

Riding the Saddle Point: asymptotics of the capacity-achieving simple decoder for bias-based traitor tracing

Sarah Ibrahimi, Boris Škorić, Jan-Jaap Oosterwijk

Abstract

We study the asymptotic-capacity-achieving score function that was recently proposed by Oosterwijk et al. for bias-based traitor tracing codes. For the bias function we choose the Dirichlet distribution with a cutoff. Using Bernstein’s inequality and Bennett’s inequality, we upper bound the false positive and false negative error probabilities. From these bounds we derive sufficient conditions for the scheme parameters. We solve these conditions in the limit of large coalition size c_0 and obtain asymptotic solutions for the cutoff, the sufficient code length and the corresponding accusation threshold. The code length converges to its asymptote approximately as $c_0^{-1/2}$, which is faster than the $c_0^{-1/3}$ of Tardos’ score function.

1 Introduction

1.1 Traitor tracing

Forensic watermarking is a means for tracing unauthorized redistribution of digital content. Before distribution, the content is modified by embedding an imperceptible watermark, which plays the role of a personalized identifier. When an unauthorized copy of the content is found, a tracing algorithm outputs a list of suspicious users, based on the watermark detected in this copy.

The most powerful attacks against watermarking are *collusion attacks*, in which multiple attackers (the ‘coalition’) combine their differently watermarked versions of the same content; the observed differences point to the locations of the hidden marks and allow for a targeted attack.

Collusion-resistant codes have been specifically designed as a defense against collusion attacks: when codewords from such a code are embedded into the content, the surviving parts of the watermark, after the collusion attack, still contain enough information to identify (some of the) attackers, provided that the coalition is not too large.

In the past two decades several types of collusion-resistant codes have been developed. The most popular type in the recent literature is the class of *bias-based* codes. These were introduced by G. Tardos in 2003. The code construction consists of two steps: first a sequence of biases is generated, one for each position in the content; then the watermark symbols for each user are randomly drawn according to these biases. The original paper [22] was followed by a flurry of activity, e.g. improved analyses [4, 8, 9, 12, 21, 26], code modifications [11, 17, 18], decoder modifications [1, 7, 15, 20] and various generalizations [6, 23, 24, 27]. The advantage of bias-based versus deterministic codes is that they can achieve the asymptotically optimal relationship $\ell \propto c_0^2$ between the sufficient code length ℓ and the coalition size c_0 to be resisted.

1.2 Capacity-achieving simple decoder

Two kinds of tracing algorithm can be distinguished: (i) *simple decoders*, which assign a score to single users independent of the watermarks of other users, and (ii) *joint decoders* [1, 7, 15], which assign scores to sets of users and are typically more powerful, but also require more computational resources. Efficient joint decoders typically employ a simple decoder as a bootstrapping step.

The performance of a traitor tracing code is often measured by looking at the sufficient code length ℓ as a function of the coalition size c_0 to be resisted and the imposed low error rate. Equivalently, one can look at the *fingerprinting rate*, which is defined as the fraction $\frac{\log_q n}{\ell}$, where q is the size of the alphabet and n is the number of users. The numerator corresponds to the number of q -ary symbols needed to point out one of the n users; the denominator is the number of symbols used to convey this ‘message’. Hence the fingerprinting rate has a natural interpretation as the fraction of codeword symbols that actually encodes the ‘message’, i.e. the identifying information that allows for tracing. The fingerprinting rate is a figure of merit that can be used to fairly compare codes which have different alphabet size. The *fingerprinting capacity*, which can be computed information-theoretically, is an upper bound on the fingerprinting rate that can be achieved against colluders who employ an optimal strategy against the tracing scheme. It was found by Boesten and Škorić [5] that the asymptotic¹ capacity is given by

$$C = \frac{q - 1}{2c_0^2 \ln q}. \quad (1)$$

Huang and Moulin [10] found the location of the corresponding asymptotic *saddlepoint*: the strongest attack is the so-called interleaving attack, and the best bias distribution is the Dirichlet distribution with concentration parameter $1/2$. (See Section 2.) For the colluders as well as the tracer it is bad to depart from the saddlepoint. If the colluders move away from it, the tracer can achieve a higher fingerprinting rate; if the tracer moves away, the colluders can launch a stronger attack which reduces the rate.

Oosterwijk et al. [19] devised a simple decoder that reaches asymptotic capacity. The possibility of such an achievement was foreseen in [10], where it was shown that the simple decoder capacity becomes equal to the joint decoder capacity as c_0 goes to infinity.

1.3 Contributions and outline

In this paper we analyze the performance of the capacity-achieving simple decoder of [19] in the Restricted Digit Model.

- Following the approach of [25], we use Bernstein’s inequality and Bennett’s inequality to upper bound the false positive and false negative error probability, respectively. From these bounds we derive conditions on the code parameters (code length, cutoff, threshold) such that the error probabilities are sufficiently low.
- We determine the asymptotics of the sufficient code length in the direct vicinity of the saddlepoint.
- We find that the optimal choice for the cutoff τ is given by $\tau \propto c_0^{-\gamma}$, with γ slightly larger than $1/2$. With this choice, the code length approaches its saddlepoint value with a correction term of order $c_0^{\gamma-1} \approx c_0^{-1/2}$. Thus, convergence to the limit is faster than in the case of the binary Tardos score, where the correction is of order $c_0^{-1/3}$ [12].
- Our analysis yields a recipe for placing the accusation threshold as a function of the innocent user score variance. This differs from the case of the Tardos score function [22, 23], where the threshold is fixed.

In Section 2 we briefly review bias-based traitor tracing, the asymptotic saddlepoint, and the asymptotic-capacity-achieving score function. We also list the inequalities of Bernstein and Bennett. In Section 3 we study the statistical properties of an innocent user’s score and the coalition’s collective score. In Section 4 we derive the bounds on the error rates and the sufficient conditions on the code parameters. The asymptotics of the sufficient code length are treated in Section 5.

¹Throughout this paper, the term asymptotic refers to the limit of large coalition size.

2 Preliminaries

2.1 Bias-based tracing using the asymptotically optimal simple decoder

2.1.1 Notation

The number of users is denoted as n , and the code length (the number of positions in the content) as ℓ . We define $[n] = \{1, \dots, n\}$. The alphabet is \mathcal{Q} , with size $|\mathcal{Q}| = q$. The symbols in the alphabet have no natural ordering. The bias in position i is denoted as $\mathbf{p}^{(i)}$. The bias is a q -dimensional vector, with components $p_\alpha^{(i)} \in [\tau, 1 - (q-1)\tau]$, $\alpha \in \mathcal{Q}$. The parameter $\tau \ll 1$ is called the cutoff. For each i the bias satisfies $|\mathbf{p}^{(i)}| = 1$, where $|\dots|$ denotes the 1-norm, i.e. $\sum_{\alpha \in \mathcal{Q}} p_\alpha^{(i)} = 1$. We will often use multi-index notation: for a scalar z , the notation \mathbf{p}^z stands for $\prod_{\alpha \in \mathcal{Q}} p_\alpha^z$; for a vector \mathbf{m} , the notation $\mathbf{p}^{\mathbf{m}}$ stands for $\prod_{\alpha \in \mathcal{Q}} p_\alpha^{m_\alpha}$. We introduce the q -component vector $\mathbf{1}_q = (1, 1, \dots, 1)$. The notation δ_{xy} stands for the Kronecker delta.

2.1.2 Code generation

The bias vectors $\mathbf{p}^{(i)}$ are drawn independently from a (truncated) Dirichlet distribution F with concentration parameter $\kappa > 0$,

$$F(\mathbf{p}) = \mathbf{p}^{-1+\kappa} / B_\tau(\kappa \mathbf{1}_q) \quad (2)$$

$$B_\tau(\kappa \mathbf{1}_q) = \int_\tau^{1-(q-1)\tau} d^q p \delta(1 - |\mathbf{p}|) \mathbf{p}^{-1+\kappa}. \quad (3)$$

The δ in the integral is a Dirac delta function; it ensures that the condition $|\mathbf{p}| = 1$ is enforced. The τ is called the cutoff parameter. Note that $p_\alpha \in [\tau, 1 - (q-1)\tau]$. Therefore $\tau \leq 1/q$ must hold, for otherwise the interval is empty (and we would get $|\mathbf{p}| > 1$).

For $\tau = 0$ the normalization constant (3) evaluates to a generalized Beta function. Let $\mathbf{z} \in (0, \infty)^q$ be a vector; then the Beta function $B(\mathbf{z})$ is defined as $B(\mathbf{z}) = [\prod_{\alpha} \Gamma(z_\alpha)] / \Gamma(\sum_{\beta} z_\beta)$, where Γ is the Gamma function. Hence $B_0(\kappa \mathbf{1}_q) = B(\kappa \mathbf{1}_q) = [\Gamma(\kappa)]^q / \Gamma(q\kappa)$.

In the asymptotic saddlepoint it holds that $\tau = 0$ and $\kappa = 1/2$. For large but finite c_0 the saddlepoint lies close to the asymptotic saddlepoint, but it is not known exactly where. It is known that for finite c_0 the optimal bias distribution is a *discrete* distribution [18, 11, 13], with a number of discrete p_α values proportional to c_0 . In spite of this, we will use the continuous probability density (2). Our motivation is that we only investigate asymptotics. The cutoff τ will depend on c_0 .

The code word assigned to user j is denoted as a row vector $X_j = (X_{j1}, \dots, X_{j\ell})$. The set of codewords is arranged in a code matrix X . The elements of the code matrix are independently generated according to the biases $\mathbf{p}^{(1)}, \dots, \mathbf{p}^{(\ell)}$ as follows: $\Pr[X_{ji} = \alpha] = p_\alpha^{(i)}$.

2.1.3 Collusion attack

The coalition is a subset $\mathcal{C} \subset [n]$ of users, with size $|\mathcal{C}| = c$. We explicitly make the distinction between the actual coalition size c and the parameter c_0 in the code construction, which is the maximum coalition size that can be resisted. The colluders see a submatrix $X_{\mathcal{C}}$ of X . The symbol ‘tallies’ are defined as follows,

$$\mathbf{m}^{(i)} = (m_\alpha^{(i)})_{\alpha \in \mathcal{Q}} \quad ; \quad m_\alpha^{(i)} = |\{j \in \mathcal{C} : X_{ji} = \alpha\}|. \quad (4)$$

In words: $m_\alpha^{(i)}$ is the number of colluders that received symbol α in position i .

Based on $X_{\mathcal{C}}$ the colluders produce an output $\mathbf{y} = (y_1, \dots, y_\ell)$. For our analysis we adopt the Restricted Digit Model as the attack model: for any $i \in [\ell]$, the output y_i is only allowed to be a symbol that the colluders have observed in position i . The strategy for choosing an output is allowed to be probabilistic. We adopt a number of frequently made assumptions about the attack strategy.

1. *Symbol symmetry.* The strategy is invariant under permutation of the alphabet for each position independently. This assumption is motivated by the lack of a natural ordering of the alphabet.
2. *Colluder symmetry.* The strategy is invariant under permutation of the colluders. (In other words, the colluders equally share the risk.) This assumption is motivated by the fact that breaking colluder symmetry will make it easier for the tracer to find at least one colluder.
3. *Position symmetry.* The same strategy is applied in each position $i \in [\ell]$, and it does not depend on any X_{jk} values with $k \neq i$. Motivation: asymptotically the optimal attack must be position-symmetric [16].

When assumptions 2 and 3 hold, the strategy can be parametrized by a set of probabilities that depend only on the ‘local’ tallies: in position i , the probability of outputting symbol y_i is a function of only $\mathbf{m}^{(i)}$. Omitting the position index, this is denoted as

$$\theta_{y|\mathbf{m}} = \Pr[\text{colluders output } y | \text{the tally is } \mathbf{m}]. \quad (5)$$

Furthermore, if assumption 1 holds as well, it is possible [21] to re-parametrize this as

$$\Psi_b(\mathbf{x}) = \theta_{y|\mathbf{m}} \text{ for } \{m_y = b, \text{ and } \mathbf{m} \text{ without the } y \text{ component is } \mathbf{x}\}. \quad (6)$$

In other words, $\Psi_b(\mathbf{x})$ is the coalition’s probability of outputting a symbol given that it has tally b , and that the other tallies are \mathbf{x} . The probability $\Psi_b(\mathbf{x})$ is invariant under permutation of \mathbf{x} .

2.1.4 Simple decoder

The tracer notices the pirated copy with watermark sequence \mathbf{y} ‘in the wild’. Based on \mathbf{y} and X he tries to find at least one colluder. The asymptotic-capacity-achieving simple decoder of [19] works as follows: for each user $j \in [n]$ a score $S_j = \sum_{i \in [\ell]} S_j^{(i)}$ is computed, where

$$S_j^{(i)} = h(X_{ji}, y_i, \mathbf{p}^{(i)}) \quad \text{with} \quad h(x, y, \mathbf{p}) = \frac{\delta_{xy}}{p_y} - 1. \quad (7)$$

Note that we normalized the function h differently from [19], by a factor $\sqrt{q-1}$, for notational brevity. The score function (7) has the special property of being ‘strongly centered’: for any \mathbf{p} and y (we are omitting the position index) the expected score of an innocent user is zero.

$$\sum_{x \in \mathcal{Q}} p_x h(x, y, \mathbf{p}) = \frac{p_y}{p_y} - \sum_{x \in \mathcal{Q}} p_x = 0. \quad (8)$$

The collective score of the coalition is written as $S_{\mathcal{C}}$,

$$S_{\mathcal{C}} = \sum_{j \in \mathcal{C}} S_j. \quad (9)$$

The tracer makes a list \mathcal{L} of ‘suspicious’ users, whose score exceeds a threshold Z ,

$$\mathcal{L} = \{j \in [n] : S_j > Z\}. \quad (10)$$

2.1.5 Measuring the performance

Two types of error can occur: a False Positive, with P_{FP} defined as the probability that a fixed innocent user gets added to \mathcal{L} ; and a False Negative, with P_{FN} defined as the probability that none of the colluders is found,

$$P_{\text{FP}} = \Pr[j \in \mathcal{L}] \text{ for fixed innocent } j \quad ; \quad P_{\text{FN}} = \Pr[\mathcal{C} \cap \mathcal{L} = \emptyset] \quad (11)$$

The tracer demands that $P_{\text{FP}} \leq \varepsilon_1$ and $P_{\text{FN}} \leq \varepsilon_2$, where ε_1 and ε_2 are constants, typically with $\varepsilon_1 \ll \varepsilon_2$.

The code length ℓ and threshold Z are often parametrized as

$$\ell = Ac_0^2 \ln \frac{1}{\varepsilon_1} \quad ; \quad Z = Bc_0 \ln \frac{1}{\varepsilon_1}. \quad (12)$$

This parametrization is motivated by the fact that asymptotically, for the Tardos code, A and B can be considered as constants. The relationship between the code length parametrization (12) and the fingerprinting rate is as follows. The rate is $R = (\log_q n)/\ell = (\ln n)/(Ac_0^2 \ln q \ln \varepsilon_1^{-1})$. Let $\eta = \Pr[\mathcal{L} \setminus \mathcal{C} \neq \emptyset]$, i.e. the probability that at least one innocent user ends up in the list \mathcal{L} . It can be shown (Lemma 6 in [25]) for $n \gg 1$, $c \ll n$ that $\eta \approx n\varepsilon_1$. Then $\ln \varepsilon_1^{-1} \approx \ln n - \ln \eta \approx \ln n$. Asymptotically the rate satisfies $R \sim 1/(Ac_0^2 \ln q)$.

Definition 2.1 *The variance of an innocent user's score, and the average and variance of the coalition score, are written as*

$$\tilde{\sigma}_{\text{inn}}^2 = \frac{1}{\ell} \sum_i \mathbb{E}(S_j^{(i)})^2 \quad \text{for arbitrary } j \notin \mathcal{C} \quad (13)$$

$$\tilde{\mu} = \frac{1}{\ell} \sum_i \mathbb{E}S_{\mathcal{C}}^{(i)} \quad (14)$$

$$\tilde{\sigma}^2 = \frac{1}{\ell} \sum_i \mathbb{E}(S_{\mathcal{C}}^{(i)})^2 - \tilde{\mu}^2. \quad (15)$$

Here \mathbb{E} stands for the expectation over all the probabilistic degrees of freedom: the biases $\mathbf{p}^{(i)}$, the code matrix X , and the coalition output \mathbf{y} . (The ‘tilde’ notation indicates that there is an average over positions.)

Remark: If assumption 3 holds (position symmetry, Section 2.1.3) then in Definition 2.1 the average over the positions is not necessary; in every position $\mathbb{E}[\dots]$ has the same value.

In this paper we introduce a rescaled version (β) of the threshold parameter B ,

$$B = \beta \tilde{\sigma}_{\text{inn}}. \quad (16)$$

It will turn out that it is more natural to use the quantity β than B .

Asymptotically the first and second moment completely determine the shape of the probability distribution of the score, for an innocent user as well as for the coalition score. (The distribution becomes Gaussian in accordance with the central limit theorem.) It was found [26] that the code length parameter (and hence the fingerprinting rate) then depend on $\tilde{\mu}$ and $\tilde{\sigma}_{\text{inn}}$ as follows,

$$A \sim \frac{2\tilde{\sigma}_{\text{inn}}^2}{\tilde{\mu}^2} \quad ; \quad R \sim \frac{\tilde{\mu}^2}{\tilde{\sigma}_{\text{inn}}^2} \cdot \frac{1}{2c_0^2 \ln q} \quad (17)$$

In the asymptotic saddlepoint, the tracer uses the bias distribution (2) with $\tau = 0$, while the coalition strategy is the *interleaving attack*, $\theta_{y|\mathbf{m}} = m_y/c$. In the asymptotic saddlepoint it holds [19] that $\tilde{\mu}^2/\tilde{\sigma}_{\text{inn}}^2 = q - 1$.

2.2 Computing expectations

Following previous work [21, 23, 25] we define (conditional) expectations as shown below. We omit the position index, and write x as shorthand for X_{ji} for a fixed innocent user $j \notin \mathcal{C}$.

$$\mathbb{E}_{\mathbf{p}}[r(\mathbf{p})] = \int_{\tau}^{1-(q-1)\tau} d^q p \delta(1 - |\mathbf{p}|) F(\mathbf{p}) r(\mathbf{p}) \quad (18)$$

$$\mathbb{E}_{x|\mathbf{p}}[r(x)] = \sum_{x \in \mathcal{Q}} p_x r(x) \quad (19)$$

$$\mathbb{E}_{\mathbf{m}|\mathbf{p}}[r(\mathbf{m})] = \sum_{\mathbf{m} \geq 0: |\mathbf{m}|=c} \binom{c}{\mathbf{m}} \mathbf{p}^{\mathbf{m}} r(\mathbf{m}) \quad (20)$$

$$\mathbb{E}_{y|\mathbf{m}}[r(y)] = \sum_{y \in \mathcal{Q}} \theta_{y|\mathbf{m}} r(y) \quad (21)$$

$$\mathbb{E}_{y|\mathbf{p}}[r(y)] = \mathbb{E}_{\mathbf{m}|\mathbf{p}} \mathbb{E}_{y|\mathbf{m}}[r(y)] = \sum_{y \in \mathcal{Q}} \left\{ \sum_{\mathbf{m} \geq 0: |\mathbf{m}|=c} \binom{c}{\mathbf{m}} \mathbf{p}^{\mathbf{m}} \theta_{y|\mathbf{m}} \right\} r(y) \quad (22)$$

$$\mathbb{E}_{\mathbf{m}}[r(\mathbf{m})] = \sum_{\mathbf{m} \geq 0: |\mathbf{m}|=c} \binom{c}{\mathbf{m}} \frac{B_{\tau}(\kappa \mathbf{1}_q + \mathbf{m})}{B_{\tau}(\kappa \mathbf{1}_q)} r(\mathbf{m}) \quad (23)$$

$$\mathbb{E}_{m_{\alpha}}[r(m_{\alpha})] = \sum_{b=0}^c P_1(b) r(b) = \sum_{b=0}^c \binom{c}{b} \frac{B_{\tau}(\kappa + b, [q-1]\kappa + c - b)}{B_{\tau}(\kappa, [q-1]\kappa)} r(b) \quad (24)$$

$$K_b = \mathbb{E}_{\mathbf{x}|b} \Psi_b(\mathbf{x}) = \sum_{\mathbf{x} \geq 0: |\mathbf{x}|=c-b} \binom{c-b}{\mathbf{x}} \frac{B(\kappa \mathbf{1}_{q-1} + \mathbf{x})}{B(\kappa \mathbf{1}_{q-1})} \Psi_b(\mathbf{x}) \quad (25)$$

Here $P_1(b)$ is a marginal probability for a single fixed symbol to have tally b . The quantity K_b is the probability, given that a certain symbol has tally b , for the colluders to output that symbol; i.e. for arbitrary fixed α , we have $K_b = \Pr[y = \alpha | m_{\alpha} = b]$. The sum rule $\sum_b P_1(b) K_b = 1/q$ holds [21], since the overall probability of outputting $y = \alpha$ is $1/q$.

2.3 Concentration inequalities

Lemma 2.2 (Bernstein's inequality [3]) *Let $a > 0$ be a constant. Let U_1, \dots, U_{ℓ} be independent zero-mean random variables, with $|U_i| \leq a$ for all i . Let $Z \geq 0$. Then*

$$\Pr \left[\sum_{i=1}^{\ell} U_i > Z \right] \leq \exp \left(- \frac{Z^2/2}{\sum_{i=1}^{\ell} \mathbb{E}[U_i^2] + aZ/3} \right). \quad (26)$$

Lemma 2.3 (Bennett's inequality [2]) *Let $b > 0$ be a constant. Let Y_1, \dots, Y_{ℓ} be independent zero-mean random variables, with $|Y_i| \leq b$ for all i . Let $s^2 = \frac{1}{\ell} \sum_{i=1}^{\ell} \mathbb{E}[Y_i^2]$. Let the function ξ be defined as*

$$\xi(v) = \int_0^v dx \ln(1+x) = (v+1) \ln(v+1) - v. \quad (27)$$

Let $T \geq 0$. Then

$$\Pr \left[\sum_{i=1}^{\ell} Y_i > T \right] \leq \exp \left(- \frac{\ell s^2}{b^2} \xi \left(\frac{b}{\ell s^2} T \right) \right). \quad (28)$$

Property 2.4 *The function ξ in Lemma 2.3 can be lower bounded as*

$$v > 0 \implies \xi(v) > v \ln \frac{v}{e}. \quad (29)$$

Proof: For $v > 0$ we have $\xi(v) = \int_0^v dx \ln(1+x) > \int_0^v dx \ln x = v \ln \frac{v}{e}$. \square

Lemma 2.5 (Weaker form of Bennett's inequality) *Let $b > 0$ be a constant. Let Y_1, \dots, Y_ℓ be independent zero-mean random variables, with $|Y_i| \leq b$ for all i . Let $s^2 = \frac{1}{\ell} \sum_{i=1}^{\ell} \mathbb{E}[Y_i^2]$. Let $T > 0$. Then*

$$\Pr \left[\sum_{i=1}^{\ell} Y_i > T \right] \leq \exp \left(-\frac{T}{b} \ln \frac{bT}{e\ell s^2} \right). \quad (30)$$

Proof: We substitute Property 2.4 in Lemma 2.3. This is allowed since the argument of ξ is positive. \square

3 Statistics of the innocent score and coalition score

We study the moments of the innocent score and coalition score in two cases: (i) Interleaving attack and arbitrary bias distribution; (ii) The bias distribution is the Dirichlet distribution with $\tau = 0$ and arbitrary concentration parameter κ ; the attack is arbitrary.

These two scenarios represent two different ways of departing from the asymptotic saddlepoint. In the first one, the bias distribution is varied. In the second one, not only the attack is varied but also a limited change of the bias distribution is allowed (κ).

The results of this section do not all contribute directly to the analysis of the sufficient code length in Section 5, but they are important in their own right since they elucidate how the score moments behave in a variety of circumstances.

3.1 General result for the moments

We investigate the first and second moment of an innocent user's score and of the coalition score. We begin with a general result for position-symmetric colluder strategies. Then we look more specifically at the Interleaving attack.

Lemma 3.1 *If the coalition is employing a position-symmetric strategy, then*

$$\tilde{\sigma}_{\text{inn}}^2 = -1 + \mathbb{E} \frac{1}{p_y} \quad (31)$$

$$\tilde{\mu} = -c + \mathbb{E} \frac{m_y}{p_y} \quad (32)$$

$$\tilde{\mu}^2 + \tilde{\sigma}^2 = \mathbb{E} \frac{(m_y - cp_y)^2}{p_y^2}. \quad (33)$$

Proof: We start from Definition 2.1. For $\tilde{\sigma}_{\text{inn}}^2$ we write $\mathbb{E}(S_j^{(i)})^2 = \mathbb{E}_{\mathbf{p}} \mathbb{E}_{y|\mathbf{p}} \mathbb{E}_{x|\mathbf{p}} (-1 + \delta_{xy}/p_y)^2 = \mathbb{E}_{\mathbf{p}} \mathbb{E}_{y|\mathbf{p}} \mathbb{E}_{x|\mathbf{p}} (1 - 2\delta_{xy}/p_y + \delta_{xy}/p_y^2) = 1 - 2\mathbb{E}_{\mathbf{p}} \mathbb{E}_{y|\mathbf{p}} 1 + \mathbb{E}_{\mathbf{p}} \mathbb{E}_{y|\mathbf{p}} 1/p_y$.

The results for $\tilde{\mu}$ and $\tilde{\sigma}$ follow directly from the fact that $S_C^{(i)} = (m_y/p_y - c) = (m_y - cp_y)/p_y$. \square

3.2 The case of the Interleaving attack

Lemma 3.2 *If the coalition is using the Interleaving attack, then*

$$\tilde{\mu}_{\text{Int}} = q - 1, \quad (\tilde{\sigma}_{\text{inn}}^2)_{\text{Int}} = q - 1, \quad \tilde{\mu}_{\text{Int}}^2 + \tilde{\sigma}_{\text{Int}}^2 = c(q - 1) - 3q + 2 + q \mathbb{E}_{\mathbf{p}} \frac{1}{p_\alpha}. \quad (34)$$

where $\alpha \in \mathcal{Q}$ is arbitrary.

Proof: For the Interleaving attack we have $\mathbb{E}[\dots] = \mathbb{E}_{\mathbf{p}} \mathbb{E}_{\mathbf{m}|\mathbf{p}} \sum_y (m_y/c) [\dots] = \sum_y \mathbb{E}_{\mathbf{p}} \mathbb{E}_{\mathbf{m}|\mathbf{p}} (\frac{m_y - cp_y}{c} + p_y) [\dots]$. We will make use of the binomial properties $\mathbb{E}_{\mathbf{m}|\mathbf{p}} m_\alpha = cp_\alpha$, $\mathbb{E}_{\mathbf{m}|\mathbf{p}} (m_\alpha - cp_\alpha)^2 = cp_\alpha(1 - p_\alpha)$ and $\mathbb{E}_{\mathbf{m}|\mathbf{p}} (m_\alpha - cp_\alpha)^3 = cp_\alpha(1 - p_\alpha)(1 - 2p_\alpha)$.

For $\tilde{\mu}$ this gives $\tilde{\mu} = \sum_y \mathbb{E}_{\mathbf{p}} \mathbb{E}_{\mathbf{m}|\mathbf{p}} \left[\frac{(m_y - cp_y)^2}{cp_y} + m_y - cp_y \right] = \mathbb{E}_{\mathbf{p}} \sum_y (1 - p_y) + 0 = q - 1$.

Furthermore $\tilde{\sigma}_{\text{inn}}^2 = -1 + \mathbb{E}_{\mathbf{p}} \sum_y \mathbb{E}_{\mathbf{m}|\mathbf{p}} \left[\frac{m_y}{cp_y} \right] = -1 + \mathbb{E}_{\mathbf{p}} \sum_y 1 = q - 1$.

Finally $\tilde{\mu}^2 + \tilde{\sigma}^2 = \sum_y \mathbb{E}_{\mathbf{p}} \mathbb{E}_{\mathbf{m}|\mathbf{p}} \left[\frac{(m_y - cp_y)^3}{cp_y^2} + \frac{(m_y - cp_y)^2}{p_y} \right] = \mathbb{E}_{\mathbf{p}} \sum_y \left[\frac{(1-p_y)(1-2p_y)}{p_y} + c(1-p_y) \right] = c(q-1) - 3q + 2 + \sum_y \mathbb{E}_{\mathbf{p}} \frac{1}{p_y}$. \square

Remark 1: Part of Lemma 3.2 ($\tilde{\mu}$ and $\tilde{\sigma}_{\text{inn}}$) was already done in [19]. We show the proof again because of our modified normalization of the score function.

Remark 2: The result for $\tilde{\mu}_{\text{Int}}$ and $(\tilde{\sigma}_{\text{inn}}^2)_{\text{Int}}$ does not depend on the bias distribution F , but $\tilde{\sigma}_{\text{Int}}$ does.

Remark 3: In the large- c limit the variance of the coalition score tends to be large, due to the $c(q-1)$ term as well as the expression $\mathbb{E}[1/p_\alpha]$ which blows up when τ becomes small.

3.3 Taking the Dirichlet distribution with cutoff $\tau = 0$

Lemma 3.3 *Let $\tau = 0$. Let the coalition use a strategy that is colluder-symmetric and position-symmetric. Then the quantities $\tilde{\mu}$ and $\tilde{\sigma}_{\text{inn}}$ can be written as*

$$\tilde{\mu}_{\tau=0} = -c + (q\kappa + c - 1) \mathbb{E}_{\mathbf{m}} \sum_{y \in \mathcal{Q}} \theta_{y|\mathbf{m}} \left[1 + \frac{1 - \kappa}{\kappa + m_y - 1} \right] \quad (35)$$

$$(\tilde{\sigma}_{\text{inn}}^2)_{\tau=0} = -1 + (q\kappa + c - 1) \mathbb{E}_{\mathbf{m}} \sum_{y \in \mathcal{Q}} \theta_{y|\mathbf{m}} \frac{1}{\kappa + m_y - 1} \quad (36)$$

Furthermore, if the colluder strategy is also symbol-symmetric, then

$$\tilde{\mu}_{\tau=0} = -c + (q\kappa + c - 1) q \sum_{b=1}^c P_1(b) K_b \frac{b}{\kappa + b - 1}, \quad (37)$$

$$(\tilde{\sigma}_{\text{inn}}^2)_{\tau=0} = -1 + (q\kappa + c - 1) q \sum_{b=1}^c P_1(b) K_b \frac{1}{\kappa + b - 1}. \quad (38)$$

Proof: We start from the expressions $\tilde{\mu} = -c + \mathbb{E}[m_y/p_y]$ and $\tilde{\sigma}_{\text{inn}}^2 = -1 + \mathbb{E}[1/p_y]$. For any function $J(m_y)$ we can write $\mathbb{E}[J(m_y)/p_y] = \mathbb{E}_{\mathbf{p}} \sum_{\mathbf{m}} \binom{c}{\mathbf{m}} \mathbf{p}^{\mathbf{m}} \sum_y \theta_{y|\mathbf{m}} J(m_y)/p_y = \sum_{\mathbf{m}} \binom{c}{\mathbf{m}} \sum_y \theta_{y|\mathbf{m}} J(m_y) \mathbb{E}_{\mathbf{p}} \mathbf{p}^{\mathbf{m}} / p_y$. For $\tau = 0$ we have

$$\mathbb{E}_{\mathbf{p}} \frac{\mathbf{p}^{\mathbf{m}}}{p_y} = \frac{B(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)}{B(\kappa \mathbf{1}_q)} = \frac{q\kappa + c - 1}{\kappa + m_y - 1} \cdot \frac{B(\kappa \mathbf{1}_q + \mathbf{m})}{B(\kappa \mathbf{1}_q)} = \frac{q\kappa + c - 1}{\kappa + m_y - 1} \mathbb{E}_{\mathbf{p}} \mathbf{p}^{\mathbf{m}}. \quad (39)$$

Setting $J(m_y) = m_y$ for $\tilde{\mu}$, and $J(m_y) = 1$ for $\tilde{\sigma}_{\text{inn}}^2$, yields (35) and (36). The final step is to notice that $\mathbb{E}[J(m_y)/p_y] = \mathbb{E}_{\mathbf{m}} \mathbb{E}_{y|\mathbf{m}} \left[\frac{q\kappa + c - 1}{\kappa + m_y - 1} J(m_y) \right]$ which can be rewritten as $q \sum_b P_1(b) K_b \frac{q\kappa + c - 1}{\kappa + b - 1} J(b)$ if the strategy is symbol-symmetric. \square

Theorem 3.4 *Let $c \gg 1$ and $\kappa \in (0, 1)$. Let the coalition use a strategy that is colluder-symmetric and position-symmetric. Then both quantities $\tilde{\mu}$ and $\tilde{\sigma}_{\text{inn}}$ are maximized by the Minority Voting attack and minimized by the Majority Voting attack.*

Proof: For $c \gg 1$ we can use the $\tau = 0$ approximation for $\tilde{\mu}$ and $\tilde{\sigma}_{\text{inn}}$, i.e. Lemma 3.3. In (35) and (36), the $\theta_{y|\mathbf{m}}$ in the y -summation multiplies a *decreasing* function of m_y . Hence the summand is maximized by outputting a symbol y with tally m_y as small as possible (but nonzero because of the Marking Assumption), and, vice versa, minimized by outputting the symbol with the largest tally. \square

Theorem 3.4 gives insight into the trade-offs that the colluders have to deal with. They want to minimize $\tilde{\mu}$ and to maximize $\tilde{\sigma}_{\text{inn}}$, since this leads to high error rates. However, the strategy that optimizes $\tilde{\mu}$ for them is the worst possible strategy regarding $\tilde{\sigma}_{\text{inn}}$, and vice versa. The Interleaving attack at the saddlepoint is ‘in the middle’ between Minority Voting and Majority Voting.

Lemma 3.5 *Let $\tau = 0$. Let the coalition use a strategy that is colluder-symmetric and position-symmetric. Then $\tilde{\mu}$ and $\tilde{\sigma}_{\text{inn}}$ can be bounded as*

$$\frac{c\kappa(q-1)}{c-1+\kappa} \leq \tilde{\mu}_{\tau=0} \leq c\left(\frac{1}{\kappa} - 1\right) + q - \frac{1}{\kappa} \quad (40)$$

$$\frac{\kappa(q-1)}{c-1+\kappa} \leq (\tilde{\sigma}_{\text{inn}}^2)_{\tau=0} \leq \frac{c}{\kappa} + q - 1 - \frac{1}{\kappa}. \quad (41)$$

Proof: For $m_y \in \{1, \dots, c\}$ we have $\frac{1}{\kappa+c-1} \leq \frac{1}{\kappa+m_y-1} \leq \frac{1}{\kappa}$. We substitute these inequalities into (35) and (36). Finally we use $\sum_y \theta_{y|\mathbf{m}} = 1$. \square

Remark: It is possible to obtain a tighter upper bound by treating the $m_y = c$ term separately in (35),(36), since then $\theta_{y|\mathbf{m}} = 1$. However, the improvement of the tightness is minimal.

4 Bounding the error probabilities

We use Bernstein's inequality and Bennett's inequality to upper bound the False Positive and False Negative error probability, respectively.

4.1 Bounding the False Positive probability

Theorem 4.1 *Let $q \geq 2$. Let the coalition use any attack strategy. Then the False Positive probability for a fixed innocent user can be bounded as*

$$P_{\text{FP}} \leq \exp \left[(\ln \varepsilon_1) \frac{\beta^2}{2A} \left(1 + \frac{\beta}{3Ac_0\tau\tilde{\sigma}_{\text{inn}}} \right)^{-1} \right]. \quad (42)$$

Proof: For *any* coalition strategy, even one that breaks the position symmetry, the single-position scores $S_j^{(i)}$ for the *innocent* user are mutually independent [22]. Hence we are allowed to use Bernstein's inequality. In Lemma 2.2 we set $U_i = S_j^{(i)}$ for the innocent user. This is allowed since $S_j^{(i)}$ has zero expectation value. We have

$$|U_i| \leq \max \left\{ \frac{1}{p_{\min}} - 1, |-1| \right\} = \max \left\{ \frac{1}{\tau} - 1, 1 \right\} = \frac{1}{\tau} - 1 < \frac{1}{\tau}. \quad (43)$$

In the last equality we used $\tau \leq 1/q$ (see Section 2.1.2). Thus we are allowed to set $a = 1/\tau$ in Lemma 2.2. Furthermore we note that by definition $\mathbb{E}[U_i^2] = \tilde{\sigma}_{\text{inn}}^2$ for all i . Lemma 2.2 then gives

$$\Pr[S_j > Z] \leq \exp \left(\frac{-Z^2/2}{\ell\tilde{\sigma}_{\text{inn}}^2 + aZ/3} \right) = \exp \left(\frac{-Z^2}{2\ell\tilde{\sigma}_{\text{inn}}^2} \cdot \frac{1}{1 + aZ/(3\ell\tilde{\sigma}_{\text{inn}}^2)} \right). \quad (44)$$

Substituting $a = 1/\tau$, $\ell = Ac_0^2 \ln \frac{1}{\varepsilon_1}$ and $Z = \beta\tilde{\sigma}_{\text{inn}}c_0 \ln \frac{1}{\varepsilon_1}$ finishes the proof. \square

Remark: In (42) we see that the bound on P_{FP} is a *decreasing* function of the product $c_0\tau$. Hence it is advantageous to set τ such that $c_0\tau \gg 1$.

Corollary 4.2 *Let $q \geq 2$ and $\tau \leq 1/2$. Let the coalition use any attack strategy. Then it holds that*

$$A \leq \frac{1}{2}\beta^2 - \frac{\beta}{3c_0\tau\tilde{\sigma}_{\text{inn}}} \implies P_{\text{FP}} \leq \varepsilon_1. \quad (45)$$

Proof: Follows directly from Theorem 4.1. \square

4.2 Bounding the False Negative probability

Theorem 4.3 *Let $q \geq 2$. Let the coalition employ a position-symmetric strategy. Let $\tilde{\mu}Ac_0 - \tilde{\sigma}_{\text{inn}}\beta c > 0$. Let τ satisfy*

$$\tau \leq c/(c + \tilde{\mu}). \quad (46)$$

Then the False Negative probability can be bounded as

$$P_{\text{FN}} \leq \exp \left[(\ln \varepsilon_1) \frac{c_0 \tau}{c} [\tilde{\mu}Ac_0 - \tilde{\sigma}_{\text{inn}}\beta c] \ln \frac{\tilde{\mu}Ac_0 - \tilde{\sigma}_{\text{inn}}\beta c}{e(\tilde{\sigma}^2/c)Ac_0\tau} \right]. \quad (47)$$

Proof: We start from

$$P_{\text{FN}} = \Pr[\forall_{j \in C} S_j < Z] < \Pr[S_C < cZ] = \Pr[\ell\tilde{\mu} - S_C > \ell\tilde{\mu} - cZ] = \Pr\left[\sum_{i=1}^{\ell} (\tilde{\mu} - S_C^{(i)}) > \ell\tilde{\mu} - cZ\right]. \quad (48)$$

Because of the assumption that the collusion attack is position-symmetric, the random variables $S_C^{(i)}$ are mutually independent. We are then allowed to use Bennett's inequality (we take the weaker form, Lemma 2.5), which we do with the following parameters: $Y_i = \tilde{\mu} - S_C^{(i)}$; $T = \ell\tilde{\mu} - cZ = (\tilde{\mu}Ac_0 - \tilde{\sigma}_{\text{inn}}\beta c)c_0 \ln \frac{1}{\varepsilon_1}$; $s^2 = \tilde{\sigma}^2$; $b = c/\tau$. The choice for b follows from

$$|Y_i| = |S_C^{(i)} - \tilde{\mu}| \leq \max\{c(\frac{1}{\tau} - 1) - \tilde{\mu}, \tilde{\mu} + c\} \leq \max\{\frac{c}{\tau}, \tilde{\mu} + c\} = \frac{c}{\tau}, \quad (49)$$

where the last equality is a consequence of the assumption (46). We can see that the T is positive from the assumption $\tilde{\mu}Ac_0 - \tilde{\sigma}_{\text{inn}}\beta c > 0$. \square

Notice that at $c \gg c_0$ Theorem 4.3 no longer applies, because the condition $\tilde{\mu}Ac_0 - \tilde{\sigma}_{\text{inn}}\beta c > 0$ cannot be satisfied. In practical terms this means that for $c > c_0$ the FN probability is no longer under control, and the colluders may evade detection with high probability.

Theorem 4.4 *Let $q \geq 2$. Let the coalition employ a position-symmetric strategy. Let $2 \leq c \leq c_0$. Let $\tilde{\mu}A - \tilde{\sigma}_{\text{inn}}\beta > 0$. Let $\tau \leq 2/(2 + \tilde{\mu})$. Then the False Negative probability can be bounded as*

$$P_{\text{FN}} \leq \exp \left[(\ln \varepsilon_1) c_0 \tau [\tilde{\mu}A - \tilde{\sigma}_{\text{inn}}\beta] \ln \frac{\tilde{\mu}A - \tilde{\sigma}_{\text{inn}}\beta}{e(\tilde{\sigma}^2/c_0)A\tau} \right]. \quad (50)$$

Proof: We start from Theorem 4.3. Due to the conditions $c \leq c_0$ and $\tilde{\mu}A - \tilde{\sigma}_{\text{inn}}\beta > 0$, the condition $\tilde{\mu}Ac_0 - \tilde{\sigma}_{\text{inn}}\beta c > 0$ in Theorem 4.3 holds. Due to $c \geq 2$ and $\tau < 2/(2 + \tilde{\mu})$, the condition (46) holds. Since all the conditions are satisfied we are allowed to apply Theorem 4.3. Finally we make use of the fact that the expression (47) is an *increasing* function of c for $c \leq c_0$. \square

Corollary 4.5 *Let $q \geq 2$. Let the coalition employ a position-symmetric strategy. Let $2 \leq c \leq c_0$. Let $\tilde{\mu}A - \tilde{\sigma}_{\text{inn}}\beta > 0$. Let $\tau \leq 2/(2 + \tilde{\mu})$. Then it holds that*

$$c_0 \tau [\tilde{\mu}A - \tilde{\sigma}_{\text{inn}}\beta] \ln \frac{\tilde{\mu}A - \tilde{\sigma}_{\text{inn}}\beta}{e(\tilde{\sigma}^2/c_0)A\tau} \geq \frac{\ln \varepsilon_2}{\ln \varepsilon_1} \implies P_{\text{FN}} \leq \varepsilon_2. \quad (51)$$

Proof: Follows directly from Theorem 4.4. \square

5 Asymptotics of the sufficient code length

The main aim of this paper is to determine the performance of the score system (7) at large but finite c_0 . The performance at ' $c_0 = \infty$ ' is known: the saddlepoint is given by the Interleaving attack, combined with the $\kappa = \frac{1}{2}$ Dirichlet distribution (with $\tau = 0$) as the bias distribution; in this saddlepoint the rate of the score system is equal to capacity. What we want to know is how the fingerprinting rate approaches capacity, and how to optimally choose the cutoff τ as a function of c_0 .

5.1 Sufficient code length

We aim for an analysis in the (unknown!) large-but-finite- c_0 saddlepoint.

- The saddlepoint (‘SP’) of the mutual information minimax game [10] is close to the asymptotic saddlepoint. The unknown strategy θ^{SP} is close to Interleaving. The unknown bias distribution $F^{\text{SP}}(\mathbf{p})$ is some *discrete* distribution close to the Dirichlet distribution. We approximate F by the continuous Dirichlet distribution with cutoff τ because this is the only available constructive approach that we know of.
- Hence, we analyze the tracing system consisting of the bias distribution (2) and the score system (7), when pitted against an unknown attack close to Interleaving. Our starting point will be the ‘sufficient’ conditions given by Corollaries 4.2 and 4.5. We know that $\tilde{\mu}^{\text{SP}} = q - 1 - \Delta\tilde{\mu}$ and $(\tilde{\sigma}_{\text{inn}}^2)^{\text{SP}} = q - 1 + \Delta\tilde{\sigma}_{\text{inn}}^2$, and we have to carefully deal with the corrections $\Delta\tilde{\mu}$ and $\Delta\tilde{\sigma}_{\text{inn}}^2$. On the other hand, the $\tilde{\sigma}$ appears only in the logarithm in (51) and hence any corrections with respect to Lemma 3.2 can be neglected.

Corollary 4.2 and the condition $\tilde{\mu}A - \tilde{\sigma}_{\text{inn}}\beta > 0$ together define an interval for the sufficient code length parameter ‘ A_{suff} ’,

$$A_{\text{suff}} \in \left(\frac{\tilde{\sigma}_{\text{inn}}}{\tilde{\mu}}\beta, \frac{1}{2}\beta^2 - \frac{\beta}{3c_0\tau\tilde{\sigma}_{\text{inn}}} \right]. \quad (52)$$

This interval exists only if

$$\beta > 2\frac{\tilde{\sigma}_{\text{inn}}}{\tilde{\mu}} + \frac{2}{3c_0\tau\tilde{\sigma}_{\text{inn}}}, \quad (53)$$

which yields

$$A_{\text{suff}} > \frac{2\tilde{\sigma}_{\text{inn}}^2}{\tilde{\mu}^2} + \frac{2}{3c_0\tau\tilde{\mu}}. \quad (54)$$

We must try to bring β and A as close as possible to the bounds (53,54) while still satisfying the condition in the left hand side of (51). We introduce the following shorthand notation,

$$\frac{\tilde{\sigma}_{\text{inn}}}{\tilde{\mu}} = \frac{1}{\sqrt{q-1}}(1+w), \quad \delta = \tilde{\mu}A - \tilde{\sigma}_{\text{inn}}\beta, \quad \frac{\tilde{\sigma}^2}{c} = q - 1 + r \quad (55)$$

where $w \ll 1$, $\delta \ll 1$, $r \ll 1$. The w will be studied in the next section. The δ we will solve approximately. The fact that r is small follows from Lemma 3.2. The expression $\mathbb{E}[1/p_\alpha]$ in (34) is of order $\tau^{\kappa-1}$; this leads to a contribution to $\tilde{\sigma}^2/c$ of order $\tau^\kappa/(c_0\tau)$, which is negligible compared to $(q-1)$ since $c_0\tau \gg 1$ (see Section 4.1).

Theorem 5.1 *Let $c_0\tau \gg 1$ and $c_0\tau^2 \ll 1$. Let the attackers employ a position-symmetric strategy. Let $2 \leq c \leq c_0$. Then the following combination of a code length parameter A and threshold parameter β is sufficient to achieve $P_{\text{FP}} \leq \varepsilon_1$ and $P_{\text{FN}} \leq \varepsilon_2$.*

$$\beta_{\text{suff}} = \frac{2}{\sqrt{q-1}} \left[1 + w + \frac{1}{3c_0\tau} + \mathcal{O}\left(\frac{w}{c_0\tau}\right) \right] \quad (56)$$

$$A_{\text{suff}} = \frac{2}{q-1} \left[1 + 2w + \frac{1}{3c_0\tau} + \frac{\ln \varepsilon_2 / \ln \varepsilon_1}{2c_0\tau \ln \frac{1}{c_0\tau^2}} + \mathcal{O}(w^2) + \mathcal{O}\left(\frac{w}{c_0\tau}\right) \right]. \quad (57)$$

Proof: The condition in (51) can be written compactly as

$$c_0\tau\delta \ln \frac{\delta}{e(q-1+r)A\tau} \geq \frac{\ln \varepsilon_2}{\ln \varepsilon_1}. \quad (58)$$

Taking the equal sign and solving for δ gives (we denote the solution as δ_0)

$$\begin{aligned}
\delta_0 &= \frac{\ln \varepsilon_2}{\ln \varepsilon_1} \cdot \frac{1}{c_0 \tau} \cdot \frac{1}{\ln \left[\frac{1}{e^{(q-1+r)A\tau}} \cdot \frac{\ln \varepsilon_2}{\ln \varepsilon_1} \cdot \frac{1}{c_0 \tau} \ln \frac{\delta_0}{e^{(q-1+r)A\tau}} \right]} \\
&= \frac{\ln \varepsilon_2}{\ln \varepsilon_1} \cdot \frac{1}{c_0 \tau} \cdot \frac{1}{\ln \left[\frac{1}{c_0 \tau^2} \right] + \ln \left[\frac{1}{e^{(q-1)A}} \frac{\ln \varepsilon_2}{\ln \varepsilon_1} \right] - \mathcal{O}(r) + \mathcal{O}(\ln \ln \frac{\delta_0}{\tau})} \\
&= \frac{\ln \varepsilon_2}{\ln \varepsilon_1} \cdot \frac{1}{c_0 \tau \ln \frac{1}{c_0 \tau^2}} \left[1 - \mathcal{O} \left(\frac{\ln \ln \frac{1}{c_0 \tau^2}}{\ln \frac{1}{c_0 \tau^2}} \right) \right] \\
&< \frac{\ln \varepsilon_2}{\ln \varepsilon_1} \cdot \frac{1}{c_0 \tau \ln \frac{1}{c_0 \tau^2}}. \tag{59}
\end{aligned}$$

We take $\delta = \frac{\ln \varepsilon_2}{\ln \varepsilon_1} \cdot \frac{1}{c_0 \tau \ln \frac{1}{c_0 \tau^2}}$ (last line of (59)), since it is a compact analytical expression that satisfies (58). We can now find the sufficient A and β . We write $\beta_{\text{suff}} = 2 \frac{\tilde{\sigma}_{\text{inn}}}{\tilde{\mu}} + \frac{2}{3c_0 \tau \tilde{\sigma}_{\text{inn}}} + \lambda$, with λ arbitrarily close to zero. Solving A from β and δ gives

$$\begin{aligned}
A_{\text{suff}} &= \beta_{\text{suff}} \frac{\tilde{\sigma}_{\text{inn}}}{\tilde{\mu}} + \frac{\delta}{\tilde{\mu}} \\
&= \frac{2}{q-1} \left[1 + 2w + \frac{1}{3c_0 \tau} + \frac{\ln \varepsilon_2 / \ln \varepsilon_1}{2c_0 \tau \ln \frac{1}{c_0 \tau^2}} + \mathcal{O}(w^2) + \mathcal{O}(\lambda) + \mathcal{O}\left(\frac{w}{c_0 \tau}\right) \right] \tag{60}
\end{aligned}$$

where we have used that $\Delta \tilde{\mu}$ and $\Delta \tilde{\sigma}_{\text{inn}}$ are of order w . Finally we note that λ is much smaller than the other high-order correction terms. \square

Note that the condition $c_0 \tau^2 \ll 1$ is required in the above proof in order to make sure that the argument of the logarithm is well-behaved, i.e. larger than 1. Hence when choosing τ we have to satisfy

- Condition 1:** $c_0 \tau \gg 1$.
- Condition 2:** $c_0 \tau^2 \ll 1$.

One way of satisfying these conditions is to set

$$\tau \propto c_0^{-\gamma} \quad \text{with} \quad \gamma \in \left(\frac{1}{2}, 1\right). \tag{61}$$

5.2 Optimization of the cutoff τ as a function of c_0

Lemma 5.2 (adapted from [19]) *Let $\Delta \theta_{y|\mathbf{m}} = \theta_{y|\mathbf{m}}^{\text{SP}} - m_y/c$. The first order and second order correction terms to $\tilde{\mu}$ and $\tilde{\sigma}_{\text{inn}}^2$ in the vicinity of the saddle point are given by*

$$\begin{aligned}
\tilde{\mu}_{(1)} &= \sum_{\mathbf{m}} \binom{c}{\mathbf{m}} \sum_{y \in \mathcal{Q}} \Delta \theta_{y|\mathbf{m}} m_y \frac{B(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)}{B(\kappa \mathbf{1}_q)} = \mathbb{E}_{\mathbf{m}} \sum_{y \in \mathcal{Q}} \Delta \theta_{y|\mathbf{m}} (1 - \kappa) \frac{c + q\kappa - 1}{m_y - (1 - \kappa)} \\
[\tilde{\sigma}_{\text{inn}}^2]_{(1)} &= \sum_{\mathbf{m}} \binom{c}{\mathbf{m}} \sum_{y \in \mathcal{Q}} \Delta \theta_{y|\mathbf{m}} \frac{B(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)}{B(\kappa \mathbf{1}_q)} = \mathbb{E}_{\mathbf{m}} \sum_{y \in \mathcal{Q}} \Delta \theta_{y|\mathbf{m}} \frac{c + q\kappa - 1}{m_y - (1 - \kappa)} \\
\tilde{\mu}_{(2)} &= \sum_{\mathbf{m}} \binom{c}{\mathbf{m}} \sum_{y \in \mathcal{Q}} \Delta \theta_{y|\mathbf{m}} m_y \left[\frac{B(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)}{B(\kappa \mathbf{1}_q)} - \frac{B_{\tau}(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)}{B_{\tau}(\kappa \mathbf{1}_q)} \right] \\
[\tilde{\sigma}_{\text{inn}}^2]_{(2)} &= \sum_{\mathbf{m}} \binom{c}{\mathbf{m}} \sum_{y \in \mathcal{Q}} \Delta \theta_{y|\mathbf{m}} \left[\frac{B(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)}{B(\kappa \mathbf{1}_q)} - \frac{B_{\tau}(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)}{B_{\tau}(\kappa \mathbf{1}_q)} \right]. \tag{62}
\end{aligned}$$

The first order correction to $\tilde{\mu}^2/\tilde{\sigma}_{\text{inn}}^2$ is zero, because of the saddlepoint. The second order correction to $\tilde{\mu}^2/\tilde{\sigma}_{\text{inn}}^2$ is given by

$$\left[\frac{\tilde{\mu}^2}{\tilde{\sigma}_{\text{inn}}^2} \right]_{(2)} = 2\tilde{\mu}_{(2)} - [\tilde{\sigma}_{\text{inn}}^2]_{(2)} + \frac{1}{q-1} (\tilde{\mu}_{(1)} - [\tilde{\sigma}_{\text{inn}}^2]_{(1)})^2 \quad (63)$$

$$= - \sum_{\mathbf{m}} \binom{c}{\mathbf{m}} \sum_{y \in \mathcal{Q}} \Delta\theta_{y|\mathbf{m}} (2m_y - 1) \frac{B_\tau(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)}{B_\tau(\kappa \mathbf{1}_q)} + \frac{\kappa^2}{q-1} ([\tilde{\sigma}_{\text{inn}}^2]_{(1)})^2. \quad (64)$$

Proof: Eqs. (62) and (63) are a slight adaptation of the saddlepoint formulas in [19], where we have substituted the saddlepoint values $\tilde{\mu} = q - 1$ and $\tilde{\sigma}_{\text{inn}}^2 = q - 1$. Note again that we have normalized the score function different from [19] by a factor $\sqrt{q-1}$. Eq. (64) follows from (63) by using (62). \square

Proposition 5.3 *The correction w is negligible compared to $\frac{1}{c_0\tau}$.*

Argumentation: The w is proportional to (63) or, differently expressed, (64). In (64) we have the $([\tilde{\sigma}_{\text{inn}}^2]_{(1)})^2$ term which is of order $(\Delta\theta)^2$. The order of magnitude of the $\sum_{\mathbf{m}}$ contribution is more difficult to determine, because the incomplete Dirichlet integral $B_\tau(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)$ is difficult to bound²; however, no matter how $B_\tau(\kappa \mathbf{1}_q + \mathbf{m} - \mathbf{e}_y)$ is behaved, the $\sum_{\mathbf{m}}$ contribution is at most of order $\Delta\theta$. Huang and Moulin [10] conjectured that $\Delta\theta = \mathcal{O}(\frac{1}{c})$, and this turned out to be consistent with their asymptotic saddlepoint analysis. If their conjecture is true, we have $w \propto \frac{1}{c_0} \ll \frac{1}{c_0\tau}$. Even if their conjecture is not true and $\Delta\theta$ scales as, for instance, $1/\sqrt{c}$, then $w \propto 1/\sqrt{c_0} \ll \frac{1}{c_0\tau}$, i.e. w is still negligible. (The latter holds because τ scales as $c_0^{-\gamma}$ with $\gamma > \frac{1}{2}$.) \square

The consequences of Proposition 5.3 are

- The optimal choice for the cutoff is to set

$$\gamma_{\text{opt}} = \frac{1}{2} + \nu \quad (65)$$

where ν denotes a very small positive number.

- The sufficient code length is then given by

$$A_{\text{suff}} = \frac{2}{q-1} [1 + \mathcal{O}(c_0^{-1/2+\nu})]. \quad (66)$$

Note that the correction term is smaller than the $\mathcal{O}(c_0^{-1/3})$ that was found [12] for Tardos's score function at $q = 2$.

6 Discussion

We have studied a q -ary bias-based collusion-resistant scheme where the score function (7) of Oosterwijk et al. [19] is used in combination with the Dirichlet distribution with a cutoff. We have used Bernstein's inequality and Bennett's inequality to upper bound the error rates. For large c_0 this leads to a sufficient code length as specified in Theorem 5.1.

Then we adopted a conjecture (based on a conjecture by Huang and Moulin) that $\Delta\theta$, the difference in strategy between the finite- c and infinite- c saddlepoint, is of order $\mathcal{O}(1/\sqrt{c})$. This leads to an optimal cutoff choice $\tau = 1/(\lambda c_0^{1/2+\nu})$, where $\lambda > 0$ is a constant, and ν is a very small positive constant. The sufficient code length is then

$$\ell_{\text{suff}} = \frac{2}{q-1} \left[1 + \lambda c_0^{-\frac{1}{2}+\nu} \left(\frac{1}{3} + \frac{1}{4} \frac{\ln \varepsilon_2}{\ln \varepsilon_1} \frac{1}{\ln(c_0^\nu \lambda)} \right) + \dots \right] c_0^2 \ln \varepsilon_1^{-1}, \quad (67)$$

²The correction to the normalization factor is known. In [25] it was found that $B_\tau(\kappa \mathbf{1}_q) = B(\kappa \mathbf{1}_q)[1 - \mathcal{O}(\tau^\kappa)]$.

and the corresponding accusation threshold is

$$Z = 2[1 + \frac{1}{3}\lambda c_0^{-\frac{1}{2}+\nu} + \dots]c_0 \ln \varepsilon_1^{-1}. \quad (68)$$

From previous work on provable bounds for bias-based codes it is clear that the bounds obtained from concentration inequalities (Markov, Bernstein, Bennett) are not tight.

As topics for future work we mention

- (i) Obtaining tighter bounds. The CSE method [21] or similar techniques may yield more precise information about the error rates.
- (ii) Studying the performance of the score function (7) further away from the asymptotic saddlepoint. This would require locating (by numerical techniques) the saddlepoint for large but finite c .
- (iii) Applying the analysis in this paper in the context of *dynamic* traitor tracing, similar to the work in [14].

Acknowledgements

We thank Benne de Weger, Jeroen Doumen and Thijs Laarhoven for useful discussions. Part of this work was supported by STW (project 10518).

References

- [1] E. Amiri and G. Tardos. High rate fingerprinting codes and the fingerprinting capacity. In *Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 336–345, 2009.
- [2] G. Bennett. Probability Inequalities for the Sum of Independent Random Variables. *Journal of the American Statistical Association*, 57(297):33–45, 1962.
- [3] S.N. Bernstein. *Theory of Probability*. 1927.
- [4] O. Blayer and T. Tassa. Improved versions of Tardos’ fingerprinting scheme. *Designs, Codes and Cryptography*, 48(1):79–103, 2008.
- [5] D. Boesten and B. Škorić. Asymptotic fingerprinting capacity for non-binary alphabets. In *Information Hiding 2011*, volume 6958 of *LNCS*, pages 1–13. Springer, 2011.
- [6] A. Charpentier, C. Fontaine, T. Furon, and I.J. Cox. An asymmetric fingerprinting scheme based on Tardos codes. In *Information Hiding*, volume 6958 of *LNCS*, pages 43–58. Springer, 2011.
- [7] A. Charpentier, F. Xie, C. Fontaine, and T. Furon. Expectation maximization decoding of Tardos probabilistic fingerprinting code. In *Media Forensics and Security*, volume 7254 of *SPIE Proceedings*, page 72540, 2009.
- [8] T. Furon, A. Guyader, and F. C erou. On the design and optimization of Tardos probabilistic fingerprinting codes. In *Information Hiding*, volume 5284 of *Lecture Notes in Computer Science*, pages 341–356. Springer, 2008.
- [9] T. Furon, L. P erez-Freire, A. Guyader, and F. C erou. Estimating the minimal length of Tardos code. In *Information Hiding*, volume 5806 of *LNCS*, pages 176–190, 2009.
- [10] Y.-W. Huang and P. Moulin. On fingerprinting capacity games for arbitrary alphabets and their asymptotics. In *IEEE International Symposium on Information Theory (ISIT) 2012*, pages 2571–2575, July 2012.
- [11] Y.W. Huang and P. Moulin. Capacity-achieving fingerprint decoding. In *IEEE Workshop on Information Forensics and Security (WIFS)*, pages 51–55, 2009.

- [12] T. Laarhoven and B.M.M. de Weger. Optimal symmetric Tardos traitor tracing schemes, 2011. <http://arxiv.org/abs/1107.3441>.
- [13] T. Laarhoven and B.M.M. de Weger. Discrete Distributions in the Tardos Scheme, Revisited. In *Information Hiding & Multimedia Security 2013*, 2013.
- [14] T. Laarhoven, J. Doumen, P. Roelse, B. Škorić, and B. de Weger. Dynamic Tardos traitor tracing schemes. *IEEE Transactions on Information Theory*, 59(7):4230–4242.
- [15] P. Meerwald and T. Furon. Towards joint Tardos decoding: the ‘Don Quixote’ algorithm. In *Information Hiding*, volume 6958 of *LNCS*, pages 28–42. Springer, 2011.
- [16] P. Moulin. Universal fingerprinting: Capacity and random-coding exponents. 2008. <http://arxiv.org/abs/0801.3837>.
- [17] K. Nuida. Short collusion-secure fingerprint codes against three pirates. In *Information Hiding*, volume 6387 of *LNCS*, pages 86–102. Springer, 2010.
- [18] K. Nuida, S. Fujitsu, M. Hagiwara, T. Kitagawa, H. Watanabe, K. Ogawa, and H. Imai. An improvement of discrete Tardos fingerprinting codes. *Des. Codes Cryptography*, 52(3):339–362, 2009.
- [19] J.-J. Oosterwijk, B. Škorić, and J. Doumen. A capacity-achieving simple decoder for bias-based traitor tracing schemes, 2013. <http://eprint.iacr.org/2013/389>.
- [20] J.-J. Oosterwijk, B. Škorić, and J. Doumen. Optimal suspicion functions for Tardos traitor tracing schemes. In *Information Hiding & Multimedia Security 2013*, 2013.
- [21] A. Simone and B. Škorić. Accusation probabilities in Tardos codes: beyond the Gaussian approximation. *Designs, Codes and Cryptography*, 63(3):379–412, 2012.
- [22] G. Tardos. Optimal probabilistic fingerprint codes. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 116–125, 2003.
- [23] B. Škorić, S. Katzenbeisser, and M.U. Celik. Symmetric Tardos fingerprinting codes for arbitrary alphabet sizes. *Designs, Codes and Cryptography*, 46(2):137–166, 2008.
- [24] B. Škorić, S. Katzenbeisser, H.G. Schaathun, and M.U. Celik. Tardos fingerprinting codes in the Combined Digit Model. *IEEE Transactions on Information Forensics and Security*, 6(3):906–919, 2011.
- [25] B. Škorić and J.-J. Oosterwijk. Binary and q-ary Tardos codes, revisited, 2012. To appear in *Designs, Codes, and Cryptography*. <http://eprint.iacr.org/2012/249>.
- [26] B. Škorić, T.U. Vladimirova, M.U. Celik, and J.C. Talstra. Tardos fingerprinting is better than we thought. *IEEE Transactions on Information Theory*, 54(8):3663–3676, 2008.
- [27] F. Xie, T. Furon, and C. Fontaine. On-off keying modulation and Tardos fingerprinting. In *Proc. 10th Workshop on Multimedia & Security (MM&Sec)*, pages 101–106. ACM, 2008.