Accusation probabilities in Tardos codes: the Gaussian approximation is better than we thought

Antonino Simone and Boris Škorić

Abstract

We study the probability distribution of user accusations in the q-ary Tardos fingerprinting system under the Marking Assumption, in the restricted digit model. In particular, we look at the applicability of the so-called Gaussian approximation, which states that accusation probabilities tend to the normal distribution when the fingerprinting code is long. We introduce a novel parametrization of the attack strategy which enables a significant speedup of numerical evaluations. We set up a method, based on power series expansions, to systematically compute the probability of accusing innocent users. The 'small parameter' in the power series is $1/m$, where m is the code length. We use our method to semi-analytically study the performance of the Tardos code against majority voting and interleaving attacks. The bias function 'shape' parameter κ strongly influences the distance between the actual probabilities and the asymptotic Gaussian curve. The impact on the collusion-reslilience of the code is shown. For some realistic parameter values, the false accusation probability is even *lower* than the Gaussian approximation predicts.

1 Introduction

1.1 Collusion attacks against forensic watermarking

Fingerprinting provides a means for tracing the origin and distribution of digital data. Before distribution of digital content, the content is modified by applying an imperceptible fingerprint, which plays the role of a personalized serial number. The fingerprint is usually embedded through a watermarking algorithm. Once an unauthorized copy of the content is found, the identity can be determined of those users who participated in the creation of the unauthorized copy. This can be done using a tracing algorithm, which outputs a list of allegedly guilty users. This process is also known as 'forensic watermarking'.

Reliable tracing of content requires security against attacks that aim to remove the embedded information from a copy. Collusion attacks, where a group of pirates collude to compare their copies, are a particular threat. As any differences between the copies have to arise from the watermarks and not the content, such a comparison gives information which can be used to remove the watermark. To counter this threat, coding theory has produced a number of collusion-resistant codes. In any practical implementation, they must be combined with some kind of embedding scheme. The resulting system has two layers [7, 15]: The coding layer determines which message to embed and protects against collusion attacks. The underlying watermarking layer hides symbols of the message in segments of the content. The symbols are either binary or from a larger alphabet. The interface between the fingerprinting code and the watermarking system is usually specified in terms of the *marking assumption* plus additional assumptions that are referred to as a 'model'. The marking assumption states that the colluders are able to perform modifications only in those content segments where the colluders received differently marked content. These segments are called detectable positions. The 'model' specifies the kind of symbol manipulations that the attackers are able to perform in detectable positions. The commonly used *restricted digit model* only allows them to choose pieces from their copies of the content, i.e. each segment of the unauthorized copy carries exactly one symbol that the attackers have available. The *unreadable* *digit model* allows for slightly stronger attacks. The attackers are also able to erase the fingerprint at detectable positions. Under the *arbitrary digit model* the attackers can put arbitrary symbols in detectable positions, while the *general digit model* additionally allows erasures at detectable positions.

1.2 Tardos codes

Many collusion resistant codes have been proposed in the literature. Most notable are the Boneh-Shaw construction [3] and the by now famous Tardos code [17]. The former construction uses a concatenation of an inner code with a random outer code, while the latter one is a fully randomized binary code. We briefly summarize some of the most important developments regarding Tardos codes.

In Tardos' original paper [17] a binary code was given achieving length $m = 100c_0^2\lceil\ln\frac{1}{\varepsilon_1}\rceil$, along with a proof that $m \propto c_0^2$ is asympotically optimal[1] for large coalitions, for all alphabet sizes. Here c_0 denotes the number of colluders that can be resisted, and ε_1 is the maximum allowed probability of accusing a fixed innocent user.

The original Tardos code construction contained two unfortunate design choices which caused the proportionality constant '100' to be so high. First, the false negative probability ε_2 (not accusing any of the guilty users) was coupled to ε_1 according to $\varepsilon_2 = \varepsilon_1^{c_0/4}$. This gives $\varepsilon_2 \ll \varepsilon_1$ which is highly unusual in the context of content distribution; a deterring effect is achieved already at $\varepsilon_2 \approx \frac{1}{2}$, while the false positive probability ($\approx n\varepsilon_1$, with n the number of users) needs to be very small. In the subsequent literature (e.g. [19, 2]) the ε_2 was decoupled from ε_1, leading to a substantial improvement of the code length.

Second, the symbols 0 and 1 were not treated on an equal footing. Only segments where the attackers produce a 1 were taken into account. This procedure ignores 50% of all the available information. A fully symbol-symmetric version of the Tardos code was given in [18], leading to a further improvement of the code length by a factor 4.

A further improvement was achieved in [13]. The Tardos code construction consists of two probabilistic steps. In the first step, a bias parameter is generated for each segment. In Tardos' original construction the probability density function (pdf) for the bias is a continuous function, suitable for arbitrary coalition size. In [13] a class of discrete distributions was given that performs better against finite coalition sizes than the original pdf.

All the above mentioned work followed the so-called 'simple decoder' approach, i.e. an accusation value is computed for each user independently, and if it exceeds a certain threshold, the user is considered suspicious. In contrast, one can also use a 'joint decoder' which considers sets of users. Amiri and Tardos [1] have given a capacity-achieving joint decoder construction for the binary code. (Capacity refers to the information-theoretic treatment [16, 12, 8] of the colluder attack as a communication channel.) However, the construction is rather impractical, requiring computations for many candidate coalitions.

In [18] the binary construction was generalized to alphabets of arbitrary size q, in the simple decoder approach. It was shown that, in the restricted digit model, the transition to a larger alphabet size has benefits beyond the mere fact that a q-ary symbol carries $\log_2 q$ bits of information.

1.3 Main topic of this paper: the Gaussian approximation

The so-called 'Gaussian approximation' or 'Gaussian assumption', introduced in [19], has been a useful tool in the analysis of Tardos codes. The assumption is that the pdf of a user's accusation value has a normal distribution. When this is the case, the statistical analysis of the code's performance can be drastically simplified; the performance is almost completely determined by a single parameter, namely the average accusation $\tilde{\mu}$ of the coalition.

[1] The proportionality $m \propto c_0^2$ was already known in the context of spread-spectrum watermarking. Kilian et al. [10] show that, if the watermarks have a component-wise normal distribution, then $\Omega(\sqrt{m/ln\ n})$ differently marked copies are required to successfully erase any mark with non-negligible probability.

The Gaussian assumption is motivated by the Central Limit Theorem (CLT): A user accusation consists of a sum of per-segment contributions, which are independent and identically distributed (i.i.d.). When many of these get added together, the result is close to normal-distributed, i.e. the pdf is very close to a Gaussian in a certain region around the average, and deviates in the tails. The longer the code becomes (i.e. the larger the coalition size c_0), the wider this central region. In [19] and [18] theoretical results were provided arguing that the central region is sufficiently wide to allow for application of the Gaussian approximation for realistic parameter choices. However, these arguments are not very precise in nature and have not been sufficiently corroborated.

In this paper we provide an in-depth analytical and numerical investigation of the Gaussian approximation. Our approach is based on the addition rule for generating functions, and a method to re-write the false accusation probability as a power series expansion with increasing powers of $1/m$.

1.4 Related work

Kuribayashi et al. [11] numerically studied the error probabilities of the binary Tardos code in the case of the majority voting attack. They used a fixed code length $m = 10^4$ and used a false accusation probability of around 10^{-8}. They found that the Gaussian approximation is valid under these circumstances.

Furon et al. [5] did a simulation-based numerical analysis of error probabilities for the binary Tardos code in the case of small coalitions and coupled false positive and false negative, $\varepsilon_2 = \varepsilon_1^{c_0/4}$. The used a rate-minimizing attack, yet combined it with the simple decoder. Their method was based on a type of rare event analysis where a rare event is split up into a chain of less rare events, each one conditioned on the previous. They found that the Tardos code performs better than expected.

In our work we decouple ε_2 from ε_1 and take $\varepsilon_2 \approx 0.5$. We stay within the simple decoder approach. Our method to compute probabilities is general, and can be applied to all alphabet sizes and parameter settings.

1.5 Contributions and outline

This paper discusses the case of the simple decoder, in the restricted digit model.

- We introduce a new parametrization of the colluder strategy in the restricted digit model. As usual in the literature, their strategy is allowed to be probabilistic. In a given content segment, they receive symbol α a number of times equal to σ_α. Under the usual symmetry assumptions, the strategy can be completely fixed by setting parameters which we denote as $\Psi_b(\vec{x})$; this is the probability that the attackers choose a symbol y that occurs $\sigma_y = b$ times, given that the rest of the symbols occur $\vec{x}$ times. The quantity $\Psi_b(\vec{x})$ does not depend on an actual symbol index, and is invariant under permutation of $\vec{x}$. This new parametrization allows us to obtain more compact expressions for e.g. the average accusation of the coalition ($\tilde{\mu}$), and the probability distribution of the accusation of innocent users.

- For nonbinary alphabets and realistic parameter choices, we show that the statistical parameter $\tilde{\mu}$ is minimized when the colluders employ a *majority voting* attack. In the Gaussian approximation, the code length scales as $m \propto c_0^2 \tilde{\mu}^{-2}$; hence, the colluders want to minimize $\tilde{\mu}$.

- We determine the pdf φ of an innocent user's accusation at a single content segment. We show that the tails of the pdf follow a power law which depends on the colluder strategy. Independent of the strategy, the right tail falls off faster than the left tail. This is an advantageous property, since positive accusation of innocent users is undesirable. The 'interleaving' colluder strategy, which has been conjectured [9] to be asymptotically optimal in the binary case, turns out to have special properties: the pdf and $\tilde{\mu}$ do not depend on the coalition size; both tails are maximally heavy.

- We compute the Fourier transform $\tilde{\varphi}$ (generating function) of φ. In the Fourier domain, the pdf of a sum of two variables is simply the product of their pdfs. Using this fact, we obtain an analytic result for the false accusation probability expressed in terms of $\tilde{\varphi}^m$, containing only a single integration.
- The integration mentioned in the previous point turns out to be rather difficult to compute numerically. In order to deal with this problem, we use a series expansion of $\tilde{\varphi}^m$ in powers of $1/m$. This yields an expression for the false accusation probability consisting of the Gaussian result plus correction terms of decreasing magnitude. The larger m is, the fewer terms are required. In the limit $m \to \infty$ the tail of a Gaussian is all that remains.
- We introduce a fast algorithm for computing strategy-dependent coefficients in the case of majority voting. We present numerical results for the majority voting and interleaving attacks. It turns out that the 'shape' parameter κ (which appears in the bias function, see Section 2) plays a major role in the speed of convergence to the Gaussian limit. The larger κ, the faster the convergence and the better the defense against the interleaving attack.

In Section 2 we briefly review the q-ary Tardos code and the Gaussian approximation, introduce some notation (including the new strategy parametrization), and give some lemmas that are needed for the computations in later sections. After these long preliminaries, we show in Section 3 that the majority voting attack minimizes the parameter $\tilde{\mu}$. In Section 4 we develop our method of systematically computing corrections to the Gaussian limit. Numerical results are shown in Sections 5 and 6.

2 Preliminaries

2.1 The q-ary Tardos code

The setting in this paper is the q-ary Tardos code in the restricted digit model. We briefly summarize the most important concepts and introduce the notation.

The length of a codeword (number of symbols) is denoted as m. The number of users who receive a codeword is n. The alphabet is $\mathcal{Q}$, with size q. Sometimes the alphabet will be referred to as $\{0, \cdots, q-1\}$ for simplicity. The notation $X_{ji} \in \mathcal{Q}$ stands for the i'th symbol in the codeword of user j. The whole matrix of embedded codewords is X.

Code generation

The code is generated by a two-step probabilistic algorithm. First, m vectors $\boldsymbol{p}^{(i)} \in [0,1]^q$ are independently drawn ($i \in [m]$) according to a distribution F, with

$$F(\boldsymbol{p}) = \delta(1 - \sum_{\beta \in \mathcal{Q}} p_\beta) \cdot \frac{1}{B(\kappa \mathbf{1}_q)} \prod_{\alpha \in \mathcal{Q}} p_\alpha^{-1+\kappa}. \tag{1}$$

Here $\mathbf{1}_q$ stands for the vector $(1, \cdots, 1)$ of length q, $\delta(\cdot)$ is the Dirac delta function ensuring that the components p_α add up to 1, and B is the generalized Beta function (also known as the Dirichlet integral). κ is a positive constant. In the case of the binary alphabet it is optimal to set $\kappa = 1/2$. For parameters $v_1, \cdots, v_n > 0$ the n-dimensional Beta function is defined as[2]

$$B(\boldsymbol{v}) := \int_0^1 \mathrm{d}x^n \; \delta(1 - \sum_{a=1}^n x_a) \prod_{b=1}^n x_b^{-1+v_b} = \frac{\prod_{a=1}^n \Gamma(v_a)}{\Gamma(\sum_{b=1}^n v_b)}. \tag{2}$$

In the second step of the code generation, all matrix elements X_{ji} are drawn independently according to the following distribution,

$$\Pr[X_{ji} = \alpha | \boldsymbol{p}^{(i)}] = p_\alpha^{(i)}. \tag{3}$$

[2] This is also known as a Dirichlet integral. The ordinary Beta function ($n = 2$) is $B(x, y) = \Gamma(x)\Gamma(y)/\Gamma(x + y)$.

Notice that the probabilities do not depend on the row index j, i.e. $\boldsymbol{p}^{(i)}$ determines the probabilities for a whole column of X.

The attack

The coalition of attackers is $\mathcal{C}$, with size $|\mathcal{C}| = c$. The part of X observed by the coalition is $X_\mathcal{C}$. In the restricted digit model, the attackers create a pirated version of the content such that segment i contains a symbol $y_i \in \mathcal{Q}$. (In contrast to other attack models, e.g. the combined digit model, where erasures and combinations of multiple symbols are allowed.) We define vectors $\boldsymbol{\sigma}^{(i)} \in \mathbb{N}^q$ as

$$\sigma_\alpha^{(i)} := |\{j \in \mathcal{C} : X_{ji} = \alpha\}| \tag{4}$$

i.e. the number of occurrences of the symbol α that the attackers see in column i. Obviously $\sum_{\alpha \in \mathcal{C}} \sigma_\alpha^{(i)} = c$. The attackers have a (probabilistic) strategy for choosing their output symbols. As usual in the literature on this subject, it is assumed that this strategy is fully column-symmetric, symbol-symmetric and attacker-symmetric. The assumption of column and symbol symmetry of the attack is motivated by the fact that these symmetries are present in the code generation and accusation algorithms, and that all columns and symbols are handled completely independently. The assumption of attacker-symmetry is motivated by (i) the row symmetry and independence of the rows in the code generation and accusation; (ii) the fact that any departure from attacker-symmetry will endanger one attacker more than the others.

The strategy is expressed as a set of probabilities $\theta_{y|\boldsymbol{\sigma}}$ that apply independently for each segment. Omitting the column index i, we have for each i

$$\Pr[\text{output } y, \text{ given } \boldsymbol{\sigma}] = \theta_{y|\boldsymbol{\sigma}}. \tag{5}$$

Due to the marking condition some of these probabilities are fixed. Let $\boldsymbol{e}_\alpha$ denote the vector $(0, \cdots, 0, 1, 0, \cdots, 0)$ with the '1' in position α. Then

$$\theta_{y|c\boldsymbol{e}_\alpha} = \delta_{y\alpha}, \tag{6}$$

where δ is the Kronecker delta.

Accusation

The watermark detector sees the symbol y_i embedded in segment i of the attacked content. Users are classified as suspicious ('accused') or not suspicious according to the following algorithm. For each user j, the so-called *accusation sum* S_j is computed,

$$S_j = \sum_{i=1}^{m} S_j^{(i)} \quad \text{where} \quad S_j^{(i)} = g_{[X_{ji}==y_i]}(p_{y_i}^{(i)}), \tag{7}$$

where the expression $[X_{ji} == y_i]$ evaluates to 1 if $X_{ji} = y_i$ and to 0 otherwise, and the functions g_0 and g_1 are defined as

$$g_1(p) = \sqrt{\frac{1-p}{p}} \quad ; \quad g_0(p) = -\sqrt{\frac{p}{1-p}}. \tag{8}$$

In words: Having the same symbol as the attacked content induces a positive contribution $g_1(p_{y_i})$ to the accusation sum, which becomes worse when y_i is unlikely to occur. Having a symbol different from y_i induces a negative amount $g_0(p_{y_i})$, which becomes more negative when y_i is likely to occur. The total accusation of the coalition is defined as $S := \sum_{j \in \mathcal{C}} S_j$.

The choice (8) of g_0, g_1 is the unique combination of functions that satisfies

$$pg_1(p) + (1-p)g_0(p) = 0 \quad ; \quad p[g_1(p)]^2 + (1-p)[g_0(p)]^2 = 1. \tag{9}$$

This choice has been shown to have optimal properties for the binary alphabet [4, 19]. Its unique properties (9) also hold for $q \geq 3$; that is the main motivation for using (8).

A user is 'accused' if his accusation sum exceeds a threshold Z. A list $\mathcal{L}$ is made of accused users,

$$\mathcal{L} = \{j : S_j > Z\}. \tag{10}$$

Performance

The 'performance' of the scheme involves four important parameters: the number of attackers that has to be resisted (c_0), the maximum tolerable false negative probability ε_2 (prob. of not catching any of the attackers),

$$\Pr[\mathcal{L} \cap \mathcal{C} = \emptyset] \leq \varepsilon_2, \tag{11}$$

the maximum tolerable false positive probability ε_1

$$\text{for fixed innocent } j: \quad \Pr[j \in \mathcal{L}] \leq \varepsilon_1, \tag{12}$$

and the length m of the code. (Note that the total probability of false positives occurring is approximately $n\varepsilon_1$.) One way of measuring how well the scheme works is to look at how big m has to be as a function of c_0, ε_1 and ε_2. The smaller m, the better the scheme. It is important to note that in forensic watermarking of AV content, a small false positive probability is the primary requirement. The false negative is far less important, since the deterring effect of forensic watermarking is preserved even for large ε_2, of the order of $1/2$. Hence m essentially becomes a function of c_0 and ε_1. In [18] an asymptotic result was obtained for large c_0,

$$m = \frac{2}{\tilde{\mu}^2} c_0^2 \ln \frac{1}{\varepsilon_1 \sqrt{2\pi}}. \tag{13}$$

Here $\tilde{\mu}$ is the expectation value of the collective accusation sum of the coalition, scaled in such a way that the dependence on m is removed: $\tilde{\mu} = \mathbb{E}[S]/m$. In the case of the binary scheme (with $\kappa = 1/2$), $\tilde{\mu} = 2/\pi \approx 0.64$. For larger alphabets the $\tilde{\mu}$ depends on the parameter κ in a complicated way; for optimal κ, the $\tilde{\mu}$ takes values from approximately 0.8 to 1.4 as q goes from 3 to 10.

2.2 The Gaussian approximation

We briefly review the analysis of error probabilities performed in [18], which leads to the result (13).

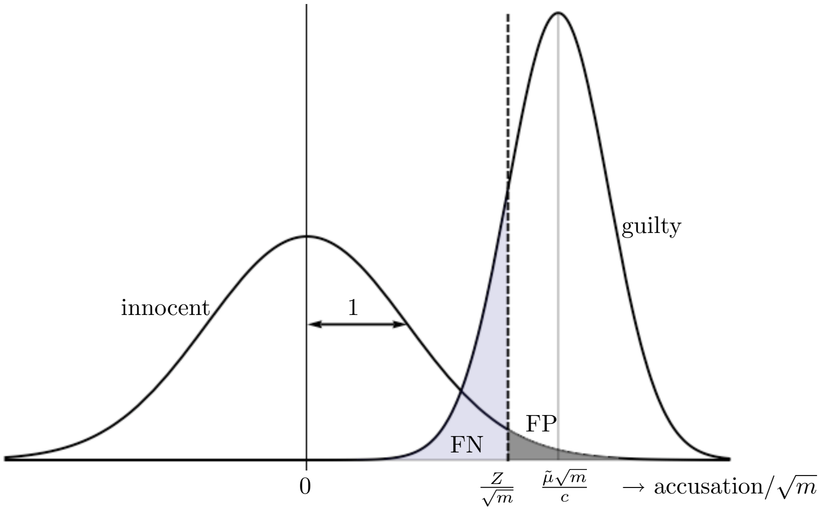

Figure 1: *Sketch of the probability distributions of $S_j/\sqrt{m}$ for some fixed innocent j, and of $S/(c\sqrt{m})$. The horizontal axis is scaled by a factor $\sqrt{m}$ so that the variance of the innocent curve is exactly 1.*

Consider, for some innocent user j, the probability distribution function (pdf) ρ_m of the quantity $S_j/\sqrt{m}$. (Note that the pdf itself depends on m.) From (9) it follows that ρ_m has zero mean and unit variance. For brevity we mow introduce the notation $\tilde{Z} = Z/\sqrt{m}$. The probability of falsely accusing j is given by

$$\int_{\tilde{Z}}^{\infty} \mathrm{d}x \; \rho_m(x) =: R_m(\tilde{Z}). \tag{14}$$

This is depicted as the shaded area 'FP' in Fig. 1. We require

$$R_m(\tilde{Z}) \leq \varepsilon_1. \tag{15}$$

Similarly, consider the probability distribution τ_m of the quantity $S/(c\sqrt{m})$, but normalized in such a way that the mean is zero and the variance is 1. The cumulative distribution function is

$$T_m(x) := \int_{-\infty}^{x} \mathrm{d}x' \; \tau_m(x'). \tag{16}$$

It was shown in [17] that $\Pr[\mathrm{FN}] \leq \Pr[S < cZ]$. Hence if $\Pr[S < cZ] \leq \varepsilon_2$ then automatically $\Pr[\mathrm{FN}] \leq \varepsilon_2$. The shaded area in Fig. 1 labeled as 'FN' is actually $\Pr[S < cZ]$, which acts as a handy bound on the FN. This area is given by $T([\tilde{Z} - \frac{\tilde{\mu}\sqrt{m}}{c}]/\frac{\tilde{\sigma}}{c}) = T(\frac{c\tilde{Z} - \tilde{\mu}\sqrt{m}}{\tilde{\sigma}})$, where $\tilde{\sigma}$ is the (scaled) standard deviation of the collective accusation, $m\tilde{\sigma}^2 := \mathbb{E}[S^2] - (\mathbb{E}[S])^2$. The requirement on the FN probability in case of c_0 attackers is then formulated as

$$T_m(\frac{c_0\tilde{Z} - \tilde{\mu}\sqrt{m}}{\tilde{\sigma}}) \leq \varepsilon_2. \tag{17}$$

The two equations (15) and (17) for given c_0, ε_1, ε_2 can be thought of as constraints in the (Z, m)-plane . It was shown that these constraints can be satisfied only if

$$m \geq \frac{1}{\tilde{\mu}^2} c_0^2 \left[R_m^{\mathrm{inv}}(\varepsilon_1) - \frac{\tilde{\sigma}}{c_0} T_m^{\mathrm{inv}}(\varepsilon_2) \right]^2 \tag{18}$$

where R_m^{inv} and T_m^{inv} are the inverse functions of R_m and T_m respectively. Note that $T_m^{\mathrm{inv}}(\varepsilon_2) < 0$ for ε_2 smaller[3] than approximately 1/2; decreasing ε_2 leads to a longer code. It was shown that the T_m^{inv} term is negligible with respect to the R_m^{inv} term if c_0 is large and/or $\varepsilon_2 \approx 1/2$. Hence, (18) in practice reduces to

$$m \geq m_{\min} \quad ; \quad m_{\min} \approx \frac{1}{\tilde{\mu}^2} c_0^2 \left[R_m^{\mathrm{inv}}(\varepsilon_1) \right]^2. \tag{19}$$

Eq. (19) in itself is not immediately useful, because R_m depends on m. In the limit of large m, however, ρ_m simply becomes a Gaussian independent of m, and R_m is the area under a Gaussian tail, which we denote as $\Omega(\tilde{Z}) = \frac{1}{2}\mathrm{Erfc}\frac{\tilde{Z}}{\sqrt{2}}$. (Here Erfc is the complementary error function.) The result (13) follows by applying the bound $[\Omega^{\mathrm{inv}}(\varepsilon_1)]^2 = [\sqrt{2}\,\mathrm{Erfc}^{\mathrm{inv}}(2\varepsilon_1)]^2 < 2\ln(\varepsilon_1\sqrt{2\pi})^{-1}$.
To the best of our knowledge, the above reasoning is the simplest argument available that yields the asymptotic relation $m \propto c_0^2$.
It was argued in [19] and [18] that m is so large that ρ_m is Gaussian even a sufficient number of standard deviations away from 0. ('Sufficient' here means that the area under the Gaussian part is at least $1 - 2\varepsilon_1$ so that the area under the right tail is estimated accurately.) The argument was based on the moments of the innocent accusation. However, a full analysis of the tails of ρ has never been done. Such a full analysis is important for the following reason. As (19) shows, it is advantageous for the attackers not only to decrease $\tilde{\mu}$, but also to modify the shape of R_m such that $R_m^{\mathrm{inv}}(\varepsilon_1)$ increases, i.e. such that the right-hand tail of the innocent's accusation probability becomes longer. How much influence their strategy has on the shape of R_m will be studied in Sections 5 and 6. If there is hardly any influence, then the value of $\tilde{\mu}$ uniquely determines $m_{\min}$, and the optimal strategy is to minimize $\tilde{\mu}$; if there is a significant influence, then the attackers' aim is to maximize the quotient $R_m^{\mathrm{inv}}(\varepsilon_1)/\tilde{\mu}$.

[3] If one is willing to set $\varepsilon_2 > 1/2$, the contribution from $T_m^{\mathrm{inv}}(\varepsilon_2)$ may even reduce the code length.

Notation	**Meaning**
$\mathcal{Q}$	the alphabet
q	alphabet size $\lvert\mathcal{Q}\rvert$
n	number of users
$\mathcal{C}$	set of colluding users
c	number of colluders $\lvert\mathcal{C}\rvert$
c_0	coalition size that the code can resist
m	code length (number of q-ary symbols)
X_{ji}	embedded symbol in segment i for user j
$\boldsymbol{p}^{(i)}$	bias vector for column i
F	distribution function of the bias vector, $\boldsymbol{p}^{(i)} \sim F$
$f(p_\alpha)$	marginal distribution of F for one component
κ	shape parameter contained in F
$\sigma_\alpha^{(i)}$	number of occurrences of symbol α in attackers' segment i
$\mathbb{P}$	probability distribution for $\boldsymbol{\sigma}$
$\mathbb{P}_1$	marginal distribution for one component of $\boldsymbol{\sigma}$
$\mathbb{P}_{q-1}$	marginal distribution for $q-1$ components of $\boldsymbol{\sigma}$
y_i	symbol in segment i of attacked content
$\theta_{y\mid\boldsymbol{\sigma}}$	prob. that attackers output symbol y, given $\boldsymbol{\sigma}$
S_j	accusation sum of user j
S	coalition accusation sum, $S = \sum_{j\in\mathcal{C}} S_j$
Z	accusation threshold
$\tilde{Z}$	$Z/\sqrt{m}$
$\mathcal{L}$	list of accused users
ε_1	max. tolerable prob. of fixed innocent user getting accused
ε_2	max. tolerable prob. of not catching any attacker
FP	false positive
FN	false negative
$\tilde{\mu}$	$\mathbb{E}[S]/m$; does not depend on m
ρ_m	prob. distribution of $S_j/\sqrt{m}$ for innocent j
R_m	area function for the right-hand tail of ρ_m
τ_m	prob. distribution of $S/(c\sqrt{m})$, normalized to zero mean and variance 1
T_m	cumulative distribution function for τ_m
φ	prob. distribution of one-segment contribution to innocent's accusation
$\Psi_b(\boldsymbol{x})$	$\theta_{y\mid\boldsymbol{\sigma}}$ when $\sigma_y = b$ and the rest of $\boldsymbol{\sigma}$ is equal to $\boldsymbol{x}$
K_b	quantity derived from $\Psi_b(\boldsymbol{x})$
$\Omega(x)$	probability mass in the right tail of a Gaussian, beyond x

2.3 Probabilities and expectation values

For given $\boldsymbol{p}$, the probability that the colluders receive symbol occurrences $\boldsymbol{\sigma}$ is the multinomial distribution. We use the following notation,

$$\mathbb{P}(\boldsymbol{\sigma}|\boldsymbol{p}) := \binom{c}{\boldsymbol{\sigma}} \prod_{\alpha\in\mathcal{Q}} p_\alpha^{\sigma_\alpha}, \tag{20}$$

where $\binom{c}{\boldsymbol{\sigma}} = c!/(\prod_\alpha \sigma_\alpha!)$. It is always implicitly understood that $\sum_\alpha \sigma_\alpha = c$. The marginal distribution for a single component σ_α is the binomial. We use the notation

$$\mathbb{P}_1(b|p) := \Pr[\sigma_\alpha = b|p_\alpha = p] = \binom{c}{b} p^b (1-p)^{c-b}. \tag{21}$$

Lemma 1 *The overall probability that the colluders receive symbol occurrences $\boldsymbol{\sigma}$ is given by*

$$\mathbb{P}(\boldsymbol{\sigma}) := \binom{c}{\boldsymbol{\sigma}} \frac{B(\kappa\mathbf{1}_q + \boldsymbol{\sigma})}{B(\kappa\mathbf{1}_q)}.$$

Proof: We have $\Pr[\boldsymbol{\sigma}] = \mathbb{E}_{\boldsymbol{p}}\mathbb{P}(\boldsymbol{\sigma}|\boldsymbol{p})$, with $\mathbb{P}(\boldsymbol{\sigma}|\boldsymbol{p})$ given by (20). The expectation $\mathbb{E}_{\boldsymbol{p}}$ stands for $\mathbb{E}_{\boldsymbol{p}}[\cdots] = \int_0^1 \mathrm{d}^q p \; F(\boldsymbol{p})(\cdots)$, with F defined in (1). The lemma follows by applying the Dirichlet integration rule (2). □

Lemma 2 *The marginal probability distribution $f(p_\alpha)$ for a single component of the vector $\boldsymbol{p}$ is*

$$f(p_\alpha) = \frac{1}{B(\kappa, \kappa[q-1])} p_\alpha^{-1+\kappa} (1-p_\alpha)^{-1+\kappa[q-1]}.$$

Proof: We have $\int_0^1 \mathrm{d}p_\alpha \; f(p_\alpha) = \int_0^1 \mathrm{d}^q p \; F(\boldsymbol{p})$. In the latter integral, we write for all $\beta \neq \alpha$: $p_\beta = (1-p_\alpha)s_\beta$, with $s_\beta \in [0,1]$. This gives $\mathrm{d}^q p = \mathrm{d}p_\alpha (1-p_\alpha)^{q-1} \mathrm{d}^{q-1}s$, and $\prod_\gamma p_\gamma^{-1+\kappa} = p_\alpha^{-1+\kappa}(1-p_\alpha)^{(q-1)(-1+\kappa)} \prod_{\beta\in\mathcal{Q}\setminus\alpha} s_\beta^{-1+\kappa}$, and $\delta(1 - \sum_{\gamma\in\mathcal{Q}} p_\gamma) = \delta([1-p_\alpha][1 - \sum_{\beta\in\mathcal{Q}\setminus\alpha} s_\beta]) = (1-p_\alpha)^{-1}\delta(1 - \sum_{\beta\in\mathcal{Q}\setminus\{\alpha\}} s_\beta)$. Combining all these ingredients, we find

$$\int_0^1 \mathrm{d}p_\alpha \; f(p_\alpha) = \int_0^1 \mathrm{d}p_\alpha \; p_\alpha^{-1+\kappa}(1-p_\alpha)^{-1+\kappa[q-1]} \frac{1}{B(\kappa\mathbf{1}_q)} \int_0^1 \mathrm{d}^{q-1}s \; \delta(1 - \sum_{\gamma\in\mathcal{Q}\setminus\alpha} s_\gamma) \prod_{\beta\in\mathcal{Q}\setminus\alpha} s_\beta^{-1+\kappa}. \tag{22}$$

The lemma follows after evaluation of the $\int \mathrm{d}^{q-1}s$ integral using (2). □

Lemma 3 *The overall marginal probability distribution for one component of $\boldsymbol{\sigma}$ is*

$$\mathbb{P}_1(b) := \Pr[\sigma_\alpha = b] = \binom{c}{b} \frac{B(\kappa + b, \kappa[q-1] + c - b)}{B(\kappa, \kappa[q-1])}.$$

Proof: We have $\Pr[\sigma_\alpha = b] = \int_0^1 \mathrm{d}p_\alpha f(p_\alpha)\mathbb{P}_1(b|p_\alpha)$ with $\mathbb{P}_1(b|p_\alpha)$ and $f(p_\alpha)$ given by (21) and Lemma 2 respectively. The integral is evaluated using (2). □

Corollary 1 *Let $\boldsymbol{\sigma}_{\setminus\alpha}$ denote the vector $\boldsymbol{\sigma}$ without the component σ_α. The probability distribution of $\boldsymbol{\sigma}_{\setminus\alpha}$ conditioned on σ_α is given by*

$$\mathbb{P}_{q-1}(\boldsymbol{x}|b) := \Pr[\boldsymbol{\sigma}_{\setminus\alpha} = \boldsymbol{x}|\sigma_\alpha = b] = \binom{c-b}{\boldsymbol{x}} \frac{B(\kappa\mathbf{1}_{q-1} + \boldsymbol{x})}{B(\kappa\mathbf{1}_{q-1})}.$$

Proof: Follows directly from Lemmas 1 and 3 by taking $\Pr[\boldsymbol{\sigma}_{\setminus\alpha} = \boldsymbol{x}|\sigma_\alpha = b] = \mathbb{P}(\boldsymbol{\sigma} = (\boldsymbol{x}, b))/\mathbb{P}_1(b)$ and simplifying the Beta functions. □

We introduce a new parametrization of the colluder strategy. For $b \in \{1, \cdots, c\}$ and $\boldsymbol{x} \in \mathbb{N}^{q-1}$, with $\sum_a x_a = c - b$, we define

$$\Psi_b(\boldsymbol{x}) := \theta_{\alpha|(\sigma_\alpha = b, \boldsymbol{\sigma}_{\setminus\alpha} = \boldsymbol{x})}. \tag{23}$$

The vector $\boldsymbol{\sigma}$ has $\sigma_\alpha = b$, and the other $q-1$ components are given by $\boldsymbol{x}$. The probability for outputting α given such a $\boldsymbol{\sigma}$ does not depend on the actual value of α, but only on b and $\boldsymbol{x}$. (In fact, it is even insensitive to permutations of $\boldsymbol{x}$.) This follows from the symbol-symmetry and attacker-symmetry of the attack strategy. In words: $\Psi_b(\boldsymbol{x})$ is the coalition's probability of outputting a symbol which for them occurs b times, with the other symbol frequencies being $\boldsymbol{x}$. In the case of the binary alphabet, $\boldsymbol{x}$ has only one component equal to $c-b$. We will then use the notation Ψ_b, with $\Psi_0 = 0$ and $\Psi_c = 1$ due to the marking condition.
Next we define

$$K_b := \mathbb{E}_{\boldsymbol{x}|b}\Psi_b(\boldsymbol{x}) = \sum_{\boldsymbol{x}} \mathbb{P}_{q-1}(\boldsymbol{x}|b)\Psi_b(\boldsymbol{x}). \tag{24}$$

It is implicit that $\sum_{\beta\in\mathcal{Q}\setminus\{\alpha\}} x_\beta = c-b$. For $q=2$ we define $K_b = \Psi_b$. (In some of the literature the notation $\theta_x := \Pr[y=1|\ \#\text{received 1s} = x]$ is used for the binary case. The relation with our notation is: $\theta_b = \Psi_b$.)
For any pirate strategy we have

$$K_0 = 0 \quad ; \quad K_c = 1 \tag{25}$$

due to the marking assumption.

Lemma 4 *The numbers K_b satisfy*

$$q\sum_{b=1}^{c} K_b \mathbb{P}_1(b) = 1.$$

Proof: The factor q can be replaced by $\sum_{y\in\mathcal{Q}}$. Using the definition (24) we get $\sum_y\sum_b K_b\mathbb{P}_1(b) = \sum_b\sum_{\boldsymbol{x}}\mathbb{P}(\boldsymbol{x},b)\cdot\sum_y\Psi_b(\boldsymbol{x}) = \sum_b\sum_{\boldsymbol{x}}\mathbb{P}(\boldsymbol{x},b)\cdot\sum_y\theta_{y|\sigma_y=b,\boldsymbol{\sigma}_{\setminus y}=\boldsymbol{x}} = \sum_b\sum_{\boldsymbol{x}}\mathbb{P}(\boldsymbol{x},b) = 1$. □

Lemma 5 *If the colluder strategy is the interleaving attack, $\theta_{y|\boldsymbol{\sigma}} = \frac{\sigma_y}{c}$, then $K_b = b/c$.*

Proof: This strategy implies $\Psi_b(\vec{x}) = b/c$ independent of $\boldsymbol{x}$. Substitute this into (24) and use the fact that the probabilities add up to 1. □

2.4 Integrals and Gamma function equalities

Lemma 6 *For $d>0$, $v>0$, the following holds*

$$\int_0^\infty \mathrm{d}u\ \frac{u^{2d-1}}{(1+u^2)^{d+v}} = \tfrac{1}{2}B(d,v).$$

Proof: Apply a change of variables $u = \sqrt{p/(1-p)}$, with $p\in[0,1]$. This gives $1+u^2 = 1/(1-p)$ and $\mathrm{d}u = \frac{1}{2}p^{-1/2}(1-p)^{-3/2}\mathrm{d}p$. The integral becomes $\frac{1}{2}\int_0^1 \mathrm{d}p\ p^{-1+d}(1-p)^{-1+v}$ which has the Dirichlet form (2). □

Lemma 7 *For $x \gg 1$, and a_1, a_2 such that $|a_1| \ll x$ and $|a_2| \ll x$, it holds that*

$$\frac{\Gamma(x+a_1)}{\Gamma(x+a_2)} = x^{a_1-a_2}[1+\mathcal{O}(\frac{1}{x})].$$

Proof: Follows directly from Stirling's approximation $\Gamma(z+1) \approx \sqrt{2\pi z}(z/e)^z$. □

Lemma 8 *Let $c \gg 1$ and $1 \ll b \leq c$. Let $\alpha_1, \alpha_2, \beta_1, \beta_2 \ll b$. Then*

$$\frac{B(b+\alpha_1, c-b+\beta_1)}{B(b+\alpha_2, c-b+\beta_2)} = (\frac{b}{c})^{\alpha_1-\alpha_2}(1-\frac{b}{c})^{\beta_1-\beta_2}[1+\mathcal{O}(\frac{1}{b})].$$

Proof: Follows directly from writing out the Beta functions in terms of Gamma functions and then applying Lemma 7. □

Definition 1 *We define $\Omega(z)$ as the probability mass in the right tail of the normal distribution beyond point z,*

$$\Omega(z) = \frac{1}{\sqrt{2\pi}} \int_z^\infty \mathrm{d}x \; e^{-x^2/2}.$$

Lemma 9 (See e.g. Eq. 9.254.1 in [6]) *For $x \in \mathbb{R}$*

$$\frac{1}{2\pi i} \int_{-\infty}^{\infty} \mathrm{d}k \; \frac{e^{ikx}}{k} e^{-k^2/2} = \tfrac{1}{2} - \Omega(x).$$

Lemma 10 (See e.g. Eq. 3.462.1 in [6]) *For $\nu > 0$ and $x \in \mathbb{R}$*

$$\int_0^\infty \mathrm{d}k \; k^{\nu-1} e^{-\frac{1}{2}k^2} e^{ikx} = \Gamma(\nu) 2^{\nu/2} H_{-\nu}(\frac{-ix}{\sqrt{2}}). \tag{26}$$

Here H is the Hermite function.

Corollary 2 *For $x \in \mathbb{R}$ and $\nu > 0$*

$$\begin{aligned}\int_{-\infty}^{\infty} \frac{\mathrm{d}k}{2\pi} (i \operatorname{sgn} k)^{\alpha-1} |k|^{\nu-1} e^{-k^2/2} e^{ikx} &= \frac{1}{\pi} \Gamma(\nu) 2^{\nu/2} \operatorname{Im} \left[i^{-\alpha} H_{-\nu}(\frac{ix}{\sqrt{2}}) \right] \\ &= \frac{2^{-(\nu-1)/2}}{(\sin \nu\pi)\sqrt{2\pi}} e^{-\frac{1}{2}x^2} \left[H_{\nu-1}(\frac{x}{\sqrt{2}}) \sin \frac{\pi}{2}(\nu - \alpha) - H_{\nu-1}(\frac{-x}{\sqrt{2}}) \sin \frac{\pi}{2}(\nu + \alpha) \right]. \end{aligned} \tag{27}$$

Proof: The first equality follows by applying Lemma 10 twice (once for the positive part of the integral, once for the negative). The second equality follows from the properties of the Hermite function (see e.g. page 1094 of [6]). □

Remark: In the case $\alpha = \nu$, the first term of the last line vanishes, yielding $-2^{-\frac{\nu}{2}} / \sqrt{\pi} e^{-\frac{1}{2}x^2} H_{\nu-1}(\frac{-x}{\sqrt{2}})$. For integer ν, the Hermite function reduces to a Hermite polynomial.

2.5 Fourier transforms

Definition 2 *Let $\chi : \mathbb{R} \to \mathbb{C}$ be a function. The Fourier transform of χ is denoted as $\tilde{\chi}$ and defined as*

$$\tilde{\chi}(k) = \int_{-\infty}^{\infty} \mathrm{d}x \; e^{-ikx} \chi(x) \qquad \textit{with } k \in \mathbb{R}.$$

Lemma 11 *If χ is a real-valued function, then $\tilde{\chi}(-k) = [\tilde{\chi}(k)]^*$.*

Proof: $[\int \mathrm{d}x \; e^{-ikx} \chi(x)]^* = \int \mathrm{d}x \; [e^{-ikx} \chi(x)]^* = \int \mathrm{d}x \; e^{ikx} \chi(x) = \tilde{\chi}(-k)$. □

Corollary 3 *If χ is a real-valued function, then the even part of $\tilde{\chi}(k)$ is $\operatorname{Re} \tilde{\chi}(k)$, and the odd part is $i \cdot \operatorname{Im} \tilde{\chi}(k)$.*

Proof: By Lemma 11, the even part is $\frac{1}{2}[\tilde{\chi}(k) + \tilde{\chi}(-k)] = \frac{1}{2}\tilde{\chi}(k) + \frac{1}{2}[\tilde{\chi}(k)]^* = \operatorname{Re} \tilde{\chi}(k)$. The odd part is $\frac{1}{2}[\tilde{\chi}(k) - \tilde{\chi}(-k)] = \frac{1}{2}\tilde{\chi}(k) - \frac{1}{2}[\tilde{\chi}(k)]^* = i \operatorname{Im} \tilde{\chi}(k)$. □

Lemma 12 *Let $\chi(x)$ be a probability distribution function, and X a random variable with $X \sim \chi$. Then*

$$\left. \frac{\partial^n \tilde{\chi}(k)}{\partial k^n} \right|_{k=0} = (-i)^n \mathbb{E}[X^n].$$

Proof: $\frac{\partial^n \tilde{\chi}(k)}{\partial k^n} = \int \mathrm{d}x \; [\frac{\partial^n}{\partial k^n} e^{-ikx}] \chi(x) = (-i^n) \int \mathrm{d}x \; x^n e^{-ikx} \chi(x)$. Setting $k = 0$ gives the result. □

Corollary 4 *Let φ be the probability distribution function of the one-symbol accusation $S_j^{(i)}$ for an innocent user j. Then its Fourier transform $\tilde{\varphi}$ has the following power series expansion,*

$$\tilde{\varphi}(k) = 1 - \tfrac{1}{2}k^2 + \text{higher powers of } k,$$

where the higher powers of k are allowed to be irrational.

Proof: We denote $u = S_j^{(i)}$ for brevity. Trivially $\mathbb{E}[u^0] = 1$. From (9) we know that $\mathbb{E}[u] = 0$ and $\mathbb{E}[u^2] = 1$. Hence by Lemma 12 we have $\tilde{\varphi}(0) = 1$, $\tilde{\varphi}'(0) = 0$ and $\tilde{\varphi}''(0) = -1$. The expansion in the corollary is consistent with these values. □

Higher orders of k do not have to be integer. In fact, if $\mathbb{E}[u^3] \neq 0$, $\mathbb{E}[u^3] < \infty$ and $\mathbb{E}[u^4] = \infty$ (as we will see is the case in Section 4.2) then there is a k^3 term in the expansion, and the lowest power of k higher than 3 lies somewhere between 3 and 4.

3 Strategy for minimizing $\tilde{\mu}$

Definition 3 *When we use the term 'majority voting' it will mean the following:*

- *If $\exists \alpha : \sigma_\alpha > \sigma_\beta$ for all $\beta \neq \alpha$, then output α. (If one symbol occurs more often than all the others, output this symbol.)*
- *If the most frequently occurring symbol is not unique, i.e. there are multiple such symbols, then output one of them uniformly at random.*

Similarly, by 'minority voting' we mean:

- *If $\exists \alpha : 1 \leq \sigma_\alpha < \sigma_\beta$ for all $\beta \neq \alpha$, then output α. (If one symbol occurs less often than all the others, output this symbol.)*
- *If the least frequently occurring symbol is not unique, i.e. there are multiple such symbols, then output one of them uniformly at random.*

3.1 Binary alphabet

The case $q = 2$ is simple. It was shown in [18] that for $\kappa > 1/2$ *minority voting* is optimal (in the sense of minimizing $\tilde{\mu}$), while for $\kappa < 1/2$ it is *majority voting*. For $\kappa = 1/2$ the strategy has no effect on $\tilde{\mu}$, whose value is then $2/\pi$.

3.2 Non-binary alphabet

In [18] the following expression was obtained for $\tilde{\mu}$ (for the case $q \geq 3$),

$$\begin{aligned} \tilde{\mu} &= \sum_{\boldsymbol{\sigma}} \mathbb{P}(\boldsymbol{\sigma}) \sum_{y \in \mathcal{Q}} \theta_{y|\boldsymbol{\sigma}} W(\sigma_y) \left\{ \tfrac{1}{2} - \kappa + \frac{\sigma_y}{c}(\kappa q - 1) \right\} \qquad (28) \\ W(b) &:= c \frac{\Gamma(b + \kappa - \frac{1}{2})}{\Gamma(b + \kappa)} \frac{\Gamma(c - b + \kappa[q-1] - \frac{1}{2})}{\Gamma(c - b + \kappa[q-1])}. \end{aligned}$$

The colluders want to minimize $\tilde{\mu}$, while the content owner wants to maximize it.

Theorem 1 *For $q \geq 3$ and $\kappa \approx 1/q$, the majority voting strategy minimizes $\tilde{\mu}$.*

Proof: The 'optimal' colluder strategy (in the sense of making $\tilde{\mu}$ as small as possible) is, for given $\boldsymbol{\sigma}$, to choose y such that the expression $W(\sigma_y)\{\frac{1}{2} - \kappa + \frac{\sigma_y}{c}(\kappa q - 1)\}$ is minimized. It was found numerically in [18] that the optimal choice of the parameter κ against this attack is slightly larger than $1/q$. Putting $\kappa \approx 1/q$ in (28), we see that the optimal attack strategy is effectively

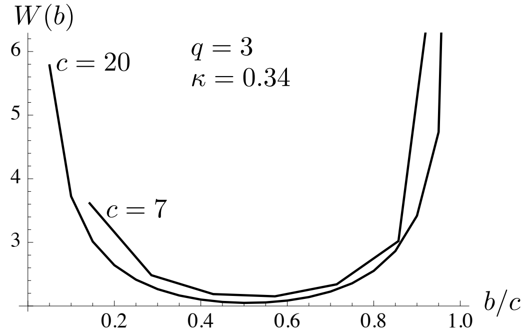

Figure 2: *Example of* $W(b)$ *for* $q = 3$, $\kappa = 0.34$.

to minimize W, i.e. the coalition chooses $y = \arg\min\{W(\sigma_\alpha)\}_{\alpha\in\mathcal{Q}}$. Numerical inspection shows that the function $W(b)$ has a minimum at $b = \lceil c/2 \rceil$ (see Fig. 2).
For large c this is easily understood: application of Lemma 7 for large b and $c-b$ gives $W(b) \approx [\frac{b}{c}(1-\frac{b}{c})]^{-1/2}$, a function with its minimum at $b = c/2$ and symmetric around this minimum. Hence the optimal strategy consists of choosing the symbol α whose σ_α is closest to $c/2$. It turns out that this is precisely the same as majority voting. This can be seen as follows. First consider the case where the 'closest to $c/2$' strategy results in $\sigma_y > c/2$. Because of the sum rule $\sum_\alpha \sigma_\alpha = c$, there can be no $\alpha \neq y$ with $\sigma_\alpha > c/2$; hence the strategy has resulted in selecting the majority symbol. Second, consider the 'closest to $c/2$' strategy yielding $\sigma_y = c/2 - \delta$, with $\delta > 0$. If there is any $\alpha \neq y$ with $\sigma_\alpha > \sigma_y$, it will have to satisfy $\sigma_\alpha \geq c/2 + \delta = c - \sigma_y$. Only the equality is allowed ($\sigma_\alpha = c - \sigma_y$) by the sum rule; it gives rise to almost the same amount of accusation as σ_y, since $W(b)$ is very close to symmetric around $c/2$. □

Theorem 2 *The quantity $\tilde{\mu}$ as defined in (28) can be written as*

$$\tilde{\mu} = q \sum_{b=1}^{c} \mathbb{P}_1(b) K_b W(b) \left\{ \tfrac{1}{2} - \kappa + \frac{b}{c}(\kappa q - 1) \right\}. \tag{29}$$

Proof: In (28) we shift the $\sum_y$ to the front and write $\mathbb{P}(\boldsymbol{\sigma}) = \Pr[\sigma_y = b]\Pr[\boldsymbol{\sigma}_{\setminus y} = \boldsymbol{x}|\sigma_y = b]$ and $\sum_{\boldsymbol{\sigma}} = \sum_b \sum_{\boldsymbol{x}}$. The $\sum_{\boldsymbol{x}}$ of $\theta_{y|\boldsymbol{\sigma}}$ yields K_b according to the definition (24). □

Corollary 5 *For $\kappa > \frac{1}{2(q-1)}$ the contribution of the $b = c$ term to $\tilde{\mu}$ vanishes in the limit of large c.*

Proof: In (29) we split off the $b = c$ term, which has $K_c = 1$ due to the marking conditon. After some rewriting of Gamma functions this yields

$$\tilde{\mu} = cq \frac{B(c + \kappa - \frac{1}{2}, \kappa[q-1] + \frac{1}{2})}{B(\kappa, \kappa[q-1])} + q \sum_{b=1}^{c-1} \mathbb{P}_1(b) K_b W(b) \left\{ \tfrac{1}{2} - \kappa + \frac{b}{c}(\kappa q - 1) \right\}. \tag{30}$$

In the limit of large c, the first term scales as $(1/c)^{\kappa[q-1]-1/2}$. For $\kappa[q-1] > \frac{1}{2}$ this vanishes asymptotically. □
Corollary 5 tells us that in the relevant case $\kappa \approx 1/q$, the contributions to $\tilde{\mu}$ work completely different than in the usual binary scheme ($q = 2, \kappa = \frac{1}{2}$). There the $b = c$ term scales as c^0 and all the $b < c$ terms are zero.

4 Statistics of the accusations

4.1 Our approach: Fourier transform

We now describe the basis of our method of computing false accusation probabilities. The whole approach is based on a single observation: *when random variables are added, the pdf of the sum is obtained by multiplying the Fourier transforms of their respective pdf's and then doing a Fourier back-transform.* In other words, if $X \sim f_1$, $Y \sim f_2$ and $Z = X + Y \sim f_3$, then $\tilde{f}_3 = \tilde{f}_1 \tilde{f}_2$. When this rule is applied to the m random variables in the accusation sum, it leads to the following result.

Theorem 3 *Let j be an innocent user. Let φ denote the pdf of $S_j^{(i)}$, with $S_j^{(i)}$ as defined in (7). Let $\tilde{\varphi}$ be the Fourier transform of φ. Then the probability that $S_j > Z$ is given by*

$$R_m(\tilde{Z}) = \frac{1}{2} + \frac{i}{2\pi} \int_{-\infty}^{\infty} \mathrm{d}k \, \frac{\exp ik\tilde{Z}}{k} \left[\tilde{\varphi}(\frac{k}{\sqrt{m}}) \right]^m . \tag{31}$$

Proof: see Appendix A. □

This result gives us a closed-form expression for $R_m(\tilde{Z})$ that contains only a single integration and a limited number of sums. (The sums are contained in the evaluation of $\tilde{\varphi}$, as will become apparent in Section 4.3.) These will have to be evaluated numerically. Note that $\Pr[S_j > 0]$ is not necessarily equal to $\frac{1}{2}$.

It turns out that numerical evaluation of the integral in (31) is difficult, because of the fast oscillations of the integrand at large k. For this reason, we have chosen for a somewhat indirect method of evaluating (31). It is based on a series expansion in powers of k. It has the advantage that the accuracy of the numerical evaluation is well under control, and that the dependence of R_m on m is visible. The disadvantage is that many terms in the expansion have to be kept.

Theorem 4 *Let j be an innocent user. Let φ have a finite third moment. Then it is possible to write*

$$\left[\tilde{\varphi}(\frac{k}{\sqrt{m}}) \right]^m = e^{-\frac{1}{2}k^2} \left[1 + \sum_{t=0}^{\infty} \omega_t(m) (i \,\mathrm{sgn}\, k)^{\alpha_t} |k|^{\nu_t} \right] , \tag{32}$$

where α_t are real numbers; the coefficients $\omega_t(m)$ are real; the powers ν_t satisfy $\nu_0 = 3$ and $\nu_{t+1} > \nu_t$. The ν_t are not necessarily integer. All the coefficients $\omega_t(m)$ are decreasing functions of m, decreasing as $m^{-\nu_t/6}$ or faster.

The probability of accusing user j is given by

$$R_m(\tilde{Z}) = \Omega(\tilde{Z}) + \frac{1}{\pi} \sum_{t=0}^{\infty} \omega_t(m) \Gamma(\nu_t) 2^{\nu_t/2} \mathrm{Im} \left[i^{-\alpha_t} H_{-\nu_t}(i\tilde{Z}/\sqrt{2}) \right] . \tag{33}$$

Here H is the Hermite function.

Proof: see Appendix B. □

The proof closely follows one of the standard proofs of the Central Limit Theorem. In the limit $m \to \infty$ all the coefficients ω_t vanish, leaving only the term $\Omega(\tilde{Z})$ which is the right-hand tail mass of the normal distribution.

For integer ν the function $H_{-\nu}$ reduces to the Hermite polynomial of order $\nu - 1$, multiplied by a factor $\exp(-\frac{1}{2}\tilde{Z}^2)$. (See Corollary 2.)

In Section 4.2 we determine the distribution φ. In Section 4.3 the Fourier transform $\tilde{\varphi}$ is computed and the leading order parameters ν_t, ω_t, α_t are derived.

b	Left tail	Right tail	$u=-0$	$u=+0$
1	$(\frac{1}{\lvert u\rvert})^{2c+1+2\kappa[q-1]}$	$(\frac{\mathbf{1}}{\mathbf{u}})^{\mathbf{5+2}\boldsymbol{\kappa}}$	$\lvert\mathbf{u}\rvert^{\mathbf{1+2}\boldsymbol{\kappa}}$	$u^{2c-3+2\kappa[q-1]}$
c	$(\frac{\mathbf{1}}{\lvert\mathbf{u}\rvert})^{\mathbf{3+2}\boldsymbol{\kappa}\mathbf{[q-1]}}$	$(\frac{1}{u})^{2c+3+2\kappa}$	$\lvert u\rvert^{2c-1+2\kappa}$	$\mathbf{u}^{\mathbf{-1+2}\boldsymbol{\kappa}\mathbf{[q-1]}}$

Table 1: Dominant powers in $\varphi(u)$ in the tails and close to $u=0$.

4.2 Distribution function of an innocent user's accusation

Theorem 5 *For an innocent user j, the distribution function φ of $S_j^{(i)}$ is given by*

$$\begin{aligned} u>0 \quad &: \quad \varphi_+(u) = \frac{2q}{B(\kappa,\kappa[q-1])}\sum_{b=1}^{c}\binom{c}{b}\frac{(u^2)^{\kappa[q-1]+c-b-\frac{1}{2}}}{(1+u^2)^{c+1+\kappa q}}K_b \\ u<0 \quad &: \quad \varphi_-(u) = \frac{2q}{B(\kappa,\kappa[q-1])}\sum_{b=1}^{c}\binom{c}{b}\frac{(u^2)^{\kappa+b-\frac{1}{2}}}{(1+u^2)^{c+1+\kappa q}}K_b. \end{aligned} \tag{34}$$

The proof is given in Appendix D. Note that all dependence on the strategy is contained in the numbers $K_b \in [0,1]$. Furthermore we see that the left tail and the right tail of $\varphi(u)$ have different power law behaviour. This is summarized in Table 1.

The right tail is dominated by the $b=1$ term; it is proportional to $(1/u)^{5+2\kappa}$. The left tail is dominated by the $b=c$ term, and is proportional to $(1/|u|)^{3+2\kappa q-2\kappa}$. It was found numerically in [18] that the 'optimal' κ (in terms of maximizing $\tilde{\mu}$) lies close to $1/q$; for such a choice of κ the left tail is heavier than the right tail.[4] Such a property is obviously beneficial for the *innocent*'s accusation. The discrepancy between the tails is even more pronounced if the attackers use the majority voting strategy (which for $q \geq 3$, $\kappa \approx 1/q$ minimizes $\tilde{\mu}$, as mentioned in Section 3). Then the right tail is dominated by the $b=\lceil c/q\rceil$ term, which behaves as $(1/u)^{3+2\lceil c/q\rceil+2\kappa}$, which for $c>q$ decreases even faster than $(1/u)^{5+2\kappa}$. From this perspective it may be better for the attackers not to use majority voting; another strategy may yield a form of the ρ curve that is better for them. The best strategy strikes a balance between decreasing $\tilde{\mu}$ and lengthening the tail of $\varphi_+(u)$.

In the binary case, it is easy to identify where the balance lies: For $\kappa \approx \frac{1}{2}$, the strategy has practically no effect on $\tilde{\mu}$, so the attackers should concentrate on lengthening the $\varphi_+(u)$ tail. This is achieved by setting Ψ_b nonzero for small values of b, e.g. interleaving or minority voting.

The behaviour of $\varphi(u)$ around $u=0$ is also noteworthy. For $u \uparrow 0$ the function is dominated by the $b=1$ contribution $|u|^{1+2\kappa}$, which has zero derivative at $u=0$. For $u \downarrow 0$ the $b=c$ term $u^{-1+2\kappa[q-1]}$ dominates; this one, however, has infinite derivative for $\kappa < 1/(q-1)$ (which is the case when e.g. $\kappa \approx 1/q$).

Corollary 6 *For an innocent user, the overall probability of positive and negative accusation are in general unequal, and are given by*

$$\begin{aligned} \Pr[u>0] &= q\sum_{b=1}^{c}K_b\mathbb{P}_1(b)\frac{b+\kappa}{c+\kappa q} \\ \Pr[u<0] &= q\sum_{b=1}^{c}K_b\mathbb{P}_1(b)\frac{c-b+\kappa[q-1]}{c+\kappa q}. \end{aligned} \tag{35}$$

[4] Notice that for $2\kappa[q-1]>1$ the absolute third moment exists: the integral $\int \mathrm{d}u\ |u|^3\varphi(u)$ is convergent in both tails. (As opposed to the binary case with $\kappa = 1/2$.) Consequently, there is a guaranteed convergence to the normal distribution when i.i.d. random variables $u_i \sim \varphi$ are added together in large numbers.

Proof: Follows by evaluating the u-integrals with Lemma 6, then applying Lemma 3 and finally rewriting the Beta functions using $B(x, y+1) = B(x,y)\frac{y}{x+y}$. □
Note that the probabilities properly add up to 1; this is readily seen from Lemma 4. Note too what happens when the colluders choose a majority voting strategy: then K_b tends to be small for small b and large for large b (see Section 5.1). The terms with large b then dominate the summations in Corollary 6, and consequently $\Pr[u > 0] > \Pr[u < 0]$. This is consistent with the fact that the left $(u < 0)$ tail is heavier: the probability mass at $u < 0$ must be further removed from $u = 0$ in order to cause $\mathbb{E}[u] = 0$.

Corollary 7 *If the colluder strategy is the interleaving attack, $\theta_{y|\boldsymbol{\sigma}} = \frac{\sigma_y}{c}$, then*

$$\begin{aligned} \varphi_+(u) &= \frac{2q}{B(\kappa, \kappa[q-1])} \frac{(u^2)^{\kappa[q-1]-\frac{1}{2}}}{(1+u^2)^{2+\kappa q}} \\ \varphi_-(u) &= \frac{2q}{B(\kappa, \kappa[q-1])} \frac{(u^2)^{\kappa+\frac{1}{2}}}{(1+u^2)^{2+\kappa q}}, \end{aligned}$$

and $\Pr[u > 0] = \frac{\kappa+1}{\kappa q+1}$*,*

Proof: The first part follows directly by applying Lemma 5 to (34) and using $\sum_{b=0}^{c} \binom{c}{b} b x^b = xc(1+x)^{c-1}$. The second part follows from taking the integral $\int_0^\infty \mathrm{d}u\ \varphi_+(u)$. □
It is interesting to note that the interleaving attack yields a $\varphi(u)$ distribution that has the heaviest possible tails for both positive and negative u (see Table 1): proportional to $(1/|u|)^{3+2\kappa[q-1]}$ for the left tail and $(1/u)^{5+2\kappa}$ for the right tail. It also has the lowest possible dominant powers around $u = 0$. Furthermore, $\varphi(u)$ has the special property that it is completely independent of c.

4.3 The Fourier transform of φ

We compute the Fourier transform of $\varphi(u)$ using the following lemma.

Lemma 13 (From [14], section 2.5.9) *Let $k \in \mathbb{R}$, $\mathrm{Re}\ v > -\frac{1}{2}$, and $d > 0$. Let the function Λ be defined as the following convergent integral,*

$$\Lambda(d, v; k) := \int_0^\infty \mathrm{d}u\ \frac{u^{2d-1}}{(u^2+1)^{v+d}} e^{iku}.$$

This integral is expressed in terms of hypergeometric ${}_1F_2$ functions as

$$\begin{aligned} \Lambda(d, v; k) &= (-ik)^{2v}\Gamma(-2v)\ {}_1F_2(v+d; v+\tfrac{1}{2}, v+1; \frac{k^2}{4}) + \tfrac{1}{2}\sum_{j=0}^{\infty} \frac{(ik)^j}{j!} B(d+\frac{j}{2}, v-\frac{j}{2}) \\ &= (-ik)^{2v}\Gamma(-2v)\ {}_1F_2(v+d; v+\tfrac{1}{2}, v+1; \frac{k^2}{4}) \\ &\quad + \tfrac{1}{2}B(d,v)\ {}_1F_2(d; \tfrac{1}{2}, 1-v; \frac{k^2}{4}) + \frac{ik}{2}B(d+\tfrac{1}{2}, v-\tfrac{1}{2})\ {}_1F_2(d+\tfrac{1}{2}; \tfrac{3}{2}, \tfrac{3}{2}-v; \frac{k^2}{4}). \end{aligned}$$

Notice that in general $\Lambda(d, v; k)$ is not an entire function of k due to the appearance of the factor k^{2v} in the first term, which for general v is not an entire function.
The hypergeometric function ${}_1F_2$ has the sum representation ${}_1F_2(\alpha; \beta_1, \beta_2; z) = \sum_{j=0}^{\infty} \frac{(\alpha)_j}{j!(\beta_1)_j(\beta_2)_j} z^j$ where $(\alpha)_j = \alpha(\alpha+1)\cdots(\alpha+j-1)$ is the Pochhammer symbol. The radius of convergence is infinity. The ${}_1F_2$ function can be evaluated by using software packages such as Mathematica.

Theorem 6 *The Fourier transform of φ is given by*

$$\tilde{\varphi}(k) = \frac{2q}{B(\kappa, \kappa[q-1])} \sum_{b=1}^{c} \binom{c}{b} K_b \cdot \Big[\Lambda(d_b, v_b; k) + \Lambda(D_b, V_b; -k)\Big],$$

with Λ *as defined in Lemma 13, and*

$$\begin{aligned} d_b = b+\kappa \quad &; \quad v_b = c-b+\kappa[q-1]+1 \\ D_b = c-b+\kappa[q-1] \quad &; \quad V_b = b+\kappa+1. \end{aligned} \tag{36}$$

Proof: The Fourier transform is defined as $\tilde{\varphi}(k) = \int_{-\infty}^{\infty}\mathrm{d}u\ \varphi(u)e^{-iku}$. We use the expression for φ given in Theorem 5. The integral for the summands in φ_+ is immediately of the form appearing in Lemma 13 and yields $\Lambda(D_b, V_b; -k)$. The integral over the φ_- terms is of the form $\int_{-\infty}^{0}\mathrm{d}u\ f(u^2)e^{-iku}$, which can be rewritten as $\int_0^{\infty}\mathrm{d}u\ f(u^2)e^{iku}$; this has the form of the integral in Lemma 13 and yields $\Lambda(d_b, v_b; k)$. □

For $q \geq 3$ and realistic κ, none of the values d_b, v_b, D_b, V_b in (36) is integer or half-integer. Hence substitution into all the Gamma functions and Pochhammers contained in the ${}_1F_2$ functions of Lemma 13 is well defined. Note that, given the summation range $1 \leq b \leq c$, the smallest possible value of v_b or V_b is $v_c = 1 + \kappa[q-1] > 1$. Hence, in a power series expansion for small k, the k^{2v} term in (36) always comes 'after' the k^3 power. In fact, for $q \geq 3$ and $\kappa \approx 1/q$ we have $2v_c \in (3, 4)$.

Corollary 8 *For* $q \geq 3$ *the leading order terms in the expansion of* $\tilde{\varphi}(k)$ *are given by*

$$\begin{aligned} \tilde{\varphi}(k) \quad = \quad & 1 - \tfrac{1}{2}k^2 + \frac{2q}{B(\kappa, \kappa[q-1])}\Bigg\{ \frac{(ik)^3}{2\cdot 3!}\sum_{b=1}^{c} K_b[B(d_b + \tfrac{3}{2}, v_b - \tfrac{3}{2}) - B(D_b + \tfrac{3}{2}, V_b - \tfrac{3}{2})] \\ & +(-ik)^{2+2\kappa[q-1]}\Gamma(-2-2\kappa[q-1]) \\ & +\frac{(ik)^4}{2\cdot 4!}\sum_{b=1}^{c} K_b[B(d_b+2, v_b-2) - B(D_b+2, V_b-2)] + (ik)^{4+2\kappa}K_1\Gamma(-4-2\kappa)\Bigg\}. \\ & +\cdots \end{aligned}$$

Proof: Follows by substituting the first expression for Λ from Lemma 13 into Theorem 6, and then cutting off the small-argument power series of the ${}_1F_2$ function after the k^0 term. □

Corollary 9 *If the colluders use the interleaving attack, then*

$$\begin{aligned} \tilde{\varphi}_{\text{inter}}(k) \quad = \quad & 1 - \tfrac{1}{2}k^2 + \frac{2q}{B(\kappa, \kappa[q-1])}\Bigg[(ik)^{4+2\kappa}\Gamma(-4-2\kappa){}_1F_2(\kappa q; \kappa + \tfrac{5}{2}, \kappa+3; \tfrac{k^2}{4}) \\ & +(-ik)^{2+2\kappa[q-1]}\Gamma(-2-2\kappa[q-1]){}_1F_2(\kappa q; \kappa[q-1]+\tfrac{3}{2}, \kappa[q-1]+2; \tfrac{k^2}{4}) \\ & +\tfrac{1}{2}\sum_{j=3}^{\infty}\frac{(ik)^j}{j!}[B(\kappa+1+\tfrac{j}{2}, \kappa[q-1]+1-\tfrac{j}{2}) + (-1)^j B(\kappa[q-1]+\tfrac{j}{2}, \kappa+2-\tfrac{j}{2})]\Bigg]. \end{aligned}$$

Proof: The Fourier integrals of the φ_+ and φ_- given in Corollary 7 are precisely of the form handled in Lemma 13, with $(d = \kappa[q-1], v = \kappa+2)$ and $(d = \kappa+1, v = \kappa[q-1]+1)$ respectively. □

5 Numerics for the majority voting strategy

We first present a fast algorithm for computing the K_b parameters in the case of majority voting. Then we show numerical results for the minimum code length required to resist a coalition of c_0 attackers who use the majority voting strategy.

5.1 Computing $\mathbf{K_b}$ for majority voting

Lemma 14 *Let the colluder strategy be majority voting. Let* $N_b \in \mathbb{N}$ *with* $N_b > \max\{c-b, bq-c\}$*, and let* t_b *and* G_{ba} *be defined as*

$$t_b = e^{i2\pi/N_b} \quad ; \quad G_{ba} = \sum_{x=0}^{b-1}\frac{\Gamma(\kappa+x)}{x!}t_b^{ax}. \tag{37}$$

Then K_b is given by

$$b < \frac{c}{q} \quad : \quad K_b = 0 \tag{38}$$

$$\frac{c}{q} \leq b < \frac{c}{2} \quad : \quad K_b = \frac{b!(c-b)!}{\Gamma(c-b+\kappa[q-1])\Gamma(b+\kappa)B(\kappa\mathbf{1}_{q-1})} \cdot \frac{1}{qN_b}\sum_{a=0}^{N_b-1} t_b^{-ac}(G_{ba})^q \cdot \left\{\left(1+\frac{\Gamma(b+\kappa)t_b^{ab}}{b!G_{ba}}\right)^q - 1\right\}. \tag{39}$$

$$b = \frac{c}{2} \quad : \quad K_{c/2} = 1 - \frac{q-1}{2}\frac{B(\kappa\mathbf{1}_{q-1}+\frac{c}{2}\boldsymbol{e}_1)}{B(\kappa\mathbf{1}_{q-1})} = 1 - \tfrac{1}{2}\frac{(1+\kappa)_{c/2-1}}{(1+\kappa[q-1])_{c/2-1}} \tag{40}$$

$$b > \frac{c}{2} \quad : \quad K_b = 1. \tag{41}$$

The proof is given in Appendix C.

These expressions look very complicated. However, they are easier to evaluate numerically than (24). Evaluation of (41) involves only two sums: for every a, the G_{ba} has fewer than $c/2$ terms, and the a-sum has N_b terms, with $N_b = \mathcal{O}(cq/2)$. The total number of terms is $\mathcal{O}(c^2q/4)$. Direct evaluation of (24) on the other hand involves a $(q-1)$-dimensional sum with $\mathcal{O}([c/2]^{q-1})$ terms, a higher power of c when $q > 3$.

Note that a large number N can be chosen that satisfies $N > \max\{c-b, bq-c\}$ for all $c/q \leq b < c/2$. Then all the N_b values in (39) can be set to N. The price one pays for this small simplification is that the sums contain more terms.

5.2 Behaviour of $\mathbf{R_m(\tilde{Z})}$ for majority voting

From all the results in the previous sections, the false accusation probability for a fixed innocent user, as a function of q, κ, c, and m, is numerically computed as follows (assuming $\varepsilon_2 \approx 1/2$). The K_b parameters are evaluated using Lemma 14. A power series expansion for $x = \tilde{\varphi}(k) - 1$ is obtained from Theorem 6. It is substituted in the series expansion of $\ln(1+x)$. Then k is replaced by $k/\sqrt{m}$ and the whole expression is multiplied by m, yielding a power series for $m \ln \tilde{\varphi}(k/\sqrt{m})$. The first term, $-\frac{1}{2}k^2$, is split off, and the rest is substituted into the power series of the exp function. The resulting series precisely yields the powers ν_t, 'angles' α_t and coefficients $\omega_t(m)$ as defined in (32). These are then used in (33) to obtain the final result.

Fig. 3 shows a typical example of the shape of the resulting curve. For low values of $\tilde{Z}$ the curve lies below the Guassian tail integral $\Omega(\tilde{Z})$, meaning that the Guassian approximation is actually pessimistic there! Then at some point the curve crosses $\Omega(\tilde{Z})$ and becomes a power-law tail.

We will use the notation $m_{\text{cross}}(\varepsilon_1)$ for the value of m where the crossover point $R_m(\tilde{Z}) = \Omega(\tilde{Z})$ lies exactly at $\Omega(\tilde{Z}) = \varepsilon_1$. For $m \geq m_{\text{cross}}(\varepsilon_1)$, the Gaussian approximation is valid (and even pessimistic) for false accusation probabilities up to ε_1. Note that $m_{\text{cross}}(\varepsilon_1)$ depends on c, q, κ and the pirate strategy. In the case of the majority voting attack, we find that m_{cross} decreases with c. This happens because the K_b parameters for majority voting (Lemma 14) kick in only at $b \geq c/q$, with $K_b = 0$ for $b < c/q$. From (34) we see that the $b = c/q$ term in $\varphi(u)$, which then is the heaviest of the contributions to the right tail, behaves as $(1/u)^{3+2\kappa+2c/q}$. Thus, the right tail becomes less heavy with increasing c, facilitating convergence to the Gaussian form.

We also find that m_{cross} increases with q. This can be understood from the same reasoning as above. The main contribution to the right tail, $(1/u)^{3+2\kappa+2c/q}$, is an increasing function of q.

It is important to remark on the number of terms that should be kept in the power series. Some general, unsurprising rules of thumb apply. For an accurate result, more terms need to be kept when $\tilde{Z}$ is increased and when m is decreased. For $m < 100$, powers larger than k^{50} are required, with rather long computation times. Less obviously, the crossover region sometimes needs more

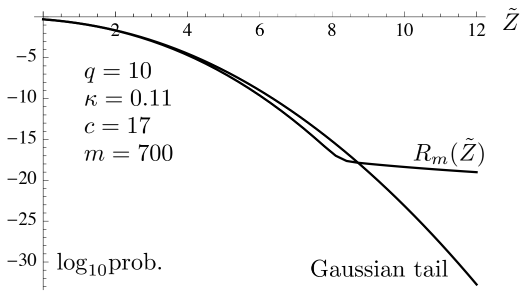

Figure 3: *Logarithmic plot of the probability $R_m(\tilde{Z})$ of accusing a fixed innocent user, as a function of the scaled threshold $\tilde{Z}$, for the majority voting attack and with parameter settings as listed in the graph.*

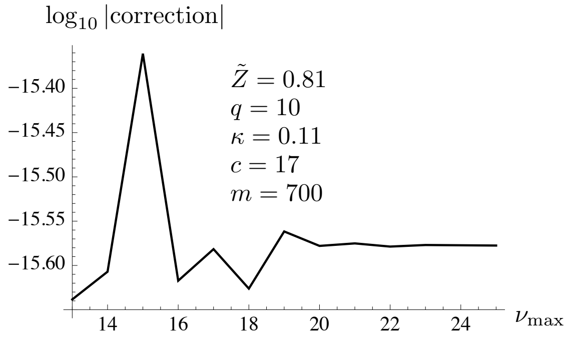

Figure 4: *Logarithmic plot of the correction to $\Omega(\tilde{Z})$ as a function of $\nu_{\max}$, the maximum power of k kept in the expansion.*

terms than other values of $\tilde{Z}$. For example, the curve in Fig. 3 requires powers up to k^{20} to get a converging result around $\tilde{Z} = 8$. An example of such convergence is shown in Fig. 4.

5.3 Sufficient code lengths for majority voting

Table 2 shows sufficient code lengths against colluders who use the majority voting strategy. The crossover values m_{cross} are also listed. We take parameters: $\varepsilon_2 \approx \frac{1}{2}$, $\kappa \approx 1/q$, with $\kappa > 1/q$. The sufficient code length m_* as a function of q, κ, c, ε_1 was determined as follows. We numerically solved the equation

$$R_m^{\text{inv}}(\varepsilon_1) = \tilde{\mu}_{\text{maj}} \frac{\sqrt{m}}{c} \tag{42}$$

for m, where $\tilde{\mu}_{\text{maj}}$ is the statistical parameter $\tilde{\mu}$ computed according to (29) for the majority voting strategy.[5] The solution gives the smallest possible value for m such that there exists a $\tilde{Z}$ satisfying $R_m(\tilde{Z}) = \varepsilon_1$ as well as $\tilde{Z} \leq \tilde{\mu}_{\text{maj}}\sqrt{m}/c$. The latter condition is required in order to have $\Pr[\text{FN}] \leq \frac{1}{2}$. (See the guilty curve in Fig. 1.)

Table 2 gives the solution m_* as well as the $\tilde{Z}$ value at the solution ($\tilde{Z}_*$), and the crossover[6] value m_{cross} as defined in Section 5.2. The proportionality constant in the relation $m_* \propto c^2 \ln(1/\varepsilon_1)$ is

[5] Note that $\tilde{\mu}_{\text{maj}}$ is only slightly larger than the $\tilde{\mu}$ of the 'optimal' $\tilde{\mu}$-reducing strategy discussed in [18], because our choice $\kappa \approx 1/q$ implies that majority voting is very close to optimal. Also note that $\tilde{\mu}_{\text{maj}}$ weakly depends on c, but is independent of m.

[6] An entry like '< 100' means that high powers of k are required in the series expansion in order to determine m_{cross} more accurately, and we did not invest the necessary time.

MAJORITY VOTING; $\varepsilon_1 = 10^{-10}$; $\varepsilon_2 \approx 0.5$						
q	κ	c	m_*	$\tilde{Z}_*$	$\frac{m_*}{c^2 \ln 1/\varepsilon_1}$	m_{cross}
3	0.34	3	$1.29 \cdot 10^3$	10.4	6.22	$11 \cdot 10^3$
		4	$7.5 \cdot 10^2$	5.91	2.04	$3 \cdot 10^2$
		5	$1.19 \cdot 10^3$	5.97	2.07	$3 \cdot 10^2$
		7	$2.41 \cdot 10^3$	6.06	2.14	$3 \cdot 10^2$
		20	$2.09 \cdot 10^4$	6.24	2.27	< 300
		80	$3.44 \cdot 10^5$	6.33	2.33	< 300
10	0.105	3	$2.48 \cdot 10^3$	21.3	12.0	$9 \cdot 10^5$
		5	$1.90 \cdot 10^3$	10.8	3.30	$3 \cdot 10^4$
		6	$1.26 \cdot 10^3$	7.30	1.52	$4 \cdot 10^3$
		7	$1.25 \cdot 10^3$	6.22	1.11	$4 \cdot 10^2$
		11	$3.16 \cdot 10^3$	6.24	1.13	< 100
		20	$1.07 \cdot 10^4$	6.29	1.16	< 100
		80	$1.75 \cdot 10^5$	6.34	1.19	< 100
16	0.066	3	$2.8 \cdot 10^3$	24.1	14	$3 \cdot 10^6$
		5	$2.36 \cdot 10^3$	12.73	4.10	$2 \cdot 10^5$
		6	$1.68 \cdot 10^3$	8.89	2.03	$2 \cdot 10^4$
		7	$1.20 \cdot 10^3$	6.42	1.06	$1.3 \cdot 10^3$
		80	$1.59 \cdot 10^5$	6.34	1.08	< 100

Table 2: *Sufficient code lengths for various alphabet and coalition sizes. The normal distribution has* $\Omega(\tilde{Z}) = 10^{-10}$ *at* $\tilde{Z} = 6.36$*.*

also shown.
Several conclusions can be drawn from the table.

- For very small coalitions the Gaussian approximation does not hold, e.g. $(q = 3, c \leq 3)$, $(q = 10, c \leq 6)$, $(q = 16, c \leq 7)$.
- Even then a decent code length $m_* \ll m_{\text{cross}}$ can often be achieved, e.g. $(q = 10, c = 5$ and $c = 6)$, $(q = 16, c = 6$ and $c = 7)$. This is possible because the R_m curve still quickly descends as a function of $\tilde{Z}$ even when $\tilde{Z}$ lies to the right of the crossover point.
- For large coalitions the Gaussian approximation holds. The proportionality constant in $m_* \propto c^2 \ln(1/\varepsilon_1)$ has a minimum as a function of c where the Gaussian regime sets in. With growing c, the $\tilde{Z}_*$ approaches 6.36, which is the value at which $\Omega(\tilde{Z}) = 10^{-10}$.

Remark: This is not the final word on the majority voting attack. Better results can probably be achieved with different choices of κ. This is left for future work.

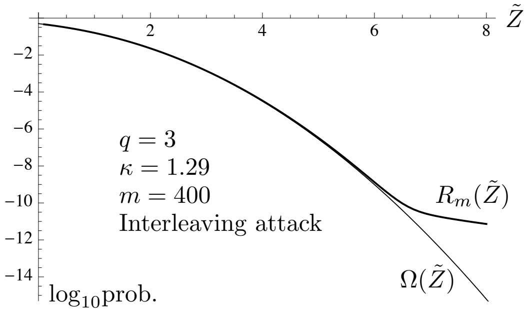

Figure 5: *Logarithmic plot of the probability $R_m(\tilde{Z})$ of accusing a fixed innocent user, as a function of the scaled threshold $\tilde{Z}$. Parameter settings as listed in the graph.*

6 Numerics for the interleaving strategy

6.1 Behaviour of $\mathbf{R_m(\tilde{Z})}$ for the interleaving attack

False positive probabilities were computed as described in Section 5.2, except for two differences: (i) The starting point for the power series in k is Corollary 9, so there is no need to compute the K_b parameters. (ii) The shape of R_m now does not depend on c.

An example is shown in Fig. 5. We have observed for $q \geq 3$ that a crossing point of the R_m and Ω curve as in Fig. 3 can occur for small κ (e.g. $q = 3$, $\kappa = 0.34$, $m = 10^4$). However, we mostly studied somewhat larger κ than in Section 5, in order to obtain shorter codes, and for these there were no crossings.

As a general rule we have observed that increasing q worsens the convergence to the Gaussian limit. We conjecture that this is caused by the faster dwindling left tail, $(1/|u|)^{3+2\kappa[q-1]}$, while the right tail remains equally heavy.

6.2 Sufficient code length for the interleaving attack

Theorem 7 *For the interleaving strategy, the $\tilde{\mu}$ parameter becomes*

$$\tilde{\mu}_{\text{inter}} = q \frac{B(\kappa + \frac{1}{2}, \kappa[q-1] + \frac{1}{2})}{B(\kappa, \kappa[q-1])}. \tag{43}$$

Proof: From the definition of $\tilde{\mu}$ it follows that it can be computed as an expectation value in a single content segment, $\tilde{\mu} = \mathbb{E}[\sigma_y g_1(p_y) + (c - \sigma_y) g_0(p_y)]$, with $\mathbb{E}$ the expectation over $\boldsymbol{p}$, $\boldsymbol{\sigma}$ and y, and g_1 and g_0 as defined in (8). The $\mathbb{E}_y(\cdots)$ expectation is given by $\sum_y \frac{\sigma_y}{c}(\cdots)$. We write

$$\frac{\sigma_y}{c} \left[\sigma_y g_1(p_y) + (c - \sigma_y) g_0(p_y)\right] = p_y \frac{\sigma_y - c p_y}{\sqrt{p_y(1-p_y)}} + \frac{1}{c} \frac{(\sigma_y - c p_y)^2}{\sqrt{p_y(1-p_y)}}. \tag{44}$$

From the properties of the multinomial distribution we get $\mathbb{E}_{\boldsymbol{\sigma}}[\sigma_y - c p_y] = 0$ and $\mathbb{E}_{\boldsymbol{\sigma}}[(\sigma_y - c p_y)^2] = c p_y (1 - p_y)$. Next, the expectation $\mathbb{E}_{\boldsymbol{p}}$ over the full vector $\boldsymbol{p}$ reduces to the expectation over the component p_y, for which we use the marginal pdf $f(p)$ (Lemma 2). This gives

$$\tilde{\mu}_{\text{inter}} = \sum_y \frac{1}{B(\kappa, \kappa[q-1])} \int_0^1 \mathrm{d}p_y \; p_y^{-1+\kappa} (1 - p_y)^{-1+\kappa[q-1]} \sqrt{p_y(1-p_y)}. \tag{45}$$

The result of the integration does not depend on y, so the $\sum_y$ yields a factor q. The integral yields $B(\kappa + \frac{1}{2}, \kappa[q-1] + \frac{1}{2})$. □

Fig. 6 shows the effect of $\tilde{\mu}_{\text{inter}}$ on the code length for various q and κ. For the interleaving attack, the factor $2/\tilde{\mu}^2$, which appears as a multiplier in the Gaussian limit expression (13) for the code

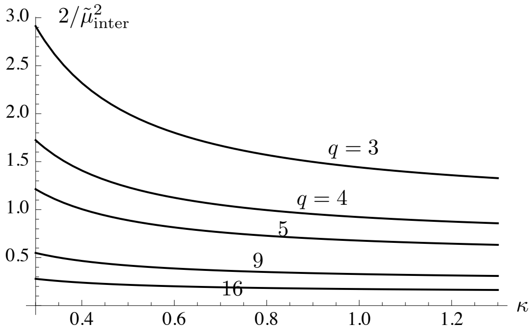

Figure 6: *The factor $2/\tilde\mu^2_{\rm inter}$ (see Eq. (13)) as a function of κ for various alphabet sizes.*

length, is a decreasing function of κ and q. Increasing the alphabet has a large impact when q is small, but very little impact when q is large.

In the case of the interleaving attack, Eq. (42) for finding the sufficient code length m_* has more structure than in the case of majority voting. To be completely explicit about the dependence on the various variables we write $\tilde\mu_{\rm inter}(q,\kappa)$ and $R_{q\kappa m}(\tilde Z)$. Since $\tilde\mu_{\rm inter}$ and $R_{q\kappa m}$ do not depend on c, it makes sense to isolate c and reorganize (42) as

$$c = \sqrt{m}\frac{\tilde\mu_{\rm inter}(q,\kappa)}{R^{\rm inv}_{q\kappa m}(\varepsilon_1)}. \tag{46}$$

This equation gives an upper bound on the coalition size that can be resisted by the code. The easiest way to handle the numerics is to choose (for fixed q, κ, ε_1) a set of values for m, and then compute $\tilde Z_*$ and c as a function of m. (The results for c are not integer in general, but it is implicitly understood that they should be rounded down.) We therefore present our results in a slightly different form than in Section 5.3.

Fig. 7 shows how $\tilde Z_*$ and m_* converge to their Gaussian limits as a function of the code length. The c on the horizontal axis is a parametrization of m, representing the coalition size that can be resisted by the code. The limiting value for $\tilde Z_*$ is $\Omega^{\rm inv}(\varepsilon_1)$. The limiting value for m_* is $m_{\rm limit} = [c\Omega^{\rm inv}(\varepsilon_1)/\tilde\mu]^2$. We have plotted the fraction $m_*/m_{\rm limit} = [\tilde Z_*/\Omega^{\rm inv}(\varepsilon_1)]^2$.

Note that the factor $[\Omega^{\rm inv}(\varepsilon_1)]^2$ in the expression for $m_{\rm limit}$ is noticeably smaller than the bound $2\ln(1/\varepsilon_1)$. This means that the code can be made even shorter. The ratio $[\Omega^{\rm inv}(\varepsilon_1)]^2/[2\ln(1/\varepsilon_1)]$ is plotted in Fig. 8. Fig. 9 shows the familiar code length proportionality constant $m/(c^2\ln\varepsilon_1^{-1})$.

The case of the binary alphabet ($q=2$) is rather special. If κ is set to $\frac12$, then the left tail of $\varphi(u)$ becomes so heavy that $\mathbb{E}(|u|^3)=\infty$, severely hampering convergence to the Gaussian limit. Tardos [17] introduced a cutoff parameter $t \ll 1$ so that $p_\alpha \in (t, 1-t)$, which curbs the tail, yielding $\mathbb{E}(|u|^3)<\infty$. (Tardos did not formulate it in this way; for him it was a technical trick that allows for the use of the Markov inequality in a crucial part of a security proof.) We do not set κ exactly to $\frac12$ and we do not use the cutoff t, but instead we consider $\kappa \geq 0.55$. This is close enough to get a good impression of the behaviour of the original Tardos code, but large enough to get numerical results quickly. For κ closer to $\frac12$ our method requires many more powers of k to be kept, leading to long computation times. We observe a difference between $q=2$ and $q\geq 3$. In the binary case, the results are better than Gaussian in a large portion of parameter space, and already at small coalition sizes. For $q\geq 3$, the Gaussian limit is approached 'from the other side', i.e. with results that are worse than the Gaussian limit.

From the numerics we conclude that the attack vs. defense game is quite complex. In the asymptotic limit, the $\tilde\mu$-minimizing strategy of [18] is the best attack; the best defense was shown to be setting κ a bit larger than $1/q$; in that regime the attack is basically majority voting. In the small c regime the interleaving attack is a potent strategy. It can be effectively defended against by choosing κ as large as possible; this facilitates convergence to the Gaussian limit and at the same time increases $\tilde{\mu}$. However, κ cannot be increased indefinitely, for otherwise the defense against other attacks becomes too weak. (The $\tilde{\mu}$-minimizing attack of [18] becomes too powerful.) Finding a balance between these effects is left for future work.

7 Summary and future work

We have analyzed the q-ary Tardos fingerprinting scheme in the restricted digit model. We have introduced a new parametrization $\Psi_b(\boldsymbol{x})$ of the attack strategy. It has the advantage that it no longer depends on any symbol index $\alpha \in \mathcal{Q}$; furthermore, it allows for pre-computation of the parameters $K_b = \mathbb{E}_{\boldsymbol{x}|b}\Psi_b(\boldsymbol{x})$

We have shown for $\kappa q \approx 1$ that the majority voting strategy minimizes $\tilde{\mu}$. We have determined the probability distribution of the accusation of an innocent user due to a single content segment. Using the Fourier approach we have used this to set up a series expansion for the systematic computation of the total accusation probability for an innocent user. As a first test of our method we have numerically evaluated our expansions for $\varepsilon_1 = 10^{-10}$ and various parameter settings. We have done this for two attacks that are of special interest, the majority voting attack and the interleaving attack. We have found that the 'shape' parameter κ plays a crucial role. When κ is chosen so as to maximize $\tilde{\mu}$ in the face of a $\tilde{\mu}$-reducing attack, then convergence to the Gaussian limit is quite bad, especially for large alphabets. Increasing κ dramatically improves the convergence. At the same time the $\tilde{\mu}$ decreases; hence, the game of attack and defense is quite complex, involving the ratio of $R_m^{\text{inv}}(\varepsilon_1)$ and $\tilde{\mu}$ instead of a single one of these parameters. A full study of general attacks, for different ε_1, is left for future work.

It would be interesting to see if the approach developed here can be applied to accusation probabilities in the joint decoder scenario.

Acknowledgements

We kindly thank Benne de Weger, Dion Boesten, Jan-Jaap Oosterwijk and Guido Janssen for useful discussions.

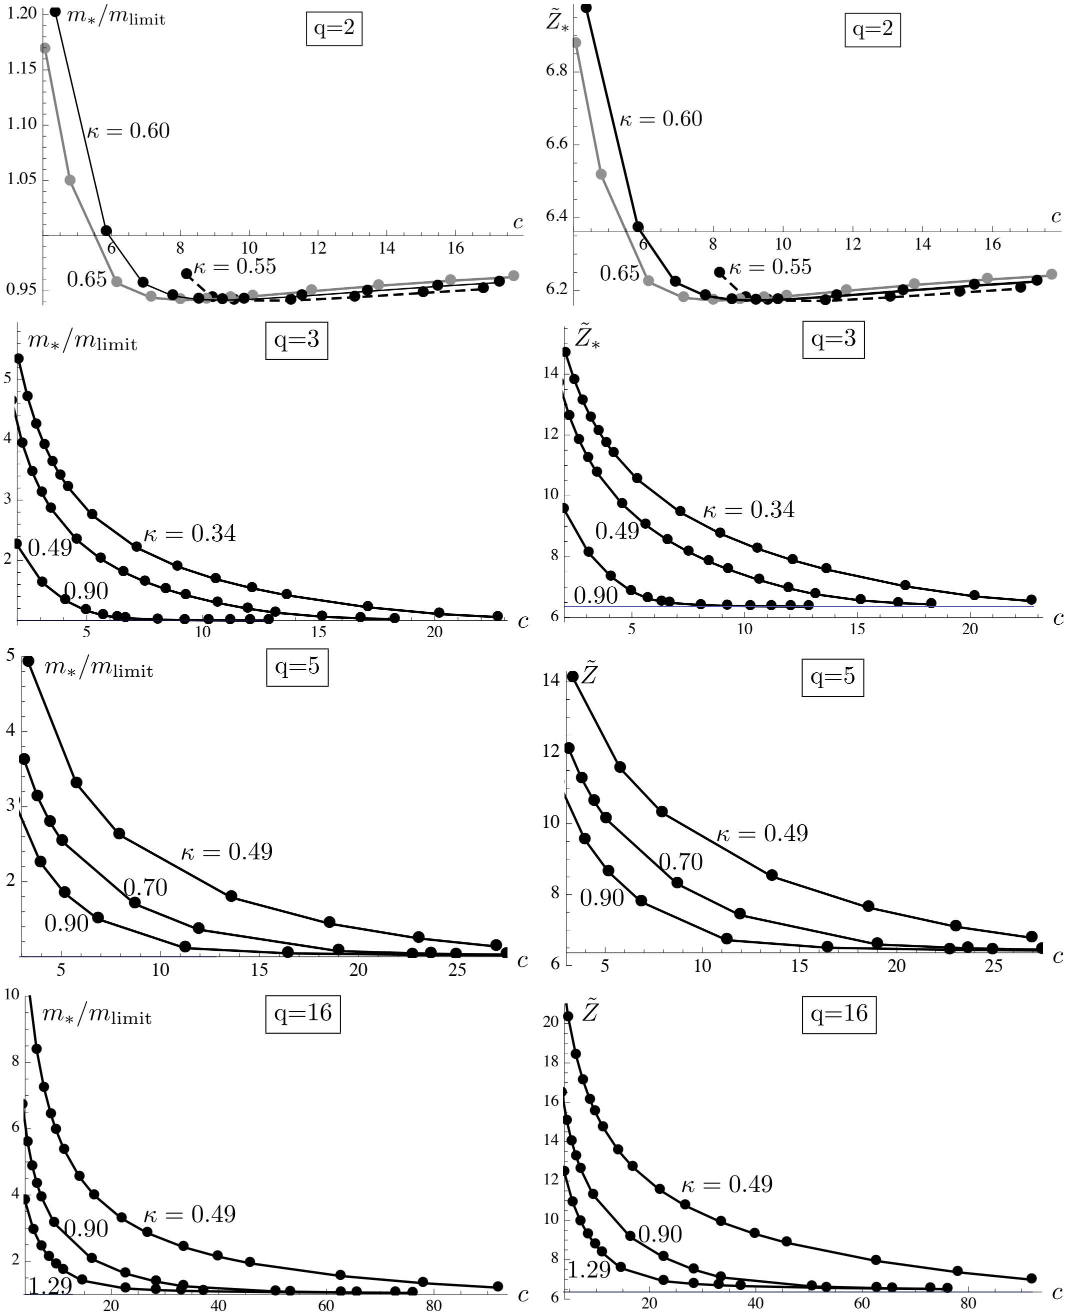

Figure 7: *Convergence to the Gaussian limit for the interleaving attack, for $\varepsilon_1 = 10^{-10}$.* **Left:** *Code length m_* compared to the Gaussian value $[c\Omega^{\text{inv}}(\varepsilon_1)/\tilde{\mu}]^2$, as a function of the coalition size c.* **Right:** *$\tilde{Z}_*$ as a function of c.*

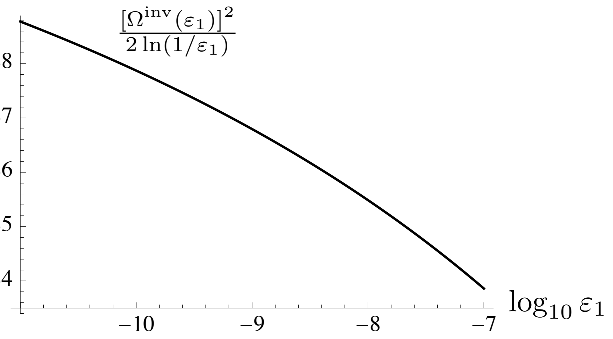

Figure 8: *The factor* $[\Omega^{\rm inv}(\varepsilon_1)]^2$ *compared to the bound* $2\ln(1/\varepsilon_1)$.

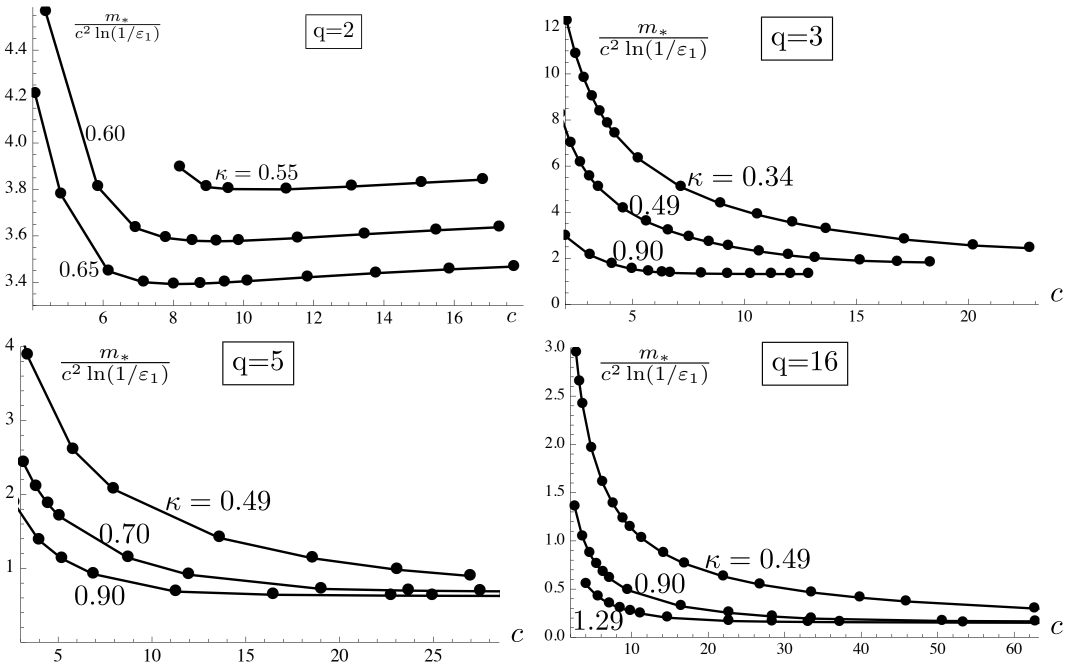

Figure 9: *Interleaving attack. The often studied proportionality constant in* $m_* \propto c^2 \ln\frac{1}{\varepsilon_1}$*, as a function of* c*, for various* q *and* κ*.*

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] O. Blayer and T. Tassa. Improved versions of Tardos' fingerprinting scheme. *Designs, Codes and Cryptography*, 48(1):79–103, 2008.

[3] D. Boneh and J. Shaw. Collusion-secure fingerprinting for digital data. *IEEE Transactions on Information Theory*, 44(5):1897–1905, 1998.

[4] T. Furon, A. Guyader, and F. Cérou. 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.

[5] T. Furon, L. Pérez-Freire, A. Guyader, and F. Cérou. Estimating the minimal length of Tardos code. In *Information Hiding 2009*, volume 5806 of *Lecture Notes in Computer Science*, pages 176–190.

[6] I.S. Gradshteyn and I.M. Ryzhik. *Table of Integrals, Series, and Products, 5th edition.* Academic Press, 1994.

[7] S. He and M. Wu. Joint coding and embedding techniques for multimedia fingerprinting. *TIFS*, 1:231–248, June 2006.

[8] Y.W. Huang and P. Moulin. Saddle-point solution of the fingerprinting capacity game under the marking assumption. In *Proc. IEEE International Symposium on Information Theory (ISIT)*, 2009.

[9] Y.W. Huang and P. Moulin. Saddle-point solution of the fingerprinting capacity game under the marking assumption, 2009. `http://arxiv.org/abs/0905.1375`.

[10] J. Kilian, F.T. Leighton, L.R. Matheson, T.G. Shamoon, R.E. Tarjan, and F. Zane. Resistance of digital watermarks to collusive attacks. In *IEEE International Symposium on Information Theory (ISIT) 1998*, page 271.

[11] M. Kuribayashi, N. Akashi, and M. Morii. On the systematic generation of Tardos's fingerprinting codes. In *International Workshop on Multimedia Signal Processing (MMSP) 2008*, pages 748–753.

[12] P. Moulin. Universal fingerprinting: Capacity and random-coding exponents. In *Preprint arXiv:0801.3837v2, avilable at `http://arxiv.org/abs/0801.3837`*, 2008.

[13] K. Nuida, M. Hagiwara, H. Watanabe, and H. Imai. Optimal probabilistic fingerprinting codes using optimal finite random variables related to numerical quadrature. *CoRR*, abs/cs/0610036, 2006.

[14] A.P. Prudnikov, Yu.A. Brychkov, and O.I. Marichev. *Integrals and Series, 4th printing*, volume 1. CRC, 1998.

[15] H.G. Schaathun. On error-correcting fingerprinting codes for use with watermarking. *Multimedia Systems*, 13(5-6):331–344, 2008.

[16] A. Somekh-Baruch and N. Merhav. On the capacity game of private fingerprinting systems under collusion attacks. *IEEE Trans. Inform. Theory*, 51:884–899, March 2005.

[17] G. Tardos. Optimal probabilistic fingerprint codes. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 116–125, 2003.

[18] 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.

[19] 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.

A Proof of Theorem 3

We have $\Pr[S_j > Z] = \Pr\left[\sum_{i=1}^m S_j^{(i)} > Z\right]$ for innocent j. The terms $S_j^{(i)}$ are independent, identically distributed random variables. This allows us to write

$$\Pr[S_j > Z] = \int_{-\infty}^{\infty} \mathrm{d}u_1 \varphi(u_1) \cdots \int_{-\infty}^{\infty} \mathrm{d}u_m \varphi(u_m)\; \Theta(u_1 + \cdots + u_m - Z). \tag{47}$$

Here Θ is the Heaviside step function. Next we use a well known integral representation of the step function,

$$\Theta(x) = \lim_{\eta\downarrow 0} \frac{1}{2\pi i} \int_{-\infty}^{\infty} \mathrm{d}\lambda\; \frac{e^{i\lambda x}}{\lambda - i\eta}. \tag{48}$$

Substituting (48) into (47) and rearranging the order of the integrations, we get

$$\begin{aligned} \Pr[S_j > Z] &= \lim_{\eta\downarrow 0} \int_{-\infty}^{\infty} \frac{\mathrm{d}\lambda}{2\pi i}\; \frac{e^{-i\lambda Z}}{\lambda - i\eta} \prod_{a=1}^{m} \left[\int_{-\infty}^{\infty} \mathrm{d}u_a\; \varphi(u_a) e^{i\lambda u_a}\right] \\ &= \lim_{\eta\downarrow 0} \int_{-\infty}^{\infty} \frac{\mathrm{d}\lambda}{2\pi i}\; \frac{e^{-i\lambda Z}}{\lambda - i\eta} \left[\tilde{\varphi}(-\lambda)\right]^m = -\lim_{\eta\downarrow 0} \int_{-\infty}^{\infty} \frac{\mathrm{d}k}{2\pi i}\; \frac{e^{ikZ/\sqrt{m}}}{k + i\eta} \left[\tilde{\varphi}(\frac{k}{\sqrt{m}})\right]^m. \end{aligned} \tag{49}$$

In the last line of (49) we changed the integration variable to $k = -\lambda\sqrt{m}$ in order to get the 'scaled' threshold $Z/\sqrt{m}$ in the integrand, which makes it easier to visualize the result using Fig. 1.

We define $D(k) = (2\pi)^{-1} e^{ikZ/\sqrt{m}} \left[\tilde{\varphi}(\frac{k}{\sqrt{m}})\right]^m$ for brevity and write $D(k) = D_{\text{even}}(k) + D_{\text{odd}}(k)$. The power expansion of D_{odd} around $k = 0$ has dominant term k^a, where $a > 0$ (Corollary 4). We write

$$\lim_{\eta\downarrow 0} \int_{-\infty}^{\infty} \mathrm{d}k\; \frac{D(k)}{k + i\eta} = \lim_{\eta\downarrow 0} \int_{-\infty}^{\infty} \mathrm{d}k\; \frac{(k - i\eta)D(k)}{k^2 + \eta^2} = \lim_{\eta\downarrow 0} \int_{-\infty}^{\infty} \mathrm{d}k\; \frac{kD_{\text{odd}}(k)}{k^2 + \eta^2} - i\pi D(0). \tag{50}$$

Here we made use of a standard representation of the delta function, $\delta(k) = \frac{1}{\pi} \lim_{\eta \to 0} \eta/(k^2 + \eta^2)$. We also used the fact that in the remaining integration the D_{even} vanishes since it gets multiplied by an odd function of k. Then we use that $a > 0$ in the power series of D_{odd}. This causes the integrand to behave like k^{-1+a} in the limit $\eta \to 0$, i.e. the integral near $k = 0$ is convergent even when η is precisely zero. Thus we can set $\eta = 0$ in this integral.

$$\Pr[S_j > Z] = i \lim_{\eta\downarrow 0} \int_{-\infty}^{\infty} \mathrm{d}k\; \frac{D(k)}{k + i\eta} = i \int_{-\infty}^{\infty} \mathrm{d}k\; \frac{D(k)}{k} + \pi D(0). \tag{51}$$

□

B Proof of Theorem 4

We start from Corollary 4 and write a general power series expansion,

$$\tilde{\varphi}(k) = 1 - \tfrac{1}{2}k^2 + \sum_{t=0}^{\infty} \gamma_t |k|^{r_t}, \tag{52}$$

where the $r_t \geq 3$ are powers and the $\gamma_t \in \mathbb{C}$ are coefficients of the form $i^{\beta_t \operatorname{sgn} k}$ times a real factor. In this expression the desired relation $\tilde{\varphi}(-k) = [\tilde{\varphi}(k)]^*$ evidently holds, and the properties $\tilde{\varphi}(0) = 1$, $\tilde{\varphi}'(0) = 0$, $\tilde{\varphi}''(0) = -1$, $|\tilde{\varphi}'''(0)| < \infty$ are clearly present. Then we write

$$\left[\tilde{\varphi}(\frac{k}{\sqrt{m}})\right]^m = \exp\left[m \ln \tilde{\varphi}(\frac{k}{\sqrt{m}})\right] = e^{-\frac{1}{2}k^2} \exp\left[m \sum_{t=0}^{\infty} (\frac{|k|}{\sqrt{m}})^{r'_t} \delta_t\right], \tag{53}$$

where the powers $r'_t \geq 3$ and coefficients $\delta_t \propto i^{\beta'_t \operatorname{sgn} k}$ are obtained (laboriously) by substituting (52) into the Taylor series for the logarithm, $\ln(1+\varepsilon) = \varepsilon - \varepsilon^2/2 + \varepsilon^3/3 - \varepsilon^4/4 + \cdots$. It is worth noting that m disappears from the k^2 term, but not from the others. Eq. (32) is obtained from (53) by using the Taylor series for the exp function,

$$\exp \varepsilon = 1 + \varepsilon + \varepsilon^2/2! + \varepsilon^3/3! + \cdots \tag{54}$$

and (again laboriously) collecting terms with equal powers of k.

Since we started out with powers $r_t \geq 3$, we end up with powers $\nu_t \geq 3$. A power $|k|^{\nu_t}$ may occur together with many different powers of m. This is seen as follows. The series expansion of $\ln \tilde{\varphi}(k/\sqrt{m})$ is a power series in $|k|/\sqrt{m}$. Then the logarithm is multiplied by m, and a power $|k|^{r'}$ always occurs together with $m^{1-r'/2}$. Next, the k-expansion of exp mixes up the powers of m. For instance, the power k^6 occurs as $m\delta_6(|k|/\sqrt{m})^6 \propto k^6 m^{-2}$ but also as a term $[m\delta_3(|k|/\sqrt{m})^3]^2/2! \propto k^6 m^{-1}$.

The 'worst case' (many factors m resulting from high powers of ε in (54)) occurs when ν_t is a multiple of 3, say $\nu_t = 3j$; there the power k^{3j} can be built up from a term $[m\delta_3(|k|/\sqrt{m})^3]^j/j!$, which is proportional to $k^{3j} m^{j-3j/2} = k^{\nu_t} m^{-\nu_t/6}$. All the j factors scale as $m(|k|/\sqrt{m})^3 = |k|^3/\sqrt{m}$. This is the least negative power of m that can occur relative to the power of k. For other powers ν_t, the 'building blocks' from which k^{ν_t} is built up cannot all scale in this way; at least one of the factors has faster decay.[7] This proves the statement about the at least $m^{-\nu_t/6}$ decay.

Finally, (33) follows by applying Lemma 9 and Corollary 2 to evaluate the integrals that arise when (32) is substituted into Theorem 3. □

C Proof of Lemma 14

C.1 The case $b < c/q$

A symbol that occurs fewer than c/q times cannot have the majority. Consider the extreme case where all the other symbols also occur b times: then the total number of symbols received by the coalition would be $q \cdot b < c$.

C.2 The case $b > c/2$

Since the colluder strategy is majority voting, we have $\Psi_b(\boldsymbol{x}) = 1$ for $b > c/2$. (This follows from the fact that none of the components x_a can exceed $c/2$ due to the sum rule $\sum_a x_a = c - b < c/2$.) The result (41) follows after substitution of $\Psi_b(\boldsymbol{x}) = 1$ into (24), summing up ($\sum_{\boldsymbol{x}}$) the probabilities to 1, and finally writing the Beta functions in terms of Gamma functions according to (2).

[7] For instance, the least negative power of m multiplying k^7 is obtained from the ε^2 term in (54) and is given by $2[m\delta_3(|k|/\sqrt{m})^3][m\delta_4(|k|/\sqrt{m})^4]/2! \propto [|k|^3/\sqrt{m}][|k|^4/m]$.

C.3 The case $b = c/2$

Now $\Psi_b(\boldsymbol{x}) = 1$ unless $x_\beta = c/2$ for some $\beta \in \{1, \cdots, q-1\}$; in that case $\Psi_b(\boldsymbol{x}) = 1/2$ since there are two equivalent symbols to choose from. We have

$$\begin{aligned} K_{c/2} &= \sum_{\boldsymbol{x}:x_\beta \neq c/2} \mathbb{P}_{q-1}(\boldsymbol{x}|\frac{c}{2}) + \sum_{a=1}^{q-1} \binom{c/2}{c/2} \frac{B(\kappa \mathbf{1}_{q-1} + \frac{c}{2}\boldsymbol{e}_a)}{B(\kappa \mathbf{1}_{q-1})} \cdot \tfrac{1}{2} \\ &= \sum_{\boldsymbol{x}} \mathbb{P}_{q-1}(\boldsymbol{x}|\frac{c}{2}) - \frac{1}{2} \sum_{a=1}^{q-1} \frac{B(\kappa \mathbf{1}_{q-1} + \frac{c}{2}\boldsymbol{e}_a)}{B(\kappa \mathbf{1}_{q-1})} \\ &= 1 - \frac{q-1}{2} \frac{B(\kappa \mathbf{1}_{q-1} + \frac{c}{2}\boldsymbol{e}_a)}{B(\kappa \mathbf{1}_{q-1})}. \end{aligned} \tag{55}$$

In the last line we used the fact that the a is arbitrary. Finally, without loss of generality we can set $a = 1$.

C.4 The case $c/q < b < c/2$

We have $\Psi_b(\boldsymbol{x}) = 0$ whenever $x_j > b$ for some index j. Hence we only have to sum over $x_j \leq b$. When $x_j < b$ for all j, then $\Psi_b(\boldsymbol{x}) = 1$. Furthermore, when there are exactly ℓ indices with $x_j = b$, then $\Psi_b(\boldsymbol{x}) = 1/(\ell + 1)$.
We reorganize the $\boldsymbol{x}$-sum in (24) to take the multiplicity ℓ into account: ℓ of the components are set to b and the leftover summation variables x_1 to $x_{q-1-\ell}$ range between 0 and $b-1$.

$$\sum_{\boldsymbol{x}} \Psi_b(\boldsymbol{x})(\cdots) \to \sum_{\ell=0}^{q-1} \frac{1}{\ell+1} \binom{q-1}{\ell} \sum_{x_1=0}^{b-1} \cdots \sum_{x_{q-1-\ell}=0}^{b-1} \delta_{\ell b + x_1 + \cdots + x_{q-1-\ell}, c-b} \, (\cdots). \tag{56}$$

Here the factor $\binom{q-1}{\ell}$ pops up because the summand in (24) is fully symmetric under permutations of $\boldsymbol{x}$. The Kronecker delta takes care of the constraint that the components of $\boldsymbol{x}$ add up to $c - b$. Notice that we let ℓ get as large as $q - 1$, even though it may be impossible to satisfy the $\boldsymbol{x}$-sum constraint for large ℓ; this is taken care of by the Kronecker delta, which sets the constraint-violating terms to zero.
Next we use a sum representation of the Kronecker δ as follows,

$$\delta_{z,0} = \frac{1}{N} \sum_{a=0}^{N-1} (e^{i2\pi/N})^{az}, \tag{57}$$

with $z = (\ell+1)b + x_1 + \cdots + x_{q-1-\ell} - c$. This is a correct representation only if N is larger than the maximum $|z|$ that can occur. Hence, in order for (57) to work for the δ in (56), N must be larger than the maximum value of $|(\ell+1)b + x_1 + \cdots + x_{q-1-\ell} - c|$ that may occur for any (b, ℓ). Taking into account that the range of b is $c/q \leq b < c/2$, and that $x_j \leq b - 1$, the bound on N as stated in the Lemma follows after some algebra.[8]
We shift the a-sum completely to the left, through the $\boldsymbol{x}$-sum and the ℓ-sum. Next we write the upper Beta function in (24), for given multiplicity ℓ, as

$$B(\kappa \mathbf{1}_{q-1} + \boldsymbol{x}) = \frac{[\Gamma(\kappa+b)]^\ell \prod_{j=1}^{q-1-\ell} \Gamma(\kappa + x_j)}{\Gamma(c - b + \kappa[q-1])}, \tag{58}$$

and the multinomial as

$$\binom{c-b}{\boldsymbol{x}} = \frac{(c-b)!}{[b!]^\ell \prod_{j=1}^{q-1-\ell} x_j!}. \tag{59}$$

[8] It is allowed to choose N as a function of b. That leads to a slightly smaller number of terms in the a-summation. We did not wish to add such a complication.

All the expressions depending on the x_j variables are fully factorized; the part of the summand that contains the x_j is given by

$$\prod_{j=1}^{q-1-\ell}\left[\sum_{x_j=0}^{b-1}\frac{\Gamma(\kappa+x_j)}{x_j!}\omega_N^{ax_j}\right]=(G_{ba})^{q-1-\ell}. \tag{60}$$

Next we evaluate the ℓ-sum analytically. It is given by

$$\sum_{\ell=0}^{q-1}\frac{1}{\ell+1}\binom{q-1}{\ell}v^\ell=\frac{(1+v)^q-1}{qv} \tag{61}$$

with

$$v=\frac{\Gamma(b+\kappa)\omega_N^{ab}}{b!G_{ba}}. \tag{62}$$

Finally the result (41) follows after some elementary rewriting. □

D Proof of Theorem 5

We start by considering the probability of a certain accusation value u occurring for an innocent user, for fixed $\boldsymbol{p}$ and y. (We omit all column indices.) There are only two discrete possibilities: (i) $g_1(p_y)$ if the user's symbol is y; this occurs with probability p_y; (ii) $g_0(p_y)$ if the user's symbol is not y; this occurs with probability $1-p_y$. Hence we can write this distribution as a sum of two delta peaks as follows,

$$\varphi(u|\boldsymbol{p},y)=p_y\delta(u-g_1(p_y))+(1-p_y)\delta(u-g_0(p_y)). \tag{63}$$

The full $\varphi(u)$, without conditioning, is obtained by taking the expectation over y and $\boldsymbol{p}$. Since the expectation over y involves the parameters $\theta_{y|\boldsymbol{\sigma}}$, the expectation over $\boldsymbol{\sigma}$ has to be done as well.

$$\varphi(u)=\mathbb{E}_{\boldsymbol{p}}\mathbb{E}_{\boldsymbol{\sigma}|\boldsymbol{p}}\sum_{y\in\mathcal{Q}}\theta_{y|\boldsymbol{\sigma}}\ \varphi(u|\boldsymbol{p},y). \tag{64}$$

Next we note that $\varphi(u|\boldsymbol{p},y)$ depends only on p_y. Hence we can write $\varphi(u|p_y)$, and

$$\varphi(u)=\sum_{y\in\mathcal{Q}}\mathbb{E}_{p_y}\mathbb{E}_{\boldsymbol{\sigma}|p_y}\theta_{y|\boldsymbol{\sigma}}\ \varphi(u|p_y)=\sum_{y\in\mathcal{Q}}\mathbb{E}_{p_y}\mathbb{E}_{\sigma_y|p_y}\mathbb{E}_{\boldsymbol{\sigma}_{\setminus y}|\sigma_y}\theta_{y|\boldsymbol{\sigma}}\ \varphi(u|p_y). \tag{65}$$

Now we use $\mathbb{E}_{\boldsymbol{\sigma}_{\setminus y}|\sigma_y}\theta_{y|\boldsymbol{\sigma}}=K_{\sigma_y}$, the binomial form (21) of $\mathbb{E}_{\sigma_y|p_y}$ and the marginal distribution of p_y (Lemma 2). The dummy summation variable σ_y is replaced by the notation b in order to stress the fact that it does not depend on y. Substitution of all these ingredients gives

$$\begin{aligned}\varphi(u) &= \sum_{y\in\mathcal{Q}}\int_0^1\mathrm{d}p_y\ f(p_y)\sum_{b=0}^{c}\binom{c}{b}p_y^b(1-p_y)^{c-b}K_b\ \varphi(u|p_y)\\ &= \frac{q}{B(\kappa,\kappa[q-1])}\sum_{b=1}^{c}\binom{c}{b}K_b\int_0^1\mathrm{d}p_y\ p_y^{-1+\kappa+b}(1-p_y)^{-1+\kappa[q-1+]c-b}\varphi(u|p_y).\end{aligned} \tag{66}$$

In the last line we have used that $K_0=0$ and that the integral over p_y yields the same result for every y. In order to evaluate the p_y-integral we have to rewrite the delta functions of (63) into the form $\delta(p_y-\cdots)$. We use the rule

$$\delta(u-w(p))=\frac{\delta(p-w^{\mathrm{inv}}(u))}{|\mathrm{d}w/\mathrm{d}p|} \tag{67}$$

for any monotonic function $w(p)$, which yields

$$\begin{aligned}\delta(u - g_1(p)) &= \Theta(u)\frac{2u}{(1+u^2)^2}\delta(p - \frac{1}{1+u^2}) \\ \delta(u - g_0(p)) &= \Theta(-u)\frac{2|u|}{(1+u^2)^2}\delta(p - \frac{u^2}{1+u^2}).\end{aligned} \tag{68}$$

After some algebra, it is then seen that the p_y-integral evaluates to

$$\frac{2}{(1+u^2)^{c+\kappa q+1}}\left[\Theta(u)(u^2)^{\kappa[q-1]+c-\sigma_y-\frac{1}{2}} + \Theta(-u)(u^2)^{\kappa+\sigma_y-\frac{1}{2}}\right]. \tag{69}$$

Splitting φ into a part containing $\Theta(u)$ and a part containing $\Theta(-u)$ finally yields the end result.
□