Asymptotic fingerprinting capacity in the Combined Digit Model

Dion Boesten and Boris Škorić

Eindhoven University of Technology

Abstract

We study the channel capacity of q-ary fingerprinting in the limit of large attacker coalitions. We extend known results by considering the Combined Digit Model, an attacker model that captures signal processing attacks such as averaging and noise addition. For $q = 2$ we give results for various attack parameter settings. For $q \geq 3$ we present the relevant equations without providing a solution. We show how the channel capacity in the Restricted Digit Model is obtained as a limiting case of the Combined Digit Model.

1 Introduction

1.1 Collusion resistant watermarking

Watermarking is a means of tracing the (re-)distribution of content. Before distribution, digital content is modified by applying an imperceptible watermark (WM), embedded using a watermarking algorithm. Once an unauthorized copy of the content is found, the WM helps to trace those users who participated in the creation of the copy. This is known as 'forensic watermarking'. Reliable tracing requires resilience against attacks that aim to remove the WM. Collusion attacks are a particular threat: multiple users cooperate, and differences between their versions of the content tell them where the WM is located. Coding theory has provided a number of collusion-resistant codes. The resulting system has two layers: The coding layer determines which message to embed, and protects against collusion attacks. The underlying watermarking layer hides symbols of the code in segments[1] of the content.

Many collusion resistant codes have been proposed in the literature. Most notable is the Tardos code [15], which achieves the asymptotically optimal proportionality $m \propto c^2$, with m the code length and c the size of the coalition of attackers. Tardos introduced a two-step stochastic procedure for generating binary codewords: (i) For each segment a bias is randomly drawn from some distribution. (ii) For each user independently, a 0 or 1 is randomly drawn for each segment using the bias for that segment. This construction was generalized to larger (q-ary) alphabets in [16].

The interface between the coding and watermarking layer is usually specified in terms of the *Marking Assumption* (MA), which states that the colluders are able to perform modifications only in those segments where they received different WM symbols. These segments are called detectable positions. Furthermore, within this class of attacks there is a classification of attacks according to the manipulations that can be performed *in the detectable positions*. In the *Restricted Digit Model* (RDM), the coalition is only allowed to pick one symbol that they received. In the *Unreadable Digit Model*, they are furthermore allowed to create an erasure. In the *Arbitrary Digit Model*, they can pick any symbol from the alphabet, even one that they did not receive (but not an erasure). The *General Digit Model* allows any symbol from the alphabet or an erasure.

For $q = 2$, all these MA attacks are equivalent. For larger alphabets, the general feeling is that realistic attacks are somewhere between the RDM and the Unreadable Digit Model. To come to

[1] The 'segments' are defined in a very broad sense.

an even more realistic attack model (also for $q = 2$) which additionally takes into account signal processing (e.g. averaging attacks and noise addition), one has to depart from the MA. Such attack models were proposed in [19] and [17] for general q, and for $q = 2$ in e.g. [7, 8].

1.2 Asymptotic channel capacity

In Tardos' scheme [15] and later improvements and generalisations (e.g. [3, 5, 9, 10, 11, 13, 14, 16, 17, 18, 19]), users are found to be innocent or guilty via an 'accusation sum', a sum of weighted per-segment contributions, computed for each user separately. The analysis of achievable performance was greatly helped by the onset of an information-theoretic treatment of anti-collusion codes. The whole class of bias-based codes can be treated as a maximin game between the watermarker and the colluders [2, 6, 12], independently played for each segment, where the payoff function is the mutual information between the symbols $x_1, \ldots, x_c$ handed to the colluders and the symbol y produced by them. In each segment (i.e. for each bias) the colluders try to minimize the payoff function using an attack strategy that depends on the (frequencies of the) received symbols $x_1, \ldots, x_c$. The watermarker tries to maximize the average payoff over the segments by setting the bias distribution.

The rate of a fingerprinting code is defined as $(\log_q n)/m$, where n is the number of users and m the code length (the number of q-ary symbols). The *fingerprinting capacity* is the maximum achievable rate. For $q = 2$ it was conjectured that the capacity is asymptotically $1/(c^2 2 \ln 2)$. The conjecture was proved in [1, 6]. Amiri and Tardos [1] developed a joint decoder scheme (for the binary case) where candidate coalitions get a score related to the mutual information between their symbols and y. This scheme achieves capacity but is computationally very expensive. Huang and Moulin [6] proved for the large-c limit (in the binary case) that the interleaving attack and Tardos's arcsine distribution are optimal.

It was shown by Boesten and Škorić [4] that the asymptotic channel capacity for q-ary alphabets in the RDM is $(q-1)/(2c^2 \ln q)$. Their proof method revealed neither the optimal attack strategy nor the optimal bias distribution.

1.3 Contributions

In this paper we study the asymptotic channel capacity of q-ary fingerprinting in the Combined Digit Model (CDM) [17], following the approach of [4]. We choose for the CDM because this model is defined for general q and captures a large range of non-MA attacks.

The CDM allows the coalition to add noise and to do averaging attacks. Given the set of symbols used in the averaging, various model parameters describe the probability of these and other symbols being detected by the watermark detector (see Sections 2.3 and 2.4 for details).

We show that the asymptotic channel capacity in the CDM can be found by solving the following problem: Find a mapping γ from the hypersphere in q dimensions to the hypersphere in 2^q dimensions, such that γ minimizes the trace of the induced metric tensor in the latter space (see Section 3). The attack parameters of the CDM give rise to non-trivial constraints on the mapping, which have to be satisfied. One of the main differences between the RDM and CDM lies in the dimension of the target space of γ, which is $q - 1$ in the RDM and $2^q - 1$ in the CDM. We show how the RDM capacity is re-obtained from the CDM setting (Section 4).

For $q \geq 3$ we have not solved the above mentioned minimization problem. For $q = 2$ we present numerical results for various attack parameter choices. The numerics involve computations of constrained geodesics, a difficult problem in general. The resulting graphs show a nontrivial dependence of the capacity on the CDM attack parameters.

2 Preliminaries

2.1 Notation

We use capital letters to represent random variables, and lowercase letters to their realizations. Vectors are denoted in boldface and the components of a vector $\boldsymbol{x}$ are written as x_i. Vectors are considered to be column vectors. The expectation over a random variable X is denoted as $\mathbb{E}_X$. The mutual information between X and Y is denoted by $I(X;Y)$, and the mutual information conditioned on a third variable Z by $I(X;Y|Z)$. The base-q logarithm is written as $\log_q$ and the natural logarithm as ln. The standard Euclidean norm of a vector $\boldsymbol{x}$ is denoted by $\|\boldsymbol{x}\|$. The Kronecker delta of two variables α and β is denoted by $\delta_{\alpha\beta}$. A sum over all possible outcomes of a random variable X is written as $\sum_x$. In order not to clutter up the notation we will often omit the set to which x belongs when it is clear from the context. We use the notation $|\mathcal{Q}|$ for the size of a set $\mathcal{Q}$.

2.2 Fingerprinting with per-segment symbol biases

Tardos [15] introduced the first fingerprinting scheme that achieves optimality in the sense of having the asymptotic behavior $m \propto c^2$. He introduced a two-step stochastic procedure for generating the codeword matrix X. Here we show the generalization to non-binary alphabets [16]. A Tardos code of length m for a number of users n over the alphabet $\mathcal{Q}$ of size q is a set of n length-m sequences of symbols from $\mathcal{Q}$ arranged in an $n \times m$ matrix X. The codeword for a user $i \in \{1, \ldots, n\}$ is the i-th row in X. The symbols in each column $j \in \{1, \ldots, m\}$ are generated in the following way. First an auxiliary bias vector $\boldsymbol{P}^{(j)} \in [0,1]^q$ with $\sum_\alpha P_\alpha^{(j)} = 1$ is generated independently for each column j, from a distribution F which is considered known to the attackers. (The $\boldsymbol{P}^{(j)}$ are sometimes referred to as 'time sharing' variables.) The result $\boldsymbol{p}^{(j)}$ is used to generate each entry X_{ij} of column j independently: $\text{Prob}\,[X_{ij} = \alpha] = p_\alpha^{(j)}$.

2.3 The collusion attack in the Combined Digit Model

Let the random variable $\Sigma_\alpha^{(j)} \in \{0, 1, \ldots, c\}$ denote the number of colluders who receive the symbol α in segment j. It holds that $\sum_{\alpha \in \mathcal{Q}} \sigma_\alpha^{(j)} = c$ for all j. (We remind the reader that outcomes of random variables are written in lowercase.) From now on we will drop the segment index j, since all segments are independent. In the *Restricted Digit Model* the colluders produce a symbol $Y \in \mathcal{Q}$ that they have seen at least once. In the *Combined Digit Model* one also allows the attackers to output a mixture of symbols. Let

$$\Omega(\boldsymbol{\Sigma}) \triangleq \{\alpha \in \mathcal{Q} \mid \Sigma_\alpha \geq 1\} \tag{1}$$

be the set of symbols that the pirates have seen in a certain segment. Then the output of the pirates is a non-empty set $\Psi \subseteq \Omega(\boldsymbol{\Sigma})$. On the watermarking level this represents a content-averaging attack where all the symbols in Ψ are used. It has been shown [12] that it is sufficient to consider a probabilistic per-segment (column) attack which does not distinguish between the different colluders. Such an attack then only depends on $\boldsymbol{\Sigma}$, and the strategy can be completely described by a set of probabilities $\theta_{\psi|\boldsymbol{\sigma}} \in [0,1]$, which are defined as

$$\theta_{\psi|\boldsymbol{\sigma}} \triangleq \text{Prob}[\Psi = \psi \mid \boldsymbol{\Sigma} = \boldsymbol{\sigma}]. \tag{2}$$

For all $\boldsymbol{\sigma}$ conservation of probability gives $\sum_\psi \theta_{\psi|\boldsymbol{\sigma}} = 1$.

2.4 Detection process in the Combined Digit Model

The Combined Digit Model also introduces a stochastic detection process. Let $|\Psi|$ be the cardinality of the output set Ψ. Then each symbol in Ψ is detected with probability $t_{|\Psi|}$. Each symbol

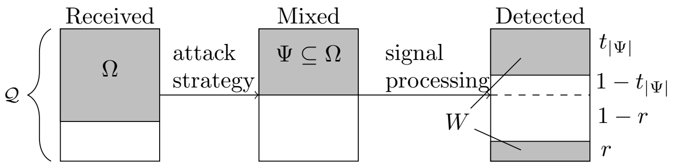

Figure 1: *Overview of the detection process in the Combined Digit Model. The detection probabilities are shown on the right.*

not in the set Ψ is detected with error probability r. The set $W \subseteq \mathcal{Q}$ indicates which symbols are detected. Note that Ψ is forced to be non-empty but $W = \emptyset$ can occur.

The numbers t_i for $i = 1, 2, \ldots, q$ are decreasing, since mixing more symbols makes it more difficult to detect the individual symbols. The overall probability of detecting a set w, given ψ, depends on r, $t_{|\psi|}$ and the sizes of the fours sets shown under 'Detected' in Fig. 1. From top to bottom, these are (i) the number of used symbols that get detected, $|\psi \cap w|$; (ii) # used symbols that *do not* get detected, $|\psi \setminus w|$; (iii) # not used symbols that are not detected, $q - |\psi \cup w|$; and (iv) # not used symbols that *are* detected due to noise, $|w \setminus \psi|$. For a given ψ, the probability that the detector outputs a detected set w is

$$\begin{aligned} M_{w|\psi} &\triangleq \text{Prob}\left[W = w \mid \Psi = \psi\right] \\ &= t_{|\psi|}^{|\psi \cap w|} \left(1 - t_{|\psi|}\right)^{|\psi \setminus w|} (1 - r)^{q - |\psi \cup w|}\, r^{|w \setminus \psi|}. \end{aligned} \tag{3}$$

These probabilities form a $2^q \times (2^q - 1)$ matrix M. In this way we can define

$$\tau_{w|\boldsymbol{\sigma}} \triangleq \text{Prob}\left[W = w \mid \boldsymbol{\Sigma} = \boldsymbol{\sigma}\right] = \sum_{\psi} M_{w|\psi} \theta_{\psi|\boldsymbol{\sigma}} = (M\theta)_{w|\boldsymbol{\sigma}}\,, \tag{4}$$

or, in matrix notation, $\tau = M\theta$. (The matrix notation for the relation (4) is novel.)

2.5 Collusion channel and fingerprinting capacity

Similarly to the RDM [4] the attack can be interpreted as a noisy channel with input $\boldsymbol{\Sigma}$ and output W. A capacity for this channel can then be defined, which gives an upper bound on the achievable code rate of a reliable fingerprinting scheme. The first step of the code generation, drawing the biases $\boldsymbol{p}$, is not considered to be a part of the channel. The fingerprinting capacity C_q^{CDM} for a coalition of size c and alphabet size q in the CDM is equal to the optimal value of the following two-player game:

$$C_q^{\text{CDM}} = \max_F \min_\theta \frac{1}{c} I(W; \boldsymbol{\Sigma} \mid \boldsymbol{P}) = \max_F \min_\theta \frac{1}{c} \int F(\boldsymbol{p}) I(W; \boldsymbol{\Sigma} \mid \boldsymbol{P} = \boldsymbol{p}) \mathrm{d}^q \boldsymbol{p}. \tag{5}$$

Here the information is measured in q-ary symbols. Our aim is to compute the fingerprinting capacity C_q^{CDM} in the limit $c \to \infty$.

The payoff function $I(W; \boldsymbol{\Sigma} \mid \boldsymbol{P})$ is linear in F and convex in τ. Because $\tau = M\theta$ is linear in θ the game is also convex in θ and we can apply Sion's Theorem:

$$\begin{aligned} \max_F \min_\theta I(W; \boldsymbol{\Sigma} \mid \boldsymbol{P}) &= \min_\theta \max_F \; I(W; \boldsymbol{\Sigma} \mid \boldsymbol{P}) \\ &= \min_\theta \max_p \; I(W; \boldsymbol{\Sigma} \mid \boldsymbol{P} = \boldsymbol{p}). \end{aligned} \tag{6}$$

In the second step we performed the maximization over F by choosing the optimum $F^*(\boldsymbol{p}) = \delta(\boldsymbol{p} - \boldsymbol{p}_{\max})$, where $\boldsymbol{p}_{\max}$ is one of the locations where $I(W; \boldsymbol{\Sigma} \mid \boldsymbol{P} = \boldsymbol{p})$ has its maximum.

3 Asymptotic analysis for general alphabet size

We are interested in how the payoff function $I(W; \boldsymbol{\Sigma} \mid \boldsymbol{P} = \boldsymbol{p})$ of the alternative game (6) behaves as c goes to infinity. Following the same approach as in [4] our starting point is the observation that the random variable $\boldsymbol{\Sigma}/c$ tends to a continuum in $[0,1]^q$ with mean $\boldsymbol{p}$. Hence we introduce a continuous strategy $\boldsymbol{h}\left(\frac{\boldsymbol{\sigma}}{c}\right)$:

$$h_\psi\left(\frac{\boldsymbol{\sigma}}{c}\right) = \lim_{c\to\infty} \theta_{\psi|\boldsymbol{\sigma}}. \tag{7}$$

We also define

$$g_w\left(\frac{\boldsymbol{\sigma}}{c}\right) = \lim_{c\to\infty} \tau_{w|\boldsymbol{\sigma}} = \sum_\psi M_{w|\psi} h_\psi\left(\frac{\boldsymbol{\sigma}}{c}\right), \tag{8}$$

which in matrix notation can be written as $\boldsymbol{g} = M\boldsymbol{h}$. The next step is to do a second order Taylor expansion of $g_w\left(\frac{\boldsymbol{\sigma}}{c}\right)$ around the point $\frac{\boldsymbol{\sigma}}{c} = \boldsymbol{p}$. This allows us to expand I in powers of $1/c$, giving (see [4])

$$I(W; \boldsymbol{\Sigma} \mid \boldsymbol{P} = \boldsymbol{p}) = \frac{T(\boldsymbol{p})}{2c\ln q} + \mathcal{O}\left(\frac{1}{c\sqrt{c}}\right) \tag{9}$$

$$T(\boldsymbol{p}) \triangleq \sum_w \frac{1}{g_w(\boldsymbol{p})} \sum_{\alpha\beta} K_{\alpha\beta} \frac{\partial g_w(\boldsymbol{p})}{\partial p_\alpha} \frac{\partial g_w(\boldsymbol{p})}{\partial p_\beta}, \tag{10}$$

where $K_{\alpha\beta} = \delta_{\alpha\beta} p_\alpha - p_\alpha p_\beta$ is the scaled covariance matrix of $\boldsymbol{\Sigma}$. The asymptotic capacity $C^{\text{CDM}}_{q,\infty}$ in the limit of $c \to \infty$ is then defined as the solution of the continuous version of the game (6):

$$C^{\text{CDM}}_{q,\infty} \triangleq \frac{1}{2c^2 \ln q} \min_{\boldsymbol{h}} \max_{\boldsymbol{p}} T(\boldsymbol{p}). \tag{11}$$

At this point we introduce the variable transformations $u_\alpha \triangleq \sqrt{p_\alpha}, \gamma_w \triangleq \sqrt{g_w}$ and also the $2^q \times q$ Jacobian matrix $J_{w\alpha}(\boldsymbol{u}) \triangleq \frac{\partial \gamma_w(\boldsymbol{u})}{\partial u_\alpha}$. This transformation means we switch to hyperspheres ($\|\boldsymbol{u}\| = 1, \|\boldsymbol{\gamma}\| = 1$) instead of the hyperplanes ($\sum_\alpha p_\alpha = 1, \sum_w g_w = 1$) that we had before. The function $\boldsymbol{\gamma}(\boldsymbol{u})$ was originally defined only on the domain $\|\boldsymbol{u}\| = 1$, but the Taylor expansion forces us to define $\boldsymbol{\gamma}$ on a larger domain, i.e. slightly away from $\|\boldsymbol{u}\| = 1$. There are many consistent ways to do this domain extension. We choose to define $\boldsymbol{\gamma}$ such that it is independent of the radial coordinate $\|\boldsymbol{u}\|$. This choice yields $J\boldsymbol{u} = 0$, which allows us to simplify $T(\boldsymbol{u})$ to

$$T(\boldsymbol{u}) = \sum_{w,\alpha} \left(\frac{\partial \gamma_w}{\partial u_\alpha}\right)^2 = \operatorname{Tr}(J^T J) = \sum_{i=1}^{q-1} \lambda_i(\boldsymbol{u}), \tag{12}$$

where $\lambda_i(\boldsymbol{u})$ are the eigenvalues of $J^T J$. Because of our choice $J\boldsymbol{u} = 0$ we already know that one of the eigenvalues is 0 with eigenvector $\boldsymbol{u}$. Hence there are only $q-1$ eigenvalues left. Note that (12) can be interpreted as the trace of a metric: if we define a metric $B_{\alpha\beta} = (\partial\boldsymbol{\gamma}/\partial u_\alpha) \cdot (\partial\boldsymbol{\gamma}/\partial u_\beta)$ in the usual way, then $T(\boldsymbol{u}) = \operatorname{Tr} B$.

We now wish to find

$$\min_{\gamma} \max_{\boldsymbol{u}} T(\boldsymbol{u}) \tag{13}$$

under the constraint

$$\gamma_w = \sqrt{g_w} = \sqrt{(M\boldsymbol{h})_w} \tag{14}$$

with M known and $\boldsymbol{h}$ satisfying

$$h_\psi \geq 0 \quad \forall\psi, \qquad\qquad \textstyle\sum_\psi h_\psi = 1. \tag{15}$$

The constraint (14) makes solving the min-max game (13) more difficult and we are unable to use the same machinery as for the RDM. The main problem is that it is no longer easy to characterize the allowed (sub)space that $\boldsymbol{\gamma}$ lives in.

For the binary alphabet we are however able to go further and compute the asymptotic capacity (see Section 5).

4 Limiting case: Restricted Digit Model

We show how the known result for the Restricted Digit Model (RDM) follows as a limiting case of the CDM.

We set $r = 0$ and $t_i = 1$ for all $i \in \{1, \cdots, q\}$. This means that there is no noise, and any symbol that the attackers use will be detected with 100% certainty. Hence $W = \Psi$. In this situation there is no gain for the attackers to use fusion, as all the fused symbols are detected and provide the content owner with more information. Their best option is to use a single symbol; hence we are back at the RDM.

Mathematically it is slightly more involved to see how the reduction to the RDM channel capacity is obtained. The matrix M becomes $\begin{pmatrix} \mathbf{0} \\ I_{2^q-1} \end{pmatrix}$ where I_{2^q-1} is the identity matrix of size $2^q - 1$. Since M has become trivial, (14) does not really represent a constraint on $\gamma_w(\boldsymbol{u})$ any more. The only difference with [4] is the dimension of the vector: $\boldsymbol{\gamma}$ has $2^q - 1$ components ($w = \emptyset$ is excluded), whereas in the RDM there were only q components. Consequently, the Jacobian J also has a larger dimension. However, the product $J^T J$ is still a $q \times q$ matrix, and the derivation in [4] can be applied in unchanged form to yield two results:

1. The solution of the min-max game satisfies $\max_{\boldsymbol{u}} T(\boldsymbol{u}) = \mathrm{Av}_{\boldsymbol{u}}[T(\boldsymbol{u})]$, i.e. the maximum is equal to the spatial average, and $T(\boldsymbol{u})$ is in fact a constant on the hypersphere $\|\boldsymbol{u}\| = 1$, with
$$T(\boldsymbol{u}) \geq (q-1) \left(\frac{\int \mathrm{d}S_{\boldsymbol{\gamma}}}{\int \mathrm{d}S_{\boldsymbol{u}}} \right)^{2/(q-1)}. \tag{16}$$
Here $\int \mathrm{d}S_{\boldsymbol{u}}$ is the $(q-1)$-dimensional 'volume' integral on the surface of the $\boldsymbol{u}$-hypersphere. The $\int \mathrm{d}S_{\boldsymbol{\gamma}}$ is the corresponding $(q-1)$-dimensional integral in the larger $(2^q - 2)$-dimensional $\boldsymbol{\gamma}$-hypersphere, with $\boldsymbol{\gamma} = \boldsymbol{\gamma}(\boldsymbol{u})$. In [4] the γ-sphere had dimension $q - 1$, and it was used that $\int \mathrm{d}S_{\boldsymbol{\gamma}} \geq \int \mathrm{d}S_{\boldsymbol{u}}$.
2. The interleaving attack yields $T(\boldsymbol{u}) = q - 1$ on the hypersphere $\|u\| = 1$.

We argue (without proof) that $\int \mathrm{d}S_{\boldsymbol{\gamma}} \geq \int \mathrm{d}S_{\boldsymbol{u}}$ still holds. This is because of the Marking Assumption, which fixes the values on the axes in $\boldsymbol{\gamma}$-space. Let $\mathbf{e}_\alpha$ be the unit vector in the α-direction. Then $\boldsymbol{u} = \mathbf{e}_\alpha \implies \boldsymbol{\gamma} = \mathbf{e}_\alpha$. These 'corner' points live in a q-dimensional subspace. It is possible to step out of that subspace for general $\boldsymbol{u}$, but doing so increases the volume $\int \mathrm{d}S_{\boldsymbol{\gamma}}$.

Thus, result #1 gives the lower bound $\max_{\boldsymbol{u}} T(\boldsymbol{u}) \geq q - 1$, while result #2 shows that there exists a strategy achieving the lower bound. The RDM channel capacity $C^{\mathrm{RDM}}_{q,\infty} = (q-1)/(2c^2 \ln q)$ follows.

Remark: If M is perturbed away from the identity matrix, then the extreme points $\boldsymbol{u} = \mathbf{e}_\alpha$ are no longer mapped to mutually orthogonal vectors $\boldsymbol{\gamma}$, but to vectors with smaller mutual angles; the reduction of the angles causes a reduction of $\int \mathrm{d}S_{\boldsymbol{\gamma}}$ and hence the channel capacity. The details are cumbersome and the general case $q \geq 3$ is left for future work.

5 Fingerprinting capacity in the CDM for $q = 2$

5.1 Solving the max-min game

For the binary alphabet $q = 2$ the expression (12) simplifies to
$$T(\boldsymbol{u}) = \mathrm{Tr}(J^T J) = \lambda(\boldsymbol{u}) \tag{17}$$

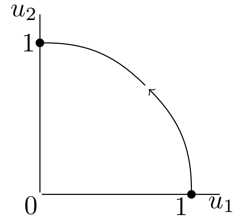

Figure 2: *The path for $\boldsymbol{u}$ is the positive quarter circle.*

since there is only one nonzero eigenvalue. Furthermore we have the relation $\mathrm{d}\boldsymbol{\gamma} = J\mathrm{d}\boldsymbol{u}$ and $\|\mathrm{d}\boldsymbol{\gamma}\| = \sqrt{\lambda}\|\mathrm{d}\boldsymbol{u}\|$ for an infinitesimal change $\mathrm{d}\boldsymbol{u}$. We proceed by rewriting

$$\begin{aligned} \max_{\boldsymbol{u}} T(\boldsymbol{u}) &= \max_{\boldsymbol{u}} \lambda(\boldsymbol{u}) = \left(\max_{\boldsymbol{u}} \sqrt{\lambda(\boldsymbol{u})}\right)^2 \\ &\geq \left(\frac{\int \sqrt{\lambda(\boldsymbol{u})}\|\mathrm{d}\boldsymbol{u}\|}{\int \|\mathrm{d}\boldsymbol{u}\|}\right)^2 = \left(\frac{\int \|\mathrm{d}\boldsymbol{\gamma}\|}{\int \|\mathrm{d}\boldsymbol{u}\|}\right)^2 \equiv \left(\frac{L_{\boldsymbol{\gamma}}}{L_{\boldsymbol{u}}}\right)^2, \end{aligned} \tag{18}$$

where the inequality results from replacing the maximum by a spatial average. The path in the integrals (see Fig. 2) is the quarter-circle $u_1^2 + u_2^2 = 1$ from $\boldsymbol{u} = (1,0)$ to $\boldsymbol{u} = (0,1)$ and hence $L_{\boldsymbol{u}} = \pi/2$.

The next step is to realize that for any curve $\gamma(\boldsymbol{u})$ we have the freedom to parameterize that curve differently in such a way that $\lambda(\boldsymbol{u})$ is constant over that curve, i.e. we are traveling at constant speed. The inequality in (18) can then be changed into an equality and we have

$$\min_{\gamma} \max_{\boldsymbol{u}} T(\boldsymbol{u}) = \frac{4}{\pi^2}(\min_{\gamma} L_{\boldsymbol{\gamma}})^2. \tag{19}$$

Hence we have reduced the problem to finding a curve $\boldsymbol{\gamma}(\boldsymbol{u})$ of minimal length with the constraint $\gamma_w(\boldsymbol{u}) = \sqrt{(M\boldsymbol{h})_w(\boldsymbol{u})}$ where $M(t_1, t_2, r)$ is given by

$$M = \begin{array}{c||c|c|c|} w\backslash\psi & \{0\} & \{1\} & \{0{,}1\} \\ \hline \emptyset & (1-t_1)(1-r) & (1-t_1)(1-r) & (1-t_2)^2 \\ \{0\} & t_1(1-r) & (1-t_1)r & t_2(1-t_2) \\ \{1\} & (1-t_1)r & t_1(1-r) & t_2(1-t_2) \\ \{0{,}1\} & t_1 r & t_1 r & t_2^2 \end{array}. \tag{20}$$

5.2 Geodesics

In general, the method to find length-minimizing curves is to solve the Euler-Lagrange differential equations for the geodesics of the appropriate metric. In our case the additional constraint $\gamma_w(\boldsymbol{u}) = \sqrt{(M\boldsymbol{h})_w(\boldsymbol{u})}$ makes things more difficult. The constraint can be interpreted in the following way. If we write $M = [m_1, m_2, m_3]$ then because of constraint (15) then we have that $\boldsymbol{g} = M\boldsymbol{h}$ is a convex combination of the three column vectors m_1, m_2, m_3. Hence the allowed space of $\boldsymbol{g}$ is anywhere inside the triangle shown in Fig. 3.

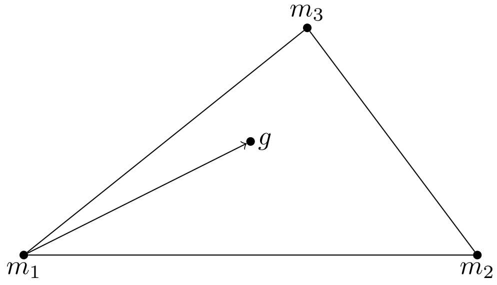

Figure 3: *The vector $\boldsymbol{g}$ is not allowed to lie outside the triangle.*

We switch from variables (u_1, u_2) to s_1, s_2 with $0 \le s_1 \le 1$ and $0 \le s_2 \le 1 - s_1$.

$$\boldsymbol{g}(s_1, s_2) \triangleq m_1 + s_1(m_2 - m_1) + s_2(m_3 - m_1). \tag{21}$$

The marking assumption gives us that $\boldsymbol{u} = (1,0) \Rightarrow \boldsymbol{h} = (1,0,0)$ and $\boldsymbol{u} = (0,1) \Rightarrow \boldsymbol{h} = (0,1,0)$. In terms of $\boldsymbol{g}(s_1, s_2)$ this means $\boldsymbol{g}(0,0) = m_1$ and $\boldsymbol{g}(1,0) = m_2$. We are looking for the shortest path from the lower left corner (m_1) of the triangle to the lower right corner (m_2).
The infinitesimal change in $\mathrm{d}\gamma_w$ in terms of $(\mathrm{d}s_1, \mathrm{d}s_2)$ is given by

$$\mathrm{d}\gamma_w = \frac{\mathrm{d}g_w}{2\sqrt{g_w}} = \frac{(m_{2,w} - m_{1,w})\mathrm{d}s_1 + (m_{3,w} - m_{1,w})\mathrm{d}s_2}{2\sqrt{g_w}}. \tag{22}$$

This allows us to define the appropriate metric $G(s_1, s_2)$,

$$\|\mathrm{d}\boldsymbol{\gamma}\|^2 = G_{11}(\mathrm{d}s_1)^2 + G_{22}(\mathrm{d}s_2)^2 + 2G_{12}\mathrm{d}s_1\mathrm{d}s_2. \tag{23}$$

We use this metric to compute the geodesics (locally distance minimizing curves). See Appendix A for the details.

5.3 Finding the shortest path

We want to find the shortest path between m_1 and m_2 that is fully inside the triangle. If a direct geodesic between these two points exists we know that it is the optimal path; but this does not always happen. We encounter three possible cases, given in Fig. 4. In case A the direct geodesic is the shortest possible path. For cases B and C the optimal paths are shown in Fig. 5.

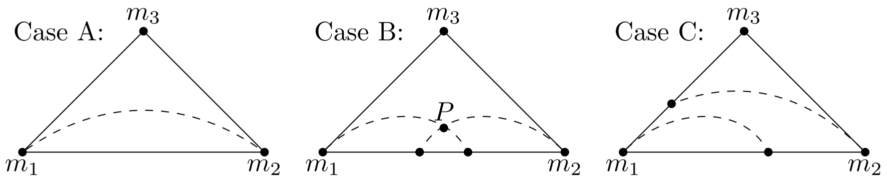

Figure 4: *In case A there exists a direct geodesic from m_1 to m_2. In case B the maximum-slope geodesics starting from m_1 and m_2 intersect in P. In case C they do not intersect.*

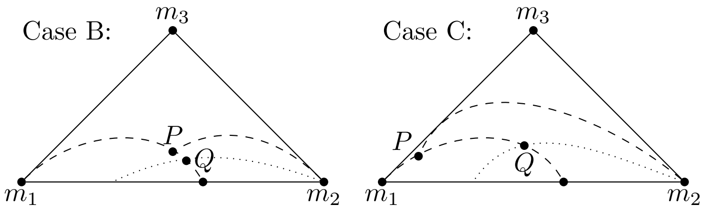

Figure 5: *The optimal path in both cases is $m_1 - P - m_2$ over the dashed lines (geodesics). In case C the geodesic from m_2 is the one which is tangent to the left side of the triangle.*

Any geodesic starting from m_2 with a smaller initial slope eventually has to cross the maximum-slope geodesic from m_1 in a point Q. From Q the optimal path to m_1 is to follow the geodesic; but when you pass P you could have done better by simply going directly from m_2 to P on the geodesic.

Once we have the optimal path we can determine its length $L_{\rm opt}$ (see Appendix) and use it to compute the capacity,

$$C^{\rm CDM}_{2,\infty} = \frac{1}{2c^2 \ln 2} \cdot \frac{4}{\pi^2} L^2_{\rm opt}. \tag{24}$$

5.4 Results

In Fig. 6 we show plots of the ratio $C = C^{\rm CDM}_{2,\infty}/C^{\rm RDM}_{2,\infty}$ between the asymptotic capacities for the CDM and the RDM as a function of the parameters t_1, t_2, r. (For the binary alphabet $\mathcal{Q} = \{0,1\}$, we have that r is the noise strength, t_1 is the probability of detecting a symbol α if the coalition used $\Psi = \{\alpha\}$, and t_2 is the probability of detecting α if the coalition used $\Psi = \{0,1\}$). Several aspects of the graphs are easily understood and yield no surprises:

- Obviously, the capacity is an increasing function of t_1 and t_2, and a decreasing function of r. When the attack options become more powerful, the capacity goes down.
- For r close to zero and t_1 close to 1, the capacity has very weak dependence on t_2. This can be understood as follows. A small value of t_2 effectively means that the attackers create an erasure, which brings us to the Unreadable Digit Model. For large t_2 is it not advantageous for them to take $\Psi = \{0,1\}$, since the detector will find both symbols, giving the tracer more information than taking $|\Psi| = 1$. The attackers will output a single symbol, which brings us back to the RDM. For $(r, t_1) \approx (0,1)$ we are close to the Marking Assumption. When the MA holds, all the attack models for $q = 2$ are equivalent.

Other behaviour is more surprising. In Fig. 6a we see a transition from linear behavior as a function of r (with almost total insensitivity to t_2) to nonlinear behavior (with dependence on t_2). The transition point depends on t_2.

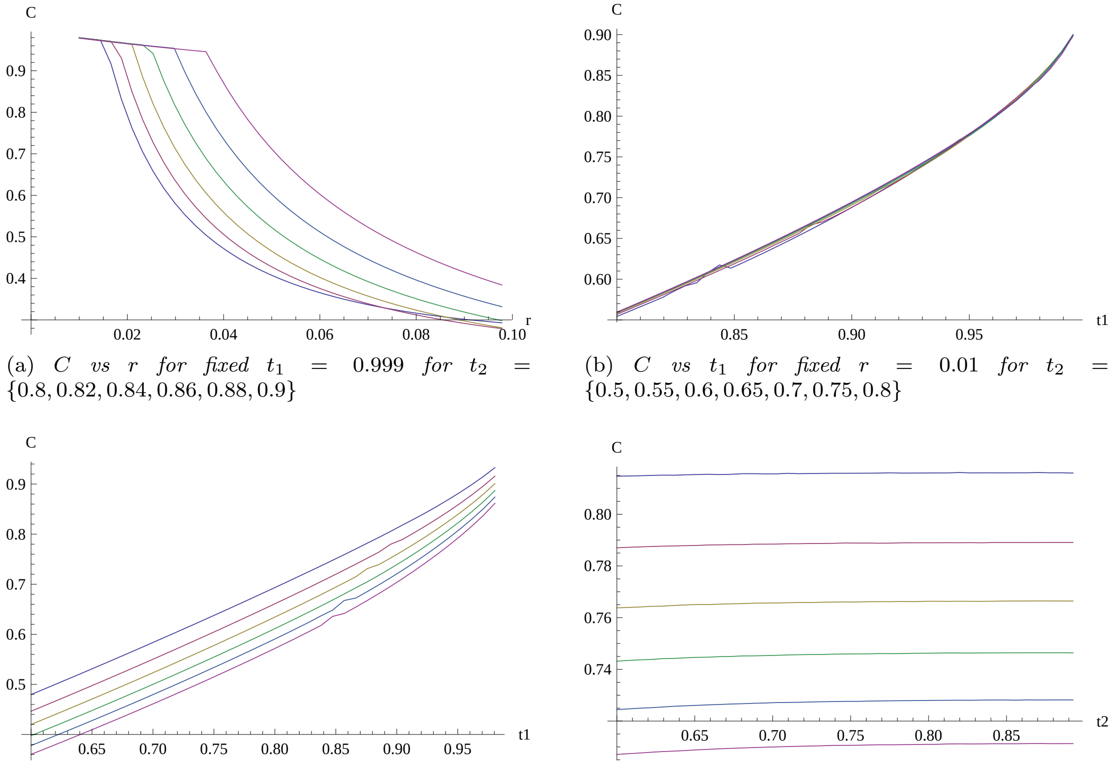

(a) *C vs r for fixed $t_1 = 0.999$ for $t_2 =$* $\{0.8, 0.82, 0.84, 0.86, 0.88, 0.9\}$

(b) *C vs t_1 for fixed $r = 0.01$ for $t_2 =$* $\{0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8\}$

(c) *C vs t_1 for fixed $t_2 = 0.5$ for $r =$* $\{0.01, 0.02, 0.03, 0.04, 0.05\}$

(d) *C vs t_2 for fixed $t_1 = 0.9$ for $r =$* $\{0.01, 0.02, 0.03, 0.04, 0.05\}$

Figure 6: *Numerics for $q = 2$. The ratio $C = C^{\rm CDM}_{2,\infty}/C^{\rm RDM}_{2,\infty}$ is plotted on the vertical axis.*

6 Discussion

We have investigated the asymptotic channel capacity $C^{\rm CDM}_{q,\infty}$ in the Combined Digit Model. For general alphabet size q it turns out to be very difficult to compute this quantity. We have shown how the asymptotic capacity for the RDM [4] follows as a limiting case of the CDM. In the general case, $C^{\rm CDM}_{q,\infty}$ can be expressed as the solution of a min-max game (13) where the payoff function is the trace of the metric induced by the mapping γ from $\sqrt{p_\alpha}$-space to $\sqrt{g_w}$-space. The CDM parameters r and $t_1, \cdots, t_q$ give rise to a constraint $\boldsymbol{g} = M\boldsymbol{h}$ on $\boldsymbol{g}$ which prevents the application of the solution method of [4]. For the binary alphabet we have shown that the problem reduces to finding a constrained geodesic between two points. Our numerical results do not contain significant surprises. They confirm the intuition in the vicinity of the Marking Assumption, $(r, t_1) \approx (0, 1)$. In this regime $C^{\rm CDM}_{2,\infty}$ is practically independent of t_2. The transitions in Fig. 6a are not intuitively clear. The study of these details and of larger alphabets $q \geq 3$ is left for future work.

References

[1] E. Amiri and G. Tardos. High rate fingerprinting codes and the fingerprinting capacity. In *SODA 2009*, pages 336–345.

[2] N.P. Anthapadmanabhan, A. Barg, and I. Dumer. Fingerprinting capacity under the marking assumption. *IEEE Transaction on Information Theory – Special Issue on Information-theoretic Security*, 54(6):2678–2689.

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

[4] D. Boesten and B. Škorić. Asymptotic fingerprinting capacity for non-binary alphabets. In *Information Hiding 2011*, volume 6958 of *LNCS*, pages 1–13. Springer, 2011.

[5] A. Charpentier, F. Xie, C. Fontaine, and T. Furon. Expectation maximization decoding of Tardos probabilistic fingerprinting code. In *Media Forensics and Security*, volume 7254 of *SPIE Proceedings*, page 72540, 2009.

[6] 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.

[7] M. Kuribayashi. Tardos's Fingerprinting Code over AWGN Channel. In R. Böhme, P.W.L. Fong, and R. Safavi-Naini, editors, *Information Hiding*, volume 6387 of *LNCS*, pages 103–117. Springer, 2010.

[8] M. Kuribayashi. A new soft decision tracing algorithm for binary fingerprinting codes. In T. Iwata and M. Nishigaki, editors, *IWSEC*, volume 7038 of *LNCS*, pages 1–15. Springer, 2011.

[9] M. Kuribayashi, N. Akashi, and M. Morii. On the systematic generation of Tardos's fingerprinting codes. In *MMSP 2008*, pages 748–753.

[10] T. Laarhoven and B.M.M. de Weger. Optimal symmetric Tardos traitor tracing schemes. 2011. `http://arxiv.org/abs/1107.3441`.

[11] P. Meerwald and T. Furon. Towards joint Tardos decoding: the 'Don Quixote' algorithm. In *Information Hiding*, volume 6958 of *LNCS*, pages 28–42. Springer, 2011.

[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, S. Fujitsu, M. Hagiwara, T. Kitagawa, H. Watanabe, K. Ogawa, and H. Imai. An improvement of discrete Tardos fingerprinting codes. *Des. Codes Cryptography*, 52(3):339–362, 2009.

[14] 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.

[15] G. Tardos. Optimal probabilistic fingerprint codes. In *STOC 2003*, pages 116–125.

[16] 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.

[17] B. Škorić, S. Katzenbeisser, H.G. Schaathun, and M.U. Celik. Tardos Fingerprinting Codes in the Combined Digit Model. *IEEE Transactions on Information Forensics and Security*, 6(3):906–919, 2011.

[18] B. Škorić, T.U. Vladimirova, M.U. Celik, and J.C. Talstra. Tardos fingerprinting is better than we thought. *IEEE Trans. on Inf. Theory*, 54(8):3663–3676, 2008.

[19] F. Xie, T. Furon, and C. Fontaine. On-off keying modulation and Tardos fingerprinting. In *MM&Sec 2008*, pages 101–106.

A Solving the geodesic equations

The metric $G(s_1, s_2)$ is a 2×2 symmetric matrix whose components can be derived from equations (22) and (23):

$$G_{ij}(s_1, s_2) = \frac{1}{4} \sum_{w} \frac{(m_{i+1,w} - m_{1,w})(m_{j+1,w} - m_{1,w})}{m_{1,w} + s_1(m_{2,w} - m_{1,w}) + s_2(m_{3,w} - m_{1,w})}, \tag{25}$$

with $i, j \in \{1, 2\}$. The Christoffel symbols Γ^i_{jk} for this metric are defined as

$$\Gamma^i_{jk} \triangleq \frac{1}{2} \sum_{m=1}^{2} G^{-1}_{im} \left(\frac{\partial G_{jm}}{\partial s_k} + \frac{\partial G_{km}}{\partial s_j} - \frac{\partial G_{jk}}{\partial s_m} \right) \tag{26}$$

where G^{-1} is the matrix inverse of G. We are looking for a shortest curve $(s_1(x), s_2(x))$ with $x \in \mathbb{R}$ from the point $(s_1, s_2) = (0, 0)$ to the point $(1, 0)$. The geodesic equations read

$$\begin{aligned}
s_1'(x) &= k_1(x) \\
s_2'(x) &= k_2(x) \\
k_1'(x) &= -\Gamma^1_{11} k_1^2(x) - 2\Gamma^1_{12} k_1(x) k_2(x) - \Gamma^1_{22} k_2^2(x) \\
k_2'(x) &= -\Gamma^2_{11} k_1^2(x) - 2\Gamma^2_{12} k_1(x) k_2(x) - \Gamma^2_{22} k_2^2(x).
\end{aligned} \tag{27}$$

Once we specify the initial conditions for $s_1(0), s_2(0), k_1(0), k_2(0)$ we can solve (27) numerically to obtain the geodesic curves starting at $(s_1(0), s_2(0))$ with initial 'velocity' vector $(k_1(0), k_2(0))$.