

PERFORMANCE OF RANDOM FINGERPRINTING CODES UNDER ARBITRARY NONLINEAR ATTACKS

Pierre Moulin and Negar Kiyavash

Beckman Inst., Coord. Sci. Lab and ECE Department
University of Illinois at Urbana-Champaign, USA

ABSTRACT

This paper analyzes the performance of arbitrary nonlinear collusion attacks on random fingerprinting codes. We derive the error exponent of the fingerprinting system, which determines the exponential decay of the error probability. A Gaussian ensemble and an expurgated Gaussian ensemble of codes are considered. The collusion attacks include order-statistics attacks as special cases. In our model, a correlation detector is used. The colluders create a noise-free forgery by applying an arbitrary nonlinear mapping to their individual copies, and next they add a Gaussian noise sequence to form the final forgery. The colluders are subject to a mean-squared distortion constraint between host and forgery. We prove that the uniform linear averaging attack outperforms all others.

Index Terms: Digital fingerprinting, coding, detection performance, nonlinear signal processing.

1. INTRODUCTION

Digital fingerprinting systems can be used for traitor tracing or digital rights management applications. A length- N real-valued signal is to be protected and distributed to M users. Some of the users (K of them) may collude and process their copies to create a *forgery* that contains only weak traces of their fingerprints. This problem was first posed by Cox *et al.* [1] who proposed the use of *Gaussian fingerprints* for this purpose. Specifically, their fingerprints were drawn randomly from an i.i.d. Gaussian distribution; the fingerprint code is shared with the detector but not revealed to the users.

A fundamental question is what are the optimal performance limits for detection of colluders. To make the problem nontrivial, one may assume embedding distortion constraints on the fingerprinter and the colluders. Example of this analysis include [2, 3] for the case of signals defined over finite alphabets, and [4, 5, 6] for the case of real-valued signals. In the latter case, an obvious (but not necessarily optimal) strategy for the colluders is to perform a uniform linear average of their copies and add i.i.d. Gaussian noise; this strategy was examined in the above papers. Possible improvements for the attackers consist of developing (nonlinear) order-statistics attacks, as proposed by Stone [7]. Computer simulation results

for seven order-statistics collusion attacks have been reported in [7, 8, 9], sometimes with conflicting findings.

Our study aims at developing a comprehensive detection-theoretic analysis of collusion attacks and identifying an optimal strategy for the colluders. The analysis is rooted in large-deviations theory. Initial results were reported in [10] for the class of order-statistics attacks, assuming a correlation detector and constraining the mean-squared distance between the host and the forgery. Under those assumptions, we proved that the uniform linear averaging strategy is optimal for the colluders in the class of order-statistics attacks. The analysis is extended in this paper to a broader class of nonlinear attacks.

In our problem setup, two random ensembles of fingerprinting codes are considered. The first one is the same as the one used by Cox [1] and other researchers and is shown to be less performant than the second one, which is an expurgated ensemble (bad codes are eliminated). The detector has access to a forgery as well as to the host signal (nonblind detection) and performs a *binary hypothesis test* on each user to determine whether that user was involved in the forgery. The cost functions in this problem are the detector's type-I and type-II probabilities of error, which the colluders want to maximize.

Throughout this paper, we use boldface uppercase letters to denote random vectors, uppercase letters for the components of the vectors, and calligraphic fonts for sets. We use the symbol \mathbb{E} to denote mathematical expectation. For any collection of samples $\{x_1, \dots, x_M\}$, we denote by $x_{\mathcal{K}} = \{x_k, k \in \mathcal{K}\}$ the restriction of this collection to its elements $m \in \mathcal{K}$. The symbols $f(N) \ll g(N)$ and $f(N) \sim g(N)$ (asymptotic equality) mean that $\lim_{N \rightarrow \infty} \frac{f(N)}{g(N)} = 0$ and $\lim_{N \rightarrow \infty} \frac{f(N)}{g(N)} = 1$, respectively. The symbol $f(N) \doteq g(N)$ denotes asymptotic equality on the exponential scale: $\ln f(N) \sim \ln g(N)$. Of course, one may have $f(N) \ll g(N)$ and $f(N) \doteq g(N)$ simultaneously. The Gaussian distribution with mean zero and variance σ^2 is denoted by $\mathcal{N}(0, \sigma^2)$.

2. PROBLEM STATEMENT

The mathematical setup of the problem is diagrammed in Fig. 1.

This research was supported in part by NSF grant CCR 03-25924.

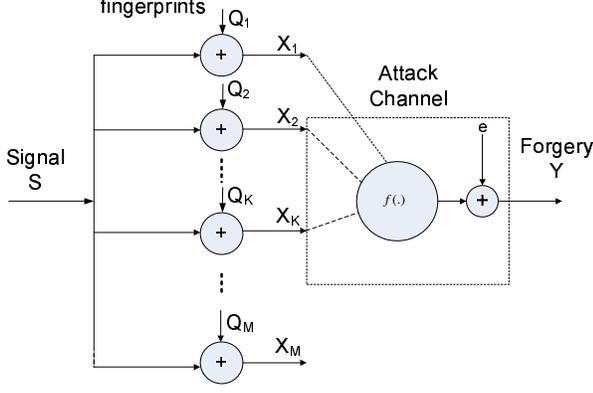


Fig. 1. The fingerprinting process and the attack channel.

2.1. Fingerprint Generation and Embedding

The host signal is a sequence $\mathbf{S} = (S(1), \dots, S(N))$ in \mathbb{R}^N , viewed as *deterministic* but *unknown* to the colluders. Fingerprints are added to \mathbf{S} , and the marked copies of the signal are distributed to M users. Specifically, user m is assigned a marked copy $\mathbf{X}_m = \mathbf{S} + \mathbf{Q}_m$ where $m \in \{1, \dots, M\}$ and $\mathbf{Q}_m \in \mathbb{R}^N$ is the fingerprint assigned to user m .

The fingerprints $\mathbf{Q}_1, \dots, \mathbf{Q}_M$ form a (N, M) fingerprinting code \mathcal{C} . The code \mathcal{C} is selected independently of \mathbf{S} from a random ensemble of codes, \mathcal{E} , such that

$$\mathbb{E}_{\mathcal{E}} [\|\mathbf{X}_m - \mathbf{S}\|^2] = \mathbb{E}_{\mathcal{E}} [\|\mathbf{Q}_m\|^2] = ND_f, \quad \forall m,$$

i.e., the expected mean-squared distortion is equal to D_f . The random ensembles \mathcal{E} of codes considered in this paper are invariant to permutations of users (m) and samples (n).

2.2. Attack Model

The attacks are of the form

$$\mathbf{Y} = f_N(\mathbf{X}_{\mathcal{K}}) + \mathbf{E} \quad (1)$$

where \mathcal{K} , the *coalition*, is the index set of the colluding users. The coalition has cardinality $K \leq M$. Moreover the noise \mathbf{E} is i.i.d. $\mathcal{N}(0, \sigma_E^2)$ and is independent of $\mathbf{X}_{\mathcal{K}}$.

The mapping $f_N : \mathbb{R}^{N|\mathcal{K}|} \rightarrow \mathbb{R}^N$ in (1) is symmetric in its arguments, i.e., any permutation of the index set \mathcal{K} does not change the value of f_N . We view f_N as a “noise-free forgery” to which noise \mathbf{E} is added to form the actual forgery, \mathbf{Y} . The symmetry requirement on f_N represents a *fairness condition*: all members of the coalition incur equal risk. Another requirement is that f_N satisfy the *separation condition*

$$f_N(\mathbf{X}_{\mathcal{K}}) = f_N(\mathbf{Q}_{\mathcal{K}}) + \mathbf{S}. \quad (2)$$

An example is the order statistic attack

$$f_N(\mathbf{X}_{\mathcal{K}})(n) = \sum_{k=1}^K a_k X_{(k)}(n), \quad 1 \leq n \leq N, \quad (3)$$

where $X_{(k)}(n)$ is the k -th order statistic of the K -vector $\mathbf{X}_{\mathcal{K}}(n)$, and f_N is parameterized by the vector $\mathbf{a} = \{a_1, \dots, a_K\}$. The fairness and separation conditions (1) and (2) are satisfied provided that $\sum a_k = 1$ and the sequence a_k is symmetric. The special case of $a_k \equiv 1/K$ reduces to the popular *uniform linear averaging attack*,

$$\bar{\mathbf{X}} = f_N(\mathbf{X}_{\mathcal{K}}) = \frac{1}{K} \sum_{k \in \mathcal{K}} \mathbf{X}_k.$$

If the attackers can retrieve the original signal \mathbf{S} , they will succeed in defeating the detector. It is therefore useful to view $f_N(\mathbf{X}_{\mathcal{K}})$ as an estimator of \mathbf{S} based on the copies available to the coalition \mathcal{K} . The mean-squared distortion of the forgery \mathbf{Y} relative to \mathbf{S} is given by

$$\mathbb{E}\|\mathbf{Y} - \mathbf{S}\|^2 = ND_c \quad (4)$$

where D_c is the average distortion per sample introduced by the coalition. Under the attack model (1) (2), we have

$$\mathbb{E}\|\mathbf{Y} - \mathbf{S}\|^2 = \mathbb{E}_{\mathcal{E}} [\|f_N(\mathbf{Q}_{\mathcal{K}})\|^2] + \mathbb{E}\|\mathbf{E}\|^2,$$

and thus $D_c \geq \sigma_E^2$. The difference

$$D_c - \sigma_E^2 = \frac{1}{N} \mathbb{E}_{\mathcal{E}} [\|f_N(\mathbf{Q}_{\mathcal{K}})\|^2] \quad (5)$$

represents the mean-squared estimation error. For the Gaussian ensemble \mathcal{E} , the mean-squared estimation error (5) is minimized by the uniform linear averaging f_N . In this case,

$$D_c = \sigma_E^2 + \frac{D_f}{K}. \quad (6)$$

2.3. Detector

We study the nonblind scenario where the host signal \mathbf{S} is available at the detector and can be subtracted from \mathbf{Y} , to form the centered content $\mathbf{Y} - \mathbf{S}$. The detector performs a binary hypothesis test to determine whether a specific user’s mark is present. The detector knows neither the mapping f_N nor even the number of colluders K . When focused on user m , our detector computes the correlation statistic $T_m(\mathbf{Y})$ below and compares it with a threshold τ :

$$T_m(\mathbf{Y}) = \mathbf{Q}_m^T(\mathbf{Y} - \mathbf{S}) = \mathbf{Q}_m^T[f_N(\mathbf{Q}_{\mathcal{K}}) + \mathbf{E}] \begin{array}{l} H_1(m) \\ \geq \tau \\ H_0(m) \end{array} \quad (7)$$

where $H_1(m)$ and $H_0(m)$ respectively denote the “guilty” and “innocent” hypotheses. The threshold τ trades off the type-I and type-II probabilities of error. The detector assumes an upper bound K_{\max} on K , and this is reflected in the choice of τ (see below).

The detector (7) does not know the mapping f_N used by the colluders or even the exact number K of colluders. However the detector’s performance generally depends on these quantities. For any given user m , the possible error events are a false positive (incorrectly declaring the user to be guilty) or a false negative (incorrectly declaring the user to be innocent). For any fixed \mathcal{K} , the corresponding type-I and type-II

error probabilities are given by $P_I(f, \mathcal{K}, m) = \Pr[T_m(\mathbf{Y}) > \tau \mid m \notin \mathcal{K}]$ and $P_{II}(f, \mathcal{K}, m) = \Pr[T_m(\mathbf{Y}) < \tau \mid m \in \mathcal{K}]$, where the average is with respect to the random ensemble \mathcal{C} of codes and the noise \mathbf{E} . By our invariance assumptions, these probabilities are independent of m and \mathcal{K} . The overall type-I and type-II error probabilities (worst case over all m, \mathcal{K}) are given by

$$\begin{aligned} P_I(K, f) &= \Pr[\max_{\mathcal{K}} \max_{m \notin \mathcal{K}} T_m(\mathbf{Y}) > \tau] \\ &\leq \binom{M}{K+1} P_I(f, \mathcal{K}, m) \end{aligned} \quad (8)$$

$$\begin{aligned} P_{II}(K, f) &= \Pr[\max_{\mathcal{K}} \max_{m \in \mathcal{K}} T_m(\mathbf{Y}) < \tau] \\ &\leq \binom{M}{K} P_{II}(f, \mathcal{K}, m) \end{aligned} \quad (9)$$

where the upper bounds follow from the union bound. Detection is said to be reliable if (8) and (9) are small enough.

2.4. Background on Large Deviations

Consider a sequence of i.i.d. random variables $U(n), 1 \leq n \leq N$, drawn from a distribution p_U with zero mean and variance σ_U^2 . Denote by $\Lambda_U(\omega) = \ln \mathbb{E}[e^{\omega U}]$ the cumulant-generating function for U . Recall that $\Lambda_U(0) = \Lambda'_U(0) = 0$ and $\Lambda''_U(0) = \sigma_U^2$. Of interest are limiting forms (as $N \rightarrow \infty$) of the probability

$$\Pr \left[\sum_{n=1}^N U(n) > Nt \right] \leq e^{-N\Lambda_U^*(t)}, \quad \forall t > 0 \quad (10)$$

where $\Lambda_U^*(t) = \sup_{\omega > 0} (t\omega - \Lambda_U(\omega))$ is the large deviations function associated with p_U [11]. By Cramer's theorem, the upper bound (10) is tight in the exponent as $N \rightarrow \infty$. If $p_U = \mathcal{N}(0, \sigma_U^2)$, then $\Lambda_U^*(t) = \frac{t^2}{2\sigma_U^2}$. Moreover, for any p_U :

$$\Lambda_U^*(t) \sim \frac{t^2}{2\sigma_U^2} \quad \text{as } t \rightarrow 0, \quad (11)$$

i.e., the exponent in (10) depends on p_U only via σ_U .

The upper bound (10) is an application of Markov's inequality and remains valid if t itself is a function of N . When t is small enough, specifically $t = N^{-1/2}a$ where $a \ll N^{1/6}$, the Central Limit Theorem (CLT) applies, and we have a sharper result, namely, the asymptotic equality

$$\Pr \left[\sum_{n=1}^N U(n) > N^{1/2}a \right] \sim \mathcal{Q} \left(\frac{a}{\sigma_U} \right) \doteq \exp \left\{ -\frac{a^2}{2\sigma_U^2} \right\}$$

where $\mathcal{Q}(x) \triangleq \int_x^\infty (2\pi)^{-1/2} \exp\{-u^2/2\} du$.

2.5. Memoryless Attacks

Lemma 1. For permutation-invariant \mathcal{C} and the correlation detector $T_m(\mathbf{Y})$, there is no loss of optimality in restricting the colluders' strategies to memoryless mappings, i.e., to f_N

of the form $f_N(X_{\mathcal{K}}) = \{f(X_{\mathcal{K}}(1)), \dots, f(X_{\mathcal{K}}(N))\}$ for some $f : \mathbb{R}^K \rightarrow \mathbb{R}$.

Lemma 2. Any mapping f satisfying the fairness and separation conditions (1) and (2) is of the form $f(X_{\mathcal{K}}) = f_{\text{ula}}(X_{\mathcal{K}}) + g(\tilde{X})$ where $f_{\text{ula}}(X_{\mathcal{K}}) = \frac{1}{K} \sum_{k \in \mathcal{K}} X_k$ denotes the uniform linear averaging mapping, \tilde{X} is the K -vector of centered order statistics: $\tilde{X}_k = X_{(k)} - \bar{X}$, $1 \leq k \leq K$, and $g : \mathbb{R}^K \rightarrow \mathbb{R}$ is an arbitrary mapping.

Example: $g(\tilde{X}) = \left[\frac{1}{2}(\tilde{X}_{(1)}^q + \tilde{X}_{(K)}^q) \right]^{1/q}$ for odd $q \in \mathbb{Z}$.

Under the assumptions above, we define a compact set \mathcal{F} of feasible mappings f ; this set is convex.

3. GAUSSIAN ENSEMBLE

Consider the Gaussian ensemble \mathcal{C} , in which random codes $\mathcal{C} = \{\mathbf{Q}_m, 1 \leq m \leq M\}$ are obtained by drawing fingerprint components $Q_m(n)$ i.i.d. $\mathcal{N}(0, D_f)$. Define the five random variables

$$\begin{aligned} Z_{0,f} &= Q_m f(Q_{\mathcal{K}}) \quad (m \notin \mathcal{K}) \\ Z_{1,f} &= Q_m f(Q_{\mathcal{K}}) \quad (m \in \mathcal{K}) \\ W &= Q_m E \\ U_{0,f} &= Z_{0,f} + W \\ U_{1,f} &= Z_{1,f} + W. \end{aligned}$$

By our assumptions on f and \mathcal{C} , the distributions of these random variables do not depend on m and \mathcal{K} . In particular, $\text{Var}(W) = D_f \sigma_E^2$. It may be shown that

$$\mathbb{E}[Z_{0,f}] = \mathbb{E}[U_{0,f}] = \mathbb{E}[W] = 0, \quad \mathbb{E}[Z_{1,f}] = \mathbb{E}[U_{1,f}] = D_f/K.$$

Finally, note that $Z_{0,f}, Z_{1,f}, W, U_{0,f}$ and $U_{1,f}$ are **non-Gaussian**, even if $f = f_{\text{ula}}$.

Proposition 1. Let $0 \leq \tau \leq \frac{ND_f}{K_{\max}}$. For the Gaussian ensemble \mathcal{C} , we have

$$\begin{aligned} P_I(f, \mathcal{K}, m) &\leq \exp \left\{ -N\Lambda_{U_{0,f}}^* \left(\frac{\tau}{N} \right) \right\} \\ P_{II}(f, \mathcal{K}, m) &\leq \exp \left\{ -N\Lambda_{U_{1,f}}^* \left(\frac{D_f}{K} - \frac{\tau}{N} \right) \right\}. \end{aligned}$$

Moreover these bounds are tight in the exponent as $N \rightarrow \infty$.

Proof. The test statistic $T_m(\mathbf{Y})$ in (7) takes the form of a sum of N i.i.d. random variables $U_{0,f}(n)$ and $U_{1,f}(n)$ under hypotheses H_0 and H_1 , respectively. Therefore $P_I(f, \mathcal{K}, m) = \Pr[T_m(\mathbf{Y}) > \tau \mid m \notin \mathcal{K}]$ and $P_{II}(f, \mathcal{K}, m) = \Pr[T_m(\mathbf{Y}) < \tau \mid m \in \mathcal{K}]$ satisfy the large-deviations bound (10). \square

Proposition 2. In the limit as $N \gg K_{\max} \geq K \rightarrow \infty$, we have the asymptotic equalities

$$\begin{aligned} \Lambda_{U_{0,f}}^* \left(\frac{\tau}{N} \right) &\sim \frac{(\tau/N)^2}{2 \text{Var}(W)} = \frac{(\tau/N)^2}{2 D_f \sigma_E^2} \leq \frac{D_f}{2\sigma_E^2 K_{\max}^2}, \\ \Lambda_{U_{1,f}}^* \left(\frac{D_f}{K} - \frac{\tau}{N} \right) &\sim \frac{1}{2 D_f \sigma_E^2} \left(\frac{D_f}{K} - \frac{\tau}{N} \right)^2 \leq \frac{D_f}{2\sigma_E^2 K_{\max}^2}. \end{aligned}$$

Proof. Recalling the range of τ in Prop. 1, we see that the arguments of the large-deviations functions $\Lambda_{U_{0,f}}^*$ and $\Lambda_{U_{1,f}}^*$ above vanish as $K \rightarrow \infty$. The claim follows from (11). \square

Prop. 1 states that the **error exponents** depend on the nonlinear mapping f selected by the colluders, and therefore detection performance **strongly** depends on f as $N \rightarrow \infty$. However, as indicated by Prop. 2, that exponential dependency vanishes for large K . For fixed values of K and N , one can resort to numerical simulations [7, 8, 9]. However, for fixed K , the Central Limit Theorem arguments advanced in [8, 9] are not applicable as $N \rightarrow \infty$. We conclude this section with Prop. 3 which establishes a fundamental relationship between N , M and K_{\max} , guaranteeing reliable detection for the random ensemble \mathcal{C} .

Proposition 3. For the Gaussian ensemble \mathcal{C} , reliable detection is guaranteed provided that $K_{\max}^2 \ln M \ll N$. *Proof:* follows from Props. 1, 2, and (8), (9).

4. EXPURGATED GAUSSIAN ENSEMBLE

The problem with the Gaussian ensemble \mathcal{C} of Sec. 3 is that error probability (which is obtained by averaging over all codes in \mathcal{C}) may be dominated by bad codes. This is a standard problem in information theory for the design of low-rate codes, for which performance is dictated by minimum-distance considerations, and the bad codes are the ones with poor minimum distance [12]. Improvements can be obtained using *expurgation* techniques, i.e., removing bad codes from the random ensemble.

We apply a similar idea to our fingerprinting problem and show that performance can indeed be improved for any finite K if we pick \mathcal{C} judiciously rather than drawing it randomly from \mathcal{C} . The derivations are much more technical than the ones given in Sec. 3 and will be presented elsewhere. The basic ideas are sketched below.

Since the code $\mathcal{C} = \{\mathbf{Q}_m\}$ is known to the detector, the quantity $\mathbf{Q}_m^T f_N(\mathbf{Q}_K)$ in (7) may be viewed as a *deterministic functional* of the unknown f rather than as a random variable. The only source of randomness in $T_m(\mathbf{Y})$ is the Gaussian noise $\mathbf{Q}_m^T \mathbf{E}$ which follows a $\mathcal{N}(0, ND_f \sigma_E^2)$ distribution.

Choose a sequence $\epsilon_N \ll 1$. Let p_N be the probability that a code drawn from the iid Gaussian distribution satisfies the conditions below for all m, \mathcal{K} , and let \mathcal{C} be the ensemble of such codes, which we call the expurgated ensemble.

$$|\mathbb{E}[T_m(\mathbf{Y})]| = |\mathbf{Q}_m^T f_N(\mathbf{Q}_K)| < \epsilon_N \tau D_f^{-1/2} \sigma_E^{-1} \quad : m \notin \mathcal{K} \quad (12)$$

$$|\mathbb{E}[T_m(\mathbf{Y})] - ND_f/K| < \epsilon_N \tau D_f^{-1/2} \sigma_E^{-1} \quad : m \in \mathcal{K}. \quad (13)$$

We have proved that p_N tends to 1 as $N \rightarrow \infty$, provided that $K_{\max}^3 \ln M \ll N$. This suggests the following procedure for selecting a code from \mathcal{C} . Pick a code randomly from the iid Gaussian ensemble and check whether this code satisfies (12) and (13). If it does, use that code. If it does not, discard it and repeat the above procedure until the it is successful. The

probability that the procedure is still unsuccessful after t trials is only $(1 - p_N)^t$.

Proposition 4. Assume $0 \leq \tau \leq \frac{ND_f}{K_{\max}}$ and $K_{\max}^3 \ln M \ll N$. For the expurgated ensemble \mathcal{C} , the type-I and type-II error probabilities satisfy

$$P_I(f, \mathcal{K}, m) \leq \mathcal{Q} \left(\frac{\sqrt{\tau}}{\sigma_E} (1 - \epsilon_N) \right) \doteq \exp \left\{ -\frac{\tau^2}{2\sigma_E^2} \right\}, \quad (14)$$

$$\begin{aligned} P_{II}(f, \mathcal{K}, m) &\leq \mathcal{Q} \left(\frac{1}{\sigma_E^2} \sqrt{(1 - \epsilon_N) \frac{ND_f}{K} - \tau} \right) \\ &\doteq \exp \left\{ -\frac{1}{2\sigma_E^2} \left(\frac{ND_f}{K} - \tau \right)^2 \right\}. \end{aligned} \quad (15)$$

Proposition 5. For the expurgated ensemble \mathcal{C} , the error exponents in (14) and (15) are minimized by $f = f_{\text{ula}}$ with $K = K_{\max}$.

Proof. Given K and σ_E , the detection bounds (14) and (15) are independent of f . Given D_c , it follows from (6) that f_{ula} with $K = K_{\max}$ simultaneously maximizes K and σ_E , and therefore minimizes the error exponents (14) and (15). \square

For any fixed K , the error exponents in (14) and (15) are uniformly better than those obtained by drawing codes from the Gaussian ensemble. The colluders can choose $f \in \mathcal{F}$ such that the exponents in Prop. 1 are worse than those for f_{ula} .

5. REFERENCES

- [1] I. J. Cox, J. Killian, F. T. Leighton and T. Shamoan, "Secure Spread Spectrum Watermarking for Multimedia," *IEEE T-IP*, Vol. 6, pp. 1673–1687, Dec. 1997. (Also NEC Tech. Rep. 95-10, 1995).
- [2] P. Moulin and J. A. O'Sullivan, "Information-Theoretic Analysis of Information Hiding," *IEEE T-IT*, Vol. 49, No. 3, pp. 563–593, 2003.
- [3] A. Somekh-Baruch and N. Merhav, "On the Capacity Game of Private Fingerprinting Systems Under Collusion Attacks," *Proc. IEEE Int. Symp. on Information Theory*, Yokohama, Japan, p. 191, July 2003.
- [4] J. Kilian, F. T. Leighton, L. R. Matheson, T. G. Shamoan, R. E. Tarjan, and F. Zane, "Resistance of digital watermarks to collusive attacks," *Proc. ISIT*, p. 271, Cambridge, MA, 1998.
- [5] F. Ergun, J. Kilian and R. Kumar, "A Note on the Bounds of Collusion Resistant Watermarks," *Proc. EUROCRYPT*, pp. 140–149, 1999.
- [6] P. Moulin and A. Briassouli, "The Gaussian Fingerprinting Game," *Proc. CISS'02*, Princeton, NJ, March 2002.
- [7] H. S. Stone, "Analysis of Attacks on Image Watermarks With Randomized Coefficients," *NEC TR 96-045*, Princeton, NJ, 1996.
- [8] H. Zhao, M. Wu, Z. Wang and K. J. R. Liu, "Forensic Analysis of Nonlinear Collusion Attacks for Multimedia Fingerprinting," *IEEE T-IP*, Vol. 14, No. 5, pp. 646–661, May 2005.
- [9] Z.J. Wang, M. Wu, H. Zhao, W. Trappe, and K.J.R. Liu, "Anti-Collusion Forensics of Multimedia Fingerprinting Using Orthogonal Modulation," *IEEE T-IP*, Vol. 14, No. 6, pp. 804–821, June 2005.
- [10] N. Kiyavash and P. Moulin, "A Framework for Optimizing Nonlinear Collusion Attacks on Fingerprinting Systems," *Proc. Conf. on Information Systems and Science*, Princeton, NJ, March 2006.
- [11] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed., Springer-Verlag, New York, 1998.
- [12] R. G. Gallager, *Information Theory and Reliable Communication*, Wiley, NY, 1968.