with false rejections, we use a threshold value τ , such that a prover is accepted if and only if the total error observed is smaller than τ . The verifier \mathcal{V} calculates the total error $\varepsilon \triangleq \sum_{i=1}^n \varepsilon_i$, and rejects the prover \mathcal{P} , if and only if $\varepsilon \geq \tau$.

The relation of ε to the challenge and response strings c and r strongly depends on the protocol. In order to make our analysis generally applicable, we make the following assumption about the protocol.

Assumption 1. *We assume there exists some $p_A \leq \mathbb{E}(\varepsilon_i | A)$, a lower bound on the expected per-round error of the attacker and some $p_U \geq \mathbb{E}(\varepsilon_i | U)$, an upper bound on the error of a legitimate user. In addition, we assume that the protocol is such that all errors are independently (but not necessarily identically) distributed.*

These bounds depend on the noise during the challenge-response phase and on the protocol under consideration. The independence assumption can be satisfied for the interaction with the user when the channel noise is independent. The interaction with the attacker requires constructing an appropriate protocol, so that learning attackers are precluded. We shall return to this issue in section 4, which discusses specific protocols.

3. Expected loss analysis

In order to perform an expected loss analysis, we must specify the potential losses arising from every aspect of the decision process defined by the protocol. For every round of the *challenge-response phase*, we suffer loss ℓ_B , due to the cost in time and energy of transmission. In addition, we suffer a loss of ℓ_A for each false acceptance and a loss ℓ_U for each false rejection.¹ Given that we perform n rounds, we

¹These losses are subjectively set to application-dependent values. Clearly, for cases where falsely authenticating an attacker the impact is severe, ℓ_A must be much greater than ℓ_U . In addition, ℓ_B must be sufficiently low that running the protocol at all is worthwhile.

define the total loss when the prover \mathcal{P} is either the legitimate user U or the attacker A as given by:

$$L = \begin{cases} n\ell_B + \ell_U, & \text{if } \varepsilon \geq \tau \text{ and } \mathcal{P} = U \\ n\ell_B + \ell_A, & \text{if } \varepsilon < \tau \text{ and } \mathcal{P} = A \\ n\ell_B, & \text{otherwise.} \end{cases} \quad (1)$$

Armed with this information, we can now embark upon an expected loss analysis. We wish to devise an algorithm that guarantees an *upper bound* on the expected loss $\mathbb{E}L$ of the authentication system. To start with, we note that the expected loss when the communicating party is an attacker A or the user U , is given respectively by:

$$\mathbb{E}(L | A) = n\ell_B + \mathbb{P}(\varepsilon < \tau | A) \cdot \ell_A + \mathbb{P}(\varepsilon \geq \tau | A) \cdot 0 \quad (2)$$

$$\mathbb{E}(L | U) = n\ell_B + \mathbb{P}(\varepsilon < \tau | U) \cdot 0 + \mathbb{P}(\varepsilon \geq \tau | U) \cdot \ell_U \quad (3)$$

The *expected loss* is in either case bounded by the *worst-case expected loss*:

$$\mathbb{L} \triangleq \max \{ \mathbb{E}(L | A), \mathbb{E}(L | U) \} \geq \mathbb{E}L \quad (4)$$

If we can find an expression that bounds both $\mathbb{E}(L | A)$ and $\mathbb{E}(L | U)$, we automatically obtain a bound on the expected loss, $\mathbb{E}L$. First, we shall see how to select a near-optimal threshold, given a fixed number of rounds of the protocol. Then we shall look at two approximate methods for selecting the number of rounds itself. These will be compared with the optimal (but computationally more expensive) solution.

3.1. Choice of threshold

We want a threshold τ such that no matter whether the prover \mathcal{P} is the attacker A or the legitimate user U the expected loss $\mathbb{E}(L | \mathcal{P})$ is as small as possible. As we *increase* the threshold τ , $\mathbb{E}(L | \mathcal{P} = U)$ *decreases*, while $\mathbb{E}(L | \mathcal{P} = A)$ *increases* and vice-versa. Intuitively, to minimise the worst-case expected loss, we can use τ such that $\mathbb{E}(L | \mathcal{P} = A, \tau) = \mathbb{E}(L | \mathcal{P} = U, \tau)$, which corresponds to an equalising strategy. We shall define a threshold τ minimising an upper bound on the worst-case expected loss in Theorem 1 on page 6. As an intermediate step, we obtain a bound on the worst-case expected loss for *any* given threshold τ . Formally, we can show the following:

Lemma 1. *Let $\varepsilon_i \in [0, 1]$ be the error of the i -th round. If, for all $i > 0$, it holds that $\mathbb{E}(\varepsilon_i | A) \geq p_A$ and $\mathbb{E}(\varepsilon_i | U) \leq p_U$, for some $p_A, p_U \in [0, 1]$ such that $np_U \leq \tau \leq np_A$, then:*

$$\begin{aligned} \mathcal{L}(n; \tau) &\triangleq n\ell_B + \max \left\{ e^{-\frac{2}{n}(np_U - \tau)^2} \ell_U, e^{-\frac{2}{n}(np_A - \tau)^2} \ell_A \right\} \\ &\geq \max \{ \mathbb{E}(L | A), \mathbb{E}(L | U) \} \geq \mathbb{E}L \end{aligned} \quad (5)$$

Proof of Lemma 1. The expected loss when $\mathcal{P} = A$, is simply:

$$\begin{aligned}\mathbb{E}(L | A) &= n\ell_B + \mathbb{P}\left(\sum_i \varepsilon_i < \tau \mid A\right) \ell_A \\ &\leq n\ell_B + \mathbb{P}\left(\sum_i \varepsilon_i - n\mathbb{E}(\varepsilon_i | A) < \tau - np_A \mid A\right) \ell_A \\ &\leq n\ell_B + e^{-\frac{2}{n}(np_A - \tau)^2} \ell_A,\end{aligned}$$

the last two steps used the fact that $\mathbb{E}(\varepsilon_i | A) \geq p_A$ and the *Hoeffding inequality* (21). Specifically, in our case, Lemma 3 (page 19) applies with $X_i = \varepsilon_i$. Then, it is easy to see that $\mu_i = \mathbb{E}(\varepsilon_i | A)$ for all i and $b_i - a_i = 1$, so $\mathbb{P}(\varepsilon < np_A + nt \mid A) \leq e^{-2nt^2}$. By setting $\tau = np_A + nt$, we obtain $t = (\tau - np_A)/n$, which we can plug into the above inequality, thus arriving at the required result.

The other case, $\mathcal{P} = U$, is handled similarly and we conclude that $\mathbb{E}(L | U) \leq n\ell_B + e^{-\frac{2}{n}(np_U - \tau)^2} \ell_U$. \square

The given upper bound is tight, as can be seen by the following lemma, which proves a matching lower bound for the example of Bernoulli-distributed errors.

Lemma 2. *Modifying our assumptions slightly, let $\varepsilon_i \in \{0, 1\}$ be the Bernoulli-distributed error of the i -th round, such that, for all $i > 0$, it holds that $\mathbb{P}(\varepsilon_i | A) = p_A$ and $\mathbb{P}(\varepsilon_i | U) = p_U$, for some $p_A, p_U \in [0, 1]$ such that $np_U \leq \tau \leq np_A$, then:*

$$\begin{aligned}\mathbb{E}L &\geq \min\{\mathbb{E}(L | A), \mathbb{E}(L | U)\} \\ &\geq n\ell_B + \left(\frac{n}{\tau}\right)^\tau \min\{\ell_A p_A^\tau (1 - p_A)^{n-\tau}, \ell_U p_U^\tau (1 - p_U)^{n-\tau}\}\end{aligned}\tag{6}$$

Proof of Lemma 2. The expected loss when $\mathcal{P} = U$, is simply:

$$\begin{aligned}\mathbb{E}(L | U) &= n\ell_B + \mathbb{P}\left(\sum_i \varepsilon_i \geq \tau \mid U\right) \ell_U \\ &\geq n\ell_B + \mathbb{P}\left(\sum_i \varepsilon_i = \tau \mid U\right) \ell_U \\ &= n\ell_B + \binom{n}{\tau} p_U^\tau (1 - p_U)^{n-\tau} \ell_U \\ &\geq n\ell_B + \left(\frac{np_U}{\tau}\right)^\tau (1 - p_U)^{n-\tau} \ell_U\end{aligned}$$

The other case, $\mathcal{P} = A$, is handled similarly and we conclude that $\mathbb{E}(L | A) \geq n\ell_B + \left(\frac{np_A}{\tau}\right)^\tau (1 - p_A)^{n-\tau} \ell_A$. \square

To see that the upper and lower bounds match, consider that $\mathcal{P} = U$ and $\tau = nc$ with $c < 1$. Then Lemmas 1 and 2 give us:

$$\begin{aligned} n\ell_B + e^{-\frac{2}{n}(np_U - \tau)^2} \ell_U &\geq n\ell_B + \left(\frac{np_U}{\tau}\right)^\tau (1 - p_U)^{n-\tau} \ell_U \\ e^{-\frac{2}{n}(np_U - \tau)^2} &\geq \left(\frac{np_U}{\tau}\right)^\tau (1 - p_U)^{n-\tau} \\ e^{-2(p_U - c)^2 n} &\geq e^{[c \ln \frac{p_U}{c} + (1-c) \ln(1-p_U)]n} \end{aligned}$$

Thus, the upper and lower bounds given for the expected loss ($\mathbb{E} L$) are of the same order with regard to n .

Having bounded the loss suffered when choosing a specific threshold, we now choose a threshold $\hat{\tau}_n^*$ that minimises the above bound for fixed n . In fact, we can show that such a threshold results in a particular upper bound on the expected loss.

Theorem 1. Let $\rho \triangleq \ell_A / \ell_U$ and select

$$\tau = \hat{\tau}_n^* \triangleq \frac{n(p_A + p_U)}{2} - \frac{\ln \rho}{4\Delta} \quad (7)$$

If $np_U \leq \tau \leq np_A$, then the expected loss $\mathbb{E} L$ is bounded by:

$$\mathbb{E}(L \mid n, \hat{\tau}_n^*) \leq \mathcal{L}_1(n) \triangleq n\ell_B + e^{-\frac{n}{2}\Delta^2} \cdot \sqrt{\ell_A \ell_U}. \quad (8)$$

with $\Delta \triangleq p_A - p_U$.

Proof of Theorem 1. Substitute (7) in the first exponential of (5) to obtain:

$$e^{-\frac{2}{n}(np_U - \hat{\tau}_n^*)^2} \ell_U = e^{-\frac{n}{2}\Delta^2} e^{-\frac{\ln^2 \rho}{8n\Delta^2}} \sqrt{\ell_A \ell_U}.$$

It is easy to see that the exact same result is obtained by substituting (7) in the second exponential of (5). Thus, both $\mathbb{E}(L \mid A)$ and $\mathbb{E}(L \mid U)$ are bounded by the same quantity and consequently, so is $\max\{\mathbb{E}(L \mid A), \mathbb{E}(L \mid U)\}$. Thus,

$$\mathcal{L}(n, \hat{\tau}_n^*) \leq n\ell_B + e^{-\frac{n}{2}\Delta^2} \sqrt{\ell_A \ell_U},$$

where we simplified the bound by noting that $\frac{\ln^2 \rho}{8n\Delta^2} > 0$. □

The intuition behind the algorithm and the analysis is that it is possible to bound the probability that A makes less errors than expected, or that U makes more than expected. For this reason, the $\hat{\tau}_n^*$ chosen in the theorem must lie between np_U and np_A . We now shall find choices for n which upper bound the overall loss.

3.2. Choice of the number of rounds

Using similar techniques to those employed for obtaining a suitable value for the threshold, we now indicate two choices for the number of rounds n and provide a matching bound on the expected loss. The first, originally mentioned in [14,15] can be used for any value of Δ . In this paper, we also introduce a significantly improved one, which can be used as long as Δ is not too small.

The first choice of n approximately minimises a simple upper bound on the worst-case expected loss. In fact, the following theorem shows that this choice guarantees a specific amount of loss.

Theorem 2. Assume $\ell_A, \ell_U, \ell_B > 0$. If we choose $\tau = \hat{\tau}_n^*$ and

$$n = \hat{n}^* \triangleq \frac{\sqrt{1 + 2CK} - 1}{C}, \quad (9)$$

where $C = \Delta^2$ and $K = \sqrt{\ell_A \ell_U} / \ell_B$, then the expected loss $\mathbb{E}L$ is bounded by:

$$\mathbb{E}(L \mid \hat{\tau}_n^*, \hat{n}^*) \leq \mathcal{L}_2 \triangleq \sqrt{8K/C} \cdot \ell_B = \frac{\sqrt{8\ell_B} (\ell_A \ell_U)^{1/4}}{\Delta}. \quad (10)$$

Proof of Theorem 2. We shall bound each one of the summands of (8) by $\sqrt{2K/C} \cdot \ell_B$. For the first term we have:

$$\begin{aligned} n\ell_B &= \frac{\sqrt{1 + 2CK} - 1}{C} \ell_B \leq \frac{\sqrt{1 + 2CK}}{C} \ell_B \\ &\leq \frac{\sqrt{2CK}}{C} \ell_B = \sqrt{\frac{2K}{C}} \ell_B. \end{aligned}$$

For the second term, by noting that $e^x \geq 1 + x$, we have:

$$\begin{aligned} \sqrt{\ell_U \ell_A} \cdot e^{-\frac{n}{2}\Delta^2} &\leq \frac{\sqrt{\ell_U \ell_A}}{1 + \frac{n}{2}\Delta^2} = \frac{K\ell_B}{1 + \frac{nC}{2}} = \frac{2K\ell_B}{1 + \sqrt{1 + 2CK}} \\ &\leq \frac{2K\ell_B}{\sqrt{1 + 2CK}} \leq \frac{2K\ell_B}{\sqrt{2CK}} = \sqrt{\frac{2K}{C}} \ell_B. \end{aligned}$$

Summing the two bounds, we obtain the required result. \square

In particular, this theorem proves that our *worst-case expected loss* \mathbb{L} grows sublinearly both with increasing round cost (with rate $O(\epsilon^{1/2})$) and with increasing authentication costs (with rate $O(\epsilon^{1/4})$). Furthermore, the *expected loss* is bounded symmetrically for both user and attacker access. Finally, there is a strong dependence on the margin Δ between the attacker and the user error rates, which is an expected result.

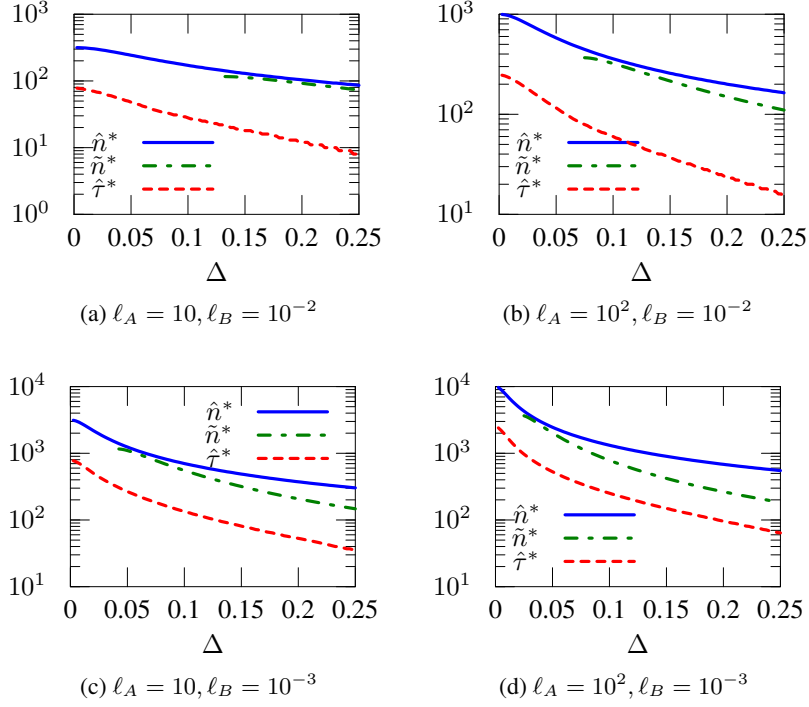


Fig. 1. Effect of gap Δ on the threshold and number of rounds, for different choices of l_A, l_B , with $l_U = 1$.

3.2.1. Improved choice for large Δ

As the above value is somewhat too conservative, especially for large values of Δ , we can obtain instead a value for the number of rounds by directly minimising (8). Observe that by differentiating the right hand side, we have:

$$\begin{aligned} \ell_B &= \frac{\Delta^2 \sqrt{\ell_A \ell_U}}{2} e^{-n \frac{\Delta^2}{2}} \\ n &= \frac{2}{\Delta^2} \ln \frac{\Delta^2 \sqrt{\ell_A \ell_U}}{2 \ell_B}, \quad \Delta > \sqrt{2 \ell_B} (\ell_A \ell_U)^{-1/4}. \end{aligned}$$

This is a good choice for the number of rounds as long as Δ is not too small. We can now prove a bound for this choice.

Theorem 3. Assume $\ell_A, \ell_U, \ell_B > 0$ and $\Delta > \sqrt{2 \ell_B} (\ell_A \ell_U)^{-1/4}$. If we choose $\tau = \hat{\tau}_n^*$ and

$$n = \tilde{n}^* \triangleq \left\lceil \frac{2}{\Delta^2} \ln \frac{\Delta^2 \sqrt{\ell_A \ell_U}}{2 \ell_B} \right\rceil, \quad (11)$$

then the expected loss is bounded by

$$\mathbb{E}(L \mid \tilde{n}^*, \hat{\tau}_n^*) \leq \mathcal{L}_3 \triangleq \ell_B \left[1 + \frac{2}{\Delta^2} \left(1 + \ln \frac{\Delta^2 \sqrt{\ell_A \ell_U}}{2 \ell_B} \right) \right]. \quad (12)$$

Proof of Theorem 3. Plugging $n = \tilde{n}^*$ into (8), we obtain

$$\mathbb{E}(L \mid \tilde{n}^*, \hat{\tau}_n^*) \leq \tilde{n}^* \ell_B + e^{-\frac{\tilde{n}^*}{2} \Delta^2} \cdot \sqrt{\ell_A \ell_U}.$$

Taking the first term, we have:

$$\tilde{n}^* \ell_B \leq \ell_B + \frac{2\ell_B}{\Delta^2} \ln \frac{\Delta^2 \sqrt{\ell_A \ell_U}}{2\ell_B}$$

where we used the fact that $\lceil x \rceil \leq x + 1$.

Taking now the second term, and using fact that $\lceil x \rceil \geq x$, we have:

$$e^{-\frac{\Delta^2}{2} \tilde{n}^*} \cdot \sqrt{\ell_A \ell_U} \leq e^{-\ln \frac{\Delta^2 \sqrt{\ell_A \ell_U}}{2\ell_B}} \cdot \sqrt{\ell_A \ell_U} = \frac{2\ell_B}{\Delta^2}.$$

Adding together the two terms, we obtain the required result. \square

In order to measure the quality of those bounds, we shall now compare them, by applying them to a family of RFID protocols.

3.2.2. The effect of problem parameters on the security variables

Since different choices for the three losses, as well as different amounts of distinguishability Δ have a direct effect on the number of rounds and threshold, it is instructive to show how these change. In particular, Figure 1 shows how the threshold, and two choices for the number of rounds, change as a function of Δ , for four different combinations of ℓ_A, ℓ_B , when $\ell_U = 1$.

In all settings, both the number of rounds and the threshold decrease as Δ increases. However, we also observe that they all increase when ℓ_B decreases (compare Fig 1a and 1c), or when ℓ_A increases (compare Fig. 1a and Fig. 1b). Indeed, the number of rounds is highest in Fig. 1d, which intuitively makes sense, since the threshold is generally proportional to the number of rounds, while the latter increases when the ratio ℓ_A/ℓ_B increases. Finally, note that although the improved number of rounds is initially quite close to the original choice, there is a significant difference for larger Δ .

In the following sections, rather than working with abstract values of Δ and with the direct expected loss calculations and bounds, we shall consider a specific family of protocols which implements such a challenge-response phase. Under assumptions on the noise and the protocol implementation, we shall show that our analysis is applicable, and we shall provide simulated experiments which compare our choices of threshold with the optimal choice, as well as a Bayesian variant.

4. Analysis of RFID thresholded protocols

Currently, the most well-known protocols employing an authentication phase without any error correction and under resource constraints are *RFID distance bounding* protocols. These are used for authentication between a reader (the verifier) and an RFID tag (the prover). These challenge-response protocols use a rapid-bit exchange phase in order to ensure that the tag is at an appropriate distance from the

reader. This way, then can avoid man-in-the-middle relay attacks, due to the time delay that these attacks introduce.

For illustrative purposes, we shall examine the properties of two such protocols, for which it is possible to derive simple expressions for p_A, p_U , given symmetric channel noise.² In general, our derivations will be biased towards simplicity of expression rather than accuracy.

We consider an exchange between the verifier \mathcal{V} and the prover \mathcal{P} , which in case of an RFID protocol are embodied as an RFID reader and an RFID tag respectively. As tags are generally underpowered (or completely unpowered) communication normally has limited range and can suffer from significant noise. To model noise in the physical medium, we assume that in any exchange between \mathcal{V} and \mathcal{P} , the former may send a symbol $x \in \mathcal{X}$, while the latter may receive a symbol $\hat{x} \in \mathcal{X}$ such that $x \neq \hat{x}$. We shall denote the probability of erroneous transmission in the data layer as: $\omega \triangleq \mathbb{P}(\hat{x} = j \mid x = i)$ for $i \neq j$. For simplicity, we shall only treat the case of symmetric channel noise such that: $\mathbb{P}(\hat{x} = y \mid x \neq y) = \frac{1}{|\mathcal{X}|-1}$, $\forall y \neq x, x, y \in \mathcal{X}$. However, all our results carry over for the case of asymmetric channel noise ω_1, ω_2 for each direction, by taking $\omega = \max_j \omega_j$. The noise is particularly important for distance-bounding protocols, as they include a challenge-response phase where single-bit messages are exchanged rapidly between the reader and tag, so as to verify that the tag is in close proximity, by taking advantage of the bounded speed of light. This phase, called the *rapid-bit exchange* phase, closely corresponds to our analysis in the previous section.

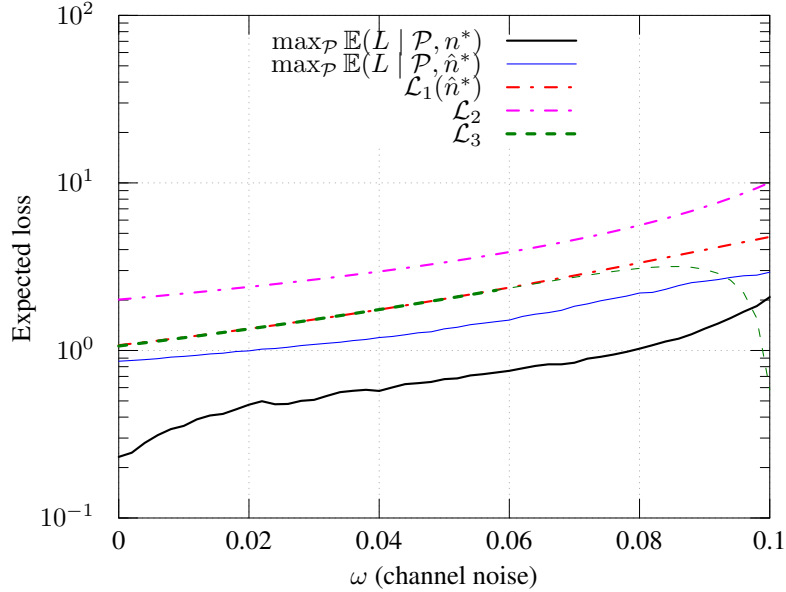


Fig. 2. The worst-case expected Loss \mathbb{L} and the bounds \mathcal{L}_1 and \mathcal{L}_2 from theorems 1 and 2 respectively vs. the channel error rate ω . The lowest line depicts the worst case expected loss when the number of rounds is chosen optimally, while the second line depicts the same thing according to our first choice of number of rounds. The alternative line \mathcal{L}_3 is only valid for large Δ and hence small amounts of noise. Thus, the region of small Δ is depicted with a thinner line.

²The asymmetric case results in more involved expressions[16] and additional parameters, but is not particularly interesting to study.

In particular, two RFID distance bounding protocols, the SWISS-KNIFE protocol [11] and the variant HITOMI [7] are additive-error authentication protocols satisfying our assumptions. Both protocols include an initialisation and a termination phase with error correction, as well as a distance-bounding (rapid-bit exchange) phase without error correction. In the initialisation phase, the session keys are agreed. The distance bounding phase is an n -round phase where a sequence of challenges $(c_t)_{t=1}^n$ and responses $(r_t)_{t=1}^n$ are exchanged without error correction. This is also called the *rapid-bit exchange* phase. However, the received challenges and responses \hat{c}_t, \hat{r}_t may differ from the ones sent due to noise. Let r_t^* denote the correct response to the i -th challenge. Then the error at time t is simply equal to 1 whenever the received response does not match the correct response, i.e. $\varepsilon_t \triangleq \mathbb{I}\{\hat{r}_t \neq r_t^*\}$. During the termination phase, the challenges and responses are sent over the error corrected channel to the verifier, who checks if they match those used in the rapid-bit exchange.

In order to satisfy Assumption 1 for both the user and attacker, we need the actual channel noise to be independent, but not necessarily identically distributed for every challenge and response. We also require the protocol to be constructed in such a way that it holds:

$$\mathbb{P}(c_t = i \mid c_1, \dots, c_{t-1}) = \frac{1}{2}, \quad \mathbb{P}(r_t = j \mid r_1, \dots, r_{t-1}, c_1, \dots, c_t) = \frac{1}{2}, \quad (13)$$

in order to ensure that the attacker has no advantage when trying to guess future challenges and responses from previous ones. The first part can be achieved by ensuring that the challenges are uniformly distributed in $\{0, 1\}$. The second part is slightly more involved, as the correct challenge depends both on a shared secret and on a cryptographic hash function using a random nonce. The hash function must be such that a uniform distribution of nonces results in a uniform bit distribution for each hash. In practice, this assumption must be relaxed to that of *computational* indistinguishability of the actual distributions from the uniform one. Finally, these probabilities must be independent of previous runs of the protocol, in order to prevent *learning* attackers (i.e. attackers that build knowledge by observing multiple runs of the protocol).

4.1. Noise-dependent bounds on the error probabilities

Similarly to [7] we can show that, for those two protocols, under symmetric channel noise ω , we can obtain the following simple bounds on the expected error: $p_A = \frac{1}{2}$ and $p_U = 2\omega$, where we note in passing that since we need $p_A \geq p_U$, so $\omega \leq \frac{1}{4}$. The user probability of error p_U follows from a simple union bound:

$$\mathbb{P}(\varepsilon_t) \leq \mathbb{P}(\hat{r}_t \neq r_t \vee \hat{c}_t \neq c_t) \leq \mathbb{P}(\hat{r}_t \neq r_t) + \mathbb{P}(\hat{c}_t \neq c_t) = 2\omega.$$

The analysis of the attacker error rate is quite straightforward.

Threat model: A naive man-in-the-middle attack would simply relay the messages between the reader and tag in the initialisation and termination phases, but would not be able to do so for the rapid-bit exchange phase, because the time delay would be too great. Instead, he would simply try and guess the correct responses, thus achieving an error rate of $1/2$. In a more sophisticated man-in-the-middle attack, the adversary may send a random string of challenges to the tag *before* the rapid-bit exchange phase begins, for which he collects the correct responses. Due to the assumptions explained in the previous section, the attacker can only guess the challenges with probability $1/2$ in each round, even if he tries to

delay guessing until he has seen some of the challenges. During the actual communication with the reader (verifier), he can replay the collected responses, if and only if the challenges happen to coincide. This would reduce his error rate to $1/4$. However, the analysed protocols also include a final *verification phase*, where the challenges used by the attacker are checked over an error-correcting channel. Consequently, the attacker's error rate is at least $1/2$.

4.2. Validation experiments

In this section, we experimentally validate our theoretical analysis for the above protocols. In the results shown here, we chose the following values for the losses: $\ell_A = 10$, $\ell_U = 1$, $\ell_B = 10^{-2}$ for illustrative purposes.

More specifically, Figure 2 depicts the *worst-case expected loss* for different choices of the number of rounds, and our near-optimal threshold, as well as the values of theoretical bounds. For the optimal number of rounds n^* we obtain the loss $\mathbb{E}(L \mid n^*)$. This is of course smaller than $\mathbb{E}(L \mid \hat{n}^*)$, the loss suffered by choosing \hat{n}^* , with the gap becoming smaller for larger error rates. Since when this occurs, the *expected loss* is very close to \mathcal{L}_1 , this implies that the bound of Theorem 1 can be further tightened for small ω . Indeed, we see that the improved bound \mathcal{L}_3 obtained for the improved number of rounds \tilde{n}^* almost matches the bound $\mathcal{L}_1(\hat{n}^*)$, but it fails to hold for large ω , i.e. for small Δ , as expected.

5. Estimating ω

One potentially major problem with the above approach is that knowledge of ω is necessary, at least in order to calculate a bound on the error probabilities for the user. In this section, we discuss how this can be done leveraging the coding performed during the initial and final phases of the protocol.

We assume some coding function $\Phi : \mathcal{X}^m \rightarrow \mathcal{X}^k$, with $k > m$, and a metric γ on \mathcal{X}^k (where usually $\mathcal{X} = \{0, 1\}$ and γ is the Hamming distance) such that:

$$\gamma_{\min} \triangleq \min \{ \gamma(\Phi(x), \Phi(y)) : x, y \in \mathcal{X}^m, x \neq y \} \quad (14)$$

is the minimum (Hamming) distance between valid codewords. For a given $x \in \mathcal{X}^m$, the source transmits $\phi = \Phi(x)$ and the sink receives $\hat{\phi}$, with $\phi, \hat{\phi} \in \mathcal{X}^k$. As before, we assume that the physical channel has a symmetric error rate $\omega = \mathbb{P}(\hat{\phi}_i \neq \phi_i)$, where ϕ_i denotes the i -th bit of ϕ . This is then decoded as $\hat{x} \triangleq \operatorname{argmin} \{ \gamma(\hat{\phi}, \Phi(y)) : y \in \mathcal{X}^m \}$. Let θ be the number of errors in the string $\hat{\phi}$, or more precisely $\theta = \gamma(\phi, \hat{\phi})$. Let $\hat{\theta} \triangleq \gamma(\Phi(\hat{x}), \hat{\phi})$ be the distance between the closest valid codeword $\Phi(\hat{x})$ and the received $\hat{\phi}$. If $\theta < (\gamma_{\min} - 1)/2$, then $\theta = \hat{\theta}$.

The crux of our method for estimating ω relies on the number of errors θ being less than $(\gamma_{\min} - 1)/2$, in which case, the estimated number of errors $\hat{\theta}$ will equal θ . Let $\hat{\omega} \triangleq \frac{\hat{\theta}}{n}$ be our empirical error rate. In that case, the expected empirical error rate equals the true error rate. More formally:

$$\mathbb{E}(\hat{\omega} \mid \theta \leq (\gamma_{\min} - 1)/2) = \omega. \quad (15)$$

If $\theta > (\gamma_{\min} - 1)/2$ then the protocol fails in any case, due to decoding errors in the initial or final phases. If not, then the above equation holds and we can obtain high probability bounds for ω via the

Hoeffding inequality (Appendix, Lemma 3). In particular, it is easy to show that, for any $\delta \in [0, 1]$:

$$\mathbb{P} \left(|\hat{\omega} - \omega| \geq \sqrt{\frac{\ln 2/\delta}{2k}} \right) \leq \delta, \quad (16)$$

by substituting the square-root term into (21), and setting $\mu_i = \omega$, $\sum X_i = \hat{\theta}$, $a_i = 0$, $b_i = 1$. Consequently, for the SWISS-KNIFE family of protocols the following values for p_A and p_U hold with probability $1 - \delta$:

$$p_A = \frac{1}{2}, \quad p_U = 2\hat{\omega} - \sqrt{\frac{2 \ln 2/\delta}{k}}. \quad (17)$$

Experimental investigations presented in the next section indicate that this choice has good performance in terms of expected loss.

Remark 1. *It is possible for an attacker to manipulate the noise estimation phase. However, doing so would only increase the security of the protocol.*

This follows directly from Theorems 1 and 2 and the fact that an increased estimated noise results in a smaller value for Δ . More precisely, the number of rounds is decreasing with Δ , and so a decrease in Δ leads to an increase in the number of rounds. The threshold itself is proportional to the number of rounds, with an offset $-\frac{\ln \rho}{4\Delta}$. Consequently, when $\rho > 1$, meaning that $\ell_A > \ell_U$, any increase in Δ would lead to a decrease in the threshold, i.e. a tolerance of a proportionally smaller number of errors.

6. Evaluation Experiments

We have performed some experiments to evaluate our methods in a more realistic setting, involving an RFID distance bounding protocol with three phases. An initialisation phase, a *rapid-bit exchange* phase and a termination phase. The rapid-bit exchange phase consists of a sequence of single-bit challenges and responses, where no error correction is performed. Consequently there is a significant probability of errors for legitimate users due to channel noise. As previously mentioned, the attacker can achieve an error probability of $\frac{1}{2}$ per bit. In our simulations, we model all channels as symmetric channels with uniform error rate ω . All the phases are executed through those channels, and the error rate is estimated through the errors made in the initial phase.³ We perform two sets of experiments. In the first set, we estimate the effect of different methods of estimating the noise on the expected loss of the protocol, for two different methods of choosing the threshold. In the second set of experiments, we compare those methods against the minimax solution for a variety of experimental conditions, keeping the noise estimation fixed.

In all cases, performance is in terms of the worst-case expected loss. We estimate this by measuring the loss incurred when a *legitimate user* U is trying to get authenticated and when an *adversary* A is trying to perform a mafia fraud attack[17]. We have estimated the *worst-case expected loss* by running 10^4 experiments for each case, obtaining a pair of estimates $\hat{\mathbb{E}}(L | A)$, $\hat{\mathbb{E}}(L | U)$ by averaging the loss L , as

³If there are more errors than can be corrected by the error correcting code, then the protocol usually fails because the session keys do not match.

defined in (1), incurred in each experiment and taking the maximum of the two. In all of the experiments shown in this section, we chose the $\ell_U = 1$ for the user loss and we used $k = 2^{10}$ for the coded messages in the initialisation phase.

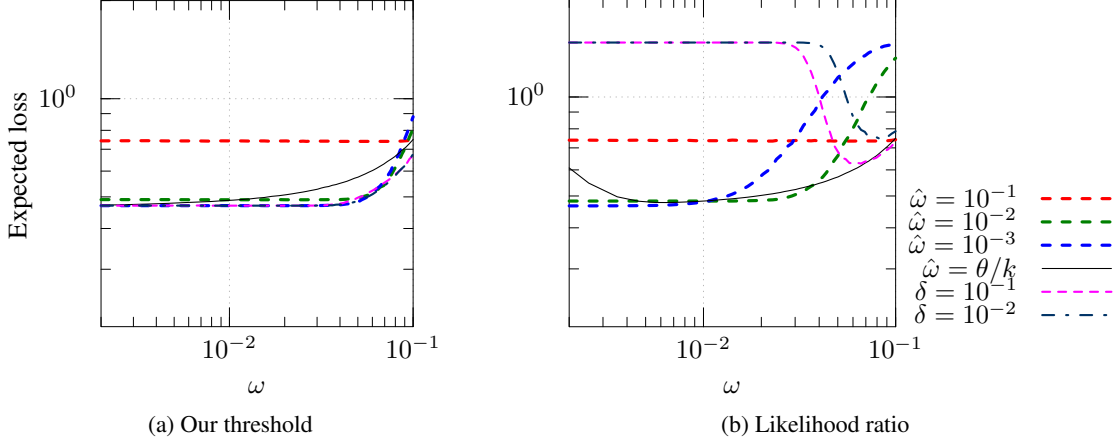


Fig. 3. The worst-case expected loss as a function of noise. We plot the evolution of the loss as noise changes, for a number of different cases. Firstly, for the case where we arbitrarily assume a noise value $\hat{\omega} \in \{10^{-1}, 10^{-2}, 10^{-3}\}$. Secondly, for an empirically estimated $\hat{\omega} = \hat{k}/n$, and finally for p_A, p_U calculated via equation (17) with $\delta \in \{10^{-1}, 10^{-2}\}$.

6.1. Evaluation of noise estimation

The actual values p_A, p_U depend on ω , which is unknown. We compare three methods for choosing p_A, p_U . Firstly, guessing a value $\hat{\omega}$ for the channel noise. Secondly, using the maximum likelihood noise estimate $\hat{\omega} = \theta/k$. In both cases, we simply use $\hat{\omega}$ as described at the beginning of Sec. 4 to obtain p_A, p_U . In the third case, we use the high-probability bounds (17) for p_A, p_U , with an arbitrary value of δ . For those experiments, we use $\ell_A = 10$, $\ell_U = 1$, $\ell_B = 10^{-2}$, while we used $k = 2^{10}$ for the coded messages in the

In the first experiment, shown in figure 3a, we use the nearly-optimal threshold and number of rounds that we have derived in our analysis. In the second experiment, shown in figure 3b, we replace our choice of threshold with a choice similar to that of Baignères et al. [18]. Their threshold is derived via a likelihood ratio test, which is asymptotically optimal (c.f. [19,20]) but they do not consider specific losses. Since in our case we have unequal losses ℓ_A and ℓ_U , we re-derive their threshold via a Bayesian test (to which a Bayesian formulation of the Neymann-Pearson lemma [19] applies) to obtain:

$$\tilde{\tau} = \frac{n \ln \frac{1-p_U}{1-p_A} - \ln \rho}{\ln \frac{1-p_U}{1-p_A} - \ln \frac{p_U}{p_A}}. \quad (18)$$

For equal losses, $\rho = 1$, and $\tilde{\tau}$ equals the threshold used in [18]. Such tests have good asymptotic properties [19,21]. Interestingly, for small Δ , the form of $\tilde{\tau}$ is similar to $\hat{\tau}_n^*$: Let \bar{p} such that $p_A = \bar{p} + \Delta/2$ and $p_U = \bar{p} - \Delta/2$. Then (18) can be approximated by $\tilde{\tau}^* = n\bar{p} - \frac{\bar{p}(1-\bar{p})}{\Delta} \ln \rho$. More details on the derivation of (18) are given in Appendix B on page 19.

Figure 3 depicts the *worst-case expected loss* \mathbb{L} as a function of the actual noise ω . Figure 3a shows \mathbb{L} using the threshold τ derived from our *expected loss* analysis (7), while in Figure 3b we use the asymptotically optimal threshold of (18). In both cases, we plot \mathbb{L} , while the actual noise ω is changing, for a number of different cases. Initially, we investigate the evolution of \mathbb{L} for three arbitrarily chosen values $\hat{\omega} \in \{10^{-1}, 10^{-2}, 10^{-3}\}$. Additionally, we examine the evolution of the *worst-case expected loss*, when the noise is empirically estimated $\hat{\omega} = \hat{\theta}/n$ and finally when p_A and p_U are calculated via equation (17) with $\delta \in \{10^{-1}, 10^{-2}\}$.

As it can be seen in Figure 3, in all cases (using ours Figure 3a or Baignères et al. [18] threshold Figure 3b) the *worst-case expected loss* is very low for small values of the actual noise and increases sharply when the actual noise approaches the value of 10^{-1} . It is interesting to see that when we use the optimistic⁴ high probability estimates for p_A, p_U , we obtain almost always better performance than simply guessing the noise, or using the plain empirical estimate $\hat{\omega}$ directly. Furthermore, using the asymptotically optimal threshold (18), we observe a deterioration in the results. Thus, our approach results in a clearly dominating performance over other methods.

As mentioned in Sec. 7, the choice of the threshold by Baignères et al. [18] is only asymptotically optimal. Ours, while not optimal, gives a worst-case expected loss guarantee for any finite sample size. Thus, it has better performance when the asymptotic approximation is not sufficiently good, which occurs when both the number of rounds n and the gap Δ are small. In addition, it is much less sensitive to the estimate of ω than the likelihood ratio threshold.

6.2. Comparison with the minimax solution

When computational considerations are not an issue, i.e. when we can perform an exhaustive search, it is possible to optimise directly for the minimax solution. That is, given particular parameters $\chi \triangleq \{p_A, p_U, \ell_A, \ell_U, \ell_B\}$ we can find a threshold τ and number of rounds n minimising the worst-case expected loss, that is:

$$\min \left\{ \max_{\mathcal{P} \in \{A, U\}} \mathbb{E}(L \mid \mathcal{P}, \tau, n, \chi) : n \in [n_{\max}], t \in [n] \right\} \quad (19)$$

Note here that for any decision rule such that $n > n_{\max} \triangleq \min\{\ell_A, \ell_U\}/\ell_B$, then our cost is trivially larger than either rejecting or accepting immediately. Consequently, we only need to consider a finite set of number of rounds to recover the minimax solution. In particular, the complexity of the solution includes $O(n_{\max}^2)$ evaluations of the maximum expected loss, with each evaluation requiring the computation of the binomial cumulative probability function. This in turn has quadratic complexity, giving an overall complexity of $O(n_{\max}^4)$.

We performed a thorough experimental comparison between the minimax solution, our suggested threshold and number of rounds, and the threshold chosen via the likelihood ratio test. In all cases, the noise was estimated with $k = 2^{10}$ and using the high-probability bound for p_A, p_U . The user loss was fixed to $\ell_A = 1$, while we varied the bit cost ℓ_B and the attacker loss ℓ_A .

The results for four different choices of those parameters, as ω varies, are shown in Figure 4. It can be seen that our choice always dominates the likelihood ratio, while of course the minimax solution is

⁴Experiments with pessimistic high probability estimates for the noise showed a significant increase in the number of rounds used, which resulted in a higher expected loss.

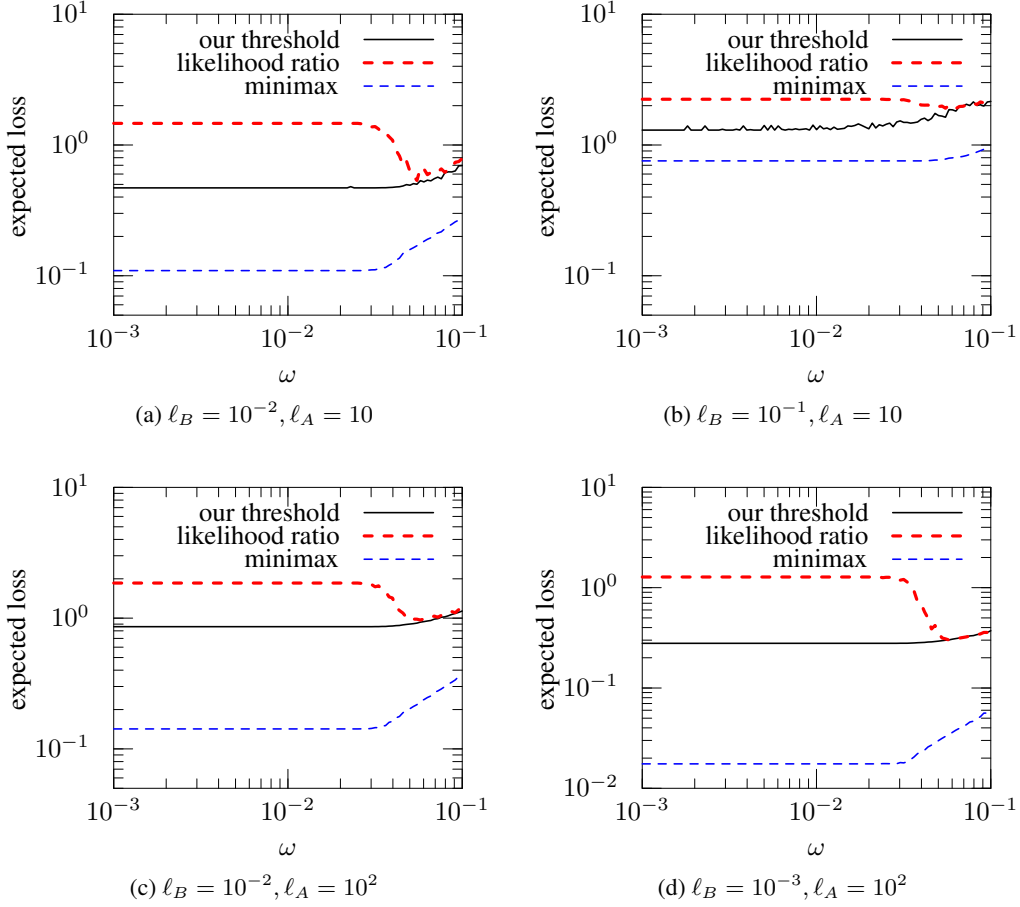


Fig. 4. Comparison of the worst-case expected loss, for: (i) the approximate number of rounds and threshold, (ii) the threshold given by the likelihood ratio test, as well (iii) the exact minimax solution for both the threshold and the number of rounds. We show results for four different cases, where we vary the bit cost ℓ_B and the attack cost ℓ_A .

clearly the best. It can also be seen that whenever the ratio between the attack cost ℓ_A and the bit cost ℓ_B is small, the three methods are closer together.

7. Conclusion

We have performed the first, to our knowledge, expected loss analysis of additive error challenge-response authentication protocols under channel constraints. Such an analysis is necessary, because of the inherent cost of increasing the number of rounds n and the need to trade off optimally the costs of incorrectly authenticating an illegitimate entity, or failing to authenticate a legitimate one, as well as the lack of an error-free channel. In order to achieve this, we made the assumption of bounds on the expected errors of a user and an attacker. In addition, we assumed that these are independent from errors in the

previous rounds.⁵ The present paper is an extension of the results originally presented in [15,14], by the inclusion of improved bounds for the choice of number of rounds, as well as new empirical results.

The *rapid-bit exchange phase* considered in parts of this paper was introduced in [3] to compute an upper bound on the distance of the prover \mathcal{P} from the verifier \mathcal{V} . This is composed of n challenge-response *rounds*, used to calculate a round-trip time and thus place a bound on the distance. Subsequently, a broad range of *distance bounding* protocols were proposed, both for RFID [9,11,10,5,7], as well as other wireless devices [23,24,6].

Hancke and Kuhn [9] were the first to indicate that since the *rapid-bit exchange* phase is taking place in a noisy channel, challenges and responses may be corrupted. Thus, a legitimate user may fail to get authenticated. Their protocol (henceforth HAKU), employed n rounds and authenticated any prover who made a number of mistakes ε less than an acceptance threshold τ , so as to reduce the number of false rejections. Using the binomial distribution and an assumption on the error rates they give expressions for the *false accept* and *false reject probability* as a function of n and τ , but they provide no further analysis. Nevertheless, they indicate that the number of challenge-response rounds n in the rapid bit exchange phase should be chosen according to the expected error rate. Kim et al. [11] extend this approach with the SWISS-KNIFE protocol by considering *three* types of errors. Finally [5], rather than using a threshold τ , proposed a protocol (henceforth ECMAD) using an error correcting code (ECC). ECMAD, which extends the MAD protocol [23], uses only k of the n total rounds for the challenges and responses. The remaining $n - k$ rounds are used to transmit the (n, k) ECC. This has the effect of achieving better security (in terms of false acceptance rates) with the same number of rounds n .

All these approaches use n rounds in the noisy authentication phase. However, they do not define the optimal n . They simply state that the probability of authenticating a user becomes much higher than the probability of authenticating an attacker as n increases. However, a large value of n is incompatible with the requirements of many applications and devices (i.e. high value of n leads to high overhead for resource-constrained devices). This can be modelled by assigning an *explicit cost* to every round, which should take into account the transmission energy, computation and time overhead. This cost has so far not been explicitly taken into consideration.

Another work that is closely related to ours is [25], which, given a *required* false acceptance and false rejection rate, provides a *lower bound* on the number of required rounds. This analysis is performed for both HAKU and ECMAD. However, it assumes that the number of rounds n would be large enough for the binomial distribution of errors to be approximately normal. Our analysis is more general, since it uses finite-sample bounds that hold for any bounded error function.

Recently, Baignères et al. [18] have given an analysis on the related topic of distinguishing between a real and a fake solver of challenge-response puzzles. More precisely, they study CAPTCHA-like protocols and provide a threshold which minimizes the probability of error in these protocols. The main differences between the analysis presented in this paper and [18] can be summarised below: (a) We perform an expected loss analysis rather than an error analysis. (b) Our bounds hold uniformly, while [18] uses an asymptotically optimal distinguisher. (c) We consider bounded errors rather than $\{0, 1\}$ errors for each challenge-response. (d) We additionally propose a method to estimate channel noise. This is of course not applicable in the context of [18], due to the different setting.

A more general work on authentication under noisy conditions was presented in [26]. This provided tight information-theoretic upper and lower bounds on the attacker's success in impersonation and substitution attacks, proving that it decreased with noise. However, this is somewhat misleading, as it ignores

⁵Consequently, error probabilities drop exponentially fast, unlike the examples given in [22].

the efficiency of the protocol. Our analysis shows that, when one considers losses due to communication overhead and false rejections of users, the expected loss increases, which is a natural result.

In this paper, we perform a detailed *expected loss* analysis for a general class of additive-error authentication protocols in a noisy channel. The analysis is performed by assigning a loss ℓ_B to each round, and losses ℓ_A, ℓ_U to false acceptance and false rejection respectively.

We show how a nearly-optimal threshold $\hat{\tau}_n^*$ for a given number of rounds n can be chosen and give *worst-case* bounds on the *expected loss* for that choice. Thus, the bounds hold no matter if the party that attempts to get authenticated is either a legitimate user U or an attacker A . This extends our previous work [7], which proposed a new *distance bounding* protocol (HITOMI) and only calculated a value for the threshold τ , without providing any bounds.

We also show how a nearly-optimal number of rounds \hat{n}^* can be chosen and give further bounds on the *expected loss*. The bounds hold for *any bounded* error function, and not only for $\{0, 1\}$ errors.⁶ Furthermore, they are valid for any n , since they are based on probability inequalities for a finite number of samples.⁷ Thus, they are considerably more general to the bounds of [25].

Finally, we provided high-probability estimates for the current noise level in the channel by leveraging the coding performed in the initial and final phases of the protocol, which takes place in a coded channel. This enables us to significantly weaken assumptions on knowledge of the noise level in the channel and in turn, provide an authentication algorithm which has low expected loss with high probability. Experimentally, we obtain uniformly superior results to guessing or direct empirical noise estimates. Finally, we repeated those experiments with an asymptotically optimal threshold similar to that used by Baignères et al.[18]. Our results indicate a significant improvement through the use of a threshold with uniform, rather than asymptotic, guarantees. Consequently, it is our view that algorithms motivated by an asymptotic analysis should be avoided in the finite-sample regime of most challenge-response authentication protocols.

Our analysis is particularly significant for areas of communications where challenges and responses are costly and where there exists significant uncertainty about the correctness of any single response.

Acknowledgements

We would like to extend thanks to the anonymous reviewers, whose insightful comments led to significant improvements in the manuscript. This work was partially supported by the Marie Curie IEF project “PPIDR: Privacy-Preserving Intrusion Detection and Response in Wireless Communications”, grant number: 252323, the Marie Curie IEF Project “ESDEMUU: Efficient sequential decision making under uncertainty”, Grant Number 237816 and the FP7 project “IM-CLeVeR - Intrinsically Motivated Cumulative Learning Versatile Robots”, grant agreement No FP7-ICT-IP-231722.

⁶In all previous proposals, there is either an error at each round, or there is not.

⁷The analysis in [25] only holds for large n , so the approximation only holds asymptotically.

Appendix

A. Useful formulas

If X_1, \dots, X_n are independent Bernoulli random variables with $X_k \in \{0, 1\}$ and $\mathbb{P}(X_k = 1) = \mu$ for all k , then

$$\mathbb{P}\left(\sum_{k=1}^n X_k \geq u\right) = \sum_{k=0}^u \binom{n}{k} \mu^k (1 - \mu)^{n-k}. \quad (20)$$

Lemma 3 (Hoeffding). *For independent random variables X_1, \dots, X_n such that $X_i \in [a_i, b_i]$, with $\mu_i \triangleq \mathbb{E} X_i$ and $t > 0$:*

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^n X_i \geq \sum_{i=1}^n \mu_i + nt\right) &\leq \exp\left(-\frac{2n^2 t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right) \\ \mathbb{P}\left(\sum_{i=1}^n X_i \leq \sum_{i=1}^n \mu_i - nt\right) &\leq \exp\left(-\frac{2n^2 t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right). \end{aligned} \quad (21)$$

B. On Bayesian hypothesis tests

One way to obtain an asymptotically minimax optimal threshold is to employ a Bayesian hypothesis test [19]. This requires defining a prior probability on the possible hypotheses. In our case, the hypothesis set is $H = \{A, U\}$, on which we define a prior probability π . For $\{0, 1\}$ errors, the probability of observing ε errors out of n observations is given by $\mathbb{P}(\varepsilon | A)$ and $\mathbb{P}(\varepsilon | U)$ for the attacker and user respectively. For concreteness, we shall assume that this follows a binomial distribution with parameters p_A, p_U in each case. Given an observed error x , the *posterior* probability of any hypothesis $h \in H$ is:

$$\pi(h | \varepsilon = x) = \frac{\mathbb{P}(\varepsilon = x | h)\pi(h)}{\sum_{h' \in H} \mathbb{P}(\varepsilon = x | h')\pi(h')}.$$

We then define a decision set $G = \{g_A, g_U\}$, where g_A means we decide that the prover is an attacker and g_U means we decide that the prover is a user. Finally, we define a loss function $L : G \times H \rightarrow \mathbb{R}$, such that $L(g, h)$ is our loss when we decide g and h is the correct hypothesis. The expected loss of decision $g \in G$, under our prior and given ε errors out of n is:

$$\mathbb{E}_\pi(L | \varepsilon, g) = \sum_{h \in H} L(g, h)\pi(h | \varepsilon),$$

where \mathbb{E}_π denotes expectation with respect to the prior π . Now define the decision function $q : \{0, 1, \dots, n\} \rightarrow G$:

$$q(\varepsilon) \triangleq \begin{cases} g_U, & \text{if } \mathbb{E}_\pi(L | \varepsilon, g_U) \leq \mathbb{E}_\pi(L | \varepsilon, g_A) \\ g_A, & \text{if } \mathbb{E}_\pi(L | \varepsilon, g_U) > \mathbb{E}_\pi(L | \varepsilon, g_A). \end{cases} \quad (22)$$

This decision function minimises $\mathbb{E}_\pi L$ by construction (c.f. [19] ch. 8). The following remark is applicable in our case:

Remark 2. Assume i.i.d errors with $\varepsilon_i \in \{0, 1\}$, so that we can use a binomial probability for $\mathbb{P}(\varepsilon | h)$. Set the loss function L to be $L(g_U, A) = \ell_A$, $L(g_A, U) = \ell_U$ and 0 otherwise. Then the decision function (22) becomes equivalent to:

$$q(\varepsilon) \triangleq \begin{cases} g_U, & \text{if } \varepsilon < \tau_b \\ g_A, & \text{if } \varepsilon \geq \tau_b, \end{cases}$$

where

$$\tau_b \triangleq \frac{n \ln \frac{1-p_U}{1-p_A} - \ln[\rho \frac{\pi(A)}{\pi(U)}]}{\ln \frac{1-p_U}{1-p_A} - \ln \frac{p_U}{p_A}}$$

Proof. We start by calculating the expected loss for either decision. First:

$$\mathbb{E}_\pi(L | \varepsilon, g_A) = \ell_U \pi(U | \varepsilon) = \frac{\ell_U \pi(U) \mathbb{P}(\varepsilon | U)}{\pi(A) \mathbb{P}(\varepsilon | A) + \pi(U) \mathbb{P}(\varepsilon | U)},$$

due to our choice of L and π . Similarly,

$$\mathbb{E}_\pi(L | \varepsilon, g_U) = \ell_A \pi(A | \varepsilon) = \frac{\ell_A \pi(A) \mathbb{P}(\varepsilon | A)}{\pi(A) \mathbb{P}(\varepsilon | A) + \pi(U) \mathbb{P}(\varepsilon | U)}.$$

Combining the above expressions, the decision function (22) can then be written so that we make decision g_U if and only if:

$$\ell_A \pi(A) \mathbb{P}(\varepsilon | A) \leq \ell_U \pi(U) \mathbb{P}(\varepsilon | U).$$

Finally, replacing (20) with means p_A, p_U respectively and taking logarithms we obtain:

$$\ln[\rho \pi(A)/\pi(U)] + \varepsilon \ln \frac{p_A}{p_U} \leq (n - \varepsilon) \ln \frac{1 - p_U}{1 - p_A},$$

as a condition for deciding g_U . With some elementary manipulations, we arrive at the required result. \square

Given the conditions of the previous remark, it is easy to see (c.f. [19] ch. 8) that the decision function q minimises the Bayes risk:

$$\mathbb{E}_\pi(L | q) = \pi(A) \mathbb{P}(\varepsilon < \tau_b | A) \ell_U + \pi(U) \mathbb{P}(\varepsilon \geq \tau_b | U) \ell_A. \quad (23)$$

Furthermore, for $\pi(A) = \pi(U) = 1/2$, we obtain (18). In addition, this choice also minimises an upper bound on the worst-case expected loss since:

$$\max_{h \in H} \mathbb{E}(L | h, q) \leq \sum_{h \in H} \mathbb{E}(L | h, q) = 2 \mathbb{E}_\pi(L | q).$$

for uniform π .

Finally, the asymptotic optimality of Bayesian testing generally follows from Bayesian *consistency* (c.f. [19] ch. 10). More specifically, [21] has proved the asymptotic optimality of Bayes solutions for hypothesis testing of the type examined here.

References

- [1] W. Stallings, *Cryptography and Network Security*. Prentice Hall, 2006.
- [2] D. R. Stinson, *Cryptography: theory and practice*. Chapman & Hall, 2006.
- [3] S. Brands and D. Chaum, “Distance-bounding protocols,” in *Proc. EUROCRYPT’93*, ser. LNCS, vol. 765, 1994, pp. 344–359.
- [4] L. Bussard and W. Bagga, “Distance-bounding proof of knowledge to avoid real-time attacks,” in *Security and Privacy in the Age of Ubiquitous Computing*, vol. 181. Springer Boston, 2005, pp. 223–238.
- [5] D. Singelée and B. Preneel, “Distance bounding in noisy environments,” in *Security and Privacy in Ad-hoc and Sensor Networks - ESAS 2007*, ser. LNCS, vol. 4572. Springer, 2007, pp. 101–115.
- [6] N. O. Tippenhauer and S. Čapkun, “ID-Based Secure Distance Bounding and Localization,” in *Computer Security - ESORICS 2009*, ser. LNCS, vol. 5789. Springer Berlin / Heidelberg, 2010, pp. 621–636.
- [7] P. Peris-Lopez, J. C. Hernandez-Castro, C. Dimitrakakis, A. Mitrokotsa, and J. M. E. Tapiador, “Shedding light on RFID distance bounding protocols and terrorist fraud attacks,” arXiv:0906.4618v2, 2010.
- [8] C. Kim and G. Avoine. (2009) RFID distance bounding protocol with mixed challenges to prevent relay attacks. Cryptology ePrint Archive, Report 2009/310.
- [9] G. Hancke and M. Kuhn, “An RFID distance bounding protocol,” in *Proc. SECURECOMM’05*, Sept. 2005, pp. 67–73.
- [10] J. Reid, J. M. G. Nieto, T. Tang, and B. Senadji, “Detecting relay attacks with timing-based protocols,” in *Proc. ASI-ACCS’07*. ACM, 2007, pp. 204–213.
- [11] C. Kim, G. Avoine, F. Koeune, F. X. Standaert, and O. Pereira, “The Swiss-knife RFID distance bounding protocol,” in *Proceedings of ICISC ’08*, ser. LNCS. Springer-Verlag, Dec. 2008.
- [12] I. Boureanu, A. Mitrokotsa, and S. Vaudenay, “Practical and provably secure distance-bounding,” in *the 16th Information Security Conference (ISC 2013)*, ser. LNCS. Springer, 2013.
- [13] ———, “Secure & Lightweight Distance-Bounding,” in *Proceedings of LIGHTSEC 2013*, ser. LNCS, vol. 8162. Springer, 2013, pp. 97–113.
- [14] C. Dimitrakakis, A. Mitrokotsa, and S. Vaudenay, “Expected loss bounds for authentication in constrained channels,” in *INFOCOM, 2012 Proceedings IEEE*, march 2012, pp. 478–485.
- [15] ———, “Expected loss analysis of thresholded authentication protocols in noisy conditions,” arXiv, Tech. Rep. 1009.0278, 2011.
- [16] A. Mitrokotsa, C. Dimitrakakis, P. Peris-Lopez, and J. C. Hernandez-Castro, “Reid et al.’s distance bounding protocol and mafia fraud attacks over noisy channels,” *IEEE Communication Letters*, vol. 14, no. 2, pp. 121–123, 2010.
- [17] Y. Desmedt, “Major security problems with the “unforgeable” (Feige)-Fiat-Shamir proofs of identity and how to overcome them,” in *Proc. SecuriCom ’88*, Mar. 1988, pp. 147–159.
- [18] T. Baignères, P. Sepahrdad, and S. Vaudenay, “Distinguishing distributions using Chernoff information,” in *ProvSec 2010*. Springer-Verlag, 2010.
- [19] M. H. DeGroot, *Optimal Statistical Decisions*. John Wiley & Sons, 1970.
- [20] H. Chernoff, “Sequential design of experiments,” *Annals of Mathematical Statistics*, vol. 30, no. 3, pp. 755–770, 1959.
- [21] ———, “A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations,” *Annals of Mathematical Statistics*, vol. 23, no. 4, pp. 493–507, 1952.
- [22] M. Bellare, R. Impagliazzo, and M. Naor, “Does parallel repetition lower the error in computationally sound protocols?” in *focs*. Published by the IEEE Computer Society, 1997, p. 374.
- [23] S. Čapkun, L. Buttyán, and J. Hubaux, “Sector: Secure tracking of node encounters in multi-hop wireless networks,” in *ACM SASN*, October 2003, pp. 21–32.
- [24] S. Čapkun and J. Hubaux, “Secure positioning in wireless networks,” *IEEE Journal on Selected Areas in Communications*, vol. 24, pp. 221–232, 2006.
- [25] D. Singelée, “Study and design of a security architecture for wireless personal area networks,” Ph.D. dissertation, Katholieke Universiteit Leuven, December 2008.
- [26] L. Lai, H. E. Gamal, and H. V. Poor, “Authentication over noisy channels,” *IEEE Transactions on Information Theory*, vol. 55, no. 2, pp. 906–916, February 2009.