

Expected loss analysis of thresholded authentication protocols in noisy conditions

Christos Dimitrakakis Aikaterini Mitrokotsa
Serge Vaudenay

September 3, 2010

Abstract

A number of authentication protocols have been proposed recently, 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. This paper describes a framework enabling an *expected loss* analysis for all the protocols in this family. Furthermore, computationally simple methods to obtain nearly optimal value of the threshold, as well as for the number of rounds is suggested. Finally, a method to adaptively select both the number of rounds and the threshold is proposed.

1 Introduction

Traditionally [16, 17], authentication is assumed to be taking place on an error-free channel, and error analysis is performed separately from cryptographic analysis of protocols. However, a number of authentication protocols have been proposed [2, 3, 7, 9, 10, 12, 13, 15, 18], where at least some part of the authentication is performed during a challenge-response phase lasting n rounds with no error correction, due to a need to detect relay attacks by timing delays. The noise necessitates the use of a tolerance threshold τ , such that a party is authenticated if the total error of its responses ϵ is below the threshold τ .

This paper introduces a general framework for analysing such protocols. We assign a cost ℓ_A to the event that we authenticate a malicious party A —which we call the *attacker*—a cost (or loss) ℓ_U to the event that we fail to authenticate a valid party U —which we call the *user*—and a cost ℓ_B for each round of the challenge-response phase. Our goal is to select n, τ so as to minimise the *expected loss* $\mathbb{E}L$.

The paper is organised as follows. Section 1.1 presents related work, while Section 2 introduces notation and thresholded authentication protocols. Section 3 contains the *expected loss* analysis under noise. In particular, Sec. 3.1 suggests a method to calculate the threshold accompanied by a finite sample loss bound, while Sec. 3.2 provides a further bound by selecting an appropriate number of rounds. These results only require that the expected error of the attacker is higher than that of the user. Section 4 applies the above analysis to a number of currently used protocols. Section 5 suggests a high-probability method for estimating the channel noise and presents the results of simulation

experiments that compare our choice of threshold with thresholds derived using asymptotic approximations. Finally, Sec. 6 concludes the paper. For completeness, the appendix provides some useful auxiliary results regarding the finite sample and the asymptotic derivations.

1.1 Related Work

In certain authentication protocols part of the communication is performed in noisy channels without employing error correction. Specifically, a *rapid-bit exchange phase* was introduced in [2] to compute an upper bound on the distance of the prover \mathcal{P} . 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 [7, 10, 12, 13, 15], as well as other wireless devices [18–20].

Hancke and Kuhn [7] 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. [10] extend this approach with the SWISS-KNIFE protocol by considering *three* types of errors. Finally [15], rather than using a threshold τ , proposed a protocol (henceforth ECMAD) using an error correcting code (ECC). ECMAD, which extends the MAD protocol [19], 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 [14], 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. [1] 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 [1] can be summarised below: (a) We perform an expected loss analysis rather than an error analysis. (b) Our bounds hold uniformly, while [1] 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 [1], due to the different setting.

A more general work on authentication under noisy conditions was presented in [11]. 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, 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.

1.2 Our contribution

In this paper, we perform a detailed *expected loss* analysis for a general class of multi-round 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 [12], 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.¹ Furthermore, they are valid for any n , since they are based on probability inequalities for a finite number of samples.² Thus, they are considerably more general to the bounds of [14].

Finally, we provide 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 take 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.

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 random variable $X \in \Omega$ when A is true. The notation \mathcal{F}_Ω will denote an appropriate σ -field on Ω . Finally, $\mathbb{I}\{A\}$ is an indicator function equal to 1 when A is true and 0 otherwise.

¹In all previous proposals, there is either an error at each round, or there is not.

²The analysis in [14] only holds for large n , so the approximation only holds asymptotically.

We consider shared secret challenge-response authentication protocols with multi-round exchanges. In such protocols, a verifier \mathcal{V} grants access to a prover \mathcal{P} , if the latter can demonstrate its identity I and possession of a shared secret $s \in \mathcal{X}^m$. The protocol has three phases: (i) An initialisation phase. (ii) A *rapid-bit exchange* phase, lasting n rounds. (iii) A termination phase. A frequent assumption is that the authentication takes place in a noise-free channel. The extension to noisy channels is done by assuming the existence of an error correcting protocol. Thus, the *error* analysis is performed separately from the *cryptographic* analysis. Here we shall integrate the two aspects of the problem by performing an *expected loss* analysis of the authentication protocol directly on the noisy channel. We shall assume that the initialisation and termination phases are fixed (due to other security considerations) and focus on the *rapid-bit exchange* phase.

Due to noise in the physical medium, 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: $\omega \triangleq \mathbb{P}(x \neq \hat{x}), \forall x, \hat{x} \in \mathcal{X}$. For simplicity, we shall only treat the case of symmetric channel noise such that: $\mathbb{P}(\hat{x} = y \mid x \neq \hat{x}) = \frac{1}{|\mathcal{X}|-1}, \forall y \neq x, x, y \in \mathcal{X}$.

2.1 Thresholded protocols

During *multi-round challenge response authentication* phase (e.g. the *rapid-bit exchange* phase in an RFID distance-bounding protocol) 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} can calculate the correct responses $R(c_i)$ and so can calculate an error ε_i for the i -th round. While the legitimate user U should also be able to calculate the correct responses, due to noise, there may be errors in the received responses. On the other hand, the attacker has to resort to guessing, so 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 need a threshold value τ , such that a prover is accepted if and only if the total error observed is smaller than τ . More precisely, we define:

Definition 1. *An additive thresholded multi-round challenge-response authentication protocol has the following parameters:*

1. A natural number $n > 0$, equal to the number of challenge-response rounds.
2. A threshold $\tau \geq 0$.
3. An error function $\varepsilon \triangleq \sum_{i=1}^n \varepsilon_i$, where $\varepsilon_i \in [0, 1]$ represents the error of the i -th round.

The verifier \mathcal{V} rejects the prover (authenticator) \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 define $p_A \leq \mathbb{E}(\varepsilon_i \mid A)$, a lower bound on the expected per-round error of the attacker and $p_U \geq \mathbb{E}(\varepsilon_i \mid U)$, an upper bound on the error of a legitimate user.

These bounds depend on the noise and on the protocol under consideration. We shall return to them in section 4.

3 Expected loss analysis

We now specify our potential losses. For every round of the *rapid-bit exchange phase*, we suffer loss ℓ_B . In addition, we suffer a loss of ℓ_A for each false acceptance and a loss ℓ_U for false rejection.³ Given that we perform n rounds, the total loss when the prover \mathcal{P} is either the legitimate user U or the attacker A is 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 $\mathbb{E}L$. 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, $\mathbb{E}L$.

The remainder of this section is organised as follows. Section 3.1 shows how a nearly-optimal threshold $\hat{\tau}_n^*$ for a given number of rounds n can be chosen and gives bounds on the expected loss for that choice. Section 3.2 shows how a nearly-optimal number of rounds \hat{n}^* can be chosen and gives further bounds.

3.1 Choice of threshold

We want to choose 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. The problem is that as we *increase* the threshold τ , $\mathbb{E}(L | \mathcal{P} = U)$ *decreases*, while $\mathbb{E}(L | \mathcal{P} = A)$ *increases*. The opposite is happening when we *decrease* the threshold τ . Thus, to minimise the worst-case expected loss, we can choose a threshold τ such that $\mathbb{E}(L | \mathcal{P} = A, \tau) = \mathbb{E}(L | \mathcal{P} = U, \tau)$. A particular choice of the threshold τ that minimises an upper bound on the worst-case expected loss is given in Theorem 1. 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:

³These 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 .

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

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

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

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

the last two steps used the fact that $\bar{z}_A \geq p_A$ and the *Hoeffding inequality* (19). Specifically, in our case, Lemma 2 (page 14) applies with $X_i = \varepsilon_i$. Then, it is easy to see that $\mu_i = \bar{z}_A$ 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 user 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

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 loss bound.

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

$$\tau = \hat{\tau}_n^* \triangleq \frac{n(p_A + p_U)}{2} - \frac{\ln \rho}{4\Delta} \quad (6)$$

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

$$\mathbb{E}(L | n, \hat{\tau}_n^*) \leq \mathcal{L}_1(n) \triangleq n\ell_B + e^{-\frac{n}{2}\Delta^2} \cdot \sqrt{\ell_A \ell_U}. \quad (7)$$

with $\Delta \triangleq p_A - p_U$.

Proof. Substitute (6) in the first exponential of (5) to obtain:

$$\begin{aligned} e^{-\frac{2}{n}(np_U - \hat{\tau}_n^*)^2} \ell_U &= \exp \left(-\frac{n}{2}\Delta^2 + \frac{1}{2} \ln \rho - \frac{\ln^2 \rho}{8n\Delta^2} + \ln \ell_U \right) \ell_U \\ &= e^{-\frac{n}{2}\Delta^2} e^{-\frac{\ln^2 \rho}{8n\Delta^2}} \sqrt{\ell_A \ell_U}. \end{aligned}$$

It is easy to see that the exact same result is obtained by substituting (6) in the second exponential of (5). Thus, both $\mathbb{E}(L | A)$ and $\mathbb{E}(L | U)$ are bounded by

the same quantity and consequently, so is $\max\{\mathbb{E}(L | A), \mathbb{E}(L | U)\}$. Thus,

$$\begin{aligned}\mathcal{L}(n, \hat{\tau}_n^*) &\leq n\ell_B + \exp\left(-\frac{n}{2}\Delta^2\right) \cdot \exp\left(-\frac{\ln^2 \rho}{8n\Delta^2}\right) \cdot \sqrt{\ell_A \ell_U} \\ &\leq n\ell_B + e^{-\frac{n}{2}\Delta^2} \sqrt{\ell_A \ell_U},\end{aligned}$$

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

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 $n\rho_U$ and $n\rho_A$. This also implies a lower bound on the number of rounds n .

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 a good choice for the number of rounds n and provide a matching bound on the expected loss.

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

$$n = \hat{n}^* \triangleq \frac{\sqrt{1+2CK} - 1}{C}, \quad (8)$$

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

$$\mathbb{E}(L | \hat{\tau}_n^*, \hat{n}^*) \leq \mathcal{L}_2 \triangleq \sqrt{8K/C} \cdot \ell_B = \frac{\sqrt{8\ell_B}(\ell_A \ell_U)^{1/4}}{\Delta}. \quad (9)$$

Proof. We shall bound each one of the summands of (7) 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:

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

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

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.

4 Analysis of RFID thresholded protocols

Currently, the only known protocols employing an authentication phase without any error correction that we are aware of are *RFID distance bounding* protocols. For that reason, we shall examine the properties of two such protocols, for which it is possible to derive expressions for p_A, p_U given a symmetric channel noise ω .

The SWISS-KNIFE protocol [10] and the variant HITOMI [12] are thresholded authentication protocols satisfying Definition 1. It is easy to show (for details see [12]) that, for those two protocols, under channel noise ω , the expected error bounds p_A, p_U are given by: $p_A = \frac{\omega+1}{2}$, $p_U = 2\omega$, where we note in passing that $p_A \geq p_U$ and so $\omega \leq \frac{1}{3}$. Finally, by substituting $\Delta = p_A - p_U = \frac{1-3\omega}{2}$ in (7), we obtain the following bound for SWISS-KNIFE and HITOMI:

$$\mathbb{E} L \leq n\ell_B + e^{-\frac{n(1-3\omega)^2}{8}} \cdot \sqrt{\ell_A \ell_U}, \quad (10)$$


We have performed a number of experiments to test the efficacy of these protocols, when used in conjunction with our suggested, as well as the optimal values of the threshold and number of rounds. In all of the experiments shown here, we chose the following values for the losses: $\ell_A = 10$, $\ell_U = 1$, $\ell_B = 10^{-2}$.

Figure 1(a) depicts the bound (10) on the expected loss, as well as the actual $\mathbb{E} L$ calculated via the binomial formula, when the threshold $\hat{\tau}_n^*$, calculated from (6), is used. We plot both the expected loss and the bound for two different channel noise levels $\omega \in \{10^{-1}, 10^{-2}\}$, where the number of rounds n varies from 1 to 256. Obviously, the bound is greater than the actual expected loss, while it approaches it exponentially fast as n increases. In addition, the losses are higher when the amount of noise increases.


Furthermore, we can see that there are minimising values of n for all cases. While they do not coincide for the bound and the actual *expected loss*, they are within a factor of two of each other. Finally, \hat{n}^* , the value of n minimising the bound, is always greater than $n^* \triangleq \operatorname{argmin}_n \max_{\mathcal{P} \in A, U} \mathbb{E}(L \mid \mathcal{P}, n)$, the value of n that minimises the *worst-case expected loss*. Since the probability of incorrect authentication always decreases with increasing n , this implies that any additional losses incurred by using \hat{n}^* is due to transmission costs only.

Figure 1(b) examines the effect of noise in more detail. In particular, it depicts the *worst-case expected loss* for the optimal number of rounds n^* , denoted by $\mathbb{E}(L \mid n^*)$ in the legend. 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 the Theorem 1 needs considerable tightening for small ω . Finally, \mathcal{L}_2 is considerably looser, and thus it is only of theoretical interest.

Finally, due to the way that the protocols under consideration generate challenges and responses, the number of rounds n must be smaller than the length k of the messages in the initialization phase and also the length of the key x . Thus, in practice we will always select a number of rounds $n = \min\{n^*, k\}$. This condition is necessary since the responses r for both protocols are calculated using an XOR operation between the secret key x and a constant value (i.e. either α or β) that has the same length k . Since these protocols are only used as examples, it is beyond the scope of this paper to propose protocols that do not suffer from this limitation.



(a) The *Expected Loss* $\mathbb{E}L$ and the bound on the *Expected Loss* \mathcal{L}_1 vs. the number of bits n exchanged during the rapid single bit exchange for various values of channel noise ω .



(b) The worst-case expected Loss L and the bounds \mathcal{L}_1 and \mathcal{L}_2 from theorems 1 and 2 respectively vs. the channel error rate ω .

Figure 1: Comparison of all losses.

5 Estimating ω

In this section, we discuss how it is possible to calculate the channel error rate ω , which is used in the expressions for p_A, p_U . This can be done by 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:

$$\gamma_{\min} \triangleq \{\gamma(\Phi(x), \Phi(y)) : x, y \in \mathcal{X}^m, x \neq y\} \quad (11)$$

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:

$$\mathbb{E}(\hat{\omega} \mid \theta \leq (\gamma_{\min} - 1)/2) = \omega. \quad (12)$$

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 2). In particular, it is easy to show that, for any $\delta \in [0, 1]$:

$$\mathbb{P}\left(|\hat{\omega} - \omega| \geq \sqrt{\frac{\ln 2/\delta}{2k}}\right) \leq \delta, \quad (13)$$

by substituting the square-root term into (19), 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 $1 - \delta$:

$$p_A = \frac{1 + \hat{\omega}}{2} + \sqrt{\frac{\ln 2/\delta}{8k}}, \quad p_U = 2\hat{\omega} - \sqrt{\frac{2 \ln 2/\delta}{k}}. \quad (14)$$

While we were unable to provide bounds on the performance of this choice, experimental investigations presented in the next section indicate that it has good performance.

5.1 Evaluation Experiments

We have performed some experiments to evaluate our methods in a more realistic setting, involving an RFID distance bounding protocol with a *rapid-bit exchange* phase. We perform simulations for two cases: Firstly, when a *legitimate user* U is trying to get authenticated and secondly, when an *adversary* A is trying

to perform a mafia fraud attack. 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 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 following values for the losses: $\ell_A = 10$, $\ell_U = 1$, $\ell_B = 10^{-2}$, while we used $k = 2^{10}$ for the coded messages in the initialisation phase.

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 $\hat{\omega} = \hat{\theta}/k$. 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 (14) for p_A, p_U , with an arbitrary value of δ .

In the first experiment, we use the nearly-optimal threshold and number of rounds that we have derived in our analysis. In the second experiment, we replace our choice of threshold with a choice similar to that of Baignères et al. [1]. Their threshold is derived via the likelihood ratio test, which is asymptotically optimal (c.f. [4, 6])

$$\hat{\tau} = \frac{n \ln \frac{1-p_U}{1-p_A}}{\ln \frac{1-p_U}{1-p_A} - \ln \frac{p_U}{p_A}} \quad (15)$$

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 [6] applies) to obtain:

$$\tilde{\tau} = \frac{n \ln \frac{1-p_U}{1-p_A} - \ln \rho}{\ln \frac{1-p_U}{1-p_A} - \ln \frac{p_U}{p_A}} \quad (16)$$

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 (16) can be approximated by:

$$\tilde{\tau}^* = n\bar{p} - \frac{\bar{p}(1-\bar{p})}{\Delta} \ln \rho \quad (17)$$

More details on the derivation of (16) are given in Appendix A.2.

Figure 2 depicts the *worst-case expected loss* \mathbb{L} as a function of the actual noise ω . Figure 2(a) shows \mathbb{L} using the threshold τ derived from our *expected loss* analysis (6), while in Figure 2(b) we use the asymptotically optimal threshold of (16). 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 (14) with $\delta \in \{10^{-1}, 10^{-2}\}$.

As it can be seen in Figure 2, in all cases (using ours Figure 2(a) or Baignères et al. [1] threshold Figure 2(b)) the *worst-case expected loss* is very low for small values of the actual noise and increases sharply when the actual noise exceeds the value of 10^{-1} . It is interesting to see that when we use the optimistic⁴ high

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.

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 (16), we observe a deterioration in the results.

As mentioned in the related work (Sec. 1.1), the choice of the threshold by Baignères et al. [1] 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.


Finally, note that in the SWISS-KNIFE and HITOMI protocols, the size of the initialisation messages is fixed, since the number of rounds n is fixed. In practice, one would have to modify these protocols in order for them to work with an arbitrary number of rounds, but this subject is beyond the scope of this paper. Our focus is mainly the *expected loss* analysis of the noisy authentication phase.

6 Conclusion


We have performed an expected loss analysis for thresholded authentication protocols under noise. This 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. More precisely, we have shown how to select a threshold and provided an upper bound on the worst-case expected loss. Additionally, we have shown how to select the number of rounds in order to tighten the loss and obtained a loss bound that holds uniformly and depends only on the error rates of the user and attacker and the individual losses. We have applied these choices of threshold and number of rounds to two representative distance bounding protocols, the SWISS-KNIFE and the HITOMI. In addition, we have presented a high-probability method for estimating the channel noise for such protocols. We have examined its performance in further simulation experiments with *unknown channel noise*, and shown that we obtain uniformly superior results to guessing or direct empirical noise estimates. Finally, we repeated those experiments with a asymptotically optimal threshold similar to that used by Baignères et al. [1]. 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.

Acknowledgements

This work was partially supported by the Marie Curie IEF project “PPIDR: Privacy-Preserving Intrusion Detection and Response in Wireless Communications”, grant number: 252323 and by the IM-CLeVeR EU FP7 Integrated Project, grant number: 231722.



(a) Our threshold



(b) Asymptotic threshold

Figure 2: 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 (14) with $\delta \in \{10^{-1}, 10^{-2}\}$.

A Auxilliary results

A.1 Useful formulas

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

$$\mathbb{P}\left(\sum_{k=1}^n X_k \geq u\right) = \sum_{k=0}^u \binom{n}{k} \mu^k (1-\mu)^{n-k}. \quad (18)$$

This probability can be bounded via *Hoeffding's inequality* [8]:

Lemma 2 (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 $t > 0$:*

$$\mathbb{P}\left(\sum_{i=1}^n X_i \geq \sum_{i=1}^n \mu_i + nt\right) = \mathbb{P}\left(\sum_{i=1}^n X_i \leq \sum_{i=1}^n \mu_i - nt\right) \leq \exp\left(-\frac{2n^2 t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right). \quad (19)$$

A.2 On asymptotic thresholds

One way to obtain an asymptotically optimal threshold is to employ a Bayesian hypothesis test [6]. 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 and it follows a binomial distribution with parameters p_A, p_U respectively. 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 $G = \{g_A, g_U, g_\emptyset\}$, where g_A means we decide that the prover is an attacker and g_U means we decide that the prover is a user and g_\emptyset means that we are undecided. 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 $q : \{0, 1, \dots, n\} \rightarrow G$:

$$q(\varepsilon) \triangleq \begin{cases} g_U, & \text{if } \mathbb{E}_\pi(L | \varepsilon, g_U) \leq \mathbb{E}_\pi(L | \varepsilon, g_A) \\ g_A, & \text{if } \mathbb{E}_\pi(L | \varepsilon, g_U) > \mathbb{E}_\pi(L | \varepsilon, g_A). \end{cases} \quad (20)$$

This decision function minimises $\mathbb{E}_\pi L$ by construction (c.f. [6] ch. 8). The following remark is applicable in our case:

Remark 1. 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 (20) 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 (20) 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 (18) 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. [6] ch. 8) that the decision function q minimises the Bayes risk:

$$\mathbb{E}_\pi(L | q) = \pi(A) \mathbb{P}(\varepsilon < \tau_b | A) \ell_U + \pi(U) \mathbb{P}(\varepsilon \geq \tau_b | U) \ell_A. \quad (21)$$

Furthermore, for $\pi(A) = \pi(U) = 1/2$, we obtain (16). 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. [6] ch. 10). More specifically, [5] has proved the asymptotic optimality of Bayes solutions for hypothesis testing of the type examined here.

References

- [1] T. Baignères, P. Sepehrdad, and S. Vaudenay. Distinguishing distributions using chernoff information. In *4th International Conference on Provable Security 2010 (ProvSec 2010)*, Malacca, Malaysia, 13-15 October 2010. Springer-Verlag.
- [2] S. Brands and D. Chaum. Distance-bounding protocols. In *Proc. EURO-CRYPT'93*, volume 765 of *LNCS*, pages 344–359, 1994.
- [3] 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*, volume 181 of *IFIP International Federation for Information Processing*, pages 223–238. Springer Boston, 2005.
- [4] H. Chernoff. Sequential design of experiments. *Annals of Mathematical Statistics*, 30(3):755–770, 1959.
- [5] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23(4):493–507, 1952.
- [6] M. H. DeGroot. *Optimal Statistical Decisions*. John Wiley & Sons, 1970.
- [7] G. Hancke and M. Kuhn. An RFID distance bounding protocol. In *Proc. SECURECOMM'05*, pages 67–73, Sept. 2005.
- [8] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, March 1963.
- [9] C.H. Kim and G. Avoine. RFID distance bounding protocol with mixed challenges to prevent relay attacks. Cryptology ePrint Archive, Report 2009/310, 2009.
- [10] C.H. Kim, G. Avoine, F. Koeune, F. X. Standaert, and O. Pereira. The Swiss-knife RFID distance bounding protocol. In *Proceedings of ICISC '08*, LNCS. Springer-Verlag, Dec. 2008.
- [11] L. Laia, H. El Gamal, and H. V. Poor. Authentication over noisy channels. *IEEE Transactions on Information Theory*, 55(2):906–916, February 2009.
- [12] 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. *IEEE Transactions on Dependable and Secure Computing*, Submitted, 2010.
- [13] J. Reid, J. M. Gonzalez Nieto, T. Tang, and B. Senadji. Detecting relay attacks with timing-based protocols. In *Proc. ASIACCS'07*, pages 204–213. ACM, 2007.
- [14] D. Singelée. *Study and Design of a Security Architecture for Wireless Personal Area Networks*. PhD thesis, Katholieke Universiteit Leuven, December 2008.

- [15] D. Singelée and B. Preneel. Distance bounding in noisy environments. In *Security and Privacy in Ad-hoc and Sensor Networks - ESAS 2007*, volume 4572 of *LNCS*, pages 101–115. Springer, 2007.
- [16] W. Stallings. *Cryptography and Network Security*. Prentice Hall, 2006.
- [17] D. R. Stinson. *Cryptography: theory and practice*. Chapman & Hall, 2006.
- [18] N. O. Tippenhauer and S. Čapkun. ID-Based Secure Distance Bounding and Localization. In *Computer Security - ESORICS 2009*, volume 5789 of *LNCS*, pages 621–636. Springer Berlin / Heidelberg, 2010.
- [19] S. Čapkun, L. Buttyán, and J. P. Hubeaux. Sector: Secure tracking of node encounters in multi-hop wireless networks. In *ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, pages 21–32, Washington, USA, October 2003. ACM Press.
- [20] S. Čapkun and J.P. Hubeaux. Secure positioning in wireless networks. *IEEE Journal on Selected Areas in Communications*, 24:221–232, 2006.
- [21] Abraham Wald. *Sequential Analysis*. John Wiley & Sons, 1947. Republished by Dover in 2004.