RISK MINIMIZATION IN TRAITORS WITHIN TRAITORS IN MULTIMEDIA FORENSICS

H. Vicky Zhao and K. J. Ray Liu

Department of Electrical and Computer Engineering
University of Maryland, College Park, MD 20742

ABSTRACT

In digital fingerprinting and multimedia forensic systems, it is possible that multiple adversaries mount attacks collectively and effectively to undermine the forensic system's traitor tracing capability. During this collusion attack, an important issue that the adversaries need to address is the fairness of attack and ensuring that all colluders share the same risk of being caught. This paper studies the dynamics among attackers in enforcing the fairness of collusion and investigates the problem of traitors within traitors, in which some selfish colluders wish to minimize their own risk while still profiting from collusion. We explore the strategies that these selfish colluders can use to further lower their probability of being detected and analyze their performance. We show that by processing their fingerprinted copies before multi-user collusion, the selfish colluders can further reduce their risk at the cost of quality degradation of their fingerprinted copies.

1. INTRODUCTION

In digital fingerprinting, unique identification information, known as "fingerprints", are embedded in each distributed copy and can be used to track the usage of multimedia data. *Multi-user collusion* is a powerful attack against digital fingerprinting, and it uses several differently marked copies of the same content to remove the identifying fingerprints [1]. To support consistent traitor tracing in multimedia forensics, the embedded fingerprints must survive multi-user collusion as well as attacks by a single adversary.

Modeling and analyzing collusion attacks help the digital rights enforcer understand the challenges in multimedia fingerprinting and design collusion secure fingerprint codes. The work in [2] studied collusion attacks on fingerprints for generic data. Observing the uniqueness of multimedia that fingerprints can be seamlessly embedded into the host signal, the fingerprint design and embedding were jointly considered in [3], and the collusion attacks on multimedia fingerprints were modeled as the averaging attacks followed by an additive noise. Collusion attacks were generalized to linear shift invariant filtering followed by an additive noise in [4]. In [1], several types of collusion attacks were studied, including a few order statistics based nonlinear collusion.

During collusion, the attackers not only share the profit from the illegal usage of multimedia, they also share the risk of being captured. Since no one is willing to take a higher risk than the others, the attackers usually agree to distribute the risk evenly among themselves and apply *fair* collusion. Most prior work assumed that the attackers keep their agreement to share the same risk during collusion. In reality, however, the assumption of fair-play does not always hold. Some selfish colluders may break the fairness agreement with others by trying to further lower their risk of being caught. The existence of such selfish colluders complicates collusion. To build a complete model of multi-user collusion, it is important to study this problem of "traitors within traitors" and understand the attackers' behavior during collusion to minimize their risk and protect their interests. As the first step in analyzing the dynamics among colluders, this paper investigates the possible strategies by the selfish colluders to minimize their risk and analyzes their performance.

This paper is organized as follows. We begin in Section 2 with the introduction of digital fingerprinting systems and the model of the traitors within traitors. Section 3 investigates the possible techniques that the selfish colluders can use to further reduce their risk and analyzes their performance. Section 4 shows the simulation results, and conclusions are drawn in Section 5.

The authors can be reached at hzhao and kjrliu@eng.umd.edu.

2. SYSTEM MODEL

2.1. Digital Fingerprinting System Model

A digital fingerprinting system usually consists of three parts: fingerprint embedding, collusion attacks, and fingerprint detection.

Fingerprint Embedding Spread spectrum embedding is widely used in multimedia fingerprinting due to its robustness against many attacks [5]. In spread spectrum embedding, for the jth frame in the video sequence represented by a vector $\mathbf{S}_j$ of length N_j, for user $\mathbf{u}^{(i)}$ in the system, the content owner generates a unique fingerprint $\mathbf{W}_j^{(i)}$ of length N_j. The fingerprinted frame j that will be distributed to $\mathbf{u}^{(i)}$ is $\mathbf{X}_j^{(i)} = \mathbf{S}_j + JND_j \cdot \mathbf{W}_j^{(i)}$, where JND_j is the just-noticeable-difference from human visual models [5] to control the energy of the embedded fingerprints.

We use Gaussian distributed fingerprints due to their resistance to a wide range of attacks, and generate $\{\mathbf{W}_j^{(i)}\}$ from distribution $\mathcal{N}(0, \sigma_W^2)$. To resist the intra-content collusion attacks on video watermarking systems [6], in each fingerprinted copy, correlated fingerprints are inserted into adjacent frames and the correlation depends on the similarity between the host frames. This is similar to the work in [7]. Furthermore, we generate fingerprints for different users independently.

Multi-user Collusion Attacks Assume that there are a total of K colluders and SC is the set containing their indices. During collusion, the colluders collect all the fingerprinted copies that they received, apply the multi-user collusion function to these copies, and generate the colluded copy $\{\mathbf{V}_j\}$ in which the originally embedded fingerprints are attenuated.

A recent investigation in [8] showed that, under the constraints that the colluded copies under different collusion have the same perceptual quality, the performance of nonlinear collusion attacks

0-7803-9134-9/05/$20.00 ©2005 IEEE

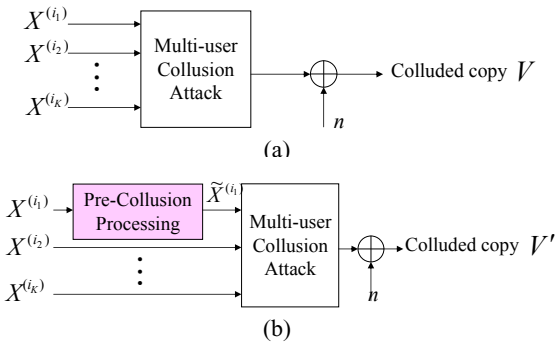

Fig. 1. (a): The collusion attack when all colluders tell each other true information of their fingerprinted copies. (b): The collusion attack when some selfish colluders want to further reduce their own probability of being detected.

is similar to that of the averaging attack. Thus, we only consider the averaging based collusion attacks here.

Fingerprint Detection In digital fingerprinting applications, the host signal can be made available to the detector and the non-blind detection is feasible. To improve the detection performance [8], we consider a non-blind detection scenario where the host signal is first removed from the test copy before colluder identification. During detection, the detector extracts the fingerprint $\mathbf{Y}_j$ from the jth frame $\mathbf{V}_j$ in the test copy. Then, he calculates the similarity between this extracted fingerprint $\{\mathbf{Y}_j\}$ and each of the original fingerprints $\{\mathbf{W}_j^{(i)}\}$, compares with a pre-determined threshold h, and outputs the estimated identities of the colluders $\widehat{SC}$.

To measure the similarity between the extracted fingerprint and the original fingerprint, given $\{\mathbf{Y}_j\}$, for each user $\mathbf{u}^{(i)}$, the detector calculates the correlation based detection statistics

$$T_N^{(i)} = \sum_j \langle \mathbf{Y}_j, \mathbf{W}_j^{(i)} \rangle / \sqrt{\sum_j ||\mathbf{W}_j^{(i)}||^2}, \qquad (1)$$

where $||\mathbf{W}_j^{(i)}||$ is the Euclidean norm of $\mathbf{W}_j^{(i)}$. For a given threshold h, the estimated colluder set is $\widehat{SC} = \{i : T_N^{(i)} > h\}$.

2.2. Traitors Within Traitors

As pointed out in Section 1, to ensure the fairness of collusion, the attackers agree to have the same probability of being detected. Most prior work assumed that the colluders keep their agreement to share the same risk and tell each other the truth information of their received copies. Figure 1 (a) shows an example of the collusion attack in this scenario. Assume that $\mathbf{X}^{(i)}$ is the fingerprinted copy that colluder $\mathbf{u}^{(i)}$ received from the content owner. In this scenario, the multi-user collusion attack is applied to $\{\mathbf{X}^{(i)}\}_{i \in SC}$, and the colluded copy equals to $\mathbf{V} = \sum_{i \in SC} \mathbf{X}^{(i)}/K + \mathbf{n}$, where $\mathbf{n}$ is an additive noise to further hinder the detection.

However, there may exist selfish colluders who wish to further lower their risk of being caught. For example, they may process their fingerprinted signals before multi-user collusion and use the processed copies instead of the originally received ones for participating in the collusion, as shown in Figure 1 (b). Without loss of generality, assume that $\mathbf{u}^{(i_1)}$ is the selfish colluder, and he received the fingerprinted copy $\mathbf{X}^{(i_1)}$ from the content owner. Based on $\mathbf{X}^{(i_1)}$, $\mathbf{u}^{(i_1)}$ generates another copy $\widetilde{\mathbf{X}}^{(i_1)}$ that is perceptually similar to $\mathbf{X}^{(i_1)}$, and uses $\widetilde{\mathbf{X}}^{(i_1)}$ during collusion. If the other colluders fail to discover the pre-collusion processing by $\mathbf{u}^{(i_1)}$, they apply the multi-user collusion attack to the copies $\widetilde{\mathbf{X}}^{(i_1)}$ and $\{\mathbf{X}^{(i)}\}_{i \in SC, i \neq i_1}$, and generate the colluded copy $\mathbf{V}' = (\widetilde{\mathbf{X}}^{(i_1)} + \sum_{i \in SC, i \neq i_1} \mathbf{X}^{(i)})/K + \mathbf{n}$, where $\mathbf{n}$ is an additive noise.

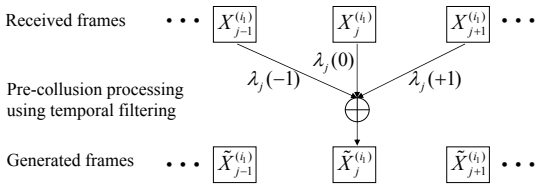

Fig. 2. Pre-collusion processing using temporal filtering.

2.3. Performance Criteria

For a selfish colluder $\mathbf{u}^{(i_1)}$, to measure the effectiveness of pre-collusion processing in reducing his risk, we fix the probability of falsely accusing an innocent user (P_{fa}), and compare $\mathbf{u}^{(i_1)}$'s probability of being captured ($P_d^{(i_1)}$) in two scenarios: when $\mathbf{u}^{(i_1)}$ does not apply pre-collusion processing (i.e., he is willing to share the risk with other colluders), and when $\mathbf{u}^{(i_1)}$ processes his fingerprinted copy before collusion. From the selfish colluder's point of view, the pre-collusion processing technique is more effective when the difference between these two probabilities is larger.

In the example shown in Figure 1 (b), in order to cover up the fact that he processed his fingerprinted copy before multi-user collusion, the selfish colluder $\mathbf{u}^{(i_1)}$ has to ensure that the newly generated copy $\widetilde{\mathbf{X}}^{(i_1)}$ has high quality when compared with $\mathbf{X}^{(i_1)}$. To measure the effect of pre-collusion processing on perceptual quality, we use the mean square error (MSE) between $\widetilde{\mathbf{X}}^{(i_1)}$ and $\mathbf{X}^{(i_1)}$, or equivalently, PSNR in image and video applications.

3. PRE-COLLUSION PROCESSING AND PERFORMANCE ANALYSIS

For a selfish colluder to further reduce his own risk, one possible solution is to attenuate the energy of the fingerprints embedded in his received copy even before multi-user collusion. An example is to replace each segment of the fingerprinted signal with another, seemingly similar segment from different regions of the content, e.g., averaging or swapping consecutive frames of similar content [6]. In this section, we take temporal filtering of adjacent frames as an example, and analyze its effects on the probability of being detected and the perceptual quality of the fingerprinted copies.

3.1. Temporal Filtering Before Multi-user Collusion

We assume that the selfish colluder $\mathbf{u}^{(i_1)}$ received fingerprinted frames $\{\mathbf{X}_j^{(i_1)}\}_{j=1,2,\ldots}$ from the content owner, and uses a simple linear interpolation to produce a temporally filtered video.[1] As shown in Figure 2, for each frame j, $\mathbf{u}^{(i_1)}$ linearly combines the current frame $\mathbf{X}_j^{(i_1)}$, the previous frame $\mathbf{X}_{j-1}^{(i_1)}$, and the next frame $\mathbf{X}_{j+1}^{(i_1)}$ with weights $\lambda_j(0)$, $\lambda_j(-1)$, and $\lambda_j(+1)$, respectively, and generates a new frame $\widetilde{\mathbf{X}}_j^{(i_1)}$ where

$$\widetilde{\mathbf{X}}_j^{(i_1)} = \lambda_j(-1) \cdot \mathbf{X}_{j-1}^{(i_1)} + \lambda_j(0) \cdot \mathbf{X}_j^{(i_1)} + \lambda_j(+1) \cdot \mathbf{X}_{j+1}^{(i_1)}. \quad (2)$$

[1] A selfish colluder can also apply more complicated motion based interpolation [9], and the analysis will be similar.

In (2), $0 \leq \lambda_j(-1), \lambda_j(0), \lambda_j(+1) \leq 1$ and $\lambda_j(-1) + \lambda_j(0) + \lambda_j(+1) = 1$. For simplicity, we let $\lambda_j(-1) = \lambda_j(+1) = (1 - \lambda_j(0))/2$. The selfish colluder $\mathbf{u}^{(i_1)}$ repeats this process for every frame in the video sequence and generates $\{\widetilde{\mathbf{X}}_j^{(i_1)}\}_{j=1,2,\cdots}$. Note that when $\lambda_j(0) = 1$, $\widetilde{\mathbf{X}}_j^{(i_1)} = \mathbf{X}_j^{(i_1)}$ and it corresponds to the scenario where $\mathbf{u}^{(i_1)}$ does not process his copy before collusion.

3.2. Performance Analysis

In this section, we analyze the quality of the newly generated frames $\{\widetilde{\mathbf{X}}_j^{(i_1)}\}$ and the selfish colluder's risk of being captured $P_d^{(i_1)}$.

Perceptual Quality If $\widetilde{\mathbf{X}}_j^{(i_1)}$ is generated as in (2), then the MSE between $\widetilde{\mathbf{X}}_j^{(i_1)}$ and $\mathbf{X}_j^{(i_1)}$ is

$$\begin{aligned} MSE_j &= ||\widetilde{\mathbf{X}}_j^{(i_1)} - \mathbf{X}_j^{(i_1)}||^2 = (\frac{1-\lambda_j(0)}{2})^2 \cdot \phi_j, \\ \text{where } \phi_j &= 4||\mathbf{X}_j^{(i_1)}||^2 + ||\mathbf{X}_{j-1}^{(i_1)}||^2 + ||\mathbf{X}_{j+1}^{(i_1)}||^2 \\ &\quad -4\langle \mathbf{X}_{j-1}^{(i_1)}, \mathbf{X}_j^{(i_1)} \rangle - 4\langle \mathbf{X}_j^{(i_1)}, \mathbf{X}_{j+1}^{(i_1)} \rangle \\ &\quad +2\langle \mathbf{X}_{j-1}^{(i_1)}, \mathbf{X}_{j+1}^{(i_1)} \rangle. \end{aligned} \quad (3)$$

In (3), $||\mathbf{X}_j^{(i_1)}||$ is the Euclidean norm of $\mathbf{X}_j^{(i_1)}$, and $\langle \mathbf{X}_{j-1}^{(i_1)}, \mathbf{X}_j^{(i_1)} \rangle$ is the correlation between $\mathbf{X}_{j-1}^{(i_1)}$ and $\mathbf{X}_j^{(i_1)}$. From (3), a larger $\lambda_j(0)$ implies a smaller MSE_j, and therefore, is preferred from the perceptual quality's point of view. Compared with $\mathbf{X}_j^{(i_1)}$, $\widetilde{\mathbf{X}}_j^{(i_1)}$ has the best possible quality when $\lambda_j(0) = 1$.

Probability of being detected We assume that there is only one selfish colluder $\mathbf{u}^{(i_1)}$ and the other colluders do not discover his pre-collusion processing actions.[2] We can show that in this scenario, the fingerprint extracted from the jth frame in the colluded copy is

$$\begin{aligned} \mathbf{Y}_j &= \frac{\lambda_j(-1) \cdot \mathbf{W}_{j-1}^{(i)} + \lambda_j(0) \cdot \mathbf{W}_j^{(i)} + \lambda_j(+1) \cdot \mathbf{W}_{j+1}^{(i)}}{K} \\ &\quad + \frac{\sum_{i \in SC, i \neq i_1} \mathbf{W}_j^{(i)}}{K} + \mathbf{d}_j, \end{aligned} \quad (4)$$

where $\mathbf{d}_j$ contains terms that are independent of the embedded fingerprints $\{\mathbf{W}_j^{(i)}\}$. For simplicity, we assume that $\mathbf{d}_j$ are i.i.d. and follow Gaussian distribution $\mathcal{N}(0,\ \sigma_n^2)$.

It is straightforward to show that given the colluder set SC and the index of the selfish colluder i_1, the detection statistics follow Gaussian distribution with mean $\mu^{(i)}$ and variance σ_n^2, i.e., $p(T_N^{(i)}|SC, i_1) \sim \mathcal{N}(\mu^{(i)}, \sigma_n^2)$. The detection statistics of an innocent user have a zero mean, and that of a guilty colluder have a positive mean. Consequently, the probability of accusing an innocent user is $P_{fa} \approx Q(h/\sigma_n)$, and the probability of capturing a guilty colluder $\mathbf{u}^{(i \in SC)}$ is $P_d^{(i)} \approx Q((h - \mu^{(i)})/\sigma_n)$. Here, $Q(\cdot)$ is the Gaussian tail function and h is the pre-determined threshold. Consequently, for fixed σ_n^2 and P_{fa}, minimizing the selfish colluder's probability of being detected is equivalent to minimizing the mean of his detection statistics.

For a selfish colluder $\mathbf{u}^{(i_1)}$,

$$\begin{aligned} \mu^{(i_1)} &= \sum_j \mu_j^{(i_1)}, \text{ where} \\ \mu_j^{(i_1)} &= \frac{\langle \mathbf{W}_{j-1}^{(i_1)}, \mathbf{W}_j^{(i_1)} \rangle + \langle \mathbf{W}_j^{(i_1)}, \mathbf{W}_{j+1}^{(i_1)} \rangle}{2K\sqrt{\sum_l ||\mathbf{W}_l^{(i_1)}||^2}} \\ &= -\lambda_j(0) \times \frac{||\mathbf{W}_j^{(i_1)}||^2 - \langle \mathbf{W}_{j-1}^{(i_1)}, \mathbf{W}_j^{(i_1)} \rangle}{2K\sqrt{\sum_l ||\mathbf{W}_l^{(i_1)}||^2}} \\ &\quad -\lambda_j(0) \times \frac{||\mathbf{W}_j^{(i_1)}||^2 - \langle \mathbf{W}_j^{(i_1)}, \mathbf{W}_{j+1}^{(i_1)} \rangle}{2K\sqrt{\sum_l ||\mathbf{W}_l^{(i_1)}||^2}}. \end{aligned} \quad (5)$$

In (5), $\langle \mathbf{W}_{j-1}^{(i)}, \mathbf{W}_j^{(i)} \rangle$ is the correlation between $\mathbf{W}_{j-1}^{(i)}$ and $\mathbf{W}_j^{(i)}$, and $\langle \mathbf{W}_j^{(i)}, \mathbf{W}_{j+1}^{(i)} \rangle$ is the correlation between $\mathbf{W}_j^{(i)}$ and $\mathbf{W}_{j+1}^{(i)}$. $\langle \mathbf{W}_{j-1}^{(i)}, \mathbf{W}_j^{(i)} \rangle \leq \langle \mathbf{W}_j^{(i)}, \mathbf{W}_j^{(i)} \rangle = ||\mathbf{W}_j^{(i)}||^2$ and $\langle \mathbf{W}_j^{(i)}, \mathbf{W}_{j+1}^{(i)} \rangle \leq ||\mathbf{W}_j^{(i)}||^2$ from the fingerprint design in Section 2.1. From (5), if $\lambda_1(0), \cdots, \lambda_{j-1}(0), \lambda_{j+1}(0), \cdots$ are fixed, $\mu^{(i_1)}$ is a non-decreasing function of $\lambda_j(0)$ and is minimized when $\lambda_j(0) = 0$. Thus, from minimizing the risk's point of view, $\mathbf{u}^{(i_1)}$ should choose a smaller $\lambda_j(0)$.

3.3. Selection of the Optimal Weight Vector

During pre-collusion processing, a selfish colluder wishes to minimize his probability of being detected while maintaining the quality of the fingerprinted copies. Thus, for a selfish colluder $\mathbf{u}^{(i_1)}$, the selection of the weight vector $\{\lambda_j(0)\}$ can be modeled as

$$\begin{aligned} &\min_{\{\lambda_j(0)\}} \ \mu^{(i_1)} \\ s.t. \quad &MSE_j \leq \varepsilon,\ 0 \leq \lambda_j(0) \leq 1,\ j = 1, 2, \cdots, \end{aligned} \quad (6)$$

where ε is the constraint on the perceptual quality of $\{\widetilde{\mathbf{X}}_j^{(i_1)}\}$. Given ϕ_j as defined in (3), we can show that the solution to (6) is: for each frame j,

$$\lambda_j^* = \max\left\{0,\ 1 - 2 \cdot \sqrt{\varepsilon/\phi_j}\right\}. \quad (7)$$

By using $\{\lambda_j^*\}$ as in (7) during temporal filtering, a selfish colluders minimizes his own probability of being detected and ensures that the newly generated frames have high perceptual quality ($MSE \leq \varepsilon$) when compared with the originally received copy.

4. SIMULATION RESULTS

In our simulations, we choose a typical video sequence "carphone", and use the first 40 frames as an example. At the content owner's side, we adopt the human visual model based spread spectrum embedding [5] and embed fingerprints in the DCT domain. The fingerprints follow Gaussian distribution $\mathcal{N}(0, 1/9)$, and fingerprints for different users are generated independently. In each fingerprinted copy, similar to the work in [7], fingerprints embedded in adjacent frames are correlated with each other, and the correlation depends on the similarity between the two host frames.

At the colluders' side, we assume that there are a total of 150 colluders. For simplicity, we assume that there is only one selfish colluder and he applies temporal filtering as in (2) during pre-collusion processing. In this paper, we adjust the power of the additive noise such that the noise term $\mathbf{d}_j$ in (4) satisfies $||\mathbf{d}_j||^2 = 2||\mathbf{W}_j^{(i)}||^2$. Other values will give the same trend.

At the detector's side, we consider a non-blind detection scenario. The detector first removes the host signal from the test copy and then applies the fingerprint detection process in Section 2.1.

Figure 3 shows the simulation results of temporally filtering adjacent frames on sequence "carphone". We compare the perceptual quality of the newly generated frames and the selfish col-

[2] The analysis is similar when there are multiple selfish colluders using temporal filtering during pre-collusion processing.

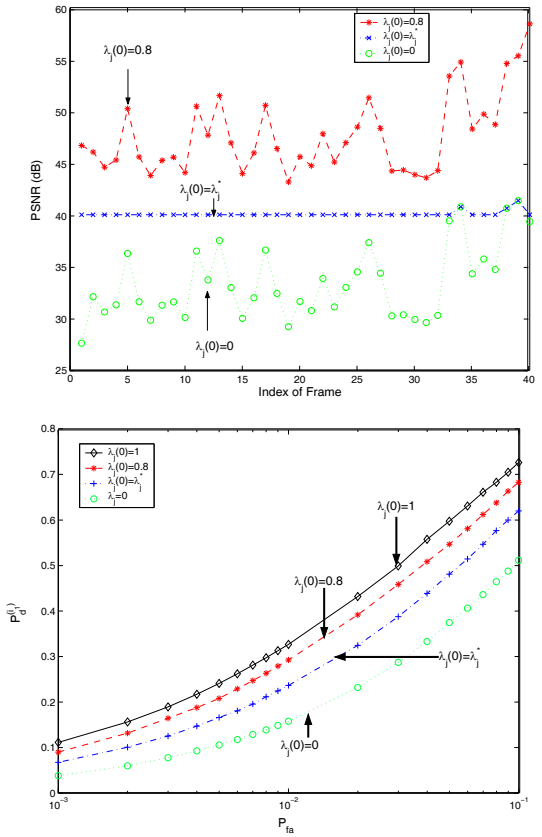

Fig. 3. Simulation results of temporal filtering of adjacent frames on sequence "carphone". (Top): PSNR of the newly generated copy compared with the originally received fingerprinted frames. (Bottom): the selfish colluder's probability of being detected.

luder's probability of being detected when $\{\lambda_j^0\}$ take different values. In Figure 3, for frame j, $PSNR_j$ is defined as PSNR of $\widetilde{\mathbf{X}}_j^{(i_1)}$ compared with $\mathbf{X}_j^{(i_1)}$. In our simulations, we let $\lambda_j(0)$ equal to 1, 0.8, λ_j^* and 0, respectively[3]. $\{\lambda_j^*\}$ are the solution of (7) in which ε is chosen to satisfy $PSNR_j \geq 40dB$ for all frames.

From Figure 3, a selfish colluder can reduce his own probability of being detected by temporally filtering adjacent frames before multi-user collusion. By choosing $\{\lambda_j^0\}$ of smaller values, the selfish colluder has a smaller probability of being detected while sacrificing the quality of the newly generated copy $\{\widetilde{\mathbf{X}}_j^{(i_1)}\}$. Thus, during pre-collusion processing, the selfish colluder has to consider the tradeoff between the risk and the perceptual quality.

Figure 4 compares the solution of $\{\lambda_j^*\}$ in (7) for different video sequences. We choose ε in (7) to satisfy $PSNR_j \geq 40dB$ for all frames in $\{\widetilde{\mathbf{X}}_j^{(i_1)}\}$. From Figure 4, for sequences that have large smooth regions and slow motion ("miss america"), a selfish colluder can choose $\{\lambda_j(0)\}$ with small values, e.g., around 0, without significant quality degradation. For moderately complicated sequences ("carphone" and "foreman"), λ_j^* is around 0.5. For sequences with fast movement and complicated scene composition ("flower"), a selfish colluder has to choose large $\{\lambda_j(0)\}$, e.g., larger than 0.9, to ensure the perceptual quality.

[3] $\lambda_j(0) = 1$ corresponds to the scenario where the selfish colluder $\mathbf{u}^{(i_1)}$ does not process his fingerprinted copy before multi-user collusion.

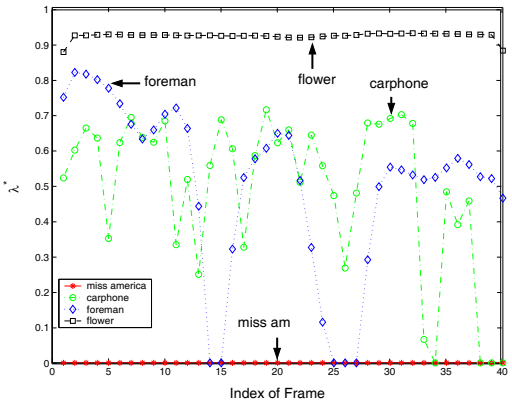

Fig. 4. λ_j^* of (7) for different sequences. ε is chosen to satisfy $PSNR_j \geq 40dB$ for all frames in $\{\widetilde{\mathbf{X}}_j^{(i_1)}\}$.

5. CONCLUSIONS

In this paper, we have studied the dynamics among attackers during collusion and investigated the possible techniques for selfish colluders to minimize their own risk of being caught while still profiting from collusion. For these selfish colluders,, we proposed to apply temporal filtering to the received copies before collusion in order to attenuate energies of the embedded fingerprints, and analyzed its performance. Our results showed that this pre-collusion processing further lowers the selfish colluder's risk of being detected at the cost of quality degradation. We then investigated the selection of the parameters in temporal filtering to minimize the selfish colluder's risk under the quality constraints, and showed that the selfish colluders should adjust these parameters according to the characteristics of video sequences.

6. REFERENCES

[1] H. Stone, "Analysis of attacks on image watermarks with randomized coefficients," Tech. Rep. 96-045, NEC Research Institute, 1996.

[2] D. Boneh and J. Shaw, "Collusion-secure fingerprinting for digital data," *IEEE Tran. on Information Theory*, vol. 44, no. 5, pp. 1897–1905, Sept. 1998.

[3] F. Ergun, J. Killian, and R. Kumar, "A note on the limits of collusion-resistant watermarks," *Advances in Cryptology – EuroCrypto '99, Lecture Notes in Computer Science*, vol. 1592, pp. 140–149, 2001.

[4] J. Su, J. Eggers, and B. Girod, "Capacity of digital watermarks subject to an optimal collusion attacks," *European Signal Processing Conference (EUSIPCO 2000)*, 2000.

[5] C. Podilchuk and W. Zeng, "Image adaptive watermarking using visual models," *IEEE Journal on Sel. Area in Comm.*, vol. 16, no. 4, pp. 525–540, May 1998.

[6] M. Swason, B. Zhu, and A. Tewfik, "Multiresolution scene-based video watermarking using perceptual models," *IEEE Journal on Sel. Area in Comm.*, vol. 16, no. 4, pp. 540–550, May 1998.

[7] K. Su, D. Kundur, and D. Hatzinakos, "Statistical invisibility for collusion-resistant digital video watermarking," *IEEE Tran. on Multimedia*, vol. 7, no. 1, pp. 43–51, Feb. 2005.

[8] Z. J. Wang, M. Wu, H. Zhao, W. Trappe, and K. J. R. Liu, "Anti-collusion forensics of multimedia fingerprinting using orthogonal modulation," *IEEE Trans. on Image Processing*, vol. 14, no. 6, June 2005.

[9] T. Chen, "Adaptive temporal interpolation using bidirectional motion estimation and compensation," *IEEE Int. Conf. on Image Processing*, Sept. 2002.